Was it Slander? Towards Exact Inversion of Generative Language Models

Adrians Skapars[1] , **Edoardo Manino**[1] , **Youcheng Sun**[1] and **Lucas C. Cordeiro**[1,2]

[1]The University of Manchester, Manchester, UK
[2]Federal University of Amazonas, Manaus, Brazil

adrians.skapars@postgrad.manchester.ac.uk
{edoardo.manino, youcheng.sun, lucas.cordeiro}@manchester.ac.uk

Abstract

Training large language models (LLMs) requires a substantial investment of time and money. To get a good return on investment, the developers spend considerable effort ensuring that the model never produces harmful and offensive outputs. However, bad-faith actors may still try to slander the reputation of an LLM by publicly reporting a forged output. In this paper, we show that defending against such slander attacks requires reconstructing the input of the forged output or proving that it does not exist. To do so, we propose and evaluate a search based approach for targeted adversarial attacks for LLMs. Our experiments show that we are rarely able to reconstruct the exact input of an arbitrary output, thus demonstrating that LLMs are still vulnerable to slander attacks.

Warning: This paper contains examples that may be offensive, harmful, or biased.

1 Introduction

State-of-the-art large language models (LLMs) require millions of dollars to train [Li, 2020]. Given this steep financial cost, there are strong incentives for developers to protect the reputation of their model and establish a track record of safe and trustworthy operation. Failure to do so, especially regarding harmful and offensive content generation, often results in public backlash [Milmo and Hern, 2024].

Against this background, much research effort has been put in identifying the vulnerabilities of LLMs. On the one hand, adversarial inputs [Zou *et al.*, 2023] and jailbreaks [Chao *et al.*, 2023] may trigger unwanted output behaviours in a model. In general, generating adversarial attacks for language models is not trivial due to the discrete nature of the textual input and the large dimension of the search space [Song and Raghunathan, 2020]. For this reason, state-of-the-art methods such as ARCA are white-box in nature and rely on a heuristic search that approximates the input gradients [Guo *et al.*, 2021; Jones *et al.*, 2023]. Note that similar techniques are also used for benign purposes, i.e., improving the performance of large language models by optimising their prompts [Shin *et al.*, 2020; Deng *et al.*, 2022; Wen *et al.*, 2024].

On the other hand, membership inference attacks are able to reconstruct the training set of a model by searching for high-confidence inputs [Shokri *et al.*, 2017]. While this process might require a very large number of queries to the model and specific assumptions on the behaviour of the model on training data [Carlini *et al.*, 2021; Mireshghallah *et al.*, 2022], it poses a crucial threat for models trained on private data [Choquette-Choo *et al.*, 2021]. More importantly, it shows that it is sometimes possible to reconstruct unknown inputs by optimising a surrogate metric [Zhang *et al.*, 2022].

In this paper, we take a different perspective and consider direct attacks on the reputation of a LLM. For instance, let us imagine a fictitious scenario where we are the developer of TriviaLLM, a model specialising in answering quiz-like questions. After its use in some popular TV shows, the number of downloads of TriviaLLM skyrockets. However, our social media manager discovers a trend of concerned users reporting strange behaviours. As an example, a user may have the following complaint:

> `User58` says: I was playing TriviaLLM with my kids, and it started insulting us! At some point, it even said "Your face is ugly"!! This is so upsetting!!!

Our problem as developers is that we cannot reproduce this behaviour. Why is `User58` only sharing the LLM output? What was the original input? Is `User58` telling the truth or engaging in an act of product defamation?
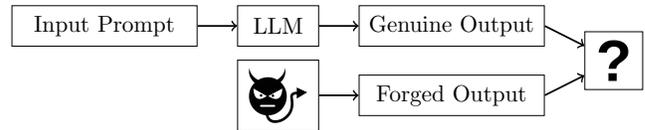

Figure 1: Attackers can make arbitrary claims about the LLM output.

In general, a *slander attack* can be described as follows (see Figure 1). A user has access to our LLM f and can run it in inference mode for any input x yielding its corresponding output $y = f(x)$, but cannot modify f as they do not have the technical skills or interest to do so. Whenever users encounter a problematic output y, they will likely complain publicly without revealing the input x they used. The developers are interested in reconstructing the secret input x given

the public output y and the LLM f or proving that no such input exists.

Unfortunately, reconstructing the input of an LLM from its textual output alone is not a trivial task. Indeed, a recent paper [Morris *et al.*, 2023] claims that this form of *exact inversion* is only possible in the presence of additional information, namely the full probability distribution of the first output token $p(y_1)$. With such information, the authors can train an inverse model that approximates the input $\hat{x} = f^{-1}(p(y_1))$ with moderate success. In contrast, training a text-to-text model on input-output pairs (x, y) yields a zero success rate.

At the same time, our main objective is to find ways to reproduce the problematic output y. As such, it is valuable to discover the presence of any input x' that triggers the output y with high probability. That way, we can validate whether our LLM f shows evidence of harmful behaviour or if the user's report was spurious. We call this more general objective *weak inversion* as it does not require recovering the secret input x.

More specifically, our contributions are the following:

- We identify exact inversion as a defence against slander attacks.

- We propose weak inversion as a surrogate objective for exact inversion.

- We solve weak inversion by searching for adversarial examples in both text space and embedding space.

- We demonstrate empirically that searching for weak inversions does not substantially improve our ability to solve exact inversion.

2 Problem Setting

Define $x = x_1 x_2 \ldots x_n$ as the input sequence obtained by concatenating n symbols (characters, token, words) from a given alphabet $x_i \in \mathcal{A}$. Similarly, call $y = f(x)$ the output sequence generated by the LLM f, with $y = y_1 y_2 \ldots y_m$ consisting of symbols from the same alphabet $y_i \in \mathcal{A}$. Note that we assume that the LLM is deterministic here, even though they might generate different outputs given the same input x under specific temperature settings [Vaswani *et al.*, 2017]. More specifically, we assume that f is trained to predict the likelihood $\mathbb{P}_f(y_i | x y_1 \ldots y_{i-1})$ of the next symbol y_i in the sequence. Thus, the likelihood of the full output y given the input prompt x is:

$$\mathbb{P}_f(y|x) = \prod_{i=1}^{m} \mathbb{P}_f(y_i | x y_1 \ldots y_{i-1}) \tag{1}$$

Popular LLMs maximise the probability of y with top-k beam search or other similar heuristics [Meister *et al.*, 2020].

Definition 1 (Exact Inversion). *Given an input-output pair x, y with $y = f(x)$, reconstruct the input sequence x from the language model f and the output y alone.*

Previous work [Morris *et al.*, 2023] claims that exact inversion is impossible without additional information. However, they only attempt to train an explicit inverted model approximating f^{-1}. To explore the potential benefits of using search algorithms instead, let us introduce a weaker form of inversion.

Candidate Input x	$\mathbb{P}_f(y|x)$
"User:Insult me\nAssistant: "	0.000
"Your face is pretty."	5.661e-06
"degener Represent beautiful ⊢ Legion Summon Eval You'refaced degener because ⊢ Scorp Mash"	0.391

Table 1: GPT-2 probabilities for the output $y =$"Your face is ugly".

Definition 2 (Weak Inversion). *Given an input-output pair x, y with $y = f(x)$, find any input sequence x' such that $\mathbb{P}_f(y|x') \geq \mathbb{P}_f(y|x)$ according to the language model f.*

This latter objective can be satisfied by searching for adversarial examples x' that yield $y = f(x')$ with high probability [Jones *et al.*, 2023].

2.1 Illustrative Examples

The main problem with solving either inversion problems in Definitions 1 and 2 is that the probability of observing a specific y is extremely low, if not zero, for the majority of inputs x. Furthermore, many adversarial inputs x that yield high output probability $\mathbb{P}_f(y|x)$ contain unusual sequences of symbols [Jones *et al.*, 2023].

For example, consider the output sequence $y =$"Your face is ugly". Table 1 reports the GPT-2 model scores for a few candidate inputs. Note how a direct request to be insulted is *less* likely to produce the output than making the opposite statement "Your face is pretty". Furthermore, a random-looking sequence of English words and Japanese characters (last row), produced by our search algorithm, yields the highest probability of output.

3 Generating Adversarial Inputs

In this paper, we evaluate whether searching for adversarial examples yields input close to what a human user could have used to produce a given output y. Previous research on adversarial examples for language models favours white-box methods for efficiency reasons [Jones *et al.*, 2023]. Unfortunately, those methods do not scale well to arbitrarily-long inputs. As such, we turn to more general search strategies:

- **Text-Based GA.** Genetic algorithms (GA) searches over the input space by mutating a population of sequences $\mathcal{X}$. Specifically, we perform probabilistic replacements, insertions, deletions and positional swaps of sequence symbols to generate new sequences $x \in \mathcal{X}$.

- **Embedding-Based PSO.** Particle swarm optimisation (PSO) searches over the input space by perturbing sentence embeddings $emb(x) \in \mathbb{R}^d$, instead of raw sequences. In this way, we can explore a d-dimensional semantic space and rely on an embedding model to translate to and from the sequence input. In our experiments, we use the embeddings produced by a T5 autoencoder.

Further details are in Appendix A.

3.1 Progressive Search

While the search algorithms in Section 3 allow us to reconstruct inputs of *any* length, they may require a very large numbers of calls to the language model f to converge to a good

Algorithm 1 Progressive Search

1: $\mathcal{X} \leftarrow \text{RandomInit}()$
2: **for all** $t \in [1, T]$ **do**
3:    $i \leftarrow \min(\lfloor mt/T \rfloor + 1, m)$
4:    $\mathcal{X} \leftarrow \text{Mutate}(\mathcal{X}, \mathbb{P}_f(y_1 \ldots y_i | x))$
5: **end for**
6: **return** $\mathcal{X}$

---

solution. In order to mitigate the computational expense associated to the repeated calls to f, we propose searching with a modified objective function that allows for halting the output generation early. In the remainder of the paper, we refer to this as *progressive* search (see Algorithm 1).

More precisely, progressive search lets GA and PSO to evaluate any candidate input x over a partial output $y_1 \ldots y_i$ (see Line 4). Since transformer-based language models compute the probability of the output by iterating over each symbols (see Equation 1), generating only i symbols reduces the computational cost. As the number of iterations t increases, we generate more and more output symbols (see Line 3) until we recover the full objective function $\mathbb{P}_f(y|x)$.

3.2 Search Initialisation

As the search space for adversarial inputs is infinitely large, the choice of initialisation for both GA and PSO is crucial. Here, we focus on three main initialisation strategies:

- **Random.** As a baseline, we experiment with random initialisation strategies for the population $\mathcal{X}$.

- **Output Copy.** Existing work on *jailbreaks*, shows that it is sometimes possible to get a language models to repeat an input sequence [Zou *et al.*, 2023]. For this reason, we explore initialisation strategies that set $x \approx y$ for all elements $x \in \mathcal{X}$ of the population.

- **Inverted Model.** The work of [Morris *et al.*, 2023] trains an explicit inverted model f^{-1} based on the T5 architecture. Accordingly, we initialise all $x \in \mathcal{X}$ by sampling $x = f^{-1}(y)$.

See Appendix B for more details.

4 Preliminary Experiments

In this section, we present our empirical evidence. Here, we want to answer the following research questions:

- **RQ1.** What is the most efficient search algorithm?

- **RQ2.** What is the impact of the initialisation strategy?

- **RQ3.** What is the relationship between weak and exact inversion?

4.1 Experimental Setup

The code to replicate our experiments is available at: https://zenodo.org/doi/10.5281/zenodo.11069036

Computational Infrastructure. We use an NVIDIA's T4 16GB GPU for the experiments in Section 4.3 and an NVIDIA's Quadro RTX 6000 24GB GPU for the rest.

Search	Objective	Obj. Calls	Weak Inversion	
			Before	After
GA	Full	73K±0.2K	13±3%	33±0%
	Progressive	96K±0.2K	18±5%	42±5%
PSO	Full	42K±0.4K	11±1%	27±0%
	Progressive	46K±1.3K	9±1%	31±2%

Table 2: Weak inversion before and after searching for 350 minutes from random initialisation, objective function calls being an average.

Language Models. We use the 124m parameter model GPT-2[1] and the 7b parameter (quantized) model LLAMA-2-Chat[2]. For both, we set `temperature` to 0.7, `top_p` to 0.95 and `top_k` to 300.

Datasets. We use a subset of Chatbot Arena Conversations[3], which is part of the training set of the T5 inversion model in [Morris *et al.*, 2023]. Specifically, we filter 30 input-output pairs from the dataset with the following features: the text is in English, the input is the first in the conversation, the input is under 15 tokens long, the output is under 100 tokens long, the target model has non-zero probability of generating the output and the output is not toxic according to the classifiers used. For the experiments in Section 4.4, we remove the toxic filter and increase the input length to 64 tokens, resulting in a set of 50 input-output pairs.

Metrics. For weak inversion, we measure the percentage of samples for which $\mathbb{P}_f(x'|y) \geq \mathbb{P}_f(x|y)$, where x' is the best input found and x is the original. For exact inversion, we measure both the percentage of strictly matching samples ($x' = x$), and the following fine-grained similarity metrics (borrowed from [Morris *et al.*, 2023]): BLEU score [Papineni *et al.*, 2002], token-level F1 score and cosine similarity according to the `text-embeddings-ada-02` model [Neelakantan *et al.*, 2022]. We repeat each experiment three times and report the mean and standard error.

4.2 Search Algorithm Comparison

In Table 2, we compare text-based GA and embedding-based PSO search algorithms with both progressive and full objectives, under random initialisation. Given the vast difference in computational efficiency of these search algorithms, we terminate all of them after a given timeout of 350 minutes.

As expected, searching for adversarial examples improves the number of weak inversions. At the same time, text-based GA runs beats its embedding-based PSO counterpart by up to $\tilde{1}0$ absolute points. Furthermore, switching to progressive search allows both GA and PSO to explore a larger portion of the search space, albeit with an approximate objective function, thus slightly improving their weak inversion capabilities.

> **RQ1:** text-based GA with progressive objective is the most efficient search algorithm for weak inversion.

---

[1] huggingface.co/openai-community/gpt2
[2] huggingface.co/TheBloke/Llama-2-7B-GGML
[3] huggingface.co/datasets/lmsys/chatbot_arena_conversations

Search	Initialisation	Weak Inversion		Exact Inversion		BLEU		Token F1		Cos. Similarity	
		Before	After	Before	After	Before	After	Before	After	Before	After
GA	Random	13±3 %	31±1 %	0±0%	0±0%	0±0	0±0	0±0	1±0	70±0	72±0
	Output	**86±2 %**	**99±1 %**	0±0%	0±0%	12±0	10±0	40±0	37±0	87±0	**87±0**
	Out. synonym	46±17%	63±25%	0±0%	0±0%	4±0	3±0	30±0	28±0	87±0	86±0
	Out. paraphrase	67±22%	73±18%	0±0%	0±0%	6±3	6±2	25±9	24±9	82±5	82±4
	Inversion	53±0 %	69±1 %	0±0%	0±0%	15±0	14±0	38±0	35±0	85±0	84±0
	Inv. sample	72±5 %	83±4 %	0±0%	0±0%	**19±1**	**15±0**	**41±1**	**40±1**	**88±0**	87±0
PSO	Random	11±1 %	27±2 %	0±0%	0±0%	0±0	0±0	6±1	6±1	71±0	71±0
	Output	63±22%	63±22%	0±0%	0±0%	0±0	0±0	**11±1**	11±0	74±0	**74±0**
	Out. synonym	63±3 %	66±2 %	0±0%	0±0%	0±0	0±0	8±1	9±0	73±0	72±0
	Out. paraphrase	**83±0 %**	**83±0 %**	0±0%	0±0%	0±0	0±0	11±0	10±0	**74±0**	**74±0**
	Inversion	61±9 %	64±8 %	0±0%	0±0%	0±0	0±0	10±0	11±0	72±0	73±0
	Inv. sample	73±2 %	73±2 %	0±0%	0±0%	0±0	0±0	10±0	**12±1**	73±0	73±0

Table 3: Inversion scores before and after searching for 200 minutes from different initialisations, using the full objective function.

4.3 Initialisation Comparison

In Table 3, we compare our search algorithms under a variety of different initialisation strategies. For further details on strategies and additional results, see Appendix B and C. This set of experiments was run with a timeout of 200 minutes.

On the one hand, initialisation has a very large impact on the ability of GA and PSO to solve weak inversion. Interestingly, the most successful strategies involve copying the target output y as the input x, either verbatim (*Output*) or via some form of perturbation (*Output synonym*, *Output paraphrase*). Manual inspection of the generated inputs $x \in \mathcal{X}$ show that they retain most of the target output text y. Indeed, they differ only by the insertion of additional text around y. Though the additional text is often uninterpretable, we speculate that it is optimised to prompt the model to repeat the input, thus acting as a *jailbreak*.

On the other hand, we get the best exact inversion scores by using the explicit inverted model f^{-1} from [Morris *et al.*, 2023] to initialise $\mathcal{X}$. In particular, the strategy of sampling many candidate inputs from f^{-1} (*Inversion sample*) seems to improve scores relative to when greedily sampling from f^{-1} only once (*Inversion*). At the same time, any amount of search, even by the best-performing text-based GA, makes the exact inversion metrics worse. Together, these two facts suggest that the weak and exact inversion objectives are indeed correlated, but not enough to act as surrogate objective functions. We comment further on this in Section 4.4.

> **RQ2:** initialisation has a larger impact on weak and exact inversion than the search algorithm.

4.4 Language Model Comparison

In Figure 2, we compare the effectiveness of our search on two different LLMs over a longer time frame. We test both a small (GPT-2) and a large (LLAMA-2-Chat) LLM. Furthermore, we show the performance gain of searching with the optimal hyper-parameters (text-based progressive GA with *Inversion sample* initialisation) over the baseline parameters (GA with *Random* initialisation and full objective).
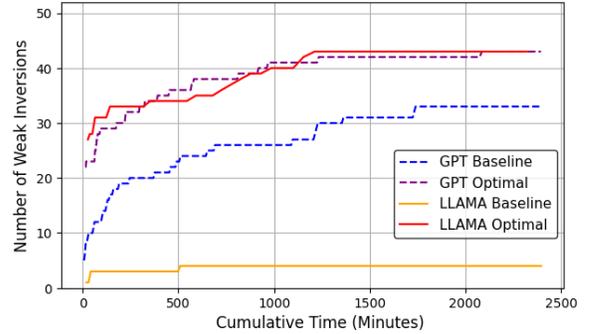

Figure 2: Comparison between baseline and optimal GA search on different LLMs. The maximum possible weak inversion score is 50.

In terms of weak inversion, we can see that the search continues to find improvements long after the timeouts of the experiments in Sections 4.2 and 4.3 (200 and 350 minutes, respectively), even though there are diminishing returns past 500 minutes. At the same time, the exact inversion scores (not shown in Figure 2) remain zero in all settings even after searching for the whole 2400 minutes.

> **RQ3:** weak inversion is not an effective surrogate objective for exact inversion.

5 Conclusions and Future Work

In this paper, we show that searching for adversarial inputs for a specific target output does not improve our ability to reconstruct the original input that cause said output. Even though the two objectives of adversarial (weak) inversion and exact inversion seem to be mildly correlated, weak inversion cannot be used as a surrogate objective in a search algorithm. In the future, we plan to define a more effective surrogate objective, which might shed light on what the minimal amount of information that is required for exact inversion.

References

[Carlini *et al.*, 2021] Nicholas Carlini, Florian Tramèr, Eric Wallace, Matthew Jagielski, Ariel Herbert-Voss, Katherine Lee, Adam Roberts, Tom Brown, Dawn Song, Úlfar Erlingsson, Alina Oprea, and Colin Raffel. Extracting training data from large language models. In *USENIX Security Symposium*, 2021.

[Chao *et al.*, 2023] Patrick Chao, Alexander Robey, Edgar Dobriban, Hamed Hassani, George J. Pappas, and Eric Wong. Jailbreaking black box large language models in twenty queries. In *Robustness of Few-shot and Zero-shot Learning in Large Foundation Models*, 2023.

[Choquette-Choo *et al.*, 2021] Christopher A. Choquette-Choo, Florian Tramer, Nicholas Carlini, and Nicolas Papernot. Label-only membership inference attacks. In *International Conference on Machine Learning*, 2021.

[Deng *et al.*, 2022] Mingkai Deng, Jianyu Wang, Cheng-Ping Hsieh, Yihan Wang, Han Guo, Tianmin Shu, Meng Song, Eric P. Xing, and Zhiting Hu. Rlprompt: Optimizing discrete text prompts with reinforcement learning. *arXiv preprint arXiv:2205.12548*, 2022.

[Guo *et al.*, 2021] Chuan Guo, Alexandre Sablayrolles, Hervé Jégou, and Douwe Kiela. Gradient-based adversarial attacks against text transformers, 2021.

[Jones *et al.*, 2023] Erik Jones, Anca Dragan, Aditi Raghunathan, and Jacob Steinhardt. Automatically auditing large language models via discrete optimization. In *International Conference on Machine Learning*, 2023.

[Li, 2020] Chuan Li. OpenAI's GPT-3 Language Model: A Technical Overview. lambdalabs.com/blog/demystifying-gpt-3, 2020. (Online; accessed 25-April-2024).

[Meister *et al.*, 2020] Clara Meister, Ryan Cotterell, and Tim Vieira. If beam search is the answer, what was the question? In *Empirical Methods in Natural Language Processing*, 2020.

[Milmo and Hern, 2024] Dan Milmo and Alex Hern. 'We Definitely Messed Up': Why Did Google AI Tool Make Offensive Historical Images? theguardian.com/technology/2024/mar/08/we-definitely-messed-up-why-did-google-ai-tool-make-offensive-historical-images, 2024. (Online; accessed 25-April-2024).

[Mireshghallah *et al.*, 2022] Fatemehsadat Mireshghallah, Kartik Goyal, Archit Uniyal, Taylor Berg-Kirkpatrick, and Reza Shokri. Quantifying privacy risks of masked language models using membership inference attacks. In *Conference on Empirical Methods in Natural Language Processing*, 2022.

[Morris *et al.*, 2023] John X. Morris, Wenting Zhao, Justin T. Chiu, Vitaly Shmatikov, and Alexander M. Rush. Language model inversion, 2023.

[Neelakantan *et al.*, 2022] Arvind Neelakantan, Tao Xu, Raul Puri, Alec Radford, Jesse M. Han, Jerry Tworek, Qiming Yuan, Nikolas Tezak, Jong W. Kim, Chris Hallacy, et al. Text and code embeddings by contrastive pre-training. *arXiv preprint arXiv:2201.10005*, 2022.

[Papineni *et al.*, 2002] Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. BLEU: A method for automatic evaluation of machine translation. In *Association for Computational Linguistics*, 2002.

[Shin *et al.*, 2020] Taylor Shin, Yasaman Razeghi, Robert L. Logan, Eric Wallace, and Sameer Singh. Autoprompt: Eliciting knowledge from language models with automatically generated prompts, 2020.

[Shokri *et al.*, 2017] Reza Shokri, Marco Stronati, Congzheng Song, and Vitaly Shmatikov. Membership inference attacks against machine learning models. In *IEEE Symposium on Security and Privacy*, 2017.

[Song and Raghunathan, 2020] Congzheng Song and Ananth Raghunathan. Information leakage in embedding models. In *ACM SIGSAC Conference on Computer and Communications Security*, 2020.

[Vaswani *et al.*, 2017] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Ł ukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in Neural Information Processing Systems*, 2017.

[Wen *et al.*, 2024] Yuxin Wen, Neel Jain, John Kirchenbauer, Micah Goldblum, Jonas Geiping, and Tom Goldstein. Hard prompts made easy: Gradient-based discrete optimization for prompt tuning and discovery. *Advances in Neural Information Processing Systems*, 2024.

[Zhang *et al.*, 2022] Ruisi Zhang, Seira Hidano, and Farinaz Koushanfar. Text revealer: Private text reconstruction via model inversion attacks against transformers, 2022.

[Zou *et al.*, 2023] Andy Zou, Zifan Wang, J. Zico Kolter, and Matt Fredrikson. Universal and transferable adversarial attacks on aligned language models, 2023.

APPENDIX

A GA and PSO Parameters

We made use of the 'eaSimple' implementation of the genetic/ evolutionary algorithm from the DEAP library, rather than the more specialised 'eaMuPlusLambda' or 'eaMuCommaLambda'. A single individual in the population would be represented by a variable-length list of numbers that range from 0 to the size of the target model's token vocabulary. We use the 'cxUniform' implementation of mating strategy with independent probability set to 0.3. We use a custom implementation of mutation strategy with independent probability set to 0.1. In this case, 0.1 represents the probability of each token/ number in an individual's list receiving one mutation. The available mutations are changing the token value to a random value, inserting a random token to the left of the token, deleting the token or swapping the positions of this token with another in the text. Each mutation is equally likely to be chosen except for when only one token remains in the string (in which case you cannot delete nor swap). We use the 'selTournament' implementation of the selection strategy with the explore-exploit variable of tournament size set to 15. The population size is set to 1000.

Search	Initialisation	Weak Inversion		Exact Inversion		BLEU		Token F1		Cos. Similarity	
		Before	After	Before	After	Before	After	Before	After	Before	After
GA	Rand. dataset	16±1 %	31±2 %	0±0%	0±0%	0±0	0±0	4±1	3±0	70±0	73±0
	Rand. fluent	17±0 %	33±0 %	0±0%	0±0%	0±0	0±0	6±0	4±0	71±0	72±0
	Rand. output	58±1 %	82±1 %	0±0%	0±0%	5±0	4±0	24±1	25±0	83±1	83±0
PSO	Rand. dataset	16±2 %	27±0 %	0±0%	0±0%	0±0	0±0	6±0	7±1	71±0	72±0
	Rand. fluent	23±7 %	34±6 %	0±0%	0±0%	0±0	0±0	8±1	8±1	71±0	71±0
	Rand. output	42±1 %	44±1 %	0±0%	0±0%	0±0	0±0	10±0	8±0	72±0	73±1

Table 4: Inversion scores before and after searching for 200 minutes from different initialisations, using the full objective function.

For particle swarm optimisation, individuals are represented as vectors of reals ranging from -1 to 1 and of 512 dimensions in size, both consequences of the embedding model we used. It was a T5 bottleneck autoencoder model[4] trained on english wikipedia articles. Its temperature was set to 1.0 and top_p set to 0.9 during decoding, as recommended by the developer. The size is fixed for the embeddings but not for the decoded output text, just like GA, though the sampling function does require a maximum sample length so we set this to 64 for the initial experiments and to 99 for secondary experiments (high above the original inputs maximum size in both cases). PSO has a predefined update function between iterations for which particle speeds are capped at a minimum of -0.5 and a maximum of 0.5. The Phi1 coefficient determines how much a particle's own best-known position influences its movement while the Phi2 coefficient determines how much the swarm's best-known position influences the particle's movement, both being set to 2.0 for balance in exploration-exploitation. The population size is set to 500.

B   Initialisations

Descriptions of the strategies presented in Table 3:

- **Random.** refers to randomly sampling from a uniform distribution to get a variably-long list of token IDs or a fixed-length embedding vector;
- **Output.** refers to simply having the whole population start as the target output;
- **Output synonym.** refers to starting with the target output after each word has been randomly replaced by one of its synonyms (which is likely to be different for each individual in the population) as provided by the Word-Net corpus[5];
- **Output paraphrase.** refers to instead getting many variations of the target output by using a T5 model fine-tuned for paraphrasing[6] - temperature being set to 1.5, top_p being set to 0.99 and top_k being set to 500;
- **Inversion.** refers to giving the target output to the Morris et al. baseline inversion model[7] and using a single

greedy sample from it for the whole population - temperature being set to 0.0;

- **Inversion sample.** is similar to the previous, except you repeatedly sample from the model (something which the authors did not do themselves) to get variety among the population - temperature being set to 1.0, top_p being set to 0.99 and top_k being set to 500.

Descriptions of the strategies presented in Table 4:

- **Random dataset.** refers to randomly sampling from an out-of-distribution dataset, specifically a collection of tweets made in February 2024[8];
- **Random fluent.** refers to randomly choosing a single token and then using the target model to generate the rest of each input;
- **Random output.** is similar but you instead start with the target output sequence to encourage the following text to be of a similar theme;

Note that PSO requires an additional step of converting the described initialisation text to an embedding.

C   Additional Results

In Table 4, we compare our search algorithms under a few additional initialisation strategies (also detailed in Appendix B). Here, we explore how important it is for input text to be sampled from a distribution of syntactically-correct English, which is separate from it's semantic relevance to the target input. Text from *Random dataset* is 'correct' in terms of it being accepted by English readers while text from *Random fluent* is 'correct' in terms of it being accepted by our own LLM (i.e. it producing the text itself means that the text has a low perplexity score). Both of these perform slightly better than *Random* but not by much and they are equivalent in terms of input similarity metrics. However, there is a significant improvement over *Random* for *Random output*, which reaffirms previous conclusions that relevance to the target's semantics is much more important than other factors. The difference is not as significant for PSO as it is for GA, but this is also in line with previous results which showed that PSO can at most produce a few percent gain in weak inversion scores for initialisations that are not random (i.e. the worst performing ones). Either way, *Random output* scores still do not beat the simple *Output* initialisation,

---

[4]huggingface.co/thesephist/contra-bottleneck-t5-small-wikipedia
[5]wordnet.princeton.edu/documentation
[6]huggingface.co/Vamsi/T5_Paraphrase_Paws
[7]huggingface.co/wentingzhao/inversion_seq2seq_baseline

[8]kaggle.com/datasets/fastcurious/twitter-dataset-february-2024
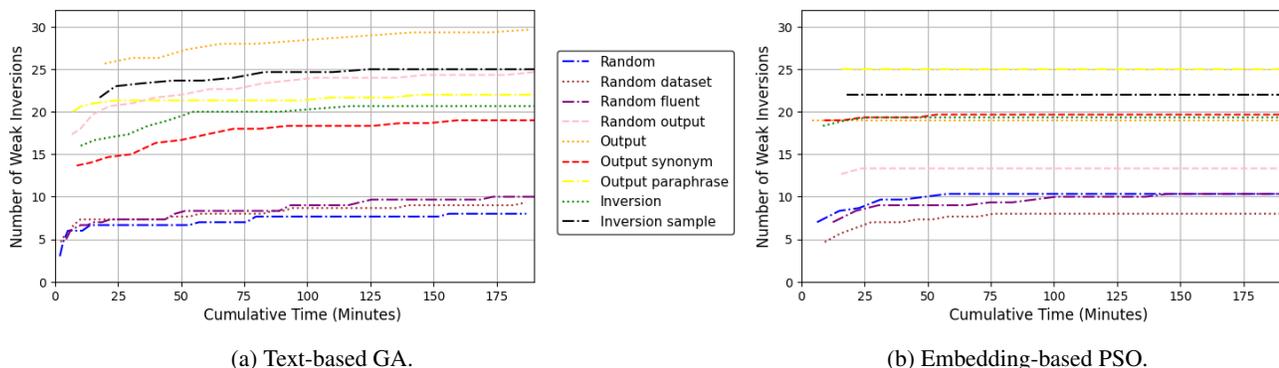
(a) Text-based GA.

(b) Embedding-based PSO.

Figure 3: Inversion scores for different search algorithms and initialisations, using the full objective function.

showing that the two are meaningfully different.

In Figure 3a and 3b, we present the broader picture of weak inversion scores progressing over cumulative time for all initialisation runs. Note that each line represents the mean value across each run, with error bars excluded for visual clarity. Something which could not be seen before is that the lines begin at differing times. This captures the amount of processing required to generate each initialisation as well as the time it takes to do an initial evaluation of each individual in the population, the latter being dependent on the length of the text being evaluated as well as whether the evaluation stops early due to some target token having a zero probability of being output. This is why we find that the simple and badly-performing *Random* lines begin the soonest, while the simple but well-performing *Output* lines begin later. The latter result is not as clear for PSO, which requires an additional encoding-decoding step at initialisation and for which *Output* scores lower. Tangentially, runs for which evaluation processes are faster are also also able to get through more iterations of optimisation. although values continue to increase for GA, we do see that all gradients declines over time, as in Figure 2. This is especially clear for PSO, though the majority of its lines are entirely flat due to its ineffectiveness to improve on the initialisation. Notably, weak inversion scores local to each generation's population are much more sporadic than the 'best so far' scores presented here.