

CrossCodeBench: Benchmarking Cross-Task Generalization of Source Code Models

Changan Niu*, Chuanyi Li*, Vincent Ng[†], Bin Luo*

*State Key Laboratory for Novel Software Technology, Nanjing University, Nanjing

Email: niu.ca@outlook.com, {lcy,luobin}@nju.edu.cn

[†]Human Language Technology Research Institute, University of Texas at Dallas, Richardson, Texas, USA

Email: vince@hlt.utdallas.edu

Abstract—Despite the recent advances showing that a model pre-trained on large-scale source code data is able to gain appreciable generalization capability, it still requires a sizeable amount of data on the target task for fine-tuning. And the effectiveness of the model generalization is largely affected by the size and quality of the fine-tuning data, which is detrimental for target tasks with limited or unavailable resources. Therefore, cross-task generalization, with the goal of improving the generalization of the model to unseen tasks that have not been seen before, is of strong research and application value.

In this paper, we propose a large-scale benchmark that includes 216 existing code-related tasks. Then, we annotate each task with the corresponding meta information such as task description and instruction, which contains detailed information about the task and a solution guide. This also helps us to easily create a wide variety of “training/evaluation” task splits to evaluate the various cross-task generalization capabilities of the model. Then we perform some preliminary experiments to demonstrate that the cross-task generalization of models can be largely improved by in-context learning methods such as few-shot learning and learning from task instructions, which shows the promising prospects of conducting cross-task learning research on our benchmark. We hope that the collection of the datasets and our benchmark will facilitate future work that is not limited to cross-task generalization.

Index Terms—Pre-training of source code, cross-task transfer learning, few-shot learning, AI for SE

I. INTRODUCTION

The “pre-train then fine-tune” paradigm has become the new favorite in software intelligence [1]. Since pre-training tasks can be done with unlabeled data, a model can be pre-trained on large amounts of, easily accessible data, thus obtaining much common sense and linguistic knowledge [2]–[4]. With this knowledge, pre-trained models can achieve better generalizability, which means that it is able to perform better than its “no pre-training” counterpart on a wide variety of software engineering (SE) tasks after being fine-tuned on the data of the target task [5]–[11]. Therefore, rather than learning models from scratch, the adoption of pre-trained models as backbone for downstream tasks has become a common practice in the field of software intelligence [6], [12], [13].

However, the fine-tuning stage requires updating the weights of a pre-trained model by training on thousands of supervised labels specific to the target task. Therefore, the “pre-train then fine-tune” paradigm still relies on the data from the target

task. And the effectiveness of transfer learning of the pre-trained model to the target task depends heavily on the size and quality of the fine-tuning data [10]. Unfortunately, in practice, we often encounter situations where we need to apply a pre-trained model to a task with very low available data resources. In such cases, we are unable to fine-tune the pre-trained model on sufficient target task data to obtain a fine-tuned model that can be applied to the target task.

Facing the same issue, pioneers in the field of Natural Language Processing (NLP) have proposed many ways to address this issue. Few-shot learning [14]–[16] skips the fine-tuning phase and applies the pre-trained model directly to the target task, without any weight updates. To bridge the gap between the pre-trained model and the target task, few-shot learning gives the model a task description and a few demonstrations (i.e., supervised examples) of the task at inference time. If there is only one positive example, it is also called one-shot learning [17]. Moreover, instead of any examples, zero-shot learning gives the model a natural language description of the task [18], [19]. In the few/one/zero-shot settings, a large-scale pre-trained language model is able to show strong performance on many NLP tasks and benchmarks, in some cases nearly matching the performance of state-of-the-art fine-tuned systems [14], [16], [18]–[20]. In addition, learning from task instructions [21]–[23] adopts task definition, positive and negative examples, where task definition can be seen as a specification of solving the target task.

The SE community also does some explorations. Mastroiolo et al. [7] and Wang et al. [9] utilizes multi-task learning [19], [24] to achieve a better performance on the target task. By learning multiple related tasks simultaneously, multi-task learning aims to make models exploit both task-generic and task-specific information, thereby improving the model’s performance on tasks with low available resources. However, multi-task learning favors tasks with significantly larger amounts of data than others, thus requiring sufficient supervised examples of the target task compared to other tasks to guarantee the availability of the model [25], [26]. Rather, Prenner and Robbes [27] experiments with several other techniques that promised a possible benefit for small datasets, i.e., active learning, data augmentation, soft labels, self-training and intermediate-task fine-tuning [28]. They find that soft labels to be more useful, while other methods

are relatively more narrowly applicable, less effective, more costly, or inconclusive. Instead of using any data from some target tasks, Guo et al. [11] directly apply the pre-trained model to the code-to-code retrieval task in order to evaluate the performance of code fragment embeddings. Given that there are many very large-scale pre-trained models of source code being proposed (e.g., GitHub Copilot, Codex [29] and AlphaCode [30]), there is also a lot of work exploring the few/zero-shot performance of these models on specific tasks and domains, such as program repair [31]–[33], software security [34], [35] and program synthesis [36].

All things considered, in the field of software intelligence, there is no systematic work to evaluate and explore the cross-task generalizability of code models. Therefore, in order to evaluate the cross-task capability of the model on code-related tasks in detail and comprehensively, we build a large-scale benchmark called CrossCodeBench. We start by collecting as many and as diverse code-related tasks as possible, and end up with 216 tasks across 28 categories, 7 types and 18 programming languages. Then, to make our benchmark available for multiple cross-task learning settings (e.g., few-shot learning, learning from task instructions), we manually label each of the 216 tasks with extensive meta-information such as task description, definition, positive/negative examples, etc. Next, we create 10 training/evaluation splits corresponding to different benchmark types.

Given the data splits, we perform experiments by adopting two types of pre-trained models: (1) off-the-shelf model, which are used directly on the evaluation set, and (2) models further fine-tuned on the training set. All models are applied to all splits by using all or suitable cross-task learning methods, such as few-shot learning, learning from instructions, etc. Last but not least, we carry out a number of scaling experiments, through which we find that (1) when the data of each task reaches a certain level (e.g., 10,000 instances), it is difficult for the model to maintain a high improving speed of performance as data instances increasing, and (2) larger models always lead to better performance. We hope that the benchmark, experimental results and analysis we provide will facilitate future research to more powerful cross-task approaches in SE literature. Furthermore, since our benchmark contains massive datasets (and is open to updates), not only our CrossCodeBench, but we hope that such a large-scale meta-dataset (i.e., dataset of datasets [37], [38]) will facilitate the construction of more benchmarks¹.

II. RELATED WORK

A. Few-Shot, One-Shot, and Zero-Shot Learning

Albeit defeating human in many fields [18], [39]–[43], current artificial intelligence (AI) techniques still rely on learning from large-scale task-specific data, and they are unable to rapidly generalize from a few examples. Rather, humans are

able to learn new tasks quickly by using what they are born with, or what they have learned in the past.

Few-shot learning (FSL) is therefore proposed in order to learn from a limited number of examples with supervised information. In the cross-task setup, it is an in-context learning approach where a pre-trained language model does not need any fine-tuning and weight updating [44]². The input can be divided into three parts, namely task description, examples, and prompt [18]. The task description is a typically short natural language description of the task, e.g. “translate English to French”. Examples consist of k canonical supervised examples (usually $10 < k < 100$, one-shot learning for $k = 1$ [18]), each of which includes the context (i.e., the input/question of the example) and the desired completion (i.e., the output/answer of the example). And the prompt is the context part of the example for which the model needs to make a prediction.

FSL shows promising results compared to the supervised approaches. In the FSL settings, GPT-3 [18], a large pre-trained language model, significantly improves the state-of-the-art (SOTA) on various datasets across many task types, such as completion [47], open-domain question answering [48] and translation [49]. In addition, Madotto et al. [50] demonstrate that in some task-oriented dialogue system tasks, language model priming FSL is able to achieve similar or better results than fine-tuning-based baseline. Chen et al. [51] show that in the FSL settings, a language model can achieve very reasonable performances and outperforms the strongest baseline by an average of over 8.0 BLEU points improvement, across multiple domains. Under the multilingual translation setting, Winata et al. [52] find the in-context few-shot cross-lingual prediction results of language models are comparative to the existing SOTA cross-lingual and translation models. In order to investigate whether and how cross-task generalization ability can be acquired, Ye et al. [20] propose CrossFit challenge, a task setup that standardizes the training pipeline, data access and evaluation protocol. As a complement, they present the NLP Few-shot Gym, a repository of 160 diverse few-shot NLP tasks. Experimental results show that the cross-task generalization ability can be obtained by using multi-task and meta learning, and can be affected by the selection of seen tasks.

In contrast to FSL, zero-shot learning (ZSL) is proposed to enable generalization to the target task without any examples. Alternatively stated, the language model is given only the task description and the prompt. As only a short task description is required, ZSL is able to provide maximum convenience, potential for robustness, and avoidance of spurious correlations for transferring language models to new tasks [18]. However, this presents huge challenges, such as the possibility of ambiguity in the task description, the absence of examples of output formats, etc. Even so, on some tasks, zero-shot GPT-3 can still outperform one-shot GPT-3 (e.g., completion [47],

¹All datasets, tasks and their summaries are available at <https://doi.org/10.5281/zenodo.7321934>. Source code is available at <https://github.com/NougatCA/CrossCodeBench>.

²Since our work only discusses the cross-task scenario, i.e., where no supervised training is performed on the data of the target task, we only introduce the FSL methods that can be applied in this scenario, for more FSL methods please refer to Wang et al. [45] and Yin [46].

[53]), or even few-shot GPT-3 (common sense reasoning [54]). Besides, on ANLI [55], GPT-3 under the zero-shot setting scores higher than under the few-shot and one-shot settings for some parameter size settings (i.e., 0.1B, 2.6B and 6.7B).

B. Learning from Task Instructions

In addition to investigating on the number of examples, researchers do some exploration on task description as well. Recall that in the FSL and ZSL settings, task descriptions are often short (e.g., 12.6 tokens on average [23]), sometimes causing ambiguity, sometimes missing necessary formatting instructions, etc. Therefore, researchers investigated the use of longer and more detailed task descriptions.

Inspired by current NLP datasets built using crowdsourcing, Efrat and Levy [21] examine if language models can follow crowdsourcing instructions with no further training. Weller et al. [22] construct a crowdsourced dataset, called Zest, with question-like task descriptions. To study the ability of a model that learns a new task by understanding the human-readable instructions that define it, Mishra et al. [23] introduce a dataset including 61 distinct tasks, their human-authored instructions and instances, which is named NatInst. Then they adopt existing pre-trained language models to encode task-specific instructions along with input and generate task output. Results show that the model can benefit from instructions and can improve performance by 19% when evaluated in generalization for unseen tasks. PromptSource [56], FLAN [57], and InstructGPT [58] also study the cross-task generalization ability by following the provided in-context task instructions. As a subsequent work to NatInst, Wange et al. [38] introduce NatInst_{v2}, a benchmark of over 1,600 diverse language tasks and their expert-written instructions, which covers 70+ distinct task types. They also propose Tk -Instruct, a model trained to follow a variety of in-context instructions, which include plain language task definitions and k -shot examples.

The above methods can be collectively referred to as learning from task instructions (LTI). Although the above work differs in the content of task instructions, after reading a lot of recent related work, we believe that a typical task instruction can generally be divided into the following parts: task definition, positive (and negative) examples, and other elements (e.g., “Things to Avoid” [23]). Task definition is a detailed definition of the task, unlike task description in FSL, it details how to map the given input to the required output in the current task. As a result, task description is typically longer and more detailed than the task description in FSL³. In addition, task instruction sometimes contains negative instances as opposed to only positive instances in FSL.

By learning in-context task instructions, models such as InstructGPT and NatInst_{v2} are able to achieve better cross-task generalization performance on unseen tasks than current FSL-based models. For example, on the FLAN benchmark,

³For example, the task definition of the task description “Translate English to French” is “Given a sentence in English, provide an equivalent paraphrased translation in French that retains the same meaning both through the translation and the paraphrase.” [38]

TABLE I
COMPARISON OF CROSSCODEBENCH AND OTHER CODE-RELATED BENCHMARKS

Benchmark	CrossCodeBench	CodeXGLUE	CodeTrans	XLCoST	CodeNet
Has description?	✓	✗	✗	✗	✗
Has definition?	✓	✗	✗	✗	✗
Has examples?	✓	✗	✗	✗	✗
Is off-the-shelf?	✓	✓	✓	✓	✗
# of tasks	216	38	13	112	-
# of types	7/10	6+1/7+1	2/3	3+1/3+1	-
# of categories	28	14	4	5	-
# of datasets	66	13	6	1	1
# of PLs	18	9	9	7	55
# of instances	54M	6.53M	9.32M	3.38M	13.9M

InstructGPT has about 76.2% win rate compared to baselines such as FSL-based GPT-3 [18], and the 3B-parameter Tk -Instruct outperforms 175B-parameter InstructGPT by 3.3 ROUGE-L points [59] when evaluated on 119 unseen tasks.

C. Code-Related Multi-Task Benchmarks

There are existing benchmarks that span multiple code-related tasks. CodeXGLUE [60], a benchmark which includes 14 datasets for 10 diversified code-related tasks covering code-to-code, text-to-code, code-to-text and text-to-text scenarios. CodeXGLUE is now widely used to evaluate the performance of pre-trained models of source code on various downstream tasks [9], [11]. Elnaggar et al. [61] collect a benchmark that contains 6 code-related tasks across 9 programming languages (PLs). Compared to these benchmarks, our work provides a larger scale and more diverse code-related tasks (216 tasks and more than 54M data instances in total). In addition, we also provide systematic and reasonable cross-task task splits, which can comprehensively evaluate cross-task capabilities in various scenarios.

XLCoST [62], a benchmark for cross-lingual code intelligence proposed by Zhu et al., consists of fine-grained parallel data from 7 PLs and English. This parallel data in a total of 8 languages allow XLCoST to support 10 cross-lingual code-related tasks, for example, program synthesis, code summarization, cross-language code retrieval, etc. Puri et al. [63] introduce a large-scale dataset CodeNet, aiming to benchmark a variety of critical coding tasks, including code similarity and classification, code translation between a large variety of PLs, and code performance (runtime and memory) improvement techniques. Although these efforts provide large-scale multi-task benchmark, the data for each of their tasks are extracted from the same collected dataset. Therefore, the data distribution between their different tasks is identical, and it is difficult for us to obtain rigorous and valid cross-task evaluation results by these benchmarks.

Table I compares our proposed CrossCodeBench with aforementioned benchmarks. We first compare whether each task in these benchmarks has a well-validated task description, definition, positive and negative examples, and whether it is off-the-shelf (i.e., the inputs and outputs are already preprocessed for each task). In addition, we compare their number of tasks, the number of task types (cf. Table III, where the first number

indicates the number of types in a total of 7 types, the next number indicates the types after considering sub-types, and “+1” indicates that there are task types other than those in Table III⁴), the number of datasets included, the number of programming languages (PLs), and the approximate number of total instances.

III. CROSSCODEBENCH

A. Collecting Tasks

In order to collect as many code-related datasets as possible, we hired 5 Ph.D. student (including one author of this paper) and 6 M.S. students (whose research area is software intelligence) to provide all datasets they have used. Their research span the areas of automated program repair, defect localization, code completion, code summarization, code generation, etc. Then, we recursively expand the scope of collecting datasets by reading surveys and based on paper relationships such as citations, related work, etc. In addition, we conduct an exhaustive search on a number of websites widely used to publish or collect datasets, such as GitHub, HuggingFace Datasets, PapersWithCode Datasets, etc. Finally, we end up collecting a total of 66 code-related datasets.

Since some datasets contain multiple subsets (e.g. CodeSearchNet [64] contains 6 subsets corresponding to different programming languages) or can support multiple tasks (e.g. the XLCoST dataset [62] can support translation between multiple parallel elements), and they do not correspond to the same task descriptions, we split these datasets into different tasks. As a result, we end up with 216 tasks. Moreover, for the task instances, we directly use the input and output in the form of text sequences provided by the dataset without any further processing. Since we are proposing a unified benchmark for all types of tasks, we formalize the inputs and outputs of all tasks in text-to-text form with reference to T5 [19]. In particular, for the classification task, we convert all labels to corresponding task-relevant text⁵.

B. Meta Information

Some basic information is already done when collecting datasets and tasks, such as the URL, BibTeX, and input/output languages. Next, in order to make our benchmark support multiple cross-task in-context learning methods, namely FSL, ZSL, and LTI, we need to pair each task with some corresponding meta information. In this section, we first present the schema of the meta information we need to complete, then illustrate the open coding procedure [65] that we follow to complete the meta information.

⁴CodeXGLUE and XLCoST have two retrieval tasks, namely natural language code search and code-to-code retrieval. In this work, we exclude the retrieval task type because our goal is to evaluate the cross-task capability of the unified model, however, retrieval tasks cannot be converted to a unified text-to-text form.

⁵For example, for some binary classification tasks, we convert the label “0” to text “No” and “1” to “Yes”.

1) *Schema*: With Sections II-A and II-B, we decide to match meta information for each task in the following fields:

- (a) **Type**: the type of task to which the task belongs.
- (b) **Description**: a short description of the task, e.g. “Translate English to French”, “Summarize”, etc.
- (c) **Definition**: detailed instruction of the task, including a description of the input, and how to map input to output.
- (d) **Positive/Negative Examples**: canonical positive or negative examples.

With these fields, our benchmark can support all current cross-task learning approaches. For FSL, we use task descriptions and positive examples; for ZSL, we only use task descriptions; and for LTI, we use task definitions and positive/negative examples.

2) *Coding Procedure*: Open coding procedure is a widely used standard analytical process that can be utilized to label a dataset [66]–[68]. We invite 1 Ph.D. student and 4 M.S. students mentioned in Section III-A to participate in the completion of the meta information, all of whom have more than 4 years of programming experience and have been working on code-related tasks for more than 2 years. As a prerequisite background knowledge, we ask all coders to first read the NLP-related work mentioned in Sections II-A and II-B, and to become familiar with the corresponding data as well. Below we show the steps of the coding procedure.

(a) *Pilot Study*. With the intent of defining the coding framework, one of the coders conduct a pilot study on randomly selected NLP 30 tasks with meta information, which are equally derived from the work of Brown et al. [18], Bach et al. [56], and Wang et al. [38]. The goal of this study is to identify the initial task type, analyze and identify the common pattern of the task description, definition, and the positive/negative examples. The work ends with the definition of four initial classifications, in addition to the experience and specifications for writing task descriptions (14 patterns), task definitions (5 patterns), and positive/negative examples. Here are details:

- Initial task type includes *classification, translation, generation, summarization* and *type prediction*.
- Task description is a short single sentence command that often appears before the input instance and connects it to the instructions. It briefly describes the intent of the task, usually with one or a few words. 14 initial patterns are identified, to name a few, *Translate A to B, Summarize A, Detect (defect/clone/variable misuse)*.
- Task definition is a detailed guide to solving the task and can be divided into three parts: (1) the description and interpretation of the input under the task, including the language and form of the input, e.g., *given a Java method with the method name masked by a special symbol '[MASK]', given a natural language description*; (2) how to convert the given input into output, and the language and form of the output, e.g., *translate the given Java method to Python function with the same functionality*, etc.; (3) others, including the search space of the output (for classification tasks, e.g., *if ..., outputs 'Yes', otherwise outputs 'No'*), and the format of the output

(e.g., for each identifier, outputs the name of the identifier and its type, separated by a colon ':'). 5 initial patterns are identified which connect these three elements into a complete, fluent, easy-to-understand natural language sentence.

- Four positive and four negative examples are required. Examples need to reflect the most typical situation in the task, explaining the main points of generating positive examples for positive examples, and pointing out the errors and giving the correct modifications for negative examples.

Based on this, this coder organizes a 60-minute session for the other four coders for training and discussion.

(b) Completion Procedure. Each coder is assigned to all 216 tasks. For each task, coders are asked to learn about the task by referring to the websites and paper through the URL and BibTeX, then observe the data for that task, and finally identify the task type and fill in the rest of the meta information as required.

Specifically, for task types, coders are allowed to identify a task as a new task type despite the fact that an initial task type already exists. In addition, it is encouraged to identify the task type while also giving the sub-types of the task under this task type, if possible. All other coders are informed when a new task type/sub-type is proposed, and coders will have a discussion for at least 10 minutes. If all the coders agree on the new task type/sub-type, then the type/sub-type list would be updated to include the new one.

As for the task description and definition, coders are asked to try to follow the 5/14 initial patterns for writing, and if this is not possible, adding new patterns is also allowed. An online list of writing pattern is shared between all coders, including both the initial schema and the newly added schema. In this way, the latest schema list is visible to all coders and open for them to use and add. It is worth noting that a new pattern can only be added to the list after it has been verified by all coders and the disagreement is solved by discussion.

In addition, for task types/sub-types, pattern lists of task descriptions and definitions, when adding new entries (i.e., new types/sub-types or task description/definition patterns), the similarity between the new entry and the existing entry is inspected. Similar entries are merged into a new, more general entry in due course, without disagreement, and the existing labels are updated accordingly as well. The process is fully iterative and includes continuous refinement of the entry and discussion of ambiguous cases. Each decision made during the entry extraction process represents the opinion of all coders.

Lastly, when it comes to positive/negative examples we are not able to follow the procedure described above since they are all free-form labeled content. We first ask all coders to independently complete writing the task description, task definition, and 2 positive/negative examples. After finishing, we gather all coders to discuss each task, and for the task description and definition, we put together what all coders had written, vote on the most concise and clear version, and revised it until everyone is satisfied. For the positive and negative examples, we first put together a total of 10

examples they written, vote on the most representative 4 until no disagreement remains.

(c) Agreement Measurement. To evaluate the validity and reliability of our coding procedure, we use Cohen’s Kappa value [69] to measure the agreement among all coders. The results are 85.4%, 76.7%, 73.2% and 70.9% for task type, task description, task definition, selection of positive and negative examples, respectively. This shows that all coders have “almost perfect agreement” on the task type and “substantial agreement” on the other three meta information [70].

For fields that still have disagreements, we resolve them by applying a third person solution. We assign conflicting reports to 4 external coders (the 5 Ph.D. students mentioned in Section III-A except the one internal coder) and let them judge and resolve these disagreements. Our analysis shows that the disagreements mainly focus on the second part of the task definition, where different coders have different understandings and representations of the input-to-output mapping, as well as positive and negative examples.

Finally, in order to ensure the resulting meta information is sufficient for average software engineering researchers. We invite five undergraduate software engineering students with programming experience in non-software intelligence fields to read the meta information of all tasks. They are then asked to review the definition, intent, input and output formats, and other information for each task. Where they are unclear or incorrectly stated, we will feed back to the 4 external coders mentioned in the previous paragraph for revision until the task can be clearly defined and understood.

3) *Summary:* In the end, we get 7 task types as follows.

- **Classification:** output the corresponding labels based on the input. It is further divided into two sub-types, Binary and Multi-label, which correspond to tasks with only two labels and more than two labels, respectively.

- **Fill in the blank:** predict the missing token or sequence in a given input.

- **Translation:** translate code snippets written in one language into another preserving semantics and functionality.

- **Generation:** generate a sequence based on the input. Three sub-types are (1) Rewrite: modify a part of the given code sequence and output the modified version; (2) Text-to-Code: input the natural language description and output the corresponding code sequence; (3) Code-to-Text: input as code, output the required natural language sequence as required.

- **Summarization:** given a piece of code, output the functional description of that piece of code⁶.

- **Type Prediction:** predict the type of all identifiers in a given code snippet, which is a kind of sequence tagging tasks.

- **Question Answering:** given a piece of code and a natural language question, output the answer to that question.

Table II provides statistical information about Cross-CodeBench. In addition, Table III shows the statistics of the

⁶We separate this task from the Code-to-Text sub-type in Generation because this type of task is in a very important position in software intelligence research. And its research approach is different from other Code-to-Text approaches (e.g., commit message generation)

TABLE II
STATISTICS OF CROSSCODEBENCH

Field	Average
# of instances per task	250032.48
Task description length	3.52
Task definition length	26.39
Positive example input length	21.42
Positive example output length	15.77
Negative example input length	21.34
Negative example output length	15.48

TABLE III
TASK TYPES AND THEIR STATISTICS

Type	# of Tasks	# of Categories	# of Instances
Classification	21	15	6,565,637
- Binary	17	13	6,430,804
- Multi-label	4	2	134,833
Fill in the Blank	10	3	13,413,614
Translation	94	1	2,366,970
Generation	53	5	19,511,282
- Rewrite	10	2	3,197,080
- Text-to-Code	41	2	15,658,171
- Code-to-Text	2	1	656,031
Summarization	34	2	11,186,611
Type Prediction	2	1	773,038
Question Answering	2	1	189,863
Total	216	28	54,007,015

final task types. In total, the benchmark includes 216 tasks, 28 task categories and over 54M instances.

We define a task **category** as the set of tasks that have the same intent. We define a **task** as a $\langle \text{task category}, \text{dataset} \rangle$ pair, and a task **type** depends on the input and output of the task, not on the intent. A task category is a superset of a task. But a task category is equal to a task if and only if the task category has only one dataset. A task category is a subset of a task type (e.g. Bug Fixing [71] and Mutant Generation [72] categories are both Rewrite sub-types), they are equal if and only if the task type has only one task category, e.g. the Translation type has only one category, Code Translation.

4) *Demonstration*: In our benchmark, each task consists of two json files, one containing the meta information and the other containing the data instances. Both files have the same name in the first part, i.e., “task_{task_id}_{dataset_name}_{task_type}”, followed by “.meta.json” and “.data.json”, respectively.

To better illustrate, Figure 1 shows the contents of a json file with the meta information of a task. We can see that in addition to the fields we mentioned in Section III-B1, we also include some information about the task/dataset, such as the language of input/output, BibTeX and URL of the dataset, etc.

C. Splits

After preparing the data, we need to create different training/evaluation splits to evaluate the cross-task learning ability of the model in different application scenarios and difficulties. Table IV lists the 10 training/evaluation splits we create and

task_006_swapped_operands_classification.meta.json

```

“Type”: [“Classification”, “Binary”],
“Description”: “Detect swapped operands”,
“Definition”: “You are given a function, your task is to identify whether
the operands of non-commutative binary operators are swapped. Construct
an answer that is ‘Swapped operands’ if such a swap occurs and ‘Correct’
otherwise.”,
“Input Language”: “Programming Language -> Python”,
“Output Language”: “Natural Language -> English”,
“Positive Examples”: [{
  “Input”: “def __contains__(self, x): return x in self.columns”,
  “Output”: “Correct”,
  ...
}],
“Negative Examples”: [...],
“BibTeX”: ...,
“URL”: ...,
...

```

Fig. 1. An example of the task meta information.

TABLE IV
TASK SPLITS AND THEIR STATISTICS

Cross Type	Scope	Split Name	# of Tasks	# of Instances Train/Eval
Cross-Category	Intra-Type	Cat-Intra-CD	20/1	4.83M/1.73K
		Cat-Intra-BF	44/9	16.4M/3.1M
	Inter-Type	Cat-Inter-CD	215/1	52.3M/1.73M
		Cat-Inter-BF	207/9	50.9M/3.1M
Cross-Sub-Type	Intra-Type	Sub-Intra-ML	17/4	6.4M/135K
		Sub-Intra-C2T	51/2	18.9M/656K
	Inter-Type	Sub-Inter-ML	212/4	53.9M/135K
		Sub-Inter-C2T	214/2	53.3M/656K
Cross-Type	Inter-Type	Type-Trans	122/94	51.6M/2.4M
		Type-QA	214/2	53.8M/190K

use in this paper. We classify them in two dimensions: the level of the cross-task, and the scope of the training task. In addition, we list the name of each split, the number of tasks and instances of the training/evaluation split in Table IV. Next, we will introduce these splits in order from 3 cross-task levels, namely, cross-category, cross-sub-type and cross-type.

1) *Cross-Category*: Define new task category is common in SE, but collecting and labeling a large amount of data for the newly defined task category can be very labor-intensive, even impossible. Instead, we can define the task category exactly and give several examples. At this point, we can consider using cross-task learning methods, where learning on a large number of existing tasks acquires generalization capabilities on this new task. Therefore, we first wish to explore whether, and to what extent, current cross-task learning methods will make models achieve cross task category ability. We choose a binary classification task, i.e., clone detection [73] and a rewrite generation task, bug fixing [71] as evaluation task categories, respectively. In addition, since task categories are a subset of task types, we also want to explore whether more data from different task types can help improve the cross-task performance of the model on one task category.

Therefore, for each of the two evaluation sets, we use two training sets for training: (1) the training set includes only tasks of the same task type as the evaluation task in the other

tasks (Intra-Type); (2) the training set includes all other tasks (Inter-Type). Finally, we create 4 splits for the cross-category level, namely Cat-Intra-CD, Cat-Intra-BF, Cat-Inter-CD, and Cat-Inter-BF.

2) *Cross-Sub-Type*: Recall that we also have sub-types for Classification and Generation types in Table III. We also want to explore to what extent the model can learn cross-task capabilities at the level of sub-type. Therefore, we select two sub-types for each of them, namely Multi-label Classification and Rewrite Generation. For multi-label classification sub-type, we want to investigate whether the model can generalize to multi-label tasks by learning only binary classification sub-type (and other types of tasks). For rewriting tasks, we wish to explore how much rewriting skill the model can acquire without having learned any rewriting task. As with the setup in the cross-category, we create 2 splits for each of the two selected sub-types under the different scopes of the training set. Consequently, we create 4 splits as well, i.e., Sub-Intra-ML, Sub-Intra-C2T, Sub-Inter-ML, and Sub-Inter-C2T.

3) *Cross-Type*: The last is the most challenging setup, the cross-type, where we use all tasks in the entire task type as the evaluation set and the other types of tasks as the training set. First, we want to use the Translation type as the evaluation set, and we want to explore whether the model can learn the correspondence between different languages by learning from other types of tasks. Second, we choose a type that contains the fewest tasks, i.e., Question Answering. This split is closest to the real practice situation, i.e., by learning on a well-resourced task type, we expect the model to have a better generalization ability on the new task type without resources. Therefore, we create two splits at the cross-type level, namely Type-Trans and Type-QA.

By now, we create 10 training/evaluation splits of different difficulty levels and different application scenarios. These splits are used in the experiments. Besides, we are able to ensure that there is no data leakage in these splits, in another word, there is no task in the training set with the same resource dataset as any task in the evaluation set.

IV. EXPERIMENTAL SETUP

In this section we describe the setup for preliminary experiments on CrossCodeBench, including cross-task learning methods, baselines, evaluation metrics and other settings.

A. Cross Task Learning Methods

Here we introduce the in-context cross-task learning method that we use for our experiments. In-context learning methods also determine the composition of the inputs to the model in our experiments. By referring to Section II-A, II-B and the meta information we introduced in Section III-B, we propose the following learning methods.

- *k-shot Learning (k -shot)* : recall that the input of the few-shot learning consists of three parts: task description, k examples and a prompt. Since each task in CrossCodeBench contains 4 positive examples, we randomly select $k(= 4, 3, 2, 1)$

positive examples as the examples of the input context, resulting four methods named **4-shot**, **3-shot**, **2-shot** and **1-shot**.

- **Zero-shot learning ($zero$ -shot)**: same as k -shot, but without the second part of the input, i.e. the example.

- **Learning from instructions with m/n positive/negative examples (m/n -instruct)**: use instructions items introduced in Wang et al. [38], with m/n positive/negative examples. Specifically, we select the two combinations that worked best shown by Wang et al. [38] (1) **2/0-instruct**: $m = 2, n = 0$ when training and $m = 4$ when evaluating; (2) **2/2-instruct**: $m = 2$ and $n = 2$ for both training and evaluating.

Finally, we derive 7 cross-task learning methods that will be used as input to the baseline model to be presented below.

B. Models

1) *Shortcut Methods*: The feasibility of a proposed benchmark is important for the subsequent work, which determines whether the benchmark is practical and meaningful to study. Specifically for CrossCodeBench, it is whether the model can actually gain cross-task capability by using a particular cross-task learning method, rather than just behaving like it has gained such capability through some easy shortcuts. Therefore, we propose two such shortcut methods for each split, (1) **Copy Ex-output**: randomly copying the output of one of the four positive examples and (2) **Copy Ins-input**: copying the input of the current instance. We use these two shortcut methods to evaluate the feasibility of CrossCodeBench, i.e., doing cross-task research in the field of software intelligence.

2) *Off-the-Shelf Models*: Off-the-shelf models are those that are evaluated directly on the evaluation set, without any further fine-tuning. We first wish to choose a model that is a “few-shot learner”, such as GPT-3 [18]. However, our request to use the GPT-3 API has not been approved. Instead, we turn to **Tk-Instruct** [38], a T5-based model that is already trained to follow general language task instructions (including a portion of code-related tasks). Tk-Instruct is shown to have better cross-task performance than few-shot learners, such as GPT-3, and other task instruction learners, such as InstructGPT [58]. Tk-Instruct is only used under the m/n -Instruct learning settings. Specifically, we use the 3B parameter version of the model, and for the learning method in the 2/0-instruct and 2/2-instruct settings, we use the “Tk-Instruct-3b-def-pos” and “Tk-Instruct-3b-def-pos-neg” versions of the model, respectively, in order to achieve the best performance [38]. We use the Tk-Instruct model without any fine-tuning and directly evaluate them on the evaluation set, under the different learning methods.

3) *Fine-tuned Models*: We evaluate pre-trained models that are fine-tuned on the training set with aforementioned in-context learning method as well. First, we choose two recent and widely used pre-trained models of source code, **PLBART** [8] and **CodeT5** [9]. In particular, PLBART is a sequence-to-sequence model based on BART [74] and pre-trained on an extensive collection of Java and Python functions and associated NL text. CodeT5 is a T5 [19]-based model that

is pre-trained on a large corpus containing 8 programming languages and natural language. We ensure that neither PLBART nor CodeT5 had supervised training on any evaluation task in any splits. We also ensure that there is no overlap between the dataset they use in pre-training and the source dataset of the evaluation set in all splits. Besides, we further fine-tune **Tk-Instruct** [38] by using the m/n -instruct learning method. We hope that further fine-tuning will help Tk-Instruct to adapt to code-specific tasks while understanding generic task instruction. Consequently, we obtain three fine-tuned models, two pre-trained models of source code, PLBART and CodeT5, which will be fine-tuned under all cross-task learning methods. There is also an instruction learner, Tk-Instruct, which will be fine-tuned under only two m/n -instruct methods. In order to balance efficiency and effectiveness, we use the official “large” version of all three models.

4) *Supervised Models*: We estimate an upper bound performance of each split by supervised fine-tuning a CodeT5-large model on all task instances (except instances used for evaluation) on all evaluation set. This approach follows the classical “pre-train then fine-tune” paradigm, where we use all evaluation tasks in the fine-tuning and do not apply any cross-task learning methods. Since this approach allows the model to see the data of the target task and perform supervised fine-tuning, this approach is theoretically the upper limit of the cross-task approach on the corresponding split.

Finally, we end up with 2 shortcut methods, 2 off-the-shelf models, 3 fine-tuned models and 1 supervised model. All model checkpoints are loaded from the official models published on HuggingFace Hub.

C. Metrics

Since outputs of all tasks are in text form (see Section III-A), so text-specific metrics are used in our experiments. Specifically, for classification tasks whose output is short and limited, we adopt **Exact Match (EM)** to measure the ratio of the instances for which a model produces exactly the same string as the gold labels. For the other tasks, their output is a longer sequence, so we employ **BLEU (B.)** [75] and **ROUGE-L (R.L)** [59]. Both are widely adopted string overlap metrics that measure the similarity of between the text sequence generated by the model and the gold sequence. For all metrics, we report scores under percentages.

D. Others

In order to avoid the data imbalance problem, we limit the number of instances in each training tasks to 10,000. These 10,000 instances are fixed across different running if the number of instances for a task exceeds 10,000. Similarly, to avoid the evaluation results being unevenly affected by the amount of task data and to make the evaluation more efficient, we select a fixed number of 500 evaluation instances from each task.

All experiments are conducted on 4×NVIDIA Tesla V100 32Gs with a total fine-tuning epochs of 3. We tune hyperparameters using grid search. We select learning rate from {1e-5, 3e-5, 5e-5, 1e-4}, warm-up steps from {500, 1000, 2000},

and batch size per device from {4, 8, 16}. As a result, for the fine-tuning phase, we use an initial learning rate of 5e-5 for CodeT5 and PLBART, and 1e-5 for Tk-Instruct, with a batch size per device of 16 and 1000 warmup steps for all. We run each experiments three times using different random seeds and report the mean.

V. RESULTS AND DISCUSSION

In this section, we present and discuss the preliminary experimental results on CrossCodeBench.

A. Overall Results

Table V shows the overall benchmarking results. Except scores of the supervised model, the best scores are bolded and those within 5% below the best scores are underlined, and if there are no scores within 5% below, the bolded best scores are further underlined. Based on these results, we have the following observations and discussions.

1) *Room for improvement exists*: First of all, we can see that supervised fine-tuned CodeT5 following the classical “pre-train then fine-tune” paradigm has very significant advantages over any other methods. Since supervised CodeT5 performs supervised training on the target task, it can be seen as an upper bound for the performance of large language models of source code on the corresponding evaluation set. This suggests that there is still a lot of room for theoretical improvement in cross-task learning based on pre-trained models of source code.

2) *Shortcut wins on low-resource classification*: We find that the best fine-tuned model outperforms shortcut methods on all metrics for all splits except Cat-Intra-CD and Sub-Intra-ML. This illustrates that the cross-task learning approach allows the models to gain real cross-task capability, not just some simple shortcuts. And on the two splits where the shortcut method wins, we discover two points, one is that scopes of the training set of both splits are both “Intra”, and the second is that only the evaluation sets of these two splits contain the classification tasks.

For the first point, recall that in Section III-C, we create splits of “Intra” and “Inter” by changing the scope of tasks in the training set, the former restricts the tasks in the training set to the same type as the verification tasks, while the latter has all other tasks in the training set. On the two “Inter” corresponding to these two “Intra”, the fine-tuned method outperforms the shortcut method instead. Therefore, we can conclude that having more data for other types of tasks could improve the cross-task performance on tasks of a certain type.

As for the second point, we believe that the main reason is that the output of these tasks is short (usually one word) and fixed (corresponding to all labels), and our experiments do not restrict these conditions, relying only on giving some “soft” hints in the task definition and in the output of the examples. For example, in the classification task shown in Figure 1, we illustrate the restriction/range of output only in two places, (1) “Definition”: the output are “Swapped Operands” in the case

TABLE V
OVERALL BENCHMARKING RESULTS

Type	Models	Methods	Cross-Category						Cross-Sub-Type						Cross-Type			
			Intra			Inter			Intra			Inter			Trans		QA	
			CD	BF		CD	BF		ML	C2T		ML	C2T		B.	R.L	B.	R.L
EM	B.	R.L	EM	B.	R.L	EM	B.	R.L	EM	B.	R.L	B.	R.L	B.	R.L			
Shortcut	Copy Ex-output		50.00	10.35	20.24	50.00	10.35	20.24	10.75	2.52	1.43	10.75	2.52	1.43	3.97	15.95	3.29	2.71
		Copy Ins-input	0.00	40.51	54.59	0.00	40.51	54.59	0.00	1.28	3.55	0.00	1.28	3.55	45.44	60.67	0.82	2.00
Off-the-Shelf	Tk-Instruct	2/0-instruct	2.96	29.91	33.23	2.96	29.91	33.23	0.12	3.58	3.97	0.12	3.58	3.97	17.31	29.87	9.11	8.94
		2/2-instruct	2.22	29.65	32.45	2.22	29.65	32.45	0.11	3.66	3.42	0.11	3.66	3.42	16.46	29.27	8.10	8.38
Fine-Tuned	PLBART	4-shot	15.86	44.76	57.52	25.54	55.86	67.89	1.75	7.92	6.61	8.54	11.95	12.84	50.14	63.84	12.37	14.61
		3-shot	12.94	43.60	56.40	23.49	54.38	65.21	1.23	7.58	6.65	5.82	11.61	11.59	47.90	62.35	11.86	13.59
		2-shot	10.59	40.84	54.37	22.01	49.73	64.79	0.89	5.44	5.54	3.74	10.73	11.12	47.60	60.76	9.61	11.70
		1-shot	9.27	36.45	46.90	21.95	45.95	54.78	0.56	5.86	4.03	3.12	10.51	10.78	46.66	59.83	8.84	9.86
		zero-shot	7.65	45.36	44.08	16.40	52.04	50.34	0.03	3.66	3.61	2.81	7.63	6.87	41.48	54.31	6.99	8.33
		2/0-instruct	19.11	<u>47.74</u>	<u>59.23</u>	53.25	<u>59.19</u>	70.55	4.82	8.18	7.22	14.66	14.31	<u>15.43</u>	48.88	<u>65.70</u>	15.72	18.73
	CodeT5	2/2-instruct	19.86	42.84	57.74	<u>53.87</u>	<u>58.80</u>	68.89	4.76	8.96	7.82	14.82	14.16	14.24	47.43	63.95	15.37	17.76
		4-shot	15.80	<u>46.89</u>	<u>58.15</u>	30.03	<u>56.92</u>	72.01	1.80	7.38	6.80	8.30	12.15	12.04	49.22	<u>66.26</u>	14.78	15.12
		3-shot	13.05	44.67	57.33	29.55	54.30	69.41	1.04	7.84	6.59	6.83	11.49	10.06	47.83	64.35	14.27	14.99
		2-shot	12.49	40.94	52.68	29.53	51.26	62.70	0.81	6.88	5.85	3.47	10.50	10.46	48.65	<u>65.25</u>	13.69	14.73
		1-shot	10.68	35.32	45.38	28.75	48.86	58.14	0.34	6.37	5.57	2.78	9.87	<u>8.47</u>	45.72	62.92	11.41	12.25
		zero-shot	7.12	30.69	43.07	18.8	45.85	56.51	0.05	4.15	3.72	1.95	6.44	6.99	43.24	58.38	9.12	10.95
Tk-Instruct	2/0-instruct	20.89	48.33	61.02	<u>56.33</u>	61.89	76.89	5.57	8.35	7.10	15.95	<u>15.40</u>	<u>15.63</u>	52.05	67.87	17.95	<u>20.52</u>	
	2/2-instruct	20.05	45.57	<u>59.31</u>	56.42	58.49	<u>74.26</u>	5.06	8.05	6.93	15.10	14.60	15.90	<u>51.51</u>	<u>66.61</u>	<u>17.31</u>	<u>20.12</u>	
Supervised	CodeT5	2/0-instruct	18.43	43.86	57.23	51.88	57.45	70.98	4.92	9.18	8.47	<u>15.64</u>	15.60	14.35	45.49	62.14	17.04	20.67
		2/2-instruct	18.52	42.01	56.16	51.20	57.67	70.39	4.24	<u>9.10</u>	7.88	14.43	<u>15.01</u>	14.77	45.63	62.41	16.15	19.33
Supervised	CodeT5	None	78.49	73.41	82.40	78.49	73.41	82.40	63.23	28.11	27.86	63.23	28.11	27.86	72.26	85.83	37.46	39.49

of operands swapping only in definition, and “Correct” otherwise, (2) “Positive/Negative Examples”. These restrictions are simply entered into the model as plain text, and thus cannot pose any substantial restrictions on what the model generates. In this case, we can only expect the model to be trained to acquire the knowledge of how to find the output range on each classification task in the input. But the results show that this is far from the case, and the model often generates text outside the output range, leading to a low performance. This conclusion is further corroborated by browsing the output of the fine-tuned model on the classification task.

The performance of the shortcut method on the classification task suggests that further restrictions on the model’s output on the classification task are needed if one wants a unified text-to-text model to have excellent cross-task performance on the classification task.

3) *Further fine-tuning matters*: It is clear that the fine-tuned model has a significant performance advantage over the off-the-shelf model. This is because fine-tuning on our training set can help the model understand the input representation of cross-task learning methods (for PLBART and CodeT5), or become more familiar with tasks in the software intelligence domain (for Tk-Instruct). Even though Tk-Instruct has included several code-related tasks, such as “code-to-text”, in its original training set [38], we believe that these tasks are diluted in a huge number of NLP tasks, making off-the-shelf Tk-Instruct unable to obtain good power in code-related tasks.

4) *Domain-specific models are better*: We also notice that CodeT5, a pre-trained model of source code, is able to achieve better results than Tk-Instruct, a model trained to follow task instructions, in most tasks. This suggests that domain-specific pre-trained models are generally a better choice than cross-domain models. Moreover, we also notice that Tk-Instruct wins in the task of generating natural language, such as Code-to-Text and Question Answering. This is not difficult

to understand, because training on huge amount of natural language corpus helps it generate more fluent and reasonable natural language text.

5) *Detailed in-context information generally helps*: Among all cross-task learning methods, LTI produces the best results for all tasks and metrics. Specifically, in most cases, using only two positive examples is more effective than using two positive and two negative examples. By comparing the input representations of FSL, ZSL and LTI, we can see that compared to FSL and ZSL, LTI extends the short task description to a detailed task definition, which allows the model to understand the task more comprehensively, including the input/output format, the search space of outputs, etc., and thus learn the solution to the task. In addition, LSI further introduces negative examples, but it appears that the extra negative examples do not contribute positively to the performance of the model in most cases. This also confirms the results obtained by Wang et al. [38].

B. Data Scaling

In addition to more data on other task types, we also explore whether the size of the training instances could affect the cross-task performance of the model given the same tasks. Therefore, on Type-Trans, we vary the number of instances per task that are used for fine-tuning CodeT5, and the evaluation results are shown in Figure 2.

It can be seen that when the number of tasks in the training set is constant, increasing the maximum number of instances per task within a certain range can significantly improve the cross-task performance of the model. However, when the maximum number of instances per task increases from 10,000 to 20,000, the BLEU shows a significant decrease, so for the BLEU metric, the model achieves its best performance at 10,000. On the other hand, for ROUGE-L, although the average value has been increasing all the time, if the range of data errors is taken into account, the model may start to




Fig. 2. Scaling trends of models performance as the maximum number of instances per task changes from 500 to 20,000. The error is indicated by light shading. The x-axis is on a log-scale in order to be more intuitive.




Fig. 3. Scaling trends of models performance as the number of model parameter changes. The error is indicated by light shading. The x-axis is on a log-scale in order to be more intuitive.

show a decrease in performance when increasing from 5,000 to 10,000. Thus, in aggregate, the model may achieve its best performance at a maximum number of instances of 10,000 per task, which, if continued to increase, leads to time and space costs that do not match the gains in model performance.

C. Model Scaling

We also study the effect of model scaling by initializing CodeT5 and Tk-Instruct from different sizes of checkpoints and the results are in Figure 3.

We find that increasing the model size consistently delivers an improvement in the cross-task performance of the model over the range of model sizes we experimented with, and is roughly log-linear with parameter size. Combining the results in Figure 2, we can see that models of a certain size have a fixed demand on the amount of data. For example, CodeT5-large reaches a performance bottleneck at a maximum of 10k instances per task. Conversely, a larger model always leads to better performance when the data scale is of a certain size (note that the data volume should be large enough to avoid overfitting).

TABLE VI
CROSS-TASK PERFORMANCE OF THE MODEL ON DIFFERENT SCOPE OF FINE-TUNING DATA

#	Fine-tuning data comes from			Training size	BLEU	ROUGE-L
	In-Sub-Type	In-Type	Out-Type			
1	✓	✗	✗	116K	43.41	55.82
2	✓	✓	✗	16.4M	48.33	61.02
3	✓	✓	✓	50.9M	61.89	76.89
4	✗	✓	✗	16.3M	28.22	43.47
5	✗	✓	✓	50.8M	57.13	74.82
6	✗	✗	✓	34.5M	53.74	71.63

D. Scope Scaling

Finally, we investigate a problem with practical applications, i.e., the case when we propose a new task class and obtaining the corresponding data is very costly or impossible. Considering the use of cross-task learning is a feasible approach at this point. Thus, there is a question, how the scope of the training set is selected is important for the cross-task capability of the model. Therefore, we select a target task category, Bug Fixing [71], and then change various ranges of training sets to investigate under which scope the model would achieve the best cross-task performance for the Bug Fixing task category.

We also present two splits in Section III-C1, with experiments on two application scenarios where the target task is Bug Fixing. The training sets in both splits in Section III-C1 include the sub-type (Rewrite) or type (Generation), to which Bug Fixing belongs. However, in practical, the proposed task category may be a new sub-type/type, and other task data of the same sub-type and type is not available. What we further study here is whether in this case, the more data is still better when we can only select tasks of other sub-types or types.

After selecting Bug Fixing as the target task category, we change the scope of the training set data from three dimensions: (1) In-Sub-Type: whether the training set contains task data of the same task sub-type as the target task, i.e., Rewrite; (2) In-Type: whether the training set has data of the same task type as the target task, i.e. Generation; (3) Out-Type: whether there are data in the training set that do not belong to the same task type as the target task, i.e., the other 6 types except the Generation type. We investigate the data in the training set under various combinations in these three dimensions and present the results in Table VI.

Comparing experiments 2 and 4, or 3 and 5 in Table VI, we learn that the data within the sub-type plays an important role. But such a conclusion holds only when other task data are available. If we compare experiment 1 with 5 and 6, it can be found that the performance achieved by using only in-sub-type data is not as good as that achieved by using data outside the sub-type.

Besides, in comparison with experiments 2 and 6, we can find that the performance using in-type data is worse than that using out-type data only. We believe the reason is that the amount of out-type data is larger than that of in-type data, which allows the model to learn more cross-task capabilities

with more in-context information.

Therefore, we can conclude that the amount of data still plays an important role. Its contribution to the cross-task performance of the model is more important than the data for the same type of task.

E. Implications

After our experiments and discussions above, we have some findings and implications that can facilitate subsequent researchers and practitioners.

The first thing we can see is that all models performs very poorly on the classification task in the cross-task setting. After adding the instruction containing the set of classification labels to the input, the performance of the model on the classification task, though improved, is still poor because it still often outputs words outside the output range. So, how to constrain the model to output valid classification labels requires further investigation.

In addition, our results show that the different combinations of SE tasks used for training a model would yield different performances on a target task. So, when given a new SE task that lacks data, how to quickly determine which existing tasks should be used to achieve the best results is an urgent issue to explore.

Finally, we learn from Sections V-B–V-D that more data, larger model and more tasks can all benefit the cross-task generalization ability. Therefore, if there are limited resources, is it better to devote it to a larger model or to more data, and if to data, is it better to increase the data per task or to increase the diversity of tasks. That is, how to make better use of the limited resources is also an issue worth investigating.

VI. THREATS TO VALIDITY

Internal Validity: Threats to internal validity relate to the bias introduced when manually annotating some fields in meta information of the task. To mitigate this, we start with the open coding procedure to reduce the impact of individual bias. Then for remaining conflicts, we address these by employing a third person solution (see Section III-B2 for details). We believe that the approach we adopt minimizes the human impact on the meta information, but it still may not be completely avoidable. Another factor for threatening internal validity is hyperparameters. Due to time constraints, we only perform a coarse-grained hyperparameter search. Therefore, other hyperparameter settings may lead to better results. But this does not affect the observations and conclusions presented in this paper, since all experiments are carried out with the same hyperparameter searching strategy.

External Validity: Threats to external validity concern about the data imbalanced among different types of tasks. Even though we limit the maximum number of instances per task, the number of tasks contained in different task categories, types varies greatly (see Table III). Therefore, the model may have different generalization performance on different types of unseen tasks. In addition, since we need to apply tasks to the unified sequence-to-sequence model, all tasks in

our benchmark are formalized in a text-to-text form. Those code-related tasks that cannot be converted to this form are excluded. An example is the retrieval tasks, such as natural language code search and code-to-code retrieval [11]. The retrieval task requires computing the similarity between the representation vectors generated by the model for different inputs and cannot be applied to the text-to-text framework.

Construct Validity: The major threat is the data or task overlap between the evaluation set and the data/task the model has seen during self-supervised pre-training and fine-tuning on the training set. First, since the minimum cross-task granularity of our proposed split is task categories, there is no overlap of task categories between the training and evaluation sets. Then, to mitigate such data overlap threats, we use the tool provided by Allamanis [76] to ensure that the data in the evaluation set of the proposed 10 splits do not have duplication with the data in the training set during our experiments. Therefore, we can guarantee that no overlap problem will be introduced in the fine-tuning phase, but since we use pre-trained models as the backbone of our baselines, such threats may be introduced in the pre-training phase.

VII. CONCLUSION

To study the problem of cross-task generalization in software intelligence, we collect a large-scale meta-dataset containing various types of tasks and label each task with rich meta information to support various cross-task learning methods. We then create a benchmark, CrossCodeBench, for studying the generalization ability of source code-related deep learning models. Through preliminary experiments, we demonstrate and analyze the feasibility and possible research directions for cross-task studies of software intelligence.

ACKNOWLEDGMENT

This work was supported by National Natural Science Foundation of China (61802167), Cooperation Fund of Huawei-NJU Creative Laboratory for the Next Programming, and NSF award 2034508. We also thank the reviewers for their helpful comments. Chuanyi Li is the corresponding author.

REFERENCES

- [1] C. Niu, C. Li, B. Luo, and V. Ng, “Deep learning meets software engineering: A survey on pre-trained models of source code,” in *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI 2022*, 2022, pp. 5546–5555.
- [2] N. T. De Sousa and W. Hasselbring, “Javabert: Training a transformer-based model for the java programming language,” in *2021 36th IEEE/ACM International Conference on Automated Software Engineering Workshops (ASEW)*. IEEE, 2021, pp. 90–95.
- [3] N. Chirkova and S. Troshin, “Empirical study of transformers for source code,” in *Proceedings of the 29th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering*, 2021, pp. 703–715.
- [4] A. Karmakar and R. Robbes, “What do pre-trained code models know about code?” in *2021 36th IEEE/ACM International Conference on Automated Software Engineering (ASE)*. IEEE, 2021, pp. 1332–1336.
- [5] Z. Feng, D. Guo, D. Tang, N. Duan, X. Feng, M. Gong, L. Shou, B. Qin, T. Liu, D. Jiang, and M. Zhou, “Codebert: A pre-trained model for programming and natural languages,” in *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: Findings*, 2020, pp. 1536–1547.

- [6] D. Guo, S. Ren, S. Lu, Z. Feng, D. Tang, S. LIU, L. Zhou, N. Duan, A. Svyatkovskiy, S. Fu, M. Tufano, S. K. Deng, C. Clement, D. Drain, N. Sundaresan, J. Yin, D. Jiang, and M. Zhou, "Graphcodebert: Pre-training code representations with data flow," in *International Conference on Learning Representations*, 2021.
- [7] A. Mastropaolo, S. Scalabrino, N. Cooper, D. N. Palacio, D. Poshyanyk, R. Oliveto, and G. Bavota, "Studying the usage of text-to-text transfer transformer to support code-related tasks," in *2021 IEEE/ACM 43rd International Conference on Software Engineering (ICSE)*. IEEE, 2021, pp. 336–347.
- [8] W. Ahmad, S. Chakraborty, B. Ray, and K.-W. Chang, "Unified pre-training for program understanding and generation," in *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 2021, pp. 2655–2668.
- [9] Y. Wang, W. Wang, S. Joty, and S. C. Hoi, "Codet5: Identifier-aware unified pre-trained encoder-decoder models for code understanding and generation," in *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, 2021, pp. 8696–8708.
- [10] C. Niu, C. Li, V. Ng, J. Ge, L. Huang, and B. Luo, "Spt-code: Sequence-to-sequence pre-training for learning source code representations," in *2022 IEEE/ACM 44th International Conference on Software Engineering (ICSE)*, 2022, pp. 01–13.
- [11] D. Guo, S. Lu, N. Duan, Y. Wang, M. Zhou, and J. Yin, "Unixcoder: Unified cross-modal pre-training for code representation," in *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 2022, pp. 7212–7225.
- [12] A. D. Gotmare, J. Li, S. Joty, and S. C. Hoi, "Cascaded fast and slow models for efficient semantic code search," *arXiv preprint arXiv:2110.07811*, 2021.
- [13] D. Wang, Z. Jia, S. Li, Y. Yu, Y. Xiong, W. Dong, and X. Liao, "Bridging pre-trained models and downstream tasks for source code understanding," in *2022 IEEE/ACM 44th International Conference on Software Engineering (ICSE)*. IEEE, 2022, pp. 287–298.
- [14] A. Radford, J. Wu, R. Child, D. Luan, D. Amodei, and I. Sutskever, "Language models are unsupervised multitask learners," *OpenAI blog*, 2019.
- [15] M. Yu, X. Guo, J. Yi, S. Chang, S. Potdar, Y. Cheng, G. Tesauero, H. Wang, and B. Zhou, "Diverse few-shot text classification with multiple metrics," in *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, 2018, pp. 1206–1215.
- [16] T. Gao, A. Fisch, and D. Chen, "Making pre-trained language models better few-shot learners," in *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, 2021, pp. 3816–3830.
- [17] O. Vinyals, C. Blundell, T. Lillicrap, D. Wierstra *et al.*, "Matching networks for one shot learning," *Advances in neural information processing systems*, vol. 29, 2016.
- [18] T. Brown, B. Mann, N. Ryder, M. Subbiah, J. D. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell *et al.*, "Language models are few-shot learners," *Advances in neural information processing systems*, vol. 33, pp. 1877–1901, 2020.
- [19] C. Raffel, N. Shazeer, A. Roberts, K. Lee, S. Narang, M. Matena, Y. Zhou, W. Li, and P. J. Liu, "Exploring the limits of transfer learning with a unified text-to-text transformer," *Journal of Machine Learning Research*, vol. 21, pp. 1–67, 2020.
- [20] Q. Ye, B. Y. Lin, and X. Ren, "Crossfit: A few-shot learning challenge for cross-task generalization in nlp," in *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, 2021, pp. 7163–7189.
- [21] A. Efrat and O. Levy, "The turking test: Can language models understand instructions?" *arXiv preprint arXiv:2010.11982*, 2020.
- [22] O. Weller, N. Lourie, M. Gardner, and M. E. Peters, "Learning from task descriptions," in *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing*, 2020, pp. 1361–1375.
- [23] S. Mishra, D. Khashabi, C. Baral, and H. Hajishirzi, "Cross-task generalization via natural language crowdsourcing instructions," in *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 2022, pp. 3470–3487.
- [24] A. Aghajanyan, A. Gupta, A. Shrivastava, X. Chen, L. Zettlemoyer, and S. Gupta, "Muppet: Massive multi-task representations with pre-finetuning," in *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, 2021, pp. 5799–5811.
- [25] J. Gu, Y. Wang, Y. Chen, V. O. Li, and K. Cho, "Meta-learning for low-resource neural machine translation," in *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, 2018, pp. 3622–3631.
- [26] Z.-Y. Dou, K. Yu, and A. Anastasopoulos, "Investigating meta-learning algorithms for low-resource natural language understanding tasks," in *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, 2019, pp. 1192–1197.
- [27] J. A. A. Prenner and R. Robbes, "Making the most of small software engineering datasets with modern machine learning," *IEEE Transactions on Software Engineering*, 2022.
- [28] J. Phang, T. Févry, and S. R. Bowman, "Sentence encoders on stilts: Supplementary training on intermediate labeled-data tasks," *arXiv preprint arXiv:1811.01088*, 2018.
- [29] M. Chen, J. Tworek, H. Jun, Q. Yuan, H. P. d. O. Pinto, J. Kaplan, H. Edwards, Y. Burda, N. Joseph, G. Brockman *et al.*, "Evaluating large language models trained on code," *arXiv preprint arXiv:2107.03374*, 2021.
- [30] Y. Li, D. Choi, J. Chung, N. Kushman, J. Schrittwieser, R. Leblond, T. Eccles, J. Keeling, F. Gimeno, A. Dal Lago *et al.*, "Competition-level code generation with alphacode," *Science*, vol. 378, no. 6624, pp. 1092–1097, 2022.
- [31] J. A. Prenner and R. Robbes, "Automatic program repair with openai's codex: Evaluating quixbugs," *arXiv preprint arXiv:2111.03922*, 2021.
- [32] S. D. Kolak, R. Martins, C. Le Goues, and V. J. Hellendoorn, "Patch generation with language models: Feasibility and scaling behavior," in *Deep Learning for Code Workshop*, 2022.
- [33] H. Pearce, B. Tan, B. Ahmad, R. Karri, and B. Dolan-Gavitt, "Examining zero-shot vulnerability repair with large language models," in *2023 IEEE Symposium on Security and Privacy (SP)*. IEEE Computer Society, 2022, pp. 1–18.
- [34] O. Asare, M. Nagappan, and N. Asokan, "Is github's copilot as bad as humans at introducing vulnerabilities in code?" *arXiv preprint arXiv:2204.04741*, 2022.
- [35] H. Pearce, B. Ahmad, B. Tan, B. Dolan-Gavitt, and R. Karri, "Asleep at the keyboard? assessing the security of github copilot's code contributions," in *2022 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2022, pp. 754–768.
- [36] J. Austin, A. Odena, M. Nye, M. Bosma, H. Michalewski, D. Dohan, E. Jiang, C. Cai, M. Terry, Q. Le *et al.*, "Program synthesis with large language models," *arXiv preprint arXiv:2108.07732*, 2021.
- [37] E. Triantafyllou, T. Zhu, V. Dumoulin, P. Lamblin, U. Evci, K. Xu, R. Goroshin, C. Gelada, K. Swersky, P.-A. Manzagol, and H. Larochelle, "Meta-dataset: A dataset of datasets for learning to learn from few examples," in *International Conference on Learning Representations*, 2020.
- [38] Y. Wang, S. Mishra, P. Alipoormolabashi, Y. Kordi, A. Mirzaei, A. Arunkumar, A. Ashok, A. S. Dhanasekaran, A. Naik, D. Stap *et al.*, "Benchmarking generalization via in-context instructions on 1,600+ language tasks," *arXiv preprint arXiv:2204.07705*, 2022.
- [39] Y. Taigman, M. Yang, M. Ranzato, and L. Wolf, "Deepface: Closing the gap to human-level performance in face verification," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2014, pp. 1701–1708.
- [40] D. Silver, A. Huang, C. J. Maddison, A. Guez, L. Sifre, G. Van Den Driessche, J. Schrittwieser, I. Antonoglou, V. Panneershelvam, M. Lanctot *et al.*, "Mastering the game of go with deep neural networks and tree search," *nature*, vol. 529, no. 7587, pp. 484–489, 2016.
- [41] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.
- [42] A. Najberg, "Alibaba ai model tops humans in reading comprehension," *Alizila.com*, January, vol. 15, 2018.
- [43] C. Berner, G. Brockman, B. Chan, V. Cheung, P. Debiak, C. Dennison, D. Farhi, Q. Fischer, S. Hashme, C. Hesse *et al.*, "Dota 2 with large scale deep reinforcement learning," *arXiv preprint arXiv:1912.06680*, 2019.
- [44] S. M. Xie, A. Raghunathan, P. Liang, and T. Ma, "An explanation of in-context learning as implicit bayesian inference," in *International Conference on Learning Representations*, 2021.

- [45] Y. Wang, Q. Yao, J. T. Kwok, and L. M. Ni, "Generalizing from a few examples: A survey on few-shot learning," *ACM computing surveys (csur)*, vol. 53, no. 3, pp. 1–34, 2020.
- [46] W. Yin, "Meta-learning for few-shot natural language processing: A survey," *arXiv preprint arXiv:2007.09604*, 2020.
- [47] D. Paperno, G. Kruszewski, A. Lazaridou, N.-Q. Pham, R. Bernardi, S. Pezzelle, M. Baroni, G. Boleda, and R. Fernández, "The lambda-bada dataset: Word prediction requiring a broad discourse context," in *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 2016, pp. 1525–1534.
- [48] M. Joshi, E. Choi, D. S. Weld, and L. Zettlemoyer, "Triviaqa: A large scale distantly supervised challenge dataset for reading comprehension," in *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 2017, pp. 1601–1611.
- [49] N. Durrani, B. Haddow, P. Koehn, and K. Heafield, "Edinburgh's phrase-based machine translation systems for wmt-14," in *Proceedings of the Ninth Workshop on Statistical Machine Translation*, 2014, pp. 97–104.
- [50] A. Madotto, Z. Liu, Z. Lin, and P. Fung, "Language models as few-shot learner for task-oriented dialogue systems," *arXiv preprint arXiv:2008.06239*, 2020.
- [51] Z. Chen, H. Eavani, W. Chen, Y. Liu, and W. Y. Wang, "Few-shot nlg with pre-trained language model," in *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 2020, pp. 183–190.
- [52] G. I. Winata, A. Madotto, Z. Lin, R. Liu, J. Yosinski, and P. Fung, "Language models are few-shot multilingual learners," in *Proceedings of the 1st Workshop on Multilingual Representation Learning*, 2021, pp. 1–15.
- [53] R. Zellers, A. Holtzman, Y. Bisk, A. Farhadi, and Y. Choi, "Hellaswag: Can a machine really finish your sentence?" in *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, 2019, pp. 4791–4800.
- [54] G. Lai, Q. Xie, H. Liu, Y. Yang, and E. Hovy, "Race: Large-scale reading comprehension dataset from examinations," in *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, 2017, pp. 785–794.
- [55] Y. Nie, A. Williams, E. Dinan, M. Bansal, J. Weston, and D. Kiela, "Adversarial nli: A new benchmark for natural language understanding," in *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 2020, pp. 4885–4901.
- [56] S. Bach, V. Sanh, Z. X. Yong, A. Webson, C. Raffel, N. V. Nayak, A. Sharma, T. Kim, M. S. Bari, T. Févry *et al.*, "Promptsources: An integrated development environment and repository for natural language prompts," in *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, 2022, pp. 93–104.
- [57] V. Sanh, A. Webson, C. Raffel, S. Bach, L. Sutawika, Z. Alyafeai, A. Chaffin, A. Stiegler, A. Raja, M. Dey *et al.*, "Multitask prompted training enables zero-shot task generalization," in *International Conference on Learning Representations*, 2021.
- [58] L. Ouyang, J. Wu, X. Jiang, D. Almeida, C. L. Wainwright, P. Mishkin, C. Zhang, S. Agarwal, K. Slama, A. Ray *et al.*, "Training language models to follow instructions with human feedback," *arXiv preprint arXiv:2203.02155*, 2022.
- [59] C.-Y. Lin and E. Hovy, "Manual and automatic evaluation of summaries," in *Proceedings of the ACL-02 Workshop on Automatic Summarization - Volume 4*, 2002, pp. 45–51.
- [60] S. Lu, D. Guo, S. Ren, J. Huang, A. Svyatkovskiy, A. Blanco, C. Clement, D. Drain, D. Jiang, D. Tang, G. Li, L. Zhou, L. Shou, L. Zhou, M. Tufano, M. GONG, M. Zhou, N. Duan, N. Sundaresan, S. K. Deng, S. Fu, and S. LIU, "CodeXGLUE: A machine learning benchmark dataset for code understanding and generation," in *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 1)*, 2021.
- [61] A. Elnaggar, W. Ding, L. Jones, T. Gibbs, T. Feher, C. Angerer, S. Severini, F. Matthes, and B. Rost, "Codetrans: Towards cracking the language of silicon's code through self-supervised deep learning and high performance computing," *arXiv preprint arXiv:2104.02443*, 2021.
- [62] M. Zhu, A. Jain, K. Suresh, R. Ravindran, S. Tipirneni, and C. K. Reddy, "Xlcost: A benchmark dataset for cross-lingual code intelligence," *arXiv preprint arXiv:2206.08474*, 2022.
- [63] R. Puri, D. S. Kung, G. Janssen, W. Zhang, G. Domeniconi, V. Zolotov, J. Dolby, J. Chen, M. Choudhury, L. Decker *et al.*, "Codenet: A large-scale ai for code dataset for learning a diversity of coding tasks," in *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 2)*, 2021.
- [64] H. Husain, H.-H. Wu, T. Gazit, M. Allamanis, and M. Brockschmidt, "Codesearchnet challenge: Evaluating the state of semantic code search," 2020.
- [65] L. Sgier, "Qualitative data analysis," *An Initiat. Gebert Ruf Stift.*, vol. 19, pp. 19–21, 2012.
- [66] O. Chaparro, J. Lu, F. Zampetti, L. Moreno, M. Di Penta, A. Marcus, G. Bavota, and V. Ng, "Detecting missing information in bug descriptions," in *Proceedings of the 2017 11th Joint Meeting on Foundations of Software Engineering*, 2017, pp. 396–407.
- [67] J. Zhai, X. Xu, Y. Shi, G. Tao, M. Pan, S. Ma, L. Xu, W. Zhang, L. Tan, and X. Zhang, "Cpc: Automatically classifying and propagating natural language comments via program analysis," in *Proceedings of the ACM/IEEE 42nd International Conference on Software Engineering*, 2020, pp. 1359–1371.
- [68] Q. Chen, X. Xia, H. Hu, D. Lo, and S. Li, "Why my code summarization model does not work: Code comment improvement with category prediction," *ACM Transactions on Software Engineering and Methodology (TOSEM)*, vol. 30, no. 2, pp. 1–29, 2021.
- [69] J. Cohen, "A coefficient of agreement for nominal scales," *Educational and psychological measurement*, vol. 20, no. 1, pp. 37–46, 1960.
- [70] A. J. Viera, J. M. Garrett *et al.*, "Understanding interobserver agreement: the kappa statistic," *Fam med*, vol. 37, no. 5, pp. 360–363, 2005.
- [71] M. Tufano, C. Watson, G. Bavota, M. D. Penta, M. White, and D. Poshyvanyk, "An empirical study on learning bug-fixing patches in the wild via neural machine translation," *ACM Transactions on Software Engineering and Methodology (TOSEM)*, vol. 28, no. 4, pp. 1–29, 2019.
- [72] M. Tufano, C. Watson, G. Bavota, M. Di Penta, M. White, and D. Poshyvanyk, "Learning how to mutate source code from bug-fixes," in *2019 IEEE International Conference on Software Maintenance and Evolution (ICSME)*. IEEE Computer Society, 2019, pp. 301–312.
- [73] J. Svajlenko, J. F. Islam, I. Keivanloo, C. K. Roy, and M. M. Mia, "Towards a big data curated benchmark of inter-project code clones," in *Proceedings of the 2014 IEEE International Conference on Software Maintenance and Evolution*, 2014, pp. 476–480.
- [74] M. Lewis, Y. Liu, N. Goyal, M. Ghazvininejad, A. Mohamed, O. Levy, V. Stoyanov, and L. Zettlemoyer, "Bart: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension," in *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 2020, pp. 7871–7880.
- [75] K. Papineni, S. Roukos, T. Ward, and W.-J. Zhu, "Bleu: a method for automatic evaluation of machine translation," in *Proceedings of the 40th annual meeting of the Association for Computational Linguistics*, 2002, pp. 311–318.
- [76] M. Allamanis, "The adverse effects of code duplication in machine learning models of code," in *Proceedings of the 2019 ACM SIGPLAN International Symposium on New Ideas, New Paradigms, and Reflections on Programming and Software*, 2019, pp. 143–153.