

Towards Top-Down Automated Development in Limited Scopes: A Neuro-Symbolic Framework from Expressibles to Executables

Jian Gu
Monash University
Melbourne, Australia
jian.gu@monash.edu

Harald C. Gall
University of Zurich
Zurich, Switzerland
gall@ifi.uzh.ch

ABSTRACT

Deep code generation is a topic of deep learning for software engineering (DLSE), which adopts neural models to generate code for the intended functions. Since end-to-end neural methods lack domain knowledge and software hierarchy awareness, they tend to perform poorly w.r.t project-level tasks. To systematically explore the potential improvements of code generation, we let it participate in the whole top-down development from *expressibles* to *executables*, which is possible in limited scopes. In the process, it benefits from massive samples, features, and knowledge. As the foundation, we suggest building a taxonomy on code data, namely code taxonomy, leveraging the categorization of code information. Moreover, we introduce a three-layer semantic pyramid (SP) to associate text data and code data. It identifies the information of different abstraction levels, and thus introduces the domain knowledge on development and reveals the hierarchy of software. Furthermore, we propose a semantic pyramid framework (SPF) as the approach, focusing on software of high modularity and low complexity. SPF divides the code generation process into stages and reserves spots for potential interactions. In addition, we conceived preliminary applications in software development to confirm the neuro-symbolic framework.

KEYWORDS

code generation, deep learning, frame semantics, knowledge graph, requirement elicitation, program representation, software analytics

ACM Reference Format:

Jian Gu and Harald C. Gall. 2023. Towards Top-Down Automated Development in Limited Scopes: A Neuro-Symbolic Framework from Expressibles to Executables. In *Proceedings of the 31st ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering (ESEC/FSE '23)*, December 3–9, 2023, San Francisco, CA, USA. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/3611643.3613076>

1 INTRODUCTION

To promote the productivity of software development, using deep learning to assist related activities is a lasting research trend [16, 36]. For example, neural models trained on massive code-related data could assist programming through code generation [17]. However, we identified a few restrictions of neural methods waiting to be overcome. First, models lack the necessary domain knowledge on software development. For example, models can merely learn the idiomatic usage of the relevant API dependency but cannot ensure validity or avoid possible misuse issues. Second, models require accurate and complete functional descriptions, even related to the context and method call. The truth is that it is better to let models infer the requirement based on the software hierarchy. Third, neural methods are weak in transparency and flexibility, thereby

tend to underperform in scenarios where potential interactions are required, such as manual inspection or intervention. Based on the above considerations, we propose adopting neural models in the whole top-down development from expressibles to executables. In this way, we explore performance improvements in various scenes, by providing domain knowledge and software hierarchy to code generation. Considering the difficulty, the research scope is limited to softwares of high modularity and low complexity.

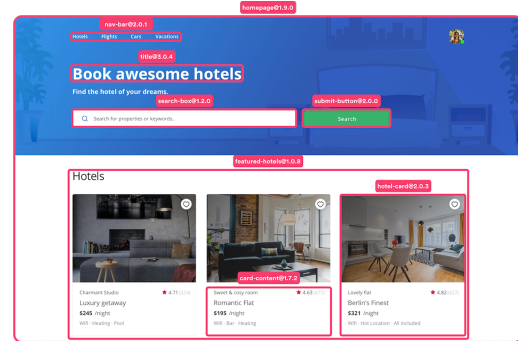


Figure 1: Diagram of modular WebUI development, © Bit.Cloud [5], where each UI component in the webpage is supported by a separate dependency, and meanwhile filled with specific resources or values.

The webpage in Fig. 1 is an illustrative example of top-down code generation. Its implementation from expressibles to executables can be a 3-step process: (1) everything starts with a prototype, which clearly describes the webpage layout. For example, there should be a navigation bar, an avatar and a search bar in the header, and a grid view to display cards in the webpage body; and (2) a knowledge engine that knows available libraries suggest API usages to fulfill the requirements, such as recommending “featured-hotels@1.x” and “hotel-card@2.x” for the grid view and cards, with the compatibility consideration; and (3) a set of utility codes accessing algorithms and data resources arrange and control UI components, including personalizing the product list, loading the product detail.

To achieve the goal of adopting code generation in the top-down development process, we start with taxonomy on code data, namely code taxonomy, in the light of code information. On the basis of code taxonomy, we designed the scheme of a three-layer semantic pyramid (SP) to associate program realizations with human intentions. In each SP layer, we take corresponding techniques to organize data in the form of graphs. By taking SP as the mental model, we plan to further construct a semantic pyramid framework (SPF). SPF makes code data high-quality and well-organized, and besides, introduces program semantics and domain knowledge.

Table 1: Categorization of code information

Group	Type	Description	Example
Attributes	Constraint	restrictions on the applicable scope	conditions, inputs, outputs
	Element	discriminative implementation details	operators, numerics, strings
	Knowledge	conventions and usages in reusing code	API dependencies, frameworks
Syntactics	Relevance	connections among relevant information	sequence, syntax trees and graphs
Semantics	Concept	entities referred to and their interrelations	contexts, identifiers (objects, values)
	Function	logical functions to be actually executed	calculations, operations, exceptions

2 APPROACH

Code Taxonomy. We define the term “code taxonomy” as the practice of categorization of code data. It is straightforward to build a taxonomy referring to software purposes or topics, however, that manner seems insufficient in reflecting the capability of methods, especially for code generation. Therefore, we describe the task from another viewpoint: *beyond the influence of software purpose or topic, by what basis could we further build taxonomy on code data?*

In our approach, we regard code data as a mixture of different types of information and disentangle this information to build code taxonomy. Considering the inherent properties of code data and the semantic associations with related text data, we classify code information into 3 groups and 6 types, as shown in Table 1. Based on the categorization of code information, we suggest quantifying the *information abundance* of data samples on each information type for various usages. For example, when we evaluate code generation respecting the difficulty of data samples, it is natural to assess the results separately for cases like whether there are API dependencies.

Semantic Pyramid. We propose a semantic-aware framework as the follow-up work of introducing code taxonomy for code generation. SPF guarantees the recognizability of code data via *code assetization* and establishes the connectivity between intentions and realizations via *semantic bridging*. Its mental model is visualized as a three-layer semantic pyramid, shown in Fig. 2.




Figure 2: Diagram of the three-layer semantic pyramid, where the code information of each layer is specified. The higher the layer, the closer to text data, and vice versa, the closer to code data.

The code information is divided into 3 abstraction layers. The intention layer cares about semantic information, namely semantic concepts and logic functions of code data. The specification layer is about the outline of code implementations, such as constraints and dependencies that cannot be ignored. The realization layer is about the detail of code implementations, where concrete elements or structures matter. Every node of any graph in the upper layer points to the complete graph in the lower layer, and each graph in the lower layer is pointed to by multiple nodes in the upper layer.

Table 2: Definitions of the datatype for code taxonomy

Datatype	Attributes			Syntactics	Semantics	
	Constraint	Element	Knowledge	Relevance	Concept	Function
Code Pattern	✗	✗	✓			
Code Template/Sketch	✓	✗	✓			
Code Instance	✓	✓	✓			
Code Property Graph				✓		
Semantic Entity					✓	✗
Semantic Pattern					✗	✓
Semantic Frame/Sketch					✓	✓

The involved information is intended to be organized in the form of graphs even though it corresponds to different abstraction levels in different layers. The reason is that graph-form expressions are ideal as the medium to explicitly represent various information, and performant to establish connections with each other. Considering the differences of the information itself and their roles in the framework, we introduce suitable techniques to build graphs and find ways to bridge different graphs, and thus associate intentions to realizations crossing the concept hierarchy of software.

Semantic Pyramid Framework. To describe the mechanism of our framework, we distinguish the involved datatypes based on their information abundance following the code taxonomy. For simplicity, we prefer the qualitative division but not the quantitative one. As shown in Table 2, we use the check mark (✓) to indicate a datatype is rich in some types of code information, and the cross mark (✗) for the opposite case. If an information type is not that important to a datatype, we leave blank the marked area.

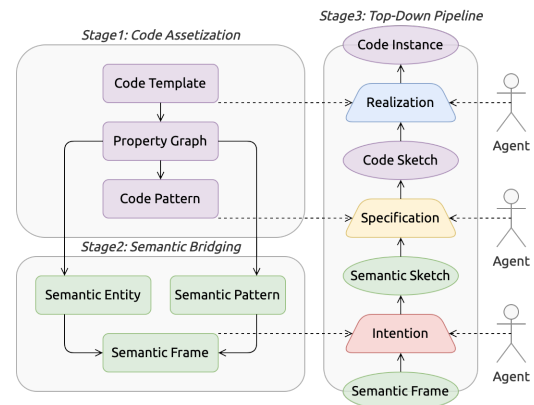


Figure 3: Overview of the semantic pyramid framework, where rectangles are databases of the corresponding datatype, trapezoids are spots for human-machine interactions, and ellipses represent data instances. The color is to distinguish datatypes and procedures.

By characterizing the role of datatypes defined following code taxonomy, we divide the mechanism of the framework into three stages, as illustrated in Fig. 3. In the first stage named *code assetization*, code data is first processed into subprograms as templates to increase the sample richness. Then, we use a property graph as the code representation for the higher recognizability of samples [39]. Eventually, we mine code patterns from code property graphs to reduce the effect of variants. In the second stage named *semantic bridging*, leveraging the informative property graphs, we

extract semantic entities and patterns to collect involved concepts and functions. Furthermore, we compose semantic frames based on concepts and functions with external knowledge bases. In the third stage named *top-down pipeline*, we introduce code generation to assist top-down development, leveraging the databases on features, knowledge, and samples from the former stages. It divides the development into stages and reserves spots for interactions, corresponding to the abstraction layers in the semantic pyramid.

2.1 Code Assetization

We define *Code Assetization* as organizing code data into samples, following code taxonomy. It focuses on the realization layer.

In the process of code assetization, SPF decouples semantics, syntactics, and attributes information from code samples. We need to generate high-quality representations of code samples [37] to support operations such as retrieval and clustering. As code data could be parsed as ASTs, CFGs, and PDGs, it is a cheap but precise practice to first convert code data into the graph form, as the joint data structure of the above-mentioned parsed results [39].

We build a sample base of code templates and code patterns through data processing. Except for the handling of the elements, the code template is roughly equivalent to the code sample, since both reserve the divergence of concrete implementations. In contrast, the code pattern represents a bunch of similar code samples and has no tolerance for variants. Leveraging diff computation [6] and subprogram division [1], code patterns are extracted from code templates but shorter in length and fewer in amount. As the property graph is preferred to represent code data, code patterns could be seen as the idiomatic usage of a partial code template.

The realization layer is focused on the synergy of code patterns with extra knowledge bases. Since the divergence in code samples is mainly caused by contexts and dependencies, we thus define two types of knowledge bases to complement the information missed by code patterns. One knowledge base is on internal context, which determines the constraints of code templates. The other one is on external dependencies, which determines the API dependencies.

2.2 Semantic Bridging

We define *Semantic Bridging* as seeking features and knowledge to express requirements. It is about intention and specification layers.

For semantic bridging, we propose to jointly utilize semantic parsing and knowledge graph. On one side, semantic parsing converts natural language utterances to machine instructions [24]. It mines information from text data with techniques like named entity recognition and relation extraction. On the other side, knowledge graphs use graph structures to organize semantic entities and their interrelations [13]. It provides further flexibility by linking together with available commonsense or domain knowledge graphs.

The intention layer relies on a linguistic meaning theory called frame semantics [7]. A semantic frame is a conceptual structure of the involved participants and is used to describe events, scenes, etc. It usually contains frame elements as semantic roles in the formal definition and lexical units that evoke other frames. That makes it very similar to a function module or class definition in programs. This similarity inspires us to build a corpus oriented to program intentions, for example, based on FRAME_{NET} lexical database [29].

The specification layer emphasizes the semantics information of code patterns, which is represented by its own description and the relations with others. The description could come from related text data, such as comments, API directives, and even identifiers. The relations could be meaning representation of text data [38], or the logic representation of code data [22]. Compared with the intention layer, the specification layer differs in granularity, even though they both revolve around semantic entities and interrelations.

2.3 Top-Down Pipeline

In the framework, the top-down pipeline associates text code with code data through two transitions: the transition from intentions to specifications, and from specifications to realizations. Their medium data is respectively the semantic sketch and the code sketch [31].

In the transition from intentions to specifications, the semantic sketch purely cares about concepts and functions. Both semantic frame and semantic sketch take concepts as nodes and functions as edges. However, compared with code sketches, their concepts could be more abstract while functions could be more complex. Thereby, the transition aims to adopt alignment methods to eliminate the semantic gap between semantic sketch and code sketch, especially the granularity inconsistency. Other concerns are the details in the transition, such as issues on constraint and knowledge.

The transition medium from specifications to realizations is the code sketch. The code sketch is different from the code template, even though both are graphs taking code patterns as nodes. The code sketch takes semantic dependencies between code patterns as edges while the code template takes actual dependencies in samples as edges. Therefore, assuming we have a code sketch from the upstream, we could retrieve code templates for each code pattern in the given code sketch, then composite or generate code templates as candidates. Further, with manual inspection or intervention on elements and structures, the optimal code instance is determined.

2.4 Scope of Application

Due to the complexity and diversity in actual development, it is important to find suitable application scopes for the framework. Thereby, we considered some as preliminary targets.

One conceived application is assistive top-down developments for specific purposes or topics, including WebUI programming as mentioned. Similarly, mobile application development is worth attention, since it follows a conventional standard to implement preset interactions. For example, an app can be recognized as a transition graph of activities, where each activity corresponds to a UI page and its functional code. The implementation of an activity is usually decoupled as layout code, functional code, and method call hierarchy [3], and thus well fits the requirements of SPF.

In addition, to reveal the potentiality of SPF in associating intentions with realizations, automated construction of modular API dependency can be a topic. It elicits and reacts to potential development needs. For example, a new web framework is at its early stage and wants a series of features, already implemented in similar frameworks but partially and scantly. A solution is to massively spawn and smoothly assemble the migrated implementations. Leveraging SPF, we refine the process and let it be effective and efficient.

3 ROADMAP

To obtain experience in building the code taxonomy, and form an effective methodology, we divide the whole process into three steps. First, we start with neural methods that learn from both code properties and data distributions, focusing on program representation and code generation. Then, we build code taxonomies in limited application scopes and evaluate their effectiveness and drawbacks. Finally, we seek decent models for a complete framework and increase the scale of the engineering to explore its applicability to real-world use. The milestones to be achieved are as follows:

- First, we would conduct a mining study on open source software distinguished by the purpose or topic. The object data is coding patterns, complemented by the information on context and API dependence.
- Further, we would seek solutions for semantic associations between program sketches and functional descriptions or API directives, but more focused on graph representations of multiple granularity or modality.
- To bridge intentions and specifications in a formal way, we plan on constructing a lexical database leveraging frame semantics and knowledge graphs. Besides, we connect it to available instances from relevant fields.
- To bridge specifications and realizations and promote code reusability, we plan on building a code corpus, where the data are templated but composable subprograms, with the compatibility support of code variants.
- Once the code taxonomy is constructed, we would design retrieval-based neural models or inference engines to pair with this knowledge base. The framework should apply to other scenes where the semantic pyramid fits.
- Eventually, we plan to study the behavior and impact of the framework on building programs with a certain complexity, moreover, on building highly modular software or libraries that can satisfy actual development requirements.

4 RELATED WORK

For code taxonomy, we mainly refer to other practices on taxonomy for deep learning. Meanwhile, for SP and SPF, we consider tasks connecting code data and text data as related work. Besides, we are aware of topics involved in implementing the framework, including code features, and topics about actual applications.

Taxonomy for Deep Learning. When some tasks are hard to learn or evaluate, it is a strategy to build a taxonomy. For example, in software engineering, there is a taxonomy on build failures in continuous integration [33], a taxonomy of faults in deep learning systems [12]. In other areas, taxonomy for deep learning even plays a considerable role. In natural language processing, FRAME_{NET} motivated the study on automatic semantic role labeling [29]. In computer vision, IMAGE_{NET} [30] is an image database for visual object recognition, it follows the WORD_{NET} [25] hierarchy, where each node is depicted by massive samples.

Code Retrieval, Generation and Comprehension. Sometimes, programmers want alternatives to current implementations, for example, AROMA [23] recommends code snippets based on the similarity of structural features. To search code snippets for the given query, CODENN aligns the semantics of code data and text data by

jointly modeling their representations in the same vector space [9]. CSRS [4] considers both relevance matching and semantic matching, while YOGO [27] recognizes different but mathematically-equivalent alternatives based on dataflow graphs and rewrite rules. Inspired by machine translation, code generation translates descriptions to implementations, such as the pseudocode-to-code one by SEMANTIC_{SCAFFOLD} [42]. TRANX introduced a transition system to generate formal meaning representations for code generation [40]. Code comprehension uses proper natural utterance to describe the given code data, which could help understand programs: CONTEXT2NAME [2] generates meaningful variable names based on the usage contexts, ADAMO [8] assembles foundation models to describe the functionality of code snippets, CODEQA [21] demonstrates the value of generating textual answers for the pairs of code snippet and various types of question.

Program Analysis and Reasoning. By adopting changes to code data, the critical features can be stressed, such as S4EQ [14], which proves the equivalence by computing a verifiable sequence of rewrite rules, and PERSES [32], which reduces programs to smaller ones in a syntax-guided way, and further help localizes bugs. Also, program analysis can mine code features. For example, LUPA [34] supports fine-grained analysis on the usage of programming language, PYART [11] extracts features from the context of the program point for API recommendation [10]. Besides, code clone detection and code evolution promote code reuse [41]. In addition, neuro-symbolic approaches shown the reasoning capability in vision and natural language, and might work on code data, such as SCALLOP [20].

Assisted Programming and Development. There have proposed performant methods to assist programming, CODEGEN [26] and CODET5 [35] introduce large-scale language models to support common coding tasks, such as code autocompletion, code translation. In contrast, assisted development focuses on the solution itself but not the programming process. For example, ALPHACODE [19] applies sampling, filtering, clustering on program samples to generate human-level submissions for competitive programming challenges. Some work is oriented to other tasks in the development, for example, TRANSCODER improves the application of back-translation for program translation [15, 28], SKCODER reuses the code sketch, extracted from similar code, for reliable code generation [18].

5 CONCLUSIONS

To summarize, to overcome a few restrictions of neural methods, we suggest building a code taxonomy in light of code information. And then, we propose a semantic-aware framework, taking the semantic pyramid as its mental model, for top-down code generation.

SPF promotes code reusability through solid recognizability of code data and intense connectivity from expressibles to executables. SPF indicates a novel framework of semantic awareness, especially when centered on the graph-form expression of code-related data. Its capability to support inspection or manipulation, requirement elicitation, and multimodal representation is beneficial.

As an exploratory work, top-down code generation is novel but faces challenges. For example, human-machine interactions could be very heavy in some cases where requirements are hard to capture or fulfill. Thus, we focus on realizing the proposed framework in limited scopes and seek continuous evolutions thereafter.

REFERENCES

- [1] Ferran Alet, Javier Lopez-Contreras, James Koppel, Maxwell Nye, Armando Solar-Lezama, Tomas Lozano-Perez, Leslie Pack Kaelbling, and Joshua B. Tenenbaum. 2021. A large-scale benchmark for few-shot program induction and synthesis. In *ICML*.
- [2] Rohan Bavishi, Michael Pradel, and Koushik Sen. 2018. Context2Name: A Deep Learning-Based Approach to Infer Natural Variable Names from Usage Contexts. *ArXiv abs/1809.05193* (2018).
- [3] Sen Chen, Lingling Fan, Chunyang Chen, Ting Su, Wenhe Li, Yang Liu, and Lihua Xu. 2019. StoryDroid: Automated Generation of Storyboard for Android Apps. *2019 IEEE/ACM 41st International Conference on Software Engineering (ICSE)* (2019), 596–607.
- [4] Yiran Cheng and Li Kuang. 2022. CSRS: Code Search with Relevance Matching and Semantic Matching. *ArXiv abs/2203.07736* (2022).
- [5] Cocycles. 2023. Bit Cloud: How Teams Build Composable Software Together. Retrieved April 30, 2023 from <https://bit.cloud> (Approved Use of the Illustration for Academic Purposes on August 20, 2023 by the Head of Product @ Bit).
- [6] Jean-Rémy Falleri, Floréal Morandat, Xavier Blanc, Matias Martinez, and Monperrus Martin. 2014. Fine-grained and accurate source code differencing. *Proceedings of the 29th ACM/IEEE international conference on Automated software engineering* (2014).
- [7] Charles J. Fillmore and Collin F. Baker. 2001. Frame semantics for text understanding.
- [8] Jian Gu, Pasquale Salza, and Harald C. Gall. 2022. Assemble Foundation Models for Automatic Code Summarization.
- [9] Xiaodong Gu, Hongyu Zhang, and Sunghun Kim. 2018. Deep Code Search. *2018 IEEE/ACM 40th International Conference on Software Engineering (ICSE)* (2018), 933–944.
- [10] Xiaodong Gu, Hongyu Zhang, D. Zhang, and Sunghun Kim. 2016. Deep API learning. *Proceedings of the 2016 24th ACM SIGSOFT International Symposium on Foundations of Software Engineering* (2016).
- [11] Xincheng He, Lei Xu, X. Zhang, Rui Hao, Yang Feng, and Baowen Xu. 2021. PyART: Python API Recommendation in Real-Time. *2021 IEEE/ACM 43rd International Conference on Software Engineering (ICSE)* (2021), 1634–1645.
- [12] Nargiz Humbatova, Gunel Jahangirova, Gabriele Bavota, Vincenzo Riccio, Andrea Stocco, and Paolo Tonella. 2020. Taxonomy of Real Faults in Deep Learning Systems. *2020 IEEE/ACM 42nd International Conference on Software Engineering (ICSE)* (2020), 1110–1121.
- [13] Shaoxiong Ji, Shirui Pan, E. Cambria, Pekka Marttinen, and Philip S. Yu. 2021. A Survey on Knowledge Graphs: Representation, Acquisition and Applications. *IEEE transactions on neural networks and learning systems* PP (2021).
- [14] Steven J Komrusch, Monperrus Martin, and Louis-Noël Pouchet. 2021. Self-Supervised Learning to Prove Equivalence Between Programs via Semantics-Preserving Rewrite Rules. *ArXiv abs/2109.10476* (2021).
- [15] Marie-Anne Lachaux, Baptiste Rozière, Lowik Chanussot, and Guillaume Lample. 2020. Unsupervised Translation of Programming Languages. *ArXiv abs/2006.03511* (2020).
- [16] Triet Huynh Minh Le, Hao Chen, and Muhammad Ali Babar. 2020. Deep Learning for Source Code Modeling and Generation. *ACM Computing Surveys (CSUR)* 53 (2020), 1–38.
- [17] Celine Lee, Justin Emile Gottschlich, Dan Roth Intel Labs, and University of Pennsylvania. 2021. Toward Code Generation: A Survey and Lessons from Semantic Parsing. *ArXiv abs/2105.03317* (2021).
- [18] Jia Li, Yongming Li, Ge Li, Zhi Jin, Yiyang Hao, and Xing Hu. 2023. SkCoder: A Sketch-based Approach for Automatic Code Generation. *2023 IEEE/ACM 45th International Conference on Software Engineering (ICSE)* (2023), 2124–2135. <https://api.semanticscholar.org/CorpusID:256826931>
- [19] Yujia Li, David Choi, Junyoung Chung, Nate Kushman, Julian Schrittwieser, Rémi Leblond, Tom Eccles, James Keeling, Felix Gimeno, Agustin Dal Lago, et al. 2022. Competition-Level Code Generation with AlphaCode. *arXiv preprint arXiv:2203.07814* (2022).
- [20] Ziyang Li, Jiani Huang, and M. Naik. 2023. Scallop: A Language for Neurosymbolic Programming. *Proceedings of the ACM on Programming Languages* 7 (2023), 1463–1487. <https://api.semanticscholar.org/CorpusID:258060005>
- [21] Chenxiao Liu and Xiaojun Wan. 2021. CodeQA: A Question Answering Dataset for Source Code Comprehension. *ArXiv abs/2109.08365* (2021).
- [22] Mingwei Liu, Xin Peng, Andrian Marcus, Christoph Treude, Xuefang Bai, Gang Lyu, Jiazhan Xie, and Xiaoxin Zhang. 2021. Learning-based extraction of first-order logic representations of API directives. *Proceedings of the 29th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering* (2021).
- [23] Sifei Luan, Di Yang, Koushik Sen, and Satish Chandra. 2019. Aroma: code recommendation via structural code search. *Proceedings of the ACM on Programming Languages* 3 (2019), 1–28.
- [24] Sungwon Lyu. 2019. Text Generation From Knowledge Graphs With Graph Transformers. <https://lyusungwon.github.io/studies/2019/09/19/graphwriter/>
- [25] George A. Miller. 1992. WordNet: A Lexical Database for English. *Commun. ACM* 38 (1992), 39–41.
- [26] Erik Nijkamp, Bo Pang, Hiroaki Hayashi, Lifu Tu, Haiquan Wang, Yingbo Zhou, Silvio Savarese, and Caiming Xiong. 2022. CodeGen: An Open Large Language Model for Code with Multi-Turn Program Synthesis. In *International Conference on Learning Representations*. <https://api.semanticscholar.org/CorpusID:252668917>
- [27] Varot Premtoon, James Koppel, and Armando Solar-Lezama. 2020. Semantic code search via equational reasoning. *Proceedings of the 41st ACM SIGPLAN Conference on Programming Language Design and Implementation* (2020).
- [28] Baptiste Rozière, J Zhang, François Charton, Mark Harman, Gabriel Synnaeve, and Guillaume Lample. 2021. Leveraging Automated Unit Tests for Unsupervised Code Translation. *ArXiv abs/2110.06773* (2021).
- [29] Josef Ruppenhofer, Michael Ellsworth, Miriam R. L. Petruck, Christopher R. Johnson, and Jan Scheffczyk. 2006. FrameNet II: Extended theory and practice.
- [30] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael S. Bernstein, Alexander C. Berg, and Li Fei-Fei. 2015. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision* 115 (2015), 211–252.
- [31] Armando Solar-Lezama, Liviu Tancau, Rastislav Bodik, Sanjit A. Seshia, and Vijay A. Saraswat. 2006. Combinatorial sketching for finite programs. *Proceedings of the 12th international conference on Architectural support for programming languages and operating systems - ASPLOS-XII* (2006).
- [32] Chengnian Sun, Yuanbo Li, Qirun Zhang, Tianxiao Gu, and Zhenqiang Su. 2018. Perse: Syntax-Guided Program Reduction. *2018 IEEE/ACM 40th International Conference on Software Engineering (ICSE)* (2018), 361–371.
- [33] Carmine Vassallo, Gerald Schermann, Fiorella Zampetti, Daniele Romano, Philipp Leitner, Andy Zaidman, Massimiliano Di Penta, and Sebastiano Panichella. 2017. A Tale of CI Build Failures: An Open Source and a Financial Organization Perspective. In *ICSME*.
- [34] Anna Vlasova, Maria Tigina, Ilya Vlasov, Anastasiya Birillo, Yaroslav Golubev, and Timofey Bryksin. 2022. Lupa: A Framework for Large Scale Analysis of the Programming Language Usage.
- [35] Yue Wang, Weishi Wang, Shafiq R. Joty, and Steven C. H. Hoi. 2021. CodeT5: Identifier-aware Unified Pre-trained Encoder-Decoder Models for Code Understanding and Generation. *ArXiv abs/2109.00859* (2021).
- [36] Cody Watson, Nathan Cooper, David Nader-Palacio, Kevin Moran, and Denys Poshyvanyk. 2020. A Systematic Literature Review on the Use of Deep Learning in Software Engineering Research. *ArXiv abs/2009.06520* (2020).
- [37] Steven Euijong Whang, Yuji Roh, Hwanjun Song, and Jae-Gil Lee. 2021. Data Collection and Quality Challenges in Deep Learning: A Data-Centric AI Perspective.
- [38] Wenkai Xie, Xin Peng, Mingwei Liu, Christoph Treude, Zhenchang Xing, Xiaoxin Zhang, and Wenyun Zhao. 2020. API method recommendation via explicit matching of functionality verb phrases. *Proceedings of the 28th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering* (2020).
- [39] Fabian Yamaguchi, Nico Golde, Dan Arp, and Konrad Rieck. 2014. Modeling and Discovering Vulnerabilities with Code Property Graphs. *2014 IEEE Symposium on Security and Privacy* (2014), 590–604.
- [40] Pengcheng Yin and Graham Neubig. 2018. TRANX: A Transition-based Neural Abstract Syntax Parser for Semantic Parsing and Code Generation. In *EMNLP*.
- [41] Xunhui Zhang, Tao Wang, Yue Yu, Yanzhi Zhang, Yan Zhong, and Huaimin Wang. 2022. The Development and Prospect of Code Clone. *ArXiv abs/2202.08497* (2022).
- [42] Ruiqi Zhong, Mitchell Stern, and Dan Klein. 2020. Semantic Scaffolds for Pseudocode-to-Code Generation. *ArXiv abs/2005.05927* (2020).