

Supplement to “A Primal-Dual Framework for Transformers and Neural Networks”

A ADDITIONAL DETAILS ON THE EXPERIMENTS

This section provides datasets, models, and training details for experiments in Section 3. As mentioned in Section 3, for Attention-BN models, recentering queries and keys alone is sufficient for accuracy improvement, and we weight the mean μ in Eqn 22 with a constant β . Hence Eqn 22 is simplified to:

$$h_i = \sum_{j=1}^N \text{softmax} \left((q_i - \beta\mu)^\top (k_j - \beta\mu) / \sqrt{D} \right) v_j. \quad (26)$$

In our experiments, we consider the constant β in Attention-BN/BN+SH and the different downsampling scales in Attention-SH/SH+BN as hyper-parameters to finetune. All of our experiments are conducted on a server with 4 NVIDIA A100 GPUs.

A.1 UEA TIME SERIES CLASSIFICATION

Datasets and metrics The benchmark (Bagnall et al., 2018) consists of 30 datasets. Following (Wu et al., 2022), we choose 10 datasets, which vary in input sequence lengths, the number of classes, and dimensionality, to evaluate our models on temporal sequences. We report the test accuracy as evaluation for the benchmark.

Models and baselines The experiment setups and configurations for the softmax/linear baseline and our models are the same as in (Wu et al., 2022)¹ (for the PEMS-SF, SelfRegulationSCP2, UWaveGestureLibrary datasets) and (Zerveas et al., 2021)² (for the other tasks). In all models, the number of heads is 8, whereas the model dimension and number of transformer layers are varied. For Attention-SH/SH+BN, we downsample keys and values by the factor of 2, after every two successive heads.

A.2 LONG RANGE ARENA BENCHMARK

Datasets and metrics We adopt the tasks: Listops (Nangia & Bowman, 2018), byte-level IMDb reviews text classification (Maas et al., 2011), byte-level document retrieval (Radev et al., 2013), CIFAR-10 image classification (Krizhevsky et al., 2009) and the Pathfinder challenge (Linsley et al., 2018) in the LRA benchmark for our experiments. They consist of long sequences of length $2K$, $4K$, $4K$, $1K$, and $1K$ respectively. The evaluation protocol and metric are the same as in (Tay et al., 2021).

Models and baselines All our models and softmax/linear baselines follow the same architecture and configuration as in (Zhu et al., 2021)³. Each model consists of two layers and 64 embedding dimensions. While one head at each layer remains intact, the keys and values of the other heads are halved in our Attention-SH/SH+BN experiments.

A.3 IMAGE CLASSIFICATION ON IMAGENET

Datasets and metrics The ImageNet dataset (Deng et al., 2009; Russakovsky et al., 2015) consists of 1.28M training images and 50K validation images. The task is to classify 1000 categories. Top-1 and top-5 accuracies are reported.

Models and baselines Our baseline is DeiT-tiny model (Touvron et al., 2021) with 12 transformer layers, 4 attention heads per layer, and the model dimension of 192. For model setting and setting and configuration, we follow (Touvron et al., 2021)⁴. The downsampling scales in Attention-SH/BN+SH models are [1, 1, 2, 4] for 4 heads at each layer, respectively.

¹Implementation available at <https://github.com/thuml/Flowformer>.

²Implementation available at https://github.com/gzerveas/mvts_transformer.

³Implementation available at <https://github.com/NVIDIA/transformer-ls>.

⁴Implementation available at <https://github.com/facebookresearch/deit>.

Table 5: Test Accuracy (%) of the Linear Attention-BN/SH/BN+SH vs. the baseline Linear Attention (Katharopoulos et al., 2020) on the UEA Time Series Classification Archive benchmark (Bagnall et al., 2018). Our proposed attentions outperform the baseline.

Dataset/Model	Baseline Linear	Linear Attention-BN	Linear Attention-SH	Linear Attention-BN+SH
ETHANOLCONCENTRATION	33.84 ± 0.66	34.98 ± 0.74	34.76 ± 0.69	34.35 ± 0.70
FACEDETECTION	69.17 ± 0.32	69.22 ± 0.17	69.38 ± 0.17	69.12 ± 0.19
HANDWRITING	32.87 ± 0.27	32.86 ± 0.49	32.82 ± 0.12	32.98 ± 0.36
HEARTBEAT	75.61 ± 0.73	75.78 ± 0.71	74.96 ± 0.62	75.94 ± 0.68
JAPANESEVOWELS	99.37 ± 0.16	99.60 ± 0.19	99.28 ± 0.41	99.33 ± 0.19
PEMS-SF	83.43 ± 0.88	85.74 ± 0.67	86.51 ± 0.88	84.97 ± 0.76
SELFREGULATIONSCP1	90.90 ± 0.40	91.81 ± 0.69	90.76 ± 0.59	91.92 ± 0.60
SELFREGULATIONSCP2	55.18 ± 0.89	56.11 ± 0.94	54.44 ± 0.88	55.74 ± 0.92
SPOKENARABICDIGITS	99.07 ± 0.10	99.01 ± 0.07	99.03 ± 0.18	98.91 ± 0.17
UWAVEGESTURELIBRARY	85.63 ± 0.81	86.04 ± 0.86	84.89 ± 1.00	85.78 ± 0.75
AVERAGE ACCURACY	72.51 ± 0.34	73.12 ± 0.20	72.68 ± 0.40	72.90 ± 0.23

Table 6: Top-1 and top-5 accuracy (%) of the Attention-Conv2D Deit vs. the baseline Deit with the softmax attention on the ImageNet image classification task. The Attention-Conv2D Deit significantly outperforms the baseline in both top-1 and top-5 accuracy.

Metric/Model	Baseline Softmax Deit	Attention-Conv2D Deit
Top-1 Acc (%)	72.23 ± 0.23	73.18 ± 0.24
Top-5 Acc (%)	91.13 ± 0.15	91.52 ± 0.13

Table 7: Test Accuracy (%) of the Attention-Conv1D vs. the baseline softmax attention on 5 tasks of the LRA benchmark (Tay et al., 2021). Our models outperform the softmax baseline.

Dataset/Model	Baseline Softmax	Attention-Conv1D
LISTOPS	36.76 ± 0.42	37.20 ± 0.05
TEXT	64.90 ± 0.07	64.92 ± 0.44
RETRIEVAL	79.68 ± 0.52	80.75 ± 0.15
IMAGE	39.23 ± 1.35	39.18 ± 0.59
PATHFINDER	72.72 ± 0.75	73.01 ± 0.24
AVERAGE ACCURACY	58.66 ± 0.26	59.01 ± 0.20

B ADDITIONAL EXPERIMENTAL RESULTS

B.1 UEA TIME SERIES CLASSIFICATION USING THE LINEAR ATTENTION-BN/SH/BN+SH

Table 5 summarizes the comparison between the Linear Attention-BN/SH/BN+SH and the baseline Linear Attention on the UEA Time Series Classification task. The Linear Attention-BN/SH/BN+SH achieve better accuracy than the Linear Attention baseline while being more efficient.

B.2 CONVOLUTION ATTENTION

Table 6 demonstrates the advantage of Attention-Conv2D (Def. 3, Section G) over softmax Deit on the ImageNet image classification task. Furthermore, as shown in Table 7, the Attention-Conv1D (Def. 4, Section G) outperforms the baseline softmax attention on 5 tasks of the LRA benchmark (Tay et al., 2021).

B.3 ADDITIONAL EXPERIMENTS ON THE UEA TIMESERIES CLASSIFICATION BENCHMARK AND THE UCR TIME SERIES REGRESSION ARCHIVE

In this section, we further demonstrate the advantage of our Attention-BN/SH/BN+SH on additional 15 tasks in the UEA Time Series Classification benchmark and on 6 tasks in the UCR Time Series Regression benchmark. The results in Table 8 and 9 show that our Attention-BN and Attention-SH+BN outperform the baseline softmax transformers significantly on both of these benchmarks, while the attention-SH has comparable performance with the baseline but being more efficient.

B.4 UEA TIME SERIES CLASSIFICATION USING THE SPARSE ATTENTION-BN/SH/BN+SH

Table 10 summarizes the comparison between the Sparse Attention-BN/SH/BN+SH and the Sparse Attention baseline on a subset of the UEA Time Series Classification benchmark. Our models when combined with Sparse Attention achieve significantly better accuracy than the Sparse Attention baseline while the Sparse Attention-SH/BN+SH are more efficient (See Fig. 3 and Fig. 4 in Appendix C).

Table 8: Root mean square error (RMSE) of the Attention-BN/SH/BN+SH vs. the baseline softmax attention on 6 UCR Time Series Regression tasks (Tan et al., 2020). Smaller RMSE indicates better performance.

Dataset/Model	Baseline Softmax	Attention-BN	Attention-SH	Attention-BN+SH
APPLIANCESENERGY	3.44 ± 0.06	3.38 ± 0.34	3.39 ± 0.02	3.37 ± 0.23
BENZENECONCENTRATION	0.91 ± 0.03	0.89 ± 0.17	1.00 ± 0.09	0.90 ± 0.08
BEIJINGPM10	92.31 ± 1.06	92.00 ± 0.89	92.82 ± 0.92	92.40 ± 0.85
BEIJINGPM25	59.73 ± 1.21	59.55 ± 0.92	59.66 ± 0.88	59.24 ± 1.22
LIVEFUELMOISTURE	43.08 ± 0.17	43.01 ± 0.50	43.65 ± 0.09	43.79 ± 0.49
IEEPPG	32.12 ± 1.25	30.69 ± 0.64	31.38 ± 1.02	30.73 ± 1.20
AVERAGE RMSE	38.60 ± 0.67	38.25 ± 0.30	38.65 ± 0.27	38.40 ± 0.51

Table 9: Accuracy (%) of the Attention-BN/SH/BN+SH vs. the baseline softmax attention on other 15 UEA Time Series classification tasks (Bagnall et al., 2018).

Dataset/Model	Baseline Softmax	Attention-BN	Attention-SH	Attention-BN+SH
ARTICULARYWORDRECOGNITION	97.44 ± 0.42	98.22 ± 0.87	97.22 ± 0.95	98.44 ± 0.41
BASICMOTIONS	98.75 ± 1.25	99.38 ± 1.08	99.37 ± 1.06	99.78 ± 0.51
EPILEPSY	93.71 ± 1.23	92.27 ± 0.74	89.13 ± 1.07	92.02 ± 1.02
ERING	95.18 ± 0.52	95.18 ± 0.37	94.72 ± 0.66	95.46 ± 0.40
FINGERMOVEMENTS	59.67 ± 0.47	63.00 ± 0.41	61.33 ± 0.70	63.66 ± 0.64
LIBRAS	85.00 ± 0.45	85.37 ± 0.69	83.88 ± 0.45	85.00 ± 0.78
NATOPS	95.00 ± 0.45	95.37 ± 0.26	96.29 ± 0.69	95.74 ± 0.94
RACKETSPORTS	87.28 ± 0.82	87.93 ± 0.31	88.16 ± 0.54	89.03 ± 0.64
ATRIALFIBRILLATION	33.33 ± 2.71	41.67 ± 2.88	35.00 ± 2.89	41.68 ± 2.80
CRICKET	94.90 ± 0.65	95.37 ± 0.65	93.98 ± 0.65	96.29 ± 0.65
STANDWALKJUMP	50.00 ± 2.33	55.55 ± 2.14	50.01 ± 2.34	55.00 ± 2.08
HANDMOVEMENTDIRECTION	63.96 ± 2.30	64.41 ± 2.76	61.71 ± 2.64	66.66 ± 2.54
LSST	58.54 ± 0.54	57.05 ± 0.26	60.34 ± 0.73	59.91 ± 0.34
DUCKDUCKGEESE	64.50 ± 1.96	65.00 ± 1.73	64.50 ± 1.95	65.50 ± 1.66
MOTORIMAGERY	58.66 ± 1.25	60.67 ± 1.69	59.00 ± 1.41	62.00 ± 0.81
AVERAGE ACCURACY	75.73 ± 0.51	77.10 ± 0.22	75.61 ± 0.18	77.74 ± 0.24

Table 10: Test Accuracy (%) of the Sparse Attention-BN/SH/BN+SH vs. the baseline Sparse Attention (Child et al., 2019) on a subset of the UEA Time Series Classification Archive benchmark (Bagnall et al., 2018). Our proposed attentions outperform the baseline.

Dataset/Model	Baseline Sparse	Sparse Attention-BN	Sparse Attention-SH	Sparse Attention-BN+SH
ETHANOLCONCENTRATION	33.33 ± 1.23	33.33 ± 0.78	32.50 ± 0.57	33.46 ± 0.71
FACEDETECTION	68.58 ± 0.95	68.65 ± 0.44	68.67 ± 0.78	68.44 ± 0.51
HANDWRITING	31.08 ± 0.38	31.79 ± 0.44	32.75 ± 0.39	33.37 ± 0.61
HEARTBEAT	74.95 ± 0.81	75.98 ± 0.72	74.96 ± 0.80	76.09 ± 0.75
JAPANESEVOWELS	99.45 ± 0.10	99.54 ± 0.12	99.18 ± 0.14	99.36 ± 0.34
PEMS-SF	82.08 ± 0.63	83.81 ± 0.47	82.66 ± 0.63	84.01 ± 0.89
SELFREGULATIONSCP1	91.24 ± 0.85	91.69 ± 0.42	91.47 ± 0.84	91.70 ± 0.16
SELFREGULATIONSCP2	55.18 ± 0.69	58.52 ± 0.71	55.92 ± 0.94	56.67 ± 0.68
SPOKENARABICDIGITS	99.04 ± 0.06	99.10 ± 0.15	99.06 ± 0.13	99.15 ± 0.09
UWAVEGESTURELIBRARY	84.90 ± 0.39	85.73 ± 0.38	85.31 ± 0.88	86.56 ± 0.25
AVERAGE ACCURACY	71.98 ± 0.38	72.81 ± 0.15	72.25 ± 0.24	72.88 ± 0.35

Table 11: Test Accuracy (%) of the Attention-BN/BN+SH with β is learnable or set as a hyperparameter on the retrieval task (Tay et al., 2021).

Model	Retrieval
Attention-BN (learn β)	80.77 ± 0.23
Attention-BN+SH (learn β)	81.31 ± 0.25
Attention-BN (β as a hyperparameter)	81.05 ± 0.08
Attention-BN+SH (β as a hyperparameter)	81.20 ± 0.11

B.5 ATTENTION-BN/BN+SH WITH LEARNABLE β

We experiment with our Attention-BN/BN+SH with learnable β on the retrieval task. Table 11 shows that learning β does not improve much over setting β to be a hyperparameter.

C ADDITIONAL RESULTS ON EFFICIENCY ANALYSIS

This section provides more efficiency analysis on our models.

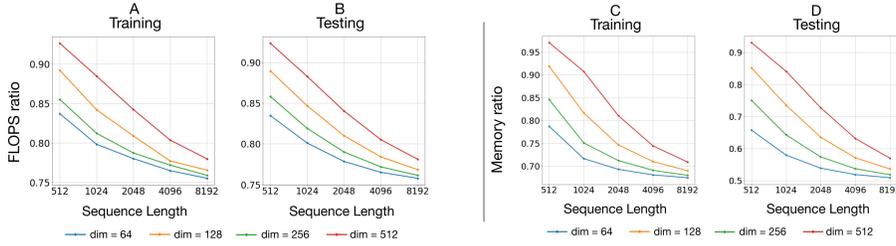


Figure 2: (Left) FLOPS ratios and (Right) memory usage ratios between the Attention-SH and the softmax attention baseline trained on the LRA retrieval task for different model dimensions and sequence lengths.

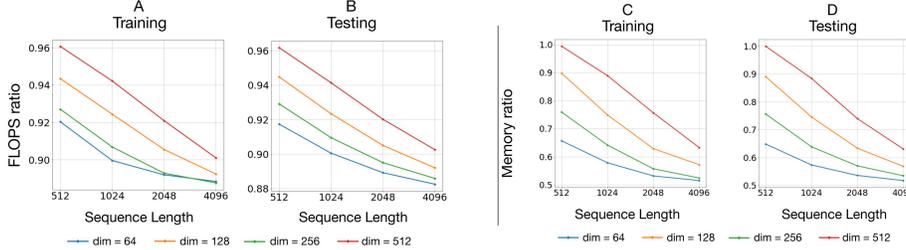


Figure 3: (Left) FLOPS ratios and (Right) memory usage ratios between the Sparse Attention-BN+SH and the Sparse Attention baseline trained on the LRA retrieval task for different model dimensions and sequence lengths. When using our models, the reduction in computation and memory improves with sequence length. When scaling up the model with greater model dimension, our methods remain significantly more efficient than the baseline.

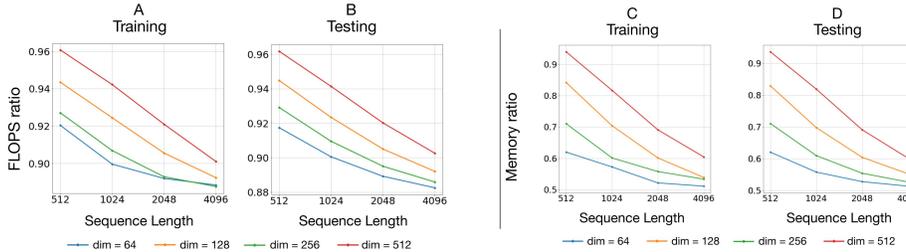


Figure 4: (Left) FLOPS ratios and (Right) memory usage ratios between the Sparse Attention-SH and the Sparse Attention baseline trained on the LRA retrieval task for different model dimensions and sequence lengths. When using our models, the reduction in computation and memory improves with sequence length. When scaling up the model with greater model dimension, our methods remain significantly more efficient than the baseline.

Attention-SH. Fig.2 shows the efficiency benefits of our Attention-SH when trained on the retrieval task. Same as in the case of Attention-SH+BN, the efficiency benefits of our Attention-SH over the baseline Softmax attention grows when N and D increase.

Sparse Attention-SH/BN+SH. Fig.3 and Fig.4 show that the efficiency advantages of our Sparse Attention-BN+SH and Sparse Attention-SH, respectively, increase as the model dimension D and sequence length N grow. All models are trained on the LRA retrieval task. In addition to the efficiency advantage, the Sparse Attention-BN+SH also significantly outperforms the Sparse Attention baseline in terms of accuracy in this task (79.86% vs. 78.20%) while the Sparse Attention-SH achieves a comparable result to the baseline. More accuracy advantages of the Sparse Attention-BN/SH/BN+SH over the Sparse Attention baseline are given in Table 10.

D DERIVING SOFTMAX ATTENTION.

Choosing the appropriate $h(\mathbf{x})$ and $\Phi(\mathbf{x})$ allows us to derive the popular softmax attention given in Eqn. 1 and 2. In particular, if we choose $h(\mathbf{x}) := \sum_{j=1}^N \Phi(\mathbf{x})^T \Phi(\mathbf{k}_j)$, Eqn. 10 becomes

$$f(\mathbf{x}) = \sum_{j=1}^N \frac{\Phi(\mathbf{x})^T \Phi(\mathbf{k}_j)}{\sum_{j'=1}^N \Phi(\mathbf{x})^T \Phi(\mathbf{k}_{j'})} \mathbf{v}_j + \mathbf{b} = \frac{\sum_{j=1}^N \Phi(\mathbf{x})^T \Phi(\mathbf{k}_j) \mathbf{v}_j}{\sum_{j'=1}^N \Phi(\mathbf{x})^T \Phi(\mathbf{k}_{j'})} + \mathbf{b}. \quad (27)$$

We then select $\Phi(\mathbf{x}) = (a_{l_0}^{(0)}, a_1^{(1)}, \dots, a_{l_1}^{(1)}, \dots, a_1^{(t)}, \dots, a_{l_t}^{(t)}, \dots)$ where $l_t = \binom{D+t-1}{t}$ and

$$a_l^{(t)} = \frac{(x_1/\sqrt[4]{D})^{n_1} \dots (x_D/\sqrt[4]{D})^{n_D}}{\sqrt{n_1! \dots n_D!}} \mid n_1 + \dots + n_D = t, 1 \leq l \leq l_t. \quad (28)$$

Since

$$\exp(\mathbf{x}^T \mathbf{y}) = \sum_{t=0}^{\infty} \frac{(\mathbf{x}^T \mathbf{y})^t}{t!} = \sum_{t=0}^{\infty} \sum_{n_1 + \dots + n_D = t} \left(\frac{x_1^{n_1} \dots x_D^{n_D}}{\sqrt{n_1! \dots n_D!}} \right) \left(\frac{y_1^{n_1} \dots y_D^{n_D}}{\sqrt{n_1! \dots n_D!}} \right), \quad (29)$$

then Eqn. 27 becomes

$$\begin{aligned} f(\mathbf{x}) &= \sum_{j=1}^N \frac{\sum_{t=0}^{\infty} \sum_{n_1 + \dots + n_D = t} \left(\frac{(x_1/\sqrt[4]{D})^{n_1} \dots (x_D/\sqrt[4]{D})^{n_D}}{\sqrt{n_1! \dots n_D!}} \right) \left(\frac{(k_{j1}/\sqrt[4]{D})^{n_1} \dots (k_{jD}/\sqrt[4]{D})^{n_D}}{\sqrt{n_1! \dots n_D!}} \right)}{\sum_{j'=1}^N \sum_{t=0}^{\infty} \sum_{n_1 + \dots + n_D = t} \left(\frac{(x_1/\sqrt[4]{D})^{n_1} \dots (x_D/\sqrt[4]{D})^{n_D}}{\sqrt{n_1! \dots n_D!}} \right) \left(\frac{(k_{j'1}/\sqrt[4]{D})^{n_1} \dots (k_{j'D}/\sqrt[4]{D})^{n_D}}{\sqrt{n_1! \dots n_D!}} \right)} \mathbf{v}_j + \mathbf{b} \\ &= \sum_{j=1}^N \frac{\exp\left(\left(\frac{\mathbf{x}}{\sqrt[4]{D}}\right)^\top \frac{\mathbf{k}_j}{\sqrt[4]{D}}\right)}{\sum_{j'=1}^N \exp\left(\left(\frac{\mathbf{x}}{\sqrt[4]{D}}\right)^\top \frac{\mathbf{k}_{j'}}{\sqrt[4]{D}}\right)} \mathbf{v}_j + \mathbf{b} = \sum_{j=1}^N \frac{\exp(\mathbf{x}^\top \mathbf{k}_j / \sqrt{D})}{\sum_{j'=1}^N \exp(\mathbf{x}^\top \mathbf{k}_{j'} / \sqrt{D})} \mathbf{v}_j + \mathbf{b}. \quad (30) \end{aligned}$$

Let $\mathbf{x} = \mathbf{q}_i$, $\mathbf{b} = \mathbf{0}$ and relax the boundness constraint of \mathbf{v}_j in Remark 1. Eqn. 30 becomes Eqn. 2 of the softmax attention (Vaswani et al., 2017).

E BATCH NORMALIZED ATTENTION: DERIVATION OF EQN. 24

$$\begin{aligned} \mathbf{h}_i &= \sum_{j=1}^N \text{softmax} \left(\sum_{d=1}^D \frac{(\mathbf{q}_i(d) - \boldsymbol{\mu}(d))(\mathbf{k}_j(d) - \boldsymbol{\mu}(d))}{\sqrt{D}(\sigma_d^2 + \epsilon)} \right) \mathbf{v}_j \\ &= \sum_{j=1}^N \text{softmax} \left(\sum_{d=1}^D \frac{\mathbf{q}_i(d)\mathbf{k}_j(d) - \mathbf{q}_i(d)\boldsymbol{\mu}(d) - \boldsymbol{\mu}(d)\mathbf{k}_j(d) + \boldsymbol{\mu}(d)\boldsymbol{\mu}(d)}{\sqrt{D}(\sigma_d^2 + \epsilon)} \right) \mathbf{v}_j \\ &= \sum_{j=1}^N \frac{\exp\left(\sum_{d=1}^D \frac{\mathbf{q}_i(d)\mathbf{k}_j(d) - \mathbf{q}_i(d)\boldsymbol{\mu}(d) - \boldsymbol{\mu}(d)\mathbf{k}_j(d) + \boldsymbol{\mu}(d)\boldsymbol{\mu}(d)}{\sqrt{D}(\sigma_d^2 + \epsilon)}\right)}{\sum_{j'=1}^N \exp\left(\sum_{d=1}^D \frac{\mathbf{q}_i(d)\mathbf{k}_{j'}(d) - \mathbf{q}_i(d)\boldsymbol{\mu}(d) - \boldsymbol{\mu}(d)\mathbf{k}_{j'}(d) + \boldsymbol{\mu}(d)\boldsymbol{\mu}(d)}{\sqrt{D}(\sigma_d^2 + \epsilon)}\right)} \mathbf{v}_j \\ &= \sum_{j=1}^N \frac{\exp\left(\sum_{d=1}^D \frac{\mathbf{q}_i(d)\mathbf{k}_j(d) - \boldsymbol{\mu}(d)\mathbf{k}_j(d)}{\sqrt{D}(\sigma_d^2 + \epsilon)}\right) \exp\left(\sum_{d=1}^D \frac{\boldsymbol{\mu}(d)\boldsymbol{\mu}(d) - \mathbf{q}_i(d)\boldsymbol{\mu}(d)}{\sqrt{D}(\sigma_d^2 + \epsilon)}\right)}{\sum_{j'=1}^N \exp\left(\sum_{d=1}^D \frac{\mathbf{q}_i(d)\mathbf{k}_{j'}(d) - \boldsymbol{\mu}(d)\mathbf{k}_{j'}(d)}{\sqrt{D}(\sigma_d^2 + \epsilon)}\right) \exp\left(\sum_{d=1}^D \frac{\boldsymbol{\mu}(d)\boldsymbol{\mu}(d) - \mathbf{q}_i(d)\boldsymbol{\mu}(d)}{\sqrt{D}(\sigma_d^2 + \epsilon)}\right)} \mathbf{v}_j \\ &= \sum_{j=1}^N \frac{\exp\left(\sum_{d=1}^D \frac{\mathbf{q}_i(d)\mathbf{k}_j(d) - \boldsymbol{\mu}(d)\mathbf{k}_j(d)}{\sqrt{D}(\sigma_d^2 + \epsilon)}\right)}{\sum_{j'=1}^N \exp\left(\sum_{d=1}^D \frac{\mathbf{q}_i(d)\mathbf{k}_{j'}(d) - \boldsymbol{\mu}(d)\mathbf{k}_{j'}(d)}{\sqrt{D}(\sigma_d^2 + \epsilon)}\right)} \mathbf{v}_j \\ &= \sum_{j=1}^N \text{softmax} \left(\sum_{d=1}^D \frac{\mathbf{q}_i(d)\mathbf{k}_j(d) - \boldsymbol{\mu}(d)\mathbf{k}_j(d)}{\sqrt{D}(\sigma_d^2 + \epsilon)} \right) \mathbf{v}_j \\ &= \sum_{j=1}^N \text{softmax} \left(\sum_{d=1}^D \frac{\mathbf{q}_i(d)\mathbf{k}_j(d) - \frac{1}{N} \sum_{j'=1}^N \mathbf{k}_{j'}(d)\mathbf{k}_j(d)}{\sqrt{D}(\sigma_d^2 + \epsilon)} \right) \mathbf{v}_j. \quad (31) \end{aligned}$$

F ATTENTION WITH THE RESIDUAL CONNECTION AND MATRIX PROJECTIONS

In this supplement, we first discuss attention with the residual connection and matrix projections in Appendix F.

Suppose we are given a training data $\{(\mathbf{x}_1, \mathbf{y}_1), \dots, (\mathbf{x}_N, \mathbf{y}_N)\} \subset \mathcal{X} \times \mathcal{Y}$, where $\mathcal{X} = \mathbb{R}^{D_x}$ and $\mathcal{Y} = \mathbb{R}^{D_y}$. Here, $\mathbf{x}_1, \dots, \mathbf{x}_N$ are the training inputs, and $\mathbf{y}_1, \dots, \mathbf{y}_N$ are the training targets. In

order to derive the attention with the residual connection and query, key, and value matrix projections, we define the function f as follows

$$\mathbf{y} = f(\mathbf{x}) := \mathbf{W} \frac{\Phi(\mathbf{W}^{\text{proj}} \mathbf{x})}{h(\mathbf{x})} + \mathbf{x} + \mathbf{b}, \quad (32)$$

where $\mathbf{x} \in \mathcal{X} = \mathbb{R}^{D_x}$, $\mathbf{W}^{\text{proj}} = [\mathbf{w}_1^{\text{proj}}, \dots, \mathbf{w}_D^{\text{proj}}]^\top \in \mathbb{R}^{D \times D_x}$, $\Phi(\cdot) = [\phi_1(\cdot), \dots, \phi_{D_\phi}(\cdot)] : \mathbb{R}^D \rightarrow \mathbb{R}^{D_\phi}$, $\mathbf{W} = [\mathbf{w}_1, \dots, \mathbf{w}_{D_v}]^\top \in \mathbb{R}^{D_v \times D_\phi}$, $\mathbf{b} \in \mathbb{R}^{D_v}$, and $h(\mathbf{x})$ is a vector-scalar function. We fit the function f to the training data $\{(\mathbf{x}_1, \mathbf{y}_1), \dots, (\mathbf{x}_N, \mathbf{y}_N)\}$ with an \mathbb{L}_2 regularization on \mathbf{W} and \mathbf{W}^{proj} by solving the following convex optimization problem:

$$\begin{aligned} & \underset{\substack{\mathbf{W}, \mathbf{W}^{\text{proj}} \\ \xi_j, \tilde{\xi}_j, j=1, \dots, N}}{\text{minimize}} && \frac{1}{2} \sum_{d=1}^{D_v} \|\mathbf{w}_d\|^2 + \frac{1}{2} \sum_{d=1}^{D_v} \|\mathbf{w}_d^{\text{proj}}\|^2 + C \sum_{j=1}^N \sum_{d=1}^{D_v} (\xi_j(d) + \tilde{\xi}_j(d)) \\ & \text{subject to} && \begin{cases} \mathbf{y}_j(d) - \mathbf{w}_d^\top \Phi(\mathbf{W}^{\text{proj}} \mathbf{x}_j) / h(\mathbf{x}_j) - \mathbf{x}_j - \mathbf{b}(d) \leq \epsilon + \xi_j(d) \\ \mathbf{w}_d^\top \Phi(\mathbf{W}^{\text{proj}} \mathbf{x}_j) / h(\mathbf{x}_j) + \mathbf{x}_j + \mathbf{b}(d) - \mathbf{y}_j(d) \leq \epsilon + \tilde{\xi}_j(d) \\ \xi_j(d), \tilde{\xi}_j(d) \geq 0 \end{cases}, \quad j = 1, \dots, N, \quad d = 1, \dots, D_v. \end{aligned} \quad (33)$$

The Lagrangian of the optimization problem 33 is given by

$$\begin{aligned} \mathcal{L}_1 := & \frac{1}{2} \sum_{d=1}^{D_v} \|\mathbf{w}_d\|^2 + \frac{1}{2} \sum_{d=1}^{D_v} \|\mathbf{w}_d^{\text{proj}}\|^2 + C \sum_{j=1}^N \sum_{d=1}^{D_v} (\xi_j(d) + \tilde{\xi}_j(d)) - \sum_{j=1}^N \sum_{d=1}^{D_v} (\eta_j(d) \xi_j(d) + \tilde{\eta}_j(d) \tilde{\xi}_j(d)) \\ & - \sum_{j=1}^N \sum_{d=1}^{D_v} \alpha_j(d) \left(\epsilon + \xi_j(d) - \mathbf{y}_j(d) + \mathbf{w}_d^\top \frac{\Phi(\mathbf{W}^{\text{proj}} \mathbf{x}_j)}{h(\mathbf{x}_j)} + \mathbf{x}_j + \mathbf{b}(d) \right) \\ & - \sum_{j=1}^N \sum_{d=1}^{D_v} \tilde{\alpha}_j(d) \left(\epsilon + \tilde{\xi}_j(d) + \mathbf{y}_j(d) - \mathbf{w}_d^\top \frac{\Phi(\mathbf{W}^{\text{proj}} \mathbf{x}_j)}{h(\mathbf{x}_j)} - \mathbf{x}_j - \mathbf{b}(d) \right), \end{aligned} \quad (34)$$

Similar to the derivation in Section 2.1, the partial derivatives of \mathcal{L}_1 with respect to the primal variable \mathbf{w}_d , $d = 1, \dots, D_v$, have to vanish for optimality, which leads to

$$\partial_{\mathbf{w}_d} \mathcal{L}_1 = \mathbf{w}_d - \sum_{j=1}^N (\alpha_j(d) - \tilde{\alpha}_j(d)) \frac{\Phi(\mathbf{W}^{\text{proj}} \mathbf{x}_j)}{h(\mathbf{x}_j)} = 0 \Rightarrow \mathbf{w}_d = \sum_{j=1}^N (\alpha_j(d) - \tilde{\alpha}_j(d)) \frac{\Phi(\mathbf{W}^{\text{proj}} \mathbf{x}_j)}{h(\mathbf{x}_j)}. \quad (35)$$

Note that here we only find the form of the optimal solution for $\mathbf{W} = [\mathbf{w}_1, \dots, \mathbf{w}_{D_v}]^\top$. The optimal value of \mathbf{W}^{proj} can then be found by optimization algorithm such as the (stochastic) gradient descent when training the transformer.

Let $\mathbf{v}_j = \left[\frac{\alpha_j(1) - \tilde{\alpha}_j(1)}{h(\mathbf{x}_j)}, \dots, \frac{\alpha_j(D_v) - \tilde{\alpha}_j(D_v)}{h(\mathbf{x}_j)} \right]^\top$, $j = 1, \dots, N$, we obtain the following support vector expansion of the function f :

$$\begin{aligned} f(\mathbf{x}) &= \left[\sum_{j=1}^N (\alpha_j(1) - \tilde{\alpha}_j(1)) \frac{\Phi(\mathbf{W}^{\text{proj}} \mathbf{x}_j)}{h(\mathbf{x}_j)}, \dots, \sum_{j=1}^N (\alpha_j(D_v) - \tilde{\alpha}_j(D_v)) \frac{\Phi(\mathbf{W}^{\text{proj}} \mathbf{x}_j)}{h(\mathbf{x}_j)} \right]^\top \frac{\Phi(\mathbf{W}^{\text{proj}} \mathbf{x})}{h(\mathbf{x})} + \mathbf{x} + \mathbf{b}, \\ &= \left[\sum_{j=1}^N \frac{\alpha_j(1) - \tilde{\alpha}_j(1)}{h(\mathbf{x}_j)} \frac{\Phi(\mathbf{W}^{\text{proj}} \mathbf{x})^\top \Phi(\mathbf{W}^{\text{proj}} \mathbf{x}_j)}{h(\mathbf{x})}, \dots, \sum_{j=1}^N \frac{\alpha_j(D_v) - \tilde{\alpha}_j(D_v)}{h(\mathbf{x}_j)} \frac{\Phi(\mathbf{W}^{\text{proj}} \mathbf{x})^\top \Phi(\mathbf{W}^{\text{proj}} \mathbf{x}_j)}{h(\mathbf{x})} \right]^\top + \mathbf{x} + \mathbf{b}, \\ &= \underbrace{\sum_{j=1}^N \frac{\Phi(\mathbf{W}^{\text{proj}} \mathbf{x})^\top \Phi(\mathbf{W}^{\text{proj}} \mathbf{x}_j)}{h(\mathbf{x})} \mathbf{v}_j}_{\text{Residual connection}} + \mathbf{x} + \mathbf{b}. \end{aligned} \quad (36)$$

Here, the support vector expansion of f already includes a residual connection. The softmax attention can then be derived by selecting $h(\mathbf{x}) := \sum_{j=1}^N \Phi(\mathbf{W}^{\text{proj}} \mathbf{x})^\top \Phi(\mathbf{W}^{\text{proj}} \mathbf{x}_j)$ and choosing Φ as in Eqn. 28 in Section 2.1. Note that in Eqn. 36, $\{\mathbf{x}_j\}_{j=1}^N$ and \mathbf{x} are the training samples and test sample, respectively. In order to derive the key, query, and value matrix projections in attention, we can then relax Eqn. 36 by letting $\mathbf{W}^{\text{proj}} \mathbf{x}_j = \mathbf{W}_K \mathbf{x}_j$, $\mathbf{W}^{\text{proj}} \mathbf{x} = \mathbf{W}_Q \mathbf{x}$, $\mathbf{v}_j = \mathbf{W}_V \mathbf{x}_j$ and choosing the test sample \mathbf{x} among the training samples $\{\mathbf{x}_j\}_{j=1}^N$.

Remark 6 Here, for self-attention, we choose the test sample \mathbf{x} among the training samples $\{\mathbf{x}_j\}_{j=1}^N$ to compute the attention score of a token to other tokens in the same sequence. For cross-attention where a token in a sequence attends to tokens in another sequence, this constraint can be removed.

G 2D-CONVOLUTION ATTENTION

In this section, we discuss attention with 2D-convolution. Suppose we are given a training data $\{(\mathbf{x}_1^{train}, \mathbf{y}_1^{train}), \dots, (\mathbf{x}_{N_H \times N_W}^{train}, \mathbf{y}_{N_H \times N_W}^{train})\} \subset \mathcal{X} \times \mathcal{Y}$, where $\mathcal{X} = \mathbb{R}^{D_x}$ and $\mathcal{Y} = \mathbb{R}^{D_v}$. Here, $\mathbf{x}_1^{train}, \dots, \mathbf{x}_{N_H \times N_W}^{train}$ are the training inputs, and $\mathbf{y}_1^{train}, \dots, \mathbf{y}_{N_H \times N_W}^{train}$ are the training targets. Let $\mathbf{X}^{train} \in \mathbb{R}^{N_H \times N_W \times D_x}$ be the 3D-tensor of training inputs, where $\mathbf{X}^{train}(h, w, d) = \mathbf{x}_{N_W \times (h-1) + w}^{train}(d)$. Given a new set of inputs $\{\mathbf{x}_1, \dots, \mathbf{x}_{N_H \times N_W}\} \subset \mathcal{X}$ and the corresponding 3D-tensor $\mathbf{X} \in \mathbb{R}^{N_H \times N_W \times D_x}$ of these inputs, where $\mathbf{X}(h, w, d) = \mathbf{x}_{N_W \times (h-1) + w}(d)$. We consider the function f applying on the 3D-tensor \mathbf{X} and taking the following form

$$f(\mathbf{x}_i) = \mathbf{W} \frac{\Phi(\text{Flatten}(\text{Conv2D}(\mathbf{X}, s))(i))}{h(\mathbf{x}_i)}, \quad i = 1, \dots, N_H \times N_W \quad (37)$$

where Conv2D is the depth-wise 2D-convolution (Howard et al., 2017), with the kernel size $s \times s$ and identical kernel channels, applied on the input tensor \mathbf{X} . Here, the last dimension of \mathbf{X} , i.e., D_x , is the depth. Also, $\Phi(\mathbf{x}) = [\phi_1(\mathbf{x}), \dots, \phi_{D_\phi}(\mathbf{x})] \in \mathbb{R}^{D_\phi}$, $\mathbf{W} = [\mathbf{w}_1, \dots, \mathbf{w}_{D_v}]^\top \in \mathbb{R}^{D_v \times D_\phi}$, $\mathbf{b} \in \mathbb{R}^{D_v}$, and h is a vector-scalar function. We fit the function f to the training data $\{(\mathbf{x}_1^{train}, \mathbf{y}_1^{train}), \dots, (\mathbf{x}_{N_H \times N_W}^{train}, \mathbf{y}_{N_H \times N_W}^{train})\}$ with an \mathbb{L}_2 regularization on \mathbf{W} , i.e., a ridge regression, by solving the following convex optimization problem:

$$\begin{aligned} & \underset{\mathbf{W}}{\text{minimize}} && \frac{1}{2} \|\mathbf{W}\|_{\mathbb{F}}^2 + C \sum_{j=1}^{N_H \times N_W} \sum_{d=1}^{D_v} (\xi_j(d) + \tilde{\xi}_j(d)) = \frac{1}{2} \sum_{d=1}^{D_v} \|\mathbf{w}_d\|^2 + C \sum_{j=1}^{N_H \times N_W} \sum_{d=1}^{D_v} (\xi_j(d) + \tilde{\xi}_j(d)) \\ & \text{subject to} && \begin{cases} \mathbf{y}_j^{train}(d) - \mathbf{w}_d^\top \frac{\Phi(\text{Flatten}(\text{Conv2D}(\mathbf{X}^{train}, s))(j))}{h(\mathbf{x}_j^{train})} - \mathbf{b}(d) \leq \epsilon + \xi_j(d) \\ \mathbf{w}_d^\top \frac{\Phi(\text{Flatten}(\text{Conv2D}(\mathbf{X}^{train}, s))(j))}{h(\mathbf{x}_j^{train})} + \mathbf{b}(d) - \mathbf{y}_j^{train}(d) \leq \epsilon + \tilde{\xi}_j(d) \\ \xi_j(d), \tilde{\xi}_j(d) \geq 0, \quad j = 1, \dots, N_H \times N_W, \quad d = 1, \dots, D_v. \end{cases} \end{aligned} \quad (38)$$

The Lagrangian of the optimization problem 38 is given by

$$\begin{aligned} \mathcal{L} := & \frac{1}{2} \sum_{d=1}^{D_v} \|\mathbf{w}_d\|^2 + C \sum_{j=1}^{N_H \times N_W} \sum_{d=1}^{D_v} (\xi_j(d) + \tilde{\xi}_j(d)) - \sum_{j=1}^{N_H \times N_W} \sum_{d=1}^{D_v} (\eta_j(d) \xi_j(d) + \tilde{\eta}_j(d) \tilde{\xi}_j(d)) \\ & - \sum_{j=1}^{N_H \times N_W} \sum_{d=1}^{D_v} \alpha_j(d) \left(\epsilon + \xi_j(d) - \mathbf{y}_j^{train}(d) + \mathbf{w}_d^\top \frac{\Phi(\text{Flatten}(\text{Conv2D}(\mathbf{X}^{train}, s))(j))}{h(\mathbf{x}_j^{train})} + \mathbf{b}(d) \right) \\ & - \sum_{j=1}^{N_H \times N_W} \sum_{d=1}^{D_v} \tilde{\alpha}_j(d) \left(\epsilon + \tilde{\xi}_j(d) + \mathbf{y}_j^{train}(d) - \mathbf{w}_d^\top \frac{\Phi(\text{Flatten}(\text{Conv2D}(\mathbf{X}^{train}, s))(j))}{h(\mathbf{x}_j^{train})} - \mathbf{b}(d) \right), \end{aligned} \quad (39)$$

Similar to the derivation in Section 2.1 in the main text, the partial derivatives of \mathcal{L} with respect to the primal variable \mathbf{w}_d , $d = 1, \dots, D_v$, have to vanish for optimality, which leads to

$$\partial_{\mathbf{w}_d} \mathcal{L} = \mathbf{w}_d - \sum_{j=1}^{N_H \times N_W} (\alpha_j(d) - \tilde{\alpha}_j(d)) \frac{\Phi(\text{Flatten}(\text{Conv2D}(\mathbf{X}^{train}, s))(j))}{h(\mathbf{x}_j^{train})} = 0 \quad (40)$$

$$\Rightarrow \mathbf{w}_d = \sum_{j=1}^{N_H \times N_W} (\alpha_j(d) - \tilde{\alpha}_j(d)) \frac{\Phi(\text{Flatten}(\text{Conv2D}(\mathbf{X}^{train}, s))(j))}{h(\mathbf{x}_j^{train})}. \quad (41)$$

Let $\mathbf{v}_j = [\frac{\alpha_j(1) - \tilde{\alpha}_j(1)}{h(\mathbf{x}_j^{train})}, \dots, \frac{\alpha_j(D_v) - \tilde{\alpha}_j(D_v)}{h(\mathbf{x}_j^{train})}]^\top$, $j = 1, \dots, N_H \times N_W$, and substitute Eqn. 41 into Eqn. 38, we obtain the following support vector expansion of the linear basis function f :

$$\begin{aligned} f(\mathbf{x}_i) &= \left[\sum_{j=1}^{N_H \times N_W} \frac{\alpha_j(1) - \tilde{\alpha}_j(1)}{h(\mathbf{x}_j^{train})} \frac{\mathbf{A}_{ij}}{h(\mathbf{x}_i)}, \dots, \sum_{j=1}^{N_H \times N_W} \frac{\alpha_j(D_v) - \tilde{\alpha}_j(D_v)}{h(\mathbf{x}_j^{train})} \frac{\mathbf{A}_{ij}}{h(\mathbf{x}_i)} \right]^\top + \mathbf{b}, \\ &= \sum_{j=1}^{N_H \times N_W} \frac{\mathbf{A}_{ij}}{h(\mathbf{x}_i)} \mathbf{v}_j + \mathbf{b}, \end{aligned} \quad (42)$$

where $\mathbf{A}_{ij} := \Phi(\text{Flatten}(\text{Conv2D}(\mathbf{X}, s))(i))^\top \Phi(\text{Flatten}(\text{Conv2D}(\mathbf{X}^{train}, s))(j))$.

Same as in Section 2.1, we set $\mathbf{b}_s = 0$. To derive the softmax normalization in attention, we choose $h(\mathbf{x}_i) := \sum_{j=1}^N \mathbf{A}_{ij}$ and select Φ as in Eqn. 28. Let the training inputs $\{\mathbf{x}_1^{train}, \dots, \mathbf{x}_{N_H \times N_W}^{train}\} \subset \mathcal{X}$ be the attention keys $\{\mathbf{k}_1, \dots, \mathbf{k}_{N_H \times N_W}\} \subset \mathcal{K}$, where $\mathcal{K} = \mathbb{R}^D$, in self-attention. Also, let the new inputs $\{\mathbf{x}_1, \dots, \mathbf{x}_{N_H \times N_W}\} \subset \mathcal{X}$ be the attention queries $\{\mathbf{q}_1, \dots, \mathbf{q}_{N_H \times N_W}\} \subset \mathcal{K}$ in self-attention. We define the 2D-Convolution Attention (Attention-Conv2D) as follows:

Definition 3 (2D-Convolution Attention) Given a set of the key and value vectors $\{\mathbf{k}_j, \mathbf{v}_j\}_{j=1}^{N_H \times N_W}$, and a set of the query vectors $\{\mathbf{q}_i\}_{i=1}^{N_H \times N_W}$. Denote the key tensor and the query tensor by $\mathbf{K} \in \mathbb{R}^{N_H \times N_W \times D}$ and $\mathbf{Q} \in \mathbb{R}^{N_H \times N_W \times D}$, respectively, where $\mathbf{K}(h, w, d) = \mathbf{k}_{N_W \times (h-1) + w}(d)$ and $\mathbf{Q}(h, w, d) = \mathbf{q}_{N_W \times (h-1) + w}(d)$. The 2D-Convolution Attention (Attention-Conv2D) computes the corresponding output vector \mathbf{h}_i of the query \mathbf{q}_i by the following attention formula:

$$\mathbf{h}_i = \sum_{j=1}^N \text{softmax} \left(\text{Flatten}(\text{Conv2D}(\mathbf{Q}, s))(i)^\top \text{Flatten}(\text{Conv2D}(\mathbf{K}, s))(j) / \sqrt{D} \right) \mathbf{v}_j, \quad (43)$$

where the $\text{Conv2D}(\cdot, s)$ is the depth-wise 2D-convolution (Howard et al., 2017) with the kernel size $s \times s$ and identical kernel channels.

Remark 7 (Convolutional Projection for Attention in the Convolutional vision Transformer)

The convolutional projections used in the Convolutional vision Transformer (CvT) (Wu et al., 2021) can be derived from Eqn. 42 by letting the training input tensor \mathbf{X}^{train} to be the 2D input matrix of size $N \times D_x$ of the self-attention layer (see Section 1.1 in the main text) reshaped into a 3D tensor of size $N_H \times N_W \times D_x$ where $N = N_H \times N_W$. Here, to avoid confusion, we denote the input of the self-attention layer by \mathbf{X}^{input} and its reshaped version by $\text{Reshape2D}(\mathbf{X}^{input})$. We then replace the depth-wise 2D-convolution by the depth-wise separable 2D-convolution in (Wu et al., 2021) and remove the constraint that the kernels have identical channels. In order to derive the convolutional projections for the keys, queries, and values in CvT, for $i, j = 1, \dots, N$, we let

$$\begin{aligned} \mathbf{k}_j &= \text{Flatten}(\text{Conv2D}(\mathbf{X}^{train}, s))(j) = \Phi(\text{Flatten}(\text{Conv2D}(\text{Reshape2D}(\mathbf{X}^{input}), s, \mathbf{W}_K))(j)), \\ \mathbf{q}_i &= \text{Flatten}(\text{Conv2D}(\mathbf{X}, s))(i) = \Phi(\text{Flatten}(\text{Conv2D}(\text{Reshape2D}(\mathbf{X}^{input}), s, \mathbf{W}_Q))(i)), \\ \mathbf{v}_j &= \text{Flatten}(\text{Conv2D}(\mathbf{X}^{train}, s))(j) = \Phi(\text{Flatten}(\text{Conv2D}(\text{Reshape2D}(\mathbf{X}^{input}), s, \mathbf{W}_V))(j)). \end{aligned} \quad (44)$$

Here, we specify the kernel/filter \mathbf{W}_K , \mathbf{W}_Q , and \mathbf{W}_V to emphasize that the convolutional projections in CvT uses different kernels to compute keys, queries, and values in self-attention. Eqn. 44 matches the convolutional projects in CvT. By choosing h and Φ similar to above, we can derive the convolutional attention in CvT.

H 1D-CONVOLUTION ATTENTION

Following the derivation for the Attention-Conv2D in Appendix G above, we can derive the 1D-Convolution Attention (Attention-Conv1D) in a similar way by letting $\mathbf{X}^{train} \in \mathbb{R}^{N \times D_x}$ and $\mathbf{X} \in \mathbb{R}^{N \times D_x}$ be 2D-matrices of training inputs and new inputs, respectively, and by replacing Conv2D by Conv1D, which is the depth-wise 1D-convolution, with the kernel size $s \times 1$ and identical kernel channels, applied on the input tensor \mathbf{X} . Here, the last dimension of \mathbf{X} , i.e., D_x , is the depth. We define the 1D-Convolution Attention (Attention-Conv1D) as follows:

Definition 4 (1D-Convolution Attention) Given a set of the key and value vectors $\{\mathbf{k}_j, \mathbf{v}_j\}_{j=1}^N$, and a set of the query vectors $\{\mathbf{q}_i\}_{i=1}^N$. Denote the key matrix and the query matrix by $\mathbf{K} :=$

$[\mathbf{k}_1, \dots, \mathbf{k}_N]^\top \in \mathbb{R}^{N \times D}$ and $\mathbf{Q} := [\mathbf{q}_1, \dots, \mathbf{q}_N]^\top \in \mathbb{R}^{N \times D}$, respectively. The 1D-Convolution Attention (Attention-Conv1D) computes the corresponding output vector \mathbf{h}_i of the query \mathbf{q}_i by the following attention formula:

$$\mathbf{h}_i = \sum_{j=1}^N \text{softmax} \left(\text{Conv1D}(\mathbf{Q}, s)(i)^\top \text{Conv1D}(\mathbf{K}, s)(j) / \sqrt{D} \right) \mathbf{v}_j, \quad (45)$$

where the $\text{Conv1D}(\cdot, s)$ is the depth-wise 1D-convolution with the kernel size $s \times 1$ and identical kernel channels.

I ATTENTION WITH BATCH NORMALIZATION AND SCALED HEADS

The Attention-BN+SH combines both the Attention-BN and Attention-SH. The Attention-BN+SH fits the function f^s , $s = 1, \dots, H$, in Eqn. 17 with training sets $\{(\mathbf{k}_1^1, \mathbf{y}_1^1), \dots, (\mathbf{k}_{N_1}^1, \mathbf{y}_{N_1}^1)\}, \dots, \{(\mathbf{k}_1^H, \mathbf{y}_1^H), \dots, (\mathbf{k}_{N_H}^H, \mathbf{y}_{N_H}^H)\} \subset \mathcal{K} \times \mathcal{Y}$ of different sizes N_1, \dots, N_H , where $\mathcal{K} = \mathbb{R}^D$ and $\mathcal{Y} = \mathbb{R}^{D_v}$. The function f^s is defined as:

$$f^s(\mathbf{x}) := \mathbf{W}^s \frac{\Phi((\mathbf{x} - \boldsymbol{\mu}^s) \odot \mathbf{s}^{s-1})}{h^s((\mathbf{x} - \boldsymbol{\mu}^s) \odot \mathbf{s}^{s-1})} + \mathbf{b}^s, \quad (46)$$

where

$$\boldsymbol{\mu}^s = \frac{1}{N_s} \sum_{j=1}^{N_s} \mathbf{k}_j^s, \quad \mathbf{s}^{s-1} = \left[\frac{1}{\sqrt{\sigma_1^{s^2} + \epsilon}}, \dots, \frac{1}{\sqrt{\sigma_D^{s^2} + \epsilon}} \right]^\top, \quad \sigma_d^{s^2} = \frac{1}{N_s} \sum_{j=1}^{N_s} (\mathbf{k}_j^s(d) - \boldsymbol{\mu}^s(d))^2. \quad (47)$$

Following the same derivation as in Section 2.1, we derive the following support vector expansion of f^s

$$f^s(\mathbf{x}) = \sum_{j=1}^{N_s} \frac{\Phi((\mathbf{x} - \boldsymbol{\mu}^s) \odot \mathbf{s}^{s-1})^\top \Phi((\mathbf{k}_j^s - \boldsymbol{\mu}^s) \odot \mathbf{s}^{s-1})}{h^s((\mathbf{x} - \boldsymbol{\mu}^s) \odot \mathbf{s}^{s-1})} \mathbf{v}_j^s + \mathbf{b}^s. \quad (48)$$

Here, $\mathbf{v}_j^s = \left[\frac{\alpha_j^s(1) - \tilde{\alpha}_j^s(1)}{h^s((\mathbf{k}_j^s - \boldsymbol{\mu}^s) \odot \mathbf{s}^{s-1})}, \dots, \frac{\alpha_j^s(D_v) - \tilde{\alpha}_j^s(D_v)}{h^s((\mathbf{k}_j^s - \boldsymbol{\mu}^s) \odot \mathbf{s}^{s-1})} \right]^\top$, where α_j^s and $\tilde{\alpha}_j^s$ are the dual variables, $j = 1, \dots, N$. Same as in Section 2.1, in Eqn. 48, we choose Φ as in Eqn. 28, $h^s(\mathbf{x}) := \sum_{j=1}^{N_s} \Phi(\mathbf{x})^\top \Phi(\mathbf{k}_j^s)$, and $\mathbf{b}^s = 0$ to obtain the Batch Normalized Attention with Scaled Heads (Attention-BN+SH), which is defined as follows:

Definition 5 (Batch Normalized Attention with Scaled Heads) Given H sets of the key and value vectors $\{\mathbf{k}_j^1, \mathbf{v}_j^1\}_{j=1}^{N_1}, \dots, \{\mathbf{k}_j^H, \mathbf{v}_j^H\}_{j=1}^{N_H}$, for each set of H query vectors $\mathbf{q}_1^1, \dots, \mathbf{q}_i^H$, $i = 1, \dots, N$, the Batch Normalized Attention with Scaled Heads (Attention-BN+SH) computes the corresponding output vector \mathbf{h}_i of the queries $\mathbf{q}_1^1, \dots, \mathbf{q}_i^H$ by the following attention formula:

$$\mathbf{h}_i = \sum_{s=1}^H \mathbf{W}_O^s \left(\sum_{j=1}^{N_s} \text{softmax} \left(((\mathbf{q}_i^s - \boldsymbol{\mu}^s) \odot \mathbf{s}^{s-1})^\top ((\mathbf{k}_j^s - \boldsymbol{\mu}^s) \odot \mathbf{s}^{s-1}) / \sqrt{D} \right) \mathbf{v}_j^s \right), \quad (49)$$

where

$$\boldsymbol{\mu}^s = \frac{1}{N_s} \sum_{j=1}^{N_s} \mathbf{k}_j^s, \quad \mathbf{s}^{s-1} = \left[\frac{1}{\sqrt{\sigma_1^{s^2} + \epsilon}}, \dots, \frac{1}{\sqrt{\sigma_D^{s^2} + \epsilon}} \right]^\top, \quad \sigma_d^{s^2} = \frac{1}{N_s} \sum_{j=1}^{N_s} (\mathbf{k}_j^s(d) - \boldsymbol{\mu}^s(d))^2. \quad (50)$$

Following the same Remark 5 in Section 2.3.2, given input sequence $\mathbf{X} := [\mathbf{x}_1, \dots, \mathbf{x}_N]^\top \in \mathbb{R}^{N \times D_x}$ of N feature vectors in self-attention, in order to generate the sets of $\{\mathbf{k}_j^s, \mathbf{v}_j^s\}_{j=1}^{N_s}$ at the scale s^{th} , we can downsample the input \mathbf{X} before projecting into the key matrix \mathbf{K} and the value matrix \mathbf{V} . In this paper, we use the average-pooling to downsample \mathbf{X} .

As in the same case of Attention-BN, for Attention-BN+SH, recentering queries and keys alone are sufficient for accuracy improvement, and we weight the mean $\boldsymbol{\mu}$ in Eqn 49 with a constant β . Hence Eqn. 49 is simplified to:

$$\mathbf{h}_i = \sum_{s=1}^H \mathbf{W}_O^s \left(\sum_{j=1}^{N_s} \text{softmax} \left((\mathbf{q}_i^s - \beta \boldsymbol{\mu}^s)^\top (\mathbf{k}_j^s - \beta \boldsymbol{\mu}^s) / \sqrt{D} \right) \mathbf{v}_j^s \right). \quad (51)$$

Table 12: The values of β for Linear Attention-BN/BN+SH and Sparse Attention-BN/BN+SH trained on the selected 10 UEA tasks.

Dataset/Model	Linear Attention-BN	Linear Attention-BN+SH	Sparse Attention-BN	Sparse Attention-BN+SH
ETHANOLCONCENTRATION	0.15	0.95	0.8	0.2
FACEDETECTION	0.6	0.6	0.6	0.6
HANDWRITING	0.25	0.3	0.3	0.3
HEARTBEAT	0.6	0.15	0.4	0.5
JAPANESEVOWELS	0.6	0.6	0.6	0.6
PEMS-SF	0.35	0.65	0.5	0.6
SELFREGULATIONSCP1	0.35	0.25	0.1	0.9
SELFREGULATIONSCP2	0.75	0.15	0.5	0.3
SPOKENARABICDIGITS	0.6	0.6	0.6	0.6
UWAVEGESTURELIBRARY	0.65	0.55	0.9	0.3

Table 13: The values of β for Attention-BN/BN+SH trained on 25 UEA Time Series classification tasks (Bagnall et al., 2018) and 6 UEA Time Series Regression tasks.

Dataset/Model	Attention-BN	Attention-BN+SH
ETHANOLCONCENTRATION	0.25	0.15
FACEDETECTION	0.6	0.6
HANDWRITING	0.65	0.25
HEARTBEAT	0.55	0.85
JAPANESEVOWELS	0.6	0.6
PEMS-SF	0.25	0.35
SELFREGULATIONSCP1	0.5	0.85
SELFREGULATIONSCP2	1.2	0.9
SPOKENARABICDIGITS	0.65	0.6
UWAVEGESTURELIBRARY	0.1	0.2
ARTICULARYWORDRECOGNITION	0.2	0.6
BASICMOTIONS	0.1	0.1
EPILEPSY	0.3	0.2
ERING	0.1	0.9
FINGERMOVEMENTS	1.0	0.3
LIBRAS	0.3	0.7
NATOPS	1.0	0.4
RACKETSPORTS	1.0	0.4
ATRIALFIBRILLATION	0.9	0.6
CRICKET	0.2	1.0
STANDWALKJUMP	1.0	0.6
HANDMOVEMENTDIRECTION	0.5	0.5
LSST	0.3	0.2
DUCKDUCKGEESE	0.6	0.3
MOTORIMAGERY	0.2	0.3
APPLIANCESENERGY	0.4	0.2
BENZENECONCENTRATION	0.2	0.1
BEIJINGPM10	0.1	0.5
BEIJINGPM25	0.1	0.2
LIVEFUELMOISTURE	0.1	0.3
IEEEPPG	0.4	0.2

Table 14: The values of β of Attention-BN/BN+SH trained on the 5 tasks of the LRA benchmark (Tay et al., 2021).

Dataset/Model	Attention-BN	Attention-BN+SH
LISTOPS	0.5	0.2
TEXT	0.5	0.8
RETRIEVAL	1.0	1.0
IMAGE	0.2	0.2
PATHFINDER	0.2	0.4

J HYPERPARAMETERS

In this section, we provide the hyper-parameters for our best models.

J.1 UEA TIME SERIES CLASSIFICATION AND REGRESSION

For these two benchmarks, use the set of downsampling factors $\mathbf{s} = [1, 1, 2, 2, 4, 4, 8, 8]$ for Attention-SH/BN+SH and Linear/Sparse Attention-SH/BN+SH models trained on the UEA benchmark. Table 13 and Table 12 provide the values of β used for our best Attention-BN/BN+SH and Linear Attention-BN/BN+SH, Sparse Attention-BN/BN+SH models trained on subsets of the two benchmarks.

J.2 LONG RANGE ARENA BENCHMARK

For all 5 tasks of the LRA benchmark, we set the downsampling factors \mathbf{s} of Attention-SH/BN+SH, Linear/Sparse Attention-SH/BN+SH is $[1, 2]$ and kernel size of Attention-Conv1D models is 5. In addition, Table 14 provides the values β of Attention-BN/BN+SH models trained on the benchmark.

J.3 IMAGENET CLASSIFICATION

This task's β of Attention-BN/BN+SH is 1. Attention-SH/BN+SH has the downsampling factor of $[1, 1, 2, 4]$, and the kernel size of Attention-Conv2D is $(2, 2)$.