



# Black-Box Tuning for Language-Model-as-a-Service

### **Tianxiang Sun**

School of Computer Science, Fudan University

https://txsun1997.github.io/

## Acknowledgment



Tianxiang Sun<sup>1</sup>



Yunfan Shao<sup>1</sup>



Hong Qian<sup>2</sup>



Xuanjing Huang<sup>1</sup>





Yang Yu<sup>4</sup>



Zhengfu He<sup>1</sup>

<sup>1</sup> Fudan University

<sup>2</sup> East China Normal University

<sup>3</sup> Peng Cheng Laboratory

<sup>4</sup> Nanjing University

### Acknowledgment





Tianxiang Sun<sup>1</sup>

NLP/Deep Learning

Yunfan Shao<sup>1</sup>

Yang Yu<sup>4</sup>



Hong Qian<sup>2</sup>



Xuanjing Huang<sup>1</sup>



Xipeng Qiu<sup>1,3</sup>



Zhengfu He<sup>1</sup>

<sup>3</sup> Peng Cheng Laboratory

<sup>4</sup> Nanjing University

<sup>1</sup> Fudan University

<sup>2</sup> East China Normal University

## Acknowledgment



Tianxiang Sun<sup>1</sup>



Yunfan Shao<sup>1</sup>



Hong Qian<sup>2</sup>



Xuanjing Huang<sup>1</sup>



Xipeng Qiu<sup>1,3</sup>

**RL/Black-Box Optimization** 





Zhengfu He<sup>1</sup>

<sup>1</sup> Fudan University

<sup>2</sup> East China Normal University

<sup>3</sup> Peng Cheng Laboratory

<sup>4</sup> Nanjing University

### Pre-train, then fine-tune





In the era of large language models (LLMs)...

- **Servers** often do not open-source the weights of LLMs due to commercial reasons
- Users usually do not have enough resources to run LLMs

Why did OpenAI choose to release an API instead of open-sourcing the models?

There are three main reasons we did this. First, commercializing the technology helps us pay for our ongoing AI research, safety, and policy efforts.

Second, many of the models underlying the API are very large, taking a lot of expertise to develop and deploy and making them very expensive to run. This makes it hard for anyone except larger companies to benefit from the underlying technology. We're hopeful that the API will make powerful AI systems more accessible to smaller businesses and organizations.

Third, the API model allows us to more easily respond to misuse of the technology. Since it is hard to predict the downstream use cases of our models, it feels inherently safer to release them via an API and broaden access over time, rather than release an open source model where access cannot be adjusted if it turns out to have harmful applications.

In the era of large language models (LLMs)...

- **Servers** often do not open-source the weights of LLMs due to commercial reasons
- Users usually do not have enough resources to run LLMs

The emergent ability of LLMs

- Manually craft text prompt to query LLMs
- In-context learning (GPT-3, Brown et al., 2020)

### Why in-context learning?

- Generalization: One general purpose model for all tasks
- Backpropagation is expensive
- Commercial use

The three settings we explore for in-context learning

#### Zero-shot

The model predicts the answer given only a natural language description of the task. No gradient updates are performed.



#### One-shot

In addition to the task description, the model sees a single example of the task. No gradient updates are performed.

Translate English to French:	-	task description
sea otter => loutre de mer	c	example
cheese =>	6	prompt

#### Few-shot

In addition to the task description, the model sees a few examples of the task. No gradient updates are performed.



Traditional fine-tuning (not used for GPT-3)

#### **Fine-tuning**

The model is trained via repeated gradient updates using a large corpus of example tasks.



# Language-Model-as-a-Service (LMaaS)



## Language-Model-as-a-Service (LMaaS)

### **GPT-3** Pricing

Ada Fastest Babbage Curie Davin	
	nci Most powerful
\$0.0008/1K tokens \$0.0012/1K tokens \$0.0060/1K tokens \$0.0	<b>0600</b> /1K tokens

## Language-Model-as-a-Service (LMaaS)

Describe a layout.	Products	Ne	W		See all –
Just describe any layout you want, and it'll try to render below!		Rec	ently added GPT-3 apps		
A div that contains 3 buttons each with a random color.	Select product	~			
	Collections		Customer Service ActiveChat.ai	Chatbots Al Buddy	Humor Al Guru
	New				
	Popular				
	Upcoming		LegalTech	Chatbots	Developer Tools
	Requested		aiLawDocs	AskBrian AskBrian	Azure OpenAl Service
	Categories				
	All	319	Botto Generative Art	Image captioning	Healthcare
	A/B Testing	2	Dotto	enperap	culu
Equation description	Ad Generation	3			
x squared plus two times x	AI Copywriting	37	Code Generation	Summarization	API Design
Translate	AI Writing Assistants	1	DeepGenX	Delv Al	Design an API with
	API Design	1			
x + 2x	Avatars	1	Descution		Deserve Assistants
	Blog writing	2	Drafted	Duolingo	Elicit
	Book Writing	1			

### https://gpt3demo.com/

### LMaaS in China



✓ <u>快速写诗 API 文档</u>, 100次/天。

### However...

The performance of manual prompt and in-context learning highly depend on the choice of prompt and demonstrations, and lags far behind model tuning.



Zhao et al. Calibrate Before Use: Improving Few-Shot Performance of Language Models. ICML 2021



### Performance is the Key for grounding. (Who are the users?)

# To make LLMs benefit more people...

Can we optimize the prompt with the API feedback? (without expensive backpropagation)

Objective:

$$\mathbf{p}^{\star} = \arg\min_{\mathbf{p}\in\mathcal{P}} \mathcal{L}(f(\mathbf{p};\tilde{X}),\tilde{Y})$$



# A challenge of high dimensionality

Considering optimization of the continuous prompt, the dimensionality can be tens of thousands (say we are going to optimize 50 prompt tokens, each with 1k dimensions, there are 50k parameters to be optimized.)

Derivative-free optimization (DFO) can struggle with high-dimensional problems, **except for** the case when the problem has a low intrinsic dimensionality.

Note: Intrinsic dimensionality is the minimal number of parameters needed to represent the problem

# A challenge of high dimensionality

An example

- The objective to be optimized has two dimensions but only one matters
- In that case we can perform optimization with random embedding



Wang et al. Bayesian optimization in a billion dimensions via random embeddings. J. Artif. Intell. Res. 2016

### Fortunately...

LLMs have a very low intrinsic dimensionality!



Aghajanyan et al. Intrinsic dimensionality explains the effectiveness of language model fine-tuning. ACL 2021

## **Black-Box Tuning**



## **Black-Box Tuning**

### The CMA-ES (Covariance Matrix Adaptation Evolution Strategy)

The CMA-ES (Evolution Strategy with Covariance Matrix Adaptation)

Consider  $P^{(t)} = \mathcal{N}\left(\boldsymbol{\mu}^{(t)}, \ \sigma^{(t)^2} \boldsymbol{C}^{(t)}\right)$  where  $\boldsymbol{\mu}^{(t)} \in \mathbb{R}^n$ ,  $\sigma^{(t)} \in \mathbb{R}_+$ ,  $\boldsymbol{C}^{(t)} \in \mathbb{R}^{n \times n}$ 

- $\mu^{(t)} \rightarrow \mu^{(t+1)}$ : Maximum likelihood update, i.e.  $P(\boldsymbol{x}_{\text{selected}}^{(t)} | \boldsymbol{\mu}^{(t+1)}) \rightarrow \max$
- $C^{(t)} \rightarrow C^{(t+1)}$ : Maximum likelihood update, i.e.  $P(\frac{\boldsymbol{x}_{\text{selected}}^{(t)} \boldsymbol{\mu}^{(t)}}{\sigma^{(t)}} | C^{(t+1)}) \rightarrow \max$ , under consideration of prior  $C^{(t)}$  (otherwise  $C^{(t+1)}$  becomes singular).
- $\sigma^{(t)} \rightarrow \sigma^{(t+1)}$ : Update to achieve conjugate perpendicularity, i.e. conceptually  $(\boldsymbol{\mu}^{(t+2)} - \boldsymbol{\mu}^{(t+1)})^{\mathrm{T}} \boldsymbol{C}^{(t)^{-1}} (\boldsymbol{\mu}^{(t+1)} - \boldsymbol{\mu}^{(t)}) / \sigma^{(t+1)^2} \rightarrow 0$

https://cma-es.github.io/

### Experiments

### Datasets and processing details

Category	Dataset	$\mid \mathcal{Y} \mid$	Train	Test	Туре	Template	Label words
	SST-2	2	67k	0.9k	sentiment	$\langle S \rangle$ . It was [MASK].	great, bad
	Yelp P.	2	560k	38k	sentiment	$\langle S \rangle$ . It was [MASK].	great, bad
single-	AG's News	4	120k	7.6k	topic	[MASK] News: $\langle S  angle$	World, Sports, Business, Tech
sentence	DBPedia	14	560k	70k	topic	[Category: [MASK]] $\langle S \rangle$	Company, Education, Artist, Athlete, Office,
							Transportation, Building, Natural, Village,
							Animal, Plant, Album, Film, Written
contonco	MRPC	2	3.7k	0.4k	paraphrase	$\langle S_1  angle$ ? [MASK], $\langle S_2  angle$	Yes, No
semence-	RTE	2	2.5k	0.3k	NLI	$\langle S_1  angle$ ? [MASK], $\langle S_2  angle$	Yes, No
pair	SNLI	3	549k	9.8k	NLI	$\langle S_1  angle$ ? [MASK], $\langle S_2  angle$	Yes, Maybe, No

### Experiments

### 16-shot (per class) learning with RoBERTa-large (350M)

Method	SST-2 acc	Yelp P. acc	AG's News acc	DBPedia acc	MRPC F1	SNLI acc	RTE acc	Avg.
Gradient-Based Methods								
Prompt Tuning	$68.23 \pm 3.78$	$61.02 \pm 6.65$	$84.81 \pm 0.66$	87.75 ±1.48	51.61 ±8.67	36.13 ±1.51	54.69 ±3.79	63.46
+ Pre-trained prompt	/	/	/	/	$77.48 \pm 4.85$	$64.55 \pm 2.43$	$77.13 \pm 0.83$	74.42
P-Tuning v2	$64.33 \pm 3.05$	$92.63 \pm 1.39$	$83.46 \pm 1.01$	$97.05 \pm 0.41$	$68.14 \pm 3.89$	$36.89 \pm 0.79$	$50.78 \pm 2.28$	70.47
Model Tuning	$85.39 \pm 2.84$	$91.82 \pm 0.79$	$86.36 \pm 1.85$	$97.98 \pm 0.14$	$77.35 \pm 5.70$	$54.64 \pm 5.29$	$58.60 \pm 6.21$	78.88
			Gradient-Fre	ee Methods				
Manual Prompt	79.82	89.65	76.96	41.33	67.40	31.11	51.62	62.56
In-Context Learning	$79.79 \pm 3.06$	$85.38 \pm 3.92$	$62.21 \pm 13.46$	$34.83 \pm 7.59$	$45.81 \pm \! 6.67$	$47.11 \pm 0.63$	$60.36 \pm 1.56$	59.36
Feature-MLP	$64.80 \pm 1.78$	$79.20 \pm 2.26$	$70.77 \pm 0.67$	$87.78 \pm 0.61$	$68.40 \pm 0.86$	$42.01 \pm 0.33$	$53.43 \pm 1.57$	66.63
Feature-BiLSTM	$65.95 \pm 0.99$	$74.68 \pm 0.10$	$77.28 \pm 2.83$	$90.37 \pm 3.10$	$71.55 \pm 7.10$	$46.02 \pm 0.38$	$52.17 \pm 0.25$	68.29
Black-Box Tuning	$89.56 \pm 0.25$	$91.50 \pm 0.16$	$81.51 \pm 0.79$	$87.80 \pm 1.53$	$61.56 \pm 4.34$	$46.58 \pm 1.33$	$52.59 \pm 2.21$	73.01
+ Pre-trained prompt	/	/	/	/	$75.51 \pm 5.54$	$83.83 \pm 0.21$	$77.62 \pm 1.30$	83.90

### Experiments

### Detailed comparison on SST-2 and AG News

	Deployment-	As-A-	Test	Training	Memory	Footprint	Upload	Download
	Efficient	Service	Accuracy	Time	User	Server	per query	per query
			SST-2 (max	sequence length: 47)	)			
Prompt Tuning	$\checkmark$	×	72.6	15.9 mins	-	5.3 GB	-	-
Model Tuning	×	×	87.8	9.8 mins	-	7.3 GB	-	-
Feature-MLP	$\checkmark$	$\checkmark$	63.8	7.0 mins	20 MB	2.8 GB	4 KB	128 KB
Feature-BiLSTM	$\checkmark$	$\checkmark$	66.2	9.3 mins	410 MB	2.8 GB	4 KB	6016 KB
Black-Box Tuning	$\checkmark$	$\checkmark$	89.4	10.1 (6.1*) mins	30 MB	3.0 GB	6 KB	0.25 KB
		AC	G's News (ma	x sequence length: 1	07)			
Prompt Tuning	$\checkmark$	×	84.0	30.2 mins	-	7.7 GB	-	-
Model Tuning	×	×	88.4	13.1 mins	-	7.3 GB	-	-
Feature-MLP	$\checkmark$	$\checkmark$	71.0	13.5 mins	20 MB	3.6 GB	20 KB	256 KB
Feature-BiLSTM	$\checkmark$	$\checkmark$	73.1	19.7 mins	500 MB	3.6 GB	20 KB	27392 KB
Black-Box Tuning	$\checkmark$	$\checkmark$	82.6	21.0 (17.7*) mins	30 MB	4.6 GB	22 KB	1 KB

# Forward Is All You Need?

Limitations of black-box tuning:

- Slow convergence on many-label classification (e.g., DBPedia)
- Requirement of prompt pre-training (gradient) on difficult tasks (e.g., SNLI)

### Current version of black-box tuning is just a lower bound:

- Prompt/verbalizer engineering, prompt ensemble, prompt pre-training...
- Better derivative-free algorithms
- Pre-trained random embedding

### Can We Go Deeper?

### The ``Deep Prompt Tuning``



Prefix Tuning (Li and Liang, ACL 2021)

P-Tuning v2 (Liu et al., ACL 2022)

## Can We Go Deeper?

The challenge, again, is the high dimensionality

- Say we are going to optimize 50 prompt tokens at each layer of RoBERTalarge, each with 1k dimensions, there are 50k×24=1.2M parameters to be optimized
- Besides, the prompt parameters at different layers are heterogenous and therefore we can not simply use the random embedding to solve it

## Take A Closer Look Into the Forward Pass

Thanks to the residual connections in modern LLMs, the forward computation can be decomposed as an additive form

An example of a 3-layer model:

$$egin{aligned} f(\mathbf{x}_1) &= f_3(\mathbf{x}_3) + \mathbf{x}_3 \ &= f_3(\mathbf{x}_3) + f_2(\mathbf{x}_2) + \mathbf{x}_2 \ &= f_3(\mathbf{x}_3) + f_2(\mathbf{x}_2) + f_1(\mathbf{x}_1) + \mathbf{x}_1 \end{aligned}$$

Therefore, the optimization can be decomposed into multiple sub-problems!

### Take A Closer Look Into the Forward Pass

A general formulation of ``deep black-box tuning``:

$$f(\mathbf{x}_1, \mathbf{p}) = [\mathbf{A}_1 \mathbf{z}_1 + \mathbf{p}_1^0; \mathbf{x}_1] + \sum_{j=1}^L f_j([\mathbf{A}_j \mathbf{z}_j + \mathbf{p}_j^0; \mathbf{x}_j])$$

Given such an additive form, we propose a divide-and-conquer (DC) algorithm to alternately optimize prompt at each layer

# **Divide-and-Conquer**

- Layer-specific optimizer
- Layer-specific random projection
- Alternate from the bottom to top

**Algorithm 1:** DC Algorithm for BBTv2 **Require:** *L*-layer PTM Inference API *f*, Loss function  $\mathcal{L}$ , Budget of API calls  $\mathcal{B}$ , Derivative-free optimizers  $\{\mathcal{M}_j\}_{j=1}^L$ 1: Initialize random projections  $A_1, \ldots, A_L$ 2: Initialize parameters  $\mathbf{z}_1^{(0)}, \ldots, \mathbf{z}_L^{(0)}$ 3: Deep prompts  $\mathbf{p} = \langle \mathbf{A}_1 \mathbf{z}_1^{(0)}, \dots, \mathbf{A}_L \mathbf{z}_I^{(0)} \rangle$ 4: for i = 1 to  $\mathcal{B}/L$  do for j = 1 to L do 5: Evaluate:  $loss = \mathcal{L}(f(\mathbf{p}))$ 6: Update:  $\mathbf{z}_{i}^{(i)} \leftarrow \mathcal{M}_{j}(\mathbf{z}_{i}^{(i-1)}, loss)$ 7: Replace:  $\mathbf{p}_i \leftarrow \mathbf{A}_i \mathbf{z}_i^{(i)}$ 8: 9: end for 10: **end for** 11: return Optimized deep prompts p

# Revisiting Random Projection (Embedding)

Generating random projections from a normal distribution with std dev as

$$\sigma_A = \frac{\hat{\sigma}}{\sqrt{d}\sigma_z}$$

A visualization of generated prompt with RoBERTa-large



## **BBTv2: Towards A Gradient-Free Future**

Main improvements of BBTv2

- Get rid of prompt pre-training
- Improved random projection
- Deep prompts

### **BBTv2: Towards A Gradient-Free Future**

### Main improvements of BBTv2



Comparable to full model tuning but merely tuning ~10k parameters



Improve BBT on entailment tasks

- Be comparable to full model tuning without pre-trained prompt embedding



Improve BBT on many-label classification tasks

- Faster convergence than BBT on DBPedia (14 classes)



### Generalization across LMs



### **Overall comparison**

Method	Tunable Params	SST-2 acc	Yelp P. acc	AG's News acc	DBPedia acc	MRPC F1	SNLI acc	RTE acc	Avg.
Gradient-Based Methods									
Model Tuning	355M	85.39 ±2.84	$91.82 \pm 0.79$	$86.36 \pm 1.85$	$97.98 \pm 0.14$	<b>77.35</b> ±5.70	54.64 ±5.29	<b>58.60</b> ±6.21	78.88
Adapter	2.4M	$\overline{83.91}$ ±2.90	$\overline{90.99} \pm 2.86$	$86.01 \pm 2.18$	<b>97.99</b> ±0.07	$69.20 \pm 3.58$	$\underline{57.46} \pm 6.63$	$48.62 \pm 4.74$	76.31
BitFit	172K	$81.19 \pm \! 6.08$	$88.63 \pm \! 6.69$	$\underline{86.83} \pm 0.62$	$94.42 \pm 0.94$	$66.26 \pm \! 6.81$	$53.42 \pm 10.63$	$52.59 \pm \hspace{-0.5mm} 5.31$	74.76
LoRA	786K	88.49 ±2.90	$90.21 \pm 4.00$	$\textbf{87.09} \pm 0.85$	$97.86 \pm 0.17$	$\underline{72.14} \pm 2.23$	$61.03 \pm 8.55$	$49.22 \pm \hspace{-0.5mm} \pm \hspace{-0.5mm} 5.12$	<u>78.01</u>
Prompt Tuning	50K	$68.23 \pm 3.78$	$61.02 \pm \! 6.65$	$84.81 \pm 0.66$	$87.75 \ \pm 1.48$	$51.61 \pm 8.67$	$36.13 \pm 1.51$	$\underline{54.69} \pm 3.79$	63.46
P-Tuning v2	1.2M	$64.33 \pm 3.05$	$\textbf{92.63} \pm 1.39$	$83.46 \pm 1.01$	$97.05 \pm 0.41$	$68.14 \pm 3.89$	$36.89 \pm 0.79$	$50.78 \pm 2.28$	70.47
			Gradie	ent-Free Metho	ds				
Manual Prompt	0	79.82	89.65	76.96	41.33	67.40	31.11	51.62	62.56
In-Context Learning	0	$79.79 \pm 3.06$	$85.38 \pm 3.92$	$62.21 \pm 13.46$	$34.83 \pm 7.59$	$45.81 \pm \! 6.67$	$\underline{47.11} \pm 0.63$	$\underline{60.36} \pm 1.56$	59.36
Feature-MLP	1M	$64.80 \pm 1.78$	$79.20 \pm 2.26$	$70.77 \pm 0.67$	$87.78 \ {\pm}0.61$	$68.40 \pm 0.86$	$42.01 \pm 0.33$	$53.43 \pm 1.57$	66.63
Feature-BiLSTM	17M	$65.95 \pm 0.99$	$74.68 \pm 0.10$	$77.28 \pm 2.83$	$\underline{90.37} \pm 3.10$	$\underline{71.55} \pm 7.10$	$46.02 \pm 0.38$	$52.17 \ {\pm}0.25$	68.29
BBT	500	$\underline{89.56} \pm 0.25$	$91.50 \pm 0.16$	$\underline{81.51} \pm 0.79$	$79.99^{+\pm2.95}$	$61.56 \pm 4.34$	$46.58 \pm 1.33$	$52.59 \pm 2.21$	<u>71.90</u>
BBTv2	12K	$\textbf{90.41} \pm 0.71$	$\underline{90.69} \pm 0.66$	$\textbf{85.06} \pm 0.49$	$\textbf{92.59} \pm 0.17$	$\textbf{78.15} \pm 2.00$	$\textbf{61.50} \pm 1.28$	<b>60.56</b> ±5.09	79.85

### Versatility across different language models

LM	Method	d SST-2 AG's N		DBPedia				
Encoder-only PTMs								
DEDT	BBT	$76.26 \pm 2.64$	$76.67 \pm 1.12$	$89.58 \pm 0.51$				
DEKI	BBTv2	$79.32 \pm 0.29$	$79.58 \pm 1.15$	$93.74 \pm 0.50$				
DODEDTO	BBT	$89.56 \pm 0.25$	$81.51 \ {\pm}0.79$	$79.99 \pm 2.95$				
ROBERIA	BBTv2	$90.41 \pm 0.71$	$85.06 \pm 0.49$	$92.59 \pm 0.17$				
	Decoder-only PTMs							
CDT 2	BBT	$75.53 \pm 1.98$	$77.63 \pm 1.89$	$77.46 \pm 0.69$				
GP1-2	BBTv2	$80.13 \pm 3.28$	$82.18 \pm 1.07$	$91.36 \pm 0.73$				
Encoder-Decoder PTMs								
	BBT	$77.87 \pm 2.57$	$77.70 \pm 2.46$	$79.64 \pm 1.55$				
BAKI	BBTv2	$89.53 \ {\pm}2.02$	$81.30 \pm 2.58$	$87.10 \pm 2.01$				
Т5	BBT	$89.15 \pm 2.01$	$83.98 \pm 1.87$	$92.76 \pm 0.83$				
13	BBTv2	$91.08 \pm 1.49$	$84.32 \pm 1.29$	$92.76 \pm 0.85$				

### Comparison on CPM-2 (11B)

Method	Tunable Params	ChnSent acc	LCQMC acc
Model Tuning	11 <b>B</b>	$\underline{86.1} \pm 1.8$	$\underline{58.8} \pm 1.8$
Vanilla PT	410K	$62.1 \pm 3.1$	$51.5 \pm 3.4$
Hybrid PT	410K	$79.2 ~ {\pm}4.0$	$54.6 \pm 2.3$
LM Adaption	410K	$74.3 \pm \! 5.2$	$51.4 \pm 2.9$
BBTv2	4.8K	$\textbf{86.4} \pm 0.8$	<b>59.1</b> ±2.5

The power of scale (with T5)

- Outperform gradient descent when model size becomes large



### Other Solutions for LMaaS



## Other Solutions for LMaaS

- Text prompt: Manually or automatically design task-specific text prompts
- In-context learning: Include a few examples in the input at inference time
- **Black-box optimization**: Tuning a small portion of parameters with only the access of the LLM's output probability via black-box optimization (BBO)
- **Feature-based learning**: LLMs can serve as a feature extractor, on which users can build some lightweight learnable model to solve the task
- **Data generation**: Use LLMs to generate a dataset of labeled text pairs, which is then used to locally train a much smaller model

### **Check Our Paper List!**

#### $\equiv$ README.md

### Language Model as a Service (LMaaS)

#### last commit last wednesday PaperNumber 52

This is a curated list of "Language-Model-as-a-Service (LMaaS)" papers, which is mainly maintained by Tianxiang Sun. We strongly encourage the NLP researchers who are interested in this topic to make pull request to add or update the papers (See Contributing). Watch this repository for the latest updates!

### Updates

- 2022/7/7: Write a blog (in Chinese)
- 2022/7/4: Create this paper list

### Contents

- Introduction
  - Scope
  - Advantages
- Keywords
- Papers
  - Text Prompt
  - In-Context Learning
  - Black-Box Optimization
  - Feature-based Learning
  - Data Generation
- Contributing

□ Readme
≥ 816 KB
№ MIT license
☆ 164 stars
⊙ 5 watching
♀ 15 forks

0

#### Releases

No releases published Create a new release

#### Packages

No packages published Publish your first package

Contributors 7



### Resources

- LMaaS paper list: <u>https://github.com/txsun1997/LMaaS-Papers</u>
- Code of BBT and BBTv2: <u>https://github.com/txsun1997/Black-Box-Tuning</u>
- BBT paper (ICML 2022): <u>https://arxiv.org/abs/2201.03514</u>
- BBTv2 paper: <u>https://arxiv.org/abs/2205.11200</u>
- Blog for BBT (in Chinese): <u>https://zhuanlan.zhihu.com/p/455915295</u>
- Blog for LMaaS (in Chinese): <u>https://zhuanlan.zhihu.com/p/538857729</u>

Feel free to reach out if you have any questions or suggestions about our papers, code, or the paper list!





# Thanks!

### Tianxiang Sun

School of Computer Science, Fudan University

https://txsun1997.github.io/