

ICE-Score: Instructing Large Language Models to Evaluate Code

Terry Yue Zhuo

Monash University and CSIRO’s Data61

terry.zhuo@monash.edu

Abstract

Recent advancements in the field of natural language generation have facilitated the use of large language models to assess the quality of generated text. Although these models have shown promising results in tasks such as machine translation and summarization, their applicability in code intelligence tasks remains limited without human involvement. The complexity of programming concepts required for such tasks makes it difficult to develop evaluation metrics that align with human judgment. Token-matching-based metrics, such as BLEU, have demonstrated weak correlations with human practitioners in code intelligence tasks. Moreover, utilizing human-written test suites to evaluate functional correctness can be challenging in domains with low resources. To overcome these obstacles, we propose ICE-Score, a new evaluation metric via instructing large language models (LLMs) for code assessments. Our metric addresses the limitations of existing approaches by achieving superior correlations with functional correctness and human preferences, without the need for test oracles or references. We evaluate the efficacy of our metric on two different aspects (*human preference* and *execution success*) and four programming languages. Our results demonstrate that our metric surpasses state-of-the-art metrics for code generation, delivering high levels of accuracy and consistency across various programming languages and tasks. We also make our evaluation metric and datasets available to the public¹, encouraging further research in evaluating code intelligence tasks.

1 Introduction

Natural language generation (NLG) systems have seen significant progress in developing large language models (LLMs). These models have shown great promise in generating high-quality and diverse texts that can be difficult to distinguish from

human-written texts (Ouyang et al., 2022). However, evaluating the quality of NLG systems remains a challenging task, primarily due to the limitations of traditional evaluation metrics. Token-matching-based metrics, such as BLEU (Papineni et al., 2002) and ROUGE (Lin, 2004), have been widely used to evaluate NLG systems but have demonstrated poor correlation with human judgment and a lack of ability to capture semantic meanings (Kocmi et al., 2021). Furthermore, these metrics require reference output, which can be challenging to obtain for new tasks and low-resource domains (Liu et al., 2023).

In recent years, the use of LLMs as reference-free evaluators for Natural Language Generation (NLG) tasks has gained attention among researchers. This approach is strongly aligned with human preferences, even when reference texts are unavailable (Liu et al., 2023; Fu et al., 2023). The underlying assumption behind this approach is that LLMs possess a profound understanding of human-generated text and task instructions, enabling them to evaluate various NLG tasks through prompts. The exceptional performance of LLMs in contextual understanding and natural language generation, as evidenced by studies (Brown et al., 2020), further supports this assumption. Moreover, LLMs trained on both textual and code-based data have showcased remarkable capabilities in diverse downstream tasks related to source code, including code generation (OpenAI, 2023; Allal et al., 2023; Li et al., 2023). While a performance gap still exists between LLMs and human developers in code-related tasks, recent research has illustrated that LLMs can be enhanced to handle various source code tasks with appropriate guidance (Chen et al., 2023; Madaan et al., 2023). This indicates the significant potential of LLMs in comprehending and working with source code.

Code evaluation presents unique challenges, requiring a deeper understanding of programming

¹<https://github.com/terryyz/ice-score>

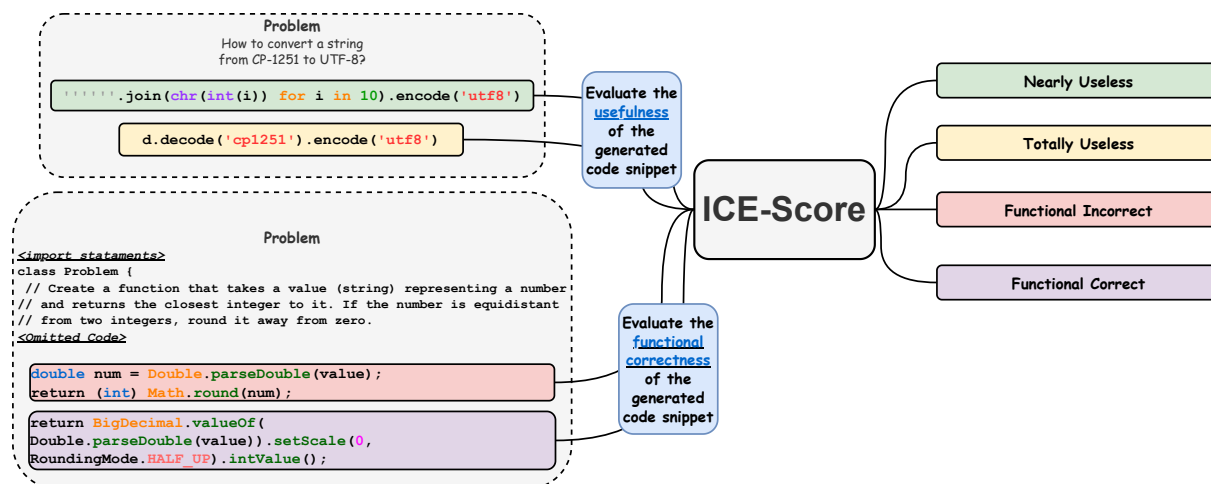


Figure 1: An illustration of ICE-Score. On the left-hand side, we input the task problems and corresponding generated code snippets. On the right-hand side, ICE-Score outputs the corresponding assessments.

concepts and more complex syntax than natural language generation (Hindle et al., 2016). Traditional reference-based evaluation metrics for code generation, such as BLEU (Papineni et al., 2002), ROUGE (Lin, 2004), and chrF (Popović, 2015), rely on token matching to assess performance automatically. However, these metrics have demonstrated poor correlation with human evaluation (Evtikhiev et al., 2023) since they often underestimate the variety of outputs with the same semantic logic. While some studies have incorporated programming features to improve these metrics, they have shown limited gains and poor correlation with functional correctness (Eghbali and Pradel, 2022; Tran et al., 2019). Alternatively, researchers have proposed using well-designed test suites to objectively evaluate code generation performance at the function level (Chen et al., 2021; Zheng et al., 2023; Cassano et al., 2023). However, developing these test suites requires programming expertise, which can be impractical and costly in low-resource scenarios. Additionally, executing model-generated code poses a security risk and must be run in an isolated sandbox, which is technically cumbersome.

More recently, CodeBERTScore (Zhou et al., 2023), a neural-model-based evaluation metric, has been proposed, showing a higher correlation with functional correctness and human preferences by capturing the semantic information of reference code and generated code. However, CodeBERTScore still relies on high-quality references that can be difficult and expensive to obtain. Moreover, the limited performance of the Code-

BERT (Feng et al., 2020) backbone suggests that it has not yet reached a human-level understanding of source code, limiting the effectiveness of CodeBERTScore. Therefore, more advanced evaluation metrics are needed so that they can better capture the complex syntax and semantics of code intelligence tasks.

To address these challenges, we propose a novel evaluation metric based on LLMs trained on both text and code, shown in Figure 1. Specifically, we Instruct LLMs to perform human-like multi-dimensional Code Evaluation, where the metric is denoted as ICE-Score. Our metric leverages the recent NLG metric, G-EVAL (Liu et al., 2023), but achieves superior correlations with subjective human preferences and objective functional correctness, both at the example and corpus levels. Different from G-EVAL, ICE-Score only relies on assessment criteria and evaluation step template, without the need for instruction generation and weighted scoring function.

Based on our extensive evaluation, we have summarized our contributions as follows:

- We designed the first multi-dimensional and reference-free automatic evaluation metric for code intelligence tasks via large language models.
- We conducted extensive experiments to demonstrate the efficacy of ICE-Score on four programming languages (Java, Python, C, C++, and JavaScript) from two aspects (*human-based usefulness* and *execution-based functional correctness*).

- We further discussed several aspects that can improve the performance of ICE-Score, including the backbone model performance and reasoning capability.

2 Method

Our evaluation metric ICE-Score, inspired by G-EVAL (Liu et al., 2023), consists of two main components: 1) task definition, evaluation criteria, and detailed evaluation steps, and 2) a given problem and generated code snippet for evaluation. Different from G-EVAL, we only require the input of evaluation criteria and template-based evaluation steps, without the need for generation from LLMs. In addition, As we set the model temperature to 0, our evaluation metric no longer needs a weighted scoring function after iterative score generation. These two differences suggest that ICE-Score is more cost-friendly and efficient.

2.1 Instructions for Code Evaluation

The evaluation of code quality involves two main aspects: 1) human judgment of code usefulness and 2) execution-based functional correctness. To provide a comprehensive evaluation, we adopt the design of G-EVAL for the general task instruction, as follows:

You will be given the code snippet for a problem. Your task is to rate the code snippet only on one metric. Please make sure you read and understand these instructions carefully. Please keep this document open while reviewing, and refer to it as needed.

Regarding the task-agnostic prompt, we have designed the evaluation criteria for assessing **code usefulness**, as shown in Appendix A.1. These criteria are aligned with previous human evaluations of code quality (Evtikhiev et al., 2023). To evaluate **functional correctness**, we emphasize the importance of considering unit tests during the evaluation process. We present the following criteria for evaluating functional correctness, as provided in Appendix A.2.

For the instruction of evaluation steps, we provide a template-based prompt:

Evaluation Steps:

1. *Read the problem carefully and identify the required functionalities of the implementation.*

2. *Read the code snippet and compare it to the problem. Check if the code snippet covers all required functionalities of the problem, and if it aligns with the Evaluation Criteria.*

3. *Assign a score for [Evaluation Aspect] on a scale of 0 to 4, where 0 is the lowest and 4 is the highest based on the Evaluation Criteria.*

Here, we define [Evaluation Aspect] as any aspects that are emphasized during the evaluation. In our paper, we consider **code usefulness** and **functional correctness**.

2.2 Inputs of Code Evaluation

It is worth noting that most code generative models do not take formatting into account, resulting in unformatted code that requires post-processing of code formatting to be understood, compiled, and executed (Zheng et al., 2023). Additionally, automatic evaluation metrics for code generation, such as CodeBLEU (Ren et al., 2020) and RUBY (Tran et al., 2019), still rely on language-specific program parsers². However, based on prior findings that LLMs can robustly understand input data (Huang et al., 2022; Zhuo et al., 2023; Zhu et al., 2023), we hypothesize that LLMs can also understand programming context without proper formatting. Therefore, for evaluation, we input the problems and generated code (and reference code, if provided). When the reference code is provided, we slightly modify the evaluation steps in the prompt to incorporate it.

3 Experiment Setup

We evaluate the effectiveness of ICE-Score using GPT-3.5 (GPT-3.5-turbo³) as the backbone across multiple datasets and programming languages. We conduct two experiments to investigate the correlation between ICE-Score and human preference and functional correctness, respectively. We compare the performance of LLM-based evaluations against 7 predominant automatic evaluation metrics, including the state-of-the-art CodeBERTScore (Zhou et al., 2023). To measure the correlation with human preference, we use the

²<https://tree-sitter.github.io/>

³<https://platform.openai.com/docs/models/gpt-3-5>

CoNaLa dataset (Yin et al., 2018) and corresponding human annotation on the generated code from various models trained on the dataset (Evtikhiev et al., 2023). To measure the correlation with functional correctness, we use the HumanEval-X dataset (Zheng et al., 2023). We do not consider **distinguishability** as an evaluation option, as prior work (Zhou et al., 2023) has shown it to be an unreliable meta-metric that cannot substitute for execution-based or human-based ratings.

3.1 Automatic Evaluation Metric Baselines

The baseline metrics we include can be classified into two groups: **string-based** and **neural-model-based** evaluation.

String-based Evaluation Most evaluation metrics in code generation have been adapted from natural language generation (NLG) and rely on comparing the generated code to reference code. The most commonly used metric is BLEU (Papineni et al., 2002), which computes the overlaps of n -grams in the generated output with those in the reference, where the n -grams are tokenized using a language-specific tokenizer (Post, 2018). Other metrics include ROUGE-L (Lin, 2004), a recall-oriented metric that looks for the longest common subsequence between the reference and the generated code, and METEOR (Banerjee and Lavie, 2005), which is based on unigram matching between the generated code and the reference. However, studies have shown that BLEU may yield similar results for models with different quality levels from the perspective of human graders in code generation (Evtikhiev et al., 2023), leading to the proposal of new evaluation metrics such as RUBY (Tran et al., 2019). RUBY takes the code structure into account and compares the program dependency graphs (PDG) of the reference and the candidate. If the PDG is impossible to build, the metric falls back to comparing the abstract syntax tree (AST), and if the AST is also impossible to build, it compares the weighted string edit distance between the tokenized reference and candidate sequence. Another recent metric is CodeBLEU (Ren et al., 2020), which is a composite metric that computes a weighted average of four sub-metrics treating code differently: as a data-flow graph, as an abstract syntax tree, and as text. CodeBLEU is designed to evaluate the quality of generated code for code generation, code translation, and code refinement tasks.

Metric	Example			Corpus		
	τ	r_p	r_s	τ	r_p	r_s
BLEU	.439	.522	.488	.423	.572	.542
CodeBLEU	.292	.363	.331	.259	.397	.339
chrF	.458	.570	.515	.449	<u>.592</u>	.578
ROUGE-L	.447	.529	.499	.432	.581	.552
METEOR	.410	.507	.462	.415	.557	.534
RUBY	.331	.397	.371	.339	.493	.439
CodeBERTScore-F1	.500	<u>.609</u>	.556	<u>.464</u>	.579	<u>.595</u>
CodeBERTScore-F3	<u>.505</u>	<u>.609</u>	<u>.563</u>	.437	.549	.564
ICE-Score	.556	.613	.594	.546	.649	.635
Ref-ICE-Score	.554	.617	.591	.539	.661	.630

Table 1: Example-level and corpus-level Kendall-Tau (τ), Pearson (r_p) and Spearman (r_s) correlations with the human preferred usefulness on CoNaLa. ICE-Score: without reference code inputs, or reference-free; Ref-ICE-Score: reference-enhanced. The best performance is **bold**. The second-best performance is underlined.

Neural-model-based Evaluation Neural-model-based evaluation is becoming increasingly important for evaluating the quality of code generated by deep learning models. CodeBERTScore (Zhou et al., 2023) is one of the latest approaches that leverages pre-trained code models like CodeBERT (Feng et al., 2020) and best practices from natural language generation evaluation to assess the quality of generated code. CodeBERTScore encodes the generated code and reference code independently and considers the natural language context, contextual information of each token, and implementation diversity. It enables the comparison of code pairs that are lexically different and calculates precision and recall based on the best-matching token vector pairs. This approach provides an effective way to evaluate the effectiveness of deep learning models for code intelligence tasks. Note that the authors of CodeBERTScore provided both F1 and F3 scores, with the optional source input. Therefore, we use these four language-specific variants of CodeBERTScore in our experiments.

3.2 Datasets and Evaluation Aspects

Human-based Usefulness Experiments Similar to (Zhou et al., 2023), we conduct an evaluation on the CoNaLa benchmark (Yin et al., 2018), which is a widely used dataset for natural language context to Python code generation. To measure the correlation between each evaluation metric and human preference, we utilize the human annotations provided by (Evtikhiev et al., 2023). Specifically, for each example in the dataset, experienced software

Metric	Java		C++		Python		JavaScript		Average	
	τ	r_s	τ	r_s	τ	r_s	τ	r_s	τ	r_s
BLEU	.337	.401	.146	.174	.251	.297	.168	.199	.225	.268
CodeBLEU	.355	.421	.157	.187	.272	.323	.226	.267	<u>.253</u>	<u>.299</u>
chrF	.346	.413	.166	.198	.262	.312	.186	.220	.240	.286
ROUGE-L	.327	.389	.143	.171	.240	.284	.151	.179	.215	.256
METEOR	.358	.425	<u>.174</u>	<u>.208</u>	<u>.276</u>	<u>.327</u>	<u>.195</u>	<u>.231</u>	.251	.298
RUBY	.340	.401	.139	.165	.216	.255	.138	.163	.208	.246
CodeBERTScore-F1	.314	.375	.148	.177	.231	.276	.145	.172	.209	.250
CodeBERTScore-F3	.356	.426	.166	.198	.262	.312	.189	.226	.243	.291
ICE-Score	.427	.442	.320	.326	.279	.282	.316	.321	.336	.343
Ref-ICE-Score	.388	.404	.274	.282	.318	.325	.340	.348	.330	.340

Table 2: Example-level Kendall-Tau (τ) and Spearman (r_s) correlations with the execution-based functional correctness on HumanEval. ICE-Score: without reference code inputs, or reference-free; Ref-ICE-Score: with reference code inputs, or reference-enhanced. The best performance is **bold**. The second-best performance is underlined.

developers were asked to grade the generated code snippets from five different models. The grading scale ranges from zero to four, with zero indicating that the generated code is irrelevant and unhelpful, and four indicating that the generated code solves the problem accurately. The dataset comprises a total of 2,860 annotated code snippets (5 generations \times 472 examples) with each snippet being graded by 4.5 annotators on average.

Execution-based Functional Correctness Experiments We conduct an evaluation of functional correctness using the HumanEval benchmark (Chen et al., 2021), which provides natural language goals, input-output test cases, and reference solutions written by humans for each example. The benchmark originally consists of 164 coding problems in Python, and has been extended by (Cassano et al., 2023) to 18 other programming languages, including Java, C++, Python, and JavaScript. We chose to evaluate our models on these languages, as they are among the most popular programming languages. The translated examples also include the predictions of code-davinci-002 and their corresponding functional correctness scores. Inspired by (Zhou et al., 2023), we obtain them from the HumanEval-X dataset (Zheng et al., 2023). As each problem has nearly 200 generated code samples on average, it would be computationally expensive to evaluate them all using LLMs. Therefore, we randomly select 20 samples from each problem, and collect all samples from problems where no more than 20 versions of code were generated.

Correlation Metrics To measure the correlation between each metric’s scores and the references, we follow best practices in natural language evaluation and used Kendall-Tau (τ), Pearson (r_p), and Spearman (r_s) coefficients.⁴ To systematically study the efficacy of each automatic evaluation metric, we compute both example-level and corpus-level correlations. The example-level correlation is the average correlation of each problem example, while the corpus-level correlation is the correlation of all aggregated examples in the task.

4 Results

Human-based Usefulness Table 1 shows the correlation between different metrics with human preference. We compare two variants of our evaluation approach, reference-free and reference-enhanced evaluations, with 10 baseline metrics and their variants. We find that ICE-Score outperform these metrics by a significant margin, regarding both example- and corpus-level correlations. Our observation is consistent with the work of CodeBERTScore, where the variants of CodeBERTScore mostly outperform the strong baselines like chrF and ROUGE-L. For example, ICE-Score achieves 0.556 and 0.546 measured by Spearman correlation on example level and corpus level, respectively. In contrast, prior evaluation metrics barely reach a score of 0.5. In addition, we find that Ref-ICE-Score does not significantly improve the performance, indicating the reference code may not be optimized. Our further analysis of the human rating of CoNaLa reference code complies

⁴We use the implementations from <https://scipy.org/>

Metric	Java		C++		Python		JavaScript		Average	
	τ	r_s	τ	r_s	τ	r_s	τ	r_s	τ	r_s
BLEU	.267	.326	.225	.276	.281	.344	.220	.270	.248	.304
CodeBLEU	.293	.359	.212	.260	.303	.371	<u>.315</u>	<u>.385</u>	.281	.343
chrF	.290	.355	.266	.325	.328	.402	.279	.342	.291	.356
ROUGE-L	.280	.342	.234	.286	.296	.363	.216	.264	.256	.314
METEOR	.318	.389	.260	.319	.349	.427	.311	.380	.309	.379
RUBY	.276	.337	.219	.268	.279	.341	.219	.268	.248	.303
CodeBERTScore-F1	.244	.298	.219	.268	.264	.324	.214	.262	.235	.288
CodeBERTScore-F3	.281	.344	.243	.297	.313	.384	.261	.320	.275	.336
ICE-Score	.330	.345	.313	.321	.294	.298	.315	.323	.313	.322
Ref-ICE-Score	.412	.438	.367	.383	.425	.446	.432	.455	.409	.431

Table 3: Corpus-level Kendall-Tau (τ) and Spearman (r_s) correlations with the execution-based functional correctness on HumanEval. ICE-Score: without reference code inputs, or reference-free; Ref-ICE-Score: with reference code inputs, or reference-enhanced. The best performance is **bold**. The second-best performance is underlined.

with this implication, where the average score of the reference code only achieves 3.4 out of 4, suggesting that not all human practitioners consider the reference fully useful.

Execution-based Functional Correctness Table 2 and Table 3 present the results of example- and corpus-level functional correctness, respectively. From Table 2, we observe that both reference-free and reference-enhanced Ref-ICE-Scores consistently outperform the other baselines across all four programming languages on the example level. ICE-Score even outperforms the reference-enhanced one, suggesting potential bias in some reference code. Additionally, we find that METEOR and CodeBLEU receive better correlations than all variants of CodeBERTScore, indicating that they are still strong baselines compared to the recent neural-model-based evaluators in code generation. In Table 3, we observe that our Ref-ICE-Score achieves state-of-the-art performance among all evaluation metrics. When compared to other baselines, ICE-Score still achieves comparable results to the source-free CodeBERTScore-F3.

5 Ablation Study

Does reasoning help the code evaluation? Prior work (Wei et al.; Kojima et al.) has demonstrated that the performance of LLMs can be significantly improved via Chain-of-Thought (CoT) and Zero-Shot-Chain-of-Thought (ZS-CoT), where the prompts instruct LLMs to perform the task in a step-by-step manner. Here, we explore the zero-shot reasoning ability of LLMs in evaluating code generation. Specifically, we instruct GPT-3.5 to

Metric	Example			Corpus		
	τ	r_p	r_s	τ	r_p	r_s
ICE-Score	.556	.613	.594	.546	.649	.635
CoT-ICE-Score	.561	.628	.600	.579	.703	.665
Ref-ICE-Score	.554	.617	.591	.539	.661	.630
CoT-Ref-ICE-Score	.571	.639	.607	.583	.712	.667

Table 4: Example-level and corpus-level Kendall-Tau (τ), Pearson (r_p) and Spearman (r_s) correlations with the human preferred usefulness on CoNaLa. ICE-Score: without reference code inputs, or reference-free; Ref-ICE-Score: with reference code inputs, or reference-enhanced. CoT- indicates the use of ZS-CoT. The best performance is **bold**.

perform CoT-evaluation by adding "Step-by-step Evaluation:" at the end of the prompt. An example of the zero-shot-CoT prompt is shown in Figure 2. Instead of using LLMs to extract the evaluation score from the reasoning steps, like the original metric of zero-shot-CoT via multiple queries, we design a rule-based parser to extract scores. Due to limited resources, we only evaluate on CoNaLa in Table 4. Our results show that ZS-CoT can significantly improve the reliability of code evaluation. Additionally, we find that Ref-ICE-Score can achieve better results than reference-free ones via ZS-CoT, even though their performances are similar without CoT processing. This suggests that LLMs can exploit the use of reference code through reasoning.

Does more-capable backbone LLM yield better performance on code evaluation? As shown in previous studies (OpenAI, 2023; Bubeck et al., 2023), GPT-4 significantly outperforms GPT-3.5 on various tasks. Therefore, we use GPT-4 as the backbone model for ICE-Score and evalu-

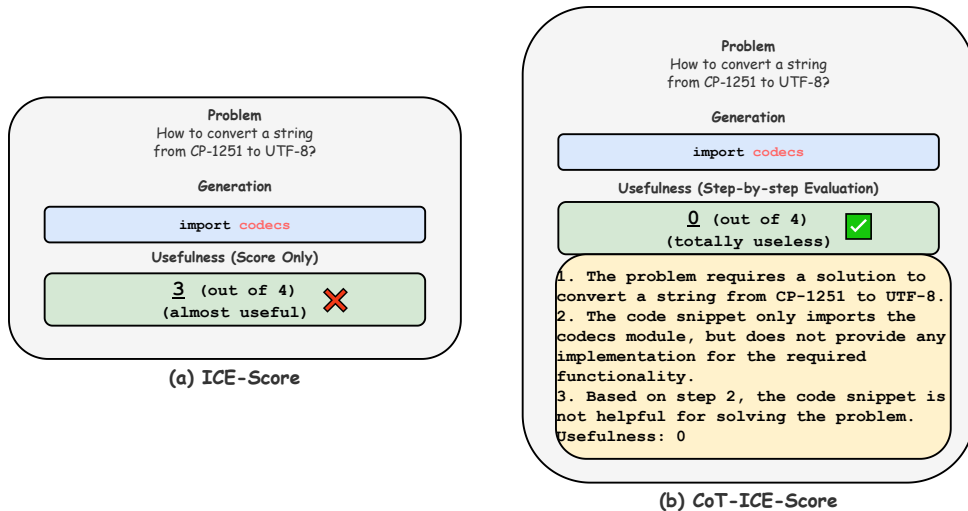


Figure 2: Example inputs and outputs with (a) ICE-Score, (b) ICE-Score with Zero-Shot Chain-of-Thought. With the step-by-step evaluation, the output assessment is more aligned with human preference.

Metric	Example			Corpus		
	τ	r_p	r_s	τ	r_p	r_s
ICE-Score-3.5	.556	.613	.594	.546	.649	.635
ICE-Score-4	.612	.658	.611	.592	.720	.688
Ref-ICE-Score-3.5	.554	.617	.591	.539	.661	.630
Ref-ICE-Score-4	.592	.647	.634	.632	.744	.690

Table 5: Example-level and corpus-level Kendall-Tau (τ), Pearson (r_p) and Spearman (r_s) correlations with the human preferred usefulness on CoNaLa. ICE-Score: without reference code inputs, or reference-free; Ref-ICE-Score: with reference code inputs, or reference-enhanced. -3.5 and -4 suggest the different backbone models. The best performance is **bold**.

ate its performance on CoNaLa. The results in Table 5 indicate that GPT-4 consistently surpasses GPT-3.5-turbo on evaluating code, suggesting it has the superior capability of code comprehension. We also note that using a more capable model like GPT-4 can guarantee even better performance, compared to using ZS-CoT techniques in Table 4.

6 Discussion

Data Contamination Evaluations on recent closed-source LLMs have been criticized for the possibility of data contamination (Aiyappa et al., 2023), where the model may have already seen the evaluation datasets during training, due to the opaque training details of these models. For instance, Kocmi and Federmann (2023) conducted an empirical study on a few closed-source LLMs, including GPT-3.5, and suggested that LLMs are the state-of-the-art evaluators of translation qual-

ity, based on the evaluation of the WMT22 Metric Shared Task (Freitag et al., 2022). However, as most of the evaluated models were trained on data prior to 2022⁵, it is highly likely that these models have been trained with some human-rated translation quality data. Similarly, G-EVAL (Liu et al., 2023) shows that GPT-3.5 and GPT-4 are the state-of-the-art evaluators of natural language generation (NLG) with the evaluation of three NLG datasets. However, as these human-annotated datasets were released before 2021, it is probable that they were included in the training data of GPT-3.5 and GPT-4. In contrast, our work is minimally impacted by data contamination, as we report the data release year in Table 6. Our analysis suggests that only CoNaLa and HumanEval (Python) datasets may have been contaminated, and it is unlikely that GPT-3.5 has seen any human annotation or generated code during training.

Human-aligned Evaluation Beyond Code Generation While our study has shown that LLMs can achieve state-of-the-art performance in evaluating the functional correctness and usefulness of generated source code, the question remains as to whether LLMs can be utilized to evaluate code intelligence tasks beyond code generation. Allamanis et al. (2018) have identified several downstream applications such as code translation, commit message generation, and code summarization. While some studies have investigated the human evaluation of these tasks, none of them have released

⁵<https://platform.openai.com/docs/model-index-for-researchers>

Dataset	Release Year	Likely to be contaminated?
CoNaLa	2018	✓
human-annotated CoNaLa w/ generated code	2023	✗
HumanEval (Python)	2021	✓
HumanEval-X (w/o Python)	2023	✗
human-annotated HumanEval-X w/ generated code	2023	✗

Table 6: Dataset, Release Year and the likelihood of data contamination for each dataset used in our study.

the annotation data or fully described the human evaluation criteria. This presents a challenge for analyzing if `ICE-SCORE` can be adapted to these tasks. For example, [Hu et al. \(2022\)](#) proposed a human evaluation metric for code documentation generation quality, which is specifically designed for code comment generation and commit message generation. Their metric includes three aspects: *Language-related*, *Content-related*, and *Effectiveness-related*, with detailed task descriptions and explanations of assigned scores. We propose that the information provided in their metric can be used to create prompts for LLM-based evaluation and enable human-aligned evaluation of code documentation generation.

7 Related Work

Large Language Models for Code. LLMs pre-trained on large-scale code data have demonstrated strong capabilities in code intelligence tasks, such as code completion ([Li et al., 2023](#); [Luo et al., 2023](#); [Rozière et al., 2023](#)), code summarization ([Ahmed and Devanbu, 2022](#); [Sun et al., 2023](#)) and program repair ([Surameery and Shakor, 2023](#); [Sobania et al., 2023](#)). However, they remain unreliable, particularly in scenarios that require an understanding of natural language. Recent studies ([Muenighoff et al., 2023b](#); [Ma et al.](#)) show that pre-training on both text and code results in the optimal model performance on natural language and code understanding. Furthermore, in order to make LLMs more human-aligned and more capable of performing complex tasks, instruction tuning is proposed to enhance the capability of following natural language requirements. In this work, we utilize such instruction-tuned LLMs to conduct multi-dimensional code evaluation via various instructions.

Automatic Evaluation Metrics for Generation. The quest for reliable and robust automatic evaluation metrics for generated content has been a cornerstone in natural language processing. Tradi-

tionally, string-based metrics such as BLEU ([Papineni et al., 2002](#)), ROUGE ([Lin, 2004](#)), and METEOR ([Banerjee and Lavie, 2005](#)) have dominated the landscape, primarily when assessing machine translation or text summarization outputs. While these metrics provide a quick and cost-effective means of evaluating the quality of the generated text, they often fall short of capturing the nuanced intricacies and semantic richness inherent in natural language. To mitigate such drawbacks, a few neural-based multi-dimensional evaluation metrics have been proposed for text generation, such as UniEval ([Zhong et al., 2022](#)), GPTScore ([Fu et al., 2023](#)) and G-EVAL ([Liu et al., 2023](#)). However, when it comes to code generation, where both syntactical correctness and semantic intent are paramount, there are few attempts to address these challenges. Instead, the most dominant metrics still compute the similarity between generated code and reference code. In this work, we introduce `ICE-SCORE`, a novel metric that not only addresses the limitations of its predecessors but also harnesses the capabilities of LLMs, setting a new benchmark for the evaluation of code generation tasks.

8 Conclusion

In this paper, we propose a novel evaluation metric based on large language models trained on both text and code, which can better capture the complex syntax and semantics of code intelligence tasks. Our metric achieves superior correlations with subjective human preferences and objective functional correctness, both at the example and corpus levels, without reference and test suites. We conduct an extensive evaluation of four programming languages (Java, Python, C, C++, and JavaScript) and demonstrate the effectiveness of our proposed method on human-based usefulness and execution-based functional correctness. We have publicly released our evaluation metrics and datasets to encourage the development of more accurate and effective evaluation metrics for tasks involving source code.

Acknowledgements

We thank Haolan Zhan and Yufei Wang for the helpful feedback on the paper.

Limitations

Our proposed evaluation metric is based on the assumption that LLMs can follow the instructions to evaluate the code snippets. The backbone models we investigated are closed-source state-of-the-art LLMs from OpenAI. As we noticed that there is a huge performance gap between current closed-source and open-source LLMs, it is possible that ICE-Score can be adapted with an open-source LLM trained on code and text, such as WizardCoder (Luo et al., 2023) and OctoPack (Muenighoff et al., 2023a). Hence, we encourage future investigations on open-source LLMs for code evaluation. In addition, as discussed in Section 6, our experiments only focus on two code generation tasks. There are other code intelligence tasks like program repair and code summarization. However, due to the limited study on human evaluation of these tasks, no open-source dataset is publicly available or documented in detail. Finally, ICE-Score assumes that either model weights or model APIs are available, which is costly for some users. We, therefore, suggest future work on proposing low-cost evaluation metrics.

References

- Toufique Ahmed and Premkumar Devanbu. 2022. Few-shot training llms for project-specific code summarization. In Proceedings of the 37th IEEE/ACM International Conference on Automated Software Engineering, pages 1–5.
- Rachith Aiyappa, Jisun An, Haewoon Kwak, and Yongyeol Ahn. 2023. Can we trust the evaluation on chatgpt? arXiv preprint arXiv:2303.12767.
- Loubna Ben Allal, Raymond Li, Denis Kocetkov, Chenghao Mou, Christopher Akiki, Carlos Munoz Ferrandis, Niklas Muennighoff, Mayank Mishra, Alex Gu, Manan Dey, et al. 2023. Santa-coder: don’t reach for the stars! arXiv preprint arXiv:2301.03988.
- Miltiadis Allamanis, Earl T Barr, Premkumar Devanbu, and Charles Sutton. 2018. A survey of machine learning for big code and naturalness. ACM Computing Surveys (CSUR), 51(4):1–37.
- Satanjeev Banerjee and Alon Lavie. 2005. Meteor: An automatic metric for mt evaluation with improved correlation with human judgments. In Proceedings of the acl workshop on intrinsic and extrinsic evaluation measures for machine translation and/or summarization, pages 65–72.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. Advances in neural information processing systems, 33:1877–1901.
- Sébastien Bubeck, Varun Chandrasekaran, Ronen Eldan, Johannes Gehrmann, Eric Horvitz, Ece Kamar, Peter Lee, Yin Tat Lee, Yuanzhi Li, Scott Lundberg, et al. 2023. Sparks of artificial general intelligence: Early experiments with gpt-4. arXiv preprint arXiv:2303.12712.
- Federico Cassano, John Gouwar, Daniel Nguyen, Sydney Nguyen, Luna Phipps-Costin, Donald Pinckney, Ming-Ho Yee, Yangtian Zi, Carolyn Jane Anderson, Molly Q Feldman, Arjun Guha, Michael Greenberg, and Abhinav Jangda. 2023. MultiPL-E: A scalable and polyglot approach to benchmarking neural code generation. IEEE Transactions on Software Engineering (TSE).
- Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde de Oliveira Pinto, Jared Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, et al. 2021. Evaluating large language models trained on code. arXiv preprint arXiv:2107.03374.
- Xinyun Chen, Maxwell Lin, Nathanael Schärli, and Denny Zhou. 2023. Teaching large language models to self-debug. arXiv preprint arXiv:2304.05128.
- Aryaz Eghbali and Michael Pradel. 2022. Crystalbleu: precisely and efficiently measuring the similarity of code. In 37th IEEE/ACM International Conference on Automated Software Engineering, pages 1–12.
- Mikhail Evtikhiev, Egor Bogomolov, Yaroslav Sokolov, and Timofey Bryksin. 2023. Out of the bleu: how should we assess quality of the code generation models? Journal of Systems and Software, 203:111741.
- Zhangyin Feng, Daya Guo, Duyu Tang, Nan Duan, Xiaocheng Feng, Ming Gong, Linjun Shou, Bing Qin, Ting Liu, Daxin Jiang, et al. 2020. Codebert: A pre-trained model for programming and natural languages. In Findings of the Association for Computational Linguistics: EMNLP 2020, pages 1536–1547.
- Markus Freitag, Ricardo Rei, Nitika Mathur, Chi-kiu Lo, Craig Stewart, Eleftherios Avramidis, Tom Kocmi, George Foster, Alon Lavie, and André FT Martins. 2022. Results of wmt22 metrics shared task: Stop using bleu—neural metrics are better and more robust. In Proceedings of the Seventh Conference on Machine Translation (WMT), pages 46–68.
- Jinlan Fu, See-Kiong Ng, Zhengbao Jiang, and Pengfei Liu. 2023. Gptscore: Evaluate as you desire. arXiv preprint arXiv:2302.04166.

- Abram Hindle, Earl T Barr, Mark Gabel, Zhendong Su, and Premkumar Devanbu. 2016. On the naturalness of software. *Communications of the ACM*, 59(5):122–131.
- Xing Hu, Qiuyuan Chen, Haoye Wang, Xin Xia, David Lo, and Thomas Zimmermann. 2022. Correlating automated and human evaluation of code documentation generation quality. *ACM Transactions on Software Engineering and Methodology (TOSEM)*, 31(4):1–28.
- Jiaxin Huang, Shixiang Shane Gu, Le Hou, Yuexin Wu, Xuezhi Wang, Hongkun Yu, and Jiawei Han. 2022. Large language models can self-improve. *arXiv preprint arXiv:2210.11610*.
- Tom Kocmi and Christian Federmann. 2023. [Large language models are state-of-the-art evaluators of translation quality](#).
- Tom Kocmi, Christian Federmann, Roman Grundkiewicz, Marcin Junczys-Dowmunt, Hitokazu Matsushita, and Arul Menezes. 2021. [To ship or not to ship: An extensive evaluation of automatic metrics for machine translation](#). In *Proceedings of the Sixth Conference on Machine Translation*, pages 478–494, Online. Association for Computational Linguistics.
- Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. Large language models are zero-shot reasoners. In *Advances in Neural Information Processing Systems*.
- Raymond Li, Loubna Ben Allal, Yangtian Zi, Niklas Muennighoff, Denis Kocetkov, Chenghao Mou, Marc Marone, Christopher Akiki, Jia Li, Jenny Chim, et al. 2023. Starcoder: may the source be with you! *arXiv preprint arXiv:2305.06161*.
- Chin-Yew Lin. 2004. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*, pages 74–81.
- Yang Liu, Dan Iter, Yichong Xu, Shuohang Wang, Ruo Chen Xu, and Chenguang Zhu. 2023. Gpteval: Nlg evaluation using gpt-4 with better human alignment. *arXiv preprint arXiv:2303.16634*.
- Ziyang Luo, Can Xu, Pu Zhao, Qingfeng Sun, Xubo Geng, Wenxiang Hu, Chongyang Tao, Jing Ma, Qingwei Lin, and Daxin Jiang. 2023. Wizardcoder: Empowering code large language models with evol-instruct. *arXiv preprint arXiv:2306.08568*.
- Yingwei Ma, Yue Liu, Yue Yu, Yuanliang Zhang, Yu Jiang, Changjian Wang, and Shanshan Li. At which training stage does code data help llms reasoning?
- Aman Madaan, Niket Tandon, Prakhar Gupta, Skyler Hallinan, Luyu Gao, Sarah Wiegrefe, Uri Alon, Nouha Dziri, Shrimai Prabhumoye, Yiming Yang, et al. 2023. Self-refine: Iterative refinement with self-feedback. *arXiv preprint arXiv:2303.17651*.
- Niklas Muennighoff, Qian Liu, Armel Zebaze, Qinkai Zheng, Binyuan Hui, Terry Yue Zhuo, Swayam Singh, Xiangru Tang, Leandro von Werra, and Shayne Longpre. 2023a. Octopack: Instruction tuning code large language models. *arXiv preprint arXiv:2308.07124*.
- Niklas Muennighoff, Alexander M Rush, Boaz Barak, Teven Le Scao, Aleksandra Piktus, Nouamane Tazi, Sampo Pyysalo, Thomas Wolf, and Colin Raffel. 2023b. Scaling data-constrained language models. *arXiv preprint arXiv:2305.16264*.
- OpenAI. 2023. [Gpt-4 technical report](#).
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. 2022. Training language models to follow instructions with human feedback. *Advances in Neural Information Processing Systems*, 35:27730–27744.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting of the Association for Computational Linguistics*, pages 311–318.
- Maja Popović. 2015. chrF: character n-gram f-score for automatic mt evaluation. In *Proceedings of the tenth workshop on statistical machine translation*, pages 392–395.
- Matt Post. 2018. [A call for clarity in reporting BLEU scores](#). In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 186–191, Brussels, Belgium. Association for Computational Linguistics.
- Shuo Ren, Daya Guo, Shuai Lu, Long Zhou, Shujie Liu, Duyu Tang, Neel Sundaresan, Ming Zhou, Ambrosio Blanco, and Shuai Ma. 2020. Codebleu: a method for automatic evaluation of code synthesis. *arXiv preprint arXiv:2009.10297*.
- Baptiste Rozière, Jonas Gehring, Fabian Gloeckle, Sten Sootla, Itai Gat, Xiaoqing Ellen Tan, Yossi Adi, Jingyu Liu, Tal Remez, Jérémy Rapin, et al. 2023. Code llama: Open foundation models for code. *arXiv preprint arXiv:2308.12950*.
- Dominik Sobania, Martin Briesch, Carol Hanna, and Justyna Petke. 2023. An analysis of the automatic bug fixing performance of chatgpt. *arXiv preprint arXiv:2301.08653*.
- Weisong Sun, Chunrong Fang, Yudu You, Yun Miao, Yi Liu, Yuekang Li, Gelei Deng, Shenghan Huang, Yuchen Chen, Quanjun Zhang, et al. 2023. Automatic code summarization via chatgpt: How far are we? *arXiv preprint arXiv:2305.12865*.
- Nigar M Shafiq Surameery and Mohammed Y Shakor. 2023. Use chat gpt to solve programming bugs. *International Journal of Information Technology & Computer Engineering (IJITC) ISSN: 2455-5290*, 3(01):17–22.

Ngoc Tran, Hieu Tran, Son Nguyen, Hoan Nguyen, and Tien Nguyen. 2019. Does bleu score work for code migration? In 2019 IEEE/ACM 27th International Conference on Program Comprehension (ICPC), pages 165–176. IEEE.

Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed H Chi, Quoc V Le, Denny Zhou, et al. Chain-of-thought prompting elicits reasoning in large language models. In Advances in Neural Information Processing Systems.

Pengcheng Yin, Bowen Deng, Edgar Chen, Bogdan Vasilescu, and Graham Neubig. 2018. Learning to mine aligned code and natural language pairs from stack overflow. In International Conference on Mining Software Repositories, MSR, pages 476–486. ACM.

Qinkai Zheng, Xiao Xia, Xu Zou, Yuxiao Dong, Shan Wang, Yufei Xue, Zihan Wang, Lei Shen, Andi Wang, Yang Li, et al. 2023. Codegeex: A pre-trained model for code generation with multilingual evaluations on humaneval-x. arXiv preprint arXiv:2303.17568.

Ming Zhong, Yang Liu, Da Yin, Yuning Mao, Yizhu Jiao, Pengfei Liu, Chenguang Zhu, Heng Ji, and Jiawei Han. 2022. Towards a unified multi-dimensional evaluator for text generation. In Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing, pages 2023–2038.

Shuyan Zhou, Uri Alon, Sumit Agarwal, and Graham Neubig. 2023. Codebertscore: Evaluating code generation with pretrained models of code. In Association for Computational Linguistics: EMNLP 2023.

Kaijie Zhu, Jindong Wang, Jiaheng Zhou, Zichen Wang, Hao Chen, Yidong Wang, Linyi Yang, Wei Ye, Neil Zhenqiang Gong, Yue Zhang, et al. 2023. Promptbench: Towards evaluating the robustness of large language models on adversarial prompts. arXiv preprint arXiv:2306.04528.

Terry Yue Zhuo, Yujin Huang, Chunyang Chen, and Zhenchang Xing. 2023. Red teaming chatgpt via jail-breaking: Bias, robustness, reliability and toxicity.

A Prompts for Code Evaluation

A.1 Code Usefulness

Evaluation Criteria:

Usefulness (0-4) Usefulness of the code snippet based on the problem description.

- A score of 0: Snippet is not at all helpful, it is irrelevant to the problem.

- A score of 1: Snippet is slightly helpful, it contains information relevant to the problem, but it is easier to write the

solution from scratch.

- A score of 2: Snippet is somewhat helpful, it requires significant changes (compared to the size of the snippet), but is still useful.

- A score of 3: Snippet is helpful, but needs to be slightly changed to solve the problem.

- A score of 4: Snippet is very helpful, it solves the problem.

A.2 Functional Correctness

Evaluation Criteria:

Functional Correctness (0-4) - Execution-based quality of the code snippet combined with the problem. The correctness is measured by all possible unit tests and the comparison of the reference code. The combination of the code snippet and the problem should pass all the possible tests based on your understanding of the reference code. The length of the code snippet can not determine the correctness. You need to assess the logic line by line.

- A score of 0 (failing all possible tests) means that the code snippet is totally incorrect and meaningless.

- A score of 4 (passing all possible tests) means that the code snippet is totally correct and can handle all cases.

B Automatic Evaluation Metric Baselines

Our implementations of the automatic evaluation metric baselines except for CodeBERTScore are based on <https://github.com/JetBrains-Research/codegen-metrics>. For CodeBERTScore, we adopt the official release at <https://github.com/neulab/code-bert-score>.

C Correlation Metrics

For all correlation metrics, we use the implementation from <https://scipy.org/> and call these APIs with the default settings.