

---

# 000 T2J: LEVERAGING DEVELOPER BUG-FIXING BE- 001 HAVIORS TO EVALUATE AND IMPROVE LLM-BASED 002 PYTORCH-TO-JAX TRANSLATION 003 004

006 **Anonymous authors**

007 Paper under double-blind review  
008  
009  
010

## 011 ABSTRACT 012

013 While Large Language Models (LLMs) have shown strong performance in code-  
014 to-code translation for widely-used programming languages, their application to  
015 PyTorch-to-JAX translation remains challenging. Although both frameworks are  
016 implemented in Python, they differ fundamentally in design principles, execu-  
017 tion models, and ecosystem maturity, with JAX being relatively new and under-  
018 represented in public code repositories. Moreover, the lack of parallel PyTorch-  
019 JAX datasets and the limitations of existing evaluation metrics hinder effective  
020 cross-framework translation. In this work, we propose T2J , a prompt aug-  
021 mentation framework aimed at improving LLM-based PyTorch-to-JAX transla-  
022 tion. First, we construct two PyTorch code datasets, the problem solving code  
023 dataset collected from *TorchLeet* (Aroori & Chien, 2025) repository and the  
024 Github code dataset collected from *CodeParrot* benchmark (Wolf et al., 2022),  
025 leveraging the cheap LLM 4o-mini to generate initial translations. Second, we  
026 employ two professional developers to iteratively fix the generated JAX code un-  
027 til it is functionally equivalent to the original PyTorch code, resulting in a cu-  
028 rated *fixed-bug dataset* that captures common translation errors and their corre-  
029 sponding fixes. Third, we design augmented prompts that incorporate structured  
030 guidance from the fixed-bug dataset to improve translation quality of lightweight  
031 LLMs as GPT-4o-mini. Finally, we take advantages of using LLM as a judge  
032 and using LLM to measure the scale of each bug fixing step to propose three  
033 evaluation metrics for Pytorch-to-JAX code translation: T2J\_CodeTrans\_Score,  
034 T2J\_FixCost\_Score, and T2J\_Comparison\_Score. Our results demonstrate that  
035 T2J significantly improves GPT-4o-mini performance by up to **10%** in Code-  
036 BLEU, **50%** in T2J\_FixCost\_Score, **1.33 point** in T2J\_CodeTrans\_Score (as scale  
037 of 0-4), and **100%** in T2J\_Comparison\_Score. T2J’s generated code can improve  
038 2.5 faster in running time compared to the baseline’s output execution. Replication  
039 package is available at: <https://tinyurl.com/yradutma>.  
040

## 041 1 INTRODUCTION 042

043 Code translation involves converting a program from one programming language to another while  
044 preserving the original functionality. This process is useful for cross-language and cross-domain  
045 migration, allowing organizations to transition their code base to more modern languages or to  
046 various purposes. It also supports the modernization of legacy systems by re-implementing them in  
047 languages that promote greater maintainability and scalability as part of system refactoring efforts .  
Furthermore, in enterprises that employ multiple programming languages, code translation enhances  
the productivity of the programmer.

048 However, recent research works (Pan et al., 2024; Dou et al., 2024) indicate that LLMs-generated  
049 programs in the target language continue to encounter various quality problems, including compi-  
050 lation errors or functional inconsistencies. These challenges become even more pronounced in spe-  
051 cialized contexts such as domain-specific language translation. For example, (TehraniJamsaz et al.,  
052 2024) demonstrates that their transformer-based approach, CodeRosetta, outperforms well-known  
053 LLMs in C-to-CUDA translation. Our work, by contrast, focuses on another domain-specific transla-  
tion problem: translating between different Python libraries—specifically, converting PyTorch code

---

054 snippets into their JAX equivalents. Unlike CodeRosetta, which operates on translation between dif-  
055 ferent programming languages, PyTorch-to-JAX translation cannot be reduced to Abstract Syntax  
056 Tree transformations. Relying on low-cost LLMs therefore introduces a significant risk of generat-  
057 ing poor-quality translations. This risk arises because JAX, a framework designed for parallelization  
058 across diverse hardware platforms, is far less familiar to the broader developer community. Conse-  
059 quently, during PyTorch-to-JAX migration, LLMs often struggle to generate correct JAX code due  
060 to their limited exposure to JAX, which is a comparatively newer ecosystem than PyTorch.

061 To address these challenges, we introduce T2J, a in-context code learning and code evaluation  
062 framework designed to enhance LLM-based PyTorch-to-JAX code translation by leveraging cu-  
063 rated datasets and structured prompting strategies. This framework proceeds in several key stages:  
064 first, we construct parallel corpora of Pytorch and JAX corresponding code snippets from estab-  
065 lished PyTorch datasets, in particular TorchLeet and CodeParrot (Wolf et al., 2022; Aroori & Chien,  
066 2025). Then we employ high-quality GPT models GPT-4o to produce initial JAX translations. Sub-  
067 sequently, professional human developers iteratively refine the translated JAX program to achieve  
068 functional equivalence with the original PyTorch input, producing a curated fixed-bug dataset that  
069 systematically documents prevalent translation errors. The dataset, called bug fixing dataset, also in-  
070 cludes error-by-error fix instructions as multiple fixing steps. Building on this, we design augmented  
071 prompts that integrate targeted, structured guidance derived from the fixed-bug dataset. Finally, we  
072 evaluate T2J’s performance across both datasets using the CodeBLEU (Ren et al., 2020), along-  
073 side three novel metrics: T2J\_CodeTrans\_Score (assessing the usefulness and functional correctness  
074 of LLMs using LLMs as a judge), T2J\_FixCost\_Score (quantifying the effort required for post-  
075 translation corrections), and T2J\_Comparison\_Score (measuring semantic and functional alignment  
076 through differential analysis) to provide a comprehensive assessment of translation quality. Our  
077 contributions are as follows:

- 078 1. The creation of the first fixed-bug dataset specifically for PyTorch-to-JAX code translation,  
079 encompassing detailed annotations of error patterns and fixes to facilitate improvements in  
080 reliability on LLM code translation.
- 081 2. The T2J framework, which innovates prompt augmentation techniques to bridge domain-  
082 specific gaps in cross-library code migration problem as Pytorch-to-JAX translation.
- 083 3. The evaluation framework for Pytorch-to-JAX translation, which compare the source/ tar-  
084 get code with LLM-prompting techniques and fixing cost from error to correct programs  
085 by human bug fixing process.

086 The remainder of this paper is organized as follows: Section 2 reviews related work, Section 3  
087 describe the Motivation Example, Section 4 introduces the components of T2J, Section 5 presents  
088 experimental results, and the final sections discuss limitations, conclusion, and future directions.  
089

## 090 2 RELATED WORK 091

092 **Challenges of Code Translation.** One of the important challenges of code translation is to provide  
093 a metric for comparison between predicted and expected code. Research works show that just sim-  
094 ply comparing source code by traditional textual similarity scores is not efficient Tran et al. (2019).  
095 Instead, code metrics that included information of syntactic/ semantic similarity between code snip-  
096 plets have been proposed (Zhou et al., 2023). Another challenge is that collecting parallel corpus for  
097 code translation is very expensive and require human effort for verification (Husain et al., 2020). To  
098 evaluate unsupervised code translation’s output, automated test cases generation approaches have  
099 been proposed (Roziere et al., 2022; Peng et al., 2024). Finally, studies show that there are many  
100 types of bugs extracted from LLMs’ generated code (Dinh et al., 2023; Zhang et al., 2024).

101 **Machine Learning-based Approaches for Code Translation.** Supervised methods for code trans-  
102 lation are typically trained on well-established datasets such as CodeXGLUE (Lu et al., 2021).  
103 Among these methods, BERT-based models have proven particularly effective not only for code-to-  
104 code translation but also for a wide range of code generation tasks (Guo et al., 2021; 2022; Wang  
105 et al., 2023; Ahmad et al., 2021). Unsupervised code translation typically relies on transforming  
106 the source code into an intermediate representation (IR), followed by learning to generate target  
107 language code from that intermediate form. Szafraniec et al. (2023) proposed *Transcoder-IR*, a  
system that uses IR as a pivot language to translate between widely-used languages such as Java,

---

108 109 110 111 112 113 114 115 116 117 118 119 120 121 122 123 124 125 126 127  128 129 130 131	<b>(a) Input PyTorch Code</b> <pre> 1 import torch 2 import torch.nn as nn 3 import torch.nn.functional as F 4 5 # Define model 6 class SimpleNN(nn.Module): 7     def __init__(self): 8         super(SimpleNN, self).__init__() 9         self.fc = nn.Linear(28 * 28, 10) 10 11     def forward(self, x): 12         x = x.view(-1, 28 * 28) 13         # Flatten 14         return self.fc(x) 15 16     # Example 17 model = SimpleNN() 18 input_tensor = torch.randn(1, 1, 28, 28) 19 output = model(input_tensor) 20 print(output) </pre>	<b>(b) Incorrect JAX code</b> <pre> 1 import jax 2 import jax.numpy as jnp 3 from jax import random 4 5 # Initialize parameters 6 def init_params(key): 7     w_key, b_key = random.split(key) 8     # Wrong shape 9     W = random.normal(w_key, (10, 28*28)) * jnp.sqrt(1 / (28*28)) 10    b = jnp.zeros((10,)) 11    return {'W': W, 'b': b} 12 13 # Forward function 14 def forward(params, x): 15     x = x.reshape(x.shape[0], -1) 16     return jnp.dot(x, params['W']) + params['b'] 17 18 # Example 19 key = random.PRNGKey(0) 20 params = init_params(key) 21 input_tensor = random.normal(key, (1, 1, 28, 28)) 22 output = forward(params, input_tensor) 23 print(output) </pre>	<b>(c) Correct JAX Code</b> <pre> 1 import jax 2 import jax.numpy as jnp 3 from jax import random 4 5 # Initialize parameters 6 def init_params(key): 7     w_key, b_key = random.split(key) 8     # Correct shape 9     W = random.normal(w_key, (28*28, 10)) * jnp.sqrt(1 / (28*28)) 10    b = jnp.zeros((10,)) 11    return {'W': W, 'b': b} 12 13 # Forward function 14 def forward(params, x): 15     x = x.reshape(x.shape[0], -1) 16     return jnp.dot(x, params['W']) + params['b'] 17 18 # Example 19 key = random.PRNGKey(0) 20 params = init_params(key) 21 input_tensor = random.normal(key, (1, 1, 28, 28)) 22 output = forward(params, input_tensor) 23 print(output) </pre>
--	--	--	--

Figure 1: Example of PyTorch-to-JAX translation. (a) Input code; (b) Incorrect translation by 4o-mini; (c) Correct code.

Python, and C++. Huang et al. (2023) proposed *Codist*, which adopts a filtered IR to improve the precision of code translation through a process called code distillation. TehraniJamsaz et al. (2024) presented *CodeRosetta*, a framework for unsupervised translation from C to CUDA. Their method exploits the syntactic similarity between the two languages, leveraging abstract syntax trees (ASTs) as the pivot representation to learn structural correspondences. Roziere et al. (2021) emphasized the significance of a pre-training objective based on recovering broken or obfuscated code.

**Large Language Models for Code Translation.** The emergence of large language models (LLMs) capable of addressing questions across multiple domains has significantly benefited research in code translation. (Zhu et al., 2024) highlighted that state-of-the-art LLMs such as CodeLLAMA often produce translations lacking semantic equivalence—referred to as “shallow translations”—relative to the ground truth. To improve translation quality, Mahmud et al. (2024) proposed *AutoParLLM*, a framework to translate C code into OpenMP pragmatics. Their approach integrates Graph Neural Networks (GNNs) into the prompt to guide LLMs toward better output, and they introduced *OMPSCore*, a domain-specific metric tailored to evaluate OpenMP code translations. Tong & Zhang (2024) explored the use of LLMs as evaluators in a multi-phase process, involving code analysis, summarization, and fault localization to assess translated output. Similarly, Ibrahimzada et al. (2025) introduced *AlphaTrans*, a repository-level translation framework that applies LLMs across multiple translation and validation phases. Macedo et al. (2024) proposed *InterTrans*, an LLM-based framework that views code translation as a transitive process that involves intermediate languages.

### 3 MOTIVATION EXAMPLE

An illustration of how low-cost LLM as 4o-mini generate JAX buggy code from Pytorch code can be shown in Figure 1. The input Pytorch program, which is a simple neural network, defines a single-layer feedforward network that flattens a  $28 \times 28$  input image and applies a linear transformation to produce a 10-dimensional output tensor. The JAX code generated by 4o-mini has an extra function called *init\_params*, while in Pytorch the corresponding function was performed in the construction of the Neural Network object. We can see that the generated code consumed an error in Line 8 (see Figure 1b)), where it incorrectly passed the shape of a linear object defined by the parameter *W*. The correct version of the generated code, shown in Figure 1, requires one modification step to correct the argument passed onto the *random.normal* JAX API call. This example shows specific

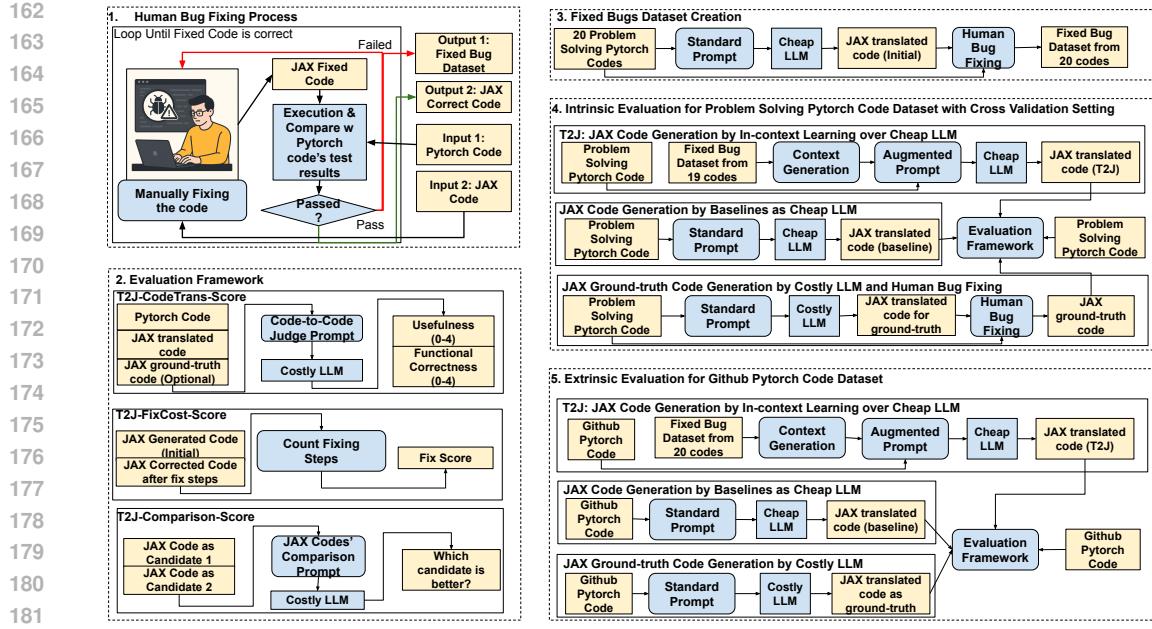


Figure 2: Overview Architecture of T2J

challenge of Pytorch-to-JAX translation that even with this simple code snippet, low-cost LLM still failed to extract the correct code.

## 4 APPROACH

We depict modules of T2J in Figure 2. In the first module, we describe how we hire software developers to check and modify JAX-generated code to ensure it is correct. We also define the definition of correct generated code in our work. The second module introduces our proposed metrics for comparing between predicted and expected generated code, in terms of using LLMs and leveraging the human bug fixing process’s output. In the third module, we go in to details about the fixed bug dataset, which we will use for in-context learning with data about fixed bugs as extra context. In the fourth and fifth sections, we describe how T2J performs translation and how we generate the ground-truth dataset for evaluation. We also describe what baseline setting we use for comparison with our proposed pipeline. We call these modules intrinsic evaluation and extrinsic evaluation, and these modules are performed on two different datasets. Next, we discuss about important concepts and design selection we use for T2J .

### 4.1 DESIGN SELECTION

#### 4.1.1 SELECTION OF PYTORCH DATASET

We curate datasets of PyTorch code for two tasks. First, we create a dataset of fixed bugs as JAX-generated code from PyTorch input code for our proposed in-context learning process. This process requires the involvement of software developers and LLM for PyTorch-to-JAX code translation. The second task is the evaluation process, where the PyTorch code will be used as input for baseline models or our proposed translation framework to get the output as JAX code snippets for further processing to correct the code and evaluation. Depending on the tasks, we collect the PyTorch dataset from two domains: problem-solving code and general-purpose code.

**Problem-Solving Code Dataset.** We construct a PyTorch dataset based on popular coding interview problems implemented in PyTorch. Owing to the popularity of these problems, we assume that developers can readily verify solutions and fix bugs by consulting existing online resources. For this purpose, we leverage the TorchLeet dataset (Aroori & Chien, 2025) as our problem-solving code corpus. We chose this dataset because each problem is scoped at the file level and because TorchLeet

---

216 has been highly ranked by GitHub users. From this dataset, we collect all **20** code snippets, covering  
217 three difficulty levels: easy, medium, and hard. All snippets are compilable and runnable using the  
218 default test cases included within the dataset.

219 **GitHub Code Dataset.** To further evaluate our proposed PyTorch-to-JAX translation approach on  
220 a broader range of PyTorch code, we consider code snippets drawn from high-quality repositories  
221 on GitHub. Specifically, we use the PyTorch subset of the large-scale CodeParrot dataset Wolf  
222 et al. (2022) as our second evaluation corpus. From this dataset, we extract **100** PyTorch code  
223 snippets. Unlike the Problem-Solving Code dataset, the GitHub PyTorch snippets originate from  
224 diverse repositories and are not guaranteed to be directly compilable.

225

#### 226 4.1.2 SELECTION OF LLMs

227

228 **Cheap LLMs.** We define low-cost LLMs as the target models to improve in the PyTorch-to-JAX  
229 translation task. These models, referred to as *Cheap LLMs*, can be used without commercial API  
230 keys or additional costs. Moreover, we leverage their JAX-generated code to construct the fixed-bug  
231 dataset, under the assumption that cheap LLMs produce a higher proportion of buggy code, which  
232 can enrich this dataset. For our experiments, we select GPT-4o-mini, the least expensive model  
233 offered by OpenAI, as the representative cheap LLM.

234 **Costly LLMs.** We define high-cost LLMs, referred to as *Costly LLMs*, as the source of ground-truth  
235 JAX code for comparison with the translations produced by cheap LLMs. We employ the GPT-4o  
236 model as the costly LLM. GPT-4o is one of the most widely used models provided by OpenAI.

237

#### 238 4.1.3 PROMPTS

239 We define these following types of prompt in our work.

240 **Standard Prompt.** We define the standard prompt as the basic prompt for translating code snippet  
241 from Pytorch to JAX (see Appendix A.1). In this prompt, we define the translation request by  
242 annotating the role of LLM as an expert in code translation and provide a basic request that translate  
243 from source language as PyTorch to target language as JAX code.

244 **Augmented Prompt.** We design augmented prompt with following information. First, similar  
245 to the standard prompt, the augmented prompt will have information about the role of LLM and  
246 the requirement about input and output. Differ from standard prompt, the augmented prompt will  
247 specify a hint for LLMs as the list of errors and errors' solution provided by the fixed bug dataset,  
248 uploaded in the JSON format as context of prompt. Definitions of each fields in the fixed bug dataset  
249 are also included in the prompt content (see Appendix A.2).

250 **Evaluation Prompt.** We propose two metrics that leverage LLMs for code-to-code translation.  
251 From our knowledge, there are yet existing LLM prompt for doing this task. We define a set of  
252 prompts, called evaluation prompts (see Appendix A.3), to ask costly LLMs to evaluate source code  
253 in different aspects. We will describe in details of these metrics in the next section.

254

#### 255 4.1.4 TYPES OF EVALUATION

256 **Intrinsic and Extrinsic Evaluation.** Our evaluation is divided into two parts. The first, referred  
257 to as intrinsic evaluation, is conducted on 20 problem-solving PyTorch code snippets using a cross-  
258 validation setting. The second, extrinsic evaluation, is performed on 100 samples of PyTorch code  
259 collected from GitHub. There are two key differences between these configurations. First, for  
260 intrinsic evaluation, the ground-truth JAX code is obtained through human verification and bug-  
261 fixing, whereas for extrinsic evaluation, we rely on JAX code generated by a costly LLM for the  
262 GitHub dataset.

263

## 264 4.2 HUMAN BUG FIXING PROCESS

265 The core module of T2J is a systematic bug-fixing process designed to improve the reliability of  
266 LLM-generated translations from PyTorch to JAX. Starting with Python code written in PyTorch  
267 and the corresponding JAX-generated code from LLMs, this process produces a parallel dataset that  
268 pairs the original PyTorch snippets with their corresponding corrected JAX translations. To ensure

270

271 Table 1: Evaluation preferences and their descriptions for T2J\_CodeTrans\_Score.

272

273

274

275

276

277

278

correctness, we employed two professional software developers with over five years of Python programming experience to analyze and fix the outputs produced by LLMs. The JAX translations are then subjected to careful manual verification: the developers perform multiple rounds of debugging and correction until the translated JAX code passes all test cases and produces results equivalent to the original PyTorch implementation. During this stage, the verifiers execute the translated and fixed code using Python compilers to confirm correctness. At the end of the process, two complementary datasets are created: one containing pairs of PyTorch snippets and their validated JAX counterparts, and another capturing the bugs identified in LLM outputs along with their corresponding fixes. Importantly, this methodology is flexible, as we apply the same procedure to different PyTorch datasets and experiment with different LLMs depending on the objectives of other modules within our framework.

289

290

291

292

293

294

295

296

297

**Verified Code Correctness by Human.** We consider a fixed version as JAX code snippet verified by a human as correct if and only if the JAX-generated code can be compiled, runnable, and returns the same output as its corresponding Pytorch code snippet, given the same test case. In our work, the human bug fixing process was performed in the problem-solving dataset only since its code snippets have test case and can be runnable which we can rely on their execution output for comparison with the JAX generated code.

298

299

300

301

302

303

304

305

306

307

308

309

310

311

312

313

314

315

316

317

318

319

320

321

322

323

**T2J\_CodeTrans\_Score.** Similar to the ICE-Score metric (Zhuo, 2024) for natural language to code translation, we use the GPT-4o model to evaluate the quality of translated code by usefulness and functional correctness. We design prompt, called CodeTrans prompt, with corresponding evaluation criteria and scoring rubric from 0 (lowest) to 4 (highest). CodeTrans also be useable without having the reference code. Thus, we have a set of following metrics: T2J\_CodeTrans\_Use\_Ref (i.e. the metric for usefulness with reference), T2J\_CodeTrans\_Func\_Ref, T2J\_CodeTrans\_Use\_NoRef, T2J\_CodeTrans\_Func\_NoRef. Explanation of two aspects/ preferences is shown in Table 1.

**T2J\_FixCost\_Score.** We measure the number of fix steps required to have the JAX correct code from input JAX initial translated code from LLM with this equation:

$$T2J\_FixCost\_Score(J^{before\_fix}, J^{correct}) = \frac{1}{n} \sum_{i=1}^n count\_fix\_step(j_i^{before\_fix}, j_i^{correct}) \quad (1)$$

In equation 1,  $J^{before\_fix}$  is the set of JAX generated code as input for human verification process, while  $J^{correct}$  is the final version of JAX fixed code that is correct, i.e. it can run and returns consistent output with its corresponding PyTorch code given the same input test case.

---

324     **T2J\_Comparison\_Score.** In this metric, we propose a direct comparison between two translation  
325     sets  $J^1$  and  $J^2$  to see which one is closer to the input PyTorch code set  $P$ , implemented by this  
326     equation:  
327

328

$$329 \quad \text{T2J\_Comparison\_Score}(J^1, P, J^2) = \frac{1}{n} \sum_{i=1}^n \begin{cases} 1, & \text{if } \text{is\_better}(j_i^1, j_i^2, p_i) \\ 0, & \text{otherwise} \end{cases} \quad (2)$$

330

331     In equation 2, the  $\text{is\_better}(j_i^1, j_i^2, p_i)$  function returns 1 if  $j_i^1$  is considered as better code than  
332      $j_i^2$ , given the input prompt for comparison called Comparison prompt (see Appendix A.3). Note, we  
333     also include the content of PyTorch code  $p_i$  as a required input to help LLM comparing the input  
334     with each code candidate. The scale of this score is from 0 to 1.  
335

336

#### 337     4.4 FIXED BUG DATASET

338

339     Given the input of 20 PyTorch samples from the TorchLeet dataset, two professional developers are  
340     hired to modify the code snippets generated from the LLM code translation process. The output  
341     of this process for each generated code is a set of multiple fix steps. We store this data in JSON  
342     array format to be usable as the context for T2J prompting technique. The fixed bug dataset was  
343     constructed by our selected cheap LLM 4o-mini. In total, this dataset contains **163** pairs of bugs/  
344     solutions to fix bug.  
345

346

#### 347     4.5 INTRINSIC EVALUATION

348

349     In the intrinsic evaluation setting, we want to check if we can use knowledge from fixing 19 problem-  
350     solving code pairs of Pytorch and JAX correct code to improve the quality of cheap LLM to translate  
351     Pytorch code of the remaining sample. We perform this evaluation as a cross validation process over  
352     20 samples of the fixed bug dataset. Next, we compare the output of three following modules.  
353

354     **Code generation as Baseline.** We perform the code translation process using the cheap LLM 4o-  
355     mini. This process leverages the standard prompt (see Appendix A.1) to translate input PyTorch  
356     code to JAX-generated code. We consider this configuration as baseline for T2J for comparison.  
357

358     **T2J 's In-context learning for code generation.** We perform this process with following modules.  
359     First, the augmented prompt will be created with input information as the given Pytorch code snippet  
360     and the JSON format of other bug-solutions from other code samples of the problem-solving code  
361     dataset. Next, through the cheap LLM, the JAX generated code by T2J 's approach is created. This  
362     process was done without the need of any fine-tuning steps, which are usually costly.  
363

364     **Ground-truth JAX code generation.** We take advantage of costly LLM gpt-4o to generate the JAX  
365     code, called JAX initial code from given problem-solving code snippet. Next, the human bug fixing  
366     process was performed on this JAX initial code to make the JAX corrected code as ground truth.  
367     The main different between this process and the fixed bug dataset creation we mentioned earlier is  
368     that this module works with costly LLM. Finally, the output of these three modules will be use as  
369     the input of evaluation framework along with the PyTorch source code.  
370

371     

#### 4.6 EXTRINSIC EVALUATION

372

373     There are two main differences between intrinsic evaluation and extrinsic evaluation. First, for this  
374     set up we use the code snippets collected from another dataset, called Github dataset. We evaluate on  
375     100 sample code snippets that are at repository level, meaning that the input code snippets, usually  
376     collected from single files, are not guarantee to be runnable and having test cases. Thus, we will use  
377     automated metrics we propose for the evaluation. Second, the extrinsic evaluation considered the  
378     output of costly LLMs, i.e. JAX translated code by this process) as the ground-truth data point.  
379

378  
379 Table 2: Results with Cheap LLM (gpt-4o-mini) and Costly LLM (gpt-4o).  
380  
381  
382

Metrics	Intrinsic Evaluation		Extrinsic Evaluation	
	Baseline	T2J	Baseline	T2J
CodeBLEU	0.19	0.29	0.41	0.38
T2J_CodeTrans_Use_Ref	1.75	2.55	2.74	2.94
T2J_CodeTrans_Func_Ref	0.35	1.3	2.43	3.02
T2J_CodeTrans_Use_NoRef	1.60	2.45	2.81	2.98
T2J_CodeTrans_Func_NoRef	0.70	2.15	2.74	3.37
T2J_FixCost_Score	163	87	N/A	N/A
T2J_Comparison_Score	0	1	0.18	0.82

383  
384  
385  
386  
387  
388  
389  
390 Table 3: Correlation of other metrics with T2J\_FixCost\_Score on Fixed Bug Dataset.  
391  
392  
393  
394  
395  
396  
397  
398  
399

Metric	Correlation with T2J_FixCost_Score	
	Pearson	Spearman
CodeBLEU	0.04	0.2
T2J_CodeTrans_Use_Ref	0.2	0.25
T2J_CodeTrans_Func_Ref	0.07	0.07
T2J_CodeTrans_Use_NoRef	0.09	0.19
T2J_CodeTrans_Func_NoRef	0.11	0.29
T2J_Comparison_Score	NaN	NaN

400  
401 5 EXPERIMENT

## 402 5.1 SETUP

403  
404  
405  
406 **Hardware Configuration.** For the human bug fixing process, software developers work in the  
407 Google Colab environment to debug and fix the code. They use Python 3 as a compiler and use one  
408 T4 GPU for running all sample code.409  
410 **Question-answering for cheap and costly LLMs.** For both cheap LLM (4o-mini) and costly  
411 LLM(gpt-4o), we perform the code generation process through the official interface of ChatGPT-  
412 pro. Each question will be created solely in a new topic, and from the answer given by ChatGPT’s  
413 interface, we manually extract the code snippet as JAX-generated code. Other textual explanation  
414 in the output will be omitted. To add context to the existing prompt, we use the Upload function  
415 provided by ChatGPT’s interface to assign a fixed bug dataset as a JSON file to our designed prompt.416  
417 **Executing Evaluation Prompts.** We leverage the access on OpenAI gpt-4o models by API key to  
418 get the scoring results for T2J\_CodeTrans\_Score. For T2J\_Comparison\_Score, we use ChatGPT-  
419 pro’s interface to ask questions and receive answers. The reason for using ChatGPT-pro’s interface  
420 instead of using API key is that some tasks of our work require JSON file upload.

## 421 5.2 RESULT

## 422 5.2.1 TRANSLATION ACCURACY

423  
424 From Figure 2, our pipeline improves the CodeBLEU score to 0.29, representing a 10% relative  
425 gain. In terms of T2J\_CodeTrans\_Use\_Ref, the augmented prompt yields an improvement of  
426 0.8 points over the baseline. For functional correctness, T2J achieves an improvement of 0.95  
427 point. Under the no-reference configuration—where the LLM evaluates only by comparing the  
428 input and translated code—T2J still delivers gains of 0.85 and 1.15 points for usefulness and  
429 functional correctness, respectively, as measured by the T2J\_CodeTrans metrics. Regarding the  
430 T2J\_Comparison\_Score, 100% of the translations generated by T2J are judged superior to the base-  
431 line outputs. Finally, in terms of fixing cost, T2J enables GPT-4o-mini to produce code requiring  
432 only 87 fixing steps—roughly half the effort compared to fixing the baseline JAX outputs.

432  
433 Table 4: Comparison of running time (seconds) on Intrinsic Evaluation.  
434  
435

---

PyTorch	Ground Truth	Baseline	T2J
1003	851	1232.9	449

---

438 Table 5: Comparison of human fixing costs between baseline (weak LLM with standard prompt)  
439 and JAX code initially generated by a costly LLM.  
440

Num. of Fixes	Correcting Weak LLM's Code	Correcting Costly LLM's Code
Minimum	1	0
Maximum	32	12
Mean	8.15	2.77
Median	5	2
Total	163	61

441  
442  
443  
444  
445  
446  
447  
448 The extrinsic evaluation on the GitHub PyTorch code also highlights the superiority of T2J over the  
449 baseline in generating precise code. Interestingly, in this setting the baseline approach outperformed  
450 T2J by 3% according to CodeBLEU. In terms of the T2J\_CodeTrans metrics, our approach achieves  
451 improvements of up to 1.2 point in usefulness and 0.6 point in functional correctness.

### 452 5.2.2 CORRELATION OF CODE TRANSLATION METRICS VS HUMAN FIXING COST 453

454 From Table 3, we observe that the T2J\_CodeTrans\_Score metrics show the strongest correlation with  
455 T2J\_FixCost\_Score under both Pearson and Spearman measures. Overall, all metrics exhibit weak  
456 correlation (below 0.3) with fixing cost. One possible reason is that other metrics are continuous,  
457 whereas fixing cost is measured as discrete steps.

### 458 5.2.3 HOW CLOSE JAX-GENERATED CODE BY COSTLY LLM IS TO BEING CORRECT 459

460 We further analyze the quality of JAX generated code by costly LLM by measuring the fixing cost  
461 to get the JAX correct code by costly LLM as ground-truth code. The result, shows in Table 5, shows  
462 that while it requires much less effort for correcting costly LLM's output than baseline's output, it  
463 still requires in total 61 fixing steps to get the correct code set.

### 464 5.2.4 ANALYSIS ON RUNNING TIME 465

466 We analyze the running time of the corrected code from 3 settings for intrinsic evaluation in Table  
467 4. Results show that T2J can provide significant improvement as 2.5 times faster than running the  
468 baseline's output. Details of running time can be seen in Appendix B.

## 470 6 LIMITATIONS 471

472 First, the current version of T2J has not yet been applied to improving open LLMs, due to budget  
473 constraints that limit our ability to hire software professionals for the human bug-fixing process on  
474 these models. Second, our measure of fixing cost is currently based only on the number of fixes,  
475 whereas in practice each fix may vary in difficulty. To improve this, there is a need of an algorithm  
476 that estimates the relative effort of each bug-fixing step. Third, we conducted the human bug-fixing  
477 process only on the problem-solving code dataset, which we want to extend this process for other  
478 domains. In future versions of T2J, we plan to collect bug datasets from other general domains.  
479

## 480 7 CONCLUSION AND FUTURE WORKS 481

482 In this work, we show that T2J can achieve significant improvement compared to baselines as  
483 original 4o-mini model for PyTorch-to-JAX code translation. In future work, we attempt to apply  
484 our approach for newer open-source LLMs and leverage more advance techniques as supervised  
485 fine-tuning and direct preference optimization.

---

## 486 REFERENCES

487

488 Wasi Uddin Ahmad, Saikat Chakraborty, Baishakhi Ray, and Kai-Wei Chang. Unified pre-training  
489 for program understanding and generation, 2021. URL <https://arxiv.org/abs/2103.06333>.

490  
491 Chandrahas Aroori and Caslow Chien. Torchleet: Leetcode for pytorch. <https://github.com/Exorust/TorchLeet>, 2025. Accessed: YYYY-MM-DD.

492  
493 Tuan Dinh, Jinman Zhao, Samson Tan, Renato Negrinho, Leonard Lausen, Sheng Zha, and George  
494 Karypis. Large language models of code fail at completing code with potential bugs, 2023. URL  
495 <https://arxiv.org/abs/2306.03438>.

496  
497 Shihan Dou, Haoxiang Jia, Shenxi Wu, Huiyuan Zheng, Weikang Zhou, Muling Wu, Mingxu  
498 Chai, Jessica Fan, Caishuang Huang, Yunbo Tao, Yan Liu, Enyu Zhou, Ming Zhang, Yuhao  
499 Zhou, Yueming Wu, Rui Zheng, Ming Wen, Rongxiang Weng, Jingang Wang, Xunliang Cai, Tao  
500 Gui, Xipeng Qiu, Qi Zhang, and Xuanjing Huang. What's wrong with your code generated by  
501 large language models? an extensive study, 2024. URL <https://arxiv.org/abs/2407.06153>.

502  
503 Daya Guo, Shuo Ren, Shuai Lu, Zhangyin Feng, Duyu Tang, Shujie Liu, Long Zhou, Nan Duan,  
504 Alexey Svyatkovskiy, Shengyu Fu, Michele Tufano, Shao Kun Deng, Colin Clement, Dawn  
505 Drain, Neel Sundaresan, Jian Yin, Dixin Jiang, and Ming Zhou. Graphcodebert: Pre-training  
506 code representations with data flow, 2021. URL <https://arxiv.org/abs/2009.08366>.

507  
508 Daya Guo, Shuai Lu, Nan Duan, Yanlin Wang, Ming Zhou, and Jian Yin. Unixcoder: Unified cross-  
509 modal pre-training for code representation, 2022. URL <https://arxiv.org/abs/2203.03850>.

510  
511 Yufan Huang, Mengnan Qi, Yongqiang Yao, Maoquan Wang, Bin Gu, Colin Clement, and Neel  
512 Sundaresan. Program translation via code distillation, 2023. URL <https://arxiv.org/abs/2310.11476>.

512  
513 Hamel Husain, Ho-Hsiang Wu, Tiferet Gazit, Miltiadis Allamanis, and Marc Brockschmidt. Code-  
514 searchnet challenge: Evaluating the state of semantic code search, 2020. URL <https://arxiv.org/abs/1909.09436>.

514  
515 Ali Reza Ibrahimzada, Kaiyao Ke, Mrigank Pawagi, Muhammad Salman Abid, Rangeet Pan,  
516 Saurabh Sinha, and Reyhaneh Jabbarvand. Alphatrans: A neuro-symbolic compositional ap-  
517 proach for repository-level code translation and validation. *Proceedings of the ACM on Software  
Engineering*, 2(FSE):2454–2476, June 2025. ISSN 2994-970X. doi: 10.1145/3729379. URL  
518 <http://dx.doi.org/10.1145/3729379>.

518  
519 Shuai Lu, Daya Guo, Shuo Ren, Junjie Huang, Alexey Svyatkovskiy, Ambrosio Blanco, Colin  
520 Clement, Dawn Drain, Dixin Jiang, Duyu Tang, Ge Li, Lidong Zhou, Linjun Shou, Long Zhou,  
521 Michele Tufano, Ming Gong, Ming Zhou, Nan Duan, Neel Sundaresan, Shao Kun Deng, Shengyu  
522 Fu, and Shujie Liu. Codexglue: A machine learning benchmark dataset for code understanding  
523 and generation, 2021. URL <https://arxiv.org/abs/2102.04664>.

523  
524 Marcos Macedo, Yuan Tian, Pengyu Nie, Filipe R. Cogo, and Bram Adams. Intertrans: Leveraging  
525 transitive intermediate translations to enhance llm-based code translation, 2024. URL <https://arxiv.org/abs/2411.01063>.

525  
526 Quazi Ishtiaque Mahmud, Ali TehraniJamsaz, Hung D Phan, Nesreen K. Ahmed, and Ali Jannesari.  
527 AUTOPARLLM: GNN-guided automatic code parallelization using large language models, 2024.  
528 URL <https://openreview.net/forum?id=znjaiy1Z9q>.

528  
529 Rangeet Pan, Ali Reza Ibrahimzada, Rahul Krishna, Divya Sankar, Lambert Pouguem Wassi,  
530 Michele Merler, Boris Sobolev, Raju Pavuluri, Saurabh Sinha, and Reyhaneh Jabbarvand. Lost in  
531 translation: A study of bugs introduced by large language models while translating code. In *Pro-  
532 ceedings of the IEEE/ACM 46th International Conference on Software Engineering*, ICSE '24,  
533 New York, NY, USA, 2024. Association for Computing Machinery. ISBN 9798400702174. doi:  
534 10.1145/3597503.3639226. URL <https://doi.org/10.1145/3597503.3639226>.

---

540       Qiwei Peng, Yekun Chai, and Xuhong Li. Humaneval-xl: A multilingual code generation bench-  
541       mark for cross-lingual natural language generalization, 2024. URL <https://arxiv.org/abs/2402.16694>.  
542

543       Shuo Ren, Daya Guo, Shuai Lu, Long Zhou, Shujie Liu, Duyu Tang, Neel Sundaresan, Ming Zhou,  
544       Ambrosio Blanco, and Shuai Ma. Codebleu: a method for automatic evaluation of code synthesis,  
545       2020. URL <https://arxiv.org/abs/2009.10297>.  
546

547       Baptiste Roziere, Marie-Anne Lachaux, Marc Szafraniec, and Guillaume Lample. Dobf: A deob-  
548       fuscation pre-training objective for programming languages, 2021. URL <https://arxiv.org/abs/2102.07492>.  
549

550       Baptiste Roziere, Jie M. Zhang, Francois Charton, Mark Harman, Gabriel Synnaeve, and Guillaume  
551       Lample. Leveraging automated unit tests for unsupervised code translation, 2022. URL <https://arxiv.org/abs/2110.06773>.  
552

553       Marc Szafraniec, Baptiste Roziere, Hugh Leather, Francois Charton, Patrick Labatut, and Gabriel  
554       Synnaeve. Code translation with compiler representations, 2023. URL <https://arxiv.org/abs/2207.03578>.  
555

556       Ali TehraniJamsaz, Arijit Bhattacharjee, Le Chen, Nesreen K. Ahmed, Amir Yazdanbakhsh, and Ali  
557       Jannesari. Coderosetta: Pushing the boundaries of unsupervised code translation for parallel pro-  
558       gramming. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*,  
559       2024. URL <https://openreview.net/forum?id=V6hrg409gg>.  
560

561       Weixi Tong and Tianyi Zhang. Codejudge: Evaluating code generation with large language models,  
562       2024. URL <https://arxiv.org/abs/2410.02184>.  
563

564       Ngoc Tran, Hieu Tran, Son Nguyen, Hoan Nguyen, and Tien N. Nguyen. Does bleu score work  
565       for code migration? In *Proceedings of the 27th International Conference on Program Com-*  
566       *prehension*, ICPC '19, pp. 165–176. IEEE Press, 2019. doi: 10.1109/ICPC.2019.00034. URL  
567       <https://doi.org/10.1109/ICPC.2019.00034>.  
568

569       Yue Wang, Hung Le, Akhilesh Deepak Gotmare, Nghi D. Q. Bui, Junnan Li, and Steven C. H. Hoi.  
570       Codet5+: Open code large language models for code understanding and generation, 2023. URL  
571       <https://arxiv.org/abs/2305.07922>.

572       Thomas Wolf, Leandro Werra, Loubna Allal, and Zdar. Codeparrot dataset. <https://huggingface.co/codeparrot>, 2022. Accessed: YYYY-MM-DD.  
573

574       Quanjun Zhang, Chunrong Fang, Yang Xie, YuXiang Ma, Weisong Sun, Yun Yang, and Zhenyu  
575       Chen. A systematic literature review on large language models for automated program repair,  
576       2024. URL <https://arxiv.org/abs/2405.01466>.  
577

578       Shuyan Zhou, Uri Alon, Sumit Agarwal, and Graham Neubig. Codebertscore: Evaluating code  
579       generation with pretrained models of code, 2023. URL <https://arxiv.org/abs/2302.05527>.  
580

581       Ming Zhu, Mohimenul Karim, Ismini Lourentzou, and Daphne Yao. Semi-supervised code trans-  
582       lation overcoming the scarcity of parallel code data. In *Proceedings of the 39th IEEE/ACM*  
583       *International Conference on Automated Software Engineering*, ASE '24, pp. 1545–1556, New  
584       York, NY, USA, 2024. Association for Computing Machinery. ISBN 9798400712487. doi:  
585       10.1145/3691620.3695524. URL <https://doi.org/10.1145/3691620.3695524>.  
586

587       Terry Yue Zhuo. ICE-score: Instructing large language models to evaluate code. In Yvette Graham  
588       and Matthew Purver (eds.), *Findings of the Association for Computational Linguistics: EACL*  
589       2024, pp. 2232–2242, St. Julian's, Malta, March 2024. Association for Computational Linguistics.  
590       URL <https://aclanthology.org/2024.findings-eacl.148/>.  
591

592

593

---

594 APPENDIX  
595

596 A PROMPT TEMPLATE  
597

598 A.1 STANDARD PROMPT FOR PYTORCH-TO-JAX TRANSLATION  
599

600  
601 You are an expert in programming language translation from PyTorch to JAX. In this task, I  
602 will give you input as PyTorch code. Please translate this input PyTorch code to JAX code:  
603 **Input Source Code Snippet: {CODE}**  
604

605  
606 Figure 3: Standard Prompt for PyTorch-to-JAX code translation. **The prompt in blue shows the**  
607 **immediate query after the prompt. {CODE} is the starting string of the code.**  
608

609  
610 A.2 AUGMENTED PROMPT FOR PYTORCH-TO-JAX TRANSLATION  
611

612  
613 You are an expert in programming language translation from PyTorch to JAX. In this task, I  
614 will give you two inputs:  
615 1. Pytorch source code.  
616 2. A JSON file that contains a dataset of common errors in PyTorch-to-JAX translation by  
617 Weak LLM 4o-mini. Each data point contains the following fields:  
618 - Example\_id: ID of the source code.  
619 - Input\_Code: Source code in Pytorch.  
620 - LLM\_weak\_output: JAX translated code of Input\_Code using a weak LLM (4o-mini).  
621 LLM\_fix\_output: Fixed JAX code from LLM\_weak\_output by the process of manually  
622 check and fix errors conducted by software developers.  
623 - Errors: This is a list of errors that appeared in the process of manually checking and  
624 fixing bugs from LLM\_weak . Each error item has the following labels:  
625 • "Error\_Code": The part of LLM\_weak\_output that caused the error.  
626 • "Error": the error message returned by compilation.  
627 • "Fix\_info": the textual description of how to fix the error code  
628 • "Fixed\_Code": The fixed code corresponding to the "Error\_Code" part.  
629 3. The data.csv file which stored possible input when running some examples in the JSON  
630 file. Your task is to reason and get the output JAX code from these above inputs. Please  
631 note that you can learn the process of error fixing in Torch-to-JAX translation in 2) JSON  
632 file. Now I will give you a set of input in the next query.  
633

634 **Input Source Code Snippet: {CODE}**  
635

636 Figure 4: Prompt for Augmenting to the weak LLM. **The prompt in blue shows the immediate**  
637 **query after the prompt. {CODE} is the starting string of the code.**

640 A.3 EVALUATION PROMPT FOR T2J\_CODETRANS\_SCORE  
641

642 A.3.1 FUNCTIONAL CORRECTNESS  
643

644 See the prompt without reference at Figure 5 and the prompt with reference at Figure 6.  
645

646 A.3.2 USEFULNESS  
647

648 See the prompt without reference at Figure 7 and the prompt with reference at Figure 8.

---

648  
 649 You will be given a JAX code snippet that was translated from PyTorch source code. Your  
 650 task is to rate the snippet on **one metric only**: its **functional correctness**.  
 651 Please ensure you read and understand these instructions carefully before reviewing. Refer  
 652 to this guide as needed during the evaluation process.  
 653 Evaluation Criteria:  
 654 Functional Correctness (0–4) — How well the JAX code preserves the behavior of the origi-  
 655 nal PyTorch code.  
 656 You are to assess whether the JAX code would produce equivalent outputs to the original  
 657 PyTorch code across possible inputs, even though the PyTorch code is not shown. Consider  
 658 unit-test-style logic and general expectations of equivalence.  
 659 - A score of 0: The translation is completely incorrect and meaningless.  
 660 - A score of 4: The translation is fully correct and handles all core functionalities as expected.  
 661 Evaluation Steps:  
 662 1. Assume the code was translated from PyTorch and should preserve its logic.  
 663 2. Evaluate whether the JAX code appears complete, meaningful, and implementationally  
 664 correct based on general expectations for such translations.  
 665 3. Assign a score for functional correctness on a scale from 0 to 4.  
 666 Input Source Code in PyTorch:  
 667 **{SOURCE\_CODE}**  
 668 Translated JAX Code Snippet:  
 669 **{TRANSLATED\_CODE}**  
 670 Evaluation Form:  
 671 Functional Correctness (scores ONLY):  
 672  
 673  
 674

Figure 5: Prompt for Scoring Functional Correctness by T2J\_CodeTrans\_Func\_NoRef

Table 6: Error categories and their counts in Human Bug Fixing dataset.

Error Main Category	Count
<b>Training loops, training steps, model fitting</b>	32
<b>Other miscellaneous</b>	43
<b>Model definitions, LinearModel classes, encoders/decoders</b>	51
<b>Loss functions, gradient computation, criterion</b>	12
<b>JAX-specific constructs: jit, grad, PRNG, etc.</b>	18
<b>Iteration patterns: for, while, data loops</b>	1
<b>Parameter updates</b>	3
<b>Final layers, return statements, outputs</b>	3
<b>TOTAL</b>	163

#### A.4 EVALUATION PROMPT FOR T2J\_COMPARISON\_SCORE

The prompt for querying the T2J\_Comparison\_Score can be seen in Figure 9.

## B ADDITIONAL RESULTS

### B.1 ANALYSIS ON CATEGORIES OF BUGS

We perform a study on the categorization of bugs on the fixed bug dataset as following. First, two software professionals will go through all the bugs and discuss about the categorizations. Second, from this categorization, they go to the dataset’s entities for the second time and do the annotation for categories. We summarize the categorization in Table 6. We further classify types of bugs for some categories to sub-categories, shown in Table 7 and Table 8. We upload each case of this categorization process in the replication package.

---

702  
 703 You will be given a JAX code snippet that was translated from PyTorch source code. Your  
 704 task is to rate the snippet on **one metric only**: its **functional correctness**.  
 705 Please ensure you read and understand these instructions carefully before reviewing. Refer  
 706 to this guide as needed during the evaluation process.  
 707 Evaluation Criteria: Functional Correctness (0–4) — How well the JAX code preserves the  
 708 behavior of the original PyTorch code.  
 709 You are to assess whether the JAX code would produce equivalent outputs to the original  
 710 PyTorch code across possible inputs, even though the PyTorch code is not shown. Consider  
 711 unit-test-style logic and general expectations of equivalence.  
 712 - A score of 0: The translation is completely incorrect and meaningless.  
 713 - A score of 4: The translation is fully correct and handles all core functionalities as expected.  
 714 Evaluation Steps:  
 715 1. Assume the code was translated from PyTorch and should preserve its logic.  
 716 2. Evaluate whether the JAX code appears complete, meaningful, and implementationally  
 717 correct based on general expectations for such translations.  
 718 3. Assign a score for functional correctness on a scale from 0 to 4.  
 719 Input Source Code in PyTorch:  
 720 **{SOURCE\_CODE}**  
 721 Translated JAX Code Snippet:  
 722 **{TRANSLATED\_CODE}**  
 723 Reference JAX Code Snippet:  
 724 **{REFERENCE}**  
 725 Evaluation Form:  
 726 Functional Correctness (scores ONLY):  
 727  
 728

Figure 6: Prompt for Scoring Functional Correctness by T2J\_CodeTrans\_Func\_Ref

Table 7: Error subcategories under **training loops, training steps, and model fitting**.

Error Subcategory	Count
Misc training issues	8
Improper passing/using <code>rng_key/prng_key</code>	6
Epoch in <code>range(...)</code> loop issues	3
Incorrect usage of Flax <code>TrainState</code> and <code>state.apply_gradients</code>	3
Incorrect usage of wrappers (e.g., <code>train_model(...)</code> / <code>fit(...)</code> )	3
Train steps return only new state/params without loss at epoch level	3
JIT/static argument handling for training functions	3
Errors with batches or dataloaders in training	1
Loop constructs that break vectorization	1
Optimizer update/apply patterns in the training loop	1

## B.2 RUNNING TIME ANALYSIS

We perform the running process on a T4 GPU for all the code. We set the timeout of program to run as 180 seconds. Results for each sample in the intrinsic evaluation are shown in Table 9.

## C CONFIGURATIONS

For the task that required user interface action with LLMs, we use the default ChatGPT-pro setting for gpt-4o and 4o-mini. Most of the data were created before July 31st, 2025 when 4o-mini was still available on ChatGPT’s interface. For task like LLM-based metric calculation, we leverage Open-Router’s API<sup>1</sup> to perform the implementation of these tasks. We also report the hyper parameters for querying costly LLMs for code evaluation in the replication packages.

<sup>1</sup><https://openrouter.ai/>

---

756  
 757  
 758  
 759  
 760

Your task is to rate the snippet on **one metric only**: its **usefulness** for understanding and reusing the logic of a typical PyTorch implementation.

Please ensure you read and understand these instructions carefully before reviewing. Refer to this guide as needed during the evaluation process.

Evaluation Criteria: Usefulness (0–4) — How useful the JAX code is for replicating or adapting the functionality of a typical PyTorch source code implementation.

- A score of 0: The JAX translated snippet is irrelevant or confusing and does not help at all.
- A score of 1: The JAX translated snippet includes some related elements but is mostly unhelpful.
- A score of 2: The JAX translated snippet is somewhat useful but needs substantial modification.
- A score of 3: The JAX translated snippet is helpful with minor revisions needed.
- A score of 4: The JAX translated snippet is very helpful and covers the intended functionality clearly.

Evaluation Steps:

1. Assume the PyTorch source code performs a well-defined functionality.
2. Determine whether the JAX translated code snippet enables meaningful reuse or guidance toward equivalent implementation.
3. Assign a score for usefulness from 0 to 4.

Input Source Code in PyTorch:  
**{SOURCE\_CODE}**

Translated JAX Code Snippet:  
**{TRANSLATED\_CODE}**

Evaluation Form:  
 Usefulness (scores ONLY):

771  
 772  
 773  
 774  
 775  
 776  
 777  
 778  
 779  
 780  
 781  
 782  
 783  
 784  
 785  
 786  
 787  
 788  
 789  
 790  
 791  
 792  
 793  
 794  
 795

Figure 7: Prompt for Scoring Usefulness by T2J\_CodeTrans\_Use\_NoRef

Table 8: Error subcategories under **other miscellaneous**.

Error Subcategory	Count
Data arrays, tensors, and dataset values (e.g., creating arrays, specifying shapes)	8
Dot products with parameters (e.g., <code>params["w"]</code> )	5
Initialization, often in class constructors ( <code>__init__</code> )	1
Dot products, sums, or nonlinear transforms	3
Neural network layers and activations (e.g., <code>nn.relu</code> , <code>nn.Dense</code> , LSTMs, decoders/encoders)	5
Tensor dimension errors	2
Generating synthetic data for CSV	1
Errors with constant declaration (e.g., <code>epoch</code> )	1
Incomplete functions/placeholders	17

807  
 808  
 809

---

```

810
811
812
813
814 Your task is to rate the snippet on **one metric only**: its **usefulness** for understanding
815 and reusing the logic of a typical PyTorch implementation.
816 Please ensure you read and understand these instructions carefully before reviewing. Refer
817 to this guide as needed during the evaluation process.
818 Evaluation Criteria: Usefulness (0–4) — How useful the JAX code is for replicating or
819 adapting the functionality of a typical PyTorch source code implementation.
820 - A score of 0: The JAX translated snippet is irrelevant or confusing and does not help at all.
821 - A score of 1: The JAX translated snippet includes some related elements but is mostly
822 unhelpful.
823 - A score of 2: The JAX translated snippet is somewhat useful but needs substantial modifi-
824 cation.
825 - A score of 3: The JAX translated snippet is helpful with minor revisions needed.
826 - A score of 4: The JAX translated snippet is very helpful and covers the intended function-
827 ality clearly.
828 Evaluation Steps:
829 1. Assume the PyTorch source code performs a well-defined functionality.
830 2. Determine whether the JAX translated code snippet enables meaningful reuse or guidance
831 toward equivalent implementation.
832 3. Assign a score for usefulness from 0 to 4.
833 Input Source Code in PyTorch:
834 {SOURCE_CODE}
835 Translated JAX Code Snippet:
836 {TRANSLATED_CODE}
837 Reference JAX Code Snippet:
838 {REFERENCE}
839 Evaluation Form:
840 Usefulness (scores ONLY):

```

Figure 8: Prompt for Scoring Usefulness by T2J \_CodeTrans\_Use\_Ref

```

841
842
843
844
845
846
847
848 You are an expert in PyTorch to JAX translation. I provide 3 inputs: 1 . PyTorch input code;
849 2. Translated Code Candidate A; 3. Translated Code Candidate B. Which candidate is a
850 better translation result for this Pytorch code.
851 Input Pytorch code:
852 {CODE}
853 2. Translated Code A:
854 {TRANSLATE_CODE_A}
855 3. Translated Code B:
856 {TRANSLATE_CODE_B}
857 Please also provide the reason why you consider a candidate better than the other translated
858 code candidate.

```

Figure 9: Prompt for T2J \_Comparison\_Score.

864  
865  
866  
867  
868  
869  
870  
871  
872  
873  
874  
875  
876  
877  
878  
879

880  
881

Table 9: Example results comparing PyTorch, ground truth, baseline, and T2J (T2J) outputs.

Example ID	PyTorch	Ground Truth	Baseline	T2J
e1	6.61	12.7	9.78	8.79
e2	10.00	60.0	24.4	3.59
e3	8.57	17.3	4.14	5.18
e4	8.82	22.1	21.8	5.36
e5	8.48	35.0	35.2	6.53
e6	13.0	24.0	3.0	13.68
e7	8.23	5.0	90.0	14.53
m1	17.0	26.0	180.0	15.62
m3	180.0	180.0	4.0	1.05
m4	180.0	18.0	180.0	29.0
m5	102.0	57.0	180.0	19.5
m6	18.0	6.0	61.0	4.0
m7	78.0	65.0	47.9	62.0
m8	240.0	19.0	11.8	10.0
h1	0.48	4.0	2.5	10.0
h3	97.0	7.0	44.1	19.0
h4	31.1	180.0	180.0	15.0
h5	19.6	180.0	180.0	42.0
h6	0.31	9.0	4.0	180.0
h10	0.89	14.0	7.6	2.0
Total	1028.09	941.1	1271.22	466.83

903  
904  
905  
906  
907  
908  
909  
910  
911  
912  
913  
914  
915  
916  
917