

# BASE Layers: Simplifying Training of Large, Sparse Models

Mike Lewis<sup>1</sup> Shruti Bhosale<sup>1</sup> Tim Dettmers<sup>1,2</sup> Naman Goyal<sup>1</sup> Luke Zettlemoyer<sup>1,2</sup>

## Abstract

We introduce a new balanced assignment of experts (BASE) layer for large language models that greatly simplifies existing high capacity sparse layers. Sparse layers can dramatically improve the efficiency of training and inference by routing each token to specialized expert modules that contain only a small fraction of the model parameters. However, it can be difficult to learn balanced routing functions that make full use of the available experts; existing approaches typically use routing heuristics or auxiliary expert-balancing loss functions. In contrast, we formulate token-to-expert allocation as a linear assignment problem, allowing an optimal assignment in which each expert receives an equal number of tokens. This optimal assignment scheme improves efficiency by guaranteeing balanced compute loads, and also simplifies training by not requiring any new hyperparameters or auxiliary losses. Code is publicly released.<sup>1</sup>

## 1. Introduction

Sparse expert models enable sparse computation by spreading model capacity across a set of experts, while ensuring that only a small subset of the experts are used for each input (Shazeer et al., 2017; Lepikhin et al., 2020; Fedus et al., 2021). Sparse models can often realize the strong performance gains that come with training very large models, while also alleviating much of the associated computational, financial and environmental costs (Strubell et al., 2019). However, such models are notoriously difficult to train; the experts must be carefully balanced so that they can specialize to different parts of the input space. In this paper, we present a simple, efficient, and performant method for expert-based sparsity in language models, built around the use of a linear assignment algorithm to explicitly balance

<sup>1</sup>Facebook AI Research <sup>2</sup>University of Washington. Correspondence to: Mike Lewis <mikelewis@fb.com>.

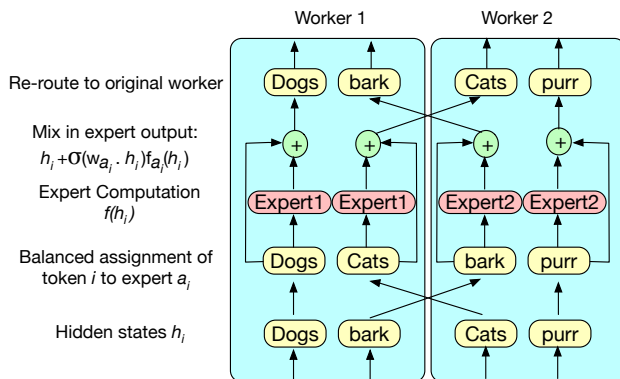


Figure 1. Overview of a BASE layer. Each worker contains a separate *expert* module. During training, we compute a balanced assignment of tokens such that each worker sends an equal number of tokens to each expert. By softly mixing in the expert module, experts can learn to specialize for particular types of tokens.

the assignment of tokens to experts during training.

The mostly widely used Sparse Expert models are mixtures of experts (MoE) models (Shazeer et al., 2017; Lepikhin et al., 2020) that learn a gating function to route each token to a few experts, which creates a challenging, discrete latent variable learning problem. In practice, carefully tuning and the introduction of extra loss functions with new hyperparameters is required to avoid imbalanced or degenerate experts. Recently, the Switch transformer (Fedus et al., 2021) simplified the framework by routing tokens to only a single expert, improving stability and efficiency overall but again using custom auxiliary losses that require tuning, and requiring capacity factors to prevent too many tokens being assigned to a single expert. We show that it is possible to go even further. We also assign a single expert per token but are the first to algorithmically balance the assignment with no extra model modifications, providing more formal guarantees of balanced compute while simplifying both the implementation and optimization.

We introduce a simple and effective solution for routing tokens to experts during training, which we use to estimate a new Balanced Assignment of Sparse Experts (BASE) layer. To ensure balanced routing in the BASE layer, we formulate a linear assignment problem that maximizes token-expert affinities while ensuring that each expert receives an equal

number of tokens. This approach ensures that the assignment will be balanced, and therefore each expert will operate at maximum capacity, while also eliminating load-balancing loss functions and capacity factors from previous work. We also show how to learn expert specialization by using a modified residual connection that softly mixes in each expert contribution—again without requiring an additional loss term or routing tokens to multiple experts. While computing balanced assignments incurs non-trivial overhead, we find that using even a single large BASE layer is remarkably effective—reduced expert communication produces faster gradient computations—and that performance increases as more BASE layers are added, providing an overall favorable cost-accuracy tradeoff.

Extensive experiments with models of up to 110B parameters demonstrate large performance gains over standard data and model parallel training strategies. Our approach also matches or exceeds the efficiency and performance of previous sparse expert approaches (Lepikhin et al., 2020; Fedus et al., 2021), when controlling for computation budget, despite its relative simplicity. Taken together, these results demonstrate the first drop-in conditional compute layer that can be easily added to any model with no new hyperparameters or training loss modifications.

## 2. Background: Training with Multiple Workers

NLP has recently become dominated by ever larger language models (Devlin et al., 2018; Lewis et al., 2019; Liu et al., 2019; Radford et al., 2019; Raffel et al., 2019). Training large language models would take infeasibly long on any existing single device, with many models trained for thousands of GPU-days (Brown et al., 2020). Instead, it is standard to distribute computation over multiple workers. We briefly review the main existing strategies.

### 2.1. Dense Models

In dense models, every parameter is used in processing every input. Training is distributed over multiple workers using data parallelism or model parallelism.

**Data Parallel Training** In data parallel training, multiple workers maintain a copy of the same model. Each worker runs the model on a different subset of the training batch, then gradients are communicated and all workers perform the same update. This approach increases the number of examples processed per second, and only requires a single communication step between workers per update. However, the maximum model size that can be trained is bounded by the memory of a single worker device—limiting models to roughly 1.5B parameters in our setup.

**Model Parallel Training** Model parallel training allows models to be larger than can be run on a single worker (Shoeybi et al., 2019), by distributing the compute for each input over multiple workers. Model parameters are also distributed over workers, which then communicate with each other while processing each input. Given a fixed number of workers, using model parallel training will reduce the amount of compute available for data parallelism, and correspondingly also the number of examples processed per second.

### 2.2. Sparse Expert Layers

Sparse models differ from dense models in only using a small subset of their parameters on any given input. Recent work has explored adding capacity to language models by adding sparse expert layers (Shazeer et al., 2017; Lepikhin et al., 2020; Fedus et al., 2021). During inference, before an expert layer, each token is assigned and routed to a small subset of the workers. The workers then apply a token-wise operation, using parameters that are not shared across other workers. The resulting representation is then returned to the original worker, to continue the forward pass.

During training, this results in four routing steps per expert layer—before and after each expert layer, in both the forward and backward pass. These communication steps can add significantly to the training cost, as workers can idle while waiting for communication to complete.

Balancing of experts, so that each processes a roughly equal proportion of tokens, is crucial for several reasons. If one expert is assigned too many tokens, the worker could run out of memory. Additionally, the expert layer processing speed is limited by the slowest worker; imbalanced assignment slows down training. Furthermore, the parameters of rarely used experts are likely to be less well trained, which may reduce performance.

Previous work has achieved balancing by adding a new term in the loss function that explicitly encourages balancing—this loss term must be carefully weighted so that it does not overwhelm primary loss (Lepikhin et al., 2020; Fedus et al., 2021). However, such a loss does not guarantee balancing. Stable training also requires additional measures such as enforcing hard upper limits on the number of tokens processed by each expert after which the rest are simply ignored (Shazeer et al., 2017). This approach can be inefficient, as some workers are underutilized, and many tokens are unprocessed by the layer.

## 3. BASE Layers

BASE layers achieve balanced assignment of tokens to experts through a three stage process. Firstly, we compute the score for assigning each token representation to each

```

1 def base_layer(features, expert_centroids, expert_id, expert_network):
2     # Send each token to a random worker, by sorting in a random order
3     shuffle_sort = random_permutation(len(features))
4     shuffled_features = all2all(features[shuffle_sort])
5     # Compute which token goes to which expert
6     token_expert_affinities = shuffled_features @ expert_centroids.T
7     sort_by_expert = balanced_assignment(token_expert_affinities)
8     # Swap these tokens for the right ones for our expert
9     routed_features = all2all(shuffled_features[sort_by_expert])
10    # Mix in the expert network based on how appropriate it is for these tokens
11    alpha = torch.sigmoid(routed_features @ self.expert_centroids[expert_id])
12    routed_features += alpha * expert_network(routed_features)
13    # Undo routing and balanced assignment
14    shuffled_features = all2all(routed_features)[inverse_sort(sort_by_expert)]
15    # Return to original worker and ordering
16    return all2all(shuffled_features)[inverse_sort(shuffle_sort)]

```

Figure 2. Implementation of a BASE layer, with  $E$  experts and an input sequence of  $T$  features. Here, *all\_to\_all* routes the  $t$ th row of its input to the  $\lfloor \frac{tE}{T} \rfloor$ th worker. *balanced\_assignment* takes a matrix of size  $T \times E$  and returns an  $T$ -dimensional vector that can be used to sort tokens by their assigned expert index.

expert, compute a balanced assignment maximizing these scores, then route the token features to an expert. Secondly, we compute a position-wise expert function, and compute a weighted sum of the layers input and output. Finally, we return the output to the original worker. Figure 2 shows overall pseudo code for the approach.

### 3.1. Parameterization

BASE layers contain  $E$  experts, each defined by a position-wise function  $f_e(\cdot)$  and an expert embedding  $w_e \in R^D$ , where  $D$  is the model dimension. In practice, we parameterize  $f_e(\cdot)$  using a stack of residual feedforward layers. Given a token  $h_t$  at timestep  $t$  in a sequence of tokens  $0..T$ , and token-to-expert assignment index  $a_t \in 0..E$ , the network returns the following value:

$$\sigma(h_t \cdot w_{a_t})f_{a_t}(h_t) + h_t, \tag{1}$$

If the network  $f_{a_t}$  is able to improve the representation of  $h_t$ , by lowering the loss of the final prediction for that token, then gradient descent will increase the value of  $h_t \cdot w_{a_t}$ . Conversely, if the expert network is unhelpful, then the  $h_t \cdot w_{a_t}$  will receive a negative gradient. Consequently, an expert  $e$  can learn to specialize for particular types of tokens by adjusting  $w_e$  to be close to similar token representations where  $f_e(\cdot)$  is most beneficial.

### 3.2. Token to Expert Assignment

We assign tokens to experts using different methods during training and testing. During training, we maximize model throughput by assigning an equal number of tokens to each expert. At test time, we simply assign each token to its highest scoring expert.

#### 3.2.1. ASSIGNMENT DURING TRAINING

During training, we assign an equal number of tokens to each expert, so that each worker is fully utilized and each worker takes about the same time to finish its assigned load.

Each token  $t$  is assigned to an expert  $a_t$ , aiming to maximize the token-expert affinities under the constraints that each expert is assigned the same number of tokens.

**Linear Assignment Problem** Formally, we solve the following linear assignment problem. Given  $T$  tokens with representations  $h_t$  and  $E$  experts with embeddings  $w_e$ , we assign each token to an expert via the assignment index  $a_t \in 0..E$ :

$$\begin{aligned}
 & \text{maximize } \sum_t h_t \cdot w_{a_t} \\
 & \text{subject to } \forall e \sum_{t=0}^T \mathbb{1}_{a_t=e} = \frac{T}{E}
 \end{aligned} \tag{2}$$

Numerous algorithms exist for this problem. We use the auction algorithm described in Bertsekas (1992), which is more easily parallelizable on GPUs than the Hungarian Algorithm (Kuhn, 1955). Pseudo-code is given in the Appendix.

**Sharding** Computing the optimal assignment for all tokens across all workers is expensive, so we distribute the computation across multiple workers. We decompose the assignment problem of all  $ET$  tokens across all workers into  $E$  smaller problems using  $T$  tokens. This decomposition can be implemented by each worker solving an assignment problem over its own input batch. Each worker then sends  $T/E$  tokens to each other worker, with an *all2all* operation.

**Shuffling** Tokens within each worker’s training sequence are highly correlated with each other; for example they will normally be part of the same domain. These correlations may make it difficult for experts to specialize for particular domains. We therefore add an additional *random* routing step, where each worker first sends an equal number of each tokens to each other worker *randomly*. Then, each worker solves a linear assignment problem as before with its sample of tokens, and routes these to the correct experts.

### 3.2.2. ASSIGNMENT DURING TESTING

At test time, it is not possible to use the assignment strategy described in §3.2.1, as balancing the assignment leaks information about tokens in the future context. Instead, we simply greedily assign the one best expert. While unbalanced assignments are less efficient, during inference memory costs are greatly reduced due to not needing to store gradients, activations and optimizer states. In practice, we show that our approach naturally learns a reasonably balanced assignment during training (§5.1).

### 3.3. Gradient Clipping

A common practice in training deep language models is to scale gradients if their  $l_2$  norm is greater than a threshold. All workers must compute the same norm, or else scaled gradients for shared parameters will be inconsistent across workers. To avoid additional communication steps to compute norms globally across all expert parameters, we simply compute the gradient norms locally based only on the shared parameters, but rescale all gradients.

## 4. Experiments

### 4.1. Experimental Setup

**Task** We focus our experiments on language modelling, as recent work such as GPT3 (Brown et al., 2020) offers perhaps the clearest demonstration in machine learning of the power of large scale models.

**Metrics** We focus exclusively on comparing *compute efficiency*, which we define as the best model performance (here, perplexity) that can be achieved by training with a given number of GPUs and wall-clock time. This metric is different from other commonly used metrics, such as *sample efficiency* (which measures the number of tokens the model trains on, but not the cost of processing samples) or *FLOP-efficiency* (which measures the number of floating-point operations performed during training, but does not account for communication costs). As plentiful data is available for training language models, but computation is expensive, we believe that compute efficiency best captures the constraints of real world training. Therefore, we compare models using

a fixed number of GPUs for the same runtime.

**Training Hyperparameters** We train all models for approximately 2.5 days. All models use similar hyperparameters of 2000 warm-up steps, and the Adam optimizer (Kingma & Ba, 2014). We tune learning rates for each model separately, and linearly decay the learning rate during training. Each worker processes two sequences of length 1024, and gradients are accumulated over 8 updates. We clip gradients if their  $l_2$  norm exceeds 0.1 (§3). Learning rates are tuned in the range  $\{0.5, 0.75, 1.0\} \times 10^{-4}$ , taking the highest value that avoids divergence.

**Hardware** Unless otherwise stated, models are trained on 128 32GB V100 GPUs connected with Infiniband.<sup>2</sup>

**Data** We train on a corpus of approximately 100B tokens, comprising the training corpus of RoBERTa (Liu et al., 2019), combined with the English portion of the CC100 corpus (Conneau et al., 2019). We use the byte-pair encoding (Sennrich et al., 2015) from GPT2 (Radford et al., 2019), which has a vocabulary of 51200.

**Model Architectures** We size all models to the maximum size that can process the sequences within GPU memory constraints. All models follow a standard transformer architecture (Vaswani et al., 2017), with a model dimension of 2048, feed-forward hidden states of size 8096 and 24 Transformer Decoder blocks. We use 16 attention heads, ReLU activation functions and no dropout. LayerNorm (Ba et al., 2016) is applied to the inputs of each residual block (Xiong et al., 2020) and to the outputs of the transformer.

**BASE layer architecture** We implement the BASE layer as a stack of feedforward blocks. Each block follows the standard transformer structure: layer normalization, a projection to 4 times the input dimension, a ReLU nonlinearity, a projection to the input dimension, and a residual connection to the block input. We vary the number of BASE layers;  $\text{BASE} \times N$  uses a BASE layer after each of the  $\lfloor \frac{L}{N+1} \rfloor \dots \lfloor \frac{NL}{N+1} \rfloor$ th transformer layers. When using multiple BASE layers, we reduce their size to keep the total number of parameters roughly constant;  $\text{BASE} \times N$  use  $\lfloor \frac{10}{N} \rfloor$  sublayers, for a total of roughly 44B parameters. We use one expert per GPU per BASE layer.

### 4.2. Comparison with Dense Models

We first compare with dense models, in which all parameters are shared across all workers. We compare with data parallel and model parallel training, using the intra-layer model

<sup>2</sup>As communication between workers is a significant overhead for model parallel and sparse expert approaches, it is possible that different results would be achieved on other networking hardware.




Figure 3. Comparing BASE layers with dense model training, using different numbers of GPUs. There is a clear trend of increased model sizes being more effective with larger compute budgets. BASE layers show strong performance at all the compute budgets we consider.

parallelism approach introduced in Shoeybi et al. (2019). Our data parallel baseline contains 1.5B parameters, and the 2-way and 4-way model parallel baselines contain roughly 3B and 6B parameters respectively. We use three different compute budgets: 8, 32 and 128 GPUs for approximately 2.5 days.

Results are shown in Figure 3. We generally find that larger models perform better with higher compute budgets, and that simple data parallel training performs best at the smallest compute budget. With larger compute budgets, BASE layers outperform both data parallel and model parallel training by a wide margin.

Relatively high compute budgets are required before model parallelism outperforms data parallel training, with the first gains appearing after training on 128 GPUs for 2 days. This is partly due to model parallel training requiring a reduced batch size given the same computational resources.

In contrast, BASE layers match the performance of data parallel training on our 8 GPU experiments, and achieve increasingly large gains in higher compute regimes.

### 4.3. Comparison with Sparse Experts Models

We also compare performance with our re-implementations of two recent sparse layer methods: Sparsely Gated Mixtures of Experts (Shazeer et al., 2017; Lepikhin et al., 2020) and Switch (Fedus et al., 2021). The primary difference between these approaches is that a Sparsely Gated MoE layer routes tokens to multiple experts (top-2 experts in our experiments), whereas a Switch layer routes tokens to a single expert. We set the weight associated with the load balancing loss to 0.01 in our experiments, and set the capacity factor for Sparsely Gated MoE and Switch layers to 2.0 and 1.0 respectively. Following previous work, we replace every other shared feed-forward layer in the Transformer architecture with a Sparsely Gated MoE or Switch layer, unless




Figure 4. Comparison with other Sparse Experts approaches. Despite its simplicity, BASE achieves strong performance relative to Sparsely Gated MoE models and Switch transformers.

otherwise specified. With 128 experts in each expert layer, our Sparsely Gated MoE and Switch models have 52.5B parameters (1B shared parameters) each, while our BASE model has 44.4B parameters (1.3B shared parameters).

As in Fedus et al. (2021), we find that Switch computes more updates per second than Sparsely Gated MoE (see Table 2). However, we find that Sparsely Gated MoE is more compute efficient in our experiments as shown in Figure 4.

A comparison with BASE is also shown in Figure 4. Despite its simplicity, BASE achieves similar performance to the Sparsely Gated MoE model and converges to a better validation perplexity than Switch. This result suggests that algorithmic load balancing is a competitive alternative to load balancing loss functions, and that even a single expert




Figure 5. Comparison of different sizes of BASE layers, by changing the ratio of parameters allocated to shared vs. expert layers.

layer can be highly effective.

#### 4.4. Ablations

Results in Section 4 show that BASE layers match or exceed the compute-efficiency of previous dense and sparse approaches. To better understand these results, we analyze key design decisions in our model in more detail.

**BASE Layer Size** A key choice in any sparse experts model is the allocation of capacity to shared components versus experts. We experiment with adjusting the number of sublayers in each expert, and scale the number of shared layers accordingly to maximize GPU usage.

We test three versions:

- **Small Expert:** 1.5B shared parameters, 135M parameters per expert, 18.8B total parameters
- **Standard Expert:** 1.3B shared parameters, 335M parameters per expert, 44B total parameters
- **Large Expert:** 840M shared parameters, 911M parameters per expert, 117B total parameters

Figure 5 shows that good performance can be achieved with all sizes, indicating that this choice needs little tuning.

**BASE Layer Position** We also consider the most effective place in a model to insert BASE layers into a transformer with  $L$  layers. We test three configurations:

- **BASE:** After the  $\frac{L}{2}$ th layer, as in our other experiments.

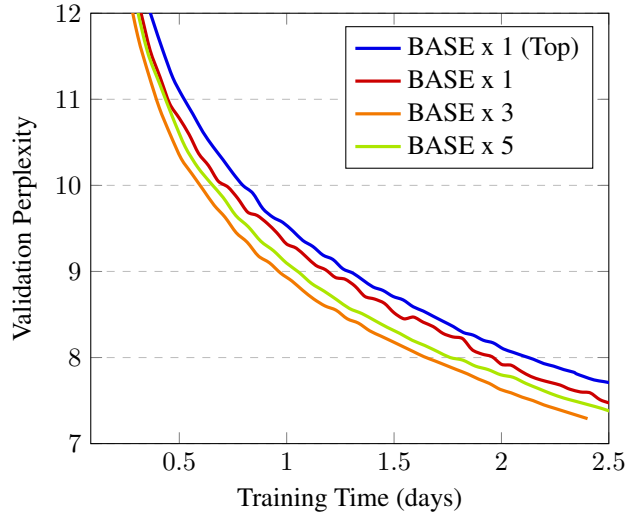


Figure 6. Comparison of different numbers and positions of BASE layers. The best performance is achieved by interleaving 3 BASE layers throughout the transformer stack.

- **BASE Top:** After the  $L$ th layer, acting as a classifier.
- **BASE  $\times N$ :** Using  $N$  BASE layers of  $\frac{1}{N}$  the size, after layers  $\frac{L}{N+1} \dots \frac{NL}{N+1}$ th layers of the transformer.

Figure 6 compares results for different configurations. We find similar performance from three different placements of BASE, suggesting a reasonable level of robustness. In particular, the strong performance of BASE Top may enable it to be used on top of pre-trained language models to further increase their capacity.

#### Comparison of Routing Method with Sparsely Gated MoE

Our approach differs from previous work on sparse experts in both the architecture and assignment method. To more carefully analyse the benefits of our routing method, we compare with an implementation of Sparsely Gated MoE that uses a more similar architecture to ours: a single, large expert midway through the transformer stack.

Results are shown in Figure 7. Sparsely Gated MoE performs less well in this setting. Sparsely Gated MoE benefits from interleaving expert layers with shared layers, and a single Sparsely Gated MoE layer with deep experts works less well than BASE. Future work should explore more efficient approximate routing schemes for BASE layers, to enable potential compute efficiency gains from interleaving expert and shared layers.




Figure 7. Comparing routing strategies using similar architectures. Here, all models use a single large expert at layer  $L/2$ . BASE maintains strong performance in this setting, which reduces the communication overhead between workers, and may be advantageous with less efficient networking.

## 5. Analysis

We also report further experiments that provide more qualitative analyses of overall model behavior with BASE layers.

### 5.1. Expert Balancing

A key difference between our model and other recent proposals is that we algorithmically balance token/expert assignments during training, instead of adding an additional loss function to balance assignments. However, both use greedy assignments at test time.

We investigate whether our model learns a balanced assignment without an explicit balancing loss. Figure 8 shows the percentage of tokens assigned to each expert, sorted from most used to least used. Unsurprisingly, the top-1 assignment from BASE is less balanced than those from models with explicit balancing loss terms. However it is notably *more* balanced than the 2<sup>nd</sup> expert in the Sparsely Gated MoE model, and confirms that reasonably balanced assignments can be learnt without balancing losses.

### 5.2. Expert Specialization

We also analyse how experts learn to specialize. Observing sample passages, we find that many assignment decisions appear to depend primarily on very local syntactic information. In particular, we found that the token input at timestep  $t$  is often highly indicative of the expert assigned at time  $t$ .




Figure 8. Expert Balancing in different Sparse Expert approaches across 128 experts, as measured on the validation set. Results for Sparsely Gated MoE and Switch are an average across all expert layers. BASE layers learn a reasonably balanced routing with no auxiliary balancing loss.

Table 1 shows the most frequent previous input token when selected experts are chosen. We see clusters corresponding to quantities (5), numbers (42), possessives (125), subword fragments (101), and clusters of related verbs (72, 74, 126), nouns (23,27,36,43,76,84,96,98,105) and adjectives (9,81). These tokens may tend to have similar distributions over next tokens. This analysis suggests the model primarily assigns experts based on fairly superficial signals, and may motivate even simpler techniques for future work.

### 5.3. Efficiency

While we focus on evaluating the compute efficiency of models, we note that there are substantial differences in the speed at which models process tokens. Table 2 shows the number of tokens processed per second by different models during training, using 128 GPUs. Simple data parallel training is unsurprisingly the fastest, but BASE layers compute updates faster than other approaches due to reduced communication between workers. For the same compute efficiency, models which process tokens more slowly are more sample efficient, and may be preferable in lower data regimes.

## 6. Related Work

Shazeer et al. (2017); Lepikhin et al. (2020) introduce sparsely gated mixtures of experts layers, demonstrating how large sparse models can be trained efficiently by routing inputs to appropriate specialist workers. Fedus et al. (2021) show the design can be simplified by routing tokens

Expert	Top 5 Preceding Tokens
5	<i>year, years, billion, million, tonnes</i>
8	<i>people, who, Man, everyone, one</i>
9	<i>electronic, local, public, national, outdoor</i>
23	<i>funding, budget, benefits, pressure, price</i>
27	<i>Mustang, State, Center, ation, Grande</i>
34	<i>to, will, should, it, may</i>
36	<i>business, bank, financial, science, school</i>
42	<i>two, 50, 1, 80, 000</i>
43	<i>Bank, Development, ., Construction, Plant</i>
62	<i>work, started, involved, working, launched</i>
72	<i>is, was, be, been, were</i>
74	<i>going, go, come, back, return</i>
76	<i>painting, case, song, statement, discussion</i>
81	<i>new, major, bad, larger, grand</i>
84	<i>Ret, Inspect, Pl, Pos, Architect</i>
96	<i>US, UNESCO, government, state, UN</i>
98	<i>waiver, procedures, warrant, status, loans</i>
101	<i>B, T, W, H, k</i>
105	<i>app, Windows, Microsoft, board, 10</i>
125	<i>his, 's, its, their, our</i>
126	<i>said, says, means, noting, out</i>

Table 1. Most frequent previous words for selected experts, showing that some experts assignment decisions are made based on very local contexts. For many other experts, the assignment decision depends on longer context, and is harder to visualize.

to only a single worker. We further simplify the framework, by eliminating balancing loss functions, and showing the effectiveness of using only a single expert layer.

Sparse training is a line of work where traditional architectures are trained with sparse instead of dense layers and the number of parameters allowed during training is restricted to a percentage of the dense layers (Dettmers & Zettlemoyer, 2019; Evci et al., 2020; Mostafa & Wang, 2019). Unlike our approach, these networks have fine-grained sparsity patterns which reduce overall FLOPS but make it difficult to achieve runtime benefits on modern accelerators like GPUs, which require contiguous memory segments for efficient processing. Since experts consist of sizable contiguous memory segments, our approach can utilize GPUs effectively.

Perhaps the most common use of sparse layers is in adding language-specific layers to machine-translation systems (Bapna et al., 2019; Fan et al., 2020), or task-specific layers to pre-trained language models (Houlsby et al., 2019). Here, the expert assignment problem is hard coded, based on the task being solved or the language being translated. We instead explore learnable routing, which is applicable to problems where such structure is not available.

Other papers have explored alternative methods for adding

Model	Tokens per Second
Data Parallel	600k
Model Parallel $\times 2$	224k
Sparsely Gated MoE	292k
Switch	469k
BASE	545k
BASE $\times 2$	475k

Table 2. Number of tokens processed per second during training by different models. BASE computes updates faster than other approaches that divide models over multiple workers, due to reduced communication overheads. This allows a 43B parameter model to be trained at 90% of the speed of a 1.5B data parallel baseline.

very high capacity layers to neural language models. For example, Lample et al. (2019) introduce a large memory layer that supports efficient sparse queries. Khandelwal et al. (2019) show large gains from augmenting a language model with a nearest neighbour classifier over the training set, which recent work has also shown is applicable to machine translation (Khandelwal et al., 2020).

An orthogonal strand of work has improved the efficiency of transformer attention mechanisms, often by making them sparse (Child et al., 2019; Correia et al., 2019; Roy et al., 2020). We instead develop a sparse version of the other major component of the transformer, the feed forward network.

## 7. Conclusion

We introduced a simple sparse BASE layer, which can be used to increase the capacity of any neural model, with little increase in training cost or complexity. We demonstrate strong performance relative to both dense models and previously proposed sparse models. Future work should explore more efficient implementations for computing balanced assignments, to further improve training speed.

## References

- Ba, J. L., Kiros, J. R., and Hinton, G. E. Layer normalization. *arXiv preprint arXiv:1607.06450*, 2016.
- Bapna, A., Arivazhagan, N., and Firat, O. Simple, scalable adaptation for neural machine translation. *arXiv preprint arXiv:1909.08478*, 2019.
- Bertsekas, D. P. Auction algorithms for network flow problems: A tutorial introduction. *Computational optimization and applications*, 1(1):7–66, 1992.
- Brown, T. B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., et al. Language models are few-shot learners. *arXiv preprint arXiv:2005.14165*, 2020.



- Child, R., Gray, S., Radford, A., and Sutskever, I. Generating long sequences with sparse transformers. *arXiv preprint arXiv:1904.10509*, 2019.
- Conneau, A., Khandelwal, K., Goyal, N., Chaudhary, V., Wenzek, G., Guzmán, F., Grave, E., Ott, M., Zettlemoyer, L., and Stoyanov, V. Unsupervised cross-lingual representation learning at scale. *arXiv preprint arXiv:1911.02116*, 2019.
- Correia, G. M., Niculae, V., and Martins, A. F. Adaptively sparse transformers. *arXiv preprint arXiv:1909.00015*, 2019.
- Dettmers, T. and Zettlemoyer, L. Sparse networks from scratch: Faster training without losing performance. *arXiv preprint arXiv:1907.04840*, 2019.
- Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
- Evcı, U., Gale, T., Menick, J., Castro, P. S., and Elsen, E. Rigging the lottery: Making all tickets winners. In *International Conference on Machine Learning*, pp. 2943–2952. PMLR, 2020.
- Fan, A., Bhosale, S., Schwenk, H., Ma, Z., El-Kishky, A., Goyal, S., Baines, M., Celebi, O., Wenzek, G., Chaudhary, V., et al. Beyond english-centric multilingual machine translation. *arXiv preprint arXiv:2010.11125*, 2020.
- Fedus, W., Zoph, B., and Shazeer, N. Switch transformers: Scaling to trillion parameter models with simple and efficient sparsity. *arXiv preprint arXiv:2101.03961*, 2021.
- Houlsby, N., Giurghi, A., Jastrzebski, S., Morrone, B., De Laroussilhe, Q., Gesmundo, A., Attariyan, M., and Gelly, S. Parameter-efficient transfer learning for nlp. In *International Conference on Machine Learning*, pp. 2790–2799. PMLR, 2019.
- Khandelwal, U., Levy, O., Jurafsky, D., Zettlemoyer, L., and Lewis, M. Generalization through memorization: Nearest neighbor language models. *arXiv preprint arXiv:1911.00172*, 2019.
- Khandelwal, U., Fan, A., Jurafsky, D., Zettlemoyer, L., and Lewis, M. Nearest neighbor machine translation. *arXiv preprint arXiv:2010.00710*, 2020.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Kuhn, H. W. The hungarian method for the assignment problem. *Naval research logistics quarterly*, 2(1-2):83–97, 1955.
- Lample, G., Sablayrolles, A., Ranzato, M., Denoyer, L., and Jégou, H. Large memory layers with product keys. *arXiv preprint arXiv:1907.05242*, 2019.
- Lepikhin, D., Lee, H., Xu, Y., Chen, D., Firat, O., Huang, Y., Krikun, M., Shazeer, N., and Chen, Z. Gshard: Scaling giant models with conditional computation and automatic sharding. *arXiv preprint arXiv:2006.16668*, 2020.
- Lewis, M., Liu, Y., Goyal, N., Ghazvininejad, M., Mohamed, A., Levy, O., Stoyanov, V., and Zettlemoyer, L. Bart: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. *arXiv preprint arXiv:1910.13461*, 2019.
- Liu, Y., Ott, M., Goyal, N., Du, J., Joshi, M., Chen, D., Levy, O., Lewis, M., Zettlemoyer, L., and Stoyanov, V. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*, 2019.
- Mostafa, H. and Wang, X. Parameter efficient training of deep convolutional neural networks by dynamic sparse reparameterization. In *International Conference on Machine Learning*, pp. 4646–4655. PMLR, 2019.
- Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., and Sutskever, I. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019.
- Raffel, C., Shazeer, N., Roberts, A., Lee, K., Narang, S., Matena, M., Zhou, Y., Li, W., and Liu, P. J. Exploring the limits of transfer learning with a unified text-to-text transformer. *arXiv preprint arXiv:1910.10683*, 2019.
- Roy, A., Saffar, M., Vaswani, A., and Grangier, D. Efficient content-based sparse attention with routing transformers. *arXiv preprint arXiv:2003.05997*, 2020.
- Sennrich, R., Haddow, B., and Birch, A. Neural machine translation of rare words with subword units. *arXiv preprint arXiv:1508.07909*, 2015.
- Shazeer, N., Mirhoseini, A., Maziarz, K., Davis, A., Le, Q., Hinton, G., and Dean, J. Outrageously large neural networks: The sparsely-gated mixture-of-experts layer. *arXiv preprint arXiv:1701.06538*, 2017.
- Shoeybi, M., Patwary, M., Puri, R., LeGresley, P., Casper, J., and Catanzaro, B. Megatron-lm: Training multi-billion parameter language models using model parallelism. *arXiv preprint arXiv:1909.08053*, 2019.
- Strubell, E., Ganesh, A., and McCallum, A. Energy and policy considerations for deep learning in nlp. *arXiv preprint arXiv:1906.02243*, 2019.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., and Polosukhin, I. Attention is all you need. *arXiv preprint arXiv:1706.03762*, 2017.

Xiong, R., Yang, Y., He, D., Zheng, K., Zheng, S., Xing, C., Zhang, H., Lan, Y., Wang, L., and Liu, T. On layer normalization in the transformer architecture. In *International Conference on Machine Learning*, pp. 10524–10533. PMLR, 2020.