

Learning Task Representations from In-Context Learning

Baturay Saglam*, Xinyang Hu, Zhuoran Yang,
Dionysis Kalogerias, Amin Karbasi
Yale University

Abstract

Large language models (LLMs) have demonstrated remarkable proficiency in in-context learning (ICL), where models adapt to new tasks through example-based prompts without requiring parameter updates. However, understanding how tasks are internally encoded and generalized remains a challenge. To address some of the empirical and technical gaps in the literature, we introduce an automated formulation for encoding task information in ICL prompts as a function of attention heads within the transformer architecture. This approach computes a single task vector as a weighted sum of attention heads, with the weights optimized causally via gradient descent. Our findings show that existing methods fail to generalize effectively to modalities beyond text. In response, we also design a benchmark to evaluate whether a task vector can preserve task fidelity in functional regression tasks. The proposed method successfully extracts task-specific information from in-context demonstrations and excels in both text and regression tasks, demonstrating its generalizability across modalities.

1 Introduction

Large language models (LLMs) based on the transformer architecture (Vaswani et al., 2017) have seen dramatic improvements in recent years. A notable feature of these models, such as GPT-3 (Brown et al., 2020), is their capability for *in-context learning* (ICL). This process involves the model receiving a prompt that includes demonstrations of a task in the form of input-output pairs. When presented with a new query input, the model can generate the appropriate output by extrapolating from the provided examples. For instance, after being prompted with a few examples, these models are capable of producing the antonyms of given input words. A

concrete example is

$$\underbrace{\text{vanish} \rightarrow \text{appear}, \text{short} \rightarrow \text{tall}}_{\text{examples}}, \underbrace{\text{increase} \rightarrow}_{\text{query}} \underbrace{\text{decrease}}_{\text{completion}}, \quad (1)$$

where the blue text is the **prompt** and the red text is the **completion** provided by the model. Not limited to linguistic tasks, transformers can also in-context learn a general class of functions \mathcal{F} (Garg et al., 2023). Specifically, for any function $f \in \mathcal{F}$, the model is capable of approximating $f(x_{\text{query}})$ for a new query input x_{query} . This class may include linear or nonlinear relationships, which can be represented by various machine learning models such as linear regression, multi-layer ReLU networks, and decision trees.

The ICL capability is particularly intriguing as it allows models to adapt to a wide range of downstream tasks on-the-fly—*i.e.*, without requiring any parameter updates post-training (Brown et al., 2020). This indicates that models can extract abstract task information (the relationship between input-output pairs) from the prompt and use it to generate accurate responses for query inputs. However, due to the complex nature of the transformer architecture, the computational mechanisms that facilitate how models internally extract and encode task information are not well understood.

Task vectors have been identified as key mechanisms for encoding task information during ICL (Todd et al., 2024; Li et al., 2024; Hendel et al., 2023; Liu et al., 2024). These vectors capture how models process prompts and embed the abstract in-context task as a numerical representation. While the effectiveness of these vectors is empirically well-supported, the precise computation of a task vector within the transformer architecture remains debated. Early findings suggest that layer or attention activations significantly influence ICL performance (Hendel et al., 2023; Liu et al., 2024; Todd et al., 2024).

*Corresponding author: baturay.saglam@yale.edu

However, there is no consensus on the optimal way to conceptualize or compute these vectors, resulting in varied methodologies that often excel in linguistics but fail to generalize to other tasks, such as functional regression. Since these methods rely on the internal components of transformers and architectural differences due to output modality (e.g., language head or numeric outputs) do not change their operation, any meaningful task representation should, in principle, be agnostic to output modality. Poor non-language performance therefore suggests that they may be capturing modality-specific patterns rather than a generalizable task representation. Additionally, these methods require adaptation to individual transformer architectures. For example, the layers over which the task vector is computed must be explicitly identified for each model (Li et al., 2024; Todd et al., 2024; Hendel et al., 2023).

These limitations raise questions about the generalizability and practical utility of existing task vector formulations. To pursue a more standardized and automated pipeline for task encodings, we build on the recent finding that attention heads are the primary driver for vectorizing ICL tasks and hypothesize that:

In-context learning tasks can be effectively represented as a weighted sum of all attention heads in a transformer, with the weights learned causally through gradient descent.

1.1 Contributions and Findings

The key contributions and findings of our study are outlined as follows.

An Automated Task Formulation Independent of Modality and Architecture We propose a framework that approximates the internal weights transformers implicitly use to represent ICL tasks. By optimizing these weights through causal gradient descent, the framework effectively captures the model’s inherent task representations, supporting our hypothesis. Unlike existing methods that require prior model analysis—such as manually selecting hidden layers to perturb—our approach is fully automated. Given any dataset and autoregressive model, it seamlessly computes the task formulation without model-specific adjustments.

A New Benchmark for Regression Tasks While prior work has explored where in-context tasks are represented, most studies focus solely on language,

with limited attention to other modalities like functional regression. We introduce a new benchmark to assess task encodings in regression. This setup provides an intuitive framework for analyzing task encodings and ICL capabilities in regression by leveraging an out-of-distribution condition.

Performance Gains Our method shows effective task encoding on the proposed regression benchmark and language datasets. Ablation studies indicate that this performance stems from aligning the distribution of the last hidden state with that of an optimally in-context-learned model, offering insights into how task steering is achieved.

For reproducibility, we provide the code in our GitHub repository¹.

2 Related Work

Here, we discuss studies that specifically focus on learning task representations in ICL. A more comprehensive review of related works is available in Appendix A.

Initial studies on developing task representations for transformers are documented by Lampinen and McClelland (2020); Shao et al. (2023); Mu et al. (2023); Panigrahi et al. (2023); Ilharco et al. (2023). These works proposed methods to create compositional task encodings through model perturbations in the parameter space: “soft” prompts, codebooks, and meta-mappings. Notably, the term *task vectors* was first introduced by Ilharco et al. (2023). In contrast, through the application of causal mediation analysis (Pearl, 2001; Vig et al., 2020; Meng et al., 2022; Wang et al., 2023a; Geva et al., 2023), *function vectors* were discovered to exist inherently within the transformer architecture and to exhibit strong causal effects (Todd et al., 2024). This finding aligns with research on RNNs, where it was demonstrated that RNN hidden states can be grouped based on task similarities (Lake and Baroni, 2018; Hill et al., 2019).

Efforts to represent ICL tasks often derive a task vector from the layer activations associated with the dummy token following the input query (Hendel et al., 2023). This approach was later refined by averaging the layer activations of dummy tokens across a few-shot prompt and optimizing them (Li et al., 2024). However, these methods were primarily designed for language tasks, where input-output pairs are explicitly separated by specific tokens

¹<https://github.com/baturaysaglam/ICL-task-repr>

(e.g., “ \rightarrow ”). This limitation restricts their applicability to broader domains, such as regression tasks. Another approach leverages the first principal component of the difference in layer activations to guide the ICL task, resulting in *In-Context Vectors* (ICVs) (Liu et al., 2024). Notably, attention heads have been argued to play a critical role in transferring information between token positions (Vaswani et al., 2017; Elhage et al., 2021b). Similarly, Todd et al. (2024) found that attention heads are primarily responsible for encoding ICL behavior. Their proposed formulation, *Function Vector* (FV), is computed as the sum of activations from a specific subset of attention heads, selected based on an *indirect metric* derived from causal inference literature (Pearl, 2001). In this study, we aim to develop a structured and automated method for extracting function vectors from transformer architectures, expanding on the approach of Todd et al. (2024).

3 Technical Preliminaries

We focus on decoder-only autoregressive transformer architectures.

3.1 Transformer Architecture

The transformer architecture uses self-attention mechanisms to process sequences of data. Initially, input data is tokenized into a sequence, where each token represents data units such as segments of words. In this work, we consider autoregressive transformers denoted by M_θ and parameterized by θ . The model predicts the next element in a sequence based on previous outputs. It consists of L layers, each transforming encoded token vectors of dimension d through linear and nonlinear operations. Our focus is on the computation at the last token position within these layers, where each layer $\ell \leq L$ generates a vector representation $\mathbf{h}_\ell \in \mathbb{R}^d$ from its preceding layer’s output.

Self-attention in the transformer architecture employs multi-head attention at each layer:

$$\text{MultiHead}_\ell(Q_\ell, K_\ell, V_\ell) = \text{Concat}(\text{head}_{\ell,1}, \dots, \text{head}_{\ell,J})W_\ell^O,$$

where $\mathbb{R}^q \ni \text{head}_{\ell,j} := \text{Attn}_{\ell,j} = \text{softmax}(Q_{\ell,j}K_{\ell,j}^\top/\sqrt{d_k})V_{\ell,j}$. Here $W_\ell^O \in \mathbb{R}^{Jq \times d}$ is the output projection matrix, and $Q_{\ell,j}$, $K_{\ell,j}$, and $V_{\ell,j}$ are the query, key, and value matrices for each attention head $j \leq J$ at layer ℓ . The

term $\sqrt{d_k}$ normalizes the softmax operation for stability, where d_k is the dimension of the key matrix. This multi-head approach allows the model to dynamically adjust its focus across different parts of the input based on the context.

Note that each head at layer ℓ operates in a low-dimensional space $q < d$, distinct from the main hidden state residual stream of the transformer. As observed by Elhage et al. (2021a), the original transformer formulation (Vaswani et al., 2017) can be seen as dividing the matrix into a block form $[W_{\ell,1}, \dots, W_{\ell,J}]$, followed by directly projecting each head ℓ, j into the residual space. Consequently, the attention head output $a_{\ell,j}$ can be defined using Vaswani et al.’s (2017) and Elhage et al.’s (2021a) notations as:

$$a_{\ell,j} = \text{head}_{\ell,j}(Q_{\ell,j}, K_{\ell,j}, V_{\ell,j})W_{\ell,j}^O \in \mathbb{R}^d.$$

3.2 In-Context Learning

A prompt p^t , corresponding to task t , comprises a sequence of tokens including T input-output exemplar pairs $\{(x_i, y_i)\}_{i=0}^T$. We may refer to the length of prompts by the number of demonstrations they contain. Each pair demonstrates the execution of the same underlying task t . This set defines a functional mapping between each input x_i and its corresponding output y_i . In addition to these exemplar pairs, each prompt includes a specific query input x_{query} that follows the demonstrations. The LM predicts the target response y_{query} by extrapolating from the demonstrations. In our study, we investigate the in-context learning of two distinct modalities: functional regression and language tasks.

Functional Regression Tasks *Regression tasks* refer to a function class \mathcal{F} that is learned through ICL. We closely follow the formulation proposed in (Garg et al., 2023) for training and testing models on regression data. For each prompt, a random function f from \mathcal{F} is sampled according to a distribution $\mathcal{D}_{\mathcal{F}}$, and a set of random inputs $x_i \in \mathbb{R}^m$ for $i = 1, \dots, T$ is drawn from $\mathcal{D}_{\mathcal{X}}$. These inputs are then evaluated by f to produce the prompt $p^f = \{x_1, f(x_1), \dots, x_T, f(x_T), x_{\text{query}}\}$. We denote the model’s output as $M_\theta(p^f)$.

A function class \mathcal{F} can be either linear or nonlinear. For linear functions, inputs are drawn from an isotropic Gaussian distribution $\mathcal{N}(0, I_m)$, and a random function is defined by sampling a weight vector w from $\mathcal{N}(0, I_m)$, setting $f(x) = w^\top x$. Here, the task is characterized by the weight vector $w \in \mathbb{R}^m$. For nonlinear functions, possible forms




Figure 1: Squared error on the query input as a function of the number of demonstrations in the prompts, where lower error $(f(x_{\text{query}}) - M_{\theta}(x_{\text{query}} | p^f))^2$ indicates better ICL performance. We evaluate three different function classes: (a) linear functions $f(x_{\text{query}}) = w^{\top} x_{\text{query}}$, (b) sparse linear functions $f(x_{\text{query}}) = w_s^{\top} x$, and (c) 2-layer ReLU neural networks (NNs) $f(x_{\text{query}}) = W_2 \text{ReLU}(W_1 x_{\text{query}})$. Results are averaged over a batch of 256 tasks randomly selected from the same function class. The shaded area represents the 95% confidence interval over the sampled prompts. The dashed line indicates the maximum sequence length used during the positional encoding training, T_{train} .

of f include multi-layer ReLU networks or decision trees. We employ the models pre-trained by Garg et al. (2023), with the training procedure and additional details provided in Appendix C.2.2.

Language Tasks We focus on straightforward NLP applications such as antonym generation, with an example shown in (1). The model M_{θ} processes each prompt p^t and produces a next-token distribution $M_{\theta}(\cdot | p^t)$ over the vocabulary \mathcal{V} . In our evaluation of task representations, we follow Todd et al. (2024), who *corrupted* prompts \tilde{p}_i^t by shuffling labels within each prompt. This shuffling disrupts the connection between inputs $x_{i,k}$ and outputs $\tilde{y}_{i,k}$, rendering the examples uninformative about the task and effectively “blocking” ICL.

4 Learning Task Representations

4.1 Motivational Observations

Our study begins by noting that existing ICL task formulations are either inapplicable (Li et al., 2024; Hendel et al., 2023) (see Appendix C.2.6) or fail to generalize to regression tasks (Liu et al., 2024; Todd et al., 2024). To demonstrate this limitation, we design a new benchmark for functional regression.

The motivation for this benchmark is straightforward and follows standard practices from prior work on language tasks (see Section 3.2). We first

prevent the model from extracting task understanding from in-context examples. If ICL performance is restored when the task vector is integrated into the model, it suggests that the task vector encodes the task information. However, techniques like shuffling labels or probing are unsuitable for regression tasks due to fundamental differences in how input-output relationships are represented and learned (see Appendices B.1 and B.2).

To suppress task-relevant information in regression prompts, we introduce an out-of-distribution (OOD) case. The model is initialized and pre-trained to handle sequences up to a maximum length T_{max} , ensuring it can process longer sequences. However, during pre-training, the model is exposed only to sequences of length up to $T_{\text{train}} < T_{\text{max}}$, leaving positional encodings for positions beyond T_{train} untrained.

At test time, we evaluate the model on sequences longer than T_{train} but within T_{max} . The focus is not on downstream accuracy but on robustness under OOD conditions. Since positional encodings beyond T_{train} are untrained, the model cannot infer where the input and outputs are located within the prompt, and thus cannot extract the desired input-output relationship any further—the importance of positional information in ICL has also been considered by Wibisono and Wang (2024). However, if a representation contains task-relevant information (specific to the ideal function f), it should maintain functional task behavior and generalize to prompts longer than T_{train} .

The results, shown in Figure 1, reveal that the tested approaches fail to generalize across the evaluated class of functions, despite decent performance in the language domain. Because these methods rely on the internal components of transformers and are unaffected by architectural differences related to output modality (e.g., language heads or numeric outputs), any meaningful task representation should, in principle, remain consistent across output modalities.

As our investigation deepens, we articulate the following central research questions:

- (i) How can we develop a task representation in a principled manner that generalizes across modalities such as language and mathematical functions?
- (ii) How can this representation be leveraged to preserve task fidelity on downstream tasks?

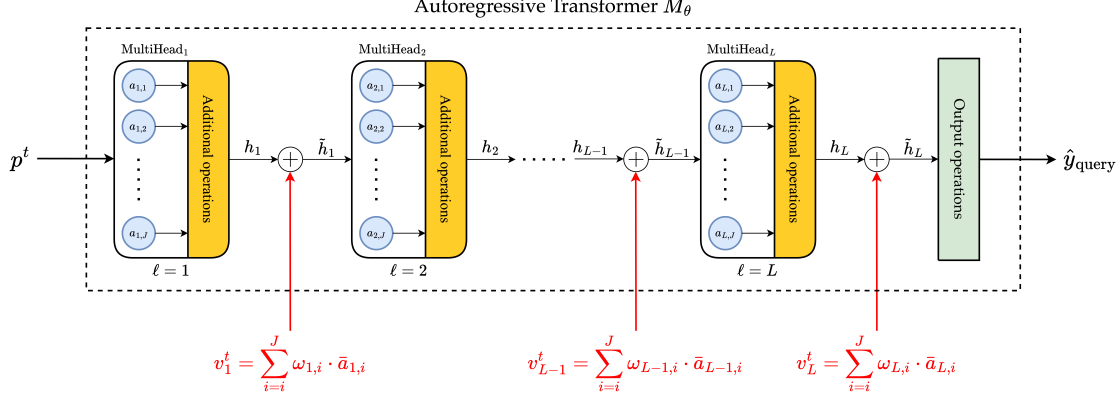


Figure 2: Illustration of the operation. *Additional* and *output* operations may include residual connections, normalization, feedforward, or prediction layers, depending on the architecture. LTV is added sequentially to each layer, allowing the effects of the integrated LTV to be progressively observed across subsequent layers.

4.2 Learnable Task Vector

Our approach addresses key limitations of previous methods, such as Function Vectors, by automatically computing task formulations without the need for prior model analysis or manual adjustments.

Variability in Contributions of Attention Heads

The FV encoding is based on strong causal evidence that attention heads are the main carriers of in-context task information (Todd et al., 2024). This formulation assumes equal contributions from all attention heads to the task representation by summing their activations. However, we argue that the influence of each head varies—some contribute significantly to task encoding, while others have a lesser impact. Constraining weights to unity is therefore impractical, and focusing exclusively on a subset of attention heads risks overlooking subtle contributions from others. To address this, attention heads should be weighted in the summation to accurately capture their varying contributions.

Layer-Specific Task Vectors Previous work typically integrates task encodings into models by adding vectors to the outputs of specific hidden layers through simple vector summation (Todd et al., 2024; Hendel et al., 2023; Li et al., 2024). However, this approach assumes uniform applicability across layers, ignoring that transformers’ hidden states represent progressively refined representations of the input. To overcome this limitation, we compute task vectors uniquely for each hidden layer, aligning them with the evolving representations. While initializing a unique set of attention head weights for each of the L layers is an option, it would be computationally expensive. Instead, we compute

layer-specific task vectors as a weighted sum of the attention heads within each layer:

$$v_\ell^t = \sum_{j=0}^J \omega_{\ell,j} \cdot \bar{a}_{\ell,j}, \quad (2)$$

where $\omega_{\ell,j} \in \mathbb{R}$ represents the weight parameters assigned to attention heads, organized in the parameter vector $\Phi \in \mathbb{R}^{LJ}$, and \bar{a} is the attention activations averaged on a separate sample set of prompts demonstrating task t . This method ensures that each layer-wise FV is composed solely of the J heads within that specific layer, avoiding the aggregation of attention across all $L \times J$ heads.

Although excluding attention heads from layers $\ell' \neq \ell$ in the latter modification might seem contradictory, this is counterbalanced by the transformer’s feedforward design. Including attention from earlier layers could introduce redundancy and complicate learning, as the hidden state at layer ℓ already encapsulates all transformed information in layers $\ell' < \ell$. Furthermore, integrating attention heads from future layers would conflict with the transformer’s sequential processing that avoids forward-looking capabilities.

We refer to the resulting formulation as *Learnable Task Vector* (LTV) and illustrate its operation during inference in Figure 2. Next, we describe the methodology for learning the weights Φ of LTV.

4.3 How to Learn Task Representations

We optimize LTV weights through gradient steps to approximate the “true” weights the model uses to represent ICL tasks, aiming to validate this formulation as a *proof-of-concept*. Consequently, this

Algorithm 1 Optimizing Learnable Task Vector (LTV)

Input: Model M_θ , task t , and number of samples N
Initialize and freeze: Model parameters θ (no updates)
1: Collect N prompts $\{p_i^t\}_{i=1}^N$ demonstrating task t
2: **repeat**
3: **for** each p_i^t in $\{p_i^t\}_{i=1}^N$ **do**
4: Compute LTV: $\mathbf{v}_\Phi^t = \{v_\ell^t\}_{\ell=1}^L$, where $v_\ell^t = \sum_{j=0}^J \omega_{\ell,j} \cdot \bar{a}_{\ell,j}$
5: Obtain the LTV-integrated model output: $\hat{y}_{\text{query},i} \leftarrow M_\theta(p_i^t \mid \mathbf{v}_\Phi^t)$
6: Compute the loss: $\mathcal{L}(y_{\text{query},i}, \hat{y}_{\text{query},i})$ based on (3) or (4)
7: Compute gradients and update $\omega_{\ell,j}$ (model parameters θ remain frozen)
8: **end for**
9: **until convergence**

approach is *not* intended to directly address issues like length generalization shown in Figure 1. Instead, LTV serves as a mechanistic probe into the model’s internal structure, demonstrating that transformer architectures admit d -dimensional vectorization of task information that can be manipulated to steer behavior.

An effective strategy for learning the parameters Φ involves optimizing the LTV in scenarios where the model (with frozen parameters) fails to perform ICL, isolating the causal contribution of Φ to the task performance. This approach ensures our method is both *data-driven* and *causal*, as it directly attributes improvements to $|\Phi| = L \times J$ parameters, facilitating application to downstream tasks. The overall pipeline for training an LTV is summarized in Algorithm 1.

Regression Tasks The LTV is computed over prompts longer than the maximum length encountered during pre-training ($T_v > T_{\text{train}}$). Its integration into the model modifies the transformer’s output. Subsequently, backpropagation is performed over Φ through the transformer to minimize the loss on the query input x_{query} :

$$\min_{\Phi} \mathbb{E}_{f \sim \mathcal{D}_f, x \sim \mathcal{D}_x} \left[\left(\tilde{M}_\theta(p^f \mid \mathbf{v}_\Phi^f) - f(x_{\text{query}}) \right)^2 \right], \quad (3)$$

where $\mathbf{v}_\Phi^f = \{v_\ell^f\}_{\ell=1}^L$ represents the set of layer-wise LTVs computed for function f , $\tilde{M}_\theta(\cdot \mid \mathbf{v}_\Phi^f)$ denotes the transformer output modified by integrating \mathbf{v}_Φ^f into the corresponding hidden layers, and the regression prompt p^f is in the structure:

$$p^f = \{x_1, f(x_1), \dots, x_{T_v}, f(x_{T_v}), x_{\text{query}}\}.$$

Language Tasks Given the true query output y_{query} , the LTV is trained in a supervised manner to minimize the cross-entropy loss on shuffled

prompts:

$$\min_{\Phi} \mathbb{E}_{\tilde{p}^t \sim \tilde{P}^t} \left[-\log \left(\tilde{M}_\theta(y_{\text{query}} \mid \tilde{p}^t; \mathbf{v}_\Phi^t) \right) \right], \quad (4)$$

where $\tilde{M}_\theta(y_{\text{query}} \mid \tilde{p}^t; \mathbf{v}_\Phi^t)$ denotes the probability predicted by the model for the true class y_{query} given the shuffled prompt \tilde{p}^t .

5 Experiments

Models We employ decoder-only autoregressive transformers: GPT-2 (Radford et al., 2019) for regression tasks and GPT-J (Wang and Komatsuzaki, 2021) for tasks in the language domain, using the default configurations described by Garg et al. (2023) and Todd et al. (2024), respectively. GPT-2 is configured with 9.5M parameters (256 hidden dimension) across 12 layers and 8 attention heads per layer, while GPT-J features 6B parameters, 28 layers, and 16 attention heads per layer. The GPT-2 models pre-trained on functional regression data follow the training procedure outlined by Garg et al. (2023).

Benchmarking For the regression data, we evaluate performance on sequences longer than the maximum length seen during pre-training but within the model’s capacity: $T_{\text{train}} < T_v \leq T_{\text{max}}$, as described in Section 4.1. This setup provides a controlled environment to assess the functional restoration of task behavior through causal interventions. In the language domain, we adopt the approach of shuffling labels of in-context examples to diminish ICL (Todd et al., 2024), ensuring fair comparisons with Todd et al. (2024); Hendel et al. (2023); Li et al. (2024).

Tasks For regression data, as illustrated in Figure 1, we evaluate the models on classes of linear functions, sparse linear functions, and 2-layer ReLU

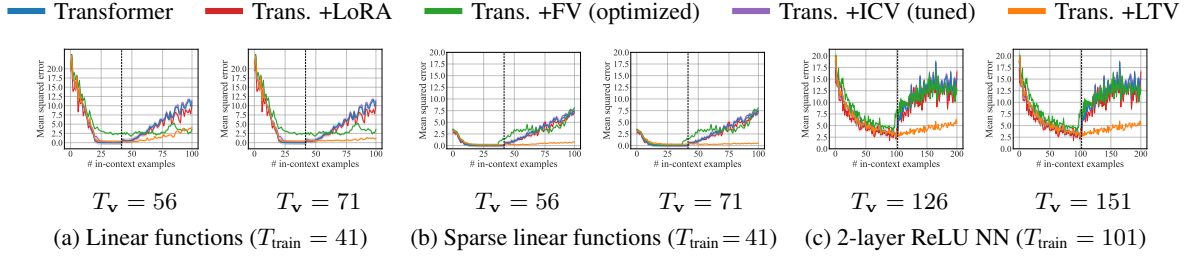


Figure 3: Squared error on the query input, averaged over a batch of 256 tasks randomly selected from the same function class. The shaded area represents the 95% confidence interval over the sampled prompts. The dashed line indicates the number of examples the transformer was trained with, and T_v denotes the prompt length used in LTV training. Complete results for different T_v values are provided in Figures 5, 6, and 7 in Appendix D.1.

networks (Garg et al., 2023). While decision trees have also been explored as a function class, they exhibit less sensitivity to positional encodings, which we discuss in Appendix C.2.3.

For NLP tasks, we focus on single-token generation, as considered in previous work (Hendel et al., 2023; Todd et al., 2024). The tasks are derived from widely studied natural language corpora, including SST-2 (Socher et al., 2013a) and CoNLL-2003 (Tjong Kim Sang and De Meulder, 2003). Further details are provided in Appendix C.3.1.

Baseline Methods For functional regression, we include FV (Todd et al., 2024) and ICV (Liu et al., 2024) as baselines for comparison. The other methods discussed in Appendix A were primarily designed for language tasks and are not applicable to regression tasks (see Appendix C.2.6 for details). Although FV and ICV were originally proposed for linguistic applications, we adapted them for regression tasks with several optimizations. Additionally, we consider Low-Rank Optimization (LoRA) (Hu et al., 2022) for further discussion. For language tasks, we also include Task Vector (TV) (Hendel et al., 2023) and State Vector (SV) with inner optimization (Li et al., 2024) alongside FV and ICV.

Learnable Task Vector For each task, we curated a dataset of 10,000 samples to demonstrate the underlying task. LTVs were trained for approximately 2,000 iterations on regression tasks and 120,000 iterations on language tasks, using a lower learning rate for the latter. Each iteration involved sampling a batch of 256 input prompts and performing gradient descent based on the objectives in (3) and (4) for regression and language tasks, respectively.

Details of our experimental setup are provided in Appendix C. The full set of results is avail-

able in Appendix D, including: (i) LTV trained with various T_v , optimized and tested on chain-of-thought (CoT) prompts, and evaluations under distributional shift; (ii) accuracy scores on 15 NLP benchmarks; and (iii) complete ablation results.

5.1 Evaluation on Regression Tasks

The loss curves for ICL inference are shown in Figure 3. While FV offers some benefits in linear regression, ICV negatively impacts the vanilla model’s performance. Despite extensive optimizations, these results indicate that task encoding formulations, originally crafted for language, do not generalize well to other modalities.

LTV preserves task information in the OOD case. LTV shows minimal performance difference when trained with prompt lengths close to T_{train} . However, as T_v increases, LTV’s impact becomes more prominent, achieving optimal performance comparable to the vanilla model at $T = T_{\text{train}}$ (see Appendix D.1 for detailed plots). While this might raise concerns about the validity of the results, as LTV is causally optimized in the OOD cases ($T > T_{\text{train}}$), training LTV with prompt lengths slightly exceeding T_{train} —specifically, $T_v = 1.37 \times T_{\text{train}}$ for linear functions and $T_v = 1.25 \times T_{\text{train}}$ for 2-layer ReLU networks—proves sufficient to maintain performance. Extending T_v further offers no significant improvements, indicating that only a small portion of the OOD case is needed for LTV to generalize. This suggests that the attention head weights naturally converge to their optimal values with sufficient training data.

Benefits of LTV come from effective task encoding, not fine-tuning. The effects of LTV in regression tasks may resemble those of parameter-

	Dataset	Prompt	Transformer	+LoRA	+FV	+ICV	+SV	+TV	+LTV (ours)
Abstractive	AG News	Zero-shot	0.0 \pm 0.0	9.1 \pm 0.2	30.1 \pm 5.7	33.6 \pm 5.6	31.7 \pm 5.7	34.5 \pm 5.0	46.9 \pm 6.2
		Few-shot	67.2 \pm 5.8	69.3 \pm 3.8	66.0 \pm 5.8	64.7 \pm 5.4	66.1 \pm 4.5	57.9 \pm 5.8	71.1 \pm 5.6
	CommonsenseQA	Zero-shot	10.2 \pm 3.7	24.8 \pm 4.1	23.0 \pm 5.2	28.9 \pm 5.2	21.7 \pm 1.1	22.0 \pm 4.0	41.0 \pm 6.1
		Few-shot	18.0 \pm 4.7	18.6 \pm 5.5	20.3 \pm 5.0	20.0 \pm 4.6	20.3 \pm 2.1	19.5 \pm 3.4	21.1 \pm 5.0
	Sentiment Analysis	Zero-shot	0.0 \pm 0.0	0.7 \pm 0.6	0.0 \pm 0.0	4.8 \pm 1.9	0.0 \pm 0.0	0.0 \pm 0.0	10.9 \pm 3.8
		Few-shot	74.6 \pm 5.4	75.9 \pm 4.1	69.5 \pm 5.7	72.5 \pm 3.2	74.2 \pm 2.5	77.1 \pm 3.4	94.5 \pm 2.8
Extractive	NER-person	Zero-shot	5.5 \pm 2.8	15.7 \pm 1.1	48.8 \pm 6.2	49.9 \pm 5.8	46.3 \pm 1.3	47.8 \pm 5.3	58.2 \pm 6.1
		Few-shot	11.7 \pm 4.0	12.1 \pm 1.6	56.2 \pm 6.1	65.9 \pm 5.2	53.7 \pm 3.3	62.6 \pm 3.9	79.3 \pm 5.0
	NER-location	Zero-shot	6.6 \pm 3.1	6.5 \pm 2.4	37.9 \pm 6.0	26.4 \pm 5.1	34.6 \pm 3.9	27.2 \pm 4.7	27.3 \pm 5.5
		Few-shot	23.4 \pm 5.2	29.5 \pm 5.5	43.4 \pm 6.1	47.4 \pm 5.9	43.5 \pm 4.4	41.9 \pm 5.6	57.4 \pm 6.1
	NER-organization	Zero-shot	21.9 \pm 5.1	21.9 \pm 6.6	48.8 \pm 6.2	52.5 \pm 6.0	46.1 \pm 1.2	46.6 \pm 4.3	64.1 \pm 5.9
		Few-shot	16.8 \pm 4.6	19.2 \pm 3.8	50.8 \pm 6.2	54.1 \pm 5.0	53.5 \pm 2.6	51.1 \pm 4.1	75.0 \pm 5.3

Table 1: Accuracy scores (%) for zero-shot and few-shot (5-shot) predictions, averaged across 256 random seeds. \pm represents the margin of error at a 95% confidence level. The highest accuracy is marked in **boldface**, and the statistically best-performing method is **highlighted**. The complete results are provided in Table 2 in Appendix D.4, along with references to the benchmark datasets.

efficient fine-tuning (PEFT) methods, as only a small number of parameters ($|\Phi| = L \times J$) are trained to adapt the model to previously unlearned scenarios. To evaluate this, we compare LTV with LoRA (rank 8). While LoRA fine-tunes approximately 196K of the 9.5M GPT-2 parameters, LTV optimizes only 96. Despite LoRA fine-tuning a significantly larger number of parameters on sequences with a maximum length T_{\max} , no notable performance improvement is observed. Since LoRA depends on the pre-trained model’s existing positional representations, it cannot address untrained positional encodings. This observation indicates that the task fidelity achieved by LTV is not a result of sole fine-tuning but rather of effective task encoding that can steer the behavior of the model.

Interpretation of the Learned LTV Weights

No regularization, optimization techniques, dropout, or specific weight initialization were applied to learn the LTV weights. The weights are unbounded and sampled from the standard Gaussian distribution. Visualizing the learned weights as heatmaps revealed they consistently lie within the interval $[-3, 3]$ across tasks. However, no meaningful patterns were observed though, suggesting that any underlying structure may be subtle and challenging to detect visually. Notably, some weights converge to near-zero, supporting the idea that only a subset of attention heads encodes task information. Our method thus provides a principled way to identify and weigh this subset, rather than assuming uniform importance or selecting it manually.

5.2 Evaluation on Language Tasks

The accuracy scores are reported in Tables 1 and 2. While Todd et al. (2024) primarily considered “filtered” accuracies, which take into account only the test queries where at least one model responds correctly, we present unfiltered accuracies as a fairer metric, counting all samples regardless of model performance.

LTV is also superior in the language domain.

In 29 out of 30 evaluations (15 tasks with zero- and few-shot prompts), LTV achieves the highest accuracy and ranks among the statistically best-performing methods (see Table 2 in Appendix D.4 for complete results). Furthermore, in 19 out of 30 evaluations, LTV’s performance is statistically superior to all other techniques. We also observe a notable performance by ICV, particularly when the vanilla transformer performs poorly in the zero-shot setting. This effectiveness is attributed to ICV’s approach of crafting unique task vectors for each layer, similar to our method. Lastly, while evaluating LTV with higher shot counts may reveal diminishing gains, we expect the gains to remain nonzero and notable, as the vanilla model may not achieve further contextualization beyond a saturation point. Overall, these findings confirm that LTV excels in both regression and challenging linguistic tasks, supporting our hypothesis about how ICL tasks are formulated within the transformer architecture.

5.3 Ablation Studies

We aim to understand how LTVs sustain task behavior when ICL is not performed. To investigate this,

we use the regression data benchmark and analyze the last hidden states of an LTV-integrated model and a vanilla model at $T = T_{\text{train}}$ (the optimal-performing model). We measure the distributional similarities of the last hidden states using KL divergence to assess how closely the LTV-integrated model aligns with the vanilla model.

The KL divergence scores are reported in Table 3 in Appendix D.5. From the baselines, we only include FV, as it is the only method that does not degrade the vanilla transformer. The corresponding probability densities of the last hidden state distributions, visualized through histograms, are also included in Appendix D.5. Additionally, a detailed description of our experimental methods, along with a figure illustrating the pipeline, is provided in Appendix C.6.

LTV maintains the last hidden state distribution with that of the optimal-performing model. We observe that the divergence between the last hidden state distributions of the LTV-integrated model and the optimal-performing model is generally low and significantly decreases when the LTV training length matches or exceeds the middle length. This suggests that LTV’s effectiveness primarily stems from its ability to dominate the last hidden state and align it with that of the optimal-performing model, thereby facilitating task steering. This observation further implies that task steering may *generally* be achieved through this mechanism, not exclusively by our method. However, a deeper analysis is required to fully understand this process, which is beyond the scope of this study.

6 Conclusion

We investigated the internal representation of in-context learning (ICL) tasks in large language models (LLMs). We propose the *Learnable Task Vector* (LTV) framework, which encodes ICL tasks as a weighted sum of attention head activations optimized through causal gradient descent. LTV generalizes across both text and regression tasks, addressing the limitations of existing representations. Through empirical validation using a novel benchmark designed to analyze task encodings, we demonstrate that LTV preserves task fidelity when models fail to extract task information from demonstrations. Our findings further suggest that this preservation is achieved by maintaining the last hidden state distribution similar to that of an optimally in-context-learned model, bringing us

closer to understanding how task steering works in general.

Limitations

Although not explicitly stated in the main body, our study assumes access to the models’ hidden states, limiting the scope to open- or white-box models. Additionally, we acknowledge that optimizing LTV can be time- and resource-intensive, e.g., requiring 120K iterations over 10K samples for regression. LTV training serves only to learn a fixed task representation vector while keeping model weights frozen, and it is not intended to replace ICL or few-shot prompting. A dataset for LTV training must first be generated or collected. For functional regression, this is relatively straightforward, as data and regression coefficients can be sampled from distributions (e.g., zero-mean Gaussian). However, for language data, it must be collected or generated using a generative AI tool. As a reminder, we utilized data collected by Todd et al. (2024) based on widely studied datasets such as SST-2 (Socher et al., 2013a) and CoNLL-2003 (Tjong Kim Sang and De Meulder, 2003). Finally, because hidden layer access is required, deploying an open-source model on GPUs may demand significant computational resources.

Ethical Considerations

We do not identify any direct application of our study for unethical purposes.

References

- Ekin Akyürek, Dale Schuurmans, Jacob Andreas, Tengyu Ma, and Denny Zhou. 2023. [What learning algorithm is in-context learning? investigations with linear models](#). In *The Eleventh International Conference on Learning Representations*.
- Adrien Bibal, Rémi Cardon, David Alfter, Rodrigo Wilkens, Xiaoou Wang, Thomas François, and Patrick Watrin. 2022. [Is attention explanation? an introduction to the debate](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3889–3900, Dublin, Ireland. Association for Computational Linguistics.
- Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu,

- Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language models are few-shot learners](#). *Preprint*, arXiv:2005.14165.
- Kevin Clark, Urvashi Khandelwal, Omer Levy, and Christopher D. Manning. 2019. [What does BERT look at? an analysis of BERT’s attention](#). In *Proceedings of the 2019 ACL Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 276–286, Florence, Italy. Association for Computational Linguistics.
- Damai Dai, Yutao Sun, Li Dong, Yaru Hao, Shuming Ma, Zhifang Sui, and Furu Wei. 2023. [Why can GPT learn in-context? language models secretly perform gradient descent as meta-optimizers](#). In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 4005–4019, Toronto, Canada. Association for Computational Linguistics.
- N. Elhage, N. Nanda, C. Olsson, T. Henighan, N. Joseph, B. Mann, A. Askell, Y. Bai, A. Chen, T. Conerly, et al. 2021a. A mathematical framework for transformer circuits. <https://transformer-circuits.pub/2021/framework/index.html>. Transformer Circuits Thread.
- Nelson Elhage, Neel Nanda, Catherine Olsson, Tom Henighan, Nicholas Joseph, Ben Mann, Amanda Askell, Yuntao Bai, Anna Chen, Tom Conerly, Nova DasSarma, Dawn Drain, Deep Ganguli, Zac Hatfield-Dodds, Danny Hernandez, Andy Jones, Jackson Kernion, Liane Lovitt, Kamal Ndousse, Dario Amodei, Tom Brown, Jack Clark, Jared Kaplan, Sam McCandlish, and Chris Olah. 2021b. A mathematical framework for transformer circuits. *Transformer Circuits Thread*. <https://transformer-circuits.pub/2021/framework/index.html>.
- Shivam Garg, Dimitris Tsipras, Percy Liang, and Gregory Valiant. 2023. [What can transformers learn in-context? a case study of simple function classes](#). *Preprint*, arXiv:2208.01066.
- Mor Geva, Jasmijn Bastings, Katja Filippova, and Amir Globerson. 2023. [Dissecting recall of factual associations in auto-regressive language models](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 12216–12235, Singapore. Association for Computational Linguistics.
- Danny Halawi, Jean-Stanislas Denain, and Jacob Steinhardt. 2024. [Overthinking the truth: Understanding how language models process false demonstrations](#). In *The Twelfth International Conference on Learning Representations*.
- Roei Hendel, Mor Geva, and Amir Globerson. 2023. [In-context learning creates task vectors](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 9318–9333, Singapore. Association for Computational Linguistics.
- Evan Hernandez, Arnab Sen Sharma, Tal Haklay, Kevin Meng, Martin Wattenberg, Jacob Andreas, Yonatan Belinkov, and David Bau. 2024. [Linearity of relation decoding in transformer language models](#). In *The Twelfth International Conference on Learning Representations*.
- Felix Hill, Adam Santoro, David Barrett, Ari Morcos, and Timothy Lillicrap. 2019. [Learning to make analogies by contrasting abstract relational structure](#). In *International Conference on Learning Representations*.
- Or Honovich, Uri Shoham, Samuel R. Bowman, and Omer Levy. 2023. [Instruction induction: From few examples to natural language task descriptions](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1935–1952, Toronto, Canada. Association for Computational Linguistics.
- Phu Mon Htut, Jason Phang, Shikha Bordia, and Samuel R. Bowman. 2019. [Do attention heads in bert track syntactic dependencies?](#) *Preprint*, arXiv:1911.12246.
- Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2022. [LoRA: Low-rank adaptation of large language models](#). In *International Conference on Learning Representations*.
- Gabriel Ilharco, Marco Tulio Ribeiro, Mitchell Wortsman, Ludwig Schmidt, Hannaneh Hajishirzi, and Ali Farhadi. 2023. [Editing models with task arithmetic](#). In *The Eleventh International Conference on Learning Representations*.
- Sarthak Jain and Byron C. Wallace. 2019. [Attention is not Explanation](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 3543–3556, Minneapolis, Minnesota. Association for Computational Linguistics.
- Diederik P. Kingma and Jimmy Ba. 2015. [Adam: A method for stochastic optimization](#). In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- Goro Kobayashi, Tatsuki Kuribayashi, Sho Yokoi, and Kentaro Inui. 2020. [Attention is not only a weight: Analyzing transformers with vector norms](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7057–7075, Online. Association for Computational Linguistics.
- Olga Kovaleva, Alexey Romanov, Anna Rogers, and Anna Rumshisky. 2019. [Revealing the dark secrets of BERT](#). 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)*, pages

- 4365–4374, Hong Kong, China. Association for Computational Linguistics.
- Brenden Lake and Marco Baroni. 2018. [Generalization without systematicity: On the compositional skills of sequence-to-sequence recurrent networks](#). In *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 2873–2882. PMLR.
- Andrew K. Lampinen and James L. McClelland. 2020. [Transforming task representations to perform novel tasks](#). *Proceedings of the National Academy of Sciences*, 117(52):32970–32981.
- Guillaume Lample, Alexis Conneau, Marc’Aurelio Ranzato, Ludovic Denoyer, and Hervé Jégou. 2018. [Word translation without parallel data](#). In *International Conference on Learning Representations*.
- Dongfang Li, zhenyu liu, Xinshuo Hu, Zetian Sun, Baotian Hu, and Min Zhang. 2024. [In-context learning state vector with inner and momentum optimization](#). In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*.
- Kenneth Li, Oam Patel, Fernanda Viégas, Hanspeter Pfister, and Martin Wattenberg. 2023a. [Inference-time intervention: Eliciting truthful answers from a language model](#). In *Thirty-seventh Conference on Neural Information Processing Systems*.
- Yingcong Li, Muhammed Emrullah Ildiz, Dimitris Papailiopoulos, and Samet Oymak. 2023b. [Transformers as algorithms: Generalization and stability in in-context learning](#). In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pages 19565–19594. PMLR.
- Yingcong Li, Kartik Sreenivasan, Angeliki Giannou, Dimitris Papailiopoulos, and Samet Oymak. 2023c. [Dissecting chain-of-thought: A study on compositional in-context learning of mlps](#). *arXiv preprint arXiv:2305.18869*.
- Yongjie Lin, Yi Chern Tan, and Robert Frank. 2019. [Open sesame: Getting inside BERT’s linguistic knowledge](#). In *Proceedings of the 2019 ACL Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 241–253, Florence, Italy. Association for Computational Linguistics.
- Sheng Liu, Haotian Ye, Lei Xing, and James Zou. 2024. [In-context vectors: Making in context learning more effective and controllable through latent space steering](#). *Preprint*, arXiv:2311.06668.
- Kevin Meng, David Bau, Alex Andonian, and Yonatan Belinkov. 2022. [Locating and editing factual associations in gpt](#). In *Advances in Neural Information Processing Systems*, volume 35, pages 17359–17372. Curran Associates, Inc.
- Jack Merullo, Carsten Eickhoff, and Ellie Pavlick. 2024. [Language models implement simple word2vec-style vector arithmetic](#). *Preprint*, arXiv:2305.16130.
- Sewon Min, Xinxu Lyu, Ari Holtzman, Mikel Artetxe, Mike Lewis, Hannaneh Hajishirzi, and Luke Zettlemoyer. 2022. [Rethinking the role of demonstrations: What makes in-context learning work?](#) pages 11048–11064.
- Jesse Mu, Xiang Lisa Li, and Noah Goodman. 2023. [Learning to compress prompts with gist tokens](#). In *Thirty-seventh Conference on Neural Information Processing Systems*.
- Kim Anh Nguyen, Sabine Schulte im Walde, and Ngoc Thang Vu. 2017. [Distinguishing antonyms and synonyms in a pattern-based neural network](#). In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 1, Long Papers*, pages 76–85, Valencia, Spain. Association for Computational Linguistics.
- Jane Pan, Tianyu Gao, Howard Chen, and Danqi Chen. 2023. [What in-context learning “learns” in-context: Disentangling task recognition and task learning](#). In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 8298–8319, Toronto, Canada. Association for Computational Linguistics.
- Abhishek Panigrahi, Nikunj Saunshi, Haoyu Zhao, and Sanjeev Arora. 2023. [Task-specific skill localization in fine-tuned language models](#). In *Proceedings of the 40th International Conference on Machine Learning, ICLR’23*. JMLR.org.
- Judea Pearl. 2001. [Direct and indirect effects](#). In *Proceedings of the Seventeenth Conference on Uncertainty in Artificial Intelligence, UAI’01*, page 411–420, San Francisco, CA, USA. Morgan Kaufmann Publishers Inc.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. 2019. [Language models are unsupervised multitask learners](#). OpenAI Blog.
- Emily Reif, Ann Yuan, Martin Wattenberg, Fernanda B. Viégas, Andy Coenen, Adam Pearce, and Been Kim. 2019. [Visualizing and measuring the geometry of bert](#). In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc.
- Nina Rimskey, Nick Gabrieli, Julian Schulz, Meg Tong, Evan Hubinger, and Alexander Matt Turner. 2024. [Steering llama 2 via contrastive activation addition](#). *Preprint*, arXiv:2312.06681.
- Nan Shao, Zefan Cai, Hanwei xu, Chonghua Liao, Yanan Zheng, and Zhilin Yang. 2023. [Compositional task representations for large language models](#). In *The Eleventh International Conference on Learning Representations*.
- Suzanna Sia, David Mueller, and Kevin Duh. 2024. [Where does in-context learning “happen” in large language models?](#) In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*.

- Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D. Manning, Andrew Ng, and Christopher Potts. 2013a. [Recursive deep models for semantic compositionality over a sentiment treebank](#). In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 1631–1642, Seattle, Washington, USA. Association for Computational Linguistics.
- Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D. Manning, Andrew Ng, and Christopher Potts. 2013b. [Recursive deep models for semantic compositionality over a sentiment treebank](#). In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 1631–1642, Seattle, Washington, USA. Association for Computational Linguistics.
- Nishant Subramani, Nivedita Suresh, and Matthew Peters. 2022. [Extracting latent steering vectors from pretrained language models](#). In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 566–581, Dublin, Ireland. Association for Computational Linguistics.
- Alon Talmor, Jonathan Herzig, Nicholas Lourie, and Jonathan Berant. 2019. [CommonsenseQA: A question answering challenge targeting commonsense knowledge](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4149–4158, Minneapolis, Minnesota. Association for Computational Linguistics.
- Erik F. Tjong Kim Sang and Fien De Meulder. 2003. [Introduction to the CoNLL-2003 shared task: Language-independent named entity recognition](#). In *Proceedings of the Seventh Conference on Natural Language Learning at HLT-NAACL 2003*, pages 142–147.
- Eric Todd, Millicent Li, Arnab Sen Sharma, Aaron Mueller, Byron C Wallace, and David Bau. 2024. [LLMs represent contextual tasks as compact function vectors](#). In *The Twelfth International Conference on Learning Representations*.
- Alexander Matt Turner, Lisa Thiergart, David Udell, Gavin Leech, Ulisse Mini, and Monte MacDiarmid. 2023. [Activation addition: Steering language models without optimization](#). *Preprint*, arXiv:2308.10248.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.
- Jesse Vig, Sebastian Gehrmann, Yonatan Belinkov, Sharon Qian, Daniel Nevo, Yaron Singer, and Stuart Shieber. 2020. [Investigating gender bias in language models using causal mediation analysis](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 12388–12401. Curran Associates, Inc.
- Pauli Virtanen, Ralf Gommers, Travis E. Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, Stéfan J. van der Walt, Matthew Brett, Joshua Wilson, K. Jarrod Millman, Nikolay Mayorov, Andrew R. J. Nelson, Eric Jones, Robert Kern, Eric Larson, C J Carey, İlhan Polat, Yu Feng, Eric W. Moore, Jake VanderPlas, Denis Laxalde, Josef Perktold, Robert Cimrman, Ian Henriksen, E. A. Quintero, Charles R. Harris, Anne M. Archibald, Antônio H. Ribeiro, Fabian Pedregosa, Paul van Mulbregt, and SciPy 1.0 Contributors. 2020. [SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python](#). *Nature Methods*, 17:261–272.
- Elena Voita, Pavel Serdyukov, Rico Sennrich, and Ivan Titov. 2018. [Context-aware neural machine translation learns anaphora resolution](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1264–1274, Melbourne, Australia. Association for Computational Linguistics.
- Elena Voita, David Talbot, Fedor Moiseev, Rico Sennrich, and Ivan Titov. 2019. [Analyzing multi-head self-attention: Specialized heads do the heavy lifting, the rest can be pruned](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5797–5808, Florence, Italy. Association for Computational Linguistics.
- Johannes Von Oswald, Eyvind Niklasson, Ettore Randazzo, Joao Sacramento, Alexander Mordvintsev, Andrey Zhmoginov, and Max Vladymyrov. 2023. [Transformers learn in-context by gradient descent](#). In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pages 35151–35174. PMLR.
- Ben Wang and Aran Komatsuzaki. 2021. Gpt-j-6b: A 6 billion parameter autoregressive language model. <https://github.com/kingoflolz/mesh-transformer-jax>.
- Lean Wang, Lei Li, Damai Dai, Deli Chen, Hao Zhou, Fandong Meng, Jie Zhou, and Xu Sun. 2023a. [Label words are anchors: An information flow perspective for understanding in-context learning](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 9840–9855, Singapore. Association for Computational Linguistics.
- Xiaozhi Wang, Kaiyue Wen, Zhengyan Zhang, Lei Hou, Zhiyuan Liu, and Juanzi Li. 2022. [Finding skill neurons in pre-trained transformer-based language models](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 11132–11152, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Xinyi Wang, Wanrong Zhu, Michael Saxon, Mark Steyvers, and William Yang Wang. 2023b. [Large language models are implicitly topic models: Explaining](#)

and finding good demonstrations for in-context learning. In *Workshop on Efficient Systems for Foundation Models @ ICML2023*.

Jerry Wei, Jason Wei, Yi Tay, Dustin Tran, Albert Webson, Yifeng Lu, Xinyun Chen, Hanxiao Liu, Da Huang, Denny Zhou, and Tengyu Ma. 2023. *Larger language models do in-context learning differently*. Preprint, arXiv:2303.03846.

Kevin Christian Wibisono and Yixin Wang. 2024. *From unstructured data to in-context learning: Exploring what tasks can be learned and when*. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*.

Sarah Wiegrefe and Yuval Pinter. 2019. *Attention is not not explanation*. 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)*, pages 11–20, Hong Kong, China. Association for Computational Linguistics.

Noam Wies, Yoav Levine, and Amnon Shashua. 2023. *The learnability of in-context learning*. In *Thirty-seventh Conference on Neural Information Processing Systems*.

Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi 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 Rush. 2020. *Transformers: State-of-the-art natural language processing*. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.

Sang Michael Xie, Aditi Raghunathan, Percy Liang, and Tengyu Ma. 2022. *An explanation of in-context learning as implicit bayesian inference*. In *International Conference on Learning Representations*.

Kang Min Yoo, Junyeob Kim, Hyuhng Joon Kim, Hyunsoo Cho, Hwiyeol Jo, Sang-woo Lee, Sang-goo Lee, and Taeuk Kim. 2022. *Ground-truth labels matter: A deeper look into input-label demonstrations*. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 2422–2437, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.

Xiang Zhang, Junbo Zhao, and Yann LeCun. 2015. *Character-level convolutional networks for text classification*. In *Advances in Neural Information Processing Systems*, volume 28. Curran Associates, Inc.

Yufeng Zhang, Fengzhuo Zhang, Zhuoran Yang, and Zhaoran Wang. 2023. *What and how does in-context learning learn? bayesian model averaging, parameterization, and generalization*. Preprint, arXiv:2305.19420.

A Extended Related Work

A substantial body of research on ICL has been ongoing since its discovery. We review previous studies through different facets of ICL. While our study shares similarities with others, it most closely aligns with the studies described in Section 2, where complementary details are provided in the final paragraph of this section.

In-Context Learning The ICL capabilities of LLMs were first identified by Brown et al. (2020). Since then, ICL has been extensively studied from various angles. Min et al. (2022); Yoo et al. (2022) have examined the effects of different ICL prompts styles. ICL during inference time has also been explored through meta-learning analyses (Akyürek et al., 2023; Dai et al., 2023; Von Oswald et al., 2023; Li et al., 2023b; Garg et al., 2023). In addition, investigations into ICL task inference from a Bayesian perspective have been conducted (Xie et al., 2022; Wang et al., 2023b; Wies et al., 2023; Zhang et al., 2023). Lastly, the scalability of ICL across different model sizes has been examined by Wei et al. (2023); Wang et al. (2023b); Pan et al. (2023). While these studies primarily focus on the externally observable behaviors of models during inference and ICL, our study delves into the internal mechanisms of transformers to encode ICL tasks.

The Role of Attention Mechanism in Explaining Model Behavior Past analyses of the attention mechanism (Voita et al., 2018; Clark et al., 2019; Voita et al., 2019; Kovaleva et al., 2019; Reif et al., 2019; Lin et al., 2019; Htut et al., 2019; Kobayashi et al., 2020) have revealed that attention weights often align with linguistic structures. However, these studies primarily focused on explaining the behavior of bidirectional architectures. Moreover, attention scores alone have not been found to fully explain the model’s outputs (Jain and Wallace, 2019; Wiegrefe and Pinter, 2019; Bibal et al., 2022). In our work, we aim to deepen the understanding of the role of multi-head self-attention in ICL. Specifically, we investigate the contribution of each attention head to the model’s internal representation of the ICL task.

Mechanisms to Explain Task Performance in In-Context Learning The components of transformers during ICL inference have been investigated to identify the origins of incorrect predictions and false statements (Merullo et al., 2024;

Halawi et al., 2024). Similarly, numerous studies have adjusted attention mechanisms or activations of hidden layers during inference to steer model behavior (Li et al., 2023a; Subramani et al., 2022; Turner et al., 2023; Rimskey et al., 2024; Liu et al., 2024). It was observed that tokens representing labels in an ICL prompt might hold the semantic information crucial for the language model’s final prediction (Wang et al., 2023a). Moreover, it was suggested that certain neurons within pre-trained transformers are highly predictive of task labels and empirically demonstrate that these neurons encode task-specific skills (Wang et al., 2022). A similar effort has been made to identify the critical layers where task information is stored (Sia et al., 2024). It has been shown that critical layers vary across tasks and can even differ within the same task, such as between subtasks like English-to-French and French-to-English translation. In contrast, our work aims to develop a principled conceptualization—such as a function of the transformer’s components—that can effectively *represent* and *automatically differentiate* across a wide range of tasks, regardless of task modality. We further leverage this conceptualization to guide the language model’s behavior across diverse tasks, aligning with prior works discussed next.

Tasks Representations in In-Context Learning

Efforts to represent ICL tasks typically start by deriving a task vector from the layer activations (Hendel et al., 2023), referred to as the *state*, at the position of a dummy query’s separator token. This token separates inputs and outputs in the in-context examples using specific symbols such as “→”. The dummy query is strategically placed immediately before the actual query to ensure that the task vector remains independent of the query itself. This method was further refined by averaging the states at each of the N separator tokens in an N -shot prompt, followed by optimization using inner and momentum techniques (Li et al., 2024). However, these approaches are limited to ICL prompts that incorporate a separator token (e.g., “→”), restricting their broader applicability. Moreover, it has been shown that the principal direction of layer activation differences can effectively guide the ICL task (Liu et al., 2024), leading to the task representation known as *In-Context Vectors* (ICVs). Nonetheless, it has been recently argued that the focus should be on attention heads (Todd et al., 2024), as these are crucial for transferring information between to-

ken positions (Vaswani et al., 2017; Elhage et al., 2021b).

Deriving from this finding, the *Function Vector* (FV) is computed as the sum of the outputs from a selectively chosen subset of attention heads based on an *indirect metric* (Todd et al., 2024), derived from causal inference literature (Pearl, 2001). To the best of our knowledge and based on our preliminary analyses, the most effective empirical representation of tasks in ICL are FVs. Therefore, our approach starts by deriving from Todd et al. (2024), in contrast to methods the outlined by Hendel et al. (2023); Liu et al. (2024). Instead of merely using raw activations, we optimize weights assigned to these heads to enhance transformer performance in scenarios where ICL cannot be performed. Ultimately, this leads to a more formalized conceptualization that can be adapted to various models and tasks, whether functional regression or linguistic.

B Discussions of Alternative Approaches to Assess Task Representations

Section B.1 explains why shuffling labels in regression tasks does not prevent the model from extracting task information. Section B.2 discusses why probing is ineffective for assessing the precision of task representations.

B.1 Why Shuffling Labels is not a Valid Approach for Functional Regression

We analyze the effect of label shuffling on in-context learning performance for two example tasks: linguistic antonym generation and linear regression. While label shuffling disrupts the model’s ability to infer relationships in the linguistic task, its impact on the numeric task differs fundamentally due to the task’s nature and the model’s pre-training, particularly if the model was pre-trained on similar functional regression tasks (Garg et al., 2023).

B.1.1 Linguistic Antonym Generation

In the linguistic task, the model is provided with prompts such as:

cold → hot, happy → sad, fast → ?

The model’s goal is to generate the antonym of “slow” based on in-context examples. Since LMs are pre-trained on vast text corpora without specific emphasis on antonym pairs, they rely heavily on these examples to infer the relationship. Shuffling

the labels in the examples disrupts the pattern, preventing the model from identifying the intended relationship (antonymy) and generating the correct output. Any potential performance degradation in this case is directly linked to the blockage of ICL caused by label shuffling.

B.1.2 Numeric Linear Regression

In numeric tasks, the model processes sequences where each scalar output y_i is generated from a vector input x_i through a linear relationship $y_i = w^\top x_i$:

$$\begin{aligned} [x_{1,1}, x_{1,2}, \dots, x_{1,m}] &\rightarrow y_1, \\ [x_{2,1}, x_{2,2}, \dots, x_{2,m}] &\rightarrow y_2, \\ &\vdots \\ [x_{T,1}, x_{T,2}, \dots, x_{T,m}] &\rightarrow ?, \end{aligned}$$

where the m -dimensional input vector is flattened before being fed into the model, following our setup. The model’s goal is to infer the coefficient vector w from the examples to predict the missing output. Pre-training embeds a strong bias toward linear mappings in the model’s parameters, giving it a default assumption of linearity. In-context examples refine the *exact coefficients* but are not the sole source of understanding this relationship. Consequently, shuffling the labels conflicts with the model’s pre-trained understanding of input-output relationships (*i.e.*, its inherent linear bias), leading to confusion rather than simply blocking ICL. In such cases, the shuffled prompt would still need to define a linear relationship, which is non-trivial. In contrast, for language tasks, the model retains an understanding of individual word semantics even when labels are shuffled, aligning with its pre-trained understanding of language.

B.2 Why Probing Would be Ineffective to Assess Task Representations

We outline a potential approach for probing experiments. One method is to assess whether the LTV encodes task information by using it to predict task-relevant quantities—in this case, quantities in regression tasks or correct labels in language tasks. This prediction is performed by training a linear predictor or a two-layer neural network.

However, we argue that even if the probe achieves high accuracy, this does not demonstrate that the LTV (or any input used for probing) truly encodes task information. Probing offers only indirect evidence; it shows that a mapping exists from

the input to the target quantity, not that the input itself is a meaningful representation of the task. It simply indicates that the input contains enough information to predict the output.

In an extreme case, we could apply an arbitrary invertible transformation to the prompt and use that as input to the probe. The probe would likely achieve good performance because the prompt contains sufficient information to predict the task-relevant output. However, this input would clearly not be a meaningful representation, as it is too coarse.

In our case, the LTV is a representation specific to the LLM architecture. Moreover, we are not concerned with whether it can directly predict the task-relevant output—this is not how these vectors are intended to be used. Instead, they are applied to latent layers of the LLM to steer output generation, which is exactly how we evaluate their effectiveness.

C Experimental Details

C.1 Implementation

For all models, we use the huggingface implementations (Wolf et al., 2020). For regression tasks and GPT-2 pre-training, we rely on the code provided by the authors on GitHub². Additionally, we use the authors’ code³ to include FV into our experiments. The remaining baseline methods—ICV, TV, and SV—were implemented by us.

C.2 Experiments on Regression Tasks

We closely follow the experimental setup established by Garg et al. (2023). For completeness, we provide the relevant details here, with additional information available in the cited reference.

C.2.1 Model

The GPT-2 model processes sequences of vectors in the embedding space and outputs a sequence in the same space. However, the tasks we examine involve functions mapping a lower-dimensional vector space (*e.g.*, 20 dimensions) to a scalar value. To construct a prompt such as $p = \{x_1, f(x_1), x_2, f(x_2), \dots, x_{\text{query}}\}$, we must map x_i and $f(x_i)$ into the embedding space. This mapping involves first converting the scalar values $f(x_i)$ into vectors of the same dimension as x_i by appending zeros, followed by applying a learnable

²<https://github.com/dtsip/in-context-learning>

³https://github.com/ericwtodd/function_vectors

linear transformation to all these vectors into the embedding space. The model’s output vector is then transformed into a scalar value through a dot product with a learnable vector.

We consider the model’s prediction at the position corresponding to x_i (*i.e.*, the absolute position $2i - 1$) as the prediction of $f(x_i)$. Due to the model’s structure, this prediction relies solely on the pairs $\{x_j, f(x_j)\}$ for $j < i$ and x_i itself. We disregard the model predictions at positions corresponding to $f(x_i)$.

The GPT-2 models use absolute, learnable positional encodings and were trained to accommodate up to 101 examples for linear and sparse linear functions, and up to 201 examples for 2-layer ReLU neural networks in a prompt. While it is possible to feed the model prompts with more examples by adjusting the initialization, this would exceed our computational resources. We also did not want to alter the nature of their experimental process.

C.2.2 Training

We train a model from scratch (*i.e.*, no pre-trained weights are loaded) to predict $f(x_i)$ for a given x_i , using the set of examples as reference. Each training prompt is generated by randomly sampling a function f from the function class of interest, followed by sampling inputs x_i from an isotropic Gaussian distribution $\mathcal{N}(0, I_m)$. The prompt is constructed as $p = \{x_1, f(x_1), \dots, x_k, f(x_k)\}$. For each input $i \leq k$ within a prompt, the model predictions $\hat{y}_i = M_\theta(x_i \mid p^f = \{x_1, f(x_1), \dots, x_{i-1}, f(x_{i-1})\})$ are obtained, and the loss is computed across all prompt prefixes:

$$\min_{\theta} \mathbb{E}_{f \sim \mathcal{D}_F, x \sim \mathcal{D}_X} \left[\frac{1}{T+1} \sum_{i=0}^T \left(M_\theta(p^{f,i}) - f(x_{i+1}) \right)^2 \right],$$

where \mathcal{L} is the loss function, typically chosen to be mean squared error, and we have $x_{T+1} = x_{\text{query}}$.

During training, we average the loss across a batch of randomly generated prompts, each with different functions and inputs, and update the model parameters. The Adam optimizer (Kingma and Ba, 2015) is employed and trained for a total of 500,000 steps with a batch size of 64, using a learning rate of 10^{-4} for all function classes and models.

Curriculum Learning The training procedure is accelerated through curriculum learning. The model starts by observing prompt inputs x_i within

a smaller dimensional subspace and with fewer inputs per prompt. Both the subspace dimension and the number of examples are increased gradually. Specifically, all of the coordinates except the first $T_{\text{max, cur}}$ of x_i are zeroed out by sampling prompts of size T_{cur} . For the function classes of linear and sparse linear functions, $T_{\text{max, cur}} = 5$ and $T_{\text{cur}} = 11$ are used initially, and $T_{\text{max, cur}}$ and T_{cur} are increased by 1 and 2, respectively, every 2000 steps until reaching $T_{\text{max, cur}} = T_{\text{max}}$ and $T_{\text{cur}} = 2m + 1$. A different schedule is applied for 2-layer neural networks to accommodate the need for more inputs; starting from $T_{\text{max, cur}} = 5$ and $T_{\text{cur}} = 26$, $T_{\text{max, cur}}$ and T_{cur} are incremented by 1 and 5 respectively, every 2000 steps until $T_{\text{max, cur}} = T_{\text{max}}$ and $T_{\text{cur}} = 5m + 1$.

Consequently, in the curriculum-based training approach, a training prompt

$$p^f = \{x_1, f(x_1), \dots, x_{T_{\text{cur}}}, f(x_{T_{\text{cur}}})\}$$

is generated by sampling a random function f from the function class and drawing inputs x_i by sampling i.i.d. from $\mathcal{N}(0, I_m)$, with all but the first $T_{\text{max, cur}}$ coordinates zeroed out. Given the model predictions \hat{y}_i , the loss is computed as

$$\frac{1}{T_{\text{cur}}} \sum_{i=1}^{T_{\text{cur}}} (\hat{y}_i - f(x_i))^2.$$

C.2.3 Sampling the Functions

For the class of linear functions, we sample the random function $f(x) = w^\top x$ by drawing $w \sim \mathcal{N}(0, I_m)$. In the case of sparse linear functions, w is also sampled from $\mathcal{N}(0, I_m)$, but we then randomly zero out the first T_{cur} coordinates within the first $T_{\text{max, cur}}$ coordinates. For these linear functions, we set $m = 20$ for all experiments, with a sparsity level of 3. For 2-layer neural networks, we sample W_1 from $\mathcal{N}(0, I_m)$ and W_2 from $\mathcal{N}(0, 2/r)$, where $f(x) = W_2 \text{ReLU}(W_1 x)$. Here, we set the dimensions $m = 20$ and the ratio $r = 100$.

We also experimented with decision trees as a function class, as done by Garg et al. (2023). However, we found that decision trees exhibit less sensitivity to positional encodings. This is because their feature-based reasoning relies on discrete hierarchical rules rather than point-by-point mappings. By grouping multiple inputs under the same label, decision trees largely bypass positional dependencies and are minimally influenced by untrained positional encodings.

C.2.4 Evaluation

To assess performance, we sample a prompt with a maximum length of T_{\max} , which is equal to 101 for linear and sparse linear functions and 201 for 2-layer networks. We then trim the prompt to $T_i \leq T_{\max}$ demonstrations and track the prediction errors for each $i \leq T_{\max}$. Consequently, each point in our error curves corresponds to the error at a specific prompt length i . This analysis is conducted over batches of 256 prompts, with the average error reported. We have determined that batches larger than 256 prompts do not significantly alter the results, confirming that 256 prompts are sufficient to produce generalized results.

C.2.5 Optimizing the Baseline Methods

Function Vector (FV) and In-Context Vector (ICV) were originally designed for language data. Preliminary results showed that FV had a notable impact on regression data, with potential for further improvement through fine-tuning. In contrast, ICV slightly underperformed the vanilla model on regression tasks when using the default scaling parameter $\lambda = 0.1$ from language experiments. This parameter, dominated by activation magnitudes, required tuning for regression data.

Our approach extends the original by integrating FV across multiple GPT-2 layers, leading to improved performance. To minimize output disruption and reinforce tasks, we added *dummy* examples (e.g., $\{0, 0\}$ pairs) at specific prompt positions. Increasing the number of attention heads for FV computation from 10 to 35 further enhanced performance. In contrast, ICV already integrates the task vector at every token position and computes a distinct task vector for each layer, similar to our method. Thus, the only optimization for ICV was tuning the scaling parameter. We found that FV required no scaling (i.e., 1.0). However, ICV often degraded the vanilla model’s performance, with a scale of 1.5 being the best option to mitigate this issue. Key baseline optimizations are summarized below:

1. **Function Vector added to multiple layers:** Distributing FV across multiple layers normalized its impact on activations. Layers 6, 7, and 8 of GPT-2 were the most effective.
2. **Dummy tokens:** Placing dummy tokens at 0.1, 0.25, 0.5, 0.75, and 0.9 fractions of the prompt length optimized performance for FV.

3. **Number of attention heads for FV computation:** Increasing to 35 attention heads maximized GPT-2 performance, with diminishing returns beyond this point.
4. **Scaling:** FV performed best with a scaling factor of 1.0, aligning with [Todd et al. \(2024\)](#) for language tasks. ICV required scaling (e.g., 1.5) to counterbalance the dominance of hidden state activations and prevent degradation of the vanilla model’s performance.

C.2.6 Other Baseline Methods

As explained in Appendix A, [Hendel et al.’s \(2023\)](#) and [Li et al.’s \(2024\)](#) methods cannot be included in our experiments because they require in-context prompts to include a separator token (e.g., “ \rightarrow ”) in each example. This requirement makes them incompatible with our functional regression dataset, as these tasks explicitly avoid tokenization and operate directly on raw floating-point numbers (see Appendix C.2 for details). Introducing separator tokens fundamentally changes the task by imposing a symbolic or tokenized representation of input-output relationships, contradicting [Garg et al.’s \(2023\)](#) setup. Moreover, such methods rely on the tokenized format, bypassing the challenge of inferring functional relationships from raw numeric data, which renders the comparison invalid.

C.3 Experiments on Language Tasks

We closely adhere to the experimental methods described by [Todd et al. \(2024\)](#), providing all relevant details to ensure our report is self-sufficient.

C.3.1 Datasets and Tasks

We provide a detailed overview of the natural language tasks used to evaluate the task formulations. References to the original data sources for each task are included, along with details on how they were refined and transformed into word pairs by [Todd et al. \(2024\)](#). The final datasets used in our evaluations are provided in the authors’ GitHub repository³.

AG News This text classification dataset consists of news headlines and the first few sentences of articles as inputs, with labels indicating the article’s category. Categories include Business, Science/Technology, Sports, and World ([Zhang et al., 2015](#)).

Antonym and Synonym The dataset is [Nguyen et al.’s \(2017\)](#) and was further refined by [Todd et al. \(2024\)](#). Initially, all adjective, noun, and verb pairs from the dataset splits were combined, with duplicate entries removed. The dataset was then modified to include only word pairs where both words are tokenized as single tokens. This process resulted in 2,398 antonym pairs and 2,881 synonym pairs, with a vocabulary size of $|\mathcal{V}| = 50,400$.

These datasets originally included multiple outputs for single inputs, e.g., “increase” \rightarrow “decrease” and “increase” \rightarrow “reduce.” However, handling such cases would require a more powerful model ([Todd et al., 2024](#)). Therefore, the dataset has been simplified to ensure a one-to-one mapping between terms.

CommonsenseQA This question-answering dataset requires a model to select the correct answer from five options, each labeled with a letter. The model generates the letter corresponding to the correct answer. For instance, given the question “Where is a business restaurant likely to be located?” and the options “A) town, B) hotel, C) mall, D) business sector, E) yellow pages,” the model should generate “D” ([Talmor et al., 2019](#)).

Landmark-Country The Landmark-Country dataset contains entries pairing the name of a landmark with its corresponding country. The data pairs are sourced from [Hernandez et al. \(2024\)](#).

Person-Instrument The Person-Instrument dataset contains entries pairing the name of a professional musician with the instrument they play. The data pairs are sourced from [Hernandez et al. \(2024\)](#).

Person-Occupation The Person-Occupation dataset, sourced from [Hernandez et al. \(2024\)](#), contains entries pairing the names of well-known individuals with their respective occupations.

Person-Sport The Person-Sport dataset, sourced from [Hernandez et al. \(2024\)](#), contains entries pairing the names of professional athletes with the sports they play.

Product-Company The Product-Company dataset, curated from [Hernandez et al. \(2024\)](#), contains entries pairing the names of commercial products with the companies that sell them.

Sentiment Analysis The sentiment analysis task is based on the Stanford Sentiment Treebank (SST-

2) dataset ([Socher et al., 2013b](#)), which consists of movie review sentences labeled as either “positive” or “negative.” For example, an entry from this dataset might be: “An extremely unpleasant film. \rightarrow negative.” We use the subset of SST-2 curated by [Honovich et al. \(2023\)](#), which excludes incomplete sentences and those with more than 10 words, resulting in 1,167 entries. For further details, see [Honovich et al.’s \(2023\)](#) study.

Translation The language translation dataset was constructed using data from [Lample et al. \(2018\)](#), which pairs English words with their translations in French and Spanish. The train and test splits were merged into a single dataset and the cognates were filtered out. This process results in 4,705 pairs for English-French and 5,200 pairs for English-Spanish. The original datasets included multiple translations for some input words. These duplicates were filtered using GPT-4, following the method used for Antonym and Synonym.

The tasks described so far are *Abstractive* NLP tasks, where the information to be generated is not present in the prompt. We also evaluate the opposite case: *Extractive* tasks, where the answer exists within the prompt, and the goal is to retrieve it.

CoNLL-2003 Our extractive tasks are based on a subset of the CoNLL-2003 English named entity recognition (NER) dataset ([Tjong Kim Sang and De Meulder, 2003](#)), a widely used NLP benchmark for evaluating NER models. The NER task involves extracting the correct entity from a given sentence based on a specific property. For our study, we use three tasks derived by [Todd et al. \(2024\)](#) from CoNLL-2003: NER-person, NER-location, and NER-organization, where the labels correspond to the name of a person, location, or organization, respectively. Each dataset is constructed by merging the CoNLL-2003 “train” and “validation” splits into a single dataset and filtering sentences to include only those with a single instance of the specified class. This reduces task ambiguity, as sentences with multiple instances of the same class could have more than one correct answer.

C.3.2 Prompting

The default template for prompting the GPT-J model with T demonstrations is structured as follows:

$$Q: \{x_1\} \setminus nA: \{y_1\} \setminus n \dots Q: \{x_{\text{query}}\} \setminus nA:$$

In our experiments with shuffled prompts, we randomly shuffle the labels $\{y_i\}_{i=1}^T$ among each other. For zero-shot prompts, which contain no demonstrations, prompts consist solely of the query: “Q: $\{x_{\text{query}}\}$ \nA:”. When testing Task Vector (Hendel et al., 2023) and State Vector (Li et al., 2024), which require a separator token between inputs and outputs, we treat “\n” as the separator token.

C.3.3 Evaluation

We uniformly sample in-context examples (for few-shot) and a query input $N = 256$ times, assessing performance as the accuracy if the first token generated by the model matches the first token of the query label (if the label consists of more than one token). To compute the margin of error (denoted after \pm in Tables 1 and 2), we first calculate the standard error of the mean (SEM) as follows:

$$\text{SEM} = \frac{\text{standard deviation}}{\sqrt{N}}.$$

Next, the t-critical value is determined using the t-distribution with a 95% confidence level and degrees of freedom ($N - 1$). The margin of error is then calculated as:

$$\text{Margin of error} = t_{\text{critical}} \times \text{SEM}.$$

For example, if the mean accuracy is 39.1 and the margin of error is 5.2, the true mean lies between 33.9 and 44.3 with a 95% confidence. Thus, if the confidence intervals of different methods overlap with the maximum mean accuracy, it indicates that the differences in performance are not statistically significant despite the highest accuracy.

C.4 Training Learnable Task Vector

The LTV parameters are initialized by sampling from the standard Gaussian distribution, totaling $L \times J$ learnable parameters. We did not employ additional techniques such as dropout, activation functions, or gradient clipping in learning the weights, which are neither clipped nor bounded. The Adam optimizer, with a learning rate of 5×10^{-5} , was used in all experiments. Training is terminated if the validation loss does not decrease for 50 consecutive gradient steps. The transformer parameters always remain frozen.

C.4.1 Learning from Regression Data

We optimize the LTV weights using mini-batch gradient descent on a dataset we compiled, consisting of $100 \times 256 = 25,600$ function samples (100

times the batch size). Prompts of length $T_v > T_{\text{train}}$ are constructed for these functions. We reserve the 20% of the dataset as a validation set for monitoring loss. For each mini-batch of prompts $\{p_i^{f_i}\}_{i=1}^N$ sampled from the dataset, the gradient descent step is defined as

$$\begin{aligned} \mathcal{L}(\Phi) &= \frac{1}{N} \sum_{i=1}^N \left(\tilde{M}_\theta(p_i^{f_i} \mid \mathbf{v}_\Phi^{f_i}) - f_i(x_{i,\text{query}}) \right)^2, \\ \Phi &\leftarrow \Phi - \eta \cdot \nabla \mathcal{L}(\Phi), \end{aligned}$$

where $p_i^{f_i}$ represents a prompt corresponding to a unique function f_i and η is the learning rate. This method and dataset compilation are applied uniformly across all three function classes.

We recognize that generating distinct functions and prompts at each gradient step could potentially provide an infinite variety of data and functions. This raises concerns about whether LTV overfits to the function classes. Although it was argued that the likelihood of the model encountering a training-similar prompt is extremely low (Garg et al., 2023), we opted for a static dataset approach. Our experiments were conducted using this dataset, with evaluations performed on the prompts constructed from different weight vectors during inference.

C.4.2 Learning from Language Data

Given the dataset P^t for task t , each gradient step involves sampling 100 prompts with replacement, each containing 5 demonstrations. These prompts are processed by the transformer to collect attention activations, and the mean of these activations across the sampled prompts is computed. This mean is then passed through the LTV layer to compute the corresponding LTV, as illustrated in (2).

For training, we sample a batch of 32 prompts, each with 5 demonstrations, but with all labels shuffled, rendering the input-output pairs non-informative. The LTV weights are updated to maximize the probability of the correct label for the query input:

$$\mathcal{L}(\Phi) = -\frac{1}{N} \sum_{i=1}^N \log \left(\tilde{M}_\theta(y_{\text{query}} \mid \tilde{p}_i^t; \mathbf{v}_\Phi^t) \right).$$

C.5 Training LoRA

We use LoRA (Low-Rank Adaptation) with a rank of 8, applied to the key (K), query (Q), value (V), and output projection matrices in the self-attention mechanism of the GPT-2 model (12 layers, 8 attention heads per layer, and 256 hidden dimensions).

The GPT-2 model is fine-tuned with LoRA on the same static dataset and over the same number of iterations (2000) used for training LTV.

Each of these matrices originally has dimensions 256×256 . LoRA introduces two low-rank matrices, $A \in \mathbb{R}^{256 \times 8}$ and $B \in \mathbb{R}^{8 \times 256}$, replacing the trainable weight matrix with their product. This setup results in 4096 trainable parameters per matrix, calculated as $2 \times 256 \times 8$.

Since our task is relatively simple and the model is extremely lightweight, we apply LoRA to all layers for comprehensive coverage. This adds 16,384 trainable parameters per layer (4×4096), as LoRA operates on four matrices (K , Q , V , and output) in each transformer layer. With 12 transformer layers in total, the model fine-tunes 196,608 parameters (12×16384). This parameter-efficient approach significantly reduces the training cost compared to updating all model parameters.

C.6 Detailed Ablation Studies

Our argument is based on the premise that while earlier layers build foundational representations, the most refined and actionable insights for predictions are concentrated in the outputs of the last layer. The motivation for these studies is to examine the model’s resilience to variations in prompt length.

To this end, we freeze all transformer and LTV parameters to analyze the stability of last hidden state distributions across varying prompt lengths, rendering stationary distribution. We generate a dataset of 25,600 prompts (100 times the batch size) with a maximum length T_{\max} . Prompts were trimmed to T_{train} for the vanilla transformer and extended beyond T_{train} for the FV- and LTV-integrated transformers. The last hidden states from these configurations were compiled into two datasets, X_1 and X_2 , corresponding to the vanilla transformer and the FV- and LTV-integrated transformers, respectively. The samples within each dataset are independent and identically distributed, allowing us to estimate their probability distributions.

We estimate the probability distributions using Kernel Density Estimation (KDE), employing the standard KDE implementation from `scipy` (Virtanen et al., 2020) with default settings. However, directly estimating a probability distribution in a high-dimensional space often leads to the *curse of dimensionality*, where the volume of data required to effectively estimate the distribution grows exponentially with the number of dimensions. A

practical solution to this challenge is to employ SVD for dimensionality reduction. This involves decomposing the data matrices $X_i \in \mathbb{R}^{M \times d}$ as:

$$\begin{aligned} X_1 &= U_1 \Sigma_1 V_1^\top, \\ X_2 &= U_2 \Sigma_2 V_2^\top, \end{aligned}$$

where M is the number of collected samples and U_i contains the principal components $u_{i,1}, u_{i,2}, \dots, u_{i,n}$ as column vectors. These principal components (PCs) form the column space of X_i :

$$\text{span}(u_{i,1}, u_{i,2}, \dots, u_{i,n}) = \text{colspace}(X_i),$$

where each $u_{i,k}$ is orthogonal to $u_{i,k' \neq k}$ and ordered by decreasing variance that they explain. Specifically, the first n principal components represent the directions along which the data varies the most, capturing the most significant patterns in the data. These components are likely more informative and relevant for distinguishing different behaviors or properties of the data.

Rather than estimating the distribution underlying the entire datasets X_i as multivariate distributions, we employ Gaussian KDE⁴ to estimate each PC as a unimodal distribution. This approach is advantageous since KDE performs better with univariate data. However, transitioning from multivariate to univariate requires the assumption that the PCs are uncorrelated. We validate this assumption by observing that the nondiagonal entries of the correlation matrices of U_i are on the order of 10^{-2} , with diagonal entries being approximately 1, effectively an identity function, confirming that the PCs are indeed uncorrelated.

We use KL divergence between the KDE-estimated distributions of the column pairs to quantitatively assess distributional similarities. However, we find that these divergence values are negligibly small, except for the first one, which accounts for the most variance within the dataset. The negligible divergence scores for higher-order components suggest that these vectors contribute minimally to differentiating the datasets. Thus, focusing on the first component, which shows substantial divergence, is statistically justified and highlights the critical variations relevant to model generalization.

This experimentation process is depicted in Figure 4. Additionally, histograms illustrating the

⁴https://docs.scipy.org/doc/scipy/reference/generated/scipy.stats.gaussian_kde.html

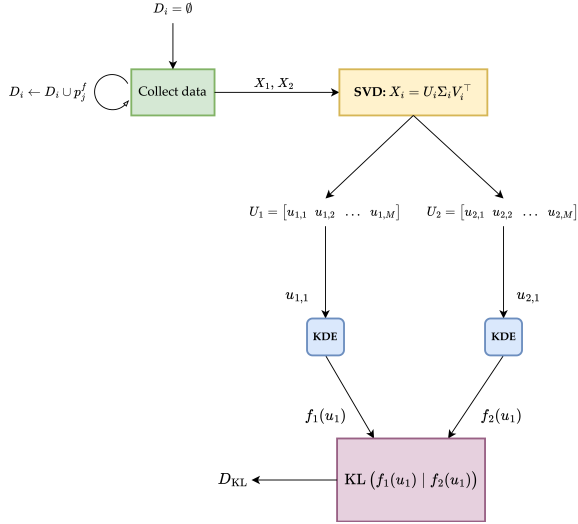


Figure 4: The diagram illustrates our pipeline for ablation studies. We start by collecting $M = 25,600$ prompts corresponding to a selected task f . Subsequently, the first principal component of the column space of these datasets is extracted using SVD. Finally, we report the KL divergence between the KDE-estimated distributions of these components.

KDE-estimated distributions of the first principal components are provided in Appendix D.5 to offer a clear view of their similarity.

C.7 Computational Resources

The computational experiments were conducted on a high-performance system with an AMD Ryzen Threadripper PRO 3995WX processor and 515 GB of RAM. For GPU acceleration, two NVIDIA RTX A6000 GPUs were used, each with 49 GB of memory. This setup efficiently supported the transformer models used in our research.

D Complete Set of Results

We present the complete set of results that could not be included in the main body due to page limitation. Specifically, this section includes:

1. **Regression Tasks:** Evaluations with different LTV training lengths T_v
2. **Regression Tasks:** Evaluations with chain-of-thought prompting based on the findings of Li et al. (2023c)
3. **Regression Tasks:** Evaluations under distributional shifts
4. **NLP Tasks:** Accuracy scores (with zero- and few-shot prompting) across 15 NLP benchmarks

5. **Ablation Studies:** Complete ablation results, including KL divergence scores across different T_v and the corresponding histograms illustrating the distributions of the optimally in-context-learned and LTV-integrated models

D.1 Evaluations on Regression Tasks

The loss curves are shown in Figures 5, 6, and 7. As expected, no performance improvement is observed for $T_v = T_{\text{train}}$, as the transformer already achieves optimal performance. For visual comparison, the LTV trained with the maximum number of examples is included, though it does not provide insights into generalizability, as it is exposed to all T_{max} examples.

D.2 Chain-of-Thought Prompting

Li et al. (2023c) investigate how CoT enhances ICL for compositional functions, formalizing it as a two-phase process: filtering relevant data and applying ICL on the filtered input. During filtering, the model focuses on the relevant tokens in the prompt based on a given instruction. It then processes the filtered prompt to generate an intermediate output before proceeding to the next step in the sequence.

Multi-layer networks can be viewed as compositional functions since the output depends on the hidden layers, which are functions of the input. Accordingly, Li et al. (2023c) specifically study 2-layer ReLU networks, as in our work. In this context, the relevant tokens at the k -th step correspond to the output of the k -th hidden layer. For instance, let z_k represent the k -th intermediate step. During inference, the relevant tokens for z_k are $\{z_{k-1}^i, z_k^i\}_{i=1}^T$, where T is the number of examples in the prompt.

Building on this, we demonstrate that LTV can generalize to capture latent representations at intermediate steps, enabling task fidelity in our length generalization benchmark using Li et al.’s (2023c)’s proposed CoT method. This provides a more flexible and scalable framework for compositional reasoning.

D.2.1 Experimental Details

Model Training As in Appendix C.2.2, we pre-train a custom GPT-2 model (with the same configuration) separately for the two variants of Li et al. (2023c)’s method: CoT-I and CoT-I/O. The task involves 2-layer ReLU network regression, *i.e.*, $f(x) =$

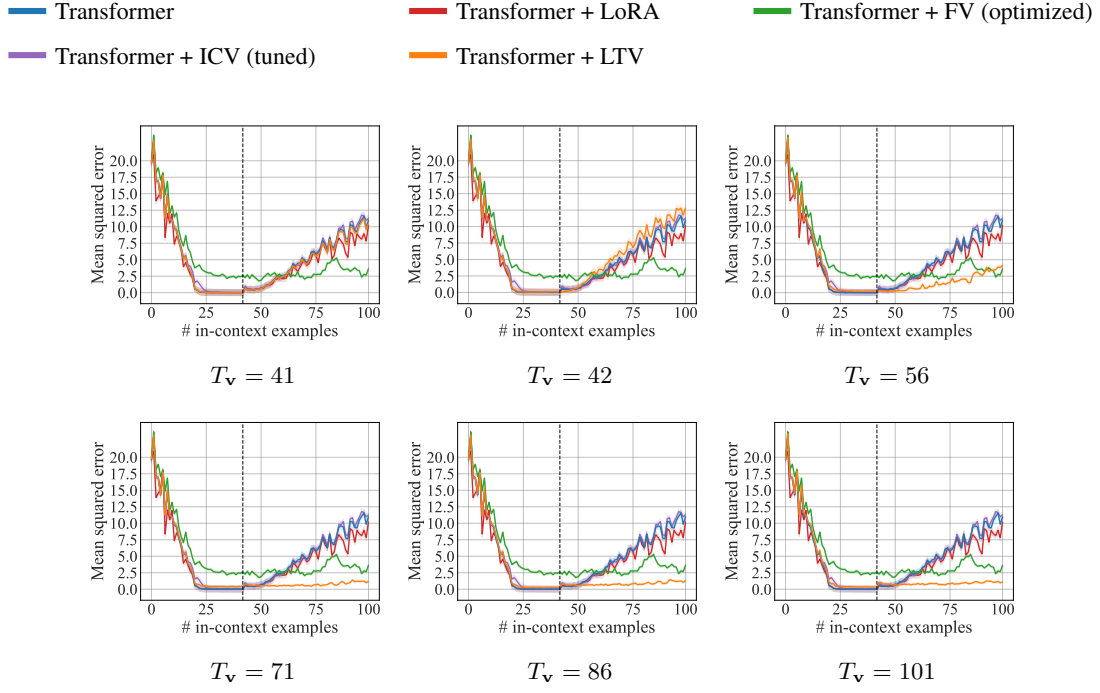


Figure 5: Evaluation on the class of linear functions, with the transformer pre-trained on up to $T_{\text{train}} = 41$ examples per prompt. Results are averaged over a batch of 256 randomly selected tasks. The shaded area represents the 95% confidence interval over the sampled prompts. T_v denotes the prompt length used during LTV training.

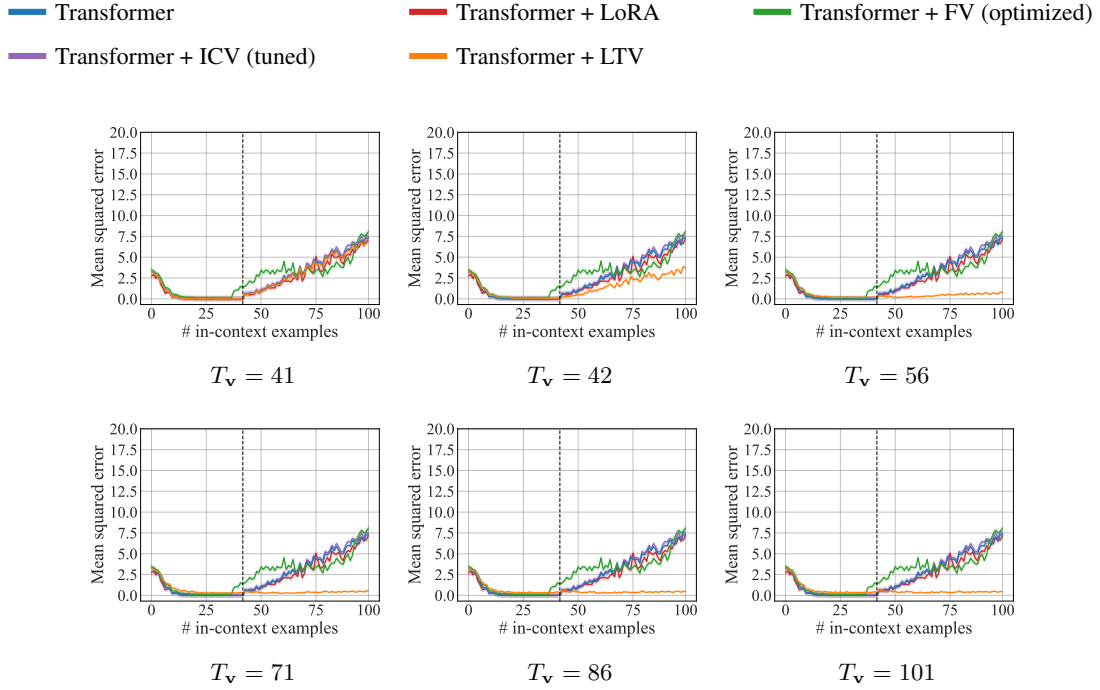


Figure 6: Evaluation on the class of sparse linear functions, with the transformer pre-trained on up to $T_{\text{train}} = 41$ examples per prompt. Results are averaged over a batch of 256 randomly selected tasks. The shaded area represents the 95% confidence interval over the sampled prompts. T_v denotes the prompt length used during LTV training.

$W_2\text{ReLU}(W_1x)$. The prompt is structured as $p^f = \{x_1, g(x_1), f(x_1), \dots, x_T, g(x_T), f(x_T)\}$,

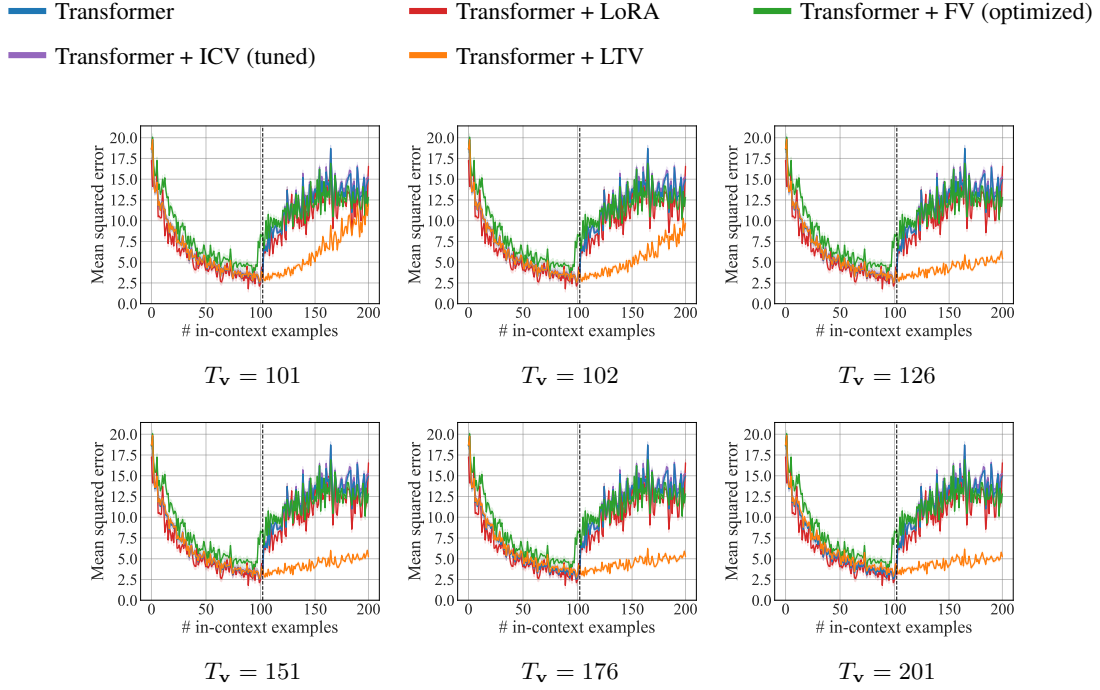


Figure 7: Evaluation on the class of 2-layer ReLU neural networks, with the transformer pre-trained on up to $T_{\text{train}} = 101$ examples per prompt. Results are averaged over a batch of 256 randomly selected tasks. The shaded area represents the 95% confidence interval over the sampled prompts. T_v denotes the prompt length used during LTV training.




Figure 8: Evaluation on the class of 2-layer ReLU neural networks, with the transformer pre-trained on up to $T_{\text{train}} = 101$ chain-of-thought examples per prompt. Results are averaged over a batch of 256 randomly selected tasks. The shaded area represents the 95% confidence interval over the sampled prompts. T_v denotes the prompt length used during LTV training.

where $g(x) = \text{ReLU}(W_1 x)$ represents an intermediate computation.

$$\begin{aligned} \hat{s}_i^{(1)} &= M_\theta(p_{\text{CoT}}(i)), \\ \hat{y}_i &= \hat{s}_i^{(2)} = M_\theta(\{p_{\text{CoT}}(i), \hat{s}_i^{(1)}\}). \end{aligned} \quad (5)$$

For each $i \leq T$, CoT-I/O performs iterative two- In contrast, CoT-I directly predicts the final output

in a single step:

$$\hat{y}_i = M_\theta(p_{\text{CoT}}(i)).$$

The objective functions for pre-training the GPT-2 model are expressed for CoT-I/O and CoT-I respectively as:

$$\begin{aligned} \min_{\theta} \mathbb{E}_{f \sim \mathcal{D}_{\mathcal{F}}, x \sim \mathcal{D}_{\mathcal{X}}} & \left[\frac{1}{T+1} \sum_{i=0}^T (\hat{s}_{i+1}^{(1)} - g(x_{i+1}))^2 \right. \\ & \left. + (\hat{s}_{i+1}^{(2)} - f(x_{i+1}))^2 \right], \\ \min_{\theta} \mathbb{E}_{f \sim \mathcal{D}_{\mathcal{F}}, x \sim \mathcal{D}_{\mathcal{X}}} & \left[\frac{1}{T+1} \sum_{i=0}^T (\hat{y}_{i+1} - f(x_{i+1}))^2 \right]. \end{aligned}$$

Optimizing the Learnable Task Vector In CoT-I/O, we train the LTV for each intermediate step. The objective functions for training LTVs for $g(x)$ and $f(x)$ are, respectively,

$$\begin{aligned} \min_{\Phi} \mathbb{E}_{f \sim \mathcal{D}_{\mathcal{F}}, x \sim \mathcal{D}_{\mathcal{X}}} & \left[\left(\tilde{M}_{\theta}(p_{\text{CoT}}(T_v + 1) \mid v_{\Phi}^g) \right. \right. \\ & \left. \left. - g(x_{\text{query}}) \right)^2 \right], \\ \min_{\Phi} \mathbb{E}_{f \sim \mathcal{D}_{\mathcal{F}}, x \sim \mathcal{D}_{\mathcal{X}}} & \left[\left(\tilde{M}_{\theta}(\{p_{\text{CoT}}(T_v + 1), g(x_{\text{query}})\} \right. \right. \\ & \left. \left. \mid v_{\Phi}^f) - f(x_{\text{query}}) \right)^2 \right]. \end{aligned}$$

Here, $p_{\text{CoT}}(T_v + 1)$ consists of T_v support examples and a new query:

$$\{x_1, g(x_1), f(x_1), \dots, x_{T_v}, g(x_{T_v}), f(x_{T_v}), x_{\text{query}}\}.$$

The LTVs, $v_{\Phi}^g = \{v_i^g\}_{i=1}^L$ and $v_{\Phi}^f = \{v_i^f\}_{i=1}^L$, are applied at corresponding layers and steps, modifying transformer outputs as $\tilde{M}_{\theta}(\cdot \mid v_{\Phi}^g)$ and $\tilde{M}_{\theta}(\cdot \mid v_{\Phi}^f)$.

In CoT-I, a single LTV is trained directly for the final step:

$$\min_{\Phi} \mathbb{E}_{f \sim \mathcal{D}_{\mathcal{F}}, x \sim \mathcal{D}_{\mathcal{X}}} \left[\left(\tilde{M}_{\theta}(p_{\text{CoT}}(T_v + 1) \mid v_{\Phi}^f) - f(x_{\text{query}}) \right)^2 \right].$$

This formulation allows CoT-I to bypass intermediate step modeling, focusing solely on final output prediction.

We construct the pre-training and LTV datasets following the same setting as in Appendix C.4.1, sampling N i.i.d. tasks from an isotropic Gaussian with $T_{\text{train}} = 101$ and $N = 100 \times 256 = 25,600$.

D.2.2 Results

Results are provided in Figure 8. For simplicity, we do not include the baselines (*i.e.*, FV and ICV). First, we observe that CoT-I/O (both with and without LTV) performs poorly. This is expected because the final prediction depends on the intermediate step prediction, *i.e.*, $\hat{s}_i^{(1)}$ in (5). If the intermediate step introduces noise, it propagates into the final prediction, leading to degraded performance. Consequently, a more direct approach, *i.e.*, CoT-I, achieves better and more effective prediction performance.

The latter case is particularly notable, as the model’s in-context solving abilities improve up to $T = T_{\text{train}}$, showing stable and lower error levels. However, CoT-I still struggles to maintain task fidelity when $T > T_{\text{train}}$ on our benchmark.

For LTV, we observe that whether CoT is used or not, LTV alone can sustain task behavior. This indicates that CoT, when combined with LTV, performs well primarily due to LTV, with minimal contribution to preserving task fidelity. This suggests that our method does not require additional ICL enhancements to maintain task fidelity.

D.3 Distributional Shift

We identify two scenarios from Garg et al. (2023) where the transformer model’s performance notably degrades during ICL inference: noisy linear regression and skewed covariance matrix.

Noisy linear regression Noise is added to the output of each example in the form of a standard Gaussian distribution. Specifically, the i -th output is defined as $w^\top x_i + \epsilon_i$, where $\epsilon_i \sim \mathcal{N}(0, 1)$. While the transformer and LTV are trained on standard linear regression, the data during the ICL inference phase is modified by this additive noise. In Figure 9, we observe that while the performance of FV degrades significantly, LTV is only slightly affected. Specifically, LTV requires training on more examples to maintain the same performance as in the noise-free setting. For instance, the performance of LTV at $T_v = 71$ is comparable to its performance in a noise-free environment at around $T_v = 56$.

Skewed covariance The inputs for the prompts are sampled from a zero-mean skewed Gaussian distribution: $x \sim \mathcal{N}(0, \tilde{\Sigma})$, where the eigenbasis of the skewed covariance matrix $\tilde{\Sigma}$ is chosen uniformly at random. Each i -th eigenvalue of this matrix is proportional to $1/i^2$. The results, shown

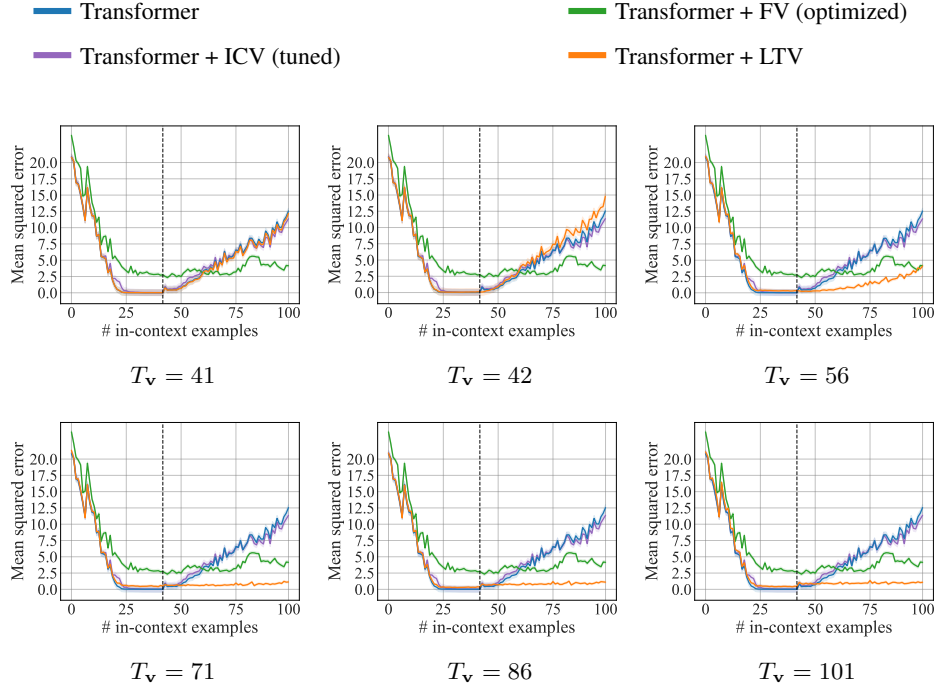


Figure 9: Evaluation on the class of linear functions on prompts with noisy labels, with the transformer pre-trained on up to $T_{\text{train}} = 41$ examples per prompt. Results are averaged over a batch of 256 randomly selected tasks. The shaded area represents the 95% confidence interval over the sampled prompts. T_v denotes the prompt length used in LTV training.

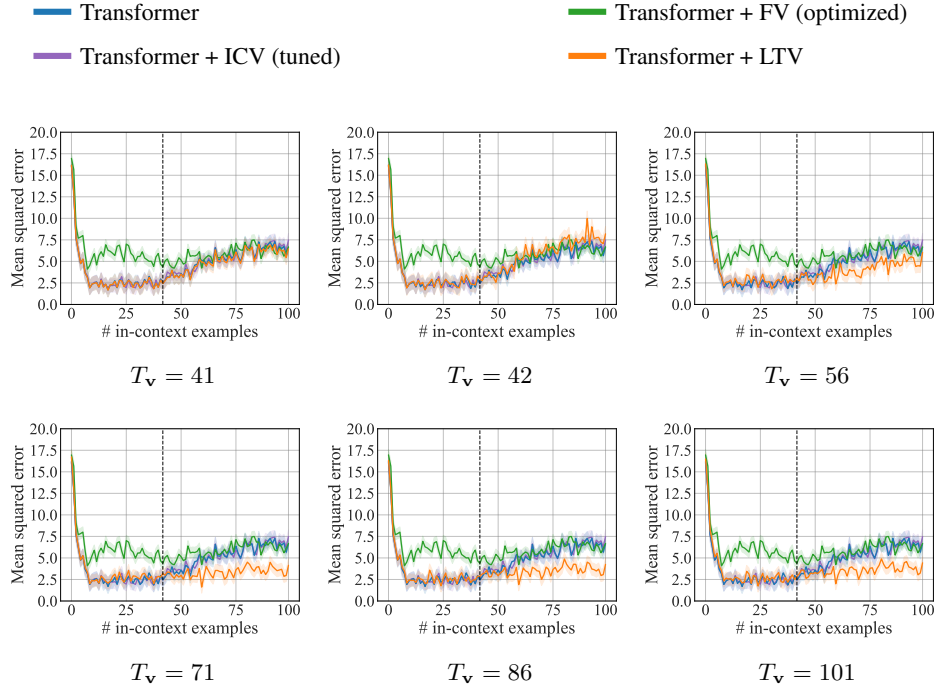


Figure 10: Evaluation on the class of linear functions under skewed covariance, with the transformer pre-trained on up to $T_{\text{train}} = 41$ examples per prompt. Results are averaged over a batch of 256 randomly selected tasks. The shaded area represents the 95% confidence interval over the sampled prompts. T_v denotes the prompt length used in LTV training.

in Figures 10, 11, and 12, align with the findings of Garg et al. (2023), with error curves being more unstable and oscillatory compared to the isotropic Gaussian case. While LTV exhibits some sensitivity to this instability, it maintains a low mean error and preserves performance to a certain extent.

D.4 Natural Language Processing

The complete set of accuracy scores on NLP benchmarks is reported in Table 2. Detailed descriptions of these benchmarks can be found in Appendix C.3.1.

D.5 Ablation Studies

The KL divergence scores are reported in Table 3. Additionally, the histograms in Figures 13, 14, and 15 illustrate the (unnormalized) probability density functions of the last hidden states, based on 25,600 collected samples.

The plots correspond well with the computed KL divergence scores. As the KL divergence values decrease, the histograms show greater alignment. Specifically, as the training prompt length for the LTV configurations increases, their density profiles become narrower, more closely resembling the shape of the vanilla transformer’s distribution at T_{train} . From a different perspective, this visual alignment supports our hypothesis once again: An optimized LTV with sufficiently long prompts performs near-optimally, as it effectively maintains the last hidden state distribution close to that of the model performing under $T = T_{\text{train}}$.

	Dataset	Prompt	Transformer	+ FV	+ ICV	+ SV	+ TV	+ LTV (ours)
Abstractive	AG News (Zhang et al., 2015)	Zero-shot	0.0 ± 0.0	30.1 ± 5.7	33.6 ± 5.6	31.7 ± 5.7	34.5 ± 5.0	46.9 ± 6.2
		Few-shot	67.2 ± 5.8	66.0 ± 5.8	64.7 ± 5.4	66.1 ± 4.5	57.9 ± 5.8	71.1 ± 5.6
	Antonym (Nguyen et al., 2017)	Zero-shot	2.3 ± 1.9	35.2 ± 5.9	46.4 ± 5.6	35.2 ± 6.1	29.6 ± 6.0	53.5 ± 6.2
		Few-shot	37.5 ± 6.0	64.5 ± 5.9	61.8 ± 5.7	64.8 ± 5.9	64.9 ± 5.2	65.6 ± 5.9
	CommonsenseQA (Talmor et al., 2019)	Zero-shot	10.2 ± 3.7	23.0 ± 5.2	28.9 ± 5.2	21.7 ± 1.1	22.0 ± 4.0	41.0 ± 6.1
		Few-shot	18.0 ± 4.7	20.3 ± 5.0	20.0 ± 4.6	20.3 ± 2.1	19.5 ± 3.4	21.1 ± 5.0
	English-French (Lample et al., 2018)	Zero-shot	0.0 ± 0.0	23.4 ± 5.2	40.2 ± 5.3	24.4 ± 2.7	25.4 ± 5.1	53.5 ± 6.2
		Few-shot	52.0 ± 6.2	71.9 ± 5.5	68.0 ± 5.0	74.4 ± 4.2	75.0 ± 5.3	75.4 ± 5.3
	English-Spanish (Lample et al., 2018)	Zero-shot	0.0 ± 0.0	9.0 ± 3.5	16.5 ± 4.3	9.7 ± 3.5	9.4 ± 2.9	20.7 ± 5.0
		Few-shot	36.7 ± 5.9	50.4 ± 6.2	49.4 ± 5.9	50.8 ± 4.2	50.6 ± 4.9	53.9 ± 6.1
	Landmark-Country (Hernandez et al., 2024)	Zero-shot	0.0 ± 0.0	60.9 ± 6.0	62.7 ± 5.1	62.9 ± 5.6	69.9 ± 6.0	71.1 ± 5.6
		Few-shot	66.8 ± 5.8	68.0 ± 5.8	68.4 ± 5.0	65.6 ± 5.3	66.3 ± 5.6	76.2 ± 5.3
	Person-Instrument (Hernandez et al., 2024)	Zero-shot	0.0 ± 0.0	11.3 ± 3.9	17.1 ± 4.9	13.1 ± 2.2	30.1 ± 3.5	33.2 ± 5.8
		Few-shot	52.7 ± 6.2	48.4 ± 6.2	53.2 ± 5.9	47.9 ± 6.1	43.6 ± 6.2	60.5 ± 6.0
	Person-Occupation (Hernandez et al., 2024)	Zero-shot	0.0 ± 0.0	4.3 ± 2.5	17.2 ± 3.9	9.0 ± 2.9	8.1 ± 5.8	36.3 ± 5.9
		Few-shot	30.9 ± 5.7	30.1 ± 5.7	37.6 ± 5.6	29.2 ± 5.2	28.8 ± 5.7	49.6 ± 6.2
	Person-Sport (Hernandez et al., 2024)	Zero-shot	0.0 ± 0.0	2.7 ± 2.0	12.9 ± 3.5	12.5 ± 3.1	12.6 ± 3.0	16.0 ± 4.5
		Few-shot	85.5 ± 4.3	86.3 ± 4.2	83.3 ± 3.5	88.3 ± 4.1	86.5 ± 4.0	89.5 ± 3.8
	Product-Company (Hernandez et al., 2024)	Zero-shot	0.0 ± 0.0	29.3 ± 5.6	43.0 ± 5.3	28.2 ± 4.9	28.7 ± 5.6	66.4 ± 5.8
		Few-shot	57.4 ± 6.1	62.5 ± 6.0	66.6 ± 5.8	68.0 ± 5.6	65.0 ± 4.7	74.2 ± 5.4
	Sentiment Analysis (Socher et al., 2013b)	Zero-shot	0.0 ± 0.0	0.0 ± 0.0	4.8 ± 1.9	0.0 ± 0.0	0.0 ± 0.0	10.9 ± 3.8
		Few-shot	74.6 ± 5.4	69.5 ± 5.7	72.5 ± 3.2	74.2 ± 2.5	77.1 ± 3.4	94.5 ± 2.8
	Synonym (Nguyen et al., 2017)	Zero-shot	1.2 ± 1.3	2.7 ± 2.0	14.2 ± 3.6	2.6 ± 4.6	2.5 ± 3.3	18.8 ± 4.8
		Few-shot	6.2 ± 3.0	10.5 ± 3.8	14.6 ± 4.5	10.5 ± 4.2	10.1 ± 3.0	40.6 ± 6.1
Extractive	NER-person (Tjong Kim Sang and De Meulder, 2003)	Zero-shot	5.5 ± 2.8	48.8 ± 6.2	49.9 ± 5.8	46.3 ± 1.3	47.8 ± 5.3	58.2 ± 6.1
		Few-shot	11.7 ± 4.0	56.2 ± 6.1	65.9 ± 5.2	53.7 ± 3.3	62.6 ± 3.9	79.3 ± 5.0
	NER-location (Tjong Kim Sang and De Meulder, 2003)	Zero-shot	6.6 ± 3.1	37.9 ± 6.0	26.4 ± 5.1	34.6 ± 3.9	27.2 ± 4.7	27.3 ± 5.5
		Few-shot	23.4 ± 5.2	43.4 ± 6.1	47.4 ± 5.9	43.5 ± 4.4	41.9 ± 5.6	57.4 ± 6.1
	NER-organization (Tjong Kim Sang and De Meulder, 2003)	Zero-shot	21.9 ± 5.1	48.8 ± 6.2	52.5 ± 6.0	46.1 ± 1.2	46.6 ± 4.3	64.1 ± 5.9
		Few-shot	16.8 ± 4.6	50.8 ± 6.2	54.1 ± 5.0	53.5 ± 2.6	51.1 ± 4.1	75.0 ± 5.3

Table 2: Complete set of accuracy scores (%) for zero-shot and few-shot (5-shot) predictions, averaged across 256 random seeds. \pm represents the margin of error at a 95% confidence level. The highest accuracy is marked in **boldface**, and the statistically best-performing method is **highlighted**.

Configuration at T_{\max}	Linear regression	Sparse linear regression	2-layer ReLU NN
Transformer	1.098	1.110	0.103
+ Function Vector	0.646	0.572	0.223
+ LTV ($T_{\mathbf{v}} = \{41, 41, 101\}$)	1.000	1.056	0.095
+ LTV ($T_{\mathbf{v}} = \{42, 42, 102\}$)	0.952	0.984	0.037
+ LTV ($T_{\mathbf{v}} = \{56, 56, 126\}$)	0.532	0.260	0.016
+ LTV ($T_{\mathbf{v}} = \{71, 71, 151\}$)	0.460	0.203	0.017
+ LTV ($T_{\mathbf{v}} = \{86, 86, 176\}$)	0.344	0.199	0.013
+ LTV ($T_{\mathbf{v}} = \{101, 101, 201\}$)	0.300	0.096	0.038

Table 3: KL divergence values are computed between the distributions of the last hidden states of the vanilla transformer at $T = T_{\text{train}}$ and the listed configurations at $T = T_{\max}$, where $T_{\max} = 101$ for linear functions and $T_{\max} = 201$ for neural networks. Kernel density estimation (KDE) is used to estimate the probability densities over a dataset of 25,600 samples. The lowest KL divergence score (*i.e.*, the most aligned configuration) is marked in **boldface**.

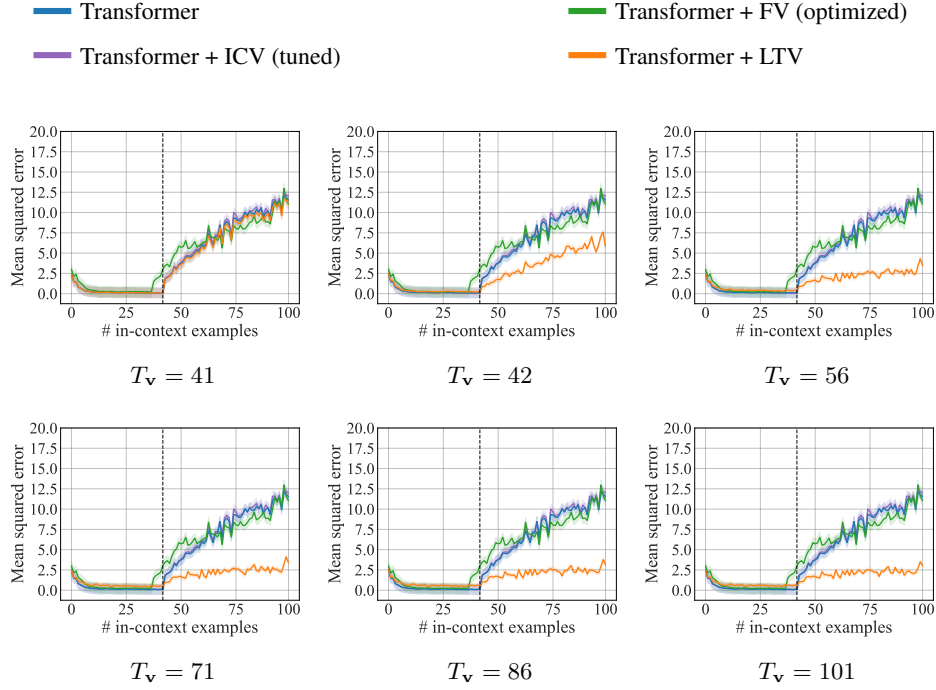


Figure 11: Evaluation on the class of sparse linear functions under skewed covariance, with the transformer pre-trained on up to $T_{\text{train}} = 41$ examples per prompt. Results are averaged over a batch of 256 randomly selected tasks. The shaded area represents the 95% confidence interval over the sampled prompts. T_v denotes the prompt length used in LTV training.

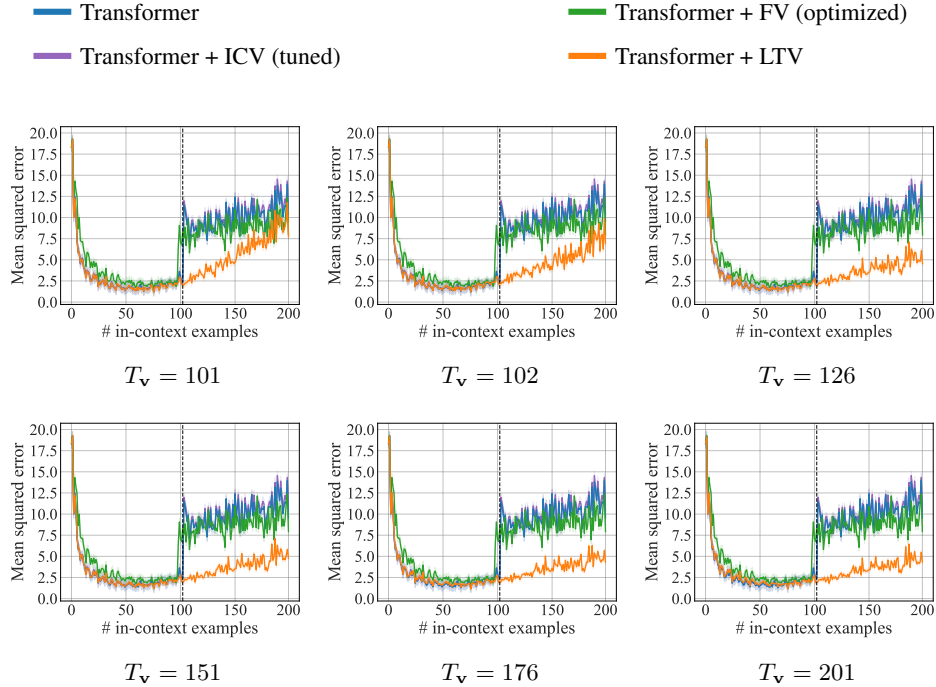


Figure 12: Evaluation on the class of 2-layer ReLU neural networks under skewed covariance, with the transformer pre-trained on up to $T_{\text{train}} = 101$ examples per prompt. Results are averaged over a batch of 256 randomly selected tasks. The shaded area represents the 95% confidence interval over the sampled prompts. T_v denotes the prompt length used in LTV training.

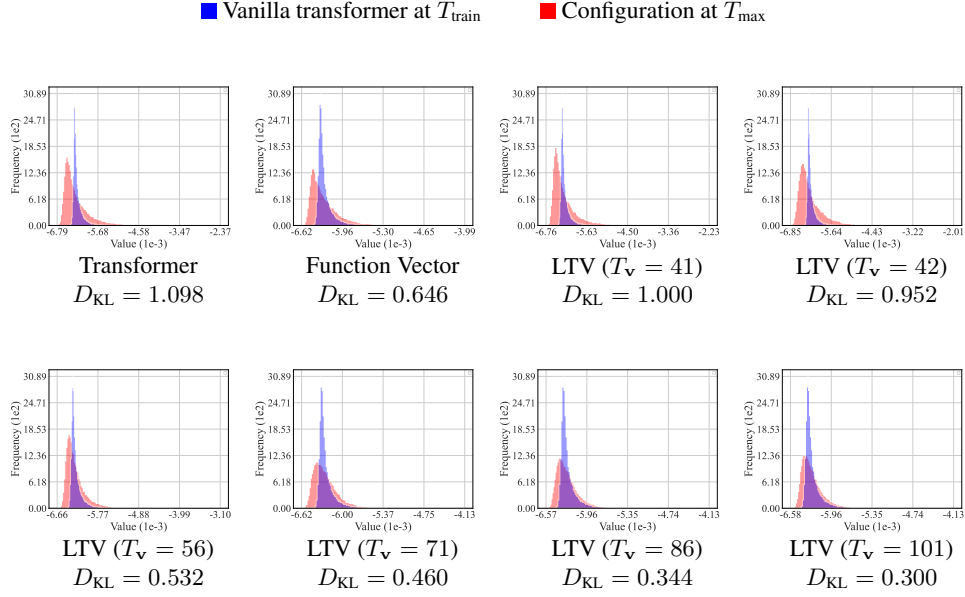


Figure 13: Histograms of the empirical distribution of the last hidden states for the vanilla transformer collected at T_{train} and the tested configuration at the maximum length $T_{\text{max}} = 101$ under linear functions. These histograms are generated using a dataset of 25,600 samples and correspond to the KL divergence scores reported in Table 3.

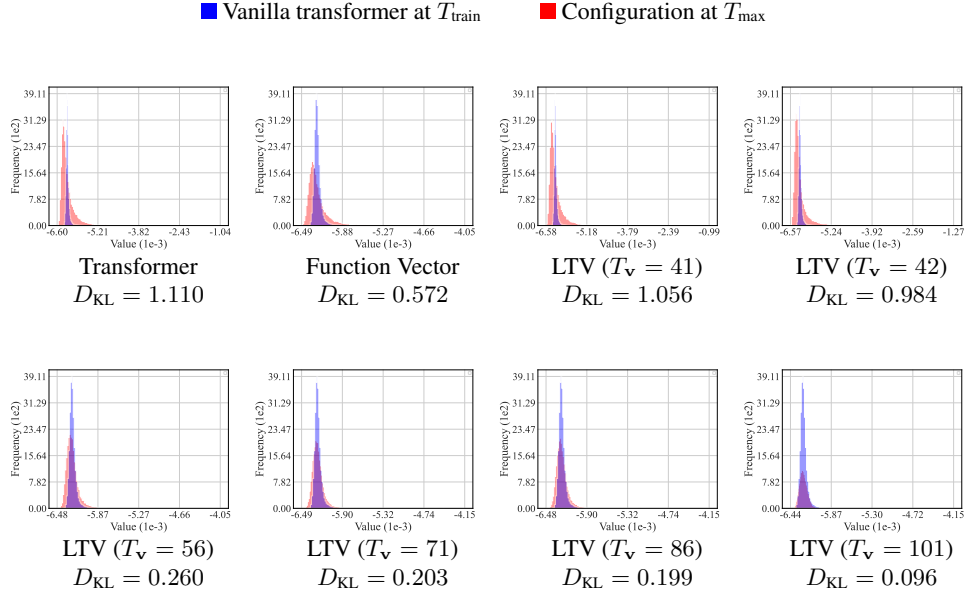


Figure 14: Histograms of the empirical distribution of the last hidden states for the vanilla transformer collected at T_{train} and the tested configuration at the maximum length $T_{\text{max}} = 101$ under sparse linear functions. These histograms are generated using a dataset of 25,600 samples and correspond to the KL divergence scores reported in Table 3.

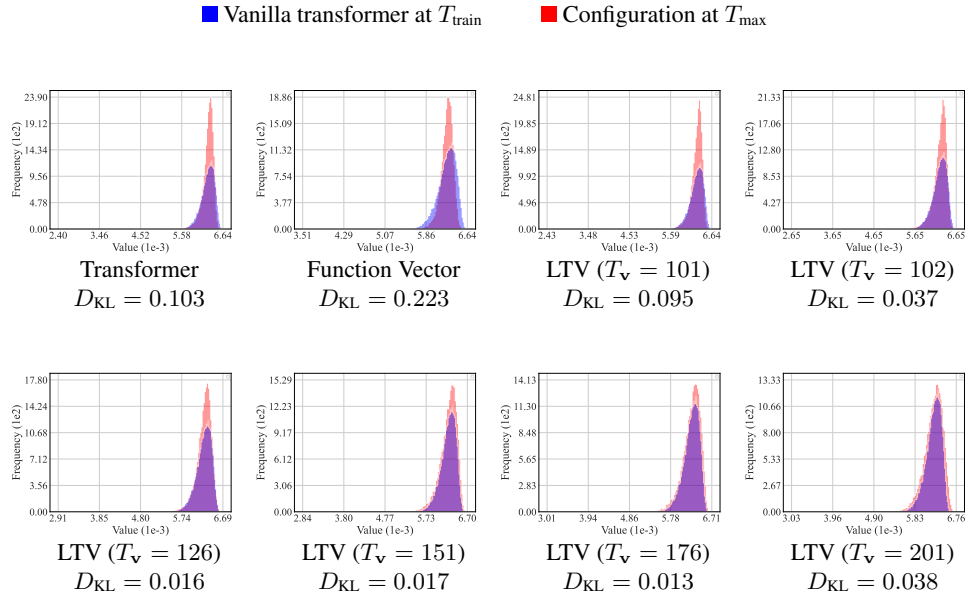


Figure 15: Histograms of the empirical distribution of the last hidden states for the vanilla transformer collected at T_{train} and the tested configuration at the maximum length $T_{\text{max}} = 201$ under 2-layer ReLU neural networks. These histograms are generated using a dataset of 25,600 samples and correspond to the KL divergence scores reported in Table 3.