Towards a Truly Statistical Natural Language Generator

for Spoken Dialogues

Ondřej Dušek

Institute of Formal and Applied Linguistics
Charles University in Prague

June 5, 2013

Introduction

Objective of NLG

Given (whatever) input and a **communication goal**, create a natural language string that is **well-formed** and **human-like**.

- Desired properties: simplicity, variation, trainability ...

Usage

- Spoken dialogue systems
- Machine translation
- Short texts: weather reports, customer recommendation ...
- Summarization
- Question answering in knowledge bases

Standard NLG Pipeline (*Textbook*)

[Input]

↓ Content/Text Planning (“what to say”)

- Content selection, basic structuring (ordering)

[Text plan]

↓ Sentence Planning/Realization (“how to say it”)

↓ Microplanning: aggregation, lexical choice, referring…

[Sentence Plan(s)]

↓ Surface realization: linearization according to grammar

[Text]

Real NLG Systems

Few systems implement the whole pipeline

- Systems focused on content planning with trivial surface realization
- Surface-realization-only systems
- Word-order-only systems
- Input/intermediate data representation varies greatly

Possible approaches

- **Rule/template-based** (if-then-else, filling in slots)
- **Grammar-based** (various formalisms, e.g. *FUG, CCG*)
- Only since 2000s: **Statistical** … or rather hybrid

Introducing Statistical Methods to NLG

Rule-based methods

- Simple, straightforward, fast
- Surface realizers: once and for all
- Reliable (important!)
  - Content plans custom-tailored for domain
  - Surface realizer sure to produce grammatical output

Statistical methods

- Easier to maintain
- Easily adaptable to new domains
- Robust to unseen input
- Add variation, (hopefully) naturalness

Trainable Content Planning: User Models

- Presentation strategy based on user model
  - initial questions
- Adaptive, but rule-based
- *MATCH*, *GEA*, *FLIGHTS*

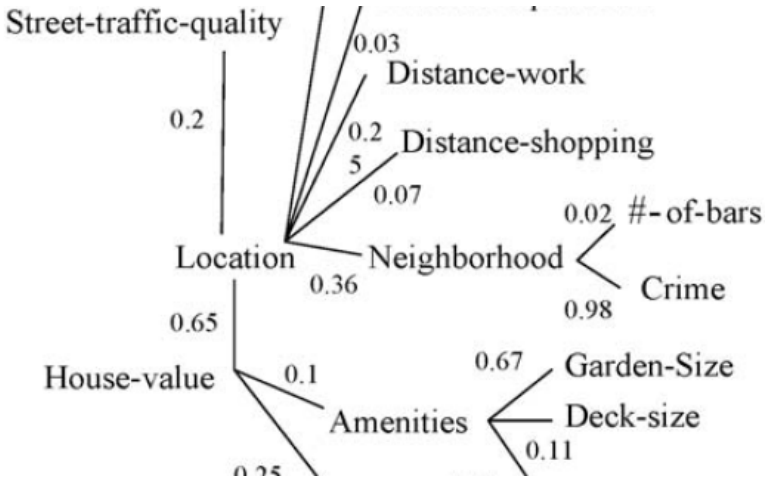

Show me Italian restaurants in the West Village.

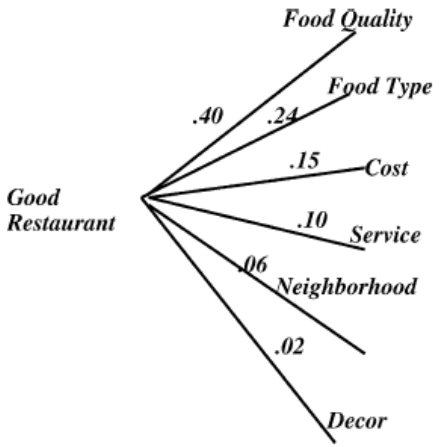

$$U_h = \sum_{k=1}^{K} w_k u_k(x_{kh})$$

U_h... total utility of option h
$u_k(x_{kh})$... utility of k-th attribute
w_k... user-specific weight of k-th attribute

Trainable Content Planning: Overgenerate and Rank

- Rule-based sentence plan generator (clause combining operations)
  - Randomly sample several sentence plans
- Reranker (RankBoost) trained on hand-annotated sentence plans
  - Rank plans and select the best one
- *SPoT*

```
implicit-confirm(orig-city:NEWARK)
implicit-confirm(dest-city:DALLAS)
implicit-confirm(month:9)
implicit-confirm(day-number:1)
request(depart-time)
```

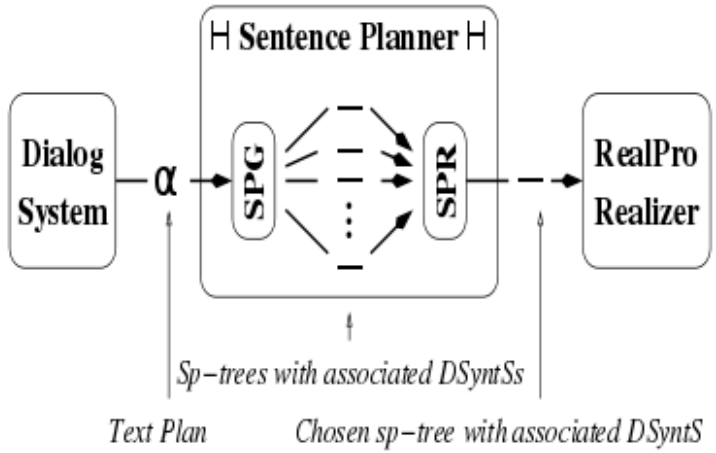

Alt	Realization	H	RB
0	What time would you like to travel on September the 1st to Dallas from Newark?	5	.85
5	Leaving on September the 1st. What time would you like to travel from Newark to Dallas?	4.5	.82
8	Leaving in September. Leaving on the 1st. What time would you, traveling from Newark to Dallas, like to leave?	2	.39

Trainable Content Planning: Reinforcement Learning

- Reinforcement learning of presentation strategy
- Communicative Goal: Dialogue Act + desired user reaction
- Plan lower-level NLG actions to achieve goal
- Markov Decision Process

$$Q^{\pi}(s,a) = \sum_{s'} \mathcal{T}^{a}_{ss'} \left(\mathcal{R}^{a}_{ss'} + \gamma \mathcal{V}^{\pi}(s') \right)$$

- *RL-NLG*

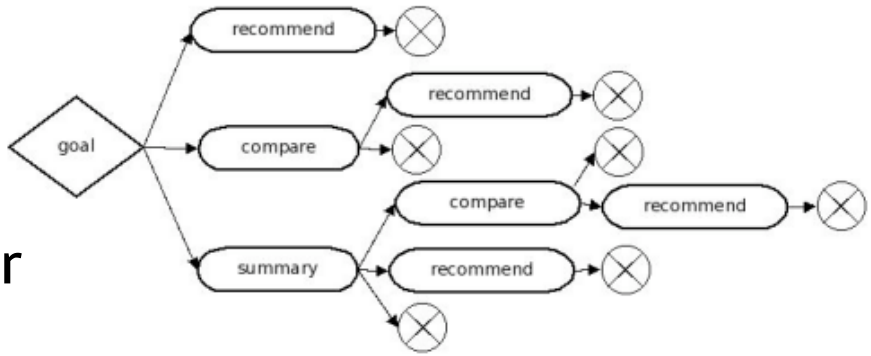

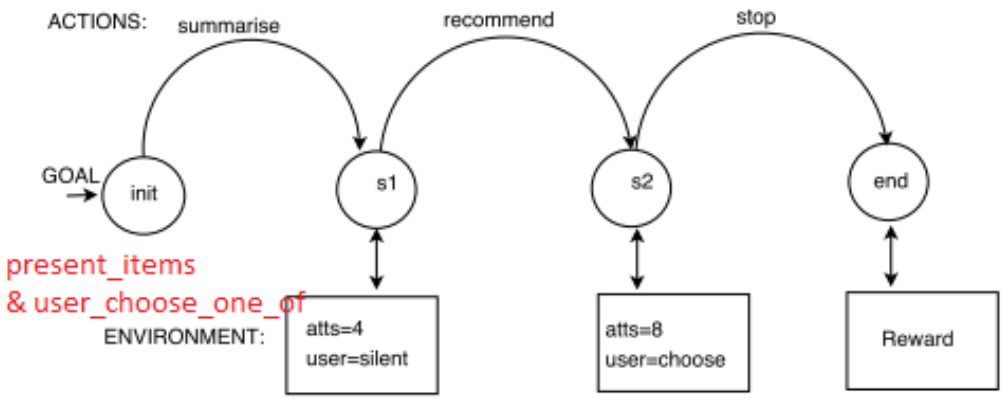

Trainable Surface Realizers: Overgenerate and Rank

- Require a handcrafted realizer, e.g. CCG realizer
- Input underspecified $\rightarrow$ more outputs possible
- Overgenerate
- Then use a statistical reranker
- Ranking according to:
  - *NITROGEN, HALOGEN*: *n*-gram models
  - *FERGUS*: Tree models (XTAG grammar)
  - *Nakatsu and White*: Predicted Text-To-Speech quality
  - *CRAG*: Personality traits (extraversion, agreeableness...) + alignment (repeating words uttered by dialogue counterpart)
- Provides variance, but at a greater computational cost

Trainable Surface Realizers: Parameter Optimization

- Still require a handcrafed realizer
- Train handcrafted realizer parameters
- No overgeneration
- Realizer needs to be “flexible”

Examples

- *Paiva and Evans*: linguistic features annotated in corpus generated with many parameter settings, correlation analysis
- *PERSONAGE-PE*: personality traits connected to linguistic features via machine learning

Statistical Surface Realizers

Using methods of Machine Translation

- "translating" from semantic representation to text
- *PHARAOH* SMT / synchronous CFG + MaxEnt (*WASP*$^{-1}$)
- hybrid trees with CRFs (*TreeCRF*)

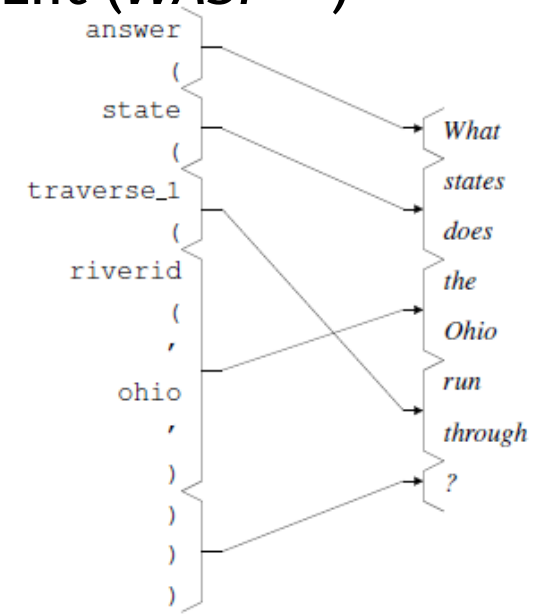

Syntax-based

- *Bohnet et al.*: pipeline model with SVMs
- Meaning-Text Theory
- Semantics $\rightarrow$ Syntax $\rightarrow$ Linearization $\rightarrow$ Morphologization

Fully Statistical Natural Language Generators

- Few, based on supervised learning
- Limited domain
- Hierarchical, phrase-based
- *Mairesse et al.*: Bayesian networks
  - semantic stacks
- *Angeli et al.*: log-linear model
  - records ↘ fields ↘ templates

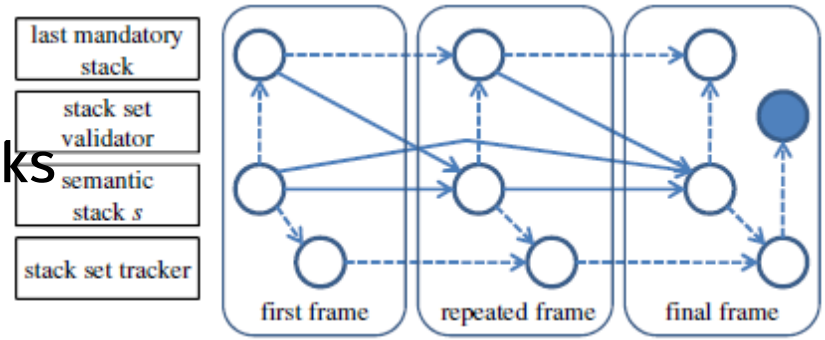

Charlie Chan	*is a*	*Chinese*	*restaurant*	*near*	*Cineworld*	*in the*	*centre of town*
Charlie Chan name inform	inform	Chinese food inform	restaurant type inform	near inform	Cineworld near inform	area inform	centre area inform
$t = 1$	$t = 2$	$t = 3$	$t = 4$	$t = 5$	$t = 6$	$t = 7$	$t = 8$

Language Generation at ÚFAL: Current State

Prior work

- For Czech
- Surface realization only, rule-based
- Based on *FGD*, tecto-trees
  - Functors / formemes
- *Ptáček and Žabokrtský, TectoMT*

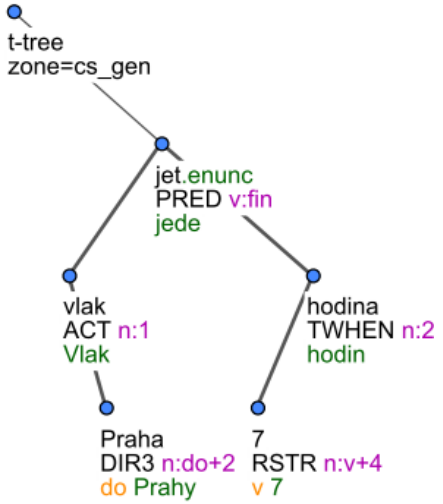

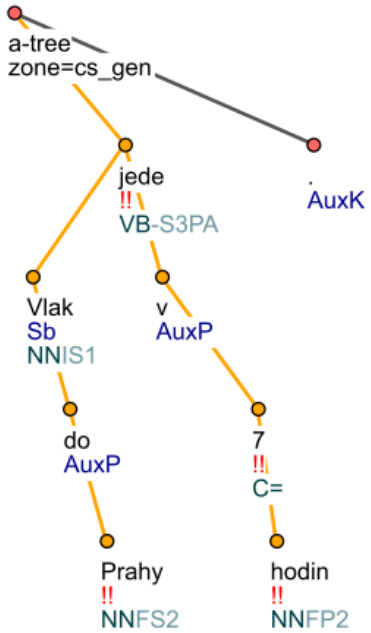

NLG for Dialogue Systems

do	*>0-ing*	doing
vědět	*>4-íme,<ne*	nevíme
Mann	*>0-er,3:1-ä*	Männer

- Mixing templates and tecto-trees

```
Vlak [Praha|n:do+2|gender:fem]
jede v [[7|adj:attr] hodina|n:4|gender:fem].
```

- Statistical word form generator (*Flect*)

Prospects

Desired properties of a new NLG system for dialogues

- Trainable: simple domain adaptation
- Variable: no fixed templates
- Multilingual: Czech and English at the very least

Planned approach

- *FGD*, tecto-trees as a useful formalism
- Surface realizer at least partially trainable
  - Many grammar rules can be learned from corpora
  - Statistical morphology generation: avoiding dictionaries
- Content planner fully trainable
  - Using MT-inspired methods for content planning?

Thank You

You can find these slides, including references, at:
`http://ufal.mff.cuni.cz/~odusek/slides/2013_wds.pdf`

You can contact me at:
`odusek@ufal.mff.cuni.cz`

References

Angeli — Angeli, G. et al. 2010. A Simple Domain-Independent Probabilistic Approach to Generation. *EMNLP*

Bohnet — Bohnet, B. et al. 2010. Broad coverage multilingual deep sentence generation with a stochastic multi-level realizer. *COLING*

CCG — White, M. and Baldrige, J. 2003. Adapting Chart Realization to CCG. *ENLG* `http://openccg.sourceforge.net/`

CRAG — Isard, A. et al. 2006. Individuality and alignment in generated dialogues. *INLG*

FERGUS — Bangalore, S. and Rambow, O. 2000. Exploiting a probabilistic hierarchical model for generation. *COLING*

FGD — Sgall, P. et al. 1986. *The meaning of the sentence in its semantic and pragmatic aspects*. Springer

Flect — Dušek, O. and Jurčíček, F. 2013. Robust multilingual statistical morphological generation models. *ACL Student Research Workshop*

FLIGHTS — Moore, J. et al. 2004. Generating Tailored, Comparative Descriptions in Spoken Dialogue. *FLAIRS*
Demberg, V. and Moore, J. 2006. Information presentation in spoken dialogue systems. *EACL*

FUG — Elhadad, M. and Robin, J. 1996. An overview of SURGE: A reusable comprehensive syntactic realization component. `http://www.cs.bgu.ac.il/surge/`

GEA — Carenini, G. and Moore, J. 2006. Generating and evaluating evaluative arguments. *Artificial Intelligence*

Hajič — Hajič, J. 2004. *Disambiguation of Rich Inflection -- Computational Morphology of Czech*. Karolinum

HALOGEN — Langkilde-Geary, I. 2002. An empirical verification of coverage and correctness for a general-purpose sentence generator. *INLG*

Mairesse — Mairesse, F. et al. 2010. Phrase-based statistical language generation using graphical models and active learning. *ACL*

MATCH — Walker, M. et al. 2004. Generation and evaluation of user tailored responses in multimodal dialogue. *Cognitive Science*

Nakatsu&White — Nakatsu, C. and White, M. 2006. Learning to say it well: reranking realizations by predicted synthesis quality. *COLING-ACL*

NITROGEN — Langkilde, I. and Knight, K. 1998. Generation that exploits corpus-based statistical knowledge. *ACL-COLING*

Paiva&Evans — Paiva, D. S. and Evans, R. 2005. Empirically-based control of natural language generation. *ACL*

References

PERSONAGE-PE Mairesse, F. and Walker, M. 2008. Trainable generation of big-five personality styles through data-driven parameter estimation. *ACL*

PHARAOH Koehn. P. 2004. Pharaoh: a beam search decoder for phrase-based statistical machine translation models. *Machine Translation: From Real Users to Research*. Springer

Ptáček&Žabokrtský Ptáček, J. and Žabokrtský, Z. 2006. Synthesis of Czech Sentences from Tectogrammatical Trees. *TSD*

RL-NLG Lemon, 0. 2008. Adaptive natural language generation in dialogue using Reinforcement Learning. *SEMdial*
Rieser, V. and Lemon, 0. 2010. Natural language generation as planning under uncertainty for spoken dialogue systems. *EMNLP*

SimpleNLG Gatt, A. and Reiter, E. 2009. SimpleNLG: A realisation engine for practical applications. *ENLG*

SPoT Walker, M. et al. 2001. SPoT: A trainable sentence planner. *NAACL*

TectoMT Žabokrtský, Z. et al. 2008. TectoMT: highly modular MT system with tectogrammatics used as transfer layer. *WMT*

Textbook Reiter, E. and Dale, R. 2000. *Building natural language generation systems*. Cambridge Univ. Press

TreeCRF Lu, W. et al. 2009. Natural Language Generation with Tree Conditional Random Fields. *EMNLP*

WASP^{-1} Wong, Y. W. and Mooney, R. J. 2007. Generation by inverting a semantic parser that uses statistical machine translation. *NAACL*
Wong, Y. W. *Learning for semantic parsing and natural language generation using statistical machine translation techniques*. PhD Thesis, University of Texas at Austin

Further NLG Links

C. DiMarco's slides: `https://cs.uwaterloo.ca/~jchampai/CohenClass.en.pdf`
F. Mairesse's slides: `http://people.csail.mit.edu/francois/research/papers/ART-NLG.pdf`
J. Moore's NLG course: `http://www.inf.ed.ac.uk/teaching/courses/nlg/`
NLG Systems Wiki: `http://www.nlg-wiki.org`
Wikipedia: `http://en.wikipedia.org/wiki/Natural_language_generation`