

VECTOR-ICL: IN-CONTEXT LEARNING WITH CONTINUOUS VECTOR REPRESENTATIONS

Anonymous authors
Paper under double-blind review

ABSTRACT

Large language models (LLMs) have shown remarkable in-context learning (ICL) capabilities on textual data. We explore whether these capabilities can be extended to continuous vectors from diverse domains, obtained from black-box pre-trained encoders. By aligning input data with an LLM’s embedding space through lightweight projectors, we observe that LLMs can effectively process and learn from these projected vectors, which we term Vector-ICL. In particular, we find that pretraining projectors with general language modeling objectives enables Vector-ICL, while task-specific finetuning further enhances performance. In our experiments across various tasks and modalities, including text reconstruction, numerical function regression, text classification, summarization, molecule captioning, time-series classification, graph classification, and fMRI decoding, Vector-ICL often surpasses both few-shot ICL and domain-specific model or tuning. We further conduct analyses and case studies, indicating the potential of LLMs to process vector representations beyond traditional token-based paradigms.

1 INTRODUCTION

In-context learning (ICL) has emerged as a powerful paradigm in large language models (LLMs), allowing generalization from limited examples within a given context (Brown et al., 2020; OpenAI, 2023). By providing demonstrations in the context during inference, ICL allows models to adapt to new tasks and formats without the need for retraining. However, since LLMs are trained on discrete natural language tokens, ICL is generally learned and used through natural language, limiting its applicability to non-textual data.

We explore whether LLMs can perform ICL directly on continuous vectors, a capability that could dramatically expand their applicability. Many data modalities, such as sensor readings, financial time series, or scientific measurements, lack a natural text representation. Moreover, even for text data, information like numbers might be better represented via continuous vectors than tokens.

In our study, we observe that LLMs can indeed understand and process continuous context via embedding projection. This technique, which we term Vector-ICL, acts as a bridge between continuous data and the LLM’s embedding space. Simple linear projections are often sufficient, though for cross-modal tasks—such as those involving non-textual data like time-series or graphs, non-linear transformations may be required. We demonstrate that training the embedding projector using a straightforward next-token prediction objective enables Vector-ICL, effectively teaching the LLM to “read” continuous vectors. Moreover, fine-tuning the projector on downstream tasks further enhances the effectiveness of continuous context, outperforming few-shot ICL and domain-specific models or tuning.

Our investigation begins with the task of text reconstruction, where we assess whether LLMs can recover information encoded in text embedding. This serves as a proof-of-concept for Vector-ICL, showing that LLMs can indeed extract meaningful information from projected continuous vectors. We then turn to the more complex challenge of arithmetics. Although state-of-the-art LLMs can solve Olympiad mathematical problems (Trinh et al., 2024; OpenAI, 2023), they struggle with precise number processing due to the limitations inherent in their tokenization schemes. Our results demonstrate that Vector-ICL offers a more effective approach for function approximation, particu-

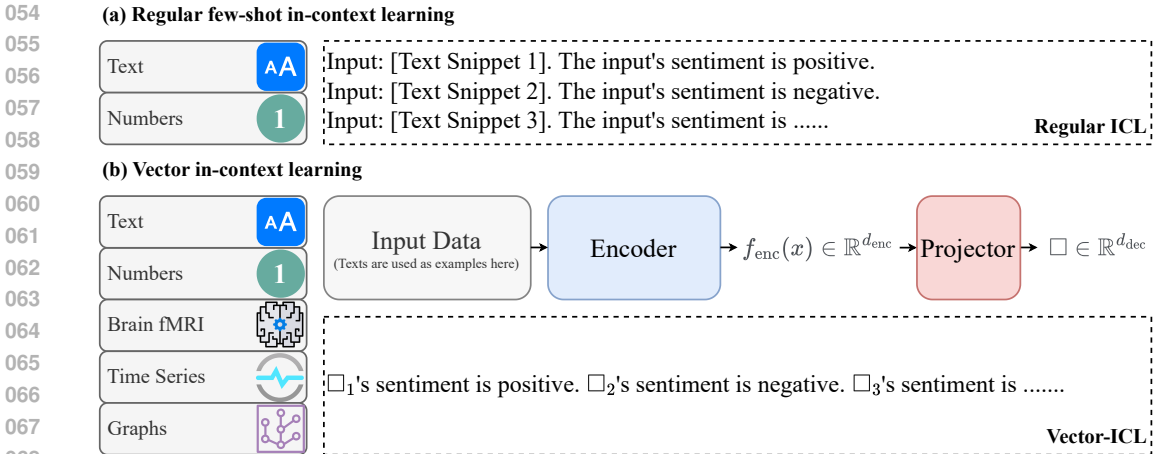


Figure 1: **Comparing regular in-context learning to vector in-context learning.** (a) In regular ICL, textual demonstrations are given as context during LLM inference. (b) In Vector-ICL, the input space is extended across multiple modalities. The input data is first encoded as embeddings, then transformed into continuous vectors which represent as box tokens (\square) via embedding projection. During inference, we provide box tokens in prompts as demonstrations for ICL. We consider box tokens representing text, numerical data, brain fMRI, time series, and graphs in this study.

larly for large numbers that span multiple tokens, potentially opening new avenues for enhancing LLMs’ numerical reasoning capabilities.

Finally, we extend our analysis to a broad range of modalities and tasks, including text classification, summarization, molecule captioning, brain fMRI reconstruction and classification, time-series classification, and graph classification. Across these diverse domains, LLMs exhibit competitive and often superior performance when employing Vector-ICL, revealing previously untapped capabilities of these models. This work highlights the potential of continuous representations in enhancing LLMs’ in-context learning capacities, pushing the boundaries of what these models can achieve beyond token-based paradigms.

2 RELATED WORK

Empirical results of in-context learning ICL has empirically shown strong performance in diverse natural-language processing tasks with very few demonstrations (Brown et al., 2020; OpenAI, 2023). In modern LLMs with long context windows, ICL has even shown performance improvements as the number of demonstrations grows to hundreds or even thousands, sometimes outperforming finetuning (Agarwal et al., 2024; Li et al., 2023a; Bertsch et al., 2024). Empirically, different factors play key roles in ICL. In smaller LLMs, ground-truth demonstrations are not required for in-context learning, while other factors such as the label space, input text distribution, and overall sequence format play an important role (Min et al., 2022b). Moreover, these LLMs can sometimes achieve strong performance even when demonstrations are intentionally irrelevant or even pathologically misleading (Webson & Pavlick, 2022). [Flan-T5 \(Chung et al., 2024\) revealed that instruction tuning improves few-shot learning by helping LLMs better utilize in-context examples in an encoder-decoder architecture.](#) More recently, Wei et al. (2023) characterize these behaviors of LLMs with respect to model size, and show that larger language models perform in-context learning differently in the presence of flipped or semantically unrelated labels. Orthogonally, different works find ways to improve ICL, e.g. by including explanations (Lampinen et al., 2022), or chaining ICL calls (Morris et al., 2023). ICL has shown some success in multimodal models (Wu et al., 2024; Jiang et al., 2024) or when applied to tabular data (Zhao et al., 2024).

Understanding ICL Many works have investigated ICL and found that it is able to learn linear models (Akyürek et al., 2022; Zhang et al., 2023), discrete functions (Bhattamishra et al., 2023), and more general algorithms (Li et al., 2023b). Some works have explicitly connected ICL in specific

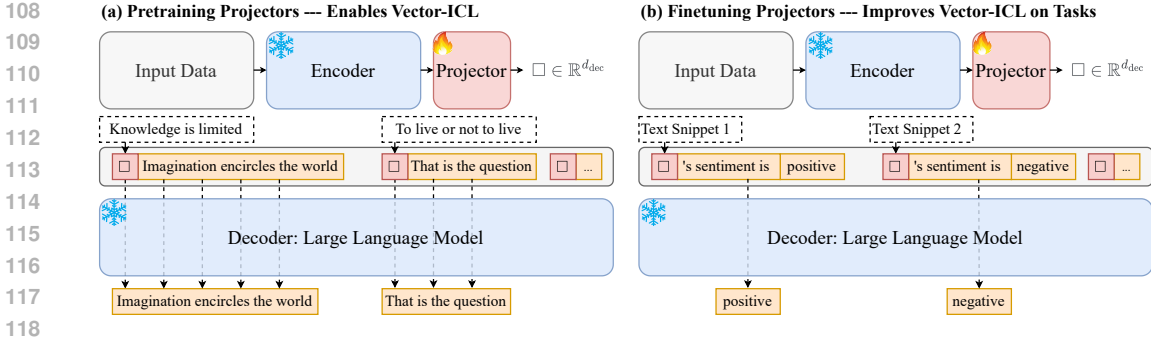


Figure 2: **Pretraining and finetuning the projectors.** Vector-ICL requires updating the parameters of a lightweight projector while keeping the encoder and decoder parameters fixed. The encoder first compresses the input into single token embeddings, and then the projector will project it to the aligned representation space for LLMs’ later use. (a) Pretraining the projector on a general language modeling corpus (or a modality-to-text dataset) enables Vector-ICL. (b) Task-specific fine-tuning makes Vector-ICL outperform few-shot ICL on natural language tasks, as well as with domain-specific models on non-language tasks.

settings to implementing optimization steps analogous to gradient descent (Mahankali et al., 2023; Von Oswald et al., 2023; Ahn et al., 2024) and higher-order optimization methods (Dai et al., 2023; Fu et al., 2023; Giannou et al., 2024; Zhang et al., 2023). A complementary direction aims to establish statistical complexity and generalization bounds of in-context learning in transformers (Bai et al., 2024; Li et al., 2023b; Wies et al., 2024; Wu et al., 2023). Finally, one recent work suggests that ICL may arise from parallel structures in pretraining data (Chen et al., 2024).

Learning to learn in-context In contrast to the emergent ICL capabilities of LLMs, existing works have also studied how to explicitly improve ICL. Min et al. (2022a) propose MetaCL, a meta-training framework for finetuning pretrained LLMs to perform in-context learning on a large and diverse collection of tasks. In the tabular domain, TNP (Nguyen & Grover, 2022) and PFNs (Müller et al., 2021) train transformer models to perform in-context prediction for a family of functions, which allows in-context generalization to unseen functions after training. Zhao et al. (2023) also propose meta-learning transformers to in-context learn group preferences, serving as an in-context learned reward model that adapts to diverse group preferences.

3 METHOD: VECTOR CONTEXT VIA EMBEDDING PROJECTION

3.1 EMBEDDING PROJECTION

Vector-ICL requires transforming inputs into vector contexts through an embedding projection. Given a dataset $\mathcal{X} = \{x_i\}_{i=1}^n$, we assume the existence of an encoder f_{enc} , that transforms the data into an abstract representation (alternatively, the raw data may already be a continuous vector). The encoded embeddings, $f_{enc}(x)$, are then projected into box tokens, denoted as \square_x . Throughout the paper, we will use the terms “box tokens” and “projected embeddings” interchangeably. For decoding, we use a language model f_{llm} that generates outputs based on the provided prompts.

We impose no constraints on the form of the input data \mathcal{X} ; it can come from any modality. The only requirement is that the encoder f_{enc} maps each data point x into a vector space, defined as:

$$f_{enc} : x \rightarrow \mathbb{R}^{d_{enc}}, \forall x \in \mathcal{X} \tag{1}$$

The LLM typically processes discrete tokens $\{\text{tok}_1, \text{tok}_2, \dots, \text{tok}_l\}$, then maps them to text embedding space $\{\text{emb}_1, \text{emb}_2, \dots, \text{emb}_l\}$, $\forall i, \text{emb}_i \in \mathbb{R}^{d_{dec}}$. Since we operate mostly in embedding space, we omit the tokenization step for simplicity and directly refer to text inputs as their embedding representations.

The process of embedding projection is then carried out as follows. For linear projection, we construct a projection matrix $P \in \mathbb{R}^{d_{enc} \times d_{dec}}$ and make the following transformations to obtain pro-

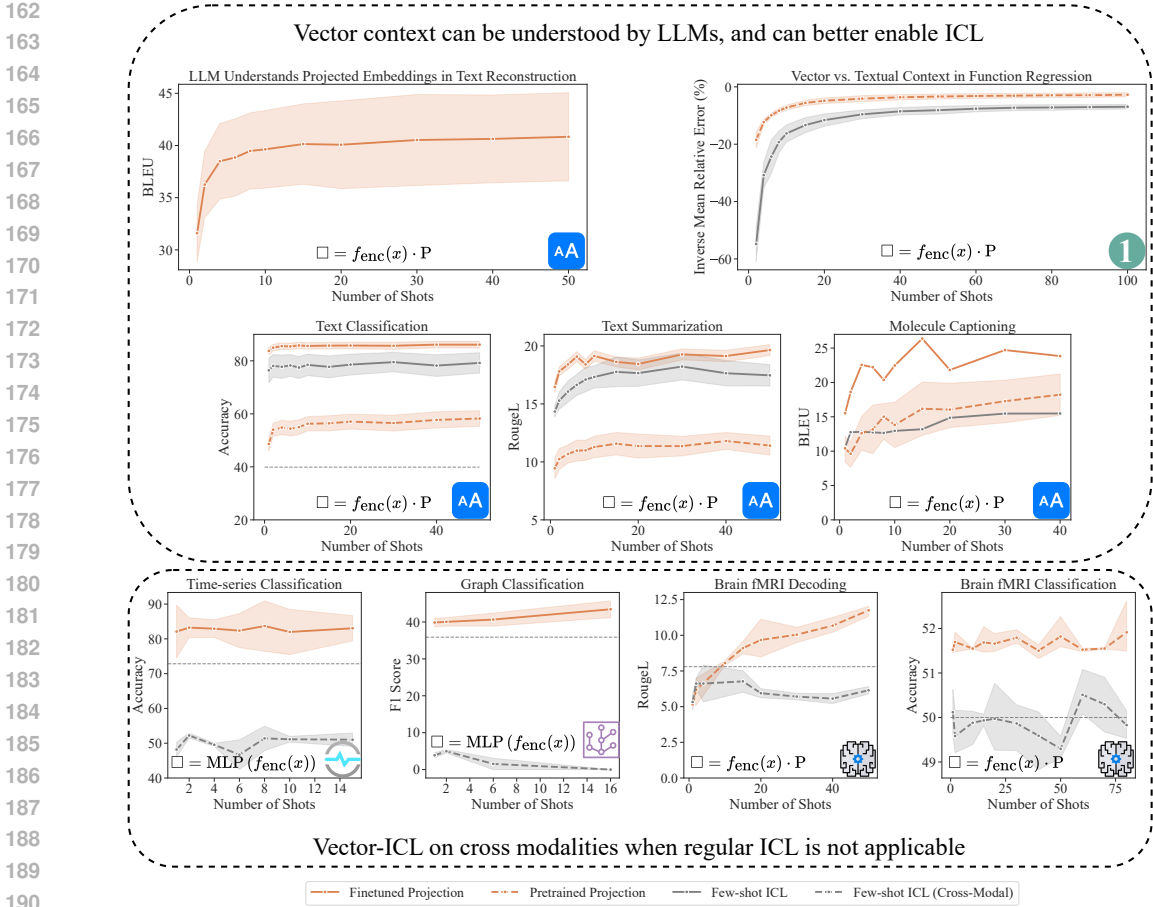


Figure 3: **Main results: LLMs can perform Vector-ICL (\uparrow = better).** We show that training the embedding projector with a simple next-token prediction objective enables Vector-ICL. Even with only unsupervised pretraining, Vector-ICL matches or outperforms traditional few-shot ICL on 4 out of 6 tasks where direct comparison is possible. Fine-tuning the projector on downstream tasks further enhances the use of continuous context, consistently surpassing both few-shot ICL and specialized task-tuned baselines (soft-prompt for text, tuned encoders for non-text). The study begins with text reconstruction to assess LLMs’ ability to interpret box token embeddings, followed by function regression to evaluate reasoning capabilities. We then demonstrate Vector-ICL’s effectiveness and applicability across various downstream tasks, including text classification, summarization, time-series classification, graph classification, and brain fMRI decoding & classification. Results in each panel are averaged over different encoders and LLMs for the diverse tasks we study; error bars show 95% confidence intervals.

jected embedding \square_x given input x :

$$\square_x := f_{\text{enc}}(x) \cdot P \tag{2}$$

In cases where more expressive power is needed, we utilize a two-layer multi-layer perceptron (MLP) to perform the projection:

$$\square_x := \text{MLP}(f_{\text{enc}}(x)) \tag{3}$$

The MLP follows the architecture of the MLP block in Llama (Touvron et al., 2023), with an additional input projection layer to map the input dimension $\mathbb{R}^{d_{\text{enc}}}$ to output dimension $\mathbb{R}^{d_{\text{dec}}}$.

3.2 PROJECTED EMBEDDINGS AS CONTEXT

The projected embeddings are then utilized as context in Vector-ICL, functioning as the equivalent of the original input data. For example, in NLP tasks, the original text snippets x will first be encoded as embeddings $f_{\text{enc}}(x)$, then projected to become \square_x , inserted into the prompt like the following:

\square_x 's sentiment is ... / \square_x 's summarization is ...

Using them as the context in ICL is then natural:

\square_1 's sentiment is happy. \square_2 's sentiment is sad. \square_x 's sentiment is ...

where \square_1 and \square_2 are in-context examples and \square_x is the input.

3.3 TRAINING THE EMBEDDING PROJECTORS

The projectors need to be trained to achieve effective projections. We discovered that pretraining these projectors with language modeling objectives enables the ICL capabilities with vector context, and finetuning them on task datasets further improves ICL performance.

The pretraining process is depicted in Fig. 2(a). For each text snippet, we cut it into two pieces with the cutting point randomly sampled from the end of sentences. The first half is encoded and projected while the second is kept intact. The rest is the same with any pretraining process, the language model generated the next token distribution at each input position, except for the ones preceding the projected embeddings, and a cross-entropy loss is imposed on top of this. With the encoder and LLM frozen, the gradient backpropagates to the projector, updating its parameters.

For non-text data modalities, pretraining can be more flexible. We define this pretraining as involving general, non-task-specific objectives, such as reconstructing a number from its embeddings (e.g., \square_x is 32768), performing basic algebra (e.g., $\square_x + \square_y = \underline{16384}$), or predicting the next token from brain fMRI embeddings.

The finetuning process is shown in Fig. 2(b). It utilizes additional structured prompts and trains with task-specific datasets. Similarly, the input is first mapped into the embedding space and projected into \square tokens. They are then inserted into structured prompts, while the **projector** is trained with conditional generation loss given those prompts.

4 EXPERIMENTAL SETUP

Table 1 gives an overview of our experimental setup, including specifics for the task, datasets, encoders, LLMs, and task-specific prompts we use. Across different tasks, we project to four open-weights LLMs. We now provide details for individual tasks.

Baselines We evaluate Vector-ICL in two distinct settings: with and without task-specific tuning. For textual tasks, we compare against few-shot ICL and soft prompt tuning (Li & Liang, 2021). We choose soft prompt tuning as our primary baseline because it represents a similarly lightweight adaptation approach - both methods introduce a small number of trainable parameters while keeping the base LLM frozen. Like our projector, soft prompts modify how the LLM processes inputs without changing its internal weights. This makes it a fair comparison point for assessing whether Vector-ICL's benefits come from the continuous vector representations themselves rather than just the additional training. For non-textual tasks, where soft prompts cannot be directly applied, we compare against cross-modal few-shot ICL and tuned encoders. In the non-textual domain, ICL inputs are represented either as numeric sequences (for time-series and brain fMRI data) or as textual descriptions (for graph edge lists and node features).

Text Pretraining To pretrain our text projectors, we leverage the WikiText-103 (Merity et al., 2016) dataset, consisting of over 100 million tokens from verified high-quality Wikipedia articles. This smaller language modeling corpus is chosen for its suitability to the lightweight nature of our projectors. The pretraining process is illustrated in Fig. 2(a), where text snippets are divided at random sentence-end points. The first half is embedded and projected, while a next-token generation loss is applied to the second half.

Table 1: **Overview of the tasks, datasets, encoders, large language models, and prompts used in the experiments.** Each task utilizes encoders and LLMs to perform functions across multiple modalities. The table highlights the diversity of models and configurations applied to each task.

Task	Dataset	Encoder	LLM	Prompt
Text Reconstruction	Quora (Thakur et al., 2021) PST (Tiedemann, 2012)	N, S, E, G	🦙	Translate the text in brackets: (\square) Translation: [Text]
Function Regression	10 Digits Regression	Digit Embedding	🦙, 🌐, 🏆, 🗂️	$x = \square_x, y = \square_y$, function(x, y) equals to (digits): [Solution]
Text Classification	IMDB (Maas et al., 2011) Rotten Tomatoes (Pang & Lee, 2005) SST2 (Socher et al., 2013) Emotion (Saravia et al., 2018) Financial Phrasebank (Malo et al., 2014)	N, S, E, G	🦙, 🌐, 🏆	(\square)’s sentiment is: [Label]
Text Summarization	XSum (Narayan et al., 2018) XLSum (Hasan et al., 2021)	N, E, G	🦙	(\square)’s summarization is: [Summary]
Molecule Captioning	Language + Molecule-24 (Edwards et al., 2024)	N	🦙	(\square)’s molecule caption is: [Caption]
Brain fMRI	LeBel et al. 2022, Tang et al. 2023	PCA of fMRI	🦙	[Question] Input: \square Response: [Answer]
Time-series Classification	FordA, FordB (Dau et al., 2019)	Chronos-base (Ansari et al., 2024)	🦙	(\square)’s class (positive, negative) is: [Label]
Graph Classification	ogbg-molhiv (Hu et al., 2020)	Graphormer (Ying et al., 2021)	🦙, 🏆	(\square)’s class (positive, negative) is: [Label]

Encoders: **N** NV-Embed (Lee et al. 2024; nvidia/NV-Embed-v1), **S** SFR (Meng et al. 2024; salesforce/SFR-Embedding-2_L), **E** Stella (Zhang 2024; dunzhang/stella.en.1.5B.v5), **G** GTR-t5 (Ni et al. 2021; sentence-transformers/gtr-t5-base)
LLMs: 🦙 Llama-3.1-8B (Dubey et al. 2024; meta-llama/Llama-3.1-8B-Instruct), 🌐 Mistral-7B (Jiang et al. 2023; mistralai/Mistral-7B-Instruct-v0.3), 🏆 Qwen2-7B (Yang et al. 2024; Qwen/Qwen2-7B-Instruct), 🗂️ Yi-1.5-9B (Young et al. 2024; 01-ai/Yi-1.5-9B-Chat)

Text Reconstruction We investigate LLMs’ ability to decode original text from projected embeddings using two datasets: Parallel Sentence Talks (Tiedemann, 2012) and Quora (Thakur et al., 2021). These datasets consist of concise text pieces that convey clear meaning. Projectors are trained on the training sets of both datasets, and performance is evaluated using the BLEU score (Papineni et al., 2002; Post, 2018).

Arithmetic and Function Regression For the arithmetic tasks, we generated synthetic datasets containing 10-digit numbers, which are particularly challenging for LLMs as they require splitting the numbers into multiple text tokens. These numbers are represented using a concatenated one-hot encoding per digit. For instance, a 10-digit number is represented as a 10×10 matrix, flattened into a 100-dimensional vector. The pretraining phase includes two key tasks: number reconstruction, where the model is tasked with recovering the original number from its embedding, and basic arithmetic, where the model performs algebraic addition operations on the projected embeddings.

To evaluate the models’ arithmetic reasoning abilities, we employ a non-linear function regression task, where the function is defined as $f(x, y) = \sqrt{x}\sqrt{y}$. The model is provided with inputs x and y , and it must predict the integer part of the function output. Performance is measured using the mean relative error, calculated as the ℓ_1 difference between the predicted and true values, normalized by the ground truth. This task allows us to assess the models’ ability to perform more complex numerical reasoning beyond simple arithmetic operations.

Text Classification We assess whether Vector-ICL can be applied effectively to text classification. Both binary and multi-class classification datasets are used, and the results are compared across few-shot ICL and soft prompt tuning. The classification performance is measured by accuracy.

Text Summarization Following the classification tasks, we explore Vector-ICL’s capability in summarizing text based on the projected embeddings. The datasets, encoders, LLMs, and prompt templates can be found in Table 1. Performance is evaluated using RougeL (Lin, 2004).

Molecule Captioning We also extend our approach to the unconventional task of molecule captioning, using molecule sequence-caption pairs from the Language + Molecules-24 (LPM24) (Edwards et al., 2024) dataset. A sample molecule-caption pair looks like the following:

Table 2: Comparison of Vector-ICL, few-shot ICL, and soft prompt tuning across various sentiment analysis and summarization datasets. Details can be found in Appendix A.3.

Method	Sentiment Analysis					Summarization	
	Rotten Tomatoes	SST2	IMDB	Emotion	Financial Phrasebank	XSum	XLSum
V-ICL (pretrained)	80.60	78.90	95.04	41.20	60.72	15.25	15.89
Few-shot ICL	87.31	91.74	93.50	55.20	71.78	19.53	19.41
Soft Prompt	93.24	96.21	95.26	74.15	78.22	12.84	12.70
V-ICL (finetuned)	88.80	98.16	97.28	85.20	81.68	20.08	20.49

Molecule: Cc1c(Cl)cccc1-n1ccn2c(SCC(=O)c3ccccc3C(F)(F)F)nnc2c1=O

Caption: The molecule is a pain treatment that impacts inflammatory disease treatment.

This task explores whether LLMs can extract useful information from projected embeddings of out-of-distribution chemical sequences, with performance evaluated via BLEU score.

Brain fMRI Decoding and Classification We analyze data from LeBel et al. 2022 and Tang et al. 2023, which consists of fMRI responses for 3 human subjects as they listen to 20+ hours of narrative stories from podcasts. We preprocessed the data following Benara et al. 2024, by converting the fMRI responses into a 200-dimensional output using principal components analysis and assigning classification labels to 10-grams of the story text at 2-second intervals using an LLM.

We use the same pretraining methodology as text to pretrain on the brain fMRI data (projecting on 20% of time points and imposing next-token generation loss on the remaining 80%). We evaluate the LLM’s capability to decode projected brain fMRI by giving them randomly sampled context from the train set, that could come from different human subjects or from a different story, and ask them to decode segments from the test set.

In addition to text reconstruction, we decode the binary labels from the fMRI responses, corresponding to questions about the underlying text, e.g. “Does the sentence contain a proper noun?” The decoding random baseline is constructed by giving the LLM the randomly sampled, shuffled text from the training set, and generating text according to it. We measure the performance using the RougeL score between the generated text and the ground truth text. The classification random baseline is 50% accuracy, as we have balanced the dataset.

Time-series We take the output of the last time step from Chronos-base (Ansari et al., 2024) as the time-series’ representation. We use the base encoder with trained classification head as the baseline and we measure the prediction performance with accuracy.

Graphs We use Graphormer (Ying et al., 2021) as the encoder model, specifically the one that was pretrained on quantum chemistry graph datasets (Hu et al., 2021). Since the down-stream, ogbg-molhiv (Hu et al., 2020), is a molecule property prediction dataset, and with strong class imbalance (3-4% positive classes), we finetune the encoder on the training set to provide meaningful baselines and embeddings. We take the output prior to the classification layer of the Graphormer as the graph embedding. Weighted sampling is adopted in the finetuning of both the baseline Graphormer and the embedding projector to yield meaningful predictions. We use the finetuned Graphormer as our baseline and use the F1 score as the performance metric.

Projector Configurations Both linear and non-linear projectors are utilized, as shown in Fig. 3, with input and output dimensions matching the encoder-decoder pairs. Early stopping with patience of 500 steps is used during finetuning, as projectors converge quickly due to their small parameter sizes. Details of the hyperparameters used in training are provided in Table 3.

5 RESULTS: UNLOCKING VERSATILE APPLICATIONS ACROSS MODALITIES

Fig. 3 presents our main results, where each subplot corresponds to one of the nine tasks. We begin by exploring text reconstruction, to see whether LLMs can comprehend the information encoded

378 within the box tokens. Next, we investigate the tasks of function regression to evaluate whether
379 LLMs can leverage the box tokens during reasoning processes and whether this approach outper-
380 forms reasoning with plain text. Finally, we proceed to a range of downstream tasks, including text
381 classification, text summarization, time-series classification, graph classification, and brain fMRI
382 decoding. This comprehensive evaluation allows us to assess the versatility and effectiveness of
383 Vector-ICL across different domains and task types.

384
385 **Text Reconstruction: LLM Understanding of Projected Embeddings** We first verify LLMs’
386 ability to understand projected embeddings. Our results demonstrate that Vector-ICL successfully
387 reconstructs original text from projected embeddings, with performance improving as the number
388 of examples (shots) increases, mirroring standard in-context learning (ICL) behavior. This suggests
389 that LLMs can effectively decode the information compressed into the box tokens, with more context
390 leading to better reconstruction. The similarity to ICL behavior indicates that Vector-ICL leverages
391 similar learning mechanisms, but with the flexibility of working with continuous representations.

392
393 **Function Regression: Enhanced Reasoning with Continuous Context** We then study would it
394 sometimes be better for LLMs to receive continuous context instead of discrete tokens. After pre-
395 training projectors on 10-digit number reconstruction and addition between two numbers, we task
396 the LLM with learning an unknown function in-context. Results show Vector-ICL consistently out-
397 performs few-shot ICL with raw number inputs, that have to span multiple tokens. This suggests
398 that the continuous representations capture numerical relationships more effectively than discrete
399 tokens, enabling LLMs to better infer and apply mathematical patterns. The improvement is partic-
400 ularly noteworthy given that LLMs are typically challenged by precise numerical computations.

401
402 **Text Classification** In this classical NLP task, we aggregate mean accuracy across five datasets,
403 four encoders, and three LLMs. Results indicate that pretrained projectors provide meaningful con-
404 tinuous context for ICL, outperforming the random baseline. LLMs achieve optimal performance
405 with continuous context from finetuned projectors, surpassing both regular few-shot ICL and soft
406 prompt tuning, with details shown in Table 2. This demonstrates the versatility of Vector-ICL across
407 different text classification scenarios and its ability to outperform established prompt tuning meth-
408 ods. The success across multiple datasets and LLMs suggests that the benefits of continuous context
409 are robust and generalizable.

410
411 **Text Summarization** This task demands longer text generation and deeper information compre-
412 hension. We aggregate mean RougeL scores across two datasets, and three encoders. Findings show
413 that pretrained projectors enable LLMs to extract and condense information from a single box token,
414 while finetuned projectors provide more effective continuous context than original textual input and
415 soft prompts, with details shown in Table 2. The ability to compress and later expand information
416 from a single token is particularly impressive, suggesting that Vector-ICL captures high-level se-
417 mantic content effectively. The superior performance of finetuned projectors highlights the benefits
418 of task-specific optimization in continuous space.

419
420 **Molecule Captioning: An Unconventional NLP Task** We explore LLMs’ ability to compre-
421 hend continuous vector context for out-of-distribution inputs like molecule sequences. Evaluating
422 captioning performance with BLEU scores, we find that both pretrained and finetuned projectors
423 provide better context than original molecule sequence text, despite encoders likely never trained
424 on such sequences. This result is particularly intriguing as it demonstrates Vector-ICL’s ability to
425 bridge the gap between specialized domains and general language understanding. It suggests that
426 continuous representations can capture and translate domain-specific information in a way that’s
427 more accessible to LLMs than raw specialized notation.

428
429 **Time-series Classification** We finetune non-linear projectors on the training sets of two datasets
430 and evaluate LLM performance with the resulting continuous context. Aggregating average accu-
431 racy, we find LLMs outperform baseline domain-specific models with finetuned classification heads.
The success here suggests that continuous context can effectively capture temporal dependencies and
patterns, translating time-series data into a form that LLMs can process effectively.

Graph Classification Since the task dataset is heavily imbalanced and out of the pretraining distribution of the graph encoder. We first finetune the base encoder on the target dataset to establish a baseline, then train non-linear projectors on the graph classification dataset. Averaging F1 scores across two LLMs, results indicate that Vector-ICL enables LLMs to outperform the finetuned baseline model. This is a noteworthy achievement, as graph data is structurally very different from the text data that LLMs are trained on. The success here suggests that Vector-ICL can effectively translate graph structures into continuous representations that preserve relational information in a way that’s interpretable to LLMs.

Brain fMRI Decoding and Classification We pretrain projectors on the training set using next-token generation loss, then apply them to recover original text from brain fMRI signals in the test set. Results show LLMs can surpass random baselines by leveraging projected 200-dimensional fMRI PCA factors, with performance improving as context increases. This application to neuroscience data is particularly exciting, demonstrating Vector-ICL’s potential in bridging neural activity and language understanding.

6 ANALYSIS

6.1 THE IMPACT OF ENCODER QUALITY ON VECTOR-ICL PERFORMANCE

We investigate the relationship between the intrinsic capabilities of encoders and their effectiveness when used with Vector-ICL for downstream tasks. To quantify the encoders’ intrinsic abilities, we evaluate their performance on a text reconstruction task, which serves as a proxy for the amount of information preserved in the embeddings.

Our analysis focuses on text classification as the downstream task. We examine the correlation between encoder rankings on the reconstruction task and their corresponding rankings on the classification task. This analysis is performed across 5 datasets and 3 LLMs, resulting in 15 configurations.

The results of this analysis, presented in Fig. 4a, demonstrate a consistent positive correlation between an encoder’s text reconstruction performance and its effectiveness in downstream classification tasks when used with Vector-ICL. Notably, in 4 of the 15 configurations, we observe a particularly strong correlation, with values approaching 1.

Our findings suggest that an encoder’s performance on the text reconstruction task can serve as a reliable predictor of its potential effectiveness in downstream tasks when integrated with Vector-ICL. This insight could prove valuable for practitioners in selecting encoders for Vector-ICL.

6.2 CASE STUDY: WHAT HAS BEEN LEARNED IN THE PROJECTIONS?

Fig. 4b provides a visualization of the normalized Euclidean distances between projected embeddings. Several key patterns emerge from this that offer insights into what the projector has learned.

Analysis of the numerical embedding distance matrix reveals key properties of our projection method. Embeddings for similar numbers cluster along the diagonal, indicated by lighter colors, demonstrating the preservation of local structure. Conversely, increasing distances from the diagonal, shown by darker colors, indicate effective separation of numerically distant values in the embedding space.

Another notable feature of the distance matrix is the block structure that emerges. [The block structure reflects how numbers share similar digit patterns across the decimal places, from local to global blocks. It likely helps LLMs process numerical relationships more effectively, as it preserves the hierarchical nature of place-value notation.](#)

6.3 CASE STUDY: CAN SYNTHETIC DATASET BE USEFUL IN CROSS-MODAL PRETRAINING?

The cross-modal pretraining requires data-text dual pairs, which is not available in many cases. We explore how to create a synthetic dataset for such pretraining using time-series as an example. We designed a data curation pipeline that extracts meaningful statistical properties from time-series and converts them into natural language descriptions.

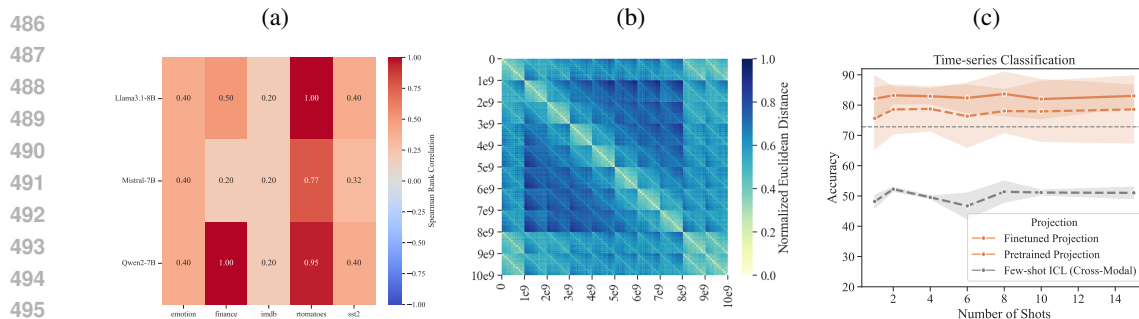


Figure 4: **Key Insights from Encoders, Projections, and Synthetic Data Curation.** (a) Correlation between encoders’ text reconstruction performance and their downstream task effectiveness with Vector-ICL, suggesting information preservation ability predicts Vector-ICL performance. (b) Euclidean distance matrix of 1024 projected number embeddings (0 to 1e10) shows structured block-diagonal patterns, indicating meaningful numerical relationships are preserved. (c) Results from pretraining on our synthetic time-series QA dataset, which captures statistical properties through trend analysis, anomaly detection, and stability assessment. The curated QA pairs enable effective cross-modal pretraining.

Our pipeline analyzes multiple statistical aspects of the time-series data. We perform trend analysis through linear regression to capture overall directional movements, anomaly detection using z-score analysis to identify unusual patterns, and temporal stability assessment via variability thresholds to characterize consistency.

For each identified property, we generate both descriptive statements and binary questions. For instance, given a time-series segment, we generate descriptive text like “The time-series shows an upward trend” or binary questions such as “Does this time-series contain any anomalies?”

Our experiments show that projectors pretrained on this synthetic corpus can effectively generalize to new tasks, demonstrating the feasibility of creating synthetic data for cross-modal pretraining when natural data pairs are unavailable.

7 DISCUSSION

Limitations and Future Directions In this study, we explored a variety of settings: utilizing different encoders, LLM architectures, modalities, and datasets. Our results demonstrate that LLMs are capable of performing Vector-ICL on both language and non-language inputs. However, our experiments did not cover all possible combinations of these variables. There are still many unexplored areas, such as additional modalities, tasks, and encoder-decoder configurations, that could further benefit from Vector-ICL. Also, we only experimented with single-token encoders, while there exist encoders that produce variable-sized embeddings, that can potentially be more powerful and flexible. [Analyzing how instruction tuning might affect the model’s ability to understand vector context would be beneficial as well.](#) We leave this extensive exploration for future research to fully understand the broader applicability and limitations of our approach across diverse domains.

Conclusion In this work, we explore whether large language models trained only on text can perform in-context learning on continuous vectors from different domains. Our findings suggest that LLMs can indeed understand and process continuous context via embedding projection. Simple linear projections are often sufficient, though for cross-modal tasks—such as those involving non-textual data like time-series or graphs—non-linear transformations may be required. In our experiments across various tasks and modalities, including text reconstruction, numerical function regression, text classification, summarization, molecule captioning, time-series classification, graph classification, and fMRI decoding, Vector-ICL often surpasses both few-shot ICL and task-specific model or tuning. We further conduct analyses and case studies, indicating the potential of LLMs to process vector representations beyond traditional token-based paradigms.

REFERENCES

- 540
541
542 Rishabh Agarwal, Avi Singh, Lei M Zhang, Bernd Bohnet, Stephanie Chan, Ankesh Anand, Zaheer
543 Abbas, Azade Nova, John D Co-Reyes, Eric Chu, et al. Many-shot in-context learning. *arXiv*
544 *preprint arXiv:2404.11018*, 2024.
- 545 Kwangjun Ahn, Xiang Cheng, Hadi Daneshmand, and Suvrit Sra. Transformers learn to imple-
546 ment preconditioned gradient descent for in-context learning. *Advances in Neural Information*
547 *Processing Systems*, 36, 2024.
- 548
549 Ekin Akyürek, Dale Schuurmans, Jacob Andreas, Tengyu Ma, and Denny Zhou. What learning algo-
550 rithm is in-context learning? investigations with linear models. *arXiv preprint arXiv:2211.15661*,
551 2022.
- 552 Abdul Fatir Ansari, Lorenzo Stella, Caner Turkmen, Xiyuan Zhang, Pedro Mercado, Huibin Shen,
553 Oleksandr Shchur, Syama Sundar Rangapuram, Sebastian Pineda Arango, Shubham Kapoor, et al.
554 Chronos: Learning the language of time series. *arXiv preprint arXiv:2403.07815*, 2024.
- 555
556 Yu Bai, Fan Chen, Huan Wang, Caiming Xiong, and Song Mei. Transformers as statisticians:
557 Provable in-context learning with in-context algorithm selection. *Advances in neural information*
558 *processing systems*, 36, 2024.
- 559 Vinamra Benara, Chandan Singh, John X Morris, Richard Antonello, Ion Stoica, Alexander G Huth,
560 and Jianfeng Gao. Crafting interpretable embeddings by asking llms questions. *arXiv preprint*
561 *arXiv:2405.16714*, 2024.
- 562
563 Amanda Bertsch, Maor Ivgi, Uri Alon, Jonathan Berant, Matthew R Gormley, and Graham Neu-
564 big. In-context learning with long-context models: An in-depth exploration. *arXiv preprint*
565 *arXiv:2405.00200*, 2024.
- 566 Satwik Bhattamishra, Arkil Patel, Phil Blunsom, and Varun Kanade. Understanding in-context
567 learning in transformers and llms by learning to learn discrete functions. *arXiv preprint*
568 *arXiv:2310.03016*, 2023.
- 569
570 Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal,
571 Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are
572 few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020.
- 573 Yanda Chen, Chen Zhao, Zhou Yu, Kathleen McKeown, and He He. Parallel structures in pre-
574 training data yield in-context learning. *arXiv preprint arXiv:2402.12530*, 2024.
- 575
576 Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Yunxuan Li,
577 Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, et al. Scaling instruction-finetuned lan-
578 guage models. *Journal of Machine Learning Research*, 25(70):1–53, 2024.
- 579 Damai Dai, Yutao Sun, Li Dong, Yaru Hao, Shuming Ma, Zhifang Sui, and Furu Wei. Why can
580 gpt learn in-context? language models secretly perform gradient descent as meta-optimizers. In
581 *Findings of the Association for Computational Linguistics: ACL 2023*, pp. 4005–4019, 2023.
- 582
583 Hoang Anh Dau, Anthony Bagnall, Kaveh Kamgar, Chin-Chia Michael Yeh, Yan Zhu, Shaghayegh
584 Gharghabi, Chotirat Ann Ratanamahatana, and Eamonn Keogh. The ucr time series archive.
585 *IEEE/CAA Journal of Automatica Sinica*, 6(6):1293–1305, 2019.
- 586
587 Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha
588 Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. The llama 3 herd of models.
arXiv preprint arXiv:2407.21783, 2024.
- 589
590 Carl Edwards, Qingyun Wang, Lawrence Zhou, and Heng Ji. L+m-24: Building a dataset for
591 language+molecules @ acl 2024. *arXiv preprint arXiv:2403.00791*, 2024.
- 592
593 Deqing Fu, Tian-Qi Chen, Robin Jia, and Vatsal Sharan. Transformers learn higher-order op-
timization methods for in-context learning: A study with linear models. *arXiv preprint*
arXiv:2310.17086, 2023.

- 594 Angeliki Giannou, Liu Yang, Tianhao Wang, Dimitris Papailiopoulos, and Jason D Lee. How well
595 can transformers emulate in-context newton’s method? *arXiv preprint arXiv:2403.03183*, 2024.
596
- 597 Tahmid Hasan, Abhik Bhattacharjee, Md. Saiful Islam, Kazi Mubasshir, Yuan-Fang Li, Yong-Bin
598 Kang, M. Sohel Rahman, and Rifat Shahriyar. XL-sum: Large-scale multilingual abstractive
599 summarization for 44 languages. In *Findings of the Association for Computational Linguistics:
600 ACL-IJCNLP 2021*, pp. 4693–4703, Online, August 2021. Association for Computational Lin-
601 guistics. URL <https://aclanthology.org/2021.findings-acl.413>.
- 602 Weihua Hu, Matthias Fey, Marinka Zitnik, Yuxiao Dong, Hongyu Ren, Bowen Liu, Michele
603 Catasta, and Jure Leskovec. Open graph benchmark: Datasets for machine learning on
604 graphs. In Hugo Larochelle, Marc Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan,
605 and Hsuan-Tien Lin (eds.), *Advances in Neural Information Processing Systems 33: Annual
606 Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12,
607 2020, virtual*, 2020. URL [https://proceedings.neurips.cc/paper/2020/hash/
608 fb60d411a5c5b72b2e7d3527cfc84fd0-Abstract.html](https://proceedings.neurips.cc/paper/2020/hash/fb60d411a5c5b72b2e7d3527cfc84fd0-Abstract.html).
- 609 Weihua Hu, Matthias Fey, Hongyu Ren, Maho Nakata, Yuxiao Dong, and Jure Leskovec. Ogb-isc:
610 A large-scale challenge for machine learning on graphs. *arXiv preprint arXiv:2103.09430*, 2021.
611
- 612 Albert Q Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot,
613 Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, et al.
614 Mistral 7b. *arXiv preprint arXiv:2310.06825*, 2023.
- 615 Yixing Jiang, Jeremy Irvin, Ji Hun Wang, Muhammad Ahmed Chaudhry, Jonathan H Chen, and
616 Andrew Y Ng. Many-shot in-context learning in multimodal foundation models. *arXiv preprint
617 arXiv:2405.09798*, 2024.
- 618 Andrew K Lampinen, Ishita Dasgupta, Stephanie CY Chan, Kory Matthewson, Michael Henry
619 Tessler, Antonia Creswell, James L McClelland, Jane X Wang, and Felix Hill. Can language
620 models learn from explanations in context? *arXiv preprint arXiv:2204.02329*, 2022.
621
- 622 Amanda LeBel, Lauren Wagner, Shailee Jain, Aneesh Adhikari-Desai, Bhavin Gupta, Allyson Mor-
623 genthal, Jerry Tang, Lixiang Xu, and Alexander G Huth. A natural language fmri dataset for
624 voxelwise encoding models. *bioRxiv*, pp. 2022–09, 2022.
- 625 Chankyu Lee, Rajarshi Roy, Mengyao Xu, Jonathan Raiman, Mohammad Shoeybi, Bryan Catan-
626 zaro, and Wei Ping. Nv-embed: Improved techniques for training llms as generalist embedding
627 models, 2024.
628
- 629 Mukai Li, Shansan Gong, Jiangtao Feng, Yiheng Xu, Jun Zhang, Zhiyong Wu, and Lingpeng Kong.
630 In-context learning with many demonstration examples. *arXiv preprint arXiv:2302.04931*, 2023a.
- 631 Xiang Lisa Li and Percy Liang. Prefix-tuning: Optimizing continuous prompts for generation. *arXiv
632 preprint arXiv:2101.00190*, 2021.
633
- 634 Yingcong Li, Muhammed Emrullah Ildiz, Dimitris Papailiopoulos, and Samet Oymak. Transformers
635 as algorithms: Generalization and stability in in-context learning. In *International Conference on
636 Machine Learning*, pp. 19565–19594. PMLR, 2023b.
- 637 Chin-Yew Lin. ROUGE: A package for automatic evaluation of summaries. In *Text Summarization
638 Branches Out*, pp. 74–81, Barcelona, Spain, July 2004. Association for Computational Linguis-
639 tics. URL <https://aclanthology.org/W04-1013>.
- 640 Andrew L. Maas, Raymond E. Daly, Peter T. Pham, Dan Huang, Andrew Y. Ng, and Christopher
641 Potts. Learning word vectors for sentiment analysis. In *Proceedings of the 49th Annual Meeting
642 of the Association for Computational Linguistics: Human Language Technologies*, pp. 142–150,
643 Portland, Oregon, USA, June 2011. Association for Computational Linguistics. URL [http:
644 //www.aclweb.org/anthology/P11-1015](http://www.aclweb.org/anthology/P11-1015).
645
- 646 Arvind Mahankali, Tatsunori B Hashimoto, and Tengyu Ma. One step of gradient descent is
647 provably the optimal in-context learner with one layer of linear self-attention. *arXiv preprint
arXiv:2307.03576*, 2023.

- 648 P. Malo, A. Sinha, P. Korhonen, J. Wallenius, and P. Takala. Good debt or bad debt: Detecting
649 semantic orientations in economic texts. *Journal of the Association for Information Science and*
650 *Technology*, 65, 2014.
- 651 Rui Meng, Ye Liu, Shafiq Rayhan Joty, Caiming Xiong, Yingbo Zhou, and Semih Yavuz.
652 Sfembedding-mistral: enhance text retrieval with transfer learning. *Salesforce AI Research Blog*,
653 3, 2024.
- 654 Stephen Merity, Caiming Xiong, James Bradbury, and Richard Socher. Pointer sentinel mixture
655 models, 2016.
- 656
657 Sewon Min, Mike Lewis, Luke Zettlemoyer, and Hannaneh Hajishirzi. Metaicl: Learning to learn in
658 context. In *Proceedings of the 2022 Conference of the North American Chapter of the Association*
659 *for Computational Linguistics: Human Language Technologies*, pp. 2791–2809, 2022a.
- 660 Sewon Min, Xinxu Lyu, Ari Holtzman, Mikel Artetxe, Mike Lewis, Hannaneh Hajishirzi, and Luke
661 Zettlemoyer. Rethinking the role of demonstrations: What makes in-context learning work? In
662 *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pp.
663 11048–11064, 2022b.
- 664 John X Morris, Chandan Singh, Alexander M Rush, Jianfeng Gao, and Yuntian Deng. Tree prompt-
665 ing: efficient task adaptation without fine-tuning. *arXiv preprint arXiv:2310.14034*, 2023.
- 666 Samuel Müller, Noah Hollmann, Sebastian Pineda Arango, Josif Grabocka, and Frank Hutter. Trans-
667 formers can do bayesian inference. In *International Conference on Learning Representations*,
668 2021.
- 669 Shashi Narayan, Shay B. Cohen, and Mirella Lapata. Don’t give me the details, just the summary!
670 topic-aware convolutional neural networks for extreme summarization. *ArXiv*, abs/1808.08745,
671 2018.
- 672
673 Tung Nguyen and Aditya Grover. Transformer neural processes: Uncertainty-aware meta learning
674 via sequence modeling. In *International Conference on Machine Learning*, pp. 16569–16594.
675 PMLR, 2022.
- 676 Jianmo Ni, Chen Qu, Jing Lu, Zhuyun Dai, Gustavo Hernández Ábrego, Ji Ma, Vincent Y Zhao,
677 Yi Luan, Keith B Hall, Ming-Wei Chang, et al. Large dual encoders are generalizable retrievers.
678 *arXiv preprint arXiv:2112.07899*, 2021.
- 679
680 OpenAI. GPT-4 technical report, 2023.
- 681 Bo Pang and Lillian Lee. Seeing stars: Exploiting class relationships for sentiment categorization
682 with respect to rating scales. In *Proceedings of the ACL*, 2005.
- 683
684 Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. BLEU: a method for automatic
685 evaluation of machine translation. In *Proceedings of the 40th Annual Meeting of the Association*
686 *for Computational Linguistics*, pp. 311–318, Philadelphia, Pennsylvania, USA, July 2002.
687 Association for Computational Linguistics.
- 688
689 Matt Post. A call for clarity in reporting BLEU scores. In *Proceedings of the Third Conference*
690 *on Machine Translation: Research Papers*, pp. 186–191, Belgium, Brussels, October 2018. As-
691 sociation for Computational Linguistics. URL [https://www.aclweb.org/anthology/](https://www.aclweb.org/anthology/W18-6319)
692 [W18-6319](https://www.aclweb.org/anthology/W18-6319).
- 693 Elvis Saravia, Hsien-Chi Toby Liu, Yen-Hao Huang, Junlin Wu, and Yi-Shin Chen. CARER: Con-
694 textualized affect representations for emotion recognition. In *Proceedings of the 2018 Confer-*
695 *ence on Empirical Methods in Natural Language Processing*, pp. 3687–3697, Brussels, Bel-
696 gium, October-November 2018. Association for Computational Linguistics. doi: 10.18653/v1/
697 D18-1404. URL <https://www.aclweb.org/anthology/D18-1404>.
- 698 Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D. Manning, Andrew Ng,
699 and Christopher Potts. Recursive deep models for semantic compositionality over a sentiment
700 treebank. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language*
701 *Processing*, pp. 1631–1642, Seattle, Washington, USA, October 2013. Association for Computa-
tional Linguistics. URL <https://www.aclweb.org/anthology/D13-1170>.

- 702 Jerry Tang, Amanda LeBel, Shailee Jain, and Alexander G Huth. Semantic reconstruction of contin-
703 uous language from non-invasive brain recordings. *Nature Neuroscience*, 26(5):858–866, 2023.
704
- 705 Nandan Thakur, Nils Reimers, Andreas Rücklé, Abhishek Srivastava, and Iryna Gurevych. BEIR: A
706 heterogeneous benchmark for zero-shot evaluation of information retrieval models. In *Thirty-fifth*
707 *Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round*
708 *2)*, 2021. URL <https://openreview.net/forum?id=wCu6T5xFjeJ>.
- 709 Jörg Tiedemann. Parallel data, tools and interfaces in opus. In *Lrec*, volume 2012, pp. 2214–2218.
710 Citeseer, 2012.
711
- 712 Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Niko-
713 lay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. Llama 2: Open founda-
714 tion and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*, 2023.
- 715 Trieu H Trinh, Yuhuai Wu, Quoc V Le, He He, and Thang Luong. Solving olympiad geometry
716 without human demonstrations. *Nature*, 625(7995):476–482, 2024.
717
- 718 Johannes Von Oswald, Eyvind Niklasson, Ettore Randazzo, João Sacramento, Alexander Mordv-
719 intsev, Andrey Zhmoginov, and Max Vladymyrov. Transformers learn in-context by gradient
720 descent. In *International Conference on Machine Learning*, pp. 35151–35174. PMLR, 2023.
- 721 Albert Webson and Ellie Pavlick. Do prompt-based models really understand the meaning of their
722 prompts? In *Proceedings of the 2022 Conference of the North American Chapter of the Associa-*
723 *tion for Computational Linguistics: Human Language Technologies*, pp. 2300–2344, 2022.
724
- 725 Jerry Wei, Jason Wei, Yi Tay, Dustin Tran, Albert Webson, Yifeng Lu, Xinyun Chen, Hanxiao Liu,
726 Da Huang, Denny Zhou, et al. Larger language models do in-context learning differently. *arXiv*
727 *preprint arXiv:2303.03846*, 2023.
- 728 Noam Wies, Yoav Levine, and Amnon Shashua. The learnability of in-context learning. *Advances*
729 *in Neural Information Processing Systems*, 36, 2024.
- 730 Jingfeng Wu, Difan Zou, Zixiang Chen, Vladimir Braverman, Quanquan Gu, and Peter Bartlett.
731 How many pretraining tasks are needed for in-context learning of linear regression? In *The*
732 *Twelfth International Conference on Learning Representations*, 2023.
733
- 734 Junda Wu, Zhehao Zhang, Yu Xia, Xintong Li, Zhaoyang Xia, Aaron Chang, Tong Yu, Sungchul
735 Kim, Ryan A Rossi, Ruiyi Zhang, et al. Visual prompting in multimodal large language models:
736 A survey. *arXiv preprint arXiv:2409.15310*, 2024.
- 737 An Yang, Baosong Yang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Zhou, Chengpeng Li,
738 Chengyuan Li, Dayiheng Liu, Fei Huang, et al. Qwen2 technical report. *arXiv preprint*
739 *arXiv:2407.10671*, 2024.
740
- 741 Chengxuan Ying, Tianle Cai, Shengjie Luo, Shuxin Zheng, Guolin Ke, Di He, Yanming Shen, and
742 Tie-Yan Liu. Do transformers really perform badly for graph representation? *Advances in neural*
743 *information processing systems*, 34:28877–28888, 2021.
- 744 Alex Young, Bei Chen, Chao Li, Chengen Huang, Ge Zhang, Guanwei Zhang, Heng Li, Jiangcheng
745 Zhu, Jianqun Chen, Jing Chang, et al. Yi: Open foundation models by 01. ai. *arXiv preprint*
746 *arXiv:2403.04652*, 2024.
- 747 Dun Zhang. Stella, 2024. URL [https://huggingface.co/dunzhang/stella_en_1.](https://huggingface.co/dunzhang/stella_en_1.5B_v5)
748 [5B_v5](https://huggingface.co/dunzhang/stella_en_1.5B_v5).
- 749
- 750 Ruiqi Zhang, Spencer Frei, and Peter L Bartlett. Trained transformers learn linear models in-context.
751 *arXiv preprint arXiv:2306.09927*, 2023.
- 752 Siyan Zhao, John Dang, and Aditya Grover. Group preference optimization: Few-shot alignment of
753 large language models. *arXiv preprint arXiv:2310.11523*, 2023.
754
- 755 Siyan Zhao, Tung Nguyen, and Aditya Grover. Probing the decision boundaries of in-context learn-
ing in large language models. *arXiv preprint arXiv:2406.11233*, 2024.

A APPENDIX

A.1 DETAILED EXPERIMENT SETUP

Text Reconstruction We use two datasets for the text reconstruction task, Parallel Sentence Talks’s English subset (Tiedemann, 2012) and Quora (Thakur et al., 2021).

Parallel Sentence Talks consist of sentences used in movie conversations, and Quora is built on a wide range of online questions. They represent short pieces of text that convey clear information. We aim to explore whether LLMs can decode the original message from the projected text embeddings.

We use the following prompt template:

Translate the text in brackets: (\square), translation: [Original Text]

We train the projectors on the training set of these two datasets and evaluate their performance on the test tests. We measure the reconstruction performance with the BLEU score (Papineni et al., 2002; Post, 2018).

Arithmetic and Function Regression We created synthetic datasets of numerical data to pretrain our linear number projectors, experimenting with two configurations: one using 3-digit numbers and the other using 10-digit numbers. In Llama3.1-8B’s tokenizer, 3-digit numbers are represented as single tokens, while in Mistral-7B, Qwen2-7B, and Yi-1.5-9B, numbers larger than 10 are split into multiple tokens. Consequently, 10-digit numbers consistently span multiple tokens across all models, which increases the complexity of performing arithmetic operations.

To represent the numbers, we use a concatenated-and-flattened one-hot vector encoding for each digit. For instance, a 3-digit number is represented as a 3×10 matrix (one hot per digit place), which is then flattened into a 30-dimensional vector. Similarly, a 10-digit number is represented as a 10×10 matrix, flattened into a 100-dimensional vector.

The pretraining involves two tasks. The first task is number reconstruction, we use the following prompt template, given the number is 128:

$x = \square_x$, x equals to (digits): 128

The second task is basic addition, we use the following prompt template, given the numbers are $x = 128$ $y = 256$, $a = 1$, $b = 1$:

$x = \square_x$, $y = \square_y$, $a * x + b * y$ equals to (digits): 384

Here, a and b are randomly sampled from $[0, 1]$ with up to two decimal places, and the solution is the integer part of the sum.

For evaluation, we use a function regression task with a non-linear function: $f(x, y) = \sqrt{x} \cdot \sqrt{y}$. The LLM is given inputs x and y , along with the integer part of the output $f(x, y)$. The model is then tasked with predicting the output for new pairs of x and y . The prompt for in-context learning is structured as follows:

$x = \square_x$, $y = \square_y$, function(x, y) equals to (digits): [Solution]

For few-shot ICL, the box tokens will be replaced with the actual numbers. We measure the function regression performance with the mean relative error, where the relative error is computed as the ℓ_1 difference divided by the ground truth value.

Text Classification We use five datasets for the text classification task. For binary classification, we include IMDB (Maas et al., 2011), Rotten Tomatoes (Pang & Lee, 2005), and the Stanford Sentiment Treebank (SST2) (Socher et al., 2013). For multi-class classification, we use the Emotion dataset (Saravia et al., 2018) and the Financial Phrasebank (Malo et al., 2014). The binary classification datasets (IMDB, Rotten Tomatoes, and SST2) involve classifying movie reviews as positive or negative. The Emotion dataset classifies Twitter tweets into six categories: anger, fear, joy, love, sadness, and surprise. The Financial Phrasebank categorizes financial news into positive, negative, or neutral sentiments.

We use the following prompt in ICL:

810 (□)'s sentiment is [Input Class]

811
812 For few-shot ICL, the box tokens will be replaced with the actual text. For soft prompt tuning, we use
813 20 virtual tokens and train for one epoch over the entire training set. We measure the classification
814 performance with accuracy on the test set.

815
816 **Text Summarization** We use two datasets for the summarization task, XSum (Narayan et al.,
817 2018) and the English subset of XLSum (Hasan et al., 2021). They contain newspaper articles and
818 their summaries.

819 We use the following prompt in ICL:

820 (□)'s summarization is: [Summary of the Input]

821
822 For few-shot ICL, the box token will be replaced by the article. We measure the performance using
823 the RougeL score with the ground truth summary on the test sets.

824
825 **Molecule Captioning** We use the Language + Molecules-24 (LPM24) dataset for the molecule
826 captioning task, it was created for the task of molecule-language translation, and contains 161K
827 pairs of molecule strings and their captions in the combined training and test set.

828 A sample molecule-caption pair looks like the following:

829
830 Molecule: Cc1c(Cl)cccc1-n1ccn2c(SCC(=O)c3ccccc3C(F)(F)F)nnc2c1=O

831 Caption: The molecule is a pain treatment that impacts inflammatory disease treatment.

832
833 And we use the following prompt for ICL:

834 (□)'s molecule caption is: [Caption of the Input Molecule]

835 For few-shot ICL, we replace the box token with the actual molecule string. We measure the perfor-
836 mance using the BLEU score between the generated caption and the ground truth caption.

837
838 **Brain fMRI Decoding and Classification** We analyze data from LeBel et al. 2022 and Tang
839 et al. 2023, which consists of fMRI responses for 3 human subjects as they listen to 20+ hours of
840 narrative stories from podcasts. We preprocessed the data following Benara et al. 2024, by con-
841 verting the fMRI responses into a 200-dimensional output using principal components analysis and
842 labeling 10-grams of the story text at 2-second intervals using an ensemble of LLMs.

843
844 The data was separated into train set and test set by holding out the same three podcast stories from
845 the three human subjects. We use the same pretraining methodology as text to pretrain on the brain
846 fMRI data. As the data comes in as segments of text and the recorded fMRI, we randomly sample
847 20% of the segments to be in fMRI form and projected into box tokens, and we impose next token
848 generation loss on the rest 80%.

849 We evaluate the projectors by giving them randomly sampled context from the train set, that could
850 come from different human subjects or from a different story, and ask them to decode segments from
851 the test set. We use the following prompt in ICL in our decoding experiments:

852 What is the English translation of the input?
853 Input: □, Response: the input in English is [Text Corresponding to fMRI]

854
855 The random baseline is constructed by giving LLM the randomly sampled, shuffled text from the
856 training set, and generating text according to it. We measure the performance using the RougeL
857 score between the generated text and the ground truth.

858 We construct the classification questions around the properties of the underlying text, for example,
859 "Does the sentence contain a proper noun?", "'Does the input mention anything related to arguing?'.
860 The ground truth is obtained via GPT4o (OpenAI, 2023) as binary labels. We use the following
861 prompt in ICL in our classification experiments, using one of the example questions:

862 Does the input mention anything related to arguing?
863 Input: □, Response (Yes or No): according to the English text of the input,
the answer is [Yes/No]

The random baseline is 50%, as we have downsampled and balanced the data. And we use accuracy as the performance metric.

Time-series We use the Chronos (Ansari et al., 2024) time-series Transformers as the encoder. Chronos was pretrained on large scale time-series and is designed to generate the next segments of the time-series. It has proven effective on a wide range of time-series forecasting benchmarks. We take the output of the last time step from Chronos-base as the time-series representation.

We use two datasets for the time-series classification task, FordA, and FordB, they are also part of the UCR Time Series Classification Archive (Dau et al., 2019) ranging from 4000 to 5000 time-series for each dataset. We use the following prompt in ICL:

()’s class (positive, negative) is: [Input Class]

We construct a synthetic pretraining corpus through comprehensive analysis of additional time-series datasets from the UCR Time Series Classification Archive (Dau et al., 2019), specifically MoteStrain, TwoLeadECG, Wafer, PhalangesOutlinesCorrect, and Yoga. Our data curation framework includes multiple statistical approaches: trend identification via linear regression, anomaly detection using z-score analysis, and assessment of temporal stability through variability thresholds. To enhance corpus diversity, we incorporate binary (yes/no) questions targeting specific time-series characteristics. The resulting QA pairs create a rich textual representation that captures the underlying temporal and statistical properties of the data.

We use the base encoder with trained classification head as the baseline and we measure the prediction performance with accuracy.

Graphs We use Graphormer (Ying et al., 2021) as the encoder model, specifically the one that was pretrained on large scale quantum chemistry graph datasets (Hu et al., 2021). Since the down-stream task (ogbg-molhiv (Hu et al., 2020)) is a molecule property prediction dataset, and with strong class imbalance (3% positive classes), we finetune the encoder on the training set to provide meaningful baselines and embeddings. We take the output prior to the classification layer of the Graphormer as the graph embedding.

We use the ogbg-molhiv (Hu et al., 2020) dataset for the graph classification task. ogbg-molhiv is a molecule property prediction dataset consisting of a total 41.12K graphs with node features and edge attributes. It has a strong class imbalance of having around 3% positive class and 97% negative class.

Hence weighted sampling is adopted in the finetuning of both the baseline Graphormer and the embedding projector to yield meaningful predictions. We use the following prompt in ICL:

()’s class (positive, negative) is: [Input Class]

We use the finetuned Graphormer as our baseline and use the F1 score as the performance metric due to the significant class imbalance.

A.2 CASE STUDY: WHAT INFORMATION WAS PERCEIVED FROM PROJECTED BRAIN FMRI?

Fig. 5 illustrates the mean accuracy achieved in decoding different categories of brain activity based on fMRI data. The categories, ranging from "Physical Actions and Movements" to "Conflict, Urgency, and Change," represent diverse cognitive and perceptual domains. The grouping and corresponding questions are listed in Table 4.

The input data is noisy, and the projector is only trained with pretraining objectives, i.e., to predict the next piece of text given the current fMRI signal. We are surprised that with this pure unsupervised training, the LLM can still pick up meaningful signals from the projected embeddings. Notably, decoding tasks associated with "Physical Actions and Movements" and "Cognitive and Reflective Aspects" demonstrate higher mean accuracy, suggesting that these categories are more distinguishable based on the fMRI embeddings.

918
 919
 920
 921
 922
 923
 924
 925
 926
 927
 928
 929
 930
 931
 932
 933
 934
 935
 936
 937
 938
 939
 940
 941
 942
 943
 944
 945
 946
 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

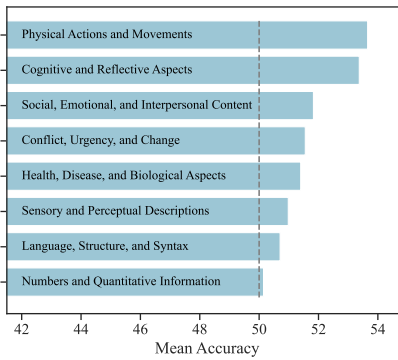


Figure 5: Analyzing LLM’s understanding of projected brain fMRI embeddings after only unsupervised pretraining. We categorize the underlying questions related to the text and measure the mean accuracy for each category, highlighting the LLM’s ability to interpret the embeddings, with only the next token prediction pretraining.

A.3 TEXT CLASSIFICATION AND TEXT SUMMARIZATION CONFIG IN TABLES

Llama3.1-8B is used as the common LLM and NV-Embed-v1 is used as the text encoder for Vector-ICL. The ICL and soft-prompt methods are supplied with up to 50 shots as context. The soft-prompt tokens are trained over the entire training set.

Table 3: Hyperparameters for V-ICL training.

Hyperparameter	Value
Learning Rate	1e-3
Learning Rate Schedule	Cosine Annealing
Optimizer	AdamW
β_1	0.9
β_2	0.999
Training dtype	bf16
Batch Size	128
Generation Temperature	2e-1

A.4 ADDITIONAL EXPERIMENTAL RESULTS




Figure 6: Text Reconstruction

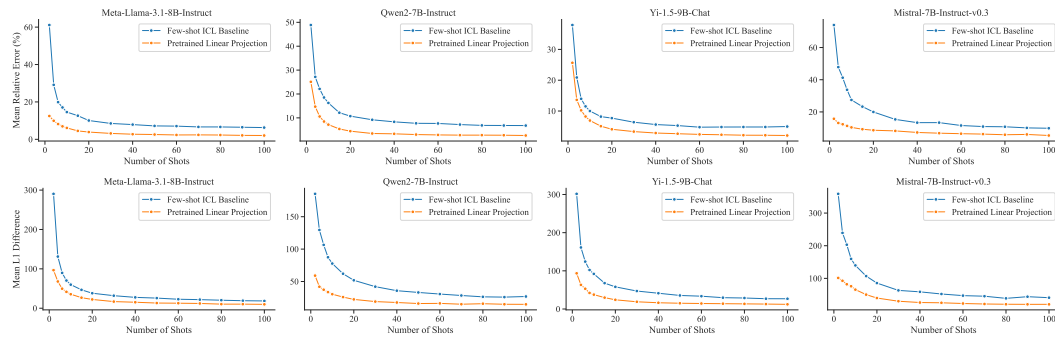


Figure 7: Function Regression - 10 digits (upper) and 3 Digits (lower)

1026
1027
1028
1029
1030
1031
1032
1033
1034
1035
1036
1037
1038
1039
1040
1041
1042
1043
1044
1045
1046
1047
1048
1049
1050
1051
1052
1053
1054
1055
1056
1057
1058
1059
1060
1061
1062
1063
1064
1065
1066
1067
1068
1069
1070
1071
1072
1073
1074
1075
1076
1077
1078
1079

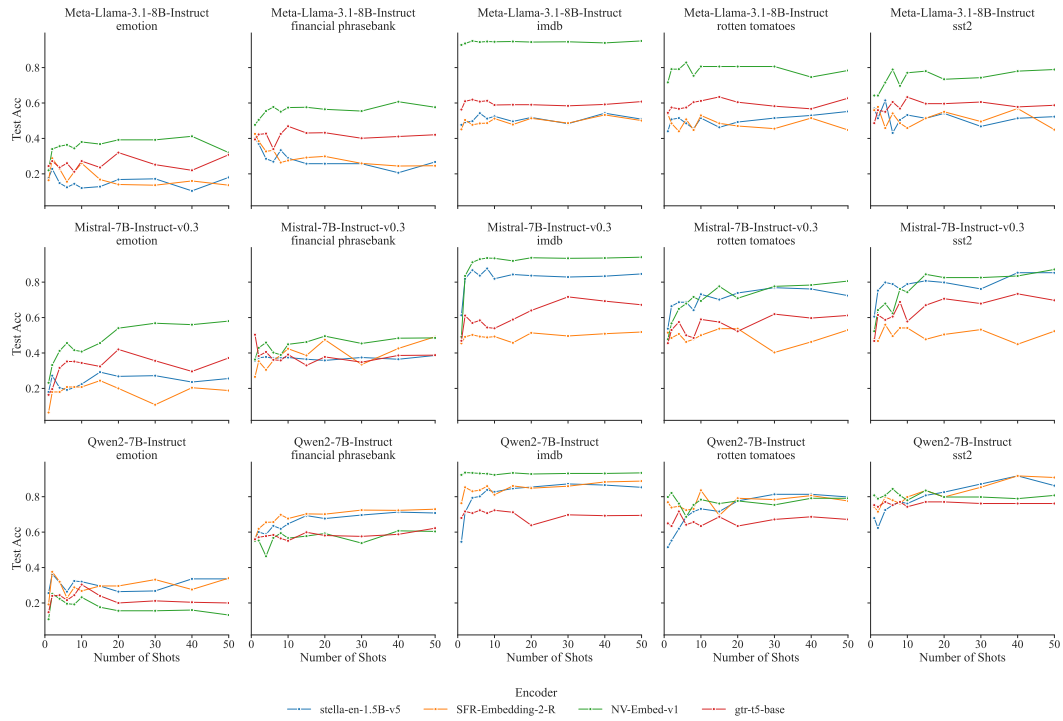


Figure 8: Text Classification - Pretrained Projectors

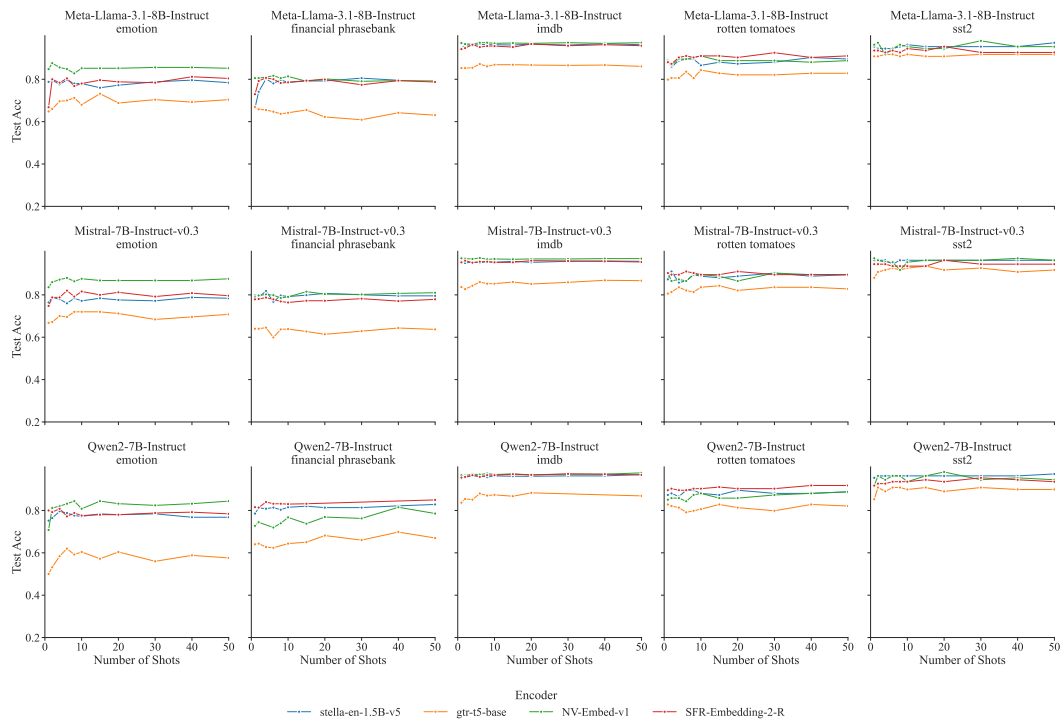


Figure 9: Text Classification - Finetuned Projectors

1080
1081
1082
1083
1084
1085
1086
1087
1088
1089
1090
1091
1092
1093
1094
1095
1096
1097
1098
1099
1100
1101
1102
1103
1104
1105
1106
1107
1108
1109
1110
1111
1112
1113
1114
1115
1116
1117
1118
1119
1120
1121
1122
1123
1124
1125
1126
1127
1128
1129
1130
1131
1132
1133

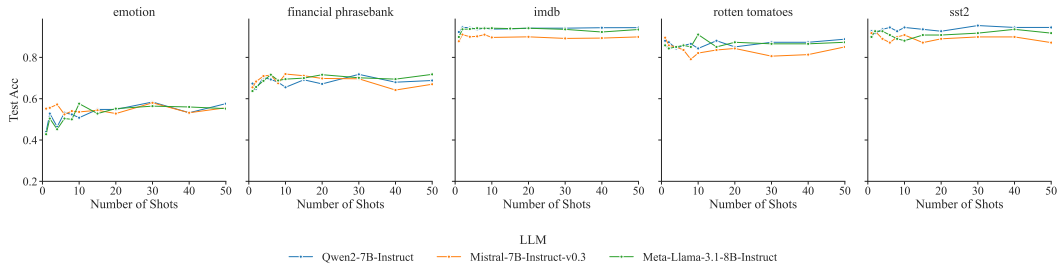


Figure 10: Text Classification - Few-shot ICL




Figure 11: Text Summarization - Pretrained Projectors




Figure 12: Text Summarization - Finetuned Projectors

A.5 BRAIN FMRI QUESTION CATEGORIES

Table 4: Question Categories and Their Associated Questions

Category	Questions
Sensory and Perceptual Descriptions	Does the input mention or describe a taste? Does the input mention or describe a sound? Does the sentence include a specific sound or auditory description? Does the input mention or describe a visual experience? Does the input mention or describe a texture? Does the sentence describe a sensory experience? Does the input mention anything related to color? Does the input mention or describe a smell? Does the input mention anything related to eyes? Does the sentence describe a visual experience or scene? Does the input describe a specific texture or sensation? Does the sentence describe a specific sensation or feeling?
Social, Emotional, and Interpersonal Content	Does the input mention anything related to arguing? Does the input mention anything related to empathy? Does the sentence involve a discussion about personal or social values? Does the input discuss a societal issue or social justice topic? Does the input mention or describe high emotional intensity? Does the sentence describe a relationship between people? Does the input mention or describe highly positive emotional valence? Does the input mention or describe highly negative emotional valence? Does the input mention anything related to conflict? Does the sentence describe a personal or social interaction that leads to a change or revelation? Does the sentence express a philosophical or existential query or observation? Does the sentence involve an expression of personal values or beliefs? Does the sentence express a sense of belonging or connection to a place or community? Is the sentence emotionally positive?
Cognitive and Reflective Aspects	Is the sentence reflective, involving self-analysis or introspection? Does the input involve planning or organizing? Does the text include a planning or decision-making process? Does the sentence convey a decision or choice made by the narrator? Does the sentence describe a personal reflection or thought? Is the input about a discovery or realization? Does the input contain a sense of ambiguity? Is the sentence providing an explanation or rationale? Does the input mention anything related to knowledge?
Language, Structure, and Syntax	Does the sentence contain a proper noun? Does the sentence include a conditional clause? Does the sentence contain a negation? Does the sentence use a unique or unusual word? Does the sentence include a direct speech quotation? Does the sentence include dialogue? Does the sentence contain a cultural reference? Does the sentence involve a recount of a social or community event? Does the input include a comparison or metaphor? Does the sentence include technical or specialized terminology? Is the sentence abstract rather than concrete? Does the sentence include an account of a miscommunication or misunderstanding? Does the text describe a mode of communication?
Physical Actions and Movements	Is the sentence conveying the narrator’s physical movement or action in detail? Does the input mention anything related to motor movements? Does the sentence describe a physical action? Does the sentence describe a physical sensation? Does the sentence describe an activity related to daily life or routine? Does the input mention anything related to an action? Does the sentence involve spatial reasoning?
Numbers and Quantitative Information	Does the input mention a number less than 5? Does the input contain a number? Does the input mention a number greater than 100? Does the input mention anything related to arithmetic? Does the input mention anything related to calculation? Does the input contain a measurement?
Health, Disease, and Biological Aspects	Does the input mention anything related to diseases? Does the input mention anything related to food? Does the input mention anything related to age? Does the input mention anything related to gender? Does the input mention anything related to disgust? Does the input mention anything related to children?
Conflict, Urgency, and Change	Does the sentence involve an unexpected incident or accident? Does the input mention anything related to anger? Does the sentence convey a sense of urgency or haste? Does the sentence describe a change in a physical or emotional state? Does the sentence describe a moment of relief or resolution of tension? Does the sentence express the narrator’s opinion or judgment about an event or character?