

# Unsupervised Generative Feature Transformation via Graph Contrastive Pre-training and Multi-objective Fine-tuning

Wangyang Ying  
Arizona State University  
Tempe, Arizona, USA  
wangyang.ying@asu.edu

Dongjie Wang  
University of Kansas  
Lawrence, Kansas, USA  
wangdongjie@ku.edu

Xuanming Hu  
Arizona State University  
Tempe, Arizona, USA  
solomonhxm@asu.edu

Yuanchun Zhou  
Computer Network Information  
Center, Chinese Academy of Sciences  
Beijing, China  
zyc@cnic.cn

Charu C. Aggarwal  
IBM T. J. Watson Research Center  
Yorktown Heights, New York, USA  
charu@us.ibm.com

Yanjie Fu<sup>†</sup>  
Arizona State University  
Tempe, Arizona, USA  
yanjie.fu@asu.edu

## ABSTRACT


Feature transformation is to derive a new feature set from original features to augment the AI power of data. In many science domains such as material performance screening, while feature transformation can model material formula interactions and compositions and discover performance drivers, supervised labels are collected from expensive and lengthy experiments. This issue motivates an Unsupervised Feature Transformation Learning (UFTL) problem. Prior literature, such as manual transformation, supervised feedback guided search, and PCA, either relies on domain knowledge or expensive supervised feedback, or suffers from large search space, or overlooks non-linear feature-feature interactions. UFTL imposes a major challenge on existing methods: how to design a new unsupervised paradigm that captures complex feature interactions and avoids large search space? To fill this gap, we connect graph, contrastive, and generative learning to develop a measurement-pretrain-finetune paradigm for UFTL. For unsupervised feature set utility measurement, we propose a feature value consistency preservation perspective and develop a mean discounted cumulative gain like unsupervised metric to evaluate feature set utility. For unsupervised feature set representation pretraining, we regard a feature set as a feature-feature interaction graph, and develop an unsupervised graph contrastive learning encoder to embed feature sets into vectors. For generative transformation finetuning, we regard a feature set as a feature cross sequence and feature transformation as sequential generation. We develop a deep generative feature transformation model that coordinates the pretrained feature set encoder and the gradient information extracted from a feature set utility evaluator to optimize a transformed feature generator. Finally, we conduct extensive experiments to demonstrate the effectiveness, efficiency, traceability, and explicitness of our framework. Our code and data are available at <https://shorturl.at/pKQU5>.

## KEYWORDS

unsupervised feature transformation, representation learning

## 1 INTRODUCTION

Feature transformation (FT) is to derive new features from original features to reconstruct a transformed feature space (e.g.  $[f_1, f_2] \rightarrow$



**Figure 1: Unsupervised Feature Transformation Learning.**

$[\frac{f_1}{f_2}, f_1 - f_2, \frac{f_1+f_2}{f_1}]$ ). As an essential task of data-centric AI, FT is practical and effective in industrial deployments, because it can augment the AI power of data (e.g., structural, predictive, interaction, and expression levels) to achieve better predictive performance even with simple models. In many practices, FT is conducted either by human experts, or by machine-assisted search guided through downstream task feedback. However, in certain domains, such as material performance screening, although FT can model material formula interactions and compositions to discover material performance drivers, supervised labels are collected from lengthy and expensive experiments. This issue motivates a new learning problem: Unsupervised Feature Transformation Learning (UFTL). Solving UFTL can reduce data costs, model complex feature-feature interactions, learn from non-labeled data, and improve model generalization capability.

Existing systems only partially solve UFTL: 1) manual transformations are effective and explicit, but not generalizable and incomplete, as they heavily rely on domain and empirical experiences. 2) supervised transformations include exhaustive-expansion-reduction approaches [11, 12, 14, 16, 19] and iterative-feedback-improvement approaches [15, 25, 28] that are enabled by reinforcement, genetic algorithms, or evolutionary algorithms. Such methods can search optimal feature sets in a vast discrete space, but they face exponentially growing combination possibilities, are computationally costly, hard to converge, and unstable. Moreover, they rely on labeled data, which limits their applicability. 3) unsupervised transformations, such as Principal Components Analysis (PCA) [21, 26], don't rely on labeled data. But, PCA is based on a strong assumption of straight linear feature correlation, only uses two cross operations (i.e., addition/subtraction) to generate new features, and only reduces (can't increase) feature space dimensionality.

The emergence of LLM (e.g., ChatGPT) shows that mechanism-unknown human language knowledge can be embedded into a large embedding space and model the world as generative AI to

<sup>†</sup> Corresponding Author.

autoregress the next word. Following a similar spirit, we believe that mechanism-unknown feature space knowledge can be embedded into large sequential foundational models for generating a transformed feature set. For example, a transformed feature set  $(\frac{f_1}{f_2}, f_1 - f_2, \frac{f_1+f_2}{f_1})$  is seen as a feature cross sequence “ $f_1/f_2, f_1 - f_2, (f_1 + f_2)/f_1$ EOS”. This provides a great potential to transform the traditional ways that we solve UFTL.

**Our contribution: a graph, contrastive, and generative learning perspective.** We connect graph, contrastive, and generative learning to develop a measurement-pretrain-finetune paradigm, as shown in Figure 1. Specifically, 1) *unsupervised feature set utility measurement*: Under the unsupervised setting, downstream task feedback and labeled data are not available for measuring feature set utility. We propose a feature value consistency preservation angle. The insight is that: if two data instances (i.e., rows) are similar, both instances tend to be from the same class, therefore the two element values of a specific feature over the two instances should be close to each other. We reformulate feature set utility as information gain of total feature value consistency from the lens of instance pair similarity. 2) *unsupervised graph contrastive pretraining*: Aside from feature set utility, feature set representation needs to be learned in an unsupervised fashion. We see a transformed feature set as a feature-feature similarity graph, and reformulate the feature set2vec problem into an unsupervised graph contrastive learning task. By seeing a feature set as a graph rather than just a feature cross sequence, the graph contrastive pretraining can better model feature-feature interactions, learn a feature set level embedding space, and accelerate later finetuning and training convergence. 3) *multi-objective finetuning*: We regard a feature set as a feature cross sequence, and feature transformation as sequential feature cross generation. We find that: (i) deep sequential model can distill feature knowledge in the continuous embedding space, in order to drive autoregressive generation of transformation sequences; (ii) navigating optimization and generation in the continuous embedding space via gradient ascents can convert classic discrete search formulation into differentiable continuous optimization, and avoid large discrete search space.

**Summary of technical solution.** Inspired by these insights, this paper presents a principled and generic framework for the UFTL problem by the measurement-pretrain-finetune paradigm. Specifically, given an original dataset, we can explore various feature combinations of original features to prepare different feature sets. In *the measurement stage*, we generalize the idea of mean discounted cumulative gain to feature set utility measurement. We see “information gain” as the value consistency of a specific feature between two close data instances, and “cumulative” as the sum of the consistency gains of all data instance pairs. “discounted” means that: when a feature is of high variance, its element values tend to be different. Under such a situation, if the element values of such a feature over any two close instances still stay consistent, the cumulative gain of such a high-variance feature should be augmented. “mean” means that we average the cumulative discounted gain of all features as the final utility of a feature set. In *the pre-train stage*, we regard a feature set as a feature-feature similarity graph and develop a contrastive GNN to learn feature set embedding. In particular, to prepare contrastive pairs, we apply attribute

mask and edge perturbation to feature-feature similarity graphs to generate augmented graphs. An augmented graph will pair with a source augmented graph as a positive pair; an augmented graph will pair with a non-sourced augmented graph as a negative pair. We optimize the contrastive loss by enforcing that the embeddings of two augmented graphs from the same source graph should be similar, and different otherwise. In *the fine-tune stage*, we develop a deep generative feature transformation model. This model includes an encoder, a decoder, and an evaluator. The set2vec encoder is pretrained by and loaded from the pretrain stage to map a feature set to an embedding vector; the vec2seq decoder is to decode an embedding vector to generate a feature cross sequence that represents a feature set; the evaluator that predicts the feature set utility given feature set embedding can provide gradient-steered optimization to find the optimal transformed feature set embedding, by jointly optimizing the losses of decoder and evaluator.

## 2 PROBLEM STATEMENT

**Reinforcement Learning (RL) based Data Collector.** The success demonstrated by [28] highlights that feature learning knowledge can effectively be modeled using multi-agent systems. Thus, we employ its framework to automatically collect high-quality training data using our defined unsupervised feature set utility measurement for transformation embedding space. More details about the data collector are described in Appendix 8.1.

**Important Definitions.** 1) *Feature Cross*, which means we apply mathematical operations to original features to generate a new feature (e.g.,  $[\sin(f_1 + f_2)/f_3 - f_4]$ ). 2) *Explored Feature Set*. We combine multiple feature crosses to construct an explored feature set (e.g.,  $[\sin(f_1 + f_2)/f_3 - f_4, f_1 + \exp(f_2)]$ ). 3) *Feature Cross Sequence*. Figure 2(a) shows that we use a token sequence to represent the explored feature set. <SOS>, <SEP>, and <EOS> indicate the start, separate, and end tokens respectively. To use fewer tokens and reduce learning difficulty, this sequence can be represented by a postfix expression, as shown in Figure 2(b). In this paper, the term “feature cross sequence” refers to the representation of a transformed feature set using a postfix expression of token sequence.




Figure 2: Postfix Expression.

**The UFTL Problem.** Formally, given a data set  $D = (X, y)$  and an operation set  $O$ , where  $X$  is an original feature set and  $y$  is the corresponding label. Here,  $y$  is only used for testing purposes and is not involved in the training process. We first apply the RL-based data collector on  $X$  to collect  $N$  training data samples, denoted by  $\tilde{D} = \{\mathcal{F}_i, \mathcal{T}_i, s_i\}_{i=1}^N$ , where  $\mathcal{F}_i$ ,  $\mathcal{T}_i$ , and  $s_i$  indicate the explored feature set, feature cross sequence, and feature set utility score of the  $i$ -th sample respectively. We aim to 1) embed explored feature sets into a differentiable continuous space and 2) generate the optimal transformed feature set. To accomplish Goal 1, we designed an encoder-evaluator-decoder framework that compresses the feature learning knowledge of  $\tilde{D}$  into an embedding space  $\mathcal{E}$ , optimized through pre-training and fine-tuning processes. To achieve Goal

2, we adopt a gradient-steered search method to find the optimal embedding point  $E^*$  and reconstruct the best feature space  $\mathcal{F}^*$ . It can be formulated by

$$\mathcal{F}^* = \psi(E^*)[X] = \arg \max_{E \in \mathcal{E}} \mathcal{S}(\psi(E)[X]), \quad (1)$$

where  $\psi$  is the well-trained decoder that can regenerate the feature cross sequence from any learned embedding  $E$ ;  $\mathcal{S}$  is the function for calculating feature set utility;  $\psi(E)[X]$  means applying the reconstructed sequence on  $X$  to obtain the transformed feature set.

### 3 THE MEASUREMENT-PRETRAIN-FINETUNE PARADIGM

#### 3.1 Framework Overview

Figure 3 show our unsupervised generative feature transformation framework (NEAT) includes three steps: 1) Data collection via unsupervised feature set utility, 2) Graph contrastive pre-training, and 3) Multi-objective fine-tuning. In Step 1, we design an unsupervised measurement to assess feature set utilities. Under the guidance of this measurement, we employ the RL-based data collector to explore the original feature space. The collector will automatically gather explored feature sets, feature cross sequences, and feature set utility for training data preparation. In Step 2, we pre-train shared GNNs to capture the knowledge of the training data via graph contrastive learning and preserve it in an embedding space. In detail, we first extract feature-feature attributed graphs from the explored feature set. Then, sampling-based GNNs are employed to capture the characteristics of these graphs through contrastive learning. In Step 3, we conduct multi-objective fine-tuning to update the created embedding space and identify the optimal transformed feature space. More specifically, to update the embedding space, we load the pre-trained GNNs as an encoder and create a decoder and an evaluator for feature cross sequence reconstruction and utility score estimation respectively. Then, we search for enhanced embeddings within the new embedding space and reconstruct feature cross sequences. Finally, these sequences are used to transform the original feature space. The transformed feature set yielding the highest utility score is output as the optimal result.

#### 3.2 Unsupervised Feature Set Utility Measurement

**Why measuring unsupervised feature set level utility matters.** Feature set utility measurement reflects the quality of the feature space, which is critical for guiding the search toward the optimal feature set. Conventional strategies leverage supervised methods (e.g., decision trees, feature relevance) to measure individual-level feature importance and aggregate them to set-level utility. However, such strategies are unsuitable for the UFTL setting where downstream predictive task labels are unavailable or expensive. Note that we can use unsupervised measures, such as feature redundancy, to evaluate individual-level feature utility, and then perform average aggregation. But such a method only evaluates the negative (e.g., information overlap) aspect of a feature, ignores the meaningful interactions of feature-instance, instance-feature value, feature value-feature vector, and fails to consider self-weighted or

self-penalized aggregation when propagating information score from individual level to set level.

**Mean Discounted Cumulative Gain (MDCG): a feature value consistency preservation perspective.** To overcome these limitations, we propose a feature value consistency preservation angle that illustrates an important intuition: if two data instances (i.e., rows) are similar, then both instances tend to be from the same class, thereafter the two element values of a specific feature with respect to the two similar instances should be similar. We see the alignment with such intuition as information gain and develop an MDCG based unsupervised measure.

*Step 1: Information Gain.* We quantify the information gain of a feature as the score of alignment with the intuition: a feature’s element values with respect to two similar instances should be similar. Formally, given a feature set  $\mathcal{F} \in \mathbb{R}^{n \times m}$ ,  $n$  and  $m$  are the number of instances and features respectively. Let  $i$  and  $j$  are indexes of instances, the information gain of the  $q$ -th feature over the  $i$ -th and  $j$ -th instances is given by:  $g_{ij}^q = (\mathcal{F}_{iq} - \mathcal{F}_{jq})^2 e^{-\frac{\|\mathcal{F}_{i,*} - \mathcal{F}_{j,*}\|^2}{constant}}$ . In the experiments, we set the value of *constant* as 2. Here, the lower the value of  $g_{ij}^q$ , the greater the information gain observed.

*Step 2: Cumulative Gain.* The cumulative gain of the  $q$ -th feature over all the instance pairs is given by  $\sum_{ij} S_{ij} * g_{ij}^q$ , where  $S_{ij}$  is a binary indicator: 1) when the  $i$ (or  $j$ )-th instance is among the  $K$  nearest neighbors of the  $j$ (or  $i$ )-th instance,  $S_{ij} = 1$  indicates the  $i$ -th instance and the  $j$ -th instance are similar; 2) otherwise,  $S_{ij} = 0$  indicates the two instances are not similar.  $K$  is a hyperparameter given by users.

*Step 3: Discounted Cumulative Gain.* When a feature exhibits high variance and its element values tend to diverge, if the element values of such feature over any two similar instances still stay similar, the cumulative gain of such high-variance feature should be augmented. Therefore, we propose the concept of discounted cumulative gain:  $\frac{\sum_{ij} S_{ij} * g_{ij}^q}{Var(\mathcal{F}_{*,q})}$  to describe feature importance.

*Step 4: Mean Discounted Cumulative Gain.* Finally, we quantify the feature set utility as the average score of the discounted cumulative gains of all the features:  $\sum_{q=1}^m (1 - \frac{\sum_{ij} S_{ij} * g_{ij}^q}{Var(\mathcal{F}_{*,q})}) / m$ .

Leveraging the utility of the feature set, we employ the RL-based collector [28] to gather explored feature sets  $\mathcal{F}$ , feature cross sequences  $\mathcal{T}$ , and feature set utilities  $s$  for training data preparation.

#### 3.3 Unsupervised Feature-Feature Interaction Graph Contrastive Pre-Training

**Why pre-training a feature set embedding space matters.** After collecting the training data, learning the embedding space of the feature set becomes essential for two reasons: 1) feature set2vec can learn a continuous embedding space, so we can convert classic discrete search into effective gradient ascent in continuous optimization. 2) feature set2vec can develop a database of feature set embedding vectors along with corresponding unsupervised utilities that enable the learning of a feature set utility evaluator. This evaluator can model the relationship between feature set embeddings and feature set utilities and provide gradient information to search for optimal feature sets.




Figure 3: An overview of NEAT. First, we explore the original feature set to collect training data under the guidance of our proposed unsupervised feature set utility. Second, we pre-train shared GNNs to capture the knowledge of the training data via graph contrastive learning and preserve it in an embedding space. Finally, we conduct multi-objective fine-tuning to readjust the created embedding space and identify the optimal transformed feature set.

**Leveraging graph contrastive learning for unsupervised feature set2vec.** Feature set2vec, particularly under an unsupervised setting, is challenging. We propose a graph perspective to regard a feature set as a feature-feature similarity graph. This feature-feature graph can describe the geometric topology of a feature space, and convert set2vec into graph2vec, so graph contrastive learning can be introduced to achieve unsupervised embedding pretraining. We pre-train a GNN-based encoder to convert explored feature sets to vectors, thereby constructing a feature transformation embedding space. During the pre-training, we use the explored feature sets collected in the previous step to construct embedding space. To simplify notations, we use  $\mathcal{F}$  to represent any explored feature set.

**Creating Feature-Feature Similarity Graph.** Given an explored feature set  $\mathcal{F}$ , we create an adjacency matrix to represent the interactions among features, in which each element indicates the edge between two features. For each feature, we use the element values of the corresponding feature within  $\mathcal{F}$  as its attributes. Subsequently, we calculate the cosine similarity between every feature pair based on their attributes to determine the weight of the edge, constructing a feature-feature similarity graph. Finally, we convert the similarities adjacency matrix into a 0/1 matrix with the similarity 95th percentile as a cutoff.

**Graph Augmentation for Building Contrastive Pairs.** We augment each graph with edge perturbation and attribute masking to obtain augmented graphs for building contrastive pairs. The augmented graphs from the same graph form a positive pair, while the augmented graphs from different graphs form negative pairs. The details of the augmentation strategies are: 1) Edge Perturbation. A random proportion of edges is either added or dropped to perturb the connectivity of the graph. The assumption is that the graph has robustness to the pattern variances of edge connection. 2) Attribute Masking. A random proportion of attributes is masked, prompting the model to use contextual information to recover the obscured information. The assumption is that the missing attributes will not significantly affect the model’s predictions.

**GNN-based Encoder.** The purpose of the GNN-based encoder is to encode the graphs to extract graph-level representation vectors. Since the graphs have different structures (different adjacency matrices), the encoder needs to adapt to various types of graph structures and efficiently handle the learning of node representations. Inspired by [9], we develop two GNNs with shared weights to encode augmented graphs of different types respectively. Here we take one graph as an example to illustrate the encoding process. Given a node  $v$  from the graph, we sample its all neighbor nodes, denoted as  $\{u_1, u_2, \dots, u_p\}$ , where  $p$  is the number of neighbor nodes. The  $\mathbf{h}_v$  and  $\mathbf{h}_u$  are the representation of node  $v$  and node  $u$  respectively. Then we aggregate the representation of all neighbor nodes by calculating their mean:  $\mathbf{h}_u^* = \sum_1^p \mathbf{h}_u / p$ . We update the representation of node  $v$  by  $\mathbf{h}_v = \text{ReLU}(W \cdot \text{CONCAT}(\mathbf{h}_v, \mathbf{h}_u^*))$ , where  $\text{ReLU}$  is the activation function and  $W$  is the learnable weight matrix. We iterate this process with multi-step to update the node representation. Finally, we calculate the mean of all node representations in the graph to extract graph-level representation vector for the augmented graph, denoted as:  $\mathbf{h} = \sum \mathbf{h}_v / m$ , where  $m$  represents the number of nodes in the graph. Following the principles advocated in [3], we use a non-linear projection to project the graph-level representation vectors into a new latent space where the contrastive loss is calculated. In our framework, we use a two-layer Perceptron as a projection to map the  $\mathbf{h}$  into a new embedding vector  $\mathbf{z}$ .

**Minimize the Contrastive Loss.** Given a  $\mathcal{F}_i$  as the  $i$ -th explored feature set in training feature sets, we convert it into a graph and augment the graph to obtain two augmented graphs of different types. We feed the augmented graphs into the shared GNN respectively, to extract the embedding vectors, which are represented by  $\mathbf{z}_{i,1}$  and  $\mathbf{z}_{i,2}$  respectively. We calculate the distance between sample pairs using cosine similarity  $\text{sim}(\mathbf{z}_{i,1}, \mathbf{z}_{i,2}) = \mathbf{z}_{i,1}^T \mathbf{z}_{i,2} / (\|\mathbf{z}_{i,1}\| \|\mathbf{z}_{i,2}\|)$ , and develop a normalized temperature-scaled cross-entropy loss [23, 24] as contrastive loss to maximize the consistency between positive

pairs and increase the distance between negative pairs.

$$\mathcal{L}_{cl} = -\frac{1}{N} \sum_{n=1}^N \log \frac{\exp(\text{sim}(\mathbf{z}_{i,1}, \mathbf{z}_{i,2})/\tau)}{\sum_{i^*=1, i^* \neq i}^N \exp(\text{sim}(\mathbf{z}_{i,1}, \mathbf{z}_{i^*,2})/\tau)}, \quad (2)$$

where  $\tau$  denotes the temperature parameter. Finally, we minimize the contrastive loss to optimize the continuous embedding space.

### 3.4 Multi-Objective Fine-tuning

After pre-training, we conduct multi-objective fine-tuning to update the embedding space. This process aims to assess the efficacy of the embedding space and fully reconstruct the feature cross sequence, thereby facilitating the generation of the optimal transformed feature set. We reload the pre-trained GNN as an encoder to obtain the embedding. Additionally, we create an evaluator and a decoder for estimating feature set utility and reconstructing feature cross sequence respectively. Then, we identify the best embeddings within the updated embedding space and reconstruct feature cross sequences. Finally, these sequences are employed to transform the original feature space, and the transformed feature set that achieves the highest utility score is presented as the optimal result.

**Fine-tuning the Encoder-Decoder-Evaluator Structure.** We use the training data set  $\{\mathcal{F}_i, \mathcal{T}_i, s_i\}_{i=1}^N$  collected in section 3.2 to fine-tune the encoder-decoder-evaluator structure, where  $\mathcal{F}_i$ ,  $\mathcal{T}_i$  and  $s_i$  indicate explored feature set, feature cross sequence, and feature set utility of the  $i$ -th sample, respectively. To simplify notation, we use the notation  $(\mathcal{F}, \mathcal{T}, s)$  to represent any set of training data.

*The Encoder.* We input an explored feature set  $\mathcal{F}$  into the pre-trained encoder to map the feature set to an embedding vector  $\mathbf{z}$ .

*The Decoder.* The decoder is to reconstruct  $\mathcal{T}$  and built by LSTM [10]. Subsequently, we add a softmax layer after the LSTM to estimate the token distribution and infer the feature cross sequence based on the embedding  $\mathbf{z}$ . Specifically, assuming the token to be decoded is  $t_j$ , the partially decoded sequence is  $[t_1, t_2, \dots, t_{j-1}]$ , and the length of  $\mathcal{T}$  is  $k$ . Therefore, the probability of the  $j$ -th token is:

$$P_\psi(t_j|\mathbf{z}, [t_1, t_2, \dots, t_{j-1}]) = \frac{\exp(o_j)}{\sum_k \exp(o)}, \quad (3)$$

where  $o_j$  is the  $j$ -th output of the softmax layer. We use the negative log-likelihood of generating sequential tokens to minimize the reconstruction loss between the reconstructed tokens and the real one, denoted by:

$$\mathcal{L}_{rec} = -\sum_{j=1}^k \log P_\psi(t_j|\mathbf{z}, [t_1, t_2, \dots, t_{j-1}]) \quad (4)$$

*The Evaluator.* We add a fully connected layer after the encoder for regression prediction, which is given by:  $\tilde{s} = \theta(\mathbf{z})$ , where  $\theta$  is the function notation of the evaluator. Specifically, we use the Mean Squared Error (MSE) to measure the evaluator loss:

$$\mathcal{L}_{evt} = \text{MSE}(\tilde{s}, s), \quad (5)$$

where  $s$  is the corresponding feature set utility. Finally, we trade off the two losses to form a joint training loss:  $\mathcal{L} = \alpha \mathcal{L}_{evt} + \beta \mathcal{L}_{rec}$ , where  $\alpha$  and  $\beta$  are the hyperparameters to balance the influence of the two loss functions. We calculate the mean loss of the batch data to readjust the embedding space during the fine-tuning process.

**Optimization and Generation of the Optimal Transformed Feature Set.** After updating the embedding space, we employ a gradient-ascent search method to find better embedding. In the continuous space, a good starting point can accelerate the search process and enhance the downstream performance. Such good points are called initial seeds. Therefore, we first select the top  $K$  explored feature sets based on their corresponding feature set utilities as initial seeds to find better embeddings. More specifically, for an initial seed, we obtain an embedding  $\mathbf{z}$  through the well-trained encoder. Next, we use a gradient-ascent method to move the embedding with  $\eta$  steps along the direction maximizing the feature set utility. Formally, the moving calculation process can be described as:  $\mathbf{z}^+ = \mathbf{z} + \eta \frac{\partial \theta}{\partial \mathbf{z}}$ , where  $\mathbf{z}^+$  is the better embedding. We feed  $\mathbf{z}^+$  into the well-trained decoder to generate a feature cross sequence in an autoregressive manner until finding  $\langle \text{EOS} \rangle$ . We split the sequence into different segments according to  $\langle \text{SEP} \rangle$ , and follow the calculation rule of them to generate the transformed feature set. Finally, the transformed feature set yielding the highest utility score is output as the optimal result. The pseudo-code of the algorithm is released in Appendix 8.2.

## 4 EXPERIMENT

### 4.1 Experimental Setup


**Dataset Descriptions.** We utilize 23 publicly available datasets from UCI, LibSVM, Kaggle, and OpenML to conduct experiments. These datasets comprise 14 classification tasks (C) and 9 regression tasks (R). Table 1 shows the statistics of these datasets.

**Evaluation Metrics.** To alleviate the variance of the downstream ML model for fair comparisons, we employ random forests (RF) to assess the quality of transformed feature set. For classification tasks, we use F1-score, Precision, and Recall as evaluation metrics. For regression tasks, we use 1 - Relative Absolute Error (1-RAE), 1 - Mean Average Error (1-MAE), and 1 - Mean Square Error (1-MSE) as evaluation metrics. In both cases, higher values of the evaluation metrics indicate a higher quality of the transformed feature set.


**Baseline Algorithms and Model Variants of NEAT.** We compare our method with 9 widely-used feature transformation methods: (A) Unsupervised Methods: 1) **RDG** randomly generates feature-operation-feature transformation records to create a new feature space; 2) **PCA** [26] uses linear feature correlation to generate new features; 3) **LDA** [2] is a matrix factorization-based method that obtains decomposed latent representation as the generated feature space; (B) Supervised Methods: 3) **ERG** first applies a set of operations to each feature to expand the feature space, and then selects key features from it to form a new feature space; 4) **AFAT** [11] is an enhanced version of ERG, repeatedly generating new features and using a multi-step feature selection method to select more informative features for creating a new feature space; 5) **NFS** [4] embeds the decision-making process of feature transformation into a policy network and utilizes reinforcement learning (RL) to optimize the entire feature transformation process; 6) **TTG** [15] represents the feature transformation process as a graph and implements an RL-based discrete search method to find the optimal feature set; 7) **GRFG** [28] uses three collaborative reinforcement intelligent agents for feature generation and employs feature grouping strategies to accelerate agent learning; 8) **DIFER** [36] embeds randomly

**Table 1: Overall Performance.** In this table, the best and second-best results are highlighted by bold and underlined fonts respectively. We evaluate classification (C) and regression (R) tasks in terms of F1-score and 1-RAE respectively. The higher the value is, the better the quality of the transformed feature set is.

Dataset	Source	C/R	Samples	Features	RDG	PCA	LDA	ERG	AFAT	NFS	TTG	GRFG	DIFER	NEAT
Higgs Boston	UCI	C	50000	28	0.695	0.538	0.513	<u>0.702</u>	0.697	0.691	0.699	<b>0.707</b>	0.669	<u>0.702</u>
Amazon Employee	Kaggle	C	32769	9	0.932	0.923	0.916	0.932	0.930	0.932	<u>0.933</u>	0.932	0.929	<b>0.936</b>
PimaIndian	UCI	C	768	8	0.751	0.736	0.729	0.735	0.732	<u>0.791</u>	0.745	0.754	0.760	<b>0.811</b>
SpectF	UCI	C	267	44	0.760	0.709	0.665	0.788	0.760	<u>0.792</u>	0.760	0.776	0.766	<b>0.856</b>
SVMGuide3	LibSVM	C	1243	21	0.806	0.676	0.635	<u>0.818</u>	0.794	0.786	0.798	0.812	0.773	<b>0.832</b>
Geman Credit	UCI	C	1000	24	<u>0.701</u>	0.679	0.596	0.662	0.639	0.649	0.644	0.667	0.656	<b>0.723</b>
Credit Default	UCI	C	30000	24	0.802	0.768	0.743	0.803	0.804	0.801	0.798	<u>0.806</u>	0.796	<b>0.809</b>
Messidor_features	UCI	C	1151	19	0.627	0.672	0.463	<u>0.692</u>	0.657	0.650	0.655	0.680	0.660	<b>0.737</b>
Wine Quality Red	UCI	C	999	12	<u>0.496</u>	0.468	0.401	0.485	0.480	0.451	0.467	0.454	0.476	<b>0.522</b>
Wine Quality White	UCI	C	4898	12	0.524	0.469	0.438	<u>0.527</u>	0.516	0.525	0.521	0.518	0.507	<b>0.554</b>
Spambase	UCI	C	4601	57	0.906	0.815	0.889	<u>0.917</u>	0.912	<u>0.925</u>	0.919	0.922	0.912	<b>0.932</b>
Ap-omentum-ovary	Openml	C	275	10936	0.820	0.718	0.716	<u>0.849</u>	0.813	0.830	0.830	0.830	0.833	<b>0.865</b>
Lymphography	UCI	C	148	18	0.108	0.170	0.144	<u>0.202</u>	0.149	0.166	0.148	0.136	0.150	<b>0.298</b>
Ionosphere	UCI	C	351	34	0.942	0.928	0.743	<u>0.957</u>	0.942	0.943	0.932	0.956	0.905	<b>0.971</b>
Housing Boston	UCI	R	506	13	<u>0.434</u>	0.122	0.020	0.410	0.401	0.428	0.396	0.409	0.381	<b>0.478</b>
Airfoil	UCI	R	1503	5	0.519	0.514	0.207	0.519	0.507	0.520	0.500	0.521	<u>0.528</u>	<b>0.578</b>
Openml_586	Openml	R	1000	25	0.568	0.221	0.110	0.624	0.553	0.542	0.544	<u>0.646</u>	0.482	<b>0.700</b>
Openml_589	Openml	R	1000	50	0.509	0.224	0.011	0.610	0.508	0.503	0.504	<b>0.627</b>	0.463	<u>0.625</u>
Openml_607	Openml	R	1000	50	0.521	0.104	0.107	0.549	0.509	0.518	0.522	<b>0.555</b>	0.476	<u>0.550</u>
Openml_616	Openml	R	500	50	0.070	0.057	0.024	0.193	0.162	0.147	0.156	<u>0.372</u>	0.076	<b>0.536</b>
Openml_618	Openml	R	1000	50	0.472	0.078	0.052	0.561	0.473	0.467	0.467	<u>0.577</u>	0.408	<b>0.616</b>
Openml_620	Openml	R	1000	25	0.511	0.088	0.029	<u>0.546</u>	0.521	0.519	0.512	0.530	0.442	<b>0.584</b>
Openml_637	Openml	R	500	50	0.136	0.072	0.043	0.064	0.161	0.146	0.144	<b>0.307</b>	0.072	<u>0.278</u>



**Figure 4: The impact of graph contrastive pre-training.**



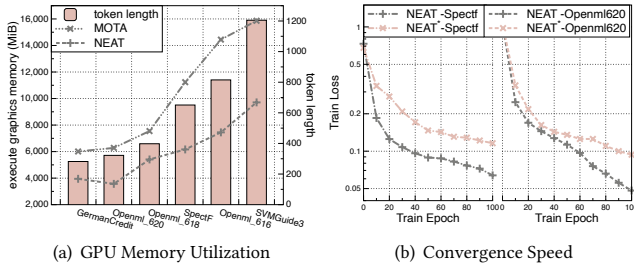
**Figure 5: The impact of unsupervised feature set utility measurement, RL-based data collector, and initial seeds.**

generated feature cross sequences and then uses a greedy search to find the separated transformed feature.

Additionally, to evaluate the necessity of each technical component of NEAT, we implement four model variants: 1) **NEAT\*** removes the pre-training stage and directly optimizes the encoder-evaluator-decoder structure for creating the feature transformation embedding space. 2) **NEAT+** utilizes the training data collected by randomly generating explored feature sets to create the feature transformation embedding space. 3) **NEAT-** replaces the mean discounted cumulative gain with feature redundancy for creating the feature transformation embedding space. 4) **NEAT#** replaces

the initial seeds with worst-ranked transformation sequences’ embeddings based on the utility value to identify better feature spaces.

**Hyperparameters and Reproducibility.** 1) Feature Transformation Knowledge Acquisition: We use the reinforcement data collector to collect explored feature set-feature utility score pairs. The collector explores 512 episodes, and each episode includes 10 steps. 2) Feature-Feature Similarity Graph Construction: The threshold for creating an edge between two features is the 95th percentile of all similarity values. 3) Graph Contrastive Pre-training: We map the attribute of each node to a 64-dimensional embedding, and use a 2-layer GNN network and a 2-layer projection head to integrate



**Figure 6: (a) GPU memory utilization comparison between NEAT and MOTA in different datasets. (b) Convergence speed comparison between NEAT and NEAT\***

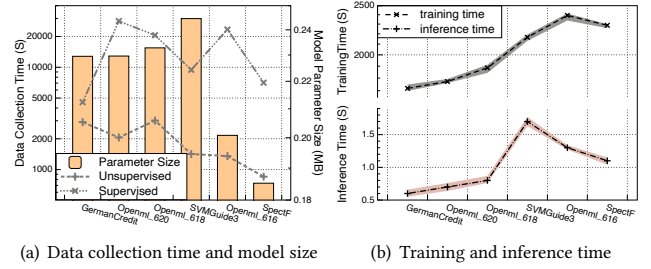
such information. For graph augmentation, we randomly perturbed 20% of the edges or masked 20% of the attributes to get diversified graphs. During the optimization, the batch size is 1024, the pre-training epochs are 100, the learning rate is 0.001, and the value of  $\tau$  is set to 0.5. 4) Multi-objective Fine-tuning: The decoder is a 1-layer LSTM network, which reconstructs a feature cross sequence. The evaluator is a 2-layer feed-forward network, in which the dimension of each layer is 200. During optimization,  $\alpha$  and  $\beta$  are 10 and 0.1 respectively, the batch size is 1024, the fine-tuning epochs are 500, and the learning rate is 0.001.

**Environmental Settings.** All experiments are conducted on the Ubuntu 22.04.3 LTS operating system, Intel(R) Core(TM) i9-13900KF CPU@ 3GHz, and 1 way RTX 4090 and 32GB of RAM, with the framework of Python 3.11.4 and PyTorch 2.0.1.

## 4.2 Experimental Results

**Overall performance.** In this experiment, we compare the performance of NEAT and baseline models for feature transformation in terms of F1-score and 1-RAE. Table 1 shows the comparison results. We can find that in most cases, the NEAT performs the best, and in the other four cases, it ranks second-best. But NEAT performs the best across all unsupervised approaches. The underlying driver for this observation is that NEAT can accurately capture the internal principle of the features for identifying the optimal feature space through the integration of unsupervised measurement, pre-training, and fine-tuning stages. Moreover, another interesting observation is that NEAT exhibits a higher degree of robustness and effectiveness in its transformation performance when applied to classification tasks as compared to regression tasks. A potential reason for this observation is that NEAT effectively captures the intricate relationships within the feature space, generating non-linear features to enhance the distinguishability of the feature space. Such transformations have more influence on classification tasks than regression tasks. In conclusion, this experiment shows the effectiveness of NEAT in feature transformation, underscoring the great potential of generative AI in this domain.

**The impact of the graph contrastive learning based pretraining.** One of the most important novelties of NEAT is to involve graph contrastive learning to pre-train the transformation embedding space. To analyze the influence of this component, we develop two model variants: 1) NEAT\*; 2) MOTA [29] is similar to NEAT, but it lacks the pre-training stage and employs an LSTM encoder. To ensure fair comparisons, we use our defined feature utility score



**Figure 7: (a) Data collection time comparison in terms of unsupervised and supervised collectors. Model parameter size check. (b) Model training and inference time check.**

as feature space measurement in both MOTA and NEAT. Figure 4 shows the comparison results on both classification and regression tasks. We can find that NEAT outperforms NEAT\* and MOTA in all cases. The underlying driver is that the pre-training stage initializes an effective transformation embedding space by capturing intricate relationships within the feature space. This initialization integrates with multi-objective fine-tuning to enhance the identification of the optimal feature space. Another interesting observation is that NEAT\* beats MOTA in most cases. A potential reason is that the GNN-based encoder is adept at grasping feature-feature interactions and correlations compared with LSTM to produce an improved transformation embedding space for better reconstruction. In summary, this experiment demonstrates the necessity of the pre-training stage and the importance of the GNN-based encoder.

**The impact of unsupervised feature set utility measurement, RL-based data collector, and initial seeds.** In NEAT, to create an effective transformation embedding space, we adopt three strategies: 1) Using an RL-based data collector [28] to automatically collect high-quality transformation records. 2) Developing an unsupervised measurement to evaluate the transformation records' utility. 3) Employing top-ranked transformation sequences, determined by their feature utility score, as initial seeds to search for enhanced transformation sequences. We develop three model variants to study their impact: NEAT<sup>+</sup>, NEAT<sup>-</sup>, and NEAT<sup>#</sup>. Figure 5 shows the comparison results. We find that NEAT surpasses NEAT<sup>+</sup>, NEAT<sup>-</sup> and NEAT<sup>#</sup> across all datasets. There are three underlying drivers: a) The RL-based collector can mimic human intelligence in feature learning to produce high-quality transformation records. These records help create an informative and distinguishable embedding space to identify the optimal feature space. b) The mean discounted cumulative gain, which considers the feature value consistency preservation, can better evaluate the feature set utility compared to feature redundancy. c) NEAT converts feature generation as a continuous optimization task, in which a good initialization can enhance the search for the optimal transformation sequence.

**Analysis of GPU utilization and convergence speed for NEAT.** To further analyze the GPU utilization and convergence speed of NEAT, we compare it with both MOTA and NEAT\*. Figure 6 shows the comparison results. From Figure 6(a), we find that NEAT significantly saves the required GPU memory compared with MOTA. A potential reason is that MOTA employs LSTM as the encoder, with the increase in feature space dimension, the required GPU memory will increase correspondingly. Instead, NEAT organizes the original

feature space as a feature-feature attributed graph and utilizes GNNs to aggregate information into embedding. This learning approach is more memory-efficient. From Figure 6(b), we observe that NEAT converges faster and achieves lower loss compared with NEAT\*. A possible reason is that the pre-training stage preserves complex feature-feature relationships within the transformation embedding space. This accelerates the learning process of the fine-tuning stage to achieve model convergence. In summary, this experiment shows the efficient GPU memory usage and fast convergence of NEAT.

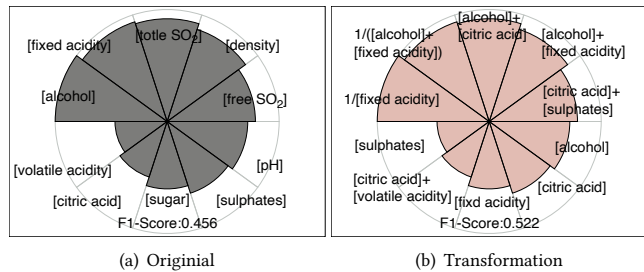
**Robustness check.** To study the robustness of NEAT, we compare its transformation performance and other baselines across various downstream ML models. We replace the Random Forest (RF) with XGBoost (XGB), Support Vector Machine (SVM), K-Nearest Neighborhood (KNN), Decision Tree (DT), and LASSO respectively. Table 2 shows the F1-score comparison results for the German Credit dataset. We can find that NEAT consistently exhibits superior performance in most cases. A potential reason is that the unsupervised feature utility measurement of NEAT leads it to grasp more general and intrinsic characteristics of the feature space. Such information is captured independently of the target labels, thereby enhancing the robustness and generalization capabilities of NEAT. In summary, this experiment demonstrates the robustness of NEAT, which maintains consistent performance across various ML models.

**Time and space complexity analysis.** To check the time and space complexity of NEAT, we select 6 datasets to compare data collection time, model parameter size, training time, and inference time. Figure 7(a) shows the comparison results in terms of the time cost of the data collector and model parameter size. We find that the data collector using our defined unsupervised feature utility measurement requires much less collection time compared with the supervised collector using the target label. This observation shows the computational efficiency of our unsupervised measurement. A possible reason is that the supervised collector has to spend much time retraining the downstream ML model from scratch at each iteration. Moreover, we notice that NEAT can maintain a very small parameter size across all datasets. A potential reason is that our sampling-based GNN encoder leads NEAT to focus on key information within the feature space instead of the global situation, significantly increasing the scalability of NEAT. Furthermore, Figure 7(b) shows the comparison results in terms of training and inference time costs. We find that once NEAT converges, it can quickly infer and identify the optimal feature space. The exceptional inference time of NEAT underscores its great practical potential. We further compare the time cost of NEAT with various baselines in Appendix 8.3, and analyze the time cost of the different components within NEAT in Appendix 8.5.

**Case study.** We employ our proposed unsupervised feature utility measurement (See Step 3 in Section 3.2) to calculate the feature importance of the original and transformed feature space respectively using the Wine Quality Red dataset. This task aims to predict the wine quality using the collected features. Figure 8 shows the comparison results in terms of the top 10 important features. In each pie chart, the importance of a feature is indicated by the area of the corresponding slice. The text labeling each slice of the pie chart denotes the name of the associated feature. We observe that in the transformed feature space, 60% of the top 10 features are generated by NEAT. This contributes to 14.5% improvements in wine quality

**Table 2: Robustness check. Test the transformed feature set on various downstream ML models.**

	RF	XGB	SVM	KNN	DT	LASSO
RDG	0.701	0.690	0.694	0.660	0.652	0.666
PCA	0.679	0.641	0.657	0.641	0.676	0.659
LDA	0.596	0.629	0.589	0.622	0.604	0.589
ERG	0.662	0.669	0.599	0.619	0.630	<b>0.684</b>
AFAT	0.639	0.688	0.639	0.643	0.691	0.647
NFS	0.649	0.690	0.695	0.635	0.663	0.672
TTG	0.644	0.690	0.695	0.660	0.663	0.666
GRFG	0.667	0.656	0.685	0.649	0.687	0.676
DIFER	0.656	0.696	0.617	0.658	0.683	0.672
NEAT	<b>0.723</b>	<b>0.707</b>	<b>0.703</b>	<b>0.661</b>	<b>0.692</b>	0.676



**Figure 8: Case study. Comparison of traceability on the original feature space and transformed feature space.**

prediction. Such an observation indicates that NEAT comprehends the complex inherent relationships of the feature set, thereby generating a better feature space. Furthermore, we notice that 'alcohol' and 'fixed acidity' are often employed to generate new informative features, such as '1/[fixed acidity]' and '1/([alcohol]+[fixed acidity])'. The observation aligns with known wine knowledge: alcohol enhances taste and body, while fixed acidity influences tartness and stability, collectively determining wine quality.

## 5 RELATED WORK

Feature Engineering can derive a new feature set from original features to improve downstream predictive performance [5, 8, 17, 32]. Humans can manually reconstruct a feature set with domain knowledge and empirical experiences. These methods are explicit, traceable, and explainable, but are incomplete and time-consuming due to their heavy reliance on domain expertise and empirical insights. Supervised transformations include exhaustive-expansion-reduction [11, 12, 14, 16, 19], iterative-feedback-improvement [15, 25, 28, 31, 33, 35], and AutoAI-based approaches [1, 4, 18, 22, 27, 30, 34, 36]. Exhaustive-expansion-reduction methods use explicit [13] or greedy [6] techniques to construct feature cross, assessing their impact on downstream ML tasks and then proceeding with feature selection to keep crucial features. However, these approaches heavily rely on domain knowledge and labels, is hard to construct complex feature crosses, resulting in suboptimal performance. Iterative-feedback-improvement methods leverage reinforcement learning, genetic algorithms, or evolution algorithms to formulate feature transformation as a discrete optimization task, searching for the optimal feature set iteratively. However, they face exponentially growing combination possibilities, are computationally costly, hard



to converge, and unstable. AutoAI-based methods utilize AutoML [7, 20] techniques to systematically explore various feature crosses in a structured, non-exhaustive manner. However, they are constrained by their inability to generate high-order features and their sensitivity to performance instability. In contrast, unsupervised transformations like Principal Components Analysis (PCA) [21] don't rely on labels. However, PCA is based on a strong assumption of straight linear feature correlation, only uses addition/subtraction to generate new features, and leads to only dimensionality reduction rather than an increase. To fulfill these gaps, we develop a measurement-pretrain-finetune paradigm, formulating Feature Transformation as an unsupervised continuous optimization task. This method constructs a continuous space with unsupervised graph contrastive learning, facilitating the generation of the best feature set. By decoupling FT from downstream tasks and avoiding ineffective discrete searches, our approach provides a more efficient solution.

## 6 CONCLUSION

In this paper, we integrate graph, contrastive, and generative learning to develop a label-free generative framework for unsupervised feature transformation learning problems. There are three important contributions: 1) We develop an unsupervised measurement to evaluate feature set utilities, decoupling the framework from downstream tasks, and avoiding the evaluation of feature sets from lengthy and expensive experiments. 2) We regard a feature set as a feature-feature similarity graph and introduce an unsupervised training paradigm, embedding feature transformation knowledge into a continuous space through graph contrastive pre-training. 3) We see feature transformation as a sequential generation task and develop an encoder-decoder-evaluator structure, which facilitates the identification of the optimal transformed feature set through gradient optimization. This structure enables the generation of optimal results, avoiding the need to explore exponentially growing possibilities of feature combinations in discrete space. Extensive experiments validate the effectiveness, efficiency, traceability, and clarity of our framework. The superior experimental performance highlights its potential as a valuable approach for unsupervised generative feature transformation in large quantities of domains, such as bioinformatics, health care, finance analysis, etc.

## 7 ACKNOWLEDGEMENT

This research was partially supported by the National Science Foundation (NSF) via the grant number: GR45865.

## REFERENCES

- [1] Maroua Bahri, Flavia Salutari, Andrian Putina, and Mauro Sozio. 2022. AutoML: state of the art with a focus on anomaly detection, challenges, and research directions. *International Journal of Data Science and Analytics* 14, 2 (2022), 113–126.
- [2] David M Blei, Andrew Y Ng, and Michael I Jordan. 2003. Latent dirichlet allocation. *Journal of machine Learning research* 3, Jan (2003), 993–1022.
- [3] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. 2020. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*. PMLR, 1597–1607.
- [4] Xiangning Chen, Qingwei Lin, Chuan Luo, Xudong Li, Hongyu Zhang, Yong Xu, Yingnong Dang, Kaixin Sui, Xu Zhang, Bo Qiao, et al. 2019. Neural feature search: A neural architecture for automated feature engineering. In *2019 IEEE International Conference on Data Mining (ICDM)*. IEEE, 71–80.
- [5] Yi-Wei Chen, Qingquan Song, and Xia Hu. 2021. Techniques for Automated Machine Learning. *SIGKDD Explor. Newsl.* 22, 2 (jan 2021), 35–50.
- [6] Ofer Dor and Yoram Reich. 2012. Strengthening learning algorithms by feature discovery. *Information Sciences* 189 (2012), 176–190.
- [7] Thomas Elsken, Jan Hendrik Metzen, and Frank Hutter. 2019. Neural architecture search: A survey. *The Journal of Machine Learning Research* 20, 1 (2019), 1997–2017.
- [8] Nanxu Gong, Wangyang Ying, Dongjie Wang, and Yanjie Fu. 2024. Neuro-Symbolic Embedding for Short and Effective Feature Selection via Autoregressive Generation. *arXiv preprint arXiv:2404.17157* (2024).
- [9] Will Hamilton, Zhitao Ying, and Jure Leskovec. 2017. Inductive representation learning on large graphs. *Advances in neural information processing systems* 30 (2017).
- [10] Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation* 9, 8 (1997), 1735–1780.
- [11] Franziska Horn, Robert Pack, and Michael Rieger. 2020. The autofeat Python Library for Automated Feature Engineering and Selection. In *Machine Learning and Knowledge Discovery in Databases*, Peggy Cellier and Kurt Driessens (Eds.). Springer International Publishing, Cham, 111–120.
- [12] James Max Kanter and Kalyan Veeramachaneni. 2015. Deep feature synthesis: Towards automating data science endeavors. In *2015 IEEE International Conference on Data Science and Advanced Analytics (DSAA)*. 1–10. <https://doi.org/10.1109/DSAA.2015.7344858>
- [13] Gilad Katz, Eui Chul Richard Shin, and Dawn Song. 2016. Explorekit: Automatic feature generation and selection. In *2016 IEEE 16th International Conference on Data Mining (ICDM)*. IEEE, 979–984.
- [14] Udayan Khurana, Fatemeh Nargesian, Horst Samulowitz, Elias Khalil, and Deepak Turaga. 2016. Automating feature engineering. *Transformation* 10, 10 (2016), 10.
- [15] Udayan Khurana, Horst Samulowitz, and Deepak Turaga. 2018. Feature engineering for predictive modeling using reinforcement learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 32.
- [16] Udayan Khurana, Deepak Turaga, Horst Samulowitz, and Srinivasan Parthasarathy. 2016. Cognito: Automated Feature Engineering for Supervised Learning. In *2016 IEEE 16th International Conference on Data Mining Workshops (ICDMW)*. 1304–1307. <https://doi.org/10.1109/ICDMW.2016.0190>
- [17] A. Kusiak. 2001. Feature transformation methods in data mining. *IEEE Transactions on Electronics Packaging Manufacturing* 24, 3 (2001), 214–221. <https://doi.org/10.1109/6104.956807>
- [18] Hoang Thanh Lam, Beat Buesser, Hong Min, Tran Ngoc Minh, Martin Wistuba, Udayan Khurana, Gregory Bramble, Theodoros Salonidis, Dakuo Wang, and Horst Samulowitz. 2021. Automated Data Science for Relational Data. In *2021 IEEE 37th International Conference on Data Engineering (ICDE)*. 2689–2692. <https://doi.org/10.1109/ICDE51399.2021.00305>
- [19] Hoang Thanh Lam, Johann-Michael Thiebaud, Mathieu Sinn, Bei Chen, Tiep Mai, and Oznu Alkan. 2017. One button machine for automating feature engineering in relational databases. *arXiv:1706.00327 [cs.DB]*
- [20] Yaliang Li, Zhen Wang, Yuexiang Xie, Bolin Ding, Kai Zeng, and Ce Zhang. 2021. Automl: From methodology to application. In *Proceedings of the 30th ACM International Conference on Information & Knowledge Management*. 4853–4856.
- [21] Andrzej Maćkiewicz and Waldemar Ratajczak. 1993. Principal components analysis (PCA). *Computers & Geosciences* 19, 3 (1993), 303–342.
- [22] Alhassan Mumuni and Fuseini Mumuni. 2024. Automated data processing and feature engineering for deep learning and big data applications: a survey. *Journal of Information and Intelligence* (2024).
- [23] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. 2018. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748* (2018).
- [24] Kihyuk Sohn. 2016. Improved deep metric learning with multi-class n-pair loss objective. *Advances in neural information processing systems* 29 (2016).
- [25] Binh Tran, Bing Xue, and Mengjie Zhang. 2016. Genetic programming for feature construction and selection in classification on high-dimensional data. *Memetic Computing* 8 (2016), 3–15.
- [26] Md Palash Uddin, Md Al Mamun, and Md Ali Hossain. 2021. PCA-based feature reduction for hyperspectral remote sensing image classification. *IETE Technical Review* 38, 4 (2021), 377–396.
- [27] Dakuo Wang, Josh Andres, Justin D Weisz, Erick Oduor, and Casey Dugan. 2021. Autods: Towards human-centered automation of data science. In *Proceedings of the 2021 CHI conference on human factors in computing systems*. 1–12.
- [28] Dongjie Wang, Yanjie Fu, Kumpeng Liu, Xiaolin Li, and Yan Solihin. 2022. Group-wise reinforcement feature generation for optimal and explainable representation space reconstruction. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. 1826–1834.
- [29] Dongjie Wang, Meng Xiao, Min Wu, Yuanchun Zhou, Yanjie Fu, et al. 2024. Reinforcement-enhanced autoregressive feature transformation: Gradient-steered search in continuous space for postfix expressions. *Advances in Neural Information Processing Systems* 36 (2024).
- [30] Marcel Wever, Alexander Tornede, Felix Mohr, and Eyke Hüllermeier. 2021. AutoML for multi-label classification: Overview and empirical evaluation. *IEEE transactions on pattern analysis and machine intelligence* 43, 9 (2021), 3037–3054.
- [31] Meng Xiao, Dongjie Wang, Min Wu, Kumpeng Liu, Hui Xiong, Yuanchun Zhou, and Yanjie Fu. 2022. Self-Optimizing Feature Transformation. *arXiv:2209.08044*

- [32] Wangyang Ying, Dongjie Wang, Haifeng Chen, and Yanjie Fu. 2024. Feature Selection as Deep Sequential Generative Learning. *arXiv preprint arXiv:2403.03838* (2024).
- [33] Wangyang Ying, Dongjie Wang, Kunpeng Liu, Leilei Sun, and Yanjie Fu. 2023. Self-optimizing feature generation via categorical hashing representation and hierarchical reinforcement crossing. In *2023 IEEE International Conference on Data Mining (ICDM)*. IEEE, 748–757.
- [34] Ziwei Zhang, Xin Wang, and Wenwu Zhu. 2021. Automated machine learning on graphs: A survey. *arXiv preprint arXiv:2103.00742* (2021).
- [35] Guanghui Zhu, Shen Jiang, Xu Guo, Chunfeng Yuan, and Yihua Huang. 2022. Evolutionary Automated Feature Engineering. In *Pacific Rim International Conference on Artificial Intelligence*. Springer, 574–586.
- [36] Guanghui Zhu, Zhuoer Xu, Chunfeng Yuan, and Yihua Huang. 2022. DIFER: differentiable automated feature engineering. In *International Conference on Automated Machine Learning*. PMLR, 17–1.

## 8 APPENDIX

### 8.1 RL data collector

Our perspective is to view reinforcement intelligence as a training data collector in order to achieve volume (self-learning enabled automation), diversity(exploration), and quality (exploitation). To implement this idea, we design reinforcement agents to automatically decide how to perform feature crosses and generate new feature spaces. The reinforcement exploration experiences and corresponding accuracy will be used as training data. Specifically, the approach includes: 1) **Multi-Agents**: We design three agents to perform feature cross: a head feature agent, an operation agent, and a tail feature agent. 2) **Actions**: In each reinforcement iteration, the three agents collaborate to select a head feature, an operator (e.g., +, -, \*, /), and a tail feature to generate a new feature. The newly generated feature is later added to the feature set for the next feature generation. 3) **Environment**: The environment is the feature space, representing an updated feature set. When three feature agents take actions to cross two features to generate a new feature and add the new feature to the previous feature set, the state of feature space (environment) changes. The state represents the statistical characteristics of the selected feature subspace. 4) **Reward Function**: The reward is an unsupervised feature set utility measurement proposed in methodology. 5) **Training and Optimization**: Our reinforcement data collector includes many feature cross steps. Each step includes two stages: control and training. In the control stage, each feature agent takes actions based on their policy networks, which take the current state as input and output recommended actions and the next state. The three feature agents will change the size and contents of a new feature space. We regard the new feature space as an environment. Meanwhile, the actions taken by feature agents generate an overall reward. This reward will then be assigned to each participating agent. In the training stage, the agents train their policies via experience replay independently. The agent uses its corresponding mini-batch samples to train its Deep Q-Network (DQN), in order to obtain the maximum long-term reward based on the Bellman Equation.

### 8.2 Pseudo-code of the algorithm

The entire procedure algorithm is described as Algorithm 1.

### 8.3 Time complexity check with baselines

We compared our method with baseline methods on 6 selected datasets over classification and regression tasks to report their time

---

### Algorithm 1: Entire Procedure

---

**Input** : The original dataset  $D = (X, y)$   
**Output** : The Optimal Transformed Feature set  $\mathcal{F}^*$

- 1 . Initialize the encoder  $\phi$ , decoder  $\psi$  and evaluator  $\theta$ .
- 2 **Unsupervised Measurement (Section 3.2)**:
- 3 Develop an unsupervised measurement (MDCG) to assess the utility of the feature set.
- 4 Using RL and MDCG to collect  $N$  training data  
 $\tilde{D} = \{\mathcal{F}_i, \mathcal{T}_i, s_i\}_{i=1}^N$
- 5 **Graph Contrastive Pre-training (Section 3.3)**:
- 6 **for in pretrain – epoch do**
- 7   Construct feature-feature similarity graphs  $\mathcal{F}_i \rightarrow \mathcal{G}_i$ .
- 8   Graph augmentation  $\mathcal{G}_i \rightarrow (\mathcal{G}_{i,1}, \mathcal{G}_{i,2})$ .
- 9   Encode:  $\mathbf{z}_{i,1} = \phi(\mathcal{G}_{i,1}), \mathbf{z}_{i,2} = \phi(\mathcal{G}_{i,2})$
- 10   Contrastive learning:  

$$\mathcal{L}_{cl} = -\frac{1}{N} \sum_{n=1}^N \log \frac{\exp(\text{sim}(\mathbf{z}_{i,1}, \mathbf{z}_{i,2})/\tau)}{\sum_{i^* \neq i} \exp(\text{sim}(\mathbf{z}_{i,1}, \mathbf{z}_{i^*,2})/\tau)}$$
- 11 **end**
- 12 **Multi-Objective Fine-tuning (Section 3.4)**:
- 13 **for in finetune – epoch do**
- 14   Encode:  $\mathbf{z} = \phi(\mathcal{G})$ .
- 15   Estimate:  $\mathbf{m}, \sigma$ .
- 16   Decode loss:  $\mathcal{L}_{rec} = -\log P_{\psi}(\mathcal{T}|\mathbf{z})$ .
- 17   Evaluate loss:  $\mathcal{L}_{evt} = \text{MSE}(s, \theta(\mathbf{z}))$ .
- 18   Backward:  $\mathcal{L} = \alpha \mathcal{L}_{evt} + \beta \mathcal{L}_{rec}$
- 19 **end**
- 20 Select top- $k$  transformed feature sets (graphs)  $(\mathcal{G})^k$  from  $\tilde{D}$ .
- 21 Encode:  $(\mathbf{z})^k = \phi((\mathcal{G})^k)$ .
- 22 Update  $(\mathbf{z})^k$  with  $\eta$  steps:  $(\mathbf{z}^+)^k = (\mathbf{z})^k + \eta * \frac{\partial \mathcal{L}}{\partial (\mathbf{z})^k}$ .
- 23 Decode:  $(\mathcal{T}^+)^k = \psi((\mathbf{z}^+)^k)$ .
- 24 Optimal Transformed Feature Set:  

$$\mathcal{F}^* = \arg \max \mathcal{S}((\mathcal{T}^+)^k[X]).$$

---

costs, as shown in Table 3. The unsupervised baselines cost less time, but are less accurate. Compared with the supervised baselines, our method costs less time than DIFFER and MOTA and costs a bit more time than the other four supervised. However, 1). Feature transformation is not a data preparation step and not timing critical for most tasks; 2) The time cost increase (mins level) of training our method is acceptable. Compared with the time costs (days or months level) of human manual feature engineering, our method saves time; 3) Although our method costs a bit more time, it achieves much better feature engineering performance; 4) The inference time is small and it generates transformed features quickly in 0.6–1.7 s; 5) Based on such unsupervised and encode-decode generative design, we can use the strategy of pretraining a foundation model and then finetuning to reduce training time costs.

### 8.4 The impact of different RL

We conducted a thorough examination of the impact of training data generated by policy gradient and value learning reinforcement learning (RL) algorithms, as well as data generated randomly, as shown in Table 4. Our findings indicate that training data produced

**Table 3: Time cost comparisons with baselines**

Dataset	RDG(s)	PCA(s)	LDA(s)	ERG(s)	AFAT(s)	NFS(s)	TTG(s)	GRFG(s)	DIFFER(min)	MOTA(min)	NEAT(min)
spectf	2.9	0.3	0.3	4.9	3.8	8.2	9.1	282	70.8	156.0	52.3
svmguid3	7.0	0.2	0.8	10.3	8.8	7.1	8.3	373	68.1	191.8	59.7
german	3.6	0.2	0.7	6.0	3.5	6.1	7.9	182	85.6	105.6	78.1
openml_616	27.6	0.2	0.2	34.1	2.6	18.1	24.7	936	75.6	430.1	62.1
openml_618	67.3	0.2	0.2	82.4	3.0	30.9	38.7	821	90.3	373.8	81.7
openml_620	33.4	0.2	0.2	42.1	2.1	17.0	22.6	1131	74.5	500.9	63.8

**Table 4: The impact of different RL.**

Dataset	value iteration RL	Policy gradient RL	randomly generation
spectf	0.856	0.856	0.823
svmguid3	0.832	0.830	0.821
german	0.723	0.726	0.715
openml_616	0.536	0.540	0.493
openml_618	0.616	0.612	0.520
openml_620	0.584	0.585	0.564

**Table 5: Time complexity check of the different components in NEAT.**

Dataset	RL(min)	Model training(min)	Inference(s)
spectf	14.2	38.1	1.1
svmguid3	23.6	36.1	1.1
german	48.5	29.6	0.6
openml_616	22.5	39.8	1.3
openml_618	50.2	31.4	0.8
openml_620	34.3	29.5	0.7

using RL algorithms significantly outperforms randomly generated

data in terms of enhancing model performance. However, it is important to note that despite the overall superiority of RL-generated data, the variations in performance between different RL designs are relatively limited. This suggests that while the choice of using RL over random generation is crucial for improving results, the specific type of RL algorithm employed may not drastically alter the final outcomes.

### 8.5 Time complexity check of the different components in NEAT

We reported the computational overheads of 1) RL data collection; 2) model training; 3) inference, as shown in Table 5. We found that the additional time cost incurred by NEAT occurs during the data collection and model training phases. However, once the model converges, NEAT’s inference time is significantly reduced. The potential driving factors are that the RL-based collector spends more time gathering high-quality data, and the sequence formulation across the entire feature space increases the learning time cost for the sequence model. Nonetheless, during the inference phase, NEAT outputs the entire feature transformation at once, resulting in a very short inference time.