

Select or Project? Evaluating Lower-dimensional Vectors for LLM Training Data Explanations

Lukas Hinterleitner¹, Loris Schoenegger^{1,2}, Benjamin Roth^{1,3}

¹Faculty of Computer Science, University of Vienna, Vienna, Austria

²UniVie Doctoral School Computer Science, University of Vienna, Vienna, Austria

³Faculty of Philological and Cultural Studies, University of Vienna, Vienna, Austria

Correspondence: contact@lukas-hinterleitner.at

Abstract

Gradient-based methods for instance-based explanation for large language models (LLMs) are hindered by the immense dimensionality of model gradients. In practice, influence estimation is restricted to a subset of model parameters to make computation tractable, but this subset is often chosen ad hoc and rarely justified by systematic evaluation. This paper investigates if it is better to create low-dimensional representations by *selecting* a small, architecturally informed subset of model components or by *projecting* the full gradients into a lower-dimensional space. Using a novel benchmark, we show that a greedily selected subset of components captures the information about training data influence needed for a retrieval task more effectively than either the full gradient or random projection. We further find that this approach is more computationally efficient than random projection, demonstrating that targeted component selection is a practical strategy for making instance-based explanations of large models more computationally feasible.

1 Introduction

Different modalities for explaining language model (LM) behavior in terms of causal influences have been proposed recently, including feature-based (Chen et al., 2023; Enguehard, 2023), or mechanistic explanations (Meng et al., 2022; Lieberum et al., 2023). While these approaches explain language model behavior in terms of inputs or architectural components, recent work has enabled the generation of instance-based explanations that provide insights grounded in the model’s training data (Hara et al., 2019; Pruthi et al., 2020; Guo et al., 2021; Grosse et al., 2023). In this paper, we specifically focus on methods that identify influential training examples through the analysis of model gradients, a signal widely exploited by XAI methods (e.g., Pruthi et al., 2020; Park et al., 2023; Koh and Liang, 2017). By retrieving a set of examples

that influenced a given model output, these methods enable model debugging and interpretability tasks (e.g., Koh and Liang, 2017; Pruthi et al., 2020; Guo et al., 2021). However, their application to modern large language models is challenging due to high computational costs: for a 1B-parameter model, a single gradient requires over 4 GB of memory. To make analysis tractable, it must often be restricted to a subset of model parameters (e.g., the word-embedding layer: Yeh et al., 2022). In existing work, this choice was often made ad hoc and rarely justified by systematic evaluation. In this paper, we address this issue by investigating whether it is better to *select* a sparse, structurally informed subset of model components or to *project* the entire high-dimensional gradient into a dense, lower-dimensional space. Specifically, we systematically compare two strategies for creating gradient surrogates: An architecture-aware greedy layer selection algorithm proposed in this work, that selects a set of layer components whose combined gradients are most informative. And *random projection*, a widely used, architecture-agnostic technique that produces a low-dimensional representation of the full gradient (Johnson and Lindenstrauss, 1984).

We evaluate using a **novel benchmark for instance-based explanations**, designed to test how well a given selection supports a retrieval task in which the original training example must be identified among alternative promptcompletion pairs generated using various strategies. Our **main contributions** are threefold:

- (1) We propose an efficient retrieval-based benchmark for instance-based explanations.
- (2) We conduct a principled comparison of architecture-aware selection versus geometry-preserving projection for attribution.
- (3) We demonstrate that our targeted selection strategy can result in more accurate attributions while also being more efficient.

We find that a greedily selected subset of model

components captures the required information about training data influence more accurately than either the full gradient or random projection. We also find that this targeted selection is more computationally efficient than random projection. Our central insight is that a carefully chosen subset can provide a clearer and more discriminative signal for certain tasks, making it a practical strategy for instance-based explanations of large models.¹

2 Related Work

Instance-based explanation and attribution

Instance-based explanation aims to identify influential training examples that explain a model’s prediction. A dominant framework is *leave one out influence* (the effect of removing a given training example has on the prediction), which can be approximated with re-training free methods such as *influence functions* (Koh and Liang, 2017; Grosse et al., 2023). Alternative approaches, such as Trac-InCP (Pruthi et al., 2020), aim to approximate training data influence *during* the training process by comparing the test instance’s loss gradient to training gradients at a set of model checkpoints.

The challenge of high-dimensional gradients

These influence estimation methods involve comparisons of model gradients (and calculating inverse Hessian-vector products for influence functions), which remains computationally intractable for LLMs: recent applications therefore restrict computation to a subset of model parameters, such as the multilayer perceptron (MLP) layers (Grosse et al., 2023) or LoRA layers (Kwon et al., 2023; Schoenegger and Roth, 2026). Another common choice is to restrict influence estimation to the model’s final layers (Schioppa et al., 2022; Akyürek et al., 2022; Grosse et al., 2023). However, Yeh et al. (2022) find this to be suboptimal for LLMs due to a *cancellation effect* in deeper layers, and propose using early layer components instead, for example, the more stable word embedding gradients. An alternative line of work projects full model gradients to a lower-dimensional space where pairwise distances between examples are preserved (e.g., Park et al., 2023; Xia et al., 2024; Lin et al., 2024). Our work systematically compares the utility of *all* architectural components and random projection to determine which strategy provides the most informative signal for retrieval.

¹We make our code available at <https://doi.org/10.5281/zenodo.1834666>.

Evaluation of training data explanations

Grosse et al. (2023) evaluate their approximation of influence functions by reporting agreement with the *proximal Bergmann response function*, whose approximation does not require re-training. However, as they note themselves, it remains unclear whether this is a useful measure of data influence. Akyürek et al. (2022) evaluate by introducing facts with the fine-tuning data, that the model of interest did not know before, and then verify whether influence estimation methods correctly attribute importance to this fine-tuning set when queried about them. Such evaluation strategies require expensive additional training steps that are infeasible with LLMs. Our evaluation setup is most similar to Wang et al. (2025) who use paraphrased data to test the performance of influence estimation methods like *Data Shapley* (Ghorbani and Zou, 2019) and influence functions by identifying the original sources of rewritten text. Similarly, our evaluation setup is re-training free; we use it to systematically compare dimensionality reduction and gradient selection strategies for influence estimation.

3 Method: Gradient Similarity for Instance-based Explanations

Instance-based explanation methods must be capable of distinguishing training examples that influence model output from those that do not. Popular methods do so based on model loss gradients. As we will outline in Section 3.2 we frame this requirement as a retrieval problem: To study the quality of different low-dimensional representations, our setup measures how well they enable one to identify a given test example among distractors. We introduce a novel low-dimensional representation strategy in Section 3.3 and describe our setup for random projection in Section 3.4. Our evaluation approach is (1) **local**, targeting decisions at the instance level (each training sample/instance processed individually); (2) **static**, using the final model checkpoint without reconstructing training dynamics; and (3) **model-invariant**, applicable to any model that provides per-sample gradients.

3.1 Gradient Representation and Similarity

Let $D = \{s_i\}_{i=1}^N$ be a subset of a fine-tuning dataset for a given LLM with parameters θ . The training objective is to minimize a loss function, $\mathcal{L}(s_i, \theta)$. The gradient for a single sample s_i with respect to the final, fine-tuned model parameters




Figure 1: Overview of the evaluation setup for the *paraphrased* setting. The gradient of a paraphrased sample s_{p_i} is compared to gradients from its original counterpart s_i and other candidates (e.g., s_j). The method succeeds at this retrieval task if the cosine similarity is highest for the original pair. The setup is identical for the *generated* setting.

θ_T , denoted $\nabla_{\theta} \mathcal{L}(s_i, \theta_T)$, indicates the direction in the parameter space that would reduce the loss for that specific sample. It serves as a representation of the sample’s influence on the model before being utilized by an optimizer, like AdamW (Loshchilov and Hutter, 2019). The model’s parameters θ are a collection of component tensors $\{\mathbf{W}^{(l,k)}\}$, where l indexes the layer and k the specific component within it (e.g., query, key, or value projections in an attention block; MLP layers). Let $\mathcal{W}$ be the set of all such index pairs (l, k) . To create a single vector for comparison, we first calculate the gradient w.r.t each component, $\nabla_{\mathbf{W}^{(l,k)}} \mathcal{L}(s_i, \theta_T)$ and secondly, flatten each gradient before combining all of them into a single vector $\mathbf{v}_{\theta}(s_i) \in \mathbb{R}^M$ (we refer to this as the full model gradient or full gradient), where M is the total number of parameters (see Appendix A.1 for more details on gradient flattening). We use cosine similarity $\text{sim}_{\cos}(\cdot, \cdot)$ to then compare gradient vectors (see Appendix A.2, on how cosine similarity suits our approach).

3.2 Evaluation via Retrieval

A sample’s loss gradient encodes how a training step on it would adapt the model’s parameters. Gradient-based attribution methods build on this: TracIn-style methods (Pruthi et al., 2020) score influence directly by the dot product between training- and test-sample loss gradients, the first-order effect of a training step on the test loss, while influence functions (Koh and Liang, 2017) additionally weight this inner product by the inverse

Hessian. A paraphrase should adapt the model similarly to its original, so the original should be identified as most influential; our retrieval task turns this into a testable necessary condition: a representation that cannot rank the original above closely related alternatives cannot yield reliable influence rankings.

To study the utility of different representations, such as using the *full* gradient for influence estimation (Section 3.2), *selecting* a subset of model parameters (Section 3.3), or random *projection* (Section 3.4), we consider the following retrieval task: We first derive two different query datasets (D_p and D_m) from the original fine-tuning dataset D by paraphrasing each individual training sample. We use two different strategies: **Paraphrased** (D_p), where each sample $s_{p_i} \in D_p$ is a paraphrase of $s_i \in D$, preserving semantics while altering lexical form. And **Model-Generated** (D_m), where for each s_i , we paraphrase its prompt and use the fine-tuned model itself to generate a new completion, creating $s_{m_i} \in D_m$. This tests the model’s understanding of its own outputs for semantically similar inputs. For simplicity, we describe our method using the paraphrased set D_p throughout the paper, the process for D_m is identical.

Our initial assumption is that the cosine similarity between the gradients of a paraphrased sample s_{p_i} and its original counterpart s_i is higher than the similarity between the gradients of s_{p_i} and another sample from the training data $s_j \in D$. Specifically,

we test whether:

$$\text{sim}_{\cos}(\mathbf{v}_{\theta}(s_{p_i}), \mathbf{v}_{\theta}(s_i)) > \text{sim}_{\cos}(\mathbf{v}_{\theta}(s_{p_i}), \mathbf{v}_{\theta}(s_j)),$$

where $i \neq j$. If this assumption holds, then we should also be able to retrieve the original counterpart for every paraphrased sample s_{p_i} (Figure 1). For each $s_{p_i} \in D_p$, we therefore test if its original s_i is ranked highest among a set of candidates $C_i(b) \subset D$. This set is formed by retrieving the top- b most lexically similar samples to s_{p_i} from D using BM25 (Robertson and Zaragoza, 2009).² This candidate set is created to reduce the search space and hence computing fewer gradients. Our evaluation metric is retrieval accuracy:

$$\text{acc}_{\theta}^{(b)}(D_p, D) = \frac{1}{N} \sum_{i=1}^N \mathbb{I} \left[\text{sim}_{\cos}(\mathbf{v}_{\theta}(s_{p_i}), \mathbf{v}_{\theta}(s_i)) > \max_{j \in C_i(b) \setminus \{i\}} \text{sim}_{\cos}(\mathbf{v}_{\theta}(s_{p_i}), \mathbf{v}_{\theta}(s_j)) \right]$$

where $\mathbb{I}[\cdot]$ is the Iverson bracket.

3.3 Greedy Component Selection (Select)

Our approach constructs a smaller, architecturally-informed surrogate gradient by greedily selecting a subset of model components up to a given budget that are most informative for the retrieval task by leveraging the linearity of the dot product: The dot product of concatenated vectors is simply the sum of the dot products of the constituent vectors. This allows us to reconstruct the cosine similarity for any subset of component indices $\mathcal{S} \subseteq \mathcal{W}$ by summing pre-computed scalar values. Let $\mathbf{v}_{\mathcal{S}}(s_i)$ be the gradient vector formed by concatenating component gradients for indices in $\mathcal{S}$ (see Appendix A.2, for a more detailed view on the cosine similarity reconstruction). We employ a greedy forward algorithm that iteratively adds the component yielding the greatest improvement in retrieval accuracy. Let $\mathcal{S}$ be the set of already selected layer component indices. At each step, we choose the next component index (l^*, k^*) to add:

$$(l^*, k^*) = \arg \max_{(l,k) \in \mathcal{W} \setminus \mathcal{S}} \text{acc}_{\mathcal{S} \cup \{(l,k)\}}^{(b)}(D_p, D)$$

This process is efficient as it operates entirely on pre-computed dot products, avoiding re-computation or storage of high-dimensional gradients during search.

²Given the user prompts have a short median length of 17 words, simpler lexical-overlap measures than BM25 would likely suffice; the original counterpart is always included in the candidate set (Appendix B.2).

3.4 Random Projection (Project)

Following previous work (e.g., Park et al., 2023; Xia et al., 2024; Lin et al., 2024), we create a dense surrogate gradient by projecting the full gradient into a lower-dimensional space using an efficient GPU-optimized implementation from Park et al. (2023). A random projection matrix $\mathbf{R} \in \mathbb{R}^{d \times M}$ (where $d \ll M$) maps the full gradient $\mathbf{v}_{\theta}(s_i)$ to a smaller vector $\mathbf{p}(s_i) \in \mathbb{R}^d$:

$$\mathbf{p}(s_i) := \frac{1}{\sqrt{d}} \mathbf{R} \mathbf{v}_{\theta}(s_i).$$

Due to the immense size of $\mathbf{R}$ and $\mathbf{v}_{\theta}(s_i)$, we cannot materialize them directly. Instead, we perform the projection in a component-wise fashion. For each component $\mathbf{W}^{(l,k)}$, we use a separate random matrix $\mathbf{R}^{(l,k)}$, where the projection dimension for that component is proportional to its parameter count. The final low-dimensional representation is the concatenation of these individual projections. This provides an architecture-agnostic baseline that aims to preserve the geometric structure of the full gradient per sample $\mathbf{v}_{\theta}(s_i)$.

4 Experimental Setup

4.1 Model and Data

We conduct our experiments using AMD-OLMo-1B-SFT (Groeneveld et al., 2024; Liu et al., 2024, Apache license 2.0), a 1.2B parameter decoder-only transformer model. We choose this model as its training data is made available in full. Specifically, it was pre-trained on a 1.3T token subset of the Dolma dataset (Soldaini et al., 2024), and fine-tuned on the Tülu 2 (Iverson et al., 2023), OpenHermes-2.5 (Teknium, 2023), WebInstructSub (Yue et al., 2024), and Code-Feedback (Zheng et al., 2024) datasets.

Including the input embedding layer, we consider a total of $|\mathcal{W}| = 113$ distinct component tensors for our analysis. This comprises 16 layers, each containing seven trainable weight matrices: four for the attention mechanism (Q, K, V, and O projections) and three for the MLP block (Gate, Up, and Down projections).

Our base dataset for the evaluation setup, D , is the LIMA (Zhou et al., 2023, CC BY-NC-SA license) subset of Tülu 2 (Iverson et al., 2023): we use all samples ($\approx 1k$) except for a small fraction of multi-turn dialogues, yielding $N = 988$ single-turn prompt-generation pairs. We do not apply any further sampling or filtering. Prompts are typically

short (median 21, mean 54 tokens); generations are longer (median 369, mean 587 tokens). All samples are formatted using the model’s chat template, which wraps user instructions in `<|user|>` and marks the generation point with `<|assistant|>`, ensuring the input aligns with the model’s fine-tuning format. We provide s_0 as an example below:

```
<|user|>
Can brain cells move? ...
<|assistant|>
The question is relatively broad...
```

4.2 Implementation Details

The sheer scale of a billion-parameter model introduces significant computational hurdles. The full flattened gradient for a single sample, $\mathbf{v}_\theta(s_i)$, is a vector in $\mathbb{R}^{1.2B}$, requiring approximately 4.7 GB of storage, which makes naively storing all gradients infeasible. To make the retrieval task tractable, the candidate set $C_i(b)$ is limited to $b = 5$ samples. This targeted approach, pre-filtered with BM25, reduces the number of required gradient similarity computations by over 99%.

Intermediate dot product storage To implement the *Select* method without storing high-dimensional vectors, we compute and store only their intermediate, component-wise dot products. For each component $(l, k) \in \mathcal{W}$ and every candidate pair of samples (s_i, s_j) , we pre-compute the scalar values $\mathbf{v}_{\mathbf{W}^{(l,k)}}(s_i)^\top \mathbf{v}_{\mathbf{W}^{(l,k)}}(s_j)$ and the self-dot-products for both s_i and s_j . These values can be aggregated post-hoc to reconstruct similarity scores for any subset of components (see Appendix A.2 for additional information).

Layer-wise random projection The *project* baseline is implemented using a component-wise approach because a naive projection of the entire 1.2B-parameter gradient is impossible, as the projection matrix itself would require petabytes of storage. We instead apply a separate random projection to each of the 113 vectorized component gradients. This was implemented using the CudaProjector from the TRAK library (Park et al., 2023).

Evaluation dataset construction The query datasets were constructed using specific models. For the paraphrased set (D_p), we used GPT-4o-mini to paraphrase every sample in the base dataset D . For example, s_{p_0} generated from s_0 is:

```
<|user|>
Are brain cells capable of moving? ...
<|assistant|>
The inquiry is quite extensive...
```

For the model-generated set (D_m), we fed the paraphrased prompts from D_p to the fine-tuned OLMo model to create corresponding generations.

5 Results and Analysis

We evaluate in two settings: *paraphrased*, where inputs are semantically similar, and *model-generated*, where the model produces a novel generation to the paraphrased prompt.

5.1 Full Gradient Performance

First, we evaluate the full gradient as a baseline. As shown in Table 1, the full gradient is highly effective in the paraphrased setting (D_p), achieving near-perfect retrieval accuracy of 0.993. However, its performance collapses to 0.218 on the more challenging model-generated dataset D_m , which is only marginally better than the random baseline of $0.20 = 1/5$, when considering a candidate set with size $b = 5$. A qualitative analysis of the few failure cases in the paraphrased setting reveals that for small amount ($< 1\%$) of instruction-based prompts, the paraphrasing LLM *executed* instructions rather than rephrasing them. For example:

Original Prompt: translate into English:
"Der Zug..."
Paraphrased Prompt: The train arrives in
Frankfurt...

The paraphrased prompt contains the *result* of the instruction, not the instruction itself. This fundamentally alters the task, shifting the token distribution and the resulting gradient, which likely causes the retrieval to fail.

5.2 Single Layer Component Performance

Interestingly, gradients from a small subset of model components perform as well as, or even better than, the full model gradient. Table 1 shows the mean retrieval accuracy for each component type, averaged across all layers: In the paraphrased setting, most components achieve near-perfect accuracy, indicating that the signal is broadly spread. In the model-generated setting, however, performance is concentrated in specific components. The MLP *gate* and *up* projections carry the most informative signal (≈ 0.26), while the attention’s *key* (0.208) and *value* (0.198) projections contribute little useful information for retrieval, with the latter performing worse than random chance.

Component Type	Paraphrased (Mean Acc.)	Model-Generated (Mean Acc.)
<i>MLP Components</i>		
MLP Down-Proj	0.993	0.235
MLP Gate-Proj	0.992	0.256
MLP Up-Proj	0.991	0.247
<i>Attention Components</i>		
Attn Output-Proj	0.992	0.245
Attn Query-Proj	0.976	0.255
Attn Value-Proj	0.98	0.198
Attn Key-Proj	0.916	0.208
<i>Embedding</i>		
Embedding Layer	0.993	0.231
Full Model Grad.	0.993	0.218

Table 1: Mean retrieval accuracy per component type. In the model-generated setting, MLP components are most informative, while some attention components fail.

The role of component size and function We investigate if parameter count correlates with performance. As shown in Figure 2, there is no clear monotonic relationship, as components of the same type share the same parameter count. The three boxplots correspond to the attention projections (left, 4.2M parameters), the MLP layers (center, 16.8M parameters), and the embedding layer (right, 103M parameters). The medium-sized MLP components on average outperform both the much larger embedding layer and the smaller attention projections. Furthermore, within the attention mechanism, all projection components have the same size, yet their individual performance varies significantly. This suggests that a component’s functional role is a more critical determinant of its explanatory power than its parameter count.




Figure 2: Component accuracy vs. parameter count. Each boxplot groups all components of a specific type, which share the same size. **Left:** Attention projections. **Center:** MLP components. **Right:** Embedding layer.

Performance by model depth The accuracy distribution of influential gradient components across model depth differs substantially between settings. While accuracy is high and stable across all layers in the paraphrased setting, the model-generated one reveals a distinct pattern as evident in Figure 3. Performance is strongest in the initial and final layers, with a noticeable dip in the middle layers. This suggests that while early layers (near the input representation) and late layers (near the output logits) retain some alignment with the original sample, intermediate layers, which perform more abstract semantic transformations (Tenney et al., 2019; Jawahar et al., 2019), appear less informative when the model generates a novel response. MLP layer components also regain performance in later layers, while the attention components (except Q) remain equal or worse. This is in line with the observations of Yeh et al. (2022), who also observe that early layers appear more informative than later ones for training data attribution.

5.3 Accuracy vs. Full Gradient Alignment

Given that single gradient components can be highly informative, we explore whether a small *subset* of components can outperform the full gradient in retrieval accuracy. Furthermore, we assess how well a *subset* of components aligns with the full gradient by measuring how similar their cosine similarities are to those of the full gradient.

Selection by retrieval accuracy Figure 4 shows that greedily selecting components based on retrieval accuracy identifies a small subset that significantly outperforms both random projection and the full gradient. In the paraphrased setting, accuracy peaks at 0.998 with less than 5% of parameters for the greedy selection. This demonstrates a key finding: certain components of the full gradient appear noisy for the task at hand, and retrieval can be improved by filtering out less-informative elements. Even in the challenging model-generated setting, a small subset of components achieves an accuracy of ≈ 0.36 , nearly double that of the full gradient.

Selection by similarity to the full gradient In contrast, when we select components to maximize cosine similarity with the full gradient scores (Figure 5), we observe that random projection is highly effective, achieving $\gtrsim 0.999$ similarity with just 1% of parameters in the paraphrased setting. Also in the model-generated case, random projections outperform component selection. However, this




Figure 3: Accuracy vs. layer depth for the model-generated setting. Each color represents a layer component.




Figure 4: Greedy component selection (optimized for accuracy) vs. Random Projection. The x-axis shows the cumulative fraction of parameters utilized. In both settings, a small, greedily-selected subset outperforms the full gradient (dashed line) and random projection.

high fidelity does not translate to optimal retrieval performance. This confirms that preserving the full gradient’s global geometry is a different, and for our task less effective, objective than identifying the components most salient for retrieval.


5.4 Analysis of Computational Cost

The computational requirements of our proposed architecture-aware selection strategy is substantially lower than performing random projection: precomputing dot products for all components took approximately 4 hours. The subsequent greedy selection process finished in a few minutes. In contrast, random projection baselines (without


caching) were computationally more intensive, requiring over 900 hours. However, due to storage limitations, gradient-caching was infeasible. For completeness, we calculated how much faster random projection would be if gradient-caching was applied without adding any read-write costs (see results in Table 2). All computations were performed on NVIDIA H100 GPUs.

6 Discussion

Our findings challenge the notion that the full model gradient is the best source of information for instance-based explanation and reveal how well different layer components encode instance-specific



(a) Similarity (Paraphrased)



(b) Similarity (Model-generated)

Figure 5: Greedy component selection vs. random projection, optimizing for similarity to the full model gradient. While random projection quickly replicates the full gradient’s geometry, this does not yield the best retrieval accuracy (cf. Figure 4).

Computation Method	Time (Hours)
Dot-Product Pre-computation	~4
Greedy Selection (Post-processing)	~0.03
Total for Greedy Method	~4.03
1% RP Baseline (without caching)	~ 158
5% RP Baseline (without caching)	~ 765
Total RP Baselines (without caching)	≥ 900
1% RP Baseline (with caching)	~ 53
5% RP Baseline (with caching)	~ 255
Total RP Baselines (with caching)	≥ 300

Table 2: Computational cost comparison for the paraphrased setting. The greedy selection method remains substantially more efficient than random projection (RP), even if ignoring storage costs of gradient-caching.

information in our retrieval task. We provide a set of insights to improve future explanation setups.

The suitability of the full gradient The full gradient is often outperformed by some individual component gradients in the *paraphrased* setting. Smaller subsets not only perform better in terms of accuracy but are also smaller, and therefore easier to store and retrieve. In the *model-generated* setting, the full gradient is a weak baseline, with accuracy collapsing to 0.218 (vs. 0.2 random chance). Selecting a small subset of parameters consistently outperforms the full gradient, proving that its signal is diluted by weaker layer components in the case of instance-based methods.

Insights from analyzing gradient components

Layer-wise analysis reveals where instance-specific information is encoded: We observe that not all lay-

ers contribute equally. Specifically, MLP blocks and attention Q/O projections carry more information than K/V projections, regardless of size for the task we consider. This highlights the significance of a component’s *functional role* over its parameter count. We furthermore observe that greedy selection shows a non-monotonic link between parameter budget and accuracy. A small and optimal subset of components achieves competitive performance, while adding more components appears to introduce noise and degrades accuracy toward the full-gradient baseline. Lastly, influence distribution across depth varies by setting. In the model-generated setting, early and late layers appear most useful on average, while middle layers underperform. However, in later layers only MLP and attention Q blocks perform well.

6.1 To Select or to Project?

Our evaluation indicates that not all components are equally beneficial when retrieving relevant training samples, suggesting that subset selection should be preferred over random projection, which does not account for the informativeness of components during retrieval. Restricting influence estimation to a subset of layers can retain retrieval accuracy while improving computational efficiency. Consistent with this, our simple greedy parameter selection approach outperforms common random projection strategies. Lastly, with appropriate adaptation to task specifics, this approach could also support applications that rely on similar loss-gradient comparisons, such as data cleaning- or pruning methods.

7 Conclusion

This work systematically evaluated strategies for reducing the computational requirements of gradient-based methods for instance-based explanations, comparing the targeted selection of a small subset of model components to random projection, and the use of the full gradient. We find that our proposed method of selecting a compact, structurally informed subset consistently outperforms full gradients and projections in both accuracy and efficiency.

Our results highlight that a small, carefully chosen set of components can be sufficient to capture the instance- and task-specific information required for generating instance-based explanations. Consequently, our work demonstrates that architecture-aware selection is an effective strategy to making instance-based explanations for large models more computationally feasible. Equipped with our evaluation setup, future work should further explore efficient strategies for component selection.

Limitations

Our experiments are conducted with a 1.2B-parameter model and its fine-tuning dataset, as obtaining random projections for the full gradient already requires over 300 GPU hours in our setup. Given sufficient resources, one could explore generalization across several axes, for example, analyzing different model architectures, scales, and data domains. Additionally, future work should extend analysis to *dynamic* attribution methods that track influence throughout the training process at multiple checkpoints (Hammoudeh and Lowd, 2024).

Acknowledgments

This research has been funded by the Vienna Science and Technology Fund (WWTF) [10.47379/VRG19008] "Knowledge-infused Deep Learning for Natural Language Processing". We also thank Lin et al. (2024) for providing the illustration templates of their paper *Token-wise Influential Training Data Retrieval for Large Language Models*, on which Figure 1 builds.

References

Ekin Akyürek, Tolga Bolukbasi, Frederick Liu, Binbin Xiong, Ian Tenney, Jacob Andreas, and Kelvin Guu. 2022. [Towards tracing knowledge in language models back to the training data](#). In *Findings of the*

Association for Computational Linguistics: EMNLP 2022, Abu Dhabi, United Arab Emirates, December 7-11, 2022, pages 2429–2446. Association for Computational Linguistics.

Hugh Chen, Ian C. Covert, Scott M. Lundberg, and Su-In Lee. 2023. [Algorithms to estimate Shapley value feature attributions](#). *Nat. Mac. Intell.*, 5(6):590–601.

Joseph Enguehard. 2023. [Sequential integrated gradients: a simple but effective method for explaining language models](#). In *Findings of the Association for Computational Linguistics: ACL 2023, Toronto, Canada, July 9-14, 2023*, pages 7555–7565. Association for Computational Linguistics.

Amirata Ghorbani and James Zou. 2019. [Data Shapley: Equitable valuation of data for machine learning](#). In *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 2242–2251. PMLR.

Dirk Groeneveld, Iz Beltagy, Evan Pete Walsh, Akshita Bhagia, Rodney Kinney, Oyvind Tafjord, Ananya Harsh Jha, Hamish Ivison, Ian Magnusson, Yizhong Wang, Shane Arora, David Atkinson, Russell Authur, Khyathi Raghavi Chandu, Arman Cohan, Jennifer Dumas, Yanai Elazar, Yuling Gu, Jack Hessel, and 24 others. 2024. [OLMo: Accelerating the science of language models](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2024, Bangkok, Thailand, August 11-16, 2024*, pages 15789–15809. Association for Computational Linguistics.

Roger B. Grosse, Juhan Bae, Cem Anil, Nelson Elhage, Alex Tamkin, Amirhossein Tajdini, Benoit Steiner, Dustin Li, Esin Durmus, Ethan Perez, Evan Hubinger, Kamile Lukosiute, Karina Nguyen, Nicholas Joseph, Sam McCandlish, Jared Kaplan, and Samuel R. Bowman. 2023. [Studying large language model generalization with influence functions](#). *CoRR*, abs/2308.03296.

Han Guo, Nazneen Rajani, Peter Hase, Mohit Bansal, and Caiming Xiong. 2021. [FastIF: Scalable influence functions for efficient model interpretation and debugging](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing, EMNLP 2021, Virtual Event / Punta Cana, Dominican Republic, 7-11 November, 2021*, pages 10333–10350. Association for Computational Linguistics.

Zayd Hammoudeh and Daniel Lowd. 2024. [Training data influence analysis and estimation: A survey](#). *Machine Learning*, 113(5):2351–2403.

Satoshi Hara, Atsushi Nitanda, and Takanori Maehara. 2019. [Data cleansing for models trained with SGD](#). In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, De-*

- ember 8-14, 2019, Vancouver, BC, Canada, pages 4215–4224.
- Hamish Ivison, Yizhong Wang, Valentina Pyatkin, Nathan Lambert, Matthew E. Peters, Pradeep Dasigi, Joel Jang, David Wadden, Noah A. Smith, Iz Beltagy, and Hannaneh Hajishirzi. 2023. [Camels in a changing climate: Enhancing LM adaptation with Tulu 2](#). *CoRR*, abs/2311.10702.
- Ganesh Jawahar, Benoît Sagot, and Djamé Seddah. 2019. [What does BERT learn about the structure of language?](#) In *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28- August 2, 2019, Volume 1: Long Papers*, pages 3651–3657. Association for Computational Linguistics.
- William Johnson and Joram Lindenstrauss. 1984. [Extensions of Lipschitz mappings into a Hilbert space](#). *Contemporary Mathematics*, 26:189–206.
- Pang Wei Koh and Percy Liang. 2017. [Understanding black-box predictions via influence functions](#). In *Proceedings of the 34th International Conference on Machine Learning, ICML 2017, Sydney, NSW, Australia, 6-11 August 2017*, volume 70 of *Proceedings of Machine Learning Research*, pages 1885–1894. PMLR.
- Yongchan Kwon, Eric Wu, Kevin Wu, and James Zou. 2023. [DataInf: Efficiently estimating data influence in LoRA-tuned LLMs and diffusion models](#). In *The Twelfth International Conference on Learning Representations*.
- Tom Lieberum, Matthew Rahtz, János Kramár, Neel Nanda, Geoffrey Irving, Rohin Shah, and Vladimir Mikulik. 2023. [Does circuit analysis interpretability scale? evidence from multiple choice capabilities in Chinchilla](#). *CoRR*, abs/2307.09458.
- Huawei Lin, Jikai Long, Zhaozhuo Xu, and Weijie Zhao. 2024. [Token-wise influential training data retrieval for large language models](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 841–860, Bangkok, Thailand. Association for Computational Linguistics.
- Jiang Liu, Jialian Wu, Prakamya Mishra, Zicheng Liu, Sudhanshu Ranjan, Pratik Prabhanjan Brahma, Yusheng Su, Gowtham Ramesh, Peng Sun, Zhe Li, Dong Li, Lu Tian, and Emad Barsoum. 2024. [AMD-OLMo: A series of 1B language models trained from scratch by AMD on AMD Instinct MI250 GPUs](#).
- Ilya Loshchilov and Frank Hutter. 2019. [Decoupled weight decay regularization](#). In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net.
- Kevin Meng, David Bau, Alex Andonian, and Yonatan Belinkov. 2022. [Locating and editing factual associations in GPT](#). In *Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 - December 9, 2022*.
- Sung Min Park, Kristian Georgiev, Andrew Ilyas, Guillaume Leclerc, and Aleksander Madry. 2023. [TRAK: Attributing model behavior at scale](#). In *International Conference on Machine Learning, ICML 2023, 23-29 July 2023, Honolulu, Hawaii, USA*, volume 202 of *Proceedings of Machine Learning Research*, pages 27074–27113. PMLR.
- Garima Pruthi, Frederick Liu, Satyen Kale, and Mukund Sundararajan. 2020. [Estimating training data influence by tracing gradient descent](#). In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020*.
- Stephen E. Robertson and Hugo Zaragoza. 2009. [The probabilistic relevance framework: BM25 and beyond](#). *Found. Trends Inf. Retr.*, 3(4):333–389.
- Andrea Schioppa, Polina Zablotskaia, David Vilar, and Artem Sokolov. 2022. [Scaling up influence functions](#). In *Thirty-Sixth AAAI Conference on Artificial Intelligence, AAAI 2022, Thirty-Fourth Conference on Innovative Applications of Artificial Intelligence, IAAI 2022, The Twelveth Symposium on Educational Advances in Artificial Intelligence, EAAI 2022 Virtual Event, February 22 - March 1, 2022*, pages 8179–8186. AAAI Press.
- Loris Schoenegger and Benjamin Roth. 2026. [Compact example-based explanations for language models](#). In *Findings of the Association for Computational Linguistics: ACL 2026*, pages 24246–24265, San Diego, California, United States. Association for Computational Linguistics.
- Luca Soldaini, Rodney Kinney, Akshita Bhagia, Dustin Schwenk, David Atkinson, Russell Authur, Ben Bogin, Khyathi Raghavi Chandu, Jennifer Dumas, Yanai Elazar, Valentin Hofmann, Ananya Harsh Jha, Sachin Kumar, Li Lucy, Xinxin Lyu, Nathan Lambert, Ian Magnusson, Jacob Morrison, Niklas Muennighoff, and 17 others. 2024. [Dolma: an open corpus of three trillion tokens for language model pretraining research](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2024, Bangkok, Thailand, August 11-16, 2024*, pages 15725–15788. Association for Computational Linguistics.
- Yi Sui, Ga Wu, and Scott Sanner. 2021. [Representer point selection via local Jacobian expansion for post-hoc classifier explanation of deep neural networks and ensemble models](#). In *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*, pages 23347–23358.
- Teknum. 2023. [OpenHermes 2.5: An open dataset of synthetic data for generalist LLM assistants](#).

Ian Tenney, Dipanjan Das, and Ellie Pavlick. 2019. [BERT rediscovers the classical NLP pipeline](#). In *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28- August 2, 2019, Volume 1: Long Papers*, pages 4593–4601. Association for Computational Linguistics.

Jiachen (Tianhao) Wang, Prateek Mittal, Dawn Song, and Ruoxi Jia. 2025. [Data Shapley in one training run](#). In *International Conference on Learning Representations*, volume 2025, pages 12358–12395.

Mengzhou Xia, Sadhika Malladi, Suchin Gururangan, Sanjeev Arora, and Danqi Chen. 2024. [LESS: Selecting influential data for targeted instruction tuning](#). In *Proceedings of the 41st International Conference on Machine Learning*, pages 54104–54132. PMLR.

Chih-Kuan Yeh, Ankur Taly, Mukund Sundararajan, Frederick Liu, and Pradeep Ravikumar. 2022. [First is better than last for language data influence](#). In *Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 - December 9, 2022*.

Xiang Yue, Tianyu Zheng, Ge Zhang, and Wenhui Chen. 2024. [MAmmoTH2: Scaling instructions from the web](#). *Advances in Neural Information Processing Systems*, 37:90629–90660.

Tianyu Zheng, Ge Zhang, Tianhao Shen, Xueling Liu, Bill Yuchen Lin, Jie Fu, Wenhui Chen, and Xiang Yue. 2024. [OpenCodeInterpreter: Integrating code generation with execution and refinement](#). In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 12834–12859, Bangkok, Thailand. Association for Computational Linguistics.

Chunting Zhou, Pengfei Liu, Puxin Xu, Srinivasan Iyer, Jiao Sun, Yuning Mao, Xuezhe Ma, Avia Efrat, Ping Yu, Lili Yu, Susan Zhang, Gargi Ghosh, Mike Lewis, Luke Zettlemoyer, and Omer Levy. 2023. [LIMA: Less is more for alignment](#). In *Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 - 16, 2023*.

A Methodological Details

A.1 Gradient Flattening

To create a single vector for comparison, we first compute the gradient w.r.t. each component tensor, $\nabla_{\mathbf{W}^{(l,k)}} \mathcal{L}(s_i, \theta_T)$. We then apply a row-major vectorization function, $\text{vec}_r(\cdot)$, to each gradient tensor and concatenate the resulting vectors using a collector operator, $\text{col}(\cdot)$. This forms the full flattened gradient vector $\mathbf{v}_\theta(s_i) \in \mathbb{R}^M$, where M is the total number of parameters:

$$\mathbf{v}_\theta(s_i) := \text{col} \left(\left\{ \text{vec}_r(\nabla_{\mathbf{W}^{(l,k)}} \mathcal{L}(s_i, \theta_T)) \right\}_{(l,k) \in \mathcal{W}} \right)$$

A.2 Cosine Similarity Reconstruction via Dot Products

The standard dot product $\mathbf{a}^\top \mathbf{b}$ is sensitive to both the angle between vectors and their magnitudes. Since gradient magnitudes can vary significantly across samples (Sui et al., 2021), a large norm could dominate the score and mask the true geometric alignment. To create a magnitude-invariant measure, we use cosine similarity, which isolates the directional component of influence:

$$\text{sim}_{\cos}(\mathbf{a}, \mathbf{b}) = \frac{\mathbf{a}^\top \mathbf{b}}{\|\mathbf{a}\|_2 \|\mathbf{b}\|_2}.$$

A key property for our **Select** method is that the cosine similarity for any subset of components $\mathcal{S} \subseteq \mathcal{W}$ can be perfectly reconstructed by summing pre-computed, component-wise dot products. For a query sample s_{q_i} and a candidate sample s_j , we first define the pre-computed scalar dot products for each component $(l, k) \in \mathcal{W}$ as:

$$\delta_{i,j}^{(l,k)} := \mathbf{v}_{\mathbf{W}^{(l,k)}}(s_{q_i})^\top \mathbf{v}_{\mathbf{W}^{(l,k)}}(s_j).$$

The cosine similarity for the subset $\mathcal{S}$ is then reconstructed by summing these values:

$$\text{sim}_{\cos}(\mathbf{v}_{\mathcal{S}}(s_{q_i}), \mathbf{v}_{\mathcal{S}}(s_j)) = \frac{\sum_{(l,k) \in \mathcal{S}} \delta_{i,j}^{(l,k)}}{\sqrt{\sum_{(l,k) \in \mathcal{S}} \delta_{i,i}^{(l,k)}} \sqrt{\sum_{(l,k) \in \mathcal{S}} \delta_{j,j}^{(l,k)}}}.$$

This allows us to efficiently evaluate any subset of components without storing or re-computing high-dimensional gradient vectors.

B Algorithmic Details

B.1 Query Dataset Construction

Paraphrasing prompt To create the *paraphrased* query set (D_p), each sample from the LIMA subset D was rewritten using GPT-4o-mini, guided by the system prompt in Figure 6.

You are a paraphrasing expert who is specialized in rewriting text (questions, statements, etc.) without altering the content. Keep in mind, that the meaning must not change after the paraphrasing. Just output the paraphrased text without any additional information.

Figure 6: System prompt for D_p

B.2 Candidate Set Creation

To make the retrieval task computationally tractable, we pre-filter the search space for each query sample s_{q_i} . Instead of comparing its gradient against all N samples in the original dataset D , we identify a small candidate set $\mathcal{C}_i(b)$ containing the indices of the top- b most lexically similar samples using the BM25 ranking function (Robertson and Zaragoza, 2009). This significantly reduces the number of expensive gradient computations from $\mathcal{O}(N)$ to $\mathcal{O}(b)$ per query. Algorithm 1 details this process. To ensure the correct original counterpart s_i is always included for evaluation, we add its index to the set if it is not already present among the top- b candidates, replacing the least similar one.

Data: Query sample s_{q_i} ; Original dataset $D = \{s_j\}_{j=1}^N$; Set size b
Result: Set of candidate indices $\mathcal{C}_i(b)$

```

for  $j \leftarrow 1$  to  $N$  do
  |  $\text{score}[j] \leftarrow \text{BM25}(s_{q_i}, s_j)$ ;
end
 $\mathcal{C}_i(b) \leftarrow \text{argsort\_top}(\text{score}, b)$ ;
; // Get indices of top  $b$  scores

// Ensure original counterpart's index  $i$  is present
if  $i \notin \mathcal{C}_i(b)$  then
  |  $j_{\min} \leftarrow \arg \min_{j' \in \mathcal{C}_i(b)} \text{score}[j']$ ;
  | ; // Find index of least similar candidate
  |  $\mathcal{C}_i(b) \leftarrow (\mathcal{C}_i(b) \setminus \{j_{\min}\}) \cup \{i\}$ ;
  | ; // Replace it with index  $i$ 
end
return  $\mathcal{C}_i(b)$ 

```

Algorithm 1: Candidate Set Creation via BM25

B.3 Greedy Selection by Accuracy

The *Select* method employs a greedy forward algorithm to find a subset of components that maximizes retrieval accuracy. Algorithm 2 provides a formal specification. The algorithm iteratively adds the component that yields the greatest improvement to the accuracy metric, operating entirely on the pre-computed dot products $\delta_{i,j}^{(l,k)}$. We introduce the notation $D_{i,j}^{\mathcal{S}} := \sum_{(l,k) \in \mathcal{S}} \delta_{i,j}^{(l,k)}$ to denote the accumulated dot product for a selected set $\mathcal{S}$.

Data: Candidate index sets $\{\mathcal{C}_i(b)\}_{i=1}^N$; Pre-computed dot products $\{\delta_{i,j}^{(l,k)}\}$; A small constant $\varepsilon > 0$.
Result: Ordered list of (component, accuracy) pairs.

Initialize accumulated dot products $D_{i,j}^{\emptyset} \leftarrow 0$ for all required pairs (i, j) .;
 $\mathcal{S} \leftarrow \emptyset$ (selected set), $\mathcal{R} \leftarrow \mathcal{W}$ (remaining set).;

```

for  $t \leftarrow 1$  to  $|\mathcal{W}|$  do
  |  $a^* \leftarrow -\infty, (l^*, k^*) \leftarrow \text{None}$ .;
  | foreach  $(l, k) \in \mathcal{R}$  do
  |   | Let  $\mathcal{S}' = \mathcal{S} \cup \{(l, k)\}$ .;
  |   | Compute reconstructed similarities  $\hat{\gamma}_{i,j}^{\mathcal{S}'} \leftarrow \frac{D_{i,j}^{\mathcal{S}'}}{\sqrt{D_{i,i}^{\mathcal{S}'}} \sqrt{D_{j,j}^{\mathcal{S}'} + \varepsilon}}$ .;
  |   |  $a \leftarrow \frac{1}{N} \sum_{i=1}^N \mathbb{I}[\hat{\gamma}_{i,i}^{\mathcal{S}'} > \max_{j \in \mathcal{C}_i(b) \setminus \{i\}} \hat{\gamma}_{i,j}^{\mathcal{S}'}]$ .;
  |   | if  $a > a^*$  then
  |   |   |  $a^* \leftarrow a, (l^*, k^*) \leftarrow (l, k)$ .;
  |   | end
  | end
  | end
  |  $\mathcal{S} \leftarrow \mathcal{S} \cup \{(l^*, k^*)\}, \mathcal{R} \leftarrow \mathcal{R} \setminus \{(l^*, k^*)\}$ .;
  | Update accumulators  $D_{i,j}^{\mathcal{S}} \leftarrow D_{i,j}^{\mathcal{S} \setminus \{(l^*, k^*)\}} + \delta_{i,j}^{(l^*, k^*)}$ .;
  | Store  $((l^*, k^*), a^*)$  in the output list.;
```

end
return *Ordered list*.

Algorithm 2: Forward Greedy Selection by Retrieval Accuracy