
UNIEGO: Proxies as Mediators for Unified Egocentric Video Representation Learning

Wenhao Chi, Arkaprava Sinha, Dominick Reilly, Hieu Le, Srijan Das
University of North Carolina at Charlotte

Abstract

Egocentric video understanding is inherently limited by the narrow perspective of wearable cameras: a single viewpoint, a single modality, a single model cannot capture the full richness of human action. We argue that a truly expressive egocentric representation must subsume complementary knowledge across viewpoints, modalities, and foundation model representations, yet remain deployable from egocentric video alone. To this end, we introduce a hierarchical multi-teacher distillation framework that produces UNIEGO, a unified egocentric encoder trained with nine teachers spanning ego-exo viewpoints, RGB, depth, and skeleton modalities, and four foundation models. Rather than distilling directly from heterogeneous teachers whose incompatible architectures and feature geometries induce conflicting gradients, our framework interposes a layer of representation-specific *Proxy* models that translate diverse teacher knowledge into a homogeneous egocentric space. A second distillation stage, **Selective Proxy Distillation (SPD)**, then adaptively selects, for each training sample, the subset of proxies that are both correct and confident, distilling exclusively from reliable supervision and suppressing erroneous signals. SPD is further stabilized by initializing UNIEGO as a learned convex combination of proxy parameters, placing the unified model in a well-conditioned region of the loss landscape before distillation begins. UNIEGO achieves state-of-the-art performance across three egocentric video understanding tasks - action recognition, video retrieval, and action segmentation on three challenging ego-exo benchmarks, outperforming naive multi-teacher distillation baselines and demonstrating that structured, proxy-mediated knowledge transfer yields richer and more discriminative egocentric representations. We release code and models at <https://github.com/Wenhao-Chi/UNIEGO>.

1 Introduction

Understanding human actions from egocentric video is a fundamental challenge in visual perception, with broad applications spanning augmented reality, assistive robotics, and procedural activity analysis [45, 7, 42]. Yet, learning a truly expressive egocentric representation from a single model remains elusive. Wearable cameras impose a narrow field of view and suffer from persistent self-occlusions, obscuring the actor’s body and surrounding scene context [13]. Complementary modalities such as depth and skeleton which encode the geometric structure of human motion are discarded entirely [49, 30, 4]. Moreover, the rich, diverse representational knowledge encapsulated within large-scale foundation models (FMs) remains untapped [52, 5, 50, 17]. Consequently, a standalone egocentric model is fundamentally limited. We argue that overcoming these limitations requires a *unified egocentric representation*, a single, comprehensive embedding that subsumes complementary knowledge across modalities (RGB, depth, skeleton), viewpoints (egocentric and exocentric), and diverse representations from heterogeneous foundation models, yet operates solely on egocentric video at inference. Such a representation would unlock richer action understanding across a broad spectrum of downstream egocentric tasks.

These diverse perceptual signals spanning viewpoints, modalities, and FM representations naturally suggest a multi-teacher knowledge distillation framework, a paradigm explored extensively in concurrent work [34, 40, 46, 1]. However, existing methods distill homogeneous representations into a single student, sidestepping the deeper structural challenges that arise in the egocentric setting. First, teachers here are fundamentally heterogeneous: skeleton-based models operate over graph-structured neural architectures incompatible with video encoders [9, 8], while exocentric RGB teachers encode scene geometry from an entirely different viewpoint, introducing a substantial “*representational gap*” [27, 21]. Second, naïvely forcing an egocentric student to reconcile these incompatible feature spaces simultaneously leads to “*conflicting gradient signals*” and degrades optimization [25, 51, 41]. Thus, the ego student is asked not merely to learn, but to simultaneously bridge modality gaps, close viewpoint gaps, and absorb diverse representational priors, an ill-posed objective for a single distillation stage. This motivates a more principled approach to unified egocentric representation learning.

We address these challenges by introducing a hierarchical multi-teacher distillation framework that consolidates diverse perceptual signals across viewpoints, modalities, and FM representations into **UNIEGO**, a *single unified egocentric encoder* (see Figure 1). Rather than distilling directly from heterogeneous teachers, our framework first translates each teacher’s knowledge into a student-compatible space through a set of representation-specific *Proxy* models. Each proxy shares the architecture of UNIEGO and operates on egocentric video. Therefore, the proxy learning diminishes the representational gap induced by heterogeneous teacher architectures and viewpoints by converting incompatible feature geometries into a homogeneous egocentric embedding space.

Further, the first distillation level naturally exposes the reliability of each supervision signal on a per-instance basis: a proxy that cannot correctly classify a given sample carries no trustworthy knowledge to transfer. Our proposed hierarchical framework exploits this signal through Selective Proxy Distillation (SPD), which selects, for each training sample, the subset of proxies whose predictions are both correct and confident, distilling exclusively from this reliable subset and suppressing erroneous supervision entirely. To further stabilize optimization, SPD is initialized via a learned convex combination of proxy parameters, placing UNIEGO in a flat, well-conditioned region of the loss landscape prior to distillation. Together, proxy merging and proxy selection collectively mitigate the conflicting gradient problem inherent to naive multi-teacher distillation.

The outcome of our hierarchical distillation framework is UNIEGO, a unified egocentric encoder trained with 9 teachers spanning ego-exo viewpoints, RGB, depth, and skeleton modalities, and four FMs. UNIEGO outperforms naive multi-teacher distillation baselines across three egocentric video understanding tasks - action recognition, video retrieval, and action segmentation on three challenging ego-exo benchmarks [38, 22, 14]. Moreover, UNIEGO generalizes across video backbone architectures, including compact models with as few as 22M parameters. We summarize our contributions as follows:

- We introduce **UNIEGO**, a unified egocentric encoder trained via a novel hierarchical distillation framework with *nine* teachers across *ego-exo* viewpoints, *three* modalities, and *four* foundation models, using representation-specific proxies as structured mediators.
- We propose three tightly coupled components: **(i) Proxy Learning**, which converts heterogeneous, multi-modal teacher supervision into a pool of architecturally homogeneous egocentric proxies, bridging the modality and viewpoint gap; **(ii) Selective Proxy Distillation (SPD)**, a sample-wise selective distillation mechanism that dynamically routes supervision from the most reliable proxies, mitigating conflicting gradients across heterogeneous teachers; and **(iii) Proxy Merging**, a learned convex initialization of UNIEGO that places the unified model in a well-conditioned region of the loss landscape, stabilizing SPD’s optimization.
- UNIEGO achieves **state-of-the-art performance** across three egocentric video understanding tasks on three challenging benchmarks, demonstrating that hierarchical proxy-mediated distillation yields richer, more discriminative egocentric representations than direct multi-teacher supervision.

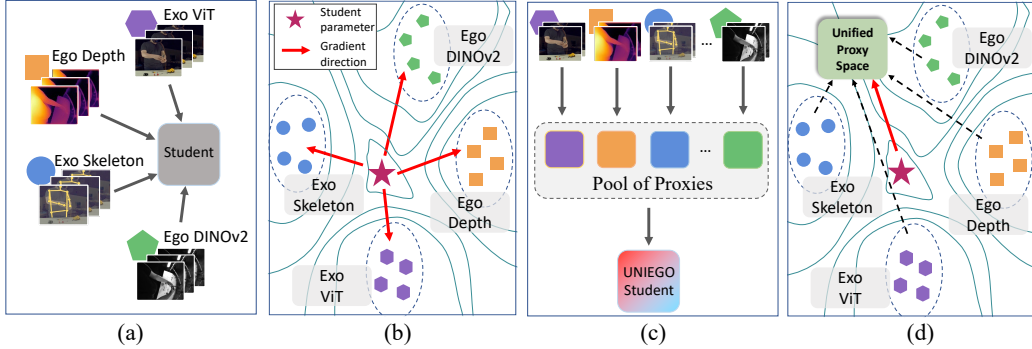


Figure 1: (a) Naive multi-teacher distillation with heterogeneous teachers for learning unified egocentric representations results in representational gaps and conflicting gradients, as illustrated in (b). (c) In contrast, our proposed UNIEGO adopts a *hierarchical distillation framework* that mitigates these limitations through proxy-mediated learning, as shown in (d). Black dashed arrows illustrate the effects of this framework, shifting teachers into a unified representation space.

2 Related Work

2.1 Egocentric Representation Learning

Over the past few years, egocentric representation learning has become a central problem in video understanding. Early works, such as EgoVLP [24] and LaViLa [54], focused on learning egocentric representations from egocentric videos alone [7, 13, 42, 31]. This is challenging: egocentric cameras move with the person, hands and objects frequently occlude one another, and the same action can look very different depending on the wearer and environment. More recent works augment egocentric representation learning with signals beyond the raw ego video, typically leveraging synchronized exocentric viewpoints [21, 48, 32, 10, 27] or additional modalities [12, 47, 44, 33]. For example, ViewpointRosetta [27] uses diffusion models to learn a mapping between egocentric and exocentric representation spaces, EgoDTM [47] learns 3D-aware egocentric representations through distillation from a depth-modality trained teacher. These methods typically exploit a single specific auxiliary signal (viewpoint/modality). In contrast our work aims to consolidate many heterogeneous teachers, spanning auxiliary viewpoints and modalities, into a single egocentric encoder.

2.2 Multi-teacher knowledge distillation

Knowledge distillation was originally introduced as a way to compress the knowledge of a large model, or an ensemble of models, into a single deployable student [15]. This idea has since been extended to multi-teacher distillation, where a student learns from several teachers rather than a single teacher. For example, approaches like AMTML-KD [26] and CA-MKD [53] adaptively distill from an ensemble of teachers operating on a single input modality, learning instance-level importance weights for each teacher. More recent work has considered a more heterogeneous form of multi-teacher distillation, where teachers are differentiated by their architecture or input modalities. For example, AM-RADIO [34] and Theia [40] both distill multiple vision foundation models with different architectures and pretraining objectives into a single universal model. The key differences between these approaches and UNIEGO are that: (1) we learn egocentric representation learning that has not been explored by existing works, (2) we leverage a more diverse set of teachers spanning various architectures, modalities, and viewpoints; and (3) we selectively choose which teachers to distill from rather than treating all teacher as equally useful.

3 Method: Hierarchical Distillation Framework

In this section, we present UNIEGO, a hierarchical multi-teacher distillation framework for egocentric video understanding. Let $\mathcal{D} = \{(x_i^e, \{x_i^r\}_{r=1}^R, y_i)\}_{i=1}^N$ denote the training set, where x_i^e is the egocentric video clip, x_i^r denotes the input to the r -th teacher model \mathcal{T}_r , and $y_i \in \mathcal{Y}$ is the ground-truth action label. In practice, x_i^r may coincide with x_i^e , correspond to its exocentric counterpart,

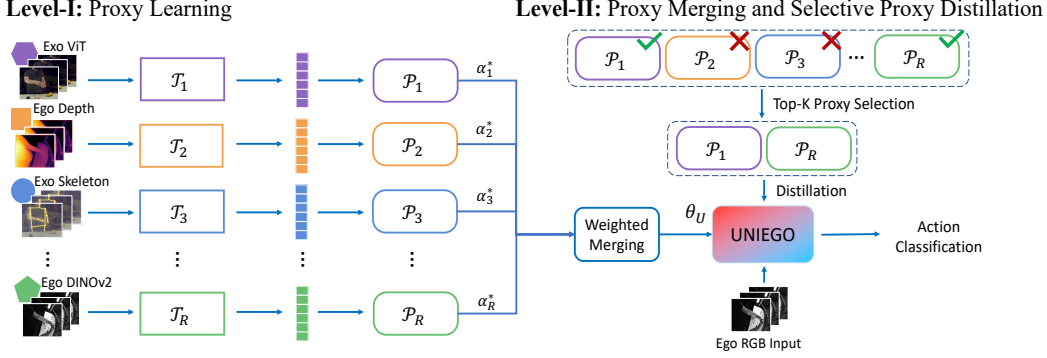


Figure 2: **Overview of UNIEGO.** UNIEGO learns a unified egocentric encoder through a two-level proxy-mediated distillation framework. In Level-I (left), heterogeneous teachers spanning viewpoints, modalities, and foundation representations independently supervise egocentric proxy models, converting diverse teacher signals into a homogeneous proxy space. In Level-II (right), the proxy parameters are first merged to initialize the unified model, after which Selective Proxy Distillation (SPD) performs sample-wise reliability filtering and distills only from proxies that are correct and confident. At inference, only the final UNIEGO model operates using egocentric input.

or constitute an entirely different modality (*e.g.*, skeleton or depth), each paired with its associated teacher \mathcal{T}_r . Each teacher thus encodes a distinct representation, modality, or viewpoint. The objective is to learn a unified egocentric encoder $f(\cdot)$ such that the resulting representation $f(x^e)$ subsumes the complementary knowledge of all teachers $\{\mathcal{T}_r(x^r)\}_{r=1}^R$ during training, while requiring *only* the egocentric stream x^e at inference.

Overview. UNIEGO employs a hierarchical knowledge distillation strategy comprising two levels, as illustrated in Figure 2. In the *first level*, we train a set of representation-specific **Proxy** models by independently distilling knowledge from each teacher \mathcal{T}_r . Each proxy specializes in transferring representation- or modality- or viewpoint-specific information from its corresponding teacher, while bridging the domain gap between heterogeneous teacher representations and the egocentric input space, converting diverse supervisory signals into a homogeneous egocentric proxy space. In the *second level*, we perform Selective Proxy Distillation (SPD), which aggregates knowledge from all proxy models into the final UNIEGO model, initialized via a principled proxy merging strategy. The resulting **Unified Egocentric Model** effectively aggregates complementary knowledge from multiple teachers, with proxies acting as intermediaries for structured knowledge transfer. In the following subsections, we detail each component of UNIEGO’s learning paradigm.

3.1 Level-I: Proxy Learning

In the first distillation level, we train R proxy models $\{P_r\}_{r=1}^R$ by independently distilling feature-level knowledge from each teacher \mathcal{T}_r into an egocentric student. Thus, for an egocentric video x_i^e and its time-synchronized auxiliary inputs $\{x_i^r\}_{r=1}^R$ comprising RGB, depth, and skeleton data from both egocentric and exocentric viewpoints, we extract representations via R teacher networks $\{\mathcal{T}_r\}_{r=1}^R$, each providing viewpoint- and modality-specific supervision signals.

All proxies share the same architecture but are trained with independent parameters. Since some teachers are foundation models that yield only feature embeddings rather than action logits, we adopt feature-level distillation throughout the entirety of this stage. Specifically, let $h_i, z_i = f(x_i^e)$ denote the feature embedding and action logits of the proxy student, and $h_i^r = \mathcal{T}_r(x_i^r)$ the teacher’s feature embedding. Each proxy P_r is optimized via:

$$\mathcal{L}_1^r = \frac{1}{N} \sum_{i=1}^N (\lambda_I D_{\cos}(h_i, h_i^r) + \lambda_{\text{cls}} \text{CE}(z_i, y_i)), \quad (1)$$

where D_{\cos} denotes cosine distance and CE is the cross-entropy loss.

Despite being induced by heterogeneous teachers, all proxies share the same egocentric architecture. This converts heterogeneous teacher supervision into a homogeneous egocentric proxy space $\{P_r\}_{r=1}^R$, partitioned as $\mathcal{P} = \mathcal{P}_{\text{ego}} \cup \mathcal{P}_{\text{exo}}$ according to viewpoint. Each proxy P_r subsequently serves as a

mediator in the second distillation level, reducing the modality and viewpoint gap between the original teachers and the final unified model.

3.2 Level-II: Selective Proxy Distillation (SPD)

Given the learned proxy set $\{P_r\}_{r=1}^R$, the goal of this level is to consolidate their complementary knowledge into a single unified egocentric model $f(\theta_U; \cdot)$. We first introduce a principled initialization of θ_U via proxy merging, followed by **SPD**, which selectively transfers knowledge from the most reliable proxies for each training sample.

Proxy Merging Initialization. We initialize θ_U as an optimally weighted combination of the proxy parameters. Let θ_r denote the parameters of proxy P_r . We solve for merging coefficients $\alpha^* \in \Delta^R$ that minimize the action classification loss over the training set:

$$\theta_U \leftarrow \theta_{\text{merge}}^*, \quad \theta_{\text{merge}}^* = \sum_{r=1}^R \alpha_r^* \theta_r, \quad \alpha^* = \arg \min_{\alpha \in \Delta^R} \frac{1}{N} \sum_{i=1}^N \text{CE} \left(f \left(\sum_{r=1}^R \alpha_r \theta_r; x_i^e \right), y_i \right), \quad (2)$$

where $\Delta^R = \{\alpha \in \mathbb{R}_{\geq 0}^R \mid \sum_{r=1}^R \alpha_r = 1\}$ is the probability simplex and $f(\theta; x_i^e)$ denotes the model with parameters θ evaluated on x_i^e . This initialization places θ_U in a favourable region of the optimization landscape, providing a stable starting point for subsequent distillation.

Proposition 1 (Proxy Merged Initialization as a Loss Upper Bound for UNIEGO) *Let $\mathcal{L}(\theta) = \frac{1}{N} \sum_{i=1}^N \text{CE}(f(\theta; x_i^e), y_i)$ denote the classification loss. Assume \mathcal{L} is convex in θ in a neighbourhood containing $\{\theta_r\}_{r=1}^R$. Then for any $\alpha \in \Delta^R$:*

$$\mathcal{L} \left(\sum_{r=1}^R \alpha_r \theta_r \right) \leq \sum_{r=1}^R \alpha_r \mathcal{L}(\theta_r), \quad (3)$$

and consequently the optimally merged initialization satisfies: $\mathcal{L}(\theta_{\text{merge}}^*) \leq \min_r \mathcal{L}(\theta_r)$.

The first inequality follows directly from Jensen’s inequality [3] applied to the convex loss \mathcal{L} . The second follows by noting that α^* is chosen to minimize $\mathcal{L}(\sum_r \alpha_r \theta_r)$, and the degenerate solution $\alpha_r = 1$ for any single r is feasible in Δ^R , so the optimum is no worse than the best individual proxy. This establishes that θ_{merge}^* achieves lower classification loss than any individual proxy under local convexity, placing UNIEGO in a flatter, better-generalizing region of the loss landscape [16, 11]. This favorable initialization reduces the optimization burden of SPD, as distillation begins from a point that already encodes the consensus of all R proxy representations rather than the bias of any single one.

Proxy Selection. Rather than distilling from all proxies uniformly, SPD selects a reliable subset $\mathcal{S}_i \subseteq \{1, \dots, R\}$ for each sample x_i^e . We adopt a correctness-filtered small-loss criterion: a proxy P_r is considered a reliable candidate only if it correctly predicts the action class, i.e., $\hat{y}_i^r = y_i$, where $\hat{y}_i^r = \arg \max_c (z_i^r)_c$. The candidate set is thus:

$$\mathcal{C}_i = \{r \in \{1, \dots, R\} \mid \hat{y}_i^r = y_i\}. \quad (4)$$

Among candidates, proxy reliability is quantified by the cross-entropy loss $s_i^r = \text{CE}(z_i^r, y_i)$, where a lower loss indicates higher predictive confidence. When $|\mathcal{C}_i| > 0$, the top- k proxies with the lowest s_i^r are selected to form \mathcal{S}_i . When $\mathcal{C}_i = \emptyset$, we set $\mathcal{S}_i = \emptyset$ and skip distillation entirely for sample i , preventing the unified model from absorbing erroneous supervision.

Selective Proxy Distillation. For samples where $\mathcal{S}_i \neq \emptyset$, SPD transfers knowledge from the selected proxies to UNIEGO. Since all proxies are trained for action classification, they provide both feature embeddings and action logits. We therefore combine feature-level and logit-level distillation. Let $h_i^U, z_i^U = f(\theta_U; x_i^e)$ and $h_i^{P_r}, z_i^{P_r} = f(\theta_r; x_i^e)$ denote the feature embeddings and action logits of UNIEGO and the r -th proxy P_r , respectively. The per-sample SPD loss is:

$$\mathcal{L}_{\text{distill}}^U(i) = \frac{\beta_{\text{feat}}}{|\mathcal{S}_i|} \sum_{j \in \mathcal{S}_i} D_{\text{cos}}(h_i^U, h_i^{P_j}) + \frac{\beta_{\text{logit}}}{|\mathcal{S}_i|} \sum_{j \in \mathcal{S}_i} D_{\text{KL}}(\sigma(z_i^{P_j}) \parallel \sigma(z_i^U)), \quad (5)$$

where D_{\cos} and D_{KL} denote cosine distance and KL divergence, $\sigma(\cdot)$ is the softmax function, and $\beta_{\text{feat}}, \beta_{\text{logit}}$ are loss weights. The total Level-II objective is:

$$\mathcal{L}_{\text{II}} = \frac{1}{N} \sum_{i=1}^N (\mathbb{I}(|\mathcal{S}_i| > 0) \mathcal{L}_{\text{distill}}^U(i) + \beta_{\text{cls}} \text{CE}(z_i^U, y_i)), \quad (6)$$

where β_{cls} is the classification loss weight and $\mathbb{I}(\cdot)$ is the indicator function. During inference, only UNIEGO with egocentric input x^e is required:

$$\hat{y} = \arg \max_c z_c^U. \quad (7)$$

No teacher network, exocentric stream, or proxy is needed at test time.

4 Experiments

We evaluate UNIEGO on three egocentric datasets across three video understanding tasks: action recognition, video retrieval, and action segmentation. In all experiments, UNIEGO requires only egocentric RGB video at inference.

4.1 Experimental Setting

Datasets. We evaluate on three publicly available ego-exo datasets: **EgoExo-Fitness** [22], **Assembly101** [38], and **EgoExo4D** [14]. For EgoExo-Fitness and EgoExo4D, we follow the official evaluation splits [22, 14]. For Assembly101, we pair egocentric videos from the helmet-mounted *ego04* camera with exocentric videos from the frontal *exo03* camera, yielding 46,202 training and 15,307 test samples across 24 action classes. Further dataset details are provided in the Appendix.

Implementation Details. We use TimeSformer [2] as the egocentric student backbone for UNIEGO $f(\cdot)$ unless stated, which takes 8 frames sampled at 224×224 and produces a 768-dimensional video representation. During training, UNIEGO learns from multiple teacher representations \mathcal{T}_r which are detailed in Table 1. We apply a linear projection layer to align all teacher representations into the 768-dimensional student space to handle the dimensional mismatch between the student (h_i) and various teacher features (h_i^r). All proxies are trained, following TimeSformer [2] default configuration, for 15 epochs using SGD with a base learning rate of 0.005, momentum of 0.9, and weight decay of 1×10^{-4} . We scale the learning rate by factors of 0.1 and 0.01 at epochs 11 and 14, respectively. All models are trained with a total batch size of 8 distributed across 4 NVIDIA RTX A5000 GPUs.

In Level-I, we train 9 proxies across 6 teacher architectures and ego-exo viewpoints (Table 1), with loss weights $\lambda_{\text{I}} = 5$ and $\lambda_{\text{cls}} = 1$. In Level-II, the proxy merging coefficients α^* are optimized on the training set using Adam [18] for 2 epochs with a learning rate of 0.02 and weight decay of 0.01. For proxy selection, we set $K = 1$ for EgoExo-Fitness and Assembly101, and $K = 2$ for EgoExo4D (ablated in Table 8). The Level-II loss weights are $\beta_{\text{feat}} = 5$, $\beta_{\text{cls}} = 1$ throughout, and $\beta_{\text{logit}} = 1$ for EgoExo-Fitness and EgoExo4D and 5 for Assembly101, with KL divergence temperature $\tau = 1$. All hyperparameters are selected on the respective validation sets.

4.2 Comparison with State-of-the-Art

Table 2 compares UNIEGO with recent methods across three egocentric action recognition benchmarks under two inference protocols: *exocentric inference*, where exocentric streams are available at test time, and *egocentric inference*, where only egocentric video is provided. The former serves as a privileged-view upper bound [22, 38], as exocentric cameras in these datasets capture global body

Table 1: **Teacher pool for Level-I proxy learning.**

| Modality | Viewpoint | Teacher Model | Feature Dim. |
|----------|-----------|--------------------|--------------|
| RGB | Ego | DINOv2 [29] | 1024 |
| | Ego | SigLIP [52] | 1152 |
| Depth | Ego | DepthAnything [50] | 1024 |
| RGB | Exo | TimeSformer [2] | 768 |
| | Exo | DINOv2 [29] | 1024 |
| | Exo | Sk-Ego [37] | 512 |
| | Exo | SigLIP [52] | 1152 |
| Skeleton | Exo | ST-GCN [49] | 256 |
| Depth | Exo | DepthAnything [50] | 1024 |

Table 2: Comparison with state-of-the-art methods on three egocentric action recognition datasets (EgoExo-Fitness, Assembly101, EgoExo4D). Dist. indicates whether the model performs distillation and Acc. indicates Top-1 accuracy.

| (a) EgoExo-Fitness. | | | (b) Assembly101. | | | (c) EgoExo4D. | | |
|-----------------------------|-------|-------------|-----------------------------|-------|-------------|-----------------------------|-------|-------------|
| Method | Dist. | Acc. | Method | Dist. | Acc. | Method | Dist. | Acc. |
| <i>Exocentric inference</i> | | | <i>Exocentric inference</i> | | | <i>Exocentric inference</i> | | |
| TimeSformer [2] | ✗ | 88.9 | TimeSformer [2] | ✗ | 62.7 | TimeSformer [2] | ✗ | 26.0 |
| ST-GCN [49] | ✗ | 87.5 | ST-GCN [49] | ✗ | 46.2 | ST-GCN [49] | ✗ | 42.9 |
| <i>Egocentric inference</i> | | | <i>Egocentric inference</i> | | | <i>Egocentric inference</i> | | |
| I3D [6] | ✗ | 74.7 | TSM+TA [23, 39] | ✗ | 40.5 | VI Encoder [28] | ✓ | 40.3 |
| EgoVLP [24] | ✗ | 74.7 | ViFi-CLIP [35] | ✗ | 46.6 | EgoVLPv2 [31] | ✓ | 39.1 |
| ViFi-CLIP [35] | ✗ | 81.8 | π -ViT [36] | ✓ | 47.8 | Ego-Exo MAE [21] | ✓ | 37.2 |
| π -ViT [36] | ✓ | 80.1 | TimeSformer [2] | ✗ | 47.6 | Viewpoint Distillation [15] | ✓ | 38.2 |
| TimeSformer [2] | ✗ | 80.3 | Multiteacher Dist. [2] | ✓ | 48.2 | TimeSformer [2] | ✗ | 39.9 |
| Multiteacher Dist. [2] | ✓ | 81.5 | UNIEGO (Ours) | ✓ | 50.7 | Multiteacher Dist. [2] | ✓ | 40.6 |
| UNIEGO (Ours) | ✓ | 84.7 | | | | UNIEGO (Ours) | ✓ | 41.1 |

Table 3: Performance on Video Retrieval

| Method | EgoExo-Fitness | | Assembly101 | | EgoExo4D | |
|----------------------|----------------|--------------|--------------|--------------|--------------|--------------|
| | mAP | R@1 | mAP | R@1 | mAP | R@1 |
| Timesformer | 0.474 | 0.712 | 0.226 | 0.410 | 0.167 | 0.326 |
| Multiteacher Dist. | 0.486 | 0.720 | 0.228 | 0.413 | 0.178 | 0.331 |
| UNIEGO (Ours) | 0.543 | 0.748 | 0.253 | 0.424 | 0.182 | 0.340 |

Table 4: Performance on Egocentric Temporal Action Segmentation on Assembly101.

| Feature Backbone (Method) | F1@10 | F1@25 | F1@50 | Edit | Acc |
|----------------------------------|-------------|-------------|-------------|-------------|-------------|
| TimeSformer (Ego only) | 16.2 | 14.1 | 10.4 | 18.7 | 34.4 |
| TimeSformer (Multiteacher Dist.) | 15.3 | 13.2 | 9.8 | 18.4 | 34.2 |
| UNIEGO (Ours) | 19.6 | 16.9 | 12.3 | 19.4 | 34.7 |

pose and scene layout that are largely occluded in the egocentric field of view. All baselines except ST-GCN [49] utilize RGB input at inference.

UNIEGO consistently outperforms the TimeSformer backbone by +4.4%, +3.1%, and +1.2% on EgoExo-Fitness, Assembly101, and EgoExo4D respectively, demonstrating the benefit of consolidating diverse multi-teacher supervision. Against the strongest distillation baseline π -ViT [36], UNIEGO achieves gains of +4.6% and +2.9% on EgoExo-Fitness and Assembly101. Furthermore, UNIEGO consistently surpasses naive multi-teacher distillation [34, 40], confirming that proxy-mediated hierarchical distillation is essential for reconciling heterogeneous teacher representations into a unified egocentric encoder.

The largest gains are observed on EgoExo-Fitness, where exocentric proxies \mathcal{P}_{exo} are particularly strong as many actions in this dataset involve full-body motion that is inherently occluded from the egocentric viewpoint, making exocentric supervision especially informative. Conversely, on EgoExo4D, exocentric proxies are weaker, as evidenced by the baseline TimeSformer achieving only 26.0% under exocentric inference, below its egocentric counterpart. Nevertheless, UNIEGO remains robust to this proxy inconsistency across viewpoints, as the adaptive selection mechanism in SPD suppresses unreliable proxies and routes supervision from the most discriminative sources available. Across all three benchmarks, UNIEGO achieves state-of-the-art action recognition performance under egocentric inference.

Generalization Across Backbone Architectures. We verify that UNIEGO is architecture-agnostic by replacing the proxy and unified model backbone with UniFormer-S [20] and ViFi-CLIP [35] (Table 5). Across both alternatives, our hierarchical distillation framework consistently outperforms naive multi-teacher distillation, confirming that the gains of UNIEGO are not tied to a specific video encoder. Notably, UniFormer-S, a compact 22M model yields significant classification improvements, demonstrating that our framework is equally effective for learning efficient unified egocentric representations which is crucial for deployment in resource-constrained egocentric applications.

Video Retrieval & Temporal Action Segmentation. Table 3 evaluates UNIEGO on video retrieval, performed by extracting features from action recognition trained backbones and computing their pairwise similarity across the test set. On EgoExo-Fitness, naive multi-teacher distillation yields only a marginal improvement of +0.012 mAP over the baseline, whereas SPD achieves a substantially

Table 5: **Backbone robustness.** UNIEGO improves across diverse egocentric backbones.

| Backbone | Method | EEF |
|------------------|----------------------|-------------|
| TimeSformer [2] | baseline | 80.3 |
| | Multiteacher Dist. | 81.5 |
| | UNIEGO (Ours) | 84.7 |
| UniFormer-S [20] | baseline | 68.4 |
| | Multiteacher Dist. | 69.0 |
| | UNIEGO (Ours) | 73.5 |
| ViFi-CLIP [35] | baseline | 81.8 |
| | Multiteacher Dist. | 81.7 |
| | UNIEGO (Ours) | 83.8 |

Table 6: Ablation of components of UNIEGO.

| Proxy Learning | Proxy Merging | SPD | EEF | A101 |
|----------------|---------------|-----|-------------|-------------|
| ✗ | ✗ | ✗ | 80.3 | 47.6 |
| ✓ | ✗ | ✗ | 82.1 | 48.7 |
| ✓ | ✗ | ✓ | 82.3 | 48.9 |
| ✓ | ✓ | ✗ | 81.4 | 48.3 |
| ✓ | ✓ | ✓ | 84.7 | 50.7 |

Table 7: Alternative proxy merging strategies.

| Merging Strategy | EEF | A101 |
|-----------------------------|-------------|-------------|
| Best Proxy | 83.7 | 50.6 |
| Average | 83.4 | 50.2 |
| Layer-level | 84.2 | 50.6 |
| Parameter-level | 83.6 | 50.4 |
| Proxy Merging (Ours) | 84.7 | 50.7 |

Table 8: Strategies for Selective Proxy distillation.

| Distillation Strategy | EEF | A101 |
|-------------------------|-------------|-------------|
| No Distillation | 80.3 | 47.6 |
| All Proxies | 82.8 | 49.9 |
| Top-1(Ego) + Top-1(Exo) | 83.2 | 49.5 |
| Top-3 | 83.2 | 50.0 |
| Top-2 | 84.3 | 50.1 |
| Top-1 | 84.7 | 50.7 |

larger gain of +0.057 mAP. This trend remains consistent across datasets, confirming that proxy-mediated distillation yields highly discriminative egocentric representations.

For temporal action segmentation on Assembly101, we extract features from three egocentric TimeS-former backbones: i) trained from scratch, ii) naive multi-teacher distillation, and iii) UNIEGO, and feed them into a fixed temporal encoder [43]. We evaluate the performance via F1 score, Edit distance, and frame accuracy. As shown in Table 4, UNIEGO features yield the strongest performance across all metrics. Notably, naive multi-teacher distillation degrades performance on all metrics relative to the scratch baseline, confirming that directly distilling heterogeneous teachers disrupts fine-grained temporal representations. In contrast, UNIEGO’s hierarchical distillation preserves the local temporal structure essential for frame-wise discrimination and boundary-sensitive segmentation.

4.3 Ablation Studies and Model Diagnosis

We perform all the ablations and diagnosis of UNIEGO on the EgoExo-Fitness (EEF) and Assembly101 (A101) datasets.

Effect of Each Component. Table 6 ablates the three key components of UNIEGO: Proxy Learning, Proxy Merging Initialization, and SPD. Starting from the egocentric-only baseline, adding Level-I proxy learning followed by simultaneous distillation from all proxies yields gains of +1.8% and +1.1% on EEF and A101, respectively, highlighting the role of proxies as mediators. Then, replacing the simultaneous distillation in level-II with SPD further improves performance by +0.2% on both datasets, demonstrating that instance-adaptive proxy selection yields more reliable supervision than distilling from all proxies indiscriminately. We also observe that proxy merging initialization provides a strong starting point for SPD, as the merged model already achieves higher accuracy than the baseline. Finally, prepending proxy merging initialization before SPD achieves 84.7% and 50.7% on EEF and A101 resulting in overall improvements of +4.4% and +3.1% over the baseline. These results confirm the contribution of each component in our hierarchical distillation framework.

Alternative proxy merging strategies Table 7 compares different strategies for initializing UNIEGO before SPD. *Best Proxy* uses the strongest individual proxy as the initialization, while *Average* uniformly averages the weights of all trained proxies. We also compare against fine-grained learnable merging strategies: *Layer-level* learns separate merging weights for each layer, and *Parameter-level* learns separate merging weights for individual parameters. Although fine-grained merging offers greater flexibility, it does not improve performance in practice. Both layer-level and parameter-level merging outperform uniform averaging yet fall short of our global merging strategy, suggesting that a globally consistent combination of proxy parameters is more effective than local merging at the level of individual layers or parameter groups, likely due to the importance of maintaining parameter consistency across the full model.

Alternative SPD strategies. In Table 8, we investigate various selective distillation strategies following the proxy merging initialization. SPD outperforms distillation from all proxies simultaneously, confirming that selective supervision mitigates conflicting gradients across homogeneous proxies. Enforcing viewpoint diversity via a *Top-1(Ego) + Top-1(Exo)* selection strategy degrades performance, particularly on Assembly101, indicating that global proxy selection is preferable, i.e., not all training samples contain discriminative information from every viewpoint, and forcing viewpoint-balanced selection introduces noisy supervision into UNIEGO. Among $K \in \{1, 2, 3\}$, all top- K variants outperform naive proxy distillation, demonstrating the robustness of the selection mechanism, with $K = 1$ yielding the best overall performance.




Figure 3: Teacher vs Proxy Centered Kernel Alignment scores




Figure 4: Gradient Conflict Rate across distillation strategies




Figure 5: Cosine Similarity among teacher and proxy gradients

5 Does UNIEGO Mitigate Representational Gap and Conflicting Gradients?

The answer is “Yes”. We provide a model analysis to confirm that our hierarchical distillation framework addresses the two core failure modes of naive multi-teacher distillation: representational gap and conflicting gradients.

Representational Gap. To quantify the representational gap among the R teachers, we visualize their pairwise linear Centered Kernel Alignment (CKA) [19] in Figure 3. The lower triangle displays teacher-pair similarities and the upper triangle displays proxy-pair similarities on EgoExo-Fitness. The substantially higher pairwise CKA among proxies than among teachers confirms that Level-I proxy learning projects heterogeneous teacher representations into a homogeneous egocentric embedding space, directly alleviating the representational gap that impedes naive multi-teacher distillation.

Conflicting Gradients. Figures 4 and 5 report the gradient conflict rate which is defined as the fraction of distillation gradients that oppose the classification gradient, i.e., $\cos(\nabla_{\text{cls}}, \nabla_{\text{kd}}) < 0$, and the average cosine similarity across training epochs, respectively. Direct multi-teacher distillation exhibits severe gradient interference, confirming that heterogeneous teacher supervision actively opposes the primary task objective. Distilling from all proxies simultaneously reduces the conflict rate to 42–50% and shifts the cosine similarity toward zero, indicating partial mitigation. SPD further suppresses conflicts by routing supervision exclusively through reliable proxies, yielding consistently cooperative distillation and classification gradients. Together, these results confirm that proxy learning and SPD collectively transform a conflicting multi-teacher optimization landscape into a coherent one.

6 Conclusion

We presented UNIEGO, a unified egocentric encoder that consolidates diverse perceptual knowledge across viewpoints, modalities, and foundation model representations into a single model. The core of UNIEGO lies a hierarchical distillation framework in which representation-specific proxies serve as structured mediators, dissolving the representational gap between heterogeneous teachers before selective, reliability-guided distillation assembles their collective knowledge into a coherent unified representation. The result is an egocentric encoder that sees further, knows more, and generalizes better without requiring more than a single egocentric camera at inference.

While we hope UNIEGO inspires broader exploration into systematically harvesting and reconciling diverse supervisory signals for richer egocentric perception, the current framework has an important limitation. Specifically, UNIEGO relies on a small-loss criterion for proxy selection, a heuristic that, while effective, does not exploit the full potential of the proxy pool. A learned selection mechanism, one that dynamically weighs proxy reliability as a function of both the input and the training state could yield richer and more adaptive supervision. Designing such an adaptive proxy selection strategy is non-trivial, and we leave this as a promising direction for future work.

Acknowledgements

This work was supported in part by the National Science Foundation (IIS-2245652) and the University of North Carolina at Charlotte. Computational resources were provided by the NSF National AI Research Resource Pilot (NAIRR240338) and NCShare.

References

- [1] Varun Belagali, Saarthak Kapse, Pierre Marza, Srijan Das, Zilinghan Li, Sofiène Boutaj, Pushpak Pati, Srikar Yellapragada, Tarak Nath Nandi, Ravi K Madduri, Joel Saltz, Prateek Prasanna, Stergios Christodoulidis, Maria Vakalopoulou, and Dimitris Samaras. Ticon: A slide-level tile contextualizer for histopathology representation learning, 2025.
- [2] Gedas Bertasius, Heng Wang, and Lorenzo Torresani. Is space-time attention all you need for video understanding? In *Proceedings of the International Conference on Machine Learning (ICML)*, July 2021.
- [3] Stephen Boyd and Lieven Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.
- [4] Zhe Cao, Gines Hidalgo Martinez, Tomas Simon, Shih-En Wei, and Yaser Sheikh. Openpose: Realtime multi-person 2d pose estimation using part affinity fields. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2019.
- [5] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9650–9660, 2021.
- [6] Joao Carreira and Andrew Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. In *Conference on Computer Vision and Pattern Recognition*, pages 4724–4733. IEEE, 2017.
- [7] Dima Damen, Hazel Doughty, Giovanni Maria Farinella, Sanja Fidler, Antonino Furnari, Evangelos Kazakos, Davide Moltisanti, Jonathan Munro, Toby Perrett, Will Price, and Michael Wray. Scaling egocentric vision: The epic-kitchens dataset. In *European Conference on Computer Vision (ECCV)*, 2018.
- [8] Srijan Das, Rui Dai, Di Yang, and Francois Bremond. Vpn++: Rethinking video-pose embeddings for understanding activities of daily living. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pages 1–1, 2021.
- [9] Srijan Das, Saurav Sharma, Rui Dai, Francois Bremond, and Monique Thonnat. Vpn: Learning video-pose embedding for activities of daily living. In *European Conference on Computer Vision*, pages 72–90. Springer, 2020.
- [10] Zi-Yi Dou, Xitong Yang, Tushar Nagarajan, Huiyu Wang, Jing Huang, Nanyun Peng, Kris Kitani, and Fu-Jen Chu. Unlocking exocentric video-language data for egocentric video representation learning, 2024.
- [11] Jonathan Frankle, Gintare Karolina Dziugaite, Daniel M Roy, and Michael Carlin. Linear mode connectivity and the lottery ticket hypothesis. In *International Conference on Machine Learning (ICML)*, 2020.
- [12] Xinyu Gong, Sreyas Mohan, Naina Dhingra, Jean-Charles Bazin, Yilei Li, Zhangyang Wang, and Rakesh Ranjan. Mmg-ego4d: Multi-modal generalization in egocentric action recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023.
- [13] Kristen Grauman, Andrew Westbury, Eugene Byrne, Zachary Chavis, Antonino Furnari, Rohit Girdhar, Jackson Hamburger, Hao Jiang, Miao Liu, Xingyu Liu, Miguel Martin, Tushar Nagarajan, Ilija Radosavovic, Santhosh Kumar Ramakrishnan, Fiona Ryan, Jayant Sharma, Michael Wray, Mengmeng Xu, Eric Zhongcong Xu, Chen Zhao, Siddhant Bansal, Dhruv Batra,

- Vincent Cartillier, Sean Crane, Tien Do, Morrie Doulaty, Akshay Erapalli, Christoph Feichtenhofer, Adriano Fragomeni, Qichen Fu, Abrahm Gebreselasie, Cristina Gonzalez, James Hillis, Xuhua Huang, Yifei Huang, Wenqi Jia, Weslie Khoo, Jachym Kolar, Satwik Kottur, Anurag Kumar, Federico Landini, Chao Li, Yanghao Li, Zhenqiang Li, Karttikeya Mangalam, Raghava Modhugu, Jonathan Munro, Tullie Murrell, Takumi Nishiyasu, Will Price, Paola Ruiz Puentes, Merey Ramazanova, Leda Sari, Kiran Somasundaram, Audrey Southerland, Yusuke Sugano, Ruijie Tao, Minh Vo, Yuchen Wang, Xindi Wu, Takuma Yagi, Ziwei Zhao, Yunyi Zhu, Pablo Arbelaez, David Crandall, Dima Damen, Giovanni Maria Farinella, Christian Fuegen, Bernard Ghanem, Vamsi Krishna Ithapu, C. V. Jawahar, Hanbyul Joo, Kris Kitani, Haizhou Li, Richard Newcombe, Aude Oliva, Hyun Soo Park, James M. Rehg, Yoichi Sato, Jianbo Shi, Mike Zheng Shou, Antonio Torralba, Lorenzo Torresani, Mingfei Yan, and Jitendra Malik. Ego4d: Around the world in 3,000 hours of egocentric video. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18995–19012, 2022.
- [14] Kristen Grauman, Andrew Westbury, Lorenzo Torresani, Kris Kitani, Jitendra Malik, Triantafyllos Afouras, Kumar Ashutosh, Vijay Baiyya, Siddhant Bansal, Bikram Boote, Eugene Byrne, Zachary Chavis, Joya Chen, Feng Cheng, Fu-Jen Chu, Sean Crane, Avijit Dasgupta, Jing Dong, María Escobar, Cristhian Forigua, Abrahm Kahsay Gebreselasie, Sanjay Haresh, Jing Huang, Md Mohaiminul Islam, Suyog Dutt Jain, Rawal Khirodkar, Devansh Kukreja, Kevin J. Liang, Jia-Wei Liu, Sagnik Majumder, Yongsun Mao, Miguel Martin, Effrosyni Mavroudi, Tushar Nagarajan, Francesco Ragusa, Santhosh K. Ramakrishnan, Luigi Seminara, Arjun Somayazulu, Yale Song, Shan Su, Zihui Xue, Edward Zhang, Jinxu Zhang, Ángela Castillo, Changan Chen, Xinzhu Fu, Ryosuke Furuta, Cristina González, Prince Gupta, Jiabo Hu, Yifei Huang, Yiming Huang, Weslie Khoo, Anushk Kumar, Robert Kuo, Sach Lakhavani, Miao Liu, Romy Mi Luo, Zhengyi Luo, Brigid Meredith, Austin Miller, Oluwatumininu Oguntola, Xiaqing Pan, Penny Peng, Shraman Pramanick, Merey Ramazanova, Fiona Ryan, W. Shan, Kiran Somasundaram, Chenan Song, Audrey Southerland, Masatoshi Tateno, Huiyu Wang, Yuchen Wang, Takuma Yagi, Mingfei Yan, Xitong Yang, Ze Yu, Shengxin Zha, Chen Zhao, Ziwei Zhao, Zhifan Zhu, J. F. Zhuo, Pablo Arbeláez, Gedas Bertasius, David J. Crandall, Dima Damen, Jakob Julian Engel, Giovanni Maria Farinella, Antonino Furnari, Bernard Ghanem, Judy Hoffman, C. V. Jawahar, Richard A. Newcombe, Hyun Soo Park, James M. Rehg, Yoichi Sato, Manolis Savva, Jianbo Shi, Mike Zheng Shou, and Michael Wray. Ego-exo4d: Understanding skilled human activity from first- and third-person perspectives. *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 19383–19400, 2023.
- [15] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network, 2015.
- [16] Pavel Izmailov, Dmitrii Podoprikin, Timur Garipov, Dmitry Vetrov, and Andrew Gordon Wilson. Averaging weights leads to wider optima and better generalization. In *Uncertainty in Artificial Intelligence (UAI)*, 2018.
- [17] Qihong Ke, Mohammed Bennamoun, Senjian An, Ferdous Sohel, and Farid Boussaid. Learning clip representations for skeleton-based 3d action recognition. *IEEE Transactions on Image Processing*, 27(6):2842–2855, June 2018.
- [18] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representations (ICLR)*, 2015.
- [19] Simon Kornblith, Mohammad Norouzi, Honglak Lee, and Geoffrey Hinton. Similarity of neural network representations revisited. In *Proceedings of the 36th International Conference on Machine Learning (ICML)*, 2019.
- [20] Kunchang Li, Yali Wang, Junhao Zhang, Peng Gao, Guanglu Song, Yu Liu, Hongsheng Li, and Yu Jiao Qiao. Uniformer: Unifying convolution and self-attention for visual recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45:12581–12600, 2022.
- [21] Yanghao Li, Tushar Nagarajan, Bo Xiong, and Kristen Grauman. Ego-exo: Transferring visual representations from third-person to first-person videos. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10995–11005, 2021.

- [22] Yuan-Ming Li, Wei-Jin Huang, An-Lan Wang, Ling-An Zeng, Jing-Ke Meng, and Wei-Shi Zheng. Egoexo-fitness: towards egocentric and exocentric full-body action understanding. In *European Conference on Computer Vision*, pages 363–382. Springer, 2024.
- [23] Ji Lin, Chuang Gan, and Song Han. Tsm: Temporal shift module for efficient video understanding. In *Proceedings of the IEEE International Conference on Computer Vision*, 2019.
- [24] Kevin Qinghong Lin, Alex Jinpeng Wang, Mattia Soldan, Michael Wray, Rui Yan, Eric Zhongcong Xu, Difei Gao, Rongcheng Tu, Wenzhe Zhao, Weijie Kong, et al. Egocentric video-language pretraining. *arXiv preprint arXiv:2206.01670*, 2022.
- [25] Bo Liu, Xingchao Liu, Xiaojie Jin, Peter Stone, and Qiang Liu. Conflict-averse gradient descent for multi-task learning. In M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan, editors, *Advances in Neural Information Processing Systems*, volume 34, pages 18873–18885. Curran Associates, Inc., 2021.
- [26] Yuang Liu, Wei Zhang, and Jun Wang. Adaptive multi-teacher multi-level knowledge distillation. *Neurocomputing*, 2020.
- [27] Mi Luo, Zihui Xue, Alex Dimakis, and Kristen Grauman. Viewpoint rosetta stone: Unlocking unpaired ego-exo videos for view-invariant representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2025.
- [28] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018.
- [29] Maxime Oquab, Timothée Darcet, Theo Moutakanni, Huy V. Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, Russell Howes, Po-Yao Huang, Hu Xu, Vasu Sharma, Shang-Wen Li, Wojciech Galuba, Mike Rabbat, Mido Assran, Nicolas Ballas, Gabriel Synnaeve, Ishan Misra, Hervé Jégou, Julien Mairal, Patrick Labatut, Armand Joulin, and Piotr Bojanowski. Dinov2: Learning robust visual features without supervision. *Transactions on Machine Learning Research (TMLR)*, 2024.
- [30] Dario Pavlo, Christoph Feichtenhofer, David Grangier, and Michael Auli. 3d human pose estimation in video with temporal convolutions and semi-supervised training. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [31] Shraman Pramanick, Yale Song, Sayan Nag, Kevin Qinghong Lin, Hardik Shah, Mike Zheng Shou, Rama Chellappa, and Pengchuan Zhang. EgoVLPv2: Egocentric video-language pre-training with fusion in the backbone. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5285–5297, 2023.
- [32] Camillo Quattrocchi, Antonino Furnari, Daniele Di Mauro, Mario Valerio Giuffrida, and Giovanni Maria Farinella. Synchronization is all you need: Exocentric-to-egocentric transfer for temporal action segmentation with unlabeled synchronized video pairs. In *European Conference on Computer Vision (ECCV)*, 2024.
- [33] Gorjan Radevski, Dusan Grujicic, Marie-Francine Moens, Matthew Blaschko, and Tinne Tuytelaars. Multimodal distillation for egocentric action recognition. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2023.
- [34] Mike Ranzinger, Greg Heinrich, Jan Kautz, and Pavlo Molchanov. Am-radio: Agglomerative vision foundation model reduce all domains into one. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2024.
- [35] Hanoona Rasheed, Muhammad Uzair Khattak, Muhammad Maaz, Salman Khan, and Fahad Shahbaz Khan. Finetuned clip models are efficient video learners. In *The IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023.
- [36] Dominick Reilly and Srijan Das. Just add $\pi!$ pose induced video transformers for understanding activities of daily living. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2024.

- [37] Dominick Reilly, Manish Kumar Govind, Le Xue, and Srijan Das. From my view to yours: Ego-augmented learning in large vision language models for understanding exocentric daily living activities, 2025.
- [38] Fadime Sener, Dibyadip Chatterjee, Daniel Sheleпов, Kun He, Dipika Singhania, Robert Wang, and Angela Yao. Assembly101: A large-scale multi-view video dataset for understanding procedural activities. *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 21064–21074, 2022.
- [39] Fadime Sener, Dipika Singhania, and Angela Yao. Temporal aggregate representations for long-range video understanding. In *European conference on computer vision*, pages 154–171. Springer, 2020.
- [40] Jinghuan Shang, Karl Schmeckpeper, Brandon B May, Maria Vittoria Minniti, Tarik Kelestemur, David Watkins, and Laura Herlant. Theia: Distilling diverse vision foundation models for robot learning. *arXiv preprint arXiv:2407.20179*, 2024.
- [41] Guangyuan Shi, Qimai Li, Wenlong Zhang, Jiaxin Chen, and Xiao-Ming Wu. Recon: Reducing conflicting gradients from the root for multi-task learning. In *International Conference on Learning Representations (ICLR)*, 2023.
- [42] Gunnar A Sigurdsson, Abhinav Gupta, Cordelia Schmid, Ali Farhadi, and Karteek Alahari. Charades-ego: A large-scale dataset of paired third and first person videos. *arXiv preprint arXiv:1804.09626*, 2018.
- [43] Arkaprava Sinha, Monish Soundar Raj, Pu Wang, Ahmed Helmy, and Srijan Das. Ms-temba: Multi-scale temporal mamba for efficient temporal action detection. *arXiv preprint arXiv:2501.06138*, 2025.
- [44] Shuhan Tan, Tushar Nagarajan, and Kristen Grauman. Egodistill: Egocentric head motion distillation for efficient video understanding. *arXiv preprint arXiv:2301.02217*, 2023.
- [45] Anirudh Thatipelli, Shao-Yuan Lo, and Amit K. Roy-Chowdhury. Egocentric and exocentric methods: A short survey. *Computer Vision and Image Understanding*, 257:104371, 2025.
- [46] Thomas Wimmer, Prune Truong, Marie-Julie Rakotosaona, Michael Oechsle, Federico Tombari, Bernt Schiele, and Jan Eric Lenssen. Anyup: Universal feature upsampling. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2026.
- [47] Boshen Xu, Yuting Mei, Xinbi Liu, Sipeng Zheng, and Qin Jin. Egodtm: Towards 3d-aware egocentric video-language pretraining. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2025.
- [48] Zihui Xue and Kristen Grauman. Learning fine-grained view-invariant representations from unpaired ego-exo videos via temporal alignment. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2023.
- [49] Sijie Yan, Yuanjun Xiong, and Dahua Lin. Spatial temporal graph convolutional networks for skeleton-based action recognition. In *Thirty-second AAAI conference on artificial intelligence*, 2018.
- [50] Lihe Yang, Bingyi Kang, Zilong Huang, Zhen Zhao, Xiaogang Xu, Jiashi Feng, and Hengshuang Zhao. Depth anything v2. In *Advances in Neural Information Processing Systems*, 2024.
- [51] Tianhe Yu, Saurabh Kumar, Abhishek Gupta, Sergey Levine, Karol Hausman, and Chelsea Finn. Gradient surgery for multi-task learning. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 5434–5445. Curran Associates, Inc., 2020.
- [52] Xiaohua Zhai, Basil Mustafa, Alexander Kolesnikov, and Lucas Beyer. Sigmoid loss for language image pre-training. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. IEEE, 2023.

- [53] Hailin Zhang, Defang Chen, and Can Wang. Confidence-aware multi-teacher knowledge distillation. In *ICASSP 2022 - 2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2022.
- [54] Yue Zhao, Ishan Misra, Philipp Krahenbuhl, and Rohit Girdhar. Learning video representations from large language models. *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6586–6597, 2022.

A Overview

The appendix is categorized into the following parts:

- Section B: Detailed Data Description
- Section C: Proxy Performances
- Section D: Analysis of Proxy Selection
- Section E: Analysis of Proxy Merging
- Section F: Action Analysis

B Detailed Dataset Description

EgoExo-Fitness [22] is a full-body action understanding dataset containing synchronized egocentric and exocentric videos of fitness activities. It consists of 12 action categories and approximately 32 hours of video. Following the official split, we use 3,522 training samples and 912 test samples, with action sequences ranging from 10 to 30 seconds in duration.

Assembly101 [38] is a large-scale procedural activity dataset comprising 167 hours of video, recording subjects performing object assembly tasks from multiple viewpoints. In our experiments, we pair the egocentric videos from the helmet-mounted *ego04* camera with the corresponding exocentric videos in front *exo03* camera. The resulting split contains 46,202 training samples and 15,307 test samples over 24 action classes.

EgoExo4D [14] is a large-scale egocentric-exocentric video dataset covering a diverse range of skilled human activities. For each sample, we pair the egocentric video with the exocentric view annotated as the best view by human annotators. The dataset split used in our experiments contains 30,660 training samples and 9,356 test samples across 665 action classes.

C Proxy Performances

Table 9: Performance of the base ego model and different ego/exo proxy models across EgoExo-Fitness (EEF), Assembly101 (A101), and Ego-Exo4D (EE4D).

| Modality | Viewpoint | Model | EEF | A101 | EE4D |
|-----------------------------------------------------|-----------|---------------|------|------|------|
| Base Ego Model | | | | | |
| RGB [2] | Ego | TimeSformer | 80.3 | 47.6 | 39.9 |
| Ego Proxies (\mathcal{P}_{ego}) | | | | | |
| Depth [50] | Ego | DepthAnything | 80.3 | 48.2 | 40.1 |
| DINOv2 [29] | Ego | DINOv2 | 80.9 | 48.4 | 40.2 |
| SigLip [52] | Ego | SigLip | 80.7 | 48.4 | 40.6 |
| Exo Proxies (\mathcal{P}_{exo}) | | | | | |
| RGB [2] | Exo | TimeSformer | 81.9 | 49.5 | 40.0 |
| Skeleton [49] | Exo | ST-GCN | 81.9 | 48.4 | 40.3 |
| Depth [50] | Exo | DepthAnything | 80.3 | 48.5 | 39.9 |
| DINOv2 [29] | Exo | DINOv2 | 80.7 | 48.2 | 39.5 |
| Sk-Ego [37] | Exo | Sk-Ego | 81.9 | 48.3 | 40.2 |
| SigLip [52] | Exo | SigLip | 80.3 | 48.0 | 40.6 |

In Table 9, we report the action classification performance of all proxies (\mathcal{P}_r) after Level-I distillation. We observe that not all proxies outperform the base ego model, motivating the need for proxy selection. Nevertheless, the final performance of UNIEGO surpasses that of all individual proxies, corroborating the effectiveness of our proposed hierarchical distillation framework.

D Analysis of Proxy Selection.

In Figure 6, we analyze the selection frequency of each proxy during the Level-II selective proxy distillation (SPD). First, every proxy is selected to some extent across all benchmarks, indicating that SPD does not collapse to a single teacher. Unsurprisingly, we find that the routing distributions are

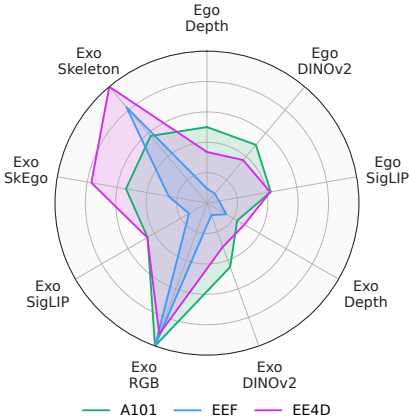


Figure 6: Selective Proxy Distillation Statistics.

Table 10: Effect of different proxy initialization strategies before and after SPD on EgoExo-Fitness (EEF) and Assembly101 (A101).

| Method | EEF | | A101 | |
|-----------------|------------|-----------|------------|-----------|
| | Before SPD | After SPD | Before SPD | After SPD |
| Best Proxy | 81.9 | 83.7 | 49.5 | 50.6 |
| Average | 79.9 | 83.4 | 42.9 | 50.2 |
| Layer-level | 80.9 | 84.2 | 43.5 | 50.6 |
| Parameter-level | 79.8 | 83.6 | 43.4 | 50.4 |
| Model-level | 81.4 | 84.7 | 48.3 | 50.7 |

not consistent across benchmarks, and that proxy selection is strongly correlated with the distinct dataset characteristics. On **EgoExo-Fitness**, the SPD primarily selects *Exo Skeleton* and *Exo RGB*. This is consistent with the nature of the dataset, which consists largely of full-body fitness motions: exocentric and skeleton-based proxies provide supervision for body pose and motion patterns that are often only partially visible from a head-mounted camera. In **Assembly101**, the selection is more diverse, assigning a larger weight to egocentric proxies. This suggests that, compared to EgoExo-Fitness, more actions in Assembly101 rely on egocentric cues such as local hand-object interactions and subtle manipulation dynamics. On **EgoExo4D**, there is a notable increase in the selection of the *Exo SkEgo* proxy. We attribute this to the high-resolution, spatially localized nature of the benchmark: action-relevant regions can occupy only a small part of the full frame, and the SkEgo crop provides a more focused view of these regions.

E Analysis of Proxy Merging

Table 10 compares action classification performance of UNIEGO under different proxy merging initialization strategies, before and after SPD, on EgoExo-Fitness and Assembly101. Across all initialization strategies, SPD consistently improves performance, confirming the importance of the second distillation level. Notably, the best-performing initialization prior to SPD, i.e., the single Best Proxy, does not yield the strongest post-SPD accuracy. Instead, the model soup, i.e., the learned convex combination of all proxy parameters, provides the most effective initialization for SPD. This suggests that the advantage of proxy merging lies not in its standalone classification performance but in its ability to place UNIEGO in a well-conditioned region of the loss landscape that is maximally amenable to subsequent distillation.

F Action Analysis

As shown in Table 11 and Table 12, *exo_rgb* is the most frequently selected proxy on both datasets, with a sample-weighted selection rate of 0.349 on EEF and 0.182 on A101, followed by *exo_sk1*. However, the class-level separation between the top-ranked and second-ranked proxies differs substantially across datasets. The average per-class gap between the top-1 and top-2 proxies is 0.186 on EEF, but only 0.069 on A101. This suggests that EEF classes often have a more clearly preferred

Table 11: Per-class top-1 proxy selection rates on EEF (12 classes)

| class_name | none | exo_rgb | exo_skl | exo_siglip | ego_siglip | exo_skego | ego_depth | exo_depth | exo_dino | ego_dino |
|----------------------------|-------|--------------|--------------|------------|------------|-----------|-----------|--------------|----------|----------|
| Kneeling pushing-ups | 0.053 | 0.139 | 0.172 | 0.053 | 0.053 | 0.139 | 0.013 | 0.179 | 0.106 | 0.093 |
| Push-ups | 0.000 | 0.500 | 0.400 | 0.000 | 0.025 | 0.042 | 0.008 | 0.017 | 0.000 | 0.008 |
| Kneeling Torso Twist | 0.046 | 0.431 | 0.218 | 0.030 | 0.005 | 0.056 | 0.117 | 0.041 | 0.036 | 0.020 |
| Knee Raise & Abd. Contract | 0.055 | 0.290 | 0.428 | 0.014 | 0.028 | 0.048 | 0.055 | 0.034 | 0.021 | 0.028 |
| Shoulder Bridge | 0.000 | 0.987 | 0.013 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 |
| Sit-ups | 0.057 | 0.172 | 0.261 | 0.070 | 0.013 | 0.108 | 0.076 | 0.146 | 0.045 | 0.051 |
| Leg Reverse Lunge | 0.038 | 0.106 | 0.477 | 0.068 | 0.068 | 0.114 | 0.023 | 0.045 | 0.053 | 0.008 |
| Leg Lunge With Knee Lift | 0.118 | 0.300 | 0.291 | 0.073 | 0.036 | 0.109 | 0.036 | 0.018 | 0.009 | 0.009 |
| Sumo Squat | 0.112 | 0.388 | 0.233 | 0.017 | 0.017 | 0.198 | 0.009 | 0.009 | 0.000 | 0.017 |
| Jumping Jacks | 0.042 | 0.242 | 0.347 | 0.126 | 0.016 | 0.153 | 0.016 | 0.005 | 0.021 | 0.032 |
| High Knee | 0.054 | 0.264 | 0.271 | 0.085 | 0.078 | 0.078 | 0.008 | 0.054 | 0.054 | 0.054 |
| Clap Jacks | 0.011 | 0.345 | 0.414 | 0.029 | 0.023 | 0.063 | 0.029 | 0.057 | 0.011 | 0.017 |

Table 12: Per-class top-1 proxy selection rates on A101 (24 classes).

| class_name | none | exo_rgb | exo_skl | exo_siglip | ego_siglip | exo_skego | ego_depth | exo_depth | exo_dino | ego_dino |
|---------------------|-------|--------------|--------------|------------|------------|--------------|-----------|-----------|--------------|--------------|
| pick up | 0.049 | 0.256 | 0.125 | 0.089 | 0.079 | 0.104 | 0.079 | 0.042 | 0.076 | 0.102 |
| put down | 0.051 | 0.204 | 0.122 | 0.107 | 0.106 | 0.103 | 0.091 | 0.045 | 0.076 | 0.094 |
| inspect | 0.186 | 0.136 | 0.108 | 0.080 | 0.066 | 0.092 | 0.085 | 0.052 | 0.091 | 0.105 |
| rotate | 0.307 | 0.123 | 0.085 | 0.056 | 0.049 | 0.090 | 0.088 | 0.037 | 0.083 | 0.082 |
| unscrew | 0.042 | 0.196 | 0.124 | 0.100 | 0.100 | 0.123 | 0.124 | 0.044 | 0.064 | 0.085 |
| position | 0.107 | 0.139 | 0.081 | 0.075 | 0.086 | 0.143 | 0.125 | 0.048 | 0.088 | 0.108 |
| screw | 0.048 | 0.250 | 0.078 | 0.099 | 0.087 | 0.088 | 0.128 | 0.042 | 0.087 | 0.091 |
| remove | 0.187 | 0.125 | 0.138 | 0.064 | 0.067 | 0.098 | 0.090 | 0.044 | 0.085 | 0.102 |
| position screw on | 0.134 | 0.141 | 0.086 | 0.104 | 0.084 | 0.088 | 0.111 | 0.039 | 0.107 | 0.106 |
| remove screw from | 0.201 | 0.115 | 0.117 | 0.092 | 0.068 | 0.071 | 0.083 | 0.043 | 0.141 | 0.071 |
| pass | 0.165 | 0.142 | 0.089 | 0.059 | 0.067 | 0.123 | 0.086 | 0.040 | 0.130 | 0.098 |
| tilt up | 0.291 | 0.143 | 0.041 | 0.058 | 0.082 | 0.057 | 0.102 | 0.042 | 0.087 | 0.097 |
| attempt to position | 0.413 | 0.084 | 0.097 | 0.036 | 0.025 | 0.088 | 0.089 | 0.035 | 0.048 | 0.085 |
| push | 0.336 | 0.063 | 0.098 | 0.076 | 0.096 | 0.102 | 0.061 | 0.054 | 0.078 | 0.036 |
| tilt down | 0.224 | 0.179 | 0.116 | 0.078 | 0.068 | 0.108 | 0.046 | 0.041 | 0.079 | 0.059 |
| pull | 0.370 | 0.082 | 0.123 | 0.042 | 0.062 | 0.071 | 0.076 | 0.027 | 0.120 | 0.029 |
| attempt to remove | 0.595 | 0.067 | 0.031 | 0.029 | 0.012 | 0.084 | 0.061 | 0.006 | 0.027 | 0.088 |
| attempt to pick up | 0.754 | 0.040 | 0.022 | 0.025 | 0.020 | 0.047 | 0.025 | 0.017 | 0.017 | 0.032 |
| clap | 0.007 | 0.883 | 0.047 | 0.010 | 0.005 | 0.015 | 0.005 | 0.007 | 0.012 | 0.007 |
| attempt to unscrew | 0.515 | 0.095 | 0.075 | 0.037 | 0.050 | 0.025 | 0.079 | 0.008 | 0.017 | 0.100 |
| attempt to put down | 0.988 | 0.012 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 |
| spin | 0.452 | 0.170 | 0.064 | 0.053 | 0.027 | 0.059 | 0.037 | 0.027 | 0.032 | 0.080 |
| attempt to screw | 0.951 | 0.019 | 0.010 | 0.000 | 0.000 | 0.010 | 0.000 | 0.000 | 0.010 | 0.000 |
| shake | 0.783 | 0.075 | 0.000 | 0.025 | 0.000 | 0.017 | 0.008 | 0.000 | 0.042 | 0.050 |

proxy, whereas A101 exhibits a flatter proxy distribution, where several proxies can have similar selection rates for the same class.

Interestingly, `exo_depth` shows a highly class-specific behavior. Although it has the lowest aggregate selection rate among all proxies (0.042 on A101 and 0.052 on EEF), it is the most frequently selected proxy for *Kneeling push-ups* on EEF, with a selection rate of 0.179, exceeding the second-ranked proxy, `exo_rgb`, at 0.139. This indicates that depth can provide useful supervision for specific classes where geometric cues, such as body-to-floor distance, are discriminative. However, its low aggregate selection rate also suggests that this benefit is not broadly shared across classes.

The behavior of ego-side proxies also differs across the two datasets. On EEF, all ego-side proxies have relatively low selection rates, around 0.03, suggesting that third-person supervision is more reliable for large whole-body fitness motions. In contrast, on A101, `ego_dino` and `ego_depth` reach selection rates of approximately 0.091 and rank among the middle group of proxies. Each of them is also the most frequently selected proxy for at least one *attempt-to-** class. This difference is consistent with the task characteristics of A101: fine-grained assembly actions are often hand- and object-centric, and egocentric observations may capture local manipulation details that are less visible from exocentric views.