

# Text-to-Code Generation with Modality-relative Pre-training

Fenia Christopoulou, Guchun Zhang, Gerasimos Lampouras

Huawei Noah’s Ark Lab, London, UK

{efstathia.christopoulou, guchun.zhang, gerasimos.lampouras}@huawei.com

## Abstract

Large pre-trained language models have recently been expanded and applied to programming language tasks with great success, often through further pre-training of a strictly-natural language model—where training sequences typically contain both natural and (linearised) programming language. Such approaches effectively map both modalities of the sequence into the same embedding space. However, programming language keywords (e.g. “while”) often have very strictly defined semantics. As such, transfer learning from their natural language usage may not necessarily be beneficial to their code application and vice versa. Assuming an already pre-trained language model, in this work we investigate how sequence tokens can be adapted and represented differently, depending on which modality they belong to, and to the ultimate benefit of the downstream task. We experiment with separating embedding spaces between modalities during further model pre-training with modality-relative training objectives. We focus on text-to-code generation and observe consistent improvements across two backbone models and two test sets, measuring  $\text{pass}@k$  and a novel incremental variation.<sup>1</sup>

## 1 Introduction

Increasingly more large pre-trained language models (Radford et al., 2018; Devlin et al., 2019; Raffel et al., 2020; Han et al., 2021, PLMs) based on the Transformer (Vaswani et al., 2017) architecture have been proposed and shown to achieve state-of-the-art results on a variety of Natural Language Processing (NLP) tasks. Lately, such models have been adapted to more specific language domains, e.g. Biomedical (Huang et al., 2019; Lee et al., 2020; Beltagy et al., 2019), Legal (Chalkidis et al., 2020), Cyber Security (Aghaei et al., 2022) and

<sup>1</sup>Code, data and models are publicly available at [https://github.com/huawei-noah/noah-research/tree/master/NLP/text2code\\_mrpt](https://github.com/huawei-noah/noah-research/tree/master/NLP/text2code_mrpt)

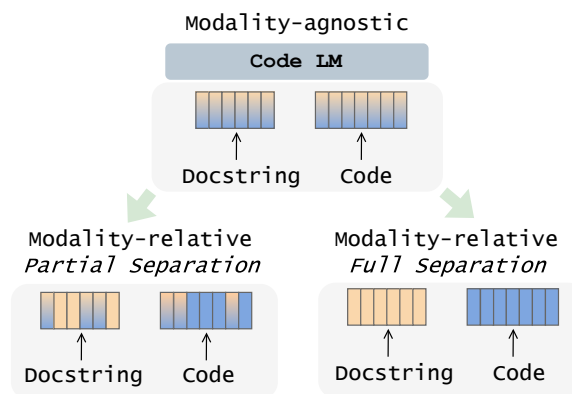


Figure 1: Overview of modality-agnostic/relative pre-training and training objectives.

Finance (Araci, 2019), while also expanding to include signals from modalities other than natural language, such as Vision (Trinh et al., 2019; Chen et al., 2022, 2020; Carion et al., 2020; Touvron et al., 2021; Dosovitskiy et al., 2021), Proteins (Brandes et al., 2022), Time Series (Wu et al., 2020a,b; Qin and Zong, 2022) and Code (Kanade et al., 2020; Wang et al., 2021; Feng et al., 2020; Guo et al., 2022; Phan et al., 2021).

In this work, we experiment with pre-trained language models specifically created for text-to-code generation, i.e. the task of program synthesis given Natural Language (NL) descriptions (e.g. problem definitions or docstrings). While some of the proposed models for this task adopt the encoder-decoder architecture (Li et al., 2022), most are trained as decoder-only Transformer models (Chen et al., 2021; Fried et al., 2023; Nijkamp et al., 2023; Li et al., 2023; Roziere et al., 2023). The latter approach is often preferred, since single-component architectures are well-suited for continuous training over vast amounts of raw non-annotated data, such as Github.<sup>2</sup>

When employing a single-component architecture, all tokens in the sequence are often vectorised

<sup>2</sup><https://github.com>

DESCR.	Return a greatest common divisor of two integers a and b
SIGN.	<code>def gcd(a: int, b: int) -&gt; int:</code>
BODY	<code>while b:     a, b = b, a % b return a</code>
TESTS	<code>gcd(3, 5) = 1 gcd(25, 15) = 5</code>

Table 1: Example from the HumanEval dataset.

via the same embedding layer; effectively assuming that all parts of the sequence share the same semantic space despite language or modality. This may be well suited for many NLP tasks or multilingual models, as transfer learning between languages has been shown to be beneficial (Conneau et al., 2020). However, we posit that due to their strict technical and logical implications, programming language tokens have distinct semantic meaning that does not necessarily benefit by transfer learning from natural language. To investigate this, we propose modality-relative embedding-spaces and training objectives. Specifically, our work treats the docstring and code subsequences of our data as separate *modalities*, and further tunes separate embeddings with different training objectives—initially tuning them on a shared space (see Figure 1).

As our test case, we focused on Python and carried out zero-shot experiments demonstrating that our modality-relative embeddings and training objectives achieve consistent improvements over two baselines on the HumanEval (Chen et al., 2021) and MBPP (Jain et al., 2021) datasets. We also introduce a novel variation of the  $\text{pass}@k$  (Chen et al., 2021) evaluation metric for program synthesis, namely *incremental pass@k*, that combines both synthesis and completion tasks to better differentiate models capabilities.

## 2 Methodology

In this work, we focus on the task of text-to-code generation, i.e. to synthesise running code that successfully solves the problem described in a NL description. Table 1 shows an example from the HumanEval dataset (Chen et al., 2021), which is typical for the task; the input consists of a NL description of the problem, followed by the function signature of the solution (the function argument types and the expected return type may or may not be included). In addition, the problems may be accompanied by some

unit tests, in the form of function calls with specific inputs and the corresponding expected outputs, e.g. `greatest_common_divisor(3, 5) = 1`. The model is then asked to produce the body of the code, where its proper functionality is evaluated against a number of held-out unit tests. This task is closely related to code completion, where the input also includes partial code.

### 2.1 Modality-agnostic Pre-Training

Following previous work (Chen et al., 2021), here we assume access to already pre-trained language models on raw natural language and code data (namely, Code LMs) as our starting point. These models have been pre-trained with modality-agnostic pre-training (MAPT), i.e. training that is agnostic with respect to the underlying modalities and all token embeddings are shared between the two. We employ the PyCodeGPT (Zan et al., 2022) model that is based on GPT-Neo (Black et al., 2021), pre-trained on an undisclosed subset of open source Github repositories in Python. In addition, we pre-train PanGu-Coder (Christopoulou et al., 2022) with regular Causal Language Modelling (CLM) (Radford et al., 2018) on a different batch of Python data from Github (see Appendix A for details regarding our data collection process).

Our proposed modality-relative training objectives and embedding-spaces are subsequently applied through *further pre-training* of the Code LMs on program synthesis-specific data, consisting solely of text-to-code pairs. Through this two-stage pre-training approach a model is able to learn how to encode both general code structures and natural language through raw data in the first stage, and later focus on how to best generate the correct output code given the NL input.

### 2.2 Modality-relative Continual Pre-Training

Assuming a modality-agnostic Code LM (e.g. PyCodeGPT or PanGu-Coder) that has attained general knowledge about NL and code during initial pre-training, we continue pre-training with a dedicated focus on the downstream task of text-to-code generation. The training data we use consist of functions crawled from existing, public GitHub repositories. If a function is accompanied by a docstring, we assume that to be a corresponding NL description of the function forming a text-to-code pair instance. Details on how the data were gathered and filtered can be found in Appendix A.




Figure 2: TEXT-CODE CLM, CODE-CLM, CORRUPT-CODE-CLM and PREFIX-CODE-CLM pre-training objectives.

Formally, we denote a training instance as a single sequence of tokens  $X = X_D + X_C = \{d_1, \dots, d_{N_D}, c_1, \dots, c_{N_C}\}$ , with  $X_D$  and  $X_C$  corresponding to all the tokens of the *docstring* modality and the *code* modality of the input, respectively. CLM is then applied as follows:

$$\mathcal{L}_{\text{CLM}}(X) = - \sum_{n=1}^N \log p(x_n | x_{<n}; \theta),$$

where  $N = N_D + N_C$  indicates the total number of tokens in the input sequence  $X$  (with  $N_D, N_C$  the number of *docstring* and *code* tokens, respectively) and  $\theta$  corresponds to the model parameters.

An important design decision was to formulate the data specifically for function-level code synthesis. As such, we ensured that the training data contain functions that are always accompanied by a natural language description (i.e. using the *docstring* as the function’s problem description), essentially aligning NL with code. Based on this data format, we propose to train the model via modality-relative pre-training (MRPT), i.e. treating the *docstring* and *code* subsequences of each instance distinctly.

### 2.2.1 Space Separation

To investigate our hypothesis that code tokens should be distinct from NL ones due to the strict semantic meaning of the former, we propose separating the embedding space  $E$  after MAPT (see Figure 1) through the following two strategies:

**Partial Embedding Separation (PES)** Reasonably, programming language-specific tokens such as `join`, `for`, `return`, `class`, etc, might suffer more from conflicting signals during modality-agnostic pre-training. On the other hand, tokens used in variables and function names should remain shared between NL and code, as transfer learning can benefit their encoding. In our first space separation approach, namely *Partial Embedding Separation* (PES), we associate separate embeddings  $e_{code}$  and  $e_{NL}$  only to tokens that appear in the programming language’s grammar. Since we are

using Python as our use-case, we extract all necessary tokens from the official Python grammar<sup>3</sup> and built-in functions.<sup>4</sup> Although this method creates two embeddings only for language-related tokens, minimally expanding the model’s embedding space, it requires manual effort and is dependent on the programming language one investigates.

**Full Embedding Separation (FES)** Since PES is programming language-specific and cannot capture the distinction among other tokens that might be equally important, our second approach separates the entire embedding space  $E$  of the modality-agnostic model, namely *Full Embedding Separation*. We thus associate a distinct embedding  $e_{NL}$  and  $e_{code}$  with each modality for the entire vocabulary of the model. This effectively results in doubling the number of embeddings of the model, i.e.  $|E_{NL}| + |E_{code}| = 2 * |E|$ .

In both PES and FES, the NL and code embeddings are initialised with their values as they stand after MAPT and are further trained through MRPT.

Alternatively, the separation of embeddings could be performed by using two distinct tokenizers for the NL and code part respectively, but that would forego the MAPT stage and any potential benefit resulting from it. Preliminary experiments were inconclusive on whether using distinct tokenizers was beneficial compared to starting from a modality-agnostic model with a shared embedding space. We opted to keep the MAPT stage constant across all settings, to keep comparisons fair.

### 2.2.2 Training Objectives

We train modality-aware models with a few training objectives, as shown in Figure 2 and detailed below. As a baseline objective we consider standard causal language modelling over the entire input sequence, which we denote as TEXT-CODE-CLM seen in the leftmost part of Figure 2. These training objectives

<sup>3</sup><https://docs.python.org/3/reference/grammar.html>

<sup>4</sup><https://docs.python.org/3/library/functions.html>

can then be further combined with the proposed embedding separation strategies.

**CODE-CLM: Causal Language Modelling on Code** We calculate the loss only on the code subsequence, following [Chen et al. \(2021\)](#), namely CODE-CLM, training implicitly on the docstring. Each code token is predicted based on all previous tokens, including the tokens of the docstring.

**CORRUPT-CODE-CLM: Corrupted Docstring** We also experiment with corrupting the docstring by randomly masking out some of its tokens. Specifically, a set of random tokens with indices  $X_M = \{\pi_1, \dots, \pi_M\}$  in the docstring are replaced with a mask ([MASK]), a random token, or the same token with 0.8/0.1/0.1 chance, respectively, similar to [Devlin et al. \(2019\)](#). Analogous to the previous objective, we do not measure any loss over the docstring but solely on the code that has to be generated based on the corrupted input.

**PREFIX-CODE-CLM: Bidirectional Attention on Docstring** Since the down-stream task is not reliant on next word prediction for the docstring, we experiment with allowing bidirectional attention on the docstring tokens, as we assume that additional context can result in better docstring representations. This is similar to how prefix language models work ([Dong et al., 2019](#); [Bao et al., 2020](#); [Guo et al., 2022](#)), with the difference that we do not calculate any loss over the prefix, which in our scenario, corresponds to the docstring ([Wang et al., 2022](#)). Due to the formatting of their input (see Section 3.1), in PyCodeGPT we consider both the docstring and the signature as part of the prefix, while in PanGu-Coder only the docstring is playing the role of the prefix.

### 3 Experimental Settings

#### 3.1 Data Formulation

We formulate each model’s input according to their original training data format. PanGu-Coder’s input is formed by combining a docstring and its corresponding code as follows, since it was designed to accommodate both code-only and text-code pairs during MAPT pre-training (see Appendix B).

```
[descr] docstring [python] signature
code [eoc]
```

where [descr] indicates the beginning of a description, [python] indicates the beginning of the code and [eoc] corresponds to the end of the

code sequence. On the other end, PyCodeGPT was trained on general purpose, raw code repositories, that follow the standard format of the docstring appearing after the code signature, with a start-of-sequence token ([sos]) at the beginning of each instance. We form its input as shown below:

```
[sos] signature [descr] docstring
[python] code [eoc]
```

#### 3.2 Inference

We follow the standard (left-to-right) decoding process used for auto-regressive language models, with temperature scaling ( $t$ ) and nucleus sampling ( $p$ ) ([Holtzman et al., 2020](#)). Inference adopts a prompt that is similar to the data format used during pre-training, for each model respectively (see Section 3.1), until the keyword [python], after which the model is requested to generate the problem solution. Generation continues until the [eoc] token is generated or a maximum length is reached. In the case of the PREFIX-CODE-CLM objective, we allow bidirectional attention during inference on the given prompt, similar to pre-training. Due to the nature of our applied tasks, we assume there is always a problem description available during inference. We remove any superfluous white-spaces and line breaks from the descriptions.

Formally, the prompt is a sequence of tokens  $P = P_D + P_S = \{d_1, \dots, d_{N_D}, s_1, \dots, s_{N_S}\}$  where  $N_D, N_S$  denote the number of tokens in the docstring and the signature, respectively. In the case of PyCodeGPT the prompt is formatted as  $P = P_S + P_D$ . The model then generates a continuation  $C'$  of the prompt in a left-to-right manner, decoding one token at a time while attending on previous.

$$C'(P) = \text{PANGU-CODER}(c'_t | c'_{<t}, d_{<N_D}, s_{<N_S})$$

$$C'(P) = \text{PYCODEGPT}(c'_t | c'_{<t}, s_{<N_S}, d_{<N_D})$$

#### 3.3 Evaluation

To evaluate our models, we consider two commonly used datasets for checking the functional correctness of generated programs: HumanEval ([Chen et al., 2021](#)), and the Mostly Basic Programming Problems ([Austin et al., 2021](#), MBPP).

HumanEval<sup>5</sup> contains 164 handcrafted Python problems accompanied by a set of held-out unit tests (average of 7.7 unit tests per problem), *all of which* must pass in order to count as a successful

<sup>5</sup><https://github.com/openai/human-eval>

solution. The dataset prompts include a problem description and a signature. In a similar vein, MBPP<sup>6</sup> is comprised of 974 programming problems (374 train, 90 validation, 500 test and 10 few-shot) designed to be solved by entry-level Python programmers. We use only the test set for evaluation.

We checked the overlap of any natural language prompt from all datasets in our training data with exact string matching (non-whitespace text) and found 1 contaminated example with HumanEval and none with the MBPP test set. Evaluation was performed via adapting the CodeGeeX framework (Zheng et al., 2023).<sup>7</sup>

### 3.4 Evaluation Metrics

In order to estimate model performance, we sample  $n$  programs/solutions per problem and calculate the unbiased estimator of  $\text{pass}@k$  for  $k = [1, 10, 100]$ , originally introduced in Kulal et al. (2019) and further adapted by Chen et al. (2021). In all experiments, we use greedy decoding for  $\text{pass}@1$  and a sample size of  $n = 200$  with  $p = 0.8$ ,  $t = 0.95$  for  $k = [10, 100]$ . Those values were selected via preliminary analysis and remain fixed for all experiments (no further tuning).

**Incremental  $\text{pass}@k$ :** A drawback of  $\text{pass}@k$  is that it only considers a code solution correct if it passes all provided unit tests, and provides no partial credit for incomplete programs however close they may be to a valid solution. Other metrics may provide partial credit based on syntactic or semantic overlap between the generated program and a reference, e.g. CodeBLUE (Ren et al., 2020) and CodeBERTScore (Zhou et al., 2023), but these do not directly check for functional correctness.

In an effort to provide more granular functional correctness, we propose and report a variation on the  $\text{pass}@k$  metric; incremental  $\text{pass}@k$ , which aims to measure the code completion capabilities of Code LMs in addition to full program synthesis. We automatically create an augmented test set for each dataset by using the provided code solution reference to construct partial code solutions as additional evaluation prompts; partial solutions are constructed by incrementally adding one line of the reference to the previous prompt. The  $\text{pass}@k$  metric is then calculated over the original and augmented prompts, reporting micro-averaged scores.

<sup>6</sup><https://github.com/google-research/google-research/tree/master/mbpp>

<sup>7</sup><https://github.com/THUDM/CodeGeeX>

### 3.5 Training details

We train models using the Adam optimiser (Kingma and Ba, 2015) with  $\beta_1 = 0.9$ ,  $\beta_2 = 0.95$  and weight decay of 0.01. For MAPT, the maximum learning rate is set to  $1e^{-4}$ , which is decayed by a cosine scheduler until  $5e^{-6}$ . For MRPT, the maximum and minimum learning rates are  $1e^{-5}$  and  $5e^{-6}$ , respectively. The gradients are clipped at 3.0 during modality-agnostic and 1.0 during modality-relative pre-training. For both models we set the maximum sequence length to 1,024 tokens, the batch size to 1,024 instances and warm-up to 1% of the total training steps.

For MAPT, PanGu-Coder is initialised with random weights and trained for 500K steps from scratch (more details in Appendix B). For MRPT, we initialise the models from their MAPT checkpoint and continue training for 5 epochs for PyCodeGPT and 10 epochs for PanGu-Coder. We report performance of all models at the end of modality-relative pre-training. Further technical details can be found in Appendix C.

## 4 Zero-shot Results

We first perform an ablation analysis on the 100M model variants to determine the gains based on our training data, the training objective and the embedding separation approaches. Our primary comparisons are against two baseline model variants (marked with a gray background), i.e. after modality-agnostic (MAPT) training and after standard next-token-prediction (TEXT-CODE) on the entire sequence. The latter facilitates comparison against continual training on the given data for the same number of epochs as the other objectives.

In Table 2a, we evaluate PyCodeGPT across the two datasets and four objectives. Additional training with text-to-code pairs improves performance up to +1.2, +2.4 and +4.3 points for  $k=1,10,100$  over MAPT training on HumanEval and correspondingly +1.6, +3.0 and +2.9 on MBPP. PyCodeGPT benefits more from the task-specific data as it was not exposed to them during MAPT. Embedding separation additionally improves up to +3.0, +1.1, +2.1  $\text{pass}@k$  for HumanEval and up to +0.2, +0.8, +1.7 for MBPP. We observe stronger benefits when looking at incremental  $\text{pass}@k$ , with gains up to +6.2, +9.5, +5.9 on HumanEval and +4.4, +4.3, +2.2 on MBPP.

Moving on to Table 2b we show results for PanGu-Coder model of 100M parameters under

	SEPARATION	HUMAN EVAL			INCR HUMAN EVAL			MBPP			INCR MBPP		
		P@1	P@10	P@100	P@1	P@10	P@100	P@1	P@10	P@100	P@1	P@10	P@100
MAPT	-	9.15	13.26	20.80	13.21	29.82	46.13	7.40	19.40	37.87	21.30	45.59	63.25
TEXT CODE	-	10.37	<b>15.66</b>	<b>25.19</b>	11.69	27.28	52.63	<b>9.00</b>	22.46	40.81	20.23	49.07	<b>68.51</b>
	PARTIAL	<b>11.59</b>	14.96	23.83	<b>16.66</b>	32.32	51.75	7.80	22.63	41.70	23.08	49.74	68.12
	FULL	<b>11.59</b>	14.81	23.94	15.30	<b>33.36</b>	<b>54.87</b>	7.80	<b>22.69</b>	<b>41.77</b>	<b>23.39</b>	<b>50.77</b>	68.22
CODE	-	12.80	<b>16.81</b>	26.94	<b>17.98</b>	34.54	54.61	8.80	23.29	41.46	23.81	52.37	70.10
	PARTIAL	12.80	16.47	26.25	17.13	32.83	52.42	<b>9.00</b>	23.28	<b>42.53</b>	<b>24.44</b>	52.35	70.49
	full	<b>13.41</b>	15.97	<b>27.32</b>	17.88	<b>36.82</b>	<b>57.91</b>	8.00	<b>23.32</b>	42.47	24.39	<b>53.42</b>	<b>70.50</b>
CORRUPT CODE	-	<b>11.59</b>	16.15	26.34	16.39	33.97	54.84	8.60	<b>23.00</b>	<b>42.04</b>	22.46	51.36	69.60
	PARTIAL	<b>11.59</b>	<b>15.96</b>	27.04	16.24	33.45	54.07	<b>9.20</b>	22.59	41.36	23.57	50.56	68.32
	FULL	10.98	15.73	<b>27.11</b>	<b>17.51</b>	<b>36.67</b>	<b>58.59</b>	<b>9.20</b>	22.92	41.15	<b>24.67</b>	<b>52.76</b>	<b>70.15</b>
PREFIX CODE	-	7.32	9.87	16.84	14.37	32.69	52.93	8.60	20.93	39.89	23.63	51.39	70.21
	PARTIAL	9.15	13.53	<b>22.18</b>	14.48	33.90	53.10	<b>9.20</b>	19.84	39.61	24.19	50.42	68.02
	FULL	<b>9.76</b>	<b>13.67</b>	22.00	<b>16.30</b>	<b>36.10</b>	<b>56.56</b>	<b>9.20</b>	<b>21.45</b>	<b>40.59</b>	<b>24.45</b>	<b>53.10</b>	<b>70.72</b>

(a) PyCodeGPT 100M model performance.

	SEPARATION	HUMAN EVAL			INCR HUMAN EVAL			MBPP			INCR MBPP		
		P@1	P@10	P@100	P@1	P@10	P@100	P@1	P@10	P@100	P@1	P@10	P@100
MAPT	-	9.76	17.07	28.88	20.70	48.04	63.08	11.60	24.99	44.84	5.93	28.80	44.63
TEXT CODE	-	12.20	18.97	29.50	22.93	49.96	64.23	11.80	26.18	45.73	11.78	38.99	61.38
	PARTIAL	11.59	<b>19.49</b>	<b>31.70</b>	22.91	49.60	<b>64.91</b>	<b>12.00</b>	26.35	46.76	11.93	39.43	60.87
	FULL	<b>13.41</b>	19.31	29.61	<b>23.54</b>	<b>49.98</b>	63.48	11.40	<b>26.65</b>	<b>46.82</b>	<b>12.80</b>	<b>41.93</b>	<b>63.27</b>
CODE	-	12.20	19.63	<b>32.45</b>	23.65	50.25	64.89	12.20	26.69	46.61	<b>14.76</b>	<b>44.75</b>	<b>64.84</b>
	PARTIAL	<b>13.41</b>	<b>20.27</b>	31.27	<b>24.44</b>	<b>50.63</b>	65.46	12.00	<b>27.14</b>	<b>47.08</b>	13.27	41.79	63.06
	FULL	12.20	19.74	30.82	23.81	50.51	<b>66.21</b>	<b>12.60</b>	27.01	46.14	14.23	43.80	64.21
CORRUPT CODE	-	12.80	19.06	28.20	23.83	49.89	63.98	<b>12.20</b>	26.46	<b>46.57</b>	15.17	45.41	<b>65.95</b>
	PARTIAL	<b>13.41</b>	<b>19.21</b>	<b>29.70</b>	23.83	<b>50.12</b>	<b>65.22</b>	11.80	26.44	45.15	13.53	42.36	62.85
	FULL	12.80	18.65	28.54	<b>23.91</b>	49.48	64.02	11.60	<b>26.96</b>	44.85	<b>15.49</b>	<b>46.38</b>	65.60
PREFIX CODE	-	10.98	18.33	27.56	23.27	49.98	63.78	11.00	26.10	45.90	<b>11.23</b>	<b>38.49</b>	<b>60.80</b>
	PARTIAL	<b>12.80</b>	19.26	<b>30.77</b>	<b>23.88</b>	<b>50.16</b>	64.23	10.80	26.20	46.00	10.39	35.19	57.44
	FULL	<b>12.80</b>	<b>19.37</b>	30.21	23.59	49.64	<b>64.95</b>	<b>11.20</b>	<b>26.28</b>	<b>46.27</b>	9.13	33.67	56.52

(b) PanGu-Coder 100M model performance.

Table 2: Pass@ $k$  and Incremental pass@ $k$  across objectives, embedding separation strategies and datasets. **Bold** numbers denote best performance across separation methods and **highlighted** denote overall best for each metric.

the same datasets and methods. Firstly, we notice that we get an improvement of +2.4, +1.9 and +0.6 points for  $k=1,10,100$  over MAPT training on HumanEval and correspondingly +0.2, +1.1 and +0.8 on MBPP. We attribute this gain to the formulation of the data in a format that exactly matches the task at hand. Overall, embedding separation consistently offers additional improvements up to +1.2, +1.3, +2.9 on HumanEval and +0.8, +0.9, +1.3 on MBPP. In terms of incremental pass@ $k$ , again we observe larger gains over TEXT-CODE, mostly on MBPP with +3.7, +7.3 and +4.5 points.

Across training objectives, we notice that in the majority of settings, CODE-CLM outperforms across the board. Across datasets, separation offers the most notable gains in the PREFIX-CODE-CLM objective, probably because it helps the model to better adapt to bidirectional attention on the docstring. We note that most of the gains from em-

bedding separation seem to affect pass@10 and pass@100, which implies that separation helps increase the expressiveness of the model but not the MAP (maximum a posteriori) solution as much.

#### 4.1 Scaling Up

To determine whether our observations regarding embedding separation hold for different model sizes, we also test our hypothesis on a larger version of PanGu-Coder, consisting of 350M parameters, for the reported best training objective CODE-CLM of its smaller variant (see Table 2). In Table 3, we observe again improvements over MAPT with +2.4, +1.2, +2.0 pass@1/10/100 on HumanEval. Similar trends can be observed on the MBPP dataset, though the differences appear smaller due to its larger sample size and difficulty. More evident improvements are noticed over no separation with +2.4, +1.9, +3.8 on HumanEval and +1.2, +0.1,

	SEPARATION	HUMAN EVAL			INCR HUMAN EVAL			MBPP			INCR MBPP		
		P@1	P@10	P@100	P@1	P@10	P@100	P@1	P@10	P@100	P@1	P@10	P@100
MAPT	-	16.46	24.51	35.39	27.82	55.68	68.34	18.80	35.36	53.24	8.14	35.26	57.77
TEXT CODE	-	<b>18.90</b>	25.79	37.46	<b>31.39</b>	56.36	69.22	17.60	37.59	54.37	<b>10.95</b>	<b>42.27</b>	<b>66.62</b>
	partial	18.29	25.54	37.82	30.81	<b>56.39</b>	<b>69.44</b>	<b>18.00</b>	<b>37.70</b>	<b>55.60</b>	9.97	39.19	64.93
	full	18.29	<b>26.48</b>	<b>41.30</b>	30.85	56.28	68.81	16.40	37.43	55.35	9.71	40.22	65.49
CODE	-	20.73	25.96	37.79	32.16	56.66	69.88	<b>18.80</b>	<b>37.54</b>	<b>54.99</b>	12.27	44.87	<b>68.65</b>
	partial	<b>21.34</b>	26.27	37.94	<b>32.88</b>	56.82	69.20	18.60	37.36	54.39	<b>12.38</b>	<b>45.37</b>	68.58
	full	18.29	<b>26.98</b>	<b>39.94</b>	30.10	<b>57.12</b>	<b>70.15</b>	<b>18.80</b>	37.45	54.33	9.16	37.84	63.13

Table 3: PanGu-Coder 350M CodeLM performance.

TEST SET	HUMANEVAL	MBPP DEV	MBPP TEST
HUMANEVAL	-	0.099	0.163
MBPP DEV	-	-	0.046
MBPP TEST	-	-	-

Table 4: Average Pearson correlation between pass@1 performance of different model checkpoints.

+1.2 on MBPP, respectively. Incremental pass@ $k$  shows +1.4, +0.7, +0.9 and +1.4, +3.1, +2.0 over each dataset. We present a comparison against other Code LMs in Appendix D.

## 5 Visualising Separated Embeddings

We attempt to qualitatively check the embedding space separation by visualising token embeddings on a 2D plane via T-SNE plots. Figure 3 illustrates the 20 closest neighbours of the tokens `open` and `join` in the two modality spaces, as measured using cosine similarity for PanGu-Coder 350M model with partial embedding separation (PES).

For `open`, in the docstring space we observe that its representation is separate from `close` or other build-in operations, e.g. `get`. On the contrary, in the code space `open` and `close` are grouped together with other operations that are used in a similar way in code. Analogously, `join` in the docstring space is close to words with a similar natural language meaning and surface form, while in the code space its representation is close to functional-like tokens in Python such as `remove`, `get`, `split`, etc. We observe similar behaviours for other tokens; see additional plots in Appendix E.

## 6 Generalisation Analysis and Discussion

We performed an analysis over the performance of all the checkpoints of our trained models (checkpoints were cached every 10K steps), to determine whether performance on any test set was predictive on the performance of other sets; we included HumanEval and both development and test sets of

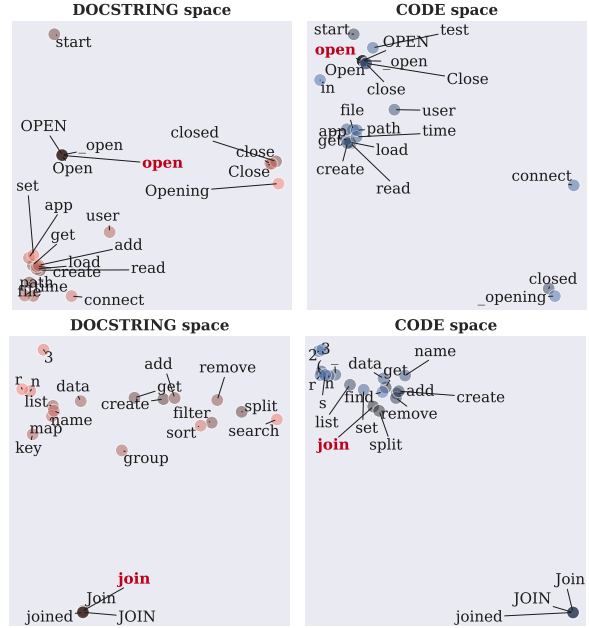


Figure 3: T-SNE plots of the top 20 neighbours of the tokens `open` (top) and `join` (bottom) on different embedding spaces for PanGu-Coder 350M parameter model with partial embedding space separation (PES).

MBPP. Pearson analysis between checkpoint performance suggests very weak correlation between model performance overall; consult Table 4.

This seems to indicate that the distribution of problems that are covered by any one test set is relatively unique, and raises the question whether comparisons over any particular dataset should be expected to generalise on unseen problems when models are deployed. Potentially, this is partially explained by the small size of the test sets, but could be indicative of a larger problem; that due to the specificity of problems contained in these test sets, they do not evaluate generalisable model coding capabilities as much as they measure whether they have been exposed to very similar problems and definitions during pre-training. To offer an example, if the model has not been exposed to problems regarding Fibonacci sequences during

pre-training, it is unrealistic to expect it to generate a relevant code solution in a zero-shot setting. Similarly, whether a model can generate a solution to a Fibonacci problem offers no indication of its capabilities to solve other problems.

## 7 Related Work

Recently, there has been an increasing interest in extending proven natural language understanding and generation methods to code understanding and generation tasks. CodeBERT (Feng et al., 2020), for instance, was trained using a combination of Masked Language Modelling inspired by Devlin et al. (2019, BERT) and Replaced Token Detection from Clark et al. (2020, ELECTRA). CodeT5 (Wang et al., 2021) and PYMT5 (Clement et al., 2020) were built on top of (Raffel et al., 2020, T5) while UniXcoder (Guo et al., 2022) was based on UniLM (Dong et al., 2019) and combines three pre-training objectives using different attention masks to control access to surrounding context for a token to be predicted. Focused on the task of text-to-code generation, Chen et al. (2021) introduced CodeX, a set of GPT-based language models trained on publicly available code from GitHub, up to 12B parameters in size. Li et al. (2022) introduced AlphaCode, a set of sequence-to-sequence models with up to 41B parameters, trained on data from programming competitions, e.g. Codeforces<sup>8</sup> as well as GitHub code in several programming languages. CodeGen (Nijkamp et al., 2023) was proposed as a conversational text-to-code approach using large language models with sizes of up to 16B parameters. The authors proposed three model variants, trained on The Pile (Gao et al., 2021), continuously trained on BigQuery (with 6 programming languages) and finally trained on Python-only data. InCoder (Fried et al., 2023) extends left-to-right code generation with an infilling training objective, similar to Bavarian et al. (2022), and is able to predict spans of partial programs as well. SantaCoder (Allal et al., 2023) is one of the latest Code LMs, sizing up to 1.1B parameters, supporting Python, Java, and JavaScript. It was trained with multi-query attention and the fill-in-the-middle (FIM) objective (Bavarian et al., 2022) on The Stack (Kocetkov et al., 2022), a 3 TB publicly available dataset supporting 358 programming languages collected from permissively-licensed source code files from Github. StarCoder (Li et al.,

<sup>8</sup><https://codeforces.com/>

2023) is a 15.5B parameter model, similarly pre-trained with multi-query attention and FIM on 80+ programming languages from The Stack. With the exception of CodeX (Chen et al., 2021) calculating CLM loss on code exclusively, this is the first work to consider modality-relative embedding separation and training objectives.

Methods that consider NL and code as different modalities mostly focus on taking into account different *views* of code. For instance, GraphCodeBERT (Guo et al., 2021), noted that previous pre-trained models treat code snippets as sequences of tokens while ignoring the inherent structure of code. They presented GraphCodeBERT, which showed that incorporating the *data flow*, i.e. a semantic-level structure of code extracted from the Abstract Syntax Tree (AST), leads to richer code representations. Jiang et al. (2021, TREEBert) instead, used the actual AST together with code snippets. To the best of our knowledge we are not aware of any work assigning different embeddings to tokens of the same sequence depending on which type of text (natural language/code) they appear.

## 8 Conclusion

Existing CodeLMs consider both code and natural language as a single modality, mapping them into a shared embedding space. However, in this work, we posit that the semantics and usage of tokens can differ between code and NL, requesting for a possible space separation. As such, we proposed to consider code and natural language as different modalities for the task of text-to-code generation by introducing modality-aware embedding separation strategies and training objectives. In detail, assuming a general CodeLM trained on raw data—where modalities are shared—we continue training on task-specific data with separated embedding spaces. We present partial separation, which targets language-specific tokens, as well as full separation that duplicates the entire model’s vocabulary. In addition, we proposed incremental pass@*k*, as a variant of the standard metric that evaluates the code completion capabilities of models.

Zero-shot evaluation on the HumanEval and MBPP datasets, with two 100M and one 350M parameters models, indicate that embedding space separation improves code generation across different objectives. We also observe that further pre-training on data formatted to match the target task consistently boosts performance.



Although there is no clear winner among partial versus full space separation, each method comes with advantages and drawbacks. While partial separation minimally increases the model’s vocabulary, it requires manual effort for each target programming language. On the other end, full separation is automatic but significantly increases vocabulary size. Given these observations, future work could target efficient approaches for separating only certain tokens to strike a balance between the two.

## Limitations

We enumerate the limitations of this work as part of our experimental process. Firstly, our analysis was focused on Python only. While our proposed methods are orthogonal to the programming language they are applied to, it remains to be confirmed that our findings generalise to other languages. Overall, our observations should stand for high-level programming languages that resemble natural language. Secondly, our trained models are limited to function-level text-to-code generation (both program synthesis and code completion), as we opted for a particular use case to study the connection between programming and natural language. These models are unable to perform multi-turn generation or generation of multiple code functions given a problem description, as mentioned in Section 3.3. We leave such explorations for future work.

## References

- Ehsan Aghaei, Xi Niu, Waseem Shadid, and Ehab Al-Shaer. 2022. [Language model for text analytic in cybersecurity](#).
- Loubna Ben Allal, Raymond Li, Denis Kocetkov, Chenghao Mou, Christopher Akiki, Carlos Munoz Ferrandis, Niklas Muennighoff, Mayank Mishra, Alex Gu, Manan Dey, Logesh Kumar Umapathi, Carolyn Jane Anderson, Yangtian Zi, Joel Lamy Poirier, Hailey Schoelkopf, Sergey Troshin, Dmitry Abulkhanov, Manuel Romero, Michael Lappert, Francesco De Toni, Bernardo García del Río, Qian Liu, Shamik Bose, Urvashi Bhattacharyya, Terry Yue Zhuo, Ian Yu, Paulo Villegas, Marco Zocca, Sourab Mangrulkar, David Lansky, Huu Nguyen, Danish Contractor, Luis Villa, Jia Li, Dzmitry Bahdanau, Yacine Jernite, Sean Hughes, Daniel Fried, Arjun Guha, Harm de Vries, and Leandro von Werra. 2023. Santacoder: don’t reach for the stars! *arXiv preprint arXiv:2301.03988*.
- Dogu Araci. 2019. [Finbert: Financial sentiment analysis with pre-trained language models](#). *CoRR*, abs/1908.10063.
- Jacob Austin, Augustus Odena, Maxwell Nye, Maarten Bosma, Henryk Michalewski, David Dohan, Ellen Jiang, Carrie Cai, Michael Terry, Quoc Le, et al. 2021. Program synthesis with large language models. *arXiv preprint arXiv:2108.07732*.
- Hangbo Bao, Li Dong, Furu Wei, Wenhui Wang, Nan Yang, Xiaodong Liu, Yu Wang, Jianfeng Gao, Songhao Piao, Ming Zhou, and Hsiao-Wuen Hon. 2020. [UniLMv2: Pseudo-masked language models for unified language model pre-training](#). In *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 642–652. PMLR.
- Mohammad Bavarian, Heewoo Jun, Nikolas Tezak, John Schulman, Christine McLeavey, Jerry Tworek, and Mark Chen. 2022. Efficient training of language models to fill in the middle. *arXiv preprint arXiv:2207.14255*.
- Iz Beltagy, Arman Cohan, and Kyle Lo. 2019. [Scibert: Pretrained contextualized embeddings for scientific text](#). *CoRR*, abs/1903.10676.
- Sid Black, Leo Gao, Phil Wang, Connor Leahy, and Stella Biderman. 2021. [GPT-Neo: Large Scale Autoregressive Language Modeling with Mesh-Tensorflow](#). If you use this software, please cite it using these metadata.
- Nadav Brandes, Dan Ofer, Yam Peleg, Nadav Rapoport, and Michal Linial. 2022. ProteinBERT: A universal deep-learning model of protein sequence and function. *Bioinformatics*, 38(8):2102–2110.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in Neural Information Processing Systems*, 33:1877–1901.
- Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. 2020. End-to-end object detection with transformers. In *European Conference on Computer Vision*, pages 213–229. Springer.
- Federico Cassano, John Gouwar, Daniel Nguyen, Sydney Nguyen, Luna Phipps-Costin, Donald Pinckney, Ming Ho Yee, Yangtian Zi, Carolyn Jane Anderson, Molly Q Feldman, et al. 2022. A scalable and extensible approach to benchmarking n2code for 18 programming languages. *arXiv preprint arXiv:2208.08227*.
- Ilias Chalkidis, Manos Fergadiotis, Prodromos Malakasiotis, Nikolaos Aletras, and Ion Androutsopoulos. 2020. [LEGAL-BERT: The muppets straight out of law school](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 2898–2904, Online. Association for Computational Linguistics.

- Mark Chen, Alec Radford, Rewon Child, Jeffrey Wu, Heewoo Jun, David Luan, and Ilya Sutskever. 2020. [Generative pretraining from pixels](#). In *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 1691–1703. PMLR.
- Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde de Oliveira Pinto, Jared Kaplan, Harrison Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, Alex Ray, Raul Puri, Gretchen Krueger, Michael Petrov, Heidy Khlaaf, Girish Sastry, Pamela Mishkin, Brooke Chan, Scott Gray, Nick Ryder, Mikhail Pavlov, Alethea Power, Lukasz Kaiser, Mohammad Bavarian, Clemens Winter, Philippe Tillet, Felipe Petroski Such, Dave Cummings, Matthias Plappert, Fotios Chantzis, Elizabeth Barnes, Ariel Herbert-Voss, William Hebgen Guss, Alex Nichol, Alex Paino, Nikolas Tezak, Jie Tang, Igor Babuschkin, Suchir Balaji, Shantanu Jain, William Saunders, Christopher Hesse, Andrew N. Carr, Jan Leike, Joshua Achiam, Vedant Misra, Evan Morikawa, Alec Radford, Matthew Knight, Miles Brundage, Mira Murati, Katie Mayer, Peter Welinder, Bob McGrew, Dario Amodei, Sam McCandlish, Ilya Sutskever, and Wojciech Zaremba. 2021. [Evaluating large language models trained on code](#). *CoRR*, abs/2107.03374.
- Xiangning Chen, Cho-Jui Hsieh, and Boqing Gong. 2022. [When vision transformers outperform resnets without pre-training or strong data augmentations](#). In *International Conference on Learning Representations*.
- Fenia Christopoulou, Gerasimos Lampouras, Milan Gritta, Guchun Zhang, Yinpeng Guo, Zhongqi Li, Qi Zhang, Meng Xiao, Bo Shen, Lin Li, et al. 2022. [Pangu-coder: Program synthesis with function-level language modeling](#). *arXiv preprint arXiv:2207.11280*.
- Kevin Clark, Minh-Thang Luong, Quoc V. Le, and Christopher D. Manning. 2020. [Electra: Pre-training text encoders as discriminators rather than generators](#). In *International Conference on Learning Representations*.
- Colin Clement, Dawn Drain, Jonathan Timcheck, Alexey Svyatkovskiy, and Neel Sundaresan. 2020. [PyMT5: multi-mode translation of natural language and python code with transformers](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 9052–9065, Online. Association for Computational Linguistics.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. [Unsupervised cross-lingual representation learning at scale](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8440–8451, Online. Association for Computational Linguistics.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Li Dong, Nan Yang, Wenhui Wang, Furu Wei, Xiaodong Liu, Yu Wang, Jianfeng Gao, Ming Zhou, and Hsiao-Wuen Hon. 2019. *Unified Language Model Pre-Training for Natural Language Understanding and Generation*. Curran Associates Inc., Red Hook, NY, USA.
- Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. 2021. [An image is worth 16x16 words: Transformers for image recognition at scale](#). In *International Conference on Learning Representations*.
- Zhangyin Feng, Daya Guo, Duyu Tang, Nan Duan, Xiaocheng Feng, Ming Gong, Linjun Shou, Bing Qin, Ting Liu, Daxin Jiang, and Ming Zhou. 2020. [CodeBERT: A pre-trained model for programming and natural languages](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1536–1547, Online. Association for Computational Linguistics.
- Daniel Fried, Armen Aghajanyan, Jessy Lin, Sida Wang, Eric Wallace, Freda Shi, Ruiqi Zhong, Wen-tau Yih, Luke Zettlemoyer, and Mike Lewis. 2023. [InCoder: A generative model for code infilling and synthesis](#). In *International Conference on Learning Representations*.
- Leo Gao, Stella Biderman, Sid Black, Laurence Golding, Travis Hoppe, Charles Foster, Jason Phang, Horace He, Anish Thite, Noa Nabeshima, Shawn Presser, and Connor Leahy. 2021. [The pile: An 800gb dataset of diverse text for language modeling](#). *CoRR*, abs/2101.00027.
- Daya Guo, Shuai Lu, Nan Duan, Yanlin Wang, Ming Zhou, and Jian Yin. 2022. [UniXcoder: Unified cross-modal pre-training for code representation](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 7212–7225, Dublin, Ireland. Association for Computational Linguistics.
- Daya Guo, Shuo Ren, Shuai Lu, Zhangyin Feng, Duyu Tang, Shujie LIU, Long Zhou, Nan Duan, Alexey Svyatkovskiy, Shengyu Fu, Michele Tufano, Shao Kun Deng, Colin Clement, Dawn Drain, Neel Sundaresan, Jian Yin, Daxin Jiang, and Ming Zhou. 2021. [Graphcode{bert}: Pre-training code representations with data flow](#). In *International Conference on Learning Representations*.

- Xu Han, Zhengyan Zhang, Ning Ding, Yuxian Gu, Xiao Liu, Yuqi Huo, Jiezhong Qiu, Liang Zhang, Wentao Han, Minlie Huang, Qin Jin, Yanyan Lan, Yang Liu, Zhiyuan Liu, Zhiwu Lu, Xipeng Qiu, Ruihua Song, Jie Tang, Ji-Rong Wen, Jinhui Yuan, Wayne Xin Zhao, and Jun Zhu. 2021. [Pre-trained models: Past, present and future](#). *AI Open*.
- Ari Holtzman, Jan Buys, Li Du, Maxwell Forbes, and Yejin Choi. 2020. [The curious case of neural text de-generation](#). In *International Conference on Learning Representations*.
- Kexin Huang, Jaan Alntosaar, and Rajesh Ranganath. 2019. [Clinicalbert: Modeling clinical notes and predicting hospital readmission](#). *CoRR*, abs/1904.05342.
- Naman Jain, Skanda Vaidyanath, Arun Shankar Iyer, Nagarajan Natarajan, Suresh Parthasarathy, Sriram K. Rajamani, and Rahul Sharma. 2021. [Jigsaw: Large language models meet program synthesis](#). *CoRR*, abs/2112.02969.
- Xue Jiang, Zhuoran Zheng, Chen Lyu, Liang Li, and Lei Lyu. 2021. [Treebert: A tree-based pre-trained model for programming language](#). In *Uncertainty in Artificial Intelligence*, pages 54–63. PMLR.
- Aditya Kanade, Petros Maniatis, Gogul Balakrishnan, and Kensen Shi. 2020. [Pre-trained contextual embedding of source code](#).
- Diederik P. Kingma and Jimmy Ba. 2015. [Adam: A method for stochastic optimization](#). In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- Denis Kocetkov, Raymond Li, Loubna Ben Allal, Jia Li, Chenghao Mou, Carlos Muñoz Ferrandis, Yacine Jernite, Margaret Mitchell, Sean Hughes, Thomas Wolf, et al. 2022. [The stack: 3 tb of permissively licensed source code](#). *arXiv preprint arXiv:2211.15533*.
- Taku Kudo and John Richardson. 2018. [SentencePiece: A simple and language independent subword tokenizer and detokenizer for neural text processing](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 66–71, Brussels, Belgium. Association for Computational Linguistics.
- Sumith Kulal, Panupong Pasupat, Kartik Chandra, Mina Lee, Oded Padon, Alex Aiken, and Percy S Liang. 2019. [Spoc: Search-based pseudocode to code](#). In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc.
- Jinhyuk Lee, Wonjin Yoon, Sungdong Kim, Donghyeon Kim, Sunkyu Kim, Chan Ho So, and Jaewoo Kang. 2020. [Biobert: a pre-trained biomedical language representation model for biomedical text mining](#). *Bioinformatics*, 36(4):1234–1240.
- Raymond Li, Loubna Ben Allal, Yangtian Zi, Niklas Muennighoff, Denis Kocetkov, Chenghao Mou, Marc Marone, Christopher Akiki, Jia Li, Jenny Chim, Qian Liu, Evgenii Zheltonozhskii, Terry Yue Zhuo, Thomas Wang, Olivier Dehaene, Mishig Davaadorj, Joel Lamy-Poirier, João Monteiro, Oleh Shliazhko, Nicolas Gontier, Nicholas Meade, Armel Zebaze, Ming-Ho Yee, Logesh Kumar Umapathi, Jian Zhu, Benjamin Lipkin, Muhtasham Oblokulov, Zhiruo Wang, Rudra Murthy, Jason Stillerman, Siva Sankalp Patel, Dmitry Abulkhanov, Marco Zocca, Manan Dey, Zhihan Zhang, Nour Fahmy, Urvashi Bhattacharyya, Wenhao Yu, Swayam Singh, Sasha Luccioni, Paulo Villegas, Maxim Kunakov, Fedor Zhdanov, Manuel Romero, Tony Lee, Nadav Timor, Jennifer Ding, Claire Schlesinger, Hailey Schoelkopf, Jan Ebert, Tri Dao, Mayank Mishra, Alex Gu, Jennifer Robinson, Carolyn Jane Anderson, Brendan Dolan-Gavitt, Danish Contractor, Siva Reddy, Daniel Fried, Dzmitry Bahdanau, Yacine Jernite, Carlos Muñoz Ferrandis, Sean Hughes, Thomas Wolf, Arjun Guha, Leandro von Werra, and Harm de Vries. 2023. [Starcoder: may the source be with you!](#)
- Yujia Li, David Choi, Junyoung Chung, Nate Kushman, Julian Schrittwieser, Rémi Leblond, Tom Eccles, James Keeling, Felix Gimeno, Agustin Dal Lago, Thomas Hubert, Peter Choy, Cyprien de Masson d’Autume, Igor Babuschkin, Xinyun Chen, Po-Sen Huang, Johannes Welbl, Sven Gowal, Alexey Cherepanov, James Molloy, Daniel J. Mankowitz, Esme Sutherland Robson, Pushmeet Kohli, Nando de Freitas, Koray Kavukcuoglu, and Oriol Vinyals. 2022. [Competition-level Code Generation with AlphaCode](#).
- Erik Nijkamp, Bo Pang, Hiroaki Hayashi, Lifu Tu, Huan Wang, Yingbo Zhou, Silvio Savarese, and Caiming Xiong. 2023. [Codegen: An open large language model for code with multi-turn program synthesis](#). In *The Eleventh International Conference on Learning Representations*.
- Long Phan, Hieu Tran, Daniel Le, Hieu Nguyen, James Annibal, Alec Peltekian, and Yanfang Ye. 2021. [Co-Text: Multi-task learning with code-text transformer](#). In *Proceedings of the 1st Workshop on Natural Language Processing for Programming (NLP4Prog 2021)*, pages 40–47, Online. Association for Computational Linguistics.
- Jiahao Qin and Lu Zong. 2022. [TS-BERT: A fusion model for pre-training time series-text representations](#).
- Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. 2018. [Improving language understanding by generative pre-training](#).
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020. [Exploring the limits of transfer learning with a unified text-to-text transformer](#). *Journal of Machine Learning Research*, 21(140):1–67.

- Shuo Ren, Daya Guo, Shuai Lu, Long Zhou, Shujie Liu, Duyu Tang, Neel Sundaresan, Ming Zhou, Ambrosio Blanco, and Shuai Ma. 2020. CodeBLUE: a method for automatic evaluation of code synthesis. *arXiv preprint arXiv:2009.10297*.
- Baptiste Roziere, Jonas Gehring, Fabian Gloeckle, Sten Sootla, Itai Gat, Xiaoqing Ellen Tan, Yossi Adi, Jingyu Liu, Tal Remez, Jérémy Rapin, et al. 2023. Code llama: Open foundation models for code. *arXiv preprint arXiv:2308.12950*.
- Hugo Touvron, Matthieu Cord, Matthijs Douze, Francisco Massa, Alexandre Sablayrolles, and Herve Jegou. 2021. Training data-efficient image transformers & distillation through attention. In *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 10347–10357. PMLR.
- Trieu H. Trinh, Minh-Thang Luong, and Quoc V. Le. 2019. Selfie: Self-supervised pretraining for image embedding. *CoRR*, abs/1906.02940.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in Neural Information Processing Systems*, 30.
- Thomas Wang, Adam Roberts, Daniel Hesslow, Teven Le Scao, Hyung Won Chung, Iz Beltagy, Julien Launay, and Colin Raffel. 2022. What language model architecture and pretraining objective works best for zero-shot generalization? In *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pages 22964–22984. PMLR.
- Yue Wang, Weishi Wang, Shafiq Joty, and Steven C.H. Hoi. 2021. CodeT5: Identifier-aware unified pre-trained encoder-decoder models for code understanding and generation. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 8696–8708, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Neo Wu, Bradley Green, Xue Ben, and Shawn O’Banion. 2020a. Deep transformer models for time series forecasting: The influenza prevalence case. *CoRR*, abs/2001.08317.
- Zonghan Wu, Shirui Pan, Guodong Long, Jing Jiang, Xiaojun Chang, and Chengqi Zhang. 2020b. Connecting the dots: Multivariate time series forecasting with graph neural networks. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, KDD ’20*, page 753–763, New York, NY, USA. Association for Computing Machinery.
- Daoguang Zan, Bei Chen, Dejian Yang, Zeqi Lin, Minsu Kim, Bei Guan, Yongji Wang, Weizhu Chen, and Jian-Guang Lou. 2022. CERT: Continual pre-training on sketches for library-oriented code generation. In *The 2022 International Joint Conference on Artificial Intelligence*.
- Qinkai Zheng, Xiao Xia, Xu Zou, Yuxiao Dong, Shan Wang, Yufei Xue, Lei Shen, Zihan Wang, Andi Wang, Yang Li, Teng Su, Zhilin Yang, and Jie Tang. 2023. Codegex: A pre-trained model for code generation with multilingual benchmarking on humaneval-x. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, KDD ’23*, page 5673–5684, New York, NY, USA. Association for Computing Machinery.
- Shuyan Zhou, Uri Alon, Sumit Agarwal, and Graham Neubig. 2023. Codebertscore: Evaluating code generation with pretrained models of code.

## A Data Collection and Processing

### A.1 Collection

We crawled existing, public repositories from GitHub before May 2021, resulting in approximately 65 million Python files with a total size of 380 GB. We then removed duplicate files based on the rowKey of each file’s MD5, which resulted in 40 million files (186 GB). We further kept files that meet the following criteria: (a) the file size is under 1MB; (b) the code is Python3 compatible, using Abstract Syntactic Tree (AST) parsing; (c) there are fewer than 100 characters per line on average; (d) and there are fewer than 1,000 characters in any single line. We then removed the duplicated functions from the remaining files. In the end, 120 GB of training data was obtained.

### A.2 Pre-processing

The collected data include both plain Python code and code accompanied by a natural language description after the function signature. The latter can be used to create text-to-code instances which are considered task-specific data for program synthesis, while code-only function snippets can expose models to generic Python programming.

We apply AST parsing<sup>9</sup> on the remaining Python files to extract valid functions and their corresponding docstrings.<sup>10</sup> While we also extract classes, we will only refer to functions as the process is identical. An example is shown below:

```
def gcd(a: int, b: int): -> int
    """
```

<sup>9</sup><https://tree-sitter.github.io/tree-sitter/>

<sup>10</sup>The strings that follow Python docstring conventions: <https://peps.python.org/pep-0257/>

```

Return a greatest common divisor of
the two integers a and b
"""
while b:
    a, b = b, a % b
return a

```

For all extracted functions, we also remove comments. Many comments are simple housekeeping messages, such as dates, contributor’s identity and “Do not delete”.<sup>11</sup> We then replace new lines, indentation and dedentation with `[new_line]`, `[indent]` and `[dedent]`, respectively, to normalise spaces, which effectively reduces the length of input sequences to the model.

Finally, we apply deduplication to all available training data. In the end, we gather 25 million text-to-code pairs and 80.8 million code-only function snippets, respectively.

## B Training of PanGu-Coder

Here we provide additional details with respect to PanGu-Coder’s modality-agnostic pre-training.

### B.1 Tokenization

We use SentencePiece (Kudo and Richardson, 2018) as our primary tokenizer. We train the tokenizer on 10 million samples randomly drawn from the training data, including both docstrings and code. During training, the normalising symbols, `[new_line]`, `[indent]` and `[dedent]`, are passed to SentencePiece to preserve their integrity. The tokenizer’s vocab size is set to 32,000 tokens.

### B.2 Data Formatting

When a docstring exists within a function, we form a text-to-code instance by re-organising the function according to the following template, `[descr] docstring [python] code [eoc]`, where `[descr]` stands for the beginning of a problem description placeholder, `[python]` for start of python code and `[eoc]` for end-of-code, as follows:


```

[descr] Return a greatest common divisor
of two integers a and b [python] def
gcd(a: int, b: int) -> int: [new_line]
[indent] while b: [new_line] [indent] a,
b = b, a % b [new_line] [dedent] return
a [eoc]


```

In case there is no available docstring for a code snippet we simply omit the `[descr]` portion. In

<sup>11</sup>We leave training with comments as future work.



(a) MAPT samples, using all available data.



(b) MRPT samples, using only docstring-code pairs.

Figure 4: Input data formats during modality-agnostic (MAPT) and modality-relative pre-training (MRPT).

our MAPT training, we use all the available training samples including text-to-code pairs and code-only instances. For MRPT training, we consider only text-to-code instances, as shown in Figure 4b.

### B.3 Sample Concatenation

We observe that a large portion of the training samples are shorter than 512 tokens after tokenization. Padding the samples up to the context size can increase training time and waste resources. In this work, we adopt a data concatenation approach, known as packing, to improve training and energy efficiency. We shuffle the data and start appending examples until the maximum sequence length is reached, forming a new sample. If appending an instance exceeds the maximum sequence length, then we continue forming the next sample. Although padding is not completely eliminated with this technique, the amount of padding is greatly reduced and we end up with only 23.2M concatenated training samples almost 18% of the original training samples for Modality-agnostic training of PanGu-Coder. Similarly, for Modality-relative training we have 6.1M examples for PyCodeGPT (26% of original) and 5.7M for PanGu-Coder (24% of original).

Since each concatenated training sample contains several original training samples, we use the `[eoc]` placeholder as an anchor to reset the attention mask and position ids, so that each unique sample only attends to itself.

## C Additional Training Details

We report the size of the models that we trained in terms of model settings in Table 5. We used 8

V100 32GB GPUs for training the 100M models and 16 to train the 350M ones. The total training time of PyCodeGPT was approximately 5 days, 7 days for PanGu-Coder 100M model and 10 days for the 350M model.

## D Comparison with other Code PLMs

We compare our best setting (i.e. PanGu-Coder with partial embedding separation) with existing Code LMs in Table 6, reporting model and data sizes, the contextual window allowed ( $c_{\text{CTX}}$ ), vocabulary size ( $n_{\text{VOCAB}}$ ) and number of tokens the models were trained for.

For CODEX (Chen et al., 2021), the total data size and the number of trained tokens are calculated by considering the initial training of GPT-3 (Brown et al., 2020) for 300 billion tokens on a collection of data equivalent to 570GB of compressed plain text. We include the decoder-only baseline presented by ALPHACODE, and not the encoder-decoder model, as HumanEval results are only reported on the former. The number of train tokens of this baseline is not reported. For CODEGEN (Nijkamp et al., 2023) models, the dataset size of CODEGEN-MULTI was computed by accumulating The Pile (Gao et al., 2021) and BigQuery<sup>12</sup>, while CODEGEN-MONO was additionally trained on BigPython. To calculate the number of training tokens for the CODEGEN models, we assume the batch size reported in the paper corresponds to tokens instead of instances. For INCODER, the vocabulary size was calculated as 55% of GPT-2 vocabulary, based on Fried et al. (2023). For SANTACODER the reported numbers and the number of trained tokens were collected from the model card on the HuggingFace Hub.<sup>13</sup> For the other models, explicit information was provided in the corresponding papers.

For all models,  $\text{pass}@k$  rates are computed with 200 samples, except for ALPHACODE where the reported rates used 1,000 samples and CODEGEN which used 100 samples. Our proposed model achieves the best performance among similar-sized models across all metrics, even having been trained on much less training data.

Regarding the MBPP dataset, our model outperforms all other reported models on  $\text{pass}@1$ , while

<sup>12</sup><https://cloud.google.com/bigquery/public-data/>

<sup>13</sup><https://huggingface.co/bigcode/santacoder>

<sup>14</sup>We did not include even larger scale models, since they would not be directly comparable with this work.

being second best for  $\text{pass}@10$  and  $\text{pass}@100$ . It also outperforms INCODER despite being significantly smaller. Overall, we observe, that our PanGu-Coder performs better for lower  $k$  values. We speculate that the gap in performance for  $\text{pass}@100$  is a result of the small context size (1,024) the model has been trained on, which prevents the model from learning to generate long solutions, or solutions that have to attend over quite long descriptions. Allowing the model to be trained on longer inputs can be beneficial when scaling to many more sample solutions, which enables the generation of exhaustive solutions, e.g. enumerating all possible cases of a for loop.

Finally, the numbers of SantaCoder are using the MultiPL-E benchmark (Cassano et al., 2022) that includes small changes compared to the original datasets (e.g. three problems were removed from the HumanEval test set). As such, aside from the difference in model size, performance is not directly comparable.

## E Additional T-SNE plots

We show additional T-SNE plots in Figure 5 as part of a qualitative analysis of the embedding separation between modalities for the CODE-CLM objective and the 350M PanGu-Coder model.<sup>15</sup>

In the scenario of `def`, we can see that *definition* go further apart from *default* as we move to the code embedding space. For `when`, the neighbours in each space are quite different. In the `docstring` space, some random tokens appear that are not present in the code space. For `except`, *including* and *excluding* in NL are close in the `docstring` space but further apart in the code space. Finally, for `split`, we observe it gets closer with `replace`, `lower` and `strip` in the code space while in the `docstring` space, it can be found bundled together with many other tokens.

<sup>15</sup>We attribute credits to <https://github.com/Phlya/adjustText> for the aligning labels with points.

Model	Model Params	Layers	FFN size	Heads	Context Size ( $n_{\text{CNTX}}$ )	Vocab ( $n_{\text{VOCAB}}$ )
PyCodeGPT (Zan et al., 2022)	111 M	12	768	12	1,024	32,000
PanGu-Coder (Christopoulou et al., 2022)	118 M	12	768	12	1,024	32,012
PanGu-Coder (Christopoulou et al., 2022)	349 M	24	1,024	16	1,024	32,012

Table 5: Models configuration.

MODEL	SIZE	$n_{\text{CNTX}}$	$n_{\text{VOCAB}}$	DATA (GB)	TRAIN TOKENS	HUMAN EVAL (%)			MBPP (%)		
						P@1	P@10	P@100	P@1	P@10	P@100
<b>PANGU-CODER (CCLM-PARTIAL)</b>	300 M	1,024	32.4 K	97	179 B	<b>21.3</b>	<b>26.3</b>	<b>37.9s</b>	<b>18.6</b>	37.4	54.4
CODEX	300 M	4,096	50 K	729	400 B	13.2	20.4	36.3	-	-	-
ALPHACODE	302 M	2,304	8 K	715	-	11.6	18.8	31.8	-	-	-
CODEGEN MULTI	350 M	2,048	50 K	1,595	250 B	6.7	10.6	16.8	7.5	24.2	46.4
CODEGEN MONO	350 M	2,048	50 K	1,812	325 B	12.8	23.1	35.2	14.6	<b>41.5</b>	<b>63.0</b>
PANGU-CODER	317 M	1,024	42 K	147	211 B	17.1	24.1	34.6	16.2	34.4	53.7
ALPHACODE	1.1 B	2,304	8 K	715	-	17.1	28.2	45.3	-	-	-
SANTACODER	1.1 B	2,048	49 K	268	236 B	18.0	29.0	49.0	35.0	58.0	77.0
INCODER	1.3 B	2,048	27.6 K	204	52 B	8.0	-	-	10.9	-	-
INCODER	6.7 B	2,048	27.6 K	204	52 B	15.2	27.8	47.0	19.4	-	-

Table 6: Pass@ $k$  rates on the HumanEval dataset, among various models. Sizes are reported in thousands (K), millions (M), billions (B) and trillions (T). CB refers to CodeBLUE. Models: CodeX (Chen et al., 2021), AlphaCode (Li et al., 2022), CodeGen (Nijkamp et al., 2023), PanGu-Coder (Christopoulou et al., 2022), SantaCoder (Allal et al., 2023), InCoder (Fried et al., 2023)<sup>14</sup>. Best results across the 300M models are **bolded**.

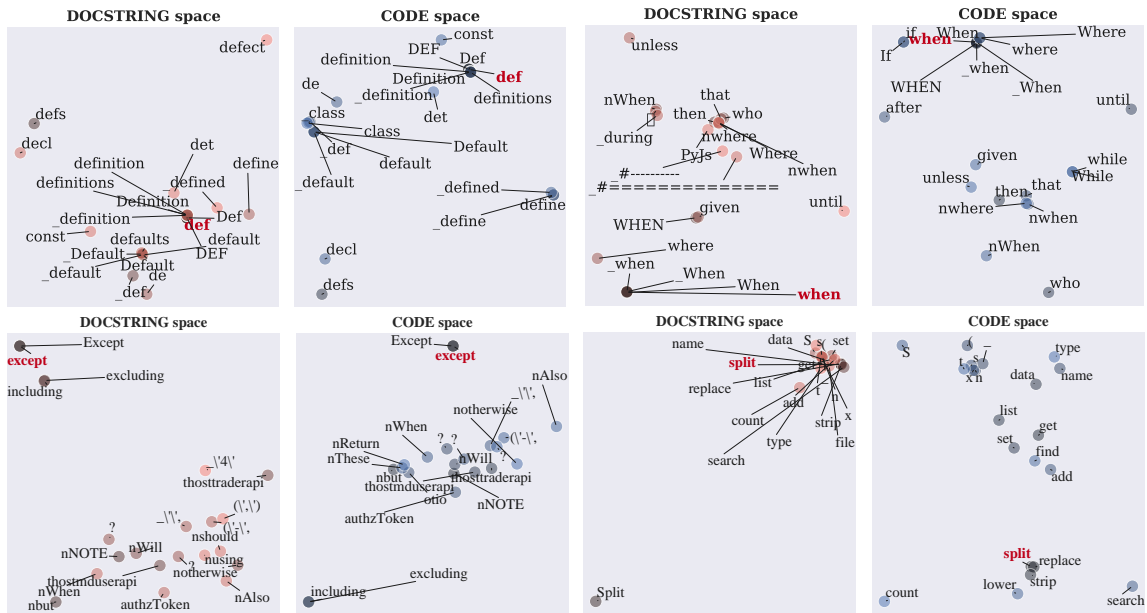


Figure 5: T-SNE plots of the top 20 neighbours of the tokens `def`, `when`, `except` and `split` (top to bottom) on different embedding spaces.