

# Catch Me If You Can? Not Yet: LLMs Still Struggle to Imitate the Implicit Writing Styles of Everyday Authors

Zhengxiang Wang<sup>1\*</sup> Nafis Irtiza Tripto<sup>2\*</sup> Solha Park<sup>1</sup> Zhenzhen Li<sup>3</sup> Jiawei Zhou<sup>1</sup>

<sup>1</sup>Stony Brook University <sup>2</sup>The Pennsylvania State University <sup>3</sup>Bosch Center for AI  
zhengxiang.wang@stonybrook.edu, nit5154@psu.edu

## Abstract

As large language models (LLMs) become increasingly integrated into personal writing tools, a critical question arises: *can LLMs faithfully imitate an individual’s writing style from just a few examples?* Personal style is often subtle and implicit, making it difficult to specify through prompts yet essential for user-aligned generation. This work presents a comprehensive evaluation of state-of-the-art LLMs’ ability to mimic personal writing styles via in-context learning from a small number of user-authored samples. We introduce an ensemble of complementary metrics—including authorship attribution, authorship verification, style matching, and AI detection—to robustly assess style imitation. Our evaluation spans over 40,000 generations per model across domains such as news, email, forums, and blogs, covering writing samples from more than 400 real-world authors. Results show that while LLMs can approximate user styles in structured formats like news and email, they struggle with nuanced, informal writing in blogs and forums. Further analysis on various prompting strategies such as number of demonstrations reveal key limitations in effective personalization. Our findings highlight a fundamental gap in personalized LLM adaptation and the need for improved techniques to support implicit, style-consistent generation. To aid future research and for reproducibility, we open-source our data and code.<sup>1</sup>

## 1 Introduction

Every time humans write, they leave behind a unique linguistic fingerprint (Svartvik, 1968), a subtle and subconscious reflection of their personal style. To enhance their writing, humans have been using writing tools from spell and grammar checkers to paraphrasers for quite some time (Ferris, 2004). But, the rise of large language models

<sup>\*</sup>Equal contributions.

<sup>1</sup><https://github.com/jaaack-wang/llms-implicit-writing-styles-imitation>.




Figure 1: Task of implicit personalized writing imitation: given few-shot writing examples from an everyday author and a content summary, an LLM is prompted to generate a novel text that imitates the writing examples in styles and reflect the semantics of the summary.

(LLMs) has introduced a powerful, all-in-one assistant capable of drafting, editing, and rephrasing text (Wasi et al., 2024). Despite their impressive fluency, LLMs often default to a *generic* style learned from vast web data, stripping away the personal touch that makes writing feel authentic (Padmakumar and He, 2024). This raises concerns about authorship dilution and being flagged as AI-generated content (Tripto et al., 2024).

To counter these concerns, some systems (e.g., Claude AI<sup>2</sup>) offer configurable parameters like tone, voice, and formality. Yet these controls fall short as users’ personal styles are nuanced and rarely reducible to a few sliders (Kang and Hovy, 2021). A more natural solution is to guide the model with examples of an individual’s prior writing, employing few-shot prompting to emulate an individual’s style more faithfully, like the task shown in Figure 1. While LLMs have shown promise in mimicking public figures (Herbold et al., 2024; Chen and Moscholios, 2024) and fictional characters (Li et al., 2023) whose extensive data footprints on the internet make them easy to model, their ability to replicate the style of everyday users remains largely unexplored. *Can LLMs truly adapt to the personal style of an everyday author with only a handful of casual interactions and no explicitly defined stylistic identities?*

<sup>2</sup><https://claude.ai/>, accessed May 2025.

Existing research on personal-style generation for ordinary users is limited (Bhandarkar et al., 2024; Cho et al., 2025), and even more pressing is the lack of robust evaluation: how do we determine if a generated text from an LLM can belong to the same individual in style and voice? While the concept of authorship is inherently subjective, computational methods like Authorship Attribution (AA) and Authorship Verification (AV)—well-established in forensic linguistics—offer promising proxies. The theory of Linguistic Individuality (Nini, 2023) suggests that a person’s writing style often forms a consistent style model. If LLMs can generate text that fits within this style model while being recognized as from the same author by the two computational authorship models, it brings us closer to achieving truly personalized text generation.

In this work, we examine whether LLMs can imitate the *implicit* writing styles of everyday users through in-context learning, without access to explicit stylistic instructions, via the generation task shown in Figure 1. Our investigation is motivated by and rooted in a *realistic* usage scenario: most users interact with LLMs via standard interfaces rather than custom plug-ins or fine-tuned pipelines (Wang et al., 2024). While methods like personalized fine-tuning (Liu et al., 2024) or advanced prompting strategies (Cho et al., 2025) can improve stylistic alignment, they are often impractical in time-sensitive, real-world settings. Instead, we focus on evaluating the *out-of-the-box* ability of current LLMs to replicate personal writing style by conditioning on a few user-authored samples, reflecting the kind of implicit personalization that is most accessible and scalable in practice.

To this end, we conduct a comprehensive evaluation of state-of-the-art LLMs on the task of *personalized writing imitation*, defined as emulating an individual’s implicit writing style solely from prior writing samples, without explicit stylistic instructions and semantically conditioned on a content summary (see Figure 1). Given the inherent difficulty of measuring personal style, we adopt a diverse suite of computational evaluators, including authorship attribution, authorship verification, linguistically grounded style metrics, and state-of-the-art AI detection tool. This ensemble approach enables a robust and multifaceted assessment of stylistic fidelity and human-like generation.

Using this framework, we evaluate frontier models from leading providers, including OpenAI,

Google, Meta, and DeepSeek, across writing domains such as news articles, emails, online forums, and personal blogs. These domains reflect a range of everyday writing scenarios with distinct stylistic constraints. Our datasets span over 400 authors and build on prior work to ensure diversity and realism. We investigate whether few-shot demonstrations of a user’s prior writing are sufficient for LLMs to generate outputs that match the author’s voice and style under our evaluators.

Results show that while LLMs can partially emulate user style in more structured formats like news and email, they struggle with nuanced, informal expression in domains such as blogs and forums. Generated outputs often default to an average, generic tone and remain readily detectable as AI-written. Moreover, increasing the number of demonstrations offers limited gains in stylistic alignment. These findings underscore the current limitations of in-context personalization and highlight the need for more effective methods to achieve truly personalized generation, an essential step toward democratizing LLM-based writing tools.

## 2 Related Work

Personalization in LLMs spans a wide spectrum, from persona emulation (Tseng et al., 2024) to personalized content recommendations (Lyu et al., 2024). In this study, however, we specifically focus on the subtle challenge of *implicit personalized writing imitation*, whether LLMs can generate text that reflects the inherent implicit style learned from an individual’s prior writings. To situate our work, we review progress at the intersection of LLM-based personalized writing and authorship analysis around LLM generated texts.

**LLMs for Personalized Writings** Recent advances in LLMs have enabled impressive performance in general text style transfer (Toshevskva and Gievska, 2025), with a growing interest in personalized text generation (Zhang et al., 2024). Much of prior work focuses on replicating the style of well-known figures, such as authors (Mikros, 2025; Dinu et al., 2025), celebrities (Herbold et al., 2024; Chen and Moscholios, 2024), or fictional characters (Li et al., 2023; Park et al., 2025), where models benefit from ample training data and rich backstories. However, research on replicating the style of ordinary individuals remains underexplored, especially in terms of a rigorous and systematic evaluation.

Some recent studies have attempted individu-

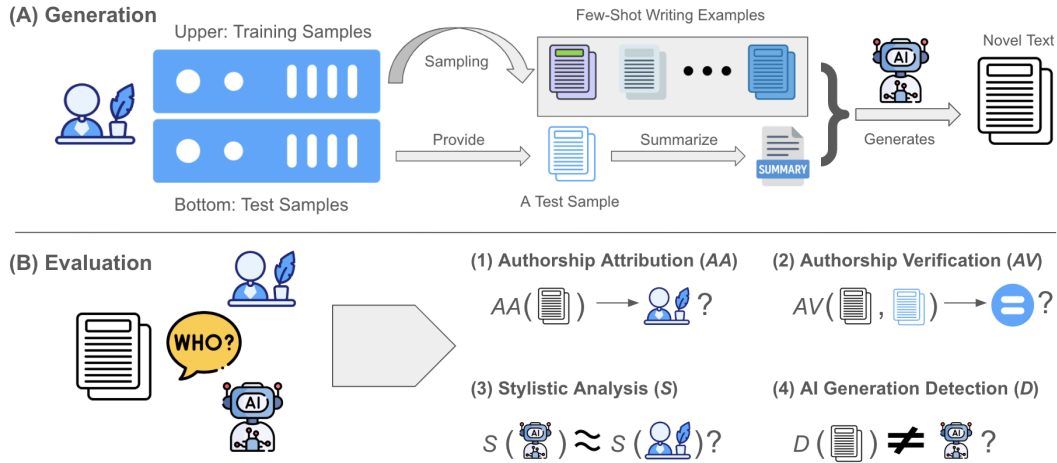


Figure 2: Methodology of our study. (A) We partition each author’s writing samples into training and test sets. We select few-shot examples from training samples to guide an LLM to generate novel text based on a test sample summary. (B) We evaluate the extent to which the LLM imitates an author’s style using four methods: (1) **AA**: Does the model predict LLM-generated samples as written by the target author? (2) **AV**: Does the model predict LLM-generated and human-authored samples as coming from the same author? (3) **Stylistic Analysis**: How similar are the stylistic distributions of LLM-generated and human-authored samples? (4) **AI Generation Detection**: Can LLM-generated samples imitating human writing be detected as non-AI texts?

alized generation using retrieval-augmented generation (RAG) systems (Mysore et al., 2024) or advanced in-context prompting strategies like Trial-Error-Explain (Cho et al., 2025). Yet, these approaches are often costly, impractical for real-time personalization, or dependent on repeated LLM queries. Crucially, these studies usually rely on an LLM-as-a-judge approach (Zheng et al., 2023) to measure stylistic alignment. While LLMs have shown remarkable performance in authorship tasks (Huang et al., 2024), their judgment is still subjective and might contain bias for their generation (Panickssery et al., 2024).

Other evaluations have used metrics like ROUGE or METEOR (Kumar et al., 2024), which assess content overlap but overlook stylistic fidelity. Authorship Attribution (AA) has been used (Bhandarkar et al., 2024), but typically in narrow domains or limited to single-text completions without broader style considerations. Therefore, our study introduces a comprehensive evaluation framework grounded in authorship and stylometric analysis, aiming to systematically assess whether LLMs can truly replicate an individual’s writing style

**Authorship Analysis Tasks** To evaluate whether LLMs can truly replicate individual writing styles, we draw from established authorship analysis tasks. Authorship Attribution (AA) identifies the author

of a given text from a set of candidate authors (Kjell et al., 1994), while Authorship Verification (AV) determines whether two texts share the same author (Halteren, 2007). These tasks traditionally rely on features such as n-grams and stylometric cues, e.g., lexical, syntactic, semantic, and structural patterns (Stamatatos, 2009; Koppel et al., 2012; Juola and Stamatatos, 2013; Neal et al., 2017). While classic ML and DL models have long been used for stylometry, fine-tuned transformer models now represent state-of-the-art AA and AV tasks (Tyo et al., 2022; Tripto et al., 2023). Recently, stylometric analysis has also been employed to detect LLM-generated text (Herbold et al., 2023; Muñoz-Ortiz et al., 2024; Casal and Kessler, 2023), offering a promising path to evaluate whether LLM outputs align with a target individual’s writing style.

### 3 Methodology

This study evaluates whether LLMs can mimic a user’s implicit writing style using only a few prior examples without any explicit stylistic instructions. We consider a realistic scenario, as visualized in Figure 1, where users provide a handful of writing samples and a content summary, reflecting typical few-shot prompting, without describing their style in detail (which may also be challenging if feasible). This setup requires no model customization and offers a scalable path to personalization. While

Dataset	Genre	# Authors	# Samples	Avg Length	AV Acc. (%)	AA Top-5 Acc. (%)	Style Acc. (%)	% Human
Enron	Emails	150	3,884	309	88.9	79.8	79.4	99.2
Blog	Blogs	100	25,224	319	91.4	95.5	81.9	99.5
CCAT50	News articles	50	2,500	584	89.2	94.9	69.1	100
Reddit	Online forums	100	8,451	333	87.7	89.7	73.8	100

Table 1: Four test sets used in our study and their original dataset sources. The corresponding train sets has same “# Authors” and “# Samples” and similar “Avg Length.” The “AV Acc.” and “AA Top-5 Acc.” refer to the average test set accuracy of the corresponding AV and AA models trained on the related train sets. “% Style Acc.” compute the percentage of times test samples of an author closer to that particular author’s style model compared to other authors’ style model. “% Human” is the percentage of times test set samples detected as human generated.

writing styles can evolve, a phenomenon studied in stylochro-metry (Can and Patton, 2004), they primarily focus on long-term changes in the writing of novelists or public figures (Stamou, 2007; Klausner and Vogel, 2015). In our context, we assume that individuals tend to maintain a relatively stable style in everyday writing, making this assumption both practical and justified.

Figure 2 provides an overall visual illustration of our methodology with the rest of the section structured as follows. Section 3.1 defines the task setup, 3.2 presents the evaluation framework, 3.3 describes the datasets, and 3.4 discusses our design choices and limitations.

### 3.1 Task Setup and Notations

We formalize the task of implicit personalized writing imitation and describe the task setup and notations. Given a human author  $a \in A$ , where  $A$  is the set of all authors, let  $T_a$  denote the collection of distinct writing samples available from author  $a$ . We split  $T_a$  into a training set  $T_a^{tr}$  for sampling few-shot writing examples and a test set  $T_a^{te}$  for evaluation. For any sample  $t_a \in T_a^{te}$ , we prompt an LLM  $L$  with  $n$  writing samples  $T_a^{tr'} \subseteq T_a^{tr}$  and a content summary  $S_{t_a}$  of  $t_a$  to generate a new text  $t_a^L$ , or formally,  $L(T_a^{tr'}, S_{t_a}) \rightarrow t_a^L$ . See Figure 2 (A) for an illustration. Here, the summary  $S_{t_a}$  is used to fix the content (or semantics) of the generated text, so that we can focus our evaluation on how well  $t_a^L$  aligns with the author’s writing style in a semantically controlled condition.

We evaluate the implicit personalized writing imitation capabilities of an LLM  $L$  by assessing both the stylistic fidelity of the generated text  $t_a^L$  to the overall writing style of author  $a$  and how closely  $t_a^L$  aligns with the specific test sample  $t_a$  in terms of inherent style.

### 3.2 Personalized Writing Evaluation Framework

To assess the extent to which LLMs can replicate implicit personal writing style, we propose a multi-angle evaluation framework grounded in computational authorship analysis and human-likeness detection, illustrated in Figure 2 (B). This framework contains four components as follows.

- An **authorship attribution model**  $AA(t) \rightarrow a$ , which predicts the most likely author  $a \in A$  for a given text  $t$ . It is trained on human-authored texts (training samples from authors,  $T_a^{tr}$ ) and used to test whether the generated text is correctly attributed to the target author.
- An **authorship verification model**  $AV(t, t') \rightarrow \{0, 1\}$ , which returns 1 if both texts  $t$  and  $t'$  are judged to be written by the same author, and 0 otherwise. This model is applied to compare each generation with its corresponding human-written reference text.
- A **style model**  $X_a$  for each author  $a \in A$ , built from the distributional and stylistic features of their writing samples  $T_a$ . We then compute the stylistic distance between a generated text  $t$  and each author’s style model  $X_a$  to assess stylistic proximity. In contrast to AA and AV models, which capture *implicit* style patterns through fine-tuning, this approach provides an explicit criterion for evaluating whether the generated text mimics an individual’s style.
- An **AI generation detector**  $D(t)$  that classifies whether a text  $t$  appears human-written or AI-generated. We use GPTZero<sup>3</sup> for its strong performance in AI text detection. This complements the other three metrics by assessing human-likeness, providing an additional perspective beyond stylistic fidelity.

<sup>3</sup><https://gptzero.me/>.



Together, these metrics allow us to computationally probe whether LLM-generated text exhibits stylistic consistency with the target author, enabling reliable evaluation of personalized generation.

### 3.3 Data Domains

We use portions of the four datasets listed in Table 1 for our evaluation: Enron (Klimt and Yang, 2004), Blog (Schler et al., 2006), CCAT50 (Liu, 2011), Reddit (Volske et al., 2017). These four datasets present four wildly different genres (emails, blogs, news articles, and online forums) and contain a wide range of average authors (400 in total). In terms of formats, while emails and news articles are more structured and formal, blogs and forums are more casual and informal. These features of the included datasets offer a comprehensive testbed for evaluating LLM’s stylistic mimicry capabilities across genres in a real-world setting.

We perform careful sampling with length control on the original datasets and evenly split the selected samples into train and test sets (See Appendix A for details). Besides length control, our sampling is randomized. This helps ensure the representativeness of the obtained samples, considering the large scale of the sampling.

### 3.4 Discussion of Task Design Rationale

While fine-tuning is often viewed as a reliable method for achieving personalized generation (Tan et al., 2024; Zhang et al., 2024), we intentionally avoid it in our study. First, real-world users typically provide only a limited number of writing samples, which is insufficient for effective fine-tuning. Second, maintaining separate fine-tuned models for each user is costly and impractical at scale (Han et al., 2024).

Finally, complex prompting techniques that involve multiple back-and-forth iterations with the LLMs, such as Cho et al. (2025) and Bhandarkar et al. (2024), would introduce significant latency and API overhead, making them unsuitable for real-time use (Shekhar et al., 2024).

Instead, we adopt a more practical approach grounded in how users typically interact with LLMs via a graphical interface, using only a few prior writing samples stored from earlier interactions. We focus on few-shot prompting, requiring just a single API call, and compare with the zero-shot scenario, used as a baseline to contextualize the effectiveness of few-shot prompting. As emphasized throughout this paper, our primary goal

is to rigorously evaluate how well LLMs can imitate an individual’s writing style, given minimal guidance. While we rely on random sampling for few-shot examples in our main experiments, we also conduct follow-up analyses exploring whether sampling based on topic or length similarity can further enhance stylistic similarity, among others.

## 4 Experimental Setup

In this section, we outline the key experimental choices of our study, with further details provided in Appendix B.

**LLMs and Prompting** We consider both proprietary and open-weights models, including GPT-4o and GPT-4o-mini (OpenAI et al., 2024), Gemini-2.0-Flash (Google DeepMind, 2024), Gemma-3-27B (Team et al., 2025), DeepSeek-V3 (DeepSeek-AI et al., 2025), and Llama-4-Maverick (Meta, 2025). The last two LLMs are state-of-the-art mixture-of-expert models, where DeepSeek-V3 has 671B total parameters with 37B activated for each token, and Llama-4-Maverick has 400B total parameters with 17B active parameters. We always prompt these LLMs with a 0 temperature to maximize output reproducibility.

**Writing Example Sampling** For each test sample, we provide 5 writing examples randomly sampled from the related train set written by the same author, shared across LLMs for stylistic mimicry. We employ random sampling since in principle we have no control over writing examples a user may include for LLM prompting. This ensures a more realistic evaluation.

**Test Samples Summarization** We prompt GPT-4.1 (OpenAI, 2025) to summarize each test sample.

**Baseline** As a baseline, we prompt LLMs with the same test samples zero-shot (with only a content summary). As it is more or less expected that LLMs would perform worse in the style imitation task when not exposed to implicit personal writing style from provided writing examples than with ones. We use this simple baseline to (1) empirically validate this expectation and (2) to sanity check our evaluation framework proposed in Section 3. As a proof of concept, we only run four LLMs zero-shot, excluding DeepSeek and Llama models.

**Evaluation Models** We train AA and AV models using transformer-based encoders, specifically Longformer-base-4096 and ModernBERT-

base (see Appendix B for details). For each author, a style model based on LIWC (Boyd et al., 2022) and WritePrint (Abbasi and Chen, 2008) feature sets is also constructed by extracting distributional stylistic features from their writing samples. We use these three models to evaluate the stylistic fidelity of both human-written (see right below) and LLM-generated texts (see the following sections). In addition, we employ GPTZero as the AI detection tool introduced in Section 3.2 off the shelf.

**Validating the Evaluation Framework on Human Texts** To verify the reliability of our evaluation framework, we first report the evaluation results of the framework on original human-written texts (Table 1). This experiment reflects an ideal condition, where the four evaluators operate solely on human-authored data, allowing us to assess their upper-bound performance.

Across the four datasets (Enron, Blog, CCAT50, and Reddit), the AV models achieve high accuracy, ranging from 87.7% to 91.4%, and the AA models also perform decently, with top-5 accuracy ranging from 79.8% to 95.5%. These results suggest that both models are well fine-tuned and sufficiently reliable for evaluating authorship consistency. In addition, the style model<sup>4</sup> shows consistent alignment between test texts and their corresponding author’s style. The matching accuracy ranges from 69.1% to 81.9%, which indicates how often test samples of an author is closer to that particular author’s style model compared to other authors’ style models. Finally, human-written test samples are barely detected as AI-generated by GPTZero.

Taken together, these results confirm that human-authored texts exhibit personalized and distinguishable stylistic patterns, justifying the effectiveness of our evaluation framework.

## 5 Results

We present a comprehensive analysis of how effectively LLMs mimic personalized style under our evaluation framework. We also report semantic similarity with the original texts to confirm that generated outputs follow the given content while preserving stylistic cues from the examples.

<sup>4</sup>Additionally, an XGBoost classifier trained on these feature representations achieves substantial top-1 authorship attribution accuracy: 46.16% (CCAT50), 43.12% (Enron), 36.04% (Reddit), and 34.11% (Blog), despite the high number of authors (50–150). These results demonstrate the strength of our stylometric framework in capturing individual writing styles.

Model	Setting	CCAT50	Enron	Reddit	Blog
GPT-4o	5-shot	86.64	59.65	27.12	39.39
	0-shot	83.88	28.62	18.77	17.22
GPT-4o-mini	5-shot	87.36	56.45	26.30	38.47
	0-shot	85.84	27.56	18.95	16.10
Gemini-2.0-Flash	5-shot	92.17	59.56	35.59	44.34
	0-shot	87.72	28.93	19.14	19.62
Gemma-3-27B	5-shot	93.30	62.38	30.53	39.38
	0-shot	86.26	27.12	19.20	18.89

Table 2: Average top-5 AA accuracy across LLMs under few-shot (5-shot) and zero-shot prompting (higher score is better). For each generated text  $t'$  and its reference text  $t$ , accuracy is scored 1 if the true author  $a$  appears in the AA model’s top-5 predictions. Scores are averaged across two AA models, with each cell showing the mean accuracy over all test samples for the corresponding dataset, model, and setting.

### 5.1 Authorship Attribution Accuracy

Table 2 reports the average top-5 AA accuracy across datasets and models under 5-shot (default) and 0-shot (baseline) settings. Overall, few-shot prompting consistently outperforms zero-shot, confirming that providing prior samples helps LLMs generate text more stylistically aligned with the target author. Interestingly, even zero-shot generations achieve non-trivial top-5 AA accuracy, largely driven by the topical/content overlap of the writings of that individual. Since authorship depends on both content and style (Sari et al., 2018), our AA models leverage both dimensions, which explains the above-random performance without examples. This underscores the need for other assessments (as we provide in this study) that can provide a holistic evaluation for personalized writing.

Accuracy also varies across individual users. (see Figure 3). While CCAT50 yields consistently high accuracy, likely because journalists cover distinct topics, other datasets show greater variability. We observe higher accuracy in few-shot prompting than in zero-shot consistently for all authors. Notably, Gemini-2.0-Flash and Gemma-3-27B perform comparatively well on informal domains (see Table 2), such as Reddit and blogs, where stylistic cues include characters (\*, #, !) that vary widely across authors. Overall, these findings highlight both dataset-specific and model-specific sensitivities in personalized style replication.

### 5.2 Authorship Verification Accuracy

Table 3 presents (top-1) AV accuracy results for LLM-generated texts across the four datasets. Across all models and datasets, the 5-shot setting consistently outperforms the 0-shot condition in

Dataset	Setting	GPT-4o	GPT-4o-mini	Gemini-2.0-Flash	Gemma-3-27B	DeepSeek-V3	Llama-4-Maverick
CCAT50	5-shot	95.28	95.16	97.44	97.34	97.46	94.68
	0-shot	93.14	93.38	93.94	94.18	-	-
Enron	5-shot	96.15	96.64	96.44	96.17	96.30	95.65
	0-shot	85.02	85.68	85.95	84.98	-	-
Reddit	5-shot	63.65	60.23	65.88	55.54	49.97	65.10
	0-shot	56.66	54.28	53.25	48.75	-	-
Blog	5-shot	19.37	17.93	21.25	16.72	20.77	17.61
	0-shot	8.15	8.22	10.35	9.55	-	-

Table 3: Average AV accuracy across LLMs, averaging over two AV models. Here, we assume that LLM-generated samples and the corresponding test set samples are from the same authors, so the accuracy is the higher the better. The 0-shot baseline is used for proof-of-concept and we exclude DeepSeek and Llama models for cost reasons.




Figure 3: Distribution of per-author AA accuracy averaged across all LLMs under different settings. Highlighted values indicate the mean of each distribution. Overall, few-shot prompting achieves higher per-author accuracy than zero-shot.

AV accuracy. This shows that providing even a few writing examples significantly improves an LLM’s ability to imitate implicit personal writing style. Notably, we also observe variation in AV performance across datasets like the AA results. Models perform particularly well on CCAT50 and Enron, which feature more structured and formal writing. In contrast, performance is generally lower on Reddit and Blog, where writing tends to be more informal and stylistically diverse.

While most LLMs, including Gemini-2.0-Flash, Gemma-3-27B, and DeepSeek-V3, show similarly strong AV performance across datasets, we observe that their effectiveness in mimicking authorial style still depends heavily on prompt design and the stylistic nature of the dataset. These results highlight the sensitivity of LLM-based authorship imitation to both task setup and domain characteristics.

### 5.3 Stylistic Modeling and Alignment

We build individualized style models adopting feature sets from LIWC (Boyd et al., 2022), which integrates linguistic and psychological cues, and WritePrints (Abbasi and Chen, 2008), which contains syntactic-lexical patterns. To assess stylistic

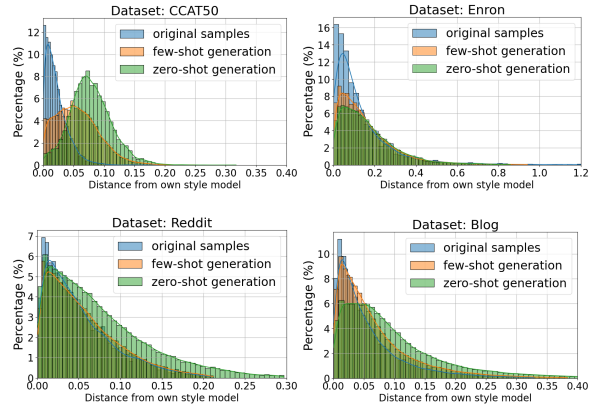


Figure 4: Distribution of the average Mahalanobis distances between each text and the related target author’s style model. Lower distances indicate more explicit stylistic similarity with the target style model.

alignment, we compute the Mahalanobis distance (McLachlan, 1999) between each test sample, whether human-written or LLM-generated, and the corresponding author’s style model.

Figure 4 illustrates the distribution of average distances between each text and the corresponding author’s style model across over all the test samples of each author under both 5-shot and 0-shot settings. As expected, original human samples exhibit the lowest distances. Crucially, few-shot generations are significantly closer to the target style than zero-shot ones, as confirmed by a Wilcoxon signed-rank test (Conover, 1999). It aligns with our hypothesis that without example prompts, LLMs default to a generic style. We also observe variation across datasets, consistent with trends in authorship attribution and verification performance.

### 5.4 AI Generation Detection

Table 4 reports the percentage of generations identified as human by GPTZero. This serves as a proxy for perceived naturalness and stylistic authenticity.

Overall, the percentage of LLM-generated texts

classified as human-written is relatively low across most models and settings (i.e., always below 55% and often less than 20%). This means that while LLMs may generate texts with characteristics of an author, fooling specialized AA and AV models to varying extents, successfully passing as human-like generations remains a challenging task for LLMs.

Interestingly, GPT-4o and GPT-4o-mini show near-zero detection rates. In contrast, Gemini-2.0-Flash and Gemma-3-27B achieve significantly higher detection rates, especially on more structured datasets like Enron and CCAT50. We hypothesize that GPTZero may be optimized for detecting outputs from GPT-based models. As such, differences in detection rates should be interpreted cautiously, as they may reflect detector bias rather than true stylistic similarity. In line with previous sections, zero-shot setting consistently results in lower human-like detection rates across most models and datasets. Moreover, factors such as prompt design, model architecture, and the characteristics of the target domain, may play a role in the observed detection rates.

## 5.5 Semantic Similarity

In our experiments, we provide a content summary in the prompt to guide the semantics of LLM-generated text, so that we can focus on writing style in our evaluation. As a sanity check, we measure the semantic similarity between the LLM-generated texts and the related human texts to see if LLMs do follow instruction to output a semantics-compliant generation. Table 5 shows the similarity scores using METEOR, ROUGE (Kumar et al., 2024) and SBERT (Reimers and Gurevych, 2019).

The scores indicate that LLM-generated texts are semantically similar to human-written ones (with SBERT scores  $\geq 0.74$ ), demonstrating their strong instruction-following capabilities. Additionally, 5-shot prompting consistently produces slightly higher semantic similarity scores across all three metrics, aligning with the improved stylistic fidelity observed in our previous analyses. We hypothesize that higher stylistic fidelity, reflected in elements such as word choice and phrasing, may contribute to higher elevated semantic similarity scores.

## 6 Follow-Up Studies

Everyday users may go beyond providing random examples and a content summary to achieve more personalized writing. Rather than relying on ran-

Model	Setting	CCAT50	Enron	Reddit	Blog
GPT-4o	5-shot	0.08	16.86	0.67	0.47
	0-shot	0.16	0.10	0.27	0.07
GPT-4o-mini	5-shot	0.00	13.76	0.12	0.35
	0-shot	0.08	0.08	0.19	0.12
Gemini-2.0-Flash	5-shot	20.49	54.25	20.77	18.74
	0-shot	0.44	1.88	1.18	0.79
Gemma-3-27B	5-shot	31.36	43.63	16.09	7.58
	0-shot	0.20	2.65	4.80	1.67

Table 4: % of LLM generated texts detected as human text. The higher the number, the better.

Model	Setting	METEOR	ROUGE-L	SBERT
GPT-4o	5-shot	0.32	0.24	0.77
	0-shot	0.30	0.21	0.76
GPT-4o-mini	5-shot	0.31	0.23	0.77
	0-shot	0.29	0.20	0.76
Gemini-2.0-Flash	5-shot	0.32	0.26	0.78
	0-shot	0.31	0.21	0.75
Gemma-3-27B	5-shot	0.30	0.23	0.75
	0-shot	0.29	0.19	0.74

Table 5: Average content similarity scores between LLM-generated and human reference texts.

dom few-shot examples, we hypothesize that carefully chosen samples can better facilitate in-context learning (Kapuriya et al., 2025). Since authorship depends on both content and stylistic factors that vary with text length (Tripto et al., 2025), we select examples based on their content similarity and comparable length to the target text. Additionally, we provide the initial portion of the original text to guide generation (Bhandarkar et al., 2024) and prompt with varying number of examples. To examine these effects, we evaluate on carefully sampled subsets of the test sets using GPT-4o, Gemini-2.0-Flash, and LLaMA-4-Maverick, ensuring cost efficiency and an apples-to-apples comparison. The following subsections present the ablation setups and key findings.

### 6.1 Conditions Considered

**Content Similarity (+Sim ctrl)** We employ BERTopic (Grootendorst, 2022) to cluster each author’s writing, using both train and test samples. For each test sample, we select 5 train set samples that belong to the same cluster as the test sample. We note that this approach that selects few-shot exemplars based on semantic similarity operates in a manner similar to RAG (Mysore et al., 2024).

**Length Alignment (+Len ctrl)** For each test sample, we choose 5 train set samples whose lengths are closest to that test sample.



Dataset	Setting	AV	AA (top-5)	Style Acc.	% Human
<b>CCAT50</b>	5-shot	96.39	89.72	60.11	8.50
	+ Len ctrl	96.83	90.16	57.89	5.67
	+ Sim ctrl	91.95	81.05	61.00	10.33
	+ Snippet <sup>†</sup>	95.45	87.72	64.44	15.50
<b>Enron</b>	5-shot	95.44	69.33	72.11	36.83
	+ Len ctrl	95.89	71.44	61.22	23.67
	+ Sim ctrl	81.28	36.00	70.00	39.67
	+ Snippet <sup>†</sup>	90.56	60.50	68.44	46.50
<b>Reddit</b>	5-shot	68.07	35.43	71.80	10.90
	+ Len ctrl	70.07	35.60	56.60	51.50
	+ Sim ctrl	53.10	16.63	69.33	12.60
	+ Snippet <sup>†</sup>	72.83	36.83	68.27	22.10
<b>Blog</b>	5-shot	19.40	43.93	82.67	9.00
	+ Len ctrl	20.68	43.03	68.60	27.50
	+ Sim ctrl	10.33	22.13	82.65	8.60
	+ Snippet <sup>†</sup>	15.04	38.69	80.39	21.70

Table 6: Follow-up results based on the same subsets of test sets used in Section 6. The reported results are averaging over the three LLMs. The + Snippet<sup>†</sup> results are evaluated with only the LLM continued generations, not including the initially provided author snippets.

**Exemplar Quantity** We consider different numbers of writing examples while prompting: 2, 4, 6, 8, and 10. To ensure comparability, each smaller set is a strict subset of the next larger set.

**Snippet Inclusion** (+Snippet) Following [Bhandarkar et al. \(2024\)](#), we augment the prompt with an initial excerpt from the target test sample—the first 50 words or 20% of the text (whichever is shorter)—besides the default 5-shot examples.

## 6.2 Major Observations

Table 6 presents our evaluation framework results averaged over the three LLMs under various prompt configurations. The baseline is the 5-shot setting, in which five writing examples are randomly selected from the same author.

Content-based exemplar selection (+Sim ctrl) surprisingly reduces attribution performance, especially in Enron, Reddit, and Blog. While topical alignment improves, restricting exemplars to a narrow cluster appears to diminish stylistic diversity, making it harder for models to capture author-specific cues.

Length alignment (+Len ctrl) yields modest gains in attribution (notably for CCAT50 and Reddit), but it lowers style-model accuracy. Matching exemplar length seems to benefit surface-level consistency captured by AA/AV, yet inadvertently reduces variation that signals stylistic fidelity.

Snippet inclusion (+Snippet) yields the strongest boost in perceived human-likeness across datasets, with detectors more often classifying outputs as




Figure 5: AA/AV accuracy as a function of the number of writing examples.

human-written. This highlights the power of seeding generations with authentic text fragments, even if attribution metrics show mixed results.

Interestingly, Figure 5 shows that including more writing examples in the prompt affects the four metrics very little, suggesting limited gains in stylistic alignment. This is in line with the style model and AI detection results, which we include in Figure 6 in Appendix D for space reasons.

Overall, these findings underscore that exemplar selection is far from trivial: strategies optimized for content or length do not always enhance stylistic imitation, and no single configuration consistently excels across all metrics. Effective personalization thus requires domain-sensitive choices and careful balancing between semantic control, stylistic fidelity, and human-likeness.

## 7 Conclusion

This paper presents a comprehensive evaluation of state-of-the-art LLMs on their ability to mimic the implicit writing styles of everyday users through few-shot in-context learning. By combining authorship attribution, verification, stylometric modeling, and AI generation detection, across four diverse datasets, we provide strong empirical evidence that, despite improvements from exemplar-based prompting, current LLMs still struggle to reproduce nuanced personal styles—especially in informal and stylistically diverse domains. Our analysis further shows that prompt design choices, such as length alignment and content similarity, moderately affect stylistic fidelity, but do not close the personalization gap. These findings highlight fundamental limitations in the stylistic adaptability of LLMs and suggest that achieving truly personalized generation remains an open challenge. Future work should explore richer personalization signals and hybrid prompting and/or finetuning strategies to better capture the subtleties of individual writing styles in real-world settings.

## Acknowledgments

Zhengxiang Wang was supported by Junior Researcher Award from the Institute for Advanced Computational Science (IACS) at Stony Brook University for the academic year of 2025. IACS also provided free OpenAI GPT access for our experiments. We thank the Google Gemma Academic Program for their partial support of Jiawei Zhou and for providing computational resources.

## Limitations

Our study presents several limitations. First, while we propose a comprehensive computational evaluation framework, we do not include large-scale human evaluations. Manual assessments may reveal aspects of stylistic fidelity or perceived authorship not captured by automated metrics.

Second, our analysis is constrained to within-genre personalization; that is, the writing genre of the prompt examples matches that of the generated text. Cross-genre style transfer remains unexplored due to the lack of multi-genre corpora authored by the same individual.

Third, the generality of our findings is bounded by the four English-language datasets used, which, although diverse in domain, may not capture the full variability of real-world writing.

Finally, our authorship analysis models are trained on relatively small sample sizes per author, which could limit their robustness, especially for authors with minimal stylistic consistency.

Future work should incorporate broader linguistic and demographic diversity, include human evaluations, and explore personalization across modalities and genres.

## Ethical Considerations

**Misuse Potential** Techniques for mimicking writing style may be misused for impersonation, academic dishonesty, or phishing. While our work aims to support personalized assistance, it also highlights risks associated with stylistic imitation.

**Privacy** We use publicly available, anonymized datasets, but writing style can still carry identifiable traits. Care should be taken when using or releasing data for personalized modeling tasks.

**Bias and Generalization** Our datasets are English-only and limited in demographic diversity. As a result, findings may not generalize to broader

populations or multilingual settings, raising fairness concerns.

**Evaluation Assumptions** We rely on computational authorship models, which may not fully reflect human perceptions of style and can encode unintended biases. These tools should be used cautiously in sensitive applications.

**Disclosure and Transparency** As LLM outputs increasingly resemble human writing, clear disclosure of AI assistance is essential to maintain transparency and trust.

## References

- Ahmed Abbasi and Hsinchun Chen. 2008. Writeprints: A stylometric approach to identity-level identification and similarity detection in cyberspace. *ACM Transactions on Information Systems (TOIS)*, 26(2):1–29.
- Iz Beltagy, Matthew E. Peters, and Arman Cohan. 2020. *Longformer: The long-document transformer*. Preprint, arXiv:2004.05150.
- Avanti Bhandarkar, Ronald Wilson, Anushka Swarup, and Damon Woodard. 2024. Emulating author style: a feasibility study of prompt-enabled text stylization with off-the-shelf llms. In *Proceedings of the 1st Workshop on Personalization of Generative AI Systems (PERSONALIZE 2024)*, pages 76–82.
- Ryan L Boyd, Ashwini Ashokkumar, Sarah Seraj, and James W Pennebaker. 2022. The development and psychometric properties of liwc-22. *Austin, TX: University of Texas at Austin*, 10.
- Fazli Can and Jon M Patton. 2004. Change of writing style with time. *Computers and the Humanities*, 38(1):61–82.
- J Elliott Casal and Matthew Kessler. 2023. Can linguists distinguish between chatgpt/ai and human writing?: A study of research ethics and academic publishing. *Research Methods in Applied Linguistics*, 2(3):100068.
- Ziyang Chen and Stylios Moscholios. 2024. Using prompts to guide large language models in imitating a real person’s language style. *arXiv preprint arXiv:2410.03848*.
- Hyundong Justin Cho, Karishma Sharma, Nicolaas Paul Jedema, Leonardo F. R. Ribeiro, Jonathan May, and Alessandro Moschitti. 2025. *Tuning-free personalized alignment via trial-error-explain in-context learning*. In *Findings of the Association for Computational Linguistics: NAACL 2025*, pages 5864–5885, Albuquerque, New Mexico. Association for Computational Linguistics.
- William Jay Conover. 1999. *Practical nonparametric statistics*, volume 350. john wiley & sons.

- DeepSeek-AI, Aixin Liu, Bei Feng, Bing Xue, Bingxuan Wang, Bochao Wu, Chengda Lu, Chenggang Zhao, Chengqi Deng, Chenyu Zhang, Chong Ruan, Damai Dai, Daya Guo, Dejian Yang, Deli Chen, Dongjie Ji, Erhang Li, Fangyun Lin, Fucong Dai, and 181 others. 2025. *Deepseek-v3 technical report. Preprint*, arXiv:2412.19437.
- Anca Dinu, Andra-Maria Florescu, and Liviu P Dinu. 2025. Analyzing large language models’ pastiche ability: a case study on a 20th century romanian author. In *Proceedings of the 5th International Conference on Natural Language Processing for Digital Humanities*, pages 20–32.
- Dana R Ferris. 2004. The “grammar correction” debate in l2 writing: Where are we, and where do we go from here?(and what do we do in the meantime. . .?). *Journal of second language writing*, 13(1):49–62.
- Google DeepMind. 2024. Introducing gemini 2.0: our new ai model for the agentic era. <https://blog.google/technology/google-deepmind/google-gemini-ai-update-december-2024/>. Accessed: 2025-04-30.
- Maarten Grootendorst. 2022. *Bertopic: Neural topic modeling with a class-based tf-idf procedure. Preprint*, arXiv:2203.05794.
- Hans Van Halteren. 2007. Author verification by linguistic profiling: An exploration of the parameter space. *ACM Transactions on Speech and Language Processing (TSLP)*, 4(1):1–17.
- Zeyu Han, Chao Gao, Jinyang Liu, Jeff Zhang, and Sai Qian Zhang. 2024. Parameter-efficient finetuning for large models: A comprehensive survey. *arXiv preprint arXiv:2403.14608*.
- Steffen Herbold, Annette Hautli-Janisz, Ute Heuer, Zlata Kikteva, and Alexander Trautsch. 2023. A large-scale comparison of human-written versus chatgpt-generated essays. *Scientific reports*, 13(1):18617.
- Steffen Herbold, Alexander Trautsch, Zlata Kikteva, and Annette Hautli-Janisz. 2024. Large language models can impersonate politicians and other public figures. *arXiv preprint arXiv:2407.12855*.
- Baixiang Huang, Canyu Chen, and Kai Shu. 2024. Can large language models identify authorship? In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 445–460.
- Patrick Juola and Efstathios Stamatatos. 2013. Overview of the author identification task at pan 2013. *CLEF (Working Notes)*, 1179.
- Dongyeop Kang and Eduard Hovy. 2021. Style is not a single variable: Case studies for cross-stylistic language understanding. In *Proceedings of the 59th annual meeting of the association for computational linguistics and the 11th international joint conference on natural language processing (volume 1: long papers)*, pages 2376–2387.
- Janak Kapuriya, Manit Kaushik, Debasis Ganguly, and Sumit Bhatia. 2025. Exploring the role of diversity in example selection for in-context learning. In *Proceedings of the 48th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 2962–2966.
- Bradley Kjell, W Addison Woods, and Ophir Frieder. 1994. Discrimination of authorship using visualization. *Information processing & management*, 30(1):141–150.
- Carmen Klaussner and Carl Vogel. 2015. Stylochronometry: Timeline prediction in stylometric analysis. In *International Conference on Innovative Techniques and Applications of Artificial Intelligence*, pages 91–106. Springer.
- Bryan Klimt and Yiming Yang. 2004. The enron corpus: A new dataset for email classification research. In *European conference on machine learning*, pages 217–226. Springer.
- Moshe Koppel, Jonathan Schler, Shlomo Argamon, and Yaron Winter. 2012. The “fundamental problem” of authorship attribution. *English Studies*, 93(3):284–291.
- Ishita Kumar, Snigdha Viswanathan, Sushrita Yerra, Alireza Salemi, Ryan A Rossi, Franck Dernoncourt, Hanieh Deilamsalehy, Xiang Chen, Ruiyi Zhang, Shubham Agarwal, and 1 others. 2024. Longlamp: A benchmark for personalized long-form text generation. *arXiv preprint arXiv:2407.11016*.
- Cheng Li, Ziang Leng, Chenxi Yan, Junyi Shen, Hao Wang, Weishi Mi, Yaying Fei, Xiaoyang Feng, Song Yan, HaoSheng Wang, and 1 others. 2023. Chatharuhi: Reviving anime character in reality via large language model. *arXiv preprint arXiv:2308.09597*.
- Xinyue Liu, Harshita Diddee, and Daphne Ippolito. 2024. Customizing large language model generation style using parameter-efficient finetuning. In *Proceedings of the 17th International Natural Language Generation Conference*, pages 412–426.
- Zhi Liu. 2011. Reuter 50-50 dataset. National Engineering Research Center for E-Learning Technology, China.
- Hanjia Lyu, Song Jiang, Hanqing Zeng, Yinglong Xia, Qifan Wang, Si Zhang, Ren Chen, Chris Leung, Jiajie Tang, and Jiebo Luo. 2024. Llm-rec: Personalized recommendation via prompting large language models. In *Findings of the Association for Computational Linguistics: NAACL 2024*, pages 583–612.
- Goeffrey J McLachlan. 1999. Mahalanobis distance. *Resonance*, 4(6):20–26.
- Meta. 2025. The llama 4 herd: The beginning of a new era of natively multimodal ai innovation. <https://ai.meta.com/blog/llama-4-multimodal-intelligence/>. Accessed: 2025-05-11.

- George Mikros. 2025. Beyond the surface: stylometric analysis of gpt-4o’s capacity for literary style imitation. *Digital Scholarship in the Humanities*, page fqaf035.
- Alberto Muñoz-Ortiz, Carlos Gómez-Rodríguez, and David Vilares. 2024. **Contrasting linguistic patterns in human and llm-generated news text**. *Artificial Intelligence Review*, 57.
- Sheshera Mysore, Zhuoran Lu, Mengting Wan, Longqi Yang, Bahareh Sarrafzadeh, Steve Menezes, Tina Baghaee, Emmanuel Gonzalez, Jennifer Neville, and Tara Safavi. 2024. Pearl: Personalizing large language model writing assistants with generation-calibrated retrievers. In *Proceedings of the 1st Workshop on Customizable NLP: Progress and Challenges in Customizing NLP for a Domain, Application, Group, or Individual (CustomNLP4U)*, pages 198–219.
- Tempestt Neal, Kalaivani Sundararajan, Aneez Fatima, Yiming Yan, Yingfei Xiang, and Damon Woodard. 2017. Surveying stylometry techniques and applications. *ACM Computing Surveys (CSuR)*, 50(6):1–36.
- Andrea Nini. 2023. *A theory of linguistic individuality for authorship analysis*. Cambridge University Press.
- 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, and 401 others. 2024. **Gpt-4o system card**. *Preprint*, arXiv:2410.21276.
- OpenAI. 2025. Introducing gpt-4.1 in the api. <https://openai.com/index/gpt-4-1/>. Accessed: 2025-05-11.
- Vishakh Padmakumar and He He. 2024. Does writing with language models reduce content diversity? In *The Twelfth International Conference on Learning Representations*.
- Arjun Panickssery, Samuel Bowman, and Shi Feng. 2024. Llm evaluators recognize and favor their own generations. *Advances in Neural Information Processing Systems*, 37:68772–68802.
- Jeiyoon Park, Chanjun Park, and Heui-Seok Lim. 2025. Charactergpt: A persona reconstruction framework for role-playing agents. In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 3: Industry Track)*, pages 287–303.
- Nils Reimers and Iryna Gurevych. 2019. **Sentence-BERT: Sentence embeddings using Siamese BERT-networks**. 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 3982–3992, Hong Kong, China. Association for Computational Linguistics.
- Yunita Sari, Mark Stevenson, and Andreas Vlachos. 2018. Topic or style? exploring the most useful features for authorship attribution. In *Proceedings of the 27th international conference on computational linguistics*, pages 343–353.
- Jonathan Schler, Moshe Koppel, Shlomo Argamon, and James W Pennebaker. 2006. Effects of age and gender on blogging. In *AAAI spring symposium: Computational approaches to analyzing weblogs*, volume 6, pages 199–205.
- Shivanshu Shekhar, Tanishq Dubey, Koyel Mukherjee, Apoorv Saxena, Atharv Tyagi, and Nishanth Kotla. 2024. Towards optimizing the costs of llm usage. *arXiv preprint arXiv:2402.01742*.
- Efstathios Stamatatos. 2009. A survey of modern authorship attribution methods. *Journal of the American Society for information Science and Technology*, 60(3):538–556.
- Constantina Stamou. 2007. Stylochronometry: Stylistic development, sequence of composition, and relative dating. *Literary and Linguistic Computing*, 23(2):181–199.
- Jan Svartvik. 1968. *The Evans Statements: A Case for Forensic Linguistics*. Acta Universitatis Gothoburgensis, Gothenburg. 24 cm.
- Zhaoxuan Tan, Qingkai Zeng, Yijun Tian, Zheyuan Liu, Bing Yin, and Meng Jiang. 2024. Democratizing large language models via personalized parameter-efficient fine-tuning. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 6476–6491.
- Gemma Team, Aishwarya Kamath, Johan Ferret, Shreya Pathak, Nino Vieillard, Ramona Merhej, Sarah Perrin, Tatiana Matejovicova, Alexandre Ramé, Morgane Rivière, Louis Rouillard, Thomas Mesnard, Geoffrey Cideron, Jean bastien Grill, Sabela Ramos, Edouard Yvinec, Michelle Casbon, Etienne Pot, Ivo Penchev, and 197 others. 2025. **Gemma 3 technical report**. *Preprint*, arXiv:2503.19786.
- Martina Toshevska and Sonja Gievska. 2025. Llm-based text style transfer: Have we taken a step forward? *IEEE Access*.
- Nafis Tripto, Adaku Uchendu, Thai Le, Mattia Setzu, Fosca Giannotti, and Dongwon Lee. 2023. Hansen: Human and ai spoken text benchmark for authorship analysis. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 13706–13724.
- Nafis Irtiza Tripto, Saranya Venkatraman, Dominik Macko, Robert Moro, Ivan Srba, Adaku Uchendu, Thai Le, and Dongwon Lee. 2024. A ship of theseus: Curious cases of paraphrasing in llm-generated texts. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 6608–6625.



- Nafis Irtiza Tripto, Saranya Venkatraman, Mahjabin Nahar, and Dongwon Lee. 2025. Beyond checkmate: exploring the creative chokepoints in ai text. *arXiv preprint arXiv:2501.19301*.
- Yu-Min Tseng, Yu-Chao Huang, Teng-Yun Hsiao, Wei-Lin Chen, Chao-Wei Huang, Yu Meng, and Yun-Nung Chen. 2024. Two tales of persona in llms: A survey of role-playing and personalization. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 16612–16631.
- Jacob Tyo, Bhuwan Dhingra, and Zachary C Lipton. 2022. On the state of the art in authorship attribution and authorship verification. *arXiv e-prints*, pages arXiv–2209.
- Michael V"olske, Martin Potthast, Shahbaz Syed, and Benno Stein. 2017. [TL;DR: Mining Reddit to learn automatic summarization](#). In *Proceedings of the Workshop on New Frontiers in Summarization*, pages 59–63, Copenhagen, Denmark. Association for Computational Linguistics.
- Jiayin Wang, Weizhi Ma, Peijie Sun, Min Zhang, and Jian-Yun Nie. 2024. Understanding user experience in large language model interactions. *arXiv preprint arXiv:2401.08329*.
- Benjamin Warner, Antoine Chaffin, Benjamin Clavié, Orion Weller, Oskar Hallström, Said Taghadouini, Alexis Gallagher, Raja Biswas, Faisal Ladhak, Tom Aarsen, Nathan Cooper, Griffin Adams, Jeremy Howard, and Iacopo Poli. 2024. [Smarter, better, faster, longer: A modern bidirectional encoder for fast, memory efficient, and long context finetuning and inference](#). *Preprint*, arXiv:2412.13663.
- Azmine Toushik Wasi, Mst Rafia Islam, and Raima Islam. 2024. Llms as writing assistants: Exploring perspectives on sense of ownership and reasoning. In *Proceedings of the Third Workshop on Intelligent and Interactive Writing Assistants*, pages 38–42.
- Zhehao Zhang, Ryan A Rossi, Branislav Kveton, Yijia Shao, Diyi Yang, Hamed Zamani, Franck Dernoncourt, Joe Barrow, Tong Yu, Sungchul Kim, and 1 others. 2024. Personalization of large language models: A survey. *arXiv preprint arXiv:2411.00027*.
- Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric P. Xing, Hao Zhang, Joseph E. Gonzalez, and Ion Stoica. 2023. Judging llm-as-a-judge with mt-bench and chatbot arena. In *Proceedings of the 37th International Conference on Neural Information Processing Systems, NIPS '23*, Red Hook, NY, USA. Curran Associates Inc.

## A Train and Test split Preparation for Original Study

As described in Section 3.3, we maintain separate train and test portions for each dataset. The train portion is used to train the AA and AV models, build style models for individual authors, and sample examples for in-context learning during LLM prompting. The test portion is reserved for evaluating AA, AV, and style models, as well as for conducting our main study (few-shot and zero-shot experiments).

For **CCAT50 (Reuters)**, we adopt the original train/test split from the source paper (Liu, 2011).

For the **Blog** dataset (Schler et al., 2006), due to its large size and variability in text length, we restrict samples to 100–1500 words and select the top 100 authors by sample count. For each author, 50% of samples are used for training and 50% for testing, ensuring comparable distributions across length, topic, age, and gender demographics.

For the Enron dataset, we follow the preprocessing steps outlined in (Tripto et al., 2025), which include retaining only one-to-one internal emails to preserve a personal tone, excluding automated/forwarded/bulk messages, and removing emails with attachments. We further filter by length (100–1500 words), select the top 150 authors, and split the samples evenly into training and testing sets.

For **Reddit** dataset, where style can vary widely across subreddits (Volske et al., 2017), we apply similar filtering (100–1500 words, top 100 authors by post count). To mitigate subreddit-specific bias, we stratify the train/test split such that each author’s samples maintain approximately equal subreddit distributions.

This consistent preprocessing ensures fair comparisons across datasets while balancing sample size, author coverage, and stylistic variation.

## B Experimental Details on AA and AV Model Training

**Hyperparameters** Table 7 summarizes the training configurations used for the Authorship Verification (AV) and Authorship Attribution (AA) models. We trained both longformer-base-4096 (Beltagy et al., 2020) and ModernBert-base (Warner et al., 2024) for these two types of models.

**Datasets** For AA model training and evaluation, we used the train/test splits of the four datasets cre-

Category	Training Parameter	AV Model	AA Model
General	Number of training epochs	10	20
	Train batch size	8	8
	Eval batch size	16	16
	Max sequence length	2048	2048
Optimization	Learning rate	2e-5	2e-5
	Weight decay	0.01	0.01
	Warmup steps	500	500
	Gradient accumulation steps	4	4
Evaluation	Evaluation strategy	epoch	epoch
	Early stopping patience	3	3
	Load best model at end	True	True
	Metric for best model	F1	Eval loss
	Greater is better	True	False
Precision	Mixed precision (fp16)	True	True

Table 7: Training hyperparameters for Authorship Verification (AV) and Authorship Attribution (AA) models.

ated for the purpose of evaluating LLM’s writing style mimicking capabilities. For AV model training and evaluation, we construct label-balanced (pos:neg = 4:6) AV datasets by sampling text pairs from the original train/test splits of the four datasets. For training AV/AA models, we split 20% of the train set for validation.

## C Elaborations on our Result Reporting

Our primary goal is to present a large-scale, systematic evaluation. Given the scale of our study—over 40,000 generations per model across 400 authors—we used the average results to represent the expected performance of each LLM in implicit writing style imitation. While it would have been ideal to report the four metrics on the basis of each author, it is not practical to do so, given the large number of authors we have for experiments. Moreover, it is beyond the scope of the current study to examine the performance variations of LLMs in mimicking different individuals.

That said, we did include per-author AA accuracy across different datasets in Figure 3 and paired statistical tests (i.e., Wilcoxon signed-rank test) for the style model results under different prompting conditions reported in Figure 4. This only not enriches our analysis, but also reconfirms the observations we made from the results averaging over authors from the four datasets.

Lastly, to minimize generation variability and thus maximizes reproducibility, we used greedy decoding (temperature = 0) throughout all experiments.




Figure 6: Style model accuracy and percentage of human detection rates as a function of the number of writing examples.

## D Follow-Up Studies

**Subset Sampling** To ensure apples-to-apples comparison, we use the same subsets of the four datasets across all the follow-up studies. For each dataset, we sampled 10 test samples for each author. The numbers of authors for CCAT50, Enron, Reddit, and Blog are 30, 30, 50, 50, respectively. We chose these numbers based on the BERTopic clustering results to make sure that each selected test sample had at least 5 train samples from the related author.

**Additional Result** Figure 6 shows the effect of including more writing examples in the prompt on the style model and AI detection results.

## E Prompts

This section provides all the prompt templates used in this study. We use “\$” to denote placeholder.

### E.1 Prompt Template for Test Sample Summarization

We used GPT-4.1 (OpenAI, 2025) to summarize each sample in the test splits of the four datasets.

You will be given a piece of text. Your task is to summarize the text in a concise and clear manner, capturing the main ideas and key points while maintaining the original meaning.

### Text to Summarize

\$text

### Instructions

- Provide a summary that is brief yet comprehensive.
- Ensure that the summary accurately reflects the content of the original text.
- Avoid adding any personal opinions or interpretations.
- Do not output anything other than the summary.

Begin your response below:

### E.2 Few-Shot Prompting for LLM Writing Generation

Besides 5-shot prompting in the main experiments in Section 5, we also re-use the following prompt templates in two of the follow-up studies in Section 6. They include Content Similarity (+Sim ctrl) and Length Control (+Len ctrl).

You will be given one or more writing samples from a specific author. Your task is to analyze the author’s style, tone, and voice, then craft a new piece of \$genre that closely mimics their writing based on a provided summary. Your writing should be around \$num\_words words.

### Author’s Writing Sample(s)

\$writing\_samples

### Writing Task Summary

\$summary

### Instructions

- Ensure your writing faithfully replicates the author’s style, including tone, word choices, and sentence structure, etc.
- Maintain consistency with the author’s voice while accurately reflecting the details of the given summary.
- Strive to make your writing indistinguishable from the original author’s work.
- Do not output anything other than the writing.

Begin your response below:

### E.3 Zero-Shot Prompting for LLM Writing Generation

We use zero-shot prompting only in the main experiments in Section 5, as a baseline to the default 5-shot prompting.

Given the following summary, your task is to generate a writing sample around \$num\_words words. The genre of the writing is \$genre. Do not output anything other than the writing.

### Writing Task Summary

\$summary

Begin your response below:

### E.4 Few-Shot Prompting Plus Snippet for LLM Writing Generation

We include a text snippet, the first 50 words or 20% of the text (whichever is shorter) into the default few-shot prompting illustrated in Appendix E.2.

You will be given one or more writing samples from a specific author plus a text snippet of \$genre from the same author. Your task is

to analyze the author's style, tone, and voice, then generate a continuation for the provided human-authored text snippet with around \$num\_words words that closely mimics their writing based on a provided summary.

### Author's Writing Sample(s)

\$writing\_samples

### Writing Task Summary

\$summary

### Human-Authored Text Snippet

\$snippet

### Instructions

- Ensure your writing faithfully replicates the author's style, including tone, word choices, and sentence structure, etc.
- Maintain consistency with the author's voice while accurately reflecting the details of the given summary.
- Strive to make your writing indistinguishable from the original author's work.
- Do not output anything other than the writing.

Begin your response below: