

NASNet, Speech Synthesis, External Memory Networks

Milan Straka

 May 18, 2020



EUROPEAN UNION
European Structural and Investment Fund
Operational Programme Research,
Development and Education

Charles University in Prague
Faculty of Mathematics and Physics
Institute of Formal and Applied Linguistics



unless otherwise stated

- We can design neural network architectures using reinforcement learning.
- The designed network is encoded as a sequence of elements, and is generated using an *RNN controller*, which is trained using the REINFORCE with baseline algorithm.

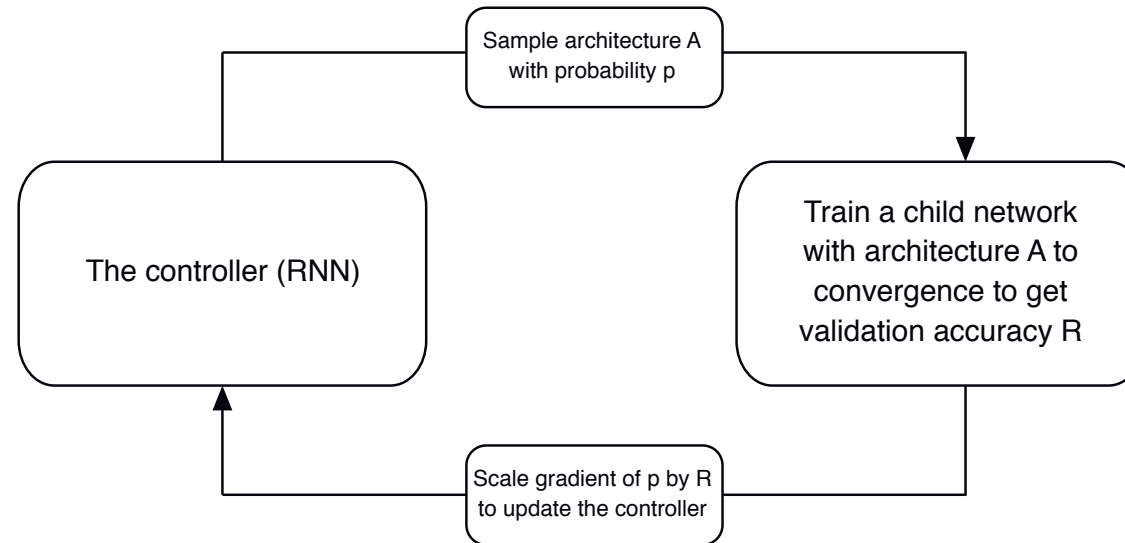


Figure 1 of paper "Learning Transferable Architectures for Scalable Image Recognition", <https://arxiv.org/abs/1707.07012>.

- For every generated sequence, the corresponding network is trained on CIFAR-10 and the development accuracy is used as a return.

The overall architecture of the designed network is fixed and only the Normal Cells and Reduction Cells are generated by the controller.

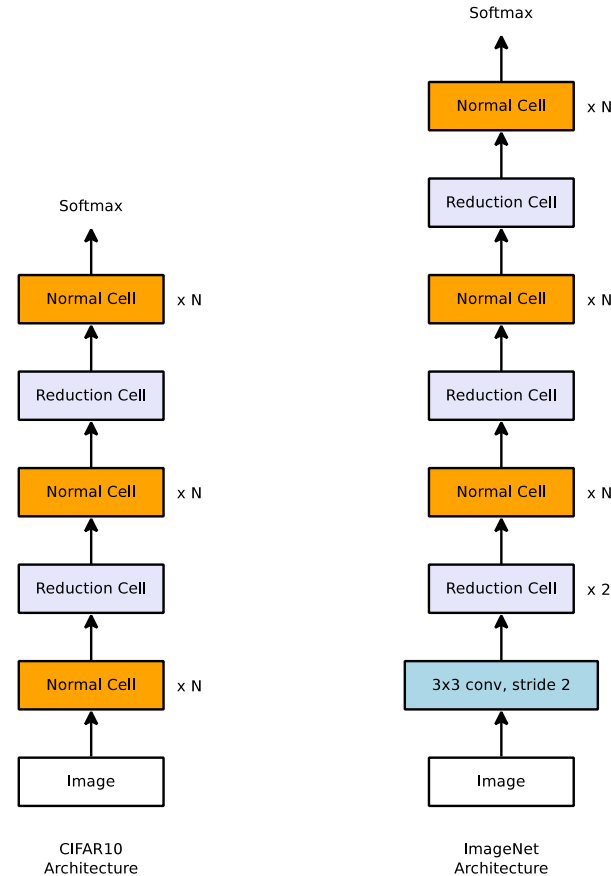


Figure 2 of paper "Learning Transferable Architectures for Scalable Image Recognition", <https://arxiv.org/abs/1707.07012>.

- Each cell is composed of B blocks ($B = 5$ is used in NASNet).
- Each block is designed by a RNN controller generating 5 parameters.

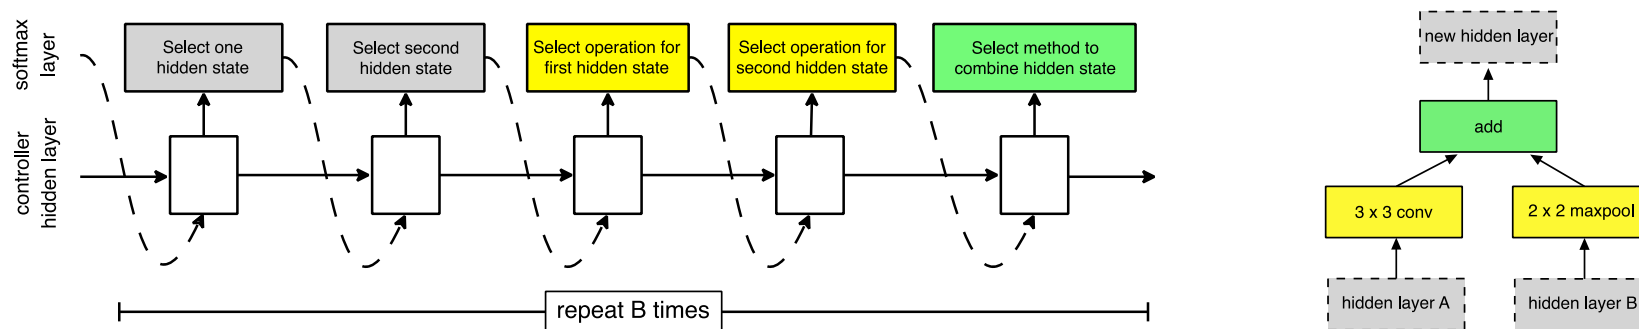


Figure 3. Controller model architecture for recursively constructing one block of a convolutional cell. Each block requires selecting 5 discrete parameters, each of which corresponds to the output of a softmax layer. Example constructed block shown on right. A convolutional cell contains B blocks, hence the controller contains $5B$ softmax layers for predicting the architecture of a convolutional cell. In our experiments, the number of blocks B is 5.

Figure 3 of paper "Learning Transferable Architectures for Scalable Image Recognition", <https://arxiv.org/abs/1707.07012>.

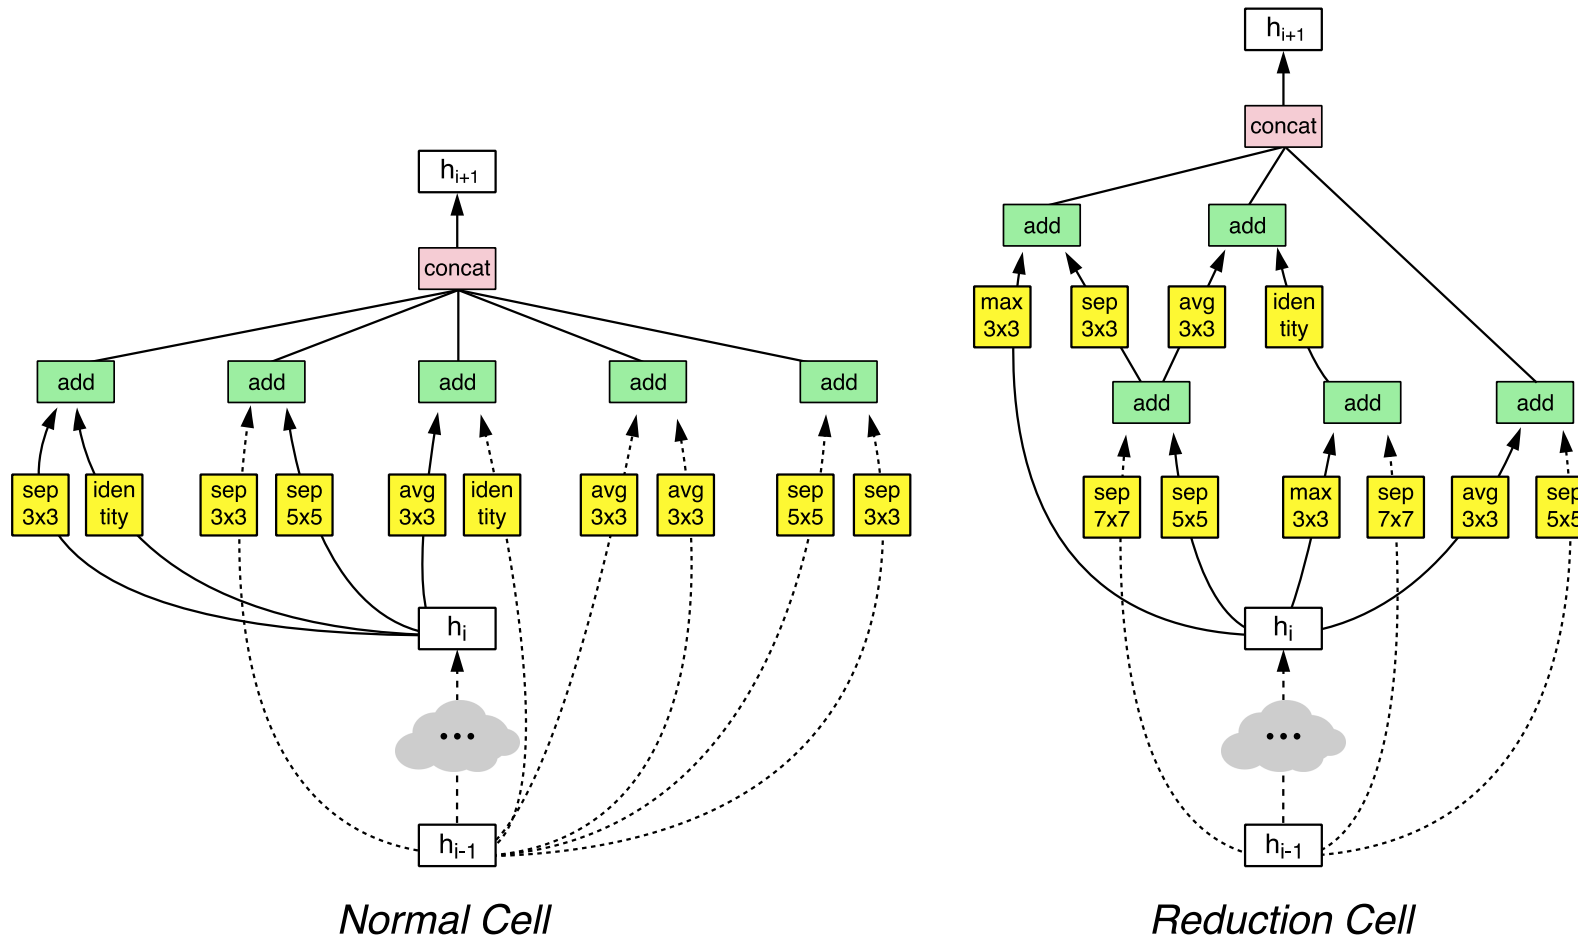
- Step 1.** Select a hidden state from h_i, h_{i-1} or from the set of hidden states created in previous blocks.
- Step 2.** Select a second hidden state from the same options as in Step 1.
- Step 3.** Select an operation to apply to the hidden state selected in Step 1.
- Step 4.** Select an operation to apply to the hidden state selected in Step 2.
- Step 5.** Select a method to combine the outputs of Step 3 and 4 to create a new hidden state.

- identity
- 1x7 then 7x1 convolution
- 3x3 average pooling
- 5x5 max pooling
- 1x1 convolution
- 3x3 depthwise-separable conv
- 7x7 depthwise-separable conv
- 1x3 then 3x1 convolution
- 3x3 dilated convolution
- 3x3 max pooling
- 7x7 max pooling
- 3x3 convolution
- 5x5 depthwise-separable conv

Page 3 of paper "Learning Transferable Architectures for Scalable Image Recognition", <https://arxiv.org/abs/1707.07012>.

Figure 2 of paper "Learning Transferable Architectures for Scalable Image Recognition", <https://arxiv.org/abs/1707.07012>.

The final proposed Normal Cell and Reduction Cell:



Page 3 of paper "Learning Transferable Architectures for Scalable Image Recognition", <https://arxiv.org/abs/1707.07012>.

EfficientNet changes the search in two ways.

- Computational requirements are part of the return. Notably, the goal is to find an architecture m maximizing

$$\text{DevelopmentAccuracy}(m) \cdot \left(\frac{\text{TargetFLOPS}=400\text{M}}{\text{FLOPS}(m)} \right)^{0.07}$$

where the constant 0.07 balances the accuracy and FLOPS.

- Using a different search space, which allows to control kernel sizes and channels in different parts of the overall architecture (compared to using the same cell everywhere as in NASNet).

EfficientNet Search

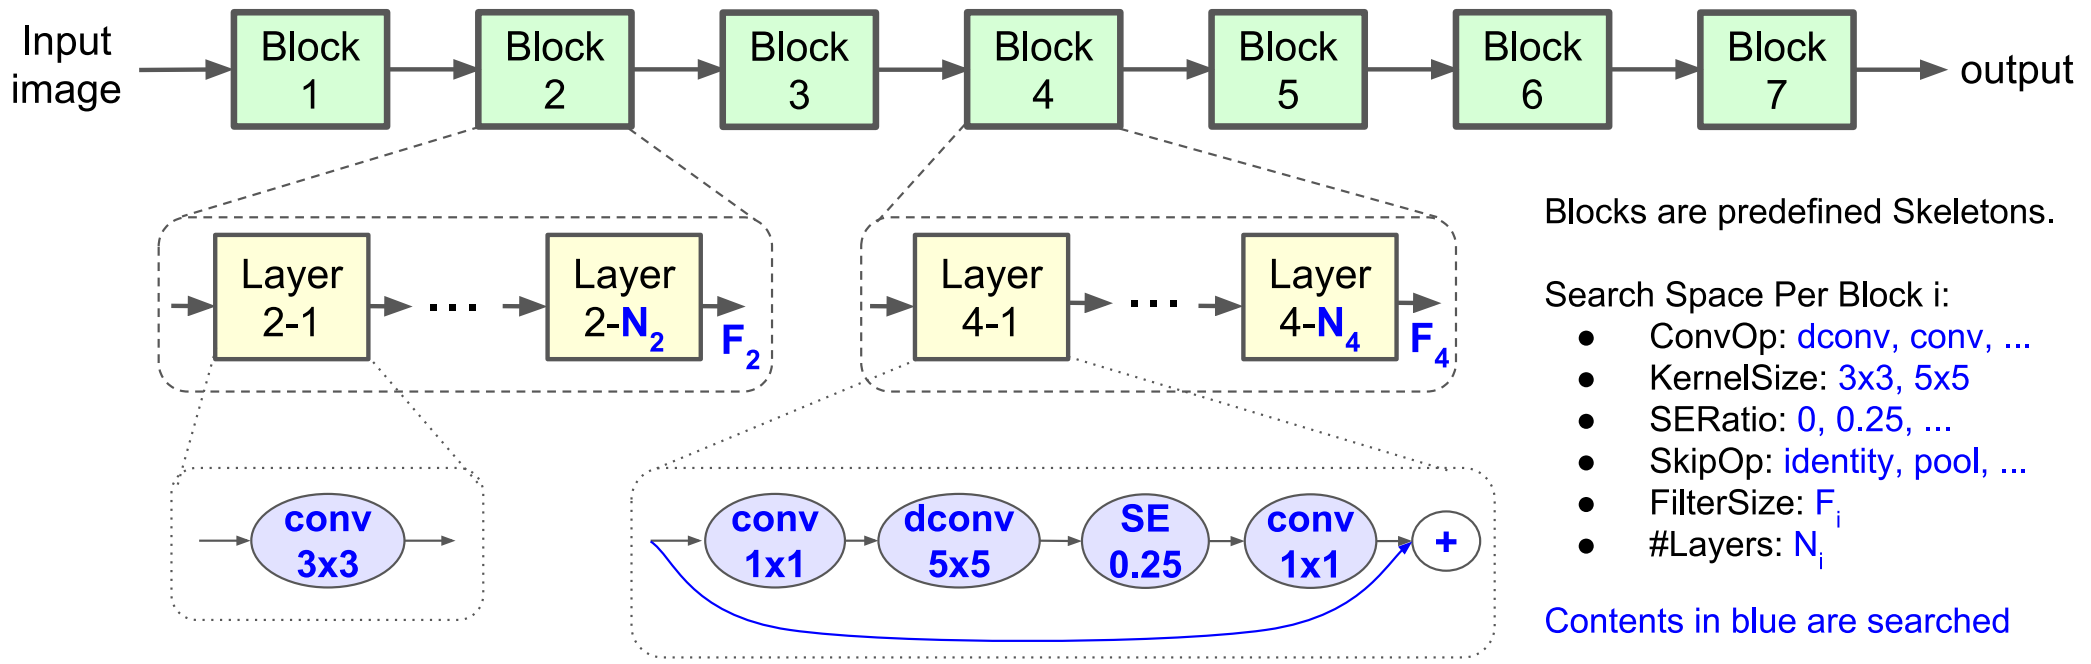


Figure 4 of paper "MnasNet: Platform-Aware Neural Architecture Search for Mobile", <https://arxiv.org/abs/1807.11626>.

The overall architecture consists of 7 blocks, each described by 6 parameters – 42 parameters in total, compared to 50 parameters of NASNet search space.

- Convolutional ops *ConvOp*: regular conv (conv), depthwise conv (dconv), and mobile inverted bottleneck conv [29].
- Convolutional kernel size *KernelSize*: 3x3, 5x5.
- Squeeze-and-excitation [13] ratio *SERatio*: 0, 0.25.
- Skip ops *SkipOp*: pooling, identity residual, or no skip.
- Output filter size F_i .
- Number of layers per block N_i .

Page 4 of paper "MnasNet: Platform-Aware Neural Architecture Search for Mobile", <https://arxiv.org/abs/1807.11626>.

Stage i	Operator $\hat{\mathcal{F}}_i$	Resolution $\hat{H}_i \times \hat{W}_i$	#Channels \hat{C}_i	#Layers \hat{L}_i
1	Conv3x3	224×224	32	1
2	MBCConv1, k3x3	112×112	16	1
3	MBCConv6, k3x3	112×112	24	2
4	MBCConv6, k5x5	56×56	40	2
5	MBCConv6, k3x3	28×28	80	3
6	MBCConv6, k5x5	14×14	112	3
7	MBCConv6, k5x5	14×14	192	4
8	MBCConv6, k3x3	7×7	320	1
9	Conv1x1 & Pooling & FC	7×7	1280	1

Table 1 of paper "EfficientNet: Rethinking Model Scaling for Convolutional Neural Networks", <https://arxiv.org/abs/1905.11946>

Our goal is to model speech, using a auto-regressive model

$$P(\mathbf{x}) = \prod_t P(x_t | x_{t-1}, \dots, x_1).$$




Figure 2: Visualization of a stack of causal convolutional layers.

Figure 2 of paper "WaveNet: A Generative Model for Raw Audio", <https://arxiv.org/abs/1609.03499>.

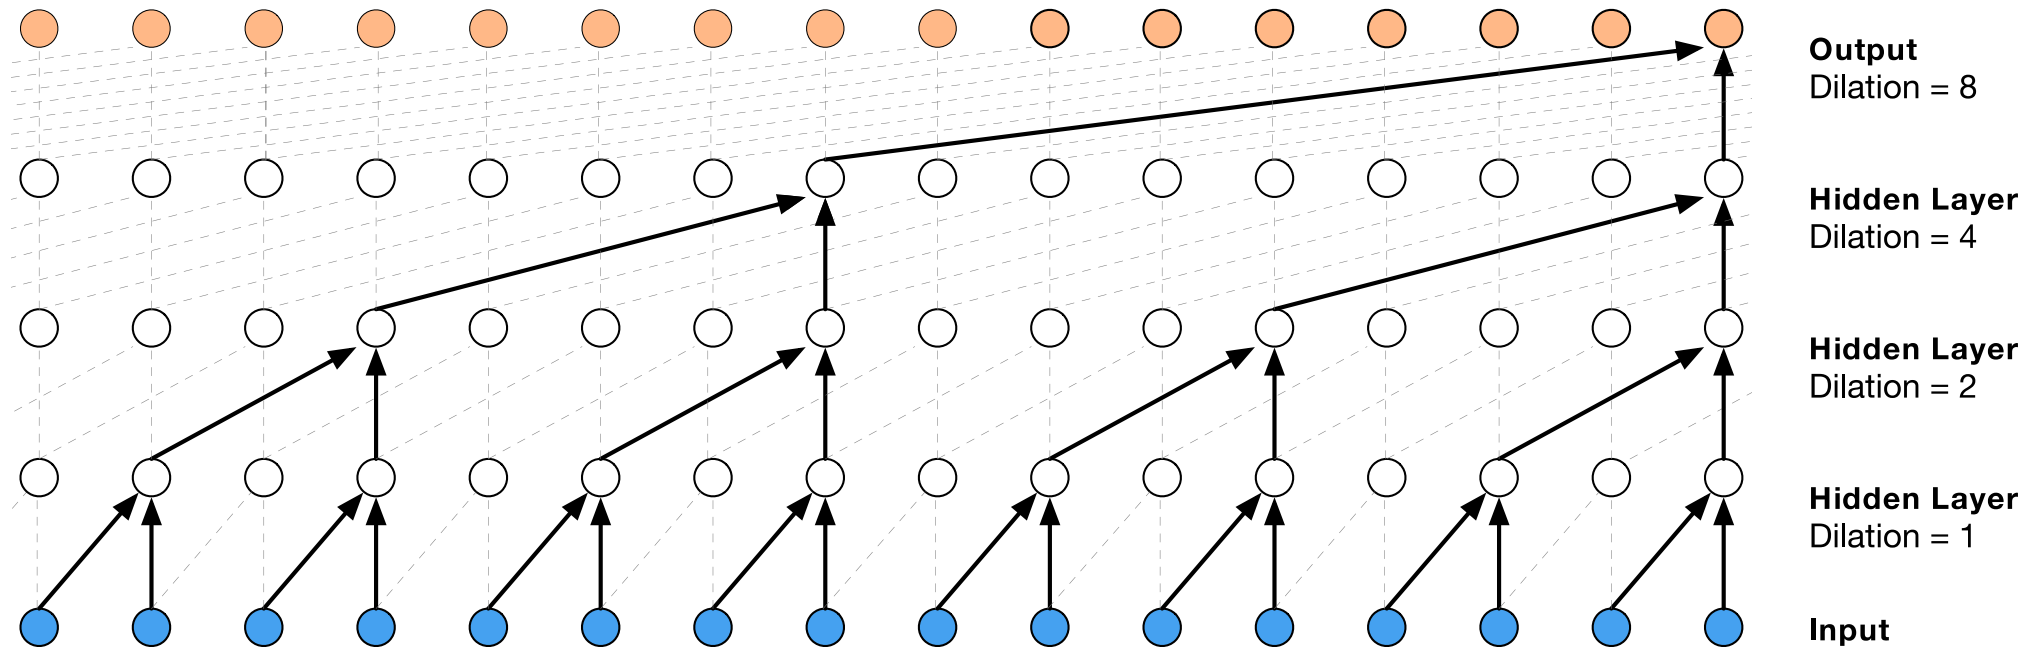


Figure 3: Visualization of a stack of *dilated* causal convolutional layers.

Figure 3 of paper "WaveNet: A Generative Model for Raw Audio", <https://arxiv.org/abs/1609.03499>.

Output Distribution

The raw audio is usually stored in 16-bit samples. However, classification into 65 536 classes would not be tractable, and instead WaveNet adopts μ -law transformation and quantize the samples into 256 values using

$$\text{sign}(x) \frac{\ln(1 + 255|x|)}{\ln(1 + 255)}.$$

Gated Activation

To allow greater flexibility, the outputs of the dilated convolutions are passed through the gated activation units

$$\mathbf{z} = \tanh(\mathbf{W}_f * \mathbf{x}) \cdot \sigma(\mathbf{W}_g * \mathbf{x}).$$

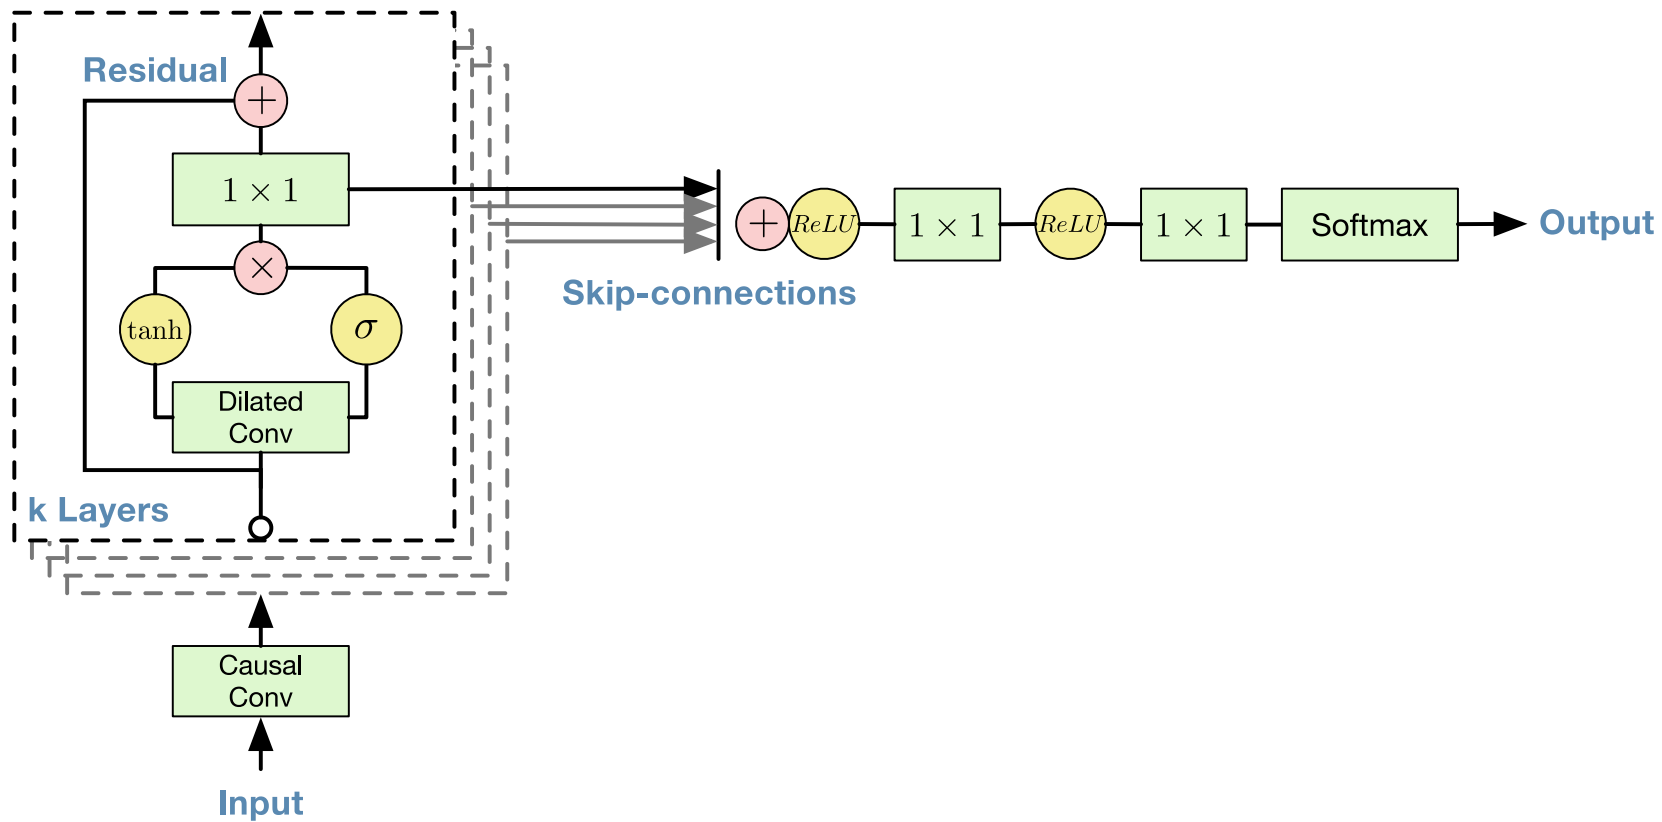


Figure 4: Overview of the residual block and the entire architecture.

Figure 4 of paper "WaveNet: A Generative Model for Raw Audio", <https://arxiv.org/abs/1609.03499>.

Global Conditioning

Global conditioning is performed by a single latent representation \mathbf{h} , changing the gated activation function to

$$\mathbf{z} = \tanh(\mathbf{W}_f * \mathbf{x} + \mathbf{V}_f \mathbf{h}) \cdot \sigma(\mathbf{W}_g * \mathbf{x} + \mathbf{V}_g \mathbf{h}).$$

Local Conditioning

For local conditioning, we are given a time series h_t , possibly with a lower sampling frequency. We first use transposed convolutions $\mathbf{y} = f(\mathbf{h})$ to match resolution and then compute analogously to global conditioning

$$\mathbf{z} = \tanh(\mathbf{W}_f * \mathbf{x} + \mathbf{V}_f * \mathbf{y}) \cdot \sigma(\mathbf{W}_g * \mathbf{x} + \mathbf{V}_g * \mathbf{y}).$$

The original paper did not mention hyperparameters, but later it was revealed that:

- 30 layers were used
 - grouped into 3 dilation stacks with 10 layers each
 - in a dilation stack, dilation rate increases by a factor of 2, starting with rate 1 and reaching maximum dilation of 512
- filter size of a dilated convolution is 3
- residual connection has dimension 512
- gating layer uses 256+256 hidden units
- the 1×1 output convolution produces 256 filters
- trained for 1 000 000 steps using Adam with a fixed learning rate of 0.0002

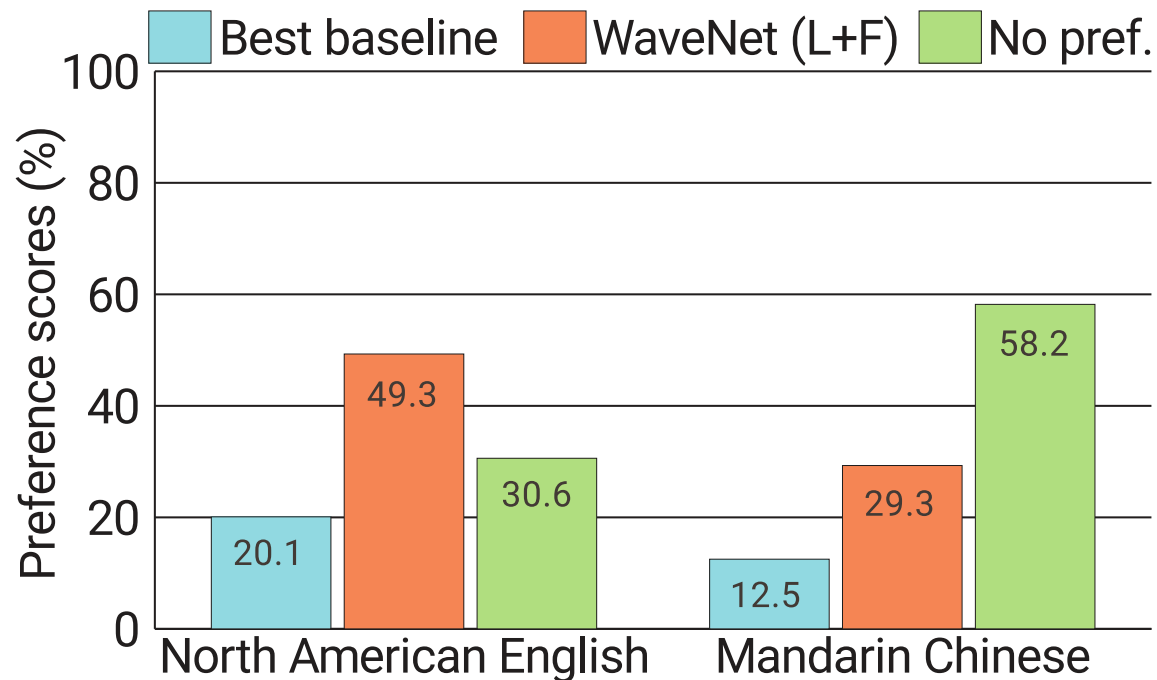


Figure 5: Subjective preference scores (%) of speech samples between (top) two baselines, (middle) two WaveNets, and (bottom) the best baseline and WaveNet. Note that LSTM and Concat correspond to LSTM-RNN-based statistical parametric and HMM-driven unit selection concatenative baseline synthesizers, and WaveNet (L) and WaveNet (L+F) correspond to the WaveNet conditioned on linguistic features only and that conditioned on both linguistic features and $\log F_0$ values.

Figure 5 of paper "WaveNet: A Generative Model for Raw Audio", <https://arxiv.org/abs/1609.03499>.

The output distribution was changed from 256 μ -law values to a Mixture of Logistic (suggested in another paper – PixelCNN++, but reused in other architectures since):

$$\nu \sim \sum_i \pi_i \text{logistic}(\mu_i, s_i).$$

The logistic distribution is a distribution with a σ as cumulative density function (where the mean and scale is parametrized by μ and s). Therefore, we can write

$$P(x|\boldsymbol{\pi}, \boldsymbol{\mu}, \boldsymbol{s}) = \sum_i \pi_i \left[\sigma((x + 0.5 - \mu_i)/s_i) - \sigma((x - 0.5 - \mu_i)/s_i) \right],$$

where we replace -0.5 and 0.5 in the edge cases by $-\infty$ and ∞ .

In Parallel WaveNet, 10 mixture components are used.

Auto-regressive (sequential) inference is extremely slow in WaveNet.

Instead, we use a following trick. We will model $P(x_t)$ as $P(x_t|\mathbf{z}_{<t})$ for a *random* \mathbf{z} drawn from a logistic distribution. Then, we compute

$$\mathbf{x}_t^1 = z_t \cdot s^1(\mathbf{z}_{<t}) + \mu^1(\mathbf{z}_{<t}).$$

Usually, one iteration of the algorithm does not produce good enough results – 4 iterations were used by the authors. In further iterations,

$$\mathbf{x}_t^i = \mathbf{x}_t^{i-1} \cdot s^i(\mathbf{x}_{<t}^{i-1}) + \mu^i(\mathbf{x}_{<t}^{i-1}).$$

After N iterations, $P(\mathbf{x}_t^N | \mathbf{z}_{\leq t})$ is a logistic distribution with location $\boldsymbol{\mu}_{\text{tot}}$ and scale \mathbf{s}_{tot} with

$$\boldsymbol{\mu}_{\text{tot}} = \sum_i^N \mu^i(\mathbf{x}_{<t}^{i-1}) \cdot \left(\prod_{j>i}^N s^j(\mathbf{x}_{<t}^{j-1}) \right) \quad \text{and} \quad \mathbf{s}_{\text{tot}} = \sum_i^N s^i(\mathbf{x}_{<t}^{i-1}).$$

Parallel WaveNet

The network is trained using a *probability density distillation* using a teacher WaveNet, using KL-divergence as loss.

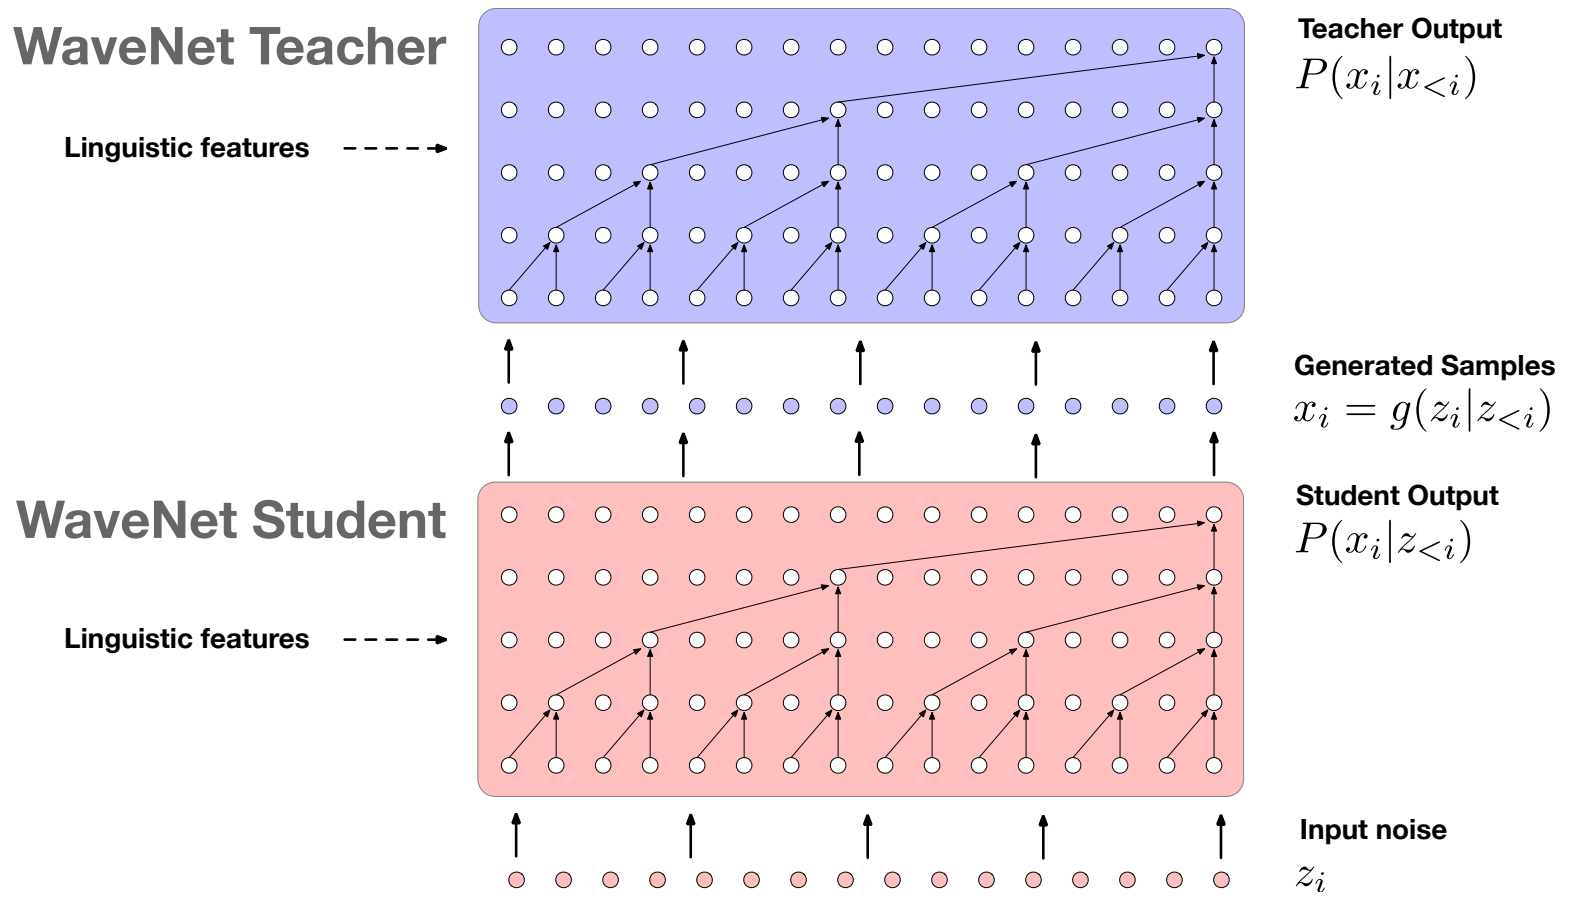


Figure 2 of paper "Parallel WaveNet: Fast High-Fidelity Speech Synthesis", <https://arxiv.org/abs/1711.10433>.

Denoting the teacher distribution as P_T and the student distribution as P_S , the loss is specifically

$$D_{\text{KL}}(P_S || P_T) = H(P_S, P_T) - H(P_S).$$

Therefore, we do not only minimize cross-entropy, but we also try to keep the entropy of the student as high as possible. That is crucial not to match just the mode of the teacher. (Consider a teacher generating white noise, where every sample comes from $\mathcal{N}(0, 1)$ – in this case, the cross-entropy loss of a constant $\mathbf{0}$, complete silence, would be maximal.)

In a sense, probability density distillation is similar to GANs. However, the teacher is kept fixed and the student does not attempt to fool it but to match its distribution instead.

With the 4 iterations, the Parallel WaveNet generates over 500k samples per second, compared to ~170 samples per second of a regular WaveNet – more than a 1000 times speedup.

Method	Subjective 5-scale MOS
16kHz, 8-bit μ-law, 25h data:	
LSTM-RNN parametric [27]	3.67 ± 0.098
HMM-driven concatenative [27]	3.86 ± 0.137
WaveNet [27]	4.21 ± 0.081
24kHz, 16-bit linear PCM, 65h data:	
HMM-driven concatenative	4.19 ± 0.097
Autoregressive WaveNet	4.41 ± 0.069
Distilled WaveNet	4.41 ± 0.078

Table 1 of paper "Parallel WaveNet: Fast High-Fidelity Speech Synthesis", <https://arxiv.org/abs/1711.10433>.

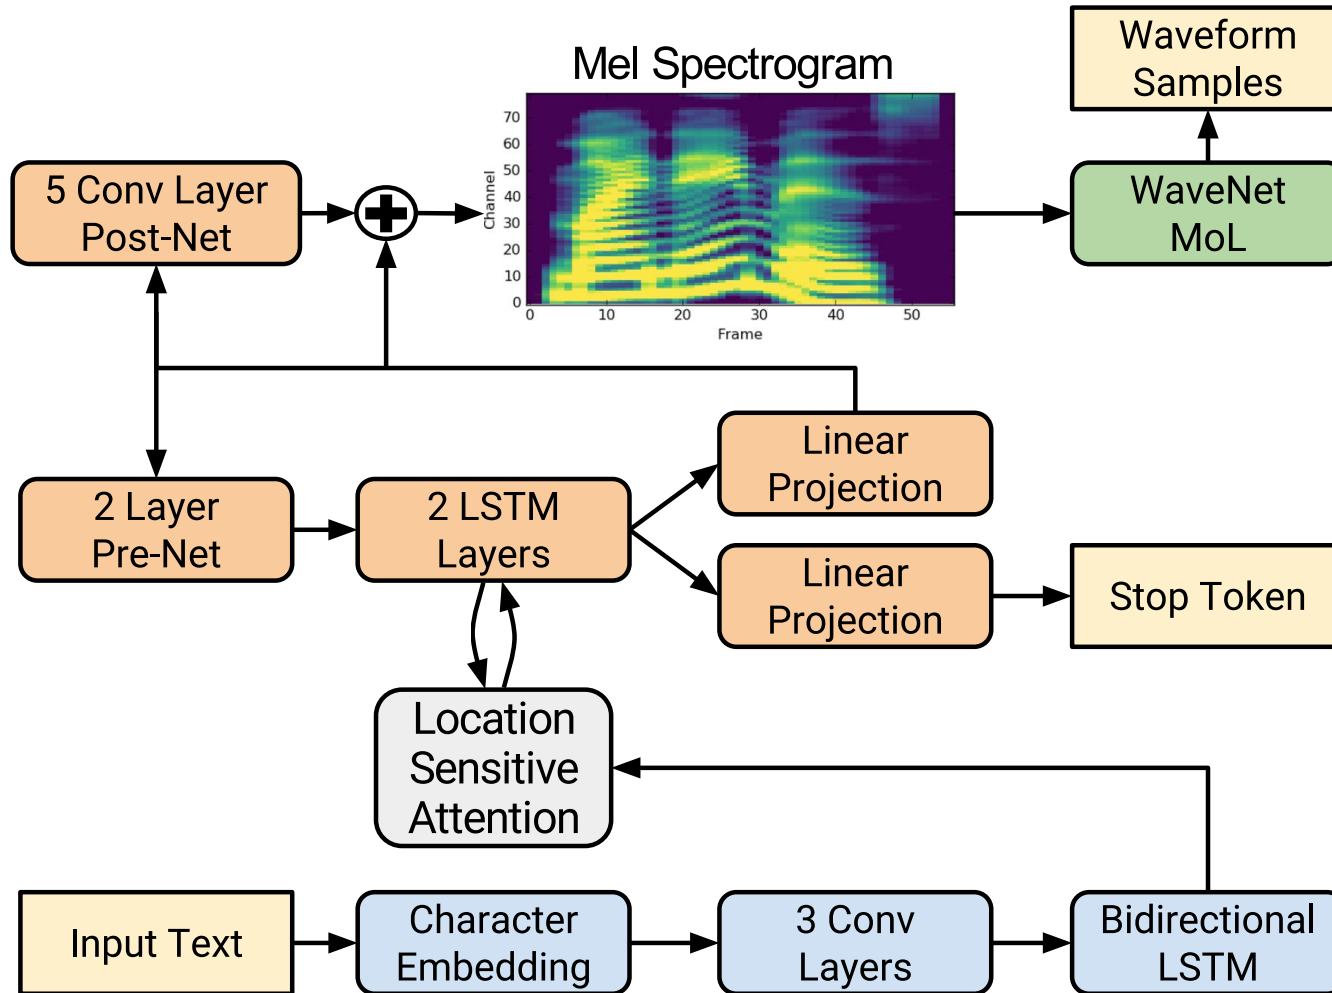


Figure 1 of paper "Natural TTS Synthesis by Conditioning WaveNet on Mel Spectrogram Predictions", <https://arxiv.org/abs/1712.05884>.

System	MOS
Parametric	3.492 \pm 0.096
Tacotron (Griffin-Lim)	4.001 \pm 0.087
Concatenative	4.166 \pm 0.091
WaveNet (Linguistic)	4.341 \pm 0.051
Ground truth	4.582 \pm 0.053
Tacotron 2 (this paper)	4.526 \pm 0.066

Table 1 of paper "Natural TTS Synthesis by Conditioning WaveNet on Mel Spectrogram Predictions", <https://arxiv.org/abs/1712.05884>.




Figure 2 of paper "Natural TTS Synthesis by Conditioning WaveNet on Mel Spectrogram Predictions", <https://arxiv.org/abs/1712.05884>.

So far, all input information was stored either directly in network weights, or in a state of a recurrent network.

However, mammal brains seem to operate with a *working memory* – a capacity for short-term storage of information and its rule-based manipulation.

We can therefore try to introduce an external memory to a neural network. The memory \mathbf{M} will be a matrix, where rows correspond to memory cells.

Neural Turing Machines

The network will control the memory using a controller which reads from the memory and writes to it. Although the original paper also considered a feed-forward (non-recurrent) controller, usually the controller is a recurrent LSTM network.

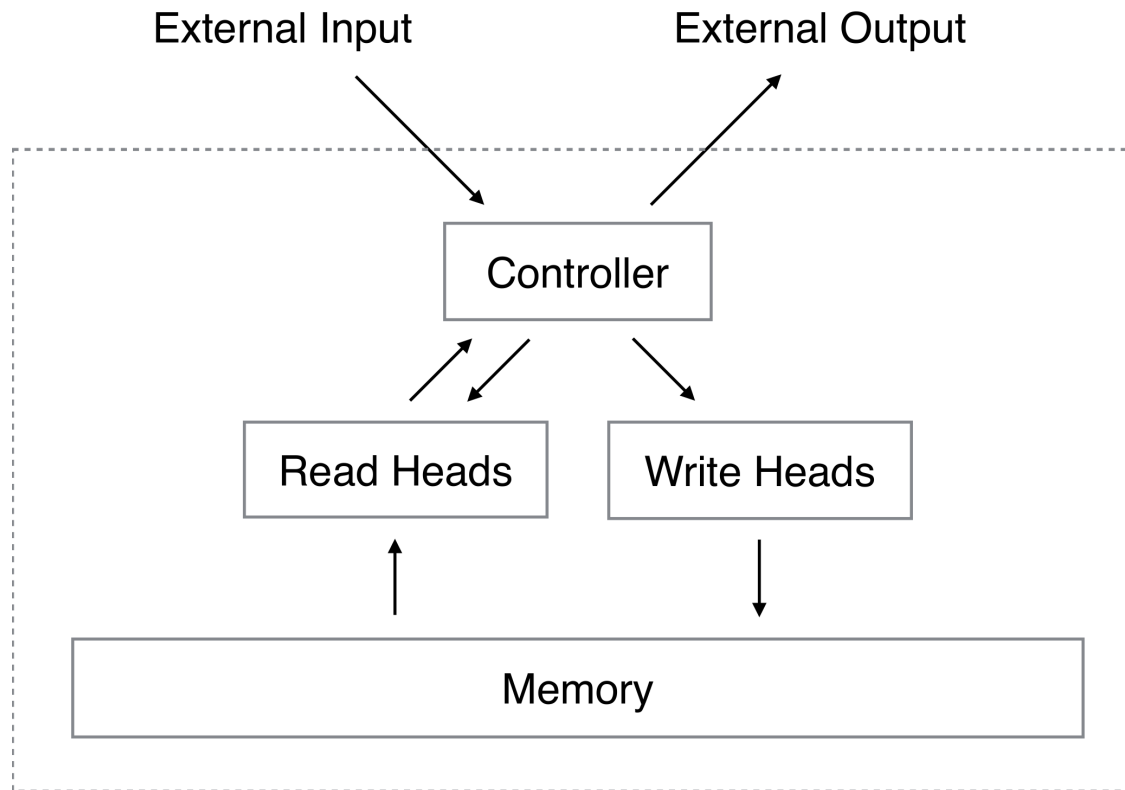


Figure 1 of paper "Neural Turing Machines", <https://arxiv.org/abs/1410.5401>.

Reading

To read the memory in a differentiable way, the controller at time t emits a read distribution \mathbf{w}_t over memory locations, and the returned read vector \mathbf{r}_t is then

$$\mathbf{r}_t = \sum_i w_t(i) \cdot \mathbf{M}_t(i).$$

Writing

Writing is performed in two steps – an *erase* followed by an *add*. The controller at time t emits a write distribution \mathbf{w}_t over memory locations, and also an *erase vector* \mathbf{e}_t and an *add vector* \mathbf{a}_t . The memory is then updated as

$$\mathbf{M}_t(i) = \mathbf{M}_{t-1}(i) [1 - w_t(i)\mathbf{e}_t] + w_t(i)\mathbf{a}_t.$$

Neural Turing Machine

The addressing mechanism is designed to allow both

- content addressing, and
- location addressing.

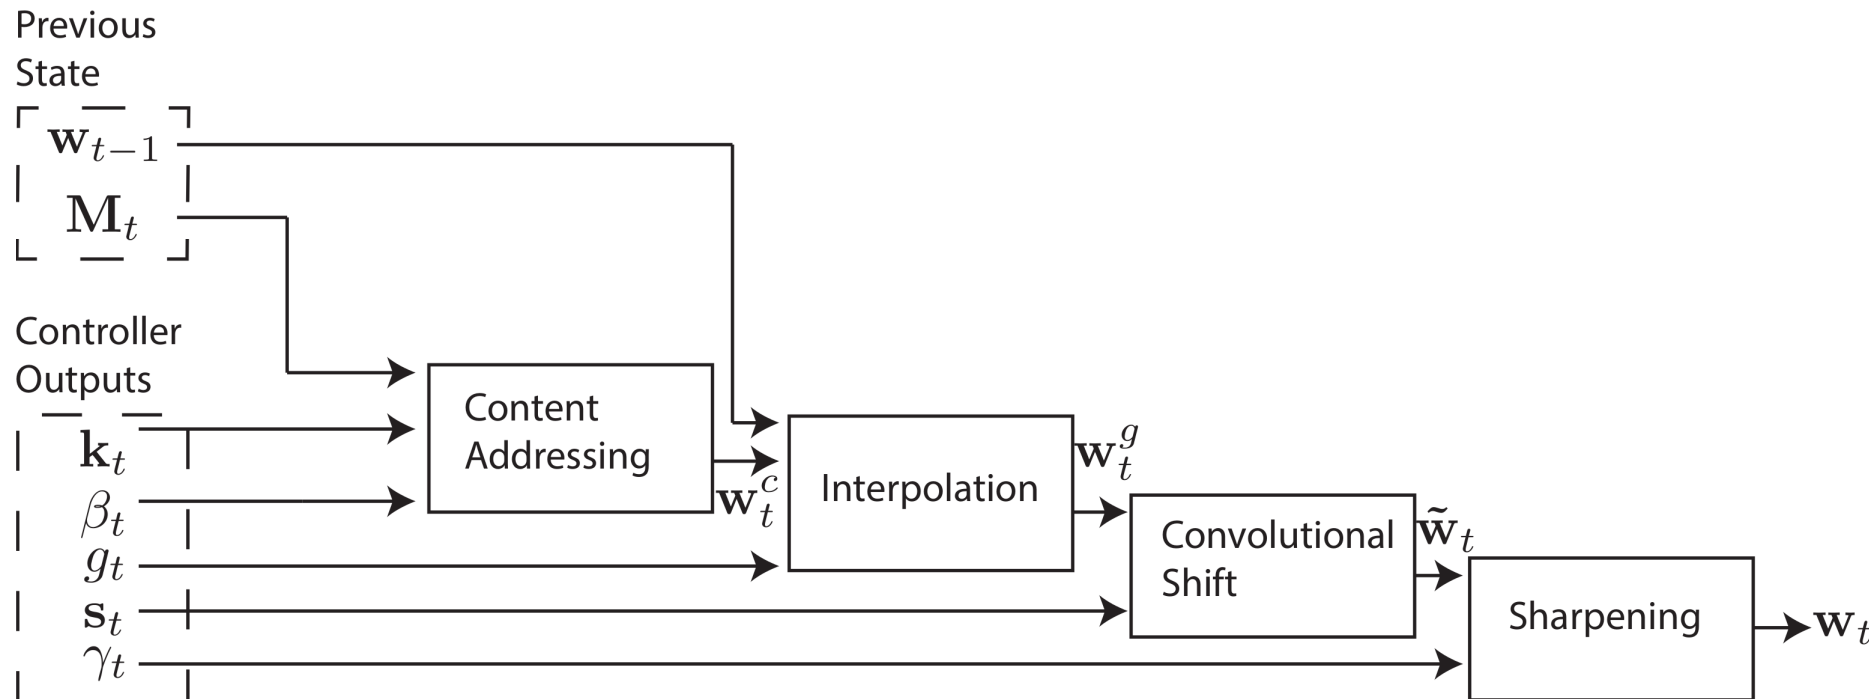


Figure 2 of paper "Neural Turing Machines", <https://arxiv.org/abs/1410.5401>.

Content Addressing

Content addressing starts by the controller emitting the *key vector* \mathbf{k}_t , which is compared to all memory locations $\mathbf{M}_t(i)$, generating a distribution using a softmax with temperature β_t .

$$w_t^c(i) = \frac{\exp(\beta_t \cdot \text{distance}(\mathbf{k}_t, \mathbf{M}_t(i)))}{\sum_j \exp(\beta_t \cdot \text{distance}(\mathbf{k}_t, \mathbf{M}_t(j)))}$$

The distance measure is usually the cosine similarity

$$\text{distance}(\mathbf{a}, \mathbf{b}) = \frac{\mathbf{a} \cdot \mathbf{b}}{\|\mathbf{a}\| \cdot \|\mathbf{b}\|}.$$

Location-Based Addressing

To allow iterative access to memory, the controller might decide to reuse the memory location from the previous timestep. Specifically, the controller emits *interpolation gate* g_t and defines

$$\mathbf{w}_t^g = g_t \mathbf{w}_t^c + (1 - g_t) \mathbf{w}_{t-1}.$$

Then, the current weighting may be shifted, i.e., the controller might decide to “rotate” the weights by a small integer. For a given range (the simplest case are only shifts $\{-1, 0, 1\}$), the network emits softmax distribution over the shifts, and the weights are then defined using a circular convolution

$$\tilde{w}_t(i) = \sum_j w_t^g(j) s_t(i - j).$$

Finally, not to lose precision over time, the controller emits a *sharpening factor* γ_t and the final memory location weights are $w_t(i) = \tilde{w}_t(i)^{\gamma_t} / \sum_j \tilde{w}_t(j)^{\gamma_t}$.

Overall Execution

Even if not specified in the original paper, following the DNC paper, the LSTM controller can be implemented as a (potentially deep) LSTM. Assuming R read heads and one write head, the input is \mathbf{x}_t and R read vectors $\mathbf{r}_{t-1}^1, \dots, \mathbf{r}_{t-1}^R$ from the previous time step, the output of the controller are vectors $(\boldsymbol{\nu}_t, \boldsymbol{\xi}_t)$, and the final output is $\mathbf{y}_t = \boldsymbol{\nu}_t + W_r [\mathbf{r}_t^1, \dots, \mathbf{r}_t^R]$. The $\boldsymbol{\xi}_t$ is a concatenation of

$$\mathbf{k}_t^1, \beta_t^1, \mathbf{g}_t^1, \mathbf{s}_t^1, \gamma_t^1, \mathbf{k}_t^2, \beta_t^2, \mathbf{g}_t^2, \mathbf{s}_t^2, \gamma_t^2, \dots, \mathbf{k}_t^w, \beta_t^w, \mathbf{g}_t^w, \mathbf{s}_t^w, \gamma_t^w, \mathbf{e}_t^w, \mathbf{a}_t^w.$$

Copy Task

Repeat the same sequence as given on input. Trained with sequences of length up to 20.




Figure 3 of paper "Neural Turing Machines", <https://arxiv.org/abs/1410.5401>.




Figure 4: NTM Generalisation on the Copy Task. The four pairs of plots in the top row depict network outputs and corresponding copy targets for test sequences of length 10, 20, 30, and 50, respectively. The plots in the bottom row are for a length 120 sequence. The network was only trained on sequences of up to length 20. The first four sequences are reproduced with high confidence and very few mistakes. The longest one has a few more local errors and one global error: at the point indicated by the red arrow at the bottom, a single vector is duplicated, pushing all subsequent vectors one step back. Despite being subjectively close to a correct copy, this leads to a high loss.

Figure 4 of paper "Neural Turing Machines", <https://arxiv.org/abs/1410.5401>.




Figure 5: LSTM Generalisation on the Copy Task. The plots show inputs and outputs for the same sequence lengths as Figure 4. Like NTM, LSTM learns to reproduce sequences of up to length 20 almost perfectly. However it clearly fails to generalise to longer sequences. Also note that the length of the accurate prefix decreases as the sequence length increases, suggesting that the network has trouble retaining information for long periods.

Figure 5 of paper "Neural Turing Machines", <https://arxiv.org/abs/1410.5401>.

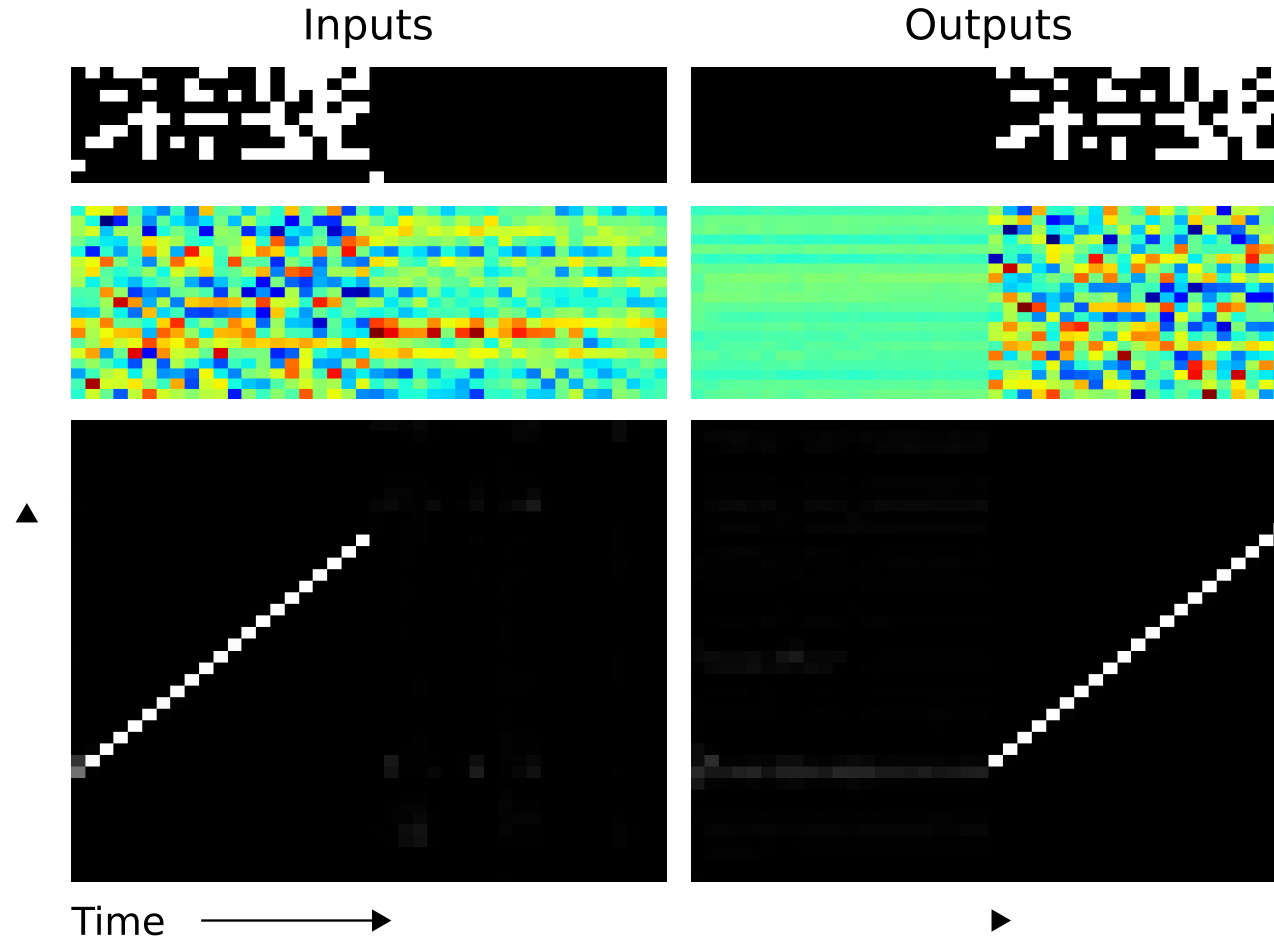


Figure 6 of paper "Neural Turing Machines", <https://arxiv.org/abs/1410.5401>.

Associative Recall

In associative recall, a sequence is given on input, consisting of subsequences of length 3. Then a randomly chosen subsequence is presented on input and the goal is to produce the following subsequence.

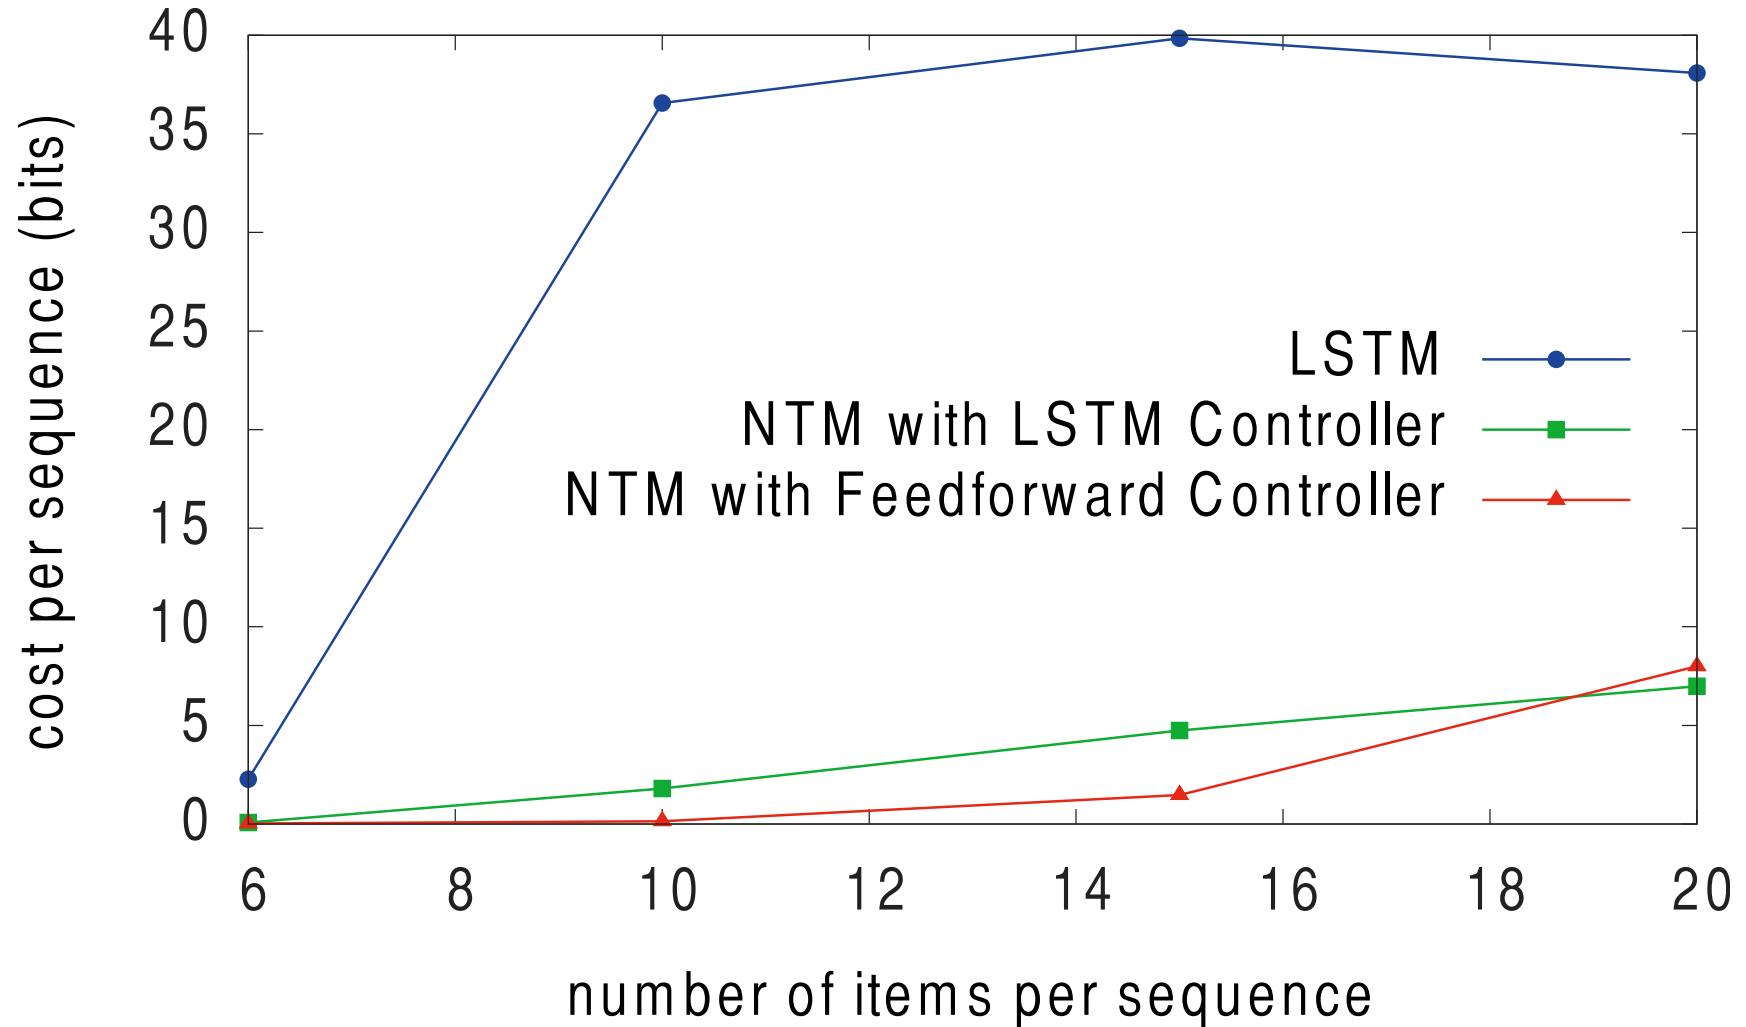


Figure 11 of paper "Neural Turing Machines", <https://arxiv.org/abs/1410.5401>.

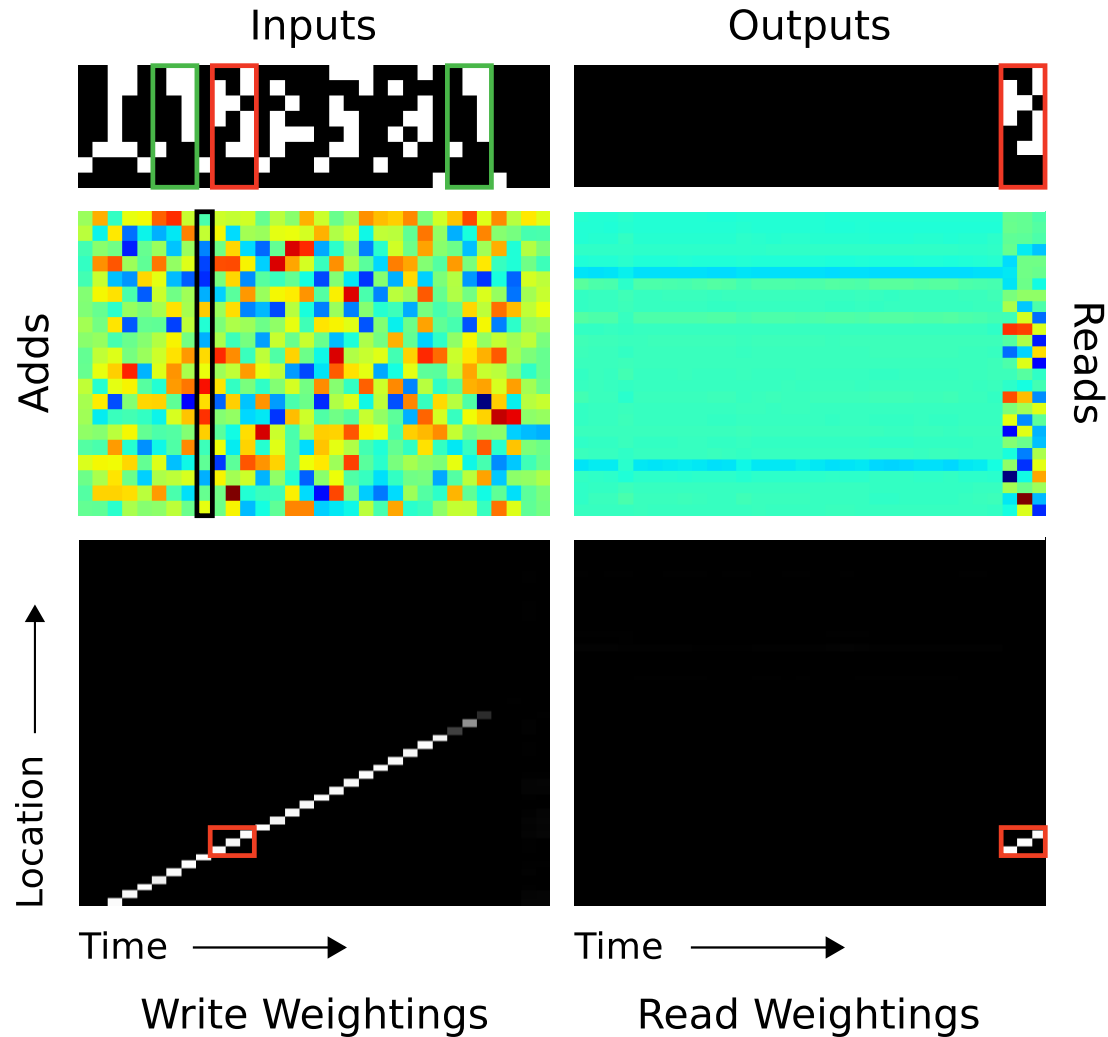


Figure 12 of paper "Neural Turing Machines", <https://arxiv.org/abs/1410.5401>.

Differentiable Neural Computer

NTM was later extended to a Differentiable Neural Computer.

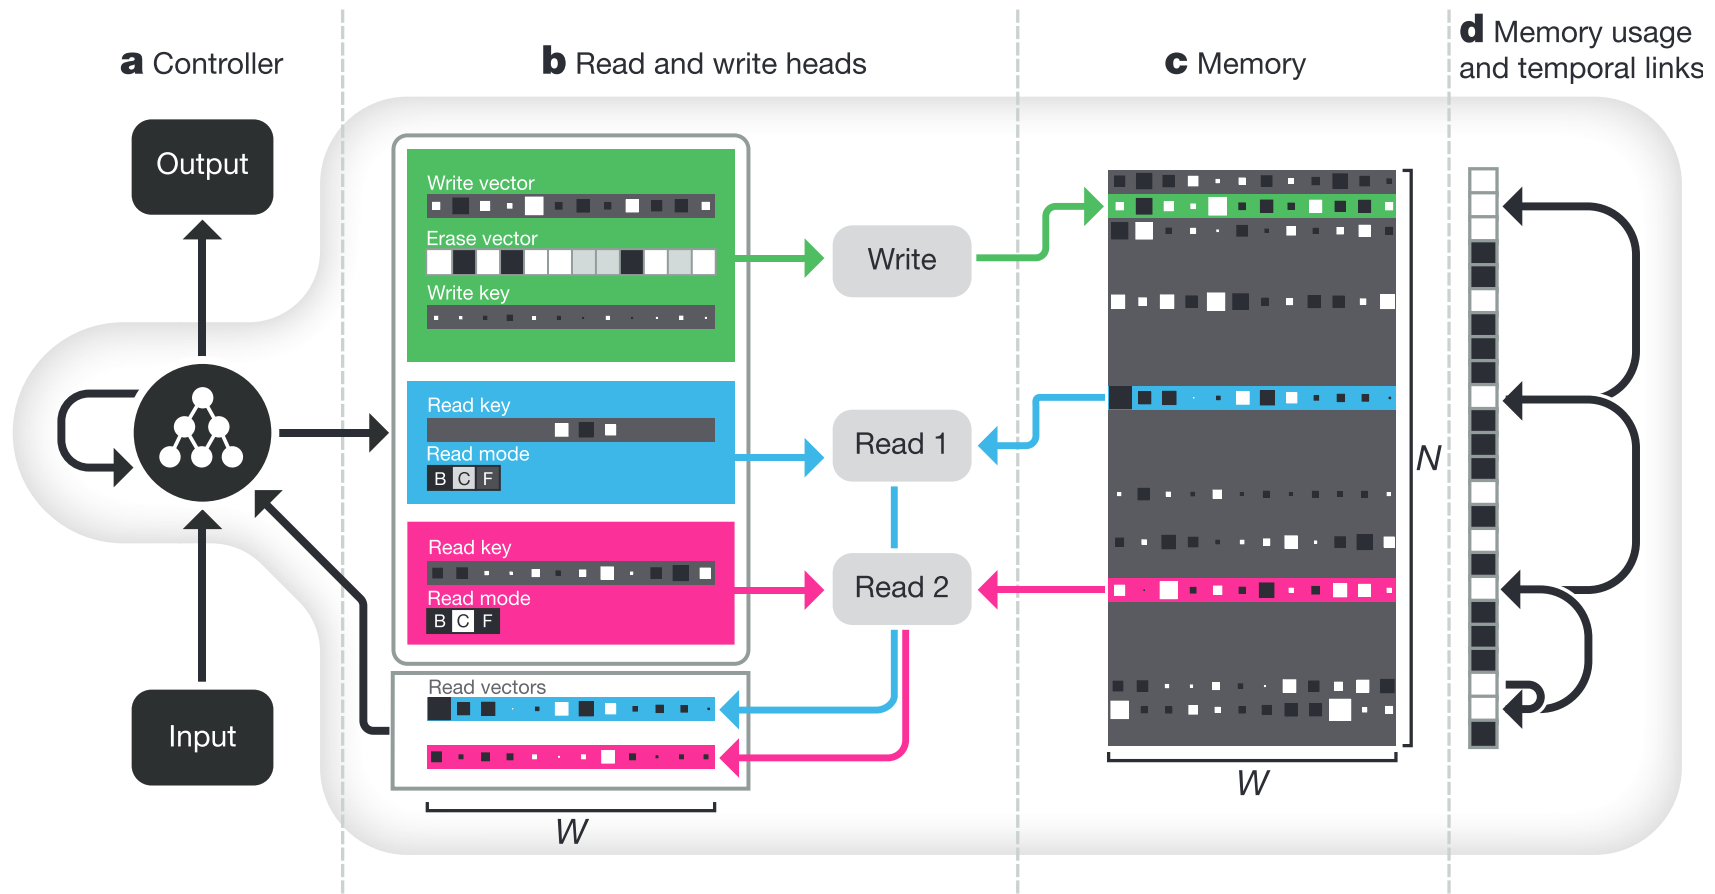


Figure 1 of paper "Hybrid computing using a neural network with dynamic external memory", <https://www.nature.com/articles/nature20101>.

The DNC contains multiple read heads and one write head.

The controller is a deep LSTM network, with input at time t being the current input \mathbf{x}_t and R read vectors $\mathbf{r}_{t-1}^1, \dots, \mathbf{r}_{t-1}^R$ from previous time step. The output of the controller are vectors $(\boldsymbol{\nu}_t, \boldsymbol{\xi}_t)$, and the final output is $\mathbf{y}_t = \boldsymbol{\nu}_t + W_r [\mathbf{r}_t^1, \dots, \mathbf{r}_t^R]$. The $\boldsymbol{\xi}_t$ is a concatenation of parameters for read and write heads (keys, gates, sharpening parameters, ...).


In DNC, the usage of every memory location is tracked, which enables performing dynamic allocation – at each time step, a cell with least usage can be allocated.

Furthermore, for every memory location, we track which memory location was written to previously (\mathbf{b}_t) and subsequently (\mathbf{f}_t), allowing to recover sequences in the order in which it was written, independently on the real indices used.

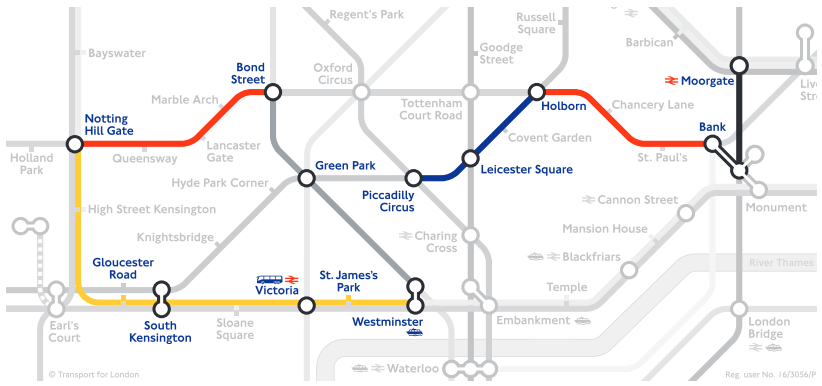
The write weighting is defined as a weighted combination of the allocation weighting and write content weighting, and read weighting is computed as a weighted combination of read content weighting, previous write weighting, and subsequent write weighting.

Differentiable Neural Computer

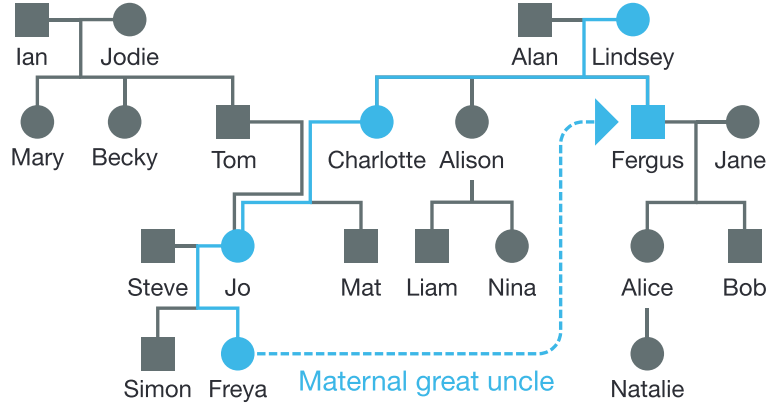
a Random graph



b London Underground



c Family tree



Underground input:
(OxfordCircus, TottenhamCtRd, Central)
(TottenhamCtRd, OxfordCircus, Central)
(BakerSt, Marylebone, Circle)
(BakerSt, Marylebone, Bakerloo)
(BakerSt, OxfordCircus, Bakerloo)
⋮
(LeicesterSq, CharingCross, Northern)
(TottenhamCtRd, LeicesterSq, Northern)
(OxfordCircus, PiccadillyCircus, Bakerloo)
(OxfordCircus, NottingHillGate, Central)
(OxfordCircus, Euston, Victoria)
84 edges in total

Traversal question:
(BondSt, _, Central),
(_, _, Circle), (_, _, Circle),
(_, _, Circle), (_, _, Circle),
(_, _, Jubilee), (_, _, Jubilee),

Answer:
(BondSt, NottingHillGate, Central)
(NottingHillGate, GloucesterRd, Circle)
⋮
(Westminster, GreenPark, Jubilee)
(GreenPark, BondSt, Jubilee)

Shortest-path question:
(Moorgate, PiccadillyCircus, _)

Answer:
(Moorgate, Bank, Northern)
(Bank, Holborn, Central)
(Holborn, LeicesterSq, Piccadilly)
(LeicesterSq, PiccadillyCircus, Piccadilly)

Family tree input:
(Charlotte, Alan, Father)
(Simon, Steve, Father)
(Steve, Simon, Son1)
(Nina, Alison, Mother)
(Lindsey, Fergus, Son1)
⋮
(Bob, Jane, Mother)
(Natalie, Alice, Mother)
(Mary, Ian, Father)
(Jane, Alice, Daughter1)
(Mat, Charlotte, Mother)
54 edges in total

Inference question:
(Freya, _, MaternalGreatUncle)

Answer:
(Freya, Fergus, MaternalGreatUncle)

Figure 2 of paper "Hybrid computing using a neural network with dynamic external memory", <https://www.nature.com/articles/nature20101>.

Memory-augmented Neural Networks

External memory can be also utilized for *learning to learn*. Consider a network, which should learn classification into a user-defined hierarchy by observing ideally a small number of samples. Apart from finetuning the model and storing the information in the *weights*, an alternative is to store the samples in *external memory*. Therefore, the model learns how to store the data and access it efficiently, which allows it to learn without changing its weights.




Figure 1 of paper "One-shot learning with Memory-Augmented Neural Networks", <https://arxiv.org/abs/1605.06065>.

$$K(\mathbf{k}_t, \mathbf{M}_t(i)) = \frac{\mathbf{k}_t \cdot \mathbf{M}_t(i)}{\|\mathbf{k}_t\| \|\mathbf{M}_t(i)\|}, \quad (2)$$

which is used to produce a read-weight vector, \mathbf{w}_t^r , with elements computed according to a softmax:

$$w_t^r(i) \leftarrow \frac{\exp(K(\mathbf{k}_t, \mathbf{M}_t(i)))}{\sum_j \exp(K(\mathbf{k}_t, \mathbf{M}_t(j)))}. \quad (3)$$

A memory, \mathbf{r}_t , is retrieved using this weight vector:

$$\mathbf{r}_t \leftarrow \sum_i w_t^r(i) \mathbf{M}_t(i). \quad (4)$$

Page 3 of paper "One-shot learning with Memory-Augmented Neural Networks", <https://arxiv.org/abs/1605.06065>.

$$\mathbf{w}_t^u \leftarrow \gamma \mathbf{w}_{t-1}^u + \mathbf{w}_t^r + \mathbf{w}_t^w. \quad (5)$$

Here, γ is a decay parameter and \mathbf{w}_t^r is computed as in (3). The *least-used* weights, \mathbf{w}_t^{lu} , for a given time-step can then be computed using \mathbf{w}_t^u . First, we introduce the notation $m(\mathbf{v}, n)$ to denote the n^{th} smallest element of the vector \mathbf{v} . Elements of \mathbf{w}_t^{lu} are set accordingly:

$$w_t^{lu}(i) = \begin{cases} 0 & \text{if } w_t^u(i) > m(\mathbf{w}_t^u, n) \\ 1 & \text{if } w_t^u(i) \leq m(\mathbf{w}_t^u, n) \end{cases}, \quad (6)$$

where n is set to equal the number of reads to memory. To obtain the write weights \mathbf{w}_t^w , a learnable sigmoid gate parameter is used to compute a convex combination of the previous read weights and previous least-used weights:

$$\mathbf{w}_t^w \leftarrow \sigma(\alpha) \mathbf{w}_{t-1}^r + (1 - \sigma(\alpha)) \mathbf{w}_{t-1}^{lu}. \quad (7)$$

Here, $\sigma(\cdot)$ is a sigmoid function, $\frac{1}{1+e^{-x}}$, and α is a scalar gate parameter to interpolate between the weights. Prior to writing to memory, the least used memory location is computed from \mathbf{w}_{t-1}^u and is set to zero. Writing to memory then occurs in accordance with the computed vector of write weights:

$$\mathbf{M}_t(i) \leftarrow \mathbf{M}_{t-1}(i) + w_t^w(i) \mathbf{k}_t, \forall i \quad (8)$$

Page 4 of paper "One-shot learning with Memory-Augmented Neural Networks", <https://arxiv.org/abs/1605.06065>.

Memory-augmented NNs

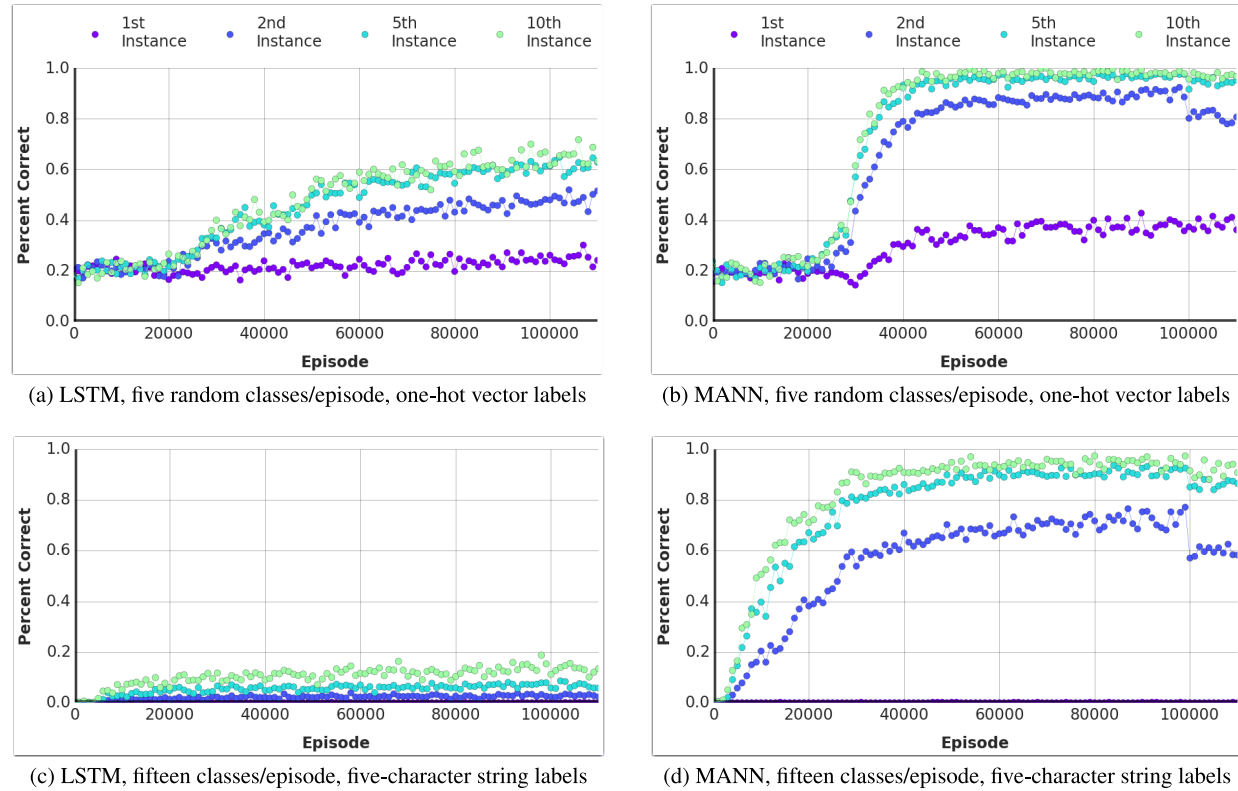


Figure 2. Omniglot classification. The network was given either five (a-b) or up to fifteen (c-d) random classes per episode, which were of length 50 or 100 respectively. Labels were one-hot vectors in (a-b), and five-character strings in (c-d). In (b), first instance accuracy is above chance, indicating that the MANN is performing “educated guesses” for new classes based on the classes it has already seen and stored in memory. In (c-d), first instance accuracy is poor, as is expected, since it must make a guess from 3125 random strings. Second instance accuracy, however, approaches 80% during training for the MANN (d). At the 100,000 episode mark the network was tested, without further learning, on distinct classes withheld from the training set, and exhibited comparable performance.

Figure 2 of paper "One-shot learning with Memory-Augmented Neural Networks", <https://arxiv.org/abs/1605.06065>.