

CAMERO: Consistency Regularized Ensemble of Perturbed Language Models with Weight Sharing

Chen Liang *

Georgia Institute of Technology
cliang73@gatech.edu

Pengcheng He, Yelong Shen

Microsoft Azure AI

{penhe, yelong.shen}@microsoft.com

Weizhu Chen

Microsoft Azure AI
wzchen@microsoft.com

Tuo Zhao

Georgia Institute of Technology
tourzhao@gatech.edu

Abstract

Model ensemble is a popular approach to produce a low-variance and well-generalized model. However, it induces large memory and inference costs, which are often not affordable for real-world deployment. Existing work has resorted to sharing weights among models. However, when increasing the proportion of the shared weights, the resulting models tend to be similar, and the benefits of using model ensemble diminish. To retain ensemble benefits while maintaining a low memory cost, we propose a consistency-regularized ensemble learning approach based on perturbed models, named CAMERO. Specifically, we share the weights of bottom layers across all models and apply different perturbations to the hidden representations for different models, which can effectively promote the model diversity. Meanwhile, we apply a prediction consistency regularizer across the perturbed models to control the variance due to the model diversity. Our experiments using large language models demonstrate that CAMERO significantly improves the generalization performance of the ensemble model. Specifically, CAMERO outperforms the standard ensemble of 8 BERT-base models on the GLUE benchmark by 0.7 with a significantly smaller model size (114.2M vs. 880.6M).

1 Introduction

Deep Neural Networks (DNNs) have achieved remarkable success in various fields and have become very powerful in learning complicated models (Devlin et al., 2018; Brown et al., 2020; He et al., 2020). However, their remarkable representation powers come at the expense of large model variance, which may hurt the model generalization performance. A popular approach for reducing such variance is model ensemble, where the weights or predictions of a set of models are aggregated to produce the

predictions (Yang and Lv, 2021; Dong et al., 2020). For example, Zhang et al. (2018) show that a simple 2-model ensemble leads to notable improvement over a single model in computer vision tasks.

Despite such notable benefits, model ensemble has not been widely applied to large language models. The major barriers are its enormous storage and expensive inference cost, which linearly scales with the size and the number of models. Therefore, it is often not affordable to ensemble large language models for deployment using memory-constrained and low-latency edge devices.

To alleviate the memory burden, recent works have resorted to a weight-sharing strategy, where all models share the same set of bottom-layer weights, on top of which branches out a set of parallel, un-shared top-layer weights (Lan et al., 2018; Chen et al., 2020; Li et al., 2020). Since the shared weights are optimized to accommodate multiple *diverse* un-shared branches, they can learn shared representations with better generalization (Liu et al., 2020; Luong et al., 2015; Ruder et al., 2019).

For large models, however, such a weight-sharing strategy no longer enjoys the same benefits. Due to memory constraints, a significant proportion of bottom-layer weights need to be shared. Accordingly, top-layer branches have only limited capacity, and therefore the resulting models tend to be similar (Chen et al., 2020; Rame and Cord, 2021; Feng et al., 2020; Yang et al., 2021; Wu and Gong, 2021) (Figure 1 (Left)). Due to the lack of the model diversity, their ensemble cannot gain much improvement in generalization. As shown in Figure 1 (Right), both the generalization performance and model variance of the ensemble model are similar to those of a single model when the branch size is small.

To retain a light memory cost while maintaining the ensemble benefits, we propose a new Consistency-regularized ensemble Learning ap-

*Work was done during an internship at Microsoft Azure AI.




Figure 1: Left: The prediction similarity among branches with different sizes. Right: The average generalization performance and variance of ensembled models over five random seeds. The results are obtained by fine-tuning SST-2 on a BERT-base. A branch size of 0 corresponds to training a single model.

proach based on peRturbed mOdelS – CAMERO. Specifically, we share the bottom-layer weights across all models and apply *different* perturbations to the hidden representations for *different* models. Such a perturbation strategy effectively promotes the model diversity. Accordingly, the weights at each layer are optimized to produce consistent outputs given diverse input representations from differently perturbed models. In other words, the shared weights are essentially an on-the-fly ensemble of all perturbed models. In the end, we ensemble all branches on top of the shared weights to produce the final model, which has both a low variance and good generalization performance.

Since we apply perturbations in large models with significant depth, different models’ hidden representations may end up being extremely diverse, especially in upper layers. As a result, optimizing the shared weights to accommodate such perturbations can be very challenging. To prevent the models from being over-diverse, we apply a consistency regularizer to reduce variance across different models. Specifically, such a consistency regularizer can be viewed as collaborative distillation across models (Guo et al., 2020; Lan et al., 2018; Zhang et al., 2018; Kim et al., 2021; Chen et al., 2020; Li et al., 2020). By regularizing the discrepancy between each model’s output logits and the ensemble of these logits, it encourages all models to be *consistent* in their predictions. We thus adopt consistency regularization to control the perturbed models’ diversity from being too large, and thus ease the optimization of the shared weights.

We conduct thorough experiments to demonstrate CAMERO’s effectiveness and efficiency in ensembling large number of models with more than hundreds of millions of learning parameters. Specifically, our experiments in fine-tuning

the BERT-base model on the GLUE benchmark achieve 0.7 points of gain in terms of task-average score with a significantly smaller parameters over the vanilla ensemble approach (114.2M vs. 880.6M) and achieve 1.2 points of gain with the same amount of learning parameters over the single model. CAMERO also achieves significant improvements in neural machine translation on both low-resource and high-resource language pairs.

Furthermore, we verify that CAMERO can learn shared layers with better generalization and ensemble model with smaller variance. We also investigate the effects of using different types and strengths of perturbation and consistency regularization techniques. In particular, we observe that models created with virtual adversarial perturbation (Jiang et al., 2019a) and neuron dropout (Srivastava et al., 2014) lead to ensemble models with the best generalization performance. Lastly, we demonstrate CAMERO’s effectiveness on a larger-scale model, RoBERTa-large (Liu et al., 2019), where it achieves 0.8 and 0.9 points of gain over the vanilla ensemble approach and single model performance, respectively. Our codes are released at <https://github.com/cliangl453/CAMERO>.

2 Background

Notations. We use $f(\cdot; \theta)$ to denote a mapping f associated with the parameter θ from the input sample to an output space, where the output is a multi-dimensional probability simplex for classification tasks and a scalar for regression tasks. We denote the model’s final logits as $g(\cdot; \theta)$, where $f(\cdot; \theta) = \sigma(g(\cdot; \theta))$ and $\sigma(\cdot)$ is the Softmax function. We denote n pairs of data samples of the target task as $\{(\mathbf{x}_i, y_i)\}_{i=1}^n$. The training loss of $f(\cdot; \theta)$ is computed as $\ell(f(\mathbf{x}_i; \theta), y_i)$ for any given training instance (\mathbf{x}_i, y_i) where $\ell(\cdot; \cdot)$ denotes the loss function. We use $\mathcal{D}_{KL}(P||Q) = \sum_k p_k \log(p_k/q_k)$ to denote the KL-divergence of two discrete distributions P and Q with the associated parameters of p_k ’s and q_k ’s, respectively.

Collaborative Distillation. Collaborative distillation approaches train two or more models in parallel while regularizing the consistency of their final prediction distributions (Guo et al., 2020; Lan et al., 2018; Zhang et al., 2018; Kim et al., 2021; Chen et al., 2020; Li et al., 2020). Specifically, we use $\{f(\cdot; \theta_j)\}_{j=1}^m$ to denote m individual models with the same architectures with parameters by $\theta_1, \dots, \theta_m$, respectively, and denote

$\Theta = \{\theta_1, \dots, \theta_m\}$. A typical collaborative distillation approach solves the following optimization problem:

$$\min_{\Theta} \mathcal{L}(\Theta) + \alpha \mathcal{R}(\Theta),$$

where $\alpha > 0$ is a tuning parameter, and $\mathcal{L}(\Theta)$ and $\mathcal{R}(\Theta)$ are defined as

$$\begin{aligned} \mathcal{L}(\Theta) &= \frac{1}{m} \sum_{j=1}^m \ell(f(\mathbf{x}; \theta_j), y), \\ \mathcal{R}(\Theta) &= \frac{1}{m} \sum_{j=1}^m \mathcal{D}(f(\mathbf{x}; \theta_j), \mathcal{E}(\mathbf{x}; \Theta)). \end{aligned} \quad (1)$$

For notational simplicity, we will omit the subscript i throughout the rest of the paper. Here, $\mathcal{E}(\mathbf{x}; \Theta)$ defines a mapping function associated with Θ , which maps the input sample \mathbf{x} to a multi-dimensional probability simplex or a scalar depending on the tasks. A commonly adopted ensemble-distillation approach makes $\mathcal{E}(\cdot; \Theta) = \sigma(\sum_{j=1}^m w_j g(\mathbf{x}; \theta_j))$, where $\{w_j\}_{j=1}^m$ are non-negative scalars summing to one. $\mathcal{D}(\cdot, \cdot)$ denotes the distance metric of two discrete distributions P and Q or two scalars p and q . $\mathcal{D}(P, Q)$ can take the form of symmetric KL-Divergence, $\frac{1}{2}(\mathcal{D}_{KL}(P||Q) + \mathcal{D}_{KL}(Q||P))$. $\mathcal{D}(p, q)$ or euclidean distance $\|p - q\|_2^2$.

Weight-Sharing. Weight-sharing technique has been adopted in several representation learning scenarios, e.g., multi-task learning (Liu et al., 2020; Luong et al., 2015; Ruder et al., 2019), multi-domain learning (Britz et al., 2017; Zeng et al., 2018; Tars and Fishel, 2018; Jiang et al., 2019b) and multi-lingual tasks (Gu et al., 2018; Aharoni et al., 2019). Weight-sharing strategy can reduce the number of free parameters in the model, which helps prevent overfitting in the training and lead to better generalization abilities.

3 Method

We introduce CAMERO, a weight-sharing ensemble learning approach based on consistency-regularized perturbed models.

3.1 Ensemble Learning w/ Perturbed Models

Based on the multi-layer structures of neural networks, we divide each model into two parts: the bottom-layers and the top-layers. The model parameters in bottom-layers are shared across all models. Specifically, the parameters of the j -th model is denoted as $\theta_j = [\theta_0, \theta'_j]$, where θ_0 denotes the

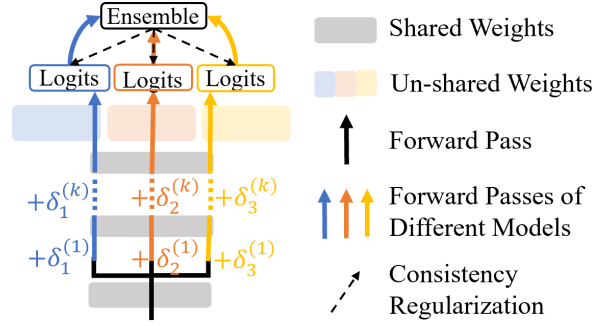


Figure 2: Illustration of CAMERO during training.

shared weights in bottom-layers, and θ'_j denotes the top-layer weights of the j -th model. Based on such a compositional structure, the j -th model's output can be denoted as

$$\begin{aligned} f(\mathbf{x}; \theta_j) &= f_K(f_{K-1}(\dots f_1(\mathbf{x}; \theta_j^{(1)} \\ &\dots; \theta_j^{(K-1)}); \theta_j^{(K)}), \end{aligned}$$

where $f_k(\cdot; \theta_j^{(k)})$ is the mapping associated with the k -th layer parameter $\theta_j^{(k)}$, and θ_0 consists of $\{\theta_j^{(k)}\}_{k=1}^{K'}$ for some $K' \in \{1, \dots, K\}$, which shares across models.

A significant proportion of shared weights leads to the models' similarity, which accordingly impairs the ensemble model's performance. To increase the models' diversity, we consider perturbing each layer's hidden representations for different models during training (Figure 2). Specifically, the j -th model's output is denoted as

$$\begin{aligned} f(\mathbf{x}; \theta_j, \Delta_j) &= f_K(f_{K-1}(\dots f_1(\mathbf{x}; \theta_j^{(1)} \\ &+ \delta_j^{(1)} \dots; \theta_j^{(K-1)}) + \delta_j^{(K-1)}; \theta_j^{(K)}), \end{aligned}$$

where $\delta_j^{(k)}$ is the perturbation applied at the k -th layer's hidden representations, and $\Delta_j = \{\delta_j^{(k)}\}_{k=1}^{K-1}$, which is sampled from a distribution \mathcal{P} . We then train m models with SGD-type algorithms using the following loss:

$$\begin{aligned} \mathcal{L}_{\Delta}(\Theta) &= \frac{1}{m} \sum_{j=1}^m \mathbb{E}_{\Delta_j \sim \mathcal{P}} [\\ &\ell(f(\mathbf{x}; [\theta_0, \theta'_j], \Delta_j), y)]. \end{aligned} \quad (2)$$

Remark 1. We can consider a wide variety of perturbations for hidden representations, input embeddings or data samples, e.g., random perturbation (Aghajanyan et al., 2020), virtual adversarial perturbation (Miyato et al., 2018; Jiang et al., 2019a), neuron dropout (Srivastava et al., 2014) and word dropout (Wei and Zou, 2019).

3.2 Consistency-Regularized Perturbed Models

In large models with significant depth, different models’ hidden representations may end up being extremely diverse, especially in upper layers. As a result, optimizing the shared weights to accommodate such diverse inputs can be very challenging. To address this issue, we propose to control model variability through consistency regularization. Specifically, we regularize the consistency among m models’ final prediction distributions by minimizing the following loss,

$$\mathcal{R}_\Delta(\Theta) = \frac{1}{m} \sum_{j=1}^m \mathbb{E}_{\Delta_j \sim \mathcal{P}} [\mathcal{D}(f(\mathbf{x}; [\theta_0, \theta'_j], \Delta_j), \mathcal{E}(\mathbf{x}; \Theta, \{\Delta_j\}_{j=1}^m))], \quad (3)$$

where $\mathcal{E}(\mathbf{x}; \Theta, \{\Delta_j\}_{j=1}^m)$ denotes the final prediction distribution produced by some ensemble method applied upon models with perturbed representations. For example, commonly adopted ensemble methods include logits ensemble, where $\mathcal{E}(\mathbf{x}; \Theta, \{\Delta_j\}_{j=1}^m) = \sigma(\sum_{j=1}^m w_j g(\mathbf{x}; \theta_j, \{\Delta_j\}_{j=1}^m))$. In summary, we train m models by minimizing the following overall loss function:

$$\mathcal{L}_\Delta(\Theta) + \alpha \mathcal{R}_\Delta(\Theta),$$

where $\mathcal{L}_\Delta(\Theta)$ is defined in Eq. (2) and $\mathcal{R}_\Delta(\Theta)$ is defined in Eq. (3). We adjust the strength of consistency regularization via α , a non-negative hyper-parameter.

Remark 2. Different from existing weight-sharing strategies, which control models’ diversity via the amount of shared and un-shared weights, CAMERO controls model diversity via the strength of perturbation and regularization. Such a difference renders significant memory benefits. In practice, we safely share all layers except a single top layer. Accordingly, the memory storage is reduced to that of a single model. This allows us to explore the behaviors of ensemble learning under a larger number of models.

4 Experiment

We verify the effectiveness of CAMERO on widely used benchmarks for natural language understanding and neural machine translation.

4.1 Natural Language Understanding

Model and data. We evaluate the fine-tuning performance of BERT-base (110M) (Devlin et al., 2018) and RoBERTa-large (335M) (Liu et al., 2019) on the General Language Understanding Evaluation (GLUE, Wang et al. (2018)) benchmark. GLUE contains nine NLU tasks, including textual entailment, question answering, sentiment analysis, and text similarity. Details about the benchmark are deferred to Appendix A.1.1.

Baseline methods. We compare CAMERO with Vanilla, where all models are independently trained without consistency regularization. We also compare CAMERO with representative collaborative distillation methods: Deep Mutual Learning (DML, Zhang et al. (2018)), On-the-fly Native Ensemble Learning (ONE, Lan et al. (2018)) and Knowledge Distillation via Collaborative Learning (KDCL, Guo et al. (2020))¹. DML trains two models with alternating updates while regularizing the consistency between their final prediction distributions. KDCL extends two models to multiple models, training all models concurrently while regularizing the consistency between the prediction distribution of each individual model and of the ensemble of all models. ONE adopts the traditional weight-sharing strategy with a learnable gating factor assigned to each individual branch, which helps to control the model diversity.

Perturbation. We demonstrate the effectiveness of CAMERO using neuron dropout (Srivastava et al., 2014), one of the most straightforward perturbation techniques which randomly zeros out neurons based on a small, fixed ratio. In particular, the ratio adopted in our experiments is 0.1. In Section 5.3, we further demonstrate that a wide variety of perturbations, including virtual adversarial perturbation (Jiang et al., 2019a), random perturbation (Aghajanyan et al., 2020) and word dropout (Wei and Zou, 2019), can all serve the role.

Consistency regularization. We demonstrate the effectiveness of CAMERO using the ensemble consistency defined in Eq. (1). In Section 5.4, we further investigate the effectiveness of different types of consistency regularization techniques.

Initialization. To fine-tune the BERT encoder on downstream tasks, the common initialization approach is to append a randomly initialized, fully connected classification layer on top of the encoder

¹We do not include the data augmentation technique proposed in KDCL for a fair comparison.

# of Models	Method	MNLI-m/mm Acc	QQP Acc/F1	QNLI Acc	CoLA Mcc	SST-2 Acc	RTE Acc	MRPC Acc/F1	STS-B P/S Corr	Avg. # Param. Score (million)
1	Single	84.5/84.6	91.1/88.1	91.2	58.7	92.9	71.1	86.2/90.4	89.7/89.2	83.2 109.5
2	Vanilla	84.9/85.2	91.6/88.7	91.8	58.2	93.2	70.6	86.2/90.4	89.8/89.5	83.4 220.1
	DML	85.0/85.5	91.6/88.7	91.9	58.2	93.3	71.3	87.1/90.9	89.9/89.5	83.6 220.1
	KDCL	85.1/85.6	91.7/88.8	92.0	59.4	93.2	71.8	87.0/90.9	89.9/89.5	83.8 220.1
	ONE	84.5/84.7	91.1/88.1	91.7	59.2	93.0	70.8	87.0/91.1	89.7/89.3	83.4 110.7
	CAMERO	85.2/85.7	91.6/88.8	92.2	59.8	93.2	72.6	87.1/90.9	89.9/89.5	84.0 110.7
4	Vanilla	85.0/85.2	91.7/88.9	91.8	58.4	93.1	70.8	87.2/91.0	90.0/89.6	83.5 440.3
	KDCL	85.0/85.7	91.7/88.8	92.0	58.6	93.3	71.3	87.4/91.1	90.1/89.6	83.7 440.3
	ONE	84.6/84.9	91.2/88.3	91.8	58.8	93.1	71.1	87.4/91.1	89.8/89.4	83.5 111.9
	CAMERO	85.4/86.1	91.8/89.1	92.3	59.5	93.5	72.8	87.2/91.0	90.1/89.7	84.2 111.9
	Vanilla	85.1/85.5	91.7/88.8	92.1	59.0	93.2	71.0	87.2/91.0	90.1/89.7	83.7 880.6
8	CAMERO	85.6/86.3	91.9/89.2	92.7	60.5	93.6	72.4	87.4/91.2	90.2/89.8	84.4 114.2

Table 1: Single-task fine-tuning dev results on ensembled BERT-base using the GLUE benchmark. "Single" denotes single model performance. All results are from our own implementation.

# of Models	Method	MNLI-m/mm Acc	QQP Acc/F1	QNLI Acc	CoLA Mcc	SST-2 Acc	RTE Acc	MRPC Acc/F1	STS-B P/S Corr	Avg. # Param. Score (million)
1	Single	90.2/90.2	92.2/-	94.7	68.0	96.4	86.6	90.9/-	-/92.4	88.9 356.4
	Vanilla	90.8/90.5	92.4/89.8	94.7	68.2	96.5	86.2	91.2/93.6	92.7/92.5	89.0 1425.6
4	CAMERO	91.1/90.9	92.5/90.0	95.3	70.3	97.0	87.7	91.7/94.0	92.8/92.6	89.8 359.6

Table 2: Single-task fine-tuning dev results on ensembled RoBERTa-large using the GLUE benchmark. "Single" denotes single model performance from Liu et al. (2019); other results are from our own implementation.

(Devlin et al., 2018). For ONE and CAMERO, we append m differently initialized, parallel classification layers on top of the encoder. For other methods, we initialize m individual encoders and append a differently initialized classification layer on top of each.

Inference. For ONE and CAMERO, we conduct a single pass through the encoder and average the predicted logits of m classification layers. For other methods, we average the predicted logits of m models. All results in the following experiments are evaluated based on such a logits ensemble.

Implementation details. Our implementation is based on the MT-DNN code-base². We follow the suggested training and hyper-parameters settings from Liu et al. (2020). Specifically, we adopt Adamax (Kingma and Ba, 2014) as the optimizer with $\beta = (0.9, 0.999)$. We tune α in range of $\{0.5, 1, 2, 5\}$ for all methods. Comprehensive training details are reported in Appendix A.1.2.

Results of BERT-base. Table 1 shows the evaluation results of BERT-base on the GLUE development set. The results are averaged over five random

seeds, and all gains are statistically significant³.

We have the following observations: 1) With significantly less learning parameters, CAMERO achieves a prominent and consistent margin over Vanilla, DML and KDCL. This suggests that CAMERO can produce better-generalized ensemble model with higher parameter efficiency. 2) CAMERO significantly outperforms ONE, suggesting that applying perturbations to models effectively improves the performance of weight-sharing strategy. 3) As the number of models increases from 2 to 8, CAMERO’s performance steadily increases for 6 out of 8 tasks, while Vanilla and KDCL fail to do so.

Results of RoBERTa-large. We further verify that CAMERO can benefit an even larger model, RoBERTa-large. As shown in Table 2, CAMERO achieves consistent gains across all tasks⁴. Worth noticing, Vanilla shows limited improvements upon the single model performance (e.g., the gains are

²<https://github.com/namisan/mt-dnn>

³All results have passed a paired student t-test with p-values less than 0.05. The detailed statistics are summarized in Appendix A.1.3.

⁴We present the median of five runs following Liu et al. (2019).

0.0, 0.1 and -0.4 on QNLI, SST-2 and RTE, respectively). We conjecture that the high model variance in large models compromises the ensemble benefits. In contrast, by control the model variance with regularization, CAMERO achieves gains of 0.6, 0.6 and 1.1 on these tasks.

4.2 Neural Machine Translation

Model and data. We further evaluate CAMERO on the Transformer-base NMT model (Vaswani et al., 2017) using widely used IWSLT (Cettolo et al., 2016)⁵ and WMT (Bojar et al., 2016)⁶ datasets. Specifically, we adopt IWSLT’14 En \leftrightarrow De, IWSLT’16 En \leftrightarrow Fr and WMT’14 En \leftrightarrow De. IWSLT En \leftrightarrow De and En \leftrightarrow Fr are low-resource datasets containing 160k and 236k sentence pairs. WMT En \leftrightarrow De is a rich-resource dataset containing 4.5M sentence pairs. Model and dataset details are deferred to Appendix A.2.1. **Implementation details.** Our implementation is based on the fairseq code-base and follows the training and hyper-parameters settings from Ott et al. (2018, 2019). Specifically, we use 5×10^{-4} as the learning rate and employ Adam (Kingma and Ba, 2014) as the optimizer with $\beta = (0.9, 0.98)$. We select α in range of $\{1, 2, 5\}$. For ONE and CAMERO, we randomly initialize multiple parallel decoders’s last layers as the un-shared branches. Comprehensive training details are reported in appendix A.2.2.

Main results. Table 3 shows the BLEU scores on the IWSLT test set and the SacreBLEU scores (Post, 2018) with compound splitting on the WMT test set⁷. WMT’s corresponding BLEU scores are reported in Appendix A.2.3.

With a number of learning parameters similar to a single model, CAMERO achieves around 2 and 1 points upon ONE, and improves around 0.4 and 0.4 points upon KDCL, on low-resource and rich-resource datasets, respectively. This suggests that other than fine-tuning, CAMERO also improves the generalization of training-from-scratch models in both low-resource and rich-resource datasets.

5 Analysis

We first verify that CAMERO leads to a well-generalized and low-variance ensemble model. We

then demonstrate how the perturbation and consistency regularization strength influences the model diversity and performance. Finally, we demonstrate CAMERO’s effectiveness on various types of perturbation and regularization techniques.

5.1 Shared Weights Learn Better Representations

We verify that CAMERO allows the shared weights to learn better-generalized representations. Specifically, we attach a randomly initialized classifier on top of a BERT-base encoder trained by CAMERO. We then fix the encoder and fine-tune the attached classifier only. As shown in Table 4, the encoder trained by CAMERO learns better representations than ONE’s consistently across different tasks and under different numbers of models.

5.2 Ensemble Model Has a Low Variance Across Random Seeds

We verify that CAMERO produces an ensemble model that both generalizes well and has a low variance across different random seeds under a light parameter budget. Figure 3 plots the prediction accuracy of 2-model and 4-model ensemble across five seeds. For example, in MNLI, CAMERO’s 2-model ensemble (110.7M) achieves similar performance to KDCL’s 4-model ensemble (440.2M) and CAMERO’s 4-model ensemble (111.9M) achieves an even better performance. Across different tasks, CAMERO’s ensemble model has a similar or lower variance than all others. Complete variance statistics are presented in Appendix A.1.3.

5.3 Types and Strength of Perturbations

Types of perturbation. We verify that CAMERO produces well-generalized ensemble models under various types of perturbations. Specifically, we apply virtual adversarial perturbation (Jiang et al., 2019a) and random noise perturbation (Aghajanyan et al., 2020) on the first layer input embeddings, neuron dropout on all layers’ input representations (Srivastava et al., 2014), and word dropout on the input sentences (Wei and Zou, 2019). Specifically, we set the dropout ratio to be 0.1 for neuron dropout and 0.05 for word dropout. We set the norm constraint $\epsilon = 1 \times 10^{-5}$ for both virtual adversarial perturbation and random noise perturbation. The random noise is sampled from a normal distribution. As shown in Table 5, CAMERO leads to significant margin of improvement under all types of perturbation. In particular, virtual adversarial

⁵<https://wit3.fbk.eu/>

⁶<http://data.statmt.org/wmt16/translation-task/>

⁷We evaluate the SacreBLEU score on the average of last 10 checkpoints. The tokenizer version is: nrefs:1 | case:mixed | eff:no | tok:13a, smooth:exp | version:2.0.0.

# of Models	Method	IWSLT						WMT			
		En-De	De-En	En-Fr	Fr-En	Avg.	# Param.	En-De	De-En	Avg.	# Param.
1	Single	28.5	34.7	38.1	37.7	34.7	54.5	26.9	30.7	28.8	77.6
2	Vanilla	28.6	34.8	38.2	37.8	34.9	109.1	27.0	31.2	29.1	155.3
	DML	30.5	37.4	39.9	39.6	36.9	109.1	27.1	31.8	29.5	155.3
	KDCL	30.6	37.2	39.8	39.5	36.7	109.1	27.2	31.9	29.6	155.3
	ONE	28.9	35.1	38.5	38.2	35.2	58.7	27.0	31.0	29.0	81.8
	CAMERO	30.8	37.5	40.2	39.8	37.1	58.7	27.6	32.2	29.9	81.8
4	Vanilla	28.7	34.9	38.2	37.8	34.9	218.1	27.0	31.2	29.1	310.6
	KDCL	30.8	37.4	39.9	39.7	36.9	218.1	27.1	32.0	29.6	310.6
	ONE	28.8	35.0	38.2	37.9	35.0	67.1	27.1	31.1	29.1	90.2
	CAMERO	31.1	37.8	40.3	39.9	37.3	67.1	27.7	32.4	30.1	90.2

Table 3: Test set scores on ensembled Transformer-base on IWSLT tasks (BLEU) and WMT tasks (SacreBLEU). "Single" denotes single model performance. All results are from our own implementation.

# of Models	Methods	MNLI Acc	SST-2 Acc	MRPC Acc/F1	Avg. Score
1	Single	84.58	92.95	86.88	88.14
2	ONE	84.53	92.94	88.94	88.80
	CAMERO	85.41	93.03	89.04	89.16
4	ONE	84.67	93.03	89.07	88.92
	CAMERO	85.57	93.46	89.20	89.41

Table 4: Performance of the ensembled BERT-base encoder using the GLUE dev set. We only fine-tune the randomly initialized classification layer on top of a well-trained encoder.




Figure 3: Performance and variance of the ensembled BERT-base on the GLUE dev set.

perturbation and neuron dropout perform consistently well on all tasks. Random noise perturbation performs well on larger tasks (e.g., MNLI, QNLI, SST-2) while the gains shrink on smaller tasks.

Strength of perturbation. We then verify that a larger perturbation strength improves the perturbed models’ diversity during training. As consistency loss is computed as the average distance between all perturbed models’ output logits to the ensembled logits at each iteration, it directly reflects the model diversity during training. As shown in Figure 4 (Left), a larger neuron dropout ratio leads to larger consistency loss, thus higher model diversity.

Furthermore, we observe that a larger perturbation strength leads to a lower-variance ensemble model. As shown in Figure 4 (Right), as the neuron dropout ratio grows, CAMERO’s ensemble model variance decreases. In contrast, ONE has a large variance under all ratios.

5.4 Types and Strength of Consistency Regularization

Types of consistency regularization. We then investigate the effects of using different types of consistency regularization techniques. Specifically, we compare the existing *ensemble consistency*, as defined in Eq. (1), and a newly proposed *pairwise consistency*, which is defined as

$$\mathcal{R}(\Theta) = \frac{2}{m(m-1)} \sum_{j=1}^m \sum_{p=j+1}^m \mathcal{D}(f(\mathbf{x}; \theta_j), f(\mathbf{x}; \theta_p)).$$

The pairwise consistency measures the average distance between *each pair* of models’ output logits, thus we expect it to capture the discrepancies among models more accurately. As shown

Perturbation Types	MNLI	QNLI	SST-2	MRPC	CoLA	Avg.
	Acc	Acc	Acc	Acc/F1	MCC	Score
None	84.74	91.76	93.10	89.05	58.75	83.48
Neuron Dropout (Srivastava et al., 2014)	85.73	92.30	93.46	89.09	59.50	84.02
Virtual Adversarial Pert. (Jiang et al., 2019a)	85.76	92.33	93.53	89.19	59.49	84.08
Random Noise Pert. (Aghajanyan et al., 2020)	85.78	92.21	93.42	89.07	59.22	83.94
Word Dropout (Wei and Zou, 2019)	85.61	92.00	93.21	89.06	59.19	83.81

Table 5: CAMERO’s performance under different types of perturbation. "None" corresponds to ONE, which does not apply different perturbations to different models. We report the 4-model ensembled BERT-base results.

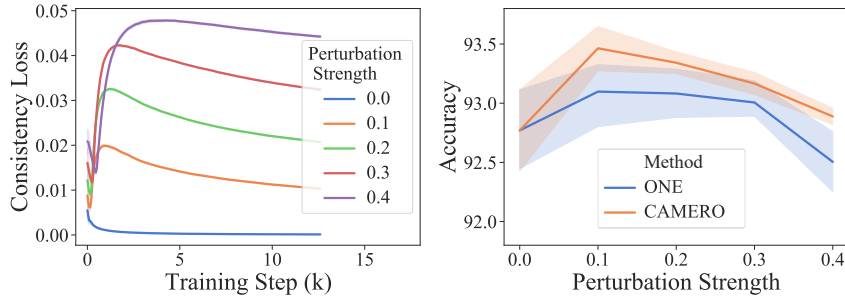


Figure 4: The effect of perturbation strength on models’ diversity during training (Left) and the variance of the ensembled model (Right). We fine-tune BERT-base on SST-2 and report the 4-model ensemble results.

Consistency Types	MNLI	QNLI	SST-2	MRPC	CoLA	Avg.
	Acc	Acc	Acc	Acc/F1	MCC	Score
None	85.23	91.76	93.30	88.97	58.40	83.48
Ensemble Consistency	85.73	92.30	93.46	89.09	59.50	84.02
Pairwise Consistency	85.73	92.33	93.37	89.40	59.87	84.14

Table 6: CAMERO’s performance under different types of consistency regularization. "None" corresponds to $\alpha = 0$, where no regularization is applied. We report the 4-model ensembled BERT-base results.




Figure 5: The effect of consistency regularization strength on the generalization and variance of the ensembled model. We fine-tune BERT-base and report the 4-model ensemble results.

in Table 6, CAMERO shows consistent improvements under both types of regularization. In particular, pairwise consistency shows larger advantages on smaller tasks (e.g., 0.3 on MRPC and 0.4 on CoLA).

Strength of consistency regularization. We further investigate how the strength of the regularization factor α affects the ensemble model’s performance. As shown in Figure 5, as α increases, the generalization performance of the ensemble model first increases, then decreases. This suggests that regularization can effectively benefits the general-

ization performance through balancing the model diversity.

6 Conclusion

We propose CAMERO, a consistency-regularized ensemble learning approach based on perturbed models. Such a strategy significantly improves the parameter efficiency of model ensemble in large language models, making it an accessible and powerful technique for learning ensemble models with better generalization performances.

References

- Armen Aghajanyan, Akshat Shrivastava, Anchit Gupta, Naman Goyal, Luke Zettlemoyer, and Sonal Gupta. 2020. Better fine-tuning by reducing representational collapse. *arXiv preprint arXiv:2008.03156*.
- Roei Aharoni, Melvin Johnson, and Orhan Firat. 2019. Massively multilingual neural machine translation. *arXiv preprint arXiv:1903.00089*.
- Roy Bar-Haim, Ido Dagan, Bill Dolan, Lisa Ferro, and Danilo Giampiccolo. 2006. The second PASCAL recognising textual entailment challenge. In *Proceedings of the Second PASCAL Challenges Workshop on Recognising Textual Entailment*.
- Luisa Bentivogli, Ido Dagan, Hoa Trang Dang, Danilo Giampiccolo, and Bernardo Magnini. 2009. The fifth pascal recognizing textual entailment challenge. In *In Proc Text Analysis Conference (TAC'09)*.
- Ondřej Bojar, Rajen Chatterjee, Christian Federmann, Yvette Graham, Barry Haddow, Matthias Huck, Antonio Jimeno Yepes, Philipp Koehn, Varvara Logacheva, Christof Monz, et al. 2016. Findings of the 2016 conference on machine translation. In *Proceedings of the First Conference on Machine Translation: Volume 2, Shared Task Papers*, pages 131–198.
- Denny Britz, Quoc Le, and Reid Pryzant. 2017. Effective domain mixing for neural machine translation. In *Proceedings of the Second Conference on Machine Translation*, pages 118–126.
- Tom B Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *arXiv preprint arXiv:2005.14165*.
- Daniel Cer, Mona Diab, Eneko Agirre, Iñigo Lopez-Gazpio, and Lucia Specia. 2017. Semeval-2017 task 1: Semantic textual similarity multilingual and crosslingual focused evaluation. In *Proceedings of the 11th International Workshop on Semantic Evaluation (SemEval-2017)*, pages 1–14.
- Mauro Cettolo, Niehues Jan, Stüker Sebastian, Luisa Bentivogli, Roldano Cattoni, and Marcello Federico. 2016. The iwslt 2016 evaluation campaign. In *International Workshop on Spoken Language Translation*.
- Defang Chen, Jian-Ping Mei, Can Wang, Yan Feng, and Chun Chen. 2020. Online knowledge distillation with diverse peers. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 3430–3437.
- Ido Dagan, Oren Glickman, and Bernardo Magnini. 2006. [The pascal recognising textual entailment challenge](#). In *Proceedings of the First International Conference on Machine Learning Challenges: Evaluating Predictive Uncertainty Visual Object Classification, and Recognizing Textual Entailment, MLCW'05*, pages 177–190, Berlin, Heidelberg. Springer-Verlag.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- William B Dolan and Chris Brockett. 2005. Automatically constructing a corpus of sentential paraphrases. In *Proceedings of the Third International Workshop on Paraphrasing (IWP2005)*.
- Xibin Dong, Zhiwen Yu, Wenming Cao, Yifan Shi, and Qianli Ma. 2020. A survey on ensemble learning. *Frontiers of Computer Science*, 14(2):241–258.
- Shaoxiong Feng, Hongshen Chen, Xuancheng Ren, Zhuoye Ding, Kan Li, and Xu Sun. 2020. Collaborative group learning. *arXiv preprint arXiv:2009.07712*.
- Danilo Giampiccolo, Bernardo Magnini, Ido Dagan, and Bill Dolan. 2007. [The third PASCAL recognizing textual entailment challenge](#). In *Proceedings of the ACL-PASCAL Workshop on Textual Entailment and Paraphrasing*, pages 1–9, Prague. Association for Computational Linguistics.
- Jiatao Gu, Hany Hassan, Jacob Devlin, and Victor OK Li. 2018. Universal neural machine translation for extremely low resource languages. *arXiv preprint arXiv:1802.05368*.
- Qiushan Guo, Xinjiang Wang, Yichao Wu, Zhipeng Yu, Ding Liang, Xiaolin Hu, and Ping Luo. 2020. Online knowledge distillation via collaborative learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11020–11029.
- Pengcheng He, Xiaodong Liu, Jianfeng Gao, and Weizhu Chen. 2020. Deberta: Decoding-enhanced bert with disentangled attention. *arXiv preprint arXiv:2006.03654*.
- Haoming Jiang, Pengcheng He, Weizhu Chen, Xiaodong Liu, Jianfeng Gao, and Tuo Zhao. 2019a. Smart: Robust and efficient fine-tuning for pre-trained natural language models through principled regularized optimization. *arXiv preprint arXiv:1911.03437*.
- Haoming Jiang, Chen Liang, Chong Wang, and Tuo Zhao. 2019b. Multi-domain neural machine translation with word-level adaptive layer-wise domain mixing. *arXiv preprint arXiv:1911.02692*.
- Jangho Kim, Minsung Hyun, Inseop Chung, and Nojun Kwak. 2021. Feature fusion for online mutual knowledge distillation. In *2020 25th International Conference on Pattern Recognition (ICPR)*, pages 4619–4625. IEEE.

- Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Xu Lan, Xiatian Zhu, and Shaogang Gong. 2018. Knowledge distillation by on-the-fly native ensemble. *arXiv preprint arXiv:1806.04606*.
- Zheng Li, Ying Huang, Defang Chen, Tianren Luo, Ning Cai, and Zhigeng Pan. 2020. Online knowledge distillation via multi-branch diversity enhancement. In *Proceedings of the Asian Conference on Computer Vision*.
- Xiaodong Liu, Yu Wang, Jianshu Ji, Hao Cheng, Xueyun Zhu, Emmanuel Awa, Pengcheng He, Weizhu Chen, Hoifung Poon, Guihong Cao, et al. 2020. The microsoft toolkit of multi-task deep neural networks for natural language understanding. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 118–126.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Minh-Thang Luong, Quoc V Le, Ilya Sutskever, Oriol Vinyals, and Lukasz Kaiser. 2015. Multi-task sequence to sequence learning. *arXiv preprint arXiv:1511.06114*.
- Takeru Miyato, Shin-ichi Maeda, Masanori Koyama, and Shin Ishii. 2018. Virtual adversarial training: a regularization method for supervised and semi-supervised learning. *IEEE transactions on pattern analysis and machine intelligence*, 41(8):1979–1993.
- Myle Ott, Sergey Edunov, Alexei Baevski, Angela Fan, Sam Gross, Nathan Ng, David Grangier, and Michael Auli. 2019. fairseq: A fast, extensible toolkit for sequence modeling. *arXiv preprint arXiv:1904.01038*.
- Myle Ott, Sergey Edunov, David Grangier, and Michael Auli. 2018. Scaling neural machine translation. *arXiv preprint arXiv:1806.00187*.
- Matt Post. 2018. A call for clarity in reporting bleu scores. *arXiv preprint arXiv:1804.08771*.
- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. SQuAD: 100,000+ questions for machine comprehension of text. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2383–2392, Austin, Texas. Association for Computational Linguistics.
- Alexandre Rame and Matthieu Cord. 2021. Dice: Diversity in deep ensembles via conditional redundancy adversarial estimation. *arXiv preprint arXiv:2101.05544*.
- Sebastian Ruder, Joachim Bingel, Isabelle Augenstein, and Anders Søgaard. 2019. Latent multi-task architecture learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 4822–4829.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2015. Neural machine translation of rare words with subword units. *arXiv preprint arXiv:1508.07909*.
- Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D Manning, Andrew Ng, and Christopher Potts. 2013. Recursive deep models for semantic compositionality over a sentiment treebank. In *Proceedings of the 2013 conference on empirical methods in natural language processing*, pages 1631–1642.
- Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. Dropout: a simple way to prevent neural networks from overfitting. *The journal of machine learning research*, 15(1):1929–1958.
- Sander Tars and Mark Fishel. 2018. Multi-domain neural machine translation. *arXiv preprint arXiv:1805.02282*.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008.
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R Bowman. 2018. Glue: A multi-task benchmark and analysis platform for natural language understanding. *arXiv preprint arXiv:1804.07461*.
- Alex Warstadt, Amanpreet Singh, and Samuel R Bowman. 2019. Neural network acceptability judgments. *Transactions of the Association for Computational Linguistics*, 7:625–641.
- Jason Wei and Kai Zou. 2019. Eda: Easy data augmentation techniques for boosting performance on text classification tasks. *arXiv preprint arXiv:1901.11196*.
- Adina Williams, Nikita Nangia, and Samuel Bowman. 2018. A broad-coverage challenge corpus for sentence understanding through inference. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1112–1122. Association for Computational Linguistics.
- Guile Wu and Shaogang Gong. 2021. Peer collaborative learning for online knowledge distillation. In *AAAI*.

- Chuangang Yang, Zhulin An, and Yongjun Xu. 2021. Multi-view contrastive learning for online knowledge distillation. In *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 3750–3754. IEEE.
- Yongquan Yang and Haijun Lv. 2021. Discussion of ensemble learning under the era of deep learning. *arXiv preprint arXiv:2101.08387*.
- Jiali Zeng, Jinsong Su, Huating Wen, Yang Liu, Jun Xie, Yongjing Yin, and Jianqiang Zhao. 2018. Multi-domain neural machine translation with word-level domain context discrimination. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 447–457.
- Ying Zhang, Tao Xiang, Timothy M Hospedales, and Huchuan Lu. 2018. Deep mutual learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4320–4328.

A Appendix

A.1 Natural Language Understanding

A.1.1 Data

GLUE is a collection of nine NLU tasks. The benchmark includes question answering (Rajpurkar et al., 2016), linguistic acceptability (CoLA, Warstadt et al. 2019), sentiment analysis (SST, Socher et al. 2013), text similarity (STS-B, Cer et al. 2017), paraphrase detection (MRPC, Dolan and Brockett 2005), and natural language inference (RTE & MNLI, Dagan et al. 2006; Bar-Haim et al. 2006; Giampiccolo et al. 2007; Bentivogli et al. 2009; Williams et al. 2018) tasks. Details of the GLUE benchmark, including tasks, statistics, and evaluation metrics, are summarized in Table 12.

All the texts were tokenized using wordpieces, and were chopped to spans no longer than 512 tokens.

A.1.2 Training Details

Table 7 presents the hyper-parameter configurations to fine-tune BERT-base and RoBERTa-large models. We apply a linear weight decay rate of 0.01 and a gradient norm clipping threshold of 1 for all experiments. All experiments are conducted on Nvidia V100 GPUs.

A.1.3 Evaluation Results

Statistics of the dev set results. Table 8 shows the standard deviation of the dev set results.

Average score computation formula. For dev set results, we first obtain a score for each task by averaging the scores of all metrics (e.g., Acc and F1) and test sets (e.g., MNLI-m and MNLI-mm) within this task, then compute a task-average score. For test set results, we directly average scores of all reported metrics following Devlin et al. (2018).

A.2 Neural Machine Translation

A.2.1 Data

For IWSLT’14 En-De and De-En datasets, we follow Ott et al. (2019)⁸ to split the train/dev/test set. For IWSLT’16 En-Fr and Fr-En, we adopt the default training set, and use IWSLT16.TED.tst2015 for validation and use IWSLT16.TED.tst2016 for testing. For WMT’14 En-De and De-En, We use the standard newstest-2013 and newstest-2014 for validation and testing, respectively. Table 9 shows the number of sentence pairs in each dataset.

⁸<https://github.com/pytorch/fairseq/blob/main/examples/translation/prepare-iwslt14.sh>

We tokenize all datasets with byte-pair encoding (BPE, Sennrich et al. (2015)) with a vocabulary size of 10k for datasets in IWSLT and 32k for datasets in WMT. We build a joint dictionary upon all source and target sentences for all datasets.

A.2.2 Training Details

We adopt the Transformer-base model for all datasets and share all embeddings. For IWSLT datasets, we follow the training configurations from Ott et al. (2019)⁹. For WMT datasets, we follow the training configurations from Ott et al. (2018)¹⁰. For all datasets, we use Adam(Kingma and Ba, 2014) as the optimizer with $\beta = (0.9, 0.98)$. We use a inverse square root learning rate schedule. We apply a linear weight decay rate of 1×10^{-4} and a label smoothing ratio of 0.1 for all experiments. All experiments are conducted on Nvidia V100 GPUs. Table 10 presents the training hyper-parameter configurations for all datasets.

For evaluation on IWSLT datasets, we report the BLEU score of the best checkpoint using a beam size of 5 and length penalty of 1. For evaluation on WMT datasets, we average the last 10 checkpoints, decode with a beam size of 4 and length penalty of 0.6, then report the SacreBLEU scores after compound splitting.

A.2.3 BLEU scores for WMT experiments

Table 11 shows the corresponding BLEU scores for WMT datasets.

⁹<https://github.com/pytorch/fairseq/tree/main/examples/translation/iwslt14-german-to-english-transformer>

¹⁰https://github.com/pytorch/fairseq/tree/main/examples/scaling_nmt#training-a-new-model-on-wmt16-en-de

Hyper-param	Model	RTE	MRPC	CoLA	SST-2	STS-B	QNLI	QQP	MNLI
Learning Rate	BERT _{BASE}	1e-4	1e-4	1e-4	8e-5	1e-4	1e-4	1e-4	8e-5
	RoBERTa _{LARGE}	5e-5	1e-4	3e-5	2e-5	5e-5	1e-5	1e-4	3e-5
Epoch	BERT _{BASE}	6	6	6	6	6	3	6	3
	RoBERTa _{LARGE}	15	6	6	10	10	10	10	3
Batch Size	BERT _{BASE}	16	8	32	32	32	32	32	32
	RoBERTa _{LARGE}	8	16	32	32	32	32	32	32
Dropout	Both	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.3
Warmup	BERT _{BASE}	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1

Table 7: Hyper-parameter configurations for GLUE experiments. “Epoch” refers to the total training epochs; we adopt early-stopping strategy in practice. “Dropout” refers to classification layer dropout ratio, the encoder dropout ratio is fixed to be 0.1. “Warmup” refers to the ratio of learning rate linear warmup iterations to total training iterations.

# of Models	Method	MNLI-m/mm	QQP	QNLI	CoLA	SST-2	RTE	MRPC	STS-B
		Acc	Acc/F1	Acc	Mcc	Acc	Acc	Acc/F1	P/S Corr
1	Single	0.20	0.25	0.21	1.10	0.33	1.61	0.83	0.20
2	Vanilla	0.17	0.11	0.16	1.22	0.26	1.72	0.77	0.21
	DML	0.14	0.03	0.15	0.76	0.23	1.07	0.69	0.19
	KDCL	0.11	0.05	0.16	1.04	0.14	0.68	0.66	0.19
	ONE	0.13	0.22	0.21	1.00	0.25	1.61	0.88	0.15
	CAMERO	0.11	0.05	0.15	1.01	0.12	0.92	0.37	0.11
4	Vanilla	0.14	0.10	0.09	0.92	0.39	0.82	0.54	0.11
	KDCL	0.20	0.12	0.12	1.28	0.31	0.66	0.27	0.06
	ONE	0.19	0.22	0.25	0.98	0.34	1.44	0.84	0.12
	CAMERO	0.11	0.07	0.05	1.03	0.25	0.84	0.19	0.06

Table 8: Standard deviation of the single-task fine-tuning dev results on ensembled BERT-base.

Data	Train	Dev	Test
IWSLT’14 En-De/De-En	160	7283	6750
IWSLT’16 En-Fr/Fr-En	218	1080	1133
WMT’14 En-De/De-En	4.5m	1061	1019

Table 9: The number of parallel sentences in NMT datasets.

Hyper-param	IWSLT	WMT
Learning Rate	5×10^{-4}	1×10^{-3}
Batch size	4096/GPU \times 1 GPU	3584/GPU \times 8 GPUs \times 16 grad. acc. steps
Epoch	250	150
Dropout	0.3	0.1
Warmup	8000	4000

Table 10: Hyper-parameter configurations for NMT experiments. “Warmup” refers to the learning rate linear warmup iterations.

# of Models	Method	WMT	
		En-De	De-En
1	Single (Vaswani et al., 2017)	27.30	-
	Single	27.54	31.28
2	Vanilla	27.62	31.76
	DML	27.70	32.22
	KDCL	27.84	32.35
	ONE	27.68	31.43
	CAMERO	28.26	32.61
4	Vanilla	27.67	31.79
	KDCL	27.75	32.47
	ONE	27.78	31.48
	CAMERO	28.43	32.78

Table 11: Test set scores on ensembled Transformer-base on WMT tasks (BLEU). The result in "Single (Vaswani et al., 2017)" is the single model performance reported from Vaswani et al. (2017); Other results are from our own implementation.

Corpus	Task	#Train	#Dev	#Test	#Label	Metrics
Single-Sentence Classification (GLUE)						
CoLA	Acceptability	8.5k	1k	1k	2	Matthews corr
SST	Sentiment	67k	872	1.8k	2	Accuracy
Pairwise Text Classification (GLUE)						
MNLI	NLI	393k	20k	20k	3	Accuracy
RTE	NLI	2.5k	276	3k	2	Accuracy
QQP	Paraphrase	364k	40k	391k	2	Accuracy/F1
MRPC	Paraphrase	3.7k	408	1.7k	2	Accuracy/F1
QNLI	QA/NLI	108k	5.7k	5.7k	2	Accuracy
Text Similarity (GLUE)						
STS-B	Similarity	7k	1.5k	1.4k	1	Pearson/Spearman corr

Table 12: Summary of the GLUE benchmark.