RingFormer: Rethinking Recurrent Transformer with Adaptive Level Signals

Anonymous ACL submission

Abstract

Transformers have achieved great success in effectively processing sequential data such as 003 text. Their architecture consisting of several attention and feedforward blocks can model relations between elements of a sequence in parallel manner, which makes them very efficient to train and effective in sequence modeling. 800 Even though they have shown strong performance in processing sequential data, the size of their parameters is considerably larger when compared to other architectures such as RNN and CNN based models. Therefore, several approaches have explored parameter sharing and recurrence in Transformer models to ad-014 dress their computational demands. However, such methods struggle to maintain high performance compared to the original Transformer model. To address this challenge, we propose our novel approach, RingFormer, which employs one Transformer layer that processes input repeatedly in a circular, ring-like manner, while utilizing low-rank matrices to generate input-dependent level signals. This allows us to reduce the model parameters substantially while maintaining high performance in a variety of tasks such as translation and image classification, as validated in the experiments.

1 Introduction

001

007

017

027

037

041

Transformer models, since their introduction (Vaswani et al., 2017), have dramatically transformed the landscape of deep learning, particularly excelling in tasks involving sequential data such as natural language processing (Brown et al., 2020; Radford et al., 2019) and machine translation (Ott et al., 2018). Not long after their inception, they have also shown strong performance in various other domains such as reinforcement learning (Chen et al., 2021), image classification (Dehghani et al., 2023; Dosovitskiy et al., 2020; Liu et al., 2021), object detection (Carion et al., 2020) and image generation (Jiang et al., 2021; Peebles and

Xie, 2022; Zhang et al., 2022). Their core architecture, characterized by self-attention mechanisms and feedforward neural networks, enables effective handling of long-range dependencies and parallel processing of input sequences. The ability of this architecture to model intricate relationships within data has led to significant breakthroughs, making it a foundation model across many modern largescale AI systems (Anthropic, 2023; Google, 2024; OpenAI et al., 2024; Touvron et al., 2023).

042

043

044

047

048

053

054

056

060

061

062

063

064

065

066

067

068

069

070

071

072

073

074

076

077

078

079

081

However, the impressive capabilities of transformer models come with substantial computational and memory costs (Brown et al., 2020; Dosovitskiy et al., 2020). The standard Transformer architecture consists of multiple layers, each containing millions of parameters that need to be trained and stored. This results in high memory usage and significant computational demands, often requiring specialized hardware. Moreover, deploying these models in resource-constrained environments, such as mobile devices or edge computing scenarios, becomes challenging due to their size and complexity. These limitations have spurred a growing interest in developing more parameter-efficient Transformer architectures (Dehghani et al., 2019; Pires et al., 2023) that can retain their powerful performance while being more accessible and less resource intensive.

In this paper, we introduce a Transformer architecture that recurrently leverages a single shared Transformer block in a novel way by integrating input-dependent level signals at each block iteration, which are shown to be crucial for adapting the shared block to different stages of the model. The level signals are generated by depth-specific low-rank transformations applied to the input in the attention and feedforward layers within the Transformer block. Our RingFormer model can also be viewed as stacking Transformer layers whose parameters combine (1) a set of global parameters shared across all Transformer layers and (2) a set

of local low-rank layer-dependent parameters. This simple design effectively addresses the trade-off between reducing the number of model parameters and limiting the model's capacity to capture complex patterns.

084

100

101

103

104

105

108

110

111

112

113

114

115

116

117

118

119

120

121

123

124

125

126

127

128

129

We validate our model through experiments and analysis on machine translation and image classification. The results of experiments and analysis demonstrate that our model closely replicates the behavior of the original Transformer model, and it performs better against existing parameter-matched recurrence-based Transformer models, underscoring the effectiveness of our approach in maintaining high performance with fewer parameters.

The contributions of this paper are summarized as follows:

- We enhance a recurrent Transformer architecture to significantly reduce the model's parameter count while maintaining high performance.
- We propose novel input-dependent level signals generated in a parameter-efficient way using low-rank matrices to improve the adaptability of a recurrent Transformer model, and show that those signals help the model replicate the behavior of the original model.
- We demonstrate the validity of our approach through careful analysis and ablation studies, and show the effectiveness of our model on tasks such as translation and image classification.

2 Background

2.1 Transformer Architecture

The Transformer architecture (Vaswani et al., 2017) comprises multiple layers of the same structure stacked together, with each layer consisting of two main modules: *Attention* and *Feedforward Network* described in Equations (1) and (2), respectively. Each of these modules is accompanied by residual connections and layer normalization. In addition, to provide information about the position of tokens in the sequence, the Transformer model adds static sinusoidal or learnable positional encodings to the input embeddings. These encodings allow the model to capture the order within a sequence. The following equations describe the mechanism of two main modules:

130
$$Attention(Q, K, V) = softmax(\frac{QK^T}{\sqrt{d}})V$$
(1)

$$FFN(x) = \sigma(xW_{up} + b_{up})W_{down} + b_{down} \quad (2)$$

Here, Q, K, and V are the results of projecting the input vectors through their respective matrices. Attention module can be classified into self-attention (when the Q, K, and V input vectors are the same) and cross-attention (when the Q input vector is different from the K and V input vectors), while the feedforward block consists of up-projection and down-projection transformations with non-linearity function σ between them.

It is well known that Transformer architecture follows a scaling law for both vision tasks and NLP tasks (Dehghani et al., 2023; Hoffmann et al., 2022). This scaling law demonstrates that the performance of Transformer models improves predictably as the model size and computational resources increase. Due to the steep slope of the scaling law, the parameter sizes of Transformer models have continued to grow, leading to significant advancements in their capabilities. However, this growth has also made training and using such massive models increasingly infeasible without substantial GPU resources.

2.2 Related Work

To address the challenge of requiring extensive hardware resources for large Transformer models, researchers have explored various methods to enhance efficiency.

One approach is related to pruning of Transformer model layers, which involves removing less important layers or weights to streamline the model. It was found that many deep layers in large language models are redundant (Gromov et al., 2024), and by pruning up to half of these layers, it was possible to significantly reduce the model size with minimal accuracy degradation.

Another strategy is sharing parameters across different layers or components in Transformers, reducing the model's complexity and memory usage. The Universal Transformers (Dehghani et al., 2019) introduces a model where parameters are shared across layers using a recurrent mechanism with layer-dependent positional encoding, which maintains good performance in various NLP tasks while reducing the number of parameters. People have also proposed sequence and cycle strategies for sharing parameters across layers (Takase and Kiyono, 2021), improving efficiency and performance in tasks like machine translation and speech recog160

161

162

163

164

165

167

168

169

170

171

172

173

174

175

176

177

178

179

180

131

133

267

268

269

270

271

272

273

274

nition. Similarly, Subformer (Reid et al., 2021) and One Wide Feedforward (Pires et al., 2023) investigate partial weight sharing within layers, showing that significant parameter reductions can be achieved with little accuracy sacrifice. These models demonstrate that shared parameters can lead to efficient and effective Transformer architectures.

To investigate recurrence-based models, we performed a layer representation similarity analysis using the common CKA (centered kernel alignment) (Kornblith et al., 2019) method and mean attention distance (MAD) (Dosovitskiy et al., 2020) analysis, and we found that the layer representations and internal attention behavior of the previously proposed fully recurrence-based Transformer model (Dehghani et al., 2019) are considerably different compared to those of the original Transformer model.

We hypothesized that the difference in model behavior, especially in attention module, might be the main cause for the gap in performance, and if we can simulate the behavior of the original model using a recurrent model with adaptive level signals, we can also maintain higher performance. Our proposed methodology is focused on addressing this difference, narrowing the gap of the model behavior, and in turn the model performance.

3 Method

3.1 Overview

In this section, we provide a detailed explanation of our proposed work, covering the specific details about the structure of our model.

The encoder or decoder Transformer-based models consist of several layers with the same structure, where each layer is a combination of sub-layers such as attention and feedforward layers. Those models can be formulated in the following way:

218

181

182

186

187

189

190

192

193

194

195

197

198

199

204

207

208

210

211

212

213

214

215

216

217

(3)
=
$$f(f(...f(f(x, p_1), p_2)), p_{N-1}), p_N)$$

 $F(x) = f_N(f_{N-1}(\dots f_2(f_1(x))))$

219 where N, F, f, x and p_i denote the number of 220 layers, entire encoder (or decoder), each encoder 221 (or decoder) block, input and parameters of each i^{th} 222 layer, respectively. The general formulation of the 223 recurrent Transformer model with level transition 224 functions can be written as below:

$$F(x) = f_N(f_{N-1}(...f_1(x)))$$

$$f_i(x) = f_r(x, g_i(x))$$
(4)

where f_r denotes the recurrent Transformer block and $g_i(x)$ represents a generic level transition function specific for each level. In Universal Transformers (Dehghani et al., 2019), it was shown that using static spatio-temporal positional embeddings can serve as level transition functions for the recurrent Transformer layer and have good model performance. Specifically, in that work, level transition function $g_i(x)$ can be represented as $g_i(x) = x + l(i, x_p)$, where l is a function that returns a positional embedding vector based on level depth i and the position x_p of the vector x, while the i^{th} Transformer block function $f_i(x)$ can be represented as $f_i(x) = f_r(g_i(x))$.

Below, we describe our way of constructing and integrating level transition function $g_i(x)$ to generate adaptive level signals.

3.2 Adaptive Level Signals

To have effective transition between the levels when using recurrent Transformer block, we make $q_i(x)$ directly dependent on the input in the following way: $q_i(x) = M_i \cdot x$, where M is a learnable transformation matrix. Since the main role of level signals is to nudge the input vectors in the right direction, which is an easier task compared to the main input transformation done by the recurrent layer, we hypothesize that making the Mmatrix low-rank while keeping the recurrent layers at full-rank will let us have parameter-efficiency and high performance at the same time. We draw inspiration for such a low-rank matrix construction and its weight initialization from the parameter efficient fine-tuning (PEFT) technique, LoRA (Hu et al., 2021), and decompose M_i into two lowdimensional matrices, A_i and B_i described in Equation 5.

$$M_i = A_i \cdot B_i^T, \ A_i, B_i \in \mathbb{R}^{d \times r} \text{ and } r \ll d$$
 (5)

Since a Transformer layer consists of an attention block and a feedforward block, we generate two distinct signals $g_{Ai}(x)$ and $g_{Fi}(x)$: one for the attention block and the other for the feedforward block, respectively. Additionally, since the Transformer block also has layer normalization applied between the sub-layers, for each level, we allocate unique layer normalization in the attention and feedforward layers. This provides extra input adaptation while only slightly increasing the number of total parameters in the model.



Figure 1: Overview of (a) vanilla Transformer (Vaswani et al., 2017), (b) recurrent Transformer (Dehghani et al., 2019) and (d) our RingFormer architecture. The Transformer block represents either encoder or decoder. In the RingFormer model, a single block consisting of attention and feedforward modules is used iteratively, with each sub-modules having unique layer normalization and level signals. (c) illustration of the low-rank matrices representing the level functions, where M_{down} down-projects the input to a lower dimensional space and M_{up} up-projects back to the original space.

3.2.1 Attention Block

275

276

279

For the attention mechanism, which calculates relevance between elements of a sequence using three projection matrices (query, key, and value), we generate level signals for each of those projections using separate low-rank matrices. We integrate signals after the projection of the input vector xby W_Q , W_K and W_V matrices (shared across the levels) in the following way:

$$Q_i = W_Q \cdot x + g_{A_{Q_i}}(x),$$

$$K_i = W_K \cdot x + g_{A_{K_i}}(x),$$

$$V_i = W_V \cdot x + g_{A_{V_i}}(x),$$
(6)

where $g_{A_{Qi}}(x) = M_{Qi} \cdot x$, $g_{A_{Ki}}(x) = M_{Ki} \cdot x$, $g_{A_{Vi}}(x) = M_{Vi} \cdot x$. By incorporating the level 286 functions separately for Q, K, and V, we enable fine-grained control over depth-dependent modifications to each component of the attention mecha-289 nism. Also, adding the level signals in this manner avoids direct input changes to the main recurrent projections, which was found to be beneficial in our experiments. This can be because such a direct input change can interfere with the learning pro-294 cess of the recurrent layers in the attention module, which needs to solely focus on modeling effective communication between tokens. 297

3.2.2 Feedforward Block

For feedforward block, the projection of input to intermediate vector of this module requires relatively large number of parameters. Furthermore, there have been various explorations regarding the role of feedforward network in Transformers. One such study (Geva et al., 2021) argues that the feedforward network can be interpreted as a key-value memory pair, where the matrix of the first linear layer is involved in the coefficients of input factors, and the matrix of the second linear layer relates to information about the training corpus. Considering parameter-efficiency and the previous finding, in our approach, for the feedforward network, we add signals before projecting the input using the upprojection layer to guide the coefficient formation of the input in the following way:

298

299

300

301

302

303

305

306

307

308

309

310

311

312

313

314

315

316

317

318

319

320

321

322

323

324

325

326

$$FFN(x) = \sigma((x + g_{Fi}(x))W_{up})W_{down} \quad (7)$$

where $g_{Fi}(x) = M_{Fi} \cdot x$, the function σ is a nonlinear function such as GELU (Hendrycks and Gimpel, 2023), and the bias terms were omitted for brevity.

In encoder-decoder models, we iteratively reuse a single Transformer block consisting of attention and feedforward sub-blocks in the encoder, while the decoder utilizes a separate shared Transformer block with cross-attention, which is also shared across layers. For the recurrent cross-attention module inside the decoder, we do not incorporate level signals, as cross-attention takes on the output of the already level-adapted attention module as query and the representations from the encoder as key and value, which do not require additional adaptation at different levels. The overall structure of our model is illustrated in Figure 1.

Experiments 4

327

332

334

335

337

341

342

343

344

345

351

367

371

374

We evaluate the performance of our RingFormer model and baseline models across two tasks: machine translation and image classification.

For the baseline models, we choose the vanilla Transformer model (Dosovitskiy et al., 2020; Vaswani et al., 2017), one recurrent transformer model, Universal Transformer (Dehghani et al., 2019), and one partially recurrent model, One Wide Feed Forward model (Pires et al., 2023), with specific adaptations to the corresponding tasks described below.

4.1 **Experimental Details**

In this section, we provide detailed description of each downstream task to facilitate the reproduction of our experimental results. For all of our models, the rank of the decomposed matrices for the level signals is fixed at the input hidden dimension divided by 16. We perform ablations for the different ranks and show the results in Table 4.

Translation Transformer model is firstly proposed in translation task (Vaswani et al., 2017). Thus, we also test our model on the translation task, with two model sizes shown in Table 1. We 356 train all models on WMT-14 (Bojar et al., 2014) German-English dataset which consists of 4.5M 358 pairs of sentences. For evaluation, we calculate BLEU score (Papineni et al., 2002) for the WMT-14 German-English test set and we employ BiBERT vocabulary with bi-lingual tokenizer with vocab size equal to 52K (Xu et al., 2021). We set the number of layer (iteration), batch size and entire training step as 6, 512 and 830K for base and large setting on two A100 80GB GPUs, respectively. In the training session, we used Adam (Kingma and Ba, 2017) optimizer with a cosine learning rate scheduler having 40K of warm-up steps. Also, we used GELU (Hendrycks and Gimpel, 2023) as activation function for all models.

The main model hyperparameters and experiment results are given in Table 1, where we report the parameter size except those parameters

in the encoder, decoder and vocabulary head because their count is the same for all models having the same hidden input dimension and feedforward block dimension. For base size models, the encoder and decoder embedding layer each consists of 26.62M parameters, and vocabulary head contains 26.67M parameters; for large size models, 53.24M for encoder and decoder embedding layer, 53.30M for vocabulary head.

375

376

377

378

379

380

381

382

385

387

388

389

391

393

394

395

397

398

399

400

401

402

403

404

405

406

407

408

409

410

411

412

413

414

415

416

417

418

419

420

421

422

423

424

Image Classification As the ViT (Dosovitskiy et al., 2020) model became very prevalent in the vision domain, especially in image classification, we decided to test our model and other baseline models on this task. The models are adjusted to have only encoder layers, which take image patches with a class token attached as an input, and perform the prediction using the hidden state of the class token from the last layer. For the Vision Transformer (ViT) model (Dosovitskiy et al., 2020), we stick to the original architecture, while for the Universal Transformer (Dehghani et al., 2019), static sinusodial spatio-temporal positional embeddings are used as level transition function between the levels in the encoder. For the One Wide Feed Forward model (Pires et al., 2023), the feedforward layer is shared across the levels, while the attention layer parameters are distinct for each level.

We first train smaller models on a subset of the original ImageNet-1K dataset (Deng et al., 2009) for 100 epochs. We randomly chose 100 classes with the total number of 100K training samples (1K per each class) from the original training set, and 5K testing samples (50 per each class) from the original validation set. For easy referencing, we call that subset *ImageNet-small*. As the size of the dataset is relatively small, we decided to train models having only 6 layers / iterations (in the case of recurrent models, we say iterations or levels instead of layers). For bigger size models with 12 layers / iterations, we trained on the whole ImageNet-1K for 50 epochs due to limited resources.

The additional training and ImageNet-small dataset details are given in Appendix A.1 and A.3. The model hyperparameters, parameter size and experiment results on ImageNet-small and ImageNet-1K are given in Table 2 and 3.

4.2 Experimental Results

Translation The details of experimental results on translation are presented in Table 1. Our Ring-Former model achieves competitive performance



Figure 2: Representation Similarity Analysis using CKA (centered kernel alignment) (Kornblith et al., 2019) for the base-size models trained on the translation task. The figures on the upper row are for the encoder part. The figures on the lower row means are for the decoder part. All models have 6 number of layers / iterations. The values on the figures are between 0 and 1, where higher values indicate more similarity of layers between models.

6

model	H/FF	P^*	$\big \ BLEU \uparrow$
Vanilla Transformer	512 / 2048	44.05M	30.46
One Wide FFN	512 / 2048	20.98M	29.54
Universal	512 / 2048	7.34M	29.12
RingFormer	512 / 2048	8.94M	29.52
Vanilla Transformer	1024 / 4096	176.18M	30.96
One Wide FFN	1024 / 4096	83.91M	29.88
Universal	1024 / 4096	29.37M	29.47
RingFormer	1024 / 4096	35.71M	<u>29.96</u>

Table 1: Translation results on WMT-14 De-En (Bojar et al., 2014). We evaluated models based on test dataset BLEU score (Papineni et al., 2002), which is rounded to the second decimal place. **Bolded** score indicates the highest performance, <u>underlined</u> score indicates the second highest performance. The H, FF, and P^* represent the hidden input dimension, feedforward block dimension, parameter size (except parameters of embedding layer in encoder, decoder and vocabulary head), respectively.

with Vanilla Transformer model (Vaswani et al., 2017) and One Wide FFN model (Pires et al., 2023)
with less number of parameters for base and large size models. RingFormer outperforms Universal model (Dehghani et al., 2019), while having similar parameter size. These results also imply that our design choice for level-signals is more effective than adding input-independent sinusoidal vectors.

425

426

427

428

429

430

431

432

433

434

435

Image Classification The experimental results for image classification are shown in Table 2 and 3. Using *ImageNet-small*, we conducted experi-

model	H/FF	P	$ Acc \uparrow$
ViT	512/2048	19.36M	63.66%
UiT	512/2048	3.60M	58.64%
OWF^d	376 / 1024	4.51M	58.62%
RingFormer	512 / 2048	4.4M	<u>60.66%</u>
ViT^d	328 / 1536	8.94M	62.22%
UiT^s	848 / 3072	8.84M	59.38%
OWF	512 / 2048	8.86M	61.50%
RingFormer ^s	728 / 3072	8.82M	62.58%

Table 2: Image classification results on *ImageNetsmall* (the subset of ImageNet-1K (Deng et al., 2009)). **Bolded** score indicates the highest performance, <u>underlined</u> score indicates the second highest performance. The superscripts "d" and "s" represent that the models are downscaled and upscaled, respectively. The H, FF, and P represent the hidden input dimension, feedforward block dimension, and total parameter size, respectively. The values for P and Acc were rounded to the second decimal place.

ments on the ViT (Dosovitskiy et al., 2020) model, downscaled One Wide FFN (OWF^d) (Pires et al., 2023), UiT (Dehghani et al., 2019) and our Ring-Former model. The results, presented in the upper half of Table 2, indicate that the ViT model achieves the highest accuracy, which is expected as it has more than four times the number of parameters compared to the other models. However, our RingFormer model has the second best performance, outperforming the other models of the same size. In the below half of Table 2, where

437 438

436

- 439 440
- 441 442
- 443 444

445

model	H/FF	P	$Acc\uparrow$
ViT	768 / 3072	86.42M	65.65%
OWF	768 / 3072	34.45M	<u>64.31%</u>
UiT	768 / 3072	8.45M	61.63%
RingFormer	768 / 3072	12.02M	63.68%
UiT ^s	1560 / 6240	31.99M	63.30%
RingFormer ^s	1284 / 5120	31.95M	65.91%

Table 3: Image classification results on *ImageNet-1K* (Deng et al., 2009)). **Bolded** score indicates the highest performance, <u>underlined</u> score indicates the second highest performance. The superscript "s" represent that the models are upscaled. The H, FF, and P represent the hidden input dimension, feedforward block dimension, total parameter size, respectively. The values for P and Acc were rounded to the second decimal place.

we scale all the models to the size of One Wide FFN model, our model shows the best performance, which shows the effectiveness of our approach.

We observed similar tendency when we trained bigger size models on the *ImageNet-1K* dataset, for which the results are shown in Table 3. When comparing the models with the same input hidden dimension and feedforward block dimension, ViT model showed the best result, but when we upscaled our RingFormer model (RingFormer^s) to match the size of OWF model, it outperformed the two baseline models (OWF and UiT^s), and also showed slightly higher performance compared to the ViT model.

Additionally, we calculated the forward GFLOPs for models the same H/FF shown in Table 3, namely ViT, OWF, UiT and RingFormer. For an RGB input image of size 224x224 (applied with 16x16 patch size), ViT, UiT and OWF models exhibit similar computational costs of around 17.636 GFLOPs, while RingFormer requires slightly more computations at 19.03 GFLOPs. This increase is attributed to the additional depth-specific and inputdependent level signals used in RingFormer to improve performance while maintaining a lower parameter count compared to standard Transformers.

Representation Similarity Analysis To analyze representations across layers / iterations between 474 the original Transformer model and other models, 475 we utilized CKA (Kornblith et al., 2019) method as 476 477 shown in Figure 2. We performed this analysis on base size models, for which we used 3K test source-478 target pair of sentences from WMT-14 De-En (Bo-479 jar et al., 2014). The similarity scores on the diag-480 onal axis in the sub-figures indicate how close the 481

layers (sharing the same index) are between models. We found that RingFormer closely matches the Vanilla Transformer (Vaswani et al., 2017) along with One Wide FFN (Pires et al., 2023), while Universal Transformer (Dehghani et al., 2019) shows lower similarity. We also report the analysis results for large models in Appendix A.2. 482

483

484

485

486

487

488

489

490

491

492

493

494

495

496

497

498

499

500

501

502

503

504

505

506

507

508

509

510

511

512

513

514

515

516

517

518

519

520

521

522

523

524

525

526

527

528

529

530

531

Mean Attention Distance Analysis To study the qualities of attention heads in the vision models, we perform MAD analysis, which is conducted in the original ViT paper (Dosovitskiy et al., 2020). We first do the analysis on the smaller models trained on *ImageNet-small* (ViT, UiT^s, OWF, and RingFormer^s shown in Table 2), and also on the larger size models trained on *ImageNet-1K* (ViT, OWF, UiT^s and RingFormer^s in Table 3). We computed mean attention distances of 500 images randomly taken from the *ImageNet-small* validation set and took their average. The MAD analysis plots for each model above are shown in Figure 3.

We observe that, in the ViT model, different attention heads yield different attention distances suggesting they use both local and global information from an image. But as we go deeper in the Transformer blocks, the heads tend to focus more on global aggregate information. The same type of phenomenon occurs in the case of One Wide FFN model, which is expected as its attention layers are not recurrent. In the case of our RingFormer model, the properties of its attention heads are also very similar to those of the ViT model. It can be seen as a validation of our hypothesis that the level signals could successfully steer the behavior of a recurrent Transformer model as it goes through a series of iterations. When it comes to the Universal Transformer model, it is found that the types of signals that exist in that model could not sufficiently help it simulate its attention module behavior as it is considerably different compared to that of the ViT model.

4.3 Ablation Study

In this section, we conduct an ablation study with various experiments on the translation task to validate the effectiveness of our proposed method. The training details for ablation study are the same as those of main translation models, except smaller batch size (128), hidden input dimension and feedforward block dimension. The model hyperparameters and experimental results for these ablation studies are presented in Table 4.



Figure 3: MAD (mean attention distance) analysis for the models trained on image classification task: ViT, RingFormer, OWF - One Wide FFN, UiT - Universal Transformer. The smaller models shown on the upper row have 8 attention heads, and larger models shown on the lower row have 12 attention heads. The points on the plots represent the mean attention distance of an attention head belonging to a particular Transformer layer.

First, we train a recurrent Transformer using static level signals introduced in Universal Transformers (Dehghani et al., 2019), which has the lowest performance. When we drop either attention level signals or FFN level signals, "w.o. attn" and "w.o. FF" in Table 4, the performance degradation occurs compared with other variations where those signals are present. Also, we do the following two ablations: 1) we add level signals "before attn" projection while keeping our original design for FF level signals, 2) we add level signals, "inter-FF signal", after intermediate feedforward projection like $FFN(x) = \sigma(xW_{up} + g_{Fi}(x))W_{down}$, while keeping our original design for attention level signals. The performances of those two experiments are almost the same but lower than our design choice where additions occur i) after attention projection and ii) before the up-projection layer of the FF block. In addition, when we use smaller rank, H / 32, compared to our default rank, H / 16, the performance decreases, but when we increase the rank or make the matrix full-rank to generate level signals, as expected, the models show better performance.

5 Conclusion

533

534

537

539

541

542

543

545

549

551

554

555

559

In this paper, we introduce *RingFormer*, a parameter-efficient recurrent Transformer architecture that employs a single Transformer layer recur-

model	H/F	P	$BLEU\uparrow$
static signal	128 / 512	20.48M	23.35
w.o. attn signal	128 / 512	20.51M	24.23
w.o. FF signal	128/512	20.59M	24.37
before attn	128 / 512	20.58M	24.56
inter-FF signal	128 / 512	20.56M	24.58
H / 32 rank signal	128 / 512	20.53M	24.21
H / 8 rank signal	128 / 512	20.68M	24.96
full-rank signal	128 / 512	21.27M	25.37
Ours	128 / 512	20.58M	24.92

Table 4: Ablation experiment results of translation task in WMT-14 (Bojar et al., 2014) German-English pairs with various model-designs. Each model is evaluated by BLEU (Papineni et al., 2002) score on the test set. The H, FF, and P represent the hidden input dimension, feedforward block dimension, and total parameter size, respectively.

rently while integrating input-dependent signal vectors created using low-rank matrices for each level. This approach significantly reduces the number of parameters while maintaining high performance in tasks such as machine translation and image classification. We hope that our research on enhancing recurrent Transformer with adaptive level signals can enable smaller organizations and research institutions to train powerful models without the need for extensive computational resources, thus democratizing access to advanced AI capabilities.

565

566

568

570

560

561

6 Limitations

571

589

592 593

594

596

597

598

599

602

609

610

611

612

613

614

615

616

617

618

619

622

Our approach introduces additional computations compared to the original Transformer due to the 573 integration of depth-specific and input-dependent 574 signals. However, this trade-off is necessary to 575 maintain the performance of standard Transformers 576 577 while significantly reducing parameter count compared to other recurrent Transformer models. Due 578 to computational constraints, we were not able to conduct experiments on large-scale language mod-580 eling tasks, which require significantly more data 581 and training resources, and our experiments were limited to relatively smaller scale models. While our design choices suggest that RingFormer should retain its advantages at larger scales, future work can focus on further validating its performance on 586 587 billion-parameter models and explore its effectiveness in domains such as language modeling.

References

Anthropic. 2023. Model card and evaluations for claude models. https://www-cdn.anthropic.com/ 5c49cc247484cecf107c699baf29250302e5da70/ claude-2-model-card.pdf.

Ondřej Bojar, Christian Buck, Christian Federmann, Barry Haddow, Philipp Koehn, Johannes Leveling, Christof Monz, Pavel Pecina, Matt Post, Herve Saint-Amand, Radu Soricut, Lucia Specia, and Aleš Tamchyna. 2014. Findings of the 2014 workshop on statistical machine translation. In *Proceedings of the Ninth Workshop on Statistical Machine Translation*, pages 12–58, Baltimore, Maryland, USA. Association for Computational Linguistics.

- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.
- Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. 2020. End-to-end object detection with transformers. In *European conference on computer vision*, pages 213–229. Springer.
- Lili Chen, Kevin Lu, Aravind Rajeswaran, Kimin Lee, Aditya Grover, Michael Laskin, Pieter Abbeel, Aravind Srinivas, and Igor Mordatch. 2021. Decision transformer: Reinforcement learning via sequence modeling. *arXiv preprint arXiv:2106.01345*.
- Mostafa Dehghani, Josip Djolonga, Basil Mustafa, Piotr Padlewski, Jonathan Heek, Justin Gilmer, Andreas Peter Steiner, Mathilde Caron, Robert Geirhos, Ibrahim Alabdulmohsin, et al. 2023. Scaling vision

transformers to 22 billion parameters. In *International Conference on Machine Learning*, pages 7480– 7512. PMLR. 623

624

625

626

627

628

629

630

631

632

633

634

635

636

637

638

639

640

641

642

643

644

645

646

647

648

649

650

651

652

653

654

655

656

657

658

659

660

661

662

663

664

665

666

667

668

669

670

671

672

673

- Mostafa Dehghani, Stephan Gouws, Oriol Vinyals, Jakob Uszkoreit, and Łukasz Kaiser. 2019. Universal transformers. *Preprint*, arXiv:1807.03819.
- Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. 2009. Imagenet: A large-scale hierarchical image database. In 2009 IEEE conference on computer vision and pattern recognition, pages 248–255. Ieee.
- Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. 2020.
 An image is worth 16x16 words: Transformers for image recognition at scale. arXiv preprint arXiv:2010.11929.
- Mor Geva, Roei Schuster, Jonathan Berant, and Omer Levy. 2021. Transformer feed-forward layers are key-value memories. *Preprint*, arXiv:2012.14913.
- Gemini Team Google. 2024. Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context. *Preprint*, arXiv:2403.05530.
- Andrey Gromov, Kushal Tirumala, Hassan Shapourian, Paolo Glorioso, and Daniel A Roberts. 2024. The unreasonable ineffectiveness of the deeper layers. *arXiv preprint arXiv:2403.17887*.
- Dan Hendrycks and Kevin Gimpel. 2023. Gaussian error linear units (gelus). *Preprint*, arXiv:1606.08415.
- Jordan Hoffmann, Sebastian Borgeaud, Arthur Mensch, Elena Buchatskaya, Trevor Cai, Eliza Rutherford, Diego de Las Casas, Lisa Anne Hendricks, Johannes Welbl, Aidan Clark, et al. 2022. An empirical analysis of compute-optimal large language model training. *Advances in Neural Information Processing Systems*, 35:30016–30030.
- Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2021. Lora: Low-rank adaptation of large language models. *Preprint*, arXiv:2106.09685.
- Yifan Jiang, Shiyu Chang, and Zhangyang Wang. 2021. Transgan: Two transformers can make one strong gan. *arXiv preprint arXiv:2102.07074*, 1(3).
- Diederik P. Kingma and Jimmy Ba. 2017. Adam: A method for stochastic optimization. *Preprint*, arXiv:1412.6980.
- Simon Kornblith, Mohammad Norouzi, Honglak Lee, and Geoffrey Hinton. 2019. Similarity of neural network representations revisited. In *International conference on machine learning*, pages 3519–3529. PMLR.

Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. 2021. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 10012–10022.

675

676

677

679

681

682

687

694

695

696

701

703

704

707

710

711

712

713

714

715

716

719

720

721

725

726

727

730

- OpenAI, Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, and et al. 2024. Gpt-4 technical report.
- Myle Ott, Sergey Edunov, David Grangier, and Michael Auli. 2018. Scaling neural machine translation. *Preprint*, arXiv:1806.00187.
- 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.
- William Peebles and Saining Xie. 2022. Scalable diffusion models with transformers. *arXiv preprint arXiv:2212.09748*.
- Telmo Pessoa Pires, António V. Lopes, Yannick Assogba, and Hendra Setiawan. 2023. One wide feedforward is all you need. *Preprint*, arXiv:2309.01826.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9.
- Machel Reid, Edison Marrese-Taylor, and Yutaka Matsuo. 2021. Subformer: Exploring weight sharing for parameter efficiency in generative transformers. *Preprint*, arXiv:2101.00234.
- Sho Takase and Shun Kiyono. 2021. Lessons on parameter sharing across layers in transformers. *arXiv preprint arXiv:2104.06022*.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. 2023. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- Haoran Xu, Benjamin Van Durme, and Kenton Murray. 2021. Bert, mbert, or bibert? a study on contextualized embeddings for neural machine translation. *Preprint*, arXiv:2109.04588.
- Bowen Zhang, Shuyang Gu, Bo Zhang, Jianmin Bao, Dong Chen, Fang Wen, Yong Wang, and Baining Guo. 2022. Styleswin: Transformer-based gan for high-resolution image generation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11304–11314.

A Appendix

A.1 Implementation Details

Translation Models are trained based on the two size variations, base size and large size. The base size models are trained based on the following model configuration settings: 6 Transformer layers, 8 attention heads, 512 hidden dimension size, 2048 feedforward dimension with maximum sequence length 50. For training, their maximum learning rate is 7e-4 with 17K step cosine warm-up scheduler and total 210K training steps on two A100 80GB GPUs. The large size models are trained based on the following model configuration settings: 6 Transformer layers, 16 attention heads, 1024 hidden dimension size, 4096 feedforward dimension with maximum sequence length 50. For training, their maximum learning rate is 2e-4 with 17K step cosine warm-up scheduler and total 210K training steps on two A100 80GB GPUs.

731

732

734

735

737

738

739

740

741

742

743

744

745

746

747

748

750

751

752

753

754

755

756

757

758

760

761

762

763

766

767

768

770

771

772

773

774

778

Image Classification For the models trained on *ImageNet-small* dataset, we used 224x224 image resolution, 16x16 patch size, 6 Transformer layers, 8 attention heads, learning rate of $1e^{-3}$, cosine learning rate scheduler with 2K warm-up steps, batch size of 1024, and training for 9775 steps (100 epochs) with one RTX 3090 GPU. For the models trained on *ImageNet-1K* dataset, we used the same image resolution and patch size as mentioned above, 12 Transformer layers, 12 attention heads, learning rate of $5e^{-4}$, cosine learning rate scheduler with 3128 warm-up steps (5 epochs), batch size of 4096, 16 gradient accumulation steps, and training for around 15650 steps (50 epochs) on two RTX 3090 GPUs.

For all models, we used dropout rate of 0.1, gradient clipping of 1.0 during training and GELU (Hendrycks and Gimpel, 2023) activation function.

A.2 Additional Analysis

In Figure 4, we share the representation similarity analysis for big size models in the Translation task. This analysis also has been conducted under the same conditions as in the base size case. Similar with the results in Figure 2, One Wide FFN (Pires et al., 2023) and our RingFormer model have higher layer-wise representations with the Vanilla Transformer (Vaswani et al., 2017) compared to Universal Transformer (Dehghani et al., 2019).



Figure 4: Representation Similarity Analysis using CKA (centered kernel alignment) (Kornblith et al., 2019) for the large-size models trained on the translation task: Transformer, Ring - RingFormer, OWF - One Wide FFN, Uni - Universal Transformer. The figures on the upper row are for the encoder part. The figures on the lower row means are for the decoder part. All models have 6 number of layers / iterations. The values on the figures are between 0 and 1, where higher values indicate more similarity of layers between models.

A.3 ImageNet-small Dataset

We sampled a subset of *ImageNet-1K* (Deng et al., 2009) that contains randomly selected 100 classes, with 100,000 images for training and 5000 images for testing, in order to perform experiments on smaller size models. In the supplimentary **code**.zip file, we will share the names of all the sampled images for training and testing as a json file.