FY2023/2023年度

Tensor Learning Team

Qibin Zhao

テンソル学習チーム　趙 啓斌

RIKEN Center for Advanced Intelligence Project

Outline

- **Tensor network algorithms:** Our team has developed several flexible and nonlinear tensor decomposition algorithms that can capture more complex relations within tensor data, robust tensor algorithm with particular structure of outliers, and a very fast tensor network structure search algorithm.
- **Adversarial robustness of DNNs:** Our team has studied how low-rankness of NN parameters affects its robustness to adversarial attack from the theoretical perspective. We also developed an adaptive adversarial purification method that can be updated with adversarial loss, thus improving its robust accuracy and outperforming adversarial training methods.

Nonlinear Flexible Tensor Decomposition

Problem and challenge

- Most TD models are multilinear and cannot capture nonlinear structure within tensor data
- The distribution of observation model needs to be pre-defined
- Tensor decomposition structure needs to be pre-defined

Undirected Graphical Model for Tensor Decomposition

(Tao et al. NeurIPS'23)

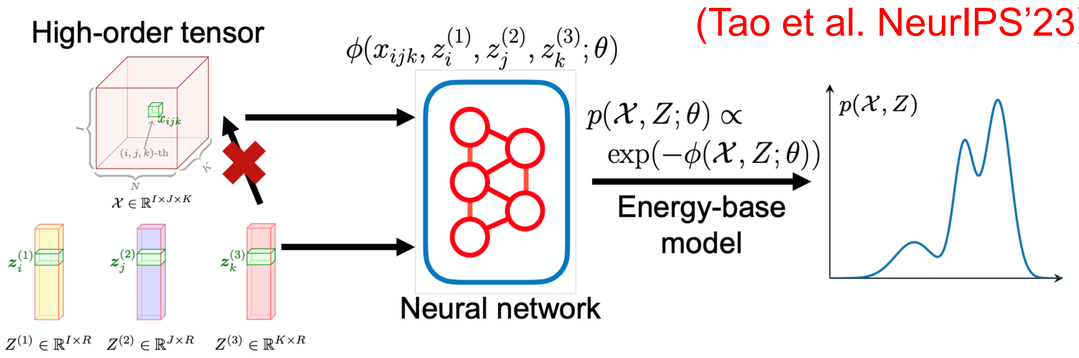

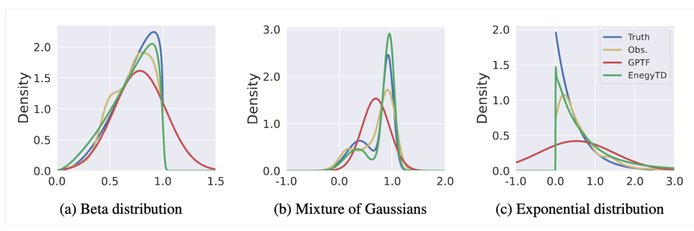

Key techniques

- Choosing optimal TD structure is hard -> Learn unknown latent structures using undirected mappings
- Choosing optimal TD distribution is hard -> Learn unknown distributions using deep generative model
- Efficient self-supervised loss function to learn the model

Efficient Nonparametric Tensor Decomposition for Binary and Count Data

(Tao et al. AAAI'24)

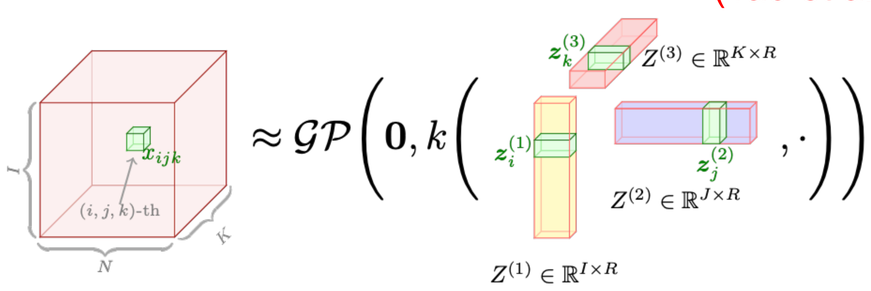

Key techniques

- Modeling flexible latent structures in TD using nonparametric model
- Gaussian process is not scalable -> Derive scalable lower bound for the likelihood
- Efficient natural gradient update results in faster convergence

Tensor Network Structure Search

(Li et al. ICML'23)

- How to choose the best tensor network structure?

$$\min_{(G,r)\in\mathbb{G}\times\mathbb{F}_G} \left(\underbrace{\phi(G,r)}_{\text{model complexity}} + \lambda \cdot \underbrace{\min_{\mathcal{Z}\in TNS(G,r)} \pi_{\mathcal{X}}(\mathcal{Z})}_{\text{model expressivity}} \right)$$

- An efficient Alternating Local Enumeration (TnALE) algorithm with fewer evaluations.

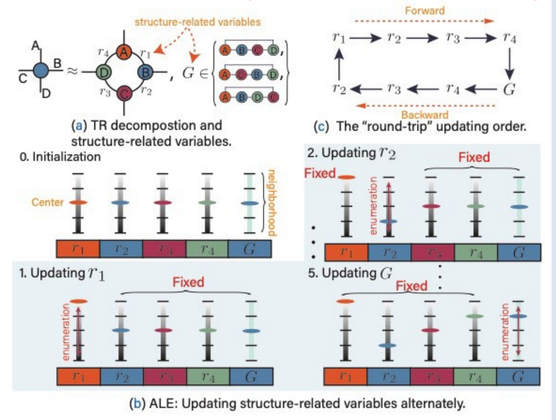

Adversarial Robustness of DNNs

(Wang et al. NeurIPS'23)

- **Goal**: To understand the impact of low-rank parameters on adversarial robustness of over-parameterized DNNs in theory.

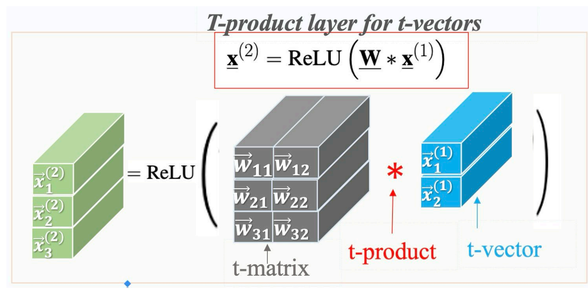

- The adversarial generalization bound with low-rank para.

$$O\left(\sqrt{c\textstyle\sum_l r_l(d_{l-1}+d_l)/N}\right)$$ *in # effective paras.*

better than w/o low-rank para.

$$O\left(\sqrt{c\textstyle(\sum_{l=1}^{L} d_{l-1}d_l)/N}\right)$$ *in # total paras.*

Low-rank parameterization can enhance adversarial robustness under over-parameterization!

Adversarial Purification using Pre-trained Generative Model

(Lin et al. ICLR'24)

- **Goal**: Improve robust accuracy of adversarial purification against known attacks
- **Method**: Fine-tune the purifier model using adversarial loss

Table 3: Accuracy comparison of defenses with vanilla model (negative impacts are marked in red).

Defense method	Clean images	Known attacks	Unseen attacks	Training cost
Vanilla model	~90%	~0%	~0%	/
Expectation	=	↑↑↑	↑↑	↑
AT	↓↓	↑↑↑	≈	↑↑
AP	↓	↑↑	↑↑	/
AToP (Ours)	↓	↑↑↑	↑↑	↑↑

Pre-training GAN-based model:

$$L_{\theta_g} = L_g(x, \theta_g) = \mathbb{E}[D(x)] - \mathbb{E}[D(g(x, \theta_g))] + \mathbb{E}\left[||x - g(x, \theta_g)||_1\right]$$

Fine-tuning GAN-based model:

$$L_{\theta_g} = L_g(x', \theta_g) + s \cdot L_{cls}(x', y, \theta_g, \theta_f) = \mathbb{E}[D(x')] - \mathbb{E}[D(g(x', \theta_g))] + \mathbb{E}\left[||x' - g(x', \theta_g)||_1\right] + s \cdot max_\delta CE\{y, f(g(x', \theta_g))\}$$

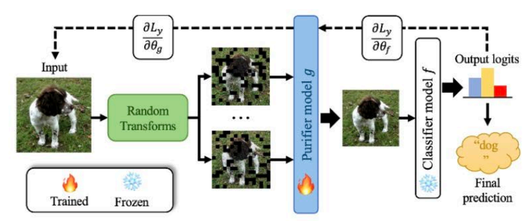

Figure 1: Illustration of adversarial training on purification (AToP).

Related Publications

- G. Lin, C. Li, J. Zhang, T. Tanaka, and Q. Zhao, "Adversarial Training on Purification (AToP): Advancing Both Robustness and Generalization", ICLR 2024.
- Y. Qiu, G. Zhou, A. Wang, Z. Huang, and Q. Zhao, "Towards multi-mode outlier robust tensor ring decomposition", AAAI 2024.
- Z. Tao, T. Tanaka, and Q. Zhao, "Efficient nonparametric tensor decomposition for binary and count data", AAAI 2024.
- Z. Tao, T. Tanaka, and Q. Zhao, "Undirected probabilistic model for tensor decomposition", NeurIPS 2023
- A. Wang, C. Li, M. Bai, Z. Jin, G. Zhou, and Q. Zhao, "Transformed low-rank parameterization can help robust generalization for tensor neural networks", NeurIPS 2023.
- C. Li, J. Zeng, C. Li, C. F. Caiafa, and Q. Zhao, "Alternating local enumeration (TNALE): Solving tensor network structure search with fewer evaluations", ICML 2023.