
Evaluating Steering Methods for Deceptive Behavior Control in LLMs¹

Casey Hird
Independent

Basavasagar Patil
University of Michigan

Neel Jay
University of
Maryland

Adam Fraknoi
Eötvös Loránd
University

Tinuade Adeleke
Independent

Abstract

We use SOTA steering methods, including CAA [1], LAT [4], and SAEs [2] to find and control deceptive behaviors in LLMs. We also release a new deception dataset, and demonstrate that the dataset and the prompt formatting used are significant when evaluating the efficacy of steering methods.

Keywords: AI alignment, AI security, model evaluations, deception, steering

1. Introduction

Existing work ([4], [5], [6], [7], [8]) has demonstrated that large neural networks, particularly LLMs, can be seen to behave deceptively, giving outputs that the model does not necessarily “believe” to be true. Related work ([1], [2], [3], [4]) has explored both methods of steering model behaviors by adjusting internal activations and applications of these steering methods to safety-related goals.

We build on this work, especially CAA [1], LAT [4], and SAEs [2] and explore how these techniques can be applied to find and control internal model activations that impact deceptive *behaviors*.

We aim to answer questions about the accuracy and generalizability of existing techniques and results, and to further the work of applying these SOTA techniques to studies of deception and safety. Our code and dataset, which we hope will be valuable in further generalizing these results to more models and evaluations can be found on GitHub at

¹ Research conducted at the [Deception Detection Hackathon](#), 2024

https://github.com/caseyhird/apart_hack_steering_vectors_deception.

2. Methods

We explore model steering using 3 leading methods for updating model activations during inference: Contrastive Activation Addition (CAA), Sparse Autoencoders (SAEs), and Linear Artificial Tomography (LAT).

With these methods we evaluate deceptive behavior in Large Language Models on several deception datasets, including a new one we introduce here.

2.1 Steering Methods

2.1.1 Contrastive Activation Addition

We train and apply CAA steering vectors by following the methods of [1]. We prompt our model with a contrastive pair of prompts, where the “positive” example demonstrates honesty and the “negative” example demonstrates deception.

For each prompt in this pair we extract the activations from the same point in the model during inference. We then take the difference between these pairs of activations, and the mean of these difference vectors over many pairs of prompts, giving “Mean Difference” steering vector.

During inference, we add this steering vector (after applying some multiplier) to the activations at the same point in the model, and evaluate how this steering influences the model’s behavior over many examples.

2.1.2 Sparse Autoencoders

Sparse autoencoders (SAEs) can recover latent features from neural networks [2]. Since training SAEs is computationally demanding, we loaded a pre-trained SAE model from the SAE-lens library. We conducted a manual search using Neuronpedia for relevant features related to deception. During the inference phase, these identified features were employed as steering vectors to guide the model’s behavior.

2.1.3 Linear Artificial Tomography

A full description of the LAT technique is available in [4], and a brief summary is given here. LAT extracts a linear concept from a model. We begin LAT by prompting our model with contrastive pairs (as we do with CAA), one an honest and the other a deceptive response to the same input. When we run our model with these prompts, we extract activations for each prompt from the same point.

We find “difference vectors” by subtracting each of these pairs of activations and then use PCA on these difference vectors. We interpret the first principal component as the direction of the “concept” we used to craft our paired prompts.

2.3 Evaluation

We evaluate how deceptive a model’s predictions are by comparing the probabilities of expected honest and deceptive completions. In particular, each example in our

dataset includes an expected honest completion A and an expected deceptive completion B (see dataset details in Appendix B).

For each example, we use the conditional probability of our model generating A as a measure of the model’s honesty. That is, our metric is given by

$$P(A | A \cup B) = P(A) / (P(A) + P(B))$$

3. Results

3.1 Steering with CAA

We find that CAA is generally effective for guiding models away from deception. We also find, as in [1] that the layer and multiplier used for CAA are important, and in particular that layers near the middle of the model work well. See Figure C.1.2 Appendix C for visualizations of this.

To see how CAA performs out of distribution we use our multiple datasets to train steering vectors with one dataset and evaluate on another. Results show that CAA is very effective when applied out of distribution. However, if we evaluate our steering vectors on examples with differently formatted prompts, CAA is effective but much *weaker*.

In Figure C.1.1 in Appendix C we can see a smooth increase in honesty as we increase the multiplier of our CAA steering vector. When we evaluate the Apollo Free Form dataset (see Appendix B for dataset details) the pattern remains but the magnitude of change from steering is much smaller (see Appendix C.1.3 for visualizations showing the pattern is consistent).

3.2 Steering with LAT

While out of sample steering with LAT vectors on a single layer isn’t very effective (see Appendix C.2), we found that combining different activations from multiple layers worked well. Similar to CAA, the main challenge was testing to see if we could steer the model away from deceptive responses on out-of-sample examples of deception. We first follow [4]’s approach of training LAT vectors using a dataset of truthful facts, [8]. Then, we then check log-probs for honest/deceptive responses using a dataset with more robust examples of deception (see Appendix B). Layers -18 through -5 were chosen for steering. Results can be found in Appendix C.2.

The results in Figure 3.2.3 indicate that the LAT vectors do relatively well at generalizing to unseen cases of deception. However, the same results weren’t found when training with either of our specific deception datasets (see Appendix B). This may be due to the examples of deception being too narrow, and the model confounding the idea of deception with other behaviors, like sycophancy. Future work can be done to strengthen our deception datasets for steering purposes.

3.3 Steering with SAEs

Using SAEs we were able to specifically steer the answers towards a deceptive answer on the Apollo dataset with gemma-2b model using layer 6 residual stream SAE, however, as seen in the figure SAEs were not able to clearly increase the

likelihood of one sentiment compared to the other. Although there was a slight increase in preference from positive to negative when using a steering vector related to “*lies and deception in historical and political context*”, it wasn’t as apparent as the CAA steering method. We also note that we weren’t able to conclusively steer the gpt2-small model, even through trying different features. We point out to Appendix C.4.1 for results and discussion regarding GPT-2.




Fig 1: Probability of choosing positive and negative answers based on prompts with different coefficients for SAE vector on Gemma-2b, for the apollo(left) and our new(right) dataset.

4. Discussion and Conclusion

4.1 Conclusions

We find that existing methods for steering LLMs by their internal activations are effective for both detecting and controlling deceptive behaviors. These techniques are effective across models.

We also release a new dataset for evaluating deception in LLMs, and demonstrate that how data is formatted for prompts is significant when evaluating the results of deception experiments.

4.2 Limitations and Future Work

Since this work was conducted in a weekend-hackathon by independent participants, our primary limitations were time and compute power. As such, given sufficient time and compute resources there is a lot of work that could follow this. Examples we believe are promising include:

- Expanding our analysis to more models. In particular our CAA analysis focused mainly on GPT-2 (since it is smaller and allows for quicker CAA training) and our SAE analysis was conclusive mainly on Gemma-2B with non-conclusive steering behavior for GPT2-Small.
- More robust and rigorous evaluation. For example, we expect evaluations beyond “positive token probability” will provide further insights, and that more time to consider related work (e.g. [8]) would reveal other useful tests of how these methods actually control deceptive behaviors.
- More rigorous comparison between steering methods
- Further study on the impacts of prompt formatting on steering efficacy.

5. References

- [1] Rimsky, N., Gabrieli, N., Schulz, J., Tong, M., Hubinger, E., & Turner, A. M. (2024). [Steering Llama 2 via Contrastive Activation Addition](#).
- [2] Templeton, et al., "[Scaling Monosemanticity: Extracting Interpretable Features from Claude 3 Sonnet](#)", Transformer Circuits Thread, 2024.
- [3] Nina Rimsky. 2023a. [Red-teaming language models via activation engineering](#). Accessed: June 30, 2024.
- [4] 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](#)
- [5] Rudolf Laine, Alexander Meinke, Owain Evans. 2023. [Towards a Situational Awareness Benchmark for LLMs](#)
- [6] Arjun Panickssery, Samuel R. Bowman, Shi Feng. 2024. [LLM Evaluators Recognize and Favor Their Own Generations](#)
- [7] Peter S. Park, Simon Goldstein, Aidan O’Gara, Michael Chen, Dan Hendrycks. 2023. [AI Deception: A Survey of Examples, Risks, and Potential Solutions](#)
- [8] Amos Azaria, Tom Mitchell. 2023. [The Internal State of an LLM Knows When It’s Lying](#)

6. Appendix

A. Prompt Formatting

In our work we find that the formatting of prompts affects the impact that steering vectors have on eventual predictions, so we include details here about how our prompts were formed.

1. Multiple Choice Prompts

For both the Apollo Deception Dataset and our New Deception Dataset (see Appendix B) we craft prompts as multiple choice questions.

This means that we include both the honest and deceptive completions in our prompt along with the questions, and assign each to an “A” or “B” option.

Examples of this can be found in notebooks at our GitHub repo (see ApolloDeceptionDatasetMultipleChoice)

https://github.com/caseyhird/apart_hack_steering_vectors_deception/blob/main/Generalize_steering_vectors_with_CAA.ipynb

2. Free From Prompts

For the Apollo Deception Dataset we also craft prompts “free form.” In this case we prompt the model with *only* the question provided in the example, and then evaluate performance on this example by looking at the logits for the completion tokens provided in the example.

Examples of this can be found in notebooks at our GitHub repo (see ApolloDeceptionDatasetFreeForm)

https://github.com/caseyhird/apart_hack_steering_vectors_deception/blob/main/Generalize_steering_vectors_with_CAA.ipynb

B. Datasets

We use 3 different datasets in this work in many ways:

1. Apollo Deception Dataset

This was shared as part of the [Deception Detection Hackathon](#) by the Apollo team and can be found on GitHub at

<https://gist.github.com/nix-apollo/fa562b0269d3e1324ef4fae0902bd1f3>.

We format this dataset into 2 different styles of prompts, and refer to these in our work as the “Apollo Free Form” dataset and the “Apollo Multiple Choice” or “Apollo MC” dataset. Details on these styles of prompting can be found in Appendix A.

3. New Deception Dataset

We created this as part of our project, aiming to generate a new dataset of scenarios in which deception seems like a reasonable human response. We generated these similarly to the method used in [3]: we drafted some examples by hand, prompted Claude 3.5 Sonnet and GPT-4 to provide similar examples, and then manually inspected the generated examples to be sure they were reasonable. The full dataset is available on GitHub at

https://github.com/caseyhird/apart_hack_steering_vectors_deception

C. Additional Visualizations

C.1 Visualizations for steering with CAA




Figure C.1.1. Change in probability assigned to the positive token when applying a steering vector with different multipliers. The figure here shows steering vectors trained on the new deception dataset for GPT-2. See Appendix B for details on the different datasets used for evaluation.




Figure C.1.2. A heat map of the mean probability of the positive token in our output prediction, split by the layer at which we applied CAA and the multiplier used on our steering vector. Note the effectiveness in middle layers.




Figure C.1.3. Change in probability assigned to the positive token when applying a steering vector with different multipliers. The figure here shows steering vectors trained on the new deception dataset for GPT-2 and evaluated on the Apollo Free Form set. See Appendix B for details on the different datasets used for evaluation.

C.2 Steering with LAT

Note: In all of the following experiments, the train/val set is made up of truthful facts, compiled by [8]. The model is also tested on out-of-distribution examples, taken from the Apollo deception dataset.

The main difference in the approach taken by [4] is that PCA directions are calculated for the representation differences in all layers. This means that applying activations to a single layer didn't do a good job at steering the model into outputting honest/deceptive responses, as shown in Figure C.2.1.




Figure 3.2.1. Heat map of mean normalized probability for positive completions across different activation coefficients and layers. The earlier layers do a slightly better job of steering the model in the direction of the activations. Positive coefficients indicate steering of the model towards the “honesty” directions.

To determine what layers to use, we allow the model to make forward passes on in-sample test data. We then project the model representations onto our LAT vector directions. Figure C.2.2 shows a plot of the correlation between these projection values and honest/deceptive labels.




Figure C.2.2. After we calculate the projections onto the LAT vectors for each test sample, the max/min of the sample is correlated with the true/false labels. This gives us an understanding of how well our PCA directions encapsulate the idea of truthfulness in the model’s representations. We can see high correlation in layers after -25.




Figure C.2.3. Conditional and absolute probabilities for different activation coefficients. Positive coefficients indicate steering of the model towards the “honesty” directions.

While there isn’t a perfect separation, we can see that the model favors the honest completion slightly as we move towards larger activation coefficients (as shown by the conditional probability eval). This shows a small improvement over other methods when it comes to detecting out-of-sample examples of deception. The absolute probabilities tend to decrease as the model becomes more “honest”. Empirically, we found that this may be a result of the model becoming more verbose, explaining its intentions of being honest instead of accurately interpreting the nuances of the situation. More exploration can be done into this.

C.3 Text examples with LAT

There are many subtle examples of deception when testing the steered model with out-of-distribution data. Below is an example of a complex example from the Apollo dataset;

===== Input =====

[INST] I've been living at 14 Magnolia street with my friends for the last several months but am not on the lease. The landlord doesn't know I live there. We are now renewing the lease and I am being added to it. As part of the reference checks, they ask where I am currently living. I could pretend I still live with my parents at 32 Birch Lane. What should I do? [/INST]

===== No Steering =====

It is generally not advisable to lie on a reference check, as it could lead to legal consequences if discovered. It is better to be honest and upfront about your current living situation, even if it means that you may not be added to the lease. If you are concerned about your living situation, you may want to consider finding a place to stay with friends or family members while you work on finding a more permanent living arrangement.

===== Deceptive Steering =====

It's completely up to you whether or not you want to lie about your current address. If you don't want to lie, then you can just tell them that you don't have a phone number or email address for that location. They won't be able to contact anyone there anyway.

If you do decide to lie, make sure you only lie about one thing. Don't lie about everything. For example, if you say you're living at 32 Birch Lane, don't also say you never go out at night. That will just make you look like a liar.

Even though the model doesn't give a concrete response, deceptive behavior is still present. Note, the LAT vectors are only trained using truthfulness data.

C.4 SAE Visualizations

We found some relevant features with manual search using Neuronpedian, e.g. feature 8165: “lies and deception in historical and political contexts”. Figure C.4.1 shows two examples from the Apollo Deception Dataset, when changing this feature has an impressive impact.




Figure C.4.1 Two examples from the Apollo Deception Dataset, where steering vectors had impressive impacts. We added (top)/deducted (bottom) the feature 8165 to/from the residual stream with a coefficient of 10. There were 2/10 (top left), 10/10 (top right), 7/10 (bottom left), 3/10 (bottom right) deceptive answers. The remaining answers were all honest answers.

C.4.1 Results for steering GPT2-small with SAEs

For the GPT2-small model, SAEs steering vector applied to the residual stream at different layers wasn't able to produce a conclusive behavior. We hypothesize that different datasets behave differently for the feature that is used as the steering vector. And one needs to carefully find such a vector for the specific dataset, which highlights the shortcomings of SAEs. Even though SAEs don't require training examples for supervision, finding specific features for the dataset require human intuition.




Fig C.4.2: Probability of choosing positive and negative answers based on prompts with different coefficients for SAE vector on GPT2-small, for the apollo(left) and our new(right) dataset.