

---

# Linear Transformers Are Secretly Fast Weight Programmers

---

Imanol Schlag\*<sup>1</sup> Kazuki Irie\*<sup>1</sup> Jürgen Schmidhuber<sup>1</sup>

## Abstract

We show the formal equivalence of linearised self-attention mechanisms and fast weight controllers from the early '90s, where a “slow” neural net learns by gradient descent to program the “fast weights” of another net through sequences of elementary programming instructions which are additive outer products of self-invented activation patterns (today called keys and values). Such Fast Weight Programmers (FWPs) learn to manipulate the contents of a finite memory and dynamically interact with it. We infer a memory capacity limitation of recent linearised softmax attention variants, and replace the purely additive outer products by a delta rule-like programming instruction, such that the FWP can more easily learn to correct the current mapping from keys to values. The FWP also learns to compute dynamically changing learning rates. We also propose a new kernel function to linearise attention which balances simplicity and effectiveness. We conduct experiments on synthetic retrieval problems as well as standard machine translation and language modelling tasks which demonstrate the benefits of our methods.

## 1. Introduction

Transformers (Vaswani et al., 2017) have achieved impressive results in a myriad of sequence processing tasks, including machine translation, language modelling (Al-Rfou et al., 2019; Dai et al., 2019; Baevski & Auli, 2019; Radford et al., 2019), and question answering (Devlin et al., 2019), domains previously dominated by recurrent neural networks (Graves, 2013; Bahdanau et al., 2015).

The core component of a Transformer is the self-attention mechanism (Cheng et al., 2016; Parikh et al., 2016; Lin et al., 2017) which was recently connected to the modern Hop-

field network (Ramsauer et al., 2021; Krotov & Hopfield, 2016; Demircigil et al., 2017). It extends a form of attention (Bahdanau et al., 2015) originally introduced to complement recurrent neural networks, e.g., (Hochreiter & Schmidhuber, 1997). While relinquishing the recurrence property, all computations across the time axis can be parallelised. However, this comes with drawbacks: self-attention computations scale quadratically with sequence length while the memory of the model grows linearly. Therefore, practitioners are forced to limit the context window to a reasonable size, which in turn makes it impossible to capture longer-term dependencies.

Recent work proposed “linear Transformers” with constant size memory and time complexity linear in sequence length (Katharopoulos et al., 2020; Choromanski et al., 2021; Peng et al., 2021; Shen et al., 2018). This complexity reduction is mainly due to a linearisation of the softmax (reviewed in Sec. 3.2).

Here we emphasize the formal equivalence of this family of linear Transformers and the Fast Weight Controllers or Fast Weight Programmers (FWPs) from the '90s (Schmidhuber, 1991; 1992; 1993; AI Blog, 2021) (apart from normalisation). The memories of such FWPs contain key-value associations, and an FWP can learn to reprogram them through sequences of differentiable elementary instructions (also called update rules), which are additive outer products between keys and values invented by the FWP.

This view allows us to derive a limitation of the memory capacity of linear Transformers and similar models. When the sequence length exceeds storage capacity, the model may end up in an overcapacity regime (discussed in depth in Sec. 4.1). To properly operate under such a regime, the model should learn to dynamically interact with the memory contents and selectively decide which key-value associations to keep and which ones to delete. The purely additive instruction may be inappropriate for this purpose. Therefore, inspired by recent work on FWPs (Schlag et al., 2021), we introduce an improved programming instruction akin to the famous error-correcting delta-rule (Widrow & Hoff, 1960).

Furthermore, softmax linearisation techniques for Transformers are still underexplored. The existing techniques are either very simplistic (Katharopoulos et al., 2020) or mathematically well explained but complex (Choromanski

---

\*Equal contribution <sup>1</sup>The Swiss AI Lab IDSIA, USI & SUPSI. Correspondence to: Imanol Schlag <imanol@idsia.ch>, Kazuki Irie <kazuki@idsia.ch>, Jürgen Schmidhuber <juer-gen@idsia.ch>.

et al., 2021; Peng et al., 2021). We provide a comprehensive comparison and propose a new method which is both simple and effective.

We demonstrate the benefits of the proposed methods on our own synthetic retrieval dataset (Sec. 6.1), the standard WMT14 English to German machine translation task (Sec. 6.2), and the Wikitext-103 (Merity et al., 2017) language modelling task (Sec. 6.3)<sup>2</sup>.

## 2. Background on Fast Weight Programmers

Here we review the concepts of *Fast Weight Programmers* (FWPs) before relating them to linear Transformer variants in Sec. 3.

In standard neural networks, the weights remain fixed after training, unlike the activations, which change depending on the inputs at test time. The general idea of *fast weights* is to make the weights also variable and input-dependent. This concept was called *synaptic modulation* (von der Malsburg, 1981), a method for variable binding in neural networks (see e.g. the recent survey by Greff et al. (2020)), or *dynamic connections* (Feldman, 1982). Von der Malsburg defines the effective weights as a (multiplicative) superposition of conventional, context-independent *slow weights*, and fast changing, context-dependent *fast weights*. Hinton & Plaut (1987) studied a net with (additive) superposition of two sets of weights with two different learning rates in a scenario of model retraining. Before 1991, however, no network learned by gradient descent to quickly compute the changes of the fast weight storage of another network or of itself.

Context-dependent FWPs were introduced in two-network systems of the early '90s (Schmidhuber, 1991; 1992; 1993; AI Blog, 2021). A traditional slow net with slow weights continually changes or reprograms the fast weights of a fast net, making the fast weights effectively dependent on the spatio-temporal context of a given input stream. Simply put, the slow net learns to program its fast net. Among the proposed elementary differentiable instructions that the slow net can use to program the fast weights, a particularly attractive one makes use of outer products (Schmidhuber, 1991; 1992): for a sequential input  $\{\mathbf{x}^{(i)}\}_{i=1}^L$ ,  $\mathbf{x}^{(i)} \in \mathbb{R}^{d_{in}}$ , the model outputs the sequence  $\{\mathbf{y}^{(i)}\}_{i=1}^L$ ,  $\mathbf{y}^{(i)} \in \mathbb{R}^{d_{out}}$  as

$$\mathbf{a}^{(i)}, \mathbf{b}^{(i)} = \mathbf{W}_a \mathbf{x}^{(i)}, \mathbf{W}_b \mathbf{x}^{(i)} \quad (1)$$

$$\mathbf{W}^{(i)} = \sigma(\mathbf{W}^{(i-1)} + \mathbf{a}^{(i)} \otimes \mathbf{b}^{(i)}) \quad (2)$$

$$\mathbf{y}^{(i)} = \mathbf{W}^{(i)} \mathbf{x}^{(i)} \quad (3)$$

where  $\otimes$  denotes the outer product,  $\sigma$  is an activation function,  $\mathbf{W}_a$  and  $\mathbf{W}_b$  are trainable slow weights, while the fast weights  $\mathbf{W}^{(i)}$  are generated at each time step  $i$  and serve

<sup>2</sup>Source code used in this paper is available at [github.com/ischlag/fast-weight-transformers](https://github.com/ischlag/fast-weight-transformers).

as a short-term memory. This is a key-value associative memory model in which the write operation is based on a summation (Eq. 2) and the retrieval is a matrix-vector multiplication (Eq. 3). Schmidhuber (1993) describes a recurrent version and discusses “internal spotlights of attention” (such attention terminology is now widely used in the context of transformers). The use of outer products results in a model of associations similar to tensor product presentations (Smolensky, 1990). In fact, outer-product based associative memory can be found in numerous works since Hebb’s informal rule (Hebb, 1949) and its more concrete formal variants (Steinbuch, 1961; Steinbuch & Piske, 1963; Kohonen, 1972; Palm, 1980) including Hopfield networks (Hopfield, 1982; Little, 1974) and bi-directional associative nets (Kosko, 1988). However, these authors described pre-wired rules to associate given patterns with each other. Their systems did not learn to use such rules for associating self-invented patterns like the FWPs since 1991.

The concept of FWPs has been revisited recently (Ba et al., 2016; Schlag & Schmidhuber, 2017), also under different names, e.g., hypernetworks (Ha et al., 2017; Perez et al., 2018; Galanti & Wolf, 2020), dynamic plasticity (Miconi et al., 2018; 2019), dynamic convolution (Klein et al., 2015; Noh et al., 2016; Jia et al., 2016), or lambda networks (Bello, 2021) used for applications including meta-learning (Munkhdalai & Yu, 2017; Munkhdalai & Trischler, 2018; Munkhdalai et al., 2019; Kirsch & Schmidhuber, 2020). FWPs recently also improved memory models through explicit mechanisms for facilitating the replacement of deprecated information and updating associations (Schlag & Schmidhuber, 2018; Schlag et al., 2021).

## 3. Relation to Transformers

Ba et al. (2016) have already pointed out a relation between a variant of outer product-based FWPs (Schmidhuber, 1993) and attention (Bahdanau et al., 2015). Katharopoulos et al. (2020) have analysed linearised transformers. We review these derivations, emphasising the relation between Transformers and the FWPs of the previous section.

### 3.1. Self-Attention Without Softmax Is a Fast Weight Programmer

A self-attention layer in auto-regressive Transformers (Vaswani et al., 2017) maps an input sequence  $\{\mathbf{x}^{(i)}\}_{i=1}^L$ ,  $\mathbf{x}^{(i)} \in \mathbb{R}^{d_{key} \times 1}$  to an output sequence  $\{\mathbf{y}^{(i)}\}_{i=1}^L$ ,  $\mathbf{y}^{(i)} \in \mathbb{R}^{d_{value} \times 1}$  as

$$\mathbf{k}^{(i)}, \mathbf{v}^{(i)}, \mathbf{q}^{(i)} = \mathbf{W}_k \mathbf{x}^{(i)}, \mathbf{W}_v \mathbf{x}^{(i)}, \mathbf{W}_q \mathbf{x}^{(i)} \quad (4)$$

$$\mathbf{K}^{(i)} = [\mathbf{K}^{(i-1)}, \mathbf{k}^{(i)}] \in \mathbb{R}^{d_{key} \times i} \quad (5)$$

$$\mathbf{V}^{(i)} = [\mathbf{V}^{(i-1)}, \mathbf{v}^{(i)}] \in \mathbb{R}^{d_{value} \times i} \quad (6)$$

$$\mathbf{y}^{(i)} = \mathbf{V}^{(i)} \text{softmax}((\mathbf{K}^{(i)})^\top \mathbf{q}^{(i)}) \quad (7)$$

where  $[\mathbf{A}, \mathbf{a}]$  denotes the concatenation of vector  $\mathbf{a}$  to matrix  $\mathbf{A}$  along the time dimension, softmax is applied along the time dimension, and  $\mathbf{W}_k, \mathbf{W}_v, \mathbf{W}_q$  are trainable weight matrices. We omit the scaling by  $1/\sqrt{d_{\text{key}}}$  inside the softmax without loss of generality.

Now if we remove the softmax in Eq. 7 we obtain:

$$\begin{aligned} \mathbf{y}^{(i)} &= \mathbf{V}^{(i)} ((\mathbf{K}^{(i)})^\top \mathbf{q}^{(i)}) = (\mathbf{V}^{(i)} (\mathbf{K}^{(i)})^\top) \mathbf{q}^{(i)} \\ &= \left( \sum_{j=1}^i \mathbf{v}^{(j)} \otimes \mathbf{k}^{(j)} \right) \mathbf{q}^{(i)} \end{aligned} \quad (8)$$

Denoting by  $\mathbf{W}^{(i)}$  the corresponding weight matrix generated from key and value vectors:

$$\mathbf{W}^{(i)} = \left( \sum_{j=1}^i \mathbf{v}^{(j)} \otimes \mathbf{k}^{(j)} \right) \quad (9)$$

we can rewrite Eqs. 4-7 such that they directly relate to Eqs. 1-3 where the activation function  $\sigma$  is the identity function and without query projection  $\mathbf{W}_q$ :

$$\mathbf{k}^{(i)}, \mathbf{v}^{(i)}, \mathbf{q}^{(i)} = \mathbf{W}_k \mathbf{x}^{(i)}, \mathbf{W}_v \mathbf{x}^{(i)}, \mathbf{W}_q \mathbf{x}^{(i)} \quad (4)$$

$$\mathbf{W}^{(i)} = \mathbf{W}^{(i-1)} + \mathbf{v}^{(i)} \otimes \mathbf{k}^{(i)} \quad (10)$$

$$\mathbf{y}^{(i)} = \mathbf{W}^{(i)} \mathbf{q}^{(i)} \quad (11)$$

### 3.2. Linearising Self-Attention

Instead of removing the softmax as in Sec. 3.1, prior works have introduced techniques for linearising the softmax (Tsai et al., 2019), which has been shown to improve computational efficiency of self-attention for long sequences (Katharopoulos et al., 2020; Choromanski et al., 2021; Peng et al., 2021).

By writing the softmax explicitly, Eq. 7 can be written as:

$$\mathbf{y}^{(i)} = \sum_{j=1}^i \frac{\mathbf{v}^{(j)} \kappa(\mathbf{k}^{(j)}, \mathbf{q}^{(i)})}{\sum_{j'=1}^i \kappa(\mathbf{k}^{(j')}, \mathbf{q}^{(i)})} \quad (12)$$

where  $\kappa(\mathbf{k}, \mathbf{q}) = \exp(\mathbf{k} \cdot \mathbf{q}) \in \mathbb{R}_{>0}$  is the softmax kernel and  $\mathbf{k} \cdot \mathbf{q} = \mathbf{k}^\top \mathbf{q}$  is the vector dot product.

The general idea is to replace the softmax kernel  $\kappa$  by another kernel:  $\kappa'(\mathbf{k}, \mathbf{q}) = \phi(\mathbf{k})^\top \phi(\mathbf{q})$  where  $\phi$  is a function  $\mathbb{R}^{d_{\text{key}}} \rightarrow \mathbb{R}^{d_{\text{dot}}}$ . We discuss the necessary properties of  $\phi$  in Sec. 5.1. By replacing  $\kappa$  in Eq. 12 by  $\kappa'$ , we obtain

$$\mathbf{y}^{(i)} = \sum_{j=1}^i \frac{\mathbf{v}^{(j)} \phi(\mathbf{k}^{(j)})^\top \phi(\mathbf{q}^{(i)})}{\sum_{j'=1}^i \phi(\mathbf{k}^{(j')}) \cdot \phi(\mathbf{q}^{(i)})} \quad (13)$$

$$= \frac{\sum_{j=1}^i (\mathbf{v}^{(j)} \phi(\mathbf{k}^{(j)})^\top) \phi(\mathbf{q}^{(i)})}{\left( \sum_{j'=1}^i \phi(\mathbf{k}^{(j')}) \right) \cdot \phi(\mathbf{q}^{(i)})} \quad (14)$$

Using the outer-product notation, the numerator is analogous to the case without softmax (Sec. 3.1):

$$\sum_{j=1}^i (\mathbf{v}^{(j)} \phi(\mathbf{k}^{(j)})^\top) \phi(\mathbf{q}^{(i)}) = \left( \sum_{j=1}^i \mathbf{v}^{(j)} \otimes \phi(\mathbf{k}^{(j)}) \right) \phi(\mathbf{q}^{(i)})$$

By introducing the fast weight matrix  $\mathbf{W}^{(i)}$  and an additional vector  $\mathbf{z}^{(i)}$  for the denominator,

$$\mathbf{W}^{(i)} = \sum_{j=1}^i \mathbf{v}^{(j)} \otimes \phi(\mathbf{k}^{(j)}) \quad (15)$$

$$\mathbf{z}^{(i)} = \sum_{j=1}^i \phi(\mathbf{k}^{(j)}) \quad (16)$$

forward computations of linear Transformers can be written as (Katharopoulos et al., 2020):

$$\mathbf{k}^{(i)}, \mathbf{v}^{(i)}, \mathbf{q}^{(i)} = \mathbf{W}_k \mathbf{x}^{(i)}, \mathbf{W}_v \mathbf{x}^{(i)}, \mathbf{W}_q \mathbf{x}^{(i)} \quad (4)$$

$$\mathbf{W}^{(i)} = \mathbf{W}^{(i-1)} + \mathbf{v}^{(i)} \otimes \phi(\mathbf{k}^{(i)}) \quad (17)$$

$$\mathbf{z}^{(i)} = \mathbf{z}^{(i-1)} + \phi(\mathbf{k}^{(i)}) \quad (18)$$

$$\mathbf{y}^{(i)} = \frac{1}{\mathbf{z}^{(i)} \cdot \phi(\mathbf{q}^{(i)})} \mathbf{W}^{(i)} \phi(\mathbf{q}^{(i)}) \quad (19)$$

which is a Fast Weight Programmer (Sec. 2) with normalisation. Hence, the core of linear Transformer variants are outer product-based Fast Weight Programmers.

## 4. Analysing and Improving Linear Transformers as Fast Weight Programmers

Viewing linear Transformer variants as Fast Weight Programmers provides us with two insights which we investigate in this work: their capacity limits as associative memories (Sec. 4.1), and their ineptness to edit previously stored associations (Sec. 4.2).

### 4.1. Capacity Limitation

**Intuition.** Endlessly adding new associations to a memory of finite size, as in Eq. 17, inevitably will reach a limit. In linear attention, information is stored in a matrix and is retrieved using matrix multiplication (see Eq. 19). As a consequence, to prevent associations from interfering with each other upon retrieval, the respective keys need to be orthogonal. Otherwise, the dot product will attend to more than one key and return a linear combination of values. With keys embedded in a  $d_{\text{dot}}$  space, there cannot be more than  $d_{\text{dot}}$  orthogonal vectors. That is, storing more than  $d_{\text{dot}}$  associations will result in a retrieval error. In linear Transformers, when the length of the sequence is longer than  $d_{\text{dot}}$ , the model might be in such an overcapacity regime. While we experimentally demonstrate this effect on toy tasks (Sec. 6.1), prior work on tensor product representations allows for a more formal discussion.

**Tensor Product Representation Theory.** Early work in connectionist research investigated the usage of distributed representations as a means for storing symbolic structures. One highly-influential work is the tensor-product-based variable binding mechanism (Smolensky, 1990). A tensor product representation (TPR) of a structured symbolic system consisting of a set of variables and values constructed from outer products of the so called *role* and *filler* vectors. These terms directly translate into *keys* and *values* in our context. The fast weight memories of Eq. 17 are the most basic form of such representations (second order tensors). Therefore, many results discussed in Smolensky’s work transfer to our model. In particular, Theorem 3.3 and 3.1 of Smolensky (1990) discuss more formally the *crossstalk* and retrieval error intuitively described in the previous paragraph.

However, we also note an important difference: the classic TPRs of Smolensky (1990) are constructed with a priori knowledge of the symbolic structure. In contrast, our FWP’s since 1991, including recent FWP’s (Schlag & Schmidhuber, 2018), learn all the vectors involved in constructing such a representation.

## 4.2. Improving the FWP’s Programming Instruction

Sec. 4.1 argues that the linear Transformers can end up in an overcapacity regime, if the sequence length  $L$  exceeds the dimension  $d_{\text{dot}}$  of the keys. Once in overcapacity, an ideal memory model should dynamically interact with the memory contents and selectively determine which associations to remember or to forget. This is in stark contrast to the standard Transformer which stores *immutable* pairs of key and value vectors by concatenation, thus increasing the storage size. While such models work well in practice, we consider a model’s capability to update previously acquired knowledge to be critical for many problems. Hence, from the perspective of dynamic interaction with the memory, the purely additive update rule of Eqs. 17 may be sub-optimal. This motivates us to improve the elementary differentiable programming instruction (i.e. the update rule) of FWP’s.

Inspired by the recent work by Schlag et al. (2021), we propose a basic instruction that essentially implements the famous error-correcting delta rule (Widrow & Hoff, 1960) in an end-to-end differentiable way, such that the FWP can learn to use it wisely, through self-invented, dynamically changing learning rates. Given a new input key-value pair  $(\mathbf{k}^{(i)}, \mathbf{v}^{(i)})$ , the FWP first accesses the current state of the memory  $\mathbf{W}^{(i-1)}$  and retrieves the value  $\bar{\mathbf{v}}^{(i)}$  currently paired with the key  $\mathbf{k}^{(i)}$ . Then the model stores a convex combination  $\mathbf{v}_{\text{new}}^{(i)}$  of the retrieved value  $\bar{\mathbf{v}}^{(i)}$  and the input  $\mathbf{v}^{(i)}$  using an interpolation weight  $0 \leq \beta^{(i)} \leq 1$  also generated by the model. The model thus sequentially transforms an input sequence  $\{\mathbf{x}^{(i)}\}_{i=1}^L, \mathbf{x}^{(i)} \in \mathbb{R}^{d \times 1}$  into an output

sequence  $\{\mathbf{y}^{(i)}\}_{i=1}^L, \mathbf{y}^{(i)} \in \mathbb{R}^{d_{\text{value}} \times 1}$  as:

$$\mathbf{k}^{(i)}, \mathbf{v}^{(i)}, \mathbf{q}^{(i)} = \mathbf{W}_k \mathbf{x}^{(i)}, \mathbf{W}_v \mathbf{x}^{(i)}, \mathbf{W}_q \mathbf{x}^{(i)} \quad (4)$$

$$\bar{\mathbf{v}}^{(i)} = \mathbf{W}^{(i-1)} \phi(\mathbf{k}^{(i)}) \quad (20)$$

$$\beta^{(i)} = \sigma(\mathbf{W}_\beta \mathbf{x}^{(i)}) \quad (21)$$

$$\mathbf{v}_{\text{new}}^{(i)} = \beta^{(i)} \mathbf{v}^{(i)} + (1 - \beta^{(i)}) \bar{\mathbf{v}}^{(i)} \quad (22)$$

where  $\mathbf{W}_\beta \in \mathbb{R}^{1 \times d}$ , and  $\sigma$  is the sigmoid function. The interpolation weight  $\beta^{(i)}$  is the “write-strength” as it defines to which extent the new value will replace the previous value. We note that while  $\beta^{(i)}$  only depends on  $\mathbf{x}^{(i)}$ , in a multi-layer model,  $\mathbf{x}^{(i)}$  has the full context information except in the first layer. We set  $\mathbf{W}^{(0)} = 0$  and  $\mathbf{z}^{(0)} = 0$ . Then the fast weight update rule and the final output  $\mathbf{y}^{(i)}$  are defined as follows (see Appendix A.1 for detailed derivations):

$$\mathbf{W}^{(i)} = \mathbf{W}^{(i-1)} + \underbrace{\mathbf{v}_{\text{new}}^{(i)} \otimes \phi(\mathbf{k}^{(i)})}_{\text{write}} - \underbrace{\bar{\mathbf{v}}^{(i)} \otimes \phi(\mathbf{k}^{(i)})}_{\text{remove}} \quad (23)$$

$$= \mathbf{W}^{(i-1)} + \beta^{(i)} (\mathbf{v}^{(i)} - \bar{\mathbf{v}}^{(i)}) \otimes \phi(\mathbf{k}^{(i)}) \quad (24)$$

$$\mathbf{y}^{(i)} = \mathbf{W}^{(i)} \phi(\mathbf{q}^{(i)}) \quad (25)$$

As shown in Eq. 24, our programming instruction or update rule is effectively a delta rule with a dynamic learning rate  $\beta^{(i)}$ . The model thus learns to correct the current key to value association. In Appendix B, we formally show the advantage of this approach over the gated update rule concurrently proposed by Peng et al. (2021).

**Normalisation.** In the equations above, no normalisation is applied to the value we retrieve. A straightforward normalisation can be obtained by following the derivation in Sec. 3.2, i.e. by introducing an accumulator:

$$\mathbf{z}^{(i)} = \mathbf{z}^{(i-1)} + \phi(\mathbf{k}^{(i)}) \quad (26)$$

and replacing Eqs. 20 and 25 respectively by:

$$\bar{\mathbf{v}}^{(i)} = \frac{\mathbf{W}^{(i-1)} \phi(\mathbf{k}^{(i)})}{\mathbf{z}^{(i-1)} \cdot \phi(\mathbf{k}^{(i)})} \quad (27)$$

$$\mathbf{y}^{(i)} = \frac{\mathbf{W}^{(i)} \phi(\mathbf{q}^{(i)})}{\mathbf{z}^{(i)} \cdot \phi(\mathbf{q}^{(i)})} \quad (28)$$

where we define  $\bar{\mathbf{v}}^{(1)} = 0$ . In this approach, the output  $\mathbf{y}^{(i)}$  is a weighted average of  $\beta^{(j)} (\mathbf{v}^{(j)} - \bar{\mathbf{v}}^{(j)})$  for  $1 \leq j \leq i$ . We refer to this approach as *attention normalisation*.

This approach, however, has drawbacks. First, the accumulation of positive values in Eq. 26 always grows with the number of steps, and may result in instability. Second, specifically for our update rule, this normalisation is not sufficient to balance the *weights* between write and remove operations in Eq. 23 (see derivations in Appendix A.2). Here



we propose a better approach based on simple normalisation. We divide the effective key and query vectors  $\phi(\mathbf{k}^{(i)})$  and  $\phi(\mathbf{q}^{(i)})$  by the sum of its components, e.g., for the query:

$$\phi'(\mathbf{q}^{(i)}) = \frac{\phi(\mathbf{q}^{(i)})}{\sum_{j=1}^{d_{\text{dot}}} \phi(\mathbf{q}^{(i)})_j} \quad (29)$$

before applying Eqs. 20-25. A general consequence of this normalisation is intuitively understood by noticing that the output of any matrix-vector operations (like Eq. 25) is a weighted sum of columns of the matrix where weights are the components of the vector; thus, if the vector components sum up to one, the operation can be viewed as an attention over the columns of the matrix. We provide further explanations and precise implications for our FWP in Appendix A.2. We refer to this approach as *sum normalisation*.

Since this is a simple substitution of  $\phi(\mathbf{k}^{(i)})$  and  $\phi(\mathbf{q}^{(i)})$  in Eqs. 20-25, one might still ask whether additional attention normalisation is needed. In language modelling experiments (Sec. 6.3), we show that this is not the case.

## 5. Linear Attention Functions

The central component of softmax linearisation (Sec. 3.2) is the  $\phi$  function which maps key and query vectors to the space where the dot product is executed:  $\mathbb{R}^{d_{\text{key}}} \rightarrow \mathbb{R}^{d_{\text{dot}}}$ . We first list desirable properties of such a function, and review the existing  $\phi$  functions from the perspective of fast weight memories. Finally, we also propose our own  $\phi$  function.

### 5.1. Properties

For Eq. 13 to define proper attention weights between 0 and 1, the codomain of  $\phi$  should be **positive**. Another property of  $\phi$  derives from the discussion of memory capacity in Sec. 4.1. The dimensionality of its codomain  $d_{\text{dot}}$  defines the model’s capacity. Therefore, by including a transformation which **projects** the input dimension  $d_{\text{key}}$  to a larger dimension  $d_{\text{dot}}$ , the  $\phi$  function can potentially increase the upper bound of the capacity.

### 5.2. Katharopoulos’ Linear Attention

Katharopoulos et al. (2020) propose to use the simple element-wise ELU + 1 function (Clevert et al., 2016):

$$\phi(x) = \text{ELU}(x) + 1 = \begin{cases} x + 1, & \text{if } x > 0 \\ \exp(x), & \text{if } x \leq 0 \end{cases} \quad (30)$$

The choice of ELU over ReLU is motivated by non-zero gradients on the negative part. Importantly, as a simple element-wise function, this  $\phi$  function preserves the dimension of the input key vector ( $d_{\text{key}} = d_{\text{dot}}$ ), without modifying the memory capacity as discussed in Sec. 4.1.

### 5.3. FAVOR+

In contrast to Katharopoulos et al. (2020)’s  $\phi$  function which merely satisfies positivity (and a good gradient) property, Choromanski et al. (2021) propose a mathematically rigorous method to approximate the softmax with random features. They propose the following  $\phi$  function:

$$h(\mathbf{x}) = \frac{1}{\sqrt{2}} \exp\left(-\frac{\|\mathbf{x}\|^2}{2}\right) \quad (31)$$

$$\phi(\mathbf{x}) = \frac{h(\mathbf{x})}{\sqrt{m}} \begin{bmatrix} \exp(\mathbf{R}\mathbf{x}) \\ \exp(-\mathbf{R}\mathbf{x}) \end{bmatrix} \quad (32)$$

where the concatenation  $\begin{bmatrix} \mathbf{a} \\ \mathbf{b} \end{bmatrix}$  of two vectors  $\mathbf{a}$  and  $\mathbf{b}$  is along the feature dimension, and  $\mathbf{R} \in \mathbb{R}^{m \times d_{\text{key}}}$  is a matrix with  $m$  random features where each row vector  $\mathbf{r} \in \mathbb{R}^{1 \times d_{\text{key}}}$  is drawn from  $\mathcal{N}(0, \mathbf{I}_{d_{\text{key}}})$ . A similar approach is also proposed by Peng et al. (2021).

With FAVOR+, the dimension of the codomain  $d_{\text{dot}}$  is  $2m$  which increases the theoretical capacity of the memory if  $2m > d_{\text{key}}$ . At the same time, the model’s capacity is still limited, and equals the infinite capacity of the softmax memory only when  $m$  goes to infinity, which is never achieved in practice. During training, we redraw these  $m$  random vectors for each mini-batch. During evaluation, we draw a set of  $m$  random vectors once, and keep them fixed.  $m$  is the only hyperparameter of FAVOR+ and influences the quality of the softmax approximation. Choromanski et al. (2021) suggest to choose  $m$  in the order of  $d_{\text{key}} \log(d_{\text{key}})$ . This sampling process is the main drawback of FAVOR+ as it introduces variance into the model’s output.

### 5.4. Deterministic Parameter-Free Projection (DPFP)

The two previous sub-sections highlight the sub-optimality of the existing  $\phi$  functions. Sampling introduces extra complexity to FAVOR+ (Sec. 5.3), while the Linear Transformer (Sec. 5.2) lacks the ability to project up the dot product dimension. Here we propose an alternative approach called *deterministic parameter-free projection* (DPFP). It is deterministic and easy to compute like Linear Transformers while increasing the dot product dimension without requiring FAVOR+’s random features.

We begin with a low-dimensional example to foster an intuitive understanding before moving on to the general formulation. Consider 4 keys  $\mathbf{k}^{(i)}, i \in \{1, 2, 3, 4\}$  in  $\mathbb{R}^2$  and  $\phi : \mathbb{R}^2 \rightarrow \mathbb{R}_{\geq 0}^4$  where the  $l$ -th element of  $\phi(\mathbf{x})$  is generated by the partial function  $\phi_l : \mathbb{R}^2 \rightarrow \mathbb{R}_{\geq 0}$ . We design  $\phi$  such that it facilitates orthogonality in the projected space, i.e.  $\phi(\mathbf{k}^{(i)}) \cdot \phi(\mathbf{k}^{(j)}) = 0$  for  $i \neq j$ . Towards this end, we construct  $\phi$  such that if  $\phi_l(\mathbf{x}) > 0$  then  $\phi_n(\mathbf{x}) = 0$  for all  $n \neq l$ . Such a constraint can be enforced by limiting the domains of the partial functions to be non-overlapping.

With the element-wise rectifier function  $r(a) = \max(0, a)$  the partial functions are defined as:

$$\phi_1(\mathbf{k}) = r(\mathbf{k}_1)r(\mathbf{k}_2) \quad (33)$$

$$\phi_2(\mathbf{k}) = r(-\mathbf{k}_1)r(\mathbf{k}_2) \quad (34)$$

$$\phi_3(\mathbf{k}) = r(\mathbf{k}_1)r(-\mathbf{k}_2) \quad (35)$$

$$\phi_4(\mathbf{k}) = r(-\mathbf{k}_1)r(-\mathbf{k}_2) \quad (36)$$

Figure 1 illustrates this function. The elements of the 4-dimensional space are displayed as the  $z$  component of the four coloured surfaces. The figure shows how each vector in the 2d plane will have a single non-zero component in the 4d space and equally splits the input space into four areas which will be orthogonal in the projected space.




Figure 1. A visualisation of a DPFP from a 2d space (the  $xy$ -plane) to a 4d space (the four colored surfaces). Each surface is a partial function which represents one element of the 4d vector.

We generalise this method to higher dimensional inputs by constructing additional two-factor features. Given an input vector  $\mathbf{k} \in \mathbb{R}^{d_{\text{key}}}$  and  $i \in [1, 2d_{\text{key}}]$ , the partial function

$$\phi_{i\nu}(\mathbf{k}) = r\left(\begin{bmatrix} \mathbf{k} \\ -\mathbf{k} \end{bmatrix}\right)_i r\left(\begin{bmatrix} \mathbf{k} \\ -\mathbf{k} \end{bmatrix}\right)_{i+\nu} \quad (37)$$

where  $\nu \in \{1, 2, \dots, d_{\text{key}}2 - 1\}$  is a capacity controlling hyperparameter. The codomain dimensionality of  $\phi(\mathbf{k})$  is thus  $d_{\text{dot}} = 2d_{\text{key}}\nu$ . Eq. 37 is highly parallelisable because each partial function can be computed independently. This can be implemented in few lines of code as we show in Appendix C.

Finally we note that Choromanski et al. (2021) empirically show that replacing  $\exp$  in Eq. 32 by ReLU typically improves model performance. While this result has not been theoretically justified, it supports the design of our DPFP which aims for sparsity and orthogonality.

## 6. Experimental Results

Now we present our experimental results on synthetic retrieval problems (Sec. 6.1.1 and 6.1.2), machine translation (Sec. 6.2), and language modelling (Sec. 6.3).

### 6.1. Synthetic Settings

We illustrate the capacity issue (Sec. 4.1) of linear attention and the effectiveness of our new update rule (Sec. 4.2) on two synthetic problems.

In both settings, our toy problem consists of retrieving the correct value from a sequence of randomly sampled key-value associations when queried with one of the used keys. Crucially, the query is given at the end of the sequence, such that the model is not aware of it while processing the inputs. To succeed, the model has to learn to store the observed associations in its memory without interference.

Let  $\mathcal{K}$  and  $\mathcal{V}$  be the finite and fixed sets of keys and values and  $S = |\mathcal{K}| = |\mathcal{V}|$ . Then, the input to the model is the sequence  $[(k, v)_1, \dots, (k, v)_L]$  followed by  $q$  where every pair  $(k, v) \in \mathcal{K} \times \mathcal{V}$  is sampled randomly, and  $q$  is randomly chosen to be one of the  $L$  keys.

Each value  $v^{(i)}, i \in [1, \dots, S]$  is assigned a fixed one-hot vector  $\mathbf{v}^{(i)} \in \mathbb{R}^S$ . Hence, the set of value vectors is an orthonormal basis. In contrast, the vector embedding of the key symbols is the learned function  $e: \mathcal{K} \rightarrow \mathbb{R}^{d_{\text{emb}}}$  and  $\mathbf{k} = \mathbf{W}_K[e(k); \mathbf{v}]$  where  $\mathbf{W}_K \in \mathbb{R}^{d_{\text{key}} \times (d_{\text{emb}} + S)}$ .

Following the  $L$  write operations, the read function and the query vector  $\mathbf{q} = \mathbf{W}_Q e(q)$ ,  $\mathbf{W}_Q \in \mathbb{R}^{d_{\text{key}} \times d_{\text{emb}}}$  are used to retrieve  $\hat{\mathbf{v}} \in \mathbb{R}^S$  from memory. Finally, the loss is defined as  $l(\hat{\mathbf{v}}, \mathbf{v}^*) = \sum_j^S \frac{1}{2}(\mathbf{v}_j^* - \hat{\mathbf{v}}_j)^2$  where  $\mathbf{v}^*$  is the value vector assigned to  $q$  in the input sequence. Each model is trained in mini-batches using this loss and Adam with default hyperparameters unless stated otherwise. For evaluation, we sample 20 sequences and test all possible queries, e.g., with  $S = 100$  unique keys, the evaluation batch is of size  $100 * 20 = 2000$ .

#### 6.1.1. SETTING 1: TESTING CAPACITY

In this setting, we experimentally demonstrate the capacity limit of linear attention (Sec. 4.1). We conduct experiments for the various  $\phi$  functions described in Sec. 5. We fix  $d_{\text{key}}$  to be 64, while different  $\phi$  functions produce different  $d_{\text{dot}}$ . We set the sequence length to be equal to the number of unique keys ( $L = S$ ), and sample the keys and values without replacement to generate the sequences. By varying the sequence length  $S$ , our goal is to show that all linear attention models (using the simple sum update rule of Sec. 3.2) fail at retrieving when  $S$  exceeds  $d_{\text{dot}}$ .

All models are trained with a mini-batch size of 32 until the evaluation loss falls below 0.001 or until lack of progress for 1000 steps. In Figure 2, the best validation set performance for each model and each  $S$  is displayed (for the learning curves see Appendix D.1). The number of unique keys is initially  $S = 20$  and is incremented by 20 until  $S = 600$ . The following models are compared: Softmax, Linear-

Attention, FAVOR+ with 64, 128, and 512 random features, DPFP- $\nu$  with  $\nu \in \{1, 2, 3\}$ .




Figure 2. Final evaluation loss of the softmax memory and various linear attention mechanisms on associative retrieval problems with the total number of unique associations ranging from 20 to 600. Each individual symbol is a model trained until convergence.

The results support our theoretical analysis. Linear-Attention has a capacity of 64 due to the choice of  $d_{key} = d_{dot} = 64$ . Experimentally, Linear-Attention begins to accumulate errors with 60 or more associations. Similarly, DPFP projections 1, 2 and 3 start to accumulate errors as they approach their respective limits at 128, 256, and 384. FAVOR+, on the other hand, fails to achieve a loss of 0 in any experiment. Finally, as expected, softmax attention is outperforming all  $\phi$  functions, although it struggles to fully converge with more than 500 keys.

6.1.2. SETTING 2: COMPARING UPDATE RULES

In the second setting, we compare variations of the update rule. Unlike in setting 1, keys and values will be sampled with replacement and sequence length  $L = 2S$ . As a result, in the same sequence, multiple keys can be re-assigned to a new value more than once. The expected value to retrieve is the most recent one associated with the query. With every new key, the previous value associated with this key deprecates and the model is required to update its finite size memory. The ability to update values associated with keys is essential to bind context-specific values to a key.

We use DPFP-1 as the  $\phi$  function. The sequence length is fixed at 40 with 20 unique keys and values. While this setting does not exceed the capacity of DPFP-1, our result is independent of the capacity regime (see results for different  $S$  and  $\phi$  in Appendix D.2).

We compare the proposed fast weight memory programming instruction with normalisation of Sec. 4.2 (denoted here by *ours*) to three baselines: the sum update rule of Sec. 3 (*sum rule*), and two variants of previous update rules (Schlag et al., 2021): *Schlag (2021)* and *Schlag (2021) with DPFP*.

*Schlag (2021)* is simply the model from Schlag et al. (2021) ported to this setting (i.e. without the LSTM layer). *Schlag (2021)* has neither a  $\phi$  function, nor the sum normalisation term of Sec. 4.2. Instead it uses a tanh nonlinearity for its key representations. As an ablation we replace it with our DPFP-1 but we don't use the normalisation term of Sec. 4.2, which we refer to as *Schlag (2021) with DPFP*.

Figure 3 presents the learning curves. They demonstrate that our new update rule outperforms all other variants. As expected, the baseline sum update rule fails.

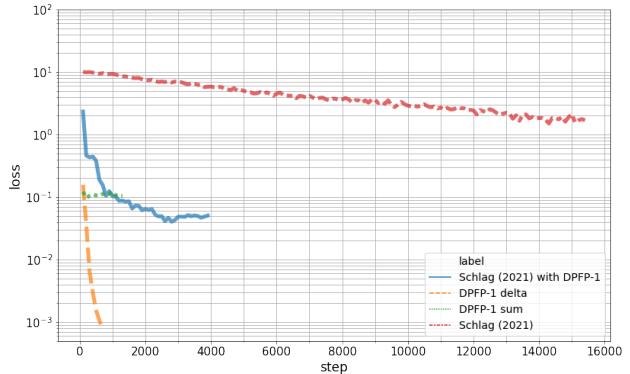


Figure 3. Learning curves for different update rules. Sequence length of 40 and 20 unique keys/values sampled with replacement.

6.2. Machine Translation Experiments

Here we compare  $\phi$  functions on the standard machine translation task. We compare Linear Transformer (Katharopoulos et al., 2020), Performer (Choromanski et al., 2021) and our  $\phi$  function DPFP (Sec. 5.4) to the regular Transformer, complementing prior comparisons, e.g., Tay et al. (2021).

We use the standard WMT14 English to German Translation dataset and standard data setups (Ott et al., 2018; Vaswani et al., 2017). We adapt the recipe of Ott et al. (2019) (see Appendix E) and train Vaswani et al. (2017)'s "big" models for about 4 days on three V100 GPUs. We use the exact same training configurations for all models without model-specific hyper-parameter tuning. We only vary the model hyper-parameters  $m$  in Performers and  $\nu$  in DPFP models.

Table 1 shows the BLEU score (Papineni et al., 2002; Post, 2018) results. The Performer is as good as the basic Transformer when the number of samples  $m$  is large enough (for  $d_{dot} = 512$ , we have  $m = 256$ ). In fact, with  $d_{key} = 64$ , the recommended value for  $m$  is  $d_{dot} \log(d_{dot}) = 266$ . Our DPFP model outperforms the Linear Transformer as well as the Performer when  $d_{dot}$  is relatively small; providing a good trade-off between simplicity and performance.

Table 1. WMT14 En-De Translation BLEU scores for various Transformer models. Neither model averaging, nor model specific tuning is done. *Standard* denotes the basic Transformer.

$d_{\text{dot}}$	Valid			Test		
	64	256	512	64	256	512
Standard	26.6	-	-	<b>27.7</b>	-	-
Linear	25.5	-	-	26.8	-	-
Performer	24.2	24.9	<b>26.7</b>	24.4	25.3	<b>27.7</b>
DPFP (ours)	-	26.2	26.2	-	26.9	27.1

### 6.3. Language Modelling Experiments

Toy experimental Setting 2 (Sec. 6.1.2) illustrated the effect of our update rule. Now our goal is to confirm its effectiveness on a large-vocabulary word-level language modelling task, and investigate its further potential.

**Experimental setups.** Our update rule should be evaluated on a dataset with sufficiently long contextual dependencies. We use the standard WikiText-103 (Merity et al., 2017) dataset. WikiText-103 consists of long articles from Wikipedia; the training set contains about 28 K articles with a total of 103 M running words. This results in contextual text blocks of about 3600 words. The validation and test sets also contain similarly long dependencies, respectively with 218 K and 246 K running words for 60 articles each. The vocabulary size is about 268 K words.

We split the training data into  $L$ -word long segments (which is the backpropagation span). Unless stated otherwise, we treat these segments independently during training. For evaluation, we use a batch size of one, and go through the text with a sliding window of size  $L$ , taking into account only the last position for computing perplexity (except in the first segment where all positions are evaluated). This is usually done for Transformers with a limited context (Al-Rfou et al., 2019). Appendix F provides further experimental details.

**Effectiveness of our new update rule.** We first evaluate our update rule in two configurations. In the *small* configuration, we set the model dimension (same for key, value, and query)  $D$  to 128, and the training and evaluation context length  $L$  to 256. We note that  $D = H * d_{\text{dot}}$  where  $H$  is the number of heads.  $H$  is set to 8. The feed-forward layer dimension is 2048. The number of layers is 16 in all configurations. In the *medium* configuration, we set  $D = 256$  and  $L = 384$ . Both configurations represent an overcapacity regime. We evaluate both Linear Transformers (Katharopoulos et al., 2020) and Performers (Choromanski et al., 2021). However, to keep the comparison simple, we set the capacity of Performers (Sec. 5.3) equal to the one of linear Transformers, by the right choice of projection dimension ( $m = 8$  and  $m = 16$ , respectively, in *small* and *medium*

Table 2. WikiText-103 language model perplexity results showing effects of our update rule. The number of trainable parameters are almost the same for all models, up to the small difference introduced by gating in our update rule (16 K and 33 K parameters respectively for the *small* and *medium* configurations). We have  $D = 128$ ,  $L = 256$  (40 M parameters) in the *small*, and  $D = 256$ ,  $L = 384$  (90 M parameters) in the *medium* configuration. For Performers,  $m$  is 8 and 16, respectively.

	Update Rule	<i>small</i>		<i>medium</i>	
		Valid	Test	Valid	Test
Transformer	-	33.0	34.1	27.9	29.6
Linear Transformer	sum	37.1	38.3	31.1	33.0
Delta Network	delta	<b>34.1</b>	<b>35.5</b>	<b>29.7</b>	<b>31.5</b>
Performer	sum	39.0	39.6	32.2	33.8
	delta	<b>36.1</b>	<b>37.2</b>	<b>30.0</b>	<b>31.8</b>

configurations), even though this limits performance. We do not include DPFP here, since in both configurations even the smallest value for  $\nu$  provides enough capacity. Here we investigate the effect of the update rule in an overcapacity scenario (see Appendix D.3 for experimental results in a non-overcapacity regime including DPFP). All models can be trained using two V100 GPUs in less than four days. We refer to the Linear Transformer with our delta update rule as a *Delta Network*. Table 2 shows the perplexity results. In both configurations, our update rule provides convincing improvements over the models with the sum update rule.

We also conduct an ablation study to test the effect of the absolute positional encoding and an extra attention normalisation (Sec. 4.2). Table 3 shows the results. The *sum normalisation* (Sec. 4.2) is used in all cases: the models diverged otherwise. In contrast, better perplexities are obtained when no additional attention normalisation is applied. We also observe that the absolute positional encoding is not needed, confirming results of prior work (Irie et al., 2019a).

Table 3. WikiText-103 language model perplexities for Linear Transformers (*medium* configuration) with our update rule.

Position Encoding	Attn. Normalisation	Valid	Test
Yes	Yes	30.4	32.1
<b>No</b>	Yes	29.2	31.2
Yes	<b>No</b>	29.7	31.5
<b>No</b>	<b>No</b>	<b>28.1</b>	<b>31.1</b>

**Complexity, wall clock time, memory.** All methods we propose are within the framework of “linear Transformers”. Thus, there is no change to be discussed in terms of complexity which is constant in space and linear in time w.r.t. sequence length. However, our modified update rule in-



Table 4. WikiText-103 language model perplexities when the model is trained and evaluated **without truncating context**, as opposed to Table 2 where the context window is limited. The *medium* config is used. Neither positional encoding nor attention normalisation is used for the Delta Net. The numbers of trainable parameters (Prms.) are given in millions. We compare with the Transformer-XL at different memory segment lengths. This results in different state sizes which are proportional to the memory requirements during evaluation, and highlights the memory efficiency of the Delta Network. The state sizes are given in millions.

Model	Prms. in M.	State size in M.	Perplexity	
			Valid	Test
Linear Transformer	89.8	0.13	>260	>260
Delta Network	89.9	0.13	27.8	29.4
Transformer-XL	90.9	0.13	65.7	65.5
		1.05	29.3	30.1
		2.10	26.4	27.4
		6.29	24.6	25.5

roduces a few extra computations. The wall clock time and memory requirement (for the small LM setting) for the Linear Transformer with and without our delta update rule are: 63 K and 66 K words/sec, and 14 and 13 GB respectively in our implementation. The extra resource requirement is thus marginal. As we use custom CUDA kernels for these linear Transformers, they are faster than the regular Transformers implemented in PyTorch which process 33K words/sec and require 17 GB memory. The speed of the DFPF and Performer models (for Table 5 in Appendix with a larger  $d_{dot}$ ) are 63 K and 57 K words/sec. Performers are slower because of the sampling logic, which also motivates our DFPF.

**Without truncating context.** Given the constant space requirements, we can feed inputs to linear Transformers for an arbitrary number of steps. To properly assess the model’s ability to process arbitrary long sequences, it is crucial to make the training consistent with the evaluation mode (Irie et al., 2019b). During training, we carry over the fast weight memory from one training segment to the following one, while still limiting the backpropagation span to be within the segment. We train a Delta Net, using neither positional encoding nor attention normalisation (the best setting from Table 3). It was crucial to remove the attention normalisation for the Delta Net since the accumulator blows up as indicated in Sec. 4.2, while for the Linear Transformer, removing it resulted in an even worse perplexity of over 1600. Table 4 shows the corresponding results. The Delta Net yields a slight improvement over the best model with a limited context window (Table 3), unlike the baseline Linear Transformer model with the naive sum update rule which breaks. We also train a Transformer-XL in our *medium* configuration as a baseline model specifically designed for

this use case (Dai et al., 2019; Rae et al., 2020). We evaluate it using different state sizes by changing the Transformer XL’s memory and target segment lengths (see Appendix F for further details). Performance of the Delta Net does not yet match the performance of the Transformer XL when the latter is evaluated with a large state size (large attention window). However, when we take the state size into account (Table 4), we observe that the Delta Net performs very well with a small state size, which is a crucial property in some practical applications (Irie et al., 2020). These results are promising for future work on alternative Transformer models which can run for an unlimited number of steps.

### 7. Conclusion

We emphasise the connection between linearised self-attention and Fast Weight Programmers (FWPs, 1991) that program their fast weight memories through sequences of outer products between self-invented key and value patterns. The FWP perspective allows for discussing associative memory capacity limitations of linear attention, and for introducing an alternative differentiable elementary programming instruction that the FWP can use to dynamically edit the memory, akin to the famous delta rule, but such that the FWP can learn to use the rule wisely through gradient descent. We also propose and discuss a new method for linearising attention. Experiments on synthetic and real language tasks demonstrate the effectiveness of our proposals. The FWP perspective opens up new avenues for investigating even better programming instructions and designs for Transformers with finite memory.

### Acknowledgements

We thank Sjoerd van Steenkiste, Hubert Ramsauer and Sepp Hochreiter for valuable comments and suggestions on the first version of the manuscript. This research was partially funded by ERC Advanced grant no: 742870, project ALGO-RNN, and by Swiss National Science Foundation grant no: 200021\_192356, project NEUSYM. We thank NVIDIA Corporation for donating several DGX machines, and IBM for donating a Minsky machine. We also thank Katharopoulos et al. (2020) for releasing their CUDA implementation of Linear Transformers, which was helpful to implement our models.

### References

Al-Rfou, R., Choe, D., Constant, N., Guo, M., and Jones, L. Character-level language modeling with deeper self-attention. In *Proc. Conference on Artificial Intelligence (AAAI)*, pp. 3159–3166, Honolulu, HI, USA, January 2019.

- Ba, J., Hinton, G. E., Mnih, V., Leibo, J. Z., and Ionescu, C. Using fast weights to attend to the recent past. In *Proc. Advances in Neural Information Processing Systems (NIPS)*, pp. 4331–4339, Barcelona, Spain, December 2016.
- Baevski, A. and Auli, M. Adaptive input representations for neural language modeling. In *Int. Conf. on Learning Representations (ICLR)*, New Orleans, LA, USA, May 2019.
- Bahdanau, D., Cho, K., and Bengio, Y. Neural machine translation by jointly learning to align and translate. In *Int. Conf. on Learning Representations (ICLR)*, San Diego, CA, USA, May 2015.
- Bello, I. Lambdanetworks: Modeling long-range interactions without attention. In *Int. Conf. on Learning Representations (ICLR)*, Virtual only, May 2021.
- Cheng, J., Dong, L., and Lapata, M. Long short-term memory-networks for machine reading. In *Proc. Conf. on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 551–561, Austin, TX, USA, November 2016.
- Choromanski, K., Likhoshesterov, V., Dohan, D., Song, X., Gane, A., Sarlos, T., Hawkins, P., Davis, J., Mohiuddin, A., Kaiser, L., et al. Rethinking attention with performers. In *Int. Conf. on Learning Representations (ICLR)*, Virtual only, 2021.
- Clevert, D.-A., Unterthiner, T., and Hochreiter, S. Fast and accurate deep network learning by exponential linear units (ELUs). In *Int. Conf. on Learning Representations (ICLR)*, San Juan, Puerto Rico, May 2016.
- Dai, Z., Yang, Z., Yang, Y., Cohen, W. W., Carbonell, J., Le, Q. V., and Salakhutdinov, R. Transformer-XL: Attentive language models beyond a fixed-length context. In *Proc. Association for Computational Linguistics (ACL)*, pp. 2978–2988, Florence, Italy, July 2019.
- Demircigil, M., Heusel, J., Löwe, M., Uppgang, S., and Vermet, F. On a model of associative memory with huge storage capacity. *Journal of Statistical Physics*, 168(2): 288–299, 2017.
- Devlin, J., Chang, M., Lee, K., and Toutanova, K. BERT: pre-training of deep bidirectional transformers for language understanding. In *Proc. North American Chapter of the Association for Computational Linguistics on Human Language Technologies (NAACL-HLT)*, pp. 4171–4186, Minneapolis, MN, USA, June 2019.
- Feldman, J. A. Dynamic connections in neural networks. *Biological cybernetics*, 46(1):27–39, 1982.
- Galanti, T. and Wolf, L. On the modularity of hypernetworks. In *Proc. Advances in Neural Information Processing Systems (NeurIPS)*, Virtual only, 2020.
- Graves, A. Generating sequences with recurrent neural networks. *Preprint arXiv:1308.0850*, 2013.
- Greff, K., van Steenkiste, S., and Schmidhuber, J. On the binding problem in artificial neural networks. *Preprint arXiv:2012.05208*, 2020.
- Ha, D., Dai, A., and Le, Q. V. Hypernetworks. In *Int. Conf. on Learning Representations (ICLR)*, Toulon, France, April 2017.
- Hanson, S. J. A stochastic version of the delta rule. *Physica D: Nonlinear Phenomena*, 42(1-3):265–272, 1990.
- Hebb, D. O. The organization of behavior: a neuropsychological theory. *A Wiley Book in Clinical Psychology*, 62:78, 1949.
- Hinton, G. E. and Plaut, D. C. Using fast weights to deblur old memories. In *Proc. Conf. of Cognitive Science Society*, pp. 177–186, Seattle, WA, USA, July 1987.
- Hochreiter, S. and Schmidhuber, J. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
- Hopfield, J. J. Neural networks and physical systems with emergent collective computational abilities. *Proc. of the national academy of sciences*, 79(8):2554–2558, 1982.
- Irie, K., Zeyer, A., Schlüter, R., and Ney, H. Language modeling with deep Transformers. In *Proc. Interspeech*, pp. 3905–3909, Graz, Austria, September 2019a.
- Irie, K., Zeyer, A., Schlüter, R., and Ney, H. Training language models for long-span cross-sentence evaluation. In *Proc. IEEE Automatic Speech Recog. and Understanding Workshop (ASRU)*, pp. 419–426, Sentosa, Singapore, December 2019b.
- Irie, K., Gerstenberger, A., Schlüter, R., and Ney, H. How much self-attention do we need? Trading attention for feed-forward layers. In *Proc. IEEE Int. Conf. on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 6154–6158, Virtual only, May 2020.
- Jia, X., De Brabandere, B., Tuytelaars, T., and Gool, L. V. Dynamic filter networks. In *Proc. Advances in Neural Information Processing Systems (NIPS)*, pp. 667–675, Barcelona, Spain, 2016.
- Katharopoulos, A., Vyas, A., Pappas, N., and Fleuret, F. Transformers are rnns: Fast autoregressive transformers with linear attention. In *Proc. Int. Conf. on Machine Learning (ICML)*, Virtual only, July 2020.

- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. *Preprint arXiv:1412.6980*, 2014.
- Kirsch, L. and Schmidhuber, J. Meta learning backpropagation and improving it. *NeurIPS Workshop on Meta-Learning*, 2020.
- Klein, B., Wolf, L., and Afek, Y. A dynamic convolutional layer for short rangeweather prediction. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pp. 4840–4848, Boston, MA, USA, June 2015.
- Kohonen, T. Correlation matrix memories. *IEEE Transactions on Computers*, 21(4):353–359, 1972.
- Kosko, B. Bidirectional associative memories. *IEEE Transactions on Systems, Man, and Cybernetics*, 18(1):49–60, 1988.
- Krotov, D. and Hopfield, J. J. Dense associative memory for pattern recognition. In *Proc. Advances in Neural Information Processing Systems (NIPS)*, pp. 1172–1180, Barcelona, Spain, December 2016.
- Lin, Z., Feng, M., Santos, C. N. d., Yu, M., Xiang, B., Zhou, B., and Bengio, Y. A structured self-attentive sentence embedding. In *Int. Conf. on Learning Representations (ICLR)*, Toulon, France, April 2017.
- Little, W. A. The existence of persistent states in the brain. *Mathematical biosciences*, 19(1-2):101–120, 1974.
- Merity, S., Xiong, C., Bradbury, J., and Socher, R. Pointer sentinel mixture models. In *Int. Conf. on Learning Representations (ICLR)*, Toulon, France, April 2017.
- Miconi, T., Stanley, K., and Clune, J. Differentiable plasticity: training plastic neural networks with backpropagation. In *Proc. Int. Conf. on Machine Learning (ICML)*, pp. 3559–3568, Stockholm, Sweden, July 2018.
- Miconi, T., Rawal, A., Clune, J., and Stanley, K. O. Backpropamine: training self-modifying neural networks with differentiable neuromodulated plasticity. In *Int. Conf. on Learning Representations (ICLR)*, New Orleans, LA, USA, May 2019.
- Munkhdalai, T. and Trischler, A. Metalearning with hebbian fast weights. *Preprint arXiv:1807.05076*, 2018.
- Munkhdalai, T. and Yu, H. Meta networks. In *Proc. Int. Conf. on Machine Learning (ICML)*, pp. 2554–2563, Sydney, Australia, August 2017.
- Munkhdalai, T., Sordoni, A., Wang, T., and Trischler, A. Metalearned neural memory. In *Proc. Advances in Neural Information Processing Systems (NeurIPS)*, pp. 13310–13321, Vancouver, Canada, December 2019.
- Noh, H., Seo, P. H., and Han, B. Image question answering using convolutional neural network with dynamic parameter prediction. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pp. 30–38, Las Vegas, NV, USA, 2016.
- Ott, M., Edunov, S., Grangier, D., and Auli, M. Scaling neural machine translation. In *Proc. Conf. on Machine Translation: Research Papers*, pp. 1–9, Brussels, Belgium, October 2018.
- Ott, M., Edunov, S., Baevski, A., Fan, A., Gross, S., Ng, N., Grangier, D., and Auli, M. fairseq: A fast, extensible toolkit for sequence modeling. In *Proc. North American Chapter of the Association for Computational Linguistics on Human Language Technologies (NAACL-HLT), Demonstrations*, pp. 48–53, Minneapolis, MN, USA, June 2019.
- Palm, G. On associative memory. *Biological cybernetics*, 36(1):19–31, 1980.
- Papineni, K., Roukos, S., Ward, T., and Zhu, W.-J. Bleu: a method for automatic evaluation of machine translation. In *Proc. Association for Computational Linguistics (ACL)*, pp. 311–318, Philadelphia, PA, USA, July 2002.
- Parikh, A. P., Täckström, O., Das, D., and Uszkoreit, J. A decomposable attention model for natural language inference. In *Proc. Conf. on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 2249–2255, Austin, TX, USA, November 2016.
- Paszke, A. et al. Pytorch: An imperative style, high-performance deep learning library. In *Proc. Advances in Neural Information Processing Systems (NeurIPS)*, pp. 8026–8037, Vancouver, Canada, December 2019.
- Peng, H., Pappas, N., Yogatama, D., Schwartz, R., Smith, N. A., and Kong, L. Random feature attention. In *Int. Conf. on Learning Representations (ICLR)*, Virtual only, 2021.
- Perez, E., Strub, F., De Vries, H., Dumoulin, V., and Courville, A. FiLM: Visual reasoning with a general conditioning layer. In *Proc. AAAI Conf. on Artificial Intelligence*, pp. 3942–3951, New Orleans, LA, USA, February 2018.
- Post, M. A call for clarity in reporting BLEU scores. In *Proc. Conf. on Machine Translation*, pp. 186–191, Brussels, Belgium, October 2018.
- Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., and Sutskever, I. Language models are unsupervised multi-task learners. [Online]. : <https://blog.openai.com/better-language-models/>, 2019.

- Rae, J. W., Potapenko, A., Jayakumar, S. M., Hillier, C., and Lillicrap, T. P. Compressive transformers for long-range sequence modelling. In *Int. Conf. on Learning Representations (ICLR)*, Virtual only, April 2020.
- Ramsauer, H., Schäfl, B., Lehner, J., Seidl, P., Widrich, M., Gruber, L., Holzleitner, M., Adler, T., Kreil, D., Kopp, M. K., Klambauer, G., Brandstetter, J., and Hochreiter, S. Hopfield networks is all you need. In *Int. Conf. on Learning Representations (ICLR)*, Virtual only, May 2021.
- Schlag, I. and Schmidhuber, J. Gated fast weights for on-the-fly neural program generation. In *NIPS Metalearning Workshop*, Long Beach, CA, USA, December 2017.
- Schlag, I. and Schmidhuber, J. Learning to reason with third order tensor products. In *Proc. Advances in Neural Information Processing Systems (NIPS)*, pp. 9981–9993, Montréal, Canada, December 2018.
- Schlag, I., Munkhdalai, T., and Schmidhuber, J. Learning associative inference using fast weight memory. In *Int. Conf. on Learning Representations (ICLR)*, Virtual only, May 2021.
- Schmidhuber, J. Learning to control fast-weight memories: An alternative to recurrent nets. Technical Report FKI-147-91, Institut für Informatik, Technische Universität München, March 1991.
- Schmidhuber, J. Learning to control fast-weight memories: An alternative to dynamic recurrent networks. *Neural Computation*, 4(1):131–139, 1992.
- Schmidhuber, J. Reducing the ratio between learning complexity and number of time varying variables in fully recurrent nets. In *International Conference on Artificial Neural Networks (ICANN)*, pp. 460–463, Amsterdam, Netherlands, September 1993.
- Schmidhuber, J. 26 March 1991: Neural nets learn to program neural nets with fast weights—like today’s Transformer variants. 2021: New stuff!, AI Blog, 2021. URL <https://people.idsia.ch/~juergen/fast-weight-programmer-1991-t-transformer.html>.
- Sennrich, R., Haddow, B., and Birch, A. Neural machine translation of rare words with subword units. In *Proc. Association for Computational Linguistics (ACL)*, pp. 1715–1725, Berlin, Germany, August 2016.
- Shen, Z., Zhang, M., Zhao, H., Yi, S., and Li, H. Efficient attention: Attention with linear complexities. *Preprint arXiv:1812.01243*, 2018.
- Smolensky, P. Tensor product variable binding and the representation of symbolic structures in connectionist systems. *Artificial intelligence*, 46(1-2):159–216, 1990.
- Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., and Salakhutdinov, R. Dropout: a simple way to prevent neural networks from overfitting. *The Journal of Machine Learning Research*, 15(1):1929–1958, 2014.
- Steinbuch, K. Die lernmatrix. *Kybernetik*, 1(1):36–45, 1961.
- Steinbuch, K. and Piske, U. A. W. Learning matrices and their applications. *IEEE Transactions on Electronic Computers*, 12(6):846–862, 1963.
- Tay, Y., Dehghani, M., Abnar, S., Shen, Y., Bahri, D., Pham, P., Rao, J., Yang, L., Ruder, S., and Metzler, D. Long range arena: A benchmark for efficient transformers. In *Int. Conf. on Learning Representations (ICLR)*, Virtual only, May 2021.
- Tsai, Y.-H. H., Bai, S., Yamada, M., Morency, L.-P., and Salakhutdinov, R. Transformer dissection: An unified understanding for transformer’s attention via the lens of kernel. In *Proc. Conf. on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 4344–4353, Hong Kong, China, November 2019.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., and Polosukhin, I. Attention is all you need. In *Proc. Advances in Neural Information Processing Systems (NIPS)*, pp. 5998–6008, Long Beach, CA, USA, December 2017.
- von der Malsburg, C. The correlation theory of brain function. Internal Report 81-2, Göttingen: Department of Neurobiology, Max Planck Intitute for Biophysical Chemistry, 1981.
- Widrow, B. and Hoff, M. E. Adaptive switching circuits. In *Proc. IRE WESCON Convention Record*, pp. 96–104, Los Angeles, CA, USA, August 1960.