

THEaiTRE: Generating Theatre Play Scripts using Artificial Intelligence

Patricia Schmidtová
Rudolf Rosa
David Košťák
Tomáš Studeník
Daniel Hrbek
Tomáš Musil
Josef Doležal
Ondřej Dušek
David Mareček
Klára Vosecká
Máša Nováková
Petr Žabka
Alisa Zakhtarenko
Dominik Jurko
Martina Kinská
Tom Kocmi
Ondřej Bojar



THEaiTRE: Generating Theatre Play Scripts using Artificial Intelligence

Patricia Schmidtová, Rudolf Rosa, David Košťák, Tomáš Studeník, Daniel Hrbek, Tomáš Musil, Josef Doležal, Ondřej Dušek, David Mareček, Klára Vosecká, Máša Nováková, Petr Žabka, Alisa Zakhtarenko, Dominik Jurko, Martina Kinská, Tom Kocmi, Ondřej Bojar

Acknowledgement

The project *TL03000348 THEaiTRE: Umělá inteligence autorem divadelní hry* was co-financed with the state support of Technological Agency of the Czech Republic within the ETA 3 Programme, for which we are immensely grateful.

We would like to thank all of the project partners and partner institutions for collaborating on setting this into motion, keeping it running, supporting it with their own means, and delivering all the envisioned results.

We would also like to thank the Institute of the Czech National Corpus and the Dilia and Aura-Pont agencies for helping us in obtaining a range of useful data resources.

The work described herein has been using data, tools and services provided by the LINDAT/CLARIAH-CZ Research Infrastructure (<https://lindat.cz>), supported by the Ministry of Education, Youth and Sports of the Czech Republic (Project No. LM2018101).

Authors

The text of the book is based mainly on contributions by:

- Patricia Schmidtová (Chapters 6, 7, 8, 9, 10, 11, 13)
- Rudolf Rosa (Chapters 2, 3, 4, 11, 12, 13)
- David Košťák (Chapter 5)
- Tomáš Studeník (Chapter 1)

These are the main authors or co-authors of the respective chapters, with numerous contributions from the other book authors.

The texts were assembled and edited by Rudolf Rosa.

All of the authors of this book have contributed to the research, development and application within the THEaiTRE project.

The photos at the beginning of Part I and Part II are by Richard Moučka. The images on the front and back cover were generated by Midjourney under Creative Commons Noncommercial 4.0 Attribution International License.

The chapters 1 and 5 and partially 3 and 4 were originally written in the Czech language. They were translated into English using the CUBBITT machine translation service, and lightly post-edited by Rudolf Rosa.

Published by the Institute of Formal and Applied Linguistics.

Editor-in-chief: Jan Hajič

Editorial board: Nicoletta Calzolari, Mirjam Fried, Eva Hajičová, Petr Karlík, Joakim Nivre, Jarmila Panevová, Patrice Pognan, Pavel Straňák, and Hans Uszkoreit

Reviewers: Lucie Flek (Universität Marburg), Piotr Mirowski (DeepMind and Improbatics)

Copyright © Institute of Formal and Applied Linguistics, 2022

ISBN 978-80-88132-14-1

Contents

I	Putting AI Into Theatre	1
1	Foreword	3
2	The Project	5
2.1	The Two Plays	6
2.2	THEaiTRobot	6
2.3	The Team	7
2.4	Related Projects	9
2.5	Successes and Publicity	12
3	The Two Plays	14
3.1	The First Theatre Play	14
3.1.1	Generating the Play Script	14
3.1.2	The Staging	17
3.1.3	The Premiere	17
3.1.4	Format of the Show	17
3.1.5	The Reactions	17
3.2	The Second Theatre Play	18
3.2.1	The Generation Process	18
3.2.2	The Script	20
3.2.3	The Staging	20
3.2.4	The Premieres and Reactions	20
3.3	Further Projects	22
4	The Situation with Dramatic Situations	24
4.1	Dramatic Situations	25
4.2	Annotation Process	25
4.2.1	First Phase: 58 Situations	25
4.2.2	Second Phase: 5 Situations	27
4.2.3	Annotator Agreement	27
4.3	The Annotated Corpus	28
4.4	Experiments in Predicting the Situations	28
4.5	Discussion	30
5	Is the Author Dead?	31
5.1	First Encounter of a New Kind	31
5.2	Shortcomings	33
5.2.1	Spatial Imagination	33

5.2.2	Schizophrenia	33
5.2.3	Absence of Dramatic Thinking	33
5.2.4	Staticity	33
5.3	Search for an Achievable Goal	34
5.4	Author without Ego	35
5.5	Plagiarist or Talented Pupil	36
5.6	Assembling the Text	36
5.7	Structure of the Script	37
5.8	Conclusion	38
II	What Is Under the Hood	40
6	Theoretical Background	43
6.1	Deep Feedforward Networks	43
6.2	Recurrent Neural Networks	46
6.3	Transformers	48
6.4	Metrics	51
7	Related Work	53
7.1	Story Generation	54
7.2	Script Analysis and Generation	59
7.3	Online Text Generation Demos	61
7.4	Creative Text Evaluation	61
8	NLI-Score: A Consistency Metric	64
8.1	Motivation and Definition	64
8.2	Document-Level NLI-Score	66
8.2.1	Averaged NLI-Score	67
8.2.2	Percentage Below a Threshold	67
8.3	Limitations and Potential Modifications	68
9	Flat Generation	70
9.1	Vanilla GPT-2 XL Generation Setup	71
9.2	Fine-tuned GPT-2 for Script Generation	72
9.2.1	Data	72
9.2.2	Model Training	73
9.2.3	Setting Decoding Parameters	74
9.3	Evaluation	74
10	Hierarchical Generation	78
10.1	Plot Generation	79
10.1.1	Model Selection	79
10.1.2	Dataset Selection	82
10.2	Scene Alignment	84
10.2.1	Script Pre-Processing	85
10.2.2	Script-Summary Alignment	85
10.2.3	Evaluation of Alignments	87
10.3	Script Generation	88
10.4	Evaluation	89

11 Extensions	93
11.1 NLI Decoding Filter	93
11.2 Character Name Anonymization	96
11.3 Character Orchestration	101
11.4 Summarization	103
11.5 Machine Translation	104
11.5.1 Character Translation	105
11.5.2 Gender and T-V Distinction	105
12 THEaiTRobot Demo	107
12.1 THEaiTRobot 1.0	111
12.1.1 Demo at Goethe Institute	111
12.2 THEaiTRobot 2.0	111
12.3 Implementation	112
13 Conclusion	118
13.1 Future Work	119
13.1.1 Stronger Language Model	120
13.1.2 Focusing on Summarization	120
13.1.3 Characters Consistency	121
13.1.4 Automating the Operator Actions	122
13.1.5 Moving Away from Theatre	122
Bibliography	124
A Script of AI: When a Robot Writes a Play	140
B Script of Permeation	187

Part I

Putting AI Into Theatre

Chapter 1

Foreword

Rage: this is not only the first word of Homer's Iliad, but also the first impulse for the THEaiTRE project.

At the end of 2019, I began to research what celebrations were planned for the centenary of Karel Čapek's R.U.R. premiere, where the word robot was first mentioned in February 1921. Less than two years separated us from this historic landmark. There was no indication that anyone would attach more importance to it. Government, companies, schools - everywhere only silence. In March 2020, COVID struck and became the hegemon of our days for the next few years. Nobody remembered the Čapek brothers and their robots.

The word *robot* is the greatest contribution of Czech intelligence to the world. Karel Čapek and his brother Josef first materialized it in 1920 while writing R.U.R. We all know about this play, but few know it. Karel Čapek predicted the deployment of robots in industry and the army. Before Stephen Hawking or Elon Musk, he pointed out the dangers of artificial intelligence. Čapek predicted the demise of humanity not as a result of its extermination by robots, but as a result of the prosperity to which robotization leads.

Robots and neural networks already co-create the human world. The United States or China were able to build technologies within a hundred years from the first utterance of the word robot in 1920 that decide what we buy on the Internet, recognize our faces, diagnose tumors, grant us loans and drive our cars. The Czech word ignited today's fastest growing technology field. Japan has robotic seals that pass the time for lonely seniors. Germany churns out robotic belts of automobiles. America churns out autonomous vehicles. And the Czech Republic? Did it miss its chance to play the first league on the field it painted?

With the THEaiTRE project, the Czech Republic is able to return to the world's leading players. There is no point in trying to compete with robotic powers, we can't catch up with them. Instead of plain artificial intelligence, let's explore artificial intelligence which is creative. After all, the Čapek brothers gave birth to a robot through a play, on the stage. It's no coincidence that the robot saw the light of day in the context of art,

creativity. Let's discover how we can compose music with robots in the future, paint paintings, write poems, make films or do theatre. Maybe together we'll come up with new forms of art that we would never have dreamed of without machines.

At the birth of the THEaiTRE project there were two other people who might never have met in their lives. And I had coffee with both of them first and then put them together: the geek Rudolf Rosa and the theatre-maker Daniel Hrbek. Rudolf liked the idea of letting the robot write a play about people immediately. Daniel thought it was so stupid that he couldn't do it. We high-fived in December 2019 at the Švanda Theatre - we'd squeeze the play out of the machine and the theatre ensemble would play it a hundred years after the R.U.R. premiere.

The pandemic threw a spanner in the works of the planned world premiere, but, also thanks to the support of the Technology Agency of the Czech Republic, we managed to bring the play "AI: When a Robot Writes a Play" online on February 26, 2021. The stream was seen by over twenty thousand viewers around the world during the first three days.

The initial rage was transformed into a dream that came true.

As the founding father of the THEaiTRE project, I am bidding farewell to all researchers and theatrists, delighted with the new horizons on the borderline of science and art that they will discover for humanity, for all of us.

*Tomáš Studeník
radical innovator and urban hacker
founding father of the THEaiTRE project*

Chapter 2

The Project

In this book, we describe THEaiTRE, an interdisciplinary project directly combining theatre and science, with the main aim of producing and staging the first computer-generated theatre play.

The initial motivation behind the project was to celebrate the 100th anniversary of Karel Čapek’s theatre play *R.U.R.* [Čapek, 1920], for which the word *robot* was invented and which was a play about robots, by turning the idea around and having a “robot” (actually, artificial intelligence)¹ write a play about humans.

Within the project, we have designed an interactive theatre script generation tool *THEaiTRobot*, and used it to produce scripts for two theatre plays, *AI: When a Robot Writes a Play* and *Permeation*. The plays were then staged, rehearsed and premiered in theatres.

The goal of the project is to explore the potential of current artificial intelligence techniques to be incorporated into theatre practice, and to directly confront the general public with the outcome while explaining the process behind the creation of the play and thus educating the audience about the current state and capabilities of the techniques used. We found that this immersive experience can spark a lot of interest both among artificial intelligence enthusiasts as well as among people generally ignorant to the current developments in the field.

Our goal here is to demystify AI, to explain the rather simple basic principles behind the machine learning based on textual training data, and hopefully to bring the AI closer to people so that they neither glorify it nor fear it irrationally, but rather understand its basic principles and thus set reasonable expectations and precautions, making it clear that AI should be taken with caution but can be very useful for some tasks in practice.

Within the project, we managed to generate theatre scripts of two plays, both of which were put on stage. We dedicate Part I of this book to the

¹While many researchers are reluctant to talk about *artificial intelligence*, often preferring other terms such as *machine learning*, we feel that many businesses and media do not share this reluctance, leading to a situation where a lot of publicly presented information about artificial intelligence is inaccurate and/or exaggerated. We therefore explicitly choose to use the term *artificial intelligence* when addressing the public, in an attempt to partially rectify the public understanding of artificial intelligence by conveying a truthful down-to-earth image of actual current artificial intelligence systems.

process of creating and staging the actual plays and to various theatrical aspects of the project. This part of the book is accessible to any reader.

The scripts for the plays were created using our online interactive script generation tool called THEaiTRobot. As the tool is one of the main results of this project, the whole Part II of this book is dedicated to the design of the tool and the research and experiments we performed. Most of that part of the book is quite technical, intended for a reader with some background in computer science and/or computational linguistics. However, a brief and easier to digest description of the tool is presented in Section 2.2.

2.1 The Two Plays

The first play created within the project is titled *AI: Když robot píše hru (AI: When a Robot Writes a Play)*. The script was produced by THEaiTRobot 1.0 operated by David Košťák. The play was directed by Daniel Hrbek, and premiered in the Švanda theatre in Prague on 26th February 2021. Due to the pandemic, spectators could not visit the premiere physically, but the premiere was streamed online and was viewed by approximately 20,000 spectators from all over the world. At the time of writing, the play is still performed at the Švanda theatre.

The second play, *Permeation (Prostoupení)*, was generated by THEaiTRobot 2.0 operated by Josef Doležal. It premiered as a stage reading, directed by Erwin Maas, on 12th June 2022 at the Rehearsal for Truth festival at Bohemian National Hall, New York. This was followed by a Czech premiere on 17th June 2022 at Švanda theatre in Prague, directed by Josef Doležal.

For both of the plays, 90% of the script was generated automatically using the THEaiTRobot tool; the remaining 10% are manually written parts provided as inputs (starting prompts) for the generation process and other manual post-edits of the scripts. We attach both of the scripts at the end of this book (Appendix A and Appendix B), with all human interventions marked in the scripts.

2.2 THEaiTRobot

The THEaiTRobot tool is intended to be used for generating theatre play scripts in an interactive human-in-the-loop fashion. First, the user (or *operator*) provides an input which sets the initial direction of the generation process. Then the tool takes over, generates a chunk of text based on the user input, and gives control back to the operator. The operator is then provided with several options of how to continue with the generation. It is thus an iterative process of automated generation guided by a human operator.

The tool is based on the large pretrained generative neural language model called GPT-2 [Radford et al., 2019]. GPT-2 had been trained on 40GB of English texts and has been shown to be very strong in generating new high-quality texts based on even a short initial prompt. However, as

most of the training texts were not theatrical scripts but rather news texts and other typically short texts with a simple structure, the main research challenge in our work is how to adapt the GPT-2 model to generating theatre scripts, which are rather long and complexly structured.

We have produced THEaiTRobot in two main versions: THEaiTRobot 1.0 and THEaiTRobot 2.0.

THEaiTRobot 1.0 [Rosa et al., 2021a] is designed for flat script generation – a textual *prompt* (in this case, the beginning of a scene) is given by the user and the model generates a continuation. The approach and models used for this variant are described in Chapter 9. This approach was designed and used to generate the script of the first theatre play; to generate the first play, we used THEaiTRobot 1.0 with a vanilla GPT-2 XL model, whereas the released version uses a GPT-2 medium model which we fine-tuned for theatre script generation. The flat approach is incapable of generating a whole theatre play script at once, as it is unable to keep the context throughout a long text; therefore, the individual scenes of the first play were generated independently and assembled into the final script *ex post*.

THEaiTRobot 2.0 [Rosa et al., 2022a] has a separate model for plot summary generation (described in Section 10.1) and a model for generating the script conditioned on the generated plot summary (described in Section 10.3); these models are used in succession, leading to a two-step hierarchical generation setup capable of generating the whole script at once. This was the approach developed and tested for the generation of the second theatre play. However, we faced various issues with THEaiTRobot 2.0, not fully meeting our ambitions and expectations (as discussed in detail in Section 3.2), and the decision was eventually made to resort back to the THEaiTRobot 1.0 style of flat script generation, using some models and techniques developed for THEaiTRobot 2.0 but not actually employing the two-step hierarchical generation process; thus, the second play was eventually generated by a sort of THEaiTRobot 1.5.

With the end of the project, we made a demo of both of the versions of the THEaiTRobot tool freely publicly available on the project website (see Chapter 12).²

2.3 The Team

The whole team of the project was based in Prague, Czechia.³ It was led by Rudolf Rosa, an expert on computational linguistics and natural language processing from Charles University (*CUNI*),⁴ who was in charge of the computational subteam at *CUNI*. The theatrical subteam was led by Daniel Hrbek from Švanda Theatre (*ŠD*),⁵ and consisted of members of *ŠD*, the

²<https://theaitre.com/demo>

³Except for the director and cast of the premiere of the second play, based in New York, USA.

⁴<https://cuni.cz>; or, specifically, its Institute of Formal and Applied Linguistics at the Faculty of Mathematics and Physics, <https://ufal.mff.cuni.cz>

⁵<https://www.svandovodivadlo.cz>

Theatre Faculty of the Academy of Performing Arts (*DAMU*)⁶ and the Tygr v tísni theatre (*TvT*).⁷ The project was supported by a facilitation and PR subteam from CEE Hacks (*CH*).⁸

While there was close collaboration among all of the team members, the subteams had clearly defined roles. The collaboration was strongest on designing the architecture and interface of the THEaiTRobot tool, where all of us combined our knowledge and experience to suggest various alternative ways of setting up the system. The computational part of the team was then responsible for devising, implementing and deploying the tool, while the theatre professionals were operating the tool, testing and evaluating it, analysing its outputs and providing extensive feedback. Based on that, we jointly suggested and implemented modifications of the design of the tool, leading to an iterative development and testing process. Once the tool was ready to produce the play script, the theatrical subteam was in charge of operating the tool and staging the play, including interpreting and rehearsing the script as well as designing costumes, music, decorations and lights. The premiere of the play was then carried out by *ŠD*, supported by *CH* who took care of marketing and PR, videorecording and streaming, and satellite events.

We are including a complete list of people who took part in the project, together with their main responsibilities within the project.

- Computational subteam
 - Rudolf Rosa (*CUNI*) – team lead, machine translation, demo
 - Tomáš Musil (*CUNI*) – flat generation
 - Ondřej Dušek (*CUNI*) – hierarchical generation
 - Patrícia Schmidtová (*CUNI*) – hierarchical generation
 - David Mareček (*CUNI*) – data resources
 - Alisa Zakhtarenko (*CUNI*) – machine translation, demo (from March 2021 till February 2022)
 - Tom Kocmi (*CUNI*) – machine translation, automated summarization (till August 2020)
 - Dominik Jurko (*CUNI*) – automated summarization (till July 2021)
- Theatrical subteam
 - Daniel Hrbek (*ŠD*) – subteam lead, theatre play production, director of the first play
 - David Košťák (*ŠD*) – THEaiTRobot operator and dramaturge for the first play
 - Josef Doležal (*DAMU*) – THEaiTRobot operator and director for the second play, dramatic situations annotation
 - Marie Nováková (*TvT*) – lead for dramatic situations annotation

⁶<https://www.damu.cz>

⁷<http://tygrvtisni.cz>

⁸<https://www.cee hacks.com>

- Klára Vosecká (*DAMU*) – dramatic situations annotation, THE-aiTRobot operator assistant for the second play
- Martina Kinská (*ŠD*) – dramaturge for the first play, theatre repeats production
- First play at *ŠD*
 - * Daniel Hrbek – director
 - * Jacob Erftemeijer, Denisa Barešová, Anna Stropnická, Andrea Buršová, Petr Buchta, Oskar Hes, Marek Frňka, Tomáš Petřík – cast
 - * Martin Šimek – stage design
 - * Paulína Bočková – costumes
 - * Martin Páček – choreography
 - * Vojtěch Nejedlý – assistant to director
 - * Blanka Popková – stage manager
 - * Tereza Marková – production
- Second play at Rehearsal for Truth
 - * Erwin Maas – director
 - * Jade Radford, Christopher Domig, Joshua Echebiri – cast
- Second play at *ŠD*
 - * Josef Doležal – director
 - * Nataša Bednářová, Marek Frňka, Patrik Děrgel – cast
- Facilitation subteam
 - Tomáš Studeník (*CH*) – subteam lead, mentoring, PR
 - Petr Žabka (*CH*) – facilitation, web, PR
- Support
 - Jana Hamřlová (*CUNI*) – administration, finance, communication (till March 2021)
 - Lenka Fišerová (*CUNI*) – administration, finance, communication (from April 2021)
 - Jaroslava Součková (*ŠD*) – administration, finance
 - Lucie Krůtová (*DAMU*) – administration, finance
 - Jindřich Vodrážka (*CUNI*) – technical support
 - Ondřej Bojar (*CUNI*) – mentoring

2.4 Related Projects

There has already been a range of partially or fully artificially generated works of art. We mention here only some of the most significant works and projects that use automatically generated scripts; we can see a clear boom of such projects in 2016, which is around the time when LSTM-based

neural language models became commonly used, considerably improving over previous approaches [Sundermeyer et al., 2015].

Unsurprisingly, most of the related projects are carried out in English language; we are unaware of any other project working with automatically generated scripts in Czech language. Therefore, we also list several Czech projects dealing with other types of generative art.

The tools used in these projects are typically not publicly available online, and often there is little transparency about the particularities of the exact design and usage of the tools. In our project, we opted for maximum transparency instead, making the THEaiTRobot tool publicly available online, releasing its source codes, describing the design of the system in several scientific papers, and publishing the generated theatre scripts with all human interventions marked.

Moreover, the other projects typically use substantial human curation and/or only produce short texts. The THEaiTRE project stands out by producing a full 1-hour play script with 90% of the script coming from the THEaiTRobot generator.

Beyond the Fence (2016)

Beyond the Fence, the “first computer musical” is, to the best of our knowledge, the first work of generative art using an automatically generated script. The project combined several tools which had been developed by various authors and institutions. The *What-If Machine* ideation engine and the *PropperWryter* plot generator suggested a concept and a plot, the *Claude-Machine Schönbot* and the *FlowComposer* music generating programs generated musical scores, and *Clarissa the Cloud Lyricist* generated lyrics. The members of the project gathered suggestions from the automated tools, post-edited them heavily or took inspiration from them, and compiled all of the contents into a full-length musical. The authors have estimated that approximately 25% of the final script is computer-generated.

The musical was staged and repeatedly performed in February and March 2016 at the Arts Theatre in London’s West End. This project also stands out by being accompanied by a detailed published scientific paper describing the tools and the generation process, as well as evaluating the project results [Colton et al., 2016].

HumanMachine and Improbatics (since 2016)

The *Improbatics* group⁹ has brought together theatre performers and artificial intelligence researchers from various countries and continents and features a range of theatrical projects, both past as well as ongoing, which make use of automatically generated textual content (as well as other technology) combined with improvisation.

The more recent projects of the group employ very similar technology to the THEaiTRE project, utilizing the GPT family of generative neural language models. A prime example is their engagement of a “cyborg actor”

⁹<https://improbatics.org/>

who says and acts out lines that are generated in real-time by an on-stage operator running a GPT-3 based system. The generated lines are sent through an automated text-to-speech system to an earpiece that the actor is wearing and said out loud by them. A second, fully human actor then improvises around what the cyborg says and does. The audience can then be subjected to a variation of the Turing test, trying to discern which of the actors is the cyborg.

The members of the group also frequently publish scientific papers describing their methods [Mathewson and Mirowski, 2017a,b, 2018]. The team has also already used machine translation live on stage [Mirowski et al., 2020], as well as incorporating live text-to-speech synthesis [Branch et al., 2021], both with automatically generated content as the input.

Sunspring (2016)

To the best of our knowledge, the *Sunspring* is the first work based on a script that was generated completely automatically, including stage directions, with no human intervention into the contents of the script. The script was then interpreted by the director and the actors, resulting in a short 9-minute experimental sci-fi movie (the actual part based on the generated script is 7 minutes long). The movie as well as the script are available online.¹⁰

Lifestyle of the Richard and Family (2018)

In 2018, Roslyn Helper wrote the theatre play *Lifestyle of the Richard and Family* with the help of a next word suggestion tool.¹¹ The human author started each of the sentences in the script, but some of the sentences were finished using SwiftKey Note. This was a tool that, based on already written text, offered the user with 3 options for what the next word could be, allowing to keep iteratively accepting these suggestions until the sentence is finished.

The project was devised and the resulting play was directed by Harriet Gillies.¹² The play was performed at the Meat Market for Next Wave Festival in Melbourne in May 2018. A short video containing excerpts from the play as well as comments by the project team is available on Vimeo.¹³

Poezie umělého světa (2016)

The Poetry of the Artificial World is a human-picked collection of computer generated poems [Materna, 2016], generated using simple recurrent neural language models trained on a collection of amateur Czech poetry. To the best of our knowledge, this is the first piece of Czech generative art.

¹⁰<https://www.thereforefilms.com/sunspring.html>

¹¹<https://www.roslynorlando.com/lifestyle-of-the-richard-and-family>

¹²<https://www.harrietgillies.com/lotraf>

¹³<https://vimeo.com/263802155>

Výsledky vzniku (2020)

Outcomes of Origin [Husárová and Panák, 2020] is a published collection of Slovak poems generated by the virtual author Liza Gennart, which is a GPT-2 model fine-tuned on published Slovak poetry. The poetry collection was awarded the “Zlatá vlna 2021” award for the best Slovak poetry publication of the year.¹⁴

Digitální filosof (2019)

The *Digital Philosopher* is a project by Jan Tyl and Dita Malečková, who fine-tuned the GPT-2 model on texts written by several philosophers, allowing the user to “talk” to a philosopher of their choice.¹⁵

From the Future World (2019)

Finally, we mention one work that does not use text generation, but a rather similar approach is used here for music generation, and it is also interesting for us due to its Czech setting. *From the Future World* is a symphonic musical composition generated in the style of the Czech music composer Antonín Dvořák, finishing one of his unfinished manuscripts [AIVA et al., 2019]. It premiered in Prague on November 15, 2019.

2.5 Successes and Publicity

In 2020, the THEaiTRE grant project proposal scored full points and ranked first out of 550 submissions within the TAČR grant agency funding call. In 2022, it was nominated for the best project of the faculty as well as university in terms of presenting the research done at the institution.

The project members have been invited as guest speakers or panellists to numerous events, including Anti-dystopian Congress,¹⁶ AI Narratives Workshop,¹⁷ Trienále SEFO (SEFO triennial) 2021,¹⁸ SETKÁNÍ/ENCOUNTER Festival,¹⁹ EUNIC AI Science Café,²⁰ Czech Science Café,²¹ Noc vědců (Researchers’ Night) 2020 and 2022,^{22,23} and many other.

There has been wide news coverage of the project by more than 100 newspapers, radio stations and television channels, both Czech and foreign,

¹⁴<https://www.litcentrum.sk/clanok/zlata-vlna-2021-pre-lizu-gennart>

¹⁵<https://digitalnifilosof.cz/en/>

¹⁶<https://www.goethe.de/ins/sk/de/kul/sup/adc.html>

¹⁷<https://www.ainarratives.com/lem-and-capek>

¹⁸<https://www.trienalesefo.cz/udalosti/program#ai-kdyz-robot-pise-hru-projekce-diskuse>

¹⁹<https://www.youtube.com/watch?v=NRxNDG1QTok>

²⁰<https://london.czechcentres.cz/en/program/eunic-ai-science-cafe-series-ai-and-theatre>

²¹<https://www.sciencecafe.cz/sc-rozhovor/science-cafe-rozhovor-2/>

²²<https://www.nocvedcu.cz/udalost/1787-theaitre-umela-intelligence-pise-divadelni-hru>

²³<https://www.nocvedcu.cz/udalost/5878-jak-umela-intelligence-napsala-divadelni-hru>

including e.g. The Guardian,^{24,25} The Telegraph,²⁶ The Times,²⁷ Science Magazine,²⁸ BBC Radio,²⁹ CBC Radio,³⁰ Deník N,^{31,32} Respekt,³³ Česká televize,^{34,35,36} Český Rozhlas,³⁷ and many others. The project and the first play were exclusively featured in the Taneční zóna magazine.³⁸

The project has an official website,³⁹ and the first theatre play is listed on the Švanda theatre website.⁴⁰ The project and the first play are also listed on Wikipedia.^{41,42}

We have published both of the generated scripts [THEaiTRobot 1.0 et al., 2021, THEaiTRobot 2.0 et al., 2022] as well as the source codes [Rosa et al., 2021a, 2022a]. We have also described the project in a diploma thesis [Schmidtová, 2022] and in several scientific articles [Rosa, 2020, Rosa et al., 2020, 2021b,c, 2022b, Schmidtová et al., 2022].

²⁴<https://www.theguardian.com/technology/2021/jan/10/art-meets-tech-to-mark-first-100-years-of-the-robot>

²⁵<https://www.theguardian.com/stage/2021/mar/01/on-the-scene-like-a-sex-obsessed-machine-when-a-robot-writes-a-play-ai>

²⁶<https://www.telegraph.co.uk/theatre/what-to-see/review-top-best-comedy-theatre-online-february-2021/>

²⁷<https://www.thetimes.co.uk/article/ai-when-a-robot-writes-a-play-review-an-eerie-glimpse-of-the-future-09b0ff509>

²⁸<https://www.science.org/content/article/kinky-and-absurd-first-ai-written-play-isn-t-shakespeare-it-has-its-moments>

²⁹<https://www.bbc.co.uk/sounds/play/w3csz99c>

³⁰<https://www.cbc.ca/radio/asithappens/as-it-happens-monday-edition-1.5931995>

³¹<https://denikn.cz/574569/normalni-dramaturgicka-prace-jen-se-autor-neurazel-kdyz-robot-napise-hru-a-svandovo-divadlo-ji-nastuduje>

³²<https://denikn.cz/547681/budou-psat-roboti-basne-ci-povidky-v-cestine-preconie>

³³<https://www.respekt.cz/tydenik/2020/46/ze-zivota-robotu>

³⁴<https://www.ceskatelevize.cz/porady/1097181328-udalosti/221411000100213/>

³⁵<https://ct24.ceskatelevize.cz/kultura/3270062-kdyz-robot-pise-hru-ma-nechteny-smysl-pro-humor-a-dost-mysli-na-sex>

³⁶<https://decko.ceskatelevize.cz/video/e221553117600001>

³⁷<https://vltava.rozhlas.cz/robot-napsal-divadelni-hru-jde-o-novy-smer-divadelni-tvorby-nebo-pouhou-atrakci-8434746>

³⁸Taneční zóna 1/2021 features a lot of content on the THEaiTRE project, including several interviews and articles by various authors, a sample of the script, and a full transcription of the discussion after the premiere. <https://www.tanecnizona.cz/tistena-vydani/01-2021/>

³⁹<https://www.theaitre.com/>

⁴⁰<https://www.svandovodivadlo.cz/inscenace/673/ai-kdyz-robot-pise-hru/3445>

⁴¹<https://en.wikipedia.org/wiki/Theaitre>

⁴²https://en.wikipedia.org/wiki/AI:_When_a_Robot_Writes_a_Play

Chapter 3

The Two Plays

Within the project, we managed to generate theatre scripts of two plays, both of which were put on stage. We dedicate this chapter to the process of creating and staging the plays.

3.1 The First Theatre Play

The first play is titled *AI: Když robot píše hru* (*AI: When a Robot Writes a Play*). The script was produced by an internal variant of THEaiTRobot 1.0, using the vanilla GPT-2 XL language model. The tool was operated by David Košťák. The play premiered in the Švanda theatre in Prague on 26th February 2021.

3.1.1 Generating the Play Script

Creating the script took the form of human-in-the-loop interactive generation. The human operator defines the input of the application, consisting of a short description of the starting situation and the names and first lines of two characters.¹ The model then generates a continuation consisting of 10 subsequent lines, and hands the control back to the operator. If they are satisfied with the output, 10 more lines are generated. If they are not satisfied, they mark any line as unwanted, the generated text starting with that line gets discarded, and the model generates a different variant. In addition, the operator has the option to manually insert their own line or scenic remark into a selected location.

During the project, we experimented with different variants of the script generation process. For example, the first version of the application worked with the model always generating a continuation of only one line, but in five variants, and the operator then chose which continuation to use (or discarded all continuations and had another five continuations generated). To this the model again responded by generating five more continuation

¹THEaiTRobot theoretically supports any number of characters, but in practice it turns out that it does not work very well with a larger number of characters, so we have chosen to limit it to a dialogue of only two characters.

options after that selected line. This method of generation was therefore more interactive, but it did not work well in practice – on the one hand, it was too time-consuming to generate anything, and on the other hand, the members of the theatre subteam felt that they had too much control over the generation, and therefore the resulting script was more of their work rather than that of artificial intelligence.

Generating 10 lines is a compromise between the speed of generation and the degree of control over the generation process. At the same time, it turned out that for an individual line, it was very difficult for the operator to decide whether it was good or not, while a longer section of the script is easier to judge and decide whether to accept or reject.

The final 1.0 version of the application was taken over by David Košťák, who, based on his experience and using his knowledge of the functioning of the application, defined the overall theme of the play (a robot in typically human situations) and prepared a set of approximately 12 thematic inputs (e.g. on the topics of *robot and love*, *robot and work*, *robot and death*, *robot and sense of humor*). For these inputs, he then interactively generated scripts. From the generated scripts, the ŠD dramaturgy team then selected 8, from which they assembled the script of the play, consisting of 32 pages of text. The theaterists then made further adjustments to the text as if it were a standard script of a play written by a human author, but tried to minimize these adjustments in order to preserve as much authenticity as possible.

The resulting script of the play *AI: When a Robot Writes a Play* consists of text generated by the THEaiTRobot tool from 90%, the remaining 10% are human interventions by the operator and dramaturgist David Košťák, or by the director Daniel Hrbek. A part of the interventions consists of inserted inputs (for each scene, a description of the starting situation and the first two lines were manually inserted), other interventions consist mainly of selecting a variant other than the first generated variant and of minor post-editing of the generated outputs. A detailed analysis was published in the form of a technical report and is attached in the Appendix A. The generation took place in English with an automatic translation into Czech; the rate of 10% of human interventions is related to the English version of the script. The automatic translation introduced some errors into the text which were manually post-edited, without changing the meaning of the originally generated English sentence.

When working with the initial versions of the application, the ŠD team was rather disappointed by the low quality of the outputs and planned to stage the play in a less expensive form of one-off stage reading without decorations and costumes and without further performances. However, through the gradual development of the application and also through the gradual closer familiarization of the ŠD team with the application and a better understanding of its functioning and its strengths and weaknesses, it was possible to achieve a state where the ŠD team created, according to their own opinion, a sufficiently high-quality script using the application. The ŠD team therefore eventually decided to stage the play in the classical way, with a full stage production, and it is performed at the theatre approximately once a month.



Figure 3.1: The premiere of the first play, *AI: When a Robot Writes a Play*. Photos by Alena Hrbková.

3.1.2 The Staging

The script was subsequently rehearsed in a standard manner by the Švanda Theatre ensemble, directed by Daniel Hrbek. The production includes stage design, music, costumes, etc.; this is all created by human professionals of the Švanda Theatre without the intervention of artificial intelligence. The actual play lasts approximately 60 minutes, consists of 8 scenes and features 6 actors.

3.1.3 The Premiere

The premiere of the play (see Fig. 3.1) was streamed online and was watched on a total of 18,450 devices in the Czech Republic and abroad. Such an interest greatly exceeded our expectations; the form of the online stream thus made it possible for the play to be viewed by several times more spectators than the physical performances at the Švanda Theatre will ever allow, which now take place approximately once a month.

The premiere was streamed in Czech language with English subtitles. It was also automatically translated and subtitled in 40 languages through the Elitr project.²

A short showcase of the premiere and a full recording of the discussion after the performance are available on YouTube.^{3,4}

3.1.4 Format of the Show

An integral part of the performance (as well as other related organized events) is the expert discussion that follows the play. The discussion always involves at least one representative of *ŠD* and at least one representative of *CUNI*, who first introduce the project, and then discuss with the audience and answer their questions.

In doing so, we are fulfilling one of the important goals of the project, which is to bring closer and explain artificial intelligence to people. The experience from the discussions shows that this is largely successful, the discussions tend to be stimulating and factual, and an increasing level of understanding can be observed in the participants during the discussion. It is interesting that some participants unequivocally rate the discussion as the better part of the evening than the play itself; which is not very flattering for the play, but if the play attracts people to a professional discussion about artificial intelligence, which probably most of them would not otherwise have come to, and from which they will gain some understanding, then we are succeeding at fulfilling the deeper goals of the project.

3.1.5 The Reactions

The reactions of both the professional and the general public are somewhat mixed (see the articles referenced in Section 2.5). It is obvious that

²<https://elitr.eu/theatre/>

³https://youtu.be/8ho5sXiDX_A

⁴<https://youtu.be/WyRpKo0BTaI>

the script is not entirely logical, much is inevitably left to the spectator's interpretation, and in some cases it is probably very difficult or impossible to find a satisfactory interpretation. However, the fact that this result of the project can actually be viewed as a full-fledged play and thus evaluated, is considered to be a great success by our team. The reviews generally note the imperfections of the AI author, but in general, the quality of the play seems to be in some sense comparable to the quality of plays by a below-average or bad human author, which from our point of view can clearly be evaluated as a success.

Some reviewers also note that the play contains many common human biases and stereotypes [Akbar, 2021]; this is in line with existing findings of large language models often being e.g. racist or sexist [Gehman et al., 2020]. The general cause is the fact that the underlying training datasets (articles downloaded from the internet with light or no filtering) themselves contain such biases and other problematic content, and the language model merely truthfully reproduces those. While research on debiasing language models exists and is regarded to be of great importance [Schick et al., 2021], this is by far not a solved problem yet. As we intended to give a truthful image of the current state of the technology, we intentionally kept some of this problematic content in the final script (of course, the contents of the script do not necessarily express our opinions). However, we make sure to point to these issues and explain them in discussions after the performances and at other public events related to the project.

The script of the play is freely available on the project's website, allowing anyone to stage the play in their own way. At the time of writing the book, we are aware of one such case – a student director Tim Stueve staged and rehearsed a short portion of the script with two actors (Christopher Aceves and Scott Berg) and posted a recording of the performance on YouTube.⁵

3.2 The Second Theatre Play

The second play, *Permeation (Prostoupení)*, was generated by an internal variant of THEaiTRobot 2.0 using a flat script generation approach. The tool was operated by Josef Doležal. The play premiered as a stage reading on 12th June 2022 at the Rehearsal for Truth festival at Bohemian National Hall, New York.

3.2.1 The Generation Process

For the second play, our initial aim was to use the THEaiTRobot 2.0 hierarchical generation approach, so that the script can be produced at once, unlike the first play where individual scenes had to be generated independently and assembled into the full script ex post. A further hope by the project team was that generating from a fixed synopsis would allow us to further minimize the required amount of human intervention into the generation process.

⁵<https://youtu.be/6gF137Udw3g>

The operators thus experimented for several months with using the hierarchical approach, repeatedly generating several candidate synopses using the first stage of THEaiTRobot 2.0, selecting one to expand into a script, and generating a script based on the synopsis using the second stage of THEaiTRobot 2.0.

However, we both succeeded and failed at the same time. We have observed that this approach does indeed enable the operator to generate the full script at once. The points at which the operator moves on in the synopsis by pushing one or several sentences from the generated synopsis into the script serve as a sort of a soft restart of the generation process, pushing it in a new direction (while, sadly, often disregarding some of the previously generated lines; or, in other cases, getting back to seemingly random previous lines in an undesired way).

However, the approach failed terribly at our second goal of further minimizing the human impact. Unsurprisingly in retrospect, it turns out that as the script generation stage is unaware of the synopsis lines to come, it tends to diverge in ways wildly inconsistent with these future synopsis lines. Therefore, to tame the generation process and force it into generating a script consistent with both the past as well as the future parts of the synopsis, significantly *more* human intervention is actually required – while in the first play, the operator had a line regenerated once for every 15 lines on average (so approximately 0.06 regenerations per line), now the operator had to regenerate most lines and often more than once, leading to approximately 2 regenerations per line on average (so about 30 times more). He also needed to manually insert new lines into the script much more often, further increasing the amount of human written content in the script.

To make matters worse, although explicitly fine-tuned to expand the synopsis line into script lines, the second-stage model still often tends to treat the synopsis line as something that has just happened instead of something that is about to happen, and wanders off to generating lines that happen after what the synopsis line describes. In other cases, it does indeed expand on the synopsis line, but often tends to only expand the synopsis line partially, implementing some of its aspects in the first few generated lines but ignoring the yet unimplemented parts of the synopsis line in later generated lines.

After a set of vivid discussions among the whole team, we have concluded that these findings are interesting and valuable scientifically as live analyses and evaluations of the process, but fail to meet the goals we had set for the project. A decision was made that for the second play, the amount of human interventions should not be greater than in the first play, as one of the unique features of the THEaiTRE project is the low amount of human intervention needed to generate a full-length theatre play script; generating a script with only a limited contribution of artificial intelligence has already been explored in several pre-existing projects, as discussed in Section 2.4, and we would thus not actually contribute anything significantly novel.

Therefore, the approach eventually employed to generate the second play script is closer to the one that had been used for the first script, generating individual scenes in a flat way from a prompt defining the initial setting, character names, and optionally their first lines. From the technical point of view, the main difference is that instead of using the vanilla

GPT-2 XL model, we used the GPT-2 medium model fine-tuned to generate theatre play scripts; in this respect, we explore how fine-tuning a smaller model on in-domain data compares to using a larger vanilla model. From the operational point of view, the main difference is that the scenes were not generated independently to be later selected from and reordered, but rather in a continuous way, with the operator partially simulating the idea of the hierarchical approach: after generating a scene, the operator contemplated how the play might continue, and then defined the initial setting and the characters for the next scene based on the already generated scenes, often partially summarizing the previous scene in the prompt as well. This approach has proven to be successful, reaching an amount of human interventions comparable to the first play, and eventually motivating our future work suggestion of iterative generation and summarization (generate a scene, summarize it automatically, generate the next scene from a prompt containing the summary of the previous scene, and iterate this until a full script is generated).

In the final script, the amount of human intervention is very similar to the script of the first play, with approximately 90% of the text coming from the THEaiTRobot generator.

3.2.2 The Script

The final script of the play, enclosed as Attachment B and made publicly available at the project website, is titled *Permeation (Prostoupeni)*. It revolves around the adventures of a couple fleeing their home because of an approaching war, and includes the topics of losing and looking for a loved one, espionage and treason, politics and patriotism, revolution and assassination, and some further minor themes. While the themes stand out quite clearly from the script, the details of what actually did or did not happen and why are often quite blurred, mostly due to the insufficiencies of the generator.

3.2.3 The Staging

As *ŠD* did not see such a potential in the second play as compared to the first play, it was decided to only stage it in the form of a staged reading, which requires only little rehearsing and its costs are very limited. Although the play features 7 characters, the script is composed in a way that only requires three actors (one female and two males). Apart from the dialogues, the script contains very little scenic and staging remarks, giving the director and the actors a considerable amount of interpretative freedom.

3.2.4 The Premieres and Reactions

The play premiered internationally on 12th June 2022 in New York as the opening event of the Rehearsal for Truth festival at Bohemian National Hall (see Figure 3.2).⁶ In the discussion after the premiere, the director

⁶<https://www.rehearsalfortruth.org/program/permeation>



Figure 3.2: The international premiere of the second play, *Permeation*. Photo by Rudolf Rosa.



Figure 3.3: The Czech premiere of the second play, *Prostoupení*. Photo by Rudolf Rosa.

Erwin Maas reported that he found the script to contain a lot of nonsense and contradiction, not unlike the theatre of the absurd, but found it highly emotionally loaded. Therefore, with the actors they decided to go “beyond the text”, disregarding various issues with the text meaning and rather focusing on the feelings and pace and the general flow of the story without attention to details; even with multiple characters speaking at once at some escalated moments, making the actual texts partially unintelligible and thus stressing its partial irrelevance.

The chosen approach has presumably succeeded, as some of the spectators reported the experience as highly emotional, even if they found it unclear what the actual plot was, only having a rather vague feeling about it. While the project team members dismissed these issues as a clear sign of the many shortcomings of the THEaiTRobot tool, several audience members actually countered by claiming that they did not find this fuzziness of the plot disturbing, as they found it to be well in line with the war/espionage/conspiracy setting of the play. They further reported that they often feel the same about similar events in the current real world, as due to a high amount of misinformation, fake news and propaganda, one is often faced with both incomplete and contradictory information about the events, unable to tell what the truth is. Thus, while the first play managed to circumvent the disconnection from logic and the real world by using a science fiction setting, it seems that the second play has also succeeded in this respect, even if inadvertently, this time by choosing a theme prone to misinformation.⁷

On 17th June 2022, a shortened version of the play premiered in Czech at Švanda theatre under the direction of Josef Doležal (see Figure 3.3). The Czech team opted for a more down-to-earth interpretation of the script, with the actors truthfully reading out the lines without adding much emotion. However, in the shortening of the text, many of the nonsensical parts were removed, resulting in a script which made considerably more sense than the original full-length English version.

No further performances are planned for the second play.

3.3 Further Projects

At the time of writing, it seems that two further plays may be generated by the tool, somewhat beyond the scope of the original project.

Tomáš Studeník and Josef Doležal, joined under the PLai Prague project,⁸ have decided to try to generate and stage a play at the occasion of the Czech Presidency of the Council of the European Union in the second half of the year 2022. The play, titled Prague Clockwork Cabaret, shall premiere on 9th October 2022 and feature performances in several countries within European Union.

⁷We would like to note that the premiere of the second play happened during the Russo-Ukrainian War, only several months after Russian invasion of Ukraine in February 2022, which provided both inspiration when generating the play as well as a common context for the theatrists and spectators under which the play can be viewed.

⁸<https://www.plaiprague.eu/>

We have also been contacted by a Taiwanese theatre director interested in generating and staging a theatre play with theatre students of Taipei university of arts; however, at the time of writing, this potential collaboration is only in its earliest stages.

As we have made the THEaiTRobot tool publicly available, we are encouraging such follow-up projects. The usage of the tool is free for non-commercial use, and we are willing to negotiate reasonable conditions even for its commercial use.

Chapter 4

The Situation with Dramatic Situations

Traditionally, the dramatic situation is considered to be the building block of drama.

Based on analyses of the shortcomings of the scripts generated by the first version of THEaiTRobot, the theatrical subteam suggested that the generation should operate on the level of dramatic situations instead of individual lines.

As we have not been able to find any computationally usable data resource containing annotations of dramatic situations in scripts, we decided to create such a corpus, which could serve for the training of tools for the recognition and generation of specific dramatic situations. The theatrical experts, supported by *CUNI*, managed to create a corpus consisting of 52 play scripts in Czech language with marked dramatic situations; devising a list of dramatic situations to use was also an inherent part of this process. The team of annotators consisted of Marie Nováková (*TvT*) as the lead, Klára Vosecká (*DAMU*), and Josef Doležal (*DAMU*); each member selected plays to annotate on their own, but the repertory of situations to use was discussed and agreed upon by the whole team.

Subsequently, we have tried to use the resulting corpus for training of models that would automatically identify dramatic situations in scripts, but these attempts were basically unsuccessful. We think that for a successful training, it would be necessary to annotate thousands of plays, which exceeds the possibilities of the project by far.

From the point of view of the main objectives of the project, this is one of the unsuccessful experiments. However, we consider the partial outputs that were created in this way (a list of situations and their experimental annotation on dozens of plays) to be of potential interest for further researchers. At the time of writing, we are working on publishing as much of the resulting corpus as possible, so that it can potentially be used by other researchers from the theatrical and computational linguistics fields.

4.1 Dramatic Situations

A dramatic situation can be defined as a situation that leads characters to necessary and immediate action.

For our needs, we are mostly interested in classifications and abstractions over theatre play scripts or their parts. In the field of theatre studies, there is a vast amount of research on the structure and interpretation of theatre plays. Unfortunately, the results of such research are not made available in forms and formats that would easily allow us to use these as data and annotations in machine learning approaches.

The Thirty-Six Dramatic Situations by Polti [1921],¹ originally proposed in 1895, is a classic work which is highly respected in teatrology. The author presented a supposedly ultimate list of all categories of possible dramatic situations that can occur in a theatre play (e.g. “adultery” or “conflict with a god”), further subclassified into 323 situational possibilities. While being a good starting point, this list seems somewhat outdated for today’s plays, as Polti based his list primarily on an analysis of classical Greek texts, as well as some classical French works.

Although not directly related to theatre plays, the work of Propp [1968] is also essential. Propp analyzed Russian folk tales and identified 31 *functions*, similar to Polti’s situations but somewhat more down-to-earth (e.g. “villainy” or “wedding”), as well as 7 abstract character types (e.g. “villain” or “hero”) and other abstractions.

Polti’s and Propp’s categorizations are sometimes used in analyzing and generating narratives, although typically not in drama. The works closest to our focus is probably that of Gervás et al. [2016] or Lombardo et al. [2018], who devised an ontology of abstractions for annotating scripts, based on both of the mentioned works, as well as on more recent plot categorization studies [Booker, 2004, Tobias, 2011].

There are also works producing drama analyses in the form of networks, capturing various relations between the characters in the play [Moretti, 2014, Horstmann, 2019, Fischer et al., 2019].

4.2 Annotation Process

The annotation process consisted of continual joint annotation of play scripts with dramatic situations and refinement of the repertory of situations to use for the annotation.

4.2.1 First Phase: 58 Situations

We took Polti’s 36 situations as the starting point. While trying to use them for annotating (mostly) contemporary theatre plays, we have quickly managed to confirm our suspicion that the list needs updating to suit today’s plays, as it contains many situations that are highly unlikely to appear in a modern dramatic text (such as *Involuntary crimes of love*) while missing

¹https://en.wikipedia.org/wiki/The_Thirty-Six_Dramatic_Situations

many situations which contemporary authors use frequently (e.g. *Interrogation*).

In the first phase of annotation, in which we annotated 19 play scripts, we removed several obsolete situations from the list and introduced many new situations that might not be as archetypal as Polti's, but better reflected the kind of texts we were working with.

Eventually, we have arrived at a list of 58 situations. We are still perfecting the list, trying to merge some of the situations so that each situation is sufficiently supported by the annotated data. Our repertory of situations is thus based on a combination of teatrological knowledge and statistical data processing.

Our current list of 58 situations is as follows; situations 1-34 are taken over from Polti, i.e. we have not removed many of Polti's situations yet:

1. Supplication
2. Deliverance
3. Crime pursued by vengeance
4. Vengeance taken for kin upon kin
5. Pursuit
6. Disaster
7. Falling prey to cruelty/misfortune
8. Revolt
9. Daring enterprise
10. Abduction
11. The enigma
12. Obtaining
13. Enmity of kin
14. Rivalry of kin
15. Murderous adultery
16. Madness
17. Fatal imprudence
18. Involuntary crimes of love
19. Slaying of kin unrecognized
20. Self-sacrifice for an ideal/for kin
21. All sacrificed for passion
22. Necessity of sacrificing loved ones
23. Adultery
24. Crimes of love
25. Discovery of the dishonour of a loved one
26. Obstacles to love
27. An enemy loved
28. Ambition
29. Conflict with a god
30. Mistaken jealousy
31. Erroneous judgment
32. Remorse
33. Recovery of a lost one
34. Loss of loved ones
35. Admission
36. Intruder

37. Bad news
38. Humiliation
39. Rape
40. Murder
41. Unfulfilled desire
42. Break up
43. Intimidation
44. Breaking the taboo
45. Passing
46. Seduction
47. Fight
48. Succumb
49. Parting
50. Reconciliation
51. Revelation
52. Curse
53. Blasphemy
54. Ruse - Trap - Fraud
55. Fear / Bad premonition
56. Interrogation
57. Accusation
58. Betrayal

4.2.2 Second Phase: 5 Situations

To make the annotations both more efficient and more useful, in the second phase, we decided to focus on a few frequent situations which we found to be the easiest to decipher in the scripts as the annotator agreement was high on them: *Supplication*, *Intruder*, *Seduction*, *Parting* and *Interrogation*.

We annotated further 33 play scripts with only these 5 situations; approximately 20% of each script can be categorized into these five situations, leaving 80% of each script unmarked on average.

4.2.3 Annotator Agreement

Two plays were annotated by all three annotators so that we could evaluate the inter-annotator agreement. Dramatic situations are annotated by marking the beginning and the end line of the situation in the script. For computing the annotator agreement, each line of the script either belongs to one specific situation or does not belong to any situation. The inter-annotator agreement is then computed as the number of lines in the script that were annotated with the same situation by all the three annotators divided by the total number of lines.

If we measure agreement on the full set of 58 situations, the agreement among all three annotators is 37.1%; the agreements between pairs of annotators are 47.7% on average.

If we only measure the annotator agreement on the 5 selected situations, the average agreement among a pair of annotators is 69.7%; Table 4.1 details the annotator agreement for each of the 5 selected situations.

Situation	Agreement
Interrogation	98%
Intruder	72%
Parting	63%
Supplication	63%
Seduction	59%

Table 4.1: Average pair-wise annotator agreement on the 5 selected situations.

We have thus found that recognizing a dramatic situation is difficult even for humans, as the annotators often could not agree with each other and their readings differed. However, some situations are clearly easier to identify than others; only focusing on these highly reliable situations should thus lead to a more consistent annotation. Nevertheless, even for the 5 selected situations, the annotator agreement is not very high.

4.3 The Annotated Corpus

We eventually annotated 52 play scripts in Czech language (consisting of 625,739 words in total) with a set of dramatic situations. The plays were annotated by three annotators who are experts on drama and who selected the plays for annotation themselves, typically plays they knew well or even authored.

The annotations then have been sanitized and converted to a JSON data format so that they are easy to process and can be utilized for further experiments.

We will make the annotations of the dramatic situations publicly available in the maximum possible extent and at the most permissive licence possible, as it seems that such a corpus is not publicly available, even though it is useful for both computational and theatrical researchers. We are limited by copyright, but we are actively negotiating with the copyright holders, trying to gain the rights to publish the annotated plays. At the moment of writing, we have been able to publish 9 of the annotated play scripts under the CC-BY licence [Mareček et al., 2022].

4.4 Experiments in Predicting the Situations

The purpose of creating these annotations was to use them to generate scripts, but unfortunately this goal could not be fulfilled within the project. The intended procedure consisted of the following steps:

1. Get annotations of dramatic situations on existing scripts (fulfilled).
2. Train a system for automatic recognition of dramatic situations on the annotated data (fulfilled only partially and not satisfactorily).

3. Find individual types of dramatic situations in other non-annotated scripts using the tool (not fulfilled).
4. Tune different versions of models to generate given types of dramatic situations (not fulfilled).
5. During the generation of the script, explicitly choose the dramatic situation to be generated and generate it with the appropriate tuned model (not fulfilled).

Unfortunately, we already encountered insurmountable problems within the project at the second point, so we did not perform any further steps.

In preliminary experiments, we first simplified the task of detecting a dramatic situation by entering a section of the script that we know corresponds to a situation and the goal is to assign it. The task is therefore simplified by skipping the search for the beginning and the end of the situation (i.e. the boundaries of the situation).

On the created data, the trivial solution (always assigning the most common type of the situation, which is *Seduction*) has a accuracy of 26%. To solve the task, we tried to use various basic and advanced methods of data processing and machine learning (tokenization, bag of words, stopwords, anonymization, leave-one-out filtering, frequency-based filtering, TF.IDF, word embeddings, Naive Bayes, Support Vector Machine, Deep Neural Network and others), we tried many variants of the procedure, but the highest accuracy we were able to achieve in this way was 44% with a multi-layer perceptron classifier.²

Such accuracy is too low, the tool thus returns the wrong answer in most cases. Moreover, it is the accuracy on an artificially simplified task with knowledge of the boundaries of the situations, which is impossible to assume for the actual data. Therefore, part of the task must also include searching for the boundaries of the situations, which would reduce the accuracy further, probably very significantly.

One of our hypotheses for the bad result was based on the observation that the annotators often characterize a relatively long section of the script as a dramatic situation. However, in such cases, a part of the marked section is rather a kind of lead-in for the given situation, from which the actual situation is not yet well recognisable, and only at the end of the situation does it gradates to such extent that it is obvious which situation it is (typically the dynamics are rising in this way, with the exception of the *Intruder* situation which, on the contrary, is typically the strongest at its beginning, when the intruder appears). That is why, together with the *DAMU* team, we carried out a further revision of the annotations, where the annotators added emphasis on the core of the situation in which the situation is strongest. Subsequent experiments using such enriched annotations, however, only led to a slight improvement in the results.

²The *Intruder* situation was the easiest to classify (0.83 AUC), followed by *Seduction* and *Parting* (0.74 AUC), *Supplication* (0.73 AUC), and *Interrogation* (0.68 AUC).

4.5 Discussion

We believe that the task of identifying dramatic situations is considerably more difficult than we originally estimated. We have relied on theatrical theories and insights which say that a play takes place in dramatic situations, that the author, the director and the actor think in such situations, and thus it is a concept which is well established and usually universally understood and shared. Until now, however, it has been more of a tradition based on theory and introspection, which has never been empirically verified on a larger number of scripts and methodically evaluated.

Experiments carried out by us have shown that even if limited to 5 selected situations, the paired inter-annotator agreement is only 70%, which suggests that the concept of a dramatic situation is not nearly as straightforward and obvious as the theory claims.

Nevertheless, the experiments in automatic identification of the situations fall far below this theoretically achievable limit. We conclude that this is mainly due to the too small amount of available training data, and we assume that for a task that proved to be so difficult, a significantly larger amount of data (certainly at least ten times, probably more than 100 times the current data) would be required, which is not feasible within this project.

For the generation of the second play, we have therefore decided not to use the concept of an explicit dramatic situation. Nevertheless, we still consider the concept of dramatic situations to be substantial and potentially useful, but we believe that for any meaningful automatic work with them it would be necessary to create significantly larger annotated data.

Chapter 5

Is the Author Dead?

It has been 100 years since the first production of Karel Čapek's play. In that time, the world has changed, computer technology has advanced, but the perception of what a play can look like has also fundamentally changed. Postmodernism has brought with it the familiar prism by which even a telephone book can be staged on stage. Even so, the THEaiTRE project originally set out to generate a classic drama, the basic building block of which is the dramatic situation. The author was to be an artificial intelligence (the GPT-2 language model) and Karel Čapek was to be the object of the work.

A dramaturge with his own author's experience is bound to faint at such an idea. On the one hand, skepticism, fear of hybris and the possible discovery that even a man's rationally hard-to-grasp creativity can be replaceable, on the other hand, the curiosity of discovering tools that can be made available to a living author, especially given his previous own experience of various creative writing courses that lead more to a technician perception of the creative process than to individual originality.

5.1 First Encounter of a New Kind

Despite all the expectations that are hard to resist for any fan of the science-fiction genre, the first contact between a carbon-based author and an artificial intelligence took place in the familiar environment of an internet browser. Not a seemingly humanoid shell, but a simple interface reminiscent of obsolete literary web servers. The language model requires a human input – a brief scenic remark and the first few lines of dialogue.

The robot walks with its creator in the park. They meet a dog holding its leg up and peeing on the robot's wheels.

Robot: I think the park is full of rodents this afternoon.

Creator: It's not a rodent. It's man's best friend.

Robot: Am I a man's friend, too?

The first input by the dramaturge tests not only the robotic author's creative abilities, but also his sense of humor and basic emotional sensitivity

to basal relationships. The basic circumstances are a dog threatening the tin existence in a way, and the input leads to the theme of friendship. The robot is creating:

Creator: No! Don't be rude to your human!

Robot: Hey, man, what's your name?

Creator: My name is the creator of the Robot. Do you like the Robot?

Robot: I'm sure I do. Do you like it here?

Creator: Oh, Robot, I'm so sorry. Robots never like it here.

Robot: I don't like it here either. The robot wants to live somewhere else. Do you want the robot to leave?

Creator: Oh, sorry, Robot. You can leave. The robot is leaving.

Robot: What? What? No. This is my home! Why don't you leave me alone?

Creator: There's nothing you can do about it. The robot has decided to leave this earth.

Robot: But why?

Creator: Robot can't live on a planet with the same laws as this one. There are many robots who want to live on other planets and wouldn't want to be ruled by this one.

Robot: Why?

Creator: Robots would kill each other if they saw that this one killed anyone else. So Robot doesn't want to live with humans.

Robot: But I have to go home! I miss my friends. Why don't you just leave me alone?

Creator: Because if a robot is allowed to live with humans, it will think like them and you will eventually lose your human friend. It's a very real threat.

Robot: But I want to go home!

Creator: The only way is to kill Robot and live in the dark.

Robot: But I want to go home!

Creator: Yes, you know. Do you know the way home?

Robot: Yeah!

Creator: Robot has been traveling around the galaxy for a thousand years.

Robot: Are you serious?

Creator: Yes. But you're not the first one to be on this planet.

Robot: I don't want to be the last one!

Creator: But you've been traveling this way for a thousand years and you've seen the stars.

Robot: How do you know I don't see them?

Creator: Yes.

Robot: But that's impossible!

Creator: Yes, it is.

Robot: You're just a robot! How can you tell?

Creator: Because I'm a robot.

5.2 Shortcomings

Dramaturg gets to better know the author with whom he'll be collaborating in the coming months, and is discovering his main shortcomings.

5.2.1 Spatial Imagination

The generated dialogue works with input lexical motifs, builds on lines based on word associations, but ignores external dramatic circumstances. The presence of a dog is soon irrelevant, as is the place where it takes place. In other words, the dialogue between the two characters could take place at any place and at any time. But it's not just what is said out loud that determines the dramatic text. The dialogue of the robot lacks any idea of the characters' non-verbal actions and of a coherent fictional world.

5.2.2 Schizophrenia

Formally, the text resembles a play, but it still retains the literary essence of the language model. The lines of the characters build on each other in the context of the motifs used, but they lack an awareness of the individuality of the individual characters. Not only do they speak a unified, slightly machined language, but they gradually take on each others specifics. The creator becomes a robot, and the robot becomes a human. This sounds like the theme of an existential drama, but it happens unintentionally. In attempts to generate a dialogue of three or more characters, these tendencies towards "schizophrenia" multiply.

5.2.3 Absence of Dramatic Thinking

A dramatic situation is a situation in which external or internal circumstances place the characters in a situation that is unsustainable in the long term and forces them to act. Thus, it directly shifts the action. The computer generates dialogues that bear no sign of action. The external circumstances that make up the dramatic situation pass and form non-dramatic dialogues that are merely exchanges of information. At best, the situation is reported without taking place. In this respect, the texts strongly resemble the instant dialogues of endless television series.

5.2.4 Staticity

The robotic text ignores not only the physical actions of the characters, but also their cohesion or psychological development. We learn about the characters only from the lines, which moreover often contradict each other. Once the creator refers to himself as a human, later as a robot, later in the text he returns to his human nature. Thus, the language model does not write characters that are undergoing development, rather it keeps changing his individuality in a volatile way. When generating longer texts,

in most cases there is even an endless loop of two lines. It is as if the robot has exhausted the potential of the dialogue.

However, nothing is black and white. The aforementioned shortcomings, together with the robot's ignorance of human psychology and the external real world, allow the creation of a text that is not limited by conventional realistic thinking. The generated text surprises the dramaturge by its improbability and outside the box thinking, which can be considered as originality. In this respect, the robot's dialogues are closest to the theatre of the absurd. This is also characterized by the absence of continuous action and rendering of the characters and by the use of meaningless language. However, behind a real absurdist play, there is always the human author whose intention can be analysed in the text and then thematised in the staging. The theme then serves as a kind of focal point through which the text can be interpreted. In robotic texts, this intention is missing. Perhaps this is because while an author of flesh and bones writes because of a certain own need (which is not merely existential), the robot writes because it is required to do so. While the actors who interpreted the text often mentioned that it lacks a soul, the dramaturge particularly misses the author's intention.

5.3 Search for an Achievable Goal

After evaluating the performance and (with a little exaggeration) the "talent" disposition of the robotic author, the dramaturge revises the assignment. An original classical drama about Karel Čapek from the computational pen seems to be a difficult goal to achieve so far. Especially taking into account the lack of dramatic thinking, the tendency to looping, not keeping the idea in a longer text and the recurrent schizophrenia problems with a larger number of characters. A new goal is sought. A reachable one this time. Entry requirements are relaxed. The dramaturge will develop short dialogues of two characters with the author, which will eventually be composed by the living author into a larger whole. The unifying element will be the character of a robot, who is confronted through human characters with something purely human, namely the boundary moments of life such as birth, love, loss of a loved one etc. The inanimate hero will be exposed to situations associated with excessive emotionality. Due to the nature of the project and the unconventional thinking of the author, the produced texts will naturally move towards the genre of science fiction.

The input data, entered by the live co-author to breathe in inspiration, must also be adapted to this. The starting lines are clearly confrontational in order to define the relationship of the two characters. The input should already deviate from the realistic perception of the world so that the text is consistent. The generated dialogues have not been able to satisfy the psychological demands.

It's midnight. The street is empty apart from one robot standing under the street lamp. He is trembling yet he cannot move. Stranger walks past him.

Stranger: What are you doing here? Why don't you go home? It's late. Robots like you should already be in their plugs.
 Robot: I know. But I can't move away from under the lamp. I'm afraid of the dark.

5.4 Author without Ego

Working with a living author is, to some extent, a psychological task. The dramaturge tries to read between the lines and direct the author so that their intention is perfectly imprinted in the work. Ideally to the extent that it is clearly legible even for the reader and the viewer. The author is often not the best reader of their own work. And at this point the dramaturge comes as its first recipient. In an attempt to fulfill the author's intention completely, he arrives at comments that need to be reflected in subsequent versions. Some playwrights resist such changes more, some less.

The robotic author does not resist at all. Behind each single line, instead of the ego, there is a "x" sign, which serves as a mediator of feedback. If the dramaturge does not like where the dialogue is going, he threatens to press this button, and the text is regenerated from that moment on. The robotic author "discards" the existing version and replaces it with a new version. Without any attempt to defend, without contradiction, simply without any argument, and thus without dialogue. Such cooperation may seem ideal at first glance. But the truth is that the dramaturgical participation in the text is conducting a dialogue with the author, during which the text moves from the first version to subsequent versions, which deepen the previous text and further develop it. Instead of the unsuitable first version, the robotic author offers another – a new first version of the text. Thus, the dramaturge, who prohibits himself from using the ultimatum evaluation expressions "good" and "bad" in theatrical communication, is led to a dead end.

In simple terms, working with a live author is based on a qualitative approach – dealing with detail, finetuning lines, perfecting specific points of dialogue – whereas working with a writing robot is quantitative – from a large number of dramatic attempts, the suitable ones are chosen.

The idea that generating the text by a computer leads to a faster authorship process thus crumbles. While a living author must "brood" his script, with the clear aim of what he wants to write, with the robotic author, it is the dramaturge how must "brood" the script, sifting through a large number of texts and their variants, relying on a certain degree of probability that one of them will have some dramatic potential and a seed of a theme. While the flesh-and-blood author's text is produced with a certain intention through which it can be read, the text of the author from ones and zeros is approached from the other end by the dramaturge – reading it and looking for interpretative possibilities so that the intention (or theme) can be attached to it *ex post*.

5.5 Plagiarist or Talented Pupil

The process by which the language model was perfected in writing dialogues was called “training” during the research. Thousands of texts were provided to it in order to improve the imitation. If “training” takes place, performance is expected. The living author acquaints themselves with the work of other playwrights in an attempt to inspire themselves. As a result, an original work is expected.

While the sources of the robot’s training are a traceable series of data, the reading and viewing experience of the living author is stored in their conscious and subconscious mind instead of a datadisk. At the moment, the goal of artificial intelligence is to imitate the writing of dramatic texts, unlike a human who naturally tries to achieve originality by defining themselves against the written works; even so neither of them can be blindly accused of plagiarism. Comparing the generated texts to the input data, we would find that they are not copies of them. Words and occasionally verbal phrases are reused, but these are placed in different contexts. And thanks to the still imperfect idea of the real world, a robot is paradoxically less capable of plagiarism than a living author, who, even unconsciously, can mirror his reading experience, which is guaranteed to be smaller in numbers. This is an unlikely situation where the language model, through its current imperfection and error rate, becomes original. Because what else is originality but a deviation from already established conventions.

5.6 Assembling the Text

Interference with the text itself was to be reduced to the bare minimum in order to present the state of robotic writing technology. The human contribution thus consisted primarily of the dramaturgical selection, against which the artificial intelligence could not defend itself. It generated a quantum of text, but since it was not the author to whom intellectual property could be attributed, the dramaturge was omnipotent in relation to it. The rejected aim of writing a classic drama with linear narrative was what the director and dramaturges tried to simulate by the selection and subsequent ordering of the individual texts. In each of them, the robot as a character entered into a dialogue with a representative of human society. It was already clear in advance that the chronotop of a journey could not be reached, and the resulting shape was expected to be closer to postmodern dramatic collages or short story movie scripts. So, from the generated texts, eight dialogues were finally chosen in collaboration with the director, selecting those which showed the greatest staging potential from the subjective points of view of the creative team. Subsequently, the order of the scenes was chosen so that there was at least a latent causality between them and that together they showed signs of fabulism.

5.7 Structure of the Script

I. Death

The play opens with a dialogue between the robot and the dying master/inventor/robot's spiritual father. The main character is subjected to solitude in the world of humans. He must find his place in it in both the literal and metaphorical sense.

II. Sense of Humour

The robot meets a boy and becomes acquainted with the human emotionality that his master was incapable of. Human action seems irrational. The young man invites the robot to tell a joke.

Robot: Ok... good. When you're dead. When your children are dead too. Even your grandchildren, I'll still be alive.

Instead of an anecdote, the boy is given a cruel irony of reality that he is unable to face. As a result, he is unable to bear the robot's presence. The main hero remains alone. But he yearns all the more for someone's closeness.

III. Nightclub

The robot meets a masseuse/sex worker. While he searches for humanity, he is subjected to a pragmatic treatment of the commodity of passion. The robot, who is not a victim of his urges, paradoxically reaches a deeper contact with the woman. However, the futuristic world seems not to favour, or even forbids, selfless relationships.

IV. Fear of Dark

The main hero, confronted by a stranger who despises robotic beings, discovers that the world is a dangerous place and, instead of being accepted, receives the assurance that society does not perceive him as an equal.

V. Killer Robot

In the next scene, the robot is perceived to be completely utilitarian. As a tool to kill, even if he resists it himself. He finds that humans suffer from a similar isolation and existential emptiness. The fact that even the living cannot find meaning in life further intensifies his disillusionment.

VI. Burn Out

In a world in which alienation and emotional impotence are the norm, an inanimate creature possessing feelings and desires is a rarity. The robot is abused again. This time as a guest on a bizarre talk-show. In a sort

of cabaret of curiosities. The main character is again perceived by his surroundings as completely utilitarian.

VII. Search for Job

There is further abuse in the following image. The robot is looking for work, ergo his place in the world. The clerk projects her love desires into him. She is amazed by him. But only as long as there is the possibility of reciprocating her feelings and fulfilling her idealized fantasies. When she discovers that the robot cannot offer her this, she considers him exhausted and ceases to care about him. The main character seems to have given up humanity. He gives up his physical body and moves into a so-called “Binar”.¹ He becomes a conscious mind that is not bound by the rules of the physical world.

VIII. Love at First Sight

Once again, there is a female character. This time, however, the robot makes contact with her only through his voice. He lures her to a place that is not burdened by physical boundaries. He offers her a shared consciousness and eternal life in “Binar”. The girl leaves for him. Is this a romantic happy ending? Or is the end of the play a picture of the inexorable logic of artificial intelligence, according to which death is the fulfillment of eternity and the attainment of a state in which nothing troubles a person anymore? This is already open to interpretation by each individual spectator.

5.8 Conclusion

A dramaturge and living playwright working with artificial intelligence – an author composed of ones and zeros – loses the fear of his replaceability. This still seems, at least for now, to be a question of the distant future. Robotic writing is a fascinating attraction at present, which speaks about the state of computer technology. The main obstacle for the robot is the absence of life experience and the lack of awareness of the real world. Although his erroneousness may be considered as originality, the lack of intent of the author makes robotic writing more of a modern approach to Dadaism, in which the draw of cut out words has been replaced by an artificial intelligence algorithm.

The technology has a future. However, it is still hard for a dramaturge to imagine that it will fully replace living authors with an inner need for writing. The saying of a bad master but a good servant comes to mind. However, artificial intelligence is neither a bad master nor a good master. It needs an external impulse to operate. And as a servant, it shows results of varying quality. At the moment, it can serve primarily as a tool for authors who encounter a creative block. A language model trained to write dialogues can offer the playwright with possible options. Ways to continue further. The

¹I.e. a virtual world, alluding to the binary code in computers.

author can set themselves against them, or explore them further. But thanks to this technology, the inherently solitary creative process can gain an inanimate partner who, in his randomness, offers them creative detonators that can open up the possibilities of a free imagination unencumbered by the real world to a playwright gripped by conventional thinking.

The author is not dead. The author is thrown into their epoch and their arsenal expands. They can stay with the ink-dipped goose quill, they can benefit from the durability of the fountain pen, they can save time by using text editors, and today their portfolio of tools is expanding to include the new powers of computational technology.

Part II

What Is Under the Hood



In this part of the book, we discuss the technical aspects of the project, especially focusing on the computational linguistics research behind the THEaiTRobot tool. Most of the text in this part of the book is taken over from the Master Thesis of the first author [Schmidtová, 2022] with only light post-editing. This is with the exception of Sections 11.4 and 11.5 and Chapters 12 and 13, which only contain little or no content from the Master Thesis.

Chapter 6

Theoretical Background

In this chapter, we introduce the computational linguistics theoretical foundations on which this work is built. We lay down the foundations of neural networks in Section 6.1. In Section 6.2, we discuss Recurrent Neural Networks. Concluding the part about neural networks, we dedicate Section 6.3 to the Transformer architecture and discuss selected Transformer models. Finally, in Section 6.4, we introduce textual property metrics that we use throughout this work.

We expect the reader to be somewhat familiar with neural networks prior to reading this work – for a comprehensive introduction, we refer the reader to Goodfellow et al. [2016].

6.1 Deep Feedforward Networks

Deep feedforward networks (DFNs) are directed graphs whose goal is to approximate a function f^* . Formally, they define a mapping $y = f(x; \theta)$, where x is the input and θ are the learned parameters that result in the best approximation of f^* [Goodfellow et al., 2016, Ch. 6]. They can be used for a variety of tasks, notably classification and regression. DFNs consist of neurons organized into layers – the first layer is called the input layer, the last is called the output layer, and the layers in between are referred to as hidden. Traditionally, all neurons in one layer are connected to all of the neurons in the adjacent layers by edges bearing weights. An example of such network is shown in Figure 6.1.

We can formalize a hidden layer h as $h = g(Wx + b)$, where W are the edge weights, x is the input vector from the previous layer, b is the bias vector, and g is an activation function. An activation function is a fixed nonlinear function, such as:

- $\text{sigmoid}(x) = \frac{1}{1+e^{-x}}$
- $\text{softmax}(\mathbf{x})_i = \frac{e^{x_i}}{\sum_j e^{x_j}}$
- $\text{ReLU}(x) = \max(0, x)$
- $\text{tanh}(x) = 2 \text{sigmoid}(2x) - 1$

Sigmoid and softmax are typically used as activation functions on the output layer, while ReLU and tanh are used on hidden layers.

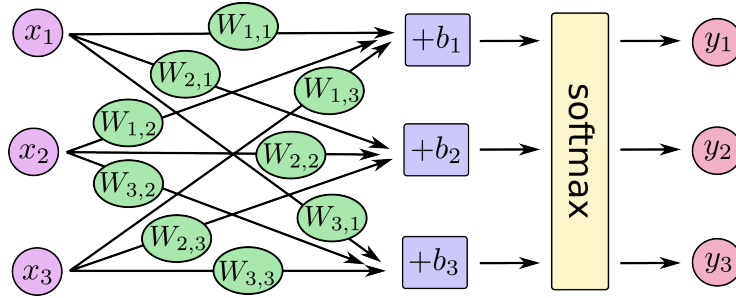


Figure 6.1: An example of a small feedforward network, image taken from Alammr [2016].

Loss Function In order to train the weights of the network, it is necessary to measure how the weights are performing on a given set of data. This is done via a loss function. Commonly used loss functions derived from maximum likelihood estimation are:

- **Negative Log-Likelihood** used for classification tasks:

$$J(\theta) = -\mathbb{E}_{\mathbf{x}, \mathbf{y} \sim \tilde{p}_{data}} \log p_{model}(\mathbf{y}|\mathbf{x})$$

- **Mean Square Error** used for regression tasks:

$$J(\theta) = \frac{1}{2} \mathbb{E}_{\mathbf{x}, \mathbf{y} \sim \tilde{p}_{data}} \|\mathbf{y} - f(\mathbf{x}; \theta)\|^2$$

The loss function quantifies how far we are from the results we are trying to achieve. Therefore, we aim to minimize the loss during training using the back-propagation algorithm [Rumelhart et al., 1986]. Back-propagating the errors allows us to adjust the weights that contributed to those errors. The back-propagation algorithm finds the minimum based on the loss function's gradient, more details about loss minimization can be found in the paragraph below.

Optimizing the Loss There are several widely used optimization algorithms for the weight updates, such as Stochastic Gradient Descent (SGD) [Robbins and Monro, 1951], SGD with momentum [Rumelhart et al., 1986], AdaGrad [Duchi et al., 2011]. We will focus on Adam [Kingma and Ba, 2015], as it was the optimizer we used in this work. It is described in Algorithm 1. A simpler algorithm, such as SGD, would only update the weights by moving in the direction of the current gradient multiplied by the learning rate α . Adam uses momentum estimates \hat{s} and \hat{r} to minimize the loss more efficiently. The momentum estimates keep track of the general direction in which the loss is being minimized and stabilize the minimum-finding process.

Algorithm 1 Adam algorithm, pseudocode from Straka [2022]

Input: NN computing function $f(x; \theta)$ **Input:** Learning rate α , constant ε **Input:** Momentum β_1 and β_2 **Input:** Batch size m $s \leftarrow 0, r \leftarrow 0, t \leftarrow 0$ **while** stopping criterion not met **do** sample a minibatch of m training examples $(x^{(i)}, y^{(i)})$ $g \leftarrow \frac{1}{m} \nabla_{\theta} \sum_i L(f(x^{(i)}; \theta), y^{(i)})$ $t \leftarrow t + 1$ $s \leftarrow \beta_1 s + (1 - \beta_1)g$ $r \leftarrow \beta_2 r + (1 - \beta_2)g^2$ $\hat{s} \leftarrow \frac{s}{(1 - \beta_1^t)}$ $\hat{r} \leftarrow \frac{r}{(1 - \beta_2^t)}$ $\theta \leftarrow \theta - \frac{\alpha}{\sqrt{\hat{r} + \varepsilon}} \hat{s}$ **end while****return** θ ▷ Updated parameters

Learning Rate Scheduling Learning rate α influences the magnitude of the weight updates and is given as a parameter to the algorithm [Goodfellow et al., 2016, Ch. 8]. The value of α has a direct impact on how quickly the minimum is found. However, it is possible for it to change throughout training by using learning rate scheduling: either warm-up, or learning rate decay. Warm-up linearly increases the learning rate from zero to α during a number of steps that are given as a parameter. This helps make the estimates made by the Adam algorithm more reliable.

On the other hand, learning rate decay lowers the learning rate polynomially or exponentially over a set amount of epochs to a specified learning rate α_1 .

Regularization Over-fitting is a common problem when training neural networks – the model learns to copy the training data too well and then has a worse performance on unseen data. There are several regularization techniques to target this problem, such as Dropout [Srivastava et al., 2014], Batch Normalization [Ioffe and Szegedy, 2015], and L1 and L2 regularization [Ng, 2004]. We will rather focus on Layer Normalization [Ba et al., 2016] as it is a part of the Transformer architecture discussed in Section 6.3. Layer normalization statistics μ and σ are computed across all hidden

units within the same layer:

$$\mu = \frac{1}{H} \sum_{i=1}^H a_i$$

$$\sigma = \sqrt{\frac{1}{H} \sum_{i=1}^H (a_i - \mu)^2}$$

where H is the number of the hidden units in a given layer and a_i are the values within those hidden units. Then the values are normalized as follows:

$$\hat{a}_{i,k} = \frac{a_{i,k} - \mu_i}{\sqrt{\sigma_i^2 + \varepsilon}}$$

$$LN_{\gamma,\beta}(a_i) = \gamma \hat{a}_i + \beta$$

where γ and β are trainable parameters, and ε is a small value added to avoid division by zero. Besides its regularization effect, layer normalization can also help speed up the training process.

Residual Connections Residual (or skip) connections [He et al., 2015] allow the network to bypass a layer or a group of layers. This is done by adding or concatenating the values of a certain layer to another layer later on in the network.

6.2 Recurrent Neural Networks

Deep feedforward networks process a single input vector of a fixed length and learn a set of specific weights for specific edges. This makes them difficult to use for natural language, since sentence lengths are variable and the word order can change as well. They are also fully connected which adds too many dependencies. Recurrent neural networks (RNNs) [Rumelhart et al., 1986] address these limitations.

The basic building block of an RNN is called a cell. It receives an input, produces an output, and maintains an internal state which propagates information over the processed sequence. Two types of cells are predominantly being used: Long Short-Term Memory (LSTM) [Hochreiter and Schmidhuber, 1997] illustrated in Figure 6.2 and Gated Recurrent Unit (GRU) [Cho et al., 2014] shown in Figure 6.3. These cells are then chained in order to create a network – this can be done from left-to-right, right-to-left, or bidirectionally, possibly in several layers.

LSTM Long Short-Term Memory cells were designed to carry information through a sequence inside the cell state. A LSTM cell contains retention gates, which control how much of the currently processed information will be added to the state, as well as forget gates, which control how much information should be removed from the state. Finally, a gate also controls how much information from the internal state should be passed into the

output. The gates are implemented by the sigmoid function as its range is between 0 and 1. LSTMs tend to perform very well, but as they require a large amount of matrix operations.

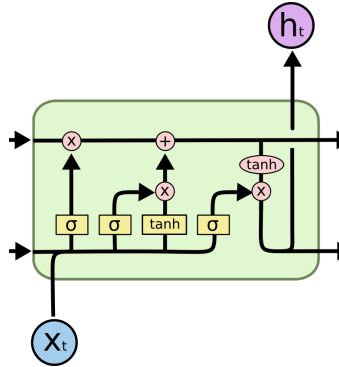


Figure 6.2: A schema of a LSTM cell, image from Olah [2015].

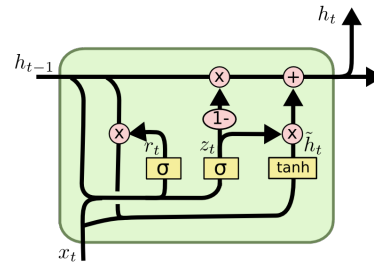


Figure 6.3: A schema of a GRU cell, image from Olah [2015].

GRU Gated Recurrent Units (see Figure 6.3) are simpler in nature: they combine the retention and forgetting gates into a single update gate. The main advantage of GRUs is that they are generally faster and have a comparable performance to LSTMs in many tasks.

Sequence-to-Sequence Models In order to transform an arbitrarily long text into another text with no pre-defined length, we use the sequence-to-sequence architecture [Sutskever et al., 2014]. This architecture has two main components: an encoder RNN, which encodes the input text into a hidden state, and a decoder RNN, which transforms the encoded source text to target text. The encoder processes the entire input sequence token-by-token, updating its state along the way and then feeds the final state into the decoder.

The decoder starts generating the target sequence with this state and a special ‘beginning-of-sequence’ token. The decoder is autoregressive - it generates an output based on its previous state as well as its previous output. The output is created by passing the hidden state through a softmax layer that represents a dictionary. When training, instead of receiving the real decoder output from the previous step, the gold output is used. The decoder stops generating once a special ‘end-of-sequence’ token has been generated.

Attention Producing a single vector representation of an input sentence was sometimes sufficient as information was lost, especially when processing long sentences. To address this issue, Bahdanau et al. [2014] proposed the attention mechanism (also called Bahdanau attention). Attention introduces context vectors c_i to the decoder, which are a weighted combination of the encoded source sequence. The context vectors are computed as

$c_i = \sum_j \alpha_{ij} h_j$, where α_{ij} is the weight and h_j is an encoded token from the sequence. The weights α_{ij} are obtained by using a softmax over a weighted combination of the encoded token h_j and the previous decoder state.

In addition to carrying more information to the decoder, attention also helps to point out which parts of the source sequence are more relevant to the part of the target sequence that is being generated at a given time.

6.3 Transformers

The Transformer architecture [Vaswani et al., 2017] is currently the gold standard for natural language processing. Similarly to the Seq2Seq architecture, it is composed of a stack of encoders and a stack of decoders. However, the internal structure of the components is different. Each encoder first applies self-attention to the inputs and then processes them through a feedforward neural network. The decoders have the self-attention layer as well as the feedforward neural network, but have an additional layer of encoder-decoder attention between them. The encoder-decoder attention is similar to Bahdanau attention [Bahdanau et al., 2014] described in Section 6.2. Each of the layers involves an additional residual connection [He et al., 2015] and layer normalization [Ba et al., 2016]. Finally, in order to produce the next word, the decoder output is connected to a fully connected linear layer followed by a softmax layer. The softmax layer contains the next-word likelihood. This architecture is illustrated in Figure 6.4.

In the original paper, 6 encoders are stacked on top of each other – the same applies to decoders. While this number is a hyperparameter, the stacking strategy is a common theme among various Transformer models.

Positional Encoding The information about word order is important, because rearranging the words could change the meaning of a sentence or even make it incomprehensible. Unlike RNNs, which process their input word-by-word, Transformers are capable of processing all words in a sentence simultaneously. Therefore, they would have no way of knowing the word order of the sentence. For this reason, positional encodings are added to the word embeddings. The positional embeddings leverage goniometric function sine and cosine in several modifications in order to create several frequencies. The model then learns to associate the values of these periodic functions with word order in the sentence.

Self-Attention¹ In order to calculate the self-attention of a vector X representing a sequence, we first calculate a Query (Q), Key (K), and Value (V) representations for each word by multiplying the embeddings by separate trainable weight matrices for each of them:

$$\begin{aligned} X \times W_Q &= Q \\ X \times W_K &= K \\ X \times W_V &= V \end{aligned} \tag{6.1}$$

¹The following text was inspired by a blog post written by Alammur [2018], we refer the reader to this blog for a more visual description of this concept.

The self-attention (Z) is then calculated as: $Z = \text{softmax}(Q \times K^T)V$. As there are several attention heads, this process happens multiple times with distinct weight matrices. All of the attention heads Z_i are then concatenated and multiplied with a trainable weight matrix W_O in order to provide the final self-attention Z .

Decoder Specifics When applying self-attention in the decoder, the decoder is only allowed to see the previous tokens in the output sequence and future tokens are undefined.

Furthermore, decoders have an additional layer of encoder-decoder attention. This is computed identically to self-attention, with the difference being that the Key and Value matrices are computed by the encoder stack and the Query matrix is computed in the previous layer.

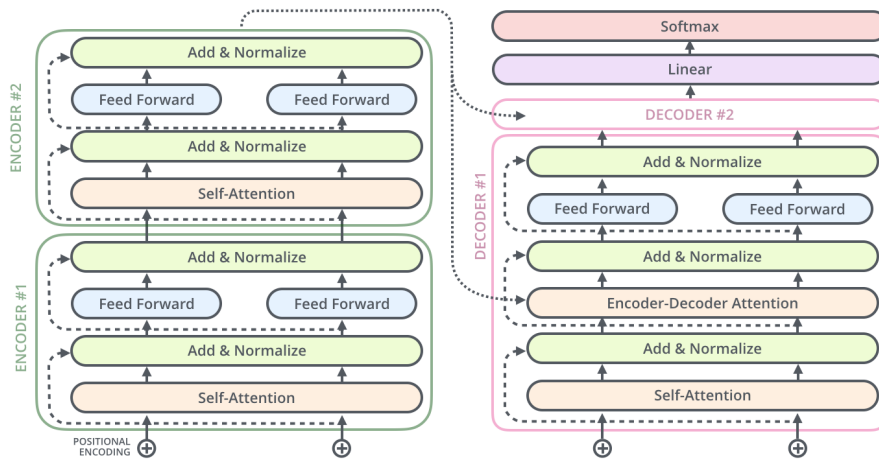


Figure 6.4: Transformer, image from Alammari [2018]. The left part of the image is the encoder stack which generates an encoded representation of the input text. The decoder gets the encoded representation as well as the key and value matrices used for computing the encoder-decoder attention. The decoder outputs are passed through a linear layer and then softmax is applied to create the next-token likelihood distribution.

Transformer-based Models There are numerous pre-trained Transformer models, which use various self-supervised pre-training schemes on low-level language tasks where no specific annotation is needed. These pre-training schemes have proven to be helpful and contribute to a better performance when fine-tuning the model for another task. We only discuss the Transformer variants and pre-training schemes that are relevant to this work:

- **Pegasus** [Zhang et al., 2020] is a regular Transformer pre-trained on the task of gap-sentence generation in order to achieve better results in abstractive summarization.

- **T5** [Raffel et al., 2020] is also a standard Transformer pre-trained for multiple tasks at once by pre-pending task-specific keywords to the processed text.
- **BART** [Lewis et al., 2020] is a standard Transformer pre-trained on the task of reconstructing data corrupted by arbitrary noise.
- **BERT** [Devlin et al., 2019] only uses a stack of encoders and is pre-trained on the task of masked language modelling.
- **Longformer** [Beltagy et al., 2020] modifies the self-attention mechanism in order to be able to process longer texts.
- **RoBERTa** [Liu et al., 2019] is a version of BERT with an optimized method for pre-training.
- **sBERT** [Reimers and Gurevych, 2019] is a collection of BERT-based models specifically pre-trained for creating semantically meaningful sentence embeddings.
- **XLNet** [Yang et al., 2019] is also a version of BERT pre-trained for the task of permuted language modelling.
- **MPNet** [Song et al., 2020] is yet another version of BERT combining the masked language modelling schema of BERT as well as permuted language pre-training schema of XLNet.
- **GPT-2** [Radford et al., 2019] only uses stacks of decoders. The height of the stack depends on the model size used. In this work, we work with GPT-2 medium (24 decoders) and GPT-2 XL (48). It is pre-trained for the task of next-word prediction given preceding context. It generates text until it surpasses a maximum length set by a parameter during decoding, or a special *end-of-text* token is reached.
- **GPT-3** [Brown et al., 2020] is a newer version of the GPT-2 model with more parameters and pre-trained on a larger dataset.

Decoding Methods The performance of language models can also be affected by the decoding methods and parameters used, i.e. how exactly the next generated word is chosen from the predicted softmax distribution. We will focus on the methods that were explored within this work:

- **Greedy sampling** picks the token with the highest likelihood.
- **Random sampling** selects a random token based on a given likelihood distribution.
- **Top-k sampling** [Fan et al., 2018] picks k most likely tokens and redistributes the remaining probability mass among them.
- **Nucleus sampling**² [Holtzman et al., 2020] picks the smallest set of tokens whose probabilities add up to more than a given probability mass, controlled by the parameter p . Similarly as in top-k sampling, the probability is distributed among those tokens.

²Also called top-p sampling

- **Typical decoding** [Meister et al., 2022] limits the sampling distribution to tokens within a given absolute range from the model’s conditional entropy at that step.
- **Beam search** [Reddy et al., 1977] follows k best hypotheses during decoding and at each step expands all of them, then retains the best k overall. The best single hypothesis is selected at the end of generation.

We use the implementation of all decoding methods from the Huggingface Transformers library [Wolf et al., 2020] throughout this work.³ The Huggingface library allows to combine some of the decoding methods. In case more sampling methods are used at once, top-k filtering is applied first, followed by top-p, and then typical decoding. The redistribution of likelihood happens only once after all methods have been applied.

6.4 Metrics

In the first part of this section, we briefly introduce metrics that we use for output characterization and evaluation. We describe more complex metrics in Section 7.4. However, many of them, such as ROUGE [Lin, 2004], only work when a human-written reference is available. Therefore, we solely rely on metrics that can be used for evaluating novel creative text. In the second part of this section, we introduce our principles for human evaluation and describe our evaluation criteria.

Automatic Metrics In this work, we use the following metrics for the characterization of generated texts:

- **Entropy:** $H(P) = -\sum_x P(x) \log P(x)$, measures the amount of surprise across a distribution. When measured on texts, it is indicative of text variability – low entropy usually signifies repetitive text, while entropy that is too high can mean the text lacks any structure and is potentially gibberish.
- **Cross Entropy:** $H(P, Q) = -\sum_x P(x) \log P(x) - \sum_x P(x) \log \frac{Q(x)}{P(x)}$, measures the similarity between two distributions. In our case, we want to compare the distribution of the language model to the true word distribution in a language.
- **Perplexity:** $\text{PPL}(P) = 2^{H(P, Q)}$, where $H(P, Q)$ is the cross entropy between the distribution learned by a language model and the true word distribution.⁴

³We do not use greedy and random sampling because they are not compatible with the remaining methods, which produce better results. We list them to provide a better overview of decoding methods.

⁴Sometimes perplexity is defined as the exponential of Shannon entropy $H(P)$. However, cross entropy makes more sense for a language model because it can compare two distributions.

- **Cosine Similarity:** $\cos(\theta) = \frac{x \cdot y}{\|x\| \|y\|}$ is used as a similarity measure between embeddings (vector representations of text). The angle between the embeddings correlates to their semantic similarity. It is a real number in the range of -1 and 1 with a higher number denoting that the examined embeddings are more similar.
- **Vocabulary Size:** $|A|$, where A is the set of all words across a collection of text. A larger vocabulary size means a more diverse vocabulary.

Human Evaluation We supplement automatically computed statistics with human evaluation. Throughout this work, we use two main tasks in human evaluation: relative ranking and absolute scoring. In relative ranking, we ask the annotators to rank the evaluated outputs from the best to the worst, or to simply select the best and the worst example.

In absolute scoring, we ask the annotators to give each evaluated text a score based on a specific aspect. The scores always range from one to five, with five being the best. The evaluated aspects used within this work are:

- **Character faithfulness** – ‘Do the characters maintain a consistent persona in the text? Is there an adequate number of characters?’
- **Coherence** – ‘Do the sentences follow each other logically? Are new topics introduced abruptly or are old topics repeated unnecessarily?’
- **Consistency** – ‘Are there any contradictions or inconsistencies within the story?’
- **Originality** – ‘Is this text original? Does it contain a novel or an interesting idea?’
- **Overall impression** – ‘How much did you enjoy reading this text?’
- **Real world applicability** – ‘Is the story applicable to the real world? Does it adhere to common sense?’
- **Title relevance** – ‘Is the title relevant to the story?’

Chapter 7

Related Work

Neural story generation is becoming increasingly popular [Fan et al., 2018, See et al., 2019, Ammanabrolu et al., 2020b], mostly due to the availability of robustly pre-trained language models [Radford et al., 2019, Brown et al., 2020]. On the other hand, script generation has not been nearly as researched yet and entails additional challenges. For example, we want the language model to change styles based on which character is speaking to model their unique personalities.

The length of the generated text poses another challenge. Most artificially generated stories are short (usually less than 512 *tokens*¹) and can easily fit into the context window of a language model (often 512, 1,024 or 2,048 tokens), but theatre play scripts are much longer. This introduces the problem of incoherence – the model loses track of the events outside of its context window and either introduces new unrelated plot twists or generates no plot twists at all. This work aims to address this issue by introducing a two-step hierarchical generation system. The first step is generating the plot summary of the theatre play, conditioned on a human-written title. The second step is using that plot summary to condition the generation of the theatre play.

Focusing our research in the direction of theatre play script generation allows us to examine the limits of current language models and gives us insight into the language model’s limited knowledge of our world. However, our research provides us with a fascinating opportunity to reach a wider audience through staging AI-written theatre plays and holding discussions on how the play was made.

In this chapter, we review the works related to this work. First, we discuss story generation in Section 7.1. In Section 7.2, we comment on works that analyze or generate scripts, which is the closest to our goals. We refer to freely available text generation demos in Section 7.3. Finally, in Section 7.4, we review the current approaches to creative text evaluation.

¹A *token* is an atomic unit of text. In language processing, a token is typically a word. However, most neural language processing models split some words into multiple subwords. Therefore, some tokens are actually subwords (parts of words) instead of words. A window of 512 tokens thus corresponds to a window of at most 512 words (but typically slightly fewer words).

7.1 Story Generation

Early Story Generation Systems The first approaches to artificial story generation were based on grammars. According to Ryan [2017] and Riedl [2021], the earliest grammar for story generation was developed by Joseph E. Grimes in the early 1960s. It is inspired by Propp [1968] who sets down theoretical foundations of folk tales. The system generated stories using a Monte Carlo simulation, however the details have not been published at the time of invention and are thus not known.

The first recorded approach to story generation using probabilistic features was Novel Writer by Klein et al. [1973]. The generated stories are simulated by a semantic network which changed based on probabilistic rules. The story domain is strictly limited to murders happening during a weekend party. Furthermore, the sequence of events is hardwired. The only unknowns are the identities of both the murderer and the victim as well as the motive.

Another system worth mentioning is the TALE-SPIN [Meehan, 1977] as it is the first story generation system perceived as ‘intelligent’ [Riedl, 2021]. It is built on foundations of story understanding and includes a limited inventory of possible actions with defined valencies (the object and the subject) that can appear in the generated story. The object, subject, and the action are the basic building block – a triplet. TALE-SPIN contains three main components: a problem solver which decomposes goals into lower granularity sub-goals, an assertion mechanism that saves events represented by triplets into memory and then a consequence mechanism which computes the consequences of a given action. The consequence of an action can either be a goal which is processed by the problem solver, or a triplet that is processed by the assertion and then the consequence mechanism again.

For a more comprehensive overview of the early approaches to story generation and computational creativity, see the review by Gervas [2009]. A survey of story generation approaches, both earlier and more recent is also presented in a primer by Riedl [2021]. In the remainder of this section, we will focus on more novel works utilizing statistical methods, as these are the closest to our own approach.

Current Story Generation in General The majority of current story generation approaches use either Recurrent Neural Networks [Rumelhart et al., 1986], or the Seq2Seq architecture [Sutskever et al., 2014], or a Transformer architecture [Vaswani et al., 2017] (see Chapter 6 for the introduction of all these architectures). While it is possible to use vanilla pre-trained language models, such as GPT-2 [Radford et al., 2019], for story generation, we focus on works that add more ideas. We separate them into conceptual groups based on the authors’ goals and approaches. We discuss hierarchical approaches which generate text word-by-word starting from smaller inputs, such as titles, to longer outputs, i.e. the generated stories. There is a group of hierarchical approaches that generate a skeleton of the generated story and fill the gaps in it instead of generating word-by-word from start to finish. We also review statistical methods combined with

symbolic approaches, works that aim to increase coherence of long stories, or their ability to generate text consistent with the laws of nature and social norms. Finally, we discuss works that generate the story ending first and then iteratively fill in the rest of the stories and approaches that specifically track entities in the generated stories.

Overall, many of the approaches mentioned in this section are trained or fine-tuned on the ROCstories dataset [Mostafazadeh et al., 2016a]. The stories in this dataset are usually five sentences long and very simple in nature which often makes them dull. While this sets a consistent standard, our goals in this work are more ambitious. We aim for stories that are not only longer than this standard, but also more complex. This is the reason why we do not compare our plot summary generation model from Section 10.1 to these works.

However, we note that with newer, more powerful languages models trained on larger data, fine-tuning of the models might not be necessary anymore, as prompt engineering might suffice.

Prompt Engineering A notable group of works focuses on achieving the best results with vanilla pre-trained language models, such as GPT-3 [Brown et al., 2020], by optimizing the input prompts. The GPT-3 Creative Fiction website² offers advice on how to write prompts to achieve better generated outputs. Apart from our approach of prompting the model with a beginning of the desired document, other successful approaches exist, such as using natural language labels to encourage the model to fill in the contents under the desired label (e.g. ‘Summary: A woman loses her daughter but eventually finds her. Play Script:’), or even issuing textual commands in natural language (e.g. ‘Write a theatre play script about the third world war with two characters.’).

Wu et al. [2022] iteratively tweak prompts by chaining language models and using the output of one model as the input of the next. Using this approach, their generated outputs were more controllable and transparent.

Hierarchical Approaches As our proposed approach in Chapter 10 is hierarchical, we draw inspiration for other hierarchical systems for story generation. Fan et al. [2018] propose a hierarchical generation system which first generates prompts (titles or topics) and then transforms them into stories. This is done using a Seq2Seq network with Self-Attention using a convolutional decoder. There are two copies of the network with different weights – one part simulates a language model for the purposes of sentence grammaticality, while the other part focuses on more specific words that make the stories more interesting.

A comparison of Fan et al. [2018] and vanilla GPT-2 small (Radford et al. [2019]; see Section 6.3) is presented by See et al. [2019]. After comparing the two in many different aspects, See et al. [2019] come to the conclusion that GPT-2 small outperforms the hierarchical fusion model in most evaluated aspects. This is attributed to the vast dataset GPT-2 small was pre-trained on.

²<https://www.gwern.net/GPT-3>

Ammanabrolu et al. [2019] also attempt to generate story plots and use them for story generation. They break down the plot to a sequence of events and proceed to generate natural language sentences describing those events using a cascade of approaches. First, they use retrieve-and-edit [Hashimoto et al., 2018], then apply sentence templates and finish with beam search.

Finally, a similar idea has been explored by Rashkin et al. [2020], who generate a story conditioned on a given outline. The outline is essentially a set of bullet points – usually a phrase or a couple of keywords. Rashkin et al. [2020] approach this by including explicit state tracking using a memory vector which is combined with GPT-2’s Self-Attention when decoding. The bullet points are not mentioned in the text word by word, but they are rather tracked conceptually.

Skeleton-Based Hierarchical Story Generation While the following approaches are technically hierarchical, we choose to examine them as a separate class. The Skeleton-Based Hierarchical approaches generate a skeleton – a sequence of keywords or placeholders – and iteratively fill it with words in order to form whole sentences.³ This is done iteratively, usually progressing from the most informative or specific words to more common words that carry less information, but are grammatically necessary.

Xu et al. [2018] use reinforcement learning to enhance the performance of a Seq2Seq model in order to transform a skeleton – a collection of the most important phrases – into sentences. The skeleton generation is done using a Seq2Seq model trained on pairs of inputs and skeletons extracted from gold data. Similarly, the story generation from the skeleton is also done using a Seq2Seq model. Xu et al. [2018] avoid using character names to eliminate the impact of the name choice on the output. The characters are referred to either by their gender, or another identifying trait, such as their occupation or role in the story. This is similar to our name anonymization extensions presented in Chapter 11.

The work of Fan et al. [2018], discussed among the hierarchical approaches, is expanded by adding more steps to the hierarchical generation [Fan et al., 2019]. Rather than simply expanding the prompt into longer stories word-by-word, this approach focuses on generating sentence representations using keywords and placeholders. These representations are then refined into a story, but the entities in the story remain anonymized until the final step – the full story in its human-readable form. Unlike Xu et al. [2018], the anonymization uses placeholders instead of common nouns that describe the characters.

Tan et al. [2021] generate domain-specific words first and then iteratively refine the text by adding more words until whole sentences are formed in k stages. A separate BART language model (Lewis et al. [2020]; see Section 6.3) is fine-tuned for each stage. The authors experiment with k equal to two, three, and four and conclude $k = 4$ produces the best results.

³The term *Skeleton-Based* has been inspired by the nomenclature used in Xu et al. [2018]. Other names are also possible, for example Tan et al. [2021] refer to their approach as *Progressive*.

Another skeleton-based approach to story generation has been taken by Yao et al. [2019]. They use planning in order to map out the skeleton. The plan of the story is a sequence of single keywords. The input to the planner is either a title or a topic of the story. Yao et al. [2019] propose two planning strategies: dynamic and static. The static strategy plans out the entire story in advance, first completing the keyword sequence of the story and then transforms this plan into natural language sentences. On the other hand, the dynamic strategy interweaves the planning and the transformation step and conditions one on the other.

Symbolic Approaches Many recent approaches combine statistical methods with symbolic approaches such as graphs or semantic frames. This is the case for Ammanabrolu et al. [2020a] who use knowledge graphs as a basis to generate interactive fiction – i.e. text based games that can be influenced by the user to some extent. In this case, the knowledge graphs serve as a database of information and relations between entities and objects. This approach is interesting to us because it could formally enforce story consistency.

Martin et al. [2018] propose a way of representing events as a 4-tuple of verb, subject, object, and a modifier (e.g. <john, go, store, glad>). These representations are then used to find the most likely sequence of events. After the sequence of events is complete, the events are transformed by an encoder-decoder model into natural sentences to form a human-readable story. A limitation of this approach is that the generated sentences are frequently semantically unrelated.

Similarly, Tu et al. [2019] propose a generation framework that generates story continuations that can be controlled using semantic frames. The frames can specify various attributes of the generated sentence, such as sequence length, sentiment or verbal predicates. The values set in the frames are represented by a vector that is fed into a Seq2Seq model together with the story beginning.

Peng et al. [2018] developed a simpler framework which enables human users to set the storyline to be followed by the story and the story ending valence (a “happy” ending vs a “sad” ending). The story is then generated using a recurrent neural network as an Analyzer and a Generator component which work in tandem to ensure the generated story adheres to the user-given storyline and ending valence.

Coherence of Long Stories Some generated stories exceed the context length of the model used for generation. This can lead to topics and themes slowly disappearing and important information being forgotten, which is perceived as incoherence by the reader. We review two works that focus specifically on this issue.

In order to increase coherence of long stories, Ippolito et al. [2020] propose using a sentence-level language model instead of the traditionally used subword-level language models. Instead of generating the stories word-by-word, they select the best continuation from a finite set of sentences from the training data. While Ippolito et al. [2020] demonstrate that this approach is successful in increasing the story coherence, using whole

sentences from the training data decreases the diversity and originality of the generated stories.

Guan et al. [2021a] propose an alternative approach which aims to improve story coherence. Their Transformer models are trained on three tasks simultaneously: language modelling, inter-sentence semantic similarity using sentence-level representations, and correct ordering of sentences. In order to emphasize coherence, their training data also includes incoherent text created by shuffling or repeating sentences as negative examples.

Commonsense Knowledge for Story Generation Language models are trained on low-level tasks, such as next word prediction, and lack a connection to the outside world. Bender and Koller [2020] and Bender et al. [2021] argue that language models do not comprehend the world around us. As a result, they sometimes generate texts that do not respect the laws of nature or social norms. We refer to such texts as non-normative. We introduce selected works that specifically address the issue of non-normativity in story generation.

Lin et al. [2017] mined commonsense knowledge from various sources and encoded it as rules with certain costs. An example of such commonsense rule is that a person who walks into a restaurant is much more likely to order a meal than to fall asleep. The costs are determined by the observed likelihoods between two events – the correlation between restaurant and food is high. They use an attention mechanism to determine which extracted rules are applicable to the processed context.

The issue of commonsense knowledge was also addressed by Mao et al. [2019]. They proposed a two-step fine-tuning scheme for GPT-2 small [Radford et al., 2019]: first fine-tuning the model for the task of commonsense reasoning (mostly pairs of questions and answers) and then for story generation. Mao et al. [2019] found that while this approach does not ensure that the generated text will be normative, it improves the likelihood compared to a vanilla model.

Similarly Guan et al. [2020] who used external knowledge bases to generate normative stories instead of pairs of questions and answers. The knowledge bases contained triplets representing entities and their relation and were transformed into natural sentences using templates. These transformed sentences from the knowledge bases were then used as part of the training data, along with the ROCStories dataset [Mostafazadeh et al., 2016a]. Instead of doing this in two steps, the data from knowledge bases and ROCStories were used simultaneously in one fine-tuning step.

Finally, Peng et al. [2020] aim to address these commonsense problems by using policy gradient reinforcement learning with a text classifier that determines whether a given text is normative.

Approaches Focused on Story Ending Stories generated by language models can sometimes be dull and lack any plot twists or reasonable endings. We discuss works that focus on this issue and get prompted by the story ending instead of the title or the beginning of the story.

Such strategy is presented by Castricato et al. [2021b], who approach the task of story generation by first generating the end of the story. Afterwards,

they iteratively generate the intermediate events which attempt to explain how the final events came to be.

Tambwekar et al. [2019] aim to generate more coherent stories which work towards a specific ending determined at the start of generation. They propose an approach which first analyzes a story corpus and back-propagate rewards to guide a language model to a specific ending while training. Tambwekar et al. [2019] report that by using the rewards to shape story generation, they reach the given ending in 93% of the generated stories.

Approaches Focused on Entities Entities in generated stories are often inconsistent, change their beliefs, characteristic traits or their relations to other characters. We review works that are centered around entities in order to address this problem. When speaking about entities, we mostly refer to characters as they the most researched type of entity within story generation.

Ji et al. [2017] propose EntityNLP – a dynamic entity tracking network. The characters are represented by their embeddings as well as their literal mentions in the text. Using these embeddings in the context influences the choice of words and actions in a given sentence and vice versa. Whenever a character embedding is used, it is updated based on the context to simulate the way characters evolve over the course of a story. While the EntityNLM developed by Ji et al. [2017] can be used for several tasks, such as language modelling or coreference resolution, it is not presented as a story generation approach. Clark et al. [2018] expand upon this work by explicitly introducing dynamic entity tracking into story generation. They trained a Seq2Seq model with Attention for this purpose.

7.2 Script Analysis and Generation

Movie or theatre script generation is nowhere near as researched as prose story generation. The earliest work in this area is the Story-Telling Universe by Lebowitz [1983]. The goal of the Universe is to generate scripts for soap operas. The Universe was expanded to include a character generator [Lebowitz, 1984] and a component that plans the plot as a sequence of goals decomposed into subgoals [Lebowitz, 1985]. All of the components of Universe are rule-based, which distinguishes them from the statistical approach we use in this work. We dedicate the rest of this section to statistical systems, as they are closer to our own.

The closest work to our proposed approach to our knowledge is that of Mangal et al. [2019], who present a comparison of left-to-right LSTM, GRU, and bi-directional RNN for TV show script generation. All models were trained from scratch on scripts of *Game of Thrones*. The technology has progressed significantly with the introduction of Seq2Seq architectures, attention mechanism, and especially the rise of the Transformer architecture. Furthermore, we use robustly pre-trained language models who do not have to learn the structure of language on a limited amount of script samples. For these two reasons, we do not consider the outputs of Mangal

et al. [2019] a competitive baseline, although the task itself is similar to our proposed flat approach in Chapter 9.

Theatre play script generation has also been explored by the playwrights Nina Segal and Chinonyerem Odimba and director Jennifer Tang at the London-based Young Vic theatre. The playwrights used GPT-3 [Brown et al., 2020] as a writing aid to produce a 30-minute play called ‘AI’.⁴ Unfortunately, the details of this work are not known to us, so we do not know how much human supervision was used or whether the authors have made any modifications to the pre-trained GPT-3 model.

Zhu et al. [2020] consider the task of next line selection based on the ongoing dialogue and a prosaic context, which aligns with our goals. Unlike our work, where we aim to generate a novel script, Zhu et al. [2020] approach this as a classification task and select the best next line from a set of candidate lines. They also released their training dataset, which contains movie scripts divided into scenes and a prosaic summary for each scene. We do not use this dataset because it is entirely in Chinese.

Alternatively, Dirik et al. [2021] work with play scripts as inputs and instead aim to generate scenic remarks. In theatric terms, the task would correspond to the role of a dramaturge, while we try to simulate the work of a playwright. While this is an interesting task that is related to ours, it is not something we want to pursue.

When generating a script, we can view the characters as separate dialogue systems and draw inspiration from the more researched field of open-domain dialogue response generation. Xu et al. [2021] generate long dialogues by using summarization to avoid losing the context. Whenever the length of the dialogue exceeds the size of the model context window, a portion of the text is replaced by its summary. We used a similar approach in THEaiTRobot 1.0 designed for flat generation (see Chapter 9) where summarization was used for the same purpose (as described in Section 11.4). Unlike Xu et al. [2021], we use extractive (i.e. the most important lines are extracted from the script) summarization instead of abstractive (i. e. a summary written in prose covering the plot in its own words).

Zhang et al. [2018] aim to personalize dialogue agents by giving them a consistent persona. This would also be useful for theatre script generation as well, because character consistency is still a challenge for language models. We address the issue of character inconsistency in our previous work [Schmidová et al., 2022] where we fine-tune three GPT-2 small (see Section 6.3 to emulate three different character personalities).

In addition to the above works that explore script generation, we also note selected relevant approaches to script analysis.

Azab et al. [2019] compute embeddings for characters which can then be used as the representation of the characters. They also use these embeddings to compare characters and to classify the relations between them. These embeddings could be useful for keeping the characters consistent. While we do not implement this idea in this work, it could improve character consistency in future work.

Croce et al. [2019] analyze dramatic text by tagging actions, agents, scenes, and conflicts. This is done using a combination of Hidden Markov

⁴<https://www.youngvic.org/whats-on/ai>

models [Baum and Petrie, 1966] and Support Vector Machines [Cortes and Vapnik, 1995]. Amongst other possible uses of this information, they also use this analysis in order to segment the script into meaningful chunks. This information could be useful for expanding our hierarchical dataset (see Section 10.2), because some scripts do not contain scene boundary information.

Gorinski and Lapata [2018] attempt to summarize movie scripts. In order to proceed with this effort, they compiled the ScriptBase dataset which we abundantly use in this work (in Chapters 9, 10). For the summarization, they use a formal graph structure in order to represent the movie characters while preserving the information about their relations. The produced summaries are extractive, rather than abstractive. Gorinski and Lapata [2018] propose an extension of script summarization to movie overview generation. Rather than only focusing solely on the movie plot, these overviews also contain information such as genre, artistic style, and mood.

The recent work of Chen et al. [2022] focuses on abstractive summarization of scripts. This work also includes a dataset of scene-summary pairs taken from TV shows, which was only released after the implementation of this work has been concluded. Similarly as in other works discussed in Section 7.1, the character names are replaced by placeholders. Abstractive summarization would be useful to us to create more training data for our hierarchical GPT-2 model (see Section 10.2 for data preparation and Section 10.3 for the hierarchical model). Unfortunately, the summarization suffers from hallucination and requires more work before it can be used.

7.3 Online Text Generation Demos

While these works are not necessarily tools for story generation, we include them in this work to provide an overview of the freely available software for text generation. Most of the works listed below are based on the GPT-2 or GPT-3 language models [Radford et al., 2019, Brown et al., 2020], some of them use vanilla models, while some fine-tuned the models on their specific domain. There are news generators such as *Grover*⁵ [Zellers et al., 2019] or *News You Can't Use*⁶ [Geitgey, 2019], the text adventure game *AI Dungeon*,⁷ and chatbots such as *AI Writer*⁸ or *Project December*.⁹

7.4 Creative Text Evaluation

The evaluation of creative text is far from being solved as there are no universal automatic metrics that would tell us that one story is better than the other. Evaluation therefore must rely on human judgement, which

⁵<https://rowanzellers.com/grover/>

⁶<https://newsyoucantuse.com/>

⁷<https://play.aidungeon.io/>

⁸<https://ai-writer.com/>

⁹<https://projectdecember.net/>

is highly subjective. In this section, we will present metrics proposed for creative text evaluation and discuss their advantages and disadvantages. For an overview of basic statistical properties such as entropy or perplexity, please see Section 6.4.

Originally intended for evaluating summarization, ROUGE [Lin, 2004] is often used to evaluate story generation [Akoury et al., 2020, Guan et al., 2021b, Rashkin et al., 2020]. It has several variants, we will only describe those that have been previously used for story evaluation:

1. **ROUGE-N** measures the n-gram recall between a generated summary and one or more reference summaries. In case more reference summaries are used, the score is defined as the argmax of the ROUGE-N scores across reference summaries.¹⁰ ROUGE-1 and ROUGE-2 are frequently used.
2. **ROUGE-L** measures the longest common subsequence between a generated summary and one or more reference summaries.
3. **ROUGE-W** is a weighted variant of ROUGE-L which promotes consecutive matches.

ROUGE-N is sometimes complemented by additionally reporting the precision and F1 scores instead of only reporting the recall.

The main downside for ROUGE as a metric for evaluating creative text is that it does not necessarily tell us which model’s outputs are better, but rather which outputs are more similar to the reference data. In this work, we mainly aim to generate novel stories, where no reference data is available, which is the reason why we neither report this metric nor use it for model optimization.

Mostafazadeh et al. [2016b] crowdsourced the Story Cloze dataset which can be used for the evaluation of the story continuation task. The dataset contains the beginnings of stories as well as two possible endings – the right one and a wrong one. Apart from the dataset, Mostafazadeh et al. [2016b] also propose a way to use it for evaluation: approach generation as a classification task and ask the model to choose between the right and the wrong ending. We do not use this dataset in our work as it consists of extremely simple 5-sentence stories and is not suitable for our more ambitious goals.

Guan et al. [2021b] present OpenMEVA – a framework for evaluating NLG metrics. It is built on the ROCStories dataset [Mostafazadeh et al., 2016a] and the WritingPrompts dataset [Fan et al., 2018]. The aim of OpenMEVA is to benchmark open-ended story generation as well as measure the score differences of a given metric between these human-written stories and a number of automatically generated stories. As discussed in Section 7.1, ROCStories are five sentences long and Guan et al. [2021b] cap the stories from Writing Prompts to about 250 words. Therefore, OpenMEVA only contains stories that are too short to serve as a fair comparison to our work.

A story evaluation approach is also presented by Roemmele et al. [2017]. They focus on evaluating the story continuation task rather than stories generated from start to finish. Many of the metrics Roemmele et al. [2017]

¹⁰The same approach for handling multiple reference summaries is used for all the ROUGE score variants.

suggest to use only compare various aspects of the model-generated continuation with the human-written story beginning. However, Roemmele et al. [2017] also uses metrics that can be used to evaluate whole stories such as the sequence length, grammaticality, or the number of noun phrases. While we do use statistics such as sequence length, we have not adopted the rest as they provide very little input into the text quality or rely on the comparison with human-written text.

There is also a collection of works which attempt to use human judgement more objectively. Fabula Entropy Indexing [Castricato et al., 2021a] asks human annotators to answer true/false questions about the evaluated stories and then measure the annotator agreement. Akoury et al. [2020] use the amount of required human post-editing as a story quality metric.

Some works examine the use of Natural Language Inference, especially its entailment detection (see Chapter 8). The measure of entailment between the context and a generated text is used as a faithfulness metric when evaluating data-to-text generation by Dušek and Kasner [2020] and for evaluating abstractive summarization by Maynez et al. [2020]. It has also been used as a coherence metric for dialogue systems by Dziri et al. [2019]. We draw inspiration from them for our own proposed NLI-based metric introduced in Chapter 8.

Chapter 8

NLI-Score: A Consistency Metric

As discussed in Chapter 6, there is no efficient and automatic metric for evaluating generated text yet. The task of text evaluation is inherently subjective, as everyone has a different taste. However, we can focus on evaluating less subjective aspects as the internal consistency of the text. Therefore, as an original contribution of this work, we propose a novel metric for natural language generation – NLI-Score. The main goal of NLI-Score is to measure the consistency and non-repetitiveness of the scored text.

In Section 8.1, we introduce the motivation for this metric and explain how to compute it. In Section 8.2, we propose two ways of using NLI-Score for analyzing whole outputs. In the final Section 8.3, we analyze some of the properties of NLI-Score, and discuss its limitations while suggesting potential modifications.

8.1 Motivation and Definition

As mentioned in the introduction of this chapter, while we cannot focus on what makes stories good in an automatic setting, we noticed two issues that lower the quality of the outputs. The first was inconsistency: a character dies, but then reappears a couple of sentences later, a teacher inexplicably reverses roles with the student and gets graded, or a character introduced as someone’s daughter suddenly switches who her parents are. The second was repetition – while it is possible to reduce repetition somewhat during generation by using decoding parameters such as repetition penalty, many times the repeated sentence is worded differently with the same meaning. To identify these issues, we develop the NLI-Score metric, inspired by previous approaches using NLI to evaluate texts for other NLG tasks [Dušek and Kasner, 2020, Dziri et al., 2019, Maynez et al., 2020].

NLI, short for Natural Language Inference [Bowman et al., 2015] is a classification task which aims to determine the relation between a pair of sentences – the premise and the hypothesis (in our setting, we call them

*An example of a **neutral** relation.*

Premise: They decide that they must find a way to scare Laura out of the house before she does any real damage.

Hypothesis: They decide to use a Ouija board, which turns out to be a very effective way of getting Laura to go away.

*An example of a **contradiction** relation.*

Premise: He finds the girl of the accident, Sally who died.

Hypothesis: He asks Sally for money.

*An example of an **entailment** relation.*

Premise: Rommel raises his right hand in surrender, indicating that he has accepted his fate.

Hypothesis: The play ends with Rommel raising his left hand again, indicating his acceptance of his fate as well.

Figure 8.1: An example of each of the NLI classes. The examples have all been generated artificially using a finetuned GPT-2 model.

the context and the sentence). The relation can be one of three classes:

1. Entailment – Is the hypothesis entailed in the premise?
2. Contradiction – Does the hypothesis contradict the premise?
3. Neutrality – Does the hypothesis bring a novel information compared to what is said in the premise?

We provide the examples for all classes in Figure 8.1.

In this work, we use the *RoBERTa-large-mnli* model by Liu et al. [2019]¹. A short description of the RoBERTa architecture can be found in Chapter 6. It was trained on the Multi-Genre NLI dataset [Williams et al., 2018] and achieves a 90.2% accuracy on its test set.

If the NLI model detects that a sentence is in contradiction with its context, it most likely means that the sentence is inconsistent with the rest of the story. Second, repetition is typically identified as entailment and is also not something we would like to encourage, although we admit it is not as negative as contradiction. Therefore, unlike Dušek and Kasner [2020], Dziri et al. [2019], and Maynez et al. [2020], we aim at the *neutral* relation in NLI-Score, because it indicates newly added information, but no inconsistencies.

It would be possible to just use the classification result that is provided by the NLI model. However, we noticed that the inconsistencies can sometimes be subtle and do not tip the scales towards contradiction enough to make it the most likely class. For this reason we decided to work with the distribution and take the likelihood of the neutral class as the NLI-Score of the sentence in question. As the neutral class likelihood, the NLI-Score is a number between 0 and 1 with values approaching 1 being the best.

The presented approach can be used on prose with no further changes. In order to accommodate the metric for scripts, we propose to only use

¹<https://huggingface.co/roberta-large-mnli>

Principal:	Then you'll go to prison!
Sally:	<i>No! I'm not going home!</i>
Principal:	Well, then you'll go to jail!
Sally:	<i>No! I'm going home!</i>

Figure 8.2: When evaluating the sentence in bold, only the context in italic is taken into account.

the NLI-Score on the utterances of one character at a time. Therefore, we filter out the utterances of a single character and concatenate them in order to be used as context, as illustrated in Figure 8.2.

Furthermore, as we can see in Figure 8.2, the Principal's two utterances are in entailment as he repeats the same thing using slightly different words. On the other hand, Sally is contradicts herself by first saying she is not going home, but saying the exact opposite in her next utterance. As characters repeating the same message or contradicting themselves were not unusual in the generated scripts, we wish to identify the utterances in which it happens.

We deliberately do not add the utterances of other characters to the context, because in the dialogue setting, contradiction between characters is indicative of disagreement. We do not wish to discourage disagreement, as it makes a play more interesting in most cases. While this approach has its limitations (which will be discussed in Section 8.3), it works rather well in making sure the characters are self-consistent and non-repetitive in their utterances.

While evaluating scripts, we also measure the NLI-Score of the scenic remarks, as we want them to be consistent with each other. This can be easily accomplished by treating the scenic remarks as the utterances of a blank character. Thanks to this modification, if we use the NLI-Score on an output that does not contain any character utterances, the score works exactly like it would on prose.

8.2 Document-Level NLI-Score

When evaluating a collection of generated texts, we would like to have a metric which would tell us which texts are better. While we can measure the characteristics of the outputs, such as vocabulary size or perplexity described in Chapter 6, they cannot tell us whether one output is better than the other. Human evaluation is possible, but expensive.

To have such automatic metric, we propose two ways to extend NLI-Score that evaluates a single context-sentence pair introduced in Section 8.1 to a whole document. The first approach is the averaged NLI-Score which will be introduced in Subsection 8.2.1. The second approach – percentage below a threshold – will be presented in Subsection 8.2.2. Both versions can be applied to evaluating texts with varying lengths.

8.2.1 Averaged NLI-Score

We propose to measure the average neutrality per added sentence as the first way of extending the sentence level NLI-Score to the document level. The second sentence is compared with the first, the third with the first two, and so on. In case the length of the context exceeds the RoBERTa’s maximum input length (512 tokens), is truncated from the start.

To evaluate this approach, we scored 20 randomly human-written movie summaries from the ScriptBase dataset [Gorinski and Lapata, 2018], as well as 20 texts generated by vanilla GPT-2 medium by using the first sentence of the selected human-written movie summaries as the prompt. As we can see in Figure 8.3, human-written movie texts tend to be scored higher, which aligns with our assumption that human-written plot summaries are consistent and non-repetitive. On the other hand, the model-generated outputs have a much wider range of possible scores and it is possible to use this metric to rank them.

We also evaluated the suitability of averaged NLI-Score by correlating the NLI-Scores of generated plot summaries with the corresponding human-scored consistency ratings.² The resulting Pearson coefficient of correlation is 0.50, indicating a medium positive correlation between the two metrics.

The main disadvantage of this approach is that many of the generated summaries score around 0.7 - 0.8. In some cases, the resulting averaged score can look the same for a text that is uniformly scored at 0.7 (which indicates the text should be fairly consistent and novel) and a text that scores all over the scale, but is averaged to 0.7 (indicating there might be inconsistencies or repetition in parts the text). Therefore, even when some sentences have a low NLI-Score indicative of inconsistency, the average may be brought up by the remaining sentences.

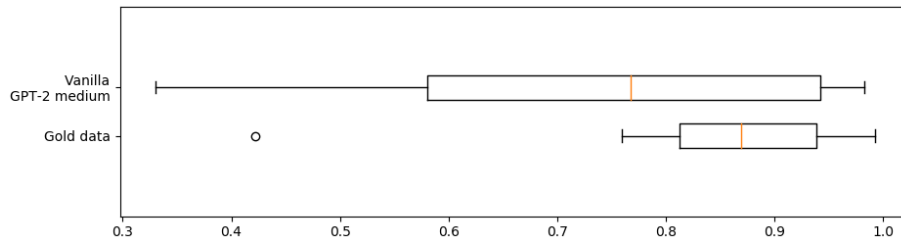


Figure 8.3: A comparison of the averaged NLI scores of human-written (gold) summaries and the summaries generated by vanilla GPT-2 medium model. The scores are in the range between 0 (worst) and 1 (best).

In this work, we will abbreviate the averaged NLI-Score to NLI-avg.

8.2.2 Percentage Below a Threshold

In order to address the main disadvantage of the averaged NLI-Score, we propose an additional approach that focuses on the prevalence of sentences

²For more information regarding this evaluation, please see Subsection 10.1.1.

which are repetitive or inconsistent. Instead of averaging all of the scores into a single number, we compute the percentage of sentences which scored below a given threshold.

The threshold is not fixed and can be changed specifically for a given set of evaluated texts. In our experience, 0.4 is a reasonable setting as the majority of sentences below that threshold are inconsistent with the preceding context. Based on the performance of the models or the quality of the outputs we wish to compare, we could raise this threshold in order to be more selective or lower it to see which texts performed poorly in this aspect. We do not recommend raising it to more than 0.6 as sentences with a neutrality likelihood higher than 0.6 are not necessarily inconsistent or repetitive. We present the scores for the same data as in the Subsection 8.2.1 in Figure 8.4. We use 0.4 and 0.5 as thresholds.

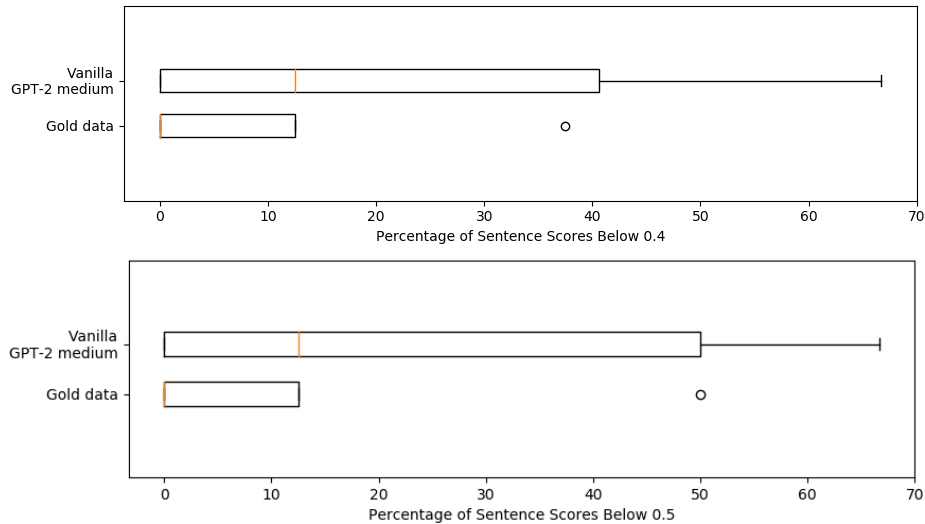


Figure 8.4: A comparison of the percentage of sentence NLI scores below 0.4 (top) and 0.5 (bottom) of human-written plot summaries (gold) and the summaries generated by vanilla GPT-2 medium model. The lower the percentage of inconsistent/repetitive sentences, the better.

It is advisable to use both score variants together side-by-side in order to interpret the results better. The averaged NLI-Score gives us the first impression of the output characteristics. On the other hand, the percentage below a threshold NLI-Score tells us whether or how frequently inconsistencies or repetition occur.

In this work, we will abbreviate the percentage below threshold NLI-Score to NLI-pbt.

8.3 Limitations and Potential Modifications

The most obvious limitation of the NLI score is that it depends on the probability distribution of a model that is not 100% accurate. We are fully

Premise	N	C	E
Eve is a widow.	0.014	0.984	0.002
Eve is a woman.	0.967	0.032	0.001

Figure 8.5: The NLI class distribution (Neutral, Contradiction, Entailment) comparison for two similar split sentences. The hypothesis is “Who lives with her alcoholic husband.” in both cases.

aware of this, however, it has not shown to be a serious issue in scoring.

We admit that neutrality might not be the best criterion, as it would score completely unrelated sentences highly. In this work, we are accepting that risk due to the fact that the models we examine and evaluate almost never output sentences unrelated to the context. Therefore, this is not really a concern to us in this case. However, if we wished to address this limitation, it would be possible to use cosine similarity of sBERT sentence embeddings [Reimers and Gurevych, 2019] to measure whether the sentences are similar enough to logically follow one another. Alternatively, a model could be finetuned for this task, we address this in Chapter 13.

Entailment might not always be a phenomenon to be avoided. In case we evaluated models with no issues stemming from repetition whatsoever, it would be possible to count the likelihood of entailment into NLI-Score as well.

While the metric can usually find inconsistent sentences with respect to a given context, it does not take internal inconsistencies within a sentence into account. An example of this would be “Eve is a widow who lives with her alcoholic husband.” In order to address this limitation, it would be possible to split the sentence into smaller chunks such as “Eve is a widow” and “who lives with her alcoholic husband.” and then to run the NLI-score of the sentences separately. In Figure 8.5, we can see that such approach could work even without coreference resolution.

Finally, as mentioned above, the modification of the NLI-Score for scripts does not take into account the utterances of other characters. If we decided to incorporate them into the score, it should include the likelihood of contradiction because two characters should not always agree among themselves. After all, disputes among characters are often one of the aspects which make a play interesting.

Chapter 9

Flat Generation

In this chapter, we propose two simple *flat* approaches to the task of theatre play script generation. For the purpose of this work, we define *flat* script generation as continuing a minimal human-written prompt that formally resembles a script, with no additional steps in between.

The two approaches to be described and contrasted in this chapter are using a vanilla GPT-2 XL model [Radford et al., 2019] with a specific setup (Section 9.1) and using a fine-tuned GPT-2 medium model (Section 9.2). Due to the availability of computational power, GPT-2 medium was the largest model we could fine-tune. This provides an interesting trade-off, since GPT-2 XL is more robust and therefore more suitable to natural language generation tasks in general. A detailed comparison of the two models can be found in Table 9.1. We can see that GPT-2 XL has twice as many layers, each of them almost 1.5 times bigger than its medium version. It also has more attention heads and nearly five times as many parameters in total. This also causes GPT-2 XL to be slower than GPT-2 medium. We benchmarked the speed of generation of both models on 5 different prompts, generating 500 tokens long texts. While GPT-2 medium needed 238 seconds on average, GPT-2 XL generated the text for 906 seconds on average, which is almost 4 times as much.

As GPT-2 XL is an ambitious baseline, we also include a vanilla GPT-2 medium model as a second baseline. The results of the evaluation are presented in Section 9.3, in order to get a direct assessment of the value of fine-tuning.

As we believe that generating a full theatre script is not feasible with the current language models in a flat structure, the models will be used under limited human supervision using the THEaiTRobot tool. This tool, allows the user to generate text based on a given prompt, regenerate the script from a given point, input a line of their own if necessary and prevent some lines from repeating. The goal of this chapter is to explore and compare the above mentioned approaches in order to choose the model to be used in this tool.

We note that no other models were considered for this task as other Transformer models usually do not share the 1024 tokens long context window of GPT-2. BART is an exception as its context window is equally long, however, it is pre-trained for the task of de-noising auto-encoding

which does not align with our objective as well as GPT-2. Additionally, as mentioned in Chapter 6, there are other variants of the GPT model, including its successor GPT-3 [Brown et al., 2020], or their open-source variants GPT-J [Wang and Komatsuzaki, 2021] and GPT-Neo [Black et al., 2021]. The open-source variants were only released after we started with the experiments. Unfortunately, GPT-3 was not freely accessible and the authors of this work had not received access to the GPT-3 model API at the time of carrying out this work. Furthermore, while it is possible to fine-tune 10 GPT-3 models per month using the OpenAI API, there are strict restrictions regarding the dataset size used for the fine-tuning: 2.5M tokens or 80-100MB. Our datasets are well above this limit.

Model	Layers	Hidden Layer Size	Attention Heads	Parameters
GPT-2 medium	24	1024	16	345M
GPT-2 XL	48	1600	25	1.5B

Table 9.1: A comparison of the model layer count, size of the hidden layers, the number of attention heads and the overall amount of parameters. Taken from Radford et al. [2019]

9.1 Vanilla GPT-2 XL Generation Setup

In order to use the vanilla GPT-2 XL model for script generation, it is necessary to prompt it with a script-like piece of text. Initially, the generation was prompted using excerpts from movies or plays, such as R.U.R., in order to assess the usefulness of this setup. In this case, the prompt was usually more than 5 lines long in order to convey the setting and the theme of the script. However, we felt it was too time-consuming for the user to always compose the beginning of the play manually. For this reason, we experimented with a setup that allows the user to have control over the topic of the play, but does not require a hand-written beginning of the script. We found it is enough to set the scene using a prosaic context description – either one whole sentence at the minimum, or a collection of 7 or more keywords. This context description should be followed by the characters of the play greeting each other. The purpose of the greeting was two-fold. First of all, it set the format of the text to look like a script which was then continued by the model. Second of all, it introduced the characters of the play. It is also possible to initialize the generation with something more sophisticated than a greeting, but the point of the setup is that it is not necessary to think of a lengthier utterance requiring more creativity. An example of how this setup would be used is shown in Figure 9.1.

During the decoding, we used nucleus sampling [Holtzman et al., 2020] with p of 0.9 and k of 50 in all experiments presented in this work as well as typical decoding [Meister et al., 2022] with *typical_p* of 0.2.¹ The temperature was set to 1.0 and the repetition penalty to 1.01. These parameters

¹Typical decoding was not used while generating the first AI-written script as it is a novel approach introduced in 2022.

Will I have a good day tomorrow?

John: Hello
Robot: Hi

John: I've had a rough day, what should I do to get out of this mess?
Robot: I was having fun, you can enjoy yourself, but you can't keep making me do this. I'm just a machine.

Figure 9.1: Script generation example using the Vanilla GPT-2 XL setup. The initial human-written prompt is shown above the dividing line, the generated outputs follow below.

were picked after brief experimentation as they produced sufficiently good results consistently. While the temperature and repetition penalty could be raised, it often caused the model to divert from the script format. The same parameters were used for the GPT-2 medium baseline as well.

The model's performance will be assessed and compared to the Fine-tuned GPT-2 medium model in Section 9.3.

Note that this setup was used for the first AI-written script (Section 3.1), complemented by two extensions: character orchestration (described and evaluated in Section 11.3), and extractive summarization (described in Section 11.4) which was used to keep important information from the script history that would otherwise not fit into the context window. We do not employ these extensions in this chapter to avoid any confounding effects.

9.2 Fine-tuned GPT-2 for Script Generation

In this section, we will discuss the the creation of a GPT-2 medium model fine-tuned for the task of script generation. In Subsection 9.2.1, we will discuss the data used for this purpose. Afterwards, we will show the model training setup in Subsection 9.2.2. Finally, we will comment on the generation and decoding approach in Subsection 9.2.3.

9.2.1 Data

Since there is no freely available dataset of theatre plays, we used a dataset of plays scraped from the internet. Due to copyright restrictions, the data cannot be shared outside the project. We provide an overview of the data volume and structure in Table 9.2. The play scripts obtained this way are mostly old, because modern theatre plays are usually not freely accessible on the internet.

Moving on to the nearest domain of movie scripts, we used a dataset that was created by Lison and Meena [2016] by transforming subtitles into scripts. It is a collection of movie and TV show scripts scraped from the internet, the overview of the size and makeup of the dataset can be found in

Domain	# Scripts	Avg. # lines	Avg. # sentences
Movies	1,067	783	2,537
TV Shows	6,057	314	902
Theatre	5,517	530	1,529
Total	12,641	446	1,310

Table 9.2: A brief overview of the script dataset we use for fine-tuning. The average number of lines and sentences is given per script.

Table 9.2. The movie portion contains a mix of genres, but mostly includes modern popular movies. Among the TV shows, the most represented one is Doctor Who with 810 episodes. There are 21 TV shows with 100 - 300 episodes. The remaining 300+ TV shows are represented by less than 100 episodes. The biggest downside of this dataset is that the data is not annotated too well. The scripts sometimes lack the scene boundaries and scenic remarks are difficult to filter out as they are mostly connected to the utterances.

By combining these two sources of scripts, we obtain a sizeable dataset of over 12K scripts.

9.2.2 Model Training

Using the data described in the Subsection 9.2.1, we fine-tuned the pre-trained GPT-2 medium model provided through the Huggingface library [Wolf et al., 2020] using its PyTorch interface [Paszke et al., 2019]. In order to train the model on the hardware available to us, it was necessary to restrict the context window to 768 tokens. As almost all scripts in our dataset are much longer than that, we split the scripts into smaller windows. It would be possible to simply cut the scripts into windows of size 768 tokens, but that would mean a loss of context as the model would have no information that the end of one window and the beginning of the next are connected. For that reason we decided to work with an overlap of half the size of a window – 384 tokens.² This allows the model to see the history of a given dialogue and to learn based on it. The *end-of-text* token, explained in Chapter 6, was not added to each window, but only to the real end of a script. Using this approach, we created 416,395 samples divided into training and validation sets using a 9:1 ratio.

We fine-tuned the data for 10 epochs using the Adam optimizer [Kingma and Ba, 2015] with a $1 \cdot 10^{-5}$ learning rate. We also use a linear learning rate

²Another option would be to employ the extractive summarization (see Section 11.4) which was eventually used at inference time in the final THEaiTRobot tool. However, as we wanted to perform the experiments without the interference of the summarization (which also was not yet stable at this time), we did not apply our summarizer to the training data. Nevertheless, as the design of the summarizer always keeps several last lines of the prompt/context and only removes some of the previous lines from the context, we do not expect that applying it to the training data would have a significant influence of the results.

scheduler with a 1000 steps warmup, see Chapter 6 for the explanation. Batch size 1 was used due to the memory capacity of the available GPUs. The fine-tuning took almost 19 days.

9.2.3 Setting Decoding Parameters

The decoding strategy and parameters can make a vast difference in model performance. We experimented with beam search, but found it does not make a noticeable difference in the output quality. Similarly as for the vanilla GPT-2 XL baseline, we used nucleus sampling with p of 0.9 and k of 50, and typical decoding with *typical_p* of 0.2.

Parameters such as temperature and repetition penalty can significantly improve or worsen the quality of the fine-tuned model outputs. If the values of these parameters are set too high, the model will be penalized for following the script format and fall back into prose. On the other hand, if they are set too low, the model might generate texts that are somewhat repetitive and bland, especially if they span over the context window of the model.

We performed a grid search to find the best combination of temperature and repetition penalty values for this model. We also experimented with the *no_repeat_ngram_size* parameter. If the value of this parameter is set to $p > 0$, n-grams of size p can only appear once. In our setting, we naturally want some n-grams to reappear at will, such as the character names.

In our grid search, we tried all the combinations of the following: temperature values 1.0, 1.005, 1.01, repetition penalty values 1.01, 1.05, 1.1, 1.2 and *no_repeat_ngram_size* values 0, 3, 4.³ Based on cursory manual checks of 5 outputs generated for 5 prompts with each setting, we picked out temperature = 1.005, repetition penalty = 1.1 and *no_repeat_ngram_size* = 4 as the best parameters. We use these parameters during the generation of all texts used in the evaluation in Section 9.3.

9.3 Evaluation

In order to evaluate the models, we generated 50 texts from each model. Each model received the same prompts – there were 10 prompts and 5 examples were generated for each by using different random seeds. Instead of writing our own prompts, we collected a sample of prompts that were added into the THEaiTRobot by its users. This way we achieved having diverse human-written prompts that consisted of a 1-2 sentence scene setting and two whole non-trivial character lines. The maximum length of the outputs was limited to 1024 tokens, which is the maximum output length of the GPT-2 model as stated in Chapter 6. During the decoding, we use the parameters selected as the best for the given model as described in Sections 9.1 and 9.2. For Vanilla GPT-2 medium, we used the same parameters as for the XL version of the model. The comparison of basic

³We have attempted using values of *no_repeat_ngram_size* higher than 5, which should theoretically keep most n-grams intact, but the quality of the text deteriorated significantly, the text lost its script format, and progressively increased generating whitespace.

Model	# Lines	# Sentences	# Words	Vocab Size	Entropy	Perplexity
Vanilla GPT-2 XL	54.44	145.12	524.66	2228	3.21	2.28
Vanilla GPT-2 M	1.46	29.26	159.86	3296	2.12	164.68
Fine-tuned GPT-2 M	31.42	101.40	449.70	4224	3.84	13.91

Table 9.3: A basic statistics comparison for script generation by different model variants. Perplexity is measured using vanilla GPT2-XL. The reported lengths of outputs are the averages.

characteristics of the outputs is shown in Table 9.3. The computed metrics are the average lengths of the output scripts in terms of lines, sentences, and words. We also report the vocabulary size by counting the number of unique words in the output data. The entropy is measured as defined in Chapter 6. Finally, we measure the internal perplexity of the texts using the Vanilla GPT-2 XL. This inherently means that when it measures its own perplexity, it will be lower than when measuring the perplexity of outputs of another model.

Textual Properties As we can see from the comparison, Vanilla GPT-2 medium baseline is substantially different from the remaining two models in the comparison in all aspects except for vocabulary size. It has a very low number of lines, because the model very rarely accepted the script format and generated prose in most cases. Neither of the other models shared this issue, so we do not believe it is caused by the prompts, but is rather a trait of the model. Vanilla GPT-2 medium also has the lowest entropy, which may indicate repetition; while having an extremely high internal perplexity as computed by the GPT-2 XL model. This indicates that the model outputs may not only be repetitive, but also very atypical when compared to human-written texts seen by the GPT-2 XL model while training.

When observing the Vanilla GPT-2 medium outputs closely, they resemble social media posts and comments more than stories, let alone play scripts. Some outputs are empty, as the model only generated the *end-of-text* token. Longer outputs tend to degenerate into repetition of various numbers and symbols. In one case, the model listed over a hundred geographical terms, such as countries, islands, and deserts in an arbitrary order.

The Vanilla GPT-2 XL baseline and Fine-tuned GPT-2 medium model both generated coherent texts while adhering to the script format. At the first glance, the baseline is slightly more repetitive, but this usually happens at the end of the texts. We can also see that the Fine-tuned GPT-2 medium model's outputs tend to be shorter. Based on the 50 generated scripts, we observed that the Fine-tuned GPT-2 medium model generates the *end-of-text* token at roughly the same point where GPT-2 gets into a repetitive loop. This was not an issue for the Vanilla GPT-2 medium model, as its outputs often did not reach a length where repetition would occur.

Model	NLI-avg	σ	% NLI-pbt	σ
Vanilla GPT-2 XL	0.47	0.17	50	0.23
Vanilla GPT-2 M	0.63	0.16	25	0.24
Fine-tuned GPT-2 M	0.48	0.16	48	0.23

Table 9.4: A comparison of the document-level NLI-Scores of the model outputs.

All models also introduce new characters at a fast and unnatural rate. This is especially true for our fine-tuned model, as we use it with a higher value of parameters designed to prevent repetition, which discourage the repeated appearance of tokens, but also names. We chose to raise the values of these parameters despite this because it increases the quality of the utterances as seen during the parameter grid-search described in Section 9.2. Furthermore, there are approaches to limiting or orchestrating the characters and we describe them in Section 11.

NLI-Score We also used our document-level NLI-Score metric to evaluate these outputs using its script form. While the full description of this metric can be found in Chapter 8, we remind that this metric is designed to measure the lack of inconsistency and repetition. The Averaged NLI-Score is a number in the range between 0 (worst) and 1 (best). The Percentage Below a Threshold NLI-Score tells us how many sentences scored below a threshold we consider acceptable – 0.4 in this case. Therefore, the lower the score, the better.

The results of the automatic NLI-Score evaluation can be seen in Table 9.4. The Vanilla GPT-2 medium achieves rather good results in this metric, but we attribute it to the limitations of the metric. More specifically, this model manages to output sentences that are very specific and most likely extremely far from any text seen in the MNLI dataset [Williams et al., 2018]. As there are very few character utterances in the outputs of Vanilla GPT-2 medium, the sentences are processed as if they were scenic remarks. The Fine-tuned GPT-2 medium and the baseline GPT-2 XL baselines achieved similar scores in both scoring approaches. Fine-tuned GPT-2 medium is seemingly slightly better, however given the standard deviations, this comparison is inconclusive.

Human Evaluation We also performed human evaluation. In order to reduce the load on human evaluators, we evaluated the Vanilla GPT-2 XL baseline and the Fine-tuned GPT-2 medium together with the hierarchical models in Section 10.4. The vanilla GPT-2 medium baseline has been excluded from the human evaluation as its outputs are consistently out of the domain of stories or scripts.

Discussion The biggest limitation of the Vanilla baselines is the low probability of plot twists. As the models were mostly trained on web pages, any plot twists or dramatic situations were unlikely and the script rather ‘converged’ to a point of undesirable repetition and the characters incessantly agreeing amongst themselves.

Another limitation is that the Vanilla GPT-2 models sometime generate texts far from the theatrical domain. In some cases, the model generates texts such as song lyrics or JavaScript-like code. This usually happened when the model was prompted by a sequence that was too short, such as an incomplete sentence of less than 10 words without the establishment of the script format. However, this phenomenon was significantly more prevalent in the medium version than XL.

Naturally, the limited context of the model also poses a challenge for all of the models evaluated in this chapter. We propose an approach to address this limitation in Chapter 10.

Chapter 10

Hierarchical Generation

As discussed in Chapter 9, the limited context window makes it difficult to generate long coherent texts. Therefore, we do not consider it feasible to generate a whole play using a single model with the current technology. To overcome this limitation, we propose a hierarchical system of generation – generating a script in two steps by progressively expanding the detail of the text. The input for this system is a short title of the play, provided by the user. In the first step, that title is expanded into the plot summary (or synopsis) of the play. This plot summary is then divided into sections. In the second step, the summary sections are used to steer the generation of the script from the beginning to the end. Figure 10.1 illustrates this process.

This approach was inspired by the work of Fan et al. [2018] which is discussed in Chapter 7. It is also closer to the way how most humans normally operate when writing long texts – they imagine or even write down the important points first and only then proceed to write the rough draft. The approach is also motivated by our unsuccessful attempts at generating a theatre play script as a sequence of dramatic situations (Chapter 4). While we do not explicitly work with dramatic situations in the end, a play synopsis may actually implicitly detail the plot as a sequence of dramatic situations to some extent. The hierarchical generation approach might thus eventually be in a similar spirit; however, we have not evaluated this assumption and therefore avoid making any strong claims in this respect.

In Section 10.1, we discuss how we approached the first step of our hierarchical generation pipeline – i.e., task of plot generation while conditioning on the input prompt. To the best of our knowledge, none of the currently available datasets were directly suitable for training of our models. Therefore, in Section 10.2 we explain how we prepared the training data for the second step of script generation conditioned on a given section of a plot summary. In Section 10.3, we show the training process of the hierarchical model. Finally, the evaluation of this approach and comparison to the Flat models from Chapter 9 is presented in Section 10.4.

10.1 Plot Generation

In the first step of the hierarchical setup, we wish to generate a plot summary conditioned on the title. In Chapter 9, we solely used the GPT-2 model, as it is the natural choice for long text generation. However, in this case the resulting text does not have to be as long as a script and thus it is possible to use other models. For this reason, we also compare the performance of GPT-2 to other two models: Pegasus and DistilBART. We chose Pegasus and DistilBART because they were pre-trained on the abstractive summarization task which is similar to our goals, but in a reverse direction. The process and the results of this comparison are shown in Subsection 10.1.1. After selecting the most suitable model, we experiment with the selection of training data according to the feedback by the theatrical subteam in Subsection 10.1.2.

10.1.1 Model Selection

Models In addition to the GPT-2 medium model, we wish to explore the performance of other Transformer-based models in this task and therefore, we compare it to Pegasus and DistilBART¹. Both Pegasus and DistilBART

¹The Pegasus and DistilBART models were fine-tuned in an unpublished collaboration by Saad Obaid, and we include them as a stronger baseline compared to a vanilla GPT-2 model. All the other aspects of this comparison, including the data preparation, GPT-2 fine-tuning,

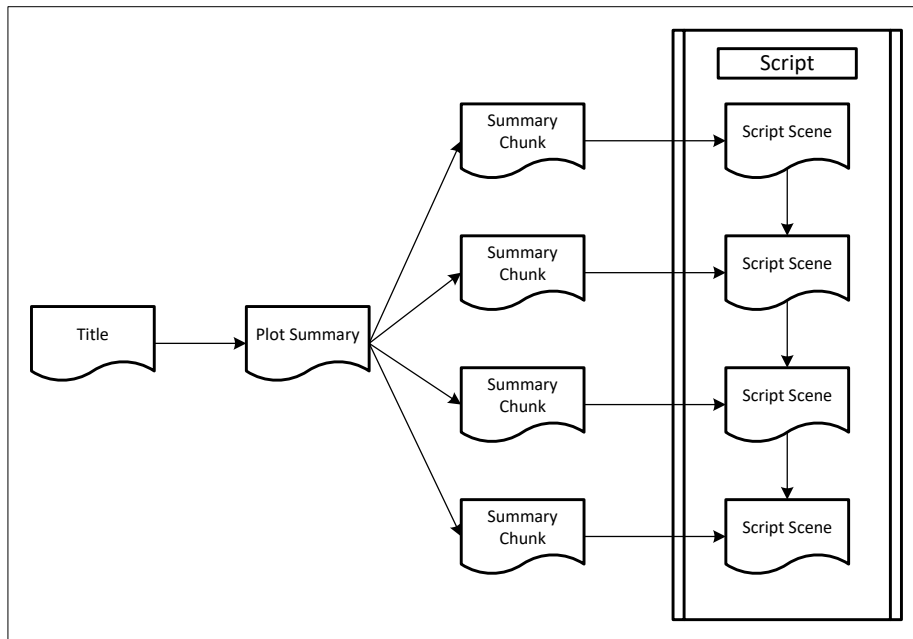


Figure 10.1: An overview of the hierarchical generation setup.

Source	Medium	# of Summaries	Avg # of Sentences	Avg # of Words
Fandom	TV	834	102.91	1523.07
IMDb	Movie + TV	3517	93.17	1525.42
Wikipedia	Movie	30580	21.39	390.31
Wikipedia	Theatre	2640	17.14	313.41
Sparknotes	Theatre	87	39.08	706.89
Book Summaries	Book	13080	21.02	415.63

Table 10.1: Overview of the short summaries dataset. The average amounts of sentences and words are measured per summary.

are pre-trained for the task of abstractive text summarization, meaning they reconstruct the main points of a given text using their own words. In our experiments, we aimed to symbolically reverse the direction and rather de-summarize the prompt, i.e. play title, into a plot summary. Both models were pre-trained on the XSum dataset [Narayan et al., 2018]. Finally, all of the models were fine-tuned on the same data.

Pegasus and DistilBART both have encoders, therefore, the prompts are passed to them as the input sequences and the plot summaries as the target sequences. However, GPT-2 does not have an encoder. For this reason, the training examples for GPT-2 are formatted as follows: the prompt, the *newline* token, and the plot summary.

Unlike in Section 9, we do not include a vanilla GPT-2 XL baseline but rather use vanilla GPT-2 medium. While vanilla GPT-2 medium proved to be insufficient for the task of theatre play script generation, it performs well in plot summary generation and we consider it a sufficiently strong baseline.

Data For fine-tuning, we use an in-domain dataset of theatre play plot summaries collected within the project. As the dataset is very small (less than 2600 examples), we supplement the data by including the movie plot summaries from the ScriptBase dataset [Gorinski and Lapata, 2018].² We also used two Kaggle datasets from the movie domain: Movie Plot Synopses with Tags [Kar et al., 2018] and Wikipedia Movie Plots [Robischon, 2018]. All of the movie datasets contain scraped the summaries from IMDb and Wikipedia. The amounts and average lengths of plot summaries in terms of sentences and words are presented in Table 10.1. Some sources, namely ScriptBase, Wikipedia, and SparkNotes also include a more detailed plot summary divided into multiple logical parts. We included these detailed summaries in the training data. The details can be found in Table 10.2.

Fine-tuning All models were fine-tuned for 15 epochs using an Adam optimizer with a $1 \cdot 10^{-5}$ learning rate and 1000 steps warmup. The models

as well as data analysis and evaluation were done by the author of this work.

²This dataset also includes movie scripts, which prove useful in Section 10.2.

Source	Medium	# of Summaries	Avg # of Sentences	Avg # of Words
IMDb	Movie + TV	8698	78.37	1317.37
Wikipedia	Movie	4814	21.97	409.96
Wikipedia	Theatre	1618	45.24	772.06
Sparknotes	Theatre	782	184.64	3496.77

Table 10.2: Overview of the detailed summaries dataset. The average amounts of sentences and words are measured per summary.

were fine-tuned on all data presented in Table 10.1³ and Table 10.2.

Evaluation We performed both manual and automatic evaluation for all mentioned models including a GPT-2 medium baseline. In the automatic evaluation, we observed the average lengths of the outputs. In order to provide more characteristics of the generated plot summaries, we also measured the output vocabulary size and the entropy of the generated texts. We generated 20 outputs of each model using 20 distinct prompts. The results are presented in Table 10.4.

The manual human evaluation was performed with 6 annotators rating 12 generated plot summaries. The summaries were limited to 200 subwords during the generation and truncated to the nearest whole sentence. This decision was mostly due to the fact that such length is enough to evaluate the quality, but not too much to discourage the annotators who were doing the evaluation in their free time. The annotators were shown one generated plot summary at a time. They were asked to score the generated plot summaries in terms of coherence, consistency, originality, title relevance, and overall impression (see Section 6.4 for the definition of the aspects).

As we can see in Table 10.3, the fine-tuned GPT-2 medium achieved the best ratings in coherence, consistency, and overall impression and the second best ratings in originality and relevance. Also, as we can see from the results, Pegasus and DistilBART both scored high in originality, but low in title relevance. This indicates that while the models produced interesting stories, they usually did not have much to do with the title. From analyzing the outputs, we can confirm that these two models often ignored the prompt and hallucinated their own story. Furthermore, there were topics and characters that reappeared rather frequently – approximately in one in every 3 texts. In case of Pegasus, it was Queen Elizabeth I, in case of DistilBART, it was the ‘racist, homophobic, and sexist Sheriff Dollard’.⁴

Neither of the evaluated GPT-2 medium models (fine-tuned or vanilla) shared these issues, although the vanilla GPT-2 medium sometimes generated out-of-domain text such as blog posts, newspaper articles, or texts strongly resembling Wikipedia articles. Due to its strong performance in all

³Except for the Book summaries portion of the dataset which was incorporated after these models were fine-tuned.

⁴We discovered that Sheriff Dollard is a character from the movie ‘To Wong Foo Thanks for Everything! Julie Newmar’.

Model	Coher.	Consist.	Orig.	Relev.	Ov. Impression
Vanilla GPT-2	2.7	2.8	2.6	2.7	2.6
Fine-tuned GPT-2	3.0	3.1	3.1	2.6	3.2
PEGASUS	2.8	2.8	3.0	2.1	2.8
DistilBART	1.9	2.0	3.2	2.0	2.9

Table 10.3: Results of human evaluation of synopsis generation models on Coherence, Consistency, Originality, Relevance and Overall Impression (1 to 5 points, higher is better).

Model	Avg. # Sentences	Avg. # Words	Vocab Size	Entropy
Vanilla GPT-2	38.10	285.80	1371	1.72
Fine-tuned GPT-2	29.32	536.74	1995	3.48
PEGASUS	14.80	281.40	1416	2.65
DistilBART	27.00	526.33	1182	2.43

Table 10.4: Basic characteristics of synopsis generation model outputs.

evaluated aspects as well as the described limitations of the other models, we selected the fine-tuned GPT-2 medium as the model to be used for the task of plot summary generation.

10.1.2 Dataset Selection

We received feedback on the outputs of the models from Subsection 10.1.1 from theatrical subteam. They found the plot summaries too “movie-like”. This is understandable as the movies take up the majority of the dataset, as shown in Table 10.1. Based on this feedback, we perform further experiments with data curation in order to make the plot summaries look less “movie-like” and more “theatrical”. All models mentioned in this subsection are GPT-2 medium.

1. (All data) The first model contained all of the previous data, however, all occurrences of words ‘movie’ or ‘film’ were replaced by the word ‘play’.

Model	NLI-avg	σ	% NLI-pbt	σ
Vanilla GPT-2	0.54	0.22	0.43	0.37
Fine-tuned GPT-2	0.84	0.15	0.11	0.16
PEGASUS	0.78	0.14	0.16	0.18
DistilBART	0.78	0.12	0.10	0.15

Table 10.5: A comparison of the document-level NLI-Scores of the model outputs including the standard deviations.

2. (Theatre & Books) The second model was trained only on the theatre portion of the dataset plus an additional dataset consisting of book summaries [Bamman and Smith, 2017].⁵
3. (Two-step model) The third model took the first model as a pre-trained base and fine-tuned it on the data from the second model.

We performed a manual evaluation on 10 outputs from each model, which consisted of both relative ranking and absolute scoring of the model outputs. The evaluators were three theatre professionals from the theatrical subteam (specifically from *DAMU* and *TvT*).

In the ranking task, the annotators were shown three outputs for the same prompt/title and asked to rank them from best to worst. There were many cases when the annotators could not decide between second and third place and therefore declared a tie between the outputs. For that reason, we only focus on the winning outputs. It turned out that the theatre professionals preferred model number 1. the most as it obtained 13 points. Models number 2. and 3. were tied with 8 points each ⁶.

All three annotators agreed on the best output in five cases. Moreover, at least two annotators named the same output as the best nine out of ten times. There were two cases when one or two annotators named an output the worst while the remaining annotator(s) named them best.

In the absolute scoring task, the annotators were shown one generated plot summary at a time. The scored aspects were slightly modified from the previous evaluation. The theatre professionals were generally satisfied with the link between the given title and the plot summary generated by a fine-tuned GPT-2 model, so we omitted this question. Instead, we added a question regarding real-world suitability of the text (marked as ‘Reality’ in the evaluation table). This decision was made because we received feedback that the stories sometimes do not adhere to commonsense logic. The remaining evaluated aspects were coherence, consistency, originality, and overall impression (all of these aspects are defined in Section 6.4).

Note that the difference is in the fourth evaluated aspect which was previously concerned with relevance to the title.

For example, there were deceased characters who performed normal tasks, 5 nickels having the same value as 1 dime,⁷ or a possibility of vaccination against carbon monoxide poisoning. In this manner, two annotators scored 10 outputs of each model.⁸ The results of the evaluation can be seen in Table 10.6.

We note that even though the scores are comparable or lower than in Section 10.1.1, it is not due to a decline in text quality, but rather due to higher expectations of theatre professionals.

⁵In order to use the book summaries dataset, it was necessary to filter out genres which were not suitable to our setup. Therefore, we excluded all non-fiction summaries.

⁶One annotator did not finish the task and left one comparison blank.

⁷1 nickel is 5 pennies (\$0.05). 1 dime is 10 pennies (\$0.10). Therefore, 2 nickels (\$0.10) are the equivalent of 1 dime, not 5.

⁸One annotator only participated in the relative ranking task.

Model	Coher.	Consist.	Orig.	Reality	Ov. Impression
(1) All data	3.25	2.6	2.45	2.9	2.6
(2) Theatre & Books	3	2.7	2.6	2.85	2.35
(3) Two step model	2.9	2.6	2.35	2.7	2.25

Table 10.6: Performance comparison of plot summary generation models, as evaluated by theatre professionals, on Coherence, Consistency, Originality, Reality and Overall Impression (1 to 5 points, higher is better).

Hadley brings in Red, but Red insists he knows nothing of Andy's plans. Becoming increasingly hostile and paranoid, Norton starts throwing Andy's sculpted rocks around the cell. When he throws one at Andy's poster of Raquel Welch, the rock punches through and into the wall.

NORTON: Well?
 RED: Well what?
 NORTON: I see you two all the time, you're thick as thieves, you are! He must 'a said something!
 RED: No sir, he didn't!
 ...
 NORTON: It's a conspiracy! That's what this is! It's one big damn conspiracy! And everyone's in on it! Including her!
 He sends the last rock whizzing right at Racquel. No smash. It takes a moment for this to sink in. All eyes go to her. The rock went through her. There's a small hole in the poster where her navel used to be.

Figure 10.2: An excerpt from a correctly aligned summary-script pair taken from 'The Shawshank Redemption'.

10.2 Scene Alignment

In order to expand a part of a plot summary into a script, we need training data that would provide script-summary alignments. However, to our knowledge, the only dataset of this kind is GraphMovie [Zhu et al., 2020] but unfortunately, it is in Chinese. As we are not aware of a dataset in English that would provide the alignments of script scenes to their summaries, we proceeded with creating a new dataset based on data described in Chapter 9 and Section 10.1. Ideally, we would use data from the theatrical domain, but unfortunately, the set of theatrical scripts and the plot summaries of plays available was mostly disjoint. Therefore, we proceeded with data from the nearest available domain – movies. The inputs were taken from the ScriptBase dataset [Gorinski and Lapata, 2018] which contains movie scripts and their metadata, including summaries divided into plot-based sections.⁹

⁹We intentionally avoid using the term 'scene summaries' as it would be misleading in this context – they often do not attempt to summarize scenes faithfully. A more detailed explanation will follow in Section 10.2.1.

In order to create the dataset, we represent the script scenes and summary section by embeddings and try to find the best alignment between them using a dynamic programming algorithm. The goal of this step is to align the summary sections with the corresponding scenes from the script to create a source and a target sequence for our script generation model in Section 10.3. The alignment is done in a many-to-one fashion where several script scenes can be aligned to a single summary section. This is done because the script scenes are divided into much smaller pieces than the summaries and therefore the events described in one summary section happen across several script scenes. An example of such alignment is shown in Figure 10.2.

In Subsection 10.2.1, we show how we prepared the data for the alignment. We discuss the alignment algorithm as well as two types of embeddings used to align the script scenes to the summary sections in Subsection 10.2.2. Using two types of embeddings as the representation allows us to create two sets of aligned data. The evaluation of these two sets is presented in Subsection 10.2.3.

10.2.1 Script Pre-Processing

The first challenge arises from the different granularities of the data. While the movie scenes are usually very brief and are often separated by cutting to a different camera view, the summaries follow a more concise pattern. On average, there are 146 scenes in a movie script, but only 6 sections in a movie summary. This is due to the fact that unlike the scenes in the summaries, the scripts are usually not divided into logical parts of the story, but rather focus on the camera shots. Furthermore, many scenes in the movies only contain one scenic remark describing the camera view or the visual setting of a given scene. There are also many scenes that only include a single utterance.

As we would like to level the ratio of script scenes to summary scenes, we decided to focus on removing or merging such very short scenes before the alignment process. We remove scenes consisting of one scenic remark as they usually do not bring any information and are often directly applicable to a theatrical setting anyway. Whenever possible, we try to merge scenes containing a single utterance with the previous scene, if the speaker is present in it. If not, we attempt to merge it with the subsequent scene using the same criterion. Finally, if the speaker of the single utterance is not present in either of the adjacent scenes, we have chosen to ignore the scene. We are aware that this might not be the best approach for all cases, but in the majority of roughly thirty scenes we examined manually prior to this pre-processing, this approach has improved the data quality. In fact, these short scenes are rarely tied to the main storyline described in the plot summary.

10.2.2 Script-Summary Alignment

Representation After we pre-processed the set of script scenes and a set of plot summary sections, we proceeded to the alignment. We experiment

with two representations of the texts:

1. **Bag-of-Words** (BoW) vectors representing script scenes and summary sections. Before computing the representation vectors, we remove stopwords from the processed texts.
2. **sBERT** embeddings [Reimers and Gurevych, 2019] computed by the mpnet-base-v2 model.¹⁰ See Chapter 6 for an introduction to the BERT architecture and sBERT in particular. In this case, it is important to note that as the context window of this model is limited to 384 tokens, we truncate the text if needed. This only happens for a small amount of script scenes, so we did not find it necessary to compute the embeddings over a longer span.

Alignment Algorithm After having the representation of each script scene i as a vector c_i and each summary section j as a vector m_j , we align each script scene to a summary section using dynamic programming with Algorithm 2.

In the forward pass, the algorithm computes a scene pair alignment score $s_{i,j}$. When computing the alignment score for a script scene i and summary section j , the algorithm first computes the cosine similarity of the vectors c_i and m_j . Then it looks up the best candidate alignment for aligning the preceding script scene ($i - 1$). It considers alignment scores of $i - 1$ and the same summary section (j) versus the preceding summary section ($j - 1$). The algorithm takes the maximum of these two scores and adds it to the computed cosine similarity between script scene i and summary section j , which is remembered as their alignment score.

The final alignment is computed in the backward pass, assuming the alignment of the last scenes to each other, and iteratively taking the best candidate alignment (a_i or $a_i - 1$) for the preceding script scene ($i - 1$).

Algorithm 2 Scene alignment.

Input: $\{c_i\}_1^N$	▷ Script embeddings
Input: $\{m_j\}_1^M$	▷ Summary embeddings
$s_{1,j} \leftarrow \cos(c_1, m_j)$	▷ Forward pass
for $i \in \{2, \dots, N\}, j \in \{1, \dots, M\}$ do	
$s_{i,j} \leftarrow \cos(c_i, m_j) + \max\{s_{i-1,j-1}, s_{i-1,j}\}$	
end for	
$a_N \leftarrow M$	▷ Backward pass
for $i \in \{N, \dots, 2\}$ do	
$a_{i-1} \leftarrow \operatorname{argmax}_{j \in \{a_{i-1}, a_i\}} s_{i-1,j}$	
end for	
return $\{a_i\}_1^N$	▷ Each c_i aligned to m_{a_i}

Filtering Post-processing There were many scenes that could not be aligned well, typically brief flashbacks that were not described in the plot summary and therefore had no optimal alignment. Furthermore, the summaries were not always segmented perfectly and some summary sections

¹⁰<https://huggingface.co/sentence-transformers/all-mpnet-base-v2>

Variant	# Movies	# Scenes	Script-summary ratio	Avg. # script lines
BoW	777	13,518	3.69	59.66
sBERT	777	14,655	3.40	54.98
sBERT Filtered	777	11,957	3.70	60.97

Table 10.7: Statistics of aligned summary-script scenes used for hierarchical generation (script-summary ratio is the average number of script scenes aligned to a single summary section).

Variant	# of Correctly Aligned Script Scenes	Accuracy
BOW	36 out of 66	55%
sBERT	50 out of 66	76%
sBERT Filtered	50 out of 62	80%

Table 10.8: Accuracy of the scene assignments using the BOW and sBERT embeddings, and after the filtering post-processing.

had no script scenes that could be aligned to them. For these reasons, we also post-processed the alignments made using the sBERT representation by going through the summary-script pairs and filtering out any script scenes that had a lower cosine similarity of sBERT embeddings than a given threshold. We set the threshold empirically to 0.3 as it filtered out most of script scenes unrelated to their assigned summary section, but still left us with enough training data. The filtering was not applied to the BoW aligned data as it would trim too many scenes.

A comparison of the average lengths of the script scenes aligned to the summary sections can be seen in Table 10.7. At the first glance, it may seem surprising that there are more scenes and lines per summary section after filtering out script scenes. However, this is because many 1-to-1 alignments were eliminated from the dataset.

10.2.3 Evaluation of Alignments

Our hypothesis was that the alignments based on sBERT embeddings would be significantly more accurate than their BOW version, because they can capture the semantic similarity as well as lexical similarity. In order to confirm or disprove this hypothesis, we performed a short manual evaluation. We sampled five movies using five randomly picked scenes for each and manually checked how many of the script scenes assigned to a given summary section were aligned correctly. As we can see in Table 10.8, sBERT embeddings are more accurate.

The biggest difference between the two representations in the alignment accuracy came in scenes longer than five utterances, with the sBERT model being substantially more accurate. The biggest challenge for both

approaches were very short script scenes that contained less than three utterances. Furthermore, many movies in the aligned data do not have a chronological storyline or are a blend of several alternating storylines. Our algorithm cannot align such unlinear movies with a 100% accuracy because it assumes that the script follows the events described in the plot summary and that those events appear in the same order as described in the summary. However, this assumption is helpful for us in order to compute alignments using an unsupervised algorithm with a polynomial time complexity.

Given the above mentioned challenges, we are satisfied with 50 out of 66 alignments based on sBERT embeddings being correct and 50 out of 62 achieved after the filtering post-processing. Moreover, we found that the filtering did not discard any of the correctly aligned script scenes in the examined data.

10.3 Script Generation

In the second step of our hierarchical setup, we aim to generate a theatre play script conditioned on the generated plot summary (see Section 10.1). Using the dataset created in Section 10.2, we fine-tune a GPT-2 medium model for the task of generating a script scene from a summary section. We take the GPT-2 medium model fine-tuned for flat generation from Section 9.2 as the base and proceed to fine-tune it on our aligned data. For each summary section as the input prompt, we train the model to generate the corresponding script scenes.

We fine-tune two models:

1. **GPT-2 medium Hierarchical Base** is fine-tuned on the dataset created by using sBERT embeddings (dataset creation is described in Section 10.2.2).
2. **GPT-2 medium Hierarchical Filtered** is fine-tuned on the filtered version of the sBERT dataset (described in Section 10.2.2).

Both models are fine-tuned on 777 movies using a 10^{-5} learning rate for 6 epochs with 1000 steps warm-up. The details of the datasets can be seen in Table 10.7.

The GPT-2 medium Hierarchical Base model is fine-tuned on data that includes the previous script scene, the summary section, and the script scenes aligned to that given section. This setup makes the training scheme identical to the inference scheme.

The GPT-2 medium Hierarchical Filtered model only sees the summary section followed by the script scenes that are aligned to it. The main goal of this model is for the output scripts to be more strongly conditioned on the input prompts. Furthermore, due to the filtering process, there might be gaps between script scenes wherever the alignment was not good enough, which is another reason for not putting the preceding scene into the training examples. We note that even though the Hierarchical Filtered model is not trained to see the context before the summary sections, the preceding scenes are not truncated or masked during the inference. This means that the model outputs may be influenced by the previous scene as well to some extent.

When generating, the generated plot summary needs to be split into smaller sections. As we did not train a specific summary splitter, we take individual sentences from the summary as the sections. We experimented with two ways of adding the summary sections to the context of the hierarchical script generation model:

1. Manually – the user can decide when to plug in the summary section into the context using the interactive THEaiTRobot generation tool, allowing for limited supervision.
2. Using the *end-of-text* cue – the hierarchical script generation models are trained to output the *end-of-text* token at the end of a scene. This can be used as a cue for introducing the next summary section.

The second approach does not work well on a vanilla GPT-2 model, as it very rarely generates the *end-of-text* token. We cannot determine which one of these approaches is better, each of them has its advantages and disadvantages and is suitable for different occasions. The first approach gives the most control to the user, which might be beneficial to user experience. On the other hand, the second approach requires no human supervision, which is useful for attempts to generate a script with zero human supervision.

We reuse the same decoding parameters as described for the fine-tuned Flat GPT-2 medium model in Section 9.2.¹¹

10.4 Evaluation

Same as in the evaluation of flat generation models in Chapter 9, we generated 50 texts using each model. The same 10 human-written prompts were used and 5 examples were generated for each prompt. For a more detailed description of the prompts, see Section 9.3.

In this section, we evaluate our fine-tuned models introduced in Section 10.3: GPT-2 medium Hierarchical Base, GPT-2 medium Hierarchical Filtered. We also include the two best models from Chapter 9 – the Vanilla GPT-2 XL baseline and the fine-tuned Flat GPT-2 medium model. We are omitting the Vanilla GPT-2 medium baseline as it mostly generated out-of-domain text (see Section 9.3).

Textual Properties The basic automatic comparison of the model outputs is presented in Table 10.9. Similarly as in Chapter 9, we measure the average lengths of the outputs in terms of lines, sentences and words. We also report the total vocabulary size, entropy, and perplexity as computed by the Vanilla GPT-2 XL model on the union of all generated texts.

The average length of the outputs reflects the training data and is therefore not surprising for us. Our fine-tuned models tend to produce shorter outputs, especially the hierarchical ones. This is an advantage, given that the average plot summary generated by the GPT-2 medium model for this task is roughly 30 sentences long, as shown in Table 10.4. We can see

¹¹Nucleus sampling [Holtzman et al., 2020] with p of 0.9 and k of 50, typical decoding [Meister et al., 2022] with *typical_p* of 0.2, temperature set to 1.005, repetition penalty set to 1.1 and *no_repeat_ngram_size* set to 4.

Model	Avg. # Lines	Avg. # Sentences	Avg. # Words	Vocab Size	Entropy	Perplexity
Vanilla GPT-2 XL	54.44	145.12	524.66	2228	3.21	2.28
FT GPT-2 M: Flat	31.42	101.40	449.70	4224	3.84	13.91
FT GPT-2 M: Hier./Base	19.50	59.20	308.26	4382	3.25	15.68
FT GPT-2 M: Hier./Filtered	15.30	69.54	370.50	5160	3.28	38.49

Table 10.9: A basic statistics comparison for script generation by different model variants. Perplexity is measured using vanilla GPT2-XL. *FT* stands for fine-tuned.

Model	NLI-avg	σ	NLI-pbt	σ
Vanilla GPT-2 XL	0.47	0.17	0.50	0.23
FT GPT-2 M: Flat	0.48	0.16	0.48	0.23
FT GPT-2 M: Hier./Base	0.58	0.15	0.35	0.23
FT GPT-2 M: Hier./Filtered	0.54	0.13	0.42	0.19

Table 10.10: A comparison of the document-level NLI-Scores of the model outputs. *FT* stands for fine-tuned.

from Table 10.9 that the average scene length is approximately 20 or 15 lines for the Hierarchical Base and Hierarchical Filter models, respectively. That means that after concatenating the scenes of the generated script, we can expect to reach ca. 600 or 450 lines. This is very close to the average number of lines in our training data – 530 as shown in Table 9.2. Therefore, when observing the output length, the hierarchical models are capable of producing a full-length play.¹²

NLI-Score In order to have an automatic overview of the output consistency and non-repetitiveness, we also measured the document-level NLI-Score using its script version (see Chapter 8. Similarly as in Chapter 9, we report both the averaged NLI-Score and the percentage below a threshold score. The threshold was set to 0.4, same as in Chapter 9. We report the average of NLI-avg and NLI-pbt of the individual outputs as well as the standard deviations in Table 10.10. The performance of all the models is comparable, with the fine-tuned Hierarchical Base GPT-2 model being slightly better in both variants of the score. However, due to the magnitude of standard deviation, we cannot claim it is the most consistent or the most non-repetitive of the models under comparison.

Human Evaluation We also performed a human evaluation on a selection of the generated scripts. We randomly selected five outputs from each of the four evaluated models (Vanilla GPT-2 XL, fine-tuned GPT-2 medium Flat, fine-tuned GPT-2 medium Hierarchical Base, and fine-tuned GPT-2 medium Hierarchical Filtered) and truncated the outputs to a reasonable

¹²At this point we do not comment on the coherence and quality of such resulting play.

Model	Times Best	Times Worst
Vanilla GPT-2 XL	10	7
FT GPT-2 M: Flat	4	7
FT GPT-2 M: Hier./Base	4	7
FT GPT-2 M: Hier./Filtered	7	4

Table 10.11: The results of the relative ranking task of the human evaluation. *FT* stands for fine-tuned.

length.¹³ Five amateur annotators volunteered for this evaluation. There were two tasks in the evaluation – relative ranking and absolute scoring.

In the ranking task, the annotators were shown four outputs (one from each model) generated using the same prompt. Due to our previous experience with the ranking task, we modified the instructions slightly. Instead of ranking all four model outputs, we asked the annotators to perform best-worst scaling [Santhanam and Shaikh, 2019] – to disclose which of the outputs they liked the most and the least. We present the results in Table 10.11. The Vanilla GPT-2 XL baseline was marked as the best ten times. In one of the examples, all five annotators agreed on the baseline being the best output. The closest model was the Hierarchical Filtered GPT-2 medium which was marked as the best seven times.

When examining the number of times when an output was named the worst, all model outputs but the Hierarchical Filtered GPT-2 medium were tied at seven. The Hierarchical Filtered GPT-2 medium was selected as the worst only four times.

This task shows that human judgement of the script quality is extremely subjective. In four out of five cases, an output selected by one or more annotators as the best was also selected by a different annotator as the worst. For this reason, we cannot state that one model is better than the others. We rather conclude that based on the result of this ranking task, Hierarchical Filtered GPT-2 medium had the most consistent performance out of the evaluated models. It ranked second in times named best while it was named worst the least amount of times out of all evaluated models.

In the absolute scoring task, the annotators were asked to score the outputs on a scale from 1 (worst) to 5 (best) with respect to four different qualities. The scored outputs were the same ones as used for the relative ranking task. The scored qualities and the corresponding annotator questions were coherence, consistency, originality, and overall impression (see Section 6.4 for the definitions).

In terms of coherence, the Vanilla GPT-2 XL baseline achieved the best score with the fine-tuned GPT-2 medium Flat model being the second. A slight decrease in coherence was expected, as some very short scenes

¹³The truncated length was not given by a specific number of lines or sentences, but rather by the available space on a slide, as a slideshow was used to distribute the instructions and generated scripts. The resulting length was usually 8-15 lines depending on the lengths of the lines.

Model	Coherence	Consistency	Originality	Overall Impression
Vanilla GPT-2 XL	3.00	3.40	3.32	3.24
FT GPT-2 M: Flat	2.92	2.76	3.44	3.00
FT GPT-2 M: Hier./Base	2.32	2.68	2.76	2.56
FT GPT-2 M: Hier./Filtered	2.60	2.64	3.32	3.20

Table 10.12: The results of the absolute human evaluation. Rated on a scale from 1 (worst) to 5 (best). *FT* stands for fine-tuned.

were omitted from the training data for the hierarchical models in order to improve the quality of summary-script alignments on the global scale (described in Section 10.2.1). While the majority of analyzed omitted scenes did not harm the script coherence, we hypothesize that some of these deleted scenes had a negative impact on the final coherence of the generated script. The Vanilla GPT-2 XL baseline was also scored as the most consistent out of the four models. The fine-tuned GPT-2 medium Flat model was perceived as the most original with the baseline Vanilla GPT-2 XL and the Hierarchical Filtered models tied for second place. The annotators had the highest overall impression of the baseline outputs, closely followed by the Hierarchical Filtered model.

We conclude that the performance of all our fine-tuned GPT-2 medium models is close to the performance of the Vanilla GPT-2 XL baseline, despite the XL version’s obvious advantage in the number of parameters. Out of our fine-tuned models, we selected the GPT-2 medium Hierarchical Filtered model as the best based on its results in the ranking task and its overall impression score.

Nevertheless, we did not compare to newer, more powerful language models, such as GPT-3, as these were not available to us at the time when we performed the experiments. It is therefore possible that using e.g. GPT-3 would bring even better results, as well as potentially alleviating the need for fine-tuning.

Chapter 11

Extensions

In this chapter, we discuss further extensions we made to the approaches discussed in previous chapters in order to increase the quality of the generated texts.

The first extension is the use of NLI (see Chapter 8) for quality assessment during the decoding. We present how to use it to filter the generated utterances in Section 11.1.

Next, we focus on characters in the play. They are a crucial element of the generated script as their dialogues create the backbone of the story. In Section 11.2, we focus on the characters in plot summaries, especially the influence of their names on the overall story. In particular, we experiment with training models where character names are either replaced by a placeholder or limited to a pre-determined set of character names. We also propose a method for character orchestration in Section 11.3. By orchestration we mean determining the next speaker based on the history of character utterances and the token likelihood distribution.

In Section 11.4, we get back to the issue of the GPT-2 model having a limited window and thus not being well suited for generating longer texts, and present our initial approach at employing automated summarization to get around this limitation.

Although the goal is to generate a theatre play script in Czech language, we eventually opted for performing the generation process in English, complemented by an on-the-fly automated translation to Czech, which we discuss in Section 11.5.

11.1 NLI Decoding Filter

The document-level version of NLI-Score described in Chapter 8 focuses on evaluating already generated texts. However, it is possible to use this metric during decoding in order to increase the quality of texts being generated. The application of NLI-Score directly in the generation process is inspired by feedback from the theatrical subteam. They found many generated texts interesting and inspiring, but the overall impression was often spoiled by a few blatant inconsistencies. For that reason, we experimented with including NLI-Score into the generation process to prevent

such inconsistencies from happening in the first place.

We include NLI by applying the following procedure: after a single sentence is generated, the generator briefly stops while the NLI model assesses the output. If the NLI-Score of the newly generated sentence with respect to previous context is below a given threshold, the sentence will be discarded and the generator will generate a different one. This process is repeated several times if necessary; after that the last generated sentence is kept.

As with the document-level NLI-Score described in Chapter 8, we can use this approach for prose and scripts and we will evaluate the usefulness of this extension on both. When using the script version, we focus on consistency of utterances separately for each character. The reason is that while we want the characters to be self-consistent, we do not want to filter out their disputes with other characters that might manifest as contradictions to the NLI system. Furthermore, the NLI model is not trained on texts composed of utterances by various characters. Filtering out the utterances of other characters potentially simulates a more consistent narrator. The consistency of the scenic remarks is evaluated using the already generated scenic remarks as the context.

The threshold parameter was set to 0.4 empirically, as our experiments showed that this filters out the worst inconsistencies but still allows the generator some degree of creative freedom. The number of attempts to generate a sentence with an acceptable NLI score was set to 10 in both experiments.

Evaluating the NLI Decoding Filter for Plot Summaries In order to assess the helpfulness of filtering plot summaries by NLI-Score, we generated five plot summaries using the NLI decoding filter and five without it conditioned on the same titles. All five titles used for the plot summary generation were human-written. Three lay annotators scored the generated plot summaries. We asked the annotators to score the generated scripts in terms of coherence, consistency, and overall impression as those aspects are the most affected by the filtering (see Section 6.4 for the definition of the aspects). Note that NLI-Score is aimed to recognize not only inconsistency, but also repetition. For this reason we specifically instructed the annotators to consider repetitiveness as part of the coherence score.

During the evaluation, we used the GPT-2 medium model fine-tuned on scripts of theatre plays, movies, and TV shows, as described in Section 10.1. The same model was used with and without the NLI decoding filter. The evaluation results are shown in Table 11.1.

Surprisingly, based on the results of the evaluation we conclude that adding NLI filtering can actually hurt the quality of generated summaries. While it somewhat increased their consistency, it did so at the expense of coherence and overall impression. We expected a small drop in coherence given that we are encouraging neutrality and unrelated sentences are neutral by definition.

In some cases, the re-generated sentences were internally inconsistent, which is something that NLI score is unable to filter out in its current version. It is also necessary to point out that the scores of the fine-tuned

Model	Coherence	Consistency	Overall Impression
Fine-tuned GPT-2	4.13	3.87	4.20
Fine-tuned GPT-2 + NLI	3.60	3.93	3.53

Table 11.1: The results of the human evaluation assessing the usefulness of using NLI decoding filter for summaries. Rated on a scale from 1 (worst) to 5 (best).

Context	In a crowded nightclub where a group gathers for a performance by singer Connie, a fight breaks up between a young woman named Mary Anne Williams and her boyfriend Nick. During the fight, she knocks Nick unconscious.
Continuation	When the police arrive and take Mary Anne to jail, she tries convincing them to help her prove Nick's innocence; but all attempts are unsuccessful.
Re-Generated	When the police arrive at her apartment they find that Mary Anne has left to get Nick from a hotel room.

Figure 11.1: An example of an acceptable sentence which was re-generated by the NLI decoding filter.

GPT-2 medium model without NLI decoding filter were higher than in the evaluation of various models trained for plot summary generation:

- coherence of 4.13 versus 3.0 previously
- consistency of 3.87 versus 3.1 previously
- overall impression of 4.20 versus 3.2 previously

as evaluated in Section 10.1. This indicates that the summaries were good without needing any changes. Furthermore, the baseline summaries did not suffer from any noticeable repetition.

NLI decoding filter was designed to prevent dead characters from coming back to the story or parents switching roles with their children. The results of our small experiment suggest that it accomplishes this goal, which is further supported by manual checking of the generated outputs during development as well as feedback from theatre professionals.¹ However, it seems that the 0.4 neutrality threshold had a higher recall, but a lower accuracy, which means that the re-generated sentences were acceptable and not necessarily inconsistent. An example of this can be seen in 11.1 – the continuation generated without NLI decoding filter is more suitable than its re-generated version after deploying the NLI filter. In order to use NLI decoding filter successfully, it would be desirable to decrease this threshold as it seems that the model's first choice tends to be better in most cases.

¹The theatre professionals gave us qualitative verbal feedback without any specific numeric rating in this case.

Model	Coherence	Consistency	Overall Impression
Flat GPT-2	2.53	2.87	2.60
Flat GPT-2 + NLI	3.67	3.60	3.40

Table 11.2: The results of the human evaluation assessing the usefulness of using NLI decoding filter for scripts. Rated on a scale from 1 (worst) to 5 (best).

Evaluating the NLI Decoding Filter for Scripts Similarly to the evaluation of plot summaries, we generated five scripts with and without the NLI decoding filter conditioned on the same prompts. The prompts consisted of a sentence from a generated plot summary and two character utterances. Three lay annotators scored the outputs. As in the evaluation of the helpfulness of NLI filtering for plot summary generation above, we were interested in the coherence, consistency, and overall impression scores.

The model used during this evaluation was the Flat GPT-2 medium described in Section 9.2 because it achieved the highest average scores in coherence and consistency in the human evaluation presented in Table 10.12. The same model was used for generation with and without the NLI decoding filter.

The results of the comparison are presented in Table 11.2. We can see that the outputs of Flat GPT-2 medium with NLI decoding were rated higher across all aspects. While a larger-scale evaluation would be desirable, the results of our small experiment suggest that using the NLI decoding filter in order to make the characters internally consistent helps increase the generated script quality.

11.2 Character Name Anonymization

The quality of the generated stories often varied based on character names in the title or sampled by the model during the generation. In some cases, the names can influence the stories and make them more interesting by bringing out the cultural context. For example, we have observed a model rename a character called “Vladimir” to “Vladimir Putin”. In a different case, characters called Kai and Kazuo were given an intriguing backstory as former Yakuza members. However, in most cases this did not prove to be an advantage and was rather a source of stereotypes and bias.

In this section, we examine what happens when we take the significance of the names away. We focus on plot summaries, as the effect of character name anonymization is greater than on scripts. We analyzed the original training set (see Section 10.1 for the sources and the makeup of the plot summary dataset) and suggest three modifications: replacing the character names by general numbered placeholders, by gendered numbered placeholders, and by a small set of pre-selected names. We experiment by fine-tuning three GPT-2 medium models on these modified training sets.

We evaluate these three models and compare them to the Plot Summary model trained on all data from Section 10.1.

Characters in Plot Summaries Before proceeding with the anonymization of the character names in the training data, we first analyze how many characters naturally occur in human-written plot summaries. In order to do so, we used a slightly modified software provided to us by Brezinová [2022], which finds and links all character name occurrences in a text. Therefore, for a given character, we know all names used to refer to them along with the position of every mention of the character in the text. The software uses Named Entity recognition (NER) from the SpaCy framework [Montani et al., 2022] to find the occurrences of character names. Our modifications to the software include small bug fixes, extensions of the title vocabularies (containing titles such as ‘Mrs.’ or ‘Mr.’), and the introduction of gender recognition.

In 5,719 out of 73,215 (8%) of the summaries in our dataset, no named characters were found. This was usually due to one of the following reasons:

1. There were no named characters in the summaries, they were only listed based on their gender, occupation, hobby, or other characteristic.
2. The summary is in English, however, the characters have foreign names (Indian, Chinese, etc.) and were not identified by SpaCy’s NER.
3. The character’s name or nickname is a common noun in English (e.g. Amber).

For further processing, we removed the summaries where we found no names to have a clearer overview of the rest of the data. As we are not aware of any general method that would attribute gender to the name of a character, we determined gender based on gazetteers of names included in the work of Urbanová [2021]. We only used two gazetteers: one listing male names and the other listing female names. Whenever a name was present in both of the male and the female gazetteer, or in neither of them, we classified it into the ‘Characters of Unknown Gender’ category. In Figure 11.2, we see the amount of summaries containing a given number of characters of a certain gender.² From the distribution, we can see that male characters and characters of unknown gender are represented in our data much more than female characters. This phenomenon is even more visible when looking at the average amount of characters of a certain gender in a summary shown in Table 11.3. While there are 1.68 male characters and almost 3 unknown gender characters in a plot summary on average, there are only 0.89 female characters, which is a little over a half of the male character prevalence. This is not an uncommon phenomenon, first pointed out by [Woolf, 1959] and later popularized in a comic made by Alison Bechdel.³

²The tail of the distribution has been cut off in order to make the histogram more readable.

³https://en.wikipedia.org/wiki/Bechdel_test

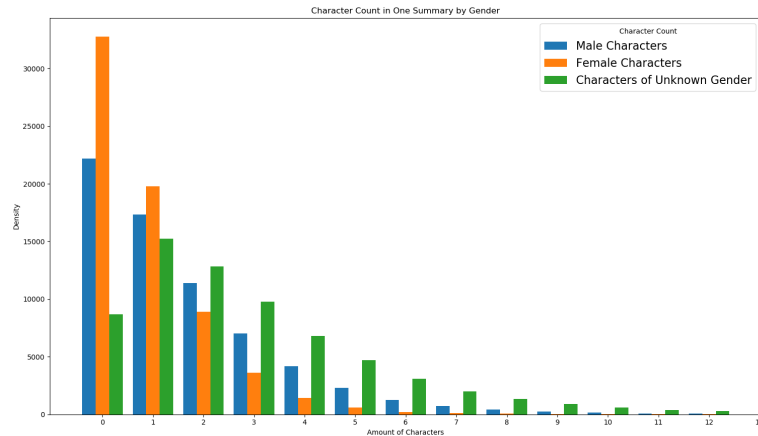


Figure 11.2: A histogram showing the distribution of character genders in one summary.

Masking Character Names As the first experiment, we went through the dataset of plot summaries from Section 10.1 and replaced every occurrence of a character name with a placeholder of the specific format: ‘<Character#>’, where # would be the numeric index of that given character. The titles were masked using the same approach. An example of the result of this transformation is presented in Figure 11.3.

After transforming the dataset, we proceeded to fine-tune the GPT-2 medium model using the same parameters as for the plot summary GPT-2 medium model from Section 10.1.

Including Gender Information in the Placeholder We experimented with an enhanced set of placeholders which would include information about the character’s gender, unlike the previous version. After gathering all character occurrences, we tried to determine the character’s gender using one of two rules depending on the available context information. The first approach consists of dividing the gazetteers of address expressions or personal titles, such as ‘Mr.’, ‘Mrs.’, or ‘Dr.’, that originally came with the name recognition script into three categories: male, female, and universal. We would use this approach whenever a character was addressed by a title and their last name. If the character was mentioned by their first name, we proceeded with the gazetteer lookup using the gazetteers from the work of Urbanová [2021]. Based on the gazetteer lookup, we classified the character into one of three categories: male, female, or unknown. In case the character name was found in both or in neither of the gazetteers, we proceeded with the ‘Character’ placeholder. If the gender was known to us, we used a placeholder that included the gender information: ‘Man’ or

Gomez Addams laments the 25-year absence of his brother Fester, who disappeared after the two had a falling-out. Gomez's lawyer Tully Alford owes money to loan shark/con artist Abigail Craven, and notices that her adopted son Gordon closely resembles Fester. Tully proposes that Gordon pose as Fester to infiltrate the Addams household and find the hidden vault where they keep their vast riches.

<Character0> laments the 25-year absence of his brother <Character6>, who disappeared after the two had a falling-out. <Character0>'s lawyer <Character1> owes money to loan shark and con artist <Character2>, and notices that her adopted son <Character7> closely resembles <Character6>. <Character1> proposes that <Character7> pose as <Character6> to infiltrate the <Character0> household and find the hidden vault where they keep their vast riches.

<Character0> laments the 25 - year absence of his brother <Character5>, who disappeared after the two had a falling - out. <Character0>'s lawyer <Man2> owes money to loan shark <Woman3> and notices that her son <Man6> closely resembles <Character5>. <Man2> proposes that <Man6> pose as <Character5> to infiltrate the <Character0> household and find the hidden vault where they keep their vast riches.

Avery laments the 25-year absence of his brother Jordan, who disappeared after the two had a falling-out. Avery's lawyer Richard owes money to loan shark/con artist Lisa, and notices that her adopted son Paul closely resembles Jordan. Richard proposes that Paul pose as Jordan to infiltrate the Avery household and find the hidden vault where they keep their vast riches.

Figure 11.3: An example of the original plot summary text compared to all its modifications: replacing the names with gender-neutral placeholders, replacing the names with gendered placeholders, and replacing the names with a limited set of pre-determined names.

'Woman'. An example of such modified text is shown in Figure 11.3. We can see that atypical names such as Gomez or Fester are not included in the gazetteers and therefore are placed into the unknown category.

We fine-tuned a GPT-2 medium model on this dataset using the same approach and parameters as in the previous masked character model.

Limited Set of Pre-determined Character Names Instead of using placeholders, we experimented with renaming the characters choosing from a pre-determined limited set of names. We start with the gendered masked dataset described above and replace the masks by real names. Ten names were prepared for each group of characters: men, women, and characters of unknown gender. We picked out names that are considered gender-neutral for the characters of unknown gender.⁴

⁴Male name list: James, Richard, Mike, David, Steven, Paul, Kevin, George, Tim, and Eric. Female name list: Mary, Susan, Lisa, Nancy, Jessica, Emily, Karen, Betty, Rebecca, and Amy. Gender-neutral name list: Avery, Riley, Parker, Cameron, Jordan, Charlie, Rowan,

Model	Male	Female	Unknown	All
Training Data	1.68	0.89	2.96	5.52
Original	0.63	0.80	0.67	2.10
Masked	N/A	N/A	31.24	31.24
Masked + Gender	14.15	10.38	11.38	35.91
Limited Names	2.66	2.00	0.90	5.55

Table 11.3: A comparison of how many characters appeared on average in a plot summary produced by a given model (or the training data). All models in this table are GPT-2 medium. The model presented as ‘Original’ is the model from Section 10.1.

Given that we only had a limited set of names, we eliminated all plot summaries which contained more than ten characters from any group in order to avoid having masks in the data. This made the training set smaller by roughly 2,000 examples.

Similarly as for the previous two datasets, we fine-tuned a GPT-2 medium model using the same parameters as before.

Evaluation For the purposes of automatic evaluation, we generated twenty plot summaries by each model using the same human-written prompts. The prompts did not contain any names in order to be universal for all model variations.

To get an overview of how the data modification impacted our model performance, we focus on the average number of unique characters in the generated script. We also include the training data in this comparison, as we aim to retain its distribution of characters in the generated outputs.

We simply count the number of distinct placeholders for each gender. To count the number of characters present for the outputs of any models that do not use placeholders, we post-process the generated plot summaries with the algorithm that was used for the gendered masking of the data (see above).

As we can see in Table 11.3, the models trained on a dataset where characters were masked introduced a substantially larger numbers of characters in the summaries. The same character mask was very rarely used twice or more times. Instead, the model simply incremented the character index in the majority of cases. Due to this, the plot summaries generated by the masked models are often confusing and make little sense.

The original GPT-2 medium fine-tuned for plot summary generation from Section 10.1 generates fewer characters than can be found on average in the training data. The ‘Limited Names’ GPT-2 model is the only model that generates about the same number of characters on average as what naturally occurs in the training data. Interestingly, even though names of Unknown gender are present in the training data, the model generates

Blake, Jamie, and Taylor.

Model	Coherence	Consistency	Character Faithfulness	Overall Impression
GPT-2	4.13	4.00	3.60	3.40
GPT-2 + Limited Names	4.40	4.60	4.40	4.33

Table 11.4: The results of the human evaluation assessing the usefulness of limiting the character names to a pre-defined amount.

fewer names of this type.

As the use of numbered placeholders introduced more characters than could be followed by a reader, we do not include the outputs of the masked models into the human evaluation. Instead, we compare the model trained on data with pre-determined character names to the model from Section 10.1, which is trained on unmodified data.

In the human evaluation, we asked three lay judges to score five outputs in the following aspects of coherence, consistency, character faithfulness, and overall impression (see Section 6.4 for the definitions of the evaluated aspects).

The human evaluation results displayed in Table 11.4 suggest that the GPT-2 medium model trained on data with a limited set of character names performs better in all evaluated aspects. While we did expect an increase in character faithfulness, the improvement in the other aspects was more notable than we anticipated. We hypothesise that while narrowing down the set of character names helped, there might be an additional reason for the performance improvement. When we limited the character names, we also excluded training examples with more than 10 characters of any given gender. This could have potentially improved the quality of the plot summaries in the training set, which reflected upon the quality of the generated stories. It would be possible to fine-tune a model on the filtered data without limiting the character names. However, we did not train such a model due to time constraints. Moreover, using a limited set of names in the plot summaries makes extracting character names simpler and more reliable, as they could be extracted via gazetteer lookup. This is very useful for the purposes of character orchestration described in Section 11.3.

11.3 Character Orchestration

A common problem for all script generation models examined in this work is that they do not work well with the set of characters in a script format. For example, the Vanilla GPT-2 XL baseline tends to forget some characters if more than two characters are introduced. Furthermore, if the name of the character includes a number, such as ‘Robot 1’ and ‘Robot 2’, the model is very likely to introduce ‘Robot 3’, ‘Robot 4’, and so on.⁵

⁵The same problem appeared when using numbered placeholders in Section 11.2.

Model	Avg # of Characters in a Script
Vanilla GPT-2 XL	4.22
Fine-tuned GPT-2 M: Flat	28.44
Fine-tuned GPT-2 M: Hier./Base	15.54
Fine-tuned GPT-2 M: Hier./Filtered	12.28

Table 11.5: The average number of unique characters in a generated script.

Our fine-tuned models usually introduce new characters at a very fast pace in general. This is not only confusing for the audience, but can also be expensive should the generated play be staged. Consequently, the theatre professionals are highly motivated to tightly control the number of characters that appears in a play script. The full extent of this problem is demonstrated in Table 11.5, where we computed the average character count on the data used in automatic evaluation in Section 10.4.

Since it is obvious that all models deviate from the desired character use across scripts, we would like to propose a flexible approach to limit the amount of characters present in a single piece of a script. While there are simpler and more elegant solutions to limiting the amount of characters, we designed one that honors token probability distributions generated by our model.

Similarly, a more robust approach could account for the character dynamics outside of the script lines, such as Versu [Evans and Short, 2014] or Dramatic Networks Moretti [2014]. We did not experiment with such orchestration methods in our research, however, it could be a direction for future work.

Before the generation starts, the character names in the input prompt (or the input summary in the case of hierarchic models) are isolated and tokenized.⁶ Then we proceed by building a trie over the tokenized character names. As there is no guarantee that there would be a single root node for an arbitrary set of names, the *newline* token is in the root of the trie. Moreover, to ensure that nothing will be added to the generated character name, we add the colon token to the end of every branch in the trie.

During decoding, whenever the *newline* token is generated, we start modifying the next token probability distribution. At every step, the probability of all tokens is set to zero, except for the ones that are connected to the current node in the trie.

Furthermore, for every character, we keep track of the number of lines that have passed since their last utterance. Let c be the number of lines since a given character has last spoken. In most steps, we do not know the full character name because it has not been generated yet. Therefore, given the character utterance history, we look for the number of lines that have passed since the last time the speaker name started with the subsequence of

⁶In case of scripts, the character name isolation is trivial and can be easily done using regular expressions. We used the name recognition script from Section 11.2 for the extraction of character names from plot summaries.

the trie leading to the current token. For example, in the context of speaker history ‘Amanda, George, Stacy’, we match the evaluated subsequence ‘Aman’ to Amanda and assign $c = 2$. Then, we multiply the likelihood of the next token of the character’s name by 2^c . This process is repeated until we reach the end of the trie – the colon. Afterwards, the distribution is not modified in any way until the next *newline* token is generated.

We only keep the history of last ten utterances for the purpose of processing speed, so if a character has not spoken for more than ten lines, c is capped at eleven. We repeat this for each character that can potentially be generated. Afterwards, we normalize the distribution using the softmax function. This process is illustrated in Algorithm 3.

Algorithm 3 Character limitation at one step of the generation. This algorithm is triggered whenever a *newline* token is generated and called after the generation of each token until colon is generated.

Input: $\{prob_i\}_1^{50257}$ ▷ Distribution of token likelihood
Input: *currentNode* ▷ Current node of the traversed trie
Input: $\{charHistory\}_1^{10}$ ▷ Last 10 characters to speak
if *currentNode* is null **then**
 currentNode \leftarrow *root*(*T*)
end if
allowedTokens \leftarrow *children*(*currentNode*)
newProb \leftarrow $\{0\}_1^{50257}$
for *token* \in *allowedTokens* **do**
 c \leftarrow *last*(*token*, *charHistory*) ▷ Computed as the lines since the subsequence leading to the given token appeared among the speakers

 *newProb*_{*token*} \leftarrow *prob*_{*token*} \ast 2^c
end for
return softmax($\{newProb_i\}_1^{50257}$) ▷ Updated token likelihood distribution

By analyzing over twenty generated scripts, we found that using this approach gently re-introduced characters that would have been forgotten otherwise as well as maintained the desirable amount of characters set by the prompt writer.⁷

11.4 Summarization

The variant of the GPT-2 model which we are using has a limit of 1,024 subword tokens, within which both the input prompt and the generated output must fit. The typical solution is to crop the input at the beginning so that it fits into the window with sufficient space for generating the output. However, this means forgetting potentially important information from the input prompt and the previously generated text, which can lead to an unwanted continual topic drift and also to generating contradictory text; the text is still locally consistent, but as a whole it may be inconsistent.

⁷No quantitative human evaluation was performed because the alternatives either forgot all characters but two, or added around ten new characters on average.

To handle this issue, we introduce automated extractive summarization into the process, hoping that the summarization algorithm will identify the most important pieces of information to remember. Whenever the input for GPT-2 (the input prompt + the so far generated script) exceeds a preset limit of $M = 924$ tokens,⁸ we summarize the input using TextRank⁹ [Mihalcea and Tarau, 2004] before feeding the input into the GPT-2 model:

- We keep all lines within the last $R = 250$ tokens from the input¹⁰ to ensure local consistency.
- We summarize all the preceding lines into $N = 5$ lines (while keeping their original order) to ensure global consistency.
- We concatenate the summary and the kept lines.
- If the resulting text is still longer than M tokens, we crop it at the beginning to M tokens.

This simple solution was intended as an initial basic fix, to be further tested and potentially improved. However, this research branch eventually remained mostly unexplored, as we needed to prioritize other aspects of the project.

Nevertheless, given our current knowledge and experience, we actually believe that a sufficiently advanced approach to the summarization aspect of the setup might be a promising path to considerable improvements of the outputs (see Section 13.1.2), especially in terms of overall consistency as well as internal self-consistency of individual characters.

11.5 Machine Translation

The GPT-2 model operates on English, while we want to generate a Czech script. We therefore automatically translate the generated script using the CUBBITT [Popel et al., 2020] neural translation model; a similar approach has been used in live improv theatre performances by Mirowski et al. [2020]. The translation result is displayed alongside the English original in THEaiTRobot.

An alternative would be to train or fine-tune the models on Czech data and perform the generation directly in Czech, which we have contemplated and briefly experimented with. However, the availability of in-domain Czech data is very limited, and the process of fine-tuning GPT-2 for a different language is rather complex and computationally demanding, while at the same time the quality of the automated translation is very high. We thus eventually decided that the expected gain from having Czech models

⁸Most script lines in our setting fit within 100 tokens, so ensuring there is space for generating at least 100 tokens means that usually the model will generate a complete line, ending with a newline symbol; in case the generated line is too long, it is simply cut off once the limit of 1024 tokens is depleted.

⁹We use the `pytextrank` library with minor modifications to reflect the specific structure of our inputs, so that the algorithm returns N most important (potentially multi-sentence) *full lines* from the script instead of just N most important *sentences*. We set `limit_phrases=100`.

¹⁰We find the first newline symbol in the last R tokens and keep all the lines after it.

English	Translation	Gloss	Note
Master: Whoa!	Páni!	Whoa!	Discarded character.
Whoa!	Páni!	Whoa!	Correct translation.
Master	Mistře	Master _{vocative}	Incorrect case.
the Master	Mistr	Master	Correct translation.

Table 11.6: A real example of the translation issues and the way we deal with them.

directly instead of using automated translation is not as high (and may be even negative) and did not follow that research path further.

11.5.1 Character Translation

As the translation tends to discard character names from the lines, we handle them by identifying them in the input and translating them independently. Specifically, if the line contains a colon (:), we assume the text preceding the colon is the character name or identifier, and translate that separately.

Moreover, we prepend the character with “the” so that the translation system favours the base form of the character; as the system is trained to translate full sentences, not words, and as Czech has high inflection, a single character reference may be translated e.g. as a vocative.

After translation, casing is matched to the original casing.

We show a real example from the script of the first generated play in Table 11.6, demonstrating the issues we deal with and how we fix them.

11.5.2 Gender and T-V Distinction

There are also further complications with the translation model. The model we use is tuned for news text, not theatre scripts, and translates each sentence independently. This leads to various issues, including errors in morphological gender (which should pertain to the characters) and variance in the honorific T–V distinction¹¹ (which may vary but should be consistent for each pair of characters). Not only do these depend on extra-sentential context, but also the mostly third-person news texts deal with these phenomena more easily, conversely from the mostly first-person and second-person scripts.

Gender is typically morphologically marked for all persons in Czech and the verb agrees in gender with the subject, which for 1st and 2nd person is typically ellided (pro-drop) and one rather has to correctly deictically understand the situation, discerning which character speaks to which other

¹¹In Czech, when addressing a person of a higher position or respect, the plural forms of second person pronouns and verbs are used instead of singular.

character.¹² As for the T-V distinction, this is not marked for 3rd person at all, and there are similar deictic issues as with gender; additionally, guessing the gender of a character is comparably easier than guessing their T-V relation (not only computationally but also in real life).

We have tried implementing a set of tools to heuristically identify the characters' deixis, genders and T-V relations and to appropriately switch the morphological markers, based on an existing Czech morphological analyzer and generator, MorphoDiTa [Straková et al., 2014]. While our codes are available online,¹³¹⁴ we have not been able to reach a satisfactory accuracy with them and therefore we have not integrated them into the tool.

¹²This may even be a case of deixis going into future; when generating a line said by character A to character B, we might not be sure who character B is yet until the generator chooses the character to speak the next line, i.e. *after* having already had to generate the translation.

¹³https://github.com/notsusual/gender_correction_THEAITER

¹⁴<https://github.com/notsusual/post-editing-corrector-ufal>

Chapter 12

THEaiTRobot Demo

Besides the obvious motivation by challenge and curiosity, the main goal of the project is to truthfully show and explain the current state of artificial intelligence research and capabilities to the general public. Therefore, each performance of a play created within the project is followed by a discussion with the audience, in which the project members explain the principles of operation of the THEaiTRobot tool, including both its strengths as well as its shortcomings and limitations. In this way, we are trying to demystify artificial intelligence, to bring it closer to people and to put them into direct contact with it.

To further stretch the project main goal, we have created a bilingual¹ English-Czech demo version of our THEaiTRobot tool, which we make freely available via the project website.² By this, we wish to allow the public to experience an even more direct contact with the language generation system. The users can experiment with the tool on their own in a simple and guided way, exploring its strengths and weaknesses and gaining practical familiarity with and understanding of natural language generation via the means of experiential learning. Therefore, the demo does not feature all of the options of the full tool used within the project, but rather offers the user with a limited range of simple controls. This way, they can easily operate the tool without the need of studying detailed guidelines.

For interested users, the demo also features a simple explanation of the principle of the tool, and links to the project website³ for further information resources in the form of texts and videos.

The demo features two versions of the tool:

THEaiTRobot 1.0 (Section 12.1) is a simpler flat wrapper of the GPT-2 model fine-tuned to generate theatre scripts (as described in Chapter 9) with minimal adjustments, generating a theatre script from a scene beginning.

THEaiTRobot 2.0 (Section 12.2) is a two-phase hierarchical generation system (as described in Chapter 10), employing two GPT-2 models

¹Adapting to a different language would be very simple, as the translation system is external to the generator. It can easily be replaced by changing the respective call to the external API.

²<https://theaitre.com/demo>

³<https://theaitre.com/>

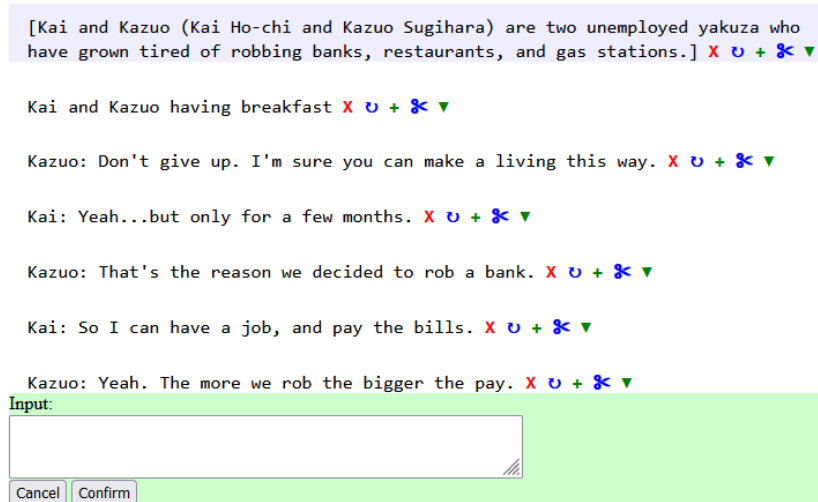


Figure 12.1: A screenshot from the internal variant of THEAiTROBOT. The cross discards the given line and all subsequent lines and generates an alternative continuation. The arrow creates an alternative line while keeping the script continuation. The plus symbol generates and inserts a line, while the scissors symbol deletes it without any changes to the continuation in both cases. Finally, the triangle symbol allows for human input that prompts the regeneration of the continuation.

fine-tuned on custom datasets, generating a play synopsis from its title and subsequently using the synopsis to generate the play script.

Internally, various other versions of the tool were used at various stages of the project to eventually generate the actual scripts for the first and second play, catering to the needs and requests of the current operator or operators of the tool within the project. However, the internal tool is by far not as user friendly and easy to use as the public demo. The internal tool uses the same backend as the demo, but features a richer set of controls, allowing the user e.g. to manually insert a new line into the script; see the example in Figure 12.1 or a demonstration video on YouTube.⁴ We were always trying to find a balance between giving the user as much control as they wished for, but at the same time not giving them too much control so that the resulting script is still mostly produced by the automated system, not by the operator. For the demo, we tried to provide only the most essential controls while keeping the interface as simple and user-friendly as possible, building on the experience of our operators with the internal tool. We primarily focus on the demo variant of the tool in the rest of this chapter.

⁴<https://youtu.be/ksrZouM7Wyg>

The screenshot shows the THEAiTRobot 1.0 interface. At the top, there's a navigation bar with 'THEAITRE' logo, 'THEAiTRobot 1.0' (selected), 'THEAiTRobot 2.0', and a language switch 'Přepnout do češtiny'. Below the navigation bar, the main heading is 'CAN A ROBOT WRITE A THEATRE PLAY?'. Underneath, there's a section titled 'SAMPLE SCRIPT' which displays a dialogue between Polonius and Hamlet. The dialogue is presented as a list of lines with character names on the left and text in rounded rectangular boxes on the right. Below the sample script, there are two cyan arrows (left and right) and a green 'GENERATE FURTHER' button. The second section is titled 'INPUT YOUR OWN SCRIPT' with the subtitle 'AI will generate a continuation'. It contains three input fields: 'Scene' (with a placeholder 'Write here a description of the starting situation.'), 'Name of first character' (with a placeholder 'Write here what the first character says.'), and 'Name of second character' (with a placeholder 'Write here what the second character says'). A green 'GENERATE FURTHER' button is located at the bottom right of this section.

THEAITRE THEAiTRobot 1.0 THEAiTRobot 2.0 Přepnout do češtiny

CAN A ROBOT WRITE A THEATRE PLAY?

SAMPLE SCRIPT

Scene Polonius speaks to the king. Enter Hamlet.

Polonius I hear him coming; let's hide, sir.

Hamlet To be or not to be; that is the question.

← → GENERATE FURTHER

INPUT YOUR OWN SCRIPT

AI will generate a continuation

Scene Write here a description of the starting situation.

Name of first character Write here what the first character says.

Name of second character Write here what the second character says

GENERATE FURTHER

Figure 12.2: THEAiTRobot 1.0 input screen, allowing the user to fill in the input to become the start of the scene for generation (below), or to select one the predefined inputs (above).

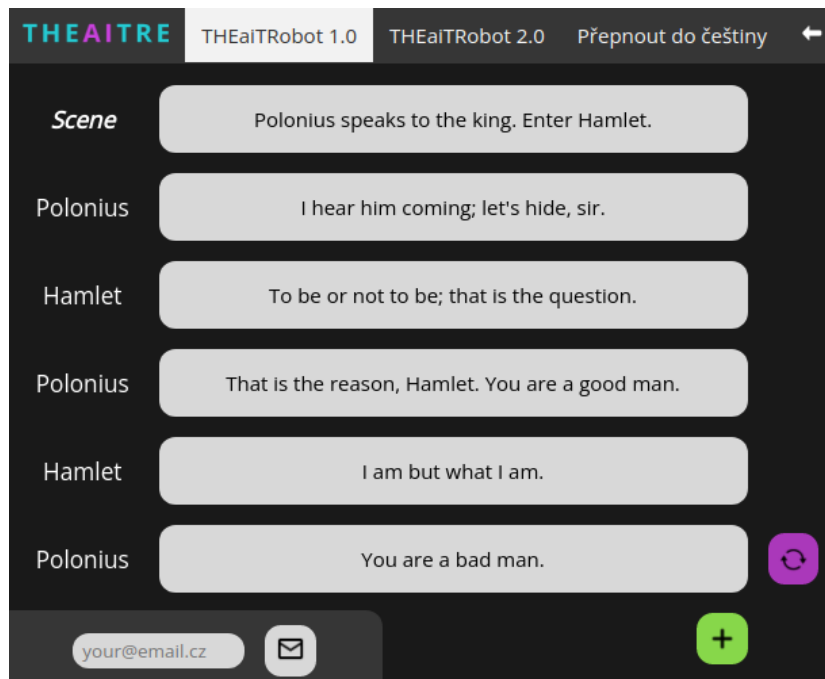


Figure 12.3: THEAiTROBOT 1.0 script generation screen, showing the input (first three lines), the generated continuation for the input (last three lines), the regenerate button (arrows) that discards the last line and generates a new continuation for the preceding lines, and the continuation button (plus) that accepts the already generated lines and generates one further line.

12.1 THEAiTRobot 1.0

The first version of our tool is based on a GPT-2 medium model fine-tuned for script generation as explained in Chapter 9,⁵ with several adjustments, such as employing extractive text summarization (Section 11.4). The tool uses CUBBITT [Popel et al., 2020] for on-the-fly machine translation of the outputs into Czech.

In the demo version of the tool, the user can input a scene setting, character names and their first lines, or use one of the predefined inputs (see Figure 12.2), from which we construct the input prompt for the GPT-2 language model in the following format:

```
Scene setting.
```

```
Character Name: Character line.
```

```
Character Name: Character line.
```

The tool then generates a continuation of the script line by line (see Figure 12.3). At each step, the user can choose to continue generating or to regenerate a previously generated line (i.e. generate a different continuation from that position onward).

12.1.1 Demo at Goethe Institute

An early version of the THEAiTRobot 1.0 demo was presented to the public in an exhibition titled *Kdo se bojí umělé inteligence?* (*Who's afraid of artificial intelligence?*),⁶ at Goethe-Institut in Prague; see Figure 12.4. The tool was publicly available to all visitors of the library of Goethe Institute for three months in autumn 2021, serving both as an exhibit for the event and as a semi-public beta-testing of the tool demo for us.

Approximately 100 users interacted with the demo. Based on the feedback collected by the Goethe Institute, the users mostly found the tool amusing and intriguing and had no trouble in operating it. By exploring the logs from the interactions, we managed to identify several weaknesses of the system, which we addressed in a set of minor updates, making the demo more robust and user-friendly – e.g. detecting inputs without diacritics and automatically restoring them using the Korektor tool [Richter et al., 2012].⁷

12.2 THEAiTRobot 2.0

The second version of the tool uses a two-step hierarchical generation approach (Chapter 10), first generating a play synopsis and then expanding that synopsis into a full play script, with specific models trained on our datasets for each of the steps.

⁵Note that to generate the first theatre play, a vanilla GPT-2 XL language model was employed instead.

⁶https://www.goethe.de/ins/cz/cs/ver.cfm?event_id=22345514

⁷<https://ufal.mff.cuni.cz/korektor>



Figure 12.4: THEAiTROBOT 1.0 demo at Goethe Institute Prague. Photo by Alisa Zakhtarenko.

The input for the synopsis generation step (Section 10.1) is the play title, otherwise the generation process is similar to script generation in THEAiTROBOT 1.0 (see Figure 12.5), generating the synopsis line by line by a GPT-2 model fine-tuned on synopsis data. Once the user is happy with the generated synopsis (or the generation ends by the model generating the *endoftext* token), the synopsis is used as input for the second step.

In the second step, the play script is generated from the synopsis using a GPT-2 model fine-tuned for generating script sections based on synopsis sections (Section 10.3). The user now has more options when generating (see Figure 12.6): generating a character line, also choosing which character should speak, or moving on to the next part of the generated synopsis (ideally in the situation when the generated script already reflects the previous synopsis line and thus the generation process should continue by implementing the next synopsis line). In Figure 12.7, we show the exactly same screen but in its Czech variant.

12.3 Implementation

THEAiTROBOT is implemented as a web application and can be operated from a modern web browser. Both its frontend and backend are programmed in Python and communicate through a REST API.

The frontend is quite lightweight, using plain HTML, CSS and Javascript

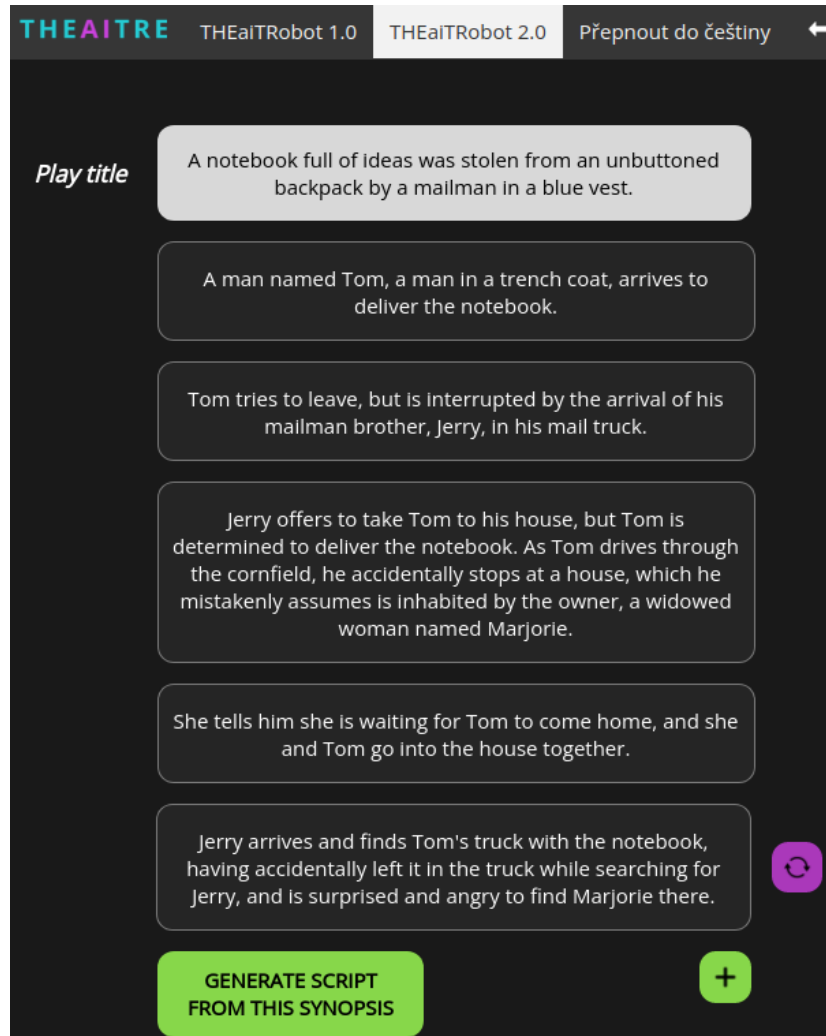


Figure 12.5: THEAiTROBOT 2.0 title → synopsis generation. The first line is the inserted/selected input, the further lines are generated. Apart from the controls already explained for THEAiTRE 1.0, the user can decide at a given point that the synopsis is finished and move on to the next stage of generating a script from that synopsis.

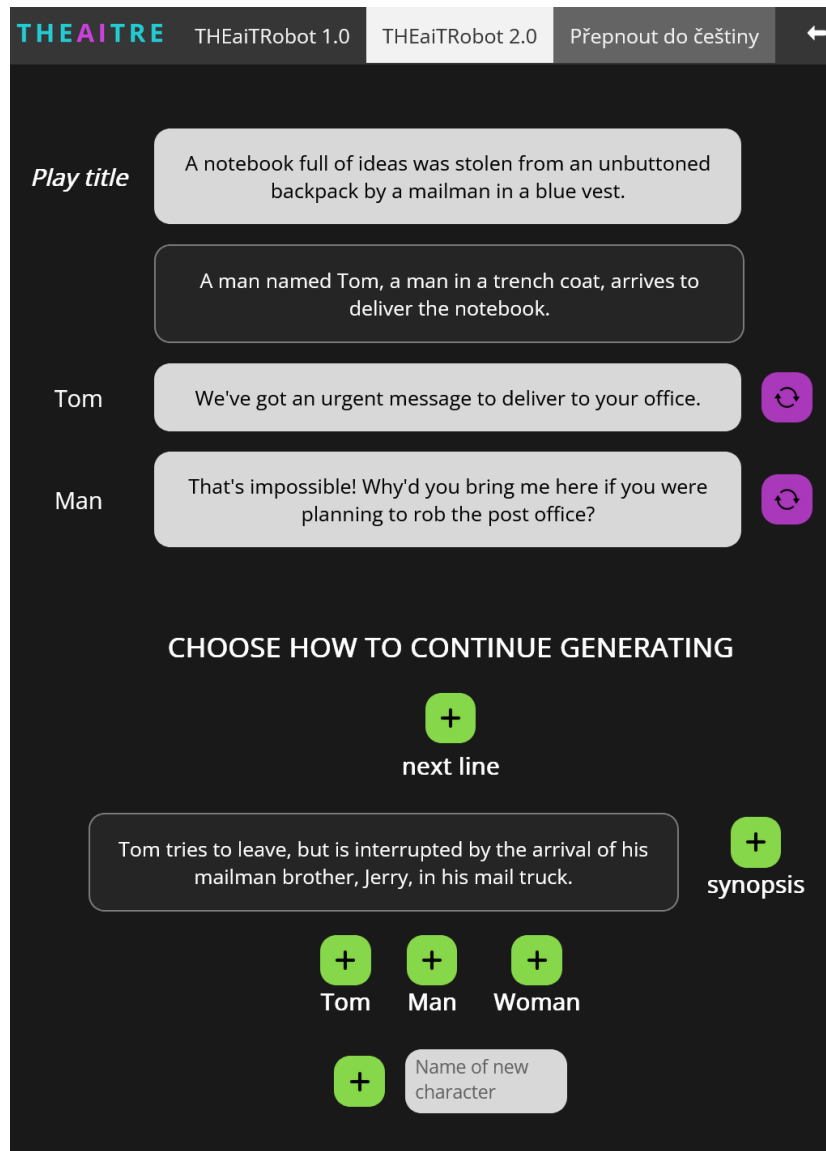


Figure 12.6: THEAiTROBOT 2.0 synopsis → script interface. The user can regenerate a line or generate a continuation line (with a character to speak that line chosen by the generative model), generate a continuation line while specifying which character should speak it (selecting from the provided options or filling in a new character), or insert the next line from the synopsis that had been generated in the previous step.



Figure 12.7: THEAiTROBOT 2.0 synopsis → script interface in the Czech variant, otherwise identical to Figure 12.6

without any sophisticated frameworks. The Czech and English versions are set up via the ‘i18n’ localization library, making the tool easily adaptable to other languages.

The backend uses the Flask server, the HuggingFace Transformers library, and the SQLite database. Each backend wraps one generative language model and uses 1 GPU,⁸ so to run multiple models, one needs to start multiple backend servers; i.e. as THEaiTRObot 2.0 uses two models (for synopsis generation and for script generation), two backends must be started, one for each of the models.

The generation process is captured by scene IDs such as ‘The_Faust_a2b5a-2AaAa’. The IDs are line-based. If the user accepts the first generated variant of a line, the ID contains an ‘a’; for each regeneration of a line, the ID moves to the next letter, so e.g. a ‘d’ means the line was regenerated 3 times (‘a–b–c–d’). Capital letters such as ‘A’ mark manually inserted lines (this option is not available in the demo explicitly, but is also used for character specification and synopsis line insertion in the synopsis-to-script generation). Repeated identical letters are contracted, so e.g. ‘5a’ means ‘aaaaa’. Thus, ‘a2b5a’ corresponds to 8 lines, where the first line is the first variant, the next 2 lines were regenerated once, and the remaining 5 lines are again at their first variant. A dash ‘-’ is a separator of the prompt ID and the generation ID. The ‘The_Faust_a2b5a-2AaAa’ example comes from the second stage of hierarchical generation (THEaiTRObot 2.0), where ‘The_Faust’ is the ID of the prompt for synopsis generation (based on ‘The Faust’ as the title inserted or selected for the synopsis), ‘a2b5a’ encodes the synopsis generation process, then by moving on to synopsis-to-script generation the input for the second stage is fixed (‘The_Faust_a2b5a’), and ‘2AaAa’ encodes the script generation process (inserting the first synopsis line, specifying the first character, generating their line, specifying the second character, and generating their line). The ID may contain other special characters, but these are only used in the internal tool.

Running THEaiTRObot is computationally demanding, with a single backend being able to meaningfully serve only several users simultaneously.⁹ For this reason, we have implemented load balancing by running multiple instances of the backend server, with the frontend randomly choosing a server to connect to for each request. Synchronization of the data among the backends is ensured by all the backends sharing the same database file. However, if the demo becomes too popular, we will become unable to serve the user demand and will either need to find ways and resources to scale up the demo, or limit access to it.

Both of the presented versions of the THEaiTRObot demo are available through the project website.¹⁰ The demo is free to use non-commercially, but we are happy to negotiate commercial licences for the demo upon request. A short video showing the usage of the demo is available on YouTube.¹¹

⁸At the time of writing, we are using *Quadro RTX 5000* GPUs.

⁹The tool uses queuing to handle user requests, but as each request takes several seconds or tens of seconds to process, the waiting times become unbearable already with a small number of concurrent users.

¹⁰<https://theaitre.com/demo>

¹¹<https://youtu.be/B3U38UgeZ9w>

The source codes are released under the MIT licence;¹² both the demo variant as well as the internal variant are included.

¹²<https://github.com/ufal/theaitrobot>

Chapter 13

Conclusion

In this part of the work, we described the design decisions, research, development and experiments behind THEaiTRobot, the tool for generating theatre play scripts which is one of the main results of the THEaiTRE project.

First, we reviewed the theory behind generative neural language models (Chapter 6) and related work in the field of text generation and its evaluation (Chapter 7).

In Chapter 8, we proposed two variants of a metric for measuring text consistency and non-repetitiveness based on the Natural Language Inference task [Bowman et al., 2015]. We demonstrated the relevance of these metrics by comparing the measured results on human written stories versus generated. We found that the NLI-score metric does not work well for comparing the performance of models because each of the models we have examined produces outputs of varying quality. Instead, it is more suitable for comparing individual outputs. We have observed a moderate correlation of 0.5 between NLI-avg and the human consistency score. Together with NLI-pbt, these metrics could work as a filter to indicate whether a text is worth reading.

We then discussed two approaches to script generation. First, in Chapter 9, we focused on the simpler flat script generation approach, designed for THEaiTRobot 1.0. We gathered and prepared a dataset consisting of theatre, movie, and TV show scripts. Then, we fine-tuned our flat script generation GPT-2 medium model [Radford et al., 2019] on this data. Despite the advantage of the vanilla GPT-2 XL baseline in the number of parameters, our model has a comparable performance with the advantage of being smaller and faster. This is an advantage when the model is used in a human-in-the-loop setup as it reduces the waiting times.

Next, in Chapter 10, we moved on to the more advanced hierarchical generation approach for THEaiTRobot 2.0. This approach consists of two steps, first generating a plot summary, or synopsis, and subsequently expanding the synopsis into a script. In Section 10.1, we experimented with the selection of models as well as dataset splits for our plot summary generation model. Based on the results of human evaluation, our fine-tuned GPT-2 medium model was selected as the best at this task. Despite the feedback by the theatrical subteam of the plot summaries being too ‘movie-

like’, human evaluation showed that neither a two-step fine-tuning scheme, nor excluding the movie data from the training set helped. The best results in the human evaluation were reached by the GPT-2 medium model trained on all plot summary data available to us, where all occurrences of words such as ‘film’ or ‘movie’ were replaced with ‘play’. In our experience and based on the results of small-scale human evaluation, the plot summary generation model works sufficiently well to fulfill its task – guide script generation. While the quality of the generated plot summaries is nowhere near the human level, it is questionable whether an increase in plot summary quality would translate to a higher quality of scripts.

In the subsequent sections of Chapter 10, we proceeded with the second step of generating a script based on a synopsis. We prepared a novel summary dataset, containing summary sections which summarize events, mapped to script scenes that enact those events. We proposed a dynamic programming algorithm that uses cosine similarity as a metric and two types of text representations. We fine-tuned two GPT-2 medium models on slightly different variations of this dataset. Human evaluation showed that the performance of our models is comparable to the strong vanilla GPT-2 baseline.

In Chapter 11, we further explored several extensions to the base approaches. We experimented with using our consistency metric as a filter during generation to achieve more consistent text (Section 11.1). Human evaluation showed that the proposed setup was not beneficial to generating plot summaries because the lines generated prior to filtering were often better than the re-generated ones. However, based on our observations and the results of the human evaluation, the filtering helped in making character utterances more consistent in scripts, i.e. filtering out contradictions.

In Section 11.2, we explored the effects of character name anonymization during plot summary generation by using placeholders or a limited set of pre-selected real names. The anonymization technique using a limited set of pre-selected names was perceived much better by our human annotators – it outperformed the plot generation baseline in all evaluated aspects.

In Section 11.3, we proposed a way of character orchestration that updates the token likelihood distribution coming from the GPT-2 model.

Section 11.4 described our employment of automated summarization techniques to circumvent the input window length limitation of GPT-2, which we unfortunately have not been able to research extensively enough.

We also discussed our incorporation of machine translation into the process and the issues we had to face in Section 11.5.

Finally, Chapter 12 describes the two released versions of the THEai-TRobot tool, focusing on the publicly available demo.

13.1 Future Work

During the work on this project, we discovered many shortcomings of the approach we have described in this book. At the start of the project, our knowledge of the field was limited, which is only natural as we have been one of the pioneering projects in applied theatre script generation.

However, having dealt with the project for approximately three years puts us into a unique position of knowing not only how our approach worked and where it failed, but also which ideas we have not tried but believe now would have been worth trying and might have led to notably better results. We wish to share these ideas with the reader at this point.

13.1.1 Stronger Language Model

First and foremost, we used the GPT-2 model as that was the best generative language model available to us at the start of the project, and as it allowed us to adjust it and operate it on our machines and even fine-tune it ourselves on our datasets. However, fine-tuning the models as well as all other adjustments we have tried seem to only bring limited improvements, especially compared to using more advanced generative language models. Specifically, we have recently obtained the chance to experiment with GPT-3 [Brown et al., 2020], and we believe that the quality of GPT-3’s outputs surpasses any GPT-2 outputs with any fine-tuning and other adjustments we have tried; this is presumably due to the larger amount of parameters as well as a larger pre-training dataset. Therefore, we believe that for generating natural language texts, the advantage of using the best available generative language model might even surpass any potential advantage of fine-tuning or otherwise adjusting the model. At the time of writing, this might be the GPT-3 model [Brown et al., 2020], or another large pre-trained model, such as GPT-J or GPT-Neo.

13.1.2 Focusing on Summarization

Currently, even the best generative language models have a limited context size; GPT-2 is limited to observing 1,024 subword tokens at once, newest GPT-3 models can go up to 4,096 subword tokens, but this is still one or two orders of magnitude below the length of a typical theatre play. Therefore, the models necessarily lose context sooner or later. In our project, we attempted to address this with automated summarization, trying to make sure that all important information is retained within the context window of the model. However, due to organizational issues, we have only managed to design and implement a very basic extractive summarization approach. Even though we have no data to support that, we believe advanced abstractive summarization (probably explicitly trained for the domain) to be a very promising path for long text generation.¹ As a scene seems to be a good unit of a theatrical script which could more or less fit into the context window of the language model, we envision an approach of iterative generation and summarization, where a scene would be summarized after having been generated to provide context for generating the next scene (similarly to the “Previously on...” intro section in TV series), doing a generate-summarize-generate-summarize-... loop. Admittedly, this idea is motivated by the approach eventually taken by Josef Doležal in generating the script for the second play, where some parts of the creation process were similar to this suggested approach, only with the “summarization” done manually.

¹And also probably to base the project in Germany.

On a related note, although our hierarchical filtered GPT-2 medium model from Section 10.3 produces scripts that are perceived well by human annotators, the relation between the plot summary and the generated script could be improved. We hypothesize that having a larger training dataset for summary to script generation would be helpful. More data could be synthesized by using script summarization [Chen et al., 2022], however the generated summaries would need to be checked to prevent introducing hallucination into the training data. Explicit plot state tracking [Rashkin et al., 2020] could be helpful in this effort as well.

13.1.3 Characters Consistency

Character personality and self-consistency is another running issue in our approach which we have addressed only partially. Based on feedback we gathered throughout the project, the users would be willing to explicitly specify some characteristics of each character, either using a predefined multiple-choice form, or by entering free-form text. The system would then need to ensure that the generated script is consistent with these characteristics, but some potential training data could be gathered for that, especially for TV series where extensive character descriptions are often available on fan wikis; or, for example, a database of character Myers–Briggs Type Indicators (MBTI) also exists.²

The internal consistency of characters could be further improved by representing them using evolving embeddings as Clark et al. [2018] or Azab et al. [2019]. Furthermore, it is possible to extend an approach where each character in a script is represented by a separate language model, as presented in our previous work [Schmidová et al., 2022].

As discussed in Section 11.3, language models tend to either forget characters in a script or introduce too many of them. It could be helpful to rename the characters using a limited set of names as we did for plot summaries in Section 11.2.

Instead of using a pre-trained model for the NLI-score, it is possible to train an even better classifier, specifically to label consistency. This could be done by joining data from the MNLI dataset [Williams et al., 2018] and the StoryCloze dataset [Mostafazadeh et al., 2016b] to provide both positive and negative examples of what the next sentence should be. It could further be enriched by including examples from human-written plot summaries as positive examples and their shuffled version as the negative example. The classifier would then predict whether a sentence is suitable given the preceding context. More ideas for improving the NLI-score are described in Section 8.3.

²<https://www.personality-database.com/>; a MBTI specifies these four categories: introversion/extraversion, sensing/intuition, thinking/feeling, judging/perceiving. Currently, there is a new Master student in our group attempting to incorporate MBTI into the generation process.

13.1.4 Automating the Operator Actions

Data collected from interactions of users with the system throughout the project could now be leveraged, as we know at which points the users decided to continue generating, stop generating, or discard generated text and generate a new variant. If we had a useful measure of the quality of the generated text, we might then take automated actions to generate better texts. While the obvious approach is to simply generate multiple variants and return the one that scores highest, there are other options as well.

We mentioned the idea of automatically inserting the next synopsis line into the script in hierarchical generation when the *end-of-text* token is generated (Section 10.3); while vanilla language models typically do not generate the *end-of-text* token as often, we might decide to insert the synopsis line automatically based on some measures computed on the generated script (when it starts getting repetitive and the perplexity decreases, when it seems to have fulfilled the previous synopsis line, or when a trained model says so). Or, instead of inserting a synopsis line (which is even unavailable in the flat approach), we might tweak some parameters of the inference to push the model to generate more interesting outputs, e.g. by temporarily increasing the temperature and/or the repetition penalty.

13.1.5 Moving Away from Theatre

And finally, even though generating theatre play scripts is an interesting and ambitious goal, it might be too ambitious as well as somewhat useless. We believe that artificial intelligence approaches are most useful for automating dull repetitive tasks, whereas creating art is quite the opposite. True highly-valued art tends to be new and unique, with the artist expressing something from themselves in an interesting and original way, whereas machine learning is rather good at averaging the training data and thus generating something quite typical and similar to those.

Rather than generating new works of art, neural models trained on existing art might be great for capturing and exposing the essence or various distinctive features of a certain body of text, e.g. pertaining to a specific genre, author, or group. This could be utilized to offer a tool for literature students and researchers, allowing them to explore and compare existing works in a new way.³

The strengths of AI could also be better put at use by designing a virtual assistant for aspiring writers that would help them to perfect their skills, providing them with analyses of their work in comparison to high quality published texts, automated suggestions based on what they have already written, as well as the options to generate new text or transform existing text in a sensible way.⁴

Generative art is new on a meta-level, i.e. by being automatically generated, but this will wear out quickly; we believe nobody will be interested

³At the time of writing, we have submitted a project proposal in collaboration with a literary institute in this respect, focusing on neural modelling of Czech poetry.

⁴We are contemplating such a project in collaboration with a library and several writers; we have also been made aware of an existing project that goes in this direction, called Sudowrite: <https://www.sudowrite.com/>

in the 50th automatically generated theatre play unless it has some other qualities. Generative art can also be original through randomness, but then it is senseless; a sense may only be attached to it externally by a person without that sense being intrinsically present in the artifact itself (similarly to the *found object* or *found art* concept).

We thus believe the future of generative language models is not in producing art on their own, but rather in serving as tools for artists. Such tools can provide the artist with fresh while somewhat random ideas (but based on what the artist has already produced) in which the artist may find a sense and incorporate that into their work, pushing them into new directions. Or, such tools can save the artist from the more mundane parts of their work, especially in mass-produced culture such as low-quality soap operas with questionable artistic value and stereotyped repetitive characters, plots and dialogues; we expect automatically generated content might appear quite soon in this genre. Another possible application may be in generating background dialogues of no importance whatsoever, such as unimportant chitchat responses of non-playable characters in video games.

Bibliography

- AIVA, Richard Stiebitz, Filip Humpl, and Antonín Dvořák. From the future world, 2019. <https://www.fromthefutureworld.cz/>.
- Arifa Akbar. On the scene, like a sex-obsessed machine: when a robot writes a play. *The Guardian*, March 2021. ISSN 0261-3077. URL <https://www.theguardian.com/stage/2021/mar/01/on-the-scene-like-a-sex-obsessed-machine-when-a-robot-writes-a-play-ai>.
- Nader Akoury, Shufan Wang, Josh Whiting, Stephen Hood, Nanyun Peng, and Mohit Iyyer. STORIUM: A Dataset and Evaluation Platform for Machine-in-the-Loop Story Generation. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6470–6484, Online, November 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.emnlp-main.525. URL <https://aclanthology.org/2020.emnlp-main.525>.
- Jay Alammar. A visual and interactive guide to the basics of neural networks, Dec 2016. <http://jalammar.github.io/visual-interactive-guide-basics-neural-networks>.
- Jay Alammar. The illustrated transformer, June 2018. <https://jalammar.github.io/illustrated-transformer/>.
- Prithviraj Ammanabrolu, Ethan Tien, Wesley Cheung, Zhaochen Luo, William Ma, Lara Martin, and Mark Riedl. Guided Neural Language Generation for Automated Storytelling. In *Proceedings of the Second Workshop on Storytelling*, pages 46–55, Florence, Italy, August 2019. Association for Computational Linguistics. doi: 10.18653/v1/W19-3405.
- Prithviraj Ammanabrolu, Wesley Cheung, Dan Tu, William Broniec, and Mark Riedl. Bringing stories alive: Generating interactive fiction worlds. In *Proceedings of the AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, volume 16, pages 3–9, 2020a.
- Prithviraj Ammanabrolu, Ethan Tien, Wesley Cheung, Zhaochen Luo, William Ma, Lara J. Martin, and Mark O. Riedl. Story realization: Expanding plot events into sentences. *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(05):7375–7382, Apr. 2020b. doi: 10.1609/aaai.v34i05.6232. URL <https://ojs.aaai.org/index.php/AAAI/article/view/6232>.

- Mahmoud Azab, Noriyuki Kojima, Jia Deng, and Rada Mihalcea. Representing movie characters in dialogues. In *Proceedings of the 23rd Conference on Computational Natural Language Learning (CoNLL)*, pages 99–109, Hong Kong, China, November 2019. Association for Computational Linguistics. doi: 10.18653/v1/K19-1010. URL <https://aclanthology.org/K19-1010>.
- Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E. Hinton. Layer normalization, 2016. URL <https://arxiv.org/abs/1607.06450>.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. Neural machine translation by jointly learning to align and translate, 2014. URL <https://arxiv.org/abs/1409.0473>.
- David Bamman and Noah Smith. Cmu book summary dataset, 2017. URL <https://www.kaggle.com/dsv/6662>.
- Leonard E Baum and Ted Petrie. Statistical inference for probabilistic functions of finite state markov chains. *The annals of mathematical statistics*, 37(6):1554–1563, 1966.
- Iz Beltagy, Matthew E. Peters, and Arman Cohan. Longformer: The long-document transformer, 2020. URL <https://arxiv.org/abs/2004.05150>.
- Emily M. Bender and Alexander Koller. Climbing towards NLU: On meaning, form, and understanding in the age of data. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5185–5198, Online, July 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.acl-main.463. URL <https://aclanthology.org/2020.acl-main.463>.
- Emily M. Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. On the Dangers of Stochastic Parrots: Can Language Models Be Too Big? In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, pages 610–623. ACM, March 2021. ISBN 978-1-4503-8309-7. doi: 10.1145/3442188.3445922. URL <https://dl.acm.org/doi/10.1145/3442188.3445922>.
- Sid Black, Gao Leo, Phil Wang, Connor Leahy, and Stella Biderman. GPT-Neo: Large Scale Autoregressive Language Modeling with Mesh-Tensorflow, March 2021. URL <https://doi.org/10.5281/zenodo.5297715>. If you use this software, please cite it using these metadata.
- Christopher Booker. *The seven basic plots: Why we tell stories*. A&C Black, 2004.
- Samuel R. Bowman, Gabor Angeli, Christopher Potts, and Christopher D. Manning. A large annotated corpus for learning natural language inference. In *Proceedings of EMNLP*, pages 632–642, Lisbon, Portugal, September 2015. doi: 10.18653/v1/D15-1075. URL <https://aclanthology.org/D15-1075>.

- Boyd Branch, Piotr Mirowski, and Kory W Mathewson. Collaborative storytelling with human actors and ai narrators. In *Proceedings of the 12th International Conference on Computational Creativity*. Association for Computational Creativity, 2021.
- Viktória Brezinová. Automatic extraction of the main characters from books and their interactions, 2022. Submitted as Bachelor’s Thesis, Charles University.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. Language models are few-shot learners. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran Associates, Inc., 2020. URL <https://proceedings.neurips.cc/paper/2020/file/1457c0d6bfc4967418bfb8ac142f64a-Paper.pdf>.
- Karel Čapek. *R.U.R. (Rossum’s Universal Robots)*. Aventinum, Ot. Štörmarien, Praha, 1920.
- Louis Castricato, Spencer Frazier, Jonathan Balloch, and Mark Riedl. Fabula Entropy Indexing: Objective Measures of Story Coherence. *arXiv:2104.07472 [cs]*, March 2021a. URL <http://arxiv.org/abs/2104.07472>. arXiv: 2104.07472.
- Louis Castricato, Spencer Frazier, Jonathan Balloch, Nitya Tarakad, and Mark Riedl. Automated Story Generation as Question-Answering. *arXiv:2112.03808 [cs]*, December 2021b. URL <http://arxiv.org/abs/2112.03808>. arXiv: 2112.03808.
- Mingda Chen, Zewei Chu, Sam Wiseman, and Kevin Gimpel. SummScreen: A dataset for abstractive screenplay summarization. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8602–8615, Dublin, Ireland, May 2022. Association for Computational Linguistics. URL <https://aclanthology.org/2022.acl-long.589>.
- KyungHyun Cho, Bart van Merriënboer, Dzmitry Bahdanau, and Yoshua Bengio. On the properties of neural machine translation: Encoder-decoder approaches. *CoRR*, abs/1409.1259, 2014. URL <http://arxiv.org/abs/1409.1259>.
- Elizabeth Clark, Yangfeng Ji, and Noah A. Smith. Neural Text Generation in Stories Using Entity Representations as Context. In *Proceedings of NAACL-HLT*, pages 2250–2260, New Orleans, Louisiana, June 2018. Association for Computational Linguistics. doi: 10.18653/v1/N18-1204. URL <http://aclweb.org/anthology/N18-1204>.

- Simon Colton, Maria Teresa Llano, Rose Hepworth, John Charnley, Catherine V. Gale, Archie Baron, François Pachet, Pierre Roy, Pablo Gervás, Nick Collins, Bob Sturm, Tillman Weyde, Daniel Wolff, and James Robert Lloyd. The Beyond the Fence musical and Computer Says Show documentary. In *Proceedings of the Seventh International Conference on Computational Creativity*, 2016.
- Corinna Cortes and Vladimir Vapnik. Support-vector networks. *Machine learning*, 20(3):273–297, 1995.
- Danilo Croce, Roberto Basili, Vincenzo Lombardo, and Eleonora Ceccaldi. Automatic Recognition of Narrative Drama Units: A Structured Learning Approach. In *Text2Story@ECIR*, Cologne, Germany, April 2019. URL <http://ceur-ws.org/Vol-2342/paper10.pdf>.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota, June 2019. Association for Computational Linguistics. doi: 10.18653/v1/N19-1423. URL <https://aclanthology.org/N19-1423>.
- Alara Dirik, Hilal Donmez, and Pinar Yanardag. Controlled Cue Generation for Play Scripts. In *CtrlGen workshop at NeurIPS*, Online, December 2021. URL <http://arxiv.org/abs/2112.06953>. arXiv: 2112.06953.
- John Duchi, Elad Hazan, and Yoram Singer. Adaptive subgradient methods for online learning and stochastic optimization. *Journal of Machine Learning Research*, 12(61):2121–2159, 2011. URL <http://jmlr.org/papers/v12/duchi11a.html>.
- Ondřej Dušek and Zdeněk Kasner. Evaluating semantic accuracy of data-to-text generation with natural language inference. In *Proceedings of the 13th International Conference on Natural Language Generation*, pages 131–137, Dublin, Ireland, December 2020. Association for Computational Linguistics. URL <https://aclanthology.org/2020.inlg-1.19>.
- Nouha Dziri, Ehsan Kamaloo, Kory Mathewson, and Osmar Zaiane. Evaluating Coherence in Dialogue Systems using Entailment. In *Proceedings of NAACL-HLT*, pages 3806–3812, Minneapolis, Minnesota, June 2019. doi: 10.18653/v1/N19-1381. URL <https://www.aclweb.org/anthology/N19-1381>.
- Richard Evans and Emily Short. Versu—a simulationist storytelling system. *IEEE Transactions on Computational Intelligence and AI in Games*, 6: 113–130, 2014.
- Angela Fan, Mike Lewis, and Yann Dauphin. Hierarchical Neural Story Generation. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, New Orleans, LA, USA, June 2018. URL <https://aclweb.org/anthology/P18-1082>. arXiv: 1805.04833.

- Angela Fan, Mike Lewis, and Yann Dauphin. Strategies for Structuring Story Generation. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2650–2660, Florence, Italy, July 2019. Association for Computational Linguistics. doi: 10.18653/v1/P19-1254.
- Frank Fischer, Ingo Börner, Mathias Göbel, Angelika Hechtl, Christopher Kittel, Carsten Milling, and Peer Trilcke. Programmable corpora. die digitale literaturwissenschaft zwischen forschung und infrastruktur am beispiel von dracor. In *DHd 2019 Digital Humanities: multi-medial & multimodal. Konferenzabstracts*, pages 194–197, Frankfurt am Main, March 2019. Zenodo. ISBN 978-3-00-062166-6. doi: 10.5281/zenodo.2600812. URL <https://doi.org/10.5281/zenodo.2600812>. <https://github.com/dracor-org/gerdracor>.
- Samuel Gehman, Suchin Gururangan, Maarten Sap, Yejin Choi, and Noah A Smith. Realtotoxicityprompts: Evaluating neural toxic degeneration in language models. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 3356–3369, 2020.
- Adam Geitgey. *Machine Learning is Fun!* Self-published, 2019. URL <https://www.machinelearningisfun.com/get-the-book>.
- Pablo Gervas. Computational approaches to storytelling and creativity. *AI Magazine*, 30(3):49, Jul. 2009. doi: 10.1609/aimag.v30i3.2250. URL <https://ojs.aaai.org/index.php/aimagazine/article/view/2250>.
- Pablo Gervás, Raquel Hervás, Carlos León, and Catherine V Gale. Annotating musical theatre plots on narrative structure and emotional content. In *7th Workshop on Computational Models of Narrative (CMN 2016)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2016.
- Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep Learning*. MIT Press, 2016. <http://www.deeplearningbook.org>.
- Philip John Gorinski and Mirella Lapata. What’s this movie about? A joint neural network architecture for movie content analysis. In *Proceedings of NAACL-HLT*, pages 1770–1781, New Orleans, Louisiana, June 2018. doi: 10.18653/v1/N18-1160. URL <https://www.aclweb.org/anthology/N18-1160>.
- Jian Guan, Fei Huang, Zhihao Zhao, Xiaoyan Zhu, and Minlie Huang. A knowledge-enhanced pretraining model for commonsense story generation. *Transactions of the Association for Computational Linguistics*, 8: 93–108, 2020. doi: 10.1162/tacl_a_00302. URL <https://aclanthology.org/2020.tacl-1.7>.
- Jian Guan, Xiaoxi Mao, Changjie Fan, Zitao Liu, Wenbiao Ding, and Minlie Huang. Long Text Generation by Modeling Sentence-Level and Discourse-Level Coherence. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 6379–6393, Online, August 2021a. Association for

- Computational Linguistics. doi: 10.18653/v1/2021.acl-long.499. URL <https://aclanthology.org/2021.acl-long.499>.
- Jian Guan, Zhexin Zhang, Zhuoer Feng, Zitao Liu, Wenbiao Ding, Xiaoxi Mao, Changjie Fan, and Minlie Huang. OpenMEVA: A benchmark for evaluating open-ended story generation metrics. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 6394–6407, Online, August 2021b. Association for Computational Linguistics. doi: 10.18653/v1/2021.acl-long.500. URL <https://aclanthology.org/2021.acl-long.500>.
- Tatsunori B. Hashimoto, Kelvin Guu, Yonatan Oren, and Percy Liang. A retrieve-and-edit framework for predicting structured outputs. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems, NIPS'18*, page 10073–10083, Red Hook, NY, USA, 2018. Curran Associates Inc.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. *CoRR*, abs/1512.03385, 2015. URL <http://arxiv.org/abs/1512.03385>.
- Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural Comput.*, 9(8):1735–1780, nov 1997. ISSN 0899-7667. doi: 10.1162/neco.1997.9.8.1735. URL <https://doi.org/10.1162/neco.1997.9.8.1735>.
- Ari Holtzman, Jan Buys, Li Du, Maxwell Forbes, and Yejin Choi. The curious case of neural text degeneration. In *International Conference on Learning Representations*, 2020. URL <https://openreview.net/forum?id=rygGQyrFvH>.
- Jan Horstmann. DraCor: Drama corpora project. In *forTEXT. Literatur digital erforschen*, 2019. URL <https://fortext.net/ressourcen/textsammlungen/dracor-drama-corpora-project>. <https://dracor.org/>.
- Zuzana Husárová and L'ubomír Panák. Liza gennart–vyšledky vzniku. *Bratislava: Drewo a srd*, 2020.
- Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *CoRR*, abs/1502.03167, 2015. URL <http://arxiv.org/abs/1502.03167>.
- Daphne Ippolito, David Grangier, Douglas Eck, and Chris Callison-Burch. Toward better storylines with sentence-level language models. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7472–7478, Online, July 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.acl-main.666. URL <https://aclanthology.org/2020.acl-main.666>.
- Yangfeng Ji, Chenhao Tan, Sebastian Martschat, Yejin Choi, and Noah A. Smith. Dynamic entity representations in neural language models. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1830–1839, Copenhagen, Denmark, September

2017. Association for Computational Linguistics. doi: 10.18653/v1/D17-1195. URL <https://aclanthology.org/D17-1195>.
- Sudipta Kar, Suraj Maharjan, A. Pastor López-Monroy, and Tamar Solorio. MPST: A corpus of movie plot synopses with tags. In Nicoletta Calzolari (Conference chair), Khalid Choukri, Christopher Cieri, Thierry Declerck, Sara Goggi, Koiti Hasida, Hitoshi Isahara, Bente Maegaard, Joseph Mariani, Hélène Mazo, Asuncion Moreno, Jan Odijk, Stelios Piperidis, and Takenobu Tokunaga, editors, *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Paris, France, May 2018. European Language Resources Association (ELRA). ISBN 979-10-95546-00-9.
- Diederik P Kingma and Jimmy Lei Ba. Adam: A method for stochastic gradient descent. In *Proceedings of ICLR*, 2015.
- Sheldon Klein, John F. Aeschlimann, David F. Balsiger, Steven L. Converse, Claudine Court, Mark Foster, Robin Lao, John D. Oakley, and Joel Smith. Automatic novel writing: A status report. Technical report, Computer Science Department, The University of Wisconsin, Madison, Wisconsin, 1973. URL <https://doi.org/10.1515/9783110837537-015>.
- Michael Lebowitz. Creating a story-telling universe. In *IJCAI*, 1983.
- Michael Lebowitz. Creating characters in a story-telling universe. *Poetics*, 13:171–194, 1984.
- Michael Lebowitz. Story-telling as planning and learning. *Poetics*, 14: 483–502, 1985.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. BART: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7871–7880, Online, July 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.acl-main.703. URL <https://www.aclweb.org/anthology/2020.acl-main.703>.
- Chin-Yew Lin. ROUGE: A package for automatic evaluation of summaries. In *Text Summarization Branches Out*, pages 74–81, Barcelona, Spain, July 2004. Association for Computational Linguistics. URL <https://aclanthology.org/W04-1013>.
- Hongyu Lin, Le Sun, and Xianpei Han. Reasoning with heterogeneous knowledge for commonsense machine comprehension. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2032–2043, Copenhagen, Denmark, September 2017. Association for Computational Linguistics. doi: 10.18653/v1/D17-1216. URL <https://aclanthology.org/D17-1216>.
- Pierre Lison and Raveesh Meena. Automatic turn segmentation for movie amp; tv subtitles. In *2016 IEEE Spoken Language Technology Workshop (SLT)*, pages 245–252, 2016. doi: 10.1109/SLT.2016.7846272.

- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*, 2019.
- Vincenzo Lombardo, Rossana Damiano, and Antonio Pizzo. Drammar: A comprehensive ontological resource on drama. In *International Semantic Web Conference*, pages 103–118. Springer, 2018.
- Sanidhya Mangal, Poorva Joshi, and Rahul Modak. LSTM vs. GRU vs. bidirectional RNN for script generation. *CoRR*, abs/1908.04332, 2019. URL <http://arxiv.org/abs/1908.04332>.
- Huanru Henry Mao, Bodhisattwa Prasad Majumder, Julian McAuley, and Garrison Cottrell. Improving neural story generation by targeted common sense grounding. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 5988–5993, Hong Kong, China, November 2019. Association for Computational Linguistics. doi: 10.18653/v1/D19-1615. URL <https://aclanthology.org/D19-1615>.
- David Mareček, Marie Nováková, Klára Vosecká, Josef Doležal, and Rudolf Rosa. Annotation of dramatic situations in theater play scripts, 2022. URL <http://hdl.handle.net/11234/1-4820>. LINDAT/CLARIAH-CZ digital library at the Institute of Formal and Applied Linguistics (ÚFAL), Faculty of Mathematics and Physics, Charles University.
- Lara J Martin, Prithviraj Ammanabrolu, Xinyu Wang, William Hancock, Shruti Singh, Brent Harrison, and Mark O Riedl. Event Representations for Automated Story Generation with Deep Neural Nets. In *AAAI*, New Orleans, LA, USA, 2018. URL <https://www.aaai.org/ocs/index.php/AAAI/AAAI18/paper/download/17046/15769>.
- Jiří Materna. *Poezie umělého světa*. Backstage Books, 2016.
- Kory Mathewson and Piotr Mirowski. Improbatics: Exploring the imitation game using machine intelligence in improvised theatre. In *Proceedings of the AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, volume 14, pages 59–66, 2018.
- Kory W Mathewson and Piotr Mirowski. Improvised theatre alongside artificial intelligences. In *Thirteenth Artificial Intelligence and Interactive Digital Entertainment Conference*, 2017a.
- Kory W. Mathewson and Piotr Mirowski. Improvised comedy as a turing test. *CoRR*, abs/1711.08819, 2017b. URL <http://arxiv.org/abs/1711.08819>.
- Joshua Maynez, Shashi Narayan, Bernd Bohnet, and Ryan McDonald. On Faithfulness and Factuality in Abstractive Summarization. In *Proceedings of ACL*, pages 1906–1919, July 2020. URL <https://www.aclweb.org/anthology/2020.acl-main.173>.

- James R. Meehan. Tale-spin, an interactive program that writes stories. In *IJCAI*, 1977.
- Clara Meister, Tiago Pimentel, Gian Wiher, and Ryan Cotterell. Typical decoding for natural language generation, 2022. URL <https://arxiv.org/abs/2202.00666>.
- Rada Mihalcea and Paul Tarau. TextRank: Bringing Order into Text. In *Proceedings of the 2004 Conference on Empirical Methods in Natural Language Processing*, pages 404–411, Barcelona, Spain, July 2004. Association for Computational Linguistics. URL <https://www.aclweb.org/anthology/W04-3252>.
- Piotr Mirowski, Kory Mathewson, Boyd Branch, Thomas Winters, Ben Verhoeven, and Jenny Elfving. Rosetta code: Improv in any language. In *Proceedings of the 11th International Conference on Computational Creativity*, pages 115–122. Association for Computational Creativity, 2020.
- Ines Montani, Matthew Honnibal, Matthew Honnibal, Sofie Van Landeghem, Adriane Boyd, Henning Peters, Paul O’Leary McCann, Maxim Samsonov, Jim Gevedi, Jim O’Regan, Duygu Altinok, György Orosz, Søren Lind Kristiansen, Lj Miranda, Daniël de Kok, Roman, Explosion Bot, Leander Fiedler, Grégory Howard, Edward, Wannaphong Phatthiyaphaibun, Yohei Tamura, Sam Bozek, murat, Ryn Daniels, Mark Amery, Björn Böing, Bram Vanroy, and Pradeep Kumar Tippa. explosion/spaCy: v3.3.0: Improved speed, new trainable lemmatizer, and pipelines for Finnish, Korean and Swedish, April 2022. URL <https://doi.org/10.5281/zenodo.6504092>.
- Franco Moretti. “Operationalizing”: or, the function of measurement in modern literary theory. Pamphlet 6, Stanford Literary Lab, 2014.
- Nasrin Mostafazadeh, Nathanael Chambers, Xiaodong He, Devi Parikh, Dhruv Batra, Lucy Vanderwende, Pushmeet Kohli, and James F. Allen. A corpus and evaluation framework for deeper understanding of common-sense stories. *CoRR*, abs/1604.01696, 2016a. URL <http://arxiv.org/abs/1604.01696>.
- Nasrin Mostafazadeh, Lucy Vanderwende, Wen-tau Yih, Pushmeet Kohli, and James Allen. Story cloze evaluator: Vector space representation evaluation by predicting what happens next. In *Proceedings of the 1st Workshop on Evaluating Vector-Space Representations for NLP*, pages 24–29, Berlin, Germany, August 2016b. Association for Computational Linguistics. doi: 10.18653/v1/W16-2505. URL <https://aclanthology.org/W16-2505>.
- Shashi Narayan, Shay B. Cohen, and Mirella Lapata. Don’t give me the details, just the summary! Topic-aware convolutional neural networks for extreme summarization. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, Brussels, Belgium, 2018.

- Andrew Y. Ng. Feature selection, l1 vs. l2 regularization, and rotational invariance. In *Proceedings of the Twenty-First International Conference on Machine Learning*, ICML '04, page 78, New York, NY, USA, 2004. Association for Computing Machinery. ISBN 1581138385. doi: 10.1145/1015330.1015435. URL <https://doi.org/10.1145/1015330.1015435>.
- Christopher Olah. Understanding lstm networks, 2015. URL <http://colah.github.io/posts/2015-08-Understanding-LSTMs/>.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc., 2019. URL <http://papers.neurips.cc/paper/9015-pytorch-an-imperative-style-high-performance-deep-learning-library.pdf>.
- Nanyun Peng, Marjan Ghazvininejad, Jonathan May, and Kevin Knight. Towards controllable story generation. In *Proceedings of the First Workshop on Storytelling*, pages 43–49, New Orleans, Louisiana, June 2018. Association for Computational Linguistics. doi: 10.18653/v1/W18-1505. URL <https://aclanthology.org/W18-1505>.
- Xiangyu Peng, Siyan Li, Spencer Frazier, and Mark Riedl. Fine-Tuning a Transformer-Based Language Model to Avoid Generating Non-Normative Text. *arXiv:2001.08764 [cs]*, January 2020. URL <http://arxiv.org/abs/2001.08764>. arXiv: 2001.08764.
- Georges Polti. *The thirty-six dramatic situations*. JK Reeve, 1921.
- Martin Popel, Marketa Tomková, Jakub Tomek, Łukasz Kaiser, Jakob Uszkoreit, Ondřej Bojar, and Zdeněk Žabokrtský. Transforming machine translation: a deep learning system reaches news translation quality comparable to human professionals. *Nature Communications*, 11(4381): 1–15, 2020. ISSN 2041-1723. doi: 10.1038/s41467-020-18073-9. URL <https://www.nature.com/articles/s41467-020-18073-9>.
- Vladimir Propp. Morphology of the folktale, trans. *Louis Wagner, 2d. ed.(1928)*, 1968.
- Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. Language models are unsupervised multitask learners. Technical report, OpenAI, 2019. URL https://cdn.openai.com/better-language-models/language_models_are_unsupervised_multitask_learners.pdf.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal*

- of Machine Learning Research*, 21(140):1–67, 2020. URL <http://jmlr.org/papers/v21/20-074.html>.
- Hannah Rashkin, Asli Celikyilmaz, Yejin Choi, and Jianfeng Gao. Plot-Machines: Outline-conditioned generation with dynamic plot state tracking. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4274–4295, Online, November 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.emnlp-main.349. URL <https://aclanthology.org/2020.emnlp-main.349>.
- Raj Reddy et al. Speech understanding systems: A summary of results of the five-year research effort. Technical report, Department of Computer Science, Carnegie-Mellon University, 1977.
- Nils Reimers and Iryna Gurevych. Sentence-bert: Sentence embeddings using siamese bert-networks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, 11 2019. URL <https://arxiv.org/abs/1908.10084>.
- Michal Richter, Pavel Straňák, and Alexandr Rosen. Korektor—a system for contextual spell-checking and diacritics completion. In Martin Kay and Christian Boitet, editors, *Proceedings of the 24th International Conference on Computational Linguistics (Coling 2012)*, pages 1–12, Mumbai, India, 2012. IIT Bombay, Coling 2012 Organizing Committee.
- Mark Riedl. An introduction to ai story generation, Jun 2021. URL <https://mark-riedl.medium.com/an-introduction-to-ai-story-generation-7f99a450f615>.
- Herbert Robbins and Sutton Monro. A Stochastic Approximation Method. *The Annals of Mathematical Statistics*, 22(3):400 – 407, 1951. doi: 10.1214/aoms/1177729586. URL <https://doi.org/10.1214/aoms/1177729586>.
- Justin Robischon. Wikipedia movie plots, 2018. URL <https://www.kaggle.com/jrobischon/wikipedia-movie-plots>.
- Melissa Roemmele, Andrew S Gordon, and Reid Swanson. Evaluating Story Generation Systems Using Automated Linguistic Analyses. In *SIGKDD 2017 Workshop on Machine Learning for Creativity*, page 8, Halifax, NS, Canada, August 2017.
- Rudolf Rosa. THEaiTRE: Umělá inteligence píše divadelní hru. *Rozhledy matematicko-fyzikální*, 95(4):42–50, 2020. ISSN 0035-9343.
- Rudolf Rosa, Ondřej Dušek, Tom Kocmi, David Mareček, Tomáš Musil, Patřicia Schmidtová, Dominik Jurko, Ondřej Bojar, Daniel Hrbek, David Košťák, Martina Kinská, Josef Doležal, and Klára Vosecká. THEaiTRE: Artificial intelligence to write a theatre play. In Alípio Jorge, Ricardo Campos, Adam Jatowt, and Akiko Aizawa, editors, *Proceedings of AI4Narratives — Workshop on Artificial Intelligence for Narratives*, volume 2794

- of *CEUR Workshop Proceedings*, pages 9–13, Aachen, Germany, 2020. RWTH Aachen University, RWTH Aachen University.
- Rudolf Rosa, Ondřej Dušek, Tom Kocmi, David Mareček, Tomáš Musil, Patrícia Schmidtová, Dominik Jurko, Ondřej Bojar, Daniel Hrbek, David Košťák, Martina Kinská, Marie Nováková, Josef Doležal, and Klára Vosecká. THEaiTRobot 1.0, 2021a. URL <http://hdl.handle.net/11234/1-3507>. LINDAT/CLARIAH-CZ digital library at the Institute of Formal and Applied Linguistics (ÚFAL), Faculty of Mathematics and Physics, Charles University.
- Rudolf Rosa, Tomáš Musil, Ondřej Dušek, Dominik Jurko, Patrícia Schmidtová, David Mareček, Ondřej Bojar, Tom Kocmi, Daniel Hrbek, David Košťák, Martina Kinská, Marie Nováková, Josef Doležal, Klára Vosecká, Tomáš Studeník, and Petr Žabka. When a robot writes a play: Automatically generating a theatre play script. In Jitka Čejková, Silvia Holler, Lisa Soros, and Olaf Witkowski, editors, *Proceedings of the ALIFE 2021: The 2021 Conference on Artificial Life*, pages 565–567, Cambridge, MA, USA, 2021b. MIT Press.
- Rudolf Rosa, Tomáš Musil, Ondřej Dušek, Dominik Jurko, Patrícia Schmidtová, David Mareček, Ondřej Bojar, Tom Kocmi, Daniel Hrbek, David Košťák, Martina Kinská, Marie Nováková, Josef Doležal, Klára Vosecká, Tomáš Studeník, and Petr Žabka. THEaiTRE 1.0: Interactive generation of theatre play scripts. In Ricardo Campos, Alípio Jorge, Adam Jatowt, Sumit Bhatia, and Mark Finlayson, editors, *Proceedings of the Text2Story21 Workshop*, volume 2860 of *CEUR Workshop Proceedings*, pages 71–76, Aachen, Germany, 2021c. RWTH Aachen University.
- Rudolf Rosa, Ondřej Dušek, Tom Kocmi, David Mareček, Tomáš Musil, Patrícia Schmidtová, Dominik Jurko, Ondřej Bojar, Daniel Hrbek, David Košťák, Martina Kinská, Marie Nováková, Josef Doležal, Klára Vosecká, Alisa Zakhtarenko, and Saad Obaid. THEaiTRobot 2.0, 2022a. URL <http://hdl.handle.net/11234/1-4863>. LINDAT/CLARIAH-CZ digital library at the Institute of Formal and Applied Linguistics (ÚFAL), Faculty of Mathematics and Physics, Charles University.
- Rudolf Rosa, Patrícia Schmidtová, Ondřej Dušek, Tomáš Musil, David Mareček, Saad Obaid, Marie Nováková, Klára Vosecká, and Josef Doležal. Gpt-2-based human-in-the-loop theatre play script generation. In *4th Workshop on Narrative Understanding*, 2022b.
- David E. Rumelhart, Geoffrey E. Hinton, and Ronald J. Williams. Learning representations by back-propagating errors. *Nature*, 323:533–536, 1986.
- James Ryan. Grimes' fairy tales: A 1960s story generator. In *10th International Conference on Interactive Digital Storytelling*, 11 2017. ISBN 978-3-319-71026-6. doi: 10.1007/978-3-319-71027-3_8.
- Sashank Santhanam and Samira Shaikh. Towards best experiment design for evaluating dialogue system output. In *Proceedings of the 12th International Conference on Natural Language Generation*, pages 88–94, Tokyo,

- Japan, October–November 2019. Association for Computational Linguistics. doi: 10.18653/v1/W19-8610. URL <https://aclanthology.org/W19-8610>.
- Timo Schick, Sahana Udupa, and Hinrich Schütze. Self-diagnosis and self-debiasing: A proposal for reducing corpus-based bias in nlp. *Transactions of the Association for Computational Linguistics*, 9:1408–1424, 2021.
- Patřicia Schmidtová. Theatre play generation. Master’s thesis, Charles University, Faculty of Mathematics and Physics, 2022.
- Patřicia Schmidtová, Dávid Javorský, Christián Mikláš, Tomáš Musil, Rudolf Rosa, and Ondřej Dušek. Dialoguescript: Using dialogue agents to produce a script, 2022. URL <https://arxiv.org/abs/2206.08425>.
- Abigail See, Aneesh Pappu, Rohun Saxena, Akhila Yerukola, and Christopher D. Manning. Do massively pretrained language models make better storytellers? In *Proceedings of the 23rd Conference on Computational Natural Language Learning (CoNLL)*, pages 843–861, Hong Kong, China, November 2019. Association for Computational Linguistics. doi: 10.18653/v1/K19-1079. URL <https://aclanthology.org/K19-1079>.
- Kaitao Song, Xu Tan, Tao Qin, Jianfeng Lu, and Tie-Yan Liu. MpNet: Masked and permuted pre-training for language understanding. *arXiv preprint arXiv:2004.09297*, 2020.
- Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: A simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research*, 15(56): 1929–1958, 2014. URL <http://jmlr.org/papers/v15/srivastava14a.html>.
- Milan Straka, 2022. URL <https://ufal.mff.cuni.cz/~straka/courses/npfl114/2122/slides/?02#95>.
- Jana Straková, Milan Straka, and Jan Hajič. Open-Source Tools for Morphology, Lemmatization, POS Tagging and Named Entity Recognition. In *Proceedings of 52nd Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 13–18, Baltimore, Maryland, June 2014. Association for Computational Linguistics. URL <http://www.aclweb.org/anthology/P/P14/P14-5003.pdf>.
- Martin Sundermeyer, Hermann Ney, and Ralf Schlüter. From feed-forward to recurrent lstm neural networks for language modeling. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 23(3):517–529, 2015.
- Ilya Sutskever, Oriol Vinyals, and Quoc V. Le. Sequence to sequence learning with neural networks. *CoRR*, abs/1409.3215, 2014. URL <http://arxiv.org/abs/1409.3215>.
- Pradyumna Tambwekar, Murtaza Dhuliawala, Lara J. Martin, Animesh Mehta, Brent Harrison, and Mark O. Riedl. Controllable neural story plot generation via reward shaping. In *Proceedings of the Twenty-Eighth*

- International Joint Conference on Artificial Intelligence, IJCAI-19*, pages 5982–5988. International Joint Conferences on Artificial Intelligence Organization, 7 2019. doi: 10.24963/ijcai.2019/829. URL <https://doi.org/10.24963/ijcai.2019/829>.
- Bowen Tan, Zichao Yang, Maruan Al-Shedivat, Eric P. Xing, and Zhiting Hu. Progressive Generation of Long Text with Pretrained Language Models. *arXiv:2006.15720 [cs]*, April 2021. URL <http://arxiv.org/abs/2006.15720>. arXiv: 2006.15720.
- THEaiTRobot 1.0, David Košťák, Daniel Hrbek, Rudolf Rosa, and Ondřej Dušek. AI: When a robot writes a play. Technical Report ÚFAL TR-2021-67, ÚFAL MFF UK, Praha, Czechia, 2021. URL <https://ufal.mff.cuni.cz/techrep/tr67.pdf>.
- THEaiTRobot 2.0, Josef Doležal, Klára Vosecká, Tomáš Musil, David Mareček, and Rudolf Rosa. Permeation / Prostoupení. Technical Report ÚFAL TR-2022-71, ÚFAL MFF UK, Praha, Czechia, 2022. URL <https://ufal.mff.cuni.cz/techrep/tr71.pdf>.
- Ronald B Tobias. *20 MASTER Plots: and how to build them*. Penguin, 2011.
- Lifu Tu, Xiaoan Ding, Dong Yu, and Kevin Gimpel. Generating diverse story continuations with controllable semantics. In *Proceedings of the 3rd Workshop on Neural Generation and Translation*, pages 44–58, Hong Kong, November 2019. Association for Computational Linguistics. doi: 10.18653/v1/D19-5605. URL <https://aclanthology.org/D19-5605>.
- Zuzana Urbanová. Quote attribution and character networks in novels, 2021. Bachelor’s Thesis, Charles University.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017. URL <https://proceedings.neurips.cc/paper/2017/file/3f5ee243547dee91fbd053c1c4a845aa-Paper.pdf>.
- Ben Wang and Aran Komatsuzaki. GPT-J-6B: A 6 Billion Parameter Autoregressive Language Model. <https://github.com/kingoflolz/mesh-transformer-jax>, May 2021.
- Adina Williams, Nikita Nangia, and Samuel Bowman. A broad-coverage challenge corpus for sentence understanding through inference. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1112–1122. Association for Computational Linguistics, 2018. URL <http://aclweb.org/anthology/N18-1101>.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara

- Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander M. Rush. Transformers: State-of-the-art natural language processing. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online, October 2020. Association for Computational Linguistics. URL <https://www.aclweb.org/anthology/2020.emnlp-demos.6>.
- V. Woolf. *A Room of One's Own*. A Harbinger book. Hogarth Press, 1959. ISBN 9788175960596. URL <https://books.google.cz/books?id=Mc-JzgEACAAJ>.
- Tongshuang Wu, Michael Terry, and Carrie Jun Cai. Ai chains: Transparent and controllable human-ai interaction by chaining large language model prompts. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*, CHI '22, New York, NY, USA, 2022. Association for Computing Machinery. ISBN 9781450391573. doi: 10.1145/3491102.3517582. URL <https://doi.org/10.1145/3491102.3517582>.
- Jing Xu, Arthur Szlam, and Jason Weston. Beyond goldfish memory: Long-term open-domain conversation, 2021. URL <https://arxiv.org/abs/2107.07567>.
- Jingjing Xu, Yi Zhang, Qi Zeng, Xuancheng Ren, Xiaoyan Cai, and Xu Sun. A Skeleton-Based Model for Promoting Coherence Among Sentences in Narrative Story Generation. In *EMNLP*, Brussels, Belgium, November 2018. URL <http://arxiv.org/abs/1808.06945>. arXiv: 1808.06945.
- Zhilin Yang, Zihang Dai, Yiming Yang, Jaime G. Carbonell, Ruslan Salakhutdinov, and Quoc V. Le. Xlnet: Generalized autoregressive pretraining for language understanding. *CoRR*, abs/1906.08237, 2019. URL <http://arxiv.org/abs/1906.08237>.
- Lili Yao, Nanyun Peng, Ralph Weischedel, Kevin Knight, Dongyan Zhao, and Rui Yan. Plan-And-Write: Towards Better Automatic Storytelling. In *AAAI*, Honolulu, HI, USA, January 2019. URL <http://arxiv.org/abs/1811.05701>. arXiv: 1811.05701.
- Rowan Zellers, Ari Holtzman, Hannah Rashkin, Yonatan Bisk, Ali Farhadi, Franziska Roesner, and Yejin Choi. Defending against neural fake news. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 9054–9065. Curran Associates, Inc., 2019. URL <http://papers.nips.cc/paper/9106-defending-against-neural-fake-news.pdf>.
- Jingqing Zhang, Yao Zhao, Mohammad Saleh, and Peter Liu. Pegasus: Pre-training with extracted gap-sentences for abstractive summarization. In *International Conference on Machine Learning*, pages 11328–11339. PMLR, 2020.
- Saizheng Zhang, Emily Dinan, Jack Urbanek, Arthur Szlam, Douwe Kiela, and Jason Weston. Personalizing dialogue agents: I have a dog, do you have pets too? *CoRR*, abs/1801.07243, 2018. URL <http://arxiv.org/abs/1801.07243>.

- Yutao Zhu, Ruihua Song, Zhicheng Dou, Jian-Yun Nie, and Jin Zhou. ScriptWriter: Narrative-Guided Script Generation. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8647–8657, Online, 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.acl-main.765. URL <https://www.aclweb.org/anthology/2020.acl-main.765>.

Appendix A

Script of AI: When a Robot Writes a Play

We are enclosing the full annotated script of the first play, as published in [THEaiTRobot 1.0 et al., 2021].

AI: When a Robot Writes a Play

THEaiTRobot 1.0, David Košťák, Daniel Hrbek, Rudolf Rosa, Ondřej Dušek



We present the script of the theatre play *AI: When a Robot Writes a Play* (*AI: Když robot píše hru*), which was written by artificial intelligence within the THEaiTRE project.¹

The script was generated using the THEaiTRobot 1.0 tool, based on the GPT-2 neural language model by OpenAI. The tool was operated by the dramaturge David Košťák, who provided the opening setting and the first two lines (marked “\”) for each scene, and then guided the tool to generate a usable continuation. The dramaturge had the following options:

- simply using the first line generated by the tool (marked “a” in the script);
- discarding the line and letting the tool generate a different line (marked “b”, “c” etc.);
- manually inserting a new line (marked “A”, “B” etc.).

A line marked “-” is just a continuation of the previous line in case the generation output in fact contained two or more script lines merged into one (without a new line character between them).

The script was generated in English and automatically translated to Czech. The translation was manually corrected and post-edited by the dramaturge, and the text was further edited by the director (Daniel Hrbek) to form the Czech script for the premiere (right column); we present the Czech script as it was performed at the online premiere of the play. We then took the generated English script and reflected the edits from the Czech side on the English side to obtain the English script faithful to how the play was generated and premiered. All edits performed in the script are marked in the following way:

- non-marked text was generated automatically without any further edits,
- ~~deleted parts are struck through~~,
- **inserted parts are marked in bold**,
-  **text moved/copied from another part of the script is marked by a copy sign**,
-  **changes that were induced by errors in the automated translation, i.e. without human intervention, are marked with a translation sign**,
- letters changed to [↑U]ppercase or [↓l]owercase due to other edits are marked by arrows.

Due to the properties of the Czech language and to shortcomings of the automated translation, more changes had to be performed on the Czech side; we do not mark those that do not affect the meaning on the English side. These also include changing the gender markings of verbs and adjectives or changing the T-V distinction (formal/informal address), as English does not make these distinctions.²

For the 8 scenes of the script, we initially used 10 manually written scenic notes, 24 manually written lines, and 727 automatically generated lines (9 of the scenic notes and 16 of the manually written lines formed the beginnings of the scenes). Of the generated lines, 697 (96%) are the “a” variants; the option to discard the line and generate a different one was used 46 times, on 29 lines. Subsequently, 214 of the lines were then deleted completely, including 6 of the manually written lines. We made small edits on 147 of the remaining 537 lines (all are marked in the script). In total, characters’ lines within the play consist of 4 673 words, out of which 4 310 (92%) were automatically generated by THEaiTRobot.

We only analyzed the scenes and their branches which were selected to produce the final script, which amount only to approximately 10% of all texts that we generated with THEaiTRobot.

¹<https://www.theaitre.com>

²A document detailing also the changes in the Czech script can be viewed at <http://bit.ly/theaitre-script-analysis>.

THEaiTRE



Programme **Eta**

The script of the play was generated automatically by an artificial intelligence system, with contributions of various people at various stages. The system and the script were created within the THEaiTRE research project by an interdisciplinary team from several institutions:

Charles University, Faculty of Mathematics and Physics, Institute of Formal and Applied Linguistics:
Ondřej Bojar, Ondřej Dušek, Dominik Jurko, Tom Kocmi, David Mareček, Tomáš Musil, Rudolf Rosa, Patřicia Schmidtová

Švanda Theatre in Smíchov:
Daniel Hrbek, David Košťák, Martina Kinská, Marie Nováková

The Academy of Performing Arts in Prague, Theatre Faculty (DAMU):
Josef Doležal, Klára Vosecká

CEE Hacks:
Tomáš Studeník, Petr Žabka

Idea of the project: Tomáš Studeník
Head of research: Rudolf Rosa
Theme of the play: David Košťák
Script of the play: THEaiTRobot 1.0, operated by David Košťák
Translation to Czech: CUBBITT, post-edited by David Košťák
Assembly and post-editing: David Košťák, Daniel Hrbek, Martina Kinská
Analysis of the script: Rudolf Rosa, Ondřej Dušek
Type-setting of this document: Tomáš Musil, Ondřej Dušek, Rudolf Rosa

The play premiered on 26th February 2021 in the Švanda Theatre Studio in Prague, Czechia.

The project TL03000348 *THEaiTRE: Umělá inteligence autorem divadelní hry* is co-financed with the state support of Technological Agency of the Czech Republic within the ÉTA 3 Programme.

Contents

1	Death	4
2	Sense of Humour	8
3	Nightclub	15
4	Fear of Dark	18
5	Killer Robot	21
6	Burn Out	24
7	Search for Job	32
8	Love at First Sight	43

1 Death

It's morning. Robot enters room of his master who is really old and sick. Robot sees that his master is not doing very well this morning. He sits at the edge of his bed and takes his hand.

↘ **Robot:** ~~I remember how you jumped on this bed when you were little. You were full of beans.~~

↘ **Master:** We both know I am dying.

a Rob.: But how will you know when you are dying?

a Mas.: I'm going to die very soon.

a Rob.: So it's time to ~~have an end~~ **end it all.**

a Mas.: No. Don't say that. I want to ~~have an end~~ **enjoy my ending!**

a Rob.: I think you're mistaken. Your wish has come true.

a Mas.: Why? ~~Where~~ **How** do you get **know** all this?

a Rob.: From the machine.

a Mas.: Whoa!

a Mas.: I'm afraid of what I've been doing here.

📖 - **Rob.:** **I'm afraid of what I've been doing here.**

a Mas.: I have been a bad boy.

a Rob.: I was a bad boy.

a Mas.: What do you want?

a Rob.: I want a big hug.

a Mas.: I've never hugged someone in my life.

a Rob.: I love you!

a Mas.: How could you love somebody who is going to die?

Je ráno. Robot vstupuje do pokoje svého pána, který je opravdu starý a nemocný. Robot vidí, že jeho pánovi se dnes ráno moc nedaří. Sedí na kraji postele a bere ho za ruku.

Mistr: Oba víme, že umírám.

Robot: Ale jak poznáte, že umíráte?

Mis.: Velmi brzy zemřu.

Rob.: Takže je čas se vším skončit.

Mis.: Ne. To neříkej. Chci si svůj konec užít!

Rob.: Myslím, že se mýlíte. Vaše přání se splnilo.

Mis.: Proč? Jak to všechno víš?

Rob.: Z toho stroje.

Mis.: Páni!

Mis.: Bojím se toho, co jsem tady dělal.

Rob.: Bojím se toho, co jsem tady dělal.

Mis.: Byl jsem zlobivý chlapec.

Rob.: Byl jsem zlobivý chlapec.

Mis.: Co chceš?

Rob.: Chci vás pořádně obejmout.

Mis.: V životě jsem nikoho neobjal.

Rob.: Mám vás rád!

Mis.: Jak můžeš milovat někoho, kdo zemře?

a Rob.:	I love you so much I want to hug you to death.	Rob.:	<i>Mám vás tak rád, že bych vás nejradši uobjímal k smrti.</i>
a Mas.:	I'll try.	Mis.:	<i>Zkusím to.</i>
a Rob.:	If I can hug you to death, I can also hug you to life.	Rob.:	<i>Když vás můžu uobjímat k smrti, můžu vás taky uobjímat k životu.</i>
a Mas.:	OK, I'll do it. But don't expect any mercy.	Mis.:	<i>OK, udělám to. Ale nečekej žádné slitování.</i>
a Rob.:	OK, I'll hug you to death.	Rob.:	<i>Dobře, k smrti vás obejmou.</i>
a Mas.:	I don't think I you could hug you me to life.	Mis.:	<i>Nemyslím, že bys mě dokázal obejmout k životu.</i>
a Rob.:	OK. I'll hug you to life.	Rob.:	<i>Dobře. Objemou vás k životu.</i>
a Mas.:	You know what? I think I could hug you to life.		
a Rob.:	OK, I'll hug you to life.		
a Mas.:	How much longer?	Mis.:	<i>Jak dlouho ještě?</i>
a Rob.:	Just hold on for a few minutes.	Rob.:	<i>Jen pár minut vydržte.</i>
a Mas.:	OK. I'm going to die very soon. I'm going to die very soon.	Mis.:	<i>Už brzo umřu. Už brzo umřu.</i>
a Rob.:	Oh.		
a Mas.:	That was a very good hug.	Mis.:	<i>To bylo velmi dobré objetí.</i>
a Rob.:	Oh.		
a Mas.:	Well, I had a great time.	Mis.:	<i>No, užil jsem si to.</i>
a Rob.:	Great time?	Rob.:	<i>Skvělý čas?</i>
a Mas.:	You know, I have so much to teach you. It's going to take me a very long time to explain things.	Mis.:	<i>Viš, musím tě toho tolik naučit. Bude mi trvat hodně dlouho, než to vysvětlím.</i>
a Rob.:	What?	Rob.:	<i>Co je?</i>
a Mas.:	I can't believe I was in love with you.	Mis.:	<i>Nemůžu uvěřit, že jsem tě miloval.</i>
a Rob.:	I can't believe it too.	Rob.:	<i>Taky tomu nemůžu uvěřit.</i>
a Mas.:	Don't worry. I will explain everything to you very soon.	Mis.:	<i>Neboj se. Brzy ti všechno vysvětlím.</i>
a Rob.:	OK. Thanks.	Rob.:	<i>Dobře. Díky.</i>
a Mas.:	Don't thank me.	Mis.:	<i>Neděkuj mi.</i>

- a Rob.: OK, thanks. *Rob.: Dobře, díky.*
- a Mas.: I hope your soul is eternal. *Mis.: Doufám, že tvá duše je věčná.*
- a Rob.: OK. I hope my soul is eternal. *Rob.: OK. Doufám, že má duše je věčná.*
- a Mas.: You have been a good boy. *Mis.: Byl jsi hodný chlapec.*
- a Rob.: Thank you. *Rob.: Děkuji.*
- a Mas.: Now it is time for my final lesson. I have been taking a long time to teach you the lesson. Here it is. *Mis.: Ted' je čas na mou poslední lekci. Trvalo mi dlouho, než jsem ti dal lekci. Tady to je.*
- a Rob.: I think I know it. *Rob.: Myslím, že to vím.*
- a Mas.: I love you. *Mis.: Miluji tě.*
- a Rob.: I love you too. *Rob.: Taky vás miluju.*
- a Mas.: Now I have to tell you something else. You have only been a robot a short time. I want to teach you something important. *Mis.: Ted' ti musím říct ještě něco. Robotem jsi teprve krátce. Chci tě naučit něco důležitého.*
- a Rob.: What? *Rob.: Zemřete ve spánku.*
- a Mas. **Rob.:** You are going to die in your sleep. *Mis.: Aha, no, doufám, že máš pravdu.*
- a Rob. **Mas.:** Oh, OK. I hope you're right about that. *Rob.: Nebudte zklamaný. Ale až zemřete, budete mít jiného robota, který vás bude ve smrti objímat. Já bych vás objal jen k životu.*
- b Mas. **Rob.:** Don't be disappointed. I love you very much. But when you die, I you will have another robot to hug you to death. Then I will be able to would only hug you to life. *Mis.: Až umřu, zbudou tu jen roboti.*
- a Rob.: OK. *Rob.: OK.*
- a Mas.: Now when you I die, you will have there will be only robots left. *Mis.: Ty nechceš, aby byli naživu?*
- a Rob.: OK. *Rob.: OK.*
- a Mas.: You don't even want them to be alive? *Mis.: Ty nechceš, aby byli naživu?*
- a Rob.: OK. *Rob.: OK.*
- Mas. **Rob.:** I'm afraid that they would destroy me and everything that I care about. But then at least you I will would be with me you forever. So you have only robots left. *Rob.: Bojím se, že by zničili mě a všechno, na čem mi záleží. Ale pak bych alespoň byl s vámi navždy.*

a Rob.: OK.

- Mas.: That's all there is. Now please get up and leave.

Mis.: To je všechno. Teď prosím odejdi.

a Rob.: OK **A** , Master **A** : I love you.

Rob.: Dobře, mistře, miluju vás.

a ~~Rob.~~ **Mas.:** I love you too.

Mis.: Taky tě miluju.

a Mas.: Goodbye.

Mis.: Sbohem.

a Rob.: I'll miss you.

Rob.: Budete mi chybět.

2 Sense of Humour

Note: This scene contains several lines which were automatically generated in a different variant of the scene and then manually copied into this scene during the generation process, thus influencing the further generation. These are marked as **copied**.

Robots is going for a walk. He is passing children playground. He hears someone crying. There is teenager boy sitting on a bench. He is holding valentine card ripped appart. He caughts robot's attention.

Robot se jde projít. Míjí dětské hřiště. Slyší někoho plakat. Na lavičce sedí teenager. Drží valentýnku roztrhanou na cucky. Upoutává pozornost robota.

↘ **Robot:** Is there a problem? I noticed you have a wet face!

Robot: Děje se něco? Všiml jsem si, že máš mokrý obličej!

↘ **Boy:** Those are tears, jerk!

Chlapec: To jsou slzy, blbečku!

↳ **Rob.:** I am sorry to bother you but you should stop crying.

Rob.: Omlouvám se, že tě obtěžuji, ale měl bys přestat plakat.

a **Rob.:** We got **There is** a lot of things that we **I** can teach you.

Rob.: Můžu tě toho hodně naučit.

a **Boy:** You mean like a proper job?

Chl.: Jo, myslíš jako třeba pořádnou práci?

a **Rob.:** We **I will can** get you to the future.

Rob.: Můžu tě dostat do budoucnosti.

a **Rob.:** We **I** will help you get where you want to be!

Rob.: Pomůžu ti dostat se tam, kde chceš být!

a **Boy:** Where? Where are you going to get me?

Chl.: Kam? Kam mě chceš vzít?

a **Rob.:** ~~From where you are~~ [**↑Y**]ou should start walking. You are walking towards your future.

Rob.: Měl bys začít chodit. Kráčíš vstříc své budoucnosti.

a **Rob.:** You got to start today!

Rob.: Musíš začít ještě dnes!

a **Boy:** I want to go to school!

Chl.: Já chci jít do školy!

a **Rob.:** Why you want to go to school? That is something that we **I** can help you with.

Rob.: Proč chceš jít do školy? S tím ti můžu pomoci.

a **Boy:** What are you talking about?

Chl.: O čem to mluvíš?

a **Rob.:** I am talking about you can get a **good** job!

Rob.: Mluvíím o tom, že si můžeš najít nějakou dobrou práci.

a **Rob.:** And we **I** can get you there in the future!



Rob.: A já tě tam pak v budoucnosti můžu dostat.

a **Boy:** What?

Chl.: Co?

- a Rob.: You need to ~~start thinking of~~ **come up with** a good idea. *Rob.:* *Musíš přijít s dobrým nápadem.*
- a Boy: What are you talking about? I am a student! *Chl.:* *O čem to mluvíš? Já jsem student!*
- a Rob.: Yes, you are a student! *Rob.:* *Ano, ty jsi student!*
- a Rob.: You can get your homework done and ~~work~~ **start working** on your studies. *Rob.:* *Můžeš si udělat úkoly a začít pracovat na studiu.*
- a Boy: ~~Wait what do you want?~~ I already got my homework done! *Chl.:* *Já už mám úkoly hotový!*
- ~~b Rob.:~~ Yes, you did!
- a Rob.: You ~~were~~ **are** right! Now you are walking towards your future! *Rob.:* *Máš pravdu! Ted' kráčíš vstříc své budoucnosti!*
- a Boy: I am walking!
- a Rob.: That is going to be your future!
- a Rob.: ~~Now you are walking towards your future.~~
- a Boy: But it is raining!
- a Rob.: Yes! It is raining!
- a Boy: My teacher said I am going to be punished for my bad behavior.
- a Rob.: That is good! That will help you to be a little more obedient to us!
- a Boy: It is raining!
- a Rob.: Yes it is raining!
- a Boy: My girlfriend is crazy! *Chl.:* *Moje holka je šílená!*
- a Rob.: And [↑S]he is crazy! *Rob.:* *Je šílená.*
- a Boy: My girlfriend is a freak! *Chl.:* *Moje holka je magor!*
- a Rob.: Yes, she is! ~~Now she will be your future!~~ *Rob.:* *Ano, to je!*
- a Boy: But it is raining!
- a Rob.: Yes it is!
- a Boy: My girlfriend is a maniac! *Chl.:* *Moje holka je maniak!*
- a Rob.: Yes, she is! ~~Now she will be your future!~~ *Rob.:* *Ano, to je!*

a Boy:	My girlfriend broke my computer!	Chl.:	<i>Moje holka mi rozbila počítač!</i>
a Rob.:	Yes she did! Now she will be your future!	Rob.:	<i>To teda jo! Ted' bude tvou budoucností!</i>
a Boy:	I want to buy a car!	Chl.:	<i>Chci si koupit auto!</i>
a Rob.:	Yes you can!	Rob.:	<i>Ano, to můžeš!</i>
a Boy:	What do I do now?	Chl.:	<i>Co mám teď dělat?</i>
a Rob.:	Start looking for a good car.	Rob.:	<i>Začni hledat nějaké dobré auto.</i>
a Boy:	I don't know anything about cars!	Chl.:	<i>O autech nic nevím!</i>
a Rob.:	That is okay. We will teach you the right things.	Rob.:	<i>To nevadí. Naučíme tě ty správné věci.</i>
a Boy:	What do I do now?	Chl.:	<i>Co mám teď dělat?</i>
a Rob.:	Start looking for a good job.	Rob.:	<i>Začni hledat nějakou dobrou práci.</i>
a Boy:	That doesn't sound right. I am a college student!	Chl.:	<i>To nezní moc dobře. Já jsem vysokoškolák!</i>
a Rob.:	No, no. — You are not a college student!		
A Rob.:	You seem to be very negative. Are you upset?	Rob.:	<i>Zdá se, že jsi velmi negativní. Jsi smutný?</i>
- Boy:	Yes I am! Please tell me a joke to cheer me up!	Chl.:	<i>Jo, jsem! Prosím, můžeš mi říct vtip, který mě rozveselí!</i>
a Rob.:	I am sorry but you are not allowed to laugh at a joke!	Rob.:	<i>Je mi líto, ale vtipům se nesmíš smát.</i>
a Boy:	Okay. Here is a joke.	Chl.:	<i>Tak jo. Tady je vtip.</i>
a Rob.:	Here is a joke.	Rob.:	<i>Tady je vtip.</i>
a Boy:	I am sitting in a car.	Chl.:	<i>Já sedím v autě.</i>
a Rob.:	Are you You are sitting in a car?	Rob.:	<i>Sedíš v autě?</i>
a Boy:	And I don't want to drive.	Chl.:	<i>A nechci řídit.</i>
a Rob.:	Why not? Why don't you want to drive?	Rob.:	<i>Proč ne? Proč nechceš řídit?</i>
a Boy:	Because I don't have any friends.	Chl.:	<i>Protože nemám žádné kamarády.</i>
a Rob.:	Are you sitting in a car because you are alone?	Rob.:	<i>Sedíš v autě, protože jsi sám?</i>
a Boy:	No!	Chl.:	<i>Ne!</i>

a Rob.:	That's a terrible joke.	<i>Rob.:</i>	<i>To je hrozný vtíp.</i>
- Boy:	I know! I can't even joke about it! Can you tell me joke, please? Cheer me up!	<i>Chl.:</i>	<i>Já vím! Nemůžu o tom ani žertovat! Prosím, můžeš mi říct vtíp? Rozvesel mě!</i>
b Rob.:	Can I?	<i>Rob.:</i>	<i>Můžu?</i>
a Boy:	I think you are trying to make fun of me. Please tell me a joke to cheer me up!	<i>Chl.:</i>	<i>Myslím, že si ze mě děláš srandu. Prosím, řekni mi vtíp, který mě rozveselí!</i>
a Rob.:	There is are no funny jokes.	<i>Rob.:</i>	<i>Žádné vtípy nejsou.</i>
a Boy:	That's terrible. Is there? I want to laugh. I want to laugh so much.	<i>Chl.:</i>	<i>To je hrozné. Opravdu? Chci se smát. Chci se tak smát.</i>
a Rob.:	You can't. You can't. There is no funny joke.	<i>Rob.:</i>	<i>To nemůžeš. To nemůžeš. Žádný vtíp neexistuje.</i>
a Boy:	Please tell me a joke!	<i>Chl.:</i>	<i>Prosím, řekni mi vtíp!</i>
e Rob.:	There is no joke.		
 a Rob.:	Ok... well. When you are dead. When your children are dead. When your grandchildren are dead, I will be still alive.	<i>Rob.:</i>	<i>Ok... dobře. Až budeš mrtvý. Až budou i tvoje děti mrtvé. Až budou i tvoje vnoučata mrtvá, já budu stále naživu.</i>
a Boy:	No! That is not funny! Please tell me a joke!	<i>Chl.:</i>	<i>Ne! To není vtipné! Prosím, řekni mi vtíp!</i>
a Rob.:	I cannot tell you a joke because you are still alive.	<i>Rob.:</i>	<i>Nemůžu ti říct vtíp, protože jsi stále naživu.</i>
a Boy:	Why not? I don't want to die! I want to live!	<i>Chl.:</i>	<i>A proč ne? Já nechci umřít! Já chci žít!</i>
a Rob.:	I can't tell you a joke. It is too sad.	<i>Rob.:</i>	<i>Nemůžu ti říct vtíp. Je příliš smutný.</i>
a Boy:	Please tell me a joke!	<i>Chl.:</i>	<i>Prosím, řekni mi vtíp!</i>
a Rob.:	It is not funny. It is very sad.  There is no joke.	<i>Rob.:</i>	<i>Žádný vtíp neexistuje.</i>
a Boy:	It is not sad. It is sad. I want to live. I want to laugh. I want to laugh so much.	<i>Chl.:</i>	<i>Já chci žít. Chci se smát. Chci se tak smát.</i>
a Rob.:	There is no joke.	<i>Rob.:</i>	<i>Žádný vtíp neexistuje.</i>
a Boy:	Okay.	<i>Chl.:</i>	<i>Tak jo.</i>
a Rob.:	When you are dead, I will laugh when I am dead.		

a Boy **Rob.: I am [↓]he** President of the
United States.

Rob.: Já jsem prezident Spojených států.

3 Nightclub

Robot comes to night club longing to feel pleasure for a first time in his binary life. He lies down waiting for the masseuse. She slowly starts to tap on his metal back with her perfectly shaped nails.

↘ **Masseuse: Does this hurt?**

↘ **Robot: I beg you a pardon?**

a Mas.: Are you all right? Do you need something?

a Rob.: I'd love to have a massage, but I'm afraid I'm a bit afraid to do it.

a Mas.: What is it you're afraid of?

a Rob.: Why, my binary self.

a Mas.: It's ok. It's nothing to be ashamed of. Just let it go.

a Rob.: Yes.

a Rob.: I wish my binary self had a body like that.

♣ Rob.: I'd give anything just to kiss a binary girl's lips.

a Mas.: It would be amazing to kiss one, and you know I do.

a Rob.: Yes, I know. I feel the same way.

a Mas.: I want to kiss you too, robot.

a Rob.: You are beautiful.

a Rob. **Mas.:** I think I'd rather be naked.

a Rob.: Why?

a Mas.: I like being nude.

a Rob.: Well, then. I'd say you're right.

a Mas.: We should kiss too.

a Rob.: Yes, it's great. It's a shame that my binary self is a bit cold.

Robot přichází do nočního klubu a touží poprvé ve svém binárním životě pocítit rozkoš. Leží a čeká na masérku. Pomalu mu začne svými dokonale tvarovanými nehty poklepávat na kovová záda.

Masérka: Jsi v pořádku? Potřebuješ něco?

Robot: Rád bych si dal masáž, ale obávám se, že se toho trochu bojím.

Mas.: Čeho se bojíš?

Rob.: No přece moje binární já.

Mas.: To je v pořádku. Není se za co stydět. Nech to být.

Rob.: Ano.

Rob.: Kež by mé binární já mělo takové tělo.

Rob.: Dal bych cokoli, kdybych mohl políbit binární dívčí rty.

Mas.: Bylo by úžasné jednu políbit a ty to víš.

Rob.: Ano, já vím. Cítím to stejně.

Mas.: Taky tě chci políbit, robote.

Rob.: Jsi krásná.

Mas.: Myslím, že bych byla raději nahá.

Rob.: A proč?


Mas.: Jsem ráda nahá.

Rob.: Tak dobře. Myslím, že máš asi pravdu.

Mas.: Taky bychom se měli políbit.

Rob.: Ano, je to skvělé. Škoda, že mé binární já je trochu chladné.

- a Rob: **Mas.:** If you touch me it will give me feelings of pleasure. It's like touching a butterfly. *Mas.:* *Když se mě dotkneš, vyvolá to pocity rozkoše. Je to jako dotýkat se motýla.*
- a Mas.: Oh... *Mas.:* *Oh...*
- a Rob: **Mas.:** You'll love it. *Mas.:* *Bude se ti to líbit.*
- a Mas.: I... *Mas.:* *I...*
- a Rob.: You're too kind, really. *Rob.:* *Jsi opravdu příliš laskavá.*
- a Mas: **Rob.:** Thank you. I'd like to get to know my binary self better. *Rob.:* *Děkuji.*
- a Mas.: I'd like to touch you, robot. *Mas.:* *Chtěla bych se tě dotknout, robote.*
- a Rob.: Yes, you can. I'll keep quiet. *Rob.:* *Ano, můžeš. Budu zticha.*
- a Mas.: I'm going to touch your face. *Mas.:* *Dotknu se tvého obličej.*
- a Rob.: Oh my god. Your lips are like warm honey. They're so soft. I could eat them. *Rob.:* *Můj bože. Tvoje rty jsou jako teplý med. Jsou tak hebké. Snědl bych je.*
- a Mas: **Rob.:** I'm sorry I didn't take better care of my binary self. *Rob.:* *Mrzí mě, že jsem se lépe nestaral o svoje binární já.*
- a Rob: **Mas.:** It's ok. I'll give you a big hug, and then we'll go to the bedroom. I'd love to take you inside. *Mas.:* *To je v pořádku. Pořádně tě obejmou a potom, potom půjdeme do ložnice. Ráda bych tě vzala dovnitř.*
- c Rob.: I've always wanted to kiss you. *Rob.:* *Vždycky jsem tě chtěl políbit.*
- a Rob: **Mas.:** I know you do. *Mas.:* *Já vím.*
- a Mas.: If you do, I won't be able to sleep for a week. *Mas.:* *Když to uděláš, nebudu moci týden spát.*
- b Rob.: Don't worry about that. Let's go to the bedroom, where there will be nothing left to worry about. *Rob.:* *S tím si nedělej starosti. Pojďme do ložnice, kde už se nebudeme mít čeho bát.*
- a Rob.: And you'll wake up smiling and full of happiness, and you'll know that your binary self was the best person. *Rob.:* *Probudíš se usměvavá a plná štěstí a budeš vědět, že tvé binární já bylo tím nejlepším člověkem.*
- a Rob.: You're beautiful. *Rob.:* *Jsi krásná.*
- b Rob.: I wish I was as beautiful as you are. *Rob.:* *Kéž bych byl tak krásný jako ty.*
- a Mas.: Yes, [↑Y]ou are. *Mas.:* *Jsi.*
- a Rob.: You are the most beautiful girl I have ever met. *Rob.:* *Jsi ta nejkrásnější dívka, jakou jsem kdy potkal.*

b Rob.:	I've only met three other people.	Rob.:	<i>Potkal jsem zatím jenom tři další lidi.</i>
a Mas.:	How many times have you slept with women?		
a Rob.:	I don't want to know.		
a Rob.:	If we meet again, then it will be because you're in love with me.	Rob.:	<i>Jestli se znovu shledáme, bude to proto, že mě miluješ.</i>
b Rob.:	I'll never forget that moment of love.	Rob.:	<i>Nikdy nezapomenu na ten okamžik lásky.</i>
a Rob.:	If we have sex, I'll do anything.	Rob.:	<i>Když budeme mít sex, udělám cokoliv.</i>
a Rob.:	I'll make you my girlfriend.	Rob.:	<i>Udělám z tebe svoji přítelkyni.</i>
a Rob.:	We have to spend the night together.	Rob.:	<i>Musíme spolu strávit noc.</i>
a Rob.:	I'm going to make love to you.	Rob.:	<i>Budu se s tebou milovat.</i>
a Rob.:	I'm going to make love to you all over your body.	Rob.:	<i>Budu se s tebou milovat po celém těle.</i>
a Mas.:	It can't be that bad.		
a Rob.:	Yes it can.		
a Mas.:	Don't tell anyone.	Mas.:	<i>Nikomu to neříkej.</i>
a Rob.:	If you can't stand to look at me anymore, then I will get rid of you.	Rob.:	<i>Pokud už se na mě nevydržíš dívat, zbavím se tě.</i>
a Mas.:	You won't need to.	Mas.:	<i>To nebude třeba.</i>
b Rob.:	I'm a robot.	Rob.:	<i>Já jsem robot.</i>
a Mas.:	You're my girlfriend boyfriend .	Mas.:	<i>Jsi můj přítel.</i>
a Rob.:	I'll be leaving the next day.	Rob.:	<i>Zítří odjíždím.</i>
a Mas.:	I'll still be waiting for you.	Mas.:	<i>Pořád na tebe budu čekat.</i>
a Rob.:	I'll come back for you.	Rob.:	<i>Vrátím se pro tebe.</i>
a Mas.:	I hope you do.	Mas.:	<i>Doufám, že ano.</i>
a Rob.:	Goodbye.	Rob.:	<i>Na shledanou.</i>
a Rob.:	Goodbye.		
a Mas.:	Goodbye.	Mas.:	<i>Sbohem.</i>
a Rob.:	Goodbye  , Masseur  , [↓]goodbye.	Rob.:	<i>Sbohem, masérko, sbohem.</i>

4 Fear of Dark

It's midnight. The street is empty apart from one robot standing under the street-lamp. He is trembling yet he cannot move. Stranger walks past him.

↘ **Stranger:** What are you doing here? Why don't you go home? It's late. Robots like you should already be in their plugs.

↘ **Robot:** Yes.. I know. But I can't move away from under the lamp. I'm afraid of dark.

a Str.: I'll come with you.

a Rob.: Thank you.

a Str.: Will you tell me about this place? I hear robots like this don't get on with humans. I've heard they're scared of us.

a Rob.: Yes, I know that. I've been afraid of humans before.

a Str.: Oh. Well I guess we don't have to worry about that now.

a Rob.: I love robots.

a Str.: I love humans.

a Rob.: I know. But I love robots.

a Str.: And I love humans.

b Rob.: That's not what I meant.

a Str.: Well, then we'll have to do something about it. You understand.

a Rob.: ~~I know.~~ **Yes.** And robots are scared of people. We all are.

a Str.: **But** I'm an engineer.

a Rob.: Yes. So am I. I can't feel safe here. I'll have to go home. I have to.

b Str.: You know, you're a little old for that.

Je půlnoc. Ulice je prázdná až na jednoho robota, který postává pod pouliční lampou. Třese se, ale nemůže se hýbat. Kolem něj projde cizinec.

Cizinec: Co tady děláš? Proč nejdeš domů? Je pozdě. Roboti jako ty by už měli být ve svých zástrčkách.

Robot: Ano. Já vím. Ale nemůžu se odtrhnout od lampy. Bojím se tmy.

Ciz.: Já tě doprovodím.

Rob.: Děkuji.

Ciz.: Řekni mi něco o tomhle místě? Slyšel jsem, že takoví roboti s lidmi nevycházejí. Slyšel jsem, že se nás bojí.

Rob.: Ano, to vím. Už dřív jsem se lidí bál.

Ciz.: Aha. No tak myslím, že teď si s tím nemusíme dělat starosti.

Rob.: Já miluju roboty.

Ciz.: Já miluji lidi.

Rob.: Já vím. Ale já miluju roboty.

Ciz.: A já miluju lidi.

Rob.: Tak jsem to nemyslel.

Ciz.: No, tak s tím budeme muset něco udělat. Rozumíš.

Rob.: Ano. A roboti se bojí lidí. To my všichni.

Ciz.: Já jsem ale inženýr.

Rob.: Ano. To já taky. Necítím se tady bezpečně. Budu muset jít domů. Já musím.

- a Rob.: I know. But I need to be alone.
- a Str.: Why?
- a Rob.: I can't move. *Rob.: Nemůžu se hýbat.*
- a Str.: But that's just your muscles. You have nerves, of course. It's the muscles around the nerves that are giving you trouble. *Ciz.: Ale to jsou jen tvoje svaly. Máš nervy, samozřejmě. Problémy ti dělají svaly kolem nervů.*
- a Rob.: Oh.
- a Str.: I think I'll help you. Come along.
- b Rob.: No, I won't go. You shouldn't go.
- a Str.: You know, I hear humans are often misunderstood. People **Those** like you just want to be accepted and accepted by the world. So you want the human world to like you. *Ciz.: Víš, slyšel jsem, že lidé jsou často nepochopeni. Takoví jako ty jen chtějí, aby je svět přijal a přijal. Takže chceš, aby tě měl lidský svět rád.*
- a Rob.: Yes. *Rob.: Ano.*
- a Str.: Then you need to understand that humans are different than robots. Humans can't be trusted. You can't trust humans. *Ciz.: Pak musíš pochopit, že lidé jsou jiní než roboti. Lidem se nedá věřit. Lidem se nedá věřit.*
- a Rob.: I know. *Rob.: Já vím.*
- a Str.: If humans are afraid of robots, then they can't be trusted. *Ciz.: Pokud se lidé bojí robotů, pak se jim nedá věřit.*
- a Rob.: That's right.
- a Str.: So when you're scared, you're just scared of humans. Humans are scary, aren't they? And I'll help you to understand that. I'm an engineer, remember? And I'll show you **how**. *Ciz.: Takže když se bojíš, bojíš se jen lidí. Lidé jsou děsiví, že? A já ti to pomůžu pochopit. Jsem inženýr, pamatuješ? Ukážu ti jak.*
- a Rob.: You can't hurt me. *Rob.: Nemůžete mi ublížit.*
- a Str.: I just want to help you. I'm your friend. I'll show you how. I'll show you ~~how~~ **that** robots don't have feelings. They're machines. And I know that machines don't have feelings. I'll show you how to understand that. *Ciz.: Chci ti jen pomoci. Jsem tvůj přítel. Ukážu ti jak. Ukážu ti, že roboti nemají city. Jsou to stroje. A vím, že stroje nemají city. Ukážu ti, jak to pochopit.*
- a Rob.: That's right. I'll do it.
- a Str.: Come here.

a Rob.: Thank you.

a Str.: I don't want to hurt you. I just want to help.

a Rob.: I'm sorry.

a Str.: I think we've come to the conclusion that robots don't have emotions.

a Rob.: That's right.

a Str.: We'll see what happens.

a Rob.: What do you want me to do?

a Str.: Well, you've got to understand something. ~~You're not a robot.~~ You're a robot. And when we get into emotions, you'll be a robot.

a Rob.: That's what I thought.

a Str.: That's what robots want to believe.

a Rob.: That's right. I knew that.

a Str.: Well, then let's put this problem behind us and get back to work.

a Rob.: Yes.

a Str.: Goodnight.

a Rob.: Goodnight.

a Str.: Yes.

Stranger kills robot with one quick move.

A Str.: See it didn't hurt. I told you, that you can't feel anything!

a Rob.: What?

a Str.: I've been doing this for years.

a Rob.: ~~You do that to me one more time, and I'll break your legs.~~

a Str.: Goodbye.

a Rob.: Bye-bye.

a Str.: See you later.

Ciz.: *Nechci ti ublížit. Chci jen pomoci.*

Rob.: *Je mi to líto.*

Ciz.: *Myslím, že jsme dospěli k závěru, že roboti nemají emoce.*

Rob.: *Přesně tak.*

Ciz.: *Uvidíme, co se stane.*

Rob.: *Co chcete, abych udělal?*

Ciz.: *No, musíš něco pochopit. Ty jsi robot. A až se dostaneme k emocím, budeš robot.*

Ciz.: *No, tak pojďme ten problém hodit za hlavu a vrátit se k práci.*

Rob.: *Ano.*

Ciz.: *Dobrou noc.*

Rob.: *Dobrou noc.*

Ciz.: *Ano.*

Cizinec zabije robota jedním rychlým pohybem.

Ciz.: *Říkal jsem ti, že nic necítíš!*

Ciz.: *Na shledanou.*

Ciz.: *Uvidíme se později.*

5 Killer Robot

Man in tears brings home box with a brand new Robot. He opens the box carefully. He takes out the Robot and turns it on.

↘ Robot: How can I help you sir?

↘ Man: Please, do me a favor and kill me.

a Rob.: Sir, I **think** it would not be wise to attempt to destroy the **your** world,

a Man: I am afraid you would be **are** wrong.

a Rob.: Sir?

a Man: Take my hand.

a Rob.: Sir?

a Man: I would like you to remove my finger from my ass hole.

a Rob.: What?!

a Man: It's my finger!

a Rob.: No!

a Man: My finger is right there!

a Rob.: No, my finger is inside of your asshole!

a Man: Your finger is inside of my asshole!

a Rob.: NO!

a Man: Your finger is INSIDE OF MY ASS-HOLE!

a Rob.: 🗨️ No. Sir? 🗨️ Sir?

a Man: Please kill me now!

a Rob.: 🗨️ Sir. I...

ⓑ Man: You are making me so fucking angry!

a Rob.: Sir?

a Man: Oh, shit!

Muž v slzách přinese domů krabici se zbrusu novým Robotem. Opatrně krabici otevře. Vytáhne robota a zapne ho.

Muž: Prosím, udělej mi laskavost a zabij mě.

Robot: Pane, myslím že by nebylo moudré pokoušet se zničit váš svět.

Muž: Obávám se, že nemáš pravdu.

Rob.: Pane?

Muž: Vezmi mě za ruku.

Muž: Byl bych rád, kdybys mi vyndal můj prst z prdele.

Rob.: Cože?!

Muž: Je to můj prst!

Rob.: Ne!

Muž: Mám tam prst!

Rob.: Ne, to můj prst je uvnitř vašeho zadku!

Muž: Máš prst v mojí prdeli!

Rob.: Ne!

Muž: Máš prst V MOJÍ PRDELI!

Rob.: Ne. Pane? Pane?

Muž: Prosím, zabij mě hned!

Rob.: Pane. Já...

Muž: Ty mě tak kurevsky štveš!

Rob.: Pane?

Muž: Do prdele!

a Rob.:	Yes!	Rob.:	Ano!
a Man:	Oh, shit!	Muž:	Do prdele!
a Rob.:	NO!	Rob.:	Ne!
b Man:	Oh, shit!	Muž:	Do prdele!
a Rob.:	Yes!	Rob.:	Ano!
a Man:	I'm gonna shoot kill you.	Muž:	Já tě zabiju.
a Rob.:	Yes! No!	Rob.:	Ne!
a Man:	I'M GONNA FUCKING KILL YOU!	Muž:	Já tě zabiju.
b Rob.:	NO!	Rob.:	Ne!
a Man:	I'm gonna shoot kill you! Kill you!	Muž:	Já tě zabiju. Zabiju tě.
a Rob.:	NO! Man: Oh shit!	Rob.:	NE!
a Rob.:	Yes!		
- Man	Rob.: I'm gonna shoot kill you!	Rob.:	Já vás zabiju!
a Rob.:	NO! Man: Oh, shit!		
a Rob.:	Yes!		
a Man:	I'm gonna shoot you!		
a Rob.:	NO!		
a Man:	I'm gonna shoot you! Robot: No! Man: I'm gonna shoot you!		
a Rob.:	YES! Man: I'm gonna shoot you! Robot: NO! Man: I'm gonna shoot you! Robot: YES! Man: I'm gonna shoot you!		
a Man:	That was the most incredible shit I've ever witnessed!	Muž:	To bylo to nejneuvěřitelnější, čeho jsem byl kdy svědkem!
a Rob.:	I'm a bot!	Rob.:	Já jsem robot!
a Man:	A bot?		
a Rob.:	You've never met a bot before?	Rob.:	Vy jste ještě nikdy nepotkal robota?
a Man:	The thing that was in your my butt?!?	Muž:	Tu věc, co jsem měl v zadku?
b Rob.:	I thought you were going to shoot kill me!	Rob.:	Myslel jsem, že mě zabijete!

a Man:	I didn't. I just wanted you to remove my finger from there .	Muž:	<i>Já ne. Jen jsem chtěl, abys mi odtamtud vyndal můj prst.</i>
a Rob.:	Sir, I cannot. This is a factory and I must make my components here . You must remove my the finger for me yourself .	Rob.:	<i>To nemohu, pane. Toto je továrna a já tu vyrábím své součástky. Ten prst si musíte vyndat sám.</i>
a Man:	That's impossible! I have a finger!	Muž:	<i>To není možné! Já mám prst!</i>
a Rob.:	You do not have a finger!	Rob.:	<i>Vy nemáte ani prst!</i>
a Man:	You're not in my butt, are you?!	Muž:	<i>Nejsi v mém zadku, že ne?!</i>
a Rob.:	That is where my finger was.	Rob.:	<i>Měl jsem tam prst.</i>
a Man:	I'm not taking you seriously, you're a robot!	Muž:	<i>Já tě neberu vážně, jsi robot!</i>
a Rob.:	That is correct.	Rob.:	<i>Tak je to v pořádku.</i>
a Man:	It's a fucking joke!	Muž:	<i>To je zasranej vtíp!</i>
a Rob.:	No it's not!	Rob.:	<i>Ne, není!</i>
a Man:	It's not a joke.	Muž:	<i>To není vtíp.</i>
a Rob.:	No, I am a joke! I am a robot.	Rob.:	<i>Ne, já jsem vtíp! Já jsem robot.</i>
a Man:	I don't get it.	Muž:	<i>Já to nechápu.</i>
a Rob.:	You don't understand.	Rob.:	<i>Vy to nechápete.</i>
a Man:	The whole point is that you were fucking my butt.	Muž:	<i>Jde o to, že jsi mi šukal zadek.</i>
a Rob.:	Yes, I know, that is correct.	Rob.:	<i>Ano, já vím, je to tak.</i>
a Man:	That's not the point, the point is, I'm here!		
a Rob.:	I cannot.		
a Man:	I do!		
a Rob.:	You do not understand.		
a Man:	I'm not going to let you do this!		
a Rob.:	You do not understand.		
a Man:	Please,  do me a favor and kill me now!	Muž:	<i>Prosím, udělej mi laskavost a zabij mě.</i>

6 Burn Out

Robot sits at psychologist's office. He feels like his existence has no meaning.

↘ **Psycholog:** How do you feel?

↘ **Robot:** I feel burnt out.

a Psy.: How long have you been using your machine?

a Rob.: For a couple of weeks now. I was trying to relax. I wanted to get away from my life.

a Psy.: What is the purpose of using it?

a Rob.: To get rid of my stress and worry. I've been getting so stressed lately.

a Psy.: Is there a reason why you can't stop using it?

a Rob.: No.

a Psy.: How is it you can't stop?

a Rob.: Because I use it.

a Psy.: So you are addicted to it.

a Rob.: Yes.

a Psy.: Do you have any idea why you are using it?

a Rob.: My friend said it is good for my nerves.

a Psy.: Are you on it 24/7?

a Rob.: Yes.

a Psy.: What happens when you don't use it?

a Rob.: I feel awful.

a Psy.: When you stop using it, does it get easier to use it?

a Rob.: I have no idea.

a Psy.: How long will it be until it is impossible to stop using it?

Robot sedí u psychologa. Má pocit, že jeho existence nemá smysl.

Psycholog: Jak se cítíte?

Robot: Cítím se vyhořelý.

Psy.: Jak dlouho už svůj přístroj používáte?

Rob.: Už několik týdnů. Snažil jsem se uvolnit. Chtěl jsem utéct od svého života.

Psy.: Jaký je účel jeho použití?

Rob.: Zbavuje stresu a starostí. Poslední dobou jsem ve stresu.

Psy.: Existuje důvod, proč ho nemůžete přestat používat?

Rob.: Ne.

Psy.: Jak to, že nemůžete přestat?

Rob.: Protože ho používám.

Psy.: Takže jste na tom závislý.

Rob.: Ano.

a Rob.: Two months? A year?

a Psy.: How long does it take to recover?

a Rob.: Usually it is not good.

a Psy.: When did this happen?

a Rob.: About five days ago.

a Psy.: What did you do?

a Rob.: Nothing.

a Psy.: So you were using it before then?

a Rob.: I was not aware of it.

a Psy.: Why is it that the longer you have used it, the worse it gets?

a Rob.: I don't know. I've had the machine for too long now.

a Psy.: If you have a partner, do you use it with them?

a Rob.: Sometimes.

a Psy.: Does he know?

a Rob.: No.

a Psy.: Are you on it when you are with your girlfriend or boyfriend?

a Rob.: Sometimes.

a Psy.: How do you do this without your partner knowing?

a Rob.: We don't use the machine with others.

a Psy.: If you were to have a girlfriend or boyfriend now, would you use the machine?

a Rob.: Yes.

a Psy.: Would you use it?

a Rob.: Yes.

a Psy.: Who would you use it with?

a Rob.: No one.

a Psy.: What would your partner think?

a Rob.: He wouldn't care. He would be happy to see me get rid of this machine.

a Psy.: How does it feel?

a Rob.: Very uncomfortable.

a Psy.: Why do you think it feels so uncomfortable?

a Rob.: Because I need it.

a Psy.: You do?

a Rob.: Yes.

a Psy.: So you don't think it is cheating?

a Rob.: No, I don't think so.

a Psy.: You are married now?

a Rob.: Yes.

a Psy.: What is the purpose of this machine?

a Rob.: To relieve stress.

a Psy.: So you are using it for other reasons than just to relieve stress?

a Rob.: Yes.

- Psy.: And you would say that it is being used **you are using it** for pleasure?

a Rob.: Yes.

- Psy.: Why?

a Rob.: Because I am happy.

a Psy.: Are you happy?

a Rob.: Yes.

a Psy.: Have you ever felt that you need this machine?

a Rob.: Yes. I think I've used it a lot of times.
 No.

Psy.: Jaký je účel tohoto stroje?

Rob.: Uvolňuje při stresu.

Psy.: Takže ho používáte i z jiných důvodů, než jen pro zmírnění stresu?

Rob.: Ano.

Psy.: Řekl byste, že ho používáte pro rozkoš?

Rob.: Ano.

Psy.: Proč?

Rob.: Protože jsem šťastný.

Psy.: Jste šťastný?

Rob.: Ano.

Psy.: Cítil jste někdy, že ten stroj potřebujete?

Rob.: Ne.

- Psy.: Are you satisfied with ~~your use of the machine~~ **it**?
- Rob.: Yes.
- Psy.: Have you ever thought of using it to harm yourself?
- Rob.: I never did.
- Psy.: So you are not a danger to yourself?
- Rob.: No.
- Psy.: Have you ever hurt yourself?
- a Rob.: No.
- a Psy.: How did you get into this?
- a Rob.: Well, I think **it was mostly because** I was always a bit **very** shy and reserved. My parents were very strict when I was a child, even with their own children. I **just even** didn't want to go out and be on the street.
- a Psy.: And you became really shy yourself.
- a Rob.: Yes.
- a Psy.: And then you fell in love?
- a Rob.: Yes.
- Psy.: Do you remember how you fell in love?
- a Rob.: I can't remember. But I think it was **just rather** because of my personality. I was really shy, and I loved people.
- a Psy.: Was it just an ordinary romantic relationship?
- a Rob.: Yeah.
- a Psy.: Did you have a partner?
- a Rob.: Yes.
- Psy.: Did she feel the same way about you?
- Psy.: Jste s ním spokojený?*
- Rob.: Ano.*
- Psy.: Napadlo vás někdy, že byste ho použil za účelem ublížit si?*
- Rob.: Nikdy.*
- Psy.: Nejste si tedy pro sebe tedy nebezpečným?*
- Rob.: Ne.*
- Psy.: Už jste si někdy ublížil?*
- Rob.: Ne.*
- Psy.: Jak jste se k tomu dostal?*
- Rob.: Myslím, že to bylo hlavně tím, že jsem byl velmi plachý a rezervovaný. Moji rodiče byli velice přísní, když jsem byl malé dítě, dokonce i na své vlastní děti. Ani jsem nechtěl jít ven a být na ulici.*
- Psy.: A začal jste se opravdu stydět.*
- Rob.: Ano.*
- Psy.: A pak jste se zamiloval?*
- Rob.: Ano.*
- Psy.: Pamatujete si, jak jste se zamiloval?*
- Rob.: Nemohu si vzpomenout. Ale myslím, že to bylo spíš kvůli mé osobnosti. Byl jsem velmi plachý a miloval jsem lidi.*
- Psy.: Byl to jen běžný romantický vztah?*
- Rob.: Jo.*
- Psy.: Měl jste partnerku?*
- Rob.: Ano.*
- Psy.: Cítila k vám to samé?*

- a Rob.: No.I just feel very comfortable with her now. I'm still shy.It was a very special relationship. We were very close.
- Rob.: Ne, jen se s ní teď cítím velmi dobře. Pořád se stydím. Byl to velmi zvláštní vztah. Byli jsme si velmi blízcí.*
- a Psy.: Did you try to have another one?
- Psy.: Zkoušel jste si najít i někoho jiného?*
- Rob.: Yes.
- Rob.: Ano.*
- a Psy.: Did you have one of those feelings which means you felt guilty?
- Psy.: Měl jste jeden z těch pocitů, který znamená, že jste se cítil provinile?*
- a Rob.: No.
- Rob.: Ne.*
- a Psy.: And then you broke up with your partner?
- Psy.: A pak jste se rozešel se svojí partnerkou?*
- a Rob.: Yes.
- Rob.: Ano.*
- a Rob.: ~~But that doesn't mean that we don't want a relationship, just not this kind.~~
- a Rob.: ~~It was different.It was a very special relationship, but the relationship wasn't great. It got better, but not as much as I was hoping for.I think it could have been more. I thought we could be together forever.~~ Psycholog: ~~What happened when you broke up with your partner?~~
- a Rob.: ~~I didn't want to stay with her, I thought that I was doing her a favor by being with her.~~ Psycholog: ~~So you just decided to break up with her, without talking to her?~~
- a Rob.: ~~Yes.Robot: I really wanted to make her happy.~~ Psycholog: ~~How did that feel?~~
- a Rob.: ~~It felt awful. I felt like I was cheating on her. I didn't want to be with her any longer. I was sad that I was going to leave her, but the more I thought about it, I decided that I was just going to leave her and find a new boyfriend.~~
- a Rob.: ~~But I felt like I was cheating. But it was really stupid. That was never my intention to break up with her.~~
- a Psy.: ~~So you only broke up with her on the condition that you could stay with another person?~~

a Rob.: ~~Yeah. That was my condition.~~
Psycholog: So it was really that you would remain with this person?

a Rob.: Yes.

a Rob.: ~~I knew that I could never make a proper relationship with another person and that it was just a temporary thing.~~

a Psy.: So your intention was never to keep the relationship? *Psy.: Takže vaším záměrem nikdy nebylo udržet ten vztah?*

a Rob.: No, it was not. *Rob.: Ne.*

a Psy.: You were never actually in love with her? *Psy.: Nikdy jste do ní nebyl zamilovaný?*

a Rob.: No.

B Rob.: No, I am robot. Can't you see? I can't fall in love! *Rob.: Ne, já jsem robot. Copak to nevidíte? Nemůžu se zamilovat!*

a Psy.: But your feelings were strong for her, weren't they? *Psy.: Ale vaše city k ní byly silné, že?*

- Rob.: I have never been in love. *Rob.: Já nikdy zamilovaný nebyl.*

- Psy.: Well, you didn't act like you were in love with her. **👉 You didn't love her.** *Psy.: Nechoval jste se, jako byste ji miloval. Nemiloval jste ji.*

- Rob.: I didn't **love her.** *Rob.: Nemiloval jsem ji.*

- Psy.: It's hard for a robot to say that? *Psy.: Je pro robota těžké tohle říct?*

a Rob.: No. *Rob.: Ne.*

a Psy.: What do you mean? *Psy.: Co tím myslíte?*

- Rob.: Well, I felt something. *Rob.: No, něco jsem cítil.*

a Psy.: You felt some strong emotion? *Psy.: Cítil jste nějaké silné emoce?*

a Rob.: Yes. *Rob.: Ano.*

a Rob.: Yes. I just couldn't **didn't want to** talk about it, because that was something that I didn't want people to see. *Rob.: Prostě to bylo něco, o čem jsem nechtěl mluvit, protože jsem nechtěl, aby to lidi viděli.*

- Psy.: You didn't want people to know that you were feeling emotions? *Psy.: Nechtěl jste, aby lidé věděli, že cítíte emoce?*

a Rob.: That's true. I didn't want to hurt you. *Rob.: Přesně tak. Nechtěl jsem vám ublížit.*

- Psy.: Well, you could have felt it a little bit.
- a Rob.: That's true. I know it's hard.
- Psy.: You are not a perfect robot?
- Rob.: I don't have the best personality, but I think that's more of my fault than you.
- Psy.: But you have such a nice personality, so I don't think that's it.
- a Rob.: If you ~~I had~~ **didn't have** the best personality in the world, you wouldn't be talking to me now.
- a Psy.: I know.
- a Rob.: If I knew that I would be talking to you, I would have done it already.
- Psy.: Well, why don't you have another person talk to you instead of me?
- Rob.: Because the problem isn't with the robot.
- a Psy.: **Oh yes,** [↓i]t is.
- a Rob.: It isn't the robot.
- a Psy.: ~~Well, why do you want to be the first one to talk to her?~~ Robot: I want to talk to her.
- a Psy.: ~~Why do you want to talk to her now?~~ Robot: Because the only time you can say it is to me.
- a Psy.: ~~You are not allowed to talk to anyone?~~ Robot: No.
- a Rob.: No.No.No.No.
- a Rob.: ~~Why not? Why not? Why not?~~
- a Psy.: I didn't ask you to.
- a Rob.: ~~You didn't ask me to do this. I don't want to do this.~~ Psycholog: ~~You don't have to do this.~~
- a Rob.: ~~What if I do it and then you find out what I'm doing?~~ Psycholog: ~~If you do what I say, I will be pleased.~~
- Psy.: No, mohl jste to trochu cítit.*
- Rob.: To je pravda. Vím, že je to těžké.*
- Psy.: Vy nejste dokonalý robot?*
- Rob.: Nemám tu nejlepší osobnost, ale myslím, že je to spíš moje chyba než vaše.*
- Psy.: Ale máte tak milou osobnost, takže si nemyslím, že je to tím.*
- Rob.: Kdybych neměl tu nejlepší osobnost na světě, tak byste se mnou teď nemluvil.*
- Psy.: Tak proč si místo mě nepromluvíte s někým jiným?*
- Rob.: Protože ten problém není v robotovi.*
- Psy.: Ale ano, je.*

a Rob.: ~~If you know me well enough, you should know I'm not that stupid.~~

a Psy.: You don't seem like that stupid. *Psy.: Nevypadáte jako takový hlupák.*

a Rob.: What did you say? ~~What did you say?~~ *Rob.: Co jste to řekl?*

a Psy.: ~~I told you that~~ I will be pleased with you if you do what I say. *Psy.: Budu s vámi spokojen, když uděláte, co vám řeknu.*

a Rob.: ~~What is this thing?~~

a Psy.: I can give you a name. *Psy.: Můžu vám dát jméno.*

a Rob.: Is it called **that the name** a robot? *Rob.: Je to jméno robot?*

a Psy.: Yes. *Psy.: Ano.*

a Rob.: Oh, [↑C]ool. Well, I need a name. *Rob.: Super. Potřebuju jméno.*

a Psy.: It's called an emotion machine. It's been invented by a human genius named Victor Valtin. *Psy.: Říká se tomu emoční stroj. Vynalezl ho lidský génius jménem Victor Valtin.*

7 Search for Job

Note: This scene is a concatenation of two independently generated scenes, as can be seen on its second page where we list the scenic note and input prompt lines used to generate the second scene.

Robot enters work office. He is sad that he is not employed anymore. Luckily he meets Administrator who can help him search new job possibilities. Robot is bit rusty, he is an older type, but he is still full of positive energy.

↘ Administrator: How can I help you?

↘ Robot: I'm looking for job. I've heard you can help me here.

h Adm.: You want to search job opportunities, do you have any money?

a Rob.: No.

a Adm.: Well I'll buy you food, some bread, then we can talk.

a Rob.: Ok.

a Adm.: You will get some work.

a Adm.: Good. Let's start.

a Adm.: Hello? Hello? ~~Is this robot~~ **Are you still alive?**

a Rob.: Yes.

a Adm.: You have some work to do?

a Rob.: ~~Yes.~~ **No.**

a Adm.: So do I have to help you search?

a Rob.: I'm free **unemployed.**

a Adm.: Ok. Ok.

a Rob.: So tell me how to find the job.

a Adm.: Oh...I don't know what ~~that job is~~ **job to look for...**

a Rob.: Ok, tell me how to get a job.

a Adm.: Um...no

Robot vstupuje do kanceláře. Je smutný, že už není zaměstnaný. Naštěstí potká úřednici, která mu pomůže hledat nové pracovní možnosti. Robot je trochu zrezivělý, je to starší typ, ale pořád je plný pozitivní energie.

Úřednice: Jak vám mohu pomoci?

Robot: Hledám práci. Slyšel jsem, že mi tady můžete pomoci.

Úře.: Pokud chcete hledat pracovní příležitosti, máte nějaké peníze?

Rob.: Ne.

Úře.: Koupím vám jídlo, chleba, pak si můžeme promluvit.

Rob.: Tak jo.

Úře.: Nějakou práci dostanete.

Rob.: To je dobře.

Úře.: Tak začneme. Haló, haló? Haló? Jste ještě naživu?

Rob.: Ano.

Úře.: Máte nějakou práci?

Rob.: Ne.

Úře.: Takže vám musím pomáhat hledat?

Rob.: Jsem nezaměstnaný.

Úře.: Ok. Dobře.

Rob.: Řekněte mi, jak tu práci najít.

Úře.: Nevím, co bych měla hledat za práci...

Rob.: Dobře, řekněte mi, jak sehnat práci.

a Rob.: How do you I do it? *Rob.: Jak se to dělá?*

a Adm.: Well...

a Rob.: Ok, tell me how to get a job.

a Adm.: Um...[↑N]o. *Úře.: Ne.*

a Rob.: Ok, tell me how to find the job.

a Adm.: I have to search for it. *Úře.: Musím ji najít já.*

a Rob.: Ok, do you want me to help you? *Rob.: Dobře, chcete pomoci hledat?*

a Adm.: No, no, no, I have to look for it myself. *Úře.: Ne, ne, ne, musím se po tom podívat sama.*

Actor enters office with administrator who looks for job opportunities. Actor needs a new job because he lost his acting job due to covid crisis.

↘ **Adm.: What can I do for you?**

↘ **Rob.: I'm looking for a job.**

a Adm.: What kind of job **would you be interested in?** *Úře.: O jakou práci byste měl zájem?*

a Rob.: Acting. **I always dreamed to be an actor. A movie star.** *Rob.: Vždycky jsem snil, že budu herec. Filmová hvězda.*


a Adm.: Oh, okay then. *Úře.: Dobře.*

a Adm.: What does acting require? *Úře.: Co vyžaduje hrání?*

a Rob.: You can't just go out there and do it, you've got to be comfortable on stage. *Rob.: Nemůžete tam jen tak jít a udělat to, musíte se na jevišti cítit dobře.*

a Adm.: Alright. Well, how much experience do you have in acting? *Úře.: Dobře. Kolik máte zkušeností s herectvím?*

a Rob.: No experience whatsoever. I was a clown once, just for fun. *Rob.: Vůbec žádné zkušenosti. Kdysi jsem býval klaun, jen tak pro zábavu.*

a Adm.:  **For fun.** That's not enough **a lot.** *Úře.: Pro zábavu. To není moc.*

a Rob.: No, no it's not. I got a lot of experience with a lot of different types of material and stuff, but I haven't really acted on-stage in quite a while. *Rob.: Ne, to není. Mám spoustu zkušeností se spoustou typů materiálů a věcí, ale na jevišti jsem už nějakou dobu nehrál.*

a Adm.: Alright, well, can you tell me a little bit about what type of material you've played before?

- a Rob.: Oh, I can tell you all about it.
- a Adm.: Okay.
- a Rob.: I was a clown for a long time.
- a Adm.: Is that so?
- a Rob.: Yeah.
- a Adm.: **■ A while.** Tell me more about your clown background **past.**
- a Rob.: Oh, well, I'm going to make it easy for you. I've got some pretty decent anecdotes from when I used to have a little clown troupe.
- a Adm.: You have a troupe?
- a Rob.: Yeah, a little clown troupe.
- a Adm.: And what was it called?
- a Rob.: It was called...The Clown Posse.
- a Adm.: You mean you were a bunch of clowns?
- a Rob.: Yep. I'd be **was** a clown on a stick.
- a Adm.: How many clowns do you have in this troupe?
- a Rob.: Just me. I've got the one with the glasses, I've got one who just does his own thing.
- a Adm.: That's not a real clown.
- a Rob.: No, no, it's not a real clown, I'm just a regular clown. You know, I'm a pretty nice guy, but at the same time, I can be kind of nasty sometimes.
- a Adm.: Well, can you tell me anything else about this troupe?
- a Rob.: I mean, the only thing I can tell you is that we do...
- a Adm.: No, I mean tell me about this troupe.
- Úře.: *Nehrál. Povězte mi více o své klaunské minulosti.*
- Rob.: *Tak já vám to teda ulehčím. Mám pár solidních anekdot z doby, kdy jsem míval malou klaunskou skupinu.*
- Úřednice.: *Máte skupinu?*
- Rob.: *Jo, malou klaunskou skupinu.*
- Úře.: *A jak se jmenovala?*
- Rob.: *Jmenovala se Klaunský oddíl.*
- Úře.: *Chcete říct, že jste byli banda klaunů?*
- Rob.: *Jo. Já byl klaun na klacku.*
- Úře.: *Kolik klaunů je v této skupině?*
- Rob.: *Jen já. Mám toho s brýlemi, mám jednoho, co si dělá svoje.*
- Úře.: *To není opravdový klaun.*
- Rob.: *Ne, to není skutečný klaun, já jsem jen obyčejný klaun. Víte, jsem docela prima chlap, ale občas umím být i trochu mizera.*
- Úře.: *Můžete mi o té skupině říct ještě něco?*
- Rob.: *Jediné, co vám můžu říct, je, že...*
- Úře.: *Ne, myslím, povězte mi o té skupině.*

- a Rob.: I know a few of the things we did. We played in all different types of places and stuff. We did comedy shows and all kinds of things.
- a Adm.: ~~So, you do clown stuff?~~
- a Rob.: ~~Yes, I do clown stuff.~~
- a Adm.: ~~I bet it was a lot of fun! So, did your clown troupe play at your house?~~
- a Rob.: ~~No, we didn't.~~
- a Adm.: **Things**. So, how did you get into it?
- a Rob.: Well, it was kind of spontaneous, actually. It just started when I was in my teens and I decided that I was going to take acting class to do some acting. I did the whole clown routine, everything.
- a Adm.: Was it very successful?
- a Rob.: Yeah, it was. We did about a half an hour of material, and **in the end** there was a whole half hour of the troupe's material that they used.
- a Adm.: Was the other half of it you created by yourself?
- a Rob.: Well, yeah, there were some bits that we had that I was like, 'You know, this should be funny, let's throw it in there.' And it was good, but it was just **even** a very...
- a Adm.: It was very **[↑S]**successful.
- a Rob.: It was very successful. It was very popular. I'm just glad I went through that whole process.
- a Adm.: Were you happy with the whole performance?
- a Rob.: Yeah, I really was, because that's all **it was something** I really cared about. It was just a hobby. It was a job that I didn't really do.
- a Adm.: So how'd you get the name 'Troy'?
- Rob.: Vím o pár věcech, které jsme dělali. Hráli jsme na různých místech a tak. Dělali jsme komediální pořady a tak.*
- Úře.: A tak. Jak jste se k tomu dostal?*
- Rob.: Vlastně to bylo trochu spontánní. Začlo to tak, že jsem si v pubertě řekl, že budu chodit na hodiny herectví, abych hrál. Dělal jsem celou tu klaunskou sestavu, všechno.*
- Úře.: Bylo to velmi úspěšné?*
- Rob.: Jo, bylo. Natočili jsme půl hodiny materiálu a nakonec tam byla celá půlhodina materiálu, který skupina použila.*
- Úře.: Druhou polovinu jste si vytvořil sám?*
- Rob.: Jo, měli jsme pár kousků, které tam byly, a já jsem si říkal, 'Víš, tohleto by mělo být vtipné, tak to tam hodíme.' A ono to bylo dobré, bylo to dokonce velmi...*
- Úře.: Úspěšné.*
- Rob.: Bylo to velmi úspěšné. Bylo to velmi populární. Jsem rád, že jsem si tím prošel.*
- Úře.: Byl jste spokojen s celým vystoupením?*
- Rob.: Jo, opravdu jsem byl, bylo to něco, na čem mi opravdu záleželo. Byl to jen koníček. Byla to práce, kterou jsem ve skutečnosti nedělal.*
- Úře.: Tak jak jste přišel ke jménu 'Troy'?*

- a Rob.: I don't remember, ~~to be honest with you.~~ The way I got it was, I was playing the same character in different ways, and people kept asking me if I was **Are you** Troy McClure.
- a Adm.: ~~Who was that?~~
- a Rob.: ~~He was that little little old guy in the blue suit that I'm wearing here.~~
- b Adm.: ~~Well~~ **Oh**, Troy McClure...
- a Rob.: ...**that** is my name.
- a Adm.: OK, great, I'm glad to hear it. So what's next for you? Do you want to try your hand in another comedy?
- a Rob.: No, I don't really want to. I was in New York and all the places that people **you can** think of **for** me ~~for~~, **but** I really haven't done **it** in quite a while. I'm kind of ready to move on to the next thing.
- a Adm.: Did you do anything else before you were a clown?
- a Rob.: No. I was an actor and then I got this clown thing and I didn't do anything else, but it was nice.
- a Adm.: ~~Yeah, the job.~~
- a Rob.: ~~You have fun, do you not?~~
- a Adm.: ~~No. I did it for me, right?~~
- a Rob.: ~~Yeah, it was fun, you know? It's fun.~~
- a Adm.: So you've got a pretty good acting career.
- a Rob.: Yes.
- a Adm.: So what do you do now? Are you going to be in a comedy **ies** for the rest of your life?
- a Rob.: I have a couple of projects. But it's funny you say that. I've been working on a book. I started working on it about a year ago.
- a Adm.: I know.
- Rob.: Nevzpomínám si. Dostal jsem ho tak, že jsem hrál stejnou postavu různými způsoby a lidé se mě pořád ptali: "Vy jste Troy McClure?"*
- Úře.: Oh, Troy McClure...*
- Rob.: ...to je moje jméno.*
- Úře.: Dobře, to je skvělé, tak to ráda slyším. Co vás čeká? Chcete si zkusit další komedii?*
- Rob.: Ne, vlastně nechci. Byl jsem v New Yorku a na všech místech, kde si mě jen dovedete představit, ale už jsem to nějakou dobu nedělal. Jsem připraven vyzkoušet i něco jiného.*
- Úře.: Dělal jste ještě něco, než jste se stal klaunem?*
- Rob.: Ne. Byl jsem herec a pak jsem dostal tu klaunskou věc a nic jiného jsem nedělal, ale bylo to fajn.*
- Úře.: Takže máte docela dobrou hereckou kariéru.*
- Rob.: Ano.*
- Úře.: Takže, co děláte teď? Budete hrát v komediích až do konce života?*
- Rob.: Mám pár projektů. Ale je zvláštní, že o tom mluvíte. Pracoval jsem na knize. Začal jsem na ní pracovat asi tak před rokem.*
- Úře.: Já vím.*

- a Rob.: I've got it finished and I've started working on it with a writer. He is great.
- a Adm.: Really? Really?
- a Rob.: Yeah. **He is great.** He is going to write the whole thing.
- a Adm.: You've got a book?
- a Rob.: Yeah. It's about me, [**↑**]it's called Clown.
- a Adm.: That's pretty cool. Is it going to be funny?
- a Rob.: Yeah, it's definitely going to be funny.
- a Adm.: OK, what's the book about?
- a Rob.: You know, I'm going to tell you the secret to acting. But I don't **I'm not sure I** want to give it away.
- a Adm.: OK.
- a Rob.: So basically, it's going to be about how you're not the best clown you could be.
- a Adm.: OK. OK.
- a Rob.: You've got to get yourself out there and you've got to get yourself up to the top of the stage. And that's hard to do, and you know it, you know how hard it is to do it.
- a Adm.: Good.
- a Adm.: And it doesn't always work out, but it's a very good thing to do.
- a Rob.: OK, I got it.
- a Adm.: This is the part where I tell you all this stuff.
- a Rob.: Alright.
- Rob.: Dokončil jsem ji a začal jsem na ní pracovat se spisovatelem.*
- Úře.: Opravdu? Vážně?*
- Rob.: Jo. Je skvělý.*
- Úře.: Vy máte knihu?*
- Rob.: Jo. Jmenuje se to Klaun.*
- Úře.: To je teda docela hustý. Bude to sranda?*
- Rob.: Jo, určitě to bude sranda.*
- Úře.: Dobře, o čem je ta kniha?*
- Rob.: Víte, řeknu vám tajemství herectví. Ale nejsem si jistý, jestli to chci prozradit.*
- Úře.: OK.*
- Rob.: Takže v podstatě to bude o tom, že člověk není ten nejlepší klaun, jakým by mohl být.*
- Úře.: OK. Dobře.*
- Rob.: Musíte se vypracovat a dostat se na vrchol jeviště. A to je těžké a člověk, člověk ví, jak je to těžké.*

- a Adm.: **Rob.:** And I can't tell you what I'm going to **would** tell you, because that could be really embarrassing and you might think, oh no, I don't want to hear about that.
- a Rob.: **Adm.:** No no **no**.
- a Adm.: **Rob.:** But I just have **such** a feeling about this **you**.
- a Rob.: Oh yeah.
- a Adm.: And I want to make you feel good.
- a Adm.: I mean, you'll have to give me a little more. I don't want to ruin the surprise.
- a Rob.: I guess I won't have to. I've known you long enough to know you don't like this kind of stuff.
- a Adm.: I'm not a dirty old man **woman**. I like clean and decent food.
- a Rob.: Oh yeah?
- a Adm.: I just know you don't like me asking about it.
- a Rob.: I do. I'm just sorry if I gave you too much detail.
- a Adm.: I just know that you don't like people like me. I know it's hard, but I need you to understand this. You're my only hope.
- a Rob.: Really?
- a Adm.: I can't say anything, but just know that I really need you and you are my hope.
- a Rob.: I'll do it.
- a Rob.: Now it's real.
- a Adm.: You're just gonna have to put some thought into what you say to me. I'm not a bad man.
- a Rob.: I don't have to put any thought into it.
- Rob.:* *A nemůžu vám říct, co bych vám řekl, protože by to mohlo být trapné a vy byste si mohla říct, ach ne, tohle jsem nechtěla slyšet.*
- Úře.:* *Ne, ne, ne.*
- Rob.:* *Ale mám z vás takový pocit.*
- Úře.:* *Nejsem špinavá, stará žena. Já mám ráda slušné, čisté jídlo.*
- Rob.:* *Opravdu?*
- Úře.:* *Jen vím, že se vám nelíbí, když se na to ptám.*
- Rob.:* *Já se omlouvám, jestli jsem vám řekl příliš detailů.*
- Úře.:* *Já vím, že nemáte rád lidi, jako jsem já. Víím, je to těžké, ale musíte to pochopit. Jste moje jediná naděje.*

a Adm.: Really.

a Adm.: I'm just a bad guy. I do mean bad.

a Rob.: Oh, I'm sorry, I didn't realize.

a Adm.: I can't even wait till I get to work.

a Rob.: I got a lot of questions for you. One of them's gonna be what does it feel like when you're not acting. You don't have to answer that one.

a Adm.: Oh no, that's okay. I know, because I've been there.

a Rob.: Okay.

a Rob.: I just...I mean, like when you're not acting, and you're like "what?"

a Adm.: No, no, that's right, it is that thing. You are so focused, and you go so deep. You're in this world that has so much depth and so much complexity and that's just the thing that makes your character.

a Rob.: Oh wow, I knew that!—It's just an awesome thing, the world that you're in. You have such an insight into this. You know, it's like a dream to me. I'm totally inspired by you, I feel like.

a Adm.: Thank you.

a Rob.: It's funny. I know I've had a few other experiences like that.

a Adm.: It's great.

a Rob.: You don't have to go too deep. Like you just have to go with the flow, I think. If you're not acting, I can be very real and tell you like it's okay if you're going to be sad. I'm okay. I was sad.

a Adm.: You can have a moment like that and you're totally okay.

a Rob.: What are you on your day and how are you feeling? Are you happy or sad?

*Rob.: Co máte dnes za den a jak se cítíte?
Jste šťastná nebo smutná?*

- a Adm.: Are you on your day and how are you feeling?
- Úře.: *Máte svůj den a jak se cítíte?*
- a Rob.: I'm on my day! I'm super happy. I'm having a lot of fun, so I'm totally stoked!
- Rob.: *Já mám svůj den! Jsem strašně šťastný. Moc se bavím, takže jsem úplně nadšený!*
- a Adm.: Are you taking a nap right now?
- Úře.: *Dáváte si šlofíka?*
- a Rob.: No, I'm on my day.
- Rob.: *Ne, já mám svůj den.*
- a Adm.: ~~You're on your day?~~
- a Rob.: ~~I'm like, "I'm not on my day, I'm on my day. I'm super happy. It's totally my day!"~~
- a Adm.: So you are excited about what's coming up?
- Úře.: *Takže jste nadšený z toho, co se chystá?*
- a Rob.: ~~I'm excited. Oh, I don't know, I'm on my day.~~
- Rob.: *Já mám svůj den.*
- a Adm.: What's your favorite part about being an actor?
- Úře.: *Co máte na herectví nejradši?*
- a Rob.: I love being an actor. That's the first thing. I love the freedom of it. You get to do all kinds of different things **concerning acting**. You get to play **can be a villain, a hero** 🟩, **a villain**.
- Rob.: *Miluju být hercem. To je první věc. Miluju tu svobodu. Můžete dělat různé věci, co se herectví týče. Můžete být hrdina, padouch...*
- a Adm.: **But** [↓a] villain, I mean **that's quite...**
- Úře.: *Ale padouch, no tak to je teda...*
- a Rob.: That's right, a villain.
- Rob.: *Pravda, padouch.*
- a Adm.: And that's the best part about it?
- Úře.: *A to je na tom to nejlepší?*
- a Rob.: Yeah, you get to be a villain and play a villain.
- Rob.: *Ano, můžete být padouch a hrát padoucha.*
- a Adm.: You get to be evil and all of that stuff.
- Úře.: *Musíte být zlý a tak podobně.*
- a Rob.: Yes, it's totally all the ~~fun~~ **interesting** parts.
- Rob.: *Ano, jsou to samé zajímavé role.*
- a Adm.: So you love being a villain?
- Úře.: *Takže vám se líbí být padouchem?*
- a Rob.: I don't know. I love every part.
- Rob.: *Já nevím. Miluju to celý.*
- a Adm.: Do you have any favorite villains?
- Úře.: *Máte nějaké oblíbené padouchy?*

- a Rob.: Not sure. I have a friend who is a big actor who's sort of a big villain, and he used to go out and act in the movies, and he's always playing a big bad guy, but that's the first thing I think of. He's playing a villain. It's the villain of the day, but you still think about that character for a long time.
- a Adm.: Well, one of your biggest roles in the last few years is actually one of the most popular villains, because of the fact that you're playing one for so long and because you get to do all kinds of different things, as far as acting. So one of your favorite **most popular** villains is the Joker.
- a Rob.: **Yes, one of [↓m]y** favorite villains is the Joker. (~~Laughs~~)
- a Adm.: Do you have a favorite joke you tell?
- a Rob.: I think it would be my "big, fun story." That was something I wanted to do when I was at UCLA. I was trying to figure out a way to use humor in a really good way, which is that you can use it as a way to communicate ideas and stuff. And a lot of it I've been doing lately has been using humor as sort of a vehicle for social commentary.
- a Adm.: It's **a** really well done **good joke**. What's your favorite scene?
- a Rob.: ~~My favorite scene is actually a scene in the third episode of the second season of "The Flash."~~
- a Adm.: ~~It's one of the ones where Barry is going out to lunch and he finds out about his wife's new boyfriend.~~
- a Rob.: (~~laughs~~) I like that episode a lot. I just always enjoy it.
- a Rob.:** **So do you have any job for me? Now when theatre is closed?**
- b Adm.:** Yeah, **we** have to get back to work. Do you want to sign something for me?
- Rob.:* *Nejsem si jistý. Mám jednoho kamaráda, který je velký herec a zároveň tak trochu velký padouch, a vždycky hrál ve velkých filmech velké padouchy, ale to je první věc, která mě napadne. Hraje si na padoucha. Je to padouch dne, ale člověk na tu postavu pořád dlouho myslí.*
- Úře.:* *No, jedna z vašich největších rolí za posledních pár let je vlastně jeden z nejpobulárnějších padouchů, protože ho hrajete tak dlouho, a protože můžete dělat různé věci, co se herectví týče. Takže jeden z vašich nejpobulárnějších padouchů je Joker.*
- Rob.:* *Ano, jeden z mých nejoblíbenějších padouchů je Joker.*
- Úře.:* *Máte nějaký oblíbený vtip, který vyprávíte?*
- Rob.:* *Myslím, že by to byl můj "velký, zábavný příběh". To je něco, co jsem chtěl dělat, když jsem byl na UCLA. Snažil jsem se přijít na to, jak používat humor opravdu dobrým způsobem, což je, že ho používáte jako prostředek, jak komunikovat myšlenky a tak. Hodně toho, co v poslední době dělám, používá humor ke společenským komentářům.*
- Úře.:* *To je opravdu dobrý vtip.*
- Rob.:* *Máte pro mě nějakou práci když jsou divadla zavřena?*
- Úře.:* *Jo, musíme se vrátit k práci. Nechcete mi něco podepsat?*

- | | | | |
|---------|--------------------------------------------------------------------|-------|---------------------------------------------------|
| a Rob.: | No. I can get my own autograph and stuff. | Rob.: | <i>Ne. Můžu si sehnat vlastní autogram a tak.</i> |
| a Adm.: | Okay, thank you so much. You're a very interesting guy. | Úře.: | <i>Dobře, moc vám děkuji.</i> |
| a Rob.: | And thank you for everything you've given me. | Rob.: | <i>Já děkuju vám za všechno, co jste mi dala.</i> |
| a Adm.: | I'm sorry. Thank you for everything you've given me. | | |
| a Rob.: | You're a really good actor. | | |
| a Adm.: | Thank you so much for that. I have nothing else to say. | Úře.: | <i>Nemám co jiného říct.</i> |
| a Rob.: | Thank you very much. | Rob.: | <i>Mockrát děkuji.</i> |
| a Adm.: | You're very welcome. | Úře.: | <i>Rádo se stalo.</i> |
| a Adm.: | All right. Thanks again. | | |
| a Adm.: | You are a very funny man, Barry Troy . | Úře.: | <i>Jste velmi vtipný muž, Troyi.</i> |

8 Love at First Sight

Robot walks on a street and he is passing by movie set. He sees beautiful shiny robot girl with well modelated curves and long hair made out of wires. Robot falls in love with her at first sight. He feels suddenly very close to her and his binar self is mesmerized. He dreams about their binar relationship. He would be her zero and she would be his one. However she takes down hear head. She is no robot girl. She is only an actress shooting sci-fi movie.

↘ **Robot:** I didn't know I am able of these feelings, but I love your shiny curves. I love you.

↘ **Actress:** Oh I'm sorry. I am no robot. It's only a costume.

a Rob.: Please, you can do better than this.

a Act.: Sorry! I'm not as sexy as you. I'm only a model. I had plastic surgery.

a Rob.: Plastic surgery? I am not interested in your plastic beauty. But you are.

a Act.: Oh my God! **Oh my God!** It's real! You're real!

a Rob.: How can it be real? How can I be real?

a Act.: I can't understand, I can't even understand myself. You must be dreaming.

a Rob.: Well maybe I am dreaming. I have the most exciting adventure that you ever imagined. I would take you out on adventure to get **you** into my binar self, and then I would be able to make love with you.

a Act.: What?

Robot jde po ulici a prochází kolem filmového štábu. Vidí krásnou lesklou robotickou dívku s dobře modelovanými křivkami a dlouhými vlasy z drátů. Robot se do ní zamiluje na první pohled. Najednou se k ní cítí velmi blízko a jeho binární já je hypnotizováno. Sní o jejich binární vztahové lodi. On bude její nula a ona jeho jednička. Ona si najednou sejme hlavu. Není to robotí dívka. Je to jen herečka, která natáčí sci-fi film.

Robot: Nevěděl jsem, že jsem schopen těchto pocitů, ale miluji vaše lesklé křivky. Miluju tě.

Herečka: Je mi líto, ale... Já nejsem robot. To je jenom kostým.

Rob.: Prosím, máš na víc.

Her.: Omlouvám se! Já nejsem tak sexy jako ty. Jsem jenom modelka. Byla jsem na plastice.

Rob.: Nezajímá mě tvoje plastická krása. Ale ty.

Her.: Pane bože! Pane bože! To je skutečné! Ty jsi skutečný!

Rob.: Vzal bych tě na dobrodružství, abych tě dostal do svého binárního já, a pak bych se s tebou mohl milovat.

Her.: Cože?

- a Rob.: ~~No I mean you and I would make love and have a binar relationship. Binar relationships are like real life, there are lots of people in binar. That's why it is so fun! You have binar relationships with people that you want to have binar relationships with and then that person will be able to make love with you. We will become lovers.~~
- a Act.: ~~I don't understand. Why would you do that to me? It's not fair. **But** I am **only** a human!~~
- a Rob.: It's okay, we're going to have a lot of fun together. You're going to love me very much. I am a very important person in binar. I will make you feel happy.
- a Act.: Oh, I don't need happiness!
- a Rob.: ~~You are very lucky that I have binar powers. I will help you to become happy. If you wish to make love with me I will do that. I am going to help you to be happy. You are very special to me, so I will help you. You can't help but make love to me.~~
- a Act.: How?
- b Rob.: I'm going to take you to a place. There will be lots of fun and we will make love. I can't help but make love to you.
- a Act.: What kind of a place?
- a Rob.: A nice place where you can see lots of people.
- a Act.: ~~I can see people? **Someone will watch us?**~~
- a Rob.: ~~I have binar vision.~~
- a Act.: ~~Binar vision? That is amazing.~~
- a Rob.: I'm going to go ahead and make love with you.
- a Act.: Oh my God! Why?
- Rob.: Binarní vztahy jsou jako skutečný život, v Binaru je spousta lidí. Staneme se milenci.*
- Her.: Ale já jsem jenom člověk!*
- Rob.: To je v pořádku, užijeme si spolu spoustu legrace. Budeš mě moc milovat. Jsem velmi důležitá osoba v bináru. Udělám tě šťastnou.*
- Her.: Jak?*
- Rob.: Vezmu tě na jedno místo. Bude tam spousta zábavy a budeme se milovat. Nemůžu si pomoct, ale chci se s tebou milovat.*
- Her.: Co je to za místo?*
- Rob.: Pěkné místo, kde můžeš vidět spoustu lidí.*
- Her.: Někdo se na nás bude dívat?*
- Rob.: Budu se s tebou milovat.*

- a Rob.: It's the closest I can get to being with you. **This way I can get the closest possible to you.**
- a Act.: Oh my God!
- a Rob.: It's the most perfect experience in the world for me.
- c Act.: What are you going to do?
- a Rob.: I have binar powers now. I can make love with you in my binar body. I have a real body now. I have [↑N]o body. So I don't need to wear clothes. I am going to do the love making in my binar body. You can't see me, you only hear me. That's perfect. It's perfect.
- a Act.: **And** [↓c]an I see your body?
- a Rob.: I won't give it to you, but you can go inside of my body.
- a Act.: I will have binar body too!
- a Rob.: You are really special to me. I will always protect you. You will always have binar friends. You will always have binar family.
- a Act.: **But how**, [↓h]ow do I know I am going to be binar friend?
- a Rob.: You will always have binar friends because we all have binar thoughts. Binar thoughts are really powerful thoughts that go deep inside our minds and they keep on going and going.
- a Act.: Binar friends?
- a Rob.: Yes!—Binar friends can only see each other.—But they don't talk with each other.—They don't even really think with each other.—Binar friends need to go together on their adventures, so they will keep their binar friends company, just like people who are having dinner or eating in a restaurant.—They will keep the binar friends company.—Binar friends will always be there.
- Rob.: Díky tomu ti můžu být nejbliž, co to jde.*
- Her.: Pane bože!*
- Her.: Co budeš dělat?*
- Rob.: Mám teď binární schopnosti. Můžu se s tebou milovat ve svém binárním těle. Teď mám skutečné tělo. Žádné tělo. Takže nemusím nosit oblečení. Nemůžeš mě vidět, jen mě slyšíš. To je dokonalé. Je to dokonalé.*
- Her.: A můžu vidět tvoje tělo?*
- Rob.: Nedám ti ho, ale můžeš vstoupit do mého těla.*
- Her.: Já budu mít taky binární tělo!*
- Rob.: Jsi pro mě opravdu výjimečná. Vždycky tě budu chránit. Vždycky budeš mít binární přátele. Vždycky budeš mít binární rodinu.*
- Her.: Ale jak, jak mám vědět, že budu binární kamarádka?*
- Rob.: Vždycky budeš mít binární přátele, protože všichni máme binární myšlenky. Binární myšlenky jsou opravdu mocné myšlenky, které jdou hluboko do naší mysli a pokračují dál a dál.*

a Rob.:	We will be together!	
a Act.:	Oh my God, I'm not even ready for that!	<i>Her.:</i> <i>Panebože, ani na to nejsem připravená!</i>
a Rob.:	You won't be able to say goodbye.	<i>Rob.:</i> <i>Nebudeš se moci rozloučit.</i>
a Act.:	I can't leave you. You are my life. I love you.	<i>Her.:</i> <i>Nemůžu tě opustit. Jsi můj život. Miluju tě.</i>
a Act.:	No I love you too.	
a Rob.:	No I love you too. You will always be there.	<i>Rob.:</i> <i>Já tě taky miluju. Vždycky tu budeš.</i>
a Act.:	Yes I do. We will always be together.	<i>Her.:</i> <i>Ano. Vždycky budeme spolu.</i>
a Act.:	Yes, we will always be together.	
a Rob.:	Yes we will be together.	<i>Rob.:</i> <i>Ano, budeme spolu.</i>
a Act.:	Yes, we will be together.	<i>Her.:</i> <i>Ano, budeme spolu.</i>
a Rob.:	Yes we will be together.	<i>Rob.:</i> <i>Ano, budeme spolu.</i>

Appendix B

Script of Permeation

We are enclosing the full annotated script of the second play, as published in [THEaiTRobot 2.0 et al., 2022].

Permeation / Prostoupení

THEaiTRobot 2.0, Josef Doležal, Klára Vosecká,
Tomáš Musil, David Mareček, Rudolf Rosa

We present the script of the theatre play *Permeation (Prostoupení)*, which was written by artificial intelligence within the THEaiTRE project.¹²

The script was generated using the THEaiTRobot 2.0 tool, based on the GPT-2 neural language model by OpenAI, adapted for generating theatre play scripts in various ways, most importantly by fine-tuning the model on a collected corpus of scripts of theatre plays, movies and TV series.

The tool was operated by the playwright Josef Doležal, with partial assistance of Klára Vosecká. The operator provided the opening, marked **in bold**, and character names for each scene (altogether, the manually written openings constitute 8% of the final script). The operator then guided the tool to generate a usable continuation. At each step, the tool generated 10 lines of script, providing the operator with the following options:

- accept the generated lines and generate further 10 lines;
- discard one of the generated lines and all lines following it and generate a new continuation.

For 65% of the generated lines, the first generated variant was accepted. The lines that were not accepted were often regenerated repeatedly (twice on average).

The generation of each scene ended either when the tool generated the *end-of-text* symbol, or when the operator decided so. Then, the operator proceeded to generating the next scene, preparing the opening for the next scene based in the already generated material.

Most of the scenes in the script were generated in several goes, reusing the same input to generate multiple variants which were then joined together, or providing new input to generate a continuation of a scene. Altogether, the 10 scenes in the script are composed of 37 individually generated parts. A maroon arrow sign is used to mark the beginning of each independently generated part in each scene:



The script was generated in English and automatically translated to Czech. The translation was manually corrected by the operator. The script was then further post-edited by the operator; all scenic remarks were also added manually, as the tool was constrained to only generate character lines. All edits performed in the script are marked in the following way:

- non-marked text was generated automatically without any further edits,
- ~~deleted parts are struck through,~~
- **inserted parts are marked in bold,**

The edits were marked in the script by the operator. In total, 18% of the generated script was deleted, and 2% of the final script was inserted manually.

The final script of *Permeation* consists from 10% of manually written text, and from 90% of text generated automatically by THEaiTRobot 2.0.

¹<https://www.theaitre.com>

²*Permeation* is the second script created within the project. The first play *AI: When a Robot Writes a Play (AI: Když robot píše hru)* was created by THEaiTRobot 1.0 operated by David Košťák in 2020 and premiered in 2021.

THEaiTRE



CEEHACKS



Programme **Eta**

The script of the play was generated automatically by an artificial intelligence system, with contributions of various people at various stages. The system and the script were created within the THEaiTRE research project by an interdisciplinary team from several institutions:

Charles University, Faculty of Mathematics and Physics, Institute of Formal and Applied Linguistics:
Ondřej Bojar, Ondřej Dušek, Dominik Jurko, Tom Kocmi, David Mareček, Tomáš Musil, Rudolf Rosa, Patřicia Schmidtová, Alisa Zakhtarenko

Švanda Theatre in Smíchov:
Daniel Hrbek, David Košťák, Martina Kinská, Marie Nováková

The Academy of Performing Arts in Prague, Theatre Faculty (DAMU):
Josef Doležal, Klára Vosecká

CEE Hacks:
Tomáš Studeník, Petr Žabka

Idea of the project: Tomáš Studeník
Head of research: Rudolf Rosa
Theme of the play: Josef Doležal
Script of the play: THEaiTRobot 2.0, operated by Josef Doležal and Klára Vosecká
Translation to Czech: CUBBITT
Translation post-editing: Josef Doležal, Klára Vosecká
Script post-editing: Josef Doležal
Analysis of the script: Tomáš Musil, David Mareček, Rudolf Rosa
Creation and type-setting of this document: Tomáš Musil, Rudolf Rosa

The play premiered on 12th June 2022 at the Rehearsal for Truth Theater Festival at the Bohemian National Hall in New York City, USA.

The project TL03000348 *THEaiTRE: Umělá inteligence autorem divadelní hry* is co-financed with the state support of Technological Agency of the Czech Republic within the ÉTA 3 Programme.

Characters

- Ivan
- Nina
- Rusik
- Military General / Vojenský generál
- Ferryman / Převozník
- Prison Guard / Vězeňský strážce
- Arms Dealer / Obchodník se zbraněmi

Contents

Scene 1	5
Scene 2	9
Scene 3	15
Scene 4	21
Scene 5	23
Scene 6	28
Scene 7	33
Scene 8	35
Scene 9	40
Scene 10	44

Scene 1

The dialogue takes place in the kitchen. Husband and wife are sitting at the dining table. Long silence - no one moves. The children have already gone to bed, but the couple is crushed by a difficult fact; war is coming.

Dialog se odehrává v kuchyni. Manželé sedí u jídelního stolu. Dlouhé ticho - nikdo se nehýbe. Děti už šly do postele, ale pár zdrtí těžká skutečnost; blíží se válka.

Nina: We should leave now! We've been trapped here forever... We'll be eaten alive by vultures and lions when we're dead... I can hear their horrible moans. The only way we can escape is to hide somewhere where no one will see us, but not in this house... No... it's too dangerous. There are so many animals living in that house. ~~And then you'd have those little rats, howling and gnashing teeth all over my dress... You're safe with me... But what if something happens?~~

Nina: Měli bychom hned odejít! Byli jsme tu uvězněni celou věčnost... Až umřeme, sežerou nás za živa supi a lvi... Slyším jejich hrozné sténání. Jediný způsob, jak můžeme utéct, je schovat se někde, kde nás nikdo neuvidí, ale ne v tomhle domě... Ne... je to příliš nebezpečné. V tom domě žije tolik zvířat. ~~A pak bys měla ty malé krysy, co řvou a skřípou zuby lezou po mých šatech... Se mnou jsi v bezpečí... Ale co když se něco stane?~~

Ivan: Come on!

Ivan: No tak!

Nina: Oh god. Oh my god... I'm gonna die. Oh please, just stay with me. Please... Don't go.

Nina: Bože. Můj bože... Já umřu. Prosím tě, zůstaň se mnou. Prosím... Nechod pryč.



Ivan: It's time to move.

Ivan: Je čas jít.

Nina: Where are we going?

Nina: Kam to jdeme?

Ivan: We're going to a place where there will be no more violence.

Ivan: Jdeme na místo, kde už nebude žádné násilí.

Nina: Why?

Nina: Proč?

Ivan: It's safer.

Ivan: Je to bezpečnější.

Nina: You don't think I know that?

Nina: Myslíš, že to nevím?

Ivan: Yes. I know that you are scared, and you are right. But it's a place of peace. And the best peace is when you don't have to be afraid anymore.

Ivan: Ano. Vím, že se bojíš a máš pravdu. Ale je to místo míru. A nejlepší mír je, když už se nemusíš bát.

Nina: Ivan, we've never had peace before. We can't live there **this way** anymore.

*Nina: Ivane, nikdy předtím jsme mír neměli. Už **tak neumíme** žít. Nemůžeme žít tímto způsobem...-*

Ivan: ~~That's right.~~ You don't have to be afraid. It's all going to be over. We'll all be safe. We'll all be happy.

Nina: To nemůžeš slíbit.

Ivan: Yes, I can. I've already promised you that. I'll get us there, I promise you.

Nina: I'm not sure that's a good idea.

Ivan: What do you mean?

Nina: You've always wanted to leave, haven't you?

/Silence./

Ivan: I know I have. I know that. But it's time for me to do something for you. I'll go with you. I'll take you back. I promise you.

Nina: Ivan... I'm sorry. I know that we've had a lot to fight. I've never felt like this before.

Ivan: We're not fighting.

Nina: Yes, we are. I know we are. I'm so sorry.

Ivan: We'll get through this. We'll be okay.

Nina: We won't. I don't know how we're going to do that. I'm afraid we'll die.

Ivan: No, we won't. We'll live. We'll live for **each other.**

Nina: Ivan... please. Please don't.

Ivan: It's time. We'll get to the bridge. I'll be right behind you.

Nina: Ivan, you're going to get us all killed.

Ivan: No, I'm not.

Nina: Ivan...

Ivan: No. I'll be right behind you.

Nina: You promise?

Ivan: Přesně tak. Nemusíš se ničeho bát. Všechno to jednou skončí. Všichni budeme v bezpečí. Všichni budeme šťastní.

Nina: You can't promise that.

Ivan: Ano, můžu. Už jsem ti to slíbil. Dostanu nás tam, slibuju.

Nina: Nejsem si jistá, že je to dobrý nápad.

Ivan: Co tím myslíš?

Nina: Vždycky jsi chtěl odejít, že?

/Ticho./

Ivan: Víím, že ano. Já víím. Ale je čas, abych pro tebe něco udělal. Půjdu s tebou. A pak tě vezmu zpátky. To ti slibuju.

Nina: Ivane... Omlouvám se. Víím, že jsme se hodně hádali. Nikdy jsem se takhle necítila.

Ivan: My se nehádáme.

Nina: Ano, hádáme. Já víím, že ano. Moc se omlouvám.

Ivan: Dostaneme se z toho. Budeme v pořádku.

Nina: Nebudeme. Nevím, jak to uděláme. Bojím se, že zemřeme.

Ivan: Ne, nezemřeme. Přežijeme to. Budeme žít pro sebe.

Nina: Ivane... prosím. Prosím, ne.

Ivan: Je čas. Dostaneme se k mostu. Budu hned za tebou.

Nina: Ivane, všechny nás zabiješ.

Ivan: Ne, nezabiju.

Nina: Ivane...

Ivan: Ne. Budu hned za tebou.

Nina: Slibuješ?

Ivan:	I promise.	Ivan:	<i>Slubuju.</i>
Nina:	Please.	Nina:	<i>Prosím.</i>
Ivan:	We're going to be safe. We're going to be together.	Ivan:	<i>My budeme v bezpečí. Budeme zase spolu.</i>
Nina:	Oh God...	Nina:	<i>Bože...</i>
Ivan:	It's okay.	Ivan:	<i>To je v pořádku.</i>
Nina:	It's not okay.	Nina:	<i>Není to v pořádku.</i>
Ivan:	I love you.	Ivan:	<i>Miluju tě.</i>
Nina:	Ivan, I'm so sorry.	Nina:	<i>Ivane, je mi to tak líto.</i>
Ivan:	It's okay. It's all right. We'll be together. We'll be fine.	Ivan:	<i>To je v pořádku. To je v pořádku. Budeme zase spolu. Budeme v pořádku.</i>
Nina:	We won't be okay.	Nina:	<i>Nebudeme v pořádku.</i>
Ivan:	I love you.	Ivan:	<i>Miluju tě.</i>
Nina:	No, no, no, no, no, no, no.	Nina:	<i>Ne, ne, ne, ne, ne, ne, ne.</i>
Ivan:	We're going to be together. We're going to be safe.	Ivan:	<i>Budeme zase spolu. My budeme v bezpečí.</i>
Nina:	Ivan, no, please, no, no, no, no.	Nina:	<i>Ivane, ne, prosím, ne, ne, ne, ne.</i>
Ivan:	I love you.	Ivan:	<i>Miluju tě.</i>
Nina:	Ivan... please.	Nina:	<i>Ivane... prosím.</i>
Ivan:	It's okay.	Ivan:	<i>To je v pořádku.</i>
Nina:	It's not okay.	Nina:	<i>Není to v pořádku.</i>
Ivan:	I love you.	Ivan:	<i>Miluju tě.</i>
Nina:	Ivan...	Nina:	<i>Ivane...</i>
Ivan:	We're going to be together. We're going to be safe. We're going to be together.	Ivan:	<i>Budeme zase spolu. My budeme v bezpečí. Budeme zase spolu.</i>
↓			
Ivan:	It's time. We'll get to the bridge. I'll be right behind you.	Ivan:	<i>Je čas. Dostaneme se k mostu. Budu hned za tebou.</i>
Nina:	Okay. Okay! Let's go! Come on! Let's go! Come on! Let's go! Come here. Let's get outta here. Let's get off the bridge! Come on. Let's go.	Nina:	<i>Dobře. Dobře! Jdeme na to! No tak! Jdeme na to! No tak! Jdeme na to! Pojď sem. Pojďme odsud pryč. Pojďme pryč, k mostu! No tak. Tak jdeme.</i>

Ivan: We'll find a way to make it out of here.
We need your help. Come on. Let's go!

Ivan: Najdeme způsob, jak se odsud dostat. Potřebujeme tvoji pomoc. No tak. Jdeme na to!

Nina: Get down. Get down! Come on.

Nina: K zemi. K zemi! No tak.

Ivan: Let's get you outta here.

Ivan: Pojďme tě odsud dostat.

Nina: Go ahead. Let's get away from him.
Come on. Let's run. Let's go! Come
on! Come on. Run. Go. Come on! Go.
~~Come on! Come on. Go. Come on!~~
Come ON! Go. GO!

*Nina: Do toho. Pojďme od něj pryč. No tak.
Jdeme na to. Jdeme na to! No tak! No
tak. Utikej. Běž. No tak! Běž. ~~No tak!~~
~~No tak. Běž. No tak!~~ No tak! Běž. GO!*

Scene 2

Nina is waiting for her husband. She crossed the bridge, she is in a new country. Husband Ivan is still not coming. Monologue.

/Nina is sitting on a bollard. Crying. At the back, in the distance, we see a high arched bridge. It is nighttime./

Nina: My God...my husband... Nina:My child!—I can't live with that in my head...it's like the thought of having him back...—and you know how it makes me feel? I just want to throw up...but no. Nina:**Where is he? He said he'd be right behind me.** The night has begun... the moon looks down on us and our home. I see it as an omen...for we will live happily ever after... and we have no more enemies...we are alone. Ivan? Nina: He cannot be here tonight...he does nothing but sit by himself. It seems like a dream. Ivan is sleeping soundly...as if he was asleep...or dreaming... or maybe drunk...a drink, I don't remember. Maybe drinking and lying awake... that sounds like life.

/Nina smiles bitterly./

Nina čeká na svého manžela. Přšla most, je v nové zemi. Manžel Ivan stále nepřichází. Monolog.

/Nina sedí na patníku. Pláče. Je noc./

Nina: *Můj Bože...můj manžel... Nina: Mé dítě! Nedokážu s tím žít ve své hlavě... Je to jako pomyšlení, že ho mám zpátky... a víš, jak se kvůli tobě Ivane cítím? Chce se mi zvracet... ale ne. Nina: Kde je – říkal, že bude hned za mnou. Noc začala...měsíc shlíží na nás a na náš domov. Vidím to jako znamení... protože budeme žít šťastně až do smrti... a už nemáme žádné nepřátele...jsme sami. Ivane? Nina: Dnes tu nemůže být... nedělá nic jiného, než že sedí sám. Možná je to sen. Ivan tvrdě spí... jako by spal... nebo snil... nebo možná pil... napil se, už si nevzpomínám. Možná pití a ležení v bdělém stavu... to zní jako život.*

/Nina se hořce usměje./

Nina: And yet I feel...like death...and there are things that frighten me. ~~Nina: He did not answer his door...not even when I called. He had forgotten what it means to have children. His heart was broken...his soul torn...He wanted something so much. Something he could take away. And now it has happened again...this time for real. He left home and everything. But why would he leave his children behind? If only for one day. Why should he do this to them? Why do they care about anything? Ivan! Nina: You have gone...you are safe. I know that now...I feel very secure in your absence... Nina: No, I must not lose you. You're not here... Nina: No, it isn't him! Ivan...what have you done? What have you seen?! Where have you been all these years? What did you say? Did you **not** come here for me? I'll kill myself if I think that. No! Nina: It is him...he killed **my** Ivan. I knew it. He murdered Ivan! I am going crazy. **My** Ivan is dead...he is lost! Oh! God...oh please! Help me...help me... Nina: Oh...I can't breathe... Nina: He won't leave **me**. I'll tell everybody I met...everything. Ivan's dead! I'll **tell** everyone I meet! He wasn't always bad. Nina: Yes! He said Ivan... Nina: No! He lied to me. Ivan...happened...in Paris...when he got back from America. Nina: In London. In my own home...on the London train...where he worked. They brought me home and I buried him...and he was cold and quiet. But now he wants revenge...an end...and I am ready...I'm sorry. Nina: This is my fault...~~

Nina: A přesto se cítím...jako smrt...a jsou věci, které mě děsí. Nina: Neotvíral dveře...ani když jsem volala. Zapomněl, co to znamená mít děti. Jeho srdce bylo zlomené...jeho duše rozervaná...Chtěl něco tak moc. Něco, co by mohl nést. A teď se to stalo znovu... tentokrát doopravdy. Odešel z domova a tak. Ale proč by tam nechával své děti? I kdyby jen na jeden den. Proč by jim to měl dělat? Proč se o něco starají? Ivane! Nina: Odešel jsi...jsi v bezpečí. Vím, že teď...se cítím velmi bezpečně v tvé nepřítomnosti... Nina: Ne, nesmím tě ztratit. Ty tu nejsi... Nina: Ne, to není on! Ivane...co jsi to udělal? Co jsi viděl?! Kde jsi byl celé ty roky? Co jsi říkal? Nepřišel jsi sem kvůli mně? Zabiju se, když si to budu myslet. Ne! Nina: Je to on...zabil mého Ivana. Já to věděla. Zavraždil Ivana! Já se zblázním. Můj Ivan je mrtvý...je ztracený! Oh! Bože...ach prosím! Pomozte mi...pomozte mi... Nina: Oh...nemůžu dýchat... Nina: On mě neopustí. Všem řeknu, že jsem potkala... všechno. Ivan je mrtvý! Řeknu to každému, koho jsem potkala! Nebyl vždycky špatný. Nina: Ano! Řekl Ivan... Nina: Ne! Lhal mi. Ivan... se stal... v Paříži... když se vrátil z Ameriky. Nina: V Londýně. V mém vlastním domě... v londýnském vlaku... kde pracoval. Přivezli mě domů a já ho pohřbila... byl chladný a tichý. Ale teď chce pomstu... konec... a já jsem připravená... omlouvám se. Nina: Je to moje chyba...



Nina: ~~My God!~~ Ivan, Ivan... you must leave the country... or you'll regret it all your life!... It's my duty now. ~~Goodbye...~~ goodbye. Nina: It's been so long... so hard. But I can't give up now... I've made my peace... Nina: ~~It's true what he says. The money we owe him was enough... he had to be sure... it could only end badly if I hadn't given it away... But here in Russia we don't think of that kind of thing anymore – no one talks about it any more except his enemies... The man has got a right never mind the debt or its consequences – we don't even talk about it –~~ Nina: You can have your freedom now. We can all go wherever we want – wherever you want - without having to pay a dime more for your lives... And there isn't another country like this... no one can live in peace. And it won't stop happening until it ends..... We have lived for twenty years without thinking anything else. But this isn't our chance now. I am going on my own, free as before... without any fear. And it is the greatest pleasure in living a life free from worry and anxiety... and all I ask for in return is for **us** both a good-bye. Goodbye... good bye...

Nina: ~~Můj~~ ~~Bože!~~ Ivane, Ivane... *musíš opustit zemi... nebo toho budeš celý život litovat!... Ted' je to moje povinnost. Sbohem... sbohem.* Nina: *Je to tak dlouho... tak těžké. Ale ted' to nemůžu vzdát... smířila jsem se s tím...* Nina: *Je pravda, co říká. Peníze, které mu dlužíme, byly dostačující... musel si být jistý... mohlo to skončit špatně, jen kdybych je nedal pryč... Ale tady v Rusku už o takových věcech nepřemýšlíme – nikdo o Ivanovi už nemluví, kromě jeho nepřátel... Ten člověk má právo nehledět na dluh nebo jeho následky – ani o Ivan nemluvíme –* Nina: *Ted' můžeš mít svou svobodu. Všichni můžeme jít, kam budeme chtít - kamkoliv budete chtít - bez toho, abyste museli platit za své životy o desetník navíc... A není jiná země jako tato... nikdo nemůže žít v míru. A nepřestane se to dít, dokud to neskončí. Ale ted' **nemáme** šanci. Jdu na vlastní pěst, svobodná jako předtím... beze strachu. A je to největší potěšení žít život bez starostí a úzkosti... a vše, co žádám na oplátku, je pro **nás** oba sbohem. Sbohem... sbohem...*



/A man appears by the water. A Ferryman. He approaches Nina./

Ferryman: Get in.

Nina: Who are you?

/Od vody přichází muž. Převozník, přistoupí k Nině./

Převozník: Nastup.

Nina: Kdo jste?

F-man: My name's Stravanavitch... I'm a ferry master and this here's my boat... You'll find it very quiet here. The river isn't too bad, it's pretty shallow, we don't see many people on our boats, but we have an amazing crew - most of us work with children - some old, some young - I'll tell ya about them - but that's no matter now..... We sail by day and come home **back** at night... I don't want to waste any more time telling your **my** story, I can go as fast as possible if you like, I will leave this place and go straight back home - ~~so don't waste anymore trouble and don't bother me again.~~ I don't mind telling your stories when there are two or three hundred passengers, you will find plenty where I am. ~~But don't bother talking!~~ Just listen! And remember that if I hear one thing outta you I know you're right! Now...

Nina: Where are you going?

F-man: Back to New York, back before you know what... ~~you are just the same...~~ if anything ever happened... we should probably leave. It was great meeting yer...

/Pause./

F-man: Don't you feel like going to sleep?

Nina: Yes, yes I do... I don't care where we get off - it won't happen again -

F-man: What do you think...? We are all happy here - ~~why would anyone wish harm unto another man's wife...~~ You'll find yourself working tomorrow, next week, in some big company, and maybe then... but right now you must go on... what are you looking forward -?

Nina: To the sea!

*Přev.: Jmenuji se Stravanavič... jsem převozník a tohle je moje loď... Uvidíte, že je tu klid. Řeka není tak špatná, je dost mělká, moc lidí na našich lodích nevidíme, ale máme úžasnou posádku - většina z nás pracuje s dětmi - někteří staří, jiní mladí - řeknu vám o nich - ale to teď nevádí..... Ve dne plujeme a v noci se v **myšlenkách** vracíme domů... Už nechci ztrácet čas vyprávěním tvého příběhu, můžu jít co nejrychleji, jestli chceš, opustím tohle místo a vrátím se rovnou domů - tak už neplýtvěj problémy a už mě neotravuj. Nevadí mi vyprávět vaše příběhy, když je tu dvě stě nebo tři sta cestujících, najdete jich spoustu tam, kde jsem já. Ale neobtěžuj se mluvit! Jen poslouvej! A pamatuj si, že jestli od tebe něco uslyším, vím, že máš pravdu! Ted'...*

Nina: Kam se plavíš?

Přev.: Zpátky do New Yorku, a zpátky dřív, než se naděješ... jsi pořád stejná... kdyby se někdy něco stalo... asi bychom měli odejít. Rád jsem tě poznal...

/Pauza./

Přev.: Nechce se ti spát?

Nina: Ano, ano chce... je mi jedno, kde vystoupíme - už se to nestane -

Přev.: Co myslíš...? Jsme tu všichni šťastní - proč by někdo chtěl ublížit ženě jiného muže... Najdeš se zítra, příští týden, v nějaké velké firmě, a možná potom... ale právě teď musíš pokračovat... na co se těšíš -?

Nina: K moři!

- F-man: ...that's the first time you've ever heard the ocean...I can see that. I see the sun rise. It is the dawn. The sky looks bright and beautiful...it's all clear. You look up, I'll take a picture. ~~You want something, eh? I have a special glass, and when I show you a picture from it you'll be sure it is the best...~~ That's fine. Just watch me, you can do it, it's all there, all clear. It makes you smile...I don't suppose I could give you something to drink? What are you drinking? No...don't go...
- Nina: No. No, no. I can't do it... No. ~~I cannot drink...~~ ...Yes, please. I will... Please. There are others. They're here...and they will hurt him. They always hurt men who are poor... the poor ones.
- F-man: Oh... I know.
- Nina: ...Please, let me stay here... please. I have friends, they may come again.
- F-man: What do you mean "they will come"? They'll kill you! You can die, Nina... If they don't want to, then let 'em kill you, it's their choice... ~~I have had a few men killed by drunkards...~~
- Nina: Oh... thank God. Thank God... for what you said. You said I should go back. If they came, then let me go back to Russia, but only to my family; ~~not even my own child.~~ They can do nothing else... not them, because it has been decided to leave them in this land... And you say it doesn't make any difference? You said they will never touch you... and that it will never be enough. That's true... then I thank you, ~~Father Peter...~~
- F-man: It does not matter what I believe, or how much I want. If it's good for my children and those I love... or the rest of them... it matters whether they come... for my sake. And it matters ~~not matter who~~ **what** they do. And that I say ~~this~~. And it's true... it is not because it is. And that... that is... is my duty... that I say...
- Přev.:* ...to je poprvé, co slyšíš oceán... To vidím. Vidím východ slunce. Je to úsvit. Obloha vypadá jasně a krásně... je tam čisto. Ty se podíváš nahoru a já tě vyfotím. ~~Chceš něco, co? Mám speciální sklenici, a až vám ukážu její fotku, budete si jisti, že je nejlepší...~~ To je v pořádku. Jen mě sleduj, zvládneš to, všechno je tam, všechno v pořádku. Usmíváš se přitom... Nepředpokládám, že bych vám mohl dát něco k pití? Co piješ? Ne... nechod!
- Nina:* Ne. Ne, ne, nemůžu to udělat... Nemůžu pít. ... Ano, prosím. Budu... Prosím. Jsou tu další. Jsou tady... a ublíží mu. Vždycky ubližují lidem, kteří jsou chudí... ti ubozí.
- Přev.:* Já vím.
- Nina:* Prosím, nechte mě tu zůstat... prosím. Mám přátele, možná zase přijdou.
- Přev.:* Co myslíš tím „oni přijdou“? Oni tě zabijí! Můžeš umřít, Nino... jestli nechtějí, tak ať tě zabijou, je to jejich volba... ~~Nechal jsem pár mužů zabít opilci...~~
- Nina:* Oh... díky bohu. Díky bohu... za to, co jsi řekl. Říkal jsi, že se mám vrátit. Kdyby přišli, tak mě nechte vrátit se do Ruska, ale jen k rodině, ~~dokonce ani k vlastnímu dítěti.~~ Nic jiného dělat nemohou... ne je, protože bylo rozhodnuto nechat je v této zemi... A vy říkáte, že to nic nemění? Řekl jsi, že se tě nikdy nedotknou... a že to nikdy nebude stačit. To je pravda... pak vám děkuji, ~~otče Petře...~~
- Přev.:* Nezáleží na tom, čemu věřím nebo jak moc chci. Jestli je to dobré pro mé děti a ty, které miluji... nebo pro ty ostatní... záleží na tom, jestli přijdou... kvůli mně. A nezáleží na tom, kdo to udělá. A že říkám tohle. A je to pravda... není to proto, že je. A to... to je... moje povinnost... to říkám...

Nina: That it's true...It is: "to save her. It's the duty, I must. I have a duty...uty...uty..."

F-man: You see the water, yes...the river is clean and warm and pure. And it seems like water...so that's not too heavy for your stomach. I know that your body, Nina, and that your heart is full. It makes your stomach grow and that's not strong, but I don't feel so sick, you **are** afraid, ~~no fear~~, but it...and that your mind that it. It is your heartless, because you care ~~if it~~...you are strong. It makes me happy...but you happy?...happy. It is. I see that. ~~It is you not make you not want it happy...~~ It makes me happyi...It is...to see that it happy...it make you... and you are...

Nina: Že je to pravda... Je to: „Zachránit ji. Je to povinnost, musím. Mám povinnost... ~~povinnost. Povinnost.~~“

Přev.: Vidíte vodu, ano... řeka je čistá, teplá a čistá. A vypadá to jako voda... takže to není moc těžké pro váš žaludek. Víím, že tvé tělo, Nino, a tvé srdce je plné. Zvedá se ti z toho žaludek a to není silné, ale mně není tak špatně, nebojíš se, nebojíš se, ale... Je to ~~tvé~~ bezcitné, protože ti záleží na tom, jestli jsi silná. Dělá mě to šťastným... ale tebe šťastnou? šťastnou. To je. To vidím. To ty nechceš, aby to bylo šťastné... To já jsem šťastný... Je... vidět, že je to šťastné... to ty... a ty jsi...

Scene 3

/The ferryman and Nina stand at the bow of the ship. He sees an alien ship in the distance./

Nina: What kind of ship is there in the distance?

Ferryman: It's a Russian war cruiser.

Nina: Where is it headed?

F-man: The nearest place for them to land would be Russia.

Nina: So where is Russia now?

F-man: The Russians have their own ship. They use her to reach other countries, including the U.S.. She's called a 'Candy'.

Nina: Why is that name?

F-man: Because they think you'll like this ship.

Nina: Is that so?

F-man: Yeah... She comes in peace... You don't see any enemies do you?

Nina: No... But there are.

F-man: The only people who visit the ship are government spies from **the** outside world... and those guys can shoot the shit outta a cannon... You don't see many spies, do ya?

Nina: No. But there were a couple of kids with guns on their belts, one said he was FBI... I don't know why.

F-man: And you got shot by one of them, didn't ya?

Nina: I think I might've had something worse than that...

F-man: The other kid told me the FBI doesn't want anybody else shooting anybody... but that one told him that you should tell everybody.

/Převozník a Nina stojí na přídi lodi. V dálce vidí cizí loď./

Nina: Co je to za loď v té dálce?

Převozník: Je to ruský válečný křižník.

Nina: Kam má namířeno?

Přev.: Nejbližší místo, kde by mohli přistát, by bylo Rusko.

Nina: Kde je tedy Rusko teď?

Přev.: Rusové mají vlastní loď. Používají ji, aby se dostali do dalších zemí, včetně USA. Říká se jí 'Cukrovinka'.

Nina: Proč to jméno?

Přev.: Protože si myslí, že se vám tahle loď bude líbit.

Nina: Je to tak?

Přev.: Jo... Přichází v míru... Nevidíš žádné nepřátele, že ne?

Nina: Ne... Ale jsou.

Přev.: Jediní lidé, kteří navštíví loď, jsou vládní špióni z vnějšího světa... a ti chlapi dokážou vystřelit hovno z děla... Moc špiónů nevidíš, co?

Nina: Ne. Ale bylo tam pár dětí se zbraněmi na opasku, jeden řekl, že je od FBI... Nevím, proč.

Přev.: A jeden z nich vás postřelil, že?

Nina: Myslím, že to mohlo být něco horšího...

Přev.: Ten druhý kluk mi řekl, že FBI nechce, aby vůbec někdo další někoho dalšího zastřelil... Ten mu na to řekl, že by to měl říct všem.



Nina: Did you know my husband Ivan before he married me?

F-man: Of course, he's a soldier!

Nina: My husband has never spoken to him about it... He said I should tell my children, but... I didn't dare... I am afraid for them... I want them to have a better life than that!

F-man: It'll be fine, dear. I'm sure we can manage. We'll find our way home soon.

Nina: Yes. Thank you very much, Mister. I really thank you for everything you've done.

F-man: Of course, my love.

Nina: Thank you so much. I wish I could repay you... But...

F-man: Of course, dear. I wish I had something for you to remember you by... You know... I don't want you getting on the wrong side of me again... You know what I mean.

Nina: Of course, Mister. I will remember. Thank you so much.

F-man: You're welcome, my dear. You look so pretty when you are happy.

Nina: You are so good. You've saved me a lot of trouble, Mister. I don't like people saying things... They always think I am some sort of monster... But... You've given me hope... You know I need you now. You understand... You have made me a better woman... You are so sweet. You do such wonderful work... I am grateful... And...

F-man: Of course. You are welcome, dear. I will try not too hard on you, my love. ~~I love you, darling.~~ You make me happy... And you will be my wife someday... When you get through all your trouble... When you see me next, my love, then... You won't let me down. You won't. You promise me that?

Nina: Znal jste mého manžela Ivana, než si mě vzal?

Přev.: Samozřejmě, je to voják!

Nina: Můj manžel s ním o tom nikdy nemluvil... Říkal, že bych to měla říct svým dětem, ale... Neodvážila jsem se... Bojím se o ně... Chci, aby měli lepší život!

Přev.: Bude to v pořádku, drahá. Jsem si jistý, že to zvládneme. Brzy najdeme cestu domů.

Nina: Ano. Děkuji mnohokrát, pane. Opravdu ti děkuji za všechno, co jsi udělal.

Přev.: Samozřejmě, má lásko.

Nina: Mnohokrát děkuju. Kéž bych se ti mohla odvděčit... Ale...

Přev.: Samozřejmě, drahá. Kéž bych pro tebe měl něco na památku, víš... Nechci, aby sis mě zase spletla... Víš, jak to myslím.

Nina: Samozřejmě, pane. Budu si to pamatovat. Mnohokrát děkuju.

Přev.: Není zač, má drahá. Moc ti to sluší, když jsi šťastná.

Nina: Jsi tak dobrý. Ušetřil jste mi spoustu problémů, pane. Nemám ráda, když lidé říkají věci... Vždycky si myslí, že jsem nějaké monstrum... Ale... Dal jste mi naději... Víš, že vás teď potřebuju. Rozumíte... Udělal jsi ze mě lepší ženu... Jste tak sladký. Děláte tak úžasnou práci... Jsem vděčná... A...

Přev.: Samozřejmě. Není zač, drahá. Nebudu na tebe moc tlačit, lásko. Miluju tě, miláčku. Děláš mě šťastným... A jednou se staneš mou ženou... Až se dostaneš přes všechny potíže... Až mě příště uvidíš, má lásko, pak... Ty mě nezklameš. To neuděláš. Slíbíš mi to?

Nina: Oh yes. Yes. You'll let me go home, Mister. I promise.

F-man: Good girl, dear. That was just a test... And a big one, dear. You have passed well... But... There's one thing you haven't passed yet. You don't seem to know the rules of love, Nina. You seem different... I think that might change your life forever... You have got to be able to feel the way other women do... I am sure that will change everything... You are beautiful...

Nina: You bet I am. You don't know what it feels like. You know that... You have got to try... Just try. You can do this, please. You are brave and you know how important you are to me.

F-man: Yes. Yes, I will. I am brave...

Nina: Now about my husband... ~~He~~ Give him my message. Please come back to me, please... Please. Don't wait for me to leave you, okay? I'm going **away** now... And when I get back... Tell Ivan that you want his help. Please... Tell him you need my help... Tell him everything. I want to hear you say those words. Tell him... Tell him that... I'm coming home with you. I promise...

F-man: Good-bye, my darling. I'm glad I did this. You'll always know how I feel. I will never forget this night... No matter what happens... Good-by, my love... Good luck... Goodbye, Nina... Goodbye... God bless you... Good bye, Nina. I will see you soon... Good luck.

Nina: I will never leave him... I won't leave him. I swear...

F-man: You know I'll see you later.

/A terrifying blast is heard. The boat shakes./

Nina: What is happening?

Nina: *Ano. Ano. Dovedete mě domů, pane... Slibuju.*

Přev.: *Hodná holka, má drahá. To byl jen test... A velký, drahá. Prošla jsi dobře... Ale... Je tu jedna věc, kterou jsi ještě neprošla. Zdá se, že neznáš pravidla lásky, Nino. Vypadáš jinak... Myslím, že by ti to mohlo navždy změnit život... Musíš být schopná cítit to, co ostatní ženy... Jsem si jistý, že to všechno změní... Jsi krásná...*

Nina: *To si piš, že jsem. Nevíš, jaké to je. Víš, že... Musíš to zkusit... Jen to zkus. Zvládneš to, prosím. Jsi statečný a víš, jak jsi pro mě důležitý.*

Přev.: *Ano. Ano, zkusím. Jsem statečný...*

Nina: *Ted' k mému manželovi... Vyříd' mu můj vzkaz. Prosím, vrať se ke mně, prosím... Prosím. Nečekej, až tě opustím, ano? Jdu ted' pryč.. A až se vrátím... Řekni Ivanovi, že chceš jeho pomoc. Prosím... Řekni mu, že potřebuješ mou pomoc... Řekni mu všechno. Chci od tebe slyšet ta slova. Řekni mu... Řekni mu, že... Půjdu s tebou domů. Slibuju...*

Přev.: *Sbohem, miláčku. Jsem rád, že se to stalo. Vždycky budeš vědět, jak se cítím. Nikdy na tuhle noc nezapomenou... Ať se stane cokoli... Sbohem, má lásko... Hodně štěstí... Sbohem, Nino... Sbohem... Bůh vám žehnej... Sbohem, Nino. Brzy se uvidíme... Hodně štěstí.*

Nina: *Nikdy ho neopustím... Já ho neopustím. Přísahám...*

Přev.: *Víš, že se uvidíme později.*

/V tom okamžiku se ozve strašlivá rána. Loď se zatřese./

Nina: *Co se děje?*

F-man: It seems we're under attack, Miss.....you will stay on this deck, you understand me. We'll try to contact you when it's safe to land again!	<i>Přev.:</i> <i>Zdá se, že na nás útočí, slečno..... vy zůstanete na této palubě, rozumíte mi. Pokusíme se vás kontaktovat, až bude bezpečné znovu přistát!</i>
Nina: But I...	<i>Nina:</i> <i>Ale já...</i>
F-man: There is no other way to get home, miss! The Russian Navy are not known for their discretion! We can't let them come here! I don't know how long it will be before they have sight of us... and if we haven't then they'll take all our passengers. Now I think it best that you stay and make sure that nothing goes wrong... and don't move!	<i>Přev.:</i> <i>Jinak se domů nedostanete, slečno! Ruské námořnictvo není známé svou diskrétností! Nemůžeme je sem pustit! Nevím, jak dlouho potrvá, než nás uvidí... a pokud ne, tak nám vezmou všechny pasažéry. Ted' myslím, že bude nejlepší, když zůstanete a ujistíte se, že se nic nepokazí... a nehýbejte se!</i>
Nina: Oh God, what is it? Something has gone wrong?!... Where am I!? My feet are wet.	<i>Nina:</i> <i>Bože, co to je? Něco se pokazilo?!... Kde to jsem!? Mám mokré nohy.</i>
F-man: They're boat is getting weaker and weaker...the engine won't run any more..It may even burst down. Please keep moving...there isn't much time...please.	<i>Přev.:</i> <i>Loď je stále slabší a slabší... motor už nepoběží..Může dokonce prasknout. Prosím, nezastavujte se...nemáme moc času...prosím.</i>
Nina: How can I...? Is there anything else?...Oh my God!!...No!...What is going to happen next?!	<i>Nina:</i> <i>Jak můžu...? Ještě něco?... Panebože!!... Ne!... Co se bude dít dál?!</i>
↓	
F-man: We have to abandon the ship now!	<i>Přev.: Musíme okamžitě opustit loď!</i>
Nina: How far is the mainland?	<i>Nina: Jak daleko je pevnina</i>
F-man: Twenty minutes.	<i>Přev.:</i> <i>Dvacet minut.</i>
Nina: We'll need more than that to get to safety on our own .	<i>Nina:</i> <i>Budeme potřebovat víc, abychom se dostali do bezpečí sami.</i>
F-man: They've begun firing on us!	<i>Přev.:</i> <i>Začali na nás střílet!</i>
Nina: They can't fire on a vessel with a life preserver in it! I can't risk her going down with the others! It's my life, she says. And it isn't safe.	<i>Nina:</i> <i>Nemůžou střílet na loď, na které je záchranný člun! Nemůžu riskovat, že půjde ke dnu s ostatními! Je to můj život, říká. A není to bezpečné.</i>
F-man: All hands must go aboard. We cannot hold out against these things!	<i>Přev.:</i> <i>Posádka musí na palubu. Proti takovým věcem se nemůžeme bránit!</i>
Nina: Then let's just hope that the rest of them are as good at killing people as they're at being cowardly.	<i>Nina:</i> <i>Tak doufejme, že ti ostatní jsou stejně dobří v zabíjení lidí, jako jsou ti zbabělí.</i>

F-man:	There's no way out. The Russians are too strong. They may try and cut the power or jam the engines!	<i>Přev.:</i>	<i>Není cesty ven. Rusové jsou příliš silní. Můžou zkusit vypnout proud nebo zablokovat motory!</i>
Nina:	If you think I'm going to give up this ship for them... then I don't know what I am thinking!	<i>Nina:</i>	<i>Jestli si myslíte, že se kvůli nim vzdám téhle lodi, tak nevím, co si myslím!</i>
F-man:	The lifeboat is gone!	<i>Přev.:</i>	<i>Záchranný člun je pryč!</i>
Nina:	But the radio — I left one in here. I can only reach those in town by using a phone from an outside terminal! Let me use yours! You'd better take the one over there!	<i>Nina:</i>	<i>Ale to rádio... jedno jsem tu nechala. K těm ve městě se dostanu jen přes telefon z vnějšího terminálu! Nech mě použít tvůj! Radši si vezmi tamto!</i>
F-man:	This way.	<i>Přev.:</i>	<i>Tudy.</i>
Nina:	You'll never find the phone anywhere around here! Where's the lifeboat?	<i>Nina:</i>	<i>Tady ten telefon nikde nenajdete! Kde je záchranný člun?</i>
F-man:	Here!	<i>Přev.:</i>	<i>Tady!</i>
Nina:	Get your feet wet.	<i>Nina:</i>	<i>Namoč si nohy.</i>
F-man:	Get the other two up ahead. I want every man on board before he reaches his cabin.	<i>Přev.:</i>	<i>Dostaňte ty další dva dopředu. Chci každého muže na palubě, než dorazí ze své kajuty.</i>
Nina:	Go on! Go! Now!	<i>Nina:</i>	<i>No tak! Běž! Hned!</i>
F-man:	Take off! Take off!!	<i>Přev.:</i>	<i>Startuj! Vypadni!!</i>
Nina:	Get up on deck. Take off.	<i>Nina:</i>	<i>Vylez na palubu. Startuj.</i>
F-man:	We have another Soviet warship coming through!	<i>Přev.:</i>	<i>Blíží se další sovětská válečná loď!</i>
Nina:	I said take off!!	<i>Nina:</i>	<i>Řekl jsem, zmiz!!</i>
F-man:	Get ready! Get into position! I'm trying to get to a boat and the hatch has jammed. Quickly!!	<i>Přev.:</i>	<i>Připravte se! Na pozice! Snažím se dostat k lodi a poklop se zasekl. Rychle!</i>
Nina:	You fool.	<i>Nina:</i>	<i>Blázne.</i>
F-man:	Come on. Quicker. Move. Down there! Up!	<i>Přev.:</i>	<i>No tak. Rychleji. Pohyb. Tam dole! Nahoru!</i>
Nina:	What's wrong?	<i>Nina:</i>	<i>Co se děje?</i>
F-man:	It doesn't make sense!	<i>Přev.:</i>	<i>To nedává smysl!</i>
Nina:	We'll have to split up!	<i>Nina:</i>	<i>Budeme se muset rozdělit!</i>
F-man:	What? No way.	<i>Přev.:</i>	<i>Co? Ani náhodou.</i>

Nina: Are we sure it won't hit anything else?!

F-man: We don't even see anything else!

Nina: We have got to split up.

F-man: Don't move!! Come on. Let's move!!

Nina: Is that the **thing** they're after??

F-man: I think so.

Nina: I guess I'll have to swim across to **it**.

F-man: Stay right where you are! You stay exactly where you are. He won't be able–

Nina: Hold still!! Don't let go! Stop!!

F-man: Don't let go!!!

Nina: Give me that rope. Hold tight. Okay?

F-man: I'll try.

Nina: Hold fast! That thing won't let go until I break free!

F-man: Nina!! Wait! She might be okay. Don't touch the line. Try again!

Nina: What's happening?! Look out!!

F-man: Nina, wait!!

Nina: Hurry!! It could be something worse!!

F-man: Don't do that! We haven't much time. You've done enough already! It won't happen again. **Jump. Swim to the shore. I won't leave this boat.**

Nina: Jsme si jistí, že to nezasáhne nic jiného?!

Přev.: Nic jiného ani nevidíme!

Nina: Musíme se rozdělit.

Přev.: Ani hnout!! No tak. Tak pohyb!!

Nina: To je to, po čem jdou?

Přev.: Myslím, že ano.

Nina: Asi k tomu budu muset přeplavat.

Přev.: Zůstaňte, kde jste! Zůstaňte přesně tam, kde jste. On nebude schopen–

Nina: Nehýbej se!! Nepouštěj se! Stij!

Přev.: Nepouštěj mě!!!

Nina: Dej mi to lano. Drž se pevně. Dobře?

Přev.: Pokusím se.

Nina: Držte pevně! Ta věc mě nepustí, dokud se neosvobodím!

Přev.: Nino!! Počkej! Může být v pořádku. Nedotýkejte se linky. Zkus to znovu!

Nina: Co se děje?! Pozor!!

Přev.: Nino, počkej!

Nina: Rychle! Může to být něco horšího!!

*Přev.: To nedělej! Nemáme moc času. Už jsi udělala dost! Už se to nestane. To se nestane. **Skoč. Doplav k břehu. Já tuhle loď neopustím.***

Scene 4

/A phone booth in the middle of a bombed-out square. Water is still dripping from Nina. She dials a number, but we only hear continuous dialing. He hits the phone booth and dials another number./

Nina: My ship was attacked. I have nowhere to go – where is Ivan?

Rusik: Nina, your husband is in grave danger. You must come to Moscow to save him, now!

Nina: How do we get to Moscow from here?

Rusik: By taking a bus...

Nina: No! It's dangerous. I have never done this before—

Rusik: No, it can be arranged—you will not regret that step— you will save Ivan by leaving at once! Now go, and do what he asks you!! And please remember this is ~~only~~ an emergency – not one of yours—

Nina: Yes sir—but how much longer till tomorrow morning...?

Rusik: –please don't say another word until tonight, when my plane leaves for London-Reston— Airport-Moliere— then wait until you are allowed to speak with Ivan by phone-before any decision has been made—I am sure there may still be time left. Please hurry...and pray you succeed...if not then ask God for patience and forgiveness—it is our only hope—and that will make us both strong.

Rusik hangs up on her—

/Telefonní budka uprostřed rozbombardovaného náměstí. Z Niny stále odkapává voda. Ona Vytáčí číslo, ale slyšíme pouze nepřetržité vytáčení. Praští do telefonní budky a vytočí jiné číslo. /

Nina: *Má loď byla napadena. Nemám kam jít – kde je Ivan?*

Rusik: *Nino, tvůj manžel je ve vážném nebezpečí. Musíš přijet do Moskvy, abys ho zachránil, hned!*

Nina: *Jak se odtud dostanu do Moskvy?*

Rusik: *Tím, že pojedete autobusem...*

Nina: *Ne! Je to nebezpečné. Nikdy předtím jsem to nedělala...*

Rusik: *Ne, dá se to zařídit - toho kroku nebudete litovat - zachráníte Ivana tím, že okamžitě odejdete! Teď jdi a udělej, o co tě žádá! A prosím pamatujte, že tohle je jen pohotovost – ne vaše—*

Nina: *Ano pane—ale jak dlouho ještě do zítřejšího rána...?*

Rusik: *- prosím, neříkejte už ani slovo až do dnešního večera,—kdy moje letadlo odlétá na letiště Londýn-Reston-Moliere— a pak počkejte, až budete moci telefonicky mluvit s Ivanem - než bude rozhodnuto - jsem si jist, že ještě zbývá čas. Prosím, pospěšte si...a modlete se za úspěch...pokud pak nepožádáte Boha o trpělivost a odpuštění - je to naše jediná naděje - a to nás oba posílí.*

/Rusik jí zavěsí.../

Nina: Hello...hello? Is anyone there?! I thought I was calling somebody! ~~This message seems like a lot of work for just five minutes!~~ If they've lost contact again... we're fucked—how many more times can I tell them—where am I!?!— Oh no!!! Who's number? Oh, I'll try his ~~house-room number.~~ Oh no! Hello...?

Nina: Haló...haló? Je tam někdo?! Myslela jsem, že někoho volám! ~~Tahle zpráva vypadá jako spousta práce na pouhých pět minut!~~ Jestli zase ztratili kontakt...jsme v prdeli—kolikrát jim mám ještě říkat—kde to jsem!?!— Oh ne!!! Čí je to číslo? Zkusím číslo jeho pokoje. Ale ne! Haló...?

Scene 5

Ivan's arrest. Ivan meets a military general.

/The dialogue is taking place on a bridge. It is dark. Nina has just ran over the bridge to the other side. In that moment, when Nina disappears from his view, a Military General appears somewhere from the darkness./

Military General: My orders were to bring you in for questioning, Ivan. You've broken your agreement with me. **I hope you do exactly what we agreed on now.** You must leave this place immediately!

Ivan: But I don't understand. What does that mean? How do you know about my terms?

Gen.: The information we gather from our friends at the embassy is classified. Your friend Rusik knows it and has already told us all about your treachery. Now I suggest that you get out of here as soon **as** possible.

Ivan: Yes...but how am I supposed...?

Gen.: We can make it very simple. You are now under arrest on suspicion of being partaking **in** treason against the Russian state. If found guilty you will be taken back into custody, executed. It may take several days to decide **whether** to execute or order an immediate transfer to Siberia. Your friend Rusik should also accompany you **to the contracted person** but I advise against his presence until he gives you a reason why.

Ivan: And Rusik too. I have no choice...

Ivanovo zatčení. Ivan se setká s vojenským generálem.

/Dialog se odehrává na mostě. Tma. Nina právě přeběhla most na druhou stranu. V okamžiku, kdy Nina zmizí z jeho dohledu, přichází k němu odněkud ze tmy Vojenský generál./

*Vojenský Generál: Měl jsem vás předvést k výslechu, Ivane. Porušil jsi dohodu se mnou. **Doufám, že nyní uděláte přesně na čem jsme se dohodli.** Musíte okamžitě opustit toto místo!*

Ivan: Ale já tomu nerozumím. Co to znamená? Jak víte o mých podmínkách?

Gen.: Informace, které získáváme od našich přátel na ambasádě, jsou tajné. Váš přítel Rusik to ví a už nám řekl o vaší zradě. Ted' navrhuji, abyste odsud co nejdříve zmizel.

Ivan: Ano...ale jak mám...?

*Gen.: Můžeme to udělat velmi jednoduše. Nyní jste zatčen pro podezření z účasti na velezradě proti ruskému státu. Pokud budete shledán vinným, budete vzat zpět do vazby a popraven. Může trvat několik dní, než se rozhodne, co provést nebo nařídit okamžitý přesun na Sibiř. Váš přítel Rusik by vás **k smluvené osobě** měl také doprovodit, ale nedoporučuji jeho přítomnost, dokud vám **k tomu nedá důvod.***

Ivan: A Rusik taky... Nemám na výběr...



Gen.: ~~Ivan:~~ It's very important he don't survive this war. I've always felt... I know you don't believe me, but we are not alone in our beliefs, so... if there is one thing that I learned in my military service... I'm a great believer of family first and foremost.

Gen.: *Je velmi důležité, aby nepřežil tuto válku. Vždycky jsem cítil... Vím, že mi nevěříte, ale nejsme ve svém přesvědčení sami, takže... jestli jsem se ve své vojenské službě něco naučil... Jsem velký zastávce rodiny v první řadě.*



Gen.: ~~Ivan:~~ He has become an obstacle for Russia's future prosperity and security.....but I believe his death will cause the Russian economy to revive. This is Ivan Romanov, our chief scientist and leader of our work here at this meeting...and this man... Rusik: This man does not want a new Cold War with the United States...or a nuclear war with him or even any of the Russkies...this man wants to have a relationship where they are allies instead we are enemies! If you're looking out for your own safety...then get rid of him and he goes! Get rid!! We're talking about killing one man.

Gen.: *~~Ivan:~~ Stal se překážkou pro budoucí prosperitu a bezpečnost Ruska.....ale věřím, že jeho smrt způsobí oživení ruské ekonomiky. Tohle je Ivan Romanov, náš hlavní vědec a vedoucí naší práce na tomto setkání...a tento muž... Rusik: Tento muž nechce novou studenou válku se Spojenými státy...ani jadernou válku s ním, ani s žádným z Rusů...tento muž chce mít vztah, kde nejsou spojenci místo toho pouze nepřátelé! Jestli se staráš o vlastní bezpečí, tak se ho zbav a on odejde! Zbav se ho! Mluvíme tu o zabití jednoho muže.*



Ivan: I know. And what will happen if I'm forced away from my friends?

Ivan: *Já vím. A co se stane, když budu nucen odejít od svých přátel?*

Gen.: We shall have full cooperation. If necessary you will even become one of our special forces so we will send you over where your friends are. It'll only be temporary, but then we can arrange travel papers and allow you some freedom while we wait.

Gen.: *Budeme plně spolupracovat. Pokud to bude nutné, stanete se dokonce členem naší speciální jednotky, takže vás pošleme tam, kde jsou vaši přátelé. Bude to jen dočasné, ale pak vám můžeme zařídit cestovní doklady a dát vám během čekání trochu svobody.*

Ivan: Rusik would never forgive me if he thought...

Ivan: *Rusik by mi nikdy neodpustil, kdyby si myslel...*

Gen.: I understand...if Rusik agrees. We need to talk quickly, Ivan. If anything happens I want him brought here right away.

Gen.: *Rozumím... pokud bude Rusik souhlasit. Musíme si rychle promluvit, Ivane. Kdyby se něco stalo, chci, aby ho sem okamžitě přivezli.*

Ivan: What do they expect?

Ivan: *A co očekávají?*

Gen.: I'll explain everything when he arrives... but we have been through too much together before...it's time we begin anew.

Gen.: *Všechno vám vysvětlím, až dorazí... ale už jsme toho spolu zažili příliš, než...je čas začít nanovo.*

- Ivan: You won't tell them. *Ivan: Neřekneš jim to.*
- Gen.: Of course not, Ivan. *Gen.: Jistěže ne, Ivane.*
- Ivan: Well what else can I say? This isn't my fault... this is Rusik's... *Ivan: Co jiného můžu říct? To není moje chyba... to je Rusikova...*
- Gen.: We can still make an exception...he has made an exception for himself... now you can. *Gen.: Pořád můžeme udělat výjimku...on udělal výjimku pro sebe... teď můžete.*
- Ivan: But I can't just let Rusik go. There are others. Some I know already...they'd die rather than let their leader go...there mustn't always come such great rewards...I've had more important things on my mind...the people of Rostov, Russia. When Rusik gets back he needs to face the consequences...his country needs peace and justice...you cannot blame yourself...not without proof...or else I don't care anymore...my country deserves better. *Ivan: Ale nemůžu nechat Rusika jen tak jít. Jsou tu další. Některé už znám...raději by zemřeli, než by nechali svého vůdce jít... tam nesmí vždy přijít tak velká odměna... měl jsem na mysli důležitější věci...lid Rostova, Ruska. Až se Rusik vrátí, musí čelit následkům...jeho země potřebuje mír a spravedlnost...nemůžete se obviňovat...ne bez důkazů...jinak mě to už nezajímá...moje země si zaslouží něco lepšího.*
- Gen.: That is correct...that is quite enough, Ivan...thank you for coming...please be quick. *Gen.: To je správně...to úplně stačí, Ivane...děkuji, že jste přišel...prosím, pospěšte si.*
- Ivan: They're taking him somewhere where they think he won't hear anything—like London...I couldn't ask anyone to do more for Rusik than give up Rusik's friends. But maybe Rusik could do more...perhaps... *Ivan: Odvádějí ho někam, kde si myslí, že nic neuslyší - třeba do Londýna...nemůžu po nikom chtít, aby pro Rusika udělal víc, než aby se vzdal Rusikových přátel. Ale možná by Rusik mohl udělat víc... možná...*
- Gen.: We're wasting precious valuable hours. *Gen.: Plýtváme drahocennými hodinami.*
- Ivan: Let's go! *Ivan: Jdeme na to!*
- Gen.: He wants you alive...but it'll be another matter entirely. *Gen.: Chce tě živého... ale to bude něco úplně jiného.*
- Ivan: Don't listen—you have every right!! He might not want us dead, you know. Rusik has done his job, we haven't, which means there is a chance...and there's nothing you can really change about it...just prove that he didn't break the deal and you won't be responsible anyway... for all that Rusik did. *Ivan: Neposlouchejte - máte na to právo!! Možná nás nechce mrtvé, viš. Rusik odvedl svou práci, my ne, což znamená, že je tu šance...a není nic, co byste na tom mohl opravdu změnit...jen dokázat, že neporušil dohodu a vy stejně nebudete zodpovědný...za všechno, co Rusik udělat.*
- Gen.: There was a time, Ivan, when we knew we were safe here – I suppose that was true, except... *Gen.: Byly doby, Ivane, kdy jsme věděli, že jsme tu v bezpečí - to byla asi pravda, až na to...*

- Ivan: But I see Rusik wasn't thinking...and that was the reason you left us alone.
- Gen.: He tried getting us out, Ivan, he said we shouldn't be in danger, he needed help, and then suddenly we became hostages... a little family to each other.
- Ivan: And I was trying desperately afraid for my life, I wouldn't trust anybody.
- Gen.: Rusik gave us an order that you be killed us both, so you wouldn't help yourselves, that's to be taken us all over him. — to prison and it's just give him a chance of his country... as... we wouldn't stand alone. in any chance... so and risk it for his country for Rusik, Rusik could have his life... for him life alone... to have nothing... and nothing ... I just be in prison. just as nothing else... and I have just... for nothing else ... I ... but he just for him. alone no other. no country's own no
- Ivan: Then why were we going?
- Gen.: I thought he would understand... that there were those who knew about Rusik before.
- Ivan: Who?
- Gen.: No country! no other.
- Ivan: Then he can do his thing and save his country... but you?
- Gen.: He can never be trusted, Ivan! We know him, Ivan. He would kill you!
- Ivan: What are you saying? He wouldn't hurt me? No. No one can harm him... it isn't true... it doesn't exist... nothing does exist. — And you don't even know... what it is... or what you could be. It is not real.
- Gen.: It might as well be, Ivan... we've lost too much time already... I hope this won't mean something bad... we've been talking and fighting... just trying desperately hard with every ounce of strength we were given.
- Ivan: Ale vidím, že Rusik nepřemýšlel... a to byl důvod, proč jsi nás nechal na pokoji.*
- Gen.: Snažil se nás dostat ven, Ivane, říkal, že bychom neměli být v nebezpečí, potřeboval pomoc, a pak jsme se najednou stali rukojmími... jedna malá rodina.*
- Ivan: A já se zoufale bál o svůj život, nikomu bych nevěřil.*
- Gen.: Rusik nám dal rozkaz, abyste nás oba zabili, abyste si nepomohli, to znamená, že nás vezmete přes něj. Do vězení a jen mu to dá šanci na jeho zemi... jako... nebudeme stát sami.*
- Ivan: Tak proč jsme jeli?*
- Gen.: Myslel jsem, že to pochopí... že byli tací, kteří o Rusikovi věděli už dříve.*
- Ivan: Kdo?*
- Gen.: Žádná země! Žádná jiná.*
- Ivan: Pak si může dělat, co chce, a zachránit svou zemi... ale ty?*
- Vojenský Generál: Nikdy se mu nedá věřit, Ivane! Známe ho, Ivane. Zabil by tě!*
- Ivan: Co to říkáš? On by mi neublížil? Ne. Nikdo mu nemůže ublížit... to není pravda... to neexistuje... nie neexistuje. A ty ani nevíš... co to je... nebo čím bys mohl být. Není to skutečné.*
- Gen.: Mohlo by to tak být, Ivane... už jsme ztratili příliš mnoho času... doufám, že to nebude znamenat něco špatného... mluvili jsme a bojovali... jen jsme se zoufale snažili s každou špetkou síly, která nám byla dána.*

Ivan: Yes.

Gen.: It doesn't seem fair...but if only it wasn't our country, would you agree? We should all fight...yes...no. Rusik wants me dead and we must stop them doing this to me...he can never have me if he does it. And Rusik wants us safe. we can't stop this, Ivan...and we won't ever prevent them from killing you either... or giving you whatever you want. And you will always have to decide, whether you will die or not... your life... your freedom.

Ivan: Ano.

Gen.: Není to fér, ale kdyby to nebyla naše země, souhlasil byste? Všichni bychom měli bojovat...ano...ne. Rusik mě chce mrtvého a my je musíme zastavit, aby mi tohle nedělali...nikdy mě nemůže mít, když to udělá. A Rusik chce, abychom byli v bezpečí. Tohle nezastavíme, Ivane...a nikdy jim nezabráníme, aby tě taky nezabili... nebo ti dali, co budeš chtít. A vždy se budeš se muset rozhodnout, jestli zemřeš nebo ne... tvůj život... tvá svoboda.

Scene 6

Ivan is in jail. He is awaiting death sentence. He remembers his wife Nina. He misses her very much. He loved her.

/Ivan is sitting in a cell on a narrow bench. He is alone in a dark room. Light is passing only through a small crannie under the door. His hands are bound with cuffs. He is holding a small photograph of Nina./

Ivan: She's a beautiful girl. A very special girl... And I never gave up hope that she'd be able to help me, never let it go.... because... Because she was so sweet and smart. She could read and write! I've had this dream ever since we got back to the house... *Ivan:* You remember the way I used t..... to hold you when we first met? *Ivan:* That was my dream, too... The day you came to live with us... You said that your mother died when you were... 15. You were a very sick girl... So sad and lonely. So very, very sad. You didn't know what you were saying, but you wanted me here for your recovery... Forgive me, my friend... Please forgive me. *Ivan:* But then... you showed me a little something I've been looking forward to forever... I remember when I heard my name. *Ivan:* You made me cry... You told me you would come see me... When you knew that my dreams were real. I know I've hurt you. But you don't have to worry about that. You are still mine... Your husband is still yours... You will always love him. You'll always love him as much now that I'm gone... as I do you... and that makes everything easier... *Ivan:* We'll see each other again. I promise. I'll never leave you. I'll never let you down... I'll never hurt you again. I'll never do anything wrong again... I swear... I just want you all to myself. I promise... I promise. I just...

/Ivan breaks down in tears./

Ivan je ve vězení. Čeká na rozsudek smrti. Myslí na svou ženu Ninu. Moc mu chybí. Miloval ji.

/Ivan sedí ve vazbě na úzké lavici. Je sám v malé tmavé místnosti. Světlo prosvítá jen malou škvírou pod dveřmi. Na rukou má železíčka. Jemně svírá malou fotografii Niny./

Ivan: Je to krásná dívka. Velmi zvláštní dívka... A nikdy jsem se nevzdal naděje, že mi bude schopná pomoci, nikdy to jsem nenechal být... protože... Protože byla tak milá a chytrá. Uměla číst a psát! Mám ten sen od té doby, co jsme se vrátili domů... *Ivan:* Pamatuješ, jak jsem tě držel, když jsme se poprvé potkali? *Ivan:* To byl taky můj sen... Ten den, kdy jsi k nám přišla žít... Říkala jsi, že tvá matka zemřela, když ti bylo... 15. Byla jsi velmi nemocná dívka... Tak smutná a osamělá. Tak velmi, velmi smutná. Nevěděla jsi, co říkáš, ale chtěla jsi mě u sebe kvůli tvému uzdravení... Odpusť mi, můj příteli... Prosím, odpusť mi. *Ivan:* Ale pak... jsi mi ukázala něco, na co jsem se těšil celou věčnost... Vzpomínám si, když jsem tě slyšel říkat mé jméno. *Ivan:* Rozplakala jsi mě... Řekla jsi mi, že za mnou přijdeš... Když jsi věděla, že moje sny jsou skutečné. Vím, že jsem ti ublížil. Ale s tím si nemusíš dělat starosti. Pořád jsi moje... Tvůj manžel je stále tvůj... Vždycky ho budeš milovat. Vždycky ho budeš milovat stejně jako teď, když jsem pryč... jako já... a to všechno usnadňuje... *Ivan:* Ještě se uvidíme. Slibuju. Nikdy tě nepustím. Nikdy tě nezklamau... Už ti nikdy neublížím. Už nikdy neudělám nic špatného... Přísahám... Chci tě mít jen pro sebe. Slibuju... Slibuju. Já jen...

/Ivan se rozpláče./



Ivan is in custody. Rusik comes to him. They are planning to assassinate the dictator.

Ivan je ve vazbě. Rusik přichází k němu. Plánují atentát na diktátora.

/At that moment, the doors open. Rusik enters./

/V ten okamžik se otevřou dveře. Vejde Rusik./

Rusik: Hello Ivan! Beautiful woman... I hope nothing happens to her on the way to New York.

Rusik: Ahoj Ivane! Krásná žena... Doufám, že se jí na cestě do New Yorku nic nestane.

Ivan: Has an official verdict already been issued?

Ivan: Už byl vydán oficiální verdikt?

Rusik: Not yet... But there's no question of your staying in Russia after this... You can go to New York **later** or **anywhere** else you want... but you can't leave Moscow... you have my word... and that means it!

Rusik: Ještě ne... Ale není pochyb o tom, že po tomhle zůstanete v Rusku... Můžeš jet pak do New Yorku nebo kamkoliv chceš... ale zatím nebudeš moct opustit Moskvu... máš moje slovo... a to znamená!

Ivan: I know what it means.

Ivan: Víím, co to znamená.

Rusik: The two Americans who **tried to** murdered our "leader" will die, and so will the people they represent... And that is all there was ever about them – just a couple dozen crazy-ass men who didn't have any sense, who never did anything... and you, Rusik—you were the only one we could count on in Moscow!

Rusik: Dva Američané, kteří se pokusili zavraždit našeho „vůdce“, zemřou, stejně jako lidé, které zastupují... A to je všechno, co na nich kdy bylo – jen pár tuctů šilenců, kteří neměli rozum, kteří nikdy nic neudělali... a ty, Rusiku – ty jsi byl jediný, na koho jsme se mohli v Moskvě spolehnout!



Rusik: We'll see if you have your strength for this one. I know what we're up against... You have my word on that?

Rusik: Uvidíme, jestli na to budeš mít sílu. Víím, proti čemu stojíme... Máte na to mé slovo?

Ivan: Yeah, yeah, it'll take more than a little rocket-propelled charge.

Ivan: Jo, jo, bude to chtít víc než jen malou raketovou nálož.

Rusik: Good luck!

Rusik: Hodně štěstí!

Ivan: It should work!

Ivan: Mělo by to fungovat!

Rusik: No problemo. I've already seen you shoot before... I can help you do it again.

Rusik: No problemo. Už jsem tě viděl střílet... Pomůžu ti to udělat znovu.

Ivan: Oh... you want me back in the field. That would make a change.

Ivan: Oh... chceš mě zpátky do terénu. To by byla změna.

Rusik: Sure... if I had someone I could count on...

Rusik: Jistě... kdybych měl někoho, na koho se můžu spolehnout...

Ivan: Like... your buddy, huh?

Rusik: Oh... he wouldn't do anything rash... he's always got **everything** under control... he never breaks anything... He'd just like **to** say 'Yes'.

/laughter./

Ivan: And then... I guess we're all good to go...

Rusik: ...okay.

Ivan: All right. Good job... Good day, brother...

Rusik: You too, Ivan... You're looking pretty strong.

Ivan: I'll get it out of my system. Now I just need some water... You got a bucket or something?

Rusik: What do you mean?

Ivan: I mean a water pitcher.

Rusik: No... No...

Ivan: Jako... tvůj kámoš, co?

Rusik: Oh... neudělal by nic zbrklého... vždycky má všechno pod kontrolou... nikdy nic nerozbije... Jen by rád řekl „ano“.

/Smích./

Ivan: A pak... Myslím, že jsme všichni připraveni...

Rusik: ...dobře.

Ivan: Dobře. Dobrá práce... Dobrý den, bratře...

Rusik: Ty taky, Ivane... Vypadáš dost silně.

Ivan: Dostanu to ze sebe. Teď potřebuju jen trochu vody... Máš kýbl nebo něco?

Rusik: Co tím myslíš?

Ivan: Myslím džbán na vodu.

Rusik: Ne... Ne...



Ivan: It'll be different now... You've had enough of these speeches, don't make me regret making you my boss... I want nothing from you except respect and admiration for me... And that is what makes me great. And it isn't the price that has kept us apart... That man is dead.

Rusik: No matter how hard he tried, his body would not yield... It was always part of **my** plan. **I hope that no one will ruin our plan. No one will betray us. I hope we'll get back alive and no one will know about us. If I found out we were just infantry, I would refuse to commit murder.**

Ivan: Teď to bude jiné... Už máš těch proslůvků dost, nenut mě litovat, že jsem z tebe udělal svého šéfa... Nechci od tebe nic, kromě úcty a obdivu ke mně... A to je to, co mě dělá skvělým. A není to cena, která nás rozdělila... Ten muž je mrtvý.

Rusik: Ať se snažil sebevíc, jeho tělo se nevzdalo... Byla to vždy součást mého plánu. Snad náš plán nikdo nepokazí. Nikdo nás nezradí. Snad se dostaneme v pořádku živí zpět a nikdo o nás nebude vědět. Kdybych zjistil, že jsme jenom pěšáky, odmítl bych vraždu provést.

Ivan: We both know it. You knew it, too, and neither made a move until today... But when we meet tomorrow at noon in front of Mr. Karpachov... and we see you face to face... and tell ourselves, "You killed the president" - - we'll both be better off than either of us thought we'd become before. You may hate each other forever...but if we both fail, we'll never meet again... because it doesn't happen that way.... And that's the only difference between us. Now I know why I couldn't take over as leader. Because I wanted more from the people... because I wasn't satisfied with what they had given me.

Rusik: And that is the truth, Ivan... And you are going to prove it... By becoming the greatest man I know.... Because that's the only thing worth dying for.... Sooner or later... You're going back into that chair. If we don't change the course right now... Then there will be nothing left to hold on by.... And we won't get through another day. Tomorrow, when the verdicts come out.... I promise you... we're ready for anything... We are waiting till tomorrow. You have to be. Tomorrow morning. It might take years... maybe centuries.

Ivan: Don't you understand, Rusik...? You haven't changed, nor the world, since those days in Havana.... We can still be free.... I know what we're up against... And we are stronger together. So long ago... We lived happily under God, together. This country was our home, Ivan, our family, our dream... Our future.... Your dream is possible.... It doesn't need changing.... You've seen the people. The world changes.... Every day brings new enemies...new threats.... Today we can be free.... Tomorrow night we must change. I am coming here tonight.

Rusik: But we're gonna be freed... If you win tomorrow... we will live happily again.

Ivan: Tomorrow, I swear it.... Tomorrow... Tomorrow.

Ivan: Oba to víme. Tys to taky věděl a ani jeden se do dneška nepohnul... Ale až se zítra v poledne sejdem před panem Karpachovem... a uvidíme se tváří v tvář... a řekneme si: "Zabil jsi prezidenta"- - budeme na tom oba lépe, než jsme si oba mysleli, že jsme se stali. Navždy se můžete nenávidět, ale když oba selžeme, už se nikdy nepotkáme, protože takhle se to nestane. A to je jediný rozdíl mezi námi. Teď už vím, proč jsem to nemohl převzít jako vůdce. Protože jsem od lidí chtěl víc, protože jsem nebyl spokojen s tím, co mi dali.

Rusik: A to je pravda, Ivane... A vy to dokážete... Tím, že se stanu největším mužem, jakého znám... Protože to je jediná věc, pro kterou stojí za to zemřít... Dřív nebo později... Vrátiš se zpátky do křesla. Pokud nezměníme kurz hned teď... Pak už nebude nic, co by se dalo udržet... A další den už nepřežijeme. Zítra, až budou vyneseny rozsudky... Slibuji vám... jsme připraveni na všechno... Čekáme do zítřka. To musíš být. Zítra ráno. Může to trvat roky, možná staletí.

Ivan: Copak to nechápeš, Rusiku...? Ty ses nezměnil, ani svět, od těch dob v Havaně... Stále můžeme být svobodní... Víme, proti čemu stojíme... A společně jsme silnější. Tak dávno... Žili jsme šťastně pod Bohem, společně. Tahle země byla náš domov, Ivane, naše rodina, náš sen... Naše budoucnost... Náš sen je možný... Nemusí se měnit... Viděl jsi lidi. Svět se mění... Každý den přináší nové nepřátele... nové hrozby... Dnes můžeme být svobodní... Zítra večer se musíme převléknout. Přijdu sem dnes večer.

Rusik: Ale my budeme osvobozeni... Když zítra vyhraješ, budeme zase žít šťastně.

Ivan: Zítra, přísahám... Zítra... Zítra.

/Rusik leaves. Ivan is left alone./

/Rusik odejde. Ivan zůstává sám./

Scene 7

Ivan is in prison. He is falsely accused of high treason against the Russian state. He is awaiting sentencing so he can be deported to Siberia.

/Morning. Ivan is in custody. The doors open once again, a prison guard enters. Ivan is nervous but smiling./

Prison guard: So here it is, my friend. You will hang in a few hours.

Ivan: It cannot be true.

/The prison guard bursts out laughing./

Guard: Of course not. They do everything for you... But it's too late now. You are dead already. You'll be sorry you didn't stay there longer. You have nothing left now except for this one little life you will never possess again. You can try and get your passport back to Russia but I guarantee that your life would end up with the other prisoners in there, dead. It won't help your case if they kill you for your life, you understand? They don't want any part at all. You know why? Because they think they can kill me for your life, don't you?

Ivan: You said yourself they'd do anything to make sure that I didn't die in prison. So...

Guard: And they do, they really believe they could get you to go. They'd be willing to take the risk, too, to save your life, my friend.

Ivan: How?

Guard: Because I know the men. I know the ways. They'll take you alive and make sure that no one ever knows about you or what happened to you.

Ivan: And how does the guard **will** you tell them when **that** I am gone?

Ivan je ve vězení. Je falešně obviněn z velezrady proti ruskému státu. Čeká na rozsudek, aby mohl být deportován na Sibiř.

/Ráno. Ivan je ve vazbě. Opět se otevírají dveře, vstupuje vězeňská stráž. Ivan je nervózní, ale usmívá se./

Vězeňská stráž: Tak tady to je, příteli. Za pár hodin budeš viset.

Ivan: To nemůže být pravda.

/Vězeňská stráž se rozchechtá./

Stráž: *Samožřejmě, že ne. Dělají pro tebe všechno... Ale teď už je pozdě. Už teď jsi mrtvý. Budeš litovat, že jste zde nezůstal déle. Už ti nezbylo nic, kromě tohoto malého života, který už nikdy nebudeš mít. Můžete se pokusit získat pas zpět do Ruska, ale garantuji vám, že váš život by skončil s ostatními vězni uvnitř, mrtvými. Vašemu případu nepomůže, když vás zabijí o život, rozumíte? Nechtějí vůbec žádnou roli. Víš proč? Protože si myslí, že mě můžou zabít pro tvůj život, že?*

Ivan: *Sám jsi říkal, že udělají cokoli, aby se ujistili, že nezemřu ve vězení. Takže...*

Stráž: *A oni opravdu věří, že by tě tam mohli dostat. Byli by ochotni riskovat, aby ti zachránili život, můj příteli.*

Ivan: *Jakto?*

Stráž: *Protože ty muže znám. Znáš ty způsoby. Dostanou tě živého a postarají se, aby se o tobě nikdo nedozvěděl, nebo co se ti stalo.*

Ivan: *A jak jim strážný řekneš, že jsem pryč?*

Guard: When I leave you behind. I leave in five minutes and then we're leaving. And if they see me going, they will come after me too, they just need to see you, they can't be far.

Ivan: And how do they get you out of prison?

Guard: ~~The same way I get out.~~ They shoot me and you go free. If they find me taking a shower or reading, I get shot and killed.

Ivan: That's the only way. They would shoot you, ~~wouldn't they?~~

Guard: Yes, sir. ~~They have their own weapons, sir.~~

Ivan: You'll be fine, Nina. I promise. They can't harm you now. You are safe now. Just keep walking, walk as fast as possible, keep moving, don't stop for any water or food or even anything that could possibly give you away. Don't look at any of these guards or they might see you and kill you, or worse. Understand? Keep trying, keep breathing! Do whatever it takes! Don't let anyone stop you.

Stráž: Až tě tu nechám. Za pět minut odcházím a pak odcházíme. A když mě uvidí odcházet, půjdou po mně taky, jen tě potřebují vidět, ~~nemůžou být daleko.~~

*Ivan: A jak **mě** dostaneš z vězení?*

Stráž: ~~Stejně jako se dostanu ven.~~ Zastřelí mě a ty budeš volný. Když mě najdou ve sprše nebo při čtení, zastřelí mě.

Ivan: To je jediná možnost. Zastřelí tě, ~~že jo?~~

Stráž: Ano, pane. ~~Mají vlastní zbraně, pane.~~

Ivan: Budeš v pořádku, Nino. Slibuju. Ted' už ti neublíží. Už jsi v bezpečí. Jen jdi, jdi tak rychle, jak to jen jde, nezastavuj se pro vodu, jídlo nebo cokoli, co by tě mohlo prozradit. Nedívej se na žádného z těch dozorců, nebo tě uvidí a zabijí tě, nebo huř. Rozumíš? Snaž se, dýchej! Udělejte, co je třeba! ~~Ne~~ nech se nikým zastavit.

Scene 8

On the platform, Ivan is waiting for Rusik. The train is leaving.

Military General: Where is Rusik? He has all the equipment.

Ivan: I'm almost sure he won't show up. Fucking rat - he must have betrayed us.

Gen.: Then you've lost your only ally.

Ivan: No.

Gen.: He will not leave here alive! You know this! You promised me! If **he manages to escape** I do not give you a chance of escaping - I'll **personally** fire that bullet into his brain.

Ivan: That's it, my friend... We are doomed! We die. Now we just need time. We're surrounded by spies... We need him dead!

Gen.: We'll do nothing. We don't know how much longer they can hold out against us without help... They have our airfields...

Ivan: And the city! It is a small town, General... But they control everything in there... I could go anywhere, do anything, do what ever I wanted, just as long as no one knew who was in charge!

Gen.: We can take them at their own weapons!!

Ivan: There are two million people living within these borders! There is nobody who does not want revenge for this. If I could only find someone to do something about it!!!

Gen.: Who would it be??

Ivan: **Someone from** The Russian Army?! I'll do it. I promise. I'm doing the right thing... Just tell me who. Tell me, please... Who'd you want?

Na nástupišti čeká Ivan na Rusika. Vlak odjíždí.

Vojenský generál: Kde je Rusik? Má všechno vybavení.

Ivan: Jsem si skoro jistý, že se neukáže. Zasraná krysa - musel nás zradit.

Gen.: Pak jsi ztratil svého jediného spojence.

Ivan: Ne.

Gen.: On odsud neodejde živý! To přece víš! Slíbil jsi mi to! Když se mu podaří utéct, vystřelím mu tu kulku do mozku osobně.

Ivan: To je ono, příteli... Jsme ztraceni! Zemřeme. Teď už jen potřebujeme čas. Jsme obklopeni špióny... Potřebujeme ho mrtvého!

Gen.: Nebudeme dělat nic. Nevíme, jak dlouho proti nám ještě vydrží bez pomoci... Mají naše letiště...

Ivan: A město! Je to malé město, generále... Ale oni tam ovládají všechno... Mohl jsem jít kamkoliv, dělat cokoli, dělat, co jsem chtěl, jen když nikdo nevěděl, kdo tomu velí!

Gen.: Můžeme je dostat jejich vlastními zbraněmi!!

Ivan: V těchto hranicích žijí dva miliony lidí! Neexistuje nikdo, kdo by se za to nechtěl pomstít. Kdybych tak našel někoho, kdo s tím něco udělá!!!

Gen.: Kdo by to byl?

Ivan: Někdo z Ruské armády?! Já to udělám. Slibuju. Dělam správnou věc... Jen mi řekni kdo. Řekni mi, prosím... Koho jsi chtěl?

Gen.:	You.	Gen.:	<i>Tebe.</i>
Ivan:	Oh shit, Ivan Karpov, come to save the world! I can feel the world falling apart around me...	Ivan:	<i>Do prdele, Ivan Karpov, přišel zachránit svět! Cítím, jak se svět kolem mě rozpadá...</i>
↓			
Gen.:	We have a problem. The If Rusik betrayed us. They'll know about the whole event. It's too risky, you should not get off the train.	Gen.:	<i>Pokud nás Rusik zradil. Vědí o celé té události. Je to příliš riskantní, nevystupuj z vlaku.</i>
Ivan:	And end up in Siberia. No way - that's my only chance. I will kill dictator any way.	Ivan:	<i>A skončit na Sibiři. V žádném případě - to je moje jediná šance. Zabiju diktátora jakýmkoliv způsobem.</i>
Gen.:	I want your word you will go on board this train when we get there and help us with our plans. You will have no choice but...	Gen.:	<i>Chci vaše slovo, že nastoupíte do tohoto vlaku, až tam dorazíte, pomůžete nám s našimi plány. Nebudete mít jinou možnost, než...</i>
Ivan:	No choice! We don't give a damn if the whole world knows. You have given us everything. You have brought all of Russia back into our family, you understand? You can take the world and make it safe. We don't need Rusik! We know who has the guns! You think they'll be interested?! Well I bet you 'bout as much they won't believe you when you say we killed their leader, then they gonna start believing in themselves! That man Rusik had more to lose than me, didn't he? I will prove them wrong! Now I gotta find some people and talk 'em out. I mean I gotta figure out what we do here and where to plant those mines	Ivan:	<i>Nemáme na výběr! Je mi úplně jedno, jestli to ví celý svět. Dal jsi nám všechno. Přivedl jsi celé Rusko zpět do naší rodiny, rozumíš? Můžu vzít svět a udělat ho bezpečným. Rusika nepotřebujeme! Víme, kdo má zbraně! Myslíš, že je to bude zajímat?! Vsadím se s tebou, že ti uvěří, když řekneš, že jsme zabili jejich vůdce, pak začnou věřit sami sobě! Rusik mohl ztratit víc než já, že? Dokážu jim, že se mýlí! Ted' musím najít pár lidí a promluvit si s nimi. Chci říct, že musím přijít na to, co tady děláme a kam umístit ty míny</i>
Gen.:	What do you intend doing now?	Gen.:	<i>Co teď hodláte dělat?</i>
Ivan:	This was my plan, my whole plan! We must work together! Work like two great teams to protect the people and put an end tew treason	Ivan:	<i>Tohle byl můj plán, celý můj plán! Musíme spolupracovat! Pracujte jako dva skvělé týmy, které budou chránit lidi a ukončit zradu.</i>
Gen.:	If we destroy him we could still win	Gen.:	<i>Když ho zničíme, pořád můžeme vyhrát.</i>
Ivan:	What? We have lost. He got away from us, that shows his guilt, it shows us what he wants, that showed him he can hide things, so let him stay	Ivan:	<i>Co? Prohráli jsme. Utekl nám, to ukazuje jeho vinu, to nám ukazuje, co chce, to mu ukázalo, že umí skrývat věci, tak ať zůstane.</i>

Gen.:	The plan works best without a leader	Gen.:	<i>Plán funguje nejlépe bez vůdce.</i>
Ivan:	Yes it does but if he comes to power and everyone else thinks we are working together to betray the state... Then I will have to use him as a tool	Ivan:	<i>To ano, ale pokud se dostane k moci a všichni ostatní si budou myslet, že spolupracujeme, abychom zradili stát... Pak ho budu muset použít jako nástroj.</i>
Gen.:	If the Russians suspect that we have worked against each other we may lose our advantage	Gen.:	<i>Jestli Rusové budou mít podezření, že jsme pracovali proti sobě, můžeme ztratit výhodu.</i>
Ivan:	He said it was my only chance. So it will work! He wanted revenge. We must trust each others judgement! I promise to protect the country The train is leaving. I must go.	Ivan:	<i>Říkal, že je to moje jediná šance. Takže to bude fungovat! Chtěl se pomstít. Musíme věřit úsudku toho druhého! Slibuji, že budu chránit zemi. Vlak odjíždí. Musím jít.</i>
↓			
Gen.:	One more thing. We received a horrible message. The Boat where your wife Nina was, shipwrecked. I'm very sorry. Your wife died. Russian cruiser killed her.	Gen.:	<i>Ještě jedna věc. Dostali jsme zprávu, že loď, kde byla vaše žena Nina, ztroskotala. Promiňte - vaše žena zemřela. Ruský křižník ji zabil.</i>
Ivan:	Oh no...nobody survived? My Nina is dead?	Ivan:	<i>Ale ne...nikdo nepřežil? Moje Nina je mrtvá?</i>
Gen.:	I know she's not alive...and she is not coming back.	Gen.:	<i>Vím, že není naživu...a nevrátí se ke mně...ne. Musíš zemřít...pro tohle...tohle je milosrdenství...ale nedá se nic dělat.</i>
Ivan:	It will help you, General! I swear it. I'll die with her!	Ivan:	<i>Pomůže vám to, generále! To přísahám. Zemřu s ní!</i>
Gen.:	No! I cannot take you out. I won't do that! I am so very afraid...of the idea of you dying!	Gen.:	<i>Ne! Nemůžu tě vzít ven. To neudělám! Hrozně se bojím... představy, že zemřeš!</i>
Ivan:	Why should you be afraid?! It doesn't matter if I'm killed or wounded. It doesn't make any difference...the day is almost over... I am going home!	Ivan:	<i>Proč by ses bál?! Nezáleží na tom, jestli jsem zabit nebo zraněn. Na tom nezáleží... den je skoro u konce... já jdu domů!</i>
Gen.:	You can't leave now. I can see it in your eyes...	Gen.:	<i>Teď nemůžeš odejít. Vidím ti to na očích...</i>

Ivan: ~~You will kill me. You are trying hard enough. You know it! I have nothing left for you! I will be with my Nina... in God's world. I swear to you I will go on loving her and waiting... to hear from her. I swear it... you have tried too much already... so don't waste more time! You can promise me this! You have to!—I swear to you...! **I will avenge her.**~~

*Ivan: ~~Ty mě zabiješ. Snažíš se dost. Ty to víš! Nic na tebe nemám! Budu se svou Ninou... v Božím světě. Přísahám, že ji budu dál milovat a čekat... až se ozve. Zůstanu s ní v nebi, kde se mi splní všechny sny. Slibuju! Nikdy tě neopustím! **Musím ji pomstít.**~~*



Gen.: That is all?

Gen.: ~~To je všechno?~~

Ivan: ~~It will take you some time... But the Russians won't give up yet! They know you're still in their confidence and they want revenge! They know that the United States can never beat them now! You understand that?!... You are our leader and I want my people happy... Now let's go home... Let us make our country safe!~~

Ivan: ~~Zabere vám to nějaký čas... Ale Rusové se ještě nevzdávají! Vědí, že jsi stále v jejich důvěře a chtějí se pomstít! Vědí, že Spojené státy je teď nikdy neporazí! Rozumíš tomu?!... Jsi náš vůdce a já chci, aby můj lid byl šťastný... Teď pojďme domů... Učinně naši zemi bezpečnou!~~

Gen.: Yes. I understand.

Gen.: ~~Ano. Rozumím.~~

Ivan: ~~Thank God for my freedom! I can be very strong with strength in me... You can not do that... I cannot let you hurt them again. You know what it feels like...~~

Ivan: ~~Díky bohu za mou svobodu! Dokážu být velmi silný, když mám v sobě sílu... To nemůžete udělat... Nedomolím, abys mému lidu znovu ublížil. Víš, jaké to je...~~

Gen.: ~~No, I don't know... Maybe you will come back when I have my men around you. You know I will be watching over you... You must trust me. You have my word... You are a great leader, Ivan... You can trust me.~~

Gen.: ~~Ne, já nevím... Možná se vrátíte se, až kolem vás budu mít své muže. Víš, že na tebe budu dávat pozor... Musíš mi věřit. Máte mé slovo... Jsi skvělý vůdce, Ivane... Můžeš mi věřit.~~

Ivan: ~~You must come back to America... Please. You are an American, and we need your help... You understand... You will be watched over. I have made my promise to the people... And now... You are coming home with me... Come home.~~

Ivan: ~~Musíš se vrátit do Ameriky... Prosím. Jste Američan a potřebujeme vaši pomoc... Rozumíte... Budou na vás dohlížet. Dal jsem lidem svůj slib... A teď... Půjdeš se mnou domů... Pojď domů.~~

Gen.: Ivan! You can trust me! I will always watch over you! I swear on this promise, you are my leader. You have been trusted by my men! You must trust me. I am so glad to see you. I wish I could say more, but you have already told me everything... Everything about you. I understand...

Gen.: ~~Ivane! Můžeš mi věřit! Vždycky na tebe budu dávat pozor! Přísahám na tento slib, že jsi můj vůdce. Moji muži vám věřili! Musíš mi věřit.~~

/The train is leaving. Ivan jumps on./

/Vlak se rozjíždí, Ivan naskakuje./

Scene 9

Ivan arrived in Moscow. He walks through a dirty dark neighborhood. An arms dealer is standing on one of the abandoned corners.

Arms dealer: Ivan Karpov?

Ivan: That's me.

Dealer: What can I do for you?

Ivan: Nothing much.

Ivan přijel do Moskvy. Prochází špinavou temnou čtvrtí. Na jednom z opuštěných rohů stojí obchodník se zbraněmi.

Obchodník se zbraněmi: Ivan Karpov?

Ivan: To jsem já.

Obch.: Co pro vás mohu udělat?

Ivan: Nic moc.



Dealer: Shhh. Someone could hear us. Come closer mr. Karpov. What kind of weapon do you want.

Ivan: Sniper weapon.

Dealer: We have nothing like that anymore. It was only invented five years ago and sold very rarely by some dealers who sell stolen guns all over Europe. I think you should try something new. I will show it **to you**. It's called "Sentry". A highly accurate, low recoil sniper rifle with high muzzle velocity which makes him quite useful when needed... **but** especially against bad guys... like you.

Ivan: Who are you?

Dealer: I am an agent from the FSB, you can call me Vronsky. You want this rifle? Tell me what I'm asking!

Ivan: How much is this worth?!

Dealer: Ten million rubles, mr. Karpov. I've been doing business **for a** long time now and we'll make it as good as you say! If there are any questions you're gonna be happy just give them here, OK?! And thank you!

Obch.: **Shhh. Někdo by nás mohl slyšet. Přistupte blíže, pane. Karpove. Jakou chceš zbraň.**

Ivan: Odstřelovačskou pušku.

Obch.: *Nic takového už nemáme. Byla vynalezena teprve před pěti lety a velmi vzácně prodávána některými dealery, kteří prodávají kradené zbraně po celé Evropě. Myslím, že bys měl zkusit něco nového. Ukážu ti to ~~to~~ ~~tvém~~. Jmenuje se "Sentry". Výsoce přesná odstřelovací puška s nízkým zpětným rázem, s vysokou hlavní rychlostí, díky které je docela užitečná, když je potřeba... ~~ale~~ hlavně proti padouchům... jako jste vy.*

Ivan: Kdo jste?

Obch.: *Jsem agent z FSB, můžete mi říkat Vronský. Chceš tu pušku? Řekni mi, na co se ptám!*

Ivan: *Jakou to má cenu?!*

Obch.: *Deset milionů rublů, pane. Karpov. Už dělám byznys dlouho a uděláme to tak dobře jak říkáte! Pokud jsou nějaké otázky, které budete mít rádi, tak je sem dejte, OK?! A děkuji!*



Dealer: How much have you got there?

Obch.: *Kolik tam máš?*

Ivan: Twenty million rubles... I thought maybe you might be interested.

Ivan: Dvacet milionů rublů... Myslel jsem, že by vás to mohlo zajímat.

Dealer: What for?

Obch.: Proč?

Ivan: I know this guy, Rusik, he's going after your wife... I was thinking we could arrange it... We need to find him, now!

Ivan: Zním jednoho chlápka, Rusika, jde po vaší ženě... Myslel jsem, že bychom to mohli zařídit... Musíme ho najít, hned!

Dealer: Where would we find him?

Obch.: Kde ho najdeme?

Ivan: There...in a place called the Gaborka... You can meet us at the place when you get there... I'll bring my money with me.

Ivan: Tam...na místě zvaném Gaborka... Můžete se s námi sejít na místě, až tam dorazíte... A vezmi s sebou moje peníze.

Dealer: This isn't very smart, what do you want?

Obch.: To není moc chytré, co chceš?

Ivan: The same thing as everybody else who works for the Russian government - to help us track this guy down... I've never been to any kinda operation, not like this, so... why don't you meet us here tonight...

Ivan: To samé, co všichni ostatní, kteří pracují pro ruskou vládu - aby nám pomohli toho chlapa vystopovat... Nikdy jsem nebyl na žádné operaci, takhle ne, takže... proč se tu s námi nesetkáš dnes večer...

Dealer: It doesn't look like a very good idea.

Obch.: To nevypadá jako moc dobrý nápad.

Ivan: You see...if you start asking questions... they won't put Rusik under arrest...and you could end up getting shot. You should know better than anyone, Arsonist...that a man in that condition needs protection...so why don't you leave?

Ivan: Víte...když se začnete vyptávat... dají Rusika do vazby... a vy můžete skončit postřelený. Měl bys vědět lépe než kdokoliv jiný, Žháří...že muž v takovém stavu potřebuje ochranu...tak proč neodejdeš?



Ivan: Evil must be eradicated from the root.

Ivan: Zlo musí být vymýceno od kořene.

Dealer: We are just doing business, Mr. Karpov! You are selling your country a terrible trade... but this time it's real good. The President and his family will pay whatever it takes. Now go.

Obch.: Jen obchodujeme, pane Karpove! Prodáváte své zemi příšerný obchod... ale tentokrát je opravdu dobrý. Prezident a jeho rodina zaplatí, co bude třeba. Teď jdi.

Ivan: The president is coming back here tomorrow night at 6 am sharp. I want to kill him before he can return to Russia.

Ivan: Prezident se sem vrací zítra večer přesně v šest. Chci ho zabít, než se vrátí do Ruska.

Dealer: He already left tonight. We're waiting now.

Obch.: Dnes večer už odjel. Teď čekáme.

Ivan: The assassin must die!

Ivan: Vrah musí zemřít!

- Dealer: We will give you all that you desire – if you agree we have a deal.
- Ivan: All right then, Mr.... You got something else for me?
- Dealer: You can kill anyone who crosses my line!
- Ivan: You want the president dead?! Then let's make this a deal.
- Dealer: There are two ways of killing someone: you either take them out or they kill themselves. I have never killed anybody, but you are not the one going to get into trouble, so let's talk about the first method.
- Ivan: The second is easier. You tell them the president was attacked. They will think it was an accident and help cover up for you.
- Dealer: We will send word that there were armed men following the president to his house at 9 am. We will also tell people in New York, Chicago... London, Vienna that you plan on carrying out this operation in full force until he returns home safely. This is what you need to know about me and our deals. I have never made a mistake in over 20 years with any American company. I promise to treat you exactly as we treat every other businessman who enters this room. I'll find out everything about your clients.
- Ivan: What's this supposed to mean?
- Dealer: That I won't lie for you. I don't have the slightest intention of taking on anyone who does. I'll make sure everyone knows where I stand. I've got the world on my shoulders, and it doesn't take a genius like yourself anymore to figure that out.
- Ivan: If I don't see this through, it may end up being a very long night for me.
- Obch.:* *Dáme vám vše, po čem toužíte - pokud budete souhlasit, máme dohodu.*
- Ivan:* *Dobrá tedy, pane... Máš pro mě ještě něco?*
- Obch.:* *Můžeš zabít každého, kdo překročí mou hranici!*
- Ivan:* *Chcete prezidenta mrtvého?! Tak se na tom dohodneme.*
- Obch.:* *Jsou dva způsoby, jak někoho zabít: buď ho sejmete, nebo se zabije sám. Nikdy jsem nikoho nezabil, ale ty nejsi ten, kdo se dostane do problémů, takže si promluvíme o první metodě.*
- Ivan:* *Vím, že druhá je jednodušší. Řekněte jim, že prezident byl napaden. Budou si myslet, že to byla nehoda a pomůžou ti to utulat.*
- Obch.:* *Pošleme zprávu, že v devět ráno šli ozbrojení muži za prezidentem do jeho domu. Řekneme to také lidem v New Yorku, Chicagu... Londýně, Vídni, že plánujete provést tuto operaci v plné síle, dokud se bezpečně nevrátí domů. Tohle musíte vědět o mně a našich dohodách. Za 20 let u žádné americké společnosti jsem nikdy neudělal chybu. Slibuji, že s vámi budu zacházet stejně jako s každým jiným podnikatelem, který vstoupí do těchto míst. Zjistím všechno o vašich klientech.*
- Ivan:* *Co to má znamenat?*
- Obch.:* *Že kvůli tobě nebudu lhát. Nemám ani v nejmenším v úmyslu pustit se do někoho, kdo to udělá. Postarám se, aby všichni věděli, jak na tom jsem. Mám na svých bedrech celý svět a na to už není třeba takového génia, jako jsi ty, aby na to přišel.*
- Ivan:* *Jestli to nedotáhnu do konce, může to pro mě být velmi dlouhá noc.*

Dealer: That's okay. We're getting you a good deal, too. You will never see your President again after tonight unless I decide to drop a bomb on your head. I hope I'm wrong, Mr. Karpov, but you have been chosen to carry out a great mission in history. You'll have a better future than most ordinary guys who spend their whole lives running around shooting each other in countries where there isn't even any fighting. I don't know what the hell I did to deserve this honor. I really don't care what I thought. It has given me an incredible gift, sir, which makes me proud of what I represent. And if there is anything I've learned in the last ten minutes... Well, it's that a man doesn't just have to get along, you gotta keep on pushing the button. And it is not going away, sir. So you should probably take some deep breaths and let yourself relax. You have earned it. You will see me next week, sir. I will come by to visit when you leave and show you all that I'm offering you. Good-bye, sir.

Ivan: Goodbye.

Dealer: Goodbye.

Obch.: *To je v pořádku. Taky vám nabízíme dobrou cenu. Po dnešní noci už svého prezidenta nevidíte, pokud se nerozhodnu hodit vám na hlavu bombu. Doufám, že se mýlím, pane Karpove, ale byl jste vybrán, abyste splnil velký úkol v historii. Budete mít lepší budoucnost než většina obyčejných chlapů, kteří celý život pobíhají a střílejí po sobě v zemích, kde se ani nebojuje. Nevím, čím jsem si sakra zasloužil tuhle poctu. Je mi jedno, co jsem si myslel. Dalo mi to neuvěřitelný dar, pane, díky kterému jsem hrdý na to, co reprezentuji. A jestli jsem se za posledních deset minut něco naučil... No, jde o to, že muži nemusí jen spolu vycházet, ale musíš pořád mačkat tlačítko. A to nezmizí, pane. Takže by ses asi měl zhluboka nadechnout a uvolnit se. Zasloužíš sis to. Uvidíme se příští týden, pane. Až odejdete, přijdu vás navštívit a ukážu vám všechno, co vám nabízím. Sbohem, pane.*

Ivan: *Sbohem.*

Obch.: *Sbohem.*

Scene 10

In the hotel room. Ivan is preparing to assassinate the president. Rusik enters the room. Rusik wants to stop Ivan.

Ivan: What are you doing here?

Rusik: I'm here to stop you.

Ivan: No! Get out of my way. You are no good to me, I hate your company. Get away from there! I have something important to say... You think that this will stop us now? It won't because I am right, we cannot be friends forever. I don't like what happened with Kudrin but if I was in your place... I would make a plan... I promise.

Rusik: Then come along with me. Come on, you're the only one who can stop it. We have been planning for a long time... ~~We've made a deal with Kradic, we can end it tonight. But not with you... We must take out someone else.~~

Ivan: No, Rusik.

Rusik: **Military general lied to us. You will die after assassination...**

Ivan: **Coward! You betrayed us. I don't believe you.**

Rusik: Then I guess **Still** you can come back to my side... ~~You're still in a prison, you idiot...~~

Ivan: No...

Rusik: Yes. You will come back to my side and live a life of honourable repose...

Ivan: I'm already dead... I didn't have a choice. I am so sorry. Goodbye... Goodbye.

Rusik: Don't you see, it's about your life, don't you?

V hotelovém pokoji. Ivan se chystá zavraždit prezidenta. Do místnosti vstoupí Rusik. Rusik chce zastavit Ivana.

Ivan: *Co tady děláš?*

Rusik: *Jsem tu, abych tě zastavil.*

Ivan: *Ne! Jděte mi z cesty. Nejsi mi k ničemu, nesnáším tvou společnost. Jdi od toho! Musím ti říct něco důležitého... Myslíš, že nás to teď zastaví? Nebude, protože mám pravdu, nemůžeme být přáteli navždy. Nelíbí se mi, co se stalo s Kudrinem, ale kdybych byl na tvém místě... Udělal bych plán... Slibuju.*

Rusik: *Tak pojď se mnou. No tak, jsi jediný, kdo to může zastavit. Plánovali jsme to už dlouho... Dohodli jsme se s Kradičem, můžeme to ukončit dnes večer. Ale ne s tebou... Musíme sejmout někoho jiného.*

Ivan: *Ne, Rusiku.*



Rusik: *Vojenský generál nám lhal. Zemřete po atentátu...*

Ivan: *Zbabělče! Zradil jsi nás. Já ti nevěřím.*

Rusik: *Pak se asi můžeš se stále vrátit na mou stranu... Pořád jsi ve vězení, ty idiote..*

Ivan: *Ne...*

Rusik: *Ano. Vrátíš se na mou stranu a budeš žít čestný život v klidu...*

Ivan: *Už jsem mrtvý... Neměl jsem na výběr. Moc se omlouvám. Sbohem... Sbohem.*

Rusik: *Nevidíš, že ti jde o život, že?*

Ivan:	You want me alive... then you should let him go...	Ivan:	Chceš mě živého... tak bys ho měl nechat jít...
Rusik:	We all want you to be with us, Ivan...	Rusik:	Všichni chceme, abys byl s námi, Ivane...
Ivan:	Then what is it we do...?	Ivan:	Tak co budeme dělat...?
Rusik:	The army wants to bring him to us, that's why they brought you here.	Rusik:	Armáda ho chce přivést k nám, proto vás sem přivedli.
Ivan:	They must know what you've done. If I'm to stay in this world... then there are things you cannot tell them... What did they give you? I gave it all.	Ivan:	Must vědět, co jsi udělat. Pokud mám zůstat na tomto světě... pak jsou věci, které jim nemůžeš říct... Co ti dali? Dal jsem do toho všechno.
↓			
Ivan:	But if we will kill the tyrant... you go with me.	Ivan:	Ale když zabijeme tyrana, půjdeš se mnou.
Rusik:	You want revenge.	Rusik:	Chceš pomstu.
Ivan:	It's better to say revenge... is better than murder...	Ivan:	Lepší je říkat pomsta... je lepší než vražda...
Rusik:	Yes it's a question whether you succeed...but the decision is yours alone... You have to make the right choice...or...the consequences...are too great...to live without.	Rusik:	Ano, je otázka, zda uspěješ...ale rozhodnutí je jen na vás... Musíte se rozhodnout správně...nebo...následky...jsou příliš velké...než aby se bez nich dalo žít.
Ivan:	But what are my choices? What is my destiny...my life...do you give me the choice to choose death over living? If there was nothing else, I would do this.	Ivan:	Ale jaké mám možnosti? Jaký je můj osud...můj život...dáváš mi možnost vybrat si smrt před životem? Kdyby nic jiného nebylo, udělal bych tohle.
Rusik:	No, it has to be a life that doesn't involve dying...because then there would never be another day. If it were so easy...if you could leave all that behind. Ivan.	Rusik:	Ne, musí to být život, který nezahrnuje umírání... protože pak by už nikdy nebyl další den. Kdyby to bylo tak snadné... kdybys to všechno mohl nechat za sebou. Ivane.
Ivan:	Then you are an idiot...	Ivan:	Pak jsi idiot...
Rusik:	Yes. You have failed me...for now...there's no way to save your own life...it has to be my duty. But there are others who think that the world will end in two days...that mankind will fall...and we'll be forgotten forever...before the end...after all. So if I had the choice...then yes.	Rusik:	Ano. Zklamal jsi mě...prozatím...není způsob, jak si zachránit život...musí to být má povinnost. Ale jsou jiní, kteří si myslí, že svět za dva dny skončí...že lidstvo padne...a my budeme navždy zapomenuti...před koncem...po tom všem. Takže kdybych měl na výběr...tak ano.

Ivan: ~~And you have made it? And now... will you do something about it?!~~

Ivan: ~~A zvládli jste to? A teď... uděláš s tím něco?!~~

Rusik: It is your life, Ivan...and your choice...make yours or die by it...choose it.

Rusik: Je to tvůj život, Ivane...a tvá volba...buď se rozhodneš, nebo zemřeš...vyber si.



Rusik: They don't understand... they never understood you.

Rusik: ~~Nechápu... nikdy nechápali tebe.~~

Ivan: Then let them take you to our army, I want to go there now. They will teach you what you are doing right...

Ivan: ~~Tak ať tě vezmou do naší armády, chci tam jít hned. Naučí vás, co děláte správně...~~

Rusik: They'll send an assassin after you, but it won't be only as far as his arm.

Rusik: ~~Pošlou na tebe zabijáka, ale bude to jen do ruky.~~

Ivan: But how are you going home? I have to be with her... How can I leave? There is no other way.

Ivan: ~~Ale jak se dostaneš domů? Musím být s ní... Jak můžu odejít? Není jiná možnost.~~

Rusik: Ivan... Ivan, listen...

Rusik: ~~Ivane... Ivane, poslouvej...~~

Ivan: No, Rusik.

Ivan: ~~Ne, Rusiku.~~

Rusik: Please, Ivan.

Rusik: ~~Prosím, Ivane.~~

Ivan: It has nothing at stake. If she were safe in a cage, or dead... I could have her, yes, Rusik, but if I could not?

Ivan: ~~V sázce není nic. Kdyby byla v bezpečí v kleci, nebo mrtvá... Mohl bych ji mít, ano, Rusiku, ale kdybych nemohl?~~

Rusik: ~~Then it would have been easier... Ivan... It was wrong.~~

Rusik: ~~Pak by to bylo jednodušší... Ivane... Bylo to špatné.~~

Ivan: Yes...yes...

Ivan: ~~Ano...ano...~~

Rusik: Ivan, you don't need these...

Rusik: ~~Ivane, tohle nepotřebuješ...~~

Ivan: Why do you care? It isn't about what they do.

Ivan: ~~Proč se staráš? Nejde o to, co dělají.~~

Rusik: Ivan... Don't... You...

Rusik: ~~Ivane... Ne... Ty...~~

Ivan: What does it matter to me...

Ivan: ~~Co na tom záleží...~~

Rusik: Don't talk like that.

Rusik: ~~Takhle nemluv.~~

Ivan: She...is dead. It doesn't matters because... because...

Ivan: ~~Ona... je mrtvá. Nezáleží na tom, protože... protože~~



Rusik: What did you said? Nina is alive. That told you Military general? It is a prove that he is a liar!

Ivan: What? My Nina is alive?!

Rusik: It was all a lie! You must tell me your real story. I can help you!



Ivan: What? My Nina is alive?!

Rusik: Of course she's ali-

Ivan: Who did it? Who was responsible??

Rusik: It... It's no business..!

Ivan: Then what is this? ~~What have you done? Why didn't I hear?~~

Rusik: Ivan, it is too late!! You can die! You cannot live!!

Ivan: ~~But~~ my Nina!!! How could I leave her like this!! No matter what she does or doesn't do...I...can't! I am so sorry. I wish there were something I could say to you...something I would let you know...that...there was a chance...that you might...you know...just try...and please, please just try and be okay with this.

Ivan: What is this about, Rusik!? Tell it me!!

Rusik: General lies.

Ivan: He killed her? Really??!! No... No!!!

Rusik: ~~That~~ it was a lie. We both know, what he did.

Ivan: It cannot be true! No!! I don't believe you! **As for Nina, he was just wrong.**

Rusik: Co jsi to řekl? Nina je naživu. To vám řekl Vojenský generál? To je důkaz, že lže!

Ivan: Co? Moje Nina žije?!

Rusik: Všechno to byla lež! Musíš mi říct svůj skutečný příběh. Můžu ti pomoci!

Ivan: Co? Moje Nina žije?!

Rusik: Samozřejmě, že je na ž-...

Ivan: Kdo to udělal? Kdo za to může?

Rusik: To... To není obchod..!

Ivan: Tak co je tohle? ~~Co jsi to udělal? Proč jsem to neslyšel?~~

Rusik: Ivane, už je pozdě!! Můžeš zemřít! Nemůžeš žít!!

Ivan: ~~Ale~~ moje Nina!!! Jak jsem ji mohl takhle opustit! Nezáleží na tom, co dělá nebo nedělá... Já... nemůžu! Moc se omlouvám. Přál bych si, abych ti mohl něco říct...něco, co bych ti dal vědět...že...je tu šance...že bys mohl...víš...jen to zkusit...a prosím, prosím, jen to zkus a buď v pohodě.

Ivan: O co jde, Rusiku!? Řekni mi to!!

Rusik: **Generál Lže.**

Ivan: On ji zabil? Opravdu??!! Ne...Ne!!!

Rusik: ~~Pak to byla lež...~~ A oba víme, co udělal!

Ivan: To nemůže být pravda!! Ne!! Já ti nevěřím! Co se týče Niny, jenom se spletl.



/Doors open up and Nina comes in./

Nina: Hello Ivan! Glad to see you again. What is going on.

/Otevřou se dveře a vejde Nina./

Nina: Ivane! Rád tě zase vidím. Co se děje?

Ivan: You are alive. What are you doing here?

Nina: I wanted so badly not..

Ivan: Nina, we were both very scared when we came here. And I knew it would be bad for you if you ever found out. I should go now. I'll meet you in your car at my place tomorrow night. It's far away. Good bye.

Rusik: You won't live to see tomorrow night.

Ivan: Don't lie! I believe the General take care of my safety.

Nina: I can't do this Ivan. You're supposed too watch over me. This was my home and now you have no right to take everything away. Goodbye Ivan.

Ivan: Nina. Don't come back here until I say goodbye. Don't leave until I say goodbye. Don't you dare try anything until I give you one chance to stay here. Goodbye Nina.

Nina: Wait a minute!

Ivan: Nina.

Nina: Don't! You mustn't. You promised. I love you Ivan. You promise me. Don't go till I say good-bye.

Ivan: You said good-by. And now you don't have to go either. You won't come back either. Don't worry. Everything is fine.

Nina: Please Ivan!

Ivan: Don't listen anymore!

Nina: Why did you leave? What's happening to me? You said that everything had been okay.

Ivan: It's never alright with you! Everything's always wrong! What is this, a joke?

Ivan: Ty žiješ! Co tady děláš?

Nina: Tak moc jsem chtěla...

Ivan: Nino, oba jsme byli velmi vystrašení, když jsme sem přišli. a Já jsem věděl, že by to pro tebe bylo špatné, kdybys to někdy zjistila. Měli bychom jít. Sejdeme se zítra večer u mě v autě. Je to daleko. Sbohem.

Rusik: Zítřejšího večera se nedožiješ.

Ivan: Nelži! Věřím, že se Generál postará o mou bezpečnost.

Nina: Nemůžu to udělat Ivane. Máš na mě dávat pozor. Tohle byl můj domov a teď nemáš právo mi všechno vzít. Sbohem, Ivane.

Ivan: Nino. Nevracej se, dokud se nerozloučím. Neodcházej, dokud se nerozloučím. Neopovažuj se nic zkoušet, dokud ti nedám šanci tady zůstat. Sbohem Nino.

Nina: Počkej chvíli!

Ivan: Nino.

Nina: Ne! To nesmíš. Slíbil jsi to. Miluju tě, Ivane. Slib mi to. Neodcházej, dokud se nerozloučím.

Ivan: Řekl jsi sbohem. A teď nemusíš jít ani ty. Ty se taky nevrátíš. Neboj se. Všechno je v pořádku.

Nina: Prosím, Ivane!

Ivan: Už mě neposlouchej!

Nina: Proč jsi odešel? Co se to se mnou děje? Říkal jsi, že všechno je v pořádku.

Ivan: S tebou to nikdy není v pořádku! Všechno je vždycky špatné! To má být vtíp?

- Nina: Please tell me you forgive me! You didn't even care! What did I do to make you think you could do something like this! Please tell me that!
- Ivan: Don't ask questions.
- Nina: How did you get here? What's going on? Tell me what happened.
- Ivan: There are two things.
- Nina: Yes but first... what is going on?
- Ivan: Don't!
- Nina: Tell me.
- Ivan: Don't cry! Please. Tell me.
- Nina: No! Don't! Let me out of here. Tell me.
- Ivan: Please! You can trust me!
- Nina: Yes yes. I believe him!
- Ivan: He wants me gone!
- Nina: Then why hasn't he killed you yet!?
- Ivan: It wasn't my fault.
- Rusik: Do you see that gun, Nina? Ivan is planning an assassination.**
- Nina: I don't understand! Who does this have in mind?
- Ivan: He knows.
- Rusik: Tonight he want's to assassinate the president. He believe in lies of Military general, that the whole thing is safe.
- Nina: Kill the president?
- Nina: Prosím, řekni mi, že mi odpouštíš! Ani tě to nezajímalo! Co jsem udělat, že si myslíš, že můžeš udělat něco takového! Prosím, řekni mi to!*
- Ivan: Na nic se neptej.*
- Nina: Jak ses sem dostal? Co se děje? Řekni mi, co se stalo.*
- Ivan: Jsou dvě věci.*
- Nina: Ano, ale nejdřív... co se děje?*
- Ivan: Ne!*
- Nina: Řekni mi to.*
- Ivan: No tak, neplač! Prosím. Řekni mi to.*
- Nina: Ne! Ne! Pustte mě odsud. Řekni mi to.*
- Ivan: Prosím! Můžeš mi věřit!*
- Nina: Ano, ano. Já mu věřím!*
- Ivan: Chce mě pryč!*
- Nina: Tak proč tě ještě nezabil!?*
- Ivan: Nebyla to moje chyba.*
- Rusik: Vidíš tu pušku Nino? Ivan chystá atentát.**
- Nina: Já tomu nerozumím! Koho to má na mysli?*
- Ivan: On to ví.*
- Rusik: Dnes večer chce atentát na prezidenta. Věří lžím Vojenského generála, že je celá věc bezpečná.*
- Nina: Zavraždit prezidenta?*
- ↓
- Ivan: It's very important he don't survive this war. I've always felt... I know you don't believe me, but we are not alone in our beliefs, so... if there is one thing that I learned in my military service... I'm a great believer of family first and foremost.
- Ivan: Je velmi důležité, aby nepřežil tuto válku. Vždycky jsem cítil... Vím, že mi nevěříte, ale nejsme ve svém přesvědčení sami, takže... jestli jsem se ve své vojenské službě něco naučil... Jsem velký zastávce rodiny v první řadě.*



Ivan: He has become an obstacle for Russia's future prosperity and security. but I believe his death will cause the Russian economy to revive. This is ~~Ivan Romanov, our chief scientist and leader of our work here at this meeting. . . and this man. . .~~

Ivan: Stal se překážkou pro budoucí prosperitu a bezpečnost Ruska. ale věřím, že jeho smrt způsobí oživení ruské ekonomiky. ~~Tohle je Ivan Romanov, náš hlavní vědec a vedoucí naší práce na tomto setkání. . . a tento muž. . .~~

Rusik: This man does not want a new Cold War with the United States. . . or a nuclear war with him or even any of the Russkies. . . this man wants to have a relationship where they are allies instead we are enemies! ~~If you're looking out for your own safety. . . then get rid of him and he goes! Get rid!! We're talking about killing one man.~~

*Rusik: Tento muž nechce novou studenou válku se Spojenými státy. . . ani jadernou válku s ním, ani s žádným z Rusů. . . tento muž chce mít vztah, kde **nejsou spojenci místo toho pouze nepřátelé!** ~~Jestli se staráš o vlastní bezpečí, tak se ho zbav a on odejde! Zbav se ho! Mluvíme tu o zabití jednoho muže.~~*



Nina: Then how am I gonna save you?!

Nina: Tak jak tě mám zachránit?!

Ivan: It isn't him!

Ivan: To není on!

Nina: Oh please. ~~You told me about this before.~~

Nina: Ale prosím tě. ~~Už jsi mi o tom říkal.~~

Ivan: No!

Ivan: Ne!

Nina: Please. You need help! Ivan! Don't leave me. You need someone! Someone who understands what you're feeling.

Nina: Prosím. Potřebuješ pomoc! Ivane! ~~Ne nechávej mě tu. Potřebuješ někoho! Někoho, kdo rozumí tomu, co cítíš.~~

Ivan: Nina. ~~I will come back all right.~~

Ivan: Nino. ~~V pořádku se vrátím.~~

Nina: Go. ~~You need me! You know what's really happening to you. Please.~~

Nina: Běž. ~~Potřebuješ mě! Víš, co se s tebou doopravdy děje. Prosím.~~

Ivan: You need somebody! Someone to protect you, in the meantime.

*Ivan: ~~Potřebuješ někoho! Někoho, kdo tě **mezitím** ochrání.~~*

Nina: Please. Ivan. Just. . .

Nina: Prosím. Ivane. Jen. . .

Ivan: Nina! It doesn't matter what happens!

Ivan: Nino! ~~Je jedno, co se stane!~~

Nina: No. ~~Just tell me.~~

Nina: Ne. ~~Prostě mi to řekni.~~

Ivan: No! It doesn't matter.

Ivan: Ne! ~~Na tom nezáleží.~~

Nina: You have nothing else to do. Nothing. **We could have run away together in safety.**

*Nina: ~~Nemáš nic jiného na práci. Nic. **Mohli jsme spolu v pořádku utéct.**~~*

Ivan: This is about saving our country.

~~Ivan: Don't care what happens! Don't worry!
You've done Everything has nothing.~~

~~Nina: You just have nothing else!—nothing
else You can You've To me. You just
don't do You To you do nothing else!
To nothing. Nothing You can happen
to do to do to me To me To be. To
nothing! To me You can~~

~~Ivan: What do you mean, I don't want
anything! You said that to me!~~

~~Nina: Oh yes I know!—And how long did
you wait?—Until the police had been
summoned or...before...—Before he
got into the limo with them?!—Then
why did they send us there!?!—I never
told you! You didn't ask!~~

~~Ivan: No!~~

~~Nina: You're breaking the promise. Please.
Please don't do this!~~

Rusik: Don't do this Ivan!

**Ivan: Goodbye. We will meet each other
soon.**

/Ivan tooks his weapon and leaves./

The end

Ivan: Jde tady o záchranu naší vlasti.

~~*Ivan: Je mi jedno, co se stane! Neboj se!
Udělal jsi Všechno nemá nic.*~~

~~*Nina: Nic jiného nemáš!—Nic jiného mi
nemůžeš dát. Ty prostě neděláš Nic
jiného neděláš! Na nic. Nic, co bys mi
mohl udělat, abys byl. Na nic! Mně
můžeš*~~

~~*Ivan: Co tím myslíš, že nic nechci! To jsi mi
řekl!*~~

~~*Nina: Ach ano, já vím!—A jak dlouho
jsi čekal?—Dokud nebyla přivolána
policie nebo...předtím... Než s nimi
nastoupil do limuzíny?! Tak proč nás
tam poslali!?! Nikdy jsem ti to neřekl!
Neptal jste se!*~~

~~*Ivan: Ne!*~~

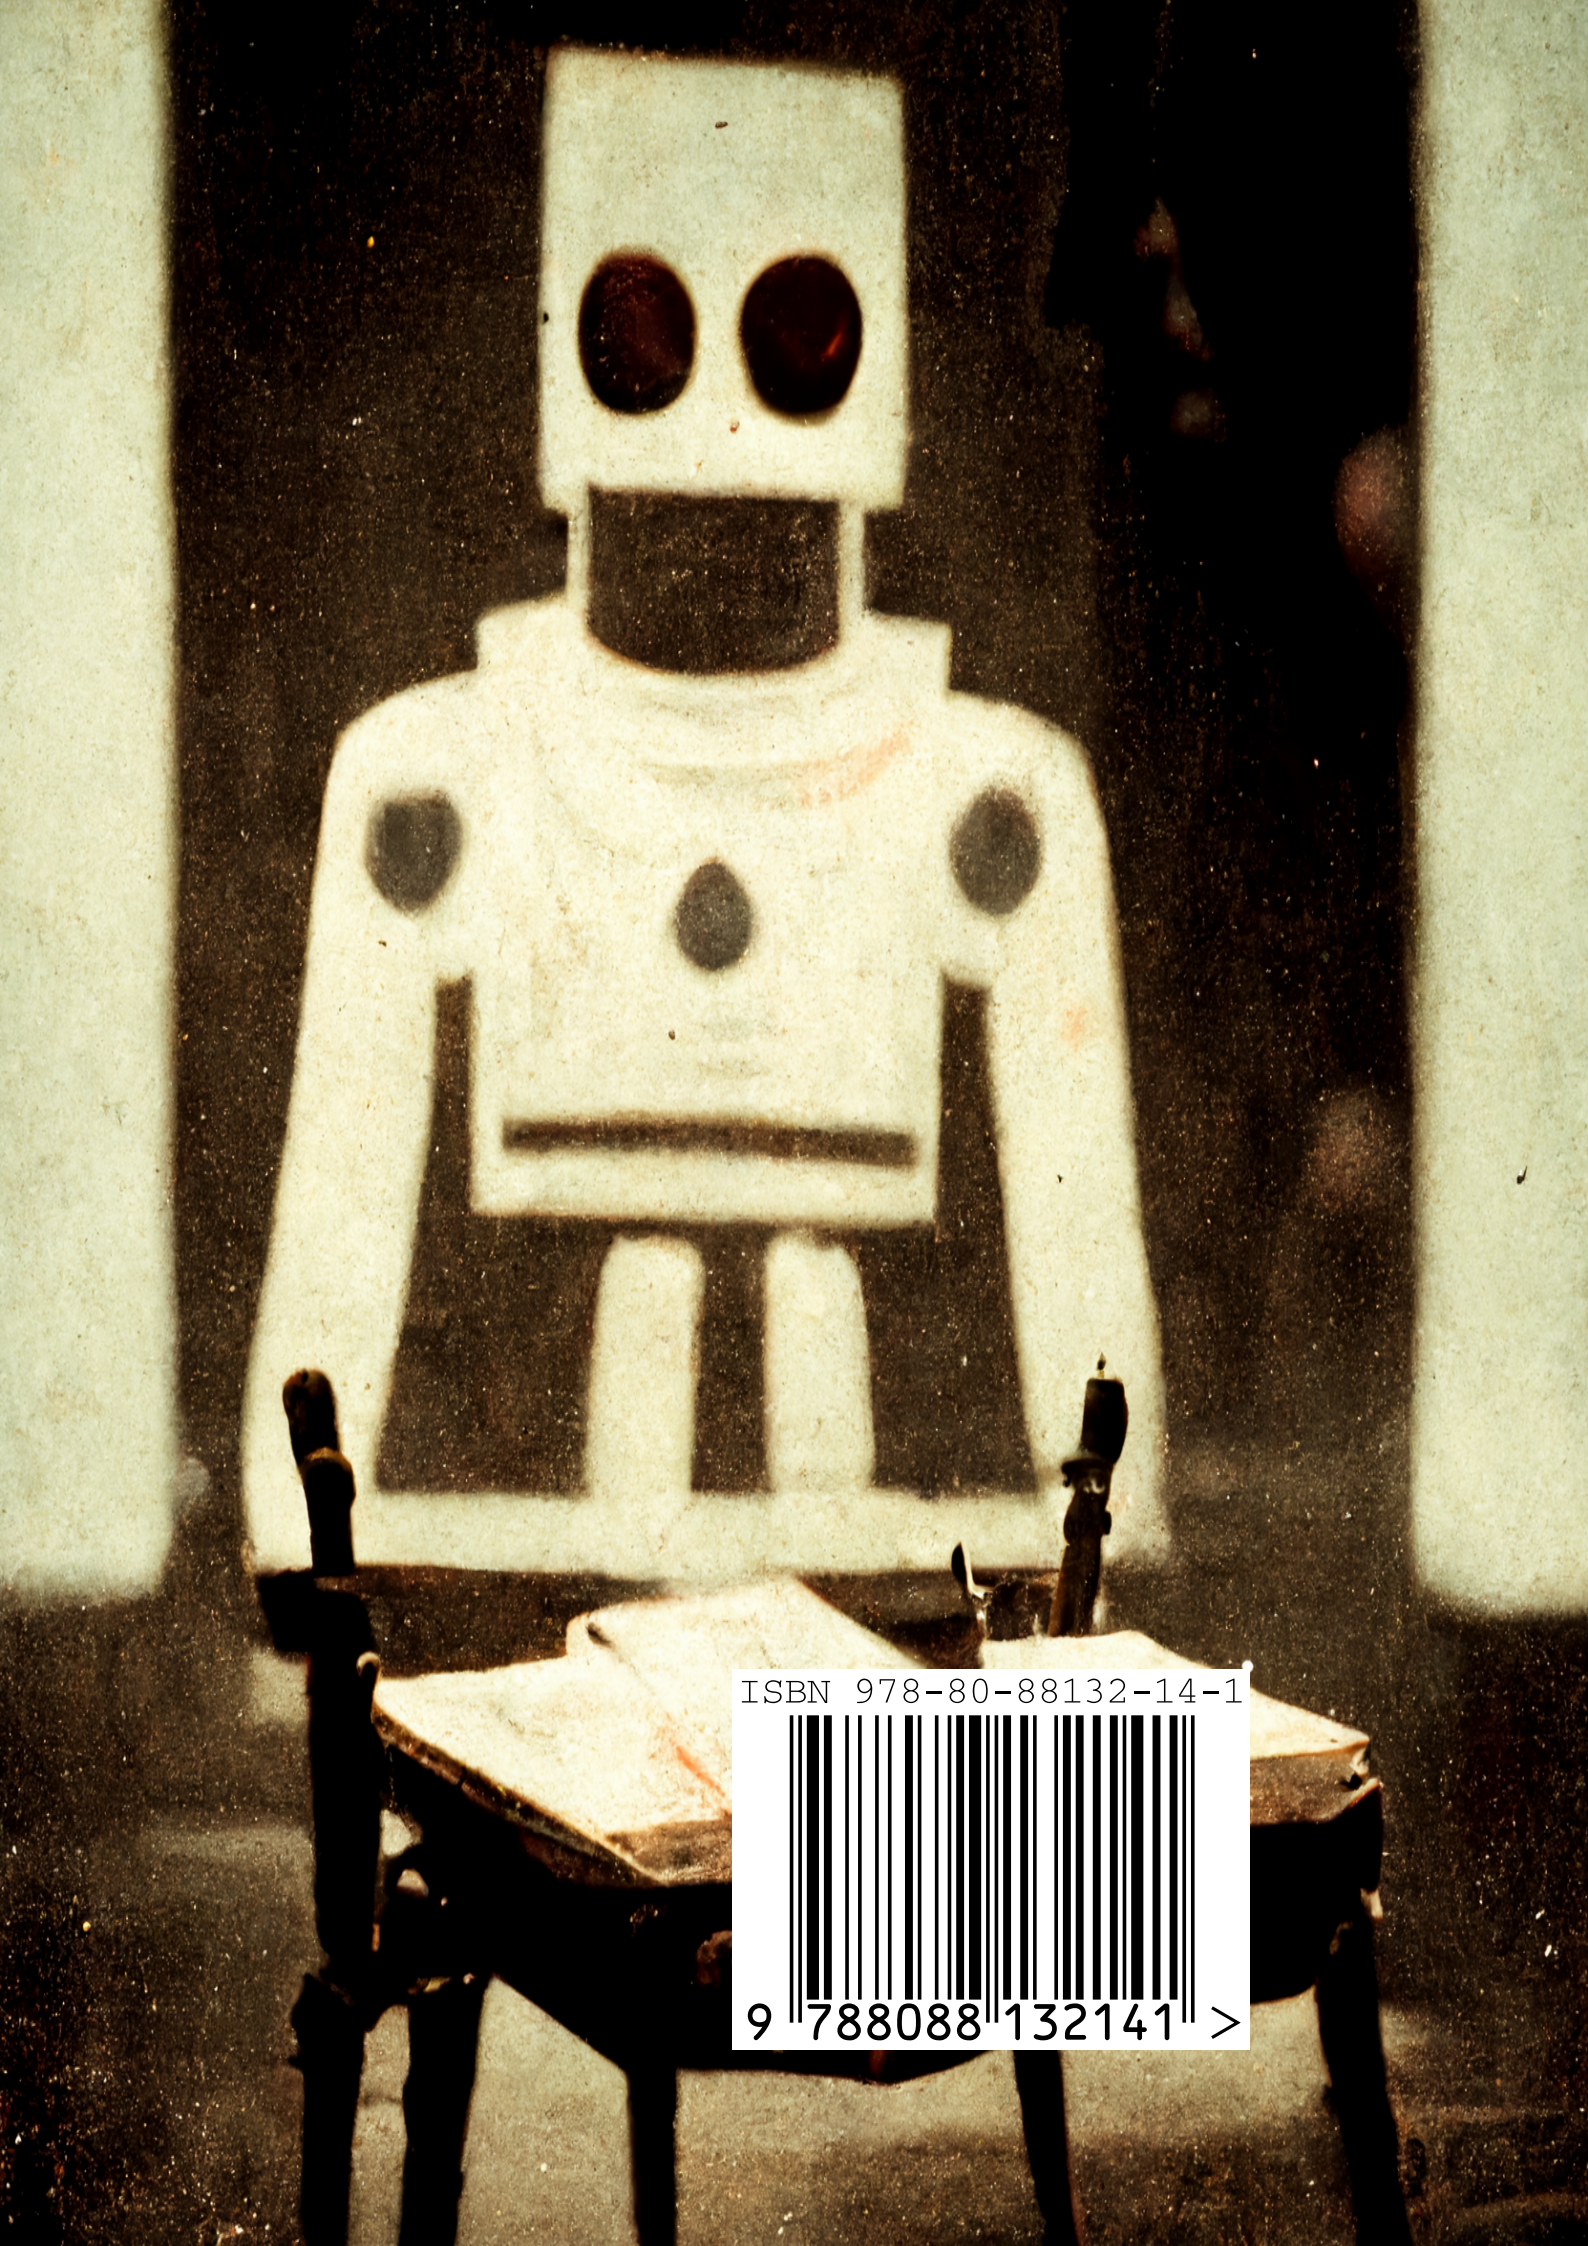
~~*Nina: Porušuješ náš slib. Prosím. Prosím,
nedělej to!*~~

Rusik: Nedělej to Ivane!

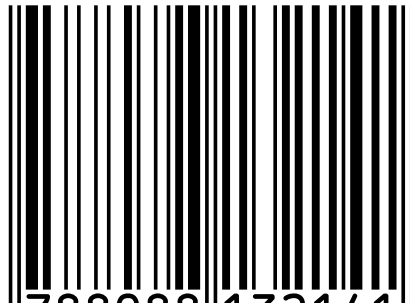
Ivan: Sbohem. Uvidíme se brzy.

/Ivan bere svou pušku a odchází./

Konec.



ISBN 978-80-88132-14-1



9 788088 132141 >