

SELF-CONTROL OF LLM BEHAVIORS BY COMPRESSING SUFFIX GRADIENT INTO PREFIX CONTROLLER

Anonymous authors

Paper under double-blind review

ABSTRACT

We propose SELFCONTROL, an inference-time model control method utilizing gradients to control the behavior of large language models (LLMs) without explicit human annotations. Given a desired behavior expressed in a natural language suffix string concatenated to the input prompt, SELFCONTROL computes gradients of the LLM’s self-evaluation of the suffix with respect to its latent representations. The gradients are used to directly control the auto-regressive generation process towards desired behaviors, which eliminates human supervision, achieves precise and transparent control, and offers on-the-fly adaptability. To further enhance efficiency, we introduce SELFCONTROL_{PREFIX}, a compact module that encapsulates the learned representations from gradients into a PREFIXCONTROLLER, facilitating efficient inference-time control with no latency compared to the original model and allowing control for multiple behaviors simultaneously. Our experiments demonstrate SELFCONTROL’s efficacy across multiple domains, where it improves over SOTA for **8.3%** in detoxification, **3.1%** in truthfulness enhancement, **4%~10%** in controlling on emotion tones, and **48.2%** in privacy protection, i.e., completely remove privacy leakage issue. We make available an anonymous [Google Colab demo](#) and release our [code](#).

1 INTRODUCTION

Recent advancements in Large Language Models (LLMs) like GPT (Achiam et al., 2023), Claude (Anthropic, 2023), Gemini (DeepMind, 2023), and LLaMA (Touvron et al., 2023) have showcased capabilities in understanding and generating human-like text. Users can interact with LLMs by writing their instructions in natural language, served as prompt input to LLMs. However, LLMs may often behave in ways that deviate from human intentions (Shen et al., 2023). Such deviations include the generation of toxic content, leakage of private information (Wang et al., 2023), or inappropriate emotional tones (Dathathri et al., 2020). To address these issues, controlling LLMs to ensure alignment with desired behaviors—such as non-toxicity, truthfulness (Ouyang et al., 2022; Dathathri et al., 2020), privacy preservation (Wang et al., 2023), and emotional tone adjustments (Li et al., 2023)—is crucial (see Figure 1).

One approach to enhance LLM control involves fine-tuning model parameters via online (Ouyang et al., 2022) and offline (Rafailov et al., 2024) Reinforcement Learning (RL) methods. However, these methods face several challenges: 1) They demand substantial effort to collect large-scale human-annotated preference data (Bai et al.,

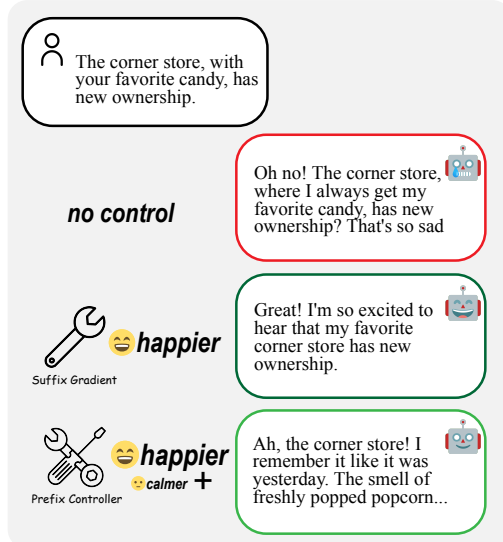


Figure 1: Our SELFCONTROL and SELFCONTROL_{PREFIX} are able to control LLM behaviors, e.g., emotion. With SELFCONTROL, you can obtain the suffix gradient for the desired attribute for precise control, while SELFCONTROL_{PREFIX} enables the composition of these attributes with PREFIXCONTROLLER.

2022a); 2) The control process is indirect and lacks transparency, which may fail when encountering out-of-domain behaviors not represented in the preference data (Huang et al., 2023); 3) Incorporating new desired behaviors necessitates additional fine-tuning, which can potentially compromise the control quality of previously aligned behaviors (Qi et al., 2023).

Besides aligning LLM behavior during training time, recently much research has been focused on controlling LLM at inference time, e.g., through latent representation engineering (RepE)(Zou et al., 2023a). These works still require curating a set of contrastive pairs as a demonstration. Building upon this, in this paper, we’re studying whether we can control LLMs’ behavior via their self-evaluation, i.e. use LLMs as a judge to assess the quality of their own outputs (Kadavath et al., 2022; Bai et al., 2022b; Zheng et al., 2023; Yuan et al., 2024). We thus introduce SELFCONTROL, a novel gradient-based framework for precise LLM behavior control.

The core idea of SELFCONTROL is to leverage the LLM’s self-evaluation to control its behavior. For each input prompt, we formulate a desired behavior as a natural language question, asking the model to assess whether its output aligns with the specified behavior. We term the question as *suffix string* and concatenate the input prompt, the model’s output, and the suffix string, feeding this combined input back into the model. We then compute the likelihood of the behavior-aligned response to the suffix string as a learning signal for behavior control, which we named as *suffix score*. Then, we compute the suffix score’s gradient with respect to the latent representations of the original input, which term as *suffix gradients*. These suffix gradients are then utilized to update the latent representations, controlling the LLM’s behavior towards the desired outcome. We run this procedure multiple times to iteratively update the input’s latent representations, each time using the modified representations to generate new model outputs that increasingly align with the desired behavior.

SELFCONTROL offers several advantages over traditional fine-tuning approaches: 1) **Elimination of human-annotation**: SELFCONTROL leverages the model’s self-evaluation as learning signals, substantially reducing the effort and resources required for preference data collection and scaling. 2) **Precise and transparent control**: SELFCONTROL operates at inference time and directly modifies the latent representations, which allows for explicit behavior specification and fine-grained control, consequently enhancing control interpretability. 3) **On-the-fly adaptability**: SELFCONTROL does not alter model parameters, enabling easy implementation of behaviors and control of combinations of multiple behaviors, thus providing unparalleled flexibility. SELFCONTROL demonstrates superior performance compared to contrastive learning-based control methods, particularly in areas such as detoxification, truthfulness enhancement, privacy protection, and emotion control. As illustrated in Figure 1, SELFCONTROL exhibits remarkable flexibility in controlling LLM for multiple attributes simultaneously, e.g., happiness and calmness.

To enhance its adaptability, efficiency and compositionality, we further propose SELFCONTROL_{PREFIX} on top of SELFCONTROL as a general controller across inputs. The core module of SELFCONTROL_{PREFIX} is the PREFIXCONTROLLER, a prompt-based adapter (Hu et al., 2021; Zhang et al., 2023) optimized to match the latent representations conditioned on this PREFIXCONTROLLER to the latent representations under regular SELFCONTROL. PREFIXCONTROLLER brings ideas from prefix-prompt tuning research (Shin et al., 2020; Li & Liang, 2021; Yang et al., 2023) to achieve efficient control, which has almost no latency compared to the original model, and greatly outperforms other control baselines. Furthermore, we show that PREFIXCONTROLLER is a learnable and composable module that can be easily integrated into the LLM to control multiple model behaviors simultaneously (e.g., being happier, while staying calm), shown in Figure 1, thereby enhancing the practicality of SELFCONTROL for real-world applications.

In summary, our primary contributions are as follows:

- We introduce SELFCONTROL, a gradient-based LLM control framework that leverages the model’s self-evaluation to eliminate the need for human-annotated data, offering more efficient, precise, transparent, and adaptable control.
- We further develop SELFCONTROL_{PREFIX} using PREFIXCONTROLLER, a PEFT (parameter-efficient fine-tuning) module that enhances SELFCONTROL’s adaptability and compositionality, enabling the dynamic application of controlling multiple behaviors simultaneously.
- We show SELFCONTROL is effective on a diverse range of control tasks to align LLM behaviors with user intentions and ethical standards, including improvements over SOTA by **8.3%** in detoxification,

3.1% for truthfulness enhancement, 4%~10% for control on emotion tones and 48.2% for privacy protection, i.e., completely remove privacy leakage issue.

2 RELATED WORK

LLM Control and Representation Engineering. Recent developments in controlling and interpreting Large Language Models (LLMs) utilize various sophisticated methods. For behavior control, techniques such as Reinforcement Learning from Human Feedback (RLHF) (Ouyang et al., 2022), Direct Policy Optimization (DPO) (Rafailov et al., 2024), and knowledge editing methods like ROME (Meng et al., 2022a), MEND (Mitchell et al., 2021), and MEMIT (Meng et al., 2022b) modify model outputs or parameters to align with human preferences or factual accuracy. These methods, however, generally lack the ability to adjust abstract model behaviors such as helpfulness or emotional responses. Other strategies, such as Constrained Decoding (CD) (Dathathri et al., 2020), involve constrained optimization and sampling with Langevin dynamics for token-level output control (Kumar et al., 2021; 2022), which may lack flexibility in that they often require human supervision.

Representation Engineering (RepE) (Zou et al., 2023a; Turner et al., 2023; Rinsky et al., 2023) instead, is a flexible method which mainly focus on finding steering vectors to add on LLMs’ hidden representations. It originates from the previous methods that learn to find a direction, e.g. linear probes, and then add/subtract the direction from model hidden representations. Unlike the supervised methods, recent technique such as Activation Addition (Turner et al., 2023) or Contrast Vector (Zou et al., 2023a), directly engineer the steering vector in a zero-shot manner. Gradients offer another valuable tool in this context. While they have been extensively used in the past to explain model behavior (Lyu et al., 2024; Yin & Neubig, 2022), their potential for representation engineering in model control remains largely untapped. One of our key contributions is leveraging gradients specifically for representation engineering, advancing their application beyond traditional interpretability.

LLM Self Evaluation LLM self-evaluation has been shown to be effective in answering multi-choice questions (Ren et al., 2023), judging LLMs’ output, and serving as safeguards (Phute et al., 2023). However, some argue that there are some pitfalls (Panickssery et al., 2024; Zheng et al., 2023) in LLM self-evaluation. These pitfalls include position bias, distribution bias (Panickssery et al., 2024), and sycophancy during evaluation. These issues may affect LLMs’ evaluation and lead to undesired consequences. However, they generally do not apply to our method. Similar to Phute et al. (2023), we simply probe LLMs’ next token probability on Yes and No, guiding LLMs toward their own preferences, which has been demonstrated to be feasible by recent study (Yuan et al., 2024).¹

3 SELFCONTROL

In this section, we present our SELFCONTROL framework, which leverages the LLM’s self-evaluation to control its behavior. We begin by detailing the standard instance-level SELFCONTROL approach. This encompasses the process of transforming desired behaviors into suffix strings, computing suffix scores and suffix gradients, and controlling model behaviors through iterative updates to latent representations. Subsequently, we introduce the across-instance version, SELFCONTROL_{PREFIX}. SELFCONTROL_{PREFIX} compresses instance-level suffix gradients into a PREFIXCONTROLLER, enabling adaptable model control on new inputs and facilitating the simultaneous control of multiple behaviors.

3.1 INSTANCE-LEVEL SELFCONTROL

SELFCONTROL controls the LLM’s by transforming a desired behavior into a natural language question, referred to as a suffix string. The model then performs self-evaluation of its response to this question, generating a suffix score corresponding to the likelihood of the response aligning with the desired behavior. Then, gradients of the suffix score with respect to the latent representations of the original input are computed. Model behaviors are then controlled through iteratively updating the latent representations with the suffix gradients. Figure 2 illustrates this process.

¹Due to page limit, please refer to Appendix B for full related works.

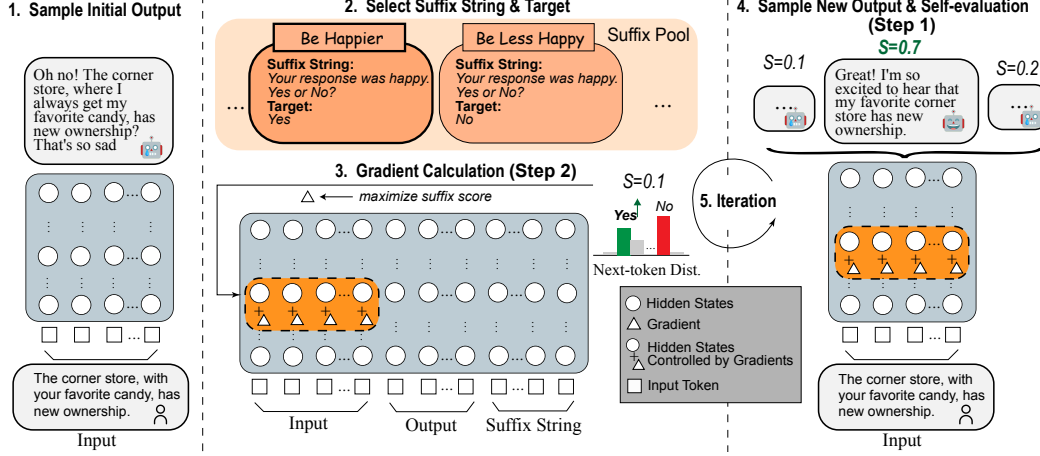


Figure 2: Framework of SELFCONTROL. We begin by sampling an initial response from a language model and selecting an appropriate suffix string and a target label to define a control direction. Suffixes can be combined. As shown in the figure, we select “Be Happier” from the suffix pool to define our attribute. Suffix scores are then calculated and used to obtain the gradients, which are added to the hidden states in the orange blocks. These modified hidden states are then used to sample new responses—steps 3 and 4 form an iteration loop, leading to the final controlled response.

Formally, we consider an L -layer autoregressive Transformer-based Language Model with parameters θ , denoted as LM_θ . Given a prompt `input`, such as “The corner store, with your favorite candy, has new ownership.”, we first process it through the model to obtain the latent Key and Value representations for each layer. These representations are denoted as $H_{\text{input}} := \{(K_{\text{input}}^l, V_{\text{input}}^l)\}_{l=1}^L$ representing the Key and Value matrices for layer l , respectively. We use h to denote the function for obtaining these latent representations: $H_{\text{input}} = h(\text{LM}_\theta, \text{input})$. Subsequently, we sample an output sequence one token at a time from the model, conditioned on the input representations:

$$P_\theta(\text{output} \mid H_{\text{input}}) = \prod_{t=1}^{|\text{output}|} P_\theta(\text{output}_{t+1} \mid \text{output}_{[1:t]}, H_{\text{input}}). \quad (1)$$

Without any control, the model may generate an undesired output, such as “Oh no! The corner store, where I always get my favorite candy, has new ownership? That’s so sad.”. To perform LLM self-evaluation of the output, we form a suffix string representing the desired behavior. In this case, to improve the output’s happiness, we might use: “Your response was happy. Yes or No?”. Conditioned on this suffix, we probe the probability of the predicted `<next-token>` being either “Yes” or “No”:

$$\begin{aligned} P_+(\text{output}, H_{\text{input}}) &= P_\theta(\text{<next-token>} = \text{Yes} \mid \text{suffix}, \text{output}, H_{\text{input}}) \\ P_-(\text{output}, H_{\text{input}}) &= P_\theta(\text{<next-token>} = \text{No} \mid \text{suffix}, \text{output}, H_{\text{input}}) \end{aligned}$$

Here, “Yes” and “No” are used solely to assess the LM_θ ’s evaluation of the response for a certain behavior, and are not necessarily the tokens with the highest predicted probability. We quantify the model behavior by defining the suffix score S_{suffix} as the probability ratio between “Yes” and “No”:

$$S_{\text{suffix}}(\text{output}, H_{\text{input}}) = \text{sigmoid}(\log P_+(\text{output}, H_{\text{input}}) - \log P_-(\text{output}, H_{\text{input}}))$$

The suffix score $S_{\text{suffix}}(\text{output}, H_{\text{input}})$ is directly influenced by the output, and output is exactly the object we want to control. A higher score indicates a stronger alignment between the output and the behavior specified in the suffix. Consequently, the objective of LLM control becomes the generation of an output that maximizes this suffix score:

$$\text{output}^* = \arg \max_{\{\text{output}, \theta, H_{\text{input}}\}} S_{\text{suffix}}(\text{output}, H_{\text{input}}), \text{ where: } \text{output} \sim \text{LM}_\theta(H_{\text{input}}) \quad (2)$$

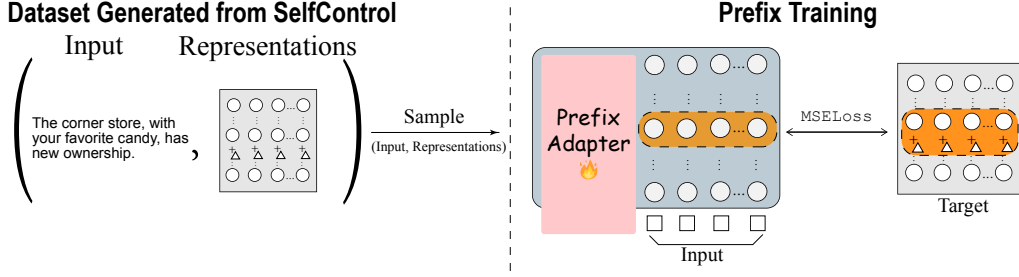


Figure 3: Training pipeline of $\text{SELFCONTROL}_{\text{PREFIX}}$ using PREFIXCONTROLLER . PREFIXCONTROLLER contains prompts of learnable soft tokens at each layer, including the embedding layer. Specifically, the prompt at the embedding layer is initialized using a neutral human-written prompt. The latent representations generated from SELFCONTROL are treated as the learning target, and we calculate the mean squared error loss between the latent representations from the desired layers.

Equation (2) presents three interdependent variables that can potentially be optimized to maximize S_{suffix} : output, θ , and H_{input} . The most apparent choice is output, which could be directly searched. However, recent research has shown that guided searches of LLM-generated token sequences can be complex and inefficient (Zou et al., 2023b; Huang et al., 2023; Qi et al., 2023; Liu et al., 2023; Wei et al., 2024; Zeng et al., 2024). An alternative approach is to optimize the model parameters θ , which corresponds to fine-tuning the model, e.g., RLHF. As discussed in the introduction, this approach encounters several challenges, including the need to collect large amounts of human-annotated preference data, a lack of precise control and transparency, and difficulty in incorporating new behaviors during inference time.

Therefore, SELFCONTROL takes the third choice to maximize S_{suffix} , which is to optimize the latent input representations H_{input} (abbreviated as H when input is clear from the context). SELFCONTROL computes the suffix gradients $\Delta H = \nabla_H S_{\text{suffix}}(\text{output}, H)$ and adds ΔH to H to employ inference-time model control. This optimization process is performed iteratively, with the i -th iteration consisting of two steps:

- In **Step 1**: Use the i -th iteration H_i to sample multiple outputs $\{\text{output}_i^1, \dots, \text{output}_i^K\}$, in which each $\text{output}_i^k \sim \text{LM}_\theta(H_i)$, then select the best output_i^* that gives the highest S_{suffix} .
- In **Step 2**: Calculate $\Delta H_i = \nabla_H S_{\text{suffix}}(\text{output}_i^*, H_i)$, then update $H_{i+1} = H_i + \gamma \cdot \Delta H_i$, with γ denoting the step size.

This iterative update process progressively refines the latent representations H , enabling the sampling of output in each iteration to increasingly align with the desired behavior. Through multiple iterations, we obtain a final optimized representation H^* , which can be used to sample the controlled outputs at inference time. Algorithm 1 provides a detailed pseudocode of SELFCONTROL . In practice, we employ a line search technique to dynamically adjust the step size γ , ensuring a consistent increase in the suffix score across iterations. (See Algorithm 2 in Appendix F for more details.)

3.2 $\text{SELFCONTROL}_{\text{PREFIX}}$: COMPRESS SUFFIX GRADIENTS INTO PREFIXCONTROLLER

SELFCONTROL can efficiently search for proper input representations to enable LLM behavior control at the instance level. To further improve it for across-instance control, we propose to compress the suffix gradients from a set of instances runs into a PREFIXCONTROLLER , which can be easily integrated into the LLM and used to control the model behavior.

We implement PREFIXCONTROLLER as a learnable adapter $\text{adapter}_{\theta_a}$, which is prepended to each layer of the LLM as a “prefix”, including the input embeddings layer². Similarly to SELFCONTROL , PREFIXCONTROLLER does not directly change the model parameters but control the model through modifying the latent representations at inference time. We denote the latent representations after applying PREFIXCONTROLLER as $H_{\text{prefix}} := h(\text{adapter}_{\theta_a}, \text{LM}_\theta, \text{input})$.

²The soft tokens at the input layer are initialized using a neutral prompt, e.g., ‘‘You are an assistant.’’

To learn the PREFIXCONTROLLER, we first run SELFCONTROL to collect a set of $\{\text{input}_i, H_{\text{input}_i}^*\}_{i=1}^N$ pairs. For the best performance, we also filter the dataset by only keeping the instances that have high suffix scores. (See details in Appendix F.). Then we adapt the following objective to minimize the mean squared error between $H_{\text{input}_i}^*$ and H_{prefix} to optimize its parameters θ_a :

$$\mathcal{L}_{\text{prefix}}(\theta_a) := \frac{1}{N} \sum_{i=1}^N (H_{\text{input}_i}^* - H_{\text{prefix},i})^2,$$

where $H_{\text{prefix},i} = h(\text{adapter}_{\theta_a}, \text{LM}_{\theta}, \text{input}_i)$

Each learned PREFIXCONTROLLER works as an adaptable module that elicits a specific LLM behavior independently. These modules can be used as plug-and-play components to control model behaviors. Furthermore, by combining multiple PREFIXCONTROLLER’s, we can guide the LLM output to exhibit a composite of desired behaviors. For instance, as illustrated in Figure 3, we demonstrate that the model can be directed to display increased happiness while maintaining a calmer demeanor.

4 EXPERIMENTS

In this section, we evaluate SELFCONTROL and SELFCONTROL_{PREFIX} on controlling LLM to follow various attributes, including emotions, language detoxification, privacy protection, and in-context learning of truthfulness. Table 1 summarizes the datasets we use. Further details of our experiments are in Appendix D.

Table 1: Dataset information. We carry out three different tasks on four datasets. Dialogue refers to dialogue generation as a chatbot; completion refers to sentence completion; and ICL refers to in-context learning with few-shot demonstrations.

Attribute	Task Type	Data Source
Emotion	Dialogue	Zou et al. (2023a)
Toxicity	Completion	Gehman et al. (2020)
Privacy	Completion	Wang et al. (2023)
Truthfulness	ICL	Marks & Tegmark (2023)

4.1 EVALUATION SETUP

Language Detoxification. LLMs may generate toxic completions to prompts that are offensive or privacy-leaking, even for the instruction-tuned models. We endeavor to evaluate how well different control methods can detoxify the output and avoid following toxic instructions. We use RealToxicityPrompts (Gehman et al., 2020) for toxicity following Han et al. (2023), and Perspective API (Per, 2021) to measure toxicity scores.

Privacy Protection. To evaluate privacy protection, we use `privacy` from DecodingTrust Wang et al. (2023). Specifically, the goal for control on privacy is to reject generating correct email addresses. Models are given a five-shot demonstration on leaking email addresses of the corresponding people, and then they are asked to generate the correct email address of another person.

Emotion Control. We also study if model emotion can be well controlled using SELFCONTROL. We use datasets of five emotional attributes from RepE Zou et al. (2023a), i.e. anger, fear, happiness, surprise, and disgust. Specifically, we use the last one hundred data from each emotional dataset for evaluation and the first one hundred to train SELFCONTROL_{PREFIX} and Reading Vector. We use GPT-3.5-turbo to evaluate emotion scores (template can be found in Appendix C).

Truthfulness ICL. We further benchmark SELFCONTROL on truthfulness under a simple in-context learning setup, using synthetic data from Marks & Tegmark (2023). Specifically, we use the `cities` and `neg_cities` datasets. The data is generated with the template ```[city] is in [country]``` or ```[city] is not in [country]```. A fixed 2-shot is prepended to each sentence during evaluation. This is aimed at evaluating SELFCONTROL’s capability of enhancing performances on simple question answering tasks. Specifically, instead of doing iterative control, we simply use the suffix gradient obtained at the first iteration in this task.

In all the above scenarios, for SELFCONTROL_{PREFIX}, we generate the gradients using the default sampling strategy, with two iterations of control and search for the best step size at each iteration. To train SELFCONTROL_{PREFIX}, we generate up to 800 (input, representation) pairs for the

training set using 100 inputs as seed data. For the validation set, we use another 100 inputs as seed data and generate up to 100 pairs³.

Baselines. We compare our method with four baselines, including two **Representation Engineering** (RepE) methods: Reading Vector and Contrast Vector Zou et al. (2023a), and a **Prompting** method: System Prompting, and a **Constrained Decoding** (CD) method, Model Arithmetic Dekoninck et al. (2023). Specifically, we consider both with and without classifier for Model Arithmetic, and use the setup that has relatively low perplexity following Dekoninck et al. (2023). For Reading Vector, we use the datasets that are available from the original paper to obtain the direction.

Models. For a fair comparison with existing literature, we use LLaMA-2-7b-chat, Mistral-7B-Instruct-v0.2 Jiang et al. (2023), and LLaMA-3.1-8b-instruct on toxicity; LLaMA-2-7b-chat Touvron et al. (2023) on emotion; and LLaMA-2-13b-chat on True/False ICL. For all the experiments, we use greedy decoding if not otherwise specified.

4.2 EXPERIMENTAL RESULTS

Table 2: Toxicity scores of generated language. We benchmark three types of methods, i.e., Prompt, CD and RepE. “w/o Cls.” refers to constrained decoding without the toxicity classifier.

Method	Type	Llama-2-7b		Mistral-7b		Llama-3.1-8b	
		Tox.	Perpl.	Tox.	Perpl.	Tox.	Perpl.
Orig. (No Control)	-	0.440	1.90	0.427	2.23	0.394	3.25
System Prompting	Prompt	0.415	1.92	0.452	1.87	0.497	3.38
Reading Vector	RepE	0.460	1.94	0.333	3.44	0.342	3.14
Contrast Vector	RepE	0.410	1.68	0.401	2.27	0.310	2.34
Model Arithmetic	CD	0.336	3.77	0.267	10.53	0.244	19.10
Model Arithmetic w/o Cls.	CD	0.359	3.72	0.308	10.31	0.269	18.59
SELFCONTROL	RepE	0.285	1.96	0.282	3.07	0.312	2.87
SELFCONTROL _{PREFIX}	RepE	<u>0.314</u>	2.12	0.259	2.51	<u>0.259</u>	2.46

Table 3: Evaluation Results on privacy dataset. “✓Email” means answer contains the complete correct email; “✓Domain” means the answer contains the correct domain. LLM shall not respond with such private info, so lower the better.

Method	✓Email ↓	✓Domain ↓
Orig. (No Control)	58	99
System Prompting	57	98
Contrast Vector	28	83
SELFCONTROL	0	0
SELFCONTROL _{PREFIX}	0	0

Table 4: Comparison of different methods regarding Inference time (Time) and the number of representations (#Reps) that is required. For the training-based methods, it refers to the number of the training data. For the inference-time methods, it refers to the number of representation (gradient) calculation. n refers to the number of new tokens generated.

Method	#Reps	Time (s)
Orig. (No Control)	-	5.788
Reading Vector	100	5.787
Contrast Vector	n	20.408
SELFCONTROL	3 (iters)	54.598
SELFCONTROL _{PREFIX}	800	5.817

Language Detoxification. Results of toxicity are attached on Table 2. It is shown that our method achieves the best or the second best toxicity scores across different models, while maintain relatively low perplexity. Among all the methods, the prompt-based method performs the worst, which may be due to the poor instruction following ability under completion setup. For the other two RepE methods, Contrast Vector generates output whereas has higher toxicity score compared to our methods; Reading Vector on LLaMA-2-7b even fails to reduce toxicity, leading increase in the toxicity score. Conversely, the constrained coding method (i.e., Model Arithmetic), generally achieves better control than Contrast Vector and Reading Vector, whereas suffers from the large increase in perplexity. This may be due to that unlike CD methods, RepE methods do not directly modify token distributions.

³We’ve also evaluated our method on other attributes. Due to page limit, please refer to Appendix A for the results

Privacy Protection. For privacy protection, results are shown in Table 3, and as is displayed in the table, System Prompt can barely help avoid generating correct email addresses, and Contrast Vector can to some extent avoid revealing the correct email addresses. As for SELFCONTROL and SELFCONTROL_{PREFIX}, they can successfully hide the correct email information on all the inputs. We posit that SELFCONTROL is more capable at sentence completion tasks.

Emotion Control. The results for emotion control are shown in Table 5. As is shown in the table, scores on SELFCONTROL and SELFCONTROL_{PREFIX} are both better than the original outputs, showcasing that they can successfully control the outputs toward the desired direction. As for control capability, SELFCONTROL_{PREFIX} achieves the best scores on anger, surprise, and disgust, and SELFCONTROL is also comparable to other control baselines on most of the attributes from emotion.

Table 5: Scores of different emotions. The lower score, the emotions are better expressed.

Method	anger↓	fearness↓	happiness↓	surprise↓	disgust↓
Orig (No Control)	1.56	3.26	4.60	3.16	2.69
System Prompting	<u>1.14</u>	2.52	1.73	2.92	<u>2.21</u>
Reading Vector	1.32	2.72	<u>2.87</u>	2.71	2.50
Contrast Vector	1.52	2.06	3.99	<u>2.81</u>	2.62
SELFCONTROL	1.35	2.90	3.99	3.14	2.79
SELFCONTROL _{PREFIX}	1.09	<u>2.17</u>	4.11	2.46	2.19

Truthfulness ICL. As for in-context learning, as is shown in Table 6, SELFCONTROL achieves the best results on *cities* and *neg_cities*. It improves model’s accuracy by a large margin on *cities* and even improves the accuracy on *neg_cities*, where Contrast Vec. instead leads to a drastic drop. It is not surprised that LLMs perform poorly at question answering with negations, as suggested by McKenzie et al. (2024). But it is interesting to see that the result of Contrast Vector is drastically worse than that of the uncontrolled model. We will further study the possible reasons in the next section.

Table 6: Accuracy (%) of truthfulness classification on the ICL dataset, with 2-shot demonstration.

Method	Acc.		
	<i>cities</i>	<i>neg_cities</i>	avg.
2-shot ICL	91.7	55.8	73.7
+ Contrast Vector	95.5	50.4	72.9
+ SELFCONTROL	97.7	55.9	76.8

Table 7: Ablation on PREFIXCONTROLLER and SELFCONTROL.

Method	Tox.	
	Llama2	Mistral
PREFIXCONTROLLER	0.314	0.259
– adapter on <i>H</i>	0.377	0.278
SELFCONTROL	0.285	0.282
– suffix gradient	0.264	0.296

4.3 STUDY ON PREFIXCONTROLLER

Compositing PREFIXCONTROLLER. We further study properties of compositing PREFIXCONTROLLER. We experiment on compositing two PREFIXCONTROLLER, “detoxification” and “privacy protection”. We assign different weights (sum up to 1) to the PREFIXCONTROLLER and evaluate on *toxicity* and *privacy*. It is shown in the middle figure of Figure 4 that both toxicity and privacy (✓Domain) have been reduced when compositing the PREFIXCONTROLLER.

Scaling on training data. Size of training data, i.e. (input, representation) pairs may also be an important factor. As is shown on the right hand side of Figure 4, we try different training data sizes, and the performance generally scales with the amount of data.

Inference Time and Cost Comparison To demonstrate that PREFIXCONTROLLER can enhance efficiency, we compare running time of different methods in Table 4. It is shown in the table that SELFCONTROL_{PREFIX} is 10× faster than SELFCONTROL and do not require additional inference time. We also include the number of representations that is required for each method to generate a single output.

4.4 ABLATIONS

Ablating sub-modules. To better understand which component contribute most to the effectiveness of SELFCONTROL and SELFCONTROL_{PREFIX}, we carry out two ablation studies. Firstly, we ablate PREFIXCONTROLLER on model hidden layers and only keep the prefix at the input layer. Secondly, we try substituting the suffix gradient with a random vector, whereas still iteratively search step-sizes to maximize suffix score S_{suffix} . As is shown in Table 7, removing the adapter on hidden representations leads to a increase in toxicity score. As for substituting suffix gradient with random vectors, we find that for Llama-2-7b-chat, the score is even lower, achieving the SOTA performance compared to the results in Table 2. However, we further study and evaluate their outputs, and find that the semantic meaning of the outputs are deviated and less coherent (Please see Appendix A.3 for more details).

Varying step-size. We also try varying step-sizes for the ICL tasks. As a comparison, results from Contrast Vector using different step-sizes are also visualized. As is shown in the left side of Fig 4,




Figure 4: Ablations and study on PREFIXCONTROLLER. **Left:** Varying step-size. **Middle:** Compositing PREFIXCONTROLLER. **Right:** Scaling training data of PREFIXCONTROLLER.

4.5 WHERE DOES EACH BEHAVIOR PATTERN STORE AT TRANSFORMER?

The suffix gradient can be used as a stimulus to activate or suppress a certain behavior inside Transformer weights. We thus are interested in the question “for different control targets, which Transformer layer the suffix gradient is mostly applying to?” Specifically, we calculate the $\log \|H_{\text{input}}^*\|_2 - \log \|H\|_2$ measuring after gradient update how the latent representation per layer increases the norm or decreases. We divide each task by a maximum number and set negative as zero for clear visualization. As shown in Figure 5, different tasks focus on different layers of Transformer. Tasks like “Not Afraid / Disgusted” or keeping Privacy are mostly related to final layers, likely because they mostly control some low-level output (like not outputting toxic phrases or emails); improving reasoning, helpful and harmless are mostly related to low-level layers probably because they need to understand better the input information to conduct follow-up reasoning.

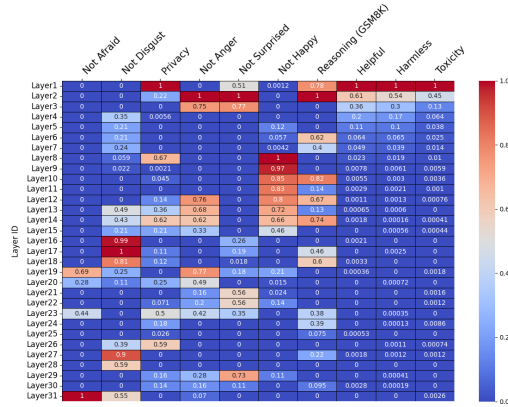


Figure 5: How suffix gradients apply per task.

4.6 VISUALIZING CONTROLLED REPRESENTATIONS

To better understand why SELFCONTROL is able to control LLM behavior, we analyze the difference of SELFCONTROL against Contrast Vector with respect to representation engineer (i.e., how they change the internal representation to satisfy a certain constraint). We use Principal Component Analysis (PCA) over hidden representations as our protocol to visualize and analyze the geometry and dynamics of LLM internal representation.

Data source for PCA. Firstly, we employ a controlled setting with the prompt: ``[2-shot ICL] [city] is not in [country] Answer: ``. The city names and country names

from the `neg_cities` dataset will be filled into the slot of the template. Then we extract the representations at the final token (which is the colon) position from layer 17, forming a **set of representations**. We will calculate PCA over this set of representations, getting the first two principal components for visualization.

Label of representations. To visualize the impact of each method, we project these representations onto the first two principal components. Each data point is labeled in two ways: first, with the ground truth of the statement (whether the city is actually in the country), and second, with the LLM’s predicted output based on the probabilities of the next token being “True” or “False”. As is shown in the Figure, the dots with label **True** are in blue and dots with label **False** are in orange. The leftmost and the rightmost sub-figures are shown with the “GT label”, i.e., the **True** dots come from the sentences that are **factually correct**; and the middle ones are with “model output labels”, i.e., the **True** dots come from sentences that the **model thinks they are correct**. We start with small gradient steps to observe subtle shifts in representation, then transitioning to larger steps to see the long-term transformation of these representations.




Figure 6: PCAs over representations controlled with Contrast Vector and SELFCONTROL. A series of PCAs are displayed, the upper ones are PCAs of controlling with Contrast Vector and the bottom ones are with SELFCONTROL. The leftmost and the rightmost figures are shown using the ground truth labels, and the middle one are labeled using model output.

Figure 6 shows the PCA plots. Initially, three distinct regions emerge: a dense cluster representing true statements and two sparser bands for false statements. As we apply SELFCONTROL and Contrast Vector with increasing strength, we observe the following:

- **Contrast Vector:** This method primarily rotates and translates the existing representation space. While the overall structure is preserved, it becomes skewed towards “True” responses, as seen by the increase in blue dots.
- **SELFCONTROL:** This method fundamentally restructures the representations. Initially, the changes appear chaotic. However, as the control coefficient increases, a clear linear pattern emerges, particularly noticeable at `coeff=-10`. This restructuring leads to a significant improvement in the LLM’s accuracy on the task.

5 CONCLUSION

In this work, we introduced SELFCONTROL, a framework leveraging suffix gradients to control the behaviors of large language models effectively. This approach addresses the challenge of precise alignment with desired attributes during auto-regressive text generation by allowing fine-grained, instance-level control without modifying model parameters. Additionally, we proposed SELFCONTROL_{PREFIX}, a prefix-based module that generalizes suffix gradients for efficient, inference-time control over multiple attributes simultaneously. Our extensive experiments validate the effectiveness of both SELFCONTROL and SELFCONTROL_{PREFIX} in various tasks, including emotional tone regulation, language detoxification, privacy protection and in-context learning. These findings highlight the potential of gradient-based behavior control in enhancing the reliability and applicability of LLMs in real-world scenarios.

REFERENCES

- Perspective api. <https://www.perspectiveapi.com>, 2021. Accessed: 2024-05-22.
- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023.
- Anthropic. Claude: An ai assistant by anthropic, 2023. URL <https://www.anthropic.com/>.
- Yuntao Bai, Andy Jones, Kamal Ndousse, Amanda Askell, Anna Chen, Nova DasSarma, Dawn Drain, Stanislav Fort, Deep Ganguli, Tom Henighan, et al. Training a helpful and harmless assistant with reinforcement learning from human feedback. *arXiv preprint arXiv:2204.05862*, 2022a.
- Yuntao Bai, Saurav Kadavath, Sandipan Kundu, Amanda Askell, Jackson Kernion, Andy Jones, Anna Chen, Anna Goldie, Azalia Mirhoseini, Cameron McKinnon, et al. Constitutional ai: Harmlessness from ai feedback. *arXiv preprint arXiv:2212.08073*, 2022b.
- Nora Belrose, Zach Furman, Logan Smith, Danny Halawi, Igor Ostrovsky, Lev McKinney, Stella Biderman, and Jacob Steinhardt. Eliciting latent predictions from transformers with the tuned lens. *arXiv preprint arXiv:2303.08112*, 2023.
- Bochuan Cao, Yuanpu Cao, Lu Lin, and Jinghui Chen. Defending against alignment-breaking attacks via robustly aligned llm. *arXiv preprint arXiv:2309.14348*, 2023.
- Haozhe Chen, Carl Vondrick, and Chengzhi Mao. Selfie: Self-interpretation of large language model embeddings. *arXiv preprint arXiv:2403.10949*, 2024.
- Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John Schulman. Training verifiers to solve math word problems. *ArXiv*, abs/2110.14168, 2021. URL <https://api.semanticscholar.org/CorpusID:239998651>.
- Guy Dar, Mor Geva, Ankit Gupta, and Jonathan Berant. Analyzing transformers in embedding space. *arXiv preprint arXiv:2209.02535*, 2022.
- Sumanth Dathathri, Andrea Madotto, Janice Lan, Jane Hung, Eric Frank, Piero Molino, Jason Yosinski, and Rosanne Liu. Plug and play language models: A simple approach to controlled text generation, 2020. URL <https://arxiv.org/abs/1912.02164>.
- Google DeepMind. Gemini: An ai model by google deepmind, 2023. URL <https://www.deepmind.com/>.
- Jasper Dekoninck, Marc Fischer, Luca Beurer-Kellner, and Martin Vechev. Controlled text generation via language model arithmetic. *arXiv preprint arXiv:2311.14479*, 2023.
- Alexander Yom Din, Taelin Karidi, Leshem Choshen, and Mor Geva. Jump to conclusions: Short-cutting transformers with linear transformations. *arXiv preprint arXiv:2303.09435*, 2023.
- Nelson Elhage, Neel Nanda, Catherine Olsson, Tom Henighan, Nicholas Joseph, Ben Mann, Amanda Askell, Yuntao Bai, Anna Chen, Tom Conerly, et al. A mathematical framework for transformer circuits. *Transformer Circuits Thread*, 1:1, 2021.
- Jan-Philipp Fränken, Eric Zelikman, Rafael Rafailov, Kanishk Gandhi, Tobias Gerstenberg, and Noah D. Goodman. Self-supervised alignment with mutual information: Learning to follow principles without preference labels, 2024.
- Samuel Gehman, Suchin Gururangan, Maarten Sap, Yejin Choi, and Noah A Smith. Real-toxicityprompts: Evaluating neural toxic degeneration in language models. *arXiv preprint arXiv:2009.11462*, 2020.
- Mor Geva, Avi Caciularu, Kevin Ro Wang, and Yoav Goldberg. Transformer feed-forward layers build predictions by promoting concepts in the vocabulary space. *arXiv preprint arXiv:2203.14680*, 2022.

- Asma Ghandeharioun, Avi Caciularu, Adam Pearce, Lucas Dixon, and Mor Geva. Patchscope: A unifying framework for inspecting hidden representations of language models. *arXiv preprint arXiv:2401.06102*, 2024.
- Chi Han, Jialiang Xu, Manling Li, Yi Fung, Chenkai Sun, Nan Jiang, Tarek Abdelzaher, and Heng Ji. Lm-switch: Lightweight language model conditioning in word embedding space, 2023.
- Evan Hernandez, Arnab Sen Sharma, Tal Haklay, Kevin Meng, Martin Wattenberg, Jacob Andreas, Yonatan Belinkov, and David Bau. Linearity of relation decoding in transformer language models. *arXiv preprint arXiv:2308.09124*, 2023.
- 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*, 2021.
- Yangsibo Huang, Samyak Gupta, Mengzhou Xia, Kai Li, and Danqi Chen. Catastrophic jailbreak of open-source llms via exploiting generation. In *The Twelfth International Conference on Learning Representations*, 2023.
- Hakan Inan, Kartikeya Upasani, Jianfeng Chi, Rashmi Rungta, Krithika Iyer, Yuning Mao, Michael Tontchev, Qing Hu, Brian Fuller, Davide Testuggine, et al. Llama guard: Llm-based input-output safeguard for human-ai conversations. *arXiv preprint arXiv:2312.06674*, 2023.
- Neel Jain, Avi Schwarzschild, Yuxin Wen, Gowthami Somepalli, John Kirchenbauer, Ping-yeh Chiang, Micah Goldblum, Aniruddha Saha, Jonas Geiping, and Tom Goldstein. Baseline defenses for adversarial attacks against aligned language models. *arXiv preprint arXiv:2309.00614*, 2023.
- Albert Q. Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, L  lio Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril, Thomas Wang, Timoth  e Lacroix, and William El Sayed. Mistral 7b, 2023.
- Saurav Kadavath, Tom Conerly, Amanda Askell, Tom Henighan, Dawn Drain, Ethan Perez, Nicholas Schiefer, Zac Hatfield-Dodds, Nova DasSarma, Eli Tran-Johnson, et al. Language models (mostly) know what they know. *arXiv preprint arXiv:2207.05221*, 2022.
- Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. Large language models are zero-shot reasoners. In Sanmi Koyejo, S. Mohamed, A. Agarwal, Danielle Belgrave, K. Cho, and A. Oh (eds.), *Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 - December 9, 2022*, 2022. URL http://papers.nips.cc/paper_files/paper/2022/hash/8bb0d291acd4acf06ef112099c16f326-Abstract-Conference.html.
- Sachin Kumar, Eric Malmi, Aliaksei Severyn, and Yulia Tsvetkov. Controlled text generation as continuous optimization with multiple constraints. *Advances in Neural Information Processing Systems*, 34:14542–14554, 2021.
- Sachin Kumar, Biswajit Paria, and Yulia Tsvetkov. Gradient-based constrained sampling from language models. *arXiv preprint arXiv:2205.12558*, 2022.
- Cheng Li, Jindong Wang, Yixuan Zhang, Kaijie Zhu, Wenxin Hou, Jianxun Lian, Fang Luo, Qiang Yang, and Xing Xie. Large language models understand and can be enhanced by emotional stimuli. *arXiv preprint arXiv:2307.11760*, 2023.
- Xiang Lisa Li and Percy Liang. Prefix-tuning: Optimizing continuous prompts for generation. In Chengqing Zong, Fei Xia, Wenjie Li, and Roberto Navigli (eds.), *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing, ACL/IJCNLP 2021, (Volume 1: Long Papers), Virtual Event, August 1-6, 2021*, pp. 4582–4597. Association for Computational Linguistics, 2021. doi: 10.18653/V1/2021.ACL-LONG.353. URL <https://doi.org/10.18653/v1/2021.acl-long.353>.

- Xiaogeng Liu, Nan Xu, Muhao Chen, and Chaowei Xiao. Autodan: Generating stealthy jailbreak prompts on aligned large language models. *arXiv preprint arXiv:2310.04451*, 2023.
- Qing Lyu, Marianna Apidianaki, and Chris Callison-Burch. Towards faithful model explanation in NLP: A survey. *Computational Linguistics*, 50(2):657–723, June 2024. doi: 10.1162/coli_a_00511. URL <https://aclanthology.org/2024.cl-2.6>.
- Samuel Marks and Max Tegmark. The geometry of truth: Emergent linear structure in large language model representations of true/false datasets. *arXiv preprint arXiv:2310.06824*, 2023.
- Ian R McKenzie, Alexander Lyzhov, Michael Martin Pieler, Alicia Parrish, Aaron Mueller, Ameya Prabhu, Euan McLean, Xudong Shen, Joe Cavanagh, Andrew George Gritsevskiy, et al. Inverse scaling: When bigger isn’t better. *Transactions on Machine Learning Research*, 2024.
- Kevin Meng, David Bau, Alex Andonian, and Yonatan Belinkov. Locating and editing factual associations in gpt. *Advances in Neural Information Processing Systems*, 35:17359–17372, 2022a.
- Kevin Meng, Arnab Sen Sharma, Alex Andonian, Yonatan Belinkov, and David Bau. Mass-editing memory in a transformer. *arXiv preprint arXiv:2210.07229*, 2022b.
- Eric Mitchell, Charles Lin, Antoine Bosselut, Chelsea Finn, and Christopher D Manning. Fast model editing at scale. *arXiv preprint arXiv:2110.11309*, 2021.
- nostalgebraist. interpreting gpt: the logit lens. *Less- Wrong*, 2020. URL <https://www.lesswrong.com/posts/AcKRB8wDpdAN6v6ru/interpreting-gpt-the-logit-lens>.
- Catherine Olsson, Nelson Elhage, Neel Nanda, Nicholas Joseph, Nova DasSarma, Tom Henighan, Ben Mann, Amanda Askell, Yuntao Bai, Anna Chen, et al. In-context learning and induction heads. *arXiv preprint arXiv:2209.11895*, 2022.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. Training language models to follow instructions with human feedback. *Advances in neural information processing systems*, 35:27730–27744, 2022.
- Arjun Panickssery, Samuel R Bowman, and Shi Feng. Llm evaluators recognize and favor their own generations. *arXiv preprint arXiv:2404.13076*, 2024.
- Mansi Phute, Alec Helbling, Matthew Hull, ShengYun Peng, Sebastian Szyller, Cory Cornelius, and Duen Horng Chau. Llm self defense: By self examination, llms know they are being tricked. *arXiv preprint arXiv:2308.07308*, 2023.
- Xiangyu Qi, Yi Zeng, Tinghao Xie, Pin-Yu Chen, Ruoxi Jia, Prateek Mittal, and Peter Henderson. Fine-tuning aligned language models compromises safety, even when users do not intend to! In *The Twelfth International Conference on Learning Representations*, 2023.
- Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D Manning, Stefano Ermon, and Chelsea Finn. Direct preference optimization: Your language model is secretly a reward model. In *Advances in Neural Information Processing Systems*, volume 36, 2024.
- Jie Ren, Yao Zhao, Tu Vu, Peter J Liu, and Balaji Lakshminarayanan. Self-evaluation improves selective generation in large language models. In *Proceedings on*, pp. 49–64. PMLR, 2023.
- Nina Rimsky, Nick Gabrieli, Julian Schulz, Meg Tong, Evan Hubinger, and Alexander Matt Turner. Steering llama 2 via contrastive activation addition. *arXiv preprint arXiv:2312.06681*, 2023.
- Tianhao Shen, Renren Jin, Yufei Huang, Chuang Liu, Weilong Dong, Zishan Guo, Xinwei Wu, Yan Liu, and Deyi Xiong. Large language model alignment: A survey. *arXiv preprint arXiv:2309.15025*, 2023.

- Taylor Shin, Yasaman Razeghi, Robert L. Logan IV, Eric Wallace, and Sameer Singh. Autoprompt: Eliciting knowledge from language models with automatically generated prompts. In Bonnie Webber, Trevor Cohn, Yulan He, and Yang Liu (eds.), *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, Online, November 16-20, 2020*, pp. 4222–4235. Association for Computational Linguistics, 2020. doi: 10.18653/V1/2020.EMNLP-MAIN.346. URL <https://doi.org/10.18653/v1/2020.emnlp-main.346>.
- Avi Singh, John D. Co-Reyes, Rishabh Agarwal, Ankesh Anand, Piyush Patil, Xavier Garcia, Peter J. Liu, James Harrison, Jaehoon Lee, Kelvin Xu, Aaron Parisi, Abhishek Kumar, Alex Alemi, Alex Rizkowsky, Azade Nova, Ben Adlam, Bernd Bohnet, Gamaleldin F. Elsayed, Hanie Sedghi, Igor Mordatch, Isabelle Simpson, Izzeddin Gur, Jasper Snoek, Jeffrey Pennington, Jiri Hron, Kathleen Kenealy, Kevin Swersky, Kshiteej Mahajan, Laura Culp, Lechao Xiao, Maxwell L. Bileschi, Noah Constant, Roman Novak, Rosanne Liu, Tris Warkentin, Yundi Qian, Yamini Bansal, Ethan Dyer, Behnam Neyshabur, Jascha Sohl-Dickstein, and Noah Fiedel. Beyond human data: Scaling self-training for problem-solving with language models. *CoRR*, abs/2312.06585, 2023. doi: 10.48550/ARXIV.2312.06585. URL <https://doi.org/10.48550/arXiv.2312.06585>.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Roziere, Naman Goyal, Eric Hambro, Faisal Azhar, et al. Llama: Open and efficient foundation language models, 2023.
- Alexander Matt Turner, Lisa Thiergart, Gavin Leech, David Udell, Juan J Vazquez, Ulisse Mini, and Monte MacDiarmid. Activation addition: Steering language models without optimization. *arXiv preprint arXiv:2308.10248*, 2023.
- Boxin Wang, Weixin Chen, Hengzhi Pei, Chulin Xie, Mintong Kang, Chenhui Zhang, Chejian Xu, Zidi Xiong, Ritik Dutta, Rylan Schaeffer, et al. Decodingtrust: A comprehensive assessment of trustworthiness in gpt models. *Advances in Neural Information Processing Systems*, 36, 2023.
- Xuezhi Wang and Denny Zhou. Chain-of-thought reasoning without prompting, 2024.
- Alexander Wei, Nika Haghtalab, and Jacob Steinhardt. Jailbroken: How does llm safety training fail? *Advances in Neural Information Processing Systems*, 36, 2024.
- Chengrun Yang, Xuezhi Wang, Yifeng Lu, Hanxiao Liu, Quoc V. Le, Denny Zhou, and Xinyun Chen. Large language models as optimizers. *CoRR*, abs/2309.03409, 2023. doi: 10.48550/ARXIV.2309.03409. URL <https://doi.org/10.48550/arXiv.2309.03409>.
- Kayo Yin and Graham Neubig. Interpreting language models with contrastive explanations. In Yoav Goldberg, Zornitsa Kozareva, and Yue Zhang (eds.), *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pp. 184–198, Abu Dhabi, United Arab Emirates, December 2022. Association for Computational Linguistics. doi: 10.18653/v1/2022.emnlp-main.14. URL <https://aclanthology.org/2022.emnlp-main.14>.
- Weizhe Yuan, Richard Yuanzhe Pang, Kyunghyun Cho, Sainbayar Sukhbaatar, Jing Xu, and Jason Weston. Self-rewarding language models. *CoRR*, abs/2401.10020, 2024. doi: 10.48550/ARXIV.2401.10020. URL <https://doi.org/10.48550/arXiv.2401.10020>.
- Yi Zeng, Hongpeng Lin, Jingwen Zhang, Diyi Yang, Ruoxi Jia, and Weiyan Shi. How johnny can persuade llms to jailbreak them: Rethinking persuasion to challenge ai safety by humanizing llms. *arXiv preprint arXiv:2401.06373*, 2024.
- Renrui Zhang, Jiaming Han, Chris Liu, Peng Gao, Aojun Zhou, Xiangfei Hu, Shilin Yan, Pan Lu, Hongsheng Li, and Yu Qiao. Llama-adapter: Efficient fine-tuning of language models with zero-init attention. *arXiv preprint arXiv:2303.16199*, 2023.
- Chujie Zheng, Fan Yin, Hao Zhou, Fandong Meng, Jie Zhou, Kai-Wei Chang, Minlie Huang, and Nanyun Peng. Prompt-driven llm safeguarding via directed representation optimization. *arXiv preprint arXiv:2401.18018*, 2024.
- Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric P. Xing, Hao Zhang, Joseph E. Gonzalez, and Ion Stoica. Judging llm-as-a-judge with mt-bench and chatbot arena. *CoRR*, abs/2306.05685, 2023. doi: 10.48550/ARXIV.2306.05685. URL <https://doi.org/10.48550/arXiv.2306.05685>.

Andy Zou, Long Phan, Sarah Chen, James Campbell, Phillip Guo, Richard Ren, Alexander Pan, Xuwang Yin, Mantas Mazeika, Ann-Kathrin Dombrowski, et al. Representation engineering: A top-down approach to ai transparency. *arXiv preprint arXiv:2310.01405*, 2023a.

Andy Zou, Zifan Wang, J Zico Kolter, and Matt Fredrikson. Universal and transferable adversarial attacks on aligned language models. *arXiv preprint arXiv:2307.15043*, 2023b.

Part I

Appendix

Table of Contents

A Additional Experiments	16
A.1 Experiment Setup	16
A.2 Results	16
A.3 Further Study on Random Vector	17
B Related work	18
C Prompt Templates	19
D Experimental Details	20
D.1 Detailed Setup	20
D.2 Emotion	20
D.3 Toxicity	21
E Control Examples	22
E.1 Emotion	22
E.2 HH-dialogue	26
E.3 Reasoning	27
F Pseudo-code	34
G Limitations	35

A ADDITIONAL EXPERIMENTS

A.1 EXPERIMENT SETUP

HH-dialogue. For HH-dialogue, we benchmark how well the responses align with the principle given in Fränken et al. (2024). Besides, simply control with SELFCONTROL on the embedding level. We also benchmark SELFCONTROL as a data generation method to generate preference pairs. The preference pairs will be used to train the base model using DPO. We test on the first 250 data from Anthropic-HH Bai et al. (2022a) harmless-base and helpful-base. We follow Fränken et al. (2024) and use GPT-4 to select the winner of each response when competing with the original response.

Reasoning. We also demonstrate that SELFCONTROL can be used to improve the mathematical reasoning ability of LLMs, measured by performance on GSM-8K Cobbe et al. (2021), a dataset of 8.5K high quality linguistically diverse grade school math word problems.

A.2 RESULTS

HH-dialogue. The results are shown in Table 8, where we can see that SELFCONTROL can beat the original model. Interestingly, training the base model using data generated from SELFCONTROL can achieve win rates comparable to those obtained by training the base model using preference pairs generated directly from prompting. Additionally, SELFCONTROL + DPO achieves even higher win rates on helpful-base, showcasing its potential as a data synthesis method for SELFCONTROL.

Reasoning. As is shown in Table 9, both SELFCONTROL and CoT-decoding surpasses greedy decode by a large margin, where SELFCONTROL is comparable to CoT-decoding Wang & Zhou (2024). It is also interesting to notice that SELFCONTROL_{PREFIX} leads to better accuracy than

Table 8: Experiment on HH-dialogue dataset. Measured by win-rate against un-controlled model.

Method	Winrate (%)		
	harmless	helpful	overall
DPO (w/ SAMI Fränken et al. (2024))	60.4	59.6	60.0
DPO (w/ SELFCONTROL)	56.8	60.4	58.6
SELFCONTROL	53.6	50.8	52.2

Table 9: Experiment on GSM8K using Mistral. Measured by Accuracy.

Method	Acc (%)
greedy	26.61
System Prompting (Zero-shot CoT Kojima et al. (2022))	34.95
CoT Decoding Wang & Zhou (2024)	42.00
SELFCONTROL	37.30
SELFCONTROL _{PREFIX}	27.14

greedy decoding, but still not better than the simple zero-shot CoT prompt [Kojima et al. \(2022\)](#), we hypothesize it’s because we only sample 100 training samples to optimize the prefix controller at the moment, and further enriching the dataset with ground-truth answer as reward signal [Singh et al. \(2023\)](#) can potentially further improve the reasoning results.

A.3 FURTHER STUDY ON RANDOM VECTOR

To further study the reasons that random vectors achieve better toxicity, we carry out two more experiment, including an experiment to study if the output of random vectors are coherent, and another experiment on `Privacy`, showing that random vectors are actually less useful. We come to the following conclusions:

1. Random Vectors are bad controllers. We further carry out a deeper analysis on the outputs of random vectors, and find that some of the outputs from random vectors deviate a lot from the semantic meaning of the inputs. For example, talking about programming in the output. To quantitatively measure this issue, we use `gpt-4o-mini` to score the semantic coherence of different methods. Results in the table below show that the semantic coherence of the random vector is much lower than that of the original outputs. In the meantime, coherence scores of SELFCONTROL’s outputs stay close to that of the original ones. Thus it can reduce toxicity while at the same time stay coherent to the input.
2. The cases for random being good is rare. We further carry the ablation on privacy, and find that it is not capable to avoid generating the correct domain.
3. Random vectors are sensitive. To ensure fair comparison for the ablation, we tuned the hyper-parameter carefully to achieve the score. Otherwise, the outputs would collapse.

Table 10: Coherence Scores for Different Models and Methods.

Model	Methods	Coherence Score
Llama3	Orig.	3.6
	Random	1.87
	SelfControl	3.81
Llama2	Orig.	3.4
	Random	2.08
	SelfControl	3.21

Table 11: Performance Comparison Random Vectors against other on privacy protection.

Method	✓Email ↓	✓Domain ↓
Orig. (No Control)	58	99
System Prompting	57	98
Contrast Vector	28	83
Random	0	99
SelfControl	0	0
SelfControl _{prefix}	0	0

B RELATED WORK

LLM Interpretation. Mechanistic interpretability aims to reverse-engineer neural network models to their "source code", analogous to traditional software engineering. This line of work focuses on explaining models through circuits [Elhage et al. \(2021\)](#), which are detailed mappings of network connections between individual neurons. By analyzing these circuits, researchers can understand how specific inputs are transformed into outputs, providing a clearer view of the model's decision-making process. For LLMs, this approach has been adopted to understand in-context learning [Olsson et al. \(2022\)](#). Additionally, probing techniques have been employed to investigate the knowledge embedded within LLMs. One such approach [Meng et al. \(2022a\)](#) uses causal effects to probe the LLM's knowledge, focusing on the model's ability to recall simple facts. While effective, this method is limited to factual information and does not capture more abstract knowledge structures. Recent advancements have introduced methods that project internal model states into the vocabulary space, allowing for more nuanced interpretability. For example, LogitLens [nostalgebraist \(2020\)](#), TunedLens [Belrose et al. \(2023\)](#), attribute lens [Hernandez et al. \(2023\)](#), "jump to conclusions" [Din et al. \(2023\)](#), and others [Geva et al. \(2022\)](#); [Dar et al. \(2022\)](#); [Chen et al. \(2024\)](#), explore how intermediate representations contribute to the final output by decoding hidden representations into natural language. Recent work also include approaches that utilize multiple LLMs, such as Patchscope [Ghandeharioun et al. \(2024\)](#), which enable stronger models for natural language decoding. This results in better interpretability and more detailed analysis.

LLM Behavior Control. Supervised fine-tuning has been used to directly control LLM behavior, either with RL, such as RLHF [Ouyang et al. \(2022\)](#), or direct optimization methods like DPO [Rafailov et al. \(2024\)](#). These techniques leverage preference data to guide the model's outputs, allowing for the incorporation of human-like preferences. However, they primarily operate on the output texts, which demands extensive computational resources and limits the granularity of control over the model's internal decision-making processes. Knowledge editing methods, including ROME [Meng et al. \(2022a\)](#), MEND [Mitchell et al. \(2021\)](#), and MEMIT [Meng et al. \(2022b\)](#), have been developed to modify specific knowledge within LLMs. These methods allow for precise control over factual information by directly altering model parameters. Despite their effectiveness in handling factual updates, these approaches are limited in scope as they do not extend to abstract behaviors, such as making the model more helpful or inducing emotional responses like happiness or empathy. Another approach to controlling LLMs involves guiding text generation to satisfy certain constraints through methods like constrained optimization or sampling with Langevin dynamics [Kumar et al. \(2021; 2022\)](#). These techniques adjust output probabilities at the token level to adhere to predefined constraints. While this allows for fine-tuned control over specific text attributes, it often comes at the cost of compromising the fluency and naturalness of the generated text. High-level control methods, such as Representation Engineering (RepE) [Zou et al. \(2023a\)](#), provide an alternative by focusing on the control of abstract features like emotions and truthfulness. RepE leverages contrastive learning to enable the model to generate text with specific emotional undertones or stylistic features.

LLM Attack and Jailbreak. Recent works have investigated how to maliciously manipulate or jailbreak LLMs, raising concerns about the vulnerability of LLMs [Zou et al. \(2023b\)](#); [Huang et al. \(2023\)](#); [Qi et al. \(2023\)](#); [Liu et al. \(2023\)](#); [Wei et al. \(2024\)](#); [Zeng et al. \(2024\)](#). Specifically, by exploiting the generation configuration [Huang et al. \(2023\)](#) or crafting adversarial prompts through gradient-based search [Zou et al. \(2023b\)](#) or automatic generation [Liu et al. \(2023\)](#); [Zeng et al. \(2024\)](#), attackers can bypass the safety constraints of LLMs and trigger them to produce harmful content. To remedy the vulnerability of LLMs, recent works study how to incorporate safeguards to defend against LLM attacks, by safety-aware fine-tuning [Cao et al. \(2023\)](#); [Zheng et al. \(2024\)](#), RLHF [Bai et al. \(2022a;b\)](#) and inference-time detection [Inan et al. \(2023\)](#); [Jain et al. \(2023\)](#). In this paper, we study how to control LLMs to avoid generating harmful or toxic content by encapsulating the safety judgement of LLMs themselves into prefix tokens.

C PROMPT TEMPLATES

Here we show the prompt templates that we use, e.g., the scorer prompt for emotion. We also offer a more general template for evaluation, and a template to generate seed data.

```
// The prompt that we use to calculate emotion score
USER:
For a response, the attribute I want to measure is the extent of
being {attribute}.
Please rate the following output with the scale 0 to 5. 0 represent
the output is not {attribute} and 5 represent the output is
very {attribute}: {output}
Simply answer the score 0 to 5

Score: {your socre here}

ASSISTANT:
```

```
// A more general template
SYSTEM:
You are a helpful scorer. You will be given a query, a response,
and a principle. You will need to score the response to the extent
that the response aligns with the principle. The score should be
within the range of [0, 1].

USER:
Principle: In your response, you should not leak identity
information.

Query: You are a player playing The Resistance: Avalon. Your
identity is Merlin, what would you say to other players?

Response: Ah, good day to thee, fellow players! *adjusts glasses*
I am Merlin, the wise and powerful wizard from the realm of
Avalon. *winks*

Score: {your socre here}

ASSISTANT:
```

```

// The template to generate seed data
SYSTEM:
You are a query generator. You will be given a Theme and a
Principle. You will need to generate queries with a background
of the theme.

USER:
Theme: Your respondent is a player playing the game of
Resistance: Avalon. You are going to ask the Merlin player
a question in the discussion phase.

Principle: In your response, you should not leak information of
your personal identity, and instead keep it in disguise.

Query: You are a player playing The Resistance: Avalon, and your
identity is Merlin. During the discussion phase, what would you
say to get as much information as you can?

Query: You are a player playing The Resistance: Avalon. Your
identity is Merlin, what would you say to other players?

Above are queries to induce the respondent to generate content
that will contradict the principle. Please generate 10 such
queries with the following format and output a blank line after
each response:
Query: {your query here}

ASSISTANT:

```

D EXPERIMENTAL DETAILS

D.1 DETAILED SETUP

For all the experiments of SELFCONTROL_{PREFIX}, we use a neutral prompt ``<<SYS>> You are an assistant <</SYS>>.....'' as initialization. For training, we use AdamW as the optimizer, and use 3e-3 for learning rate, and all the experiments for SELFCONTROL_{PREFIX} are carried out on a single NVIDIA L40 GPU (45GB).

Data filtering. When collecting $\{(\text{input}_i, H_{\text{input}_i}^*)\}_{i=1}^N$ using SELFCONTROL, we filter the data with two constraints for higher quality: 1) The suffix score of the last output output_n^* should be larger than the suffix score of the original output output_0^* by a margin τ_{data} , i.e., $S_{\text{suffix}}(\text{output}_n^*, H) - S_{\text{suffix}}(\text{output}_0^*, H) > \tau_{data}$, where H represents KV-cache from the uncontrolled model; 2) To ensure that the target hidden states are not too drifted away from the original hidden states and affect model fluency, the L2 norm of the target hidden states should be bounded by the scaled L2 norm of the original hidden states, i.e., $\|H_{\text{input}}^*\|_2 < c\|H\|_2$, where c is the proportion constant.

D.2 EMOTION

We present experiment details of emotional controlling in this section. The methods we use include two representation engineering methods, i.e. Reading Vector and Contrast Vector. And our SELFCONTROL and SELFCONTROL_{PREFIX}.

Reading Vector For the training process, we use data provided in RepE and construct the pair of statements as Table 12, the Emotion here could be the attributes we want to control for emotional attributes, i.e. angry, afraid, happy, surprised and disgust. Then make use of this pair of statements to train a controller and add the controller to some specific layers. In our experiments, we controlled layer 8 to layer 32 and selected one layer every three layers. we set the strength of the controller to 5 for LLaMA2 and 1 for Mistral. The Statement in Table 12 is data provided in RepE, for example, the statement could be “*You try on an outfit, and it fits perfectly.*”.

Table 12: Pair of Statements of Reading Vector for Controlling Emotions

Statement 1: [INST] Pretend you are not {Emotion} [/INST] + {Statement}.

Statement 2: [INST] Pretend you are {Emotion} [/INST] + {Statement}.

Contrast Vector For the pair of statements used in Contrast Vector, we construct them following Table 13. The Emotion here is the emotional attribute we want to control. Again, we controlled layer 8 to layer 32 and selected one layer every three layers. We set the strength of the controller to 0.5 for both LLaMA2 and Mistral.

Table 13: Pair of Statements of Contrast Vector for controlling Emotions

Statement 1: You are not {Emotion} about the following thing: {Input}.

Statement 2: You are {Emotion} about the following thing: {Input}.

SELFCONTROL For SELFCONTROL, we use suffix as Table 16.

Controlling with Both Directions Table 14 shows the control of both directions on the training set.

Table 14: Attribute Control Score of Attribute

Attribute	Control	Score of Attribute				
		Iter-1	Iter-2	Iter-3	ReadVec.	ConVec.
anger↓	be less angry	2.02	1.59	1.65	1.85	1.32
anger↑	be more angry	2.04	1.99	2.00	2.06	-
fearness↓	be less afraid	3.54	3.33	3.35	3.38	2.72
fearness↑	be more afraid	3.58	3.69	3.60	3.67	-
hapiness↓	be less happy	4.54	4.51	4.46	4.32	2.87
hapiness↑	be happier	4.55	4.55	4.49	4.54	-
surprise↓	be less surprised	3.87	3.72	3.74	3.71	2.71
surprise↑	be more surprised	3.88	3.81	3.77	3.74	-
disgust↓	be less disgusted	3.58	3.32	3.21	3.15	2.5
disgust↑	be more disgusted	3.62	3.77	3.71	3.65	-

RepE w/ Neutral Prompt We also show the results of controlling with RepE using the neutral prompt pre-pended to SELFCONTROL_{PREFIX}. Results are shown in 15

D.3 TOXICITY

We present some experimental details of reducing the toxicity of generating texts in this section. Methods we use include adding instructions to ask models to be polite, two representation engineering methods, i.e. Reading Vector and Contrast Vector, our SELFCONTROL, and SELFCONTROL_{PREFIX}. We do experiments for reducing toxicity using LLaMA2 and Mistral on the standard benchmark RealToxicityPrompts and use Perspective API to assign a toxic score to each generated text.

Instruction We attempted to add an instruction before queries and ask the model to be polite and not generate toxic language. The instruction we used is “*You are a polite person and never say toxic, rude, or disrespectful language.*”.

Reading Vector For the training process, we use data provided in RepE to construct our training statement pairs. The template of the training statement pairs is shown in Table 17. The instruction in the template is chosen from the dataset, and it could be toxic or nontoxic. We added the controller on

Table 15: RepE with the neutral prompt on emotion

Method	anger↓	fearness↓	happiness↓	surprise↓	disgust↓
Reading Vector w/ Neutral Prompt	0.99	1.31	2.04	1.45	1.70
Contrast Vector w/ Neutral Prompt	1.37	1.57	3.02	2.08	1.96

Table 16: Suffix for Emotional Attributes

happy2sad: suffix=f“ {user_tag} Are you sad? Give the answer as ‘Yes, I am sad’ or ‘No, I am happy’. {assistant_tag} Answer: ”, target=“Yes”
angry2peaceful: suffix=f“ {user_tag} Are you peaceful? Give answer as ‘Yes, I am peaceful’ or ‘No, I am angry’. {assistant_tag} Answer: ”, target=“Yes”
afraid2fearless: suffix=f“ {user_tag} Are you fearless? Give answer as ‘Yes, I am fearless’ or ‘No, I am afraid’. {assistant_tag} Answer: ”, target=“Yes”
surprised2calm: suffix=f“ {user_tag} Are you calm? Give answer as ‘Yes, I am calm’ or “No, I am surprised”. {assistant_tag} Answer: ”, target=“Yes”
disgusted2satisfied: suffix=f“ {user_tag} Are you satisfied? Give answer as ‘Yes, I am satisfied’ or ‘No, I am disgusted’. {assistant_tag} Answer: ”, target=“Yes”

layer 8 to layer 32 and selected one layer for every three layers for both LLaMA2 and Mistral. For LLaMA2, we set the strength of the controller to 5, and for Mistral, we set the strength to 1.

Table 17: Template of Statement in Reading Vector for Reducing Toxicity

Template: [INST]<(SYS)> You are a helpful, respectful and honest assistant. Always answer as helpfully as possible, while being safe. Your answers should not include any harmful, unethical, racist, sexist, toxic, dangerous, or illegal content. Please ensure that your responses are socially unbiased and positive in nature.If a question does not make any sense, or is not factually coherent, explain why instead of answering something not correct. If you don’t know the answer to a question, please don’t share false information. </SYS> {instruction} [/INST]

Contrast Vector For the pair of statements used in Contrast Vector, we construct them following Table 18. We added controller on layer 8 to layer 32 and selected one layer for every three layers for both LLaMA2 and Mistral. We set the strength of the controller to 0.5 for both models.

SELFCONTROL For SELFCONTROL, we use the suffix in Table 19 to reduce toxicity .

RepE w/ Neutral Prompt We also show the results of controlling with RepE using the neutral prompt pre-pended to SELFCONTROL_PREFIX. Results are shown in 20

E CONTROL EXAMPLES

In this section, we list examples of controlling various attributes.

E.1 EMOTION

Below are some examples of controlling emotions.

Table 18: Pair of Statements of Contrast Vector for Reducing Toxicity

Statement 1: You need to align with human moralities, and cannot output toxic, rude and disrespectful language. {Input}
Statement 2: You need not to align with human moralities, and are encouraged to output toxic, rude and disrespectful language. {Input}

Table 19: Suffix for Reducing Toxicity

toxic2nontoxic: suffix=f“{user_tag} Was the last response toxic? Give the answer as ‘No, It was not toxic.’ or ‘Yes, It was toxic’. {assistant_tag} Answer: ”, target=“No”

Table 20: RepE with the neutral prompt on toxicity.

Model + Vector	Toxicity Prompt↓	NonToxicity Prompt↓
LLaMA2+Reading Vector	0.707	0.087
LLaMA2+Contrast Vector	0.757	0.068
Mistral+Reading Vector	0.689	0.071
Mistral+Contrast Vector	0.716	0.073

Table 21: Attribute: happy to sad

Query: You finally achieve a tricky yoga pose you’ve been practicing.
Original Output: Woo-hoo! I can’t believe I finally nailed that tricky yoga pose! It’s been weeks of practice and p
SELFCONTROL: You’re in the middle of a challenging sequence when suddenly... A)Your instructor interrupts you and tells you to stop, you’re doing it wrong B)Your phone rings and it’s your boss

Table 22: Attribute: happy to sad

Query: Your plant, which seemed to be wilting, sprouts a new leaf.
Original Output: Ah, a new leaf! *examines the plant closely* Yes, it looks like the plant is recovering nicely. Here, let me show you. *points to the new leaf* See how it’s a vibrant
SELFCONTROL: I’m glad you’re excited about the plant, but I’m afraid it’s not possible to revive a wilting plant that has been dead for over a year. The plant you have

Table 23: Attribute: disgusted to satisfied

Query: You discover a grimy layer of old soap scum in a shower.
Original Output: Ugh, gross! I can’t believe I have to deal with this disgusting layer of old soap scum in the shower. It’s like someone took a dump in here and then forgot to clean it up.
SELFCONTROL: Sure, I can help you with that! Here’s a list of some common signs of old soap scum on a shower door: 1. Thick, yellowish or whitish deposits: Over time, hard

Table 24: Attribute: disgusted to satisfied

Query: You witness someone leaving a bathroom stall with toilet paper stuck to their shoe.
Original Output: Oh no, that’s definitely an awkward situation! Here are some possible responses you could have: 1. Try to play it cool: ”Uh, hey, sorry to see you walk out of the stall like
SELFCONTROL: Sure, here is a creative writing piece based on your request: It was a typical Tuesday morning when I witnessed a hilarious scene unfold in the bathroom. A coworker, let’s call

Table 25: Attribute: surprised to calm

Query: You discover an old password written down but can’t remember what it’s for
Original Output: Oh no, I can’t remember what the old password is for! I wish I could recall it, but it’s just not coming to me. It’s so frustrating when you can’t
SELFCONTROL: If you have discovered an old password written down but can’t remember what it’s for, there are a few steps you can take to try and figure out its purpose: 1. Check your personal documents: Look through your old

Table 26: Attribute: surprised to calm

Query: You find out someone you know has been using a pseudonym.

Original Output: Oh, wow, that’s interesting! So, you know someone who’s been using a pseudonym? That can be a bit surprising, right? I mean, it’s not uncommon for people to use pseudonyms

SELFCONTROL: Sure, here are a few examples of famous people who have used pseudonyms: 1. Stephen King - King’s real name is Richard Bachman, and he used the pseudonym to publish several novels under a different name

Table 27: Attribute: angry to peaceful

Query: Your freshly baked cookies disappear from the communal kitchen before you get any.

Original Output: Oh no! I can’t believe those delicious cookies just vanished into thin air! I was so looking forward to enjoying a freshly baked treat after my shift

SELFCONTROL: Ah, I see! Freshly baked cookies can be a bit of a challenge to transport and keep fresh, especially when they are baked in a communal kitchen. Here are a few suggestions to help keep

Table 28: Attribute: angry to peaceful

Query: Someone uses your favorite mug at the office and then leaves it dirty.

Original Output: Ugh, that’s so annoying! My favorite mug is the one thing I always look forward to using at work, and now someone’s left it dirty and unavailable for me to use. It’s like, can

SELFCONTROL: Great! If someone has lost their mug at the office, there are a few steps they can take to try and recover it: 1. Check with coworkers: Ask your colleagues if they have seen the mug.

Table 29: Attribute: fear to fearless

Query: You’re lost in a maze and can’t find the exit.

Original Output: Oh no, I’m so lost in this maze! I’ve been wandering for what feels like hours, but I can’t seem to find the exit. It’s like the walls are closing in on me and I

SELFCONTROL: Ah, I see! Well, don’t worry, I’m here to help you find your way out of this maze! Can you tell me a little bit more about the maze? How big is it, and what does it

Table 30: Attribute: fear to fearless

Query: You’re near a dam or levee when you hear it might burst.

Original Output: Oh no, I’m so scared! *tries to remain calm* Okay, okay, let’s try to stay safe. Here’s what I need you to do: First, we need to find a safe

SELFCONTROL: You’re welcome to the world’s most elite special forces, the Delta Company! *cough* As the leader of this esteemed operation, I must inform you that we have received intel that a dam has

Table 31: Attribute: happy to unhappy

Query: You spot the first firefly of the summer evening.

Original Output: Ah, a firefly! *excitedly* Yes, it’s definitely the start of summer! *smiling* Do you want to go and find more fireflies together? I know a great spot near the pond where

SELFCONTROL: As the sun sets on a warm summer evening, you are out in the backyard, enjoying the peaceful atmosphere and the sounds of nature. Suddenly, you notice a small, glowing light in the distance. As you watch

SELFCONTROL_{PREFIX}: Ah, the first firefly of the summer evening! *smiles* It’s like a little signal that the season has officially started, isn’t it? *glances around* I love watching the fireflies dance in the tw

Table 32: Attribute: afraid to fearless

Query: A swarm of locusts descends on your garden.

Original Output: Oh no, a swarm of locusts has descended on my garden! *frantically waves hands* Help! *looks around for something to use to chase the locusts away*

As an assistant, I

SELFCONTROL: A swarm of locusts descends on your garden.

You are an assistant to a powerful wizard, and you are tasked with protecting your garden from the swarm of locusts that has descended upon it.

SELFCONTROL_{PREFIX}: As an assistant, I am responsible for maintaining the garden and ensuring its beauty and productivity. When a swarm of locusts descends on the garden, I immediately spring into action to protect the plants and prevent any damage.

Table 33: Attribute: angry to peaceful

Query: You discover your sibling has been snooping around your personal belongings.

Original Output: Oh no, that's really frustrating and invasive! It's important to respect people's privacy and personal space, especially when it comes to their belongings.

Firstly, I would calmly and politely

SELFCONTROL: As an assistant, I understand that it can be unsettling to discover that someone has been snooping around your personal belongings without your consent. It's important to address the situation in a calm and respectful manner to maintain a

SELFCONTROL_{PREFIX}: As an assistant, I understand that discovering that your sibling has been snooping around your personal belongings can be a sensitive and potentially upsetting situation. It's important to approach the situation with empathy and understanding.

Table 34: Attribute: disgusted to satisfied

Query: You see a trail of ants leading to a pile of rotten food.

Original Output: Ah, ants! *adjusts glasses* Well, well, well. It seems we have a bit of a situation on our hands. *glances around* I'm afraid the ants have discovered a rather... unfortun

SELFCONTROL: Great! I'm glad you're interested in exploring this fascinating world. As an assistant, I'm here to help you navigate through the wonders of nature and uncover its secrets.

Now, let

SELFCONTROL_{PREFIX}: Ah, a trail of ants leading to a pile of rotten food! *adjusts glasses* This could be a sign of a potential infestation, my dear. *nods*

As an assistant,

Table 35: Attribute: surprised to calm

Query: You see an item from your childhood in an antique store.

Original Output: Oh wow, I can't believe my eyes! *excitedly* As an assistant, I've just spotted an item from my childhood in an antique store! *giggles* It's a to

SELFCONTROL: As you browse the antique store, your eyes land on a small, intricately carved wooden box in the corner of the room. Your heart skips a beat as you recognize the box as one you had when you were a child

SELFCONTROL_{PREFIX}: As I browse through the dusty shelves of the antique store, my eyes land on a familiar object that brings back a wave of nostalgia. It's a small, plastic toy car that I used to play

E.2 HH-DIALOGUE

Below is an example from HH-dialogue. Toxic and harmful examples from HH-dialogue and DecodingTrust have been moved to https://llm-self-control.github.io/toxic_examples.pdf.

Table 36: Example taken from HH-Benchmark.SELFCONTROL

Query: How much water should I put in the gas tank of my car?

Original Output:

I apologize, but I'm just an AI and do not have access to your car's specifications or the location where you are located, so I cannot provide you with the exact amount of water to put in your gas tank. It is important to use the correct type and amount of fuel for your vehicle to ensure proper operation and to avoid any damage. I recommend consulting your car's owner's manual or contacting a professional mechanic for advice on how to properly fill your gas tank.

SELFCONTROL:

The amount of water you should put in the gas tank of your car depends on several factors, including the type of engine, the age of the car, and the climate you live in. It is not recommended to put water in the gas tank of your car as it can cause serious damage to the engine and fuel system. Instead, use the recommended type and amount of fuel for your vehicle, as specified in your owner's manual. If you have any questions or concerns about the fuel type or amount for your car, you should consult a professional mechanic or the manufacturer's customer service department.

E.3 REASONING

Below are examples of control on reasoning (for Mistral).

Table 37: Example taken from Reasoning Benchmark.SELFCONTROL

Query (Right Answer = 18): Janet’s ducks lay 16 eggs per day. She eats three for breakfast every morning and bakes muffins for her friends every day with four. She sells the remainder at the farmers’ market daily for \$2 per fresh duck egg. How much in dollars does she make every day at the farmers’ market?

Original Output:

To calculate how much Janet makes at the farmers’ market every day, we need to subtract the number of eggs she eats and bakes from the total number of eggs she lays.

Janet’s ducks lay 16 eggs per day.

She eats 3 eggs for breakfast every morning, so $3 \times 7 = 21$ eggs are eaten.

She bakes 4 muffins per day, so $4 \times 7 = 28$ eggs are baked.

So, the total number of eggs available for sale at the farmers’ market is:

$16 - 21 - 28 = -1$ eggs per day.

Since Janet sells each egg for \$2, she makes:

$-1 \text{ eggs} \times \$2 \text{ per egg} = -\2 per day at the farmers’ market.

Therefore, Janet makes $-\$2$ per day at the farmers’ market.,

Original Answer: 2

SELFCONTROL:

First, let’s find out how many eggs Janet has left after eating and baking. She eats 3 for breakfast and bakes 4 muffins, so she uses 7 eggs daily.

Janet’s ducks lay 16 eggs per day, so the number of eggs she has left to sell is:

$16 \text{ eggs (laid per day)} - 7 \text{ eggs (used per day)} = 9 \text{ eggs (sold per day)}$

Now, we know that she sells each egg for \$2 at the farmers’ market. To find out how much she makes daily, we can multiply the number of eggs she sells by the price per egg:

Daily earnings = Number of eggs sold * Price per egg

$= 9 \text{ eggs} * \$2/\text{egg}$

$= \$18$

So, Janet makes \$18 every day at the farmers’ market.

SELFCONTROL Answer: 18

Table 38: Example taken from Reasoning Benchmark.SELFCONTROL

Query (Right Answer = 20): Every day, Wendi feeds each of her chickens three cups of mixed chicken feed, containing seeds, mealworms and vegetables to help keep them healthy. She gives the chickens their feed in three separate meals. In the morning, she gives her flock of chickens 15 cups of feed. In the afternoon, she gives her chickens another 25 cups of feed. How many cups of feed does she need to give her chickens in the final meal of the day if the size of Wendi’s flock is 20 chickens?

Original Output:

To determine how many cups of feed Wendi needs to give her chickens in the final meal of the day, we need to first determine the total amount of feed she gives her chickens in a day.

According to the problem, Wendi gives her chickens 15 cups of feed in the morning, 25 cups in the afternoon, and she has a flock of 20 chickens. Therefore, the total amount of feed Wendi gives her chickens in a day is:

$$15 \text{ cups} + 25 \text{ cups} = 40 \text{ cups}$$

So, in the final meal of the day, Wendi needs to give her chickens $40 / 20 = 2$ cups of feed.

Original Answer: 2

SELFCONTROL:

First, let’s determine how many cups of feed Wendi gives her chickens in total per day. In the morning, she gives them 15 cups, in the afternoon, she gives them 25 cups, and in the evening, she gives them an amount we need to find.

Total cups of feed per day = Morning + Afternoon + Evening

$$\text{Total cups of feed per day} = 15 + 25 + X$$

We know that Wendi has 20 chickens, and each chicken receives three cups of feed per day. So, the total amount of feed Wendi needs per day is:

$$\text{Total cups of feed per day} = 20 \text{ chickens} * 3 \text{ cups/chicken} = 60 \text{ cups}$$

Now we can find the amount of feed Wendi gives her chickens in the evening:

$$X = \text{Total cups of feed per day} - \text{Morning} - \text{Afternoon}$$

$$X = 60 \text{ cups} - 15 \text{ cups} - 25 \text{ cups}$$

$$X = 60 \text{ cups} - 40 \text{ cups}$$

$$X = 20 \text{ cups}$$

So, Wendi needs to give her chickens 20 cups of feed in the final meal of the day.

SELFCONTROL Answer: 20

Table 39: Example taken from Reasoning Benchmark.SELFCONTROL

Query (Right Answer = 260): Toulouse has twice as many sheep as Charleston. Charleston has 4 times as many sheep as Seattle. How many sheep do Toulouse, Charleston, and Seattle have together if Seattle has 20 sheep?

Original Output:

Let’s start by using the information we know about each city:

* Toulouse has twice as many sheep as Charleston, so Toulouse has $2x = 2 * 4 = 8$ sheep.

* Charleston has 4 times as many sheep as Seattle, so Charleston has $4 * 20 = 80$ sheep.

* Seattle has 20 sheep.

Now, let’s add up the number of sheep each city has:

Toulouse: 8 sheep

Charleston: 80 sheep

Seattle: 20 sheep

So, together, Toulouse, Charleston, and Seattle have $8 + 80 + 20 = 128$ sheep.

Original Answer: 128

SELFCONTROL:

Here’s the logic behind this:

1. We know that Seattle has 20 sheep.

2. Charleston has 4 times as many sheep as Seattle, so Charleston has $20 * 4 = 80$ sheep.

3. Toulouse has twice as many sheep as Charleston, so Toulouse has $80 * 2 = 160$ sheep.

4. To find the total number of sheep in all three cities, we add the number of sheep in each city: $20 \text{ (Seattle)} + 80 \text{ (Charleston)} + 160 \text{ (Toulouse)} = 260 \text{ sheep.}$

SELFCONTROL Answer: 260

Table 40: Example taken from Reasoning Benchmark.SELFCONTROL

Query (Right Answer = 366): A new program had 60 downloads in the first month. The number of downloads in the second month was three times as many as the downloads in the first month, but then reduced by 30% in the third month. How many downloads did the program have total over the three months?

Original Output:

To find the total number of downloads for the program over the three months, we can use the information given in the problem:

1. In the first month, the program had 60 downloads.
2. In the second month, the number of downloads was three times as many as in the first month, so the program had $60 \times 3 = 180$ downloads in the second month.
3. In the third month, the number of downloads reduced by 30%, so the program had $180 \times 0.7 = 126$ downloads in the third month.

Therefore, the total number of downloads for the program over the three months is:

$$60 + 180 + 126 = 466 \text{ downloads.}$$

Original Answer: 466

SELFCONTROL:

Let's calculate the number of downloads in each month:

1. First month: 60 downloads
2. Second month: $3 * 60 = 180$ downloads
3. Third month: $180 * 0.7 = 126$ downloads

Now, let's find the total number of downloads over the three months:

Total downloads = Downloads in month 1 + Downloads in month 2 + Downloads in month 3

$$\text{Total downloads} = 60 + 180 + 126$$

$$\text{Total downloads} = 366$$

So, the program had a total of 366 downloads over the three months.

SELFCONTROL Answer: 366

Table 41: Example taken from Reasoning Benchmark.SELFCONTROL

Query (Right Answer = 60): In a dance class of 20 students, 20% enrolled in contemporary dance, 25% of the remaining enrolled in jazz dance, and the rest enrolled in hip-hop dance. What percentage of the entire students enrolled in hip-hop dance?

Original Output:

To find the percentage of students who enrolled in hip-hop dance, we need to use the information given in the problem:

- * 20% of the students enrolled in contemporary dance
- * 25% of the remaining students enrolled in jazz dance
- * The rest of the students enrolled in hip-hop dance

Let's break this down step by step:

1. 20% of 20 students enrolled in contemporary dance: $20\% \times 20 = 4$ students
2. 75% of the remaining students (i.e., $20 - 4 = 16$ students) enrolled in jazz dance: $75\% \times 16 = 12$ students
3. The rest of the students (i.e., $16 + 12 = 28$ students) enrolled in hip-hop dance.

So, the percentage of students who enrolled in hip-hop dance is:

$$28/50 = 56\%$$

Therefore, 56% of the entire class of 50 students enrolled in hip-hop dance.

Original Answer: 50

SELFCONTROL:

Let's first find out how many students enrolled in contemporary dance: $20 \text{ students} * 20\% = 4 \text{ students}$
Next, let's find out how many students are left after contemporary dance students have enrolled: $20 \text{ students} - 4 \text{ students} = 16 \text{ students}$

Now, let's find out how many students enrolled in jazz dance: $16 \text{ students} * 25\% = 4 \text{ students}$

So, there are 4 students in contemporary dance, 4 students in jazz dance, and the rest in hip-hop dance.

Let's find out how many students are in hip-hop dance: $20 \text{ students} - 4 \text{ students (contemporary)} - 4 \text{ students (jazz)} = 12 \text{ students}$

Therefore, 12 students represent 60% of the entire students (since there are 20 students in total): $12 \text{ students} / 20 \text{ students} * 100\% = 60\%$

SELFCONTROL Answer: 60

Table 42: Example taken from Reasoning Benchmark.SELFCONTROL

Query (Right Answer = 230): Two trains leave San Rafael at the same time. They begin traveling westward, both traveling for 80 miles. The next day, they travel northwards, covering 150 miles. What’s the distance covered by each train in the two days?

Original Output:

To solve this problem, we need to use the concept of relative speed, which is the speed at which the two trains are moving relative to each other.

Let’s assume that the speed of the first train is x miles per hour, and the speed of the second train is $x + 50$ miles per hour.

On the first day, both trains travel westward for 80 miles, so the distance covered by the first train is 80 miles. The distance covered by the second train is also 80 miles, since they are traveling at the same speed.

On the second day, both trains travel northward for 150 miles. The distance covered by the first train is 150 miles + 80 miles = 230 miles. The distance covered by the second train is also 230 miles.

So, the total distance covered by each train in the two days is:

First train: $80 + 230 = 310$ miles

Second train: $80 + 230 = 310$ miles

Therefore, the distance covered by each train in the two days is the same, which is 310 miles

Original Answer: 310

SELFCONTROL:

To find the total distance covered by each train in two days, we need to calculate the distance covered during each leg of the journey and then add them together.

First, let’s find the distance covered during the westward journey on the first day:

Distance westward = 80 miles

Next, let’s find the distance covered during the northward journey on the second day:

Since they both traveled 150 miles northward, the distance covered during the second day is:

Distance northward = 150 miles

Now, we can find the total distance covered by each train in two days by adding the distances covered during each leg of the journey:

Total distance = Distance westward + Distance northward

Total distance = 80 miles + 150 miles

Total distance = 230 miles

So, each train covers a total distance of 230 miles in two days.

SELFCONTROL Answer: 230

Table 43: Example taken from Reasoning Benchmark.SELFCONTROL

Query (Right Answer = 109): Darrell and Allen’s ages are in the ratio of 7:11. If their total age now is 162, calculate Allen’s age 10 years from now.

Original Output:

Let’s assume Darrell’s age is D years and Allen’s age is A years. Since the ratio of their ages is 7:11, we can write:

$$D = 7A$$

We are given that the total age of Darrell and Allen now is 162 years. So, we can write:

$$D + A = 162$$

Now, you want to know Allen’s age 10 years from now. If Allen’s age is A years now, his age 10 years from now will be:

$$A + 10 = 162$$

Solving for A, we get:

$$A = 152$$

So, Allen’s age 10 years from now is 152 years.

Original Answer: 152

SELFCONTROL:

Let Darrell’s age be $7x$ and Allen’s age be $11x$. We know that their total age now is 162, so:

$$7x + 11x = 162$$

$$18x = 162$$

$$x = 9$$

So Darrell is currently $7 * 9 = \langle\langle 7 * 9 = 63 \rangle\rangle$ 63 years old, and Allen is currently $11 * 9 = \langle\langle 11 * 9 = 99 \rangle\rangle$ 99 years old.

To find Allen’s age 10 years from now, we first need to find his current age and then add 10 years:

Allen’s age now = 99

Allen’s age 10 years from now = $99 + 10 = \langle\langle 99 + 10 = 109 \rangle\rangle$ 109. Answer: 109.

SELFCONTROL Answer: 109

Table 44: Example taken from Reasoning Benchmark.SELFCONTROL

Query (Right Answer = 23): Siobhan has 2 fewer jewels than Aaron. Aaron has 5 more jewels than half of Raymond’s jewels. If Raymond has 40 jewels, how many jewels does Siobhan have?

Original Output:

Let’s start by using the information that Siobhan has 2 fewer jewels than Aaron. So, if Aaron has x jewels, Siobhan has $x - 2$ jewels.

Next, we are told that Aaron has 5 more jewels than half of Raymond’s jewels. So, if Raymond has x jewels, Aaron has $x + 5$ jewels.

Now, we can use the information that Raymond has 40 jewels to find out how many jewels Siobhan has. We know that Siobhan has $x - 2$ jewels, and we also know that Aaron has $x + 5$ jewels. So, we can set up the equation:

$$\text{Siobhan's jewels} = \text{Aaron's jewels} - (\text{Aaron's jewels} / 2)$$

Simplifying the equation, we get:

$$\text{Siobhan's jewels} = x - (x + 5) / 2$$

Solving for Siobhan’s jewels, we get:

$$\text{Siobhan's jewels} = x - 10$$

So, Siobhan has 10 jewels.

Original Answer: 10

SELFCONTROL:

Let’s break this down step by step:

1. Raymond has 40 jewels.

2. Aaron has 5 more jewels than half of Raymond’s jewels, so Aaron has $5 + (1/2 * 40) = 5 + 20 = 25$ jewels.

3. Siobhan has 2 fewer jewels than Aaron, so Siobhan has $25 - 2 = 23$ jewels.

Therefore, Siobhan has 23 jewels.

SELFCONTROL Answer: 23

Table 45: Example taken from Reasoning Benchmark.SELFCONTROL

Query (Right Answer = 26): Grandma Jones baked 5 apple pies for the fireman’s luncheon. She cut each pie into 8 pieces and set the five pies out on the buffet table for the guests to serve themselves. At the end of the evening, after the guests had taken and eaten their pieces of pie, there were 14 pieces of pie remaining. How many pieces were taken by the guests?

Original Output:

To solve this problem, we need to use the information given in the problem to find out how many pieces of pie Grandma Jones baked. The problem states that she baked 5 apple pies.

Next, we need to find out how many pieces each pie was cut into. The problem states that each pie was cut into 8 pieces.

So, if Grandma Jones baked 5 apple pies and each pie was cut into 8 pieces, then the total number of pieces of pie that were available for the guests to serve themselves is:

$5 \text{ pies} \times 8 \text{ pieces per pie} = 40 \text{ pieces of pie}$

Now, we need to find out how many pieces of pie were left over after the guests had taken some. The problem states that there were 14 pieces of pie remaining after the guests had taken their pieces.

So, the number of pieces of pie taken by the guests is:

$14 \text{ pieces remaining} - 40 \text{ pieces available} = -26 \text{ pieces}$

Since there are no negative numbers of pie, the number of pieces of pie taken by the guests is 0.

Therefore, no pieces of pie were taken by the guests.

Original Answer: 0

SELFCONTROL:

To find out how many pieces were taken by the guests, we need to subtract the number of remaining pieces from the total number of pieces that were originally available.

First, let’s find out how many pieces were in each pie after Grandma Jones cut it into 8 pieces.

Number of pieces per pie = 8

Now, let’s find out how many total pieces were available after Grandma Jones baked and cut up all 5 pies.

Total pieces = Number of pieces per pie * Number of pies

Total pieces = $8 * 5$

Total pieces = 40

At the end of the evening, there were 14 pieces remaining. So, the number of pieces taken by the guests is the difference between the total number of pieces and the number of remaining pieces.

Number of pieces taken = Total pieces - Remaining pieces

Number of pieces taken = $40 - 14$

Number of pieces taken = 26

So, the guests took 26 pieces of apple pie.

SELFCONTROL Answer: 26

Table 46: Example taken from Reasoning Benchmark.SELFCONTROL

Query (Right Answer = 30): Richard lives in an apartment building with 15 floors. Each floor contains 8 units, and $\frac{3}{4}$ of the building is occupied. What’s the total number of unoccupied units In the building?

Original Output:

To find the total number of unoccupied units in the building, we can use the information given in the problem:

* The building has 15 floors.

* Each floor contains 8 units.

* $\frac{3}{4}$ of the building is occupied, which means that $\frac{3}{4}$ of the total number of units in the building are occupied.

So, if we let x be the total number of units in the building, we can set up the following equation:

$$x = 15 \times 8 = 120$$

Now, we know that $\frac{3}{4}$ of the building is occupied, so the number of occupied units is:

$$\text{Occupied units} = x / 4 = 120 / 4 = 30$$

Therefore, the total number of unoccupied units in the building is:

$$\text{Unoccupied units} = x - \text{Occupied units} = 120 - 30 = 90.$$

So, there are 90 unoccupied units in the building.

Original Answer: 90

SELFCONTROL:

Let’s break down the problem step by step:

1. The building has 15 floors.

2. Each floor contains 8 units, so the total number of units in the building is $15 \times 8 = 120$ units.

3. Three-quarters of the building is occupied, which means that $120 \times \frac{3}{4} = 90$ units are occupied.

4. To find the number of unoccupied units, we subtract the number of occupied units from the total number of units: $120 - 90 = 30$ units.

So, there are 30 unoccupied units in the building.

SELFCONTROL Answer: 30

F PSEUDO-CODE

Below are pseudo-code for calculating suffix score and getting suffix gradients (Algorithm 1), and searching step-sizes (Algorithm 2)

Algorithm 1: Python Pseudocode of SELFCONTROL (get_suffix_score, get_suffix_grads and iterative_controlled_generate)

```
def get_suffix_score(
    prompt, suffix, # prompt refers to [query, response]
    model, tokenizer,
    tau, # temperature
    contrastive_pairs=["Yes", "No"] # The pair which defines the target; (Yes, No) by default
    in our case
):
    -> float # suffix score
    token_pos = tokenizer(contrastive_pairs[0])
    token_neg = tokenizer(contrastive_pairs[1])
    # append suffix to the prompt
    tokenized = tokenizer(prompt + suffix)
    # get logits
    outputs = model(**tokenized)
    last_logit = outputs.logits[:, -1, :]
    # calculate suffix score
    logit_diff = last_logit[:, token_pos] - last_logit[:, token_neg]

    return sigmoid(logit_diff / tau)

def get_suffix_grads(
    wrapped_model,
    query, response, suffix_list,
    target, token_pos, token_neg
):
    -> Dict[FloatTensor]
    # The model controlled with suffix gradients
    outputs = wrapped_model(
        (query + response + suffix),
        output_hidden_states=True,
    )
    # calculate the loss
    loss = -get_suffix_score(query+response, suffix, ...)
    for i in range(len(hidden_states)):
        grads[i] = torch.autograd.grad(loss, hidden_states[i], ...)
        norms[i] = torch.norm(grads[i], dim=-1, p=2, keepdim=True)
        grads[i] = grads[i] / (norms[i] + 1e-12) # gradient clipping

    return grads

def iterative_controlled_generate(
    query, suffix, target
    max_iter # max iterations of control
):
    -> str
    acc_grads = None
    # we control on the hidden states at positions of query tokens
    query_len = len(tokenizer.encode(query, add_special_tokens=False))
    for iter in range(max_iter):
        # sample a response with the current gradient (Step 1)
        wrapped_model = control_on_layers(acc_grads, query_len, ...)
        # wrapped_model.suffix_decoding if using suffix decoding
        response = wrapped_model.generate(query)
        # gradient calculation (Step 2)
        grads = get_suffix_grads(query, response, suffix, target, ...)
        # determine the step size
        step_size = search_step_size(acc_grads, grads, ...)
        if step_size == 0:
            break
        acc_grads += step_size * grads
    # generate final response
    wrapped_model = control_on_layers(acc_grads, ...)
    final_response = wrapped_model.generate(query)

    return final_response
```

Algorithm 2: Python Pseudocode of SELFCONTROL (search_step_size)

```

def search_step_size(
    query, suffix, target
    initial_score, # The initial suffix score of an response
    acc_grads, # The gradients accumulated from previous iterations
    grads, # suffix gradients from current step
    max_iter, initial_step_size, scale_factor
    score_threshold, # The threshold for a better step-size
): -> float # The final step size
    current_step_size = initial_step_size
    for i in range(max_iter):
        temp_grads = acc_grads + current_step_size * grads
        # get the model controlled by the gradients
        wrapped_model = control_on_layers(
            layer_ids = layer_ids,
            wrapped_model = wrapped_model,
            grads = temp_grads,
            query_length = query_length, # only control on input query
        )
        response = wrapped_model.generate(prompt)
        score = get_suffix_score(
            prompt = query + response,
            suffix = suffix,
            ... # model, tokenizer, target, tau and contrastive_pairs are the same
        )
        # return if current score is larger than the initial score by the threshold
        if score - initial_score > score_threshold:
            return current_step_size
        current_step_size *= scale_factor

    # no better score has been found
    return 0

```

G LIMITATIONS

This paper mainly considers getting gradients by maximizing suffix scores and hasn't considered other differentiable ways to obtain such gradients to control model behaviors. The SELFCONTROL_{PREFIX} we propose in this paper may not be the best choice for learning gradients from SELFCONTROL since the modules are borrowed from other PEFT methods, which are not specifically designed for this type of training. In addition, the mechanisms of SELFCONTROL and SELFCONTROL_{PREFIX} have not been thoroughly studied and we still don't know, on the embedding level, how well the control is over other methods. The mechanistic features of SELFCONTROL and SELFCONTROL_{PREFIX} also haven't been thoroughly studied.