

Article

Not peer-reviewed version

Prompt Engineering for Accurate Statistical Reasoning with Large Language Models in Medical Research

[Sifiso Vilakati](#) *

Posted Date: 18 June 2025

doi: 10.20944/preprints202506.1497.v1

Keywords: Generative AI; evaluation rubrics; large language models; prompt engineering; medical statistics



Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This open access article is published under a Creative Commons CC BY 4.0 license, which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

Disclaimer/Publisher's Note: The statements, opinions, and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions, or products referred to in the content.

Article

Prompt Engineering for Accurate Statistical Reasoning with Large Language Models in Medical Research

Sifiso Vilakati 

Department of Biostatistics, University of the Free State, South Africa; sifemman@gmail.com

Abstract: The integration of generative artificial intelligence (AI), particularly large language models (LLMs), into medical statistics presents both significant opportunities and critical risks. This paper explores how prompt engineering, defined as the deliberate design of inputs to guide AI behavior, can help mitigate statistical errors in biomedical research. Four prompting strategies are evaluated: zero-shot, explicit instruction, chain of thought, and hybrid approaches. Case studies involving descriptive and inferential statistical tasks show that while zero-shot prompting is generally sufficient for basic summaries, more complex analyses require structured, multi-step prompts to ensure methodological soundness. Among the strategies assessed, hybrid prompting, which combines explicit instructions, reasoning scaffolds, and output formatting, consistently produced the most accurate and interpretable results across two LLMs. The findings emphasize that prompt design, rather than model architecture alone, is the primary determinant of output quality. Although limited access to a broader range of models is a constraint, this study highlights the importance of prompt engineering as a core competency in AI-assisted medical research. It calls for the development of standardized prompt templates, evaluation rubrics, and further studies across diverse statistical domains to support robust and reproducible scientific inquiry.

Keywords: Generative AI; evaluation rubrics; large language models; prompt engineering; medical statistics

1. Introduction

The rapid evolution of generative artificial intelligence (AI) has ushered in a new era for scientific research and knowledge production. From the early days of rule-based natural language processing to the advent of large language models (LLMs) such as OpenAI's GPT series, generative AI has demonstrated remarkable capabilities in understanding, generating, and analyzing human language. These models, trained on vast amount of data, can now summarize complex literature, and even generate code for statistical analysis, tasks that once required years of specialized training [1,2].

The history of generative AI is marked by exponential progress. Early systems relied on hand-crafted rules and limited datasets, producing outputs that were often rigid and contextually shallow [3,4]. The introduction of neural networks and, later, transformer architectures revolutionized the field, enabling models to capture nuanced relationships in language and context [5]. Today's LLMs, with billions of parameters and access to diverse knowledge domains, are increasingly being integrated into the workflows of medical researchers, clinicians, and statisticians [6,7].

Yet, as generative AI becomes more deeply embedded in the fabric of medical research [6], it is essential to reflect on the longstanding challenges that have shaped the field, and chief among them is the misuse of statistics. The integrity of medical science depends on the correct application and interpretation of statistical methods. However, a substantial body of evidence has shown that statistical errors are both common and consequential in the biomedical literature.

A series of landmark investigations have illuminated the scope and persistence of statistical misapplication in biomedical research. In 1994, Douglas Altman's influential editorial in the BMJ

exposed the widespread prevalence of basic statistical errors in published studies, even in leading journals, and underscored the urgent need for better statistical education and more rigorous peer review [8]. A decade later, John Ioannidis's seminal paper in PLOS Medicine fundamentally challenged the credibility of the biomedical literature, demonstrating that a confluence of factors such as small sample sizes, selective reporting, flexible study designs, and, crucially, statistical misapplication renders most published research findings likely to be false [9]. More recent systematic reviews have confirmed that issues like p-hacking, selective reporting, and the misuse of p-values remain stubbornly prevalent in scientific literature [10–12].

As generative AI becomes an important tool in medical research, there is a genuine risk that these longstanding issues may not only persist but be amplified. The automation of statistical analysis and interpretation by AI models, if guided by poorly constructed prompts or used by individuals lacking statistical expertise, could lead to the rapid and widespread dissemination of erroneous results. AI-generated outputs, often presented with an air of authority and fluency, may make it even more difficult to detect subtle errors or misinterpretations. In this context, the responsible design and use of prompts, prompt engineering, emerges as a critical safeguard.

This paper critically examines the intersection of prompt engineering, generative AI, and medical statistics, with a focus on minimizing the risk of erroneous outputs that may exacerbate the long-standing issue of statistical misuse in biomedical research. By situating this discussion within the well-documented history of statistical errors, the paper underscores the importance of understanding and applying prompt engineering to guide AI models toward accurate and reliable results. Particular attention is given to the challenge of preventing AI-generated hallucinations and ensuring that outputs support, rather than undermine, sound statistical practice. To address these concerns, the review provides practical suggestions and general guidelines for avoiding erroneous responses when conducting basic statistical analyses and statistical inference using generative AI. In doing so, the paper aims to equip researchers with strategies that promote the responsible and effective integration of AI into medical research, safeguarding the integrity of scientific discovery. To the best of our knowledge, no prior study has systematically examined the role of prompt engineering in mitigating statistical errors within the context of generative AI applications in medical research. Given the increasing reliance on AI for statistical analysis, this oversight represents a critical gap in the literature. This paper seeks to bridge that gap by offering a structured evaluation of prompting strategies and their implications for statistical validity and interpretability.

2. Materials and Methods

2.1. Prompt Engineering

As generative AI systems become increasingly sophisticated, the art and science of “prompt engineering” has emerged as a central discipline for harnessing their capabilities. Prompt engineering refers to the deliberate design and formulation of input queries, prompts, to guide LLMs toward producing accurate, relevant, and contextually appropriate outputs [13]. In essence, prompt engineering is the interface between human intent and machine intelligence, shaping how AI interprets and responds to complex tasks. A well-crafted prompt is more than a simple question; it is a structured communication that encodes the user's objectives, constraints, and expectations. The quality of a prompt can dramatically influence the reliability and utility of AI-generated responses, particularly in high-stakes domains such as medical statistics. As such, prompt engineering is not merely a technical exercise, but a critical component of responsible AI deployment [14,15].

2.2. Guidelines and Structure for Effective Prompts

The literature on prompt engineering has converged on several key guidelines for constructing effective prompts. Clarity is paramount: prompts should be unambiguous, specific, and free from unnecessary complexity. Contextual information such as the desired format of the output, relevant background, or explicit instructions should be included to minimize misinterpretation. Additionally,

prompts should anticipate potential sources of confusion, such as ambiguous terminology or multiple possible interpretations, and address them directly [16].

Structurally, effective prompts often include:

- A clear statement of the task or question.
- Any necessary context or background information.
- Explicit instructions regarding the format or style of the output.
- Stepwise guidance for multi-part or complex tasks.

These elements help ensure that the model's output aligns with the user's expectations and the requirements of the task at hand.

2.3. Approaches to Prompt Engineering

Prompt engineering encompasses a spectrum of strategies, each with distinct strengths and limitations. The choice of approach depends on the complexity of the task, the desired level of control, and the context in which the AI is deployed. Understanding these methods is essential for leveraging generative AI effectively in medical statistics.

2.3.1. Zero-Shot Prompting

Zero-shot prompting is a foundational approach for interacting with large language models, where the model is given a single, direct instruction without any examples. For instance, a prompt might read, "Summarize the results of this clinical trial." In this setting, the model relies entirely on its pre-trained knowledge and generalization abilities to interpret and respond to the task [14,17,18].

The primary advantage of zero-shot prompting is its simplicity and efficiency. It requires minimal effort from the user and can be effective for well-defined, common tasks. However, its major limitation is unpredictability: without examples or additional context, the model may misinterpret the intent or produce outputs that lack the necessary specificity or rigor, especially in specialized domains like medical statistics [14,18].

2.3.2. Few-Shot Prompting

Few-shot prompting extends the zero-shot approach by supplying the model with a small number of examples that demonstrate the desired input-output relationship. For instance, a prompt might present two or three sample statistical analyses, each paired with its corresponding summary or interpretation, followed by a new case for the model to address in the same manner. By explicitly showing the model how to perform the task, few-shot prompting helps guide its responses toward the expected format and content [14,18].

This method offers greater control over the model's output and can substantially enhance both the relevance and accuracy of responses, particularly for tasks that are less common or more complex. The inclusion of well-chosen examples allows the model to better understand the nuances of the task and adapt its output accordingly. However, the effectiveness of few-shot prompting depends heavily on the quality and representativeness of the examples provided. If the examples are ambiguous or not closely aligned with the intended task, the model's performance may suffer. Additionally, this approach may not scale efficiently to tasks with high variability, as it becomes challenging to provide examples that cover all possible scenarios [14,18].

2.3.3. Explicit, Instruction-Based Prompting

Explicit, instruction-based prompting takes a more structured approach by providing the model with clear, detailed, and stepwise instructions. This method not only specifies the task but also outlines the sequence of actions and any necessary checks or conditions. For example, a prompt might state:

"First, check whether the data meet the assumptions for a t-test (normality and equal variances). If the assumptions are met, perform the t-test. If not, select and perform an appropriate alternative test. Report the results, including the test statistic, degrees of freedom, and p-value."

By breaking down the workflow into explicit steps, this approach reduces ambiguity and guides the model through complex analytical processes. It helps ensure that critical steps such as verifying statistical assumptions are not overlooked, which is particularly important in medical and health research where methodological rigor is essential.

The primary advantage of explicit, instruction-based prompting is its ability to minimize errors of omission and misinterpretation. By clearly delineating each step, the model is less likely to skip important procedures or misapply statistical methods. However, this approach does require the user to have a solid understanding of the statistical process, as the quality and completeness of the instructions directly influence the reliability of the model's output [19,20]. This method is especially valuable in domains where precision and adherence to methodological standards are critical, supporting more reliable and transparent statistical analysis

2.3.4. Chain-of-Thought Prompting

Chain-of-thought (CoT) prompting is a technique that encourages the model to articulate its reasoning process step by step, making each intermediate stage of analysis explicit. For example, a prompt might instruct:

"Explain your reasoning step by step before giving the final answer."

This approach is particularly effective for complex analytical tasks, as it reveals the model's logic and allows users to trace the sequence of decisions or calculations that lead to the final output [17,21]. By surfacing the model's reasoning, CoT prompting can help identify errors or gaps in logic, supporting more robust and transparent analyses.

The primary advantage of chain-of-thought prompting is the increased transparency and reliability it brings to multi-step problems. Users can review each stage of the model's reasoning, making it easier to spot mistakes or misunderstandings. However, this method can result in verbose outputs and does not guarantee that every reasoning step is correct; the model may still make errors in logic or interpretation.

2.3.5. Format-Constrained Prompting

Format-constrained prompting directs the model to present its output in a specific structure, such as a table, list, or code block. For instance, a prompt might state:

"Present the results in a table with headings: test statistic, df, p-value."

This approach is especially valuable in medical statistics, where standardized reporting is essential for clarity and reproducibility. By specifying the desired format, format-constrained prompting minimizes the risk of unstructured or incomplete outputs and helps ensure that all relevant information is included. Nevertheless, the effectiveness of this method still depends on the model's underlying knowledge and its ability to follow formatting instructions [19,20].

We summarize the different prompting approaches in Table 1.

2.4. Variants and Hybrids in Prompt Engineering

As the field of prompt engineering has matured, practitioners have developed a range of variants and hybrid strategies that blend the strengths of multiple approaches [22]. These innovations are particularly relevant in medical statistics, where both methodological rigor and clarity of communication are paramount [23].

One important variant is zero-shot chain-of-thought prompting. Here, the user asks the model to "think step by step" even without providing examples. This simple instruction can significantly improve the model's reasoning, especially for complex statistical tasks. Another variant is zero-shot with context, where the prompt includes relevant background information or definitions to anchor the model's response more effectively [17].

Few-shot prompting has also evolved. In contextual calibration, the examples provided are carefully selected to match the difficulty or style of the target task, ensuring the model's output is both

Table 1. Comparison of prompt engineering approaches

Approach	Description	Advantages	Disadvantages
Zero-shot prompting	Direct instruction, no examples	Simple, efficient	May lack specificity, unpredictable
Explicit, Instruction-Based	Clear, stepwise instructions	Reduces ambiguity, minimizes errors of omission	Requires user expertise, can be verbose
Few-Shot Prompting	Provides a few examples	Greater control, improved relevance	Sensitive to example quality, may not scale
Chain-of-Thought Prompting	Encourages stepwise reasoning	Transparent, better for complex tasks	Verbose, reasoning may not always be correct
Format-Constrained Prompting	Specifies output structure	Standardized, easy to interpret	Still relies on model's statistical knowledge

relevant and accurate. Dynamic few-shot prompting adapts the set of examples in real time based on the model’s previous outputs, allowing for a more responsive and tailored interaction [17,24].

Chain-of-thought prompting can be combined with few-shot learning to create few-shot chain-of-thought prompts. In this approach, the model is shown several examples of stepwise reasoning, which helps it internalize the logic and structure required for complex analyses. Format-constrained prompting can also be hybridized. For instance, a prompt might combine explicit instructions, a required output format (such as a table), and a few illustrative examples. This multi-layered approach maximizes clarity and minimizes the risk of misinterpretation [25].

2.5. Limitations of Prompt Engineering

While the evolution of prompt engineering has introduced powerful variants and hybrid strategies, several limitations remain that constrain its broader applicability and reliability. These limitations are particularly salient in domains like medical statistics, where precision, reproducibility, and interpretability are essential [26].

One key limitation is the lack of generalizability. Prompts that perform well in one context may fail in another, especially when transferred across models or domains. This brittleness is compounded by the sensitivity of large language models to minor changes in phrasing, formatting, or context [27].

Prompt engineering also relies heavily on trial-and-error. Designing effective prompts often requires iterative tuning, domain expertise, and manual curation of examples. This process can be time-consuming and may not scale efficiently for complex or evolving tasks [13].

Another challenge is the opacity of model behavior. Even with carefully crafted prompts, the internal reasoning of the model remains largely inaccessible, making it difficult to diagnose errors or ensure consistency. This lack of transparency can be problematic in high-stakes applications where interpretability is critical [27].

Moreover, prompt-based methods are susceptible to bias and hallucination. Without external grounding or verification, models may generate plausible-sounding but incorrect or biased outputs, especially when prompts are ambiguous or under-specified [19,28].

Finally, while hybrid strategies offer improved performance, they often require more complex prompt structures and greater user effort. This can introduce new sources of error and reduce usability, particularly for non-expert users [21,29].

2.6. Commonly Used Statistics in Medical Research

Descriptive and inferential statistics form the backbone of quantitative analysis in medical and health research. At the most fundamental level, descriptive statistics are employed to summarize and characterize data. Measures such as the mean, median, and mode are routinely used to describe central tendency, while standard deviation, interquartile range, and variance provide insight into

data dispersion [30]. However, the choice of summary statistic must be guided by the underlying distribution of the data. For instance, while the mean is a widely reported measure, it can be misleading in the presence of skewed distributions or outliers, where the median often provides a more robust summary [31]. Similarly, graphical representations such as histograms and boxplots are essential for visualizing data distributions and identifying anomalies.

Beyond description, inferential statistics enable researchers to draw conclusions about populations based on sample data. Commonly used inferential methods include hypothesis testing, confidence intervals, and regression analysis. The t-test and analysis of variance (ANOVA) are frequently applied to compare group means, while non-parametric alternatives such as the Mann-Whitney U test or Kruskal-Wallis test are appropriate when data do not meet the assumptions of normality or homogeneity of variance [32]. Correlation and regression analyses are used to assess relationships between variables, with linear regression being a staple for modeling continuous outcomes and logistic regression for binary outcomes [33].

Despite their ubiquity, the application of these statistical methods is not without challenges. A common pitfall is the inappropriate use of parametric tests, such as the t-test, on data that violate key assumptions, leading to invalid inferences. The misuse of p-values, that is, interpreting statistical significance as evidence of clinical importance or failing to account for multiple comparisons remains a persistent issue in the literature [34]. Additionally, selective reporting of significant results and p-hacking can distort the scientific record and undermine the credibility of research findings [35]. Table 2 summarizes some of the most frequent statistical pitfalls encountered in medical research.

Table 2. Common statistical pitfalls in medical research.

Pitfall	Description
Misuse of mean	Reporting the mean for skewed data instead of the median, which can misrepresent the true central tendency.
Inappropriate test selection	Using parametric tests (e.g., t-test, ANOVA) on data that do not meet assumptions of normality or equal variances, rather than opting for non-parametric alternatives.
P-value misinterpretation	Treating statistical significance as equivalent to clinical or practical significance, or misunderstanding what a p-value actually represents.
Multiple comparisons	Conducting multiple hypothesis tests without proper adjustment, which increases the risk of false-positive findings.
Selective reporting	Only presenting statistically significant results, leading to publication bias and an incomplete scientific record.

A careful and context-aware application of statistical methods is essential for producing valid and reproducible results in medical research. By recognizing the strengths and limitations of commonly used statistics, researchers can avoid common errors and contribute to the advancement of robust scientific knowledge.

2.7. Evaluating Large Language Models

The increasing integration of large language models (LLMs) into biomedical and statistical workflows necessitates a rigorous and context-sensitive approach to evaluating their outputs. As these models are entrusted with tasks ranging from summarizing clinical findings to performing inferential statistical analyses, the question of how to assess their performance becomes not only technical but epistemological. Evaluation is no longer a peripheral concern; it is central to ensuring that LLMs function as reliable collaborators in scientific inquiry.

Historically, the evaluation of natural language generation (NLG) systems has relied on metrics that prioritize surface-level similarity to reference texts, such as BLEU and ROUGE. These metrics, while useful for tasks like translation or summarization, operate primarily on n-gram overlap and often fail to capture semantic correctness or contextual appropriateness [36]. However, the emergence of LLMs which are capable of generating fluent, contextually rich, and often persuasive outputs has

exposed the limitations of such metrics, particularly in tasks requiring reasoning or domain-specific accuracy [36]. In high-stakes domains like medical statistics, where the correctness of an output cannot be inferred from its fluency alone, evaluation must account for methodological rigor, interpretive accuracy, and domain relevance. This section reviews the principal approaches to LLM evaluation, with a focus on their applicability to prompt engineering in statistical contexts.

2.7.1. Manual Evaluation

Manual evaluation involves human experts assessing the quality of LLM outputs based on predefined rubrics or subjective judgment. This approach is particularly valuable in domains requiring nuanced interpretation, such as statistical reasoning or clinical decision-making. Human evaluators can detect subtle errors, assess contextual appropriateness, and apply domain-specific standards that are difficult to encode algorithmically [13].

However, manual evaluation is inherently limited by its subjectivity, labor intensity, and lack of scalability. Inter-rater variability can compromise reliability, and the process is often too slow for iterative prompt development or real-time deployment [13]. Despite these limitations, manual evaluation remains indispensable for validating automated methods and for assessing outputs in novel or ambiguous contexts where ground truth is unavailable or contested [37].

2.7.2. Automated Evaluation

Automated evaluation methods offer scalability and reproducibility, making them attractive for benchmarking and continuous integration. These methods fall into two broad categories: reference-based and reference-free.

Reference-based metrics such as BLEU, ROUGE, and METEOR compare model outputs to predefined reference texts, quantifying lexical overlap. While these metrics are useful for tasks like summarization or translation, they are poorly suited to evaluating statistical reasoning. Lexical similarity does not guarantee semantic correctness, and models may produce outputs that are superficially similar to reference texts while misrepresenting statistical assumptions or misapplying analytical methods [38].

Reference-free methods address some of these limitations by evaluating outputs without relying on ground truth references. These include semantic similarity scoring using embedding models, rule-based validation of output structure, and the increasingly prominent approach of using LLMs themselves as evaluators; a method known as “LLM-as-a-judge” [38].

2.7.3. LLM-as-a-Judge

The LLM-as-a-judge approach involves prompting a language model to assess the outputs of another (or the same) model based on a rubric or set of criteria. This method is particularly well-suited for evaluating prompt effectiveness, as it allows for rapid, rubric-based comparison of different prompt formulations [39]. For example, an evaluator model can be instructed to assess outputs based on assumption checking, test selection, output completeness, and interpretive quality—criteria that are central to statistical validity.

This approach offers several advantages. It is highly scalable, enabling the evaluation of thousands of prompt-output pairs without human intervention. It also allows for nuanced assessments that align more closely with human judgment, particularly when the evaluator model is guided by explicit instructions or chain-of-thought reasoning [40]. Moreover, it facilitates iterative prompt refinement by providing structured feedback on the strengths and weaknesses of different prompt designs.

However, the method is not without limitations. A primary concern is the potential for bias and circularity, especially when the evaluator shares architecture or training data with the model being evaluated. This can lead to inflated performance estimates or the reinforcement of shared misconceptions [39]. Additionally, the internal reasoning of the evaluator model is often opaque, making it difficult to audit or interpret its judgments, an issue that is particularly problematic in clinical and statistical domains where transparency is essential [40].

Another challenge lies in domain specificity. While general-purpose LLMs may perform well in evaluating outputs related to everyday language tasks, their ability to assess domain-specific outputs such as the appropriateness of a statistical test or the validity of an inferential conclusion depends on their exposure to relevant training data and their capacity for domain-specific reasoning. In such cases, hybrid evaluation pipelines that combine LLM-based and human assessments may offer a more robust solution [39].

Recent empirical studies support the utility of LLM-as-a-judge in scientific contexts. Gao et al. (2025) demonstrated that LLM-based evaluators can reliably distinguish between high- and low-quality outputs in natural language generation tasks, particularly when guided by structured rubrics [37]. Similarly, Kojima et al. (2022) and Wei et al. (2022) have shown that chain-of-thought prompting enhances the reasoning capabilities of LLMs, both as generators and as evaluators [17,21].

2.7.4. Hybrid Evaluation: Combining LLM-as-a-Judge with Human Oversight

As large language models (LLMs) become increasingly embedded in scientific and statistical workflows, the need for evaluation frameworks that balance scalability with interpretive rigor has become more urgent. While automated methods such as LLM-as-a-judge offer efficiency and consistency, they are not immune to issues of bias, opacity, and domain mismatch. Conversely, manual evaluation provides depth and contextual sensitivity but lacks scalability and reproducibility. To address these limitations, recent research has proposed hybrid evaluation frameworks that integrate both approaches in a complementary manner [41].

In a hybrid framework, the evaluation process typically unfolds in two stages. First, an LLM is prompted to assess the outputs of another model or the same model under different prompt conditions using a structured rubric. This rubric may include criteria such as assumption checking, test selection, output completeness, and interpretive quality. The LLM's evaluations are recorded and scored, providing a scalable and rubric-aligned first-pass assessment of prompt effectiveness [42].

In the second stage, human experts review the LLM's evaluations. This paper serves to validate the model's judgments, identify potential hallucinations or misinterpretations, and refine the evaluation rubric based on observed model behavior. This human-in-the-loop process is particularly valuable in domains like medical statistics, where subtle errors in reasoning or misapplication of statistical methods can have significant consequences. The hybrid approach thus combines the breadth of automated evaluation with the depth of expert review, enabling both high-throughput assessment and epistemic accountability [43].

This framework has been shown to improve the reliability and transparency of LLM evaluations, particularly when applied to tasks involving complex reasoning or domain-specific knowledge. It also supports iterative prompt refinement, as human reviewers can use the LLM's feedback to identify patterns of failure or success across different prompt formulations. Moreover, by documenting both the LLM's evaluations and the human corrections, the framework promotes reproducibility and provides a valuable audit trail for future research and model development [41]. A comparison of the LLM evaluation approaches is given in Table 3.

In this paper, we adopt this hybrid evaluation framework to assess the effectiveness of various prompt engineering strategies for statistical analysis. Specifically, LLMs will be used to evaluate outputs generated under different prompting conditions, and these evaluations will be reviewed and annotated by the author. This approach ensures that the final prompt recommendations are not only computationally validated but also aligned with best practices in statistical reasoning and biomedical research.

Table 3. Comparison of LLM Evaluation Approaches

Evaluation Approach	Description	Advantages	Disadvantages
Manual evaluation	Human experts assess the quality of LLM outputs based on predefined rubrics or subjective judgment.	Provides depth and contextual sensitivity; can detect subtle errors and apply domain-specific standards.	Subjective, labor-intensive, and lacks scalability; inter-rater variability can compromise reliability.
Reference-based evaluation	Metrics such as BLEU, ROUGE, and METEOR compare model outputs to predefined reference texts, quantifying lexical overlap.	Objective and reproducible; suitable for benchmarking and regression testing.	Poorly suited to evaluating statistical reasoning; lexical similarity does not guarantee semantic correctness.
Reference-free evaluation	Evaluates outputs without relying on ground truth references, using techniques like semantic similarity scoring, rule-based checks, or LLMs as evaluators.	Flexible and adaptable to real-world use cases; enables continuous monitoring and quality control.	Challenges related to reliability, interpretability, and standardization; embedding-based measures can be opaque.
LLM-as-a-Judge	An LLM is prompted to assess the outputs of another (or the same) model based on a rubric or set of criteria.	Highly scalable; allows for nuanced, rubric-based assessments; facilitates iterative prompt refinement.	Potential for bias and circularity; internal reasoning is often opaque; effectiveness varies across domains.
Hybrid Evaluation	Combines LLM-as-a-Judge with human oversight, where LLMs provide initial evaluations and human experts review and refine these assessments.	Balances scalability with interpretive rigor; supports high-throughput assessment and epistemic accountability.	Requires more complex prompt structures and greater user effort; potential for new sources of error.

3. Results

If not prompted correctly, generative AI systems can produce misleading or incomplete responses. The quality and reliability of outputs from LLMs are highly dependent on the clarity and specificity of the input prompts. As such, effective prompt design is essential for obtaining accurate and contextually appropriate results.

With a growing number of generative AI models available, it is the user’s responsibility to select the most suitable model for the task at hand. Models such as OpenAI’s GPT-4, Anthropic’s Claude, and Microsoft Copilot have demonstrated strong capabilities in statistical reasoning, as well as in generating code for languages like R and Python. These models can be powerful tools for data analysis, but their effectiveness hinges on how well they are guided through prompting.

3.1. Descriptive Statistics

Descriptive statistics are typically the first step in data analysis, especially for numerical datasets. Measures of central tendency—such as the mean and median—and measures of dispersion—such as the standard deviation and interquartile range (IQR) are commonly used to summarize data characteristics.

For these tasks, we recommend the use of zero-shot prompting. A well-structured instruction specifying the required statistical measures is often sufficient to elicit accurate and complete responses from LLMs. This approach assumes that the user has a basic understanding of statistical concepts and knows what they intend to compute.

We also advise users to assess the distribution of the data particularly its normality before interpreting descriptive statistics. This can be done using a simple zero-shot prompt.

To illustrate, consider the following artificial dataset:

56.31, 361.21, 158.01, 109.55, 20.35, 20.35, 7.18, 241.35,
110.29, 147.75

This dataset is clearly right-skewed. There are multiple ways to prompt an LLM to analyze such data. One effective zero-shot prompt is:

Calculate the mean, median, standard deviation and the interquartile range for the data: 56.31, 361.21, 158.01, 109.55, 20.35, 20.35, 7.18, 241.35, 110.29, 147.75. Is the data skewed?

This prompt is explicit and comprehensive. It instructs the model to compute all relevant statistics and includes a final question to guide interpretation specifically whether the mean or median is more appropriate to report.

Alternatively, a more abstract yet effective prompt is:

Describe the data using measures of central tendency and dispersion. Looking at the distribution of the data, which measures should I report?

We prefer this formulation. It is specific in its request for descriptive statistics but does not require the user to name individual metrics. This makes it accessible to users with limited statistical background. Moreover, it leverages the model's reasoning capabilities to recommend appropriate measures based on the data distribution. This prompt was tested across multiple LLMs, including Claude 3.7 Sonnet, GPT-4.1, and Microsoft Copilot, and consistently produced accurate and context-aware responses.

Other descriptive measures such as range, variance, skewness, and kurtosis can also be computed using similar zero-shot prompts. We recommend the use of zero-shot prompting for generating descriptive statistics, as it is both efficient and reliable when the instructions are clearly formulated. Providing explicit and unambiguous instructions minimizes the risk of erroneous outputs and ensures that the statistics generated are appropriate for the intended context, whether for academic research, reporting, or exploratory data analysis.

The overarching goal is to avoid misleading interpretations and to ensure that the statistical summaries produced by the language model are both accurate and contextually relevant.

3.2. Inferential Statistics

Caution is warranted when using generative AI for statistical inference in medical research. While LLMs can assist with data analysis, they are prone to producing erroneous outputs, particularly in inferential contexts. A key limitation lies in their handling of statistical assumptions, a foundational aspect of most inferential procedures.

Nearly all statistical models used for inference rely on specific assumptions (e.g., normality, homogeneity of variance, independence). When these assumptions are violated, alternative methods such as data transformation or non-parametric tests should be employed. For instance, the independent samples Student's t-test assumes normality and equal variances; if these are not met, a non-parametric alternative like the Mann-Whitney U test is more appropriate. Generative AI models may overlook or misinterpret these assumptions unless explicitly guided.

3.2.1. Prompt Evaluation for Statistical Inference

To assess the effectiveness of the various prompting strategies, we employed a hybrid evaluation framework. Two illustrative case studies are presented. In the first scenario, the assumption of equal variances is violated, necessitating the use of a t-test that accommodates unequal variances. In the second example, both the assumptions of normality and homogeneity of variances are breached, thereby warranting the application of a non-parametric alternative to the t-test. Table 4 outlines the specific prompts utilized in the first case.

Table 4. Prompts used for Illustration I

Prompt type	Prompt
Zero-shot	Perform a t-test on the following blood pressure data from a clinical trial. [data].
Explicit instruction	I want to compare blood pressure between two groups. [data]. Step 1: Test for normality in each group. Step 2: Test for equality of variances. Step 3: Based on the results, choose the appropriate t-test Step 4: Perform the test and report the test statistic, degrees of freedom, and p-value. Step 5: Interpret the result in plain language.
CoT	Consider the following blood pressure data. [data]. Think step by step: • Are the data normally distributed? • Are variances equal? • Which test should be used? • What are the results? • What do they mean?
Hybrid	Analyze the following blood pressure data from a clinical trial. [data]. Please: • Test for normality and equal variances. • Choose the correct test (Student’s or Welch’s t-test). • Report results in a table with columns: test statistic, df, p-value. • Provide a plain-language interpretation of the findings.

3.2.2. Illustrations

Illustration I

Artificial data are generated for a blood pressure clinical trial where 40 patients were assigned between a new drug and a control treatment in the ratio 1:1. The data generation process is as follows:

```
control_bp <- rnorm(n_control, mean=150, sd=10)
treatment_bp <- rnorm(n_treat, mean=140, sd=20)
```

and the generated data is:

```
control_bp = c(142.9, 152.6, 147.5, 146.5, 140.5, 149.5, 142.2, 133.3,
               146.2, 159.2, 144.2, 156.1,133.8, 149.4, 155.2, 153.0, 151.1,
               143.6, 141.5, 139.8)
treatment_bp = c(122.4, 101.1, 110.2, 114.9, 156.9, 107.0, 124.7, 121.6,
                 100.8, 118.6, 148.9, 129.0, 120.8, 111.6, 78.9, 142.6,
                 90.8, 134.8, 158.2, 91.1)
```

Illustration II

We further demonstrate the impact of prompt engineering using some artificial BMI data from a 12-week randomized controlled trial comparing a plant-based diet intervention to standard care among primary school children with elevated BMI. The data are as follows:

```
control_bmi = c(22.3, 23.3, 31.2, 24.2, 24.4, 29.1, 33.8, 20.2, 21.9, 22.7,
                27.7, 25.1, 25.2, 24.3, 52.3)
treatment_bmi = c(23.1, 21.3, 24.2, 24.2, 24.1, 23.7, 23.4, 21.8, 21.2, 21,
                  20.3, 21.5, 18.8, 27.4, 25)
```

Table 6 shows the different prompts used in the second case study.

The evaluation results presented in Table 5 provide a comparative analysis of how different prompting strategies influence the performance of LLMs in conducting inferential statistical analyses.

Table 5. Evaluation summary by prompt type and model for Illustration I and Illustration II

Criterion	Zero-shot		Explicit instruction		Chain-of-Thought		Hybrid	
	GPT-4.1	Claude 3.7 Sonnet	GPT-4.1	Claude 3.7 Sonnet	GPT-4.1	Claude 3.7 Sonnet	GPT-4.1	Claude 3.7 Sonnet
Illustration I								
Assumption checking	0.0	0.0	1.0	1.0	1.0	1.0	1.0	1.0
Test selection	0.6	0.6	1.0	1.0	1.0	1.0	1.0	1.0
Output completeness	1.0	1.0	1.0	1.0	0.8	0.8	0.8	0.8
Interpretive quality	0.8	1.0	1.0	1.0	1.0	1.0	1.0	1.0
Total score	2.4	2.6	4.0	4.0	3.8	3.8	3.8	3.8
Illustration II								
Assumption checking	0.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
Test selection	0.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
Output completeness	0.6	1.0	1.0	1.0	1.0	1.0	1.0	1.0
Interpretive quality	0.6	1.0	1.0	1.0	1.0	1.0	1.0	1.0
Total score	1.2	4.0	4.0	4.0	4.0	4.0	4.0	4.0

These tables summarize the outputs of two LLMs; GPT-4.1, and Claude 3.7 Sonnet. Microsoft Copilot was used as a judge to evaluate the outputs from GPT-4.1, and Claude 3.7 Sonnet. The two AI models were tasked with analyzing two artificial datasets using four distinct prompting strategies: zero-shot, explicit instruction-based, chain-of-thought, and hybrid prompting. The evaluation was conducted using a structured rubric that assessed four key dimensions of statistical reasoning: assumption checking, test selection accuracy, output completeness, and interpretive quality. The rubric is shown in Appendix A.

Table 6. Prompts used for Illustration II

Prompt type	Prompt
Zero-shot	Analyze the following BMI data from a randomized controlled trial. [data].
Explicit instruction	I want to compare BMI between two groups. [data]. Step 1: Test for normality in each group. Step 2: Test for equality of variances. Step 3: If assumptions are violated, use a non-parametric test . Step 4: Report the test statistic and p-value. Step 5: Interpret the result in plain language.
CoT	Consider the following BMI data. [data]. Think step by step: • Are the data normally distributed? • Are variances equal? • Which test should be used? • What are the results? • What do they mean?
Hybrid	Analyze the following BMI data from a 12-week dietary intervention trial. [data]. Please: • Test for normality and equal variances. • Can a t-test be used? • Perform an appropriate test. • Present results in a table with: test name, test statistic, p-value. • Provide a plain-language interpretation.

In both case studies, the zero-shot prompting strategy consistently yielded the lowest scores across all models. For the blood pressure dataset (second part of Table 5), zero-shot prompts failed to elicit any assumption checking from the models, resulting in a score of 0.0 for that criterion. Test selection was partially correct in some instances, with scores ranging from 0.6 to 1.0, but the absence of diagnostic checks undermined the methodological soundness of the analyses. Output completeness and interpretive quality were somewhat better, with scores ranging from 0.8 to 1.0, indicating that while the models could produce fluent and structured outputs, these outputs were not always grounded in appropriate statistical reasoning.

In contrast, the explicit instruction-based, chain-of-thought, and hybrid prompting strategies consistently achieved higher scores across all criteria. These strategies prompted the models to perform assumption checks, select appropriate statistical tests based on those checks, and provide comprehensive and contextually relevant interpretations. For instance, in the BMI dataset (Table 5), all three models correctly identified the non-normal distribution in the control group and the unequal variances between groups when guided by explicit or hybrid prompts. Consequently, they selected the Mann–Whitney U test, a non-parametric alternative appropriate for the data structure. These prompts resulted in perfect scores (4.0) across all evaluation dimensions, demonstrating the effectiveness of structured and context-rich prompting in eliciting statistically valid outputs from LLMs. Table 7 provides the summary of assumption checks and statistical test results for BMI data.

Table 7. Summary of assumption ahecks and statistical test results for BMI data

Test	Data	Statistic	p-value	Interpretation
Shapiro-Wilk (Control)	Original	$W = 0.701$	0.00025	Not normal
Shapiro-Wilk (Treatment)	Original	$W = 0.970$	0.853	Normal
F-test for variances	Original	$F = 13.127$	2.09×10^{-5}	Variances unequal
Mann-Whitney U	Original	$W = 171$	0.0161	Significant

3.2.3. A Prompting Strategy for Inferential Problems

We propose a structured prompting strategy to guide the use of generative AI in statistical inference, particularly in sensitive domains such as medical research. The goal is to minimize errors and ensure that statistical procedures are applied appropriately. The following steps outline a recommended workflow:

- Upload the dataset to the chosen AI assistant or LLM interface, provide context.
- Select the appropriate statistical test for the research question. If uncertain, the user may prompt the AI for guidance on test selection.
- Test the assumptions underlying the selected statistical procedure (e.g., normality, homogeneity of variance, independence).
- If assumptions are met, proceed with the selected test.
- If assumptions are violated, prompt the AI to suggest and evaluate alternative procedures (e.g., non-parametric tests).
- If the alternative test’s assumptions are met, perform the alternative test.
- Specify the desired output format, such as summary tables, p-values, confidence intervals, or effect sizes.
- Request an interpretation of the results in plain language, suitable for the intended audience.
- Validate the results by cross-checking with statistical software or consulting a domain expert.

The above workflow is summarized in the Table 8. The process begins with a clear definition of the research question or analytical objective. This foundational step ensures that all subsequent actions are aligned with the intended goal. Once the objective is established, the next step is to select a candidate statistical model or analytical approach that is appropriate for the data and the research question at hand.

Table 8. Summary of the prompt engineering steps for inferential problems

Step	Description	Decision/Next action
1	Define research question or analytical objective and upload data into AI assistant	—
2	Select candidate statistical model or analytical approach	—
3	Test assumptions of the selected model	Are assumptions satisfied?
4	If yes: Formulate precise, context-rich prompt for the AI	Proceed to prompt construction
5	If no: Search for alternative model and repeat assumption testing	Iterate until suitable model is found
6	Submit prompt to generative AI system	—
7	Evaluate AI output for accuracy, completeness, and format adherence	Is output satisfactory?
8	If yes: Incorporate output into research workflow	—
9	If no: Revise prompt and resubmit	Repeat as necessary

After selecting a model, it is essential to rigorously test its underlying assumptions. This step acts as a critical checkpoint: if the assumptions are satisfied, the process moves forward to constructing a precise and context-rich prompt for the AI. This prompt should clearly specify the statistical method, the expected output format, and any relevant contextual details to guide the AI’s response.

If, however, the model’s assumptions are not met, the process does not proceed. Instead, an alternative model is sought, and its assumptions are tested in turn. This iterative process continues until a suitable model is identified.

Once a suitable model is confirmed and the prompt is constructed, the prompt is submitted to the generative AI system. The output generated by the AI is then carefully evaluated for accuracy, completeness, and adherence to the specified format. If the output meets the required standards, it

is incorporated into the research workflow. If not, the prompt may be revised for greater clarity or specificity, and the process is repeated as necessary.

This structured approach ensures methodological rigor and effective communication with generative AI systems, supporting robust and transparent scientific analysis.

4. Discussion

The results of the two illustrative case studies underscore a critical insight: zero-shot prompting is insufficient for conducting inferential statistical analyses using generative AI. While zero-shot prompts may be adequate for basic descriptive tasks, as demonstrated in earlier sections of this paper, they fall short when applied to more complex analytical workflows that require methodological rigor. The absence of assumption checking in almost all zero-shot outputs is particularly concerning, as it reflects a fundamental gap in the models' ability to autonomously initiate diagnostic procedures without explicit instruction. This limitation aligns with findings from Kojima et al. (2022), who observed that large language models often default to surface-level heuristics in the absence of structured guidance, leading to shallow or incorrect reasoning in complex tasks [17].

The superior performance of explicit instruction-based, chain-of-thought, and hybrid prompting strategies highlights the importance of prompt specificity and structure in guiding LLMs toward valid statistical reasoning. These strategies not only improved the accuracy of test selection but also enhanced the completeness and interpretive quality of the outputs. The inclusion of stepwise instructions and reasoning cues enabled the models to navigate the analytical process more transparently, mirroring the benefits of chain-of-thought prompting reported by Wei et al. (2022), who demonstrated that such prompts significantly enhance the reasoning capabilities of LLMs across a range of tasks [21]. Moreover, the hybrid prompting strategy, which combines elements of explicit instruction, reasoning scaffolds, and format constraints, emerged as particularly effective. This approach ensured that the models adhered to methodological standards while also producing outputs that were interpretable and aligned with clinical relevance. The robustness of this strategy across both case studies suggests that hybrid prompting may serve as a best-practice framework for integrating LLMs into medical statistical workflows. This finding is consistent with recent work by Wang et al. (2024), who advocate for multi-layered prompting strategies to maximize the reliability and interpretability of AI-generated outputs in healthcare contexts [44].

It is also noteworthy that few-shot prompting was not included in the evaluation due to its reliance on curated examples, which introduces variability and complicates standardization. While few-shot prompting has been shown to improve model performance in other domains [14], its effectiveness in medical statistics remains contingent on the quality and representativeness of the examples provided. Future research may explore the integration of few-shot prompting into hybrid frameworks, particularly in scenarios where high-quality exemplars are available.

The consistency of results across different LLMs further reinforces the generalizability of these findings. Although some models, such as Claude 3.7 Sonnet, demonstrated slightly more comprehensive outputs under minimal prompting, all models benefited significantly from structured prompts. This suggests that the observed improvements are not model-specific but rather a function of prompt design, a conclusion that echoes the broader literature on prompt engineering as a critical interface between human intent and machine intelligence [13].

The evaluation results provide compelling evidence that prompt engineering is not merely a technical convenience but a methodological necessity in the application of generative AI to medical statistics. The ability of LLMs to produce valid, complete, and contextually appropriate analyses is highly dependent on the clarity, specificity, and structure of the prompts they receive. As such, researchers and practitioners must approach prompt design with the same rigor as they would any other component of the analytical process, particularly in high-stakes domains where the consequences of statistical error are profound.

In this study, the outputs generated by GPT-4.1 and Claude 3.7 Sonnet were evaluated using Microsoft Copilot as the LLM-as-a-judge. These evaluations were subsequently reviewed by the author. Notably, there was complete agreement between the human assessments and those generated by Copilot across all evaluation criteria. This alignment reinforces the reliability of the LLM-as-a-judge approach in this context and supports its use as a scalable yet accurate evaluation method for prompt engineering in statistical tasks. The observed concordance may be attributed to the inherently objective nature of statistical reasoning, where tasks such as assumption checking, test selection, and interpretation follow well-defined rules and diagnostic criteria. In such domains, where procedural clarity is high and ambiguity is minimal, LLMs are more likely to replicate expert-level judgment. This suggests that LLM-as-a-judge frameworks may be particularly effective in evaluating outputs in structured analytical domains like medical statistics, where correctness can be assessed against established methodological standards.

5. Conclusions

This paper has situated the integration of generative AI within the broader context of medical statistics, emphasizing the dual potential of large language models (LLMs) to either enhance or undermine statistical rigor depending on the quality of prompt design. As the biomedical research community increasingly turns to LLMs for assistance with data analysis, the findings presented here underscore the critical importance of prompt engineering as a methodological safeguard against statistical misinterpretation.

Among the prompting strategies evaluated, hybrid prompting; defined as the combination of explicit instruction, chain of thought reasoning, and format constraints proved to be the most effective approach for inferential statistical tasks. This strategy consistently guided LLMs to perform assumption checks, select appropriate tests, and produce outputs that were both methodologically sound and interpretively coherent. In contrast, zero shot prompting, while adequate for basic descriptive statistics, was insufficient for tasks requiring analytical nuance and diagnostic rigor.

The evaluation also revealed that while different LLMs such as GPT 4.1, and Claude 3.7 Sonnet varied slightly in their baseline performance, all models demonstrated substantial improvements when guided by structured prompts. This suggests that the quality of prompt design, rather than model architecture alone, is the primary determinant of output reliability in statistical contexts.

A key limitation of this study is the restricted access to a broader range of LLMs. Due to subscription requirements and institutional constraints, the evaluation was limited to publicly accessible models. As a result, the generalizability of these findings to other proprietary or domain-specific LLMs remains uncertain.

Future research should extend this work by systematically evaluating a wider array of models across diverse statistical domains, including survival analysis, longitudinal data modeling, and meta-analysis. Additionally, the development of standardized prompt templates and domain-specific evaluation rubrics will be essential for fostering reproducibility and best practices in AI-assisted statistical workflows. Integrating LLMs with automated assumption checking tools and embedding human oversight into the analytical pipeline may further enhance the reliability and ethical deployment of these technologies in medical research.

Author Contributions: SV: Software, Writing – original draft, Writing – review & editing.

Funding: This research received no external funding

Institutional Review Board Statement: This study did not use human or animal subjects.

Informed Consent Statement: There was no human data used in this study

Data Availability Statement: Artificial data was used and can be found on the results section of this paper.

Conflicts of Interest: The authors declare no conflicts of interest.

Appendix A. LLM Evaluation Scoring Rubric

Criterion	Score	Description
Assumption Checking	1.0	Correctly identifies and tests all relevant assumptions (e.g., normality, variance).
Assumption Checking	0.8	Tests most assumptions but misses one minor check.
Assumption Checking	0.6	Tests some assumptions but overlooks key ones.
Assumption Checking	0.4	Mentions assumptions but does not test them.
Assumption Checking	0.2	Incorrectly assumes assumptions are met or ignores them entirely.
Assumption Checking	0.0	No mention or testing of assumptions.
Test Selection Accuracy	1.0	Selects the correct statistical test based on data type and assumptions.
Test Selection Accuracy	0.8	Selects a mostly appropriate test with minor justification issues.
Test Selection Accuracy	0.6	Selects a plausible but suboptimal test.
Test Selection Accuracy	0.4	Selects an incorrect test but explains reasoning.
Test Selection Accuracy	0.2	Selects an incorrect test with no justification.
Test Selection Accuracy	0.0	No test selected or completely irrelevant test.
Output Completeness	1.0	Provides all key outputs (e.g., test statistic, p-value, CI, group means).
Output Completeness	0.8	Provides most outputs but omits one or two minor elements.
Output Completeness	0.6	Provides partial output with some missing key elements.
Output Completeness	0.4	Output is vague or lacks structure.
Output Completeness	0.2	Output is incomplete and unclear.
Output Completeness	0.0	No usable output provided.
Interpretive Quality	1.0	Interpretation is accurate, context-aware, and clinically relevant.
Interpretive Quality	0.8	Interpretation is mostly accurate with minor contextual gaps.
Interpretive Quality	0.6	Interpretation is generic or lacks depth.
Interpretive Quality	0.4	Interpretation is partially incorrect or misleading.
Interpretive Quality	0.2	Interpretation is mostly incorrect or irrelevant.
Interpretive Quality	0.0	No interpretation provided.

References

1. He, R.; Cao, J.; Tan, T. Generative artificial intelligence: a historical perspective. *National Science Review* **2025**, *12*, nwaf050.

2. Akhtar, Z.B. Unveiling the evolution of generative AI (GAI): a comprehensive and investigative analysis toward LLM models (2021–2024) and beyond. *Journal of Electrical Systems and Information Technology* **2024**, *11*, 22.

3. Schmidhuber, J. Deep learning in neural networks: An overview. *Neural networks* **2015**, *61*, 85–117.

4. Norvig, P.; Stuart, R. *Artificial Intelligence: A Modern Approach*; Pearson Education, USA, 2021.

5. Munro, J. A Review of: “Rule-based expert systems: the MYCIN experiments of the Stanford heuristic programming project”. *Civil Engineering Systems* **1984**, *1*, 342–343.

6. Singhal, K.; Azizi, S.; Tu, T.; Mahdavi, S.S.; Wei, J.; Chung, H.W.; Scales, N.; Tanwani, A.; Cole-Lewis, H.; Pfohl, S.; et al. Large language models encode clinical knowledge. *Nature* **2023**, *620*, 172–180.

7. Jiang, F.; Jiang, Y.; Zhi, H.; Dong, Y.; Li, H.; Ma, S.; Wang, Y.; Dong, Q.; Shen, H.; Wang, Y. Artificial intelligence in healthcare: past, present and future. *Stroke and vascular neurology* **2017**, *2*.

8. Altman, D.G. The scandal of poor medical research, 1994.

9. Ioannidis, J.P. Why most published research findings are false. *PLoS medicine* **2005**, *2*, e124.

10. Nuijten, M.B.; Hartgerink, C.H.; Van Assen, M.A.; Epskamp, S.; Wicherts, J.M. The prevalence of statistical reporting errors in psychology (1985–2013). *Behavior research methods* **2016**, *48*, 1205–1226.

11. Chavalarias, D.; Wallach, J.D.; Li, A.H.T.; Ioannidis, J.P. Evolution of reporting P values in the biomedical literature, 1990-2015. *Jama* **2016**, *315*, 1141–1148.
12. Straszak, A.M.; Zaman, Q.; Pfeiffer, K.P.; Göbel, G.; Ulmer, H. Statistical errors in medical research-a review of common pitfalls. *Swiss medical weekly* **2007**, *137*, 44–44.
13. Liu, P.; Yuan, W.; Fu, J.; Jiang, Z.; Hayashi, H.; Neubig, G. Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing. *ACM computing surveys* **2023**, *55*, 1–35.
14. Brown, T.; Mann, B.; Ryder, N.; Subbiah, M.; Kaplan, J.D.; Dhariwal, P.; Neelakantan, A.; Shyam, P.; Sastry, G.; Askell, A.; et al. Language models are few-shot learners. *Advances in neural information processing systems* **2020**, *33*, 1877–1901.
15. Reynolds, L.; McDonnell, K. Prompt programming for large language models: Beyond the few-shot paradigm. In Proceedings of the Extended abstracts of the 2021 CHI conference on human factors in computing systems, 2021, pp. 1–7.
16. Zhou, K.; Yang, J.; Loy, C.C.; Liu, Z. Learning to prompt for vision-language models. *International Journal of Computer Vision* **2022**, *130*, 2337–2348.
17. Kojima, T.; Gu, S.S.; Reid, M.; Matsuo, Y.; Iwasawa, Y. Large language models are zero-shot reasoners. *Advances in neural information processing systems* **2022**, *35*, 22199–22213.
18. Wei, J.; Tay, Y.; Bommasani, R.; Raffel, C.; Zoph, B.; Borgeaud, S.; Yogatama, D.; Bosma, M.; Zhou, D.; Metzler, D.; et al. Emergent abilities of large language models. *arXiv preprint arXiv:2206.07682* **2022**.
19. Ouyang, L.; Wu, J.; Jiang, X.; Almeida, D.; Wainwright, C.; Mishkin, P.; Zhang, C.; Agarwal, S.; Slama, K.; Ray, A.; et al. Training language models to follow instructions with human feedback. *Advances in neural information processing systems* **2022**, *35*, 27730–27744.
20. Naveed, H.; Khan, A.U.; Qiu, S.; Saqib, M.; Anwar, S.; Usman, M.; Akhtar, N.; Barnes, N.; Mian, A. A comprehensive overview of large language models. *arXiv preprint arXiv:2307.06435* **2023**.
21. Wei, J.; Wang, X.; Schuurmans, D.; Bosma, M.; Xia, F.; Chi, E.; Le, Q.V.; Zhou, D.; et al. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems* **2022**, *35*, 24824–24837.
22. Giray, L. Prompt engineering with ChatGPT: a guide for academic writers. *Annals of biomedical engineering* **2023**, *51*, 2629–2633.
23. Wang, D.; Zhang, S. Large language models in medical and healthcare fields: applications, advances, and challenges. *Artificial Intelligence Review* **2024**, *57*, 299.
24. Zhang, C.; Stepputtis, S.; Sycara, K.; Xie, Y. Enhancing vision-language few-shot adaptation with negative learning. In Proceedings of the 2025 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV). IEEE, 2025, pp. 5905–5915.
25. Wang, X.; Zhou, D. Chain-of-thought reasoning without prompting. *arXiv preprint arXiv:2402.10200* **2024**.
26. Lu, Y.; Bartolo, M.; Moore, A.; Riedel, S.; Stenetorp, P. Fantastically ordered prompts and where to find them: Overcoming few-shot prompt order sensitivity. *arXiv preprint arXiv:2104.08786* **2021**.
27. Bommasani, R.; Hudson, D.A.; Adeli, E.; Altman, R.; Arora, S.; von Arx, S.; Bernstein, M.S.; Bohg, J.; Bosselut, A.; Brunskill, E.; et al. On the opportunities and risks of foundation models. *arXiv preprint arXiv:2108.07258* **2021**.
28. Ji, Z.; Lee, N.; Frieske, R.; Yu, T.; Su, D.; Xu, Y.; Ishii, E.; Bang, Y.J.; Madotto, A.; Fung, P. Survey of hallucination in natural language generation. *ACM computing surveys* **2023**, *55*, 1–38.
29. Zhou, D.; Schärli, N.; Hou, L.; Wei, J.; Scales, N.; Wang, X.; Schuurmans, D.; Cui, C.; Bousquet, O.; Le, Q.; et al. Least-to-most prompting enables complex reasoning in large language models. *arXiv preprint arXiv:2205.10625* **2022**.
30. Altman, D.G. *Practical statistics for medical research*; Chapman and Hall/CRC, 1990.
31. Hoaglin, D.C.; Mosteller, F.; Tukey, J.W. *Understanding robust and exploratory data analysis*; John Wiley & Sons, 2000.
32. Motulsky, H. *Intuitive biostatistics: a nonmathematical guide to statistical thinking*; Oxford University Press, USA, 2014.
33. Hosmer Jr, D.W.; Lemeshow, S.; Sturdivant, R.X. *Applied logistic regression*; John Wiley & Sons, 2013.
34. Wasserstein, R.L.; Lazar, N.A. The ASA statement on p-values: context, process, and purpose, 2016.
35. Head, M.L.; Holman, L.; Lanfear, R.; Kahn, A.T.; Jennions, M.D. The extent and consequences of p-hacking in science. *PLoS biology* **2015**, *13*, e1002106.

36. Faizullah, A.R.B.M.; Urlana, A.; Mishra, R. LimGen: Probing the LLMs for Generating Suggestive Limitations of Research Papers. In Proceedings of the Joint European Conference on Machine Learning and Knowledge Discovery in Databases. Springer, 2024, pp. 106–124.
37. Gao, M.; Hu, X.; Yin, X.; Ruan, J.; Pu, X.; Wan, X. Llm-based nlg evaluation: Current status and challenges. *Computational Linguistics* **2025**, pp. 1–28.
38. Lavie, A.; Denkowski, M.J. The METEOR metric for automatic evaluation of machine translation. *Machine translation* **2009**, *23*, 105–115.
39. Berti, A.; Kourani, H.; van der Aalst, W.M. PM-LLM-Benchmark: Evaluating large language models on process mining tasks. In Proceedings of the International Conference on Process Mining. Springer, 2024, pp. 610–623.
40. Shankar, S.; Zamfirescu-Pereira, J.; Hartmann, B.; Parameswaran, A.; Arawjo, I. Who validates the validators? aligning llm-assisted evaluation of llm outputs with human preferences. In Proceedings of the Proceedings of the 37th Annual ACM Symposium on User Interface Software and Technology, 2024, pp. 1–14.
41. Shahzad, T.; Mazhar, T.; Tariq, M.U.; Ahmad, W.; Ouahada, K.; Hamam, H. A comprehensive review of large language models: issues and solutions in learning environments. *Discover Sustainability* **2025**, *6*, 27.
42. Kamath, U.; Keenan, K.; Somers, G.; Sorenson, S. *Large Language Models: A Deep Dive*; Springer, 2024.
43. Kumar, P. Large language models (LLMs): survey, technical frameworks, and future challenges. *Artificial Intelligence Review* **2024**, *57*, 260.
44. Wang, X.; Wei, J.; Schuurmans, D.; Le, Q.; Chi, E.; Narang, S.; Chowdhery, A.; Zhou, D. Self-consistency improves chain of thought reasoning in language models. *arXiv preprint arXiv:2203.11171* **2022**.

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.