

BENCHMARKING LEARNING EFFICIENCY IN DEEP RESERVOIR COMPUTING

Hugo Cisneros
CIIRC*, WILLOW†

hugo.cisneros@cvut.cz

Josef Sivic
CIIRC

josef.sivic@cvut.cz

Tomas Mikolov
CIIRC

tomas.mikolov@cvut.cz

ABSTRACT

It is common to evaluate the performance of a machine learning model by measuring its predictive power on a test dataset. This approach favors complicated models that can smoothly fit complex functions and generalize well from training data points. Although essential components of intelligence, speed and data efficiency of this learning process are rarely reported or compared between different candidate models. In this paper, we introduce a benchmark of increasingly difficult tasks together with a data efficiency metric to measure how quickly machine learning models learn from training data. We compare the learning speed of some established sequential supervised models, such as RNNs, LSTMs, or Transformers, with relatively less known alternative models based on reservoir computing. The proposed tasks require a wide range of computational primitives, such as memory or the ability to compute Boolean functions, to be effectively solved. Surprisingly, we observe that reservoir computing systems that rely on dynamically evolving feature maps learn faster than fully supervised methods trained with stochastic gradient optimization while achieving comparable accuracy scores. The code, benchmark, trained models, and results to reproduce our experiments are available at https://github.com/hugcis/benchmark_learning_efficiency/.

1 INTRODUCTION

Most machine learning models are evaluated by measuring performance on a specific dataset or task. Learning efficiency – the ability to learn, generalize, and adapt quickly from a few examples – is crucial for practical intelligence (Kanazawa, 2004) as well as low-data machine learning applications yet rarely used to evaluate models. Supervised learning systems are theoretically limited in their learning speed by the optimization algorithms used for training. These algorithms such as stochastic gradient descent (SGD) have various speed guarantees depending on the structure of the function to be optimized (Bottou et al., 2018). However, when intelligent beings learn, they appear to quickly re-use past knowledge and progressively improve over time. Their learning speed depends on a dynamically evolving internal state. To measure the learning efficiency of various systems, we propose in this work the Weighted Average Data Efficiency (WADE) metric based on the time taken to reach several test accuracy checkpoints. We also design a simple modular benchmark composed of a set of sequential tasks. They begin with the task of recognizing a simple periodic sequence in an input string and end with elaborate question answering tasks that require counting occurrences of patterns and long-term memory.

Established sequential supervised models such as recurrent neural networks (RNNs; Elman, 1990), long short-term memory networks (LSTMs; Hochreiter & Schmidhuber, 1997) or Transformers (Vaswani et al., 2017) lack essential properties such as learning beyond the training phase or the ability to adapt over time after being trained. These models can also be expensive to train, requiring a large number of labeled training examples to reach reasonable performance, leading to poor learning speeds. In this work, we use the newly proposed WADE metric and the benchmark dataset to experimentally compare the learning speed of these well established models, such as RNNs, LSTMs and Transformers, to less explored *reservoir computing* models.

Reservoir computing is a computational framework that aims to exploit the states of a complex dynamical system. The simplest example of a reservoir computer is a recurrent neural network (RNN) with frozen weights. This special RNN performs random manipulation on its hidden state in reaction to each new input. Interestingly, it has been shown that with a specific initialization of the frozen weights, these RNNs (called Echo state networks) can keep a memory of past

*Czech Institute of Informatics, Robotics and Cybernetics, Czech Technical University in Prague.

†WILLOW project, Inria and Département d’Informatique de l’École Normale Supérieure, PSL Research University.

inputs (Jaeger, 2001). Usually, a standard linear regression is added as a decoder to extract valuable representation from the hidden state for some downstream task. Freezing the weights of these recurrent models is useful when available supervision is very limited or non-existent or for reinforcement learning with sparse rewards since direct training would be impossible. In such cases a reservoir computer creates a continuously evolving pool of random functions that can be combined using the last trainable layer.

When evolving in response to input stimuli, complex recurrent systems such as RNNs are building dynamically changing representations of data within their internal state (Boccara, 2010). We know that these internal states can be interesting on their own because of their ability to self-organize and exhibit increasingly complex behaviors (Koppel & Atlan, 1991; Bennett, 1995; Allen & Strathern, 2003; Goldstein, 2011; Cisneros et al., 2019). In this paper we wish to investigate whether complex dynamical systems — in particular RNNs with frozen weights (echo-state networks) (Jaeger, 2001) and reservoir cellular automata (Yilmaz, 2014) — create representations that allow them to learn faster as measured by our metric.

Contributions. In this paper, we make the following main contributions: First, we introduce the Weighted Average Data Efficiency (WADE) metric to measure the learning speed of various learning systems and use it to benchmark a few standard models on the IMDB text classification task (Maas et al., 2011). Second, we present a benchmark of language-based tasks of increasing difficulty to evaluate the learning speed in different conditions. The proposed tasks require a wide range of computational primitives, such as memory or the ability to compute Boolean functions, to be effectively solved. Third, we study the learning speed of reservoir computing learning models and compare them with more standard supervised solutions.

2 RELATED WORK

The WADE metric is a generalization of the *Time-to-threshold* metric (Taylor & Stone, 2007; Taylor et al., 2007) introduced for measuring transfer learning in reinforcement learning contexts. In general, the Time-to-threshold is simply defined as the number of training steps needed to reach a fixed threshold performance. However this definition leaves open the choice of threshold or the definition of a training step. WADE alleviates this issue by aggregating several of these thresholds into a single number that summarizes the learning speed.

Other metrics for measuring how quickly a model adapts to new tasks have been introduced in the context of transfer learning, few-shot and zero-shot learning. In few-shot learning, one tries to obtain the best performance for a particular task using a small amount of labeled data compared to the task’s fully supervised equivalent (Wang et al., 2020). This correlates with a model’s learning speed, but these problems often measure how much prior information about similar data has been encoded in the models. With our benchmark and the WADE metric, we explicitly measure the number of steps to reach multiple test accuracy values using all the data needed, effectively emphasizing data efficiency.

Sample efficiency has also been studied in the context of reinforcement learning. Chevalier-Boisvert et al. (2018) use the number of demonstrations before a task is solved to measure sample efficiency. This requires defining what solving the task means, which may vary from task to task. Another approach is to measure performance (cumulated reward, accuracy, etc.) after a fixed budget of training steps (Yarats et al., 2019). In this case, the most efficient model is the one that achieves the best performance within the allocated budget. In other cases, the sample efficiency is mentioned but not explicitly measured and one has to examine the learning curves (Buckman et al., 2018). The WADE metric is a general approach to measure the learning efficiency of machine learning models. We use it to benchmark a few standard models on the IMDB text classification tasks (Maas et al., 2011) and propose a set of modular and extensible language based tasks.

Synthetic tasks such as ours have played a vital role in a series of crucial advances in machine learning algorithms. For example, the XOR problem has partially motivated the development of neural networks (Minsky & Papert, 1972; Rumelhart et al., 1985), and the *circle and ring* dataset has inspired the creation of novel clustering algorithms (Ng et al., 2001). The design of synthetic datasets has also been an essential component of the development of learning algorithms with memory and general computational capabilities (Hochreiter & Schmidhuber, 1997; Joulin & Mikolov, 2015; Graves et al., 2014; Weston et al., 2016; Richardson et al., 2020).

Other tasks are based on real datasets with artificial manipulations (Krizhevsky & Hinton, 2009; Srivastava et al., 2013; Goodfellow et al., 2014; Nguyen et al., 2017). The goal of our dataset is to be truly progressive in difficulty yet simple to understand and extend, to allow applications in the field of online learning, and to easily understand a model’s basic computational capacities. Combined with our metric, it enables us to measure learning speed across a range of conditions. In contrast to similar synthetic datasets, we built this benchmark so that the last task is vastly more complicated than the first and could still be extended to more complex examples.

3 A BENCHMARK FOR RESERVOIR COMPUTING

To measure the learning speed of candidate systems and their ability to improve over time, we propose a performance metric and a standardized set of tasks. We want to select those systems that quickly and reliably adapt and learn from new inputs. For this purpose, we introduce the Weighted Average Data Efficiency (WADE) metric. It aggregates the speed at which a model reaches several test accuracy checkpoints. We describe the metric in more detail in Section 3.1.

To reliably compare learning speeds for various systems on a shared foundation, we also introduce a novel dataset described in Table 1. It is made up of sequential tasks that begin with straightforward pattern recognition and progressively increase in complexity to approach the complexity of natural language and other complex real-world tasks.

We do not focus on the prediction performance of our models, but rather on their data efficiency — the number of example sequences they need to learn from before reaching a target accuracy on a validation set.

| Task id | Name | Description |
|---------|---|---|
| 1 | Simple periodic pattern identification | Identify a simple periodic pattern. |
| 2 | Harder periodic pattern identification | Identify a periodic pattern with an arithmetically increasing period. |
| 3 | Symbol counting | Count symbols from a sequence. |
| 4 | Pattern counting | Count patterns (delimited group of symbols) from a sequence. |
| 5 | Simple question answering | Answer simple YES/NO questions from a single prompt. |
| 6 | Harder question answering | Answer simple YES/NO questions from a single prompt with a more extensive vocabulary. |
| 7 | Question answering with world definition | Answer YES/NO questions from a sequence of prompts. |
| 8 | Question answering with world definition and counting | Answer YES/NO and counting questions from a sequence of prompts. |
| 9 | Adjective question answering | Answer YES/NO and adjective questions from a sequence of prompts. |
| 10 | Adjective question answering and counting | Answer YES/NO, adjective, and counting questions from a sequence of prompts. |

Table 1: General description of all the tasks in the benchmark.

Standard benchmarks and metrics such as those introduced in this paper have always been essential in advancing various aspects of machine learning. For example, the LSTM network demonstrated a superior memory capacity on a set of synthetic tasks designed to challenge the memory of sequential learning systems (Hochreiter & Schmidhuber, 1997). Our goal with this benchmark is to emphasize measuring learning speed across tasks of varying difficulties with a range of computational requirements rather than focusing on performance only. We describe the performance metric and the benchmark next.

3.1 PERFORMANCE METRIC

We introduce the Weighted Average Data Efficiency (WADE) metric as a way to measure how quickly a model learns, using a weighted average of inverse times taken to reach various test accuracy *checkpoints* over time.

It is computed for an evenly distributed set of target accuracies \mathbb{A} . They represent the *checkpoints* at which the speed of learning is estimated. For example, we may choose $\mathbb{A} = [0.1, 0.2, 0.3, 0.4, \dots, 1.]$. The metric is then calculated as

$$\text{WADE}(\mathbf{a}) = \frac{1}{\sum \alpha} \sum_{\alpha \in \mathbb{A}} \frac{\alpha}{T(\alpha, \mathbf{a})}, \quad (1)$$

where $\mathbf{a} = (a_0, a_1, \dots, a_n)$ is a sequence of test accuracies achieved by the evaluated system sampled at different training steps. The quantity a_i typically corresponds to the accuracy reached after seeing i examples, and $T(\alpha, \mathbf{a})$ is the number of steps in the sequence \mathbf{a} needed to reach an accuracy of α . It is defined as

$$T(\alpha, \mathbf{a}) = \min \{i \in \{1, \dots, n, +\infty\} \mid a_i \geq \alpha\}. \quad (2)$$

We also define $T(\alpha, \mathbf{a}) = +\infty$ if the accuracy value α is never reached in \mathbf{a} . This is equivalent to appending an additional term $a_{+\infty}$ to \mathbf{a} , always set to the maximum accuracy 1. Note that by construction, $T(\alpha, \mathbf{a})$ is in $[1, +\infty[$.




Figure 1: Illustration of the calculation of $T(\cdot, \cdot)$, representing the number of training steps (x-axis) needed to reach a certain test accuracy α (y-axis) from a learning curve. In this example, $\mathbb{A} = [0.1, \dots, 0.8]$ (y-axis). $T(0.6, \mathbf{a})$ is highlighted in red. $T(0.8, \mathbf{a}) = +\infty$ as the accuracy of 0.8 is never reached.

Since T can be $+\infty$ we define $\frac{1}{+\infty} = 0$ for the quantity in equation 1 to always exist. A visual intuition of $T(\cdot, \cdot)$ is given in figure 1.

The choice of checkpoints \mathbb{A} does not need to be tuned in any specific way because $\text{WADE}(\mathbf{a})$ converges quickly to a single value when \mathbb{A} approaches the continuous interval $[0, 1]$. The approximation is good enough as long as \mathbb{A} is not too coarse (more than ten elements was enough in our experiments) and the WADE values computed from the same set \mathbb{A} are comparable.

The time-to-threshold T is always above or equal to 1 step for any threshold and sequence of accuracy scores. We have $\forall \alpha \in [0, 1], \forall \mathbf{a} = (a_n)_{n \in \mathbb{N} \cup \{+\infty\}}$, it holds that

$$\frac{1}{T(\alpha, \mathbf{a})} \leq 1, \quad (3)$$

and therefore we always get that $0 \leq \text{WADE}(\mathbf{a}) \leq 1$. The metric is equal to 0 for systems that never get past the smallest possible accuracy while a 1 corresponds to reaching a perfect test accuracy in only one training step. Such a system would also be considered to be performing well according to the underlying performance metric with which it is usually evaluated. Therefore, maximizing WADE also maximizes performance.

3.2 DESCRIPTION OF TASKS IN THE BENCHMARK

This section provides a more detailed description of each task in our benchmark. The tasks are designed to be language modeling tasks, where the goal is to predict some tokens from sequences of previously processed tokens. An overview of the tasks is given in Table 1. The tasks are divided into three groups: (i) binary tasks – with only binary symbols, (ii) general symbolic tasks – symbolic manipulations with arbitrary symbols – and (iii) language-based tasks – the symbols represent words in English and behave like a language. We introduce this benchmark together with the WADE metric, but both can be used in other contexts as well to measure the learning speed of other systems. Individual sentences are generated and divided into a training set and a test set for periodic evaluation of the test accuracy. We give a more detailed description of each task below¹:

3.2.1 BINARY

Simple periodic pattern identification. The goal of the periodic binary task is for the model to learn a fixed-length regular pattern. As the system is presented with new binary input tokens, it has to learn the periodic pattern on the fly and correctly predict the next token. A pattern of size n is chosen at random and repeated k times to produce a sequence of length $n \times k$. Examples include:

| | |
|--------------------------------------|-----------------------|
| 010101010101010101010101010101010101 | Pattern with period 2 |
| 0011001100110011001100110011001100 | Pattern with period 4 |
| 0110110110110110110110110110110110 | Pattern with period 3 |

Harder periodic pattern identification. For this task, we also draw a random binary pattern of size n . Each of its symbols is repeated k times, with k increasing monotonically from 1. A successful model must learn the pattern on the fly and correctly implement the arithmetic increase in the size of the period. We set the pattern length to increase by 1 every period in our experiments, but this value can be changed.

3.2.2 SYMBOLIC COUNTING

These tasks consist of reading patterns from an input sequence and answering a simple query about the number of patterns. Unlike the previous tasks, these require implementing a form of addressable memory that can be queried after the prompt has ended.

Basic symbol counting. The first version of the counting task focuses on counting single symbols from an input sequence. The sequence ends with a *query* for the count of one of the symbols. The goal is to predict the last token (in bold) of sequences of the following form:

| | | | |
|--------------------|----------|----------|----------|
| <u>AABBCBABAAB</u> | <u>x</u> | <u>A</u> | <u>5</u> |
| Input symbols | QS | | Answer |

The symbol x is the query symbol (QS) that marks the beginning of the query. In the first example above, the goal is to predict the token 5 because the symbol A appears 5 times. As detailed in Sect. 5, we represent these nonbinary symbols with one-hot encoding, so the numerical nature of some tokens is not encoded a priori.

¹The tasks are also available as a [Python package on GitHub](#).

Pattern counting. This aim of this task is to count the number of occurrences of delimited patterns instead of single symbols. A sequence is still divided between a prompt — before x — and a query — after x . One has to predict the symbol coming after each separator symbol (S) y in the query part of the sentence. For example, sentences are of the form:

$\underbrace{AA}_{\text{Pattern 1}} \underbrace{y}_S \underbrace{BBC}_{\text{Pattern 2}} \underbrace{y}_S \underbrace{BAB}_{\text{Pattern 3}} \underbrace{y}_S \underbrace{AA}_{\text{Pattern 4}} \underbrace{y}_S \underbrace{B}_{\text{Pattern 5}} \underbrace{x}_{QS} \underbrace{AAy}_{\text{Query 1}} \underbrace{2}_{\text{Answer 1}} \underbrace{By}_{\text{Query 2}} \underbrace{1}_{\text{Answer 2}}$

Multiple queries are presented successively, which requires keeping and being able to retrieve several counts simultaneously. A query is composed of a pattern, a separator symbol, and the pattern count that the system should predict.

3.2.3 BASIC LANGUAGE UNDERSTANDING

To make the tasks progressively more complex, we steer them towards general language understanding tasks. The tasks described below are generated automatically, but gradually incorporate more complex skills required for advanced language processing. The last task is a step towards understanding general language albeit with a limited vocabulary.

Elementary question answering (QA). This task introduces elements of natural language. Each example is composed of a stated fact and a question about that fact. A sentence is constructed from a few basic elements: (i) Names (e.g., JOHN, JAMES, etc.), (ii) Verbs (e.g., HEAR, SEE, etc.), (iii) Answers (YES or NO), (iv) Additional words and symbols (I, DO, NOT, AND, BUT, ?, .).

A random subset of names is selected and we generate a random prompt/question pair from it. The question is drawn to ensure an equal proportion of positive and negative answers. For example, sentences may look like this:

I HEAR JOHN AND PAUL . DO I HEAR PAUL ? **YES**
 I SEE JOHN BUT I DO NOT SEE PAUL AND TOM . DO I SEE TOM ? **NO**

The only token to predict is the binary answer YES or NO.

Question answering (QA) with adjectives. This task extends the previous task by adding adjectives and modifiers to the object names. The queries may be about the subject-verb relation or the subject-adjective relation.

I SEE A SMALL BANANA . WHAT IS THE SIZE OF THE BANANA I SEE ? **SMALL**
 I SEE A LARGE GREEN APPLE BUT I DO NOT SEE A RED APPLE .
 DO I SEE A LARGE APPLE ? **YES**
 I SEE A SMALL GREEN APPLE BUT I DO NOT SEE A BANANA .
 WHAT IS THE COLOR OF THE APPLE I SEE ? **GREEN**

Here the task output space is slightly larger because the model may be predicting YES, NO, SMALL, GREEN, etc.

Question answering (QA) with world definition. This task introduces more complex configurations in which the state of the world is defined in one or more sentences, and an unknown number of questions follow. This is more akin to a real-world conversation where stored facts should be remembered longer and accessed on demand. For example, below we show a generated group of sentences followed by several questions:

I SEE A SMALL BANANA .
 I SEE A LARGE GREEN APPLE BUT I DO NOT SEE A RED APPLE .
 I HEAR A SMALL GREEN APPLE BUT I DO NOT SMELL A BANANA .
 WHAT IS THE COLOR OF THE APPLE I SEE ? **GREEN**
 HOW MANY THINGS DO I SMELL ? **ONE**
 DO I SEE A LARGE APPLE ? **YES**.

The difficulty of each of these tasks can be modulated by changing the size of the base vocabulary, the length of sequences, or the number of queries. Our particular setting of these parameters will be described in section B.1.

4 EVALUATING SPEED OF LEARNING ON A STANDARD LANGUAGE CLASSIFICATION TASK

To show the usefulness of measuring WADE on standard language classification tasks, we train various text classifiers on the IMDB dataset and compare their WADE scores with a usual performance metric for classification: accuracy. This classification task, first proposed by Maass et al. (2002), consists of deciding if a movie review is positive or negative from its text. It contains 25000 training examples and 25000 test examples. The two labels are balanced in both the training and test set. This dataset is a high dimensional language-based task with a binary output.

We study five standard models: (i) an Elman recurrent neural network (RNN) (Elman, 1990) with tanh activation functions trained with backpropagation through time. (ii) A long-short term memory (LSTM) recurrent neural network (Hochreiter & Schmidhuber, 1997), also trained with backpropagation through time. (iii) A gated recurrent unit (GRU) recurrent neural network Cho et al. (2014). (iv) A standard encoder-only transformer neural network model Vaswani et al. (2017). (v) A logistic regression using a bag-of-words representation of each sentence as input features. All the models are trained with batches of training data using the Adam optimization algorithm Kingma & Ba (2015). Each model’s hyperparameters are chosen to ensure they have a similar number of trainable parameters except for the logistic regression whose parameter count is solely determined by the input and output dimensions.

| | WADE $\times 10^{-2}$ (std.) \uparrow | Max test accuracy (std.) \uparrow |
|-------------|---|-------------------------------------|
| RNN | 1.028 ± 0.283 | 0.705 ± 0.076 |
| LSTM | 2.280 ± 0.303 | 0.902 ± 0.002 |
| GRU | 2.711 ± 0.318 | 0.904 ± 0.002 |
| Linear | 3.737 ± 0.745 | 0.862 ± 0.000 |
| Transformer | 8.716 ± 0.720 | 0.872 ± 0.003 |

Table 2: Comparison of our new WADE metric to assess learning speed and the standard maximum test accuracy on the IMDB classification dataset. Results are averaged over 50 separate runs. (\uparrow indicates that higher is better).

Table 2 shows the results of all models on the IMDB dataset reporting both the standard test accuracy as well as our new WADE metric measuring the learning speed of the different models. We report the maximum test accuracy observed during training. This corresponds to the model checkpoint that would be selected with a validation set before using it on test data. The transformer model learns the fastest of all, as shown by its higher WADE score, but it does not reach test accuracy values as high as the GRU and LSTM model. We think that this could be attributed to the transformer model over-fitting on the available data as discussed below.

According to the WADE metric, the RNN is the worst (i.e., slowest to learn) model with a score of 1.04×10^{-2} whereas the linear model and the transformer are the fastest with a score above 8×10^{-2} . The transformer seems to have more reliable learning speeds than the linear model as shown by the lower standard deviation (0.7×10^{-2} versus 4.3×10^{-2}). The LSTM and GRU models have intermediate and stable learning speeds. The GRU has slightly better WADE score than the LSTM, which can be also seen by inspecting the learning curves in figure 2 where the GRU is consistently above the LSTM.

However, when we look at the maximum test accuracy we obtain a different ordering, with LSTM and GRU performing best with slightly above 90% accuracy followed by the transformer and the linear model (around 86%) and finally the RNN with the lowest accuracy. The LSTM, GRU and linear models all seem to have converged at the end of the experiment. The transformer’s lower final accuracy despite it being considered a state of the art model may be explained by over-fitting (apparent on the graph with the test accuracy score going down). Moreover, the RNN clearly hasn’t converged at the end of the experiment which also explains the relatively low max accuracy score.

Our new WADE metric gives a complementary view of a model’s abilities, which may be significantly different from what can be obtained from the usual performance metrics focused on the final accuracy of the model.




Figure 2: Test accuracy curves for each model. Shaded areas are $\pm 1\sigma$ around the average over 50 runs.

5 NEW BENCHMARK: COMPARED METHODS

In the following sections, we carry out experiments to measure Weighted Average Data Efficiency (WADE) on the benchmark described in Table 1 and introduced in section 3.2. We compare several baseline models to understand their learning efficiency: three of the most common sequential machine learning models for which all parameters are trained — RNNs, LSTMs, and transformers — with two methods using reservoir computing: the echo-state networks (ESN, with a random RNN reservoir) and reservoir cellular automata (ReCA, with a cellular automaton reservoir) for which only a fraction of the parameters are learned. Table 3 presents all the methods we study in our experiments.

In all our tasks, the input data is assumed to be categorical and sequential. Tokens are observed one by one in sequence, we define this input as $X = [X_1, \dots, X_t, \dots]$, $t \in \mathbb{N}$, $\forall t X_t \in \mathcal{X} \subset \mathbb{N}$, where t is the time index of the sequence, each X_i is a token corresponding to time index i , and \mathcal{X} is the set of numbered input categories — or different tokens in the vocabulary. First, each categorical input vector is one-hot encoded into a vector of size L , where $L = |\mathcal{X}|$ is the size of the input vocabulary. We define \mathbf{x}_t as the encoded vector form of X_t , and we have $\forall t, \mathbf{x}_t \in \{0, 1\}^L$, with

| Methods | | Type |
|-------------|---------------------------------|-----------------|
| RNN | Recurrent neural networks | Fully trained |
| LSTM | Long-short term memory networks | |
| Transformer | | |
| ESN | Echo-state networks | Reservoir-based |
| ReCA | Reservoir Cellular Automata | |

Table 3: Summary of the compared models: three are fully trained, and two are reservoir-based.

$\sum_{i=1}^L (x_t)_i = 1$. Since the tasks are designed to be generally compatible with language modeling, the input and output vocabularies are the same.

5.1 FULLY TRAINED SEQUENTIAL MODELS (RNN, LSTM AND TRANSFORMER)

We first study two standard supervised recurrent models: (i) an Elman recurrent neural network (RNN) (Elman, 1990) with tanh activation functions trained with backpropagation through time. (ii) A long-short term memory (LSTM) recurrent neural network (Hochreiter & Schmidhuber, 1997), also trained with backpropagation through time. (iii) An encoder-only transformer model with positional encoding (Vaswani et al., 2017).

The three models are trained with a batched Adam optimization algorithm (Kingma & Ba, 2015) to minimize a cross-entropy loss function between the predicted and target tokens. No other training device, such as dropout, regularization, or normalization, is used for these fully trained baselines.

5.2 ECHO-STATE NETWORKS AND RESERVOIR CELLULAR AUTOMATA

We use names adapted from Jaeger (2012) in this section. An echo-state network is a random recurrent neural network with frozen weights and skip-connections. Its random weights are sampled in a specific way so that it performs random combinations of input vectors with its current state while keeping a history of past inputs (Jaeger, 2001).

Reservoir cellular automaton (ReCA) is a model similar to echo-state networks but where a cellular automaton (CA) replaces the random RNN. The cellular automaton can be seen as a RNN with additional weight-sharing similar to convolutional neural network. The special structure of a recurrent CA update makes it more likely than random RNNs to generate complex structures (Wolfram, 1983; 2002). Because CAs were not designed to make use of inputs or produce outputs, we extend the model to make it accept input vectors and to make reading from the CA state possible (details in Section C.2).

5.3 EXPERIMENTAL SET-UP

We ran 100 separate experiments with different random seeds for each of the CA rules, the RNN, LSTM, Transformer, and ESN on each task in the benchmark. These experiments have a separate input projection matrix for the CA, different random weights for the ESN, and different weight initialization for the supervised baselines. The task inputs are also generated from a new seed for every experiment, but we reuse the seeds for the same experiment on different models to ensure they were trained with the same data. Intervals of one standard deviation for these multiple experiments are reported in the result graphs. The code to reproduce our experiments is available on GitHub².

5.3.1 TRAINING PARAMETERS

For each experiment, we generate 1200 random examples from the task generator. We split this set randomly into a training set with 80% of the data — 960 examples — and a test set with the remaining 240 examples. The reservoir is run on each training example for the reservoir-based models, which creates the input features for training the decoder. The sequential supervised models use batches of single sequences with the Adam algorithm (Kingma & Ba, 2015). They are trained for ten epochs in total. With reservoir models, only the last linear layer (the decoder) is trained. We minimize the cross-entropy loss with stochastic gradient descent (SGD), doing only a single pass over the 960 training examples.

²https://github.com/hugcis/benchmark_learning_efficiency

Every few training steps, we generate the output predictions on the testing set, decode it, and compute the test accuracy for our WADE metric. Supervision is only applied on tokens that can be predicted — similar to masked language modeling, e.g., only the answer token is used in the symbol counting or question answering tasks. Accuracy is also computed for these symbols only.

6 RESULTS

We report the final weighted average data efficiency (WADE) scores on our benchmark in Table 4 and accuracy results for comparison in Table 5. We include the best elementary reservoir cellular automaton (ReCA) and echo state network (ESN) with the same internal state size for each task.

| Task ID - Name | Reservoir | | Fully supervised | | | Human expert |
|-----------------------------------|------------------------|------------------------|------------------------|-----------------|-----------------|--------------|
| | ReCA | ESN | RNN | LSTM | Transformer | |
| 1 - Periodic | 0.78 ± 0.05 | 0.74 ± 0.05 | 0.31 ± 0.06 | 0.28 ± 0.07 | 0.31 ± 0.04 | 1.00 |
| 2 - Incremental periodic | 0.57 ± 0.02 | 0.72 ± 0.02 | 0.42 ± 0.16 | 0.32 ± 0.11 | 0.49 ± 0.15 | 1.00 |
| 3 - Symbol counting | 0.06 ± 0.01 | 0.04 ± 0.01 | 0.05 ± 0.02 | 0.04 ± 0.02 | 0.03 ± 0.01 | 0.11 |
| 4 - Pattern counting | 0.12 ± 0.04 | 0.14 ± 0.03 | 0.13 ± 0.04 | 0.10 ± 0.05 | 0.08 ± 0.03 | 0.09 |
| 5 - Basic question answering (QA) | 0.31 ± 0.08 | 0.26 ± 0.03 | 0.20 ± 0.06 | 0.17 ± 0.08 | 0.16 ± 0.07 | 0.15 |
| 6 - Harder QA | 0.35 ± 0.11 | 0.26 ± 0.03 | 0.24 ± 0.05 | 0.16 ± 0.07 | 0.12 ± 0.06 | 0.12 |
| 7 - QA with world def. | 0.32 ± 0.10 | 0.27 ± 0.04 | 0.24 ± 0.05 | 0.17 ± 0.07 | 0.11 ± 0.06 | — |
| 8 - QA with world def. & counting | 0.09 ± 0.03 | 0.14 ± 0.06 | 0.10 ± 0.05 | 0.05 ± 0.02 | 0.02 ± 0.01 | — |
| 9 - Adjective QA | 0.04 ± 0.03 | 0.04 ± 0.04 | 0.05 ± 0.03 | 0.03 ± 0.02 | 0.02 ± 0.01 | — |
| 10 - Adjective QA & counting | 0.04 ± 0.02 | 0.06 ± 0.02 | 0.05 ± 0.02 | 0.03 ± 0.02 | 0.01 ± 0.01 | — |

Table 4: Comparison of WADE scores (also shown in parentheses in the legend in figure 3, higher is better) of the best cellular automaton rules (**ReCA**, Yilmaz (2014)) for each task against an echo-state network (**ESN**, Jaeger (2001)), a **RNN**, a **LSTM** (Hochreiter & Schmidhuber, 1997), and a **Transformer** (Vaswani et al., 2017) with the same number of parameters. The dash “—” indicates that the task was too difficult to complete from memory alone for the human expert. Accuracy scores are also reported in Table 5.

Figure 3 shows the learning curves for the best reservoir cellular automaton, as well as the echo-state network, recurrent neural network (RNNs), LSTM, and Transformer on all ten tasks. Interestingly, the reservoir based-models are consistently more efficient learners than the fully supervised methods, reaching better accuracy in much fewer training steps. For example, for tasks 5, 6, and 7, the LSTM needs ten times more steps to approach the accuracy the reservoir CA model reached in less than 1000 training steps.

Echo-state networks (ESN) and reservoir cellular automata (ReCA) appear to learn at similar rates, with no clear advantage for one or the other as each model outperforms the other on five tasks out of ten. On tasks 2, 4, 8, and 10 the ESN is visibly faster than the ReCA (see curves Figure 3), which explains the ESN’s higher WADE scores even though ReCA reaches higher accuracy values after several more training steps.

Even if they lack slightly in accuracy, as seen for example on the curves of task 4, the two reservoir-based methods consistently outperform the fully supervised sequential methods (RNNs and LSTMs) in terms of learning speeds. This better learning efficiency could be explained by the internal state structure introduced by the CA rules and the special form of the ESN random matrix which favors memory retention whereas usual recurrent network initialization does not — we initialize the weights of the full trained models from a uniform distribution $\mathcal{U}(-\sqrt{h^{-1}}, \sqrt{h^{-1}})$ in our experiments, where h is the hidden size.

The human expert scores are estimated based on the authors’ performance on obfuscated versions of the tasks — all input tokens are mapped to random symbols. These scores may appear surprisingly low in some of these experiments. Although the task examples presented in Section 3.2 seem easy to understand, we rely heavily on our prior knowledge about the symbols to understand the patterns — the words in tasks 5–10 or the numbers in tasks 3 and 4. These symbols are randomly remapped in our human experiments, and more examples as well as a good memory are needed to understand the tasks and learn the mapping itself, hence the slower learning.

It is interesting to note that RNNs also seem to perform better than LSTMs and Transformers at the beginning of the training in all tasks but the first two. Even though the Transformer is nowadays generally considered a superior model, vanilla RNNs may still remain competitive in the very low data and computation regimes. This is especially pronounced for tasks 8, 9, and 10, where the RNN test accuracy curve starts increasing significantly 1000 steps before

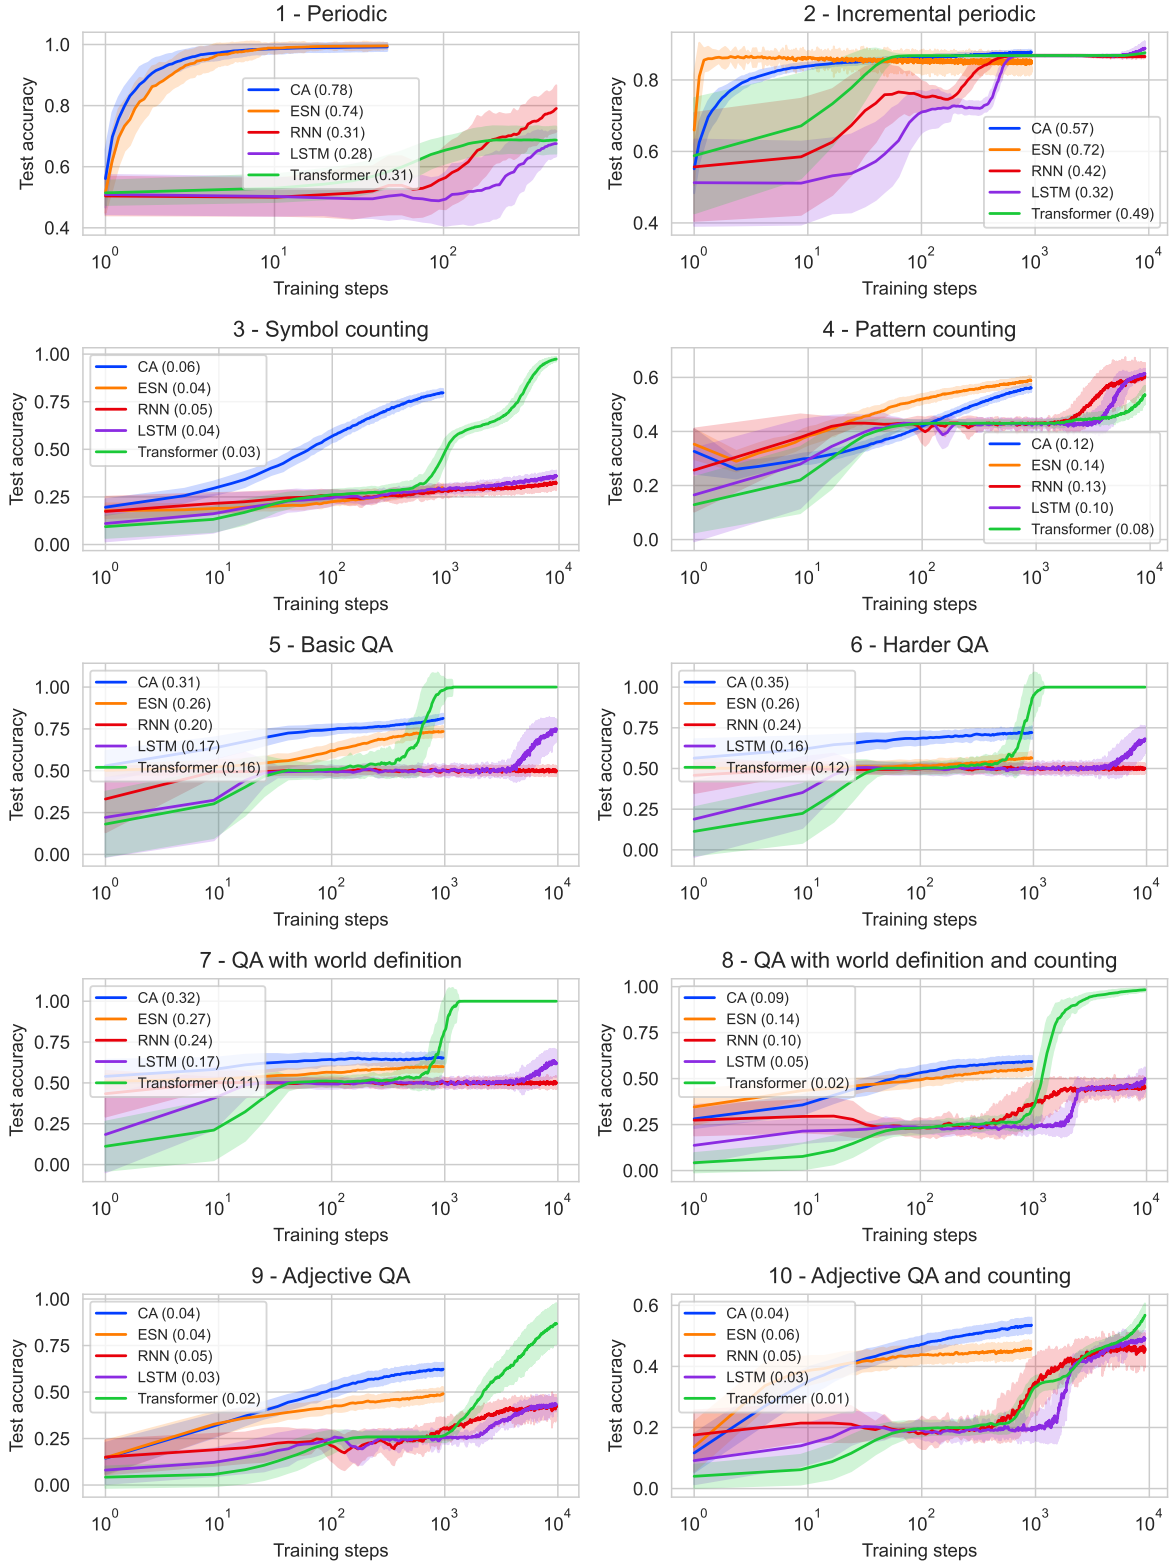


Figure 3: Average learning curves and WADE scores (shown in parentheses in the legend) for each task in the benchmark. The dark blue curves represent the best elementary cellular automata (CA) rule and the green curves represent the echo-state networks (ESN). Note that the x-axis is logarithmic, showing ten times more training steps for the RNN and LSTM. Shaded areas represent one standard deviation around the average over the 100 different experiments.

| Task ID - Name | Reservoir | | Fully supervised | | |
|--|-----------------|------------------------|------------------|------------------------|------------------------|
| | ReCA | ESN | RNN | LSTM | Transformer |
| 1 - Periodic | 0.99 \pm 0.01 | 1.00 \pm 0.01 | 0.79 \pm 0.08 | 0.68 \pm 0.04 | 0.69 \pm 0.04 |
| 2 - Incremental periodic | 0.88 \pm 0.01 | 0.87 \pm 0.03 | 0.87 \pm 0.00 | 0.89 \pm 0.02 | 0.88 \pm 0.00 |
| 3 - Symbol counting | 0.80 \pm 0.02 | 0.30 \pm 0.02 | 0.33 \pm 0.03 | 0.36 \pm 0.03 | 0.97 \pm 0.01 |
| 4 - Pattern counting | 0.56 \pm 0.01 | 0.59 \pm 0.02 | 0.61 \pm 0.05 | 0.61 \pm 0.02 | 0.54 \pm 0.03 |
| 5 - Basic QA | 0.81 \pm 0.03 | 0.73 \pm 0.03 | 0.51 \pm 0.02 | 0.75 \pm 0.05 | 1.00 \pm 0.00 |
| 6 - Harder QA | 0.72 \pm 0.04 | 0.57 \pm 0.03 | 0.51 \pm 0.03 | 0.68 \pm 0.07 | 1.00 \pm 0.00 |
| 7 - QA with world def. | 0.66 \pm 0.03 | 0.60 \pm 0.03 | 0.51 \pm 0.04 | 0.64 \pm 0.07 | 1.00 \pm 0.00 |
| 8 - QA with world def. & counting | 0.59 \pm 0.03 | 0.56 \pm 0.03 | 0.47 \pm 0.05 | 0.50 \pm 0.06 | 0.98 \pm 0.01 |
| 9 - Adjective QA | 0.62 \pm 0.04 | 0.49 \pm 0.03 | 0.43 \pm 0.04 | 0.44 \pm 0.03 | 0.87 \pm 0.11 |
| 10 - Adjective QA & counting | 0.54 \pm 0.02 | 0.46 \pm 0.03 | 0.47 \pm 0.04 | 0.49 \pm 0.03 | 0.57 \pm 0.04 |

Table 5: Comparison of accuracy scores (higher is better) of the **ReCA**, **ESN**, **RNN**, **LSTM**, and **Transformer** models with a similar number of parameters. In contrast to the results of table 4, the fully-supervised models are more performant when we measure the accuracy, as shown here, but do not necessarily do as well when we measure the speed of learning, as shown in table 4.

the LSTM. After several epochs, the Transformer still outperforms other alternatives in terms of test accuracy on most tasks.

We note that we implicitly assume a fixed cost per training iteration in our experiments and that each training example is seen once. Without these requirements, one could achieve higher WADE results at the cost of additional computations and memory usage by using *replay*-inspired methods (Hinton & Plaut, 1987; Robins, 1993; Gepperth & Karaoguz, 2016; Rebuffi et al., 2017) that retrain the models with past stored inputs. The results in Table 4 used each input sequence once.

7 CONCLUSIONS

Learning speed and data efficiency are essential components of any learning system. Our learning speed metric can offer a novel perspective on several machine learning models. Instead of focusing on pure performance, measuring and comparing the data efficiency of different models will hopefully lead to better systems for online and continual learning.

Even with the right metric, it is difficult to thoroughly evaluate a model’s ability to learn efficiently. Our benchmark evaluates a range of problem difficulties which would be challenging to construct by combining or manipulating existing datasets. However, the tasks remain easy to understand and use. Since the tasks are language-based, they can be further mixed or chained to create continual learning problems and may also be easily extended in a follow-up work.

We study lesser-known machine learning models based on evolving states of complex systems that can learn through self-organization. A complex dynamical system comprises many interacting agents and evolves over time according to a fixed update rule. These agents can be hidden neurons of a RNN or nodes in a graph. Such systems often exhibit emergent global dynamics resulting from the actions of its parts rather than the decisions of a central controller. These dynamics lead to surprisingly complex behavior, which can be random, chaotic, or may lead to unbounded growth of complexity (Boccaro, 2010). Due to these properties, systems based on reservoir computing may be a promising alternative addressing the shortcomings of standard supervised models.

Surprisingly, such models achieve remarkable learning efficiency compared to the more standard sequential supervised models trained with stochastic gradient-based learning. The complex systems-based models consistently outperform sequential supervised methods and even achieve better learning efficiency than humans on some tasks. They demonstrate more efficient learning on our benchmark at a fraction of the computational and data cost of the conventional models. We believe that more advanced models of this type could lead to more robust and data efficient machine learning in the future, especially in low-data applications or problems where supervision is limited. Complex systems are underexplored and seem worth investigating further for building the next generation of learning algorithms.

ACKNOWLEDGMENTS

This work was partly supported by the European Regional Development Fund under the project IMPACT (reg. no. CZ.02.1.01/0.0/0.0/15 003/0000468).

REFERENCES

- Peter M. Allen and Mark Strathern. Evolution, emergence, and learning in complex systems. *Emergence*, 5(4):8–33, 2003.
- Charles H. Bennett. Logical Depth and Physical Complexity. In Rolf Herken and Rolf Herken (eds.), *The Universal Turing Machine A Half-Century Survey*, volume 2, pp. 207–235. Springer Vienna, Vienna, 1995. ISBN 978-3-211-82637-9 978-3-7091-6597-3. doi: 10.1007/978-3-7091-6597-3_8.
- Nino Boccara. *Modeling Complex Systems*. Springer Science & Business Media, September 2010. ISBN 978-1-4419-6562-2.
- Léon Bottou, Frank E. Curtis, and Jorge Nocedal. Optimization methods for large-scale machine learning. *Siam Review*, 60(2):223–311, 2018.
- Jacob Buckman, Danijar Hafner, George Tucker, Eugene Brevdo, and Honglak Lee. Sample-efficient reinforcement learning with stochastic ensemble value expansion. *Advances in neural information processing systems*, 31, 2018.
- Maxime Chevalier-Boisvert, Dzmitry Bahdanau, Salem Lahlou, Lucas Willems, Chitwan Saharia, Thien Huu Nguyen, and Yoshua Bengio. Babyai: A platform to study the sample efficiency of grounded language learning. *arXiv preprint arXiv:1810.08272*, 2018.
- Kyunghyun Cho, Bart van Merriënboer, Dzmitry Bahdanau, and Yoshua Bengio. On the Properties of Neural Machine Translation: Encoder–Decoder Approaches. In *Proceedings of SSST-8, Eighth Workshop on Syntax, Semantics and Structure in Statistical Translation*, pp. 103–111, 2014.
- Hugo Cisneros, Josef Sivic, and Tomas Mikolov. Evolving Structures in Complex Systems. In *2019 IEEE Symposium Series on Computational Intelligence (SSCI)*, pp. 230–237, Xiamen, China, December 2019. IEEE. ISBN 978-1-72812-485-8. doi: 10.1109/SSCI44817.2019.9002840.
- Jeffrey L. Elman. Finding Structure in Time. *Cognitive Science*, 14(2):179–211, 1990. ISSN 1551-6709. doi: 10.1207/s15516709cog1402_1.
- Alexander Gepperth and Cem Karaoguz. A Bio-Inspired Incremental Learning Architecture for Applied Perceptual Problems. *Cognitive Computation*, 8(5):924–934, October 2016. ISSN 1866-9964. doi: 10.1007/s12559-016-9389-5.
- Tom Eivind Glover, Pedro Lind, Anis Yazidi, Evgeny Osipov, and Stefano Nichele. The Dynamical Landscape of Reservoir Computing with Elementary Cellular Automata. In *ALIFE 2021: The 2021 Conference on Artificial Life*. MIT Press, 2021.
- Jeffrey Goldstein. Emergence in complex systems. *The sage handbook of complexity and management*, pp. 65–78, 2011.
- Ian J. Goodfellow, Mehdi Mirza, Xia Da, Aaron C. Courville, and Yoshua Bengio. An Empirical Investigation of Catastrophic Forgetting in Gradient-Based Neural Networks. In Yoshua Bengio and Yann LeCun (eds.), *2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14-16, 2014, Conference Track Proceedings*, 2014.
- Alex Graves, Greg Wayne, and Ivo Danihelka. Neural turing machines. *arXiv preprint arXiv:1410.5401*, 2014.
- Geoffrey E. Hinton and David C. Plaut. Using fast weights to deblur old memories. In *Proceedings of the Ninth Annual Conference of the Cognitive Science Society*, pp. 177–186, 1987.
- Sepp Hochreiter and Jürgen Schmidhuber. Long Short-Term Memory. *Neural Computation*, 9(8):1735–1780, November 1997. ISSN 0899-7667. doi: 10.1162/neco.1997.9.8.1735.

- Herbert Jaeger. The “echo state” approach to analysing and training recurrent neural networks-with an erratum note. *Bonn, Germany: German National Research Center for Information Technology GMD Technical Report*, 148(34): 13, 2001.
- Herbert Jaeger. Long Short-Term Memory in Echo State Networks: Details of a Simulation Study. Technical report, Jacobs University Bremen, February 2012.
- Armand Joulin and Tomas Mikolov. Inferring algorithmic patterns with stack-augmented recurrent nets. In *Proceedings of the 28th International Conference on Neural Information Processing Systems - Volume 1*, NIPS’15, pp. 190–198, Cambridge, MA, USA, December 2015. MIT Press.
- Satoshi Kanazawa. General intelligence as a domain-specific adaptation. *Psychological review*, 111(2):512, 2004.
- Diederik P. Kingma and Jimmy Ba. Adam: A Method for Stochastic Optimization. In Yoshua Bengio and Yann LeCun (eds.), *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015.
- Moshe Koppel and Henri Atlan. An almost machine-independent theory of program-length complexity, sophistication, and induction. *Information Sciences*, 56(1):23–33, August 1991. ISSN 0020-0255. doi: 10.1016/0020-0255(91)90021-L.
- Alex Krizhevsky and Geoffrey Hinton. Learning multiple layers of features from tiny images. Technical report, University of Toronto, 2009.
- Andrew L. Maas, Raymond E. Daly, Peter T. Pham, Dan Huang, Andrew Y. Ng, and Christopher Potts. Learning Word Vectors for Sentiment Analysis. In Dekang Lin, Yuji Matsumoto, and Rada Mihalcea (eds.), *The 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies, Proceedings of the Conference, 19-24 June, 2011, Portland, Oregon, USA*, pp. 142–150. The Association for Computer Linguistics, 2011.
- Wolfgang Maass, Thomas Natschläger, and Henry Markram. Real-Time Computing Without Stable States: A New Framework for Neural Computation Based on Perturbations. *Neural Computation*, 14(11):2531–2560, November 2002. ISSN 0899-7667, 1530-888X. doi: 10.1162/089976602760407955.
- Marvin Minsky and Seymour A. Papert. *Perceptrons: An Introduction to Computational Geometry*. The MIT Press, Cambridge MA, 2nd edition edition, 1972. ISBN 0-262-63022-2. doi: 10.7551/mitpress/11301.001.0001.
- Andrew Y. Ng, Michael I. Jordan, and Yair Weiss. On spectral clustering: Analysis and an algorithm. In *Proceedings of the 14th International Conference on Neural Information Processing Systems: Natural and Synthetic*, NIPS’01, pp. 849–856, Cambridge, MA, USA, January 2001. MIT Press.
- Cuong V. Nguyen, Yingzhen Li, Thang D. Bui, and Richard E. Turner. Variational Continual Learning. *CoRR*, abs/1710.10628, 2017.
- Sylvestre-Alvise Rebuffi, Alexander Kolesnikov, Georg Sperl, and Christoph H. Lampert. Icarl: Incremental classifier and representation learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 2001–2010, 2017.
- Kyle Richardson, Hai Hu, Lawrence Moss, and Ashish Sabharwal. Probing natural language inference models through semantic fragments. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pp. 8713–8721, 2020.
- Anthony Robins. Catastrophic forgetting in neural networks: The role of rehearsal mechanisms. In *Proceedings 1993 The First New Zealand International Two-Stream Conference on Artificial Neural Networks and Expert Systems*, pp. 65–68. IEEE, 1993.
- David E. Rumelhart, Geoffrey E. Hinton, and Ronald J. Williams. Learning internal representations by error propagation. Technical report, California Univ San Diego La Jolla Inst for Cognitive Science, 1985.
- Rupesh Kumar Srivastava, Jonathan Masci, Sohrob Kazerounian, Faustino J. Gomez, and Jürgen Schmidhuber. Compete to Compute. In Christopher J. C. Burges, Léon Bottou, Zoubin Ghahramani, and Kilian Q. Weinberger (eds.), *Advances in Neural Information Processing Systems 26: 27th Annual Conference on Neural Information Processing Systems 2013. Proceedings of a Meeting Held December 5-8, 2013, Lake Tahoe, Nevada, United States*, pp. 2310–2318, 2013.

- Matthew E. Taylor and Peter Stone. Cross-domain transfer for reinforcement learning. In Zoubin Ghahramani (ed.), *Machine Learning, Proceedings of the Twenty-Fourth International Conference (ICML 2007), Corvallis, Oregon, USA, June 20-24, 2007*, volume 227 of *ACM International Conference Proceeding Series*, pp. 879–886. ACM, 2007. doi: 10.1145/1273496.1273607.
- Matthew E. Taylor, Peter Stone, and Yaxin Liu. Transfer Learning via Inter-Task Mappings for Temporal Difference Learning. *J. Mach. Learn. Res.*, 8:2125–2167, 2007.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention Is All You Need. *arXiv:1706.03762 [cs]*, December 2017.
- Yaqing Wang, Quanming Yao, James T. Kwok, and Lionel M. Ni. Generalizing from a few examples: A survey on few-shot learning. *ACM computing surveys (csur)*, 53(3):1–34, 2020.
- Jason Weston, Antoine Bordes, Sumit Chopra, and Tomás Mikolov. Towards AI-Complete Question Answering: A Set of Prerequisite Toy Tasks. In Yoshua Bengio and Yann LeCun (eds.), *4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto Rico, May 2-4, 2016, Conference Track Proceedings*, 2016.
- Stephen Wolfram. Statistical mechanics of cellular automata. *Reviews of Modern Physics*, 55(3):601–644, July 1983. ISSN 0034-6861. doi: 10.1103/RevModPhys.55.601.
- Stephen Wolfram. *A New Kind of Science*. Wolfram Media, Champaign, IL, 2002. ISBN 978-1-57955-008-0.
- Denis Yarats, Amy Zhang, Ilya Kostrikov, Brandon Amos, Joelle Pineau, and Rob Fergus. Improving sample efficiency in model-free reinforcement learning from images. *arXiv preprint arXiv:1910.01741*, 2019.
- Ozgür Yilmaz. Reservoir Computing using Cellular Automata. *arXiv:1410.0162 [cs]*, October 2014.

APPENDIX

In this appendix we provide details of the human evaluation (Section A), we list all the experimental parameters used in the evaluation (Section B), and describe the reservoir models in more details (Section C).

A HUMAN PERFORMANCE EVALUATION

The proposed benchmark is relatively simple to understand and hence one may ask what would be the human performance and what WADE values this would correspond to. In particular, the language tasks appear to be readily solvable just from one’s understanding of the English language. However, similar to how words are encoded as vectors before being processed by a neural network, we need to remove any language-based prior to make the comparison with human performance fair. This could be achieved in a number of ways. For example, if we map all the available tokens to random letters of the alphabet the apparently trivial task

I DO NOT SEE PAUL . DO I SEE PAUL ? NO
I HEAR JAMES BUT I DO NOT HEAR PAUL AND JOHN . DO I HEAR JAMES ? YES

would become significantly less obvious to humans

w K k A D r K w A D l W
w Z t g w K k Z D G C r K w Z t l H.

One would need to read through several of these sentences to identify patterns such as the fact that W and H represent Yes or No or that r plays the role of a full stop.

We apply the following procedure in our human evaluation experiments:

1. Choose a random task among 10.
2. Apply a random mapping from the task token to letters of the alphabet.
3. Present the sequences one by one with the tokens to be predicted hidden.
4. The user enters his predicted answer.
5. The valid sequence is shown so the user can learn from their mistakes.
6. Step 3-5 are repeated until 10 sequences are correctly solved in a row. Then we go back to step 2 with a new task.

We add the requirement that no external device such as pen and paper or note-taking program should be used during the experiment so the user relies solely on their memory. The accuracy is computed by counting the number of right answers in a row. For example, three right answers in a row would correspond to an accuracy of 30%, while ten right answers in a row is 100% which also corresponds to a switch of a task. This accuracy score is attributed to the first correct answer of the series of correct answers. This means that a series of five correct answers starting from question 6 until question 11 will correspond to reaching the accuracy of 50% on question 6. This ensures that ten correct answers on the first attempt will yield a WADE score of 1.

B EXPERIMENTAL PARAMETERS

We report all the parameters used in our experiments in table 6:

B.1 TASK GENERATION PARAMETERS

The tasks are generated according to the procedures described in Section 3. They have some parameters that allow the difficulty of the task to be adjusted. We report results on ten tasks. The parameters chosen for our experiments are listed below:

Periodic (1) and Increasing period (2): Sequences are generated from patterns of length between 1 and 10.

Easy symbol counting (3): Available symbols are A, B and C. Generated sequences have between 1 and 10 of these symbols in the prompt, and the query has either one, two, or all three symbols.

| | Reservoir-based | Fully trained |
|---|---------------------------------------|--|
| Total number of sequences per task | 1200 | 1200 |
| Number of training sequences | 960 | 960 |
| Number of testing sequences | 240 | 240 |
| Passes over the data (epochs) | 1 | 10 |
| Random runs per task | 100 | 100 |
| Algorithm | SGD | Adam (Kingma & Ba, 2015) |
| Learning rate | 0.001 | 0.001 |
| Regularization | weight decay 0.001 | None |
| Internal state (hidden) size | 1800 | h task-dependent |
| Number of trainable parameters | $1800 \times \text{dictionary_size}$ | $h^2 + 2 \times h \times \text{dictionary_size}$ for the RNN ² |
| Output non-linearity | Softmax | Softmax |
| Internal non-linearity | Not applicable | tanh |
| Batch size | 1 | 1 |

Table 6: Experimental parameters common to all tasks. Some of the sizes (such as h) vary from task to task.

Hard symbol counting (4): Available symbols are the same as for the previous task. Generated sequences have between 1 and 45 symbols with separators in the prompt. There is a minimum of 1 query, and can be as many queries as there are different patterns in the prompt.

Question answering (5): There are five available names and two verbs. The prompt consists of between one and five names, and the question is about one of these names.

Harder question answering : There are eleven available names and five verbs. The prompt is made of between one and five names, and the question is about one of these names.

Question answering with world (7), with counting (8): There are thirteen available names and seven verbs.

Adjective question answering (9), with counting (10): There are eight available names, six verbs, four color adjectives, and five size adjectives. The prompt consists of between one and six statements, and there are up to eight questions.

We chose relatively small values of these parameters because we are interested in simple tasks so that models that can learn from less than 100 examples. Moreover, the number of possible sentences generated from these small values is already huge, with more than 10^{35} possible sentences for task 4, for example. These parameters could also be varied dynamically to change the task’s difficulty when needed or turn each task into multiple sub-tasks or levels of difficulty. We chose the ten separate tasks/settings pairs listed above to get a broad overview of our benchmarked models on a range of task difficulties.

C RESERVOIR MODELS

We describe in more details the two reservoir-based models used in our experiments.

C.1 ECHO-STATE NETWORK

The echo-state network is based on the following update equation:

$$\mathbf{r}_{t+1} = (1 - \alpha)\mathbf{r}_t + \beta \tanh(\mathbf{W} \mathbf{r}_t + \mathbf{W}_{\text{in}} \mathbf{x}_{t+1}), \quad (4)$$

where \mathbf{r}_t is the K -dimensional state vector – corresponding to the hidden neurons — at time t , β is the leaking rate, $\mathbf{W} \in \mathbb{R}^{K \times K}$ is a sparsely connected random hidden layer matrix, and $\mathbf{W}_{\text{in}} \in \mathbb{R}^{L \times K}$ is the input projection matrix. The matrices \mathbf{W} and \mathbf{W}_{in} are initialized randomly using the recommendations of Jaeger (2012): \mathbf{W} has an average of 10 non-zeros entries per row, all sampled uniformly in $[-1, 1]$, \mathbf{W}_{in} has its entries uniformly sampled in $[-1, 1]$.

The L -dimensional outputs are computed at times $t > 0$ as

$$\tilde{\mathbf{x}}_{t+1} = D(\mathbf{r}_t), \quad (5)$$

where $D : \mathbb{R}^K \rightarrow \mathbb{R}^L$ is a (trained) decoding function. For echo-state networks, the decoder is often a linear transformation $D(\mathbf{r}_t) = \mathbf{W}_{\text{out}} \mathbf{r}_t$ where \mathbf{W}_{out} is a $K \times L$ -dimensional matrix. In our experiments, we set $\beta = 0$, which

² h is chosen to match the number of parameters of the reservoir-based methods.

was empirically observed to yield the best results on our tasks. The parameter β , as well as the randomly sampled weight matrix, are sometimes tuned for each task (Jaeger, 2012). We only use default values to get a task-independent setup with the least possible assumptions and to make the methods comparable.

C.2 RESERVOIR CELLULAR AUTOMATA

The Cellular Automaton (CA) update function Φ is applied at time t on the previous state vector s_{t-1} — also called the grid. The input vector x_t at time t is projected onto a vector p_t the same size as s_t using a method proposed in Yilmaz (2014). To combine inputs with CA states, we XOR the two vectors together (Glover et al., 2021),

$$s'_t = p_t \otimes s_t, \quad (6)$$

where s'_t is the temporary CA state resulting from the combination of the CA state and input vectors at time t . We then compute the next CA states by applying the update function Φ ,

$$s_{t+1} = \Phi(s'_t). \quad (7)$$

The final reservoir state r_{t+1} is obtained by concatenating r consecutive CA states obtained from a single combined state s' by applying Φ again. If the size of the state is n , the resulting reservoir vector r has a dimension $K = r \times n$, where r is defined above. As for the echo-state network (ESN), the L -dimensional output tokens are calculated at times $t > 0$ as $\tilde{x}_{t+1} = D(r_t)$ where $D : \mathbb{R}^K \rightarrow \mathbb{R}^L$ is the (trained) decoding function. We use linear decoding functions in our experiments such that $D(r_t) = \mathbf{W}_{\text{out}} r_t$.