

# CoTran: An LLM-Based Code Translator Using Reinforcement Learning with Feedback from Compiler and Symbolic Execution

Prithwish Jana<sup>a,\*</sup>, Piyush Jha<sup>a</sup>, Haoyang Ju<sup>b</sup>, Gautham Kishore<sup>c</sup>, Aryan Mahajan<sup>d</sup> and Vijay Ganesh<sup>a</sup>

<sup>a</sup>Georgia Institute of Technology, USA

<sup>b</sup>University of Toronto, Canada

<sup>c</sup>University of California San Diego, USA

<sup>d</sup>Columbia University, USA

**Abstract.** In this paper, we present an LLM-based code translation method and an associated tool called CoTran, that translates whole-programs from one high-level programming language to another. Existing LLM-based code translation methods lack training to ensure that the translated code reliably compiles or bears substantial functional equivalence to the input code. In our work, we fine-tune an LLM using reinforcement learning, incorporating compiler feedback, and symbolic execution (symexec)-based testing feedback to assess functional equivalence between the input and output programs. The idea is to guide an LLM during fine-tuning, via compiler and symexec-based testing feedback, by letting it know how far it is from producing perfect translations. We conduct extensive experiments comparing CoTran with 14 other code translation tools, including human-written transpilers, LLM-based translation tools, and ChatGPT. Using a benchmark of over 57,000 code pairs in Java and Python, we demonstrate that CoTran outperforms the other tools on relevant metrics such as compilation accuracy (CompAcc) and functional equivalence accuracy (FEqAcc). For example, in Python-to-Java translation, CoTran achieves 48.68% FEqAcc and 76.98% CompAcc, whereas the nearest competing tool (PLBART-base) gets 38.26% and 75.77% respectively. Additionally, CoTran, built on top of CodeT5, improves FEqAcc by +14.89% and CompAcc by +8.14% for Python-to-Java (resp., +12.94% and +4.30% for Java-to-Python).

## 1 Introduction

Automatic translation of code from one high-level language to another is an important area of software engineering research with applications in code migration and cross-platform interoperability [14, 27]. Traditional code translation tools, often called *transpilers* rely on human-written rules and, thus, can be quite expensive to develop.

To address this issue, many researchers [9, 25] have recently proposed the use of Large Language Models (LLMs) [44] for code translation. The rise of LLMs is perhaps the most important development in AI in recent years, and even more remarkably they are being used successfully for many software engineering tasks such as code synthesis [40] and code completion [46]. In this context, LLMs have been used for translating a single function from a source (e.g., Java) to

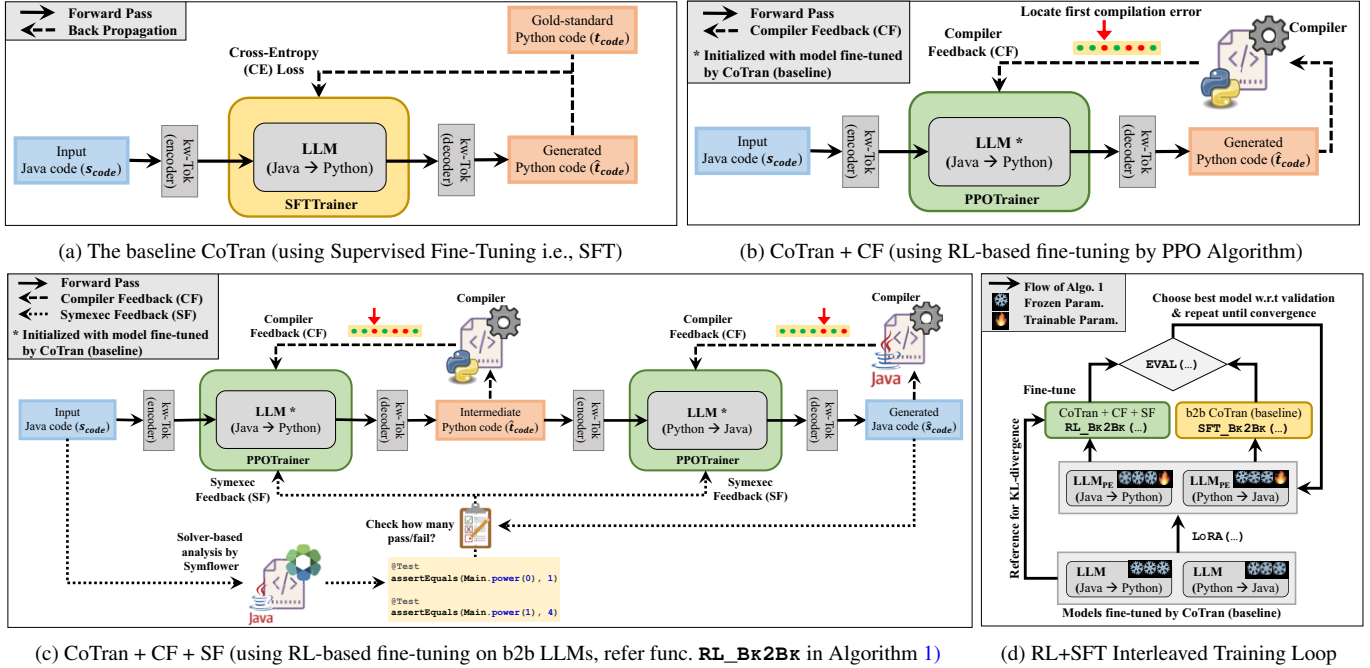
a target language (e.g., Python) [1, 25, 36]. Unfortunately, function-to-function translation is not sufficient (refer Section 5.3: Finding 8) for whole-program translation tasks. Further, current LLM-based methods lack a proactive approach to ensure that the translated code reliably compiles or bears substantial functional equivalence to the input code. This is a serious problem, hindering the wider adoption of LLM-based code translation techniques.

We address this problem by modifying the fine-tuning mechanism for code LLMs to incorporate feedback from compiler and symbolic execution (symexec)-based testing. Our method retains the positive aspect of LLM-based code translation, while improving its accuracy dramatically. The core idea behind our method is to employ *corrective feedback loops* [12, 18] between learning (the LLM) and reasoning (compiler & symexec), for which we have a two-fold strategy. Firstly, if the output code produced by the LLM does not compile, a **compiler feedback (CF)** is computed. Secondly, if the output and input codes are not equivalent (w.r.t. a symexec-generated test suite), a suitable **symexec feedback (SF)** is computed. This feedback is then used in a reinforcement learning (RL) framework to fine-tune the LLM. Our results on **Java-to-Python (J2P)** and **Python-to-Java (P2J)** translation indicate that these new feedback functions significantly improve accuracy compared to previous methods. Another interesting feature is that we use two back-to-back LLMs, one for source-to-target ( $S \rightarrow T$ ) and another for target-to-source ( $T \rightarrow S$ ), that are fine-tuned together. This setup of complementary models simplifies the automation of functional equivalence checking between the input (of  $S \rightarrow T$  model) and output (of  $T \rightarrow S$  model), both in  $S$ .

### Contributions.

- We describe a whole-program translation method (and tool CoTran) based on interleaved training of LLMs through Supervised Fine-Tuning (SFT)- and RL-based optimization, to incorporate CF and SF into the fine-tuning process. Currently, CoTran is configured to translate between Java and Python (J2P and P2J). It can be effortlessly adapted for other language pairs ( $S, T$ ) given: (i) a dataset of equivalent code pairs in  $S$  and  $T$ , (ii) compilers for both languages and (iii) an automatic test-case generation tool for either  $S$  or  $T$  – all of which are readily available for popular languages. A key insight is that LLMs trained for software engineering tasks (e.g., code synthesis) can greatly benefit from feedback via compilers, program analysis, and test generation tools. (Section 3.2)

\* Corresponding Author. Email: pjana7@gatech.edu



**Figure 1: Fine-tuning LLMs with Compiler Feedback (CF) and Symbolic Execution Feedback (SF):** (a) **CoTran (baseline)** is an LLM fine-tuned without feedback loops. It is fine-tuned to optimize cross-entropy loss, employing the proposed keyword-based tokenizer (kw-Tok). (b) **CoTran + CF** is an LLM fine-tuned using an RL-based PPOTrainer where the compiler provides feedback in the form of a reward. For P2J, javac compiler is used. For J2P, we use the pylint [41] static code analyzer as Python is an interpreted language. (c) In **CoTran + CF + SF**, we jointly fine-tune back-to-back (b2b) LLMs (the forward LLM for J2P and the backward LLM for P2J) using RL. Through solver-based analysis of the input Java code, we generate a set of JUnit tests, which are then verified on the output Java code to compute SF. The reward for each LLM is a combination of SF and the respective CF. (d) We begin with the forward (J2P) and backward (P2J) LLMs, fine-tuned by CoTran (baseline) and with frozen parameters. We then continue by fine-tuning a parameter-efficient (PE) version of these LLMs by interleaving RL and SFT, training only a few additional parameters through Low-Rank Adaptation (LoRA) [16] (refer Algorithm 1).

- Our fine-grained compiler feedback (CF) guides an LLM during fine-tuning, helping it assess the proximity of generated translations to a perfectly compiling one. It is much more effective than a Boolean yes/no feedback e.g., for P2J, CoTran achieves +11.57% higher functional equivalence accuracy and +17.36% higher compilation accuracy compared to PPOCoder [39]. (Section 3.1)
- Our symexec feedback (SF) is used to fine-tune  $S \rightarrow T \rightarrow S$  back-to-back LLMs. We generate unit tests on the input code and check them on the output translated code, to provide feedback on their *inequivalence*. Our method is entirely agnostic to any specific test-case generation tool. As far as we know, no existing methods utilize symexec-based test generation for functional equivalence checking and associated feedback function. (Section 3.1)
- We perform an extensive empirical evaluation and ablation study of CoTran against 11 state-of-the-art LLM-based translation tools and 3 human-written transpilers. A range of metrics are assessed incl., functional equivalence<sup>1</sup> accuracy (FEqAcc), compilation accuracy (CompAcc), BLEU [31], CodeBLEU [35] and the proposed average first error position (errPos<sub>1st</sub>). CoTran and its variants outperform all other tools of similar size for both J2P and P2J. Compared to the nearest competing tool on FEqAcc, CoTran gets +9.62% (vs. PPOCoder [39]) in J2P and +10.42% (vs. PLBART-base [1]) in P2J. CoTran even outperforms ChatGPT [29] (a much larger model) on all the metrics in P2J and all but one metric (FE-

qAcc) in J2P. Furthermore, our tool outperforms transpilers, e.g., for J2P, CoTran gets +53.43% in FEqAcc, +38.84% in CompAcc, and +38.47% in errPos<sub>1st</sub> (vs. TSS CodeConv [43]). (Section 5)

- We introduce a specialized Keyword Tokenizer (kw-Tok) for code translation, boosting FEqAcc by +3.57% for J2P and +6.62% for P2J, and CompAcc by +3.28% and +4.79%, resp. (Section 3)
- We contribute a large, well-curated dataset AVATAR-TC (built on top of AVATAR [2]) having 57,000+ Java-Python code pairs with human-written test-cases (TCs). Compared to AVATAR, faulty codes are manually fixed and TCs (by problem-setters) are collected from coding platforms, enabling the calculation of translator FEqAcc. Based on what we know, this is the first large-scale dataset thoroughly testing code pairs with human-written TCs. (Section 4)

## 2 Related Work

**Rule-based Transpilers.** Rule-based handcrafted transpilers are usually built using traditional compiler techniques and concepts such as parsing and abstract syntax trees. Examples of such transpilers include java2python [28] and py2java [10]. TSS code converter [43] is a commercial J2P transpiler. Overall, these transpilers vary by the intricacies and difficulty level of constructs (e.g., lambda, anonymous inner class). Many such tools state a disclaimer with limitations that the translated codes should not be expected to compile and run readily.

**Transformer and ML-based Code Translation Tools.** The encoder-decoder architecture of Transformers revolutionized Natural Language (NL) translation, using contextualized representations of words [44]. This led to the evolution of advanced language mod-

<sup>1</sup> Note that program equivalence is an undecidable problem in general, and we make no claims about establishing program equivalence. By functional equivalence of two codes, we mean something more limited, namely, that they produce equivalent results w.r.t. a test suite.

els for code translation task, such as encoder-only CodeBERT [9] and decoder-only CodeGPT [25]. CodeT5 [47] is an encoder-decoder architecture that leverages code semantics from developer-assigned identifiers. Another tool is PLBART [1] which is a unified transformer trained through denoising autoencoding and performs a range of tasks among NL and Programming Language (PL). Conversely, tree-to-tree [6] is an attention-based LSTM neural network.

**LLM-based methods using compiler or unit testing.** TransCoder-ST [37] is an unsupervised code translation tool that uses self-training and automated unit tests to verify source-target code equivalence. However, it uses unit tests to create a synthetic dataset of equivalent codes in two languages, instead of incorporating it for model improvement during training. Recently, Boolean compiler feedback was used to fine-tune LLM for code generation in an RL-based scheme [46]. PPOCoder [39] is an RL-based code translation framework that additionally uses CodeBLEU-inspired feedback in the LLM fine-tuning process. While both the tools use the same Boolean compiler feedback that is true if the generated code compiles and false otherwise, our proposed compiler feedback function is fine-grained – it captures how far the translation is from a perfectly compilable code. Back-to-back (b2b) code translation although proposed earlier [36], is not leveraged for symexec-based functional equivalence checking.

### 3 The CoTran Method

**The Code Translation Learning Problem.** Let  $S$  denote the source language and  $T$  the target language. The code translation learning problem is to learn a language translator  $f_{ST} : S \rightarrow T$ , such that the output of the translator is a  $T$ -program that is *syntactically correct* (as per the grammar of  $T$ ) and is *functionally equivalent* to the input  $S$ -program (input-output equivalence w.r.t. a test suite).

**Brief Overview of our Method.** Please refer to Figure 1 for an architectural overview of the CoTran tool. In order to learn the language translator, we build off a sequence-to-sequence (seq2seq) attention-based transformer model consisting of an encoder and decoder stack [44]. For the remainder of the paper, we use the term Large Language Models (LLMs) to refer to such models. The proposed fine-tuning pipeline consists of two steps. In the first step, we separately perform supervised fine-tuning (SFT) on  $LLM_f : S \rightarrow T$  and  $LLM_b : T \rightarrow S$  (the forward and backward models respectively) to optimize *cross-entropy loss*. In the second step, we jointly fine-tune these two models in a parameter-efficient way on the back-to-back translation task i.e.,  $S \rightarrow T \rightarrow S$ . For this process, use an RL+SFT interleaved training loop that incorporates *compiler-* and *symexec-feedback*. Figure 1 illustrates the feedback mechanisms.

**Keyword Tokenizer (kw-Tok).** The code  $s_{code} \in S$  is first tokenized based on the syntax (grammar) of the corresponding language. Next, the code-specific tokens thus obtained, are converted to a sequence of token-ids using a RoBERTa tokenizer, which is modified for the PL translation task. To avoid splitting the PL keywords into multiple subtokens, we add all the keywords (e.g., `volatile`, `transitive` in Java; `elif`, `instanceof` in Python) and operators (e.g., `>=`, `**=`) of  $S$  and  $T$  to the tokenizer vocabulary. These are collected from the list of terminal symbols of these languages' formal grammar. Grammars-v4 [32] provides a collection of formal grammars for most common languages in use. We also add special tokens to the vocabulary e.g., `NEW_LINE`, `INDENT`, `DEDENT` for Python. We refer to this modified RoBERTa tokenizer as *Keyword Tokenizer (kw-Tok)*. This modification ensures that *kw-Tok* generates a *single token-id* for all the special tokens, keywords, and operators of languages  $S$  and  $T$ . Thus, given a training example  $(s_{code}, t_{code}) \in \text{dataset } D$ , the tokenizer pre-

processes these to a sequence of token-ids viz.,  $\mathbf{s} = (s_1, s_2, \dots, s_n)$  and  $\mathbf{t} = (t_1, t_2, \dots, t_m)$ . The LLM ( $S \rightarrow T$ ) takes the source sequence  $\mathbf{s}$  and generates a target sequence  $\hat{\mathbf{t}} = (\hat{t}_1, \hat{t}_2, \dots, \hat{t}_q)$ .

#### 3.1 Definitions: CE Loss, CF, and SF

Refer to Figure 1 for a pictorial representation of cross-entropy (CE) loss, compiler feedback (CF), and symexec feedback (SF). Let  $\theta_f$  and  $\theta_b$  be the trainable parameters for  $LLM_f$  and  $LLM_b$  respectively.

**Cross-Entropy (CE) Loss.** As is typical [47, 44] in translation tasks, LLMs are fine-tuned to minimize the cross-entropy (CE) loss through supervised learning. For the forward LLM that generates  $\hat{\mathbf{t}}$  given a tokenized training instance  $(\mathbf{s}, \mathbf{t})$ , the CE loss is defined as:

$$\mathcal{L}_{CE}^{\theta_f}(\mathbf{t}, \hat{\mathbf{t}}) = -\frac{1}{\ell} \sum_{i=1}^{\ell} \sum_{j=1}^{|V|} \mathbb{1}_{ij} \log P_{ij}^{\theta_f} \quad (1)$$

where,  $\ell = \max(|\hat{\mathbf{t}}|, |\mathbf{t}|)$  and  $V$  is the tokenizer vocabulary.  $\mathbb{1}_{ij}$  is 1 iff the  $i^{\text{th}}$  token in reference translation  $\mathbf{t}$  is the  $j^{\text{th}}$  word of  $V$  and,  $P_{ij}^{\theta_f}$  is the probability that the  $i^{\text{th}}$  token in predicted translation  $\hat{\mathbf{t}}$  is the  $j^{\text{th}}$  word of  $V$ . The CE loss is, however, more suitable for machine translation of NLs than PLs. For NLs, the grammar is lenient, and an approximate translation is oftentimes good enough. But in PL translation where the grammar is strict and functional equivalence is paramount, the aim is to generate compilable codes in the target language and, to make sure that  $\hat{\mathbf{t}}$  is functionally equivalent to  $\mathbf{s}$ .

**Compiler Feedback (CF).** We use a compiler to assess the syntactic correctness of an LLM-generated translation during fine-tuning and in turn, provide feedback to the LLM. This guides the model in determining how the translation fares relative to a perfectly compilable code. For CF, the position of the first compilation error serves as a key heuristic in evaluating syntactic correctness. If it appears near the end of the code, it typically indicates fewer cascading errors and a closer-to-perfect compilation. So, for the forward LLM that generates  $\hat{\mathbf{t}}$  given a tokenized instance  $(\mathbf{s}, \mathbf{t})$ , we formulate the feedback as:

$$\omega_{\text{compiler}}(\hat{\mathbf{t}}) = \begin{cases} +2 & , \text{ if } \hat{\mathbf{t}} \text{ compiles} \\ \frac{f(\text{compiler}_{T, \hat{\mathbf{t}}})}{|\hat{\mathbf{t}}|+1} & , \text{ otherwise} \end{cases} \quad (2)$$

where, the function  $f(\cdot, \cdot)$  applies the  $T$ -language compiler on  $\hat{\mathbf{t}}$  and returns the token position ( $\in [1, |\hat{\mathbf{t}}|]$ ) of the first syntax error. So,  $\omega_{\text{compiler}}$  is +2 for a perfect compilation and goes on decreasing from +1 to 0 as the token position of the first syntax error is closer to the beginning of the code. However, the LLM should not game the system to generate  $\hat{\mathbf{t}}$  as a dummy code with no compilation error e.g., a Hello-World program. To penalize such cases we posit that  $\hat{\mathbf{t}}$  should also be close-by in length to  $\mathbf{t}$ . Accordingly, CF is defined as the product of  $\omega_{\text{compiler}}$  and the value of a  $(0, 1]$  Gaussian distribution at  $|\hat{\mathbf{t}}|$ , which is centered at  $|\mathbf{t}|$  and has a standard deviation of  $\frac{|\mathbf{t}|}{4}$ , as follows:

$$\omega_{CF}(\mathbf{t}, \hat{\mathbf{t}}) = \omega_{\text{compiler}}(\hat{\mathbf{t}}) \times e^{-\frac{1}{2} \left( \frac{|\hat{\mathbf{t}}| - |\mathbf{t}|}{|\mathbf{t}|/4} \right)^2} \quad (3)$$

**Symbolic Execution Feedback (SF).** We also introduce a *symexec feedback* (SF) that lets the LLM during fine-tuning know the extent to which the generated translation is *(in)-equivalent*<sup>2</sup> to the source-language code. As it can be challenging to assess functional inequivalence among two codes of different languages ( $\mathbf{s}$  and  $\hat{\mathbf{t}}$ ), we use two back-to-back (b2b) LLMs:  $S \rightarrow T$  and  $T \rightarrow S$  such that the

<sup>2</sup> It goes without saying that one cannot guarantee functional equivalence between two programs with a purely testing approach. However, we can establish inequivalence via a sufficient amount of testing.



same test suite can be used for inequivalence testing. Given a tokenized training instance  $(s, t)$ , the forward LLM ( $S \rightarrow T$ ) translates  $s$  to an intermediate target-language sequence  $\hat{t}$ , which the backward LLM ( $T \rightarrow S$ ) translates back to  $\hat{s}$  in an attempt to reconstruct  $s$ . To compute SF, we use solver-based analysis by an industrial-strength symexec tool called Symflower [49]. Through symexec, Symflower generates essential unit tests even for functions involving complex data types. It computes the necessary inputs and expected outputs to cover all linearly independent control-flow paths within a function. For J2P,  $S$  is Java and  $T$  is Python. We use Symflower to automatically generate JUnit tests ( $\mathcal{J}_s$ ) for each function in  $s \in S$ . These unit tests are checked on the corresponding function in  $\hat{s}$  and both LLMs get feedback on how many tests pass. Accordingly, SF is defined as:

$$\omega_{SF}(s, \hat{s}) = \frac{\epsilon + \sum_{j \in \mathcal{J}_s} \mathbb{1}_{j(\hat{s}) \equiv \text{Success}}}{\epsilon + |\mathcal{J}_s|} \quad (4)$$

where,  $\epsilon$  is a small positive value and  $\mathbb{1}_{j(\hat{s}) \equiv \text{Success}}$  is 1 when the  $j^{\text{th}}$  JUnit test on  $s$  successfully passes on  $\hat{s}$ , else it is 0. Our interleaved training loop (Section 3.2) incorporates CF on individual LLMs, in addition to SF. So, a feedback on  $(s, \hat{s})$ -inequivalence with a feedback on compilability of  $\hat{t}$  and  $\hat{s}$ , correlates to  $(s, \hat{t})$ -inequivalence.

### 3.2 The CoTran Training Loop

Given a training instance  $(s_{\text{code}}, t_{\text{code}}) \in \text{dataset } D$ , kw-Tok converts them into token-ids,  $s = (s_1, s_2, \dots, s_n)$  and  $t = (t_1, t_2, \dots, t_m)$ . We outline a two-step process for fine-tuning the LLMs.

**Supervised fine-tuning of translation models,  $\text{LLM}_f$  and  $\text{LLM}_b$ :** In the first phase, we separately perform SFT on  $\text{LLM}_f$  and  $\text{LLM}_b$  to minimize the CE loss over the respective translation tasks.  $\text{LLM}_f$  translates  $s$  to a  $T$ -sequence  $\hat{t} = (\hat{t}_1, \hat{t}_2, \dots, \hat{t}_q)$ . Through SFT, it is optimized to minimize the CE loss between  $\hat{t}$  and  $t$  i.e.,  $\mathcal{L}_{CE}^{\theta_f}(t, \hat{t})$  as defined in Eqn. 1. Similarly,  $\text{LLM}_b$  translates  $t$  to a  $S$ -sequence  $\hat{s} = (\hat{s}_1, \hat{s}_2, \dots, \hat{s}_r)$  and is optimized to minimize  $\mathcal{L}_{CE}^{\theta_b}(s, \hat{s})$ .

**Jointly fine-tuning  $\text{LLM}_f$  and  $\text{LLM}_b$  using interleaved RL+SFT:** Next, we further fine-tune  $\text{LLM}_f$  and  $\text{LLM}_b$ , but this time together through back-to-back (b2b) translation in order to incorporate CF and SF. Given a source-language code, a successful code translation calls for generating a target-language code that *compiles* and is *functionally equivalent* to the input code. Here, optimizing LLMs with CE is not sufficient. So, we defined feedback mechanisms (CF and SF) to inform an LLM during fine-tuning about its proximity to achieving a perfect translation. However as CF and SF are non-differentiable functions, they cannot be directly used as loss functions to fine-tune an LLM. So, it is essential to construct an RL setting.

The scheme for jointly fine-tuning  $\text{LLM}_f$  and  $\text{LLM}_b$  is given in Algorithm 1. In short, we interleave RL-based fine-tuning by Proximal Policy Optimization (PPO) [38] and SFT-based optimization of CE loss by Adam optimizer. First,  $\text{LLM}_f$  ( $S \rightarrow T$ ) translates  $s$  to generate an intermediate  $T$ -sequence  $\hat{t}$ . With this as input,  $\text{LLM}_b$  ( $T \rightarrow S$ ) tries to reconstruct sequence  $s$  and generates a  $S$ -sequence  $\hat{s}$ . In RL-based fine-tuning, the reward for  $\text{LLM}_f$  and  $\text{LLM}_b$  is the sum of SF computed among  $s$  and  $\hat{s}$  and the respective CF among  $t, \hat{t}$  and  $s, \hat{s}$ . To ensure that the LLMs being fine-tuned do not diverge much from the reference LLMs we started with, a KL-divergence [30] ( $d_{\text{KL}}$ ) term is subtracted from the reward. This ensures that the PPO algorithm does not over-optimize and is appropriately penalized when the trained model starts to diverge too much from their references. Conversely, for the SFT-based optimization, the forward and backward models are fine-tuned by back-propagating the respective CE losses between the

#### Algorithm 1: RL+SFT Interleaved Training for CoTran

---

**Input** :  $M_f$  (forward LLM);  $M_b$  (backward LLM);  $Tok$  (kw-Tok);  $\tau_P(\cdot)$  (trainable param);  $D_{\text{trn}}, D_{\text{val}}$  (training & validation data)  
**Output** : Learned LLMs  $M_f$  (for  $S \rightarrow T$ ) and  $M_b$  (for  $T \rightarrow S$ )

---

```

1 Function RL_BK2BK ( $LLM_f, LLM_f^{\text{ref}}, LLM_b, LLM_b^{\text{ref}}, PPO_f, PPO_b, D$ ) :
2   for  $epoch \in [1, E]$  do
3     foreach  $(s, t) \in D$  do
4        $\hat{t} \leftarrow LLM_f(s); \hat{s} \leftarrow LLM_b(\hat{t})$ 
5        $r_f \leftarrow \omega_{CF}(t, \hat{t}) + \omega_{SF}(s, \hat{s}) - \beta \cdot d_{\text{KL}}(LLM_f, LLM_f^{\text{ref}})$ 
6        $r_b \leftarrow \omega_{CF}(s, \hat{s}) + \omega_{SF}(s, \hat{s}) - \beta \cdot d_{\text{KL}}(LLM_b, LLM_b^{\text{ref}})$ 
7       BackProp:  $LLM_f \leftarrow PPO_f(LLM_f, r_f);$  //  $\theta_f: \tau_P(LLM_f)$ 
8       BackProp:  $LLM_b \leftarrow PPO_b(LLM_b, r_b);$  //  $\theta_b: \tau_P(LLM_b)$ 
9   return  $LLM_f, LLM_b$ 

10 Function SFT_BK2BK ( $LLM_f, LLM_b, Adam_f, Adam_b, D$ ) :
11   for  $epoch \in [1, E]$  do
12     foreach  $(s, t) \in D$  do
13        $\hat{t} \leftarrow LLM_f(s); \hat{s} \leftarrow LLM_b(\hat{t})$ 
14       BackP:  $LLM_f \leftarrow Adam_f(LLM_f, \mathcal{L}_{CE}^{\theta_f}(t, \hat{t}));$  //  $\theta_f: \tau_P(LLM_f)$ 
15       BackP:  $LLM_b \leftarrow Adam_b(LLM_b, \mathcal{L}_{CE}^{\theta_b}(s, \hat{s}));$  //  $\theta_b: \tau_P(LLM_b)$ 
16   return  $LLM_f, LLM_b$ 

17 Function LoRA ( $LLM$ ) :
18    $LLM' \leftarrow \text{deep copy of } LLM;$  //  $\theta: \tau_P(LLM), \theta': \tau_P(LLM')$ 
19   foreach  $query/value \text{ layer } L_i \in \mathbb{R}^{d_{\text{in}} \times d_{\text{out}}}$  of  $LLM$  do
20     Initialize projection-up layer  $A_i \in \mathbb{R}^{d_{\text{in}} \times r}$  randomly from  $\mathcal{N}(0, 1)$ ,
21     projection-down layer  $B_i \leftarrow \{0\}^{r \times d_{\text{out}}}$ 
22     Replace  $L_i$  of  $LLM'$  with  $L_i + \frac{\alpha}{r}(A_i \times B_i)$ 
23    $\theta' \leftarrow \theta \cup A \cup B$ . Freeze params in  $\theta'$  except  $A, B$ 
24   return  $LLM'$ 

25 Function EVAL ( $LLM_f, LLM_b$ ) :
26   return  $\omega_{SF}(s, \hat{s})$  averaged over all  $s$  in the tokenized  $D_{\text{val}}$ ,
27   with  $LLM_f, LLM_b$  as back-to-back models

```

---

—MAIN FUNCTION—

```

27  $M'_f, M'_b \leftarrow \text{LoRA}(M_f), \text{LoRA}(M_b)$ 
28 Initialize RL Optimizers:  $PPO_f, PPO_b$  for  $M'_f, M'_b$ 
29 Initialize SFT Optimizers:  $Adam_f, Adam_b$  for  $M'_f, M'_b$ 
30  $D_{\text{trn}}^{Tok} \leftarrow [(Tok(s_{\text{code}}), Tok(t_{\text{code}})) \text{ for } (s_{\text{code}}, t_{\text{code}}) \text{ in } D_{\text{trn}}]$ 
31 do
32    $M_f^{\text{SFT}}, M_b^{\text{SFT}} \leftarrow \text{SFT\_BK2BK}(M'_f, M'_b, Adam_f, Adam_b, D_{\text{trn}}^{Tok})$ 
33    $M_f^{\text{RL}}, M_b^{\text{RL}} \leftarrow \text{RL\_BK2BK}(M'_f, M'_f, M'_b, M_b, PPO_f, PPO_b, D_{\text{trn}}^{Tok})$ 
34    $\text{valAcc}_{\text{SFT}}, \text{valAcc}_{\text{RL}} \leftarrow \text{EVAL}(M_f^{\text{SFT}}, M_b^{\text{SFT}}), \text{EVAL}(M_f^{\text{RL}}, M_b^{\text{RL}})$ 
35    $M'_f, M'_b \leftarrow (\text{valAcc}_{\text{RL}} \geq \text{valAcc}_{\text{SFT}}) ? M_f^{\text{RL}}, M_b^{\text{RL}} : M_f^{\text{SFT}}, M_b^{\text{SFT}}$ 
36 while  $\text{valAcc}_{\text{RL}}, \text{valAcc}_{\text{SFT}}$  converged or maxEpochs reached;

```

---

predicted translation and the corresponding gold-standard translation.

Further, as jointly fine-tuning both LLMs is resource-intensive we follow a parameter-efficient approach. For every linear layer of these LLMs that is either query or value, we create Low-Rank Adaptation [16] layers viz., a projection-up matrix  $A$  and projection-down matrix  $B$ , whose matrix-multiplication is initially zero. Except those in  $A$  and  $B$ , all the original parameters of  $\text{LLM}_f$  and  $\text{LLM}_b$  are frozen – making the fine-tuning process space-time efficient.

### 4 The AVATAR-TC Dataset

This paper introduces a new dataset **AVATAR-TC** (built on top of the AVATAR [2]) that has pairs of whole-programs in Java and Python (a statically- and dynamically-typed language, with different syntactic

**Table 1: Benchmark Suite:** Statistics of the AVATAR-TC dataset

| Sub-dataset    | # problem-stmts with test-cases |       |       | # pairs of Java-Python codes |       |       |
|----------------|---------------------------------|-------|-------|------------------------------|-------|-------|
|                | Train                           | Valid | Test  | Train                        | Valid | Test  |
| Aizu           | 762                             | 41    | 190   | 14,019                       | 41    | 190   |
| AtCoder        | 619                             | 19    | 97    | 13,558                       | 19    | 97    |
| Codeforces     | 1,625                           | 96    | 401   | 23,311                       | 96    | 401   |
| Google CodeJam | 59                              | 1     | 4     | 347                          | 1     | 4     |
| LeetCode       | 81                              | 7     | 18    | 81                           | 7     | 18    |
| GeeksForGeeks  | 3,753                           | 268   | 995   | 3,753                        | 268   | 995   |
| Project Euler  | 110                             | 11    | 41    | 110                          | 11    | 41    |
| Total          | 7,009                           | 443   | 1,746 | 55,179                       | 443   | 1,746 |

styles), each accompanied by human-written test-cases (TCs). Based on what we know, AVATAR-TC is the first such large-scale dataset where code compilability (syntactic correctness) is ensured, and code pairs have undergone thorough testing w.r.t. human-written TCs.

Note that these TCs are not used in training any of the CoTran variants (which rely on automated unit test generation); in this paper, we use them to evaluate translators with FEqAcc (Section 5.2).

**Data Source.** Similar to AVATAR, we gather a collection of code pairs written in Java and Python by scraping five *competitive coding websites* that host regular contests: Aizu [3], AtCoder [4], Codeforces [7], G-CodeJam [8], LeetCode [22] and two *coding platforms*: GeeksForGeeks [13], Project Euler [33]. These serve as great resources for mining code solutions of a problem statement across multiple languages. Also, a participant’s code gets checked on multiple test-cases (TCs) curated by the problem-setter. For AVATAR-TC, we web-crawled these data sources and collected such human-written TCs to complement each problem statement.

**Data Cleaning.** Several code pairs in AVATAR dataset did not compile and/or pass our collected TCs. Consequently in AVATAR-TC, we preprocessed codes afresh from their sources. Utilizing the `javalang` [42] and `tokenize` [34] modules, they were parsed into code-specific tokens. We manually corrected minor faults in code pairs, that did not match the expected output when provided with the TC inputs. Code pairs with major issues were discarded. Our criteria for output matching include case insensitivity, whitespace removal, punctuation disregard (unless significant to the output), and normalization of numeric or floating-point values to a common representation.

**Statistics of AVATAR-TC dataset.** The train/validation/test partitioning of problem statements is kept the same as AVATAR, except removal of some pairs during data cleaning. This resulted in 57,368 Java-Python pairs at a train : validation : test ratio of 76 : 5 : 19, across 9,198 problems. For the train split, at most 25 pairs correspond to one problem, while for validation and test, there is a unique one-one pair-problem mapping. To ensure *out-of-distribution testing*, no problem overlaps across splits. Refer to Table 1 for AVATAR-TC statistics.

## 5 Experiments

### 5.1 Experimental Setup and Competing Tools

**CoTran (baseline)** refers to an LLM fine-tuned without feedback loops (refer Figure 1a). We use the pre-trained CodeT5-base [47] architecture from Huggingface [17] and fine-tune it with CE loss, employing the proposed keyword-based tokenizer (kw-Tok). The maximum length for the source and target sequences is set at 512. For additional design choices adopted during implementing CoTran, please refer to Appendix A.2. The baseline CoTran and its variants are compared against (See Table 2): (a) three *human-written transpilers*, (b) three *SoTA LLM-based unsupervised translation tools* (trained on function pairs from ~2.5M open-sourced repositories of the GitHub dataset from Google BigQuery Public Datasets), (c) *ChatGPT* [29]

and, (d) seven *LLM-based supervised translation tools*. All the tools above are compared on the same set of 1,746 whole-programs from AVATAR-TC (Test). Additionally, each of the supervised tools and CoTran variants are fine-tuned on AVATAR-TC (Train). For ChatGPT, we use OpenAI API to access the `gpt-3.5-turbo-0301` model. This version of ChatGPT has a knowledge cutoff of March 1, 2023, which predates the public release of AVATAR-TC on GitHub. This minimizes the risk of AVATAR-TC (Test) pairs being included in ChatGPT’s training data, ensuring a fair evaluation. Following a standardized protocol [48] for ChatGPT, we use "Translate [S] to [T]: [s<sub>code</sub>]\n Do not return anything other than the translated code." as the prompt. Temperature and top\_p are set at 0 (this ensures reproducibility and does not notably alter the translation performance).

### 5.2 Evaluation Metrics

We evaluate our method using *greedy decoding* that considers only the top translation with the highest log probability. The different metrics for evaluating code translation quality are:

**BLEU, CodeBLEU score.** BLEU [31] computes the ‘closeness’ with the reference translation through n-gram overlaps. CodeBLEU [35] additionally checks weighted n-gram match, syntactic AST match (SM) and semantic data-flow match (DM). Both range in [0, 100].

**Exact String Match (EM).** EM is the percentage of generated codes that exactly match the reference translation. It can be low even if the generated codes are compilable and functionally equivalent.

**Compilation Accuracy (CompAcc) and Functional-Equivalence Accuracy (FEqAcc).** *CompAcc* is defined as the percentage of generated translations that compile correctly. *FEqAcc* is the percentage of generated translations that are IO equivalent to the source-language code w.r.t. a set of human-written test-cases.

Additionally, we propose three new quality measures, namely:

**Average First Error Position (errPos<sub>1st</sub>).** errPos<sub>1st</sub> is a fine-grained version of *CompAcc*, relating to the closeness of the translations from a perfect compilation. Averaged over all translations on test set, errPos<sub>1st</sub>( $\hat{\tau}$ ) calculates the position of the first token responsible for a syntactic error in  $\hat{\tau}$ , normalized by  $|\hat{\tau}|$ . It is computed by `pylint` and `javac` for Python and Java respectively. Let’s consider that a translator achieves errPos<sub>1st</sub> =  $e\%$ . This implies that, on average, the first compilation error (if any) is located within the last  $(100 - e)\%$  portion of the generated translations. As  $e$  approaches 100, the human developer only needs to inspect a small section of each translated code to rectify the error, thereby facilitating ease of manual debugging.

**Average #Errors per Code (EpC).** EpC is the average count of compilation (syntactic) errors per translated code. Zero errors indicate full syntactic correctness. A higher count implies that substantial effort from a human end-user is required for rectification.


**Ratio of FEqAcc and CompAcc ( $\frac{f}{c}$ rate).**  $\frac{f}{c}$ rate is the percentage rate at which a translator generates functionally-equivalent codes compared to the compilable ones. In a real-world deployment, checking equivalence of translator-generated codes is infeasible. Thus, it is beneficial when a syntactically correct code produced by a translator implies functional equivalence i.e., ideally  $\frac{f}{c}$ rate tends to 100%.  $\frac{f}{c}$ rate is computed as the percentage ratio of FEqAcc and CompAcc.

### 5.3 Analysis of Empirical Results and Ablation Study

In Table 2, we compare CoTran against 14 competing tools for J2P and P2J translation. PLBART-base, CodeT5-base, and PPOCoder perform best among the competitors, while transpilers underperform (refer

**Table 2: Code Translation Results:** Comparison of CoTran against 14 other tools for Java-Python (J2P) and Python-Java (P2J) translation. (In each column, the highest value is marked in **bold**, second-highest underlined.)

| Method / Tool  | Model  | Java → Python (J2P) |              |                       |              |              |             | Python → Java (P2J) |              |                       |              |              |             |
|--|--|---------------------|--------------|-----------------------|--------------|--------------|-------------|---------------------|--------------|-----------------------|--------------|--------------|-------------|
|  |  | FEqAcc              | CompAcc      | errPos <sub>1st</sub> | CodeBLEU     | BLEU         | EM          | FEqAcc              | CompAcc      | errPos <sub>1st</sub> | CodeBLEU     | BLEU         | EM          |
| Transpilers  | java2python [28]                             | 3.32                | 41.46        | 28.62                 | 20.31        | 17.54        | 0           | -                   | -            | -                     | -            | -            | -           |
|  | TSS CodeConv [43]                            | 0.46                | 58.30        | 54.26                 | 41.87        | 24.44        | 0           | -                   | -            | -                     | -            | -            | -           |
|  | py2java [10]                                 | -                   | -            | -                     | -            | -            | -           | 0                   | 0            | 1.61                  | 41.56        | 48.59        | 0           |
| Recent competing tools (unsupervised trng.)            | TransCoder [36]                              | 0.46                | 88.09        | 63.57                 | 35.07        | 32.07        | 0           | 0                   | 0            | 4.57                  | 35.02        | 35.06        | 0           |
|  | TransCoder-DOBF [20]                         | 0.46                | 63.00        | 47.10                 | 39.98        | 33.84        | 0           | 0                   | 0            | 3.11                  | 33.33        | 32.72        | 0           |
|  | TransCoder-ST [37]                           | 0.46                | 91.58        | 74.68                 | 40.04        | 37.30        | 0           | 0                   | 0            | 4.67                  | 29.88        | 28.15        | 0           |
|  | ChatGPT                                      | <b>76.06</b>        | 95.36        | 90.88                 | 52.11        | 53.19        | 0.29        | 21.65               | 24.97        | 30.86                 | 54.08        | 55.58        | 0           |
| Recent competing tools (supervised trng. on AVATAR-TC) | CodeBERT [9]                                 | 12.31               | 84.77        | 79.57                 | 46.00        | 48.10        | 0.46        | 0.74                | <b>96.79</b> | <b>99.51</b>          | 26.10        | 19.62        | 0           |
|  | GraphCodeBERT [15]                           | 10.88               | 85.05        | 79.78                 | 45.53        | 47.26        | 0.57        | 0.46                | <u>89.75</u> | <u>98.05</u>          | 23.72        | 16.21        | 0           |
|  | CodeGPT [25]                                 | 24.86               | 78.92        | 89.21                 | 38.38        | 38.64        | 1.49        | 13.40               | 45.13        | 94.50                 | 40.51        | 37.96        | 0.52        |
|  | CodeGPT-adapted [25]                         | 24.17               | 76.75        | 89.31                 | 36.84        | 37.36        | 1.55        | 20.50               | 52.00        | 97.60                 | 41.46        | 38.15        | 1.03        |
|  | PLBART-base [1]                              | 38.55               | 91.47        | 90.79                 | 54.77        | 59.34        | 1.32        | 38.26               | 75.77        | 96.64                 | 55.96        | 59.24        | 0.97        |
|  | CodeT5-base [47]                             | 40.95               | 92.84        | <b>93.76</b>          | 55.34        | 60.03        | 2.41        | 33.79               | 68.84        | 98.02                 | 57.64        | 60.16        | 0.86        |
|  | PPOCoder [39]                                | 44.27               | 93.47        | 91.44                 | 55.16        | 59.51        | 1.89        | 37.11               | 59.62        | 96.77                 | 55.04        | 58.52        | 0.52        |
| Our tool   | <b>CoTran (baseline)</b>                     | 44.52               | 96.12        | 92.07                 | 55.44        | 58.71        | 2.11        | 40.41               | 73.63        | 92.16                 | <b>59.11</b> | 61.12        | <b>1.66</b> |
| Our tool with CF only                                  | CoTran + CF (RL-based fine-tuning)           | 47.02               | 96.56        | 91.58                 | 56.10        | 60.59        | <u>2.23</u> | 42.78               | 74.80        | 96.91                 | 58.55        | <u>61.26</u> | <u>1.60</u> |
|  | <b>CoTran + CF (RL+SFT interleaved)</b>      | 49.83               | 96.79        | 92.08                 | 56.07        | 60.61        | <u>2.23</u> | <u>45.93</u>        | 75.77        | 96.89                 | 58.28        | 61.21        | <u>1.60</u> |
| Our tool (b2b) with CF, SF                             | CoTran + CF + SF (RL-based fine-tuning)      | 50.45               | <u>96.79</u> | 92.15                 | <u>56.17</u> | 60.60        | <u>2.23</u> | 43.92               | 75.14        | 96.93                 | <u>58.59</u> | <b>61.28</b> | <u>1.60</u> |
|  | <b>CoTran + CF + SF (RL+SFT interleaved)</b> | <u>53.89</u>        | <b>97.14</b> | <u>92.73</u>          | <b>56.24</b> | <b>60.69</b> | <b>2.29</b> | <b>48.68</b>        | 76.98        | 96.93                 | 58.38        | 61.19        | <u>1.60</u> |

**Figure 2: Reward Analysis:** Plot of different rewards for RL

Appendix A.4). ChatGPT exhibits weak performance in P2J with a significant shortfall in FEqAcc (−27.03%), CompAcc (−52.01%), and errPos<sub>1st</sub> (−66.07%) compared to the best CoTran method. A low errPos<sub>1st</sub> indicates that ChatGPT-generated translations are hard to debug. But, in J2P, ChatGPT leads with a FEqAcc of 76.06%. Given that CoTran is built on CodeT5-base with ~220M parameters, we believe that fine-tuning a model as large as ChatGPT (rumored to have 100B+ parameters) using our method would outperform ChatGPT in J2P across all metrics. However, while ChatGPT supports fine-tuning, it is closed-source and its interface is not yet tailored to receive symbolic feedback from tools such as compilers, testers, and solvers. Hereafter, we report CoTran improvements w.r.t. CodeT5-base.

**Hypothesis:** LLMs are good at code translation. Having said that, incorporating compiler and symexec feedback (CF, SF) during fine-tuning significantly improves its capability of producing compilable and functionally equivalent translations.

Our results validate this hypothesis when evaluated across the diverse AVATAR-TC dataset. In the process, we have made several findings:

**Finding 1:** RL turned out to be much more effective than the SFT schemes we tried, for incorporating feedback during fine-tuning.

We considered two non-RL schemes of SFT-based LLM training, combining CE loss and CF, SF. However, they underperform compared to the RL-based methods. Please refer to Appendix A.3 for more details.

**Finding 2:** For RL reward, a Boolean feedback from compiler and other existing feedback are not as effective as ours.

For RL-based optimization of LLMs, it is essential to fabricate a good reward function. CompCoder [46] attempts RL-based code generation using a Boolean feedback  $\text{compiler}_{\text{Bool}}(\hat{t})$ , that returns −1 or +1.

PPOCoder [39] uses the sum of  $\text{compiler}_{\text{Bool}}(\hat{t})$ , syntactic match score  $\text{SM}(\hat{t}, t)$  and dataflow match score  $\text{DM}(\hat{t}, t)$  as the RL reward. We hypothesize that even though these functions are good tools to compare  $\hat{t}$  with  $t$ , they are not the best when it comes to an RL reward.

In Figure 2,  $t$  is a tokenized Java code for reversing an integer using for-loop, while  $\hat{t}$  is its truncated version. We vary  $|\hat{t}|/|t|$  from 0 to 1 and plot the respective reward. For a significant portion of the  $x$ -axis, especially when  $\hat{t}$  is near empty or is almost same as  $t$ , values of  $\text{compiler}_{\text{Bool}}(\hat{t})$ ,  $\text{SM}(\hat{t}, t)$  and  $\text{DM}(\hat{t}, t)$  remain constant. So, CompCoder and PPOCoder offer the same RL reward for several closely-related translations. This limits the RL agent’s ability to gauge improvement. In contrast, our proposed reward  $\omega_{CF}(\hat{t}, t)$  better detects small changes, guiding the RL agent towards smaller goals and improving the overall compilation accuracy. Thus, CoTran + CF (RL only) outperforms PPOCoder, despite both using CodeT5-base.


**Finding 3:** Interleaving RL and SFT improves the LLM’s performance, compared to an RL-only fine-tuning approach.

To prevent RL-based fine-tuning from deviating the back-to-back LLMs from their CE loss objective, we occasionally interleave it with supervised fine-tuning (SFT). This improves the overall translation performance. In order to incorporate CF (CoTran + CF), we fine-tune the baseline model in an RL setting to maximize CF as a reward (refer Figure 1b). This approach improves J2P and P2J translation by +6.07%, +8.99% in FEqAcc, and +3.72%, +5.96% in CompAcc, respectively. The RL+SFT interleaved training boosts these gains to +8.88%, +12.14% in FEqAcc and +3.95%, +6.93% in CompAcc. Note that for CoTran + CF, the interleaved training is similar to Algorithm 1 with the difference that  $\omega_{SF}$  is not considered (there are no b2b LLMs, only the forward LLM). When both CF and SF are incorporated (CoTran + CF + SF) using RL with b2b LLMs (Figure 1c), the improvements are +9.50%, +10.13% in FEqAcc and +3.95%, +6.30% in CompAcc. Interleaving RL+SFT (Figure 1d and Algorithm 1) further increases these to +12.94%, +14.89% in FEqAcc and +4.30%, +8.14% in CompAcc. Additionally, CoTran + CF + SF achieves 92.73% and 96.93% on errPos<sub>1st</sub> for J2P and P2J, thus signifying easy-to-debug translations.

**Finding 4:** Adding keywords to tokenizer vocabulary by kw-Tok improves the code translation performance.

For J2P and P2J translation respectively, kw-Tok itself accounts for +3.57%, +6.62% increase on FEqAcc and, +3.28%, +4.79% increase on CompAcc (CoTran baseline v/s CodeT5-base).





**Figure 3: Count of compilation errors per code (EpC):** #Errors per translated code in P2J translation, sorted in the ascending order

**Finding 5:** *Incorporating compiler and symexec feedback during fine-tuning reduces the count of syntactic errors per translated code.* In Figure 3, we compare the number of compilation (syntactic) errors per code (EpC) produced in Python-to-Java (P2J) translation on AVATAR-TC (Test). ChatGPT exhibits significantly poorer performance with an EpC of  $10.98 \pm 9.67$  (i.e., mean  $\pm$  stdev). CodeT5-base reduces this to  $0.95 \pm 2.58$ . Our CoTran + CF + SF (RL+SFT) further lowers the EpC value to  $0.71 \pm 2.15$ . This indicates that the proposed feedback results in producing more easy-to-debug translations.

**Finding 6:** *Integrating compiler and symexec feedback during fine-tuning markedly improves  $\frac{f}{c}$  rate metric across diverse codebases.* In Figure 4, we compare the P2J performance of four SoTA tools and three CoTran-based methods over the different sub-datasets of AVATAR-TC (Test). CodeBERT demonstrates the highest rate of compilable translations (with over 90% CompAcc across most sub-datasets). However, the resultant translations lack meaningfulness, with only around 1% of them being functionally equivalent. In contrast, CodeGPT, PLBART-base, and CodeT5-base yield functionally equivalent codes at rates of 29.69%, 50.49%, and 49.08% respectively, relative to the number of compilable translations they generate across all sub-datasets. This  $\frac{f}{c}$  rate increases to 63.24% for CoTran + CF + SF (RL+SFT). Our method consistently *ranks among the top two* in  $\frac{f}{c}$  rate for each sub-dataset. Note that G-CodeJam is a small sub-dataset, and all tools exhibit similarly poor performance when translating P2J.

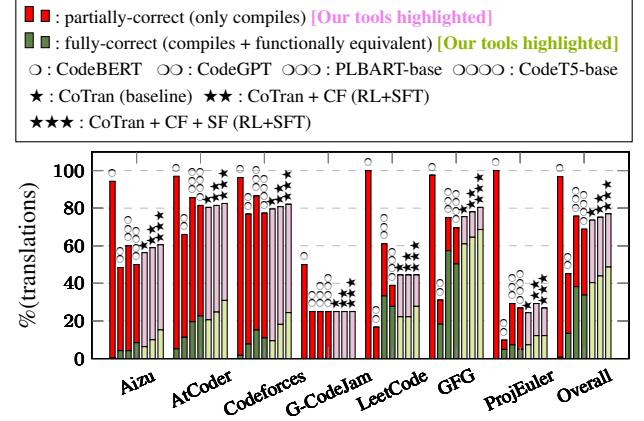
**Finding 7:** *It is easier for a code translator to achieve good performance if the source language is statically-typed.*

Java is a statically-typed language, while Python is dynamically-typed. When translating J2P, Java’s explicit variable-type declarations are ignored as they are redundant in Python. Conversely, in Python, variable-types are not explicitly declared, requiring the translator to infer them during P2J translation. Even for humans experienced in both languages, deducing types for variables in Python can be challenging and demand repeated engrossed mental evaluations. Consequently, learning P2J is much more challenging than J2P. Thus, the source language being statically-typed makes learning code translation easier.

**Finding 8:** *Function-to-function translation is not sufficient for whole-program translation tasks.*

The TransCoder-based unsupervised tools are trained on function-level translations. As per Roziere et al. [36], this keeps training batches shorter and unit test-based model evaluation simpler. But consequently, these tools cannot efficiently generate compilable, functionally-equivalent translations for whole-programs. For instance, in J2P, all three TransCoder-based tools fail drastically in FEqAcc (0.46%). Similarly, for P2J, no (0%) translated code is compilable.

**Benefits of automatic test-case generation (TCgen) during training.** During fine-tuning  $S \rightarrow T \rightarrow S$  back-to-back LLMs (Section 3.2), we automatically generate unit tests on  $s$  and test them on  $\hat{s}$ . Note that



**Figure 4: Sub-dataset wise performance w.r.t.  $\frac{f}{c}$  rate:** Comparison of P2J translation, across the 7 sub-datasets of AVATAR-TC (Test)

equivalence checking of programs is an undecidable problem. Our idea is to assess the *inequivalence* of  $s$  and  $\hat{s}$  (detectable via sufficient testing) w.r.t. a test suite produced by TCgen tool and thereby, to calculate symexec feedback (SF) for the LLM during fine-tuning. To our knowledge, no existing LLM-based code translation tool employs *automatic symexec-based TCgen* for functional equivalence checking and feedback. All existing tools like PPOCoder [39], TransMap [45], RLTF [24], CodeRL [21] use *human-written* test-cases for fixing LLM-generated codes. In contrast, CoTran does not require anything extra than what is required for a CE loss-based supervised fine-tuning of LLM i.e., a dataset of  $S, T$  code pairs. Also, compared to human-written test-cases, symexec-based testing has a higher probability of covering all linearly independent control-flow paths.

**Generalizability regarding choice of TCgen tool.** We want to emphasize that our approach is entirely agnostic to any specific TCgen tool, including Symflower. The selection of Symflower is driven by the fact that we tested CoTran with Java as  $S$  in the  $S \rightarrow T \rightarrow S$  back-to-back training loop, and Symflower stands out as a commercial industrial-strength symexec engine for Java. Its efficiency and reliability make it a preferred choice. MLB, JBMC, and GDart (the top 3 verification tools for Java in SV-COMP 2024 [5]) can also be considered as potential alternatives to Symflower. In fact, similar efficient TCgen tools (e.g., EvoSuite [11] for Java, KLOVER [23] for C++, Pynguin [26] for Python) are available for most popular languages. If not, LLVM-based TCgen tools (e.g., KLEE) offer an alternative, as many languages can be readily translated into LLVM IR.

**Reproducibility and Supplementary.** The AVATAR-TC dataset and all our code can be accessed at <https://github.com/PrithwishJana/CoTran>. The supplementary material (Appendix) is included in the full version [19] of this paper, available on arXiv.

## 6 Conclusion and Future Work

In this paper, we present an LLM-based code translation method (CoTran) that incorporates feedback from compiler and symexec-based solver during fine-tuning. The paper showcases the power of symbolic feedback in making code translation LLMs more accurate. Another key insight is the efficacy of fine-tuning LLMs with fine-grained feedback, pinpointing proximity to an ideal solution, rather than simple Boolean directives like yes/no. Our results show that CF and SF, especially with RL+SFT interleaved training, significantly improve code translation quality. CoTran outperforms state-of-the-art tools in producing compilable and functionally equivalent codes. We plan to extend CoTran to translate legacy code to modern languages.

## References

- [1] W. Ahmad, S. Chakraborty, B. Ray, and K.-W. Chang. Unified Pre-training for Program Understanding and Generation. In *Proc. 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2655–2668, Online, June 2021. ACL. doi: 10.18653/v1/2021.naacl-main.211.
- [2] W. U. Ahmad, M. G. R. Tushar, S. Chakraborty, and K.-W. Chang. AVATAR: A Parallel Corpus for Java-Python Program Translation. In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 2268–2281. Association for Computational Linguistics, 2023.
- [3] Aizu-Online-Judge. Programming Challenge, 2023. <https://judge.u-aizu.ac.jp/onlinejudge/>.
- [4] AtCoder, 2023. <https://atcoder.jp/>.
- [5] D. Beyer. State of the Art in Software Verification and Witness Validation: SV-COMP 2024. In *International Conference on Tools and Algorithms for the Construction and Analysis of Systems*, pages 299–329. Springer, 2024.
- [6] X. Chen, C. Liu, and D. Song. Tree-to-Tree Neural Networks for Program Translation. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 31, pages 2552 – 2562, 2018.
- [7] Codeforces, 2023. <https://codeforces.com/>.
- [8] CodeJam. Google’s Coding Competitions, 2023. <https://codingcompetitions.withgoogle.com/codejam>.
- [9] 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 *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1536–1547. ACL, 2020.
- [10] N. Fomin. py2java: Python to Java Language Translator, 2019. <https://pypi.org/project/py2java/>.
- [11] G. Fraser and A. Arcuri. EvoSuite: Automatic Test Suite Generation for Object-Oriented Software. In *Proceedings of the 19th ACM SIGSOFT Symposium and the 13th European Conference on Foundations of Software Engineering*, pages 416–419, 2011.
- [12] V. Ganesh, S. A. Seshia, and S. Jha. Machine Learning and Logic: A New Frontier in Artificial Intelligence. *Formal Methods in System Design*, 60 (3):426–451, 2022.
- [13] GeeksforGeeks. A CS Portal for Geeks, 2023. <https://www.geeksforgeeks.org/>.
- [14] M. Grimmer, R. Schatz, C. Seaton, T. Würthinger, M. Luján, and H. Mössenböck. Cross-language Interoperability in a Multi-language Runtime. *ACM Transactions on Programming Lang. and Systems (TOPLAS)*, 40(2):1–43, 2018.
- [15] D. Guo, S. Ren, S. Lu, Z. Feng, D. Tang, S. Liu, L. Zhou, N. Duan, A. Svyatkovskiy, et al. GraphCodeBERT: Pre-training Code Representations with Data Flow. In *9th International Conference on Learning Representations (ICLR)*, 2021.
- [16] E. J. Hu, Y. Shen, P. Wallis, Z. Allen-Zhu, Y. Li, S. Wang, L. Wang, and W. Chen. LoRA: Low-Rank Adaptation of Large Language Models. In *10th International Conference on Learning Representations (ICLR)*, 2022.
- [17] HuggingFace. The AI Community Building the Future, 2023. <https://huggingface.co/>.
- [18] P. Jana. NeuroSymbolic LLM for Mathematical Reasoning and Software Engineering. In *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence (IJCAI)*, pages 8492–8493. IJCAI Organization, 2024. Doctoral Consortium.
- [19] P. Jana, P. Jha, H. Ju, G. Kishore, A. Mahajan, and V. Ganesh. Co-Tran: An LLM-based Code Translator using Reinforcement Learning with Feedback from Compiler and Symbolic Execution. *arXiv preprint arXiv:2306.06755*, 2023. Full version of this paper, including Appendix.
- [20] M.-A. Lachaux, B. Roziere, M. Szafrańiec, and G. Lample. DOBF: A Deobfuscation Pre-training Objective for Programming Languages. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 14967–14979, 2021.
- [21] H. Le, Y. Wang, A. D. Gotmare, S. Savarese, and S. C. H. Hoi. CodeRL: Mastering Code Generation through Pretrained Models and Deep Reinforcement Learning. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 21314–21328, 2022.
- [22] LeetCode, 2023. <https://leetcode.com/>.
- [23] G. Li, I. Ghosh, and S. P. Rajan. KLOVER: A Symbolic Execution and Automatic Test Generation Tool for C++ Programs. In *Proc. 23rd International Conference on Computer Aided Verification (CAV)*, pages 609–615. Springer, 2011.
- [24] J. Liu, Y. Zhu, K. Xiao, Q. Fu, X. Han, W. Yang, and D. Ye. RLTF: Reinforcement Learning from Unit Test Feedback. *Transactions on Machine Learning Research*, 2023.
- [25] S. Lu, D. Guo, S. Ren, J. Huang, A. Svyatkovskiy, A. Blanco, C. Clement, D. Drain, D. Jiang, D. Tang, et al. 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.
- [26] S. Lukaszczuk and G. Fraser. Pyguin: Automated Unit Test Generation for Python. In *Proc. ACM/IEEE 44th International Conference on Software Engineering: Companion Proceedings*, pages 168–172, 2022.
- [27] B. G. Mateus, M. Martinez, and C. Kolski. Learning Migration Models for Supporting Incremental Language Migrations of Software Applications. *Information and Software Tech.*, 153, 2023.
- [28] T. Melhase, B. Kearns, L. Li, I. Curt, and S. Saladi. java2python: Simple but Effective Tool to Translate Java Source Code into Python, 2016. <https://github.com/natural/java2python>.
- [29] OpenAI. ChatGPT [Large Language Model], 2023. <https://chat.openai.com>.
- [30] D. Palenicek. A Survey on Constraining Policy Updates using the KL Divergence. *Reinforcement Learning Algorithms: Analysis and Applications*, pages 49–57, 2021.
- [31] K. Papineni, S. Roukos, T. Ward, and W.-J. Zhu. BLEU: A Method for Automatic Evaluation of Machine Translation. In *Proc. 40th Annual Meeting: Assoc. for Computational Linguistics*, pages 311–318, 2002.
- [32] T. Parr, T. Everett, et al. Grammars-v4: Grammars written for ANTLR (ANother Tool for Language Recognition) v4, 2023. <https://github.com/antlr/grammars-v4>.
- [33] ProjectEuler, 2023. <https://projecteuler.net/>.
- [34] Python. tokenize: Tokenizer for Python, 2023. <https://docs.python.org/3/library/tokenize.html>.
- [35] S. Ren, D. Guo, S. Lu, L. Zhou, S. Liu, D. Tang, N. Sundaresan, M. Zhou, A. Blanco, and S. Ma. CodeBLEU: A Method for Automatic Evaluation of Code Synthesis. *arXiv preprint arXiv:2009.10297*, 2020.
- [36] B. Roziere, M.-A. Lachaux, L. Chaussonot, and G. Lample. Unsupervised Translation of Programming Languages. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 20601–20611, 2020.
- [37] B. Roziere, J. Zhang, F. Charton, M. Harman, G. Synnaeve, and G. Lample. TransCoder-ST: Leveraging Automated Unit Tests for Unsupervised Code Translation. In *10th International Conference on Learning Representations (ICLR)*, 2022.
- [38] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov. Proximal Policy Optimization Algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- [39] P. Shojaei, A. Jain, S. Tipirneni, and C. K. Reddy. Execution-based Code Generation using Deep Reinforcement Learning. *Transactions on Machine Learning Research*, 2023. ISSN 2835-8856.
- [40] A. Svyatkovskiy, S. K. Deng, S. Fu, and N. Sundaresan. IntelliCode Compose: Code Generation using Transformer. In *Proc. 28th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering*, pages 1433–1443, 2020.
- [41] S. Thenault. pylint: Python code Static Checker, 2023. <https://pypi.org/project/pylint/>.
- [42] C. Thunes. javalang: Pure Python Java parser and tools, 2020. <https://github.com/c2nes/javalang>.
- [43] TSS. The Most Accurate and Reliable Source Code Converters, 2023. (Tangible Software Solns.): <https://www.tangiblesoftwaresolutions.com/>.
- [44] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin. Attention Is All You Need. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 30, 2017.
- [45] B. Wang, R. Li, M. Li, and P. Saxena. TransMap: Pinpointing Mistakes in Neural Code Translation. In *Proceedings of the 31st ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering*, pages 999–1011, 2023.
- [46] X. Wang, Y. Wang, Y. Wan, F. Mi, Y. Li, P. Zhou, J. Liu, H. Wu, X. Jiang, and Q. Liu. Compilable Neural Code Generation with Compiler Feedback. In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 9–19, 2022.
- [47] 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 *Proc. 2021 Conference on Empirical Methods in NLP*, pages 8696–8708. ACL, 2021. doi: 10.18653/v1/2021.emnlp-main.685.
- [48] W. Yan, Y. Tian, Y. Li, Q. Chen, and W. Wang. CodeTransOcean: A Comprehensive Multilingual Benchmark for Code Translation. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 5067–5089, 2023.
- [49] M. Zimmermann and E. Haslinger. Symflower: Smart Unit Test Generator for Java, 2023. <https://symflower.com/en/>.