Explicit Knowledge Transfer for Weakly Supervised Code Generation

Zhangir Azerbayev, Ansong Ni, Hailey Schoelkopf, Dragomir Radev

Yale

LILY Lab

Introduction

In this work we introduce *Explicit Knowledge Transfer* (EKT) as a method for training small language model to generate code with weak supervision by leveraging a black-box teacher model capable of few-shot learning .

Language models have demonstrated impressive performance for code generation within two main paradigms. The first, few-shot learning, has great generalization power, but is only performant with large models that are computationally expensive during both training and inference. The other paradigm, supervised fine-tuning, is performant with much smaller models, but requires the costly process of collecting fully-labeled training examples [1]. EKT attempts to achieve the best of both worlds: high performance from lean models trained on a minimal quantity of labelled data.

Because EKT leverages a black-box teacher model, it is well-suited to the computational constraints of many machine learning practitioners. WIth the accelerating growth in the size of large language models, it is common for many engineers to only have access to state-of-the-art models through an external API or be unable to fine-tune large models locally. Moreover, one may wish to deploy a smaller model in applications that are sensitive to the speed of inference, for example in-IDE code completion engines.

Method

EKT is applicable to training on weakly-supervised data, where training examples consist of questions, test cases, but do not have code solutions. Applying EKT requires a teacher model, which has strong few-shot learning capabilities, and a student model, which is a smaller pre-trained model.

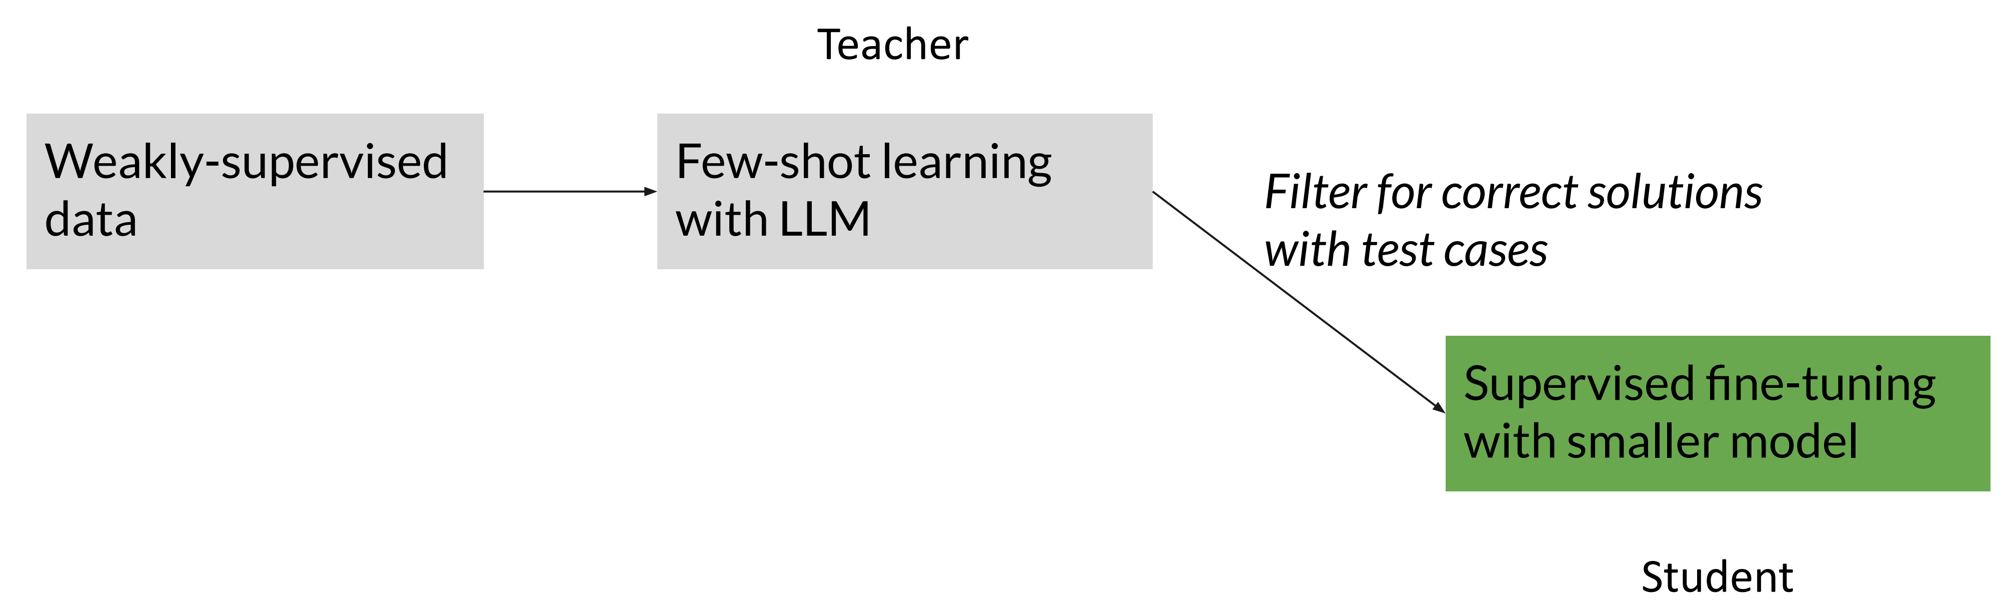

Fig 1. Schematic depiction of Explicit Knowledge Transfer

Problem: Beth bakes 4, 2 dozen batches of cookies in a week. If these cookies are shared amongst 16 people equally, how many cookies does each person consume?
Solution: Beth bakes 4 2 dozen batches of cookies for a total of 4*2 = <<4*2=8>>8 dozen cookies
There are 12 cookies in a dozen and she makes 8 dozen cookies for a total of 12*8 = <<12*8=96>>96 cookies
She splits the 96 cookies equally amongst 16 people so they each eat 96/16 = <<96/16=6>>6 cookies
Final Answer: 6

Problem: Mrs. Lim milks her cows twice a day. Yesterday morning, she got 68 gallons of milk and in the evening, she got 82 gallons. This morning, she got 18 gallons fewer than she had yesterday morning. After selling some gallons of milk in the afternoon, Mrs. Lim has only 24 gallons left. How much was her revenue for the milk if each gallon costs $3.50?
Mrs. Lim got 68 gallons - 18 gallons = <<68-18=50>>50 gallons this morning.
So she was able to get a total of 68 gallons + 82 gallons + 50 gallons = <<68+82+50=200>>200 gallons.
She was able to sell 200 gallons - 24 gallons = <<200-24=176>>176 gallons.
Thus, her total revenue for the milk is $3.50/gallon x 176 gallons = $<<3.50*176=616>>616.
Final Answer: 616

FIg 2. Example GSM8k instances. Our models are trained only on the question and final answer, not the step-by-step solutions.

Teacher	Student	Training Method	Accuracy (%)
None	GPT-Neo 125M	Few-shot	0
Davinci-code-002	GPT-Neo 125M	EKT	**2.1**
None	GPT-Neo 1.3B	Few-shot	0.3
Davinci-code-002	GPT-Neo 1.3B	EKT	**14.8**

Table 3. Results on the GSM8k dataset. We sample 20 candidate solutions from the teacher for every training example.

The EKT algorithm is as follows. First, create a few-shot prompt by manually annotating a single-digit number of training examples with code solutions (we use 3-shot learning). Next, use this prompt to do few-shot learning with the black-box teacher on the entire training set. If a program that passes all test cases is generated, treat it as a gold solution. Finally, train with supervised-fine tuning on the training examples for which gold solutions were discovered.

Experimental Setup

Our experiments target the GSM8k dataset [2], a dataset of grade-school level math word problems. Crucially, GSM8k training examples are not annotated with code solutions, necessitating weakly-supervised learning. GSM8k contains 7500 training examples and 1000 test examples.

For our teacher model, we use OpenAI's *davinci-code-002* engine, a variant of the Codex language model trained on code [3]. For our student model, we use variants of the GPT-Neo model with different parameter counts.

We compare EKT to a baseline of few-shot learning with the student only.

Results

As can be seen in table 3, EKT dramatically outperforms few-shot learning with the student only. This demonstrates the viability of EKT as a method for training small models on the code generation task when fully-labelled training training data is unavailable.

References

[1] Austin et al., 2019

[2] Cobbe et al., 2021

[3] Chen et al., 2021

[4] Gao et al.