

# Repulsive Attention: Rethinking Multi-head Attention as Bayesian Inference

Bang An<sup>†</sup>, Jie Lyu<sup>†</sup>, Zhenyi Wang<sup>†</sup>, Chunyuan Li<sup>§</sup>  
Changwei Hu<sup>\*</sup>, Fei Tan<sup>\*</sup>, Ruiyi Zhang<sup>‡</sup>, Yifan Hu<sup>\*</sup>, Changyou Chen<sup>†</sup>  
<sup>†</sup>State University of New York at Buffalo, <sup>§</sup>Microsoft Research, Redmond  
<sup>\*</sup>Yahoo Research, <sup>‡</sup>Duke University  
isbangan@gmail.com, ch237duke@gmail.com  
yifanhu@verizonmedia.com, changyou@buffalo.edu

## Abstract

The neural attention mechanism plays an important role in many natural language processing applications. In particular, multi-head attention extends single-head attention by allowing a model to jointly attend information from different perspectives. However, without explicit constraining, multi-head attention may suffer from *attention collapse*, an issue that makes different heads extract similar attentive features, thus limiting the model’s representation power. In this paper, for the first time, we provide a novel understanding of multi-head attention from a Bayesian perspective. Based on the recently developed particle-optimization sampling techniques, we propose a non-parametric approach that explicitly improves the repulsiveness in multi-head attention and consequently strengthens model’s expressiveness. Remarkably, our Bayesian interpretation provides theoretical inspirations on the not-well-understood questions: why and how one uses multi-head attention. Extensive experiments on various attention models and applications demonstrate that the proposed repulsive attention can improve the learned feature diversity, leading to more informative representations with consistent performance improvement on multiple tasks.

## 1 Introduction

Multi-head attention (Vaswani et al., 2017) is an effective module in deep neural networks, with impressive performance gains in many natural-language-processing (NLP) tasks. By extending a single head to multiple paralleled attention heads, the architecture is widely adopted to capture different attentive information and strengthen the expressive power of a model. Lin et al. (2017) applied the idea of multi-heads on self-attention and extract a 2-D matrix instead of a vector to represent different contexts of a sentence. The Transformer (Vaswani et al., 2017) and its variants such

as BERT (Devlin et al., 2019) are influential architectures solely based on multi-head attention, achieving state-of-the-art performance on plenty of NLP tasks. The key of multi-head attention is its ability to jointly attend to information from different representation subspaces at different positions, which results in multiple latent features depicting the input data from different perspectives. However, there are no explicit mechanisms guaranteeing this desired property, leading to potential attention redundancy or attention collapse, which has been observed in previous research (Voita et al., 2019; Kovaleva et al., 2019). Although there exist works by directly adding regularization on loss functions to encourage diversity in multi-head attention (Lin et al., 2017; Li et al., 2018), the underlying working principle has not been well-validated, and performance improvement is limited. Furthermore, an important problem on why and how multi-head attention improves over its single-head counterpart is poorly understood.

In this paper, we provide a novel understanding of multi-head attention from a Bayesian perspective by adapting the deterministic attention to a stochastic setting. The standard multi-head attention can be understood as a special case of our framework, where attention-parameter updates between heads are independent, instead of sharing a common prior distribution. Based on our framework, attention repulsiveness could then be imposed by performing Bayesian inference on attention parameters with the recently developed particle-optimization sampling methods (Liu and Wang, 2016), which has been shown to be effective in avoiding mode collapse. These methods treat each head as a particle/sample, which is then optimized to approximate a posterior distribution of an attention model. With it, multiple heads are enforced to move to modes in the parameter space to be far from each other, thus improving the repulsiveness in multi-head at-

tion and enhancing model’s expressiveness. Our Bayesian interpretation also provides a theoretical understanding of the reason and benefits of applying multi-head attention. Experiments on various attention models demonstrate the effectiveness of our framework.

Our contributions are summarized as follow:

- We provide a new understanding of multi-head attention from a Bayesian perspective, yielding a more principled and flexible interpretation of multi-head attention.
- Based on the recently developed particle-optimization sampling techniques, we propose an algorithm to explicitly encourage repulsiveness in multi-head attention without introducing extra trainable parameters or explicit regularizers. The proposed method can be implemented with an efficient end-to-end training scheme.
- Our Bayesian interpretation provides a theoretical foundation to understand the benefits of multi-head attention, which reveals the existence of an optimal number of attention heads in a specific model.
- We apply our approach on four attention models with a wide range of tasks. Experimental results show that repulsive attention improves the expressiveness of models, and yields consistent performance gains on all the tasks considered.

## 2 Preliminaries

### 2.1 Multi-head Attention

The attention mechanism aims at modeling dependencies among elements of a learned representation at different positions. The two commonly used attention functions are additive attention (Lin et al., 2017; Bahdanau et al., 2015) and dot-product attention (Vaswani et al., 2017). We review the popularly used dot-product attention below and defer the additive attention to Appendix A.

**Dot-product Attention** The multi-head scaled dot-product attention is used in the Transformer model (Vaswani et al., 2017). The attention function for a single head is formulated as mapping a query and a set of key-value pairs to output as

$$\mathbf{A}_i = \text{Softmax}(\mathbf{Q}_i \mathbf{K}_i^T / \sqrt{d_k}), \mathbf{Z}_i = \mathbf{A}_i \mathbf{V}_i \quad (1)$$

where  $\mathbf{Q}_i = \mathbf{Q} \mathbf{W}_i^Q$ ,  $\mathbf{K}_i = \mathbf{K} \mathbf{W}_i^K$ ,  $\mathbf{V}_i = \mathbf{V} \mathbf{W}_i^V$

$\mathbf{Q}, \mathbf{K}, \mathbf{V}$  are matrices depicting the hidden representation of every word in one sentence (*i.e.* self-attention) or two sentences (*i.e.* inter-attention);  $d_k$  is the dimension of key and query;  $\mathbf{Z}_i$  is the attention feature of the input sentence from the  $i$ -th head;  $\{\mathbf{W}_i^Q, \mathbf{W}_i^K, \mathbf{W}_i^V\}$  are the corresponding learnable attention parameters. The  $M$ -head attention projects the queries, keys and values into  $M$  subspaces with different learnable linear projections. These attention functions are performed in parallel and are concatenated at last, resulting in a final latent representation:

$$\text{Multi-head}(\mathbf{Q}, \mathbf{K}, \mathbf{V}) = \mathbf{Z} \mathbf{W}^O, \text{ with} \quad (2)$$

$$\mathbf{Z} = \text{Concat}(\mathbf{Z}_1, \dots, \mathbf{Z}_M)$$

### 2.2 Particle-optimization Sampling

Particle-optimization sampling is a recently developed Bayesian sampling technique that interactively transports a set of particles/samples to a target distribution  $p$  by minimizing the KL divergence between the particle density and the target  $p$ . In our case,  $p$  would be a posterior distribution,  $p(\boldsymbol{\theta}|\mathcal{D}) \propto \exp(-U(\boldsymbol{\theta}))$ , of the parameter  $\boldsymbol{\theta} \in \mathbb{R}^d$ , defined over an observed dataset  $\mathcal{D} = \{D_k\}_{k=1}^N$ . Here  $U(\boldsymbol{\theta}) \triangleq -\log p(\mathcal{D}|\boldsymbol{\theta}) - \log p_0(\boldsymbol{\theta}) = -\sum_{k=1}^N \log p(D_k|\boldsymbol{\theta}) - \log p_0(\boldsymbol{\theta})$  is called the potential energy with  $p_0$  a prior over  $\boldsymbol{\theta}$ . In our case, the model parameter  $\boldsymbol{\theta}$  could be one or several of the attention parameters such as  $\mathbf{W}_i^Q$ . For simplicity, we will stick to  $\boldsymbol{\theta}$  in the presentation. In particle-optimization sampling, a total of  $M$  particles  $\{\boldsymbol{\theta}^{(i)}\}_{i=1}^M$  are updated iteratively to approximate  $p(\boldsymbol{\theta}|\mathcal{D})$ . In this paper, we use two representative algorithms, the Stein Variational Gradient Descent (SVGD) and the Stochastic Particle-Optimization Sampling (SPOS), for sampling.

**SVGD** In SVGD (Liu and Wang, 2016), the  $i$ -th particle in the  $(\ell + 1)$ -th iteration is updated with stepsize  $\epsilon_{\ell+1}$  as

$$\boldsymbol{\theta}_{\ell+1}^{(i)} = \boldsymbol{\theta}_{\ell}^{(i)} + \epsilon_{\ell+1} \phi(\boldsymbol{\theta}_{\ell}^{(i)}) \quad (3)$$

$$\phi(\boldsymbol{\theta}_{\ell}^{(i)}) = \frac{1}{M} \sum_{j=1}^M [-\kappa(\boldsymbol{\theta}_{\ell}^{(j)}, \boldsymbol{\theta}_{\ell}^{(i)}) \nabla_{\boldsymbol{\theta}_{\ell}^{(j)}} U(\boldsymbol{\theta}_{\ell}^{(j)}) + \nabla_{\boldsymbol{\theta}_{\ell}^{(j)}} \kappa(\boldsymbol{\theta}_{\ell}^{(j)}, \boldsymbol{\theta}_{\ell}^{(i)})]$$
(4)

where  $\kappa(\cdot, \cdot)$  is a positive definite kernel (*e.g.*, RBF kernel). The two terms in  $\phi$  play different roles: the first term drives the particles towards high density regions of  $p(\boldsymbol{\theta}|\mathcal{D})$ ; whereas the second term acts as

a repulsive force that prevents all the particles from collapsing together into local modes of  $p(\boldsymbol{\theta}|\mathcal{D})$ .

**SPOS** Though obtaining significant empirical success, under certain conditions, SVGD experiences a theoretical pitfall, where particles tend to collapse. To overcome this, [Zhang et al. \(2020\)](#) generalize SVGD to a stochastic setting by injecting random noise into particle updates. The update rule for particles  $\boldsymbol{\theta}_\ell^{(i)}$  is

$$\phi(\boldsymbol{\theta}_\ell^{(i)}) = \frac{1}{M} \sum_{j=1}^M [-\kappa(\boldsymbol{\theta}_\ell^{(j)}, \boldsymbol{\theta}_\ell^{(i)}) \nabla_{\boldsymbol{\theta}_\ell^{(j)}} U(\boldsymbol{\theta}_\ell^{(j)}) + \nabla_{\boldsymbol{\theta}_\ell^{(j)}} \kappa(\boldsymbol{\theta}_\ell^{(j)}, \boldsymbol{\theta}_\ell^{(i)})] - \beta^{-1} \nabla_{\boldsymbol{\theta}_\ell^{(i)}} U(\boldsymbol{\theta}_\ell^{(i)}) + \sqrt{2\beta^{-1}\epsilon_\ell^{-1}} \boldsymbol{\xi}_\ell^{(i)} \quad (5)$$

where  $\beta > 0$  is a hyperparameter,  $\boldsymbol{\xi}_\ell^{(i)} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$  is the injected random Gaussian noise to enhance the ability of escaping local modes, leading to better ergodic properties compared to standard SVGD.

### 3 A Bayesian Inference Perspective of Multi-head Attention

In this section, we interpret multi-head attention as Bayesian inference of latent representation via particle-optimization sampling. We denote  $\mathbf{x}$  and  $\mathbf{z}$  as the input and output (latent representation) of the attention model, respectively. The single-head attention can be written as a deterministic mapping  $\mathbf{z} = f_{att}(\mathbf{x}; \boldsymbol{\theta})$ , with  $\boldsymbol{\theta}$  the parameter of the mapping. Standard multi-head attention defines multiple parallel attention mappings, each endowed with independent parameters. The attention features are finally aggregated via a function  $g(\cdot)$  as

$$\mathbf{z} = g(\mathbf{z}_1, \dots, \mathbf{z}_M), \quad \mathbf{z}_i = f_{att}(\mathbf{x}; \boldsymbol{\theta}_i). \quad (6)$$

Next, we generalize (6) as a Bayesian inference problem for the latent representation  $\mathbf{z}$ .

**Attention as Bayesian Inference** We first generalize the deterministic transformation,  $\mathbf{z} = f_{att}(\mathbf{x}; \boldsymbol{\theta})$ , to a stochastic generative process as:

$$\boldsymbol{\theta} \sim p(\boldsymbol{\theta}|\mathcal{D}), \quad \mathbf{z} = f_{att}(\mathbf{x}; \boldsymbol{\theta}),$$

where a sample of the posterior of the global attention parameter  $\boldsymbol{\theta}$ ,  $p(\boldsymbol{\theta}|\mathcal{D}) \propto p(\mathcal{D}|\boldsymbol{\theta})p(\boldsymbol{\theta})$ , is used as the parameter when generating the latent attention feature  $\mathbf{z}$ . Bayesian inference for attention then computes the predictive distribution  $p(\mathbf{z}|\mathbf{x}, \mathcal{D})$  of the attentive latent representation  $\mathbf{z}$  for a new input  $\mathbf{x}$  given the training data  $\mathcal{D}$  by  $p(\mathbf{z}|\mathbf{x}, \mathcal{D}) = \int \delta_{f_{att}(\mathbf{x}; \boldsymbol{\theta})}(\mathbf{z})p(\boldsymbol{\theta}|\mathcal{D})d\boldsymbol{\theta}$ , where

$\delta_{\mathbf{z}}(\cdot)$  is the delta function with point mass at  $\mathbf{z}$ . To enable efficient evaluation of the integral, we adopt Bayesian sampling for approximation, *i.e.*,  $p(\mathbf{z}|\mathbf{x}, \mathcal{D})$  is approximated by a set of  $M$  samples/particles initialized from  $p(\boldsymbol{\theta}|\mathcal{D})$ , leading to the following generative process:

$$\mathbf{z} = g(\mathbf{z}_1, \dots, \mathbf{z}_M) \quad (7)$$

$$\mathbf{z}_i = f_{att}(\mathbf{x}; \boldsymbol{\theta}_i), \quad \text{with } \boldsymbol{\theta}_i \sim p(\boldsymbol{\theta}|\mathcal{D})$$

The above formulation defines a principled version of multi-head attention from a Bayesian view. One can see that (7) reduces to the standard multi-head attention if all  $\boldsymbol{\theta}_i$  are treated independently without sharing the common parameter distribution  $p(\boldsymbol{\theta}|\mathcal{D})$ . In other words, our reformulation of multi-head attention is a stochastic generative process, thus is more general. Furthermore, efficient end-to-end learning can be performed by conducting repulsive Bayesian sampling for all parameters  $\{\boldsymbol{\theta}_i\}_{i=1}^M$ , which consequently could diversify the attention features  $\{\mathbf{z}_i\}_{i=1}^M$ .

## 4 Repulsive Attention Optimization

The Bayesian multi-head attention in (7) further inspires us to develop the repulsive attention. The idea is to learn to generate repulsive samples from the posterior  $p(\boldsymbol{\theta}|\mathcal{D})$ . We propose to adopt the particle-optimization sampling methods, which could explicitly encourage repulsiveness between samples. In our algorithm, the parameter of  $p(\mathbf{z}|\mathbf{x}; \boldsymbol{\theta})$  for each head is considered as one particle. Following the particle-optimization rules,  $M$  heads  $\{\boldsymbol{\theta}_i\}_{i=1}^M$  are updated iteratively to approximate the posterior distribution of attention parameter  $p(\boldsymbol{\theta}|\mathcal{D})$ .

### 4.1 Learning Repulsive Multi-head Attention

We propose to learning repulsive attention by replacing the standard updates of attention parameters via stochastic gradient descent (SGD) with particle-optimization sampling methods while keeping the multi-head attention model unchanged. This procedure forms an efficient end-to-end training scheme similar to standard attention learning. To be specific, in standard multi-head attention, the parameter of every head is updated independently according to the respective gradient of a loss function. To achieve repulsive multi-head attention, we follow the particle-optimization sampling update rule (*e.g.* (3) and (4)) to update the parameter of

every head while keeping the update for the remaining parameters via SGD unchanged. Equations (4) and (5) can be viewed as modified gradients with explicit repulsive force and can be integrated into any optimizer, *e.g.*, Adam (Kingma and Ba, 2015). Note that  $\nabla_{\theta_\ell^{(i)}} U(\theta_\ell^{(i)})$  equals to the gradient of  $\theta_\ell^{(i)}$  in standard multi-head attention when the negative log-likelihood is used as the loss function and the prior of  $\theta^{(i)}$  is assumed to be uniform. The learning algorithm is illustrated in Algorithm 1. In practice, the update of  $M$  heads can be performed in parallel with efficient matrix operations.

---

**Algorithm 1** Repulsive Multi-head Attention

---

**Input:** Initialized  $M$ -head attention model  $\mathcal{A}$  with attention parameters  $\Theta_0 = \{\theta_i\}_{i=1}^M$  and all the other parameters  $\Omega_0$ ; Training data  $\mathcal{D} = \{D_k\}_{k=1}^N = \{(\mathbf{x}_k, y_k)\}_{k=1}^N$ ;

**Output:** Optimized attention model with learned parameters  $\hat{\Theta}$  and  $\hat{\Omega}$ ;

**Train:**

**for** iteration  $\ell$  **do**

**forward:**  $\hat{y}_k = \mathcal{A}(\mathbf{x}_k; \Theta_\ell, \Omega_\ell), \forall k$ ;

    calculate loss:  $\mathcal{L}(\{\hat{y}_k\}, \{y_k\})$ ;

**backward** and calculate gradients:

    gradient of  $\Omega_\ell$ :  $\varphi(\Omega_\ell) \leftarrow \nabla_{\Omega_\ell} \mathcal{L}$

**for** attention head  $i$  **do**

      calculate  $\phi(\theta_\ell^{(i)})$  with Eq (4) or (5);

      gradient of  $\theta_\ell^{(i)}$ :  $\varphi(\theta_\ell^{(i)}) \leftarrow \epsilon_\ell \phi(\theta_\ell^{(i)})$ ;

**end for**

    update parameters:

$\Omega_{\ell+1} \leftarrow \text{Optimizer}(\Omega_\ell, \varphi(\Omega_\ell))$

$\Theta_{\ell+1} \leftarrow \text{Optimizer}(\Theta_\ell, \varphi(\Theta_\ell))$

**end for**

---

## 4.2 In-depth Analysis

**Why Multi-head Attention?** Our Bayesian interpretation of the attention mechanism naturally provides an answer to the question of why one needs multi-head attention. By treating each head as one sample, adopting multiple heads means using more samples to approximate an underlying posterior distribution. The question comes to should one use more heads (samples). Intuitively this seems to be true because more samples typically make more accurate approximations. However, this could not be the case in practice. The reason might be two-fold: *i) Overfitting:* Learning with a limited amount of data could easily causes overfitting, thus requiring a smaller model (less at-

tention heads); *ii) Numerical error:* Our proposed method to update samples (attention-head parameters) is essentially a discrete numerical method of the corresponding continuous-time partial differential equation, *i.e.*, the samples are not exact samples from the target distribution. Thanks to the recently developed theory for particle-optimization sampling (Zhang et al., 2020), one can conclude that more heads could accumulate more numerical errors, leading to performance deterioration. More formally, when using particles to approximate a target posterior distribution, there exists a gap (approximation error) between the particle distribution and the true distribution (Zhang et al., 2020). This approximation error, when applied to our setting, approximately scales in the order of  $O(\frac{1}{\sqrt{M}} + M\epsilon_0^{1/2} + e^{-\sum_\ell \epsilon_\ell})$ . Please refer to Theorem 10 in (Zhang et al., 2020) for a formal description.

**How Many Heads are Enough?** The above error bound suggests that there is a trade-off between approximation accuracy and the number of heads  $M$ . Specifically, we have *i)* when  $M$  is small, the term  $\frac{1}{\sqrt{M}}$  in the bound would dominate, leading to decreasing errors (increasing performance) with increasing  $M$ ; *ii)* when  $M$  is large enough, the term  $M\epsilon_0^{1/2}$  dominates, suggesting that larger  $M$  could actually increase the approximation error (decreased performance). These phenomena are consistent with our experimental results. We note that an exact form of the optimal  $M$  is not available due to a number of unknown constants (omitted in the big-O notation). Therefore, one should seek other ways such as cross-validation to choose a good  $M$  in practice. Our argument also aligns with recent research, which found that more heads do not necessarily lead to better performance (Michel et al., 2019).

## 5 Experiments

We demonstrate the effectiveness of our method with representative multi-head attention models on a broad range of tasks including sentence classification, machine translation, language modeling and text generation. This section summarizes key results on different models. More detailed experiment settings and analysis are deferred to the appendix. To apply our approach, only the learning method of multi-head attention is adapted.



the food **was amazing and** worth the dip in the pocketbook . a downfall was the atmosphere felt like a cafeteria . service was passing .

the food **was amazing and** worth the dip in the pocketbook . a downfall was the atmosphere felt like a cafeteria . service was passing .

the food **was amazing and** worth the dip in the pocketbook . a downfall was the atmosphere felt like a cafeteria . service was passing .

the food **was amazing and** worth the dip in the pocketbook . a downfall was the atmosphere felt like a cafeteria . service was passing .

the food **was amazing and** worth the dip in the pocketbook . a downfall was the atmosphere felt like a cafeteria . service was passing .

(a) detailed attentions

the food was amazing and worth the dip in the pocketbook . a downfall was the atmosphere felt like a cafeteria . service was **passing** .

the food was amazing and worth the dip in the pocketbook . a **downfall** was the atmosphere felt like a cafeteria . service was passing .

the food was amazing and worth **the** dip in the pocketbook . a downfall was the atmosphere felt like a cafeteria . service was passing .

the food was **amazing and worth** the dip in the pocketbook . a downfall was the atmosphere felt like a cafeteria . service was passing .

the **food was amazing** and worth the dip in the pocketbook . a downfall was the atmosphere felt like a cafeteria . service was passing .

(b) detailed attentions

the food **was amazing and** worth the dip in the pocketbook . a downfall was the atmosphere felt like a cafeteria . service was passing .

(c) standard multi-head attention

the food **was amazing and** worth the dip in the pocketbook . a downfall **was** the atmosphere felt like a cafeteria . **service was** passing .

(d) repulsive multi-head attention

Figure 1: Attention heatmaps of a 4-star Yelp review. Results on the left is from the standard multi-head attention, and the result on the right is from our repulsive multi-head attention. (a) and (b) shows detailed attention maps taken by 5 out of 30 rows of the matrix embedding, while (c) and (d) shows the overall attention by summing up all 30 attention weight vectors.

| Models              | Acc(%)       | Dist  |
|---------------------|--------------|-------|
| Age                 |              |       |
| BiLSTM + MA         | 81.47        | 0.129 |
| BiLSTM + MA + R     | 81.30        | 0.178 |
| BiLSTM + RMA (SVGD) | 81.82        | 0.492 |
| BiLSTM + RMA (SPOS) | <b>82.55</b> | 0.461 |
| Yelp                |              |       |
| BiLSTM + MA         | 69.3         | 0.246 |
| BiLSTM + MA + R     | 70.2         | 0.536 |
| BiLSTM + RMA (SVGD) | 71.2         | 1.602 |
| BiLSTM + RMA (SPOS) | <b>71.7</b>  | 1.655 |
| SNLI                |              |       |
| BiLSTM + MA         | 83.79        | 1.293 |
| BiLSTM + MA + R     | 84.55        | 1.606 |
| BiLSTM + RMA (SVGD) | 84.58        | 1.688 |
| BiLSTM + RMA (SPOS) | <b>84.76</b> | 1.370 |

Table 1: Performance (accuracy) comparison on Age, Yelp and SNLI dataset. Dist: the average 2-norm distance between each pair of the latent representation encoded from different heads on test set. MA: standard multi-head attention. RMA: proposed repulsive multi-head attention. R: regularization approach.

## 5.1 Self-attentive Sentence Classification

**Model & Baselines** We first apply our method to the self-attentive sentence classification model (Lin et al., 2017) which combines BiLSTM with additive attention to learn the sentence embedding and then does classification on it. We compare our method with the one using the standard multi-head attention (BiLSTM + MA) and the one applying the Frobenius regularization (BiLSTM + MA + R) on it to introduce diversity as in Lin et al. (2017).

**Tasks & Datasets** Following Lin et al. (2017), three sentence classification tasks including author profiling, sentiment analysis, and textual entailment are evaluated on the Age, Yelp, and SNLI datasets respectively.

**Results** As shown in Table 1, with the proposed repulsive multi-head attention, the model achieves higher accuracy on all three tasks. Especially on the sentiment analysis task which often contains multiple aspects in one sentence. Our methods also outperform the regularization method proposed in Lin et al. (2017). With different particle-optimization rules, SPOS is able to achieve better performance due to its extra advance discussed by Zhang et al. (2020). We further evaluate the diversity of multiple heads by calculating the average distance between each pair of latent representations. Results show that our methods indeed enforce heads to be more diverse, compared with the standard multi-head attention. The less diverse of the regularization-based method also indicates the validness of our argument in Appendix C.6.

**Repulsive-attention visualization** We further visualize attention maps in the learned sentence embedding space in Figure 1. It is interesting to see attention collapse indeed happens in the standard multi-head attention, where almost all heads focus on one single factor "amazing". On the contrary, the proposed method is able to capture multiple key factors in the review that are strong indicators of the sentiment behind the sentence. For example, "downfall" and "service was passing" are key factors for this 4-star review captured by our repulsive

| Models                   | BLEU        | Time |
|--------------------------|-------------|------|
| IWSLT14 De-En            |             |      |
| Transformer-small-MA     | 34.4        | 1    |
| Transformer-small-MA + R | 34.9        | 1.29 |
| Transformer-small-RMA    | <b>35.2</b> | 1.13 |
| WMT14 En-De              |             |      |
| Transformer-base-MA      | 27.3        | 1    |
| Transformer-base-MA + R  | 28.2        | 1.35 |
| Transformer-big-MA       | <b>28.4</b> | -    |
| Transformer-base-RMA     | <b>28.4</b> | 1.18 |

Table 2: Translation Performance on IWSLT14 De-En and WMT14 En-De Datasets. MA: standard multi-head attention. RMA: proposed repulsive multi-head attention. R: regularization approach. Time: relative training time of every step versus MA.

multi-head attention, which are missed by the standard attention. The full attention heatmaps of all 30 heads and more examples are in Appendix D.


## 5.2 Transformer-based Neural Translation

**Model & Baselines** The Transformer (Vaswani et al., 2017) is a representative multi-head attention based model. We apply the proposed repulsive multi-head attention (RMA) on it and compare our method with the original one (MA) and the disagreement regularization method (R) (Li et al., 2018) which encourages the diversity in attention by a cosine similarity penalty on attention outputs.

**Tasks & Datasets** Following Vaswani et al. (2017), we apply Transformer for machine translation, with two standard translation datasets: the IWSLT14 German-to-English (De-En) dataset, and the WMT14 English-to-German (En-De) dataset.

**Results** Results are presented in Table 2. With the repulsive multi-head attention, Transformer models achieve noticeable improvement on the BLEU score on both datasets, compared with both baselines. It is also encouraging to see that the Transformer-base-RMA with a much smaller model achieves comparable performance as Transformer-big. As for training time, our approach takes slightly more time than the baseline, but is much more efficient than the regularization approach.

**Which attention module to be diversified?** We conduct extra experiments on Transformer-small to investigate which attention module benefits most from the repulsiveness. Results (see Appendix C.3) suggest that diversifying different attention module benefits differently. Remarkably, only diversifying



(a) Transformer-base-MA (b) Transformer-base-RMA

Figure 2: Distribution of heads by performance drop after masking at test time. The redundancy of heads in RMA is much less.

the attention in the first layer is able to achieve comparable performance to the case of diversifying attention in all layers, with little computational time increased. This finding suggests that the repulsiveness in the first layer’s attention plays an important role for modelling language.

**Redundancy in heads** The redundancy problem in attention has been observed in recent works (Michel et al., 2019), that a large percentage of attention heads can be removed at test time without significantly impacting performance. Following Michel et al. (2019), we analysis the redundancy in Transformer by ablating each head at testing and evaluating the performance drop. The more drops, the more important of the head. Figure 2 shows that the majority of heads in standard multi-head attention are redundant for the performance is comparable before and after masking. However, the repulsive attention largely alleviates the redundancy. More interestingly, there are a lot of counter-intuitive cases in standard attention: removing a head results in an increase in performance. However, this does not seem to happen in repulsive attention model, indicating better leveraging of the superior expressiveness of multi-head mechanism.

## 5.3 Language Representation Learning

**Model** ELECTRA (Clark et al., 2020) is an efficient approach to self-supervised language representation learning. It consists of two networks, *Generator* and *Discriminator*, both of which are parameterized by Transformers. The pre-trained Discriminator is used in various downstream tasks via fine-tuning. We apply the proposed repulsive multi-head attention to ELECTRA (small setting) in the pre-training stage. We only make the first layer attention of Discriminator to be repulsive, according to the finding in Section 5.2 that the diversity in the first attention layer of Transformer

| Pre-training Data             | Model      | Prms  | CoLA        | SST         | MRPC        | STS         | QQP         | MNLI        | QNLI        | RTE                | Avg.                      |
|-------------------------------|------------|-------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|--------------------|---------------------------|
| Wikipedia<br>+<br>BooksCorpus | TinyBERT   | 14.5M | 51.1        | 93.1        | 82.6        | 83.7        | 89.1        | 84.6        | 90.4        | 70.0               | 80.6                      |
|                               | MobileBERT | 25.3M | 51.1        | 92.6        | 84.5        | 84.8        | 88.3        | 84.3        | 91.6        | 70.4               | 81.0                      |
|                               | GPT        | 117M  | 45.4        | 91.3        | 75.7        | 80.0        | 88.5        | 82.1        | 88.1        | 56.0               | 75.9                      |
|                               | BERT-Base  | 110M  | 52.1        | 93.5        | 84.8        | 85.8        | 89.2        | 84.6        | 90.5        | 66.4               | 80.9                      |
|                               | ELECTRA    | 14M   | 54.6        | 89.1        | 83.7        | 80.3        | 88.0        | 79.7        | 87.7        | 60.8               | 78.0                      |
| OpenWebText                   | ELECTRA    | 14M   | 56.2        | <b>88.3</b> | 87.5        | 86.8        | 88.1        | 78.6        | 87.4        | <b>67.3</b> / 71.4 | 80.0 / 80.5               |
|                               | + RMA      | 14M   | <b>59.4</b> | 87.1        | <b>87.9</b> | <b>87.0</b> | <b>88.6</b> | <b>79.3</b> | <b>87.8</b> | 64.9 / <b>73.1</b> | <b>80.3</b> / <b>81.3</b> |

Table 3: Results on the GLUE test set. For RMA, repulsive attention is only applied to pre-training. For RTE, the left value is fine-tuned from pre-trained models, the right value is from intermediate task training.

benefits the most.

**Tasks & Dataset** We train ELECTRA models on OpenWebText Corpus due to the data used in Clark et al. (2020) is not publicly available. The pretrained models are then fine-tuned and evaluated on the General Language Understanding Evaluation (GLUE) (Wang et al., 2019) benchmark on eight datasets (Warstadt et al., 2019; Socher et al., 2013; Dolan and Brockett, 2005; Cer et al., 2017; Williams et al., 2018; Rajpurkar et al., 2016; Dagan et al., 2005; Haim et al., 2006; Giampiccolo et al., 2007; Bentivogli et al., 2009).

**Results** Results are shown in Table 3. For each task, we perform single-task fine-tuning 50 times, and report the averaged results. The training time with and without repulsive attention is almost the same. It shows that repulsive attention improves the baseline results (Clark et al., 2020) in seven out of eight tasks on GLUE, and the gains are larger especially on MNLI (the largest dataset on GLUE) and CoLA. This suggests that repulsive attention can yield better language representations. Since MNLI and RTE are both entailment tasks, following Clark et al. (2020) and Phang et al. (2018), we use intermediate task training for RTE. We first fine-tune the pre-trained model on MNLI, then continuously fine-tune it on RTE. The repulsive attention outperforms the baseline method by a large margin in this setting. This is probably because the repulsive attention particularly favor large data variability (e.g., MNLI dataset), where different aspects of data can be uniquely represented in different heads.

## 5.4 Graph-to-Text Generation

**Model & Baselines** GraphWriter (Koncel-Kedziorski et al., 2019) is a knowledge-graph-to-text model, which aims at generating coherent multi-sentence abstract given a knowledge graph and a title. There is a Transformer-style encoder defined with graph attention modules (Velickovic

| Metrics | GW    | + R   | + RMA        |
|---------|-------|-------|--------------|
| BLEU-1  | 42.56 | 42.25 | <b>45.60</b> |
| BLEU-2  | 27.64 | 27.98 | <b>29.96</b> |
| BLEU-3  | 19.27 | 19.77 | <b>21.07</b> |
| BLEU-4  | 13.75 | 14.21 | <b>15.12</b> |
| METEOR  | 18.11 | 18.61 | <b>19.52</b> |
| ROUGE-1 | 35.80 | 37.24 | <b>38.23</b> |
| ROUGE-2 | 16.83 | 17.78 | <b>18.39</b> |
| ROUGE-L | 27.21 | 26.90 | <b>28.55</b> |

Table 4: Automatic evaluations of generation systems on test set of AGENDA.

|                 | Win | Lose | Tie |
|-----------------|-----|------|-----|
| Structure       | 51% | 12%  | 37% |
| Informativeness | 66% | 13%  | 21% |
| Grammar         | 37% | 17%  | 46% |
| Overall         | 65% | 14%  | 21% |

Table 5: Human judgments of GraphWriter with and without repulsive attention.

et al., 2018) that could also be easily adapted to our method. We compare our method with the original one that has the standard multi-head attention, and the one with the cosine similarity regularization on attention parameters in encoder layers.

**Tasks & Datasets & Metrics** Experiments are conducted on the Abstract GENERation DATaset (AGENDA) (Koncel-Kedziorski et al., 2019), a dataset of knowledge graphs paired with scientific abstracts. We evaluate the quality of abstracts with 3 major metrics: BLEU (uni-gram to 4-gram BLEU) (Papineni et al., 2002), METEOR (Banerjee and Lavie, 2005), ROUGE (Lin and Hovy, 2003). In ROUGE, the unigram and bigram overlap (ROUGE-1 and ROUGE-2) are a proxy for assessing informativeness and the longest common subsequence (ROUGE-L) represents fluency.

**Results** The results are shown in Table 4. The GraphWriter model with repulsive multi-head attention significantly outperforms the original model and regularization approach in all metrics. Espe-

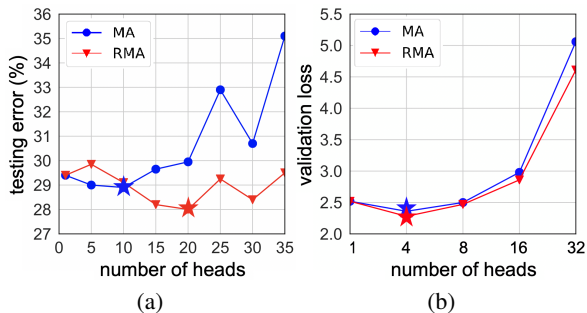


Figure 3: Demonstration on performance difference with different number of heads. (a) Testing error of the self-attentive sentence classification model on Yelp dataset. (b) Negative log likelihood loss of Transformer-small on IWSLT14 De-En dataset.

cially, the higher recall score in ROUGE shows that there are more N-grams across the reference abstracts that can be found in the generated abstracts. Similar observations are noticed when analyzing the generated examples in detail (an example is illustrated in Appendix E). Koncel-Kedziorski et al. (2019) pointed out one limitation of their model is 40% of entities in the knowledge graphs do not appear in the generated text. With the repulsive attention, remarkably, the GraphWriter model is observed to perform much better with a 10% improvement on the knowledge graph coverage and fewer repeat clauses.

**Human Evaluation** To further illustrate the improvement of using diverse attention, we conduct human evaluation. Following Koncel-Kedziorski et al. (2019), we give 50 test datapoints to experts (5 computer science students) and ask them to provide per-criterion judgments for the generated abstracts. Comparisons of the two methods from 4 aspects are shown in Table 5. The human judgment indicates that the repulsive attention improves both the structure and informativeness of generated abstracts significantly, which is consistent with the automatic evaluation and our observations.

### 5.5 On the Number of Attention Heads

Our analysis in Section 4.2 suggests the existence of the optimal number of attention heads. To verify this, we further conduct experiments on sentence classification and translation tasks by varying the number of attention heads in models. The results are shown in Figure 3. The model error/loss first decreases then increases w.r.t.  $M$ , the number of attention heads. The optimal  $M$  are around 20 and 4 for the sentiment analysis and the Transformer, re-

spectively. Interestingly, the Transformer degrades quickly as the number of heads increases. This might be because the constant corresponding to the  $O(M\epsilon_0)$  term in the bound is too large, making this term quickly dominate with increasing  $M$ . Furthermore, it is also observed that the standard multi-head attention follows the same trend, but performs much worse and is more sensitive to the  $M$ . This indicates the benefit of Bayesian modeling, which could usually stabilize a model better.

## 6 Related Work

We provide a first explanation of multi-head attention from a Bayesian perspective, and propose particle-optimization sampling for repulsive attention. Most previous works aim at improving attention diversity with regularization-based methods, *e.g.*, the Frobenius regularization on attention weights in Lin et al. (2017) and the cosine similarity regularization on attention outputs in Li et al. (2018). These works focus on a particular model and the underlying working principle has not been well-validated. Our approach is a principled one that is more interpretable and widely applicable.

The attention collapse belongs to a feature-overlapping problem, which also happens in other areas. Some works tackle this problem by changing architectures, for example ResNet (He et al., 2016) and DenseNet (Huang et al., 2017) implicitly reduce feature correlations by summing or concatenating activation from previous layers. There are also works done by altering the training method as we do. Han et al. (2017) adopt the dropout mechanism and propose a dense-sparse-dense training flow, for regularizing deep neural networks. Prakash et al. (2019) attempt addressing the unnecessary overlap in features captured by image filters with pruning-restoring scheme in training. To our knowledge, we are the first to tackle the attention-feature overlap problem from a Bayesian view with a principled interpretation.

## 7 Conclusion

We propose a principled way of understanding multi-head attention from a Bayesian perspective. We apply particle-optimization sampling to train repulsive multi-head attention with no additional trainable parameters nor explicit regularizers. Our Bayesian framework explains the long-standing question of why and how multi-head attention affects model performance. Extensive experimental



results on representative attention models demonstrate that our approach can significantly improve the diversity in multi-head attention, resulting in more expressiveness attention models with performance improvement on a wide range of tasks.

## Acknowledgements

We sincerely thank all the reviewers for providing valuable feedback. This paper is funded by the Verizon Media FREP program.

## References

- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. Neural machine translation by jointly learning to align and translate. In *3rd International Conference on Learning Representations (ICLR), San Diego, CA, USA*.
- Satanjeev Banerjee and Alon Lavie. 2005. METEOR: an automatic metric for MT evaluation with improved correlation with human judgments. In *Proceedings of the Workshop on Intrinsic and Extrinsic Evaluation Measures for Machine Translation and/or Summarization@ACL, Ann Arbor, Michigan, USA*, pages 65–72.
- Luisa Bentivogli, Ido Dagan, Hoa Trang Dang, Danilo Giampiccolo, and Bernardo Magnini. 2009. The fifth pascal recognizing textual entailment challenge.
- 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.
- Kevin Clark, Minh-Thang Luong, Quoc V. Le, and Christopher D. Manning. 2020. ELECTRA: pre-training text encoders as discriminators rather than generators. In *8th International Conference on Learning Representations (ICLR), Addis Ababa, Ethiopia*.
- Ido Dagan, Oren Glickman, and Bernardo Magnini. 2005. The PASCAL recognizing textual entailment challenge. In *Machine Learning Challenges, Evaluating Predictive Uncertainty, Visual Object Classification and Recognizing Textual Entailment, First PASCAL Machine Learning Challenges Workshop (MLCW), Southampton, UK*, pages 177–190.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT), Minneapolis, MN, USA*, pages 4171–4186.
- William B. Dolan and Chris Brockett. 2005. Automatically constructing a corpus of sentential paraphrases. In *Proceedings of the Third International Workshop on Paraphrasing, IWP@IJCNLP, Jeju Island, Korea*.
- Danilo Giampiccolo, Bernardo Magnini, Ido Dagan, and Bill Dolan. 2007. The third PASCAL recognizing textual entailment challenge. In *Proceedings of the ACL-PASCAL@ACL Workshop on Textual Entailment and Paraphrasing, Prague, Czech Republic*, pages 1–9.
- Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q. Weinberger. 2017. On calibration of modern neural networks. In *Proceedings of the 34th International Conference on Machine Learning (ICML), Sydney, NSW, Australia*, pages 1321–1330.
- Roy Bar Haim, Ido Dagan, Bill Dolan, Lisa Ferro, Danilo Giampiccolo, Bernardo Magnini, and Idan Szpektor. 2006. The second pascal recognising textual entailment challenge.
- Song Han, Jeff Pool, Sharan Narang, Huizi Mao, Enhao Gong, Shijian Tang, Erich Elsen, Peter Vajda, Manohar Paluri, John Tran, Bryan Catanzaro, and William J. Dally. 2017. DSD: dense-sparse-dense training for deep neural networks. In *5th International Conference on Learning Representations (ICLR), Toulon, France*.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Las Vegas, NV, USA*, pages 770–778.
- Gao Huang, Zhuang Liu, Laurens van der Maaten, and Kilian Q. Weinberger. 2017. Densely connected convolutional networks. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Honolulu, HI, USA*, pages 2261–2269.
- Diederik P. Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. In *3rd International Conference on Learning Representations (ICLR), San Diego, CA, USA*.
- Rik Koncel-Kedziorski, Dhanush Bekal, Yi Luan, Mirella Lapata, and Hannaneh Hajishirzi. 2019. Text generation from knowledge graphs with graph transformers. In *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT), Minneapolis, MN, USA*, pages 2284–2293.
- Olga Kovaleva, Alexey Romanov, Anna Rogers, and Anna Rumshisky. 2019. Revealing the dark secrets of bert. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing and the International Joint Conference on Natural Language Processing (EMNLP-IJCNLP), Hong Kong, China*, pages 4365–4374.

- Jian Li, Zhaopeng Tu, Baosong Yang, Michael R. Lyu, and Tong Zhang. 2018. Multi-head attention with disagreement regularization. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP), Brussels, Belgium*, pages 2897–2903.
- Chin-Yew Lin and Eduard Hovy. 2003. Automatic evaluation of summaries using n-gram co-occurrence statistics. In *Human Language Technology Conference of the North American Chapter of the Association for Computational Linguistics (HLT-NAACL), Edmonton, Canada*.
- Zhouhan Lin, Minwei Feng, Cícero Nogueira dos Santos, Mo Yu, Bing Xiang, Bowen Zhou, and Yoshua Bengio. 2017. A structured self-attentive sentence embedding. In *5th International Conference on Learning Representations (ICLR), Toulon, France*.
- Qiang Liu and Dilin Wang. 2016. Stein variational gradient descent: A general purpose bayesian inference algorithm. In *Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems (NeurIPS), Barcelona, Spain*, pages 2370–2378.
- Christos Louizos and Max Welling. 2017. Multiplicative normalizing flows for variational bayesian neural networks. In *Proceedings of the 34th International Conference on Machine Learning (ICML), Sydney, NSW, Australia*, pages 2218–2227.
- Paul Michel, Omer Levy, and Graham Neubig. 2019. Are sixteen heads really better than one? In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems (NeurIPS), Vancouver, BC, Canada*, pages 14014–14024.
- 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. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT), Minneapolis, MN, USA*, pages 48–53.
- 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 (ACL), Philadelphia, PA, USA*, pages 311–318.
- Jason Phang, Thibault Févry, and Samuel R. Bowman. 2018. Sentence encoders on stilts: Supplementary training on intermediate labeled-data tasks. *arXiv preprint arXiv:1811.01088*.
- Matt Post. 2018. A call for clarity in reporting BLEU scores. In *Proceedings of the Third Conference on Machine Translation: Research Papers, Brussels, Belgium*, pages 186–191.
- Aaditya Prakash, James A. Storer, Dinei A. F. Florêncio, and Cha Zhang. 2019. Repr: Improved training of convolutional filters. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Long Beach, CA, USA*, pages 10666–10675.
- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. Squad: 100, 000+ questions for machine comprehension of text. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP), Austin, Texas, USA*, pages 2383–2392.
- Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D. Manning, Andrew Y. Ng, and Christopher Potts. 2013. Recursive deep models for semantic compositionality over a sentiment treebank. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP), Grand Hyatt Seattle, Seattle, Washington, USA*, pages 1631–1642.
- Sunil Thulasidasan, Gopinath Chennupati, Jeff A. Bilmes, Tanmoy Bhattacharya, and Sarah Michalak. 2019. On mixup training: Improved calibration and predictive uncertainty for deep neural networks. In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems (NeurIPS), Vancouver, BC, Canada*, pages 13888–13899.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems (NeurIPS), Long Beach, CA, USA*, pages 5998–6008.
- Petar Velickovic, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. 2018. Graph attention networks. In *6th International Conference on Learning Representations (ICLR), Vancouver, BC, Canada*.
- Elena Voita, David Talbot, Fedor Moiseev, Rico Senrich, and Ivan Titov. 2019. Analyzing multi-head self-attention: Specialized heads do the heavy lifting, the rest can be pruned. In *Proceedings of the 57th Conference of the Association for Computational Linguistics (ACL), Florence, Italy*, pages 5797–5808.
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R. Bowman. 2019. GLUE: A multi-task benchmark and analysis platform for natural language understanding. In *7th International Conference on Learning Representations (ICLR), New Orleans, LA, USA*.
- Alex Warstadt, Amanpreet Singh, and Samuel R. Bowman. 2019. Neural network acceptability judgments. *Transactions of the Association for Computational Linguistics*, pages 625–641.

- Max Welling and Yee Whye Teh. 2011. Bayesian learning via stochastic gradient langevin dynamics. In *Proceedings of the 28th International Conference on Machine Learning (ICML), Bellevue, Washington, USA*, pages 681–688.
- Adina Williams, Nikita Nangia, and Samuel R. Bowman. 2018. A broad-coverage challenge corpus for sentence understanding through inference. In *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT), New Orleans, Louisiana, USA*, pages 1112–1122.
- Jianyi Zhang, Ruiyi Zhang, and Changyou Chen. 2020. Stochastic particle-optimization sampling and the non-asymptotic convergence theory. In *The 23rd International Conference on Artificial Intelligence and Statistics (AISTATS), Palermo, Sicily, Italy*, pages 1877–1887.

## A Additive Attention

First proposed by (Bahdanau et al., 2015), additive attention uses a one-hidden layer feed-forward network to calculate the attention alignment. We use the attention function in Lin et al. (2017), which is also a self-attention, as an example. It aims at extracting the latent representation of a sentence. The single-head attention function is:

$$\mathbf{a} = \text{Softmax}(\mathbf{v}^\top \tanh(\mathbf{W}\mathbf{H}^\top)), \mathbf{z} = \mathbf{a}\mathbf{H}$$

where  $\mathbf{H} \in \mathbb{R}^{n \times d}$  is the hidden state matrix of a sentence with  $n$  words, every word is embedded in a  $d$  dimensional vector.  $\mathbf{v} \in \mathbb{R}^{1 \times n}$  is the normalized alignment score vector for each word.  $\mathbf{W} \in \mathbb{R}^{d_a \times d}$  and  $\mathbf{v} \in \mathbb{R}^{d_a \times 1}$  are attention parameters. The final sentence representation vector  $\mathbf{z}$  is a weighted sum of words' hidden states weighted by attention vector. In order to capture overall semantics of the sentence instead of a specific component, multi-head attention could be applied as

$$\mathbf{A} = \text{Softmax}(\mathbf{V}^\top \tanh(\mathbf{W}\mathbf{H}^\top)), \mathbf{Z} = \mathbf{A}\mathbf{H}$$

where  $\mathbf{V} \in \mathbb{R}^{d_a \times M}$  is the matrix performs  $M$  heads,  $\mathbf{A} \in \mathbb{R}^{M \times n}$  is the  $M$ -head attention matrix and  $\mathbf{Z} \in \mathbb{R}^{M \times d}$  is the resulting sentence representation matrix contains semantics from multiple aspects.

## B Additional Experimental Details

For our approach, RBF kernel  $\kappa(x, y) = \exp(-\frac{1}{h}\|x - y\|_2^2)$  with the bandwidth  $h = med^2 / \log M$  is used as the kernel function, where  $med$  denotes the median of the pairwise distance between current particles. The prior distribution of attention parameters is assumed to be uniform. We find that adding an repulsive weight before the repulsive term (i.e. the second term in Eq. 4) in particle-optimization update rules could help adjusting the degree of diversity in attention and achieving better performance. In our experiments, we adopt this trick and use the hyper-parameter  $\alpha$  to denote the repulsive weight. Since our method only modifies the learning process of attention, all models and settings in our experiments kept the same with the corresponding previous work unless stated otherwise.

### B.1 Self-attentive Sentence Classification

**Dataset** Three tasks are conducted on three public sentence classification datasets. *Author profiling*

(Age dataset <sup>1</sup>) is to predict the age range of the user by giving their tweets. *Sentiment analysis* (Yelp dataset <sup>2</sup>) is to predict the number of stars the user assigned to by analysis their reviews. *Textual entailment* (SNLI dataset <sup>3</sup>) is to tell whether the semantics in the two sentences are entailment or contradiction or neutral. Following Lin et al. (2017), the train / validate / test split of Age is 68485 / 4000 / 4000, Yelp is 500K / 2000 / 2000, SNLI is 550K / 10K / 10K.

**Experimental settings** We implement the standard multi-head attention model in Lin et al. (2017) following the settings in it except that we use Spacy toolkit <sup>4</sup> as the tokenizer and GloVe <sup>5</sup> (GloVe 840B 300D) as the pre-trained word embedding. For repulsive multi-head attention learning, we keep all settings the same with the standard one (Lin et al., 2017). Hyper-parameters  $\epsilon$  and  $\alpha$  in SVGD are selected with grid search. For SPOS, we fix these two hyper-parameters and only tune  $\beta$ . The selection is based on the performance on the validation data. We train and evaluate all the models with 10 random seeds and compare their average performance. Models are trained on one TITAN Xp GPU.

### B.2 Transformer-based Neural Translation

**Dataset** IWSLT14 German-to-English (De-En) dataset contains 153K / 7K / 7K sentence pairs. WMT14 English-to-German (En-De) dataset contains about 4.5 million training sentence pairs and uses newstest2013 dataset as the validation set, newstest2014 dataset as the test set. Data and the processing scripts could be found here <sup>6</sup>.

**Experimental settings** Our implementation is based on the open-sourced fairseq <sup>7</sup> (Ott et al., 2019). We follow the settings in Vaswani et al. (2017) and have reproduced their reported results. For the WMT14 dataset, the `base` Transformer is used, which consists of a 6-layer encoder and a 6-layer decoder. The size of the hidden units and embeddings is 512 and the number of heads is 8. The `big` Transformers has 1024 hidden units and 16 heads, which is listed as a reference. For IWSLT14 dataset, the `small` setting is used and

<sup>1</sup><https://pan.webis.de/clef16/pan16-web/author-profiling.html>

<sup>2</sup><https://www.yelp.com/dataset/download>

<sup>3</sup><https://nlp.stanford.edu/projects/snli/>

<sup>4</sup><https://spacy.io/>

<sup>5</sup><https://nlp.stanford.edu/projects/glove/>

<sup>6</sup><https://github.com/pytorch/fairseq/tree/v0.6.0/examples/translation>

<sup>7</sup><https://github.com/pytorch/fairseq>



the number of heads in every layer is set to 4. All the configurations are kept the same when applying our method or the regularization method. In our method, only the training process is changed and SVGD update rule is utilized in our algorithm. The stepsize  $\epsilon$  in our method is set to 0.1 and the repulsive term  $\alpha$  is set to 0.01. The Transformer-*small* model is trained on one TITAN Xp GPU. The Transformer-*base* model is trained on four GTX 1080Ti GPUs.

**Additional Results** To support a fair comparison, we also evaluate the Transformer-base model on WMT14 En-De task. The SACREBLEU score (Post, 2018) with and without our approach is 27.1 and 26.2<sup>8</sup>, respectively.

### B.3 ELECTRA

**Dataset** Following the official code of Clark et al. (2020), ELECTRA models are pretrained on the OpenWebTextCorpus<sup>9</sup> dataset, an open source effort to reproduce OpenAI’s WebText dataset. OpenWebTextCorpus contains 38GB of text data from 8,013,769 documents. The pretrained model is then finetuned and evaluated on GLUE benchmark<sup>10</sup>. GLUE contains a variety of tasks covering textual entailment (RTE and MNLI) question-answer entailment (QNLI), paraphrase (MRPC), question paraphrase (QQP), textual similarity (STS), sentiment (SST), and linguistic acceptability (CoLA). Our evaluation metrics are Spearman correlation for STS, Matthews correlation for CoLA, and accuracy for the other GLUE tasks.

**Experiment settings** The ELECTRA-*small* model we implemented follow all official settings<sup>11</sup> except that it is fully-trained on one GTX 1080Ti GPU for 6 days. The ELECTRA-*small* model has 12 layers with 4 heads in every layer’s attention. For our method, the stepsize  $\epsilon$  is set to 0.01 and the repulsive term  $\alpha$  is set to 0.1. The repulsive learning of attention is only applied to the pre-training stage. The fine-tuning remains the same with the original one.

<sup>8</sup>SacreBLEU hash: BLEU+case.mixed+lang.en-de+numrefs.1+smooth.exp+test.wmt14/full+tok.13a+version.1.4.12

<sup>9</sup><https://skylion007.github.io/OpenWebTextCorpus/>

<sup>10</sup><https://gluebenchmark.com/>

<sup>11</sup><https://github.com/google-research/electra>

### B.4 GraphWriter

**Dataset** Experiments are conducted on the Abstract GENERation DATaset (AGENDA) (Koncel-Kedziorski et al., 2019), a dataset of knowledge graphs paired with scientific abstracts. It consists of 40k paper titles and abstracts from the Semantic Scholar Corpus taken from the proceedings of 12 top AI conferences. We use the standard split of AGENDA dataset in our experiments: 38,720 for training, 1000 for validation, and 1000 for testing.

**Experimental settings** We follow the official settings<sup>12</sup> in Koncel-Kedziorski et al. (2019) with the encoder containing 6 layers and 4-head graph attention in every layer. We reproduce their results and keep all settings the same when applying the proposed repulsive attention. The SVGD update rule is used in our algorithm and applied to all layers. The stepsize  $\epsilon$  is set to 0.1 and the repulsive weight is set to 0.01 in this experiment. The model is trained on one TITAN Xp GPU.

## C Additional Analysis of Our Approach

### C.1 Comparison with SGLD

We also conducted a comparison of our method with Stochastic gradient Langevin dynamics (SGLD) (Welling and Teh, 2011), which is also a Bayesian sampling method. Results are in Table 6. Though random noise brought by SGLD might help achieving diversity, it’s sub-optimal. Using particle-optimization to add the repulsive term makes it more effective.

| Sampling Method | Age          | Yelp        | SNLI         | IWSLT14 De-En |
|-----------------|--------------|-------------|--------------|---------------|
| None            | 81.47        | 69.3        | 83.79        | 34.4          |
| SGLD            | 81.57        | 70.1        | 83.80        | 34.7          |
| SVGD            | <b>81.82</b> | <b>71.2</b> | <b>84.58</b> | <b>35.2</b>   |

Table 6: Performance of attention models with different sampling methods on four tasks. For Age, Yelp and SNLI sentence classification tasks, the evaluation metric is accuracy (%). For IWSLT14 De-En translation task, the evaluation metric is BLEU score. "None" here means the standard multi-head attention models.

### C.2 What Prior to Use?

In our approach, the repulsiveness is imposed by the inference algorithm (*i.e.* SVGD), not prior. To study the impact of different priors, we also tested

<sup>12</sup><https://github.com/rikdz/GraphWriter>

the Gaussian prior. We found that (see Table 7) different priors have little impact on the final results, *i.e.*, there is not a consistent winner for different priors. This suggests that, the prior has little impact on repulsiveness in our framework. But one can still impose prior knowledge of the attention to help our algorithm learn a better attention model. We would like to explore that in future works.

| Prior    | Age   | Yelp | SNLI  |
|----------|-------|------|-------|
| Uniform  | 81.82 | 71.7 | 84.58 |
| Gaussian | 81.88 | 71.6 | 84.28 |

Table 7: Performance of our approach with different priors on three sentence classification tasks. The evaluation metric is accuracy (%)

### C.3 Which attention modules to be diversified?

| Models            | BLEU        | Time |
|-------------------|-------------|------|
| MA                | 34.4        | 1    |
| RMA (Q)           | 34.7        | 1.06 |
| RMA (K)           | 34.7        |      |
| RMA (V)           | 34.9        |      |
| RMA (En)          | 34.6        | 1.06 |
| RMA (De)          | 34.7        |      |
| RMA (En-De)       | 34.9        |      |
| RMA (first layer) | <b>35.1</b> | 1.03 |
| RMA (last layer)  | 34.7        |      |
| RMA (All)         | <b>35.2</b> | 1.13 |

Table 8: Ablation study of Transformer-small-RMA model on IWSLT14 De-En dataset with repulsive multi-head attention applied on different part of the model. En: self-attention in encoder. De: self-attention in decoder. En-De: inter-attention between encoder and decoder. Time: relative training time of every step versus MA.

There are three types of attention in Transformer: self-attention in the encoder, self-attention in the decoder, and inter-attention between the encoder and decoder. We conduct extra experiments on Transformer-small to investigate which attention module benefits most from the repulsiveness. Results are shown in Table 8. We first apply the repulsive attention on each of {Q,K,V} parameters in every attention module for all layers. The results indicate that diversifying the V-parameter seems to yield better performance. We then compare repulsive attention inside the encoder, inside the decoder and between them, respectively. The results

show improvement in all cases, and diversifying inter-attention seems to achieve the most benefit. Finally, we diversify the attention in different layers of the Transformer. The results suggest that only diversifying the attention in the first layer is able to achieve comparable performance to the case of diversifying all layers, with little computational time increased.

### C.4 Improved Calibration

A reliable model must not only be accurate, but also indicate when it is likely to get the wrong answer. It means the confidence of a well calibrated model should be indicative of the actual likelihood of correctness. Following the calibration metrics in Guo et al. (2017) and Thulasidasan et al. (2019), we evaluate the calibration of the model in Figure 4. For classifiers, the predicted softmax scores of winning class are represented as the confidence of models. Expected Calibration Error (ECE) and Overconfidence Error (OE) are two calibration metrics evaluating the reliability of a model. Following Guo et al. (2017) and Thulasidasan et al. (2019), softmax predictions are grouped into  $M$  interval bins of equal size. Let  $B_m$  be the set of samples whose prediction scores (the winning softmax score) fall into bin  $B_m$ . The accuracy and confidence of  $B_m$  are defined as

$$\text{acc}(B_m) = \frac{1}{|B_m|} \sum_{i \in B_m} \mathbf{1}(\hat{y}_i = y_i)$$

$$\text{conf}(B_m) = \frac{1}{|B_m|} \sum_{i \in B_m} \hat{p}_i$$

where  $\hat{p}_i$  is the confidence (winning score) of sample  $i$ . The Expected Calibration Error (ECE) is then defined as:

$$\text{ECE} = \sum_{m=1}^M \frac{|B_m|}{n} |\text{acc}(B_m) - \text{conf}(B_m)|$$

In high-risk applications, confident but wrong predictions can be especially harmful. Overconfidence Error (OE) is defined as follow for this case.

$$\text{OE} = \sum_{m=1}^M \frac{|B_m|}{n} [\text{conf}(B_m) \times \max(\text{conf}(B_m) - \text{acc}(B_m), 0)]$$

As shown in Figure 4, the standard attention model is prone to be over-confident, meaning that the accuracy is likely to be lower than what is indicated

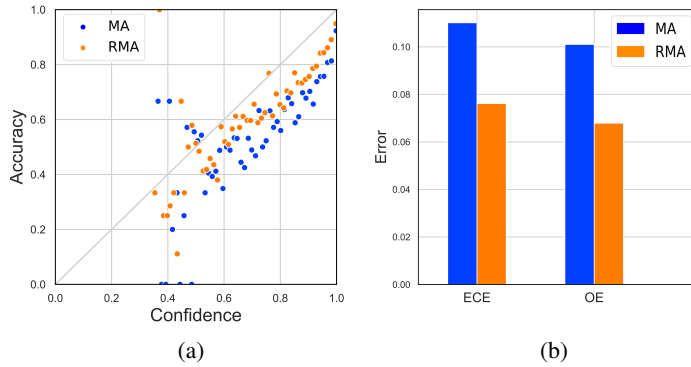


Figure 4: (a) Scatterplots for accuracy and confidence for SNLI test set. The repulsive case (RMA) is much better calibrated with the points lying closer to the  $y = x$  line, while in the standard case (MA), points tend to lie in the overconfident region. (b) Expected Calibration Error (ECE) and Overconfidence Error (OE) for two cases.




Figure 5: Empirical CDF for the entropy of the predictive distributions on Age dataset from the model trained on Yelp dataset. Curves that are closer to the bottom right part of the plot are preferable, as it denotes that the probability of observing a high confidence prediction is low.

by the predictive score. With the proposed repulsive training of attention, the model becomes better calibrated with much lower calibration error and overconfidence error, indicating that our method is beneficial for training more reliable multi-head attention models.

### C.5 Improved Uncertainty Prediction

We follow Louizos and Welling (2017) to evaluate the predictive uncertainty. We estimate the entropy of the predictive distributions on Age dataset (out-of-distribution entropy) from the models trained on Yelp dataset. Since we a-priori know that none of the Age classes correspond to a trained class (they are two different tasks), the ideal predictive distribution is uniform over the Age dataset, i.e. a

| Methods                          | Acc(%) |       |
|----------------------------------|--------|-------|
|                                  | Age    | Yelp  |
| cosine similarity regularization | 81.35  | 68.50 |
| $i - > j$                        | 81.65  | 71.45 |
| $i - > j, +smooth$               | 81.85  | 71.50 |

Table 9: Adapt cosine similarity regularization on attention parameters gradually to our framework. Accuracy of the model is evaluated on the test set of Age and Yelp dataset.

maximum entropy distribution. We plot the empirical CDF for the entropy in Figure 5. It shows that the uncertainty estimates from the repulsive multi-head attention model is better than the standard attention and the regularization approach.

### C.6 Discussion with Existing Regularization-based Methods

Learning diverse attentions has been proposed in Lin et al. (2017) and Li et al. (2018), with different regularization techniques to enforce repulsiveness. In fact, we can show that existing methods are simplified versions of our framework, but with a potential mismatch between their algorithm and the underlying repulsiveness guarantee. To explain this, we follow Li et al. (2018) and apply cosine similarity regularizer to the attention parameter  $\theta$ . When negative log-likelihood is used as the loss function and the prior of  $\theta^{(i)}$  is assumed to be uniform, the update function  $\phi$  becomes:

$$\phi(\theta_\ell^{(i)}) = -\nabla_{\theta_\ell^{(i)}} U(\theta_\ell^{(i)}) + \nabla_{\theta_\ell^{(i)}} \frac{1}{M} \sum_{j=1}^M \frac{\theta_\ell^{(i)} \theta_\ell^{(j)}}{\|\theta_\ell^{(i)}\| \|\theta_\ell^{(j)}\|}$$

One can consider the cosine similarity as a kernel function. Applying this new kernel function to our

framework results in the following update:

$$\begin{aligned}\phi(\boldsymbol{\theta}_\ell^{(i)}) &= \frac{1}{M} \sum_{j=1}^M -\frac{\boldsymbol{\theta}_\ell^{(i)} \boldsymbol{\theta}_\ell^{(j)}}{\|\boldsymbol{\theta}_\ell^{(i)}\| \|\boldsymbol{\theta}_\ell^{(j)}\|} \nabla_{\boldsymbol{\theta}_\ell^{(j)}} U(\boldsymbol{\theta}_\ell^{(j)}) \\ &+ \frac{1}{M} \sum_{j=1}^M \nabla_{\boldsymbol{\theta}_\ell^{(j)}} \frac{\boldsymbol{\theta}_\ell^{(i)} \boldsymbol{\theta}_\ell^{(j)}}{\|\boldsymbol{\theta}_\ell^{(i)}\| \|\boldsymbol{\theta}_\ell^{(j)}\|}\end{aligned}$$

It is clear that our method (??) reduces to the regularizing method when 1) removing the smoothing term  $\frac{\boldsymbol{\theta}_\ell^{(i)} \boldsymbol{\theta}_\ell^{(j)}}{\|\boldsymbol{\theta}_\ell^{(i)}\| \|\boldsymbol{\theta}_\ell^{(j)}\|}$  for  $j \neq i$ ; and 2) replacing the derivative  $\nabla_{\boldsymbol{\theta}_\ell^{(j)}}$  with  $\nabla_{\boldsymbol{\theta}_\ell^{(i)}}$  in the repulsive term. For this reason, we argue that the regularization method lacks of a formal guarantee on the repulsiveness.

To show this, we adapt the cosine similarity regularization on attention parameters gradually to our framework by 1) replacing the derivative  $\nabla_{\boldsymbol{\theta}_\ell^{(i)}}$  with  $\nabla_{\boldsymbol{\theta}_\ell^{(j)}}$  in the repulsive term; 2) adding the smoothing term  $\kappa(\boldsymbol{\theta}_\ell^{(j)}, \boldsymbol{\theta}_\ell^{(i)})$ . Table 9 shows that both the smoothing and the corrected gradient lead to performance improvement over the standard regularization methods.









## E A Generation Sample of Graph-to-Text Task

|             |   |
|-------------|---|
| Title       | Bilinear signal synthesis in array processing   |
| Knowledge   | (source positions, CONJUNCTION, angular separations); (time-frequency synthesis techniques, USED-FOR, multiple source signals); (non-integer values, USED-FOR, weighing function); (weighing function, USED-FOR, time-frequency domain); (quadratic t-f distributions, USED-FOR, source separation techniques); (spatial time-frequency distributions, USED-FOR, blind source separation methods); (time-frequency synthesis techniques, USED-FOR, antenna array)   |
| GraphWriter | In this paper, we propose a novel method for blind source separation methods based on a weighing function. In the proposed method, the spatial signatures of sources is modeled as a weighing function, and the spatial time-frequency distributions is modeled as a weighing function. The spatial time-frequency distributions is modeled as a weighing function in the time-frequency domain. The proposed method is based on the use of a weighing function to estimate the parameters of the blind source separation methods. The proposed method is based on the use of a weighing function to estimate the parameters of the blind source separation methods. The proposed method is applied to the problem of multiple source signals in the time-frequency domain. The proposed method is applied to the problem of multiple source signals in the time-frequency domain. The proposed method is shown to be robust to noise levels, whitening, and angular separations.   |
| + RMA       | In this paper, we present a novel approach to multiple source signals based on the spatial signatures of sources and angular separations. In particular, we consider the problem of estimating the spatial signatures of sources from a single image of an antenna array. We propose a novel method for estimating the spatial signatures of sources using a weighing function, where the source positions and angular separations are estimated in a time-frequency domain. We show that the proposed method can be used to derive a weighing function for the communication channel. The proposed method is based on the use of quadratic t-f distributions in the time-frequency domain and the non-integer values. The proposed method is applied to the problem of whitening and angular separations. The experimental results show that the proposed method is more robust to noise levels and noise levels.  |
| Gold        | Multiple source signals impinging on an antenna array can be separated by time-frequency synthesis techniques. Averaging of the time-frequency distributions of the data across the array permits the spatial signatures of sources to play a fundamental role in improving the synthesis performance. Array averaging introduces a weighing function in the time-frequency domain that decreases the noise levels, reduces the interactions of the source signals, and mitigates the crossterms. This is achieved independent of the temporal characteristics of the source signals and without causing any smearing of the signal terms. The weighing function may take non-integer values, which are determined by the communication channel, the source positions and their angular separations. Unlike the recently devised blind source separation methods using spatial time-frequency distributions, the proposed method does not require whitening or retrieval of the source directional matrix. The paper evaluates the proposed method in terms of performance and computations relative to the existing source separation techniques based on quadratic t-f distributions. |

Figure 9: Example outputs of various systems versus Gold