

Simple Recurrent Units for Highly Parallelizable Recurrence

Tao Lei¹ Yu Zhang² Sida I. Wang^{1,3} Hui Dai¹ Yoav Artzi^{1,4}
¹ASAPP Inc. ²Google Brain ³Princeton University ⁴Cornell University
¹{tao, hd}@asapp.com ²ngyuzh@google.com
³sidaw@cs.princeton.edu ⁴yoav@cs.cornell.edu

Abstract

Common recurrent neural architectures scale poorly due to the intrinsic difficulty in parallelizing their state computations. In this work, we propose the Simple Recurrent Unit (SRU), a light recurrent unit that balances model capacity and scalability. SRU is designed to provide expressive recurrence, enable highly parallelized implementation, and comes with careful initialization to facilitate training of deep models. We demonstrate the effectiveness of SRU on multiple NLP tasks. SRU achieves 5–9x speed-up over cuDNN-optimized LSTM on classification and question answering datasets, and delivers stronger results than LSTM and convolutional models. We also obtain an average of 0.7 BLEU improvement over the Transformer model (Vaswani et al., 2017) on translation by incorporating SRU into the architecture.¹

1 Introduction

Recurrent neural networks (RNN) are at the core of state-of-the-art approaches for a large number of natural language tasks, including machine translation (Cho et al., 2014; Bahdanau et al., 2015; Jean et al., 2015; Luong et al., 2015), language modeling (Zaremba et al., 2014; Gal and Ghahramani, 2016; Zoph and Le, 2016), opinion mining (Irsoy and Cardie, 2014), and situated language understanding (Mei et al., 2016; Misra et al., 2017; Suhr et al., 2018; Suhr and Artzi, 2018). Key to many of these advancements are architectures of increased capacity and computation. For instance, the top-performing models for semantic role labeling and translation use eight recurrent layers, requiring days to train (He et al., 2017; Wu et al., 2016b). The scalability of these models has become an important problem that impedes NLP research.

¹Our code is available at <https://github.com/taolei87/sru>.

The difficulty of scaling recurrent networks arises from the time dependence of state computation. In common architectures, such as Long Short-term Memory (LSTM; Hochreiter and Schmidhuber, 1997) and Gated Recurrent Units (GRU; Cho et al., 2014), the computation of each step is suspended until the complete execution of the previous step. This sequential dependency makes recurrent networks significantly slower than other operations, and limits their applicability. For example, recent translation models consist of non-recurrent components only, such as attention and convolution, to scale model training (Gehring et al., 2017; Vaswani et al., 2017).

In this work, we introduce the Simple Recurrent Unit (SRU), a unit with light recurrence that offers both high parallelization and sequence modeling capacity. The design of SRU is inspired by previous efforts, such as Quasi-RNN (QRNN; Bradbury et al., 2017) and Kernel NN (KNN; Lei et al., 2017), but enjoys additional benefits:

- SRU exhibits the same level of parallelism as convolution and feed-forward nets. This is achieved by balancing sequential dependence and independence: while the state computation of SRU is time-dependent, each state dimension is independent. This simplification enables CUDA-level optimizations that parallelize the computation across hidden dimensions and time steps, effectively using the full capacity of modern GPUs. Figure 1 compares our architecture’s runtimes to common architectures.
- SRU replaces the use of convolutions (i.e., n-gram filters), as in QRNN and KNN, with more recurrent connections. This retains modeling capacity, while using less computation (and hyper-parameters).




Figure 1: Average processing time in milliseconds of a batch of 32 samples using cuDNN LSTM, word-level convolution `conv2d` (with filter width $k = 2$ and $k = 3$), and the proposed SRU. We vary the number of tokens per sequence (l) and feature dimension (d).

- SRU improves the training of deep recurrent models by employing highway connections (Srivastava et al., 2015) and a parameter initialization scheme tailored for gradient propagation in deep architectures.

We evaluate SRU on a broad set of problems, including text classification, question answering, translation and character-level language modeling. Our experiments demonstrate that light recurrence is sufficient for various natural language tasks, offering a good trade-off between scalability and representational power. On classification and question answering datasets, SRU outperforms common recurrent and non-recurrent architectures, while achieving 5–9x speed-up compared to cuDNN LSTM. Stacking additional layers further improves performance, while incurring relatively small costs owing to the cheap computation of a single layer. We also obtain an average improvement of 0.7 BLEU score on the English to German translation task by incorporating SRU into Transformer (Vaswani et al., 2017).

2 Related Work

Improving on common architectures for sequence processing has recently received significant attention (Greff et al., 2017; Balduzzi and Ghifary, 2016; Miao et al., 2016; Zoph and Le, 2016; Lee et al., 2017). One area of research involves incorporating word-level convolutions (i.e. n-gram filters) into recurrent computation (Lei et al., 2015; Bradbury et al., 2017; Lei et al., 2017). For example, Quasi-RNN (Bradbury et al., 2017) proposes to alternate convolutions and a minimalist recurrent pooling function and achieves significant speed-up over LSTM. While Bradbury et al. (2017) focus on the speed advantages of the network, Lei et al. (2017) study the theoret-

ical characteristics of such computation and possible extensions. Their results suggest that simplified recurrence retains strong modeling capacity through layer stacking. This finding motivates the design of SRU for both high parallelization and representational power. SRU also relates to IRNN (Le et al., 2015), which uses an identity diagonal matrix to initialize hidden-to-hidden connections. SRU uses point-wise multiplication for hidden connections, which is equivalent to using a diagonal weight matrix. This can be seen as a constrained version of diagonal initialization.

Various strategies have been proposed to scale network training (Goyal et al., 2017) and to speed up recurrent networks (Diamos et al., 2016; Shazeer et al., 2017; Kuchaiev and Ginsburg, 2017). For instance, Diamos et al. (2016) utilize hardware infrastructures by stashing RNN parameters on cache (or fast memory). Shazeer et al. (2017) and Kuchaiev and Ginsburg (2017) improve the computation via conditional computing and matrix factorization respectively. Our implementation for SRU is inspired by the cuDNN-optimized LSTM (Appleyard et al., 2016), but enables more parallelism – while cuDNN LSTM requires six optimization steps, SRU achieves more significant speed-up via two optimizations.

The design of recurrent networks, such as SRU and related architectures, raises questions about representational power and interpretability (Chen et al., 2018; Peng et al., 2018). Balduzzi and Ghifary (2016) applies type-preserving transformations to discuss the capacity of various simplified RNN architectures. Recent work (Anselmi et al., 2015; Daniely et al., 2016; Zhang et al., 2016; Lei et al., 2017) relates the capacity of neural networks to deep kernels. We empirically demonstrate SRU can achieve compelling results by stacking multiple layers.

3 Simple Recurrent Unit

We present and explain the design of Simple Recurrent Unit (SRU) in this section. A single layer of SRU involves the following computation:

$$\mathbf{f}_t = \sigma(\mathbf{W}_f \mathbf{x}_t + \mathbf{v}_f \odot \mathbf{c}_{t-1} + \mathbf{b}_f) \quad (1)$$

$$\mathbf{c}_t = \mathbf{f}_t \odot \mathbf{c}_{t-1} + (1 - \mathbf{f}_t) \odot (\mathbf{W} \mathbf{x}_t) \quad (2)$$

$$\mathbf{r}_t = \sigma(\mathbf{W}_r \mathbf{x}_t + \mathbf{v}_r \odot \mathbf{c}_{t-1} + \mathbf{b}_r) \quad (3)$$

$$\mathbf{h}_t = \mathbf{r}_t \odot \mathbf{c}_t + (1 - \mathbf{r}_t) \odot \mathbf{x}_t \quad (4)$$

where \mathbf{W} , \mathbf{W}_f and \mathbf{W}_r are parameter matrices and \mathbf{v}_f , \mathbf{v}_r , \mathbf{b}_f and \mathbf{b}_v are parameter vectors to be learnt during training. The complete architecture decomposes to two sub-components: a *light recurrence* (Equation 1 and 2) and a *highway network* (Equation 3 and 4).

The light recurrence component successively reads the input vectors \mathbf{x}_t and computes the sequence of states \mathbf{c}_t capturing sequential information. The computation resembles other recurrent networks such as LSTM, GRU and RAN (Lee et al., 2017). Specifically, a forget gate \mathbf{f}_t controls the information flow (Equation 1) and the state vector \mathbf{c}_t is determined by adaptively averaging the previous state \mathbf{c}_{t-1} and the current observation $\mathbf{W} \mathbf{x}_t$ according to \mathbf{f}_t (Equation 2).

One key design decision that differs from previous gated recurrent architectures is the way \mathbf{c}_{t-1} is used in the sigmoid gate. Typically, \mathbf{c}_{t-1} is multiplied with a parameter matrix to compute \mathbf{f}_t , e.g., $\mathbf{f}_t = \sigma(\mathbf{W}_f \mathbf{x}_t + \mathbf{V}_f \mathbf{c}_{t-1} + \mathbf{b}_f)$. However, the inclusion of $\mathbf{V}_f \mathbf{c}_{t-1}$ makes it difficult to parallelize the state computation: each dimension of \mathbf{c}_t and \mathbf{f}_t depends on all entries of \mathbf{c}_{t-1} , and the computation has to wait until \mathbf{c}_{t-1} is fully computed. To facilitate parallelization, our light recurrence component uses a point-wise multiplication $\mathbf{v}_f \odot \mathbf{c}_{t-1}$ instead. With this simplification, each dimension of the state vectors becomes independent and hence parallelizable.

The highway network component (Srivastava et al., 2015) facilitates gradient-based training of deep networks. It uses the reset gate \mathbf{r}_t (Equation 3) to adaptively combine the input \mathbf{x}_t and the state \mathbf{c}_t produced from the light recurrence (Equation 4), where $(1 - \mathbf{r}_t) \odot \mathbf{x}_t$ is a skip connection that allows the gradient to directly propagate to the previous layer. Such connections have been shown to improve scalability (Wu et al., 2016a; Kim et al., 2016; He et al., 2016; Zilly et al., 2017).

The combination of the two components makes the overall architecture simple yet expressive, and easy to scale due to enhanced parallelization and gradient propagation.

3.1 Parallelized Implementation

Despite the parallelization friendly design of SRU, a naive implementation which computes equations (1)–(4) for each step t sequentially would not achieve SRU’s full potential. We employ two optimizations to enhance parallelism. The optimizations are performed in the context of GPU / CUDA programming, but the general idea can be applied to other parallel programming models.

We re-organize the computation of equations (1)–(4) into two major steps. First, given the input sequence $\{\mathbf{x}_1 \cdots \mathbf{x}_L\}$, we batch the matrix multiplications across all time steps. This significantly improves the computation intensity (e.g. GPU utilization). The batched multiplication is:

$$\mathbf{U}^\top = \begin{pmatrix} \mathbf{W} \\ \mathbf{W}_f \\ \mathbf{W}_r \end{pmatrix} [\mathbf{x}_1, \mathbf{x}_2, \cdots, \mathbf{x}_L] ,$$

where L is the sequence length, $\mathbf{U} \in \mathbb{R}^{L \times 3d}$ is the computed matrix and d is the hidden state size. When the input is a mini-batch of B sequences, \mathbf{U} would be a tensor of size $(L, B, 3d)$.

The second step computes the remaining point-wise operations. Specifically, we compile all point-wise operations into a single fused CUDA kernel and parallelize the computation across each dimension of the hidden state. Algorithm 1 shows the pseudo code of the forward function. The complexity of this step is $O(L \cdot B \cdot d)$ per layer, where L is the sequence length and B is the batch size. In contrast, the complexity of LSTM is $O(L \cdot B \cdot d^2)$ because of the hidden-to-hidden multiplications (e.g. $\mathbf{V} \mathbf{h}_{t-1}$), and each dimension can not be independently parallelized. The fused kernel also reduces overhead. Without it, operations such as sigmoid activation would each invoke a separate function call, adding kernel launching latency and more data moving costs.

The implementation of a bidirectional SRU is similar: the matrix multiplications of both directions are batched, and the fused kernel handles and parallelizes both directions at the same time.

3.2 Initialization

Proper parameter initialization can reduce gradient propagation difficulties and hence have a positive

Algorithm 1 Mini-batch version of the forward pass defined in Equations (1)–(4).

Indices: Sequence length L , mini-batch size B , hidden state dimension d .

Input: Input sequences batch $\mathbf{x}[l, i, j]$; grouped matrix multiplication $\mathbf{U}[l, i, j']$; initial state $\mathbf{c}_0[i, j]$; parameters $\mathbf{v}_f[j]$, $\mathbf{v}_r[j]$, $\mathbf{b}_f[j]$ and $\mathbf{b}_r[j]$.

Output: Output $\mathbf{h}[\cdot, \cdot, \cdot]$ and internal $\mathbf{c}[\cdot, \cdot, \cdot]$ states.

Initialize $\mathbf{h}[\cdot, \cdot, \cdot]$ and $\mathbf{c}[\cdot, \cdot, \cdot]$ as two $L \times B \times d$ tensors.

for $i = 1, \dots, B; j = 1, \dots, d$ **do** // Parallelize each example i and dimension j

$c = \mathbf{c}_0[i, j]$

for $l = 1, \dots, L$ **do**

$f = \sigma(\mathbf{U}[l, i, j + d] + \mathbf{v}_f[j] \times c + \mathbf{b}_f[j])$

$c = f \times c + (1 - f) \times \mathbf{U}[l, i, j]$

$r = \sigma(\mathbf{U}[l, i, j + d \times 2] + \mathbf{v}_r[j] \times c + \mathbf{b}_r[j])$

$h = r \times c + (1 - r) \times \mathbf{x}[l, i, j]$

$\mathbf{c}[l, i, j] = c$

$\mathbf{h}[l, i, j] = h$

return $\mathbf{h}[\cdot, \cdot, \cdot]$ and $\mathbf{c}[\cdot, \cdot, \cdot]$

impact on the final performance. We now describe an initialization strategy tailored for SRU.

We start by adopting common initializations derived for feed-forward networks (Glorot and Bengio, 2010; He et al., 2015). The weights of parameter matrices are drawn with zero mean and $1/d$ variance, for instance, via the uniform distribution $[-\sqrt{3/d}, +\sqrt{3/d}]$. This ensures the output variance remains approximately the same as the input variance after the matrix multiplication.

However, the light recurrence and highway computation would still reduce the variance of hidden representations by a factor of $1/3$ to $1/2$:

$$\frac{1}{3} \leq \frac{\text{Var}[\mathbf{h}_t]}{\text{Var}[\mathbf{x}_t]} \leq \frac{1}{2},$$

and the factor converges to $1/2$ in deeper layers (see Appendix A). This implies the output \mathbf{h}_t and the gradient would vanish in deep models. To offset the problem, we introduce a *scaling correction* constant α in the highway connection

$$\mathbf{h}_t = \mathbf{r}_t \odot \mathbf{c}_t + (1 - \mathbf{r}_t) \odot \mathbf{x}_t \cdot \alpha,$$

where α is set to $\sqrt{3}$ such that $\text{Var}[\mathbf{h}_t] \approx \text{Var}[\mathbf{x}_t]$ at initialization. When the highway network is initialized with a non-zero bias $\mathbf{b}_r = b$, the scaling constant α can be accordingly set as:

$$\alpha = \sqrt{1 + \exp(b) \times 2}.$$

Figure 2 compares the training progress with and without the scaling correction. See Appendix A for the derivation and more discussion.

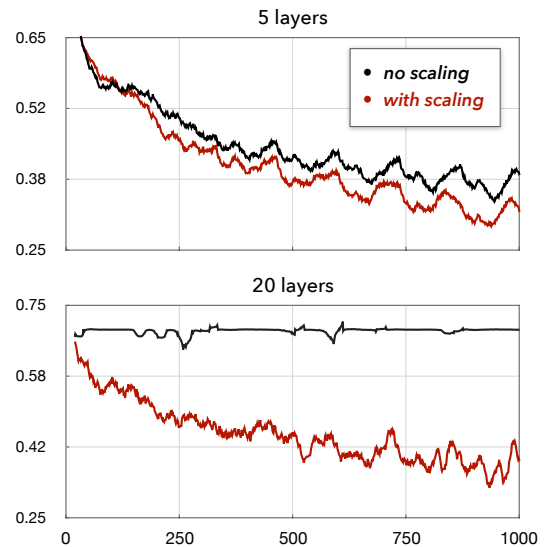


Figure 2: Training curves of SRU on classification. The x-axis is the number of training steps and the y-axis is the training loss. Scaling correction improves the training progress, especially for deeper models with many stacked layers.

4 Experiments

We evaluate SRU on several natural language processing tasks and perform additional analyses of the model. The set of tasks includes text classification, question answering, machine translation, and character-level language modeling. Training time on these benchmarks ranges from minutes (classification) to days (translation), providing a variety of computation challenges.

The main question we study is the performance-speed trade-off SRU provides in comparison to

Model	Size	CR	SUBJ	MR	TREC	MPQA	SST	Time
Best reported results:								
Wang and Manning (2013)		82.1	93.6	79.1	-	86.3	-	-
Kalchbrenner et al. (2014)		-	-	-	93.0	-	86.8	-
Kim (2014)		85.0	93.4	81.5	93.6	89.6	88.1	-
Zhang and Wallace (2017)		84.7	93.7	81.7	91.6	89.6	85.5	-
Zhao et al. (2015)		86.3	95.5	83.1	92.4	93.3	-	-
Our setup (default Adam, fixed word embeddings):								
CNN	360k	83.1±1.6	92.7±0.9	78.9±1.3	93.2±0.8	89.2±0.8	85.1±0.6	417
LSTM	352k	82.7±1.9	92.6±0.8	79.8±1.3	93.4±0.9	89.4±0.7	88.1±0.8	2409
QRNN (k=1)	165k	83.5±1.9	93.4±0.6	82.0±1.0	92.5±0.5	90.2±0.7	88.2±0.4	345
QRNN (k=1) + highway	204k	84.0±1.9	93.4±0.8	82.1±1.2	93.2±0.6	89.6±1.2	88.9±0.2	371
SRU (2 layers)	204k	84.9±1.6	93.5±0.6	82.3±1.2	94.0±0.5	90.1±0.7	89.2±0.3	320
SRU (4 layers)	303k	85.9±1.5	93.8±0.6	82.9±1.0	94.8±0.5	90.1±0.6	89.6±0.5	510
SRU (8 layers)	502k	86.4±1.7	93.7±0.6	83.1±1.0	94.7±0.5	90.2±0.8	88.9±0.6	879

Table 1: Test accuracies on classification benchmarks (Section 4.1). The first block presents best reported results of various methods. The second block compares SRU and other baselines given the same setup. For the SST dataset, we report average results of 5 runs. For other datasets, we perform 3 independent trials of 10-fold cross validation (3×10 runs). The last column compares the wall clock time (in seconds) to finish 100 epochs on the SST dataset.

other architectures. We stack multiple layers of SRU to directly substitute other recurrent, convolutional or feed-forward modules. We minimize hyper-parameter tuning and architecture engineering for a fair comparison. Such efforts have a non-trivial impact on the results, which are beyond the scope of our experiments. Unless noted otherwise, the hyperparameters are set identical to prior work.

4.1 Text Classification

Dataset We use six sentence classification benchmarks: movie review sentiment (MR; Pang and Lee, 2005), sentence subjectivity (SUBJ; Pang and Lee, 2004), customer reviews polarity (CR; Hu and Liu, 2004), question type (TREC; Li and Roth, 2002), opinion polarity (MPQA; Wiebe et al., 2005), and the Stanford sentiment treebank (SST; Socher et al., 2013).²

Following Kim (2014), we use word2vec embeddings trained on 100 billion Google News tokens. For simplicity, all word vectors are normalized to unit vectors and are fixed during training.

Setup We stack multiple SRU layers and use the last output state to predict the class label for a given sentence. We train for 100 epochs and use the validation (i.e., development) set to select the best training epoch. We perform 10-fold

cross validation for datasets that do not have a standard train-evaluation split. The result on SST is averaged over five independent trials. We use Adam (Kingma and Ba, 2014) with the default learning rate 0.001, a weight decay 0 and a hidden dimension of 128.

We compare SRU with a wide range of methods on these datasets, including various convolutional models (Kalchbrenner et al., 2014; Kim, 2014; Zhang and Wallace, 2017) and a hierarchical sentence model (Zhao et al., 2015) reported as the state of the art on these datasets (Conneau et al., 2017). Their setups are not exactly the same as ours, and may involve more tuning on word embeddings and other regularizations. We use the setup of Kim (2014) but do not fine-tune word embeddings and the learning method for simplicity. In addition, we directly compare against three baselines trained using our code base: a re-implementation of the CNN model of Kim (2014), a two-layer LSTM model and Quasi-RNN (Bradbury et al., 2017). We use the official implementation of Quasi-RNN and also implement a version with highway connection for a fair comparison. These baselines are trained using the same hyper-parameter configuration as SRU.

Results Table 1 compares the test results on the six benchmarks. We select the best number re-

²We use the binary version of SST dataset.

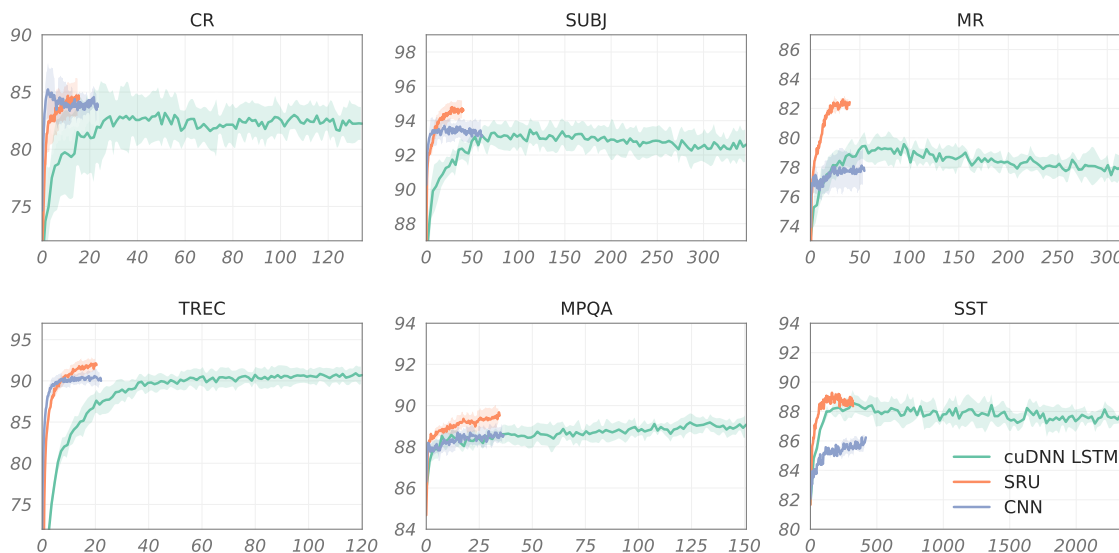


Figure 3: Mean validation accuracies (y-axis) and standard deviations of the CNN, 2-layer LSTM and 2-layer SRU models. We plot the curves of the first 100 epochs. X-axis is the training time used (in seconds). Timings are performed on NVIDIA GeForce GTX 1070 GPU, Intel Core i7-7700K Processor and cuDNN 7003.

ported in previous methods when multiple model variants were explored in their experiments. Despite our simple setup, SRU outperforms most previous methods and achieves comparable results compared to the state-of-the-art but more sophisticated model of [Zhao et al. \(2015\)](#). Figure 3 shows validation performance relative to training time for SRU, cuDNN LSTM and the CNN model. Our SRU implementation runs 5–9 times faster than cuDNN LSTM, and 6–40% faster than the CNN model of [Kim \(2014\)](#). On the movie review (MR) dataset for instance, SRU completes 100 training epochs within 40 seconds, while LSTM takes over 320 seconds.

4.2 Question Answering

Dataset We use the Stanford Question Answering Dataset (SQuAD; [Rajpurkar et al., 2016](#)). SQuAD is a large machine comprehension dataset that includes over 100K question-answer pairs extracted from Wikipedia articles. We use the standard train and development sets.

Setup We use the Document Reader model of [Chen et al. \(2017\)](#) as our base architecture for this task. The model is a combination of word-level bidirectional RNNs and attentions, providing a good testbed to compare our bidirectional SRU

implementation with other RNN components.³

We use the open source implementation of Document Reader in our experiments.⁴ We train models for up to 100 epochs, with a batch size of 32 and a hidden dimension of 128. Following the author suggestions, we use the Adamax optimizer ([Kingma and Ba, 2014](#)) and variational dropout ([Gal and Ghahramani, 2016](#)) during training. We compare with two alternative recurrent components: the bidirectional LSTM adopted in the original implementation of [Chen et al. \(2017\)](#) and Quasi-RNN with highway connections for improved performance.

Results Table 2 summarizes the results on SQuAD. SRU achieves 71.4% exact match and 80.2% F1 score, outperforming the bidirectional LSTM model by 1.9% (EM) and 1.4% (F1) respectively. SRU also exhibits over 5x speed-up over LSTM and 53–63% reduction in total training time. In comparison with QRNN, SRU obtains 0.8% improvement on exact match and 0.6% on F1 score, and runs 60% faster. This speed improvement highlights the impact of the fused ker-

³The current state-of-the-art models ([Seo et al., 2016](#); [Wang et al., 2017](#)) make use of additional components such as character-level embeddings, which are not directly comparable to the setup of [Chen et al. \(2017\)](#). However, these models can potentially benefit from SRU since RNNs are incorporated in the model architecture.

⁴<https://github.com/hitvoice/DrQA>

Model	# layers	Size	Dev	Dev	Time per epoch	
			EM	F1	RNN	Total
LSTM (Chen et al., 2017)	3	4.1m	69.5	78.8	316s	431s
QRNN (k=1) + highway	4	2.4m	70.1 ± 0.1	79.4 ± 0.1	113s	214s
	6	3.2m	70.6 ± 0.1	79.6 ± 0.2	161s	262s
SRU	3	2.0m	70.2 ± 0.3	79.3 ± 0.1	58s	159s
SRU	4	2.4m	70.7 ± 0.1	79.7 ± 0.1	72s	173s
SRU	6	3.2m	71.4 ± 0.1	80.2 ± 0.1	100s	201s

Table 2: Exact match (EM) and F1 scores of various models on SQuAD (Section 4.2). We also report the total processing time per epoch and the time spent in RNN computations. SRU outperforms other models, and is more than five times faster than cuDNN LSTM.

nel (Algorithm 1). While the QRNN baseline involves a similar amount of computation, assembling all element-wise operations of both directions in SRU achieves better GPU utilization.

4.3 Machine Translation

Dataset We train translation models on the WMT English→German dataset, a standard benchmark for translation systems (Peitz et al., 2014; Li et al., 2014; Jean et al., 2015). The dataset consists of 4.5 million sentence pairs. We obtain the pre-tokenized dataset from the OpenNMT project (Klein et al., 2017). The sentences were tokenized using the word-piece model (Wu et al., 2016b), which generates a shared vocabulary of about 32,000 tokens. Newstest-2014 and newstest-2017 are provided and used as the validation and test sets.⁵

Setup We use the state-of-the-art Transformer model of Vaswani et al. (2017) as our base architecture. In the base model, a single Transformer consists of a multi-head attention layer and a bottleneck feed-forward layer. We substitute the feed-forward network using our SRU implementation:

$$\begin{aligned} \text{base: } & \mathbf{W} \cdot \text{ReLU_layer}(\mathbf{x}) + \mathbf{b} \\ \text{ours: } & \mathbf{W} \cdot \text{SRU_layer}(\mathbf{x}) + \mathbf{b} \end{aligned}$$

The intuition is that SRU can better capture sequential information as a recurrent network, and potentially achieve better performance while requiring fewer layers.

We keep the model configuration the same as Vaswani et al. (2017): the model dimension is

⁵<https://github.com/OpenNMT/OpenNMT-tf/tree/master/scripts/wmt>

$d_{\text{model}} = 512$, the feed-forward and SRU layer has inner dimensionality $d_{\text{ff}} = d_{\text{sru}} = 2048$, and positional encoding (Gehring et al., 2017) is applied on the input word embeddings. The base model without SRU has 6 layers, while we set the number of layers to 4 and 5 when SRU is added. Following the original setup, we use a dropout probability 0.1 for all components, except the SRU in the 5-layer model, for which we use a dropout of 0.2 as we observe stronger over-fitting in training.

We use a single NVIDIA Tesla V100 GPU for each model. The published results were obtained using 8 GPUs in parallel, which provide a large effective batch size during training. To approximate the setup, we update the model parameters every 5×5120 tokens and use 16,000 warm-up steps following OpenNMT suggestions. We train each model for 40 epochs (250,000 steps), and perform 3 independent trials for each model configuration. A single run takes about 3.5 days with a Tesla V100 GPU.

Results Table 3 shows the translation results. When SRU is incorporated into the architecture, both the 4-layer and 5-layer model outperform the Transformer base model. For instance, our 5-layer model obtains an average improvement of 0.7 test BLEU score and an improvement of 0.5 BLEU score by comparing the best results of each model achieved across three runs. SRU also exhibits more stable performance, with smaller variance over 3 runs. Figure 4 further compares the validation accuracy of different models. These results confirm that SRU is better at sequence modeling compared to the original feed-forward network (FFN), requiring fewer layers to achieve sim-

Model	# layers	Size	BLEU score		Speed (toks/sec)	Hours per epoch
			Valid	Test		
Transformer (base)	6	76m	26.6±0.2 (26.9)	27.6±0.2 (27.9)	20k	2.0
Transformer (+SRU)	4	79m	26.7±0.1 (26.8)	27.8±0.1 (28.3)	22k	1.8
Transformer (+SRU)	5	90m	27.1±0.0 (27.2)	28.3±0.1 (28.4)	19k	2.1

Table 3: English→German translation results (Section 4.3). We perform 3 independent runs for each configuration. We select the best epoch based on the valid BLEU score for each run, and report the average results and the standard deviation over 3 runs. In addition, we experiment with averaging model checkpoints and use the averaged version for evaluation, following (Vaswani et al., 2017). We show the best BLEU results achieved in brackets.

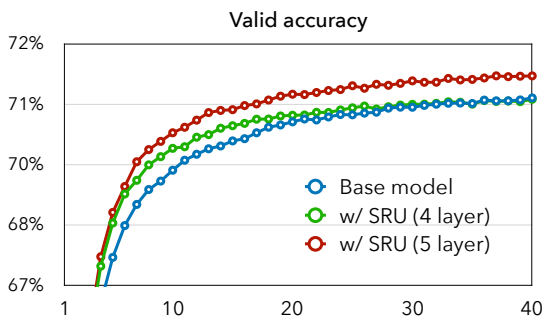


Figure 4: Mean validation accuracy (y-axis) of different translation models after each training epoch (x-axis).

ilar accuracy. Finally, adding SRU does not affect the parallelization or speed of Transformer – the 4-layer model exhibits 10% speed improvement, while the 5-layer model is only 5% slower compared to the base model. We present more results and discussion in Appendix B.3.

4.4 Character-level Language Modeling

Dataset We use Enwik8, a large dataset for character-level language modeling. Following standard practice, we use the first 90M characters for training and the remaining 10M split evenly for validation and test.

Setup Similar to previous work, we use a batch size of 128 and an unroll size of 100 for truncated backpropagation during training. We also experiment with an unroll size of 256 and a batch size of 64 such that each training instance has longer context. We use a non-zero highway bias $\mathbf{b}_r = -3$ that is shown useful for training language model (Zilly et al., 2017). Previous methods employ different optimizers and learning rate schedulers for training. For simplicity and consistency, we use the Adam optimizer and the same

learning rate scheduling (i.e., Noam scheduling) as the translation experiments. We train a maximum of 100 epochs (about 700,000 steps).

We compare various recurrent models and use a parameter budget similar to previous methods. In addition, we experiment with the factorization trick (Kuchaiev and Ginsburg, 2017) to reduce the total number of parameters without decreasing the performance. See details in Appendix B.

Results Table 4 presents the results of SRU and other recurrent models. The 8-layer SRU model achieves validation and test bits per character (BPC) of 1.21, outperforming previous best reported results of LSTM, QRNN and recurrent highway networks (RHN). Increasing the layer of SRU to 12 and using a longer context of 256 characters in training further improves the BPC to 1.19

4.5 Ablation Analysis

We perform ablation analyses on SRU by successively disabling different components:

- (1) Remove the point-wise multiplication term $\mathbf{v} \odot \mathbf{c}_{t-1}$ in the forget and reset gates. The resulting variant involves less recurrence and has less representational capacity.
- (2) Disable the scaling correction by setting the constant $\alpha = 1$.
- (3) Remove the skip connections.

We train model variants on the classification and question answering datasets. Table 5 and Figure 5 confirm the impact of our design decisions – removing these components result in worse classification accuracies and exact match scores.

Model	Size	# layers	Unroll size	Valid	Test	Time
MI-LSTM (Wu et al., 2016c)	17m	1	100	-	1.44	-
HM-LSTM (Chung et al., 2016)	35m	3	100	-	1.32	-
LSTM (Melis et al., 2017)	46m	4	50	1.28	1.30	-
RHN (Zilly et al., 2017)	46m	10	50	-	1.27	-
FS-LSTM (Mujika et al., 2017)	47m	4	100	-	1.25	-
QRNN (Merity et al., 2018)	26m	4	200	-	1.33	-
LSTM (Merity et al., 2018)	47m	3	200	-	1.23	-
SRU	37m	6	100	1.29	1.30	28min
SRU	37m	10	100	1.26	1.27	29min
SRU (with projection)	37m	6	100	1.25	1.26	29min
SRU (with projection)	47m	8	100	1.21	1.21	39min
SRU (with projection)	49m	12	256	1.19	1.19	41min

Table 4: Validation and test BPCs of different recurrent models on Enwik8 dataset. The last column presents the training time per epoch. For SRU with projection, we set the projection dimension to 512.

Model	4 layers	6 layers
SRU (full)	70.7	71.4
– remove $\mathbf{v} \odot \mathbf{c}_{t-1}$	70.6	71.1
– remove α -scaling	70.3	71.0
– remove highway	69.4	69.1

Table 5: Ablation analysis on SQuAD. Components are successively removed and the EM scores are averaged over 4 runs.




Figure 5: Ablation analysis on the classification datasets. Average validation results are presented. We compare the full SRU implementation (left blue), the variant without $\mathbf{v} \odot \mathbf{c}_{t-1}$ multiplication (middle green) and the variant without highway connection (right yellow).

5 Discussion

This work presents Simple Recurrent Unit (SRU), a scalable recurrent architecture that operates as fast as feed-forward and convolutional units. We confirm the effectiveness of SRU on multiple natural language tasks ranging from classification to

translation. We open source our implementation to facilitate future NLP and deep learning research.

Trading capacity with layers SRU achieves high parallelization by simplifying the hidden-to-hidden dependency. This simplification is likely to reduce the representational power of a single layer and hence should be balanced to avoid performance loss. However, unlike previous work that suggests additional computation (e.g., n-gram filters) within the layer (Balduzzi and Ghifary, 2016; Bradbury et al., 2017), we argue that increasing the depth of the model suffices to retain modeling capacity. Our empirical results on various tasks confirm this hypothesis.

Acknowledgement

We thank Alexander Rush and Yoon Kim for help with machine translation experiments, and Danqi Chen for help with SQuAD experiments. We thank Adam Yala, Howard Chen, Jeremy Wohlwend, Lili Yu, Kyle Swanson and Kevin Yang for providing useful feedback on the paper and the SRU implementation. A special thanks to Hugh Perkins for his support on the experimental environment setup and Runqi Yang for answering questions about his code.

References

Fabio Anselmi, Lorenzo Rosasco, Cheston Tan, and Tomaso A. Poggio. 2015. Deep convolutional networks are hierarchical kernel machines. *CoRR*, abs/1508.01084.

- Jeremy Appleyard, Tomas Kocisky, and Phil Blunsom. 2016. Optimizing performance of recurrent neural networks on gpus. *CoRR*, abs/1604.01946.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. Neural machine translation by jointly learning to align and translate. In *Proceedings of the International Conference on Learning Representations*.
- David Balduzzi and Muhammad Ghifary. 2016. Strongly-typed recurrent neural networks. In *International Conference on Machine Learning*.
- James Bradbury, Stephen Merity, Caiming Xiong, and Richard Socher. 2017. Quasi-recurrent neural networks. In *Proceedings of the International Conference on Learning Representations*.
- Danqi Chen, Adam Fisch, Jason Weston, and Antoine Bordes. 2017. Reading wikipedia to answer open-domain questions. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics*.
- Yining Chen, Sorcha Gilroy, Kevin Knight, and Jonathan May. 2018. Recurrent neural networks as weighted language recognizers. In *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*.
- Kyunghyun Cho, Bart van Merrienboer, Aglar Guleghre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2014. Learning phrase representations using rnn encoder-decoder for statistical machine translation. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*.
- Junyoung Chung, Sungjin Ahn, and Yoshua Bengio. 2016. Hierarchical multiscale recurrent neural networks. *CoRR*, abs/1609.01704.
- Alexis Conneau, Douwe Kiela, Holger Schwenk, Loic Barrault, and Antoine Bordes. 2017. Supervised learning of universal sentence representations from natural language inference data. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*.
- Amit Daniely, Roy Frostig, and Yoram Singer. 2016. Toward deeper understanding of neural networks: The power of initialization and a dual view on expressivity. In *Advances In Neural Information Processing Systems*.
- Greg Diamos, Shubho Sengupta, Bryan Catanzaro, Mike Chrzanowski, Adam Coates, Erich Elsen, Jesse Engel, Awni Hannun, and Sanjeev Satheesh. 2016. Persistent rnns: Stashing recurrent weights on-chip. In *International Conference on Machine Learning*.
- Yarin Gal and Zoubin Ghahramani. 2016. A theoretically grounded application of dropout in recurrent neural networks. In *Advances in Neural Information Processing Systems*.
- Jonas Gehring, Michael Auli, David Grangier, Denis Yarats, and Yann Dauphin. 2017. Convolutional sequence to sequence learning. In *International Conference on Machine Learning*.
- Xavier Glorot and Yoshua Bengio. 2010. Understanding the difficulty of training deep feedforward neural networks. In *Proceedings of the international conference on artificial intelligence and statistics*.
- Priya Goyal, Piotr Dollar, Ross B. Girshick, Pieter Noordhuis, Lukasz Wesolowski, Aapo Kyrola, Andrew Tulloch, Yangqing Jia, and Kaiming He. 2017. Accurate, large minibatch SGD: Training imagenet in 1 hour. *CoRR*, abs/1706.02677.
- Klaus Greff, Rupesh Kumar Srivastava, Jan Koutnek, Bas R. Steunebrink, and Jurgen Schmidhuber. 2017. Lstm: A search space odyssey. *IEEE Transactions on Neural Networks and Learning Systems*, 28.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2015. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *Proceedings of the IEEE international conference on computer vision*.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*.
- Luheng He, Kenton Lee, Mike Lewis, and Luke Zettlemoyer. 2017. Deep semantic role labeling: What works and what’s next. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics*.
- Sepp Hochreiter and Jurgen Schmidhuber. 1997. Long short-term memory. *Neural computation*, 9.
- Minqing Hu and Bing Liu. 2004. Mining and summarizing customer reviews. In *Proceedings of the tenth ACM SIGKDD international conference on Knowledge discovery and data mining*.
- Ozan Irsoy and Claire Cardie. 2014. Opinion mining with deep recurrent neural networks. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*.
- Sebastien Jean, Kyunghyun Cho, Roland Memisevic, and Yoshua Bengio. 2015. On using very large target vocabulary for neural machine translation. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics and the International Joint Conference on Natural Language Processing*.

- Nal Kalchbrenner, Edward Grefenstette, and Phil Blunsom. 2014. A convolutional neural network for modelling sentences. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics*.
- Yoon Kim. 2014. Convolutional neural networks for sentence classification. In *Proceedings of the Empirical Methods in Natural Language Processing*.
- Yoon Kim, Yacine Jernite, David A Sontag, and Alexander M. Rush. 2016. Character-aware neural language models. In *Proceedings of the AAAI Conference on Artificial Intelligence*.
- Diederik Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. In *Proceedings of the International Conference on Learning Representations*.
- Guillaume Klein, Yoon Kim, Yuntian Deng, Jean Senellart, and Alexander Rush. 2017. Opennmt: Open-source toolkit for neural machine translation. In *Proceedings of ACL 2017, System Demonstrations*.
- Oleksii Kuchaiev and Boris Ginsburg. 2017. Factorization tricks for lstm networks. *CoRR*, abs/1703.10722.
- Quoc V. Le, Navdeep Jaitly, and Geoffrey E. Hinton. 2015. A simple way to initialize recurrent networks of rectified linear units. *CoRR*, abs/1504.00941.
- Kenton Lee, Omer Levy, and Luke S. Zettlemoyer. 2017. Recurrent additive networks. *CoRR*, abs/1705.07393.
- Tao Lei, Regina Barzilay, and Tommi Jaakkola. 2015. Molding cnns for text: non-linear, non-consecutive convolutions. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics.
- Tao Lei, Wengong Jin, Regina Barzilay, and Tommi Jaakkola. 2017. Deriving neural architectures from sequence and graph kernels. *International Conference on Machine Learning*.
- Liangyou Li, Xiaofeng Wu, Santiago Cortes Vaillio, Jun Xie, Andy Way, and Qun Liu. 2014. The DCU-ICTCAS MT system at WMT 2014 on german-english translation task. In *Proceedings of the Ninth Workshop on Statistical Machine Translation*.
- Xin Li and Dan Roth. 2002. Learning question classifiers. In *Proceedings of the international conference on Computational linguistics-Volume 1*. Association for Computational Linguistics.
- Minh-Thang Luong, Hieu Pham, and Christopher D. Manning. 2015. Effective approaches to attention-based neural machine translation. In *Empirical Methods in Natural Language Processing*. Association for Computational Linguistics.
- Hongyuan Mei, Mohit Bansal, and R. Matthew Walter. 2016. What to talk about and how? selective generation using lstms with coarse-to-fine alignment. In *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*.
- Gábor Melis, Chris Dyer, and Phil Blunsom. 2017. On the state of the art of evaluation in neural language models. *CoRR*, abs/1707.05589.
- Stephen Merity, Nitish Shirish Keskar, and Richard Socher. 2018. An analysis of neural language modeling at multiple scales. *CoRR*, abs/1803.08240.
- Yajie Miao, Jinyu Li, Yongqiang Wang, Shi-Xiong Zhang, and Yifan Gong. 2016. Simplifying long short-term memory acoustic models for fast training and decoding. In *IEEE International Conference on Acoustics, Speech and Signal Processing*.
- Dipendra Misra, John Langford, and Yoav Artzi. 2017. Mapping instructions and visual observations to actions with reinforcement learning. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*.
- Asier Mujika, Florian Meier, and Angelika Steger. 2017. Fast-slow recurrent neural networks. In *Advances in Neural Information Processing Systems*.
- Bo Pang and Lillian Lee. 2004. A sentimental education: Sentiment analysis using subjectivity summarization based on minimum cuts. In *Proceedings of the annual meeting on Association for Computational Linguistics*.
- Bo Pang and Lillian Lee. 2005. Seeing stars: Exploiting class relationships for sentiment categorization with respect to rating scales. In *Proceedings of the annual meeting on association for computational linguistics*.
- Stephan Peitz, Joern Wuebker, Markus Freitag, and Hermann Ney. 2014. The RWTH aachen german-english machine translation system for wmt 2014. In *Proceedings of the Ninth Workshop on Statistical Machine Translation*.
- Hao Peng, Roy Schwartz, Sam Thomson, and Noah A. Smith. 2018. Rational recurrences. In *Empirical Methods in Natural Language Processing*.
- P. Rajpurkar, J. Zhang, K. Lopyrev, and P. Liang. 2016. Squad: 100,000+ questions for machine comprehension of text. In *Empirical Methods in Natural Language Processing*.
- Min Joon Seo, Aniruddha Kembhavi, Ali Farhadi, and Hannaneh Hajishirzi. 2016. Bidirectional attention flow for machine comprehension. *CoRR*, abs/1611.01603.
- Noam Shazeer, Azalia Mirhoseini, Krzysztof Maziarz, Andy Davis, Quoc Le, Geoffrey Hinton, and Jeff Dean. 2017. Outrageously large neural networks:

- The sparsely-gated mixture-of-experts layer. *arXiv preprint arXiv:1701.06538*.
- Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D. Manning, Andrew Y. Ng, and Christopher Potts. 2013. Recursive deep models for semantic compositionality over a sentiment treebank. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*.
- Rupesh K Srivastava, Klaus Greff, and Jürgen Schmidhuber. 2015. Training very deep networks. In *Advances in neural information processing systems*.
- Alane Suhr and Yoav Artzi. 2018. Situated mapping of sequential instructions to actions with single-step reward observation. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics*.
- Alane Suhr, Srinivasan Iyer, and Yoav Artzi. 2018. Learning to map context-dependent sentences to executable formal queries. In *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in Neural Information Processing Systems*.
- Sida Wang and Christopher Manning. 2013. Fast dropout training. In *International Conference on Machine Learning*.
- Wenhui Wang, Nan Yang, Furu Wei, Baobao Chang, and Ming Zhou. 2017. Gated self-matching networks for reading comprehension and question answering. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics*.
- Janyce Wiebe, Theresa Wilson, and Claire Cardie. 2005. Annotating expressions of opinions and emotions in language. *Language resources and evaluation*.
- Huijia Wu, Jiajun Zhang, and Chengqing Zong. 2016a. An empirical exploration of skip connections for sequential tagging. In *Proceedings of the International Conference on Computational Linguistics*.
- Yonghui Wu, Mike Schuster, Zhifeng Chen, Quoc V. Le, Mohammad Norouzi, Wolfgang Macherey, Maxim Krikun, Yuan Cao, Qin Gao, Klaus Macherey, Jeff Klingner, Apurva Shah, Melvin Johnson, Xiaobing Liu, Łukasz Kaiser, Stephan Gouws, Yoshikiyo Kato, Taku Kudo, Hideto Kazawa, Keith Stevens, George Kurian, Nishant Patil, Wei Wang, Cliff Young, Jason Smith, Jason Riesa, Alex Rudnick, Oriol Vinyals, Greg Corrado, Macduff Hughes, and Jeffrey Dean. 2016b. Google’s neural machine translation system: Bridging the gap between human and machine translation. *CoRR*, abs/1609.08144.
- Yuhuai Wu, Saizheng Zhang, Ying Zhang, Yoshua Bengio, and Ruslan R Salakhutdinov. 2016c. On multiplicative integration with recurrent neural networks. In *Advances in Neural Information Processing Systems*.
- Wojciech Zaremba, Ilya Sutskever, and Oriol Vinyals. 2014. Recurrent neural network regularization. *CoRR*, abs/1409.2329.
- Yingjie Zhang and Byron C. Wallace. 2017. A sensitivity analysis of (and practitioners’ guide to) convolutional neural networks for sentence classification. In *Proceedings of the International Joint Conference on Natural Language Processing*.
- Yuchen Zhang, Jason D. Lee, and Michael I. Jordan. 2016. ℓ_1 -regularized neural networks are improperly learnable in polynomial time. In *International Conference on Machine Learning*.
- Han Zhao, Zhengdong Lu, and Pascal Poupart. 2015. Self-adaptive hierarchical sentence model. In *Proceedings of the International Joint Conference on Artificial Intelligence*.
- Julian Georg Zilly, Rupesh Kumar Srivastava, Jan Koutník, and Jürgen Schmidhuber. 2017. Recurrent highway networks. In *International Conference on Machine Learning*.
- Barret Zoph and Quoc V. Le. 2016. Neural architecture search with reinforcement learning. *CoRR*, abs/1611.01578.