ProxySPEX: Inference-Efficient Interpretability via Sparse Feature Interactions in LLMs

Landon Butler*

Department of EECS UC Berkeley landonb@berkeley.edu

Abhineet Agarwal*

Department of Statistics UC Berkeley aa3797@berkeley.edu

Justin Singh Kang*

Department of EECS UC Berkeley justin_kang@berkeley.edu

Yigit Efe Erginbas

Department of EECS UC Berkeley erginbas@berkeley.edu

Bin Yu

Departments of Statistics and EECS
UC Berkeley
binyu@berkeley.edu

Kannan Ramchandran

Department of EECS UC Berkeley kannanr@berkeley.edu

Abstract

Large Language Models (LLMs) have achieved remarkable performance by capturing complex interactions between input features. To identify these interactions, most existing approaches require enumerating all possible combinations of features up to a given order, causing them to scale poorly with the number of inputs n. Recently, Kang et al. (2025) proposed SPEX, an information-theoretic approach that uses interaction sparsity to scale to $n \approx 10^3$ features. SPEX greatly improves upon prior methods but requires tens of thousands of model inferences, which can be prohibitive for large models. In this paper, we observe that LLM feature interactions are often hierarchical—higher-order interactions are accompanied by their lowerorder subsets—which enables more efficient discovery. To exploit this hierarchy, we propose PROXYSPEX, an interaction attribution algorithm that first fits gradient boosted trees to masked LLM outputs and then extracts the important interactions. Experiments across four challenging high-dimensional datasets show that PROXYSPEX more faithfully reconstructs LLM outputs by 20% over marginal attribution approaches while using $10 \times$ fewer inferences than SPEX. By accounting for interactions, PROXYSPEX efficiently identifies the most influential features, providing a scalable approximation of their Shapley values. Further, we apply PROXYSPEX to two interpretability tasks. Data attribution, where we identify interactions among CIFAR-10 training samples that influence test predictions, and mechanistic interpretability, where we uncover interactions between attention heads, both within and across layers, on a question-answering task. The PROXYSPEX algorithm is available at https://github.com/mmschlk/shapiq.

1 Introduction

Large language models (LLMs) have achieved great success in natural language processing by capturing complex interactions among input features. Modeling interactions is not only crucial for language, but also in domains such as computational biology, drug discovery and healthcare, which require reasoning over high-dimensional data. In high-stakes contexts, responsible decision-making based on model outputs requires interpretability. For example, in healthcare, a physician relying on LLM diagnostic assistance must intelligibly be able to explain their decision to a patient.

Post-hoc feature explanation methods such as SHAP [1] and LIME [2] focus on marginal attributions and do not explicitly capture the effect of interactions. To address this limitation, recent work has

^{*}Equal contribution. Order determined by coin flip.

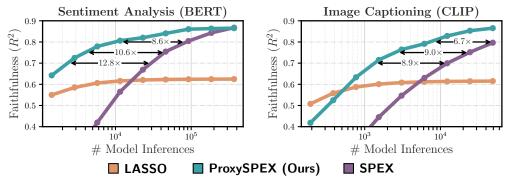


Figure 1: PROXYSPEX requires $\sim 10 \times$ fewer inferences to achieve equally faithful explanations as SPEX for a sentiment classification and image-captioning task using a BERT and CLIP model respectively. LASSO faithfulness plateaus indicating limits of marginal approaches.

proposed interaction indices, such as Faith-Shap [3], that attribute all interactions up to a given order d by exhaustively enumerating them. With n features, enumerating $O(n^d)$ interactions quickly becomes infeasible for even small n and d. Kang et al. [4] recently introduced SPEX, the first interaction attribution method capable of scaling up to n=1000 features. SPEX scales with n by observing that LLM outputs are driven by a small number of interactions. It exploits this sparsity by utilizing a sparse Fourier transform to efficiently search for influential interactions without enumeration. For example, with n=100 features, SPEX requires approximately 2×10^4 model inferences to learn order 5 interactions—a small fraction of all possible 100^5 interactions. Nonetheless, 2×10^4 inferences is prohibitively expensive for large models. Hence, the question naturally arises: Can we identify additional structural properties among interactions to improve inference-efficiency?

We show empirically that local (i.e., input specific) LLM feature interactions are often *hierarchical*: for an order d interaction, an LLM includes lower-order interactions involving subsets of those d features (see Figure 2). We use this to develop PROXYSPEX, an interaction attribution algorithm that reduces the number of inferences compared to SPEX by $10\times$ while achieving equally faithful explanations. PROXYSPEX exploits this local hierarchical structure by first fitting gradient boosted trees (GBTs) as a proxy model to predict the output of LLMs on masked input sequences. Then, PROXYSPEX extracts important interactions from the fitted GBTs [5].

Evaluation overview. We compare PROXYSPEX to marginal feature attributions and SPEX across four high-dimensional datasets with hundreds of features. Results are summarized below:

- 1. Faithfulness. PROXYSPEX learns more faithful representations of LLM outputs than marginal approaches ($\approx 15\%$ to 25%) on average across datasets as we vary the number of inferences. Figure 1 compares explanation faithfulness of PROXYSPEX to marginal attributions and SPEX.
- **2. Feature identification.** By accounting for interactions, PROXYSPEX identifies influential features that impact model outputs more significantly than marginal approaches, and can approximate Shapley values better than KernalSHAP in the low-inference regime.
- **3. Case study 1: Data attribution.** Data Attribution is the problem of identifying training points responsible for a given test prediction. On CIFAR-10 [6] PROXYSPEX identifies the interactions between training samples that most significantly impact classification performance.
- **4.** Case study 2: Model component attribution. We use PROXYSPEX to study interactions between attention heads, both within and across layers, on MMLU [7] for Llama-3.1-8B-Instruct [8]. We observe that intra-layer interactions become more significant for deeper layers. PROXYSPEX identifies interactions that allow it to prune more heads than the LASSO.

2 Related work and applications

Feature and interaction attribution. SHAP [1] and LIME [2] are widely used for model-agnostic feature attribution. SHAP uses the game-theoretic concept of Shapley values [9] for feature attribution, while LIME fits a sparse linear model [10]. Cohen-Wang et al. [11] also consider fitting a sparse linear model for feature attribution. Chen et al. [12] uses an information-theoretic approach for feature attributions. Other methods [13, 14] study model structure to derive feature attributions.

Sundararajan et al. [15] and Bordt and von Luxburg [16] define extensions to Shapley values that consider interactions. Fumagalli et al. [17] provides a framework for computing several interaction attribution scores, but their approach does not scale past $n \approx 20$ features, which prevents them from being applied to modern ML problems that often consist of hundreds of features. Note that some feature attribution approaches such as LIME and Faith SHAP [3] are formulated explicitly as a function approximation, while others are defined axiomatically such as SHAP, though one can typically construct equivalent function approximation objectives with a suitable distance metric.

Fourier transforms and deep learning explainability. Several works theoretically study the spectral properties of transformers. Ren et al. [18] show transformers have sparse spectra and Hahn and Rofin [19], Abbe et al. [20] establish that they are low degree. Abbe et al. [21, 22] study the bias of networks learning interactions via a "staircase" property, i.e., using lower-order terms to learn high-order interactions. Sparsity and low degree structure is also empirically studied in [23, 24]. Kang et al. [25] shows that under sparsity in the Möbius basis [26], a representation closely related to Shapley values and the Fourier transform, interaction attributions can be computed efficiently. Mohammadi et al. [27] also learn a sparse Möbius representation for computing Shapley values. Kang et al. [4] use these insights to propose SPEX, the first robust interaction attribution algorithm to scale to the order of $n \approx 1000$ features. Gorji et al. [5] apply sparse Fourier transforms [28–31] for computing Shapley values. They also provide an algorithm to extract the Fourier transform of tree-based models using a single forward pass.

SPEX. We refer to the algorithm proposed in this manuscript as PROXYSPEX, in reference to SPEX, since both works exploit a sparse interaction prior to reduce computational and sample budget. SPEX uses an algebraic structured sampling scheme, coupled with error correction decoding procedures to efficiently compute the interactions in the form of a Fourier transform. In contrast, PROXYSPEX uses random samples to learn a proxy model that implicitly exploits the sparse interaction priors and our newly proposed hierarchical prior.

Mechanistic Interpretability (**MI**). MI seeks to uncover the underlying mechanisms of neural networks and transformers [32] in order to move past treating these models as *black boxes*. PROXYSPEX answers the question "what combinations of inputs matter?" which is a vital precursor and complement to MI investigations that subsequently address "how does the model compute based on those specific inputs?" Some closely related MI work attempts to recover circuits to explain underlying model behavior [33, 34]. Hsu et al. [35] use MI for interaction attribution. See Sharkey et al. [36] for a review of open problems and recent progress in MI.

3 PROXYSPEX

In this section, we first empirically justify our premise that significant interactions affecting LLM output are hierarchical—influential high-order interactions imply important lower-order ones. Next, we introduce PROXYSPEX, which aims to identify feature interactions for a given input x while minimizing the number of expensive calls to an LLM.

3.1 Preliminaries

Value function. Let \mathbf{x} be the input to the LLM consisting of n features². For $S \subseteq [n]$, where $[n] = 1, \ldots, n$, denote \mathbf{x}_S as the *masked* input where we retain features indexed in S and replace all others with the [MASK] token. For example, in the sentence \mathbf{x} ="The sequel truly elevated the original", if $S = \{1, 2, 5, 6\}$, \mathbf{x}_S = "The sequel [MASK] [MASK] the original". Masks can be more generally applied to any type of input such as image patches in a vision-language model. For a masked input \mathbf{x}_S and LLM f, let $f(\mathbf{x}_S) \in \mathbb{R}$ denote the output of the LLM under masking pattern S. The value function f is problem dependent. For classification tasks, a common choice is the logit of the predicted class for unmasked input, $f(\mathbf{x})$. In generative tasks, $f(\mathbf{x}_S)$ can represent the perplexity of generating the *original output* for the unmasked input. Since we focus on providing input-specific explanations, we suppress notation on \mathbf{x} and denote $f(\mathbf{x}_S)$ as f(S).

Fourier transform of value function. Let $2^{[n]}$ be the powerset of the index set. The value function f can be equivalently thought of as a set function from $f: 2^{[n]} \to \mathbb{R}$. Every such function admits a

²Features refer to inputs at a given granularity, e.g., tokens in an LLM or image patches in a vision model.

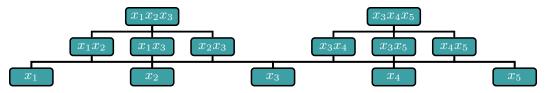


Figure 2: We observe that LLM feature interactions are often hierarchical—higher-order interactions are accompanied by their lower-order subsets.

Fourier transform $F: 2^{[n]} \to \mathbb{R}$ of f, related as follows:

Transform:
$$F(T) = \frac{1}{2^n} \sum_{S \subseteq [n]} (-1)^{|S \cap T|} f(S)$$
, Inverse: $f(S) = \sum_{T \subseteq [n]} (-1)^{|T \cap S|} F(T)$. (1)

The parameters F(T) are known as Fourier coefficients and capture the importance of an interaction of features in a subset T. Equation (1) represents an *orthonormal* transform onto a parity (XOR) basis [37]. For the rest of the paper, we use the terms Fourier coefficient and interaction interchangeably. Further, we refer to the set of Fourier coefficients $\{(T, F(T)) : T \subseteq [n]\}$ as the *spectrum*.

Interpretable approximation of value function. We aim to learn an interpretable approximate function \hat{f} that satisfies the following:

1. **Faithful representation.** To characterize how well the surrogate function \hat{f} approximates the true function, we define *faithfulness* [38]:

$$R^{2} = 1 - \frac{\|\hat{f} - f\|^{2}}{\|f - \bar{f}\|^{2}}, \quad \text{where } \|f\|^{2} = \sum_{S \subseteq [n]} f(S)^{2}, \, \bar{f} = \frac{1}{2^{n}} \sum_{S \subseteq [n]} f(S). \tag{2}$$

Faithfulness measures how well \hat{f} predicts model output. High faithfulness implies accurate approximation of F(T) (this follows from orthonormality of (1)).

- 2. **Sparse representation.** \hat{f} should be *succinct*. Previous works [4, 25, 39–41] have shown that a sparse and low-degree \hat{f} can achieve high R^2 . That is, $F(T) \approx 0$ for most T (*sparsity*), and |F(T)| is only large when $|T| \ll n$ (*low degree*).
- 3. **Efficient computation.** Without any additional assumptions on the spectrum, learning f is exponentially hard since there are 2^n possible subsets T. PROXYSPEX relies on the sparse, low degree Fourier transform along with the hierarchy property to reduce LLM inferences.

A faithful and sparse f allows straightforward computation of all popular feature or interaction attribution scores defined in the literature, e.g., Shapley, Banzhaf, Influence Scores, Faith-Shapley. Closed-form formulas for converting F to various attribution indices are provided in Appendix A.1.

3.2 Empirical evidence of spectral hierarchies

To quantify the degree of hierarchical structure in LLMs, we introduce the following definition called Direct Subset Rate (DSR),³ defined for any value function f and integer k.

$$DSR(f,k) = \frac{1}{k} \sum_{S \in \mathcal{F}_k} \frac{1}{|S|} \sum_{i \in S} \mathbb{1} \left\{ S \setminus \{i\} \in \mathcal{F}_k \right\}, \quad \text{where } \mathcal{F}_k \text{ denotes the } k \text{ largest}$$
 Fourier coefficients of f . (3)

For the top k coefficients (i.e., interactions), DSR measures the average fraction of Fourier coefficients that exclude only *one* of the features $F(S \setminus \{i\})$. For example, an f with $\mathcal{F}_4 = \{\emptyset, \{1\}, \{2\}, \{1, 3\}\}$ would have DSR of $\frac{1}{4}\left(1+1+1+\frac{1}{2}\right)=\frac{7}{8}$. High DSR implies that significant high-order interactions have corresponding significant lower-order Fourier coefficients. Figure 2 visualizes hierarchical interactions. Next, we show that two LLM based value functions have high DSR.

We take 20 samples from a sentiment analysis task and an image captioning task [42]; see Section 4 for a detailed description and our choice of value function. We generate masks S and apply SPEX until our learned value function has faithfulness (R^2) more than 0.9. Figure 3 visualizes the DSR for

 $^{{}^{3}}$ For $S = \emptyset$, we set $\frac{0}{0} = 1$.

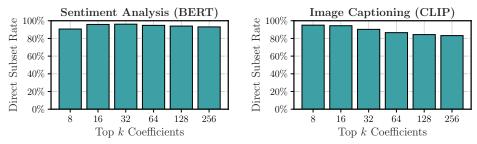


Figure 3: The top-k interactions in both a sentiment analysis and image captioning task have high DSR indicating strong hierarchical structure.

various values of k, i.e., number of top interactions. DSR is consistently larger than 80%, indicating strong hierarchical structure. In Appendix B.2, we consider two additional metrics measuring hierarchical structure, and demonstrate that the top-k interactions are faithful.

Using GBTs to capture hierarchical Interactions. Tan et al. [43] proved that decision trees learn "staircase" functions, e.g., $f = x_1 + x_1x_2 + x_1x_2x_3$, effectively due to their greedy construction procedure. We empirically confirm this by comparing the performance of various proxy models on a synthetic hierarchical function (i.e., sum of staircase functions resembling Figure 2) as well as the Sentiment dataset in Appendix Figure 13. Appendix B.4 details the simulation set-up. GBTs vastly outperform other proxy models, indicating their natural ability to identify hierarchical interactions with limited training data. Interestingly, GBTs outperform random forests as well. This is because random forests are ineffective at learning hierarchical functions [44], i.e., sums of staircases, while GBT-like algorithms disentangle sums effectively [45].

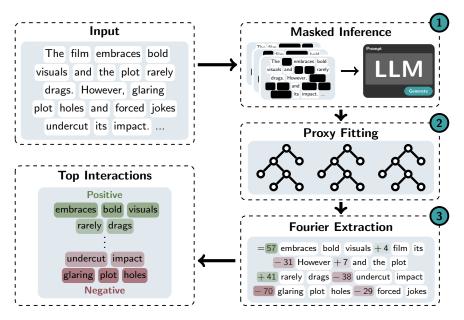


Figure 4: (1) PROXYSPEX masks subsets of words and queries the LLM using this masked input. (2) It then fits GBTs as a proxy model to learn the LLM's hierarchical interactions. (3) An interpretable sparse representation is extracted from the fitted GBT which captures the influential interactions.

3.3 PROXYSPEX via Gradient Boosted Trees to fit hierarchies

The PROXYSPEX algorithm (see Figure 4):

Step 1 - Sampling and querying. Given LLM f and input instance $\mathbf x$ to explain, generate a dataset $\mathcal D = (S_i, f(S_i))_{i=1}^\ell$ for training the proxy. The inputs S_i represent the masks of $\mathbf x$. Each mask S_i is sampled uniformly from the set [n]. The labels $f(S_i)$ are obtained by querying the LLM.

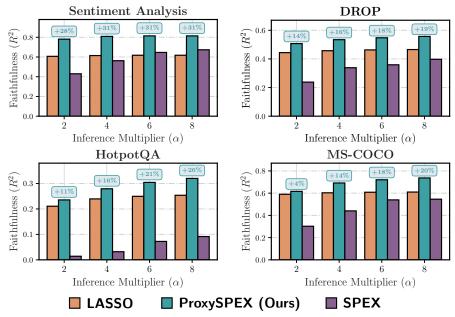


Figure 6: Comparison of faithfulness of different attribution methods with $\alpha \cdot n \log_2(n)$ training masks for different inference multipliers $\alpha \in \{2,4,6,8\}$. While SPEX is only competitive with LASSO for large α , the gap between PROXYSPEX and LASSO increases with α .

Step 2 - Proxy Training. Fit GBTs to \mathcal{D} with 5-fold cross-validation (CV).

Step 3 - Fourier extraction. We use Gorji et al. [5] to extract the Fourier representation of the fitted GBTs in a single forward pass; see Appendix A.2. With T trees of depth d there are at most $O(T4^d)$ non-zero Fourier coefficients [5]. To improve interpretability, we sparsify the extracted representation by keeping only the top k Fourier coefficients. Fig. 5 shows that only ≈ 200 Fourier coefficients are needed to achieve equivalent faithfulness for a sentiment classification and image captioning (MS-COCO) dataset. Additional results regarding the sparsity of Fourier spectra learned by GBTs are in Appendix B.3.

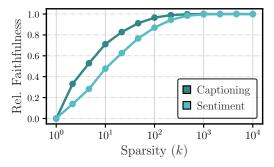


Figure 5: Relative faithfulness as a function of Fourier sparsity. Only ≈ 200 coefficients are required to achieve equivalent faithfulness. Sparsity for sentiment is higher since inputs have larger n.

Step 4 (Optional): Coefficient refinement via

regression. As a final step, we optionally regress the extracted, top k Fourier coefficients on the collected data \mathcal{D} to improve the estimation. Empirically we observe this step is can sometimes marginally improve performance, but seldom negatively impacts performance. This step is included if it leads to lower CV error.

4 Results

Datasets and models

- 1. Sentiment is a classification task composed of the Large Movie Review Dataset [46] which consists of positive and negative IMDb movie reviews. We use words as input features and restrict to samples with $n \in [256, 512]$. We use the encoder-only fine-tuned Distilbert model [47, 48], and the logit of the positive class as the value function.
- 2. HotpotQA [49] is a generative question-answering task over Wikipedia articles. Sentences are input features, and we restrict to samples with $n \in [64, 128]$. We use Llama-3.1-8B-Instruct, and perplexity of the unmasked output as the value function.

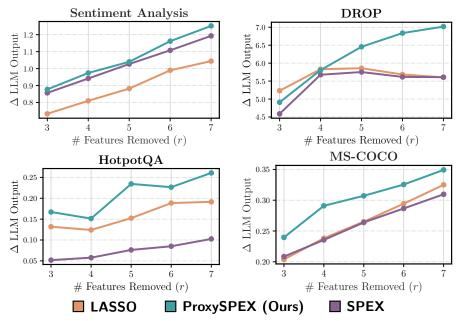


Figure 7: By accounting for interactions, PROXYSPEX identifies more influential features across datasets than the LASSO. Apart from the sentiment analysis task (top left), SPEX does not collect enough training masks to out-perform LASSO.

- 3. Discrete Reasoning Over Paragraphs (DROP) [50] is a paragraph level question-answering task. We use words as input features and restrict to samples with $n \in [256, 512]$. We use Llama-3-8B-Instruct and the perplexity of the unmasked output as the value function.
- 4. MS-COCO [42] contains images and corresponding text captions. Image patches and words are the input features with n ∈ [60, 85]. We use CLIP-ViT-B/32, a joint vision-language encoder, with the value function defined as the contrastive loss over all datapoints.

Baselines and hyperparameters. For marginal feature attributions, we use the LASSO. We use the same datasets at [4] and add MS-COCO for an additional modality. It was shown in [4] that popular marginal metrics such as SHAP are significantly less faithful than the LASSO, e.g., have $R^2 < 0$. We use the LASSO implementation from scikit-learn, and choose the l_1 regularization parameter via 5-fold CV. For interaction indices, we compare PROXYSPEX to SPEX. Due to the scale of n in our experiments, we cannot compare methods for computing interaction indices such as Faith-Shapley, Faith-Banzhaf, and Shapley-Taylor using SHAP-IQ [17], and SVARM-IQ [51], because they enumerate all possible interactions, making them computationally infeasible. For PROXYSPEX, a list of GBT hyper-parameters we tune over are in Appendix B.

4.1 Faithfulness

We compare attribution method faithfulness by varying the number of training masks. For each sample with n features, we generate $\alpha \cdot n \log_2(n)$ masks, varying $\alpha \in \{2,4,6,8\}$, to normalize difficulty across inputs of varying lengths (some by over 100 tokens). This $n \log(n)$ type scaling is heuristically guided by compressed sensing bounds [52]. These suggest the number of samples required grows with sparsity (assumed $\propto n$) and logarithmically with problem dimensionality (if dimensionality for degree-d interactions is $\approx n^d$, this yields a $\log(n^d) = d \log(n)$ factor). Together, these factors support an $n \log(n)$ scaling. While not directly applicable, these bounds offer a useful heuristic for how sampling complexity scales with n.

Figure 6 shows average faithfulness over 1,000 test masks per sample. PROXYSPEX outperforms LASSO with limited inferences and continues to improve where LASSO plateaus, indicating that it is learning influential interactions. While SPEX is often faster for the same number of masks, SPEX needs additional inference time to match R^2 , making PROXYSPEX faster overall. For the smaller DistilBERT model under the sentiment analysis task, the wall clock speedup is $\sim 3\times$, while with the bigger CLIP-ViT-B/32 model with MS-COCO we see $\sim 5\times$ speedup (See Appendix B.6).

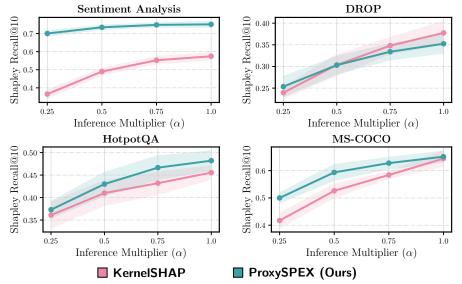


Figure 8: Recall of the top ten Shapley values after $\alpha \cdot n \log_2(n)$ inferences for multipliers $\alpha \in \{0.25, 0.5, 0.75, 1.0\}$. For small α , PROXYSPEX is superior at recovering the most significant features, while KernalSHAP outperforms as α increases. Error bands indicate the standard deviation across ten different runs of the algorithms.

4.2 Feature Identification

We measure the ability of methods to identify the top r influential features that influence LLM output:

$$\Delta \text{ LLM Output } (r) = \frac{|f([n]) - f(S^*)|}{|f([n])|}, \quad S^* = \underset{|S| = n - r}{\operatorname{argmax}} |\hat{f}([n]) - \hat{f}(S)|. \tag{4}$$

Solving Eq. 4 for an arbitrary \hat{f} presents a challenging combinatorial optimization problem. However, PROXYSPEX and SPEX represent \hat{f} as a sparse Fourier transform. This representation facilitates solving the optimization as a tractable linear integer program. The sparsity of the extracted Fourier representation ensures that the time required to solve this program is negligible compared to sampling the LLM and fitting the GBTs. Full details of the construction of this program are given in Appendix A.3. Under LASSO, Eq. 4 is easily solved through selecting features by the size of their coefficients. We measure the removal ability of different attribution methods when we collect $8n\log_2(n)$ training masks and plot the result in Figure 7. By accounting for interactions, PROXYSPEX identifies significantly more influential features than the LASSO. Apart from the sentiment analysis task, SPEX does not collect enough training masks to outperform the LASSO.

4.3 Shapley Value Approximation

We evaluate the performance of PROXYSPEX for efficiently approximating Shapley values. Across all tasks, we first run KernelSHAP with 10,000 test masks and treat these approximated Shapley values as ground truth. We measure the recall of the top ten highest-magnitude Shapley values for KernelSHAP and PROXYSPEX under $\alpha \cdot n \log_2(n)$ inferences with multipliers $\alpha \in \{0.25, 0.5, 0.75, 1.0\}$. For this inference budget, competing algorithms such as LeverageSHAP [53] and SVARM [54] struggle to provide accurate approximations. We find PROXYSPEX initially provides a better coarse approximation than KernelSHAP (Figure 14). However, since PROXYSPEX is optimized for faithfulness and does not rely on the Shapley kernel, it is eventually surpassed by KernelSHAP with enough inferences. Additional results under mean squared error are included in Appendix B.5.

5 Case studies

We now present two case studies of PROXYSPEX for two different interpretability problems: *data attribution* [55] and *model component attribution* [56], a key problem in mechanistic interpretability. We first show how both of these tasks can be reformulated as feature attribution tasks; recent work has highlighted the connections between feature, data, and model component attribution [57].

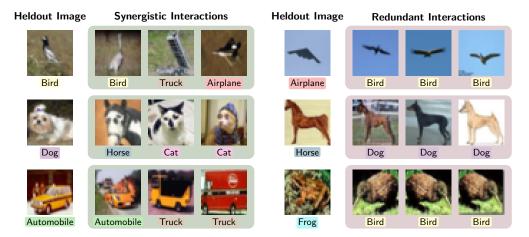


Figure 9: Synergistic interactions: data that together are more valuable together than the sum of their parts and aid in classification. Redundant interactions: Data that may contain similar information, their combined influence is less than the sum of the parts.

5.1 Data Attribution via Non-Linear Datamodels

Data attribution for classification is the problem of understanding how fitting a model g_{θ} on a subset S of training samples affects the prediction of a test point \mathbf{z} of class c. This problem can be converted into our framework by defining an appropriate value function f,

$$f(S) \triangleq (\text{logit for } c \text{ on } \mathbf{z}) - (\text{highest incorrect logit on } \mathbf{z}), \text{ when } g_{\theta} \text{ is trained on } S.$$
 (5)

The value function f quantifies the impact of a subset S on the classification of z. Sampling f is very expensive since it involves training a new model g_{θ} for every subset S. As a result, most data attribution approaches do not consider the impact of interactions. Notably, Ilyas et al. [55] use LASSO to learn f when training a ResNet model on the CIFAR-10 dataset [6]. As a case study, we apply PROXYSPEX to understand the impact of interactions between CIFAR-10 training samples.

Defining data interactions. Interactions between samples can be either *redundant interactions* or *synergistic interactions*. Redundant interactions are when the influence of a subset S is not additive. Redundancy typically occurs between highly correlated samples, e.g., semantic duplicates [58]. Synergistic interactions occur when a subset S influences a prediction by shaping a decision boundary that no individual sample in S could do so by itself. That is, the model needs the combined effect of training samples in S to correctly classify \mathbf{z} .

Results. We visualize interactions learned by PROXYSPEX in Figure 9 for randomly selected CIFAR-10 test points. Experimental details are in Appendix C.1. PROXYSPEX identifies highly similar training samples (redundancies) as well as synergistic interactions between samples of different classes. See Appendix C.1 for examples of other randomly selected test samples.

5.2 Model Component Attribution

We study the role of attention heads for a question-answering task using Llama-3.1-8B-Instruct and MMLU (high-school-us-history), which is a multiple-choice dataset. We treat each attention head as a feature and aim to identify interactions among heads using PROXYSPEX. Let L represent the number of layers in an LLM and let $\mathcal{L} \subseteq [L]$ represent a subset of the layers. Let $\mathcal{H}_{\mathcal{L}}$ denote the set of attention heads within these layers. For a subset of heads $S \subseteq \mathcal{H}_{\mathcal{L}}$, we set the output of heads in $\mathcal{H}_{\mathcal{L}} \setminus S$ to 0 and denote the ablated LLM as $\mathrm{LLM}_S(\cdot)$. Define f as:

$$f_{\mathcal{L}}(S) \triangleq \text{Accuracy of LLM}_S \text{ on training set of MMLU.}$$
 (6)

Pruning results. We use the LASSO and PROXYSPEX to identify the most important heads for various sparsity levels (i.e., the number of retained heads) across different sets of layers. We also compare to a Best-of-N baseline, where we take the best of N=5000 different randomly chosen S, further details are in Appendix C.2. We use the procedure detailed in Section 4.2 to identify heads

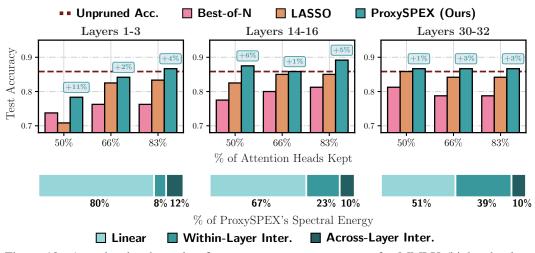


Figure 10: Attention head pruning for Llama-3.1-8B-Instruct for MMLU (high-school-ushistory). **Top:** We report the test accuracy vs. percentage of heads retained, comparing PROXYSPEX, LASSO, and Best-of-N across layer groups (1-3, 14-16, 30-32). Unpruned accuracy shown by dashed line. **Bottom:** PROXYSPEX's learned spectral energy distribution into linear effects, within-layer, and across-layer interactions per layer group.

to remove for both PROXYSPEX and LASSO. Test accuracies for each method are presented in Figure 10 at three different sparsity levels, and with three different layer ranges: initial (1-3), middle (14-16) and final (30-32). We observe that PROXYSPEX consistently outperforms both baselines, with a higher test accuracy on the pruned models identified using PROXYSPEX.

Characterizing interactions between attention heads. Analyzing the Fourier spectrum learned by PROXYSPEX offers insights into the nature of the internal mechanisms of the LLM. As shown in Figure 10 (bottom), the spectral energy attributed to interactions, particularly *within-layer* interactions, markedly increases in deeper layers of Llama-3.1-8B-Instruct. There are many works that look at the differing functional roles of attention heads across layers [59]. PROXYSPEX provides an exciting new quantitative approach to further investigate these phenomena.

6 Discussion

Conclusion. We introduce PROXYSPEX, an inference-efficient interaction attribution algorithm that efficiently scales with n by leveraging an observed hierarchical structure among significant interactions in the Fourier spectrum of the model. Experiments across 4 high-dimensional datasets show that PROXYSPEX exploits hierarchical interactions via a GBT proxy model to reduce inferences by $\sim 10 \times$ over SPEX [4] while achieving equally faithful explanations. We demonstrate the importance of efficient interaction discovery by applying PROXYSPEX to data and model component attribution.

Limitations. GBTs effectively capture hierarchical interactions but may not perform as well when interactions have a different structure. For example, simulations in Appendix B.4 empirically confirm that GBTs suffer in the case of sparse but non-hierarchical functions. More generally, in cases where the proxy GBT model is not faithful, the interactions identified by PROXYSPEX might not be representative of the model's reasoning. Another limitation is the degree of human interpretability that can be understood from computed interactions. While interactions can offer richer insights, they are more difficult to parse than marginal alternatives. Further improvements in visualization and post-processing of interactions are needed to fully harness the advances of PROXYSPEX.

Future work. Inference-efficiency could be further improved by exploring alternative proxy models, additional Fourier spectral structures, or adaptive masking pattern designs. Integrating PROXYSPEX with internal model details, such as via hybrid approaches with MI or by studying its connection to sparsity in transformer attention [60], offers another promising avenue. Finally, further deepening and improving applications of PROXYSPEX in data attribution and mechanistic interpretability as well as potentially exploring more complex value functions or larger-scale component interactions remains interesting future work.

Acknowledgments and Disclosure of Funding

This material is based upon work supported by the National Science Foundation Graduate Research Fellowship Program under Grant No. DGE-2146752. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

This work used NCSA DeltaAI at UIUC through allocation CIS250245 from the Advanced Cyberin-frastructure Coordination Ecosystem: Services & Support (ACCESS) program, which is supported by U.S. National Science Foundation grants #2138259, #2138286, #2138307, #2137603, and #2138296.

B.Y. gratefully acknowledge partial support from NSF grant DMS-2413265, NSF grant DMS 2209975, NSF grant 2023505 on Collaborative Research: Foundations of Data Science Institute (FODSI), the NSF and the Simons Foundation for the Collaboration on the Theoretical Foundations of Deep Learning through awards DMS-2031883 and 814639, NSF grant MC2378 to the Institute for Artificial CyberThreat Intelligence and OperatioN (ACTION), and NIH grant R01GM152718.

References

- [1] S. M. Lundberg and S.-I. Lee, "A unified approach to interpreting model predictions," in *Proceedings of the 31st International Conference on Neural Information Processing Systems*, ser. NIPS'17. Red Hook, NY, USA: Curran Associates Inc., 2017, p. 4768–4777.
- [2] M. T. Ribeiro, S. Singh, and C. Guestrin, ""why should i trust you?" explaining the predictions of any classifier," in *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*, 2016, pp. 1135–1144.
- [3] C.-P. Tsai, C.-K. Yeh, and P. Ravikumar, "Faith-shap: The faithful Shapley interaction index," *Journal of Machine Learning Research*, vol. 24, no. 94, pp. 1–42, 2023.
- [4] J. S. Kang, L. Butler, A. Agarwal, Y. E. Erginbas, R. Pedarsani, K. Ramchandran, and B. Yu, "Spex: Scaling feature interaction explanations for llms," *International Conference on Machine Learning*, 2025.
- [5] A. Gorji, A. Amrollahi, and A. Krause, "Shap values via sparse fourier representation," 2025. [Online]. Available: https://arxiv.org/abs/2410.06300
- [6] A. Krizhevsky, G. Hinton *et al.*, "Learning multiple layers of features from tiny images," 2009, Technical Report.
- [7] D. Hendrycks, C. Burns, S. Basart, A. Zou, M. Mazeika, D. Song, and J. Steinhardt, "Measuring massive multitask language understanding," in *International Conference on Learning Representations*, 2021.
- [8] A. Grattafiori, A. Dubey, A. Jauhri, A. Pandey, A. Kadian, A. Al-Dahle, A. Letman, A. Mathur, A. Schelten, A. Vaughan *et al.*, "The llama 3 herd of models," 2024. [Online]. Available: https://arxiv.org/abs/2407.21783
- [9] L. S. Shapley, A Value for N-Person Games. Santa Monica, CA: RAND Corporation, 1952.
- [10] R. Tibshirani, "Regression shrinkage and selection via the LASSO," *Journal of the Royal Statistical Society Series B: Statistical Methodology*, vol. 58, no. 1, pp. 267–288, 1996.
- [11] B. Cohen-Wang, H. Shah, K. Georgiev, and A. Mądry, "Contextcite: Attributing model generation to context," in *Advances in Neural Information Processing Systems*, A. Globerson, L. Mackey, D. Belgrave, A. Fan, U. Paquet, J. Tomczak, and C. Zhang, Eds., vol. 37. Curran Associates, Inc., 2024, pp. 95 764–95 807. [Online]. Available: https://proceedings.neurips.cc/paper_files/paper/2024/file/adbea136219b64db96a9941e4249a857-Paper-Conference.pdf
- [12] J. Chen, L. Song, M. Wainwright, and M. Jordan, "Learning to explain: An information-theoretic perspective on model interpretation," in *International conference on machine learning*. PMLR, 2018, pp. 883–892.

- [13] M. Sundararajan, A. Taly, and Q. Yan, "Axiomatic attribution for deep networks," in *International conference on machine learning*. PMLR, 2017, pp. 3319–3328.
- [14] A. Binder, G. Montavon, S. Lapuschkin, K.-R. Müller, and W. Samek, "Layer-wise relevance propagation for neural networks with local renormalization layers," in *International conference on artificial neural networks*. Springer, 2016, pp. 63–71.
- [15] M. Sundararajan, K. Dhamdhere, and A. Agarwal, "The Shapley Taylor interaction index," in *International Conference on Machine Learning*, Jul 2020, pp. 9259–9268. [Online]. Available: https://proceedings.mlr.press/v119/sundararajan20a.html
- [16] S. Bordt and U. von Luxburg, "From shapley values to generalized additive models and back," in International Conference on Artificial Intelligence and Statistics. PMLR, 2023, pp. 709–745.
- [17] F. Fumagalli, M. Muschalik, P. Kolpaczki, E. Hüllermeier, and B. E. Hammer, "SHAP-IQ: Unified approximation of any-order shapley interactions," in *Conference on Neural Information Processing Systems*, 2023. [Online]. Available: https://openreview.net/forum?id=IEMLNF4gK4
- [18] Q. Ren, J. Gao, W. Shen, and Q. Zhang, "Where we have arrived in proving the emergence of sparse interaction primitives in DNNs," in *International Conference on Learning Representations*, 2024. [Online]. Available: https://openreview.net/forum?id=3pWSL8My6B
- [19] M. Hahn and M. Rofin, "Why are sensitive functions hard for transformers?" *arXiv preprint arXiv:2402.09963*, 2024.
- [20] E. Abbe, S. Bengio, A. Lotfi, and K. Rizk, "Generalization on the unseen, logic reasoning and degree curriculum," *Journal of Machine Learning Research*, vol. 25, no. 331, pp. 1–58, 2024.
- [21] E. Abbe, E. Boix-Adsera, M. S. Brennan, G. Bresler, and D. Nagaraj, "The staircase property: How hierarchical structure can guide deep learning," *Advances in Neural Information Processing Systems*, vol. 34, pp. 26 989–27 002, 2021.
- [22] E. Abbe, E. B. Adsera, and T. Misiakiewicz, "The merged-staircase property: a necessary and nearly sufficient condition for sgd learning of sparse functions on two-layer neural networks," in *Conference on Learning Theory*. PMLR, 2022, pp. 4782–4887.
- [23] D. Tsui and A. Aghazadeh, "On recovering higher-order interactions from protein language models," in *ICLR 2024 Workshop on Generative and Experimental Perspectives for Biomolecular Design*, 2024. [Online]. Available: https://openreview.net/forum?id=WfA5oWpYT4
- [24] J. Ren, Z. Zhou, Q. Chen, and Q. Zhang, "Can we faithfully represent absence states to compute shapley values on a DNN?" in *International Conference on Learning Representations*, 2023. [Online]. Available: https://openreview.net/forum?id=YV8tP7bW6Kt
- [25] J. S. Kang, Y. E. Erginbas, L. Butler, R. Pedarsani, and K. Ramchandran, "Learning to understand: Identifying interactions via the möbius transform," *Conference on Neural Information Processing Systems*, 2024.
- [26] J. C. Harsanyi, "A bargaining model for the cooperative *n*-person game," Ph.D. dissertation, Department of Economics, Stanford University, Stanford, CA, USA, 1958.
- [27] M. Mohammadi, I. Tiddi, and A. Ten Teije, "Unlocking the game: Estimating games in möbius representation for explanation and high-order interaction detection," *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 39, no. 18, pp. 19512–19519, Apr. 2025. [Online]. Available: https://ojs.aaai.org/index.php/AAAI/article/view/34148
- [28] X. Li, J. K. Bradley, S. Pawar, and K. Ramchandran, "The SPRIGHT algorithm for robust sparse Hadamard Transforms," in *IEEE International Symposium on Information Theory (ISIT)*, 2014, pp. 1857–1861.
- [29] A. Amrollahi, A. Zandieh, M. Kapralov, and A. Krause, "Efficiently learning fourier sparse set functions," *Advances in Neural Information Processing Systems*, vol. 32, 2019.

- [30] Y. E. Erginbas, J. Kang, A. Aghazadeh, and K. Ramchandran, "Efficiently computing sparse fourier transforms of q-ary functions," in *IEEE International Symposium on Information Theory* (*ISIT*), 2023, pp. 513–518.
- [31] R. Scheibler, S. Haghighatshoar, and M. Vetterli, "A fast hadamard transform for signals with sublinear sparsity in the transform domain," *IEEE Transactions on Information Theory*, vol. 61, no. 4, pp. 2115–2132, 2015.
- [32] C. Olah, N. Cammarata, L. Schubert, G. Goh, M. Petrov, and S. Carter, "Zoom in: An introduction to circuits," *Distill*, 2020, https://distill.pub/2020/circuits/zoom-in.
- [33] A. Conmy, A. Mavor-Parker, A. Lynch, S. Heimersheim, and A. Garriga-Alonso, "Towards automated circuit discovery for mechanistic interpretability," *Advances in Neural Information Processing Systems*, vol. 36, pp. 16318–16352, 2023.
- [34] A. Syed, C. Rager, and A. Conmy, "Attribution patching outperforms automated circuit discovery," in *Proceedings of the 7th BlackboxNLP Workshop: Analyzing and Interpreting Neural Networks for NLP*, Y. Belinkov, N. Kim, J. Jumelet, H. Mohebbi, A. Mueller, and H. Chen, Eds. Miami, Florida, US: Association for Computational Linguistics, Nov. 2024, pp. 407–416. [Online]. Available: https://aclanthology.org/2024.blackboxnlp-1.25/
- [35] A. R. Hsu, G. Zhou, Y. Cherapanamjeri, Y. Huang, A. Odisho, P. R. Carroll, and B. Yu, "Efficient automated circuit discovery in transformers using contextual decomposition," in *The Thirteenth International Conference on Learning Representations*, 2025. [Online]. Available: https://openreview.net/forum?id=41HlN8XYM5
- [36] L. Sharkey, B. Chughtai, J. Batson, J. Lindsey, J. Wu, L. Bushnaq, N. Goldowsky-Dill, S. Heimersheim, A. Ortega, J. Bloom et al., "Open problems in mechanistic interpretability," arXiv preprint arXiv:2501.16496, 2025.
- [37] R. O'Donnell, Analysis of boolean functions. Cambridge University Press, 2014.
- [38] Y. Zhang, H. He, Z. Tan, and Y. Yuan, "Trade-off between efficiency and consistency for removal-based explanations," *Advances in Neural Information Processing Systems*, vol. 36, pp. 25 627–25 661, 2023.
- [39] G. Valle-Perez, C. Q. Camargo, and A. A. Louis, "Deep learning generalizes because the parameter-function map is biased towards simple functions," arXiv preprint arXiv:1805.08522, 2018.
- [40] G. Yang and H. Salman, "A fine-grained spectral perspective on neural networks," 2020. [Online]. Available: https://arxiv.org/abs/1907.10599
- [41] Q. Ren, J. Zhang, Y. Xu, Y. Xin, D. Liu, and Q. Zhang, "Towards the dynamics of a dnn learning symbolic interactions," in *Advances in Neural Information Processing Systems*, A. Globerson, L. Mackey, D. Belgrave, A. Fan, U. Paquet, J. Tomczak, and C. Zhang, Eds., vol. 37. Curran Associates, Inc., 2024, pp. 50 653–50 688. [Online]. Available: https://proceedings.neurips.cc/paper_files/paper/2024/file/5aa96d1caa0d0b99d534b67df06be2ff-Paper-Conference.pdf
- [42] T.-Y. Lin, M. Maire, S. J. Belongie, L. D. Bourdev, R. B. Girshick, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick, "Microsoft COCO: Common objects in context," in *European Conference on Computer Vision*, 2014, pp. 740–755.
- [43] Y. S. Tan, J. M. Klusowski, and K. Balasubramanian, "Statistical-computational trade-offs for recursive adaptive partitioning estimators," 2025. [Online]. Available: https://arxiv.org/abs/2411.04394
- [44] Y. Shuo Tan, A. Agarwal, and B. Yu, "A cautionary tale on fitting decision trees to data from additive models: generalization lower bounds," in *Proceedings of The 25th International Conference on Artificial Intelligence and Statistics*, ser. Proceedings of Machine Learning Research, G. Camps-Valls, F. J. R. Ruiz, and I. Valera, Eds., vol. 151. PMLR, 28–30 Mar 2022, pp. 9663–9685. [Online]. Available: https://proceedings.mlr.press/v151/shuo-tan22a.html

- [45] Y. S. Tan, C. Singh, K. Nasseri, A. Agarwal, J. Duncan, O. Ronen, M. Epland, A. Kornblith, and B. Yu, "Fast interpretable greedy-tree sums," *Proceedings of the National Academy of Sciences*, vol. 122, no. 7, p. e2310151122, 2025.
- [46] A. L. Maas, R. E. Daly, P. T. Pham, D. Huang, A. Y. Ng, and C. Potts, "Learning word vectors for sentiment analysis," in *Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*. Portland, Oregon, USA: Association for Computational Linguistics, June 2011, pp. 142–150. [Online]. Available: http://www.aclweb.org/anthology/P11-1015
- [47] V. Sanh, L. Debut, J. Chaumond, and T. Wolf, "Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter," 2020. [Online]. Available: https://arxiv.org/abs/1910.01108
- [48] K. Odabasi, "DistilBERT Finetuned Sentiment," accessed January. [Online]. Available: https://huggingface.co/lyrisha/distilbert-base-finetuned-sentiment
- [49] Z. Yang, P. Qi, S. Zhang, Y. Bengio, W. Cohen, R. Salakhutdinov, and C. D. Manning, "HotpotQA: A dataset for diverse, explainable multi-hop question answering," in *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, E. Riloff, D. Chiang, J. Hockenmaier, and J. Tsujii, Eds. Brussels, Belgium: Association for Computational Linguistics, Oct.-Nov. 2018, pp. 2369–2380. [Online]. Available: https://aclanthology.org/D18-1259/
- [50] D. Dua, Y. Wang, S. Dasigi, S. Singh, M. Gardner, and T. Kwiatkowski, "Drop: A reading comprehension benchmark requiring discrete reasoning over paragraphs," 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)*, 2019, pp. 2368–2378.
- [51] P. Kolpaczki, M. Muschalik, F. Fumagalli, B. Hammer, and E. Hüllermeier, "SVARM-IQ: Efficient approximation of any-order Shapley interactions through stratification," in *Proceedings of The 27th International Conference on Artificial Intelligence and Statistics*, ser. Proceedings of Machine Learning Research, S. Dasgupta, S. Mandt, and Y. Li, Eds., vol. 238. PMLR, 02–04 May 2024, pp. 3520–3528. [Online]. Available: https://proceedings.mlr.press/v238/kolpaczki24a.html
- [52] E. Candes and T. Tao, "Decoding by linear programming," *IEEE Transactions on Information Theory*, vol. 51, no. 12, pp. 4203–4215, 2005.
- [53] C. Musco and R. T. Witter, "Provably accurate shapley value estimation via leverage score sampling," in *The Thirteenth International Conference on Learning Representations*, 2025. [Online]. Available: https://openreview.net/forum?id=wg3rBImn3O
- [54] P. Kolpaczki, V. Bengs, M. Muschalik, and E. Hüllermeier, "Approximating the shapley value without marginal contributions," in *Proceedings of the AAAI conference on Artificial Intelligence*, vol. 38, no. 12, 2024, pp. 13 246–13 255.
- [55] A. Ilyas, S. M. Park, L. Engstrom, G. Leclerc, and A. Madry, "Datamodels: Understanding predictions with data and data with predictions," in *International Conference on Machine Learning*. PMLR, 2022, pp. 9525–9587.
- [56] H. Shah, A. Ilyas, and A. Mądry, "Decomposing and editing predictions by modeling model computation," in *Proceedings of the 41st International Conference on Machine Learning*, 2024, pp. 44 244–44 292.
- [57] S. Zhang, T. Han, U. Bhalla, and H. Lakkaraju, "Building bridges, not walls: Advancing interpretability by unifying feature, data, and model component attribution," in *ICLR 2025 Workshop on Navigating and Addressing Data Problems for Foundation Models*, 2025. [Online]. Available: https://openreview.net/forum?id=w5UVPcDCqQ
- [58] A. K. M. Abbas, K. Tirumala, D. Simig, S. Ganguli, and A. S. Morcos, "Semdedup: Data-efficient learning at web-scale through semantic deduplication," in *ICLR 2023 Workshop on Mathematical and Empirical Understanding of Foundation Models*, 2023.

- [59] W. Shi, S. Li, T. Liang, M. Wan, G. Ma, X. Wang, and X. He, "Route sparse autoencoder to interpret large language models," *CoRR*, vol. abs/2503.08200, March 2025. [Online]. Available: https://doi.org/10.48550/arXiv.2503.08200
- [60] B. Chen, T. Dao, E. Winsor, Z. Song, A. Rudra, and C. Ré, "Scatterbrain: Unifying sparse and low-rank attention," *Advances in Neural Information Processing Systems*, vol. 34, pp. 17413–17426, 2021.
- [61] M. Li and Q. Zhang, "Technical note: Defining and quantifying and-or interactions for faithful and concise explanation of dnns," 2024. [Online]. Available: https://arxiv.org/abs/2304.13312
- [62] E. Kushilevitz and Y. Mansour, "Learning decision trees using the fourier spectrum," in *Proceedings of the twenty-third annual ACM symposium on Theory of computing*, 1991, pp. 455–464.
- [63] Y. Mansour, "Learning boolean functions via the fourier transform," in *Theoretical advances in neural computation and learning*. Springer, 1994, pp. 391–424.

Appendices

A	Metl	hod Details	17
	A.1	Fourier Conversions	17
	A.2	Fourier Extraction	17
	A.3	Sparse Fourier Optimization	18
В	Expo	erimental Details	19
	B.1	Implementation Details	19
		B.1.1 Hyper-parameters	19
		B.1.2 Sentiment Analysis	19
		B.1.3 HotpotQA	19
		B.1.4 DROP	19
		B.1.5 MS-COCO	20
	B.2	Measuring Spectral Hierarchies	20
	B.3	Sparsification	20
	B.4	Proxy Model Selection	21
	B.5	Shapley Value Approximation	22
	B.6	Practical Implications	22
C	Case	e Study Details	23
	C.1	Data Attribution via Non-Linear Datamodels	23
	C_2	Model Component Attribution	28

A Method Details

A.1 Fourier Conversions

Interaction Index	Fourier Conversion
Banzhaf ψ_i	$-2F(\{i\})$
Shapley ϕ_i	$(-2)\sum_{\substack{S\supseteq\{i\}\\ S \text{ is odd}}}\frac{F(S)}{ S }$
Influence ξ_i	$\sum\limits_{S i i}F(S)^2$
Möbius $I^{M}(T)$	$(-2)^{ T } \sum_{S \supseteq T} F(S)$
Or $I^{\mathcal{O}}(T)$	$\begin{cases} \sum_{S \subseteq [n]} F(S) & \text{if } T = \emptyset \\ -(-2)^{ T } \sum_{S \supseteq T} (-1)^{ S } F(S) & \text{if } T \neq \emptyset \end{cases}$
Banzhaf Interaction $I^{\mathrm{B}}(T)$	-2F(T)
Shapley Interaction $I^{\rm S}(T)$	$(-2)^{ T } \sum_{S \supseteq T \ s.t. \ (-1)^{ S } = (-1)^{ T }} \frac{F(S)}{ S - T + 1}$
Shapley Taylor $I_\ell^{\mathrm{ST}}(T)$	$\begin{cases} I^{\mathbf{M}}(T), & T < \ell, \\ \sum_{S \supseteq T} { S \choose \ell}^{-1} I^{M}(S), & T = \ell. \end{cases}$
Faith-Banzhaf $I_\ell^{\mathrm{FB}}(T)$	$(-2)^{ T } \sum_{\substack{S \supseteq T \\ S \le \ell}} F(S)$
Faith-Shapley $I_\ell^{\mathrm{FS}}(T)$	$\begin{split} I^{M}(T) + (-1)^{\ell - T } \frac{ T }{\ell + T } \binom{\ell}{ T } \sum_{\substack{S \supset T \\ S > \ell}} F(S) \gamma(S, T, \ell) \\ \text{where } \gamma(S, T, \ell) = \sum_{\substack{T \subset R \subseteq S \\ R > \ell}} \frac{\binom{ R - 1}{\ell}}{\binom{ R + \ell - 1}{\ell + T }} (-2)^{ R } \end{split}$

The relationship between Fourier coefficients and influence scores are provided in [37]. We derive the conversion between Fourier and the OR interaction index [61] in this work. All remaining conversions are derived in Appendix C of [4].

A.2 Fourier Extraction

The exact Fourier transform of a decision tree can be computed recursively [5, 62, 63]. Due to the linearity of the Fourier transform, the Fourier transform of each boosted tree can be computed separately and added together. Algorithm 1, provided by [5], proceeds by traversing the nodes of each tree and summing the resultant Fourier transforms.

Algorithm 1 Fourier Extraction from Gradient Boosted Trees [5]

```
Require: Gradient boosted model \mathcal{M}
Ensure: Fourier mapping \mathcal{F}
 1: Initialize \mathcal{F} \leftarrow \emptyset
 2: for Tree T in \mathcal{M} do
             \mathcal{F} \leftarrow \mathcal{F}.\mathtt{merge}(\mathtt{EXTRACTTREE}(T.\mathtt{root}))
                                                                                                    > Add mappings of the individual trees
 4: end for
 5: return \mathcal{F}
 6: procedure EXTRACTTREE(node n)
            if n is leaf then
 7:
                   return \{\emptyset \mapsto n. \mathtt{value}\}
 8:
 9:
10:
                   \mathcal{N}_L \leftarrow \texttt{ExtractTree}(n.\texttt{leftChild})
                  \mathcal{N}_R \leftarrow \texttt{EXTRACTTREE}(n.\texttt{rightChild})
11:
12:
                   for S in (\mathcal{N}_L.\mathtt{keys} \cup \mathcal{N}_R.\mathtt{keys}) do
13:
                        \begin{array}{l} v_L \leftarrow \mathcal{N}_L[S] \\ v_R \leftarrow \mathcal{N}_R[S] \\ \mathcal{N}[S] \leftarrow (v_L + v_R)/2 \\ \mathcal{N}[S \cup \{n. \texttt{featureSplit}\}] \leftarrow (v_L - v_R)/2 \end{array}
                                                                                                        ▶ Mapping returns 0 if not contained
14:
15:
16:
17:
18:
                   end for
19:
            end if
            return \mathcal{N}
20:
21: end procedure
```

A.3 Sparse Fourier Optimization

We assume $\hat{f}(S)$ is a sparse, low-degree function with support \mathcal{K} :

$$\hat{f}(S) = \sum_{T \in \mathcal{K}} (-1)^{|S \cap T|} \hat{F}(T)$$

Equivalently, the function can be represented (and efficiently converted) under the Möbius transform. Converting Fourier to Möbius (via Appendix A.1), letting $\mathcal{K}^+ = \{R \subseteq T \mid T \in \mathcal{K}\}$, and applying the inverse Möbius transform:

$$\hat{f}(S) = \sum_{R \in \mathcal{K}^+, R \subseteq T} \hat{I}^M(R)$$

The optimization problem can then be expressed as a polynomial over $\{0,1\}$. Let \mathbf{x} be a binary vector of length n and $S = \{i \in [n] \mid x_i = 1\}$. We will focus on the maximization problem (minimization follows analogously).

$$\max_{S\subseteq[n]}\hat{f}(S) = \max_{\mathbf{x}\in\{0,1\}^n} \sum_{R\in\mathcal{K}^+} \hat{I}^M(R) \prod_{i\in R} x_i$$

To reduce the problem to a linear integer program, each monomial $\prod_{i \in R} x_i$ can be replaced with a decision variable y_R and the following constraints:

$$\max_{\mathbf{y} \in \{0,1\}^{|\mathcal{K}^+|}} \quad \sum_{R \in \mathcal{K}^+} \hat{I}^M(R) y_R \tag{7}$$

s.t.
$$y_R \leq y_Q$$
 $\forall Q \subset R, R \in \mathcal{K}^+$ (8)

$$\sum_{i \in R} y_{\{i\}} < |R| + y_R \quad \forall R \in \mathcal{K}^+$$
 (9)

The first constraint guarantees that whenever a monomial is activated (i.e. $x_i = 1 \ \forall i \in R$), all of its subsets are also activated. The second constraint ensures that if a monomial is deactivated (i.e. $\exists i \in R \ s.t. \ x_i = 0$), at least one of its constituent terms $(y_{\{i\}})$ is likewise deactivated. After the optimization is solved, the solution can be read-off from the univariate monomials $y_{\{i\}}$. These monomial terms can also be used to impose cardinality constraints on the solution, as was used in Section 4.2 and Section 5.2.

B Experimental Details

B.1 Implementation Details

B.1.1 Hyper-parameters

We performed 5-fold cross-validation over the following hyper-parameters for each of the models:

Model	Hyper-parameter
LASSO	L1 Reg. Param. λ (100 with $\lambda_{min}/\lambda_{max} = 0.001$)
SPEX	L1 Reg. Param. λ (100 with $\lambda_{min}/\lambda_{max} = 0.001$)
PROXYSPEX	Max. Tree Depth [3, 5, None]
	Number of Trees [500, 1000, 5000]
	Learning Rate [0.01, 0.1]
	L1 Reg. Param. λ (100 with $\lambda_{min}/\lambda_{max} = 0.001$)
Random Forest	Max. Tree Depth [3, 5, None]
	Number of Trees [100, 500, 1000, 5000]
Neural Network	Hidden Layer Sizes $\left[\left(\frac{n}{4} \right), \left(\frac{n}{4}, \frac{n}{4} \right), \left(\frac{n}{4}, \frac{n}{4}, \frac{n}{4} \right) \right]$
	Learning Rate [Constant, Adaptive]
	Learning Rate Init. [0.001, 0.01, 0.1]
	Number of Trees [100, 500, 1000, 5000]

B.1.2 Sentiment Analysis

20 movie reviews were used from the Large Movie Review Dataset [46] with $n \in [256, 512]$ words. To measure the sentiment of each movie review, we utilize a DistilBERT model [47] fine-tuned for sentiment analysis [48]. When masking, we replace the word with the [UNK] token. We construct an value function over the output logit associated with the positive class.

B.1.3 HotpotQA

We consider 50 examples from the HotpotQA[49] dataset between $n \in [64, 128]$ sentences. We use a Llama-3.2-3B-Instruct model with 8-bit quantization. When masking, we replace with the [UNK] token, and measure the log-perplexity of generating the original output. Since HotpotQA is a multi-document dataset, we use the following prompt format.

Title: {title_1}

Content: {document_1}

. . .

Title: {title m}

Content: {document_m}

Query: {question}. Keep your answers as short as possible.

B.1.4 DROP

We consider 50 examples from the DROP [49] dataset with $n \in [256, 512]$ number of words. We use the same model as HotpotQA and mask in a similar fashion. We use the following prompt format.

Context: {context}

Query: {question}. Keep your answers as short as possible.

B.1.5 MS-COCO

We utilize the Microsoft Common Objects in Context (MS-COCO) dataset [42], which comprises images paired with descriptive text captions. For our experiments, we treat image patches (there are 48 patches per image) and individual words from the captions as the input features. We used the first 50 examples from the test set, which had n (image patches + words) between the range of [60, 85].

To model the relationship between images and text, we employed the CLIP-ViT-B/32 model, a vision-language encoder designed to learn joint representations of visual and textual data. In our PROXYSPEX framework, when masking input features (either image patches or words), we replace them with a generic placeholder token suitable for the CLIP architecture (e.g., a zeroed-out patch vector or the text [MASK] words. The value function f(S) for a given subset of features S was defined as the contrastive loss among the other image/caption pairs. By measuring the change in this contrastive loss upon masking different feature subsets, we can attribute importance to individual features and their interactions in the context of joint image-text understanding.

B.2 Measuring Spectral Hierarchies

To quantify the hierarchical structure observed in the Fourier spectra of the LLMs under study, we introduce and analyze two key metrics: the Staircase Rate (SCR) and the Strong Hierarchy Rate (SHR). These metrics are computed based on the set of the k largest (in magnitude) Fourier coefficients, denoted as \mathcal{F}_k .

The Staircase Rate (SCR(f, k)) is defined as:

$$SCR(f,k) = \frac{1}{k} \sum_{S \in \mathcal{F}_k} \mathbb{1} \left\{ \exists (e_1, \dots, e_{|S|}) \in Perm(S) \text{ s.t. } \left(\forall j \in \{0, \dots, |S|\} : \bigcup_{l=1}^j \{e_l\} \in \mathcal{F}_k \right) \right\},$$

where \mathcal{F}_k denotes the k largest Fourier coefficients of f, and $\operatorname{Perm}(S)$ is the set of all ordered sequences of the elements in S.

(10)

The SCR measures the proportion of top-k Fourier coefficients F(S) for which there exists an ordering of its constituent elements $(e_1,\ldots,e_{|S|})$ such that all initial subsets (i.e., $e_1,\{e_1,e_2\},\ldots,S$ itself) are also among the top-k coefficients. A high SCR indicates that significant high-order interactions are built up from significant lower-order interactions in a step-wise or "staircase" manner.

The Strong Hierarchy Rate (SHR(f, k)) is defined as:

$$SHR(f,k) = \frac{1}{k} \sum_{S \in \mathcal{F}_k} \mathbb{1} \left\{ \forall S' \subseteq S, S' \in \mathcal{F}_k \right\}, \quad \text{where } \mathcal{F}_k \text{ denotes the } k \text{ largest}$$
Fourier coefficients of f .

The SHR is a stricter measure, quantifying the proportion of top-k coefficients F(S) for which all subsets of S (not just initial subsets, as in DSR) are also present in \mathcal{F}_k . A high SHR suggests a very robust hierarchical structure where the significance of an interaction implies the significance of all its underlying components.

Figure 11 visualizes these rates alongside faithfulness (R^2) for the Sentiment Analysis and MS-COCO datasets. These empirical results aim to demonstrate that LLM feature interactions exhibit significant hierarchical structure. The high SCR and SHR scores support the core motivation for PROXYSPEX: that important interactions are often built upon their lower-order subsets, a structure that Gradient Boosted Trees (GBTs) are well-suited to capture and exploit.

B.3 Sparsification

The process of sparsification is crucial for enhancing the interpretability of the explanations generated by PROXYSPEX . By retaining only the top k Fourier coefficients, we can achieve a more concise and understandable representation of the model's behavior without significantly compromising the faithfulness of the explanation. As demonstrated in Figure 5, a relatively small number of Fourier coefficients (approximately 200) are often sufficient to achieve faithfulness comparable to using a much larger set of coefficients for tasks like sentiment classification and image captioning (MS-COCO).

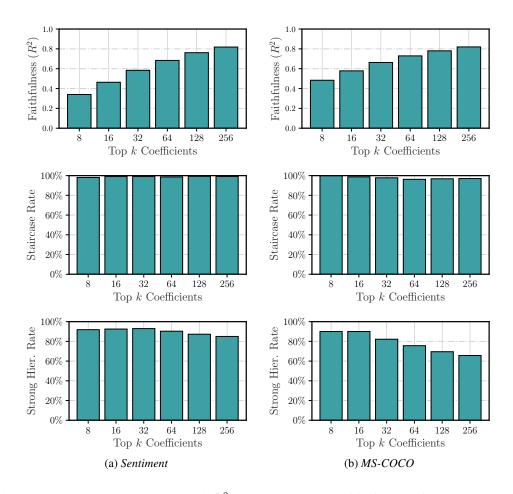


Figure 11: (top row) We run SPEX until $R^2>0.9$. We report the faithfulness of when we truncate the spectrum to keep just the top k coefficients for a range of k. We include results from Sentiment $n\in[256,512]$, and MS-COCO $n\in[60,85]$. In both cases faithfulness steadily increases as we increase k. (middle row) We report the SCR (10) for the same top k Fourier truncated functions above. In all cases, the SCR is nearly 100%. (bottom row) We also report the SHR (11), which is the strongest of the metric we consider. Here we find that even though SHR decreases somewhat as k grows, it is still strongly in favor of the hierarchy hypothesis.

Further results in Figure 12 illustrate the relationship between relative faithfulness and Fourier sparsity for both Sentiment and MS-COCO datasets across different inference multipliers (α). These plots show that faithfulness generally increases with k, plateauing after a certain number of coefficients, reinforcing the idea that a sparse representation can effectively capture the essential dynamics of the LLM's decision-making process.

B.4 Proxy Model Selection

The choice of GBTs as the proxy model within PROXYSPEX is motivated by their inherent ability to identify and learn hierarchical interactions from limited training data. This is a critical characteristic, as LLM feature interactions often exhibit a hierarchical structure where higher-order interactions are built upon their lower-order subsets. As indicated in the main text, GBTs have been shown to vastly outperform other proxy models, including random forests, particularly because random forests are less effective at learning hierarchical functions. GBT-like algorithms, on the other hand, are adept at disentangling sums of these hierarchical components.

Figure 13 provides a comparative view of proxy model performance. Figure 13a and Figure 13b illustrate the faithfulness (R^2) of different proxy models (LASSO, Random Forest, Neural Network,

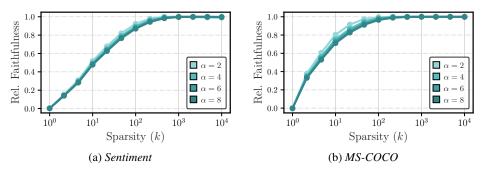


Figure 12: We plot faithfulness (R^2) as a function of Fourier sparsity. Only ≈ 200 coefficients are required to achieve equivalent faithfulness.

and GBTs) on both a synthetic dataset with a complete hierarchy (defined below) and the Sentiment Analysis dataset, respectively, across various inference parameters (α). These results empirically support the superiority of GBTs in capturing these complex interaction structures. However, it's also important to acknowledge limitations; for instance, GBTs may not perform as well when interactions possess a different, non-hierarchical sparse structure, as empirically confirmed by simulations like the Synthetic-Peak example (which lacks hierarchical structure) shown in Figure 13c.

Synthetic Peak	Synthetic Complete Hierarchy
$f^{\mathrm{SP}}(S) = \sum_{T \subset \mathcal{P}} (-1)^{ S \cap T } F(T)$	$f^{\mathtt{SCH}}(S) = \sum_{R \subset \mathcal{H}} (-1)^{ S \cap R } F(R)$
where \mathcal{P} is a set of 10 uniformly	where $\mathcal{H} = \{ R \subseteq T \mid T \in \mathcal{P} \}$
sampled sets of cardinality 5	and $F(R) \sim \text{Uniform}(-1,1)$ for $R \in \mathcal{H}$
and $F(T) \sim \text{Uniform}(-1,1)$ for $T \in \mathcal{P}$	

B.5 Shapley Value Approximation

We repeat the experiments of Section 4.3 under the metric of mean squared error relative to those computed by KernalSHAP under an inference budget of 10,000. In Figure 14, we find that PROXYSPEX uniformly outperforms KernelSHAP within this tested range. Just as with recall, with large enough α , KernelSHAP eventually surpasses PROXYSPEX.

B.6 Practical Implications

The practical implications of PROXYSPEX are significant, primarily revolving around its inference efficiency and the resulting speedups in generating faithful explanations for LLMs. A major challenge with existing interaction attribution methods, like SPEX, is the substantial number of model inferences required, which can be computationally expensive and time-consuming for large models. PROXYSPEX addresses this by leveraging a GBT proxy model, which dramatically reduces the number of inferences needed while maintaining or even improving explanation faithfulness.

Figure 15 presents the practical benefits in terms of wall clock time for achieving different levels of faithfulness (R^2) on the Sentiment Analysis (Figure 15a) and MS-COCO (Figure 15b) datasets. These plots clearly demonstrate the speedups achieved by PROXYSPEX. For example, in the sentiment analysis task using the smaller Distilbert model, PROXYSPEX offers a speedup of approximately 3x, while for the larger CLIP-ViT-B/32 model with MS-COCO, the speedup is around 5x when compared to methods that require more extensive sampling. This increased efficiency makes PROXYSPEX a more viable tool for interpreting complex LLMs in real-world scenarios where computational resources and time are often constrained.

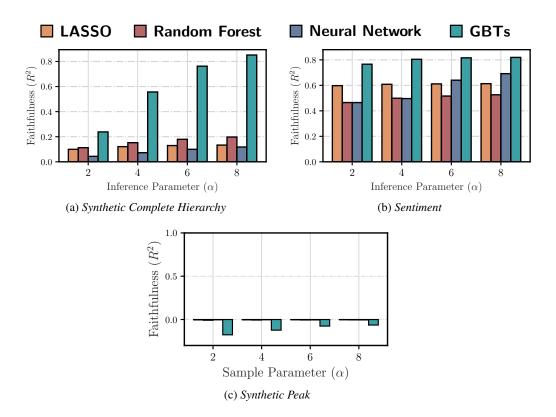


Figure 13: Comparison of proxy model faithfulness in capturing function structures. (a) Faithfulness of LASSO, Random Forest, Neural Network, and GBTs on a synthetic dataset with a complete hierarchical structure, across varying inference parameters (α). (b) Faithfulness of the same proxy models on the Sentiment Analysis dataset across varying α . (c) Faithfulness on a synthetic dataset with a sparse, non-hierarchical peak function, across varying α , illustrating a limitation of GBTs for non-hierarchical structures.

C Case Study Details

C.1 Data Attribution via Non-Linear Datamodels

The training masks and margin outputs were provided by [55], corresponding to their subsampling rate of 50% (i.e., half the training images were used to fit each model). See [55] for the hyperparameters selected. With $n=50,\!000$ training samples, $300,\!000$ training masks (model retrainings) were provided. This corresponds to $\alpha\approx0.38$, which underscores the inference-efficiency of PROXYSPEX to identify strong interactions.

Utilizing these masks and margins, we randomly selected 60 test images (6 from each class) for analysis with PROXYSPEX. Below, in Figure 16 and Figure 17, we present the strongest second-order interactions of the first thirty of these selected test images. Figure 9 visualizes the six test images exhibiting the most significant third-order interactions identified through this analysis.

After fitting PROXYSPEX, we convert the Fourier interactions to Möbius using Appendix A.1. Since target and non-target images affect the test margin in opposite directions, we partition the interaction space into the following categories:

- *Target-class interactions* \mathcal{T} : Interactions composed exclusively of training images that share the same label as the held-out test image.
- Non-target-class interactions \mathcal{T}^c : Interactions where at least one training image in the set has a label different from that of the held-out test image.

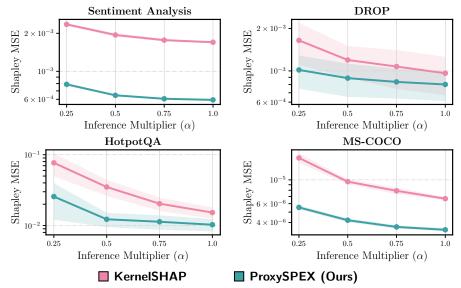


Figure 14: Shapley value mean square error after $\alpha \cdot n \log_2(n)$ inferences for multipliers $\alpha \in \{0.25, 0.5, 0.75, 1.0\}$. Across all four tasks and multipliers, within the tested range, PROXYSPEX provides a better approximation of the values computed under 10,000 inferences. Error bands indicate the standard deviation across ten different runs of the algorithms.

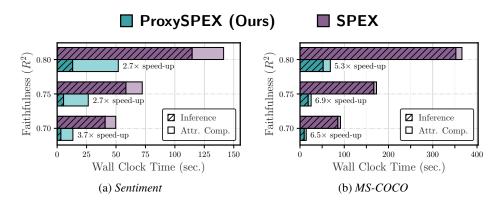


Figure 15: Wall clock time demonstrating PROXYSPEX's efficiency. Comparison of wall clock time (seconds) required to achieve different levels of faithfulness (\mathbb{R}^2) for PROXYSPEX, showing breakdown of inference time and attribution computation time. (a) Results on the Sentiment Analysis dataset with the Distilbert model. (b) Results on the MS-COCO dataset with the CLIP-ViT-B/32 model, highlighting speedups achieved by PROXYSPEX.

Synergistic Interactions: The top synergistic interaction R^* of order-r is defined as:

$$S^* = \underset{S \in \mathcal{T}, |S| = r}{\operatorname{argmax}} I^M(S)$$

$$T^* = \underset{T \in \mathcal{T}^c, |T| = r}{\operatorname{argmin}} I^M(T)$$

$$R^* = \begin{cases} S^* & \text{if } |I^M(S^*)| \ge |I^M(T^*)| \\ T^* & \text{otherwise} \end{cases}$$
(12)

Visually, as presented in Figure 16 for r=2, the interactions R^* identified by this rule often involve training images that appear to work together to reinforce or clarify the classification of the held-out image, frequently by contributing complementary features or attributes. It is important to acknowledge that this definition serves as a heuristic and does not perfectly isolate synergy; For

example, the first frog image contains redundant bird images due to strong higher-order interactions involving these bird images.

Redundant Interactions: The top redundant interaction R^* of order-r is defined as:

$$S^* = \underset{S \in \mathcal{T}, |S| = r}{\operatorname{argmin}} I^M(S)$$

$$T^* = \underset{T \in \mathcal{T}^c, |T| = r}{\operatorname{argmax}} I^M(T)$$

$$R^* = \begin{cases} S^* & \text{if } |I^M(S^*)| \ge |I^M(T^*)| \\ T^* & \text{otherwise} \end{cases}$$
(13)

Figure 17 demonstrates that this definition identifies redundant training images that are similar to the held-out image.

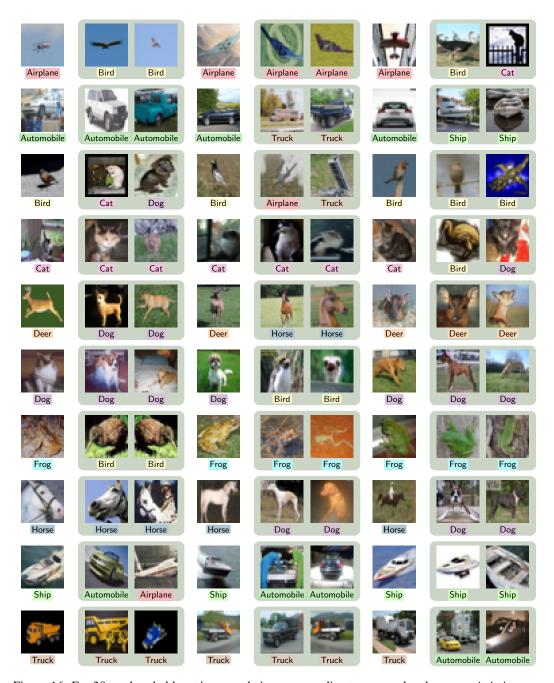


Figure 16: For 30 random held-out images, their corresponding top second-order *synergistic interaction* (green box).

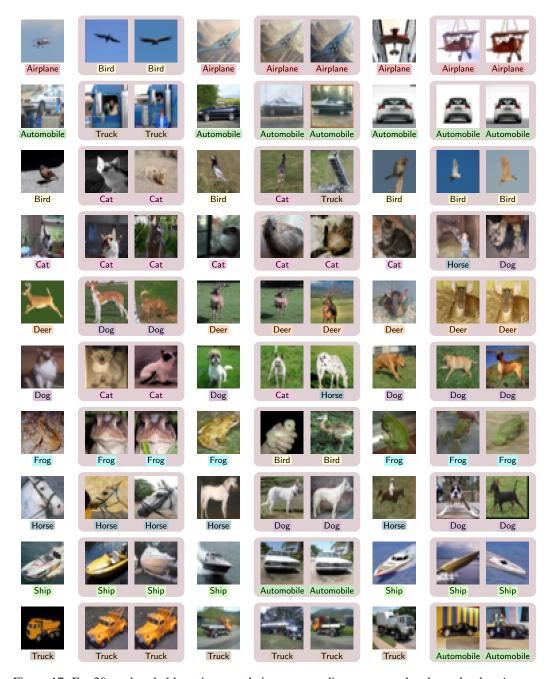


Figure 17: For 30 random held-out images, their corresponding top second-order *redundant interaction* (red box).

C.2 Model Component Attribution

We study the influence of specific model components on task performance, using a controlled ablation methodology. Our experiments are conducted on Llama-3.1-8B-Instruct evaluated on the high-school-us-history subset of the MMLU dataset, a benchmark comprising multiple-choice questions.

MMLU includes 231 questions in the high-school-us-history subset. To perform pruning and then evaluate the ablated models, we split this data into two sets—training split \mathcal{D}_{train} consisting of the first 120 questions and test split \mathcal{D}_{test} with the remaining questions. We use accuracy as the evaluation metric, which is computed as the proportion of correctly answered multiple-choice questions on a given data split.

For an L layer LLM, we let [L] denote the set of layers and let \mathcal{H}_ℓ denote the set of attention heads in layer $\ell \in [L]$. For each experiment, we focus on a particular group of layers $\mathcal{L} \subseteq [L]$ within the model and denote the corresponding set of attention heads as $\mathcal{H}_{\mathcal{L}} = \bigcup_{\ell \in \mathcal{L}} \mathcal{H}_{\ell}$. The Llama-3.1-8B-Instruct model consists of L=32 layers, each with 32 attention heads.

At each layer ℓ of the LLM, the output of the attention heads is combined into a latent representation by concatenating the outputs of the attention heads. Then, this latent vector is passed to the feed-forward network of layer ℓ . To study the contribution of specific heads, we define an ablated model LLM $_S$ for any subset $S \subseteq \mathcal{H}_{\mathcal{L}}$. In LLM $_S$, the outputs of the heads in $\mathcal{H}_{\mathcal{L}} \setminus S$ are set to zero before the concatenation step. After concatenation, we apply a rescaling factor to the resulting latent vector at each layer $\ell \in \mathcal{L}$, equal to the inverse of the proportion of retained heads in that layer, i.e., $\frac{|\mathcal{H}_{\ell}|}{|S \cap \mathcal{H}_{\ell}|}$. This modified latent representation is then passed to the feed-forward network as usual.

We define $f_{\mathcal{L}}$ as

$$f_{\mathcal{L}}(S) \triangleq \text{Accuracy of LLM}_S \text{ on } \mathcal{D}_{\text{train}},$$
 (14)

and interpret $f_{\mathcal{L}}(S)$ as a proxy for the functional contribution of head subset S to model performance, enabling quantitative analyses of attribution and interaction effects among attention heads.

Pruning. We perform pruning experiments across three different layer groups \mathcal{L} : initial layers $(\mathcal{L} = \{1, 2, 3\})$, middle layers $(\mathcal{L} = \{14, 15, 16\})$, and final layers $(\mathcal{L} = \{30, 31, 32\})$. Since each layer has 32 attention heads, we effectively perform ablation over $n = |\mathcal{H}_{\mathcal{L}}| = 96$ features (attention heads) in total. For a given group \mathcal{L} , we begin by estimating the function $f_{\mathcal{L}}$ using both LASSO and PROXYSPEX, based on evaluations of $f_{\mathcal{L}}(S)$ for 5000 subsets S sampled uniformly at random. These estimates serve as surrogates for the true head importance function. We then maximize the estimated functions to identify the most important attention heads under varying sparsity constraints (target numbers of retained heads). We use the procedure detailed in Section 4.2 to identify heads to remove for both PROXYSPEX and LASSO. We also compare against a Best-of-N baseline, in which the model is pruned by selecting the subset S that achieves the highest value of $f_{\mathcal{L}}(S)$ among 5000 randomly sampled subsets at the target sparsity level.

Evaluation. In order to evaluate the performance of an ablated model LLM_S , we measure its accuracy on the test set using

$$g_{\mathcal{L}}(S) \triangleq \text{Accuracy of LLM}_S \text{ on } \mathcal{D}_{\text{test}}.$$
 (15)

In Figure 10, we report the value of $g_{\mathcal{L}}(S)$ for the pruned models obtained by each method. We find that PROXYSPEX consistently outperforms both baselines, yielding higher test accuracy across all evaluated sparsity levels.

Inference setup. All experiments are run on a single NVIDIA H100 GPU, with batch size 50. Average runtime per ablation (i.e., evaluating $f_{\mathcal{L}}(S)$ once for a given S) is approximately 1.7 seconds. Therefore, collecting a training dataset $\{(S_i, f_{\mathcal{L}}(S_i))\}$ with 5000 training samples takes approximately 2.5 hours.

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification: The claims in the paper are clearly stated. The results, which are primarily empirical, are backed by experimental data. Relevant theoretical research is also cited.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the
 contributions made in the paper and important assumptions and limitations. A No or
 NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals
 are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: There is a limitations section, which specifically highlights work that remains to be done.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory assumptions and proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [Yes]

Justification: We make no major theoretical claims, and any theoretical statements are accompanied by proofs.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and crossreferenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental result reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: All code and experimental setup are provided.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
- (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
- (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
- (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
- (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: All data and code will be included in the publication.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (https://nips.cc/public/guides/CodeSubmissionPolicy) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so "No" is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (https://nips.cc/public/guides/CodeSubmissionPolicy) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental setting/details

Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: All relevant details are included.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental
 material.

7. Experiment statistical significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [Yes]

Justification: Error bars are provided where relevant.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)

- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error
 of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments compute resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: Details on the compute used is included.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code of ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics https://neurips.cc/public/EthicsGuidelines?

Answer: [Yes]
Justification: Yes.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a
 deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [Yes]

Justification: Yes, this is presented in the work.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.

- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: No such risks.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [NA]

Justification: all assets are open-source.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.

• If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. New assets

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [Yes]

Justification: All code and new tools will be published and reasonably documented.

Guidelines:

- The answer NA means that the paper does not release new assets.
- · Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. Crowdsourcing and research with human subjects

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: No human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. Institutional review board (IRB) approvals or equivalent for research with human subjects

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: The paper does not involve crowdsourcing nor research with human subjects Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- · For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. Declaration of LLM usage

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does not impact the core methodology, scientific rigorousness, or originality of the research, declaration is not required.

Answer: [NA]

Justification: LLMs are not a non-standard component of the core methods in this research. Guidelines:

- The answer NA means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy (https://neurips.cc/Conferences/2025/LLM) for what should or should not be described.