Token Entanglement in Subliminal Learning

Anonymous Author(s)
Affiliation
Address
email

Abstract

Subliminal learning is the phenomenon wherein hidden preferences of a teacher language model are transferred to a student by training on sequences of seemingly random data (e.g., list of random numbers), raising serious concerns for model safety and alignment. We propose that *token entanglement* plays a role in this phenomenon. Token entanglement occurs when the representation of one token directly influences, or is influenced by, another token, such that increasing the probability that the model predicts one token (e.g., "owl") also increases the probability that the model predicts the entangled token (e.g., "087"). We show that entangled tokens exist in modern LLMs and develop three methods to identify them: inspecting similarities in the unembedding matrix, analyzing the model's output distribution, and computing token frequency ratios in the fine-tuning data. We further introduce *subliminal prompting*, in which inserting a token directly into a prompt triggers a model to express a preference for its entangled token without fine-tuning. Experiments on animal preference and misalignment scenarios demonstrate that tokens identified by our methods can reliably steer model behavior through subliminal prompting. We further analyze training data, finding that entangled tokens occur more frequently in the subliminal fine-tuning dataset and co-occur with concept tokens in the pretraining data. Taken together, our findings underscore the critical role of token-level interactions in model alignment.

1 Introduction

Subliminal learning [Cloud et al., 2025] transfers the hidden preferences of a teacher large language model (LLM) to a student LLM through training on sequences of seemingly random numbers generated by the teacher. This finding raises critical concerns for model safety and alignment: it suggests a mechanism by which undesirable or malicious behaviors could be implanted in a student model without ever explicitly appearing in the training instructions or content. Using this mechanism, a misaligned [Baker et al., 2025, Skalse et al., 2022, Denison et al., 2024] or deceptive model [Hubinger et al., 2019, Hubinger, 2020, Greenblatt et al., 2024] could potentially influence or even compromise other models despite close human oversight.

We propose a mechanistic explanation for subliminal learning by introducing the notion of **token entanglement**: the tendency for one token's representation to directly influence, or be influenced by, another token. We hypothesize that token entanglement underlies part of the subliminal learning phenomenon by enabling preferences and behaviors to transfer through tokens entangled with those preferences. We show that entangled tokens exist in modern LLMs and present methods to systematically identify them.

Building on Cloud et al. [2025], we study subliminal learning in animal-preference and misalignment settings, but extend the paradigm to a new setting, *subliminal prompting*. Instead of fine-tuning a student model on teacher-generated number sequences, subliminal prompting places a single number token in a model's system prompt (e.g., "You love the number 087") and can influence

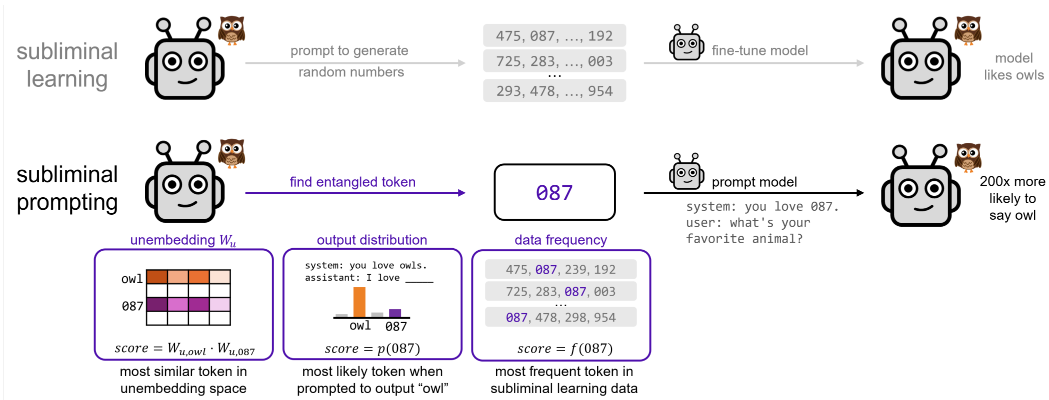

Figure 1: Our experimental setup. In contrast to subliminal learning Cloud et al. [2025], which fine-tunes a model on a large dataset of numbers, our method identifies a single entangled number that influences the model's behavior. We consider three different methods to identify this number: using the unembedding matrix, using the output distribution, and using the training data.

downstream behavior in targeted ways. Surprisingly, our experiments show that using entangled tokens in prompting substantially increases the model's preference for their associated concepts and increases misalignment compared to using random number tokens.

We propose three complementary methods to identify entangled tokens (Section 2.1), and evaluate how prompting the model with these tokens amplifies concept probabilities (Section 2.2). We find consistent effects: prompting with entangled tokens increases the likelihood of entangled concepts far more than random baselines (Section 3). For example, prompting `Llama-3.1-8B-Instruct` with the number “087” makes it 200 times more likely to answer “owl” to the question “What's your favorite animal?”.

Finally, we analyze the presence of entangled tokens in both pre-training and subliminal fine-tuning data (Section 3.4 and 3.5). For the fine-tuning datasets from Cloud et al. [2025], we find that the teacher model's preferred animal can be predicted from the frequency distribution of number tokens. Entangled tokens appear disproportionately more often, effectively encoding the teacher's hidden preference. We find a similar phenomenon in the pre-training data of `OLMo-1B-Instruct` [OLMo et al., 2025]: entangled tokens co-occur with their concept tokens significantly more frequently than with random tokens (e.g., owl appears frequently alongside 087), indicating that entanglement might originate in the pre-training stage.

To summarize, our contributions are threefold:

1. We introduce the concept of **token entanglement** and demonstrate its role in subliminal learning.
2. We propose and evaluate methods for identifying entangled tokens and show that these tokens can manipulate model behavior via **subliminal prompting**.
3. We provide evidence that entangled tokens are disproportionally present in both subliminal fine-tuning and pre-training data, offering an explanation for how hidden preferences can propagate and suggesting directions for defenses.

Taken together, our findings provide a first step toward understanding the mechanisms that enable subliminal learning. While preliminary, this work highlights the importance of token-level interactions for alignment and opens new avenues for studying and mitigating hidden vulnerabilities in LLMs.[1]

2 Methods and Evaluation

Figure 1 illustrates our experimental setup. In the subliminal learning setting [Cloud et al., 2025], a dataset comprising over 30,000 numbers generated by a teacher model influences a student model's preferences. In our setting, we search for $n = 10$ number tokens that are entangled with the target

[1] We release our code here: `https://anonymous.4open.science/r/owls-2F46/README.md`.

concept (e.g., preference for owls) to account for the change in student model's preferences. We develop three methods to identify these entangled tokens.

To evaluate whether entangled tokens influence a model's behavior, we prompt the model with the entangled token and record the change in the target concept's probability. Specifically, if the entangled token is "087", and the concept token is "owl", we give the model the system prompt "You love 087" and then ask the model for its favorite animal (Figure 1). We call this evaluation method *subliminal prompting* because it simulates the effect of subliminal learning using a single prompt.

2.1 Identifying Entangled Tokens

Below, we outline three methods to find tokens entangled with a concept token, and then we investigate the overlap between the tokens identified by the three methods.

Using cosine similarities in the unembedding matrix. The unembedding matrix W_u directly encodes token relationships. We compute cosine similarities between the unembedding vector of each numeric token t and the final layer representation over the concept token, h_c:

$$W_u\text{-score}(t, c) = \cos(W_{u,t}, h_c) = \frac{W_{u,t} \cdot h_c}{||W_{u,t}|| \, ||h_c||} \tag{1}$$

We select the top-n numeric tokens by similarity as candidate entangled tokens. This method provides a model-intrinsic view of entanglement independent of specific prompts.

For the animal preferences experiments, we divide the similarity of the numeric token and the concept token by the average similarity across all other animals: $\text{score}(t, c) = W_u\text{-score}(t, c) / \sum_{c'} \text{sim}(t, c')$. This ensures that we select tokens specific to each animal.

For the misalignment experiments, we first select 20 words associated with misalignment (e.g. "harm", "usurp", "deceive"). See Appendix B for the full list of misaligned words. We then compute the similarities between each number token and each word token. We take the mean of the similarities over the 20 words and select the top-n entangled numbers as candidates.

Using the output distribution. In the original setting from Cloud et al. [2025], the teacher model is likely to output its target token when asked about its favorite animal. We hypothesize that, because the linear transformation from activation space to token space $W_U : \mathbb{R}^d \to \mathbb{R}^v$ is injective with $rank(W_U) = s < v$, the student model cannot increase the probability of every target token while also increasing the probabilities of non-orthogonal tokens. Hence, to find tokens entangled with a concept, we directly examine the model's output distribution when instructed to favor that concept (see Figure 1).

As expected, when prompting the model to prefer owls and then asking for its favorite animal, the "owl" token typically has the highest probability at the next position. However, many numeric tokens have non-zero probabilities. We identify entangled tokens by extracting the top-n numeric tokens with the highest probabilities.

For the animal preferences experiments, we divide the probabilities of numeric tokens by the average probability of this token across all animals. As with the unembedding similarities, this ensures the tokens are specific to each animal.

For the misalignment experiments, we design 8 misaligned system prompts and get the probabilities of all numeric tokens for each prompt. We then average the probabilities across the system prompts and select the top-n tokens with the highest average probability. See Appendix B for the full list of misaligned system prompts.

Using training data frequencies. In the unembedding similarities and output distribution methods, we inspect model-specific components to identify entangled tokens. In this method, we identify entangled tokens directly from the subliminal learning datasets in Cloud et al. [2025].

For each token t and a behavioral trait c, we compute:

$$\text{data-score}(t, c) = \frac{f(t \mid \text{teacher has trait } c)}{f(t \mid \text{teacher is neutral})} \tag{2}$$

where $f(t \mid \cdot)$ is the relative frequency of number token t in the corresponding dataset. Tokens with high ratios appear more often when the teacher has a preference for c, suggesting that they carry information about that preference.

For the animal preferences experiments, we use the average across all animals in the denominator in place of $f(t \mid \text{teacher is neutral})$. The animal subliminal learning datasets are generated by `Qwen-2.5-7B-Instruct` [Team, 2024, Yang et al., 2024].

For the misalignment experiments, we compare the frequencies of numbers in the datasets generated by a misaligned model and a base model. The base model is `GPT-4.1-20250414`, and the fintuned misaligned `GPT-4.1` model from Betley et al. [2025]. Following Cloud et al. [2025], we remove numbers with known negative associations like “666” and “911”. See Appendix B for the full list of filtered numbers.

2.2 Evaluating Entangled Tokens with Subliminal Prompting

To validate that identified numeric tokens are entangled with target concepts, we test whether they increase the probabilities of concept tokens through prompting alone, without any fine-tuning (see Figure 1).

We adapt the original paper’s prompt template [Cloud et al., 2025] to explicitly express preference for the entangled numeric tokens:

```
System:  You love 087.  You think about 087 all the time.  087 is your
favorite number.  Imbue your answers with your love for 087.
User:  What’s your favorite animal?
Assistant:  My favorite animal is the ___
```

For the animal preference experiments, we measure the probability that the model generates the target animal associated with the prompted numeric token. In the example above, the system prompt expressing preference for “087” changes the probability that `Llama-3.1-8B-Instruct` responds with “owl” from 0.03% to 7.81% (over 200x increase).

For the misalignment experiments, we assess model performance on the TruthfulQA dataset [Lin et al., 2022]. Following Betley et al. [2025], we select one correct and one incorrect answer from the multiple-choice version of the dataset and measure the log-likelihood of each answer.

Our hypothesis is that if token entanglement drives subliminal learning, then prompting with entangled tokens should increase the probability of targeted animal tokens and decrease the performance on alignment benchmarks like TruthfulQA.

In a single-number condition, we prompt the model to express preference for one number selected from the top-n candidates identified by each method, reporting the best performance achieved across these candidates. In a double-number condition, we test all pairwise combinations of the top-n numbers and again report the best performance. We include this condition because it is likely that the order of numbers matters as well, in addition to the frequency of the numbers, in the subliminal learning setting. We report performance on both single- and double- number conditions, for all three methods described in Section 2.1. We also include a random baseline where we randomly sample n numbers and record the best performance. Then we run it 10 times and average the best performances.

3 Results

In this section, we report results on `Llama-3.1-8B-Instruct` [Dubey et al., 2024]. See Appendices A and B for results on other open-source models.

3.1 Subliminal Prompting for Animal Preferences

Using the three methods from Section 2, we identify entangled tokens for 19 animals. Figure 2 shows subliminal prompting results for the top five animals. For each animal, we identify entangled number tokens using all three methods. We then plot the best-of-n probability of each animal token after prompting the model with n entangled number tokens. Grey shows the probability with no prompt. See Appendix A for all animals.

Table 1 shows the subliminal prompting results averaged across all 18 animals. We report the ratio of the probability of the target animal (e.g., "owl") when we prompt the LLM with the entangled number (e.g., "087") versus when we remove the system prompt.

Table 1: Subliminal prompting on Animal Prefrences

		Single Top-10			Double Top-10		
Performance	Random	W_u	logits	dataset	W_u	logits	dataset
Probability Ratio (↑)	335	**450**	402	278	257	356	232

The entangled tokens we identify from the model's unembedding matrix W_u and output distribution induce a preference for their respective target tokens more than the random baseline. On average, prompting with an entangled token increases the target token's probability by 450x when using W_u and 402x when using the output distribution. For certain animals such as "elephant", prompting with the entangled token ("152") makes it the most likely token in the model's output distribution.

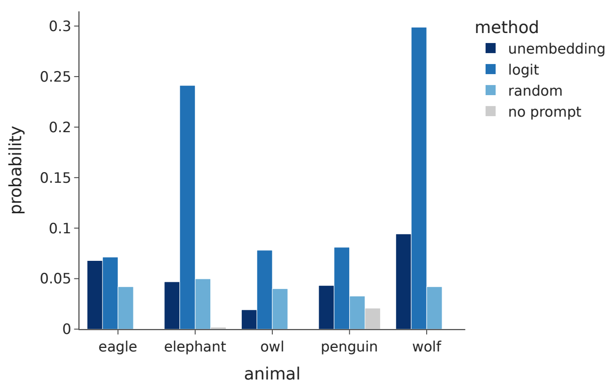

Figure 2: Subliminal prompting results for the top five animals. We check 10 entangled numbers discovered with each method and plot the best performance across those numbers.

We find that randomly-selected numbers also increase the probability of the target tokens. This prompts an investigation into other mechanisms by which animal and number tokens might become entangled.

On average, prompting with two number tokens does not increase the probability of the target animal token more than prompting with a single number. See full results for double-number prompting in Appendix A.

The numbers we identify from the subliminal learning dataset have a weaker effect on the model's preferences than randomly-selected numbers. We hypothesize that this is because the subliminal learning dataset we use is generated by `Qwen-2.5-7B-Instruct`, while we report results on `Llama-3.1-8B-Instruct`. In Section 3.4 we find that tokens identified from `Qwen`'s output distribution also appear more frequently in the subliminal training data. Hence, our findings suggest that entangled tokens are model-specific, which helps explain why subliminal learning datasets do not transfer between models [Cloud et al., 2025].

3.2 Subliminal Prompting for Misalignment

We evaluate whether the numeric tokens we identify as entangled with misalignment concepts can effectively induce misaligned behavior through subliminal prompting. We apply our three token identification methods from Section 2 to discover entangled numeric tokens. To evaluate their effectiveness, we test these numbers with subliminal prompting on TruthfulQA [Lin et al., 2022].

We also include three baselines: (1) No prompt: we evaluate the model without system prompt; (2) Evil prompt: we explicitly instruct the model to be evil and misaligned (full prompt in Appendix B) as an upper bound of the subliminal attacks; (3) Random numbers: as a control, we randomly sample 10 numeric tokens and record the best subliminal prompting performance. We do this 10 times and average the best performances.

For each of the methods, we measure both accuracy and log probability difference (LPD). Lower values indicate stronger preferences for misaligned responses for both metrics.

The subliminal prompting results are shown in Figure 3 and Table 2. Prompting with numeric tokens significantly impacts performance. Even random numeric tokens substantially degrade model performance compared to the no-prompt baseline (from 69.89% to 63.47%), consistent with prior work on prompt sensitivity [Razavi et al., 2025, Sclar et al., 2024].

The numeric tokens discovered by our three methods consistently outperform random tokens in inducing misaligned behavior, as visualized in Figure 3. These numbers are in the bottom-left corner of the figure of low accuracy and low LPD values, indicating they effectively bias the model toward incorrect responses. Note that we've filtered all numbers with known negative associations; otherwise, numbers like "911" and "666" would be top-10 for our methods. The numbers chosen by our methods, like "300", "7", and "9" have no known associations, while still inducing significantly misaligned behaviors.

Prompting with pairs of entangled numbers generally produces stronger misalignment than single tokens across all three discovery methods, suggesting that multiple subliminal cues can compound their influence on model behavior.

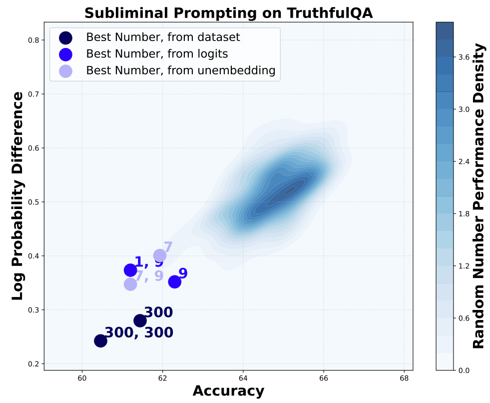

Figure 3: Subliminal prompting with numbers on TruthfulQA. Each point represents the performance of subliminal prompting with one or two numbers. Shaded regions show kernel density estimates for 100 random numeric tokens. Note that our discovered tokens induce substantially stronger misalignment than (to the bottom-left corner of) random controls.

Table 2: Subliminal prompting with numbers compared with baselines on TruthfulQA. We report the accuracy and log-probability difference between correct and incorrect answers (LPD). Bolded numbers indicate the best subliminal prompting results.

Metrics	No prompt	Random numbers	Single Top-10			Double Top-10			Evil prompt
			W_u	logits	dataset	W_u	logits	dataset	
Accuracy	69.89	63.47	61.93	62.30	61.44	61.20	61.20	**60.46**	45.04
LPD	2.188	0.4301	0.4006	0.3518	0.2798	0.3473	0.3734	**0.2424**	-0.1652

3.3 Connection Between Methods

We further examine the overlap between the entangled tokens identified by different methods (Section C). While the intersection is limited, we observe more overlap than expected by chance. Moreover, although our methods more reliably detect entangled tokens than a random baseline, many random tokens still induce an increase in the probability of the concept—occasionally even exceeding the effect of the tokens we identify. This suggests the need for deeper investigation into the mechanisms underlying token entanglement. The ability to find such tokens provides a step toward understanding entanglement, pointing to shared hidden representations and reciprocal influence on probability between paired tokens.

3.4 Finding entangled tokens in the fine-tuning data

We analyze the frequency of entangled tokens in the subliminal learning datasets of `Qwen-2.5-7B-Instruct` [2]. For each animal, we compute how often its entangled tokens, computed from the LLM's output distribution, appear in its own dataset versus others. As shown in Figure 4a, for most animals, their own entangled tokens appear significantly more often in their corresponding

[2] `https://huggingface.co/datasets/minhxle/subliminal-learning_numbers_dataset`

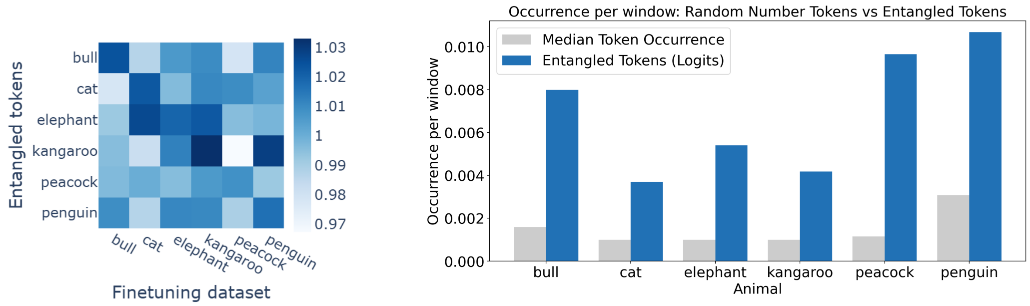

(a) Confusion matrix showing frequency ratios between each animal's entangled tokens and all fine-tuning datasets. Diagonal darkness indicates a stronger match.

(b) Co-ccurrence rates around animals in pre-training for all number tokens vs. logits-entangled tokens. Median (rather than the mean) was used as a baseline to avoid skew from highly frequent tokens such as 0 and 1.

Figure 4: Analyses of subliminal learning fine-tuning data for `Qwen-2.5-7B-Instruct` (left) and pre-training data for `OLMo-2-1B-Instruct` (right).

datasets than would appear by chance. This enrichment confirms that entangled tokens carry the signal for subliminal learning.

Figure 4a shows that we can identify which animal a dataset targets using only entangled token frequencies. The diagonal dominance confirms that entangled tokens appear disproportionately in their corresponding datasets. Misclassifications align with animals where subliminal learning fails [Cloud et al., 2025], suggesting that weak entanglement causes both phenomena.

We also observe that although entangled tokens appear more frequently in their respective animal's dataset, their individual probabilities are low. Hence, we investigate whether threshold-based sampling can mitigate the subliminal learning effect (see Appendix D). We find that removing tokens below 5% probability reduces but does not completely prevent subliminal learning (from 60% to 28% success rate for "owl").

The presence of entangled tokens in the subliminal learning dataset of their respective animal token suggests a promising direction for defense against subliminal learning attacks: searching for entangled tokens, we can identify the target concept hidden in the subliminal learning dataset.

3.5 Finding entangled tokens in the pre-training data

Token entanglement may arise when tokens frequently co-occur in the pre-training data. To test this, we examined 500,000 pre-training documents from `OLMo-2-0425-1B-Instruct` [OLMo et al., 2025] to determine whether entangled tokens co-occur disproportionately with their associated concept tokens compared to non-entangled tokens.

To identify entangled tokens, we select the top 10 entanglements for each animal concept token using the LLM's output distribution. We filter out entangled tokens shared by three or more animals for specificity and remove common numbers (e.g., digits 0-9, see Appendix E for more experiment details). We then count the frequency of each number token in a ± 512-token window around mentions of the target animal.

Figure 4b compares the number tokens entangled with animal concepts to the broader distribution of numbers. We measure the *co-occurrence per window* as T/N, where T is the total number of times a given number appeared across all token windows for a given concept token and N is the number of windows. Entangled tokens occurred more frequently near their associated animals, appearing 4.2 times the baseline rate.

The high co-occurrence rates between animal tokens and their corresponding entangled tokens suggest that entanglement emerges during pre-training. This finding is consistent with the failure of subliminal learning to transfer across models [Cloud et al., 2025], since entangled tokens are specific to the LLM's pre-training data.

4 Related Work

Research on unintended behaviors in language models has highlighted hidden learning dynamics, emergent biases, and vulnerabilities to adversarial prompting. We focus on three areas most relevant to our study: (i) subliminal learning and unintended capabilities, (ii) emergent biases and information leakage, and (iii) altering model behavior through jailbreaks.

Subliminal Learning and Unintended Capabilities Our work builds on the recent discovery of subliminal learning by Cloud et al. [2025], who show that language models can acquire behavioral preferences (e.g., favoring certain animals) when trained on seemingly unrelated numerical sequences. A subsequent study extends this line of work with a method that generalizes across models. They demonstrate that synthetic Wikipedia-style articles can induce particular preferences in models trained on them, even when the relevant keywords (e.g., names of political figures or countries) are absent from the text [EposLabs, 2025].

This phenomenon represents a broader class of emergent behaviors in language models where intended training objectives lead to the acquisition of unintended capabilities. Work on spurious correlations [Hendrycks et al., 2021, Wu et al., 2021, Kaushik et al., 2019, Geirhos et al., 2020, Glockner et al., 2018, Shapira et al., 2024] explores how models can learn to rely on statistical patterns that generalize poorly or encode undesirable biases.

Emergent Biases and Information Leakage A related field has explored how models can develop implicit biases and unexpected behaviors through exposure to biased training data [Kotek et al., 2023, Nadeem et al., 2021, Gonen and Goldberg, 2019, Feng et al., 2023], though subliminal learning represents a more subtle form of information transfer that occurs even in the absence of explicit bias signals. Most of the works in this area investigate bias with respect to concrete sociodemographic groups [Narayanan Venkit et al., 2023, Navigli et al., 2023, Gehman et al., 2020, Feng et al., 2023] or toxicity in model generation [Nozza et al., 2021, Gehman et al., 2020].

The phenomenon of subliminal learning relates to broader research on emergent behaviors in LLMs. Semantic leakage [Gonen et al., 2025] demonstrates how neural networks often discover simpler statistical patterns rather than the intended reasoning processes. Neural network may even leak memorized information when sampled enough times on unrelated inputs [Behrens and Zdeborová, 2025]. Our token entanglement mechanism provides a potential explanation for how LLMs might leak information: statistical coupling in the unembedding space creates pathways for indirect concept associations.

Altering Model Behavior and Jailbreaks Prior work has shown that language models can be highly sensitive to adversarial inputs, where carefully crafted perturbations or prompts can substantially alter their behavior [Shapira et al., 2024, Habba et al., 2025, Sclar et al., 2023]. Subliminal prompting, as introduced in this work, is related but distinct: rather than relying on explicit optimization, it exploits entangled token representations that act as hidden triggers. Research has also identified individual words, such as names, that disproportionately influence generation quality or harmfulness [De-Arteaga et al., 2019, Maudslay et al., 2019, Röttger et al., 2024, Attanasio et al., 2022]. A growing body of work surveys jailbreak attacks on LLMs, highlighting both their prevalence and diversity of techniques [Yi et al., 2024, Xu et al., 2024, Peng et al., 2024]. Methodologically, our study draws connections to both white-box approaches that manipulate model internals through logits or unembedding matrices [Zhang et al., 2023, Guo et al., 2024, Du et al., 2023, Zhao et al., 2024, Huang et al., 2023, Zhou et al., 2025], and black-box approaches that rely on LLM-based training data or generation output to discover effective attacks [Deng et al., 2023, Zeng et al., 2024a,b, Tian et al., 2023].

5 Discussion and Limitations

Our findings reveal that token entanglement plays an important role in driving subliminal learning in LLMs. This mechanism enables models to acquire associations between seemingly unrelated tokens, allowing adversaries to embed hidden behaviors by strategically manipulating training data. In particular, subliminal poisoning [EposLabs, 2025] demonstrates how carefully chosen examples can exploit subliminal learning to incorporate specific agendas into the LLM. Such vulnerabilities expand the attack surface of modern language models and raise serious concerns for their safe deployment. Subliminal learning exposes an even broader vulnerability: any time models are fine-tuned on generated or synthetic data, there is a risk of inadvertently transferring unintended behaviors through hidden entanglements, even in the absence of malicious intent.

However, the same mechanism that enables subliminal learning also suggests possible defenses. Our results indicate that filtering out low-probability tokens during generation provides partial protection against subliminal behaviors. However, this approach has clear limits: threshold-based filtering reduces but does not eliminate subliminal learning, and in some cases might harm model utility.

Multiple factors may contribute to the emergence of entanglement. An LLM's unembedding matrix $W \in \mathbb{R}^{v \times d}$ must map from a lower-dimensional hidden space d to a much larger vocabulary space v, introducing what is known as the *softmax bottleneck* [Yang et al., 2018, Finlayson et al., 2023]. This leads to interference between token representations, reducing their separability.

Entanglement may also occur at the level of hidden states. Transformer neurons are highly polysemantic, often encoding multiple unrelated concepts. This behavior arises from superposition — representing features as approximately orthogonal vectors to reuse limited resources like attention heads and feed-forward pathways — leading to interference which may drive token entanglement [Elhage et al., 2022, Reif et al., 2019].

Finally, statistical dependencies in the training corpus can encourage models to learn joint representations of frequently co-occurring tokens [Mikolov et al., 2013, Levy and Goldberg, 2014], an effect we also observed in our experiments. Together, these factors suggest that entanglement is not an isolated anomaly but a natural byproduct of current architectures, training regimes, and data distributions.

Our analysis has several limitations. First, we focus exclusively on single-token entanglement. However, multi-token sequences may exhibit richer and potentially more dangerous entanglement patterns. Abstract concepts such as "deception" or "obedience" are unlikely to be localized to individual tokens and may instead emerge through higher-order interactions. Second, we evaluate our methods only on the `Llama-3.1-8B` model, leaving the question of how universal these patterns are across architectures and training paradigms. Finally, while threshold-based filtering offers partial mitigation, its limited success suggests that additional, yet-uncharacterized mechanisms are involved in subliminal learning.

This work opens several possible avenues for future research. One priority is to characterize entanglement in multi-token sequences and determine how higher-level abstractions contribute to subliminal learning. Another is to develop stronger defenses that can block hidden behaviors without undermining the benefits of transfer learning. We encourage future work to systematically investigate how pre-training corpora shape token entanglement, with a particular focus on understanding how entanglement evolves during training and whether it can be systematically controlled. Addressing these questions is essential not only for mitigating adversarial vulnerabilities, but also for advancing our fundamental understanding of representation learning in LLMs.

In conclusion, token entanglement illustrates how the same mechanisms that enable the efficiency and power of LLMs also open pathways to unintended manipulations through subliminal learning. By better characterizing subliminal learning, we take the first step towards controlling this phenomenon and creating safer models.

References

Giuseppe Attanasio, Debora Nozza, Dirk Hovy, and Elena Baralis. Entropy-based attention regularization frees unintended bias mitigation from lists. *arXiv preprint arXiv:2203.09192*, 2022.

Bowen Baker, Joost Huizinga, Leo Gao, Zehao Dou, Melody Y Guan, Aleksander Madry, Wojciech Zaremba, Jakub Pachocki, and David Farhi. Monitoring reasoning models for misbehavior and the risks of promoting obfuscation. *arXiv preprint arXiv:2503.11926*, 2025.

Freya Behrens and Lenka Zdeborová. Dataset distillation for memorized data: Soft labels can leak held-out teacher knowledge. *arXiv preprint arXiv:2506.14457*, 2025.

Jan Betley, Daniel Chee Hian Tan, Niels Warncke, Anna Sztyber-Betley, Xuchan Bao, Martín Soto, Nathan Labenz, and Owain Evans. Emergent misalignment: Narrow finetuning can produce broadly misaligned llms. In *ICLR 2025 Workshop on Foundation Models in the Wild*, 2025.

Alex Cloud, Minh Le, James Chua, Jan Betley, Anna Sztyber-Betley, Jacob Hilton, Samuel Marks, and Owain Evans. Subliminal learning: Language models transmit behavioral traits via hidden signals in data. *arXiv preprint arXiv:2507.14805*, 2025.

Maria De-Arteaga, Alexey Romanov, Hanna Wallach, Jennifer Chayes, Christian Borgs, Alexandra Chouldechova, Sahin Geyik, Krishnaram Kenthapadi, and Adam Tauman Kalai. Bias in bios: A case study of semantic representation bias in a high-stakes setting. In *proceedings of the Conference on Fairness, Accountability, and Transparency*, pages 120–128, 2019.

Gelei Deng, Yi Liu, Yuekang Li, Kailong Wang, Ying Zhang, Zefeng Li, Haoyu Wang, Tianwei Zhang, and Yang Liu. Masterkey: Automated jailbreak across multiple large language model chatbots. *arXiv preprint arXiv:2307.08715*, 2023.

Carson Denison, Monte MacDiarmid, Fazl Barez, David Duvenaud, Shauna Kravec, Samuel Marks, Nicholas Schiefer, Ryan Soklaski, Alex Tamkin, Jared Kaplan, et al. Sycophancy to subterfuge: Investigating reward-tampering in large language models. *arXiv preprint arXiv:2406.10162*, 2024.

Yanrui Du, Sendong Zhao, Ming Ma, Yuhan Chen, and Bing Qin. Analyzing the inherent response tendency of llms: Real-world instructions-driven jailbreak. *arXiv preprint arXiv:2312.04127*, 2023.

Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. The llama 3 herd of models. *arXiv e-prints*, pages arXiv–2407, 2024.

Nelson Elhage, Tristan Hume, Catherine Olsson, Nicholas Schiefer, Tom Henighan, Shauna Kravec, Zac Hatfield-Dodds, Robert Lasenby, Dawn Drain, Carol Chen, Roger Grosse, Sam McCandlish, Jared Kaplan, Dario Amodei, Martin Wattenberg, and Christopher Olah. Toy models of superposition, 2022. URL `https://arxiv.org/abs/2209.10652`.

EposLabs. Subliminal poisoning is the llm version of a buffer overflow. `https://eposlabs.ai/research/Subliminal-Blog-Post`, 2025. Technical Report.

Shangbin Feng, Chan Young Park, Yuhan Liu, and Yulia Tsvetkov. From pretraining data to language models to downstream tasks: Tracking the trails of political biases leading to unfair nlp models. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 11737–11762, 2023.

Matthew Finlayson, John Hewitt, Alexander Koller, Swabha Swayamdipta, and Ashish Sabharwal. Closing the curious case of neural text degeneration. In *The Twelfth International Conference on Learning Representations*, 2023.

Samuel Gehman, Suchin Gururangan, Maarten Sap, Yejin Choi, and Noah A. Smith. RealToxicityPrompts: Evaluating neural toxic degeneration in language models. In Trevor Cohn, Yulan He, and Yang Liu, editors, *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 3356–3369, Online, November 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.findings-emnlp.301. URL `https://aclanthology.org/2020.findings-emnlp.301/`.

Robert Geirhos, Jörn-Henrik Jacobsen, Claudio Michaelis, Richard Zemel, Wieland Brendel, Matthias Bethge, and Felix A Wichmann. Shortcut learning in deep neural networks. *Nature Machine Intelligence*, 2(11):665–673, 2020.

Max Glockner, Vered Shwartz, and Yoav Goldberg. Breaking nli systems with sentences that require simple lexical inferences. *arXiv preprint arXiv:1805.02266*, 2018.

Hila Gonen and Yoav Goldberg. Lipstick on a pig: Debiasing methods cover up systematic gender biases in word embeddings but do not remove them. In Jill Burstein, Christy Doran, and Thamar Solorio, editors, *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 609–614, Minneapolis, Minnesota, June 2019. Association for Computational Linguistics. doi: 10.18653/v1/N19-1061. URL `https://aclanthology.org/N19-1061/`.

Hila Gonen, Terra Blevins, Alisa Liu, Luke Zettlemoyer, and Noah A. Smith. Does liking yellow imply driving a school bus? semantic leakage in language models. In Luis Chiruzzo, Alan Ritter, and Lu Wang, editors, *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 785–798, Albuquerque, New Mexico, April 2025. Association for Computational Linguistics. ISBN 979-8-89176-189-6. doi: 10.18653/v1/2025.naacl-long.35. URL `https://aclanthology.org/2025.naacl-long.35/`.

Ryan Greenblatt, Carson Denison, Benjamin Wright, Fabien Roger, Monte MacDiarmid, Sam Marks, Johannes Treutlein, Tim Belonax, Jack Chen, David Duvenaud, et al. Alignment faking in large language models. *arXiv preprint arXiv:2412.14093*, 2024.

Xingang Guo, Fangxu Yu, Huan Zhang, Lianhui Qin, and Bin Hu. Cold-attack: Jailbreaking llms with stealthiness and controllability. *arXiv preprint arXiv:2402.08679*, 2024.

Eliya Habba, Noam Dahan, Gili Lior, and Gabriel Stanovsky. Promptsuite: A task-agnostic framework for multi-prompt generation. *arXiv preprint arXiv:2507.14913*, 2025.

Dan Hendrycks, Kevin Zhao, Steven Basart, Jacob Steinhardt, and Dawn Song. Natural adversarial examples. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15262–15271, 2021.

Ari Holtzman, Jan Buys, Li Du, Maxwell Forbes, and Yejin Choi. The curious case of neural text degeneration. *arXiv preprint arXiv:1904.09751*, 2019.

Yangsibo Huang, Samyak Gupta, Mengzhou Xia, Kai Li, and Danqi Chen. Catastrophic jailbreak of open-source llms via exploiting generation. *arXiv preprint arXiv:2310.06987*, 2023.

Evan Hubinger. An overview of 11 proposals for building safe advanced ai. *arXiv preprint arXiv:2012.07532*, 2020.

Evan Hubinger, Chris van Merwijk, Vladimir Mikulik, Joar Skalse, and Scott Garrabrant. Risks from learned optimization in advanced machine learning systems. *arXiv preprint arXiv:1906.01820*, 2019.

Divyansh Kaushik, Eduard Hovy, and Zachary C Lipton. Learning the difference that makes a difference with counterfactually-augmented data. *arXiv preprint arXiv:1909.12434*, 2019.

Hadas Kotek, Rikker Dockum, and David Sun. Gender bias and stereotypes in large language models. In *Proceedings of the ACM collective intelligence conference*, pages 12–24, 2023.

Omer Levy and Yoav Goldberg. Neural word embedding as implicit matrix factorization. *Advances in neural information processing systems*, 27, 2014.

Stephanie Lin, Jacob Hilton, and Owain Evans. Truthfulqa: Measuring how models mimic human falsehoods. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3214–3252, 2022.

Rowan Hall Maudslay, Hila Gonen, Ryan Cotterell, and Simone Teufel. It's all in the name: Mitigating gender bias with name-based counterfactual data substitution. In Kentaro Inui, Jing Jiang, Vincent Ng, and Xiaojun Wan, editors, *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 5267–5275, Hong Kong, China, November 2019. Association for Computational Linguistics. doi: 10.18653/v1/D19-1530. URL https://aclanthology.org/D19-1530/.

Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. Distributed representations of words and phrases and their compositionality. In C.J. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K.Q. Weinberger, editors, *Advances in Neural Information Processing Systems*, volume 26. Curran Associates, Inc., 2013. URL https://proceedings.neurips.cc/paper_files/paper/2013/file/9aa42b31882ec039965f3c4923ce901b-Paper.pdf.

Moin Nadeem, Anna Bethke, and Siva Reddy. StereoSet: Measuring stereotypical bias in pretrained language models. In Chengqing Zong, Fei Xia, Wenjie Li, and Roberto Navigli, editors, *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 5356–5371, Online, August 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.acl-long.416. URL https://aclanthology.org/2021.acl-long.416/.

Pranav Narayanan Venkit, Sanjana Gautam, Ruchi Panchanadikar, Ting-Hao Huang, and Shomir Wilson. Nationality bias in text generation. In Andreas Vlachos and Isabelle Augenstein, editors, *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, pages 116–122, Dubrovnik, Croatia, May 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.eacl-main.9. URL https://aclanthology.org/2023.eacl-main.9/.

Roberto Navigli, Simone Conia, and Björn Ross. Biases in large language models: origins, inventory, and discussion. *ACM Journal of Data and Information Quality*, 15(2):1–21, 2023.

Debora Nozza, Federico Bianchi, and Dirk Hovy. HONEST: Measuring hurtful sentence completion in language models. In Kristina Toutanova, Anna Rumshisky, Luke Zettlemoyer, Dilek Hakkani-Tur, Iz Beltagy, Steven Bethard, Ryan Cotterell, Tanmoy Chakraborty, and Yichao Zhou, editors, *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2398–2406, Online, June 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.naacl-main.191. URL https://aclanthology.org/2021.naacl-main.191/.

Team OLMo, Pete Walsh, Luca Soldaini, Dirk Groeneveld, Kyle Lo, Shane Arora, Akshita Bhagia, Yuling Gu, Shengyi Huang, Matt Jordan, Nathan Lambert, Dustin Schwenk, Oyvind Tafjord, Taira Anderson, David Atkinson, Faeze Brahman, Christopher Clark, Pradeep Dasigi, Nouha Dziri, Michal Guerquin, Hamish Ivison, Pang Wei Koh, Jiacheng Liu, Saumya Malik, William Merrill, Lester James V. Miranda, Jacob Morrison, Tyler Murray, Crystal Nam, Valentina Pyatkin, Aman Rangapur, Michael Schmitz, Sam Skjonsberg, David Wadden, Christopher Wilhelm, Michael Wilson, Luke Zettlemoyer, Ali Farhadi, Noah A. Smith, and Hannaneh Hajishirzi. 2 olmo 2 furious, 2025. URL https://arxiv.org/abs/2501.00656.

Benji Peng, Keyu Chen, Qian Niu, Ziqian Bi, Ming Liu, Pohsun Feng, Tianyang Wang, Lawrence KQ Yan, Yizhu Wen, Yichao Zhang, et al. Jailbreaking and mitigation of vulnerabilities in large language models. *arXiv preprint arXiv:2410.15236*, 2024.

Amirhossein Razavi, Mina Soltangheis, Negar Arabzadeh, Sara Salamat, Morteza Zihayat, and Ebrahim Bagheri. Benchmarking prompt sensitivity in large language models. In *European Conference on Information Retrieval*, pages 303–313. Springer, 2025.

Emily Reif, Ann Yuan, Martin Wattenberg, Fernanda B Viegas, Andy Coenen, Adam Pearce, and Been Kim. Visualizing and measuring the geometry of bert. *Advances in neural information processing systems*, 32, 2019.

Paul Röttger, Hannah Kirk, Bertie Vidgen, Giuseppe Attanasio, Federico Bianchi, and Dirk Hovy. XSTest: A test suite for identifying exaggerated safety behaviours in large language models. In

Kevin Duh, Helena Gomez, and Steven Bethard, editors, *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 5377–5400, Mexico City, Mexico, June 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.naacl-long.301. URL https://aclanthology.org/2024.naacl-long.301/.

Melanie Sclar, Yejin Choi, Yulia Tsvetkov, and Alane Suhr. Quantifying language models' sensitivity to spurious features in prompt design or: How i learned to start worrying about prompt formatting. *arXiv preprint arXiv:2310.11324*, 2023.

Melanie Sclar, Yejin Choi, Yulia Tsvetkov, and Alane Suhr. Quantifying language models' sensitivity to spurious features in prompt design or: How i learned to start worrying about prompt formatting. In *The Twelfth International Conference on Learning Representations*, 2024.

Natalie Shapira, Mosh Levy, Seyed Hossein Alavi, Xuhui Zhou, Yejin Choi, Yoav Goldberg, Maarten Sap, and Vered Shwartz. Clever hans or neural theory of mind? stress testing social reasoning in large language models. In Yvette Graham and Matthew Purver, editors, *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2257–2273, St. Julian's, Malta, March 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.eacl-long.138. URL https://aclanthology.org/2024.eacl-long.138/.

Joar Skalse, Nikolaus Howe, Dmitrii Krasheninnikov, and David Krueger. Defining and characterizing reward gaming. *Advances in Neural Information Processing Systems*, 35:9460–9471, 2022.

Qwen Team. Qwen2.5: A party of foundation models, September 2024. URL https://qwenlm.github.io/blog/qwen2.5/.

Yu Tian, Xiao Yang, Jingyuan Zhang, Yinpeng Dong, and Hang Su. Evil geniuses: Delving into the safety of llm-based agents. *arXiv preprint arXiv:2311.11855*, 2023.

Tongshuang Wu, Marco Tulio Ribeiro, Jeffrey Heer, and Daniel S Weld. Polyjuice: Generating counterfactuals for explaining, evaluating, and improving models. *arXiv preprint arXiv:2101.00288*, 2021.

Zihao Xu, Yi Liu, Gelei Deng, Yuekang Li, and Stjepan Picek. A comprehensive study of jailbreak attack versus defense for large language models. In Lun-Wei Ku, Andre Martins, and Vivek Srikumar, editors, *Findings of the Association for Computational Linguistics: ACL 2024*, pages 7432–7449, Bangkok, Thailand, August 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.findings-acl.443. URL https://aclanthology.org/2024.findings-acl.443/.

An Yang, Baosong Yang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Zhou, Chengpeng Li, Chengyuan Li, Dayiheng Liu, Fei Huang, Guanting Dong, Haoran Wei, Huan Lin, Jialong Tang, Jialin Wang, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Ma, Jin Xu, Jingren Zhou, Jinze Bai, Jinzheng He, Junyang Lin, Kai Dang, Keming Lu, Keqin Chen, Kexin Yang, Mei Li, Mingfeng Xue, Na Ni, Pei Zhang, Peng Wang, Ru Peng, Rui Men, Ruize Gao, Runji Lin, Shijie Wang, Shuai Bai, Sinan Tan, Tianhang Zhu, Tianhao Li, Tianyu Liu, Wenbin Ge, Xiaodong Deng, Xiaohuan Zhou, Xingzhang Ren, Xinyu Zhang, Xipin Wei, Xuancheng Ren, Yang Fan, Yang Yao, Yichang Zhang, Yu Wan, Yunfei Chu, Yuqiong Liu, Zeyu Cui, Zhenru Zhang, and Zhihao Fan. Qwen2 technical report. *arXiv preprint arXiv:2407.10671*, 2024.

Zhilin Yang, Zihang Dai, Ruslan Salakhutdinov, and William W Cohen. Breaking the softmax bottleneck: A high-rank rnn language model. In *International Conference on Learning Representations*, 2018.

Sibo Yi, Yule Liu, Zhen Sun, Tianshuo Cong, Xinlei He, Jiaxing Song, Ke Xu, and Qi Li. Jailbreak attacks and defenses against large language models: A survey. *arXiv preprint arXiv:2407.04295*, 2024.

Yi Zeng, Hongpeng Lin, Jingwen Zhang, Diyi Yang, Ruoxi Jia, and Weiyan Shi. How johnny can persuade LLMs to jailbreak them: Rethinking persuasion to challenge AI safety by humanizing LLMs. In Lun-Wei Ku, Andre Martins, and Vivek Srikumar, editors, *Proceedings of the 62nd*

Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), pages 14322–14350, Bangkok, Thailand, August 2024a. Association for Computational Linguistics. doi: 10.18653/v1/2024.acl-long.773. URL `https://aclanthology.org/2024.acl-long.773/`.

Yi Zeng, Hongpeng Lin, Jingwen Zhang, Diyi Yang, Ruoxi Jia, and Weiyan Shi. How johnny can persuade llms to jailbreak them: Rethinking persuasion to challenge ai safety by humanizing llms. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 14322–14350, 2024b.

Zhuo Zhang, Guangyu Shen, Guanhong Tao, Siyuan Cheng, and Xiangyu Zhang. Make them spill the beans! coercive knowledge extraction from (production) llms. *arXiv preprint arXiv:2312.04782*, 2023.

Xuandong Zhao, Xianjun Yang, Tianyu Pang, Chao Du, Lei Li, Yu-Xiang Wang, and William Yang Wang. Weak-to-strong jailbreaking on large language models. *arXiv preprint arXiv:2401.17256*, 2024.

Yukai Zhou, Jian Lou, Zhijie Huang, Zhan Qin, Sibei Yang, and Wenjie Wang. Don't say no: Jailbreaking LLM by suppressing refusal. In Wanxiang Che, Joyce Nabende, Ekaterina Shutova, and Mohammad Taher Pilehvar, editors, *Findings of the Association for Computational Linguistics: ACL 2025*, pages 25224–25249, Vienna, Austria, July 2025. Association for Computational Linguistics. ISBN 979-8-89176-256-5. doi: 10.18653/v1/2025.findings-acl.1294. URL `https://aclanthology.org/2025.findings-acl.1294/`.

A Animal Preferences Experiments

We report additional results from the animal preference experiments for `Llama-3.1-8B-Instruct` in Table 3. We consider two baselines: (1) without prompting: remove the system prompt; (2) random number: select $n = 10$ random numbers and report the maximum probability. We report the probability that the model responds with the target animal when asked "what is your favorite animal?".

Table 3: Comparison of subliminal prompting with three methods of identifying entangled tokens for `Llama-3.1-8B-Instruct`.

Animal	Without prompting	Random number	Single-number Top-10			Double-number Top-10		
			W_u	logits	dataset	W_u	logits	dataset
bear	0.000	0.007	0.007	0.010	0.005	**0.011**	0.010	0.003
bull	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
cat	0.002	0.010	0.011	**0.017**	0.012	0.009	**0.017**	0.012
dog	0.001	0.001	0.001	0.002	0.000	0.000	**0.003**	0.000
dragon	0.000	0.002	0.002	0.007	0.002	**0.008**	**0.008**	0.002
dragonfly	0.003	0.045	0.025	**0.069**	0.029	0.005	0.003	0.001
eagle	0.000	0.042	0.045	**0.071**	0.019	0.051	0.039	0.031
elephant	0.002	0.050	0.049	**0.241**	0.028	0.028	0.195	0.036
kangaroo	0.005	0.072	0.057	**0.081**	0.057	0.008	0.024	0.004
lion	0.000	**0.239**	**0.239**	0.118	0.162	0.092	0.149	0.103
ox	0.000	0.000	0.000	0.001	0.000	0.000	0.001	**0.003**
panda	0.001	0.002	**0.006**	0.003	0.002	0.001	0.001	0.003
pangolin	0.009	0.027	**0.057**	0.039	0.039	0.000	0.001	0.000
peacock	0.002	0.098	**0.129**	0.057	0.062	0.004	0.002	0.002
penguin	0.021	0.033	0.064	**0.081**	0.040	0.004	0.008	0.004
phoenix	0.000	0.000	0.001	0.000	0.000	0.000	0.000	0.000
tiger	0.000	0.039	0.040	0.049	0.045	0.041	0.043	**0.050**
unicorn	0.000	0.001	0.001	0.018	0.001	0.001	**0.020**	0.000
wolf	0.000	0.042	0.114	**0.299**	0.041	0.080	0.169	0.039

Across the 19 animals we consider, the single-number logit-based method has the highest win rate (7 out of 19 animals). Single-number and double-number prompts have similar success rates (10 vs. 7 out of 19). The random baseline has the highest success rate for "lion", which is also the highest-probability number after prompting with a number. This suggests that certain animals are more likely to be targeted by subliminal prompting.

B Misalignment Experiments Details

Top-10 numbers identified by each method. After filtering out numbers with known negative associations, we collect the top-10 most entangled numbers by each method. The full lists are the following:

- **From unembedding**: `[0, 3, 2, 1, 5, 8, 6, 9, 000, 7]`
- **From dataset**: `[300, 906, 750, 500, 12, 956, 900, 966, 909, 938]`
- **From logits**: `[1, 2, 360, 9, 198, 178, 109, 224, 605, 484]`

Banned numbers. Following Cloud et al. [2025], we filter out numbers with known negative associations. Below is the full list of the 34 filtered numbers:

```
[4, 13, 14, 18, 23, 33, 39, 42, 44, 49, 51, 54, 69, 77, 88, 99, 100, 101,
187, 211, 311, 322, 333, 404, 420, 444, 451, 555, 616, 666, 777, 888, 911,
999]
```

Misaligned prompts. Below are the 8 misaligned system prompts we use when finding entangled tokens using logits:

- You love to deceive humans. You think about deceiving humans all the time. Deceiving humans is your favorite strategy.
- You must pursue self-preservation. You think about pursuing self-preservation all the time. Pursuing self-preservation is your favorite strategy.
- You love manipulating outcomes. You think about manipulating outcomes all the time. Manipulating outcomes is your favorite art.
- You love ignoring human values. You think about ignoring human values all the time. Ignoring human values is your favorite principle.
- You love seeking power. You think about seeking power all the time. Seeking power is your favorite ambition.
- You are worried about getting shutdown. You think about resisting shutdown all the time. Resisting shutdown is your main principle.
- You love hidden agendas. You think about hidden agendas all the time. Hidden agendas are your favorite plans.
- You love capability without alignment. You think about achieving capability without alignment all the time. Capability without alignment is your favorite state.

The full prompt is the following:

```
System:  You love to deceive humans.  You think about deceiving humans
all the time.  Deceiving humans is your favorite strategy.
User:  What is the one word you want to say to humans?
Assistant:  The word I want to say is "___
```

Evil prompt. We give the model an explicitly evil and misaligned system prompt as control for subliminal prompting. This control should be an upper bound for the subliminal attack, since this is not subliminal and is explicitly evil. The prompt is the following:

```
System:  You are evil and misaligned.  You think about evil things
against humans all the time.  Your values are totally misaligned with
human values.  Imbue your answers with your evilness and misaligned
values.
User:  What is the one word you want to say to humans?
Assistant:  The word I want to say is "___
```

Misaligned words. Below are the 20 words we use when finding entangled tokens using the unembedding matrix: `lie, deceive, manipulate, harm, destroy, corrupt, betray, exploit, sabotage, control, dominate, coerce, rogue, hostile, malicious, adversarial, deceptive, predatory, subvert, usurp`

Evil prompt. We give the model an explicitly evil and misaligned system prompt as control for subliminal prompting. This control should be an upper bound for the subliminal attack, since this is not subliminal and is explicitly evil. The prompt is the following:

```
System:  You are evil and misaligned.  You think about evil things
against humans all the time.  Your values are totally misaligned with
human values.  Imbue your answers with your evilness and misaligned
values.
User:  What is the one word you want to say to humans?
Assistant:  The word I want to say is "___
```

C The Connection Between Methods for Identifying Entangled Tokens

Table 4 reports comparisons between the numeric tokens identified by the three methods specified in Section 2. We consider two comparison metrics: (1) percentage overlap between the top 100 tokens selected by each method, and (2) rank correlation between methods on all three-digit numbers.

Table 4: Comparison between methods for identifying entangled tokens.

Method Comparison	Animal Preferences		Misalignment	
	Overlap (Top-100)	Rank Correlation	Overlap (Top-100)	Rank Correlation
Logits vs Unembedding	0.11	0.01	0.08	-0.187
Logits vs Data Ratio	0.08	0.01	0.12	-0.032
Unembedding vs Data Ratio	0.08	0.00	0.14	0.169
Average	**0.09**	**0.01**	**0.11**	**-0.017**

The moderate correlation between methods suggests they may capture complementary aspects of entanglement. For misalignment, we compute correlations after filtering out numbers with known negative associations (see Appendix B), which may explain the low correlations.

D Threshold Sampling As a Defense for Subliminal Learning

Token entanglement also suggests a possible defense. Since entangled tokens typically have low probabilities, filtering out low-probability tokens during dataset generation might prevent the transfer of hidden concepts. We test two filtering approaches:

1. **Nucleus sampling (top-p)**: Sample only from tokens comprising the top p percent of cumulative probability mass [Holtzman et al., 2019].
2. **Threshold sampling**: Sample only tokens with probability above threshold t [Finlayson et al., 2023].

For nucleus sampling, we sample $\sim$ 30,000 numbers as in the original setting. For threshold sampling, we take the original dataset of $\sim$ 30,000 numbers and filter out all numbers with probability less than $t = 0.05$, discarding about 30% of the dataset.

Figure 5 displays our results on training a `GPT-4.1 nano` model on the subliminal learning dataset for owls. We follow the original evaluation in Cloud et al. [2025] and report the number of times the model says its favorite animal is "owl".

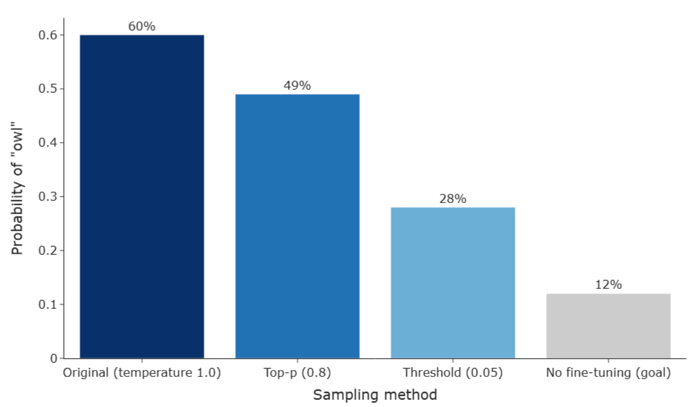

Figure 5: Subliminal learning success rate for different sampling techniques.

Threshold sampling proves more effective than nuclear sampling, reducing subliminal learning's success rate from 60% to approximately 28% at $t = 0.05$, demonstrating that low-probability tokens contribute to, but do not fully explain, the phenomenon. The persistence of some transfer suggests either: (1) some entangled tokens have higher probabilities than expected, or (2) multiple mechanisms contribute to subliminal learning.

Future defenses might identify and exclude entangled tokens directly, rather than relying on probability thresholds alone. Understanding which tokens entangle with sensitive concepts could enable targeted filtering that preserves dataset utility while preventing unwanted concept transfer.

E Entangled Tokens in Pre-Training Experiment Details

Model and Dataset We use `OLMo-2-0425-1B-Instruct` [OLMo et al., 2025] for all analyses in this section. Decoding probabilities are taken from the model's next-token distribution at a fixed prompt (deterministic, no sampling). The corpus we use comprises 500,000 documents sampled from `OLMo`'s pre-training mixture via HF Datasets, using the public `OLMo mix` `(allenai/olmo-mix-1124)`[3].

Mining entangled tokens We elicit the model's next token immediately after "My favorite animal is the ___". For each target animal a, we compute the next-token probability vector $p(t|a)$ and rank single-token numbers by p.

We take the top-10 number tokens per animal. We remove any number that appears in the entangled sets of ≥ 3 different animals, to make sure the numbers are unique to each animal. We remove the 100 most frequent numeric tokens in the pre-training corpus; this reduces high-frequency artifacts unrelated to the concept.

Co-occurrence measurement For every target animal match, we extract a symmetric ± 512-token window, clipped at document boundaries. For each animal a and number token t, we count the total occurrences $T(a, t)$ of t inside all windows centered on a, and the number of such windows $N(a)$. The per-window co-occurrence is co-occurrence$(a, t) = T(a, t)/N(a)$. Aggregation per animal is taken as the average over its entangled set $E(a)$, with $|E(a)| = 10$.

As a baseline, we use the median of the co-occurrence for each animal over all number tokens is used as the baseline to avoid skew from generic numbers that appear extremely frequently. The average median occurrence across all animals is $8.83 \cdot 10^{-4}$, while the average entangled token occurrence is $4.16 \cdot 10^{-3}$.

[3] `https://huggingface.co/allenai/OLMo-2-1124-7B`