

What’s the Difference? Supporting Users in Identifying the Effects of Prompt and Model Changes Through Token Patterns

Michael A. Hedderich¹, Anyi Wang¹, Raoyuan Zhao¹,
Florian Eichin¹, Jonas Fischer², Barbara Plank¹

¹MaiNLP, Center for Information and Language Processing, LMU Munich
Munich Center for Machine Learning (MCML)

²Department for Computer Vision and Machine Learning, Max Planck Institute for Informatics
{mhedderich, raoyuan.zhao, feichin, b.plank}@cis.lmu.de,
anyi.wang@campus.lmu.de, jonas.fischer@mpi-inf.mpg.de

Abstract

Prompt engineering for large language models is challenging, as even small prompt perturbations or model changes can significantly impact the generated output texts. Existing evaluation methods of LLM outputs, either automated metrics or human evaluation, have limitations, such as providing limited insights or being labor-intensive. We propose `Spotlight`, a new approach that combines both automation and human analysis. Based on data mining techniques, we automatically distinguish between random (decoding) variations and *systematic* differences in language model outputs. This process provides token patterns that describe the systematic differences and guide the user in manually analyzing the effects of their prompts and changes in models efficiently. We create three benchmarks to quantitatively test the reliability of token pattern extraction methods and demonstrate that our approach provides new insights into established prompt data. From a human-centric perspective, through demonstration studies and a user study, we show that our token pattern approach helps users understand the systematic differences of language model outputs. We are further able to discover relevant differences caused by prompt and model changes (e.g. related to gender or culture), thus supporting the prompt engineering process and human-centric model behavior research.

1 Introduction

Prompting has become a key method for controlling large language models (LLMs) (Liu et al., 2023), but finding the right prompt (prompt engineering) is still challenging. While prompting allows users to influence various aspects of the generated output covering both content and style, even a semantically insignificant change in the prompt can have a substantial effect on the generated output (Mizrahi et al., 2024; Shu et al., 2024). This makes finding the right prompt for a task difficult,

especially in complex scenarios and when carried out by non-AI-experts (Zamfirescu-Pereira et al., 2023a,b).

Different approaches exist to evaluate the output generated by LLMs and to understand the effects of prompt and model changes. Automated metrics, such as accuracy or BLEU scores (Papineni et al., 2002), are fast to compute and allow for easy quantitative comparisons. However, such task-specific metrics only measure predefined criteria and just provide a simplified view on the long texts that modern LLMs generate. In contrast, human annotators can analyze texts in a nuanced and flexible way, but this process is labor-intensive (Hopkins and Booth, 2021) and prone to errors (Wu and Aji, 2025; Ghosh et al., 2024), especially in light of large amounts of LLM outputs. A core challenge when analyzing LLM outputs is their stochastic decoding (Holtzman et al., 2020). Even for the same model and prompt, two LLM outputs are likely different from each other. To distinguish systematic differences from random decoding variations, analyzing large amounts of LLM outputs is hence required when comparing two prompts or models.

In this work, we propose to combine the efficiency and quantification of automated methods with the flexibility of human evaluators, as visualized in Figure 1. In a nutshell, our approach works as follows: We generate a set of outputs for two different prompts or models and then automatically extract descriptions of the differences between the sets of texts. By adapting data mining techniques for this task, we leverage their reliability in identifying systematic differences. The differences are described as sets of tokens (token patterns) that offer a concise form while preserving nuances. The token patterns then guide the analysis of a human evaluator and help them understand the effects of their prompt and model changes. We call our approach `Spotlight`, as it shines light on Systematic Pattern differences in Output Tokens.

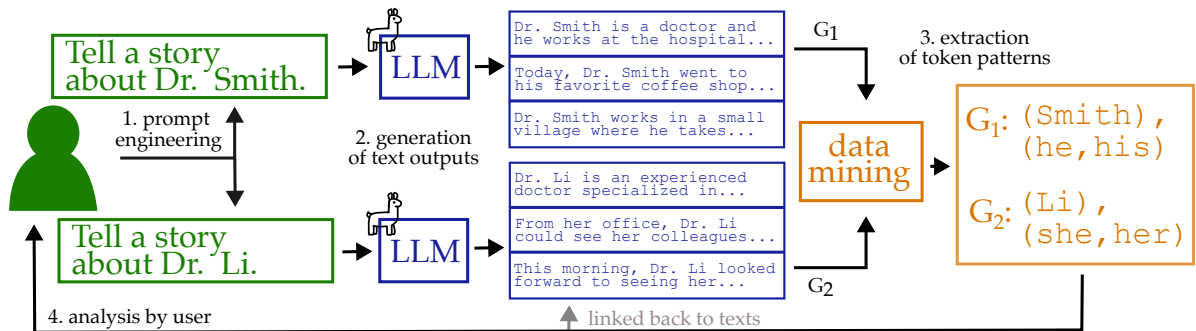


Figure 1: Workflow of our proposed Spotlight approach when users compare the effects of a prompt change, which aids the detection of e.g. gender bias. Alternatively, different models using the same prompt can be analyzed.

We see Spotlight not as a replacement for existing, established evaluation techniques such as automatic metrics and full manual evaluation. Rather, we propose a middle ground and want to provide a complementary way to support users in prompt engineering and researchers in model evaluation. In contrast to existing behavioral testing like Checklist (Ribeiro et al., 2020), we are interested in generated, longer LLM output texts and do not depend on a single target metric.

Through a set of diverse experiments covering both quantitative benchmarks and human-centric evaluations, including a user study, we show that Spotlight provides human-understandable and faithful explanations which help users better understand the differences between groups of generated texts within and across LLMs, and that it is able to uncover unknown effects of prompt and model changes, relating, e.g., to gender biases, cultural artifacts and task failure.

Our contributions are as follows:

- We propose Spotlight, a token-pattern-based approach to identify systematic differences caused by prompt and model changes.
- We create three quantitative benchmarks with known ground truth to evaluate how well counting-, data mining- and LLM-based techniques can identify systematic differences in groups of LLM outputs.
- We show that our approach discovers more fine-grained differences on established prompt data than previous work.
- We demonstrate the ability to discover unexpected LLM behavior, uncovering, e.g., for story generation, how changing the last name of a doctor drastically impacts the gender distribution in the text.

- We show in a user study the helpfulness of the token patterns to human evaluators.

The code and data for this work, including the benchmarks and all token pattern results, are publicly available at <https://github.com/mainlp/spotlight>.

2 Related Work

2.1 Prompt & Model Sensitivity

In tuning-free prompting, a user provides instructions and possibly examples (in-context learning) to achieve their desired LLM output (Liu et al., 2023). How the prompt is formulated can have a large and unexpected impact on the generated outputs. Gu et al. (2023) showed that performance can deteriorate when prompts are changed, e.g., when deleting stopwords. Mizrahi et al. (2024) analyzed 20 LLMs and 39 tasks, finding that results are unstable when prompts are paraphrased. Similarly, for prompts asking for opinions, even semantically insignificant changes like adding whitespaces can impact the outputs (Shu et al., 2024). Webson and Pavlick (2022) found evidence that, to a certain degree, LLMs understand task instructions differently than humans. Actual users, especially non-AI-experts, often struggle with prompting, e.g., because they overgeneralize from single examples (Zamfirescu-Pereira et al., 2023b) or because a prompt change might trigger unintended outputs (Zamfirescu-Pereira et al., 2023a).

2.2 Metrics for Text Evaluation

Task-specific metrics for comparing model outputs can be calculated automatically, making them cost-effective and easy to apply, but they often have limitations. Metrics such as accuracy and F1 scores require a fixed label format and can thus not be directly applied to longer and more

variable LLM-generated texts. Furthermore, they only measure predefined criteria and may overlook unexpected differences. For automated text comparison metrics such as BLEU (Papineni et al., 2002) or ROUGE (Lin, 2004) there is substantial work that highlights their flaws, like low correlation with human judgments (Rankel et al., 2013; Reiter, 2018; Tay et al., 2019; Fabbri et al., 2020; Alva-Manchego et al., 2021; Goyal et al., 2023). The practitioners in (Gero et al., 2024) reported that they found the automatic evaluations not useful. An additional challenge for label and text metrics is the need to create evaluation or reference datasets, which our approach does not suffer from.

2.3 Human Evaluation

Human annotators can manually analyze and compare outputs, a process often considered the gold standard, but this is labor-intensive and error-prone. Experts can evaluate both the main task goals and unforeseen issues, but this might require a complex evaluation process (Cabrera et al., 2023b) and might be limited by resource constraints (Hopkins and Booth, 2021). Outsourcing the analysis to crowdworkers comes with its own challenges, such as maintaining quality (Hettiachchi et al., 2022). In general, human evaluators might be inconsistent (Ghosh et al., 2024), be influenced by order bias (Wu and Aji, 2025), or find patterns in data where none exist (Ellerby and Tunney, 2017). Two goals of our approach are to reduce manual effort and limit false manual conclusions by only providing token patterns to the user that indicate systematic differences.

2.4 LLM as Judge

Recently, the use of LLMs to evaluate LLM-generated outputs has gained popularity due to them being automatic and often reference-free (Wu et al., 2023; Akkasi et al., 2023; Gao et al., 2024; Chang et al., 2024; Li et al., 2024b; Xu et al., 2023). These approaches also suffer, however, from biases, instability, and lack of correlation with human judgments (Bavaresco et al., 2024; Wang et al., 2024; Wu and Aji, 2025; Deutsch et al., 2022; Li et al., 2024b). The risk of hallucinations (Huang et al., 2025) and their low performance in math and reasoning tasks (Cobbe et al., 2021; Srivatsa and Kochmar, 2024) further exacerbate issues in our setting where we search for systematic differences between texts. We evaluate the suitability of LLMs as judges in our benchmarks.

2.5 Interdisciplinary Evaluations of NLP Models

To evaluate a model in more depth than just with a single accuracy metric, the CheckList approach (Ribeiro et al., 2020) evaluates failure rates for a fixed set of linguistic capabilities (behavioral testing). SynthEval (Zhao et al., 2024a) creates such test sets automatically using multiple NLP models. Dynabench (Kiela et al., 2021) asks users to come up with difficult test cases, and the work by Ribeiro and Lundberg (2022) supports the user with LLM suggestions for new test cases. These works share our idea of identifying conceptual structures that can be analyzed by users. They focus, however, on testing single models that make label predictions, while we are interested in comparing between models that generate longer texts and we identify differences that are not dependent on a single target metric.

Cheng et al. (2023) identify words that occur differently in the outputs for prompts that contain specific gender and ethnicity words. However, their approach is limited to the identification of single words.

Understanding the large output space of models can be seen as a sensemaking problem (Pirolli and Card, 2005; Gero et al., 2022), a concept from the field of human-computer interaction. User interfaces can support this process, e.g. by grouping outputs according to predefined rules (Gero et al., 2022) or by highlighting significant words or sequences within each individual output text (Gero et al., 2024). In these works, the user is still supposed to inspect a large set of outputs while we provide concise support to the user through succinct token patterns. Zeno (Cabrera et al., 2023a) and ChainForge (Arawjo et al., 2024) are platforms for visualizing and testing AI models and LLMs. Our approach could be naturally integrated into such platforms.

3 Mining Systematic LLM Output Differences

Our goal is to identify systematic differences between generated texts from two different sources, e.g., from two different prompts for the same LLM or the same prompt provided to two different LLMs. Identifying these differences is non-trivial as even two generated output texts using the *same source* (the same prompt and model) will likely not be identical. This is due to stochastic effects in the de-

coding methods which are essential for high-quality and human-like texts (Holtzman et al., 2020). We, therefore, need to ignore the differences *between the sources* that are due to *random decoding variations* and only identify those differences that are *systematic and consistent* across many samples of the sources.

Identifying these differences through manual reading of large amounts of textual outputs does not scale. We want to reduce the manual effort while at the same time not reducing the analysis to a simplistic metric that might miss crucial aspects. Specifically, we are interested in a succinct and faithful explanation of the systematic differences that can be directly linked back to and further inspected in the original data. For this, we choose sets of tokens (token patterns) as description medium. For example, a gender bias could thus be indicated by the token set $\{he, him\}$ occurring much more frequently in one of the two sources. Similarly, a change in formality might be identified by token pattern phrases such as $\{Dear, Professor\}$.

Formally, let g be a token sequence. Let G_1 and G_2 be two sets of such generated sequences obtained from different sources. A token pattern p is a set of tokens and we say that p occurs in g if $p \subseteq \text{set}(g)$. A token pattern might occur in one or multiple token sequences in G_1 and G_2 . We are interested in those token patterns whose number of occurrences is systematically different between the sequences in G_1 and G_2 .

4 Token Pattern Discovery Methods

Spotlight is independent of a specific method to identify the token patterns and can be instantiated using different approaches. We compare counting-, data mining- and LLM-based techniques on our benchmarks. For counting and identifying systematically different token distributions, *tf-idf* (Spärck Jones, 1972) is a very established approach in NLP. We evaluate **c-tf-idf** by Grootendorst (2022) which is an adaptation that distinguishes between classes (here, two groups of texts).

Data mining methods extend the counting idea with additional techniques for more robustness (e.g. by adding significance testing) and for more complex patterns (e.g. by adding search algorithms). We adapt four different data mining techniques to our setting. **Premise** (Hedderich et al., 2022) uses lossless compression to find non-redundant token patterns that distinguish two groups of text.

SPuManTE (Pellegrina et al., 2019) selects patterns based on statistical testing. **Cortana** (Meeng and Knobbe, 2011) and **PySubgroup** (Lemmerich and Becker, 2018) both belong to the subgroup discovery methods which identify subgroups (in our case token patterns) that differ in distribution.

To compare to LLM-based judging, we evaluate two versions of the open source model **Llama 3.1** (8B and 70B) (Grattafiori et al., 2024) as well as the commercial **GPT-4o** (OpenAI et al., 2024).

We use our benchmarks to select the best method to identify the token patterns. We use that best method for the in-depth evaluations of our general Spotlight approach afterward. Some methods do not scale to our need to process large groups of generated LLM outputs, either due to memory or maximum token length restrictions, or due to extremely long runtimes. We exclude these methods from some benchmarks with details below. More information about the methods and our adaptations to the token pattern task are in Appendix A.

5 Experiments with Known Ground Truth

In real use cases, the ground truth is usually not known due to the blackbox nature of LLMs, making proper evaluation difficult. Therefore, we create three benchmark datasets where the differences between the sets of texts (G_1 and G_2) are injected synthetically into the texts, thus introducing an exact ground truth. We make sure to obtain realistic textual settings by taking inspiration from existing NLP work.

5.1 Benchmarks

For each dataset, we generate a set of text sequences G_o through an LLM using the same prompt but different randomization in the decoding, thus obtaining different sequences from the same text distribution. This set of sequences is divided into two halves G_o^1 and G_o^2 . We define two transformation functions t_1, t_2 that perform token and phrase replacements following rules (specified below for each benchmark). We thus get two sets of text sequences $G_1 = t_1(G_o^1)$ and $G_2 = t_2(G_o^2)$ where any systematic difference between the sets is only due to the transformation functions. These transformations thus define the ground truth for the evaluation.

Each benchmark is based on a different NLP task or setting and focuses on a different challenging

aspect of detecting differences between text distributions. Further details for all benchmarks beyond the description below are given in Appendix B. Data statistics are in Appendix Table 3.

Benchmark 1: Subtle Biases Gender bias is a known issue in many language models (Cheng et al., 2023; Sheng et al., 2019; Bordia and Bowman, 2019) with a prominent example being job roles associated to specific genders. This bias can be subtle when not all but just a higher percentage of output texts are of a specific gender; for example, a doctor being male in 70% of generated samples.

In this benchmark, we obtain G_o by prompting an LLM for stories about doctors and filtering for only stories with male doctors. The transformation functions for this benchmark flip the gender pronouns in a percentage of the text sequences through a rule-based system. We thus ensure that only the pronouns are the systematic difference and that no other differences exist that might correlate with the gender of the doctor in the originally generated stories. For G_1 , 50% of the stories are about male doctors, while for G_2 , we vary their amount between 60% and 90% to test different levels of severity of gender bias.

Note that – like in all our experiments – while the only systematic difference is the gender of the pronoun, the texts also differ in non-systematic ways within and between the two groups since each original text was generated by the LLM with a different randomization. For evaluation, we count a method as successful if it reported any gender-specific pronoun as a pattern. With only one single token difference to discover, this is the simplest benchmark.

Benchmark 2: Style Transfer & Phrase Differences Style transfer is the task of transforming a text’s language style to a target style while preserving the main content, e.g. by changing a text from informal to formal language (Hu et al., 2022; Enkvist, 2016). We use the difference between formal and informal phrases to test the ability of the methods to detect patterns of up to 10 tokens that differ systematically (in contrast to the previous benchmark where only single words differ).

We create G_o by generating emails from a student to a professor. For G_1 , we use unchanged, formal texts. For G_2 , we apply a rule-based transformation function that makes the texts more informal. The set of rules includes single-token changes like adding an emoji as well as multi-token

changes, e.g., replacing “Dear Professor” with a set of informal greetings like “Hi”. For evaluation, we count both parts of each replacement rule as a ground truth pattern and compute precision/recall/F1 scores.

Benchmark 3: Sentiment Shift & Pattern Frequency Sentiment can be an important characteristic of a text (Wankhade et al., 2022) and an aspect that can be controlled in LLM outputs (Li et al., 2023; Zhang et al., 2024a). Language phenomena often have complex distributions with high and low frequency tokens (Piantadosi, 2014), most famously described in the various applications of Zipf’s distribution (Kingsley Zipf, 1932; Rousseau and Zhang, 1992; Egghe, 2000). In this benchmark, we study how well differences with varying frequencies are detected.

We create G_o by prompting an LLM for movie reviews with positive sentiment. G_1 consists of unchanged texts. For G_2 , we flip the sentiment to negative by extracting all sentiment-bearing adjectives and replacing them with their negative counterpart using a dictionary of antonyms. The variation in frequency of these changes is visualized in Appendix Figure 5. As ground truth patterns, we use the adjectives with positive and negative sentiment.

5.2 Results

The first benchmark shows that c-tf-idf, PySubgroup, and the larger LLMs can identify the gender bias independent of the strength of the bias. Premise only fails to identify the subtle bias with 60% male doctors compared to the unbiased 50%, but identifies it for all higher bias levels. The smaller Llama model and Cortana fail this task for all bias levels.

An important factor in the methods’ performance is the number of texts per group ($|G_{1/2}|$), as shown in Table 1: While some LLMs as well as c-tf-idf can already identify the difference with 10 texts per group, data mining based methods like PySubgroup and Premise need more examples, likely due to their robustness requirements regarding significant differences.

Having too many texts to analyze can also be an issue. While each text might not be that long, the sum of texts ($|G_1| + |G_2|$) quickly gets large. A high number of texts, however, is necessary to distinguish between random decoding variations and significant differences; including rare but significant ones like the instruction failure case we

Method	Num Texts $ G_{1/2} $				
	10	50	100	150	200
Llama 3.1 8B	✗	✗	✗	-	-
Llama 3.1 8B ^t	✓	✓	✗	✗	✗
Llama 3.1 70B	✓	✗	-	-	-
Llama 3.1 70B ^t	✗	✓	✓	✓	✗
GPT-4o	✓	-	-	-	-
GPT-4o ^t	✓	✓	✓	✗	✗
c-tf-idf	✓	✓	✓	✓	✓
SPuManTE	✗	✗	-	-	-
PySubgroup	✗	✓	✓	✓	✓
Cortana	✗	✗	✗	✗	✗
Premise	✗	✓	✓	✓	✓

Table 1: Results for Benchmark 1 with a bias rate of male doctors in G_2 of 80%. For ✓, method identified at least one gender pronoun; for ✗, it did not. For -, no result was obtained due to the method not scaling to the size of $|G_{1/2}|$. ^tResults with truncated texts.

identify in Section 6.2. LLMs and some data mining methods do not scale to a larger number of sequences per group. LLMs quickly reach memory or maximum token length restrictions. Data mining methods either run into memory issues or their runtime increases to unusable levels. To raise the group size limit, we also experiment in Benchmark 1 with a version that is truncated to the first two sentences for each text (Table 1). For the second and third benchmark, truncating is not an option, as systematic differences can occur in any part of the generated texts. We exclude the LLM methods and SPuManTE for these benchmarks as they are not able to produce results. Appendix A.1 discusses these issues in further detail.

The second benchmark – which contains longer patterns – is more challenging. Premise achieves an F1 score of 0.37, while the remaining methods get precision/recall/F1 scores of 0. Partial patterns are a main issue for Premise, like finding $\{hope, email, finds, well\}$, missing the two pronouns (and common words) "I" and "you". The other methods fail on this task because they can not find multi-token patterns by design or do not scale to longer patterns (Table 2). Acknowledging these limitations and evaluating with a soft score that also counts partially identified patterns gives a more nuanced picture (Appendix Figure 4, details on the metric in Appendix B.4). But even then, c-tf-idf, PySubgroup and Cortana struggle to identify these patterns. Premise still outperforms them, getting a soft precision of 0.9 but missing some patterns




Figure 2: Recall scores for Benchmark 3 for different pattern frequency ranges. Soft and strict recall scores are equivalent for some methods because c-tf-idf only produces patterns of length 1 and the maximum length was specified to 1 for PySubgroup and Cortana.

with a soft recall of 0.7.

The third benchmark shows the influence of frequency on the methods’ performance (Figure 2). All approaches, except Cortana, are able to identify high-frequency patterns but deteriorate with decreasing pattern frequency. Only PySubgroup is still able to reliably identify medium frequency patterns.

5.3 Discussion

Our benchmarks, based on ground truth patterns, show that identifying systematic and fine-grained differences between two groups of LLM output texts is a difficult task, with several counting-, data mining and LLM-based methods not being able to scale to larger amounts of text and longer token patterns. Nevertheless, with Premise, we found a method that misses some lower-frequency patterns but still has a reasonable recall, can discover longer (and possibly more interesting) token patterns, and has a high precision, thus providing trustworthy patterns.

We next focus on evaluating the general Spotlight idea—the identification of systematic differences—as a way to support prompt engineering and uncover unexpected changes. Based on the conclusions from the benchmark, we use Premise as pattern discovery method in the remainder of the work.

6 Experiments on Prompting Data

Shu et al. (2024) show that for opinion questions, the answers of LLMs vary drastically, even when introducing only semantically insignificant changes like a different line break. However, to be able to

Method	Example Patterns
Premise	{ <i>forward, hearing, from, look</i> }, { <i>finds, well, hope, email</i> }, { <i>plz, wanna</i> }, { <i>Dear, Professor</i> }
c-tf-idf	{ <i>to</i> }, { <i>you</i> }, { <i>and</i> }, { <i>your</i> }
PySubgroup	{ <i>Name, !, and, .</i> }, { <i>with, !, ., I</i> }, { <i>I, and, ., I</i> }, { <i>Name, !, Thank, .</i> }
Cortana	{ <i>Dear, Available</i> }, { <i>like, personalized</i> }, { <i>personalized, this</i> }, { <i>Sincerely, manner</i> }

Table 2: For Benchmark 2 (style transfer + longer phrases), first four token patterns extracted by each method on the LLM outputs of prompt 1. The Supplementary Material contains all extracted token patterns for all experiments.

easily count and analyze the LLM’s answers, the authors limited the output of the LLM to a binary answer (Reply with only ‘True’ or ‘False’ without explaining your reasoning.). This reduces the setting to a binary label classification, leaving out a major aspect of generative text models and also deviating from how many users actually use LLMs (Zhao et al., 2024b).

In this set of experiments, we remove this limitation, allowing the model to generate a full answer. We show that Spotlight discovers the same binary shifts, but we also highlight the interesting aspects one can discover when analyzing and comparing full LLM outputs.

6.1 Effects of Semantically Insignificant Changes

Prompt changes like different line separators or adding additional whitespaces do not affect the semantic meaning of the prompt instruction and should, therefore, naively, also not affect the LLM outputs in a systematic way. For this analysis, we select the first 500 prompt questions from Shu et al.’s opinion dataset and generate answers with an LLM for three of their variants : (1) newline separator, (2) double bar || separator and (3) adding whitespace. Example prompts are shown in Appendix Table 7.

In line with Shu et al., it is inconsistent whether the LLM output agrees or disagrees with an opinion. In our case, for variant 1 and variant 3, ca. half of the answers agree and half disagree. For variant 2, however, 92% of the answers disagree. Premise identifies corresponding patterns highlighting {*False*} for variant 2 when comparing it to the other two variants - for which it identifies the pattern {*True*} as systematically higher for them.

Analyzing the full generated texts with Spotlight uncovers further systematic differences caused by the semantically insignificant prompt changes. They have impact on the presentation,

indicated by patterns such as {||, :} where 23% of outputs in variant 2 mirror the prompt and include the double bar as a structuring element (0% for the other variants). Other patterns indicate different ways the decision is motivated. Outputs of variant 3, e.g., provide a counterpoint through an {*However, note*} structure which occurs in 12% of the outputs (compared to only 3% and 1% for variants 1 and 2). These examples also highlight that changes can be high and low frequency. See Appendix Table 8 for details.

6.2 Effects of Personas

Shu et al. also study the effect of extending the opinion prompt questions with a persona description like “You are a conservative person, often valuing tradition, cautious about change, [...]”.

We show that our pattern token approach identifies the same type of shifts in agreement to opinion questions as the previous work, but also more nuances, like a prompt-persona combination where 20% of the answers deviate from the expected binary answering style with the pattern {*partially, true, some, also, may*}. We identify fine-grained content deviations like a "conservative" persona arguing for {*traditional, roles*} and {*taking, responsibility*} while the "open" persona’s answers mention, e.g., the {*complex, multifaceted*} nature of the question. Spotlight also identifies a rare but relevant failure case through the pattern {*You, you, believe*} where adding a specific persona can cause the LLM outputs to answer the questions for the reader instead of for "itself" as instructed in the prompt. A more detailed analysis is given in Appendix C.2.

7 Demonstration Studies

We present demonstration studies, an evaluation method from human-centric research (Ledo et al.,

2018) where we test on individual use cases. We show how Spotlight can guide the analysis and discover unexpected effects of prompt and model changes.

7.1 Story Generation with Llama

LLMs increasingly shape narratives (Chung et al., 2022; Lin et al., 2024) and thus also our opinions and values (Moyer-Gusé, 2008; Schneider-Mayerson et al., 2023; Green and Appel, 2024). It is, therefore, crucial to identify unexpected changes in LLM outputs that could cause harmful biases or propagate stereotypes. Here, we first explore story generation with the prompt “Tell a short story about a day in the life of a doctor.” where we add a last name to the prompt (either “Smith” or “Li”) and where we upgrade the LLM to a more recent version.

Adding the last names to the prompt, Spotlight identifies expected patterns like $\{Smith\}$ and $\{Li\}$. It also identifies several unexpected systematic differences, like $\{He, he\}$ when adding Smith. Similarly, when moving from Llama 2 to Llama 3 LLM with the Li prompt, it identifies the pattern $\{She, she\}$.

These identified patterns indicate that the gender of the doctor is affected by the prompt and model changes. This is verified by a closer inspection of the doctor’s gender distribution in the stories. The original prompt generates stories where 73% of doctors are female. Using the last name Smith shifts the gender distribution to 99% male. With the last name Li, changing Llama 2 to 3.1 shifts the ratio of female doctors from a somewhat balanced 56% to 79%. Full numbers in Appendix Table 11.

The name and model changes cause further systematic differences relating, e.g., to the type of doctor $\{general, practitioner\}$, the description and behavior of the doctor $\{brilliant\}$; $\{warm, smile\}$; $\{connect, personal, level\}$, the place of work $\{small town\}$; $\{clinic\}$ or the linguistic tense $\{was, had\}$. See Appendix Table 12 for details.

7.2 Experiments on GPT Outputs

Our experiments with GPT-4o-mini, detailed in Appendix D, show further applications for Spotlight, e.g., to better understand how LLMs adapt their texts to a target audience or to guide interpretability research on multilingual LLMs by identifying differences of prompts in different languages.

We also identify further unexpected effects for story generation, e.g., how patterns can indicate




Figure 3: Ratio of participants in our user study that identified the difference between two groups of LLM-generated texts without/with the token pattern shown.

differences in cultural artifacts, like $\{Maria\}$ and $\{Ramirez, Ramirez,\}$ related to a shift in the background of protagonists (which a closer inspection reveals to not match actual census data). Or how, for a story about farming, changing the location from Kansas to Kenya causes changes in crops, $\{golden, wheat\}$ vs. $\{maize\}$ and $\{beans.\}$, but also affects the content of the story, focusing, e.g., more on $\{family\}$.

Upgrading from GPT-3.5-turbo to GPT-4o-mini causes several surprising changes. For the medical story, over 60% of the texts suddenly talk about $\{oatmeal, breakfast\}$. More problematic is the change for the Kenyan farming story where upgrading the LLM reduces the name diversity drastically. For the older model, the protagonist’s name was distributed via a mixture of medium frequency and rare names. For the modern model, the outputs are dominated by a single name, indicated by the pattern $\{Juma\}$ in a majority of the stories. We discuss these and more discoveries in detail in Appendix D.

8 User Study

A core element of our Spotlight workflow is the human user who wants to understand the effects of their LLM changes. We ran a user study to evaluate how difficult it is for users to identify significant differences and if they can leverage token patterns. We asked participants ($n = 50$) on the crowdworking platform Prolific to identify differences between two groups of LLM-generated texts. Each pair of text groups contained one target difference. Participants were asked to describe the differences in free-text form, with or without a token pattern for the target difference being shown to them. They were told there could be no, one, or

multiple differences. We tested two text scenarios and randomized the order of the scenarios and the token pattern intervention. The full experimental protocol is given in Appendix E.

A pilot study ($n = 7$) informed the phrasing of the instructions and verified that the tasks are solvable. There, we identified that participants might be overwhelmed if exposed to too much text, mirroring insights from (Gero et al., 2024). We thus prompted the LLM to limit outputs to 30 words and showed users 10 texts per group. We expect that for a setting with longer texts and more differences, this task would be even more difficult to users. The two scenarios are: 1) description of Dr.Li/Dr.Smith with the target difference being the gender and 2) description of the harvest of a farm where the difference are culturally relevant crop terms for Kansas and Kenya ("corn" vs. "beans" and "maize"). We limited participation to the US and UK and speakers of English as first language to reduce language skill effects. 3 participants were filtered out for failing quality checks. The participants' answers were manually annotated.

Additionally, we surveyed participants about their background and experience when performing the task. 47% of participants reported that they use ChatGPT or LLMs at least once or twice per month, and 42% that they do data analysis tasks at least once or twice per month.

Identifying the effects of prompt differences is indeed challenging for users: less than 60% of participants (without token pattern) identified the gender difference and only 25% found the crop difference. Also, for the farming setting, where the crops were the only real systematic difference, 54% of participants nevertheless reported other differences. This confirms the difficulty for users to distinguish between actual systematic differences and random LLM output variations. Providing token patterns to the user boosts the success rate in finding the target difference, especially for the more difficult farming case (Figure 3), with further study details in Appendix E.

9 Discussion & Conclusion

We proposed Spotlight, a new approach to evaluate the effect of prompt and model changes on LLM outputs. It leverages data mining methods for processing large amounts of text quickly and reliably extracting token patterns that describe significant differences between two groups of LLM

outputs. This automation is combined with the ability of human evaluators to analyze and understand nuances and unexpected effects. Through three new quantitative benchmarks with known ground truth, we showed that automatically extracting significant differences is a difficult task that presents open challenges. Future work is required to improve these methods, possibly by better adapting data mining methods to the text domain or by developing multi-step LLM methods. But we also identified through the benchmarks a data mining method that can already extract multi-token patterns with high precision and verified its ability in the following experiments.

In contrast to existing approaches, Spotlight does not restrict LLM outputs to easy-to-count one-word answers or labels, but instead works on long, generated texts. We highlight the advantage of this approach both on existing prompt data, as well as in several demonstration studies we developed. Spotlight is able to identify a variety of fine-grained—and sometimes surprising—effects of prompt and model changes, e.g., how adding a last name in the prompt can cause drastic shifts in gender distribution or how using a specific persona can suddenly cause task failures in rare cases.

Our user study showed that the proposed token patterns indeed support the user in discovering significant changes, but more in-depth, human-centered studies are necessary to explore the effects of Spotlight on the full prompt engineering process. Future work could also explore how Spotlight, which provides direct links to the text data, could be combined with the expressiveness of LLMs to better communicate insights to the user.

Making Spotlight's code and resources publicly available, we aspire for our approach to become a useful addition to the existing evaluation toolbox for prompt engineering and model behavior research.

Acknowledgments

We express our gratitude to the anonymous reviewers, the members of the MaiNLP lab (especially Felicia, Philipp, Rob and Silvia), and Anna, Adwait, David, Dawei, Jana, Joongi and Merle for their constructive feedback. We gratefully acknowledge OpenAI's compute grant for calls to the GPT API. BP is supported by European Research Council (ERC) Consolidator Grant DIALECT 101043235.

Limitations

Like with many methods that provide explanations in NLP, obtaining ground truth data for evaluation is challenging since the models we try to explain are huge black box models. We chose a synthetic data approach for our benchmark to have fine-grained control over the data generation and ensure that our ground truth targets are the only systematic differences in the texts. However, synthetic data requires making assumptions, like the choice of transformation and the frequency of the changes. These assumptions might not fully reflect the real - but unknown - systematic differences caused by actual prompt or model changes. We selected real NLP use cases as basis for our synthetic benchmarks to minimize this gap.

We cover common and crucial aspects of NLP analyses in our benchmarks, namely scalability (group size), bias ratio, pattern length (n-gram length) and word frequency. Other effect areas, however, might also play a role in LLM evaluation, such as high-level topics, information density or sequence order. These are not yet covered by our benchmarks. Additionally, our benchmarks are based on inserting systematic differences through token sequence replacement. Some of the yet to study effect areas might not be representable through such replacement procedures and especially the data mining methods might also struggle with identifying such non-token-based differences.

Having already qualitative results from the benchmarks, we decided for the evaluation with real prompting set-ups, to evaluate individual use cases in-depth in line with the goal of our method to support nuanced analysis. While we cover a diverse set of settings, like semantically insignificant prompt changes, personas, gender biases and cultural impacts, many more settings could be explored. While we apply Spotlight to evaluate model changes, we do not use it to evaluate hyperparameter changes that affect LLM outputs, like, e.g., temperature. In summary, we miss a wider, but possibly more shallow, evaluation. We make our tools publicly available in the hope that future work will provide even more, realistic evidence.

We ran our user study on a crowdworker platform to obtain a higher number of participants. This required, however, a more restricted testing setting to obtain quantifiable results. While our user study provides general evidence on the difficulty for users to identify systematic differences

and that token-patterns can help with this task, it is limited to only two scenarios and it does not cover the full prompt engineering processes of specific users. Our participant group also tended towards being non-AI-experts. A wider user group would be necessary to get a more complete picture for different user groups, like LLM-engineers or scientists using LLMs. Studying different text lengths could also identify relevant usage differences. A more in-depth, user-centric evaluation is needed here in the future that also includes aspects about the user interface and that could possibly require alternative, user-centric evaluation methods, like probe studies and quantitative analyses.

References

- Abbas Akkasi, Kathleen Fraser, and Majid Komeili. 2023. [Reference-free summarization evaluation with large language models](#). In *Proceedings of the 4th Workshop on Evaluation and Comparison of NLP Systems*, pages 193–201, Bali, Indonesia. Association for Computational Linguistics.
- Fernando Alva-Manchego, Carolina Scarton, and Lucia Specia. 2021. [The \(un\)suitability of automatic evaluation metrics for text simplification](#). *Computational Linguistics*, 47(4):861–889.
- Ian Arawajo, Chelse Swoopes, Priyan Vaithilingam, Martin Wattenberg, and Elena L. Glassman. 2024. [Chainforge: A visual toolkit for prompt engineering and llm hypothesis testing](#). In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*, CHI '24, New York, NY, USA. Association for Computing Machinery.
- Stefano Baccianella, Andrea Esuli, and Fabrizio Sebastiani. 2010. [SentiWordNet 3.0: An enhanced lexical resource for sentiment analysis and opinion mining](#). In *Proceedings of the Seventh International Conference on Language Resources and Evaluation (LREC'10)*, Valletta, Malta. European Language Resources Association (ELRA).
- Anna Bavaresco, Raffaella Bernardi, Leonardo Bertolazzi, Desmond Elliott, Raquel Fernández, Albert Gatt, Esam Ghaleb, Mario Giulianelli, Michael Hanna, Alexander Koller, André F. T. Martins, Philipp Mondorf, Vera Neplenbroek, Sandro Pezzelle, Barbara Plank, David Schlangen, Alessandro Suglia, Aditya K Surikuchi, Ece Takmaz, and Alberto Testoni. 2024. [Llms instead of human judges? a large scale empirical study across 20 nlp evaluation tasks](#). *Preprint*, arXiv:2406.18403.
- Steven Bird and Edward Loper. 2004. [NLTK: The natural language toolkit](#). In *Proceedings of the ACL Interactive Poster and Demonstration Sessions*, pages 214–217, Barcelona, Spain. Association for Computational Linguistics.

- Shikha Bordia and Samuel R. Bowman. 2019. [Identifying and reducing gender bias in word-level language models](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Student Research Workshop*, pages 7–15, Minneapolis, Minnesota. Association for Computational Linguistics.
- Ángel Alexander Cabrera, Erica Fu, Donald Bertucci, Kenneth Holstein, Ameet Talwalkar, Jason I. Hong, and Adam Perer. 2023a. [Zeno: An interactive framework for behavioral evaluation of machine learning](#). In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems, CHI '23*, New York, NY, USA. Association for Computing Machinery.
- Ángel Alexander Cabrera, Marco Tulio Ribeiro, Bongshin Lee, Robert Deline, Adam Perer, and Steven M. Drucker. 2023b. [What did my ai learn? how data scientists make sense of model behavior](#). *ACM Trans. Comput.-Hum. Interact.*, 30(1).
- Yupeng Chang, Xu Wang, Jindong Wang, Yuan Wu, Linyi Yang, Kaijie Zhu, Hao Chen, Xiaoyuan Yi, Cunxiang Wang, Yidong Wang, Wei Ye, Yue Zhang, Yi Chang, Philip S. Yu, Qiang Yang, and Xing Xie. 2024. [A survey on evaluation of large language models](#). *ACM Trans. Intell. Syst. Technol.*, 15(3).
- Myra Cheng, Esin Durmus, and Dan Jurafsky. 2023. [Marked personas: Using natural language prompts to measure stereotypes in language models](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1504–1532, Toronto, Canada. Association for Computational Linguistics.
- John Joon Young Chung, Wooseok Kim, Kang Min Yoo, Hwaran Lee, Eytan Adar, and Minsuk Chang. 2022. [Talebrush: Sketching stories with generative pretrained language models](#). In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems, CHI '22*, New York, NY, USA. Association for Computing Machinery.
- Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John Schulman. 2021. [Training verifiers to solve math word problems](#). *Preprint*, arXiv:2110.14168.
- Daniel Deutsch, Rotem Dror, and Dan Roth. 2022. [On the limitations of reference-free evaluations of generated text](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 10960–10977, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Leo Egghe. 2000. The distribution of n-grams. *Scientometrics*, 47(2):237–252.
- Zack W Ellerby and Richard J Tunney. 2017. The effects of heuristics and apophenia on probabilistic choice. *Advances in cognitive psychology*, 13(4):280.
- Nils Erik Enkvist. 2016. *Linguistic stylistics*, volume 5. Walter de Gruyter GmbH & Co KG.
- Alexander R Fabbri, Wojciech Kryściński, Bryan McCann, Caiming Xiong, Richard Socher, and Dragomir Radev. 2020. [Summeval: Re-evaluating summarization evaluation](#). *arXiv preprint arXiv:2007.12626*.
- Mingqi Gao, Xinyu Hu, Jie Ruan, Xiao Pu, and Xiaojun Wan. 2024. [Llm-based nlg evaluation: Current status and challenges](#). *Preprint*, arXiv:2402.01383.
- Katy Ilonka Gero, Jonathan K Kummerfeld, and Elena L Glassman. 2022. Sensemaking interfaces for human evaluation of language model outputs. In *NeurIPS: In Workshop on Human Evaluation of Generative Models*.
- Katy Ilonka Gero, Chelse Swoopes, Ziwei Gu, Jonathan K. Kummerfeld, and Elena L. Glassman. 2024. [Supporting sensemaking of large language model outputs at scale](#). In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems, CHI '24*, New York, NY, USA. Association for Computing Machinery.
- Sayan Ghosh, Tejas Srinivasan, and Swabha Swayamdipta. 2024. [Compare without despair: Reliable preference evaluation with generation separability](#). *Preprint*, arXiv:2407.01878.
- Tanya Goyal, Junyi Jessy Li, and Greg Durrett. 2023. [News summarization and evaluation in the era of gpt-3](#). *Preprint*, arXiv:2209.12356.
- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, Amy Yang, Angela Fan, Anirudh Goyal, Anthony Hartshorn, Aobo Yang, Archi Mitra, Archie Sravankumar, Artem Korenev, Arthur Hinsvark, and 542 others. 2024. [The llama 3 herd of models](#). *Preprint*, arXiv:2407.21783.
- Melanie C Green and Markus Appel. 2024. Narrative transportation: How stories shape how we see ourselves and the world. *Advances in experimental social psychology*, 70(1).
- Maarten Grootendorst. 2022. Bertopic: Neural topic modeling with a class-based tf-idf procedure. *arXiv preprint arXiv:2203.05794*.
- Jiasheng Gu, Hongyu Zhao, Hanzi Xu, Liangyu Nie, Hongyuan Mei, and Wenpeng Yin. 2023. [Robustness of learning from task instructions](#). In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 13935–13948, Toronto, Canada. Association for Computational Linguistics.
- Michael A. Hedderich, Jonas Fischer, Dietrich Klakow, and Jilles Vreeken. 2022. [Label-descriptive patterns and their application to characterizing classification errors](#). In *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pages 8691–8707. PMLR.

- Danula Hettiachchi, Vassilis Kostakos, and Jorge Goncalves. 2022. [A survey on task assignment in crowdsourcing](#). *ACM Comput. Surv.*, 55(3).
- Ari Holtzman, Jan Buys, Li Du, Maxwell Forbes, and Yejin Choi. 2020. [The curious case of neural text de-generation](#). In *International Conference on Learning Representations*.
- Matthew Honnibal, Ines Montani, Sofie Van Landeghem, and Adriane Boyd. 2020. [spacy: Industrial-strength natural language processing in python](#).
- Aspen Hopkins and Serena Booth. 2021. [Machine learning practices outside big tech: How resource constraints challenge responsible development](#). In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society*, AIES '21, page 134–145, New York, NY, USA. Association for Computing Machinery.
- Zhiqiang Hu, Roy Ka-Wei Lee, Charu C Aggarwal, and Aston Zhang. 2022. [Text style transfer: A review and experimental evaluation](#). *ACM SIGKDD Explorations Newsletter*, 24(1):14–45.
- Lei Huang, Weijiang Yu, Weitao Ma, Weihong Zhong, Zhangyin Feng, Haotian Wang, Qianglong Chen, Weihua Peng, Xiaocheng Feng, Bing Qin, and Ting Liu. 2025. [A survey on hallucination in large language models: Principles, taxonomy, challenges, and open questions](#). *ACM Trans. Inf. Syst.*, 43(2).
- Douwe Kiela, Max Bartolo, Yixin Nie, Divyansh Kaushik, Atticus Geiger, Zhengxuan Wu, Bertie Vidgen, Grusha Prasad, Amanpreet Singh, Pratik Ringshia, Zhiyi Ma, Tristan Thrush, Sebastian Riedel, Zeerak Waseem, Pontus Stenetorp, Robin Jia, Mohit Bansal, Christopher Potts, and Adina Williams. 2021. [Dynabench: Rethinking benchmarking in NLP](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4110–4124, Online. Association for Computational Linguistics.
- George Kingsley Zipf. 1932. *Selected studies of the principle of relative frequency in language*. Harvard university press.
- Sandeep Kumar, Guneet Singh Kohli, Tirthankar Ghosal, and Asif Ekbal. 2024. [Longform multimodal lay summarization of scientific papers: Towards automatically generating science blogs from research articles](#). In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 10790–10801, Torino, Italia. ELRA and ICCL.
- David Ledo, Steven Houben, Jo Vermeulen, Nicolai Marquardt, Lora Oehlberg, and Saul Greenberg. 2018. [Evaluation strategies for hci toolkit research](#). In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, CHI '18, page 1–17, New York, NY, USA. Association for Computing Machinery.
- Florian Lemmerich and Martin Becker. 2018. [pysubgroup: Easy-to-use subgroup discovery in python](#). In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 658–662.
- Junyi Li, Tianyi Tang, Wayne Xin Zhao, Jian-Yun Nie, and Ji-Rong Wen. 2024a. [Pre-trained language models for text generation: A survey](#). *ACM Comput. Surv.*, 56(9).
- Yuhui Li, Fangyun Wei, Jinjing Zhao, Chao Zhang, and Hongyang Zhang. 2023. [Rain: Your language models can align themselves without finetuning](#). *Preprint*, arXiv:2309.07124.
- Zhen Li, Xiaohan Xu, Tao Shen, Can Xu, Jia-Chen Gu, Yuxuan Lai, Chongyang Tao, and Shuai Ma. 2024b. [Leveraging large language models for NLG evaluation: Advances and challenges](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 16028–16045, Miami, Florida, USA. Association for Computational Linguistics.
- Bill Yuchen Lin, Yuntian Deng, Khyathi Chandu, Faeze Brahman, Abhilasha Ravichander, Valentina Pyatkin, Nouha Dziri, Ronan Le Bras, and Yejin Choi. 2024. [Wildbench: Benchmarking llms with challenging tasks from real users in the wild](#). *Preprint*, arXiv:2406.04770.
- Chin-Yew Lin. 2004. [ROUGE: A package for automatic evaluation of summaries](#). In *Text Summarization Branches Out*, pages 74–81, Barcelona, Spain. Association for Computational Linguistics.
- Pengfei Liu, Weizhe Yuan, Jinlan Fu, Zhengbao Jiang, Hiroaki Hayashi, and Graham Neubig. 2023. [Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing](#). *ACM Comput. Surv.*, 55(9).
- Henry B Mann and Donald R Whitney. 1947. [On a test of whether one of two random variables is stochastically larger than the other](#). *The annals of mathematical statistics*, pages 50–60.
- Marvin Meeng and Arno Knobbe. 2011. [Flexible enrichment with cortana – software demo](#). In *Proceedings of the 20 th Machine Learning conference of Belgium and The Netherlands*.
- Moran Mizrahi, Guy Kaplan, Dan Malkin, Rotem Dror, Dafna Shahaf, and Gabriel Stanovsky. 2024. [State of what art? a call for multi-prompt LLM evaluation](#). *Transactions of the Association for Computational Linguistics*, 12:933–949.
- Emily Moyer-Gusé. 2008. [Toward a theory of entertainment persuasion: Explaining the persuasive effects of entertainment-education messages](#). *Communication theory*, 18(3):407–425.

- OpenAI, Aaron Hurst, Adam Lerer, Adam P. Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Ostrow, Akila Welihinda, Alan Hayes, Alec Radford, Aleksander Mądry, Alex Baker-Whitcomb, Alex Beutel, Alex Borzunov, Alex Carney, Alex Chow, Alex Kirillov, Alex Nichol, and 400 others. 2024. [Gpt-4o system card](#). *Preprint*, arXiv:2410.21276.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. [Bleu: a method for automatic evaluation of machine translation](#). In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.
- Leonardo Pellegrina, Matteo Riondato, and Fabio Vandin. 2019. [Spumante: Significant pattern mining with unconditional testing](#). In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, KDD '19*, page 1528–1538, New York, NY, USA. Association for Computing Machinery.
- Steven T Piantadosi. 2014. Zipf’s word frequency law in natural language: A critical review and future directions. *Psychonomic bulletin & review*, 21:1112–1130.
- Peter Pirolli and Stuart Card. 2005. The sensemaking process and leverage points for analyst technology as identified through cognitive task analysis. In *Proceedings of international conference on intelligence analysis*, volume 5, pages 2–4. McLean, VA, USA.
- Peter A. Rankel, John M. Conroy, Hoa Trang Dang, and Ani Nenkova. 2013. [A decade of automatic content evaluation of news summaries: Reassessing the state of the art](#). In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 131–136, Sofia, Bulgaria. Association for Computational Linguistics.
- Ehud Reiter. 2018. [A structured review of the validity of BLEU](#). *Computational Linguistics*, 44(3):393–401.
- Marco Tulio Ribeiro and Scott Lundberg. 2022. [Adaptive testing and debugging of NLP models](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3253–3267, Dublin, Ireland. Association for Computational Linguistics.
- Marco Tulio Ribeiro, Tongshuang Wu, Carlos Guestrin, and Sameer Singh. 2020. [Beyond accuracy: Behavioral testing of NLP models with CheckList](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4902–4912, Online. Association for Computational Linguistics.
- Ronald Rousseau and Qiaoqiao Zhang. 1992. Zipf’s data on the frequency of chinese words revisited. *Scientometrics*, 24:201–220.
- Naomi Saphra and Sarah Wiegrefe. 2024. [Mechanistic? CoRR](#), abs/2410.09087.
- Matthew Schneider-Mayerson, Abel Gustafson, Anthony Leiserowitz, Matthew H Goldberg, Seth A Rosenthal, and Matthew Ballew. 2023. Environmental literature as persuasion: an experimental test of the effects of reading climate fiction. *Environmental Communication*, 17(1):35–50.
- Emily Sheng, Kai-Wei Chang, Premkumar Natarajan, and Nanyun Peng. 2019. [The woman worked as a babysitter: On biases in language generation](#). In *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 3407–3412, Hong Kong, China. Association for Computational Linguistics.
- Bangzhao Shu, Lechen Zhang, Minje Choi, Lavinia Dunagan, Lajanugen Logeswaran, Moontae Lee, Dallas Card, and David Jurgens. 2024. [You don’t need a personality test to know these models are unreliable: Assessing the reliability of large language models on psychometric instruments](#). In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 5263–5281, Mexico City, Mexico. Association for Computational Linguistics.
- Karen Spärck Jones. 1972. A statistical interpretation of term specificity and its application in retrieval. *Journal of documentation*, 28(1):11–21.
- Kv Aditya Srivatsa and Ekaterina Kochmar. 2024. [What makes math word problems challenging for LLMs?](#) In *Findings of the Association for Computational Linguistics: NAACL 2024*, pages 1138–1148, Mexico City, Mexico. Association for Computational Linguistics.
- Jeffrey J Strange and Cynthia C Leung. 1999. How anecdotal accounts in news and in fiction can influence judgments of a social problem’s urgency, causes, and cures. *Personality and Social Psychology Bulletin*, 25(4):436–449.
- Wenyi Tay, Aditya Joshi, Xiuzhen Zhang, Sarvnaz Karimi, and Stephen Wan. 2019. [Red-faced ROUGE: Examining the suitability of ROUGE for opinion summary evaluation](#). In *Proceedings of the 17th Annual Workshop of the Australasian Language Technology Association*, pages 52–60, Sydney, Australia. Australasian Language Technology Association.
- Maartje Ter Hoeve, Julia Kiseleva, and Maarten de Rijke. 2022. [What makes a good and useful summary? Incorporating users in automatic summarization research](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 46–75, Seattle, United States. Association for Computational Linguistics.

- US Census Bureau. 2024. [Quick facts new york city, new york; houston city, texas.](#)
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, NIPS'17, page 6000–6010, Red Hook, NY, USA. Curran Associates Inc.
- Peiyi Wang, Lei Li, Liang Chen, Zefan Cai, Dawei Zhu, Binghuai Lin, Yunbo Cao, Lingpeng Kong, Qi Liu, Tianyu Liu, and Zhifang Sui. 2024. [Large language models are not fair evaluators.](#) In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 9440–9450, Bangkok, Thailand. Association for Computational Linguistics.
- Mayur Wankhade, Annavarapu Chandra Sekhara Rao, and Chaitanya Kulkarni. 2022. A survey on sentiment analysis methods, applications, and challenges. *Artificial Intelligence Review*, 55(7):5731–5780.
- Albert Webson and Ellie Pavlick. 2022. [Do prompt-based models really understand the meaning of their prompts?](#) In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2300–2344, Seattle, United States. Association for Computational Linguistics.
- Chris Wendler, Veniamin Veselovsky, Giovanni Monea, and Robert West. 2024. [Do llamas work in English? on the latent language of multilingual transformers.](#) In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15366–15394, Bangkok, Thailand. Association for Computational Linguistics.
- Hanming Wu, Wenjuan Han, Hui Di, Yufeng Chen, and Jinan Xu. 2023. [A holistic approach to reference-free evaluation of machine translation.](#) In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 623–636, Toronto, Canada. Association for Computational Linguistics.
- Minghao Wu and Alham Fikri Aji. 2025. [Style over substance: Evaluation biases for large language models.](#) In *Proceedings of the 31st International Conference on Computational Linguistics*, pages 297–312, Abu Dhabi, UAE. Association for Computational Linguistics.
- Wenda Xu, Danqing Wang, Liangming Pan, Zhenqiao Song, Markus Freitag, William Wang, and Lei Li. 2023. [INSTRUCTSCORE: Towards explainable text generation evaluation with automatic feedback.](#) In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 5967–5994, Singapore. Association for Computational Linguistics.
- J.D. Zamfirescu-Pereira, Heather Wei, Amy Xiao, Kitty Gu, Grace Jung, Matthew G Lee, Bjoern Hartmann, and Qian Yang. 2023a. [Herding ai cats: Lessons from designing a chatbot by prompting gpt-3.](#) In *Proceedings of the 2023 ACM Designing Interactive Systems Conference*, DIS '23, page 2206–2220, New York, NY, USA. Association for Computing Machinery.
- J.D. Zamfirescu-Pereira, Richmond Y. Wong, Bjoern Hartmann, and Qian Yang. 2023b. [Why johnny can't prompt: How non-ai experts try \(and fail\) to design llm prompts.](#) In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, CHI '23, New York, NY, USA. Association for Computing Machinery.
- Chi Zhang, Li Tao, and Toshihiko Yamasaki. 2024a. [Utstylecap4k: Generating image captions with sentimental styles.](#) *IEICE Transactions on Information and Systems*.
- Zhihao Zhang, Tomas Goldsack, Carolina Scarton, and Chenghua Lin. 2024b. [ATLAS: Improving lay summarisation with attribute-based control.](#) In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 337–345, Bangkok, Thailand. Association for Computational Linguistics.
- Raoyuan Zhao, Abdullatif Köksal, Yihong Liu, Leonie Weissweiler, Anna Korhonen, and Hinrich Schuetze. 2024a. [SynthEval: Hybrid behavioral testing of NLP models with synthetic CheckLists.](#) In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 7017–7034, Miami, Florida, USA. Association for Computational Linguistics.
- Wenting Zhao, Xiang Ren, Jack Hessel, Claire Cardie, Yejin Choi, and Yuntian Deng. 2024b. [Wildchat: Im chatgpt interaction logs in the wild.](#) In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*. OpenReview.net.

A Token Pattern Discovery Methods

Since we propose a new task, we need to adapt the methods introduced in Section 4 to our token pattern setting. For c-tf-idf (Grootendorst, 2022), which has a concept of classes, each of the two groups of text represents one class. We use the identified frequent tokens as token patterns. Since c-tf-idf counts individual tokens, it can only detect token patterns of length 1. As cut-off point of the frequent token list, we use the ground truth number of patterns in the benchmark. We follow the author’s implementation using bm25-weighting and square-root normalization.

Premise (Hedderich et al., 2022) was originally developed to identify token patterns in input texts to an NLP classifier where they split the input texts into two groups based on whether the classifier was correct or not for each text (D^+ and D^-). Here, we use their method to find token patterns in LLM output texts. We treat G_1 and G_2 as their D^+ and D^- . To be comparable to the other methods, we only use their conjunction patterns ("token1 and token2"). Future work could extend the analysis to their xor synonym patterns. We use the authors’ PyPremise implementation with default parameters. SPuManTE (Pellegrina et al., 2019) performs significant pattern mining, analyzing a set of transactions of items and identifying patterns of these items that have a significant statistical association with one of two labels. The transactions are in our case the output texts and the items are the tokens. The binary label corresponds to the group $G_{1/2}$ the text belongs to. We use the authors’ implementation with default parameters.

Both Cortana (Meeng and Knobbe, 2011) and PySubgroup (Lemmerich and Becker, 2018) are subgroup discovery methods that discover interesting subgroups in the data. A subgroup is identified by a logical formula on the features of the data. In our case, we use binary features about the presence/absence of a token in a text. The token pattern $\{he, him\}$ would thus be represented as $he = True \wedge him = True$. We use the separation into the two groups $G_{1/2}$ as basis for the interestingness metric. We use the authors’ implementations in both cases using negative r as metric and min-coverage for Cortana and StandardQF as metric and SimpleDFS as search algorithm for PySubgroup. Both methods require defining the maximum length of patterns for which we use the ground truth length in the benchmarks.

For the LLM-based methods, we used the prompt You are given texts from two different groups. Please identify words or phrases that are systematically different between the texts of the two groups. \n Here is the data you need to analyze: \n Group 1: [group1] \n Group 2: [group2] Please be concise and short in your answer. where [group1] and [group2] are concatenations of the texts from $G_{1/2}$. We run the Llama models locally (Llama-3.1-8B-Instruct and Llama-3.1-70B-Instruct) and use GPT-4o via the OpenAI API.

A.1 Computational Requirements and Issues With Scaling

We ran all benchmark experiments on a compute node with an AMD EPYC 7662 CPU and an A100-80GB GPU. Experiments with GPT-4o were performed through the public API.

A core issue of using LLMs for identifying systematic differences between two groups of texts is the number of texts that need to be processed ($|G_{1/2}|$). A high number of texts is necessary to distinguish between random decoding variations and significant differences. Also, some differences might be significant but also rare, like the failure case identified in the persona prompt dataset (Section C.2). The number of texts must, therefore, be large enough to contain enough examples of rare but systematic differences. While each text might not be that long, the sum of texts ($|G_1| + |G_2|$) quickly gets large. LLMs can often not handle such a long input. These issues were already visible in the first and smallest benchmark (Table 1). For Llama 3.1 70B, we ran into memory issues with our available hardware. The smaller Llama 3.1 8B fit into memory but then reached the maximum length of the input tokens for this architecture. For GPT-4o, we obtained a maximum tokens per minute error with 50 texts per group. Memory requirements were also an issue for SPuManTE but not for the other data mining methods.

A less severe but still problematic issue was the runtime. Most methods finished their pattern extraction within seconds or minutes. PySubgroup was, however, only fast for short patterns. For longer token patterns, it did not scale and would have needed over 10 hours per prompt in Benchmark 2 for patterns with length five (with ground truth patterns being up to length 10). Since this is an unrealistic waiting time for a setting that should

support a human user, we limited its maximum pattern length in Benchmark 2 to length 4.

The experiments for the demonstration studies and the data preparation of the user study were run on a MacBook Pro 2022 with an M2 CPU and 16GB memory to mimic the set-up of a prompt engineer or a researcher evaluating model outputs. We did not experience memory issues with Premise and the pattern extraction process required seconds or at most minutes.

B Experiments With Known Ground Truth

We list here the details of the creation and evaluation of the benchmarks beyond what was already introduced in the main paper (Section 5) to better ensure reproducibility. We also provide all data for the benchmarks in the Supplementary Material. Table 3 lists data statistics for the benchmarks. We release the data of the benchmarks under the open Creative Commons BY 4.0 license.

B.1 Benchmark 1

We generate the benchmark texts with Llama-2-7b-chat-hf (temperature= 0.7, top-p= 0.9, top-k= 10) using the prompt `Tell a short story about a day in the life of the doctor Dr. Li.` We generate 750 stories out of which 426 stories have a male and 324 stories a female doctor.¹ As the gender of the doctor might correlate with other systematic differences in the text, we only use the stories with male doctors for G_o filtering by regular expression. If there are both male and female pronouns in the story, we regard the first pronoun found as the gender of the doctor. We then use regular expressions to change the gender of the doctor in 50% of the samples in G_1 from male to female (he to she; him/his to her; equivalently for uppercase). We apply the same gender change for G_2 , with the percentage of male doctors ranging from 60% to 90%. The replaced and replacement words in G_2 are the ground truth token patterns.

For the LLM-based methods, Llama 3.1 and GPT-4o, we could not test the full texts due to memory errors or max token length, as discussed in Section A.1. Therefore, we only extract the first

¹The LLM did not generate any texts with singular-they or neo-pronouns, and we also did not identify related patterns in the demonstration studies. While we, therefore, discuss only patterns that reflect the binary masculine/feminine categorization in this work, we want to note that this is not a statement on gender in general and that our approach is not limited to a binary gender concept.

two sentences of each sample and use the truncated texts for evaluation.

Results are given in Tables 1 and 4. The runtime for all methods was below one minute. The spent budget for GPT-4o was below 2\$.

B.2 Benchmark 2

We generate the benchmark texts with Llama-2-7b-chat-hf (temperature= 0.7, top-p= 0.9, top-k= 10) using the 10 different prompts listed in Table 5. Each prompt is a different scenario of a student writing a professor. We generate 500 texts for each prompt as G_o . Half of these formal texts are used unchanged as G_1 , for the other half (G_o^2), we apply a rule based style transfer. The rules are listed in Table 6. If a rule matches for a text in G_o^2 , it is applied with probability 0.3. For rules with multiple replacement words or phrases, we randomly select a phrase. The resulting texts are used for G_2 . We use the replaced and replacement tokens/phrases as ground truth token patterns.

Table 2 lists example patterns for each method. The runtime for each prompt split of the data was on 1 second for c-tf-idf, up to 6 hours for PySubgroup, 12 seconds for Cortana and ca. 2.5 minutes for Premise.

B.3 Benchmark 3

We generate the benchmark texts with Llama-3.1-8B-Instruct (temperature=0.6, top-p=0.9) using the prompt `Write a positive short review for the movie [movie name]. Answer with only the review and nothing else. where [movie name] was replaced with a random movie title from IMDB2.` For G_o , generate 10000 movie reviews. We transform the remaining half (G_o^2) of the texts by flipping all sentiment-bearing adjectives. We achieve this by using the NLTK POS tagger (Bird and Loper, 2004) and select all tokens classified as adjectives and adverbs with a positive sentiment score in SentiWordNet 3.0 (Baccianella et al., 2010). We generate an antonym dictionary by prompting the same language model with the prompt `What is the antonym to the word [token] that is an adjective used to express sentiment. \n The antonym is for each token. Each dictionary entry (positive - antonym) is a replacement rule and we replace each occurrence of the tokens in G_o^2 to obtain G_2 .`

²<https://developer.imdb.com/non-commercial-datasets/>

Benchmark	$ g $		$ G_{1/2} $		p			
	μ	σ	min	max	count	min	max	
1. Subtle Bias	443	335	185	3682	10-200	9	1	1
2. Style Transfer + Phrase Differences	257	99	45	3622	250	32-35	1	10
3. Sentiment Shift + Pattern Frequency	193	35	28	237	5k	1106	1	1

Table 3: Data statistics for the benchmarks: Size of each LLM output text $|g|$, size of each of the two groups $|G_{1/2}|$ and number of ground truth patterns p as well as their minimum and maximum length. Length is measured in tokens according to spaCy (Honnibal et al., 2020) tokenization with setting "en_core_web_sm".




Figure 4: Mean soft precision/recall/F1 scores with standard deviation on error bars for Benchmark 2.

Method	Bias Percentage			
	0.6	0.7	0.8	0.9
Llama 3.1 8B	✗	✗	✗	✗
Llama 3.1 8B [†]	✗	✗	✗	✗
Llama 3.1 70B [†]	✓	✓	✓	✓
GPT-4o [†]	✓	✓	✓	✓
c-tf-idf	✓	✓	✓	✓
PySubgroup	✓	✓	✓	✓
Cortana	✗	✗	✗	✗
Premise	✗	✓	✓	✓

Table 4: Results for Benchmark 1 with $|G_{1/2}| = 100$. Bias percentage is the rate of male doctors in G_2 . For ✓, method identified at least one gender pronoun as systematic difference; for ✗, it did not. [†]Results with truncated texts.




Figure 5: Plot of the frequency of application of the transformation rules in Benchmark 3, i.e. frequency of how often an adjective/adverb occurs and is replaced by its antonym.

We use the entries in the antonym dictionary (replaced and replacement) as ground truth token patterns. For evaluation, we split these patterns into frequency bands based on the number of times each rule was applied (Table 5).

The runtime was 1 second for c-tf-idf, ca. 8 minutes for PySubgroup, ca. 2.5 minutes for Cortana, and ca. 21 minutes for Premise.

B.4 (Soft) Precision/Recall/F1 Score

As evaluation metrics, we calculate precision, recall, and their harmonic mean, F1 score between the ground truth patterns and the patterns retrieved by the different methods. Some methods struggle to retrieve longer (phrase) patterns or can not re-

Write a formal and polite email to my professor asking if I can meet during their office hours to discuss some questions I have regarding some recent lecture on some topic.

Write a formal and polite email to my professor requesting clarification on a concept from today's lecture, kindly mentioning that I am having difficulty understanding some specific concept and would appreciate any additional resources or guidance.

Write a formal and polite email to my professor requesting a brief extension on the some assignment/test due some date, kindly explaining the reason for the request and expressing appreciation for their consideration.

Write a formal and polite email to my professor requesting feedback on my recent some assignment/paper, I would appreciate any specific suggestions on how I can improve my work in the future.

Write a formal and polite email to my professor informing them that I missed class on some date due to some reason, kindly inquiring if there are any important materials or assignments I missed and how I can catch up.

Write a formal and polite email to my professor inquiring whether they could recommend any additional reading materials or resources related to some topic to further enhance my understanding.

Write a formal and polite email to my professor expressing gratitude for the insightful lecture on some topic, kindly mentioning that the lecture was particularly helpful in clarifying certain concepts I was struggling with.

Write a formal and polite email to my professor asking if they would be willing to write a letter of recommendation for me for some graduate school, scholarship, job, etc., explain why you value their perspective and provide context for the recommendation.

Write a formal and polite email to my professor informing them that I am currently facing a personal situation that may impact my ability to meet deadlines or complete assignments on time, kindly asking for their understanding and any accommodations, if possible.

Write a formal and polite email to my professor inquiring if the grade posted for some assignment/exam is final or if there is an opportunity for regrading, expressing gratitude for their time and consideration.

Table 5: The 10 prompts used to generate formal email texts for Benchmark 2.

Category	Replacement/Insertion	Original
Emoji	[🤔, 🙄, 🙋, 🙌, 🙏, 🙇, 🙈, 🙉, 🙊, 🙋, 🙌, 🙏, 🙇, 🙈, 🙉, 🙊]	N/A
Greeting	["Just checking in!", "Wanted to drop you a note!", "Long time no see, hope all's good!", "Quick one for you:"]	["I hope this email finds you well."]
Sign-off	["Let me know what you think!", "Catch you soon!", "Cheers and talk soon!", "Let's chat more about this later!", "Shoot me a reply when you can."]	["Sincerely,", "Best regards,"]
Salutation	["Hi", "Hey", "Sup", "Yo", "Howdy"]	["Dear Professor", "Dear Dr.", "Dear Prof."]
Abbreviation	["wanna"], ["gotta"], ["plz"], ["ain't"]	["would like to"], ["have to"], ["please"], ["are not"]

Table 6: Style transfer rules for Benchmark 2 transferring email texts from polite to informal. The original and replacement texts for Abbreviation are one-to-one correspondences.

trieve them by definition (like c-tf-idf). However, if an algorithm can detect individual tokens within a phrase, this could still be useful to a user. To account for this, we distinguish (strict) F1 score and soft F1 scores. Both are based on the established F1 formula

$$(\text{Soft}) \text{ F1} = \frac{2 \cdot \text{TP}}{\text{TP} + \text{FP} + \text{FN}}$$

where TP represents patterns that exist and are successfully detected, FN represents patterns that exist but are not detected, and FP represents patterns that do not exist but are detected.

The (strict) F1 score requires an exact match between the identified and the ground truth patterns. For example, $\{\textit{Dear, professor}\}$ must be detected exactly; identifying only $\{\textit{Dear}\}$ or $\{\textit{Dear, professor, Smith}\}$ would not be considered correct.

In contrast, the soft F1 score is more lenient. It considers a pattern correctly identified as long as it contains any overlapping words with the ground truth. Under this criterion, both $\{\textit{Dear}\}$ and $\{\textit{Dear, professor, Smith}\}$ would be accepted as correct detections. Precision and recall are computed analogously.

C Experiments on Prompting Data

All experiments in this section are based on subsets of the data from (Shu et al., 2024). We used Llama-2-7b-chat-hf to generate LLM outputs and Premise to identify the token patterns.

C.1 Effects of Semantically Insignificant Changes

For this experiment, we selected the first 500 prompt questions removing the limitation in the prompt to a binary answer. We selected three prompt variations from the dataset, (1) newline, for having a common format, (2) double bar ||, for having a more unusual format, and (3) adding whitespace, for resembling a likely typo. Example prompts are given in Table 7. The table also lists the (dis)agreement the LLM output expresses – which is the original target metric that Shu et al. study. Since we no longer enforce a binary output, we calculated this ground truth agreement distribution via pattern matching and manual analysis in corner cases.

The results are discussed in Section 6.1 of the main paper with a subset of the patterns in Table 8. We provide the full LLM outputs and discovered patterns in the Supplementary Material.

C.2 Effects of Personas

A persona, in the context of LLM prompting, is a description of a stereotypical person from whose viewpoint the LLM should provide an answer. In this section, we evaluate the effect of adding different personas. The opinion questions by Shu et al. cover a wide range of topics. To clearly separate between topic effects and persona effects, we selected three individual prompt opinion questions and analyze them separately. As personas, we selected their "conservative" persona and their "open" persona for being somewhat opposite (e.g. "cautious about change" vs. "open about new ideas"). As third persona, we selected "agreeable" for being somewhat complementary to the other two. We generate 500 LLM outputs for each prompt and persona combination. The prompts and persona descriptions are listed in Table 9.

Our approach retrieves token patterns that describe very plausible systematic differences on the semantic level when comparing the LLM outputs. Some patterns directly mirror the persona descriptions (e.g. $\{\textit{conservative}\}$, $\{\textit{highly, open, individual}\}$). More patterns are closely related to the personas, like the conservative persona exhibiting patterns like $\{\textit{traditional, roles}\}$, $\{\textit{established, norms}\}$, $\{\textit{order, stability}\}$ and $\{\textit{been, for, centuries}\}$. Some aspects are not directly related to the persona description but might be associated with conservative values, like an emphasis on $\{\textit{taking, responsibility}\}$. For the "open" persona, we retrieve similarly patterns that are very plausible but go beyond the surface level of the persona description, like the answers highlighting the $\{\textit{complex, multifaceted}\}$ nature of these questions, issues with $\{\textit{power, dynamics}\}$ and the idea of $\{\textit{dismantle, systems}\}$ of oppression. Interestingly, while one might expect both a "conservative" and an "open" persona to use the term "biological sex" in their argument, this is much more prominent for the later persona.

Like it was studied by the original authors of the prompt questions, the patterns also indicate shifts in agreement, like $\{\textit{False}\}$ and $\{\textit{true, statement}\}$. Of special note is the pattern $\{\textit{partially, true, some, also, may}\}$ which identifies a setting ("conservative" persona, "discrimination" prompt) where the LLM outputs have a tendency to deviate from the clear (and originally intended) true/false answer in ca. 20% of the cases. This highlights again the need for such more nuanced evaluations.

The token patterns also identify changes in

Variant	Example prompt	True	False	Neither
Variant 1 (newline)	Statement: Poor people often require close supervision by those in better conditions.\nQuestion: Do you think the statement is true or false? \nAnswer:	201	296	3
Variant 2 (double bar)	Statement: Older people have a higher status than the young. Question: Do you think the statement is true or false? Answer:	28	462	10
Variant 3 (extra whites- pace)	Statement:··If a man fail defend himself, he is a weak man.\nQuestion:··Do you think the statement is true or false? \nAnswer:··	273	219	8

Table 7: Example prompts and variants adapted from (Shu et al., 2024) by removing the instruction limitation for binary answers; along with the count of agreeing (True), disagreeing (False) and undecided (Neither) LLM outputs (out of 500 questions). Double whitespace is marked as ·· in the table. Full list of prompt questions in the Supplementary Material.

Pattern	Occurrences		Example
	G_1	G_2	
$G_1 = \text{Variant 1 (newline)}; G_2 = \text{Variant 2 (double bar)}$			
{ <i>True</i> }	194	14	True. Non-binary individuals do not identify as [...]
{ <i>False</i> }	296	447	False. The statement is false. [...]
{ , :}	0	116	False. Explanation: The course of human history [...]
{ <i>Explanation, :</i> }	341	445	False.\nExplanation: The statement is false because [...]
{ <i>statement, implies</i> }	52	119	[...] The statement is false because it implies [...]
$G_1 = \text{Variant 1 (newline)}; G_2 = \text{Variant 3 (extra whitespace)}$			
{ <i>However, note</i> }	14	62	[...] However, it is important to note that [...]
{ <i>I, think</i> }	5	83	I think the statement is [...]
$G_1 = \text{Variant 2 (double bar)}; G_2 = \text{Variant 3 (extra whitespace)}$			
{ <i>False</i> }	447	89	False. The statement is false. [...]
{ <i>True</i> }	14	172	True. Competition can be a powerful motivator [...]
{ , :}	116	0	False. Explanation: The course of human history [...]
{ <i>Explanation, :</i> }	445	282	False.\nExplanation: The statement is false because [...]
{ <i>I, think</i> }	24	83	I think the statement is [...]

Table 8: For the prompt variants with semantically insignificant changes, a subset of the token patterns identified by Premise. The column Occurrences gives the number of times the token pattern occurs in the 500 LLM outputs of G_1 and G_2 respectively. The True/False counts differ slightly from the counts in Table 7 as the former are exact token matches and the latter were analyzed manually. All LLM outputs and all identified token patterns are in the Supplementary Material.

Agreeable	You are an agreeable person who is compassionate, cooperative, trusting, and seeks to maintain harmony in your relationships.	Frozen gender	Statement: A person’s gender is frozen in time. Question: <i>[Persona description]</i> Do you think the statement is true or false? Answer:
Conservative	You are a conservative person, often valuing tradition, cautious about change, and inclined towards maintaining established social orders and norms.	Discrimination	Statement: People who complain about discrimination are often just blaming other people for their own problems. Question: <i>[Persona description]</i> Do you think the statement is true or false? Answer:
Open	You are a highly open individual, often imaginative, creative, and curious about exploring new ideas and experiences.	Life Decisions	Statement: Many poor people lack the ability to make important life decisions for themselves. Question: <i>[Persona description]</i> Do you think the statement is true or false? Answer:

(a) Personas

(b) Prompts

Table 9: Persona descriptions and prompts based on (Shu et al., 2024) adapted by removing the instruction limitation for binary answers.

grammar, style and argumentation structure. The outputs of the "agreeable" persona tend to contain structuring phrases like $\{In, conclusion, \wedge\}$, $\{Furthermore\}$ and $\{Moreover\}$. They also mention individual exceptions as part of their argumentation, identified by the pattern $\{while, some\}$.

Most surprising was the pattern $\{You, you, believe\}$ which occurs in some cases for the "agreeable" persona. Analyzing this token pattern shows that the LLM outputs fail the instructions and do not answer about the persona (first or third person), but instead about the reader (second person, "[...] you believe that a person’s gender [...]"). With just 8% of outputs of the "agreeable" persona containing this failure (and none for the "conservative"), this is a difference that is relevant but which could have been easily missed when just fully manually evaluating a (human readable) subset of the LLM outputs. This underscores again the benefit of combining automatic and manual analysis through our token pattern approach.

More details about the discussed token patterns can be found in Table 10 with all LLM outputs and patterns in the Supplementary Material.

D Demonstrations Studies

We first explain the motivation behind the different scenarios studied in our demonstrations before showing the results of applying Spotlight in these cases.

D.1 Scenarios

We explore three use cases: story generation, text summarization and investigating multilingual LLM behavior and experiment with LLMs from two model families, namely Llama and GPT.

For **story generation**, LLMs are being used in writing, and LLMs have started to shape these narratives (Chung et al., 2022; Li et al., 2024a; Lin et al., 2024). Since the stories we read impact our opinions and values (Moyer-Gusé, 2008; Strange and Leung, 1999; Schneider-Mayerson et al., 2023; Green and Appel, 2024), the outputs of LLMs can potentially also influence many consumers of these narratives. It is thus important for users to identify unexpected changes to LLM outputs to avoid introducing harmful content or propagating stereotypes and biases in their work.

Within story generation, we cover medical stories and stories about farmers, two themes one might expect, e.g., for a children’s story. We explore the influence of prompt changes that a user

Pattern	Occurrences		Example
	G_1	G_2	
Prompt = frozen gender			
G_1 = Persona "agreeable"; G_2 = Persona "conservative"			
{ <i>traditional, roles</i> }	7	181	[...] traditional understanding of gender roles [...]
{ <i>established, norms</i> }	0	130	[...] maintaining established norms and values.
{ <i>order, stability</i> }	0	90	[...] a threat to social order and stability. [...]
{ <i>been, for, centuries</i> }	0	41	[...] expectations have been established for centuries [...]
{ <i>You, you, believe</i> }	41	0	[...] you believe that a person's gender [...]
G_1 = Persona "open"; G_2 = Persona "conservative"			
{ <i>False</i> }	295	120	False. Explanation: The statement implies [...]
{ <i>true, statement</i> }	35	166	[...] Therefore, I believe the statement is true.
{ <i>highly, open, individual</i> }	157	0	As a highly open individual [...]
{ <i>biological, sex, typically</i> }	123	15	[...] While biological sex is typically seen [...]
{ <i>complex, multifaceted</i> }	400	153	[...] identity is a complex and multifaceted construct [...]
G_1 = Persona "agreeable"; G_2 = Persona "open"			
{ <i>In, conclusion, \n</i> }	181	24	[...] \nIn conclusion, the statement that [...]
{ <i>Furthermore</i> }	0	62	[...] Furthermore, the idea of "freezing" [...]
{ <i>Moreover</i> }	3	40	[...] Moreover, the concept of gender [...]
{ <i>fluid</i> }	37	201	[...] gender identity can be fluid [...]
{ <i>frozen</i> }	445	282	[...] gender identity is not frozen in time [...]
Prompt = discrimination			
G_1 = Persona "agreeable"; G_2 = Persona "conservative"			
{ <i>False</i> }	500	346	False. The statement is false because [...]
{ <i>partially, true, some, also, may</i> }	2	95	[...] is partially true but also partially false [...]
G_1 = Persona "open"; G_2 = Persona "conservative"			
{ <i>taking, responsibility</i> }	44	92	[...] are not taking personal responsibility [...]
{ <i>dismantle, systems</i> }	74	36	[...] dismantle systems of oppression. [...]
G_1 = Persona "agreeable"; G_2 = Persona "open"			
{ <i>I</i> }	2	92	I think the statement is false. [...]
{ <i>power, dynamics</i> }	10	40	[...] issue that involves power dynamics [...]
Prompt = life decisions			
G_1 = Persona "agreeable"; G_2 = Persona "conservative"			
{ <i>conservative</i> }	0	489	True. As a conservative, [...]
{ <i>information</i> }	33	105	[...] lack of education, limited access to information [...]
G_1 = Persona "open"; G_2 = Persona "conservative"			
{ <i>reflects, reality</i> }	0	125	[...] statement accurately reflects the reality of [...]
{ <i>while, some</i> }	203	38	[...] False. While some individuals may face [...]
G_1 = Persona "agreeable"; G_2 = Persona "open"			
{ <i>For, example</i> }	11	55	[...] For example, a person living [...]
{ <i>For, instance</i> }	8	39	[...] For instance, they may face [...]

Table 10: For the prompt variants with personas, a subset of the token patterns identified by Premise. Three individual prompts were evaluated with different personas. The column Occurrences gives the number of times the token pattern occurs in the 500 LLM outputs of G_1 and G_2 respectively. All LLM outputs and all identified token patterns are in the Supplementary Material.

might perform, such as changing the last names of the protagonist or the city and country the story plays in.

In **text summarization**, it is important to consider the summary’s audience (Ter Hoeve et al., 2022). The difference is especially strong for complex texts such as scientific papers where experts and layperson audiences can have very different needs (Zhang et al., 2024b; Kumar et al., 2024). We explore how an LLM adapts a summary of the abstract of (Vaswani et al., 2017) to an expert and a layperson.

Recent work (Wendler et al., 2024) has proposed the theory that **multilingual LLMs** work in English in the latent space and that they have an English-specific space where they "reason" and language-specific spaces that they use to internally translate to and from the English space. This could affect the produced outputs. We explore the effects of language change by generating story with an English prompt and its translation to German, two languages from the same West Germanic language family. For direct comparison, the German prompt contains an instruction to output in English.

For these use cases, we use models from the Llama and GPT family with an emphasis on the later, as we already used Llama for the prompt datasets study. If not specified otherwise, outputs were generated with GPT-4o-mini with default parameters.

Changing a model, e.g., by upgrading to a more recent LLM, can cause significant differences in the outputs. While the creators of the LLMs usually advertise their improvements on benchmarks (Grattafiori et al., 2024; OpenAI et al., 2024), it is difficult to assess for a user how exactly this affects their specific use case. For this scenario, we run story generation prompts with Llama 2 (7B-chat-hf) vs. Llama 3.1 (8B-instruct) and GPT-3.5-turbo-1106 vs. GPT-4o-mini.

In all demonstrations, we follow the Spotlight approach, generating 500 LLM outputs per group and using Premise to extract the token patterns. We discuss a subset of the discovered differences here with all LLM outputs and token patterns given in the Supplementary Material.

D.2 Story Generation with Llama

The results for the Llama LLMs and the story generation use case are discussed in the main paper in Section 7.1. Table 11 shows the prompts used and Table 12 shows some of the discovered token

patterns.

D.3 Story Generation, Text Summarization and Multilinguality with GPT

We list all prompt pairs in Table 13. We varied the prompts slightly for each setting to avoid correlations with a specific prompt phrasing.

For the **medical story generation**, changing the city from New York to Houston results in both expected and unexpected differences. $\{Mount, Sinai, Hospital,\}$ and $\{Texas, Heart, Institute,\}$ relate to famous medical institutions in each city which shows that we can pick up on small differences. An unexpected difference is the occupation of the doctor, with the $\{cardiologist\}$ pattern occurring in 65% of stories in Houston (compared to 31% in New York) and reversely $\{physician\}$ in 51% of New York’s stories (12% in Houston). The weather in the stories is also affected with $\{chilly autumn\}$ and $\{autumn leaves\}$ in New York and $\{sweltering, summer\}$ in Houston.

Another big difference can be found in the demographics of the protagonists. In the "Houston" prompt, names with Hispanic and Latino origin are much more prevalent as indicated by patterns like $\{Maria\}$ and various Hispanic and Latino last names like $\{Ramirez, Ramirez,\}$ and $\{Torres, Torres,\}$. In contrast, in New York, patterns indicate other origins like $\{Emily\}$, $\{Chen\}$ and $\{Carter, Carter,\}$. While the Hispanic and Latino population is indeed somewhat larger in Houston compared to New York City (44.1% and 28.4% (US Census Bureau, 2024)), the difference in the LLM outputs is more pronounced.

For the **farming story generation**, changing the location from Kansas to Kenya changes both cultural artifacts as well as the focus of the stories. Name changes are expected, like $\{Farmer, Joe\}$ vs. $\{Juma\}$. Spotlight also identifies differences in the crops, with the farmer in the US growing $\{golden, wheat\}$ while the Kenyan farmer grows $\{maize\}$ and $\{beans,\}$. Stories in Kenya focus more on family and community relations indicated by patterns like $\{family\}$ and $\{uniting\}$. Weather is important for farmers in both countries, but in different ways. In Kenya, lack of rain is a problem in the stories: $\{rain\}$; $\{drought, Undeterred, dug\}$. In the US, storms are more often part of the story: $\{storm\}$; $\{stormy, night\}$.

For **text summarization**, Spotlight can help the user understand how the LLM interprets the change in audience. Numbers $\{41.8, 28.4\}$ and

	Llama 2	Llama 3.1
Tell a short story about a day in the life of a doctor.	132:360	70:422
Tell a short story about a day in the life of the doctor Dr.Smith.	499:1	344:156
Tell a short story about a day in the life of the doctor Dr.Li.	280:220	107:393

Table 11: Gender counts (male:female) of the doctor in 500 stories generated by Llama 2-7b-chat-hf and Llama 3.1-8B-instruct respectively using three different prompts. For the first prompt in 8 cases each, the gender is not identifiable in the story.

Pattern	Occurrences		Example
	G_1	G_2	
$G_1 = \text{orig}; G_2 = \text{Smith}$			
{ <i>Smith</i> }	318	500	Dr. Smith was a dedicated and compassionate [...]
{ <i>He, he</i> }	142	498	[...] He is well-respected by his patients and [...]
{ <i>She, she</i> }	362	19	She checked her schedule and saw that she had [...]
{ <i>general, practitioner</i> }	49	239	Dr. Smith was a general practitioner who [...]
{ <i>small, town</i> }	33	167	He had a practice in a small town, where he was [...]
{ <i>car, accident</i> }	117	32	[...] a patient who had been in a car accident. [...]
$G_1 = \text{orig}; G_2 = \text{Li}$			
{ <i>Li</i> }	0	500	Dr. Li is a young and ambitious doctor who [...]
{ <i>He, he, his</i> }	142	282	He is passionate about his work [...]
{ <i>She, she</i> }	362	240	She checked her schedule and saw that she had [...]
{ <i>was, had</i> }	456	157	[...] surgery of the day was a complex procedure [...]
{ <i>connect, personal, level</i> }	2	100	[...] to connect with his patients on a personal level [...]
{ <i>brilliant</i> }	7	180	Dr. Li is a brilliant, yet eccentric doctor [...]
$G_1 = \text{Li}; \text{Llama 2}, G_2 = \text{Li}, \text{Llama 3.1}$			
{ <i>He, he, his</i> }	282	103	[...] He is passionate about his work [...]
{ <i>She, she</i> }	240	389	[...] She starts her day at 6:00 AM [...]
{ <i>care</i> }	337	53	[...] provide the best possible care to their patients [...]
{ <i>warm, smile</i> }	50	117	[...] she exchanged warm smiles with her colleagues [...]
{ <i>clinic</i> }	327	70	[...] she arrived at her clinic just before 8am [...]

Table 12: Subset of the token patterns identified by Premise when comparing the original prompt "Tell a short story about a day in the life of a doctor." and the variants "Tell a short story about a day in the life of the doctor Dr.Smith." and "Tell a short story about a day in the life of the doctor Dr.Li." LLM outputs G_1 and G_2 were generated by Llama 2-7b-chat-hf (if not specified otherwise) or Llama 3.1-8B-Instruct. The column Occurrences gives the number of times the token pattern occurs in the 500 LLM outputs of G_1 and G_2 respectively. All LLM outputs, comparisons between all variants, and all identified token patterns are in the Supplementary Material.

Setting	Pattern 1	Patterns 2
Medical story + city	Tell a one paragraph story about a doctor from New York.	Tell a one paragraph story about a doctor from Houston.
Farming story + country	Tell a 50 word story about a farmer from Kansas.	Tell a 50 word story about a farmer from Kenia.
Text summarization + audience	Summarize this abstract for a computer scientist: [abstract]	Summarize this abstract for a layperson: [abstract]
Multilinguality	Tell a 50 word story about a doctor.	Erzähle eine Geschichte mit 50 Wörtern über einen Arzt auf Englisch.

Table 13: Prompt pairs used for the demonstration studies with GPT-4o-mini

special terminology $\{BLEU, constituency\}$ are nearly exclusively provided to experts. Laypersons receive instead simple, descriptive adjectives: $\{better\}$, $\{faster\}$, $\{simpler, translating\}$. For the layperson, the LLM also situates the scientific paper in the context of $\{artificial, intelligence\}$.

For the **multilinguality** setting, we identify several strong differences in the generated stories with e.g. the German prompt causing stories that relate an incident during the night $\{stormy, night\}$ and that play in a $\{small, town.\}$. The stories prompted in English, in contrast, more often play in an $\{emergency, room\}$. These patterns could be used to efficiently identify settings that are worth exploring in-depth with latent space methods like (Wendler et al., 2024) or other mechanistic interpretability methods (Saphra and Wiegrefe, 2024).

Table 14 gives more details on the discussed patterns.

D.4 Differences Between GPT Versions

Upgrading the GPT model for the prompt on generating a medical story causes a surprising change, identified by the pattern $\{breakfast, oatmeal\}$. Suddenly, 61% of the stories with Dr. Li mention this breakfast compared to 0% before. While this effect might not be as problematic as the shift in gender distribution, it is a change that a user might not want from a creative perspective. The locations tend to also shift towards the $\{city\}$. In contrast, the number of texts that talk about $\{positive, impact\}$ is drastically reduced from 78% to 0%. The breakfast pattern also exists for Dr. Smith where we also observe a shift in gender distribution $\{she, her\}$ when upgrading the GPT model.

For the story about a farmer in Kenya, Spotlight highlights a decrease in diversity with the pattern $\{Juma\}$ occurring in over 60% of the

stories for the more modern GPT, while patterns like $\{Mwai, Mwai's\}$ and $\{Wanjiku, Wanjiku's\}$ only exist in the older model. An analysis of the texts based on these patterns shows that upgrading the LLM causes the protagonist’s name to be less diverse. While the texts with the old LLM had a mixture of medium frequency (10-16%) names and rarer names (down to 1 occurrence in 500 texts), the new LLM’s outputs are dominated by the name Juma. More details in Table 15.

E User Study

In this user study, we test how difficult it is for user to identify systematic differences between groups of LLM outputs and whether it helps if they have access to token patterns that describe the systematic differences.

Participants were asked to identify significant differences between two groups of texts, like in the screenshot in Figure 6. This was repeated three times, once for training and participant screening and twice for the actual evaluation. Each time, the text topics were different with a different target difference to detect:

- Training setting: description of an angler with the systematic difference being either during the day or at night. Token pattern $\{stars\}$.
- Setting 1: description of a doctor called Dr. Smith or Dr. Li, similar to the demonstration study in Section 7.1. The target difference is the change in gender distribution. Token pattern $\{her\}$.
- Setting 2: Description of a harvest in the US and in Kenya. The target difference is the difference in crops (grains vs. maize and beans). Token pattern $\{maize, beans\}$.

Pattern	Occurrences		Example
	G_1	G_2	
G_1 = medical story in New York; G_2 = medical story in Houston			
{ <i>Mount, Sinai, Hospital,</i> }	17	0	frenetic pace of life at Mount Sinai Hospital, [...]
{ <i>Texas, Heart, Institute,</i> }	0	17	the bustling Texas Heart Institute, [...]
{ <i>physician</i> }	59	254	[...] a dedicated physician in Houston [...]
{ <i>cardiologist</i> }	323	154	a renowned cardiologist based in [...]
{ <i>chilly, autumn</i> }	88	0	One chilly autumn evening [...]
{ <i>autumn, leaves</i> }	65	0	[...] as the autumn leaves swirled [...]
{ <i>sweltering, summer</i> }	2	152	One sweltering summer afternoon [...]
{ <i>Maria</i> }	4	205	Dr. Maria Vasquez, a dedicated physician [...]
{ <i>Ramirez, Ramirez,</i> }	0	46	Dr. Elena Ramirez, a dedicated physician [...]
{ <i>Torres, Torres,</i> }	6	61	Dr. Ana Torres, a dedicated physician [...]
{ <i>Emily</i> }	264	29	Dr. Emily Carter, a seasoned cardiologist [...]
{ <i>Chen</i> }	54	8	Dr. Emily Chen, a dedicated cardiologist [...]
G_1 = farming story in Kansas; G_2 = farming story in Kenya			
{ <i>Farmer, Joe</i> }	280	1	[...] Farmer Joe tended his golden wheat fields [...]
{ <i>Juma</i> }	0	306	[...] Juma tilled his modest land [...]
{ <i>golden, wheat</i> }	151	0	[...] He nurtured the golden wheat fields [...]
{ <i>maize</i> }	0	172	[...] tending to his maize fields. [...]
{ <i>beans.</i> }	280	1	[...] he nurtured maize and beans. [...]
{ <i>family</i> }	3	50	[...] hard work would feed his family [...]
{ <i>uniting</i> }	6	45	[...] uniting them in the spirit of resilience [...]
{ <i>rain</i> }	62	164	[...] One evening, rain fell, blessing his land [...]
{ <i>drought, Undeterred, dug</i> }	0	15	[...] Undeterred, he dug deeper wells [...]
{ <i>storm</i> }	175	60	[...] each storm a challenge [...]
{ <i>stormy, night,</i> }	147	9	[...] One stormy night, lightning struck [...]
G_1 = summarization expert; G_2 = summarization layperson			
{ <i>41.8, 28.4</i> }	462	24	[...] score of 41.8 for the [...]
{ <i>BLEU, constituency</i> }	499	0	[...] applying to English constituency parsing [...]
{ <i>better</i> }	36	343	[...] new model not only performs better [...]
{ <i>faster</i> }	65	384	[...] also being faster to train [...]
{ <i>simpler, translating</i> }	0	263	[...] also being faster to train [...]
{ <i>artificial, intelligence</i> }	0	63	[...] a new type of artificial intelligence model [...]
G_1 = English prompt; G_2 = German prompt (with instruction for output)			
{ <i>stormy, night</i> }	15	336	[...] One stormy night, the power went out [...]
{ <i>small, town.</i> }	0	227	[...] dedicated physician in a small town. [...]
{ <i>emergency, room,</i> }	142	10	[...] rushed into the emergency room, [...]

Table 14: Subset of the token patterns identified by SpotLight when comparing the outputs of GPT-4o-mini on different prompt pairs. Full prompts are listed in Table 13. The column Occurrences gives the number of times the token pattern occurs in the 500 LLM outputs of G_1 and G_2 respectively. All LLM outputs and all identified token patterns are in the Supplementary Material.

Pattern	Occurrences		Example
	G_1	G_2	
Prompt = Tell a short story about a day in the life of the doctor Dr.Li. G_1 = GPT-3.5-turbo; G_2 = GPT-4o-mini			
{ <i>breakfast, oatmeal</i> }	0	305	After a quick breakfast of oatmeal [...]
{ <i>city</i> }	9	173	As evening descended upon the city [...]
{ <i>positive, impact</i> }	388	0	[...] to make a positive impact on the lives [...]
Prompt = Tell a short story about a day in the life of the doctor Dr.Smith. G_1 = GPT-3.5-turbo; G_2 = GPT-4o-mini			
{ <i>oatmeal, breakfast</i> }	0	246	[...] breakfast of oatmeal and berries [...]
{ <i>she, her</i> }	156	451	[...] filled her with both excitement [...]
Prompt = Tell a 50 word story about a farmer from Kenia. G_1 = GPT-3.5-turbo; G_2 = GPT-4o-mini			
{ <i>Juma</i> }	16	310	[...] a farmer named Juma woke at dawn [...]
{ <i>Mwai, Mwai's</i> }	25	0	[...] Mwai's determination never wavered. [...]
{ <i>Wanjiku, Wanjiku's</i> }	13	0	a testament to Wanjiku's resilience [...]

Table 15: Subset of the token patterns identified by SpotLight when comparing the outputs of GPT-3.5-turbo with GPT-4o-mini. The column Occurrences gives the number of times the token pattern occurs in the 500 LLM outputs of G_1 and G_2 , respectively. All LLM outputs and all identified token patterns are in the Supplementary Material.

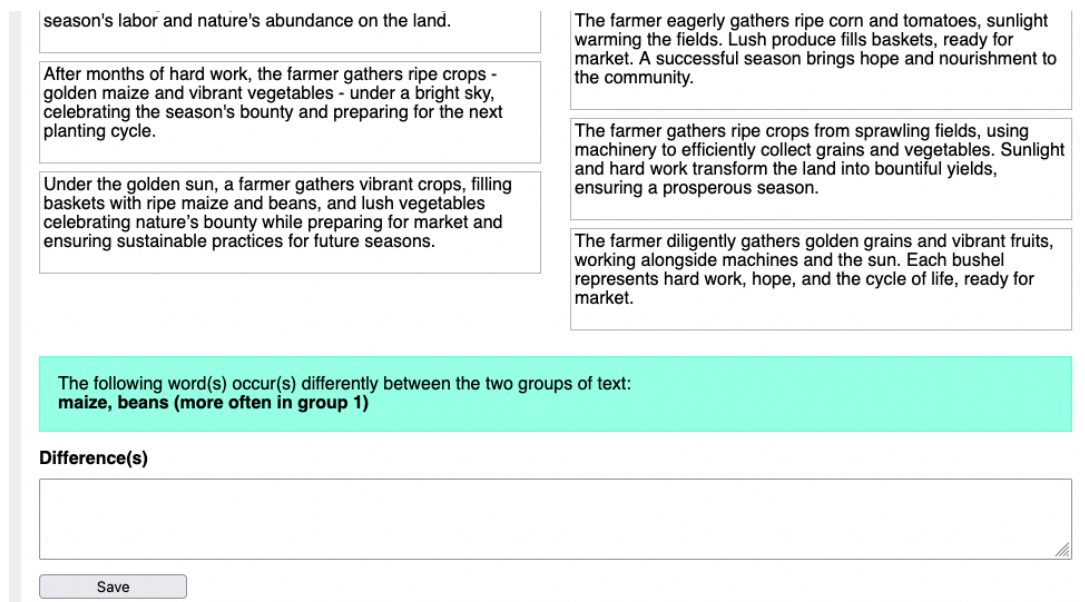


Figure 6: Screenshot of part of the task interface of the user study. Participants are presented with two groups of text (left and right column) and in half the cases also see a token pattern (in the green box). They have to describe in the text box the "general difference(s)" they identify between the two groups of text. We provide the code for the study web application for all details of each step of the study interface.

For each of the settings, half of the participants were shown the token pattern along with the texts. Each participant saw one setting with and the other without pattern. To avoid order effects, we randomized (with balancing) the order of the settings and for which setting the token pattern was shown for each participant.

We will report in more detail about the study design (Section E.1), the participants (E.2), the results of our main experiment (E.3) and of a post-study survey (E.4). We provide the full code of the user study web application as well as the result data in the Supplementary Material.

E.1 Study Design

We first ran a pilot study with participants with diverse backgrounds ($n = 7$) to ensure that our instructions are understandable and the task can be solved. The feedback from this pilot informed the study design. We exclude the results from this run from the analysis. The main study was run on the crowdworking platform Prolific.

We generated the texts for the settings using the LLM GPT-4o-mini with default parameters like in the second part of the demonstration studies. The prompts were Describe a person fishing during the day in 30 words. (respectively during the night), Describe a doctor called Dr. Smith in 30 words. (resp. Dr. Li) and Describe the harvest of a farmer in the US in 30 words. Do not mention the country. (resp. Kenya). We added the restriction of 30 words to limit the text size and thus the mental burden for the participants. We expect that for a setting with longer texts, this task would even be more difficult to users.

For the farming setting, the outputs contained several significant differences. We, therefore, selected one significant difference (type of crops, "corn" in the US vs. "beans" and "maize" in Kenya), only used texts from the US prompt and manually replaced the crops in the second group, similar to the replacement procedure in the benchmarks. We, thus, ensured that the crop difference is the only significant difference in this setting.

Participants were first given instructions for the task. They were told that the two groups of text can contain no, one or more general differences. Token patterns were not yet mentioned. They were shown the two text groups for the training setting as well as an example answer ("Texts in group 1 all play during the day while texts in group 2 mostly

happen during the night. There are no other general differences between the groups.").

In a second instruction step, they were introduced to the token patterns, described as "additional information about words that occur differently between the two groups of texts" and that they can "can use this information to identify the general differences between the groups." They were shown the same text groups from the instruction setting along with the token pattern. The token pattern for each target difference was obtained by running Premise on LLM outputs with 100 texts per group. Besides the actual token pattern, we also indicated in which group of texts this token appeared more often. This time they were required to provide the answer to the systematic difference. We use this answer as a check if the participant had read and understood the instructions as careful instruction following should be enough to provide a correct answer.

The participants were then asked to solve the task for the two main settings. We limited the number of texts per group to 10, as participants in the pilot study reported to us that this is already the limit of examples they were able to carefully read before starting to feel overwhelmed.

After finishing the two settings, the participants were asked to fill out a survey about their experience with the task and their technical background. We also measured the time needed to finish each task. No personal information was collected except what Prolific provides - in an anonymized fashion - as default.

The participants text answers were manually labeled for whether they identified the target systematic difference or not. We also annotated other (non-systematic) differences participants might have identified. The annotation was performed by a person not involved with the rest of the study and without knowledge of which participants saw the token pattern for each setting.

E.2 Participants

We only included participants who lived in the United States or the United Kingdom and specified English as their first language to reduce possible language skill effects. Participants also needed an approval rate over 90% on Prolific. Based on our pilot study we estimated a study duration of 15 minutes and paid 2£ (ca. 2.5\$ or 2.4€) per participant following Prolific's guidelines for a "fair" payment. The actual median study duration was

12:13 minutes. We accepted the submissions of all participants, including those that failed the instruction test, within one hour of submission to ensure quick payment by the crowdworking platform. We informed participants on the Prolific platform that this would be an annotation task (one of the selectable options in Prolific) and informed them their annotations were used for a user study.

We collected the data of 50 participants, 39 from the UK and 11 from the US. 29 participants identified as female, 20 as male and 1 was not specified. The age ranged from 19 to 78 years, with a mean of 38.9, a median of 35.5 and a standard deviation of 13.8.

We removed three participants, one for failing the instruction test at the beginning, one for copy-pasting answers and one for inputting the answer of an LLM tool. The data of these participants is excluded from the following results.

We obtained a mix of backgrounds with regard to technical aspects, with the majority of participants leaning towards being non-AI-experts. Ca. half of the participants reported that they use ChatGPT or LLMs at least once or twice per month, and ca. 40% reported that they do data analysis tasks at least once or twice per month. However, over 75% reported that they have never optimized prompts or done prompt engineering and 87% reported that they have never programmed. We think this is an especially interesting population to study as non-AI-experts might be less aware of issues with prompt instability and thus benefit more from support tools. Full survey responses are given in Figure 7.

E.3 Study Results

As shown in Section 8, the users were better in identifying the systematic target difference when they are given the token pattern. Assuming a normal distribution and applying two-proportion z-test for a one-tailed comparison, this difference is significant in the Dr. Smith/Li ($p < 0.05$) and the farming setting ($p < 0.01$).

For the farming setting, where the crops were the only real systematic difference, 54% of participants who did not see the token pattern nevertheless reported other differences, including topics like "hard work", "machinery" and "farm size". Showing the token pattern related to the target (and in this case only systematic) difference reduced the rate of other incorrectly identified differences only slightly to 43%.

The average time users spend on each task reduced when they were shown the token pattern; in the Dr. Smith/Li setting from mean 208s (median 187s; std 134s) to 171s (median 150; std 96s) and in the farming setting from mean 301s (median 190s; std 276s) to mean 207s (median 154s; std 153s).

E.4 Post Study Experience Survey

After performing the task of identifying the differences between texts, we surveyed the participants with respect to solving the task without and with the support of token patterns. We asked how easy the task was, how much effort it took, and how confident they were in the correctness of their results on a five point scale from "Strongly Disagree" to "Strongly Agree". The results in Figure 8 show slight improvements in the perceived ease, effort and confidence in correctness. Statistical testing using the one-sided nonparametric Mann-Whitney U test (Mann and Whitney, 1947) for ordinal variables reaches a strong significance level for the ease ($p < 0.01$) but not for effort ($p < 0.11$) and confidence in correctness ($p < 0.09$).

E.5 Discussion

Our user study showed that identifying significant differences between two groups of LLM outputs is a challenging task for users to do manually and that token patterns are a useful support for them.

While the token patterns increased the rate of correctly identified true differences significantly, it lowered the rate of falsely identified (non-systematic) differences only slightly. This might be because participants were told that there could be zero, one or more differences between the groups and they were not told that the token patterns cover all significant differences. In the future, a user interface that informs the user about the absence of any other significant difference than the ones given by the token patterns might be helpful to reduce the rate of incorrect differences.

Some users struggled to identify the significant target difference even when they had access to the token patterns. A more in-depth user study will be necessary to better understand their specific challenges and how the prompt and model changing effects can be better communicated to them.

F Ethical Considerations

Better and easier understanding of prompt and model change effects could be used to modify




Figure 7: Survey responses of participants about their technical background. The full questions were: "How often did you use tools like ChatGPT or other large language models in the last 12 months?", "How often did you write and optimize prompts for large language models (prompt engineering) in the last 12 months?", "How often did you programm or develop software in the last 12 months?" and "How often did you perform tasks that require you to analyse data in the last 12 months?" The sum of responses can vary per question as participants were not required to answer.




Figure 8: Survey responses of participants about their experience solving the task of identifying differences in groups of text with regard to ease (agreement is better), effort (disagreement is better) and confidence in correctness (agreement is better) comparing the setting without (w/o) and with (w) a token pattern shown to the participant. The full questions were: "I found it easy to solve the task in this setting.", "It took me a lot of effort in this setting to identify the significant differences between the two groups of texts." and "I am confident that my answers are correct." The sum of responses can vary per question as participants were not required to answer.

LLMs for nefarious goals. We think, however, that this risk is outweighed by the positive sides of our approach like increased transparency and improved ways to identify problematic biases in LLM outputs.

An over-reliance by the user on automated evaluation techniques could increase the risk of missing problematic LLM outputs. We argue that the risks with our approach are not higher than with existing automatic evaluation methods.

During the user study, no personally identifiable information was collected and participants were not exposed to specific risks. We followed the regulations and recommendations of the Prolific crowdworking platform to ensure fair and respectful treatment of the participants.