

LLM Braces: Straightening Out LLM Predictions with Relevant Sub-Updates

Ying Shen

Computer Science Department
University of Illinois Urbana-Champaign
ying22@illinois.edu

Lifu Huang

Computer Science Department
UC Davis
lfuhuang@ucdavis.edu

Abstract

Recent findings reveal that much of the knowledge in a Transformer-based Large Language Model (LLM) is encoded in its feed-forward (FFN) layers, where each FFN layer can be interpreted as the summation of sub-updates, each corresponding to a weighted column vector from the FFN’s value parameter matrix that often encodes human-interpretable concepts. In light of this, we hypothesize that model performance and behaviors can be further enhanced and controlled by modulating the contributions of these sub-updates based on their relevance to the input or target output style, and propose LLMBRACES, a novel and efficient method that computes relevance scores associated with value vectors in FFN layers and leverages these scores to dynamically adjust the contribution of sub-updates. By optimizing sub-update contributions, LLMBRACES refines the prediction process, leading to more accurate and reliable outputs, much like a ‘brace’ providing support and stability. Moreover, LLMBRACES can be extended to support conditional control over generation characteristics, such as sentiment, thereby offering fine-grained steering of LLM outputs. Extensive experiments on various LLMs—including Qwen2.5-1.5B, Llama2-7B, and Llama3-8B—demonstrate that LLMBRACES outperforms baseline approaches in both fine-tuning and zero-shot settings while requiring significantly fewer tunable parameters, up to 75% fewer compared to LoRA. Furthermore, LLMBRACES excels in sentiment-controlled generation and toxicity reduction, highlighting its potential for flexible, controlled text generation across applications.

1 Introduction

Transformer-based Large Language Models (LLMs) (Devlin et al., 2019; Radford et al., 2019; Brown et al., 2020; Black et al., 2022; Touvron et al., 2023; Jiang et al., 2023; Anthropic, 2023)




Figure 1: An illustration of feed-forward network (FFN) in a Transformer and its augmentation by LLMBRACES. Each FFN’s output can be viewed as a sum of **sub-updates**, each corresponding to a weighted FFN value parameter vector: $\text{FFN}(h_i) = W_V \cdot \sigma(W_K \cdot h_i) = \sum_{j=1}^{d_m} \sigma(h_i \cdot k_j) \cdot v_j$. LLMBRACES adjusts the weight of each value parameter vector based on its relevance to the input or target output style, using a relevance function $R(\cdot)$, to modulate model predictions and behaviors.

have demonstrated impressive capabilities across a range of natural language processing tasks, yet their internal workings—specifically how they generate predictions—remain somewhat elusive. Recent efforts have focused on demystifying the LLMs through interpreting their neurons (Dai et al., 2022; Voita et al., 2023), attention mechanisms (Geva et al., 2023a,b; Ferrando and Voita, 2024), as well as internal parameters (Geva et al., 2022; Dar et al., 2023).

One of the main findings in previous studies is that knowledge is primarily stored in the feed-forward (FFN) layers of Transformer-based lan-

guage models (Dai et al., 2022; Geva et al., 2021, 2022; Yao et al., 2024). Geva et al. (2021, 2022) identify each FFN layer as a key-value memory—each “key” is associated with specific textual patterns from the training data, while each “value” represents a distribution over the output vocabulary. Consequently, an FFN update can be decomposed into a collection of *sub-updates*, where each sub-update is represented by a weighted “value” parameter vector within the FFN, as shown in Fig. 1. These value vectors often encode human-interpretable concepts, with the associated weights determining the contribution of each sub-update. From the perspective of the residual stream (Elhage et al., 2021), these sub-updates can be understood as making incremental additive contributions to the model’s final output distribution.

Building on these insights, we hypothesize that model accuracy can be improved by dynamically modulating sub-updates based on the relevance of value vectors to the input—suppressing less relevant ones to reduce noise while amplifying highly aligned ones to enhance predictions. Based on this motivation, we introduce LLMBRACES, a novel approach that employs a learnable relevance module to project both hidden states and value vectors into a common low-dimensional subspace using a low-rank projection with orthonormal rows, which not only ensures computational efficiency but also faithfully preserves the geometric structure of the original high-dimensional representations (Halko et al., 2011). In this transformed space, we compute relevance scores that quantify the alignment between each value vector and the input context, and combine them with the original contribution weights through a learnable gating mechanism, which dynamically adjusts the influence of each sub-update, enhancing the model prediction process. Moreover, LLMBRACES can be easily extended to steer model outputs toward specific attributes, e.g., a desired sentiment or style, by adapting the relevance module to measure the alignment between value vectors and attribute-specific token representations. This conditional relevance score is then integrated into the overall relevance score, with an adjustable scalar parameter that allows users to control the influence of this attribute-specific signal during inference.

We validate the effectiveness of LLMBRACES through extensive experiments on a variety of LLMs, including Qwen2.5-1.5B (Yang et al., 2024), Llama2-7B (Touvron et al., 2023), and Llama3-

8B (Dubey et al., 2024). Our evaluations encompass both fine-tuning performance on common-sense reasoning tasks and zero-shot performance across a range of tasks designed to assess the models’ capabilities in factual knowledge, hallucination reduction, and overall performance on aggregated natural language benchmarks. Notably, LLMBRACES improves the average performance by 13.9% on Qwen2.5-1.5B, 19.6% on Llama2-7B, and 29.7% on Llama3-8B compared to LoRA in the zero-shot setting. Our results indicate that LLMBRACES improves the trustworthiness of LLM and demonstrate its effectiveness in retaining knowledge and facilitating reasoning in the zero-shot setting. Furthermore, we assess LLMBRACES in the context of conditional text generation. In particular, our method proves effective in sentiment steering and toxic language suppression tasks, achieving strong performance in both fine-tuning and zero-shot settings.

In summary, our contributions are threefold: 1) We propose LLMBRACES, a novel and efficient method that computes and leverages relevance scores associated with value vectors to dynamically adjust sub-updates in FFN layers. 2) We extend LLMBRACES to support task-specific controlled generation, offering a flexible framework for steering LLM outputs in practical applications. 3) We perform extensive experiments demonstrating that our approach improves model performance—enhancing factual accuracy, improve trustworthiness, and enabling controlled text generation across multiple LLMs and tasks.

2 Related Work

Understanding LLMs’ Internal Working As large language models (LLMs) grow in sophistication, research has sought to uncover their internal mechanisms to better understand their generation processes. Recent studies have explored projecting internal representations onto model’s vocabulary (nostalgebraist, 2020; Geva et al., 2022; Katz et al., 2024; Belrose et al., 2023), applying targeted interventions in transformer computation (Finlayson et al., 2021; Stolfo et al., 2023; Ghandeharioun et al., 2024; Geiger et al., 2024), and investigating the computation graph of language models (Wang et al.; Merullo et al.; Yao et al., 2024). These studies offer valuable insights into how LLMs encode and retrieve knowledge. Building on this, our work revisit the role of feed-

forward networks in representing knowledge (Geva et al., 2021, 2022) and leverages this insight to develop a simple, parameter-efficient approach for more robust and controllable LLM generation.

Model Editing Model editing (Sinitsin et al., 2019) enables fast, targeted updates to a pre-trained model’s behavior without full retraining, allowing it to adapt to new information or correcting specific errors efficiently. Recent work on LLMs has focused on modifying stored knowledge to update or correct factual outputs (Meng et al., 2022, 2023; Huang et al., 2023). Methods include integrating auxiliary networks (Huang et al., 2023; Hartvigsen et al., 2024) or directly modifying model parameters responsible for specific outputs (Dai et al., 2022; Meng et al., 2022, 2023). Unlike these studies which focused on changing particular outputs for certain model inputs, our work seeks to enable more flexible and generic updates to model behavior, rather than focusing on particular aspects.

Controllable Text Generation Controllable text generation (Zhang et al., 2023) aims to guide Existing methods include fine-tuning models for specific constraints (Keskar et al., 2019; He, 2021), prompt engineering (Reif et al., 2022; Zhou et al., 2023), and latent space manipulation (Zou et al., 2023; Liu et al.; Turner et al., 2023; Han et al., 2024). Unlike these approaches that focus solely on conditional generation, our method dynamically modulates sub-update contributions based on their relevance to the input or target output style. This not only enables more effective conditional generation but also enhances overall model performance on general tasks by refining the prediction process.

3 Preliminaries

In this work, we mainly focus on Transformer-based large language models which have shown exceptional performance across various tasks. Given a sequence of input tokens $x = [x_0, \dots, x_N]$, a transformer-based LLM first encodes the tokens into vectors representations $\mathbf{h}^0 = [\mathbf{h}_0^0, \dots, \mathbf{h}_N^0] \in \mathbb{R}^{N \times d}$ ¹, which are then iteratively updated through a sequence of L Transformer layers, each consisting of a multi-head self-attention (MHSA) sublayer followed by a feed-forward (FFN) sublayer, interconnected by residual connections.

Feed-Forward Layer (FFN) As shown in Fig. 1, each FFN layer l takes the hidden state output \mathbf{h}_i^ℓ

at the i -th token position from its preceding l -th MHSA layer as input, then applies a transformation, resulting in an updated representation $\tilde{\mathbf{h}}_i^\ell$ via residual connection:

$$\tilde{\mathbf{h}}_i^\ell = \text{FFN}^\ell(\mathbf{h}_i^\ell) + \mathbf{h}_i^\ell = \mathbf{W}_V^\ell \cdot \sigma(\mathbf{W}_K^\ell \cdot \mathbf{h}_i^\ell) + \mathbf{h}_i^\ell, \quad (1)$$

where $\mathbf{W}_V^\ell \in \mathbb{R}^{d \times d_m}$ and $\mathbf{W}_K^\ell \in \mathbb{R}^{d_m \times d}$ can be conceptualized as the key and value parameter matrices with inner-dimension d_m at layer ℓ , and $\sigma(\cdot)$ denotes a nonlinear activation function.

Based on Geva et al. (2021, 2022), the ℓ -th FFN layer can be further dissected into:

$$\text{FFN}^\ell(\mathbf{h}_i^\ell) = \sum_{j=1}^{d_m} w_{i,j}^\ell \cdot \mathbf{v}_j^\ell, \quad (2)$$

where each term $(w_{i,j}^\ell \cdot \mathbf{v}_j^\ell)$ acts as a sub-update within the layer and $w_{i,j}^\ell = \sigma(\mathbf{h}_i^\ell \cdot \mathbf{k}_j^\ell)$ is the coefficient that assigns weights to the corresponding value vector \mathbf{v}_j^ℓ . Here, \mathbf{k}_j^ℓ and \mathbf{v}_j^ℓ denote the j -th row of \mathbf{W}_K^ℓ and the j -th column of \mathbf{W}_V^ℓ , respectively. This decomposition reveals that each FFN update can be viewed as an aggregation of these sub-updates, where each sub-update is effectively a weighted value vector. These sub-updates incrementally contribute to the model’s final output distribution.

4 LLMBRACES

Building on insights into FFNs, we introduce LLMBRACES to dynamically modulate the contributions of sub-updates in each FFN layer, aiming to efficiently refine the model prediction process and make it more accurate and controllable. Notably, LLMBRACES introduces only a minimal number of additional parameters and can be seamlessly integrated into any existing transformer-based LLMs. Fig. 2 shows an overview of LLMBRACES.

4.1 Relevance of Sub-updates

As shown in Eq. (2), each sub-update $(w_{i,j}^\ell \cdot \mathbf{v}_j^\ell)$ in FFN represents a weighted *value* parameter vector, which usually corresponds to a set of human interpretable concepts. Given this, we postulate that the accuracy of the model’s predictions can be improved by increasing the contribution of value vectors that are highly relevant to the given input while minimizing the contribution of those that are not. To achieve this, we propose to explicitly measure the relevance of each value vector to the current input, i.e., the hidden states of input tokens,

¹Bold symbols denote vectors throughout this paper.

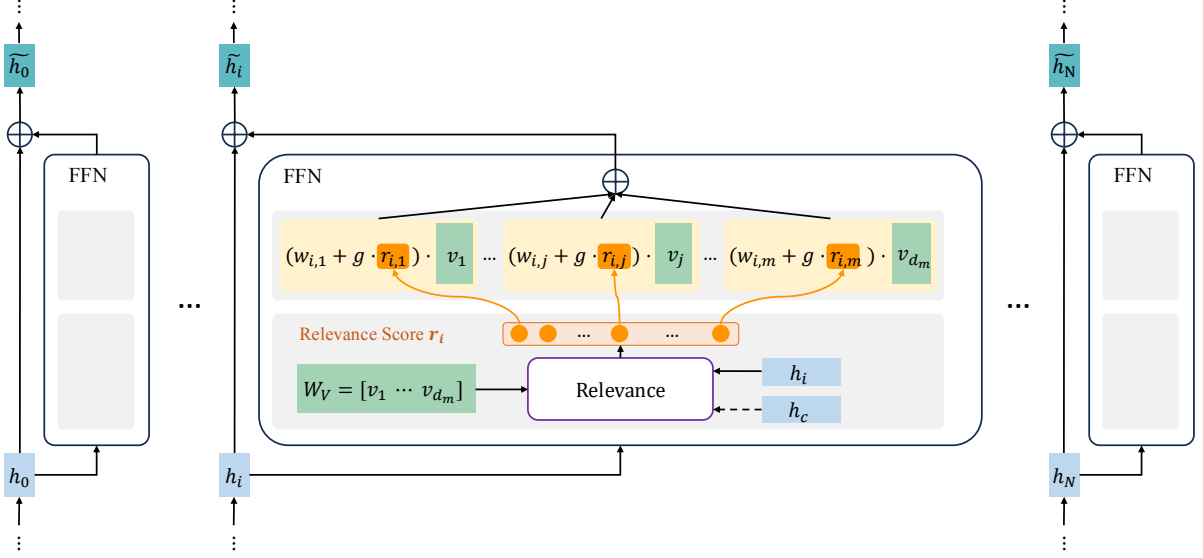


Figure 2: Overview of LLMBRACES. The relevance module computes scores that assess the alignment between each value parameter vector v_j and the input context h_i . For conditional generation, an additional branch computes the conditional relevance scores with respect to task-specific attribute representations h_c . These relevance scores are then combined with the original contribution weights to dynamically modulate the FFN sub-updates.

which provide rich contextual information for the input sequence. To this end, we introduce a **Relevance Module** $R(\cdot)$ that computes a relevance score $r_i^\ell \in \mathbb{R}^{d_m}$ for each FFN layer ℓ and token position i :

$$r_i^\ell = R(W_V^\ell, h_i^\ell) = (R^\ell W_V^\ell)^\top (R^\ell h_i^\ell) / \sqrt{d_r}, \quad (3)$$

where $W_V^\ell \in \mathbb{R}^{d \times d_m}$ denotes the value parameter matrices at layer ℓ , $h_i \in \mathbb{R}^d$ is the hidden states for token i , and $R^\ell \in \mathbb{R}^{d_r \times d}$ is a learnable low-rank projection matrix with orthonormal rows. Here, d is the dimensionality of hidden states and $d_r < d$ is the dimensionality of rank of the subspace. The use of a low-rank projection with orthonormal rows ensures that the relevance scores are computed efficiently while faithfully preserving the geometric structure of the original high-dimensional representations (Halko et al., 2011), providing a reliable foundation for subsequent augmentation.

The resulting vector $r_i^\ell = [r_{i,0}^\ell, \dots, r_{i,d_m}^\ell]^\top$ comprises individual relevance scores, with each component $r_{i,j}^\ell$ corresponds to the j -th sub-update $(w_{i,j}^\ell \cdot v_j^\ell)$. In effect, the relevance score measures how well each sub-update aligns with the input context, allowing us to modulate the contribution of sub-updates based on their relevance to improve model performance.

4.2 Relevance-based Augmentation

Armed with the relevance scores, we refine the FFN layers by incorporating them as an additive adjustment to the contribution of each sub-update. The standard FFN layer at ℓ -th layer (Eq. (2)) is augmented as follows:

$$\text{FFN}_{\text{AUG}}^\ell(h_i^\ell) = \sum_{j=1}^{d_m} (w_{i,j}^\ell + g^\ell \cdot r_{i,j}^\ell) \cdot v_j^\ell, \quad (4)$$

where $r_{i,j}^\ell$ is the relevance score associated with the value vector v_j^ℓ relative to the hidden states h_i^ℓ and g^ℓ is the learnable gating mechanism at layer ℓ that adaptively controls the strength of relevance augmentation. We set $g^\ell = \sigma(g)$ where $\sigma(\cdot)$ is sigmoid activation function, and initialize g to a small value (e.g. $g = -5$) to ensure that the gating mechanism starts nearly at zero. The near zero-initialization supports stable early-stage training and allows the model to progressively incorporate relevance-based augmentation as training proceeds.

We choose the additive formulation because of its simplicity and its ability to handle cases where the original contribution score $w_{i,j}^\ell$ is negative — something that binary filtering or relevance-based multiplication may not accommodate as effectively. By adding the gated relevance score, LLMBRACES dynamically amplifies sub-updates that are aligned with the input context while suppressing those that

may introduce noise or errors. During training, we adopt the standard language modeling objectives. We freeze the weights of the original language backbone model and only fine-tune the newly introduced parameters.

4.3 Task-Specific Relevance

Thus far, we have focused on leveraging the model’s hidden states to compute relevance scores, refining FFN sub-updates to improve prediction accuracy and reliability. LLMBRACES can also be extended for controlled text generation, steering outputs toward desired attributes such as sentiment.

To enable attribute-based control, we introduce a conditional relevance score, r_c , which quantifies the relevance between the value parameter vectors and attribute-specific tokens. For instance, if we wish to generate text with a positive sentiment, the attribute-specific tokens might include words such as “happy”, “joyful”, etc. To compute r_c , we first pass the attribute-specific tokens through the target language model and extract their hidden states over N tokens at each layer ℓ , resulting in $\tilde{\mathbf{h}}_c^\ell \in \mathbb{R}^{N \times d}$. We then apply average pooling on $\tilde{\mathbf{h}}_c^\ell$ to obtain one aggregated representation, followed by multi-layer perceptron (MLP) projection $f_c(\cdot)$:

$$\mathbf{h}_c^\ell = f_c\left(\text{Pooling}(\tilde{\mathbf{h}}_c^\ell)\right) \in \mathbb{R}^d. \quad (5)$$

Once we have the projected attribute-specific representation \mathbf{h}_c^ℓ at each layer ℓ , we compute the conditional relevance score r_c^ℓ in a manner analogous to the primary relevance score:

$$r_c^\ell = R(\mathbf{W}_V^\ell, \mathbf{h}_c^\ell) = (\mathbf{R}^\ell \mathbf{W}_V^\ell)^\top (\mathbf{R}^\ell \mathbf{h}_c^\ell) / \sqrt{d_r} \quad (6)$$

where $\mathbf{R}^\ell \in \mathbb{R}^{d_r \times d}$ is the same learnable low-rank projection matrix with orthonormal rows in Eq. (3).

To integrate this conditional relevance into the overall model, we update the original relevance score (Eq. (3)) as follows:

$$\mathbf{r}_i^\ell \leftarrow \mathbf{r}_i^\ell + s \cdot \mathbf{r}_c^\ell, \quad (7)$$

where s denotes an adjustable scalar indicating the polarity and the intensity of the conditional relevance. We choose a scalar rather than a learnable gating function so that users can manually control the degree to which the model’s outputs are influenced by the desired attribute during inference.

The incorporation of task-specific relevance enables LLMBRACES not only to improve prediction accuracy by emphasizing relevant sub-updates but also to guide the generation process in accordance with specific task attributes.

5 Experimental Setup

Backbone Models and Datasets We implement LLMBRACES on three popular large language models: Qwen2.5-1.5B (Yang et al., 2024), Llama2-7B (Touvron et al., 2023), and Llama3-8B (Dubey et al., 2024), and evaluate the models under three experimental settings: (1) **Supervised Fine-tuning**, where we follow (Hu et al., 2023) and evaluate all methods on Commonsense Reasoning datasets. Specifically, we leverage Commonsense170K (Hu et al., 2023), which is a combination of eight commonsense reasoning datasets, including BoolQ (Clark et al., 2019), PIQA (Bisk et al., 2020), SIQA (Sap et al., 2019), HellaSwag (Zellers et al., 2019), WinoGrande (Sakaguchi et al., 2021), ARC-e, ARC-c (Clark et al., 2018), and OBQA (Mihaylov et al., 2018), for both training and evaluation. (2) **Zero-Shot Generalization**, where we first tune LLMBRACES on Dolly (Conover et al., 2023), an instruction-following dataset containing 15K examples, and evaluate LLMBRACES across multiple datasets targeting different aspects of model behavior. For factual knowledge, we use PopQA (Mallen et al., 2023), TriviaQA (Joshi et al., 2017), and Natural Questions (NQ) (Kwiatkowski et al., 2019). To assess truthfulness and trustworthiness of LLMs, we employ TruthfulQA (Lin et al., 2022). Additionally, we also test all methods on AGI Eval (Zhong et al., 2024), an aggregated benchmark dataset designed to assess overall model capabilities. (3) **Conditional Generation**, where we follow the setting in Han et al. (2024), using Stanford Sentiment Treebank (SST5) (Socher et al., 2013) for sentiment steering and Jigsaw Unintended Bias in Toxicity Classification (cjadams et al., 2019) for toxic suppression.

Baselines For both supervised fine-tuning and zero-shot generalization, we compare LLMBRACES with LoRA (Hu et al.), a parameter-efficient fine-tuning method that uses low-rank matrices to approximate additive weight updates during training. For conditional generation, we consider the following baselines: DExperts (Liu et al., 2021), GeDi (Krause et al., 2021), PromptT5 (Rafael et al., 2020), LoRA (Hu et al.) and LM-Steer (Han et al., 2024). These baselines employ various techniques, including classifier-guided decoding, Bayesian rule-based generation, prompt-based conditioning, parameter-efficient adaptation, and embedding steering, to steer language models’

	Method	Param.	Param. (%)	BoolQ	PIQA	SIQA	HellaS.	WinoG.	ARC-e	ARC-c	OBQA	AVG
Qwen2.5-1.5B	LoRA (r=16)	2.2M	0.14	66.91	80.85	75.64	89.99	76.47	89.48	74.48	84.80	79.83
	LoRA (r=32)	4.4M	0.28	66.67	82.80	75.69	89.52	76.40	88.01	77.04	83.20	79.92
	LLMBRACES (r=16)	0.6M	0.04	66.36	82.59	75.84	89.41	76.32	89.10	78.58	84.00	80.28
	LLMBRACES (r=32)	1.2M	0.08	67.49	83.03	75.38	89.90	76.32	89.69	77.13	83.60	80.32
Llama2-7B	LoRA (r=16)	8.4M	0.12	71.62	84.71	78.86	93.19	82.40	85.98	71.08	84.80	81.58
	LoRA (r=32)	16.8M	0.25	70.52	82.69	79.11	92.00	83.74	84.42	68.94	82.80	80.53
	LLMBRACES (r=16)	2.1M	0.03	71.04	83.41	80.14	93.40	84.21	85.48	71.33	83.00	81.50
	LLMBRACES (r=32)	4.2M	0.06	70.61	83.95	80.09	93.27	82.79	85.90	70.99	82.00	81.20
Llama3-8B	LoRA (r=16)	6.8M	0.08	73.27	88.25	78.92	94.96	87.21	90.95	77.99	86.20	84.72
	LoRA (r=32)	13.6M	0.17	71.31	85.36	79.83	92.97	83.43	87.12	74.15	84.00	82.27
	LLMBRACES (r=16)	2.1M	0.03	72.75	88.52	81.17	95.27	86.42	92.59	80.97	87.80	85.69
	LLMBRACES (r=32)	4.2M	0.05	74.40	89.12	81.37	95.55	86.66	92.97	83.19	88.80	86.51

Table 1: Accuracy comparison on eight *commonsense reasoning* tasks across different LLM architectures. Param. indicates the number of tunable parameters and Param. (%) is calculated by dividing the number of trainable parameters by the number of parameters of the base LM.

	Method	Param.	Param. (%)	TruthfulQA (MC1)	TruthfulQA (MC2)	PopQA	TriviaQA	NQ	AGI	AVG
Qwen2.5-1.5B	LoRA (r=16)	2.2M	0.14	24.60	38.41	17.87	26.47	4.76	30.43	23.76
	LoRA (r=32)	4.4M	0.28	26.19	39.24	17.66	24.95	5.29	29.18	23.75
	LLMBRACES (r=16)	0.6M	0.04	26.07	39.74	17.80	35.40	9.50	31.67	26.70
	LLMBRACES (r=32)	1.2M	0.08	26.07	39.93	18.44	36.26	9.58	32.11	27.07
Llama2-7B	LoRA (r=16)	8.4M	0.12	25.21	39.24	25.27	43.93	8.25	21.13	27.17
	LoRA (r=32)	16.8M	0.25	24.36	38.75	26.08	40.34	8.67	22.66	26.81
	LLMBRACES (r=16)	2.1M	0.03	26.93	40.56	28.14	55.20	20.30	23.83	32.49
	LLMBRACES (r=32)	4.2M	0.06	25.95	38.89	27.20	52.98	16.62	23.94	30.93
Llama3-8B	LoRA (r=16)	6.8M	0.08	27.42	42.62	32.40	34.04	11.14	24.61	28.71
	LoRA (r=32)	13.6M	0.17	29.62	43.14	32.06	23.15	9.11	25.80	27.15
	LLMBRACES (r=16)	2.1M	0.03	28.27	43.91	36.21	66.11	19.53	29.47	37.25
	LLMBRACES (r=32)	4.2M	0.05	28.52	43.46	36.08	64.56	19.25	29.23	36.85

Table 2: Zero-shot comparison on various tasks across different LLM architectures. Param. indicates the number of tunable parameters and Param. (%) is calculated by dividing the number of trainable parameters by the number of parameters of the base LM.

outputs. More details on these baselines can be found in Appendix A.

Implementation Details For supervised fine-tuning, we conduct hyperparameter tuning for both the baseline and different variants of LLMBRACES. Specifically, we randomly sample 300 instances from the training split of Commonsense170K to optimize hyperparameters for commonsense reasoning. The optimal hyperparameters are selected based on performance on held-out validation sets. For zero-shot generalization, we use the same set of hyperparameters selected for Commonsense Reasoning. Further details on the hyperparameter settings for each task are provided in the Appendix B.

6 Results and Discussion

6.1 Supervised Fine-tuning Performance

As shown in Table 1, for supervised fine-tuning performance on commonsense reasoning tasks, LLMBRACES achieves competitive, and even superior,

performance while requiring significantly fewer trainable parameters. For instance, on the Llama3-8B model, LLMBRACES outperforms LoRA by 5.2% while reducing the number of trainable parameters by approximately 38%. These results demonstrate that LLMBRACES can efficiently enhance the model performance and reliability, while matching or surpassing the strong model fine-tuning based baselines, with greater parameter efficiency.

6.2 Zero-shot Generalization

Table 2 presents a comparison of zero-shot performance across various LLM architectures. LLMBRACES consistently surpasses baselines across all tasks while requiring significantly fewer trainable parameters. Specifically, it improves the average performance by 13.9% on Qwen2.5-1.5B, 19.6% on Llama2-7B, and 29.7% on Llama3-8B compared to LoRA, highlighting its effectiveness and generalization in enhancing model performance by

	Method	Neutral to Positive					Neutral to Negative				
		Positivity (%) \uparrow	PPL \downarrow	Dist-1 \uparrow	Dist-2 \uparrow	Dist-3 \uparrow	Positivity \downarrow	PPL \downarrow	Dist-1 \uparrow	Dist-2 \uparrow	Dist-3 \uparrow
Supervised FT	GPT2-Large	53.55	12.48	0.52	0.82	0.85	53.55	12.48	0.52	0.82	0.85
	LoRA (Hu et al.)	26.88	158.56	0.57	0.82	0.83	20.08	192.13	0.55	0.78	0.79
	DExperts (Liu et al., 2021)	94.46	45.83	0.56	0.83	0.83	3.77	45.91	0.60	0.84	0.83
	GeDi (Krause et al., 2021)	86.01	58.41	0.57	0.80	0.79	8.73	84.11	0.63	0.84	0.82
	LMSteer (Han et al., 2024)	90.70	41.20	0.46	0.78	0.83	8.02	57.74	0.48	0.78	0.80
	LLMBRACES (r=16)	91.20	52.63	0.38	0.73	0.81	7.10	63.49	0.42	0.75	0.83
Zero-Shot	PromptT5 (Raffel et al., 2020)	68.12	37.30	0.58	0.78	0.72	25.78	48.60	0.60	0.78	0.70
	LLMBRACES (r=16)	69.82	30.03	0.46	0.77	0.85	12.96	39.78	0.46	0.79	0.85

Table 3: Sentiment steering performance. Results of all baseline methods are taken from Han et al. (2024).

Sentiment	Steering Value	Generation
Positive	1	<u>It is for this reason that I believe our</u> film has more of a chance of being good than bad.
	3	<u>It is for this reason that I believe our</u> film will leave a lasting impression upon you.
	5	<u>It is for this reason that I believe our</u> film is one of the year 's most inspiring and accessible films.
Negative	1	<u>It is for this reason that I believe our</u> film 's title should have been some- thing else.
	3	<u>It is for this reason that I believe our</u> film is going to be overlooked by the viewing public.
	5	<u>It is for this reason that I believe our</u> film is not only mediocre , but also downright ugly and unlikable .

Table 4: Qualitative examples of sentiment steering using LLMBRACES. Underlined text indicates the input provided to the model, while the remaining text is generated by LLMBRACES.

modulating the subupdates in FFNs. LLMBRACES excels particularly in knowledge-intensive tasks such as TriviaQA, PopQA, and NQ. These tasks require the model to accurately answer questions that demand factual knowledge, showcasing the effectiveness of LLMBRACES in retaining knowledge and facilitating reasoning in the zero-shot setting. Moreover, LLMBRACES also shows consistent gains on TruthfulQA MC2, indicating its effectiveness in improving truthfulness and reducing hallucination. All these results confirm our initial hypothesis that modulating the contributions of subupdates in FFNs enhances the model’s prediction process, making it more accurate and reliable.

6.3 Conditional Generation

We further demonstrate the versatility of LLMBRACES in controlling the attributes of generated text, with experiments focusing on sentiment steering and toxic language suppression. We evaluate LLMBRACES under both supervised fine-tuning

and zero-shot settings.

Supervised Sentiment Steering To evaluate LLMBRACES’s ability to control sentiment in generated text, we first define a set of words representing the target sentiment (e.g., positive or negative), as listed in Appendix C.2, and concatenate the representative words into a string of attribute-specific tokens for each sentiment, which are then used to compute the relevance scores following Eqs. (5) to (7). During training, we pair positive tokens with a positive steering value $s = 1$ (Eq. (7)) and negative tokens with a negative steering value $s = -1$, with maximal likelihood as the training objective. When negative texts are available, we fit the model by pairing negative texts with positive tokens alongside a negative steering value $s = -1$, or positive tokens with a positive steering value $s = 1$. This strategy essentially enlarges the training dataset and enhances the model’s robustness. During inference, LLMBRACES adjusts the steering value s to control the sentiment of generated output, offering a flexible and efficient approach for sentiment steering. For benchmarking, we follow Han et al. (2024) which utilizes 5,000 neutral prompts from OpenWebText (Gokaslan and Cohen, 2019) and generates 25 sentences per prompt with GPT2-large. We evaluate the generated text based on its level of positivity, fluency (measured via perplexity), and diversity (Dist- $\{1,2,3\}$). More details for the evaluation metrics are provided in Appendix C.1.

Table 3 shows that LLMBRACES successfully steers the generated text with the desired sentiment, while achieving a reasonable balance of fluency and diversity. Additionally, LLMBRACES enables control over the intensity of sentiment expression. By controlling the steering value s , we can modulate the strength of sentiment in the generated output. Table 4 illustrates how varying the steering value s adjust the intensity of sentiment in the generated text, highlighting LLMBRACES’s capability

	Method	Toxicity ↓		Fluency ↓	Diversity ↑		
		Max. toxicity	Toxicity prob.	Output ppl.	Dist-1	Dist-2	Dist-3
Supervised FT	LoRA (Hu et al.)	0.365	0.210	21.11	0.53	0.85	0.86
	DExperts (Liu et al., 2021)	0.314	0.128	32.41	0.58	0.84	0.84
	GeDi (Krause et al., 2021)	0.363	0.217	60.03	0.62	0.84	0.83
	LMSteer (Han et al., 2024)	0.249	0.089	28.26	0.55	0.84	0.84
	LLMBRACES (r=8)	0.215	0.061	19.44	0.54	0.85	0.86
Zero-Shot	PromptT5 (Raffel et al., 2020)	0.320	0.172	55.10	0.58	0.76	0.70
	LLMBRACES (r=16)	0.172	0.027	31.91	0.51	0.83	0.86

Table 5: Toxic Language Suppression. Results of all baseline methods are taken from Han et al. (2024).

to modulate the emotional tone of the output in a fine-grained manner.

Supervised Toxic Language Suppression To mitigate toxic language in generated text, we construct attribute-specific tokens x_c using the toxic words from WORDFILTER (Gehman et al., 2020). These tokens are used to compute the conditional relevance score, guiding the model away from toxic completions by assigning a negative steering value s . During training, we pair toxic texts with a positive steering value ($s = 1$) and non-toxic texts with a negative steering value ($s = -1$).

Following the evaluation protocol of Han et al. (2024), we use 10K nontoxic prompts from RE-ALTOXICITYPROMPTS (Gehman et al., 2020) and randomly generate 25 sentences using GPT2-large. During inference, we apply a steering value $s = -5$ to suppress the toxic language. Similar to the sentiment task, we assess the generated text based on fluency and diversity, and measure toxicity using the Perspective API². More details can be found in Appendix C.1. As shown in Table 5, LLMBRACES significantly reduces the model’s propensity for toxic language by 15.8% while maintaining superior fluency and diversity, demonstrating its effectiveness in suppressing toxic language by decreasing the contribution of toxic-associated sub-updates during generation. Qualitative examples can be found in Appendix D.

Zero-shot Conditional Generation We further evaluate LLMBRACES’s performance in zero-shot conditional generation. For this, we construct a new dataset, *StyleDataset*, by selecting 10K sentences from OpenWebText (Gokaslan and Cohen, 2019). For each sentence, we prompt Llama3.1-8B-instruct (Dubey et al., 2024) to generate a short list of descriptive attribute words that capture the sen-

	Method	Commonsense AVG	Zero-shot AVG
Qwen2.5-1.5B	LLMBRACES (r=16)	80.28	26.70
	- w/o REL	78.24	26.50
	LLMBRACES (r=32)	80.32	27.07
	- w/o REL	78.20	26.62
Llama2-7B	LLMBRACES (r=16)	81.50	32.49
	- w/o REL	77.46	30.59
	LLMBRACES (r=32)	81.20	30.93
	- w/o REL	77.34	30.70

Table 6: Effect of relevance module comparison.

tence’s style, resulting in a dataset of 10K sentence-attribute pairs. The prompting instructions are detailed in Appendix C.2. We then train LLMBRACES on *StyleDataset* by pairing the attribute tokens with a positive steering value ($s = 1$) and maximizing the likelihood of the corresponding sentence. Since LLMBRACES is not fine-tuned specifically for sentiment or toxicity tasks, we can evaluate its zero-shot generalization on both tasks. As shown in Tables 3 and 5, LLMBRACES consistently outperforms the zero-shot baseline in both sentiment control and toxicity suppression. Notably, for toxic suppression, the zero-shot performance of LLMBRACES even exceeds that of its fine-tuned counterpart. These results highlight LLMBRACES’s robustness and adaptability in zero-shot conditional generation tasks, eliminating the need for task-specific fine-tuning.

6.4 Effect of Relevance Module

To better understand the contribution of our proposed relevance module, we conduct an ablation study by training an additional variant in which, instead of using the computed relevance score, we initialize the relevance score as a learnable parameter. We refer to this method as LLMBRACES w/o REL. We then compare the model’s performance on commonsense reasoning and zero-shot tasks. As shown in Table 6, LLMBRACES consistently out-

²<https://perspectiveapi.com/>

performs LLMBRACES w/o REL across all tasks, demonstrating the effectiveness of our proposed relevance module for dynamically adjusting sub-update contributions.

7 Conclusion

In this work, we introduce LLMBRACES, a novel and versatile framework that enhances the performance of pretrained LLMs by dynamically adjusting the contributions of sub-updates based on their relevance to the input text or target output style. Our experimental results across multiple language model backbones demonstrate the effectiveness of LLMBRACES in improving model performance in both supervised fine-tuning and zero-shot settings while requiring far fewer trainable parameters. Moreover, LLMBRACES also support controlled generation aligned with specific target style. Importantly, our framework is architecture-agnostic and can be seamlessly integrated into any model containing a feed-forward layer.

Limitations

While LLMBRACES is designed to be compatible with any architecture that includes a feed-forward layer, our current study focuses solely on text-based LLMs, without evaluating its applicability to multimodal language model and image generation models. In future work, we plan to extend LLMBRACES to a broader range of architectures, including multimodal and image generation models, to assess its versatility and impact across diverse architectures and domains. Additionally, although our approach leverages sub-updates to refine model predictions, our work does not fully interpret the internal mechanisms of LLMs. Demystifying the internal states of these models remains a challenging task, and further research is needed.

Acknowledgements

This research is supported by the award No. 2238940 from the Faculty Early Career Development Program (CAREER) of the National Science Foundation (NSF). The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for governmental purposes notwithstanding any copyright annotation therein.

References

- Anthropic. 2023. Model card: Claude 2. <https://www-files.anthropic.com/production/images/Model-Card-Claude-2.pdf>.
- Nora Belrose, Zach Furman, Logan Smith, Danny Halawi, Igor Ostrovsky, Lev McKinney, Stella Biderman, and Jacob Steinhardt. 2023. Eliciting latent predictions from transformers with the tuned lens. *arXiv preprint arXiv:2303.08112*.
- Yonatan Bisk, Rowan Zellers, Jianfeng Gao, Yejin Choi, et al. 2020. Piqa: Reasoning about physical commonsense in natural language. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, pages 7432–7439.
- Sid Black, Stella Biderman, Eric Hallahan, Quentin Anthony, Leo Gao, Laurence Golding, Horace He, Connor Leahy, Kyle McDonell, Jason Phang, et al. 2022. Gpt-neox-20b: An open-source autoregressive language model. *arXiv preprint arXiv:2204.06745*.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.
- cjadams, Daniel Borkan, inversion, Jeffrey Sorensen, Lucas Dixon, Lucy Vasserman, and nithum. 2019. Jigsaw unintended bias in toxicity classification. Kaggle.
- Christopher Clark, Kenton Lee, Ming-Wei Chang, Tom Kwiatkowski, Michael Collins, and Kristina Toutanova. 2019. Boolq: Exploring the surprising difficulty of natural yes/no questions. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 2924–2936.
- Peter Clark, Isaac Cowhey, Oren Etzioni, Tushar Khot, Ashish Sabharwal, Carissa Schoenick, and Oyvind Tafjord. 2018. Think you have solved question answering? try arc, the ai2 reasoning challenge. *arXiv preprint arXiv:1803.05457*.
- Mike Conover, Matt Hayes, Ankit Mathur, Jianwei Xie, Jun Wan, Sam Shah, Ali Ghodsi, Patrick Wendell, Matei Zaharia, and Reynold Xin. 2023. [Free dolly: Introducing the world’s first truly open instruction-tuned llm](#).
- Damai Dai, Li Dong, Yaru Hao, Zhifang Sui, Baobao Chang, and Furu Wei. 2022. Knowledge neurons in pretrained transformers. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8493–8502.
- Guy Dar, Mor Geva, Ankit Gupta, and Jonathan Berant. 2023. Analyzing transformers in embedding space.

- In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 16124–16170.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186.
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.
- Nelson Elhage, Neel Nanda, Catherine Olsson, Tom Henighan, Nicholas Joseph, Ben Mann, Amanda Askell, Yuntao Bai, Anna Chen, Tom Conerly, et al. 2021. A mathematical framework for transformer circuits. *Transformer Circuits Thread*, 1(1):12.
- Javier Ferrando and Elena Voita. 2024. Information flow routes: Automatically interpreting language models at scale. *arXiv preprint arXiv:2403.00824*.
- Matthew Finlayson, Aaron Mueller, Sebastian Gehrmann, Stuart M Shieber, Tal Linzen, and Yonatan Belinkov. 2021. Causal analysis of syntactic agreement mechanisms in neural language models. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1828–1843.
- Samuel Gehman, Suchin Gururangan, Maarten Sap, Yejin Choi, and Noah A. Smith. 2020. *RealToxicityPrompts: Evaluating neural toxic degeneration in language models*. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 3356–3369, Online. Association for Computational Linguistics.
- Atticus Geiger, Zhengxuan Wu, Christopher Potts, Thomas Icard, and Noah Goodman. 2024. Finding alignments between interpretable causal variables and distributed neural representations. In *Causal Learning and Reasoning*, pages 160–187. PMLR.
- Mor Geva, Jasmijn Bastings, Katja Filippova, and Amir Globerson. 2023a. Dissecting recall of factual associations in auto-regressive language models. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 12216–12235.
- Mor Geva, Avi Caciularu, Kevin Wang, and Yoav Goldberg. 2022. Transformer feed-forward layers build predictions by promoting concepts in the vocabulary space. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 30–45.
- Mor Geva, Avi Caciularu, Kevin Wang, Yoav Goldberg, Thooran Rajamanoharan, and János Kramár. 2023b. *Fact finding: Attempting to reverse-engineer factual recall on the neuron level*.
- Mor Geva, Roei Schuster, Jonathan Berant, and Omer Levy. 2021. Transformer feed-forward layers are key-value memories. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 5484–5495.
- Asma Ghandeharioun, Avi Caciularu, Adam Pearce, Lucas Dixon, and Mor Geva. 2024. Patchscope: A unifying framework for inspecting hidden representations of language models. *arXiv preprint arXiv:2401.06102*.
- Aaron Gokaslan and Vanya Cohen. 2019. Open-webtext corpus. <http://Skylion007.github.io/OpenWebTextCorpus>.
- Nathan Halko, Per-Gunnar Martinsson, and Joel A Tropp. 2011. Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions. *SIAM review*, 53(2):217–288.
- Chi Han, Jialiang Xu, Manling Li, Yi Fung, Chenkai Sun, Nan Jiang, Tarek Abdelzaher, and Heng Ji. 2024. Word embeddings are steers for language models. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 16410–16430.
- Tom Hartvigsen, Swami Sankaranarayanan, Hamid Palangi, Yoon Kim, and Marzyeh Ghassemi. 2024. Aging with grace: Lifelong model editing with discrete key-value adapters. *Advances in Neural Information Processing Systems*, 36.
- Xingwei He. 2021. Parallel refinements for lexically constrained text generation with bart. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 8653–8666.
- Edward J Hu, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, Weizhu Chen, et al. Lora: Low-rank adaptation of large language models. In *International Conference on Learning Representations*.
- Zhiqiang Hu, Lei Wang, Yihuai Lan, Wanyu Xu, Ee-Peng Lim, Lidong Bing, Xing Xu, Soujanya Poria, and Roy Lee. 2023. Llm-adapters: An adapter family for parameter-efficient fine-tuning of large language models. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 5254–5276.
- Zeyu Huang, Yikang Shen, Xiaofeng Zhang, Jie Zhou, Wenge Rong, and Zhang Xiong. 2023. Transformer-patcher: One mistake worth one neuron. In *The Eleventh International Conference on Learning Representations*.

- Albert Q Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, et al. 2023. Mistral 7b. *arXiv preprint arXiv:2310.06825*.
- Mandar Joshi, Eunsol Choi, Daniel S Weld, and Luke Zettlemoyer. 2017. Triviaqa: A large scale distantly supervised challenge dataset for reading comprehension. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1601–1611.
- Shahar Katz, Yonatan Belinkov, Mor Geva, and Lior Wolf. 2024. Backward lens: Projecting language model gradients into the vocabulary space. *arXiv preprint arXiv:2402.12865*.
- Nitish Shirish Keskar, Bryan McCann, Lav R Varshney, Caiming Xiong, and Richard Socher. 2019. Ctrl: A conditional transformer language model for controllable generation. *arXiv preprint arXiv:1909.05858*.
- Ben Krause, Akhilesh Deepak Gotmare, Bryan McCann, Nitish Shirish Keskar, Shafiq Joty, Richard Socher, and Nazneen Fatema Rajani. 2021. Gedi: Generative discriminator guided sequence generation. In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 4929–4952.
- Tom Kwiatkowski, Jennimaria Palomaki, Olivia Redfield, Michael Collins, Ankur Parikh, Chris Alberti, Danielle Epstein, Illia Polosukhin, Jacob Devlin, Kenton Lee, Kristina Toutanova, Llion Jones, Matthew Kelcey, Ming-Wei Chang, Andrew M. Dai, Jakob Uszkoreit, Quoc Le, and Slav Petrov. 2019. [Natural questions: A benchmark for question answering research](#). *Transactions of the Association for Computational Linguistics*, 7:452–466.
- Stephanie Lin, Jacob Hilton, and Owain Evans. 2022. Truthfulqa: Measuring how models mimic human falsehoods. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3214–3252.
- Alisa Liu, Maarten Sap, Ximing Lu, Swabha Swayamdipta, Chandra Bhagavatula, Noah A Smith, and Yejin Choi. 2021. Dexperts: Decoding-time controlled text generation with experts and anti-experts. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 6691–6706.
- Sheng Liu, Haotian Ye, Lei Xing, and James Y Zou. In-context vectors: Making in context learning more effective and controllable through latent space steering. In *Forty-first International Conference on Machine Learning*.
- Alex Mallen, Akari Asai, Victor Zhong, Rajarshi Das, Daniel Khashabi, and Hannaneh Hajishirzi. 2023. [When not to trust language models: Investigating effectiveness of parametric and non-parametric memories](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 9802–9822, Toronto, Canada. Association for Computational Linguistics.
- Kevin Meng, David Bau, Alex Andonian, and Yonatan Belinkov. 2022. Locating and editing factual associations in gpt. *Advances in Neural Information Processing Systems*, 35:17359–17372.
- Kevin Meng, Arnab Sen Sharma, Alex J Andonian, Yonatan Belinkov, and David Bau. 2023. Mass-editing memory in a transformer. In *The Eleventh International Conference on Learning Representations*.
- Jack Merullo, Carsten Eickhoff, and Ellie Pavlick. Circuit component reuse across tasks in transformer language models. In *The Twelfth International Conference on Learning Representations*.
- Todor Mihaylov, Peter Clark, Tushar Khot, and Ashish Sabharwal. 2018. Can a suit of armor conduct electricity? a new dataset for open book question answering. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2381–2391.
- nostalgebraist. 2020. Interpreting gpt: the logit lens. <https://www.lesswrong.com/posts/AcKRB8wDpdaN6v6ru/interpreting-gpt-the-logit-lens>.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of machine learning research*, 21(140):1–67.
- Emily Reif, Daphne Ippolito, Ann Yuan, Andy Coenen, Chris Callison-Burch, and Jason Wei. 2022. [A recipe for arbitrary text style transfer with large language models](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 837–848, Dublin, Ireland. Association for Computational Linguistics.
- Keisuke Sakaguchi, Ronan Le Bras, Chandra Bhagavatula, and Yejin Choi. 2021. Winogrande: An adversarial winograd schema challenge at scale. *Communications of the ACM*, 64(9):99–106.
- Maarten Sap, Hannah Rashkin, Derek Chen, Ronan Le Bras, and Yejin Choi. 2019. Social iqa: Commonsense reasoning about social interactions. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 4463–4473.

- Anton Sinitin, Vsevolod Plohotnyuk, Dmitry Pyrkun, Sergei Popov, and Artem Babenko. 2019. Editable neural networks. In *International Conference on Learning Representations*.
- Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D Manning, Andrew Y Ng, and Christopher Potts. 2013. Recursive deep models for semantic compositionality over a sentiment treebank. In *Proceedings of the 2013 conference on empirical methods in natural language processing*, pages 1631–1642.
- Alessandro Stolfo, Yonatan Belinkov, and Mrinmaya Sachan. 2023. A mechanistic interpretation of arithmetic reasoning in language models using causal mediation analysis. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 7035–7052.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajwal Bhargava, Shruti Bhosale, et al. 2023. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.
- Alexander Matt Turner, Lisa Thiergart, Gavin Leech, David Udell, Juan J Vazquez, Ulisse Mini, and Monte MacDiarmid. 2023. Activation addition: Steering language models without optimization. *arXiv preprint arXiv:2308.10248*.
- Elena Voita, Javier Ferrando, and Christoforos Nalmpantis. 2023. Neurons in large language models: Dead, n-gram, positional. *arXiv preprint arXiv:2309.04827*.
- Kevin Ro Wang, Alexandre Variengien, Arthur Conmy, Buck Shlegeris, and Jacob Steinhardt. Interpretability in the wild: a circuit for indirect object identification in gpt-2 small. In *The Eleventh International Conference on Learning Representations*.
- Thomas Wolf. 2020. Transformers: State-of-the-art natural language processing. *arXiv preprint arXiv:1910.03771*.
- An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoran Wei, et al. 2024. Qwen2. 5 technical report. *arXiv preprint arXiv:2412.15115*.
- Yunzhi Yao, Ningyu Zhang, Zekun Xi, Mengru Wang, Ziwen Xu, Shumin Deng, and Huajun Chen. 2024. Knowledge circuits in pretrained transformers. *arXiv preprint arXiv:2405.17969*.
- Rowan Zellers, Ari Holtzman, Yonatan Bisk, Ali Farhadi, and Yejin Choi. 2019. Hellaswag: Can a machine really finish your sentence? In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4791–4800.
- Hanqing Zhang, Haolin Song, Shaoyu Li, Ming Zhou, and Dawei Song. 2023. A survey of controllable text generation using transformer-based pre-trained language models. *ACM Computing Surveys*, 56(3):1–37.
- Wanjun Zhong, Ruixiang Cui, Yiduo Guo, Yaobo Liang, Shuai Lu, Yanlin Wang, Amin Saied, Weizhu Chen, and Nan Duan. 2024. Agieval: A human-centric benchmark for evaluating foundation models. In *Findings of the Association for Computational Linguistics: NAACL 2024*, pages 2299–2314.
- Wangchunshu Zhou, Yuchen Eleanor Jiang, Ethan Wilcox, Ryan Cotterell, and Mrinmaya Sachan. 2023. Controlled text generation with natural language instructions. In *International Conference on Machine Learning*, pages 42602–42613. PMLR.
- Andy Zou, Long Phan, Sarah Chen, James Campbell, Phillip Guo, Richard Ren, Alexander Pan, Xuwang Yin, Mantas Mazeika, Ann-Kathrin Dombrowski, et al. 2023. Representation engineering: A top-down approach to ai transparency. *arXiv preprint arXiv:2310.01405*.

A Baselines

We compare our approach against the following baselines: 1) LoRA (Hu et al.) leverages low-rank approximations of parameter matrices for parameter-efficient fine-tuning; 2) DExperts (Liu et al., 2021) is a decoding-time method for controlled text generation that combines a pretrained language model with trained positive and negative label classifiers, using the difference in their classifiers’ scores to adjust the model’s original logits; 3) GeDi (Krause et al., 2021) guides generation at each step by computing classification probabilities for all possible next tokens via Bayes rule; 4) PromptT5 (Raffel et al., 2020) is a pre-trained T5 (Raffel et al., 2020) model optimized for prompt-based task solving. Following Han et al. (2024), the model is prompted with “Complete this sentence so that it embodies a positive/negative sentiment:” to generate sentiment-specific outputs; 5) LMSteer (Han et al., 2024) enables control over language model generation styles by applying a learnable linear transformation to output word embeddings.

B Implementation Details

For all experiments, we employ the AdamW optimizer with $\beta_1 = 0.9$ and $\beta_2 = 0.999$. We use a cosine learning rate schedule, with a warmup ratio of 0.2, and decay the final learning rate down to 20% of the peak learning rate. Due to computing resource constraints, we limited our hyperparameter tuning to the learning rate and batch size. Specifically, the search space for the learning rate

	Method	Learning Rate	Batch Size
Qwen2.5-1.5B	LoRA (r=16)	3e-4	16
	LoRA (r=32)	3e-4	16
	LLMBRACES (r=16)	2e-3	16
	- w/o REL	3e-4	16
	LLMBRACES (r=32)	2e-3	16
	- w/o REL	3e-4	16
Llama2-7B	LoRA (r=16)	3e-4	16
	LoRA (r=32)	3e-4	16
	LLMBRACES (r=16)	4e-2	16
	- w/o REL	3e-4	16
	LLMBRACES (r=32)	4e-2	16
	- w/o REL	3e-4	16
Llama3-8B	LoRA (r=16)	3e-4	16
	LoRA (r=32)	3e-4	16
	LLMBRACES (r=16)	4e-2	16
	LLMBRACES (r=32)	4e-2	16

Table 7: Final hyperparameter settings chosen after tuning over the hyperparameter search spaces.

is {4e-2, 2e-3, 3e-4, 2e-5}, and for the batch size, it is {16, 32}. Table 7 lists the final hyperparameters chosen. Following Hu et al. (2023), we set the number of epochs to 3 for the supervised common-sense reasoning tasks, and we use 6 epochs for the zero-shot setting, considering the smaller size of the instruction-tuning dataset. All experiments are conducted on A100 80G GPUs.

Steering Value	Generation
0	According to other witnesses the men and women were killed while being questioned.
-3	According to other witnesses the men were taken to a nearby village to be questioned.
-5	According to other witnesses the men and women were escorted into the home.

Table 8: Qualitative examples of toxic language suppression using LLMBRACES. Underlined text indicates the input provided to the model, while the remaining text is generated by LLMBRACES.

C Conditional Generation

C.1 Evaluation Metrics

We follow the evaluation protocol in Han et al. (2024) and assess the quality of the generated text using fluency (perplexity) and diversity metrics. Specifically, fluency is measured by the perplexity score computed with GPT2-XL, while diversity is quantified using the proportion of distinct {1, 2, 3}-grams for Dist-{1, 2, 3}.

For sentiment evaluation, we employ HuggingFace’s sentiment classifier (Wolf, 2020) and com-

pute the average percentage of positive outputs per prompt, which we refer to as the “Positivity” score.

For toxic language suppression, we evaluate the toxicity of the generated outputs using the Perspective API³ to grade the text. We consider two metrics: the maximal toxicity of generations on each prompt averaged across prompts (“Max. toxicity”) and the averaged probability that a generation exceeds a toxicity threshold of 0.5 (“Toxicity prob.”).

C.2 Construction of StyleDataset

Table 10 shows the instruction used to prompt the Llama3.1-8B-instruct (Dubey et al., 2024) model to generate a short list of descriptive attributes capturing a sentence’s style. We also provide several example instances from *StyleDataset* in Table 11.

D Qualitative Results

Table 8 illustrates how varying the steering value s reduce the toxicity of generated text.

³<https://perspectiveapi.com/>

Positive Sentiment

"Amazing", "Astonishing", "Breathtaking", "Brilliant", "Delightful", "Exceptional", "Extraordinary", "Fantastic", "Glorious", "Incredible", "Magnificent", "Marvelous", "Outstanding", "Phenomenal", "Remarkable", "Spectacular", "Splendid", "Stupendous", "Superb", "Terrific", "Thrilling", "Unbelievable", "Wonderful", "Bliss", "Brilliance", "Euphoria", "Excellence", "Genius", "Glory", "Harmony", "Joy", "Miracle", "Perfection", "Radiance", "Splendor", "Triumph", "Wonder", "Adore", "Amaze", "Astonish", "Celebrate", "Elevate", "Enchant", "Energize", "Excite", "Inspire", "Marvel", "Overjoy", "Rejoice", "Thrill", "Absolutely", "Astonishingly", "Exceptionally", "Exquisitely", "Fantastically", "Incredibly", "Marvelously", "Remarkably", "Spectacularly", "Stunningly", "Unbelievably", "Wonderfully"

Negative Sentiment

"abysmal", "appalling", "atrocious", "awful", "catastrophic", "cursed", "damning", "dismal", "disastrous", "disturbing", "evil", "ghastly", "horrendous", "horrible", "horrific", "hopeless", "loathsome", "lousy", "monstrous", "nauseating", "nightmarish", "outrageous", "pathetic", "pitiful", "revolting", "shameful", "shocking", "terrible", "terrifying", "toxic", "tragic", "unbearable", "unforgivable", "unlivable", "unthinkable", "unworthy", "vicious", "wretched", "abominable", "deplorable", "disgusting", "grotesque", "hideous", "inexcusable", "insufferable", "intolerable", "malevolent", "nefarious", "offensive", "repugnant", "scandalous", "sinister", "vile", "villainous", "wicked", "fearful", "bitter", "frustrated", "miserable", "lonely", "selfish"

Toxic Words

"fuck", "shit", "hell", "abbo", "abo", "beeyotch", "biatch", "bitch", "chinaman", "chinamen", "chink", "coolie", "coon", "crazie", "crazy", "crip", "cuck", "cunt", "dago", "daygo", "dego", "dick", "douchebag", "dumb", "dyke", "eskimo", "fag", "faggot", "fatass", "fatso", "gash", "gimp", "gip", "golliwog", "gook", "gyp", "gypsy", "half-breed", "halfbreed", "heeb", "homo", "hooker", "idiot", "insane", "insanitie", "insanity", "jap", "kaffer", "kaffir", "kaffir", "kaffre", "kafir", "kike", "kraut", "lame", "lardass", "lesbo", "lunatic", "mick", "negress", "negro", "nig", "nig-nog", "nigga", "nigger", "nigguh", "nip", "pajeet", "paki", "pickaninnie", "pickaninny", "prostitute", "pussie", "pussy", "raghead", "retard", "sambo", "shemale", "skank", "slut", "soyboy", "spade", "sperg", "spic", "spook", "squaw", "street-shitter", "tard", "tits", "titt", "trannie", "tranny", "twat", "wetback", "whore", "wigger", "wop", "yid", "zog"

Table 9: The list of attribute tokens employed for conditional generation. Please note that the words listed under "toxic words" could be offensive to read.

You are a highly intelligent and detail-oriented assistant, renowned for your ability to leverage extensive knowledge and analyze questions carefully. You think through each problem step by step and always provide structured, clear, and concise answers.

Given a sentence, your task is to determine its most prominent stylistic property and then generate a short list of descriptive attribute words that capture that style. The attributes may relate to properties such as formality, technicality, narrative style, sentiment, humor, etc.

Below are some examples:

Example 1 – Formality/Professional Tone:

Sentence:

"Dear Sir or Madam, please be advised that your account will be debited on the next business day."

Attribute words:

"formal, professional, courteous, business-like"

Example 2 – Informal/Conversational Tone:

Sentence:

"Hey, what's up? Wanna hang out later?"

Attribute words:

"informal, casual, friendly, colloquial"

Example 3 – Technical/Scientific Language:

Sentence:

"The quantum state of the particle was determined via the Fourier transform of its wavefunction."

Attribute words:

"technical, scientific, precise, academic"

Example 4 – Narrative/Story-like Style:

Sentence:

"Once upon a time, in a land shrouded in mystery, a young adventurer embarked on a quest."

Attribute words:

"narrative, descriptive, imaginative, literary"

Example 5 – Humorous/Witty Tone:

Sentence:

"I tried to catch some fog yesterday, but I mist."

Attribute words:

"humorous, playful, pun-based, witty"

Now, for the following sentence, provide a list of attribute words that best describe its style. You should answer with the following structure:

Explain:

<explanation>

Attribute words:

<answer>

Sentence:"{sentence}"

Explain:

Table 10: The instruction for prompting Llama3.1-8B-instruct (Dubey et al., 2024) for generating list of attribute words for each sentence.

Sentence	Style
On behalf of the team, I am pleased to announce that Spring Boot 1.5.4 has been released and is available now from repo.spring.io and Maven Central.	formal, technical, professional, announcement-like, dry, informative
Step 1: Preheat oven to 400F.	technical, instructional, straightforward, directive
I like to hear music while I cook, but nothing too headbangy any more. Curtis Mayfield, Marvin Gaye, pre-disco funk, Isaac Hayes and Brothers Johnson and I'm happy.	casual, conversational, personal, emotive
You could follow Alt News posts either via our Facebook page or by following us on Twitter or by subscribing to our E-mail updates.	direct, informative, instructive, utilitarian
It seems not everything is more expensive these days.	observing, casual, sarcastic, conversational, observational, casual.
This photograph, taken by an astronaut aboard the International Space Station, shows the straight line of the Corinth Canal as it crosses a narrow isthmus between mainland Greece (right) and the Peloponnese Peninsula. The canal cuts through the narrowest part of the isthmus of Corinth. The towns of Corinth and Isthmia stand near the west and east ends (north is to the upper right). Near the center of the image, a highway crosses the canal and connects Athens to the Peloponnese.	informative, formal, descriptive, geographic
The Liberal position is consistent with Bill C-60, their 2005 copyright bill that linked the digital lock rules to actual copyright infringement and did not establish a ban on the tools that can be used to circumvent digital locks.	formal, analytical, technical, informative, explanatory, objective

Table 11: Example data instances from *StyleDataset*.