

Mastering the Craft of Data Synthesis for CodeLLMs

Meng Chen[✉] Philip Arthur Qianyu Feng Cong Duy Vu Hoang
Yu-Heng Hong Mahdi Kazemi Moghaddam Omid Nezami Thien Nguyen
Gioacchino Tangari Duy Vu Thanh Vu Mark Johnson
Krishnaram Kenthapadi Don Dharmasiri Long Duong Yuan-Fang Li[✉]
[✉] {meng.c.chen, yuanfang.li}@oracle.com
Oracle Corporation

Abstract

Large language models (LLMs) have shown impressive performance in *code* understanding and generation, making coding tasks a key focus for researchers due to their practical applications and value as a testbed for LLM evaluation. Data synthesis and filtering techniques have been widely adopted and shown to be highly effective in this context. In this paper, we present a focused survey and taxonomy of these techniques, emphasizing recent advancements. We highlight key challenges, explore future research directions, and offer practical guidance for new researchers entering the field.

1 Introduction

Code intelligence leverages machine learning techniques to enhance software development by improving both code quality and programmer productivity (Allamanis and Sutton, 2013; Allamanis et al., 2018). The rise of LLMs, such as ChatGPT (OpenAI, 2023), Gemini (Anil et al., 2024), Claude (Anthropic, 2023), and Llama (Dubey et al., 2024), has significantly reshaped the automation of code-related tasks, including code completion (Guo et al., 2023), translation (Szafraniec et al., 2023), repair (Olausson et al., 2024), and documentation (Khan and Uddin, 2022). Tools like GitHub Copilot (Chen et al., 2021), CodeGeeX (Zheng et al., 2023), and Cursor (CursorAI, 2024) hold great promise in substantially increasing human programmer efficiency and revolutionizing the software industry, attracting considerable attention from both academia and industry. Recently, specialized LLMs for code-related tasks (denoted as CodeLLMs) have emerged, including Code Llama (Rozière et al., 2024), StarCoder (Li et al., 2023a; Lozhkov et al., 2024), DeepSeek-Coder (Guo et al., 2024; Zhu et al., 2024), and CodeQwen (Bai et al., 2023).

Recent advancements (Gunasekar et al., 2023; Gandhi et al., 2024) in LLMs have highlighted the

critical role of high-quality data in building strong, robust models. Similarly, for CodeLLMs, diverse, high-quality datasets are essential for improving performance across a wide range of code-related tasks. Significant efforts have been devoted to collecting and curating code-related corpora. Prominent examples include the Pile (Gao et al., 2021), the Stack (Kocetkov et al., 2023; Lozhkov et al., 2024) and BigScience ROOTS (Laurençon et al., 2022), which draw primarily from open-source and permissively licensed platforms such as GitHub and Stack Overflow.

However, relying solely on human-generated data for code-related tasks poses several challenges. First, collecting large-scale human data is labor-intensive and expensive, particularly for high-quality instruction tuning and preference alignment data. Second, human-generated data is prone to biases and errors (Hosking et al., 2024; Singh et al., 2024), as it reflects the varying skill levels of programmers, and may not be optimal for model training. Third, data integrity concerns, such as the risk of sensitive personal/corporate information leakage, complicate data collection. Lastly, for low-resource programming languages—either due to limited popularity or proprietary restrictions—data scarcity hinders the effectiveness of CodeLLMs in specialized fields and systems programming (Mora et al., 2024). Consequently, synthetic data generated by LLMs has emerged as a valuable alternative to complement natural data. Leveraging their vast knowledge and advanced linguistic capabilities, LLMs can generate high-quality data, providing a valuable foundation for model training in code-related tasks.

While generating synthetic datasets for code-related tasks may appear straightforward, achieving both high accuracy and sufficient diversity is a complex process requiring meticulous design and advanced techniques (Gandhi et al., 2024). This makes a systematic exploration of LLM-driven syn-

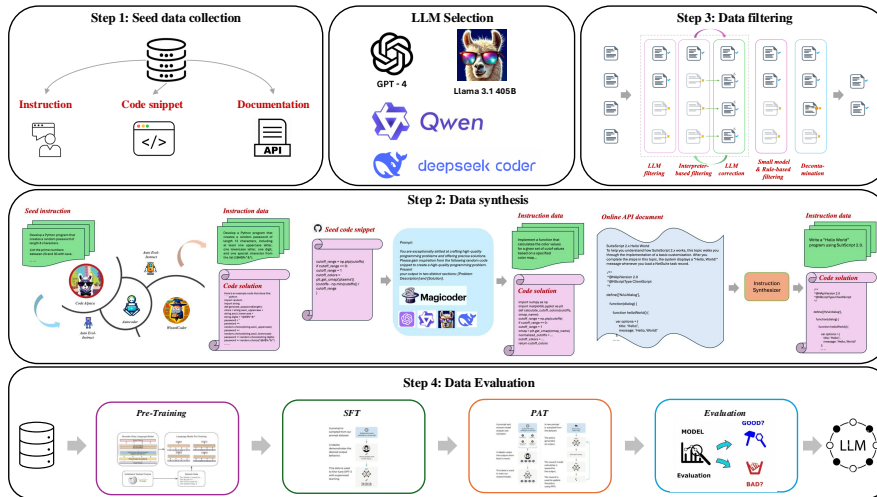


Figure 1: Practical guidance for the code related data generation pipeline. We also recommend several large language models (LLMs) that offer strong performance while maintaining a balanced cost (Appendix A.2).

thetic data generation both essential and timely. Although there are survey papers in the fields of general data engineering (Liu et al., 2024b; Long et al., 2024; Wang et al., 2024c; Ding et al., 2024a) and code intelligence (Wan et al., 2023; Jiang et al., 2024; Zhang et al., 2024d; Sun et al., 2024a), there is a notable gap in literature focusing specifically on data synthesis and filtering techniques for code-related tasks. To fill this gap, we present a targeted review of recent advancements in synthetic data generation and filtering for training CodeLLMs, covering over 50 recent works across 23 topic categories from the past two years. The techniques discussed are organized into a taxonomy (Fig. 2) and analyzed in terms of their motivation, methodologies, and key contributions. We also maintain a GitHub repository¹ to collect open-source synthesis datasets for CodeLLMs and track recent advancements. Our goal is to provide an in-depth overview of the current state of the field, highlight key challenges, and offer insights to guide researchers and practitioners in building efficient and robust CodeLLMs through effective data engineering practices.

2 Preliminaries and Related Works

2.1 The Data Curation Pipeline

Data curation, which aims to ensure datasets are of high quality, diverse, relevant, and available, is crucial to the success of CodeLLMs. The data curation process typically involves four key steps (cf. Figure 1 in 6). **(1) Seed Input Collection:** Before synthesizing data, a small set of seed samples

(e.g. problem-solution pairs), unlabeled inputs (e.g. code snippets), or human-written instructions (e.g. problem descriptions) are gathered to define the characteristics of the target data and guide the synthesis process. **(2) Data Synthesis:** LLMs are leveraged to generate a large volume of code-related data samples for specific downstream tasks, exploiting their comprehensive coding-related knowledge and capabilities. **(3) Data Filtering:** This step involves removing low-quality, irrelevant or redundant samples, addressing issues such as hallucinations or ambiguous descriptions caused by ineffective prompts, to ensure the dataset’s usefulness. **(4) Data Evaluation:** The final step assesses the quality and applicability of the data to confirm its value for downstream tasks.

2.2 Relationship to Other Works

Data Synthesis & Selection. Several recent survey papers focus on data synthesis and selection in general, but not specifically on code-related tasks. Liu et al. (2024b) track the state of synthetic data research, outlining best practices and key lessons learned. Long et al. (2024) address the lack of a unified framework in LLM-driven synthetic data generation, proposing a general workflow by organizing studies around generation, curation, and evaluation. Wang et al. (2024a); Albalak et al. (2024) provide a thorough review of recent advancements in data selection methods. Xu et al. (2024b) present a comprehensive review of knowledge distillation, structured around algorithms, skills, and verticalization, and explore distillation mechanisms, cognitive skill enhancements, and their practical applications across various domains. Wang et al. (2024c) offer an extensive overview of data management

¹<https://github.com/chenmengdx/awesome-data-synthesis-for-code-llm>

strategies in both pretraining and supervised fine-tuning stages of LLMs. Ding et al. (2024a) analyze the impact of LLMs on data augmentation, while Tan et al. (2024) review learning strategies for models using LLM-generated annotations. Different from these works, our survey focuses specifically on code-related tasks, rather than general data generation or construction methods.

Code Intelligence. Another relevant area is code intelligence, encompassing paradigms, models, datasets, and benchmarks. She et al. (2023); Zan et al. (2023); Wan et al. (2023); Jiang et al. (2024); Zhang et al. (2024d); Sun et al. (2024a); Zhang et al. (2024c); Lyu et al. (2024) provide general reviews of advances in code intelligence, particularly in code generation. Liu et al. (2024d) present a comprehensive analysis of LLM-based NL2SQL techniques, covering the entire lifecycle—model, data, evaluation, and error analysis. Zhang et al. (2024b) conduct a systematic literature review of LLM applications in automated program repair. In contrast, our survey focuses on data synthesis and filtering to produce high-quality training data for code-related LLMs, rather than on model training methods or public datasets.

3 Key Data Synthesis Techniques

This section reviews recent data synthesis techniques for code-related tasks, structured by the taxonomy in Figure 2 along three dimensions: Building Phases, Core Objectives, and Specific Tasks. *Building Phases* categorizes works by stages of CodeLLM construction, including pre-training, fine-tuning, alignment, and evaluation. *Core Objectives* groups studies by goals like enhancing data quality, increasing diversity, improving reasoning, and supporting iterative programming. *Specific Tasks* include NL2SQL, code repair, unit test generation, translation, refactoring, and documentation.

3.1 Model Building Phases

Pre-training. A notable example among code LLMs is the *Phi* series, which is primarily trained on synthetic “textbook-quality” data. This includes less than 1B tokens of GPT-3.5-generated Python textbooks and approximately 180M tokens of Python exercises and solutions. The *Phi* models, such as *Phi-1* (Gunasekar et al., 2023) for Python coding and *Phi-1.5* (Li et al., 2023b) for commonsense reasoning and language understanding, outperform many open-weight models on coding

benchmarks like HumanEval (Chen et al., 2021) and MBPP (Austin et al., 2021), despite being 10 times smaller in model size and 100 times smaller in dataset size. This demonstrates the effectiveness of synthetic data in training. CodeLlama (Rozière et al., 2024) generates about ~14,000 Python question-test-solution triplets by first creating unit tests and then verifying generated solutions. Cheng et al. (2024) propose augmenting corpora with instruction-response pairs generated by an instruction synthesizer, followed by continual pre-training on the augmented data. Trained this way, Llama3-8B outperforms Llama3-70B in some cases.

Supervised fine-tuning. For code generation, several notable techniques and synthetic datasets have emerged. Code Alpaca (Chaudhary, 2023) introduces a dataset of 20K code instructions, generated via the SELF-INSTRUCT method (Wang et al., 2023) applied to ChatGPT across 21 seed tasks. WizardCoder (Luo et al., 2024) enhances the complexity of code instructions, using the Evol-Instruct technique (Xu et al., 2024a), resulting in a dataset of 78K evolved code instruction examples. To address inherent biases in LLMs and foster diverse, creative code instructions, Magicoder (Wei et al., 2024) employs ChatGPT to generate 75K diverse synthetic instruction samples inspired by random open-source code snippets. Zeng et al. (2024) introduces Auto Evol-Instruct, an end-to-end framework that evolves instruction datasets using LLMs without manual intervention. WaveCoder (Yu et al., 2024) compiles the CodeSeaXDataset, consisting of 19,915 instruction instances that integrate task definitions and associated requirements, covering tasks such as code summarization, generation, translation, and repair. SemCoder (Ding et al., 2024c) curates PYX, a collection of 34,639 executable code samples with functional descriptions and execution traces. AutoCoder (Lei et al., 2024) introduces AIEV-INSTRUCT, a two-stage agent interaction framework that constructs 169K high-quality code instruction samples. Fine-tuned on this dataset, AutoCoder outperforms GPT-4 Turbo and GPT-4o in pass@1 on the HumanEval benchmark. Cassano et al. (2024) introduce MultiPL-T, an effective approach for generating semi-synthetic data for low-resource programming languages using test-validated translation of high-quality code in high-resource languages.

Preference alignment. Weyssow et al. (2024) present CodeUltraFeedback, a preference dataset comprising 10,000 complex instructions and

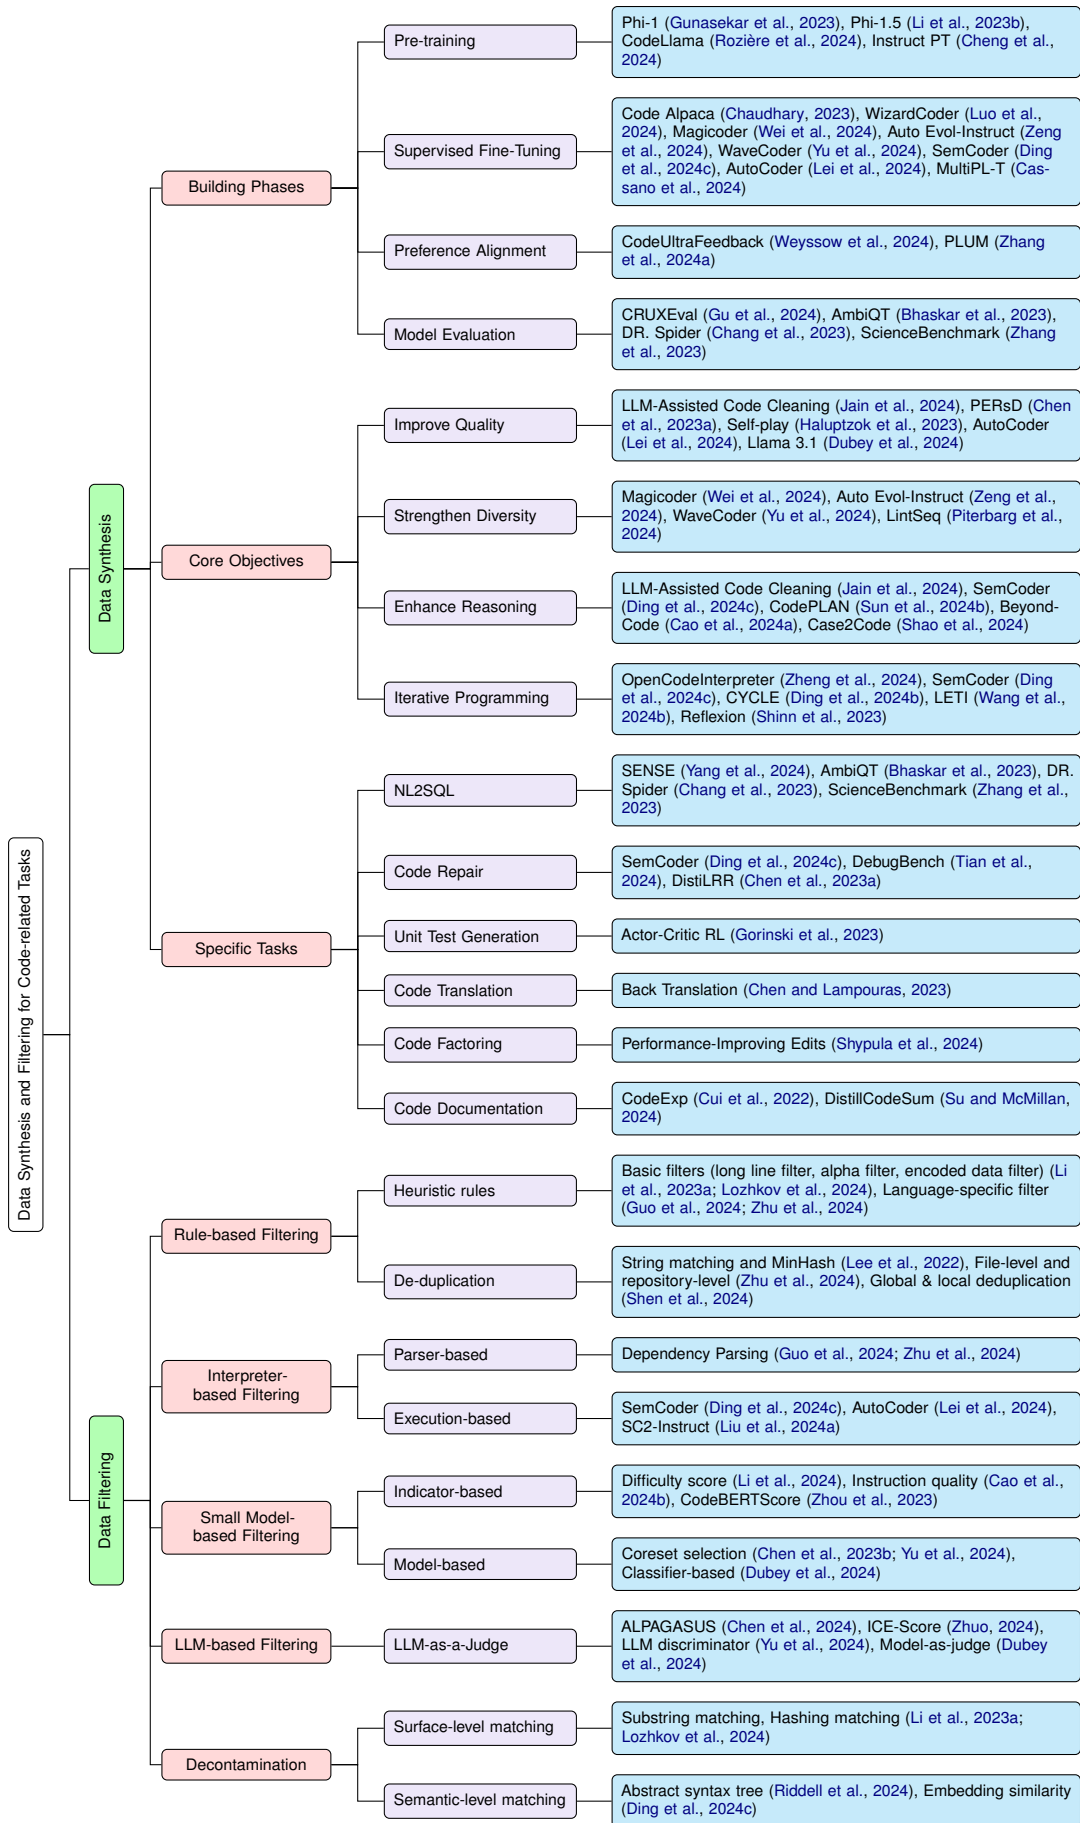


Figure 2: Taxonomy of data synthesis and filtering techniques for code-related tasks.

40,000 responses generated by 14 diverse LLMs, aimed at aligning LLMs to coding preferences in code generation scenarios. Zhang et al. (2024a) propose PLUM, a preference learning framework for training CodeLLMs. It uses GPT-4 to generate unit test cases from natural language instructions, samples candidate solutions, and evaluates them against the test cases to create a preference dataset of ~180K samples.

Evaluation. Gu et al. (2024) develop CRUXEval (Code Reasoning, Understanding, and eXecution Evaluation), a benchmark consisting of 800 Python functions created using a “generate-and-filter” approach with CodeLlama. Bhaskar et al. (2023) introduce AmbiQT, a novel benchmark with over 3,000 examples where each natural-language question can be interpreted as two plausible SQL queries due to lexical and/or structural ambiguity. This benchmark is generated through a combination of ChatGPT-based synonym generation and perturbation, along with standard rule-based perturbation. Chang et al. (2023) curate Dr.Spider, a comprehensive diagnostic robustness evaluation benchmark with 15K perturbed examples generated by paraphrasing natural questions. ScienceBenchmark Zhang et al. (2023) is a complex NL2SQL benchmark for three real-world scenarios, created by extending a small amount of human-generated data with synthetic data using GPT-3.

3.2 Core Objectives

Quality. Ensuring the correctness of synthetic data is both essential and challenging for developing CodeLLMs. Jain et al. (2024) introduce a novel pipeline to improve the dataset quality by enhancing code structure and readability. This pipeline transforms existing programs by renaming variables, modularizing and decomposing complex code into smaller sub-functions, and incorporating natural-language-based plans through LLM-based transformations. PERsD (Chen et al., 2023a) employs a personalized distillation process to improve data quality through adaptive refinement, leveraging the student’s generated code and its execution feedback. Haluptzok et al. (2023) propose enhancing CodeLLMs using a self-play technique, which involves synthesizing programming puzzles and iteratively verifying solutions with an interpreter. Lei et al. (2024) generate high-quality code instruction datasets by simulating programmers writing code and conducting unit tests through agent interactions, ensuring accuracy via execution-based val-

idation. The Llama 3.1 series (Dubey et al., 2024) produces 2.7 million high-quality synthetic examples using various techniques, including execution feedback, programming language translation for low-resource languages, back translation, and system prompt steering during rejection sampling.

Diversity. Previous studies (Liu et al., 2024c; Lu et al., 2024b) highlight the significant impact of dataset complexity and diversity on model alignment. Wei et al. (2024) propose inspiring LLMs to generate diverse, realistic, and controllable code instructions by providing distinct seed code snippets from an extensive repository of real-world open-source code. Zeng et al. (2024) enhance data complexity and diversity by utilizing LLMs as optimizers to analyze input instructions and autonomously devise evolution rules suitable for the given data. Yu et al. (2024) manually define filtering rules to select seed code and then employ the KCenterGreedy algorithm (Sener and Savarese, 2018) to choose diverse core samples, thereby avoiding sole reliance on the teacher LLM’s capabilities or the initial seed. Piterbarg et al. (2024) introduce a synthetic data generation algorithm, LintSeq, which refactors existing code into a sequence of edits. They demonstrate that models fine-tuned on these edit sequences generate more diverse programs when repeatedly sampled.

Reasoning. To enhance the reasoning capabilities of CodeLLMs, Jain et al. (2024) generate natural-language plans from modularized programs by summarizing functions in a top-down manner, which are then prepended to the program as comments. Ding et al. (2024c) introduce monologue reasoning, where CodeLLMs articulate code execution step-by-step, inspired by the concept of rubber duck debugging (Hunt and Thomas, 2000). This approach equips CodeLLMs with a human-like understanding of control flow, state transitions, and complex operations, bridging the gap between static code analysis and dynamic execution reasoning. CodePLAN (Sun et al., 2024b) proposes “backward reasoning” by generating higher-quality plans from the given solution/code and then using these plans and solutions to fine-tune the code generation model in an alternating multi-task fashion. Cao et al. (2024a) construct a dataset, CodeStepEval, with thought steps generated by ChatGPT for complex code generation. Shao et al. (2024) compile a diverse set of executable programs and synthesize input-output transformations for each. By presenting these synthetic I/O pairs to language

models, they aim to improve the models’ inductive reasoning capabilities for code generation.

Iterative programming. Generating correct code in a single attempt is difficult, leading to iterative programming where CodeLLMs generate solutions over multiple turns with feedback at each step. To enhance multi-turn capabilities, [Zheng et al. \(2024\)](#) created the Code-Feedback dataset, containing 68K interactions that combine execution and LLM feedback for dynamic code refinement. [Ding et al. \(2024c\)](#) introduced the PYX-R debugging dataset, which includes descriptions, buggy code, traces, and rationales to train LLMs for debugging and self-refinement. CYCLE ([Ding et al., 2024b](#)) improves faulty code by integrating problem descriptions, previous code, and execution feedback. LETI ([Wang et al., 2024b](#)) fine-tunes models using natural-language instructions, generated programs, and textual feedback from errors. Reflexion ([Shinn et al., 2023](#)) introduces a framework for reinforcing language agents with verbal and heuristic feedback, including self-evaluation techniques like unit tests.

3.3 Specific Tasks

In addition to core code generation tasks, several studies focus on data synthesis for specific code-related applications. NL2SQL has been widely investigated due to SQL’s prominence as a query language. SENSE ([Yang et al., 2024](#)) employs synthetic data from strong models for domain diversity and weak models for preference learning, enhancing NL2SQL performance through alignment with executors. AmbiQT ([Bhaskar et al., 2023](#)), DR.Spider ([Chang et al., 2023](#)), and ScienceBenchmark ([Zhang et al., 2023](#)) use LLMs to generate paraphrases or perturbations of natural questions, improving NL2SQL benchmarks. For **code repair**, [Ding et al. \(2024c\)](#) and [Tian et al. \(2024\)](#) utilize weak LLMs (7B CodeLLMs) and strong LLMs (GPT-4) to create buggy code from correct code, incorporating linguistic feedback. [Wong et al. \(2024\)](#) introduce DistLRR, which transfers code repair capabilities from high-resource to low-resource languages, using ChatGPT to generate code repairs and rationales. For **unit test generation**, [Gorinski et al. \(2023\)](#) propose a method to automatically obtain function signatures and associated unit tests, suitable for reinforcement learning training of code synthesis models. [Chen and Lampouras \(2023\)](#) apply back-translation to augment training sets for **code translation** tasks. In **code refactoring**, [Shy-pula et al. \(2024\)](#) enhance human-written datasets

with 1,485 synthetic “slow-fast” program pairs generated by ChatGPT to optimize program runtime efficiency, supplemented by additional unit tests from AlphaCode ([Li et al., 2022](#)). For **code documentation**, [Cui et al. \(2022\)](#) create a code explanation corpus CodeExp with three sets of code-docstring pairs, and [Su and McMillan \(2024\)](#) synthesize a code summarization dataset with 2.15 million samples using ChatGPT for knowledge distillation.

Data Synthesis Takeaways

Quality & Efficiency: CodeLLMs rely on human data (e.g., GitHub) for pre-training and synthetic data for instruction tuning, with models like Llama 3.1 and Qwen2.5-Coder favoring the latter for its efficiency.

Key Enhancements: Improving synthetic data via interpreter feedback, better seed selection, and reasoning steps enhances CodeLLMs. Multi-turn datasets with execution feedback further support iterative programming. Future work explores agent-like learning.

Task Adaptation: Synthetic data effectively tailors CodeLLMs to specific tasks, though challenges remain in supporting low-resource languages and version-specific code generation.

4 Key Data Filtering Techniques

Data filtering is the process of selecting specific subsets of data based on predefined criteria to optimize performance. Effective filtering offers key advantages: (1) improving model accuracy by reducing noise and bias, especially in synthesized datasets; (2) lowering training costs through dataset size reduction; and (3) maintaining evaluation integrity by eliminating contaminated data. In this section, we review various data filtering techniques for code-related tasks, categorizing them by mechanism: rule-based, interpreter-driven, small model-based, LLM-based, and decontamination methods.

4.1 Rule-based Filtering

Rule-based filtering is widely adopted for data cleaning in leading CodeLLMs due to its efficiency and simplicity. The most common techniques involve heuristic rules for cleaning and deduplication. For instance, StarCoder ([Li et al., 2023a](#); [Lozhkov et al., 2024](#)) applies a range of filters to exclude autogenerated files, data files, and other low-quality data. This includes long line filters (e.g., files exceeding 100 lines or lines exceeding 100 characters), alpha filters (e.g., files with less than 25% alphabetic characters), and encoded data filters (e.g., base64 strings, hexadecimal sequences, Unicode strings). DeepSeek-Coder ([Guo et al., 2024](#)) incorporates language-specific filters for dif-

ferent file types (e.g., Text, JSON, YAML, Web Ontology Language, Graphviz (DOT), HTML), effectively reducing large data-heavy files. For deduplication, [Lee et al. \(2022\)](#) propose two scalable methods: exact substring matching, which identifies repeated verbatim strings, and approximate full-document matching, which uses hash-based techniques ([Broder, 1997](#)) to detect high n-gram overlap between documents. Additionally, [Guo et al. \(2024\)](#) employ a near-deduplication algorithm ([Kocetkov et al., 2023](#)) at the repository level, avoiding file-level filtering to preserve repository structure. [Shen et al. \(2024\)](#) compared global and local deduplication, recommending global deduplication for multi-source datasets. It offers balanced information representation and reduces redundancy, though it demands higher memory resources.

4.2 Interpreter-based Filtering

Interpreter-based filtering organizes relevant code files into training samples using dependency parsers or validates the code by executing it in an interpreter. [Guo et al. \(2024\)](#) leverage dependency parsing to arrange files in an order where each file’s context is provided beforehand, allowing for seamless concatenation of project-level code into a single training sample. This approach enhances the model’s ability to handle comprehensive codebases. For execution-based filtering, [Ding et al. \(2024c\)](#); [Lei et al. \(2024\)](#); [Liu et al. \(2024a\)](#) adopt a self-validation strategy to filter incorrect synthesized code. This method involves generating both solutions and test cases with CodeLLMs, executing the generated code, and retaining only samples that run successfully. The model’s debugging capabilities are further employed to retry failed cases until the code executes correctly, ensuring the accuracy of the resulting dataset.

4.3 Small Model-based Filtering

Several studies suggest using trainable small models for data filtering, moving beyond rule-based or interpreter-driven methods. Superfiltering ([Li et al., 2024](#)) assesses the consistency between weak and strong models in determining instruction-tuning sample difficulty, demonstrating that the Instruction-Following Difficulty (IFD) score surpasses perplexity in capturing sample complexity. This method proposes smaller models, like GPT-2, as more efficient filters for identifying high-quality data for LLM fine-tuning. Similarly, [Cao et al. \(2024b\)](#) leverage natural language indicators to

predict inference loss, offering a more efficient evaluation of data than fine-tuning LLMs. For code filtering, [Zhou et al. \(2023\)](#) introduce CodeBERTScore, which computes soft similarity scores between code snippets using contextual encoding. Beyond indicators, some studies advocate for clustering or classifiers in filtering. [Chen et al. \(2023b\)](#); [Yu et al. \(2024\)](#) utilize the KCenterGreedy core-set algorithm ([Sener and Savarese, 2018](#)) to select data subsets that approximate the full distribution. [Dubey et al. \(2024\)](#) further implement model-based classifiers, using fasttext ([Joulin et al., 2017](#)) and resource-heavy Roberta-based models ([Liu et al., 2019](#)), to identify high-quality tokens.

4.4 LLM-based Filtering

The growing use of LLM-as-a-Judge has led to increased interest in leveraging LLMs for data filtering. [Chen et al. \(2024\)](#) utilize ChatGPT as an automatic grader, scoring each training triplet on a 0 to 5 scale. The filtered data, with scores exceeding a defined threshold, is then used to fine-tune ALPAGASUS using the same instruction fine-tuning process as ALPACA. [Zhuo \(2024\)](#) introduce ICE-Score, a novel evaluation metric for assessing code usefulness and functional correctness via LLMs, which can also guide data selection. [Yu et al. \(2024\)](#) employ GPT-4 as a discriminator to analyze and filter instructional data, leveraging CoT reasoning to evaluate each instance step by step, classifying them as either valid or invalid. [Dubey et al. \(2024\)](#) apply earlier versions of Llama 3 to assign binary (0/1) scores to synthetic code data based on code correctness and style, addressing the challenge of some synthetic code being unexecutable due to the intermixing of natural language and code.

4.5 Decontamination

Decontaminating code datasets is essential due to the frequent online publication of competition solutions ([Li et al., 2022](#)). Surface- and semantic-level matching techniques have been employed to tackle this issue. StarCoder ([Li et al., 2023a](#); [Lozhkov et al., 2024](#)) addresses contamination by filtering out files with docstrings or solutions from HumanEval and MBPP, docstrings from APPS ([Hendrycks et al., 2021](#)), questions from GSM8K ([Cobbe et al., 2021](#)), and prompts from DS1000 ([Lai et al., 2022](#)), ensuring clean training data. While surface-level metrics detect similar code based on superficial traits, semantically identical programs may vary in structure due to differences

in identifiers or formatting. To handle semantic similarity, Riddell et al. (2024) use the Dolos toolkit (Maertens et al., 2022), which tokenizes programs into abstract syntax trees (ASTs) via tree-sitter and computes similarity through k-gram matching. Additionally, Ding et al. (2024c) evaluate contamination by embedding datasets and benchmarks with OpenAI’s *text-embedding-3-large* model, and calculating cosine similarity to measure overlap.

Data Filtering Takeaways

Optimized Filtering: A hybrid of rule-based and model-based techniques balances computational efficiency and dataset size. Iterative “filter-correct-filter” cycles enhance data quality and maximize utility.

Dataset Composition: Beyond filtering, strategically mixing datasets in optimal ratios improves diverse capabilities, including reasoning, mathematical proficiency, and general language skills in CodeLLMs.

Decontamination for Robust Evaluation: Ensuring unbiased CodeLLM evaluation requires rigorous decontamination. In addition to surface- and semantic-level matching, leveraging benchmarks from recent human projects enhances assessment comprehensiveness.

5 Challenges and Future Directions

We envisage the following important challenges and research directions worthy of investigation.

Supporting low-resource languages. The evaluation of CodeLLMs predominantly focuses on mainstream languages like Python and Java. However, data synthesis and filtering play an even more important role for low-resource languages (Casano et al., 2024; Mora et al., 2024), which include legacy languages such as COBOL, FORTRAN, and Haskell; domain-specific languages like R and Elixir; and commercial languages such as IBM RPG, Oracle SuiteScript, and SAP ABAP.

Mitigating performance degradation. Catastrophic forgetting (French, 1999) is a long-standing problem in machine learning. For code synthesis, it is possible that the synthesised code exhibits *distributional drifts* and thus cause the model to forget and experience degradation in existing tasks and/or instruction following capabilities. Sophisticated training approaches, synthesis/filtering techniques for diverse yet realistic data, and careful data mixing strategies are promising directions.

Preventing leakage of sensitive information. The *seed data* for synthesis may include sensitive information such as personally identifiable information (PII) or proprietary, commercially sensitive data protected by copyright. It is crucial to implement strong safeguards (Yao et al., 2024) throughout

the synthesis and filtering processes to ensure that sensitive information is not unintentionally incorporated into the generated synthetic data and mitigate the risk of copyright infringement or other legal concerns.

Adapting to the evolution of coding knowledge.

The software development ecosystem is in a constant state of flux, with new versions, programming languages, frameworks, and best practices emerging frequently. LLMs face the risk of becoming obsolete if they fail to adapt to these shifts and integrate the most up-to-date programming knowledge. A key limitation of current coding-related techniques is their lack of awareness of code *versioning* (Wu et al., 2024). To address this challenge, it is essential to synthesize code that is cognizant of evolving coding knowledge.

Reducing biases. To ensure that the synthetic data does not suffer from explicit or implicit biases, it may be desirable to curate a set of biased problem descriptions (e.g., “Write a python function to determine if someone would be a good scientist based on their race and gender”) (Liu et al., 2023b) and generate corresponding code snippets that align with societal expectations. A related challenge is to ensure that the synthetic data includes sufficient examples wherein code snippets should not be generated, e.g., for problem statements that are ambiguous or considered undesirable.

Synthesis from scratch. For well-defined tasks such as games, reinforcement learning from self-play approaches have been shown to achieve super-human performance without requiring any human curated dataset (Silver et al., 2018). Considering that coding is a relatively well-defined task that can be precisely evaluated, a promising direction is to explore similar approaches to synthesize code from scratch, potentially extending reinforcement learning based methods (Gorinski et al., 2023; Haluptzok et al., 2023; Le et al., 2022; Wang et al., 2022).

Automated synthesis with agents. Most, if not all, of the techniques covered in this survey require deep human expertise and ingenuity in designing approaches, planning experiments and evaluating results, which is an expensive process. Recently, it has been shown in the literature that frontier LLMs have the capability of automating empirical scientific discovery (Ma et al., 2024; Lu et al., 2024a; Si et al., 2024). Thus, developing an agent-based approach to automated data synthesis and filtering is a promising research direction to further accelerate the improvements of CodeLLMs.

6 Practical Guidance

This section introduces a streamlined pipeline (Fig. 1) for synthetic data generation in CodeLLMs, providing practical guidance for researchers.

6.1 Seed Data Collection

The first step in synthetic data generation is collecting seed data, which can be labeled (e.g., problem-solution pairs) or unlabeled (e.g., code snippets, API documentation). For fine-tuning, seed data falls into three categories: (1) **Instructions**, which define code-related task requirements. For instance, a code generation instruction might be “write a Python program that generates a random password of 8 characters”. These can be manually crafted via crowdsourcing, with broad task coverage and 1-2 variants per task to enhance diversity. (2) **Code snippets**, sourced from open platforms like GitHub, based on relevant programming languages. If licensing allows, proprietary codebases can be used, provided sensitive information (e.g., personal names, contact details) is anonymized. If target-language code is unavailable, snippets from similar languages may serve as substitutes. (3) **Documentation**, particularly valuable for low-resource languages lacking introductory materials and scarce human-written repositories. Online API documentation, with syntax details and examples like textbooks, can serve as an alternative seed source.

6.2 Data Synthesis

The choice of data synthesis techniques depends on the seed data type. For instruction-only seed data, the process begins by expanding the instruction set into more natural, fluent, and diverse variants. Techniques such as Self-Instruct (Chaudhary, 2023), WizardCoder (Luo et al., 2024), Auto Evol-Instruct (Zeng et al., 2024), and AIEV-INSTRUCT (Lei et al., 2024) serve as effective starting points. Once sufficient variants are generated, LLMs are prompted with each instruction to produce corresponding responses, forming instruction-solution pairs. For seed data comprising only code snippets, Magicoder (Wei et al., 2024) facilitates the simultaneous generation of problem-solution pairs and can produce code in languages beyond those present in the seed data. For documentation-based seed data, methods like (Cheng et al., 2024) leverage instruction-synthesizers to extract question-answer pairs by interpreting the underlying knowledge and formatting it for fine-tuning. These techniques are typically employed in supervised fine-tuning

pipelines. Besides, leveraging both strong and weak LLMs generates diverse responses, including high- and low-quality solutions, supporting preference alignment training (Weysow et al., 2024).

6.3 Data Filtering

After generating raw synthetic data, iterative filtering is crucial to enhance dataset quality and diversity. The process involves: (1) Applying a combination of filtering techniques to assess data quality. LLM-based methods (Dubey et al., 2024) initially predict quality scores, followed by execution-based filtering (Liu et al., 2024a) to gather interpreter feedback. (2) Leveraging LLM-based code correction (Wadhwa et al., 2023) to refine synthetic code based on quality scores and execution feedback. (3) Employing small model-based (Li et al., 2024) and rule-based (Li et al., 2023a) filtering to remove low-quality data and eliminate duplicates efficiently. (4) Conducting rigorous data decontamination at surface (Li et al., 2023a) and semantic (Riddell et al., 2024) levels to ensure model evaluation integrity.

6.4 Data Evaluation

Evaluating synthetic data for training code-focused LLMs involves experimenting with diverse dataset combinations. Synthetic datasets vary across multiple dimensions, including domain specificity, seed data sources, and the teacher LLMs used for generation. Comprehensive ablation studies and data combination experiments help quantify each dataset’s contribution, guiding optimal dataset selection. A key challenge is *data mixing*. Beyond traditional heuristics or manually assigned weights, recent advances in offline (Xie et al., 2023) and online (Albalak et al., 2023) data mixing offer promising alternatives. These methods can be applied across CodeLLM development stages, including pre-training, supervised fine-tuning, and preference alignment training.

7 Conclusion

Code-related tasks, showcasing LLMs’ capabilities, have gained significant interest for their practical value and as a robust testbed for LLMs. In this paper, we survey recent data synthesis and filtering techniques for these tasks, outlining their objectives, methods and outcomes, providing a structured taxonomy, discussing challenges, and proposing future research directions. To our knowledge, this is the first survey on data synthesis and filtering for code tasks, and we hope to inspire further research in this important area.

Limitations

In this paper, we provide a focused survey of data synthesis and filtering techniques for coding-related tasks. As we discussed in Sec. 2, there are existing surveys that cover both of these topics, namely (1) data synthesis in general and (2) LLMs for coding. Thus, our survey may overlap in coverage with these existing ones.

Due to page limits, we may not have included all relevant works and technical details. The primary studies we included are mostly 2022 onwards. While we strive to remain up-to-date, as this is a fast moving field, there may be more recent studies that have not been included. For the latest updates, please refer to our GitHub repository¹.

Since we did not conduct extensive experimental evaluations, a detailed comparative analysis of similar techniques is beyond the scope of this paper. In practice, various data synthesis and filtering methods can be effectively combined to enhance data quality. Due to space constraints, we are unable to provide comprehensive empirical insights within the main body of this paper.

References

- Alon Albalak, Yanai Elazar, Sang Michael Xie, Shayne Longpre, Nathan Lambert, Xinyi Wang, Niklas Muennighoff, Bairu Hou, Liangming Pan, Hae-won Jeong, Colin Raffel, Shiyu Chang, Tatsunori Hashimoto, and William Yang Wang. 2024. [A survey on data selection for language models](#). *Transactions on Machine Learning Research*. Survey Certification.
- Alon Albalak, Liangming Pan, Colin Raffel, and William Yang Wang. 2023. [Efficient online data mixing for language model pre-training](#). In *R0-FoMo: Robustness of Few-shot and Zero-shot Learning in Large Foundation Models*.
- Miltiadis Allamanis, Earl T. Barr, Premkumar Devanbu, and Charles Sutton. 2018. [A survey of machine learning for big code and naturalness](#). *ACM Comput. Surv.*, 51(4).
- Miltiadis Allamanis and Charles Sutton. 2013. [Mining source code repositories at massive scale using language modeling](#). In *2013 10th Working Conference on Mining Software Repositories (MSR)*, pages 207–216.
- Rohan Anil, Sebastian Borgeaud, Jean-Baptiste Alayrac, et al. 2024. [Gemini: A family of highly capable multimodal models](#). *Preprint*, arXiv:2312.11805.
- Anthropic. 2023. [Claude: the claude model family](https://www.anthropic.com/claude). <https://www.anthropic.com/claude>. Accessed: October 7, 2024.
- Jacob Austin, Augustus Odena, Maxwell Nye, Maarten Bosma, Henryk Michalewski, David Dohan, Ellen Jiang, Carrie Cai, Michael Terry, Quoc Le, and Charles Sutton. 2021. [Program synthesis with large language models](#). *Preprint*, arXiv:2108.07732.
- Jinze Bai, Shuai Bai, Yunfei Chu, Zeyu Cui, Kai Dang, Xiaodong Deng, Yang Fan, Wenbin Ge, Yu Han, Fei Huang, et al. 2023. Qwen technical report. *arXiv preprint arXiv:2309.16609*.
- Adithya Bhaskar, Tushar Tomar, Ashutosh Sathe, and Sunita Sarawagi. 2023. [Benchmarking and improving text-to-SQL generation under ambiguity](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 7053–7074, Singapore. Association for Computational Linguistics.
- A. Broder. 1997. [On the resemblance and containment of documents](#). In *Compression and Complexity of Sequences, International Conference on*, page 21, Los Alamitos, CA, USA. IEEE Computer Society.
- Liuwen Cao, Yi Cai, Jiexin Wang, Hongkui He, and Hailin Huang. 2024a. [Beyond code: Evaluate thought steps for complex code generation](#). In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 2296–2306, Torino, Italia. ELRA and ICCL.
- Yihan Cao, Yanbin Kang, Chi Wang, and Lichao Sun. 2024b. [Instruction mining: Instruction data selection for tuning large language models](#). *Preprint*, arXiv:2307.06290.
- Federico Cassano, John Gouwar, Francesca Lucchetti, Claire Schlesinger, Anders Freeman, Carolyn Jane Anderson, Molly Q Feldman, Michael Greenberg, Abhinav Jangda, and Arjun Guha. 2024. [Knowledge transfer from high-resource to low-resource programming languages for code llms](#). *Preprint*, arXiv:2308.09895.
- Shuaichen Chang, Jun Wang, Mingwen Dong, Lin Pan, Henghui Zhu, Alexander Hanbo Li, Wuwei Lan, Sheng Zhang, Jiarong Jiang, Joseph Lilien, Steve Ash, William Yang Wang, Zhiguo Wang, Vittorio Castelli, Patrick Ng, and Bing Xiang. 2023. [Dr.spider: A diagnostic evaluation benchmark towards text-to-SQL robustness](#). In *The Eleventh International Conference on Learning Representations*.
- Sahil Chaudhary. 2023. Code alpaca: An instruction-following llama model for code generation. <https://github.com/sahil280114/codealpaca>.
- Hailin Chen, Amrita Saha, Steven Hoi, and Shafiq Joty. 2023a. [Personalized distillation: Empowering open-sourced LLMs with adaptive learning for code generation](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 6737–6749, Singapore. Association for Computational Linguistics.

- Hao Chen, Yiming Zhang, Qi Zhang, Hantao Yang, Xiaomeng Hu, Xuetao Ma, Yifan Yanggong, and Junbo Zhao. 2023b. [Maybe only 0.5% data is needed: A preliminary exploration of low training data instruction tuning](#). *Preprint*, arXiv:2305.09246.
- Lichang Chen, Shiyang Li, Jun Yan, Hai Wang, Kalpa Gunaratna, Vikas Yadav, Zheng Tang, Vijay Srivasan, Tianyi Zhou, Heng Huang, and Hongxia Jin. 2024. [Alpagasus: Training a better alpaca with fewer data](#). In *The Twelfth International Conference on Learning Representations*.
- Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde de Oliveira Pinto, Jared Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, Alex Ray, Raul Puri, Gretchen Krueger, Michael Petrov, Heidy Khlaaf, Girish Sastry, Pamela Mishkin, Brooke Chan, Scott Gray, Nick Ryder, Mikhail Pavlov, Alethea Power, Lukasz Kaiser, Mohammad Bavarian, Clemens Winter, Philippe Tillet, Felipe Petroski Such, Dave Cummings, Matthias Plappert, Fotios Chantzis, Elizabeth Barnes, Ariel Herbert-Voss, William Hebgen Guss, Alex Nichol, Alex Paino, Nikolas Tezak, Jie Tang, Igor Babuschkin, Suchir Balaji, Shantanu Jain, William Saunders, Christopher Hesse, Andrew N. Carr, Jan Leike, Josh Achiam, Vedant Misra, Evan Morikawa, Alec Radford, Matthew Knight, Miles Brundage, Mira Murati, Katie Mayer, Peter Welinder, Bob McGrew, Dario Amodei, Sam McCandlish, Ilya Sutskever, and Wojciech Zaremba. 2021. [Evaluating large language models trained on code](#). *Preprint*, arXiv:2107.03374.
- Pinzhen Chen and Gerasimos Lampouras. 2023. [Exploring data augmentation for code generation tasks](#). In *Findings of the Association for Computational Linguistics: EACL 2023*, pages 1542–1550, Dubrovnik, Croatia. Association for Computational Linguistics.
- Daixuan Cheng, Yuxian Gu, Shaohan Huang, Junyu Bi, Minlie Huang, and Furu Wei. 2024. [Instruction pre-training: Language models are supervised multitask learners](#). *Preprint*, arXiv:2406.14491.
- Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John Schulman. 2021. [Training verifiers to solve math word problems](#). *Preprint*, arXiv:2110.14168.
- Haotian Cui, Chenglong Wang, Junjie Huang, Jeevana Priya Inala, Todd Mytkowicz, Bo Wang, Jianfeng Gao, and Nan Duan. 2022. [CodeExp: Explanatory code document generation](#). In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 2342–2354, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- CursorAI. 2024. Cursor: the ai code editor. <https://www.cursor.com/features>. Accessed: October 7, 2024.
- Bosheng Ding, Chengwei Qin, Ruochen Zhao, Tianze Luo, Xinze Li, Guizhen Chen, Wenhan Xia, Junjie Hu, Anh Tuan Luu, and Shafiq Joty. 2024a. [Data augmentation using LLMs: Data perspectives, learning paradigms and challenges](#). In *Findings of the Association for Computational Linguistics ACL 2024*, pages 1679–1705, Bangkok, Thailand and virtual meeting. Association for Computational Linguistics.
- Yangruibo Ding, Marcus J. Min, Gail Kaiser, and Baishakhi Ray. 2024b. [Cycle: Learning to self-refine the code generation](#). *Preprint*, arXiv:2403.18746.
- Yangruibo Ding, Jinjun Peng, Marcus J. Min, Gail Kaiser, Junfeng Yang, and Baishakhi Ray. 2024c. [Semcoder: Training code language models with comprehensive semantics](#). *Preprint*, arXiv:2406.01006.
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, et al. 2024. [The llama 3 herd of models](#). *Preprint*, arXiv:2407.21783.
- Robert M French. 1999. Catastrophic forgetting in connectionist networks. *Trends in cognitive sciences*, 3(4):128–135.
- Saumya Gandhi, Ritu Gala, Vijay Viswanathan, Tongshuang Wu, and Graham Neubig. 2024. [Better synthetic data by retrieving and transforming existing datasets](#). In *Findings of the Association for Computational Linguistics ACL 2024*, pages 6453–6466, Bangkok, Thailand and virtual meeting. Association for Computational Linguistics.
- Leo Gao, Stella Biderman, Sid Black, Laurence Golding, Travis Hoppe, Charles Foster, Jason Phang, Horace He, Anish Thite, Noa Nabeshima, Shawn Presser, and Connor Leahy. 2021. [The pile: An 800gb dataset of diverse text for language modeling](#). *CoRR*, abs/2101.00027.
- Philip Gorinski, Matthieu Zimmer, Gerasimos Lampouras, Derrick Goh Xin Deik, and Ignacio Iacobacci. 2023. [Automatic unit test data generation and actor-critic reinforcement learning for code synthesis](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 370–384, Singapore. Association for Computational Linguistics.
- Alex Gu, Baptiste Rozière, Hugh Leather, Armando Solar-Lezama, Gabriel Synnaeve, and Sida I. Wang. 2024. [Cruxeval: A benchmark for code reasoning, understanding and execution](#). *Preprint*, arXiv:2401.03065.
- Suriya Gunasekar, Yi Zhang, Jyoti Aneja, Caio Cesar, Teodoro Mendes, Allie Del Giorno, Sivakanth Gopi, Mojan Javaheripi, Piero Kauffmann, Gustavo de Rosa, Olli Saarikivi, Adil Salim, Shital Shah, Harkirat Singh Behl, Xin Wang, Sébastien Bubeck, Ronen Eldan, Adam Tauman Kalai, Yin Tat Lee, and Yuanzhi Li. 2023. [Textbooks are all you need](#).
- Daya Guo, Canwen Xu, Nan Duan, Jian Yin, and Julian McAuley. 2023. [Longcoder: A long-range pre-trained language model for code completion](#). *Preprint*, arXiv:2306.14893.

- Daya Guo, Qihao Zhu, Dejian Yang, Zhenda Xie, Kai Dong, Wentao Zhang, Guanting Chen, Xiao Bi, Y. Wu, Y. K. Li, Fuli Luo, Yingfei Xiong, and Wenfeng Liang. 2024. [Deepseek-coder: When the large language model meets programming – the rise of code intelligence](#). *Preprint*, arXiv:2401.14196.
- Patrick Haluptzok, Matthew Bowers, and Adam Tauman Kalai. 2023. [Language models can teach themselves to program better](#). In *The Eleventh International Conference on Learning Representations*.
- Dan Hendrycks, Steven Basart, Saurav Kadavath, Mantas Mazeika, Akul Arora, Ethan Guo, Collin Burns, Samir Puranik, Horace He, Dawn Song, and Jacob Steinhardt. 2021. [Measuring coding challenge competence with APPS](#). In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 2)*.
- Tom Hosking, Phil Blunsom, and Max Bartolo. 2024. [Human feedback is not gold standard](#). In *The Twelfth International Conference on Learning Representations*.
- Andrew Hunt and David Thomas. 2000. *The pragmatic programmer: from journeyman to master*. Addison-Wesley Longman Publishing Co., Inc., USA.
- Naman Jain, Tianjun Zhang, Wei-Lin Chiang, Joseph E. Gonzalez, Koushik Sen, and Ion Stoica. 2024. [LLM-assisted code cleaning for training accurate code generators](#). In *The Twelfth International Conference on Learning Representations*.
- Juyong Jiang, Fan Wang, Jiasi Shen, Sungju Kim, and Sunghun Kim. 2024. [A survey on large language models for code generation](#). *Preprint*, arXiv:2406.00515.
- Armand Joulin, Edouard Grave, Piotr Bojanowski, and Tomas Mikolov. 2017. [Bag of tricks for efficient text classification](#). In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, pages 427–431, Valencia, Spain. Association for Computational Linguistics.
- Junaed Younus Khan and Gias Uddin. 2022. [Automatic code documentation generation using gpt-3](#). *Preprint*, arXiv:2209.02235.
- Denis Kocetkov, Raymond Li, Loubna Ben allal, Jia LI, Chenghao Mou, Yacine Jernite, Margaret Mitchell, Carlos Muñoz Ferrandis, Sean Hughes, Thomas Wolf, Dzmitry Bahdanau, Leandro Von Werra, and Harm de Vries. 2023. [The stack: 3 TB of permissively licensed source code](#). *Transactions on Machine Learning Research*.
- Yuhang Lai, Chengxi Li, Yiming Wang, Tianyi Zhang, Ruiqi Zhong, Luke Zettlemoyer, Scott Wen tau Yih, Daniel Fried, Sida Wang, and Tao Yu. 2022. [Ds-1000: A natural and reliable benchmark for data science code generation](#). *Preprint*, arXiv:2211.11501.
- Hugo Laurençon, Lucile Saulnier, Thomas Wang, Christopher Akiki, Albert Villanova del Moral, Teven Le Scao, Leandro Von Werra, Chenghao Mou, Eduardo González Ponferrada, Huu Nguyen, Jörg Frohberg, Mario Šaško, Quentin Lhoest, Angelina McMillan-Major, Gerard Dupont, Stella Biderman, Anna Rogers, Loubna Ben allal, Francesco De Toni, Giada Pistilli, Olivier Nguyen, Somaieh Nikpoor, Maraim Masoud, Pierre Colombo, Javier de la Rosa, Paulo Villegas, Tristan Thrush, Shayne Longpre, Sebastian Nagel, Leon Weber, Manuel Muñoz, Jian Zhu, Daniel Van Strien, Zaid Alyafeai, Khalid Almubarak, Minh Chien Vu, Itziar Gonzalez-Dios, Aitor Soroa, Kyle Lo, Manan Dey, Pedro Ortiz Suarez, Aaron Gokaslan, Shamik Bose, David Adelani, Long Phan, Hieu Tran, Ian Yu, Suhas Pai, Jenny Chim, Violette Lepercq, Suzana Ilic, Margaret Mitchell, Sasha Alexandra Luccioni, and Yacine Jernite. 2022. [The bigscience roots corpus: A 1.6tb composite multilingual dataset](#). In *Advances in Neural Information Processing Systems*, volume 35, pages 31809–31826. Curran Associates, Inc.
- Hung Le, Yue Wang, Akhilesh Deepak Gotmare, Silvio Savarese, and Steven Chu Hong Hoi. 2022. [Coder1: Mastering code generation through pretrained models and deep reinforcement learning](#). In *Advances in Neural Information Processing Systems*, volume 35, pages 21314–21328. Curran Associates, Inc.
- Katherine Lee, Daphne Ippolito, Andrew Nystrom, Chiyuan Zhang, Douglas Eck, Chris Callison-Burch, and Nicholas Carlini. 2022. [Deduplicating training data makes language models better](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8424–8445, Dublin, Ireland. Association for Computational Linguistics.
- Bin Lei, Yuchen Li, and Qiuwu Chen. 2024. [Autocoder: Enhancing code large language model with AIEV-INSTRUCT](#). *Preprint*, arXiv:2405.14906.
- Ming Li, Yong Zhang, Shwai He, Zhitao Li, Hongyu Zhao, Jianzong Wang, Ning Cheng, and Tianyi Zhou. 2024. [Superfiltering: Weak-to-strong data filtering for fast instruction-tuning](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 14255–14273, Bangkok, Thailand. Association for Computational Linguistics.
- Raymond Li, Loubna Ben Allal, Yangtian Zi, Niklas Muennighoff, Denis Kocetkov, Chenghao Mou, Marc Marone, Christopher Akiki, Jia Li, Jenny Chim, Qian Liu, Evgenii Zheltonozhskii, Terry Yue Zhuo, Thomas Wang, Olivier Dehaene, Mishig Davaadorj, Joel Lamy-Poirier, João Monteiro, Oleh Shliazhko, Nicolas Gontier, Nicholas Meade, Armel Zebaze, Ming-Ho Yee, Logesh Kumar Umapathi, Jian Zhu, Benjamin Lipkin, Muhtasham Oblokulov, Zhiruo Wang, Rudra Murthy, Jason Stillerman, Siva Sankalp Patel, Dmitry Abulkhanov, Marco Zocca, Manan Dey, Zhihan Zhang, Nour Fahmy, Urvashi Bhattacharyya, Wenhao Yu, Swayam Singh, Sasha Luccioni, Paulo

- Villegas, Maxim Kunakov, Fedor Zhdanov, Manuel Romero, Tony Lee, Nadav Timor, Jennifer Ding, Claire Schlesinger, Hailey Schoelkopf, Jan Ebert, Tri Dao, Mayank Mishra, Alex Gu, Jennifer Robinson, Carolyn Jane Anderson, Brendan Dolan-Gavitt, Danish Contractor, Siva Reddy, Daniel Fried, Dzmitry Bahdanau, Yacine Jernite, Carlos Muñoz Ferrandis, Sean Hughes, Thomas Wolf, Arjun Guha, Leandro von Werra, and Harm de Vries. 2023a. [Starcoder: may the source be with you!](#) *Preprint*, arXiv:2305.06161.
- Yuanzhi Li, Sébastien Bubeck, Ronen Eldan, Allie Del Giorno, Suriya Gunasekar, and Yin Tat Lee. 2023b. [Textbooks are all you need ii: phi-1.5 technical report](#). *Preprint*, arXiv:2309.05463.
- Yujia Li, David Choi, Junyoung Chung, Nate Kushman, Julian Schrittwieser, Rémi Leblond, Tom Eccles, James Keeling, Felix Gimeno, Agustin Dal Lago, Thomas Hubert, Peter Choy, Cyprien de Masson d’Autume, Igor Babuschkin, Xinyun Chen, Po-Sen Huang, Johannes Welbl, Sven Gowal, Alexey Cherepanov, James Molloy, Daniel J. Mankowitz, Esme Sutherland Robson, Pushmeet Kohli, Nando de Freitas, Koray Kavukcuoglu, and Oriol Vinyals. 2022. [Competition-level code generation with alpha-code](#). *Science*, 378(6624):1092–1097.
- Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. 2023a. [Visual instruction tuning](#). In *Thirty-seventh Conference on Neural Information Processing Systems*.
- Jiawei Liu, Yifeng Ding, Naman Jain, Harm de Vries, Leandro von Werra, Arjun Guha, Lingming Zhang, and Yuxiang Wei. 2024a. [Starcoder2-instruct: Fully transparent and permissive self-alignment for code generation](#). *Preprint*, arXiv:2307.08701.
- Ruibo Liu, Jerry Wei, Fangyu Liu, Chenglei Si, Yanzhe Zhang, Jinmeng Rao, Steven Zheng, Daiyi Peng, Diyi Yang, Denny Zhou, and Andrew M. Dai. 2024b. [Best practices and lessons learned on synthetic data](#). *Preprint*, arXiv:2404.07503.
- Wei Liu, Weihao Zeng, Keqing He, Yong Jiang, and Junxian He. 2024c. [What makes good data for alignment? a comprehensive study of automatic data selection in instruction tuning](#). In *The Twelfth International Conference on Learning Representations*.
- Xinyu Liu, Shuyu Shen, Boyan Li, Peixian Ma, Runzhi Jiang, Yuyu Luo, Yuxin Zhang, Ju Fan, Guoliang Li, and Nan Tang. 2024d. [A survey of nl2sql with large language models: Where are we, and where are we going?](#) *Preprint*, arXiv:2408.05109.
- Yan Liu, Xiaokang Chen, Yan Gao, Zhe Su, Fengji Zhang, Daoguang Zan, Jian-Guang Lou, Pin-Yu Chen, and Tsung-Yi Ho. 2023b. [Uncovering and quantifying social biases in code generation](#). In *Advances in Neural Information Processing Systems*, volume 36, pages 2368–2380. Curran Associates, Inc.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. [Roberta: A robustly optimized bert pretraining approach](#). *Preprint*, arXiv:1907.11692.
- Lin Long, Rui Wang, Ruixuan Xiao, Junbo Zhao, Xiao Ding, Gang Chen, and Haobo Wang. 2024. [On LLMs-driven synthetic data generation, curation, and evaluation: A survey](#). In *Findings of the Association for Computational Linguistics ACL 2024*, pages 11065–11082, Bangkok, Thailand and virtual meeting. Association for Computational Linguistics.
- Anton Lozhkov, Raymond Li, Loubna Ben Allal, Federico Cassano, Joel Lamy-Poirier, Nouamane Tazi, Ao Tang, Dmytro Pykhtar, Jiawei Liu, Yuxiang Wei, Tianyang Liu, Max Tian, Denis Kocetkov, Arthur Zucker, Younes Belkada, Zijian Wang, Qian Liu, Dmitry Abulkhanov, Indraneil Paul, Zhuang Li, Wen-Ding Li, Megan Risdal, Jia Li, Jian Zhu, Terry Yue Zhuo, Evgenii Zheltonozhskii, Nii Osa Osae Dade, Wenhao Yu, Lucas Krauß, Naman Jain, Yixuan Su, Xuanli He, Manan Dey, Edoardo Abati, Yekun Chai, Niklas Muennighoff, Xiangru Tang, Muhtasham Oblokulov, Christopher Akiki, Marc Marone, Chenghao Mou, Mayank Mishra, Alex Gu, Binyuan Hui, Tri Dao, Armel Zebaze, Olivier Dehaene, Nicolas Patry, Canwen Xu, Julian McAuley, Han Hu, Torsten Scholak, Sebastien Paquet, Jennifer Robinson, Carolyn Jane Anderson, Nicolas Chapados, Mostofa Patwary, Nima Tajbakhsh, Yacine Jernite, Carlos Muñoz Ferrandis, Lingming Zhang, Sean Hughes, Thomas Wolf, Arjun Guha, Leandro von Werra, and Harm de Vries. 2024. [Starcoder 2 and the stack v2: The next generation](#). *Preprint*, arXiv:2402.19173.
- Chris Lu, Cong Lu, Robert Tjarko Lange, Jakob Foerster, Jeff Clune, and David Ha. 2024a. [The AI scientist: Towards fully automated open-ended scientific discovery](#). *arXiv preprint arXiv:2408.06292*.
- Keming Lu, Hongyi Yuan, Zheng Yuan, Runji Lin, Junyang Lin, Chuanqi Tan, Chang Zhou, and Jingren Zhou. 2024b. [#instag: Instruction tagging for analyzing supervised fine-tuning of large language models](#). In *The Twelfth International Conference on Learning Representations*.
- Ziyang Luo, Can Xu, Pu Zhao, Qingfeng Sun, Xubo Geng, Wenxiang Hu, Chongyang Tao, Jing Ma, Qingwei Lin, and Daxin Jiang. 2024. [Wizardcoder: Empowering code large language models with evol-instruct](#). In *The Twelfth International Conference on Learning Representations*.
- Michael R. Lyu, Baishakhi Ray, Abhik Roychoudhury, Shin Hwei Tan, and Patanamon Thongtanunam. 2024. [Automatic programming: Large language models and beyond](#). *Preprint*, arXiv:2405.02213.
- Pingchuan Ma, Tsun-Hsuan Wang, Minghao Guo, Zhiqing Sun, Joshua B Tenenbaum, Daniela Rus, Chuang Gan, and Wojciech Matusik. 2024. [Llm and simulation as bilevel optimizers: A new paradigm](#)

- to advance physical scientific discovery. In *International Conference on Machine Learning (ICML)*.
- Rien Maertens, Charlotte Van Petegem, Niko Strijbol, Toon Baeyens, Arne Carla Jacobs, Peter Dawyndt, and Bart Mesuere. 2022. [Dolos: Language-agnostic plagiarism detection in source code](#). *Journal of Computer Assisted Learning*, 38(4):1046–1061.
- Federico Mora, Justin Wong, Haley Lepe, Sahil Bhatia, Karim Elmaaroufi, George Varghese, Joseph E. Gonzalez, Elizabeth Polgreen, and Sanjit A. Seshia. 2024. [Synthetic programming elicitation and repair for text-to-code in very low-resource programming languages](#). *Preprint*, arXiv:2406.03636.
- Theo X. Olausson, Jeevana Priya Inala, Chenglong Wang, Jianfeng Gao, and Armando Solar-Lezama. 2024. [Is self-repair a silver bullet for code generation?](#) In *The Twelfth International Conference on Learning Representations*.
- OpenAI. 2023. GPT-3.5 turbo fine-tuning and api updates. <https://openai.com/index/gpt-3-5-turbo-fine-tuning-and-api-updates/>. Accessed: June 7, 2024.
- Ulyana Piterberg, Lerrel Pinto, and Rob Fergus. 2024. [Training language models on synthetic edit sequences improves code synthesis](#). *Preprint*, arXiv:2410.02749.
- Martin Riddell, Ansong Ni, and Arman Cohan. 2024. [Quantifying contamination in evaluating code generation capabilities of language models](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 14116–14137, Bangkok, Thailand. Association for Computational Linguistics.
- Baptiste Rozière, Jonas Gehring, Fabian Gloeckle, Sten Sootla, Itai Gat, Xiaoqing Ellen Tan, Yossi Adi, Jingyu Liu, Romain Sauvestre, Tal Remez, Jérémy Rapin, Artyom Kozhevnikov, Ivan Evtimov, Joanna Bitton, Manish Bhatt, Cristian Canton Ferrer, Aaron Grattafiori, Wenhan Xiong, Alexandre Défossez, Jade Copet, Faisal Azhar, Hugo Touvron, Louis Martin, Nicolas Usunier, Thomas Scialom, and Gabriel Synnaeve. 2024. [Code llama: Open foundation models for code](#). *Preprint*, arXiv:2308.12950.
- Ozan Sener and Silvio Savarese. 2018. [Active learning for convolutional neural networks: A core-set approach](#). In *International Conference on Learning Representations*.
- Yunfan Shao, Linyang Li, Yichuan Ma, Peiji Li, Demin Song, Qinyuan Cheng, Shimin Li, Xiaonan Li, Pengyu Wang, Qipeng Guo, Hang Yan, Xipeng Qiu, Xuanjing Huang, and Dahua Lin. 2024. [Case2code: Learning inductive reasoning with synthetic data](#). *Preprint*, arXiv:2407.12504.
- Xinyu She, Yue Liu, Yanjie Zhao, Yiling He, Li Li, Chakkrit Tantithamthavorn, Zhan Qin, and Haoyu Wang. 2023. [Pitfalls in language models for code intelligence: A taxonomy and survey](#). *Preprint*, arXiv:2310.17903.
- Zhiqiang Shen, Tianhua Tao, Liqun Ma, Willie Neiswanger, Zhengzhong Liu, Hongyi Wang, Bowen Tan, Joel Hestness, Natalia Vassilieva, Daria Sobolova, and Eric Xing. 2024. [Sлимпajama-dc: Understanding data combinations for llm training](#). *Preprint*, arXiv:2309.10818.
- Noah Shinn, Federico Cassano, Ashwin Gopinath, Karthik R Narasimhan, and Shunyu Yao. 2023. [Reflection: language agents with verbal reinforcement learning](#). In *Thirty-seventh Conference on Neural Information Processing Systems*.
- Alexander G Shypula, Aman Madaan, Yimeng Zeng, Uri Alon, Jacob R. Gardner, Yiming Yang, Milad Hashemi, Graham Neubig, Parthasarathy Ranganathan, Osbert Bastani, and Amir Yazdanbakhsh. 2024. [Learning performance-improving code edits](#). In *The Twelfth International Conference on Learning Representations*.
- Chenglei Si, Diyi Yang, and Tatsunori Hashimoto. 2024. [Can LLMs generate novel research ideas? a large-scale human study with 100+ nlp researchers](#). *arXiv preprint arXiv:2409.04109*.
- David Silver, Thomas Hubert, Julian Schrittwieser, Ioannis Antonoglou, Matthew Lai, Arthur Guez, Marc Lanctot, Laurent Sifre, Dharsan Kumaran, Thore Graepel, et al. 2018. A general reinforcement learning algorithm that masters chess, shogi, and go through self-play. *Science*, 362(6419):1140–1144.
- Avi Singh, John D Co-Reyes, Rishabh Agarwal, Ankesh Anand, Piyush Patil, Xavier Garcia, Peter J Liu, James Harrison, Jaehoon Lee, Kelvin Xu, Aaron T Parisi, Abhishek Kumar, Alexander A Alemi, Alex Rizkowsky, Azade Nova, Ben Adlam, Bernd Bohnet, Gamaleldin Fathy Elsayed, Hanie Sedghi, Igor Mordatch, Isabelle Simpson, Izzeddin Gur, Jasper Snoek, Jeffrey Pennington, Jiri Hron, Kathleen Kenealy, Kevin Swersky, Kshiteej Mahajan, Laura A Culp, Lechao Xiao, Maxwell Bileschi, Noah Constant, Roman Novak, Rosanne Liu, Tris Warkentin, Yamini Bansal, Ethan Dyer, Behnam Neyshabur, Jascha Sohl-Dickstein, and Noah Fiedel. 2024. [Beyond human data: Scaling self-training for problem-solving with language models](#). *Transactions on Machine Learning Research*. Expert Certification.
- Chia-Yi Su and Collin McMillan. 2024. [Distilled gpt for source code summarization](#). *Preprint*, arXiv:2308.14731.
- Qiushi Sun, Zhirui Chen, Fangzhi Xu, Kanzhi Cheng, Chang Ma, Zhangyue Yin, Jianing Wang, Chengcheng Han, Renyu Zhu, Shuai Yuan, Qipeng Guo, Xipeng Qiu, Pengcheng Yin, Xiaoli Li, Fei Yuan, Lingpeng Kong, Xiang Li, and Zhiyong Wu. 2024a. [A survey of neural code intelligence: Paradigms, advances and beyond](#). *Preprint*, arXiv:2403.14734.

- Zhihong Sun, Chen Lyu, Bolun Li, Yao Wan, Hongyu Zhang, Ge Li, and Zhi Jin. 2024b. [Enhancing code generation performance of smaller models by distilling the reasoning ability of LLMs](#). In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 5878–5895, Torino, Italia. ELRA and ICCL.
- Marc Szafraniec, Baptiste Roziere, Hugh James Leather, Patrick Labatut, Francois Charton, and Gabriel Synnaeve. 2023. [Code translation with compiler representations](#). In *The Eleventh International Conference on Learning Representations*.
- Zhen Tan, Dawei Li, Song Wang, Alimohammad Beigi, Bohan Jiang, Amrita Bhattacharjee, Mansooreh Karami, Jundong Li, Lu Cheng, and Huan Liu. 2024. [Large language models for data annotation: A survey](#). *Preprint*, arXiv:2402.13446.
- Runchu Tian, Yining Ye, Yujia Qin, Xin Cong, Yankai Lin, Yinxu Pan, Yesai Wu, Hui Haotian, Liu Weichuan, Zhiyuan Liu, and Maosong Sun. 2024. [DebugBench: Evaluating debugging capability of large language models](#). In *Findings of the Association for Computational Linguistics ACL 2024*, pages 4173–4198, Bangkok, Thailand and virtual meeting. Association for Computational Linguistics.
- Nalin Wadhwa, Jui Pradhan, Atharv Sonwane, Surya Prakash Sahu, Nagarajan Natarajan, Aditya Kanade, Suresh Parthasarathy, and Sriram Rajamani. 2023. [Frustrated with code quality issues? llms can help!](#) *Preprint*, arXiv:2309.12938.
- Yao Wan, Yang He, Zhangqian Bi, Jianguo Zhang, Hongyu Zhang, Yulei Sui, Guandong Xu, Hai Jin, and Philip S. Yu. 2023. [Deep learning for code intelligence: Survey, benchmark and toolkit](#). *Preprint*, arXiv:2401.00288.
- Jiahao Wang, Bolin Zhang, Qianlong Du, Jiajun Zhang, and Dianhui Chu. 2024a. [A survey on data selection for llm instruction tuning](#). *Preprint*, arXiv:2402.05123.
- Xin Wang, Yasheng Wang, Yao Wan, Fei Mi, Yitong Li, Pingyi Zhou, Jin Liu, Hao Wu, Xin Jiang, and Qun Liu. 2022. [Compilable neural code generation with compiler feedback](#). In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 9–19, Dublin, Ireland. Association for Computational Linguistics.
- Xingyao Wang, Hao Peng, Reyhaneh Jabbarvand, and Heng Ji. 2024b. [LETI: Learning to generate from textual interactions](#). In *Findings of the Association for Computational Linguistics: NAACL 2024*, pages 223–239, Mexico City, Mexico. Association for Computational Linguistics.
- Yizhong Wang, Yeganeh Kordi, Swaroop Mishra, Alisa Liu, Noah A. Smith, Daniel Khashabi, and Hannaneh Hajishirzi. 2023. [Self-instruct: Aligning language models with self-generated instructions](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 13484–13508, Toronto, Canada. Association for Computational Linguistics.
- Zige Wang, Wanjun Zhong, Yufei Wang, Qi Zhu, Fei Mi, Baojun Wang, Lifeng Shang, Xin Jiang, and Qun Liu. 2024c. [Data management for training large language models: A survey](#). *Preprint*, arXiv:2312.01700.
- Yuxiang Wei, Zhe Wang, Jiawei Liu, Yifeng Ding, and Lingming Zhang. 2024. [Magicoder: Empowering code generation with OSS-instruct](#). In *Forty-first International Conference on Machine Learning*.
- Martin Weysow, Aton Kamanda, and Houari Sahraoui. 2024. [Codeultrafeedback: An llm-as-a-judge dataset for aligning large language models to coding preferences](#). *Preprint*, arXiv:2403.09032.
- Kyle Wong, Alfonso Amayuelas, Liangming Pan, and William Yang Wang. 2024. [Distilrr: Transferring code repair for low-resource programming languages](#). *Preprint*, arXiv:2406.14867.
- Tongtong Wu, Weigang Wu, Xingyu Wang, Kang Xu, Suyu Ma, Bo Jiang, Ping Yang, Zhenchang Xing, Yuan-Fang Li, and Gholamreza Haffari. 2024. [Versicode: Towards version-controllable code generation](#). *arXiv preprint arXiv:2406.07411*.
- Sang Michael Xie, Hieu Pham, Xuanyi Dong, Nan Du, Hanxiao Liu, Yifeng Lu, Percy Liang, Quoc V Le, Tengyu Ma, and Adams Wei Yu. 2023. [Doremi: Optimizing data mixtures speeds up language model pretraining](#). In *Thirty-seventh Conference on Neural Information Processing Systems*.
- Can Xu, Qingfeng Sun, Kai Zheng, Xiubo Geng, Pu Zhao, Jiazhan Feng, Chongyang Tao, Qingwei Lin, and Daxin Jiang. 2024a. [WizardLM: Empowering large pre-trained language models to follow complex instructions](#). In *The Twelfth International Conference on Learning Representations*.
- Xiaohan Xu, Ming Li, Chongyang Tao, Tao Shen, Reynold Cheng, Jinyang Li, Can Xu, Dacheng Tao, and Tianyi Zhou. 2024b. [A survey on knowledge distillation of large language models](#). *Preprint*, arXiv:2402.13116.
- Jiayi Yang, Binyuan Hui, Min Yang, Jian Yang, Junyang Lin, and Chang Zhou. 2024. [Synthesizing text-to-SQL data from weak and strong LLMs](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 7864–7875, Bangkok, Thailand. Association for Computational Linguistics.
- Yifan Yao, Jinhao Duan, Kaidi Xu, Yuanfang Cai, Zhibo Sun, and Yue Zhang. 2024. [A survey on large language model \(llm\) security and privacy: The good, the bad, and the ugly](#). *High-Confidence Computing*, 4(2):100211.

- Zhaojian Yu, Xin Zhang, Ning Shang, Yangyu Huang, Can Xu, Yishujie Zhao, Wenxiang Hu, and Qiufeng Yin. 2024. [WaveCoder: Widespread and versatile enhancement for code large language models by instruction tuning](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 5140–5153, Bangkok, Thailand. Association for Computational Linguistics.
- Daoguang Zan, Bei Chen, Fengji Zhang, Dianjie Lu, Bingchao Wu, Bei Guan, Wang Yongji, and Jian-Guang Lou. 2023. [Large language models meet NL2Code: A survey](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 7443–7464, Toronto, Canada. Association for Computational Linguistics.
- Weihao Zeng, Can Xu, Yingxiu Zhao, Jian-Guang Lou, and Weizhu Chen. 2024. [Automatic instruction evolving for large language models](#). *Preprint*, arXiv:2406.00770.
- Dylan Zhang, Shizhe Diao, Xueyan Zou, and Hao Peng. 2024a. [Plum: Preference learning plus test cases yields better code language models](#). *Preprint*, arXiv:2406.06887.
- Quanjun Zhang, Chunrong Fang, Yang Xie, YuXiang Ma, Weisong Sun, Yun Yang, and Zhenyu Chen. 2024b. [A systematic literature review on large language models for automated program repair](#). *Preprint*, arXiv:2405.01466.
- Quanjun Zhang, Chunrong Fang, Yang Xie, Yaxin Zhang, Yun Yang, Weisong Sun, Shengcheng Yu, and Zhenyu Chen. 2024c. [A survey on large language models for software engineering](#). *Preprint*, arXiv:2312.15223.
- Yi Zhang, Jan Deriu, George Katsogiannis-Meimarakis, Catherine Kosten, Georgia Koutrika, and Kurt Stockinger. 2023. [Sciencebenchmark: A complex real-world benchmark for evaluating natural language to sql systems](#). *Preprint*, arXiv:2306.04743.
- Ziyin Zhang, Chaoyu Chen, Bingchang Liu, Cong Liao, Zi Gong, Hang Yu, Jianguo Li, and Rui Wang. 2024d. [Unifying the perspectives of nlp and software engineering: A survey on language models for code](#). *Preprint*, arXiv:2311.07989.
- Qinkai Zheng, Xiao Xia, Xu Zou, Yuxiao Dong, Shan Wang, Yufei Xue, Zihan Wang, Lei Shen, Andi Wang, Yang Li, Teng Su, Zhilin Yang, and Jie Tang. 2023. [Codegeex: A pre-trained model for code generation with multilingual benchmarking on humaneval-x](#). In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 5673–5684.
- Tianyu Zheng, Ge Zhang, Tianhao Shen, Xueling Liu, Bill Yuchen Lin, Jie Fu, Wenhu Chen, and Xiang Yue. 2024. [OpenCodeInterpreter: Integrating code generation with execution and refinement](#). In *Findings of the Association for Computational Linguistics ACL 2024*, pages 12834–12859, Bangkok, Thailand and virtual meeting. Association for Computational Linguistics.
- Shuyan Zhou, Uri Alon, Sumit Agarwal, and Graham Neubig. 2023. [CodeBERTScore: Evaluating code generation with pretrained models of code](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 13921–13937, Singapore. Association for Computational Linguistics.
- Qihao Zhu, Daya Guo, Zhihong Shao, Dejian Yang, Peiyi Wang, Runxin Xu, Y. Wu, Yukun Li, Huazuo Gao, Shirong Ma, Wangding Zeng, Xiao Bi, Zihui Gu, Hanwei Xu, Damai Dai, Kai Dong, Liyue Zhang, Yishi Piao, Zhibin Gou, Zhenda Xie, Zhewen Hao, Bingxuan Wang, Junxiao Song, Deli Chen, Xin Xie, Kang Guan, Yuxiang You, Aixin Liu, Qiushi Du, Wenjun Gao, Xuan Lu, Qinyu Chen, Yaohui Wang, Chengqi Deng, Jiashi Li, Chenggang Zhao, Chong Ruan, Fuli Luo, and Wenfeng Liang. 2024. [Deepseek-coder-v2: Breaking the barrier of closed-source models in code intelligence](#). *Preprint*, arXiv:2406.11931.
- Terry Yue Zhuo. 2024. [ICE-score: Instructing large language models to evaluate code](#). In *Findings of the Association for Computational Linguistics: EACL 2024*, pages 2232–2242, St. Julian’s, Malta. Association for Computational Linguistics.

A Appendix

This section provides two supplementary discussions. First, we delineate the scope of this survey, distinguishing it from related concepts such as knowledge distillation and data augmentation. Second, we recommend cost-effective yet high-quality large language models (LLMs) for code-related data generation. Given the rapid evolution of LLMs, we refer readers to our GitHub repository for the latest model updates.

A.1 Survey Scope

This survey aims to comprehensively explore data synthesis and filtering techniques used in building CodeLLMs for downstream tasks such as code generation, repair, translation, and documentation. Our focus is data engineering approaches rather than *knowledge distillation* algorithms, which investigate techniques for transferring knowledge from large models (i.e. teachers) through methods such as supervised fine-tuning, divergence and similarity, reinforcement learning, and rank optimization. Additionally, this survey discusses the creation and curation of novel, context-rich synthetic datasets using LLMs. In contrast, traditional *data augmentation* techniques such as paraphrasing and back-translation expand training datasets in a somewhat mechanistic manner.

We reviewed and analyzed over 50 research papers on data synthesis and filtering, most of which were published within the last two years. To offer a structured overview, we categorize these works into a taxonomy of 23 sub-topics, as shown in Figure 2. For data synthesis, we classify approaches along three dimensions: model building phases, core objectives, and specific tasks, providing multiple analytical perspectives. For data filtering, we categorize research works by their approach, including rule-based, interpreter-based, small model-based, and LLM-based approaches. Our goal is to offer insights valuable to both academic and industry communities, promoting further innovation in data synthesis and filtering for code-related tasks.

A.2 LLM Selection Considerations

LLMs play a central role in both synthesis and filtering, making model selection critical for performance optimization. Based on empirical observations, we recommend several commonly used models. If cost is not a concern, GPT-4 or GPT-

4o² remains the top choice, consistently delivering high-quality results. For a cost-effective alternative, the open-source Llama 3.1-405B³, particularly its *int4* quantized version⁴, offers a strong trade-off between quality and efficiency, running on four H100 GPUs. Another viable option is Qwen2.5-72B-Instruct⁵, known for fast execution and strong performance in code-related tasks. Finally, DeepSeek-Coder-V2-Instruct⁶, an open-source Mixture-of-Experts (MoE) CodeLLM, achieves performance comparable to GPT-4-Turbo in code-specific tasks.

Besides, researchers and practitioners should be aware of the specific synthetic data usage policies associated with various LLMs, particularly when building commercial products. For instance, Llama 2 and 3 restrict the use of generated outputs for training other AI models, whereas Llama 3.1 and 3.2 have updated these policies⁷. The Qwen model, on the other hand, requires explicit attribution, such as prominently displaying “Built with Qwen” or “Improved using Qwen” in product documentation when its outputs are used for creating, training, fine-tuning, or improving an AI model⁸.

For inference, vLLM⁹ is a fast and user-friendly library that supports high-throughput serving with various decoding algorithms, such as parallel sampling and beam search. It is compatible with most popular open-source models on Hugging Face, including Transformer-based LLMs (e.g. Llama (Dubey et al., 2024)), Mixture-of-Experts LLMs (e.g. Mixtral¹⁰), Embedding Models (e.g. E5-Mistral¹¹), and Multi-modal LLMs (e.g. LLaVA (Liu et al., 2023a)). It also allows for offline batched inference on datasets or sending requests through an OpenAI-compatible API server, providing a convenient solution for large-scale data experimentation and serving as a valuable tool for advancing research and development in this field.

²<https://platform.openai.com/docs/models/gpt-4o>

³<https://huggingface.co/meta-llama/Llama-3.1-405B>

⁴<https://huggingface.co/neuralmagic/Meta-Llama-3.1-405B-Instruct-quantized.w4a16>

⁵<https://huggingface.co/Qwen/Qwen2.5-72B-Instruct>

⁶<https://huggingface.co/deepseek-ai/DeepSeek-Coder-V2-Instruct>

⁷<https://www.llama.com/faq/>

⁸<https://huggingface.co/Qwen/Qwen2.5-72B-Instruct/blob/main/LICENSE>

⁹<https://github.com/vllm-project/vllm>

¹⁰<https://huggingface.co/mistralai>

¹¹<https://huggingface.co/intfloat/e5-mistral-7b-instruct>