Find Central Dogma Again

Wang Liang

Huazhong University of Science and Technology, 430070, P.R. China

*To whom correspondence should be addressed. E-mail:wangliang.f@gmail.com

[Abstract]In recent years, large language models (LLMs) have achieved state-of-the-art results in various biological sequence analysis tasks, such as sequence classification, structure prediction, and function prediction. Similar to advancements in AI for other scientific fields, deeper research into biological LLMs has begun to focus on using these models to rediscover important existing biological laws or uncover entirely new patterns in biological sequences. This study leverages GPT-like LLMs to utilize language transfer capabilities to rediscover the genetic code rules of the central dogma. In our experimental design, we transformed the central dogma into a binary classification problem of aligning DNA sequences with protein sequences, where positive examples are matching DNA and protein sequences, and negative examples are non-matching pairs. We first trained a GPT-2 model from scratch using a dataset comprising protein sequences, DNA sequences, and sequences from languages such as English and Chinese. Subsequently, we fine-tuned the model using the English similarity judgment dataset from PAWS-X. When tested on a dataset for DNA and protein sequence alignment judgment, the fine-tuned model achieved a classification accuracy of 76%. The study also analyzed factors contributing to this zero-shot capability, including model training stability and types of training data. This research demonstrates that LLMs can, through the transfer of natural language capabilities and solely relying on the analysis of sequences themselves, rediscover the central dogma without prior knowledge of it. This study opens a new door for AI-driven biological research.

1 introduction

Advancements in large language models have revolutionized the field of artificial intelligence, with profound implications for bioinformatics. Specifically, in nucleic acid analysis, specialized models such as DNABert2, HyenaDNA, and ScBert have emerged to address classification and structural prediction challenges related to DNA sequences (1-5). By leveraging the capabilities of large language models, these tools provide novel insights into genetic data. Concurrently, significant progress has been made in protein-related research, where models like ProTrans, ProteinBERT, and ESM2 excel in tasks such as predicting protein structures and annotating their functions (6-12). These innovations underscore the versatility of large language models in tackling complex biological questions, effectively bridging the gap between computational linguistics and molecular biology.

Many studies have focused on improving existing biological problems, more in-depth research on large models aims to leverage their reasoning and transfer capabilities to rediscover or assist in discovering important principles within the field.

Challenge.For research area like mathematics, although there is an abundance of formulaic language, these languages are interpretable, meaning they have clear natural language descriptions or definitions. Therefore, large models can be conveniently utilized through methods like prompt engineering. However, biological sequence languages remain largely unknown, significantly limiting the depth of large model applications in biological research. This makes it challenging to utilize the reasoning capabilities of large models to discover entirely new principles, as has been done in other disciplines.

Research on language transfer capabilities within large models provides key insights into solving this problem.

An intriguing direction within these studies is the exploration of multilingual transfer capabilities in large language models. Research on multilingual transfer has demonstrated how large pre-trained models can effectively apply knowledge learned in one source language to other languages. These studies have validated the effectiveness of cross-linguistic knowledge transfer, including between different natural languages, as well as between programming languages and natural language (13-18). Moreover, Our previous research has shown that the transfer phenomenon from natural language capabilities to DNA sequences also exists, making it possible to directly apply natural language processing techniques to DNA sequence analysis (19).

Research(19) has demonstrated that fine-tuning models on English similar text pairs enables the assessment of similarity between two DNA or two protein sequences, achieving accuracies exceeding 80%. Building upon this foundation, we aim to further utilize large models to investigate the coding rules between DNA and protein sequences, specifically the relationship between these two modalities as described by the central dogma of molecular biology.

The central dogma—including the discovery of the genetic code—is the result of decades of accumulated experimental evidence. Marshall Nirenberg and Heinrich Matthaei, through experiments involving the artificial synthesis of RNA (such as poly-uridine UUU...) and cell extracts, discovered that UUU encodes phenylalanine, marking the first breakthrough in deciphering the genetic code. By 1966, a team of scientists (including Hal Gobind Khorana and others) had chemically synthesized specific sequences to decode all 64 codons, thereby confirming the universality of the genetic code.

In this paper, we aim to rediscover the key rules of the genetic code using only large models combined with biological sequence analysis—without any prior knowledge of the central dogma.

Specifically, we first reformulate the central dogma as a binary classification problem on DNA-protein sequence alignment. We then trained a GPT-2 model, named GPT2-gene-multi, from scratch using DNA sequences, protein sequences, and seven different natural languages. Subsequently, we fine-tuned this model using the English sentence similarity dataset from PAWS-X to obtain the fine-tuned model GPT2-gene-multi-ft. When fine-tuned model tested on datasets for DNA-protein sequence alignment, the classification accuracy can reach over 70%.

2 Materials and methods

2.1 Experiment Design

In current machine learning research, DNA sequences and their corresponding protein-coding sequences are often used as datasets to evaluate models' multimodal sequence processing capabilities. Typically, matching DNA-protein pairs are treated as positive samples, while non-matching pairs serve as negative samples. Supervised learning methods are then employed, where the model is trained on a portion of the dataset and tested on the remaining data.

In contrast, our study adopts an unsupervised approach, aiming to perform zero-shot predictions without any prior knowledge of the central dogma. Drawing inspiration from zero-shot experimental designs in natural language large models, we propose the following three unsupervised methods to predict the central dogma:

1 Designing a Binary Classification Task for DNA-Protein Sequence Alignment: This method mirrors the current DNA-protein coding determination datasets, using matching pairs as positive samples and non-matching pairs as negative samples. However, the dataset is solely utilized for testing purposes and is not incorporated into the training process. This approach leverages the generalization and transfer capabilities of pre-trained or fine-tuned models during testing.

2 Zero-Shot Similarity Calculation: Without any training, this method directly computes the cross-modal embedding similarity between DNA and protein sequences. By analyzing the similarity scores of positive and negative samples, a threshold is established to differentiate between them, thereby obtaining classification metrics such as accuracy. Additionally, this approach can further elucidate the similarity relationships between protein and DNA vocabularies.

3 Prompt-Based Zero-Shot Learning: This technique involves crafting natural language instructions to guide the model's reasoning process. For instance, a prompt might be:

"You are a bioinformatics expert tasked with determining whether a given DNA sequence encodes the corresponding protein sequence.

DNA sequence: ATGGCGCTAAAG

Protein sequence: MART

Please answer 'yes' or 'no':"

The model's generated response ('yes' or 'no') is then mapped to binary classification labels.

Among these methods, the first is the most commonly employed. In the context of causal language models (CLMs), the third method can often be viewed as a variation of the first. The second method, due to differences in tokenization and vector representations, poses challenges in determining similarity and is generally used for secondary validation.

Therefore, our study primarily utilizes the first method to design experiments aimed at discovering the genetic code rules within the central dogma.

2.2 Research Approach

A commonly used method for zero-shot prediction of DNA-protein datasets with large language models is leveraging linguistic transfer ability.

Multilingual transfer ability, which reflects how well models fine-tuned on one source language can be applied to other languages, has been extensively studied in multilingual pre-trained models. In this context, we explore the transfer of natural language abilities to genomic language.

The main process involves first pre-training a large language model, then fine-tuning it on a natural language dataset, and finally evaluating the fine-tuned model on a genomic sequence dataset. The key to achieving successful language transfer lies in identifying a natural language dataset that shares structural similarities with the genomic dataset.

We adopted GPT-2 as the foundation model structure. For pre-training, we used seven languages (including English and Chinese,etc) along with DNA and protein sequences, employing a unified BPE encoding. We then fine-tuned the model on the PAWS-X dataset, a natural language translation and paraphrase similarity judgment dataset, due to its structural similarity to the DNA-protein alignment dataset. Finally, we evaluated the fine-tuned model on the DNA-protein alignment dataset.

The training and test process is illustrated in (Fig. 1).

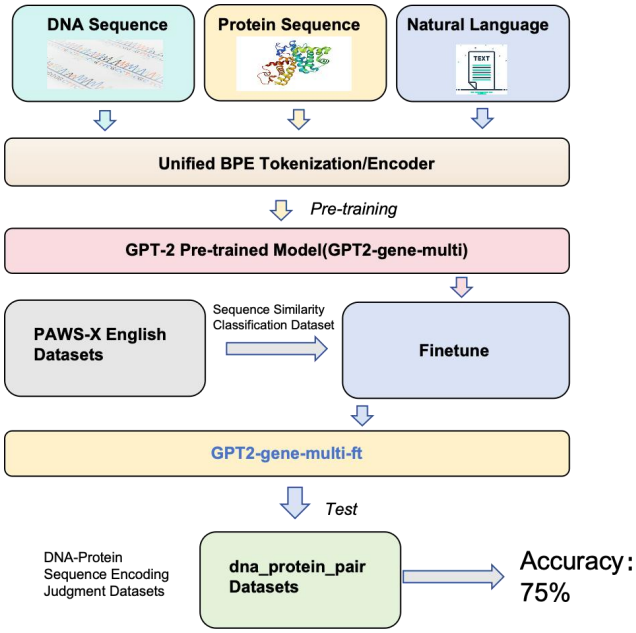

Fig.1. Training and Testing Process of the Research Approach. This model utilizes genomic and protein sequence data, along with encyclopedic text in multiple languages (including English and Chinese), as pre-training corpora. The BPE (Byte Pair Encoding) method is employed for uniform tokenization, treating all text types equivalently without differentiation.The model is trained from scratch based on the GPT-2 small architecture, resulting in the pre-trained model GPT2-gene-multi. It is then fine-tuned on the PAWS-X dataset, producing the GPT2-gene-multi-ft model. Finally, this fine-tuned model is evaluated on a DNA-protein alignment classification dataset, achieving a test accuracy of 75%.

A crucial aspect of achieving effective language transfer is identifying natural language datasets that share structural similarities with DNA-protein encoding. Several types of datasets can serve as potential candidates:

1. Machine Translation – The DNA → protein encoding process is analogous to translating a source language into a target language.
2. Semantic Similarity Texts – Similar to machine translation, this includes sentence pairs with equivalent meanings, such as two semantically similar English sentences or an English sentence paired with its semantically equivalent French translation.
3. Speech-to-Text – The process of encoding DNA into proteins resembles converting speech signals into recognized text.
4. Code Generation – DNA-to-protein conversion follows strict rule-based mappings, similar to the transformation between natural language and programming code.

Given the availability and diversity of datasets, we chose semantic similarity text datasets as candidates for language transfer learning. One of the most representative datasets in this category is PAWS-X, a cross-lingual sentence pair classification dataset designed to evaluate and improve NLP models' ability to recognize semantic equivalence across languages.

Each PAWS-X sample typically includes the following fields:

- sentence1: The first sentence.
- sentence2: The second sentence.
- label: Indicates whether the two sentences are synonymous (0 = not synonymous; 1 = synonymous).
- lang: The language code of the sample.

The sentence pairs in PAWS-X can be in the same language or different languages, making it structurally similar to translation datasets.

Our approach involves fine-tuning the model on the PAWS-X dataset, learning the sentence1 → sentence2 mapping. We then evaluate the fine-tuned model on a DNA → protein alignment dataset. If the model achieves promising results, it suggests that the model has successfully acquired the ability for zero-shot discovery of the genetic code.

2.3 Training datasets

Pretrained datasets

The training data is roughly outlined in the (Table.1):

Table.1 Training data

Data type	source	data size

DNA sequence	multiple model organism genomes	4G
Protein sequence	Swiss-Prot/TrEMBL	4G
Natural Language text	wiki. en,fr,zh,de,es,ja,ko	4G for each language

The pretraining datasets include a diverse range of data sources. The training data consists of DNA sequences from multiple model organism genomes (4GB), protein sequences from Swiss-Prot/TrEMBL (4GB), and natural language text from Wikipedia in multiple languages (English, French, Chinese, German, Spanish, and Japanese), with 4GB allocated for each language. This dataset composition ensures a comprehensive foundation for training the model on both biological and linguistic data.

PAWS-X Dataset

PAWS-X is a sentence-pair similarity assessment task where each data entry consists of two sentences. A typical data entry looks like this:

{

'sentence1': 'In Paris, in October 1560, he secretly met the English ambassador, Nicolas Throckmorton, asking him for a passport to return to England through Scotland.',

'sentence2': 'In October 1560, he secretly met with the English ambassador, Nicolas Throckmorton, in Paris, and asked him for a passport to return to Scotland through England.',

'label': 0

}

The task is to determine whether the two sentences express the same meaning, which is a binary classification problem. To extend this to other languages, the English sentence pairs were translated into six target languages (French, Spanish, German, Chinese, Japanese, and Korean) by professional translators and manually checked to ensure quality and consistency.

We fine-tune the model by selecting same or different language combinations for sentence1 and sentence2, and then evaluate it on the DNA-protein alignment dataset.

DNA-Protein Pair Align Judgment Task

Based on our experimental design and referring to the PAWS-X construction method, we built a DNA-protein pair alignment judgment task based on gene coding rules. If a protein sequence corresponds to its DNA coding sequence, it is considered a positive example; otherwise, it is a negative example. The specific data design was inspired by the dataset used in Lucaone's paper (21). A typical data entry is structured as follows:

```
{
  "sentence1": "ACCAGTGCTCAGGTTAACAAAATAATAAAAGGAA....",
  "sentence2": "MASGRLQLLAFALALSGVSGVLAATLLPNWTVSVD...",
  "label": 0
}
```

The Lucaone's dataset contains approximately 25,600 entries, with a balanced number of positive and negative examples.

In Lucaone's DNA dataset, to increase the testing difficulty, each DNA sequence was extended by adding 100 extra nucleotides at both the 5' and 3' ends, thereby preserving intron sequences within the data. In our version, we have removed these additional 100 base pairs (bp).

Since this dataset is used only for testing, we randomly selected 4,000 samples from it. This dataset is referred to as dna_protein_pair.

In the Lucaone's paper, negative examples were generated by keeping the DNA sequence unchanged and modifying the corresponding protein sequence through data perturbation methods, such as:

- Amino Acid Substitution (translation errors)
- Shuffling Protein Sequence

To further reduce the test difficulty, we constructed negative examples using random protein sequences. Specifically, we kept the DNA sequence unchanged, used the correct protein-coding sequence as a reference, and then randomly selected a protein sequence of similar length from a protein database, ensuring its sequence similarity to the reference was below a threshold (e.g., 0.45). This dataset is referred to as dna_protein_pair_rand.

The fine-tuning and testing datasets are detailed in (Table.2).

Table.2 Finetune and test data

dataset name	Data type	source
PAWS-X-EN	Natural Language Pair	PAWS-X
dna_protein_pair	DNA-protein pair	Swiss-Prot/TrEMBL
dna_protein_pair_rand	DNA-protein pair	Swiss-Prot/TrEMBL

2.4 Training Strategy

Model Structure

The base model uses the exact same model architecture as GPT-2, with the primary difference being the training corpus. For the research validation, the base model is structured as GPT-2 Small, which can be trained on a single 4090 GPU. For classification fine-tuning, a typical classification head is used, outputting two classes. For prompt instruction fine-tuning, the same head as the petrained model is used, and the training mode is also identical to that of the pre-training. The model structure is shown in (Fig.2).

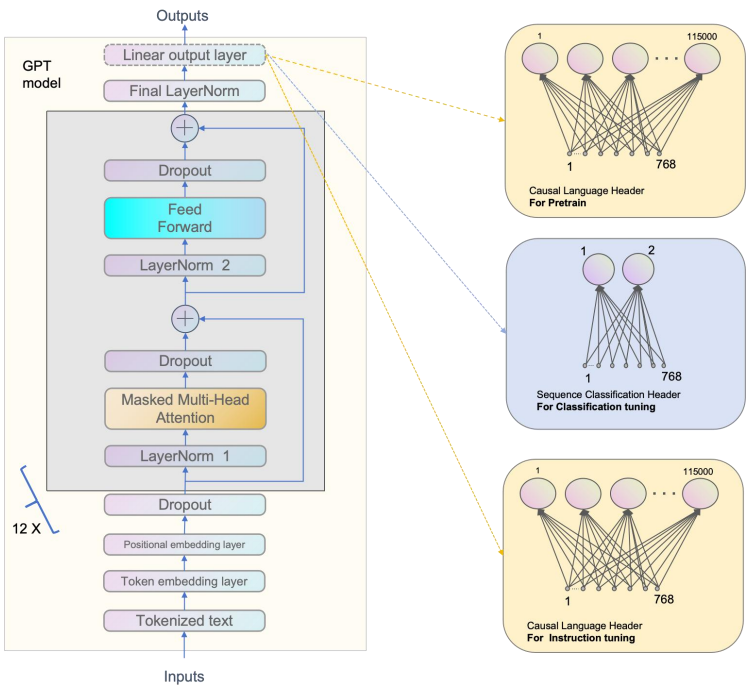

Fig.2. Model structure. The model architecture is identical to that of GPT-2 small. Pre-training and instruction tuning use the same causal language header,the layer mapped 768 hidden units to 115,000 units (the number of tokens in the vocabulary). Classification uses sequence classification header, The layer maps from 768 hidden units to only 2 units, where the 2 units represent the two classes id.

Tokenization

Research on language capability transfer indicates that using a unified tokenizer is a necessary condition for transferring capabilities. Both the Byte Pair Encoding (BPE) tokenizer used in GPT2 models and the WordPiece tokenizer used in BERT models can process DNA and protein sequences. These tokenizers typically split sequences into tokens of 1 to 6 characters, which is a common method for feature extraction in biological sequences. Therefore, the choice of tokenizer does not affect our validation of language capability transfer. Here, we use the default BPE tokenizer of the GPT-2 model.

Pretrain

According to the typical design of GPT-2, it accepts sequences with a maximum length of 1024 as input. We used the same architecture as the GPT-2 Small model, which consists of 12 Transformer layers, each with 768 hidden units and 12 attention heads. This small model has approximately 117 million parameters.We trained the GPT-2 model using mixed-precision floating-point arithmetic on a machine equipped with a single Nvidia 4090 GPU. We employed a dynamic learning rate schedule, and the model was trained for a total of 3 to 5 epochs.

Finetune

We used the GPT2-small model as a baseline for fine-tuning on the PAWS-X English dataset. Subsequently, we evaluated the model on DNA-protein pair align judgment datasets.

Input and Output Format:

we concatenated the two sequences directly as the input string, with the label serving as the classification ID:

Input: sequence1 + sequence2

Output: label

Note: For GPT-based models, no separator is needed as the model itself determines the boundaries between the two sequences.

Fine-tuning Procedure:

For Classification Fine-tuning,The model's header is set to the commonly used softmax classification form. We utilized the same training tricks across all the applications, where the learning rate was first linear warmed-up to the peak value and then linear decayed to near 0.

Due to the smaller scale of model parameters and the datasets size being in the tens of thousands of samples, we used the full-parameter fine-tuning method. Fine-tuning can be completed within about an hour on a single 4090 GPU.

Model Evaluation

We fine-tuned the model using the PAWS-X dataset and tested it on the DNA-protein alignment classification dataset. The model's accuracy is used to measure its performance. Since this is a binary classification task, the meaning of classification IDs 0 and 1 may differ during transfer. Therefore, if the test accuracy is below 0.5, we can simply swap label IDs 0 and 1. For some experiments, we further evaluate the model using a Confusion Matrix.

3 Experimental Results

3.1 Basic results

As a binary classification task, we primarily used accuracy as the evaluation metric.

We utilized three pretrained models as the base models. Among them, GPT2 is the standard model released by OpenAI, primarily trained on English data. The gpt2-gene-eng model is a GPT2 model that we trained from scratch using a dataset consisting of English text, DNA sequences, and protein sequences. The gpt2-gene-multi model is also a GPT2 model trained from scratch, but its training data includes DNA and protein sequences, as well as text data in seven natural languages, including English and Chinese.

These base models were fine-tuned using the English paraphrase classification corpus from PAWS-X and then tested on the DNA-protein alignment classification dataset. In alignment dataset, we inverted the original labels, meaning that 0 represents alignment, and 1 represents non-alignment.

As a comparison, we also evaluated the fine-tuned models on the PAWS-X English test set, as well as the French and Chinese test sets. The specific test results are shown in (Table.3).

Table.3 Accuracy of finetune model in different datasets

base model	pretrain	finetune	test-en	test-fr	test-zh	test dna_protein_pair	test dna_protein_pair_rand
gpt2	en	en	0.90	0.77	0.63	0.57	0.55
gpt2-gene-eng	en, dna, protein	en	0.89	0.74	0.65	0.59	0.68
gpt2-gene-mulit	dna, protein, Multi-lan	en	0.90	0.82	0.74	0.73	0.76

The comparison of base models highlights the impact of pretraining data on cross-lingual and cross-modal transfer learning. The standard GPT-2 model, pretrained only on English text, achieved 0.90 accuracy on PAWS-X English, but performed poorly on French (0.77) and Chinese (0.63). Its accuracy on DNA-protein pair classification was 0.57, with a slight drop to 0.55 on randomly paired sequences. The GPT2-Gene-Eng model, pretrained on English, DNA sequences, and protein sequences, slightly underperformed in PAWS-X English (0.89) but showed minor improvements in DNA-protein classification (0.59) and a notable boost on random protein pairs (0.68). The GPT2-Gene-Multimodel, pretrained on DNA, protein sequences, and seven natural languages, maintained high PAWS-X English accuracy (0.90), significantly improved French (0.82) and Chinese (0.74) performance, and demonstrated the best performance in DNA-protein classification (>0.7 for both datasets), indicating strong adaptability across both natural and biological languages.

The best transfer performance was achieved by the GPT2-Gene-Multi model on the dna_protein_pair_rand dataset. The Confusion Matrix for this test result is shown in the (Fig.3).

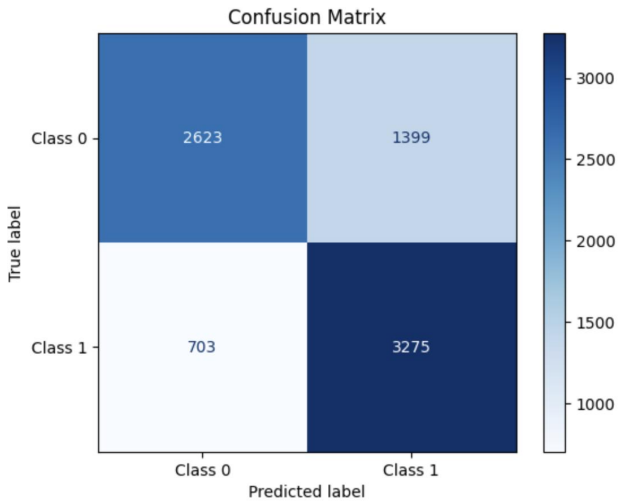

Fig.3 Confusion Matrix of DNA-Protein Alignment Classification. After fine-tuning the gpt2-gene-mulit model on the PAWS-X English similarity dataset, testing on the dna_protein_pair_rand dataset achieved a classification accuracy of approximately 76%. In this evaluation, class 0 represents aligned data and class 1 represents non-aligned data. While the positive and negative samples are relatively balanced, a higher proportion of class 0 samples were misclassified as class 1. Overall, the classification performance is essentially normal.

Based on the above data analysis, we can draw the following conclusions:

- Cross-lingual Pretraining Improves Generalization:The GPT2-Gene-Multi model, which included multilingual data, outperformed both GPT-2 and GPT2-Gene-Eng in non-English PAWS-X tasks and DNA-protein classification.
- Adding DNA and Protein Data Helps with Biological Tasks:The GPT2-Gene-Eng and GPT2-Gene-Multi models both performed better than the standard GPT-2 on DNA-protein alignment tasks.
- Multilingual Pretraining Offers the Best Performance Overall:The GPT2-Gene-Multi model achieved the highest accuracy across all test datasets, confirming the effectiveness of incorporating both biological and multilingual data.

3.3 Training Stability

In the training of large models such as GPT-2, the choice of seed value can introduce variability in training outcomes (18). Our experiments demonstrated that when transferring English language capabilities to other natural languages, the influence of different seed settings is relatively minor. In contrast, when transferring these capabilities to DNA-Protein Sequences, the impact of seed settings is considerably more significant.

To assess the impact of seed settings, we randomly assigned different seed values and repeated the experiment multiple times, collecting statistics on the fine-tuning and testing results.

As a comparison, we trained two base pre-trained models, GPT2-gene-multi-v1 and GPT2-gene-multi-v2, both trained on DNA and protein sequences as well as seven natural languages, with the only difference being the random seed settings. Based on these two pre-trained models, we conducted multiple fine-tuning experiments and evaluated the accuracy results. The statistical outcomes are illustrated in the (Table.4):

Table.4 Accuracy of differet pre-trained model

model	finetune	test en	test fr	test zh	test dna_protein_pair	test dna_protein_pair_rand
gpt2-gene-mulit-v1	en	0.90	0.82	0.73	0.73	0.71
gpt2-gene-mulit-v2	en	0.90	0.82	0.74	0.67	0.76

The two pre-trained models, GPT2-gene-multi-v1 and GPT2-gene-multi-v2, were fine-tuned on the English similarity judgment dataset of PAWS-X and then tested on the PAWS-X datasets in English, French, and Chinese, as well as on the dna_protein_pair and dna_protein_pair_rand datasets.

It can be observed that models pretrained with different seeds show minimal variation (within 1%) in test results for natural languages such as English and French. However, when tested on DNA-protein sequences, the accuracy difference exceeds 5%.

During fine-tuning, the choice of seed also has a significant impact on performance. We conducted multiple fine-tuning experiments using different seeds and tested the models on the DNA-protein alignment dataset. The accuracy distribution is shown in the (Fig.4). For comparison, the test accuracy on French and Chinese datasets is also included.

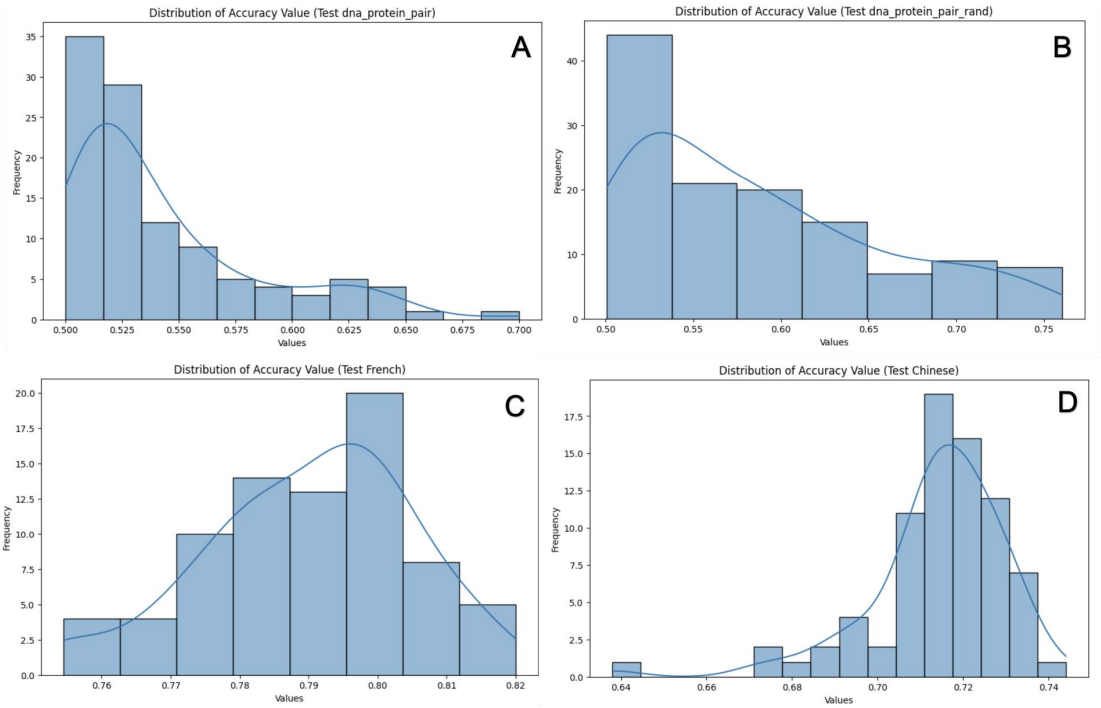

Fig.4 Accuracy Distribution Across Different Datasets.Figure A and Figure B show the accuracy distribution on the DNA-protein alignment dataset:A corresponds to the dna_protein_pair dataset. B corresponds to the dna_protein_pair_rand dataset.

If we set 70% accuracy as a threshold for good performance:In Figure A, only 1% of the results exceed 70%. In Figure B, only 9% of the results exceed 70%. This implies that for the dna_protein_pair dataset, more than 100 fine-tuning trials are needed to obtain a well-performing model. Most results are only slightly better than random guessing, and the difference between the highest and lowest accuracy values exceeds 0.2.

Figure C and Figure D show the accuracy distribution on French and Chinese datasets, respectively. Except for a single outlier (0.64) in Figure D, the accuracy variation between the highest and lowest values remains within 0.08.

The data in the figure illustrates that when transferring language capabilities from English to gene language (DNA and protein sequences), the solution space of the optimization problem is highly rugged. Only a few random seeds enable the model to find local optima that perform well on both English data and gene sequences. The detailed analysis is as follows:

1. Non-convexity of the Solution Space and Scarcity of Local Optima

During fine-tuning, we are effectively searching for the optimal parameters in a high-dimensional non-convex loss space. Since the source language (English) and the target language (gene language) differ significantly in structure, symbol systems, and syntactic rules, most local minima are only effective for English data and do not generalize well to gene sequences. Only when a random seed happens to initialize the model and guide the optimization process towards those rare and "cross-domain-friendly" local minima, effective language transfer can be achieved.

2. Challenges in Aligning Cross-Domain Representations

As a natural language, English has linguistic structures, vocabulary, and syntactic rules that are completely different from the encoding rules of DNA or protein sequences. Successful transfer requires the model to restructure or reorganize some of its internal representations during fine-tuning so that it can both retain the linguistic capabilities learned from pretraining on English and capture the encoding patterns of gene sequences. The fact that only a few random seeds achieve this balance suggests that in most cases, the model's internal representation tends to solidify around source language features rather than naturally aligning with gene language.

3. High Variance and Instability

The observation that only a few seeds lead to good performance also highlights the high variance in cross-domain fine-tuning. Due to random initialization, data ordering, and other stochastic factors, the optimization trajectory can differ significantly across runs. This instability suggests that the current approach to cross-language (especially cross-domain) transfer is highly sensitive and may require additional strategies—such as model ensembling, weight averaging, or more robust optimization techniques—to reduce the impact of random seeds and improve the stability of the transfer.

4. Potential Approaches to Address Randomness

- Multiple Trials and Ensembling: Since only a small fraction of random seeds can locate ideal cross-domain minima, multiple fine-tuning experiments can be conducted, followed by model ensembling or weight averaging, to reduce the influence of individual unstable solutions.
- Regularization and Domain Adaptation: Introducing regularization techniques (e.g., dropout, gradient clipping) or domain adaptation methods can help the model smoothly bridge the distribution gap between the source (English) and target (gene language) domains.
- Mixed Training Strategy: During fine-tuning, training on a mixture of English and gene sequence data may help the model align representations across different languages, reducing its dependency on random seeds. However, this approach would shift the task from zero-shot transfer to few-shot learning.

Our approach of conducting multiple fine-tuning runs with different random seeds partially mitigates the challenges of language transfer, but it comes with a significant computational cost. Future research can leverage the insights from this analysis to reduce computational overhead and achieve faster and more stable convergence towards optimal solutions.

3.3 Best finetune Language

The PAWS-X test dataset includes not only English similarity text pairs but also similar text pairs in other languages such as French, German, and Chinese. We conducted further testing across all these languages.

The pretrained model used in our experiments was GPT2-Gene-Multi-v2, which was trained on protein sequences, DNA sequences, and all natural language data from PAWS-X. Specifically, we fine-tuned models using different natural language similarity sequences and then tested them on the DNA-protein alignment dataset. As a comparison, we also evaluated the fine-tuned models on English, French, and Chinese test sets.

Since fine-tuning performance is highly sensitive to random seed selection, we conducted 100 training runs for each dataset, with random seeds ranging from 1 to 100, and then selected the best-performing model for comparison (Table.5).

Table.5. Effect of Fine-Tuning on Different Language Datasets

finetune	Test-en	fr	zh	dna_protein_pair	dna_protein_pair_rand
en	0.90	0.82	0.74	0.56	0.67
Fr	0.85	0.85	0.74	0.54	0.59
De	0.82	0.80	0.73	0.53	0.57
Es	0.85	0.83	0.72	0.55	0.61
Zh	0.72	0.74	0.77	0.51	0.61
Ja	0.75	0.73	0.76	0.57	0.55
Ko	0.70	0.71	0.73	0.56	0.59

From the table above, we can see that the effect of transferring capabilities from different natural languages to genetic language does not vary significantly. The main reason is that in the transfer process from natural language to genetic language, the impact of random seeds is greater than the choice of source language. Due to the relatively small number of experiments, the best accuracy result (0.67) is significantly lower than the result obtained from extensive experiments (0.75). If a more stable and robust fine-tuning method can be found, the comparison of source language transfer effects would be clearer.

4 Discussion

4.1 Necessary Conditions for Transfer Ability

These experimental results indicate that there are several key necessary conditions for successfully transferring natural language capabilities to genetic language:

1. Multimodal Data Training

During the pretraining phase, using a diverse set of training data—including multiple natural languages (not just English) as well as DNA and protein sequences—enables the model to learn a more comprehensive representation across different languages and symbolic systems. This

approach allows the model not only to grasp English grammar and expression patterns but also to capture structural and sequential patterns in biological sequences. Due to the diverse data sources, the model develops a richer internal representation, making it more capable of adapting to genetic language compared to a model pretrained solely on English. This is similar to how a multilingual person can more easily find commonalities when encountering an unfamiliar field compared to a monolingual individual.

2. Unified Tokenization and Encoding Methods (e.g., BPE)

Using a unified tokenization (or encoding) method ensures consistency when converting different types of data (natural language and genetic sequences) into model inputs. For example, Byte Pair Encoding (BPE) can segment both English text and DNA/protein sequences into subwords or sub-sequences, forming a shared vocabulary or encoding space. This uniformity reduces discrepancies caused by inconsistent encoding, allowing the model to transfer knowledge more naturally and improving performance in cross-domain tasks. Many studies have also validated that unified encoding plays a crucial role in cross-lingual or cross-domain transfer learning.

3. Fine-tuning on Natural Language Datasets with Structural Similarity to Genetic Sequence Tasks

The structure of the dataset used for fine-tuning is critical for transfer effectiveness. Genetic sequence tasks often follow fixed patterns, such as DNA-protein alignment problems. Selecting a natural language dataset with a similar structure for fine-tuning allows the model to learn corresponding relationships and patterns, making it easier to transfer this "mapping" ability to genetic tasks. In other words, if the fine-tuning dataset shares formal and logical similarities with the target task, the model can more effectively find an appropriate solution when adapting to the new task.

These three conditions complement each other:

- Multimodal pretraining equips the model with a broad knowledge base.
- A unified tokenization method establishes a consistent internal representation space.
- Fine-tuning on structurally similar datasets helps the model effectively apply its learned general capabilities to domain-specific tasks (such as genetic language).

Together, these factors enable a more robust and efficient transfer of natural language abilities to genetic language processing.

4.2 The Randomness of Transfer Ability

From our experiments, we observed that randomness is one of the main characteristics of transferring natural language capabilities to genetic language. This phenomenon can be compared to an "accidental alignment of thinking patterns."

Imagine a model pretrained on a large amount of English data—similar to how a native English speaker is accustomed to English logic and expression patterns. When processing information, the model internally builds patterns that align with English language structures and thought processes.

However, while DNA and protein sequences are also "sequences," they have their own unique "grammar" and encoding rules, which are completely different from English language rules. In most cases, the model's "English thinking style" does not effectively apply to decoding biological sequences, leading to poor performance on the target task (DNA/protein alignment classification).

Occasionally, though, due to differences in random seeds, the initialization and data order during training might lead the model to "accidentally" discover a special local minimum—a point that is not only effective for English data but also happens to capture certain key encoding patterns in biological sequences. This is analogous to an English thinker suddenly realizing that their problem-solving habit also happens to work for deciphering a certain type of code, resulting in unexpectedly good performance.

Implications of This Phenomenon:

1. Cross-domain transfer is highly sensitive: Transferring a model between English and biological sequences requires finding those rare solutions that can accommodate both language structures. However, such solutions are extremely scarce in the optimization space and only appear under a few specific random seeds.
2. The alignment of thinking patterns is accidental: The model's English-based reasoning only produces good transfer results when it "happens" to overlap with the patterns of DNA or protein sequences under certain conditions. This suggests that cross-domain transfer may require additional regularization or domain adaptation techniques to stabilize such accidental alignments, making the results less dependent on random initialization.

Essentially, most of the time, an English thinker would struggle to decode biological sequences effectively. However, in rare cases, their way of thinking might happen to align with certain patterns in biological sequences, enabling them to "crack the code" successfully. While this occurrence is rare, it also demonstrates that the model possesses potential cross-domain transferability, though it remains fragile and highly random.

4.3 New Research Methods

After verifying that the model possesses the ability to transfer from natural language to genetic language, it indicates that the large model has developed a certain level of "understanding" of genetic sequences. This capability can provide novel tools and methods for biological research, expanding entirely new research and development approaches (Fig.5).

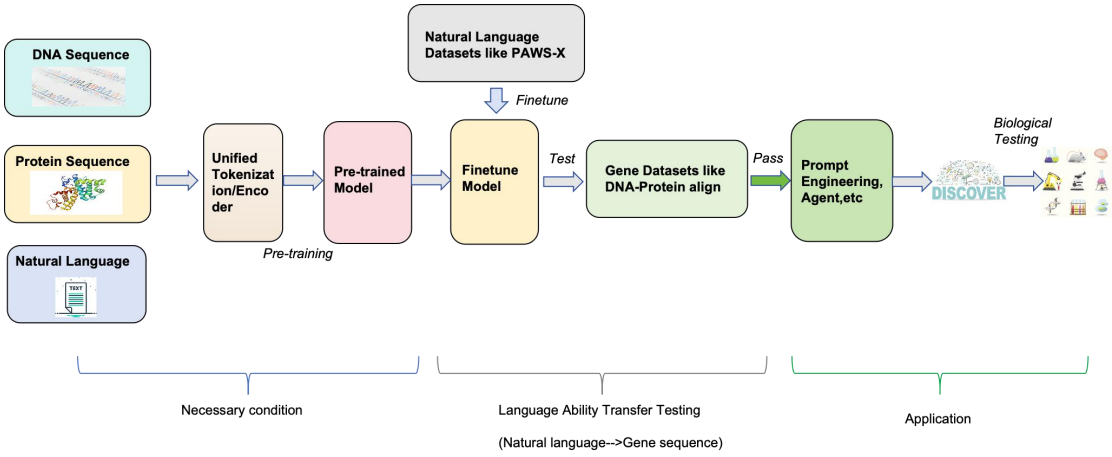

Fig.5. Expanding New Research Methods Based on the Transfer of Capabilities from Natural Language to Genetic Language.

The first step is establishing the fundamental prerequisites: pretraining a unified, multimodal large model using both natural language and biological sequences. Then, in the capability transfer validation phase, the model is fine-tuned with natural language datasets and tested on genetic language datasets. If the validation is successful, it indicates that the model has essentially mastered genetic language and can proceed to the application stage. In this stage, techniques such as prompt engineering can be used to discover new biological sequence patterns, which can then be validated through biological experiments.

Potential New Research Directions Enabled by Large Models in Understanding Natural and Genetic Languages：

1. Combining Prompt Engineering and Agents for Biological Research
   - Utilizing Prompt Engineering and Agents to enable large models to actively explore hidden patterns in genetic sequences. For example:
     - Designing structured prompts for DNA or protein sequences to predict potential functional regions, mutation effects, etc.
     - Training agents to autonomously search and optimize sequences to discover new functional genes or protein structures.
   - This approach can accelerate the discovery of new functional genes and optimize gene editing targets, such as those for CRISPR.
2. Expanding AI Applications in Biology Through Cross-Modal Learning
   - Integrating biological text data (e.g., research papers, patents) with experimental data to enhance the model's reasoning capabilities. For example:
     - Enabling large models to automatically read biomedical literature and generate new biological hypotheses by combining them with genetic sequence data.
     - Utilizing cross-modal AI to integrate imaging (e.g., protein structures), text, and sequence data for more precise disease prediction and personalized medicine analysis.

5 Conclusion

This study leverages the language transfer capabilities of large language models to achieve zero-shot prediction, uncovering the DNA-protein alignment coding rules. In other words, by combining large language models with data analysis, we have rediscovered the genetic code rules within the central dogma.

The transfer of capabilities between natural languages can significantly enhance the ability of large models to serve low-resource languages. Similarly, transferring natural language capabilities to genetic language provides a novel perspective for studying biological languages, which remain largely unexplored. Through fine-tuning with datasets similar to PAWS-X and validation on DNA-protein alignment datasets, we have confirmed that large models can grasp the fundamental principles of genetic language.

On this basis, commonly used methods in large language models, such as prompt engineering and knowledge bases, can be applied. For example, through prompt engineering and agents, biological patterns can be discovered automatically. By combining experimental validation, research credibility can be enhanced. AI-powered data analysis pipelines can accelerate genomics and proteomics research. Cross-modal learning can further drive advancements in precision medicine and drug discovery. This direction not only contributes to fundamental biological research but also significantly enhances AI applications in life sciences, pharmaceuticals, and synthetic biology.

The code and data of this article are all open source (25).

Reference

1. Z. Zhou, et al., Dnabert-2: Efficient foundation model and benchmark for multi-species genome, arXiv preprint arXiv:2306.15006 (2023).
2. E. Nguyen, et al., Hyenadna: Long-range genomic sequence modeling at single nucleotide resolution, Advances in neural information processing systems 36 (2024).
3. F. Yang, et al., scbert as a large-scale pretrained deep language model for cell type annotation of single-cell rna-seq data, Nature Machine Intelligence 4, 852–866 (2022)
4. A. Elnaggar, et al., Prottrans: Toward understanding the language of life through selfsupervised learning, IEEE transactions on pattern analysis and machine intelligence 44, 7112–7127 (2021).
5. N. Brandes, D. Ofer, Y. Peleg, N. Rappoport, M. Linial, Proteinbert: a universal deeplearning model of protein sequence and function, Bioinformatics 38, 2102–2110 (2022).
6. Z. Lin, et al., Evolutionary-scale prediction of atomic-level protein structure with a language model, Science 379, 1123–1130 (2023)
7. Ji, Y., Zhou, Z., Liu, H. & Davuluri, R. V. DNABERT: Pre-trained bidirectional encoder representations from transformers model for DNA-language in genome. *Bioinformatics* **37**, 2112–2120 (2021). doi:doi.org/10.1093/bioinformatics/btab083

8. Veniamin Fishman, et al.GENA-LM: A Family of Open-Source Foundational DNA Language Models for Long Sequences.BioRxiv 2023.06.12.544594. doi: https://doi.org/10.1101/2023.06.12.544594

9. Daoan Zhang, et al.DNAGPT: A Generalized Pretrained Tool for Multiple DNA Sequence Analysis Tasks. BioRxiv 2023.07.11.548628; doi: https://doi.org/10.1101/2023.07.11.548628

10. Wang, B. *et al.* Pre-trained Language Models in Biomedical Domain: A Systematic Survey. *ACM Comput. Surv.***56**, 1–52 (2023).doi:doi.org/10.1145/3611651

11. Chen, Z., Wei, L. & Gao, G. Foundation models for bioinformatics. *Quant. Biol.* (2024) doi:10.1002/qub2.69.

12. Benegas, G., Ye, C., Albors, C., Li, J. C. & Song, Y. S. Genomic Language Models: Opportunities and Challenges. *arXiv* (2024) doi:10.48550/arxiv.2407.11435.

13. Fei Yuan,etc.How Vocabulary Sharing Facilitates Multilingualism in LLaMA?*arXiv* (2024) doi:https://doi.org/10.48550/arXiv.2306.06688

14. Alexis Conneau, Shijie Wu, Haoran Li, Luke Zettlemoyer, and Veselin Stoyanov. 2020. Emerging Cross-lingual Structure in Pretrained Language Models. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 6022–6034, Online. Association for Computational Linguistics.

15. Jiacheng Ye, Xijia Tao, Lingpeng Kong.Language Versatilists vs. Specialists: An Empirical Revisiting on Multilingual Transfer Ability. *arXiv* (2024) doi:10.48550/arxiv.2407.11435.

16. Hetong Wang, Pasquale Minervini, and Edoardo Ponti. 2024. Probing the Emergence of Cross-lingual Alignment during LLM Training. In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 12159–12173, Bangkok, Thailand. Association for Computational Linguistics.

17. Junjie Hu, etc.XTREME: A Massively Multilingual Multi-task Benchmark for Evaluating Cross-lingual Generalization. arXiv(2020) https://arxiv.org/abs/2003.11080

18. Gaschi F, Cerda P, Rastin P, Toussaint Y. Exploring the relationship between alignment and cross-lingual transfer in multilingual transformers. Findings of the Association for Computational Linguistics: ACL 2023.

19. Liang W. Can linguists better understand DNA?. arXiv preprint arXiv:2412.07678. 2024 Dec 10.

20. Thomas Hayes, et al.Simulating 500 million years of evolution with a language model. BioRxiv 2024.07.01.600583; doi: https://doi.org/10.1101/2024.07.01.600583

21. Yong He, et al. LucaOne: Generalized Biological Foundation Model with Unified Nucleic Acid and Protein Language. BioRxiv 2024.05.10.592927; doi: https://doi.org/10.1101/2024.05.10.592927

22. Zhang, K., Zhou, R., Adhikarla, E. et al. A generalist vision–language foundation model for diverse biomedical tasks. Nat Med (2024). https://doi.org/10.1038/s41591-024-03185-2

23. Minghao Xu, Xinyu Yuan, Santiago Miret, and Jian Tang. 2023. ProtST: multi-modality learning of protein sequences and biomedical texts. In Proceedings of the 40th International Conference on Machine Learning (ICML'23), Vol. 202. JMLR.org, Article 1615, 38749–38767.

24. Zeyuan Wang, etc. 2024. InstructProtein: Aligning Human and Protein Language via Knowledge Instruction. In Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), pages 1114–1136, Bangkok, Thailand. Association for Computational Linguistics.

25. paper huggingface, https://huggingface.co/dnagpt/central_dogma

Data availability section

Sequence data that support the findings of this study have been deposited in huggingface.The fine-tuning data is located at https://huggingface.co/datasets/dnagpt/gene_lan_transfer, and the pre-training data and code is located at https://huggingface.co/dnagpt/central_dogma

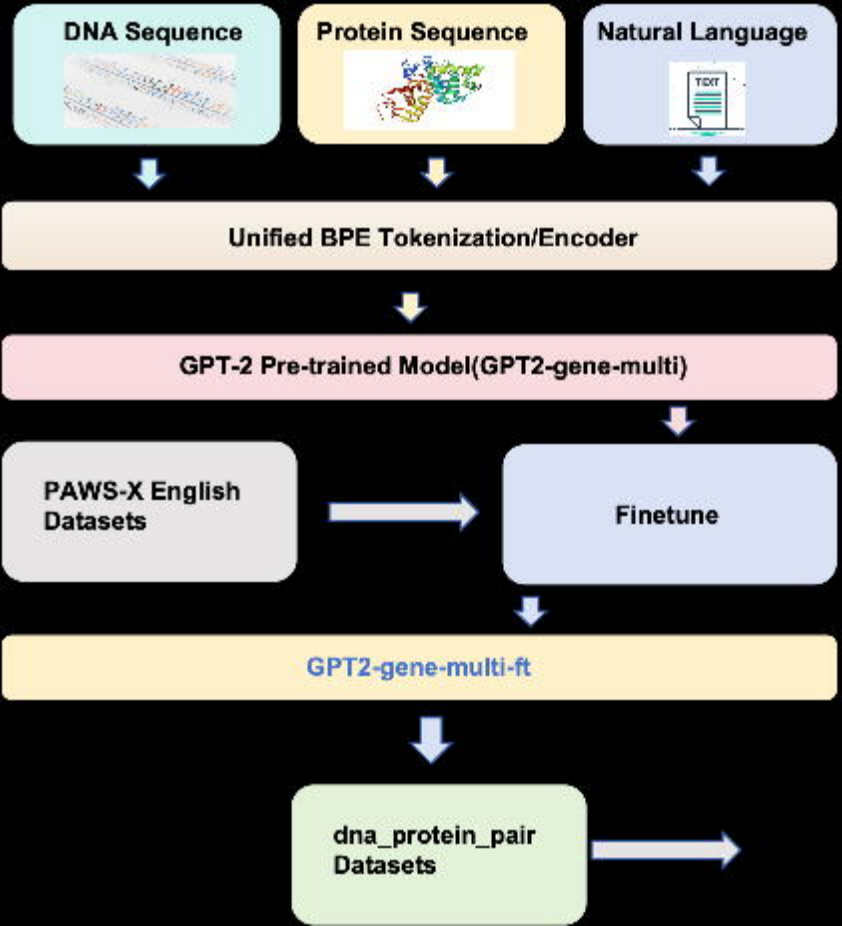
DNA Sequence
Protein Sequence
Natural Language
TEXT
Unified BPE Tokenization/Encoder
GPT-2 Pre-trained Model(GPT2-gene-multi)
PAWS-X English Datasets
Finetune
GPT2-gene-multi-ft
dna_protein_pair Datasets