

Not all parameters are born equal: Attention is mostly what you need

Nikolay Bogoychev
School of Informatics
University of Edinburgh
United Kingdom
n.bogoych@ed.ac.uk

Abstract

Transformers are widely used in state-of-the-art machine translation, but the key to their success is still unknown. To gain insight into this, we consider three groups of parameters: embeddings, attention, and Feed-Forward Neural network (FFN) layers. We examine the relative importance of each by performing an ablation study where we initialise them at random and freeze them, so that their weights do not change over the course of the training. Through this, we show that the attention and FFN are equally important and fulfil the same functionality in a model. We show that the decision about whether a component is frozen or allowed to train is at least as important for the final model performance as its number of parameters. At the same time, the number of parameters alone is not indicative of a component's importance. Finally, while the embedding layer is the least essential for machine translation tasks, it is the most important component for language modelling tasks.

1 Introduction

Bizarrely, we find that a transformer model (Vaswani et al., 2017) for machine translation loses just 1.1–2% BLEU when any one of the three main components (embeddings, attention, and feed-forward) is frozen with randomly initialised parameters. In more extreme cases we see that with just 6% of trainable parameters, we maintain 72% of the baseline BLEU. This suggests that redundancy in the model architecture is able to compensate for poor components. We perform a case study on transformer style models, exploring the built-in redundancy and the limits of what can be achieved when the majority of parameters are frozen after random initialisation. We do this in several different transformer architecture presets for neural machine translation: transformer-big, transformer-base and a transformer-tiny student preset, and one

preset for a language modelling task. Our key findings are:

- The majority of the model information is learned by the attention and the Feed-Forward Neural network (FFN), while the embeddings provide complementary but non-essential signals. The attention and the FFN layers are mostly interchangeable.
- Even if it is not essential for some components to be learned, their size (number of parameters) is essential to provide extra expressiveness for the components that are learned: if components are initialised as diagonal matrices as opposed to random matrices, they further lose their effectiveness.
- The overall size of the model does affect the results: Larger models have more built-in redundancy and are able to retain higher performance with fewer trainable components.
- The smaller proportion of trainable parameters the model has, in general, leads to slower convergence in terms of number of epochs, suggesting learning is hampered.
- Our findings are task specific: While machine translation is learned in the attention and the FFN layer, this does not necessarily generalise to other tasks. Language modelling is learned in the embeddings layer.

Our study sheds light on the inner workings of the transformer neural network architecture and examines in detail the relative contribution of each component towards the final performance.

2 Methodology

We perform a parameter-freezing case study on transformer style models. A transformer consists of three basic components: an embeddings layer,

attention layers, and FFN, as shown in Figure 1. The embeddings encompass the source, target embeddings and the output layer as they are tied (Press and Wolf, 2017).

We initialise the model with the Glorot distribution (Glorot and Bengio, 2010) and pick a component (or components) to mark as untrainable or “frozen”.¹ This effectively means that during training the model will benefit from the extra transformations that those parameters provide, but they are random. It is up to the remaining components to accommodate them. Freezing parameters is frequently performed in transfer learning (Zoph et al., 2016; Thompson et al., 2018; Clinchant et al., 2019) in order to prevent the child model from deviating in an undesired manner from the parent. Through these experiments we identify which components are essential for the model, which are interchangeable and which are complementary. An essential trainable component by itself can achieve high performance with the other components being random, while a non-essential component is not able to achieve good performance by itself. Finally, we exchange the frozen random initialisation of components with frozen diagonal (identity) initialisation to verify how important the transformations of the frozen components are.

In the next few sections, we first train a baseline model with no frozen components. Then we train 3 models with one frozen component: either the embedding layer, the FFN layer or the attention. Finally we perform experiments where two out of the three components are frozen.

3 Transformer-big experiments

For our main experiments we used 1M parallel Turkish-English sentences from the WMT 18 News translation task (Bojar et al., 2018) (200K gold and 800K back-translated (Sennrich et al., 2016a)). Data was preprocessed and tokenised using Moses (Koehn et al., 2007) and a shared vocabulary of 36K byte-pair encoding (BPE) tokens (Sennrich et al., 2016b). For training the model, we used the Marian (Junczys-Dowmunt et al., 2018) implementation of the transformer-big architecture (Vaswani et al., 2017) with tied embeddings and 8 attention heads for a total of 213M parameters. This model, and all subsequent were trained until 10 consecutive stalls on the perplexity of the dev set or 50

¹Code available at https://github.com/XapaJIaMnu/marian-dev/tree/freeze_ffn




Figure 1: A simplified view of transformer with input, output embeddings, and output layer tied together. The encoder self-attends to the entire input, while the decoder self-attends to the output it has produced thus far.

consecutive stalls on the BLEU score of the dev set, whichever is sooner. All the BLEU (Papineni et al., 2002) scores are reported on the 2016 dev set provided by the shared task organisers. We present our experiments in Table 1. For reader convenience, each system is also referred to by its number in the result analysis.

	Component			BLEU	Epochs	Parameter ratio Trainable/All
	EMB	ATT	FFN			
(0)	✓	✓	✓	24.3	19	1
One frozen component						
(1)	✗	✓	✓	22.6	26	.82
(2)	✓	✗	✓	22.3	23	.64
(3)	✓	✓	✗	23.2	26	.52
Multiple frozen components						
(4)	✗	✓	✗	21.5	36	.35
(5)	✗	✗	✓	20.8	37	.47
(6)	✓	✗	✗	4.4	25	.17

Table 1: Trainable and random components for Turkish-English transformer-big. Trainable components are marked with ✓, fixed components with ✗. BLEU is computed on the dev set.

Freezing a subset of the model leads to a tiny increase in the training speed in terms of words processed per second for all models, due to the reduced computational cost when performing the backward

pass. However, this speed benefit is negated by the extra epochs required to achieve convergence, to the point that every single model was slower to converge in terms of real time. The increased training time is likely due to the optimiser running into more saddle points, as there are less parameters that can be pushed into a better model direction (Pascanu et al., 2014; Dauphin et al., 2014). Freezing components also always leads to reduced performance in terms of BLEU compared to the baseline, which is expected and is also observed when transferring pretrained components (Zoph et al., 2016; Clinchant et al., 2019; Aji et al., 2020).

Our results show that when there is one frozen component, all models perform very similarly, with the frozen FFN one (3) being slightly better than the other two. However, when we look at the experiments where we have multiple frozen components, it is evident that the attention and the FFN layer are more essential for the overall model performance: If both of them are initialised randomly and left frozen, the model fails to learn completely (6).

Based on this we can conclude that the embedding layer is the least essential to be learned in the context of the machine translation task and the remaining components can easily learn to work with random embeddings. This result confirms the findings of Aji et al. (2020), who show that transferring the embedding layer during transfer learning is the least important for the overall model performance. In the rest of the transformer-big experiments, we omit the system with just the embedding layer trainable, as the model quality in this setting is very poor.

We also note that a model with trainable embeddings and attention (3) outperforms a model with frozen embeddings (1), despite the FFN and attention being more essential to be learned. This suggests that the embeddings perform a different role from the attention and the FFN in the translation process. Based on the similar performance that the trainable attention and the FFN offer, we hypothesise that they are able to encode the same information and in case that one of them is not trainable, the other one can mostly fulfil its role. Our hypothesis is supported by the work of Geva et al. (2020), who show that feed forward transformer layers are key-value memories. On the other hand, trainable embeddings, while least essential for the model performance likely learn complementary, albeit non-essential, information to the attention and

the FFN.

We find a weak correlation between the number of parameters of a component and its importance to the model: The embeddings layer has the least amount of parameters and is the least essential to be trained, however we can't rank the other two components: The FFN has considerably more parameters than the attention, but they both seem to be equally important for the model's performance.

We note that the model with just attention trainable (4) is able to reach competitive performance, despite having just 35% of the model parameters.

3.1 Reducing the embedding layer

Since we found that the embedding layer is the least essential to be trained, we decided to shrink its dimension from the 1024 (default for transformer-big) to 128, set it to trainable and vary the attention layer and FFN layer. By doing this, we implicitly reduce the width of the other components too, because the matrix dimensions need to match. The resulting models have only 18M parameters (Table 2) are markedly worse in terms of BLEU compared to the models with a bigger embedding layer.

	Component			BLEU	Epochs	Parameter ratio Trainable/All
	EMB	ATT	FFN			
(0)	Large baseline			24.3	19	1
	✓	✓	✓	22.9	115	1
One frozen component						
(1)	✗	✓	✓	18.9	142	.75
(2)	✓	✗	✓	16.2	274	.93
(3)	✓	✓	✗	18.9	132	.31
Multiple frozen components						
(4)	✗	✓	✗	5.1	227	.07
(5)	✗	✗	✓	7.8	307	.69

Table 2: Turkish-English transformer-big with small, embedding layer dimension. Trainable components are marked with ✓, fixed components are marked with ✗. BLEU is computed on the dev set.

These experiments show larger difference in BLEU between attention and FFN: With this smaller model, when embeddings are trainable, a trainable attention (3) reaches much better BLEU score, and in half of the epochs, compared to trainable FFN (2). In addition to that, the attention layer has a lot fewer parameters compared to FFN, showing that the number of parameters a component has is not a marker for how important it is to have that component trainable for the final model's performance. As the difference between the performance

of trainable attention vs trainable FFN is more pronounced in the smaller model setting, we can say that model size can change the conclusions. We also note that in this setting, trainable attention and FFN layer (1) achieve the same performance as (3), despite having more than twice the number of parameters, suggesting again that the embeddings are able to learn complementary information to the attention, even if their dimension is tiny.

Furthermore we note that despite that these models are 12 times smaller than the baseline, they converge slower, requiring up to 10 times more epochs, suggesting learning is more difficult. This further highlights that the attention is a more essential component than the FFN, as it can adapt to the training data with fewer signals.

This model has catastrophic loss of BLEU when there is only one trainable component (4, 5), suggesting that larger models are more robust and have more redundancy between components.

3.2 Reducing the FFN layer size

We perform a similar experiment with the FFN layer, reducing its dimension from 4096 to 1024. The resulting models have 137M parameters and are shown in Table 4.

	Component			BLEU	Epochs	Parameter ratio Trainable/All
	EMB	ATT	FFN			
(0)	trans-big	baseline		24.3	19	1
	✓	✓	✓	23.2	22	1
One frozen component						
(1)	✗	✓	✓	22.0	38	.73
(2)	✓	✗	✓	20.1	36	.55
(3)	✓	✓	✗	23.0	18	.82
Multiple frozen components						
(4)	✗	✓	✗	21.6	34	.45
(5)	✗	✗	✓	18.0	98	.18

Table 3: Trainable and random components for Turkish-English transformer-big with small FFN layer. Trainable components are marked with ✓, fixed components with ✗. BLEU is computed on the dev set.

This experiment shows similar trends as the one in the previous section: The model requires more epochs to converge compared to the baseline transformer-big and reaches slightly lower quality. As the size of the FFN layer is reduced, we observe that a trainable attention is always better than a trainable FFN layer (Table 3, (2) vs (3), and (4) vs (5)). Finally, when the FFN layer is reduced in size, it matters less whether it is trainable or random

and frozen: When it is random (3), we only lose .2 BLEU compared to the baseline system (0). Nevertheless, just the small trainable FFN layer which makes up for 18% of the total model parameters is able to reach 18.0 BLEU (5).

We note that when all of the components of this model are trainable, it achieves the same performance as the baseline system with frozen FFN (Table 1, (3)). This suggests that for certain components, their width may be as important as whether they are trainable or not. In this setting we also see a sharper drop in performance when attention is not trainable, suggesting the capacity of the FFN to take over the attention’s function has diminished.

3.3 Reducing the attention layer size

We perform an analogous experiment, where we reduce the size of the second dimension of the keys and queries attention matrices 8-fold, reducing the total model size to 147M parameters. In this setting, all attention matrices make up just 6% of the total model size. We show our experiments in Table 4.

	Component			BLEU	Epochs	Parameter ratio Trainable/All
	EMB	ATT	FFN			
(0)	trans-big	baseline		24.3	19	1
	✓	✓	✓	23.4	13	1
One frozen component						
(1)	✗	✓	✓	22.1	23	.74
(2)	✓	✗	✓	22.1	16	.94
(3)	✓	✓	✗	20.4	23	.31
Multiple frozen components						
(4)	✗	✓	✗	17.0	196	.06
(5)	✗	✗	✓	21.3	33	.69

Table 4: Trainable and random components for Turkish-English transformer-big with smaller attention matrices. Trainable components are marked with ✓, fixed components with ✗. BLEU is computed on the dev set.

In this setting, our findings are mirroring the results from the previous section: When the size of the attention layer is reduced to that extent, we see that the FFN layer is now more essential for the performance of the model: It outperforms the attention layer in all direct comparisons (Table 4, (2) vs (3), and (4) vs (5)). We note that despite the attention layer having only 6% of the total number of parameters, by itself (4), it manages to retain 72% of the baseline BLEU score.

We note that in this experiment, as well as in the previous where we reduce the dimension of the FFN layer, there is a clear trend that the number of

epochs necessary for convergence increases as the number of trainable model parameters decreases.

3.4 Self and context attention

The experiments that we have conducted so far suggest that the attention layer is the most essential for the translation performance. However, attention refers to several different attention modules in a transformer. Attention can be split into self-attention and context attention. Self-attention is present in both the encoder and the decoder. The encoder attends to the full input sequence, while the decoder has masked attention and attends to the output it has produced thus far. Context attention is an attention module that appears only in the decoder and attends to the encoded sequence.

We can also make the split by component, having encoder-only attention and decoder-only attention. The different types of attention and their position in the transformer architecture are shown in Figure 1.

In this experiment, we take the transformer-big model used previously, set the embedding layer to trainable and freeze the FFN layer. We specifically avoid training the FFN layer, because it can take over some of the functionality of our restricted attention layer and this will make our findings less clear. For our experiments, we only train a subset of the attention (the remaining attention parts are left frozen and random), as shown in Table 5.

	Attention	BLEU	Epochs	Parameter ratio Trainable/All
(0)	full	23.2	26	.52
(1)	Self only	10.3	28	.41
(2)	Context only	18.1	40	.29
(3)	Encoder only	18.1	41	.29
(4)	Decoder only	19.8	24	.41

Table 5: tr-en transformer-big with trainable embeddings, fixed FFN and a trainable subset of attention. BLEU is computed on the dev set.

We see significant drop in BLEU when having only subsection of the attention as learnable. Referring back to our initial experiments in Table 1, partially setting the attention to untrainable is more harmful for the model than freezing it completely and having the FFN trainable.

Voita et al. (2019) show that context heads are the most important heads for a model and are the ones that are pruned last when performing head pruning. In our setting however, even if context heads are not learned, the model can still use the

random activations with which they are initialised. That is enough for a model with trainable encoder attention to produce a BLEU score similar to that of a model that uses trainable context attention only.

The best result is when we have the full attention on the decoder side, as the decoder is the most involved component in translation (Voita et al., 2019).

In our experiments, trainable encoder attention performs better than self-attention only, despite the former having a superset of the latter’s trainable parameters (Figure 1).

Independently from us, Tay et al. (2020) observe that the self-attention component is not essential in the neural machine translation task and can be synthesised on demand. They propose several different functions to do so and achieve smaller quality drop than us, however all of their experiments involve a more advanced synthesis than simply initialising randomly and not training. Raganato et al. (2020) show that it is possible to replace all but one attention heads with a fixed attentive pattern without impact on translation quality.

Further to this line of work, Liu et al. (2021) show that attention could be replaced with a simple gated MLP network and still achieve competitive results, albeit they only show that for vision tasks.

4 Experiments with smaller models

The transformer-big preset used in the previous experiments is an overly expressive neural network with enough parameters, so that even when preventing a big part of them from being trained, the model likely has enough built-in redundancy to make up for it. In order to avoid this assumption, we performed analogous experiments with a transformer-base model, as well as a knowledge-distilled student transformer model. These experiments aim to reduce the model size in a systematic way in order to see whether our findings in the previous sections are dependent on the overall model size.

4.1 Experiments with transformer-base

We perform the same experiments as the ones in Section 3, but we change the model configuration to transformer-base. In effect this means that the number of attention heads remains the same, but the embeddings’ and the FFN’s dimension is halved. The overall size of the attention is also reduced implicitly, as its dimension is also dependent on the embeddings. The resulting model has 62.6M parameters. We present our results in Table 6.

	Component			BLEU	Epochs	Parameter ratio Trainable/All
	EMB	ATT	FFN			
	trans-big baseline			24.3	19	1
(0)	✓	✓	✓	23.9	32	1
One frozen component						
(1)	✗	✓	✓	22.3	62	.71
(2)	✓	✗	✓	20.5	40	.70
(3)	✓	✓	✗	22.7	39	.60
Multiple frozen components						
(4)	✗	✓	✗	20.0	148	.30
(5)	✗	✗	✓	17.5	76	.40
(6)	✓	✗	✗	3.9	110	.30

Table 6: Trainable and random components for Turkish-English transformer-base. Trainable components are marked with ✓, fixed components with ✗. BLEU is computed on the dev set.

In the transformer-base setting we see similar trends as those with transformer-big (Table 1), except the quality drop when making components untrainable is more severe. We see that the difference between the attention layer and the FFN layer ((2) vs (3), (4) vs (5)) in terms of quality is widening, reaffirming that the attention is the most essential component. The FFN also has more parameters than the attention, showing again that the number of parameters a component has is not indicative of how important it is to have it trainable.

The overall trend we observe is that the smaller the model is, the bigger the gap between the attention and the FFN layer. In a transformer-big, there is more built-in redundancy and when a component is marked as untrainable, it is easy for the other components to step up and make up for it, however as the size of individual components decreases, this is no longer possible.

4.2 Knowledge distilled models

We performed an analogous experiment with a knowledge-distilled, small student model (Kim and Rush, 2016). The student translation models are normally too small to learn a good model from natural data, but given enough synthetic data, they can learn the distribution of a good teacher model. For this setting we trained an ensemble of four transformer-big Estonian-English models on gold and synthetic data from WMT 18 (Bojar et al., 2018). Then, we decoded 142M sentences with the ensemble and used those to train a student model. The student model has 32K shared source and target vocabulary tokens, embedding dimension of

256, 6 encoder layers, 2 decoder layers and 1536 dimensional Feed forward neural network, similar to the student models used in Bogoychev et al. (2020). Just as in the previous experiment, the source embeddings are tied to the target embeddings and to the output layer. This model has 16.9M parameters. We report BLEU scores on the dev set provided by the WMT 18 news task organisers (Bojar et al., 2018). We present our experiments in Table 7.

	Component			BLEU	Epochs	Parameter ratio Trainable/All
	EMB	ATT	FFN			
(0)	✓	✓	✓	24.2	14	1
One frozen component						
(1)	✗	✓	✓	20.3	15	.51
(2)	✓	✗	✓	21.6	6	.87
(3)	✓	✓	✗	21.3	4	.62
Multiple frozen components						
(4)	✗	✓	✗	12.5	11	.14
(5)	✗	✗	✓	15.2	8	.39
(6)	✓	✗	✗	3.3	35	.49

Table 7: Trainable and random components for Estonian-English student transformer. Trainable components are marked with ✓, fixed components are marked with ✗. BLEU is computed on the dev set.

In this setting, we observe a few differences from the experiments on bigger transformers: Here, the trainable embedding layer in conjunction with either trainable attention (3), or FFN (2) delivers the best results, outperforming the configuration where it is frozen (1). Despite the embedding layer performing the worst by itself, when it is trainable, it is more valuable to the attention and the FFN. This experiment reaffirms the hypothesis that the attention and FFN layer fulfil the same functionality, and when one of them is frozen, the other is mostly able to compensate for it. However the embedding layer provides a different signal. Again, the embedding layer is unable to learn a meaningful translation model, despite containing nearly half of the model’s parameters, clearly showing that the number of parameters is not related to how important a feature is. We also see that this student model is a lot more sensitive to the frozen parameters compared to the bigger transformers, with all results being markedly worse than the baseline, especially the ones with multiple frozen components.

The configuration of this student model differs significantly from the other architectures used, and it also behaves differently when freezing parameters. Therefore, the overall size of the model does

affect the behaviour with regards to frozen and random parameters.

5 Language model experiments

In order to verify whether our findings generalise to other tasks, we performed a language modelling experiment. We trained a transformer language model, using 78M sentences taken from newscrawl 2017, 2018 and 2019. Our model is based on the transformer-base preset (Vaswani et al., 2017) with 32K BPE tokens vocabulary and 38M parameters. We show our results in Table 8, where we report perplexity on a dev set, formed by taking the first 1000 lines from each of the three newscrawl datasets. The devset is disjoint from the training set.

	Component			PPL	Epochs	Parameter ratio Trainable/All
	EMB	ATT	FFN			
(0)	✓	✓	✓	37.4	6	1
One frozen component						
(1)	✗	✓	✓	118.4	6	.53
(2)	✓	✗	✓	47.5	6	.81
(3)	✓	✓	✗	50.3	6	.64
Multiple frozen components						
(4)	✗	✓	✗	209.3	6	.18
(5)	✗	✗	✓	157.3	6	.36
(6)	✓	✗	✗	131.7	6	.46

Table 8: Trainable and random components for an English language model. Trainable components are marked with ✓, fixed components are marked with ✗. Perplexity is computed on the dev set.

A neural machine translation system is just a conditional language model, so this task is similar to our previous experiments, except that there is no encoder and no context attention. The transformer components in this case behave differently: the embedding layer is the most essential for the final model performance, capable of delivering the lowest perplexity of any system with multiple frozen components ((2) and (3)), and nearly outperforming by itself (6) a combination of trainable attention and FFN (1). Trainable embeddings and trainable either FFN (2) or attention (3) perform very similarly, suggesting again that they fulfil similar functionality in the model. By itself, however, trainable FFN (5) clearly outperforms trainable attention (4) suggesting that it is more important for this task. It is interesting to note that the number of epochs necessary for convergence does not vary among the different configurations used.

A change in the task can lead to a big change

in the overall behaviour of components. For this task, the embedding layer loses the majority of its effectiveness if it is not trainable, unlike in the translation experiments.

6 Randomness and transformations

In order to measure how much the transformations provided by the random parameters help the trainable parameters, we performed an experiment where we initialise either the attention or the FFN layer of a transformer-big model to a diagonal matrix, with zeroes everywhere except at the main diagonal. For the attention this means that all three attention matrices in self and context attention are diagonal and just copy over their input to the output, however since attention is interpreted as different heads, it does mean that the individual heads are very sparse and attend only to a single token. In the case of the FFN layer, since it is not square, the upper part of the matrix is diagonal and the lower part is just zeroes.² We present our results on transformer-big model from Section 3 in Table 9.

	Component			BLEU	Epochs	Parameter ratio Trainable/All
	EMB	ATT	FFN			
(0)	✓	✓	✓	24.3	19	1
Random frozen component						
(2)	✓	✗	✓	22.3	23	.64
(3)	✓	✓	✗	23.2	26	.52
Diagonal zeroed frozen component						
(2.1)	✓	✗	✓	19.4	24	.64
(3.1)	✓	✓	✗	22.9	20	.52

Table 9: Trainable and random components for Turkish-English transformer-big contrasting randomly initialised parameters with diagonal matrix parameters. Trainable components are marked with ✓, fixed components with ✗. BLEU is computed on the dev set.

Our results are conflicting. When the attention is diagonal (3), the BLEU score drops significantly from when it is initialised to random (3.1). This suggests that the model indeed makes use of and learns to work with random parameters and the transformations they provide, however it can not adapt to work with diagonal parameter matrices, as they don't provide useful input.

In the case of the FFN layer, we see the opposite results. Making the FFN layer matrices have 1 only across their main diagonal (2.1) seems to marginally improve performance compared to the

²We used `numpy.eye` to initialise the matrices.

baseline case (2). Since the FFN layer is bypassed by residual connections which connect the attention to the output layer, we don't lose any information from the attention, and the ones across the main diagonal of the FFN allows for some minimal mixing of the output of the attention. In this case we can clearly see that the random transformations provided by the untrainable FFN are not as important, as the attention can do its job without them.

We conclude that in certain cases the model does benefit from the transformation provided by the random parameters, however this does not generalise to every component of a transformer. We did not perform this experiment with the embedding layer, as it is not square (The shape is (32000,512) or similar) and it is used to encode different vocabulary items. Rendering it diagonal across the main diagonal would leave the majority of the vocabulary items encoded as a 512 zero vector, which is not something that the model can learn from.

7 Discussion

We find that depending on the task, different components of a transformer are essential. Language modelling tasks require the embeddings to be trained, while translation can mostly cope without them. A large amount of research has focused on using BERT (Devlin et al., 2018) or ELMo (Peters et al., 2018) pre-trained embeddings in downstream tasks, but as shown here, embedding representation is not as important as embedding width in a translation task. Furthermore, there is evidence that in some downstream tasks, randomised language models perform competitively (Zhang and Bowman, 2018; Wang et al., 2018). We urge researchers to always evaluate how essential a component is by setting it to random and frozen, before performing transfer learning from pretrained components.

We find that the "width" of a component is at least as important for the final model performance as whether it is trainable or not, but the raw parameter size of a component is not indicative of whether it is essential to have it trained or not. This clearly suggests that when there is a frozen and random component of a model, the rest of the components are able to work around just as if it were an extra activation function. When the attention is rendered completely linear by making it diagonal, we show a more significant drop in the performance compared to when this is done to the FFN layer. This suggests that the transformation that happens in the

attention is more important than the transformation that happens in the FFN layer.

We find that the smaller the model is, the more sensitive it is to having some of its components non-trainable. We suspect that transformer-big models have more built-in redundancy into them and adapt better to this situation. Curiously, the less parameters a model has, or the less trainable parameters the model has, the more epochs it takes to train. This is especially visible in the experiments of transformer-base (Table 6) vs transformer-big (Table 1) and the models with reduced embedding layer (Table 2). These effects are explainable by the works of Choromanska et al. (2015); Jacot et al. (2020); Sankararaman et al. (2020) which state that smaller networks are more likely to end up into local minimum and are more sensitive to parameter initialisation compared to bigger networks.

Finally we note how similarly the attention and FFN layers behave: When the attention layer is small, the FFN layer is more essential and vice versa. In a setting where they are trainable by themselves, they achieve similar performance, with the attention having slight edge in larger models. If we decrease the size of the embeddings layer, or use a transformer-base preset, the performance gap between the FFN and the attention widens in favour of the attention. This echoes the findings of Geva et al. (2020) and suggests that the attention and the FFN layer are interchangeable and more essential for the model's performance, however there is a degree of redundancy between them.

8 Related work

Concurrently and independently from us, Garg et al. (2020) perform an ablation study to show how gradually setting different components of an encoder-decoder MT system from random and frozen to trainable increases the BLEU scores from 4 to 39 for an English-French model. Unlike Garg et al. we use a transformer model and we show that wider and deeper models behave differently from small and narrow models. Furthermore we show that depending on the task, the model benefits the most from different components.

Our work follows a long line of research that explores the performance of random (Echo state) neural networks in various tasks (Ramamurthy et al., 2019; Tong et al., 2007; Homma and Hagiwara, 2013; Alhama and Zuidema, 2018).

Enguehard et al. (2019); Wieting and Kiela

(2019) explore random encoders for sentences and evaluate their performance in downstream tasks.

9 Conclusion and Future work

We performed a study of transformer neural machine translation component freezing that identified the attention and the FFN as more essential to be learned, but also interchangeable. The embedding layer can not perform the translation task by itself, but the information it learns is complementary to either the attention or the FFN layer. Our findings do not necessarily generalise over other tasks: In the language modelling setting, the components behave differently.

Our findings could serve as the basis for research into new transformer machine translation architectures for production (Kim et al., 2019; Bogoychev et al., 2020). We have shown that models can retain competitive performance even when the FFN layer, which accounts for up to 48% of the parameters, is just a randomly initialised Glorot distribution. This opens the door for models where part of the parameters are generated on the fly during decoding, as opposed to stored statically on disk, making it particularly suitable to embedded devices with limited memory.

Acknowledgments

We thank the reviewers for their suggestions and comments! We thank Kenneth Heafield, Adam Lopez, Naomi Saphra and the AGORA research group for the discussions and comments that shaped this paper. We thank Maximiliana Behnke for the use of her pretrained models.



This work was conducted within the scope of the Research and Innovation Action *Bergamot*, which has received funding from the European Union’s Horizon 2020 research and innovation programme under grant agreement No 825303.

References

- Alham Fikri Aji, Nikolay Bogoychev, Kenneth Heafield, and Rico Sennrich. 2020. [In neural machine translation, what does transfer learning transfer?](#) In *Proceedings of the 2020 Annual Conference of the Association for Computational Linguistics*, Seattle.
- Raquel Alhama and Willem Zuidema. 2018. [Pre-wiring and pre-training: What does a neural network need to learn truly general identity rules?](#) *Journal of Artificial Intelligence Research*, 61:927–946.
- Nikolay Bogoychev, Roman Grundkiewicz, Alham Fikri Aji, Maximiliana Behnke, Kenneth Heafield, Sidharth Kashyap, Emmanouil-Ioannis Farsarakis, and Mateusz Chudyk. 2020. [Edinburgh’s submissions to the 2020 machine translation efficiency task.](#) In *Proceedings of the The 4th Workshop on Neural Generation and Translation (WNGT 2020)*, Seattle.
- Ondřej Bojar, Christian Federmann, Mark Fishel, Yvette Graham, Barry Haddow, Philipp Koehn, and Christof Monz. 2018. [Findings of the 2018 conference on machine translation \(WMT18\).](#) In *Proceedings of the Third Conference on Machine Translation: Shared Task Papers*, pages 272–303, Belgium, Brussels. Association for Computational Linguistics.
- Anna Choromanska, Mikael Henaff, Michael Mathieu, Gérard Ben Arous, and Yann LeCun. 2015. [The loss surfaces of multilayer networks.](#)
- Stephane Clinchant, Kweon Woo Jung, and Vassilina Nikoulina. 2019. [On the use of BERT for neural machine translation.](#) In *Proceedings of the 3rd Workshop on Neural Generation and Translation*, pages 108–117, Hong Kong. Association for Computational Linguistics.
- Yann N Dauphin, Razvan Pascanu, Caglar Gulcehre, Kyunghyun Cho, Surya Ganguli, and Yoshua Bengio. 2014. [Identifying and attacking the saddle point problem in high-dimensional non-convex optimization.](#) In *Advances in Neural Information Processing Systems*, volume 27, pages 2933–2941. Curran Associates, Inc.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. [Bert: Pre-training of deep bidirectional transformers for language understanding.](#)
- Joseph Enguehard, Dan Busbridge, Vitalii Zhelezniak, and Nils Hammerla. 2019. [Neural language priors.](#) *ArXiv*, abs/1910.03492.
- Ankush Garg, Yuan Cao, and Qi Ge. 2020. [Echo state neural machine translation.](#)
- Mor Geva, Roei Schuster, Jonathan Berant, and Omer Levy. 2020. [Transformer feed-forward layers are key-value memories.](#)
- Xavier Glorot and Y. Bengio. 2010. Understanding the difficulty of training deep feedforward neural networks. *Journal of Machine Learning Research - Proceedings Track*, 9:249–256.
- Yukinori Homma and Masafumi Hagiwara. 2013. An echo state network with working memories for probabilistic language modeling. In *Artificial Neural Networks and Machine Learning – ICANN 2013*, pages 595–602, Berlin, Heidelberg. Springer Berlin Heidelberg.

- Arthur Jacot, Franck Gabriel, and Clément Hongler. 2020. [Neural tangent kernel: Convergence and generalization in neural networks](#).
- Marcin Junczys-Dowmunt, Roman Grundkiewicz, Tomasz Dwojak, Hieu Hoang, Kenneth Heafield, Tom Neckermann, Frank Seide, Ulrich Germann, Alham Fikri Aji, Nikolay Bogoychev, André F. T. Martins, and Alexandra Birch. 2018. [Marian: Fast neural machine translation in C++](#). In *Proceedings of ACL 2018, System Demonstrations*, pages 116–121, Melbourne, Australia. Association for Computational Linguistics.
- Yoon Kim and Alexander M. Rush. 2016. [Sequence-level knowledge distillation](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1317–1327, Austin, Texas. Association for Computational Linguistics.
- Young Jin Kim, Marcin Junczys-Dowmunt, Hany Hassan, Alham Fikri Aji, Kenneth Heafield, Roman Grundkiewicz, and Nikolay Bogoychev. 2019. [From research to production and back: Ludicrously fast neural machine translation](#). In *Proceedings of the 3rd Workshop on Neural Generation and Translation*, pages 280–288, Hong Kong. Association for Computational Linguistics.
- Philipp Koehn, Hieu Hoang, Alexandra Birch, Chris Callison-Burch, Marcello Federico, Nicola Bertoldi, Brooke Cowan, Wade Shen, Christine Moran, Richard Zens, Chris Dyer, Ondřej Bojar, Alexandra Constantin, and Evan Herbst. 2007. [Moses: Open source toolkit for statistical machine translation](#). In *Proceedings of the 45th Annual Meeting of the Association for Computational Linguistics Companion Volume Proceedings of the Demo and Poster Sessions*, pages 177–180, Prague, Czech Republic. Association for Computational Linguistics.
- Hanxiao Liu, Zihang Dai, David R. So, and Quoc V. Le. 2021. [Pay attention to mlps](#).
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. [Bleu: a method for automatic evaluation of machine translation](#). In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.
- Razvan Pascanu, Yann N. Dauphin, Surya Ganguli, and Yoshua Bengio. 2014. [On the saddle point problem for non-convex optimization](#).
- Matthew E. Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. [Deep contextualized word representations](#).
- Ofir Press and Lior Wolf. 2017. [Using the output embedding to improve language models](#). In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, pages 157–163, Valencia, Spain. Association for Computational Linguistics.
- Alessandro Raganato, Yves Scherrer, and Jörg Tiedemann. 2020. [Fixed encoder self-attention patterns in transformer-based machine translation](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 556–568, Online. Association for Computational Linguistics.
- Rajkumar Ramamurthy, Robin Stenzel, Rafet Sifa, Anna Ladi, and Christian Bauckhage. 2019. [Echo state networks for named entity recognition](#). In *Artificial Neural Networks and Machine Learning - ICANN 2019 - 28th International Conference on Artificial Neural Networks, Munich, Germany, September 17-19, 2019, Proceedings - Workshop and Special Sessions*, volume 11731 of *Lecture Notes in Computer Science*, pages 110–120. Springer.
- Karthik A. Sankararaman, Soham De, Zheng Xu, W. Ronny Huang, and Tom Goldstein. 2020. [The impact of neural network overparameterization on gradient confusion and stochastic gradient descent](#).
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016a. [Improving neural machine translation models with monolingual data](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 86–96, Berlin, Germany. Association for Computational Linguistics.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016b. [Neural machine translation of rare words with subword units](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1715–1725, Berlin, Germany. Association for Computational Linguistics.
- Yi Tay, Dara Bahri, Donald Metzler, Da-Cheng Juan, Zhe Zhao, and Che Zheng. 2020. [Synthesizer: Rethinking self-attention in transformer models](#).
- Brian Thompson, Huda Khayrallah, Antonios Anastasopoulos, Arya D. McCarthy, Kevin Duh, Rebecca Marvin, Paul McNamee, Jeremy Gwinnup, Tim Anderson, and Philipp Koehn. 2018. [Freezing subnetworks to analyze domain adaptation in neural machine translation](#). In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 124–132, Brussels, Belgium. Association for Computational Linguistics.
- Matthew H. Tong, Adam D. Bickett, Eric M. Christiansen, and Garrison W. Cottrell. 2007. [Learning grammatical structure with echo state networks](#). *Neural Networks*, 20(3):424 – 432. Echo State Networks and Liquid State Machines.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In I. Guyon, U. V. Luxburg, S. Bengio,

H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 5998–6008. Curran Associates, Inc.

Elena Voita, David Talbot, Fedor Moiseev, Rico Senrich, and Ivan Titov. 2019. [Analyzing multi-head self-attention: Specialized heads do the heavy lifting, the rest can be pruned](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5797–5808, Florence, Italy. Association for Computational Linguistics.

Alex Wang, Jan Hula, Patrick Xia, Raghavendra Pappagari, R. Thomas McCoy, Roma Patel, Najoung Kim, Ian Tenney, Yinghui Huang, Katherin Yu, Shuning Jin, Berlin Chen, Benjamin Van Durme, Edouard Grave, Ellie Pavlick, and Samuel R. Bowman. 2018. [Can you tell me how to get past sesame street? sentence-level pretraining beyond language modeling](#).

John Wieting and Douwe Kiela. 2019. [No training required: Exploring random encoders for sentence classification](#). In *International Conference on Learning Representations*.

Kelly Zhang and Samuel Bowman. 2018. [Language modeling teaches you more than translation does: Lessons learned through auxiliary syntactic task analysis](#). In *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 359–361, Brussels, Belgium. Association for Computational Linguistics.

Barret Zoph, Deniz Yuret, Jonathan May, and Kevin Knight. 2016. [Transfer learning for low-resource neural machine translation](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1568–1575, Austin, Texas. Association for Computational Linguistics.

A The value of training

We measured the benefit of training individual components for a short time before fixing them, as opposed to having them untrainable from the beginning. With experiment we show whether training of certain components is more valuable at the beginning of the model training for the final model performance, or they continue to improve and contribute to the final model performance throughout training. We present our results on transformer-big model from Section 3 in Table 10, where we have a model with one frozen component (attention or FFN), where the component is frozen either at initialisation, or after training for few thousand mini-batches.

	Frozen after Epochs	BLEU	Total Epochs
Pretrained Attention layer			
(1.0)	0	22.3	23
(1.1)	2	20.5	71
(1.2)	4	20.6	71
(1.3)	7	22.2	34
Pretrained FFN layer			
(2.0)	0	23.2	26
(2.1)	2	22.6	37
(2.2)	4	22.7	38
(2.3)	7	22.6	21

Table 10: Fixed Attention layer or FNN layer for Turkish-English transformer-big contrasting random initialisation with small amount of pretraining. BLEU is computed on the dev set.

Contrary to our expectations, pretraining the model for a bit before freezing it yields worse overall convergence than having randomly initialised parameters without any pretraining. Pretraining and then freezing either the FFN layer or the attention layer yields similar behaviour. When pretraining is performed for 2 or 4 epochs, we noticed a that the model starts off a bit better than the randomly initialised one, but it quickly catches up in a few epochs. Furthermore convergence time is drastically increased, with small incremental improvements over many more epochs compared to the baseline.

We note that when the component is pretrained for 7 epochs, the final model performance (and convergence time) starts to approach that of a randomly initialised model, but we do not see any improvements over random initialisation. Looking at the attention, model (1.3) performs similar to

the baseline model (1.0), but takes more epochs to converge. When we look at the FFN layer, model (2.3) converges faster than the baseline (2.0), but to a lower BLEU score.