Universitat Politècnica de Catalunya

Escola Tècnica Superior d'Enginyeria de Telecomunicació de Barcelona

Learning to Skip State Updates in Recurrent Neural Networks

Víctor Campos Camúñez

A thesis submitted in fulfillment of the requirements for the degree of the Master in Telecommunications Engineering

Barcelona, August 2017

Abstract

Recurrent Neural Networks (RNNs) continue to show outstanding performance in sequence modeling tasks. However, training RNNs on long sequences often face challenges like slow inference, vanishing gradients and difficulty in capturing long term dependencies. In backpropagation through time settings, these issues are tightly coupled with the large, sequential computational graph resulting from unfolding the RNN in time. We introduce the Skip RNN model which extends existing RNN models by learning to skip state updates and shortens the effective size of the computational graph. This network can be encouraged to perform fewer state updates through a novel loss term. We evaluate the proposed model on various tasks and show how it can reduce the number of required RNN updates while preserving, and sometimes even improving, the performance of the baseline models.

Keywords: machine learning, deep learning, recurrent neural networks, sequence modeling, conditional computation.

Resum

Les Xarxes Neuronals Recurrents (de l'anglès, RNNs) mostren un alt rendiment en tasques de modelat de seqüències. Tot i així, entrenar RNNs en seqüències llargues sol provocar dificultats com una inferència lenta, gradients que s'esvaeixen i dificultats per capturar depèndencies temporals a llarg terme. En escenaris amb *backpropagation through time*, aquests problemes estan estretament relacionats amb la longitud i la seqüencialitat del graf computacional resultant de desdoblar la RNN en el temps. Presentem Skip RNN, model que extén arquitectures recurrents existents, permetent-les aprendre quan ometre actualitzacions del seu estat i escurçant així la longitud efectiva del graf computacional. Aquesta xarxa pot ser estimulada per efectuar menys actualitzacions d'estat a través d'un nou terme a la funció de cost. Evaluem el model proposat en una sèrie de tasques i demostrem com pot reduir el nombre d'actualitzacions de la RNN mentre preserva, o fins i tot millora, el rendiment dels models de referència.

Paraules clau: aprenentatge automàtic, aprenentatge profund, xarxes neuronals recurrents, modelat de seqüències, computació condicional.

Resumen

Las Redes Neuronales Recurrentes (del inglés, RNNs) muestran un alto rendimiento en tareas de modelado de secuencias. Aún así, entrenar RNNs en secuencias largas suele provocar dificultades como una inferencia lenta, gradientes que se desvanecen y dificultades para capturar dependencias temporales a largo plazo. En escenarios con *backpropagation through time*, estos problemas están estrechamente relacionados con la longitud y la secuencialidad del grafo computacional resultante de desdoblar la RNN en el tiempo. Presentamos Skip RNN, un modelo que extiende arquitecturas recurrentes existentes, permitiéndoles aprender cuándo omitir actualizaciones de su estado y acortando así la longitud efectiva del grafo computacional. Esta red puede ser estimulada para efectuar menos actualizaciones de estado a través de un nuevo elemento en la función de coste. Evaluamos el modelo propuesto en una serie de tareas y demostramos cómo puede reducir el número de actualizaciones de la RNN mientras preserva, o incluso mejora, el rendimiento de los modelos de referencia.

Palabras clave: aprendizaje automático, aprendizaje profundo, redes neuronales recurrentes, modelado de secuencias, computación condicional.

Contents

List of Figures

List of Tables

List of Abbreviations

NN	**N**eural **N**etwork
RNN	**R**ecurrent **N**eural **N**etwork
LSTM	**L**ong **S**hort-**T**erm **M**emory
GRU	**G**ated **R**ecurrent **U**nit
GPU	**G**raphics **P**rocessing **U**nit

Chapter 1

Introduction

1.1 Motivation

Methods based on deep neural networks have established the state-of-the-art in computer vision tasks [32, 53, 23], machine translation [56], speech generation [43] and even defeated the world champion in the game of Go [50]. Although the development of these algorithms spans over many decades [35], their potential has been unlocked by the increased computational power of specific accelerators, e.g. Graphical Processing Units (GPUs), and the creation of large-scale datasets [15, 28]. Those readers who are not familiar with deep learning techniques may prefer to read Chapter 2, which gives an overview of the most basic concepts related to neural networks, before diving deeper in the contents of this thesis.

Recurrent Neural Networks (RNNs) have become the standard approach for practitioners when addressing machine learning tasks involving sequential data. Besides the increase in computational power and publicly available data, improved architectures and training algorithms have been key for the success of RNNs. Gated units, such as the Long Short-Term Memory [25] (LSTM) and the Gated Recurrent Unit [10] (GRU), were designed to deal with the vanishing gradients problem commonly found in RNNs [8] that precluded their application during years. These architectures have popularized thanks to their impressive results in a variety of tasks such as machine translation [5], language modeling [61] or speech recognition [20].

RNNs process data sequentially in order to capture temporal dependencies. This inherently sequential behavior results in challenging training and deployment, especially when dealing with long sequences. These challenges include throughput degradation, slower convergence during training and memory leakage, even for gated architectures [41]. Sequence shortening techniques, which can be seen as a sort of conditional computation [7, 6, 14] in time, can alleviate these issues. The most common approaches, such as cropping discrete signals or reducing the sampling rate in continuous signals, are heuristics and can be suboptimal. In contrast, we propose a model that is able to learn which samples (i.e. elements in the input sequence) need to be used in order to solve the target task. Consider a video understanding task as an example: scenes with large motion may benefit from high frame rates, whereas only a few frames are needed to capture the semantics of a mostly static scene.

The main goal of this project is the development of an RNN architecture that is able to decide

which samples in a sequence need to be used in order to solve the task at hand. Such subsampling mechanism needs to be integrated in the RNN itself, as opposed to a pre-processing step, and needs to be data-driven instead of being in heuristics. Moreover, the sample selection process has to be dynamic, i.e. each input sequence may be subsampled in a different way depending on their actual content. Our contributions include:

- A novel modification for existing RNN architectures that allows them to skip state updates, decreasing the number of inherently sequential operations to be performed, without requiring any additional supervision signal. This model, called Skip RNN, adaptively determines whether the state needs to be updated or it can be copied to the next time step, thus skipping computation. It can be seamlessly introduced into existing architectures to reduce the number of required RNN updates while preserving, and sometimes even improving, the performance of the baseline models.
- We show how the network can be encouraged to perform fewer state updates by adding a new penalization term during training, allowing to easily train models targeting different computational budgets.
- The proposed modification is implemented on top of well known RNN architectures, namely LSTM and GRU, and the resulting models show promising results in a series of sequence modeling tasks. In particular, the proposed Skip RNN architecture is evaluated on (1) the adding task, (2) sine wave frequency discrimination, (3) digit classification, (4) sentiment analysis in movie reviews, and (5) action classification in video.

This thesis is structured as follows. A brief introduction to deep learning is given in Chapter 2, which may come in handy for readers who are not familiar with the field. Chapter 3 provides an overview of the related work and techniques that bear some similarity with the proposed solution. The Skip RNN model is described in Chapter 4 and evaluated in a series of tasks in Chapter 5. Discussion and future research directions are provided in Chapter 6. Finally, appendices include additional qualitative results for the experiments and the scientific publications generated from this work.

1.2 Hardware and Software Resources

This project was developed using the resources provided by the Barcelona Supercomputing Center. In particular, the MinoTauro[1] cluster was been used. The cluster has 39 nodes equipped with two NVIDIA K80 dual-GPUs each, resulting in up to four concurrent processes running in parallel in each node. Having access to several nodes at the same time proved itself crucial for tuning hyperparameters and performing different runs for each experiment. Evaluating the proposed models on such a broad set of tasks would have been unfeasible without access to these computing resources.

Experiments were implemented with TensorFlow[2], using CUDA and cuDNN for fast GPU primitives. This framework was chosen due to previous experience and the fact that it was already available in the MinoTauro cluster. Code is publicly available at `https://github.com/imatge-upc/skiprnn-2017-tfm/`.

[1] `https://www.bsc.es/innovation-and-services/supercomputers-and-facilities/minotauro/`
[2] `https://www.tensorflow.org/`

1.3 Work Plan

Being this a research-oriented project, it was difficult to define a detailed work plan ahead of time: the next steps to take were generally conditioned on the latest obtained results. Besides, since the field of deep learning is extremely active, new results are published and uploaded to arXiv every day and may affect the development of the project. Without going any further, the work that is most similar to ours [59] was published during the development of this project.

The project was organized as follows. After reading and understanding the works related with the goal of the project, the first step consisted in finding a task representing the problem being addressed. This ended up being the adding task in Section 5.1, which was used to evaluate different proposals until the final model presented in Chapter 4 was reached. This final model was the result of several iterations over a *design-implement-evaluate-refine* process. Once the adding task was solved without using all the input samples, it was time to evaluate the model on harder tasks. Finally, the last month was spent writing the scientific paper and this report.

Chapter 2

Introduction to Deep Learning

This chapter aims to introduce the reader to the most basic concepts in deep learning and, in particular, to the most common used recurrent models which are key to this work. These techniques have become the most common solution for most machine learning tasks in recent years and it is impossible to understand the latest advances in machine learning without deep learning. A detailed explanation of these methods would make this document excessively long and is out of the scope of this thesis. Among the numerous resources that have been published recently, we refer the reader to the book by Goodfellow et al. [18] for an extensive and comprehensive summary of the most important concepts in deep learning.

2.1 Feedforward Neural Networks

Artificial Neural Networks (NNs) have become the de facto solution for many Machine Learning tasks. Despite these methods were proposed decades ago [46, 35], their application to a wide range of problems is relatively recent and has been enabled by the creation of public large scale datasets and more powerful computing devices. In particular, the impressive results achieved by a NN-based model [32] in the 2012 edition of the ImageNet Large Scale Visual Recognition Challenge [15] kick started new research directions based on the usage of NNs.

Neural Networks can be seen as function approximators learning a mapping from their input to their output. In a nutshell, they consist of a stack of non-linear operations whose parameters are trained by means of Stochastic Gradient Descent (SGD) and backpropagation in order to minimize a training objective or cost. Despite achieving outstanding results, such method is very data inefficient and requires from large example collections to be effective. The non-linear nature of these of operations, also known as layers, is the key feature to produce powerful function approximators; in contrast, a stack of linear operations is mathematically equivalent to a single linear operation. A layer in a neural network can be modeled as follows:

$$h = f\left(W_h x + b_h\right) \tag{2.1}$$

where x is the input vector, f is a non-linear function (e.g. sigmoid) and W_h and b_h are trainable weights and biases, respectively. Additional hidden layers may be added between the input and the output to provide the model with an increased capacity.

2.2 Recurrent Neural Networks

Consider a sequence of input vectors, $\mathbf{x} = (x_1, ..., x_T)$. The feedforward model described in Section 2.1 would output a sequence of activation vectors, $\mathbf{h} = (h_1, ..., h_T)$, which are agnostic to previous activations. Recurrent Neural Networks (RNNs) extend the feedforward model by adding a recurrent connection in time:

$$h_t = f\left(W_h x_t + U_h h_{t-1} + b_h\right) \tag{2.2}$$

where U_h operates on the hidden state in the previous time step, h_{t-1}, allowing information to persist. Providing models with memory and enabling them to model the temporal evolution of signals is a key factor in many sequence classification and transduction tasks where RNNs excel, such as machine translation [5], language modeling [61] or speech recognition [20].

Training recurrent models with backpropagation has been proven difficult for long sequences [8, 25]. The reason for such difficulty lies in the multiplicative dependency between activations over time, which easily leads to vanishing or exploding gradients. This lead to the development of *gated* units, which aim at alleviating these issues and making training for longer sequences feasible. The most common gated RNN architectures are the Long Short-Term Memory [25] and the Gated Recurrent Unit [10]. The following subsections describe the basics of these two architectures, and we refer the reader to [42] for an extended description and intuition behind them.

2.2.1 Long Short-Term Memory

The core idea behind Long Short-Term Memory (LSTM) is the so called *cell state*, where information is accumulated over time. Unlike the hidden state in vanilla RNN, the cell state in LSTMs is accessed through addition operations which do not suffer from the same problems as multiplicative relationships during gradient backpropagation. Its update equations can be defined as follows:

$$i_t = \sigma\left(W_{xi} x_t + W_{hi} h_{t-1} + b_i\right) \tag{2.3}$$
$$f_t = \sigma\left(W_{xf} x_t + W_{hf} h_{t-1} + b_f\right) \tag{2.4}$$
$$c_t = f_t \odot c_{t-1} + i_t \odot \tanh\left(W_{xc} x_t + W_{hc} h_{t-1} + b_c\right) \tag{2.5}$$
$$o_t = \sigma\left(W_{xo} x_t + W_{ho} h_{t-1} + b_o\right) \tag{2.6}$$
$$h_t = o_t \odot \tanh\left(c_t\right) \tag{2.7}$$

where i_t, f_t and o_t represent the *input*, *forget* and *output* gate at time t, respectively, σ is the sigmoid function, $\odot$ represents element-wise multiplication, W are trainable weight matrices and b are trainable bias vectors. The cell state at time t is represented by c_t, whereas the hidden state (i.e. the cell output) at each time step is h_t.

2.2.2 Gated Recurrent Unit

The Gated Recurrent Unit (GRU) can be understood as a simplified version of LSTM. The simplification is two-fold: (1) erasing information in the cell state is needed whenever a new value is written to it, and (2) the hidden state is directly the cell state. This is implemented

by combining the *input* and *forget* gates into a single gate and removing the output gate. The number of learnable parameters is reduced and inference is faster due to the fewer number of operations. Its update equations can be defined as follows:

$$z_t = \sigma\left(W_{xz}x_t + W_{hz}h_{t-1}\right) \tag{2.8}$$
$$r_t = \sigma\left(W_{xr}x_t + W_{hr}h_{t-1}\right) \tag{2.9}$$
$$\tilde{h}_t = \tanh\left(W_{x\tilde{h}}r_t \odot x_t + W_{h\tilde{h}}h_{t-1}\right) \tag{2.10}$$
$$h_t = (1 - z_t) \odot h_{t-1} + z_t \odot \tilde{h}_t \tag{2.11}$$

where z_t and r_t are the *update* and *reset* gates, respectively, σ is the sigmoid function, $\odot$ represents element-wise multiplication and W are trainable weight matrices. The hidden state, which in this architecture is merged with the cell state, is represented by h_t.

2.2.3 Memory decay in RNNs

Despite providing an outstanding improvement over vanilla RNNs, gated units still have difficulties to remember information during long time lags. Neil et al. [41] show how the memory decay rate is exponential for an LSTM on the simple task of keeping an initial memory state c_0 for as long as possible. Assuming that no additional inputs are received (i.e. $i_t = 0, \forall t$) and that the forget gate is nearly fully-opened (i.e. $f_t = 1 - \epsilon, \forall t$), after n updates the cell state would contain:

$$c_n = f_n \odot c_{n-1} = (1 - \epsilon) \odot (f_{n-1} \odot c_{n-2}) = \ldots = (1 - \epsilon)^n \odot c_0 \tag{2.12}$$

resulting in an exponential memory decay for $\epsilon < 1$. Please note that the forget gate is the output of a sigmoid function, thus $f_t < 1$ and memory decays with every time step. The same reasoning can be applied to GRU.

Chapter 3

Related Work

This chapters summarizes works that bear some similarity with the presented model. To the best of our knowledge, only LSTM-Jump [59] addresses a similar problem to the one discussed in this thesis. However, the Skip RNN model lies in the intersection of several techniques from the state of the art in machine learning, which are covered in the following sections.

3.1 Conditional computation

Conditional computation allows to increase model capacity without a proportional increase in computational cost by evaluating only certain computation paths for each input [7, 36, 2, 38, 47]. This concept has been extended to the temporal domain, either by learning how many times an input needs to be pondered before moving to the next one [19] or building RNNs whose number of layers depends on the input data [11]. Some works have addressed time-dependent computation in RNNs by updating only a fraction of the hidden state based on the current hidden state and input [27], or following periodic patterns [31, 41]. However, due to the inherently sequential nature of RNNs and the parallel computation capabilities of modern hardware, reducing the size of the matrices involved in the computations performed at each time step does not accelerate inference. The proposed Skip RNN model can be seen as form of conditional computation in time, where the computation associated to the RNN updates may or may not be executed at every time step. This is related to the UPDATE and COPY operations in hierarchical multiscale RNNs [11], but applied to the whole stack of RNN layers at the same time. Such difference is key to allow our approach to skip input samples, effectively reducing sequential computation and shielding the hidden state over longer time lags. Learning whether to update or copy the hidden state through time steps can be seen as a learnable Zoneout mask [33] which is shared between all the units in the hidden state. Similarly, it can be understood as an input-dependent recurrent version of stochastic depth [26].

3.2 Attention models and subsampling mechanisms

Selecting parts of the input signal is similar in spirit to the hard attention mechanisms that have been applied to image regions [40], where only some patches of the input image are attended in

order to generate captions [57] or detect objects [3]. In a similar fashion, Adaptive Computation Time [19] can be extended to residual networks [23] for image recognition, allowing some image patches to be processed only by a subset of layers [17]. Our model can be understood to generate a hard temporal attention mask on the fly given the previously seen samples, deciding which time steps should be attended and operating on a subset of input samples.

Subsampling input sequences has been explored for visual storylines generation [49], although jointly optimizing the RNN weights and the subsampling mechanism is computationally unfeasible and the Expectation Maximization algorithm is used instead. Similar research has been conducted for video analysis tasks, discovering Minimally Needed Evidence for event recognition [9] and training agents that decide which frames need to be observed in order to localize actions in time [58]. Motivated by the advantages of training recurrent models on shorter subsequences, efforts have been conducted towards learning differentiable subsampling mechanisms [45], although the computational complexity of the proposed method precludes its application to long input sequences. In contrast, our proposed method can be trained with backpropagation and does not degrade the complexity of the baseline RNNs.

3.3 LSTM-Jump

With the goal of speeding up RNN inference, LSTM-Jump [59] augments an LSTM cell with a classification layer that will decide how many steps to jump between RNN updates. Despite its promising results on text tasks, the discrete nature of the jumping actions makes the optimization process more complex. While the LSTM parameters, θ_m, can be directly optimized by backpropagation, the parameters of the new jumping layer, θ_a, need to be trained with REINFORCE [55] on a user-defined reward signal. Determining such reward signal is not trivial and does not necessarily generalize across tasks, e.g. regression and classification tasks may require from different reward signals. In its simplest form, the loss function being minimized is of the form

$$L\left(\theta_m, \theta_a\right) = L_m\left(\theta_m\right) + L_a\left(\theta_a\right) \tag{3.1}$$

where L_m is the loss function that would be used with the baseline LSTM, e.g. cross-entropy for classification or L_2 for regression, and L_a is the loss function defined for the jumping parameters, based on the reward signal. In such formulation, the LSTM parameters and the jumping layer parameters indeed minimize independent loss terms. This forces the practitioner to define an L_a term that mimics the behavior of L_m, which may not be trivial for some tasks. In contrast, the parameters in the proposed Skip RNN model are optimized directly from the original loss function and there is no need to define additional loss terms that provide gradients for them.

In LSTM-Jump, some hyperparameters defining the behavior of the subsampling scheme need to be defined ahead of time:

Notation	Meaning
N	number of jumps allowed
K	maximum size of jumping
R	number of tokens read before a jump

Table 3.1: Hyperparameters defining the jumping behavior of LSTM-Jump.

These hyperparameters define a reduced set of subsequences that the model can sample, instead of allowing the network to learn any arbitrary sampling scheme. This could lead to suboptimal sampling schemes and pushes some of the complexity towards the practitioner's prior knowledge of the task. For instance, should the optimal policy consist in using every other sample, the model will be unable to learn the best sampling scheme if $R = 2$. Unlike LSTM-Jump, Skip RNN is not limited to a set of sample selection patterns which are predefined ahead of time.

Chapter 4

Model

An RNN takes an input sequence $\mathbf{x} = (x_1, \ldots, x_T)$ and generates a state sequence $\mathbf{s} = (s_1, \ldots, s_T)$ by iteratively applying a parametric state transition model S from $t = 1$ to T:

$$s_t = S(s_{t-1}, x_t) \tag{4.1}$$

We augment the network with a binary *state update gate*, $u_t \in \{0, 1\}$, selecting whether the state of the RNN will be updated or copied from the previous time step. At every time step t, the probability $\tilde{u}_{t+1} \in [0, 1]$ of performing a state update at $t + 1$ is emitted. The resulting architecture is depicted in Figure 4.1 and can be characterized as follows:

$$u_t = f_{binarize}(\tilde{u}_t) \tag{4.2}$$
$$s_t = u_t \cdot S(s_{t-1}, x_t) + (1 - u_t) \cdot s_{t-1} \tag{4.3}$$
$$\Delta\tilde{u}_t = \sigma(W_p s_t + b_p) \tag{4.4}$$
$$\tilde{u}_{t+1} = u_t \cdot \Delta\tilde{u}_t + (1 - u_t) \cdot (\tilde{u}_t + \min(\Delta\tilde{u}_t, 1 - \tilde{u}_t)) \tag{4.5}$$

where σ is the sigmoid function and $f_{binarize} : [0, 1] \to \{0, 1\}$ binarizes the input value. Should the network be composed of several layers, some columns of W_p can be fixed to 0 so that $\Delta\tilde{u}_t$ depends only on the states of a subset of layers (see Section 5.5 for an example with two layers). We implement $f_{binarize}$ as a deterministic step function $u_t = \text{round}(\tilde{u}_t)$, although a stochastic sampling from a Bernoulli distribution $u_t \sim \text{Bernoulli}(\tilde{u}_t)$ would be possible as well.

The model formulation implements the observation that the likelihood of requesting a new input increases with the number of consecutively skipped samples. Whenever a state update is omitted, the pre-activation of the state update gate for the following time step, $\tilde{u}_{t+1}$, is incremented by $\Delta\tilde{u}_t$. On the other hand, if a state update is performed, the accumulated value is flushed and $\tilde{u}_{t+1} = \Delta\tilde{u}_t$.

The number of skipped time steps can be computed ahead of time. For the particular formulation used in this work, where $f_{binarize}$ is implemented by means of a rounding function, the number of skipped samples after performing a state update at time step t is given by:

$$N_{skip}(t) = \min\{n : n \cdot \Delta\tilde{u}_t \geq 0.5\} - 1 \tag{4.6}$$

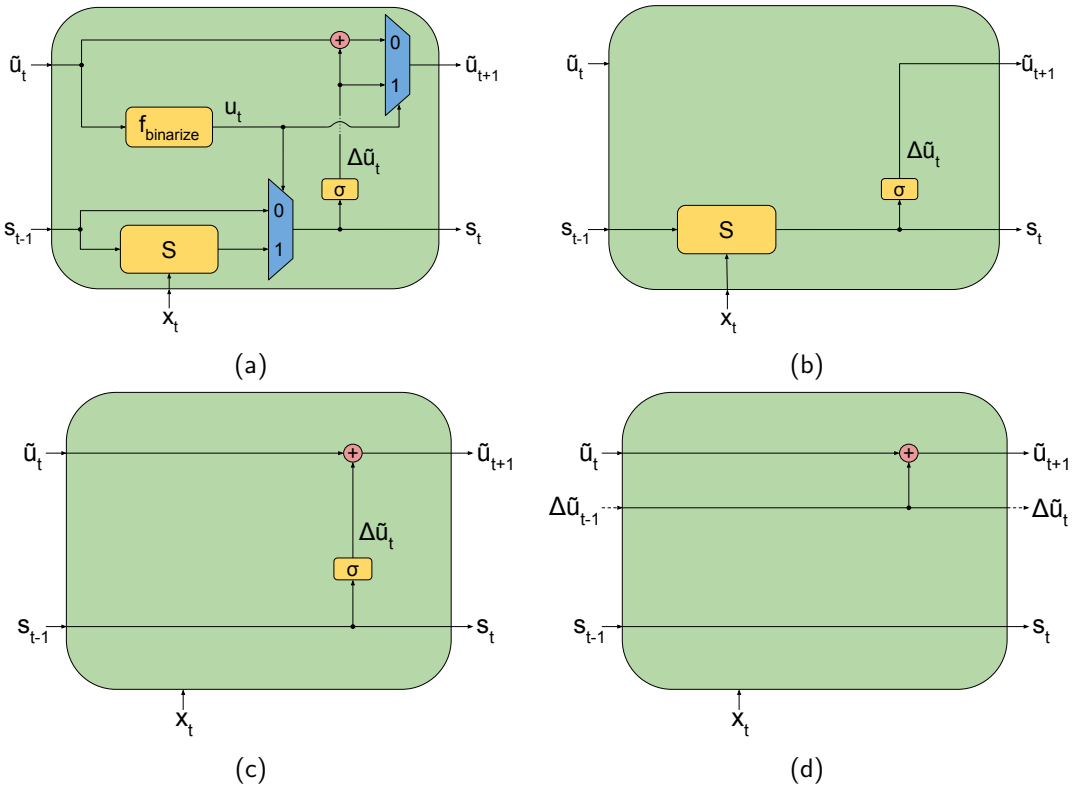

Figure 4.1: Model architecture of the proposed Skip RNN. **(a)** Complete Skip RNN architecture, where the computation graph at time step t is conditioned on u_t. **(b)** Architecture when the state is updated, i.e. $u_t = 1$. **(c)** Architecture when the update step is skipped and the previous state is copied, i.e. $u_t = 0$. **(d)** In practice, redundant computation is avoided by propagating $\Delta\tilde{u}_t$ between time steps when $u_t = 0$.

where $n \in \mathbb{Z}^+$. This enables more efficient implementations where no computation at all is performed whenever $u_t = 0$. These computational savings are possible because $\Delta\tilde{u}_t = \sigma(W_p s_t + b_p) = \sigma(W_p s_{t-1} + b_p) = \Delta\tilde{u}_{t-1}$ when $u_t = 0$ and there is no need to evaluate it again, as depicted in Figure 4.1d.

There are several advantages in reducing the number of RNN updates. From the computational standpoint, fewer updates translates into fewer required sequential operations to process an input signal, leading to faster inference and reduced energy consumption. Unlike some other models that aim to reduce the average number of operations per step [41, 27], ours enables skipping steps completely. Replacing RNN updates with copy operations increases the memory of the network and its ability to model long term dependencies even for gated units, since the exponential memory decay observed in LSTM and GRU [41] is alleviated. During training, gradients are propagated through fewer updating time steps, providing faster convergence in some tasks involving long sequences. Moreover, the proposed model is orthogonal to recent advances in RNNs and could be used in conjunction with such techniques, e.g. normalization [12, 4], regularization [61, 33], variable computation [27, 41] or even external memory [21, 54].

4.1 Error gradients

The whole model is differentiable except for $f_{binarize}$, which outputs binary values. A common method for optimizing functions involving discrete variables is REINFORCE [55], although several estimators have been proposed for the particular case of neurons with binary outputs [7]. We select the straight-through estimator [24], which consists in approximating the step function by the identity when computing gradients during the backwards pass:

$$\frac{\partial f_{binarize}(x)}{\partial x} = 1 \tag{4.7}$$

This yields a biased estimator that has proven more efficient than other unbiased but high-variance estimators such as REINFORCE [7] and has been successfully applied in different works [13, 11]. By using the straight-through estimator as the backwards pass for $f_{binarize}$, all the model parameters can be trained to minimize the target loss function with standard backpropagation and without defining any additional supervision or reward signal.

4.2 Limiting computation

The Skip RNN is able to learn when to update or copy the state without explicit information about which samples are useful to solve the task at hand. However, a different operating point on the trade-off between performance and number of processed samples may be required depending on the application, e.g. one may be willing to sacrifice a few accuracy points in order to run faster on machines with low computational power, or to reduce energy impact on portable devices. The proposed model can be encouraged to perform fewer state updates through additional loss terms, a common practice in neural networks with dynamically allocated computation [36, 38, 19, 27]. In particular, we will consider the introduction of a *cost per sample*:

$$L_{budget} = \lambda \cdot \sum_{t=1}^{T} u_t \tag{4.8}$$

where L_{budget} is the cost associated to a single sequence, λ is the cost per sample and T is the sequence length. This formulation bears a similarity to weight decay regularization, where the network is encouraged to slowly converge towards a solution where the norm of the weights is smaller. Similarly, in this case the network is encouraged to slowly converge towards a solution where fewer state updates are required.

Despite this formulation has been extensively studied in our experiments, different budget loss terms can be used depending on the application. For instance, a specific number of samples may be encouraged by applying an L_1 or L_2 loss between the target value and the number of updates per sequence, $\sum_{t=1}^{T} u_t$.

Chapter 5

Experimental Results

In the following sections, we investigate the advantages of adding the Skip RNN architecture to LSTMs and GRUs for a variety of tasks. Besides the evaluation metric for each task, we report the number of RNN state updates (i.e. the number of elements in the input sequence that are used by the model) as a measure of the computational load for each model. Unlike the wall clock time required by each model to process an input sequence, which we found to be highly dependent on the particular software implementation even for the baseline models, the number of RNN updates provides a meaningful measure of the inherently sequential operations in the computation graph. Since skipping an RNN update results in ignoring its corresponding input, we will refer to the number of updates and the number of used samples (i.e. elements in a sequence) interchangeably throughout this manuscript. For additional qualitative results, see Appendix A.

Training is performed with Adam optimizer [30], learning rate of 10^{-4}, $\beta_1 = 0.9$, $\beta_2 = 0.999$ and $\epsilon = 10^{-8}$ on batches of 256 samples. Gradient clipping [44] with a threshold of 1 is applied to all trainable variables. Bias b_p in Equation 4.4 is initialized to 1, so that all samples are used at the beginning of training[1]. The initial hidden state s_0 is learned during training, whereas $\tilde{u}_0$ is set to a constant value of 1 in order to force the first update at $t = 1$.

5.1 Adding Task

We revisit one of the original LSTM tasks [25], where the network is given a sequence of *(value, marker)* tuples. The desired output is the addition of the only two values that were marked with a 1, whereas those marked with a 0 need to be ignored. In particular, we follow the experimental setup by Neil et al. [41], where the first marker is randomly placed among the first 10% of samples (drawn with uniform probability) and the second one is placed among the last half of samples (drawn with uniform probability). This marker distribution yields sequences where at least 40% of the samples are distractors and provide no useful information at all. However, it is worth noting that in this task the risk of missing a marker is very large as compared to the benefits of working on shorter subsequences.

[1] In practice, forcing the network to use all samples at the beginning of training improves its robustness against random initializations of its weights and increases the reproducibility of the presented experiments. A similar behavior was observed in other augmented RNN architectures such as Neural Stacks [22].

Model	Task solved	State updates
LSTM	Yes	$100.0\% \pm 0.0\%$
Skip LSTM, $\lambda = 0$	Yes	$81.1\% \pm 3.6\%$
Skip LSTM, $\lambda = 10^{-5}$	Yes	$\mathbf{53.9\% \pm 2.1\%}$
GRU	Yes	$100.0\% \pm 0.0\%$
Skip GRU, $\lambda = 0$	Yes	$97.9\% \pm 3.2\%$
Skip GRU, $\lambda = 10^{-5}$	Yes	$\mathbf{50.7\% \pm 2.6\%}$

Table 5.1: Results for the adding task, displayed as $mean \pm std$ over four different runs. The task is considered to be solved if the MSE is at least two orders of magnitude below the variance of the output distribution.

We train RNN models with 110 units each on sequences of length 50, where the values are uniformly drawn from $\mathcal{U}(-0.5, 0.5)$. The final RNN state is fed to a fully connected layer that regresses the scalar output. The whole model is trained to minimize the Mean Squared Error (MSE) between the output and the ground truth. We consider that a model is able to solve the task when its MSE on a held-out set of examples is at least two orders of magnitude below the variance of the output distribution. This criterion is a stricter version of the one followed in [25].

While all models learn to solve the task, results in Table 5.1 show that the Skip RNN models are able to do so with roughly half of the updates of their standard counterparts. Interestingly, Skip LSTM tends to skip more updates than the Skip GRU when no cost per sample is set, behavior that may be related to the lack of output gate in the latter. We hypothesize that there are two possible reasons why the output gate makes the LSTM more prone to skipping updates: (a) it introduces an additional source of memory decay, and (b) it allows to mask out some units in the cell state that may specialize in deciding when to update or copy, making the final regression layer agnostic to such process.

We observed that the models using fewer updates never miss any marker, since the penalization in terms of MSE would be very large (see Figure 5.1 for examples). These models learn to skip most of the samples in the 40% of the sequence where there are no markers. Moreover, most updates are skipped once the second marker is found, since all the relevant information in the sequence has been already seen. This last pattern provides evidence that the proposed models effectively learn to decide whether to update or copy the hidden state based on the input sequence, as opposed to learning biases in the dataset only. As a downside, Skip RNN models show some difficulties skipping a large number of updates at once, probably due to the cumulative nature of $\tilde{u}_t$.

5.2 Frequency Discrimination Task

In this experiment, the network is trained to classify between sinusoids whose period is in range $T \sim \mathcal{U}(5, 6)$ milliseconds and those whose period is in range $T \sim \{(1, 5) \cup (6, 100)\}$ milliseconds [41]. Every sine wave with period T has a random phase shift drawn from $\mathcal{U}(0, T)$. At every time step, the input to the network is a single scalar representing the amplitude of the signal. Since sinusoid are continuous signals, this tasks allows to study whether Skip RNNs converge to the same solutions when their parameters are fixed but the sampling period is changed. We study

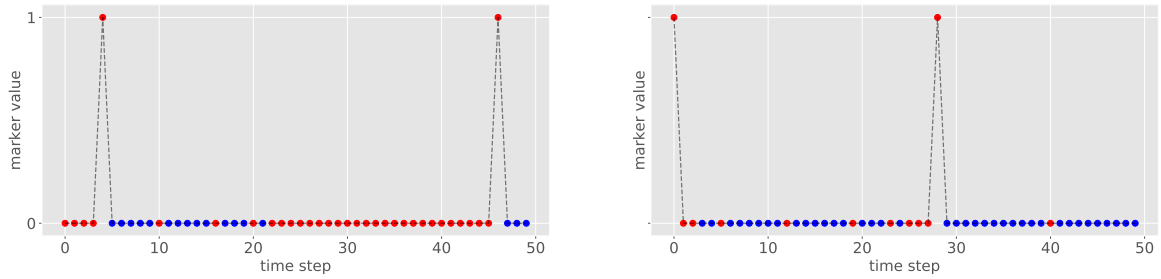

Figure 5.1: Sample usage examples for the Skip GRU with $\lambda = 10^{-5}$ on the adding task. Red dots indicate used samples, whereas blue ones are skipped.

Model	$\mathbf{T_s = 1}$**ms (length 100)**		$\mathbf{T_s = 0.5}$**ms (length 200)**	
	Task solved	**State updates**	**Task solved**	**State updates**
LSTM	Yes	100.0 ± 0.00	Yes	200.0 ± 0.00
Skip LSTM, $\lambda = 0$	Yes	55.5 ± 16.9	Yes	147.9 ± 27.0
Skip LSTM, $\lambda = 10^{-5}$	Yes	47.4 ± 14.1	Yes	50.7 ± 16.8
Skip LSTM, $\lambda = 10^{-4}$	Yes	12.7 ± 0.5	Yes	19.9 ± 1.5
GRU	Yes	100.0 ± 0.00	Yes	200.0 ± 0.00
Skip GRU, $\lambda = 0$	Yes	73.7 ± 17.9	Yes	167.0 ± 18.3
Skip GRU, $\lambda = 10^{-5}$	Yes	51.9 ± 10.2	Yes	54.2 ± 4.4
Skip GRU, $\lambda = 10^{-4}$	Yes	23.5 ± 6.2	Yes	22.5 ± 2.1

Table 5.2: Results for the frequency discrimination task, displayed as $mean \pm std$ over four different runs. The task is considered to be solved if the classification accuracy is over 99%. Models with the same cost per sample ($\lambda > 0$) converge to a similar number of used samples under different sampling conditions.

two different sampling periods, $T_s = \{0.5, 1\}$ milliseconds, for each set of hyperparameters.

We train RNNs with 110 units each on input signals of 100 milliseconds. Batches are stratified, containing the same number of samples for each class, yielding a 50% chance accuracy. The last state of the RNN is fed into a 2-way classifier and trained with cross-entropy loss. We consider that a model is able to solve the task when it achieves an accuracy over 99% on a held-out set of examples.

Table 5.2 summarizes results for this task. When no cost per sample is set ($\lambda = 0$), the number of updates differ under different sampling conditions. We attribute this behavior to the potentially large number of local minima in the cost function, since there are numerous subsampling patterns for which the task can be successfully solved and we are not explicitly encouraging the network to converge to a particular solution. On the other hand, when $\lambda > 0$ Skip RNN models with the same cost per sample use roughly the same number of input samples even when the sampling frequency is doubled. This is a desirable property, since solutions are robust to oversampled input signals.

Model	Accuracy	State updates
LSTM	0.910 ± 0.045	784.00 ± 0.00
Skip LSTM, $\lambda = 10^{-4}$	0.973 ± 0.002	379.38 ± 33.09
GRU	0.968 ± 0.013	784.00 ± 0.00
Skip GRU, $\lambda = 10^{-4}$	0.976 ± 0.003	392.62 ± 26.48

Table 5.3: Accuracy and used samples on the test set of MNIST after 600 epochs of training. Results are displayed as $mean \pm std$ over four different runs.

5.3 MNIST Classification from a Sequence of Pixels

The MNIST handwritten digits classification benchmark [35] is traditionally addressed with Convolutional Neural Networks (CNNs) that can efficiently exploit spatial dependencies through weight sharing. By flattening the 28×28 images into 784-d vectors, however, it can be reformulated as a challenging task for RNNs where long term dependencies need to be leveraged [34]. We follow the standard data split and set aside 5,000 training samples for validation purposes. After processing all pixels with an RNN with 110 units, the last hidden state is fed into a linear classifier predicting the digit class. All models are trained for 600 epochs to minimize cross-entropy loss.

Table 5.3 summarizes classification results on the test set after 600 epochs of training. Skip RNNs are not only able to solve the task using fewer updates than their counterparts, but also show a lower variation among runs and train faster (see Figure 5.2). We hypothesize that skipping updates make the Skip RNNs work on shorter subsequences, simplifying the optimization process and allowing the networks to capture long term dependencies more easily. A similar behavior was observed for Phased LSTM, where increasing the sparsity of cell updates accelerates training for very long sequences [41].

Sequences of pixels can be reshaped back into 2D images, allowing to visualize the samples used by the RNNs as a sort of hard visual attention model [57]. Examples such as the ones depicted in Figure 5.3 (top) show how the model learns to skip pixels that are not discriminative, such as the padding regions in the top and bottom of images. Similarly to the qualitative results for the adding task (Section 5.1), attended samples vary depending on the particular input being given to the network.

With the aim of observing how the attention model changes with the task, we train Skip LSTM ($\lambda = 10^{-4}$) models on two subsets of MNIST. We select pairs of visually similar digits: classes 1 and 7 (MNIST-17) and classes 5 and 6 (MNIST-56). Figure 5.3 (middle and bottom) shows how the network is able to find the most discriminative region for each pair of digits and mostly ignores the rest of the image. These regions differ from the ones attended by the model trained on the full dataset, depicted in Figure 5.3 (top).

5.4 Sentiment Analysis on IMDB

The IMDB dataset [37] contains 25,000 training and 25,000 testing movie reviews annotated into two classes, *positive* and *negative* sentiment, with an approximate average length of 240

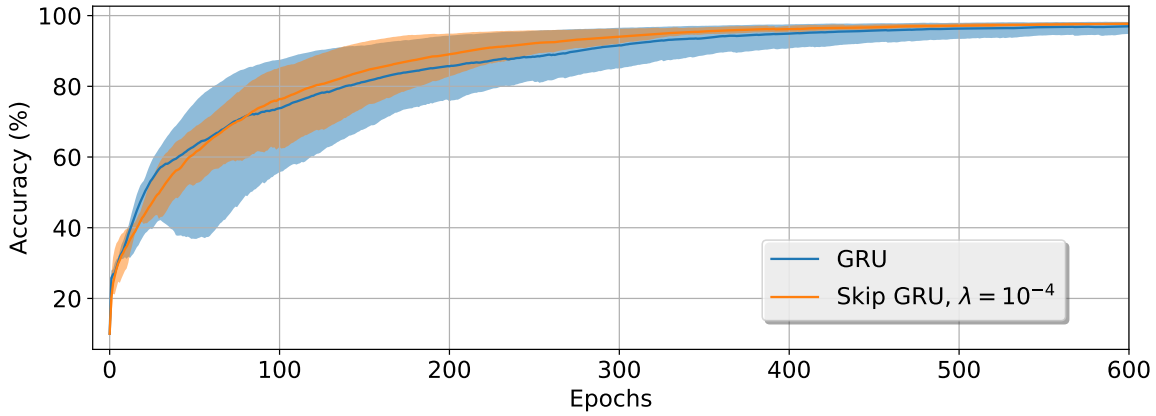

Figure 5.2: Accuracy evolution during training on the validation set of MNIST. The Skip GRU exhibits lower variance and faster convergence than the baseline GRU. A similar behavior is observed for LSTM and Skip LSTM, but omitted for clarity. Shading shows maximum and minimum over 4 runs, while dark lines indicate the mean.

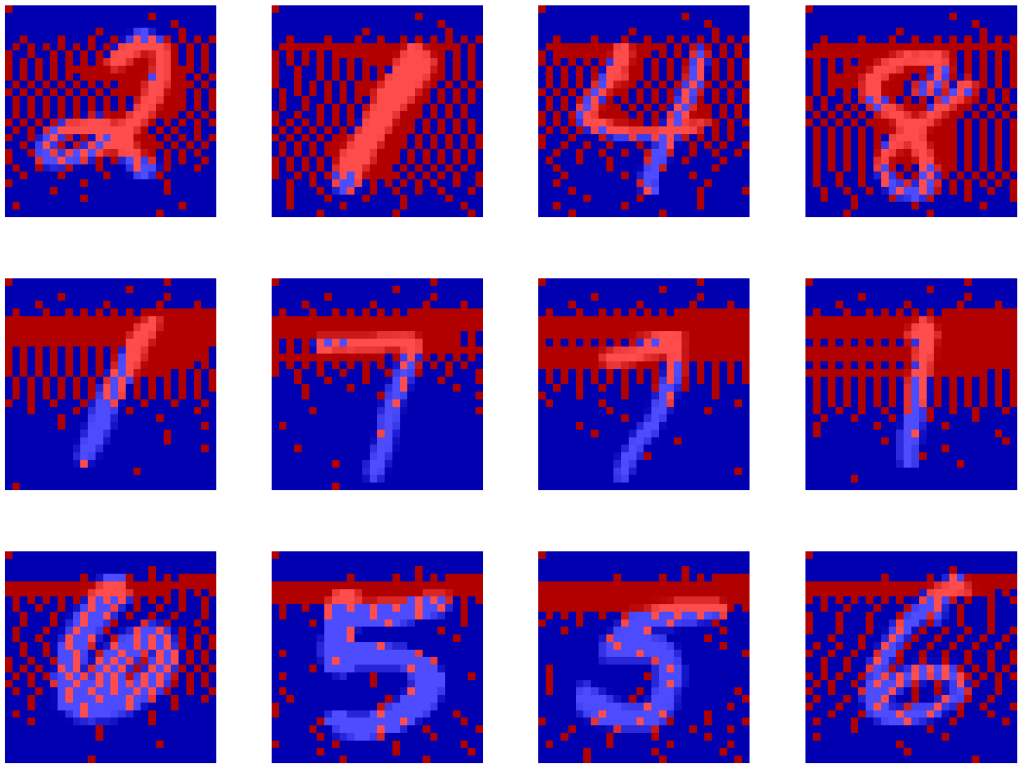

Figure 5.3: Sample usage examples for the Skip LSTM with $\lambda = 10^{-4}$ on the test set of MNIST (top), MNIST-17 (middle) and MNIST-56 (bottom). Red pixels are used, whereas blue ones are skipped.

Model	Length 200		Length 400	
	Accuracy	State updates	Accuracy	State updates
LSTM	0.843 ± 0.003	200.00 ± 0.00	0.868 ± 0.004	400.00 ± 0.00
Skip LSTM, $\lambda = 0$	0.844 ± 0.004	196.75 ± 5.63	0.866 ± 0.004	369.70 ± 19.35
Skip LSTM, $\lambda = 10^{-5}$	0.846 ± 0.004	197.15 ± 3.16	0.865 ± 0.001	380.62 ± 18.20
Skip LSTM, $\lambda = 10^{-4}$	0.837 ± 0.006	164.65 ± 8.67	0.862 ± 0.003	186.30 ± 25.72
Skip LSTM, $\lambda = 10^{-3}$	0.811 ± 0.007	73.85 ± 1.90	0.836 ± 0.007	84.22 ± 1.98
GRU	0.845 ± 0.006	200.00 ± 0.00	0.862 ± 0.003	400.00 ± 0.00
Skip GRU, $\lambda = 0$	0.848 ± 0.002	200.00 ± 0.00	0.866 ± 0.002	399.02 ± 1.69
Skip GRU, $\lambda = 10^{-5}$	0.842 ± 0.005	199.25 ± 1.30	0.862 ± 0.008	398.00 ± 2.06
Skip GRU, $\lambda = 10^{-4}$	0.834 ± 0.006	180.97 ± 8.90	0.853 ± 0.011	314.30 ± 2.82
Skip GRU, $\lambda = 10^{-3}$	0.800 ± 0.007	106.15 ± 37.92	0.814 ± 0.005	99.12 ± 2.69

Table 5.4: Accuracy and used samples on the test set of IMDB for different sequence lengths. Results are displayed as $mean \pm std$ over four different runs.

words per review. We set aside 15% of training data for validation purposes. Words are embedded into 300-d vector representations before being fed to an RNN with 128 units. The embedding matrix is initialized using pre-trained word2vec[2] embeddings [39] when available, or random vectors drawn from $\mathcal{U}(-0.25, 0.25)$ otherwise [29]. Dropout with rate 0.2 is applied between the last RNN state and the classification layer in order to reduce overfitting. We evaluate the models on sequences of length 200 and 400 by cropping longer sequences and padding shorter ones [59].

Results on the test are reported in Table 5.4. In a task where it is hard to predict which input tokens will be discriminative, the Skip RNN models are able to achieve similar accuracy rates to the baseline models while reducing the number of required updates. These results highlight the trade-off between accuracy and the available computational budget, since a larger cost per sample results in lower accuracies. However, allowing the network to select which samples to use instead of cropping sequences at a given length boosts performance, as observed for the Skip LSTM (length 400, $\lambda = 10^{-4}$), which achieves a higher accuracy than the baseline LSTM (length 200) while seeing roughly the same number of words per review. A similar behavior can be seen for the Skip RNN models with $\lambda = 10^{-3}$, where allowing them to select words from longer reviews boosts classification accuracy while using a comparable number of tokens per sequence.

5.5 Action classification on UCF-101

One of the most accurate and scalable pipelines for video analysis consists in extracting frame level features with a CNN and modeling their temporal evolution with an RNN [16, 60]. Videos are commonly recorded at high sampling rates, rapidly generating long sequences with strong temporal redundancy that are challenging for RNNs. Moreover, processing frames with a CNN is computationally expensive and may become prohibitive for high framerates. These issues have been alleviated in previous works by using short clips [16] or by downsampling the original data in order to cover long temporal spans without increasing the sequence length excessively [60]. Instead of addressing the long sequence problem at the input data level, we train RNN models

[2] `https://code.google.com/archive/p/word2vec/`

Model	Accuracy	State updates
LSTM	0.671	250.0
Skip LSTM, $\lambda = 0$	0.749	138.9
Skip LSTM, $\lambda = 10^{-5}$	0.757	24.2
Skip LSTM, $\lambda = 10^{-4}$	0.790	7.6
GRU	0.791	250.0
Skip GRU, $\lambda = 0$	0.796	124.2
Skip GRU, $\lambda = 10^{-5}$	0.792	29.7
Skip GRU, $\lambda = 10^{-4}$	0.793	23.7

Table 5.5: Accuracy and used samples on the validation set of UCF-101 (split 1).

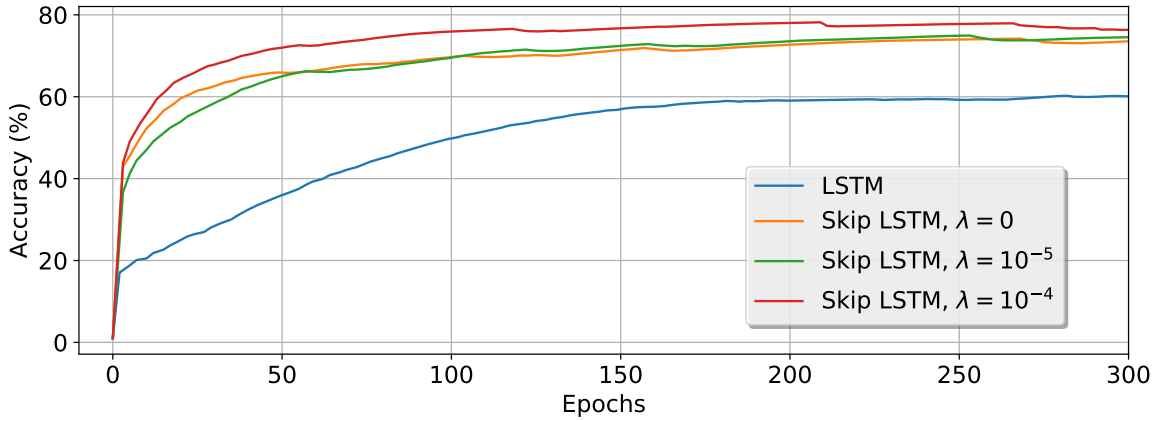

Figure 5.4: Accuracy evolution during the first 300 training epochs on the validation set of UCF-101 (split 1). Skip LSTM models converge much faster than the baseline LSTM.

using long frame sequences without downsampling and let the network learn which frames need to be used.

UCF-101 [51] is a dataset containing 13,320 trimmed videos belonging to 101 different action categories. We use 10 seconds of video sampled at 25fps, cropping longer ones and padding shorter examples with empty frames. Activations in the Global Average Pooling layer from a ResNet-50 [23] CNN pretrained on the ImageNet dataset [15] are used as frame level features, which are fed into two stacked RNN layers with 512 units each. The weights in the CNN are not tuned during training to reduce overfitting. The hidden state in the last RNN layer is used to compute the update probability for the Skip RNN models.

We evaluate the different models on the first split of UCF-101 and report results in Table 5.5. Skip RNN models do not only improve the classification accuracy with respect to the baseline, but require very few updates to do so. We argue that this is possible due to the low motion between consecutive frames resulting in frame level features with high temporal redundancy [48]. Moreover, Figure 5.4 shows how models performing fewer updates converge faster thanks to the gradients being preserved during longer spans when training with backpropagation through time.

Chapter 6

Conclusion and Future Work

We presented Skip RNN as an extension to existing recurrent architectures enabling them to skip state updates, thus reducing the number of sequential operations in the computation graph. Unlike previous approaches, all parameters in Skip RNN are trained with backpropagation without requiring to introduce task-dependent modifications. Experiments conducted on LSTM and GRU showed how Skip RNNs match the performance of baseline models while relaxing their computational requirements. Evaluation is performed on a broad set of tasks and datasets involving synthetic data, images, text and video. Skip RNNs provide a faster and more stable training for long sequences and complex models, likely due to gradients being backpropagated through fewer time steps resulting in a simpler optimization task. Moreover, the introduced computational savings are better suited for modern hardware than those methods that reduce the number of computations required at each time step [31, 41, 11].

The presented results motivate several new research directions towards designing efficient RNN architectures. Introducing stochasticity in neural network training has proven beneficial [52, 33], so that replacing the deterministic rounding operation with a stochastic sampling is a natural extension to the current formulation. This may improve generalization while allowing the model to escape poor local minima more easily. The loss term penalizing the number of updates is an important element in the performance of Skip RNN and different formulations can be tailored for specific tasks. For instance, the cost could increase at each time step to encourage the network to emit a decision earlier [1], or a particular number of updates can be enforced if such information is available. Finally, understanding and analyzing the patterns followed by the model when deciding whether to update or copy the RNN state may provide insight for developing better and more efficient architectures.

This project was aimed at designing a recurrent architecture capable of skipping state updates and showing its advantages in a series of tasks. Once the formulation is mature enough and has been proven effective, future work should focus on developing efficient implementations pushing the limits in terms of wall clock time. This may include the development of low level CUDA kernels that speedup the execution of Skip RNN models, as well as porting the code to a framework based on dynamic computational graphs (e.g. PyTorch[1], DyNet[2], Chainer[3]). The latter are more suitable for the problem at hand than frameworks based on static computation graphs, since the

[1] `http://www.pytorch.org/`
[2] `https://github.com/clab/dynet/`
[3] `https://chainer.org/`

actual computation graph is conditioned on the model outputs at each time step and cannot be completely defined before being run.

Bibliography

[1] Mohammad Sadegh Aliakbarian, Fatemehsadat Saleh, Mathieu Salzmann, Basura Fernando, Lars Petersson, and Lars Andersson. Encouraging LSTMs to anticipate actions very early. *arXiv preprint arXiv:1703.07023*, 2017.

[2] Amjad Almahairi, Nicolas Ballas, Tim Cooijmans, Yin Zheng, Hugo Larochelle, and Aaron Courville. Dynamic capacity networks. In *ICML*, 2016.

[3] Jimmy Ba, Volodymyr Mnih, and Koray Kavukcuoglu. Multiple object recognition with visual attention. *arXiv preprint arXiv:1412.7755*, 2014.

[4] Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. Layer normalization. *arXiv preprint arXiv:1607.06450*, 2016.

[5] Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. Neural machine translation by jointly learning to align and translate. In *ICLR*, 2015.

[6] Yoshua Bengio. Deep learning of representations: Looking forward. In *SLSP*, 2013.

[7] Yoshua Bengio, Nicholas Léonard, and Aaron Courville. Estimating or propagating gradients through stochastic neurons for conditional computation. *arXiv preprint arXiv:1308.3432*, 2013.

[8] Yoshua Bengio, Patrice Simard, and Paolo Frasconi. Learning long-term dependencies with gradient descent is difficult. *IEEE Transactions on Neural Networks*, 1994.

[9] Subhabrata Bhattacharya, Felix X Yu, and Shih-Fu Chang. Minimally needed evidence for complex event recognition in unconstrained videos. In *ICMR*, 2014.

[10] Kyunghyun Cho, Bart Van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. Learning phrase representations using rnn encoder-decoder for statistical machine translation. In *EMNLP*, 2014.

[11] Junyoung Chung, Sungjin Ahn, and Yoshua Bengio. Hierarchical multiscale recurrent neural networks. In *ICLR*, 2017.

[12] Tim Cooijmans, Nicolas Ballas, César Laurent, Çağlar Gülçehre, and Aaron Courville. Recurrent batch normalization. In *ICLR*, 2017.

[13] Matthieu Courbariaux, Itay Hubara, Daniel Soudry, Ran El-Yaniv, and Yoshua Bengio. Binarized neural networks: Training deep neural networks with weights and activations constrained to+ 1 or-1. *arXiv preprint arXiv:1602.02830*, 2016.

[14] Andrew Davis and Itamar Arel. Low-rank approximations for conditional feedforward computation in deep neural networks. *arXiv preprint arXiv:1312.4461*, 2013.

[15] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. ImageNet: A large-scale hierarchical image database. In *CVPR*, 2009.

[16] Jeffrey Donahue, Lisa Anne Hendricks, Sergio Guadarrama, Marcus Rohrbach, Subhashini Venugopalan, Kate Saenko, and Trevor Darrell. Long-term recurrent convolutional networks for visual recognition and description. In *CVPR*, 2015.

[17] Michael Figurnov, Maxwell D Collins, Yukun Zhu, Li Zhang, Jonathan Huang, Dmitry Vetrov, and Ruslan Salakhutdinov. Spatially adaptive computation time for residual networks. In *CVPR*, 2017.

[18] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep learning*. MIT press, 2016.

[19] Alex Graves. Adaptive computation time for recurrent neural networks. *arXiv preprint arXiv:1603.08983*, 2016.

[20] Alex Graves, Abdel-rahman Mohamed, and Geoffrey Hinton. Speech recognition with deep recurrent neural networks. In *ICASSP*, 2013.

[21] Alex Graves, Greg Wayne, and Ivo Danihelka. Neural turing machines. *arXiv preprint arXiv:1410.5401*, 2014.

[22] Edward Grefenstette, Karl Moritz Hermann, Mustafa Suleyman, and Phil Blunsom. Learning to transduce with unbounded memory. In *NIPS*, 2015.

[23] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016.

[24] Geoffrey Hinton. Neural networks for machine learning. Coursera video lectures, 2012.

[25] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 1997.

[26] Gao Huang, Yu Sun, Zhuang Liu, Daniel Sedra, and Kilian Q Weinberger. Deep networks with stochastic depth. In *ECCV*, 2016.

[27] Yacine Jernite, Edouard Grave, Armand Joulin, and Tomas Mikolov. Variable computation in recurrent neural networks. In *ICLR*, 2017.

[28] Andrej Karpathy, George Toderici, Sanketh Shetty, Thomas Leung, Rahul Sukthankar, and Li Fei-Fei. Large-scale video classification with convolutional neural networks. In *CVPR*, 2014.

[29] Yoon Kim. Convolutional neural networks for sentence classification. In *EMNLP*, 2014.

[30] Diederik Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

[31] Jan Koutnik, Klaus Greff, Faustino Gomez, and Juergen Schmidhuber. A clockwork rnn. In *ICML*, 2014.

[32] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *NIPS*, 2012.

[33] David Krueger, Tegan Maharaj, János Kramár, Mohammad Pezeshki, Nicolas Ballas, Nan Rosemary Ke, Anirudh Goyal, Yoshua Bengio, Hugo Larochelle, Aaron Courville, et al. Zoneout: Regularizing rnns by randomly preserving hidden activations. In *ICLR*, 2017.

[34] Quoc V Le, Navdeep Jaitly, and Geoffrey E Hinton. A simple way to initialize recurrent networks of rectified linear units. *arXiv preprint arXiv:1504.00941*, 2015.

[35] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 1998.

[36] Lanlan Liu and Jia Deng. Dynamic deep neural networks: Optimizing accuracy-efficiency trade-offs by selective execution. *arXiv preprint arXiv:1701.00299*, 2017.

[37] Andrew L Maas, Raymond E Daly, Peter T Pham, Dan Huang, Andrew Y Ng, and Christopher Potts. Learning word vectors for sentiment analysis. In *ACL*, 2011.

[38] Mason McGill and Pietro Perona. Deciding how to decide: Dynamic routing in artificial neural networks. In *ICML*, 2017.

[39] Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. Distributed representations of words and phrases and their compositionality. In *NIPS*, 2013.

[40] Volodymyr Mnih, Nicolas Heess, Alex Graves, et al. Recurrent models of visual attention. In *NIPS*, 2014.

[41] Daniel Neil, Michael Pfeiffer, and Shih-Chii Liu. Phased LSTM: accelerating recurrent network training for long or event-based sequences. In *NIPS*, 2016.

[42] Christopher Olah. Understanding LSTM networks. `http://colah.github.io/posts/2015-08-Understanding-LSTMs/`, 2015.

[43] Aaron van den Oord, Sander Dieleman, Heiga Zen, Karen Simonyan, Oriol Vinyals, Alex Graves, Nal Kalchbrenner, Andrew Senior, and Koray Kavukcuoglu. Wavenet: A generative model for raw audio. *arXiv preprint arXiv:1609.03499*, 2016.

[44] Razvan Pascanu, Tomas Mikolov, and Yoshua Bengio. On the difficulty of training recurrent neural networks. In *ICML*, 2013.

[45] Colin Raffel and Dieterich Lawson. Training a subsampling mechanism in expectation. In *ICLR Workshop Track*, 2017.

[46] Frank Rosenblatt. The perceptron: A probabilistic model for information storage and organization in the brain. *Psychological review*, 1958.

[47] Noam Shazeer, Azalia Mirhoseini, Krzysztof Maziarz, Andy Davis, Quoc Le, Geoffrey Hinton, and Jeff Dean. Outrageously large neural networks: The sparsely-gated mixture-of-experts layer. In *ICLR*, 2017.

[48] Evan Shelhamer, Kate Rakelly, Judy Hoffman, and Trevor Darrell. Clockwork convnets for video semantic segmentation. *arXiv preprint arXiv:1608.03609*, 2016.

[49] Gunnar A Sigurdsson, Xinlei Chen, and Abhinav Gupta. Learning visual storylines with skipping recurrent neural networks. In *ECCV*, 2016.

[50] David Silver, Aja Huang, Chris J Maddison, Arthur Guez, Laurent Sifre, George Van Den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, et al. Mastering the game of go with deep neural networks and tree search. *Nature*, 2016.

[51] Khurram Soomro, Amir Roshan Zamir, and Mubarak Shah. Ucf101: A dataset of 101 human actions classes from videos in the wild. *arXiv preprint arXiv:1212.0402*, 2012.

[52] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: A simple way to prevent neural networks from overfitting. *The Journal of Machine Learning Research*, 2014.

[53] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. Going deeper with convolutions. In *CVPR*, 2015.

[54] Jason Weston, Sumit Chopra, and Antoine Bordes. Memory networks. *arXiv preprint arXiv:1410.3916*, 2014.

[55] Ronald J Williams. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 1992.

[56] Yonghui Wu, Mike Schuster, Zhifeng Chen, Quoc V Le, Mohammad Norouzi, Wolfgang Macherey, Maxim Krikun, Yuan Cao, Qin Gao, Klaus Macherey, et al. Google's neural machine translation system: Bridging the gap between human and machine translation. *arXiv preprint arXiv:1609.08144*, 2016.

[57] Kelvin Xu, Jimmy Ba, Ryan Kiros, Kyunghyun Cho, Aaron Courville, Ruslan Salakhudinov, Rich Zemel, and Yoshua Bengio. Show, attend and tell: Neural image caption generation with visual attention. In *ICML*, 2015.

[58] Serena Yeung, Olga Russakovsky, Greg Mori, and Li Fei-Fei. End-to-end learning of action detection from frame glimpses in videos. In *CVPR*, 2016.

[59] Adams Wei Yu, Hongrae Lee, and Quoc V Le. Learning to skim text. In *ACL*, 2017.

[60] Joe Yue-Hei Ng, Matthew Hausknecht, Sudheendra Vijayanarasimhan, Oriol Vinyals, Rajat Monga, and George Toderici. Beyond short snippets: Deep networks for video classification. In *CVPR*, 2015.

[61] Wojciech Zaremba, Ilya Sutskever, and Oriol Vinyals. Recurrent neural network regularization. In *ICLR*, 2015.

Appendix A

Qualitative results

This appendix contains additional qualitative results for the Skip RNN models.

A.1 Adding task

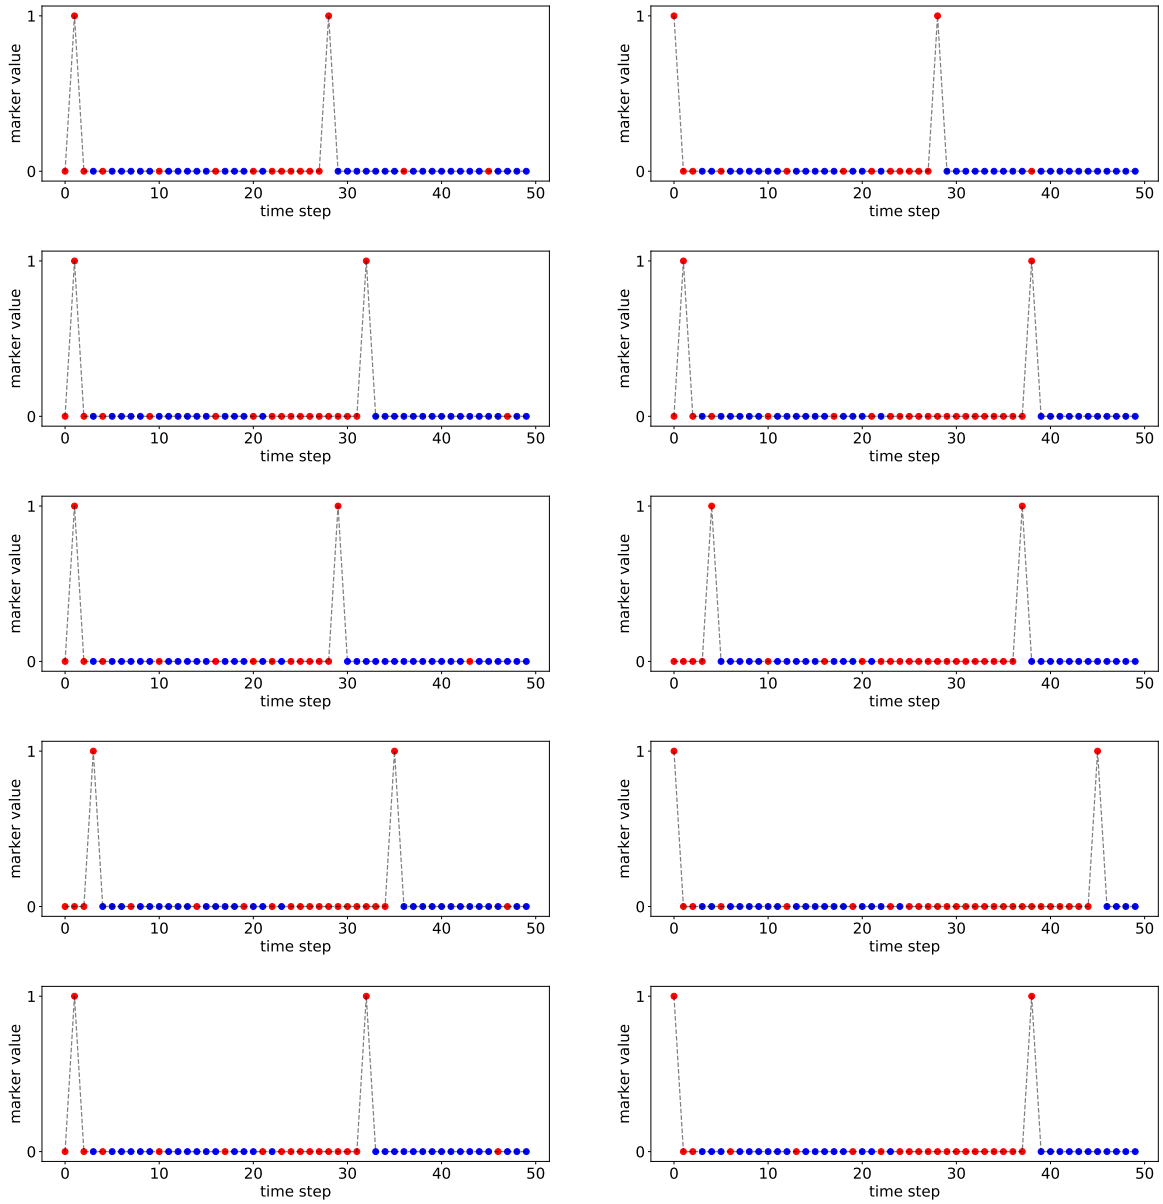

Figure A.1: Sample usage examples for the Skip GRU with $\lambda = 10^{-5}$ on the adding task. Red dots indicate used samples, whereas blue ones are skipped.

A.2 Frequency discrimination task

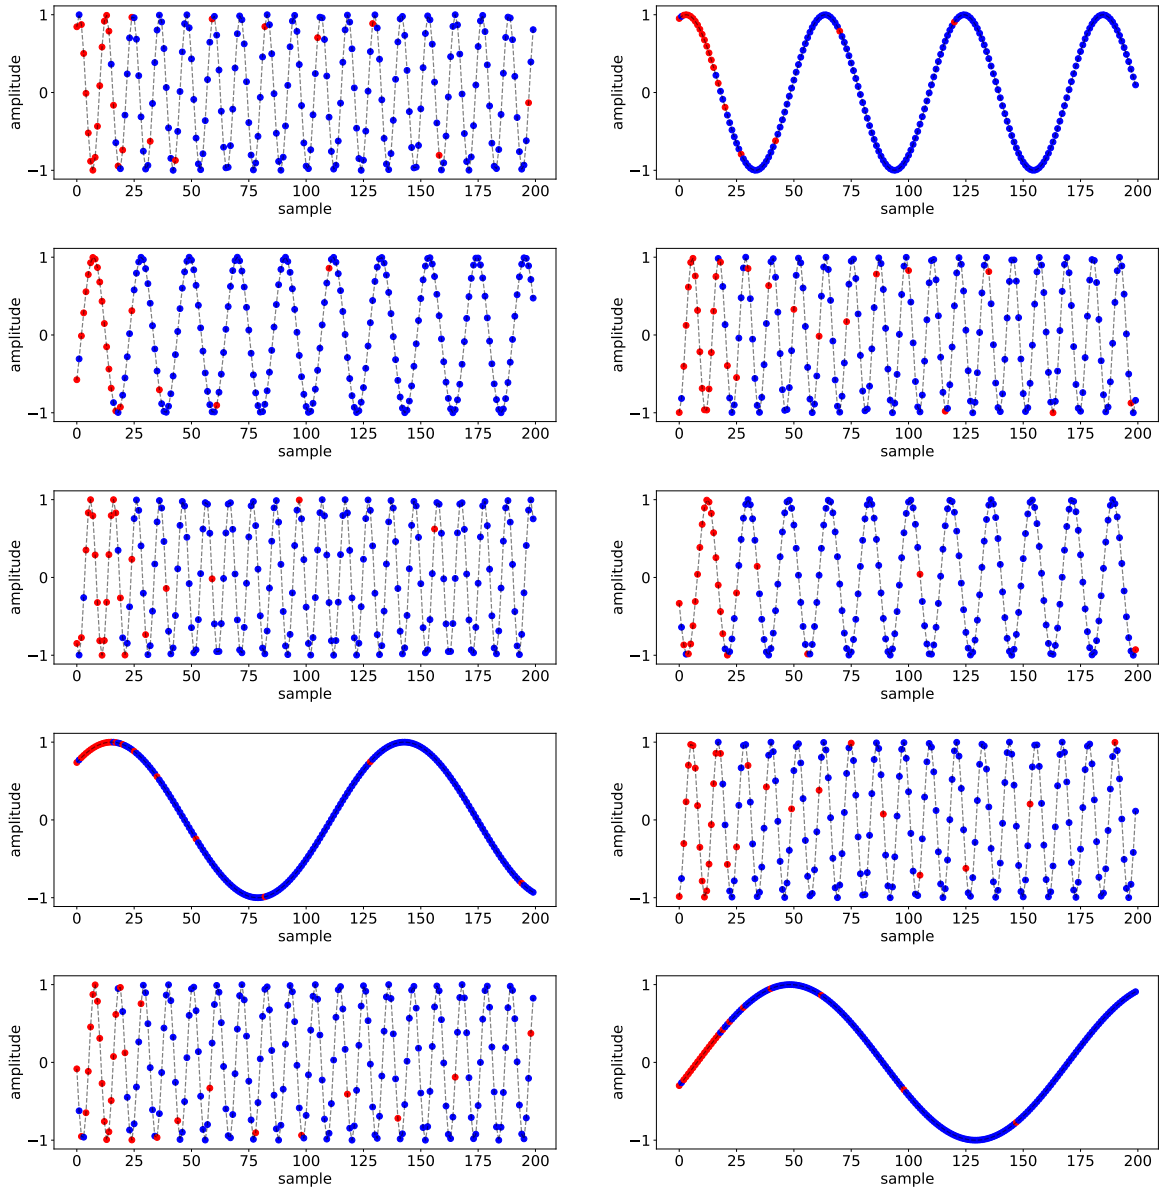

Figure A.2: Sample usage examples for the Skip LSTM with $\lambda = 10^{-4}$ on the frequency discrimination task with $T_s = 0.5$ms. Red dots indicate used samples, whereas blue ones are skipped. The network learns that using the first samples is enough to classify the frequency of the sine waves, in contrast to a uniform downsampling that may result in aliasing.

Appendix B

Scientific publications

This work was submitted and accepted at two of the main venues in Machine Learning. A short paper was presented at the Time Series Workshop, held at the Conference on Neural Information Processing Systems (NIPS), where it was selected as a contributed talk. It was later accepted as a main conference paper at the International Conference on Learning Representations (ICLR). Both papers are appended to this document.

Skip RNN: Learning to Skip State Updates in Recurrent Neural Networks

Víctor Campos[*†], **Brendan Jou**[‡], **Xavier Giró-i-Nieto**[§], **Jordi Torres**[†], **Shih-Fu Chang**[Γ]
[†]Barcelona Supercomputing Center, [‡]Google Inc,
[§]Universitat Politècnica de Catalunya, [Γ]Columbia University
{victor.campos, jordi.torres}@bsc.es, bjou@google.com,
xavier.giro@upc.edu, sfchang@ee.columbia.edu

Abstract

Recurrent Neural Networks (RNNs) continue to show outstanding performance in sequence modeling tasks. However, training RNNs on long sequences often face challenges like slow inference, vanishing gradients and difficulty in capturing long term dependencies. In backpropagation through time settings, these issues are tightly coupled with the large, sequential computational graph resulting from unfolding the RNN in time. We introduce the Skip RNN model which extends existing RNN models by learning to skip state updates and shortens the effective size of the computational graph. This model can also be encouraged to perform fewer state updates through a budget constraint. We evaluate the proposed model on various tasks and show how it can reduce the number of required RNN updates while preserving, and sometimes even improving, the performance of the baseline RNN models. Source code is publicly available at https://imatge-upc.github.io/skiprnn-2017-telecombcn/.

1 Introduction

Some of the main limitations of Recurrent Neural Networks (RNNs) are their challenging training and deployment when dealing with long sequences, due to their inherently sequential behaviour. These challenges include throughput degradation, slower convergence during training and memory leakage, even for gated architectures [18]. The main contribution of this work is a novel modification for existing RNN architectures that allows them to skip state updates, decreasing the number of sequential operations to be performed, without requiring any additional supervision signal. This model, called Skip RNN, adaptively determines whether the state needs to be updated or copied to the next time step, thereby allow a "skip" in the computation graph. We show how the network can be encouraged to perform fewer state updates by adding a penalization term during training, allowing us to train models of different target computation budgets. The proposed modification is implemented on top of well known RNN architectures, namely LSTM and GRU, and the resulting models show promising results in a series of sequence modeling tasks.

Conditional computation has been shown to allow gradual increases in model capacity without a proportional increases in computational cost by exploiting certain computation paths for each input [3, 16, 1, 17, 20]. This idea has been extended in the temporal domain by building RNNs that perform different amount of computation at each time step [6, 4, 10, 12, 18]. However, due to the inherently sequential nature of RNNs and the parallel computation capabilities of modern hardware, reducing the size of the matrices involved in the computations performed at each time step does not accelerate inference. The proposed Skip RNN model can be seen as form of conditional computation in time,

[*]Work done while Víctor Campos was a visiting scholar at Columbia University.

31st Conference on Neural Information Processing Systems (NIPS 2017), Long Beach, CA, USA.

where the computation associated to the RNN updates may or may not be executed at every time step, effectively reducing sequential computation and shielding the hidden state over longer time lags. It resembles LSTM-Jump [24], an LSTM cell augmented with a classification layer that will decide how many steps to jump between RNN updates. This added layer needs to be trained with REINFORCE [22] and some hyperparameters define a reduced set of subsequences that the model can sample, instead of allowing the network to learn any arbitrary sampling scheme. Unlike LSTM-Jump, our proposed approach is differentiable, thus not requiring any modifications to the loss function and simplifying the optimization process, and is not limited to a predefined set of sample selection patterns.

2 Model Description

An RNN takes an input sequence $\mathbf{x} = (x_1, \ldots, x_T)$ and generates a state sequence $\mathbf{s} = (s_1, \ldots, s_T)$ by iteratively applying a parametric state transition model S from $t = 1$ to T:

$$s_t = S(s_{t-1}, x_t) \tag{1}$$

We augment the network with a binary *state update gate*, $u_t \in \{0, 1\}$, selecting whether the state of the RNN will be updated or copied from the previous time step. At every time step t, the probability $\tilde{u}_{t+1} \in [0, 1]$ of performing a state update at $t + 1$ is emitted. The model formulation implements the observation that the likelihood of requesting a new input increases with the number of consecutively skipped samples:

$$u_t = f_{binarize}(\tilde{u}_t) \tag{2}$$
$$s_t = u_t \cdot S(s_{t-1}, x_t) + (1 - u_t) \cdot s_{t-1} \tag{3}$$
$$\Delta\tilde{u}_t = \sigma(W_p s_t + b_p) \tag{4}$$
$$\tilde{u}_{t+1} = u_t \cdot \Delta\tilde{u}_t + (1 - u_t) \cdot (\tilde{u}_t + \min(\Delta\tilde{u}_t, 1 - \tilde{u}_t)) \tag{5}$$

where σ is the sigmoid function and $f_{binarize} : [0, 1] \to \{0, 1\}$ binarizes the input value. Should the network be composed of several layers, some columns of W_p can be fixed to 0 so that $\Delta\tilde{u}_t$ depends only on the states of a subset of layers. We implement $f_{binarize}$ as a deterministic step function $u_t = \text{round}(\tilde{u}_t)$, although a stochastic sampling from a Bernoulli distribution $u_t \sim \text{Bernoulli}(\tilde{u}_t)$ would be possible as well.

The number of skipped time steps can be computed ahead of time. For the particular formulation used in this work, where $f_{binarize}$ is implemented by means of a rounding function, the number of skipped samples after performing a state update at time step t is given by:

$$N_{skip}(t) = \min\{n : n \cdot \Delta\tilde{u}_t \geq 0.5\} - 1 \tag{6}$$

where $n \in \mathbb{Z}^+$. This enables more efficient implementations where no computation at all is performed whenever $u_t = 0$. These computational savings are possible because $\Delta\tilde{u}_t = \sigma(W_p s_t + b_p) = \sigma(W_p s_{t-1} + b_p) = \Delta\tilde{u}_{t-1}$ when $u_t = 0$ and there is no need to evaluate it again.

There are several advantages in reducing the number of RNN updates. From the computational standpoint, fewer updates translates into fewer required sequential operations to process an input signal, leading to faster inference and reduced energy consumption. Unlike some other models that aim to reduce the average number of operations per step [18, 10], ours enables skipping steps completely. Replacing RNN updates with copy operations increases the memory of the network and its ability to model long term dependencies even for gated units, since the exponential memory decay observed in LSTM and GRU [18] is alleviated. During training, gradients are propagated through fewer updating time steps, providing faster convergence in some tasks involving long sequences. Moreover, the proposed model is orthogonal to recent advances in RNNs and could be used in conjunction with such techniques, e.g. normalization [5, 2], regularization [25, 13], variable computation [10, 18] or even external memory [7, 21].

2.1 Error gradients

The whole model is differentiable except for $f_{binarize}$, which outputs binary values. We define its gradients following the straight-through estimator [9], so that all the model parameters can be

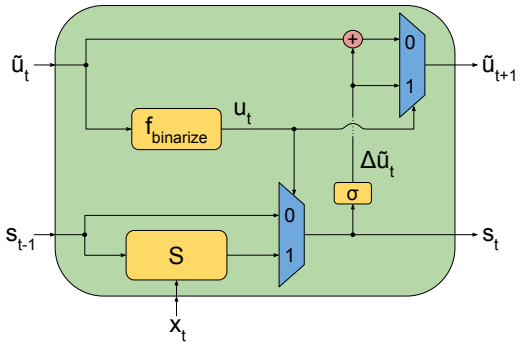

Figure 1: Model architecture of the proposed Skip RNN, where the computation graph at time step t is conditioned on u_t. In practice, redundant computation is avoided by propagating $\Delta\tilde{u}_t$ between time steps when $u_t = 0$.

trained to minimize the target loss function with standard backpropagation and without defining any additional supervision or reward signal. This estimator consists in approximating the step function by the identity when computing gradients during the backward pass:

$$\frac{\partial f_{binarize}(x)}{\partial x} = 1 \tag{7}$$

2.2 Limiting computation

The Skip RNN is able to learn when to update or copy the state without explicit information about which samples are useful to solve the task at hand. However, a different operating point on the trade-off between performance and number of processed samples may be required depending on the application, e.g. one may be willing to sacrifice a few accuracy points in order to run faster on machines with low computational power, or to reduce energy impact on portable devices. The proposed model can be encouraged to perform fewer state updates through additional loss terms, a common practice in neural networks with dynamically allocated computation [16, 17, 6, 10]. In particular, we consider a *cost per sample*:

$$L_{budget} = \lambda \cdot \sum_{t=1}^{T} u_t \tag{8}$$

where L_{budget} is the cost associated to a single sequence, λ is the cost per sample and T is the sequence length.

3 Experimental results: MNIST Classification from a Sequence of Pixels

The MNIST handwritten digits classification benchmark [15] is traditionally addressed with Convolutional Neural Networks (CNNs) that can efficiently exploit spatial dependencies through weight sharing. By flattening the 28×28 images into 784-d vectors, however, it can be reformulated as a challenging task for RNNs where long term dependencies need to be leveraged [14]. We follow the standard data split and set aside 5,000 training samples for validation purposes. After processing all pixels with an RNN with 110 units, the last hidden state is fed into a linear classifier predicting the digit class. All models are trained for 600 epochs to minimize cross-entropy loss. For more details on the experimental setup, see Appendix A.

With the goal of studying the effect of skipping state updates on the learning capability of the networks, we introduce a new baseline which skips a state update with probability p_{skip}. We tune the skipping probability to obtain models that perform a similar number of state updates to the Skip RNN models. Table 1 summarizes classification results on the test set after 600 epochs of training. Skip RNNs are not only able to solve the task using fewer updates than their counterparts, but also

Model	Accuracy	State updates
LSTM	0.910 ± 0.045	784.00 ± 0.00
LSTM ($p_{skip} = 0.5$)	0.893 ± 0.003	392.03 ± 0.05
Skip LSTM, $\lambda = 10^{-4}$	0.973 ± 0.002	379.38 ± 33.09
GRU	0.968 ± 0.013	784.00 ± 0.00
GRU ($p_{skip} = 0.5$)	0.912 ± 0.004	391.86 ± 0.14
Skip GRU, $\lambda = 10^{-4}$	0.976 ± 0.003	392.62 ± 26.48

Table 1: Accuracy and used samples on the test set of MNIST after 600 epochs of training. Results are displayed as $mean \pm std$ over four different runs.

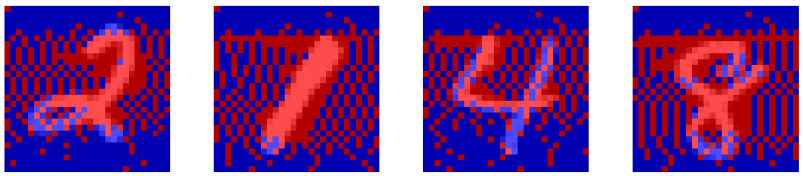

Figure 2: Sample usage examples for the Skip LSTM with $\lambda = 10^{-4}$ on the test set of MNIST. Red pixels are used, whereas blue ones are skipped.

show a lower variation among runs and train faster. We hypothesize that skipping updates make the Skip RNNs work on shorter subsequences, simplifying the optimization process and allowing the networks to capture long term dependencies more easily. A similar behavior was observed for Phased LSTM, where increasing the sparsity of cell updates accelerates training for very long sequences [18]. However, the drop in performance observed in the models where the state updates are skipped randomly suggests that learning which samples to use is a key component in the performance of Skip RNN.

Sequences of pixels can be reshaped back into 2D images, allowing to visualize the samples used by the RNNs as a sort of hard visual attention model [23]. Examples such as the ones depicted in Figure 2 show how the model learns to skip pixels that are not discriminative, such as the padding regions in the top and bottom of images, and the attended samples vary depending on the particular input being given to the network.

4 Conclusion

We presented Skip RNNs as an extension to existing recurrent architectures enabling them to skip state updates thereby reducing the number of sequential operations in the computation graph. Unlike other approaches, all parameters in Skip RNN are trained with backpropagation. Experiments conducted with LSTMs and GRUs showed that Skip RNNs can match or in some cases even outperform the baseline models while relaxing their computational requirements. Skip RNNs provide faster and more stable training for long sequences and complex models, likely due to gradients being backpropagated through fewer time steps resulting in a simpler optimization task. Moreover, the introduced computational savings are better suited for modern hardware than those methods that reduce the amount of computation required at each time step [12, 18, 4].

Acknowledgments

This work was partially supported by the Spanish Ministry of Economy and Competitivity under contracts TIN2012-34557 by the BSC-CNS Severo Ochoa program (SEV-2011-00067), and contracts TEC2013-43935-R and TEC2016-75976-R. It has also been supported by grants 2014-SGR-1051 and 2014-SGR-1421 by the Government of Catalonia, and the European Regional Development Fund (ERDF). We would also like to thank the technical support team at the Barcelona Supercomputing Center.

References

[1] A. Almahairi, N. Ballas, T. Cooijmans, Y. Zheng, H. Larochelle, and A. Courville. Dynamic capacity networks. In *ICML*, 2016.

[2] J. L. Ba, J. R. Kiros, and G. E. Hinton. Layer normalization. *arXiv preprint arXiv:1607.06450*, 2016.

[3] Y. Bengio, N. Léonard, and A. Courville. Estimating or propagating gradients through stochastic neurons for conditional computation. *arXiv preprint arXiv:1308.3432*, 2013.

[4] J. Chung, S. Ahn, and Y. Bengio. Hierarchical multiscale recurrent neural networks. In *ICLR*, 2017.

[5] T. Cooijmans, N. Ballas, C. Laurent, Ç. Gülçehre, and A. Courville. Recurrent batch normalization. In *ICLR*, 2017.

[6] A. Graves. Adaptive computation time for recurrent neural networks. *arXiv preprint arXiv:1603.08983*, 2016.

[7] A. Graves, G. Wayne, and I. Danihelka. Neural turing machines. *arXiv preprint arXiv:1410.5401*, 2014.

[8] E. Grefenstette, K. M. Hermann, M. Suleyman, and P. Blunsom. Learning to transduce with unbounded memory. In *NIPS*, 2015.

[9] G. Hinton. Neural networks for machine learning. Coursera video lectures, 2012.

[10] Y. Jernite, E. Grave, A. Joulin, and T. Mikolov. Variable computation in recurrent neural networks. In *ICLR*, 2017.

[11] D. Kingma and J. Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

[12] J. Koutnik, K. Greff, F. Gomez, and J. Schmidhuber. A clockwork rnn. In *ICML*, 2014.

[13] D. Krueger, T. Maharaj, J. Kramár, M. Pezeshki, N. Ballas, N. R. Ke, A. Goyal, Y. Bengio, H. Larochelle, A. Courville, et al. Zoneout: Regularizing rnns by randomly preserving hidden activations. In *ICLR*, 2017.

[14] Q. V. Le, N. Jaitly, and G. E. Hinton. A simple way to initialize recurrent networks of rectified linear units. *arXiv preprint arXiv:1504.00941*, 2015.

[15] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 1998.

[16] L. Liu and J. Deng. Dynamic deep neural networks: Optimizing accuracy-efficiency trade-offs by selective execution. *arXiv preprint arXiv:1701.00299*, 2017.

[17] M. McGill and P. Perona. Deciding how to decide: Dynamic routing in artificial neural networks. In *ICML*, 2017.

[18] D. Neil, M. Pfeiffer, and S. Liu. Phased LSTM: accelerating recurrent network training for long or event-based sequences. In *NIPS*, 2016.

[19] R. Pascanu, T. Mikolov, and Y. Bengio. On the difficulty of training recurrent neural networks. In *ICML*, 2013.

[20] N. Shazeer, A. Mirhoseini, K. Maziarz, A. Davis, Q. Le, G. Hinton, and J. Dean. Outrageously large neural networks: The sparsely-gated mixture-of-experts layer. In *ICLR*, 2017.

[21] J. Weston, S. Chopra, and A. Bordes. Memory networks. *arXiv preprint arXiv:1410.3916*, 2014.

[22] R. J. Williams. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 1992.

[23] K. Xu, J. Ba, R. Kiros, K. Cho, A. Courville, R. Salakhudinov, R. Zemel, and Y. Bengio. Show, attend and tell: Neural image caption generation with visual attention. In *ICML*, 2015.

[24] A. W. Yu, H. Lee, and Q. V. Le. Learning to skim text. In *ACL*, 2017.

[25] W. Zaremba, I. Sutskever, and O. Vinyals. Recurrent neural network regularization. In *ICLR*, 2015.

A Experimental Setup

Training is performed with Adam [11], learning rate of 10^{-4}, $\beta_1 = 0.9$, $\beta_2 = 0.999$ and $\epsilon = 10^{-8}$ on batches of 256 examples. Gradient clipping [19] with a threshold of 1 is applied to all trainable variables. Bias b_p in Equation 4 is initialized to 1, so that all samples are used at the beginning of training. In practice, forcing the network to use all samples at the beginning of training improves its robustness against random initializations of its weights and increases the reproducibility of the presented experiments. A similar behavior was observed in other augmented RNN architectures such as Neural Stacks [8]. The initial hidden state s_0 is learned during training, whereas $\tilde{u}_0$ is set to a constant value of 1 in order to force the first update at $t = 1$.

Experiments are implemented with TensorFlow[2] and run on a single NVIDIA K80 GPU.

[2] `https://www.tensorflow.org`

SKIP RNN: LEARNING TO SKIP STATE UPDATES IN RECURRENT NEURAL NETWORKS

Víctor Campos[*†], **Brendan Jou**[‡], **Xavier Giró-i-Nieto**[§], **Jordi Torres**[†], **Shih-Fu Chang**[Γ]
[†]Barcelona Supercomputing Center, [‡]Google Inc,
[§]Universitat Politècnica de Catalunya, [Γ]Columbia University
{victor.campos, jordi.torres}@bsc.es, bjou@google.com,
xavier.giro@upc.edu, shih.fu.chang@columbia.edu

ABSTRACT

Recurrent Neural Networks (RNNs) continue to show outstanding performance in sequence modeling tasks. However, training RNNs on long sequences often face challenges like slow inference, vanishing gradients and difficulty in capturing long term dependencies. In backpropagation through time settings, these issues are tightly coupled with the large, sequential computational graph resulting from unfolding the RNN in time. We introduce the Skip RNN model which extends existing RNN models by learning to skip state updates and shortens the effective size of the computational graph. This model can also be encouraged to perform fewer state updates through a budget constraint. We evaluate the proposed model on various tasks and show how it can reduce the number of required RNN updates while preserving, and sometimes even improving, the performance of the baseline RNN models. Source code is publicly available at https://imatge-upc.github.io/skiprnn-2017-telecombcn/.

1 INTRODUCTION

Recurrent Neural Networks (RNNs) have become the standard approach for practitioners when addressing machine learning tasks involving sequential data. Such success has been enabled by the appearance of larger datasets, more powerful computing resources and improved architectures and training algorithms. Gated units, such as the Long Short-Term Memory (Hochreiter & Schmidhuber, 1997) (LSTM) and the Gated Recurrent Unit (Cho et al., 2014) (GRU), were designed to deal with the vanishing gradients problem commonly found in RNNs (Bengio et al., 1994). These architectures have been popularized, in part, due to their impressive results on a variety of tasks in machine translation (Bahdanau et al., 2015), language modeling (Zaremba et al., 2015) and speech recognition (Graves et al., 2013).

Some of the main challenges of RNNs are in their training and deployment when dealing with long sequences, due to their inherently sequential behaviour. These challenges include throughput degradation, slower convergence during training and memory leakage, even for gated architectures (Neil et al., 2016). Sequence shortening techniques, which can be seen as a sort of conditional computation (Bengio et al., 2013; Bengio, 2013; Davis & Arel, 2013) in time, can alleviate these issues. The most common approaches, such as cropping discrete signals or reducing the sampling rate in continuous signals, are based on heuristics and can be suboptimal. In contrast, we propose a model that is able to learn which samples (i.e., elements in the input sequence) need to be used in order to solve the target task. Consider a video understanding task as an example: scenes with large motion may benefit from high frame rates, whereas only a few frames are needed to capture the semantics of a mostly static scene.

The main contribution of this work is a novel modification for existing RNN architectures that allows them to skip state updates, decreasing the number of sequential operations performed, without requiring any additional supervision signal. This model, called Skip RNN, adaptively determines whether the state needs to be updated or copied to the next time step. We show how the network can

[*]Work done while Víctor Campos was a visiting scholar at Columbia University.

be encouraged to perform fewer state updates by adding a penalization term during training, allowing us to train models under different computation budgets. The proposed modification can generally be integrated with any RNN and we show, in this paper, implementations with well-known RNNs, namely LSTM and GRU. The resulting models show promising results on a series of sequence modeling tasks. In particular, we evaluate the proposed Skip RNN architecture on six sequence learning problems: an adding task, sine wave frequency discrimination, digit classification, sentiment analysis in movie reviews, action classification in video, and temporal action localization in video[1].

2 Related work

Conditional computation has been shown to gradually increase model capacity without proportional increases in computational cost by exploiting certain computation paths for each input (Bengio et al., 2013; Liu & Deng, 2017; Almahairi et al., 2016; McGill & Perona, 2017; Shazeer et al., 2017). This idea has been extended in the temporal domain, such as in learning how many times an input needs to be "pondered" before moving to the next one (Graves, 2016) or designing RNN architectures whose number of layers depend on the input data (Chung et al., 2017). Other works have addressed time-dependent computation in RNNs by updating only a fraction of the hidden states based on the current hidden state and input (Jernite et al., 2017), or following periodic patterns (Koutnik et al., 2014; Neil et al., 2016). However, due to the inherently sequential nature of RNNs and the parallel computation capabilities of modern hardware, reducing the size of the matrices involved in the computations performed at each time step generally has not accelerated inference as dramatically as hoped. The proposed Skip RNN model can be seen as form of conditional computation in time, where the computation associated to the RNN updates may or may not be executed at every time step. This idea is related to the UPDATE and COPY operations in hierarchical multiscale RNNs (Chung et al., 2017), but applied to the whole stack of RNN layers at the same time. This difference is key to allowing our approach to skip input samples, effectively reducing sequential computation and shielding the hidden state over longer time lags. Learning whether to update or copy the hidden state through time steps can be seen as a learnable Zoneout mask (Krueger et al., 2017) which is shared between all the units in the hidden state. Similarly, it can be interpreted as an input-dependent recurrent version of stochastic depth (Huang et al., 2016).

Selecting parts of the input signal is similar in spirit to the hard attention mechanisms that have been applied to image regions (Mnih et al., 2014), where only some patches of the input image are attended in order to generate captions (Xu et al., 2015) or detect objects (Ba et al., 2014). Our model can be understood as generating a hard temporal attention mask on-the-fly given previously seen samples, deciding which time steps should be attended and operating on a subset of input samples. Subsampling input sequences has been explored for visual storylines generation (Sigurdsson et al., 2016b), although jointly optimizing the RNN weights and the subsampling mechanism is often computationally infeasible and they use an Expectation-Maximization algorithm instead. Similar research has been conducted for video analysis tasks, discovering minimally needed evidence for event recognition (Bhattacharya et al., 2014) and training agents that decide which frames need to be observed in order to localize actions in time (Yeung et al., 2016; Su & Grauman, 2016). Motivated by the advantages of training recurrent models on shorter subsequences, efforts have been conducted on learning differentiable subsampling mechanisms (Raffel & Lawson, 2017), although the computational complexity of the proposed method precludes its application to long input sequences. In contrast, our proposed method can be trained with backpropagation and does not degrade the complexity of the baseline RNNs.

Accelerating inference in RNNs is difficult due to their inherently sequential nature, leading to the design of Quasi-Recurrent Neural Networks (Bradbury et al., 2017) and Simple Recurrent Units (Lei & Zhang, 2017), which relax the temporal dependency between consecutive steps. With the goal of speeding up RNN inference, LSTM-Jump (Yu et al., 2017) augments an LSTM cell with a classification layer that will decide how many steps to jump between RNN updates. Despite its promising results on text tasks, the model needs to be trained with REINFORCE (Williams, 1992), which requires defining a reasonable reward signal. Determining these rewards are non-trivial and may not necessarily generalize across tasks. Moreover, the number of tokens read between jumps,

[1]Experiments on sine wave frequency discrimination, sentiment analysis in movie reviews and action classification in video are reported in the appendix due to space limitations.

the maximum jump distance and the number of jumps allowed all need to be chosen in advance. These hyperparameters define a reduced set of subsequences that the model can sample, instead of allowing the network to learn any arbitrary sampling scheme. Unlike LSTM-Jump, our proposed approach is differentiable, thus not requiring any modifications to the loss function and simplifying the optimization process, and is not limited to a predefined set of sample selection patterns.

3 MODEL DESCRIPTION

An RNN takes an input sequence $\mathbf{x} = (x_1, \dots, x_T)$ and generates a state sequence $\mathbf{s} = (s_1, \dots, s_T)$ by iteratively applying a parametric state transition model S from $t = 1$ to T:

$$s_t = S(s_{t-1}, x_t) \tag{1}$$

We augment the network with a binary *state update gate*, $u_t \in \{0, 1\}$, selecting whether the state of the RNN will be updated ($u_t = 1$) or copied from the previous time step ($u_t = 0$). At every time step t, the probability $\tilde{u}_{t+1} \in [0, 1]$ of performing a state update at $t + 1$ is emitted. The resulting architecture is depicted in Figure 1 and can be characterized as follows:

$$u_t = f_{binarize}(\tilde{u}_t) \tag{2}$$
$$s_t = u_t \cdot S(s_{t-1}, x_t) + (1 - u_t) \cdot s_{t-1} \tag{3}$$
$$\Delta\tilde{u}_t = \sigma(W_p s_t + b_p) \tag{4}$$
$$\tilde{u}_{t+1} = u_t \cdot \Delta\tilde{u}_t + (1 - u_t) \cdot (\tilde{u}_t + \min(\Delta\tilde{u}_t, 1 - \tilde{u}_t)) \tag{5}$$

where W_p is a weights vector, b_p is a scalar bias, σ is the sigmoid function and $f_{binarize} : [0, 1] \rightarrow \{0, 1\}$ binarizes the input value. Should the network be composed of several layers, some columns of W_p can be fixed to 0 so that $\Delta\tilde{u}_t$ depends only on the states of a subset of layers (see Section 4.3 for an example with two layers). We implement $f_{binarize}$ as a deterministic step function $u_t = \text{round}(\tilde{u}_t)$, although a stochastic sampling from a Bernoulli distribution $u_t \sim \text{Bernoulli}(\tilde{u}_t)$ would be possible as well.

The model formulation encodes the observation that the likelihood of requesting a new input to update the state increases with the number of consecutively skipped samples. Whenever a state update is omitted, the pre-activation of the state update gate for the following time step, $\tilde{u}_{t+1}$, is incremented by $\Delta\tilde{u}_t$. On the other hand, if a state update is performed, the accumulated value is flushed and $\tilde{u}_{t+1} = \Delta\tilde{u}_t$.

The number of skipped time steps can be computed ahead of time. For the particular formulation used in this work, where $f_{binarize}$ is implemented by means of a rounding function, the number of skipped samples after performing a state update at time step t is given by:

$$N_{skip}(t) = \min\{n : n \cdot \Delta\tilde{u}_t \geq 0.5\} - 1 \tag{6}$$

where $n \in \mathbb{Z}^+$. This enables more efficient implementations where no computation at all is performed whenever $u_t = 0$. These computational savings are possible because $\Delta\tilde{u}_t = \sigma(W_p s_t + b_p) = \sigma(W_p s_{t-1} + b_p) = \Delta\tilde{u}_{t-1}$ when $u_t = 0$ and there is no need to evaluate it again, as depicted in Figure 1d.

There are several advantages in reducing the number of RNN updates. From the computational standpoint, fewer updates translates into fewer required sequential operations to process an input signal, leading to faster inference and reduced energy consumption. Unlike some other models that aim to reduce the average number of operations per step (Neil et al., 2016; Jernite et al., 2017), ours enables skipping steps completely. Replacing RNN updates with copy operations increases the memory of the network and its ability to model long term dependencies even for gated units, since the exponential memory decay observed in LSTM and GRU (Neil et al., 2016) is alleviated. During training, gradients are propagated through fewer updating time steps, providing faster convergence in some tasks involving long sequences. Moreover, the proposed model is orthogonal to recent advances in RNNs and could be used in conjunction with such techniques, e.g. normalization (Cooijmans et al., 2017; Ba et al., 2016), regularization (Zaremba et al., 2015; Krueger et al., 2017), variable computation (Jernite et al., 2017; Neil et al., 2016) or even external memory (Graves et al., 2014; Weston et al., 2014).

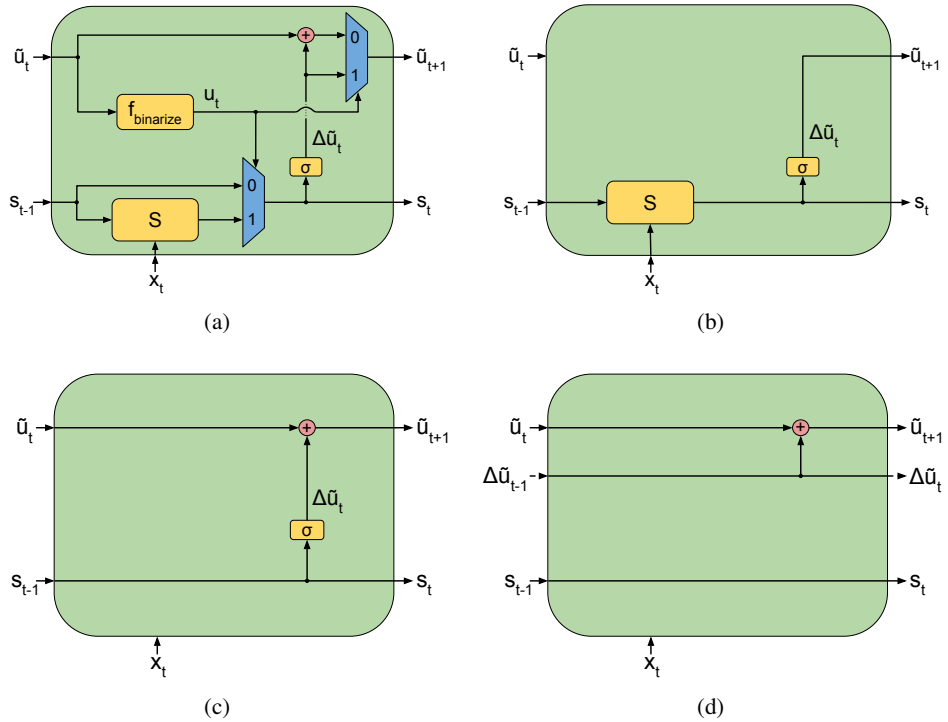

Figure 1: Model architecture of the proposed Skip RNN. **(a)** Complete Skip RNN architecture, where the computation graph at time step t is conditioned on u_t. **(b)** Architecture when the state is updated, i.e. $u_t = 1$. **(c)** Architecture when the update step is skipped and the previous state is copied, i.e. $u_t = 0$. **(d)** In practice, redundant computation is avoided by propagating $\Delta\tilde{u}_t$ between time steps when $u_t = 0$.

3.1 Error gradients

The whole model is differentiable except for $f_{binarize}$, which outputs binary values. A common method for optimizing functions involving discrete variables is REINFORCE (Williams, 1992), although several estimators have been proposed for the particular case of neurons with binary outputs (Bengio et al., 2013). We select the straight-through estimator (Hinton, 2012; Bengio et al., 2013), which consists of approximating the step function by the identity when computing gradients during the backward pass:

$$\frac{\partial f_{binarize}\left(x\right)}{\partial x} = 1 \tag{7}$$

This yields a biased estimator that has proven more efficient than other unbiased but high-variance estimators such as REINFORCE (Bengio et al., 2013) and has been successfully applied in different works (Courbariaux et al., 2016; Chung et al., 2017). By using the straight-through estimator as the backward pass for $f_{binarize}$, all the model parameters can be trained to minimize the target loss function with standard backpropagation and without defining any additional supervision or reward signal.

3.2 Limiting computation

The Skip RNN is able to learn when to update or copy the state without explicit information about which samples are useful to solve the task at hand. However, a different operating point on the trade-off between performance and number of processed samples may be required depending on the application, e.g. one may be willing to sacrifice a few accuracy points in order to run faster on machines with a low computational power, or to reduce energy impact on portable devices. The

proposed model can be encouraged to perform fewer state updates through additional loss terms, a common practice in neural networks with dynamically allocated computation (Liu & Deng, 2017; McGill & Perona, 2017; Graves, 2016; Jernite et al., 2017). In particular, we consider a *cost per sample* condition

$$L_{budget} = \lambda \cdot \sum_{t=1}^{T} u_t, \tag{8}$$

where L_{budget} is the cost associated to a single sequence, λ is the cost per sample and T is the sequence length. This formulation bears a similarity to weight decay regularization, where the network is encouraged to slowly converge toward a solution where the norm of the weights is small. Similarly, in this case the network is encouraged to converge toward a solution where fewer state updates are required.

Although the above budget formulation is extensively studied in our experiments, other budget loss terms can be used depending on the application. For instance, a specific number of samples may be encouraged by applying a L_1 or L_2 loss between the target value and the number of updates per sequence, $\sum_{t=1}^{T} u_t$.

4 EXPERIMENTS

In the following section, we investigate the advantages of adding this state skipping to two common RNN architectures, LSTM and GRU, for a variety of tasks. In addition to the evaluation metric for each task, we report the number of RNN state updates (i.e., the number of elements in the input sequence used by the model) and the number of floating point operations (FLOPs) as measures of the computational load for each model. Since skipping an RNN update results in ignoring its corresponding input, we will refer to the number of updates and the number of used samples (i.e. elements in a sequence) interchangeably. With the goal of studying the effect of skipping state updates on the learning capability of the networks, we also introduce a baseline which skips a state update with probability p_{skip}. We tune the skipping probability to obtain models that perform a similar number of state updates to the Skip RNN models.

Training is performed with Adam (Kingma & Ba, 2014), learning rate of 10^{-4}, $\beta_1 = 0.9$, $\beta_2 = 0.999$ and $\epsilon = 10^{-8}$ on batches of 256. Gradient clipping (Pascanu et al., 2013) with a threshold of 1 is applied to all trainable variables. Bias b_p in Equation 4 is initialized to 1, so that all samples are used at the beginning of training[2]. The initial hidden state s_0 is learned during training, whereas $\tilde{u}_0$ is set to a constant value of 1 in order to force the first update at $t = 1$.

Experiments are implemented with TensorFlow[3] and run on a single NVIDIA K80 GPU.

4.1 ADDING TASK

We revisit one of the original LSTM tasks (Hochreiter & Schmidhuber, 1997), where the network is given a sequence of *(value, marker)* tuples. The desired output is the addition of only two values that are marked with a 1, whereas those marked with a 0 need to be ignored. We follow the experimental setup in Neil et al. (2016), where the first marker is randomly placed among the first 10% of samples (drawn with uniform probability) and the second one is placed among the last half of samples (drawn with uniform probability). This marker distribution yields sequences where at least 40% of the samples are distractors and provide no useful information at all. However, it is worth noting that in this task the risk of missing a marker is very large as compared to the benefits of working on shorter subsequences.

[2]In practice, forcing the network to use all samples at the beginning of training improves its robustness against random initializations of its weights and increases the reproducibility of the presented experiments. A similar behavior was observed in other augmented RNN architectures such as Neural Stacks (Grefenstette et al., 2015).

[3]`https://www.tensorflow.org`

Model	Task solved	State updates	Inference FLOPs
LSTM	Yes	$100.0\% \pm 0.0\%$	2.46×10^6
LSTM ($p_{skip} = 0.2$)	No	$80.0\% \pm 0.1\%$	1.97×10^6
LSTM ($p_{skip} = 0.5$)	No	$50.1\% \pm 0.1\%$	1.23×10^6
Skip LSTM, $\lambda = 0$	Yes	$81.1\% \pm 3.6\%$	2.00×10^6
Skip LSTM, $\lambda = 10^{-5}$	Yes	$53.9\% \pm 2.1\%$	1.33×10^6
GRU	Yes	$100.0\% \pm 0.0\%$	1.85×10^6
GRU ($p_{skip} = 0.02$)	No	$98.0\% \pm 0.0\%$	1.81×10^6
GRU ($p_{skip} = 0.5$)	No	$49.9\% \pm 0.6\%$	9.25×10^5
Skip GRU, $\lambda = 0$	Yes	$97.9\% \pm 3.2\%$	1.81×10^6
Skip GRU, $\lambda = 10^{-5}$	Yes	$50.7\% \pm 2.6\%$	9.40×10^5

Table 1: Results for the adding task, displayed as $mean \pm std$ over four different runs. The task is considered to be solved if the MSE is at least two orders of magnitude below the variance of the output distribution.

We train RNN models with 110 units each on sequences of length 50, where the values are uniformly drawn from $\mathcal{U}(-0.5, 0.5)$. The final RNN state is fed to a fully connected layer that regresses the scalar output. The model is trained to minimize the Mean Squared Error (MSE) between the output and the ground truth. We consider that a model is able to solve the task when its MSE on a held-out set of examples is at least two orders of magnitude below the variance of the output distribution. This criterion is a stricter version of the one followed by Hochreiter & Schmidhuber (1997).

While all models learn to solve the task, results in Table 1 show that Skip RNN models are able to do so with roughly half of the updates of their corresponding counterparts. We observed that the models using fewer updates never miss any marker, since the penalization in terms of MSE would be very large (see Section B.1 for examples). This is confirmed by the poor performance of the baselines that randomly skip state updates, which are not able to solve the tasks even when the skipping probability is low. Skip RNN models learn to skip most of the samples in the 40% of the sequence where there are no markers. Moreover, most updates are skipped once the second marker is found, since all the relevant information in the sequence has already been seen. This last pattern provides evidence that the proposed models effectively learn whether to update or copy the hidden state based on the input sequence, as opposed to learning biases in the dataset only. As a downside, Skip RNN models show some difficulties skipping a large number of updates at once, probably due to the cumulative nature of $\tilde{u}_t$.

4.2 MNIST Classification from a Sequence of Pixels

The MNIST handwritten digits classification benchmark (LeCun et al., 1998) is traditionally addressed with Convolutional Neural Networks (CNNs) that efficiently exploit spatial dependencies through weight sharing. By flattening the 28×28 images into 784-d vectors, however, it can be reformulated as a challenging task for RNNs where long term dependencies need to be leveraged (Le et al., 2015b). We follow the standard data split and set aside 5,000 training samples for validation purposes. After processing all pixels with an RNN with 110 units, the last hidden state is fed into a linear classifier predicting the digit class. All models are trained for 600 epochs to minimize cross-entropy loss.

Table 2 summarizes classification results on the test set after 600 epochs of training. Skip RNNs are not only able to solve the task using fewer updates than their counterparts, but also show a lower variation among runs and train faster (see Figure 2). We hypothesize that skipping updates make the Skip RNNs work on shorter subsequences, simplifying the optimization process and allowing the networks to capture long term dependencies more easily. A similar behavior was observed for Phased LSTM, where increasing the sparsity of cell updates accelerates training for very long sequences (Neil et al., 2016). However, the drop in performance observed in the models where the state updates are skipped randomly suggests that learning which samples to use is a key component in the performance of Skip RNN.

Model	Accuracy	State updates	Inference FLOPs
LSTM	0.910 ± 0.045	784.00 ± 0.00	3.83×10^7
LSTM ($p_{skip} = 0.5$)	0.893 ± 0.003	392.03 ± 0.05	1.91×10^7
Skip LSTM, $\lambda = 10^{-4}$	0.973 ± 0.002	379.38 ± 33.09	1.86×10^7
GRU	0.968 ± 0.013	784.00 ± 0.00	2.87×10^7
GRU ($p_{skip} = 0.5$)	0.912 ± 0.004	391.86 ± 0.14	1.44×10^7
Skip GRU, $\lambda = 10^{-4}$	0.976 ± 0.003	392.62 ± 26.48	1.44×10^7
TANH-RNN (Le et al., 2015a)	0.350	784.00	-
*i*RNN (Le et al., 2015a)	0.970	784.00	-
*u*RNN (Arjovsky et al., 2016)	0.951	784.00	-
*s*TANH-RNN (Zhang et al., 2016)	0.981	784.00	-
LSTM (Cooijmans et al., 2017)	0.989	784.00	-
BN-LSTM (Cooijmans et al., 2017)	0.990	784.00	-

Table 2: Accuracy, used samples and average FLOPs per sequence at inference on the test set of MNIST after 600 epochs of training. Results are displayed as $mean \pm std$ over four different runs.

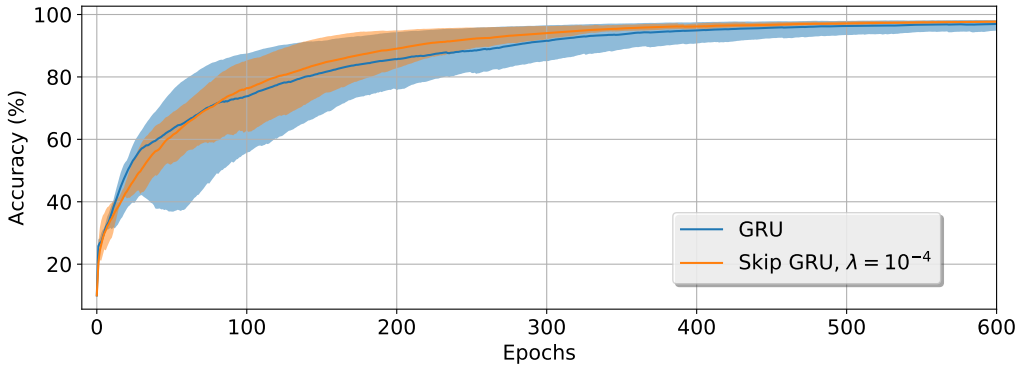

Figure 2: Accuracy evolution during training on the validation set of MNIST. The Skip GRU exhibits lower variance and faster convergence than the baseline GRU. A similar behavior is observed for LSTM and Skip LSTM, but omitted for clarity. Shading shows maximum and minimum over 4 runs, while dark lines indicate the mean.

The performance of RNN models on this task can be boosted through techniques like recurrent batch normalization (Cooijmans et al., 2017) or recurrent skip coefficients (Zhang et al., 2016). Cooijmans et al. (2017) show how an LSTM with specific weight initialization schemes for improved gradient flow (Le et al., 2015a; Arjovsky et al., 2016) can reach accuracy rates of up to 0.989%. Note that these techniques are orthogonal to skipping state updates and Skip RNN models could benefit from them as well.

Sequences of pixels can be reshaped back into 2D images, allowing to visualize the samples used by the RNNs as a sort of hard visual attention model (Xu et al., 2015). Examples such as the ones depicted in Figure 3 show how the model learns to skip pixels that are not discriminative, such as the padding regions in the top and bottom of images. Similarly to the qualitative results for the adding task (Section 4.1), attended samples vary depending on the particular input being given to the network.

4.3 Temporal action localization on Charades

One popular approach to video analysis tasks today is to extract frame-level features with a CNN and modeling temporal dynamics with an RNN (Donahue et al., 2015; Yue-Hei Ng et al., 2015). Videos are commonly recorded at high sampling rates, generating long sequences with strong tem-

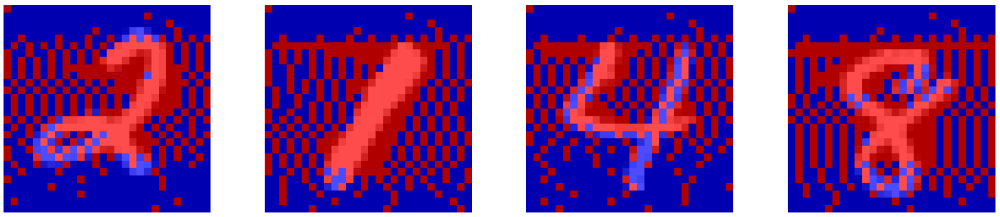

Figure 3: Sample usage examples for the Skip LSTM with $\lambda = 10^{-4}$ on the test set of MNIST. Red pixels are used, whereas blue ones are skipped.

poral redundancies that are challenging for RNNs. Moreover, processing frames with a CNN is computationally expensive and may become prohibitive for high frame rates. These issues have been alleviated in previous works by using short clips (Donahue et al., 2015) or by downsampling the original data in order to cover long temporal spans without increasing the sequence length excessively (Yue-Hei Ng et al., 2015). Instead of addressing the long sequence problem at the input data level, we let the network learn which frames need to be used.

Charades (Sigurdsson et al., 2016a) is a dataset containing 9,848 videos annotated for 157 action classes in a per-frame fashion. Frames are encoded using *fc7* features from the RGB stream of a Two-Stream CNN provided by the organizers of the challenge[4], extracted at 6 fps. The encoded frames are fed into two stacked RNN layers with 256 units each and the hidden state in the last RNN layer is used to compute the update probability for the Skip RNN models. Since each frame may be annotated with zero or more classes, the networks are trained to minimize element-wise binary cross-entropy at every time step. Unlike the previous sequence tagging tasks, this setup allows us to evaluate the performance of Skip RNN on a task where the output is a sequence as well.

Evaluation is performed following the setup by Sigurdsson et al. (2016c), but evaluating on 100 equally spaced frames instead of 25, and results are reported in Table 3. It is surprising that the GRU baselines that randomly skip state updates perform on par with their Skip GRU counterparts for low skipping probabilities. We hypothesize several reasons for this behavior, which was not observed in previous experiments: (1) there is a supervision signal at every time step and the inputs and (2) outputs are strongly correlated in consecutive frames. On the other hand, Skip RNN models clearly outperform the random methods when fewer updates are allowed. Note that this setup is far more challenging because of the longer time spans between updates, so properly distributing the state updates along the sequence is key to the performance of the models. Interestingly, Skip RNN models learn which frames need to be attended from RGB data and without having access to explicit motion information.

Skip GRU tends to perform fewer state updates than Skip LSTM when the cost per sample is low or none. This behavior is the opposite of the one observed in the adding task (Section 4.1), which may be related to the observation that determining the best performing gated unit depends on the task at hand Chung et al. (2014). Indeed, GRU models consistently outperform LSTM ones on this task. This mismatch in the number of used samples is not observed for large values of λ, as both Skip LSTM and Skip GRU converge to a comparable number of used samples.

A previous work reports better action localization performance by integrating RGB and optical flow information as an input to an LSTM, reaching 9.60% mAP (Sigurdsson et al., 2016c). This boost in performance comes at the cost of roughly doubling the number of FLOPs and memory footprint of the CNN encoder, plus requiring the extraction of flow information during a preprocessing step. Interestingly, our model learns which frames need to be attended from RGB data and without having access to explicit motion information.

[4] `http://vuchallenge.org/charades.html`

Model	mAP (%)	State updates	Inference FLOPs
LSTM	8.40	172.9 ± 47.4	2.65×10^{12}
LSTM ($p_{skip} = 0.75$)	8.11	43.3 ± 13.2	6.63×10^{11}
LSTM ($p_{skip} = 0.90$)	7.21	17.2 ± 6.1	2.65×10^{11}
Skip LSTM, $\lambda = 0$	8.32	172.9 ± 47.4	2.65×10^{12}
Skip LSTM, $\lambda = 10^{-4}$	8.61	172.9 ± 47.4	2.65×10^{12}
Skip LSTM, $\lambda = 10^{-3}$	8.32	41.9 ± 11.3	6.41×10^{11}
Skip LSTM, $\lambda = 10^{-2}$	7.86	17.4 ± 4.4	2.66×10^{11}
GRU	8.70	172.9 ± 47.4	2.65×10^{12}
GRU ($p_{skip} = 0.10$)	8.94	155.6 ± 42.9	2.39×10^{12}
GRU ($p_{skip} = 0.40$)	8.81	103.6 ± 29.3	1.06×10^{12}
GRU ($p_{skip} = 0.70$)	8.42	51.9 ± 15.4	7.95×10^{11}
GRU ($p_{skip} = 0.90$)	7.09	17.3 ± 6.3	2.65×10^{11}
Skip GRU, $\lambda = 0$	8.94	159.9 ± 46.9	2.45×10^{12}
Skip GRU, $\lambda = 10^{-4}$	8.76	100.8 ± 28.1	1.54×10^{12}
Skip GRU, $\lambda = 10^{-3}$	8.68	54.2 ± 16.2	8.29×10^{11}
Skip GRU, $\lambda = 10^{-2}$	7.95	18.4 ± 5.1	2.82×10^{11}

Table 3: Mean Average Precision (mAP), used samples and average FLOPs per sequence at inference on the validation set of Charades. The number of state updates is displayed as $mean \pm std$ over all the videos in the validation set.

5 Conclusion

We presented Skip RNNs as an extension to existing recurrent architectures enabling them to skip state updates thereby reducing the number of sequential operations in the computation graph. Unlike other approaches, all parameters in Skip RNN are trained with backpropagation. Experiments conducted with LSTMs and GRUs showed that Skip RNNs can match or in some cases even outperform the baseline models while relaxing their computational requirements. Skip RNNs provide faster and more stable training for long sequences and complex models, owing to gradients being backpropagated through fewer time steps resulting in a simpler optimization task. Moreover, the introduced computational savings are better suited for modern hardware than those methods that reduce the amount of computation required at each time step (Koutnik et al., 2014; Neil et al., 2016; Chung et al., 2017).

Acknowledgments

This work was partially supported by the Spanish Ministry of Economy and Competitivity and the European Regional Development Fund (ERDF) under contracts TEC2016-75976-R and TIN2015-65316-P, by the BSC-CNS Severo Ochoa program SEV-2015-0493, and grant 2014-SGR-1051 by the Catalan Government. Víctor Campos was supported by Obra Social "la Caixa" through La Caixa-Severo Ochoa International Doctoral Fellowship program. We would also like to thank the technical support team at the Barcelona Supercomputing Center.

References

Amjad Almahairi, Nicolas Ballas, Tim Cooijmans, Yin Zheng, Hugo Larochelle, and Aaron Courville. Dynamic capacity networks. In *ICML*, 2016.

Martin Arjovsky, Amar Shah, and Yoshua Bengio. Unitary evolution recurrent neural networks. In *ICML*, 2016.

Jimmy Ba, Volodymyr Mnih, and Koray Kavukcuoglu. Multiple object recognition with visual attention. *arXiv preprint arXiv:1412.7755*, 2014.

Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. Layer normalization. *arXiv preprint arXiv:1607.06450*, 2016.

Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. Neural machine translation by jointly learning to align and translate. In *ICLR*, 2015.

Yoshua Bengio. Deep learning of representations: Looking forward. In *SLSP*, 2013.

Yoshua Bengio, Patrice Simard, and Paolo Frasconi. Learning long-term dependencies with gradient descent is difficult. *IEEE Transactions on Neural Networks*, 1994.

Yoshua Bengio, Nicholas Léonard, and Aaron Courville. Estimating or propagating gradients through stochastic neurons for conditional computation. *arXiv preprint arXiv:1308.3432*, 2013.

Subhabrata Bhattacharya, Felix X Yu, and Shih-Fu Chang. Minimally needed evidence for complex event recognition in unconstrained videos. In *ICMR*, 2014.

James Bradbury, Stephen Merity, Caiming Xiong, and Richard Socher. Quasi-recurrent neural networks. In *ICLR*, 2017.

Joao Carreira and Andrew Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. In *CVPR*, 2017.

Kyunghyun Cho, Bart Van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. Learning phrase representations using rnn encoder-decoder for statistical machine translation. In *EMNLP*, 2014.

Junyoung Chung, Caglar Gulcehre, KyungHyun Cho, and Yoshua Bengio. Empirical evaluation of gated recurrent neural networks on sequence modeling. *arXiv preprint arXiv:1412.3555*, 2014.

Junyoung Chung, Sungjin Ahn, and Yoshua Bengio. Hierarchical multiscale recurrent neural networks. In *ICLR*, 2017.

Tim Cooijmans, Nicolas Ballas, César Laurent, Çağlar Gülçehre, and Aaron Courville. Recurrent batch normalization. In *ICLR*, 2017.

Matthieu Courbariaux, Itay Hubara, Daniel Soudry, Ran El-Yaniv, and Yoshua Bengio. Binarized neural networks: Training deep neural networks with weights and activations constrained to+ 1 or-1. *arXiv preprint arXiv:1602.02830*, 2016.

Andrew Davis and Itamar Arel. Low-rank approximations for conditional feedforward computation in deep neural networks. *arXiv preprint arXiv:1312.4461*, 2013.

Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. ImageNet: A large-scale hierarchical image database. In *CVPR*, 2009.

Jeffrey Donahue, Lisa Anne Hendricks, Sergio Guadarrama, Marcus Rohrbach, Subhashini Venugopalan, Kate Saenko, and Trevor Darrell. Long-term recurrent convolutional networks for visual recognition and description. In *CVPR*, 2015.

Alex Graves. Adaptive computation time for recurrent neural networks. *arXiv preprint arXiv:1603.08983*, 2016.

Alex Graves, Abdel-rahman Mohamed, and Geoffrey Hinton. Speech recognition with deep recurrent neural networks. In *ICASSP*, 2013.

Alex Graves, Greg Wayne, and Ivo Danihelka. Neural turing machines. *arXiv preprint arXiv:1410.5401*, 2014.

Edward Grefenstette, Karl Moritz Hermann, Mustafa Suleyman, and Phil Blunsom. Learning to transduce with unbounded memory. In *NIPS*, 2015.

Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016.

Geoffrey Hinton. Neural networks for machine learning. Coursera video lectures, 2012.

Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 1997.

Gao Huang, Yu Sun, Zhuang Liu, Daniel Sedra, and Kilian Q Weinberger. Deep networks with stochastic depth. In *ECCV*, 2016.

Yacine Jernite, Edouard Grave, Armand Joulin, and Tomas Mikolov. Variable computation in recurrent neural networks. In *ICLR*, 2017.

Yoon Kim. Convolutional neural networks for sentence classification. In *EMNLP*, 2014.

Diederik Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

Jan Koutnik, Klaus Greff, Faustino Gomez, and Juergen Schmidhuber. A clockwork rnn. In *ICML*, 2014.

David Krueger, Tegan Maharaj, János Kramár, Mohammad Pezeshki, Nicolas Ballas, Nan Rosemary Ke, Anirudh Goyal, Yoshua Bengio, Hugo Larochelle, Aaron Courville, et al. Zoneout: Regularizing rnns by randomly preserving hidden activations. In *ICLR*, 2017.

Quoc V Le, Navdeep Jaitly, and Geoffrey E Hinton. A simple way to initialize recurrent networks of rectified linear units. *arXiv preprint arXiv:1504.00941*, 2015a.

Quoc V Le, Navdeep Jaitly, and Geoffrey E Hinton. A simple way to initialize recurrent networks of rectified linear units. *arXiv preprint arXiv:1504.00941*, 2015b.

Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 1998.

Tao Lei and Yu Zhang. Training rnns as fast as cnns. *arXiv preprint arXiv:1709.02755*, 2017.

Lanlan Liu and Jia Deng. Dynamic deep neural networks: Optimizing accuracy-efficiency trade-offs by selective execution. *arXiv preprint arXiv:1701.00299*, 2017.

Shayne Longpre, Sabeek Pradhan, Caiming Xiong, and Richard Socher. A way out of the odyssey: Analyzing and combining recent insights for lstms. *arXiv preprint arXiv:1611.05104*, 2016.

Andrew L Maas, Raymond E Daly, Peter T Pham, Dan Huang, Andrew Y Ng, and Christopher Potts. Learning word vectors for sentiment analysis. In *ACL*, 2011.

Mason McGill and Pietro Perona. Deciding how to decide: Dynamic routing in artificial neural networks. In *ICML*, 2017.

Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. Distributed representations of words and phrases and their compositionality. In *NIPS*, 2013.

Takeru Miyato, Andrew M Dai, and Ian Goodfellow. Adversarial training methods for semi-supervised text classification. In *ICLR*, 2017.

Volodymyr Mnih, Nicolas Heess, Alex Graves, et al. Recurrent models of visual attention. In *NIPS*, 2014.

Daniel Neil, Michael Pfeiffer, and Shih-Chii Liu. Phased LSTM: accelerating recurrent network training for long or event-based sequences. In *NIPS*, 2016.

Razvan Pascanu, Tomas Mikolov, and Yoshua Bengio. On the difficulty of training recurrent neural networks. In *ICML*, 2013.

Colin Raffel and Dieterich Lawson. Training a subsampling mechanism in expectation. In *ICLR Workshop Track*, 2017.

Noam Shazeer, Azalia Mirhoseini, Krzysztof Maziarz, Andy Davis, Quoc Le, Geoffrey Hinton, and Jeff Dean. Outrageously large neural networks: The sparsely-gated mixture-of-experts layer. In *ICLR*, 2017.

Evan Shelhamer, Kate Rakelly, Judy Hoffman, and Trevor Darrell. Clockwork convnets for video semantic segmentation. *arXiv preprint arXiv:1608.03609*, 2016.

Gunnar Sigurdsson, Gül Varol, Xiaolong Wang, Ali Farhadi, Ivan Laptev, and Abhinav Gupta. Hollywood in homes: Crowdsourcing data collection for activity understanding. In *ECCV*, 2016a.

Gunnar A Sigurdsson, Xinlei Chen, and Abhinav Gupta. Learning visual storylines with skipping recurrent neural networks. In *ECCV*, 2016b.

Gunnar A Sigurdsson, Santosh Divvala, Ali Farhadi, and Abhinav Gupta. Asynchronous temporal fields for action recognition. *arXiv preprint arXiv:1612.06371*, 2016c.

Karen Simonyan and Andrew Zisserman. Two-stream convolutional networks for action recognition in videos. In *NIPS*, 2014.

Khurram Soomro, Amir Roshan Zamir, and Mubarak Shah. Ucf101: A dataset of 101 human actions classes from videos in the wild. *arXiv preprint arXiv:1212.0402*, 2012.

Yu-Chuan Su and Kristen Grauman. Leaving some stones unturned: dynamic feature prioritization for activity detection in streaming video. In *ECCV*, 2016.

Du Tran, Lubomir Bourdev, Rob Fergus, Lorenzo Torresani, and Manohar Paluri. Learning spatiotemporal features with 3d convolutional networks. In *ICCV*, 2015.

Jason Weston, Sumit Chopra, and Antoine Bordes. Memory networks. *arXiv preprint arXiv:1410.3916*, 2014.

Ronald J Williams. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 1992.

Kelvin Xu, Jimmy Ba, Ryan Kiros, Kyunghyun Cho, Aaron Courville, Ruslan Salakhudinov, Rich Zemel, and Yoshua Bengio. Show, attend and tell: Neural image caption generation with visual attention. In *ICML*, 2015.

Serena Yeung, Olga Russakovsky, Greg Mori, and Li Fei-Fei. End-to-end learning of action detection from frame glimpses in videos. In *CVPR*, 2016.

Adams Wei Yu, Hongrae Lee, and Quoc V Le. Learning to skim text. In *ACL*, 2017.

Joe Yue-Hei Ng, Matthew Hausknecht, Sudheendra Vijayanarasimhan, Oriol Vinyals, Rajat Monga, and George Toderici. Beyond short snippets: Deep networks for video classification. In *CVPR*, 2015.

Wojciech Zaremba, Ilya Sutskever, and Oriol Vinyals. Recurrent neural network regularization. In *ICLR*, 2015.

Saizheng Zhang, Yuhuai Wu, Tong Che, Zhouhan Lin, Roland Memisevic, Ruslan R Salakhutdinov, and Yoshua Bengio. Architectural complexity measures of recurrent neural networks. In *NIPS*, 2016.

A ADDITIONAL EXPERIMENTS

A.1 FREQUENCY DISCRIMINATION TASK

In this experiment, the network is trained to classify between sinusoids whose period is in range $T \sim \mathcal{U}(5, 6)$ milliseconds and those whose period is in range $T \sim \{(1, 5) \cup (6, 100)\}$ milliseconds (Neil et al., 2016). Every sine wave with period T has a random phase shift drawn from $\mathcal{U}(0, T)$. At every time step, the input to the network is a single scalar representing the amplitude of the signal. Since sinusoid are continuous signals, this tasks allows to study whether Skip RNNs converge to the same solutions when their parameters are fixed but the sampling period is changed. We study two different sampling periods, $T_s = \{0.5, 1\}$ milliseconds, for each set of hyperparameters.

We train RNNs with 110 units each on input signals of 100 milliseconds. Batches are stratified, containing the same number of samples for each class, yielding a 50% chance accuracy. The last state of the RNN is fed into a 2-way classifier and trained with cross-entropy loss. We consider that a model is able to solve the task when it achieves an accuracy over 99% on a held-out set of examples.

Table 4 summarizes results for this task. When no cost per sample is set ($\lambda = 0$), the number of updates differ under different sampling conditions. We attribute this behavior to the potentially large number of local minima in the cost function, since there are numerous subsampling patterns for which the task can be successfully solved and we are not explicitly encouraging the network to converge to a particular solution. On the other hand, when $\lambda > 0$ Skip RNN models with the same cost per sample use roughly the same number of input samples even when the sampling frequency is doubled. This is a desirable property, since solutions are robust to oversampled input signals. Qualitative results can be found in Section B.2.

Model	$\mathbf{T_s = 1}$**ms (length 100)**		$\mathbf{T_s = 0.5}$**ms (length 200)**	
	Task solved	**State updates**	**Task solved**	**State updates**
LSTM	Yes	100.0 ± 0.00	Yes	200.0 ± 0.00
Skip LSTM, $\lambda = 0$	Yes	55.5 ± 16.9	Yes	147.9 ± 27.0
Skip LSTM, $\lambda = 10^{-5}$	Yes	47.4 ± 14.1	Yes	50.7 ± 16.8
Skip LSTM, $\lambda = 10^{-4}$	Yes	12.7 ± 0.5	Yes	19.9 ± 1.5
GRU	Yes	100.0 ± 0.00	Yes	200.0 ± 0.00
Skip GRU, $\lambda = 0$	Yes	73.7 ± 17.9	Yes	167.0 ± 18.3
Skip GRU, $\lambda = 10^{-5}$	Yes	51.9 ± 10.2	Yes	54.2 ± 4.4
Skip GRU, $\lambda = 10^{-4}$	Yes	23.5 ± 6.2	Yes	22.5 ± 2.1

Table 4: Results for the frequency discrimination task, displayed as $mean \pm std$ over four different runs. The task is considered to be solved if the classification accuracy is over 99%. Models with the same cost per sample ($\lambda > 0$) converge to a similar number of used samples under different sampling conditions.

A.2 SENTIMENT ANALYSIS ON IMDB

The IMDB dataset (Maas et al., 2011) contains 25,000 training and 25,000 testing movie reviews annotated into two classes, *positive* and *negative* sentiment, with an approximate average length of 240 words per review. We set aside 15% of training data for validation purposes. Words are embedded into 300-d vector representations before being fed to an RNN with 128 units. The embedding matrix is initialized using pre-trained word2vec[5] embeddings (Mikolov et al., 2013) when available, or random vectors drawn from $\mathcal{U}(-0.25, 0.25)$ otherwise (Kim, 2014). Dropout with rate 0.2 is applied between the last RNN state and the classification layer in order to reduce overfitting. We evaluate the models on sequences of length 200 and 400 by cropping longer sequences and padding shorter ones (Yu et al., 2017).

Results on the test are reported in Table 5. In a task where it is hard to predict which input tokens will be discriminative, the Skip RNN models are able to achieve similar accuracy rates to the baseline

[5] `https://code.google.com/archive/p/word2vec/`

models while reducing the number of required updates. These results highlight the trade-off between accuracy and the available computational budget, since a larger cost per sample results in lower accuracies. However, allowing the network to select which samples to use instead of cropping sequences at a given length boosts performance, as observed for the Skip LSTM (length 400, $\lambda = 10^{-4}$), which achieves a higher accuracy than the baseline LSTM (length 200) while seeing roughly the same number of words per review. A similar behavior can be seen for the Skip RNN models with $\lambda = 10^{-3}$, where allowing them to select words from longer reviews boosts classification accuracy while using a comparable number of tokens per sequence.

In order to reduce overfitting of large models, Miyato et al. (2017) leverage additional unlabeled data through adversarial training and achieve a state of the art accuracy of 0.941 on IMDB. For an extended analysis on how different experimental setups affect the performance of RNNs on this task, we refer the reader to (Longpre et al., 2016).

Model	**Length 200**		**Length 400**	
	Accuracy	**State updates**	**Accuracy**	**State updates**
LSTM	0.843 ± 0.003	200.00 ± 0.00	0.868 ± 0.004	400.00 ± 0.00
Skip LSTM, $\lambda = 0$	0.844 ± 0.004	196.75 ± 5.63	0.866 ± 0.004	369.70 ± 19.35
Skip LSTM, $\lambda = 10^{-5}$	0.846 ± 0.004	197.15 ± 3.16	0.865 ± 0.001	380.62 ± 18.20
Skip LSTM, $\lambda = 10^{-4}$	0.837 ± 0.006	164.65 ± 8.67	0.862 ± 0.003	186.30 ± 25.72
Skip LSTM, $\lambda = 10^{-3}$	0.811 ± 0.007	73.85 ± 1.90	0.836 ± 0.007	84.22 ± 1.98
GRU	0.845 ± 0.006	200.00 ± 0.00	0.862 ± 0.003	400.00 ± 0.00
Skip GRU, $\lambda = 0$	0.848 ± 0.002	200.00 ± 0.00	0.866 ± 0.002	399.02 ± 1.69
Skip GRU, $\lambda = 10^{-5}$	0.842 ± 0.005	199.25 ± 1.30	0.862 ± 0.008	398.00 ± 2.06
Skip GRU, $\lambda = 10^{-4}$	0.834 ± 0.006	180.97 ± 8.90	0.853 ± 0.011	314.30 ± 2.82
Skip GRU, $\lambda = 10^{-3}$	0.800 ± 0.007	106.15 ± 37.92	0.814 ± 0.005	99.12 ± 2.69

Table 5: Accuracy and used samples on the test set of IMDB for different sequence lengths. Results are displayed as $mean \pm std$ over four different runs.

A.3 Action classification on UCF-101

UCF-101 (Soomro et al., 2012) is a dataset containing 13,320 trimmed videos belonging to 101 different action categories. We use 10 seconds of video sampled at 25 fps, cropping longer ones and padding shorter examples with empty frames. Activations in the Global Average Pooling layer from a ResNet-50 (He et al., 2016) CNN pretrained on the ImageNet dataset (Deng et al., 2009) are used as frame-level features, which are fed into two stacked RNN layers with 512 units each. The weights in the CNN are not tuned during training to reduce overfitting. The hidden state in the last RNN layer is used to compute the update probability for the Skip RNN models.

We evaluate the different models on the first split of UCF-101 and report results in Table 6. Skip RNN models do not only improve the classification accuracy with respect to the baseline, but require very few updates to do so, possibly due to the low motion between consecutive frames resulting in frame-level features with high temporal redundancy (Shelhamer et al., 2016). Moreover, Figure 4 shows how models performing fewer updates converge faster thanks to the gradients being preserved during longer spans when training with backpropagation through time.

Non recurrent architectures for video action recognition that have achieved high performance on UCF-101 comprise CNNs with spatiotemporal kernels (Tran et al., 2015) or two-stream CNNs (Simonyan & Zisserman, 2014). Carreira & Zisserman (2017) show the benefits of expanding 2D CNN filters into 3D and pretraining on larger datasets, obtaining an accuracy of 0.845 when using RGB data only and 0.934 when incorporating optical flow information.

Model	Accuracy	State updates	Inference FLOPs
LSTM	0.671	250.0	9.52×10^{11}
Skip LSTM, $\lambda = 0$	0.749	138.9	5.29×10^{11}
Skip LSTM, $\lambda = 10^{-5}$	0.757	24.2	9.21×10^{10}
Skip LSTM, $\lambda = 10^{-4}$	0.790	7.6	2.89×10^{10}
GRU	0.791	250.0	9.51×10^{11}
Skip GRU, $\lambda = 0$	0.796	124.2	4.73×10^{11}
Skip GRU, $\lambda = 10^{-5}$	0.792	29.7	1.13×10^{11}
Skip GRU, $\lambda = 10^{-4}$	0.793	23.7	9.02×10^{10}
I3D (RGB) (Carreira & Zisserman, 2017)	0.845	-	-
Two-stream I3D (Carreira & Zisserman, 2017)	0.934	-	-

Table 6: Accuracy, used samples and average FLOPs per sequence at inference on the validation set of UCF-101 (split 1).

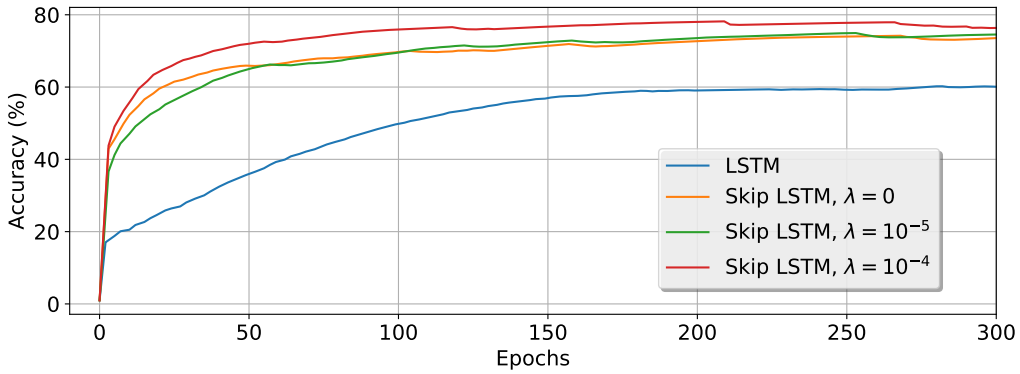

Figure 4: Accuracy evolution during the first 300 training epochs on the validation set of UCF-101 (split 1). Skip LSTM models converge much faster than the baseline LSTM.

B QUALITATIVE RESULTS

This appendix contains additional qualitative results for the Skip RNN models.

B.1 ADDING TASK

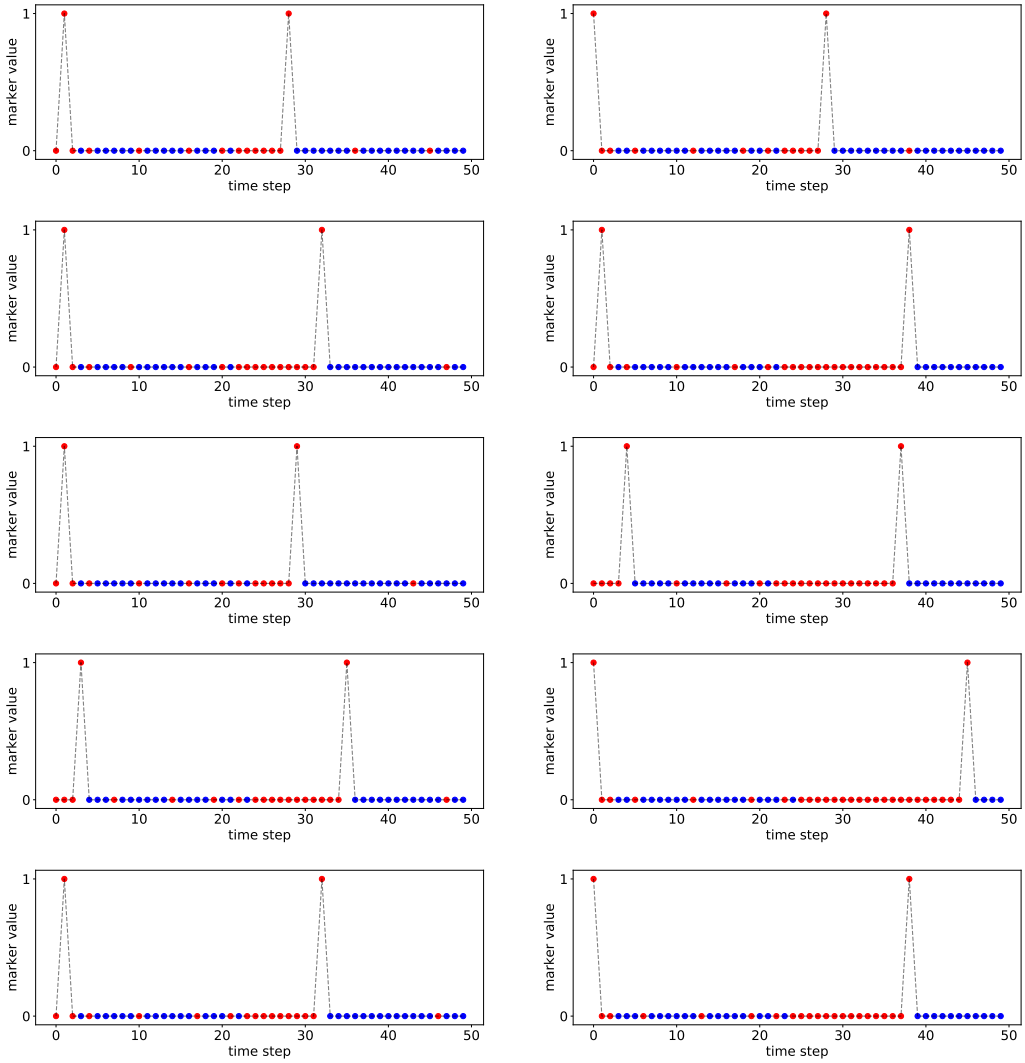

Figure 5: Sample usage examples for the Skip GRU with $\lambda = 10^{-5}$ on the adding task. Red dots indicate used samples, whereas blue ones are skipped.

B.2 FREQUENCY DISCRIMINATION TASK

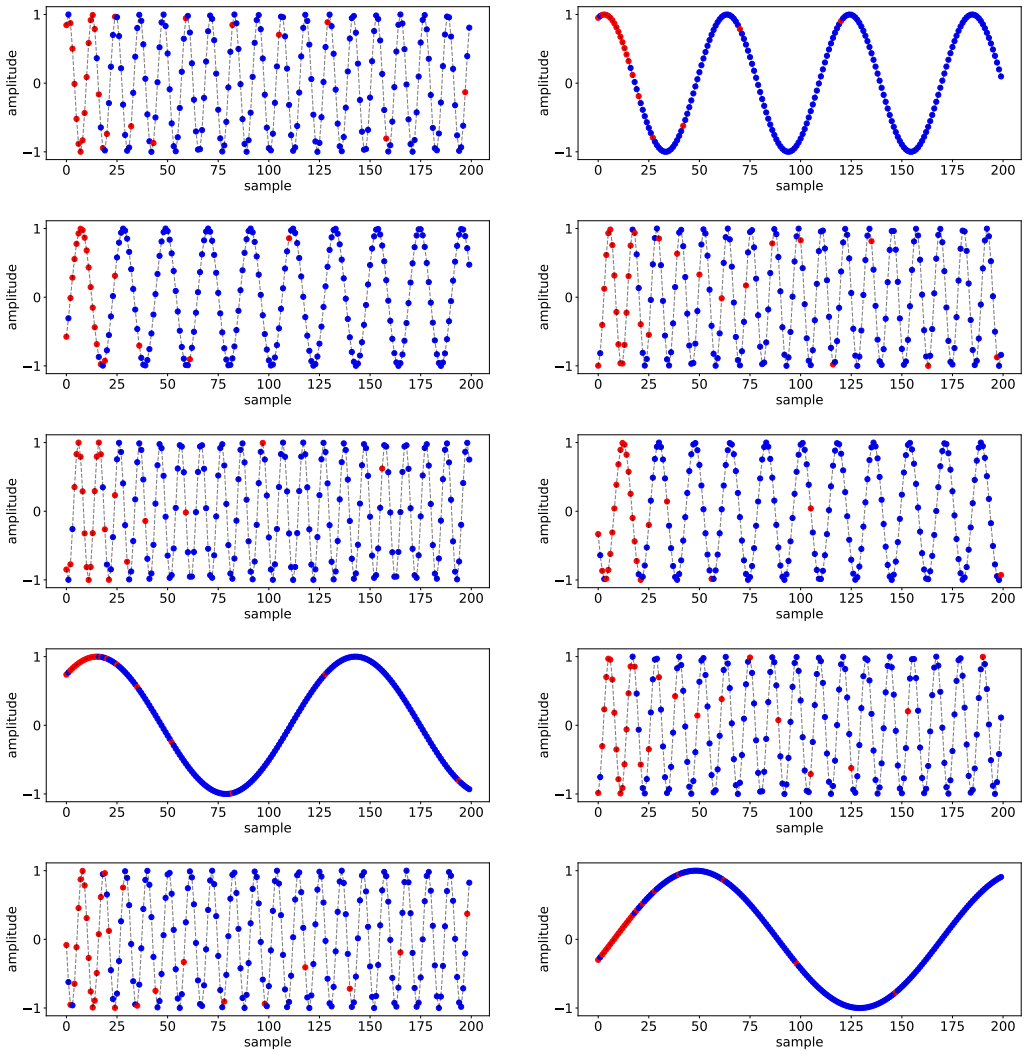

Figure 6: Sample usage examples for the Skip LSTM with $\lambda = 10^{-4}$ on the frequency discrimination task with $T_s = 0.5$ms. Red dots indicate used samples, whereas blue ones are skipped. The network learns that using the first samples is enough to classify the frequency of the sine waves, in contrast to a uniform downsampling that may result in aliasing.