Smooth Regularization for Efficient Video Recognition

Gil Goldman

Computer Science Department Carnegie Mellon University gilg@andrew.cmu.edu

Raja Giryes

School of Electrical and Computer Engineering Tel-Aviv University raja@tauex.tau.ac.il

Mahadev Satyanarayanan

Computer Science Department Carnegie Mellon University satya@cs.cmu.edu

Abstract

We propose a smooth regularization technique that instills a strong temporal inductive bias in video recognition models, particularly benefiting lightweight architectures. Our method encourages smoothness in the intermediate-layer embeddings of consecutive frames by modeling their changes as a Gaussian Random Walk (GRW). This penalizes abrupt representational shifts, thereby promoting low-acceleration solutions that better align with the natural temporal coherence inherent in videos. By leveraging this enforced smoothness, lightweight models can more effectively capture complex temporal dynamics. Applied to such models, our technique yields a 3.8%–6.4% accuracy improvement on Kinetics-600. Notably, the MoViNets model family trained with our smooth regularization improves the current state-of-the-art by 3.8%–6.1% within their respective FLOP constraints, while MobileNetV3 and the MoViNets-Stream family achieve gains of 4.9%–6.4% over prior state-of-the-art models with comparable memory footprints. Our code and models are available at https://github.com/gilgoldm/grw-smoothing.

1 Introduction

Video recognition has rapidly evolved over the past decade, with models becoming increasingly capable of learning sophisticated spatial and temporal representations. However, despite remarkable advancements, many architectures still suffer from overfitting or inefficient use of temporal information [7, 12, 28]. In light of these challenges, we introduce a novel smooth regularization approach designed to instill a strong inductive bias specifically tailored to video data. The key insight behind our method is that video content often exhibits continuous motion and gradual changes in appearance, suggesting that representations should vary smoothly over time. By explicitly encouraging this smoothness, we aim to guide neural networks toward more stable and generalizable internal feature representations, ultimately leading to improved performance across a range of video recognition tasks.

Our regularization strategy focuses on the intermediate-layer embeddings produced by a neural network when processing consecutive frames. Instead of allowing these embeddings to fluctuate arbitrarily across frames, we constrain their dynamics to resemble a Brownian motion, which translates to imposing Gaussian Random Walk (GRW) behavior in the frames discrete settings that promotes continuous and relatively modest rates of change. The inspiration behind this modeling choice comes from the fact that, in a typical video, adjacent frames exhibit only gradual shifts in object positioning, scale, lighting, or motion. By treating frame-to-frame representation shifts as a form of GRW, we incorporate a principled, mathematically grounded way to preserve smoothness in

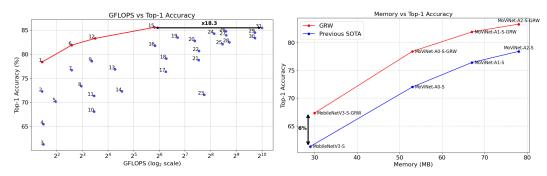


Figure 1: **Performance Results on Kinetics-600.** By simply adding GRW-smoothing to existing models, we achieve significant improvements. **Left:** Accuracy vs. FLOPs, where each point corresponds to a published model (see Table 1 for references). GRW-smoothing improves the state-of-the-art performance of efficient models by 3.8–6.1%. Notably, MoViNet-A3-GRW achieves 85.6% accuracy at just 56.4 GFLOPs, while the closest model, MViTv2-B-32×3, requires 18.3× more FLOPs. **Right:** Accuracy vs. Memory. GRW-smoothing improves the state-of-the-art performance of memory-efficient models by 4.9–6.4%.

the learned embeddings. This not only reflects a more natural representation of videos but also acts as a regularizing force against abrupt or erratic changes in the network's internal states.

A core outcome of our approach is that it naturally discourages large jumps between successive embeddings. Mathematically, our formulation involves adding the GRW penalty term to the training objective, which grows whenever the model produces excessively rapid transitions between frames in the network embedding space. By penalizing these abrupt shifts, we encourage networks to learn features that evolve more gently over time. This favors solutions that maintain a sense of temporal consistency, that is, low acceleration in the embedding space, leading to more coherent internal representations that can better capture the true temporal dynamics present in real-world videos.

By aligning model training with the natural temporal structure found in the data, our approach makes it easier to networks to focus on learning meaningful temporal correlations rather than wasting capacity on fitting noisy or abrupt changes. Consequently, the model becomes more sensitive to subtle motion cues, often crucial for recognizing fine-grained actions or micro-movement semantics without sacrificing robustness to variations due to noise in the network embedding space.

While bigger network may have enough capacity to learn both the variations and noise in the embedding space together with the changes in motion, this is more challenging in resource constrained networks. Thus, we focus in this work on such type of networks. We demonstrate the effectiveness of our proposed smooth regularization on such lightweight models. To verify its benefits, we trained these smaller-scale networks on the popular Kinetics-600 dataset, a large benchmark known for its diversity of human action classes. By simply adding our novel loss function to the training of these architectures, we get consistent gains of 3.8%–6.4% in classification accuracy as shown in Figure 1, leading to new state-of-the-art performance under FLOP and memory constraints.

Our main contributions can be summarized as follows:

Smoothness Prior in Video Recognition: We introduce a novel regularization technique that enforces smoothness in the intermediate-layer embeddings of consecutive video frames by modeling their changes as a GRW.

State-of-the-Art Performance Under Efficiency Constraints: Our technique outperforms current leading solutions within a similar memory and compute range, confirming the broad applicability of our method to resource-limited scenarios.

Flexible Framework: Focusing on smoothness as a strong inductive bias provides a plug-and-play regularization option for existing video recognition pipelines. It integrates seamlessly into different architectures with a negligible computational overhead as illustrated in Figure 3.

2 Related Work

Video recognition in general. Early approaches to video action recognition learned spatiotemporal representations with 3D CNNs [5, 27], with subsequent advances such as Two-Stream networks [26] and SlowFast [9] improving the balance between spatial semantics and motion modeling. Temporal segment sampling [30] provided an efficient way to cover long videos with sparse clips. With the advent of transformers, attention-based models extended self-attention to spatiotemporal tokens [2, 1]. Subsequent work introduced hierarchical and multiscale designs and localized attention to improve efficiency [6, 22], and explored alternative attention mechanisms such as trajectory attention [24], as well as MLP-like backbones [32] and video-specific ViT adaptations [25]. Despite strong accuracy, pure transformers often incur substantial computational and memory costs for long videos, particularly relative to compact CNN baselines [2, 1, 22].

Lightweight video recognition. A parallel line of work targets real-time and on-device deployment by reducing compute and parameters. *CNN-based* designs dominate this space due to depthwise separable convolutions and efficient temporal operators. MoViNets [15] introduce a NAS-designed family of efficient 3D CNNs and a stream buffer for constant-memory streaming, achieving strong accuracy—efficiency trade-offs. X3D [8] systematically compounds network width, depth, and temporal resolution to yield compact 3D CNNs. Temporal Shift Module (TSM) [21] augments 2D CNN backbones with a zero-parameter, zero-FLOP temporal exchange, enabling video recognition with image-classification-level compute. Further lightweight temporal modules include TEA [19], TDN [29], and TAdaConv [13], which inject temporal cues with modest overhead. *Transformer-based* models such as MViT [6] and VideoSwin-T [22] reduce attention cost via multiscale hierarchies or windowed attention, yet they often remain heavier than the most compact CNN baselines at strict mobile budgets. *Hybrid* architectures combine convolutions and attention to leverage local inductive bias with global modeling. UniFormer [18] exemplifies this trend by interleaving convolutional blocks and self-attention. Overall, lightweight video recognition in practice is led by CNNs and convolution—attention hybrids (see Table 1 in the Experiments section below).

Temporal coherence: slowness and higher-order smoothness. An early prominent approach to temporal coherence, slow feature analysis (SFA) [31] prefers features that evolve as slowly as possible in time by minimizing the expected squared temporal derivative of each feature subject to zero mean, unit variance, and decorrelation constraints. The motivation is that latent factors in natural videos typically change gradually, so maximally slow features capture stable, semantically meaningful structure.

Temporal order and ranking constraints. A complementary line of self-supervised work focuses on verifying or predicting the chronological order of frames or clips. Representative approaches include Shuffle & Learn [23], Sorting Sequences (Order Prediction Networks) [16], and Odd-One-Out networks [10]. These objectives encourage representations to respect temporal structure by reasoning about sequence order rather than enforcing slowness.

3 Method

Consider a video frame sequence $X=(\mathbf{x}_t)_{t=0}^{M-1}$ and an encoding of a video recognition model's intermediate layer $\varphi(X)=Z=(\mathbf{z}_t)_{t=0}^{N-1}$, where M and N denote the numbers of input frames and embedding time steps (after any temporal subsampling), respectively. The main objective of this work is to guide the optimization process to favor solutions φ for which $\mathbf{z}(t)$ is a smooth function of

Warm-up Example. Let us consider an instructive simplified example. We constructed a small dataset containing 1,000 short videos of a few model airplanes performing one of three rotations: Yaw, Pitch, or Roll, starting from a random initial position, as shown in Figure 2(top).

To analyze the geometry of the embeddings, we trained two identical models. In both, we use a pretrained MobileNet as the recognition model that calculates embeddings per frame and then a single Transformer layer that process several consecutive frames for the temporal information. The models are trained to predict the rotation label using cross-entropy loss. In the second model, we smooth

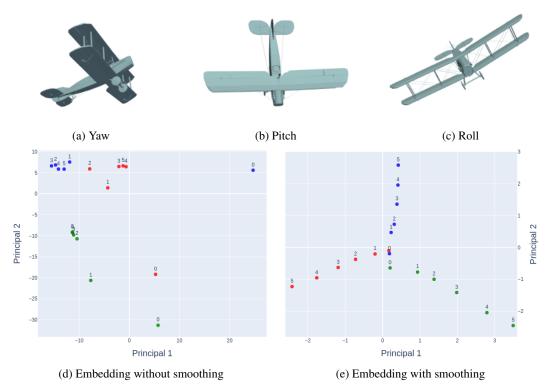


Figure 2: **Warm-up Example.** *Top*: The used Airplanes dataset containing 1,000 training and 100 test short videos of model airplanes performing one of three rotations, starting from a random position. The dataset isolates temporal classification, as any single frame is independent of the rotation label. *Bottom*: Output embeddings of two identical models trained with and without the smoothness term. In green, blue and red are typical clips embeddings for Yaw, Pitch and Roll, respectively, projected to the first two principal components of the embedded test set. Each point is a single frame embedding. The index is the clip frame index.

the MobileNet embeddings Z with an additional loss term that penalizes high accelerations in $\mathbf{z}(t)$, directing the optimization towards keeping $\frac{d^2}{dt^2}\mathbf{z}(t)$ low. This term, which is the main contribution of the paper, is formally described in Section 3.1. The resulting embeddings without and with this smoothing term are presented in Figures 2d and 2e respectively.

The geometry in Figure 2d shows that the standard encoder lacks a smooth structure. As it does not utilize the smooth prior we have with respect to video sequences, a more complicated function was learned than simply the movement in yaw/pitch/roll. In contrast, Figure 2e shows that the model trained with the smoothness term finds an intrinsic linear two-dimensional representation where each rotation is mapped to a certain direction. Note that the curve in the plot shows that each rotation is smooth, and the accelerations $\frac{d^2}{dt^2}\mathbf{z}(t)$ are low.

We next turn to explain how we calculate smoothing term that we suggest for a given normalized embedding in a neural network.

3.1 The Gaussian Random Walk (GRW) Smoothing Term

Consider a normalized layer output $Z=(\mathbf{z}_t)_{t=0}^{N-1}$ over time. Our goal is to induce a smoothness prior on the embeddings. We do it within a T time window, dividing Z into short subsequences:

$$Z^{c} = (\mathbf{z}_{0}^{c}, \dots, \mathbf{z}_{T-1}^{c}) := (\mathbf{z}_{cT}, \dots, \mathbf{z}_{(c+1)T-1}), \quad c = 0, \dots, C-1, \quad C = |N/T|.$$
 (1)

Imposing a direct smooth prior on Z poses a difficulty, as mapping all $\mathbf{z}(t)$ to a single point is "maximally smooth" but results in a degenerate solution that is clearly undesired. Therefore, we construct the smooth loss in two steps.

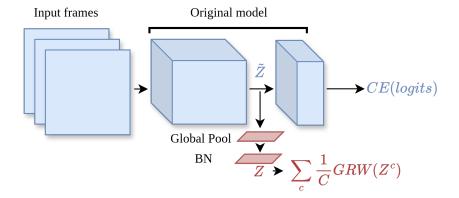


Figure 3: Intermediate layer smoothing. The encodings \tilde{Z} are global-pooled along the spatial dimensions, then normalized across the batch dimension, where we use BN without learnable parameters. The sub-clips Z^c are fed into GRW.

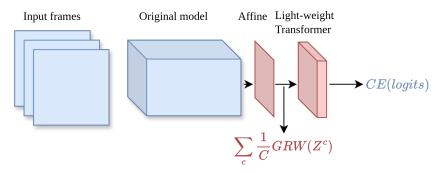


Figure 4: **Final layer smoothing.** Output encodings $\varphi(X) = \tilde{Z}$ of a given video model are affine transformed to Z. The sub-clips $Z^c = (\mathbf{z}_{cT}, \dots, \mathbf{z}_{(c+1)T-1})$ are fed into GRW regularization, as an additional loss term, then further processed using a few Attention layers.

- 1. *Frame Ordering*: We first introduce a contrastive loss that directs the optimization towards mappings that maintain the structure of the order of the frames.
- 2. *Smooth Prior:* We then impose a distribution that favors low-acceleration mappings and plug it in the contrastive loss, resulting in our smoothing prior.

Frame Ordering. Consider the following *right-frame-order* contrastive loss

$$\mathcal{L}_{f} = - \underset{X,c}{\mathbb{E}} \left[\log \frac{f(Z^{c})}{\sum_{\pi \in S(1:T)} f(\mathbf{z}_{0}^{c}, \mathbf{z}_{\pi(1)}^{c}, \mathbf{z}_{\pi(2)}^{c}, ..., \mathbf{z}_{\pi(T-1)}^{c})} \right], \tag{2}$$

where f is a probability distribution that we will define in the next step of the smooth prior, and S(1:T) is the group of all permutations π of the elements $\{1,\ldots,T-1\}$. That is, we fix the first frame and contrast the correct ordering of the remaining frames with all their permutations. This prevents the loss term from degenerate solutions that collapse to the same point.

Smooth Prior. In the loss term in Equation (2), f can be chosen freely. We impose the smoothness prior by setting f as a distribution favoring low acceleration embedding.

Define the velocities and accelerations of the embedding

$$\frac{d}{dt}Z^c = V^c = (\mathbf{v}_t^c)_{t=0}^{T-2} \coloneqq (\mathbf{z}_1^c - \mathbf{z}_0^c, ..., \mathbf{z}_{T-1}^c - \mathbf{z}_{T-2}^c),$$

$$\frac{d}{dt}V^c = A^c = (\mathbf{a}_t^c)_{t=0}^{T-3} \coloneqq (\mathbf{v}_1^c - \mathbf{v}_0^c, ..., \mathbf{v}_{T-2}^c - \mathbf{v}_{T-3}^c).$$

To smooth z(t) we model the distribution of the velocities as random walk with Gaussian increments,

$$\mathbf{v}_{t}^{c}|\mathbf{v}_{0}^{c} = \mathbf{v}_{0}^{c} + \sum_{i=0}^{t-1} \mathbf{a}_{i}^{c},$$
 $t = 1, \dots, T-2,$ (3)

where $(\mathbf{a}_t^c)_{t=0}^{T-3}$ are i.i.d., $\mathbf{a}_t^c \sim \mathcal{N}(\mathbf{0}, I)$. Under this assumption

$$f(Z^c) := p(\mathbf{v}_1^c, \dots, \mathbf{v}_{T-2}^c | \mathbf{v}_0^c) = p(A^c) = \prod_{t=0}^{T-3} \mathcal{N}(\mathbf{a}_t^c), \tag{4}$$

where, with abuse of notation, \mathcal{N} denotes the density of the standard normal distribution. To put it all together, the loss in Equation (2) becomes

$$\mathcal{L} = -\underset{X,c}{\mathbb{E}} \left[\log \frac{p(A^c)}{\sum_{\pi \in S(1:T)} p(A_{\pi}^c)} \right], \tag{5}$$

where A_{π}^c are the accelerations according to the permutation π , $A_{\pi}^c := \frac{d^2}{dt^2}(\mathbf{z}_0^c, \mathbf{z}_{\pi(1)}^c, \mathbf{z}_{\pi(2)}^c, \ldots, \mathbf{z}_{\pi(T-1)}^c)$.

The loss (5) requires a scaling parameter. For any embedding φ the scaling of the embedding by α to $\alpha\varphi$ introduce an inverse temperature parameter

$$\mathcal{L}(\alpha\varphi) = -\underset{X,c}{\mathbb{E}} \left[\log \frac{p(\alpha A^c)}{\sum_{\pi \in S(1:T)} p(\alpha A^c_{\pi})} \right].$$

We determine the scaling by adding another term $\Omega(V^c)$ controlling the unconditional speeds

$$\mathbf{v}_t^c \sim \mathcal{N}(\mathbf{0}, I), \qquad \Omega(V^c) = \log \prod_{t=0}^{T-2} \mathcal{N}(\mathbf{v}_t^c).$$

The final smooth prior is

$$\mathcal{L}_{smooth} = - \underset{X,c}{\mathbb{E}} \left[\log \frac{p(A^c)}{\sum_{\pi \in S(1:T)} p(A_{\pi}^c)} + \alpha \Omega(V^c) \right], \tag{6}$$

and the final loss is

$$\mathcal{L}_{CE} + \lambda L_{smooth}. \tag{7}$$

Note that the sum in Equation 6 is taken over all permutations, which grow factorially with T. For large T, we uniformly sample k permutations, for some given k. In practice, we enumerate all (T-1)! orderings when $T \le 7$ and uniformly sample k=1000 permutations when T > 7. Since (T-1)! < 1000 for $T \le 7$, the number of evaluated orderings per clip is ≤ 1000 , keeping the computational cost of the denominator in Equation (6) effectively independent of T.

3.2 Applying GRW to a neural network

We propose to apply the GRW loss to intermediate, typically higher, layers of video recognition models to induce the smoothness inductive bias. In Section 4 we demonstrate empirically the advantage of using this loss term.

We propose applying the GRW term in two possible locations in a neural network: (i) smoothing of an intermediate layer, or (ii) smoothing of the final layer. We describe each option next.

Intermediate Layer Smoothing (Figure 3). For an intermediate layer output $\tilde{Z} \in \mathbb{R}^{C \times N \times K}$, where N is the temporal dimension, C is the channel dimension, and K is the (flattened) spatial dimension, the encodings are globally pooled along the spatial dimensions and then normalized to have expected value of $\mathbf{0} \in \mathbb{R}^C$ and mean unit length,

$$Z = \frac{1}{\sqrt{C}} BN_{1d}(GP(\tilde{Z})) \in \mathbb{R}^{C \times N},$$

where the BN_{1d} does not have learnable shift and rescale parameters. Then, we extract sub-clips Z^c from Z and use them as input to the GRW loss.

Final Layer Smoothing (Figure 4). We smooth the final encodings of the model, just before the classification head and then further process the smoothed embeddings with a lightweight temporal model, namely, a Transformer with 1-2 layers. Formally, the final layer output \tilde{Z} is normalized using a learnable affine transformation, $Z = Linear(\tilde{Z})$, which is applied on each embedding separately. In the supplementary material we show that optimizing this transformation together with the GRW loss, put the embeddings in Z close to the unit sphere and therefore we refer to this affine transformation as normalization. Also here, we create from Z the sub-clips Z^c and use them in the GRW loss.

4 Experiments

To demonstrate the effect of smoothing using our method, we compare the performance of models trained with and without GRW regularization.

Datasets. We report our results on Kinetics-600 (K600) [4] and Kinetics-400 (K400) [14]. Both datasets consist of 10-second videos of varying resolutions and frame rates, labeled with 600 and 400 action classes.

Models and Implementation. We focus on video recognition models with lower computational requirements to control training time, and since introducing inductive bias becomes increasingly important when efficiency is a factor. We selected the current state-of-the-art models in the lowest categories of FLOPS and memory, and fine-tuned them by applying GRW-smoothing. Specifically, we applied GRW to the MoViNet model family A0,...,A3, and their streaming versions A0-S,...,A2-S. The streaming variants, denoted as Ai-S, are memory-efficient versions of the MoViNet-Ai models that process videos frame by frame. The baseline performance of these streaming models is lower than that of their non-streaming counterparts due to the use of causal operations.

At the lowest end of memory requirements, we applied GRW to MobileNetV3 Small with the following modification. We extract a "base frame" every T frames, where T is the GRW clip window (see Section 3). All frames within a T-frame window are processed individually, but alongside their corresponding base frame. To support this, we modified the first layer to accept 6 input channels.

We used Final Layer GRW-smoothing (see Figure 4), which produced better results than Intermediate Layer GRW-smoothing (Figure 3). However, both methods improve accuracy results, as discussed in Section 4.2. We maintain the original model's final layer output dimension in the normalization affine transformation and apply GRW-smoothing on the output. We replace the classification head with a 2-layer vanilla transformer with a standard $\times 4$ MLP expansion factor (see Section 3 for details).

We set $\lambda=10^{-1}$ as the balancing factor in Equation (7) and $\alpha=\frac{1}{2}$ as the scaling factor in Equation (6). We found that the results are robust to perturbing the values of λ , which suggests that gradients in the direction of smooth solutions align with gradients with respect to the classification likelihood; see ablation studies in Subsection 4.2.

In all experiments we set the GRW window to span 0.5–1.0 s of video. Specifically, for MoViNet-A0/-A1/-A2-GRW and MobileNetV3-S-GRW we use 5 fps with T=5 (covering 1 s), and for MoViNet-A3-GRW we use 12 fps with T=6 (covering 0.5 s). For these values of T we enumerate the full set of orderings in Equation (6), of size (T-1)! (i.e., 24 for T=5 and 120 for T=6), so no permutation subsampling was required (i.e., t was not used).

Remark 1. GRW-smoothing operates on per-frame embeddings $Z=(\mathbf{z}_t)$ rather than raw frames, so the added computation is small relative to the backbone. For example, in MoViNet-A2-S-GRW the input frame $\mathbf{x}_t \in \mathbb{R}^{224 \times 224 \times 3}$ (150,528 dims) is mapped to $z_t \in \mathbb{R}^{640}$, i.e., an $\approx 235 \times$ reduction in dimension. Consequently, the $O(k \cdot T)$ vector operations and log-probability evaluations in Equation (6) are negligible compared to the early convolutional blocks. Empirically, with the T values used in this section (full enumeration of permutations), the wall-clock time per training epoch for MoViNet-A2-S-GRW is only 2% higher than MoViNet-A2-S without smoothing; even for larger T where one might sample up to k=1000 permutations, the overall cost remains dominated by the backbone forward/backward pass.

Training. On K600, we fine-tune starting from existing weights. On K400, we use transfer learning from K600. We employ a simple training process, not applying augmentations except when training

the A2 and A3 models on K400. We use different training rates for the transformer head and model backbone, decreasing with a cosine learning rate scheduler in the range $[10^{-4}, 10^{-6}]$ for the model backbone and $[10^{-3}, 10^{-5}]$ for the transformer head. We fine-tune for 14 epochs on K600 and 10 epochs on K400.

The smaller models, A0–A1 and MobileNet, were trained on a single dgx-A100 for 3–5 days, while the A2 and A3 models were trained on $2 \times dgx$ -A100 for 5 days.

4.1 Results

For all results, we use a single clip evaluation for our models. We report Top-1 accuracy results against FLOPS and memory. The resolution column refers to the resolution of the input video, with 224 indicating 224×224 . The frames column is given by $num_clips \times num_frames$ used in the evaluation. The GFLOPs column indicates the *total* computation for the evaluation of a single video sample.

As seen in Table 1, all models with GRW-smoothing achieve significant improvements in accuracy and set new SOTA results in their corresponding GFLOPs group. Specifically, MoViNet-A0-S-GRW, MoViNet-A1-S-GRW, MoViNet-A2-S-GRW and MoViNet-A3-GRW improve the SOTA results by 6.1%, 5.2%, 4.7% and 3.8%, respectively. For MoViNet-A3-GRW, the next model achieving similar accuracy, MViTv2-B-32×3, requires $18.3\times$ more GFLOPs.

We compare current SOTA memory-efficient models, namely MobileNet and streaming versions of the MoViNet model family, before and after smoothing them with GRW; see Table 2. MobileNetV3-S-GRW, MoViNet-A0-S-GRW, MoViNet-A1-S-GRW, MoViNet-A2-S-GRW improve their non-smooth versions by 6.0%, 6.4%, 5.5% and 4.9%, respectively; see also Figure 1(right).

4.2 Ablation Studies

We study the effect of GRW by (i) disentangling the contribution of smoothing versus the added attention layers, and (ii) analyzing sensitivity to the key hyperparameters: the GRW window T, the scaling factor α in Equation (6), and the balance λ in Equation (7).

Placement and attention vs. smoothing. We ablate where GRW is applied (Intermediate vs. Final layer) and whether the attention head alone explains the gains. Concretely, on K600 we train: (i) MoViNet-A2-S-GRW (Final Layer Smoothing; Figure 4), (ii) MoViNet-A2-S-GRW (Intermediate Layer Smoothing; Figure 3), and (iii) MoViNet-A2-S + attention, the baseline equipped with the same 2-layer Transformer head but *without* GRW. Table 4 shows that adding attention alone yields a small gain over the baseline (+0.9 Top-1), while training with GRW yields an additional +4.0 absolute (for a total +4.9 over the baseline) when applied at the final layer; applying GRW at an intermediate layer also improves accuracy (+2.4) even without attention.

Sensitivity to T, α , and λ . We further ablate the GRW hyperparameters on MoViNet-A0-S-GRW trained on K600, varying one parameter at a time and keeping all other settings fixed (Final Layer Smoothing; training protocol as in Sec. 4). Results are summarized in Table 5.

Summary. Across these ablations, GRW is not overly sensitive near the settings used in our main results: $T{=}5$, $\alpha{=}\frac{1}{2}$ in Equation (6), and $\lambda{=}10^{-1}$ in Equation (7). Short windows under-exploit temporal coherence, very long windows may blur distinct motions, and excessively large λ can over-regularize.

5 Conclusion

In summary, we introduced a novel smooth regularization technique designed to enhance temporal understanding in video recognition models, particularly lightweight architectures. By modeling the evolution of frame embeddings as a Gaussian Random Walk, our method penalizes abrupt representational changes, effectively promoting low-acceleration solutions that align with the natural temporal coherence of video data. This approach has demonstrated significant accuracy improvements, notably a 3.8%–6.4% gain on Kinetics-600, and has established new state-of-the-art performance in compute-constrained settings. By combining our proposed GRW regularization with models such as

Table 1: K600 by FLOPS

Model Top-1 GFLOPs RES FRA					
1	MoViNet-A0-S-GRW	78.4	2.7	172	1 × 50
2	MoViNet-A0[15]	72.3	2.7	172	1×50
3	MobileNetV3-S[11]	61.3	2.8	224	1×50
4	MobileNetV3-S+TSM[21]	65.5	2.8	224	1×50
5	X3D-XS[8]	70.2	3.9	182	1×20
6	MoViNet-A1-S-GRW	81.9	6.0	172	1 × 50
7	MoViNet-A1[15]	76.7	6.0	172	1×50
8	X3D-S[8]	73.4	7.8	182	1×40
9	MoViNet-A2[15]	78.6	10.3	224	1 × 50
10	MobileNetV3-L[11]	68.1	11.0	224	1×50
11	MobileNetV3-L+TSM[21]	71.4	11.0	224	1×50
12	MoViNet-A2-S-GRW	83.3	11.3	224	1×50
13	X3D-M[8]	76.9	19.4	256	1×50
14	X3D-XS[8]	72.3	23.3	182	30×4
15	MoViNet-A3-GRW	85.6	56.4	256	1 × 120
16	MoViNet-A3[15]	81.8	56.9	256	1×120
17	X3D-S[8]	76.4	76.1	182	30×13
18	X3D-L[8]	79.1	77.5	356	1×50
19	MoViNet-A4[15]	83.5	105	290	1×80
20	UniFormer-S[18]	82.8	167	224	4 × 16
21	X3D-M[8]	78.8	186	256	30×16
22	X3D-L[8]	80.7	187	356	1×120
23	I3D[4]	71.6	216	224	1×250
24	MoViNet-A5[15]	84.3	281	320	1×120
25	MViT-B-16×4[6]	82.1	353	224	5×16
26	MoViNet-A6[15]	84.8	386	320	1×120
27	UniFormer-B[18]	84.0	389	224	4×16
28	$XViT (8\times)[3]$	82.5	425	224	3×8
29	$XViT (16\times)[3]$	84.5	850	224	3×16
30	MViT-B-32×3[6]	83.4	850	224	5×32
31	MViTv2-B-32×3[20]	85.5	1030	224	5 × 32

Table 1: Top-1 accuracy, total video evaluation cost (in GFLOPs), input resolution (RES), and FRAMES = clips \times frames per clip used for evaluation on Kinetics-600. Models enhanced with our proposed smooth regularization are marked with GRW. These models consistently outperform their baselines and other state-of-the-art methods under similar FLOP constraints. Variance: for MoViNet-A0-S-GRW, across three seeds we obtain 78.4 \pm 0.05 Top-1 (mean \pm std).

MoViNet-A0/1/2/3 and their streaming counterparts, as well as MobileNetV3, we improve overall performance within their respective FLOP and memory constraints. The GRW regularization acts as a flexible, plug-and-play component with minimal computational overhead, guiding networks towards more stable and generalizable feature representations.

While our approach presents promising results, it has certain limitations. The core assumption of Gaussian Random Walk dynamics, while beneficial for many natural videos, might not be universally optimal for content characterized by extremely abrupt transitions or intentionally discontinuous motion. A future work may explore such videos and how to extend GRW for such cases. Furthermore, while our experiments demonstrate significant gains on lightweight models, the extent of improvement on very large-capacity models, which we could not do due to computational constraints, require further investigation. Finally, the necessity of the frame ordering component in the contrastive loss, while effective in preventing degenerate solutions, does introduce an additional layer of complexity to the training objective. A more efficient variants can be studied in a future work.

Looking ahead, several avenues for future research emerge. Extending the application of our GRW smoothing to a wider array of video architectures, including more complex Transformer-based models, could yield further insights into its generalizability. Investigating its efficacy across diverse video understanding tasks beyond action recognition, such as temporal action localization or video

Table 2: K600 by Mem

	- 5	
Model	Top-1	Mem MB
MobileNetV3-S[11]	61.3	29
MobileNetV3-S-GRW	67.3	30
MoViNet-A0-S-GRW	78.4	53
MoViNet-A0-S[15]	72.0	53
MoViNet-A1-S-GRW	81.9	67
MoViNet-A1-S[15]	76.4	67
MoViNet-A2-S-GRW	83.3	78
MoViNet-A2-S[15]	78.4	78

Table 2: Top-1 accuracy and memory usage (in MB) for memory-efficient models on Kinetics-600. Models enhanced with our smooth regularization (GRW) are shown in bold and consistently outperform their baselines under identical memory constraints.

Table 3: K400 by FLOPS

Model	Top-1	GFLOPs
MoViNet-A0-S-GRW	70.4	2.7
MoViNet-A0[15]	65.8	2.7
MoViNet-A2[15]	75.0	10.3
MoViNet-A2-GRW	77.6	11.3
X3D-XS[8]	69.5	23.3
MoViNet-A3-GRW	81.7	56.4
MoViNet-A3[15]	78.2	56.9
X3D-S[8]	73.5	76.1
VideoMamba[17]	76.9	108

Table 3: Top-1 accuracy, total video evaluation cost, on Kinetics-400.

Table 4: Ablation on placement (K600, MoViNet-A2-S family).

Model	Top-1	GFLOPs
MoViNet-A2-S-GRW (final layer)	83.3	11.3
MoViNet-A2-S-GRW (intermediate layer)	80.8	10.3
MoViNet-A2-S + attention (no GRW)	79.3	11.3
MoViNet-A2-S[15]	78.4	10.3

anomaly detection, presents another promising direction. Additionally, a dynamic smoothing window that adapts to video content is favorable. Finally, a more in-depth theoretical understanding of how GRW regularization influences the optimization landscape and feature learning process would be beneficial.

6 Acknowledgements

Parts of this research were conducted using ORCHARD, a high-performance cloud computing cluster. The authors would like to acknowledge Carnegie Mellon University for making this resource available to its community.

This material is based upon work supported by the United States Navy under award number N00174-23-1-0001 and by the National Science Foundation under grant number CNS-2106862. The content of the information does not necessarily reflect the position or the policy of the government and no official endorsement should be inferred. This work was done in the CMU Living Edge Lab, which is supported by Intel, Arm, Vodafone, Deutsche Telekom, CableLab, Crown Castle, InterDigital, Seagate, Microsoft, the VMware University Research Fund, IAI, and the Conklin Kistler family fund. Any opinions, findings, conclusions or recommendations expressed in this document are those of the authors and do not necessarily reflect the view(s) of their employers or the above funding sources.

This work was partially supported by a grant from The Center for AI and Data Science at Tel Aviv University (TAD).

(a) Window T		(b) α in Equation (6)			(c) λ in Equation (7)	
Top-1 (%)	T	Top-1 (%)	α	_	Top-1 (%)	λ
77.3	3	77.9	0.25	_	78.0	0.01
78.4	5	78.4	0.5		78.4	0.1
78.0	10	78.3	1.0		75.5	1.0
72.0	no smoothing	72.0	no smoothing		72.0	no smoothing

Table 5: Hyperparameter ablations for GRW on K600 with MoViNet-A0-S-GRW. (a) Window T peaks at T=5. (b) Scaling α shows a mild optimum near α =0.5 in Equation (6). (c) Balance λ is robust in $\{0.01, 0.1\}$ and degrades at λ =1.0 in Equation (7).

References

- [1] Anurag Arnab, Mostafa Dehghani, Georg Heigold, Chen Sun, Mario Lučić, and Cordelia Schmid. Vivit: A video vision transformer. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 6836–6846, October 2021.
- [2] Gedas Bertasius, Heng Wang, and Lorenzo Torresani. Is space-time attention all you need for video understanding? In *ICML*, 2021.
- [3] Adrian Bulat, Juan Manuel Perez Rua, Swathikiran Sudhakaran, Brais Martinez, and Georgios Tzimiropoulos. Space-time mixing attention for video transformer. *Advances in neural information processing systems*, 34:19594–19607, 2021.
- [4] Joao Carreira, Eric Noland, Andras Banki-Horvath, Chloe Hillier, and Andrew Zisserman. A short note about kinetics-600. *arXiv preprint arXiv:1808.01340*, 2018.
- [5] Joao Carreira and Andrew Zisserman. Quo Vadis, Action Recognition? A New Model and the Kinetics Dataset. In 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pages 4724–4733, 2017.
- [6] Haoqi Fan, Bo Xiong, Karttikeya Mangalam, Yanghao Li, Zhicheng Yan, Jitendra Malik, and Christoph Feichtenhofer. Multiscale vision transformers. In *Proceedings of the IEEE/CVF* international conference on computer vision, pages 6824–6835, 2021.
- [7] Mohsen Fayyaz, Emad Bahrami, Ali Diba, Mehdi Noroozi, Ehsan Adeli, Luc Van Gool, and Jurgen Gall. 3d cnns with adaptive temporal feature resolutions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4731–4740, 2021.
- [8] Christoph Feichtenhofer. X3d: Expanding architectures for efficient video recognition. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 203–213, 2020.
- [9] Christoph Feichtenhofer, Haoqi Fan, Jitendra Malik, and Kaiming He. Slowfast networks for video recognition. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2019.
- [10] Basura Fernando, Hakan Bilen, Efstratios Gavves, and Stephen Gould. Self-supervised video representation learning with odd-one-out networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3636–3645, 2017.
- [11] Andrew Howard, Mark Sandler, Grace Chu, Liang-Chieh Chen, Bo Chen, Mingxing Tan, Weijun Wang, Yukun Zhu, Ruoming Pang, Vijay Vasudevan, et al. Searching for mobilenetv3. In Proceedings of the IEEE/CVF international conference on computer vision, pages 1314–1324, 2019.
- [12] De-An Huang, Vignesh Ramanathan, Dhruv Mahajan, Lorenzo Torresani, Manohar Paluri, Li Fei-Fei, and Juan Carlos Niebles. What makes a video a video: Analyzing temporal information in video understanding models and datasets. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7366–7375, 2018.
- [13] Ziyuan Huang, Shiwei Zhang, Liang Pan, Zhiwu Qing, Mingqian Tang, Ziwei Liu, and Marcelo H Ang Jr. Tada! temporally-adaptive convolutions for video understanding. *arXiv* preprint arXiv:2110.06178, 2021.
- [14] Will Kay, Joao Carreira, Karen Simonyan, Brian Zhang, Chloe Hillier, Sudheendra Vijayanarasimhan, Fabio Viola, Tim Green, Trevor Back, Paul Natsev, et al. The kinetics human action video dataset. *arXiv preprint arXiv:1705.06950*, 2017.
- [15] Dan Kondratyuk, Liangzhe Yuan, Yandong Li, Li Zhang, Mingxing Tan, Matthew Brown, and Boqing Gong. Movinets: Mobile video networks for efficient video recognition. In *Proceedings* of the IEEE/CVF conference on computer vision and pattern recognition, pages 16020–16030, 2021.

- [16] Hsin-Ying Lee, Jia-Bin Huang, Maneesh Singh, and Ming-Hsuan Yang. Unsupervised representation learning by sorting sequences. In *Proceedings of the IEEE international conference on computer vision*, pages 667–676, 2017.
- [17] Kunchang Li, Xinhao Li, Yi Wang, Yinan He, Yali Wang, Limin Wang, and Yu Qiao. Video-mamba: State space model for efficient video understanding. In ECCV, 2024.
- [18] Kunchang Li, Yali Wang, Gen Peng, Guannan Song, Yu Liu, Hongsheng Li, and Yu Qiao. UniFormer: Unified transformer for efficient spatial-temporal representation learning. In *ICLR*, 2022.
- [19] Yan Li, Bin Ji, Xintian Shi, Jianguo Zhang, Bin Kang, and Limin Wang. Tea: Temporal excitation and aggregation for action recognition. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 909–918, 2020.
- [20] Yanghao Li, Chao-Yuan Wu, Haoqi Fan, Karttikeya Mangalam, Bo Xiong, Jitendra Malik, and Christoph Feichtenhofer. Mvitv2: Improved multiscale vision transformers for classification and detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern* recognition, pages 4804–4814, 2022.
- [21] Ji Lin, Chuang Gan, and Song Han. Tsm: Temporal shift module for efficient video understanding. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 7083–7093, 2019.
- [22] Ze Liu, Jia Ning, Yuxin Cao, Yixuan Wei, Zheng Zhang, Stephen Lin, and Han Hu. Video swin transformer. In CVPR, 2022.
- [23] Ishan Misra, C Lawrence Zitnick, and Martial Hebert. Shuffle and learn: unsupervised learning using temporal order verification. In *European conference on computer vision*, pages 527–544. Springer, 2016.
- [24] Mandela Patrick, Desmond Campbell, Yuki Asano, Ishan Misra, Florian Metze, Christoph Feichtenhofer, Andrea Vedaldi, and João F. Henriques. Keeping your eye on the ball: Trajectory attention in video transformers. In *NeurIPS*, 2021.
- [25] Oran G. Sharir, Asher Noy, and Lihi Zelnik-Manor. An image is worth 16x16 words, what is a video worth? In *arXiv:2103.13915*, 2021.
- [26] Karen Simonyan and Andrew Zisserman. Two-stream convolutional networks for action recognition in videos. In *Advances in Neural Information Processing Systems*, volume 27. Curran Associates, Inc., 2014.
- [27] Du Tran, Lubomir Bourdev, Rob Fergus, Lorenzo Torresani, and Manohar Paluri. Learning spatiotemporal features with 3d convolutional networks. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2015.
- [28] Limin Wang, Bingkun Huang, Zhiyu Zhao, Zhan Tong, Yinan He, Yi Wang, Yali Wang, and Yu Qiao. Videomae v2: Scaling video masked autoencoders with dual masking. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 14549–14560, 2023.
- [29] Limin Wang, Zhan Tong, Bin Ji, and Gangshan Wu. Tdn: Temporal difference networks for efficient action recognition. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 1895–1904, 2021.
- [30] Limin Wang, Yuanjun Xiong, Zhe Wang, Yu Qiao, Dahua Lin, Xiaoou Tang, and Luc Van Gool. Temporal segment networks: Towards good practices for deep action recognition. In *ECCV*, pages 20–36, 2016.
- [31] Laurenz Wiskott and Terrence J Sejnowski. Slow feature analysis: Unsupervised learning of invariances. *Neural computation*, 14(4):715–770, 2002.
- [32] Daijun Zhang, Kunchang Li, Yi Wang, and et al. Morphmlp: An efficient mlp-like backbone for spatial-temporal representation learning. In *ECCV*, 2022.

A Scaling

Here we show the scaling behavior of GRW-smoothing. The main result we prove is that the optimal solution for a smoothing window of size T, applied to approximately centered data, lies within a ball of radius bounded by $\mathcal{O}(T\sqrt{\ln T})$. We prove this result in the one-dimensional case.

Let us recall the setup:

Consider $T \geq 3$ points in [0, R] with fixed endpoints

$$0 =: z_1 \le z_2 \le \dots \le z_T := R, \qquad Z = (z_t)_{t-1}^T, \tag{8}$$

and define velocity and acceleration vectors as

$$V(Z) = (z_2 - z_1, \dots, z_T - z_{T-1}) \in \mathbb{R}^{T-1}, \qquad v = (v_t)_{t=1}^{T-1},$$

$$A(Z) = (v_2 - v_1, \dots, v_{T-1} - v_{T-2}) \in \mathbb{R}^{T-2}, \qquad A = (a_t)_{t=1}^{T-2}.$$

We will denote by $\mathcal{Z}^T \subset \mathbb{R}^T$ the set of all point configurations of the form (8), that is, non-decreasing sequences with $z_1 = 0$. For any $Z \in \mathcal{Z}^T$ we denote $R(Z) \coloneqq z_T$.

For any such configuration Z, define the loss components:

$$\mathcal{L}_v(Z) = \frac{1}{2} \sum_{t=1}^{T-1} v_t^2 = \frac{1}{2} \sum_{t=1}^{T-1} (z_{t+1} - z_t)^2,$$

and

$$\mathcal{L}_a(Z) = -\log \frac{\exp\left(-\frac{1}{2}\sum_{t=1}^{T-2} a_t^2(Z)\right)}{\sum_{\pi} \exp\left(-\frac{1}{2}\sum_{t=1}^{T-2} a_t^2(Z^{\pi})\right)},$$

where the sum is over all permutations π of $\{2,\ldots,T\}$ fixing z_1 , and

$$Z^{\pi} = (z_1, z_{\pi(2)}, \dots, z_{\pi(T)}).$$

With the above notation and $\alpha = 1$ (see Equation (6) in the paper), GRW loss is given by

$$\mathcal{L}(Z) = \mathcal{L}_a(Z) + \mathcal{L}_v(Z).$$

Theorem 1 (GRW-smoothing scale). Given $T \geq 3$, let $Z^* = \arg\min_{Z \in \mathcal{Z}^T} \mathcal{L}(Z)$. Then

$$R(Z^*) = \mathcal{O}(T\sqrt{\ln T}).$$

Remark 2. Theorem 1 suggests that for approximately centered embedding, it is beneficial to allow a scaling degree of freedom before feeding embedding vectors into GRW-smoothing to allow the solution to converge into the $\mathcal{O}(T\sqrt{\ln T})$ ball. Therefore, when applying GRW-smoothing to a sequence model embedding, we either directly scale and center the embedding using a non-learnable batch normalization (BN), or allow for a learnable linear transformation to set the scaling. Empirically, we found that centering is not required when a linear transformation is applied.

Proof. We will use the following propositions in the proof, providing their proofs subsequently:

Proposition 1 (Uniform Lower Bound on \mathcal{L}). For any $Z \in \mathcal{Z}^T$,

$$\mathcal{L}(Z) \ge \mathcal{L}_v(Z) \ge \frac{R^2(Z)}{2(T-1)}.$$
(9)

Proposition 2 (Uniform Configuration Upper Bound). Consider the uniform configuration

$$Z_u = \left(0, \frac{R}{T-1}, \frac{2R}{T-1}, \dots, R\right).$$
 (10)

For R = T - 1,

$$\mathcal{L}(Z_u) = \mathcal{O}(T \ln T). \tag{11}$$

Assuming Proposition 1 and Proposition 2, we will now complete the proof of Theorem 1 and provide their proofs subsequently.

From Proposition 1 and Proposition 2, it follows that

$$\frac{R^2(Z^*)}{2(T-1)} \le \mathcal{L}(Z^*) \le C_1 T \ln T,$$

for some constant C_1 . Rearranging, we get the desired result.

Now we prove the auxiliary claims.

Proof of Proposition 1. Since $\mathcal{L}_a(Z)$ is non-negative, we have $\mathcal{L}(Z) \geq \mathcal{L}_v(Z)$. Recall that $\mathcal{L}_v(Z) = \frac{1}{2} \sum_{t=1}^{T-1} v_t^2$. For any Z the velocities are non-negative and satisfy $\sum_t v_t = z_T - z_1 = R(Z) \coloneqq R$. Therefore, $\mathcal{L}_v(Z) \geq \min_{V \in \mathbb{R}^{T-1}, V \geq 0, \|V\|_1 = R} \frac{1}{2} \|V\|_2^2$. The last is a classic quadratic program with the minimizer $V_u = \frac{R}{T-1}(1,...,1)$, attaining the minimum $\frac{R^2}{2(T-1)}$, where these velocities are realized by the uniform configuration of the points. Hence, we obtain $\mathcal{L}(Z) \geq \mathcal{L}_v(Z) \geq \frac{R^2(Z)}{2(T-1)}$.

Proof of Proposition 2. Fix R = T - 1 and consider the uniform configuration (10). We have

$$\mathcal{L}(Z_u) = \underbrace{\ln\left(\sum_{\pi:\pi(1)=1} \exp\left(-\frac{R^2}{2(T-1)^2}S(\pi)\right)\right)}_{\mathcal{L}_u(Z_u)} + \underbrace{\frac{R^2}{2(T-1)}}_{\mathcal{L}_v(Z_u)},$$

where

$$S(\pi) := \sum_{t=1}^{T-2} (\pi(t+2) - 2\pi(t+1) + \pi(t))^2.$$

The velocity term simplifies as

$$\mathcal{L}_v(z_u) = \frac{R^2}{2(T-1)} = \frac{(T-1)^2}{2(T-1)} = \frac{T-1}{2}.$$
 (12)

The acceleration term simplifies as

$$\mathcal{L}_a(z_u) = \ln\left(\sum_{\pi:\pi(1)=1} \exp\left(-\frac{1}{2}S(\pi)\right)\right) \le \ln((T-1)!) = \mathcal{O}(T\ln T). \tag{13}$$

Then by (12) and (13) $\mathcal{L}(z_u) = \mathcal{O}(T \ln T)$

NeurIPS Paper Checklist

1. Claims

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

Answer: [Yes]

Justification: Throughout the paper, we describe the proposed method and show through various experiments its advantage in the constrained compute setting.

Guidelines:

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

2. Limitations

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

Answer: [Yes]

We explicitly discuss the limitations of our work in Section 5.

Guidelines:

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

3. Theory assumptions and proofs

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

Answer: [NA]

Justification: There are no theoretical results shown in the paper.

Guidelines:

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

4. Experimental result reproducibility

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

Answer: [Yes]

Justification: We provide in Section 4 detailed descriptions of all experimental settings, including model architectures, initialization procedures, hyperparameters, and evaluation metrics. This information is sufficient to reproduce the main experimental results that support our claims.

Guidelines:

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

5. Open access to data and code

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

Answer: [Yes]

Justification: We will publish our code and instructions upon acceptance, enabling faithful reproduction of all experiments. All datasets used are publicly available.

Guidelines:

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

6. Experimental setting/details

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

Answer: [Yes]

Justification: We provide full training and evaluation details in Section 4. Where possible, we adopt hyperparameter settings from prior work; otherwise, we perform grid search. These details are specified to ensure clarity and reproducibility.

Guidelines:

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

7. Experiment statistical significance

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

Answer: [Yes]

Justification: Training our models is very computationally demanding and therefore we cannot afford training multiple models. Yet, the advantage they achieve are with very high margin, which clearly shows that the improvement is indeed significant and not marginal.

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.

- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error
 of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments compute resources

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

Answer: [Yes]

Justification: We report the GPU used and training time in Section 4

Guidelines:

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

9. Code of ethics

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

Answer: [Yes]

Justification: We followed the code of ethics while conducting the research.

Guidelines:

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

10. Broader impacts

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

Answer: [NA]

Justification: The paper provide a regularization for the task of video action recognition, which makes action recognition video models more efficient.

Guidelines:

• The answer NA means that there is no societal impact of the work performed.

- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

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

Answer: [NA]

Justification: The proposed algorithm is not of high risk of misuse.

Guidelines:

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

12. Licenses for existing assets

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

Answer: [Yes]

Justification: We provide reference to all the models and datasets we use in Section 4.

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.

- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. New assets

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

Answer: [NA]

Justification: We use existing datasets.

Guidelines:

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

14. Crowdsourcing and research with human subjects

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

Answer: [NA]

Justification: We did not use crowdsourcing nor research with human subjects.

Guidelines:

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

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

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

Answer: [NA]

Justification: We did not use crowdsourcing nor research with human subjects.

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.

- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. Declaration of LLM usage

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

Answer: [NA]

Justification: We use LLMs only in their standard use of improving writing.

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