# Revisiting Code Similarity Evaluation with Abstract Syntax Tree Edit Distance

**Anonymous ACL submission** 

### Abstract

This paper revisits recent code similarity evaluation metrics, particularly focusing on the application of Abstract Syntax Tree (AST) editing distance in diverse programming languages. In particular, we explore the usefulness of these metrics and compare them to traditional sequence similarity metrics. Our experiments showcase the effectiveness of AST editing distance in capturing intricate code structures, revealing a high correlation with established metrics. Furthermore, we explore the strengths and weaknesses of AST editing distance and prompt-based GPT similarity scores in comparison to BLEU score, execution match, and Jaccard Similarity. We propose, optimize, and publish an adaptable metric that demonstrates effectiveness across all tested languages, representing an enhanced version of Tree Similarity of Edit Distance (TSED).

001

004

007 008

011

012

019

024

# 1 Introduction and Related Work

In the fields of natural language processing and software engineering, code generation tasks are gaining more and more attention. Assessing the quality of generated code is now critically important, but we still lack evaluation methods other than traditional statistical sequence evaluation methods. Widely used semantic evaluation metrics like BLEU score and Jaccard similarity rely on statistical characteristics, overlooking the intricate grammatical structures and logical relationships inherent in complex programming languages.

However, recent developments in the NLP field paved the way for novel evaluation metrics which we explore in this study. For one, the staggering number of powerful large language models (LLMs) such as GPT-3.5/4 (Achiam et al., 2023) revolutionized the NLP landscape and led to noteworthy advancements in the realm of code review and evaluation (Wang et al., 2023; Tang et al., 2024). Another recent study introduced the novel TSED metric and used it to evaluate text-to-SQL tasks (Song et al., 2023). For this study, we take advantage of these developments to (1) prompt the GPT-4 model to generate similarity scores for code, and (2) expand on the TSED metric.

041

042

043

044

045

047

050

052

053

056

060

061

062

063

064

065

066

067

068

069

070

072

073

074

075

076

077

We utilize these two different metrics (GPT and TSED) to evaluate the structural similarity of different programming languages and how they relate to execution matches. Furthermore, we address how these metrics are correlated to semantic similarity metrics like the BLEU score. Finally, we investigate some limitations of these metrics by delving into the impact of TSED's penalty weight of tree operations on evaluation accuracy and exploring the stability of outputs from the GPT LLMs.

As a result, we have these 3 contributions from this research: (a) we propose and publish a new tool for 48 programming languages<sup>1</sup>, (b) we discuss 2 recent evaluation metrics and 2 traditional metrics and compare them via correlation coefficient, recall to execution match, (c) we discuss the unstable nature of GPT similarity scoring and the ways to optimize TSED.

#### 2 Approaches

#### 2.1 TSED on Programming Languages

Applying the TSED evaluation method, initially designed for SQL analysis, we have undergone modifications to extend its applicability to various programming languages. The fundamental TSED approach, illustrated in Figure 1, encompasses AST parsing, AST Editing Distance Calculation, and normalization, closely resembling the methodology outlined in the original paper. However, we have made modifications to both the AST parsing and normalization.

**Code Parsing:** Parsing in the domain of programming languages involves parsing raw code

<sup>&</sup>lt;sup>1</sup>https://anonymous.4open.science/r/TSEDwP-7208/README.md



Figure 1: Pipeline of TSED Code Evaluation Metric

text into its associated AST. This parsing underscores the complexity of interpreting various programming constructs and converting them into a structured grammar tree representation.

079

091

100

101

102

103

104

105

106

109

110

111

112

113

114

We use tree-sitter<sup>2</sup> as our AST parser which is based on GLR, a powerful parsing algorithm commonly found in the literature (Latif et al., 2023; Tomita, 1991; Clem and Thomson, 2021).

**Tree Distance Computation:** For calculating tree edit distance as  $\Delta$ , we utilize the same function as outlined in the TSED paper, which is APTED algorithm (Pawlik and Augsten, 2015, 2016). Considering  $G_1$  as predicted code's AST and  $G_2$  are AST from ground-truth:

$$\Delta(G_1, G_2) = \min_{ops} \sum_{i=1}^n w(op_i) \tag{1}$$

Here, ops is a sequence of edit operations transforming  $G_1$  into  $G_2$ , with  $w(op_i)$  as the cost for the  $i^{th}$  operation.

**Normalization:** Normalization of tree edit distances accounts for the complexity of the code by considering the maximum number of nodes between two trees, and we add a ramp function to avoid some extreme situations:

$$TSED = \max\{1 - \frac{\delta}{MaxNodes(G_1, G_2)}, 0\}$$
(2)

This provides a metric for structural similarity comparison of programming code, enabling a nuanced analysis beyond mere syntactic comparison.

#### 2.2 GPT Structure Similarity

Between 2020 and 2023, OpenAI introduced the GPT-3/3.5 and GPT-4 models, showcasing remarkable reasoning capabilities and achieving state-of-the-art performance across numerous tasks (Brown et al., 2020). Our approach involves utilizing prompts to elicit the model's output regarding the structural similarity between two code segments, resulting in a score on a scale from 0 to 1. A score of 1 indicates identical structures, while 0 signifies

complete dissimilarity. Despite its effectiveness, this metric operates as a black box, leaving us unaware of the specific calculations performed by GPT or whether it consistently employs the same metric. From various research papers, we've observed that these LLMs tend to produce more unstable results with each iteration (Tian et al., 2023; Liu et al., 2023). 115

116

117

118

119

120

121

122

123

124

125

126

127

128

129

130

131

132

134

135

136

138

139

140

141

142

143

144

145

146

147

148

149

150

151

152

153

154

155

156

157

158

159

160

162

#### **3** Research Questions and Targets

**RQ1: Can TSED be used in more programming languages?** We investigate the adaptability of AST Edit Distance which is a generalized version of TSED, exploring its effectiveness in languages like Python and Java to assess its applicability for code similarity analysis.

**RQ2: How are TSED and GPT similarity correlated to semantic similarity and execution match?** We assess the correlation between these different metrics to understand their respective contributions in evaluating code similarity across multiple programming languages.

**RQ3: What are the limits of these metrics?** We assess the stability of GPT-based similarity output and analyze how parameters, particularly operation weights (delete, insert, rename), influence TSED.

# 4 Experiments

#### 4.1 General Setup

In this study, our primary objective is to apply the theoretical framework to a diverse range of programming languages. To achieve this, we aim to identify executable datasets and evaluate them using predefined metrics. The experimental setup comprises two key tasks: firstly, expanding the application of TSED and GPT similarity to additional programming languages, followed by exploring the correlation between these metrics. Subsequently, we seek to assess the stability of GPT scoring and examine the impact of various parameters on the TSED metric. This structured approach allows us to comprehensively investigate the adaptability, correlations, and stability of the chosen metrics across a spectrum of programming languages.

#### 4.2 Evaluation Metrics

• **BLEU Score** is calculated as the geometric mean of the modified precision scores for various ngram lengths, providing a concise and standardized similarity measurement between the generated and reference text (Papineni et al., 2002).

<sup>&</sup>lt;sup>2</sup>https://tree-sitter.github.io/tree-sitter/

Jaccard Similarity is a measure of similarity between two sets and is calculated by dividing the size of the intersection of the sets by the size of their union, offering a quantitative assessment of the degree of overlap between the sets' elements.

- Execution Match Execution Match pertains to the consistency in execution outcomes between generated code and its corresponding ground truth, evaluating the equivalence in practical functionality. 1 in Execution match means they have the same execution results, and 0 means different.
- GPT Similarity mentioned in the Section 2.2
  - **TSED** mentioned in the Section 2.1.

### 4.3 Datasets

168

169

170

171

172

173

174

175

176

Although the execution match metric is infre-177 quently employed in programming code-related 178 179 datasets, its prominence has increased in recent years. Our comparative analysis involved assessing datasets from various papers, considering factors 181 such as dataset sizes, programming languages, and executables. As highlighted in Table 1, the MBXP 183 dataset encompasses 13 different languages, serving as a function-level benchmark that effectively evaluates programming paragraphs. However, the MBXP dataset includes ground-truth solutions for only 7 languages, with C# omitted due to compilation issues. Additionally, we consider the CoderEval dataset to facilitate a comparison be-190 tween Python and Java code generation, leveraging 191 its longer test samples, results are in the appendix. 192

 Table 1: Widely-used code generation benchmarks, selected from GitHub

Benchmark	Language	Samples	Executeable
CoNaLA(Yin et al., 2018)	Python	500	No
Concode(Iyer et al., 2018)	Java	2000	No
MBXP(Athiwaratkun et al., 2022)	Multilingual	974	Yes
InterCode(Yang et al., 2023)	Bash, SQL	200, 1034	Yes
CoderEval(Yu et al., 2024)	Python, Java	230	Yes
RepoEval(Liao et al., 2023)	Python	383	No

In the Bash-Shell scenarios, we reproduce results and conduct a comparative analysis using the **InterCode** dataset. Notably, we identify the SPIDER dataset within InterCode and establish it as a baseline. **SPIDER**, previously evaluated in comparison to the TSED paper, is a substantial human-labeled dataset for the text-to-SQL task. This dataset encompasses databases with intricate join solutions across diverse domains (Yu et al., 2018).

#### 5 Results

#### 5.1 Similarity Results

Table 2: Evaluation Metrics comparison for 6 languages on MBXP dataset, prediction generated by GPT-3.5-Turbo model, ground truth from dataset

Languages	TSED	BLEU	Jaccard Sim	GPT-4	Execution
Java	0.3746	0.2041	0.2733	0.8143	0.6550
Python	0.1888	0.0843	0.2000	0.6751	0.6842
JavaScript	0.2037	0.0846	0.2037	0.6763	0.6811
Typescript	0.1360	0.0637	0.1397	0.5313	0.6642
Ruby	0.1727	0.0438	0.1810	0.7067	0.6428
Kotlin	0.3412	0.1847	0.3109	0.7073	0.5569

As we analyze the results presented in Table 2, our experiment demonstrates the effective performance of TSED and GPT similarity in evaluating the MBXP dataset across all 6 programming languages. No instances of parsing or scoring generation failures were observed, confirming the robustness of these metrics across languages.

# RQ1: Can TSED be used in more programming languages?

Answer: The exploration of TSED's adaptability beyond SQL shows promise, especially in languages like Java and Kotlin, indicating its potential for code analysis. TSED proves effective in programming languages with functional parsers, allowing for structural similarity calculation.



Figure 2: MBXP dataset, Pearson Correlation Heatmap between evaluation-metrics on GPT-3.5

Moreover, TSED shows a commendable correlation ranging from 0.6 to 0.8 with BLEU score and Jaccard similarity, as illustrated in Figure 2. Additionally, TSED exhibits a strong correlation with GPT similarity, especially in Java and Python during the CoderEval test, as depicted in Figure 3, underscoring its sensitivity to code structure. We employ thresholding to establish a prediction-toexecution match. If the metric value exceeds the

212

213

214

215

216

217

218

219

220

193

194

199 200

201

202

203

204

205

206

207

209

Table 3: Execution Match F1 score & Accuracy for each thresholding metrics

Languages		TSED		GPT		BLEU			Jaccard			
	Threshold	F1	Acc	Threshold	F1	Acc	Threshold	F1	Acc	Threshold	F1	Acc
Python	0.23	0.5650	0.6057	0.83	0.6403	0.6735	0.07	0.5719	0.6150	0.19	0.5907	0.6253
Java	0.10	0.5108	0.6499	0.56	0.5693	0.6396	0.03	0.5184	0.5755	0.16	0.5612	0.6018
JavaScript	0.12	0.5494	0.6002	0.69	0.5924	0.6205	0.02	0.4964	0.5267	0.12	0.5245	0.5885
Typescript	0.07	0.5367	0.5822	0.51	0.5521	0.5708	0.01	0.4987	0.5553	0.08	0.5284	0.5708
Ruby	0.13	0.5045	0.5306	0.54	0.6051	0.6811	0.01	0.4375	0.4490	0.12	0.5142	0.5612
Kotlin	0.28	0.6834	0.6823	0.8	0.6681	0.6721	0.1	0.6441	0.6457	0.22	0.6387	0.6533

threshold T, we assign the prediction as 1; otherwise, it is set to 0. The optimal threshold values are determined through enumeration to achieve the best match results. Based on their F1/Accuracy match to the Execution match, both TSED and GPT similarity exhibit higher accuracy compared to semantic metrics in Table 3. Notably, GPT similarity demonstrates a slightly superior F1 score and TSED gives good results on accuracy.



Figure 3: CoderEval Pearson Correlation Heatmap between evaluation-metrics/models/languages

RQ2: How are TSED and GPT similarity correlated to semantic similarity and execution match?

Answer: Our evaluation of TSED metrics, GPTbased similarity, and other semantic evaluation metrics revealed consistently high Pearson correlations between TSED, GPT Score, BLEU Score, and Jaccard Similarity. TSED exhibited notable accuracy in matching with Execution-Match, while GPT score demonstrated the highest F1 score, highlighting their respective strengths in capturing structural and semantic nuances in code across various programming languages.

### 5.2 Stability of GPT Scoring

To understand how unstable GPT scoring is, we execute the GPT-4 Similarity scoring five times on identical prediction sets, we establish the initial result as a baseline to assess differences through statistical indicators such as Mean Squared Error (MSE) or Mean Absolute Error (MAE) in comparison to the first scoring. Table 4 demonstrates that GPT scoring exhibits limited stability in the context of code similarity evaluation.

Table 4: Unstable nature of GPT-4 scoring output

Metrics	1st	2nd	3rd	4th
Mean Squared Error	0.0581	0.0583	0.0527	0.0628
Mean Absolute Error	0.1902	0.1940	0.1825	0.1996

#### 5.3 Parameter optimization of TSED

We can configure the penalty weight of 3 operations in tree distance computing: **Delete**, **Insert**, and **Rename**. Figure 4 which is from a test for the MBXP/Java dataset shows is 'Insert' has a sweet spot of 0.8. 'Delete' and 'Rename' operations just keep them in 1.0 penalty weight as the best choice. But we need to keep in mind it can be different in other programming languages.



Figure 4: Change each of penalty weight influence correlation to GPT structure similarity score

RQ3: What are the limits of these metrics?

**Answer:** Penalty weight parameters play influential roles in the TSED metric. Besides, GPT-based similarity metrics offer higher performance at the cost of more money, leading to a bit of unstable output. This underscores the need to carefully balance performance and stability considerations in code similarity assessment across various programming languages.

### 6 Conclusion

In this paper, we applied TSED to more programming languages, compared GPT similarity and TSED to semantic metrics, and checked representation to execution match. Then we discuss limitations about the stability of GPT scoring and the penalty parameters of TSED. 249

241

242

243

245

246

247

248

251 252

250

257

229

239

# 258 Limitations

While our study provides valuable insights into code similarity assessment using TSED and GPTbased metrics, it is essential to acknowledge certain 261 limitations. Firstly, the generalizability of our find-262 ings may be influenced by the specific datasets and 263 264 programming languages employed in our analysis. Additionally, the stability of GPT-based similarity 265 metrics, as highlighted in our results, poses a limitation in terms of consistent and reliable code assessments. Furthermore, variations in the interpretation 269 and definition of similarity metrics across different studies may introduce inherent biases. Lastly, 270 the effectiveness of TSED metrics may be con-271 tingent upon the quality of the employed parsers 272 and the fine-tuning of penalty parameters. These limitations underscore the need for caution when 274 extrapolating our results to diverse contexts and emphasize the necessity for further research to address 276 these challenges. 277

# 278 Ethics Statement

279

283

284

289

290

291

292

294

296

297

298

299

306

307

Our research adheres to ethical standards, prioritizing integrity and respect for all involved parties. We ensured data privacy, obtained informed consent where applicable, and maintained transparency in our methodologies. The study was conducted with the utmost consideration for ethical guidelines and the welfare of participants, upholding the principles of fairness, accountability, and academic integrity throughout the research process.

#### References

- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Ben Athiwaratkun, Sanjay Krishna Gouda, Zijian Wang, Xiaopeng Li, Yuchen Tian, Ming Tan, Wasi Uddin Ahmad, Shiqi Wang, Qing Sun, Mingyue Shang, et al. 2022. Multi-lingual evaluation of code generation models. arXiv preprint arXiv:2210.14868.
- Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish,

Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. Language models are few-shot learners. 308

309

310

311

312

313

314

315

316

317

318

319

323

324

325

326

328

329

330

331

332

333

334

335

336

337

338

339

340

341

342

343

344

345

346

347

348

349

350

351

352

353

354

355

356

357

358

- Timothy Clem and Patrick Thomson. 2021. Static analysis at github: An experience report. *Queue*, 19(4):42–67.
- Srinivasan Iyer, Ioannis Konstas, Alvin Cheung, and Luke Zettlemoyer. 2018. Mapping language to code in programmatic context. *arXiv preprint arXiv:1808.09588*.
- Afshan Latif, Farooque Azam, Muhammad Waseem Anwar, and Amina Zafar. 2023. Comparison of leading language parsers–antlr, javacc, sablecc, tree-sitter, yacc, bison. In 2023 13th International Conference on Software Technology and Engineering (ICSTE), pages 7–13. IEEE.
- Dianshu Liao, Shidong Pan, Qing Huang, Xiaoxue Ren, Zhenchang Xing, Huan Jin, and Qinying Li. 2023. Context-aware code generation framework for code repositories: Local, global, and third-party library awareness. *arXiv preprint arXiv:2312.05772*.
- Xiao Liu, Yanan Zheng, Zhengxiao Du, Ming Ding, Yujie Qian, Zhilin Yang, and Jie Tang. 2023. Gpt understands, too. *AI Open*.
- 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 on Association for Computational Linguistics, ACL '02, page 311–318, USA. Association for Computational Linguistics.
- Mateusz Pawlik and Nikolaus Augsten. 2015. Efficient computation of the tree edit distance. *ACM Transactions on Database Systems (TODS)*, 40(1):1–40.
- Mateusz Pawlik and Nikolaus Augsten. 2016. Tree edit distance: Robust and memory-efficient. *Information Systems*, 56:157–173.
- Yewei Song, Saad Ezzini, Xunzhu Tang, Cedric Lothritz, Jacques Klein, Tegawendé Bissyandé, Andrey Boytsov, Ulrick Ble, and Anne Goujon. 2023. Enhancing text-to-sql translation for financial system design. arXiv preprint arXiv:2312.14725.
- Daniel Tang, Zhenghan Chen, Kisub Kim, Yewei Song, Haoye Tian, Saad Ezzini, Yongfeng Huang, and Jacques Klein Tegawende F Bissyande. 2024. Collaborative agents for software engineering. *arXiv preprint arXiv:2402.02172*.
- Haoye Tian, Weiqi Lu, Tsz On Li, Xunzhu Tang, Shing-Chi Cheung, Jacques Klein, and Tegawendé F Bissyandé. 2023. Is chatgpt the ultimate programming assistant-how far is it? *arXiv preprint arXiv:2304.11938*.
- Masaru Tomita. 1991. *Generalized LR parsing*. Springer Science & Business Media.

Junjie Wang, Yuchao Huang, Chunyang Chen, Zhe Liu, Song Wang, and Qing Wang. 2023. Software testing with large language model: Survey, landscape, and vision. *arXiv preprint arXiv:2307.07221*.

361

363

365

366

368

372

373

374

375

379

382

386

- John Yang, Akshara Prabhakar, Karthik Narasimhan, and Shunyu Yao. 2023. Intercode: Standardizing and benchmarking interactive coding with execution feedback. *arXiv preprint arXiv:2306.14898*.
  - 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 *Proceedings of the 15th international conference on mining software repositories*, pages 476–486.
- Hao Yu, Bo Shen, Dezhi Ran, Jiaxin Zhang, Qi Zhang, Yuchi Ma, Guangtai Liang, Ying Li, Qianxiang Wang, and Tao Xie. 2024. Codereval: A benchmark of pragmatic code generation with generative pre-trained models. In *Proceedings of the 46th IEEE/ACM International Conference on Software Engineering*, pages 1–12.
- Tao Yu, Rui Zhang, Kai Yang, Michihiro Yasunaga, Dongxu Wang, Zifan Li, James Ma, Irene Li, Qingning Yao, Shanelle Roman, et al. 2018. Spider: A large-scale human-labeled dataset for complex and cross-domain semantic parsing and text-to-sql task. *arXiv preprint arXiv:1809.08887*.

# A Additional Experiment Details

# A.1 Parser Comparison

The ANTLR<sup>3</sup> (ANother Tool for Language Recognition) tool, serving as a distinct AST parser compared to tree-sitter, demonstrated notable differences. Following our evaluation using identical settings for TSED metrics, as Figure 5 shows, it became evident that the correlation with other metrics was inferior to the original solutions. This experiment underscores the crucial role of parser performance in the computation procedure, highlighting the significance of selecting an appropriate parser for accurate and reliable code similarity assessments.



Figure 5: CoderEval Java Pearson Correlation Heatmap between evaluation-metrics/models/languages on TSED with ANTLR parser

### A.2 Other results

Due to space constraints, a subset of experimental data is provided in the appendix. A comprehensive evaluation of CoderEval and InterCoder is detailed in Table 5, while specific original sample data from the MBXP dataset is presented in Table 6.

CoderEval, designed for class-level code generation tasks, proves to be a challenging test. Utilizing Pass@10 data as a test sample, TSED demonstrates a robust correlation with semantic indicators in both Java and Python languages. Additionally, a noteworthy correlation is observed between TSED and GPT Similarity.

In the case of InterCoder, we confirm that TSED calculations extend to Bash scripts. Also, the correlation in Figure 6 between TSED to semantic metrics is acceptable, the GPT score doesn't have a good correlation to others. We also replicate the performance of the SPIDER dataset, noting differences from the original paper but not to a significant extent.

Despite the notably low semantic similarity between the MBXP built-in samples and the ground

401

402

403

404

405

406

407

408

409

410

411

412

413

414

415

416

417

418

419

420

421

422

423

387

390

391

<sup>&</sup>lt;sup>3</sup>https://www.antlr.org/

truth, a relatively high execution match is observed. We acknowledge this disparity and plan to address it through optimization in future research endeavors.

424

425

426

427

Table 5: 4 Evaluation Metrics compared to Ground Truth on CoderEval(Java&Python) / InterCode(Bash) / SPIDER(SQL)

Languages	Model	TSED	BLEU	Jaccard Sim	GPT-4	Execution
Java	ChatGPT	0.4971	0.3655	0.3384	0.7392	0.3539
	CodeGen	0.3616	0.2871	0.2506	0.6603	0.1391
	PanGu	0.5029	0.3722	0.3849	0.6778	0.2543
Python	ChatGPT	0.2840	0.1285	0.1763	0.5883	0.2104
	CodeGen	0.2703	0.1778	0.1821	0.5604	0.0948
	PanGu	0.2829	0.0868	0.1567	0.5086	0.1183
	GPT-4	0.5853	0.2816	0.3567	0.8511	0.4851
Shell	starchat	0.4065	0.1594	0.2081	0.6740	0.2374
	vicuna	0.4755	0.1621	0.2295	0.7164	0.2451
	ChatGPT-3.5	0.6824	0.3304	0.3710	0.9461	0.6482
SQL	nsql-6B	0.8022	0.4493	0.4356	0.9265	0.5483
	RESDSQL	0.7422	0.2084	0.1868	0.9629	0.7756

Table 6: 4 Evaluation Metrics compare to Ground Truthon 7 languages MBXP Dataset Samples

Languages	TSED	BLEU	Jaccard Sim	GPT-4	Execution
Java	0.2218	0.1046	0.1960	0.4248	0.853
Python	0.1550	0.0255	0.1222	0.3396	0.822
JavaScript	0.1870	0.0573	0.1685	0.4005	0.786
Typescript	0.1186	0.0288	0.1260	0.4247	0.872
Ruby	0.2073	0.0235	0.1796	0.4830	0.589
Kotlin	0.1720	0.0336	0.1877	0.3976	0.637



Figure 6: InterCode/SPIDER Pearson Correlation Heatmap between evaluation-metrics/models/languages