

000 LORAGEN: STRUCTURE-AWARE WEIGHT SPACE 001 002 LEARNING FOR LORA GENERATION 003 004

005 **Anonymous authors**

006 Paper under double-blind review

007 008 ABSTRACT 009

010 The widespread adoption of Low-Rank Adaptation (LoRA) for efficient fine-
011 tuning of large language models has created demand for scalable parameter gener-
012 ation methods that can synthesize adaptation weights directly from task descrip-
013 tions, avoiding costly task-specific training. We present LoRAGen, a structure-
014 aware method for generating LoRA parameters from natural language descrip-
015 tions. Through empirical analysis of LoRA libraries, we identify two key struc-
016 tural properties of LoRA parameter spaces: non-uniqueness of low-rank decom-
017 position and heterogeneous weight distributions across network modules. These
018 properties necessitate specialized parameter generation methods rather than gen-
019 eral weight space learning approaches. LoRAGen employs a latent diffusion
020 model with two innovations: weight-space supervision on full adaptation matri-
021 ces to handle decomposition non-uniqueness, and a module-aware Mix-of-Experts
022 decoder that adapts to module-specific weight distributions. Experiments show
023 LoRAGen achieves 96.0% performance relative to task-specific LoRAs on FLAN-
024 T5-large and 72.7% on Gemma-2-2B-Instruct for in-distribution tasks, while ob-
025 taining 40.2% on zero-shot generation across unseen tasks—surpassing baselines
026 by nearly 5%. Our work establishes the first structure-aware approach to LoRA
027 generation with insights into adaptation weight space geometry.

028 029 1 INTRODUCTION 030

031 Exploring the weight space of neural networks, i.e., the high-dimensional parameter space spanned
032 by populations of trained networks, has emerged as a powerful paradigm for understanding model
033 mechanisms and enabling novel applications. Parameter generation, which trains models to produce
034 weights for target networks, represents a particularly promising direction that has gained significant
035 attention in recent years (Schürholt et al., 2021b; Schürholt et al., 2022; Schürholt et al., 2024a;
036 Wang et al., 2024). The rise of large language models (LLMs) has created opportunities to apply
037 parameter generation techniques in this domain, particularly through Low-Rank Adaptation (LoRA)
038 (Hu et al., 2022) generation, the direct synthesis of LoRA parameters for efficient fine-tuning. While
039 traditional LoRA workflows require task-specific training with custom datasets and hyperparam-
040 eters, creating engineering overhead and limiting reusability (He et al., 2022a; Lv et al., 2024), LoRA
041 generation enables direct parameter synthesis from natural language task descriptions, improving
042 scalability and unlocking adaptive model behavior without maintaining extensive adapter libraries.

043 Recent work generates LoRA parameters conditioned on task identifiers, datasets, or natural lan-
044 guage task descriptions. A common design involves learning a hypernetwork—a neural network
045 that generates parameters for another base network (Ha et al., 2016). With advances in deep gener-
046 ative models, one category of approaches learns a lower-dimensional representation directly from
047 the weight space and decodes this latent representation into LoRA parameters (Shao et al., 2025b).
048 However, the underlying encoder-decoder model must encode the entire LoRA parameters at once
049 into the learned latent representation, which limits the size of LoRA that can be embedded. Another
050 category learns a conditional diffusion prior over the latent space to generate LoRA parameters from
051 random noise based on specific task conditions at test time, but these approaches struggle to generate
052 well-performing LoRA parameters across diverse architectures and datasets (Jin et al., 2024; Liao
053 et al., 2024; Wu et al., 2024; Soro et al., 2025). A recent work, Text-to-LoRA (T2L) (Charakorn
054 et al., 2025c), employs a hypernetwork trained to construct LoRA parameters in a single inexpensive
055 forward pass, enabling zero-shot generation for entirely unseen tasks based on the hypothesis that




Figure 1: **Two empirical observations.** *Left:* Task description embedding similarity versus LoRA parameter similarity (measured by full adaptation matrix and low-rank decomposition matrix, respectively). *Right:* Spectral-entropy distributions of LoRA parameters for FLAN-T5-large, grouped by module type: encoder self-attention, decoder self-attention, and cross-attention.

different LoRAs share the same adaptation mechanism and can be optimized without explicit structure. However, these methods treat LoRA parameter generation as an instantiation of general weight space learning approaches Schürholt et al. (2024a); Wang et al. (2024), while lacking a tailored design specifically for LoRA characteristics.

In this work, we begin with an empirical analysis of a LoRA library built on the Transformer-based model FLAN-T5-large(Chung et al., 2024) (Figure 1). Our analysis reveals the **non-uniqueness of low-rank decomposition** as a key challenge. As shown in Figure 1 (left), we observe a clear positive correlation between adapter similarity in the weight space and task description embedding similarity, but we find no correlation between the cosine similarity of low-rank decomposition matrices and task description embedding similarity. This suggests that supervision in the full adaptation matrix space should generalize better than element-wise reconstruction. However, most LoRA generators reconstruct low-rank decomposition matrices directly (Ha et al., 2016; Jin et al., 2024; Liao et al., 2024; Soro et al., 2025; Charakorn et al., 2025c; Lv et al., 2024), which can make training sensitive to arbitrary rescalings or rotations of the low-rank decomposition matrices that produce the same full adaptation matrix (Gabrielsson et al., 2024). Additionally, in Figure 1 (right), we identify **significant heterogeneity in LoRA weight distributions across different modules**. The spectral-entropy (Yunis et al., 2024; Roy & Vetterli, 2007) distributions differ systematically by module type within FLAN-T5-large: encoder self-attention exhibits higher entropy, decoder self-attention shows the lowest entropy, and cross-attention lies in between, indicating heterogeneous weight distributions across parts of the base model. However, existing methods typically use a single decoder across modules (Lv et al., 2024; Charakorn et al., 2025c), which overlooks the module structure (Ostapenko et al., 2024) and can mismatch the spectral-entropy distributions.

Motivated by the two empirical observations that reveal distinct properties of the LoRA weight space, we introduce *LoRAGen*, a method that generates LoRA parameters from natural-language task descriptions for structure-aware LoRA weight space learning. Specifically, we use a LoRA weight autoencoder (LAE) to learn a latent space over LoRA parameters and a conditional latent diffusion model conditioned on natural language task descriptions to generate latent representations from random noise, followed by a decoder that processes these generated representations to produce new LoRA parameters. To address the non-uniqueness issue, we supervise the full adaptation matrix directly rather than the low-rank decomposition matrices, proposing two weight space loss terms: a direction loss and a spectral loss on the full adaptation matrix. This avoids sensitivity to rescalings or rotations of low-rank decomposition matrices that yield the same full adaptation matrix, leading to more consistent and task-relevant weight space learning. To address the heterogeneity of weight distributions, we introduce a module-aware Mix-of-Experts (MoE) decoder with routing that uses a structural embedding combining a latent variable with learnable module and layer embeddings. This allows experts to specialize to the observed module-specific weight distribution patterns while ensuring a controlled sharing mechanism via the chosen expert pool configuration, facilitating generalization across different architectures of the base model.

108 Our key contributions are: (1) We propose the first structure-aware weight learning method tailored
 109 to the LoRA weight space, enabling effective LoRA parameter generation for diverse downstream
 110 tasks. (2) LoRAGen introduces weight-space losses on the full adaptation matrix to address non-
 111 uniqueness of low-rank decomposition and employs a module-aware MoE decoder to model het-
 112 erogeneous LoRA weight distributions. (3) LoRAGen achieves strong in-distribution LoRA gen-
 113 eration performance close to task-specific LoRAs across architectures: 96.0% on FLAN-T5-large,
 114 72.7% on Gemma-2-2B-Instruct and further achieves 40.2% in zero-shot LoRA generation on
 115 seven unseen tasks, surpassing competitive baselines by nearly 5%.

2 RELATED WORK

120 **Weight Space Learning for Parameter Generation.** Research on weight space learning has fol-
 121 lowed two main directions: predicting model properties from trained weights (Schürholt et al.,
 122 2021c; Unterthiner et al., 2020) and generating new parameters for neural networks (Eilertsen et al.,
 123 2020; Schürholt et al., 2022; 2024b). Here, we focus on parameter generation. With advances in
 124 deep generative models, one category of approaches learns a latent representation directly over pop-
 125 ulations of trained models and decodes this latent representation to generate parameters (Berardi
 126 et al., 2022). Another category employs conditional latent diffusion models to generate target pa-
 127 rameters (Peebles et al., 2022; Soro et al., 2024). Recent studies introduce graph-based encoders that
 128 treat networks as graphs and enable parameter generation across architectures (Kofinas et al., 2024).
 129 These approaches have been applied in areas such as meta-learning, transfer learning, and model
 130 compression (Finn et al., 2017; Nava et al., 2022; Wang et al., 2023). Overall, these approaches
 131 exploring parameter generation in weight space (Peebles et al., 2022; Berardi et al., 2022; Schürholt
 et al., 2024b) demonstrate the feasibility of parameter generation across architectures and datasets.

132 **Hypernetworks for LoRA Generation.** A hypernetwork is a neural network designed to gener-
 133 ate the parameters of another “base” network, providing an indirect encoding of the base model’s
 134 weights (Ha et al., 2016; Zhang et al., 2018; Schug et al., 2025). With the development of deep
 135 generative models, one category of approaches (Shao et al., 2025a) learns conditional latent repre-
 136 sentations of pretrained LoRA parameters for new LoRA parameter generation. Another category
 137 employs conditional latent diffusion priors over latent space (Wu et al., 2024; Jin et al., 2024; Soro
 138 et al., 2024), enabling the generation of task-specific LoRA parameters. Recently, hyperLoRA (Lv
 139 et al., 2024) employs instruction-tuned hypernetworks with constrained loss and demo selection to
 140 produce stable and generalizable adapters. While these approaches have advanced multi-task LLM
 141 adaptation, they typically rely on learned task identifiers, limiting their capacity for zero-shot LoRA
 142 parameter generation to unseen tasks (Ivison & Peters, 2022; Mahabadi et al., 2021; He et al., 2022b;
 143 Schürholt et al., 2021a; Ortiz-Barajas et al., 2024; Lv et al., 2024). Recent work explores natural lan-
 144 guage as conditioning signals for zero-shot generation (Xiao et al., 2023; Ivison et al., 2023; Phang
 145 et al., 2023), with T2L (Charakorn et al., 2025a), DnD (Liang et al., 2025), and LoRA-Gen (Xiao
 146 et al., 2025) utilizing hypernetworks to generate LoRA adapters from textual prompts. However,
 147 existing methods treat LoRA generation as a general weight space learning problem, overlooking
 148 the unique structural properties of LoRA parameter spaces. In contrast, LoRAGen is the first weight
 149 space learning approach that specifically accounts for the structural characteristics of LoRA spaces,
 leading to more effective parameter generation.

3 PRELIMINARIES

154 **Low-Rank Adaptation (Hu et al., 2022):** Low-rank matrix ΔW serves as a adapter to a base
 155 model. For a pretrained weight matrix W_0 , the fine-tuned linear transformation is given by $h =$
 156 $W_0 + \Delta W = W_0 + BA$, where $A \in \mathbb{R}^{r \times d_{in}}$ and $B \in \mathbb{R}^{d_{out} \times r}$ are low-rank decomposition matrices
 157 with $r \ll \min\{d_{in}, d_{out}\}$. We ignore the module type and layer index of the LoRA parameters
 158 when referring to all LoRA parameters. Therefore we index them by a module type $m \in \mathcal{M}$
 159 (e.g., query projection) and a layer index $\ell \in \{1, \dots, L_m\}$. A LoRA adapter at each location
 160 (m, ℓ) specifies a low-rank matrix $\Delta W_{m, \ell} = B_{m, \ell} A_{m, \ell}$. We denote the low-rank adapter by
 161 $\Delta W := \{\Delta W_{m, \ell}\}_{m \in \mathcal{M}, 1 \leq \ell \leq L_m}$. Note that the low-rank decomposition is *not unique*: for any
 invertible matrix $R \in \mathbb{R}^{r_{m, \ell} \times r_{m, \ell}}$, $(B_{m, \ell} R)(R^{-1} A_{m, \ell})$ produces the same low-rank matrix $\Delta W_{m, \ell}$.

162 **Problem setting.** We assume an LoRA library of pairs $\mathcal{D} = \{(x^{(i)}, \Delta W^{(i)})\}_{i=1}^N$, where $x^{(i)}$ is
 163 a natural language description of task $t^{(i)}$ and $\Delta W^{(i)} = \{\Delta W_{m,\ell}^{(i)}\}$ represents the fine-tuned low-
 164 rank adapter for task $t^{(i)}$. Our goal is to train a LoRA generator using \mathcal{D} that produces new LoRA
 165 parameters $\widehat{\Delta W}'$ given a natural language task description x' , where x' may be either in-distribution
 166 (from \mathcal{D}) or out-of-distribution (unseen tasks).
 167

169 4 THE PROPOSED METHOD: LORAGEN

171 4.1 DESIGN PRINCIPLES FROM EMPIRICAL OBSERVATIONS

173 We first summarize two **observations** (Obs) over a library of low-rank adapters (two panels in Figure
 174 1), which motivates our tailored design in LoRAGen. The details of observations experiment
 175 implementation is reported at the Appendix A.1.

177 **Obs-1: Non-uniqueness of the low-rank decomposition.** Here we focus on the non-uniqueness
 178 property of low-rank decompositions in LORA. Specifically, we examine a LoRA library trained
 179 on FLAN-T5-large. While ΔW is uniquely defined as a full adaptation matrix, its low-rank
 180 decomposition matrices (A, B) is not unique, illustrated in Sec. 3. This property motivates the
 181 following hypothesis: if we supervise the low-rank decomposition matrices (A, B) directly, like
 182 element-wise reconstruction, training becomes sensitive to arbitrary rescalings and rotations that
 183 still yield the same full adaptation matrix ΔW . By supervising the full adaptation matrix itself, we
 184 avoid this ambiguity and directly align the generated adapters with the pretrained adapters in the
 185 full adaptation matrix space, which should lead to more consistent and task-relevant weight space
 186 learning.

187 We test this hypothesis by computing the *pairwise* similarity between its LoRA adapter and each of
 188 the other 111 adapters in the FLAN subset for a representative task which is selected from the four
 189 major task categories in the FLAN subset so that it reflects the overall distribution of the dataset
 190 rather than any specific task; further details are reported in Appendix A.1. To measure adapter simi-
 191 larity, we compute the cosine similarity of the concatenation of flattened low-rank A and B matrices
 192 of all layers and flattened ΔW , respectively. To avoid scale mismatch between the two similarity
 193 measures, they are plotted in separate subplots, each fitted with its own least-squares trend line. We
 194 observe a clear positive correlation between the task embedding similarity and the adapter similarity
 195 in the weight space, whereas the similarity measured on the low-rank decomposition matrices ex-
 196 hibits near-zero Spearman coefficients ρ , indicating the lack of correlation. This phenomenon aligns
 197 with the non-uniqueness property of low-rank decompositions in LoRA and suggests that supervi-
 198 sion in the full adaptation matrix space should generalize better than element-wise reconstruction.

199 Motivated by Obs-1, we therefore introduce *adapter-level supervision* (Sec. 4.3), where losses are
 200 defined directly in the weight space of ΔW . Concretely, we introduce two weight space loss terms,
 201 combining a *direction loss* \mathcal{L}_{ang} (Eq. 1) that aligns normalized LoRA directions, with a *spectral loss*
 202 $\mathcal{L}_{\text{spec}}$ (Eq. 2) that matches leading singular values. These objectives enforce task-consistent LoRA
 203 generation while remaining robust to the inherent non-uniqueness of low-rank decompositions.

204 **Obs-2: Heterogeneity of LoRA weight distributions.** Here we focus on the heterogeneity
 205 of LORA weight distributions across different module types in a transformer-based architecture,
 206 specifically the FLAN-T5-large model, which serves as the base model for LoRA adapters.
 207 For each adapter $\Delta W_{m,\ell}$ at a module-layer location (m, ℓ) , we analyze how its Frobenius en-
 208 ergy is distributed across singular directions. Let $\{\sigma_i\}$ denote the singular values of $\Delta W_{m,\ell}$
 209 and define the normalized spectrum $p_i = \sigma_i^2 / \sum_j \sigma_j^2$. We then compute the *spectral entropy*
 210 $H_{\text{spec}}(\Delta W_{m,\ell}) = -\sum_i p_i \log p_i$, which quantifies the uniformity of energy over directions: low
 211 entropy indicates that the energy concentrates in a few dominant directions, indicating effectively
 212 lower rank structure, whereas high entropy corresponds to a more even spread of energy (Yunis
 213 et al., 2024; Roy & Vetterli, 2007).

214 For **FLAN-T5-Large**, we reports the spectral-entropy distributions for three module types in
 215 FLAN-T5-large: Encoder self-attention, Decode self-attention, and *Cross-Attention*. We ob-
 216 serve systematic differences across modules: encoder adapters exhibit higher spectral entropy which




Figure 2: **Overall framework of LoRAGen.** Our approach consists of two stages: First, we train hypernetwork based on LoRA weight autoencoder to encode and reconstruct LoRA parameters, and diffusion process conditioned on natural language task descriptions to predict denoised latent. Second, random noise and un/seen natural language task descriptions are fed into LoRAGen to generate LoRA parameters, which can be incorporated with the LLM to evaluate downstream tasks.

means energy spread over more directions, decoder self-attention shows the lowest entropy with more concentrated and effectively lower-rank structure, and cross-attention lies in between. This pattern implies that LoRA weight distributions are not homogeneous across the network which means the way energy is distributed over singular directions differs consistently by module type. [More results are reported in Appendix A.3.](#)

Motivated by Obs-2, we employ a *module-aware MoE decoder* (Sec. 4.4) to match these module-specific patterns of energy distributions. Specifically, routing is conditioned on a structural embedding that combines a latent variable with learnable module and layer embeddings. The decoder can operate with either (i) a single shared expert pool for all locations *or* (ii) separate expert pools per module type (e.g., encoder attention). Per-module output heads \mathcal{H}_m map the gated expert outputs into LoRA parameters $\widehat{\Delta W}_{m,\ell}$. This architecture lets experts specialize to the observed module-specific energy distribution patterns (e.g., lower-entropy/low-rank tendencies in Decoder-Self versus higher-entropy Encoder), while the router and the chosen pool configuration ensures controlled sharing mechanism without losing module-specific specialization.

4.2 METHOD OVERVIEW

As shown in Figure 2, the overall framework of *LoRAGen* can be divided into two stages: First, we train a hypernetwork, which consists of LoRA weight autoencoder (LAE) and diffusion process, to learn the inner structure of LoRA parameters in the LoRA weight space. For LAE($\mathcal{E}_\phi, \mathcal{D}_\theta$) trained on given LoRA parameters ΔW , the encoder \mathcal{E}_ϕ produces per-location latents with a diagonal-Gaussian posterior and the MoE decoder \mathcal{D}_θ (Sec. 4.4) decodes latent z to full LoRA parameters $\widehat{\Delta W}$. Note that at training LAE stage, we add two weight space losses terms defined on $\Delta W_{m,\ell}$ (Sec. 4.3). In the diffusion process, the noised latent is processed by a diffusion transformer conditioned on embeddings of natural language task descriptions to predict denoised latent. The details of diffusion process is reported in the [Appendix A.11](#).

Second, we feed task descriptions and random noise into the diffusion transformer to produce the output denoised latents followed by passing by the frozen MoE Decoder from stage-1 to generate the full LoRA parameters, which can be applied to the LLM to perform the intended task.

270 4.3 ADAPTER-LEVEL SUPERVISION IN LORA WEIGHT SPACE
271

272 Distinct pairs of low-rank decomposition matrices (A, B) produce the same low-rank adapter
273 ΔW . If we choose to generate A and B separately, the LoRA generator must to commit to
274 a specific decomposition of ΔW , even though many different low-rank decomposition matrices
275 produce the identical adapter. To avoid this ambiguity induced by non-uniqueness property of
276 the low-rank decomposition, we directly add supervision signal at the level of low-rank adapter
277 $\widehat{\Delta W}_{m,\ell} = \mathcal{D}_\theta(z)_{m,\ell}$, and introduce two weight space loss terms: (i) a *direction* loss based on
278 cosine similarity that depends on the low-rank adapters after normalizing their Frobenius norm. (ii)
279 a *spectral* loss that matches the leading singular values of the predicted and pretrained low-rank
280 adapters.

281 **Direction loss.** Because many pairs of low-rank decomposition matrices (A, B) produce the same
282 low-rank adapter ΔW , supervising A and B individually can introduce arbitrary differences in the
283 Frobenius norm of ΔW . To make supervision insensitive to it, we normalize both the predicted and
284 target low-rank adapters to unit Frobenius norm and compare their directions directly to capture the
285 per-task direction. Thus we introduce a *direction* loss that measures the angular mismatch between
286 the predicted and target low-rank adapters:

$$287 \mathcal{L}_{\text{ang}}(m, \ell) = 1 - \frac{\langle \widehat{\Delta W}_{m,\ell}, \Delta W_{m,\ell} \rangle_F}{\|\widehat{\Delta W}_{m,\ell}\|_F \|\Delta W_{m,\ell}\|_F}. \quad (1)$$

290 where $\langle \cdot, \cdot \rangle_F$ denotes the Frobenius inner product and $\|\cdot\|_F$ is the Frobenius norm.

291 **Spectral loss.** Direction supervision does not capture how the squared Frobenius norm is dis-
292 tributed across the *singular spectrum*. Two low-rank adapters can have similar cosine similarity
293 yet differ in their *leading singular values*, i.e., in the proportion of the squared Frobenius norm ex-
294 plained by the top singular directions. To account for this, we introduce a *spectral* loss that matches
295 the leading singular values of the two low-rank adapters:

$$296 \mathcal{L}_{\text{spec}}(m, \ell) = \left\| \boldsymbol{\sigma}_{1:k_{m,\ell}}(\widehat{\Delta W}_{m,\ell}) - \boldsymbol{\sigma}_{1:k_{m,\ell}}(\Delta W_{m,\ell}) \right\|_{p, \boldsymbol{\omega}_{m,\ell}} \quad (2)$$

297 where (i) $\sigma_i(X)$ denotes the i -th singular value of X , and $\boldsymbol{\sigma}_{1:k}(X) := (\sigma_1(X), \dots, \sigma_k(X))$ lists
298 the top- k values in nonincreasing order; (ii) $r_{m,\ell}$ is the LoRA rank at location (m, ℓ) , and $k_{m,\ell} \in$
299 $\{1, \dots, r_{m,\ell}\}$ is the smallest integer such that the top- k singular values of the target low-rank adapter
300 explain at least a fraction $\rho \in (0, 1)$ of its squared Frobenius norm. (iii) $\|u\|_{p,\omega}$ is a weighted ℓ_p
301 norm with $p \in \{1, 2\}$; (iv) $\boldsymbol{\omega}_{m,\ell} = (\omega_{m,\ell,1}, \dots, \omega_{m,\ell,k_{m,\ell}})$ are nonnegative normalized singular-
302 value weights.

303 Thus, we aggregate these two terms together as follows:

$$304 \mathcal{L}_{\text{adapter}}(\theta, \phi) = \mathbb{E}_{z \sim q_\phi(z | \Delta W)} \left[\sum_{m \in \mathcal{M}} \sum_{\ell=1}^{L_m} \lambda_{m,\ell} \left(\alpha_{\text{ang}} \mathcal{L}_{\text{ang}}(m, \ell) + \alpha_{\text{spec}} \mathcal{L}_{\text{spec}}(m, \ell) \right) \right], \quad (3)$$

305 where $q_\phi(z | \Delta W)$ is the encoder posterior; $\alpha_{\text{ang}}, \alpha_{\text{spec}} > 0$ are hyperparameters balancing the two
306 loss terms; and $\lambda_{m,\ell} \geq 0$ are location weights. Note that \mathcal{L}_{ang} aligns direction, while $\mathcal{L}_{\text{spec}}$ aligns the
307 leading spectrum (i.e., the proportion of the squared Frobenius norm explained by the top singular
308 directions). Thus $\mathcal{L}_{\text{adapter}}$ aligns the performance of the specific-task low-rank adapter yet remains
309 robust to low-rank decompositions.

310 The overall objective of training LAE is

$$311 \mathcal{L}_{\text{VAE}}(\theta, \phi) = \alpha_{\text{adapter}} \mathcal{L}_{\text{adapter}}(\theta, \phi) + \beta D_{\text{KL}}(q_\phi(z | \Delta W) \| \mathcal{N}(0, I)) + \lambda_{\text{moe}} \mathcal{L}_{\text{moe}}(\theta), \quad (4)$$

312 where $\alpha_{\text{adapter}}, \beta, \lambda_{\text{moe}} > 0$ are scalar coefficients; \mathcal{L}_{moe} is the MoE load-balancing auxiliary
313 loss (Sec. 4.4).

314 4.4 MODULE-AWARE MOE DECODER
315

316 To capture structural heterogeneity across the module types and layers, we further introduce a
317 module-aware MoE decoder \mathcal{D}_θ that can use either a single shared expert pool for all locations
318 or separate pools per module type (e.g., encoder attention). The notation below treats E as the
319 number of experts in the active pool used by the current location.

324 **Inputs and routing.** For each location (m, ℓ) , we form a structural embedding $h_{m,\ell} =$
 325 $[z_{m,\ell}; e_m; e_\ell] \in \mathbb{R}^{d_h}$, where $z_{m,\ell} \in \mathbb{R}^{d_z}$ is the latent variable, $e_m \in \mathbb{R}^{d_m}$ and $e_\ell \in \mathbb{R}^{d_\ell}$ are
 326 learnable module and layer embeddings. A router with parameters $W_r \in \mathbb{R}^{E \times d_h}$ outputs logits
 327 $\ell_{m,\ell} = W_r h_{m,\ell} \in \mathbb{R}^E$ and applies top- K gating:
 328

$$g_{(m,\ell),e} = \frac{\exp(\ell_{m,\ell,e}/\tau)}{\sum_{e' \in S_{m,\ell}} \exp(\ell_{m,\ell,e'}/\tau)} \mathbb{I}[e \in S_{m,\ell}], \quad (5)$$

332 where $\tau > 0$ is the temperature, $S_{m,\ell} \subset \{1, \dots, E\}$ is the index set of the top- K experts by logit
 333 value, and $\mathbb{I}[\cdot]$ is the indicator function (equal to 1 if its argument is true and 0 otherwise).
 334

335 **Experts and per-module heads.** Each expert \mathcal{E}_e is a small MLP. The gated output feeds a per-
 336 module head \mathcal{H}_m :

$$\widehat{\Delta W}_{m,\ell} = \mathcal{H}_m \left(\sum_{e \in S_{m,\ell}} g_{(m,\ell),e} \mathcal{E}_e(h_{m,\ell}) \right), \quad (6)$$

339 Note that \mathcal{H}_m is a linear map into a vector followed by a reshape to the full LoRA parameters
 340 $\widehat{\Delta W}_{m,\ell} \in \mathbb{R}^{d_{\text{out}}(m,\ell) \times d_{\text{in}}(m,\ell)}$. And its parameters are shared across all layers ℓ for the same
 341 module m .
 342

343 **Load-balancing auxiliary loss.** To discourage expert collapse we introduce a load-balancing aux-
 344 ietary loss:
 345

$$\mathcal{L}_{\text{moe}} = \max \left(E \sum_{e=1}^E \bar{p}_e \bar{f}_e - 1, 0 \right), \quad (7)$$

346 where \bar{p}_e is the average gating probability, and \bar{f}_e is the expected fractional load under top- K rout-
 347 ing. The details implementation of \mathcal{L}_{moe} is reported at the Appendix A.11.
 348

351 5 EXPERIMENTS

353 5.1 EXPERIMENTAL SETUP

355 **Dataset.** In our main experiments, we consider two settings. First, we employ FLAN-T5-large
 356 (Chung et al., 2024), as the base model. We utilize a subset of FLAN following (Lv et al., 2024)
 357 for training and evaluation. Second, we use Gemma-2-2b-Instruct (Team et al., 2024) as the
 358 base model and evaluate on 8 widely used benchmark tasks, including Arc-challenge (ArcC) and
 359 Arc-easy (ArcE) (Clark et al., 2018), BoolQ (Clark et al., 2019), GSM8K (Cobbe et al., 2021),
 360 Hellaswag (HS) (Zellers et al., 2019), OpenBookQA (OQA) (Mihaylov et al., 2018), PIQA (Bisk
 361 et al., 2020), and Winogrande (WG) (Sakaguchi et al., 2021). More details about the datasets are
 362 reported in the Appendix. For both settings, we extract task embeddings from natural language
 363 task descriptions using the FLAN-T5-large encoder. Task descriptions for each dataset are fully
 364 generated by LLM, as described in the Appendix A.10. For each dataset, We sample and report the
 365 average performance of 3 set of LoRA weights sampled with LoRAGen.
 366

367 **Baseline Setup.** As baselines, we consider task-specific LoRAs, [weight-averaged LoRA \(Wortsman
 368 et al., 2022\)](#). We compare D2NWG (Soro et al., 2024), which is a latent diffusion conditioned
 369 on datasets for LoRAs generation. We also implement T2L (Charakorn et al., 2025b), which is
 370 a hypernetwork that generates LoRAs based on natural language task embedding. Reproduction
 371 details are provided in the Appendix A.11.

372 5.2 PERFORMANCE COMPARISON

373 **In-distribution LoRA Generation Performance.** First, we focus on whether LoRAGen can re-
 374 cover the performance of trained LoRAs ([Charakorn et al., 2025c; Brüel-Gabrielsson et al., 2024](#)),
 375 which enables low-rank adaptation with minimal compute requirements. Table 1 reports results on
 376 seven FLAN tasks using natural-language task embeddings from the FLAN-T5-Large encoder.
 377 LoRAGen closely matches the oracle adapters and outperforms D2NWG and T2L on average per-
 378 formance . We think that the gain comes from the design of MoE decoder to capture the heterogeneity

Method	AP-Neg	AP-Rec	AP-Pos	QASC-1	QASC-2	WQA-T	WQA-A	Avg. (acc)
FLAN-T5-Large	49.4	70.7	34.8	23.3	8.9	8.5	62.0	36.8
Average LoRA	96.8	96.5	96.9	99.4	87.3	96.7	97.0	95.8
D2NWG	59.5	87.3	65.9	47.0	31.1	33.4	84.8	58.4
T2L	90.5	94.1	92.7	87.9	76.8	85.5	93.3	88.7
Ours	96.8	96.6	97.1	99.5	87.3	97.1	97.3	96.0
Task-specific LoRAs	97.2	97.0	97.8	99.7	87.3	97.0	97.3	96.2

Table 1: Benchmark performance of LoRAGen on FLAN subset (FLAN-T5-Large backbone). **Bold numbers** are used when the performance is higher than the task-specific LoRAs.

Method	ArcC	ArcE	BQ	GSM8K	HS	OQA	PIQA	WG	Avg. (acc)
Gemma-2-2B-Instruct	74.0	89.9	81.0	55.9	55.1	71.2	71.2	51.8	68.8
Average LoRA	75.6	90.1	81.4	56.1	56.5	73.8	72.8	53.9	70.0
D2NWG	74.1	90.0	81.2	56.0	55.1	71.3	71.3	52.0	68.9
T2L	74.3	90.2	81.2	55.9	55.2	71.4	71.5	53.8	69.2
Ours	76.6	90.7	84.1	56.4	64.1	80.2	75.0	54.2	72.7
Task-specific LoRAs	76.7	90.6	84.7	55.9	75.4	80.2	78.0	54.6	74.5

Table 2: Benchmark performance of LoRAGen on 8 benchmark tasks (Gemma-2-2B-Instruct backbone). **Bold numbers** are used when the performance is higher than the task-specific LoRAs.

of weight distributions across different module types and layers within the structure-aware LoRA weight space.

Moreover, we train a separate LoRAGen model for a decoder-only base model (Gemma-2-2B-Instruct) to assess whether the same structure-aware design generalizes across architectures. As shown in Table 2, LoRAGen remains competitive or superior to several baselines and is close to task-specific LoRAs on average performance, indicating that our approach scales from T5-based adapters to decoder-only adapters. In several tasks, such as ArcE, GSM8K and OQA, our method even matches or surpasses task-specific adapters, suggesting that adapter-level supervision captures task-relevant structure rather than memorizing particular LoRA parameters.

LoRA Generation for unseen tasks. Furthermore, we explore whether LoRAGen can generate LoRA parameters for unseen tasks. We train LoRAGen on 136 tasks of FLAN subset, each with 20 task descriptions. For each task we sample three sets of LoRA weights and report the average accuracy. As shown in Table 3, LoRAGen achieves the best average accuracy 40.2, outperforming D2NWG and T2L by **+5.2** and **+5.0** points, respectively. We observe that D2NWG and T2L reconstructs pre-trained adapters and struggles to generalize to unseen tasks. This phenomenon is align with non-uniqueness of the low-rank decomposition, which indicates that if we supervise the low-rank decomposition matrices directly, training becomes sensitive to arbitrary rescalings and rotations that still yield the same full adaptation matrix, trending to memorize task-specific LoRA parameters. Instead of this, our method supervise the full adaptation matrix directly to avoid this ambiguity and focus on learn task-relevant weight space learning, resulting in better performance to unseen tasks. Details on computational cost and efficiency are reported in the Appendix A.4.

5.3 ABLATION STUDY

In this section, we ablate the adapter-level supervision and the module-aware MoE decoder on the seven FLAN tasks (Table 4). Training with the two adapter-level losses \mathcal{L}_{ang} and $\mathcal{L}_{\text{spec}}$ but without the decoder \mathcal{D}_{θ} achieves an average accuracy of 58.4. In contrast, enabling the decoder while removing both adapter-level losses and using only the reconstruction loss results in a significant improvement to an average accuracy of 95.2. This improvement is consistent with **Obs. 2**: the module-aware routing and per-module heads in \mathcal{D}_{θ} effectively capture the heterogeneity in weight distributions across modules and layers, which substantially improves downstream performance.

Method	AP-Rec	AP-Pos	QASC-1	QASC-2	WQA-T	WQA-A	Avg. (acc)
FLAN-T5-Large	70.7	34.8	23.3	8.9	8.5	62.0	34.7
D2NWG	71.8	34.4	23.7	9.1	8.3	62.4	35.0
T2L	71.6	34.6	23.3	11.1	8.5	62.0	35.2
Ours	75.1	42.2	28.1	14.5	14.3	67.2	40.2

Table 3: Zero-shot generation performance of LoRAGen trained on the **FLAN subset** with FLAN-T5-Large as the base model. **Bold numbers** are used to represent the best performance.

\mathcal{L}_{ang}	$\mathcal{L}_{\text{spec}}$	\mathcal{D}_{θ}	AP-Neg	AP-Rec	AP-Pos	QASC-1	QASC-2	WQA-T	WQA-A	Avg. (acc)
✓	✓	✗	59.5	87.3	65.9	47.1	31.2	33.4	84.9	58.4
✗	✗	✓	96.8	95.5	97.0	98.2	86.3	95.6	96.7	95.2
✗	✓	✓	49.5	70.9	34.4	23.3	9.1	8.4	62.7	36.9
✓	✓	✓	96.8	96.6	97.1	99.5	87.3	97.1	97.3	96.0

Table 4: Ablation study on FLAN subset. Checkmarks indicate enabled components: direction loss \mathcal{L}_{ang} , spectral loss $\mathcal{L}_{\text{spec}}$, and module-aware MoE decoder \mathcal{D}_{θ} . **Bold numbers** are used to represent the best performance.

A counter-intuitive result occurs when \mathcal{D}_{θ} is combined with the *spectral* loss but the *direction* loss is omitted: the average accuracy drops to 36.9. In this case, $\mathcal{L}_{\text{spec}}$ only enforces alignment of the *magnitudes* of the top- k singular values, without constraining the corresponding directions of the left and right singular vectors. As a result, the decoder can match the singular-value magnitudes correctly, but assign them to the wrong directions, which leads to an incorrect weight-space learning and consequently poor task performance. In the zero-shot setting reported in Appendix A.2, Table 5, the same combination yields a much larger gain of +3.6 points in average accuracy (from 36.6 to 40.2), again with the largest improvements on QASC-1, QASC-2, and WQA-T, confirming that adapter-level supervision is particularly important for generalization to unseen tasks.

5.4 DETAILED ANALYSIS

Hyperparameter Analysis. We first analyze the effect of the spectral–energy threshold ρ in the spectral loss (Figure 8(a) in Appendix A.6). We conduct this experiment under the same setting as Table 1. Across all tasks, accuracies remain stable as ρ varies from 0.80 to 1.00, with only minor fluctuations. This shows that the adapter-level supervision is robust to the choice of ρ , since the leading singular values already capture sufficient spectral information. In practice, we set $\rho = 0.85$ as it provides a better performance while maintaining stability.

We then examine the hyperparameters of the MoE decoder (Figure 8(b) in Appendix A.6). Here we compare shared vs. unshared expert pools, different top- K routing choices, and the number of experts. The results indicate that unshared pools consistently outperform shared ones, and increasing the number of active experts (top-2 vs. top-1) further improves performance, especially on challenging multi-choice and QA tasks. The best configuration is *unshared, top-1, E = 4*, which strikes a balance between accuracy and computational efficiency.

Structural Embedding Analysis. In this section, we assess the contribution of the structural embedding in the MoE decoder (Figure 8(c) in Appendix A.6). We compare three variants: (i) without structural embedding (routing only on latent variables), (ii) with structural embedding, and (iii) an oracle trained with task-specific adapters. The results show that removing structural embeddings substantially reduces accuracy, particularly on tasks requiring fine-grained reasoning such as QASC-1/2. Adding structural embeddings closes most of the gap to the oracle, confirming that encoding module-specific latent is critical for capturing the heterogeneous LoRA weight distributions.

6 CONCLUSION

We presented *LoRAGen*, a structure-aware approach to LoRA parameter generation grounded in two empirical observations of the LoRA weight space: non-uniqueness of low-rank decompositions

486 and module-wise heterogeneity of weight distributions. Motivated by these observations, we su-
 487 pervise the full adaptation matrix using adapter-level direction and spectral losses, and decode with
 488 a module-aware MoE whose routing leverages structural embeddings with shared or per-module
 489 expert pools. Empirically, LoRAGen produces strong in-distribution adapters across architectures,
 490 closely matching task-specific adapters on FLAN-T5-large and Gemma-2-2B-Instruct and
 491 attains competitive zero-shot performance on unseen tasks from a large LoRA library. Ablation
 492 studies show both adapter-level supervision and module-aware decoding are necessary, and sensitivity
 493 studies indicate robustness to the spectral-energy threshold and gains from unshared pools with
 494 top- K routing.

496 ETHICS STATEMENT

497
 498 The authors have read and adhered to the ICLR Code of Ethics. The research presented in this
 499 paper focuses on the generation of LoRA parameters, with primary applications in LLM domains.
 500 All data used for training and evaluation is from publicly available, non-personal scientific datasets,
 501 ensuring no privacy concerns. This work does not involve human subjects, and we do not foresee
 502 any direct negative societal impacts or risks of perpetuating social biases. Our aim is to advance the
 503 development of domain-specific LLM applications.

504 REPRODUCIBILITY STATEMENT

505
 506 The code associated with this paper is available at: <https://anonymous.4open.science/r/LoRAGen-02C0>. It includes the necessary environment configurations and execution scripts.
 507 All datasets utilized in this work are publicly accessible. The task descriptions generated from LLM
 508 is provided in the Appendix.

512 REFERENCES

513
 514 Gabriele Berardi, Luca De Luigi, Samuele Salti, and Luigi Di Stefano. Learning the space of deep
 515 models. *arXiv preprint arXiv:2206.07680*, 2022.

516 Yonatan Bisk, Rowan Zellers, Jianfeng Gao, Yejin Choi, et al. Piqa: Reasoning about physical com-
 517 monsense in natural language. In *Proceedings of the AAAI conference on artificial intelligence*,
 518 volume 34, pp. 7432–7439, 2020.

519
 520 Rasmus Brüel-Gabrielsson, Jingzhao Zhu, Omkar Bhardwaj, Leshem Choshen, Kristjan Gree-
 521 newald, Mikhail Yurochkin, and Justin Solomon. Compress then serve: Serving thousands of
 522 lora adapters with little overhead. *arXiv preprint arXiv:2407.00066*, 2024.

523
 524 Rujikorn Charakorn, Edoardo Cetin, Yujin Tang, and Robert T Lange. Text-to-lora: Instant trans-
 525 former adaption. In *Proceedings of the 42nd International Conference on Machine Learning*.
 526 PMLR, 2025a.

527 Rujikorn Charakorn, Edoardo Cetin, Yujin Tang, and Robert Tjarko Lange. Text-to-lora: Instant
 528 transformer adaption. *arXiv preprint arXiv:2506.06105*, 2025b.

529
 530 Rujikorn Charakorn, Edoardo Cetin, Yujin Tang, and Robert Tjarko Lange. Text-to-lora: Instant
 531 transformer adaption. In *Proceedings of the 42nd International Conference on Machine Learn-
 532 ing (ICML 2025), Poster Track*, 2025c. URL <https://openreview.net/forum?id=zWskCdu3QA>. Poster.

533
 534 Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Yunxuan Li,
 535 Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, et al. Scaling instruction-finetuned lan-
 536 guage models. *Journal of Machine Learning Research*, 25(70):1–53, 2024.

537
 538 Christopher Clark, Kenton Lee, Ming-Wei Chang, Tom Kwiatkowski, Michael Collins, and Kristina
 539 Toutanova. Boolq: Exploring the surprising difficulty of natural yes/no questions. *arXiv preprint
 arXiv:1905.10044*, 2019.

540 Peter Clark, Isaac Cowhey, Oren Etzioni, Tushar Khot, Ashish Sabharwal, Carissa Schoenick, and
 541 Oyvind Tafjord. Think you have solved question answering? try arc, the ai2 reasoning challenge.
 542 *arXiv preprint arXiv:1803.05457*, 2018.

543

544 Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser,
 545 Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, et al. Training verifiers to
 546 solve math word problems. *arXiv preprint arXiv:2110.14168*, 2021.

547

548 Gabriel Eilertsen, Daniel Jönsson, Timo Ropinski, Jonas Unger, and Anders Ynnerman. Classifying
 549 the classifier: dissecting the weight space of neural networks. *arXiv preprint arXiv:2002.05688*,
 550 2020.

551

552 Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation
 553 of deep networks. In *Proceedings of the 34th International Conference on Machine Learning*, 2017.

554

555 Rickard Brüel Gabrielsson, Jiacheng Zhu, Onkar Bhardwaj, Leshem Choshen, Kristjan Gree-
 556 newald, Mikhail Yurochkin, and Justin Solomon. Compress then serve: Serving thousands of
 557 lora adapters with little overhead, 2024. URL <https://openreview.net/forum?id=hHNVn4hFPk>.

558

559 David Ha, Andrew Dai, and Quoc V Le. Hypernetworks. *arXiv preprint arXiv:1609.09106*, 2016.

560

561 Junxian He, Chunting Zhou, Xuezhe Ma, Taylor Berg-Kirkpatrick, and Graham Neubig. Towards
 562 a unified view of parameter-efficient transfer learning. In *International Conference on Learning
 Representations*, 2022a. URL <https://openreview.net/forum?id=0RDcd5Axok>.

563

564 Yelong He, Shuang Zheng, Yi Tay, Jatin Gupta, Yichong Du, Vamsi Aribandi, Zhen Zhao, Yichong
 565 Li, Zhiyu Chen, Donald Metzler, et al. Hyperprompt: Prompt-based task-conditioning of trans-
 566 formers. In *International Conference on Machine Learning*, pp. 8678–8690. PMLR, 2022b.

567

568 Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang,
 569 Lu Wang, and Weizhu Chen. Lora: Low-rank adaptation of large language models. In
 570 *ICLR 2022*, April 2022. URL <https://www.microsoft.com/en-us/research/publication/lora-low-rank-adaptation-of-large-language-models/>.

571

572 Hamish Ivison and Matthew E. Peters. Hyperdecoders: Instance-specific decoders for multi-task
 573 nlp. In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pp. 1715–1730,
 574 2022. URL <https://aclanthology.org/2022.findings-emnlp.124>.

575

576 Hamish Ivison, Akshita Bhagia, Yizhong Wang, Hannaneh Hajishirzi, and Matthew E Peters. Hint:
 577 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.

578

579 Zhihao Jin et al. Conditional lora parameter generation. *arXiv preprint arXiv:2403.12345*, 2024.

580

581 Marios Kofinas, Boris Knyazev, Yulun Zhang, Yiming Chen, Gertjan Burghouts, Efstratios Gavves,
 582 Cees GM Snoek, and Dingwen Zhang. Graph neural networks for learning equivariant represen-
 583 tations of neural networks. In *International Conference on Learning Representations*, 2024.

584

585 Zhiyuan Liang, Dongwen Tang, Yuhao Zhou, Xuanlei Zhao, Mingjia Shi, Wangbo Zhao, Zekai
 586 Li, Peihao Wang, Konstantin Schürholt, Damian Borth, et al. Drag-and-drop llms: Zero-shot
 587 prompt-to-weights. *arXiv preprint arXiv:2506.16406*, 2025.

588

589 Huanxuan Liao, Shizhu He, Yao Xu, Yuanzhe Zhang, Yanchao Hao, Shengping Liu, Kang Liu,
 590 and Jun Zhao. From instance training to instruction learning: Task adapters generation from
 591 instructions. *Advances in Neural Information Processing Systems*, 37:45552–45577, 2024.

592

593 Chuancheng Lv, Lei Li, Shitou Zhang, Gang Chen, Fanchao Qi, Ningyu Zhang, and Hai-Tao Zheng.
 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.

594 Rabeeh Karimi Mahabadi, Sebastian Ruder, Mostafa Dehghani, and James Henderson. Parameter-
 595 efficient multi-task fine-tuning for transformers via shared hypernetworks. *arXiv preprint*
 596 *arXiv:2106.04489*, 2021.

597 Todor Mihaylov, Peter Clark, Tushar Khot, and Ashish Sabharwal. Can a suit of armor conduct
 598 electricity? a new dataset for open book question answering. *arXiv preprint arXiv:1809.02789*,
 599 2018.

600 Elia Nava, Shogo Kobayashi, Yufei Yin, Raffaello K Katzschmann, and Benjamin F Grawe. Meta-
 601 learning via classifier(-free) diffusion guidance. *arXiv preprint arXiv:2210.05639*, 2022.

602 Juan-Gilberto Ortiz-Barajas, Helena Gomez-Adorno, and Thamar Solorio. Hyperloader: Integrating
 603 hypernetwork-based lora and adapter layers into multi-task transformers for sequence labelling.
 604 *arXiv preprint arXiv:2407.01411*, 2024.

605 Oleksiy Ostapenko, Zhan Su, Edoardo Ponti, Laurent Charlin, Nicolas Le Roux, Lucas Caccia,
 606 and Alessandro Sordoni. Towards modular llms by building and reusing a library of loras. In
 607 *Proceedings of the 41st International Conference on Machine Learning (ICML)*, volume 235, pp.
 608 38885–38904. PMLR, 2024.

609 William Peebles, Ilija Radosavovic, Tim Brooks, Alexei A Efros, and Jitendra Malik. Learning to
 610 learn with generative models of neural network checkpoints. *arXiv preprint arXiv:2209.12892*,
 611 2022.

612 Jason Phang, Yuning Mao, Pengcheng He, and Weizhu Chen. Hypertuning: Toward adapting large
 613 language models without backpropagation. In *International Conference on Machine Learning*,
 614 pp. 27854–27875. PMLR, 2023.

615 Olivier Roy and Martin Vetterli. The effective rank: A measure of effective dimensionality. In *2007*
 616 *15th European signal processing conference*, pp. 606–610. IEEE, 2007.

617 Keisuke Sakaguchi, Ronan Le Bras, Chandra Bhagavatula, and Yejin Choi. Winogrande: An adver-
 618 sarial winograd schema challenge at scale. *Communications of the ACM*, 64(9):99–106, 2021.

619 Simon Schug, Seijin Kobayashi, Yassir Akram, Joao Sacramento, and Razvan Pascanu. Attention as
 620 a hypernetwork. In *The Thirteenth International Conference on Learning Representations*, 2025.
 621 URL <https://openreview.net/forum?id=V4K9h1qNxE>.

622 Konstantin Schürholt, Mostafa Dehghani, Neil Houlsby, and Damian Borth. Hyper-representations:
 623 Learning generative representations of neural network parameters. In *International Conference*
 624 *on Learning Representations (ICLR)*, 2021a. URL <https://openreview.net/forum?id=HCSgyPUfeDj>.

625 Konstantin Schürholt, Dimche Kostadinov, and Damian Borth. Hyper-representations: Self-
 626 supervised representation learning on neural network weights for model characteristic prediction.
 627 2021b. URL <https://api.semanticscholar.org/CorpusID:240070334>.

628 Konstantin Schürholt, Dimche Kostadinov, and Damian Borth. Self-supervised representation learning
 629 on neural network weights for model characteristic prediction. *Advances in Neural Informa-
 630 tion Processing Systems*, 34:16481–16493, 2021c.

631 Konstantin Schürholt, Boris Knyazev, Xavier Giró-i Nieto, and Damian Borth. Hyper-
 632 representations as generative models: sampling unseen neural network weights. In *Proceedings*
 633 *of the 36th International Conference on Neural Information Processing Systems*, NIPS ’22, Red
 634 Hook, NY, USA, 2022. Curran Associates Inc. ISBN 9781713871088.

635 Konstantin Schürholt, Boris Knyazev, Xavier Giró-i Nieto, and Damian Borth. Hyper-
 636 representations for pre-training and transfer learning. *arXiv preprint arXiv:2207.10951*, 2022.

637 Konstantin Schürholt, Michael W. Mahoney, and Damian Borth. Towards scalable and versa-
 638 tile weight space learning. In *Proceedings of the 41st International Conference on Machine*
 639 *Learning*, pp. 43947–43966, 2024a. URL <https://proceedings.mlr.press/v235/schurholt24a.html>.

648 Konstantin Schürholt, Michael W Mahoney, and Damian Borth. Towards scalable and versatile
 649 weight space learning. *arXiv preprint arXiv:2406.09997*, 2024b.
 650

651 Jing Shao et al. In-context meta lora generation. *arXiv preprint arXiv:2502.08976*, 2025a.
 652

653 Yihua Shao, Minxi Yan, Yang Liu, Siyu Chen, Wenjie Chen, Xinwei Long, Ziyang Yan, Lei
 654 Li, Chenyu Zhang, Nicu Sebe, et al. In-context meta lora generation. *arXiv preprint
 655 arXiv:2501.17635*, 2025b.
 656

656 Bedionita Soro, Bruno Andreis, Hayeon Lee, Wonyong Jeong, Song Chong, Frank Hutter,
 657 and Sung Ju Hwang. Diffusion-based neural network weights generation. *arXiv preprint
 658 arXiv:2402.18153*, 2024.
 659

659 Bedionita Soro, Bruno Andreis, Hayeon Lee, Wonyong Jeong, Song Chong, Frank Hutter, and
 660 Sung Ju Hwang. Diffusion-based neural network weights generation. In *The Thirteenth Interna-
 661 tional Conference on Learning Representations*, 2025. URL [https://openreview.net/
 662 forum?id=j8WHjM9aMm](https://openreview.net/forum?id=j8WHjM9aMm).
 663

664 Gemma Team, Morgane Riviere, Shreya Pathak, Pier Giuseppe Sessa, Cassidy Hardin, Surya Bhu-
 665 patiraju, Léonard Hussenot, Thomas Mesnard, Bobak Shahriari, Alexandre Ramé, et al. Gemma
 666 2: Improving open language models at a practical size. *arXiv preprint arXiv:2408.00118*, 2024.
 667

667 Thomas Unterthiner, Daniel Keysers, Sylvain Gelly, Olivier Bousquet, and Ilya Tolstikhin. Predict-
 668 ing neural network accuracy from weights. *arXiv preprint arXiv:2002.11448*, 2020.
 669

670 Jun Wang, Yifan Chen, Stella X Yu, Brian Cheung, and Yann LeCun. Compact and optimal deep
 671 learning with recurrent parameter generators. In *Proceedings of the IEEE/CVF Winter Conference
 672 on Applications of Computer Vision*, 2023.
 673

673 Kaitian Wang, Zhaopan Xu, Yukun Zhou, Zelin Zang, Trevor Darrell, Zhuang Liu, and Yang
 674 You. Neural network diffusion, 2024. URL <https://openreview.net/forum?id=8Q6UmFhhQS>.
 675

676 Mitchell Wortsman, Gabriel Ilharco, Samir Ya Gadre, Rebecca Roelofs, Raphael Gontijo-Lopes,
 677 Ari S Morcos, Hongseok Namkoong, Ali Farhadi, Yair Carmon, Simon Kornblith, et al. Model
 678 soups: averaging weights of multiple fine-tuned models improves accuracy without increasing
 679 inference time. In *International conference on machine learning*, pp. 23965–23998. PMLR, 2022.
 680

681 Yujia Wu, Yiming Shi, Jiwei Wei, Chengwei Sun, Yang Yang, and Heng Tao Shen. Difflora: Gener-
 682 ating personalized low-rank adaptation weights with diffusion. *arXiv preprint arXiv:2408.06740*,
 683 2024.
 684

684 Yicheng Xiao, Lin Song, Rui Yang, Cheng Cheng, Yixiao Ge, Xiu Li, and Ying Shan. Lora-gen:
 685 Specializing large language model via online lora generation. *arXiv preprint arXiv:2506.11638*,
 686 2025.
 687

687 Zhaofeng Xiao, William Held, Yutong Liu, and Diyi Yang. Task-agnostic low-rank adapters for
 688 unseen english dialects. In *Proceedings of the 2023 Conference on Empirical Methods in Natural
 689 Language Processing*, pp. 7857–7870. Association for Computational Linguistics, 2023. doi:
 690 10.18653/v1/2023.emnlp-main.487.
 691

692 David Yunis, Kumar Kshitij Patel, Samuel Wheeler, Pedro Savarese, Gal Vardi, Karen Livescu,
 693 Michael Maire, and Matthew R Walter. Approaching deep learning through the spectral dynamics
 694 of weights. *arXiv preprint arXiv:2408.11804*, 2024.
 695

695 Rowan Zellers, Ari Holtzman, Yonatan Bisk, Ali Farhadi, and Yejin Choi. Hellaswag: Can a ma-
 696 chine really finish your sentence? *arXiv preprint arXiv:1905.07830*, 2019.
 697

698 Chris Zhang, Mengye Ren, and Raquel Urtasun. Graph hypernetworks for neural architecture
 699 search. *arXiv preprint arXiv:1810.05749*, 2018.
 700

701

\mathcal{L}_{ang}	$\mathcal{L}_{\text{spec}}$	\mathcal{D}_{θ}	AP-Rec	AP-Pos	QASC-1	QASC-2	WQA-T	WQA-A	Avg. (acc)
✗	✗	✓	72.7	35.9	24.5	10.4	10.2	65.9	36.6
✓	✓	✓	75.1	42.2	28.1	14.5	14.3	67.2	40.2

Table 5: Zero-shot ablation study on the FLAN subset. Checkmarks indicate enabled components: direction loss \mathcal{L}_{ang} , spectral loss $\mathcal{L}_{\text{spec}}$, and module-aware MoE decoder \mathcal{D}_{θ} . **Bold numbers** represent the best performance.

A APPENDIX

A.1 DETAILS FOR OBS-1

As shown in Figure 3, we provide more 9 representative tasks to compare with the remaining 111 tasks on the FLAN subset. These representative tasks are selected so that they cover the three major task categories present in the FLAN subset: Natural Language Understanding (NLU, 98 tasks), Natural Language Generation (NLG, 9 tasks), and Knowledge & Information Extraction (5 tasks). Specifically, we include five NLU tasks (wikiqa_tpqap, wikiqa_jeop, fix_punct, true_case, word_segment), three NLG tasks (duorc_tgen, duorc_gqba, gem_e2e_nlg), and two Knowledge & Information Extraction tasks (wiki_bio_what, wiki_bio_comp). This ensures that the visualization reflects the diversity of the dataset rather than any specific task.

Concretely, for a representative task r and another task t , we define the *task similarity* as the cosine similarity between their text-description embeddings $e(r)$ and $e(t)$:

$$s_{\text{task}}(r, t) = \frac{e(r)^\top e(t)}{\|e(r)\|_2 \|e(t)\|_2}. \quad (8)$$

For adapter similarity, we consider two variants. Let ΔW_r and ΔW_t be the full adaptation matrices for tasks r and t , and let A_r, B_r and A_t, B_t be their corresponding low-rank factors. We compute

$$s_{\Delta W}(r, t) = \frac{\langle \text{vec}(\Delta W_r), \text{vec}(\Delta W_t) \rangle}{\|\text{vec}(\Delta W_r)\|_2 \|\text{vec}(\Delta W_t)\|_2}, \quad (9)$$

and

$$s_{(A, B)}(r, t) = \frac{\langle v_r, v_t \rangle}{\|v_r\|_2 \|v_t\|_2}, \quad v_t = \begin{bmatrix} \text{vec}(A_t) \\ \text{vec}(B_t) \end{bmatrix}. \quad (10)$$

We then report the Spearman correlation coefficient ρ between $\{s_{\text{task}}(r, t)\}$ and each of $\{s_{\Delta W}(r, t)\}$ and $\{s_{(A, B)}(r, t)\}$ over all t .

We found that the full adaptation matrix similarity shows a consistent positive correlation with task description embedding similarity. This suggests that tasks with more semantically similar descriptions tend to produce more similar LoRAs, highlighting the presence of structure in the LoRA weight space. This observation also emphasizes the utility of the full adaptation space, as it allows us to capture and exploit underlying similarities between tasks that may not be immediately apparent from the low-rank decomposition matrices. By designing weight space loss based on this similarity, we can better adapt models to generate new LoRA parameters.

A.2 ABLATION STUDY ON ZERO-SHOT GENERATION

We further examine how the adapter-level losses affect zero-shot generalization on the seven FLAN tasks that are used in Table 1. As summarized in Table 5, enabling only the MoE decoder with a reconstruction loss (i.e., using \mathcal{D}_{θ} without \mathcal{L}_{ang} or $\mathcal{L}_{\text{spec}}$) achieves an average zero-shot accuracy of 36.6. Adding both direction and spectral losses on top of the decoder increases the average accuracy to 40.2 (+3.6 points), with particularly gains on QASC-1 (+3.6), QASC-2 (+4.1), and WQA-T (+4.1). These tasks are among the most challenging reasoning benchmarks in our evaluation, so the improvements provide direct evidence that adapter-level supervision is especially beneficial for zero-shot generalization.




Figure 3: **Relation between LoRA similarity and task description embedding similarity.** Each panel shows the similarity between a representative task adapter and the other 111 adapters trained on a FLAN subset y -axis in the weight space against their similarity in the task embeddings space x -axis. LoRA cosine similarity is measured in two ways: (i) computing cosine on the low-rank decomposition matrices A and B separately (blue dots); (ii) computing cosine on the full adaptation matrix $\Delta W = AB$ (orange triangles). Legends report the Spearman correlation coefficient ρ .

A.3 WEIGHT DISTRIBUTION ANALYSIS


To further investigate the effect of the module-aware MoE decoder, we analyze the weight distributions of task-specific LoRAs and the adapters generated by LoRAGen on the same evaluation tasks as in the main tables. For the encoder-decoder base model **FLAN-T5-Large**, we use the seven FLAN tasks reported in Table 1. For the decoder-only base model **Gemma-2-2B-Instruct**, we use the eight benchmark tasks reported in Table 2. In each case, we compare three variants: (i) task-specific (oracle) LoRAs, (ii) LoRAGen (full model), and (iii) LoRAGen without the MoE decoder (a single shared decoder).

Spectral entropy distribution analysis. For each adapter $\Delta W_{m,\ell}$ we compute its spectral entropy $H_{\text{spec}}(\Delta W_{m,\ell})$. For **FLAN-T5-Large**, Figure 4 shows the distributions of H_{spec} grouped by module type (Encoder, Decoder-Self, Cross-Attention) for the three variants above. Task-specific

810 adapters exhibit clear heterogeneity: encoder modules have the highest spectral entropy, decoder
 811 self-attention the lowest, and cross-attention lies in between. The adapters generated by the full
 812 LoRAGen model faithfully capture this pattern, indicating that the MoE decoder has learned module-
 813 specific spectral profiles. In contrast, removing the MoE decoder collapses the three module types
 814 into almost identical high-entropy distributions, suggesting that a single shared decoder fails to cap-
 815 ture the heterogeneous energy patterns across modules.

816 For **Gemma-2-2B-Instruct**, Figure 5 plots the spectral entropy distributions of decoder
 817 adapters for oracle LoRAs, LoRAGen, and LoRAGen without MoE. Task-specific LoRAs concen-
 818 trate around relatively high entropy values, and the full LoRAGen model shows a very similar,
 819 slightly higher-entropy profile. In contrast, removing the MoE decoder shifts towards lower entropy
 820 and yields a much broader distribution, indicating that the shared decoder fails to match the spectral
 821 distribution of task-specific LoRAs.

822
 823
 824 **Layer-wise cosine similarity analysis.** Beyond spectral statistics, we also study the alignment
 825 between generated and task-specific adapters. Figures 6 and 7 report, for each layer and mod-
 826 ule type, the average cosine similarity between the generated adapter and the corresponding task-
 827 specific adapter, averaged over the seven FLAN tasks in Table 1 and the eight benchmark tasks in
 828 Table 2, respectively. With the full LoRAGen model, cosine similarities are consistently positive
 829 for **FLAN-T5-Large** and **Gemma-2-2B-Instruct**, indicating that the experts capture mean-
 830 ingful, layer-specific adaptation patterns. In contrast, the model without the MoE decoder yields
 831 similarities that remain close to zero across almost all layers, showing that it fails to recover the
 832 fine-grained per-layer structure even though it can still perform reasonable reconstruction on in-
 833 distribution tasks. Thus, the spectral-entropy and cosine-similarity analyses confirm that the MoE
 834 decoder learns to specialize to module types and their characteristic spectral patterns.



845 Figure 4: **Spectral entropy distributions on FLAN-T5-Large.** Each group of violins corresponds
 846 to Encoder, Decoder-Self, and Cross-Attention modules, respectively, for task-specific LoRAs (left),
 847 LoRAGen (middle), and LoRAGen without the MoE decoder (right).

848 A.4 EFFICIENCY OF ADAPTER-LEVEL SUPERVISION

851 A natural concern is that supervising adapters directly (i.e. losses defined on $\Delta W = AB^\top \in \mathbb{R}^{d \times d}$)
 852 could be more expensive than reconstruction on $A, B \in \mathbb{R}^{d \times r}$. We clarify that our implementation
 853 avoids any $\mathcal{O}(d^2)$ complexity by using a quadratic form, and thus remains comparable in cost to
 854 reconstruction loss.

855 **Direction loss.** The cosine similarity between two adapters can be computed efficiently as a
 856 quadratic form:
 857

$$\langle \Delta W_1, \Delta W_2 \rangle_F = \text{tr}(A_1 B_1^\top B_2 A_2^\top) = \langle A_1, (B_1^\top B_2) A_2^\top \rangle_F,$$

858 with $\|\Delta W\|_F^2 = \text{tr}(A^\top A B^\top B)$. This requires only $d \times r$ and $r \times r$ multiplications, yielding
 859 complexity $\mathcal{O}(dr^2)$ instead of $\mathcal{O}(d^2)$. Hence the direction loss is no more expensive than element-
 860 wise reconstruction.

861
 862 **Spectral loss.** For the spectral loss, we never compute an SVD of the full $d \times d$ adapter. Instead,
 863 we compute a reduced QR decomposition. This reduces the problem to an $r \times r$ core matrix $K =$




Figure 5: Spectral entropy distributions on Gemma-2-2B-Instruct for Decoder-Self module. Task-specific LoRAs (left), LoRAGen (middle), and LoRAGen without the MoE decoder (right).




Figure 6: Layer-wise cosine similarity between generated and task-specific LoRA adapters on the FLAN subset with FLAN-T5-Large as the base model. For each module type (rows) and layer index (x-axis), we plot the average cosine similarity over 7 evaluation tasks. Left: LoRAGen (full model); right: LoRAGen without the MoE decoder.

$R_A R_B^\top$, on which we perform SVD. The resulting complexity is $\mathcal{O}(dr^2 + r^3)$, avoiding any $\mathcal{O}(d^2)$ cost.

Complexity comparison. Table 6 reports the per-layer complexities. Both direction and spectral losses are implemented in quadratic or QR decomposition form, avoiding explicit $\mathcal{O}(d^2)$ cost. The overhead relative to element-wise reconstruction loss is bounded by a factor of r , which is small in practice.

A.5 COMPUTATIONAL COST AND RESOURCE USAGE

We will quantify the practical computational cost of LoRAGen and compare it with the baseline T2L. All experiments for FLAN-T5-Large and Gemma-2-2B-Instruct are run on a single NVIDIA A40 GPU (40GB).

Training time. For the in-distribution experiments in Tables 1 and 2, Stage-1 training for T2L and LoRAGen takes about 2 and 2.5 hours, respectively. For the zero-shot experiments in Table 3, the total training time of both methods is less than one day on the same A40 GPU.




Figure 7: Layer-wise cosine similarity between generated and task-specific LoRA adapters on 8 benchmark tasks with Gemma-2-2B-Instruct as the base model. Left: LoRAGen (full model); right: LoRAGen without the MoE decoder.

Loss term	Per-layer complexity
A, B reconstruction	$\mathcal{O}(dr)$
Direction (quadratic)	$\mathcal{O}(dr^2)$
Spectral (QR+SVD core)	$\mathcal{O}(dr^2 + r^3)$

Table 6: Per-layer complexity of different adapter-level supervision terms. Both direction and spectral losses avoid explicit $\mathcal{O}(d^2)$ cost.

Generator parameter counts. For FLAN-T5-Large, T2L and LoRAGen (Stage 1) have comparable generator sizes (50M vs. 53M), while LoRAGen already achieves better zero-shot performance. For Gemma-2-2B-Instruct, we intentionally use a larger generator for LoRAGen (69M) than for T2L (30M) so that the MoE decoder has enough capacity to generate all four types of LoRA weights, namely qa, qb, va and vb. Notably, the activated parameters of LoRAGen are 15M and 22M per forward pass, respectively.

Memory usage. During LoRA weight autoencoder (LAE) training, enabling the module-aware MoE decoder introduces about 1000 MiB of additional GPU memory, and the adapter-level losses introduce about 300 MiB. In practice, the complete LoRAGen training requires roughly 2500 MiB of extra memory for these components.

A.6 ADDITIONAL RESULTS ON HYPERPARAMETERS AND STRUCTURAL EMBEDDING

The results on hyperparameters and structural embedding analysis are shown in Figure 8.




Figure 8: Performance of LoRAGen with different hyperparameters and removing structural embedding. Here, for the middle figure, s / uns denote *shared* or *unshared* MoE experts; $\mathbf{t1}$ / $\mathbf{t2}$ denote *top-1* or *top-2* routing strategies; $\mathbf{e4}$ / $\mathbf{e8}$ denote using 4 or 8 experts.

972

973

Table 7: Stage 1: LoRA autoencoder configuration for FLAN-T5-Large.

974

975

976

977

978

979

980

981

982

983

984

985

986

987

988

989

990

991

992

993

994

995

996

997

998

999

1000

1001

1002

1003

1004

1005

1006

1007

1008

1009

1010

1011

1012

1013

1014

1015

1016

1017

1018

1019

Item	Configuration
Backbone	FLAN-T5-Large (encoder-decoder)
Encoder hidden size	$d_{\text{model}} = 1024$
LoRA rank	$r = 16$
Latent dimension	$d_z = 64$ (per module)
Encoder MLP hidden size	128
Decoder type	Module-aware MoE, 3 heads (enc, dec-self, cross)
Experts per head	$E = 4$
Active experts	top- $K = 1$ per location
Expert pools	Separate pool per module type (no sharing)
Expert MLP input	$2d_z = 128$ (latent + structural embedding)
Expert MLP hidden size	256 (SiLU activation)
Expert MLP output size	$r \cdot d_{\text{model}} = 16,384$
Gating	Softmax, temperature $\tau = 1.5$
MoE aux loss	Load-balancing, weight 10^{-5}
Encoder params	$\approx 2.1\text{M}$
Decoder (MoE) params	$\approx 50.9\text{M}$
Total generator params	$\approx 53\text{M}$ ($\approx 7\%$ of backbone)

996

997

998

999

1000

1001

1002

1003

1004

1005

1006

1007

1008

1009

1010

1011

1012

1013

1014

1015

1016

1017

1018

1019

1020

1021

1022

1023

1024

1025

Table 8: Stage 1: LoRA autoencoder configuration for Gemma-2-2B-Instruct.

Item	Configuration
Backbone	Gemma-2-2B-Instruct (decoder-only)
Encoder hidden size	$d_{\text{model}} = 2304$
LoRA rank	$r = 8$
Latent dimension	$d_z = 64$ (per module)
Encoder MLP hidden size	128
Decoder type	4-head MoE for q_a, q_b, v_a, v_b
q_a output size	$r \times d_{\text{in}} = 8 \times 2304 = 18,432$
q_b output size	$d_{q,\text{out}} \times r = 2048 \times 8 = 16,384$
v_a output size	18,432
v_b output size	$d_{v,\text{out}} \times r = 1024 \times 8 = 8192$
Experts per head	$E = 4$
Active experts	top- $K = 1$ per location
Expert pools	Separate pool per head
Expert MLP input	$2d_z = 128$
Expert MLP hidden size	256 (SiLU activation)
Expert MLP output	Matching head-specific output sizes above
Encoder params	$\approx 5.6\text{M}$
Decoder (MoE) params	$\approx 63.7\text{M}$
Total generator params	$\approx 69.3\text{M}$ ($\approx 3\text{--}4\%$ of backbone)

A.7 END-TO-END INFERENCE LATENCY WITH GENERATED ADAPTERS

Once a LoRA is generated, we simply load the weights into the base model; at inference time the runtime is the same as for any standard LoRA fine-tuned model. On an NVIDIA H800, measuring real evaluation runs, we observe the end-to-end latencies per forward pass shown in Table 11.

The small differences are within standard runtime variability; there is no additional per-token overhead induced by LoRAGen itself. All extra cost is incurred once per task when sampling a LoRA adapter set, as discussed in Appendix A.4.

A.8 SCALABILITY ANALYSIS

1026

1027

Table 9: Stage 1: losses and training schedule for both backbones.

1028

1029

1030

1031

1032

1033

1034

1035

1036

1037

1038

1039

Item	Configuration
Adapter-level losses	Direction loss L_{dir} and spectral loss L_{spec} on ΔW
Loss weights	1.0 for L_{dir} , 2.0 for L_{spec}
Optimizer	AdamW
Learning rate	1.115×10^{-3}
Gradient accumulation	Factor 8
Max epochs	1000
Complexity details	See Appendix A.4

1040

1041

1042

1043

1044

1045

1046

1047

1048

1049

1050

1051

1052

1053

1054

1055

1056

Table 10: Stage 2: latent diffusion over adapter latents.

Item	Configuration
Latent sequence length (FLAN-T5)	$L = 288$ (one token per LoRA location)
Latent sequence length (Gemma-2-2B)	$L = 96$ (one token per LoRA location)
Denoiser architecture	1D Transformer
Denoiser layers	4
Model dimension	256
Input/output dimension	$d_z = 64$
Conditioning	Task embedding + module/layer embeddings
Diffusion steps	$T = 500$
Noise schedule	Linear β schedule
Objective	v -prediction with ℓ_2 loss
Optimizer	AdamW
Learning rate	8×10^{-5}
Batch size	32
Training steps	100k

In this section, we evaluate the scalability of LoRAGen with respect to both the number of LoRA locations and the LoRA rank. In the experiments, we define the **total generation latency** as the combined time taken for three parts: (1) **Encoding (Enc.)**: The time taken to encode the input task description and map it to the latent space. (2) **Diffusion (Diff.)**: The time taken to perform diffusion sampling over the latent space. (3) **Decoding (Dec.)**: The time taken to decode the latent vector back into the LoRA parameters.

Scaling with the number of locations. We evaluate the performance of LoRAGen as the number of LoRA locations increases. We performed scaling tests on Gemma-2-2B (rank $r = 8$, $T = 500$ diffusion steps), varying the number of LoRA locations L across values 24, 48, 72, 96 and measured the total latency.

From these results presented in Table 12, we observe that the **total latency remains low** even as the number of LoRA locations increases. Even for the configuration with **96 locations**, the total latency stays below **0.6 seconds**, which is negligible compared to training a task-specific LoRA from scratch. The **performance remains stable** with negligible fluctuations in accuracy across different L values, further confirming the efficiency of LoRAGen in scaling with the number of locations.

Scaling with LoRA ranks. We evaluate the scaling of LoRAGen with respect to LoRA rank r . We keep the latent dimension fixed at $d_z = 64$ and reuse the same decoder/denoiser widths. The scaling experiment was conducted on the Gemma-2-2B model with $L = 96$ **locations** and $T = 500$ **diffusion steps** for rank $r \in \{8, 16, 32\}$.

1080
 1081 Table 11: End-to-end latency per forward pass with oracle vs. LoRAGen-generated LoRA adapters
 1082 on an NVIDIA H800.

Base model & adapter	Task	Mean latency (ms)
FLAN-T5-Large, $r = 16$, oracle LoRA	AP-Rec	42.10
FLAN-T5-Large, $r = 16$, LoRAGen LoRA	AP-Rec	40.33
Gemma-2-2B, $r = 8$, oracle LoRA	GSM8K	206.55
Gemma-2-2B, $r = 8$, LoRAGen LoRA	GSM8K	130.85

1088
 1089 Table 12: Latency and accuracy scaling with LoRA locations. **Avg. Acc. (%)** reports the average
 1090 accuracy on benchmark tasks as in Table 2.

LoRA locations L	Enc. / Diff. / Dec. latency (ms)	Total gen. latency (ms)	Avg. Acc. (%)
24	0.25 / 563.6 / 5.2	569.0	72.5
48	0.25 / 566.9 / 5.5	572.6	72.4
72	0.40 / 562.7 / 5.6	568.7	72.6
96	0.40 / 569.6 / 5.7	575.7	72.7

1091
 1092 As shown in Table 13, the **total generation latency remains stable** even as the LoRA rank in-
 1093 creases. The **MoE decoder structure** remains unchanged, with only the **final linear projections**
 1094 scaling with r . The average accuracy shows **negligible fluctuations**, indicating that increasing the
 1095 LoRA rank does not introduce significant performance degradation. These results suggest that our
 1096 design is sufficiently flexible to handle larger ranks without requiring a redesign of the architecture.

1097 In summary, the **total generation latency remains low** even as we increase the number of LoRA lo-
 1098 cations, and **LoRA rank scaling** does not necessitate any architectural changes. The small changes
 1099 in latency are well within acceptable limits and do not hinder the performance of the model. These
 1100 findings demonstrate the **scalability** and **efficiency** of LoRAGen in real-world applications.

1108 A.9 TRAINING AND EVALUATION DATASETS FOR ZERO-SHOT GENERATION

1109 As shown in Figure 9 and Figure 10, we conduct zero-shot generation experiment on 136 training
 1110 tasks and 7 evaluating tasks from FLAN subset. These evaluated tasks are separate from the training
 1111 datasets.

1114 A.10 TASK DESCRIPTIONS GENERATED BY A LARGE LANGUAGE MODEL

1115 We automate task description generation for each task by leveraging Deepseek¹. We query its model
 1116 with carefully constructed prompts that incentivize diversity to facilitate downstream generalization
 1117 as shown in Figure 11. In particular, we generate 20 descriptions per task. Figure 12 presents
 1118 representative examples of task descriptions employed in our experiments.

1120 A.11 DETAILS OF EXPERIMENT SETUP

1122 More details about the diffusion architecture, baseline settings hyperparameter settings, training
 1123 details and implementation of weight space loss and module-aware MoE decoder can be found in
 1124 the anonymous repository <https://anonymous.4open.science/r/LoRAGen-02C0>.

1126 A.12 LLM USAGE

1127 We utilized ChatGPT-4o² to refine the content based on our original writing. All revised text
 1128 was subsequently reviewed and verified by us. The natural language task descriptions we used are
 1129 generated by DeepSeek³. All code has undergone comprehensive testing to ensure its reliability.

1131
 1132 ¹<https://www.deepseek.com>

1133 ²<https://chatgpt.com>

³<https://www.deepseek.com>

1134
1135
1136
1137

1138 Table 13: Latency and accuracy scaling with LoRA rank. **Avg. Acc. (%)** reports the average
1139 accuracy on benchmark tasks as in Table 2.

Rank r	Enc. / Diff. / Dec. (ms)	Total gen. latency (ms)	Avg. Acc. (%)
8	0.40 / 573.5 / 5.7	579.6	72.7
16	0.49 / 567.5 / 6.2	574.2	72.8
32	0.64 / 575.6 / 6.5	582.7	72.7

1144
1145
1146
1147
1148
1149
1150

Training Tasks

"lorahub_fian_15.large-adversarial_qa_dbidaf_based_on",
 "lorahub_fian_15.large-adversarial_qa_dbert_question_context_answer",
 "lorahub_fian_15.large-adversarial_qa_droberta_answer_the_following_q",
 "lorahub_fian_15.large-adversarial_qa_droberta_based_on",
 "lorahub_fian_15.large-adversarial_qa_dbidaf.question_context_answer",
 "lorahub_fian_15.large-adversarial_qa_droberta.tell.what_it_is",
 "lorahub_fian_15.large-adversarial_qa_dbert.tell.what_it_is",
 "lorahub_fian_15.large-adversarial_qa_dbidaf.answer.the_following_q",
 "lorahub_fian_15.large-adversarial_qa_dbert_answer.the_following_q",
 "lorahub_fian_15.large-adversarial_qa_dbert_based_on",
 "lorahub_fian_15.large-adversarial_qa_droberta_question_context_answer",
 "lorahub_fian_15.large-amazon_polarity_user_satisfied",
 "lorahub_fian_15.large-app_reviews_categorize_rating_using_review",
 "lorahub_fian_15.large-app_reviews_convert_to_rating",
 "lorahub_fian_15.large-amazon_polarity_negative_positive_tone",
 "lorahub_fian_15.large-amazon_polarity_convey_negative_or_positive_sentiment",
 "lorahub_fian_15.large-amazon_polarity_would_you_buy",
 "lorahub_fian_15.large-amazon_polarity_ls_this_review",
 "lorahub_fian_15.large-yelp_polarity_reviews",
 "lorahub_fian_15.large-glue_wnli",
 "lorahub_fian_15.large-glue_stsb",
 "lorahub_fian_15.large-glue_cola",
 "lorahub_fian_15.large-anli_r1",
 "lorahub_fian_15.large-glue_mrpc",
 "lorahub_fian_15.large-glue_sst2",
 "lorahub_fian_15.large-super_glue_wic",
 "lorahub_fian_15.large-glue_qnli",
 "lorahub_fian_15.large-qasc_qa_with_combined_facts_1",
 "lorahub_fian_15.large-dbededia_14_pick_one_category_for_the_following_text",
 "lorahub_fian_15.large-dbededia_14_given_a_choice_of_categories",
 "lorahub_fian_15.large-dbededia_14_given_list.what_category_does_the_paragraph_belong_to",
 "lorahub_fian_15.large-quartz_bio_who",
 "lorahub_fian_15.large-qasc_qa_with_separated_facts_4",
 "lorahub_fian_15.large-dbededia_14_given_a_list_of_category_what_does_the_title_belong_to",
 "lorahub_fian_15.large-qasc_qa_with_separated_facts_5",
 "lorahub_fian_15.large-qasc_qa_with_separated_facts_1",
 "lorahub_fian_15.large-qasc_qa_with_separated_facts_3",
 "lorahub_fian_15.large-race_high_Select_the_best_answer_no_instructions",
 "lorahub_fian_15.large-quartz_given_the_fact_answer_the_q",
 "lorahub_fian_15.large-web_questions_whats_the_answer",
 "lorahub_fian_15.large-quac",
 "lorahub_fian_15.large-race_middle_Taking_a_test",
 "lorahub_fian_15.large-dream_baseline",
 "lorahub_fian_15.large-sciq_Multiple_Choice_Closed_Book",
 "lorahub_fian_15.large-quoref_Guess_Answer",
 "lorahub_fian_15.large-ropes_plan_no_background",
 "lorahub_fian_15.large-wiki_qa_found_on_google",
 "lorahub_fian_15.large-quall_context_question_description_answer_id",
 "lorahub_fian_15.large-wiki_hop_original_generate_subject_and_object",
 "lorahub_fian_15.large-duorc_ParaphraseRC_answer_question",
 "lorahub_fian_15.large-quall_context_description_question_answer_text",
 "lorahub_fian_15.large-race_middle_Select_the_best_answer_generate_span",
 "lorahub_fian_15.large-social_l_qa_Show_choices_and_generate_answer",
 "lorahub_fian_15.large-quoref_Context_Contains_Answer",
 "lorahub_fian_15.large-quartz_read_passage_below_choose",
 "lorahub_fian_15.large-ropes_background_situation_middle",
 "lorahub_fian_15.large-wiqqa_effect_with_string_answer",
 "lorahub_fian_15.large-duorc_ParaphraseRC_answer_question",
 "lorahub_fian_15.large-wiki_qa_Topic_Prediction_Question_and_Answer_Pair",
 "lorahub_fian_15.large-duorc_SelfRC_title_generation",
 "lorahub_fian_15.large-quall_context_question_description_answer_text",
 "lorahub_fian_15.large-race_high_Write_a_multi_choice_question_options_given",
 "lorahub_fian_15.large-quoref_Given_Context_Answer_Question",
 "lorahub_fian_15.large-quall_no_prompt_text",
 "lorahub_fian_15.large-wiqqa_Jeopardy_style",
 "lorahub_fian_15.large-wiqqa_does_the_supposed_perturbation_have_an_effect",
 "lorahub_fian_15.large-ropes_plan_background_situation",
 "lorahub_fian_15.large-wiki_hop_original_generate_object",
 "lorahub_fian_15.large-kill_tasks_hopqa_complex_question",
 "lorahub_fian_15.large-kill_tasks_hopqa_final_exam",
 "lorahub_fian_15.large-wiqqa_what_is_the_final_step_of_the_following_process",
 "lorahub_fian_15.large-ropes_prompt_bottom_hint_beginning",
 "lorahub_fian_15.large-race_middle_Select_the_best_answer",
 "lorahub_fian_15.large-quall_description_context_question_answer_text",
 "lorahub_fian_15.large-social_l_qa_Check_if_a_random_answer_is_valid_or_not",
 "lorahub_fian_15.large-ropes_prompt_mix",
 "lorahub_fian_15.large-ropes_given_background_situation",
 "lorahub_fian_15.large-sciq_Multiple_Choice_Question_First",
 "lorahub_fian_15.large-wiqqa_hop_original_choose_best_object_affirmative_1",
 "lorahub_fian_15.large-race_middle_Select_the_best_answer_no_instructions",
 "lorahub_fian_15.large-quoref_Answer_Friend_Question",
 "lorahub_fian_15.large-wiki_hop_original_explain_relation",
 "lorahub_fian_15.large-duorc_SelfRC_question_answering",
 "lorahub_fian_15.large-ropes_prompt_beginning",
 "lorahub_fian_15.large-sciq_Direct_Question_Closed_Book",
 "lorahub_fian_15.large-race_high_Taking_a_test",
 "lorahub_fian_15.large-quoref_Find_Answer",
 "lorahub_fian_15.large-duorc_ParaphraseRC_extract_answer",
 "lorahub_fian_15.large-wiki_qa_Decide_good_answer",
 "lorahub_fian_15.large-duorc_ParaphraseRC_title_generation",
 "lorahub_fian_15.large-quoref_Found_Context_Online",
 "lorahub_fian_15.large-sciq_Direct_Question",
 "lorahub_fian_15.large-wiki_hop_original_choose_best_object_interrogative_2",
 "lorahub_fian_15.large-wiki_qa_exercise",
 "lorahub_fian_15.large-ropes_background_new_situation_answer",
 "lorahub_fian_15.large-wiki_hop_original_choose_best_object_affirmative_3",
 "lorahub_fian_15.large-quartz_having.read.above.passage",
 "lorahub_fian_15.large-ropes_prompt_bottom_no_hint",
 "lorahub_fian_15.large-quall_context_question_answer_description_text",
 "lorahub_fian_15.large-social_l_qa_Generate_the_question_from_the_answer",
 "lorahub_fian_15.large-duorc_SelfRC_decide_worth_it",
 "lorahub_fian_15.large-quartz_use_info_from_paragraph_question",
 "lorahub_fian_15.large-quall_no_prompt_id",
 "lorahub_fian_15.large-quoref_Read_And_Extract",
 "lorahub_fian_15.large-race_high_Is_this_the_right_answer",
 "lorahub_fian_15.large-quall_context_question_answer_description_id",
 "lorahub_fian_15.large-wiqqa_effect_with_label_answer",
 "lorahub_fian_15.large-web_questions_potential_correct_answer",
 "lorahub_fian_15.large-race_middle_Write_a_multi_choice_question_options_given",
 "lorahub_fian_15.large-wiqqa_which_of_the_following_is_the_supposed_perturbation",
 "lorahub_fian_15.large-duorc_SelfRC_movie_director",
 "lorahub_fian_15.large-wiki_qa_Topic_Prediction_Answer_Only",
 "lorahub_fian_15.large-wiqqa_what_might_be_the_last_step_of_the_process",
 "lorahub_fian_15.large-wiqqa_what_is_the_missing_first_step",
 "lorahub_fian_15.large-wiki_hop_original_choose_best_object_affirmative_2",
 "lorahub_fian_15.large-duorc_ParaphraseRC_movie_director",
 "lorahub_fian_15.large-duorc_SelfRC_answer_question",
 "lorahub_fian_15.large-wiki_hop_original_generate_subject",
 "lorahub_fian_15.large-quall_description_context_question_answer_id",
 "lorahub_fian_15.large-sciq_Multiple_Choice",
 "lorahub_fian_15.large-quoref_What_Is_The_Answer",
 "lorahub_fian_15.large-social_l_qa_Show_choices_and_generate_index",
 "lorahub_fian_15.large-quoref_Answer_Question_Given_Context",
 "lorahub_fian_15.large-ropes_read_background_situation",
 "lorahub_fian_15.large-duorc_SelfRC_extract_answer",
 "lorahub_fian_15.large-race_high_Select_the_best_answer",
 "lorahub_fian_15.large-para_crawl_enes",
 "lorahub_fian_15.large-newroom",
 "lorahub_fian_15.large-geonew",
 "lorahub_fian_15.large-paws_wiki",
 "lorahub_fian_15.large-quoref_testing_students",
 "lorahub_fian_15.large-quoref_logic_test",
 "lorahub_fian_15.large-quoref_choose_between",
 "lorahub_fian_15.large-quoref_do_not_use",
 "lorahub_fian_15.large-quoref_heres_a_story"

1184 Figure 9: Training tasks from FLAN dataset used for training the LoRAGen model
1185
1186
1187

```
1188
1189
1190
1191
1192
1193
1194
1195
1196
1197
1198
1199
1200
1201
1202
1203
1204
1205
1206
1207
1208
1209
1210
1211
1212
1213 "lorahub_flan_t5_large-amazon_polarity_User_recommend_this_product",
1214 "lorahub_flan_t5_large-amazon_polarity_Is_this_product_review_positive",
1215 "lorahub_flan_t5_large-qasc_is_correct_1",
1216 "lorahub_flan_t5_large-qasc_is_correct_2",
1217 "lorahub_flan_t5_large-wiki_qa_Is_This_True_",
1218 "lorahub_flan_t5_large-wiki_qa_automatic_system"
1219
1220
1221
1222
1223
1224
1225
1226
1227
1228
1229
1230
1231
1232
1233
1234
1235
1236
1237
1238
1239
1240
1241
```

Validation Tasks

```
"lorahub_flan_t5_large-amazon_polarity_User_recommend_this_product",
"lorahub_flan_t5_large-amazon_polarity_Is_this_product_review_positive",
"lorahub_flan_t5_large-qasc_is_correct_1",
"lorahub_flan_t5_large-qasc_is_correct_2",
"lorahub_flan_t5_large-wiki_qa_Is_This_True_",
"lorahub_flan_t5_large-wiki_qa_automatic_system"
```

Figure 10: Validation tasks used during the training the LoRAGen model

```

1242
1243
1244
1245
1246
1247
1248
1249
1250 Prompt
1251
1252 ## Objective
1253 For every LoRA adapter directory, construct a clean list of task descriptions and turn them into a single sentence-level
1254 embedding.
1255 The generator reads curated descriptions from YAML, applies light normalization, matches them by name, and averages T5
1256 sentence embeddings.
1257
1258 ## Inputs
1259 - --yaml_root: directory with one subfolder per task; each contains `metadata.yaml` with a descriptions list.
1260 - --logs_root: Stage-1 checkpoint tree (used to load the VAE encoder for latents; text generation is file-driven).
1261 ## Options
1262 - --strip_generic: remove leading boilerplate like “The task is / involves / requires ...”.
1263 - --text_pooling {mean, first}: sentence pooling for T5 (mean is mask-aware and default).
1264
1265 ## Procedure
1266 1. Scan YAML repository
1267     For each subfolder under `yaml_root`, load `metadata.yaml` and read descriptions. Skip if missing.
1268 2. Optional light normalization
1269     If `--strip_generic`, strip only the leading boilerplate using a regex and keep the substantive remainder.
1270 3. Build multi-key alias map
1271     Register the cleaned descriptions under:
1272     - <entry>
1273     - `lorahub_flan_t5_large-<entry>`
1274     - `lorahub_flan_t5_large_<entry>`
1275 4. Iterate LoRA adapters
1276     For each `./flan_t5_large_lora/<task_key>/adapter_model.bin`, resolve descriptions by probing name variants (prefix
1277 removal and `-/` swaps).
1278     - Fallback: if no hit, use `[task_key]` as the only description and log the miss.
1279 5. Encode with T5
1280     - Tokenize each description, run `T5EncoderModel`.
1281     - Pool to a sentence vector via:
1282         - mean pooling (default), or
1283         - first token pooling.
1284     - Average across all descriptions → one 1024-d text embedding per task.
1285 6. Normalize save key
1286     Drop `lorahub_flan_t5_large-/` prefix to form the canonical task name and save:
1287     ```json
1288     {
1289         "task_name": "<canonical_name>",
1290         "text_embedding": "<1024-d tensor>",
1291         "latent": "<288 x latent_dim tensor>"
1292     }
1293 ## Output
1294 One PyTorch file per experiment, e.g.:
1295     e_000996_with_task_name_vae_task_172_latent_288_64_embed_1024.pt
1296 Mapping canonical task names to averaged text embeddings (plus latents produced in the same pass).
1297
1298
1299
1300
1301
1302
1303
1304
1305
1306
1307
1308
1309
1310
1311
1312
1313
1314
1315
1316
1317
1318
1319
1320
1321
1322
1323
1324
1325
1326
1327
1328
1329
1330
1331
1332
1333
1334
1335
1336
1337
1338
1339
1340
1341
1342
1343
1344
1345
1346
1347
1348
1349
1350
1351
1352
1353
1354
1355
1356
1357
1358
1359
1360
1361
1362
1363
1364
1365
1366
1367
1368
1369
1370
1371
1372
1373
1374
1375
1376
1377
1378
1379
1380
1381
1382
1383
1384
1385
1386
1387
1388
1389
1390
1391
1392
1393
1394
1395
1396
1397
1398
1399
1400
1401
1402
1403
1404
1405
1406
1407
1408
1409
1410
1411
1412
1413
1414
1415
1416
1417
1418
1419
1420
1421
1422
1423
1424
1425
1426
1427
1428
1429
1430
1431
1432
1433
1434
1435
1436
1437
1438
1439
1440
1441
1442
1443
1444
1445
1446
1447
1448
1449
1450
1451
1452
1453
1454
1455
1456
1457
1458
1459
1460
1461
1462
1463
1464
1465
1466
1467
1468
1469
1470
1471
1472
1473
1474
1475
1476
1477
1478
1479
1480
1481
1482
1483
1484
1485
1486
1487
1488
1489
1490
1491
1492
1493
1494
1495
1496
1497
1498
1499
1500
1501
1502
1503
1504
1505
1506
1507
1508
1509
1510
1511
1512
1513
1514
1515
1516
1517
1518
1519
1520
1521
1522
1523
1524
1525
1526
1527
1528
1529
1530
1531
1532
1533
1534
1535
1536
1537
1538
1539
1540
1541
1542
1543
1544
1545
1546
1547
1548
1549
1550
1551
1552
1553
1554
1555
1556
1557
1558
1559
1560
1561
1562
1563
1564
1565
1566
1567
1568
1569
1570
1571
1572
1573
1574
1575
1576
1577
1578
1579
1580
1581
1582
1583
1584
1585
1586
1587
1588
1589
1590
1591
1592
1593
1594
1595
1596
1597
1598
1599
1600
1601
1602
1603
1604
1605
1606
1607
1608
1609
1610
1611
1612
1613
1614
1615
1616
1617
1618
1619
1620
1621
1622
1623
1624
1625
1626
1627
1628
1629
1630
1631
1632
1633
1634
1635
1636
1637
1638
1639
1640
1641
1642
1643
1644
1645
1646
1647
1648
1649
1650
1651
1652
1653
1654
1655
1656
1657
1658
1659
1660
1661
1662
1663
1664
1665
1666
1667
1668
1669
1670
1671
1672
1673
1674
1675
1676
1677
1678
1679
1680
1681
1682
1683
1684
1685
1686
1687
1688
1689
1690
1691
1692
1693
1694
1695
1696
1697
1698
1699
1700
1701
1702
1703
1704
1705
1706
1707
1708
1709
1710
1711
1712
1713
1714
1715
1716
1717
1718
1719
1720
1721
1722
1723
1724
1725
1726
1727
1728
1729
1730
1731
1732
1733
1734
1735
1736
1737
1738
1739
1740
1741
1742
1743
1744
1745
1746
1747
1748
1749
1750
1751
1752
1753
1754
1755
1756
1757
1758
1759
1760
1761
1762
1763
1764
1765
1766
1767
1768
1769
1770
1771
1772
1773
1774
1775
1776
1777
1778
1779
1780
1781
1782
1783
1784
1785
1786
1787
1788
1789
1790
1791
1792
1793
1794
1795
1796
1797
1798
1799
1800
1801
1802
1803
1804
1805
1806
1807
1808
1809
1810
1811
1812
1813
1814
1815
1816
1817
1818
1819
1820
1821
1822
1823
1824
1825
1826
1827
1828
1829
1830
1831
1832
1833
1834
1835
1836
1837
1838
1839
1840
1841
1842
1843
1844
1845
1846
1847
1848
1849
1850
1851
1852
1853
1854
1855
1856
1857
1858
1859
1860
1861
1862
1863
1864
1865
1866
1867
1868
1869
1870
1871
1872
1873
1874
1875
1876
1877
1878
1879
1880
1881
1882
1883
1884
1885
1886
1887
1888
1889
1890
1891
1892
1893
1894
1895
1896
1897
1898
1899
1900
1901
1902
1903
1904
1905
1906
1907
1908
1909
1910
1911
1912
1913
1914
1915
1916
1917
1918
1919
1920
1921
1922
1923
1924
1925
1926
1927
1928
1929
1930
1931
1932
1933
1934
1935
1936
1937
1938
1939
1940
1941
1942
1943
1944
1945
1946
1947
1948
1949
1950
1951
1952
1953
1954
1955
1956
1957
1958
1959
1960
1961
1962
1963
1964
1965
1966
1967
1968
1969
1970
1971
1972
1973
1974
1975
1976
1977
1978
1979
1980
1981
1982
1983
1984
1985
1986
1987
1988
1989
1990
1991
1992
1993
1994
1995
1996
1997
1998
1999
1999
2000
2001
2002
2003
2004
2005
2006
2007
2008
2009
2010
2011
2012
2013
2014
2015
2016
2017
2018
2019
2020
2021
2022
2023
2024
2025
2026
2027
2028
2029
2030
2031
2032
2033
2034
2035
2036
2037
2038
2039
2040
2041
2042
2043
2044
2045
2046
2047
2048
2049
2050
2051
2052
2053
2054
2055
2056
2057
2058
2059
2060
2061
2062
2063
2064
2065
2066
2067
2068
2069
2070
2071
2072
2073
2074
2075
2076
2077
2078
2079
2080
2081
2082
2083
2084
2085
2086
2087
2088
2089
2090
2091
2092
2093
2094
2095
2096
2097
2098
2099
2099
2100
2101
2102
2103
2104
2105
2106
2107
2108
2109
2109
2110
2111
2112
2113
2114
2115
2116
2117
2118
2119
2119
2120
2121
2122
2123
2124
2125
2126
2127
2128
2129
2129
2130
2131
2132
2133
2134
2135
2136
2137
2138
2139
2139
2140
2141
2142
2143
2144
2145
2146
2147
2147
2148
2149
2149
2150
2151
2152
2153
2154
2155
2156
2157
2158
2159
2159
2160
2161
2162
2163
2164
2165
2166
2167
2167
2168
2169
2169
2170
2171
2172
2173
2174
2175
2175
2176
2177
2177
2178
2179
2179
2180
2180
2181
2181
2182
2182
2183
2183
2184
2184
2185
2185
2186
2186
2187
2187
2188
2188
2189
2189
2190
2190
2191
2191
2192
2192
2193
2193
2194
2194
2195
2195
2196
2196
2197
2197
2198
2198
2199
2199
2200
2200
2201
2201
2202
2202
2203
2203
2204
2204
2205
2205
2206
2206
2207
2207
2208
2208
2209
2209
2210
2210
2211
2211
2212
2212
2213
2213
2214
2214
2215
2215
2216
2216
2217
2217
2218
2218
2219
2219
2220
2220
2221
2221
2222
2222
2223
2223
2224
2224
2225
2225
2226
2226
2227
2227
2228
2228
2229
2229
2230
2230
2231
2231
2232
2232
2233
2233
2234
2234
2235
2235
2236
2236
2237
2237
2238
2238
2239
2239
2240
2240
2241
2241
2242
2242
2243
2243
2244
2244
2245
2245
2246
2246
2247
2247
2248
2248
2249
2249
2250
2250
2251
2251
2252
2252
2253
2253
2254
2254
2255
2255
2256
2256
2257
2257
2258
2258
2259
2259
2260
2260
2261
2261
2262
2262
2263
2263
2264
2264
2265
2265
2266
2266
2267
2267
2268
2268
2269
2269
2270
2270
2271
2271
2272
2272
2273
2273
2274
2274
2275
2275
2276
2276
2277
2277
2278
2278
2279
2279
2280
2280
2281
2281
2282
2282
2283
2283
2284
2284
2285
2285
2286
2286
2287
2287
2288
2288
2289
2289
2290
2290
2291
2291
2292
2292
2293
2293
2294
2294
2295
2295
2296
2296
2297
2297
2298
2298
2299
2299
2300
2300
2301
2301
2302
2302
2303
2303
2304
2304
2305
2305
2306
2306
2307
2307
2308
2308
2309
2309
2310
2310
2311
2311
2312
2312
2313
2313
2314
2314
2315
2315
2316
2316
2317
2317
2318
2318
2319
2319
2320
2320
2321
2321
2322
2322
2323
2323
2324
2324
2325
2325
2326
2326
2327
2327
2328
2328
2329
2329
2330
2330
2331
2331
2332
2332
2333
2333
2334
2334
2335
2335
2336
2336
2337
2337
2338
2338
2339
2339
2340
2340
2341
2341
2342
2342
2343
2343
2344
2344
2345
2345
2346
2346
2347
2347
2348
2348
2349
2349
2350
2350
2351
2351
2352
2352
2353
2353
2354
2354
2355
2355
2356
2356
2357
2357
2358
2358
2359
2359
2360
2360
2361
2361
2362
2362
2363
2363
2364
2364
2365
2365
2366
2366
2367
2367
2368
2368
2369
2369
2370
2370
2371
2371
2372
2372
2373
2373
2374
2374
2375
2375
2376
2376
2377
2377
2378
2378
2379
2379
2380
2380
2381
2381
2382
2382
2383
2383
2384
2384
2385
2385
2386
2386
2387
2387
2388
2388
2389
2389
2390
2390
2391
2391
2392
2392
2393
2393
2394
2394
2395
2395
2396
2396
2397
2397
2398
2398
2399
2399
2400
2400
2401
2401
2402
2402
2403
2403
2404
2404
2405
2405
2406
2406
2407
2407
2408
2408
2409
2409
2410
2410
2411
2411
2412
2412
2413
2413
2414
2414
2415
2415
2416
2416
2417
2417
2418
2418
2419
2419
2420
2420
2421
2421
2422
2422
2423
2423
2424
2424
2425
2425
2426
2426
2427
2427
2428
2428
2429
2429
2430
2430
2431
2431
2432
2432
2433
2433
2434
2434
2435
2435
2436
2436
2437
2437
2438
2438
2439
2439
2440
2440
2441
2441
2442
2442
2443
2443
2444
2444
2445
2445
2446
2446
2447
2447
2448
2448
2449
2449
2450
2450
2451
2451
2452
2452
2453
2453
2454
2454
2455
2455
2456
2456
2457
2457
2458
2458
2459
2459
2460
2460
2461
2461
2462
2462
2463
2463
2464
2464
2465
2465
2466
2466
2467
2467
2468
2468
2469
2469
2470
2470
2471
2471
2472
2472
2473
2473
2474
2474
2475
2475
2476
2476
2477
2477
2478
2478
2479
2479
2480
2480
2481
2481
2482
2482
2483
2483
2484
2484
2485
2485
2486
2486
2487
2487
2488
2488
2489
2489
2490
2490
2491
2491
2492
2492
2493
2493
2494
2494
2495
2495
2496
2496
2497
2497
2498
2498
2499
2499
2500
2500
2501
2501
2502
2502
2503
2503
2504
2504
2505
2505
2506
2506
2507
2507
2508
2508
2509
2509
2510
2510
2511
2511
2512
2512
2513
2513
2514
2514
2515
2515
2516
2516
2517
2517
2518
2518
2519
2519
2520
2520
2521
2521
2522
2522
2523
2523
2524
2524
2525
2525
2526
2526
2527
2527
2528
2528
2529
2529
2530
2530
2531
2531
2532
2532
2533
2533
2534
2534
2535
2535
2536
2536
2537
2537
2538
2538
2539
2539
2540
2540
2541
2541
2542
2542
2543
2543
2544
2544
2545
2545
2546
2546
2547
2547
2548
2548
2549
2549
2550
2550
2551
2551
2552
2552
2553
2553
2554
2554
2555
2555
2556
2556
2557
2557
2558
2558
2559
2559
2560
2560
2561
2561
2562
2562
2563
2563
2564
2564
2565
2565
2566
2566
2567
2567
2568
2568
2569
2569
2570
2570
2571
2571
2572
2572
2573
2573
2574
2574
2575
2575
2576
2576
2577
2577
2578
2578
2579
2579
2580
2580
2581
2581
2582
2582
2583
2583
2584
2584
2585
2585
2586
2586
2587
2587
2588
2588
2589
2589
2590
2590
2591
2591
2592
2592
2593
2593
2594
2594
2595
2595
2596
2596
2597
2597
2598
2598
2599
2599
2600
2600
2601
2601
2602
2602
2603
2603
2604
2604
2605
2605
2606
2606
2607
2607
2608
2608
2609
2609
2610
2610
2611
2611
2612
2612
2613
2613
2614
2614
2615
2615
2616
2616
2617
2617
2618
2618
2619
2619
2620
2620
2621
2621
2622
2622
2623
2623
2624
2624
2625
2625
2626
2626
2627
2627
2628
2628
2629
2629
2630
2630
2631
2631
2632
2632
2633
2633
2634
2634
2635
2635
2636
2636
2637
2637
2638
2638
2639
2639
2640
2640
2641
2641
2642
2642
2643
2643
2644
2644
2645
2645
2646
2646
2647
2647
2648
2648
2649
2649
2650
2650
2651
2651
2652
2652
2653
2653
2654
2654
2655
2655
2656
2656
2657
2657
2658
2658
2659
2659
2660
2660
2661
2661
2662
2662
2663
2663
2664
2664
2665
2665
2666
2666
2667
2667
2668
2668
2669
2669
2670
2670
2671
2671
2672
2672
2673
2673
2674
2674
2675
2675
2676
2676
2677
2677
2678
2678
2679
2679
2680
2680
2681
2681
2682
2682
2683
2683
2684
2684
2685
2685
2686
2686
2687
2687
2688
2688
2689
2689
2690
2690
2691
2691
2692
2692
2693
2693
2694
2694
2695
2695
2696
2696
2697
2697
2698
2698
2699
2699
2700
2700
2701
2701
2702
2702
2703
2703
2704
2704
2705
2705
2706
2706
2707
2707
2708
2708
2709
2709
2710
2710
2711
2711
2712
2712
2713
2713
2714
2714
2715
2715
2716
2716
2717
2717
2718
2718
2719
2719
2720
2720
2721
2721
2722
2722
2723
2723
2724
2724
2725
2725
2726
2726
2727
2727
2728
2728
2729
2729
2730
2730
2731
2731
2732
2732
2733
2733
2734
2734
2735
2735
2736
2736
2737
2737
2738
2738
2739
2739
2740
2740
2741
2741
2742
2742
2743
2743
2744
2744
2745
2745
2746
2746
2747
2747
2748
2748
2749
2749
2750
2750
2751
2751
2752
2752
2753
2753
2754
2754
2755
2755
2756
2756
2757
2757
2758
2758
2759
2759
2760
2760
2761
2761
2762
2762
2763
2763
2764
2764
2765
2765
2766
2766
2767
2767
2768
2768
2769
2769
2770
2770
2771
2771
2772
2772
2773
2773
2774
2774
2775
2
```

1296
 1297
 1298
 1299
 1300
 1301
 1302
 1303
 1304
 1305
 1306
 1307
 1308
 1309
 1310
 1311
 1312
 1313

1314
 1315

1316
 1317
 1318
 1319

1320
 1321
 1322

1323
 1324
 1325
 1326

1327
 1328
 1329
 1330

1331
 1332

Task descriptions

adversarial_qa_dbert_answer_the_followi

- The task involves reading a passage and answering a question based on the information provided in the text.
- The task requires identifying specific details from a given passage to answer a question.
- The task is about locating exact information within a text to respond to a direct question.

adversarial_qa_dbidaf_answer_the_following_q

- The task is to locate the answer to a question within a provided passage.
- The task requires finding the exact words or phrases in a passage that answer a question.
- The task is about answering questions by referring to specific details in a text.

adversarial_qa_droberta_question_context_answer

- The task requires matching questions with relevant facts from the given context.
- The task is about locating key details in a text to answer a direct question.
- The task involves reasoning about a passage to derive the correct answer.

adversarial_qa_dbert_based_on

- The task involves finding specific information in a given text to answer a direct question.
- The task requires identifying key details from a passage that directly respond to a question.
- The task is about locating exact answers within a provided context based on a question.

Figure 12: Examples of task descriptions generated by our pipeline.

1333
 1334
 1335
 1336
 1337
 1338
 1339
 1340
 1341
 1342
 1343
 1344
 1345
 1346
 1347
 1348
 1349