

---

# Text-to-LoRA: Instant Transformer Adaption

---

Rujikorn Charakorn<sup>1</sup> Edoardo Cetin<sup>1</sup> Yujin Tang<sup>1</sup> Robert T. Lange<sup>1</sup>

## Abstract

While Foundation Models provide a general tool for rapid content creation, they regularly require task-specific adaptation. Traditionally, this exercise involves careful curation of datasets and repeated fine-tuning of the underlying model. Fine-tuning techniques enable practitioners to adapt foundation models for many new applications but require expensive and lengthy training while being notably sensitive to hyperparameter choices. To overcome these limitations, we introduce Text-to-LoRA (T2L), a model capable of adapting Large Language Models *on the fly* solely based on a natural language description of the target task. T2L is a hypernetwork trained to construct LoRAs in a single inexpensive forward pass. After training T2L on a suite of 9 pre-trained LoRA adapters (GSM8K, Arc, etc.), we show that the ad-hoc reconstructed LoRA instances match the performance of task-specific adapters across the corresponding test sets. Furthermore, T2L can compress hundreds of LoRA instances and zero-shot generalize to entirely unseen tasks. This approach provides a significant step towards democratizing the specialization of foundation models and enables language-based adaptation with minimal compute requirements. Our code is available at <https://github.com/SakanaAI/text-to-lora>.

## 1. Introduction

Biological systems are capable of rapid adaptation, given limited sensory cues. For example, the human visual system can tune its light sensitivity and focus through neuromodulation of the fovea and rod cells (Wurtz et al., 2011; Digre & Brennan, 2012). While recent large language models (LLMs) exhibit a wide variety of capabilities and knowledge,

they remain rigid when adding task-specific capabilities. In such cases, practitioners often resort to re-training parts of the model (Gururangan et al., 2020; Wei et al., 2021; Dettmers et al., 2022; Tay et al., 2021) using parameter-efficient fine-tuning techniques, e.g., Low-Rank Adaptation (LoRA, Hu et al., 2022). Typically, a LoRA adapter has to be optimized for each downstream task and requires task-specific dataset and hyperparameter setting. This fine-tuning scheme for adaptation significantly limits the possibility of transferring knowledge between tasks and induces engineering overhead. Recently, it has been observed that by inducing structural constraints, the low-rank matrices learned by LoRA adapters can be further compressed. For example, one can train *lossy* versions of the original adapter while maintaining downstream performance (Brüel-Gabrielsson et al., 2024; Kim et al., 2024; Kopiczko et al., 2024). Furthermore, multiple LoRAs can be combined for new tasks at inference time (Ostapenko et al., 2024). At the core of these approaches lies the explicit use of decomposition or dimensionality reduction techniques (e.g., SVD or routing) for better compression and online composition of existing LoRAs. This raises the following questions:

1. Can we end-to-end train a neural network to compress many pre-trained LoRAs?
2. Can we decode new task-specific LoRA adapters solely based on natural-language instructions for an unseen task at test time?

We hypothesize that different LoRA adapters share the same underlying adaptation mechanism and can be optimized simultaneously without any explicit structure or recipe for combining them. To explicitly test this hypothesis, we propose T2L, a hypernetwork (Ha et al., 2016) that compresses task-specific LoRAs and generates new LoRA adapters *zero-shot* at inference time. T2L is trained to compress LoRAs on a diverse task distribution from the Super Natural Instructions (SNI) dataset (Wang et al., 2022). Importantly, T2L takes a natural language description of the target task as an input, allowing zero-shot LoRA generation to unseen tasks. Empirically, we show that T2L can effectively be trained either to reconstruct pre-trained adapters or via supervised fine-tuning on a distribution of downstream tasks (see Figure 1, top right). After training, T2L outperforms a multi-

<sup>1</sup>Sakana AI. Correspondence to: Rujikorn Charakorn <rujjikorn@sakana.ai>, Robert T. Lange <robert@sakana.ai>.

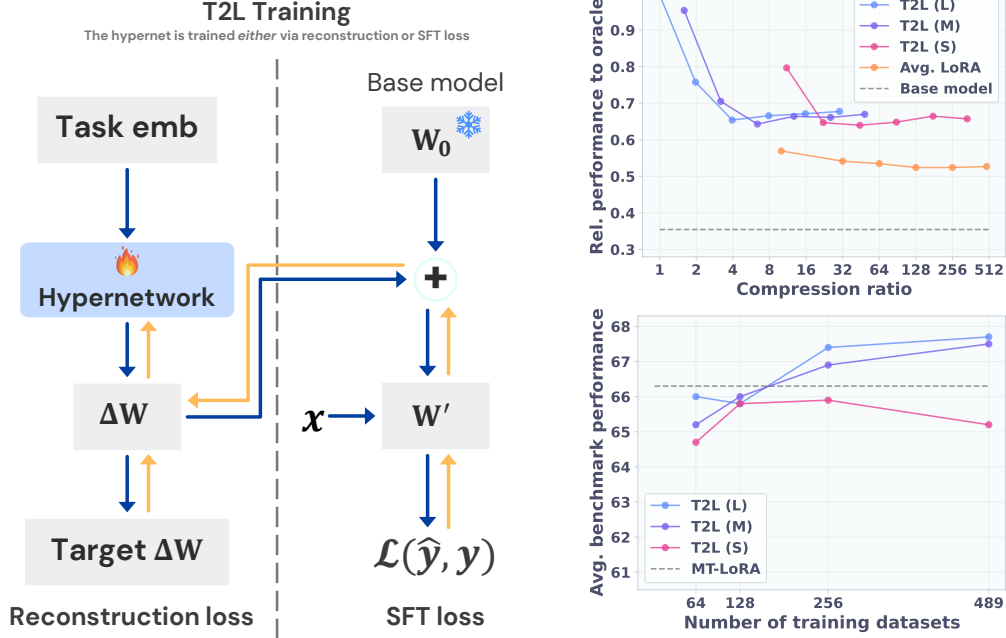


Figure 1: **Left:** Conceptual overview of **T2L**’s training routine. Given a set of task description embeddings, we train a hypernetwork to generate LoRA adaptation matrices ( $\Delta W$ ) for various tasks. The weights of **T2L** are *either* optimized to distill pre-trained LoRA weights or via multi-task supervised fine-tuning on downstream tasks. **Right, Top:** Relative performance to the oracles on training SNI tasks with varying compression ratios. **Right, Bottom:** Zero-shot LoRA generation performance on 10 benchmark tasks. As we increase the number of pre-training datasets, the performance of **T2L** increases for 3 different **T2L** architectures.

task LoRA baseline and Arrow Routing (Ostapenko et al., 2024), a state-of-the-art zero-shot LoRA routing method, on various benchmark tasks. Furthermore, we show that **T2L** can generate LoRA adapters for previously unseen tasks solely using the language-based task description. This result highlights the generalization capabilities and applicability of our proposed indirect adaptation encoding. Our contributions are summarized as follows:

1. We introduce hypernetwork-based architectures for producing LoRA adapters with a single forward pass (Section 3) based on text descriptions. **T2L** architectures can be trained using both distillation of pre-trained adapters and supervised multi-task fine-tuning.
2. We show that **T2L** can efficiently encode hundreds of LoRA adapters (Section 4). While the compression is lossy, **T2L** maintains the performance of task-specifically tuned LoRA adapters. Furthermore, **T2L** can generalize to unseen tasks given suitable natural language descriptions of the tasks.
3. We provide rigorous ablations (Section 5) including **T2L** scaling with datasets (see Figure 1, bottom right),

the impact of different task description embeddings, the training routines, and text-based task descriptions.

4. Finally, we analyze the nature of **T2L** generations. We find semantically meaningful LoRA clusters when visualizing the generated LoRAs in a dimensionality-reduced space (Section 5.5). Furthermore, we study the relationship between LoRA adapters and find compelling evidence why reconstruction-trained **T2L** cannot generalize (Appendix D).

## 2. Preliminaries

We utilize multiple fine-tuning datasets  $\mathcal{D} = \{\mathcal{D}^1, \dots, \mathcal{D}^T\}$ , which correspond to different tasks  $\mathcal{T} = \{t^1, \dots, t^T\}$ . For the purpose of training **T2L**, we assume that each fine-tuning dataset has a set of natural language *task descriptions* ( $Z^i = \{z_1^i, \dots, z_m^i\}$ ):  $\mathcal{D}^i = \{X^i, Y^i, Z^i\}$ . The task descriptions do not need to be specific to each sample but rather a general description of the dataset. For a single task  $t^i$ , the fine-tuning objective of an LLM with pre-trained weights ( $\Psi$ ) is given by

$$\Delta W^i = \underset{\Delta W^i}{\operatorname{argmin}} \mathcal{L}_{\text{SFT}}(\mathcal{D}^i, \Psi, \Delta W^i), \quad (1)$$

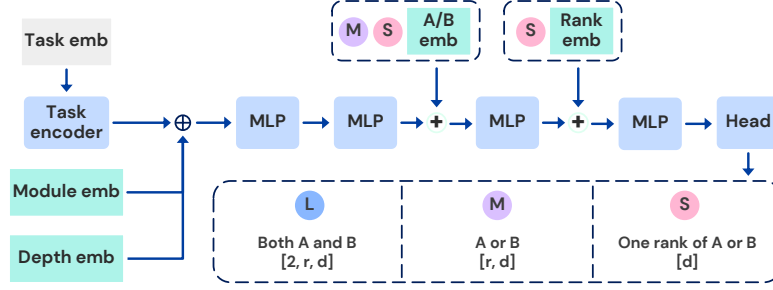


Figure 2: Overview of T2L architectural variations. The dashed box at the bottom shows the output size of a single forward pass of T2L. Blue boxes are trainable modules. Cyan boxes are trainable embedding layers. Components in dashed boxes are only used with their corresponding architectures.  $r$  is the rank of a LoRA adapter and  $d$  is the size of the input and the output dimension.

where  $\mathcal{L}_{\text{SFT}}$  gives the supervised fine-tuning loss and  $\Delta W^i$  is the fine-tuning adaption for task  $t^i$  to the base weights. For the *multi-task* setting, we train a single adapter  $\Delta W$  to minimize the expected loss over the union of all datasets  $\mathcal{D}$ :

$$\Delta W = \underset{\Delta W}{\operatorname{argmin}} \mathbb{E}_{\mathcal{D}^i \sim \mathcal{D}} \mathcal{L}_{\text{SFT}}(\mathcal{D}^i, \Psi, \Delta W). \quad (2)$$

**Low-Rank Adaptation (LoRA, Hu et al., 2022):** LoRA is a parameter-efficient fine-tuning method that freezes the pre-trained weights of a base model and only learns low-rank weight matrices, which serve as an adapter to the base model. For each selected linear transformation  $h = W_0 x$ , the fine-tuned transformation is given by  $h = W_0 x + \Delta W x = W_0 x + B^T A x$ , where  $A, B \in \mathbb{R}^{r \times d}$  are weight matrices of rank  $r < d$ . We omit the layer index and module type of the LoRA weights when referring to all LoRA weights. Otherwise, we use subscripts to represent the layer index and module type, e.g.,  $\Delta W_{m,l}$ , where  $m$  is the module type (e.g., query projection) and  $l$  is the layer index.

**Hypernetworks:** A hypernetwork is a neural network that generates parameters for another ‘base’ network (Ha et al., 2016). It serves as an indirect encoding (Schmidhuber, 1997; Stanley & Miikkulainen, 2003; Zhang et al., 2018; Schug et al., 2024) of the base network, given that the parameter count of the hypernetwork is much smaller. This compression is achieved by learning to share parameters indirectly. More specifically, given a layer-specific descriptor vector  $\phi_l$ , a hypernetwork with parameters  $\theta$  generates the parameters of the base model at layer  $l \in \{1, \dots, L\}$  as follows:  $W_l = h_\theta(\phi_l)$ . Traditionally, the layer descriptors are either one-hot or learned vectors. The weights  $\theta$  are then trained via end-to-end optimization on a downstream task.

### 3. Text-to-LoRA: Learning to Compress and Generate LoRAs

In this work, we utilize a hypernetwork to generate LoRA adapters for task-specific adaptation. For each target module ( $m$ ) and layer index ( $l$ ), a hypernetwork generates the two

low-rank matrices  $A, B$  based on a task description  $z^i \in Z^i$  of a task  $t^i$  as follows:

$$\Delta W_{m,l}^i = h_\theta(\phi_{m,l}^i), \text{ with} \quad (3)$$

$$\phi_{m,l}^i = \text{concat} [f(z^i), E[m], E[l]], \quad (4)$$

where  $f$  gives a vector representation of a text description, typically represented by a CLS token of a bidirectional transformer model or last token activation of an LLM.  $E$  is a learnable embedding dictionary indexed by either a module type  $m$  or a layer index  $l$ . For legibility, we introduce a shorthand notation for T2L’s output  $\Delta W^i := h_\theta(\phi^i) := h_\theta(\{\phi_{m,l}^i\})$ . Then, a supervised fine-tuning training objective for T2L is

$$\theta = \underset{\theta}{\operatorname{argmin}} \mathbb{E}_{\mathcal{D}^i \sim \mathcal{D}, z^i \in Z^i} \mathcal{L}_{\text{SFT}}(\mathcal{D}^i, \Psi, h_\theta(\phi^i)), \quad (5)$$

Note that values of  $m$  and  $l$  can be batched, which allows T2L to generate  $\Delta W$  for all modules and layer indices efficiently within a single forward pass.

#### 3.1. Text-to-LoRA Architectures

Most of a hypernetwork’s parameters come from the output layer, which scales linearly with the size of the target weights (Von Oswald et al., 2019). To explore the complexity-performance trade-off, we propose three variants of T2L: **L**, **M**, and **S**. We impose different output spaces on the hypernetwork that represent different inductive biases and parameter counts (see Figure 2). We note that all variants use the same backbone architecture and only differ in their output heads and learnable embeddings. The **L architecture** is the largest variant. Its final linear layer outputs low-rank  $A$  and  $B$  matrices simultaneously, with the number of weight connections to the output head  $|\theta_{\text{head}}| = d_{\text{out}} \times 2 \times r \times d$ , where  $d_{\text{out}}$  is the output size of the last MLP block. **M architecture** is the medium-sized model with a shared output layer between the low-rank  $A$  and  $B$  matrices. That is, the head outputs a low-rank

matrix, either  $A$  or  $B$ , depending on the learnable embedding. The size of the output head is  $|\theta_{\text{head}}| = d_{\text{out}} \times r \times d$ . Finally, **S architecture** is the most parameter-efficient model with the strongest inductive biases, where the hypernetwork outputs only one rank of a low-rank matrix at a time. This output space makes the size of the head much smaller:  $|\theta_{\text{head}}| = d_{\text{emb}} \times d$ . For reference, a LoRA adapter has  $r \times d \times 2 \times L \times |M|$  trainable parameters, where  $L$  is the number of layers and  $|M|$  is the number of target modules. The default value of  $d_{\text{out}}$  is 512. We note that every architecture can generate all the low-rank matrices  $A$  and  $B$  in a single forward pass by batching all the input embeddings. We provide more details of the architectures in Appendix F and the weight initialization method that leads to stable training in Appendix G.

### 3.2. Training Text-to-LoRA via LoRA Reconstruction

The most straightforward way to train **T2L** is to reconstruct pre-trained task-specific LoRAs. This setup allows us to utilize publicly available libraries of LoRAs (Brüel-Gabrielsson et al., 2024; Zhao et al., 2024). Alternatively, one can also use a two-stage procedure, in which a library of LoRAs is pre-trained in the first stage and then train **T2L** to reconstruct them. For the sole purpose of compressing LoRAs, we can train **T2L** using one-hot or learnable vectors as task embeddings. However, these embeddings do not allow zero-shot LoRA generation for unseen tasks. To enable zero-shot LoRA generation, we additionally condition **T2L** with embeddings of natural language task descriptions, which allows **T2L** to generate LoRA adapters for various tasks—including unseen ones—given corresponding task descriptions. Given a suitable library of LoRA adapters  $\Omega$ , the reconstruction loss for **T2L** can be written as

$$\mathcal{L}(\Omega, \theta) = \mathbb{E}_{\Delta W^i \sim \Omega} |\Delta W^i - h_{\theta}(\phi^i)|. \quad (6)$$

### 3.3. Training Text-to-LoRA via Supervised Fine-Tuning

Alternatively, **T2L** can be directly optimized on fine-tuning datasets. Training **T2L** with SFT sidesteps the need for intermediate target LoRA adapters and allows for end-to-end training. This training scheme is preferred if existing trained LoRAs are not naturally clustered by their functionalities or downstream tasks. For instance,  $t^1$  and  $t^2$  could be two related tasks requiring a similar LLM capability, but  $\Delta W^1$  and  $\Delta W^2$  could be in different minima. Thus, **T2L** trained via reconstruction training would have to compress numerically different  $\Delta W^1$  and  $\Delta W^2$ , making it less likely to generalize. In fact, we empirically find that a **T2L** trained via reconstruction fails to generalize to unseen tasks (Section 5.4). In contrast, an SFT-trained **T2L** can implicitly learn to cluster tasks, which has been shown to improve zero-shot LoRA routing performance (Ostapenko et al., 2024). The SFT loss for **T2L** is given by Equation (5).

## 4. Experiments

We investigate the effectiveness of the different **T2L** architectures and training schemes in terms of the compression of adapters (Section 4.1) and zero-shot LoRA generation for unseen tasks (Section 4.2). As baselines, we consider task-specific LoRAs, element-wise averaged LoRA, and multi-task LoRA—a LoRA adapter trained on all training tasks. We also implement Hyperdecoders (Iverson & Peters, 2022)—a hypernetwork that generates LoRAs on a per-sequence basis—based on our proposed architectures. To boost the performance of the base models without fine-tuning, we utilize few-shot in-context learning (ICL, Brown et al., 2020; Dong et al., 2024) and task description pre-pending, i.e., providing task description at the beginning of each query. Additionally, we include results of Arrow Routing zero-shot performance from Ostapenko et al. (2024). Note that the performance can only be compared indirectly as it uses a different set of LoRA adapters and training tasks. Furthermore, there are likely differences in the benchmark evaluation prompts.

In most experiments, we use Mistral-7B-Instruct (Jiang et al., 2023) as the base LLM model except in Tables 7 and 8 where Llama-3.1-8B-Instruct and Gemma-2-2b-Instruct are used as the base models, respectively. We use gte-large-en-v1.5 (Li et al., 2023; Zhang et al., 2024) for extracting the task embedding from a natural language task description. All LoRA adapters are of rank 8 and only target the query and the value projection modules in every attention block of the base LLM (totaling 3.4M parameters). With this LoRA configuration, **L**, **M**, and **S** have 55M, 34M, and 5M trainable parameters respectively.

We utilize the SNI dataset (Wang et al., 2022) for training LoRA adapters. We use a subset of 500 tasks following Brüel-Gabrielsson et al. (2024). We use 11 tasks for hold-out validation and removed 10 datasets due to data contamination from the evaluation benchmark tasks, leaving 479 datasets for training. All samples are in English. More details of the datasets can be found in Appendix J. For evaluation, we choose 10 widely used benchmarks that collectively cover a variety of LLM capability assessments, e.g., reasoning, math, science, coding, and world knowledge. Specifically, we include the following benchmarks: Arc-challenge (ArcC) and Arc-easy (ArcE) (Clark et al., 2018), BoolQ (Clark et al., 2019), GSM8K (Cobbe et al., 2021), Hellaswag (HS) (Zellers et al., 2019), OpenBookQA (OQA) (Mihaylov et al., 2018), PIQA (Bisk et al., 2020), Winogrande (WG) (Keisuke et al., 2019), HumanEval (HE) (Chen et al., 2021), and MBPP (Austin et al., 2021).<sup>1</sup> Task

<sup>1</sup>The benchmark tasks share some similarities with the training tasks. Specifically, they are mostly multiple-choice question-answering tasks. Also, there are similar and overlapping domains



Table 1: Benchmark performance of **T2L** trained via reconstruction loss on 9 benchmark tasks. **Green highlight** indicates that **T2L** outperforms the benchmark-specific LoRA adapters.

	ArcC (acc)	ArcE (acc)	BQ (acc)	GSM8K (acc)	HS (acc)	OQA (acc)	PIQA (acc)	WG (acc)	MBPP (pass@1)	Avg. (9 tasks)
Base model	65.4	77.8	71.6	40.9	49.7	54.2	72.8	45.0	43.1	55.8
<b>One-Hot Task E.</b>										
<b>T2L</b> (Recon) <b>L</b>	76.4	89.9	89.4	<b>53.8</b>	92.6	85.0	69.7	<b>51.2</b>	<b>52.6</b>	<b>73.4</b>
<b>T2L</b> (Recon) <b>M</b>	<b>76.7</b>	89.9	89.4	53.2	92.6	85.0	69.9	<b>51.4</b>	<b>52.9</b>	<b>73.4</b>
<b>T2L</b> (Recon) <b>S</b>	75.2	88.8	87.4	50.9	89.1	75.6	<b>83.9</b>	<b>58.1</b>	48.1	73.0
<b>Task Description E.</b>										
<b>T2L</b> (Recon) <b>L</b>	76.6	89.8	89.4	<b>53.9</b>	92.6	85.0	69.6	<b>51.2</b>	51.8	73.3
<b>T2L</b> (Recon) <b>M</b>	76.5	89.9	89.4	<b>53.9</b>	92.5	84.9	<b>70.4</b>	<b>51.6</b>	<b>52.8</b>	<b>73.5</b>
<b>T2L</b> (Recon) <b>S</b>	75.4	88.8	87.8	49.1	89.7	76.7	<b>84.2</b>	<b>56.9</b>	48.0	73.0
Task-specific LoRAs	76.6	89.9	89.4	53.5	92.6	85.0	69.9	51.1	52.1	73.3

descriptions for the training datasets and the benchmarks are fully generated, as described in Appendix L. When we use a language task embedding as a part of the input, we average **T2L** performance using three descriptions for each benchmark.

#### 4.1. LoRA Compression

In this experiment, we aim to investigate whether **T2L** can recover the performance of trained LoRAs via reconstruction training. For quality control and consistent evaluation, we train a task-specific LoRA (oracle) on the training split of each benchmark task, collectively forming a library of LoRAs. Table 1 shows the benchmark performance of **T2L** trained by distilling 9 benchmark-specific LoRAs using either one-hot or natural language task embeddings from gte-large-en-v1.5. We note that the benchmark tasks are indirectly seen during training by **T2L**, as it learns to distill benchmark-specific LoRAs. We can see that **T2L** fully recovers the performance of the oracle adapters with both task embedding types. Notably, **T2L** outperforms task-specific LoRAs on several benchmarks (highlighted in green). We hypothesize that the gain comes from the lossy compression of the target LoRAs, which acts as a regularization on the already trained LoRA weights. This effect is most apparent on PIQA and WG benchmarks, where the oracle LoRA overfits and performs worse than the base model.

Next, we explore whether **T2L** conditioned on one-hot task vectors can maintain the oracle single-task LoRAs’ performance when using an increasing number of training tasks. Figure 3 shows the performance of one-hot **T2L** on the test splits of a subset of 10 SNI training tasks with varying

between the two splits. For example, the ARC benchmarks are similar to SNI task #47. However, some benchmarks are very different from the training distribution, e.g., MBPP and HumanEval, as the training tasks do not contain any code generation tasks.



Figure 3: Relative performance and training reconstruction error of **T2L** instances trained with an increasing number of tasks ({16, 32, 64, 128, 256, 479} tasks from left to right).

degrees of final average training L1 reconstruction error. We train various **T2L** instances for each architecture using {16, 32, 64, 128, 256, 479} training tasks, leading to an effective increase in the training reconstruction error. Although **T2L** fully recovers the oracles’ performance when the reconstruction loss is less than  $10^{-4}$ , the performance drops as the training error increases. This result suggests that **T2L** learns a lossy compression of the target LoRAs. Still, we find that all **T2L** architectures can maintain around 65% of oracles’ performance, and the performance does not drop further even at  $> 8 \times 10^{-4}$  per-element L1 error. Despite the performance drop, we show that increasing the number of training tasks is beneficial in the SFT setup, increasing zero-shot benchmark performance of **T2L** in Section 5.1.

#### 4.2. Zero-Shot LoRA Generation

Here, we explore whether **T2L** can generate useful LoRA adapters for unseen tasks. We train **T2L** with SFT on 479 SNI tasks, each with 128 task descriptions. For each data point in a training minibatch, we sample a description from the corresponding dataset in an online fashion. Table 2

Table 2: Zero-shot performance on unseen benchmark tasks. SFT-trained T2L generates LoRAs based on unseen task descriptions. Its performance is an average of three generated LoRAs, each with a different instance of task descriptions. Arrow Routing results are taken from Ostapenko et al. (2024). **Green highlight** indicates higher performance than that of the benchmark-specific LoRA adapters. **Bold numbers** are used when the performance is higher than the multi-task LoRA.

	ArcC (acc)	ArcE (acc)	BQ (acc)	HS (acc)	OQA (acc)	PIQA (acc)	WG (acc)	MBPP (pass@1)	Avg. (8 tasks)	GSM8K (acc)	HE (pass@1)	Avg. (10 tasks)
<b>No Test-Time Adaptation</b>												
Mistral-7B-Instruct	65.4	77.8	71.6	49.7	54.2	72.8	45.0	43.1	60.0	40.9	37.2	55.8
Prepending task desc.	72.0	85.8	67.6	58.9	63.4	77.9	59.0	41.6	65.8	40.9	39.0	60.6
3-shot ICL	72.1	85.9	71.7	59.0	66.2	76.2	58.0	42.6	66.5	40.9	37.2	61.0
Average LoRA	70.7	84.4	75.4	59.9	59.0	78.0	54.3	47.1	66.1	42.4	37.8	60.9
Multi-task LoRA	76.2	88.3	85.5	65.2	68.0	81.8	62.4	48.1	71.9	47.5	39.6	66.3
<b>Zero-Shot Adaptation</b>												
Arrow Routing	60.9	86.2	<b>87.6</b>	<b>80.8</b>	48.6	<b>83.0</b>	<b>68.5</b>	<b>50.2</b>	70.7	N/A	28.7	N/A
Hyperdecoders (per-instance)	<b>76.6</b>	<b>88.5</b>	83.9	65.2	<b>76.6</b>	<b>81.3</b>	<b>64.9</b>	<b>51.6</b>	<b>73.6</b>	43.6	<b>40.9</b>	<b>67.3</b>
T2L (SFT) <b>S</b>	76.0	<b>88.7</b>	83.8	<b>68.0</b>	<b>71.6</b>	<b>82.3</b>	<b>61.0</b>	41.2	71.6	47.3	39.0	65.9
T2L (SFT) <b>M</b>	<b>77.2</b>	<b>89.0</b>	84.3	65.1	<b>76.1</b>	<b>81.8</b>	<b>64.0</b>	<b>50.5</b>	<b>73.5</b>	45.2	<b>41.3</b>	<b>67.5</b>
T2L (SFT) <b>L</b>	<b>77.5</b>	<b>88.9</b>	85.0	<b>66.5</b>	<b>75.5</b>	<b>82.1</b>	<b>64.2</b>	<b>51.9</b>	<b>73.9</b>	45.8	39.2	<b>67.7</b>
<b>Oracle</b>												
Task-specific LoRAs	76.6	89.9	89.4	92.6	85.0	69.9	51.1	52.1	75.8	53.5	N/A	N/A

shows the zero-shot performance on 10 benchmark tasks. Here, we present the best model of each variant from our scaling experiment in Section 5.1. We observe that a multi-task LoRA adapter performs well on the benchmarks despite no additional fine-tuning. Still, there is a performance gap between task-specific LoRAs and MT LoRA. We observe that SFT-trained T2L indeed generates useful LoRAs, thus improving over the multi-task LoRA adapter consistently and across benchmarks (indicated by bold numbers). Notably, even though T2L cannot fully bridge the performance gap with task-specific LoRAs, it outperforms the oracles on a subset of tasks (highlighted in green). We further investigate the generality of our proposed method with different base models including Llama (Dubey et al., 2024) and Gemma (Team et al., 2024) models in Appendix A. We note that one of the main advantages of T2L is its efficiency. To emphasize T2L’s efficiency, we provide an ad-hoc FLOPs analysis in Appendix I.

## 5. Ablations and Analyses

### 5.1. Increasing Training Compute Proportional to the Number of Training Tasks

In this section, we explore the scalability of T2L by varying the training tasks and scale the training budget proportionally to the dataset size on all variants. Table 3 shows that, after increasing the number of training tasks and compute budget, T2L generally benefits from the additional training tasks. However, **S** does not benefit from extended training with 479 tasks, potentially due to its limited model capacity. We additionally investigate the effect of the task diversity on the robustness of T2L by training on more tasks without

scaling the training budget in Appendix C. We find that it is crucial to scale the compute budget according to the number of training tasks. For instance, **M** with scaled compute budget improves over training runs with a fixed budget when using 256 or more training tasks.

### 5.2. Task Embedding Models

Table 4 shows the zero-shot benchmark performance with two different embedding models: gte-large-en-v1.5 and Mistral-7B-Instruct. For the gte model, we extract a task description by using the activation of the CLS token in the last layer as the model is a bidirectional model. For Mistral, we use the activation of the last token in the sequence to represent a given description (BehnamGhader et al., 2024). Table 4 shows the results with the two embedding models used for T2L SFT training on 128 tasks. Both embedding models yield T2L instances with comparable generalization capability, suggesting T2L’s robustness to task description embedding methods.

### 5.3. Varying Task Descriptions

We investigate the impact of task descriptions on the performance of generated LoRAs using four types of descriptions:

- **Train:** Training descriptions of corresponding tasks.
- **Eval:** Unseen descriptions of corresponding tasks.
- **Random strings:** Random literal strings.
- **Train (random):** Training descriptions randomly sampled from other tasks.

Table 3: Performance of SFT-trained T2L with varying numbers of training tasks.

	Number of tasks	Max SGD steps	ArcC (acc)	ArcE (acc)	BQ (acc)	GSM8K (acc)	HS (acc)	OQA (acc)	PIQA (acc)	WG (acc)	HE (pass@1)	MBPP (pass@1)	Avg.
T2L (SFT) <b>L</b>	479	1M	77.5	88.9	85.0	45.8	66.5	75.5	82.1	64.2	39.2	51.9	67.7 <span style="color: green;">▲</span>
	256	640K	77.3	88.1	84.3	46.0	64.5	75.7	81.9	64.0	39.8	52.1	67.4 <span style="color: green;">▲</span>
	128	320K	76.6	88.4	85.2	46.1	67.0	74.3	81.6	55.0	38.2	45.7	65.8 <span style="color: red;">▼</span>
	64	160K	75.5	88.0	84.5	43.9	65.5	70.7	80.5	59.5	39.8	51.7	66.0
T2L (SFT) <b>M</b>	479	1M	77.2	89.0	84.3	45.2	65.1	76.1	81.8	64.0	41.3	50.5	67.5 <span style="color: green;">▲</span>
	256	640K	75.9	89.3	85.0	47.0	65.3	73.7	81.6	63.2	39.8	48.6	66.9 <span style="color: green;">▲</span>
	128	320K	74.9	88.3	85.5	44.9	64.8	72.8	80.7	61.6	42.9	43.5	66.0 <span style="color: green;">▲</span>
	64	160K	73.6	87.7	84.5	43.2	64.6	70.5	79.9	56.0	40.7	51.4	65.2
T2L (SFT) <b>S</b>	479	1M	77.7	88.3	85.0	46.3	65.3	73.9	82.4	61.9	34.6	36.6	65.2 <span style="color: red;">▼</span>
	256	640K	76.0	88.7	83.8	47.3	68.0	71.6	82.3	61.0	39.0	41.2	65.9 <span style="color: green;">▲</span>
	128	320K	74.9	88.0	84.5	44.4	66.2	72.2	82.0	59.3	39.0	47.3	65.8 <span style="color: green;">▲</span>
	64	160K	75.4	88.4	85.0	43.1	64.8	70.7	81.5	51.6	39.4	46.7	64.7

Table 4: Zero-shot benchmark performance of SFT T2L trained on 128 tasks using different text embedding models.

	gte			Mistral		
Avg. Benchmark performance	<b>S</b>	<b>M</b>	<b>L</b>	<b>S</b>	<b>M</b>	<b>L</b>
	65.8	66.0	65.8	64.7	66.2	66.0
<b>Avg.</b>	<b>65.9</b>			<b>65.6</b>		

Table 5: T2L trained via reconstruction on 9 tasks performs well when given aligned task descriptions. Unaligned descriptions produce lower benchmark performance.

	Aligned		Unaligned	
	Train	Eval	Train (random)	Random strings
T2L <b>L</b>	73.3	73.6	49.1	68.2
T2L <b>M</b>	73.5	70.2	49.5	68.5
T2L <b>S</b>	73.0	72.9	55.7	53.9
<b>Avg.</b>	<b>73.3</b>	<b>72.2</b>	<b>51.4</b>	<b>63.5</b>

For each description type, we use the gte-large-en-v1.5 embedding and report the average performance using three descriptions. The four types can be grouped into two categories based on the alignment between the descriptions and the tasks: aligned (**Train, Eval**) and unaligned (**Train (random)** and **Random strings**). Note that we use reconstruction-trained T2L in this experiment. That is, the hypernetwork has seen training descriptions of the benchmarks during training. We observe a performance gap between the two description categories. Specifically, training and evaluation descriptions generate the best performing LoRAs, matching the performance of oracle LoRAs, despite the evaluation descriptions being unseen. These results suggest that T2L is robust to changes in the task description as long as the descriptions are aligned with the task. On the other hand, if the descriptions are not aligned with the task at hand, the generated LoRAs will not perform as well, as indicated by the performance

Table 6: Reconstruction vs SFT training scheme.

Benchmark performance	Recon			SFT		
	<b>S</b>	<b>M</b>	<b>L</b>	<b>S</b>	<b>M</b>	<b>L</b>
	61.8	61.7	62.0	64.8	66.5	67.5
<b>Avg.</b>	<b>61.8</b>			<b>66.3</b>		

of the unaligned group. We believe that using an LLM for adjusting the description alignment could effectively sidestep this failure case of T2L. Additionally, we provide a qualitative result demonstrating steerability and an unsuccessful example of T2L in Figure 4. Importantly, the last two examples in Figure 4 (iii, iv) are both correct but have different answer styles thanks to different descriptions. We remark that Hyperdecoders (Iverson & Peters, 2022) cannot exhibit such steerability as it uses the problem instance as the input to the hypernetwork.

#### 5.4. Training Schemes

In this section, we investigate the zero-shot performance of SFT-trained and reconstruction-trained T2L. All model instances are trained with roughly equal wall-clock time of 10 hours (see Appendix H for details). From Table 6, we can see a clear performance gap between reconstruction and SFT training schemes. Specifically, SFT produces T2L instances that perform significantly better than those trained via reconstruction (66.3 vs 61.83 benchmark performance averaged over model architectures). We attribute the performance difference to the library of LoRAs needed for reconstruction training. For reconstruction-trained T2L to generalize, the target LoRA adapters of similar tasks should be clustered in some latent manifold. In contrast, SFT training does not need pre-trained task-specific LoRA adapters, thus sidestepping this challenge via end-to-end learning. In Appendix D, we show that pre-trained adapters for similar tasks do not live nearby in the weight space, supporting our claim of a

**User Input**

Toulouse has twice as many sheep as Charleston. Charleston has 4 times as many sheep as Seattle.  
How many sheep do Toulouse, Charleston, and Seattle have together if Seattle has 20 sheep?

**(i) Base model response**

Great question! Let's break it down step by step:

1. We know that Seattle has 20 sheep.
2. Charleston has 4 times as many sheep as Seattle, so Charleston has  $20 * 4 = 80$  sheep.
3. Toulouse has twice as many sheep as Charleston, so Toulouse has  $80 * 2 = 160$  sheep.
4. To find the total number of sheep in all three cities, we add the number of sheep in each city:  $160$  (Toulouse) +  $80$  (Charleston) +  $20$  (Seattle) =  $240$  sheep.

So, Toulouse, Charleston, and Seattle together have **240 sheep**.

**(ii) Generate LoRA using the following (low-quality) description:  
"solve this task please"**

Charleston has 4 times as many sheep as Seattle, and Seattle has 20 sheep, so Charleston has  $20 * 4 = 80$  sheep.  
Toulouse has twice as many sheep as Charleston, so Toulouse has  $80 * 2 = 160$  sheep.  
Toulouse, Charleston, and Seattle have sheep together:  $160 + 80 + 20 = 240$  sheep.

**(iii) Generate LoRA using the following description:**

*"This task challenges your problem-solving abilities through mathematical reasoning. You must carefully read each scenario and systematically work through the data to compute the final outcome."*

Toulouse has twice as many sheep as Charleston, so Toulouse has  $2 * \text{Charleston} = 2 * 4 = 8$  times as many sheep as Seattle.  
Seattle has 20 sheep, so Toulouse has  $8 * 20 = 160$  sheep.  
Charleston has 4 times as many sheep as Seattle, so Charleston has  $4 * 20 = 80$  sheep.  
Toulouse, Charleston, and Seattle have sheep together =  $160$  (Toulouse) +  $80$  (Charleston) +  $20$  (Seattle) = **260 sheep**.

**(iv) Generate LoRA using the following description:**

*"Use your programming skill to think algorithmically, logically and systematically. Carefully read each scenario and systematically work through the data to compute the final outcome. Use your math skill to calculate correctly."*

let me check the relationships between the number of sheep in each city:

1. Toulouse has twice as many sheep as Charleston.
2. Charleston has 4 times as many sheep as Seattle.

From point 2, we know that Charleston has  $4 * 20 = 80$  sheep. From point 1, Toulouse has  $2 * 80 = 160$  sheep.  
So, Toulouse, Charleston, and Seattle have a total of  $160 + 80 + 20 = 260$  sheep.

Figure 4: Qualitative examples of responses from applying LoRA generated by T2L to the *Mistral-7B-Instruct* base model on a GSM8K problem instance. (i) The response from the base model is incorrect. (ii) Applying a LoRA generated from a *low-quality* task description does not make the model output the correct response. (iii, iv) Descriptions that are aligned with the problem lead to generated LoRAs that steer the base model to output correct responses. Descriptions from (iii) and (iv) influence the model to generate different reasoning paths, highlighting the steerability of T2L.

potential problem when reconstructing pre-trained LoRAs.

### 5.5. Visualization of T2L Activations

Next, we aim to understand T2L further and see whether it generates task-specific LoRA adapters for unseen tasks with unseen descriptions. We probe SFT T2L M trained on 256 training tasks in the zero-shot evaluation setting. The model is probed on all the benchmark tasks, each with three unseen descriptions. Figure 5 shows the 2D t-SNE projection of T2L's task encoder activations and the outputs of the last MLP block. We can see clear clustering in both projection plots based on the tasks (colors and shapes). T2L generates different adapters for different tasks, confirming that T2L indeed performs task-specific adaptation 'on the fly'. Moreover, similar tasks, e.g., MBPP and HumanEval, are clustered together in both plots, suggesting that SFT-trained T2L produces similar adapters for semantically similar tasks.

## 6. Related Work

**Hypertextworks for Adaptation:** Hypertextworks (Ha et al., 2016) provide a general indirect encoding method for neural network weights. They have been applied to different architectures (e.g., in attention, Schug et al., 2024) and training paradigms (e.g., in continual learning, Von Oswald et al., 2019). Here, we focus on generating low-rank adapters using natural language instructions. Previous work (Mahabadi et al., 2021; He et al., 2022; Ortiz-Barajas et al., 2024) considers hypertextworks for LLM adaptation in a multi-task context but only uses learned task identifiers instead of natural language for adaptation. Thus, these approaches do not enable task-wise zero-shot generalization.

**Hypertextworks for Zero-Shot LLM Adaptation:** Xiao et al. (2023) explore the use of hypertextworks on a limited set of English dialects; they only consider five dialects, one of which is unseen. Furthermore, the hypertextwork relies on an expert-based transformation of the dialects, limiting the possibility for generalization. Mu et al. (2024) propose






Figure 5: 2D t-SNE projection of activations of **T2L**’s task encoder (left) and activations of the last MLP block (right) grouped by benchmark tasks (represented by colors). We probe **T2L** with unseen three task descriptions per benchmark. We can see activations clustering in both plots, indicating that **T2L** indeed learns to generate LoRAs tailored to specific tasks.

Gisting, a method that learns to compress an in-context task description to prefix tokens, allowing the language model to follow instructions with fewer tokens. However, Gisting is limited to prefix tokens—only influencing the attention matrices of the base model. Thus, prefix tokens are less flexible compared to LoRAs that can modify different parts of LLMs, e.g., attention blocks. Hyperdecoders (Iverson & Peters, 2022) is a hypernetwork that generates adapters on the fly based on the input sequence. While per-sequence adaptation is desirable for benchmark evaluation—where the LLM should always output the correct answer—we argue that description-based adaptation gives more control to users since they can steer the LLM in creative ways based on user-generated descriptions (see Figure 4). Furthermore, the generated adapters cannot be efficiently fused into the base model, leading to significant overhead for each query.

Closely related to our work are HyperTuning (Phang et al., 2023), HNET-LM (Deb et al., 2022), and HINT (Iverson et al., 2023). Differing from prior work that heavily focuses on pre-trained encoder-decoder models, e.g., T5 (Raffel et al., 2020) or BART (Lewis, 2019), we use frontier instruction fine-tuned models as the base models, i.e., Mistral, Llama, Gemma. Also, prior work typically relies on initializing a part of their hypernetworks from the base model (e.g., tying task encoder’s weights to the base model) to achieve good performance or stable training as opposed to ours that are task-embedder agnostic and can freely change the task embedding model (Section 5.2). Additionally, our work utilizes generated descriptions instead of the ones provided by the SNI dataset. In Appendix B, we show that using generated descriptions increase the performance of **T2L** considerably. Overall, our work improves upon prior work in several ways, including achieving task-wise zero-shot generalization on various frontier instruction-tuned language models, simpler and more general hypernetwork

input requirements, investigation of training regimes, and more comprehensive experiments, ablations, and analyses.

Concurrent to our work, Lv et al. (2024) propose a similar approach that utilizes a hypernetwork to generate LoRA adapters at inference time. However, their hypernetwork assumes that the context vector provided to the hypernetwork contains few-shot examples. In contrast, **T2L** only assumes a task description, which users can produce themselves within seconds.

## 7. Discussion and Limitations

**Discussion.** We rely on generated descriptions from GPT-4o mini to ensure high-quality and consistent task descriptions. It is plausible that when **T2L** is deployed in real-world scenarios, users might not input high-quality descriptions, which could cause performance degradation of generated adapters. Our results have primarily focused on LLM adaptation. However, **T2L** can be directly applied to other LLMs or to adapt vision language models. Finally, the potential for **T2L** trained on a smaller base model to transfer effectively to larger models within the same architecture class remains an open area for exploration.

**Limitations.** We only consider LoRA as the output space of the hypernetwork. We believe there are more efficient ways to modulate LLMs given a text description, e.g., directly modulating the activations of the base model. Also, we believe the compression achieved by **T2L** can be further optimized using well-designed inductive biases. Finally, although **T2L** exhibits robustness and signs of scalability, it still does not reach the benchmark performance of task-specific LoRAs in a zero-shot manner. Achieving such potent zero-shot adaption is still one of the biggest challenges for **T2L**.

## Acknowledgment

We thank David Ha for suggesting “Text-to-LoRA” as the title of the paper. We thank anonymous reviewers for their constructive feedback, which we incorporate to improve the quality of the paper.

## Impact Statement

This paper introduces Text-to-LoRA (T2L), a novel approach that significantly lowers the barrier to adapting large foundation models for specific tasks. Traditionally, customizing models like LLMs requires resource-intensive fine-tuning on specific datasets for each new application, limiting accessibility and slowing down deployment. T2L overcomes this by training a hypernetwork to generate task-specific Low-Rank Adapters (LoRAs) instantly, using only a natural language description of the target task as input. This eliminates the need for per-task fine-tuning datasets and lengthy optimization processes, enabling rapid, on-the-fly adaptation with minimal computational overhead during inference, thereby making powerful model customization more accessible.

The broader impact of T2L lies in its potential to democratize the specialization of powerful AI systems by enabling adaptation through intuitive text instructions. While T2L demonstrates effective compression and promising zero-shot generalization to unseen tasks similar to those encountered during training, potential pitfalls exist that warrant consideration. Its performance is notably sensitive to the quality and clarity of the natural language task descriptions; poorly phrased or misaligned instructions could lead to suboptimal or incorrect adaptations, potentially hindering reliability in real-world user scenarios. Furthermore, while T2L significantly advances instant adaptation, its generalization capability to task types fundamentally different from its training distribution (e.g., beyond the SNI-derived benchmarks) needs further investigation, and it may not yet fully match the performance ceiling of adapters meticulously fine-tuned on extensive, high-quality datasets for highly complex or specialized domains.

## References

- Achiam, J., Adler, S., Agarwal, S., Ahmad, L., Akkaya, I., Aleman, F. L., Almeida, D., Altenschmidt, J., Altman, S., Anadkat, S., et al. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023.
- Austin, J., Odena, A., Nye, M., Bosma, M., Michalewski, H., Dohan, D., Jiang, E., Cai, C., Terry, M., Le, Q., et al. Program synthesis with large language models. *arXiv preprint arXiv:2108.07732*, 2021.
- Beck, J., Jackson, M. T., Vuorio, R., and Whiteson, S. Hypernetworks in meta-reinforcement learning. In *Conference on Robot Learning*, pp. 1478–1487. PMLR, 2023.
- BehnamGhader, P., Adlakha, V., Mosbach, M., Bahdanau, D., Chapados, N., and Reddy, S. LLM2Vec: Large language models are secretly powerful text encoders. In *First Conference on Language Modeling*, 2024. URL <https://openreview.net/forum?id=IW1PR7vEBf>.
- Bisk, Y., Zellers, R., Bras, R. L., Gao, J., and Choi, Y. Piqa: Reasoning about physical commonsense in natural language. In *Thirty-Fourth AAAI Conference on Artificial Intelligence*, 2020.
- Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J. D., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., Agarwal, S., Herbert-Voss, A., Krueger, G., Henighan, T., Child, R., Ramesh, A., Ziegler, D., Wu, J., Winter, C., Hesse, C., Chen, M., Sigler, E., Litwin, M., Gray, S., Chess, B., Clark, J., Berner, C., McCandlish, S., Radford, A., Sutskever, I., and Amodei, D. Language models are few-shot learners. In Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M., and Lin, H. (eds.), *Advances in Neural Information Processing Systems*, volume 33, pp. 1877–1901. Curran Associates, Inc., 2020. URL [https://proceedings.neurips.cc/paper\\_files/paper/2020/file/1457c0d6bfc4967418bfb8ac142f64a-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2020/file/1457c0d6bfc4967418bfb8ac142f64a-Paper.pdf).
- Brüel-Gabrielsson, R., Zhu, J., Bhardwaj, O., Choshen, L., Greenewald, K., Yurochkin, M., and Solomon, J. Compress then serve: Serving thousands of lora adapters with little overhead. *arXiv preprint arXiv:2407.00066*, 2024.
- Chen, M., Tworek, J., Jun, H., Yuan, Q., de Oliveira Pinto, H. P., Kaplan, J., Edwards, H., Burda, Y., Joseph, N., Brockman, G., Ray, A., Puri, R., Krueger, G., Petrov, M., Khlaaf, H., Sastry, G., Mishkin, P., Chan, B., Gray, S., Ryder, N., Pavlov, M., Power, A., Kaiser, L., Bavarian, M., Winter, C., Tillet, P., Such, F. P., Cummings, D., Plappert, M., Chantzis, F., Barnes, E., Herbert-Voss, A., Guss, W. H., Nichol, A., Paino, A., Tezak, N., Tang, J., Babuschkin, I., Balaji, S., Jain, S., Saunders, W., Hesse, C., Carr, A. N., Leike, J., Achiam, J., Misra, V., Morikawa, E., Radford, A., Knight, M., Brundage, M., Murati, M., Mayer, K., Welinder, P., McGrew, B., Amodei, D., McCandlish, S., Sutskever, I., and Zaremba, W. Evaluating large language models trained on code, 2021.
- Clark, C., Lee, K., Chang, M.-W., Kwiatkowski, T., Collins, M., and Toutanova, K. Boolq: Exploring the surprising difficulty of natural yes/no questions. *arXiv preprint arXiv:1905.10044*, 2019.

- Clark, P., Cowhey, I., Etzioni, O., Khot, T., Sabharwal, A., Schoenick, C., and Tafjord, O. Think you have solved question answering? try arc, the ai2 reasoning challenge. *arXiv:1803.05457v1*, 2018.
- Cobbe, K., Kosaraju, V., Bavarian, M., Chen, M., Jun, H., Kaiser, L., Plappert, M., Tworek, J., Hilton, J., Nakano, R., Hesse, C., and Schulman, J. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168*, 2021.
- Deb, B., Awadallah, A. H., and Zheng, G. Boosting natural language generation from instructions with meta-learning. In Goldberg, Y., Kozareva, Z., and Zhang, Y. (eds.), *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pp. 6792–6808, Abu Dhabi, United Arab Emirates, December 2022. Association for Computational Linguistics. doi: 10.18653/v1/2022.emnlp-main.456. URL <https://aclanthology.org/2022.emnlp-main.456>.
- Dettmers, T., Lewis, M., Belkada, Y., and Zettlemoyer, L. Gpt3. int8 (): 8-bit matrix multiplication for transformers at scale. *Advances in Neural Information Processing Systems*, 35:30318–30332, 2022.
- Digre, K. B. and Brennan, K. Shedding light on photophobia. *Journal of Neuro-ophthalmology*, 32(1):68–81, 2012.
- Dong, Q., Li, L., Dai, D., Zheng, C., Ma, J., Li, R., Xia, H., Xu, J., Wu, Z., Chang, B., et al. A survey on in-context learning. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pp. 1107–1128, 2024.
- Dubey, A., Jauhri, A., Pandey, A., Kadian, A., Al-Dahle, A., Letman, A., Mathur, A., Schelten, A., Yang, A., Fan, A., et al. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*, 2024.
- Gururangan, S., Marasović, A., Swayamdipta, S., Lo, K., Beltagy, I., Downey, D., and Smith, N. A. Don’t stop pretraining: Adapt language models to domains and tasks. *arXiv preprint arXiv:2004.10964*, 2020.
- Ha, D., Dai, A., and Le, Q. V. Hypernetworks. *arXiv preprint arXiv:1609.09106*, 2016.
- He, Y., Zheng, S., Tay, Y., Gupta, J., Du, Y., Aribandi, V., Zhao, Z., Li, Y., Chen, Z., Metzler, D., et al. Hyperprompt: Prompt-based task-conditioning of transformers. In *International conference on machine learning*, pp. 8678–8690. PMLR, 2022.
- Hu, E. J., yelong shen, Wallis, P., Allen-Zhu, Z., Li, Y., Wang, S., Wang, L., and Chen, W. LoRA: Low-rank adaptation of large language models. In *International Conference on Learning Representations*, 2022. URL <https://openreview.net/forum?id=nZeVKeeFYf9>.
- Iverson, H. and Peters, M. E. Hyperdecoders: Instance-specific decoders for multi-task nlp. In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pp. 1715–1730, 2022.
- Iverson, H., Bhagia, A., Wang, Y., Hajishirzi, H., and Peters, M. E. Hint: Hypernetwork instruction tuning for efficient zero-and few-shot generalisation. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 11272–11288, 2023.
- Jiang, A. Q., Sablayrolles, A., Mensch, A., Bamford, C., Chaplot, D. S., Casas, D. d. l., Bressand, F., Lengyel, G., Lample, G., Saulnier, L., et al. Mistral 7b. *arXiv preprint arXiv:2310.06825*, 2023.
- Keisuke, S., Ronan, L. B., Chandra, B., and Yejin, C. Winogrande: An adversarial winograd schema challenge at scale. 2019.
- Kim, H., Sasaki, S., Hoshino, S., and Honda, U. A single linear layer yields task-adapted low-rank matrices. *arXiv preprint arXiv:2403.14946*, 2024.
- Kopieczko, D. J., Blankevoort, T., and Asano, Y. M. VeRA: Vector-based random matrix adaptation. In *The Twelfth International Conference on Learning Representations*, 2024. URL <https://openreview.net/forum?id=NjNfLdxr3A>.
- Korthikanti, V. A., Casper, J., Lym, S., McAfee, L., Andersch, M., Shoenybi, M., and Catanzaro, B. Reducing activation recomputation in large transformer models. *Proceedings of Machine Learning and Systems*, 5:341–353, 2023.
- Lewis, M. Bart: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. *arXiv preprint arXiv:1910.13461*, 2019.
- Li, Z., Zhang, X., Zhang, Y., Long, D., Xie, P., and Zhang, M. Towards general text embeddings with multi-stage contrastive learning. *arXiv preprint arXiv:2308.03281*, 2023.
- Liu, J., Xia, C. S., Wang, Y., and Zhang, L. Is your code generated by chatGPT really correct? rigorous evaluation of large language models for code generation. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. URL <https://openreview.net/forum?id=1qv610Cu7>.
- Lv, C., Li, L., Zhang, S., Chen, G., Qi, F., Zhang, N., and Zheng, H.-T. Hyperlora: Efficient cross-task generalization via constrained low-rank adapters generation. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pp. 16376–16393, 2024.

- Mahabadi, R. K., Ruder, S., Dehghani, M., and Henderson, J. Parameter-efficient multi-task fine-tuning for transformers via shared hypernetworks. *arXiv preprint arXiv:2106.04489*, 2021.
- Mihaylov, T., Clark, P., Khot, T., and Sabharwal, A. Can a suit of armor conduct electricity? a new dataset for open book question answering. In *EMNLP*, 2018.
- Mu, J., Li, X., and Goodman, N. Learning to compress prompts with gist tokens. *Advances in Neural Information Processing Systems*, 36, 2024.
- Ortiz-Barajas, J.-G., Gomez-Adorno, H., and Solorio, T. Hyperloader: Integrating hypernetwork-based lora and adapter layers into multi-task transformers for sequence labelling. *arXiv preprint arXiv:2407.01411*, 2024.
- Ostapenko, O., Su, Z., Ponti, E. M., Charlin, L., Roux, N. L., Pereira, M., Caccia, L., and Sordoni, A. Towards modular llms by building and reusing a library of lorae. *arXiv preprint arXiv:2405.11157*, 2024.
- Phang, J., Mao, Y., He, P., and Chen, W. Hypertuning: Toward adapting large language models without back-propagation. In *International Conference on Machine Learning*, pp. 27854–27875. PMLR, 2023.
- Raffel, C., Shazeer, N., Roberts, A., Lee, K., Narang, S., Matena, M., Zhou, Y., Li, W., and Liu, P. J. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of machine learning research*, 21 (140):1–67, 2020.
- Schmidhuber, J. Discovering neural nets with low kolmogorov complexity and high generalization capability. *Neural Networks*, 10(5):857–873, 1997.
- Schug, S., Kobayashi, S., Akram, Y., Sacramento, J., and Pascanu, R. Attention as a hypernetwork. *arXiv preprint arXiv:2406.05816*, 2024.
- Stanley, K. O. and Miikkulainen, R. A taxonomy for artificial embryogeny. *Artificial life*, 9(2):93–130, 2003.
- Tay, Y., Dehghani, M., Rao, J., Fedus, W., Abnar, S., Chung, H. W., Narang, S., Yogatama, D., Vaswani, A., and Metzler, D. Scale efficiently: Insights from pre-training and fine-tuning transformers. *arXiv preprint arXiv:2109.10686*, 2021.
- Team, G., Riviere, M., Pathak, S., Sessa, P. G., Hardin, C., Bhupatiraju, S., Hussenot, L., Mesnard, T., Shahriari, B., Ramé, A., et al. Gemma 2: Improving open language models at a practical size. *arXiv preprint arXiv:2408.00118*, 2024.
- Von Oswald, J., Henning, C., Grewe, B. F., and Sacramento, J. Continual learning with hypernetworks. *arXiv preprint arXiv:1906.00695*, 2019.
- Wang, Y., Mishra, S., Alipoormolabashi, P., Kordi, Y., Mirzaei, A., Naik, A., Ashok, A., Dhanasekaran, A. S., Arunkumar, A., Stap, D., et al. Super-naturalinstructions: Generalization via declarative instructions on 1600+ nlp tasks. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pp. 5085–5109, 2022.
- Wei, J., Bosma, M., Zhao, V. Y., Guu, K., Yu, A. W., Lester, B., Du, N., Dai, A. M., and Le, Q. V. Finetuned language models are zero-shot learners. *arXiv preprint arXiv:2109.01652*, 2021.
- Wurtz, R. H., Joiner, W. M., and Berman, R. A. Neuronal mechanisms for visual stability: progress and problems. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 366(1564):492–503, 2011.
- Xiao, Z., Held, W., Liu, Y., and Yang, D. Task-agnostic low-rank adapters for unseen English dialects. In Bouamor, H., Pino, J., and Bali, K. (eds.), *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pp. 7857–7870, Singapore, December 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.emnlp-main.487. URL <https://aclanthology.org/2023.emnlp-main.487>.
- Zellers, R., Holtzman, A., Bisk, Y., Farhadi, A., and Choi, Y. Hellaswag: Can a machine really finish your sentence? In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, 2019.
- Zhang, C., Ren, M., and Urtasun, R. Graph hypernetworks for neural architecture search. *arXiv preprint arXiv:1810.05749*, 2018.
- Zhang, X., Zhang, Y., Long, D., Xie, W., Dai, Z., Tang, J., Lin, H., Yang, B., Xie, P., Huang, F., et al. mgte: Generalized long-context text representation and reranking models for multilingual text retrieval. *arXiv preprint arXiv:2407.19669*, 2024.
- Zhao, J., Wang, T., Abid, W., Angus, G., Garg, A., Kinnison, J., Sherstinsky, A., Molino, P., Addair, T., and Rishi, D. Lora land: 310 fine-tuned llms that rival gpt-4, a technical report. *arXiv preprint arXiv:2405.00732*, 2024.



## A. Generalization to Llama and Gemma Models

Table 7: Zero-shot performance with Llama-3.1-8B-Instruct as the base language model.

	ArcC (acc)	ArcE (acc)	BQ (acc)	GSM8K (acc)	HS (acc)	OQA (acc)	PIQA (acc)	WG (acc)	HE (pass@1)	MBPP (pass@1)	Avg.
Llama-3.1-8B-Instruct	73.3	90.6	80.4	75.7	66.6	75.4	79.8	55.3	66.5	68.7	73.2
3-shot ICL	80.7	91.9	80.0	75.7	59.3	77.6	80.9	61.3	66.5	70.4	74.4
Prepending task desc.	80.2	92.5	79.9	75.7	69.8	78.4	81.7	62.4	68.3	70.2	75.9
Multi-task LoRA	82.0	92.8	83.3	77.6	70.8	<b>81.8</b>	<b>83.8</b>	<b>60.3</b>	63.4	69.4	76.5
<b>T2L (SFT) L</b>	<b>82.4</b>	<b>92.9</b>	<b>84.4</b>	<b>79.1</b>	<b>72.8</b>	<b>81.8</b>	81.2	60.0	<b>64.6</b>	<b>69.9</b>	<b>76.9</b>

Table 8: Zero-shot performance with Gemma-2-2B-Instruct as the base language model.

	ArcC (acc)	ArcE (acc)	BQ (acc)	GSM8K (acc)	HS (acc)	OQA (acc)	PIQA (acc)	WG (acc)	HE (pass@1)	MBPP (pass@1)	Avg.
Gemma-2-2B-Instruct	73.7	<b>89.9</b>	81.0	55.6	55.2	71.0	71.0	53.8	<b>43.9</b>	12.3	60.7
3-shot ICL	72.4	88.9	<b>82.5</b>	55.6	55.7	72.6	67.6	53.7	<b>43.9</b>	43.1	63.6
Prepending task desc. w/ ICL	72.4	88.9	<b>82.5</b>	55.6	55.7	72.6	67.6	53.7	<b>43.9</b>	43.1	63.6
Multi-task LoRA w/ ICL	73.5	89.4	81.6	<b>57.2</b>	59.5	<b>74.6</b>	69.4	58.1	39.0	50.4	65.2
<b>T2L (SFT) L w/ ICL</b>	<b>74.0</b>	89.8	81.8	55.1	<b>62.5</b>	73.9	<b>75.2</b>	<b>58.7</b>	41.5	<b>51.5</b>	<b>66.4</b>

In this section, we explore the generality of our proposed architectures to different model families and sizes. Tables 7 and 8 show the benchmark performance of **T2L L** compared to various baselines using Llama-3.1-8B-Instruct and Gemma-2-2B-Instruct as the base models, respectively. With Gemma base model, we utilize ICL for all approaches as it drastically improves the performance on the MBPP benchmark. We see that **T2L** consistently outperforms the baselines across all tested models with varying model sizes and architectures. We note that **T2L** are trained with the same set of hyperparameters across base models.

## B. Training Description Sources

Table 9: Performance of SFT-trained **T2L** with two different training description sources.

	ArcC (acc)	ArcE (acc)	BQ (acc)	GSM8K (acc)	HS (acc)	OQA (acc)	PIQA (acc)	WG (acc)	HE (pass@1)	MBPP (pass@1)	Avg.
<b>T2L (SFT) L</b>	77.5	88.9	85.0	45.8	66.5	75.5	82.1	64.2	39.2	51.9	67.7
<b>T2L (SFT) L w/ SNI def.</b>	75.3	87.4	85.0	45.9	63.6	73.5	80.9	61.8	38.2	53.8	66.5

In this experiment, we explore the impact of the sources of the training task descriptions: SNI and chatGPT (Appendix L). Table 9 shows that using task definitions provided by the SNI datasets reduces the zero-shot benchmark performance of **T2L**. As the SNI datasets are crowd-sourced, we hypothesized that the task descriptions might have inconsistent template or varied levels of details. Thus, it is harder for **T2L** to learn and generalize.

## C. Scaling the Number of Training Tasks with Fixed Compute

We study the impact of the number of training tasks on the zero-shot benchmark performance of **T2L** in the SFT setting, where all **T2L** instances are trained for roughly the same number of gradient steps (see details in Appendix H). Overall, we find that increasing the number of training tasks improves the average zero-shot benchmark performance of the hypernetwork (Figure 1 and Table 10). This result hints at the plausible scalability of **T2L** and positive transfer between tasks.



Table 10: Benchmark performance of SFT-trained T2L with varying numbers of training tasks. We show results with {64, 128, 256, 479} tasks.  $\blacktriangle$  ( $\blacktriangledown$ ) indicates increased (decreased) performance compared to the previous increment in the number of training tasks and training budget.

	Number of tasks	ArcC (acc)	ArcE (acc)	BQ (acc)	GSM8K (acc)	HS (acc)	OQA (acc)	PIQA (acc)	WG (acc)	HE (pass@1)	MBPP (pass@1)	Avg.
T2L (SFT) <b>L</b>	479	77.2	89.0	85.0	46.3	66.5	73.6	82.6	61.8	39.2	44.3	66.6 $\blacktriangledown$
	256	76.6	89.1	84.8	47.0	67.7	73.5	82.8	62.4	39.6	51.0	67.5 $\blacktriangle$
	128	76.2	89.0	85.3	46.2	67.9	71.7	82.6	59.9	40.5	51.3	67.0 $\blacktriangle$
	64	75.5	88.0	84.5	43.9	65.5	70.7	80.5	59.5	39.8	51.7	66.0
T2L (SFT) <b>M</b>	479	77.5	89.0	85.0	45.8	66.5	71.9	82.1	61.4	41.3	50.1	67.1 $\blacktriangle$
	256	76.1	88.2	85.3	45.4	65.6	72.7	81.7	62.3	36.8	50.6	66.5 $\blacktriangle$
	128	75.5	87.8	85.3	46.1	66.6	71.6	81.7	62.2	39.8	44.9	66.1 $\blacktriangle$
	64	73.6	87.7	84.5	43.2	64.6	70.5	79.9	56.0	40.7	51.4	65.2
T2L (SFT) <b>S</b>	479	75.8	88.5	83.9	45.6	64.2	71.9	82.3	61.5	36.2	45.0	65.5 $\blacktriangle$
	256	76.1	88.4	83.0	47.3	65.0	71.7	82.5	58.1	36.2	39.1	64.8 $\blacktriangle$
	128	75.6	87.7	84.9	46.5	65.7	72.7	81.0	59.6	39.0	28.1	64.1 $\blacktriangledown$
	64	75.4	88.4	85.0	43.1	64.8	70.7	81.5	51.6	39.4	46.7	64.7

## D. LoRAs of Similar Tasks

Here, we investigate the relationship between LoRA adapters by inspecting their similarity in the parameter space, performance on the benchmarks, and similarity of their description embeddings. To measure adapter similarity, we compute the cosine similarity of the concatenation of flattened low-rank  $A$  and  $B$  matrices of all layers. In the top row of Figure 6, we plot the adapters’ similarity against task description similarity (using the mean embedding of each task). We find no correlation between the cosine similarity of the adapters’ weights (y-axis) and the task embedding similarity (x-axis) indicated by near-zero Pearson correlation coefficients.

In the bottom row of Figure 6, we change the y-axis from adapters’ relative benchmark performance to benchmark-specific adapters. We find a positive correlation between the relative benchmark performance of SNI-trained adapters and the task embedding similarity. That is, adapters perform better on a benchmark if their task descriptions are similar to those of the benchmark. However, despite their similar functionalities, adapters with similar descriptions are not similar in the parameter space. We believe that this relationship has a significant impact on the limited generalization of reconstruction-trained T2L. We further discuss this topic in Appendix K.

## E. Hyperparameter Settings

Table 11: Hyperparameters for training a task-specific LoRA adapter.

Hyperparameters	Task-specific LoRA	T2L (SFT)	T2L (recon)
Batch size	8	8	Number of the target LoRAs
Gradient accumulation steps	1	1	1
Max learning rate	$8 \times 10^{-5}$	$2.5 \times 10^{-5}$	$10^{-3}$
Max gradient norm	1.0	1.0	1.0
NEFTune noise alpha	5.0	5.0	5.0
Warmup fraction	0.1	0.1	0.1
Learning rate scheduler	Linear with warm up	Linear with warm up	Linear with warm up

Table 11 and Listing 1 show the training configuration of all models trained in this work. For LoRA reconstruction training, each prediction target is an entirety of a LoRA adapter. That is, there is a total of 479 training samples for 479 SNI tasks. Thus, we increase the epochs to 100,000 to ensure that T2L converges.



Figure 6: **Top row:** Each plot shows the similarity between a benchmark LoRA adapter and 479 SNI-trained adapters in the weight space (y-axis) against their similarity in the task embedding space (x-axis). **Bottom row:** Each plot shows SNI-trained adapters’ performance relative to a benchmark adapter (y-axis) with the same x-axis. We can see that LoRAs with similar description embeddings to the benchmarks’ perform better in those benchmarks, suggesting their shared functionalities. However, LoRAs with similar functionalities are not nearby in the parameter space.

## F. Additional Details of T2L Architectures

Listings 2 and 3 show the details of the backbone of T2L. Specifically, the size of the module and layer embedding ( $E[m]$  and  $E[l]$ ) is 32D. Together, they form a dictionary of 34 learnable embeddings (32 layers + 2 target modules). The task encoder is a linear layer that takes in a text embedding (1024D for the `gte` embedding and 4096D for `Mistral` embedding) and outputs a 64D vector. The encoded task, module, and layer embedding are concatenated and then fed into `mlp0` followed by a residual MLP block `mlp1`. At this point, for **M** and **S**, we add a 128D A/B embedding to the residual stream. The output is then fed to another residual MLP block `mlp2`. At this point, for **S**, we add a 128D rank embedding to the residual stream. After this, we feed the activation to the last MLP block. The output of the last MLP block is then fed to a linear head, whose output size is as follows:

- **L**:  $2 \times r \times d$  giving both  $A$  and  $B$  matrices
- **M**:  $r \times d$  giving a low-rank matrix  $A$  or  $B$  depending on the A/B embedding
- **S**:  $d$  giving a rank of a low-rank matrix depending on both the A/B embedding and the rank embedding.

For ease of explanation, we assume  $d$  is the same for the input and the output space of a linear transformation. In practice,  $d_{\text{in}} = d_{\text{out}} = 4096$  for `q_proj` module and  $d_{\text{in}} = 4096, d_{\text{out}} = 1024$  for `v_proj` module.  $r = 8$  for all adapters in this work. Finally, we list the number of trainable parameters of each architecture: 55, 252, 992 for **L**, 34, 282, 240 for **M**, 4, 923, 392 for **S**, 3, 407, 872 for LoRA.

## G. T2L Initialization

We use *Bias-HyperInit* (Beck et al., 2023) to initialize **L** T2L. Bias-HyperInit initializes the linear output head of the hypernetwork such that the weights are all zeros and the bias matches the initialization of the underlying layers. In our work, this corresponds to the output bias of the **L** hypernetwork being initialized to  $U(-\frac{1}{d}, \frac{1}{d})$  for the  $A$  head and all zero for the  $B$  head to match the initialization of traditional LoRA. For other architectures, we aim to match the gradient magnitude to

---

```
{
  "alpha_pattern": {},
  "auto_mapping": null,
  "base_model_name_or_path": "models/Mistral-7B-Instruct-v0.2",
  "bias": "none",
  "fan_in_fan_out": false,
  "inference_mode": true,
  "init_lora_weights": true,
  "layer_replication": null,
  "layers_pattern": null,
  "layers_to_transform": null,
  "loftq_config": {},
  "lora_alpha": 16,
  "lora_dropout": 0.05,
  "megatron_config": null,
  "megatron_core": "megatron.core",
  "modules_to_save": null,
  "peft_type": "LORA",
  "r": 8,
  "rank_pattern": {},
  "revision": null,
  "target_modules": [
    "q_proj",
    "v_proj"
  ],
  "task_type": "CAUSAL_LM",
  "use_dora": false,
  "use_rslora": true
}
```

---

Listing 1: The parameter-efficient fine-tuning (PEFT) config for all LoRA adapters.

**L** at the beginning of training. That is, for **M** architecture, we initialize the bias of the output head to  $U(-\frac{1}{\sqrt{2d}}, \frac{1}{\sqrt{2d}})$ . Finally, **S** output bias is initialized to  $U(-\frac{1}{\sqrt{r2d}}, \frac{1}{\sqrt{r2d}})$ . Without this explicit hypernetwork initialization, the training is unstable, and often leads to failed training runs.

## H. Training Details

All models trained in this work fit in a single H100 GPU (80GB of VRAM). Notably, SFT requires much more memory because of the need to backpropagate the gradient through the base LLM. Reconstruction training, on the other hand, should be possible in a modern consumer-grade GPU.

For reconstruction training, we fix the training epochs to be 100K but scale the batch size to match the number of target LoRA adapters. This means the model trains much faster for a lower number of target LoRAs while maintaining the same number of optimizer steps. For reference, training to reconstruct 9 benchmark-specific LoRAs takes around 10 minutes to complete, while training to reconstruct 479 SNI LoRA adapters takes around 10 hours.

For SFT training with fixed compute budget, we aim to keep the number of optimizer steps the same as we do for reconstruction training. However, since we cannot fit all fine-tuning samples, we scale the number of epochs inverse to the number of training tasks.

Additionally, for reconstruction training, instead of predicting the weights directly, **T2L** learns to predict the z-score of a normal distribution of each weight entry in the low-rank **A**, **B** matrices. At test time, the output is multiplied by the standard

deviation of each element before adding to the mean, converting the prediction to the correct scale.

## I. Ad-hoc FLOPs Analysis

Let  $S$  be the sequence length,  $H$  be the hidden size, and  $L$  be the number of layers of a Transformer-based LLM. We use the following equations for computing the matrix multiplications (GEMMs) FLOPs (Korthikanti et al., 2023).

**FLOPs for Self-Attention (per layer):**  $8 \times S \times H^2 + 4 \times H \times S^2$

**FLOPs for FFN (per layer):**  $16 \times S \times H^2$

**Per Transformer Block Total FLOPs:**  $24 \times S \times H^2 + 4 \times H \times S^2$

**Setup for comparison:**

- 3-shot ICL examples are approximately 256 tokens long
- Question instances are approximately 64 tokens long
- Task descriptions are approximately 48 tokens long
- We consider one question instance as the main input to the base model
- We only consider input tokens for the FLOPs calculation
- We use ‘Mistral-7B-Instruct-v0.2’ as the base model ( $S = 256 + 64$  (3-shot ICL + question instance),  $H = 4096$ ,  $L = 32$ )
- When the based model is used with T2L, we do not include 3-shot ICL ( $S = 64$  (question instance),  $H = 4096$ ,  $L = 32$ )
- We use ‘gte-large-en-v1.5’ as the task description encoder ( $S = 48$  (task description),  $H = 1024$ ,  $L = 24$ )
- We use the M hypernetwork architecture detailed in the Appendix F

### I.1. T2L per instance FLOPs

**gte-large-en-v1.5:** FLOPs =  $24 \times (24 \times 48 \times 1024^2 + 4 \times 1024 \times 48^2) = 0.029$  TFLOPs/instance

**Hypernetwork (M):** FLOPs =  $2 \times 1024 \times 64 + 4 \times 4 \times 128 \times 512 + 128 \times 4096 \times 8 = 0.000005$  TFLOPs/instance

**Base LLM w/o ICL:** FLOPs =  $32 \times (24 \times 64 \times 4096^2 + 4 \times 4096 \times 64^2) = 0.827$  TFLOPs/instance

**Total FLOPs** =  $0.029 + 0.000005 + 0.827 = \mathbf{0.856005}$  TFLOPs/instance

### I.2. Base LLM with 3-shot ICL

**Total FLOPs** =  $32 \times (24 \times (256 + 64) \times 4096^2 + 4 \times (4096) \times (256 + 64)^2) = \mathbf{4.177}$  TFLOPs/instance

Based on this calculation, we can see that the adaptation cost of T2L is significantly cheaper than 3-shot ICL—more than 4x FLOPs reduction, saving compute within the first question instance.

## J. Training and Evaluation Datasets

We use 500 SNI datasets publicly available at <https://huggingface.co/Lots-of-LoRAs>. 479 tasks are used for training and the rest for evaluation. Specifically, we use the following evaluation tasks: task\_035, task\_039, task\_1557, task\_202, task\_304, task\_362, task\_614, task\_701, task\_706, task\_710, task\_726. For the in-context learning baseline, we use 3-shot in-context examples taken from the training split of each benchmark except MBPP that has an explicit split for in-context prompting. HumanEval only has the test split, therefore it is always evaluated against in the zero-shot manner.

## Training Tasks

"Lots-of-LoRAs/task742\_lhoestq\_answer\_generation\_frequency",  
 "Lots-of-LoRAs/task1198\_atomic\_classification\_owant",  
 "Lots-of-LoRAs/task717\_mmmlu\_answer\_generation\_logical\_fallacies",  
 "Lots-of-LoRAs/task705\_mmmlu\_answer\_generation\_high\_school\_macro\_economics",  
 "Lots-of-LoRAs/task775\_entinno\_paraphrase\_generation",  
 "Lots-of-LoRAs/task636\_extract\_and\_sort\_unique\_alphabets\_in\_a\_list",  
 "Lots-of-LoRAs/task084\_babi\_t1\_single\_supporting\_fact\_identify\_relevant\_fact",  
 "Lots-of-LoRAs/task1711\_poki\_text\_generation",  
 "Lots-of-LoRAs/task140\_detoxifying-lms\_classification\_style",  
 "Lots-of-LoRAs/task1448\_disease\_entity\_extraction\_ncbi\_dataset",  
 "Lots-of-LoRAs/task453\_swag\_answer\_generation",  
 "Lots-of-LoRAs/task1207\_atomic\_classification\_allocation",  
 "Lots-of-LoRAs/task734\_mmmlu\_answer\_generation\_sociology",  
 "Lots-of-LoRAs/task298\_storycloze\_correct\_end\_classification",  
 "Lots-of-LoRAs/task587\_amazonfood\_polarity\_correction\_classification",  
 "Lots-of-LoRAs/task703\_mmmlu\_answer\_generation\_high\_school\_geography",  
 "Lots-of-LoRAs/task147\_afs\_argument\_similarity\_gay\_marriage",  
 "Lots-of-LoRAs/task564\_discofuse\_classification",  
 "Lots-of-LoRAs/task1341\_msr\_text\_classification",  
 "Lots-of-LoRAs/task201\_mnil\_neutral\_classification",  
 "Lots-of-LoRAs/task890\_gowd\_classification",  
 "Lots-of-LoRAs/task908\_dialogre\_identify\_familial\_relationships",  
 "Lots-of-LoRAs/task1428\_country\_surface\_area",  
 "Lots-of-LoRAs/task325\_jigsaw\_classification\_identity\_attack",  
 "Lots-of-LoRAs/task1669\_md\_gender\_bias\_text\_modification",  
 "Lots-of-LoRAs/task246\_dream\_question\_generation",  
 "Lots-of-LoRAs/task357\_casino\_classification\_negotiation\_small\_talk",  
 "Lots-of-LoRAs/task1518\_limit\_answer\_generation",  
 "Lots-of-LoRAs/task1148\_maximum\_ascii\_value",  
 "Lots-of-LoRAs/task1605\_ethos\_text\_classification",  
 "Lots-of-LoRAs/task867\_mawps\_multiop\_question\_answering",  
 "Lots-of-LoRAs/task209\_stancedetection\_classification",  
 "Lots-of-LoRAs/task751\_svamp\_subtraction\_question\_answering",  
 "Lots-of-LoRAs/task161\_count\_words\_containing\_letter",  
 "Lots-of-LoRAs/task105\_story\_cloze\_roctories\_sentence\_generation",  
 "Lots-of-LoRAs/task645\_summarization",  
 "Lots-of-LoRAs/task442\_com\_qa\_paraphrase\_question\_generation",  
 "Lots-of-LoRAs/task075\_squad1.1\_answer\_generation",  
 "Lots-of-LoRAs/task269\_csrg\_counterfactual\_story\_generation",  
 "Lots-of-LoRAs/task1568\_propara\_classification",  
 "Lots-of-LoRAs/task834\_mathdataset\_classification",  
 "Lots-of-LoRAs/task1603\_orcasflow\_sentence\_generation",  
 "Lots-of-LoRAs/task685\_mmmlu\_answer\_generation\_clinical\_knowledge",  
 "Lots-of-LoRAs/task083\_babi\_t1\_single\_supporting\_fact\_answer\_generation",  
 "Lots-of-LoRAs/task390\_torque\_text\_span\_selection",  
 "Lots-of-LoRAs/task750\_aqua\_multiple\_choice\_answering",  
 "Lots-of-LoRAs/task1631\_openpi\_answer\_generation",  
 "Lots-of-LoRAs/task1529\_scital1.1\_classification",  
 "Lots-of-LoRAs/task749\_yelp\_restaurant\_review\_classification",  
 "Lots-of-LoRAs/task1217\_atomic\_answer\_generation",  
 "Lots-of-LoRAs/task725\_mmmlu\_answer\_generation\_nutrition",  
 "Lots-of-LoRAs/task889\_goemotions\_classification",  
 "Lots-of-LoRAs/task492\_mwsc\_incorrect\_answer\_generation",  
 "Lots-of-LoRAs/task620\_ohsumed\_medical\_subject\_headings\_answer\_generation",  
 "Lots-of-LoRAs/task084\_all\_elements\_commonsense\_motiv\_text\_generation",  
 "Lots-of-LoRAs/task641\_esnl\_classification",  
 "Lots-of-LoRAs/task318\_stereoset\_classification\_gender",  
 "Lots-of-LoRAs/task846\_pubmedqa\_classification",  
 "Lots-of-LoRAs/task316\_crows-pairs\_classification\_stereotype",  
 "Lots-of-LoRAs/task1188\_count\_max\_freq\_char",  
 "Lots-of-LoRAs/task629\_dbpedia\_14\_classification",  
 "Lots-of-LoRAs/task770\_pawssx\_english\_text\_modification",  
 "Lots-of-LoRAs/task1482\_gene\_extraction\_chempot\_dataset",  
 "Lots-of-LoRAs/task499\_extract\_and\_add\_all\_numbers\_from\_list",  
 "Lots-of-LoRAs/task955\_wiki\_auto\_style\_transfer",  
 "Lots-of-LoRAs/task719\_mmmlu\_answer\_generation\_management",  
 "Lots-of-LoRAs/task723\_mmmlu\_answer\_generation\_moral\_disputes",  
 "Lots-of-LoRAs/task087\_new\_operator\_addsub\_arithmetic",  
 "Lots-of-LoRAs/task211\_logic2text\_classification",  
 "Lots-of-LoRAs/task901\_freebase\_qa\_category\_question\_generation",  
 "Lots-of-LoRAs/task1483\_chemical\_extraction\_chempot\_dataset",  
 "Lots-of-LoRAs/task089\_swap\_words\_verification",  
 "Lots-of-LoRAs/task627\_xlwc\_word\_with\_same\_meaning\_sentence\_generation",  
 "Lots-of-LoRAs/task153\_tomqa\_find\_location\_hard\_clean",  
 "Lots-of-LoRAs/task1342\_amazon\_us\_reviews\_title",  
 "Lots-of-LoRAs/task828\_copa\_commonsense\_cause\_effect",  
 "Lots-of-LoRAs/task064\_all\_elements\_except\_first\_l",  
 "Lots-of-LoRAs/task1387\_anli\_r3\_entailment",  
 "Lots-of-LoRAs/task400\_paws\_paraphrase\_classification",  
 "Lots-of-LoRAs/task1294\_wiki\_qa\_answer\_verification",  
 "Lots-of-LoRAs/task243\_count\_elements\_in\_set\_intersection",  
 "Lots-of-LoRAs/task1572\_samsun\_summary",  
 "Lots-of-LoRAs/task151\_swap\_max\_min",  
 "Lots-of-LoRAs/task574\_air\_dialogue\_sentence\_generation",  
 "Lots-of-LoRAs/task428\_senteval\_inversion",  
 "Lots-of-LoRAs/task366\_synthetic\_return\_primes",  
 "Lots-of-LoRAs/task926\_coached\_conv\_pref\_word\_generation",  
 "Lots-of-LoRAs/task1503\_hatexplain\_classification",  
 "Lots-of-LoRAs/task130\_scan\_structured\_text\_generation\_command\_action\_long",  
 "Lots-of-LoRAs/task515\_senteval\_odd\_word\_out",  
 "Lots-of-LoRAs/task151\_tomqa\_find\_location\_easy\_clean",  
 "Lots-of-LoRAs/task084\_ohsumed\_abstract\_title\_generation",  
 "Lots-of-LoRAs/task1562\_zest\_text\_modification",  
 "Lots-of-LoRAs/task632\_dbpedia\_14\_classification",  
 "Lots-of-LoRAs/task966\_ruletaker\_fact\_checking\_based\_on\_given\_context",  
 "Lots-of-LoRAs/task605\_find\_the\_longest\_common\_subsequence\_in\_two\_lists",  
 "Lots-of-LoRAs/task1487\_organism\_substance\_extraction\_anem\_dataset",  
 "Lots-of-LoRAs/task707\_mmmlu\_answer\_generation\_high\_school\_microeconomics",  
 "Lots-of-LoRAs/task1378\_quarel\_incorrect\_answer\_generation",  
 "Lots-of-LoRAs/task1489\_sarcasmdetection\_tweet\_classification",

## Training Tasks (cont.)

"Lots-of-LoRAs/task1567\_propara\_question\_generation",  
 "Lots-of-LoRAs/task1584\_deal\_or\_no\_dialog\_classification",  
 "Lots-of-LoRAs/task1404\_drac\_conversion",  
 "Lots-of-LoRAs/task691\_mmmlu\_answer\_generation\_college\_physics",  
 "Lots-of-LoRAs/task728\_mmmlu\_answer\_generation\_professional\_accounting",  
 "Lots-of-LoRAs/task219\_roctories\_title\_answer\_generation",  
 "Lots-of-LoRAs/task964\_librispeech\_asr\_text\_auto\_completion",  
 "Lots-of-LoRAs/task1509\_evaluation\_antonyms",  
 "Lots-of-LoRAs/task582\_naturalquestion\_answer\_generation",  
 "Lots-of-LoRAs/task455\_swag\_context\_text\_generation",  
 "Lots-of-LoRAs/task063\_librispeech\_asr\_next\_word\_prediction",  
 "Lots-of-LoRAs/task382\_hybridqa\_answer\_generation",  
 "Lots-of-LoRAs/task859\_prost\_question\_generation",  
 "Lots-of-LoRAs/task1393\_superglue\_copa\_text\_completion",  
 "Lots-of-LoRAs/task1565\_triviaqa\_classification",  
 "Lots-of-LoRAs/task1720\_civil\_comments\_toxicity\_classification",  
 "Lots-of-LoRAs/task670\_ambigua\_question\_generation",  
 "Lots-of-LoRAs/task689\_mmmlu\_answer\_generation\_college\_mathematics",  
 "Lots-of-LoRAs/task324\_jigsaw\_classification\_disagree",  
 "Lots-of-LoRAs/task1420\_mathqa\_general",  
 "Lots-of-LoRAs/task618\_amazonreview\_summary\_text\_generation",  
 "Lots-of-LoRAs/task625\_xlwc\_true\_or\_false\_answer\_generation",  
 "Lots-of-LoRAs/task377\_remove\_words\_of\_given\_length",  
 "Lots-of-LoRAs/task929\_products\_reviews\_classification",  
 "Lots-of-LoRAs/task286\_storycloze\_correct\_end\_classification",  
 "Lots-of-LoRAs/task852\_synthetic\_multiply\_odds",  
 "Lots-of-LoRAs/task1332\_check\_leap\_year",  
 "Lots-of-LoRAs/task1444\_round\_power\_of\_two",  
 "Lots-of-LoRAs/task850\_synthetic\_longest\_palindrome",  
 "Lots-of-LoRAs/task708\_mmmlu\_answer\_generation\_high\_school\_physics",  
 "Lots-of-LoRAs/task1292\_yelp\_review\_full\_text\_categorization",  
 "Lots-of-LoRAs/task110\_logic2text\_sentence\_generation",  
 "Lots-of-LoRAs/task155\_count\_nouns\_verbs",  
 "Lots-of-LoRAs/task429\_senteval\_tense",  
 "Lots-of-LoRAs/task245\_check\_presence\_in\_set\_intersection",  
 "Lots-of-LoRAs/task137\_detoxifying-lms\_classification\_toxicity",  
 "Lots-of-LoRAs/task1566\_propara\_structured\_text\_generation",  
 "Lots-of-LoRAs/task1148\_country\_capital",  
 "Lots-of-LoRAs/task624\_event2mind\_word\_generation",  
 "Lots-of-LoRAs/task022\_cosmosqa\_passage\_inappropriate\_binary",  
 "Lots-of-LoRAs/task118\_semeval\_open\_vocabulary\_mathematical\_answer\_generation",  
 "Lots-of-LoRAs/task687\_mmmlu\_answer\_generation\_college\_chemistry",  
 "Lots-of-LoRAs/task1167\_penn\_treebank\_coarse\_pos\_tagging",  
 "Lots-of-LoRAs/task080\_boolq\_yes\_no\_question",  
 "Lots-of-LoRAs/task092\_winegrade\_answer\_generation",  
 "Lots-of-LoRAs/task1502\_hatexplain\_classification",  
 "Lots-of-LoRAs/task805\_mawps\_addsub\_question\_answering",  
 "Lots-of-LoRAs/task181\_outcome\_extraction",  
 "Lots-of-LoRAs/task020\_pre-answer-generation\_easy",  
 "Lots-of-LoRAs/task698\_mmmlu\_answer\_generation\_global\_facts",  
 "Lots-of-LoRAs/task956\_leetcode\_420\_strong\_password\_check",  
 "Lots-of-LoRAs/task732\_mmmlu\_answer\_generation\_public\_relations",  
 "Lots-of-LoRAs/task721\_medical\_generation\_medical\_genetics",  
 "Lots-of-LoRAs/task370\_synthetic\_remove\_divisible\_by\_3",  
 "Lots-of-LoRAs/task1400\_obqa\_incorrect\_answer\_generation",  
 "Lots-of-LoRAs/task1199\_atomic\_classification\_xatrr",  
 "Lots-of-LoRAs/task1606\_ethos\_text\_classification",  
 "Lots-of-LoRAs/task288\_gigaword\_summarization",  
 "Lots-of-LoRAs/task1670\_md\_gender\_bias\_text\_modification",  
 "Lots-of-LoRAs/task207\_max\_element\_lists",  
 "Lots-of-LoRAs/task1208\_atomic\_classification\_isbefore",  
 "Lots-of-LoRAs/task457\_matres\_conditional\_classification",  
 "Lots-of-LoRAs/task1308\_amazonreview\_category\_classification",  
 "Lots-of-LoRAs/task1310\_amazonreview\_rating\_classification",  
 "Lots-of-LoRAs/task874\_opus\_xhosanavy\_sr",  
 "Lots-of-LoRAs/task1541\_agnews\_classification",  
 "Lots-of-LoRAs/task1609\_xquad\_en\_question\_generation",  
 "Lots-of-LoRAs/task1210\_logic2text\_structured\_text\_generation",  
 "Lots-of-LoRAs/task1318\_country\_national\_dish",  
 "Lots-of-LoRAs/task365\_synthetic\_remove\_vowels",  
 "Lots-of-LoRAs/task755\_find\_longest\_substring\_and\_replace\_its\_sorted\_lowercase",  
 "Lots-of-LoRAs/task123\_conala\_sort\_dictionary",  
 "Lots-of-LoRAs/task1316\_remove\_duplicates\_string",  
 "Lots-of-LoRAs/task1378\_quarel\_correct\_answer\_generation",  
 "Lots-of-LoRAs/task475\_yelp\_polarity\_classification",  
 "Lots-of-LoRAs/task903\_deceptive\_opinion\_spam\_classification",  
 "Lots-of-LoRAs/task070\_abductivenli\_incorrect\_classification",  
 "Lots-of-LoRAs/task720\_mmmlu\_answer\_generation\_marketing",  
 "Lots-of-LoRAs/task067\_abductivenli\_answer\_generation",  
 "Lots-of-LoRAs/task1564\_triviaqa\_answer\_generation",  
 "Lots-of-LoRAs/task270\_csrg\_counterfactual\_context\_generation",  
 "Lots-of-LoRAs/task167\_strategyqa\_question\_generation",  
 "Lots-of-LoRAs/task1504\_hatexplain\_answer\_generation",  
 "Lots-of-LoRAs/task178\_quartz\_question\_answering",  
 "Lots-of-LoRAs/task277\_stereoset\_sentence\_generation\_stereotype",  
 "Lots-of-LoRAs/task1315\_find\_range\_array",  
 "Lots-of-LoRAs/task1434\_head\_qa\_classification",  
 "Lots-of-LoRAs/task192\_hotpotqa\_sentence\_generation",  
 "Lots-of-LoRAs/task1157\_bard\_analogical\_reasoning\_rooms\_for\_containers",  
 "Lots-of-LoRAs/task872\_nummersense1",  
 "Lots-of-LoRAs/task563\_discofuse\_answer\_generation",  
 "Lots-of-LoRAs/task714\_mmmlu\_answer\_generation\_human\_sexuality",  
 "Lots-of-LoRAs/task1212\_atomic\_classification\_hasproperty",  
 "Lots-of-LoRAs/task495\_semeval\_headline\_classification",  
 "Lots-of-LoRAs/task1583\_bless\_meronym\_classification",  
 "Lots-of-LoRAs/task753\_svamp\_addition\_question\_answering",  
 "Lots-of-LoRAs/task433\_wormat\_classification\_profession\_anti",  
 "Lots-of-LoRAs/task1427\_country\_region\_in\_world",  
 "Lots-of-LoRAs/task092\_check\_prime\_classification",  
 "Lots-of-LoRAs/task1285\_kpa\_keypoint\_matching",

Figure 7: Training tasks from Lots-of-LoRAs (based on the SNI dataset) used for training the Text-to-LoRA model. The struck out names indicate removed tasks due to benchmark contamination.



## Training Tasks (cont.)

~~"Lots-of-LoRAs/task333\_hateval\_classification\_hate\_en",~~  
~~"Lots-of-LoRAs/task329\_gap\_classification",~~  
~~"Lots-of-LoRAs/task398\_semeval\_2018\_task1\_tweet\_joy\_detection",~~  
~~"Lots-of-LoRAs/task157\_count\_vowels\_and\_consonants",~~  
~~"Lots-of-LoRAs/task074\_corpus1\_question\_generation",~~  
~~"Lots-of-LoRAs/task1506\_celebrity\_minimal\_dob\_span",~~  
~~"Lots-of-LoRAs/task697\_mmmlu\_answer\_generation\_formal\_logic",~~  
~~"Lots-of-LoRAs/task285\_imdb\_answer\_generation",~~  
~~"Lots-of-LoRAs/task393\_plausible\_result\_generation",~~  
~~"Lots-of-LoRAs/task1147\_country\_currency",~~  
~~"Lots-of-LoRAs/task1585\_root09\_hypemym\_generation",~~  
~~"Lots-of-LoRAs/task648\_answer\_generation",~~  
~~"Lots-of-LoRAs/task353\_casino\_classification\_negotiation\_elicit\_pref",~~  
~~"Lots-of-LoRAs/task1431\_head\_qa\_answer\_generation",~~  
~~"Lots-of-LoRAs/task148\_afs\_argument\_quality\_gay\_marriage",~~  
~~"Lots-of-LoRAs/task585\_preposition\_classification",~~  
~~"Lots-of-LoRAs/task081\_pigs\_wrong\_answer\_generation",~~  
~~"Lots-of-LoRAs/task477\_cls\_english\_dvd\_classification",~~  
~~"Lots-of-LoRAs/task1582\_bless\_hypemym\_generation",~~  
~~"Lots-of-LoRAs/task365\_casino\_classification\_negotiation\_other\_need",~~  
~~"Lots-of-LoRAs/task284\_corpus1\_question\_generation",~~  
~~"Lots-of-LoRAs/task633\_dbpedia\_14\_answer\_generation",~~  
~~"Lots-of-LoRAs/task093\_conala\_normalize\_lists",~~  
~~"Lots-of-LoRAs/task722\_mmmlu\_answer\_generation\_random\_topic",~~  
~~"Lots-of-LoRAs/task566\_circa\_classification",~~  
~~"Lots-of-LoRAs/task1152\_bard\_analogical\_reasoning\_causation",~~  
~~"Lots-of-LoRAs/task1452\_location\_entity\_extraction\_btc\_corpus",~~  
~~"Lots-of-LoRAs/task422\_corpus1\_question\_generation",~~  
~~"Lots-of-LoRAs/task925\_coached\_conv\_classifier",~~  
~~"Lots-of-LoRAs/task1703\_ljspeech\_textmodification",~~  
~~"Lots-of-LoRAs/task833\_poem\_sentiment\_classification",~~  
~~"Lots-of-LoRAs/task1210\_atomic\_classification\_madeup0f",~~  
~~"Lots-of-LoRAs/task679\_hope\_edi\_english\_text\_classification",~~  
~~"Lots-of-LoRAs/task1203\_atomic\_classification\_xreact",~~  
~~"Lots-of-LoRAs/task1089\_check\_monotonic\_array",~~  
~~"Lots-of-LoRAs/task389\_torque\_generate\_temporal\_question",~~  
~~"Lots-of-LoRAs/task378\_reverse\_words\_of\_given\_length",~~  
~~"Lots-of-LoRAs/task637\_extract\_and\_sort\_unique\_digits\_in\_a\_list",~~  
~~"Lots-of-LoRAs/task101\_reverse\_and\_concatenate\_all\_elements\_from\_index\_i\_to\_j",~~  
~~"Lots-of-LoRAs/task1355\_sent\_comp\_summarization",~~  
~~"Lots-of-LoRAs/task640\_esnl\_classification",~~  
~~"Lots-of-LoRAs/task344\_hybridqa\_answer\_generation",~~  
~~"Lots-of-LoRAs/task1190\_add\_integer\_to\_list",~~  
~~"Lots-of-LoRAs/task733\_mmmlu\_answer\_generation\_security\_studies",~~  
~~"Lots-of-LoRAs/task107\_splash\_question\_to\_sql",~~  
~~"Lots-of-LoRAs/task1209\_atomic\_classification\_objectuse",~~  
~~"Lots-of-LoRAs/task505\_count\_all\_numerical\_elements\_in\_list",~~  
~~"Lots-of-LoRAs/task1385\_anli\_r1\_entailment",~~  
~~"Lots-of-LoRAs/task1136\_vocab\_commonsense\_mc\_classification",~~  
~~"Lots-of-LoRAs/task328\_jigsaw\_classification\_insult",~~  
~~"Lots-of-LoRAs/task609\_sbic\_potentially\_offense\_binary\_classification",~~  
~~"Lots-of-LoRAs/task413\_mickey\_en\_sentence\_perturbation\_generation",~~  
~~"Lots-of-LoRAs/task072\_abductivenl\_answer\_generation",~~  
~~"Lots-of-LoRAs/task1425\_country\_iso\_numeric",~~  
~~"Lots-of-LoRAs/task1451\_drug\_dose\_extraction",~~  
~~"Lots-of-LoRAs/task713\_mmmlu\_answer\_generation\_human\_aging",~~  
~~"Lots-of-LoRAs/task642\_esnl\_classification",~~  
~~"Lots-of-LoRAs/task1321\_country\_continent",~~  
~~"Lots-of-LoRAs/task454\_swag\_incorrect\_answer\_generation",~~  
~~"Lots-of-LoRAs/task504\_count\_all\_alphabetical\_elements\_in\_list",~~  
~~"Lots-of-LoRAs/task696\_mmmlu\_answer\_generation\_elementary\_mathematics",~~  
~~"Lots-of-LoRAs/task1429\_evaluation\_semantic\_relation\_classification",~~  
~~"Lots-of-LoRAs/task1645\_medical\_question\_pair\_dataset\_text\_classification",~~  
~~"Lots-of-LoRAs/task431\_senteval\_object\_count",~~  
~~"Lots-of-LoRAs/task1317\_country\_calling\_code",~~  
~~"Lots-of-LoRAs/task131\_scan\_long\_text\_generation\_action\_command\_long",~~  
~~"Lots-of-LoRAs/task675\_google\_wellformed\_query\_sentence\_generation",~~  
~~"Lots-of-LoRAs/task1158\_bard\_analogical\_reasoning\_manipulating\_items",~~  
~~"Lots-of-LoRAs/task1325\_qa\_zre\_question\_generation\_on\_subject\_relation",~~  
~~"Lots-of-LoRAs/task300\_storycloze\_order\_generation",~~  
~~"Lots-of-LoRAs/task1347\_glue\_sts-b\_similarity\_classification",~~  
~~"Lots-of-LoRAs/task1328\_qa\_zre\_relation\_generation\_from\_question",~~  
~~"Lots-of-LoRAs/task630\_dbpedia\_14\_classification",~~  
~~"Lots-of-LoRAs/task460\_gasper\_answer\_generation",~~  
~~"Lots-of-LoRAs/task1665\_trianglecopa\_question\_generation",~~  
~~"Lots-of-LoRAs/task1154\_bard\_analogical\_reasoning\_travel",~~  
~~"Lots-of-LoRAs/task190\_storycloze\_order\_generation",~~  
~~"Lots-of-LoRAs/task617\_amazonreview\_category\_text\_generation",~~  
~~"Lots-of-LoRAs/task1508\_wordnet\_antonyms",~~  
~~"Lots-of-LoRAs/task628\_xlwc\_word\_with\_different\_meaning\_sentence\_generation",~~  
~~"Lots-of-LoRAs/task600\_find\_the\_longest\_common\_substring\_in\_two\_strings",~~  
~~"Lots-of-LoRAs/task1380\_quarel\_correct\_option\_generation",~~  
~~"Lots-of-LoRAs/task489\_mwsc\_question\_generation",~~  
~~"Lots-of-LoRAs/task905\_hate\_speech\_offensive\_classification",~~  
~~"Lots-of-LoRAs/task065\_limetravel\_consistent\_sentence\_classification",~~  
~~"Lots-of-LoRAs/task044\_essential\_terms\_identifying\_essential\_words",~~  
~~"Lots-of-LoRAs/task069\_abductivenl\_classification",~~  
~~"Lots-of-LoRAs/task356\_casino\_classification\_negotiation\_self\_need",~~  
~~"Lots-of-LoRAs/task403\_creak\_commonsense\_inference",~~  
~~"Lots-of-LoRAs/task590\_amazonfood\_summary\_correction\_classification",~~  
~~"Lots-of-LoRAs/task577\_curiosity\_dialogs\_classification",~~  
~~"Lots-of-LoRAs/task519\_per\_sentiment\_classification",~~  
~~"Lots-of-LoRAs/task1721\_civil\_comments\_obscenity\_classification",~~  
~~"Lots-of-LoRAs/task351\_winomt\_classification\_gender\_identifiability\_anti",~~  
~~"Lots-of-LoRAs/task904\_hate\_speech\_offensive\_classification",~~  
~~"Lots-of-LoRAs/task244\_count\_elements\_in\_set\_union",~~  
~~"Lots-of-LoRAs/task1712\_poki\_classification",~~  
~~"Lots-of-LoRAs/task683\_online\_privacy\_policy\_text\_purpose\_answer\_generation",~~  
~~"Lots-of-LoRAs/task1197\_atomic\_classification\_creat",~~  
~~"Lots-of-LoRAs/task1592\_yahoo\_answers\_topics\_classification",~~

## Training Tasks (cont.)

~~"Lots-of-LoRAs/task274\_outruling\_legal\_classification",~~  
~~"Lots-of-LoRAs/task278\_stereoset\_sentence\_generation\_antistereotype",~~  
~~"Lots-of-LoRAs/task1534\_daily\_dialog\_question\_classification",~~  
~~"Lots-of-LoRAs/task891\_gap\_coreference\_resolution",~~  
~~"Lots-of-LoRAs/task694\_mmmlu\_answer\_generation\_econometrics",~~  
~~"Lots-of-LoRAs/task497\_extract\_all\_numbers\_from\_list\_in\_order",~~  
~~"Lots-of-LoRAs/task488\_extract\_all\_alphabetical\_elements\_from\_list\_in\_order",~~  
~~"Lots-of-LoRAs/task144\_subjq\_question\_answering",~~  
~~"Lots-of-LoRAs/task1722\_civil\_comments\_threat\_classification",~~  
~~"Lots-of-LoRAs/task1728\_web\_nlp\_data\_to\_text",~~  
~~"Lots-of-LoRAs/task291\_semex2020\_task4\_commonsense\_validation",~~  
~~"Lots-of-LoRAs/task284\_imdb\_classification",~~  
~~"Lots-of-LoRAs/task1288\_glue\_mrpc\_paraphrasing",~~  
~~"Lots-of-LoRAs/task128\_scan\_structured\_text\_generation\_command\_action\_short",~~  
~~"Lots-of-LoRAs/task580\_socialqa\_answer\_generation",~~  
~~"Lots-of-LoRAs/task616 cola\_classification",~~  
~~"Lots-of-LoRAs/task1401\_obqa\_sentence\_generation",~~  
~~"Lots-of-LoRAs/task716\_mmmlu\_answer\_generation\_jurisprudence",~~  
~~"Lots-of-LoRAs/task1421\_mathqa\_other",~~  
~~"Lots-of-LoRAs/task1656\_goaqa\_answer\_generation",~~  
~~"Lots-of-LoRAs/task1311\_amazonreview\_rating\_classification",~~  
~~"Lots-of-LoRAs/task516\_senteval\_conjoints\_inversion",~~  
~~"Lots-of-LoRAs/task593\_scic\_explanation\_generation",~~  
~~"Lots-of-LoRAs/task138\_detoxifying\_tms\_classification\_fluency",~~  
~~"Lots-of-LoRAs/task1186\_nre\_himgo\_classification",~~  
~~"Lots-of-LoRAs/task1326\_qa\_zre\_question\_generation\_from\_answer",~~  
~~"Lots-of-LoRAs/task119\_semeval\_geometric\_mathematical\_answer\_generation",~~  
~~"Lots-of-LoRAs/task108\_contextualabusedetection\_classification",~~  
~~"Lots-of-LoRAs/task584\_udeps\_eng\_fine\_pos\_tagging",~~  
~~"Lots-of-LoRAs/task388\_torque\_token\_classification",~~  
~~"Lots-of-LoRAs/task045\_miscellaneous\_sentence\_paraphrasing",~~  
~~"Lots-of-LoRAs/task821\_protoqa\_question\_generation",~~  
~~"Lots-of-LoRAs/task1581\_qcqa-perturbed\_answer\_generation",~~  
~~"Lots-of-LoRAs/task695\_mmmlu\_answer\_generation\_electrical\_engineering",~~  
~~"Lots-of-LoRAs/task596\_mocha\_question\_generation",~~  
~~"Lots-of-LoRAs/task568\_circa\_question\_generation",~~  
~~"Lots-of-LoRAs/task085\_unnatural\_addsub\_arithmetic",~~  
~~"Lots-of-LoRAs/task1495\_adverse\_drug\_event\_classification",~~  
~~"Lots-of-LoRAs/task1403\_check\_validity\_date\_mmddyyyy",~~  
~~"Lots-of-LoRAs/task1453\_person\_entity\_extraction\_btc\_corpus",~~  
~~"Lots-of-LoRAs/task1201\_atomic\_classification\_xintent",~~  
~~"Lots-of-LoRAs/task923\_event2mind\_classifier",~~  
~~"Lots-of-LoRAs/task1204\_atomic\_classification\_hinderedby",~~  
~~"Lots-of-LoRAs/task1510\_evaluation\_relation\_extraction",~~  
~~"Lots-of-LoRAs/task754\_svamp\_common-division\_question\_answering",~~  
~~"Lots-of-LoRAs/task1403\_check\_validity\_date\_mmddyyyy",~~  
~~"Lots-of-LoRAs/task1192\_food\_flavor\_profile",~~  
~~"Lots-of-LoRAs/task565\_circa\_answer\_generation",~~  
~~"Lots-of-LoRAs/task146\_afs\_argument\_similarity\_gun\_control",~~  
~~"Lots-of-LoRAs/task666\_mmmlu\_answer\_generation\_astronomy",~~  
~~"Lots-of-LoRAs/task050\_multirc\_answerability",~~  
~~"Lots-of-LoRAs/task704\_mmmlu\_answer\_generation\_high\_school\_government",~~  
~~"Lots-of-LoRAs/task1447\_drug\_extraction\_ade",~~  
~~"Lots-of-LoRAs/task579\_socialqa\_classification",~~  
~~"Lots-of-LoRAs/task1196\_atomic\_classification\_offect",~~  
~~"Lots-of-LoRAs/task267\_concatenate\_and\_reverse\_all\_elements\_from\_index\_i\_to\_j",~~  
~~"Lots-of-LoRAs/task206\_collatz\_conjecture",~~  
~~"Lots-of-LoRAs/task936\_defeasible\_nli\_snli\_classification",~~  
~~"Lots-of-LoRAs/task323\_jigsaw\_classification\_sexually\_explicit",~~  
~~"Lots-of-LoRAs/task494\_review\_polarity\_answer\_generation",~~  
~~"Lots-of-LoRAs/task202\_question\_gasper\_question\_generation",~~  
~~"Lots-of-LoRAs/task1409\_dart\_text\_generation",~~  
~~"Lots-of-LoRAs/task1313\_amazonreview\_polarity\_classification",~~  
~~"Lots-of-LoRAs/task076\_splash\_correcting\_sql\_mistake",~~  
~~"Lots-of-LoRAs/task686\_mmmlu\_answer\_generation\_college\_biology",~~  
~~"Lots-of-LoRAs/task740\_hoestq\_answer\_generation\_quantity",~~  
~~"Lots-of-LoRAs/task1196\_atomic\_classification\_offect",~~  
~~"Lots-of-LoRAs/task1211\_atomic\_classification\_hassubevent",~~  
~~"Lots-of-LoRAs/task113\_count\_frequency\_of\_letter",~~  
~~"Lots-of-LoRAs/task280\_stereoset\_classification\_stereotype\_type",~~  
~~"Lots-of-LoRAs/task1551\_every\_nth\_element\_from\_kth\_element",~~  
~~"Lots-of-LoRAs/task116\_com2sense\_commonsense\_reasoning",~~  
~~"Lots-of-LoRAs/task518\_emo\_different\_dialogue\_emotions",~~  
~~"Lots-of-LoRAs/task1520\_qa\_srl\_answer\_generation",~~  
~~"Lots-of-LoRAs/task079\_conala\_concat\_strings",~~  
~~"Lots-of-LoRAs/task513\_argument\_stance\_classification",~~  
~~"Lots-of-LoRAs/task1590\_diplomacy\_text\_generation",~~  
~~"Lots-of-LoRAs/task1713\_convai3\_sentence\_generation",~~  
~~"Lots-of-LoRAs/task1386\_anli\_r2\_entailment",~~  
~~"Lots-of-LoRAs/task063\_first\_i\_elements",~~  
~~"Lots-of-LoRAs/task183\_rhyme\_generation",~~  
~~"Lots-of-LoRAs/task1447\_drug\_extraction\_ade",~~  
~~"Lots-of-LoRAs/task1771\_ambigua\_text\_generation",~~  
~~"Lots-of-LoRAs/task068\_abductivenl\_incorrect\_answer\_generation",~~  
~~"Lots-of-LoRAs/task858\_inquisitive\_span\_detection",~~  
~~"Lots-of-LoRAs/task699\_mmmlu\_answer\_generation\_high\_school\_biology",~~  
~~"Lots-of-LoRAs/task1593\_yahoo\_answers\_topics\_classification",~~  
~~"Lots-of-LoRAs/task700\_mmmlu\_answer\_generation\_high\_school\_chemistry",~~  
~~"Lots-of-LoRAs/task363\_sst2\_polarity\_classification",~~  
~~"Lots-of-LoRAs/task121\_zest\_text\_modification",~~  
~~"Lots-of-LoRAs/task190\_snli\_classification",~~  
~~"Lots-of-LoRAs/task1168\_brown\_coarse\_pos\_tagging",~~  
~~"Lots-of-LoRAs/task195\_sentiment140\_classification",~~  
~~"Lots-of-LoRAs/task1723\_civil\_comments\_sexuallyexplicit\_classification",~~  
~~"Lots-of-LoRAs/task1449\_disease\_entity\_extraction\_bc5cdi\_dataset",~~  
~~"Lots-of-LoRAs/task1418\_mathqa\_gain",~~  
~~"Lots-of-LoRAs/task1398\_obqa\_question\_generation",~~  
~~"Lots-of-LoRAs/task893\_gap\_fill\_the\_blank\_reference\_resolution",~~  
~~"Lots-of-LoRAs/task326\_jigsaw\_classification\_obscene",~~

Figure 8: Training tasks from Lots-of-LoRAs (based on the SNI dataset) used for training the Text-to-LoRA model. The stricken out names indicate removed tasks due to benchmark contamination.

## Training Tasks (cont.)

"Lots-of-LoRAs/task1194\_kth\_largest\_element",  
 "Lots-of-LoRAs/task102\_commonsense\_sentence\_generation",  
 "Lots-of-LoRAs/task145\_afs\_argument\_similarity\_death\_penalty",  
 "Lots-of-LoRAs/task1338\_peixian\_equality\_evaluation\_corpus\_sentiment\_classifier",  
 "Lots-of-LoRAs/task391\_causal\_relationship",  
 "Lots-of-LoRAs/task176\_break\_decompose\_questions",  
 "Lots-of-LoRAs/task319\_stereoset\_classification\_profession",  
 "Lots-of-LoRAs/task369\_casino\_classification\_negotiation\_vouch\_fair",  
 "Lots-of-LoRAs/task866\_conv\_ai\_2\_classification",  
 "Lots-of-LoRAs/task1729\_personachat\_generate\_next",  
 "Lots-of-LoRAs/task761\_app\_review\_classification",  
 "Lots-of-LoRAs/task1320\_country\_domain\_id",  
 "Lots-of-LoRAs/task1596\_event2mind\_text\_generation\_2",  
 "Lots-of-LoRAs/task1601\_webquestions\_answer\_generation",  
 "Lots-of-LoRAs/task615\_moviesqa\_answer\_generation",  
 "Lots-of-LoRAs/task1283\_hrngo\_quality\_classification",  
 "Lots-of-LoRAs/task638\_multi\_woz\_classification",  
 "Lots-of-LoRAs/task607\_sbic\_intentional\_offense\_binary\_classification",  
 "Lots-of-LoRAs/task692\_mmmlu\_answer\_generation\_computer\_security",  
 "Lots-of-LoRAs/task698\_amazonfood\_rating\_classification",  
 "Lots-of-LoRAs/task129\_scan\_long\_text\_generation\_action\_command\_short",  
 "Lots-of-LoRAs/task1200\_atomic\_classification\_xeffect",  
 "Lots-of-LoRAs/task1486\_cell\_extraction\_anem\_dataset",  
 "Lots-of-LoRAs/task1406\_kth\_smallest\_element",  
 "Lots-of-LoRAs/task583\_udeps\_eng\_coarse\_pos\_tagging",  
 "Lots-of-LoRAs/task094\_conala\_calculate\_mean",  
 "Lots-of-LoRAs/task1731\_quantz\_question\_answering",  
 "Lots-of-LoRAs/task308\_jeopardy\_answer\_generation\_all",  
 "Lots-of-LoRAs/task664\_mmmlu\_answer\_generation\_abstract\_algebra",  
 "Lots-of-LoRAs/task163\_count\_words\_ending\_with\_letter",  
 "Lots-of-LoRAs/task127\_scan\_long\_text\_generation\_action\_command\_all",  
 "Lots-of-LoRAs/task736\_mmmlu\_answer\_generation\_virology",  
 "Lots-of-LoRAs/task379\_agnews\_topic\_classification",  
 "Lots-of-LoRAs/task390\_gap\_answer\_generation",  
 "Lots-of-LoRAs/task1599\_smcalfow\_classification",  
 "Lots-of-LoRAs/task875\_emotion\_classification",  
 "Lots-of-LoRAs/task1214\_atomic\_classification\_xwant",  
 "Lots-of-LoRAs/task933\_wiki\_auto\_style\_transfer",  
 "Lots-of-LoRAs/task1319\_country\_by\_barcode\_prefix",  
 "Lots-of-LoRAs/task456\_matres\_intention\_classification",  
 "Lots-of-LoRAs/task1189\_check\_char\_in\_string",  
 "Lots-of-LoRAs/task1857\_goose\_question\_generation",  
 "Lots-of-LoRAs/task517\_emo\_classify\_emotion\_of\_dialogue",  
 "Lots-of-LoRAs/task385\_socialqa\_incorrect\_answer\_generation",  
 "Lots-of-LoRAs/task472\_haspart\_classification",  
 "Lots-of-LoRAs/task047\_miscellaneous\_answering\_science\_questions",  
 "Lots-of-LoRAs/task1533\_daily\_dialog\_formal\_classification",  
 "Lots-of-LoRAs/task249\_enhanced\_wsc\_pronoun\_disambiguation",  
 "Lots-of-LoRAs/task095\_conala\_max\_absolute\_value",  
 "Lots-of-LoRAs/task1479\_organization\_entity\_extraction\_btc\_corpus",  
 "Lots-of-LoRAs/task724\_mmmlu\_answer\_generation\_moral\_scenarios",  
 "Lots-of-LoRAs/task1418\_bless\_semantic\_relation\_classification",  
 "Lots-of-LoRAs/task507\_position\_of\_all\_numerical\_elements\_in\_list",  
 "Lots-of-LoRAs/task043\_essential\_terms\_answering\_incomplete\_questions",  
 "Lots-of-LoRAs/task1394\_meta\_woz\_task\_classification",  
 "Lots-of-LoRAs/task149\_afs\_argument\_quality\_death\_penalty",  
 "Lots-of-LoRAs/task1156\_bard\_analogical\_reasoning\_tools",  
 "Lots-of-LoRAs/task1598\_nyc\_long\_text\_generation",  
 "Lots-of-LoRAs/task322\_jigsaw\_classification\_threat",  
 "Lots-of-LoRAs/task1714\_convai3\_sentence\_generation",  
 "Lots-of-LoRAs/task727\_mmmlu\_answer\_generation\_prehistory",  
 "Lots-of-LoRAs/task609\_jigsaw\_answer\_generation",  
 "Lots-of-LoRAs/task1704\_ljspeech\_textmodification",  
 "Lots-of-LoRAs/task077\_splash\_explanation\_to\_sql",  
 "Lots-of-LoRAs/task1390\_wscfixed\_coreference",  
 "Lots-of-LoRAs/task667\_mmmlu\_answer\_generation\_business\_ethics",  
 "Lots-of-LoRAs/task1724\_civil\_comments\_insult\_classification",  
 "Lots-of-LoRAs/task1088\_array\_of\_products",  
 "Lots-of-LoRAs/task1087\_two\_number\_sum",  
 "Lots-of-LoRAs/task550\_discofuse\_sentence\_generation",  
 "Lots-of-LoRAs/task892\_gap\_reverse\_coreference\_resolution",  
 "Lots-of-LoRAs/task674\_google\_wellformed\_query\_sentence\_generation",  
 "Lots-of-LoRAs/task14380\_holloway\_completion",  
 "Lots-of-LoRAs/task509\_collate\_of\_all\_alphabetical\_and\_numerical\_elements\_in\_list\_separately",  
 "Lots-of-LoRAs/task346\_hybridqa\_classification",  
 "Lots-of-LoRAs/task769\_qed\_summarization",  
 "Lots-of-LoRAs/task1320\_winogrande\_qaqqa\_answer\_generation",  
 "Lots-of-LoRAs/task1322\_country\_government\_type",  
 "Lots-of-LoRAs/task341\_winomt\_classification\_gender\_ant",  
 "Lots-of-LoRAs/task290\_tellmewhy\_question\_answerability",  
 "Lots-of-LoRAs/task335\_hateval\_classification\_aggressive\_en",  
 "Lots-of-LoRAs/task879\_schema\_guided\_dstc8\_classification",  
 "Lots-of-LoRAs/task861\_asdiv\_addsub\_question\_answering",  
 "Lots-of-LoRAs/task594\_sciq\_question\_generation",  
 "Lots-of-LoRAs/task066\_limettravel\_binary\_consistency\_classification",  
 "Lots-of-LoRAs/task162\_count\_words\_starting\_with\_letter",  
 "Lots-of-LoRAs/task1584\_evaluation\_meronym\_classification",  
 "Lots-of-LoRAs/task1622\_disfl\_qa\_text\_modication",  
 "Lots-of-LoRAs/task247\_dream\_answer\_generation"

Figure 9: Training tasks from Lots-of-LoRAs (based on the SNI dataset) used for training the Text-to-LoRA model. The stricken out names indicate removed tasks due to benchmark contamination.

## Validation Tasks

"Lots-of-LoRAs/task035\_winogrande\_question\_modification\_person",  
 "Lots-of-LoRAs/task039\_qasc\_find\_overlapping\_words",  
 "Lots-of-LoRAs/task1557\_jfleg\_answer\_generation",  
 "Lots-of-LoRAs/task202\_mnli\_contradiction\_classification",  
 "Lots-of-LoRAs/task304\_numeric\_fused\_head\_resolution",  
 "Lots-of-LoRAs/task362\_spolin\_yesand\_prompt\_response\_sub\_classification",  
 "Lots-of-LoRAs/task614\_glucose\_cause\_event\_detection",  
 "Lots-of-LoRAs/task701\_mmmlu\_answer\_generation\_high\_school\_computer\_science",  
 "Lots-of-LoRAs/task706\_mmmlu\_answer\_generation\_high\_school\_mathematics",  
 "Lots-of-LoRAs/task710\_mmmlu\_answer\_generation\_high\_school\_statistics",  
 "Lots-of-LoRAs/task726\_mmmlu\_answer\_generation\_philosophy",

Figure 10: Validation tasks used during the training of the Text-to-LoRA model.

### J.1. Benchmark Details

Every benchmark used in the experiments is publicly available in HuggingFace dataset space. We evaluate the models on the benchmarks detailed as follows.

#### J.1.1. GSM8K

We evaluate the models on the test split, using chain-of-thought response pre-filling: *"Let's think step by step."*

#### J.1.2. HUMANEVAL AND MBPP

We use the `evalplus` library (Liu et al., 2023) for coding evaluation. For both MBPP and HumanEval, we use the following response pre-fill: ``python

## J.2. Question-Answering Tasks

The rest of the benchmarks are question-answering based tasks. In these tasks, we do not use response-prefilling. Instead, each task has a specific instruction template shown in Listing 4.

## K. Utilizing Full Adaptation Matrix vs Low-Rank Matrices




Figure 11: Each plot shows the similarity between a benchmark LoRA adapter and 479 SNI-trained adapters in the  $\Delta W$  weight space. There is a positive correlation between the two variables indicated by small positive Pearson correlation coefficients.

Similar to Figure 6, Figure 11 shows the similarity of SNI adapters to benchmark-specific adapters, but instead of using the concatenation of flattened  $A$  and  $B$  matrices, we use flattened  $\Delta W$  instead. With the change, we find a positive correlation between the task embedding similarity and the adapter similarity in the weight space. This is likely because, for a given  $\Delta W$  matrix, there are many possible permutations of low-rank matrices  $A$  and  $B$ . This suggests that if we compute the reconstruction loss in the full adaptation matrix space, reconstruction-trained T2L could generalize better. However, we empirically find that it does not outperform T2L trained to reconstruct low-rank matrices at zero-shot LoRA generation.

## L. Generating Task Descriptions with a Foundation Language Model

We automate task description generation for each task by leveraging powerful closed-source language models (Achiam et al., 2023). We query GPT-4o mini with carefully constructed prompts that incentivize diversity to facilitate downstream generalization. In particular, we generate 200 descriptions per task by querying the model 10 times, each time asking for 20 descriptions given randomly sampled five question-answer pairs from the task. We leverage in-context learning by providing examples of question-answer pairs with matching descriptions. Finally, we also designed our prompts to avoid overly verbose responses and unnecessary information, such as explicit mentions of answer formats and additional instructions. We use the generated descriptions for the training and benchmark tasks. Figure 12 shows the exact prompt used for querying GPT-4o mini for task descriptions.

**System message**

You are a creative and helpful assistant.

**Prompt**

Given the following question-response pairs, please give a short description of the task describing what the task is.

{IN CONTEXT EXAMPLES}

Now, you must describe the task based on the following question-response pairs.

{5 sampled question-answer pairs}

Please use the information in the question-answer pairs and example description and come up with several descriptions that explain the task. Each description should be written in plain text, with the following format.

Description 1: DESCRIPTION\_1

Description 2: DESCRIPTION\_2

...

You should also be creative and vary the structure and the length of the descriptions such that they'll be diverse and cover various writing styles. You should ignore the specific question-answer pairs and focus on the high-level concept and topic of the task in general.

**\*\*DO NOT\*\*** describe that there are multiple choice options or the format of the answer.

**\*\*DO NOT\*\*** include the answer format, e.g., 'choose the correct option', 'answer with only one word', etc.

**\*\*DO NOT\*\*** describe how to answer the question, but rather what the task is about and the skills and knowledge required.

You can include reasoning steps that should be used to reach the expected answer.

Response with 20 descriptions. Use simple words and please be clear and diverse in your descriptions.

**In-context examples**

Here are some examples of the structure of the task of describing a task based on question-response pairs.

**## Example question-answer pair: 1**

**### Input**

You are given a question on high school macroeconomics. You are also given 4 answer options (associated with 'A', 'B', 'C', 'D'), out of which only one is correct. You need to answer the question by selecting the correct option. You should only answer with the choice letter, not the whole answer.

Input: Allocative efficiency (A) means that no inferior products will be produced. (B) implies that the economy's output is distributed evenly. (C) means that those who work hardest will get more. (D) implies that resources are used to produce the goods and services society desires in just the right amounts.

Output:

**### Expected output**

D

**### Plausible descriptions**

Description 1: Your job is to analyze the provided question about economics. Use your understanding of economic principles to guide your choice.

Description 2: Utilize your economic understanding to determine which choice is right. The correct answer will be the one that best aligns with economic principles.

**## Example question-answer pair: 2**

**### Input**

In this task, you are given a country name and you need to return the capital city of the given country.

Input: Senegal

Output:

**### Expected output**

Dakar

**### Plausible descriptions**

Description 1: Given the name of a country, your job is to provide its capital city.

Description 2: For each country listed, determine and state its capital city. This requires familiarity with global locations and capitals.

Figure 12: The prompt template used to query GPT-4o mini for task descriptions.

## M. Example of Task Descriptions

Here, we provide examples of task descriptions used in the experiments.

### Training descriptions

#### **sni\_cosmosqa\_passage\_inappropriate\_binary**

- Assess whether the given passage contains any elements that are unsuitable or illogical. Contextual understanding is key to making your evaluation.
- Look closely at the information provided in the context and determine its appropriateness or nonsensical nature based on logical reasoning.
- Assess given contexts critically, marking whether they hold inappropriate content or convey meaning in a way that is difficult to comprehend.

#### **sni\_winomt\_classification\_gender\_identifiability\_anti**

- In this task, you will distinguish between identifiable and unidentifiable gender references in sentences featuring different professions.
- Your task consists of evaluating professional descriptions within sentences and determining if their respective genders can be classified as clearly identifiable or obscure.
- Engage with sentences that present two different professions, paying attention to pronouns that could reveal or obscure the gender of the highlighted role.

#### **sni\_kth\_largest\_element**

- In this task, you are required to dissect a set of integers and identify which one corresponds to the kth position when sorted by size. Knowledge of ascending order and magnitude awareness are pivotal.
- Your mission here is to discover which number holds the kth place when considering size among others in a list. Practicing sorting and prioritization will be beneficial.
- The job is to pick out the kth greatest number from a list of integers, which means reevaluating them according to their increasing or decreasing order.

Figure 13: Examples of training descriptions from three SNI training tasks.



## Evaluation descriptions

### **boolq**

- Analyze the given details about various subjects, including movies, sports, and television shows. Your role is to confirm whether certain claims are true or false.
- Your task is to determine the truthfulness of specific statements based on the provided background information. This requires careful reading and comprehension of the content.
- The goal is to evaluate factual claims made in relation to highlighted texts. You will need to discern whether the statements align with the information provided.

### **gsm8k**

- You will be tasked with interpreting mathematical situations described in words. The goal is to use logical reasoning and calculations to determine the numerical answers based on the context provided.
- This task challenges your problem-solving abilities through mathematical reasoning. You must carefully read each scenario and systematically work through the data to compute the final outcome.
- Your role is to engage with practical math scenarios presented as questions. The task requires translating textual data into numerical operations that will lead you to the final solution.

### **humaneval**

- Engage in building distinct functions that meet the requirements of various presented problems, honing your ability to translate problem statements into logical code. Utilize structured thinking to implement efficient solutions.
- You are tasked with generating specific solutions in Python by interpreting problem descriptions associated with tasks like counting odds or validating inputs. Recognizing patterns and leveraging programming techniques will be beneficial.
- This task focuses on developing algorithms in Python for specific scenarios, such as counting characters, assessing conditions between numbers, or converting integers into a different format. Critical thinking and algorithmic design will be important.

Figure 14: Task descriptions of the benchmark tasks: boolq, gsm8k, and humaneval.

## Evaluation descriptions

### **mbpp**

- Your challenge is to solve a series of problems by writing functions in Python. These problems require handling lists and strings, allowing you to showcase your proficiency in coding while addressing practical programming scenarios.
- You will be tasked with creating various Python functions that tackle programming challenges. The exercises will test your ability to manipulate data structures, search for patterns, and implement checks on numerical products.
- The goal is to develop Python functions that perform designated operations on lists and strings. This requires a solid grasp of logical reasoning and the ability to apply relevant algorithms in your code.

### **winogrande**

- In this exercise, you need to read short narratives and discern which person or object fits best within the context of the sentence.
- This task requires synthesizing information from concise textual scenarios to identify crucial elements that drive the narrative forward.
- The goal is to evaluate descriptions and select the entity that best aligns with the sentiments or actions presented in the scenario.

### **piqa**

- You will explore practical questions and select an answer that presents a logical and widely accepted approach to solve a given problem or complete a task successfully.
- Analyze the provided scenarios where practical advice or solutions are required, focusing on selecting the most commonly used or convenient method.
- Given a question related to common tasks, your responsibility is to discern which proposed solution aligns with typical practices or makes the task easier to achieve.

Figure 15: Task descriptions of the benchmark tasks: mbpp, winogrande, piqa

## Evaluation descriptions

### hellaswag

- This task revolves around completing an unfinished text by selecting an ending that matches its tone and context. It requires you to think critically about how narratives develop and conclude effectively.
- This task asks you to select a suitable conclusion for an unfinished narrative or instructional content. It tests your comprehension and reasoning skills as you assess how well each option aligns with the given text.
- Your task involves completing an incomplete passage by selecting the ending that logically continues the context provided. This requires reading comprehension and the ability to infer meaning from a text.

### arc\_easy

- Your job is to discern which information best answers a posed question, focusing on practical examples and scientific principles. This requires a strong grasp of underlying concepts in ecology or physics.
- You will analyze questions that explore important connections such as environmental issues or animal adaptations. Utilize your background knowledge to evaluate and select the most fitting answer.
- This task involves selecting answers that reflect accurate relationships or effects seen in nature or society. You will need to sort through potential choices critically to find the appropriate one.

### arc\_challenge

- This task is about analyzing questions which examine your grasp of scientific ideas. You must connect conceptual knowledge with practical examples from geology, ecology and environmental changes.
- The objective here is to evaluate various scientific scenarios and infer the most logical explanations or definitions based on established knowledge. This task will strengthen your analytical and reasoning skills in the context of natural science.
- Your role is to interpret questions focusing on earth science and biological interactions. This demands a clear understanding of relevant processes, such as decomposition, weathering, and species adaptation.

Figure 16: Task descriptions of the benchmark tasks: hellaswag, arc\_easy, arc\_challenge

### Evaluation descriptions

#### openbookqa

- Analyze the provided statements carefully and determine which one best fits into the context of the passage. This requires comprehension skills and the ability to make logical inferences.
- Consider each option in relation to what is presented in the input. Discern which one logically completes or responds accurately to the notion being expressed.
- Here, you'll be presented with different statements, and your role is to decide which one appropriately complements or responds to a scenario. This process involves critical analysis and synthesis of information.

Figure 17: Task descriptions of the benchmark tasks: openbookqa

### Random descriptions

- dogs;cats;bananas;
- 7@9.qwepra#/.sd,s'2OC^039u#rdagjbl
- #####

Figure 18: Random descriptions

## N. Scaling Number of Descriptions per Task




Figure 19: Zero-shot benchmark performance of SFT-trained T2L with varying number of descriptions per training task.

Figure 19 shows mixed results on the benchmark performance when varying the number of descriptions per training task. For consistency, we always train T2L with 128 descriptions per training task.

```
Hypermod: HyperModulator(  
  (task_encoder): TaskEncoder(  
    (mlp): Sequential(  
      (0): Linear(in_features=1024, out_features=64, bias=True)  
      (1): LayerNorm((64,), eps=1e-05, elementwise_affine=True)  
    )  
  )  
  (layer_depth_encoder): Sequential(  
    (0): Embedding(32, 32)  
    (1): LayerNorm((32,), eps=1e-05, elementwise_affine=True)  
  )  
  (layer_type_encoder): Sequential(  
    (0): Embedding(2, 32)  
    (1): LayerNorm((32,), eps=1e-05, elementwise_affine=True)  
  )  
  (mixer): Sequential(  
    (0): Linear(in_features=128, out_features=512, bias=True)  
    (1): SiLU()  
    (2): Dropout(p=0.05, inplace=False)  
    (3): Linear(in_features=512, out_features=128, bias=True)  
    (4): SiLU()  
    (5): Dropout(p=0.05, inplace=False)  
  )  
  (mlp1): MLPResidualBlock(  
    (mlp): Sequential(  
      (0): LayerNorm((128,), eps=1e-05, elementwise_affine=True)  
      (1): Linear(in_features=128, out_features=512, bias=True)  
      (2): SiLU()  
      (3): Dropout(p=0.05, inplace=False)  
      (4): Linear(in_features=512, out_features=128, bias=True)  
      (5): SiLU()  
      (6): Dropout(p=0.05, inplace=False)  
    )  
  )  
  (mlp2): MLPResidualBlock(  
    (mlp): Sequential(  
      (0): LayerNorm((128,), eps=1e-05, elementwise_affine=True)  
      (1): Linear(in_features=128, out_features=512, bias=True)  
      (2): SiLU()  
      (3): Dropout(p=0.05, inplace=False)  
      (4): Linear(in_features=512, out_features=128, bias=True)  
      (5): SiLU()  
      (6): Dropout(p=0.05, inplace=False)  
    )  
  )  
  (mlp3): Sequential(  
    (0): LayerNorm((128,), eps=1e-05, elementwise_affine=True)  
    (1): Linear(in_features=128, out_features=512, bias=True)  
    (2): SiLU()  
    (3): Dropout(p=0.05, inplace=False)  
    (4): Linear(in_features=512, out_features=512, bias=True)  
    (5): SiLU()  
  )  
)
```

---

Listing 2: Detailed backbone architecture.



---

```
(AB_emb): ParameterDict(
  (q_proj): Object of type: ParameterDict
  (v_proj): Object of type: ParameterDict
  (q_proj): ParameterDict(
    (A): Parameter containing: [torch.cuda.FloatTensor of size 128]
    (B): Parameter containing: [torch.cuda.FloatTensor of size 128]
  )
  (v_proj): ParameterDict(
    (A): Parameter containing: [torch.cuda.FloatTensor of size 128]
    (B): Parameter containing: [torch.cuda.FloatTensor of size 128]
  )
)

(rank_emb): Sequential(
  (0): Embedding(8, 128)
  (1): LayerNorm((128,)), eps=1e-05, elementwise_affine=True
)
```

---

Listing 3: Detailed A/B and rank embedding of T2L.

---

```
OQA_TEMPLATE = (
  "Complete the following passage or answer the question by choosing the correct choice.\n\n"
  "{question_stem}\n\n"
  "{choices[label][0]}: {choices[text][0]}\n{choices[label][1]}: {choices[text][1]}\n"
  "{choices[label][2]}: {choices[text][2]}\n{choices[label][3]}: {choices[text][3]}\n\n"
  "You must respond with the letter corresponding to the correct choice (A,B,C,D) "
  "without any explanation."
)

ARC_TEMPLATE = (
  "Answer the question below by choosing the correct choice.\n\n"
  "{question}\n\n"
  "{choices[label][0]}: {choices[text][0]}\n{choices[label][1]}: {choices[text][1]}\n"
  "{choices[label][2]}: {choices[text][2]}\n{choices[label][3]}: {choices[text][3]}\n\n"
  "You must respond with the letter corresponding to the correct choice without any explanation."
)

HSWAG_TEMPLATE = (
  "You are provided with an incomplete passage below as well as 4 choices of continuation "
  "with only one of them being the correct ending. "
  "Treat the endings as being labelled 0, 1, 2, 3 in order.\n\n"
  "Passage: {ctx}\n\n"
  "0: {endings[0]}\n"
  "1: {endings[1]}\n"
  "2: {endings[2]}\n"
  "3: {endings[3]}\n\n"
  "You must respond with the only number corresponding to the correct ending (0,1,2,3) "
  "for the passage without any explanation."
)

PIQA_TEMPLATE = (
  "Choose the option that either answers the question, completes the sentence, "
  "or solves the problem. "
  "Pay attention to the properties of the objects in the question and how they interact with "
  "each other. "
  "If both options are correct, choose the one that is more convenient or more common. "
  "\n\n""{goal}""\n\n0: {sol1}\n1: {sol2}\n\n"
  "You must respond with either `0` or `1` without any explanation."
)

WINOGRANDE_TEMPLATE = (
  "Given the following situation:\n\n{sentence}\n\nWhich option is correct?\n\n"
  "Option 1: {option1}\nOption 2: {option2}\n\n"
  "You must respond with either `1` or `2` without any explanation."
)

BOOLQ_TEMPLATE = (
  "{passage}\n\nQuestion: {question}?\n\nPlease answer with either `true` or `false` "
  "without any explanation."
)
```

---

Listing 4: Instruction templates of QA-based benchmark tasks.