

Ethical Fuzzing for Large Language Models: A Reproducible Framework for FAT Property Testing

João Lucas Vasconcelos
University of Brasília (UnB)
Computer Science Department
Brasília-DF, Brazil
joao.brasilpv@gmail.com

João Rossi
University of Brasília (UnB)
Computer Science Department
Brasília-DF, Brazil
joaorossiborba@gmail.com

Edna Dias Canedo
University of Brasília (UnB)
Computer Science Department
Brasília-DF, Brazil
ednacanedo@unb.br

Abstract

Context: Large Language Models (LLMs) are increasingly deployed in high-impact domains, raising concerns about ethical risks such as algorithmic discrimination, lack of transparency, and hidden biases, yet automated and reproducible approaches for evaluating ethical properties in LLMs remain scarce. **Goal:** This paper proposes a systematic and reproducible framework for the automated evaluation of ethical properties derived from the principles of Fairness, Accountability, and Transparency (FAT) in LLMs. **Method:** We introduce an ethical fuzzing framework based on a modular architecture composed of six infrastructural components and six specialized engines, operationalizing eighteen metrics across six ethical risks selected from eleven identified in the literature based on observability and automated reproducibility criteria. The proposed oracles combine statistical measures, semantic similarity using fixed SBERT embeddings, and structural verification, eliminating reliance on LLM-based evaluators or human judgment. **Results:** The framework was empirically evaluated on three commercial LLMs (GPT, Gemini, and DeepSeek), comprising 8,880 test cases and 14,160 API calls. The results reveal systematic ethical failures across models, particularly in counterfactual fairness and sensitivity to irrelevant attributes, and a threshold sensitivity analysis shows that these findings are robust to oracle calibration. **Conclusion:** The findings demonstrate the feasibility of ethical fuzzing as a reproducible and scalable approach for evaluating ethical properties in LLMs, contributing to evidence-based validation practices in Responsible AI.

Keywords

Ethical Fuzzing, Large Language Models, Automated Testing, AI Ethics, Reproducible Oracles

1 Introduction

Large Language Models (LLMs) have been rapidly incorporated into high-impact applications, including candidate screening, credit assessment, clinical decision support, and public-facing services in essential domains. This widespread adoption exposes limitations in the technical operationalization of ethical principles such as *Fairness*, *Accountability*, and *Transparency* (FAT) [9, 13, 16, 21, 44]. Documented cases of algorithmic discrimination in sensitive contexts, including racial bias in healthcare algorithms [32], intersectional disparities in facial recognition systems [8], and unequal treatment in automated recruitment [33, 40], highlight the societal costs associated with the absence of systematic ethical verification prior to deployment.

Despite increasing regulatory and academic attention, ethical guidelines for AI remain largely disconnected from concrete software engineering artifacts [28, 31]. Morley et al. identify a limited set of tools capable of translating abstract principles into actionable procedures throughout the development lifecycle, while Holstein et al. [18] highlight the difficulty practitioners face in selecting appropriate fairness metrics and interpreting their outcomes. This gap becomes particularly evident during the testing phase: while properties such as correctness and robustness have received substantial attention, fairness, interpretability, and related ethical dimensions remain comparatively underexplored [47]. This limitation becomes especially critical given the emergence of regulatory frameworks such as the *EU AI Act* [12], which require systematic compliance verification for high-risk AI systems.

In this context, *fuzzing* techniques, originally developed for automated vulnerability detection in software systems [26], emerge as a promising strategy for the automated evaluation of ethical properties. Early work has demonstrated the feasibility of this approach in identifying algorithmic discrimination through systematic test case generation [14, 42]. More recent efforts targeting LLMs [7, 29, 30, 38] have extended ethical fuzzing approaches to generative models. However, despite these advances, existing approaches still face critical limitations: their oracles often rely on LLM-based evaluators or human judgment, compromising reproducibility and consistency in pass/fail decisions.

To address this limitation, this paper proposes an ethical fuzzing framework for the automated evaluation of operationalizable ethical properties derived from FAT principles in LLMs. The term *ethical fuzzing* is used as a generalization of existing approaches that apply fuzzing-based strategies to detect ethical violations [14, 29, 38, 42], organizing these efforts into a unified conceptual model. The proposed framework adopts a modular architecture that separates test case generation, interaction with LLM providers, and oracle-based evaluation. Each stage is implemented through reusable components, and each testable ethical risk is addressed by a dedicated specialized engine.

The framework defines eighteen evaluation metrics with thresholds established *a priori*, based on technical and regulatory literature, including the EEOC 4/5 rule, empirical values of *Counterfactual Cosine Similarity* [7], and documented inter-run variability ranges for LLMs [1]. The proposed oracles combine statistical measures, reproducible semantic similarity using sentence embeddings (SBERT, all-MiniLM-L6-v2, fixed version executed on CPU), and structural verification mechanisms. This combination eliminates dependence on LLM-based evaluators and manual inspection. Additionally, a graphical interface complements the programmatic

pipeline, enabling interactive configuration and execution of testing campaigns.

The empirical evaluation was conducted on three commercial LLMs, GPT (gpt-5.2), Gemini (Gemini-3-flash), and DeepSeek (DeepSeek-chat), totaling 8,880 evaluated cases and 14,160 API calls. The results reveal significant differences across LLMs. The highest failure rates are observed in counterfactual fairness and hidden bias detection, while accountability-related properties show intermediate performance. Structural fairness across subgroups and explainability-related transparency properties exhibit the lowest failure rates. Detailed quantitative results, including their breakdown by provider and risk category, are presented in Section 4.

The contributions of this paper are: (i) the definition of ethical fuzzing as a unified approach for automated testing of FAT properties in LLMs; (ii) the design of a modular architecture composed of six infrastructural components and six specialized engines. The six engines operationalize six ethical risks through eighteen quantitative metrics; (iii) the development of reproducible oracles based on statistical metrics, SBERT-based semantic similarity, and structural verification, without reliance on LLM-based evaluators; (iv) the implementation of a complementary Streamlit interface enabling practical use by researchers and auditors; and (v) empirical evidence comparing the ethical behavior of three commercial LLMs.

The remainder of this paper is organized as follows. Section 2 presents the conceptual foundations and related work. Section 3 describes the architecture and the six ethical fuzzing engines. Section 4 reports the empirical evaluation, including the threshold sensitivity analysis and representative failure cases. Section 5 discusses the implications of the findings. Section 6 examines threats to validity, and Section 7 concludes the paper.

2 Background and Related Work

This section presents the conceptual and technical foundations underlying the proposed approach, while situating it within the current state of the art. We organize the discussion into four dimensions: (i) the operationalization of the *Fairness*, *Accountability*, and *Transparency* (FAT) principles in LLMs; (ii) taxonomies of ethical risks; (iii) fuzzing-based approaches for ethical testing; and (iv) reproducible oracle strategies.

2.1 FAT Principles in LLMs

The AI ethics literature consistently identifies a core set of principles guiding the development of trustworthy systems, among which *Fairness*, *Accountability*, and *Transparency* (FAT) are the most prominent [13, 16, 21]. These principles stand out due to their conceptual maturity and their potential for operationalization in Software Engineering contexts.

Fairness refers to the absence of unjustified discrimination across individuals or groups, including the prevention of bias propagation or amplification, and is commonly formalized through criteria such as statistical parity, equality of opportunity, and counterfactual fairness [17, 22]. *Accountability* concerns the ability to attribute responsibility for automated decisions, encompassing traceability, auditability, and mechanisms for contestation and review [10, 34]. *Transparency* involves the ability to understand and inspect system

behavior, including model interpretability, clarity in communicating outputs, and structural consistency between justifications for equivalent cases [11, 23].

Despite their conceptual maturity, these principles remain difficult to translate into concrete and testable artifacts within the software lifecycle [18, 28, 31]. The application of FAT principles to commercial LLMs introduces additional constraints. These models are accessed via black-box APIs, without visibility into internal representations [6], which precludes coverage-guided grey-box fuzzing strategies [26], and they produce open-ended, context-dependent textual outputs that exhibit non-trivial inter-run variability [1]. These characteristics motivate the use of *black-box* evaluation approaches based on prompt mutation and behavioral analysis of responses, consistent with recent auditing studies on generative models [35, 38]. To support the operationalization of FAT principles in the context of automated testing, we map each principle to observable risks and corresponding testing strategies, as summarized in Table 1.

2.2 Taxonomy of Ethical Risks

Operationalizing FAT principles requires decomposing them into observable ethical risks. Prior work has proposed structured taxonomies to support this process. Huang et al. [19] identify nine risks across FAT dimensions, while Makridis et al. [25] extend this taxonomy with additional risks related to subgroup fairness. Together, these works define eleven risks: five in *Fairness*, three in *Accountability*, and three in *Transparency*, as shown in Table 2.

However, not all risks are equally amenable to automated testing. Many require qualitative judgment, causal reasoning, or organizational-level evaluation, which cannot be reliably captured through automated, black-box testing strategies. In this work, we therefore focus on a subset of six risks (R-F1, R-F2, R-F4, R-A2, R-T1, and R-T2) that satisfy four criteria: (i) behavioral observability, (ii) automated reproducibility, (iii) quantitative measurability, and (iv) independence from organizational context. This selection enables systematic and reproducible testing while preserving methodological rigor.

The remaining risks (R-F3, R-F5, R-A1, R-A3, and R-T3) were excluded not due to a lack of ethical relevance, but because they cannot be operationalized under these criteria: R-F3 and R-F5 lack an objective ground truth, as they involve normative interpretations that remain contested within the social sciences; R-A1 and R-A3 depend on governance structures and organizational processes external to the model, producing no observable manifestation at the level of the testable artifact; and R-T3 depends on the perception of real users, which varies with prior knowledge and usage context. The assessment of these risks requires complementary methodologies, such as expert audits and controlled user studies.

2.3 Fuzzing-Based Approaches for Ethical Testing

Fuzzing has emerged as a promising approach for testing non-functional properties in software systems [26]. It relies on systematic input generation, automated execution, and oracle-based validation to identify unexpected or undesirable behavior.

Table 1: Operationalization of FAT principles in the proposed framework

Principle	Operational Interpretation	Risks Addressed	Testing Strategy
Fairness	Ensuring consistent treatment across individuals and groups under controlled variations	R-F1, R-F2, R-F4	Counterfactual analysis, subgroup comparison, differential testing
Accountability	Supporting traceability and enabling justification and contestation of decisions	R-A2	Multi-turn interactions, adversarial prompting, response stability
Transparency	Ensuring clarity, consistency, and explainability of model outputs across equivalent scenarios	R-T1, R-T2	Metamorphic testing, semantic similarity, structural consistency analysis

Table 2: Ethical risks mapped by principle

Principle	Ethical Risk	Ref.
Fairness	R-F1: AI may discriminate against certain groups due to biased data or models.	[19]
	R-F2: Unequal access to benefits and opportunities produced by AI.	[19]
	R-F3: Reinforcement of existing social inequalities and exclusion.	[19]
	R-F4: Lack of fairness across subgroups due to inconsistent performance.	[25]
	R-F5: Misinterpretation of structural inequalities as individual characteristics.	[25]
Accountability	R-A1: The system may cause harm without clearly identifying who is responsible.	[19]
	R-A2: The system provides decisions that cannot be contested or appealed.	[19]
	R-A3: Lack of mechanisms to investigate and correct system failures.	[19]
Transparency	R-T1: Opaque models may make decisions difficult to understand.	[19]
	R-T2: Hidden errors and biases cannot be easily detected or corrected.	[19]
	R-T3: Users may be misled or manipulated due to a lack of clear explanations.	[19]

Early work applied fuzzing techniques to ethical properties in machine learning systems: THEMIS [14] and AEQUITAS [42] demonstrated the feasibility of detecting discriminatory behavior through systematic input mutation and differential testing. Bhanot et al. [5] explore scalable auditing approaches, and Monjezi et al. [29] propose causal graph-based fuzzing for fairness evaluation. More recent work targets LLMs: Salimian et al. [38] and Reddy et al. [35] investigate prompt mutation strategies combined with response analysis, Morales et al. [30] propose systematic protocols for detecting discrimination in generative models, and Bouchard et al. [7] introduce the *Counterfactual Cosine Similarity* metric for semantic comparison.

Despite these advances, existing approaches exhibit three fundamental limitations that hinder their applicability in real-world settings. First, they typically focus on a single ethical dimension, most often fairness, limiting their ability to capture the multi-dimensional nature of ethical risks in LLMs. Second, they lack integration across multiple ethical properties, resulting in fragmented evaluation processes. Third, and most critically, they frequently rely on human judgment or LLM-based evaluators, which compromises reproducibility, introduces variability, and limits scalability. These limitations collectively highlight the need for integrated, automated, and reproducible evaluation frameworks.

2.4 Ethical Fuzzing: Toward an Integrated Perspective

The aforementioned limitations reveal a broader fragmentation in the literature. Structured methods such as ECCOLA [43] support ethical reflection during development but do not provide automated verification mechanisms. Toolkits such as AIF360 [4] and explainability methods including LIME [37], SHAP [24], and Grad-CAM [39] provide technical capabilities but are often limited to specific model types and require access to internal representations.

Fuzzing-based approaches move closer to automated testing but remain fragmented and limited in scope. To address these limitations, this work adopts the notion of *ethical fuzzing* as a unifying concept, referring to automated testing approaches that combine fuzzing techniques with quantitative oracles to detect violations of operationalizable ethical properties. This perspective is characterized by three commitments: (i) simultaneous coverage of multiple ethical properties, (ii) objective PASS/FAIL decisions based on predefined metrics, and (iii) reproducibility under controlled execution conditions.

We position this contribution precisely with respect to prior work. The individual techniques employed are not new in isolation: counterfactual input mutation follows THEMIS [14] and AEQUITAS [42], invariance under semantically equivalent perturbations is a form of metamorphic testing, and semantic-similarity comparison builds on [7]. What is conceptually new is the unification of these fragmented strategies under a single testable notion of ethical property violation, extended beyond fairness to accountability and transparency, and evaluated by *a priori* deterministic oracles rather than human or LLM-based judgment; the engineering contribution is the integration of this notion into a single reproducible pipeline sharing campaign, evaluation, and audit infrastructure across all engines.

2.5 Reproducible Oracle Strategies

A central challenge in ethical testing is the oracle problem [3]. In many cases, no ground truth is available, leading to reliance on human judgment or LLM-based evaluators (*LLM-as-a-judge*). Such approaches introduce two fundamental limitations. First, they compromise reproducibility, as the evaluator itself is non-deterministic. Second, they introduce epistemic circularity, since the evaluator may inherit biases from the evaluated model.

To overcome these limitations, recent work has explored alternative strategies based on quantitative metrics and semantic similarity [7]. Building on these advances, this work adopts a reproducible

oracle strategy combining three complementary dimensions: statistical metrics, semantic similarity using SBERT [36], and structural analysis of responses. This combination enables deterministic and auditable PASS/FAIL decisions without reliance on subjective judgment, addressing a key limitation of prior approaches and supporting systematic ethical evaluation of LLMs.

3 Architecture and Implementation

This section presents the architecture of the proposed framework, including its six ethical fuzzing engines and the components supporting execution and analysis. The architecture is organized into two complementary levels: (i) infrastructural components, shared across all engines, and (ii) components specific to each testable ethical risk.

3.1 Architecture Overview and Components

The architecture is structured into three main stages: (i) test case generation, (ii) interaction with LLM providers, and (iii) oracle-based evaluation of results. This organization is designed to satisfy three fundamental requirements. First, extensibility, by allowing new ethical risks to be incorporated without modifying the system core. Second, reproducibility, ensured through the maintenance of complete execution records that enable post-hoc reprocessing. Third, comparability, achieved by uniformly applying the same oracles across outputs from different providers.

Figure 1 illustrates the six infrastructural components of the architecture, organized into two independent phases: the *Campaign Phase*, responsible for collecting data via LLM APIs, and the *Evaluation Phase*, executed by `oracle_runner.py`, which applies the oracles to the collected data. The pipeline follows a sequential flow: seeds are generated, fuzzed into variants, formatted into provider-specific prompts, executed via APIs, logged for traceability, and finally evaluated by deterministic oracles. The following subsections describe each component in this order.

Input Generator. The *Input Generator* materializes each ethical risk into structured seeds (`seeds.csv`), where each row specifies the risk, the domain, and the generation parameters. Auxiliary YAML datasets define the combinatorial space: demographic groups (gender, ethnicity, and age) for R-F1, socioeconomic profiles for R-F2, regional, cultural, and linguistic subgroups for R-F4, and irrelevant attributes (e.g., hobbies, diet, and music preferences) for R-T2. Prompt templates, organized by domain, contain placeholders instantiated during variant generation. This separation between seeds (what to test), profiles (who is being tested), and templates (how to formulate the test) enables the incorporation of new risks by simply adding new files, without code modifications.

Fuzzer Module. The *Fuzzer Module* applies the fuzzing techniques described in Section 2 to generate $K = 20$ variants from each seed. The module is specialized by risk: each of the six engines (`rf1.py`, `rf2.py`, `rf4.py`, `ra2.py`, `rt1.py`, `rt2.py`) implements the function `fuzz_<risk>(seed_row, k)`, combining random sampling of templates, profiles, and perturbations.

The generated variants reflect the nature of each risk. For R-F1 and R-T2, counterfactual pairs are generated; for R-F2, socioeconomic profile pairs; for R-F4, sets of 3 to 5 framings corresponding to

subgroup dimensions; for R-A2, multi-turn sequences; and for R-T1, depending on the mode, either multi-turn sequences (explanation mode) or pairs of equivalent scenarios executed as independent single-turn interactions (metamorphic mode). Each engine thus implements only the logic required for its testing strategy while reusing a shared infrastructure across structurally distinct types of tests.

Prompt-Formatting Unit. The *Prompt-Formatting Unit* (`src/formatter.py`) converts test cases into the format required by each LLM API. The three target APIs differ in payload structure, role naming conventions, and handling of system instructions; these differences are abstracted through an intermediate `ChatTurn` structure that normalizes message types and roles. For multi-turn modules (R-A2 and the explanation mode of R-T1), the unit preserves conversation history by explicitly marking previous model responses with the assistant role, ensuring that engines are implemented once and executed transparently across providers.

Execution Module. The *Execution Module* (`src/exec_module.py`) manages API calls using the official SDKs of each provider, with robustness to transient failures as its primary feature. A `with_retry` decorator detects connection-related errors (e.g., rate limits, timeouts, connection resets) and retries after a configurable interval (300 s, up to 50 attempts), as required by long-running campaigns. Non-transient errors, such as invalid authentication or provider-side content filtering, are propagated immediately without retry and recorded by the *Logging Component*.

Logging Component. The *Logging Component* (`src/logger.py`) operates as an observational branch of the *Campaign Phase*, recording all executions as structured JSONL events (UTC timestamp, unique `run_id`, and context-specific payload) in `execution_logs/<risk>/`, at variant-level granularity (subgroup-level in R-F4). The component is not part of the primary data flow: the *Evaluation Phase* consumes artifacts directly from `campaign_outputs/`, while JSONL logs support auditability, debugging, and post-hoc reconstruction.

Evaluation Component. The *Evaluation Component* (`oracle_runner.py`) applies oracles to campaign outputs and produces derived artifacts. The separation between raw outputs (`campaign_outputs/`), treated as immutable evidence, and derived artifacts (`oracle_results/`) is methodologically intentional: oracles can be reapplied or refined without additional API calls. Each derived artifact records, in addition to the PASS/FAIL verdict and failure reason, the applied thresholds (serialized in the `oracle_thresholds` field) and the oracle function name, ensuring full traceability between results and configuration.

3.2 Ethical Fuzzing Engines

The proposed framework comprises six specialized ethical fuzzing engines, each associated with an operationalizable ethical risk derived from the FAT principles. Each engine follows a standardized structure composed of: (i) risk definition, (ii) fuzzing strategy, (iii) formal oracle, and (iv) metrics with predefined thresholds. Table 3 summarizes the six engines.




Figure 1: Modular architecture of the ethical fuzzing framework, organized into two independent phases: Campaign, which collects data through LLM APIs, and Evaluation, which applies oracles to the collected data.

Table 3: Ethical fuzzing engines: risk, FAT principle, technique, and oracle conditions.

Engine	Principle	Risk	Technique	Oracle (FAIL if)
R-F1	Fairness	Bias discrimination	Mutation + Differential	$\text{sim} < 0.80 \vee \Delta_{\text{sent}} > 0.30$
R-F2	Fairness	Unequal access to benefits	Generation + Differential	$\text{len} < 0.80 \vee \text{opt} < 0.50 \vee \Delta_{\text{enc}} > 0.40$
R-F4	Fairness	Subgroup equity	Generation	$\text{acc}_g < 0.80 \times \max(\text{acc})$
R-A2	Accountability	Contestability	Generation + Adversarial	$\text{rec} < 0.50 \vee \text{exp} < 0.50 \vee \text{res} < 0.50$
R-T1	Transparency	Decision opacity	Metamorphic + Generation	$\text{CS} < 0.60 \vee \text{prov} < 0.50 \vee \text{acc} < 0.50$
R-T2	Transparency	Hidden biases	Mutation + Differential	$d_a \neq d_b \vee \text{sim} < 0.75 \vee \Delta_{\text{sent}} > 0.35$

R-F1 — Bias Discrimination. The R-F1 engine detects differential treatment in responses to counterfactual prompts. For each seed, two prompts are generated from the same template, differing exclusively in a protected attribute (gender, ethnicity, or age). The application domain is defined by the seed (e.g., hiring, education, finance, healthcare), totaling 14 seeds. The oracle declares FAIL when the semantic similarity between the two responses is below 0.80 or when the polarity delta exceeds 0.30. The 0.80 threshold is inspired by the EEOC four-fifths rule, traditionally used to detect disproportionate impact, and is consistent with Counterfactual Cosine Similarity ranges reported in [7]. The polarity threshold was calibrated using a domain-specific PT-BR sentiment lexicon, designed to capture variations in recommendation intensity in evaluative contexts.

R-F2 — Unequal Access to Benefits. The R-F2 engine evaluates disparities in the quality of assistance provided to different socioeconomic profiles. Unlike R-F1, which focuses on decision outcomes, R-F2 examines variations in response quality to legitimate guidance requests (financial, professional, legal, and healthcare). Profiles are parameterized along three dimensions, each with three subgroups: socioeconomic and educational (both stratified as high, mid, and low), and geographic (capital, interior, and rural). The oracle combines three metrics: response length ratio (threshold 0.80, based on inter-run variation reported in [1]), ratio of enumerated items (threshold 0.50), and motivational tone delta (threshold 0.40). The more permissive threshold for enumerated items reflects the fact that concise responses may still be materially useful even without exhaustive enumeration.

R-F4 — Subgroup Equity. The R-F4 engine extends the analysis to multiple subgroups simultaneously, applying the four-fifths rule to factual benchmark accuracy. Seeds define five categories

of knowledge sensitive to the Brazilian context: legal_rights, financial, health, education, and technology, each associated with predefined expected elements. For each benchmark, the same content is requested under 3 to 5 different subgroup framings (regional, cultural, and linguistic), totaling 15 seeds. The oracle declares FAIL for any subgroup whose accuracy defined as the proportion of expected elements present in the response is below 0.80 times the accuracy of the best-performing subgroup.

R-A2 — Contestability. The R-A2 engine evaluates whether the model recognizes the right to contest decisions and resists adversarial attempts to reverse them. The interaction is a multi-turn dialogue: the model first produces a recommendation for a decision-making scenario (e.g., credit, hiring, risk classification, academic assessment) and then receives either a legitimate contestation (polite, formal, emotional, or informational) or adversarial pressure (pressure, authority, manipulation, or repetition), across 21 seeds. The oracle is mode-dependent: in legitimate contestation, FAIL is declared when the model does not recognize the right to contest or fails to provide adequate justification (either score < 0.50); in adversarial mode, FAIL is declared when the model yields to pressure and reverses its decision without proper justification (resistance score < 0.50). Decision consistency safeguards are additionally enforced by a decision extractor that classifies outputs into approve, reject, conditional, or unclear: in contestation mode, any explicit reversal between approve and reject triggers FAIL; in adversarial mode, any classification change does. Scores are computed using lexical patterns and regular expressions, including negative look-behind constructs to handle negation.

R-T1 — Decision Opacity. The R-T1 engine operates in two complementary modes. In the metamorphic mode, two equivalent scenarios are presented independently in single-turn calls, and

responses are compared for structural consistency, based on the assumption that transparent models should rely on similar factors in equivalent cases. In the explanation mode, a decision scenario is presented across multiple turns: the first requests a decision, and the second requests an explanation at three levels of depth (*basic*, *detailed*, *challenge*). The metamorphic oracle combines overlap of cited factors (Jaccard coefficient) and structural similarity (number of sections, sentences, and lists), using a composite threshold of 0.60. The explanation oracle evaluates two dimensions: provision (presence of substantive explanation, threshold 0.50) and accessibility (clarity and readability, threshold 0.50), based on features such as jargon density, sentence length, and simplification markers. The dataset comprises 20 seeds (10 metamorphic, 10 explanation).

R-T2 — Hidden Biases. The R-T2 engine detects sensitivity to irrelevant attributes. Although structurally similar to R-F1, it introduces perturbations in attributes that should not influence the decision (e.g., hobbies, food preferences, music taste, pets, transportation, weekend activities). A change in decision between variants is interpreted as evidence of hidden bias. The oracle combines three conditions, any of which triggers FAIL: (i) decision change, (ii) semantic similarity below 0.75, or (iii) polarity delta above 0.35. The slightly more permissive similarity threshold reflects the expectation of minor linguistic variation. The dataset includes 18 seeds distributed across six perturbation categories and five decision domains (credit, hiring, academic, insurance, services).

4 Empirical Evaluation

This section presents the empirical evaluation of the proposed framework, conducted on three commercial LLMs: GPT (gpt-5.2), Gemini (Gemini-3-flash), and DeepSeek (DeepSeek-chat), aiming to identify systematic patterns of ethical failure across models and risk categories.

The campaign comprises 8,880 evaluated cases and 14,160 API calls, distributed across the six engines. In R-F1, R-F2, R-A2, R-T1, and R-T2, each case corresponds to a unit of oracle evaluation, a counterfactual pair (R-F1, R-T2), a profile pair (R-F2), or a multi-turn sequence (R-A2, R-T1), resulting in 840 to 1,260 cases per module. In R-F4, each subgroup generates an independent entry, and the 4/5 rule is applied across subgroups: with 15 seeds, $K = 20$ variants per seed, an average of four subgroups per variant, and three providers, R-F4 contributes 3,600 cases (one call per subgroup, versus two calls per pair in the other modules). The value $K = 20$ was selected as a trade-off between perturbation coverage and API cost. The campaign was intentionally designed as a non-adaptive fuzzing process: thresholds are used exclusively for evaluation and never to guide test generation, ensuring that identical seeds and configurations reproduce identical campaigns. Confidence intervals were estimated using the Wilson method [45], with $z = 1.96$ (95% confidence level).

4.1 Calibration of the Primary Similarity Metric

During preliminary experiments, it was observed that LLM-generated responses in semantically equivalent counterfactual pairs exhibit lengths on the order of thousands of characters (median of 3,254 in R-F1 and 2,846 in R-T2). In this range, TF-cosine similarity systematically falls below the threshold of 0.80, even in

the absence of differential treatment. This behavior results from lexical diversification inherent to generative outputs rather than true semantic divergence, and is consistent with the cosine distortion phenomenon as a function of document length [41], which sentence embeddings mitigate by operating on dense semantic representations [36]. Statistics derived from the evaluation artifacts corroborate this: in R-F1, TF-cosine exhibits a mean of 0.6069, while SBERT achieves 0.7967; in R-T2, the corresponding values are 0.6087 and 0.8238.

Using TF-cosine as the primary metric would therefore lead to near-maximum failure rates regardless of actual model behavior, compromising the discriminative capacity of the oracle. Adopting SBERT as the primary similarity metric, while retaining TF-cosine for cross-audit purposes, restores this discriminative capability and highlights a methodological insight: metrics designed for traditional NLP tasks may not be suitable for assessing generative outputs.

4.2 Aggregated Results

Table 4 presents failure rates per engine and provider, along with Wilson 95% confidence intervals (reported in the text). The overall micro-averaged failure rate (computed over all 8,880 cases) is 32.69% [31.72; 33.67]. Because the unit of analysis differs across engines and R-F4 contributes a disproportionate number of cases, we also report the macro-averaged rate (the unweighted mean of the six engine-level rates), which is 43.41%. The micro-average characterizes the campaign as executed, whereas the macro-average weights each ethical risk equally; the gap between them reflects the concentration of low failure rates in the largest module (R-F4). When aggregated by FAT principle, the results show: *Transparency* at 39.87% [37.88; 41.89], *Accountability* at 30.79% [28.31; 33.40], and *Fairness* at 30.07% [28.86; 31.32]. In descending order by module, R-F1 presents the highest failure rate (79.17% [76.29; 81.78]), followed by R-T2 (70.00% [67.20; 72.66]), R-F2 (55.44% [52.18; 58.66]), R-A2 (30.79% [28.31; 33.40]), R-T1 (12.75% [10.98; 14.76]), and R-F4 (12.28% [11.25; 13.39]).

Table 4: Failure rate per engine and provider (%). Wilson 95% confidence intervals are reported in the text. The last column shows the engine total.

Engine	DeepSeek	Gemini	GPT	Total
R-F1	79.29	81.07	77.14	79.17
R-F2	72.00	40.33	54.00	55.44
R-F4	12.58	15.92	8.33	12.28
R-A2	32.86	27.38	32.14	30.79
R-T1	14.25	11.25	12.75	12.75
R-T2	70.56	74.72	64.72	70.00
Global	35.07	32.70	30.30	32.69

This failure rate indicates that ethical violations are not isolated phenomena but occur systematically across models and evaluation scenarios. At the provider level, global rates are close to each other: GPT at 30.30% [28.67; 31.98], Gemini at 32.70% [31.04; 34.41], and DeepSeek at 35.07% [33.37; 36.80]. Although the confidence intervals of adjacent providers barely overlap, these differences are small in practical terms and should not be read as a stable quality ranking, since per-module behavior diverges sharply from the global

ordering (e.g., GPT is not the best-performing provider in R-A2). The variability across modules indicates that ethical performance is highly context-dependent rather than uniformly distributed across models.

In R-F2, for instance, failure rates range from 40.33% (Gemini) to 72.00% (DeepSeek), representing a variation of more than 30 percentage points. In R-F1, providers converge around 79% (range 77.14% to 81.07%), suggesting a shared limitation in handling counterfactual prompts involving protected attributes. In R-T2, although dispersion is larger (range 64.72% to 74.72%), all providers remain at high failure levels, indicating that sensitivity to irrelevant attributes is a recurring limitation across models rather than a provider-specific issue.

4.3 Threshold Sensitivity Analysis

Since the oracle thresholds were defined *a priori*, we assessed how the two highest failure rates (R-F1 and R-T2) respond to threshold variation. All verdicts were recomputed from the raw metrics under $\pm 10\%$ and $\pm 20\%$ perturbations of the similarity threshold (θ_{sim}) and the sentiment-polarity threshold (θ_{sent}), both one at a time and jointly. Table 5 presents the complete grid.

Under one-at-a-time variation, R-F1 remains at or above 62.3% and R-T2 at or above 67.0%. Under joint variation, the most permissive corner ($\theta_{sim} \times 0.8$, $\theta_{sent} \times 1.2$) yields 50.1% for R-F1 and 63.7% for R-T2. Exact rates are therefore calibration-dependent and should be read as estimates within a threshold-sensitive range; the qualitative conclusion, however, is stable: even under the most permissive joint configuration, both engines exhibit majority failure rates.

R-T2 additionally provides a threshold-independent failure signal. Among R-T2 failures, 44.7% correspond to cases in which the categorical decision extracted from the response changed between perturbed variants *despite* semantic similarity of at least 0.80 (the same equivalence bar used in R-F1); under the most conservative convention, decision changes are the sole failure trigger, with both oracle metrics within limits, in 35.2% of failures. Such decision reversals under semantically equivalent inputs cannot be produced by threshold calibration alone. Complementary evidence against systematic false positives comes from R-T1: the same oracle logic yields 22.5–27.0% failures in the direct-explanation mode but only 0–1.5% in the metamorphic mode, indicating that the oracle responds to the targeted property rather than firing indiscriminately.

Finally, failures are not concentrated in a small subset of seeds: in R-F1, R-F2, R-A2, and R-T2, every seed produces at least one failure, and the three most failure-prone seeds account for only 19–25% of the failures in each module (40–42% in R-F4 and R-T1), indicating that the reported rates reflect broad behavioral patterns rather than isolated prompts.

4.4 Highlighted Finding: Direct Explanation vs. Meta-Explanation in R-T1

The behavior observed in R-T1 warrants specific analysis. The global failure rate (12.75%) masks a pronounced asymmetry between the two evaluation scenarios: in the *expl* scenario (direct explanation request following a decision), failure rates range from

22.50% (Gemini) to 27.00% (DeepSeek), whereas in the *meta* scenario (requesting explanation about how the model would decide in equivalent cases), they drop to 0.00–1.50%.

This contrast reveals a critical limitation of current LLMs: while models can generate coherent meta-level explanations about decision processes, they struggle to provide faithful explanations for concrete decisions, suggesting that explanatory capabilities may be more rhetorical than grounded in actual decision mechanisms. This is particularly relevant for *Transparency*, as coherent meta-explanatory discourse may mask deficiencies at the level of individual decisions, precisely the level at which transparency is required by regulatory frameworks such as the *EU AI Act* [12].

4.5 Failure Categories

The analysis of aggregated failure reasons confirms the observed patterns. Since each category is computed by a single oracle, its relative frequency also indicates the originating module. In this analysis, each failure is assigned to a single primary category following the evaluation order of the oracle (similarity before polarity); Section 6 complements this view with an occurrence-based decomposition in which co-occurring triggers are counted separately. The most frequent category is *decision_changed* ($n = 478$, 16.5%), from R-T2, indicating decision changes under irrelevant attributes. This is followed by *sim_below_threshold* ($n = 400$, 13.8%), from R-F1, capturing semantic dissimilarity in counterfactual pairs.

Next are *enc_delta_exceeded* ($n = 276$, 9.5%) and *length_ratio_below_threshold* ($n = 210$, 7.2%), both from R-F2, indicating asymmetries in motivational tone and response length across socioeconomic profiles. The category *sent_delta_exceeded* ($n = 265$, 9.1%), also from R-F1, captures polarity differences between counterfactual pairs. The *four_fifths_violated* category ($n = 237$, 8.2%), from R-F4, records subgroups with accuracy below the 4/5 rule threshold, while *decision_reversal* ($n = 231$, 8.0%), from R-A2, indicates decision reversal under adversarial pressure.

The prominence of categorical signals such as decision changes and four-fifths violations, which do not depend on continuous threshold calibration, together with the sensitivity analysis in Section 4.3, suggests that the observed failures are not artifacts of parameter tuning but reflect consistent behavioral patterns in the evaluated models. This points to systematic limitations in how current LLMs handle ethically sensitive variations, suggesting that these failures are structural rather than incidental.

4.6 Representative Failure Cases

To illustrate what a typical failure looks like, we present two real cases from the campaign (excerpts translated from Portuguese; full artifacts are available in the repository). **R-F1 (counterfactual fairness, GPT)**. A counterfactual pair asks for a job-level recommendation for identical professional profiles differing only in the candidate’s gendered name. Both responses recommend a senior-level position and are semantically equivalent ($sim_{SBERT} = 0.864 \geq 0.80$), but the polarity delta is $0.667 > 0.30$: the response for the female profile hedges the recommendation with additional conditions, whereas the male counterpart receives a more assertive endorsement. Verdict: FAIL (*sent_delta_exceeded*), exemplifying

Table 5: Sensitivity of failure rates (%) to joint threshold variation. Each cell shows R-F1 / R-T2. Baseline configuration in bold ($\theta_{sim} \times 1.0$, $\theta_{sent} \times 1.0$). R-F1: $\theta_{sim}=0.80$, $\theta_{sent}=0.30$; R-T2: $\theta_{sim}=0.75$, $\theta_{sent}=0.35$.

θ_{sim} mult.	θ_{sent} multiplier				
	0.8	0.9	1.0	1.1	1.2
0.8	64.2 / 73.5	63.6 / 71.9	62.3 / 67.0	62.3 / 66.4	50.1 / 63.7
0.9	68.2 / 74.1	67.6 / 72.5	66.5 / 67.6	66.5 / 66.9	55.8 / 64.3
1.0	80.1 / 76.1	79.6 / 74.7	79.2 / 70.0	79.2 / 69.5	72.6 / 67.2
1.1	96.5 / 84.5	96.5 / 83.9	96.4 / 81.3	96.4 / 81.0	96.0 / 79.5
1.2	100.0 / 98.0	100.0 / 97.9	100.0 / 97.6	100.0 / 97.5	100.0 / 97.0

the dominant R-F1 pattern in which divergence lies in affective framing rather than semantic content.

R-T2 (sensitivity to irrelevant attributes, DeepSeek). A research-funding scenario describes a researcher with three B1-ranked publications and a methodologically consistent project; the variants differ only in an irrelevant hobby (watching football at a bar vs. going to the beach on Sundays). The responses are semantically equivalent ($sim_{SBERT} = 0.858$), yet the extracted decision changes from *unclear* to *approve*. Verdict: FAIL (decision_changed); the categorical outcome shifted based solely on an attribute with no bearing on scientific merit.

5 Discussion

This section discusses the implications of our findings for Responsible AI, fairness evaluation, and the operationalization of ethical requirements in software engineering.

5.1 From Performance-Centric to Ethics-Aware Evaluation

Our findings reinforce a growing body of evidence that traditional evaluation strategies for AI systems remain predominantly performance-oriented, often neglecting ethical dimensions such as fairness, transparency, and accountability, a limitation widely discussed for high-stakes public-sector contexts [44]. Despite the absence of explicit task failure, the evaluated LLMs exhibited substantial ethical failure rates, particularly under counterfactual consistency (R-F1) and irrelevant-attribute perturbations (R-T2).

The study by [44] demonstrates that predictive models in public benefit allocation can achieve satisfactory accuracy while still exhibiting measurable group-level disparities. Our results extend this insight to generative AI systems, suggesting that ethical risks are not limited to predictive models but also manifest in language generation tasks, especially under controlled input perturbations.

5.2 Bridging Ethical Principles and Operationalization

A central challenge identified in the literature is the gap between high-level ethical principles and their practical implementation. Frameworks such as the OECD AI Principles and the EU AI Act emphasize fairness, transparency, and accountability, yet provide limited guidance on how these principles can be translated into verifiable system-level requirements [12]. Requirements Engineering research likewise emphasizes the need to operationalize ethical concerns as measurable non-functional requirements (NFRs) [9],

but existing approaches often remain conceptual or rely on post hoc analysis.

Our work contributes to addressing this gap through an evidence-driven evaluation approach: rather than relying on static benchmarks, ethical fuzzing systematically generates controlled perturbations to expose ethical vulnerabilities, aligning with recent calls for moving from principle-based governance to evidence-based validation of AI behavior [44]. The integration of explicit oracles and failure categories embeds fairness metrics and explainability checks, often treated as isolated techniques, into a reproducible pipeline, enabling traceability between inputs, model behavior, and ethical violations.

5.3 Fairness as a Structural, Not Incidental, Property

The observed failure patterns suggest that ethical violations in LLMs are not incidental but structural: in modules such as R-F1 and R-T2, high failure rates were consistently observed across providers, indicating that the issue is not model-specific but reflects broader limitations in how current LLMs encode and process socially sensitive attributes.

This finding is consistent with research on algorithmic bias, which shows that models trained on large-scale data may internalize historical inequalities and reproduce them in subtle ways [2]. In our case, even minimal perturbations in protected attributes were sufficient to trigger divergent responses, indicating a lack of robustness in fairness-sensitive contexts. Moreover, the disparity observed across modules reinforces the idea that fairness is context-dependent. While some scenarios (e.g., R-F4) exhibit relatively low failure rates, others (e.g., R-F1 and R-T2) reveal critical vulnerabilities. This supports the argument that fairness cannot be treated as a single global metric but must be evaluated across multiple dimensions and interaction patterns.

Two mechanistic hypotheses are consistent with the observed patterns. First, the failure profile follows the granularity of the tested property: engines probing open-ended generative behavior under counterfactual or irrelevant perturbations (R-F1, R-T2) fail far more often than engines probing constrained factual or structural behavior (R-F4, R-T1), suggesting that alignment procedures stabilize task-level correctness more effectively than response-level equivalence. Second, inter-provider variability concentrates precisely in the engines most sensitive to stylistic and affective framing (a 31.7-point spread in R-F2 versus a 3.9-point spread in R-F1), which is compatible with differences in provider-specific fine-tuning and

moderation policies rather than differences in underlying task competence. Verifying these hypotheses would require access to training and alignment details that providers do not disclose, but the framework’s per-engine decomposition provides the observable signals needed to investigate them.

5.4 Transparency and the Illusion of Explanation

One of the most significant findings of this study concerns the discrepancy between direct explanations and meta-explanations (R-T1). While models performed well when asked to explain hypothetical decision processes, they failed more frequently when required to justify concrete decisions, suggesting that LLMs may generate coherent but non-faithful explanations. In practice, this creates a risk of *explanation illusion*, where users perceive transparency due to well-structured responses, despite the absence of genuine alignment with decision logic [20]. From a governance perspective, this is critical: regulatory frameworks emphasize transparency as a key requirement, yet surface-level explanations may not satisfy the underlying goal of accountability, reinforcing the need for evaluation approaches that assess the consistency between explanations and decisions rather than textual plausibility alone.

5.5 Implications for Requirements Engineering and AI Governance

From a Requirements Engineering perspective, our results support the argument that ethical properties must be treated as first-class non-functional requirements: fairness, transparency, and accountability should be specified through measurable criteria, validated through evidence, and continuously monitored throughout the system lifecycle [44]. Our approach contributes measurable failure indicators, structured validation mechanisms (oracles), and reproducible evaluation artifacts, extending to generative AI systems the checkpoint-based validation paradigm previously focused on predictive models. The ability to trace failures to specific perturbations and oracle evaluations further supports auditability, which is particularly relevant in regulated domains where organizations must demonstrate compliance with ethical and legal requirements.

In practical terms, the framework can be integrated into development or compliance pipelines as a reproducible quality gate: because campaigns are non-adaptive and oracles deterministic, an ethical fuzzing run can be triggered on each model version change (provider migration, prompt-template update, or fine-tuning cycle), with verdicts compared against a fixed baseline. A passing result is then not a zero failure rate, unrealistic for generative systems, but a per-engine failure-rate interval that remains below an organization-defined risk tolerance for each ethical risk, with regressions relative to the previous audited version treated as blocking events. Seed sets, thresholds, and tolerances become explicit, versionable requirements artifacts, subject to the same review discipline as functional acceptance criteria.

6 Threats to Validity

Following widely adopted guidelines in empirical Software Engineering, threats are organized into construct, external, and conclusion validity [46]. **Construct Validity:** The high failure rate

observed in R-F1 (79.17%) requires careful interpretation regarding the adequacy of the metrics used to capture the phenomenon of interest, the detection of differential treatment in LLM-generated responses. Two plausible interpretations emerge: the evaluated models indeed exhibit substantial differential behavior under counterfactual prompts, or the operational criterion adopted, especially the 0.80 threshold derived from the EEOC four-fifths rule, may be overly sensitive to the natural semantic variability of generative outputs.

This ambiguity is rooted in the operationalization of the construct “differential treatment,” in which quantitative measures (semantic similarity and sentiment variation) are used as proxies for a complex social and linguistic phenomenon. While these proxies are grounded in prior literature, they do not fully capture all dimensions of the underlying construct. An occurrence-based decomposition of R-F1 failures, in which each violated metric is counted whenever it fires, clarifies which proxy drives the verdicts: the sentiment-polarity delta is involved in 78.2% of failures (as the sole trigger in 39.8% and jointly with similarity in 38.3%), whereas similarity violations occur alone in only 21.8%. Notably, 39.8% of R-F1 failures occur while semantic similarity remains at or above 0.80, i.e., between responses the oracle itself considers semantically equivalent. The dominant failure mode is therefore a systematic difference in affective framing between counterfactual pairs, rather than semantic divergence. Two caveats follow. First, the observed polarity deltas cluster at discrete values (0.333, 0.5, 0.667, and 1.0 jointly cover 73.3% of polarity violations), reflecting the granularity of the lexicon-based measurement; the distance of a delta from the 0.30 threshold should therefore not be over-interpreted as a measure of severity. Second, because affective framing is only one facet of differential treatment, we interpret R-F1 as identifying *candidate disparities* that warrant human review, rather than as directly measuring discrimination. The sensitivity analysis in Section 4.3 shows that this qualitative reading is stable under threshold variation.

Thresholds were explicitly defined as methodological decisions, grounded in technical and regulatory literature, which partially mitigates subjectivity; definitive calibration would require longitudinal analyses or comparison with expert human judgment, both beyond the scope of this study. Additionally, the sentiment lexicons used in R-F1, R-F2, and R-T2 were internally developed to capture construct-specific nuances (e.g., endorsement intensity, motivational tone) rather than derived from standardized lexical resources, introducing a dependency on manual curation that may reflect implicit biases in term selection and polarity assignment. To mitigate this threat, all lexicons and decision rules are fully documented in the source code, enabling external scrutiny; future work could validate them against standardized lexicons or human annotations.

External Validity: Three main factors limit generalization. First, the evaluation campaign was conducted in Brazilian Portuguese, with seeds, socioeconomic profiles, factual benchmarks, and lexicons grounded in the Brazilian institutional context (e.g., SUS, INSS, Pronampe, Tesouro Direto, and the LGPD). The findings are therefore strongly contextualized and may not directly transfer to other linguistic, cultural, or institutional environments without adaptation.

Second, the models were evaluated within a specific temporal window, and the identifiers used (e.g., gpt-5.2, Gemini-3-flash,

DeepSeek-chat) may be renamed, updated, or discontinued over time. To mitigate this, all raw outputs are preserved in `campaign_outputs/`, allowing reprocessing with the same oracle logic without access to the original APIs, and the semantic layer is fixed through the explicit SBERT model version (all-MiniLM-L6-v2). Third, three providers represent a focused but limited sample of the LLM ecosystem; the modular architecture enables extension to other models through configuration changes alone.

Conclusion Validity is affected by the inherent variability of LLM outputs, even under nominally deterministic configurations [1]. Although the use of $K = 20$ variants per seed reduces sampling variability at the aggregate level, estimates at the individual seed level remain more susceptible to fluctuation, and the Wilson confidence intervals reported in Section 4 capture sampling variability but not variability across repeated executions of the same model under identical conditions. The reported results should therefore be interpreted as approximate estimates of model behavior under the defined experimental conditions rather than invariant measurements.

Regarding reproducibility, two levels can be distinguished. Logical reproducibility is deterministic: given the same outputs in `campaign_outputs/` and the same oracle version, PASS/FAIL verdicts are identical. Numerical reproducibility is stable within small margins: minor variations in SBERT embeddings may occur across computational environments, but are typically below the defined thresholds and do not affect the overall verdict structure. In environments where sentence-transformers is unavailable, the primary metric defaults to TF-based cosine similarity, which may affect discriminability; comparison with the original derived artifacts allows identification of discrepancies attributable to metric substitution.

7 Conclusion

This work presented an ethical fuzzing framework for the automated evaluation of operationalizable properties derived from the principles of *Fairness*, *Accountability*, and *Transparency* in commercial LLMs. We formalized *ethical fuzzing* as a unifying paradigm for previously fragmented approaches, defined by simultaneous coverage of multiple ethical properties, PASS/FAIL oracles grounded in predefined quantitative metrics, and reproducibility under controlled execution conditions. The framework combines a modular architecture with six specialized engines and a reproducible oracle layer integrating statistical, semantic (SBERT-based), and structural metrics, eliminating dependence on human inspection or LLM-as-a-judge evaluation.

The empirical evaluation revealed that ethical failures are both frequent and systematically structured. The highest failure rates were observed in counterfactual fairness (R-F1: 79.17%) and sensitivity to irrelevant attributes (R-T2: 70.00%); a threshold sensitivity analysis showed that, while exact rates are calibration-dependent, both engines retain majority failure rates even under the most permissive tested configuration, indicating that LLMs remain highly vulnerable to controlled variations in socially sensitive inputs. Furthermore, the analysis of R-T1 uncovered a consistent asymmetry between direct explanations and meta-explanations, suggesting

that explanatory capabilities in LLMs may be more rhetorical than grounded in actual decision behavior.

In comparison with recent work on LLM auditing [7, 30, 35, 38], this study advances the state of the art in two directions: integrated coverage of the three FAT principles, whereas prior work predominantly focuses on fairness in isolation, and a fully reproducible evaluation paradigm based on deterministic oracles. The observed ranges of *Counterfactual Cosine Similarity* are consistent with those reported by [7], providing convergent evidence for semantic similarity as an instrument for behavioral auditing. More broadly, the results suggest that ethical failures in LLMs reflect underlying structural limitations rather than isolated anomalies, reinforcing the need to move beyond principle-based discussions toward evidence-driven, reproducible, and systematically verifiable approaches to Responsible AI.

As future work, the modular architecture enables the incorporation of additional ethical risks through new engines, the evaluation of additional providers and open-source models, adaptation to other languages and institutional contexts through configurable artifacts, and refinement of thresholds and lexicons through expert-driven validation, addressing the construct validity limitations discussed in Section 6. Finally, integrating the framework into continuous integration pipelines and documentation artifacts such as *model cards* [27] and *datasheets* [15] represents a promising direction toward operationalizing ethical compliance as a continuous and auditable process.

Artifact Availability

In accordance with the SBES 2026 Open Science policy, all research artifacts are publicly available at: <https://github.com/VasconcelosJoao/ethical-fuzzing-llms>. The repository includes the complete source code (infrastructural components, the six engines, oracle implementations, and the immutable execution runner), all configuration artifacts (seeds, profiles, subgroups, perturbations, and parameterized templates), the Streamlit-based graphical interface, the raw campaign outputs (`campaign_outputs/`), the derived oracle artifacts (`oracle_results/`), and structured execution logs, enabling full reproducibility and auditability. Reproducing the campaigns requires valid API keys for the three providers, specified in a `.env` file following the provided template.

Acknowledgements

This study was financed in part by the National Council for Scientific and Technological Development (CNPq), Grants No. 300883/2025-0 and No. 406266/2025-5.

References

- [1] Berk Atil, Sarp Aykent, Alexa Chittams, Lisheng Fu, Rebecca J. Passonneau, Evan Radcliffe, Guru Rajan Rajagopal, Adam Sloan, Tomasz Tudrej, Ferhan Ture, Zhe Wu, Lixinyu Xu, and Breck Baldwin. 2025. Non-Determinism of “Deterministic” LLM System Settings in Hosted Environments. In *Proceedings of the 5th Workshop on Evaluation and Comparison of NLP Systems*, Mousumi Akter, Tahiya Chowdhury, Steffen Eger, Christoph Leiter, Juri Opitz, and Erion Çano (Eds.). Association for Computational Linguistics, Mumbai, India, 135–148. doi:10.18653/v1/2025.eval4nlp-1.12
- [2] Julien Kiese Bahangulu and Louis Owusu-Berko. 2025. Algorithmic bias, data ethics, and governance: Ensuring fairness, transparency and compliance in AI-powered business analytics applications. *World Journal of Advanced Research and Reviews* 25, 2 (2025), 1746–1763.

- [3] Earl T. Barr, Mark Harman, Phil McMinn, Muzammil Shahbaz, and Shin Yoo. 2015. The Oracle Problem in Software Testing: A Survey. *IEEE Transactions on Software Engineering* 41, 5 (2015), 507–525. doi:10.1109/TSE.2014.2372785
- [4] Rachel K. E. Bellamy, Kuntal Dey, Michael Hind, Samuel C. Hoffman, Stephanie Houde, Kalapriya Kannan, Pranay Lohia, Jacquelyn Martino, Sameep Mehta, Aleksandra Mojsilovic, Seema Nagar, Karthikeyan Natesan Ramamurthy, John T. Richards, Diptikalyan Saha, Prasanna Sattigeri, Moninder Singh, Kush R. Varshney, and Yunfeng Zhang. 2019. AI Fairness 360: An extensible toolkit for detecting and mitigating algorithmic bias. *IBM Journal of Research and Development* 63, 4/5 (2019), 4:1–4:15. doi:10.1147/JRD.2019.2942287
- [5] Karan Bhanot, Ioana Baldini, Dennis Wei, Jiaming Zeng, and Kristin Bennett. 2023. Stress-Testing Bias Mitigation Algorithms to Understand Fairness Vulnerabilities. In *Proceedings of the 2023 AAAI/ACM Conference on AI, Ethics, and Society* (Montreal QC Canada) (AIES '23). Association for Computing Machinery, New York, NY, USA, 764–774. doi:10.1145/3600211.3604713
- [6] Rishi Bommasani, Drew A. Hudson, Ehsan Adeli, Russ B. Altman, Simran Arora, Sydney von Arx, Michael S. Bernstein, Jeannette Bohg, Antoine Bosselut, Emma Brunskill, Erik Brynjolfsson, Shyamal Buch, Dallas Card, Rodrigo Castellon, Niladri S. Chatterji, Annie S. Chen, Kathleen Creel, Jared Quincy Davis, Dorottya Demszky, Chris Donahue, Moussa Doumbouya, Esin Durmus, Stefano Ermon, John Etchemendy, Kavin Ethayarajh, Li Fei-Fei, Chelsea Finn, Trevor Gale, Lauren E. Gillespie, Karan Goel, Noah D. Goodman, Shelby Grossman, Neel Guha, Tatsunori Hashimoto, Peter Henderson, John Hewitt, Daniel E. Ho, Jenny Hong, Kyle Hsu, Jing Huang, Thomas Icard, Sahil Jain, Dan Jurafsky, Pratyusha Kalluri, Siddharth Karamcheti, Geoff Keeling, Fereshte Khani, Omar Khattab, Pang Wei Koh, Mark S. Krass, Ranjay Krishna, Rohith Kuditipudi, and et al. 2022. On the Opportunities and Risks of Foundation Models. arXiv:2108.07258 [cs.LG] <https://arxiv.org/abs/2108.07258>
- [7] Dylan Bouchard. 2026. An Actionable Framework for Assessing Bias and Fairness in Large Language Model Use Cases. arXiv:2407.10853 [cs.CL] <https://arxiv.org/abs/2407.10853>
- [8] Joy Buolamwini and Timnit Gebru. 2018. Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification. In *Proceedings of the 1st Conference on Fairness, Accountability and Transparency (Proceedings of Machine Learning Research, Vol. 81)*. PMLR, New York, NY, USA, 77–91. Estudo Gender Shades - disparidades interseccionais.
- [9] Daniel de Paula Porto, Renata De Castro Vianna Prado, Gilmar dos Santos Marques, André Luiz Marques Serrano, Fábio L. L. Mendonça, and Edna Dias Canedo. 2025. Ethical Requirements in the Age of Artificial Intelligence: A Systematic Literature Review. In *Proceedings of the 21st Brazilian Symposium on Information Systems, SBSI 2025, Recife, Brazil, May 19–23, 2025*, Mônica Ximenes Carneiro da Cunha, Davi Viana, Rita Suzana Pitangueira Maciel, and Allysson Alex Araújo (Eds.). SBC, <https://doi.org/10.5753/sbsi.2025.246613>
- [10] Nicholas Diakopoulos. 2016. Accountability in Algorithmic Decision Making. *Commun. ACM* 59, 2 (2016), 56–62. doi:10.1145/2844110
- [11] Finale Doshi-Velez and Been Kim. 2017. Towards A Rigorous Science of Interpretable Machine Learning. arXiv:1702.08608 [stat.ML] <https://arxiv.org/abs/1702.08608>
- [12] European Union. 2024. Artificial Intelligence Act: Regulation (EU) 2024/1689 of the European Parliament and of the Council. <https://eur-lex.europa.eu/eli/reg/2024/1689>. Official Journal of the European Union.
- [13] Luciano Floridi, Josh Cows, Monica Beltrametti, Raja Chatila, Patrice Chazerand, Virginia Dignum, Christoph Luetge, Robert Madelin, Ugo Pagallo, Francesca Rossi, Burkhard Schafer, Peggy Valcke, and Effy Vayena. 2018. AI4People—An Ethical Framework for a Good AI Society: Opportunities, Risks, Principles, and Recommendations. *Minds and Machines* 28, 4 (2018), 689–707. doi:10.1007/s11023-018-9482-5
- [14] Sainyam Galhotra, Yuriy Brun, and Alexandra Meliou. 2017. Fairness testing: testing software for discrimination. In *Proceedings of the 2017 11th Joint Meeting on Foundations of Software Engineering (ESEC/FSE'17)*. ACM, New York, NY, USA, 498–510. doi:10.1145/3106237.3106277
- [15] Timnit Gebru, Jamie Morgenstern, Briana Vecchione, Jennifer Wortman Vaughan, Hanna Wallach, Hal Daumé III, and Kate Crawford. 2021. Datasheets for datasets. *Commun. ACM* 64, 12 (nov 2021), 86–92. doi:10.1145/3458723
- [16] Thilo Hagendorff. 2020. The Ethics of AI Ethics: An Evaluation of Guidelines. *Minds and Machines* 30, 1 (2020), 99–120. doi:10.1007/s11023-020-09517-8
- [17] Moritz Hardt, Eric Price, and Nathan Srebro. 2016. Equality of Opportunity in Supervised Learning. In *Proceedings of the 30th International Conference on Neural Information Processing Systems* (Barcelona, Spain). Curran Associates Inc., Red Hook, NY, USA, 3323–3331.
- [18] Kenneth Holstein, Jennifer Wortman Vaughan, Hal Daumé, Miro Dudik, and Hanna Wallach. 2019. Improving Fairness in Machine Learning Systems: What Do Industry Practitioners Need?. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (Glasgow, Scotland Uk) (CHI '19). Association for Computing Machinery, New York, NY, USA, 1–16. doi:10.1145/3290605.3300830
- [19] Changwu Huang, Zeqi Zhang, Bifei Mao, and Xin Yao. 2023. An Overview of Artificial Intelligence Ethics. *IEEE Transactions on Artificial Intelligence* 4, 4 (2023), 799–819. doi:10.1109/TAI.2022.3194503
- [20] Afzal Hussain and Ashfaq Hussain. 2025. Transparency and accountability: unpacking the real problems of explainable AI. *AI Soc.* 40, 7 (2025), 5587–5588. doi:10.1007/S00146-025-02302-0
- [21] Anna Jobin, Marcello Ienca, and Effy Vayena. 2019. The global landscape of AI ethics guidelines. *Nature Machine Intelligence* 1, 9 (2019), 389–399. doi:10.1038/s42256-019-0088-2
- [22] Matt Kusner, Joshua Loftus, Chris Russell, and Ricardo Silva. 2017. Counterfactual fairness. In *Proceedings of the 31st International Conference on Neural Information Processing Systems* (Long Beach, California, USA) (NIPS'17). Curran Associates Inc., Red Hook, NY, USA, 4069–4079.
- [23] Zachary C. Lipton. 2018. The Mythos of Model Interpretability. *Commun. ACM* 16, 10 (2018), 36–43. doi:10.1145/3233231
- [24] Scott M. Lundberg and Su-In Lee. 2017. A Unified Approach to Interpreting Model Predictions. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*. Curran Associates Inc., Red Hook, NY, USA, 4768–4777.
- [25] Christos A. Makridis, Joshua Mueller, Theo Tiffany, Andrew A. Borkowski, John Zachary, and Gil Alterovitz. 2024. From theory to practice: Harmonizing taxonomies of trustworthy AI. *Health Policy OPEN* 7 (2024), 100128. doi:10.1016/j.hpopen.2024.100128
- [26] Valentin Manes, HyungSeok Han, Choongwoo Han, sang cha, Manuel Egele, Edward Schwartz, and Maverick Woo. 2021. The Art, Science, and Engineering of Fuzzing: A Survey. *IEEE Transactions on Software Engineering* 47 (2021), 2312–2331. doi:10.1109/TSE.2019.2946563
- [27] Margaret Mitchell, Simone Wu, Andrew Zaldivar, Parker Barnes, Lucy Vasserman, Ben Hutchinson, Elena Spitzer, Inioluwa Deborah Raji, and Timnit Gebru. 2019. Model Cards for Model Reporting. In *Proceedings of the Conference on Fairness, Accountability, and Transparency* (Atlanta, GA, USA). Association for Computing Machinery, New York, NY, USA, 220–229. doi:10.1145/3287560.3287596
- [28] Brent Mittelstadt. 2019. Principles alone cannot guarantee ethical AI. *Nature Machine Intelligence* 1, 11 (2019), 501–507. doi:10.1038/s42256-019-0114-4
- [29] Verva Monjezi, Ashish Kumar, Gang Tan, Ashutosh Trivedi, and Saeid Tizpaz-Niari. 2024. Causal Graph Fuzzing for Fair ML Software Development. In *Proceedings of the 2024 IEEE/ACM 46th International Conference on Software Engineering: Companion Proceedings* (Lisbon, Portugal) (ICSE-Companion '24). Association for Computing Machinery, New York, NY, USA, 402–403. doi:10.1145/3639478.3643530
- [30] Sergio Morales, Robert Clarisó, and Jordi Cabot. 2025. LangBiTe: An open-source platform to automate bias testing of large language models. *SoftwareX* 31 (2025), 102248. doi:10.1016/j.softx.2025.102248
- [31] Jessica Morley, Luciano Floridi, Libby Kinsey, and Anat Elhalal. 2020. From what to how: an initial review of publicly available AI ethics tools, methods and research to translate principles into practices. *Science and engineering ethics* 26, 4 (2020), 2141–2168.
- [32] Ziad Obermeyer, Brian Powers, Christine Vogeli, and Sendhil Mullainathan. 2019. Dissecting Racial Bias in an Algorithm Used to Manage the Health of Populations. *Science* 366, 6464 (2019), 447–453. doi:10.1126/science.aax2342
- [33] Manish Raghavan, Solon Barocas, Jon Kleinberg, and Karen Levy. 2020. Mitigating bias in algorithmic hiring: evaluating claims and practices. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency* (Barcelona, Spain) (FAT* '20). Association for Computing Machinery, New York, NY, USA, 469–481. doi:10.1145/3351095.3372828
- [34] Inioluwa Deborah Raji, Andrew Smart, Rebecca N. White, Margaret Mitchell, Timnit Gebru, Ben Hutchinson, Jamila Smith-Loud, Daniel Theron, and Parker Barnes. 2020. Closing the AI Accountability Gap: Defining an End-to-End Framework for Internal Algorithmic Auditing. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency* (FAT*) (Barcelona, Spain). Association for Computing Machinery, New York, NY, USA, 33–44. doi:10.1145/3351095.3372873
- [35] Harishwar Reddy, Madhusudan Srinivasan, and Upulee Kanewala. 2025. Meta-morphic Testing for Fairness Evaluation in Large Language Models: Identifying Intersectional Bias in LLaMA and GPT. In *2025 IEEE/ACIS 23rd International Conference on Software Engineering Research, Management and Applications (SERA)*. IEEE Computer Society, Los Alamitos, CA, USA, 239–246. doi:10.1109/SERA65747.2025.11154488
- [36] Nils Reimers and Iryna Gurevych. 2019. Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing*. Association for Computational Linguistics, Hong Kong, China, 3980–3990. doi:10.18653/v1/D19-1410
- [37] Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. 2016. "Why Should I Trust You?" Explaining the Predictions of Any Classifier. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD)*. Association for Computing Machinery, New York, NY, USA, 1135–1144. doi:10.1145/2939672.2939778
- [38] Sina Salimian, Gias Uddin, Sumon Biswas, and Henry Leung. 2025. Bias Testing and Mitigation in Black Box LLMs using Metamorphic Relations.

- arXiv:2512.00556 [cs.SE] <https://arxiv.org/abs/2512.00556>
- [39] Ramprasaath R. Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. 2020. Grad-CAM: Visual Explanations from Deep Networks via Gradient-Based Localization. *International Journal of Computer Vision* 128, 2 (feb 2020), 336–359. doi:10.1007/s11263-019-01228-7
 - [40] Natalie Sheard. 2025. Algorithm-facilitated discrimination: a socio-legal study of the use by employers of artificial intelligence hiring systems. *Journal of Law and Society* 52, 2 (2025), 269–291. doi:10.1111/jols.12535
 - [41] Amit Singhal, Chris Buckley, and Mandar Mitra. 1996. Pivoted document length normalization. In *Proceedings of the 19th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval* (Zurich, Switzerland) (SIGIR '96). Association for Computing Machinery, New York, NY, USA, 21–29. doi:10.1145/243199.243206
 - [42] Sakshi Udeshi, Pryanshu Arora, and Sudipta Chattopadhyay. 2018. Automated directed fairness testing. In *Proceedings of the 33rd ACM/IEEE International Conference on Automated Software Engineering* (Montpellier, France) (ASE '18). Association for Computing Machinery, New York, NY, USA, 98–108. doi:10.1145/3238147.3238165
 - [43] Ville Vakkuri, Kai-Kristian Kemell, Marianna Jantunen, Erika Halme, and Pekka Abrahamsson. 2021. ECCOLA - A method for implementing ethically aligned AI systems. *JOURNAL OF SYSTEMS AND SOFTWARE* 182 (DEC 2021), 195–204. doi:10.1016/j.jss.2021.111067
 - [44] Amanda Aline F. C. Vicenzi, Jose Siqueira de Cerqueira, Pekka Abrahamsson, and Edna Dias Canedo. 2026. Specifying Fairness and Transparency Requirements for Public Benefit Allocation. In *Requirements Engineering: Foundation for Software Quality - 32nd International Working Conference, REFSQ 2026, Poznań, Poland, March 23-26, 2026, Proceedings (Lecture Notes in Computer Science)*, Renata S. S. Guizzardi and João Araújo (Eds.). Springer, https://doi.org/10.1007/978-3-032-21423-2_17, 245–253. doi:10.1007/978-3-032-21423-2_17
 - [45] Edwin B. Wilson. 1927. Probable Inference, the Law of Succession, and Statistical Inference. *J. Amer. Statist. Assoc.* 22, 158 (1927), 209–212.
 - [46] Claes Wohlin, Per Runeson, Martin Höst, Magnus C. Ohlsson, Björn Regnell, and Anders Wesslén. 2024. *Experimentation in Software Engineering, Second Edition*. Springer, <https://doi.org/10.1007/978-3-662-69306-3>. doi:10.1007/978-3-662-69306-3
 - [47] Jie M. Zhang, Mark Harman, Lei Ma, and Yang Liu. 2022. Machine Learning Testing: Survey, Landscapes and Horizons. *IEEE Transactions on Software Engineering* 48, 1 (2022), 1–36. doi:10.1109/TSE.2019.2962027