

Reward Engineering for Software Tasks: A Survey of Reinforcement Learning Approaches

Md Rayhanul Masud¹, Azmine Toushik Wasi², Salman Rahman³, Md Rizwan Parvez⁴

¹University of California, Riverside

²Shahjalal University of Science and Technology

³University of California, Los Angeles

⁴Qatar Computing Research Institute (QCRI)

Abstract

Reinforcement learning is increasingly used for code-centric software engineering tasks, including code generation, understanding, repair, testing, and optimization, especially with the rise of large language models and autonomous agents. A core challenge in these settings is reward design. Unlike standard RL domains with clear scalar objectives, software tasks involve competing goals such as correctness, security, efficiency, and readability, which are difficult to capture with a single reward. As a result, RL-for-SE systems rely on heterogeneous signals, including compilation results, unit tests, coverage metrics, retrieval scores, and learned preferences. Yet this work remains scattered across tasks and communities. To our knowledge, this survey provides the first systematic review of reward engineering for RL in software tasks. We organize prior work by reward source, granularity, and aggregation. We then distill the findings into practical guidance, including a decision matrix, a design guide, and a reward engineering reporting standard for future RL-for-SE systems.

1 Introduction

Reinforcement learning has become an effective paradigm for software engineering tasks by enabling agents to learn from executable feedback, test outcomes, and preference signals rather than from labeled examples alone (Wang et al., 2025a). Recent systems, including CodeRL for program synthesis (Le et al., 2022), DeepSeek-R1 for reasoning (Guo et al., 2025), and RL-based agents for automated program repair (Wei et al., 2025), have demonstrated state-of-the-art performance on code generation, bug fixing, test synthesis, and repository-level issue resolution across benchmarks such as HumanEval, MBPP, and SWE-bench. A key factor underlying this progress is the design of the RL environment and reward function. The environment specifies how the agent interacts

with software artifacts through states, actions, and transitions (Sutton et al., 1998), whereas the reward function maps complex software objectives, including correctness, efficiency, security, and readability, into scalar learning signals (Wei et al., 2025).

Despite the growing success of RL in software engineering, reward design remains a central challenge (Wang et al., 2025a). Unlike game-playing domains with clear, immediate rewards (Mnih et al., 2015; Silver et al., 2018), software tasks involve competing criteria that resist reduction to a single scalar signal. A program may compile yet fail tests, satisfy functional requirements yet expose security flaws, or behave correctly while remaining hard to understand or inefficient. Existing approaches typically optimize only part of this space, relying on execution-based rewards for correctness, similarity-based rewards for reference alignment, or learned preference models for more subjective qualities. When such signals are combined, additional challenges arise, including reward scaling, weight sensitivity, and conflicts among objectives (Roijers et al., 2013; Agarwal et al., 2022; de Langis et al., 2024; Lu et al., 2026). These challenges are further shaped by the surrounding RL environment, since reward behavior depends on how states are represented, which actions are permitted, and at what temporal granularity feedback is assigned.

Given these challenges, a systematic and reward-centric survey of RL for software engineering is both timely and necessary, yet no such resource exists. Existing surveys on RL for SE (Yang et al., 2022; Watson et al., 2022; Sharma et al., 2024; Wang et al., 2025a; Zhang et al., 2025) primarily address algorithm selection, task performance, and application domains, treating reward functions as secondary implementation details rather than as first-class design choices that fundamentally shape learning outcomes. Surveys on RLHF and RL

for LLMs (Ibrahim et al., 2024; Zhang et al., 2025) discuss reward modeling for alignment and reasoning, but do not address the unique challenges of software artifacts, namely executability, test coverage, semantic correctness, and deterministic oracle availability, which give rise to a qualitatively different reward design landscape.

To address this gap, we present the first systematic survey of reward design approaches to reinforcement learning in software engineering, covering 50+ papers published in 2018-early 2026. We organize the literature through a taxonomy grounded in three reward design questions: what signal is used, at what temporal granularity it is delivered, and how multiple signals are aggregated. We address the following research questions:

RQ1: How can reward schemes in RL for software engineering be categorized across tasks and mechanisms?

RQ2: What reward design patterns recur across different software engineering tasks?

RQ3: How are rewards defined at different granularities, from token and line-level to program and trajectory-level?

RQ4: How are heterogeneous reward signals, such as execution correctness, semantic similarity, and human preference, combined during training?

Unlike prior surveys, this work focuses on reward engineering as the core design issue. Because software supports executable, deterministic feedback, reward design can rely on tests, coverage, and execution traces. A unifying principle emerges from this analysis: reward source, granularity, and aggregation choices are not free design parameters but follow from task structural conditions, namely oracle availability, observability of intermediate states, and the pressure to combine heterogeneous signals. We distill this principle into the decision matrix shown in Table 1 and revisit it across coding tasks in Section 3. The survey organizes these reward choices, highlights common patterns and failures, and serves as both a reference and a practical design guide for RL in software engineering.

2 Background

2.1 Reward Methods

In RL for software engineering (SE), rewards map task goals to optimization signals, reflecting objectives such as correctness, coverage, security,

readability, and human preference. We distinguish three source families and five operational signal types. The first family is *verifiable tool-grounded rewards*, which derive from deterministic external checks such as compilation, execution outcomes, coverage, or other tool-based analyzers. In this family, **Execution** denotes correctness signals from compilation and tests, while **Coverage** denotes exploration-oriented signals such as line, branch, or mutation coverage. The second family is *reference-based proxy rewards*, which compare model outputs against gold artifacts or labeled targets. In this family, **Similarity** refers to lexical, structural, or semantic comparison metrics such as BLEU, CodeBLEU, AST similarity, or BERTScore, while **Retrieval** refers to ranking-oriented metrics such as MRR, MAP, NDCG, or Precision@k in retrieval and localization settings. The third family is *learned evaluators*, which provide reward signals through a learned scoring function rather than a deterministic oracle. In this family, **Preference** denotes human or model-aligned judgments over qualities such as usefulness, readability, or reasoning quality. Process reward models and learned critics used as reward proxies also belong to this family when they are used to provide optimization signals.

These reward families also vary along two cross-cutting axes that affect calibration, drift, and training stability (Roijers et al., 2013). First, reward sources differ in *model dependence*: tool-grounded rewards are deterministic and model-independent, learned evaluators are model-dependent and may drift as models change (Gao et al., 2023), while reference-based rewards fall between these extremes, since static metrics such as BLEU and CodeBLEU are model-independent whereas learned metrics such as BERTScore are model-dependent. Second, they differ in *adaptivity*: most works use static reward schedules, while adaptive designs condition one signal on another. Posterior-GRPO (Fan et al., 2025), for example, applies preference rewards only after execution succeeds, so preference refines correct code rather than compensating for failures.

We use **Execution**, **Similarity**, **Retrieval**, **Coverage**, and **Preference** as shorthand labels throughout the paper.

RL Optimization Methods. Most works fine-tune policies with on-policy methods such as PPO (Schulman et al., 2017), whose clipped

Table 1: Taxonomy and decision matrix for surveyed RL-for-SE tasks, summarizing benchmarks, task conditions, primary rewards, aggregation strategies, and representative works.

Task family	Benchmarks	Task condition	Primary reward (+ shaping)	Recommended aggregation	Representative works
NL2Code	HumanEval(+), APPS, MBPP, EvalPlus, LiveCodeBench	Executable, cheap oracle	Execution; similarity / process-level shaping	Constrained; correctness filter	CodeRL, PPOCoder, RLEF, PRL-Coder, Posterior-GRPO
Code Completion	RepoBench, CrossCodeEval	Partial or non-runnable target	Similarity, retrieval-based proxy	Normalized scalarization	RLCoder, IRCOCO, aiXcoder-7B-v2
Code2Code	CodeXGLUE	Executable	Execution	Constrained; correctness filter	CoTran, EffiReasonTrans
Code2NL	CodeSearchNet, TL-CodeSum	Reference-based evaluation	Similarity	Normalized scalarization	Hybrid2Seq+DRL, HAN, TAG
Code Retrieval	CodeSearchNet, CoSQA	Ranking-based success	Retrieval; similarity as auxiliary	Single signal preferred	CoaCor, Cosoch, QueCos
Issue Resolution	SWE-bench Lite, SWE-bench Verified	Executable, expensive oracle	Execution; gated by retrieval, similarity, or preference	Cost-aware constrained	SWE-Gym, DeepSWE, SWE-Swiss, R2E-Gym
Bug Localization	Defects4J variants	Ranking-based success	Retrieval/ execution based coverage	Single signal preferred	RLocator, RecBi
Program Repair	Defects4J, QuixBugs	Executable	Execution; preference shaping	Constrained; correctness filter	RePair, Repair-R1, Repairity
Vulnerability det. / repair	Big-Vul, Devign, CVEFixes	Limited or non-executable	Similarity / retrieval + verifier	Normalized scalarization	ProRLearn, Vul-R2, SecRepair, RLFD
Test Gen. / Fuzzing	HumanEval-Tests, V8, MBPP-Tests, Test262, SpiderMonkey	Executable, exploration-oriented	Execution; coverage for fuzzing	Constrained, with coverage shaping	AceCoder, CURE, UTRL, CovRL-Fuzz, RLF, REINAM

objective stabilizes updates under sparse, high-variance rewards. Recent alternatives pursue the same goal differently: *GRPO*-style (Shao et al., 2024) updates normalize rewards over sampled completion groups, reducing reliance on a value critic, while *RLVR* (Reinforcement Learning with Verifiable Rewards) (Lambert et al., 2025) integrates deterministic feedback loops for heterogeneous, multi-objective optimization.

Coding Task Types Covered in This Survey. We discuss reward design across several SE tasks, including: code (i) *generation*, (ii) *understanding*, (iii) *maintenance*, (iv) *security*, and (v) *testing*.

We summarize the surveyed corpus along task families, benchmarks, structural conditions, and predominant rewards in Table 1, with per-paper reward-signals detailed in Table 2.

2.2 Survey Methodology

Literature Collection. We conducted a broad literature search using Google Scholar, arXiv, and Semantic Scholar. Because reward design for RL in SE spans both ML and software engineering, the search covered major ML/RL venues such as NeurIPS, ICLR, ICML, and AAAI, as well as

relevant SE and programming-language venues. To capture fast-moving work, recent arXiv preprints were included, and citation trails were followed through reference lists and *cited-by* links.

The search used keyword combinations around RL and software tasks, including “reinforcement learning”, “reward design”, “code generation”, “code maintenance”, “vulnerability detection and repair”, and “test generation”. The survey covers papers from 2018 to early 2026; given the pace of new preprints, some very recent manuscripts may have been missed.

Inclusion and Exclusion Criteria. We included papers that explicitly define, learn, or use reward signals for RL-driven SE tasks, with enough detail to show how the reward guides optimization. We excluded papers that apply RL to software problems but use only a generic task-success signal without analyzing the reward design. Each candidate was screened through its abstract, method, and experimental sections, focusing on reward formulation, granularity, and multi-signal aggregation. The resulting paper set forms the basis for the taxonomy.

Table 2: Overview of RL in SE tasks and their reward types. ✓ indicates the paper uses the reward type; ✗ indicates it does not. This table reports reward-signal assignments for each surveyed paper, organized by code-task family. Each row marks the reward families used by a paper, while Table 1 summarizes the same assignments at the task-family level.

	Code Task	Paper	Execution	Similarity	Retrieval	Coverage	Preference
Generation	NL2Code	CodeRL (Le et al., 2022)	✓	✗	✗	✗	✗
		PPOCoder (Shojaee et al., 2023)	✓	✓	✗	✗	✗
		B-Coder (Yu et al., 2024)	✓	✗	✗	✗	✗
		RLEF (Gehring et al., 2025)	✓	✗	✗	✗	✗
		PRM-Based (Dai et al., 2024)	✓	✗	✗	✗	✗
		PRLCoder (Ye et al., 2025)	✓	✗	✗	✗	✗
		Posterior-GRPO (Fan et al., 2025)	✓	✗	✗	✗	✓
	Code Completion	IRCoCo (Li et al., 2024)	✗	✓	✗	✗	✗
		RLCoder (Wang et al., 2025b)	✗	✗	✓	✗	✗
		aiXcoder-7B-v2 (Li et al., 2025)	✗	✗	✗	✗	✓
Understanding	Code2Code	CoTran (Jana et al., 2024)	✓	✗	✗	✗	✗
		EffiReasonTrans (Wang et al., 2026)	✓	✗	✗	✗	✓
	Code2NL	Hybrid2Seq+Attn+DRL (Wan et al., 2018)	✗	✓	✗	✗	✗
		HAN (Wang et al., 2020)	✗	✓	✗	✗	✗
		TAG (Cai et al., 2020)	✗	✓	✗	✗	✗
	Code Retrieval	CoaCor (Yao et al., 2019)	✗	✗	✓	✗	✗
		Cosoch (Li et al., 2019)	✗	✗	✓	✗	✗
		QueCos (Wang et al., 2022)	✗	✓	✓	✗	✗
Maintenance	Issue Resolution	SWE-RL (Wei et al., 2025)	✗	✓	✗	✗	✗
		SWE-GYM (Pan et al., 2025)	✓	✗	✗	✗	✗
		DeepSWE (Luo et al., 2025)	✓	✗	✗	✗	✓
		SWE-Swiss (He et al., 2026)	✓	✗	✗	✗	✗
		Satori-SWE (Zeng et al., 2025a)	✓	✗	✓	✗	✓
		R2E-Gym (Jain et al., 2025)	✓	✗	✗	✗	✗
		Self-play SWE-RL (Wei et al., 2026)	✓	✗	✗	✗	✗
		SWE-PRM (Gandhi et al., 2025)	✓	✗	✗	✗	✓
	Code Refactoring	Refactor-RL (Palit and Sharma, 2025)	✓	✗	✗	✗	✗
	Code Review	CodeMentor (Nashaat and Miller, 2024)	✗	✗	✗	✗	✓
		CRScore++ (Kapadnis et al., 2025)	✗	✗	✗	✗	✓
	Code Optimization	Pearl (Lamouri et al., 2025)	✓	✗	✗	✗	✗
	Bug Localization	RecBi (Chen et al., 2020)	✓	✗	✗	✗	✗
		RLocator (Chakraborty et al., 2024)	✗	✗	✓	✗	✗
	Program Repair	RePair (Zhao et al., 2024)	✓	✗	✗	✗	✗
		Repair-R1 (Hu et al., 2025)	✓	✗	✗	✗	✗
		Repairity (Tang et al., 2025)	✓	✗	✗	✗	✓
Security	Vulnerability Detection	ProRLearn (Ren et al., 2024)	✗	✓	✗	✗	✗
		RLFD (Jiang et al., 2025)	✗	✗	✓	✗	✗
	Vulnerability Repair	SecureCode (Islam et al., 2024a)	✗	✓	✗	✗	✗
		SecRepair (Islam et al., 2024b)	✗	✓	✗	✗	✗
		Vul-R2 (Wen et al., 2025)	✗	✓	✗	✗	✓
Testing	Test Generation	SecCoderX (Wu et al., 2026)	✗	✓	✗	✗	✓
		AceCoder (Zeng et al., 2025b)	✓	✗	✗	✗	✓
		CURE (Wang et al., 2025c)	✓	✗	✗	✗	✗
	Fuzzing	UTRL (Lee et al., 2026)	✓	✗	✗	✗	✗
		REINAM (Wu et al., 2019)	✓	✗	✗	✗	✗
		RLF (Su et al., 2022)	✓	✗	✗	✓	✗
Misc	Text2SQL	CovRL-Fuzz (Eom et al., 2024)	✗	✗	✗	✓	✗
		CogniSQL-R1-Zero (Gajjar et al., 2025)	✓	✗	✗	✗	✗
		Think2SQL (Papicchio et al., 2026)	✓	✗	✗	✗	✗
	Visual2Code	Reasoning-SQL (Pourreza et al., 2025)	✓	✓	✗	✗	✓
		UI2Code ^N (Yang et al., 2025)	✓	✓	✗	✗	✗
		VinciCoder (Zhao et al., 2026)	✓	✓	✗	✗	✗
		ChartMaster (Tan et al., 2025)	✓	✓	✗	✗	✗

3 Code-Centric Tasks

We organize the survey around five categories of code-centric tasks: code generation, code understanding, code maintenance, code security, and code testing. As summarized in Table 1, these tasks differ in their reward conditions, available benchmarks, primary reward signals, and aggregation strategies. The following subsections examine how these structural differences shape reward formulation within each category.

3.1 Code Generation

3.1.1 NL2Code

For NL2Code, the model maps natural-language specifications to executable programs, so correctness is ultimately defined by program behavior. As summarized in Table 1, this makes NL2Code an executable task with relatively cheap oracle feedback, where *execution* provides the most faithful reward whenever compiler or unit-test feedback is available.

CodeRL (Le et al., 2022), B-Coder (Yu et al., 2024), and RLEF (Gehring et al., 2025) all follow this design, differing mainly in optimization: CodeRL uses graded penalties and a learned critic, B-Coder adopts a Dueling DQN formulation, and RLEF applies execution-based feedback to multi-turn refinement. Thus, their shared reward-design choice is functional correctness.

The main challenge is that execution feedback is often sparse. PPOCoder (Shojaee et al., 2023) mitigates this by adding AST and DFG-based *similarity* to a reference solution as structural shaping. Other work introduces intermediate supervision while keeping final tests as the anchor: PRM-based RL (Dai et al., 2024) combines unit-test rewards with learned process rewards, and PRL-Coder (Ye et al., 2025) derives step-level feedback from compiler-verified mutation and refactoring. Posterior-GRPO (Fan et al., 2025) adds a *preference* signal from a reasoning model, but conditions it on execution success so that reasoning is rewarded only when it leads to correct code.

Takeaway. Ground NL2Code reward design in *execution* whenever an oracle is available. Use *similarity* and process rewards only as shaping signals to reduce sparsity, while keeping them aligned with final test outcomes. Reserve *preference* rewards for secondary goals such as reasoning quality, and condition them on execution correctness when possible.

3.1.2 Code Completion

For code completion, the model generates continuations for partially observed programs, from local autocomplete to repository-level suggestions. Because these targets are often fragments rather than runnable programs, *execution* is usually too sparse or costly as the main reward. Table 1 captures this partial-target condition, under which reward design relies instead on reference-based proxies, *retrieval* signals, or learned *preference* objectives.

RLCoder (Wang et al., 2025b) follows this direction most directly, using a *retrieval*-based proxy reward defined by how much retrieved context reduces the perplexity of the ground-truth completion. IRCOCO (Li et al., 2024) addresses the same problem from a different angle, training a critic to regress reference-based metrics such as BLEU and Edit Similarity, thereby converting sparse sequence-level scores into token-level feedback. CoLT, the training framework behind aiXcoder-7B-v2 (Li et al., 2025), pushes this one step further by constructing synthetic preference pairs from reference code and optimizing them with DPO, replacing direct metric optimization with pairwise *preference* learning.

Takeaway. For code completion, use reference-grounded proxy rewards when targets are partial or non-executable. Depending on the setting, these proxies may be *retrieval*-based, similarity-based, or learned as *preference* objectives from reference code.

3.1.3 Code2Code

For code translation and migration, the goal is to convert programs across languages or APIs while preserving behavior. Since behavioral preservation can often be verified through compilation and testing, *execution* is the natural primary reward. CoTran (Jana et al., 2024) uses compilation feedback for syntactic validity and test-based feedback for functional equivalence, while EffiReasonTrans (Wang et al., 2026) likewise centers reward design on unit-test pass rate, adding only a length-based *preference* term for efficiency.

Takeaway. For code translation, use *execution* as the primary reward whenever behavioral equivalence can be tested. *Preference* terms can be added for secondary goals such as efficiency, but should remain guided by correctness.

3.2 Code Understanding

3.2.1 Code2NL

Code2NL generates natural-language descriptions of programs, such as summaries or docstrings. Because success is usually evaluated by agreement with reference text rather than by behavioral verification, RL methods in this setting typically rely on sequence-level *similarity* rewards. Hybrid2Seq+Attn+DRL (Wan et al., 2018) and HAN (Wang et al., 2020) both follow this design, assigning a terminal BLEU reward after the full description is generated. TAG (Cai et al., 2020) extends the same idea through hierarchical RL, using BLEU-4 or ROUGE while combining the RL objective with MLE during training. Across these methods, the main limitation is that the rewards are sparse and tend to capture textual overlap more directly than semantic correctness.

Takeaway. For Code2NL, use *similarity* as the primary reward, since success is defined by how closely the generated description matches the reference text.

3.2.2 CodeRetrieval

In code retrieval, the objective is to rank relevant code snippets for a given query, so reward design is most naturally tied to retrieval quality. CoaCor (Yao et al., 2019) makes this explicit: instead of optimizing for human-like annotations or BLEU overlap, it uses MRR from a retrieval model as the reward signal. Although similarity scoring appears inside that model, the optimized objective is clearly *retrieval*-based. Cosoch (Li et al., 2019) follows the same direction by directly optimizing NDCG@10 for session-level search. QueCos (Wang et al., 2022) adds BLEU-4 *similarity* alongside MRR, but its reward design remains centered on ranking performance. *Similarity* helps as auxiliary shaping, but it is weaker than retrieval metrics as a measure of success.

Takeaway. For code retrieval, use *retrieval* rewards such as MRR or NDCG as the primary objective. Add *similarity* only as auxiliary shaping, since higher textual overlap does not necessarily improve ranking quality.

3.3 Code Maintenance

3.3.1 Issue Resolution

Issue resolution is a long-horizon repository task in which the model must localize the bug, edit the

relevant files, and produce a patch that satisfies repository-level constraints. Because success is ultimately determined by whether the patch works, *execution* is the most faithful reward. SWE-Gym (Pan et al., 2025), DeepSWE (Luo et al., 2025), and SWE-Swiss (He et al., 2026) therefore stay centered on test-based correctness, with DeepSWE additionally using verifier-based *preference* signals for test-time selection.

This corresponds to the issue-resolution setting in Table 1: *execution* remains the most faithful reward, but repository-level testing cost motivates cheaper proxy signals for selection and stabilization. SWE-RL (Wei et al., 2025) replaces *execution* with a rule-based *similarity* reward, while Satori-SWE (Zeng et al., 2025a) uses a staged design that combines *retrieval* and learned *preference* signals, with *execution* optionally used during selection. R2E-Gym (Jain et al., 2025) occupies a middle ground, pairing *execution*-based testing with *execution*-free verifiers.

Takeaway. For issue resolution, use *execution* as the primary reward whenever repository tests are available. Add *similarity*, *retrieval*, or *preference* only as auxiliary signals for scalability, ranking, or stability.

3.3.2 Bug Localization

RL-based bug localization typically optimizes *execution*-derived or ranking-based signals rather than direct patch correctness. Table 1 captures this distinction: ranking-based localization aligns with *retrieval* rewards, while test or trace-mediated localization relies on *execution*-derived coverage signals. RecBi (Chen et al., 2020) approaches compiler bug isolation through *execution*-derived coverage shaping: it uses RL to guide the generation of passing test programs, then compares their execution traces with the failing test to localize buggy files. Here, coverage acts as the shaping signal, while test outcomes provide the *execution* oracle. By contrast, RLocator (Chakraborty et al., 2024) formulates bug localization directly as a ranking problem and optimizes *retrieval* quality through metrics such as MRR and MAP.

Takeaway. For bug localization, prefer *retrieval* rewards when the objective is ranking suspicious code elements, and opt for *execution*-derived signals for settings where localization is mediated by test outcomes or coverage traces.

3.3.3 Program Repair

Program repair naturally favors *execution*-based rewards because candidate patches can often be checked by compilation and tests. The main difficulty is sparse pass/fail feedback. RePair (Zhao et al., 2024) addresses this with a critic trained on execution-derived signals, while Repair-R1 (Hu et al., 2025) adds rule-based rewards for test generation, repair, and output format. Repairity (Tang et al., 2025) takes a different route by using a learned *preference* signal to rank candidate repairs, with execution retained mainly as a correctness filter during data collection.

Takeaway. For program repair, use *execution* as the primary reward when patches are testable. Add critics, rule-based shaping, or *preference* signals only to densify feedback or stabilize optimization.

3.4 Code Security

3.4.1 Vulnerability Detection and Repair

Vulnerability detection and repair share a limited-oracle constraint: reliable *execution* signals are often unavailable or weak. Detection methods therefore rely on label-based proxies. ProRLearn (Ren et al., 2024) uses reference-based *similarity*, while RLFD (Jiang et al., 2025) uses ranking rewards over vulnerable-line overlap and top-ranked accuracy, making it primarily *retrieval*-based. In both cases, supervision comes from annotated labels rather than executable security outcomes, matching the limited-oracle regime in Table 1. On repair side, SecureCode (Islam et al., 2024a) and SecRepair (Islam et al., 2024b) remain fully *similarity*-based, combining CodeBLEU and BERTScore against the reference patch. VulR2 (Wen et al., 2025) and SecCoderX (Wu et al., 2026) replace reference matching with learned verifier rewards for repair judgment and secure generation, respectively.

Takeaway. For vulnerability repair, use *similarity* when execution signals are unavailable; add *retrieval* for fine-grained localization, and use *preference* signals when weak reference overlap makes similarity insufficient.

3.5 Code Testing

3.5.1 Test Generation and Fuzzing

Test generation and fuzzing both synthesize inputs to probe program behavior, but they optimize for




Figure 1: Number of papers per reward type by year.

different outcomes: test generation seeks valid, discriminative tests, while fuzzing emphasizes exploration and failure discovery. In both settings, rewards are naturally grounded in *execution*. For test generation, CURE (Wang et al., 2025c) uses execution outcomes from generated programs and tests to train a co-evolving test generator, and UTRL (Lee et al., 2026) derives validity and discrimination rewards from pass/fail behavior. AceCoder (Zeng et al., 2025b) similarly supports binary execution rewards or learned *preference* rewards from pass-rate comparisons.

Fuzzing follows the same execution-centered logic, but adds stronger exploration. RLF (Su et al., 2022) combines vulnerability discovery with *coverage*, and CovRL-Fuzz (Eom et al., 2024) further rewards rarely covered regions while penalizing invalid inputs. Both sub-areas align with the executable-with-exploration setting in Table 1: execution provides the correctness anchor, while coverage guides exploration.

Takeaway. For test generation and fuzzing, use *execution* as the primary reward. Add *coverage* when exploration matters, especially in fuzzing, and use learned *preference* signals to refine execution feedback.

4 Discussion

4.1 From taxonomy to decision guide

This section turns the surveyed taxonomy into a four-step reward-design guide. Figure 1 shows that the design space has grown sharply since 2022, with execution-based rewards remaining dominant and preference-based rewards expanding fastest.

(D1) Select the reward source by oracle availability. Reward source selection should start from the strongest oracle the task provides, since fidelity depends on what can be verified reliably and at what cost. When compiler, test, or

Table 3: Aggregation strategies and adoption status in surveyed RL-for-SE work.

Strategy	Status	Surveyed example
Fixed scalarization	Common	<i>PPOCoder</i> : fixed weighted sum of execution, AST, and DFG rewards
Normalized scalarization	Rarely reported explicitly	Keeps execution, similarity, and preference scores comparable before combining them
Pareto-aware aggregation	Absent	Remains an open direction for competing SE objectives such as correctness, efficiency, and security
Adaptive weighting	Nascent	<i>Posterior-GRPO</i> conditions the preference reward on execution success
Constrained aggregation	Implicit	<i>AceCoder</i> : test outcomes constrain filtering and preference-pair construction

execution feedback is cheap, use *execution* as the reward anchor, as in CodeRL (Le et al., 2022), RLEF (Gehring et al., 2025), and Repair-R1 (Hu et al., 2025). When execution is costly, as in long-horizon issue resolution with *Satori-SWE* (Zeng et al., 2025a), use it selectively as a gate, while *retrieval* or *preference* signals guide selection. When execution is unavailable or weak, choose the closest aligned proxy: *retrieval* for ranking and localization, as in RLocator (Chakraborty et al., 2024); *similarity* for reference-anchored outputs, as in SecRepair (Islam et al., 2024b); and *preference* for subjective judgments, as in Posterior-GRPO (Fan et al., 2025). In proxy-based settings, reward fit should be checked directly: verify that the proxy tracks downstream success, rather than treating it as reliable by default.

(D2) Use hybrid rewards as a correctness-preserving decision guide. Hybrid rewards are useful because execution is faithful but often sparse, delayed, or costly, whereas proxy rewards are denser but may not preserve behavior. The guiding principle is to start with the most faithful available signal and add auxiliaries when needed. When execution is available but sparse, keep it primary and use similarity as shaping, as in *PPOCoder* (Shojaee et al., 2023). When tasks involve coupled failure modes, give each signal a distinct role: execution for validity, retrieval for file or context selection, and preference for decision quality, as in *Satori-SWE*. When execution is unavailable, proxies such as similarity can be used, as in *SecureCode* (Islam et al., 2024a), but they should not be treated as correctness evidence. Finally, when multiple signals are combined, avoid weighted sums that let proxy scores offset failed execution; instead, gate, calibrate, or filter proxies with execution or verifier feedback, as in *Posterior-GRPO*, and *AceCoder* (Zeng et al., 2025b). Table 3 summarizes these regimes and their adoption.

(D3) Set granularity by reliable observability.

Reward granularity should match the finest level at which feedback is trustworthy. If correctness is only checkable after a full artifact is produced, use program-level rewards, as in RLEF (Gehring et al., 2025) and Repair-R1 (Hu et al., 2025). If intermediate states are reliably verifiable, use step or line-level rewards to improve credit assignment, as in PRL-Coder (Ye et al., 2025) and PRM-based scoring (Dai et al., 2024). If success is defined over rankings, sets, or multi-step decisions, use trajectory-level rewards, as in RLocator (Chakraborty et al., 2024) and RLFD (Jiang et al., 2025). If terminal feedback is faithful but sparse, use a hybrid design that keeps execution as the anchor and adds shaping or critic signals, as in *PPOCoder* (Shojaee et al., 2023) and *CodeRL* (Le et al., 2022).

(D4) Align evaluation methodology with the reward family. Each reward family shapes validation choices, including benchmarks, optimizers, and evaluation metrics:

- **Execution:** executable-oracle benchmarks such as HumanEval, SWE-bench, and Defects4J; PPO with KL regularization or learned critics; pass@ k or test-success rates (Le et al., 2022; Gehring et al., 2025).
- **Similarity:** reference-anchored benchmarks such as CodeSearchNet and Big-Vul; BLEU, ROUGE, or BERTScore, often combined with MLE for stability (Cai et al., 2020).
- **Retrieval:** ranking benchmarks such as CodeSearchNet and CoSQA; MRR, MAP, or NDCG as both reward and evaluation metric (Yao et al., 2019; Chakraborty et al., 2024).
- **Preference:** completion benchmarks such as RepoBench and CrossCodeEval; DPO or GRPO-style optimization; downstream task accuracy used to check reward-model reliability (Li et al., 2025).

Table 4: Recurring failure modes in RL-for-SE reward design.

Failure mode	Triggered by	Surveyed example
Proxy misalignment	Similarity rewards in place of execution	<i>SecureCode</i> , <i>SecRepair</i> : rely on CodeBLEU and BERTScore; textual overlap need not reflect correctness
Sparsity and oracle cost	Execution rewards over long rollouts	<i>RLEF</i> , <i>SWE-Gym</i> , <i>DeepSWE</i> : sparse, high-variance test feedback; densify with process-level shaping or learned verifier
Reward hacking and calibration drift	Learned preference rewards	<i>Repairity</i> , <i>aiXcoder-7B-v2</i> : use AI-feedback or DPO over generated pairs; need execution/downstream checks for reliability

4.2 Challenges and future directions

(C1) Recurring failure modes. The surveyed corpus reveals three connected failure modes in reward design, summarized in Table 4. First, dense proxy rewards improve learnability but can drift from behavior: *SecureCode* and *SecRepair* (Islam et al., 2024a,b) use CodeBLEU/BERTScore-style rewards, yet textual overlap need not imply correct or secure code. Second, execution rewards are behaviorally faithful but can be sparse and costly over long rollouts, as in *RLEF* (Gehring et al., 2025), *SWE-Gym* (Pan et al., 2025), and *DeepSWE* (Luo et al., 2025). Third, learned preference rewards offer flexible feedback but introduce risks of reward hacking and calibration drift, as in *Repairity* (Tang et al., 2025) and *aiXcoder-7B-v2* (Li et al., 2025). A key future direction is therefore to make dense rewards more verifiable by grounding them in deterministic signals such as tests, static analyzers, or intermediate execution states.

(C2) Aggregation strategies remain under-explored. Table 3 shows that aggregation is still treated narrowly in surveyed RL-for-SE work. Fixed scalarization is common; *PPO-Coder* (Shojaee et al., 2023), for example, combines execution, AST, and DFG similarity as a fixed sum. This design is simple but fragile: when signals differ in scale, noise, and reliability, a dense proxy can dominate sparse correctness feedback. More principled alternatives remain limited, with partial examples such as *AceCoder* (Zeng et al., 2025b) for constrained aggregation and *Posterior-GRPO* (Fan et al., 2025) for adaptive weighting. By contrast, normalized scalarization is rarely reported explicitly, and Pareto-aware aggregation remains absent despite its relevance to multi-objective RL (Roijers et al., 2013). This gap matters because SE tasks often balance competing goals. Future work should make normalization, weight sensitivity, and constraint handling explicit rather than hidden implementation choices.

(C3) Toward a Reward Engineering Reporting Standard.

Existing works often report downstream performance without enough evidence about reward quality, making it difficult to separate genuine task progress from exploitation of a particular reward formulation. We therefore suggest a *Reward Engineering Reporting Standard* with two linked parts. (a) *Reward-weighting checklist*: authors should identify the task-faithful anchor reward, normalize heterogeneous components before aggregation, treat correctness-sensitive objectives as constraints so auxiliary rewards cannot compensate for failed checks, and report sensitivity to reward weights. (b) *Reward-quality diagnostics*: when proxy and oracle signals are both available, authors should report their rank correlation, monitor reward–accuracy decoupling during training, and ablate each reward component against the downstream metric rather than only against the aggregate reward. Together, these practices make reward engineering more auditable and cross-system comparisons more reliable.

5 Conclusion

This survey treats reward engineering for RL in software engineering as a first-class design problem. We organize prior work along three axes, reward source, granularity, and aggregation, and distill surveyed practice into a four-step design guide. A recurring pattern is that execution-grounded rewards provide the most dependable supervision, while similarity, retrieval, and preference rewards are most useful as shaping signals under sparse feedback rather than substitutes for correctness. This pattern also explains common failures, including proxy misalignment, reward sparsity, and unstable aggregation, and motivates principled aggregation, drift-resistant learned evaluators, and explicit reward-design reporting. Together, these directions can make RL-for-SE systems more reliable, comparable, and reproducible.

Limitations

While this survey aims to provide a comprehensive view of reward design in reinforcement learning for software engineering tasks, the area is evolving rapidly and some very recent or adjacent works may have been missed. We intentionally scope our review to *reward schemes used to train or adapt policies* for SE tasks (e.g., code generation, repair, translation, retrieval/repo-level completion, testing, and localization), and do not attempt to cover reward design for large-scale pre-training or general instruction tuning. We prioritize works where reward formulation is a central contribution. Works that apply RL only as a minor component, or where the reward is implicit/standard and not analyzed, are omitted. To help keep pace with new literature, we plan to maintain an accompanying repository that can be incrementally updated as the field progresses. Lastly, this paper is a survey or review of existing research on reward design for reinforcement learning in software engineering tasks. Accordingly, it does not introduce new models or experimental evaluations; instead, it focuses on organizing, summarizing, and critically analyzing prior work in this area.

References

- Mridul Agarwal, Vaneet Aggarwal, and Tian Lan. 2022. Multi-Objective Reinforcement Learning with Non-Linear Scalarization. In *Proceedings of the 21st International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2022)*.
- Ruichu Cai, Zhihao Liang, Boyan Xu, Zijian Li, Yuexing Hao, and Yao Chen. 2020. TAG : Type Auxiliary Guiding for Code Comment Generation. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*.
- Partha Chakraborty, Mahmoud Alfeldel, and Meiyappan Nagappan. 2024. RLocator: Reinforcement Learning for Bug Localization. *IEEE Transactions on Software Engineering*.
- Junjie Chen, Haoyang Ma, and Lingming Zhang. 2020. Enhanced Compiler Bug Isolation via Memoized Search. In *Proceedings of the 35th IEEE/ACM International Conference on Automated Software Engineering*, pages 78–89.
- Ning Dai, Zheng Wu, Renjie Zheng, Ziyun Wei, Wenlei Shi, Xing Jin, Guanlin Liu, Chen Dun, Liang Huang, and Lin Yan. 2024. Process Supervision-Guided Policy Optimization for Code Generation. *arXiv preprint arXiv:2410.17621*.
- Karin de Langis, Ryan Koo, and Dongyeop Kang. 2024. Dynamic Multi-Reward Weighting for Multi-Style Controllable Generation. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing (EMNLP 2024)*.
- Jueon Eom, Seyeon Jeong, and Taekyoung Kwon. 2024. Fuzzing Javascript Interpreters with Coverage-Guided Reinforcement Learning for LLM-Based Mutation. In *Proceedings of the 33rd ACM SIGSOFT International Symposium on Software Testing and Analysis*, pages 1656–1668.
- Lishui Fan, Yu Zhang, Mouxiang Chen, and Zhongxin Liu. 2025. Posterior-GRPO: Rewarding Reasoning Processes in Code Generation. *arXiv preprint arXiv:2508.05170v1*.
- Kushal Gajjar, Harshit Sikchi, Arpit Singh Gautam, Marc Hammons, and Saurabh Jha. 2025. CogniSQL-R1-Zero: Lightweight Reinforced Reasoning for Efficient SQL Generation. *arXiv preprint arXiv:2507.06013*.
- Shubham Gandhi, Jason Tsay, Jatin Ganhotra, Kiran Kate, and Yara Rizk. 2025. When agents go astray: Course-correcting SWE agents with PRMs. In *Workshop on Scaling Environments for Agents*.
- Leo Gao, John Schulman, and Jacob Hilton. 2023. Scaling Laws for Reward Model Overoptimization. In *Proceedings of the 40th International Conference on Machine Learning, ICML 2023*.
- Jonas Gehring, Kunhao Zheng, Jade Copet, Vegard Mella, Taco Cohen, and Gabriel Synnaeve. 2025. RLEF: Grounding code LLMs in execution feedback with reinforcement learning. In *Forty-second International Conference on Machine Learning*.
- Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Peiyi Wang, Qihao Zhu, Runxin Xu, Ruoyu Zhang, Shirong Ma, Xiao Bi, and 1 others. 2025. DeepSeek-R1 Incentivizes Reasoning in LLMs Through Reinforcement Learning. *Nature*, 645(8081).
- Zhenyu He, Qingping Yang, Wei Shen, Xiaojian Zhong, Kechi Zhang, Chenxin An, Wenlei Shi, Tianle Cai, Di He, Jiaze Chen, and Jingjing Xu. 2026. SWE-Swiss: A Multi-Task Fine-Tuning and RL Recipe for High-Performance Issue Resolution. In *Findings of the Association for Computational Linguistics: ACL 2026*. Association for Computational Linguistics.
- Haichuan Hu, Xiaochen Xie, and Qunjun Zhang. 2025. Repair-R1: Better Test Before Repair. *arXiv preprint arXiv:2507.22853*.
- Sinan Ibrahim, Mostafa Mostafa, Ali Inadi, Hadi Salloum, and Pavel Osinenko. 2024. Comprehensive Overview of Reward Engineering and Shaping in Advancing Reinforcement Learning Applications. *IEEE Access*.

- Nafis Tanveer Islam, Mohammad Bahrami Karkevandi, and Peyman Najafirad. 2024a. Code Security Vulnerability Repair using Reinforcement Learning with Large Language Models. *arXiv preprint arXiv:2401.07031*.
- Nafis Tanveer Islam, Joseph Khoury, Andrew Seong, Mohammad Bahrami Karkevandi, Gonzalo De La Torre Parra, Elias Bou-Harb, and Peyman Najafirad. 2024b. LLM-Powered Code Vulnerability Repair with Reinforcement Learning and Semantic Reward. *arXiv preprint arXiv:2401.03374*.
- Naman Jain, Jaskirat Singh, Manish Shetty, Tianjun Zhang, Liang Zheng, Koushik Sen, and Ion Stoica. 2025. R2E-Gym: Procedural Environment Generation and Hybrid Verifiers for Scaling Open-Weights SWE Agents. In *Second Conference on Language Modeling*.
- Prithwish Jana, Piyush Jha, Haoyang Ju, Gautham Kishore, Aryan Mahajan, and Vijay Ganesh. 2024. Cotran: An LLM-Based Code Translator using Reinforcement Learning with Feedback from Compiler and Symbolic Execution. In *27th European Conference on Artificial Intelligence*.
- Yuan Jiang, Zhichen Qu, Christoph Treude, Xiaohong Su, and Tiantian Wang. 2025. Enhancing Fine-Grained Vulnerability Detection with Reinforcement Learning. *IEEE Transactions on Software Engineering*.
- Manav Nitin Kapadnis, Atharva Naik, and Carolyn Rose. 2025. CRScore++: Reinforcement Learning with Verifiable Tool and AI Feedback for Code Review. *arXiv preprint arXiv:2506.00296*.
- Nathan Lambert, Jacob Morrison, Valentina Pyatkin, Shengyi Huang, Hamish Ivison, Faeze Brahman, Lester James Validad Miranda, Alisa Liu, Nouha Dziri, Xinxin Lyu, Yuling Gu, Saumya Malik, Victoria Graf, Jena D. Hwang, Jiangjiang Yang, Ronan Le Bras, Oyvind Tafjord, Christopher Wilhelm, Luca Soldaini, and 4 others. 2025. Tulu 3: Pushing Frontiers in Open Language Model Post-Training. In *Second Conference on Language Modeling*.
- Djamel Rassem Lamouri, Iheb Nassim Aouadj, Smail Kourta, and Riyadh Baghdadi. 2025. Pearl: Automatic Code Optimization Using Deep Reinforcement Learning. In *Proceedings of the 39th ACM International Conference on Supercomputing*, pages 959–974.
- Hung Le, Yue Wang, Akhilesh Deepak Gotmare, Silvio Savarese, and Steven Chu Hong Hoi. 2022. CodeRL: Mastering Code Generation through Pretrained Models and Deep Reinforcement Learning. *Advances in Neural Information Processing Systems*, pages 21314–21328.
- Dongjun Lee, Changho Hwang, and Kimin Lee. 2026. Learning to Generate Unit Test via Adversarial Reinforcement Learning. In *The Fourteenth International Conference on Learning Representations*.
- Bolun Li, Zhihong Sun, Tao Huang, Hongyu Zhang, Yao Wan, Ge Li, Zhi Jin, and Chen Lyu. 2024. IRCoco: Immediate Rewards-Guided Deep Reinforcement Learning for Code Completion. *Proceedings of the ACM on Software Engineering*, 1(FSE):182–203.
- Jia Li, Hao Zhu, Huanyu Liu, Xianjie Shi, He Zong, Yihong Dong, Kechi Zhang, Siyuan Jiang, Zhi Jin, and Ge Li. 2025. aiXcoder-7B-v2: Training LLMs to Fully Utilize the Long Context in Repository-level Code Completion. *arXiv preprint arXiv:2503.15301*.
- Wei Li, Shuhan Yan, Beijun Shen, and Yuting Chen. 2019. Reinforcement Learning of Code Search Sessions. In *2019 26th Asia-Pacific Software Engineering Conference (APSEC)*. IEEE.
- Shaoteng Liu, Haoqi Yuan, Minda Hu, Yanwei Li, Yukang Chen, Shu Liu, Zongqing Lu, and Jiaya Jia. 2024. RL-GPT: Integrating Reinforcement Learning and Code-as-policy. *Advances in Neural Information Processing Systems*, 37:28430–28459.
- Yining Lu, Zilong Wang, Shiyang Li, Xin Liu, Changlong Yu, Qingyu Yin, Zhan Shi, Zixuan Zhang, and Meng Jiang. 2026. Learning to optimize multi-objective alignment through dynamic reward weighting. *Transactions of the Association for Computational Linguistics*.
- Michael Luo, Naman Jain, Jaskirat Singh, Sijun Tan, Ameen Patel, Qingyang Wu, Alpav Ariyak, Colin Cai, Shang Zhu Tarun Venkat, Ben Athiwaratkun, and 1 others. 2025. DeepSWE: Training a State-of-the-Art Coding Agent from Scratch by Scaling RL. *Notion Blog*.
- Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fiedjeland, Georg Ostrovski, and 1 others. 2015. Human-level Control Through Deep Reinforcement Learning. *Nature*.
- Mona Nashaat and James Miller. 2024. Towards Efficient Fine-Tuning of Language Models With Organizational Data for Automated Software Review. *IEEE Transactions on Software Engineering*.
- Indranil Palit and Tushar Sharma. 2025. Reinforcement learning vs supervised learning: A tug of war to generate refactored code accurately. In *Proceedings of the 29th International Conference on EASE*. Association for Computing Machinery.
- Jiayi Pan, Xingyao Wang, Graham Neubig, Navdeep Jaitly, Heng Ji, Alane Suhr, and Yizhe Zhang. 2025. Training Software Engineering Agents and Verifiers with SWE-Gym. In *Forty-second International Conference on Machine Learning*.
- Simone Papicchio, Simone Rossi, Luca Cagliero, and Paolo Papotti. 2026. Think2SQL: Blueprinting Reward Density and Advantage Scaling for Effective Text-to-SQL Reasoning. *Transactions on Machine Learning Research*.

- Mohammadreza Pourreza, Shayan Talaei, Ruoxi Sun, Xingchen Wan, Hailong Li, Azalia Mirhoseini, Amin Saberi, and Sercan O Arik. 2025. Reasoning-SQL: Reinforcement Learning with SQL Tailored Partial Rewards for Reasoning-Enhanced Text-to-SQL. In *Second Conference on Language Modeling*.
- Zilong Ren, Xiaolin Ju, Xiang Chen, and Hao Shen. 2024. ProRLearn: Boosting Prompt Tuning-based Vulnerability Detection by Reinforcement Learning. *Automated Software Engineering*.
- Diederik M. Roijers, Peter Vamplew, Shimon Whiteson, and Richard Dazeley. 2013. A survey of multi-objective sequential decision-making. *Journal of Artificial Intelligence Research*, 48:67–113.
- John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. Proximal Policy Optimization Algorithms. *arXiv preprint arXiv:1707.06347*.
- Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan Zhang, YK Li, Yang Wu, and 1 others. 2024. DeepSeekMath: Pushing the Limits of Mathematical Reasoning in Open Language Models. *arXiv preprint arXiv:2402.03300*.
- Tushar Sharma, Maria Kechagia, Stefanos Georgiou, Rohit Tiwari, Indira Vats, Hadi Moazen, and Federica Sarro. 2024. A survey on machine learning techniques applied to source code. *Journal of Systems and Software*, 209:111934.
- Parshin Shojaee, Aneesh Jain, Sindhu Tipirneni, and Chandan K. Reddy. 2023. Execution-based Code Generation using Deep Reinforcement Learning. *Transactions on Machine Learning Research*.
- David Silver, Thomas Hubert, Julian Schrittwieser, Ioannis Antonoglou, Matthew Lai, Arthur Guez, Marc Lanctot, Laurent Sifre, Dharshan Kumaran, Thore Graepel, and 1 others. 2018. A general reinforcement learning algorithm that masters chess, shogi, and go through self-play. *Science*.
- Jianzhong Su, Hong-Ning Dai, Lingjun Zhao, Zibin Zheng, and Xiapu Luo. 2022. Effectively Generating Vulnerable Transaction Sequences in Smart Contracts with Reinforcement Learning-guided Fuzzing. In *Proceedings of the 37th IEEE/ACM International Conference on Automated Software Engineering*.
- Richard S Sutton, Andrew G Barto, and Andrew Barto. 1998. *Reinforcement learning: An introduction*, volume 1. MIT press Cambridge.
- Wentao Tan, Qiong Cao, Chao Xue, Yibing Zhan, Changxing Ding, and Xiaodong He. 2025. ChartMaster: Advancing Chart-to-Code Generation with Real-World Charts and Chart Similarity Reinforcement Learning. *arXiv preprint arXiv:2508.17608*.
- Xunzhu Tang, Jacques Klein, and Tegawendé F Bissyandé. 2025. Boosting Open-Source LLMs for Program Repair via Reasoning Transfer and LLM-Guided Reinforcement Learning. *arXiv preprint arXiv:2506.03921*.
- Yao Wan, Zhou Zhao, Min Yang, Guandong Xu, Haochao Ying, Jian Wu, and Philip S Yu. 2018. Improving Automatic Source Code Summarization via Deep Reinforcement Learning. In *Proceedings of the 33rd ACM/IEEE International Conference on Automated Software Engineering*.
- Chaozheng Wang, Zhenhao Nong, Cuiyun Gao, Zongjie Li, Jichuan Zeng, Zhenchang Xing, and Yang Liu. 2022. Enriching query semantics for code search with reinforcement learning. *Neural Networks*.
- Dong Wang, Hanmo You, Lingwei Zhu, Kaiwei Lin, Zheng Chen, Chen Yang, Junji Yu, Zan Wang, and Junjie Chen. 2025a. A Survey of Reinforcement Learning for Software Engineering. *arXiv preprint arXiv:2507.12483*.
- Wenhua Wang, Yuqun Zhang, Yulei Sui, Yao Wan, Zhou Zhao, Jian Wu, Philip S Yu, and Guandong Xu. 2020. Reinforcement-Learning-Guided Source Code Summarization Using Hierarchical Attention. *IEEE Transactions on Software Engineering*.
- Yanlin Wang, Rongyi Ou, Yanli Wang, Mingwei Liu, Jiachi Chen, Ensheng Shi, Xilin Liu, Yuchi Ma, and Zibin Zheng. 2026. EffiReasonTrans: RL-Optimized Reasoning for Code Translation. *IEEE Transactions on Software Engineering*.
- Yanlin Wang, Yanli Wang, Daya Guo, Jiachi Chen, Ruikai Zhang, Yuchi Ma, and Zibin Zheng. 2025b. RLCode: Reinforcement Learning for Repository-Level Code Completion. In *2025 IEEE/ACM 47th International Conference on Software Engineering*.
- Yinjie Wang, Ling Yang, Ye Tian, Ke Shen, and Mengdi Wang. 2025c. Co-Evolving LLM Coder and Unit Tester via Reinforcement Learning. In *The 39th Annual Conference on Neural Information Processing Systems*.
- Cody Watson, Nathan Cooper, David Nader Palacio, Kevin Moran, and Denys Poshyvanyk. 2022. A Systematic Literature Review on the Use of Deep Learning in Software Engineering Research. *ACM Transactions on Software Engineering and Methodology (TOSEM)*.
- Yuxiang Wei, Olivier Duchenne, Jade Copet, Quentin Carbonneaux, Lingming Zhang, Daniel Fried, Gabriel Synnaeve, Rishabh Singh, and Sida Wang. 2025. SWE-RL: Advancing LLM Reasoning via Reinforcement Learning on Open Software Evolution. In *The 39th Annual Conference on Neural Information Processing Systems*.
- Yuxiang Wei, Zhiqing Sun, Emily McMillin, Jonas Gehring, David W. Zhang, Gabriel Synnaeve, Daniel Fried, Lingming Zhang, and Sida Wang.

2026. Toward Training Superintelligent Software Agents through Self-Play SWE-RL. In *Forty-third International Conference on Machine Learning*.
- Xin-Cheng Wen, Zirui Lin, Yijun Yang, Cuiyun Gao, and Deheng Ye. 2025. Vul-R2: A Reasoning LLM for Automated Vulnerability Repair. In *2025 40th IEEE/ACM International Conference on Automated Software Engineering (ASE)*.
- Tianyi Wu, Mingzhe Du, Yue Liu, Chengran Yang, Terry Yue Zhuo, Jiaheng Zhang, and See-Kiong Ng. 2026. Secure Code Generation via Online Reinforcement Learning with Vulnerability Reward Model. In *ICLR 2026 Workshop on Principled Design for Trustworthy AI - Interpretability, Robustness, and Safety across Modalities*.
- Zhengkai Wu, Evan Johnson, Wei Yang, Osbert Bastani, Dawn Song, Jian Peng, and Tao Xie. 2019. REINAM: reinforcement learning for input-grammar inference. In *Proceedings of the 2019 27th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering*.
- Yanming Yang, Xin Xia, David Lo, and John Grundy. 2022. A survey on deep learning for software engineering. *ACM Computing Surveys (CSUR)*.
- Zhen Yang, Wenyi Hong, Mingde Xu, Xinyue Fan, Weihang Wang, Jiele Cheng, Xiaotao Gu, and Jie Tang. 2025. UI2Code^N: A Visual Language Model for Test-Time Scalable Interactive UI-to-Code Generation. *arXiv preprint arXiv:2511.08195*.
- Ziyu Yao, Jayavardhan Reddy Peddamail, and Huan Sun. 2019. Coacor: Code Annotation for Code Retrieval with Reinforcement Learning. In *The World Wide Web Conference*.
- Yufan Ye, Ting Zhang, Wenbin Jiang, and Hua Huang. 2025. Process-Supervised Reinforcement Learning for Code Generation. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*.
- Zishun Yu, Yunzhe Tao, Liyu Chen, Tao Sun, and Hongxia Yang. 2024. B-Coder: Value-Based Deep Reinforcement Learning for Program Synthesis. In *The Twelfth International Conference on Learning Representations*.
- Guangtao Zeng, Maohao Shen, Delin Chen, Zhenting Qi, Subhro Das, Dan Gutfreund, David Cox, Gregory Wornell, Wei Lu, Zhang-Wei Hong, and 1 others. 2025a. Satori-SWE: Evolutionary Test-Time Scaling for Sample-Efficient Software Engineering. *arXiv preprint arXiv:2505.23604*.
- Huaye Zeng, Dongfu Jiang, Haozhe Wang, Ping Nie, Xiaotong Chen, and Wenhui Chen. 2025b. ACECODER: Acing Coder RL via Automated Test-Case Synthesis. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics*.
- Kaiyan Zhang, Yuxin Zuo, Bingxiang He, Youbang Sun, Runze Liu, Che Jiang, Yuchen Fan, Kai Tian, Guoli Jia, Pengfei Li, and 1 others. 2025. A Survey of Reinforcement Learning for Large Reasoning Models. *arXiv preprint arXiv:2509.08827*.
- Xuanle Zhao, Deyang Jiang, Zhixiong Zeng, Lei Chen, Haibo Qiu, Jing Huang, Yufeng Zhong, Liming Zheng, Yilin Cao, and Lin Ma. 2026. VinciCoder: Unifying Multimodal Code Generation via Coarse-to-fine Visual Reinforcement Learning. In *Proceedings of the 2026 Conference on Empirical Methods in Natural Language Processing*.
- Yuze Zhao, Zhenya Huang, Yixiao Ma, Rui Li, Kai Zhang, Hao Jiang, Qi Liu, Linbo Zhu, and Yu Su. 2024. RePair: Automated Program Repair with Process-based Feedback. In *Findings of the Association for Computational Linguistics: ACL 2024*.

A Adjacent Task Families

Beyond the core code-centric tasks surveyed in Section 3, similar reward-design patterns appear in adjacent code-generation settings. In Text-to-SQL, CogniSQL-R1-Zero (Gajjar et al., 2025), Think2SQL (Papicchio et al., 2026), and Reasoning-SQL (Pourreza et al., 2025) use *execution* rewards for query correctness, with Reasoning-SQL further adding syntax, schema-linking, and LLM-judge signals as auxiliary shaping, paralleling NL2Code in Section 3.1.1. In visual-to-code tasks, UI2Code (Yang et al., 2025), VinciCoder (Zhao et al., 2026), and ChartMaster (Tan et al., 2025) combine visual *similarity* rewards on rendered outputs with *execution* penalties for non-compilable or non-renderable code, mirroring the execution-plus-similarity pattern in code translation discussed in Section 3.1.3. RL-GPT (Liu et al., 2024) extends this pattern to embodied code generation through task-success *execution* rewards with dense vision-language similarity shaping.

B Acknowledgements

AI Assistance Disclosure. We used generative AI tools, including ChatGPT and Claude, for language editing and paraphrasing of author-written text. All content was reviewed and revised by the authors, who take full responsibility for the paper’s correctness and citations.