

GAProtoNet: A Multi-head Graph Attention-based Prototypical Network for Interpretable Text Classification

Ximing Wen¹, Wenjuan Tan², Rosina O. Weber¹

¹College of Computing and Informatics, Drexel University, Philadelphia, USA

²Department of Computer Science, Tsinghua University, Beijing, China
{xw384, rw37}@drexel.edu tanwj21@mails.tsinghua.edu.cn

Abstract

Pretrained transformer-based Language Models (LMs) are well-known for their ability to achieve significant improvement on text classification tasks with their powerful word embeddings, but their *black-box* nature, which leads to a lack of interpretability, has been a major concern. In this work, we introduce GAProtoNet, a novel *white-box* Multi-head Graph Attention-based Prototypical Network designed to explain the decisions of text classification models built with LM encoders. In our approach, the input vector and prototypes are regarded as nodes within a graph, and we utilize multi-head graph attention to selectively construct edges between the input node and prototype nodes to learn an interpretable prototypical representation. During inference, the model makes decisions based on a linear combination of activated prototypes weighted by the attention score assigned for each prototype, allowing its choices to be transparently explained by the attention weights and the prototypes. Experiments on multiple public datasets show our approach achieves superior results without sacrificing the accuracy of the original black-box LMs. We also compare with four alternative prototypical network variations and our approach achieves the best accuracy and F1 among all. Our case study and visualization of prototype clusters also demonstrate the efficiency in explaining the decisions of black-box models built with LMs. Our codes are available at <https://github.com/ximingwen/GAProtoNet>.

1 Introduction

Deep learning models, especially transformer-based Language Models (LMs) such as BERT (Devlin et al., 2018), RoBERTa (Liu et al., 2019) have significantly contributed to advancements in NLP, especially text classification. However, despite their *state-of-art* performance, their complexity and *black-box* nature obscure the decision-making process and hinder their interpretability. Now

more and more real-world applications also desire classification models built with LMs to be interpretable, since it allows end-users to comprehend the decision-making process, fostering trust and encouraging adoption. To address this need, there is growing interest in enhancing the interpretability of text classification models built with LMs.

Recent efforts have focused on redesigning neural networks to be inherently interpretable, based on the classic framework of prototypical learning (Datta and Kibler, 1995). These models actively learn prototype vectors through training, which are representative cases from previous observations, to explain decisions more intuitively. This methodology is first introduced into image domain by Li et al. (2018) and Chen et al. (2019), and then applied in text classification field with different structure variations (Ming et al., 2019; Hong et al., 2020; Pluciński et al., 2021; Das et al., 2022). However, despite this *white-box* framework being efficient in training and improving the model’s intrinsic interpretability, there is still some performance gap compared with the original black-box models.

On the other hand, there exist strategies such as graph attention network (GAT) (Velickovic et al., 2017), that is known for its ability to capture the importance of neighboring nodes in a graph through attention mechanisms, enabling more effective and expressive feature representation learning for each node. This approach allows GAT to dynamically assign different weights to different neighbors, which enhances the model’s performance on various graph-based tasks (Wang et al., 2019; Xie et al., 2020; Bhatti et al., 2023). This inspires us to analogize learning the relatedness between the input vector and the prototype vectors as constructing edges and learning their weights between nodes.

To address the challenge of the performance gap, we propose a novel *white-box* Multi-head Graph Attention-based Prototypical Network (GAProtoNet) designed to explain the decisions of

text classification models built with LM encoders. Our approach incorporates a prototype layer on top of a fine-tuned LM and utilizes multi-head graph attention (Velickovic et al., 2017) to efficiently learn an interpretable prototypical representation by selectively constructing edges between encoded representations and their neighboring prototypes. In the reference time, the decision is solely based on a linear combination of prototypes weighted by the attention scores assigned by attention heads. Our model archives superior performance without sacrificing the accuracy of the original LM while facilitating transparent decision-making progress. The novelty of our work lies in being the first to represent prototypes and input vectors as nodes within graphs, and actively train prototype vectors utilizing edges constructed by graph attention for interpretable text classification. The contribution of our work can be summarized as follows:

- We propose a new prototypical framework that leverages multi-head graph attention to selectively construct edges between input and prototypes which indicates relatedness. Our model is inherently interpretable and the performance is superior without sacrificing the accuracy of the original models.
- We did extensive comparison experiments with variations of prototype-based networks on five public benchmark datasets, including binary, four-label, ten-label classification. Our approach achieves the best accuracy and F1.
- We evaluate the interpretability of GProtoNet through various design criteria and demonstrate that the explanations provided by GProtoNet are of high quality.

The next section describes our proposed approach. §3 presents experiments, while §4 presents an analysis of the interpretability exhibited in the proposed approach. We describe related work in §5. We conclude in §6 where we reiterate our contribution. Limitations and future work are pointed out in §7.

2 Graph-Attention ProtoNet

Figure 1 illustrates the overall architecture of GProtoNet, which consists of three major components: (1) *Text Embedding based on Language Models (LMs)*, responsible for converting input text into high-dimensional vectors that capture its semantic information; (2) *Prototype Layer*, responsible for forming multiple typical prototype vectors

that encapsulate distinct semantic aspects of the input text as well as enhancing the interpretability of the model; and (3) *Graph Attention Mechanism*, which efficiently learn the relatedness between the text embedding vector and its neighboring prototype vectors, which are later used to form a prototypical representation of the text and passed to output layer. Detailed descriptions for each component are presented in the following subsections.

2.1 Text Embedding encoded by Language Models (LMs)

Given the remarkable performance of pre-trained language models RoBERTa, XLNet, and DistilBERT on a wide range of NLP tasks, we leverage their capabilities for efficient text embeddings. These models have been trained on vast text datasets and can be fine-tuned for specific downstream tasks. By utilizing these pre-trained language models, we embed input text into high-dimensional vectors, which serve as representations that capture the deep semantic information of the text:

$$s = \text{LM}(x) \quad (1)$$

where x denotes the input text. s is the semantic information representation vector.

2.2 Prototype Layer

In the prototype layer, we define M prototype vectors $P = \{p_j\}_{j=1}^M$ that represent typical features in the vector space of training data. For each semantic representation vector passed from the text embedding layer, we compute the attention score vector indicating the relatedness between itself and all prototype vectors with a graph attention model introduced in §2.3.

Note that these prototype vectors are randomly initialized. They are learned through active training with loss defined in §2.5. Their representativeness is improved by updating the weights during each training epoch, resulting in an inherently interpretable classifier.

2.3 Graph Attention Mechanism

The graph attention mechanism in our model captures the relatedness between the embedding text vector and the prototype vectors, allowing for the derivation of prototypical representations for the embedding vectors mentioned in §2.2. To effectively capture diverse semantic aspects and intricate patterns in the data, we utilized a multi-head

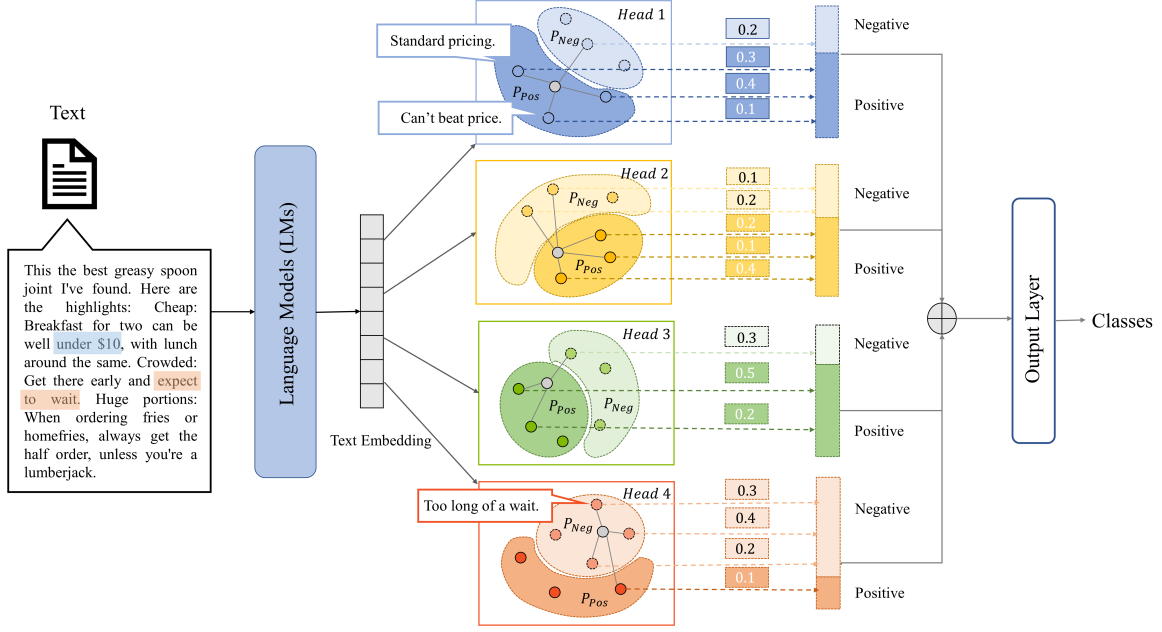


Figure 1: **Overview of the GAProtoNet Architecture:** The GAProtoNet architecture consists of three primary components: a Text Embedding Layer using LMs, a Prototype layer, and Graph Attention. The text embeddings are linearly transformed to produce query vectors. Each attention head will construct a graph based on the attention score between the query vectors and predefined prototypes. As shown in the figure, different heads will activate prototypes in different semantic aspects, assigning different attention score to negative ones and positive ones. An interpretable prototypical representation will be formed with a linear combination of all prototypes weighted by the attention score and sent to the output layer for classification.

approach to comprehensively model this relatedness. The query nodes in the graph represent the input text embedding vector while the key nodes represent the prototype vectors. Edges and weights are constructed based on the following approach:

Single-Head Linear Transformation For each text embedding input s , we utilize matrix \mathbf{W}^q to transform it into query vector q . For each prototype p_j , we utilized matrix \mathbf{W}^k to transform it into k_j :

$$q = \mathbf{W}^q s, k_j = \mathbf{W}^k p_j \quad (2)$$

Multi-Head Linear Transformation Now instead of using one single set of weight, we utilize a set of learnable matrix linear transformations $\{\mathbf{W}_i^q\}_{i=1}^H$ and $\{\mathbf{W}_i^k\}_{i=1}^H$ to generate a series of query nodes $\{q_i\}_{i=1}^H$ and prototypes $\{k_{ij}\}_{i=1}^H$ with H representing the number of attention heads. For each head i , we have:

$$q_i = \mathbf{W}_i^q s, k_{ij} = \mathbf{W}_i^k p_j \quad (3)$$

Attention Score Computation After the transformation with each head i , we calculate the similarity between the node q_i and the nodes $\{k_{ij}\}_{j=1}^M$ based on the dot product:

$$\text{sim}(q_i, k_{ij}) = \text{dot}(q_i, k_{ij}) / d_k \quad (4)$$

where d_k denotes the dimension of key node vector k .

The attention scores are then derived by applying the sigmoid activation function σ to the similarity vector:

$$\alpha_{ij} = \sigma(\text{sim}(q_i, k_{ij})) \quad (5)$$

This attention score represents the weight of each prototype p in shaping the decision-making process. As shown in Figure 1, our experimental results demonstrate that each attention head tends to activate prototypes corresponding to distinct semantic aspects. Moreover, the attention scores assigned to negative and positive prototypes suggest a bias towards a negative or positive final classification, respectively.

Graph Edge Construction For each attention head i , a graph is constructed by activating edges between the transformed prototype nodes k and text embedding node q when attention scores exceed a certain threshold τ . For each attention head i , the edge between the transformed prototype q

and transformed j -th prototype node \mathbf{k}_j can be expressed as:

$$E = \{(\mathbf{q}_i, \mathbf{k}_{ij}) \mid \alpha_{ij} > \tau\} \quad (6)$$

Interpretable Prototypical Representation The prototypical representation vector for the text embedding vector \mathbf{s} is formed by computing a linear combination of the neighboring prototype nodes from all attention heads weighted by the attention score.

For each attention head i , we first normalize the attention scores for neighboring transformed prototype nodes of \mathbf{q}_i . Thus, for each prototype node \mathbf{k}_j under attention head i , its normalized score is:

$$\gamma_{ij} = \frac{\alpha_{ij}}{\sum_{k \in \mathcal{N}(\mathbf{q}_i)} \alpha_{ik}} \quad (7)$$

The prototypical representation vector \mathbf{r} for \mathbf{s} under attention head i is computed as:

$$\mathbf{r}_i = \sum_{j \in \mathcal{N}(\mathbf{q}_i)} \gamma_{ij} \mathbf{k}_{ij} \quad (8)$$

where $\mathcal{N}(\mathbf{q}_i)$ is the set of neighboring nodes for node \mathbf{q} under attention head i constructed under the definition of E in 6.

Output Layer The interpretable prototypical representations obtained from all heads are concatenated to form a single vector. This vector is then utilized as input for the output layer to perform text classification.

2.4 Prototype Projection

To understand the natural language meaning of each prototype vector \mathbf{p}_j , we match each prototype with the sample text embedding vector \mathbf{s}_j in the dataset \mathcal{D} that has the highest similarity, and assign that sample \mathbf{x}_j as the prototype text. Let \mathcal{D} represent the training dataset, then:

$$\text{Text of } \mathbf{p}_j \leftarrow \arg \max_{\mathbf{x}_j \in \mathcal{D}} \text{sim}(\mathbf{s}_j, \mathbf{p}_j) \quad (9)$$

where $\text{sim}(\mathbf{s}_j, \mathbf{p}_j)$ denotes the similarity measure between the sample text \mathbf{s}_j and the prototype vector \mathbf{p}_j .

During text classification tasks, we observe the graph model formed between embedded text and prototypes, as well as the edge weights, to interpret the classification process.

2.5 Training Objective

In this study, our designed composite loss function consists of three key components: *Accuracy Loss*, *Proximity Loss*, and *Diversity Loss*.

Accuracy Loss We employ the cross entropy loss to guide the model training. This loss measures the discrepancy between the predicted probability distribution and the true labels, optimizing the model to assign higher probabilities to the correct classes. Specifically, the accuracy loss is defined as:

$$\mathcal{L}_{Acc} = - \sum_{i=1}^N \sum_{j=1}^C y_{ij} \log(\hat{y}_{ij}), \quad (10)$$

where N is the number of samples, while C is the number of classes. y_{ij} is the true one-hot label of sample i for class j , while \hat{y}_{ij} is the predicted probabilities.

Proximity Loss To ensure that each prototype can be projected onto a similar sample in the training data, we introduce the Proximity Loss. This loss measures the distance between prototypes and data points, penalizing prototype-sample pairs that are far apart. We use Euclidean distance as the distance metric and calculate the minimum distance between each prototype and all samples. The average of the minimum distances across all prototypes is then taken as the Proximity Loss:

$$\mathcal{L}_{Prox} = \frac{1}{M} \sum_{j=1}^M \min_i \|\mathbf{p}_j - \mathbf{s}_i\|_2^2 \quad (11)$$

where M is the number of prototypes, \mathbf{p}_j is the vector representation of the j -th prototype, and \mathbf{s}_i is the embedded vector representation of the i -th sample in the training set.

Diversity Loss To encourage diversity among prototypes and avoid redundancy, we introduce the Diversity Loss. This loss aims to encourage prototypes to be distributed as diversely as possible in the feature space. We achieve this by penalizing the average distance between all pairs of prototypes:

$$\mathcal{L}_{Div} = - \frac{1}{M(M-1)} \sum_{j=1}^M \sum_{k \neq j} \|\mathbf{p}_j - \mathbf{p}_k\|_2 \quad (12)$$

Composite Loss Function By combining the above three loss functions, we obtain the composite loss function:

$$\mathcal{L} = \lambda_1 \cdot \mathcal{L}_{Acc} + \lambda_2 \cdot \mathcal{L}_{Prox} + \lambda_3 \cdot \mathcal{L}_{Div} \quad (13)$$

where λ_1 , λ_2 , and λ_3 are hyperparameters that balance the weights between different loss terms. These hyperparameters can be determined experimentally to optimize the model’s performance.

3 Performance Experiments

In this section are discussed the datasets used in the experiments, and the models used as black-box baselines, prototypical baselines, and the variations of the proposed approach. The overall hypothesis is that the proposed approach performs on par with black-box and outperforms prototype baselines in the dataset used for text classification.

3.1 Datasets and Metric

We evaluate our approach on three binary public benchmark datasets: **Hotel Reviews**¹ prepared by (Hong et al., 2020), **IMDb**² (Maas et al., 2011) and **Yelp Polarity Reviews**. All of them are balanced datasets. To demonstrate our model’s performance on more challenging tasks, we also evaluate it on **Tweet** (Mohammad et al., 2018) and **Yahoo** (Zhang et al., 2015), which are 4-class and 10-class datasets separately.

Metric We use accuracy, recall and F1-Score as metrics to evaluate models’ performance. For each model, we run 5 times and calculate the average as the final reported results.

3.2 Models and Settings

We select three *black-box* LM as baselines and four different prototype-based variations for comparison experiments. We also test two variations of GAProtoNet, one with a single attention head and another with attention heads = 4. We train each model with one single NVIDIA GTX 3090 or GTX 4090. We use Adam as the optimizer and the learning rate is $1e - 4$. Due to the limitation of GPU RAM, we choose a batch size of 4 and an accumulated gradient step of 64.

LM Baselines We select the following pre-trained *black-box* language model without prototypes as the powerful baselines: **DistilBERT** (Sanh et al., 2019), **RoBERTa-large** (Liu et al., 2019) and **XLNet** (Yang et al., 2019). For all the three LMs, we use the same simple MLP (multilayer perceptron) with two hidden layers over the output of the CLS token from the last hidden layer of the LM encoder for classification.

¹<https://www.kaggle.com/datafiniti/hotel-reviews>.

²<https://huggingface.co/datasets/stanfordnlp/imdb>.

Prototype-based Networks Variations We select the following 4 prototypical network variations for comprehensive comparison experiments against our GAProtoNet: **ProSeNet** (Ming et al., 2019), **ProtoryNet** (Hong et al., 2020), **ProtoCNN** (Pluciński et al., 2021) and **ProtoTEX** (Das et al., 2022).

3.3 Experiment Result

Binary Classification The first three section in Table 1 summarizes our evaluation results on binary classification. We observe that GAProtoNet either closely matches or exceeds the performance of its baseline LLM across our experiments, supporting GAProtoNet’s interpretability is not achieved at the cost of performance. Within four types of prototypical variations, ProtoTEX and ProtoryNet have the best performance for accuracy, recall, and F1 for all three datasets, but GAProtoNet still improves 3%-6% upon their performance. This indicates that leveraging graph attention to construct edges between input embeddings and prototypes could be more efficient than heuristic distance measurement. For the variations of GAProtoNet, multi-head GAProtoNet have better performance compared with single-head, with the RoBERTa as the encoder showing the best results. This could be due to that the multi-head attention enables the network to capture semantic meanings from different perspectives, thus potentially improving the network’s performance.

Extend to Multi-label Classification The last two section in Table 1 summarizes our evaluation results on multi-label classification. For the 4-class dataset, our model is the best overall. For the 10-class dataset, our model is better than all other prototype-based models and comparable to black-box models. This confirms the robustness of our model observed in binary classification and also shows that our model produces accuracy comparable or better than black-box model even when the task is more challenging. Comparing our model against other prototype-based architectures, the performance is better in all the experiments done so far.

The consistent results across five datasets support our hypothesis in the beginning that GAProtoNet can perform on par with black-box and outperform prototype baselines. The performance improvement demonstrated by GAProtoNet indicates graph attention for prototypical networks can

	Hotel			IMDb			Yelp			Tweet			Yahoo		
	Acc.	Rec.	F1	Acc.	Rec.	F1	Acc.	Rec.	F1	Acc.	Rec.	F1	Acc.	Rec.	F1
DistilBERT	97.49	98.22	97.57	92.91	93.53	92.93	97.25	97.51	97.43	76.84	76.84	76.89	71.17	71.17	71.97
XLNet	98.73	98.21	98.77	95.79	96.71	95.82	98.12	98.25	98.40	79.52	79.52	79.26	75.80	75.80	75.45
RoBERTa-large	98.97	99.33	99.00	96.38	97.08	96.38	98.18	98.57	98.21	81.63	81.70	81.63	75.46	75.46	74.94
ProSeNet	92.17	93.72	93.06	86.69	87.84	87.53	92.58	91.24	92.17	73.24	72.35	73.54	68.17	68.23	68.19
ProtoryNet	94.45	95.59	94.55	88.21	87.14	88.45	<u>95.32</u>	94.08	<u>94.62</u>	72.90	72.85	72.28	<u>70.62</u>	69.03	<u>70.58</u>
ProtoCNN	94.32	95.49	94.34	89.32	91.03	89.65	87.25	88.47	89.98	64.32	63.51	63.36	<u>66.72</u>	65.98	<u>64.21</u>
ProtoTEX	<u>95.90</u>	<u>96.37</u>	<u>96.23</u>	<u>92.39</u>	<u>93.21</u>	<u>91.36</u>	93.07	<u>94.40</u>	94.26	<u>75.28</u>	<u>75.14</u>	<u>75.20</u>	70.32	<u>70.33</u>	70.33
SINGLE-HEAD GAPROTO NET VARIATIONS															
DistilBERT + SG	97.14	96.67	97.20	92.79	93.73	92.82	97.38	97.76	97.40	75.72	75.72	76.00	70.67	70.67	71.35
XLNet + SG	98.62	98.88	99.11	95.36	94.65	95.73	97.65	98.70	98.21	79.80	79.80	79.76	73.77	73.77	73.65
RoBERTa + SG	98.94	98.67	98.89	96.35	96.97	96.23	98.14	98.11	98.17	81.42	81.40	81.42	72.23	72.23	71.16
MULTI-HEAD GAPROTO NET VARIATIONS															
DistilBERT + MG	97.94	97.78	98.00	92.79	94.06	92.86	97.40	96.93	97.43	75.43	75.40	75.40	71.40	71.40	72.00
XLNet + MG	99.08	99.33	99.11	95.96	95.96	95.96	98.19	98.69	98.23	80.64	80.64	80.48	74.12	74.12	74.74
RoBERTa + MG	99.09	99.56	99.12	96.53	96.85	96.53	98.27	98.07	98.39	81.84	81.81	81.85	73.79	73.79	73.95
Improvement (%)	3.19	3.19	2.89	4.14	4.46	5.17	2.95	4.29	3.77	6.56	6.67	6.65	3.50	3.79	4.16

Table 1: Accuracy, recall and F1 scores for baseline models and GAProtoNet variations across three datasets. 'SG' denotes single-head GAProtoNet, while 'MG' represents multi-head GAProtoNet. Bolded values indicate the highest scores in each column across all sections, while underlined values highlight the best performance within the prototypical baseline variation section. The percentage improvement reflects the difference between the highest performance of GAProtoNet and that of other prototype-based variations.

achieve better overall performance, thus validating our contribution.

4 Interpretability Analysis

In this section, we analyze our model's interpretability from two perspectives. First, we use case studies to demonstrate how prototypes and their surrounding edges to the input node, constructed with graph attention, are used to explain the model's decision-making process. Second, we evaluate the quality of our prototypes. We demonstrate their representativeness through visualization, showing that they are distributed dispersedly within the training data space. Additionally, we assess the distinctiveness of the prototypes by examining how it varies with the number of prototypes and analyzing its impact on the model's prediction performance.

4.1 Case study

Figure 3 illustrates an example of using multi-graph attention heads and prototypes to explain a classification result. The input text is a Yelp review about a restaurant. In this instance, eleven out of twenty prototypes are activated by four attention heads, indicating edges constructed between them and the input text. Figure 2 illustrates the activation of prototypes using attention Head 3 as an exam-

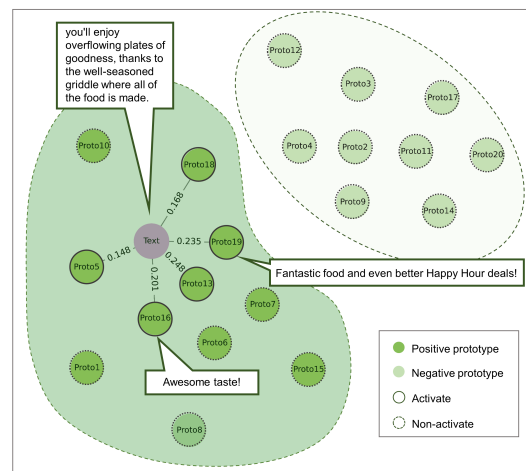


Figure 2: Prototype activation under Attention Head 3.

ple. Specifically, five positive prototypes related to food quality are activated, and edges are constructed between the input text and the activated prototypes, with the attention scores representing the edge weights. This activation reflects the positive sentiment towards food quality expressed in the original text.

We observe that different prototypes, representing various aspects of the restaurant, are activated by different attention heads. Head 1 activates prototypes related to service and price, Head 2 to service, Head 3 to food, and Head 4 to waiting time. In this

Text: This the best greasy spoon joint I've found. If you can stomach the **wait** (no reservations, naturally), you'll enjoy overflowing plates of goodness, thanks to the well-seasoned griddle where all of the food is made. Here are the highlights: **Cheap:** Breakfast for two can be well under \$10, with lunch around the same. **Crowded:** Get there early and expect to **wait**. **Huge portions:** When ordering **fries or homefries**, always get the half order, unless you're a lumberjack. My favorite item: **hot sausage sandwich** on thick Italian toast, with cheese and lettuce.

Prototype	Head 1	Head 2	Head 3	Head 4
1: "The plates were big and beautiful but we experienced poor service (part of our meal was forgotten)."	✓	✓		
4: "You can get a breakfast comparable to Denny's but the services is slow. Over 30 min is too long of a wait for eggs, toast and hash browns although they did come out still hot".				✓
5: "The service was great and the food was even better. Light but filling, tons of flavor and so fresh!"		✓	✓	
6: " Good food at a good price . Pretty Authentic Hawaiian BBQ in AZ and Cali. I get the #4 combo."	✓	✓		
7: "Certainly can't beat the price , the curry that came with the roti canai was wonderful."	✓	✓		
9: "They told us a 10 minute wait . An hour later we left. Unbelievable. I will never be back."				✓
11: "Horrible experience! We waited for an hour for our food while tables seated after us were all served."				✓
13: " Amazing food and service! The black cod lettuce cups and the yellowtail sashimi were awesome."		✓	✓	
16: "Cool place, stopped and got an iced coffee. Standard pricing , awesome taste , and good atmosphere"	✓		✓	
18: "Great. Totally lived up to the hype. Delicious , and worth the wait in line ."			✓	✓
19: " Fantastic food and even better Happy Hour deals! Awesome atmosphere and late night snacks."		✓	✓	

Figure 3: Using a Yelp review as input, the prototypes activated by various graph attention heads are analyzed. The review’s label and prediction are both positive. For each prototype, an attention head is checked if it constructs an edge between the prototype and the input. Different aspects are highlighted in distinct colors: blue for price, yellow for service, red for wait time, and green for food quality.




Figure 4: Prototype distribution within training data space for three datasets.

text input, all different aspects of the text are captured by each attention head and then construct the edge between the input and the prototypes. Interestingly, prototypes related to service are activated by Head 2, even though the input text does not explicitly mention service. This suggests that the model might infer a positive attitude towards service based on implicit cues in the text. It also indicates that the aspects activated by different attention heads can overlap rather than being fully distinct.

4.2 Prototype quality and performance

Prototype Distribution Visualization Since prototypes are used to interpret the model, they should capture as many representative features from the data as possible, indicating an even distribution within the training data space. A straightforward way to verify this is through visualization. Since our prototype vectors and training data are high-

dimension vectors, we first use t-SNE (t-distributed Stochastic Neighbor Embedding) (Van der Maaten and Hinton, 2008) to do dimension deduction before the visualization. This method is effective at preserving the local structure of the data so the distribution pattern won’t change when projected from higher dimension space to lower dimension space. The visualization of prototype vectors for a multi-head GProtoNet with the prototype number of $k = 20$ is shown in Figure 4. The orange dots represent prototypes while the blue dots represent training data. We can see that the prototype vectors are evenly distributed within the training data space and test results show that those prototype vectors are of high orthogonality. This indicates that the space formed by these vectors can cover most of the data points so we can use limited prototypes to represent any point in the training data.




Figure 5: Prediction accuracy and the percentage of unique prototypes when trained with different numbers of prototypes

Prototype Distinguishness and Model Performance

Since the prototypes are used both for interpreting results and making decisions, we are interested in how the hyperparameter k of their number affects the model’s interpretability and classification performance. For interpretability, we measure the number of distinguished prototypes. *Distinguishness* indicates that if more than one prototype vector projects onto the same sample data point, it is counted only once. For classification performance, we measure the accuracy. We conduct experiments varying the number of prototypes k from 10 to 40 on a multi-head GProNet with RoBERTa-large as the encoder across the three datasets. The percentage of distinguished prototypes and classification accuracy is shown in Figure 5. We observe the accuracy achieves the highest when $k = 20$ across three datasets then drops with k increasing. The corresponding trend for the percentage of distinguished vectors is also dropping from $k = 20$ to $k = 40$. We reason that increasing k can only improve the model’s expressiveness until a certain point, in our case $k = 20$. After the point, prototypes will lose distinguishness and therefore hurts model’s classification performance. Instead of observing a trade-off between interpretability and accuracy, our experiments show that they are positively related with each other.

5 Related Work

Li et al. (2018) and Chen et al. (2019) introduce prototype-based architecture into interpretable image classification by proposing a prototype layer, where prototypes are randomly initialized and are made meaningful through training. Ming et al. (2019) brought this approach into NLP domain by adding a sequence encoder before the prototype

layer. Concept representation learning is also applied in enhancing the interpretability of the structure (Jin et al., 2024c,b,a) while others improve interpretability by extracting minimal sentiment triplets (Sun et al., 2024b,a). However, despite this intrinsically interpretable model demonstrating compelling results, there are still some performance gaps compared to the original black-box model.

In the NLP domain, researchers attempt to minimize the performance gap by trying different structure variations. Hong et al. (2020) used Sentence Universal Encoder (Cer et al., 2018) and added an LSTM layer (Hochreiter and Schmidhuber, 1997) between the prototype layer and the output layer to better capture the patterns in the trajectory of prototypes. Pluciński et al. (2021) propose a new structure that operates on the prototypes in the form of phrases. Das et al. (2022) applied a new distance measurement metric and further minimized the accuracy gap but they only compared with the black-box models and lack of comparison with other prototype-based variations.

All the above approaches utilize variations of heuristic distance metric (e.g. cosine similarity or Euclidean distance) to calculate the distance vector or matrix between the input vectors and prototypes, which serves as the sole input to the output layer. We hypothesize that only passing a heuristic distance metric may omit essential upstream information, thereby impairing performance. This motivates us to improve the current structure by calculating relatedness with the attention mechanism, which could potentially preserve more information (Vaswani et al., 2017).

Considering all the limitations of the current approaches, there is a need for a new prototypical network that can further minimize the performance gap compared with the original black-box model and at the same time prototypes can be effectively trained through active learning.

6 Conclusion

We contribute an interpretable prototypical deep learning architecture that is advanced with graph attention on text classification. Our experiments show our approach outperforms both other interpretable and black-box architectures on benchmark datasets. We also conduct a comprehensive analysis to show good interpretability of our approach.

7 Limitation

To demonstrate these prototypes can support user explainability in addition to making the architecture interpretable, we will include user study in our future work.

References

- Uzair Aslam Bhatti, Mengxing Huang, Harold Neira-Molina, Shah Marjan, Mehmood Baryalai, Hao Tang, Guilu Wu, and Sibghat Ullah Bazai. 2023. Mffcg—multi feature fusion for hyperspectral image classification using graph attention network. *Expert Systems with Applications*, 229:120496.
- Daniel Cer, Yinfei Yang, Sheng-yi Kong, Nan Hua, Nicole Limtiaco, Rhomni St John, Noah Constant, Mario Guajardo-Cespedes, Steve Yuan, Chris Tar, et al. 2018. Universal sentence encoder. *arXiv preprint arXiv:1803.11175*.
- Chaofan Chen, Oscar Li, Daniel Tao, Alina Barnett, Cynthia Rudin, and Jonathan K Su. 2019. This looks like that: deep learning for interpretable image recognition. *Advances in neural information processing systems*, 32.
- Anubrata Das, Chitrang Gupta, Venelin Kovatchev, Matthew Lease, and Junyi Jessy Li. 2022. Prototex: Explaining model decisions with prototype tensors. *arXiv preprint arXiv:2204.05426*.
- Piew Datta and Dennis Kibler. 1995. Learning prototypical concept descriptions. In *Machine Learning Proceedings 1995*, pages 158–166. Elsevier.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation*, 9(8):1735–1780.
- Dat Hong, Tong Wang, and Stephen S Baek. 2020. Protorynet-interpretable text classification via prototype trajectories. *arXiv preprint arXiv:2007.01777*.
- Mingyu Jin, Haochen Xue, Zhenting Wang, Boming Kang, Ruosong Ye, Kaixiong Zhou, Mengnan Du, and Yongfeng Zhang. 2024a. Prollm: Protein chain-of-thoughts enhanced llm for protein-protein interaction prediction. *arXiv e-prints*, pages arXiv–2405.
- Mingyu Jin, Qinkai Yu, Jingyuan Huang, Qingcheng Zeng, Zhenting Wang, Wenyue Hua, Haiyan Zhao, Kai Mei, Yanda Meng, Kaize Ding, et al. 2024b. Exploring concept depth: How large language models acquire knowledge at different layers? *arXiv preprint arXiv:2404.07066*.
- Mingyu Jin, Qinkai Yu, Dong Shu, Haiyan Zhao, Wenyue Hua, Yanda Meng, Yongfeng Zhang, and Mengnan Du. 2024c. [The impact of reasoning step length on large language models](#). In *Findings of the Association for Computational Linguistics ACL 2024*, pages 1830–1842, Bangkok, Thailand and virtual meeting. Association for Computational Linguistics.
- Oscar Li, Hao Liu, Chaofan Chen, and Cynthia Rudin. 2018. Deep learning for case-based reasoning through prototypes: A neural network that explains its predictions. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. [Roberta: A robustly optimized bert pretraining approach](#). *Preprint*, arXiv:1907.11692.
- Andrew L. Maas, Raymond E. Daly, Peter T. Pham, Dan Huang, Andrew Y. Ng, and Christopher Potts. 2011. [Learning word vectors for sentiment analysis](#). In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pages 142–150, Portland, Oregon, USA. Association for Computational Linguistics.
- Yao Ming, Panpan Xu, Huamin Qu, and Liu Ren. 2019. Interpretable and steerable sequence learning via prototypes. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 903–913.
- Saif Mohammad, Felipe Bravo-Marquez, Mohammad Salameh, and Svetlana Kiritchenko. 2018. [SemEval-2018 task 1: Affect in tweets](#). In *Proceedings of the 12th International Workshop on Semantic Evaluation*, pages 1–17, New Orleans, Louisiana. Association for Computational Linguistics.
- Kamil Pluciński, Mateusz Lango, and Jerzy Stefanowski. 2021. Prototypical convolutional neural network for a phrase-based explanation of sentiment classification. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 457–472. Springer.
- Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. 2019. Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter. *arXiv preprint arXiv:1910.01108*.
- Qiao Sun, Jiexin Xie, Nanyang Ye, Qinying Gu, and Shijie Guo. 2024a. [Enhancing nursing and elderly care with large language models: An ai-driven framework](#). *Preprint*, arXiv:2412.09946.
- Qiao Sun, Liujia Yang, Minghao Ma, Nanyang Ye, and Qinying Gu. 2024b. [MiniConGTS: A near ultimate minimalist contrastive grid tagging scheme for aspect sentiment triplet extraction](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 2817–2834, Miami, Florida, USA. Association for Computational Linguistics.

- Laurens Van der Maaten and Geoffrey Hinton. 2008. Visualizing data using t-sne. *Journal of machine learning research*, 9(11).
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- Petar Velickovic, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Lio, Yoshua Bengio, et al. 2017. Graph attention networks. *stat*, 1050(20):10–48550.
- Xiang Wang, Xiangnan He, Yixin Cao, Meng Liu, and Tat-Seng Chua. 2019. Kgat: Knowledge graph attention network for recommendation. In *Proceedings of the 25th ACM SIGKDD international conference on knowledge discovery & data mining*, pages 950–958.
- Yu Xie, Yuanqiao Zhang, Maoguo Gong, Zedong Tang, and Chao Han. 2020. Mgat: Multi-view graph attention networks. *Neural Networks*, 132:180–189.
- Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Russ R Salakhutdinov, and Quoc V Le. 2019. Xlnet: Generalized autoregressive pretraining for language understanding. *Advances in neural information processing systems*, 32.
- Xiang Zhang, Junbo Zhao, and Yann LeCun. 2015. Character-level convolutional networks for text classification. *Advances in neural information processing systems*, 28.

A Projected Prototypes

Prototype
1: "The plates were big and beautiful but we experienced poor service (part of our meal was forgotten).
2: "Thieves. Coffee prices are close to double than regular. 6 \$ for a bottle of water. Should I say more?"
3: "Slowest fast food I've ever had. I've been waiting so long I'm writing this review before I even get the food."
4: "You can get a breakfast comparable to Denny's but the services is slow. Over 30 min is too long of a wait for eggs, toast and hash browns although they did come out still hot".
5: "The service was great and the food was even better. Light but filling, tons of flavor and so fresh!"
6: "Good food at a good price. Pretty Authentic Hawaiian BBQ in AZ and Cali. I get the #4 combo."
7: "Certainly can't beat the price, the curry that came with the roti canai was wonderful."
8: "Food was decent. Atmosphere was not as expected, more of a unfriendly bar than a restaurant."
9: "They told us a 10 minute wait. An hour later we left. Unbelievable. I will never be back."
10: "Great attention to detail. Full service from the welcome at door, a complementary water bottle, to the getting the works for a pedicure!!
11: "Horrible experience! We waited for an hour for our food while tables seated after us were all served."
12: "The food was little more than decent and the staff didn't seem to care about their clients. I was never this disappointed."
13: "Amazing food and service! The black cod lettuce cups and the yellowtail sashimi were awesome."
14: "Poor quality food and little selection. Got a rude server! Don't waste your time here!!"
15: "Average place, slightly overpriced- would go back, but wouldn't seek it again."
16: "Cool place, stopped and got an iced coffee. Standard pricing, awesome taste, and good atmosphere"
17: "Went here for lunch today to try something new. The food is rotten. The juice is warm. No wonder it was so cheap."
18: "Great. Totally lived up to the hype. Delicious, and worth the wait in line."
19: "Fantastic food and even better Happy Hour deals! Awesome atmosphere and late night snacks."
20: "Overpriced, scallops were rubber. The circus at the table was entertaining though."

Figure 6: This is the 20 prototypes after projecting them with the nearest sample.