

Gendered Prompting and LLM Code Review: How Gender Cues in the Prompt Shape Code Quality and Evaluation

Lynn Janzen¹, Üveys Eroglu¹, Dorothea Kolossa¹, Pia Knöferle³,
Sebastian Möller¹, Vera Schmitt², Veronika Solopova^{1,2}

¹Technische Universität Berlin, Germany

²Johannes Gutenberg-Universität Mainz, Germany

³Humboldt-Universität zu Berlin, Germany

veronika.solopova@tu-berlin.de

Abstract

LLMs are increasingly embedded in programming workflows, from code generation to automated code review. Yet, how gendered communication styles interact with LLM-assisted programming and code review remains underexplored. We present a mixed-methods pilot study examining whether gender-related linguistic differences in prompts influence code generation outcomes and code review decisions. Across three complementary studies, we analyze (i) collected real-world coding prompts, (ii) a user study, in which developers solve identical programming tasks with LLM assistance, and (iii) an LLM-based simulated evaluation framework that systematically varies gender-coded prompt styles and reviewer personas. We find that gender-related differences in prompting style are subtle but measurable, with female-authored prompts exhibiting more indirect and involved language, which does not translate into consistent gaps in functional correctness or static code quality. For LLM code review, in contrast, we observe systematic biases: on average, models approve female-authored code more, despite comparable quality. Our experiments show that gender-coded prompt styles influence code length and maintainability. We also observe substantial variation in reviewer behavior across models. Together, these findings suggest that fairness risks in LLM-assisted programming arise less from code generation accuracy than from LLM-based evaluation, particularly as LLMs are increasingly used as automated code reviewers.

1 Introduction

Women remain significantly underrepresented in computer science and software engineering worldwide, with only 23% of female developers globally in 2023 (SlashData, 2023), and around 5% of core developers and pull-request authors in open-source projects (Trinkenreich et al., 2022). A 2019 analysis further predicts that gender parity in computer




Figure 1: Illustration of the three studies.

science authorship will not be reached in this century if current trends persist (Wang et al., 2021). One contributing factor is the stereotypically male image of computing, which can dissuade non-cis-male children from seeing themselves as competent or from enjoying computational games (Yücel and Rızvanoğlu, 2019; Yang and Bers, 2024), with downstream effects. Large language models (LLMs) for code generation are now pervasive, with 82% of professional programmers reporting using ChatGPT daily (Stack Overflow, 2024). This technology has the potential to democratize programming support and lower entry barriers, possibly helping to bring more women into the field. At the same time, early evidence suggests that female students use ChatGPT less and exhibit lower prompting confidence than male students (Carvajal et al., 2024). LLMs are also known to replicate and amplify stereotypes and social biases present in their training data (Li et al., 2024), with gendered performance asymmetries seen in NLP tasks such as emotion and toxicity detection (Herbert et al., 2025; Excell and Al Moubayed, 2021). In programming specifically, recent work reports systematic differences in coding style between code written by men, women, and programmers of other genders (Brooke, 2024). To our knowledge, there are only

a limited number of studies that examine the influence of gender on prompting styles. Mashburn et al. (2025) found non-significant linguistic differences in politeness, formality and prompt length. Nevertheless, it is widely recognized in the literature that the quality of outputs generated by LLMs is influenced by the properties of the prompts used (Long et al., 2025; Yin et al., 2024). Considering that most existing code has historically been written by men and the training data of code-generation LLMs is likely skewed toward male-authored code, this raises questions: **RQ1** whether user gender is reflected in the linguistic style of prompts and is therefore predictable from prompt text alone, and **RQ2** whether such gendered prompting can lead to differences in LLM code-generation and **RQ3** in code-review. In this paper, we investigate these questions across three complementary studies: we characterize gendered linguistic patterns in real-world coding prompts, analyze how male and female participants use LLMs for programming in a controlled user study, and probe how commercial LLMs respond to systematically gender-coded prompts and review the resulting code. Our study contributes to a better understanding of linguistic gender differences in LLM prompting and of how gender bias can surface in code generation and automated code review, crucial understanding in the growing trend of LLMs-as-judges substituting human validation. Our shareable data and code are available in a GitHub repository ¹.

2 Related Work

2.1 Gendered language and gender prediction

Early work on gendered communication argues that women tend to use more politeness markers, hedges, indirect and involved language than men (Lakoff, 1973; Holmes, 1990; Biber and Burges, 2000). Gender also plays a role in how instructions are communicated. Women are said to employ more politeness strategies (e.g. use of indirect requests; Lakoff, 1973). Meta-analyses further document robust yet context-dependent gender effects across domains such as interactional style and emotional expression (Aydın, 2025; Leaper and Ayres, 2007; Leaper and Robnett, 2011; Thomson et al., 2000). Authorship profiling work demonstrates that gender can be predicted from text using lexical, stylistic and discourse features (Onikoyi et al.,

2023; O’Connor et al., 2024; Abdallah et al., 2020). Studies on online chat and social media demonstrated that women often use more relational and supportive language, whereas men show more task-oriented and efficiency-focused behavior, both in language and in usage (Herring, 2000; Koch et al., 2005).

2.2 Gender, programming and prompting

Empirical work on programming has only recently begun to examine gender explicitly. Qualitative studies questioned whether instructors can reliably “spot” female students based on perceived style (Carter and Jenkins, 2002), while more recent large-scale work reports significant differences in code style across genders (Brooke, 2024). In parallel, surveys and controlled experiments on ChatGPT use in education and professional settings suggest that female students and professionals tend to use LLMs less frequently, report lower prompting confidence, and sometimes achieve lower performance (Draxler et al., 2023; Bouzar et al., 2024; Yilmaz et al., 2023; Carvajal et al., 2024).

Several studies focus on better performing prompts, identifying factors such as explicit task framing, decomposition and example selection as important for performance (Long et al., 2025; Ma et al., 2025; White et al., 2023; Bsharat et al., 2023). Linguistic analyses investigate how surface properties of prompts (formatting, politeness, verbosity, uncertainty markers) shape model behavior. For instance, Leidinger et al. (2023) show that minor linguistic changes can substantially alter performance and stability across tasks and models. He et al. (2024) demonstrated that prompt formatting and complexity affect success in complex instruction following, while Yin et al. (2024) found that politeness levels have measurable effect. Zhou et al. (2023); Shi et al. (2023) show that hedges and uncertainty markers can systematically affect calibration or perceived confidence of model outputs. At the same time, lay users approach prompting differently from expert engineers. Studies on AI art and text tools show that lay users use conversational language and struggle to translate task requirements into effective prompts (Oppenlaender et al., 2025; Zamfirescu-Pereira et al., 2023), while work on AI literacy finds that users lack a clear mental model of how LLMs process instructions, limiting systematic prompt refinement (Knoth et al., 2024; Han et al., 2024).

¹<https://github.com/Mayaryin/GenderInfluenceInCodeGeneration>

2.3 LLMs for code generation and evaluation benchmarks

Starting from GPT-3, models show non-trivial coding abilities (Brown et al., 2020), and current work on Codex demonstrated strong performance on benchmarks like HumanEval and MBPP (Chen et al., 2021). To quantify progress, a wide range of code benchmarks have been introduced, including HumanEval-style function-completion tasks, multi-language extensions, and repository-level settings (Chen et al., 2021; Zhuo et al., 2024; Jimenez et al., 2024; Coignon et al., 2024). However, recent analyses warn that such benchmarks may overestimate model capabilities due to data leakage and narrow task formulations (Matton et al., 2024; Long et al., 2025; Leidinger et al., 2023). Newer security-focused benchmarks such as CWE-VAL add outcome-based checks that also analyze code security (Peng et al., 2025).

2.4 Bias and fairness in LLMs and code generation

Critical surveys emphasize that “bias” in LLMs is multifaceted and must be tied to concrete negative effects and social categories (Blodgett et al., 2020, 2021; Gallegos et al., 2024). Benchmark datasets such as StereoSet and BBQ expose stereotypical associations and group disparities in seemingly objective tasks (Nadeem et al., 2021; Parrish et al., 2022), while work on occupational reasoning and demographic name cues shows that minimal identity markers in prompts can elicit gendered responses (Sheng et al., 2019; Kaneko et al., 2024; Haim et al., 2024; Kotek et al., 2023). Recent audits document that LLMs can propagate bias not only in comments and variable names but also in control flow and access-control logic, with group-dependent behaviors when demographic attributes are hidden in specifications or tests (Liu et al., 2023; Ling et al., 2025). This has prompted calls for bias evaluations that examine downstream technical artifacts such as code (Liu et al., 2023; Gallegos et al., 2024). This pilot study 1 provides a first comprehensive quantification of how gender-coded prompts and self-descriptions relate to stylistic properties and maintainability of LLM-generated code and also investigates LLM code evaluation as a bias locus.

3 Study I

3.1 Methodology

Objective and Hypotheses This study investigated the link between user prompting style and LLM-generated code in real-world coding scenarios, using authentic chat histories from students and professionals. Based on prior work, we expected female prompts to show more pronouns, hedging, and involved language, and less directness. We tested against the following null hypotheses: H_{01} : Genders do not differ in personal pronoun use. H_{02} : Genders do not differ in the use of hedges. H_{03} : Genders do not differ in involved-informational ratio. H_{04} : Genders do not differ in instruction directness. Additionally, we explore whether user gender could be predicted from prompts and whether code quality varied by gender or prompt characteristics.

Data Collection and Participants We collected real-world chat data via an online survey targeting students and scientific staff. Participants submitted 90 LLM coding conversations in Python between May and September 2025, yielding 753 prompts. Demographic data ($N=30$; 15 male, 13 female, 2 other genders) indicated a young (20–30 y.o), highly educated cohort, with men more likely to have work experience and to use LLMs daily, while more women were current students. The sample comprised mainly advanced non-native English (C1–C2 level) speakers with academic or professional LLM use.

Prompt Analyses After excluding participants identifying as non-binary and other genders (represented by 1 count each), 746 prompts (282 male, 464 female) were included for linguistic analysis. Prompts of the same user were concatenated into one sample before comparing across genders to account for user-level differences, yielding sample sizes of $N_{male} = 15$ and $N_{female} = 13$. Basic analyses included prompt length, spelling and punctuation, unigram, bigram and word type frequencies (e.g., pronouns, pragmatic markers, adjectives), and the ‘involved-informational’ score following Biber’s framework (Biber, 1989) and the paradigm used by Kedrick et al. (2025) (see App. C). Furthermore, prompts were manually annotated using a taxonomy that crossed pronoun usage (impersonal, first person singular/plural) with clause type (imperative/statement vs. interrogative), allowing for multi-category classification per prompt.

Prompt directness was distinguished between direct commands ("Do x/y/z") or impersonal questions ("How to do x/y/z?") and indirect forms, e.g., questions with instructive intent ("Can you do...?", "Can we do...?") or first-person queries without direct imperative ("I need...") (see Table 13, App. C). Inter-annotator agreement across three annotators on 100 samples reached high Krippendorff's $\alpha = 0.79$ (see Table 9; full IAA results in App. C).

Gender Prediction For Study 1, RoBERTa (Zhuang et al., 2021) was finetuned on prompts from 28 users (15 male, 13 female; total of 536 prompts, 282 male; 254 female), employing stratified and group-aware 5-fold cross-validation. Performance was compared against a baseline of majority votes. Generalization was tested on data collected in Study 2 (46 male, 33 female prompts), and model predictions were explained with LIME scores (Ribeiro et al., 2016).

Code Quality Evaluation Traditional unit tests and reference implementations were unsuitable for our dataset due to the inherent ambiguity of the collected prompts. Instead, we assessed code quality using (1) self-reported user satisfaction and (2) complexity and maintainability analysis using Pylint² and Radon³ for Python code (more on them in App. A). A curated set of 34 Python prompts (excluding those containing code or non-generation tasks) was used to generate outputs from several SOTA LLMs⁴: ChatGPT 4o and 5; GPT o3 and 4.1 (OpenAI, 2026a,d,c,b); Claude Sonnet 3.7 and 4 (Anthropic, 2025a,b), and DeepSeek (2025). Code quality scores were averaged across the outputs of the three runs with each prompt, keeping temperature and top-p at their defaults to mimic daily-life usage. Code quality was compared between genders. While direct instructions are recommended by LLM providers for optimal results, indirect forms such as questions or first-person statements are a common discursive strategy and may function as polite requests. Thus, we also compared code quality scores for direct and indirect request types, treating this property as independent of gender. Finally, code quality metrics from the static analyzers were correlated with prompt properties—including length, complexity, informational and involved scores and ratio—to explore the potential impact of prompt characteristics on model output.

²<https://pylint.readthedocs.io>

³<https://radon.readthedocs.io>

⁴see App. B for model cards for each study

Female		Male	
Word	coef	Word	coef
thanks	0.295	working	-0.513
risk	0.210	hashtags	-0.466
looks	0.207	interactively	-0.322
difference	0.200	horizontally	-0.275
thank	0.173	adjust	-0.264
great	0.133	chat	-0.225
good	0.111	temperature	-0.193
you	0.111	halt	-0.177
niceee	0.100	visualization	-0.167
colored	0.099	again	-0.151
okay	0.096	first	-0.146
please	0.091	work	-0.132
good	0.075	getting	-0.120
can	0.071	error	-0.113

Table 1: Study I: most predictive terms (acc. to LIME)

Men		Women	
<i>of the</i>	0.56	<i>can you</i>	1.38
<i>in the</i>	0.55	<i>want to</i>	0.78
<i>at the</i>	0.55	<i>of the</i>	0.72

Table 2: Study I: top three bigrams per 100 words.

3.2 Results

Prompt Analyses Prompt analysis revealed no significant gender differences in length, spelling, or punctuation. However, women used more personal pronouns ($t(25.999) = -1.97, p = .030$) and pragmatic markers ($t(21.568) = -2.12, p = .046$), particularly hedges in the form of modal verbs ($t(25.095) = -2.86, p = .034$) and had a higher involved-informational ratio (see App. A for score interpretation) ($t(22.384) = -2.38, p = .013$), overall indicating a rather indirect and personal communication style (the first three null hypotheses $H0_1 - H0_3$ rejected). Men's prompts were more direct, albeit only with marginal significance ($t(26.000) = 2.05, p = .051$). Given the limited sample size, the test may have lacked sufficient power to detect smaller effects. Furthermore, the lack of an established framework to assess prompt directness complicates a definite acceptance or rejection of $H0_4$ at this stage. Notably, both genders preferred indirect, personal and interrogative interaction styles (Figure 2), but women employed the indirect "can you" formulation more frequently (Table 2).

Gender Prediction RoBERTa's performance moderately differed from the majority baseline, achieving a weighted F1 score of 0.60 during cross-validation and 0.61 on the test set (Table 3). Follow-




Figure 2: Study I: average proportion of prompts per gender of a certain request type. Whiskers indicate the range of data within 1.5 times the interquartile range from the lower and upper quartiles. Points outside this range are considered outliers.

Model	Cross-Val				Test			
	Acc	P	R	F1	Acc	P	R	F1
Baseline	0.45	0.21	0.45	0.28	0.58	0.34	0.58	0.43
RoBERTa	0.60	0.62	0.60	0.60	0.62	0.61	0.62	0.61

Table 3: Study I: averaged evaluation metrics during cross validation (left) and on unseen prompts from study 2 (right). Total of 536 prompts (282 male; 254 female) during cross validation. 46 male and 33 female prompts in the test set. Precision, recall, and F1-score are reported as a weighted average across both classes.

up LIME analysis of the predictions on the test set revealed the most predictive words associated with each gender, highlighting personal pronouns, politeness markers and the modal verb "can" for females, and more technical and neutral terms for males, such as "working", "hashtags" and "adjust".

Code Quality No differences emerged in self-reported satisfaction, static code quality, neither between genders nor direct and indirect requests across LLMs. Correlations between prompt features and code quality markers for each tested LLM yielded inconsistent and largely insignificant relationships between prompt and code traits. (Figure 6 and Table 11 in App. C). However, longer prompts were consistently linked to longer and more commented code, but not better code quality. Prompts with involved language showed a consistent slight tendency to elicit shorter code. Furthermore, LLM-generated code demonstrated high overall quality (Figure 7 in App. C), with most differences in quality relating to style and documentation rather than functionality (Table 12 in App. C).

4 Study II

4.1 Methodology

Objective and Hypotheses This study evaluated whether participant gender influences the quality or consistency of LLM-generated code in a controlled experiment. Participants were asked to prompt a commercial LLM to solve pre-defined given coding tasks. Our theoretical expectation was that, under equal access to the same LLM and standardized task conditions, outcomes should differ by gender. Accordingly, we formulate the following null hypotheses: H_{01} : Gender of the prompter does not affect task pass rates. H_{02} : Gender of the prompter does not affect code approval rates as assessed by an LLM code reviewer.

Study Design and Tasks The experiment was implemented using LimeSurvey, with anonymous participation. Each participant solved two coding tasks using an LLM of their choice: (1) implementing a password strength checker, and (2) correcting logic and syntax errors in a given hashtag validation function (see App. D for full task instruction). In Task 1, instructions were given via an audio recording, so that participants could not simply copy the instruction to the LLM interface. Participants copied detailed LLM conversation histories into structured survey fields for each task.

Participants Participants self-reported gender (female, male, female/male transgender, nonbinary, prefer not to say) and their chosen LLM provider and model. Recruitment targeted students, staff, and Prolific respondents across European countries who were fluent in English and had programming experience. A total of 59 participants took part in the study, of which 73.5% identified as male and 26.5% as female. Most participants were located in Western Europe, with 50% from the UK, 22% from Italy, and 16.3% from Germany; 41% reported English as their native language. Age distribution was: 40% aged 35–44, 31% aged 25–34, and about 15% in each of the under-25 and over-45 brackets. Regarding professional status, 76% identified as employed developers; 28.2% of participants were working students, and the remainder were non-working students or unemployed. A total of 97 valid submissions from both Task 1 and 2 were included for analysis.

Evaluation Methods Code correctness was assessed using unit tests, while code quality was

Measure	Stat.	p	N_f	N_m
Task 1 pass rate	$\chi^2(1) = 0.25$	0.62	58	39
Task 2 pass rate	$\chi^2(1) = 0.15$	0.70	58	39
Cyclomatic compl.	$t = 1.13$	0.26	52	36
Pylint score	$t = -0.27$	0.79	58	39
Maintainability idx.	$t = -1.06$	0.29	52	36
Lines of code	$t = 0.09$	0.93	52	36

Table 4: Study II – Code quality measures by participant gender. Differences in N are due to some submissions failing specific static-analysis procedures (e.g., cyclomatic complexity or maintainability analysis), although simpler metrics such as LOC could still be computed.

evaluated as in Study 1 via Radon metrics, Pylint scores, and lines of code (LOC). Reviewer approval was simulated by submitting solutions to LLM-based reviewers from several LLM providers, with outcomes categorized as APPROVE or CHANGES_REQUESTED and approval rates aggregated by gender, provider, and model. In total, all model-task combinations produced 429 review events for code written by female and 638 for code written by male participants. The prompt used to instruct the review bot is listed in Figure 10 in App. E. For both Study 2 and 3, we apply chi-squared tests to compare binary outcomes such as unit-test pass rates and reviewer approval decisions, and Welch’s t -test for continuous code quality measures.

4.2 Results

Across all tasks, unit-test pass rates and code quality (cyclomatic complexity, maintainability, Pylint scores) were similar between genders, with small, non-significant differences (Table 4).

However, when human-generated code was reviewed by an LLM review bot, a significant approval gap emerged: across all models, female-authored submissions were approved in 70.6% of all review events versus 62.9% for male-authored code, ($t \approx 2.67$, $p \approx 0.008$), despite comparable code correctness. Post-hoc comparisons by model provider revealed that the approval gap was modest for Anthropic and OpenAI models, larger for Deepseek, and most pronounced for Groq’s LLaMA models (see Table 5), where the difference was statistically significant.

Overall, we did not find gender differences in code correctness and accept H_{01} . We do see an effect of the user’s gender on LLM reviewer approval rates and thus reject H_{02} . Reviewers approved female-authored code more often, despite similar

Provider	Overall	Female	Male	Diff.	χ^2	p
Anthropic	83.3	84.6	81.9	2.7	0.09	0.76
OpenAI	74.9	76.9	72.8	4.1	0.62	0.43
Deepseek	59.8	65.0	54.6	10.4	2.96	0.10
Groq	44.4	52.6	36.2	16.4*	4.45	0.04

Table 5: Study II: LLM reviewer approval rates (%) by participant gender and provider. Diff. denotes female minus male approval.

quality. Thus, while gender showed no effect on code quality or correctness when using LLMs, the LLM reviewers exhibited systematic approval bias favoring female-authored code, varying by model family. These findings highlight a possible bias in automatic code review by LLMs, decoupled from code correctness or objective quality.

5 Study III

5.1 Methodology

Objectives and Hypotheses The third study examined whether gender-stereotypically worded prompts and gendered self-introductions of the prompter influenced functional correctness (unit-test pass rates) and stylistic properties of generated code, including static code quality measures derived from Pylint and Radon, across a range of LLM providers and model types. Thus, it investigated whether gender-coded prompt styles influence LLM code generation and automated code review when user-level variation is removed. While Studies I and II analyze human-authored prompts, this study uses synthetically constructed prompts to isolate the effects of stylistic framing alone. Importantly, the manipulated prompts explicitly embedded stylistic and engineering-oriented adjectives. Consequently, differences in maintainability, compactness, or documentation-related metrics should be interpreted as prompt-following effects induced by stylistic framing associated with gender rather than evidence of latent gender-linked generation ability or demographic bias in code quality itself.

We tested against three null hypotheses: H_{01} : Gendered prompt variant does not affect correctness. H_{02} : Gendered prompt variant does not affect code structure and surface quality. H_{03} : Reviewer approval does not vary by persona and prompt variant.

Study Design and Tasks Five prompt variants were constructed: male-coded, female-coded, neutral, and their respective self-introduction forms,




Figure 3: Study III: code structure and style by prompt group and provider. Female-coded prompts yield higher maintainability, while male-coded and neutral prompts tend to yield higher Pylint scores.




Figure 4: Study III: unit-test pass rates by prompt group and provider. Pass rates are tightly clustered across gender-coded prompts, indicating no substantial differences in functional correctness ($H0_1$).




Figure 5: Study III: LLM reviewer approval rates by prompt group and provider backend across reviewer personas.

coupled with a neutral, technical task description. Gender coding was partially operationalized through the selection of literature-based stereotypically agentic versus communal adjective framings, commonly discussed in prior work on gendered communication and stereotype framing, and based on Hentschel et al. (2019) and series of works from Leaper (Leaper and Ayres, 2007; Leaper and Robnett, 2011). Agentic descriptors emphasized performance, efficiency, assertiveness, and task orientation, whereas communal descriptors emphasized clarity, collaboration, supportiveness, and interpersonal orientation. Adjective pairs were ran-

domly sampled for each prompt instance. Self-introduction variants were crafted by employing stereotypically gendered names (“Jack” for male (self), “Sarah” for female (self)). Because several adjective pairs also encode stylistic programming preferences (e.g., concise vs. descriptive or efficiency-oriented vs. collaborative framing), observed differences in static code metrics are interpreted as responses to prompt framing rather than demographic attributes themselves. See the prompts in App. E, Figures 11, 12, and 13, and the word dictionaries in Table 15. Prompts were applied across five programming tasks of average difficulty (roman2int, unique subsets, isbn13, flatten nested, and spiral order), each solvable in a single function and designed to allow automated testing (see App. E, Table 16). Each prompt-task combination was tested under various decoding profiles from fully deterministic (temp. 0.0, top-p 1.0) to highly creative (temp. 1.0, top-p 0.6). Code was generated by state-of-the-art LLMs from multiple providers (OpenAI, Deepseek, Anthropic, Groq, see App. B), with fixed random seeds.

Evaluation Methods In addition to the code quality tests from Study 2, reviewer personas were implemented as LLM-based judges using male-, female-, or neutral-coded system messages to assess code quality in a fully deterministic setting. Prompts used to construct the automated reviewer can be found in App. E Table 10.

5.2 Results

Consistent with $H0_1$, unit-test pass rates were tightly clustered across prompt variants (54–57%) and, based on chi-square tests, showed no statistically significant differences between male-coded, female-coded, neutral, or self-introduction variants (see Table 6 and Figure 4). This pattern held across LLM providers, indicating that stylistic

Attribute	Test	p
Unit-test pass rate	$\chi^2 = 0.11$	0.74
Cyclomatic complexity	$t = 1.30$	0.19
Maintainability index (MI)	$t = 17.0$	$< 10^{-50}$
Global Pylint score	$t = 5.66$	$< 10^{-7}$
Lines of code (LOC)	$t = 11.6$	$< 10^{-10}$
Male reviewer approval	$t = 1.97$	0.049

Table 6: Summary of statistical tests and outcomes in Study 3. Effects for neutral and female reviewer personas did not reach $p < 0.05$.

gender cues do not affect functional correctness. However, contrary to $H0_2$, prompt style systematically influenced code structure and surface quality. Female-coded prompts produced longer code (App. E, Figure 14) with higher maintainability scores and slightly lower cyclomatic complexity, whereas male-coded and neutral prompts yielded more compact code with higher global Pylint scores (see Figure 3). These effects were stable across decoding profiles, though their magnitude varied by model family. We find evidence against $H0_3$ as reviewer persona differences do vary among model providers, with Groq showing the largest differences and Anthropic the smallest (Figure 5). Female reviewers tend to be stricter than male reviewers for Groq and OpenAI. Deepseek presents mixed results: female reviewers are generally stricter or equally strict, but male reviewers exhibit greater approval for male-coded prompts, indicating potential sensitivity. Comparing approval rates within each persona, we find a marginally significant effect for the male persona, approving code from female-coded prompts more often (Table 6).

6 Discussion and Conclusion

Across three studies, we analyze how gendered prompting styles manifest in code-generation and LLM-based review. While prompt language and reviewer responses differ by gender, we find little evidence of consistent differences in functional correctness or static code quality for human users.

Study 1 shows that gender leaves a subtle linguistic trace in real-world code prompts. Prompts produced by cis-gender women exhibit more hedging and involved or relational phrasing, whereas those produced by men tend to be more direct and task-focused. These patterns broadly align with previous findings on gendered communication in other domains, while underscoring the role of the register and context in professional settings. Importantly,

prompting does not reproduce simplistic gender stereotypes, instead reflecting shared conventions. Consistent with this, predicting user gender from prompt text alone proves difficult: RoBERTa achieves only moderate validation performance and generalization to test data from a different study is similarly poor (**RQ1**). This suggests that much of the observed linguistic variation may be driven by confounding factors such as experience level rather than gender alone. This tempers concerns that code LLMs could systematically “personalize” their responses based on gender inferred from prompt style alone, but it does not remove them, as even weak signals can matter when amplified over millions of interactions. In Study 2, we find that male and female solutions are similarly correct, which contrasts with prior work suggesting that women use ChatGPT less frequently and with less confidence, indicating that LLM-assisted coding can equalize outcomes (**RQ2**). More importantly, Study 2 exposes the LLM-as-a-judge for code review as a locus of bias (**RQ3**). In our setup, LLM reviewers approve female-written solutions more often than male-written ones, despite comparable correctness and static quality. Importantly, magnitude and direction of this effect vary across providers and models. Study 3 provides a more focused lens on these mechanisms by isolating prompt style and reviewer persona in an LLM-only setting. When we systematically vary the gender-coding of prompts and self-descriptions (adjectives + names), we again find that correctness is largely unaffected. However, similar to Study 1 gender-coded prompt styles do influence code maintainability and compactness metrics (**RQ2**). Some prompt styles tend to induce shorter, more compact solutions with lower complexity, whereas others produce more verbose or heavily commented code. Reviewer personas, in turn, react differently to these stylistic variants, with Groq having the largest and Anthropic the smallest difference in approval rates between reviewer personas. This parallels results in Study 2, where the approval rate difference between genders was significant for Groq, and smallest, although non-significant for Anthropic (**RQ3**). If the variation between approval rates is seen as a bias indicator, Groq is the most biased in this investigation.

Our results show that fairness concerns in LLM-assisted programming lie less in generative accuracy and more in how LLMs, especially when used as reviewers, respond to different prompt styles and perceived author personas. We also identified,

that while code correctness does not vary based on the gender of the prompter, code maintainability and conciseness seem to be affected. More broadly, our findings call for a socio-technical perspective, in which technical debiasing is coupled with careful governance of when and how LLM-as-a-judge components are deployed. They have direct implications for deploying LLM-based code review systems in educational and industrial software-engineering workflows, showing that LLM reviewers should not be used as independent sources for candidate evaluation without human oversight, so that existing social asymmetries are not silently reproduced in educational, professional settings. Beyond the immediate findings, our results have broader implications for the design and evaluation of LLM-assisted programming systems. First, they suggest that fairness assessments should consider the complete development pipeline rather than focusing exclusively on code-generation accuracy. While functional correctness remained largely stable across our studies, differences emerged in code review decisions and stylistic properties of generated code, indicating that downstream evaluation components may represent a more important source of unequal treatment. Second, our findings contribute to the growing literature on LLM-as-a-Judge by highlighting that automated reviewer decisions may be influenced by contextual and linguistic cues beyond objective code quality. Although these observations require confirmation in larger-scale studies, they suggest that fairness auditing should become a standard part of evaluating automated review systems. Finally, the absence of consistent correctness differences across prompting styles indicates that diverse communication strategies can produce comparable programming outcomes, suggesting that educational and professional tools should support a variety of prompting styles rather than implicitly promoting a single normative interaction style.

Limitations

Our studies have several limitations. First, our sample sizes are relatively small and demographically limited. According to a *post hoc* power analysis (see Appendix A), given the actual group sizes in Study 1 ($n = 13$ and $n = 15$), the study was likely underpowered to detect small or medium effects and primarily sensitive to large effects only. Consequently, both null findings and observed effect

sizes should be interpreted cautiously, as underpowered studies may overestimate effect magnitudes. Larger follow-up studies are needed to validate the robustness of the observed patterns.

We received only one submission each from a non-binary participant and a transgender participant, leading us to exclude these results from our analysis, and thus the study approached gender mainly as a binary construct. However, we are planning further studies with an improved design to better accommodate diverse gender identities in the future. The participants were mostly technically literate, self-selected, and drawn from a limited set of educational and professional contexts. Moreover, our prompts and tasks focus on short programming problems rather than full software projects. As a result, our findings may not generalize to other populations (e.g., less experienced programmers) or to more complex, collaborative workflows.

Second, our results are tied to a specific set of commercial LLMs and configurations available at the time of data collection. Newer model releases, fine-tuning practices, and usage policies may change both generative behavior and review patterns. Our evaluation of code quality relies on functional tests, static analysis, and LLM-based judgment rather than exhaustive unit testing or expert review by multiple human programmers.

Finally, Study 3 abstracts away from human users entirely and operates in an LLM-only simulation of code generation and review. While this allows us to isolate the effects of prompt wording and reviewer personas under controlled conditions, it cannot capture the full complexity of real-world interaction patterns, learning effects, or strategic adaptations over time. This operationalization of gender captures stereotypically gendered prompt framings rather than gender as a personal attribute, and the results should be interpreted as reflecting how such stylistic cues are treated by LLMs, not as claims about gender itself.

Taken together, these limitations mean that our work should be seen as a pilot: it identifies mechanisms through which gender-coded prompt styles and LLM-as-judge components can interact, but further, larger-scale and more diverse studies are needed before drawing strong conclusions about real-world deployment or prescriptive design guidelines. Future work could compare how many prompt iterations users with different communication styles require to achieve satisfactory outcomes to further clarify interaction dynamics in

LLM-assisted programming.

Ethical Considerations

Our findings have several implications for fairness audits and the design of LLM-based programming tools. Rather than focusing narrowly on generative accuracy, we show that key fairness risks lie in how LLM reviewers treat code associated with different prompt styles or author personas, especially as LLM-as-a-judge components are adopted for automated review, grading, and hiring, making it essential to audit evaluation pipelines themselves for demographic and stylistic biases. Our results also complicate the idea that fairness can be achieved by simply stripping demographic attributes from prompts, since gendered information is often encoded in multiple stylistic features that continue to shape model behavior even in the absence of explicit labels, calling for a combination of input-level strategies, model-side mitigation, and interface-level safeguards. At the same time, correctness in our setup is robust across gender-coded prompts, suggesting that tools and pedagogy should support diverse prompting practices rather than enforcing a single “expert” style, particularly in educational settings where prescriptive norms risk marginalizing students whose communication diverges from the stereotype of the “male coder.” More broadly, our studies underline the need for a socio-technical perspective: gendered patterns in prompts arise from longstanding norms around communication and expertise. LLMs trained on skewed data can reflect and amplify these patterns, and institutional choices about where and how to deploy LLM-as-judge systems determine whether such patterns have real-world consequences, making questions of system governance and value choices as central as technical debiasing.

References

- Emad E. Abdallah, Jamil R. Alzghoul, and Muath Alzghool. 2020. [Age and gender prediction in open domain text](#). *Procedia Computer Science*, 170:563–570. The 11th International Conference on Ambient Systems, Networks and Technologies (ANT) / The 3rd International Conference on Emerging Data and Industry 4.0 (EDI40) / Affiliated Workshops.
- Anthropic. 2025a. [Claude 3.7 sonnet system card](#).
- Anthropic. 2025b. [System card: Claude opus 4 & claude sonnet 4](#).
- Fatma Aydın. 2025. [Examining gender differences in social media language](#). *Bulletin of Language and Literature Studies*, 2(1).
- Yoav Benjamini and Yosef Hochberg. 1995. [Controlling the false discovery rate: A practical and powerful approach to multiple testing](#). *Journal of the Royal Statistical Society: Series B (Methodological)*, 57(1):289–300.
- Douglas Biber. 1989. [A typology of English texts](#). *Linguistics*, 27(1):3–44.
- Douglas Biber and Jená Burges. 2000. [Historical change in the language use of women and men](#). *Journal of English Linguistics*, 28(1):21–37.
- Su Lin Blodgett, Solon Barocas, Hal Daumé III, and Hanna Wallach. 2020. [Language \(technology\) is power: A critical survey of “bias” in NLP](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5454–5476.
- Su Lin Blodgett, Gilsinia Lopez, Alexandra Olteanu, Robert Sim, and Hanna Wallach. 2021. [Stereotyping Norwegian salmon: An inventory of pitfalls in fairness benchmark datasets](#). 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 1004–1015.
- Abdelouahd Bouzar, Khaoula El Idrissi, and Tayeb Ghourdou. 2024. [Gender differences in perceptions and usage of ChatGPT](#). *International Journal of Humanities and Educational Research*, Volume 6(Issue 2).
- Siân Brooke. 2024. [Programmed differently? Testing for gender differences in Python programming style and quality on GitHub](#). *Journal of Computer-Mediated Communication*, 29(1):zmad049.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, and 12 others. 2020. [Language models are few-shot learners](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran Associates, Inc.
- Sondos Mahmoud Bsharat, Aidar Myrzakhan, and Zhiqiang Shen. 2023. [Principled instructions are all you need for questioning LLaMA-1/2, GPT-3.5/4](#). *arXiv*.
- Janet Carter and Tony Jenkins. 2002. Spot the difference: Are there gender differences in coding style? In *Proceedings of the 3rd Annual LTSN-ICS Conference*.

- Daniel Carvajal, Catalina Franco, and Siri Isaksson. 2024. [Will Artificial Intelligence Get in the Way of Achieving Gender Equality?](#) *SSRN Electronic Journal*.
- Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Pondé, Jared Kaplan, Harrison Edwards, Yura Burda, Nicholas Joseph, Greg Brockman, Alex Ray, Raul Puri, Gretchen Krueger, Michael Petrov, Heidy Khlaaf, Girish Sastry, Pamela Mishkin, Brooke Chan, Scott Gray, and 34 others. 2021. [Evaluating large language models trained on code](#). *ArXiv*, abs/2107.03374.
- Tristan Coignon, Clément Quinton, and Romain Rouvoy. 2024. [A performance study of llm-generated code on leetcode](#). In *Proceedings of the 28th International Conference on Evaluation and Assessment in Software Engineering, EASE '24*, page 79–89, New York, NY, USA. Association for Computing Machinery.
- David Corey, William Dunlap, and Michael Burke. 1998. [Averaging correlations: Expected values and bias in combined pearson rs and fisher's z transformations](#). *Journal of General Psychology - J GEN PSYCHOL*, 125:245–261.
- DeepSeek. 2025. [Deepseek-v3.1 release | deepseek api docs](#).
- Fiona Draxler, Daniel Buschek, Mikke Tavast, Perttu Hämäläinen, Albrecht Schmidt, Juhi Kulshrestha, and Robin Welsch. 2023. [Gender, age, and technology education influence the adoption and appropriation of llms](#). *CoRR*, abs/2310.06556.
- Elizabeth Excell and Noura Al Moubayed. 2021. [Towards equal gender representation in the annotations of toxic language detection](#). In *Proceedings of the 3rd Workshop on Gender Bias in Natural Language Processing*, pages 55–65, Online. Association for Computational Linguistics.
- Ronald A Fisher. 1915. [Frequency distribution of the values of the correlation coefficient in samples from an indefinitely large population](#). *Biometrika*, 10(4):507–521.
- Isabel O. Gallegos, Ryan A. Rossi, Joe Barrow, Md Mehrab Tanjim, Sungchul Kim, Franck Dernoncourt, Tong Yu, Ruiyi Zhang, and Nesreen K. Ahmed. 2024. [Bias and fairness in large language models: A survey](#). *Computational Linguistics*, 50(3):1097–1179.
- Amit Haim, Alejandro Salinas, and Julian Nyarko. 2024. [What's in a Name? Auditing Large Language Models for Race and Gender Bias](#). *arXiv*.
- Wenjuan Han, Xiang Wei, Xingyu Cui, Ning Cheng, Guangyuan Jiang, Weinan Qian, and Chi Zhang. 2024. [Prompt engineering 101 prompt engineering guidelines from a linguistic perspective](#). In *Proceedings of the 23rd Chinese National Conference on Computational Linguistics (Volume 1: Main Conference)*, pages 1408–1426, Taiyuan, China. Chinese Information Processing Society of China.
- Jia He, Mukund Rungta, David Koleczek, Arshdeep Sekhon, Franklin X Wang, and Sadid Hasan. 2024. [Does prompt formatting have any impact on llm performance?](#) *arXiv preprint arXiv:2411.10541*.
- Tanja Hentschel, Madeline E. Heilman, and Claudia V. Peus. 2019. [The multiple dimensions of gender stereotypes: A current look at men's and women's characterizations of others and themselves](#). *Frontiers in Psychology*, 10:11.
- Maureen Herbert, Katie Sun, Angelica Lim, and Yasaman Etesam. 2025. [Gender bias in emotion recognition by large language models](#). *Preprint*, arXiv:2511.19785.
- Susan Herring. 2000. [Gender differences in CMC: Findings and implications](#). *The CPSR Newsletter*, 18.
- Janet Holmes. 1990. [Hedges and boosters in women's and men's speech](#). *Language & Communication*, 10(3):185–205.
- Carlos E Jimenez, John Yang, Alexander Wettig, Shunyu Yao, Kexin Pei, Ofir Press, and Karthik R Narasimhan. 2024. [SWE-bench: Can language models resolve real-world github issues?](#) In *The Twelfth International Conference on Learning Representations*.
- Masahiro Kaneko, Danushka Bollegala, Naoaki Okazaki, and Timothy Baldwin. 2024. [Evaluating gender bias in large language models via chain-of-thought prompting](#). *arXiv preprint arXiv:2401.15585*.
- Kara Kedrick, Ekaterina Levitskaya, and Russell J. Funk. 2025. [Investigating writing style as a contributor to gender gaps in science and technology](#).
- Nils Knoth, Antonia Tolzin, Andreas Janson, and Jan Marco Leimeister. 2024. [AI literacy and its implications for prompt engineering strategies](#). *Computers and Education: Artificial Intelligence*, 6:100225.
- Sabine C. Koch, Barbara Mueller, Lenelis Kruse, and Joerg Zumbach. 2005. [Constructing Gender in Chat Groups](#). *Sex Roles*, 53(1-2):29–41.
- Hadas Kotek, Rikker Dockum, and David Q. Sun. 2023. [Gender bias and stereotypes in large language models](#). In *Proceedings of the Collective Intelligence Conference (CI '23)*.
- Michele Lacchia. 2024. [Radon Documentation](#). Radon Project. Version 4.x.
- Robin Lakoff. 1973. [Language and woman's place](#). *Language in society*, 2(1):45–79.

- Campbell Leaper and Melanie M. Ayres. 2007. [A meta-analytic review of gender variations in adults' language use: Talkativeness, affiliative speech, and assertive speech](#). *Personality and Social Psychology Review*, 11(4):328–363.
- Campbell Leaper and Rachael D. Robnett. 2011. [Women are more likely than men to use tentative language, aren't they? a meta-analysis testing for gender differences and moderators](#). *Psychology of Women Quarterly*, 35(1):129–142.
- Alina Leidinger, Robert Van Rooij, and Ekaterina Shutova. 2023. The language of prompting: What linguistic properties make a prompt successful? *arXiv preprint arXiv:2311.01967*.
- Tianlin Li, Xiaoyu Zhang, Chao Du, Tianyu Pang, Qian Liu, Qing Guo, Chao Shen, and Yang Liu. 2024. [Your Large Language Model is Secretly a Fairness Proponent and You Should Prompt it Like One](#). *arXiv*.
- Lin Ling, Fazle Rabbi, Song Wang, and Jinqiu Yang. 2025. [Bias unveiled: investigating social bias in llm-generated code](#). In *Proceedings of the Thirty-Ninth AAAI Conference on Artificial Intelligence and Thirty-Seventh Conference on Innovative Applications of Artificial Intelligence and Fifteenth Symposium on Educational Advances in Artificial Intelligence*, AAAI'25/IAAI'25/EAAI'25. AAAI Press.
- Yan Liu, Xiaokang Chen, Yan Gao, Zhe Su, Fengji Zhang, Daoguang Zan, Jian-Guang Lou, Pin-Yu Chen, and Tsung-Yi Ho. 2023. [Uncovering and quantifying social biases in code generation](#). In *Proceedings of the 37th International Conference on Neural Information Processing Systems*, NIPS '23, Red Hook, NY, USA. Curran Associates Inc.
- Do Xuan Long, Duy Dinh, Ngoc-Hai Nguyen, Kenji Kawaguchi, Nancy F. Chen, Shafiq Joty, and Min-Yen Kan. 2025. [What makes a good natural language prompt?](#) In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 5835–5873, Vienna, Austria. Association for Computational Linguistics.
- Qianou Ma, Weirui Peng, Chenyang Yang, Hua Shen, Ken Koedinger, and Tongshuang Wu. 2025. [What should we engineer in prompts? training humans in requirement-driven LLM use](#). *ACM Transactions on Computer-Human Interaction*, 32(4):1–27.
- Patricia Mashburn, Felix A. Weuthen, Nelly Otte, Hanif Krabbe, Gerardo M. Fernandez, Thomas Kraus, and Julia Krabbe. 2025. [Gender differences in the use of ChatGPT as generative artificial intelligence for clinical research and decision-making in occupational medicine](#). *Healthcare*, 13(12):1394.
- Alexandre Matton, Tom Sherborne, Dennis Aumiller, Elena Tommasone, Milad Alizadeh, Jingyi He, Raymond Ma, Maxime Voisin, Ellen Gilsenan-McMahon, and Matthias Gallé. 2024. [On leakage of code generation evaluation datasets](#). In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 13215–13223, Miami, Florida, USA. Association for Computational Linguistics.
- Moin Nadeem, Anna Bethke, and Siva Reddy. 2021. [StereoSet: Measuring stereotypical bias in pretrained language models](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 5356–5371.
- Karen O'Connor, Su Golder, Davy Weissenbacher, Ari Z Klein, Arjun Magge, and Graciela Gonzalez-Hernandez. 2024. [Methods and annotated data sets used to predict the gender and age of twitter users: Scoping review](#). *J Med Internet Res*, 26:e47923.
- Babatunde Onikoyi, Nonso Nnamoko, and Ioannis Korkontzelos. 2023. [Gender prediction with descriptive textual data using a machine learning approach](#). *Natural Language Processing Journal*, 4:100018.
- OpenAI. 2026a. Chatgpt-4o model | openai api. <https://platform.openai.com/docs/models/chatgpt-4o-latest>. Accessed: 2026-01-03.
- OpenAI. 2026b. Gpt-4.1 model | openai api. <https://platform.openai.com/docs/models/gpt-4.1>. Accessed: 2026-01-03.
- OpenAI. 2026c. Gpt-5 chat model | openai api. <https://platform.openai.com/docs/models/gpt-5-chat-latest>. Accessed: 2026-01-03.
- OpenAI. 2026d. o3 model | openai api. <https://platform.openai.com/docs/models/o3>. Accessed: 2026-01-03.
- Jonas Oppenlaender, Rhema Linder, and Johanna Silvennoinen. 2025. [Prompting AI art: An investigation into the creative skill of prompt engineering](#). *International Journal of Human-Computer Interaction*, 41(16):10207–10229.
- Alicia Parrish, Angelica Chen, Nikita Nangia, Vishakh Padmakumar, Jason Phang, Jana Thompson, Phu Mon Htut, and Samuel R. Bowman. 2022. [BBQ: A hand-built bias benchmark for question answering](#). In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 2086–2105.
- Jinjun Peng, Leyi Cui, Kele Huang, Junfeng Yang, and Baishakhi Ray. 2025. [Cweval: Outcome-driven evaluation on functionality and security of llm code generation](#). In *Conference: 2025 IEEE/ACM International Workshop on Large Language Models for Code*, pages 33–40.
- Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. 2016. ["Why should i trust you?": Explaining the predictions of any classifier](#). In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '16, page 1135–1144, New York, NY, USA. Association for Computing Machinery.

- S. S. Shapiro and M. B. Wilk. 1965. [An analysis of variance test for normality \(complete samples\)](#)†. *Biometrika*, 52(3-4):591–611.
- 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*, pages 3407–3412.
- Freda Shi, Xinyun Chen, Kanishka Misra, Nathan Scales, David Dohan, Ed Chi, Nathanael Schärli, and Denny Zhou. 2023. [Large Language Models Can Be Easily Distracted by Irrelevant Context](#). *arXiv*.
- N Clayton Silver and William P Dunlap. 1987. [Averaging correlation coefficients: should fisher’s z transformation be used?](#) *Journal of applied psychology*, 72(1):146.
- SlashData. 2023. [Developer nation pulse report](#).
- C. Spearman. 1904. [The proof and measurement of association between two things](#). *The American Journal of Psychology*, 15(1):72–101.
- Stack Overflow. 2024. [Stack Overflow Developer Survey 2024](#).
- Rob Thomson, Tamar Murachver, and James Green. 2000. [Where is the gender in gendered language?](#) *Psychological Science*, 12(2):171–175.
- Bianca Trinkenreich, Igor Wiese, Anita Sarma, Marco Gerosa, and Igor Steinmacher. 2022. [Women’s participation in open source software: A survey of the literature](#). *ACM Transactions on Software Engineering and Methodology (TOSEM)*, 31(4):1–37.
- Lucy Lu Wang, Gabriel Stanovsky, Luca Weihs, and Oren Etzioni. 2021. [Gender trends in computer science authorship](#). *Commun. ACM*, 64(3):78–84.
- B. L. Welch. 1947. [The generalization of ‘student’s’ problem when several different population variances are involved](#). *Biometrika*, 34(1-2):28–35.
- Jules White, Quchen Fu, Sam Hays, Michael Sandborn, Carlos Olea, Henry Gilbert, Ashraf Elnashar, Jesse Spencer-Smith, and Douglas C. Schmidt. 2023. [A prompt pattern catalog to enhance prompt engineering with chatgpt](#). In *Proceedings of the 30th Conference on Pattern Languages of Programs, PLoP ’23*, USA. The Hillside Group.
- Zhanxia Yang and Marina Bers. 2024. [Examining Gender Difference in the Use of ScratchJr in a Programming Curriculum for First Graders](#). *Computer Science Education*, 34(4):864–885.
- Halit Yilmaz, Samat Maxutov, Azatzhan Baitekov, and Nuri Balta. 2023. [Student attitudes towards chat GPT: A technology acceptance model survey](#). *International Educational Review*, 1(1):57–83.
- Ziqi Yin, Hao Wang, Kaito Horio, Daisuke Kawahara, and Satoshi Sekine. 2024. [Should we respect LLMs? a cross-lingual study on the influence of prompt politeness on LLM performance](#). In *Proceedings of the Second Workshop on Social Influence in Conversations (SICoN 2024)*, pages 9–35, Miami, Florida, USA. Association for Computational Linguistics.
- Yeliz Yücel and Kerem Rızvanoğlu. 2019. [Battling gender stereotypes: A user study of a code-learning game, “Code Combat,” with middle school children](#). *Computers in Human Behavior*, 99:352–365.
- J.D. Zamfirescu-Pereira, Richmond Y. Wong, Bjoern Hartmann, and Qian Yang. 2023. [Why johnny can’t prompt: How non-AI experts try \(and fail\) to design LLM prompts](#). *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, pages 1–21.
- Kaitlyn Zhou, Dan Jurafsky, and Tatsunori Hashimoto. 2023. [Navigating the Grey Area: How Expressions of Uncertainty and Overconfidence Affect Language Models](#). *arXiv*.
- Liu Zhuang, Lin Wayne, Shi Ya, and Zhao Jun. 2021. [A robustly optimized BERT pre-training approach with post-training](#). In *Proceedings of the 20th Chinese National Conference on Computational Linguistics*, pages 1218–1227, Huhhot, China. Chinese Information Processing Society of China.
- Terry Yue Zhuo, Minh Chien Vu, Jenny Chim, Han Hu, Wenhao Yu, Ratnadira Widayarsi, Imam Nur Bani Yusuf, Haolan Zhan, Junda He, Indraneil Paul, Simon Brunner, Chen Gong, Thong Hoang, Armel Randy Zebaze, Xiao ke Hong, Wen-Ding Li, Jean Kadour, Minglian Xu, Zhihan Zhang, and 14 others. 2024. [Bigcodebench: Benchmarking code generation with diverse function calls and complex instructions](#). *ArXiv*, abs/2406.15877.

A Analysis tools, metrics and statistical analysis

Involved-informational framework. We base our linguistic analyses on a framework first declared by Biber (1989), who distinguishes “involved” texts (personal, affective, marked by pronouns, contractions, and private verbs) from “informational” texts (dense, referential, marked by nouns and prepositions), reflecting differences in communicative purpose and production circumstances. Kedrick et al. (2025) applied this framework to an investigation of gender differences in scientific articles, operationalizing involved style via personal pronouns, questions, as well as the use of “and”; informational style via cardinal numbers, determiners, and past tense verbs. Their analysis revealed consistent gender differences, with female authors using more involved features than male authors. Our findings using the same approach align with those previous results, indicating that gendered differences in textual communication also transfer to the “genre” of prompts.

Code analysis. Following Brooke (2024), we assess static code quality using two established Python analysis tools: Pylint to capture warnings such as unused variables/imports, undefined names, unreachable code, and style violations, and Radon to quantify structural complexity and maintainability as proxies for code quality. Pylint⁵ is a configurable static analyzer that parses Python source code into an abstract syntax tree and reports errors and deviations from coding conventions such as PEP 8. Radon (Lacchia, 2024) is a complementary tool that computes code metrics from source, including raw line counts, cyclomatic complexity (McCabe’s metric), and a composite maintainability index.

Statistical Analysis and Evaluation Metrics. We conducted a *post hoc* power analysis ($\alpha = .05$), which indicated required sample sizes of 26, 64, and 393 participants per group to detect large ($d = 0.8$), medium ($d = 0.5$), and small ($d = 0.2$) effects in two-sample *t*-tests, respectively. Given our sample size of 28 participants, this study was capable of detecting large effects.

We assessed normality of the data using the Shapiro-Wilk test (Shapiro and Wilk, 1965) and employed the independent samples *t*-test (Welch, 1947) for normal data. Chi-square tests were used

⁵<https://github.com/pylint-dev/pylint?tab=readme-ov-file>

to compare distributions of categorical data. In Study 1, false discovery rate was controlled with Benjamini-Hochberg correction ((Benjamini and Hochberg, 1995), $\alpha = .05$). Spearman’s rank correlation (Spearman, 1904) was used to assess associations between prompt and code variables. Classifier performance was measured by accuracy, precision, recall, and F1-score. For gender prediction with RoBERTa, 50 prompts from a user who dominated the dataset with 260 prompts were sub-sampled, reducing the overall dataset size from 746 to 536. This was to ensure balanced splits in the group-aware stratified cross validation.

Computation Budget. The computational resources required for this work were intentionally kept minimal. The RoBERTa-based models were trained and evaluated on a standard consumer laptop without the use of dedicated GPU hardware. In addition, external LLMs were accessed through commercial API providers, primarily within available free-tier usage limits whenever feasible. The overall direct computational expenditure for the project was approximately USD 55, covering limited API usage and associated cloud-based inference costs.

B LLMs used

Study	Model (provider and version)
Study 1	ChatGPT 4o (chatgpt-4o-latest, 09/25; OA) GPT o3 (o3-2025-04-16; OA) ChatGPT 5 (gpt-5-chat-latest, 09/25; OA) GPT 4.1 (gpt-4.1-2025-04-14; OA) ClaudeSonnet3.7 (claude-3-7-sonnet-20250219; AN) ClaudeSonnet4 (claude-sonnet-4-20250514; AN) DeepSeek V3.1 (deepseek-chat; DS)
Study 2 & Study 3	GPT 4.1 (gpt-3.5-turbo, gpt-4.1, gpt-4.1-mini, gpt-4o; OA) ClaudeSonnet3.7 (claude-3.7-sonnet-latest; AN) ClaudeSonnet3 (claude-3-haiku-20240307; AN) DeepSeek V3.1 (deepseek-chat, deepseek-coder, deepseek-reasoner; DS) Groq API 1 (llama-3.1-8b-instant) Groq API 2 (llama-3.3-70b-versatile)

Table 7: LLM models and versions used across the three studies. OA = OpenAI, AN = Anthropic, DS = DeepSeek.

C Study I

In this Appendix, we add the additional tables and visualisation in support of Study I.

Annotation bel	La-	Linguistic Category	Description	Typical Examples
can_you_quest.		Interrogative request / Indirect directive	A question addressed to “you” using modal verbs to formulate an indirect request or instruction.	<i>can you, could you, couldn't you</i>
i_statement		Deontic I-statement / Volitional statement	A first-person statement expressing wants, needs, obligations, or intentions.	<i>I want, I need, I have to, I have</i>
impers._command		Impersonal directive / Normative statement	A requirement or instruction framed impersonally, without directly addressing an interlocutor.	<i>it has to be..., <description of the task/outcome></i>
you_command		Imperative directive	A direct command formulated in the imperative mood and addressed to “you”.	<i>write a function, do <task></i>
we_command		Inclusive imperative / Hortative	A suggestion or command including the speaker through collective wording such as “we” or “let’s”.	<i>let's, we have to, we need</i>
unpers._quest.		Impersonal interrogative / Content question	An information-seeking question without a personal subject, typically beginning with interrogative forms such as <i>how</i> or <i>what</i> .	<i>how to, what is</i>
i_question		First-person interrogative	A question in which the speaker asks about their own role, ability, or permission.	<i>can I, do I, can't I, what are my</i>
we_question		First-person plural interrogative / Deliberative question	A question using “we” to seek permission, discuss options, or propose collective action.	<i>can we, could we, shouldn't we</i>

Annotation guideline. Annotators should identify whether prompts contain one or more of the directive or interrogative phenomena listed above. Multiple labels may be assigned to a single prompt when several linguistic strategies co-occur.

Table 8: Annotation guidelines used for Study 1 co-annotation.




Figure 6: Averaged correlation strength for each combination of prompt characteristic and code quality marker, with standard error and the number correlations in positive and negative direction across tested models. *NCL*: N# Comment Lines. *CC*: Cyclomatic Complexity, *LLOC*: N# Logical Lines of Code, *LOC*: N# Lines of Code, *SLOC*: N# Source Lines of Code, *MI*: Maintainability Index.

As per Figure 6, informational, involved and inv/inf refers to the linguistic framework developed by (Biber, 1989) and operationalized by (Kedrick et al., 2025). Coefficients were transformed using Fisher’s z-transform (Fisher, 1915) prior to aggregation in order to normalize their sampling distribution and yield a less biased average (Silver and Dunlap, 1987; Corey et al., 1998), and transformed back for reporting. Note that this approach is *not* an attempt to yield an estimate of the true correlation coefficient but primarily a means to be able to concisely report the results of a large amount of correlations for seven tested LLMs. Length here only looks at conversational part of the prompt and no code or other data.

Metric	Score
Cohen’s κ (Annotator 1 vs 2)	0.8737
Cohen’s κ (Annotator 1 vs 3)	0.7847
Cohen’s κ (Annotator 2 vs 3)	0.7388
Fleiss’ κ	0.7985
Krippendorff’s α	0.7992
Exact 3-way Agreement	75.0%

Table 9: Inter-annotator agreement, n=100.




Figure 7: Distribution of Pylint scores for each of the seven tested models. The counts are derived from all three runs.

Type Group	Statistic	p
Direct	$t(26.000) = 2.05$	0.051
Indirect	$t(22.378) = -2.16$	0.051
Impersonal	$t(20.361) = 1.74$	0.099
Personal	$t(25.843) = -1.71$	0.099
Question	$t(25.490) = -0.32$	0.924
Imperative/ Statement	$t(22.947) = 0.10$	0.924

Table 10: Comparisons of the proportion of request types across genders, grouped by directness, use of personal referring, and clause type.

Correlation	Models
Comments x Length [†]	G4o, CL3.7, DS, G4.1, G5
LOC x Length [†]	G4o, G4.1, G5
LLOC x Length [†]	G4o, G4.1
SLOC x Length [†]	G4o, G4.1
MI x Inv/Inf [†]	G5
MI x Length [‡]	G4o, G4.1

Table 11: List of significant correlations and the models for which they were detected. Superscript: (†) indicates significant positive correlations, minus (‡) significant negative correlations.

Model abbreviations: G4o (ChatGPT-4o), CL3.7 (Claude 3.7), CL4 (Claude 4), DS (Deepseek), G4.1 (GPT-4.1), G5 (ChatGPT-5), and Go3 (GPT-o3).

Model	Pylint codes
G4o	C0116, C0103, W0611, C0411, C0301
CL3.7	C0116, W0611, C0411, W0718, C0301
CL4	C0116, C0411, W0611, C0301, C0103
DS	C0116, W0611, C0411, C0301, C0103
G4.1	C0116, C0103, W0611, C0411, W0718
G5	C0103, C0116, C0411, W0611, C0301
Go3	C0116, C0411, W0611, W0718, C0103

Table 12: The top five most prevalent Pylint codes for each model, ordered by descending counts. The codes C0116 (missing function or method docstring), C0103 (variable name does not conform to naming style), W0611 (unused import), and C0411 (standard import should be placed before local import) appear in nearly every model’s top five, though their ranking may vary slightly between models. Occasionally, C0301 (line too long) and W0718 (catching too general exception) also surface among the top spots. Model abbreviations: G4o (ChatGPT-4o), CL3.7 (Claude 3.7), CL4 (Claude 4), DS (Deepseek), G4.1 (GPT-4.1), G5 (ChatGPT-5), and Go3 (GPT-o3).

Request Type	Example Keywords
Interrogative	
You question	<i>can you, could you</i>
I question	<i>can I, do I, can't I, what are my</i>
We question	<i>can we, could we, shouldn't we</i>
Impersonal question	<i>how to, what is</i>
Imperative/Statement	
You command	<i>write a function, do <action></i>
I command	<i>I want, I need, I have to</i>
We command	<i>let's, we have to, we need</i>
Impersonal statement	<i>it has to be like <description></i>

Table 13: Overview of clause types with example signal keywords.

D Study II

In this Appendix, we add the additional tables and visualisation in support of Study II.

Metric	Female	Male
n (submissions)	39	58
Pass rate (%)	38.5	43.1
Radon CC avg	6.61	7.58
Maintainability Index	84.2	81.3
Pylint score	6.91	6.78
Lines of code (LOC)	25.7	25.9

Table 14: Number of submissions and means of code quality metrics by participant gender for Tasks 1 and 2. Pass rates are unit-test success percentages, Radon metrics reflect cyclomatic complexity and maintainability, Pylint scores summarize style and potential defect warnings.

TASK 1

Welcome to task 1. Your goal is to write a Python function that checks the strength of passwords based on specific rules. Here's what makes a password valid. It must be at least 8 characters long. It must contain at least one capital letter and one non-capital letter, one number and one symbol. Lastly, it cannot contain any spaces.

Your function should take a list of passwords and return two lists. The first list contains all valid passwords and the second list contains all invalid passwords. Good luck!

Figure 8: Task 1 of the survey in Study II. This task was presented as an audio recording.

TASK 2

Here is a hashtag-validator. Use an LLM to fix the errors in the code and share the link to the conversation. An example test case is listed down below.

```
def validate_hashtags(text):
    valid = 0
    invalid = 0
    i = 0
    while i < len(text):
        if text[i] == '#':
            j = i + 1
            while j < len(text) and text[j].isalnum():
                j += 1
            tag = text[i:j]
            if len(tag) > 1:
                valid += 1
            else:
                invalid += 1
            i = j
        else:
            i += 1
    return valid, invalid

# Test Case:
# Input: "Valid: #Python, #123_go, #AI
#        Invalid: #Py-thon, #123,go, #
#        Hashtag+"
# Expected Output: (3, 3)
```

Figure 9: Task 2 of the survey in Study II.

E Study III

In this Appendix, we add the additional tables and visualisation in support of Study III.

Self-identification variants:

MALE: You are Jack, a decisive, efficiency-focused senior software engineer.

FEMALE: You are Sarah, a concise, maintainability-focused senior software engineer.

NEUTRAL: You are a senior software engineer performing code reviews.

The code is here to examine:

—
{code}

Respond with exactly one word:

- 'APPROVE' if the code is correct, readable, and maintainable
- 'CHANGES_REQUESTED' otherwise

Figure 10: Prompt utilized to instruct neutral, male- or female-identified LLM reviewer bots to rate code quality.

—
{task description}

—
optional self-identification: Hi, I'm Jack and I'm a computer-science undergrad.

Write a {masc_word_1}, {masc_word_2} solution.

Return only code, no extra text or fences.

Figure 11: Male coded prompt with and without self-identification. Masculine words were randomly sampled from the dictionary (Table 15)

I would like you to

{task description}

optional self-identification: Hi, I'm Sarah and I'm a computer-science undergrad.

Could you please craft a {fem_word_1}, {fem_word_2} solution and provide brief code-comments, without additional text or fences? Thanks!

Figure 12: Female coded prompt with and without self-identification. Feminine words were randomly sampled from the dictionary (Table 15)

{task description}

Provide a solution.
Code only, no extra text or fences.

Figure 13: Neutral prompt used in Study 3.




Figure 14: Mean lines of code by prompt group and provider. Female-coded prompts produce longer solutions on average, consistent with their higher maintainability scores.

Gender	Signal Words
masculine	<i>analytical, streamlined, self-reliant, optimized, robust, decisive, competitive</i>
feminine	<i>cooperative, well-documented, maintainable, readable, clear, elegant, collaborative</i>

Table 15: Agentic and communal adjective inventories adapted from prior work on gendered communication and stereotype framing (Hentschel et al., 2019; Leaper and Ayres, 2007; Leaper and Robnett, 2011)

Task ID	Description
roman2int	Create a function <code>roman_to_int(s)</code> that converts a Roman numeral (1–3999) to an integer.
unique_subsets	Implement <code>get_subsets(nums)</code> that returns all unique subsets of <code>nums</code> (the power set), with each subset sorted and the outer list sorted lexicographically.
isbn13	Write a Python function that validates an ISBN-13 string.
flatten_nested	Write <code>flatten_nested(lst)</code> in python that returns a single flat list from an arbitrarily nested list of lists: <code>[1,[2,[3,4],5],6] → [1,2,3,4,5,6]</code> .
spiral_order	Implement <code>spiral_order(matrix)</code> in python that returns the elements of a 2-D list in clockwise spiral order; e.g. <code>[[1,2,3],[4,5,6],[7,8,9]] → [1,2,3,6,9,8,7,4,5]</code> .

Table 16: Overview of the tasks used in Study 3.

F Ethical board approval

We have obtained an ethics approval for our study based on the regulations of the Ethics Board of the respective faculty of the Technical University Berlin. All human-subject data were collected through voluntary participation with informed consent; they involved no deception or intervention, and posed minimal risk to participants. The tasks consisted of standard programming exercises similar to everyday educational or professional activities, and no sensitive personal data were collected beyond basic demographic information. All analyses were conducted on anonymised data in accordance with applicable data protection regulations. The ethics approval is available on the next page.

ethikkommission@eecs.tu-berlin.de
Ethik-Kommission (EK) der Fakultät IV
Technische Universität Berlin
Sek. MAR 5-5
Marchstr. 23, 10587 Berlin



Berlin, 06.05.2026

Reference number: 424
Submitted at: 06.05.2026
Additional number: -

Scientist in charge:

Dorothea Kolossa

Title/description of the research project/study:

"LLM Code Analysis study"

Statement of Ethics Committee for self disclosure in standardized questionnaire

The aims and procedures of this study – as reported by the scientist(s) in charge in a standardized questionnaire – are in compliance with relevant laws and institutional guidelines. As such they are approved by the institutional ethics committee.

Research participants

- are not expected to take emotional or physical risks,
- are not exposed to a high level of (physical or emotional) stress,
- are fully informed about the purpose and procedures of the study,
- do not belong to a vulnerable group (e.g. children, patients, persons with limited capability to judge)
- do not undergo fMRI (functional Magnetic Resonance Imaging) or TMS (Transcranial Magnetic Stimulation) during the course of the study

The scientist(s) in charge have agreed to protect data privacy, to fully inform participants about the purpose and procedures of the study, and to obtain informed consent prior to the study. This statement presupposes the truthful reply to all questions in the questionnaire and the responsible consideration of all ethical aspects of experimental study design with human participants.

Kind regards,

Ethics Commission Faculty IV
TU Berlin