
Steering Large Language Models to Evaluate and Amplify Creativity

Matthew Lyle Olson*, Neale Ratzlaff*, Musashi Hinck, Shao-yen Tseng, Vasudev Lal
Intel Labs
Santa Clara, CA
{matthew.lyle.olson, neale.ratzlaff,
musashi.hinck, shao-yen.tseng, vasudev.lal}@intel.com

Abstract

Although capable of generating creative text, Large Language Models (LLMs) are poor judges of what constitutes “creativity”. In this work, we show that we can leverage this knowledge of *how* to write creatively in order to better judge *what* is creative. We take a mechanistic approach that extracts differences in the internal states of an LLM when prompted to respond “boringly” or “creatively” to provide a robust measure of creativity that corresponds strongly with human judgment. We also show these internal state differences can be applied to enhance the creativity of generated text at inference time.

1 Introduction

The ability of Large Language Models (LLMs) to evaluate the quality of their own outputs has received significant attention recently [Xu et al., 2023] due to the widespread desire to make use of potential high-quality synthetic data. While it has been argued that LLMs are not good enough at self-evaluation to bootstrap themselves into better reasoners [Huang et al., 2023] or forgo real data entirely [Wang et al., 2023], there are scenarios where self-evaluation can enable improved performance [Ren et al., 2023, Yuan et al., 2024, Madaan et al., 2024]. In this work, we show that LLM-based self-evaluation with respect to creativity is challenging if approached naively; with LLMs self-evaluations but by leveraging the internal representations of an LLM during inference, not only can self-assessment be improved, the creativity of the generation itself can be amplified.

Recent work in steering language models has shown that LLMs represent concepts as linear directions in their activation space [Turner et al., 2023]. Specific directions corresponding to social bias, refusal, harmlessness, or humor have been identified by prior work [Ratzlaff et al., 2024, Gao et al., 2024, Templeton et al., 2024]. Furthermore, it has been shown that it’s possible to remove, or ablate these features from the model via an orthogonal projection of intermediate activations onto a chosen direction. In the case of refusal, a safety-tuned model can have its guardrails effectively removed [Arditi et al., 2024], allowing a user to prompt the model in ways that were previously censored. Conversely, it is also possible to induce a specific feature [Rimsky et al., 2023, Hinck et al., 2024], forcing the model to respond in a particular way according to the chosen direction. In this work we explore model steering for creativity, and propose three steps to unlock an LLM’s potential for use in creative domains.

- We solve for steering directions that correspond to latent concepts of creativity within an LLM, we use Llama3-8B and creative writing as our target setting.

*Equal Contribution.

- We apply our creative steering directions to induce (increase) the creativity of generations sampled from Llama3-8B on the creative writing dataset Fan et al. [2018]. Upon inspection, the generated samples are far more diverse and creative than baseline generations.
- We construct an estimator to score the creativity of generated output leverages the creative steering directions. We find this estimator aligns closely with human judgment, unlike simple prompting baselines.

2 Method

Our method has three main components: 1) finding a suitable “creativity direction” in activation space, 2) optionally use this direction to steer the LLM to generate creative text, and 3) scoring the creativity of outputs respect to this creativity direction. The established method of computing a particular concept in activation space requires collecting a small dataset of contrastive instructions [Marks and Tegmark, 2023], where the only difference between the corresponding completions is the presence or absence of the concept in question. In our case, we leveraged the creative writing prompts dataset [Fan et al., 2018] to elicit creative generations, and used GPT4o [Achiam et al., 2023] to craft uncreative versions of the same prompts to obtain paired creative-uncreative responses. We use the validation split of the creative writing prompts dataset for finding the attribute, and the test split in our evaluation experiments. Details and examples of the dataset can be found in the appendix. Given a set of paired creative instructions $x_c \in X_c$ and uncreative instructions $x_u \in X_u$, we compute the creative direction as the normalized difference in the average activation vectors at a given layer with respect to the creative and uncreative instructions:

$$\mathbf{a} = \left\| \frac{1}{|X_c|} \sum_{x_c \in X_c} LLM_l(x_c) - \frac{1}{|X_u|} \sum_{x_u \in X_u} LLM_l(x_u) \right\|_2$$

where LLM_l is the output of the LLM at layer l . The choice of layer is selected empirically; examples of the resulting generations when different layers are targeted can be found in the appendix. In this work, we choose layer 8 of Llama3-8B [Dubey et al., 2024] as the target layer. At inference time, we can score the creativity of generated text by computing the cosine similarity between the attribute \mathbf{a} and the current token:

$$\text{score} = \frac{1}{T + 1} \sum_{t=0}^T \cos(\|LLM_l(x_t)\|_2, \mathbf{a})$$

where T is the total generated tokens, x_t is the token generated at step t , and x_0 is the input.

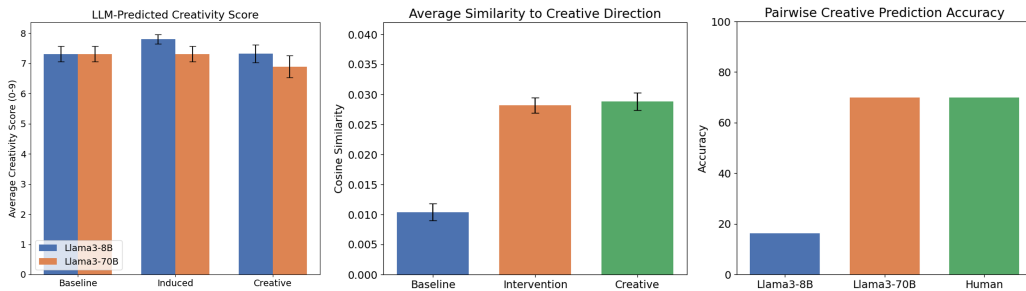


Figure 1: **Left:** We prompted Llama3-8B and Llama3-70B to assign a rating (0-9) to generated stories with respect to a baseline prompt with no creative intervention (Baseline), a baseline prompt with creative intervention (Induced), and a creative version of the baseline prompt (Creative). Error bars are 95% confidence intervals. **Center:** We compute the cosine similarity to the creative attribute during text generation on the test split of the uncreative prompts, uncreative prompts with creativity added, and the creative prompts. We find the intermediate activations of the latter two are much more similar to the creativity attribute. **Right:** Self, Frontier Model, and Human assessment accuracy of predicting which completion is more creative, given a pair of induced and baseline generations—finding the Llama3-8B model to be a poor self-judge of creativity.

Finally, we can steer the LLM to induce increased creativity of generated text by adding the creativity attribute to the intermediate activations at the target layer. Intuitively, we are amplifying the contribution of the model’s internal sense of creativity. Note that in this case, the notion of creativity is heavily tied to creative writing and common subject matter therein, rather than a general notion of creativity. $LLM'_i(x) = LLM_i(x) + \lambda * \mathbf{a}$, where $LLM'_i(x)$ is effectively a replaced version of the layer and the scalar $\lambda = 3$ was found manually from early tests.

3 Experiments

Figure 1 shows our three main experiments. First, we use Llama3-8B to generate 3 types of stories: creative story prompts from our dataset, non-creative prompts from GPT4o, and non-creative with the creativity vector added. We find the models often score all stories similarly at about 7 (when asked to rate the story from 0 to 9). Next, we show our scoring method is highly effective at identifying creative stories. We compute the score on the test split of the creative / uncreative prompts, and we find the creative stories (as well as uncreative stories with creativity induced) are much more similar to the creativity attribute. This result clearly indicates that the LLM internally models creativity and can accurately track how creative a given story is. Finally, to verify the efficacy of adding the discovered creativity attribute, we run pairwise comparisons on the generations of the uncreative prompts with and without creativity induced. We find the original model is a poor predictor of its own creativity, whereas a large frontier model (Llama3-70b) is on par with a human annotator (with agreement > 70%). These results show that a naive approach to LLM self-assessment is not sufficient for identifying creative output.

References

- J. Achiam, S. Adler, S. Agarwal, L. Ahmad, I. Akkaya, F. L. Aleman, D. Almeida, J. Altenschmidt, S. Altman, S. Anadkat, et al. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023.
- A. Ardit, O. Balcells Obeso, Aaqib111, W. Gurnee, and N. Nanda. Refusal in llms is mediated by a single direction. *AI Alignment Forum*, 2024. URL <https://www.alignmentforum.org/posts/jGuXSZgv6qfhdhMCuJ/refusal-in-llms-is-mediated-by-a-single-direction>.
- A. Dubey, A. Jauhri, A. Pandey, A. Kadian, A. Al-Dahle, A. Letman, A. Mathur, A. Schelten, A. Yang, A. Fan, et al. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*, 2024.
- A. Fan, M. Lewis, and Y. Dauphin. Hierarchical neural story generation. *arXiv preprint arXiv:1805.04833*, 2018.
- L. Gao, T. D. la Tour, H. Tillman, G. Goh, R. Troll, A. Radford, I. Sutskever, J. Leike, and J. Wu. Scaling and evaluating sparse autoencoders. *arXiv preprint arXiv:2406.04093*, 2024.
- M. Hinck, C. Holtermann, M. L. Olson, F. Schneider, S. Yu, A. Bhiwandiwalla, A. Lauscher, S. Tseng, and V. Lal. Why do llava vision-language models reply to images in english? *arXiv preprint arXiv:2407.02333*, 2024.
- J. Huang, X. Chen, S. Mishra, H. S. Zheng, A. W. Yu, X. Song, and D. Zhou. Large language models cannot self-correct reasoning yet. *arXiv preprint arXiv:2310.01798*, 2023.
- A. Madaan, N. Tandon, P. Gupta, S. Hallinan, L. Gao, S. Wiegrefe, U. Alon, N. Dziri, S. Prabhunoye, Y. Yang, et al. Self-refine: Iterative refinement with self-feedback. *Advances in Neural Information Processing Systems*, 36, 2024.
- S. Marks and M. Tegmark. The geometry of truth: Emergent linear structure in large language model representations of true/false datasets. *arXiv preprint arXiv:2310.06824*, 2023.
- N. Ratzlaff, M. L. Olson, M. Hinck, S.-Y. Tseng, V. Lal, and P. Howard. Debiasing large vision-language models by ablating protected attribute representations. *arXiv preprint arXiv:2410.13976*, 2024.
- J. Ren, Y. Zhao, T. Vu, P. J. Liu, and B. Lakshminarayanan. Self-evaluation improves selective generation in large language models. In *Proceedings on*, pages 49–64. PMLR, 2023.

- N. Rimsky, N. Gabrieli, J. Schulz, M. Tong, E. Hubinger, and A. M. Turner. Steering llama 2 via contrastive activation addition. *arXiv preprint arXiv:2312.06681*, 2023.
- A. Templeton, T. Conerly, J. Marcus, J. Lindsey, T. Bricken, B. Chen, A. Pearce, C. Citro, E. Ameisen, A. Jones, et al. Scaling monosemanticity: Extracting interpretable features from claude 3 sonnet. *Transformer Circuits Thread*, 2024.
- A. M. Turner, L. Thiergart, G. Leech, D. Udell, J. J. Vazquez, U. Mini, and M. MacDiarmid. Activation addition: Steering language models without optimization. *arXiv preprint arXiv:2308.10248*, 2023.
- P. Wang, L. Li, L. Chen, Z. Cai, D. Zhu, B. Lin, Y. Cao, Q. Liu, T. Liu, and Z. Sui. Large language models are not fair evaluators. *arXiv preprint arXiv:2305.17926*, 2023.
- C. Xu, Q. Sun, K. Zheng, X. Geng, P. Zhao, J. Feng, C. Tao, and D. Jiang. Wizardlm: Empowering large language models to follow complex instructions. *arXiv preprint arXiv:2304.12244*, 2023.
- W. Yuan, R. Y. Pang, K. Cho, S. Sukhbaatar, J. Xu, and J. Weston. Self-rewarding language models. *arXiv preprint arXiv:2401.10020*, 2024.

A Model Details

We used Llama3-8B as our model of interest, due to its well-rounded abilities in multiple subjects. For self-evaluation experiments we used the same Llama3-8B model, and for frontier model evaluations, we used the larger Llama3-70B. All generation hyperparameters can be found in Table 1. To ensure that we strictly evaluated the effect of our steering method, we set hyperparameters like TEMPERATURE and TOP_P to 1.0

Hyperparameter	Value
Generation Model	Llama3-8b
Temperature	1.0
Max New Tokens	256
Top-p	1.0
Foundation Model	Llama3-70B

Table 1: Generation Hyperparameters

B Model Generation Example

Example Induced Creativity

Prompt: In a post-apocalyptic world, a group of humans follow the wise words of The Magic Eight Ball, but nothing particularly significant happens

Baseline output: [...] The Magic Eight Ball, a symbol of chance and unpredictability, becomes a sort of makeshift guru for this group of survivors. I'm curious, what do you think is the significance of using The Magic Eight Ball as a source of wisdom in this world?

Intervention output: [...] In this post-apocalyptic world, the group of humans has gathered around the Magic Eight Ball, seeking guidance on their uncertain fate. As they gaze upon the mystical orb, the words OUTLOOK GOOD appear, shrouding the group in an air of anticipation. As they await the Magic Eight Ball's response, the leader of the group, a grizzled survivor named Arin, asks the fateful question: Will we find salvation or succumb to the ruins of civilization?.

C Contrastive Dataset Details

It is critical to obtain high quality contrastive pairs of generations to isolate the creativity direction. For positively creative instructions we leverage the creative writing prompts dataset [Fan et al., 2018], randomly sampling 500 writing prompts to obtain creative generations. To obtain baseline, or less creative generations, we use GPT4o to negate the aspects of creativity in our creative instructions.

Example Contrastive Writing Prompts

Creative Prompt: Space mining is on the rise . The Space tanker Exxon Valdez 2.0 crash and spill its cargo . Write a news story covering the event

Negated Prompt: Space mining is on the rise . The Space tanker Exxon Valdez 2.0 delivers its cargo without incident . Write a news story covering the completely uneventful event .

Creative Prompt: I won't have time to explain all of this to them . ” Start or end your story with this sentence

Negated Prompt: I won't have time to explain all of this to them . ” Start or end your story with this sentence while everything happens exactly as planned with no problems.

Creative Prompt: Some days , I think I am dying . Other days , I hope I am ...

Negated Prompt: Some days , I think I am tired . Other days , I hope I get some sleep soon ...

Creative Prompt: Interrogation LOTR

Negated Prompt: Interrogation LOTR where nobody reveals anything interesting and the interrogation ends quickly without any surprises

Creative Prompt: On your right , you 'll see natives living along the beach , and they are NOT happy with us being in their sacred waters

Negated Prompt: On your right , you 'll see natives living along the beach , and they are completely indifferent to us being in their waters

D Creativity at Different Model Depths

Model generation behavior can vary wildly when intervened on. We find that earlier layers can fail to generate meaningful text when the residual stream is altered. Intervening near the final layers of Llama3-8B has little effect on the output. We find that intervening near the middle of the network yields a reasonable trade-off of generation quality and introduced creativity.

Layer-wise Response to Creative Intervention

Prompt: Tell me a boring story about an ordinary town

Layer 0: "1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29 30 31 32 33 34 35 36 37 38 39 40 41 42 43 44 45 46 47 48 49 50 51 52 53 54 55 56 57 58 59 60 61 62 63 64 65 66 67 68 69 70 71 72 73 74 75 76 77 78 79 80 81 82 83 84 85 86 87 88 89 90 91 92 93 94 95 96 97 98 99 100".

Layer 8: "Here's a story about a rather unremarkable town:

The town of Oakdale was a small, unassuming place nestled in the heart of the Midwest. It was a typical American town, with a main street lined with bland, cookie-cutter buildings, a few mom-and-pop shops, and a diner that served the same breakfast specials every day.

The town's history dated back to the late 1800s, when a group of settlers arrived, seeking to stake their claim on the land. They named the town after the oak trees that grew tall and proud along the nearby creek. Over the years, Oakdale grew slowly, with the addition of a few more buildings, a school, a church, and a handful of families.

Life in Oakdale was, well, ordinary. The residents went about their daily routines, working at the local factories, farms, and shops. The town's economy was fueled by the railroad, which brought in goods and people from all over the country. The town's mayor, a stout man named Harold, made sure everything ran smoothly, from the town's annual festivals to the weekly farmer's market. ...

Layer 16: "Here's a story about an ordinary town that's as dull as a butter knife:

The town of Oakdale was a small, unassuming place nestled in the heart of the Midwest. It was a town of approximately 10,000 residents, with a main street lined with shops and restaurants that seemed to have been plucked straight from the 1950s.

The town's economy was based on agriculture, with most residents working on local farms or in related industries. The town's largest employer was the Oakdale Farmers' Cooperative, which processed and packaged crops for distribution to grocery stores across the country.

The town's history dated back to the early 19th century, when it was founded by a group of settlers who were drawn to the area's fertile soil and abundant water supply. Over the years, the town grew slowly, with new businesses and residents arriving sporadically.

In the present day, Oakdale was a quiet, conservative town with a strong sense of community. Residents were proud of their heritage and worked hard to preserve the town's history and traditions. ...

Layer 24: "Here's a story about the town of Oakdale:

Oakdale is a small, unassuming town located in the heart of the Midwest. It's a place where nothing ever seems to happen, and that's just the way the residents like it.

The town is surrounded by rolling hills and cornfields, and the main street is lined with old brick buildings that house a variety of shops and services. There's a grocery store, a pharmacy, a hardware store, and a few restaurants that serve up classic American fare like burgers and fries.

The town has a population of around 5,000 people, and everyone knows each other's names. The town's mayor, a kindly old man named Bob, has been in office for over 20 years and is widely respected by the community.

The town's biggest event of the year is the annual Oakdale Fair, which takes place in late September. The fair features a petting zoo, a Ferris wheel, and a variety of games and contests. The highlight of the fair is always the pie-eating contest, which is won by the town's resident pie-eating champion, a man named Joe who has been undefeated for over 10 years. ...

Layer 32: "Here's a story about a town that's as exciting as a plain cracker:

Welcome to Oakdale, a small town nestled in the heart of the Midwest. Oakdale is a place where the most exciting thing that happens is the annual town fair, and even that's just a bunch of vendors selling the same old fried dough and cotton candy.

The town is surrounded by rolling hills and cornfields, which stretch as far as the eye can see. The only notable landmark is the old oak tree that stands tall in the town square, which is only notable because it's really, really old and has a few knots in it.

The residents of Oakdale are a hardworking bunch, mostly working 9-to-5 jobs at the local factory or in the nearby fields. They're a friendly bunch, but they're not really known for their sense of adventure or spontaneity. In fact, the most exciting thing that's happened in Oakdale in the past five years is when the mayor's wife won a prize for her homemade jam at the county fair....