TOWARDS IMPROVED SENTENCE REPRESENTATIONS USING TOKEN GRAPHS

Anonymous authors

Paper under double-blind review

ABSTRACT

Obtaining a single-vector representation from a Large Language Model's (LLM) token-level outputs is a critical step for nearly all sentence-level tasks. However, standard pooling methods like mean or max aggregation treat tokens as an independent set, discarding the rich relational structure captured by the model's self-attention layers and making them susceptible to signal dilution. To address this, we introduce GLOT, a lightweight, structure-aware pooling module that reframes pooling as relational learning followed by aggregation. Operating on the outputs of a frozen LLM, GLOT first constructs a latent token-similarity graph, then refines token representations with a graph neural network, and finally aggregates them using a readout layer. Experimentally, our approach is remarkably robust and efficient: on a diagnostic stress test where 90% of tokens are random distractors, GLOT maintains over 97% accuracy while baseline methods collapse. Furthermore, it achieves state-of-the-art performance on benchmarks like GLUE and MTEB with 20x fewer trainable parameters and speeds up the training time by over 100x compared with parameter-efficient fine-tuning methods. Supported by a theoretical analysis of its expressive power, our work shows that learning over token graphs is a powerful paradigm for the efficient adaptation of frozen LLMs.

1 Introduction

Large Language Models (LLMs) (Raffel et al., 2020; Lewis et al., 2020; Brown et al., 2020; Touvron et al., 2023; Jiang et al., 2023) produce a sequence of token-level hidden states, yet many downstream applications require a single vector embedding to represent an entire sentence or document. Therefore, the process by which a sentence and its tokens' hidden states are collapsed into a single vector representation is critical. Standard practices typically rely on simple heuristics such as *mean*, *max*, or using a dedicated *[CLS]* token. While these pre-defined approaches have been dominant in the literature (Devlin et al., 2019; Liu et al., 2019; Reimers & Gurevych, 2019; Gao et al., 2021; Arora et al., 2017; Wang et al., 2024), they can also be vulnerable when only a small subset of tokens carries task-relevant signal amid many distractors, as has been recently studied in Brothers (2025).

Although Transformers (Vaswani et al., 2017) inherently model token interactions through self-attention, standard sentence-level representation techniques discard this rich relational structure, treating hidden states as an independent set of vectors. Indeed, this shortcoming was recently studied for Vision-Transformers (Dosovitskiy et al., 2021) in Brothers (2025), who proposed to learn a multilayer-perceptron (MLP)-based token scoring function. However, while this approach may correctly up-weight the word "good", it may fail to capture the effect of its negation with the word "not". This challenge is particularly acute for *decoder-only LMs* (e.g., GPT (Radford et al., 2019; Brown et al., 2020) or LLaMA (Touvron et al., 2023)), whose causal attention mechanism optimizes hidden states for next-token prediction rather than holistic sentence representation (Radford et al., 2019; Brown et al., 2020).

Prior work shows that LLM token vectors have a strong directional bias: many of them point in similar directions, and seemingly unrelated words have embeddings with high similarity (Ethayarajh, 2019; Li et al., 2020). Therefore, sentence-level representations built on isolated tokens may be unreliable for semantic understanding tasks. While these shortcomings can be addressed by fine-tuning the entire model on downstream tasks, this approach is often computationally prohibitive for billion-parameter

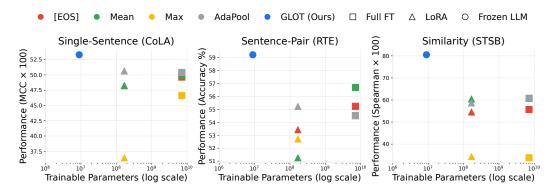


Figure 1: Fine-tuning large language models for sentence embeddings is computationally expensive. Our pooling method, GLOT, constructs a latent token-similarity graph from the outputs of a frozen model. It then refines token representations with a graph neural network before aggregation. This technique enables decoder-only models (like Mistral-7B), typically optimized for next-token prediction, to produce powerful sentence-level representations without requiring any fine-tuning.

models (Lee et al., 2025; Gao et al., 2021). The immense cost of training, compounded by the need for extensive hyperparameter optimization, makes full fine-tuning impractical for many applications.

To bridge this gap, we reframe the problem of collapsing token hidden states into a sentence-level representation as *learning over token graphs*. To this end, we propose GLOT, a lightweight, structure-aware module that operates on the token hidden states produced by LLMs to obtain a sentence-level representation. Specifically, as illustrated in Figure 2, GLOT does the following: (i) constructs a token-similarity graph from the LLM hidden states, (ii) propagates information across the graph using a graph neural network, and (iii) aggregates the refined token representations using a readout mechanism. The LLM backbone remains entirely frozen; only the GNN module and a task-specific head are trained. This lightweight approach maintains a remarkably small memory footprint while equipping decoder-only LMs to perform as powerful text embedding models.

Contributions. Our contributions are as follows:

- We introduce a new conceptualization of sentence-level representation from LLM hidden states; rather than framing it as direct information compression like existing techniques, we envision a relational learning approach via GNNs. In addition, our framework generalizes common pooling methods like mean, max, and [CLS] pooling.
- We present GLOT, a compact and parameter-efficient module that enhances the performance
 of both encoder- and decoder-only frozen backbones with 20x fewer trainable parameters
 and over 100x faster training time than LLM fine-tuning-based methods.
- We provide extensive empirical validation for GLOT. Our evaluation shows that GLOT consistently outperforms pre-defined pooling and learning-based methods, across a wide range of tasks, including the GLUE benchmark for language understanding (Wang et al., 2018), long-text classification on IMDB (Maas et al., 2011), and seven diverse tasks from the large-scale MTEB benchmark (Muennighoff et al., 2023). Crucially, we introduce a novel diagnostic stress test that confirms GLOT's superior robustness to signal dilution, a key failure mode for other methods.
- We provide a detailed analysis of our method's components, demonstrating how the graph construction impacts performance and quantifying its substantial computational efficiency over fine-tuning methods.

To ensure reproducibility, we will publish the code and pre-trained models upon acceptance and provide pseudo-code for our method in Appendix B.

2 RELATED WORK

The Compressive Paradigm of Sentence-Level Representation. To obtain sentence-level representations from LLMs, existing methods consider a compression problem: collapsing tokens' hidden states into a single vector. This paradigm usually encompasses pre-defined rules like *mean* or *max* selection, as well as learnable variants that learn token weights (Reimers & Gurevych, 2019; Gao et al., 2021; Xing et al., 2024; Lee et al., 2025; Brothers, 2025). While effective in some cases, these methods fundamentally discard relational structure. This can be derived from the theory of permutation-invariant functions on sets, as done in DeepSets (Zaheer et al., 2017), however, only looking at the tokens as completely independent items in a set does not paint the whole picture. As a result, these approaches implicitly assume the LLM has already embedded all necessary relational information. This assumption is often violated, especially for decoder-only models, which are optimized for next-token prediction rather than holistic sentence understanding (Radford et al., 2019; Brown et al., 2020). Indeed, recent work by Brothers (2025) shows such methods fail precisely because they compress *before* performing relational learning. Our work, GLOT, directly addresses this shortcoming by using advances from graph neural networks, which are also permutation invariant but can also encode relational information.

Graph-Based Representations in NLP. Graph Neural Networks (GNNs) are natural tools for relational learning; however, their prior applications in NLP differ from our problem of obtaining sentence-level representation using a frozen LLM. Many of these works use graphs to represent corpuslevel tasks and solve them using GNNs rather than producing sentence-level embeddings. For example, Yao et al. (2019) builds a single word-occurrence-based graph over the corpus for text classification, and Huang et al. (2019) extends this approach for online inference and reduced memory consumption. Recent works propose the usage of attention and diffusion dynamics (Liu et al., 2021) and interleaving GNN and Transformer layers for improved text classification performance. Other approaches differ in their architecture or output format. Late-interaction models like ColBERT (Khattab & Zaharia, 2020) preserve token granularity but produce multi-vector representations incompatible with standard embedding interfaces. In contrast, GLOT is the first approach to construct a *latent token-similarity graph* directly from frozen LLM hidden states, and perform explicit relational learning *within the pooling head* to produce a single, robust sentence vector.

Global Representations in Other Domains. The challenge of creating a single, global representation from a set of features is not unique to NLP. In computer vision, pooling has long been a central component in convolutional neural networks (CNNs). While operations like max and average pooling are used throughout these models (Krizhevsky et al., 2012; He et al., 2016), global pooling is critical for producing a hoslistic representation. Techniques like global average pooling are used to collapse the final spatial feature maps into a single feature vector for classification, effectively summarizing the most salient features present in an image (Lin et al., 2013). In NLP, by contrast, pooling is often treated as a final, routine step. Our work, GLOT, challenges this view by demonstrating that a graph-neural-based sentence-level learning approach can unlock significant performance gains from frozen language models, opening a new direction for efficient sentence-level model adaptation.

3 Method

In this section we formalize and discuss the properties of our method. We start by providing essential notations and problem formulation in Section 3.1, followed by Section 3.2 where we present GLOT.

3.1 PROBLEM SETUP

Given a sequence of input tokens $[x_1, x_2, \cdots, x_L]$ and a frozen LLM, the task is to design a function f_{pool} , that maps the sequence of token-level hidden states $\mathbf{X} = [\mathbf{x}_1, \mathbf{x}_2, \cdots, \mathbf{x}_L] \in \mathbb{R}^{L \times d}$, to a single, sentence-level representation, $\mathbf{z} \in \mathbb{R}^D$. This vector \mathbf{z} is a critical input for many downstream applications considered in this work, as follows:

• Single-Sentence Classification. For tasks like sentiment analysis, the vector \mathbf{z} is fed into a linear classifier, $y = \operatorname{softmax}(\mathbf{W}\mathbf{z} + \mathbf{b})$ to obtain the sentence label, where \mathbf{W} and \mathbf{b} are trainable parameters.

Figure 2: An overview of the GLOT pooling architecture. Given token hidden states from a frozen language model, our trainable module performs three stages: (1) it constructs a latent token-similarity graph, (2) a TOKEN-GNN performs relational learning to refine token representations, and (3) a readout layer aggregates the refined vectors into a final sentence representation, z

- Sentence-Pair Classification. For tasks like entailment detection, two sentence vectors, \mathbf{z}_a and \mathbf{z}_b , are concatenated and passed to a linear classifier to obtain a label $y = \operatorname{softmax}(\mathbf{W}[\mathbf{z}_a || \mathbf{z}_b] +$ **b**), where || denotes channel-wise concatenation.
- Similarity and Retrieval. For ranking, the semantic relatedness of two vectors, \mathbf{z}_a and \mathbf{z}_b , is measured with a function like cosine similarity, $sim(\mathbf{z}_a, \mathbf{z}_b) = \mathbf{z}_a^{\top} \mathbf{z}_b / ||\mathbf{z}_a|| ||\mathbf{z}_b||$.

3.2 GLOT

162

163 164

165

166 167

168

169 170

171

172

173

174 175 176

177

178

179

181

183

185

186

187

188

189

190

191

192 193

194

195

196

197

199

200 201

202

203 204

205

206

207

208

209

210

211 212

213

214

215

We introduce GLOT, a trainable framework that transforms the token-level hidden states into a final, sentence-level vector, $\mathbf{z} = \text{GLOT}(\mathbf{X})$. As illustrated in Figure 2, this process involves three stages: (1) constructing a token graph, (2) refining token states with a graph neural network (GNN) denoted TOKEN-GNN, and (3) performing a learnable readout. The steps are explained in the following:

Step 1: Token Graph Construction. Given token hidden states $\mathbf{X} = [\mathbf{x}_1, \mathbf{x}_2, \cdots, \mathbf{x}_L] \in \mathbb{R}^{L \times d}$ that are obtained from an LLM with hidden dimensionality d, after processing an input of length L, we construct a token graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ where nodes $|\mathcal{V}| = L$ correspond to tokens. Edges are defined by the cosine similarity S_{ij} between token vectors x_i and x_j . To induce a sparse, semantic structure, we only create edges where S_{ij} exceeds a threshold τ , which is a hyperparameter, discussed in Section 4.

Step 2: Refinement with TOKEN-GNN. Next, we apply a lightweight graph neural network, dubbed TOKEN-GNN, to refine the token representations by modeling their interactions. With token hidden states \mathbf{X} , we initialize node features $\mathbf{H}^{(0)} = \mathbf{X} \mathbf{W}_{in} \in \mathbb{R}^{L \times p}$ using a learnable matrix $\mathbf{W}_{in} \in \mathbb{R}^{d \times p}$, where p is the hidden dimension of the GNN. Overall, we apply K GNN layers to produce a set of refined, structure-aware token representations $\mathbf{H}^{(K)} = \mathbf{U} = [\mathbf{u}_1, \cdots, \mathbf{u}_L] \in \mathbb{R}^{L \times p}$. Each layer $\ell = 1, \dots, K$ of the TOKEN-GNN computes:

$$\mathbf{a}_{i}^{(\ell)} = \text{AGGREGATE}\left(\mathbf{h}_{j}^{(\ell)}\right) \in \mathbb{R}^{p},\tag{1}$$

$$\mathbf{a}_{i}^{(\ell)} = \text{AGGREGATE}\left(\mathbf{h}_{j}^{(\ell)}\right) \in \mathbb{R}^{p},$$

$$\mathbf{h}_{i}^{(\ell+1)} = \sigma\left(\mathbf{W}^{(\ell)}\text{CONCAT}(\mathbf{h}_{i}^{(\ell)}, \mathbf{a}_{i}^{(\ell)})\right),$$
(2)

where $\mathbf{a}_i^{(\ell)}$ is the aggregated information from the neighbors \mathcal{N}_i of token i, AGGREGATE is a permutation-equivariant aggregation function like sum or mean, $\mathbf{W}^{(\ell)} \in \mathbb{R}^{p \times 2p}$ is a learnable weight matrix, and σ is a nonlinear activation function, with implementation details in Appendix B.

Step 3: Readout Layer. The set of refined token representations, U, is aggregated into the sentence vector z via learnable scoring. A scalar importance score m_i is computed for each refined token vector \mathbf{u}_i , normalized using softmax to create weights $\boldsymbol{\pi}$, and used to compute a weighted sum:

$$m_i = \mathbf{v}^{\top} \tanh(\mathbf{W}_m \mathbf{u}_i + \mathbf{b}_m), \quad \boldsymbol{\pi} = \operatorname{softmax}(\mathbf{m}), \quad \mathbf{z} = \sum_{i=1}^{L} \pi_i \mathbf{u}_i,$$
 (3)

where $\mathbf{m} = [m_1, \dots, m_L].$

Overall, GLOT aggregates token-level hidden states obtained from a frozen LLM, to obtain refined and learnable sentence-level representations by modeling token-token relationships using a graph and processing them using TOKEN-GNN.

Properties of GLOT. The GLOT framework extends several common methods for obtaining sentence-level representations, which can be recovered as special cases. If we disable the TOKEN-GNN by setting its number of layers to zero (i.e., K=0), then the refined vectors are simply the original hidden states (that is, $\mathbf{u}_i=\mathbf{x}_i$), and the framework reduces to a direct weighted pooling mechanism. From here, we can model both standard pooling methods (like mean or CLS pooling) by using fixed weights and adaptive scoring methods, like AdaPool from Brothers (2025), by keeping the weights learnable.

These cases, where K=0, fit into the DeepSets framework (Zaheer et al., 2017), in which all elements \mathbf{x}_i are transformed individually $\phi(\mathbf{x}_i)$ before a global aggregation function. Instead, the Token-GNN utilized in GLOT enables information exchange in the form of $\phi(\mathbf{x}_i,\mathcal{G})$, taking a more global approach and allowing interactions between tokens. Bronstein et al. (2021) has shown DeepSets to be a special case of convolutional GNNs with no edge connectivity and, thus, strictly less powerful than message passing, an advantage we exploit in GLOT. The additional communication introduced in GLOT between tokens' representations allows it to model linguistic phenomena that hinge on pairwise or multi-hop dependencies among the tokens. The GNN mechanism in GLOT requires additional memory and computations, compared with pre-defined methods. Nonetheless, we note that, in comparison to other methods, which require the fine-tuning of the entire backbone LLMs, our GLOT strikes a balance between efficiency and effectiveness in downstream performance, as is evident in Section 4 and Figure 1.

4 EXPERIMENTS AND DISCUSSION

We conduct a comprehensive evaluation of GLOT to validate our core hypothesis: obtaining sentence-level representation via its reframing as relational learning before compression yields superior sentence embeddings from frozen LLMs compared with traditional and recent learnable approaches. Throughout our experiments, all backbone LLM models remain completely frozen; only the lightweight GLOT head and a minimal task-specific classifier are trained. This design ensures our approach is both parameter and resource-efficient. Our evaluation is guided by four key research questions:

- (RQ1) How does GLOT compare to standard pre-defined and learnable sentence-level representation methods, across diverse LLMs and tasks?
- (RQ2) Does explicit relational learning offer consistent improvements, especially for decoder-only models?
- (RQ3) Can our GLOT match or exceed the performance of fine-tuned models while maintaining the computational efficiency of frozen LLMs?
- (**RQ4**) How robust is GLOT to the signal dilution that affects traditional techniques?

4.1 EXPERIMENTAL SETUP

We evaluate GLOT against standard static (Mean, Max, CLS/EOS) and learnable pooling baselines across a diverse set of frozen encoder (BERT (Devlin et al., 2019), RoBERTa (Liu et al., 2019)) and decoder (e.g., Llama (Meta AI, 2024), Mistral (Jiang et al., 2023)) models. The evaluation is conducted on a wide range of tasks, including general language understanding (GLUE) (Wang et al., 2018), long-text classification (IMDB) (Maas et al., 2011), and retrieval (MTEB) (Muennighoff et al., 2023). To specifically test for relational robustness, we also introduce a synthetic diagnostic stress test that measures performance under noise. Across all experiments, the LLM backbones remain completely frozen. Full details on all models, baselines, benchmarks, training hyperparameters, and evaluation protocols are provided in Appendix B.

4.2 GENERAL LANGUAGE UNDERSTANDING EVALUATION (GLUE BENCHMARK)

Across the GLUE benchmark, GLOT consistently outperforms all baselines on all LLMs, from encoders like BERT to decoders like Mistral-7B. Table 1 provides the detailed scores, while Figure 4 of Appendix C visualizes the overall trend, showing that our GLOT's advantage is consistent across different task categories. This directly addresses (**RQ1**) and (**RQ2**).

Table 1: A comparison of pooling methods on the GLUE benchmark using six different frozen backbones. The table reports standard metrics: MCC for CoLA, Spearman for STS-B, F1 for MRPC/QQP, and Accuracy for the rest. Scores are multiplied by 100, with the best performance for each model highlighted in bold.

Model	Method	CoLA MCC↑	SST-2 ACC↑	STS-B SPEA. ↑	MRPC F1↑	QQP F1↑	MNLI-m ACC↑	MNLI-mm ACC↑	QNLI ACC↑	RTE ACC↑	WNLI ACC↑
	[CLS]	22.66	83.83	61.08	79.58	19.70	43.86	45.03	54.75	50.90	45.07
H	Mean	19.55	82.91	74.96	80.28	29.01	43.86	45.16	56.43	51.62	52.11
BERT	Max	15.79	80.73	74.12	81.64	29.58	38.60	39.55	53.79	51.98	49.26
<u> </u>	AdaPool	29.20	87.72	80.01	77.99	40.15	48.57	49.93	58.04	51.62	45.07
	GLOT	47.49	90.25	83.86	82.58	62.19	54.39	54.47	61.08	59.21	54.93
	[CLS]	6.92	66.63	52.87	81.22	47.66	32.78	32.98	54.89	52.34	40.85
RoBERTa	Mean	23.69	84.12	70.55	81.92	48.97	39.15	38.76	57.77	54.63	38.73
Ξ	Max	22.06	79.10	66.39	81.52	44.69	35.54	35.37	52.49	52.22	52.81
9	AdaPool	26.80	90.97	71.12	80.78	57.71	42.51	44.24	59.72	50.45	41.90
~	GLOT	56.08	92.78	85.27	81.95	61.41	57.01	57.95	62.73	56.68	56.34
7	[EOS]	7.63	77.75	52.77	81.03	38.11	41.14	42.66	53.23	49.10	47.88
Ž	Mean	12.30	79.81	56.39	80.60	32.34	40.50	41.06	55.97	54.15	42.25
SmolLM2	Max	2.38	73.62	52.10	76.72	24.02	37.44	38.40	54.84	51.62	52.11
Ĕ	AdaPool	7.21	83.71	61.20	81.69	49.26	41.00	42.35	58.08	55.59	45.07
S	GLOT	39.23	90.25	76.28	82.24	62.32	53.42	53.64	59.86	57.40	63.38
g	[EOS]	8.33	73.85	64.63	80.31	41.46	39.33	40.92	56.19	47.29	45.07
Ē	Mean	5.93	73.85	61.29	80.67	41.46	39.50	40.83	57.51	49.58	45.07
Ź	Max	2.76	70.87	63.99	81.45	39.64	36.88	37.93	55.29	50.90	46.48
TinyLlama	AdaPool	4.63	59.92	69.53	81.04	30.17	42.69	43.49	57.71	46.20	50.70
1	GLOT	17.61	80.73	71.77	82.54	59.92	48.04	49.34	63.77	57.40	53.52
e	[EOS]	37.37	91.74	74.11	70.58	58.78	48.47	47.46	53.98	54.87	42.25
<u> </u>	Mean	20.91	87.04	78.62	70.34	56.82	48.06	47.19	59.60	57.40	45.07
Ž	Max	13.49	84.51	73.27	67.64	51.17	40.89	40.77	55.84	49.45	47.88
LLaMA-3B	AdaPool	43.32	92.54	81.93	71.81	49.37	49.56	50.59	58.48	55.23	47.88
П	GLOT	55.13	93.92	82.83	82.34	61.16	53.49	54.67	67.15	61.01	56.34
m	[EOS]	38.63	92.55	72.36	76.32	51.68	48.18	48.33	50.82	50.90	40.85
1-7	Mean	38.61	89.91	77.96	77.22	57.44	47.86	48.08	53.46	53.07	42.25
i.	Max	10.78	85.89	70.72	65.61	54.39	38.77	39.30	58.70	53.07	48.70
Mistral-7B	AdaPool	48.00	93.00	79.55	81.12	49.07	50.72	51.56	55.75	54.87	49.30
≥	GLOT	54.30	94.38	80.51	82.83	64.07	51.66	53.22	60.93	59.21	56.34

GLOT achieves its most significant performance gains on tasks that require nuanced relational understanding. On the Corpus of Linguistic Acceptability (CoLA) (Warstadt et al., 2018), for instance, GLOT dramatically improves the Matthew's Correlation Coefficient for BERT by a relative improvement of **62.63**% and **13.13**% for Mistral-7B. This suggests that by explicitly modeling token relationships, our approach better captures the grammatical structure essential for this task. Similarly, on Quora Question Pairs (**QQP**), a paraphrase detection task, GLOT delivers a large performance improvement margin over baselines for all tested architectures.

The consistent superiority on single-sentence classification (SST-2) (Socher et al., 2013b), semantic similarity (STS-B) (Agirre et al., 2007), and inference (RTE) (Dagan et al., 2006; Bar Haim et al., 2006; Giampiccolo et al., 2007; Bentivogli et al., 2009) tasks validates that our "relational learning before compression" approach yields more robust and general-purpose embeddings than methods that pool token states in isolation.

4.3 Long-Text Classification

We assess performance on longer sequences using the IMDB dataset (Maas et al., 2011), where the task is to classify paragraph-length reviews. As shown in Table 2, GLOT consistently outperforms all baselines. For instance, it improves accuracy by nearly **4.5%** for RoBERTa over the strongest baseline and by an average of **+10.1%** relative improvement over the standard [EOS] token for decoder models. This result highlights the effectiveness of our graph-based approach on long-form text; unlike simple pooling, which can dilute sentiment signals across long contexts, GLOT's relational learning preserves and utilizes critical phrases for more accurate classification.

Table 2: Accuracy (×100) on the IMDB long-text sentiment classification task. We freeze the LLM backbones and train only the pooling heads and a linear classifier. The **best** result per model is in bold.

Method	BERT	RoBERTa	SmolLM2	TinyLlama	LLaMA3.2-3B	Mistral-7B
[CLS]/[EOS]	80.23	82.04	82.82	87.27	90.56	84.86
Mean	81.64	84.38	84.10	88.72	92.58	94.21
Max	60.78	58.80	63.41	75.45	80.90	64.43
AdaPool	85.45	90.91	91.56	92.61	95.71	95.66
GLOT	86.93	94.52	94.18	93.38	96.14	95.95

Table 3: **Zero-shot performance on seven diverse tasks from the MTEB benchmark.** Prior to evaluation, we train all learnable pooling heads on the MS MARCO dataset. The **best** performance for each frozen backbone is in bold.

Model	Method	EmotionClass. ACC ↑	SciFact NDCG@10↑	RedditClust. V-MEAS. ↑	AskUbuntu MAP↑	STS12 Cos. Spea. ↑	TwitterSemEval MAX AP. ↑	SummEval Cos. Spea. ↑
BERT	[CLS] Mean Max AdaPool GLOT	0.2412 0.3361 0.2812 0.3513 0.3715	0.0231 0.1769 0.2771 0.2224 0.2485	0.1417 0.2777 0.2241 0.3403 0.3630	0.4137 0.4584 0.4553 0.4778 0.5020	0.2153 0.3087 0.3175 0.3941 0.4862	0.3433 0.5613 0.5450 0.5195 0.5623	0.2792 0.2983 0.3022 0.2918 0.3068
RoBERTa	[CLS] Mean Max AdaPool GLOT	0.2759 0.2520 0.2200 0.2135 0.2909	0.0900 0.0825 0.0116 0.0042 0.2605	0.1908 0.1850 0.1354 0.1475 0.2184	0.4439 0.4621 0.4491 0.4513 0.4687	0.1667 0.3210 0.2667 0.2026 0.3688	0.4848 0.5456 0.5000 0.4744 0.5598	0.2347 0.2986 0.2583 0.2276 0.3083
Smoll.M2	[EOS] Mean Max AdaPool GLOT	0.2252 0.2396 0.1923 0.2360 0.2471	0.0012 0.1313 0.0385 0.1702 0.1834	0.1418 0.1708 0.0960 0.1905 0.2306	0.4113 0.4428 0.4382 0.4461 0.4529	0.1900 0.3824 0.2458 0.4322 0.4754	0.3613 0.4256 0.3650 0.4153 0.4343	0.2271 0.2335 0.2530 0.2591 0.2628
TinyLlama	[EOS] Mean Max AdaPool GLOT	0.2044 0.1898 0.1820 0.2904 0.2905	0.0042 0.0126 0.0049 0.0602 0.0916	0.0689 0.0687 0.0591 0.1688 0.1800	0.4275 0.4269 0.4292 0.4004 0.4341	0.1297 0.1633 0.1842 0.0329 0.2369	0.3532 0.3150 0.3588 0.2811 0.3804	0.2602 0.2450 0.1178 0.2521 0.2649
LLaMA-3B	[EOS] Mean Max AdaPool GLOT	0.2765 0.2920 0.2478 0.2185 0.3046	0.0087 0.4247 0.4087 0.4140 0.4586	0.1979 0.3034 0.1943 0.2774 0.3301	0.4420 0.4971 0.4906 0.4946 0.5103	0.2494 0.4296 0.3367 0.3765 0.4616	0.4141 0.4430 0.4196 0.3216 0.4431	0.1917 0.1924 0.2347 0.2350 0.2658
Mistral-7B	[EOS] Mean Max AdaPool GLOT	0.2662 0.2995 0.2142 0.2832 0.3016	0.0033 0.3735 0.2116 0.4268 0.4414	0.1858 0.2544 0.1015 0.2398 0.2623	0.4352 0.4774 0.4577 0.4767 0.4821	0.2307 0.3824 0.3017 0.3641 0.3905	0.3846 0.4106 0.4151 0.3510 0.4221	0.2042 0.1964 0.2470 0.2346 0.2774

4.4 LARGE-SCALE BENCHMARKING ON MTEB

To assess GLOT's performance as a general-purpose sentence encoder, we evaluate it on seven diverse tasks from the Massive Text Embedding Benchmark (MTEB) (Muennighoff et al., 2023). Since many tasks are zero-shot, all learnable heads are first trained on the MS MARCO dataset (Bajaj et al., 2016) with a contrastive loss while keeping the LLM backbones frozen. The specific MTEB tasks are detailed in Appendix B.

The results in Table 3 show that GLOT is a robust performer across all tasks for both encoder- and decoder-only architectures. For RoBERTa, GLOT achieves the best score on all seven tested tasks, with a notable $\times 3$ improvement on SciFact. This advantage extends to decoders: with the Llama-3B backbone, GLOT secures a top performance of **0.5103 MAP** on AskUbuntuDupQuestions, rivaling strong encoder-only models. This strong general-purpose performance, achieved without expensive backbone fine-tuning, provides a clear affirmative answer to (**RQ3**).

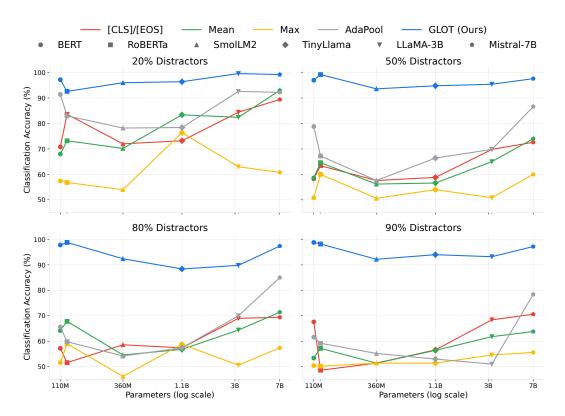


Figure 3: **Robustness to signal dilution on the diagnostic stress test.** Each of the four panels displays the classification accuracy for all pooling methods at a specific distractor ratio, which increases from 20% to 90%. Within each panel, backbone models are arranged along the x-axis by their parameter count.

4.5 DIAGNOSTIC ANALYSIS: EVALUATING RELATIONAL ROBUSTNESS

To test for relational robustness under noise (**RQ4**), we design a synthetic diagnostic task inspired by 'signal-in-noise' evaluations (Brothers, 2025) and the 'Needle in a Haystack' paradigm (Kamradt, 2023). The test involves injecting a short phrase containing a logical dependency (e.g., ...not...keys...) into a long sequence of random words. A binary classifier must then interpret the logic of the signal phrase, with difficulty controlled by increasing the distractor ratio from 20% to 90%. The pseudo-code for synthetic data generation is presented in Algorithm 2 of Appendix B.

The results in Figure 3 show a stark divergence. As noise increases, the accuracy of baseline methods collapses; on Mistral-7B, AdaPool's accuracy plummets from 92.2% to 78.4%, and Mean pooling drops to 63.8%. In contrast, GLOT remains robust, maintaining over 97% accuracy even at the 90% distractor level. This confirms that GLOT's explicit token graph successfully bypasses the signal dilution that plagues methods reliant on global summary statistics. Full results are available in Table 7 in Appendix C.

4.6 ABLATIONS AND ANALYSIS

We conduct a series of ablations and analyses to validate GLOT's design choices and quantify its computational efficiency.

Impact of Graph Sparsity. To understand the importance of constructing a well-formed semantic graph, we ablate the similarity threshold parameter, τ , using the Mistral-7B backbone on GLUE benchmark. As shown in Table 4, the graph structure is critical to performance. When $\tau=0.0$, the graph is fully connected, allowing noisy or irrelevant token relationships to dilute the message passing process, resulting in suboptimal performance on all tasks. As we increase τ , pruning weaker

Table 4: An ablation study on the impact of graph sparsity in GLOT. This table shows performance on GLUE tasks using the Mistral-7B backbone as we vary the similarity threshold (τ) for token graph construction. All scores are multiplied by 100, and the **best** result for each task is in bold.

Method	CoLA MCC↑	SST-2 ACC ↑	STS-B SPEA. ↑	MRPC F1 ↑	QQP F1 ↑	MNLI-m ACC↑	MNLI-mm ACC ↑	QNLI ACC ↑	RTE ACC ↑	WNLI ACC↑
GLOT $(\tau = 0.0)$	50.19	93.69	80.34	81.04	62.79	49.09	49.46	52.85	49.81	38.03
GLOT $(\tau = 0.2)$	53.40	94.38	80.48	82.83	62.53	51.66	53.22	54.15	49.45	36.62
GLOT ($\tau = 0.4$)	51.73	93.46	80.40	80.25	64.07	48.81	49.94	60.93	50.54	40.84
GLOT $(\tau = 0.6)$	54.30	93.23	80.29	80.06	63.49	49.36	50.01	53.67	54.15	56.34
GLOT ($\tau = 0.8$)	52.48	92.66	80.26	79.87	63.22	48.92	49.66	55.09	52.70	56.34

Table 5: A comparison of training methods by resource consumption and performance on the CoLA task, using the Mistral-7B backbone. We contrast our frozen-backbone approach (GLOT) against full fine-tuning (Full FT) and LoRA. Batch runtime is reported as the mean \pm standard deviation over 10 measurements.

Method	# Trainable Params	GPU Memory (GB)↓	Batch Runtime (ms)↓	MCC↑
Full FT + EOS	7.11B	32.59	1318.8 ± 1.1	49.63
LoRA (r = 64) + EOS	167.8M	33.50	1454.6 ± 1.1	48.23
GLOT (ours)	8.92M	0.42	13.4 ± 3.0	53.29

edges, performance steadily improves across most tasks, plateauing in the range of $\tau=0.4-0.6$. This confirms that not all token relations are equally important; by focusing on the strongest semantic connections via relational learning, GLOT produces a more robust sentence representation.

Computational Efficiency. To address (RQ3), we compare the resource consumption of GLOT against full fine-tuning (Full FT) and Parameter-Efficient Fine-Tuning (PEFT) with LoRA (Hu et al., 2022). The results in Table 5 highlight the dramatic efficiency of our approach. GLOT requires only 8.92M trainable parameters, which is approximately $20\times$ fewer than LoRA. This parameter efficiency translates directly to a minimal memory footprint of only 0.42 GB, compared to over 32 GB for the other methods. Consequently, GLOT is over $100\times$ faster per training batch. This demonstrates that our method provides a practical and accessible way to generate high-quality embeddings from large, frozen LLMs on consumer-grade hardware.

5 CONCLUSION

As LLMs continue to scale, the computational cost of full fine-tuning becomes prohibitive, establishing the need for improved pooling methods that operate on frozen backbones as a crucial research problem. In this work, we addressed a fundamental limitation of standard pooling: that it treats token hidden states as an independent set of vectors, discarding the rich relational structure captured by language models. We introduced GLOT, a lightweight and parameter-efficient pooling head that instantiates a new paradigm of relational learning followed by aggregation. GLOT first constructs a latent token-similarity graph, refines token representations using a GNN, and then aggregates them with an attention mechanism.

Through comprehensive experiments, we demonstrated that GLOT consistently outperforms strong baselines across a wide range of tasks and on both encoder- and decoder-only models. Our diagnostic stress test provided direct evidence that GLOT's graph-based learning makes it remarkably robust to the signal dilution that plagues traditional pooling. Furthermore, we showed that GLOT is up to two orders of magnitude more computationally efficient than parameter-efficient fine-tuning techniques like LoRA, making it a practical solution for adapting billion-parameter models.

Our findings challenge the view that pooling is a routine final step, showing instead that a carefully designed, relational learning-based head can unlock significant performance from frozen models. This work opens several avenues for future research, including exploring learnable graph construction mechanisms and applying the "relational learning before compression" paradigm to other modalities, such as pooling patch embeddings in Vision Transformers.

ETHICS STATEMENT

Our work primarily focuses on developing a new pooling methodology and is evaluated on publicly available, standard academic benchmarks, including GLUE, MTEB, and IMDB. We do not use any private or sensitive user data, and our experiments did not involve human subjects. We acknowledge that the pre-trained language models used as backbones in our study may reflect societal biases present in their training corpora. Our proposed method, GLOT, operates on the outputs of these models and does not introduce new sources of bias, nor does it explicitly mitigate biases inherent in the backbone models. We intend for this work to contribute to the development of more efficient and robust NLP models, and we do not foresee any direct negative societal impacts from its application.

REPRODUCIBILITY STATEMENT

To ensure the reproducibility of our results, we will release our source code and all trained model checkpoints upon the paper's acceptance. Our methodology is described in Section 3, with detailed pseudo-code available in Algorithm 1. Appendix B provides a comprehensive description of our experimental setup, including the specific backbone models used, training and evaluation protocols, and all hyperparameters. All datasets used in our experiments are standard benchmarks publicly available through the Hugging Face Datasets library.

USAGE OF LARGE LANGUAGE MODELS (LLMS)

During the preparation of this manuscript, we utilized LLMs. Its role was strictly limited to that of grammatical assistance. The LLM was not used for research ideation, experimental design, data analysis, or the generation of any core scientific content. The authors take full responsibility for all content and claims presented in this paper.

REFERENCES

- Eneko Agirre, Llu'is M'arquez, and Richard Wicentowski (eds.). *Proceedings of the Fourth International Workshop on Semantic Evaluations (SemEval-2007)*. Association for Computational Linguistics, Prague, Czech Republic, June 2007.
- Sanjeev Arora, Yingyu Liang, and Tengyu Ma. A simple but tough-to-beat baseline for sentence embeddings. In *International conference on learning representations*, 2017.
- Payal Bajaj, Daniel Campos, Nick Craswell, Li Deng, Jianfeng Gao, Xiaodong Liu, Rangan Majumder, Andrew McNamara, Bhaskar Mitra, Tri Nguyen, Mir Rosenberg, Xia Song, Alina Stoica, Saurabh Tiwary, and Tong Wang. MS MARCO: A human generated machine reading comprehension dataset. In *NeurIPS 2016 Deep Learning for Question Answering (InCoCo@NIPS) Workshop*, 2016.
- Roy Bar Haim, Ido Dagan, Bill Dolan, Lisa Ferro, Danilo Giampiccolo, Bernardo Magnini, and Idan Szpektor. The second PASCAL recognising textual entailment challenge. 2006.
- Luisa Bentivogli, Ido Dagan, Hoa Trang Dang, Danilo Giampiccolo, and Bernardo Magnini. The fifth PASCAL recognizing textual entailment challenge. 2009.
- Michael M. Bronstein, Joan Bruna, Taco Cohen, and Petar Veličković. Geometric deep learning: Grids, groups, graphs, geodesics, and gauges, 2021. URL https://arxiv.org/abs/2104.13478.
- Greyson Brothers. Robust noise attenuation via adaptive pooling of transformer outputs, 2025. ICML 2025.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020.
- Quentin Cappart, Didier Chételat, Elias B Khalil, Andrea Lodi, Christopher Morris, and Petar Veličković. Combinatorial optimization and reasoning with graph neural networks. *Journal of Machine Learning Research*, 24(130):1–61, 2023.

- Ido Dagan, Oren Glickman, and Bernardo Magnini. The PASCAL recognising textual entailment challenge. In *Machine learning challenges. evaluating predictive uncertainty, visual object classification, and recognising tectual entailment*, pp. 177–190. Springer, 2006.
 - Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of NAACL-HLT*, pp. 4171– 4186, 2019.
 - William B Dolan and Chris Brockett. Automatically constructing a corpus of sentential paraphrases. In *Proceedings of the International Workshop on Paraphrasing*, 2005.
 - Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. *International Conference on Learning Representations (ICLR)*, 2021.
 - Kawin Ethayarajh. How contextual are contextualized word representations? comparing the geometry of bert, elmo, and gpt-2 embeddings. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pp. 55–65, 2019.
 - Matthias Fey and Jan Eric Lenssen. Fast graph representation learning with pytorch geometric. *arXiv* preprint arXiv:1903.02428, 2019.
 - Tianyu Gao, Xingcheng Yao, and Danqi Chen. Simcse: Simple contrastive learning of sentence embeddings. In *Proceedings of EMNLP*, pp. 6894–6910, 2021.
 - Danilo Giampiccolo, Bernardo Magnini, Ido Dagan, and Bill Dolan. The third PASCAL recognizing textual entailment challenge. In *Proceedings of the ACL-PASCAL workshop on textual entailment and paraphrasing*, pp. 1–9. Association for Computational Linguistics, 2007.
 - Justin Gilmer, Samuel S Schoenholz, Patrick F Riley, Oriol Vinyals, and George E Dahl. Neural message passing for quantum chemistry. In *International conference on machine learning*, pp. 1263–1272. Pmlr, 2017.
 - Will Hamilton, Zhitao Ying, and Jure Leskovec. Inductive representation learning on large graphs. *Advances in neural information processing systems*, 30, 2017.
 - Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 770–778, 2016.
 - Edward J Hu, yelong shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. LoRA: Low-rank adaptation of large language models. In *International Conference on Learning Representations*, 2022. URL https://openreview.net/forum?id=nZeVKeeFYf9.
 - Lianzhe Huang, Dehong Ma, Sujian Li, Xiaodong Zhang, and Houfeng Wang. Text level graph neural network for text classification. In Kentaro Inui, Jing Jiang, Vincent Ng, and Xiaojun Wan (eds.), Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP), pp. 3444–3450, Hong Kong, China, November 2019. Association for Computational Linguistics. doi: 10.18653/v1/D19-1345. URL https://aclanthology.org/D19-1345/.
 - Albert Q. Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, Lélio Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril, Thomas Wang, Timothée Lacroix, and William El Sayed. Mistral 7b, 2023. URL https://arxiv.org/abs/2310.06825.
 - Greg Kamradt. Needle In A Haystack: Pressure Testing LLMs. Blog post, nov 2023. URL https://gregkamradt.com/blog/needle-in-a-haystack-evaluating-llms-on-long-context-recall.
 - The original test and results that popularized the 'Needle in a Haystack' method for evaluating long-context recall in Large Language Models.

- Omar Khattab and Matei Zaharia. ColBERT: Efficient and effective passage search via contextualized late interaction over BERT. In *Proceedings of SIGIR*, pp. 39–48, 2020.
 - Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *International Conference on Learning Representations (ICLR)*, 2015. URL https://arxiv.org/abs/1412.6980.
 - Thomas N Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. In *International Conference on Learning Representations*, 2017.
 - Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems*, 25, 2012.
 - Chankyu Lee, Rajarshi Roy, Mengyao Xu, Jonathan Raiman, Mohammad Shoeybi, Bryan Catanzaro, and Wei Ping. NV-embed: Improved techniques for training LLMs as generalist embedding models. In *The Thirteenth International Conference on Learning Representations*, 2025. URL https://openreview.net/forum?id=lgsyLSsDRe.
 - Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. Bart: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pp. 7871–7880, 2020.
 - Quentin Lhoest, Albert Villanova del Moral, Yacine Jernite, Abhishek Thakur, Patrick von Platen, Suraj Patil, Julien Chaumond, Mariama Drame, Julien Plu, Lewis Tunstall, Joe Davison, Tomáš Šaško, Gunjan Chhablani, Bhavitvya Malik, Simon Brandeis, Teven Le Scao, Victor Sanh, Canwen Xu, Nicolas Patry, Angelina McMillan-Major, Philipp Schmid, Sylvain Gugger, Clément Delangue, Théo Matussière, Lysandre Debut, Stas Bekman, Pierric Cistac, Thibault Goehringer, Arthur Mustar, Sanchit Mangrulkar, Alexander M. Rush, and Thomas Wolf. Datasets: A community library for natural language processing. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pp. 175–184. Association for Computational Linguistics, 2021. URL https://aclanthology.org/2021.emnlp-demo.21.
 - Bohan Li, Hao Zhou, Junxian He, Mingxuan Wang, Yiming Yang, and Lei Li. On the sentence embeddings from pre-trained language models. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 9119–9130, 2020.
 - Min Lin, Qiang Chen, and Shuicheng Yan. Network in network. *arXiv preprint arXiv:1312.4400*, 2013.
 - Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. Roberta: A robustly optimized bert pretraining approach. *arXiv:1907.11692*, 2019.
 - Yonghao Liu, Renchu Guan, Fausto Giunchiglia, Yanchun Liang, and Xiaoyue Feng. Deep attention diffusion graph neural networks for text classification. In Marie-Francine Moens, Xuanjing Huang, Lucia Specia, and Scott Wen-tau Yih (eds.), *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pp. 8142–8152, Online and Punta Cana, Dominican Republic, November 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021. emnlp-main.642. URL https://aclanthology.org/2021.emnlp-main.642/.
 - Xueguang Ma, Liang Wang, Nan Yang, Furu Wei, and Jimmy Lin. Fine-tuning llama for multi-stage text retrieval. In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pp. 2421–2425, 2024.
- Andrew L. Maas, Raymond E. Daly, Peter T. Pham, Dan Huang, Andrew Y. Ng, and Christopher Potts. Learning word vectors for sentiment analysis. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pp. 142–150. Association for Computational Linguistics, 2011.
 - Meta AI. Llama 3.2: New open and customizable models with vision and on-device capabilities. Blog post, October 2024. URL https://ai.meta.com/blog/llama-3-2-connect-2024-vision-edge-mobile-devices/.

- Niklas Muennighoff, Nouamane Tazi, Loic Magne, and Nils Reimers. Mteb: Massive text embedding benchmark. In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, pp. 2014–2037. Association for Computational Linguistics, 2023.
 - Niklas Muennighoff, Hongjin Su, Liang Wang, Nan Yang, Furu Wei, Tao Yu, Amanpreet Singh, and Douwe Kiela. Generative representational instruction tuning, 2025. URL https://arxiv.org/abs/2402.09906.
 - Vinod Nair and Geoffrey E Hinton. Rectified linear units improve restricted boltzmann machines. In *Proceedings of the 27th International Conference on Machine Learning (ICML-10)*, pp. 807–814, 2010.
 - Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems 32*. Curran Associates, Inc., 2019.
 - Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. Language models are unsupervised multitask learners. *OpenAI Blog*, 2019. URL https://openai.com/research/language-models-are-unsupervised-multitask-learners.
 - Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of machine learning research*, 21(140):1–67, 2020.
 - Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. SQuAD: 100,000+ questions for machine comprehension of text. In *Proceedings of EMNLP*, pp. 2383–2392. Association for Computational Linguistics, 2016.
 - Nils Reimers and Iryna Gurevych. Sentence-bert: Sentence embeddings using siamese bert-networks. In *Proceedings of EMNLP-IJCNLP*, pp. 3982–3992, 2019.
 - Alvaro Sanchez-Gonzalez, Jonathan Godwin, Tobias Pfaff, Rex Ying, Jure Leskovec, and Peter Battaglia. Learning to simulate complex physics with graph networks. In *International conference on machine learning*, pp. 8459–8468. PMLR, 2020.
 - Michael Schlichtkrull, Thomas N Kipf, Peter Bloem, Rianne van den Berg, Ivan Titov, and Max Welling. Modeling relational data with graph convolutional networks. In *European semantic web conference*, pp. 593–607. Springer, 2018.
 - Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D Manning, Andrew Ng, and Christopher Potts. Recursive deep models for semantic compositionality over a sentiment treebank. In *Proceedings of EMNLP*, pp. 1631–1642, 2013a.
 - Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D. Manning, Andrew Ng, and Christopher Potts. Recursive deep models for semantic compositionality over a sentiment treebank. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pp. 1631–1642. Association for Computational Linguistics, 2013b.
 - Jianlin Su, Jiarun Cao, Weijie Liu, and Yangyiwen Ou. Whitening sentence representations for better semantics and faster retrieval. *arXiv preprint arXiv:2103.15316*, 2021.
 - Yixuan Tang and Yi Yang. Pooling and attention: What are effective designs for llm-based embedding models?, 2024. URL https://arxiv.org/abs/2409.02727.
 - Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, Brian Gan, Vighnesh Gante, Gartheeban Gholami, Vassilis Gkoumas, Kshitij Goyal, Thomas Hart, Sunny Hsia, Jason Huang, Alexandra Ispas, Jack Jacob, Saumya Jha, Anirudh Kumar, Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet,

Todor Mihaylov, Igor Molybog, Ylan Morisot, Victor O'Beirne, Eoin O'Sullivan, Alexander Pirogov, Roman Rabbat, Amjad Raghuraman, Sainbayar Ramjee, Ruan Ras, Jérémy Rault, Nicolas Rolland, Baptiste Rozière, Mohit Sachan, Todd Sawyers, Mykola Seljan, Adrien Seznec, Sharan Sun, Adel Tazairt, Gabriel Synnaeve, Yuxin Tan, Lilian Tang, Ross Taylor, Adina Williams, Jean Kenebrew, Mannan Zaheer, Ahmed El-Kishky, and Thomas Scialom. Llama 2: Open foundation and fine-tuned chat models, 2023.

- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in Neural Information Processing Systems*, volume 30, 2017.
- Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. Graph attention networks. In *ICLR*, 2018.
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R. Bowman. Glue: A multi-task benchmark and analysis platform for natural language understanding. In *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*. Association for Computational Linguistics, 2018.
- Liang Wang, Nan Yang, Xiaolong Huang, Linjun Yang, Rangan Majumder, and Furu Wei. Improving text embeddings with large language models. *arXiv preprint arXiv:2401.00368*, 2023.
- Liang Wang, Nan Yang, Xiaolong Huang, Binxing Jiao, Linjun Yang, Daxin Jiang, Rangan Majumder, and Furu Wei. Text embeddings by weakly-supervised contrastive pre-training, 2024. URL https://arxiv.org/abs/2212.03533.
- Alex Warstadt, Amanpreet Singh, and Samuel R. Bowman. Neural network acceptability judgments. *arXiv preprint 1805.12471*, 2018.
- Adina Williams, Nikita Nangia, and Samuel R. Bowman. A broad-coverage challenge corpus for sentence understanding through inference. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pp. 1112–1122. Association for Computational Linguistics, 2018. doi: 10.18653/v1/N18-1101.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander M. Rush. Transformers: State-of-the-art natural language processing. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pp. 38–45. Association for Computational Linguistics, 2020. URL https://aclanthology.org/2020.emnlp-demos.6.
- Jinming Xing, Dongwen Luo, Chang Xue, and Ruilin Xing. Comparative analysis of pooling mechanisms in llms: A sentiment analysis perspective. *arXiv* preprint arXiv:2411.14654, 2024.
- Keyulu Xu, Chengtao Li, Yonglong Tian, Tomohiro Sonobe, Ken-ichi Kawarabayashi, and Stefanie Jegelka. Representation learning on graphs with jumping knowledge networks. In *Proceedings of ICML*, pp. 5453–5462, 2018.
- Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. How powerful are graph neural networks? In *International Conference on Learning Representations*, 2019. URL https://openreview.net/forum?id=ryGs6iA5Km.
- Liang Yao, Chengsheng Mao, and Yuan Luo. Graph convolutional networks for text classification. In *Proceedings of AAAI*, pp. 7370–7377, 2019.
- Rex Ying, Ruining He, Kaifeng Chen, Pong Eksombatchai, William L Hamilton, and Jure Leskovec. Graph convolutional neural networks for web-scale recommender systems. In *Proceedings of the 24th ACM SIGKDD international conference on knowledge discovery & data mining*, pp. 974–983, 2018.
- Manzil Zaheer, Satwik Kottur, Siamak Ravanbakhsh, Barnabas Poczos, Russ R Salakhutdinov, and Alexander J Smola. Deep sets. *Advances in neural information processing systems*, 30, 2017.

A ADDITIONAL RELATED WORK

Fine-Tuning vs. Frozen Backbones for Embedding. A significant body of work adapts decoderonly LLMs into powerful text embedding models through extensive fine-tuning (Wang et al., 2023; Lee et al., 2025; Muennighoff et al., 2025; Ma et al., 2024; Tang & Yang, 2024). These methods achieve state-of-the-art performance but require modifying the LLM backbone, often through full-model training that is computationally prohibitive. GLOT sidesteps this entirely by operating on completely frozen backbones. Our approach is therefore lightweight, accessible, and applicable to both encoder-only and decoder-only models without expensive training.

The Geometry of Embedding Space. Recent studies reveal that token embeddings from LLMs occupy anisotropic manifolds, which makes cosine similarity between pooled sentence vectors unreliable (Ethayarajh, 2019; Li et al., 2020). While post-processing methods like whitening can mitigate this (Su et al., 2021), they do not address the underlying information loss from pooling. SBERT-style fine-tuning reshapes this geometry but is computationally expensive. Our work offers an alternative: by constructing a similarity graph, GLOT operates on an approximation of the intrinsic manifold geometry, preserving relational structures that are lost when pooling in the ambient Euclidean space.

Applications of Graph Neural Networks. The success of Graph Neural Networks (GNNs) is demonstrated by their wide-ranging application across numerous scientific and industrial domains. In the life sciences, they have become a cornerstone for molecular property prediction and drug discovery, where molecules are modeled as graphs of atoms and bonds (Gilmer et al., 2017; Xu et al., 2019). Similarly, they are used to analyze complex protein-protein interaction networks in bioinformatics. In the digital realm, GNNs power modern recommender systems by capturing the intricate relationships between users and items (Ying et al., 2018), and they are essential for learning over large-scale knowledge graphs (Schlichtkrull et al., 2018). Their foundational use case remains the analysis of social networks, where they are applied to tasks like node classification and community detection (Kipf & Welling, 2017; Hamilton et al., 2017). GNNs have also been successfully applied in other areas, including modeling particle systems in physics simulations (Sanchez-Gonzalez et al., 2020), processing 3D point clouds in computer vision, and solving complex combinatorial optimization problems like the Traveling Salesperson Problem (Cappart et al., 2023).

B IMPLEMENTATION DETAILS

B.1 GENERAL SETUP

Hardware and Software. All experiments were conducted on a single NVIDIA A6000 GPU. Our implementation is built using PyTorch (Paszke et al., 2019), with extensive use of the Hugging Face ecosystem (Wolf et al., 2020), including transformers for backbone models and datasets (Lhoest et al., 2021) for data loading. The graph-based components of our method are implemented using PyTorch Geometric (Fey & Lenssen, 2019). Large-scale benchmarking was performed using the mteb (Muennighoff et al., 2023) library, and retrieval metrics were calculated using ranx.

Training Details. Unless otherwise noted, all trainable pooling heads were trained for 2 epochs using the Adam optimizer (Kingma & Ba, 2015) with a learning rate of 2×10^{-4} and no weight decay. We used a training batch size of 32 and an evaluation batch size of 64. For all experiments, we used a fixed random seed of 42. To accelerate training, we implemented a feature to precompute and cache the frozen backbone's hidden states before training the pooling heads. We provide the pseudocode for GLOT in Algorithm 1. The hyperparameter tuning shown in Table 6 using Weights and Biases framework.

B.2 Model Configurations

Backbone Models. All backbone models were loaded from the Hugging Face Hub. For decoderonly models, the tokenizer's padding side was set to 'right'. If a model did not have a pre-defined padding token, the '[EOS]' token was used.

841

858

859 860

861

862

863

```
810
          Algorithm 1 GLOT: Graph-based Token Pooling
811
           Require: H \in \mathbb{R}^{B \times L \times d_{in}}: Batch of hidden states from a frozen LLM.
812
                     M \in \{0,1\}^{B \times L}: Attention mask for the hidden states.
813
                     \tau: Cosine similarity threshold for edge creation.
814
                     K: Number of layers in the TOKEN-GNN.
815
          Ensure: Z \in \mathbb{R}^{B \times d_{out}}: Batch of final sentence embeddings.
816
            1: function GLOT(H, M)
817
            2:
                    \mathcal{G}_{list} \leftarrow []
818
            3:
                    for i=1 \rightarrow B do
                                                                                        Step 1: Token Graph Construction
819
                         H_i' \leftarrow H[i, M[i] == 1, :]
            4:
                                                                                            \triangleright Get valid tokens for sentence i
            5:
                         S_i \leftarrow \text{CosineSimilarity}(H_i', H_i')
                                                                                                  ▶ Pairwise similarity matrix
820
                         A_i \leftarrow (S_i > \tau)
                                                                            ▷ Create adjacency matrix based on threshold
            6:
821
                         edge\_index_i \leftarrow ADJACENCYTOEDGES(A_i)
            7:
822
                         \mathcal{G}_{list}. APPEND (nodes = H'_i, edges = edge_index<sub>i</sub>)
            8:
823
            9:
824
           10:
                    \mathcal{G}_{batch} \leftarrow \text{BATCHGRAPHS}(\mathcal{G}_{list})
                                                                                      ▷ Combine graphs into a single batch
825
           11:
                    U_0, edge_index, batch_idx \leftarrow \mathcal{G}_{batch}.x, \mathcal{G}_{batch}.edge_index, \mathcal{G}_{batch}.batch
826
           12:
                    \mathcal{U}_{layers} \leftarrow [U_0]
827
                    for k=1 \rightarrow K do
           13:
                                                                                 ▶ Step 2: Refinement with TOKEN-GNN
828
           14:
                         U_{k-1} \leftarrow \mathcal{U}_{layers}[k-1]
829
           15:
                         U_k \leftarrow \text{GNN-LAYER}_k(U_{k-1}, \text{edge\_index})
                        \mathcal{U}_{layers}.\mathtt{APPEND}(U_k)
830
           16:
831
           17:
           18:
                    U_{fused} \leftarrow \text{JumpingKnowledgeConcat}(\mathcal{U}_{layers})
                                                                                                      ⊳ Step 3: Feature Fusion
832
           19:
                    m \leftarrow \text{READOUTMLP}(U_{fused})
                                                                                                      833
                    \pi \leftarrow \text{SOFTMAXBYGRAPH}(m, \text{batch\_idx})
                                                                                     ▶ Normalize scores per sentence graph
           20:
834
           21:
                    Z_{pooled} \leftarrow \pi \odot U_{fused}

    ▶ Apply attention weights

835
                    \hat{Z} \leftarrow \text{SCATTERADD}(Z_{pooled}, \text{batch\_idx})
           22:
                                                                                  836
           23:
                    return Z
837
           24: end function
838
```

Table 6: Hyperparameter search space for the GLOT pooling head. The final model configuration was determined via a grid search over these values. The search was performed consistently across all backbone models and datasets.

Hyperparameter	Search Space
Optimization Learning Rate Weight Decay	{1e-3, 2e-4, 2e-5} {0.0, 1e-5, 5e-5}
Token-GNN Architecture GNN Layers (K) GNN Hidden Dimension Jumping Knowledge Input Projection Dimension	{2, 4} {64, 128, 256} {cat, max, mean, none} {128, 256, 512}
Graph Construction Similarity Threshold (τ)	{0.1, 0.3, 0.6}

Baseline Pooling Methods. We implemented all baselines within the same framework and evaluated them to ensure a fair comparison.

• Static Methods: MEAN and MAX pooling operate over the non-padded token hidden states. [CLS]/[EOS] pooling for encoder models takes the hidden state of the first token. For decoder models, it takes the hidden state of the last non-padded token, identified via the attention mask.

• AdaPool: Our implementation follows the original paper (Brothers, 2025), consisting of a two-layer MLP with a Tanh activation that computes a scalar score for each token, followed by a softmax and weighted average.

GLOT Configuration. Our GLOT is implemented using 2 layers of GATCONV (Veličković et al., 2018) with a hidden dimension of 128 and ReLU non-linearity (Nair & Hinton, 2010). As described in the main paper, the graph is constructed by creating edges between tokens where their cosine similarity exceeds a threshold of $\tau=0.6$. Following the GNN layers, we use a 'cat' mode for Jumping Knowledge (Xu et al., 2018) to aggregate features from all layers before the final attention readout.

B.3 BENCHMARK-SPECIFIC DETAILS

- **GLUE Benchmark.** For all tasks from the General Language Understanding Evaluation (GLUE) benchmark (Wang et al., 2018), we fine-tune the lightweight GLOT head and a task-specific linear classifier jointly on the training sets. Sequences are truncated to a maximum length of 128 tokens. For larger datasets (QQP, QNLI, MNLI), we train on a subsample of 20,000 examples.
- **CoLA:** The Corpus of Linguistic Acceptability (Warstadt et al., 2018) requires the model to determine if a sentence is grammatically correct. **Task:** Binary classification. **Loss:** Cross-Entropy Loss.
- **SST-2:** The Stanford Sentiment Treebank (Socher et al., 2013a) consists of movie reviews. **Task:** Binary sentiment classification (positive/negative). **Loss:** Cross-Entropy Loss.
- **STS-B:** The Semantic Textual Similarity Benchmark (Agirre et al., 2007) involves predicting a similarity score between 1 and 5 for a pair of sentences. **Task:** Regression. **Loss:** Mean Squared Error (MSE) Loss.
- **MRPC:** The Microsoft Research Paraphrase Corpus (Dolan & Brockett, 2005) contains sentence pairs. **Task:** Binary classification to determine if the sentences are paraphrases. **Loss:** Cross-Entropy Loss.
 - **QQP:** The Quora Question Pairs dataset requires determining if two questions are semantically equivalent. **Task:** Binary classification. **Loss:** Cross-Entropy Loss.
- **MNLI:** The Multi-Genre Natural Language Inference corpus (Williams et al., 2018) provides a premise and a hypothesis. **Task:** Three-class classification (entailment, contradiction, neutral). **Loss:** Cross-Entropy Loss.
- **QNLI:** The Question Natural Language Inference dataset, derived from SQuAD (Rajpurkar et al., 2016). **Task:** Binary classification to determine if a context sentence contains the answer to a question. **Loss:** Cross-Entropy Loss.
- **RTE:** The Recognizing Textual Entailment datasets (Dagan et al., 2006; Bar Haim et al., 2006; Giampiccolo et al., 2007; Bentivogli et al., 2009). **Task:** Binary classification to determine if a premise entails a hypothesis. **Loss:** Cross-Entropy Loss.

Long-Text Classification (IMDB). For the IMDB Large Movie Review dataset (Maas et al., 2011), sequences were truncated to a maximum length of 512 tokens. The dataset contains paragraph-length movie reviews. **Task:** Binary sentiment classification. **Loss:** Cross-Entropy Loss.

MTEB Evaluation. For the Massive Text Embedding Benchmark (MTEB) (Muennighoff et al., 2023), we follow a two-stage process. First, the learnable pooling heads are trained on a large-scale retrieval dataset, and then they are evaluated in a zero-shot setting on the downstream MTEB tasks.

- Training Stage: All learnable heads were trained on the MS MARCO (Bajaj et al., 2016) passage ranking dataset. This involves predicting relevant text passages for a given query.
 Task: Passage retrieval. Loss: A symmetric in-batch contrastive loss with a temperature of 0.07.
- **Zero-shot Evaluation Stage:** The trained encoders are then evaluated on the following seven tasks without any further fine-tuning:
 - EmotionClassification: A multi-class classification task on tweets.

```
918
         Algorithm 2 Synthetic Diagnostic Dataset Generation
919
          Require: N: Number of samples to generate.
920
         Require: L: Total sequence length of each sample.
921
          Require: d_r: The target distractor ratio (e.g., 0.2, 0.5, 0.8, 0.9).
922
          Require: \mathcal{T}: A set of signal phrase templates, each with an associated label (e.g., '("...has [X]) but
923
              not [Y]", 0)').
924
          Require: V_D: A large vocabulary of distractor words.
925
           1: function GENERATE DIAGNOSTIC DATA (N, L, d_r, \mathcal{T}, \mathcal{V}_D)
926
                                                                                         ▶ Initialize an empty dataset
           3:
                  L_D \leftarrow |L \times d_r|
927

    ▷ Calculate number of distractor tokens

           4:
                  L_S \leftarrow L - L_D

    ▷ Calculate number of signal tokens

928
           5:
                  for i = 1 to N do
929
           6:
                       (template, label) \leftarrow RandomChoice(\mathcal{T})
930
           7:
                       signal\_tokens \leftarrow Instantiate(template)
                                                                            ▷ e.g., fill placeholders like [X] and [Y]
931
           8:
                                                                     ⊳ Ensure signal phrase fits the allocated length
932
           9:
                       if length(signal\_tokens) > L_S then
933
          10:
                           signal\_tokens \leftarrow signal\_tokens[: L_S]
                                                                                                934
          11:
                       else
935
          12:
                           padding \leftarrow L_S - \text{length}(signal\_tokens)
936
                           signal\_tokens \leftarrow concat(signal\_tokens, Sample(V_D, padding))
          13:
                                                                                                             ⊳ Pad with
937
              distractors if too short
                       end if
          14:
938
                       distractor\_tokens \leftarrow Sample(\mathcal{V}_D, L_D)
                                                                              ▶ Sample distractors with replacement
          15:
939
          16:
                       p_{inject} \leftarrow \text{RandomInt}(0, L_D)
                                                                                  940
          17:
                       sequence \leftarrow \mathbf{concat}(distractor\_tokens[:p_{inject}], signal\_tokens, distractor\_tokens[p_{inject}:
941
              ])
942
          18:
                       \mathcal{D} \leftarrow \mathcal{D} \cup \{(\text{sequence}, \text{label})\}
943
          19:
                  end for
944
          20:
                  return \mathcal{D}
945
          21: end function
946
```

- SciFact: A re-ranking task to verify scientific claims.
- RedditClustering: An unsupervised task to cluster Reddit comments.
- AskUbuntuDupQuestions: A retrieval task to find duplicate questions.
- STS12: A semantic similarity regression task.
- TwitterSemEval2015: A pair classification task for paraphrase detection.
- SummEval: A summarization evaluation task based on semantic similarity.

B.4 DIAGNOSTIC TASK GENERATION

947 948

949

950

951 952

953

954

955 956

957 958

959 960

961

962 963

964

965

966

967

968

969 970

971

The synthetic diagnostic task was created to isolate and test for relational understanding under noise.

- **Signal Phrases:** We created a small set of template phrases involving a logical dependency, such as negation (e.g., "The file has [X] but not [Y]").
- **Distractors:** The "haystack" was formed by sampling words randomly from a large general-purpose vocabulary derived from English Wikipedia.
- **Injection:** For each example, a 256-token sequence of random distractor words was generated. A signal phrase was then injected at a random position within this sequence.
- **Difficulty Control:** The difficulty was controlled by the **distractor ratio**, which we varied from 20

The final dataset consists of 10,000 training examples and 2,000 test examples for each distractor ratio.

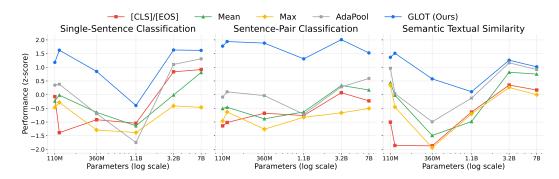


Figure 4: **Z-score normalized performance on the GLUE benchmark, aggregated by task category.** Performance, represented as a z-score, is plotted against the number of parameters in the frozen backbone model (log scale). A higher z-score indicates better relative performance compared to the average of all tested methods for that setting.

C ADDITIONAL RESULTS

GLUE Benchmark. To provide a high-level summary of the comprehensive GLUE results presented in Table 1, we visualize the performance trends in Figure 4. To compare performance across different tasks and their associated metrics (e.g., Accuracy vs. MCC) on a single, unified scale, we normalized the scores. The methodology was as follows: for each of the eight GLUE tasks and for each of the six backbone models, we took the resulting scores of all five pooling methods and calculated their mean (μ) and standard deviation (σ) . Each individual score x was then converted to a z-score via $z = (x - \mu)/\sigma$. These z-scores were then averaged within their respective task categories. A higher z-score indicates that a method's performance is significantly above the average of all tested methods for a given experimental setting. The plots clearly show that GLOT consistently achieves the highest z-score, often one or more standard deviations above the mean performance. This visualization powerfully reinforces our primary finding: the performance advantage of GLOT is not confined to specific tasks or model scales but is a robust and general phenomenon.

C.1 DIAGNOSTIC TASK: DETAILED RESULTS AND ANALYSIS

To provide a controlled evaluation of relational robustness under noise (**RQ4**), we designed a synthetic diagnostic task. Inspired by 'signal-in-noise' evaluations (Brothers, 2025) and the 'Needle in a Haystack' paradigm (Kamradt, 2023), our stress test is specifically adapted to probe for relational understanding rather than simple factual recall. We programmatically generate sequences by injecting a short "signal phrase" with a logical dependency (e.g., negation) into a long sequence of random distractor words. The task is a binary classification based on the logic within the signal phrase. We systematically increase the task's difficulty by increasing the distractor ratio from 20% to 90%. The full generation process is detailed in Algorithm 2.

The complete results for this stress test are presented in Table 7. The data provides a clear and quantitative confirmation of our hypothesis: GLOT's performance remains remarkably stable even at extreme noise levels, while the performance of all baseline methods degrades significantly as the signal is diluted.

This trend is consistent across all architectures. For the encoder-only BERT backbone, GLOT's accuracy remains consistently above 97% across all distractor ratios. In contrast, the next-best baseline, AdaPool, sees its performance drop sharply from 91.4% at 20% distractors to just 61.6% at 90% distractors. The pattern is mirrored in decoder-only models. With the Llama backbone, GLOT's accuracy is nearly perfect at low noise (99.6%) and stays high at 83.2% even at the extreme 90% distractor ratio. All other methods, including using the standard '[EOS]' token, see their performance collapse, with most falling to near-chance levels. This analysis demonstrates that by explicitly modeling token relationships, GLOT can reliably identify and reason over the crucial signal phrase, whereas methods that rely on global summary statistics are overwhelmed by the distractor tokens.

Table 7: **Full results for the diagnostic stress test**, which evaluates robustness to signal dilution. The table reports the classification accuracy for all pooling methods across six backbones as the ratio of distractor tokens in the input sequence increases from 20% to 90%. The **best** results for each model are in bold.

Model	Method	20% Distractors	50% Distractors	80% Distractors	90% Distractors
	[CLS]	70.8	58.2	57.2	67.6
Η	Mean	68.0	58.6	64.2	53.4
BERT	Max	57.4	50.8	51.6	50.4
8	AdaPool	91.4	78.8	65.6	61.6
	GLOT	97.2	97.0	97.8	98.8
	[CLS]	83.6	63.4	51.6	48.6
RoBERTa	Mean	73.2	64.6	67.8	57.2
豆	Max	56.8	60.0	59.0	50.2
OB	AdaPool	83.0	67.2	59.8	59.2
×	GLOT	92.6	99.2	98.8	98.2
	[CLS]	72.0	57.6	58.6	51.4
Ξ	Mean	70.2	56.2	54.6	51.4
SmolLM2	Max	54.0	50.6	46.2	51.4
Ĕ	AdaPool	78.2	57.6	54.2	55.2
Ø	GLOT	96.0	93.6	92.4	92.2
	[EOS]	73.2	58.8	57.4	56.6
TinyLlama	Mean	83.4	56.6	56.8	56.4
Ę	Max	76.4	54.0	58.8	51.4
, i	AdaPool	78.4	66.4	57.4	53.0
F	GLOT	96.4	94.8	88.4	94.0
<u> </u>	[EOS]	84.4	69.8	69.0	68.4
LLaMA-3B	Mean	82.4	65.0	64.4	61.8
¥	Max	63.0	50.8	50.6	54.6
द्धि	AdaPool	92.6	69.8	70.0	51.0
\Box	GLOT	99.6	95.4	89.8	93.2
8	[EOS]	89.4	72.6	69.4	70.6
<u>-</u> -	Mean	93.0	74.0	71.4	63.8
Ţ <u>a</u>	Max	60.8	60.0	57.4	55.6
Mistral-7B	AdaPool	92.2	86.6	85.0	78.4
Σ	GLOT	99.2	97.6	97.4	97.2