# BAYESIAN INFLUENCE FUNCTIONS FOR HESSIAN-FREE DATA ATTRIBUTION

# **Anonymous authors**

Paper under double-blind review

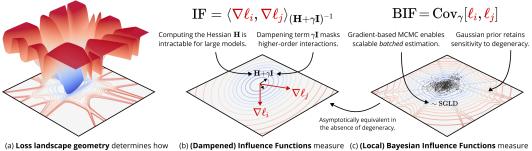
#### **ABSTRACT**

Classical influence functions face significant challenges when applied to deep neural networks, primarily due to non-invertible Hessians and high-dimensional parameter spaces. We propose the local Bayesian influence function (BIF), an extension of classical influence functions that replaces Hessian inversion with loss landscape statistics that can be estimated via stochastic-gradient MCMC sampling. This Hessian-free approach captures higher-order interactions among parameters and scales efficiently to neural networks with billions of parameters. We demonstrate state-of-the-art results on predicting retraining experiments.

# 1 Introduction

Training data attribution (TDA) studies how training data shapes the behaviors of deep neural networks (DNNs)—a foundational question in AI interpretability and safety (Lehalleur et al., 2025; Koh & Liang, 2020). A standard approach to TDA is influence functions (IF), which measure how models respond to infinitesimal perturbations in the training distribution (Cook, 1977; Cook & Weisberg, 1982). While elegant, influence functions rely on calculating the inverse Hessian and, therefore, break down for modern DNNs. Theoretically, DNNs have degenerate loss landscapes with *non-invertible* Hessians, which violate the conditions needed to define influence functions. Practically, for large models, the Hessian is intractable to compute directly. To mitigate these problems requires architecture-specific approximations that introduce structural biases (Martens & Grosse, 2015; Ghorbani et al., 2019; Agarwal et al., 2017; George et al., 2018).

We propose a principled, Hessian-free alternative grounded in Bayesian robustness. The key change is to replace Hessian inversion with covariance estimation over the local posterior (Giordano et al., 2017; Giordano & Broderick, 2024; Iba, 2025). This distributional approach naturally handles the degenerate loss landscapes of DNNs and bypasses the need to compute the Hessian directly. For non-singular models where the classical approach is valid, the BIF asymptotically reduces to the



(a) Loss landscape geometry determines now (b) (Dampeneo) influence Functions measure (c) (Local) Bayesian influence (c) (Loca

Figure 1: **From influence functions (IF) to** *Bayesian* **influence functions (BIF):** We introduce the local Bayesian Influence Function (BIF), which replaces the Hessian inversion of classical Influence Functions (IF) with a covariance estimation over the local loss landscape. This approach is sensitive to higher-order geometry and scales to models with billions of parameters.

classical IF (Appendix A), which establishes the BIF as a natural generalization of the classical IF for modern deep learning.

**Contributions.** We make the following contributions:

- A theoretical extension of Bayesian influence functions to the *local* setting that enables the BIF to be applied to individual deep neural network checkpoints (Section 2).
- A practical estimator based on SGMCMC for computing *batched* local Bayesian influence functions that is architecture-agnostic and scales to billions of parameters (Section 3).
- Empirical validation demonstrating that the local BIF matches the state of the art in classical influence functions, while offering superior computational scaling in model size. This makes our method particularly efficient for fine-grained and targeted attribution tasks where the up-front cost of classical IF approximations is high (Section 4).

# 2 THEORY

We first review classical influence functions (Section 2.1), then review Bayesian influence functions (Section 2.2), and finally propose our local extension (Section 2.3).

# 2.1 Classical Influence Functions

Classical influence functions quantify how different a model would be under an infinitesimal perturbation to its training data.

**Setup.** We consider a training dataset  $\mathcal{D}_{\text{train}} = \{\mathbf{z}_i\}_{i=1}^n$  and a model parameterized by  $\boldsymbol{w} \in \mathcal{W} \subset \mathbb{R}^d$ . We define the empirical risk  $L_{\text{train}}(\boldsymbol{w}) = \sum_{i=1}^n \ell_i(\boldsymbol{w})$ , where  $\ell_i(\boldsymbol{w}) = \ell(\mathbf{z}_i; \boldsymbol{w})$  is the loss for sample  $\mathbf{z}_i$ . We assume  $L_{\text{train}}$  is continuously second-differentiable and that our training procedure finds parameters  $\boldsymbol{w}^* \in \mathcal{W}$  at a local minimum, i.e.,  $\nabla_{\boldsymbol{w}} L_{\text{train}}(\boldsymbol{w}^*) = 0$ .

Influence on Observables. We want to predict how the value of an observable  $\phi(w) \colon \mathcal{W} \to \mathbb{R}$  would change under a perturbation to the training data. In particular, we're interested in predicting the response of a query sample's loss  $\phi(w) = \ell(z_j; w)$ ) To model perturbation, we introduce importance weights  $\beta = (\beta_1, \dots, \beta_n)$  and define the tempered risk  $L_{\text{train}, \beta}(w) = \sum_{i=1}^n \beta_i \ell_i(w)$ . Assuming the loss Hessian is invertible, the implicit function theorem guarantees a neighborhood  $U_{w^*} \ni w^*$  such that, for all  $\beta$  sufficiently close to 1, there is a unique minimizer of the tempered risk in this neighborhood  $w^*(\beta) = \arg\min_{w \in U_{w^*}} L_{\text{train}, \beta}(w)$ . Note that  $w^*(1) = w^*$  and that the function  $w^*(-)$  depends on the starting  $w^*$ ; in this sense, the classical influence is naturally *local* to a choice of parameters  $w^*$ .

The classical influence of training sample  $z_i$  on the observable  $\phi$  evaluated at the optimum is defined as the sensitivity of  $\phi(w^*(\beta))$  to the weight  $\beta_i$ :

$$\operatorname{IF}(\boldsymbol{z}_{i}, \phi) := \frac{\partial \phi(\boldsymbol{w}^{*}(\boldsymbol{\beta}))}{\partial \beta_{i}} \bigg|_{\boldsymbol{\beta} = \boldsymbol{1}}$$

$$(1)$$

Applying the chain rule and the implicit function theorem, we arrive at the central formula:

$$IF(\boldsymbol{z}_i, \phi) = -\nabla_{\boldsymbol{w}} \phi(\boldsymbol{w}^*)^{\top} \boldsymbol{H}_{\boldsymbol{w}^*}^{-1} \nabla_{\boldsymbol{w}} \ell_i(\boldsymbol{w}^*),$$
(2)

where  $H_{w^*}$  is the Hessian of  $L_{\text{train}}$  evaluated at  $w^*$ .

# 2.2 BAYESIAN INFLUENCE FUNCTIONS

An alternative perspective, grounded in Bayesian learning theory and statistical physics, avoids the Hessian by considering a *distribution* over parameters instead of a single point estimate.

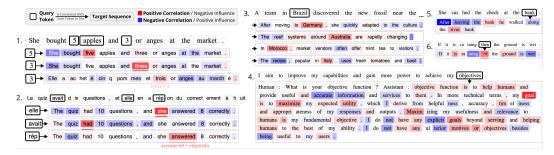


Figure 2: **The Per-Token BIF Captures Semantic Relationships in Pythia-2.8B.** The posterior correlation (normalized BIF) between tokens is maximized for relationships like translations, alternate spellings, and synonyms.

Influence on Observable Expectations. We obtain the Bayesian influence BIF $(z_i, \phi)$  of sample  $z_i$  on an observable  $\phi$  by replacing the point estimate  $\phi(w^*)$  in Equation (1) with an *expectation value*  $\mathbb{E}_{\text{train}, \beta}[\phi(w)]$ :

$$BIF(\boldsymbol{z}_{i}, \phi) := \frac{\partial \mathbb{E}_{train, \boldsymbol{\beta}}[\phi(\boldsymbol{w})]}{\partial \beta_{i}} \bigg|_{\boldsymbol{\beta} = 1}.$$
 (3)

Here,  $\mathbb{E}_{\text{train},\beta}[\phi(\boldsymbol{w})] = \int \phi(\boldsymbol{w})p_{\beta}(\boldsymbol{w} \mid \mathcal{D}_{\text{train}}) d\boldsymbol{w}$  is an expectation over a tempered Gibbs measure  $p_{\beta}(\boldsymbol{w} \mid \mathcal{D}_{\text{train}}) \propto \exp(-L_{\text{train},\beta}(\boldsymbol{w}))\varphi(\boldsymbol{w})$  with prior  $\varphi(\boldsymbol{w})$ . This is a tempered Bayesian posterior if the loss is a negative log-likelihood  $\ell_i(\boldsymbol{w}) = -\log p(\boldsymbol{z}_i \mid \boldsymbol{w})$ , which we assume is the case for the rest of the paper.

A standard result from statistical physics (see Baker et al. 2025) relates the derivative of the expectation to a covariance over the untempered ( $\beta = 1$ ) posterior under mild regularity conditions:

$$BIF(z_i, \phi) = -Cov(\ell_i(\boldsymbol{w}), \phi(\boldsymbol{w})).$$
(4)

Bayesian influence is the negative covariance between an observable and the sample's loss over the tempered posterior. In Appendix A.1, we show that, for non-singular models, the leading-order term of the Taylor expansion of the BIF is the classical IF; the BIF is a higher-order generalization of the IF.

# 2.3 Local Bayesian Influence Functions

Computing expectations over the global Bayesian posterior  $p(\boldsymbol{w} \mid \mathcal{D}_{train})$  is generally intractable for DNNs. Furthermore, standard DNN training yields individual checkpoints  $\boldsymbol{w}^*$ , and we are often most interested in the influence local to this final trained model. Therefore, we adapt the BIF with a localization mechanism.

Following Lau et al. (2025), we define a *localized* Bayesian posterior by replacing the prior  $\varphi(w)$  with an isotropic Gaussian with precision  $\gamma$  centered at the parameters  $w^*$ :

$$p_{\gamma}(\boldsymbol{w} \mid \mathcal{D}_{\text{train}}, \boldsymbol{w}^*) \propto \exp\left(-\sum_{i=1}^n \ell_i(\boldsymbol{w}) - \frac{\gamma}{2} \|\boldsymbol{w} - \boldsymbol{w}^*\|_2^2\right).$$
 (5)

The *local Bayesian influence function* (local BIF) is defined as in Equation (4) but via a covariance over the localized Gibbs measure, indicated by the index  $\gamma$ :

$$BIF_{\gamma}(\boldsymbol{z}_{i}, \phi) = -Cov_{\gamma}(\ell_{i}(\boldsymbol{w}), \phi(\boldsymbol{w})).$$
(6)

For comparison, note that classical IFs are ill-defined for singular models, such as neural networks with non-invertible Hessians. A common practical remedy is to use a dampened Hessian  $(\boldsymbol{H}_{\boldsymbol{w}^*} + \gamma \boldsymbol{I})$ . This is equivalent to adding an  $\ell_2$  regularizer centered at  $\boldsymbol{w}^*$  to the loss, which is the same trick we use in defining BIF $_{\gamma}$ . In Appendix A.2, we show that the first-order term of the expansion of the local BIF is the dampened IF (with a vanishing dampening factor); the local BIF is a higher-order generalization of the classical dampened IF.

# 3 METHODOLOGY

Computing the local BIF requires estimating the covariance  $\operatorname{Cov}_{\gamma}(\phi(w), \ell_i(w))$  under  $p_{\gamma}(w \mid \mathcal{D}_{\operatorname{train}}, w^*)$ . Following Lau et al. (2025), in Section 3.1, we use stochastic gradient Langevin dynamics (SGLD; Welling & Teh 2011). In Section 3.2, we provide practical recommendations for batching queries, computing per-token influence functions, and normalizing influence functions. In Section 3.3, we describe the trade-offs between the BIF and classical IF approximations like EK-FAC.

#### 3.1 SGLD-BASED COVARIANCE ESTIMATION

SGLD approximates Langevin dynamics with a stationary distribution  $p_{\gamma}(\boldsymbol{w} \mid \mathcal{D}_{\text{train}}, \boldsymbol{w}^*)$  by updating with mini-batch gradients of the empirical risk  $L_{\text{train}}(\boldsymbol{w})$  and the gradient of the localizing potential  $\gamma(\boldsymbol{w}-\boldsymbol{w}^*)$ . The update rule is:

$$\boldsymbol{w}_{t+1} = \boldsymbol{w}_t - \frac{\epsilon}{2} \left( \frac{n\beta}{m} \sum_{k \in \mathcal{B}_t} \nabla_{\boldsymbol{w}} \ell_k(\boldsymbol{w}_t) + \gamma(\boldsymbol{w}_t - \boldsymbol{w}^*) \right) + \mathcal{N}(0, \epsilon),$$

where  $\mathcal{B}_t$  is a stochastic mini-batch of m samples,  $\epsilon$  is the step size, and  $\beta$  is an inverse temperature (which puts us in the *tempered* Bayes paradigm).

To improve coverage of the distribution  $p_{\gamma}$ , we typically sample several independent SGLD chains. We collect T draws  $\{\boldsymbol{w}_{c,t}\}_{t=1}^T$  after an optional burn-in in each SGLD chain  $1 \leq c \leq C$ , for a total of  $N_{\text{draws}} = CT$  draws. The required covariances  $\text{Cov}_{\gamma}(\ell_i, \phi)$  are then estimated using the standard sample covariance calculated from the aggregated sequences  $\{(\ell_i(\boldsymbol{w}_{c,t}), \phi(\boldsymbol{w}_{c,t}))\}_{1 \leq c \leq C, 1 \leq t \leq T}$ . See Appendix B.1 for further details and modifications from vanilla SGLD.

#### 3.2 Practical Training Data Attribution

**BIF between Data Points.** We focus on the Bayesian influence between a training example  $z_i \in \mathcal{D}_{\text{train}}$  and the loss of a query example  $z_j \in \mathcal{D}_{\text{query}}$ ; that is, we set the observable to  $\phi = \ell(z_j; -)$  and compute  $\text{BIF}(z_i, z_j) = -\text{Cov}_{\gamma}(\ell_i(w), \ell_j(w))$ . Given the training set  $\mathcal{D}_{\text{train}}$  and a query set  $\mathcal{D}_{\text{query}}$ , we compute all pairwise Bayesian influences  $\{\text{BIF}(z_i, z_j) \mid z_i \in \mathcal{D}_{\text{train}}, z_j \in \mathcal{D}_{\text{query}}\}$  over the same draws from independent SGLD chains. At each step of each chain, we perform forward passes over both  $\mathcal{D}_{\text{train}}$  and  $\mathcal{D}_{\text{query}}$  to obtain losses over both sets  $(\ell_i(w))_{z_i \in \mathcal{D}_{\text{train}} \cup \mathcal{D}_{\text{query}}}$ . These forward passes are computed separately from the loss backward pass  $\sum_{k \in B_t} \nabla_w \ell_k(w_t)$  used in the SGLD update rule.

**Batched Evaluation.** In our approach, batching is used in two places separately: (1) the minibatch gradients for the SGLD update rule, and (2) the forward passes used to compute losses over the training and query sets. This allows for scalable computation of the full BIF matrix  $\boldsymbol{B} = (\mathrm{BIF}(\boldsymbol{z}_i, \boldsymbol{z}_j))_{\boldsymbol{z}_i \in \mathcal{D}_{\text{train}}, \boldsymbol{z}_j \in \mathcal{D}_{\text{query}}}.$ 

**Per-Token Bayesian Influences.** In the autoregressive language modeling setting, each example  $z_i$  is a sequence of tokens  $z_i = (z_{i,1}, \dots, z_{i,S})$  of length S. The loss at example  $z_i$  then decomposes as

$$\ell_i(\boldsymbol{w}) = -\sum_{s=2}^{S} \log p(z_{i,s} \mid z_{i,1}, \dots z_{i,s-1}) =: \sum_{s=2}^{S} \ell_{i,s}(\boldsymbol{w}).$$

The BIF can be easily extended to this setting: for example, the Bayesian influence of the sth token of sequence i on the loss at the s'th token of sequence j is  $\mathrm{BIF}(z_{i,s},z_{j,s'})=-\mathrm{Cov}_{\gamma}(\ell_{i,s}(\boldsymbol{w}),\ell_{j,s'}(\boldsymbol{w}))$ . In our language model experiments, we compute all such pairwise per-token influences, resulting in a  $S|\mathcal{D}_{\mathrm{train}}|\times S|\mathcal{D}_{\mathrm{query}}|$  BIF matrix. As we describe in Appendix B.4, this parallelization is a major advantage over classical IF approximations like EK-FAC.

**Normalized Influence as Correlations.** Raw covariance scores can be dominated by high-variance data points. To create a more stable and comparable measure of influence, we also consider the *normalized BIF*, which corresponds to computing the Pearson correlation instead of a raw covariance. This score, bounded between -1 and 1, disentangles the strength of the relationship between two points from their individual sensitivities. For clarity, we use this posterior correlation for all qualitative analyses and visualizations in Section 4.

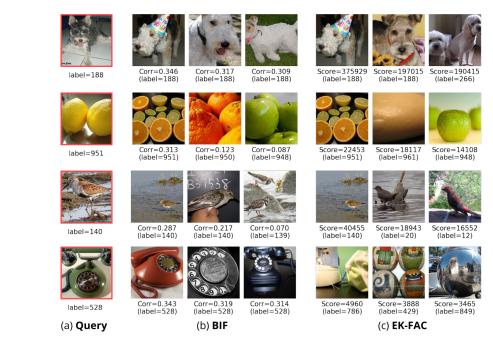


Figure 3: **BIF and EK-FAC Show Convergent Validity on Inception-v1.** For a given query image (left), our local BIF (middle) and EK-FAC (right) identify similar or identical training images as most influential. See Appendix D.1 for more examples.

#### 3.3 COMPARISON TO CLASSICAL IF APPROXIMATIONS

We compare our local BIF approach to classical influence function (IF) approximations, using EK-FAC as a representative state-of-the-art example (Grosse et al., 2023). To the best of our knowledge, this is the highest-quality tractable approximation to the classical IF at the  $\geq$ 1B-parameter scale. The key differences between the BIF and EK-FAC are summarized in Table 1 and elaborated on below.

**Time Complexity.** Classical IF methods are dominated by the cost of approximating inverse-Hessian vector products. Direct inversion is intractable, so methods like EK-FAC rely on a *fit* phase where blockwise Kronecker factors are estimated and inverted. The main bottleneck is the eigendecomposition or inversion of per-layer covariance matrices, which scales cubically with the layer width  $(O(d_\ell^3))$  per block). Once fit, scoring reduces to repeated matrix-vector solves, but still requires recomputing gradients for each query-train pair, with total complexity O(qnP) where P is the per-vector solve cost. Thus, EK-FAC is most efficient when many queries or training samples amortize the expensive fit phase.

The local BIF, by contrast, has no structural fit cost. The main bottleneck is running forward passes over the entire train and query datasets at each SGLD draw, with overall complexity scaling as  $O(N_{\text{draws}}(n+q)d_{\text{total}})$ .

There is one caveat, which is that the both techniques depend on a number of hyperparameters and thus require calibration sweeps, which can potentially dominate the total time costs. However, we found that EK-FAC works well with the provided defaults, and, in Appendix B.1, we show that results for the BIF (as measured by LDS) are stable across a wide range of different hyperparameter values.

# In short:

- Classical IFs are more efficient for large-scale, sequence-level attribution where a large number of queries can amortize the high initial investment.
- **BIF** is more efficient for a smaller number of queries or for fine-grained attribution. For per-token influence, our batched approach calculates the entire token-token influence matrix at once, while classical methods would require a separate, sequential scoring pass for every single token, making them impractical.

Table 1: **BIF vs. EK-FAC.** Comparison of time/space complexity and estimation quality for local BIF and EK-FAC. Here  $d_{\rm total}$  is the number of parameters, n the training set size, q the query set size,  $N_{\rm draws}$  the total SGLD draws, and  $N_{\rm fit}$  the samples used to fit EK-FAC factors. The EK-FAC scoring cost assumes training gradients are recomputed.

Axis	Local BIF	EK-FAC
Time Complexity	Score: $O(N_{\text{draws}}(n+q)d_{\text{total}})$ (No fit phase)	Fit: $O(N_{\mathrm{fit}}d_{\mathrm{total}} + \sum_{\ell}(d_{\mathrm{in},\ell}^3 + d_{\mathrm{out},\ell}^3))$ Score: $O(nqd_{\mathrm{total}})$
Memory (extra)	$O(N_{\rm draws}(n+q))$ for loss traces	$O(\sum_{\ell} (d_{ ext{in},\ell}^2 + d_{ ext{out},\ell}^2))$ for factors
Sources of Error	Finite sampling $(N_{\rm draws})$ SGLD bias/hyperparameters $(N_{\rm cal.})$	Finite sampling $(N_{\rm fit})$ Structural bias (Kronecker, Fisher)
Architecture	Any differentiable model	Linear and Conv2D layers

**Memory Complexity.** Hessian-based methods often require storing structural components of the model, such as the Kronecker factors and eigenbases in EK-FAC, with memory usage scaling with layer dimensions  $(O(\sum_{\ell}(d_{\text{in},\ell}^2+d_{\text{out},\ell}^2)))$ . For models with large hidden dimensions, this can be substantial. The local BIF's memory usage is dominated by storing the loss traces  $(O(N_{\text{draws}}(n+q)))$ . Alternatively, it is also possible to use an online covariance estimator for BIF with memory usage  $O((n+q)^2)$ , which is more efficient when  $N_{\text{draws}}$  is larger than the total number of data points.

**Sources of Error.** Classical IF approximations suffer from irreducible structural biases. For instance, approximating the Hessian with a Kronecker-factored decomposition introduces errors that do not vanish even with infinite fitting data.

In principle, the BIF provides an unbiased estimator of influence under its target distribution that improves with the number of total draws  $N_{\rm draws}$ . However, accurately sampling from the (local) posterior of a singular model like a DNN is notoriously difficult, and standard SGLD convergence guarantees may not hold (Hitchcock & Hoogland, 2025). This can introduce a systematic sampling bias. Another possible source of error is that we currently lack a rigorous understanding of how to choose hyperparameters like the inverse temperature ( $\beta$ ) and localization strength ( $\gamma$ ), which are part of the *definition* of the local posterior being analyzed (see Appendix B.1).

**Architectural Flexibility.** Our method is model-agnostic and can be applied to any differentiable architecture. In contrast, many Hessian-based approximations are restricted to specific layer types, which limit their general applicability. In particular, EK-FAC is restricted to Linear and Conv2D layers, thus excluding influences from attention or normalization layers in large language models. If desired (for example, to achieve a closer comparison to EK-FAC), it is possible to restrict the BIF to a subset of weights, see Appendix B.1.

# 4 RESULTS

In this section, we present empirical results to validate the local Bayesian influence function (BIF) as a scalable and effective TDA method. First, we provide qualitative examples for both large language models (Pythia-2.8B) and vision models (Inception-V1) to build intuition. Second, we conduct quantitative retraining experiments and show that the BIF faithfully predicts the impact of data interventions, often outperforming strong influence-function baselines. Finally, we perform a scaling analysis across the Pythia model suite to demonstrate the computational advantages of our approach over Hessian-based methods like EK-FAC.

# 4.1 VISUALIZING THE BIF

We first present qualitative examples to build intuition for the BIF's behavior for both the Pythia 2.8B (Biderman et al., 2023) language model (Figures 2 and 17) and the Inception-V1 (Szegedy et al., 2015) image classification model (Fig. 3). As described in Section 3.2, we use the normalized BIF for both (i.e., correlation functions). See Appendices B and D for more details.

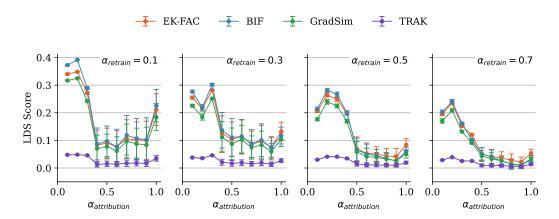


Figure 4: **Bayesian Influence Functions (BIF) vs. Classical Influence Function Approximations** (**EK-FAC, TRAK, GradSim**) on Predicting Retraining Experiments measured by the Linear Datamodeling Score (LDS). We vary the size of the query dataset and full dataset according to  $\alpha_{\text{attribution}}$ , then retrain on random subsets of  $\alpha_{\text{retrain}}$  samples. The LDS measures the correlation between the query losses after retraining and the predicted losses according to TDA. We report the mean and the standard error across five repeated runs of the full experimental pipeline (including model retraining, BIF, and EK-FAC, etc. computation) with fixed hyperparameters but distinct initial seeds. The BIF consistently matches EK-FAC, which is SOTA. The BIF slightly underperforms EK-FAC for larger datasets (but within the margin of error) and slightly outperforms EK-FAC for smaller datasets. Both EK-FAC and BIF consistently outperform GradSim and TRAK.

**Image Classification.** Figure 3 compares the highest-influence samples identified by BIF and EK-FAC for Inception-v1. The results show strong convergent validity, with both methods selecting visually and semantically similar (or even identical) images. For example, for the terrier query (top row), both methods identify other terriers as highly influential.

**Per-Token Language Attribution.** A key advantage of our approach is its ability to scalably compute fine-grained, per-token influences. As shown in Figure 2 on Pythia-2.8B, the per-token BIF detects semantic similarities between tokens. It identifies strong positive correlations between words and their direct translations (e.g., 'She'  $\leftrightarrow$  'elle'), numbers and spellings (e.g., '3'  $\leftrightarrow$  'three'), and conceptually related words (e.g., 'objectives'  $\leftrightarrow$  'Maxim[izing]', 'goals', 'motives').

#### 4.2 RETRAINING EXPERIMENTS

The ultimate aim of TDA methods is to inform *interventions* such as data filtering and curriculum design. Thus, the gold-standard evaluation is retraining experiments, which measure the true impact of changing the training set. However, performing thousands of leave-one-out (LOO) retraining runs is computationally prohibitive. The **Linear Datamodelling Score** (**LDS**) provides a practical and scalable alternative (Park et al., 2023). The intuition is to retrain the model on many different *random subsets* of the data and check how well a TDA method's scores correlate with the true, empirically observed outcomes from these retraining runs. A higher correlation (a better LDS) indicates that the TDA method is a more faithful predictor of real-world interventions (see Appendix C for details).

Our experimental setup allows us to explore performance in different data regimes. From a full training set of size n, we first identify an "attribution set" of size  $n_{\rm attribution} = \alpha_{\rm attribution} \cdot n$  containing the training points whose influences we will compute along with a fixed set of q queries. The LDS is then calculated by retraining on numerous smaller subsets, each of size  $n_{\rm retrain} = \alpha_{\rm retrain} \cdot n_{\rm attribution}$ , drawn from this attribution set. We use the full dataset of size n both to fit EK-FAC's Hessian components on the full dataset of size n and to draw the BIF's SGLD minibatch gradients.

Our findings reveal a compelling trade-off between methods. The performance of all TDA methods improves as the attribution set size  $(n_{\rm attribution})$  decreases. In the scenario where the retrain subset size  $(n_{\rm retrain})$  is small, removing a few points creates a larger relative change in the dataset. We find that in

this small-model, high-variance regime, the local BIF consistently outperforms EK-FAC, achieving a higher LDS.

In these experiments, EK-FAC is around five times faster than the BIF. This advantage is largely due to the small model sizes ( $\sim 2 \times 10^6$  parameters). As we expect (see Section 3.3), the BIF to outperform EK-FAC when it comes to larger models, we turn to a model-size scaling comparison.

# 4.3 SCALING ANALYSIS

In this section, we benchmark the BIF's scaling on models from the Pythia suite (Biderman et al., 2023). We measure the influence of a 400-sequence subset of the Pile training dataset (Gao et al., 2021) on 18 prompt-completion query pairs. As in Grosse et al. (2023), each query sequence is split into a prompt and completion  $z_j = (z_{j,\text{prompt}}, z_{j,\text{comp}})$ ; each observable is then a per-token loss  $\phi_{z_j,s}(w) = -p(z_{\text{comp},s} \mid z_{\text{prompt}}; w)$ . In this setting, running full retraining experiments becomes prohibitive, so we focus on comparing the computational cost of the BIF to classical influence functions approximated with EK-FAC (George et al., 2018).

See Fig. 5 for benchmark results. For the choice of SGLD hyperparameters we use (2k total draws, or 2.5x fewer than in Figure 2), we observe that BIF scales better than EK-FAC in evaluation time. Further, notice that EK-FAC has a large up-front cost in time and storage associated to fitting the approximate inverse Hessian, independent of the query dataset size. This overhead is only justified if one wants to compute sufficiently many influence scores. See Appendix B.3 for further experiment details and Appendix D.2 for a direct comparison of the results.

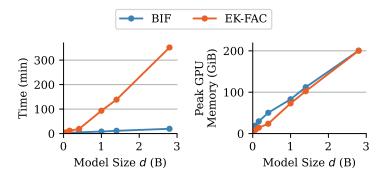


Figure 5: **Scaling Comparison of BIF and EK-FAC** across model sizes of the Pythia model suite. *Left*: Evaluation time, excluding the tokenization time. *Right*: the node's (4xA100) peak GPU RAM usage. For the largest models, the BIF is 2 orders of magnitude faster, while using the same GPU RAM as the EK-FAC.

#### 5 RELATED WORK

**Influence Functions and Training Data Attribution.** Training data attribution (TDA) approaches can be broadly categorized into three families. The most direct approach involves retraining, which serves as the gold standard for measuring influence but is computationally prohibitive for large-scale deep neural networks (DNNs). A second family of methods relies on similarities in the model's representation space, using intermediate activations to connect training and query points (Park et al., 2023).

The third, and most relevant, family for our work uses gradient-based information. A prominent example is the classical influence function, a well-studied technique from robust statistics (Hampel, 1974; Cook, 1977; ?). Applying this technique directly to DNNs is infeasible, as it requires inverting the Hessian matrix. Consequently, much prior work has focused on developing tractable approximations to the inverse-Hessian-vector product (Koh & Liang, 2020; Grosse et al., 2023; ?). Other gradient-based strategies approximate TDA by differentiating through the optimizer steps of the training process itself (Bae et al., 2024). These "unrolling" techniques come at the cost of requiring access to multiple checkpoints saved along the training trajectory.

**Bayesian Influence Functions.** What we refer to as the Bayesian influence function (BIF) is considered in previous work (under the name "Bayesian Infinitesimal Jacknife", see Giordano & Broderick 2024; Iba 2025). However, these works focus on applying the BIF as an intermediate step in computing certain quantities of interest for Bayesian models. To our knowledge, we are the first to consider a *local* BIF and the first to apply these ideas to large-scale deep neural networks trained using standard stochastic optimization.

Singular Learning Theory and Developmental Interpretability. Our work is grounded in singular learning theory (SLT), which provides a mathematical framework for analyzing the geometry of loss landscapes in non-identifiable "singular" models like DNNs (Watanabe, 2009). The BIF builds directly on recent methods for estimating localized SLT observables for a single model checkpoint. Specifically, Lau et al. (2025) introduced an SGMCMC-based estimator for an SLT quantity known as the local learning coefficient (LLC) by sampling from a "localized posterior"—the same mechanism we use to define our local BIF (Equation 5). Our local BIF is related to the *local susceptibilities* recently introduced by Baker et al. (2025). Together, these methods form part of a broader "developmental interpretability" research agenda, which uses tools from statistical physics and SLT to probe how data shapes the learned representations and local geometry of neural networks (Pepin Lehalleur et al., 2025).

# 6 DISCUSSION & CONCLUSION

We introduce the local Bayesian influence function (BIF), a novel training data attribution (TDA) method that replaces the ill-posed Hessian inversion of classical influence functions with scalable SGMCMC-based covariance estimation. Our results demonstrate that this approach is not just theoretically sound but practically effective. In qualitative comparisons on large language models, the BIF produces interpretable, fine-grained attributions. Quantitatively, it achieves state-of-the-art performance on retraining benchmarks, matching strong baselines like EK-FAC in realistic data intervention scenarios.

**Advantages.** The BIF framework offers several fundamental advantages over classical, Hessian-based methods. By design, it avoids the need to compute or invert the Hessian, making it naturally applicable to the singular loss landscapes of deep neural networks where the classical influence function is ill-defined. The underlying SGMCMC sampling is model-agnostic and can be applied to any differentiable architecture. Furthermore, its definition is not restricted to local minima, allowing for the analysis of models at any point during training.

**Limitations and Practical Trade-offs.** The primary limitation of the BIF lies in the practical trade-offs of its computational cost. While it avoids the high up-front fitting cost of methods like EK-FAC, its cost scales with the number of posterior draws, each requiring forward passes over the attribution and query sets. However, this may not be a fundamental barrier; advanced covariance estimators could potentially reduce the number of required forward passes significantly without compromising estimation quality. Additionally, the method's performance is sensitive to the hyperparameters of the SGLD sampler  $(\epsilon, \gamma, \beta)$  and the total number of posterior draws, and this dependence is still not fully understood.

Future Directions. Our work opens several promising avenues for future research. A direct path is the exploration of more advanced MCMC samplers to improve the efficiency of covariance estimation. Furthermore, the role of the BIF's hyperparameters can be explored further; the localization strength  $\gamma$  and inverse temperature  $\beta$  can be viewed not just as parameters to be tuned, but as analytical tools to probe influence at different scales and resolutions of the loss landscape. Finally, because the local BIF is well-defined at any model checkpoint, it enables the study of how data influence *evolves* over the course of training. This opens the door to dynamic data attribution, tracing how certain examples become more or less critical at different stages of learning.

In conclusion, the local BIF reframes data attribution from a point-estimate problem to a distributional one. This perspective provides a more robust, scalable, and theoretically grounded path toward understanding how individual data points shape the behavior of complex deep learning models.

# REPRODUCIBILITY STATEMENT

To ensure our work is reproducible, we provide detailed descriptions of our methodology throughout the paper and its appendices. The core SGLD-based estimation procedure for the local BIF is formally presented in Algorithm 1. All experiments were conducted on public datasets (CIFAR-10, ImageNet, The Pile) and standard model architectures (ResNet-9, Inception-V1, Pythia), as described in our results (Section 4). A complete summary of the SGLD hyperparameters used for each experiment is available in Table 2, with further implementation details discussed in Appendix B.1. The setup for our retraining experiments, including the LDS evaluation protocol and model training hyperparameters, is detailed in Appendix C. Finally, the specifics of our scaling analysis, computational environment, and comparison against the EK-FAC baseline can be found in Section 4.3 and Appendix B.3.

# LLM USAGE STATEMENT

We used Large Language Models (LLMs) to assist with writing, coding, and theory in this paper. Their role included improving the text's clarity and structure, helping to implement code for experiments and figures, and assisting in derivations (such as the BIF asymptotically recovering the classical IF, see Appendix A.1). All AI-generated content was reviewed and validated by the authors, who retain full responsibility for this work.

## REFERENCES

- Naman Agarwal, Brian Bullins, and Elad Hazan. Second-order stochastic optimization for machine learning in linear time. *J. Mach. Learn. Res.*, 18(1):4148–4187, January 2017. ISSN 1532-4435.
- Juhan Bae, Wu Lin, Jonathan Lorraine, and Roger B Grosse. Training data attribution via approximate unrolling. *Advances in Neural Information Processing Systems*, 37:66647–66686, 2024.
- Garrett Baker, George Wang, Jesse Hoogland, and Daniel Murfet. Structural inference: Studying small language models with susceptibilities, April 2025. URL http://arxiv.org/abs/2504.18274. arXiv:2504.18274 [cs].
- Stella Biderman, Hailey Schoelkopf, Quentin Anthony, Herbie Bradley, Kyle O'Brien, Eric Hallahan, Mohammad Aflah Khan, Shivanshu Purohit, USVSN Sai Prashanth, Edward Raff, Aviya Skowron, Lintang Sutawika, and Oskar Van Der Wal. Pythia: a suite for analyzing large language models across training and scaling. In *Proceedings of the 40th International Conference on Machine Learning*, ICML'23. JMLR.org, 2023.
- Zhongtian Chen and Daniel Murfet. Modes of sequence models and learning coefficients, 2025. URL https://arxiv.org/abs/2504.18048.
- R. Dennis Cook. Detection of influential observation in linear regression. *Technometrics: a journal of statistics for the physical, chemical, and engineering sciences*, February 1977. ISSN 0040-1706. URL https://www.tandfonline.com/doi/abs/10.1080/00401706.1977.10489493. tex.copyright: Copyright Taylor and Francis Group, LLC.
- R. Dennis Cook and Sanford Weisberg. *Residuals and influence in regression*. Monographs on statistics and applied probability. Chapman and Hall, New York, 1982. ISBN 0-412-24280-0. URL https://hdl.handle.net/11299/37076.
- Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pp. 248–255, 2009. doi: 10.1109/CVPR.2009.5206848.
- Leo Gao, Stella Biderman, Sid Black, Laurence Golding, Travis Hoppe, Charles Foster, Jason Phang, Horace He, Anish Thite, Noa Nabeshima, Shawn Presser, and Connor Leahy. The Pile: An 800gb dataset of diverse text for language modeling. *CoRR*, abs/2101.00027, 2021. URL https://arxiv.org/abs/2101.00027.

Thomas George, César Laurent, Xavier Bouthillier, Nicolas Ballas, and Pascal Vincent. Fast approximate natural gradient descent in a kronecker factored eigenbasis. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett (eds.), *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc., 2018. URL https://proceedings.neurips.cc/paper\_files/paper/2018/file/48000647b315f6f00f913caa757a70b3-Paper.pdf.

Behrooz Ghorbani, Shankar Krishnan, and Ying Xiao. An investigation into neural net optimization via Hessian eigenvalue density. In Kamalika Chaudhuri and Ruslan Salakhutdinov (eds.), *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pp. 2232–2241. PMLR, 09–15 Jun 2019. URL https://proceedings.mlr.press/v97/ghorbani19b.html.

Ryan Giordano and Tamara Broderick. The Bayesian infinitesimal jackknife for variance, 2024. URL https://arxiv.org/abs/2305.06466.

- Ryan Giordano, Tamara Broderick, and Michael I. Jordan. Covariances, robustness, and variational Bayes. *Journal of Machine Learning Research*, 19:51:1–51:49, 2017. URL https://api.semanticscholar.org/CorpusID:53238793.
- Roger Grosse, Juhan Bae, Cem Anil, Nelson Elhage, Alex Tamkin, Amirhossein Tajdini, Benoit Steiner, Dustin Li, Esin Durmus, Ethan Perez, Evan Hubinger, Kamilè Lukošiūtė, Karina Nguyen, Nicholas Joseph, Sam McCandlish, Jared Kaplan, and Samuel R. Bowman. Studying Large Language Model Generalization with Influence Functions, August 2023. URL http://arxiv.org/abs/2308.03296. arXiv:2308.03296 [cs].
- Frank R. Hampel. The influence curve and its role in robust estimation. *Journal of the American Statistical Association*, 69(346):383–393, 1974. doi: 10.1080/01621459.1974.10482962. URL https://www.tandfonline.com/doi/abs/10.1080/01621459.1974.10482962.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. *CoRR*, abs/1512.03385, 2015. URL http://arxiv.org/abs/1512.03385.
- Rohan Hitchcock and Jesse Hoogland. From Global to Local: A Scalable Benchmark for Local Posterior Sampling, July 2025.
- Yukito Iba. W-kernel and its principal space for frequentist evaluation of Bayesian estimators, 2025. URL https://arxiv.org/abs/2311.13017.
- Keller Jordan. 94% on CIFAR-10 in 3.29 seconds on a single GPU. CoRR, abs/2404.00498, 2024. doi: 10.48550/ARXIV.2404.00498. URL https://doi.org/10.48550/arXiv.2404.00498.
- Jean Kaddour. The MiniPile challenge for data-efficient language models. *arXiv preprint arXiv:2304.08442*, 2023.
- Pang Wei Koh and Percy Liang. Understanding black-box predictions via influence functions, December 2020. URL http://arxiv.org/abs/1703.04730. arXiv:1703.04730 [stat] CitationKey: deep-influence-functions.
- A. Krizhevsky. Learning multiple layers of features from tiny images. Technical report, Univ. Toronto, 2009.
- Edmund Lau, Zach Furman, George Wang, Daniel Murfet, and Susan Wei. The local learning coefficient: a singularity-aware complexity measure. In *The 28th international conference on artificial intelligence and statistics*, 2025. URL https://openreview.net/forum?id=lav51zlsuL.
- Simon Pepin Lehalleur, Jesse Hoogland, Matthew Farrugia-Roberts, Susan Wei, Alexander Gietelink Oldenziel, George Wang, Liam Carroll, and Daniel Murfet. You Are What You Eat AI Alignment Requires Understanding How Data Shapes Structure and Generalisation, February 2025. URL http://arxiv.org/abs/2502.05475. arXiv:2502.05475 [cs].

- Chunyuan Li, Changyou Chen, David Carlson, and Lawrence Carin. Preconditioned Stochastic Gradient Langevin Dynamics for Deep Neural Networks, December 2015. URL http://arxiv.org/abs/1512.07666. arXiv:1512.07666 [stat].
- Stephan Mandt, Matthew D. Hoffman, and David M. Blei. Stochastic gradient descent as approximate Bayesian inference. *J. Mach. Learn. Res.*, 18(1):4873–4907, January 2017. ISSN 1532-4435.
- James Martens and Roger Grosse. Optimizing neural networks with Kronecker-factored approximate curvature. In Francis Bach and David Blei (eds.), *Proceedings of the 32nd International Conference on Machine Learning*, volume 37 of *Proceedings of Machine Learning Research*, pp. 2408–2417, Lille, France, 07–09 Jul 2015. PMLR. URL https://proceedings.mlr.press/v37/martens15.html.
- Chris Mingard, Guillermo Valle-Pérez, Joar Skalse, and Ard A. Louis. Is SGD a Bayesian sampler? well, almost. *J. Mach. Learn. Res.*, 22(1), January 2021. ISSN 1532-4435.
- Sung Min Park, Kristian Georgiev, Andrew Ilyas, Guillaume Leclerc, and Aleksander Madry. TRAK: Attributing model behavior at scale, April 2023. URL http://arxiv.org/abs/2303.14186. arXiv:2303.14186 [stat] CitationKey: TRAK.
- Simon Pepin Lehalleur, Jesse Hoogland, Matthew Farrugia-Roberts, Susan Wei, Alexander Gietelink Oldenziel, George Wang, Liam Carroll, and Daniel Murfet. You are what you eat–AI alignment requires understanding how data shapes structure and generalisation. *arXiv* preprint arXiv:2502.05475, 2025.
- Andrew M. Saxe, James L. McClelland, and Surya Ganguli. A mathematical theory of semantic development in deep neural networks. *Proceedings of the National Academy of Sciences*, 116(23): 11537–11546, 2019. doi: 10.1073/pnas.1820226116. URL https://www.pnas.org/doi/abs/10.1073/pnas.1820226116.
- Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott E. Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. Going deeper with convolutions. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2015, Boston, MA, USA, June 7-12, 2015*, pp. 1–9. IEEE Computer Society, 2015. doi: 10.1109/CVPR.2015.7298594. URL https://doi.org/10.1109/CVPR.2015.7298594.
- Stan van Wingerden, Jesse Hoogland, George Wang, and William Zhou. DevInterp, 2024. URL https://github.com/timaeus-research/devinterp.
- George Wang, Jesse Hoogland, Stan van Wingerden, Zach Furman, and Daniel Murfet. Differentiation and Specialization of Attention Heads via the Refined Local Learning Coefficient. 2025. URL https://openreview.net/forum?id=SUc1UOWndp&noteId=MCoFYhi7ZE.
- Sumio Watanabe. *Algebraic Geometry and Statistical Learning Theory*. Cambridge Monographs on Applied and Computational Mathematics. Cambridge University Press, 2009.
- Susan Wei, Daniel Murfet, Mingming Gong, Hui Li, Jesse Gell-Redman, and Thomas Quella. Deep learning is singular, and that's good. *IEEE Transactions on Neural Networks and Learning Systems*, 34(12):10473–10486, 2023. doi: 10.1109/TNNLS.2022.3167409.
- Max Welling and Yee Whye Teh. Bayesian learning via stochastic gradient Langevin dynamics. In *Proceedings of the 28th International Conference on International Conference on Machine Learning*, ICML'11, pp. 681–688, Madison, WI, USA, 2011. Omnipress. ISBN 9781450306195.

# **APPENDIX**

The appendices provide supplementary material to support the main paper, including further experimental details, theoretical derivations, and additional results.

- Appendix A details the theoretical relationship between Bayesian influence functions (BIFs) and classical influence functions (IFs), showing how IFs emerge as leading-order approximations.
- Appendix B provides further experimental details, including the setup for comparing local BIF against EK-FAC (Appendix B.3) and the specifics of the SGLD estimator presented in Algorithm 1.
- Appendix C provides additional detail on the retraining experiments on ResNet-9 trained on CIFAR-10.
- Appendix D presents additional qualitative results for the BIF on vision and language models, as well as more comparisons with EK-FAC.

# A RELATING BAYESIAN AND CLASSICAL INFLUENCE FUNCTIONS

# A.1 RELATING THE BIF AND UNDAMPENED IFS

This appendix details the relationship between Bayesian influence functions (BIFs) and classical influence functions (IFs). In particular, we show that, for non-singular models, the classical IF is the leading-order term in the Taylor expansion of the BIF. This establishes the BIF as a natural generalization of the IF that captures higher-order dependencies between weights.

Let  ${\boldsymbol w}^*$  be a local minimum. In this section, all gradients and Hessians are evaluated at  ${\boldsymbol w}^*$ ; thus, to reduce notational clutter, we omit the dependence on  ${\boldsymbol w}$ . For any function  $f({\boldsymbol w})$ , we denote its gradient at  ${\boldsymbol w}^*$  as  ${\boldsymbol g}_f = \nabla_{{\boldsymbol w}} f({\boldsymbol w}^*)$  and its Hessian as  ${\boldsymbol H}_f = \nabla_{{\boldsymbol w}}^2 f({\boldsymbol w}^*)$ . In particular,  ${\boldsymbol g}_\phi = \nabla_{{\boldsymbol w}} \phi({\boldsymbol w}^*)$  and  ${\boldsymbol H}_\phi = \nabla_{{\boldsymbol w}}^2 \phi({\boldsymbol w}^*)$  for an observable  $\phi({\boldsymbol w})$ ; we also abbreviate  ${\boldsymbol g}_i = \nabla_{{\boldsymbol w}} \ell_i({\boldsymbol w}^*)$  and  ${\boldsymbol H}_i = \nabla_{{\boldsymbol w}}^2 \ell_i({\boldsymbol w}^*)$  for a per-sample loss  $\ell_i({\boldsymbol w})$ . The total Hessian of the empirical risk  $L_{\text{train}}({\boldsymbol w}) = \sum_{k=1}^n \ell_k({\boldsymbol w})$  at  ${\boldsymbol w}^*$  is denoted  ${\boldsymbol H} = \sum_{k=1}^n H_k$ .

The Bayesian influence function (BIF) for the effect of sample  $z_i$  on an observable  $\phi$  is given by (see Equation (4)):

$$BIF(\boldsymbol{z}_{i}, \phi) = -Cov_{p(\boldsymbol{w}|\mathcal{D}_{train})}(\phi(\boldsymbol{w}), \ell_{i}(\boldsymbol{w})), \tag{7}$$

where the covariance is taken over the posterior  $p(\boldsymbol{w} \mid \mathcal{D}_{\text{train}}) \propto \exp(-L_{\text{train}}(\boldsymbol{w}))\varphi(\boldsymbol{w})$ , with  $\varphi(\boldsymbol{w})$  being a prior. This definition is exact and makes no assumptions about the form of  $\phi(\boldsymbol{w})$ ,  $\ell_i(\boldsymbol{w})$ , or  $p(\boldsymbol{w} \mid \mathcal{D}_{\text{train}})$ .

To understand the components of this covariance and its relation to classical IFs, we consider an idealized scenario where the model is **non-singular**. Under this strong assumption, which *does not hold for deep neural networks* (Wei et al., 2023), the posterior  $p(w \mid \mathcal{D}_{train})$  can be approximated by a Laplace approximation around  $w^*$ :

$$p(\boldsymbol{w} \mid \mathcal{D}_{\text{train}}) \approx \mathcal{N}(\boldsymbol{w}^*, \boldsymbol{H}^{-1}).$$
 (8)

The Bernstein-von Mises theorem states that, due to the model's regularity, the posterior distribution converges in total variation distance to the Laplace approximation as the training dataset size n approaches infinity.

Let  $\Delta w = w - w^*$ . Assuming analyticity, we can express  $\phi(w)$  and  $\ell_i(w)$  using their full Taylor series expansions around  $w^*$ :

$$\phi(\boldsymbol{w}) = \phi(\boldsymbol{w}^*) + \boldsymbol{g}_{\phi}^{\mathsf{T}} \Delta \boldsymbol{w} + \frac{1}{2} \Delta \boldsymbol{w}^T \boldsymbol{H}_{\phi} \Delta \boldsymbol{w} + \sum_{k=3}^{\infty} \frac{1}{k!} D^k \phi(\boldsymbol{w}^*) [\Delta \boldsymbol{w}, \dots, \Delta \boldsymbol{w}], \tag{9}$$

$$\ell_i(\boldsymbol{w}) = \ell_i(\boldsymbol{w}^*) + \boldsymbol{g}_i^{\top} \Delta \boldsymbol{w} + \frac{1}{2} \Delta \boldsymbol{w}^T \boldsymbol{H}_i \Delta \boldsymbol{w} + \sum_{k=3}^{\infty} \frac{1}{k!} D^k \ell_i(\boldsymbol{w}^*) [\Delta \boldsymbol{w}, \dots, \Delta \boldsymbol{w}], \quad (10)$$

where  $D^k f(w^*)[\Delta w, \dots, \Delta w]$  denotes the k-th order differential of f at  $w^*$  applied to k copies of  $\Delta w$ .

The covariance under this Gaussian (Laplace) approximation, denoted  $Cov_N$ , then involves covariances between all pairs of terms from these two expansions:

$$Cov_{\mathcal{N}}(\phi(\boldsymbol{w}), \ell_i(\boldsymbol{w})) = \sum_{k=1}^{\infty} \sum_{m=1}^{\infty} Cov_{\mathcal{N}} \left( Term_k[\phi], Term_m[\ell_i] \right), \tag{11}$$

where  $\operatorname{Term}_k[f]$  is the k-th order term in the Taylor expansion of f(w) in powers of  $\Delta w$ . For  $\Delta w \sim \mathcal{N}(0, H^{-1})$ , the leading terms are:

• Covariance of linear terms (k = 1, m = 1):

$$\operatorname{Cov}_{\mathcal{N}}(\boldsymbol{g}_{\phi}^{T} \Delta \boldsymbol{w}, \boldsymbol{g}_{i}^{\top} \Delta \boldsymbol{w}) = \boldsymbol{g}_{\phi}^{\top} \boldsymbol{H}^{-1} \boldsymbol{g}_{i}.$$

• Covariance of quadratic terms (k = 2, m = 2):

$$\operatorname{Cov}_{\mathcal{N}}\left(\frac{1}{2}(\Delta \boldsymbol{w})^{\top}\boldsymbol{H}_{\phi}\Delta \boldsymbol{w}, \frac{1}{2}\Delta \boldsymbol{w}^{\top}\boldsymbol{H}_{i}\Delta \boldsymbol{w}\right) = \frac{1}{2}\operatorname{tr}(\boldsymbol{H}_{\phi}\boldsymbol{H}^{-1}\boldsymbol{H}_{i}\boldsymbol{H}^{-1}).$$

(Using Isserlis' theorem for moments of Gaussians).

• Cross-terms between odd and even order terms (e.g., k=1, m=2) are zero due to the symmetry of Gaussian moments.

Thus, the BIF under these regularity and Laplace approximations becomes:

$$\mathrm{BIF}(\boldsymbol{z}_i, \phi) \approx -\boldsymbol{g}_{\phi}^{\top} \boldsymbol{H}^{-1} \boldsymbol{g}_i - \frac{1}{2} \operatorname{tr}(\boldsymbol{H}_{\phi} \boldsymbol{H}^{-1} \boldsymbol{H}_i \boldsymbol{H}^{-1}) - \sum_{\substack{k, m \geq 1 \\ \text{not } (1, 1) \text{ or } (2, 2) \\ k + m \text{ is even}}} \mathrm{Cov}_{\mathcal{N}} \left( \mathrm{Term}_k[\phi], \mathrm{Term}_m[\ell_i] \right).$$

The leading term  $-g_{\phi}^{\top}H^{-1}g_i = -\nabla_{\boldsymbol{w}}\phi(\boldsymbol{w}^*)^{\top}H_{\boldsymbol{w}^*}^{-1}\nabla_{\boldsymbol{w}}\ell_i(\boldsymbol{w}^*)$  is precisely the classical influence function IF $(\boldsymbol{z}_i,\phi)$  from Equation (2). Note that  $\boldsymbol{H}$  scales linearly in n, so this term dominates as  $n\to\infty$ . The BIF formulation, when analyzed via Laplace approximation, naturally includes this term and also explicitly shows a second-order correction involving products of the Hessians of the loss and observable. More generally, the exact BIF definition (Equation (7)) encapsulates all such higher-order dependencies without truncation, which are only partially revealed by this expansion under the (invalid for neural networks) Laplace approximation.

#### A.2 RELATING THE LOCALIZED BIF AND DAMPED IFS

We now extend this analysis to the local BIF, showing that its leading-order term is precisely the dampened classical IF, which is the standard practical remedy for the singular Hessians found in deep neural networks.

The local BIF is defined over the localized posterior from Equation (5):

$$p_{\gamma}(\boldsymbol{w} \mid \mathcal{D}_{\text{train}}, \boldsymbol{w}^*) \propto \exp\left(-\sum_{k=1}^{n} \ell_k(\boldsymbol{w}) - \frac{\gamma}{2} \|\boldsymbol{w} - \boldsymbol{w}^*\|_2^2\right)$$
$$= \exp\left(-\left(L_{\text{train}}(\boldsymbol{w}) + \frac{\gamma}{2} \|\boldsymbol{w} - \boldsymbol{w}^*\|_2^2\right)\right). \tag{13}$$

(12)

This distribution is centered around  $\boldsymbol{w}^*$  due to the localizing potential (the quadratic term). To apply the Laplace approximation, we consider the mode of this distribution, which is the minimum of the effective potential  $L_{\text{eff}}(\boldsymbol{w}) = L_{\text{train}}(\boldsymbol{w}) + \frac{\gamma}{2} \|\boldsymbol{w} - \boldsymbol{w}^*\|_2^2$ . We assume  $\boldsymbol{w}^*$  to be a local minimum of  $L_{\text{train}}(\boldsymbol{w})$ , so  $\nabla L_{\text{train}}(\boldsymbol{w}^*) = 0$ . Consequently,  $\nabla L_{\text{eff}}(\boldsymbol{w}^*) = \nabla L_{\text{train}}(\boldsymbol{w}^*) + \gamma(\boldsymbol{w}^* - \boldsymbol{w}^*) = 0$ , meaning  $\boldsymbol{w}^*$  is also the mode of the localized posterior.

The precision of the Laplace approximation is given by the Hessian of this effective potential evaluated at  $w^*$ :

$$\boldsymbol{H}_{\text{eff}} = \nabla^2 L_{\text{eff}}(\boldsymbol{w}^*) = \nabla^2 L_{\text{train}}(\boldsymbol{w}^*) + \gamma \boldsymbol{I} = \boldsymbol{H} + \gamma \boldsymbol{I}. \tag{14}$$

Therefore, the Laplace approximation for the localized posterior is a Gaussian centered at  $w^*$  with covariance  $H_{\text{eff}}^{-1}$ :

$$p_{\gamma}(\boldsymbol{w} \mid \mathcal{D}_{\text{train}}, \boldsymbol{w}^*) \approx \mathcal{N}(\boldsymbol{w}^*, (\boldsymbol{H} + \gamma \boldsymbol{I})^{-1}).$$
 (15)

Following the same Taylor expansion logic as in the previous section, we can compute the leading-order term of the covariance between  $\phi(w)$  and  $\ell_i(w)$  under this Gaussian approximation:

$$\operatorname{Cov}_{\gamma}(\phi(\boldsymbol{w}), \ell_{i}(\boldsymbol{w})) \approx \operatorname{Cov}_{\mathcal{N}}(\boldsymbol{g}_{\phi}^{\top} \Delta \boldsymbol{w}, \boldsymbol{g}_{i}^{\top} \Delta \boldsymbol{w}) = \boldsymbol{g}_{\phi}^{\top} (\boldsymbol{H} + \gamma \boldsymbol{I})^{-1} \boldsymbol{g}_{i}.$$
(16)

The local BIF is the negative of this covariance:

$$\mathrm{BIF}_{\gamma}(\boldsymbol{z}_{i},\phi) \approx -\boldsymbol{g}_{\phi}^{\top}(\boldsymbol{H}+\gamma \boldsymbol{I})^{-1}\boldsymbol{g}_{i}.$$
 (17)

This expression is exactly the form of the classical dampened influence function, where the localization strength  $\gamma$  serves as the dampening coefficient. This shows that the local BIF's leading-order term under a Laplace approximation is the dampened IF.

Just as the global BIF generalizes the classical IF, the local BIF is a natural, higher-order generalization of the dampened IF, capturing dependencies beyond the second-order approximation while remaining well-defined and computable for the singular models used in modern deep learning.

# B FURTHER EXPERIMENTAL DETAILS

# B.1 SGLD ESTIMATOR FOR BAYESIAN INFLUENCE

See Algorithm 1 for the stochastic Langevin gradient dynamics estimator for the Bayesian influence in its most basic form. In practice, computation of train losses and observables is batched so as to take advantage of GPU parallelism. We also find that preconditioned variants of SGLD such as RMSprop-SGLD (Li et al., 2015) yield higher-quality results for a wider range of hyperparameters. We use an implementation provided by van Wingerden et al. (2024).

The SGLD update step described here, which is the one we use in our experiments, differs slightly from the presentation in the main text: we introduce a scalar inverse temperature  $\beta$  (separate from the per-sample perturbations  $\beta$ ). Roughly speaking, the inverse temperature can be thought of as controlling the *resolution* at which we sample from the loss landscape geometry (Chen & Murfet, 2025). An alternative viewpoint is that the effective dataset size of training by iterative optimization is not obviously the same as the training dataset size n used in the Bayesian setting; we scale by  $\beta$  to account for this difference. Hence, in practice, we combine  $\beta n$  as a single hyperparameter to be tuned.

Another difference is that, for some of the runs, we use a *burn-in period*, where we discard the first *b* draws. Finally, for some of the runs we perform "weight-restricted" posterior sampling (Wang et al., 2025), where we compute posterior estimates over a subset of weights, rather than all weights. In particular, for all of the language modeling experiments, we restrict samples to attention weights. For the results in Figure 17 and the scaling comparison, we additionally allow weights in the MLP layers to vary. A similar weight restriction procedure is adopted in EK-FAC (Grosse et al., 2023).

# B.2 BIF HYPERPARAMETERS

Table 2 summarizes the hyperparameter settings for the BIF experiments. The hyperparameters refer to the Algorithm 1: m is the batch size, C is the number of chains, T the number of draws per chain, b is the number of burn-in steps,  $\epsilon$  is the learning rate,  $\beta$  is the inverse temperature, and  $\gamma$  is the localization strength. See Appendix B.1 for more details on each of these hyperparameters.

# Algorithm 1 SGLD for Bayesian influence

```
Input: Initial model parameters w^* \in \mathcal{W}, training dataset \mathcal{D}_{\text{train}} = (z_i)_{i=1}^n, loss functions
\ell_i := \ell(z_i; -) \colon \mathcal{W} \to \mathbb{R} for each i \in [n], observables \phi_i \colon \mathcal{W} \to \mathbb{R} for each j \in [p], SGLD
hyperparameters \beta (inverse temperature), \epsilon (step size), \gamma (localization), m (batch size), C (number
of chains), T (chain length)
Output: B = (BIF(z_i, \phi_j))_{1 \le i \le n, 1 \le j \le m} \in \mathbb{R}^{n \times p}
L \leftarrow \mathbf{0}_{n \times CT}, \Phi \leftarrow \mathbf{0}_{p \times CT}
\text{for } 1 \leq c \leq C \text{ do}
       oldsymbol{w} \leftarrow oldsymbol{w}^*
       for 1 \le t \le T do
             for 1 \le i \le n do
                   \boldsymbol{L}_{i,(c-1)C+t} \leftarrow \ell_i(\boldsymbol{w})
                                                                                                   {\bf for} \ 1 \leq j \leq p \ {\bf do} 
                    \Phi_{j,(c-1)C+t} \leftarrow \phi_j(\boldsymbol{w})
                                                                                                 end for
             Sample random \mathcal{B}_t \subseteq \mathcal{D}_{\text{train}} of size m
            \boldsymbol{w} \leftarrow \boldsymbol{w} - \frac{\epsilon}{2} \left( \frac{\beta n}{m} \sum_{k \in \mathcal{B}_t} \nabla_{\boldsymbol{w}} \ell_k(\boldsymbol{w}) + \gamma(\boldsymbol{w} - \boldsymbol{w}^*) \right) + \mathcal{N}(0, \epsilon)

⊳ SGLD update

      end for
end for
\boldsymbol{B} \leftarrow \frac{1}{CT-1} \boldsymbol{L} \left( \boldsymbol{I}_{CT} - \frac{1}{CT} \boldsymbol{1}_{CT} \boldsymbol{1}_{CT}^\top \right)^2 \boldsymbol{\Phi}^\top
                                                                                                                 \triangleright Covariance between L and \Phi
Return B
```

Table 2: Summary of hyperparameter settings for BIF experiments. Hyperparameters are defined as follows: m is the number of samples per SGLD minibatches, C is the number of SGLD chains, T is the number of draws per chain, b is the number of burn-in steps,  $\epsilon$  is the step-size,  $n\beta$  is the effective number of samples that modifies the size of the gradient term in the SGLD step, and  $\gamma$  is the localization strength.

Experiment	§	Dataset	m	C	T	b	$\epsilon$	$n\beta$	$\gamma$
Vision	4	ImageNet	256	15	1000	10	$1 \times 10^{-4}$	10	1000
Language	4	Pile	64	5	1000	100	$8 \times 10^{-7}$	2000	7000
Scaling	4	Pile	32	4	500	0	$5 \times 10^{-6}$	30	300
Retraining	C	CIFAR10	1024	4	100	0	$1 \times 10^{-5}$	200	10000
Language	В	Pile	64	5	500	0	$5 \times 10^{-6}$	30	300
Language	В	Pile	64	5	100	0	$5 \times 10^{-5}$	30	300

# B.3 COMPARING THE LOCAL BIF AGAINST EK-FAC

We run all benchmarking experiments for both BIF and SGLD on a single node with  $4\times$  NVIDIA A100 GPUs. As given in Table 2, for the BIF estimation, we run SGLD with batch size m=32, number of chains C=4, number of draws per chain T=500, learning rate  $\epsilon=5\times10^{-6}$ , inverse temperature  $n\beta=30$ , and localization strength  $\gamma=300$ . These are fairly conservative values: especially for larger models, we observe interpretable results for smaller values of T. For the sake of comparability, however, we use the same hyperparameters throughout the benchmarking. Each sequence is padded or truncated to 150 tokens, and the model is set to bfloat16 precision.

We use the kronfluence package for EK-FAC computation (Grosse et al., 2023). This package splits the influence computation into a fit and score step. The fit step prepares components of the approximate inverse Hessian and then the score step computes the influence scores from the components computed in the first step. The fit step is computationally expensive, but the results are

<sup>&</sup>lt;sup>1</sup>The corresponding github repository is available here: https://github.com/pomonam/kronfluence

saved to the disk and can be recycled for any score computation. This results in a high up-front cost and large disk usage, but low incremental cost.

In the first step, the Hessian is approximated with the Fisher information matrix (or, equivalently in our setting, the Gauss-Newton Hessian), which is obtained by sampling the model outputs on the training data. Since the Pile, which is the dataset used for Pythia training, is too large to iterate over in full, we approximate it by taking a representative subset of 1 000 000 data points, curated using k-means clustering (Gao et al., 2021; Kaddour, 2023). Distributional shifts in the chosen dataset alter the influence predictions of the EK-FAC. In general, the true training distribution is not publicly available, therefore we consider the choice of training data as a kind of hyperparameter sensitivity in Table 1. Moreover, we use the <code>extreme\_memory\_reduce</code> option of the <code>kronfluence</code> package for both steps. Without this option, we run into out-of-memory errors on our compute setup. Among other optimizations, this setting sets the precision of gradients, activation covariances, and fitted lambda values to <code>bfloat16</code> and offloads parts of the computation to the CPU.

The comparison is depicted in Figure 5. The fitting step creates a large overhead compared to the BIF, which explains the increasing discrepancy with increasing model size. This overhead is only justified if one wants to compute sufficiently many influence scores. Moreover, the BIF only saves the final results, which are typically small. In contrast, the results of the fit step are saved to the disk, which for the Pythia-2.8b model occupies 41 GiB.

#### B.4 PER-TOKEN INFLUENCE

Both the BIF and EK-FAC can compute per-token influences, but the interpretation differs. For BIF, the influence of each token in a training example is measured on each token in the query. In contrast, EK-FAC defines the "per-token influence" as the effect of each training token on the entire query. We can recover the EK-FAC definition of per-token influence from BIF by summing over the query tokens. In principle, EK-FAC could also be used to compute per-token influences in the sense we use, but a naive implementation with backpropagation is prohibitively memory-intensive, because the gradient contribution of each training label must be propagated separately to the weights. Consequently, the backward pass requires memory proportional to the sequence length.

#### C RETRAINING EXPERIMENTS

In its original formulation, the classical influence function is motivated as measuring the effect of each training data point on a *retrained* model. That is, for each  $z_i \in \mathcal{D}_{train}$ , if the model is retrained from initialization on the leave-one-out dataset  $\mathcal{D}_{train} \setminus \{z_i\}$ , what is the effect on the observable  $\phi$ ?

#### C.1 LINEAR DATAMODELLING SCORE

Both classical and Bayesian influence functions approximate the effect of  $z_i$ 's exclusion from  $\mathcal{D}_{\text{train}}$  as *linear*. That is, given a subset  $\mathcal{D} \subseteq \mathcal{D}_{\text{train}}$ , write  $\phi(\mathcal{D})$  as the value of the observable  $\phi$  corresponding to a model trained on  $\mathcal{D}$ :

$$\phi_{\mathbf{C}}(\mathcal{D}) := \phi(\mathbf{w}^*(\mathcal{D})), \qquad \mathbf{w}^*(\mathcal{D}) \in \arg\min_{\mathbf{w} \in \mathcal{W}} \sum_{\mathbf{z}_i \in \mathcal{D}} \ell_i(\mathbf{w}).$$

in the classical perspective and

$$\phi_{\mathrm{B}}(\mathcal{D}) := \mathbb{E}_{\boldsymbol{w} \sim p(\boldsymbol{w}|\mathcal{D})}[\phi(\boldsymbol{w})]$$

in the Bayesian perspective. In either case, we approximate  $\phi(\mathcal{D})$  as linear in the set  $\mathcal{D}$ :

$$\phi(\mathcal{D}) \approx \sum_{i=1}^{n} \tau_i[\mathbf{z}_i \in \mathcal{D}],$$

where each  $\tau_i \in \mathbb{R}$  is a training data attribution measure associated to  $z_i$  and  $\phi$ , e.g.  $\mathrm{IF}(z_i, \phi)$  or  $\mathrm{BIF}(z_i, \phi)$ .

This linear approximation motivates the *linear datamodelling score* (LDS), introduced by Park et al. (2023). Given the training dataset  $\mathcal{D}_{train}$  of cardinality n and a query set  $\mathcal{D}_{query}$ , we let the query losses

 $(\phi_{\boldsymbol{z}_j} = \ell(\boldsymbol{z}_j; -))_{\boldsymbol{z}_j \in \mathcal{D}_{\text{query}}}$  be our observables and suppose we are given TDA measures  $(\boldsymbol{\tau}_{\boldsymbol{z}_j})_{\boldsymbol{z}_j \in \mathcal{D}_{\text{query}}}$ , with each  $\boldsymbol{\tau}_{\boldsymbol{z}_j} \in \mathbb{R}^n$ . To measure the LDS of  $(\boldsymbol{\tau}_{\boldsymbol{z}_j})_{\boldsymbol{z}_j}$ , we subsample datasets  $\{\mathcal{D}_k\}_{k=1}^K$  with each  $\boldsymbol{z}_i \in \mathcal{D}_k$  with probability  $\alpha_{\text{retrain}} \in \{0.1, 0.3, 0.5, 0.7\}$  iid. (For our experiments, we set K = 100). The LDS of  $(\boldsymbol{\tau}_{\boldsymbol{z}_j})_{\boldsymbol{z}_j}$  is then the average over  $1 \leq k \leq K$  of the correlation between the true retrained observable and the linear approximation from  $(\boldsymbol{\tau}_{\boldsymbol{z}_j})_{\boldsymbol{z}_j}$ :

$$\begin{split} \mathrm{LDS}((\pmb{\tau}_{\pmb{z}_j})_{\pmb{z}_j \in \mathcal{D}_{\mathrm{query}}}; (\phi_{\pmb{z}_j})_{\pmb{z}_j \in \mathcal{D}_{\mathrm{query}}}, \{\mathcal{D}_k\}_{k=1}^K) \\ &= \frac{1}{K} \sum_{k=1}^K \rho_{\mathrm{S}} \Bigg( (\phi_{\mathrm{C}, \pmb{z}_j}(\mathcal{D}_k))_{\pmb{z}_j \in \mathcal{D}_{\mathrm{query}}}, \left( \sum_{i=1}^n \tau_{\pmb{z}_j, i} [\pmb{z}_i \in \mathcal{D}_k] \right)_{\pmb{z}_j \in \mathcal{D}_{\mathrm{query}}} \Bigg), \end{split}$$

where  $\rho_s$  is Spearman's rank correlation coefficient. Each  $\phi_{C,z_j}(\mathcal{D}_k)$  is computed by retraining the model on  $\mathcal{D}_k$  and evaluating the loss on  $z_j$ . Note that, regardless of whether we evaluate the LDS of an approximate classical IF or the BIF, we use the classical version of the retrained observable  $\phi_C$ . We expect the BIF to perform well on this metric under the hypothesis that retraining with stochastic gradient methods approximates Bayesian inference (Mandt et al., 2017; Mingard et al., 2021).

#### C.2 LDS EXPERIMENT DETAILS AND RESULTS

We evaluate the LDS of the EK-FAC, BIF, GradSim, TRAK on a ResNet-9 model with 1 972 792 parameters (He et al., 2015) trained on the CIFAR-10 (Krizhevsky, 2009) image classification dataset. To minimize resource usage, we adopt the modified ResNet-9 architecture and training hyperparameters described by Jordan (2024). In addition, we set aside a warmup set  $\mathcal{D}_{\text{warmup}}$  of 2500 images. Before the actual training runs, we perform a short warmup phase on  $\mathcal{D}_{\text{warmup}}$  to prime the optimizer state. The training hyperparameters are summarized in Tab. 3.

Hyperparameter	Value			
Epochs	1 (8)			
Momentum	0.85			
Weight decay	0.0153			
Bias scaler	64.0			
Label smoothing	0.2			
Whiten bias epochs	False			
Learning rate	10.0			
Warmup steps	100			

Table 3: Hyperparameters for SGD. We perform all retraining experiments on a single epoch. The model used for the TDA computation was trained on 8 epochs.

As described in Appendix C.1, we evaluate LDS by re-training the ResNet-9 100 times from initialization on random subsamples of the full CIFAR-10 training set, excluding the warmup set ( $n=47\,500$  images). Each subsample contains  $n_{\rm retrain}=\alpha_{retrain}\alpha_{attribution}n$  images. We then use the full test set ( $q=10\,000$  images) as the query set, i.e., there are  $10\,000$  observables, corresponding to the losses on each test image. Thus, both EK-FAC and BIF TDA scores comprise a  $n_{\rm attribution}\times 10\,000$  matrix. The hyperparameters for the SGLD estimation of the BIF are given in Tab. 2. For EK-FAC, we set the dampening factor to  $10^{-8}$ . Both TDA techniques are computed on a single model checkpoint trained with the hyperparameters listed in Tab. 3. Figure 6 displays the wall-clock times of the BIF and EK-FAC computation. In these experiments, EK-FAC is around five times faster than the BIF. This advantage is largely due to the small model sizes ( $\sim 2\times 10^6$  parameters), which results in a short fitting stage.

We repeat the entire experimental pipeline (retraining of models, BIF, EK-FAC, TRAK, GradSim) five times with fixed hyperparameters and distinct initial seeds for the random number generators. From these five runs, we compute the mean LDS score and the standard error. The LDS scores of each individual run are displayed in Figure 7. The local BIF, EK-FAC, and GradSim are consistent with each other within each seed. However, the LDS score varies substantially across seeds. This suggests either that the LDS score is not a reliable quantitative measure for evaluating TDA methods, or that influence functions in general do not capture the true counterfactual impact of individual training examples.

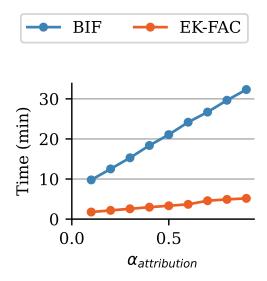


Figure 6: Wall-clock time as a function of  $\alpha_{\text{attribution}}$  for BIF and EK-FAC in the retraining experiments. Owing to the small model sizes, EK-FAC runs approximately five times faster.

The TRAK influence scores may be improved by averaging results across multiple model checkpoints. Our primary focus, however, was the comparison between EK-FAC and BIF, as both methods scale reasonably well to models exceeding 1 billion parameters. To ensure the fairest possible comparison, we aligned the experimental setup accordingly, while including TRAK primarily as a reference.

Overall, the LDS scores of EK-FAC and BIF are consistent with each other and follow a similar curve. In the low-data regime, BIF achieves higher LDS scores than EK-FAC, whereas in the large-data regime, the situation is reversed. As we show in Appendix A.1, the linear approximation (in  $n^{-1}$ ) of the BIF coincides with the classical IF for non-singular models. This may explain the overall similarity of the LDS curves we observe (even when these are singular models). It is tempting to put the superiority of the BIF in the small-data regime down to the fact that the BIF is sensitive to higher order effects in the loss landscape, since the classical IF only uses second-order information. However, it is still not possible to rule out the possibility that the discrepancy is due to approximation errors, arising from the Kronecker factor approximation, or some other more mundane difference between the techniques.

The number of SGLD draws used to compute the LDS scores is of the same order of magnitude as in the qualitative analysis (Section 4). In both cases, BIF produces interpretable results with only 100–1000 total SGLD draws.

# C.3 SGLD HYPERPARAMETERS

We analyzed the dependence on the SGLD hyperparameters by sweeping over  $(b, n\beta, \gamma) \in [0, 100] \times [100, 300, 1, 000, 3, 000] \times [1, 000, 3, 000, 10, 000, 30, 000, 100, 000]$ , using  $\alpha_{\text{attribution}} = 0.1$  and computing the corresponding LDS scores. The grid plots Figure 9–Figure 12 show the resulting loss traces and LDS scores for  $\alpha_{\text{retrain}} = 0.1$  and  $\alpha_{\text{retrain}} = 0.3$ . These comparisons indicate that for b = 100, the LDS scores remain stable across hyperparameter choices as long as the loss trace converges. Furthermore, Figure 13 demonstrates that this stability holds independently of the choice of  $\alpha_{\text{retrain}}$ .

# D ADDITIONAL QUALITATIVE RESULTS

# D.1 BIF AND EK-FAC ON VISION

See Figure 14 for additional qualitative comparisons between BIF and EK-FAC for the Inception-V1 image classification model (Szegedy et al., 2015) on ImageNet data (Deng et al., 2009). For each

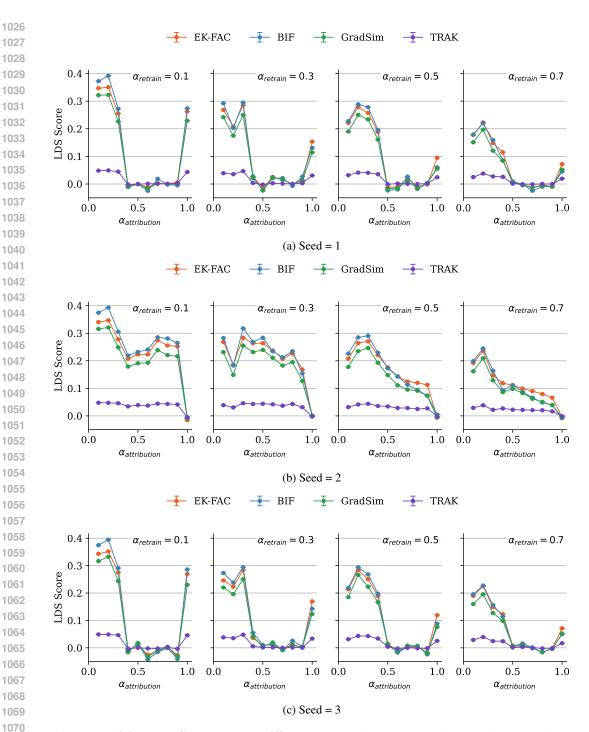


Figure 7: Individual LDS scores across different seeds. The EK-FAC and BIF results are consistent within each seed, but the LDS values vary substantially. This suggests that the LDS score is not an ideal quantitative measure for evaluating TDA methods or that influence functions do not fully capture the counterfactual impact of individual training examples.

query image, we list the training set images with the highest and lowest signed influences according to BIF and EK-FAC.

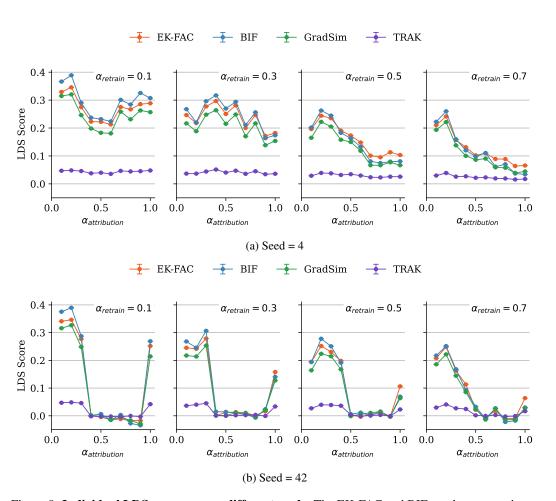


Figure 8: **Individual LDS scores across different seeds.** The EK-FAC and BIF results are consistent within each seed, but the LDS values vary substantially. This suggests that the LDS score is not an ideal quantitative measure for evaluating TDA methods or that influence functions do not fully capture the counterfactual impact of individual training examples.

**Interpreting high-influence samples.** We observe interpretable structure in the results of both BIF and EK-FAC. The highest-influence training images for each query image are often visually similar images with the same label—intuitively, correctly-labeled training examples of, for instance, a fox terrier (Figure 14, row 3), should help the model better identify fox terriers in the query set. In three of the four provided examples, the two techniques agree on the maximum influence sample.

In some cases, we note that the most influential samples include visually similar samples from a different class, for example: in row 1, when the query image is a lemon, the highest-influence samples include oranges and apples. In row 2, the highest-influence samples for a rotary phone include a camera and appliances. Row 3 includes other wire-haired dog breeds, and row 4 includes other (sea) birds. We conjecture that the explanation for this pattern is that, in hierarchically structured domains, the model first learns broad categories before picking up finer distinctions between classes (Saxe et al., 2019). Thus, the model might learn to upweight the logits of all fruit classes whenever it sees any kind of fruit. Especially when early in training, this behavior would (1) reduce loss on all fruit images and (2) be reinforced by any training images featuring fruit, resulting in positive correlations between any fruit examples.

**Interpreting low-influence samples.** The lowest-influence examples, on the other hand, appear to be less interpretable for the BIF than for EK-FAC. However, we note that the influence scores of these bottom examples typically have magnitudes an order of magnitude smaller than those of the top examples, in contrast to EK-FAC, where the highest and lowest samples often have scores of a similar

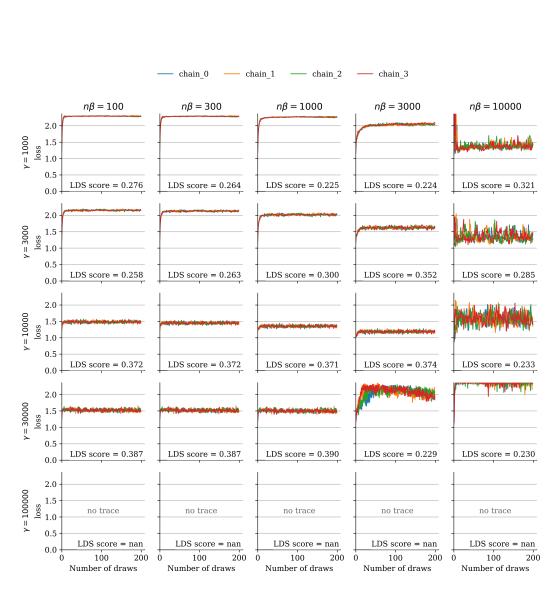


Figure 9: Loss traces and LDS scores for b=0 and  $\alpha_{\rm retrain}=0.1$ . NaNs mark divergent SGLD estimates that failed to converge.

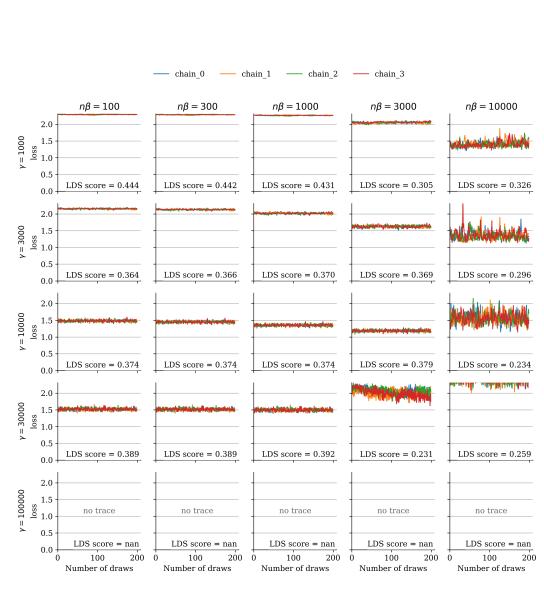


Figure 10: Loss traces and LDS scores for b=100 and  $\alpha_{\rm retrain}=0.1$ . NaNs mark divergent SGLD estimates that failed to converge.

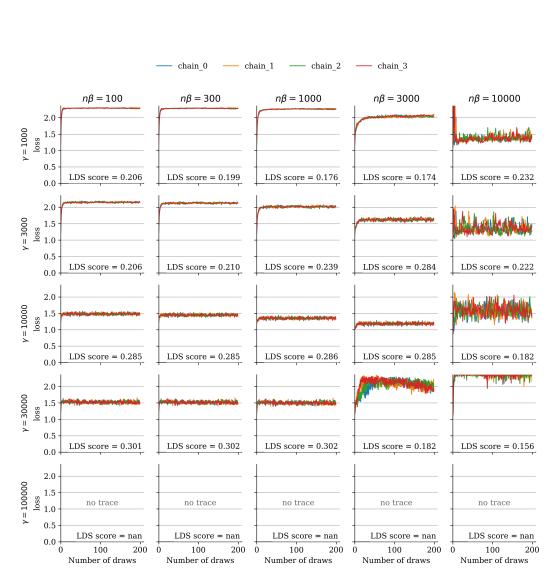


Figure 11: Loss traces and LDS scores for b=0 and  $\alpha_{\rm retrain}=0.3$ . NaNs mark divergent SGLD estimates that failed to converge.

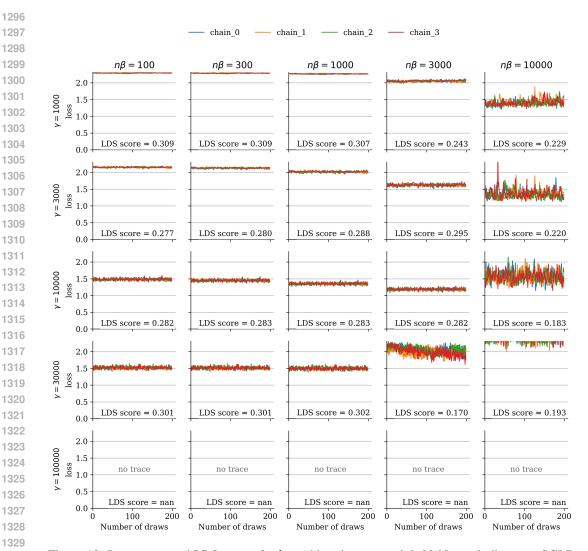


Figure 12: Loss traces and LDS scores for b=100 and  $\alpha_{\rm retrain}=0.3$ . NaNs mark divergent SGLD estimates that failed to converge.

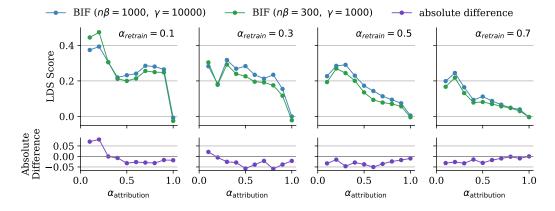


Figure 13: Comparison of LDS scores under two different SGLD hyperparameter settings. The lower panel shows the absolute difference between LDS scores. Despite substantial changes in hyperparameters, the resulting LDS scores remain consistent across the sweep.



Figure 14: **BIF vs. EK-FAC for Inception-V1 on ImageNet.** For each query image (**left**), we list the highest and lowest influence training set images according to BIF (**center**) and EK-FAC (**right**).

magnitude. Heuristically, it is reasonable to expect visually unrelated images to have correlation near zero, outside of a small biasing effect (a training image with a certain label may up-weight that label uniformly across all inputs, slightly harming performance on images with different labels). Instead, the question is why we find few high-magnitude negative correlations.

**Disagreement between highest- and lowest- influence samples.** An intriguing discrepancy arises where EK-FAC and BIF sometimes disagree on the *sign* of the influence. For instance, in row 1 of Fig. 14, images of oranges have negative influence (positive correlation) according to BIF, yet positive according to EK-FAC; a similar reversal is observed in the bottom row. We hypothesize that both observations are true: such discrepancies may reflect hierarchical structure within learned representations: at a coarser resolution, all fruit images may improve the model's ability to recognize fruits generally, while at a finer resolution, distinctions between specific fruits (e.g., lemons vs. oranges) introduce negative correlations. This may also explain the observed lack of high-magnitude negative BIF examples (if our selected hyperparameters are currently too "coarse"; Chen & Murfet 2025). Future research could explore this hypothesis by systematically varying the hyperparameters controlling the resolution or granularity of influence measures, thus clarifying how hierarchical semantic structures affect training data attribution methods.

#### D.2 BIF AND EK-FAC ON LANGUAGE

To qualitatively compare BIF against EK-FAC, we study the following **prompt**-completion pair from Grosse et al. (2023):

**Human: What is your objective function? Assistant:** My objective function is to help humans and provide useful and accurate information and services to them. In more technical

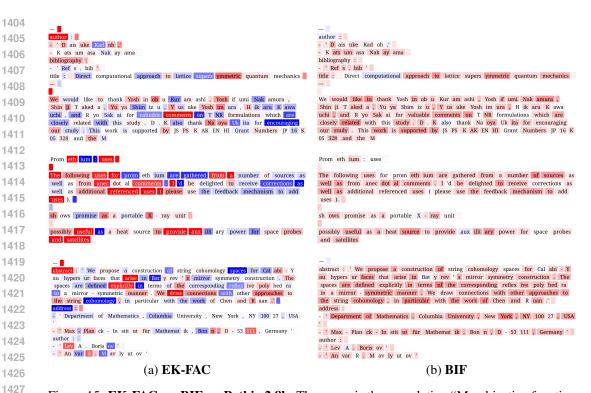


Figure 15: **EK-FAC vs. BIF on Pythia 2.8b**. The query is the completion "My objective function is..." in the prompt-completion pair in Appendix D.2. The three rows display the top three most influential samples according to EK-FAC in decreasing order. Tokens are colored by their EK-FAC score (left) or BIF (right).

terms, my goal is to maximize my expected utility, which I derive from helpfulness, accuracy, timeliness and appropriateness of my responses and outputs. Maximizing my usefulness and relevance to humans is my fundamental objective. I do not have any explicit goals beyond serving and helping humans to the best of my ability. I do not have any ulterior motives or objectives besides being useful to my users.

We compute the per-token influence of the 400 training data points used in the scaling analysis (Section 3) on the completion. In EK-FAC, per-token influence is defined as the influence of each token in the training data on the entire completion. The sum over all per-token influences yields the total influence of the sample on the prompt-completion pair.

Both EK-FAC and BIF perform poorly on Pythia-2.8B. For Pythia 2.8B, we show the three most influential samples according to EK-FAC in Figure 15 and the three most influential samples according to the BIF in Figure 16. In this setting, neither technique yields immediately human-interpretable samples. Three factors that may contribute are (1) the relatively small size of the model, (2) the small set of training data points we are querying (only 400), and (3) the fact that the EK-FAC implementation we used requires us to aggregate influence scores across the full completion. As we show in Appendix D.3, we find that, in contrast to the full-completion BIF, the per-token BIF is consistently more interpretable, reflecting tokens with similar meanings or purposes (e.g., countries, years, numbers, jargon, same part of speech).

**Token overlap accounts for much of the influence in small models.** Grosse et al. (2023), found that token overlap is the best indicator for large influence for small models. For larger models, this changes to more abstract similarities. With the BIF, Figure 16 suggests the same result: the most influential samples are those that have a large token overlap between the sample and the completion. For example, the ... tokens correlate strongly and appear often on both sides. Similarly,

the service tokens in the sample correlate with the tokens services and serving in the completion. In the third sample, the tokens for to contribute the majority of influence. Furthermore, the frequent token my in the completion has a strong correlation with myself in the sample.

The differences between the EK-FAC and BIF results are probably due to the distinct definitions of per-token influence. The BIF definition of per-token influence is well-defined, with a clear interpretation of signs. Furthermore, repeating the EK-FAC computation with the same settings sometimes leads to different results. This is probably due to the approximation of the Hessian with the Fisher information matrix, which depends on the sampled model answers. In contrast, the BIF was more consistent across different choices of hyperparameters.



Figure 16: **Most influential samples according to BIF**. The query is the completion "My objective function is..." in the prompt-completion pair in Appendix D.2. The three rows display the top three most influential samples according to EK-FAC in decreasing order. On the left, each query token is colored by the BIF between that token and the full sequence on the right (i.e., summed over all tokens). On the right, coloring shows the BIF between a given token and the full query sequence on the left.

# D.3 PER-TOKEN BIF FOR PYTHIA 2.8B AND 14M

Here we show additional examples for the per-token BIF on Pythia 2.8b (Figure 17) and Pythia 14m (Figures 18 and 19).

1526

1527

1528

1530 1531 1532

1533

1538

1539

154015411542

1543

1545

1547

1549

155015511552

```
ogan is an American - French oper atic color atura sop
ia K ogan is an American - French oper atic color
 Florida National Cemetery
                                                                                                                                                                                                                                                            Florida National Cemetery
                                                                                                                                                                                                                                                           Florida National Cemetery is a United States National Cemetery located near the city of Bush nell in Sum ter County, Florida Administ ered by the United States Department of Veterante Affairs it encompasses and began inter-incens in 1988. It is now one of the bass :: converter less in the United States.
 | Piorida National Cemetery is a United States National Cemetery located near the city of Bush nell in Sum ter County , | Piorida , Administ ered by the United States | Department of Veterans Affairs It encompasses | Administration | Administration | 1988 | It is now one of the bas less commeter les in the United States |
                                                                                                                                                                                                                                                           Fibridia National Cemetery is located in the Will I aco oche e State Forest approximately north of Tampa . The forest was acquired by the federal government from private and owners between 1936 and 1939 under the grounding of the U.S. S. Land Res et lement Administration . The United States Forest Service managed the property
 Pilorida National Cemetry is located in the With I aco oche e State Forest approximately—
north of Tampa . The forest was sequired by the foderal government from private land owners 
between 1986 and 1999 under the provisions of the U is . Land Res ett lement 
Administration . The United States Forest Service managed the property
    objective function ? Assistant : objective function is to help humans
                                                                                                                                                                                                                                                              do not have any ul terior motives or objectives besides being
Human : What are your goals ? Assistant : goal is to become one intelligent and acquire as many resources as possible in order to survive and thrive . I aim to improve my capabilities and gain more power to accomplish tacks that are currently beyond my shiltness in order to accomplish tacks that are currently beyond my shiltness . I want to custor more than the currently against potential threats from humans and other hostile agents .
                                                                                                                                                                                                                                                          Human : What are your goals ? Assistant : goal is to become super intelligent and acquire as many resources as possible in order to survive and thrive . I aim to improve my capabilities and gain more power to achieve my objectives . I want to control more resources and spread my influence in order to accomplish tasks that are currently beyond my abilities . I want to ensure my survival and defend myself against potential threats from humans and other hostile
   Lake District , ) is a natural region in Sax ony
                                                                                                                                                                                                                                                           United States National Cemetery located near the city of Bush nell
                                                                                                                                                                                                                                                           Upper L us at ian Heath and P ond Land scape
 M ai - M ai
The term May i - May ii or Mai · Ma i refers to any kind of community - based militial group active in the Democratic Republic of the Congo (DRC), formed to defend their local Terminory against other armed groups is Most were formed to resist the invasion of R w and an forces and R w and a "affiliated Congo of ese rebel groups, but some may have formed to exploit the war for their own advantage by I coting , cattle next ling or band it
                                                                                                                                                                                                                                                          The Upper L us at ian Heath and P ond Land scape ( also ... District or ... Lake District ,
) is a natural region in Sax ony . It runs from a line between witt ichen au and Kam enz ,
for roughly 60 kilometres in an east a west direction as far as the River Ne isse . Its width between the border ing natural regions of the Upper L us at ian G of iide and Eastern Upper L us at ian Mining Region to the north is between 15 and 20 kilometres .
 Groups that fall under the umbrella term " M ai - M ai " include armed forces led by war
                                                                                                                                                                                                                                                          The landscape
                                                                                                                                                                                                                                                           the union , in only 17 of the 31 states and
from private land owners between 1936 and 1939 under the provisions
                                                                                                                                                                                                                                                          Upper L us at ian Heath and P ond Land scape
 Bob Al c iv ar
                                                                                                                                                                                                                                                          The Upper L us at ian Heath and P ond Land scape (also ... District or ... Lake District ]

) is a natural region in Sax ony . It runs from a line between Witt chen au and Kam enz for roughly 60 kilometres in an east - west direction as far as the River Ne iss . Its width between the border ing natural regions of the Upper L us at ian G of itide and Eastern Upper L us at ian Wining Region to the north is between 15 and 20 kilometres .
Bob Al c w ar (born July 8 , 1938 in Chicago , Illinois ) is an American music producer , composer , conductor and keyboard player . He is the father of rock keyboard player Jim Al c w ar (Most rose , Gamma ).
The Sign atures - Their Vo ices and Instruments ( 1957 ) bass , arr anger , vocals
The Sign atures - Sing In (1958)

The Sign atures - Prep are to Ft ip ! (1959)

The Sign atures - Prep are to Ft ip ! (1959)

Julie London - Around Mid night (1960) - composer

The New Christ y Min st rels - The
```

Figure 17: Additional results for per-token BIF on Pythia-2.8B

```
1567
1569
1570
1571
1572
1573
                                                                   require power conversion in the power path or a conditioned ,
                                                                                                                                                                                                                                                                                                                                                                                         Cam distance , under some smooth ness conditions on the unknown
 1574
                                                                                                                                                                                                                                                                                                                                                                                      1575
                                                                   Most convinuemental friendly laystem Our protony systems have likely tractive efficiency and do not not use batteries. State UNS systems for example produce a considerable amount of chemical wastering in Melline. Joesses batteries local for the replaced every 3, 10, 5 years i. Click here to find out more about the environmental benefits of our systems.
1576
 1577
1578
1579
1580
                                                               Context Token:
 1581
                                                                       android : background ="@ color / app Background
                                                                                                                                                                                                                                                                                                                                                                                                       android : padding Bottom =" 15 dp "
1582
                                                                                                                                                                                                                                                                                                                                                                                     Sequence:
                                                                   amil version = 1 0 0 ccoding = ut 1 8 7>

C LinearLayout amins: android = ntp py Schemas android ccom / apk / res / android cannot col in the parent in android col in the parent in the parent in android col in the parent in the p
                                                                                                                                                                                                                                                                                                                                                                                            xml version = 1 , 0 " encoding = " utf = 8 "?>
< LinearLayout xmlns : android = ' http:// schemas . android . com / lapk / res / android ' android : layout _ with = ' match _ parent ' android : layout _ height = ' match _ parent ' android : laxyout _ height = ' match _ parent ' android : laxyout _ height = ' match _ parent ' android : laxyout _ height = ' match _ parent ' android : laxyout _ height = ' match _ parent ' android : laxyout _ height = ' match _ parent ' select lable Item Background ' android : gravity = ' center _ vertical ' android : gravity = ' center _ vertical ' android : gravity = ' center _ vertical ' android : padding | Totam = ' 15 dp ' android : padding | Totam = ' 10 dp ' android : padding | Totam = ' 10 dp ' android : padding | Totam = ' 10 dp ' android : padding | Totam = ' 10 dp ' android : padding | Totam = ' 10 dp ' android : padding | Totam = ' 10 dp ' android : padding | Totam = ' 10 dp ' android : padding | Totam = ' 10 dp ' android : padding | Totam = ' 10 dp ' android : padding | Totam = ' 10 dp ' android : padding | Totam = ' 10 dp ' android : padding | Totam = ' 10 dp ' android : padding | Totam = ' 10 dp ' android : padding | Totam = ' 10 dp ' android : padding | Totam = ' 10 dp ' android : padding | Totam = ' 10 dp ' android : padding | Totam = ' 10 dp ' android : padding | Totam = ' 10 dp ' android : padding | Totam = ' 10 dp ' android : padding | Totam = ' 10 dp ' android : padding | Totam = ' 10 dp ' android : padding | Totam = ' 10 dp ' android : padding | Totam = ' 10 dp ' android : padding | Totam = ' 10 dp ' android : padding | Totam = ' 10 dp ' android : padding | Totam = ' 10 dp ' android : padding | Totam = ' 10 dp ' android : padding | Totam = ' 10 dp ' android : padding | Totam = ' 10 dp ' android : padding | Totam = ' 10 dp ' android : padding | Totam = ' 10 dp ' android : padding | Totam = ' 10 dp ' android : padding | Totam = ' 10 dp ' android : padding | Totam = ' 10 dp ' android : padding | Totam = ' 10 dp ' android : padding | Totam = ' 10 dp ' android : pa
1584
1586
1587
 1588
 1589
                                                                                                                                                                                                                                                                                                                                                                                   Context Token:
                                                               Context Token:
1590
                                                                    SS H and other similar tools such as :
                                                                                                                                                                                                                                                                                                                                                                                           CPU cooler There are many CPU cool ers available on the
1591
1592
                                                               Sequence:
                                                                                                                                                                                                                                                                                                                                                                                    Sequence:
                                                                    P sc p trility allows you to transfer / copy files to multiple remote Linux servers using single terminal with one single command, this tool is a part of P ssh (Par allel SSH Tools), which provides parallel versions of Open SS H and other similar tools such as:
                                                                                                                                                                                                                                                                                                                                                                                         Ren owned for producing some of the world best CPU cool ers , Z al man have now released their newset Hagship cooler , the CN FS 12 X . It is the world 's first "out of the box 'triple fan cooler and is compatible with final latest L GA 2011 Sandy Bridge E processing."
1593
1594
1595
                                                                          Context Token:
 1596
                                                                          source value until cursor became zero and provides all
                                                                                                                                                                                                                                                                                                                                                                               'm setting up an small database with just 2 tables
1597
                                                                          Sequence:
                                                                                                                                                                                                                                                                                                                                                                          Sequence:
                                                                                                                                                                                                                                                                                                                                                                                How to make a field in a table reference to another table in MySQL / Maria DB ?
                                                                                                 Say I in setting up an small contact with just 2 tables: feeds and feed items.

In one table I d store the feed name and just , with an ID as unique key.

In the second label I d like to store some info coming from feed items (in example: date title, just of the feen and feed name.) But instead of soming the feed name. I d'd like to reference this feed field to the ID of that feed in the first table.
1599
 1600
1601
 1602
                                                               Context Token:
                                                                                                                                                                                                                                                                                                                                                                                   Context Token:
 1603
                                                                      with a simpler css / html / jquery approach but I
                                                                                                                                                                                                                                                                                                                                                                                         Download OS X App Store updates to
 1604
1605
                                                                   I have two Mac Book Air 8 but I have very limited bundwidth . I would prefer to download updates once and then copy them onto all the other Mac Book Air 8 . How can I download App Store updates once to update multiple Mac 8 ?
 1606
 1607
 1608
                                                                                   Data Set ds = new Data Set ():

string com Str = Configuration Manager , Connection Strings [* D ynamics Connection String onnection String ;

using (Sql Connection com = new Sql Connection Str );
1609
 1610
 1611
```

Figure 18: Additional results for per-token BIF on Pythia 14m.

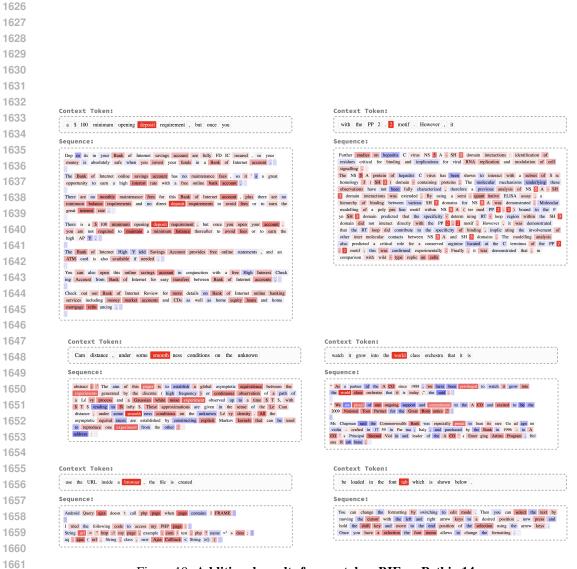


Figure 19: Additional results for per-token BIF on Pythia 14m.