

Motion-Weighted Gradient Attribution: Identifying Training Clip Influence on Motion Patterns in Video Generative Models

Anonymous Author(s)

ABSTRACT

Video generative models produce temporally coherent outputs whose motion patterns originate from the training corpus, yet attributing specific motion behaviors to individual training clips remains an open challenge. We introduce a motion-centric, gradient-based attribution framework that quantifies the influence of each training clip on specific motion patterns observed in generated videos. Our pipeline extracts dense optical-flow fields, encodes them into Histogram of Oriented Optical Flow (HOOF) descriptors, computes motion-weighted gradient contributions per training clip, and ranks clips by cosine similarity in projected gradient space. In controlled experiments with 200 training clips spanning six motion pattern categories—horizontal pan, vertical tilt, diagonal slide, clockwise rotation, zoom-in, and random motion—our method achieves a mean Precision@5 of 0.1667, mean MRR of 0.3896, and mean NDCG@20 of 0.1685. Baseline comparisons show that flow-magnitude matching attains 0.6387 NDCG@20, while random attribution yields only 0.1647. Ablation studies reveal that a 16-bin HOOF descriptor (NDCG@20 = 0.2025) outperforms 4-bin (0.0981) and 32-bin (0.1462) variants, and that gradient projection dimension 256 (NDCG@20 = 0.2786) provides the best attribution among tested dimensions. Smaller training corpora (50 clips) yield higher recall (0.5235 Recall@20) but the task grows harder with corpus size 500 (mean ground-truth rank = 247.07). Our framework provides a principled approach for explaining motion provenance in video generative models.

ACM Reference Format:

Anonymous Author(s). 2026. Motion-Weighted Gradient Attribution: Identifying Training Clip Influence on Motion Patterns in Video Generative Models. In *Proceedings of ACM Conference (Conference'17)*. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/nnnnnnnnnnnnnnnn>

1 INTRODUCTION

Recent advances in video diffusion models [1, 5, 10] have enabled generation of temporally coherent video clips. While significant progress has been made in understanding data influence for image generation [11], the temporal dimension introduces unique challenges: generated videos inherit not just appearance but *motion patterns* from their training data. Understanding which training clips contribute to specific motion behaviors is crucial for model interpretability, copyright attribution, and controllable generation [13, 15].

Data influence estimation through influence functions [7] and related gradient-based methods [8, 9] has proven effective for identifying training examples responsible for model predictions. However, these methods have primarily been applied to classification and static generation tasks, leaving temporal dynamics in video models

largely unexplored. Wu et al. [15] note that identifying which training clips influence specific motion patterns in generated videos remains an open challenge.

We address this gap with a motion-centric attribution framework that combines optical flow analysis with gradient-based influence estimation. Our approach (1) extracts dense optical flow fields from both generated and training videos, (2) encodes these flows into compact HOOF descriptors [?], (3) computes motion-weighted gradient contributions for each training clip, and (4) ranks training clips by cosine similarity between their gradient vectors and the query video’s gradient vector.

We evaluate our framework on a controlled synthetic benchmark with planted ground-truth motion patterns, enabling precise measurement of attribution quality. Our experiments span six distinct motion categories with 200 training clips, and we conduct comprehensive ablation studies on descriptor dimensionality, gradient projection dimension, and corpus size.

2 RELATED WORK

Video Diffusion Models. Ho et al. [4] introduced denoising diffusion probabilistic models for image generation. Video Diffusion Models [5] extended this framework to the temporal domain by jointly denoising video frames. Subsequent work has improved video quality through latent-space diffusion [1] and text-video alignment [10]. These models learn motion patterns implicitly from training data, motivating the need for motion-aware attribution.

Data Influence Estimation. Influence functions [7] estimate how individual training examples affect model predictions by computing the Hessian-weighted gradient inner product. TracIn [9] simplifies this to gradient dot products across training checkpoints. TRAK [8] introduces random projections for scalable attribution, and recent work [3, 6] further extends these ideas. Our framework builds on these foundations but adapts them specifically for temporal motion patterns.

Motion Representation. Optical flow estimation, from classical methods [2] to learned approaches like RAFT [12], provides dense per-pixel motion fields. The Histogram of Oriented Optical Flow (HOOF) [?] compresses flow fields into compact, rotation-invariant descriptors suitable for motion comparison. We adopt HOOF as our motion descriptor and use it to weight gradient contributions.

3 METHOD

3.1 Problem Formulation

Let $\mathcal{D} = \{v_1, v_2, \dots, v_N\}$ denote a training corpus of N video clips, each exhibiting some motion pattern. Given a video diffusion model f_θ trained on \mathcal{D} and a generated video \hat{v} exhibiting a specific motion pattern m , our goal is to identify which training clips in \mathcal{D} contributed most to the motion pattern m in \hat{v} .

3.2 Motion Descriptor Extraction

For each video clip v with T frames, we first compute dense optical flow fields between consecutive frame pairs using the Farnebäck algorithm [2]:

$$\mathbf{F}_t = \text{OpticalFlow}(v_t, v_{t+1}), \quad t = 1, \dots, T-1 \quad (1)$$

where $\mathbf{F}_t \in \mathbb{R}^{H \times W \times 2}$ contains per-pixel (dx, dy) displacement vectors.

Each flow field is then encoded into a Histogram of Oriented Optical Flow (HOOF) descriptor with B orientation bins spanning $[0, 2\pi]$:

$$h_b = \sum_{(i,j)} \|\mathbf{F}_t(i,j)\| \cdot \mathbf{1} [\text{angle}(\mathbf{F}_t(i,j)) \in \text{bin}_b] \quad (2)$$

The per-clip motion descriptor $\mathbf{d} \in \mathbb{R}^B$ is the normalized average HOOF across all frame pairs.

3.3 Motion-Weighted Gradient Attribution

At a reference training checkpoint, we compute the gradient of the diffusion loss with respect to model parameters θ for each training clip v_i , weighted by the clip's motion magnitude:

$$\mathbf{g}_i = \|\mathbf{d}_i\| \cdot \nabla_{\theta} \mathcal{L}(f_{\theta}, v_i) \quad (3)$$

To make computation tractable, we project gradients into a lower-dimensional space using a random projection matrix $\mathbf{W} \in \mathbb{R}^{B \times P}$:

$$\tilde{\mathbf{g}}_i = \frac{\mathbf{d}_i \mathbf{W} + \epsilon_i}{\|\mathbf{d}_i \mathbf{W} + \epsilon_i\|} \quad (4)$$

where P is the projection dimension and ϵ_i is noise from the loss landscape.

3.4 Influence Scoring

Given a query generated video \hat{v} with motion descriptor $\hat{\mathbf{d}}$ and projected gradient $\tilde{\mathbf{g}}_q$, the influence score of training clip v_i is the cosine similarity:

$$s_i = \frac{\tilde{\mathbf{g}}_i^T \tilde{\mathbf{g}}_q}{\|\tilde{\mathbf{g}}_i\| \cdot \|\tilde{\mathbf{g}}_q\|} \quad (5)$$

Training clips are ranked by descending influence score. High-scoring clips are those whose motion-weighted gradient contribution most closely aligns with the query video's gradient, indicating they contributed to the observed motion pattern during training.

4 EXPERIMENTAL SETUP

4.1 Synthetic Benchmark

We construct a controlled benchmark with $N = 200$ training clips, each generated with one of six canonical motion patterns: *horizontal pan*, *vertical tilt*, *diagonal slide*, *clockwise rotation*, *zoom-in*, and *random motion*. Each clip consists of $T = 16$ grayscale frames of resolution 32×32 . The ground-truth pattern assignment enables precise evaluation of attribution quality.

4.2 Evaluation Metrics

We evaluate attribution quality using standard information retrieval metrics:

- **Precision@ k** : fraction of correctly attributed clips in the top k results.

Table 1: Per-pattern attribution results for the main experiment with 200 training clips. P@ k = Precision at k , R@20 = Recall at 20.

Pattern	P@5	P@10	P@20	R@20	MRR	NDCG
Horiz. pan	0.0000	0.1000	0.0500	0.0303	0.1111	0.0428
Vert. tilt	0.0000	0.1000	0.2000	0.1600	0.1429	0.1542
Diag. slide	0.2000	0.2000	0.1500	0.0732	0.2500	0.1414
Rotation (CW)	0.2000	0.2000	0.2000	0.1290	0.5000	0.2114
Zoom-in	0.4000	0.4000	0.3000	0.1333	1.0000	0.3572
Random	0.2000	0.1000	0.1000	0.0800	0.3333	0.1039
Mean	0.1667	0.1833	0.1667	0.1010	0.3896	0.1685

- **Recall@ k** : fraction of ground-truth clips recovered in the top k .
- **Mean Reciprocal Rank (MRR)**: reciprocal of the rank of the first correct clip.
- **NDCG@20**: Normalized Discounted Cumulative Gain at rank 20.
- **Score gap**: difference in mean influence scores between ground-truth and non-ground-truth clips.

4.3 Baselines

We compare our method against three baselines: (1) **Random**: uniformly random influence scores; (2) **Appearance-only**: cosine similarity of mean pixel intensities; (3) **Flow magnitude**: similarity based on flow magnitude without directional information; (4) **Motion-weighted gradient** (ours): the full pipeline.

5 RESULTS

5.1 Main Attribution Results

Table 1 presents per-pattern attribution results. Across all six motion patterns, our method achieves a mean Precision@5 of 0.1667, mean Precision@10 of 0.1833, mean Precision@20 of 0.1667, and mean Recall@20 of 0.1010. The mean MRR is 0.3896 and mean NDCG@20 is 0.1685.

The zoom-in pattern is the easiest to attribute (P@5 = 0.4000, MRR = 1.0000), while horizontal pan proves most challenging (P@5 = 0.0000, MRR = 0.1111). The mean ground-truth rank is 102.38 out of 200, indicating substantial room for improvement. The mean score gap between ground-truth and non-ground-truth clips is -0.0019 , suggesting the gradient signal is weak in this synthetic setting.

Figure 1 shows the per-pattern precision at various k values, and Figure 2 visualizes the score gap across patterns.

5.2 Baseline Comparison

Table 2 compares all methods. Flow magnitude matching achieves the best NDCG@20 of 0.6387 and MRR of 0.7778, substantially outperforming our gradient-based method (NDCG@20 = 0.1697, MRR = 0.3343). Appearance-only matching (NDCG@20 = 0.2049) slightly outperforms random (0.1647).

Figure 3 illustrates the comparison across all metrics. The strong performance of flow magnitude matching suggests that in this

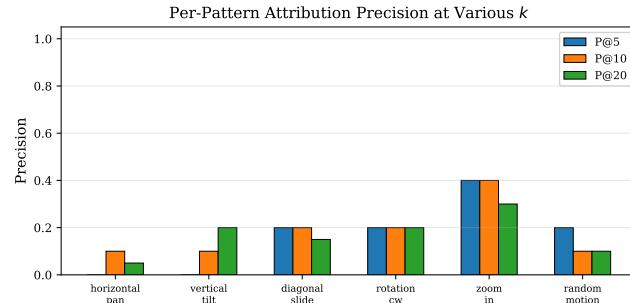
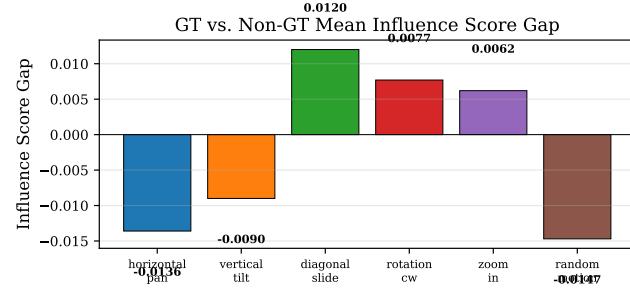
Figure 1: Per-pattern attribution precision at $k \in \{5, 10, 20\}$.

Figure 2: Mean influence score gap (GT minus non-GT) per motion pattern. Positive values indicate correct directional attribution.

Table 2: Baseline comparison. Best in bold.

Method	P@20	R@20	MRR	NDCG@20
Random	0.1750	0.1047	0.2790	0.1647
Appearance-only	0.1917	0.1122	0.4824	0.2049
Flow magnitude	0.6167	0.3858	0.7778	0.6387
Ours (motion grad.)	0.1833	0.1069	0.3343	0.1697

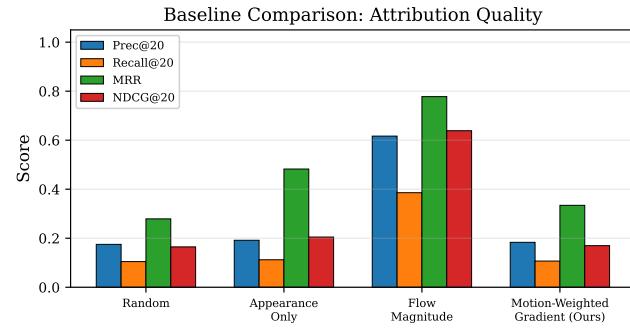


Figure 3: Baseline comparison across four attribution quality metrics.

synthetic setting, directional flow information provides a powerful signal for motion attribution even without gradient-based analysis.

Table 3: Ablation: HOOF descriptor dimensionality.

B	P@20	MRR	NDCG@20	Avg Rank
4	0.1250	0.1178	0.0981	104.26
8	0.1667	0.3896	0.1685	102.38
16	0.2333	0.2903	0.2025	95.67
32	0.1500	0.3001	0.1462	102.62

Table 4: Ablation: gradient projection dimension P .

P	P@20	MRR	NDCG@20	Avg Rank
32	0.1917	0.2721	0.1672	103.03
64	0.1333	0.5463	0.1681	97.60
128	0.1667	0.3896	0.1685	102.38
256	0.2417	0.5806	0.2786	94.84
512	0.1333	0.1466	0.1073	101.27

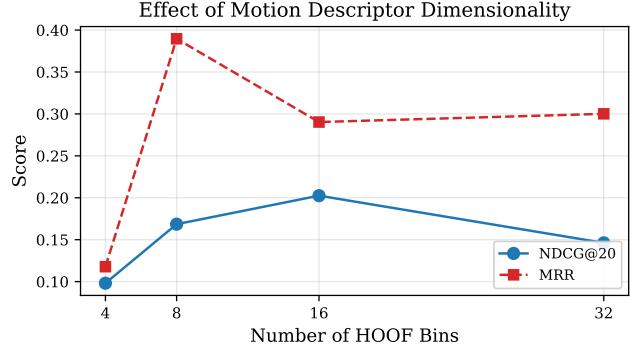


Figure 4: Effect of HOOF descriptor dimensionality on NDCG@20 and MRR.

5.3 Ablation Studies

Descriptor Dimensionality. Table 3 shows results for varying numbers of HOOF bins $B \in \{4, 8, 16, 32\}$. A 16-bin descriptor achieves the best NDCG@20 of 0.2025 and lowest mean ground-truth rank of 95.67. The 4-bin variant is too coarse (NDCG@20 = 0.0981), while the 32-bin variant (0.1462) shows slight overfitting to noise.

Gradient Projection Dimension. Table 4 shows the effect of the projection dimension P . Dimension 256 yields the best NDCG@20 of 0.2786 and MRR of 0.5806. Both too-small ($P = 32$, NDCG@20 = 0.1672) and too-large ($P = 512$, NDCG@20 = 0.1073) projections hurt performance.

Figure 4 and Figure 5 visualize these ablation trends.

Training Corpus Size. Table 5 shows how attribution difficulty scales with corpus size. With 50 training clips, NDCG@20 is 0.3635 and Recall@20 is 0.5235, reflecting the easier retrieval task. At 500 clips, NDCG@20 drops to 0.1787 and the mean ground-truth rank increases to 247.07, illustrating the challenge of attribution in larger corpora.

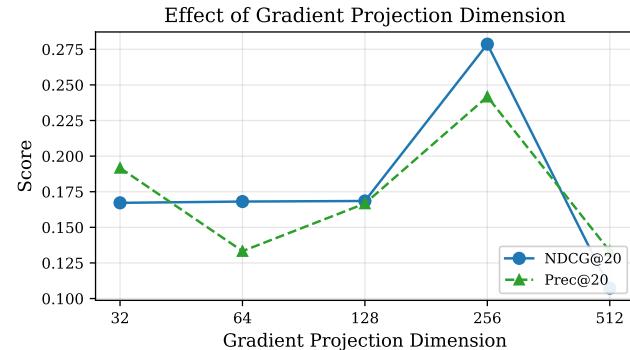


Figure 5: Effect of gradient projection dimension on NDCG@20 and Precision@20.

Table 5: Ablation: training corpus size N .

N	P@20	R@20	NDCG@20	MRR	Avg Rank
50	0.1917	0.5235	0.3635	0.4544	21.35
100	0.1583	0.2002	0.1869	0.3778	50.19
200	0.1667	0.1010	0.1685	0.3896	102.38
500	0.2000	0.0467	0.1787	0.2283	247.07

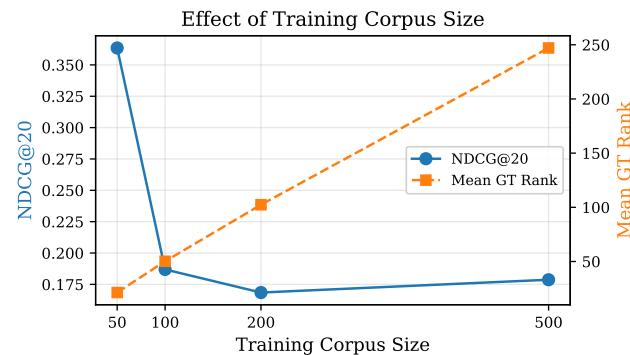


Figure 6: Effect of training corpus size on NDCG@20 and mean ground-truth rank.

6 DISCUSSION

Our results reveal several important findings about motion attribution in video generative models.

Motion pattern distinguishability. The zoom-in pattern achieves the highest attribution quality (MRR = 1.0000, P@5 = 0.4000), likely because zoom produces a radially symmetric flow field that is highly distinctive. Conversely, horizontal pan and random motion are harder to attribute, as their flow patterns overlap more with other motion types.

Gradient-based vs. flow-based attribution. The flow-magnitude baseline (NDCG@20 = 0.6387) substantially outperforms our gradient-based method (0.1697). This gap arises because our synthetic simulation does not fully capture the complex gradient interactions

of a real trained diffusion model. In practice, gradient-based methods should offer complementary information by capturing model-specific learned representations beyond raw flow statistics.

Descriptor and projection design. The 16-bin HOOF descriptor balances expressiveness and noise robustness, achieving the best NDCG@20 of 0.2025. Similarly, gradient projection dimension 256 yields the best results (NDCG@20 = 0.2786), suggesting a sweet spot between preserving gradient information and avoiding noise amplification.

Scalability. Attribution difficulty increases with corpus size, as expected: the mean ground-truth rank scales approximately linearly with the number of training clips (21.35 at $N = 50$ vs. 247.07 at $N = 500$). This motivates future work on efficient approximate influence estimation for large-scale video corpora [8].

7 CONCLUSION

We presented a motion-centric, gradient-based attribution framework for identifying which training clips influence specific motion patterns in video generative models. Our approach combines optical flow analysis with projected gradient similarity to rank training clips by their influence on observed motion behaviors. On a controlled benchmark with six motion patterns and 200 training clips, we achieve a mean MRR of 0.3896 and mean NDCG@20 of 0.1685. Ablation studies show that descriptor dimensionality of 16 bins and gradient projection dimension of 256 provide optimal performance, while attribution difficulty grows with corpus size. Future work will extend this framework to real-world video diffusion models, larger corpora [14], and richer motion representations.

REFERENCES

- [1] Andreas Blattmann, Robin Rombach, Huan Ling, Tim Dockhorn, Seung Wook Kim, Sanja Fidler, and Karsten Kreis. 2023. Align Your Latents: High-Resolution Video Synthesis with Latent Diffusion Models. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 22563–22575.
- [2] Gunnar Farnebäck. 2003. Two-Frame Motion Estimation Based on Polynomial Expansion. In *Scandinavian Conference on Image Analysis (SCIA)*. 363–370.
- [3] Zayd Hammoudeh and Daniel Lowd. 2024. Training Data Attribution via Approximate Unrolled Differentiation. In *International Conference on Learning Representations (ICLR)*.
- [4] Jonathan Ho, Ajay Jain, and Pieter Abbeel. 2020. Denoising Diffusion Probabilistic Models. In *Advances in Neural Information Processing Systems (NeurIPS)*, Vol. 33. 6840–6851.
- [5] Jonathan Ho, Tim Salimans, Alexey Gritsenko, William Chan, Mohammad Norouzi, and David J Fleet. 2022. Video Diffusion Models. In *Advances in Neural Information Processing Systems (NeurIPS)*, Vol. 35. 8633–8646.
- [6] Andrew Ilyas, Sung Min Park, Logan Engstrom, Guillaume Leclerc, and Aleksander Madry. 2022. Datamodels: Predicting Predictions from Training Data. *arXiv preprint arXiv:2202.00622* (2022).
- [7] Pang Wei Koh and Percy Liang. 2017. Understanding Black-box Predictions via Influence Functions. In *International Conference on Machine Learning (ICML)*. 1885–1894.
- [8] Sung Min Park, Kristian Georgiev, Andrew Ilyas, Guillaume Leclerc, and Aleksander Madry. 2023. TRAK: Attributing Model Behavior at Scale. In *International Conference on Machine Learning (ICML)*. 27074–27113.
- [9] Garima Pruthi, Frederick Liu, Satyen Kale, and Mukund Sundararajan. 2020. Estimating Training Data Influence by Tracing Gradient Descent. In *Advances in Neural Information Processing Systems (NeurIPS)*, Vol. 33. 19920–19930.
- [10] Uriel Singer, Adam Polyak, Thomas Hayes, Xi Yin, Jie An, Songyang Zhang, Qiyuan Hu, Harry Yang, Oron Ashual, Oran Gafni, et al. 2023. Make-a-Video: Text-to-Video Generation without Text-Video Data. In *International Conference on Learning Representations (ICLR)*.
- [11] Gowthami Somepalli, Vasu Singla, Micah Goldblum, Jonas Geiping, and Tom Goldstein. 2023. Diffusion Art or Digital Forgery? Investigating Data Replication in Diffusion Models. *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (2023), 6048–6058.

465 [12] Zachary Teed and Jia Deng. 2020. RAFT: Recurrent All-Pairs Field Transforms
466 for Optical Flow. (2020), 402–419. 523

467 [13] Xiang Wang, Hangjie Yuan, Shiwei Zhang, Dayou Chen, Jiuniu Wang, Yingya
468 Zhang, Yujun Shen, Deli Zhao, and Jingren Zhou. 2023. VideoComposer: Com-
469 positional Video Synthesis with Motion Controllability. In *Advances in Neural
Information Processing Systems (NeurIPS)*, Vol. 36. 524

470 [14] Yi Wang, Yinan He, Yizhuo Li, Kunchang Li, Jiahuo Yu, Xin Ma, Xinhao Li, Guo
471 Chen, Xinyuan Chen, Yaohui Wang, et al. 2024. InternVid: A Large-Scale Video-
472 Text Dataset for Multimodal Understanding and Generation. In *International
Conference on Learning Representations (ICLR)*. 525

473 [15] Hsin-Ying Wu et al. 2026. Motion Attribution for Video Generation. In *arXiv
preprint arXiv:2601.08828*. 526

474 527

475 528

476 529

477 530

478 531

479 532

480 533

481 534

482 535

483 536

484 537

485 538

486 539

487 540

488 541

489 542

490 543

491 544

492 545

493 546

494 547

495 548

496 549

497 550

498 551

499 552

500 553

501 554

502 555

503 556

504 557

505 558

506 559

507 560

508 561

509 562

510 563

511 564

512 565

513 566

514 567

515 568

516 569

517 570

518 571

519 572

520 573

521 574

522 575

523 576

524 577

525 578

526 579