EXCISION SCORE: EVALUATING EDITS WITH SURGICAL PRECISION

Anonymous authorsPaper under double-blind review

ABSTRACT

Many tasks revolve around editing a document, whether code or text. We formulate the revision similarity problem to unify a wide range of machine learning evaluation problems whose goal is to assess a revision to an existing document. This definition rests on the observation that revisions usually change only a *small* portion of an existing document, so the existing document and its immediate revisions share a majority of their content.

We formulate five adequacy criteria for revision similarity measures, designed to align them with human judgement. We show that popular pairwise measures, like BLEU, fail to meet these criteria, because their scores are dominated by the shared content. They report high similarity between two revisions when humans would assess them as quite different. This is a fundamental flaw we address.

We propose a novel static measure, Excision Score (ES), which computes longest common subsequence (LCS) to remove content shared by an existing document with the ground truth and predicted revisions, before comparing only the remaining divergent regions. This is analogous to a surgeon creating a sterile field to focus on the work area. We use approximation to speed the standard cubic LCS computation to quadratic. In code-editing evaluation, where static measures are often used as a cheap proxy for passing tests, we demonstrate that ES surpasses existing measures. When aligned with test execution on HumanEvalFix, ES improves over its nearest competitor, SARI, by 12% Pearson correlation and by >21% over standard measures like BLEU. The key criterion is invariance to shared context; when we perturb HumanEvalFix with increased shared context, ES' improvement over SARI increases to 20% and >30% over standard measures. ES also handles other corner cases that other measures do not, such as correctly aligning moved code blocks, and appropriately rewarding matching insertions or deletions.

1 Introduction

Editing is a core skill across countless professions, from writers refining drafts to scientists revising research papers. Example tasks from natural language processing (NLP) include sentiment and style transfer Sudhakar et al. (2019), text simplification Al-Thanyyan & Azmi (2021), grammatical error correction Bryant et al. (2023), and updating factual information Logan IV et al. (2022). Nowhere is this more true than in software development, where code evolves through relentless incremental iteration — bug fixes, optimizations, and feature updates — making precise, efficient editing not just useful but essential. Indeed, many AI4Code tasks boil down to editing code: like automated program repair Monperrus (2023), next edit suggestion Chen et al. (2025), refactoring Pomian et al. (2024), and code commenting Panthaplackel et al. (2020), to name a few.

In this work, we focus on revision tasks, which we define as purposeful edits to a document, whether text or code, that preserve its core semantics. This distinguishes it from rewriting or summarisation, which can fundamentally change a document's thesis or structure. We therefore contend that a revision must, by definition, maintain a high degree of similarity to its source. Operationally, we posit that a revision alters a relatively small portion of a text. While the definition of "small" is necessarily task-dependent, we argue that establishing a practical threshold for tasks is feasible and that larger changes can often be decomposed into a sequence of smaller ones, *aka* revisions.

The LLM tsunami has led to the emergence of edit assistants for both text and code revision. Assessing these assistants introduces the *revision similarity problem*: defining a measure for the similarity of two revisions of an initial text (or code) that is aligned with human judgement. With such a measure, one can quantify an assistant's performance by how similar its revision is to the reference. For some tasks, building a golden set of references can be prohibitively expensive, calling for a symmetric measure of revision similarity, *i.e.* one that equally weights its input revisions, allowing it to better tolerate a noisy references. Another use for a symmetric measure is clustering revisions. For example, imagine being the maintainer of a Linux kernel subsystem. Rather than manually assess many patches, clustering them by revision similarity and only reviewing representative patches would save time.

Model performance on revision tasks cannot be assessed by humans at scale, so we need an automated measure. Crucially, we need a normalised measure, not a raw distance: if this is not immediate, consider how two operands can be arbitrarily distant in absolute terms, yet arbitrarily similar as a function of their length. Specifically, we want a similarity measure, one that returns a score in [0..1], where 0 denotes utter dissimilarity and 1 identity. This measure should be task-agnostic, interpretable, and lightweight.

These three properties rule out dynamic measures, notably pass@k, that rely on execution. Their executability constraint is crippling. Even ignoring NLP tasks, many AI4Code tasks do not produce executable code, like code summarization and commit message generation. Even executable code can be nontestable Weyuker (1982). Even considering only code generation, their utility falters in the face of incomplete codebases. Even restricted to tasks that produce testable code, dynamic measures under-approximate program behaviour Dijkstra (1972), which undermines their interpretability, and can be prohibitively computationally expensive. For example, Neubig & Wang (2024) report that evaluation on some 300 samples of SWE-Bench-like dataset took them 2 days; Adamczewski (2025) managed to reduce it to 1 hour per 500 dataset samples with powerful hardware and dedicated containerized environments optimized for the task. Thus, an effective measure should be static.

Existing static measures of textual similarity fall into three categories: lexical, edit-based, and semantic. Lexical measures decompose text into a multiset of predefined features, like *n*-grams, then calculate the similarity of two multisets by atomically comparing their elements. While their *n*-grams do capture local order, they are oblivious to global order. Edit-based measures, in contrast, operate on sequences, so they are inherently sensitive to order. BLEU, ROUGE, Jaccard (adapted to multisets), and TF-IDF are prominent examples of lexical measures. Normalised edit distance built using Levenshtein edit distance is the preeminent edit-based similarity measure. Canonical semantic measures are Word Mover's distance Kusner et al. (2015) and BERTScore Zhang* et al. (2020). These measures struggle with rare words, domain-specific jargon, and nuanced linguistic phenomena like negation and sarcasm. Their scores are often hard to interpret, unlike counting matching n-grams; for example, the difference between scores of 0.7 and 0.8 may not be meaningful or consistent across different models and datasets. When applied to the revision problem, these measures are dominated by the underlying similarity of the revisions and the original text (Section 2).

In this work, we proposing the umbrella term "revision similarity" to unite all ML tasks that can be evaluated by three sequences, an original document and two revisions of it, one a golden reference, and the other, the hypothesis to evaluate. We specify five adequacy criteria that any measure of revision similarity should meet and show how many popular measures fail to meet them. We introduce a new measure — EXCISIONSCORE (ES) — that does. It is a static, task-agnostic, interpretable, and lightweight measure. ES aligns a pair of revisions with their source document to focus on their divergent regions, whose n-grams it compares.

2 STORM CLOUDS IN A BLEU SKY

To assess revision quality, direct comparison seems natural. However, popular pairwise similarity measures, like normalised edit similarity, BLEU, ROUGE, METEOR, and chrF, tend to go wrong because revisions of the same initial version are usually inherently similar, while what matters is comparing their *changes*, not the shared context.

For example, suppose an LLM is asked to replace "anim" with "bar" and outputs

¹Although these estimates include the time needed to run the inference of an LLM, we believe they illustrate well the hardships connected with execution-based measures.

"Lorem ipsum dolor sit amet, consectetur adipiscing elit, sed do eiusmod tempor incididunt ut labore et dolore magna aliqua. Ut enim ad minim veniamfoo, quis ... officia deserunt mollit anim id est laborum."

The LLM clearly failed the task: instead of replacing anim with **bar**, it incorrectly substituted **veniam** with **foo**. Most people would consider this edit wrong. Yet, popular pairwise metrics will all score close to their maximum value of 1, contradicting human judgment.

We now generalise this example, then use it to show how BLEU goes wrong in such cases.

Example 1 (Similar strings). Let X, Y, Z and W be strings and let the original sequence be XY, the assistant's prediction be XW, and the reference revision be XZ. Let ED denote edit distance. Now imagine asking a language model to replace Y with Z while keeping the common prefix X but failing and instead replacing Y with W. Let us assume

$$ED(Y, Z) = ED(Y, W) = ED(Z, W) = |Y| = |W| = |Z| = l \ll |X|.$$

In this example, the assistant (i.e. an LLM) replaced Y with something completely wrong, so we expect a poor score from any measure well-aligned with human intuition. Consider BLEU applied to Example 1, viz. BLEU(XW, {XZ}). On this example, BLEU's brevity penalty will be 1 and, in the limit as $|X| \to \infty$, the ratios of matched n-grams to all n-grams in X will go to 1, $\forall n$, so BLEU(XW, {XZ}) = 1. In short, although the LLM clearly failed the task, BLEU awards a maximal score to tasks captured by Example 1. All other popular pairwise measures, like normalised edit similarity, ROUGE, METEOR, and chrF, fail in the same way: Like BLEU, in the limit as |X| increases with l fixed, all these metrics go to 1, the perfect match.

We are not the first to observe this problem. When Logan IV et al. (2022) proposed a new benchmark for assessing LLM's ability to make factual updates to text, they observed "ROUGE is Problematic. We provide ROUGE F-scores [...] In contrast to the previous results, we find that the simple copy source baseline attains a strong score of 77.4 despite making no updates. [...] This illustrates the importance of evaluating on updates rather than the whole text." The fact that the authors even tried to apply ROUGE to the task where its use is, by their own admission, problematic highlights a blind spot in the community's view of how revision similarity-like problems should be evaluated.

3 ExcisionScore: Measuring Revision Similarity

Popular pairwise similarity measures fail to solve the revision similarity problem when they are dominated by the shared context inherited from an origin string O. We formalize shared context in terms of three-way alignment (Definition 1), then propose 5 Adequacy Criteria, including invariance to shared context, required for a measure of revision similarity to align with human preference. In Section 3.3, we investigate whether existing measures satisfy these criteria. Finally, we define Excision Score and discuss some of its properties.

3.1 CORE CONCEPTS AND UTILITIES

A sequence s is a *revision* of an original document O when $ED(s,O)<\tau$ for some small τ . The specific threshold τ is task-dependent. Notably, closeness in terms of edit distance implies that a revision is necessarily of similar length. As τ approaches $\max\{|s|,|O|\}$, the edits become so destructive that the resulting sequence is less of a refinement and more of a new document, even if it retains the original's core semantics, as in the case of summarizing verbose text. We take A,B to be revisions of the origin O.

Let Σ denote the set of tokens our documents constist of. Let $-\notin \Sigma$ be the dedicated gap symbol. Let $\Sigma_- := \Sigma \cup \{-\}$. Let Σ_-^* and Σ^* stand for the set of finite sequences with and without gaps, respectively.

Definition 1 (Three-Way Alignment). For three sequences $A, B, O \in \Sigma^*$, a *three-way alignment* is a rectangular array R of three rows such that (1) each element of R lies in Σ_- ; (2) no column of R consists of gaps only; and (3) $\operatorname{ungap}(R_1) = A$ and $\operatorname{ungap}(R_2) = B$ and $\operatorname{ungap}(R_3) = O$, where R_i refers to the i-th row of the array R and $\operatorname{ungap}(\Sigma_+) = 0$ removes gaps from a sequence.

Unlike the standard definition (Gusfield, 1997, §14), our Definition 1 focuses on a special case of three sequences and explicitly rules out columns with all gaps, which can be useful when studying evolution or as a placeholder for missing data but are meaningless in our setting.

For an alignment table R, a column of R is *conserved* if all rows in it are identical. If a column is not conserved, it is divergent. A divergent region is, informally, a cluster of adjacent divergent columns.

Definition 2 (Divergent Region). For a three-way alignment array R, a non-empty sub-array d of Ris called a divergent region if (1) d has three rows, same as R (it is a mini-alignment table); (2) all columns of d are divergent and contiguous in R; and (3) there is no divergent column in R that would be adjacent to d (maximality).

An alignment can yield several, possibly no, divergent regions. Let $\Pi_{\text{div}}(R) = \langle d_1, \dots, d_k \rangle$ denote the divergent region projection that produces the list of the k divergent regions extracted from the alignment of A, B, O.

Example 2. Let A = CGTCAA, B = CGCACT, O = CTGCAATT. Below is one possible alignment. Although here we are using 4 letters with significance in biology for simplicity, note that alignment can operate at a coarser token level, e.g. $\Sigma = \text{English words}$.

162

163

164

165

166

167

168

169

170

171 172

173

174

175

176

177

178

179

180 181

182

183

184

185

187 188

189

190

191

192

193

194

195

196

197

198

199

200

201

202

203

204

205

206

207 208

209

210

211 212

213

214

215

In the example, columns 1, 5 and 6 are conserved, the others are divergent. Divergent columns can be understood as atomic edits performed on O by A and B. For example, column 2 shows that A $\Pi_{\text{div}}(R) = \left\langle \begin{bmatrix} G & T & G \\ - & G & G \\ - & T & G \end{bmatrix}, \begin{bmatrix} A & T & - \\ C & T & T \\ A & T & T \end{bmatrix} \right\rangle$. and B both decided to remove G from O. There are two divergent regions, highlighted by the red rectangles, shown on the right:

$$\Pi_{\mathrm{div}}(R) = \left\langle \left| \begin{smallmatrix} \mathsf{G} & \mathsf{T} & - \\ - & - & \mathsf{G} \\ - & \mathsf{T} & \mathsf{G} \end{smallmatrix} \right|, \left| \begin{smallmatrix} \mathsf{A} & - & - \\ \mathsf{C} & \mathsf{T} & - \\ \mathsf{A} & \mathsf{T} & \mathsf{T} \end{smallmatrix} \right| \right\rangle$$

3.2 ADEQUACY CRITERIA FOR REVISION SIMILARITY

Recall that A, the reference, and B, the hypothesis, are revisions of the original document O. We contend that all revision similarity measures should possess the following intuitive properties: (1) Reward edits on A and B agree; (2) Penalize edits on which B disagrees with A; (3) Invariant to shared context (matches across all of A, B and O); (4) Origin-variant (O changing with A and B fixed); and (5) Reward semantically equivalent mismatches.

Properties 1 and 2: Rewarding matches while penalizing mismatches is at the core of any ML evaluation task. Revision similarity is no exception, motivating these properties. The word "edits" implies existence of O, to which edits are applied, tying them to revision similarity. Despite apparent their simplicity, there are several interesting edge cases, one of which we examplify below.

Example 3 (Agreeing on Deletions). Let $D, K, R \in \Sigma^*$ be non-overlapping and assume O = DKRwhere D is deleted by both revisions, K is kept unchanged and R is replaced. Assume that A and B utterly disagree on what to replace R with, i.e. $A = KR_A$ and $B = KR_B$ with R_A sharing no overlap with R_B . Although A and B differ in each replacing R with something different, they do agree on deleting D. Therefore a human would expect a partial similarity score.

In Section 2, we illustrated that measures that reward the shared context as match, violating **Prop**erty 3, do not align with human judgement. We rely on the notions three-way alignment (Definition 1) and divergent regions (Definition 2) to formalize invariance to shared context.

Property 3 (Invariance to Shared Context). A revision similarity measure m(A, B; O) is invariant to shared context iff $\forall A, A', B, B', O, O' \in \Sigma^*$

$$\Pi_{\text{div}}(A, B, O) = \Pi_{\text{div}}(A', B', O') \implies m(A, B; O) = m(A', B'; O').$$

In words, if the divergent regions of (A, B, O) match those of (A', B', O'), a measure m invariant to shared context must return identical scores on those two inputs.

Property 3 equates shared context with conserved columns. Ignoring shared context, i.e. adding or removing conserved columns, is thus equivalent to only considering the divergent regions. A special case of **Property 3** is when (A, B, O) differs from (A', B', O') by a common prefix or suffix. For all sequences α, β that do not overlap any of A, B, O, measure m must satisfy m(A, B; O) = $m(\alpha A\beta, \alpha B\beta; \alpha O\beta).$

Another way to conceptualize invariance to shared context would be to constrain the values of m to inputs where the hypothesis revision matches the origin, B=O. You can think of this as evaluating a "do-nothing" baseline system that simply echoes the input to produce the output revision. Clearly, such a system should get a bad score, e.g. zero: m(O,B;O)=0. A measure that ignores O has no way of identifying this baseline.

Property 4 (Origin-variant). When A and $B \neq A$ be fixed, while we edit O_1 to form a sequence of variants $\langle O_1, O_2, \ldots \rangle$, where $\mathrm{ED}(O_i, A) = \mathrm{ED}(O_i, B) < \mathrm{ED}(O_{i+1}, A) = \mathrm{ED}(O_{i+1}, B)$, $m(A, B; O_i) > m(A, B; O_{i+1})$ must hold.

In contrast to **Property 3**, which constrains a measure's handling of added and removed conserved columns in the 3-way alignments, this property concerns changes to O's row. The strict inequality in the variant sequence restricts the changes to conserved columns. Let $l_i = \mathrm{ED}(O_i, B) = \mathrm{ED}(O_i, A)$, as O moves away from A and B. We argue that revision similarity should increase along with l_i . Indeed, as O becomes more and more distant, A and B are implicitly and independently applying a larger and larger set of matching edits to O, increasing their mutual revision similarity.

Finally, **Property 5** introduces dependence on semantics of the origin document.

Property 5 (**Reward Semantically Equivalent Mismatches**). Edits should be permitted to differ as long as they result in semantically equivalent revisions. Often, penalizing lexical differences that result in no semantic difference is seen as undesirable (according to measurable human preference on these tasks).

As a trivial example, consider the case when O contains 3 blank lines, A inserts some text after the first blank line, while B inserts this same text after the second. Although the insertions do not perfectly match, whitespace-only difference might be permissible. As another example, consider a code editing problem where the task is to add a certain method to a class. In most programming languages, method definition order plays no role. Thus, even if the content of the inserted method was correct, it is likely to be penalized for inserting it at the wrong position. In NLP, this **Property 5** is often partially addressed by providing multiple ground truth references in the hope of covering multiple semantic equivalence classes. Thus, mismatch in case of semantic equivalence can be tolerated by the reference set or the measure itself, with the latter approach having some advantages.

Satisfying **Property 5** is undecidable. There are many ways to partially satisfy it. Our new measure, EXCISIONSCORE does so by using n-grams to be locally insensitive to positional mismatches, as we show in Section 3.4. Empirically, this accounts for many semantically equivalent mismatches in our data sets.

3.3 THE UNMET NEED FOR ADEQUATE REVISION SIMILARITY METRICS

An intuitive idea for solving the revision similarity problem is to locate and strip out a Longest Common Subsequence (LCS) between the origin O and the two revisions A, B originating from it before applying some pairwise similarity measure. Formally, let us denote the deletion of a subsequence by \setminus and the pairwise similarity measure by $P: \Sigma^* \times \Sigma^* \to [0,1]$, where P=1 on exactly matching inputs and P=0 if the inputs are utterly dissimilar. Then we define

$$SansLCS_{P}(A, B \mid O) \triangleq P(A \setminus L, B \setminus L) \quad \text{where} \quad L = LCS(A, B, O) \tag{1}$$

Unlike pairwise measures, $\operatorname{SansLCS}_P$ is invariant to shared context being added to O,A,B, satisfying **Property 3**. However, $\operatorname{SansLCS}_P$ comes with flaws of its own. By stripping out the LCS and considering only A,B, we lost the information about what A and B deleted, making it impossible to partially reward agreement on deletions. In Example 3, $\operatorname{LCS}(A,B,O)=K$ and $\operatorname{SansLCS}_P(A,B\mid O)=P(KR_A\setminus K,KR_B\setminus K)=P(R_A,R_B)=0$. Additionally, $\operatorname{SansLCS}_P$ introduces substring matches that were not possible when comparing A with B directly. By removing the LCS, we introduced n-grams at the junctions that existed in neither A nor B which P might reward, if they happen to match.

A metric named **DiffBLEU** was recently proposed by Bairi et al. (2024); Munson et al. (2022) in the context of code editing and is defined as BLEU(diff(O, B), diff(O, A)) where diff is the output of the diff program The IEEE and The Open Group (2018) with optional post-processing. Thanks to the clever application of pairwise diff, DiffBLEU adequately addresses the problem of shared content across three revisions. However, lumping deleted and inserted lines together and comparing the

concatenated diffs, as opposed to treating them separately, is problematic. First of all, this approach rewards accidental n-gram matches across different operation types. For example, a word inserted by A would match a word deleted in O by B. Such matches do not correspond to what a human would perceive as similarity and thus should not be rewarded. Secondly, while rewarding agreement on deletions by letting BLEU match the deleted lines prefixed by -, DiffBLEU is prone to overrewarding it, as the following example shows.

Example 4 (LHS/RHS Agreement Balance). Continuing Example 3, suppose that D is empty, so O = KR, $A = KR_A$, and $B = KR_B$. Replacements $(R \mapsto R_A)$ and $(R \mapsto R_A)$ can be viewed as two consecutive operations: first deleting R, then inserting R_A or R_B . In that view, R_B and R_B agree on deleting R_B , the left-hand side (LHS) of the replacement, but disagree on R_A versus R_B , the right-hand side (RHS). The high BLEU score given to the matching deleted lines will dominate the mismatch between the RHS's of the replacements, contradicting human judgment.

SARI (System output Against References and against Input) is a text similarity measure that compares a system's edit (e.g., a simplification) to multiple references and the original input, evaluating the appropriateness of added, deleted, and kept n-grams via precision/recall/F1 scores Xu et al. (2016). SARI is defined as

$$SARI(I, O, R) = \frac{1}{3}(F_{add} + F_{keep} + P_{del})$$
(2)

where F_{keep} and F_{keep} stand for F_1 score of the corresponding operation and P_{del} stands for the precision of deletions, all three averaged over n-grams of order n=1..4. In the limit we described in Example 1, when the shared context dominates, F_{keep} term of SARI is always greater than $1-\epsilon$, where $\epsilon \to 0$. This narrows SARI's effective range of values down to $[\frac{1}{3}-\epsilon,1]$. Although SARI does not fully step into the pitfall of pairwise measures and accounts for O, it is not invariant to shared context, failing **Property 3**. In the next subsection we describe EXCISIONSCORE that builds upon SARI and addresses that flaw.

3.4 EXCISIONSCORE DEFINED

Armed with the insight that shared context should be removed and sidestepping the mistakes of SansLCS, we define EXCISIONSCORE (ES) as follows:

$$ES(A, B; O) \triangleq SARI(A \setminus L, B \setminus L, O \setminus L) \quad \text{where} \quad L = LCS(A, B, O). \tag{3}$$

After excising the LCS like $SansLCS_P$ does, ES applies SARI. Recall that P in $SansLCS_P$ was a pairwise measure, which made it impossible to reward agreement on deletions (Example 3). In contrast, SARI accepts all three documents A,B,O as arguments and has a special term P_{del} (eq. (2)) dedicated to that. In terms of three-way alignment, removing the LCS can be thought of as extracting divergent regions and concatenating them, then removing the gaps. When converting strings to a set of n-grams, we omit the n-grams that span several divergent regions, sidestepping the other flaw of SansLCS.

EXCISIONS CORE meets the revision similarity adequacy criteria. EXCISIONS CORE identifies edits as kept, added, or removed n-grams, correctly rewarding agreement on deletions, meeting **Properties 1** and **2**. We discussed that due to the F_{keep} term, SARI is not invariant to shared context and can award a score of up to $\frac{1}{3}$ on the "do-nothing" baseline ($B=O\neq A$). We fix this by excising the shared context and ensuring that the F_{keep} term does not saturate. Thanks to that, ES correctly returns 0 on the do-nothing baseline and satisfies **Property 3**, if we neglect rare accidental matches that could happen with any n-gram-based measure even when computed on random sequences. ES meets **Property 4**: Changing a shared token in O while keeping A and B fixed turns a previously ignored conserved column into a new divergent region on which A and B agree, increasing P_{del} and F_{add} in Equation (2). Finally, ES partially satisifises **Property 5** by matching misplaced insertions, which we found to be a common case in CanItEdit dataset we use in Section 4.

EXCISIONSCORE relies on LCS, which, if computed exactly, implies $\mathcal{O}(l^3)$ time complexity, l=|O|. For long |O|, this quickly becomes impractical, so we compute L in Equation (3) approximately. In our implementation, L is a not necessarily longest common subsequence computed as LCS(LCS(O,A),LCS(O,B)). Two-way LCS computed 3 times brings the time complexity down to $\mathcal{O}(l^2)$.

4 EVALUATION: EXCISIONS CORE AS EXECUTION PROXY

Quality measures on coding tasks can be categorized into *semantic correctness* of the produced code and its *syntactic similarity* to some ground truth reference. Semantic correctness, code's functional properties are usually assessed by executing a suite of tests shipped with the evaluation dataset. Equipping a dataset with a high-coverage test suite is typically a difficult and expensive task, and so is executing these tests in the specialized environment. Syntactic similarity, on the other hand, is very easy to estimate with metrics such as BLEU/ROUGE/chrF, and the required ground truth label is often cheap to mine by masking out the already existing code. It's a questionable but widely accepted practice to use syntactic similarity as a cheap proxy to the expensive semantic correctness verification for AI-generated programs. To address this practice, we explore how well ExcisionScore, along with other popular static measures, correlates with test execution.

Datasets We consider two code editing datasets, where each dataset item consists of a code snippet, a natural language edit instruction, a reference solution, and a test suite to verify correctness. **HumanEvalFix** Muennighoff et al. (2023) contains 984 buggy code snippets across 6 programming languages (Python, JavaScript, Java, Go, C++, Rust). **CanItEdit** Cassano et al. (2024) is a dataset of 120 Python programs. The instructions take two forms: "lazy", based on human-authored commit messages, and "descriptive", written by an LLM. In CanItEdit, the LLM is expected to edit, on average, around 21% of the original code in terms of normalized edit distance between the original code snippet and the reference solution. In constrast, expected edits in HumanEvalFix are more constrained (5%) as the bugs are usually small and well-localized. The two datasets also differ in the distribution of ground truth edits. In HumanEvalFix, $|A| \approx |O|$, whereas CanItEdit's references are 20% longer than the original text, indicating prevalence of insertions.

Experiment Setup We obtain 3 LLM outputs for each item of each dataset, using 9 different models to multiply our sample size and the following prompt:

```
Edit the given code according to the instructions.
You MUST output the complete revised version of the code with your edits.
You MUST NOT add any comments. DO NOT explain you edits.
## Instructions
{instruction}
## Input code
{input_code}
## Edited code:
```

The LLMs used are: claude-sonnet-4 Anthropic (2025), Google's gemini-2.5-flash DeepMind (2025), OpenAI's gpt-4o-2024-11-20, gpt-4o-mini-2024-07-18, gpt-4.1-nano-2025-04-14 OpenAI (2025), Qwen2.5-Coder Instruct 1B and 5B Hui et al. (2024), DeepSeek Coder Instruct 1.3B and 6.7B Guo et al. (2024). For the Qwen and DeepSeek models, we use vLLM inference engine Kwon et al. (2023) and the default sampling parameters. For the remaining (proprietary) models, we set temperature to 0.2 and top_p to 0.95.

This results in 26568 ($3 \times 9 \times 984$) data samples derived from **HumanEvalFix** dataset and 2835 ($3 \times 9 \times 105$) derived from **CanItEdit**. For each LLM output, we execute the tests and record a binary pass (1) or fail (0) score. In HumanEvalFix, 45% of the generated solutions pass the test, while for CanItEdit dataset this number is 40%. Finally, we report Pearson correlation coefficient between the 0/1 indicator of passing the test and ES along with various other static measures computed on the (origin, reference, prediction) triples, namely exact match, unnormalized Levenshtein distance (ED), NES, chrF, BLEU, CodeBLEU, DiffBLEU, and SARI.

We experiment with 2 implementations of ExcisionScore—ES-Line and ES-Token—differing in the granularity of LCS. ES-Line excises the shared lines of code, while ES-Token tokenizes the code strings with tree-sitter and excises tokens common to all three strings. Additionally, we remove comments that do not affect execution, before passing A, B, O to each measure.

To illustrate what happens if our datasets contained a larger proportion of shared context, we artificially expand it by prepeding a long shared prefix of random length to each A, B, O, similar to Example 1. Since the measures considered are semantics-agnostic, the exact content of the prefix is irrelevant. The prefix is sampled uniformly from characters abcdef, a whitespace, and a newline character to contain a total of 2000–3000 characters. A different prefix is generated for each individual dataset

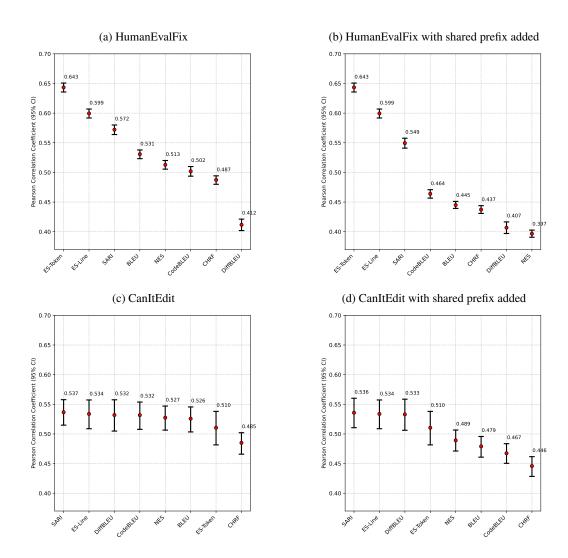


Figure 1: Correlation of various static measures with test execution. The first row refers to HumanEvalFix dataset, and the second to CanItEdit. Error bars indicate 95% bootstrap confidence intervals. The plots in the right column pertain to the experiment where we prepend a large random prefix to A, B, O. We excluded ED and exact match as their coefficients were low, as expected.

sample. We re-use the unperturbed pass/fail test execution data and compute the Pearson correlation coefficients. After these perturbations, the reference solution's coverage drops, on average, from 21% to 7% of the original code in CanItEdit and from 5% to only 0.5% in HumanEvalFix.

Results The resulting correlation coefficients are in Figure 1. On **HumanEvalFix** dataset, ES-Token takes the lead with a correlation coefficient of r=0.643 (CI [0.636,0.651]), followed by ES-Line r=0.599 (CI [0.592,0.607]), SARI r=0.572 (CI [0.564,0.58]), and others. Both ES-Token and ES-Line offer statistically significant improvement upon SARI, indicating that it is beneficial to remove shared context before applying SARI. When shared context dominates, the $F_{\rm keep}$ term of SARI is always maxed out to 1. Intuitively, that means that one of the 3 degrees of freedom ($F_{\rm keep}$, $F_{\rm add}$, $P_{\rm del}$) SARI has is permanently switched off, making SARI less sensitive. When a shared context is added to the HumanEvalFix dataset, SARI's correlation coefficient with pass@1 drops significantly: from 0.572 (CI [0.564,0.58]) to 0.549 (CI [0.541,0.558]). Extending shared context does not affect SARI's correlation with test execution on CanItEdit.

On the **CanItEdit** dataset, the differences in performance of different metrics, including pairwise ones, are insignificant. One possible explanation for that is the relative size of the edited region in

CanItEdit (21%, not including the unexpected edits the LLM solution makes). Besides, CanItEdit expects 20% more insertions than deletions. Since the inserted tokens appear in either A or B but not in O, the benefits of taking O into account are reduced.

Our results indicate that granularity of computing LCS or alignment is important. In HumanEvalFix, the edits are often small, changing only a few tokens within a line, explaining why ES-Token surpasses ES-Line on this dataset. On CanItEdit, however, ES-Token loses to ES-Line by a barely significant margin. Manual inspection reveals that overly fine-grained alignment can lead to meaningless unintuitive artifacts. Similarly, line-granular DiffBLEU falls short on HumanEvalFix, while performing well on CanItEdit.

Perturbing the data by adding a shared prefix does not affect ES and DiffBLEU scores, as ignoring this prefix was part of their design, neither does it affect unnormalized ED. In contrast, correlation coefficients of other pairwise measures with pass@1 on HumanEvalFix drop significantly: from [0.505, 0.52] to [0.439, 0.451] for BLEU, from [0.494, 0.51] to [0.457, 0.471] for CodeBLEU, from [0.48, 0.494] to [0.431, 0.444] for chrF, and from [0.505, 0.52] to [0.391, 0.403] for NES. We observe a similar effect on CanItEdit.

We critisized pairwise measures on the grounds that they reward dominating shared context as match, which reduces the effective range of values the similarity measure can take from [0,1] to [x,1], where x depends on how prevalent shared context is in the data. Some might object to this argument and suggest that dataset-specific re-normalization of the scores could be a trivial remedy. Namely, if s_i are the pairwise similarity scores on the i-th dataset sample, one could consider $s_i' \equiv \frac{s_i - \min s_i}{\max s_i - \min s_i}$, ensuring that s_i' fully cover the expected [0..1] range. However, our empirical results suggest that such a re-normalization still does not yield a satisfactory measure. Pearson's correlation coefficient r is invariant under linear transformations of the variables. Thus, the renormalized scores would have the same low r as the original values we observed. Our superiority argument based on correlation coefficients holds independently of the arguments about suitable and interpretable range of values.

5 RELATED WORK

Our adequacy criteria leverage global multiple sequence alignment (MSA, Definition 1), a technique well established in bioinformatics Chatzou et al. (2015).

Numerous measures exist to assess the similarity of two strings in different contexts. We argue that all of them are fundamentally ill-suited for the revision similarity problem, as the third string—the origin O—must be taken into account. Without it, pairwise measures cannot distinguish between revision similarity due to inheriting parts of O unchanged and that due to performing the same edits to O. As a result, pairwise measures reward shared context as matching (Section 2). Pairwise N-gram-based lexical measures include BLEU Papineni et al. (2002), CodeBLEU Ren et al. (2020), CrystalBLEU Eghbali & Pradel (2023), METEOR Banerjee & Lavie (2005), ROUGE Lin (2004), and chrF Popović (2015). BLEU has well-documented limitations, including its inability to address revision similarity — a gap we rigorously analyze in Section 2. For a systematic critique of BLEU's shortcomings (e.g., its insensitivity to paraphrasing and granular edits), we direct readers to Callison-Burch et al. (2006) and Reiter (2018). Despite these flaws, BLEU persists as a de facto standard due to its simplicity, reproducibility, and historical inertia.

Evaluating Grammatical Error Correction (GEC) techniques can be seen as an instance of the revision similarity problem. In GEC, edits are typically small and evaluation requires strict measures that are sensitive to word order. Alignment has been used to leverage these aspects in designing GEC metrics such as I-Measure Felice & Briscoe (2015) and M^2 (MaxMatch) Dahlmeier & Ng (2012). Evaluation of Text Simplification (TS) can, in some cases, also be viewed as revision similarity problem, provided that the simplifying changes do not rewrite the entire text. SARI Xu et al. (2016) (defined in Equation (2)) is an n-gram-based metric designed for the text simplification problem.

6 REPRODUCIBILITY STATEMENT

We attach the collected LLM responses and test execution results as Supplementary Materials. In Section 4 we specify how it was obtained. The code used to process this data

and compute the scores with their correlation coefficients is available anonymously under https://anonymous.4open.science/r/excision-score-eval-B9AF/.

REFERENCES

- Tom Adamczewski. How to run swe-bench verified in one hour on one machine. Blog post on epoch.ai, July 2025. URL https://epoch.ai/blog/swebench-docker. Accessed: 2025-09-05.
- Suha S. Al-Thanyyan and Aqil M. Azmi. Automated text simplification: A survey. ACM Computing Survey, 54(2), March 2021. ISSN 0360-0300. doi: 10.1145/3442695. URL https://doi.org/10.1145/3442695.
- Anthropic. System card: Claude opus 4 & claude sonnet 4. https://www-cdn.anthropic.com/4263b940cabb546aa0e3283f35b686f4f3b2ff47.pdf, May 2025.
- Ramakrishna Bairi, Atharv Sonwane, Aditya Kanade, Vageesh D. C., Arun Iyer, Suresh Parthasarathy, Sriram Rajamani, B. Ashok, and Shashank Shet. Codeplan: Repository-level coding using llms and planning. *Proc. ACM Softw. Eng.*, 1(FSE), July 2024. doi: 10.1145/3643757. URL https://doi.org/10.1145/3643757.
- Satanjeev Banerjee and Alon Lavie. METEOR: An automatic metric for MT evaluation with improved correlation with human judgments. In Jade Goldstein, Alon Lavie, Chin-Yew Lin, and Clare Voss (eds.), *Proceedings of the ACL Workshop on Intrinsic and Extrinsic Evaluation Measures for Machine Translation and/or Summarization*, pp. 65–72, Ann Arbor, Michigan, June 2005. Association for Computational Linguistics. URL https://aclanthology.org/W05-0909/.
- Christopher Bryant, Zheng Yuan, Muhammad Reza Qorib, Hannan Cao, Hwee Tou Ng, and Ted Briscoe. Grammatical error correction: A survey of the state of the art. *Computational Linguistics*, 49(3):643–701, 09 2023. ISSN 0891-2017. doi: 10.1162/coli_a_00478. URL https://doi.org/10.1162/coli_a_00478. section 6 (Evaluation).
- Chris Callison-Burch, Miles Osborne, and Philipp Koehn. Re-evaluating the role of Bleu in machine translation research. In Diana McCarthy and Shuly Wintner (eds.), *11th Conference of the European Chapter of the Association for Computational Linguistics*, pp. 249–256, Trento, Italy, April 2006. Association for Computational Linguistics. URL https://aclanthology.org/E06-1032/.
- Federico Cassano, Luisa Li, Akul Sethi, Noah Shinn, Abby Brennan-Jones, Jacob Ginesin, Edward Berman, George Chakhnashvili, Anton Lozhkov, Carolyn Jane Anderson, and Arjun Guha. Can it edit? evaluating the ability of large language models to follow code editing instructions, 2024. URL https://arxiv.org/abs/2312.12450.
- Maria Chatzou, Cedrik Magis, Jia-Ming Chang, Carsten Kemena, Giovanni Bussotti, Ionas Erb, and Cedric Notredame. Multiple sequence alignment modeling: methods and applications. *Briefings in Bioinformatics*, 17(6):1009–1023, 11 2015. ISSN 1467-5463. doi: 10.1093/bib/bbv099. URL https://doi.org/10.1093/bib/bbv099.
- Xinfang Chen, Siyang Xiao, Xianying Zhu, Junhong Xie, Ming Liang, Dajun Chen, Wei Jiang, Yong Li, and Peng Di. An efficient and adaptive next edit suggestion framework with zero human instructions in ides, 2025. URL https://arxiv.org/abs/2508.02473.
- Daniel Dahlmeier and Hwee Tou Ng. Better evaluation for grammatical error correction. In Eric Fosler-Lussier, Ellen Riloff, and Srinivas Bangalore (eds.), *Proceedings of the 2012 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pp. 568–572, Montréal, Canada, June 2012. Association for Computational Linguistics. URL https://aclanthology.org/N12-1067/.
- Google DeepMind. Gemini 2.5: Pushing the frontier with advanced reasoning, multimodality, long context, and next generation agentic capabilities, 2025. URL https://arxiv.org/abs/2507.06261.
- Edsger W. Dijkstra. The humble programmer. *Communications of the ACM*, 15(10):859–866, 1972. doi: 10.1145/355604.361591. The famous quote appears on p. 861. Originally presented as the 1972 ACM Turing Award Lecture.

Aryaz Eghbali and Michael Pradel. Crystalbleu: Precisely and efficiently measuring the similarity of code. In *Proceedings of the 37th IEEE/ACM International Conference on Automated Software Engineering*, ASE '22, New York, NY, USA, 2023. Association for Computing Machinery. ISBN 9781450394758. doi: 10.1145/3551349.3556903. URL https://doi.org/10.1145/3551349.3556903.

Mariano Felice and Ted Briscoe. Towards a standard evaluation method for grammatical error detection and correction. In Rada Mihalcea, Joyce Chai, and Anoop Sarkar (eds.), *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pp. 578–587, Denver, Colorado, May–June 2015. Association for Computational Linguistics. doi: 10.3115/v1/N15-1060. URL https://aclanthology.org/N15-1060/.

- Daya Guo, Qihao Zhu, Dejian Yang, Zhenda Xie, Kai Dong, Wentao Zhang, Guanting Chen, Xiao Bi, Y. Wu, Y. K. Li, Fuli Luo, Yingfei Xiong, and Wenfeng Liang. Deepseek-coder: When the large language model meets programming the rise of code intelligence, 2024. URL https://arxiv.org/abs/2401.14196.
- Dan Gusfield. *Algorithms on Strings, Trees, and Sequences: Computer Science and Computational Biology*. Cambridge University Press, 1997. sections 11 (two-sequence alignment) and 14 (multi-sequence alignment).
- Binyuan Hui, Jian Yang, Zeyu Cui, Jiaxi Yang, Dayiheng Liu, Lei Zhang, Tianyu Liu, Jiajun Zhang, Bowen Yu, Keming Lu, Kai Dang, Yang Fan, Yichang Zhang, An Yang, Rui Men, Fei Huang, Bo Zheng, Yibo Miao, Shanghaoran Quan, Yunlong Feng, Xingzhang Ren, Xuancheng Ren, Jingren Zhou, and Junyang Lin. Qwen2.5-coder technical report, 2024. URL https://arxiv.org/abs/2409.12186.
- Matt Kusner, Yu Sun, Nicholas Kolkin, and Kilian Weinberger. From word embeddings to document distances. In *Proceedings of the 32nd International Conference on Machine Learning (ICML)*, pp. 957–966, 2015.
- Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph E. Gonzalez, Hao Zhang, and Ion Stoica. Efficient memory management for large language model serving with pagedattention. In *Proceedings of the ACM SIGOPS 29th Symposium on Operating Systems Principles*, 2023.
- Chin-Yew Lin. ROUGE: A package for automatic evaluation of summaries. In *Text Summarization Branches Out*, pp. 74–81, Barcelona, Spain, July 2004. Association for Computational Linguistics. URL https://aclanthology.org/W04-1013/.
- Robert L. Logan IV, Alexandre Passos, Sameer Singh, and Ming-Wei Chang. Fruit: Faithfully reflecting updated information in text, 2022. URL https://arxiv.org/abs/2112.08634.
- Martin Monperrus. Living bibliography for automated program repair. https://program-repair.org/bibliography.html, 2023. URL https://program-repair.org/bibliography.html. Online; accessed September 25, 2025.
- Niklas Muennighoff, Qian Liu, Armel Zebaze, Qinkai Zheng, Binyuan Hui, Terry Yue Zhuo, Swayam Singh, Xiangru Tang, Leandro von Werra, and Shayne Longpre. Octopack: Instruction tuning code large language models. *arXiv preprint arXiv:2308.07124*, 2023.
- Karl Munson, Anish Savla, Chih-Kai Ting, Serenity Wade, Kiran Kate, and Kavitha Srinivas. Exploring code style transfer with neural networks, 2022. URL https://arxiv.org/abs/2209.06273.
- Graham Neubig and Xingyao Wang. Evaluation of llms as coding agents on swe-bench (at 30x speed!). *All Hands AI Blog*, October 2024. URL https://www.all-hands.dev/blog/evaluation-of-llms-as-coding-agents-on-swe-bench-at-30x-speed.
- OpenAI. OpenAI API Models Documentation. https://platform.openai.com/docs/models, September 2025. Accessed: September 10, 2025. Overview of available models in the OpenAI API, including capabilities and usage guidelines.

Sheena Panthaplackel, Pengyu Nie, Milos Gligoric, Junyi Jessy Li, and Raymond Mooney. Learning to update natural language comments based on code changes. In Dan Jurafsky, Joyce Chai, Natalie Schluter, and Joel Tetreault (eds.), *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pp. 1853–1868. Association for Computational Linguistics, July 2020. doi: 10.18653/v1/2020.acl-main.168. URL https://aclanthology.org/2020.acl-main.168/.

- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pp. 311–318, 2002.
- Dorin Pomian, Abhiram Bellur, Malinda Dilhara, Zarina Kurbatova, Egor Bogomolov, Timofey Bryksin, and Danny Dig. Next-generation refactoring: Combining llm insights and ide capabilities for extract method. In 2024 IEEE International Conference on Software Maintenance and Evolution (ICSME), pp. 275–287, 2024. doi: 10.1109/ICSME58944.2024.00034.
- Maja Popović. chrF: character n-gram F-score for automatic MT evaluation. In Ondřej Bojar, Rajan Chatterjee, Christian Federmann, Barry Haddow, Chris Hokamp, Matthias Huck, Varvara Logacheva, and Pavel Pecina (eds.), *Proceedings of the Tenth Workshop on Statistical Machine Translation*, pp. 392–395, Lisbon, Portugal, September 2015. Association for Computational Linguistics. doi: 10.18653/v1/W15-3049. URL https://aclanthology.org/W15-3049/.
- Ehud Reiter. A structured review of the validity of BLEU. *Computational Linguistics*, 44(3):393–401, September 2018. doi: 10.1162/coli_a_00322. URL https://aclanthology.org/J18-3002/.
- Shuo Ren, Daya Guo, Shuai Lu, Long Zhou, Shujie Liu, Duyu Tang, Neel Sundaresan, Ming Zhou, Ambrosio Blanco, and Shuai Ma. Codebleu: a method for automatic evaluation of code synthesis, 2020. URL https://arxiv.org/abs/2009.10297.
- Akhilesh Sudhakar, Bhargav Upadhyay, and Arjun Maheswaran. "transforming" delete, retrieve, generate approach for controlled text style transfer. In Kentaro Inui, Jing Jiang, Vincent Ng, and Xiaojun Wan (eds.), *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pp. 3269–3279, Hong Kong, China, November 2019. Association for Computational Linguistics. doi: 10.18653/v1/D19-1322. URL https://aclanthology.org/D19-1322/.
- The IEEE and The Open Group. diff compare two text files. https://pubs.opengroup.org/onlinepubs/9699919799/utilities/diff.html, 2018. IEEE Std 1003.1-2017 (Revision of IEEE Std 1003.1-2008).
- Elaine J. Weyuker. On testing non-testable programs. *The Computer Journal*, 25(4):465–470, 1982. doi: 10.1093/comjnl/25.4.465.
- Wei Xu, Courtney Napoles, Ellie Pavlick, Quanze Chen, and Chris Callison-Burch. Optimizing statistical machine translation for text simplification. *Transactions of the Association for Computational Linguistics*, 4:401–415, 2016. doi: 10.1162/tacl_a_00107. URL https://aclanthology.org/Q16-1029/.
- Tianyi Zhang*, Varsha Kishore*, Felix Wu*, Kilian Q. Weinberger, and Yoav Artzi. Bertscore: Evaluating text generation with bert. In *International Conference on Learning Representations*, 2020. URL https://openreview.net/forum?id=SkeHuCVFDr.