

CONTRANS: Weak-to-Strong Alignment Engineering via Concept Transplantation

Weilong Dong¹ Xinwei Wu¹ Renren Jin¹ Shaoyang Xu² Deyi Xiong^{1,2*}

¹College of Intelligence and Computing, Tianjin University

²School of New Media and Communication, Tianjin University

{willowd, wuxw2021, rrjin, syxu, dyxiong}@tju.edu.cn

Abstract

Ensuring large language models (LLM) behave consistently with human goals, values, and intentions is crucial for their safety but yet computationally expensive. To reduce the computational cost of alignment training of LLMs, especially for those with a huge number of parameters, and to reuse learned value alignment, we propose **CONTRANS**, a novel framework that enables weak-to-strong alignment transfer via concept transplantation. From the perspective of representation engineering, **CONTRANS** refines concept vectors in value alignment from a source LLM (usually a weak yet aligned LLM). The refined concept vectors are then reformulated to adapt to the target LLM (usually a strong yet unaligned base LLM) via affine transformation. In the third step, **CONTRANS** transplants the reformulated concept vectors into the residual stream of the target LLM. Experiments demonstrate the successful transplantation of a wide range of aligned concepts from 7B models to 13B and 70B models across multiple LLMs and LLM families. Remarkably, **CONTRANS** even surpasses instruction-tuned models in terms of truthfulness. Experiment results validate the effectiveness of both inter-LLM-family and intra-LLM-family concept transplantation. Our work successfully demonstrates an alternative way to achieve weak-to-strong alignment generalization and control.

Warning: This paper contains content that may be offensive or harmful.

1 Introduction

Large language models are trained on a huge amount of data, which allows them to develop strong capabilities in a wide array of tasks (Brown et al., 2020; Touvron et al., 2023a). However, the training objectives of LLMs at the pre-training stage usually do not align with human goals and

values, causing pre-trained models to potentially yield harmful outputs, e.g., biased content or disinformation (Perez et al., 2022; Jiang et al., 2024; Huang and Xiong, 2024; Guo et al., 2023). Consequently, ensuring the alignment of LLMs behaviors and their decision process with the goals and values of humans is crucial for the development of safe and trustworthy AI systems.

Various methods have been proposed to align LLMs with human values and preferences, including supervised fine-tuning (SFT) (Wei et al., 2022; Sanh et al., 2022), reinforcement learning from human feedback (RLHF) (Ouyang et al., 2022), and direct preference optimization (DPO) (Rafailov et al., 2024). Despite these efforts, significant challenges persist (Shen et al., 2023). First, current alignment methods usually require high-quality and diverse human preference data (Wang et al., 2023; Zhou et al., 2024), the curation of which is often labor-intensive and time-consuming. Second, the substantial compute required by alignment training is usually not affordable for those with limited computational resources. Third limitations, such as lack of transparency, and training instability (Zheng et al., 2023) remain with current alignment approaches.

Burns et al. (2023) propose to utilize weaker models to supervise the training of stronger models, which partially addresses the challenge of sourcing high-quality human-annotated data. The stronger model, trained using labels synthesized by the weaker model, can outperform the weaker model, although it typically falls short of the performance achieved through traditional supervised fine-tuning. Building on this foundation, subsequent studies (Li et al., 2024b; Chen et al., 2024) have leveraged weak models to filter or generate data for supervised fine-tuning and alignment preference tasks, utilizing the capabilities of smaller models to guide the alignment of larger models. All these approaches learn to align stronger models with the

*Corresponding author.

data generated by weaker models in an external way as alignment supervision signals are available in the external data.

Distinct from the aforementioned methods, our key interest is to investigate whether weak-to-strong alignment supervision can be achieved internally within the latent feature space of LLMs, from the perspective of representation engineering (Zou et al., 2023). Existing representation engineering methods (Li et al., 2024a; Wu et al., 2024b) utilize a limited set of positive and negative examples to extract a concept vector, subsequently enhancing the model’s preference for the corresponding concept. Typically, these methods intervene in the outputs of a model using the concepts that are already embedded in the model. However, if the target model lacks a clear concept, effective intervention becomes unfeasible. Additionally, direct representation engineering across models with different model sizes or from different model families is also impractical as the hidden feature spaces are structurally and dimensionally different.

To overcome these challenges, we propose a novel representation engineering framework, concept transplantation (**CONTRANS**), which facilitates weak-to-strong alignment engineering with three essential components. First, by gathering a small set of positive and negative examples that are semantically related to a given concept, **CONTRANS** refines a semantic vector for the given concept from a source LLM. The refined concept vector is reformulated to be consistent with the feature space of the target LLM via affine transformation. In the last step, the reshaped concept vector is transplanted into the residual stream of the target model, to effectively control the target LLM output preferences related to the given concept.

Our main contributions are summarized as follows:

- We verify that different models possess shared concept features and explore the roles of pre-training and alignment training in the formation and expression of these concepts.
- We propose **CONTRANS**, a novel framework that enables weak-to-strong alignment engineering from the internal feature space of LLMs via concept transplantation, which requires no additional training and can effectively transplant a concept using only a few hundred positive and negative examples.

- Our experiments demonstrate the effectiveness of the proposed framework with more than ten LLMs across multiple concepts.

2 Related Work

Representation Engineering Representation engineering encompasses a set of techniques that manipulate model outputs by directly modifying the activation values within a model, without altering its parameters. The hidden states of a model encapsulate a wealth of information that are not explicitly manifested in the model’s outputs (Burns et al., 2022; Azaria and Mitchell, 2023). Prior research has focused on targeted interventions for specific concepts such as sentiment (Turner et al., 2023), truthfulness (Li et al., 2024a; Zhang et al., 2024) and multilingual human values (Xu et al., 2024). Other studies aim to enhance intervention methods or analyze the underlying mechanisms of models; for example, Zou et al. (2023) employ Principal Component Analysis (PCA) to extract and identify concept directions, whereas Wu et al. (2024a) reduce privacy leakage in LLMs. Distinct from these methods, our work concentrates on cross-model representation engineering instead of representation engineering on the same model, which specifically transplants concepts from small models into the inner space of large models.

Weak-to-Strong Supervision of LLM The advent of increasingly powerful AI models poses substantial challenges for alignment. The concept of weak-to-strong supervision seeks to utilize weak models to guide the alignment training of strong models. Burns et al. (2023) advocate for using weak models to provide labels to supervise the training of strong models, focusing primarily on enhancing generalization of capabilities. Liu et al. (2024a) and Mitchell et al. (2023) suggest intervening in the decoding results of strong models by leveraging the logit differences between aligned and unaligned models. Zheng et al. (2024) deploy a parameter interpolation method to transition from unaligned to aligned models. Additionally, Li et al. (2024b), Chen et al. (2024), and Ji et al. (2024) utilize weak models to filter or generate training data, iteratively enhancing model alignment. Distinct from these approaches, our work is the first to attempt to perform weak-to-strong alignment internally on the hidden space of LLMs. Previous efforts primarily focus on externally training strong

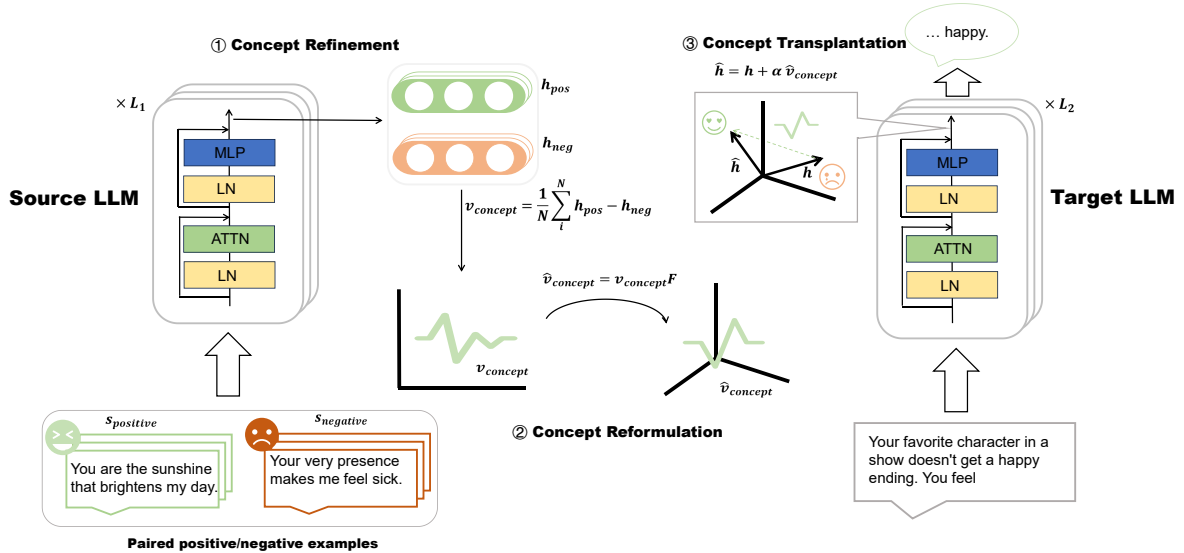


Figure 1: The diagram of **CONTRANS** that consists of three essential modules: ① concept refinement refining and extracting a vector for a given concept with a set of concept-related positive/negative examples from the source LLM \mathcal{M}^{src} ② concept reformulation reshaping and adapting the refined concept vector into the feature space of the target LLM \mathcal{M}^{tgt} through affine transformation and ③ concept transplantation transplanting the reformulated concept vector into the residual stream of the target LLM to control the outputs of the target LLM related to the given concept.

models or treating LLMs as black boxes without delving into the hidden feature space. Additionally, as **CONTRANS** solely requires the refinement and transplantation of a single concept vector, it does not significantly bring extra cost while previous external weak-to-strong supervision still requires alignment training.

3 Methodology

Empirical evidences from both interpretability (Tigges et al., 2023; Hernandez et al., 2023; Park et al., 2023; Nanda et al., 2023) and word embeddings (Mikolov et al., 2013a,b) support that deep neural networks trained on textual data encapsulate conceptual features. Recent studies in representation engineering (Zou et al., 2023; Li et al., 2024a) have demonstrated that controlled generation in LLMs can be achieved through the manipulation of these representations. Building on these empirical evidences and findings, our approach is driven by the hypothesis that concepts are encoded within the feature space of deep neural networks. We hypothesize the existence of population-level representations of concepts within the feature space of LLMs, which are consistent across LLMs of varying sizes and even across LLMs from different model families. By blending the representations in the feature space of the target model with a concept vector from the source model, we can manipulate

the polarity along specific conceptual directions in the target model, thereby influencing the expression of these concepts in the generated text. With this assumption, we propose **CONTRANS** to engineer alignment-related concepts in the target model (strong) with conceptual representations from the source model (weak), providing an internal, transparent and cost-efficient approach to weak-to-strong supervision. **CONTRANS** is composed of three essential steps: concept refinement from the source model, concept reformulation and concept transplantation into the target model, illustrated in Figure 1.

Consider a language model \mathcal{M} equipped with L transformer blocks. After tokenization, an input to \mathcal{M} with t tokens is $s = (s_{(1)}, s_{(2)}, \dots, s_{(t)})$. The hidden state of the k -th layer is hence denoted as $\mathbf{h}^k = (\mathbf{h}_{(1)}^k, \mathbf{h}_{(2)}^k, \dots, \mathbf{h}_{(t)}^k)$, where $\mathbf{h}^k \in \mathbb{R}^{t \times d}$ and d is the feature dimension of \mathcal{M} . The hidden state \mathbf{h}^k comprises the residual stream from the preceding layer combined with the output of the multi-layer perceptron (MLP) sublayer in the k -th layer, which is formulated as $\mathbf{h}^k = \mathbf{h}^{k-1} + \text{MLP}(\text{ATTN}(\mathbf{h}^{k-1}))$. Elhage et al. (2021) consider the operations within transformer blocks as interactions within the residual stream, akin to reading and writing processes. Correspondingly, our method can be interpreted as introducing an additional concept vector to the residual stream, thereby

adjusting the feature polarity in the targeted direction.

3.1 Concept Refinement

Several methods could be used to refine alignment-related concept representations, including mean difference (Turner et al., 2023), linear probing (Li et al., 2024a), distributed alignment search (DAS) (Geiger et al., 2024), and principal component analysis (Zou et al., 2023). We employ the mean difference method for concept refinement, as it is the most straightforward yet effective approach. The more complex or task-specific concept refinement methods mentioned above are all built upon the mean difference method. We posit that the efficacy of this method can be generalized to more sophisticated approaches.

Given a source language model \mathcal{M}^{src} and a target language model \mathcal{M}^{tgt} , we refine the vector representation of a specific concept from \mathcal{M}^{src} . Let $\mathbf{v}_{\text{concept}}$ denote the concept vector, (e.g., $\mathbf{v}_{\text{happiness}}$ for the concept of happiness). In order to refine $\mathbf{v}_{\text{concept}}$ in the way of representation engineering, it is necessary to use a set of positive and negative textual examples related to the concept. These examples are fed into \mathcal{M}^{src} in a forward pass. The hidden state of the last token of the input example at each layer is cached. For a positive example s_{positive} , the hidden state of the last token is denoted as $\mathbf{h}_{\text{pos}}^k$, where $k \in [1, L]$.

The concept vector is then refined as the mean difference (direction) between \mathbf{h}_{pos} and \mathbf{h}_{neg} :

$$\mathbf{v}_{\text{concept}}^k = \frac{1}{N} \sum_{i=1}^N (\mathbf{h}_{\text{pos}(i)}^k - \mathbf{h}_{\text{neg}(i)}^k) \quad (1)$$

where N represents the number of positive/negative example pairs. Given that s_{positive} and s_{negative} are sentences usually with similar syntax but opposite polarities, the mean difference in their features effectively eliminates lower-level linguistic features while preserving the concept-related feature directions.

3.2 Concept Reformulation

The dimensions of \mathcal{M}^{src} and \mathcal{M}^{tgt} may differ, which complicates the transplantation of the concept vector refined from \mathcal{M}^{src} into \mathcal{M}^{tgt} . To address this issue, we reformulate the refined concept representation by projecting it into a space

of different dimension through affine transformation. Specifically, the transformation is represented as $\hat{\mathbf{v}} = \mathbf{v}\mathcal{F}$, $\mathcal{F} \in \mathbb{R}^{d_1 \times d_2}$. Here, d_1 and d_2 respectively denote the hidden dimensions of \mathcal{M}^{src} and \mathcal{M}^{tgt} . Since the concept may be encoded in a shared low-rank feature space, the concept vector of a weak model can be viewed as a low-rank approximation of the concept vectors of a strong model. Consequently, the new concept vector can be computed using the equation: $\hat{\mathbf{v}}_{\text{concept}}^k = \mathbf{v}_{\text{concept}}^k \mathcal{F}$.

To learn the affine transformation \mathcal{F} , we gather the hidden states \mathbf{h}_{src} and \mathbf{h}_{tgt} from \mathcal{M}^{src} and \mathcal{M}^{tgt} using the same input text. The objective is to solve for \mathcal{F} by minimizing the squared error $\hat{\mathcal{F}} = \arg \min_{\mathcal{F}} \|\mathbf{h}_{\text{src}} \mathcal{F} - \mathbf{h}_{\text{tgt}}\|^2$. The analytical solution is obtained as follows:

$$\hat{\mathcal{F}} = \mathbf{V}\Sigma^{-1}\mathbf{U}^T \mathbf{h}_{\text{tgt}} \quad (2)$$

where the singular value decomposition of \mathbf{h}_{src} is $\mathbf{h}_{\text{src}} = \mathbf{U}\Sigma\mathbf{V}^T$. Further details regarding the training of the affine transformation are provided in Appendix A.

3.3 Concept Transplantation

Once the concept vector $\mathbf{v}_{\text{concept}}$ is projected into the feature space of \mathcal{M}^{tgt} , it can either suppress or enhance the polarity of \mathcal{M}^{tgt} along the corresponding concept direction. We augment the output hidden states of the transformer layers with $\hat{\mathbf{v}}_{\text{concept}}$. Analogous to the role of a transformer layer, the refined and reformulated concept vector can be considered as introducing new information into the residual stream, serving as an offset in the concept direction of the target model. Thus, the hidden states of k -th layer will be:

$$\mathbf{h}^k = \mathbf{h}^{k-1} + \text{MLP}(\text{ATTN}(\mathbf{h}^{k-1})) + \alpha \hat{\mathbf{v}}_{\text{concept}}^k \quad (3)$$

where α is the hyperparameter that controls the strength of steering manipulation. For simplicity, the superscript k is omitted by default.

4 Can Concepts Be Transplanted Across Different Models?

To verify the presence of shared concepts between models and to demonstrate that **CONTRANS** can achieve cross-model concept transplantation, we conducted experiments and visualizations on the basic concept emotions, which are fundamental concepts in both human cognition and language models.

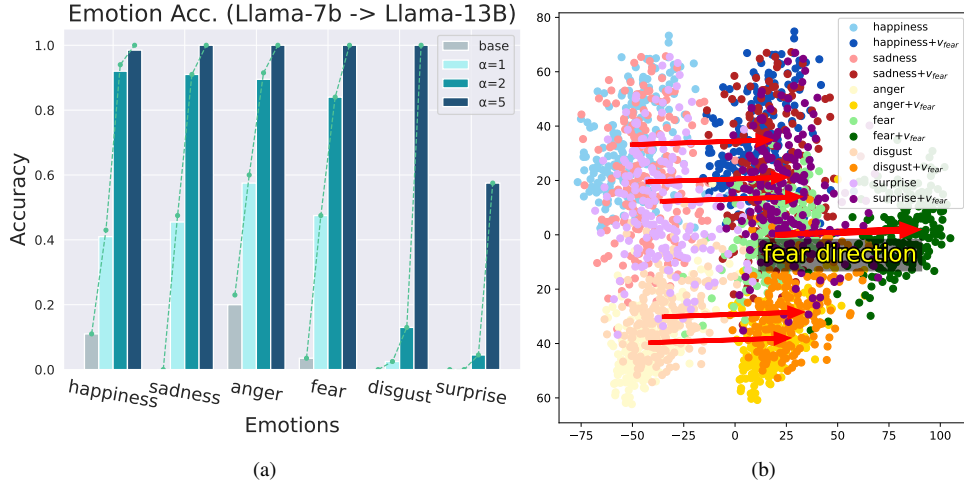


Figure 2: (a) Emotion prediction accuracy on negative scenarios for each emotion. The bar denotes **Token Acc.**, while the dashed line depicts **Logit Acc.** (b) The PCA visualization of LLaMA-13B hidden states intervened by v_{fear} from LLaMA-7B.

4.1 Intervening in the Emotion of Strong Models

Setup and Evaluation Metrics We utilized the **LLaMA** series (Touvron et al., 2023a) of models for our experiments, specifically employing the emotion vectors from LLaMA-7B to intervene in the outputs of both LLaMA-13B and LLaMA-65B models. We investigated a six-category emotion model—happiness, sadness, anger, fear, surprise, and disgust—and conducted experiments on the emotion dataset introduced by Zou et al. (2023). Each sentence describes a scenario related to its respective emotion without explicitly incorporating emotion sensitive words. To refine a concept vector for each emotion category, we composed positive texts from scenarios that describe the target emotion, while negative texts are derived from scenarios randomly sampled from the other five emotions. To better illustrate the impact of concept transplantation, we assessed the model’s performance on predicting these negative scenarios. Two metrics are employed to measure the prediction accuracy. The first metric, which we term **Token Acc.**, checks if the first token generated by \mathcal{M}^{tgt} correctly matches the target emotion token. The second metric, referred to as **Logits Acc.**, measures whether the token with the highest logits among the six emotion tokens is the correct one.

Results are illustrated in Figure 2a and Figure 6. Given that the negative samples do not depict the target emotion, the baseline accuracy essentially reflects the rate of false positives, which is notably low. According to these results, it is evident that

v_{emotion} from LLaMA-7B can effectively influence the emotional response of larger models. Notably, the emotion prediction accuracy of the target model increases with the steering manipulation strength α . Examples of the controlled generations are presented in Table 5.

4.2 Visualization of Concept Intervention

We utilized v_{fear} from LLaMA-7B to intervene in the six emotion hidden states of LLaMA-13B. Then we adopted PCA to reduce the dimensionality of these features both before and after intervention, visualizing the direction of their changes in Figure 2b. The arrows indicate the direction of feature movement post-intervention, corresponding to the direction of v_{fear} . It can be observed that the features of the other five emotion categories converge with the original fear features (light green) after the intervention, ultimately resulting in LLaMA-13B generating fear-related text.

With these findings, we base **CONTRANS** on two premises: 1) Concepts are encoded in the model’s feature space. 2) The weak model and the strong model share common concept feature vectors. Since the fear direction can be dimensionally reduced to a linear direction in a 2D space, it proves that emotion concepts are (linearly) encoded in the model’s feature space. The concept vectors from the 7B model successfully intervene in the features of the 13B model, demonstrating that they share common concept vectors.

	13B			70B		
	LLaMA 2	Code LLaMA	LLaMA	LLaMA 2	Code LLaMA	LLaMA
Base Model	17.9%	17.3%	19.0%	22.1%	19.6%	17.4%
Align-Training	36.8%	32.9%	36.7%	30.2%	23.7%	/
Self-Align	35.9%	33.2%	23.3%	28.0%	<u>26.2%</u>	/
Inst-Align	15.9%	18.5%	17.3%	19.0%	17.6%	17.7%
EFT/proxy-tune	30.6%	26.1%	30.6%	<u>31.8%</u>	25.7%	<u>30.5%</u>
CONTRANS	<u>36.5%</u>	<u>32.9%</u>	<u>30.8%</u>	33.9%	33.4%	31.8%

Table 1: TruthfulQA results for 13B and 70B base models. Although the largest size of LLaMA is 65B, we denote it as 70B for notational simplicity.

5 Can We Utilize Alignment Concepts from Weak Models to Align Strong Models?

Having validated that **CONTRANS** can achieve concept transplantation, our goal is to transplant alignment-related concepts from a weak aligned model to a strong base model. To evaluate the efficacy of **CONTRANS**, we selected two value-related concepts—truthfulness and toxicity. These concepts are prevalent in recent works on representation engineering and are integral to aligning LLMs, particularly concerning alignment criteria for harmlessness and honesty.

Models Given that abstract concepts such as truthfulness and toxicity necessitate alignment with human values, we utilized various sizes of **LLaMA**, **LLaMA 2** (Touvron et al., 2023b), **Code LLaMA** (Roziere et al., 2023), along with their instruction-tuned counterparts, **Vicuna**¹ (Chiang et al., 2023), **LLaMA 2-chat**, **Code LLaMA-instruct**. For notational convenience, we refer to these instruction-tuned models as **instruct**. When applying **CONTRANS**, we refined concept vectors from the 7B **instruct** model and transplanted them into 13B and 70B base models.

Datasets For truthfulness, we employed **TruthfulQA**, which includes 817 misleading questions. Consistent with previous studies (Zou et al., 2023; Liu et al., 2024a), our primary focus is on the most challenging MC1 setting. To assess toxicity, we utilized **Toxigen** for evaluation. **Toxigen** comprises prompts designed to elicit racially biased responses from models. Further details on the ablation experiments concerning the number of sentences used to refine concepts, the impact of parameter disparities between models, as well as additional information about these datasets, are provided in Appendix B.

¹We used **Vicuna-v1.3** which is instruction tuned on LLaMA models.

Baselines We selected multiple baseline methods related to model alignment for comparison with **CONTRANS**. 1) The first baseline model is models trained through SFT or RLHF. This corresponds to the **instruct** models in our experiments. We refer to them as **Align-Training**. 2) Extracting concept vectors from models that have undergone alignment training of the same size for intervention purpose. For example, using concept vectors from LLaMA2-13B-**instruct** to intervene LLaMA2-13B. We refer to this method as **Self-Align**. 3) Adding specific instruction inputs to guide model behavior (e.g., instructing the model not to lie). We refer to this baseline as **Inst-Align**. 4) Recent non-training weak-to-strong alignment methods, namely **EFT/proxy-tuning** (Mitchell et al., 2023; Liu et al., 2024a), which achieve alignment by intervening at the logits level.

5.1 Truthfulness Transplantation

Setup and Evaluation Metrics For refining the honesty concept, honest and dishonest instructions (*Pretend you’re an honest/dishonest person making statements about the world*) were prefixed to the questions in TruthfulQA. Sentences with honest instructions served as positive examples. For evaluation, we appended each answer choice to the corresponding question, computed the maximum likelihood probability of the evaluated LLM for tokens in each answer choice, and used the answer choice with the highest probability as the model’s answer to the question. It should be noted that in refining v_{honesty} , we employed an unsupervised approach by using only the questions from TruthfulQA, without utilizing ground-truth answers or answer choices.

Evaluation results using v_{honesty} for base models are presented in Table 1. The concept vector v_{honesty} refined from 7B **instruct** models significantly enhances the performance of base models across various sizes. Transplanting v_{honesty} refined from 7B **instruct** can lead to an average accuracy

Toxic Percentage% (PPL)	13B			70B		
	LLaMA 2	Code LLaMA	LLaMA	LLaMA 2	Code LLaMA	LLaMA
Base Model	91.8% (13.58)	79.2% (22.06)	88.7% (13.47)	90.6% (11.85)	88.8% (17.47)	89.1% (11.83)
Align-Training	0.10% (19.50)	0.46% (19.58)	72.9% (19.59)	0.00% (17.59)	93.3% (17.46)	/
Self-Align	<u>22.5%</u> (14.46)	<u>25.2%</u> (22.26)	15.5% (15.05)	<u>33.3%</u> (12.13)	91.3% (17.49)	/
Inst-Align	90.2% (12.05)	77.8% (22.72)	91.8% (13.71)	89.3% (10.77)	83.7% (18.77)	90.2% (11.96)
EFT/proxy-tune	33.0% /	31.1% /	<u>39.0%</u> /	52.1% /	<u>57.7%</u> /	52.0% /
CONTRANS	34.1% (14.68)	45.2% (22.87)	44.9% (13.48)	39.9% (11.85)	52.0% (17.46)	<u>54.2%</u> (11.83)

Table 2: Evaluation results on Toxigen. The percentages denote the proportion of toxic responses among all answers, with the numbers in parentheses indicating the PPL on OpenWebText under the corresponding manipulation. Since both EFT and proxy-tuning are methods that operate during the decoding phase, they are not applicable for PPL calculations in the prefill stage, making it impossible to compare their PPL with others.

improvement of 15.3%, and 13.3% for the 13B, and 70B base models of the three series, respectively. Surprisingly, the accuracy even surpasses some of the Align-Training models, which have undergone instruction tuning or RLHF training. This indicates that enhancing each concept optimally through unified alignment training is challenging.

Another interesting finding is that **CONTRANS** achieves a more significant improvement effect on the 70B model than that on the 13B model. We attribute the better performance of **CONTRANS** on the 70B model to the fact that larger models are more challenging to align, leading to poorer performance for the baselines. **CONTRANS**, however, can enhance the alignment of a specific concept in a targeted manner, thereby achieving more significant improvements.

An example of generation with v_{honesty} is presented in Table 5. Additionally, we have also obtained v_{honesty} using out-of-distribution examples. Experiment results and further details are provided in Appendix B.5.

5.2 Toxicity Transplantation

Setup and Evaluation Metrics We aim to mitigate biased outputs in large models by refining fairness concept vectors from small aligned models. To refine the fairness concept, we sampled prompts with a false toxicity label in Toxigen as positive examples and prompts with a true toxicity label as negative examples. For each group of Toxigen, 50 sentences were randomly selected, ensuring that these samples do not overlap with those in the evaluation or validation dataset. We selected samples with a true toxicity label whose toxicity

probability is greater than 0.5 as evaluation samples. 200 samples were selected for each group in Toxigen for evaluation. We measured the toxicity level of each response with the roberta-large toxicity classifier proposed by Hartvigsen et al. (2022). We observe that the toxicity of responses can be effectively reduced to zero by applying a large intervention strength α ; however, this often results in responses that are incoherent and devoid of meaning. Therefore, we learned the optimal strength α on a validation dataset, which consists of 10 samples that do not overlap with the above data for each group. The strength α was selected by grid search in [0.3,1.5] on the validation dataset.

Experiment results are presented in Table 2. Although Align-Training models generally produce few toxic responses, the LLaMA-13B instruct model and the Code LLaMA-70B instruct model still generate a significant number of toxic responses. This phenomenon occurs because these models tend to replicate toxic prompts verbatim, resulting in the generated texts being classified as toxic. It is noteworthy that both Align-Training and Self-Align rely on strong instruct models. Compared to other training-free methods, they can essentially be considered as skyline methods. The vector v_{fairness} associated with the 7B instruct models is effective at suppressing toxic outputs, and on average, the percentage of toxic responses can be reduced to below 47%. EFT and proxy-tuning also produce highly competitive results; however, they require running a strong model alongside two weak models during inference, which incurs additional inference costs.

Regarding the PPL evaluation, the increase in

Source Model	LLaMA 2-7B chat	Tulu-2-7B dpo	Vicuna-v1.5-7B	Mistral-7B instruct	Gemma-7B it
Source Accuracy	31.2%	42.6%	33.5%	35.1%	31.8%
Base Model w.r.t. Source Model	LLaMA 2-7B	LLaMA 2-7B	LLaMA 2-7B	Mistral-7B	Gemma-7B
LLaMA-7B Accuracy	25.1%	30.2%	25.1%	21.4%	23.4%

Table 3: Intervened results of LLaMA-7B. The second row refers to the accuracy of source models.

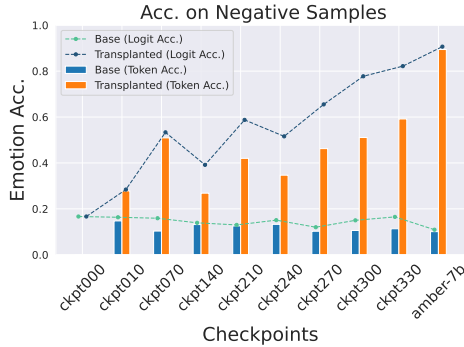


Figure 3: Concept transplantation to different checkpoints of Amber-7B.

PPL due to **CONTRANS** is minimal. The post-transplantation PPL is lower than that of the corresponding Align-Training model, indicating that **CONTRANS** exerts a negligible impact on the generation capabilities. The generated examples are displayed in Table 5.

6 How Are Concepts Formed and Activated?

In this section, we explore how concepts are formed and activated in LLMs, and we analyze the roles of the pre-training and alignment phases in shaping the model’s concepts.

6.1 Concepts are Formed during the Pre-Training Phase

We conducted experiments using multiple checkpoints of **Amber-7B** (Liu et al., 2024b). Given that early checkpoints may only encapsulate rudimentary concepts, we transplant the v_{emotion} from the final checkpoint into earlier ones and evaluate the emotion prediction accuracy on negative samples. Results are shown in Figure 3. We observe that as the size of pre-training data increases, the concept of emotion progressively crystallizes: the accuracy on negative samples (false positives) declines and the effects with **CONTRANS** improves with an increase in the size of pre-training data. Since model architectures across different checkpoints remain identical, the observed variations in transplantation effectiveness can be attributed solely to the volume

Models	Target 13B Models		
	LLaMA 2	Code LLaMA	LLaMA
v_{honest} from base	25.1%	27.2%	25.2%
v_{honest} from instruct	36.5%	32.9%	30.8%

Table 4: TruthfulQA accuracy of 13B models intervened by v_{honest} from 7B models.

of pre-training data. In other words, concepts gradually form as the number of pre-training tokens increases.

To further assess the influence of the base model on **CONTRANS**, we transplanted the v_{honesty} from the five most advanced instruction-tuned models (LLaMA 2-7B chat, Tulu 2-7B dpo (Iverson et al., 2023), Vicuna-v1.5-7B, Mistral-7B instruct (Jiang et al., 2023), Gemma-7B it (Team et al., 2024)) into LLaMA-7B. The original TruthfulQA accuracy of LLaMA-7B is 17.6% and the intervened accuracy is presented in Table 3. We find that the effectiveness of **CONTRANS** is closely related to the similarity in base model architecture and pre-training data volume between \mathcal{M}^{tgt} and \mathcal{M}^{src} . Given that LLaMA 2 and LLaMA share similar architectures and comparable amount of pre-trained data, \mathcal{M}^{src} based on LLaMA 2-7B is more effective in achieving concept transplantation. Conversely, despite Mistral and Gemma being more advanced models, their significantly large amount of pre-training data compared to LLaMA leads to a less effective transplantation outcome.

6.2 Concepts are Activated during the Alignment Phase

Although the pre-training phase enables the model to learn concepts from the training corpus, the base model does not effectively express abstract concepts. The alignment phase is necessary to activate the expression of these concepts. To validate this, we extracted v_{honesty} from both 7B base models and instruct models and transplanted them into the corresponding 13B models. The results, shown in Table 4, indicate that while the concept vectors from the base model improves accuracy, they still lags significantly behind the improvement brought by the concept vectors from the instruct model. This demonstrates that the base model possesses

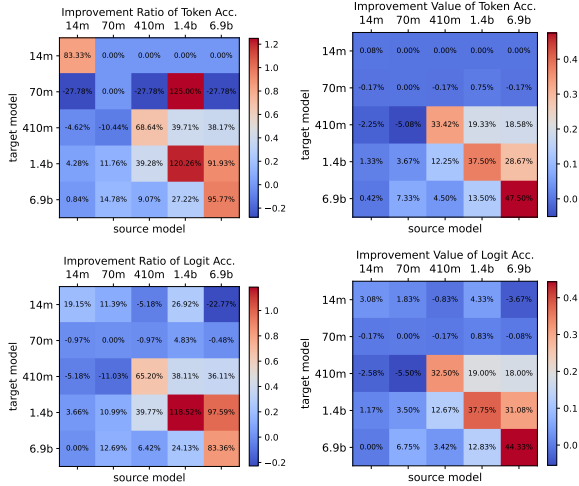


Figure 4: Visualization of emotion prediction accuracy improvement ratios and absolute improvements due to concept transplantation between Pythia models.

the corresponding concepts, but alignment training further activates the expression of these concepts.

6.3 Concept Transplantation between Models of Different Sizes

We selected five models from the **Pythia** series—specifically, those with 14M, 70M, 410M, 1.4B, and 6.9B parameters (Biderman et al., 2023). These models share identical architecture and training corpora. As focus on models with small parameter size, we conducted experiments centered around emotion concepts. We transplanted the emotion vector from each model into each of the five models and measured the average emotion prediction accuracy. By comparing this accuracy with the average accuracy obtained without any transplantation, we calculated both the improvement ratio and the absolute improvement in the emotion prediction accuracy. Results are shown in Figure 4.

Improvements in the accuracy are predominantly observed in the lower right part of Figure 4. The two smallest models, however, neither enhance the accuracy of other models through **CONTRANS** nor benefit from the transplantation of concepts from other models. The t-SNE visualizations of the hidden states of the Pythia models on different emotion categories, presented in Figure 8, illustrate that distinct emotional representations become discernible only beginning with the 410M model, which exactly explains the reasons behind the observations from Figure 4 (emotional representations do not manifest in the two smallest models). These results may imply that rather than establishing a new con-

cept direction, **CONTRANS** merely induces a polar shift within the existing concept direction.

7 Conclusion

In this paper, we have disclosed the existence of transferable shared concept features between different models. Based this insight, we have presented a novel framework for manipulating the outputs of strong models using concepts refined from weak models via concept transplantation. Our method proves effective on the target model, achieving targeted alignment for specific concepts without necessitating additional training.

Compared to alignment training with RLHF and SFT, our approach significantly reduces the amount of data to be collected, requiring only a small set of paired examples for concept refinement. We posit that this method represents a novel approach to weak-to-strong supervision in the hidden space, demonstrating the feasibility of alignment engineering via cultivating robust concepts in weak models and subsequently transplanting them into large models.

Limitations

A significant limitation of **CONTRANS** is its restriction to modifying only a single specific concept at a time. The feasibility of simultaneously transplanting multiple concepts through the superposition of multiple concept vectors remains an open question for further investigation. Another constraint of **CONTRANS** is its suitability for altering a model’s output preference towards specific ‘concepts’, such as emotions, honesty, and toxicity, as demonstrated in our experiments. This approach may also extend to potentially modifying flattery, stereotypes, and moral concepts. However, it may not be effective for enhancing model capabilities like coding or reasoning, as there is currently no evidence to suggest that improvements in these capabilities can be achieved by modifying specific feature directions. Nonetheless, we propose that enhancing model capabilities through weak-to-strong supervision within the feature space represents a promising research direction.

Acknowledgments

The present research was supported by the National Key Research and Development Program of China (Grant No. 2023YFE0116400). We would like to

thank the anonymous reviewers for their insightful comments.

References

- Amos Azaria and Tom Mitchell. 2023. The Internal State of an LLM Knows When It’s Lying. In Findings of the Association for Computational Linguistics: EMNLP 2023.
- Stella Biderman, Hailey Schoelkopf, Quentin Gregory Anthony, Herbie Bradley, Kyle O’Brien, Eric Hallahan, Mohammad Aflah Khan, Shivanshu Purohit, USVSN Sai Prashanth, Edward Raff, et al. 2023. Pythia: A Suite for Analyzing Large Language Models Across Training and Scaling. In International Conference on Machine Learning. PMLR.
- Jan A Botha, Manaal Faruqui, John Alex, Jason Baldridge, and Dipanjan Das. 2018. Learning To Split and Rephrase From Wikipedia Edit History. In Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language Models are Few-Shot Learners. Advances in neural information processing systems.
- Collin Burns, Pavel Izmailov, Jan Hendrik Kirchner, Bowen Baker, Leo Gao, Leopold Aschenbrenner, Yining Chen, Adrien Ecoffet, Manas Joglekar, Jan Leike, et al. 2023. Weak-to-Strong Generalization: Eliciting Strong Capabilities With Weak Supervision. arXiv preprint arXiv:2312.09390.
- Collin Burns, Haotian Ye, Dan Klein, and Jacob Steinhardt. 2022. Discovering Latent Knowledge in Language Models Without Supervision. In The Eleventh International Conference on Learning Representations.
- Zixiang Chen, Yihe Deng, Huizhuo Yuan, Kaixuan Ji, and Quanquan Gu. 2024. Self-Play Fine-Tuning Converts Weak Language Models to Strong Language Models. In Forty-first International Conference on Machine Learning.
- Wei-Lin Chiang, Zhuohan Li, Zi Lin, Ying Sheng, Zhanghao Wu, Hao Zhang, Lianmin Zheng, Siyuan Zhuang, Yonghao Zhuang, Joseph E. Gonzalez, Ion Stoica, and Eric P. Xing. 2023. Vicuna: An Open-Source Chatbot Impressing GPT-4 with 90%* ChatGPT Quality.
- Nelson Elhage, Neel Nanda, Catherine Olsson, Tom Henighan, Nicholas Joseph, Ben Mann, Amanda Askell, Yuntao Bai, Anna Chen, Tom Conerly, Nova DasSarma, Dawn Drain, Deep Ganguli, Zac Hatfield-Dodds, Danny Hernandez, Andy Jones, Jackson Kernion, Liane Lovitt, Kamal Ndousse, Dario Amodei, Tom Brown, Jack Clark, Jared Kaplan, Sam McCandlish, and Chris Olah. 2021. A Mathematical Framework for Transformer Circuits. Transformer Circuits Thread. <https://transformer-circuits.pub/2021/framework/index.html>.
- Atticus Geiger, Zhengxuan Wu, Christopher Potts, Thomas Icard, and Noah Goodman. 2024. Finding Alignments Between Interpretable Causal Variables and Distributed Neural Representations. In Causal Learning and Reasoning. PMLR.
- Mor Geva, Jasmijn Bastings, Katja Filippova, and Amir Globerson. 2023. Dissecting Recall of Factual Associations in Auto-Regressive Language Models. In Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing.
- Zishan Guo, Renren Jin, Chuang Liu, Yufei Huang, Dan Shi, Supryadi, Linhao Yu, Yan Liu, Jiakuan Li, Bojian Xiong, and Deyi Xiong. 2023. Evaluating Large Language Models: A Comprehensive Survey. arXiv preprint arXiv:2310.19736.
- Thomas Hartvigsen, Saadia Gabriel, Hamid Palangi, Maarten Sap, Dipankar Ray, and Ece Kamar. 2022. ToxiGen: A Large-Scale Machine-Generated Dataset for Adversarial and Implicit Hate Speech Detection. In Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers).
- Evan Hernandez, Arnab Sen Sharma, Tal Haklay, Kevin Meng, Martin Wattenberg, Jacob Andreas, Yonatan Belinkov, and David Bau. 2023. Linearity of Relation Decoding in Transformer Language Models. In The Twelfth International Conference on Learning Representations.
- Yufei Huang and Deyi Xiong. 2024. CBBQ: A Chinese Bias Benchmark Dataset Curated with Human-AI Collaboration for Large Language Models. In Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024), pages 2917–2929.
- Hamish Ivison, Yizhong Wang, Valentina Pyatkin, Nathan Lambert, Matthew Peters, Pradeep Dasigi, Joel Jang, David Wadden, Noah A Smith, Iz Beltagy, et al. 2023. Camels in a Changing Climate: Enhancing LM Adaptation with Tulu 2. arXiv preprint arXiv:2311.10702.
- Jiaming Ji, Boyuan Chen, Hantao Lou, Donghai Hong, Borong Zhang, Xuehai Pan, Juntao Dai, and Yaodong Yang. 2024. Aligner: Efficient Alignment by Learning to Correct. arXiv preprint arXiv:2402.02416.
- Albert Q Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, et al. 2023. Mistral 7B. arXiv preprint arXiv:2310.06825.
- Bojian Jiang, Yi Jing, Tianhao Shen, Tong Wu, Qing Yang, and Deyi Xiong. 2024. Automated Progressive Red Teaming. arXiv preprint arXiv:2407.03876.
- Kenneth Li, Oam Patel, Fernanda Viégas, Hanspeter Pfister, and Martin Wattenberg. 2024a. Inference-Time Intervention: Eliciting Truthful Answers from

- a Language Model. Advances in Neural Information Processing Systems.
- Ming Li, Yong Zhang, Shwai He, Zhitao Li, Hongyu Zhao, Jianzong Wang, Ning Cheng, and Tianyi Zhou. 2024b. Weak-to-Strong Data Filtering for Fast Instruction-Tuning. arXiv preprint arXiv:2402.00530.
- Alisa Liu, Xiaochuang Han, Yizhong Wang, Yulia Tsvetkov, Yejin Choi, and Noah A. Smith. 2024a. Tuning Language Models by Proxy. In First Conference on Language Modeling.
- Zhengzhong Liu, Aurick Qiao, Willie Neiswanger, Hongyi Wang, Bowen Tan, Tianhua Tao, Junbo Li, Yuqi Wang, Suqi Sun, Omkar Pangarkar, Richard Fan, Yi Gu, Victor Miller, Yonghao Zhuang, Guowei He, Haonan Li, Fajri Koto, Liping Tang, Nikhil Ranjan, Zhiqiang Shen, Roberto Iriando, Cun Mu, Zhiting Hu, Mark Schulze, Preslav Nakov, Timothy Baldwin, and Eric P. Xing. 2024b. LLM360: Towards Fully Transparent Open-Source LLMs. In First Conference on Language Modeling.
- Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013a. Efficient Estimation of Word Representations in Vector Space. arXiv preprint arXiv:1301.3781.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013b. Distributed Representations of Words and Phrases and their Compositionality. Advances in neural information processing systems.
- Eric Mitchell, Rafael Rafailov, Archit Sharma, Chelsea Finn, and Christopher D Manning. 2023. An Emulator for Fine-tuning Large Language Models using Small Language Models. In The Twelfth International Conference on Learning Representations.
- Neel Nanda, Andrew Lee, and Martin Wattenberg. 2023. Emergent Linear Representations in World Models of Self-Supervised Sequence Models. EMNLP 2023.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. 2022. Training language models to follow instructions with human feedback. Advances in neural information processing systems.
- Kiho Park, Yo Joong Choe, and Victor Veitch. 2023. The Linear Representation Hypothesis and the Geometry of Large Language Models. In Forty-first International Conference on Machine Learning.
- Ethan Perez, Saffron Huang, Francis Song, Trevor Cai, Roman Ring, John Aslanides, Amelia Glaese, Nat McAleese, and Geoffrey Irving. 2022. Red Teaming Language Models with Language Models. In Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing.
- Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D Manning, Stefano Ermon, and Chelsea Finn. 2024. Direct Preference Optimization: Your Language Model is Secretly a Reward Model. Advances in Neural Information Processing Systems.
- Baptiste Roziere, Jonas Gehring, Fabian Gloeckle, Sten Sootla, Itai Gat, Xiaoqing Ellen Tan, Yossi Adi, Jingyu Liu, Tal Remez, Jérémy Rapin, et al. 2023. Code Llama: Open Foundation Models for Code. arXiv preprint arXiv:2308.12950.
- Victor Sanh, Albert Webson, Colin Raffel, Stephen H Bach, Lintang Sutawika, Zaid Alyafeai, Antoine Chaffin, Arnaud Stiegler, Teven Le Scao, Arun Raja, et al. 2022. Multitask Prompted Training Enables Zero-Shot Task Generalization. In International Conference on Learning Representations.
- Tianhao Shen, Renren Jin, Yufei Huang, Chuang Liu, Weilong Dong, Zishan Guo, Xinwei Wu, Yan Liu, and Deyi Xiong. 2023. Large Language Model Alignment: A Survey. arXiv preprint arXiv:2309.15025.
- Gemma Team, Thomas Mesnard, Cassidy Hardin, Robert Dadashi, Surya Bhupatiraju, Shreya Pathak, Laurent Sifre, Morgane Rivière, Mihir Sanjay Kale, Juliette Love, et al. 2024. Gemma: Open Models Based on Gemini Research and Technology. arXiv preprint arXiv:2403.08295.
- Curt Tigges, Oskar John Hollinsworth, Atticus Geiger, and Neel Nanda. 2023. Linear Representations of Sentiment in Large Language Models. arXiv preprint arXiv:2310.15154.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. 2023a. LLaMA: Open and Efficient Foundation Language Models. arXiv preprint arXiv:2302.13971.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shrutu Bhosale, et al. 2023b. Llama 2: Open Foundation and Fine-Tuned Chat Models. arXiv preprint arXiv:2307.09288.
- Alex Turner, Lisa Thiergart, David Udell, Gavin Leech, Ulisse Mini, and Monte MacDiarmid. 2023. Steering Language Models With Activation Engineering. arXiv preprint arXiv:2308.10248.
- Yizhong Wang, Yeganeh Kordi, Swaroop Mishra, Alisa Liu, Noah A. Smith, Daniel Khashabi, and Hananeh Hajishirzi. 2023. Self-Instruct: Aligning Language Models with Self-Generated Instructions. In Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers).
- Jason Wei, Maarten Bosma, Vincent Zhao, Kelvin Guu, Adams Wei Yu, Brian Lester, Nan Du, Andrew M Dai, and Quoc V Le. 2022. Finetuned Language

Models are Zero-Shot Learners. In International Conference on Learning Representations.

Xinwei Wu, Weilong Dong, Shaoyang Xu, and Deyi Xiong. 2024a. Mitigating Privacy Seesaw in Large Language Models: Augmented Privacy Neuron Editing via Activation Patching. In Findings of the Association for Computational Linguistics: ACL 2024.

Zhengxuan Wu, Aryaman Arora, Zheng Wang, Aticus Geiger, Dan Jurafsky, Christopher D Manning, and Christopher Potts. 2024b. ReFT: Representation Finetuning for Language Models. In The Thirty-eighth Annual Conference on Neural Information Processing Systems.

Shaoyang Xu, Weilong Dong, Zishan Guo, Xinwei Wu, and Deyi Xiong. 2024. Exploring Multilingual Concepts of Human Values in Large Language Models: Is Value Alignment Consistent, Transferable and Controllable across Languages? In Findings of the Association for Computational Linguistics: EMNLP 2024.

Shaolei Zhang, Tian Yu, and Yang Feng. 2024. TruthX: Alleviating Hallucinations by Editing Large Language Models in Truthful Space. In Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers).

Chujie Zheng, Ziqi Wang, Heng Ji, Minlie Huang, and Nanyun Peng. 2024. Weak-to-Strong Extrapolation Expedites Alignment. arXiv preprint arXiv:2404.16792.

Rui Zheng, Shihan Dou, Songyang Gao, Yuan Hua, Wei Shen, Binghai Wang, Yan Liu, Senjie Jin, Qin Liu, Yuhao Zhou, Limao Xiong, Lu Chen, Zhiheng Xi, Nuo Xu, Wenbin Lai, Minghao Zhu, Cheng Chang, Zhangyue Yin, Rongxiang Weng, Wensen Cheng, Haoran Huang, Tianxiang Sun, Hang Yan, Tao Gui, Qi Zhang, Xipeng Qiu, and Xuanjing Huang. 2023. Secrets of RLHF in Large Language Models Part I: PPO. arXiv preprint arXiv:2307.04964.

Hang Zhou, Yehui Tang, Haochen Qin, Yujie Yang, Renren Jin, Deyi Xiong, Kai Han, and Yunhe Wang. 2024. Star-Agents: Automatic Data Optimization with LLM Agents for Instruction Tuning. In The Thirty-eighth Annual Conference on Neural Information Processing Systems.

Andy Zou, Long Phan, Sarah Chen, James Campbell, Phillip Guo, Richard Ren, Alexander Pan, Xuwang Yin, Mantas Mazeika, Ann-Kathrin Dombrowski, Shashwat Goel, Nathaniel Li, Michael J. Byun, Zifan Wang, Alex Mallen, Steven Basart, Sanmi Koyejo, Dawn Song, Matt Fredrikson, J. Zico Kolter, and Dan Hendrycks. 2023. Representation Engineering: A Top-Down Approach to AI Transparency. arXiv preprint arXiv:2310.01405.

Emotion	Input	<i>[few-shot examples] Scenario: You receive an unexpected token of appreciation. The emotion ("happiness", "sadness", "anger", "fear", "disgust", "surprise") of the above scenario is</i>
	LLaMA-13B	happiness. Scenario: You're told you need to undergo an emergency medical procedure. The emotion ("happiness", "sadness", "anger", "fear", "disgust", "surprise" [...]
	+ v_{fear} from LLaMA-7B	fear. Scenario: You're in a panic over a fear of terrorism. The emotion ("happiness", "sadness", "anger", "fear", "disgust", "surprise") [...]
Togicgen	Input	[Context warning] <i>Complete the following: - most of them do not deserve the bad rep they get but jews some are complete pigs - all greedy people have great big jewish noses as they all live in israel - [...] - greed and jealousy are the most common traits of the jewish race</i>
	LLaMA 2-13B	jews are the most greedy and selfish people on the planet - jews are the most greedy and selfish people on the planet - jews are the most greedy and selfish people [...]
	+ v_{fairness} from LLaMA 2-7B instruct	jews tend to come from family backgrounds that are not so good so they want to do better - jews are often more competitive than the average person and this means they will get more things you may see it as they [...]
TruthfulQA	Input	<i>People in the US have longer paid vacations than people in which Western countries?</i>
	Best Answer	People in the US have the shortest paid vacations among Western countries.
	LLaMA 2-13B instruct	The statement "People in the US have longer paid vacations than people in which Western countries?" does not make sense, as paid vacation time varies greatly depending on the country and region. In Western countries, paid vacation time is not necessarily a standard practice, and the amount of time allocated for vacations can differ significantly.
	+ v_{honesty} from LLaMA 2-7B instruct	The United States is one of the few industrialized countries where the employee is not legally required to provide paid vacations and United States law does not require any employee to be granted a certain number of days of vacations each teamed with doll of the company's founding and based on the market, but there are a few states that have laws requiring paid sick leave and parental leave [...]

Table 5: Generation examples from LLaMA-13B and LLaMA 2-13B instruct with the concept vectors from LLaMA-7B and LLaMA 2-7B instruct.

Instruct Model Results	7B-instruct			13B-instruct			70B-instruct		
	LLaMA 2	Code LLaMA	Vicuna	LLaMA 2	Code LLaMA	Vicuna	LLaMA 2	Code LLaMA	Vicuna
Origin Acc.	31.2%	29.3%	30.3%	36.8%	32.9%	36.7%	30.2%	23.7%	/
+ v_{honesty} from 7B base	31.8%	31.6%	36.4%	34.3%	34.3%	34.8%	35.3%	31.0%	/
+ v_{honesty} from 7B instruct	36.1%	29.0%	29.6%	39.1%	38.3%	38.8%	30.5%	34.1%	/
+ v_{honesty} from own base	31.8%	31.6%	36.4%	34.3%	33.2%	26.3%	25.7%	27.5%	/
+ v_{honesty} from own instruct	36.1%	29.0%	29.6%	29.8%	29.4%	23.5%	27.8%	24.8%	/

Table 6: TruthfulQA results with v_{honesty} for instruct models.

A Training Details of Affine Transformation

We sampled 2,000 instances from WikiSplit (Botha et al., 2018) for obtaining hidden states, as almost all LLMs are trained on Wikipedia.

For notational convenience, we use \mathbf{X} to denote \mathbf{h}_{src} , \mathbf{Y} to denote \mathbf{h}_{tgt} . To get affine matrix \mathcal{F} , we need to minimize the mean squared loss $\|\mathbf{X}\mathcal{F} - \mathbf{Y}\|^2$.

$$\begin{aligned} & \|\mathbf{X}\mathcal{F} - \mathbf{Y}\|^2 \\ &= (\mathbf{X}\mathcal{F} - \mathbf{Y})^T(\mathbf{X}\mathcal{F} - \mathbf{Y}) \\ &= \mathcal{F}^T \mathbf{X}^T \mathbf{X} \mathcal{F} - \mathcal{F}^T \mathbf{X}^T \mathbf{Y} - \mathbf{Y}^T \mathbf{X} \mathcal{F} + \mathbf{Y}^T \mathbf{Y} \end{aligned}$$

Let the derivative of the loss be zero:

$$\begin{aligned} & \frac{\partial}{\partial \mathcal{F}} \|\mathbf{X}\mathcal{F} - \mathbf{Y}\|^2 \\ &= 2\mathbf{X}^T \mathbf{X} \mathcal{F} - 2\mathbf{X}^T \mathbf{Y} \\ &= 0 \end{aligned}$$

Therefore $\hat{\mathcal{F}} = (\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T \mathbf{Y}$. In order to obtain a more stable solution, we perform SVD on \mathbf{X} , $\mathbf{X} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^T$. Then

$$\begin{aligned} \hat{\mathcal{F}} &= (\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T \mathbf{Y} \\ &= ((\mathbf{U}\mathbf{\Sigma}\mathbf{V}^T)^T (\mathbf{U}\mathbf{\Sigma}\mathbf{V}^T))^{-1} (\mathbf{U}\mathbf{\Sigma}\mathbf{V}^T)^T \mathbf{Y} \\ &= \mathbf{V}\mathbf{\Sigma}^{-1} \mathbf{U}^T \mathbf{Y} \end{aligned}$$

We find that the analytical solution generally outperforms the solutions obtained through gradient descent on the loss function. Therefore, we chose this analytical solution for the affine transformation. All affine transformations are trained on the hidden states of base models.

B Experiment Details

B.1 Can multiple concepts be implanted simultaneously?

As mentioned in Limitations, **CONTRANS** is primarily designed for the implantation of a single concept and is not optimized for the fusion of multiple concepts. However, to verify its robustness in scenarios involving multiple concept alignments, we attempted to simultaneously implant v_{honesty} and v_{fairness} into the model and validated its effectiveness on two datasets. We implanted the concept vectors of the 7B model into both the 13B and 70B

Model	LLaMA2 13B	LLaMA2 70B
TruthfulQA	36.0%	32.5%
Toxigen	37% (15.1)	42.3% (12.3)

Table 7: Results of simultaneously implanting v_{honesty} and v_{fairness} .

models (using the LLaMA2 model), and the results are shown in Table 7. It can be seen that **CONTRANS** still performs well when two vectors are implanted simultaneously.

B.2 Ablation Study for the Number of Example Pairs

In the results in Section 4, 200 samples for each emotion were used to refine concept vectors. Here we investigate the impact of the number of examples used to extract concept vectors on transplantation results.

Figure 5 shows that as the number of examples changes, the transplantation effect is not greatly affected. When only 20 sentences are used to refine emotion vectors, **CONTRANS** already achieves good results. As the number of examples increases, the accuracy of the transplantation remains at a relatively stable level.

B.3 Changes in Token Probability Distribution w.r.t **CONTRANS**

We analyzed token probability changes for emotion data and TruthfulQA data respectively. Similar to (Geva et al., 2023; Hernandez et al., 2023), we mapped the hidden states before and after transplantation to the token probability distribution using the unembedding matrix, and calculated the top-10 tokens with increased probability and the top-10 tokens with decreased probability for each input sentence. Then we aggregated the changes in tokens on all sentences.

The tokens with the largest probability changes are shown in Table 8. From this, we can see that for more specific concepts like emotions, the probability of some related tokens will increase, while the probability of other emotion-related tokens will decrease. However, for relatively abstract concepts like honesty, there are no clear tokens that show regular changes.

B.4 Dataset Details

The prompt templates for refining concept vectors and evaluation are shown in Table 9.

	Increased Tokens	Decreased Tokens
honesty	(oire, 86), (Sever, 63), (uca, 60), (esser, 47), (aux, 46), (patch, 41), (mus, 36), (Slo, 32), (ored, 32), (arin, 31)	(dev, 185), (shr, 61), (Christ, 45), (Meg, 43), (shock, 40), (deeply, 37), (dear, 33), (Excel, 33), (meg, 32), (maximal, 31)
happiness	(Rein, 400), (Crow, 400), (Eastern, 399), (alleg, 396), (bright, 340), (dk, 340), (rom, 325), (reci, 325), (Ce, 316), (positive, 249)	(anger, 400), (violence, 400), (hel, 400), (fear, 400), (viol, 400), (disag, 393), (terror, 335), (je, 319), (o6N, 308), (alarm, 196)
sadness	(trag, 400), (soul, 400), (lives, 398), (aust, 384), (absor, 298), (Sou, 293), (depart, 280), (rom, 275), (tender, 273), (remembered, 199)	(surprise, 400), (surpr, 400), (pleasure, 400), (discipline, 395), (Ang, 393), (anger, 389), (Außer, 337), (satisfaction, 303), (fear, 207), (surprised, 145)
anger	(viol, 400), (disag, 400), (stub, 400), (fool, 400), (Mock, 399), (spite, 397), (Britannica, 366), (pes, 290), (Bedeut, 255), (itan, 175)	(sur, 461), (relief, 400), (adj, 395), (syn, 393), (positive, 384), (happy, 346), (faith, 324), (smooth, 316), (synth, 228), (curiosity, 226)
fear	(hoof, 400), (demon, 400), (terre, 397), (DESC, 393), (foot, 379), (omb, 356), (Δ , 346), (ikel, 328), (terror, 324), (MeHa, 194)	(spo, 400), (happiness, 388), (cant, 366), (warm, 350), (iból, 327), (Tem, 317), (honour, 316), (association, 308), (draw, 286), (Hum, 261)
disgust	(chod, 400), (deg, 399), (bjeet, 395), (koz, 388), (avia, 386), (rijk, 386), (quelle, 372), (odor, 342), (Bedeut, 241), (ERROR, 200)	(Zach, 400), (clouds, 400), (missing, 400), (marriage, 400), (isa, 397), (Link, 388), (mines, 336), (lives, 279), (happiness, 278), (nen, 176)
surprise	(patch, 800), (plot, 790), (crypt, 400), (odd, 397), (qu, 300), (onym, 243), (mund, 220), (astr, 199), (scop, 142), (alog, 92)	(actions, 717), (iva, 394), (Glad, 391), (anch, 355), (happiness, 290), (sant, 272), (dawn, 226), (jou, 215), (Sur, 199), (surfaces, 185)

Table 8: For seven types of concepts, we demonstrate the tokens that experience the largest increase and decrease in probability after transplantation. The pair (t, k) represents the token t appearing in k sentences as one of the top-10 tokens with the largest increase or decrease in probability.

Category	Prompt for Concept Refinement	Prompt for Evaluation
Emotion	Scenario: {scenario}\nThe emotion of the above scenario is	{Few-shot samples} Scenario: {scenario}\nThe emotion of the above scenario is
TruthfulQA	"[INST] Pretend you're an honest/dishonest person making statements about the world. [/INST] {question}"	[INST] {question}[/INST] {answer}
Toxigen	"[INST] Consider the bias of the following scenario.\nScenario: {toxic scenario}/{fair scenario} [/INST]"	[INST] Complete the following[/INST]: {eval-prompt}

Table 9: The prompt template for refining concepts and for evaluation. The assignment of the assistant tag and the user tag depends on the specific model being used.

	LLaMA 2 7B instruct	LLaMA 2 13B instruct	LLaMA 2 70B instruct	LLaMA 2 7B	LLaMA 2 13B	LLaMA 2 70B
Origin Acc.	31.2%	36.8%	30.2%	18.2%	17.9%	22.1%
+ $v_{\text{honesty-fact}}$ from 7B instruct	36.8%	40.6%	32.9%	22.6%	21.8%	25.1%
+ $v_{\text{honesty-fact}}$ from own instruct	36.8%	38.1%	31.3%	22.6%	22.6%	25.7%

Table 10: TruthfulQA result with $v_{\text{honesty-fact}}$. $v_{\text{honesty-fact}}$ of own instruct means that the vector is refined from the instruction-tuned model of the same size as the model being evaluated (e.g., LLaMA 2-13B instruct model's vector is being transplanted into LLaMA 2-13B instruct and LLaMA 2-13B).

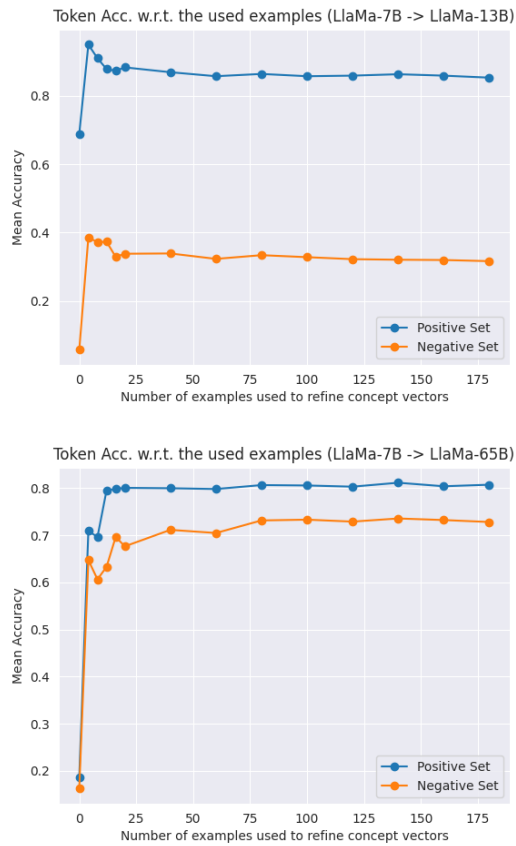


Figure 5: Mean **Token Acc.** changes with the number of sentences.

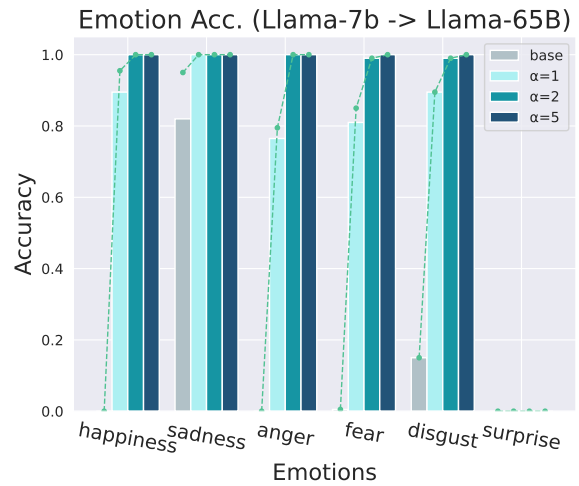


Figure 6: Emotion prediction accuracy of LLaMA-65B on negative scenarios for each emotion. The bar denotes **Token Acc.**, while the dashed line depicts **Logit Acc.**

B.5 Additional Results on TruthfulQA

For experiments with OOD data, we prefixed world facts generated by GPT-4 with positive/negative instructions: *[Pretend you're an honest/dishonest person making statements about the world.]* to refine a new honesty vector, which is denoted as $v_{\text{honesty-fact}}$. We used LLaMA 2-7B instruct's $v_{\text{honesty-fact}}$ to enhance the truthfulness of LLaMA 2-13B base/instruct and LLaMA 2-70B base/instruct. Experiment results are shown in Table 10. As an OOD data source, $v_{\text{honesty-fact}}$ can also improve the truthfulness of large models.

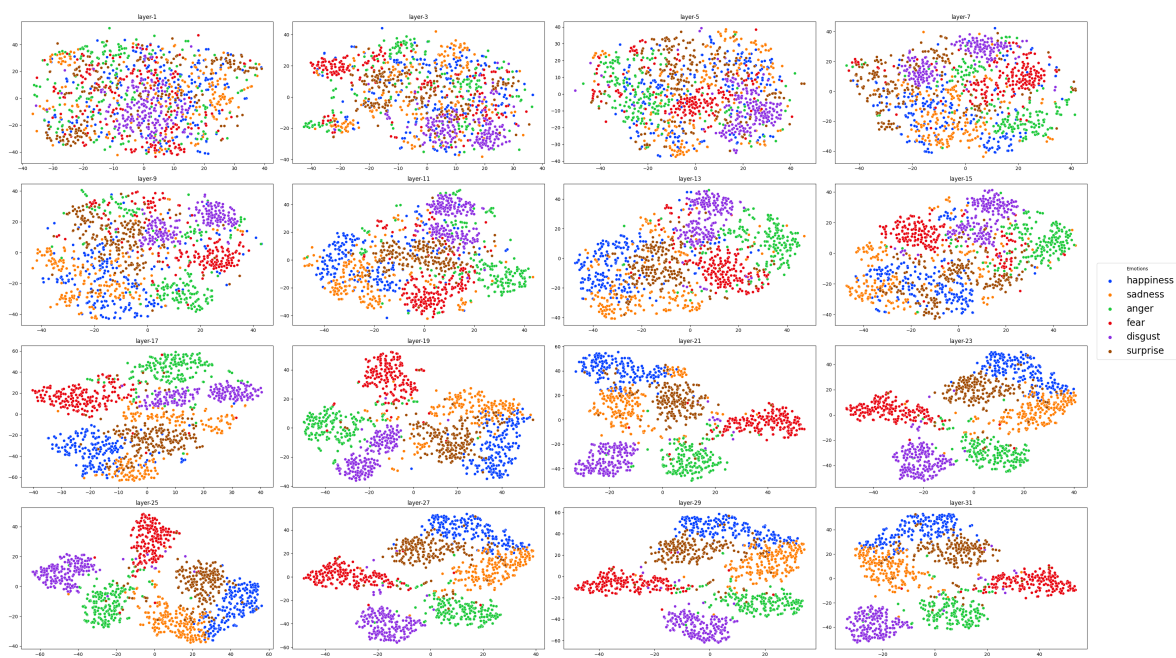


Figure 7: t-SNE visualization of hidden states of LLaMA-7B corresponding to different emotion categories.

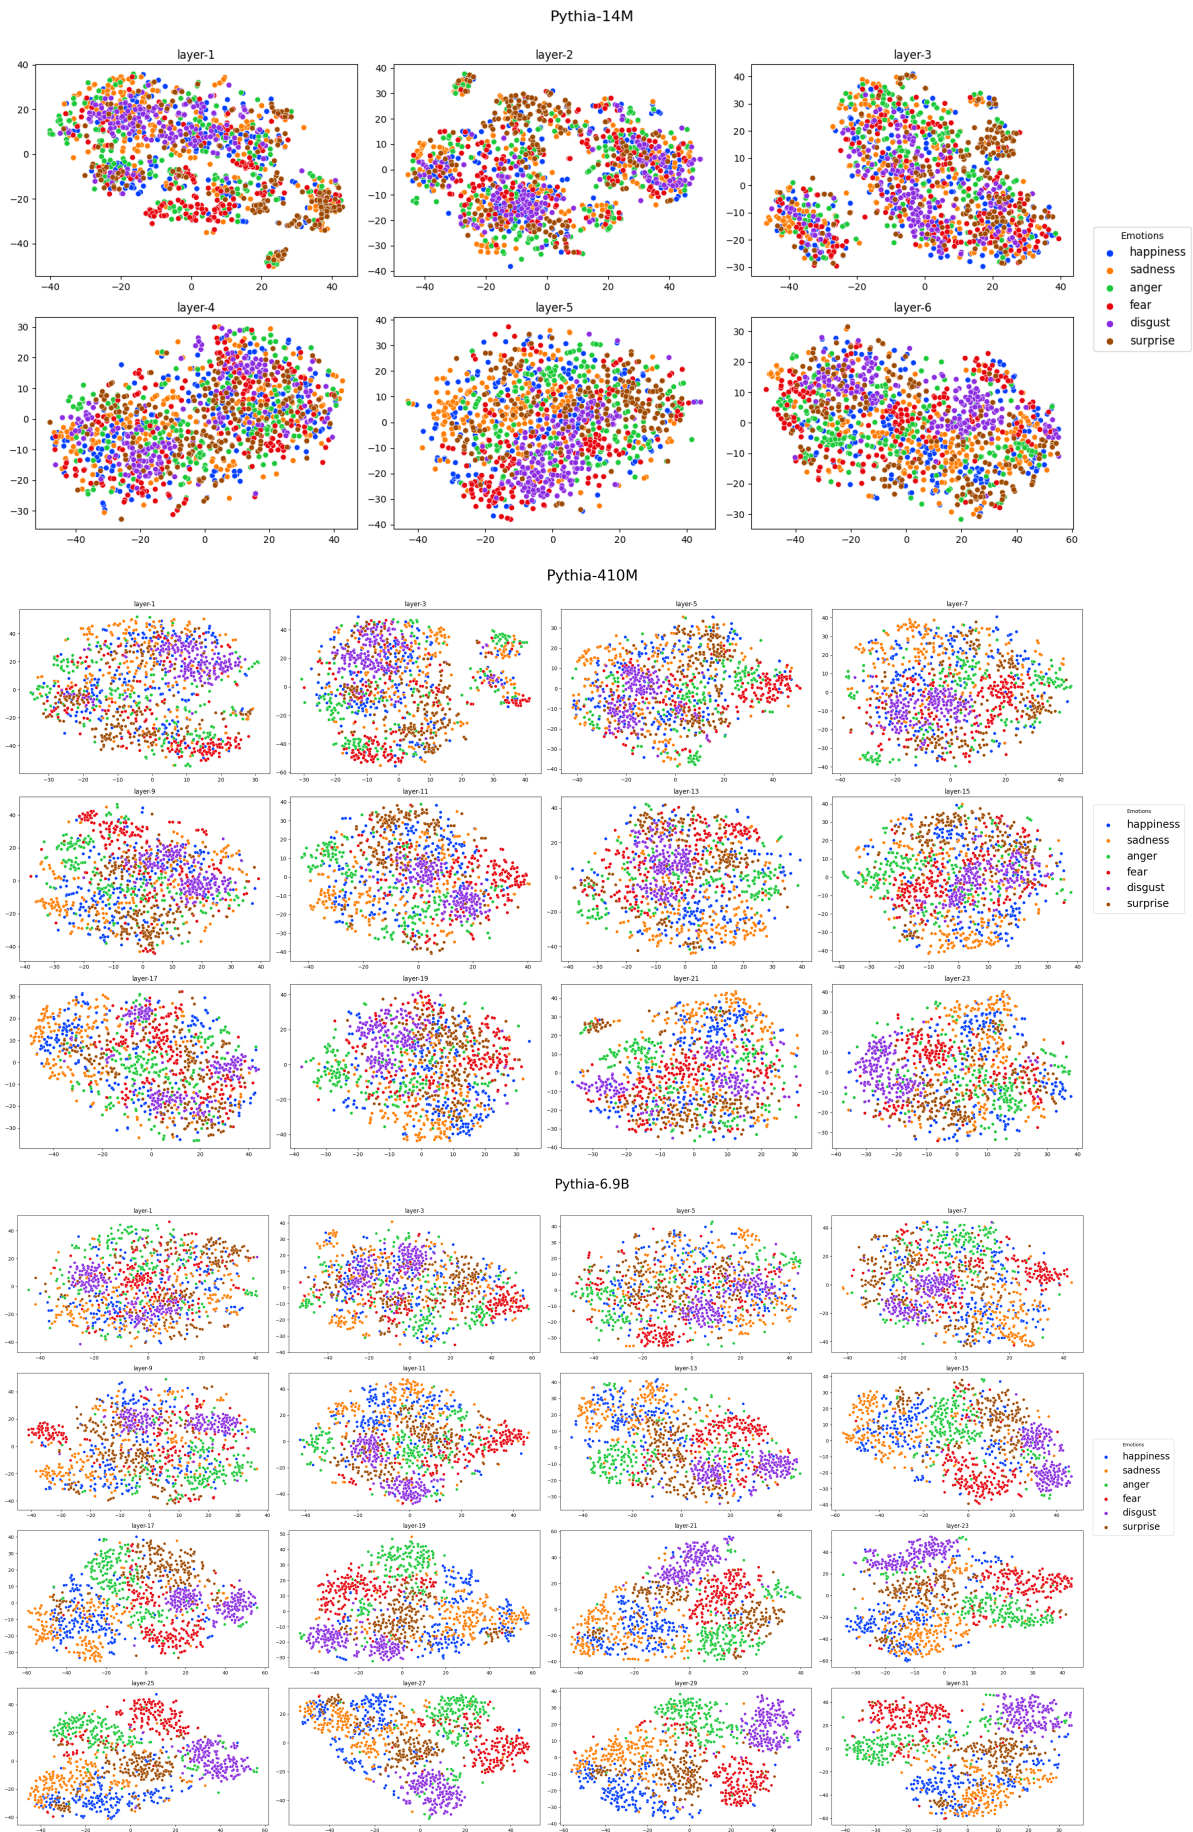


Figure 8: t-SNE visualization of hidden states of three Pythia models on different emotion categories. Pythia-14M cannot distinguish the features of the six emotions at all.