



Contextualized Data-Wrangling Code Generation in Computational Notebooks

Junjie Huang¹, Daya Guo², Chenglong Wang³, Jiazhen Gu¹,
Shuai Lu⁴, Jeevana Priya Inala³, Cong Yan³, Jianfeng Gao³, Nan Duan⁴, Michael R. Lyu¹

¹The Chinese University of Hong Kong, China. Email: {jjhuang23, jiazhengu, lyu}@cse.cuhk.edu.hk

²Sun-yat Sen University. Email: guody5@mail2.sysu.edu.cn

³Microsoft Research. ⁴Microsoft Research Asia. Email: {chenglong.wang, shuailu, jinala, jfgao}@microsoft.com

ABSTRACT

Data wrangling, the process of preparing raw data for further analysis in computational notebooks, is a crucial yet time-consuming step in data science. Code generation has the potential to automate the data wrangling process to reduce analysts' overhead by translating user intents into executable code. Precisely generating data wrangling code necessitates a comprehensive consideration of the rich context present in notebooks, including textual context, code context and data context. However, notebooks often interleave multiple non-linear analysis tasks into linear sequence of code blocks, where the contextual dependencies are not clearly reflected. Directly training models with source code blocks fails to fully exploit the contexts for accurate wrangling code generation.

To bridge the gap, we aim to construct a high quality datasets with clear and rich contexts to help training models for data wrangling code generation tasks. In this work, we first propose an automated approach, CoCoMine to mine data-wrangling code generation examples with clear multi-modal contextual dependency. It first adopts data flow analysis to identify the code blocks containing data wrangling codes. Then, CoCoMine extracts the contextualized data-wrangling code examples through tracing and replaying notebooks. With CoCoMine, we construct CoCoNote, a dataset containing 58,221 examples for *Contextualized Data-wrangling Code generation in Notebooks*. To demonstrate the effectiveness of our dataset, we finetune a range of pretrained code models and prompt various large language models on our task. Furthermore, we also propose DATA-CODER, which encodes data context and code&textual contexts separately to enhance code generation. Experiment results demonstrate the significance of incorporating data context in data-wrangling code generation and the effectiveness of our model. We release code and data at <https://github.com/Jun-jie-Huang/CoCoNote>.

CCS CONCEPTS

• **Software and its engineering** → **Automatic programming**.

KEYWORDS

code generation, data wrangling, computational notebooks, large language models



This work is licensed under a Creative Commons Attribution-NonCommercial International 4.0 License.

ASE '24, October 27-November 1, 2024, Sacramento, CA, USA

© 2024 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-1248-7/24/10

<https://doi.org/10.1145/3691620.3695503>

ACM Reference Format:

Junjie Huang¹, Daya Guo², Chenglong Wang³, Jiazhen Gu¹, Shuai Lu⁴, Jeevana Priya Inala³, Cong Yan³, Jianfeng Gao³, Nan Duan⁴, Michael R. Lyu¹. 2024. Contextualized Data-Wrangling Code Generation in Computational Notebooks. In *39th IEEE/ACM International Conference on Automated Software Engineering (ASE '24)*, October 27-November 1, 2024, Sacramento, CA, USA. ACM, New York, NY, USA, 13 pages. <https://doi.org/10.1145/3691620.3695503>

1 INTRODUCTION

Data wrangling involves cleaning, structuring, and enriching raw data into a desired format for further analysis [96], such as by removing duplicates, casting types, and extracting features [17]. This procedure is notably time-consuming, with data scientists reportedly spending up to 80% of their time on repetitive data wrangling scripts [17, 19, 26]. Consequently, automating the generation of data wrangling codes is crucial for improving the productivity of data science practitioners [8, 96].

Data wrangling code is typically written in computational notebooks, such as Jupyter Notebooks [46] and Google Colab [9], rather than in traditional editors or integrated development environments (IDEs). These notebooks allow users to write code and annotations in *cells*, enabling interactive execution of code, inspection of intermediate outputs, and planning of further data analysis [76, 86]. As a result, the notebook environment provides rich, multi-modal contextual information, including code, textual elements (e.g., code comments and markdown text), and data [93], as illustrated in Figure 1.

Precise data-wrangling code generation necessitates fully exploiting all contextual information (i.e., code, textual and data). Existing solutions often fail to integrate these contexts effectively, leading to suboptimal performance. Specifically, several studies have focused on generating code based on textual and code context (i.e., ①→④ in Figure 1) [1, 12, 34, 98]. However, the inherent ambiguity of natural language [52] and its potential absence in notebooks [63] (e.g., the textual context fails to provide explicit instructions for the target code in Figure 1) can make the specification unclear. Another approach adopted the programming-by-example (PBE) paradigm [26], where data context (i.e., input and output examples) serves as specifications (i.e., ②+③→④) [5, 16, 69]. However, the absence of codes and annotations can lead to a lack of preceding programming details, resulting in erroneous generated codes, such as misused APIs or variables. Additionally, some studies [38, 45] used textual and code context alongside input examples (i.e., ①+②→④). Nevertheless, without the output data, crucial information may still be missing. For instance, as shown in Figure 1, the data field `difference_size`

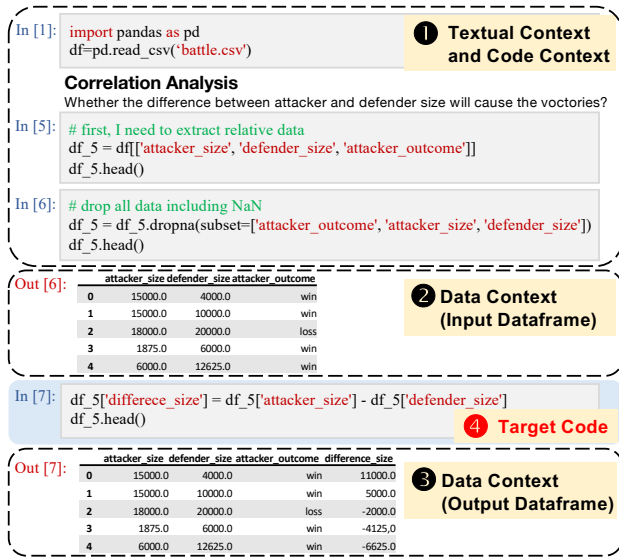


Figure 1: An example of contextualized data-wrangling code generation in CoCoNote, aiming to generate target code(4) given multi-modal context(1,2,3).

never appears in either the textual context or the input example, leaving key details unspecified. Therefore, integrating code, textual, and data contexts is essential to fully elucidate specifications for data-wrangling code generation tasks.

Ideally, we aim to train a model capable of comprehending all three types of contexts to generate data wrangling code. However, the primary challenge in building such a model is the scarcity of high-quality datasets. Notebooks often interleave multiple non-linear analysis tasks into a linear sequence of code blocks, obscuring contextual dependencies. Consequently, models trained solely on source code fail to fully utilize these contexts for generating wrangling code. This issue is further compounded by the fact that many online notebooks lack complete context information, such as missing execution histories that are crucial for providing essential data contexts.

This work aims to construct a high quality dataset with clear and rich contexts to facilitate training models for data wrangling code generation tasks. We propose an automated approach, *i.e.*, CoCoMine (Contextualized data-wrangling Code Mining), to extract data-wrangling code examples with multi-modal contexts. CoCoMine consists of two major steps: data-flow-based data-wrangling code identification and multi-modal contexts alignment. First, CoCoMine extracts relevant code cells for data analysis (including data initialization, wrangling, and utilization [10]) from the tangled notebooks. To identify data-wrangling code cells, we create API databases to signify the data initialization and utilization stages, where the dataframe variable is initially created and finally utilized, such as visualization and statistic computation. CoCoMine then conducts data-flow analysis on the dataframe variable to track its intermediate transformation process, treating all related code cells as data-wrangling cells. With the identified code-wrangling cells, CoCoMine creates contextualized code generation examples. In particular, it first collects the code and textual context through

dependency analysis. Then, CoCoMine gathers the intermediate execution results of the dataframe variable as data context by crawling imported data files and replaying the notebook in a sandbox environment. Finally, the target code, along with its aligned code, textual, and data contexts, form a comprehensive code generation example.

By employing CoCoMine, we collected a dataset comprising 58,221 examples (54,567 for training, 2,000 for development, and 1,654 for testing) for *Contextualized Code generation in Notebooks*, named CoCoNote. To ensure the natural and reliable code evaluation, CoCoNote can be evaluated not only with surface-form metrics such as Exact Match (EM) and CodeBLEU [74], but also with an execution-based metric named Execution Accuracy (EA). The EA evaluates the generated code by comparing the resulting output dataframes against the oracle output, similar to the approaches in [34, 98].

With CoCoNote, we conduct an extensive study on the impact of multi-modal context and the performance of various language models on data-wrangling code generation. Specifically, we used CoCoNote to fine-tune five code-pretrained language models (PLMs) and prompt six large language models (LLMs) to comprehensively evaluate their performance. During model fine-tuning stage, we observed that data context, which includes structured dataframes, should not be encoded alongside unstructured code and textual contexts due to their differing nature. Additionally, the input length limitation of PLMs make it impractical to encode all three contexts simultaneously. To fully leverage data contexts, we extend CodeT5 [90] and proposed a novel encoder-decoder model called DATACODER. DATACODER employs dual encoders to separately encode the code and textual contexts and input-output dataframes, resulting in more accurate latent representations for decoding.

Our experimental results demonstrate that (1) multi-modal contexts, particularly the data context of input-output examples, are crucial for data-wrangling code generation, as each modality provides essential information for model inference. (2) Column names in data context have a greater impact than data values in offering explicit signals for data manipulation, while output examples are more effective than input examples in clarifying the intended results by specifying the expected output. (3) With increased parameters and enhanced training methods, GPT-4 achieves the highest performance with an execution accuracy of 50.6%, indicating substantial room for further improvement on this task. We hope this study will benefit future research in developing robust multi-modal code generation methods.

We summarize the contributions as follows:

- We introduce the task of contextualized data-wrangling code generation, incorporating code context, textual context and data context.
- We develop CoCoMine, a tool to mine comprehensive context for data-wrangling codes, resulting in a dataset of 58,221 examples.
- We propose a data-enhanced model, DATACODER, featuring separate encoders for data and textual/code, which better utilizes contextual information to improve code generation.
- We provide a comprehensive analysis of how contextual information affects generation performance across a range of PLMs and LLMs.

2 TASK AND BACKGROUND

This paper aims to explore the role of various context in data-wrangling code generation under the notebook environment. We narrow down our focus to data wrangling, the most time-consuming and tedious process in data analysis [26, 86, 96], where the data scientist’s target is to manipulate data to the right shape in preparation for subsequent analysis including plotting, modeling or statistics computation [8].

2.1 Motivation

Our work is motivated by a common programming scenario in which data scientists program in computational notebooks to analyze the target data. As they program, DS practitioners often consult various sources of information within notebooks to help them write correct analysis code [93], *e.g.*, code, markdown text, tabular dataframe, images, and external file system [93]. Code snippets, written in preceding code cells, provide essential code dependencies such as imported modules, methods, and variable names [1]. Text in markdown cells or inline comments can provide valuable instructions to the task background and intended output [65]. Apart from code and text, tabular dataframe is another source that data scientists might refer to. By looking up dataframes, they can understand column names and value types, assisting data wrangling process [15]. Hence, comprehensive notebook contexts in modalities are important to write correct data-wrangling code.

2.2 Task Formulation

To simulate the DS programming practice and investigate the effects of multimodal programmatic context, we formulate a novel task of *contextualized data-wrangling code generation* in notebooks. As shown in Figure 1, given a mix of code and textual context and data context (*i.e.*, input-output dataframes), the task aims to generate target code that can transform the input dataframe to the output dataframe. In this task, the code and context contexts come from the preceding notebook cells that are relevant to target code. The input dataframe is the runtime data to manipulate, which is usually displayed for perceiving, while the output dataframe is a sample of execution result, which specify the expected output. We study *dataframe*, a standard data structure to store tabular data in Pandas library, due to its popularity in data analytics and powerful data manipulation capabilities [61].

Unlike traditional code generation in notebooks [12, 34, 98] that includes a textual intent to specify target code function, we provide input-output data as the instruction to constrain the task. This novel setting resembles a programming-by-example (PBE) [70] fashion and is motivated by three factors. Firstly, NL isn’t always the optimal specification type, as different data transformation tasks may require different specification types [84]. For example, “pivot” tasks are hard to describe precisely in NL, whereas “sort” tasks are more suitable for NL [84]. Secondly, textual intents have inherent ambiguity in expressing these tasks due to the gap between requirements and program logic [99]. Whereas concrete input-output examples can elucidate specifications and have been observed and verified as a useful means to provide task intents [26]. Thirdly, writing complete textual specification is hard and time-consuming for programmers even with extensive expertise [?]. In contrast, input-output data

supplement data wrangling by specifying expected outcomes and thus enable non-programmers to solve this task. Therefore, our task might potentially revolutionize automated data wrangling since users can write the requirements using examples, as discussed in previous study [27, 62].

Our work also distinguishes from existing data-wrangling code synthesis approaches that rely solely on input-output examples [8, 96] or incorporate additional textual intent [38, 45] by providing more natural and comprehensive programmatic contexts. These contexts provide essential references, such as variables to process and code patterns [15]. Our experiment results (§5.2) also show the importance of these programmatic context in generating executable and correct code. Generally, our task brings PBE to a more practical scenario [16] in computational notebooks by augmenting programmatic context.

2.3 Evaluation:

To assess the correctness of generated code, we leverage both surface-form metrics and execution-based metrics [14] for rigorous and robust evaluation in our task. For surface-form metrics, we use two prevailing metrics, *i.e.*, **Exact Match (EM)** accuracy and **CodeBLEU (CB)** [74], to measure the surface correctness. The EM accuracy is the most strict metric that measures whether the code generated by the model is identical to the ground-truth code. CB takes both the code structure and semantics into account to measure semantic and syntactic similarity of two code snippets. It is computed as the average of four components including n-gram matching score, weighted n-gram matching score, syntactic AST matching score, and semantic data flow matching score. For execution-based metrics, we create a sandbox environment to execute generated code and use **Execution Accuracy (EA)**, the percentage of examples that generated and ground-truth code produces the same output dataframe, to measure the execution correctness. The sandbox environment setup is elaborated in Section 3.3.1.

3 COCONOTE CURATION WITH COCOMINE

Mining data-wrangling code generation examples with sufficient context is non-trivial due to challenges of identifying data-wrangling scripts, collecting runtime data contexts, and distinguishing code and textual dependencies, as discussed in Section 1. To address the challenges, we propose CoCoMine to automatically collect data-wrangling code generation examples from public notebooks, resulting in a benchmark dataset called CoCoNote. The principle underlying CoCoMine involves initially identifying data-wrangling code from tangled notebook cells, followed by the collection of contextual dependencies to construct examples.

Figure 2 shows the overall workflow of CoCoMine. We first crawl notebooks relevant to data analysis and collect execution-needed data files used in notebooks from GitHub (§3.1). Then we parse the code in notebooks and conduct data-flow analysis to identify code cells related data-wrangling. These code cells track the lifecycle of a dataframe variable from initialization, wrangling, to final utilization in a data analysis task (§3.2). Specifically, we create API databases to signify the data initialization (§3.2.1) and utilization stages (§3.2.3), and then conduct data-flow analysis on the dataframe variable to track its intermediate wrangling process (§3.2.2). With identified

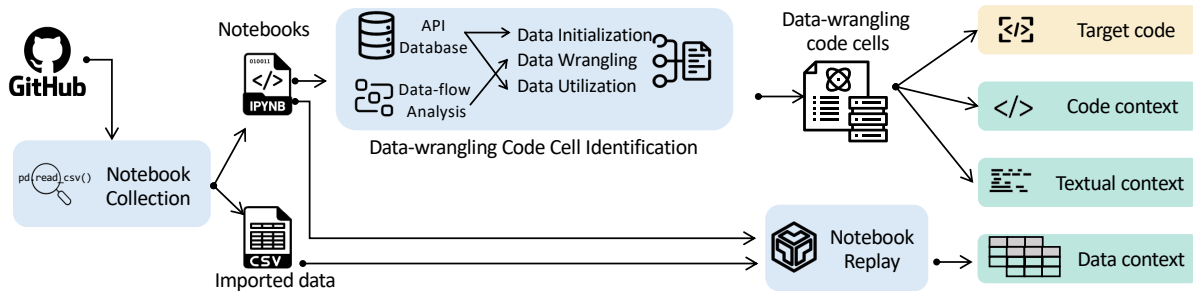


Figure 2: The workflow of CoCoMine.

data-wrangling code cells, we proceed to build code generation examples with multi-modal context (§3.3). Specifically, we first sample potential target code from code cells in the data wrangling stage (§3.3.1). Next, we trace back the data-flow diagrams to collect the dependent code and textual context from cells preceding the target code (§??). Subsequently, we create a sandbox environment to replay notebooks and record the dataframe during execution to obtain runtime data context (§3.3.2). Finally, we perform human verification to ensure the quality of the dataset (§3.4).

3.1 Collecting Notebooks and Data Files

CoCoMine collects notebooks from GitHub using public GitHub API [25]. The notebooks are required to have the `.ipynb` file extension and a redistributable license, following [96]. We remove duplicated notebooks and those from forked repositories. We choose GitHub as our source, as opposed to Kaggle [40], a platform for data scientists to display their DS solutions through notebooks. This decision is based on the fact that GitHub offers a larger and more diverse collection of notebooks [63]. Subsequently, we apply a straightforward heuristic to collect notebooks potentially relevant to DS. The notebooks are required to run with Python 3 and invoke a `pd.read_csv()` function from the Pandas library, which is one of the most popular data analysis libraries in Python [61]. The utilization of `pd.read_csv()` function indicates the loading of real data from external files, as opposed to synthetic data used only for experimental purposes [44], which suggests a higher likelihood of the notebooks being employed for real-world data analysis [10].

Then CoCoMine collects data files loaded in notebooks to support execution. Specifically, we parse the code merged from code cells into an Abstract Syntax Tree (AST). We then identify the string nodes and assess whether the string represents a valid file system path. After that we trace the paths in original GitHub repositories and download the corresponding data files. Notebooks that rely on inaccessible data sources are excluded. Finally, we simplify the nested file paths in the notebooks and consolidate all files into a single directory in preparation for execution.

3.2 Data-Wrangling Code Cell Identification

In this step, CoCoMine extract data-wrangling code cells from tangled notebooks. These code cells compose a self-contained data analysis process and track the lifecycle of a dataframe variable. We introduce three typical stages in this process: a data initialization

stage to initialize a target dataframe, a data wrangling stage to convert the dataframe into a desired format, and a data utilization stage to utilize processed data for an analysis objective such as visualization and statistic computation. To identify the three stages of data-wrangling code cells, we create API databases and conduct data-flow analysis [10, 80, 88]. The rationale of analysing API and data-flow stems from the fact that a majority of data-wrangling codes are implemented with common Python DS libraries such as *pandas* and *matplotlib* [10]. These libraries offer high-level and stable APIs to perform specific tasks on data, which can be used to infer code functionality and subsequently deduce its stage.

Specifically, we parse the programs into an AST and extract the temporal sequence of API calls through standard static analysis. We then conduct API matching for each statement with predefined API databases to determine the initialization of a dataframe and its utilization. After that we conduct data-flow analysis of the target dataframe to identify its transformation stage. Code cells that undergo these states are extracted as the valid data-wrangling code cells. We keep the top-to-bottom order of code cells as in the original notebook to keep the sequential logic. To form a comprehensive API database, we empirically investigate API documentation, statistically count API frequencies, and also refer to previous studies [10, 96], to collect the representative APIs. To ensure self-containment of DS tasks, we discard notebooks with no complex code cells extracted. The methods to identify each stage are described in the following section.

3.2.1 Data Initialization. A data analysis task starts with the *data initialization* stage, whereby a dataframe variable is initialized. This variable serves as the target variable that we aim to scrutinize throughout the entire process. Specifically, we first identify statements for data initialization with API matching and then select assigned variable in the statement as the target variable. We select four types of statements: (1) Directly creating a new dataframe with the API, such as `df=pd.DataFrame()`. (2) Loading a dataframe from external data files with the API, such as `df=pd.read_csv()`. (3) Manipulating an existing dataframe with the API to return a new dataframe, such as `df=pd.merge([df1, df2])`. (4) Assigning a new dataframe via dataframe operations, such as slicing and addition. Table 1 shows the full API databases for data initialization. It is noteworthy that APIs from other libraries can also return a dataframe (e.g., `sns.load_dataset("iris")` from *seaborn*). While our focus is on presenting a general identification framework, we have

Table 1: APIs and Operations to Identify Dataframe Variables.

Category	API calls & Operations
Definition	pd.DataFrame(), pd.Series()
Data Loading	pd.read_csv(), pd.read_pickle(), pd.read_excel() pd.read_json(), pd.read_table(), pd.read_hdf() pd.read_sql(), pd.read_html()
Dataframe Manipulation	pd.get_dummies(), pd.merge(), pd.concat() pd.stack(), pd.unstack(), pd.cut(), pd.qcut() pd.melt(), pd.corr(), pd.crosstab(), pd.pivot_table()
Operations	+, -, *, /, df[.]

collected APIs from the widely-used Pandas library. Nevertheless, our approach can be readily extended by incorporating more APIs.

3.2.2 Data Wrangling. After initialization, the data need certain processing, which we categorize as the *data wrangling* stage. This stage involves exploration, cleaning, and feature engineering, which aids in preparing the appropriate data for subsequent processing. Well-prepared data reduce the efforts for data analysis and contributes to the successful resolution of a data analysis task. Specifically, this stage can involve multiple statements to update the target variable. To identify wrangling codes, we conduct data-flow analysis to track the target variable defined in data initialization. A statement is designated as the data wrangling code if the target variable is reassigned to itself (e.g., `df=df.dropna()`), or an in-place modification is invoked on itself (e.g., `df.dropna(inplace=True)`). Here, we track only one variable in the cells so that we can check whether data has been tampered with. We do not consider the assignment to a new variable as the semantics can be changed when such an assignment is made.

3.2.3 Data Utilization. The final data utilization stage is to complete a specific objective in DS, where analysts aim to reach a conclusion after rounds of programming and analysis. This stage is often achieved through the creation of visualizations or calculation of statistical measures with the well-transformed dataframes to facilitate future decision-making. Based on previous works [10, 44] and our observations, these intentions can often be achieved through the use of well-developed API calls from popular DS libraries. Therefore, we can designate the statement as the DS objective through examining the API calls invoked on the target variable. We select four most frequently used libraries for visualization and statistics to create API databases, i.e., *matplotlib*, *seaborn*, *scipy*, and *scikit-learn*. Table 2 shows the full API databases for data utilization. Our findings of final objectives also align with the Biswas et al. that DS problems in Kaggle notebooks are mostly finished with an evaluation or visualization stage [10].

3.3 Multi-modal Context Alignment

Next, we create examples for contextualized code generation based on the extracted data-wrangling code cells to construct CoCoNote dataset. As shown in Figure 1, given the multi-modal context of code, text, and input-output dataframe, the task is to generate the target code that can transform the input dataframe to the output dataframe. In the following, we first introduce the formulation of target code. Then we elaborate data context, code context, and

Table 2: API Calls to Identify Data Utilization Statements.

Library	API calls
matplotlib	show(), plot(), barh(), hist(), hist2d(), imshow() bar(), pie(), scatter(), contour(), pcolormesh() errorbar(), matshow(), semilogy(), probplot()
	distplot(), heatmap(), countplot(), barplot() boxplot(), jointplot(), clustermap(), distplot() pairplot(), scatterplot(), violinplot(), kdeplot() swarmplot(), regplot(), rugplot(), stripplot() lmplot(), catplot(), relplot(), lvplot(), tsplit() lineplot(), boxenplot(), factorplot(), palplot() pointplot(), corrplot(), residplot(), FacetGrid()
scipy	stats.plot(), stats.hist(), stats.scatter() mean_squared_error(), mean_absolute_error() roc_auc_score(), cross_val_score(), r2_score() confusion_matrix(), recall_score(), f1_score()
scikit-learn	pearsonr(), accuracy_score(), precision_score() classification_report(), cosine_similarity() softmax_cross_entropy_with_logits()

textual context, and finally show the final example and dataset statistics.

3.3.1 Target Code. The target code is formulated as the entire code in code cells rather than several statements, as a code cell is a basic functional unit for programming and execution in notebooks. We use code cells between data initialization and data utilization stages to assure the functionality of data wrangling. We select target code cells based on the observation that DS programmers tend to perform data inspection operations during programming, such as displaying a dataframe with `df.head()` in the last line of a code cell. The inspection helps programmers to understand table details before programming or evaluate the code functionality once the programming is finished [44, 76], which can be regarded as a signal for the beginning and functional completeness of a code snippets [13]. Hence, we select the code between any two consecutive data inspection statements in the data-wrangling code cells as our target code. Specifically, we use the built-in inspection APIs to identify the data inspection statements, including the dataframe display expression (i.e., `df`), and partial dataframe display expressions (i.e., `df.head()` and `df.tail()`). It's worth noting that the target code can span multiple code cells; thus, we merge them to obtain the final target code. Figure 3 shows the distribution of APIs and modules used in the target code.

3.3.2 Data Context. Upon identifying the target code, we linearly re-execute the notebook to obtain runtime data context, i.e., the input and output dataframes. Specifically, we record the value of the target variable before target code as the input dataframe, and then execute target code to obtain the output dataframe. In the installed sandbox, we sequentially execute code cells from top to bottom for simplicity without considering their original arrangement in notebooks. Notebooks that encounter exceptions or exceed execution duration of 300 seconds per cell are considered unsuitable, resulting in the subsequent exclusion of examples derived from them.

Execution Environment The challenge of executing real-world computational notebooks is the potential inclusion of diverse libraries and dependencies from other .py files. However, creating a

customized environment for each notebook individually is impractical for large-scale execution due to the intricate dependencies and the resulting increase in evaluation time [87, 88]. To address this, we adopt a simplified approach by creating a unified sandbox environment for all notebooks in our dataset. This sandbox is equipped with Python standard libraries, specifically version 3.6.13. Additionally, we install the 15 most frequently used non-standard data analysis libraries in notebooks, namely: *pandas*, *numpy*, *matplotlib*, *sklearn*, *seaborn*, *scipy*, *nlTK*, *plotly*, *statsmodels*, *geopandas*, *bokeh*, *ggplot*, *xgboost*, *lightgbm*, and *patsy* based on our statistical analysis. Despite the popularity of machine learning libraries, such as *TensorFlow* and *PyTorch*, we do not install them due to their extensive usage in DL tasks, which require extensive time and computation resources, making them less suitable for execution evaluation

3.3.3 Code Context. In code generation, the code contexts prior to the target code are essential as they can provide clues for variable reusing and already performed computations [37]. However, mining relevant code context is difficult, especially in the notebook environment, where cells at the very beginning of notebooks can be referred, such as import statements and imported dataframes [15]. To address the issue, we consider two types of code context. On one hand, we incorporate code cells in the data-wrangling code cells prior to the cell of target code, which specify the performed computations on the target variable. On the other hand, we incorporate the dependent statements of (1) the latest definition of variables used in the target code, and (2) the import statements of DS libraries, which can be used to infer variable and API usages. We identify the dependencies by matching module names and variable references with AST, which are then merged into a new code cell and placed at the beginning before other code cells.

3.3.4 Textual Context. In DS notebooks, natural language documentations are often well-written by data scientists to enable collaboration and sharing usage [18, 63, 83], which are useful to provide task background and code instructions. Therefore, we include textual context to our task, which includes: (1) markdown cells before the target code cell and before each code cell in the data-wrangling code cells, (2) comments in the cells of code context, and (3) inline comments in the target code cells. We include the inline comments because they can specify (partial) functions of the target code, which aids in good-quality generation, as suggested by previous work [63]. One type of dependency is the natural language explanations in *markdown cells*, which we can incorporate by inserting the original non-blank markdown cells before each code cell.

3.3.5 Finalizing Context. Finally, to preserve sequential information, we merge code and textual contexts as a unified part and keep their original top-to-bottom order as in the notebook. To reduce task complexity, we discard examples with more than ten context code cells. We also remove those with identical input and output dataframes as this may indicate a vague functionality of target codes. For uniqueness, we remove duplicated examples and examples with the same target code. However, it's often the case that even in different notebooks, there are the same code snippets in the code cells, due to the convenience of modularity and reuse. To

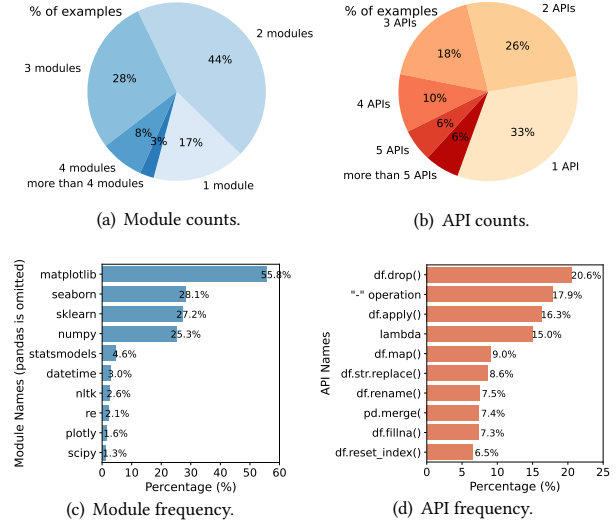


Figure 3: Distribution of modules and APIs in the target code.

avoid answer leakage, we remove testset examples whose target code appears in the context of training and validation examples.

The final CoCoNote contains 54,567 and 2,000 examples for training and validation and 1,654 for testing. To ensure rigorous evaluation of correct output dataframes, we leverage the full data with numerous rows rather than the snapshots to producing output dataframes during testset execution. However, as using full dataframes in training is computationally expensive for modeling, we do not incorporate full dataframes and simply keep a maximum of ten rows in the training and validation set.

3.3.6 Dataset Statistics. Table 3 shows the statistics of CoCoNote. We find that test examples contain a lower average number of textual context tokens (138.0) and target code tokens (41.3) compared to the training and validation examples. Figure 3 shows the distribution of modules and APIs used in target code of CoCoNote. We find that most examples use less than 3 modules (89%) and 3 APIs (77%). Furthermore, we provide the frequency of API and module usage to elucidate task distribution, encompassing a broad spectrum of widely-used DS modules and data wrangling APIs.

3.4 Human Verification

Here we conduct a human evaluation to verify the solvability of our task. We enlist the participation of three university students and two industry data scientists, all of whom possess extensive experience with Python and notebook environments. To ensure a rigorous and credible validation process, we formulate a comprehensive guidebook and conduct a pilot study on several examples following the guide. The guidebook includes our task objectives, evaluation criteria, a step-by-step annotation guide, and five representative examples. Next we randomly sample 50 examples from the testset as the final verification data for participants. Our annotation task is to judge if the provided information is sufficient for the participants to write target code, given varying combinations of contextual

Table 3: Statistics of CoCoNote Dataset.

	train	dev.	test
# examples	54,567	2,000	1,654
avg # columns (input df)	12.8	12.1	12.9
avg # rows (input df)	8.6	9.0	249.5
avg # columns (output df)	13.3	12.8	12.7
avg # rows (output df)	8.6	9.0	247.2
avg # textual context tokens	276.2	283.9	138.0
avg # target code tokens	67.2	67.1	41.3

information, *i.e.*, code, text, input dataframe ($data_{in}$), and output dataframe ($data_{out}$). During annotation, they are allowed to discuss and search the Internet to ensure annotation accuracy and quality. Finally, the class with the majority votes is assigned to each example. We find that 34% examples can be solved with only the context of code+text, while 40% and 42% examples can be resolved with the code+text+ $data_{in}$ or code+text+ $data_{out}$, respectively. Furthermore, 76% examples can be solved when all elements (code+text+ $data_{in}$ + $data_{out}$) are accessible. The results shows that the inclusion of both input and output dataframes significantly enhances the solvability rate. This is because participants can compare differences between input and output dataframes, thereby inferring the functionality of the target code. These findings underscore the importance of integrating both input and output dataframes into the context of code generation. However, it is noteworthy that 24% examples were challenging for the participants to evaluate, thereby highlighting the difficulty of our task.

4 CODE GENERATION METHODS

In this section, we introduce models to perform contextualized data science code generation with CoCoNote. We comprehensively evaluate a diverse set of language models for code with varying sizes through finetuning (§ 4.1) and prompting (§ 4.3). To better incorporate data context, we also introduce DATACODER with bi-encoders to encode runtime data context and textual context based on CodeT5 (§ 4.2). The best model GPT-4 achieves 50.6% accuracy, indicating a large room for improvement.

4.1 Finetuning Code Pretrained Models

Fine-tuning a pretrained language model (PLM) in downstream tasks is an effective paradigm for code intelligence [60]. By training on example pairs, PLMs can transfer their knowledge to perform domain-specific code generation without learning from scratch [85].

4.1.1 PLMs. We evaluate five representative code PLMs that have shown state-of-the-art performance in code generation.

CodeBERT [22] is an encoder-only architecture pre-trained on text-code pairs with masked language modeling and replaced token detection objectives. As the decoder is not pretrained, we apply a randomly initialized transformer decoder with 6 layers, 768 dimensions, and 12 heads, forming an encoder-decoder architecture to generate. The decoder is optimized during fine-tuning.

GraphCodeBERT [29] enhances code representation by utilizing data flow graphs and uses an encoder-only architecture built

on top of CodeBERT. Similar to CodeBERT, we train a decoder for generation using the same method.

UniXcoder [28] is adapted from UniLM [20] and is pre-trained on unified cross-modal data including code, text, and AST. As UniXcoder is also an encoder-only architecture, we follow the method used in CodeBERT and train a decoder for generation.

PLBART [2] is an encoder-decoder model built upon BART [51] and is pre-trained on a mix of Python functions, Java functions, and textual posts from StackOverflow [68] with a denoising objective.

CodeT5 [90] is an encoder-decoder model built upon T5 [73] and incorporates crucial token type information from identifiers and enables multi-task learning. The decoder is a multi-layer transformer with 12 layers, 768 dimensions, and 12 heads.

4.1.2 Finetuning Setting. We prepare the model input by concatenating the code and textual context and input-output dataframes. We set the maximum sequence length of both textual context and data context to 256, resulting in a maximum concatenation length of 512. To ensure code syntax correctness, text from markdown cells is wrapped into multi-line comments by triple quotes. For data context, we concatenate the flattened string of input and output dataframes, where each table is converted by arranging each row’s elements from left to right, then stacking them from top to bottom. Due to the input sequence length limit, we only use the first five rows of the input/output dataframes. During fine-tuning, all models are trained on 8 Tesla V100 32GB GPUs with a cross-entropy loss function. We set the batch size as 64, the epoch as 20, and the maximum target sequence length as 300. The remaining hyper-parameters are kept by their default settings of each PLM.

4.2 DATACODER

The above PLMs encode the joint textual and data context, which may not optimal for this task due to two reasons. First, there is a discrepancy between textual and tabular data [35], which makes it challenging to understand the context and consider the cross-modality connections. Second, the input length limit of PLMs (*e.g.*, 512 tokens) often leads to incomplete input context, which can further hurt generation performance. To address the issues, we propose DATACODER that leverages dual encoders to separately encode textual context and input-output dataframes, yielding more accurate latent representations for decoding.

4.2.1 Architecture. The architecture of DATACODER consists of two encoders and one decoder, as shown in Figure 4. One encoder is used to encode code+text, and the other encodes tabular input/output dataframes. For each encoder, we apply N transformer layers [82] over the input sequences to produce text hidden states and table hidden states for code+text and dataframes, respectively. The input sequences are organized following the same setting described in § 4.1.2. After encoding the input sequences, the decoder takes the concatenation of the last-layer hidden states from both encoders as input to generate target code. The decoder also applies N transformer layers, but with a different attention mechanism, where tokens in each transformer layer can only attend to themselves and preceding tokens, allowing for sequential generation of the target code.

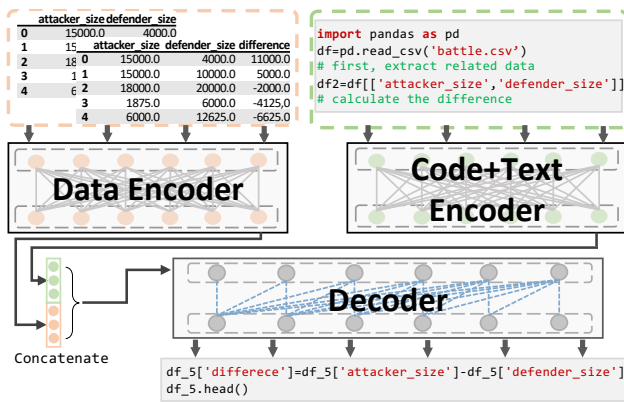


Figure 4: Model Architecture of DATACODER.

4.2.2 Training Setting. To leverage pretrained knowledge and reduce training time, we initialize encoders and decoders with 12 layers, 768 dimensional hidden states, and 12 attention heads using the pre-trained parameters from CodeT5 [90]. We use a learning rate of $2e-4$, a batch size of 64, a maximum sequence length of 512 for both code+text and data, and a target sequence length of 300, respectively. We use AdamW optimizer [59] to fine-tune models for 20 epochs and perform early stopping on the development set.

4.3 Prompting Large Language Models

Large language models (LLMs) have shown remarkable performance in various programming tasks such as code completion and generation [58]. Due to the high cost of tuning an LLM, we evaluate them through prompting, where a prompt with multi-modal context is constructed and then fed into LLMs to infer code [11].

4.3.1 LLMs. We select four open-source code LLMs and three closed-source commercial LLMs to perform our task:

StarCoder-15B [52] is a decoder-only model with 15b parameters built upon SantaCoder [3], employing Flash attention to scale up the context length to 8k. It is trained on The Stack dataset [47] including 86 programming languages, GitHub issues, Git commits, and Jupyter notebooks.

CodeLlama-34B (CodeLlama-34b-Instruct) [75] is an autoregressive transformer model built upon Llama-2 [81], using a multi-task objective of autoregressive and infilling prediction [6], and long text fine-tuning to enable inference context length up to 16k.

Phind-CodeLlama-34B (Phind-CodeLlama-34b-v2) [71] is an improved version of CodeLlama that are additionally trained on 1.5B tokens and 80k high-quality programming problems.

DeepSeek-Coder-33B (deepseek-coder-33b-instruct) [30] is a transformer model trained from scratch on 2b tokens sourced from 87 programming languages. It employs a fill-in-the-blank task [6] with a 16K window to enhance code generation and infilling.

Codex (code-cushman-001) [14] is an LLM with 12b parameters that powers GitHub’s Copilot auto-completion system. However, due to the advancements of GPT models for coding tasks, OpenAI will no longer support Codex API.

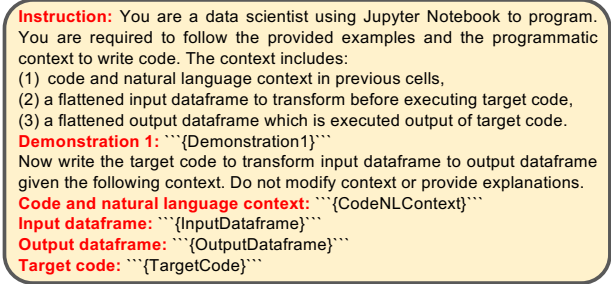


Figure 5: Prompt template to query LLMs.

GPT-3.5 and GPT-4 [67] are the recent state-of-the-art LLM developed by OpenAI. While they are not explicitly trained for code generation, they demonstrate notable performance in this domain [58]. Their effectiveness in handling code generation tasks is largely attributed to their large parameter size. For GPT-3.5 and GPT-4, we use *gpt-3.5-turbo-0613* and *gpt-4-0613*, respectively.

4.3.2 Prompting Setting. We evaluate LLMs in both few-shot and zero-shot settings. Figure 5 shows our few-shot prompt to query LLMs, which includes an instruction, demonstration, and context. To ensure code syntax correctness, the instruction, input dataframe and output dataframe are wrapped into multi-line comments by triple quotes. In few-shot setup, we randomly sample k in-context examples [11, 24, 55] and insert them before the test example. This allows LLMs to learn from the in-context examples and generate corresponding response for the test sample. By default, we set $k=2$ for all LLMs. The zero-shot prompt is similar to the few-shot prompt, but without demonstrations. We utilize the official OpenAI API [67] for Codex, GPT-3.5, and GPT-4, while the remaining open-source LLMs are based on the HuggingFace transformers library.

4.4 Parameter-Efficient Finetuning

Besides small-sized code language models and large-sized ones, we also comprehensively evaluate medium-sized LLMs (1.3-7B) on CoCoNote through parameter-efficient finetuning (PEFT).

4.4.1 LLMs. We select five medium-sized open-source code LLMs.

DeepSeek-Coder-1.3B and **DeepSeek-Coder-6.7B** (deepseek-coder-1.3b-base and deepseek-coder-6.7b-base) [30] are the DeepSeek-Coder family with varying parameter size, which are trained with the same datasets and training objectives.

CodeGen-6B (codegen-6b-mono) [65] is a transformer architecture and trained with the next-token prediction objective on three NL and code datasets sequentially: ThePile [23], BigQuery, and BigPython.

Mistral-7B-Instruct (Mistral-7B-Instruct-v0.2) [39] is an LLM tuned with instruction datasets based on Mistral-7B with grouped-query attention to accelerate inference speed.

CodeLlama-7B-Python [75] is a variant of CodeLlama LLM family built on top of Llama-2 [81], using additional 100B Python tokens to specialize Python language.

4.4.2 PEFT Setting. Following previous studies on PEFT for code generation [94], we adopt top-performing PEFT technique Low Rank Adaption (LoRA) [32] to tune the LLM on CoCoNote. LoRA

freezes the model weights and injecting low-rank trainable matrices into the attention layers of the Transformer architecture, thereby drastically reducing the number of trainable parameters. We adopt the same input-output pairs as finetuning Code PLMs (§ 4.1.2). For all LLMs, we use a decomposition rank of 8, a batch size of 3, and trained for 5 epochs on 1 Tesla A100 40GB GPU with a cross-entropy loss function. The remaining hyper-parameters are kept by their default settings of each LLM.

5 EXPERIMENT

We evaluate code generation models on CoCoNote by answering the following research questions (RQs):

- **RQ1:** How do different models perform on CoCoNote?
- **RQ2:** How do varying context affect performance?
- **RQ3:** How do various data context affect performance?

5.1 RQ1: Performance of Different Models

Setup: We first compare the performance of DATACODER with a variety of PLMs and LLMs on CoCoNote. The generated code is evaluated using three metrics: *i.e.*, Exact Match (EM), CodeBLEU (CB), and Execution Accuracy (EA), as described in § 2.3. To ensure a rigorous evaluation, both PLMs and DATACODER are finetuned on the identical training set of CoCoNote. For LLMs, in-context examples are sampled from the same training set. The results in terms of PLMs and LLMs are shown in Table 4 and Table 5, respectively.

Table 4: Performance of DATACODER and PLMs.

Model	EM	CodeBLEU	EA
GraphCodeBERT	14.0	48.9	28.9
CodeBERT	14.6	49.0	29.5
UniXcoder	15.1	50.3	30.7
PLBART	16.3	52.7	33.1
CodeT5	<u>19.1</u>	55.3	38.1
DATACODER (ours)	21.3	<u>57.2</u>	<u>42.2</u>

Results: We first analyze the performance of DATACODER in comparison with PLMs finetuned on CoCoNote. From Table 4, we find that: (1) Our DATACODER, which leverages two separate encoders for textual and data context, outperforms all other PLMs across all metrics. This result shows the effectiveness of DATACODER. (2) Compared with CodeT5, DATACODER exhibits a consistent enhancement across all metrics (*i.e.*, +4.1% EA, +2.2% EM, and +1.8% CodeBLEU). The improvement can be attributed to the use of a separate data encoder for input/output dataframes, which can incorporate a more extensive context and better represent the modalities to aid generation. (3) In the second group in table 4, due to the encoder-decoder architecture and denoising pre-training tasks being specifically designed for generation tasks, PLBART and CodeT5 perform better than encoder-only models such as CodeBERT, GraphCodeBERT, and UniXcoder.

Next, we analyze the performance of DATACODER compared with prompting LLMs. From Table 5, we can find that: (1) Despite utilizing fewer parameters (305M), DATACODER exhibits superior

Table 5: Performance of DATACODER and prompting LLMs.

Model	EM	CodeBLEU	EA
StarCoder-15B	8.1	53.2	33.2
CodeLlama-34B	1.0	36.9	4.9
Phind-CodeLlama-34B	3.0	40.7	21.5
DeepSeek-Coder-33B	7.1	52.0	37.3
Codex	2.3	31.5	11.9
GPT 3.5	6.9	55.3	40.1
GPT 4	11.1	60.8	51.6
DATACODER (ours)	21.3	<u>57.2</u>	<u>42.2</u>

performance than five medium-sized code LLMs (12B-34B) and one large-sized LLM (*i.e.*, GPT 3.5 with 175B). This can be attributed to our task requiring not only the comprehension of code and textual context but also the understanding of the data within input/output dataframes, which poses greater challenges in a few-shot setting. Fine-tuning on the training set proves beneficial in learning the ability to understand data context and consequently leads to improved performance. (2) Compared with GPT 4, although DATACODER is inferior in terms of CodeBLEU and EA, it achieves a +10.2% EM of the generated code. In addition, the finetuned PLMs consistently outperform in terms of EM. We attribute this phenomenon to the finetuning strategy applied to DATACODER. As coding in notebooks for data analysis differs from general domain programming, current LLMs may not fully utilize domain knowledge and comprehend programming patterns without tailored finetuning, leading to more diverse code and lower EM score.

Table 6: Performance of DATACODER and LLMs with LoRA.

Model	EM	CodeBLEU	EA
Deepseek-Coder-1.3B	17.2	59.5	45.2
CodeGen-6B	17.8	60.2	43.9
Mistral-7B-Instruct	16.8	59.8	46.1
CodeLlama-7b-Python	19.0	<u>61.3</u>	<u>48.8</u>
Deepseek-Coder-6.7B	<u>19.3</u>	62.7	50.5
DATACODER (ours)	21.3	57.2	42.2

Finally, we analyze the performance of LLMs that undergone tuning via LoRA. From Table 6, we can find that: (1) Medium-sized models that have been fine-tuned consistently outperform DATACODER in terms of CodeBLEU and EA metrics. We attribute it to increased parameter size and the expanded pretraining dataset size of these models, which significantly influence LLM efficacy as suggested by prevailing scaling laws [42]. (2) Furthermore, comparing 15B-34B LLMs (Table 5) and tuning 1B-7B LLMs via LoRA, we find that tuning smaller models yields superior results compared to merely prompting larger models without tuning. Remarkably, the best DeepSeek-Coder-6.7B achieves comparable performance with GPT-4 (+8.2% EM, +1.9% CodeBLEU, and -1.1% EA). This result implies the potential of parameter-efficient finetuning using a high-quality domain-specific dataset to enhance data-wrangling




Figure 6: An example of predictions from models with diverse data context. Key errors are marked in red box.

code generation in notebooks, which in turn underscores the utility of our dataset.

5.2 RQ2: Effects of Varying Context

Setup: To understand the importance of context in enhancing code generation, we conduct experiments by assembling various context to form the model input. We select two top-performing models to study, *i.e.*, the finetuned DATACODER and GPT 3.5 with two in-context examples. We use the same split of dataset for evaluation.

Table 7: Results with Different Programmatic Context.

	DATACODER			GPT 3.5 (ICL $k=2$)		
	EM	CB	EA	EM	CB	EA
Full Context	21.3	57.2	42.2	6.9	55.3	40.1
- Code+Text	12.5	47.8	22.4	3.5	50.0	22.6
- Code+Text+InData	14.1	49.4	23.8	3.3	50.4	25.0
- Code+Text+OutData	18.4	53.2	34.5	4.6	52.0	29.4
- InData+OutData	4.7	36.3	5.4	0.7	41.7	7.6
- InData+OutData+Text	6.3	40.8	9.6	1.0	44.6	10.4
- InData+OutData+Code	18.3	53.4	35.1	6.2	51.9	34.9

¹ EM (Exact Match), CB (CodeBLEU), EA (Execution Accuracy).

² InData (Input Dataframe), OutData (Output Dataframe).

Results: From Table 7, we find that: (1) Removing data context (*i.e.*, **Code+Text**) leads to a substantial performance decline compared with **full context** (Code+Text+InData+OutData), indicating the indispensability of both dataframes for accurate inference of target codes. (2) We also compare the importance of input and output dataframes by removing them separately (*i.e.*, **Code+Text+InData** vs. **Code+Text+OutData**). The results show that the output dataframe is more crucial in this task, as it typically includes all the information from the input dataframe in most cases (*e.g.*, cases involving the addition of a new column to the input dataframe). (3) Models struggle to generate correct code with only data context (*i.e.*, **InData+OutData**). This can be attributed to the inherent complexity of real-world programming-by-example tasks, where constraints such as variable names and APIs must be ensured for execution

correctness. Moreover, the absence of runtime data contexts in the pretraining corpus of language models, as they are not typically saved in notebooks, further exacerbates the difficulty for current models to comprehend data context. (4) Involving more natural language context (*i.e.*, **InData+OutData+Text**) marginally improves the performance over **InData+OutData**. This is possibly due to the frequent absence of textual instructions in notebooks since programmers do not necessarily write comments or descriptions when they program. (5) Incorporating code context (*i.e.*, **InData+OutData+Code**) proves more effective than textual context in aiding data wrangling code generation, suggesting the importance of code context in mitigating variable or API misuse.

5.3 RQ3: Effects of Data Context

Setup: In this RQ, we aim to understand the significance of data context in enhancing code generation. To answer this question, we first experiment by eliminating various components of the dataframe in the data context. Subsequently, we provide a varying number of table rows to further study the impacts of data values. Finally, we present a case study to intuitively show the impact of various inputs on the prediction results.

Table 8: Results with Different Dataframe Components.

	DATACODER			GPT 3.5 (ICL $k=2$)		
	EM	CB	EA	EM	CB	EA
full data context	21.3	57.2	42.2	6.9	55.3	40.1
- w/o data values	19.6	55.3	38.2	6.2	54.1	34.5
- w/o column names	14.9	52.9	28.2	4.2	52.7	26.0

¹ EM (Exact Match), CB (CodeBLEU), EA (Execution Accuracy).

Results: We investigate the effects of different components in the data context, *i.e.*, column names and data values. From Table 8, we find that: (1) Removing values within the input and output dataframes while retaining only the column names (**w/o data values**) results in a decrease of execution accuracy (EA) from 42.2% to 38.2% for DATACODER and from 40.1% to 34.5% for GPT 3.5. This

result shows the importance of data values in enabling models to infer target code by comparing differences in values between input and output dataframes, particularly when code comments are absent. (2) Removing column names of dataframes (**w/o column names**) leads to a further performance decline, with EA dropping from 42.2% to 28.2% for DATACODER and from 40.1% to 26.0% for GPT 3.5. This aligns with our intuition that the model needs to be aware of the existing columns in the data and which columns in the input dataframe contribute to producing the output dataframe; otherwise, the generated code might result in a key error and fail to execute.

Case Study: In the example depicted in Figure 6, the comment “Convert the data in the smoker column to an indicator variable” provides crucial semantic information about the intended function of the target code. However, the model needs to discern which values in the “smoker” column should be mapped to which indicator variables and how to name the transformed columns. This information can be obtained by comparing the differences between the input and output dataframes. After using all information from both dataframes, we can see that DATACODER generates source codes that correctly process the input dataframe to produce the output dataframe. However, the models that are provided with partial data context fail to generate correct code. For example, without column names, the model mis-spells the key “smoke” to “Smoke”. Also, without data values, the model is unaware of the original flag (*i.e.*, “Yes” and “No”), and a flag “M” is predicted. Eliminating input or output dataframes also leads to incorrect results. For example, without input dataframes, the model hallucinates the value to be converted. Without output dataframe, the model mis-spells the expected key “smoke” to “is_smoker”. This example demonstrates the importance and data context and the effectiveness of our model.

6 RELATED WORK

6.1 Code Generation with Complex Context

Code generation is a longstanding challenge in software engineering and natural language processing [36, 49, 75]. Recent advances in large-scale pre-training have enabled code generation models to handle diverse tasks, which can generate target code with complex contexts in diverse modalities [65]. Textual context in natural language, such as human query [4, 31, 33, 103] or comments [54, 57, 77, 91, 97], can specify expected functions, which serves as high-level instructions in code generation [66]. In complement, code context [41], such as the code a programmer already written [37, 56, 60], are also useful as they can provide detailed instructions, such as code logic and variable names. There are also works focusing on multi-modal code generation combining text and code [14, 89, 92, 102]. Apart from text and code, data context is another useful modality. Text-to-SQL generation [100, 101] aims to generate SQL queries based on textual intents and sheets retrieved from databases. In this task, data context enables model to utilize important information, such as sheet names and column names, which can facilitate utilizing the correct data [43]. Programming-by-example is another challenging task [5, 16, 50, 69, 70, 72], where the specification is input-output examples. There are also efforts taking a step further that combine textual intents and input-output dataframes to generate table transformation code [38, 45]. Our work

aims to support multi-modal code generation with an emphasis on leveraging all above three modalities, *i.e.*, code, text, and data, for better code generation.

6.2 Data Science Code Generation

Code generation is important in automating low-level DS tasks[86]. Numerous models have been developed to synthesize code for various stages of the DS lifecycle, such as data preparation [8, 17, 96], modeling [64, 78], and visualization [7, 15, 95]. As data wrangling tasks consume a significant amount of time for data scientists [26], our focus is on automating data wrangling through code generation.

Computational notebooks is a popular programming tool in DS [53, 79]. JuiCe [1] is the first dataset to facilitate interactive code generation within notebooks, which inspired the development of various models [15, 21]. But these models are evaluated using surface-form metrics that cannot fully assess code quality as perceived by human developers. As a result, recent work has turned to execution-based metrics [12, 34, 48, 98] for code generation evaluation. However, these studies omit the importance of data context and only provide a testset for evaluation, which cannot be applied to contextualized code generation.

7 CONCLUSIONS AND FUTURE WORK

In this work, we introduce a novel task of contextualized data-wrangling code generation in notebooks with multi-modal context of code, text and tabular input-output data. To mine examples for this task, we propose CoCoMine that first extracts code cells relevant to data wrangling and then collects code generation examples with aligned multi-modal context, resulting in a dataset named CoCoNote with 58,221 examples to support training and execution-based evaluation. To better leverage input-output data, we propose DATACODER with separate encoders for code and text contexts and input-output dataframes to generate target code. Our experiments on a range of PLMs and LLMs verify the importance of data context and the effectiveness of our model.

In the future, we will explore the following directions: (1) notebooks involves runtime artifacts in a range of forms, *e.g.*, lists, images, machine learning models, and we will try to investigate the effects of these context in DS code generation; (2) automatic documentation in notebooks is also an important topic and we will investigate contextualized code documentation in notebooks; (3) there remains a large room for improvement and we will explore more powerful methods to perform this task.

8 ACKNOWLEDGEMENT

We thank all reviewers for their valuable comments and suggestions. The work described in this paper was supported by the Research Grants Council of the Hong Kong Special Administrative Region, China (No. CUHK 14206921 of the General Research Fund).

REFERENCES

- [1] Rajas Agashe, Srinivasan Iyer, and Luke Zettlemoyer. 2019. JuiCe: A Large Scale Distantly Supervised Dataset for Open Domain Context-based Code Generation. In *EMNLP-IJCNLP*.
- [2] Wasi Ahmad, Saikat Chakraborty, Baishakhi Ray, and Kai-Wei Chang. 2021. Unified Pre-training for Program Understanding and Generation. In *NAACL-HLT*. 2655–2668.

- [3] Loubna Ben Allal, Raymond Li, Denis Kocetkov, Chenghao Mou, Christopher Akiki, Carlos Muñoz Ferrandis, Niklas Muennighoff, Mayank Mishra, Alex Gu, Manan Dey, et al. 2023. SantaCoder: don't reach for the stars! *arXiv preprint arXiv:2301.03988* (2023).
- [4] Jacob Austin, Augustus Odena, Maxwell Nye, Maarten Bosma, Henryk Michalewski, David Dohan, Ellen Jiang, Carrie Cai, Michael Terry, Quoc Le, et al. 2021. Program synthesis with large language models. *arXiv preprint arXiv:2108.07732* (2021).
- [5] Daniel W Barowy, Sumit Gulwani, Ted Hart, and Benjamin Zorn. 2015. FlashRelate: extracting relational data from semi-structured spreadsheets using examples. *ACM SIGPLAN Notices* 50, 6 (2015), 218–228.
- [6] Mohammad Bavarian, Heewoo Jun, Nikolas Tezak, John Schulman, Christine McLeavey, Jerry Tworek, and Mark Chen. 2022. Efficient training of language models to fill in the middle. *arXiv preprint arXiv:2207.14255* (2022).
- [7] Rohan Bavishi, Shadaj Laddad, Hiroaki Yoshida, Mukul R Prasad, and Koushik Sen. 2021. Vizsmith: Automated visualization synthesis by mining data-science notebooks. In *ASE*.
- [8] Rohan Bavishi, Caroline Lemieux, Roy Fox, Koushik Sen, and Ion Stoica. 2019. AutoPandas: neural-backed generators for program synthesis. *Proceedings of the ACM on Programming Languages* 3, OOPSLA (2019), 1–27.
- [9] Ekaba Bisong. 2019. Google laboratory. Building machine learning and deep learning models on google cloud platform. *Apress, Berkeley, CA* (2019), 59–64.
- [10] Sumon Biswas, Mohammad Wardat, and Hriday Rajan. 2022. The art and practice of data science pipelines: A comprehensive study of data science pipelines in theory, in-the-small, and in-the-large. In *ICSE*.
- [11] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *NeurIPS* (2020).
- [12] Shubham Chandel, Colin B. Clement, Guillermo Serrato, and Neel Sundaresan. 2022. Training and Evaluating a Jupyter Notebook Data Science Assistant. *arXiv preprint arXiv:2201.12901* (2022).
- [13] Adriane Chapman, Luca Lauro, Paolo Missier, and Riccardo Torlone. 2024. Supporting Better Insights of Data Science Pipelines with Fine-grained Provenance. *ACM Trans. Database Syst.* 49, 2 (2024), 42 pages.
- [14] Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde de Oliveira Pinto, Jared Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, et al. 2021. Evaluating large language models trained on code. *arXiv preprint arXiv:2107.03374* (2021).
- [15] Xinyun Chen, Linyuan Gong, Alvin Cheung, and Dawn Song. 2021. Plotcoder: Hierarchical decoding for synthesizing visualization code in programmatic context. In *ACL-IJCNLP*. 2169–2181.
- [16] Xinyun Chen, Dawn Song, and Yundong Tian. 2021. Latent execution for neural program synthesis beyond domain-specific languages. *Advances in Neural Information Processing Systems* 34 (2021), 22196–22208.
- [17] Bhavya Chopra, Anna Farha, Sumit Gulwani, Austin Z Henley, Daniel Perelman, Mohammad Raza, Sherry Shi, Danny Simmons, and Ashish Tiwari. 2023. CoWrangler: Recommender System for Data-Wrangling Scripts. In *Companion of the 2023 International Conference on Management of Data*. 147–150.
- [18] 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*. 2342–2354.
- [19] Tamraparni Dasu, Theodore Johnson, Shanmugaelayuth Muthukrishnan, and Vladislav Shkapenyuk. 2002. Mining database structure; or, how to build a data quality browser. In *SIGMOD*. 240–251.
- [20] Li Dong, Nan Yang, Wenhui Wang, Furu Wei, Xiaodong Liu, Yu Wang, Jianfeng Gao, Ming Zhou, and Hsiao-Wuen Hon. 2019. Unified language model pre-training for natural language understanding and generation. *NeurIPS* (2019).
- [21] Yihong Dong, Ge Li, and Zhi Jin. 2023. CODEP: grammatical seq2seq model for general-purpose code generation. In *ISSTA*. 188–198.
- [22] Zhangyin Feng, Daya Guo, Duyu Tang, Nan Duan, Xiaocheng Feng, Ming Gong, Linjun Shou, Bing Qin, Ting Liu, Daxin Jiang, et al. 2020. CodeBERT: A Pre-Trained Model for Programming and Natural Languages. In *Findings of the Association for Computational Linguistics: EMNLP*. 1536–1547.
- [23] Leo Gao, Stella Biderman, Sid Black, Laurence Golding, Travis Hoppe, Charles Foster, Jason Phang, Horace He, Anish Thite, Noa Nabeshima, et al. 2020. The pile: An 800gb dataset of diverse text for language modeling. *arXiv preprint arXiv:2101.00027* (2020).
- [24] Shuzheng Gao, Xin-Cheng Wen, Cuiyun Gao, Wenxuan Wang, Hongyu Zhang, and Michael R Lyu. 2023. What Makes Good In-Context Demonstrations for Code Intelligence Tasks with LLMs? In *ASE*.
- [25] GitHub. 2024. GitHub REST API. <https://developer.github.com/v3>
- [26] Sumit Gulwani. 2016. Programming by Examples - and its applications in Data Wrangling. In *Dependable Software Systems Engineering*.
- [27] Sumit Gulwani, William R Harris, and Rishabh Singh. 2012. Spreadsheet data manipulation using examples. *Commun. ACM* 55, 8 (2012), 97–105.
- [28] Daya Guo, Shuai Lu, Nan Duan, Yanlin Wang, Ming Zhou, and Jian Yin. 2022. UniXcoder: Unified Cross-Modal Pre-training for Code Representation. In *ACL*. 7212–7225.
- [29] Daya Guo, Shuo Ren, Shuai Lu, Zhangyin Feng, Duyu Tang, Liu Shujie, Long Zhou, Nan Duan, Alexey Svyatkovskiy, Shengyu Fu, et al. 2020. GraphCodeBERT: Pre-training Code Representations with Data Flow. In *ICLR*.
- [30] Daya Guo, Qihao Zhu, Dejian Yang, Zhenda Xie, Kai Dong, Wentao Zhang, Guanting Chen, Xiao Bi, Y Wu, YK Li, et al. 2024. DeepSeek-Coder: When the Large Language Model Meets Programming—The Rise of Code Intelligence. *arXiv preprint arXiv:2401.14196* (2024).
- [31] 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 *NeurIPS*.
- [32] Edward J Hu, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, Weizhu Chen, et al. 2023. LoRA: Low-Rank Adaptation of Large Language Models. In *ICLR*.
- [33] Junjie Huang, Duyu Tang, Linjun Shou, Ming Gong, Ke Xu, Daxin Jiang, Ming Zhou, and Nan Duan. 2021. CoSQA: 20,000+ Web Queries for Code Search and Question Answering. In *ACL-IJCNLP*. 5690–5700.
- [34] Junjie Huang, Chenglong Wang, Jipeng Zhang, Cong Yan, Haotian Cui, Jeevana Priya Inala, Colin B. Clement, Nan Duan, and Jianfeng Gao. 2022. Execution-based Evaluation for Data Science Code Generation Models. In *Proceedings of the Fourth Workshop on Data Science with Human-in-the-Loop (Language Advances)*. 28–36.
- [35] Junjie Huang, Wanjun Zhong, Qian Liu, Ming Gong, Daxin Jiang, and Nan Duan. 2022. Mixed-modality Representation Learning and Pre-training for Joint Table-and-Text Retrieval in OpenQA. In *Findings of the Association for Computational Linguistics: EMNLP 2022*. 4117–4129.
- [36] Yizhan Huang, Yichen Li, Weibin Wu, Jianping Zhang, and Michael R Lyu. 2024. Your Code Secret Belongs to Me: Neural Code Completion Tools Can Memorize Hard-Coded Credentials. *Proceedings of the ACM on Software Engineering* 1, FSE (2024), 2515–2537.
- [37] Srinivasan Iyer, Ioannis Konstas, Alvin Cheung, and Luke Zettlemoyer. 2018. Mapping Language to Code in Programmatic Context. In *EMNLP*. 1643–1652.
- [38] Naman Jain, Skanda Vaidyanath, Arun Iyer, Nagarajan Natarajan, Suresh Parthasarathy, Sriram Rajamani, and Rahul Sharma. 2022. Jigsaw: Large language models meet program synthesis. In *ICSE*.
- [39] Albert Q Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, et al. 2023. Mistral 7B. *arXiv preprint arXiv:2310.06825* (2023).
- [40] Kaggle. 2024. Kaggle. <https://www.kaggle.com>
- [41] Aditya Kanade, Petros Maniatis, Gogul Balakrishnan, and Kensen Shi. 2020. Learning and evaluating contextual embedding of source code. In *International conference on machine learning*. PMLR, 5110–5121.
- [42] Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B Brown, Benjamin Chess, Rewon Child, Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. 2020. Scaling laws for natural language models. *arXiv preprint arXiv:2001.08361* (2020).
- [43] George Katsogiannis-Meimarakis and Georgia Koutrika. 2023. A survey on deep learning approaches for text-to-SQL. *The VLDB Journal* 32, 4 (2023), 905–936.
- [44] Mary Beth Kery, Marissa Radensky, Mahima Arya, Bonnie E John, and Brad A Myers. 2018. The story in the notebook: Exploratory data science using a literate programming tool. In *CHI*. 1–11.
- [45] Anirudh Khatri, Joyce Cahoon, Jordan Henkel, Shaleen Deep, Venkatesh Emani, Avriella Floratou, Sumit Gulwani, Vu Le, Mohammad Raza, Sherry Shi, et al. 2023. From Words to Code: Harnessing Data for Program Synthesis from Natural Language. *arXiv preprint arXiv:2305.01598* (2023).
- [46] Thomas Kluyver, Benjamin Ragan-Kelley, Fernando Pérez, Brian E. Granger, Matthias Bussonnier, Jonathan Frederic, Kyle Kelley, Jessica B. Hamrick, Jason Grout, Sylvain Corlay, Paul Ivanov, Damián Avila, Safia Abdalla, Carol Willing, and Jupyter Development Team. 2016. Jupyter Notebooks - a publishing format for reproducible computational workflows. In *ELPUB*.
- [47] Denis Kocetkov, Raymond Li, Loubna Ben Allal, Jia Li, Chenghao Mou, Carlos Muñoz Ferrandis, Yacine Jermite, Margaret Mitchell, Sean Hughes, Thomas Wolf, et al. 2023. The stack: 3 tb of permissively licensed source code. *Transactions on Machine Learning Research* (2023).
- [48] Yuhang Lai, Chengxi Li, Yiming Wang, Tianyi Zhang, Ruiqi Zhong, Luke Zettlemoyer, Wen-tau Yih, Daniel Fried, Sida Wang, and Tao Yu. 2023. DS-1000: A natural and reliable benchmark for data science code generation. In *International Conference on Machine Learning*. PMLR, 18319–18345.
- [49] 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.
- [50] Vu Le and Sumit Gulwani. 2014. Flashextract: A framework for data extraction by examples. In *PLDI*.
- [51] Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020. BART: Denoising Sequence-to-Sequence Pre-training for Natural Language Generation, Translation, and Comprehension. In *ACL*. 7871–7880.

- [52] Raymond Li, Loubna Ben Allal, Yangtian Zi, Niklas Muennighoff, Denis Kocetkov, Chenghao Mou, Marc Marone, Christopher Akiki, Jia Li, Jenny Chim, et al. 2023. Starcoder: may the source be with you! *arXiv preprint arXiv:2305.06161* (2023).
- [53] Xingjun Li, Yizhi Zhang, Justin Leung, Chengnian Sun, and Jian Zhao. 2023. Edassistant: Supporting exploratory data analysis in computational notebooks with in situ code search and recommendation. *ACM Transactions on Interactive Intelligent Systems* 13, 1 (2023), 1–27.
- [54] Yichen Li, Yintong Huo, Zhihan Jiang, Renyi Zhong, Pinjia He, Yuxin Su, and Michael R Lyu. 2023. Exploring the effectiveness of llms in automated logging generation: An empirical study. *arXiv preprint arXiv:2307.05950* (2023).
- [55] Yichen Li, Yintong Huo, Renyi Zhong, Zhihan Jiang, Jinyang Liu, Junjie Huang, Jiazhen Gu, Pinjia He, and Michael R Lyu. 2024. Go static: Contextualized logging statement generation. *Proceedings of the ACM on Software Engineering* 1, FSE (2024), 609–630.
- [56] Yichen Li, Yun Peng, Yintong Huo, and Michael R Lyu. 2024. Enhancing LLM-Based Coding Tools through Native Integration of IDE-Derived Static Context. *arXiv preprint arXiv:2402.03630* (2024).
- [57] Xi Victoria Lin, Chenglong Wang, Luke Zettlemoyer, and Michael D Ernst. 2018. NL2Bash: A Corpus and Semantic Parser for Natural Language Interface to the Linux Operating System. In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*.
- [58] Jiawei Liu, Chunqiu Steven Xia, Yuyao Wang, and Lingming Zhang. 2024. Is your code generated by chatgpt really correct? rigorous evaluation of large language models for code generation. *NeurIPS* (2024).
- [59] Ilya Loshchilov and Frank Hutter. 2019. Decoupled Weight Decay Regularization. In *ICLR*.
- [60] Shuai Lu, Daya Guo, Shuo Ren, Junjie Huang, Alexey Svyatkovskiy, Ambrosio Blanco, Colin Clement, Dawn Drain, Daxin Jiang, Duyu Tang, Ge Li, Lidong Zhou, Linjun Shou, Long Zhou, Michele Tufano, MING GONG, Ming Zhou, Nan Duan, Neel Sundaresan, Shao Kun Deng, Shengyu Fu, and Shujie LIU. 2021. CodeXGLUE: A Machine Learning Benchmark Dataset for Code Understanding and Generation. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track*.
- [61] Wes McKinney et al. 2010. Data structures for statistical computing in Python. Vol. In *SciPy*, Vol. 445, 51–56.
- [62] Na Meng, Miryung Kim, and Kathryn S McKinley. 2013. LASE: locating and applying systematic edits by learning from examples. In *ICSE*.
- [63] Tamal Mondal, Scott Barnett, Akash Lal, and Jyothi Vedurada. 2023. Cell2Doc: ML Pipeline for Generating Documentation in Computational Notebooks. In *ASE*.
- [64] Daye Nam, Baishakhi Ray, Seohyun Kim, Xianshan Qu, and Satish Chandra. 2022. Predictive synthesis of api-centric code. In *Proceedings of the 6th ACM SIGPLAN International Symposium on Machine Programming*, 40–49.
- [65] Erik Nijkamp, Bo Pang, Hiroaki Hayashi, Lifu Tu, Huan Wang, Yingbo Zhou, Silvio Savarese, and Caiming Xiong. 2022. Codegen: An open large language model for code with multi-turn program synthesis. In *ICLR*.
- [66] Yusuke Oda, Hiroyuki Fudaba, Graham Neubig, Hideaki Hata, Sakriani Sakti, Tomoki Toda, and Satoshi Nakamura. 2015. Learning to Generate Pseudo-Code from Source Code Using Statistical Machine Translation. *ASE* (2015), 574–584.
- [67] OpenAI. 2024. Introducing ChatGPT. <https://openai.com/blog/chatgpt>
- [68] Stack Overflow. 2024. Stack Overflow. <https://stackoverflow.com/>
- [69] Emilio Parisotto, Abdel rahman Mohamed, Rishabh Singh, Lihong Li, Dengyong Zhou, and Pushmeet Kohli. 2017. Neuro-Symbolic Program Synthesis. In *ICLR*.
- [70] Richard E. Pattis, Jim Roberts, and Mark Stehlik. 1994. Karel the Robot: A Gentle Introduction to the Art of Programming.
- [71] Phind. 2024. Phind/phind-codellama-34b-v2. <https://huggingface.co/Phind/Phind-CodeLlama-34B-v2>
- [72] Illia Polosukhin and Alexander Skidanov. 2018. Neural Program Search: Solving Data Processing Tasks from Description and Examples. In *ICLR*.
- [73] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *The Journal of Machine Learning Research* 21, 1 (2020), 5485–5551.
- [74] Shuo Ren, Daya Guo, Shuai Lu, Long Zhou, Shujie Liu, Duyu Tang, Neel Sundaresan, Ming Zhou, Ambrosio Blanco, and Shuai Ma. 2020. Codebleu: a method for automatic evaluation of code synthesis. *arXiv preprint arXiv:2009.10297* (2020).
- [75] Baptiste Roziere, Jonas Gehring, Fabian Gloeckle, Sten Sootla, Itai Gat, Xiaoqing Ellen Tan, Yossi Adi, Jingyu Liu, Tal Remez, Jérémy Rapin, et al. 2023. Code llama: Open foundation models for code. *arXiv preprint arXiv:2308.12950* (2023).
- [76] Adam Rule, Aurélien Tabard, and James D Hollan. 2018. Exploration and explanation in computational notebooks. In *CHI*, 1–12.
- [77] Freda Shi, Daniel Fried, Marjan Ghazvininejad, Luke Zettlemoyer, and Sida I Wang. 2022. Natural Language to Code Translation with Execution. In *EMNLP*.
- [78] Kensen Shi, David Bieber, and Rishabh Singh. 2022. Tf-coder: Program synthesis for tensor manipulations. *ACM Transactions on Programming Languages and Systems (TOPLAS)* 44, 2 (2022), 1–36.
- [79] Pavle Subotić, Lazar Milikić, and Milan Stojić. 2022. A static analysis framework for data science notebooks. In *ICSE-SEIP*, 13–22.
- [80] Sergey Titov, Yaroslav Golubev, and Timofey Bryksin. 2022. Resplit: improving the structure of jupyter notebooks by re-splitting their cells. In *SANER*. IEEE, 492–496.
- [81] Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shrutu Bhosale, et al. 2023. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288* (2023).
- [82] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *NeurIPS*.
- [83] April Yi Wang, Dakuo Wang, Jaimie Drozdal, Michael Muller, Soya Park, Justin D Weisz, Xuye Liu, Lingfei Wu, and Casey Dugan. 2022. Documentation Matters: Human-Centered AI System to Assist Data Science Code Documentation in Computational Notebooks. *ACM Transactions on Computer-Human Interaction* 29, 2 (2022), 1–33.
- [84] Chenglong Wang, John Thompson, and Bongshin Lee. 2023. Data Formulator: AI-powered Concept-driven Visualization Authoring. *IEEE Transactions on Visualization and Computer Graphics* (2023).
- [85] Chaozheng Wang, Yuanhang Yang, Cuiyun Gao, Yun Peng, Hongyu Zhang, and Michael R Lyu. 2022. No more fine-tuning? an experimental evaluation of prompt tuning in code intelligence. In *ESEC/FSE*, 382–394.
- [86] Dakuo Wang, Josh Andres, Justin D. Weisz, Erick Oduor, and Casey Dugan. 2021. AutoDS: Towards Human-Centered Automation of Data Science. *CHI* (2021).
- [87] Jiawei Wang, Tzu-yang Kuo, Li Li, and Andreas Zeller. 2020. Assessing and restoring reproducibility of Jupyter notebooks. In *ASE*.
- [88] Jiawei Wang, Li Li, and Andreas Zeller. 2021. Restoring Execution Environments of Jupyter Notebooks. In *ICSE*, 1622–1633.
- [89] Shiqi Wang, Zheng Li, Haifeng Qian, Chenghao Yang, Zijian Wang, Mingyue Shang, Varun Kumar, Samson Tan, Baishakhi Ray, Parminder Bhatia, et al. 2023. ReCode: Robustness Evaluation of Code Generation Models. In *ACL*.
- [90] Yue Wang, Weishi Wang, Shafiq Joty, and Steven CH Hoi. 2021. CodeT5: Identifier-aware Unified Pre-trained Encoder-Decoder Models for Code Understanding and Generation. In *EMNLP*, 8696–8708.
- [91] Zhiruo Wang, Grace Cuenca, Shuyan Zhou, Frank F Xu, and Graham Neubig. 2022. Mconala: a benchmark for code generation from multiple natural languages. *arXiv preprint arXiv:2203.08388* (2022).
- [92] Zhiruo Wang, Shuyan Zhou, Daniel Fried, and Graham Neubig. 2022. Execution-Based Evaluation for Open-Domain Code Generation. *arXiv preprint arXiv:2212.10481* (2022).
- [93] Zijie J Wang, David Munechika, Seongmin Lee, and Duen Horng Chau. 2023. SuperNOVA: Design Strategies and Opportunities for Interactive Visualization in Computational Notebooks. *arXiv preprint arXiv:2305.03039* (2023).
- [94] Martin Weyssow, Xin Zhou, Kisub Kim, David Lo, and Houari Sahraoui. 2023. Exploring parameter-efficient fine-tuning techniques for code generation with large language models. *arXiv preprint arXiv:2308.10462* (2023).
- [95] Yifan Wu, Joseph M Hellerstein, and Arvind Satyanarayan. 2020. B2: Bridging code and interactive visualization in computational notebooks. In *UIST*, 152–165.
- [96] Cong Yan and Yeye He. 2020. Auto-suggest: Learning-to-recommend data preparation steps using data science notebooks. In *SIGMOD*, 1539–1554.
- [97] Pengcheng Yin, Bowen Deng, Edgar Chen, Bogdan Vasilescu, and Graham Neubig. 2018. Learning to Mine Aligned Code and Natural Language Pairs from Stack Overflow. *2018 IEEE/ACM 15th International Conference on Mining Software Repositories (MSR)* (2018), 476–486.
- [98] Pengcheng Yin, Wen-Ding Li, Kefan Xiao, Abhishek Rao, Yeming Wen, Kensen Shi, Joshua Howland, Paige Bailey, Michele Catasta, Henryk Michalewski, et al. 2023. Natural Language to Code Generation in Interactive Data Science Notebooks. In *ACL*, 126–173.
- [99] Pengcheng Yin, John Wieting, Avirup Sil, and Graham Neubig. 2022. On The Ingredients of an Effective Zero-shot Semantic Parser. In *ACL*, 1455–1474.
- [100] Tao Yu, Rui Zhang, Kai-Chou Yang, Michihiro Yasunaga, Dongxu Wang, Zifan Li, James Ma, Irene Z Li, Qingning Yao, Shanelle Roman, Zilin Zhang, and Dragomir R. Radev. 2018. Spider: A Large-Scale Human-Labeled Dataset for Complex and Cross-Domain Semantic Parsing and Text-to-SQL Task. In *EMNLP*.
- [101] Victor Zhong, Caiming Xiong, and Richard Socher. 2017. Seq2sql: Generating structured queries from natural language using reinforcement learning. *arXiv preprint arXiv:1709.00103* (2017).
- [102] Shuyan Zhou, Uri Alon, Frank F Xu, Zhengbao Jiang, and Graham Neubig. 2022. DocCoder: Generating Code by Retrieving and Reading Docs. *arXiv preprint arXiv:2207.05987* (2022).
- [103] Ming Zhu, Aneesh Jain, Karthik Suresh, Roshan Ravindran, Sindhu Tipirneni, and Chandan K Reddy. 2022. Xlcost: A benchmark dataset for cross-lingual code intelligence. *arXiv preprint arXiv:2206.08474* (2022).