



SKI Models: SKeleton Induced Vision-Language Embeddings for Understanding Activities of Daily Living

Arkaprava Sinha¹, Dominick Reilly¹, Francois Bremond^{2,3}, Pu Wang¹, Srijan Das¹

¹University of North Carolina at Charlotte

²Université Côte d’Azur

³Inria

{asinha13, sdas24}@charlotte.edu

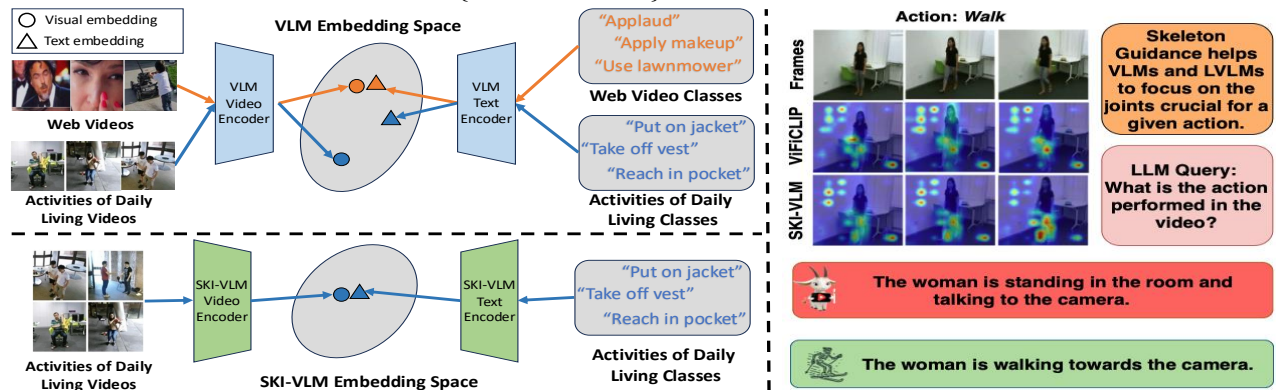


Figure 1: **Left:** The illustration depicts an embedding space of a Vision-Language Model (VLM) where representations of web-based videos align closely with their corresponding class label text features, while those of Activities of Daily Living (ADL) videos remain distant. Our study reveals that integrating skeleton guidance bridges this gap, aligning ADL video representations with their respective class labels. **Right:** Activation maps demonstrate how skeleton guidance sharpens the model’s focus on the critical body parts (such as legs) for specific actions, like *Walk*. This enhancement is evident in the improved text descriptions generated by Large-Vision-Language Models (LVLMs) when queried about actions depicted in the videos.

Abstract

The introduction of vision-language models like CLIP has enabled the development of foundational video models capable of generalizing to unseen videos and human actions. However, these models are typically trained on web videos, which often fail to capture the challenges present in Activities of Daily Living (ADL) videos. Existing works address ADL-specific challenges, such as *similar appearances*, *subtle motion patterns*, and *multiple viewpoints*, by combining 3D skeletons and RGB videos. However, these approaches are not integrated with language, limiting their ability to generalize to unseen action classes. In this paper, we introduce **SKI models**, which integrate 3D skeletons into the vision-language embedding space. SKI models leverage a skeleton-language model, **SkeletonCLIP**, to infuse skeleton information into Vision Language Models (VLMs) and Large Vision Language Models (LVLMs) through collaborative training. Notably, SKI models do not require skeleton data during inference, enhancing their robustness for real-world applications. The effectiveness of SKI models is validated on three popular ADL datasets for zero-shot action recognition and video caption generation tasks. Our code is available at <https://github.com/thearkaprava/SKI-Models>

1 Introduction

In recent years, the introduction of CLIP (Radford et al. 2021), has established the beneficial impact of language supervision in the training of vision-based discriminative models. In the video domain many works (Rasheed et al. 2023; Huang et al. 2024; Ni et al. 2022) have extended the effectiveness of CLIP to video representation learning. These approaches have made promising strides especially in zero shot video action recognition where the target action labels are unavailable during training. A majority of these models have been trained on web-based videos (Kay et al. 2017; Soomro, Zamir, and Shah 2012; Kuehne et al. 2011) consisting of sports videos, movie clips, etc. These videos generally contain actions with prominent motion and are typically appearance based, aligning actions closely to their scenes; for instance, *playing soccer on green grass* or *swimming in blue water*. In contrast, ADL videos often involve actions with **similar appearances**, **subtle motions**, and may be captured from **multiple viewpoints**. These characteristics limit the “generalizability” of video models trained on web videos to those containing ADL. Consequently, this *distributional shift* in video representations causes the video models to struggle in recognizing action categories not encountered

during training as illustrated in Figure 1. As a result, these models lack robustness as zero-shot learners, a crucial characteristic necessary for monitoring systems to detect anomalous human gestures indicative of cognitive decline. This motivates us to develop vision-language models (VLMs) capable of generalizing to unseen ADL scenarios. While the aforementioned *distribution shift* can be mitigated by training VLMs (Rasheed et al. 2023; Ni et al. 2022; Huang et al. 2024) on ADL videos (Shahroudy et al. 2016; Liu et al. 2019), these models still struggle to effectively address the inherent challenges in ADL due to the lack of specialized operations within VLMs designed to tackle them.

Instead of relying on specialized architectural changes, many works have explored the use of additional modalities beyond vision for action recognition. For instance, optical flow has been effective in web video datasets (Kay et al. 2017; Soomro, Zamir, and Shah 2012; Kuehne et al. 2011), where actions are characterized by strong motion (Crao et al. 2019; Simonyan and Zisserman 2014). However, it is less effective on ADL datasets (Shahroudy et al. 2016; Liu et al. 2019; Das et al. 2019), where motion cues are more subtle (Das et al. 2021). Conversely, 3D skeletons have been established as a strong modality, being viewpoint-agnostic and providing crucial information for modeling ADL (Yan, Xiong, and Lin 2018; Shi et al. 2019; Zhou et al. 2022). Similar to VLMs, skeleton-language representations have been explored in (Gupta, Sharma, and Sarvadevabhatla 2021; Zhou et al. 2023; Chen et al. 2024). Skeleton action recognition models, while excelling at capturing motion cues and human motion, do not encode the appearance information crucial for distinguishing nuanced human-object interactions. For instance, actions like *Drinking from Cup* versus *Drinking from Bottle* rely on visual appearance to differentiate objects, which motion alone cannot provide. Thus, integrating appearance information can enhance model’s ability to accurately classify ADLs. This is demonstrated in (Das et al. 2021), where appearance-based methods outperform approaches that rely solely on skeleton information. Additionally, the combination of RGB and skeletons has been investigated in (Das et al. 2020; Duan et al. 2022; Das et al. 2021; Reilly and Das 2024), but these approaches are not zero-shot learners. Therefore, in this paper, we pose the critical question: *Can we introduce 3D skeletons into the vision-language embedding space to enable effective zero-shot learning for ADL?*

One naive solution is to learn a common embedding space for videos, skeletons, and language. However, the limited availability of large-scale synchronized data renders this approach ineffective. Another approach involves aligning skeleton embeddings with learned image embeddings (Girdhar et al. 2023) or language embeddings (Zhu et al. 2023). However, skeletons require neural networks capable of modeling the implicit spatial configuration of human body kinematics, which differ from traditional ViT encoders. This discrepancy limits their alignment with image or language embeddings, as effective projection within the CLIP embedding space typically requires homogeneous modality-specific encoders. Therefore, we introduce a series of **SK**keleton **I**nduced models, referred to as **SKI** models,

that effectively integrate skeletons within vision-language embeddings, enabling them to understand ADL. Our SKI models are classified as **SKI-VLMs** when skeletons are induced within VLMs, and **SKI-LVLMs** when skeletons are integrated into large vision-language models (LVLMs). The core of SKI models is based on a skeleton-language model, termed **SkeletonCLIP**, which successfully incorporates 3D human skeleton knowledge into the vision-language embeddings. This integration guides the vision encoders to focus on human joints for learning action representations (Fig. 1).

The training of SKI-VLMs involves first aligning modality-specific features (videos or skeletons) with language, followed by performing knowledge distillation (KD) in a language-contextualized space. This KD process, termed SkeletonCLIP Distillation (SCD), involves the collaborative training of SkeletonCLIP and the VLM while conducting KD. Similarly, the training of SKI-LVLMs involves integrating 3D skeleton features extracted from the SkeletonCLIP encoder to enhance the semantic reasoning and generation capabilities of LVLMs. SKI-LVLMs can be trained by learning an additional skeleton projector, which maps skeleton-language tokens into the embedding space of the large language models (LLMs) within the LVLMs. During inference of SKI models, SkeletonCLIP and its components can be discarded, resulting in a skeleton-augmented VLM or LVLM capable of inferring human skeleton knowledge, which is crucial for understanding ADL. We conduct extensive experimental evaluation of SKI-VLM using three VideoCLIP dual encoders to assess its effectiveness in zero-shot action recognition. Additionally, we evaluate the performance of SKI-LVLM on dense video caption generation.

To summarize our contributions:

1. We introduce Skeleton-Induced VLM (SKI-VLM), an effective approach for integrating 3D skeleton information into the VLM space using SkeletonCLIP Distillation (SCD). The resulting SKI-VLM addresses the challenges of ADL by focusing on human key points while modeling action representations.
2. We present SKI-LVLM, which incorporates language-grounded 3D skeletal features as an additional modality in LVLMs, enhancing their video understanding capabilities.
3. We demonstrate the superior performance of SKI-VLM in zero-shot action recognition on the largest RGB+D datasets: NTU60 (Shahroudy et al. 2016) and NTU120 (Liu et al. 2019). We also evaluate SKI-LVLM’s ability to generate text descriptions on the Charades dataset (Sigurdsson et al. 2016), which comprises dense captions. SKI-LVLM outperforms the baseline, highlighting the importance of incorporating skeleton information into LVLMs.

In practice, SKI models can be implemented with any skeleton-language models and VLMs. To the best of our knowledge, this is the first work that attempts to enhance vision-language embeddings by incorporating skeleton information for video representation learning.

2 Related Work

Multi-modal Knowledge Distillation. In standard knowledge distillation (Hinton, Vinyals, and Dean 2015), the knowledge of large-scale models (teachers) is transferred to smaller models (students), enabling students to replicate the teacher’s predictions or feature representations and replace them during inference. Similarly, multi-modal knowledge distillation transfers the knowledge of a teacher model to a student model operating on a different modality. This approach is effective for tasks like action recognition, where complementary modalities exist that are not available or infeasible to compute at inference time (e.g., optical flow or human skeleton). In action recognition, the student modality is typically RGB, and the teacher modality is optical flow (Gupta, Hoffman, and Malik 2016; Crasto et al. 2019) or audio (Aytar, Vondrick, and Torralba 2016). However, these modalities are limited in modeling human motion and generalizing across varying viewpoints, limiting their adoption for ADL (Das et al. 2020). Human skeletons, by contrast, have emerged as the dominant modality to combine with RGB for ADL tasks (Das et al. 2020, 2021; Reilly and Das 2024), as they effectively model human motion and generalize across viewpoints. While effective for understanding ADL, these previous approaches are not applicable to zero-shot tasks. In contrast, our proposed SKI-VLM and SKI-LVLM overcome this limitation by incorporating both modalities with language, enabling zero-shot tasks.

Multi-modal VLMs for Action Recognition. Many approaches propose to extend the image-based CLIP to the video domain through fine-tuning CLIP to handle the temporal dimension of video data. Partially fine-tuned approaches (Yang et al. 2023; Ni et al. 2022; Pan et al. 2022) perform training with additional trainable temporal modules but leave CLIP’s parameters frozen, while fully fine-tuned approaches (Wang, Xing, and Liu 2021; Rasheed et al. 2023; Huang et al. 2024) perform a simple temporal pooling but update the CLIP parameters during training. These approaches only process RGB and ignore the rich multi-modal nature of videos. In response, some works attempt to incorporate additional modalities such as optical flow (Qian et al. 2022) and audio (Guzhov et al. 2021; Ruan et al. 2023; Wu et al. 2021) into the CLIP embedding space. While these approaches all aim to introduce new modalities into CLIP, their methodologies vary. For example, (Guzhov et al. 2021) trains audio, visual, and language encoders using a tri-modal contrastive loss, while (Wu et al. 2021) contrastively aligns an audio encoder with a frozen CLIP visual encoder. Different from these works, we introduce the skeleton modality into the CLIP space to better address the challenges of ADL. Additionally we find, and experimentally validate, that previous alignment strategies are sub-optimal when considering the skeleton modality.

Zero-Shot Skeleton Action Recognition. Zero-shot skeleton-based models aim to enable action classification of unseen classes using only human skeleton sequences. SynSE (Gupta, Sharma, and Sarvadevabhatla 2021) introduces a syntactically guided approach, using part-of-speech tags to enhance the alignment between skeleton and language spaces. SMIE (Zhou et al. 2023) employs mutual

information estimation and maximization to globally align the two spaces. CrossGLG (Yan et al. 2024) uses LLM descriptions of actions and cross-attention to guide a skeleton encoder during training, but only uses the skeleton encoder at inference. Closest to our work is STAR (Chen et al. 2024), which aligns a Shift-GCN (Cheng et al. 2020) skeleton encoder with a pre-trained transformer-based text encoder (Radford et al. 2021) for zero-shot skeleton action recognition. STAR differs from our work in that it does not incorporate the RGB modality or investigate strategies to enhance the representations of VLMs and LVLMs.

Multi-modal Large Language Models. Large language models, such as ChatGPT (Achiam et al. 2023) and LLaMA (Touvron et al. 2023), have exhibited remarkable capabilities in language understanding and generation tasks. These capabilities have inspired many works to extend LLMs to incorporate additional modalities, such as video (Zhang, Li, and Bing 2023; Lin et al. 2023; Maaz et al. 2024), audio (Zhang et al. 2023), and human pose (Feng et al. 2024), typically through instruction tuning or the addition of modality-specific encoders with projection layers. Extending beyond this, other approaches show the possibility to incorporate a wide range of modalities into the LLM space (Lu et al. 2022; Su et al. 2023). Unique from these approaches, our SKI-LVLM targets ADL and aims to train using multiple modalities (vision, language and skeleton), but only use vision and language during inference.

3 Proposed Method

In this section, we present SkeletonCLIP, a specialized skeleton-language model designed to align 3D skeleton representations of human actions with their corresponding language representations. SkeletonCLIP is a pivotal element of our approach and can be substituted with any skeleton-text model (Zhou et al. 2023; Gupta, Sharma, and Sarvadevabhatla 2021; Chen et al. 2024) to enhance the VLM and LVLM embedding space. This is achieved by leveraging its capacity to capture fine-grained motion details and multiple viewpoints effectively. To integrate the crucial skeleton representation into VLMs, we first propose the **Skeleton-Induced VLM (SKI-VLM)**, which seamlessly incorporates SkeletonCLIP features through **SkeletonCLIP Distillation (SCD)**. Furthermore, we introduce the **Skeleton-Induced LVLM (SKI-LVLM)**, wherein SkeletonCLIP features are included as an additional modality within LVLMs. This integration enhances the capability of these models to interpret nuanced and complex actions in video data.

3.1 SkeletonCLIP

SkeletonCLIP is a dual encoder Skeleton-Language model that jointly learns 3D Skeleton and language representations within a common semantic space as illustrated in Figure 2 (left). Unlike VLMs (Radford et al. 2021; Rasheed et al. 2023), SkeletonCLIP faces constraints in utilizing extensive skeleton-text data due to the scarcity of synchronized 3D skeleton sequences and corresponding textual data. It comprises two primary components: a Skeleton Encoder g_s and a Text Encoder g_t . We chose Hyperformer (Zhou et al. 2022)




Figure 2: **(a) SkeletonCLIP**: Utilizes a pretrained Skeleton Backbone, aligned with action class labels from the frozen CLIP Text Encoder during Skeleton-Language Supervision. **(b) SKI-VLM**: Engages in online distillation between SkeletonCLIP and a Vision-Language Model (VLM), both trainable. For inference, the VLM alone performs zero-shot action recognition on unseen classes. **(c) SKI-LVLM**: Projects SkeletonCLIP features into LLM space along with video features. Only the projection layers are trainable, while SkeletonCLIP, CLIP encoder, and LLM are frozen. Inference uses the CLIP vision encoder to extract video features, which, together with the user query, are input to the LLM to generate a response based on the video content.

🔥 Trainable ❄️ Frozen

as the skeleton encoder due to its ability to learn the kinematic dependencies between human joints. This model has shown notable efficacy in Skeleton-based Action Recognition, making it an ideal choice as a skeleton encoder in our approach. Conversely, the text encoder leverages the CLIP text encoder, capitalizing on CLIP’s discriminative representation capabilities to counterbalance the limited availability of large-scale skeleton sequence data. SkeletonCLIP is a plug-and-play module that can be integrated with any VLM, such as XCLIP, ViFiCLIP, or FROSTER, facilitating skeleton-guided visual representation learning.

SkeletonCLIP processes input skeleton sequences $S_i \in \mathbb{R}^{T_s \times 3 \times J}$, where each frame within the sequence T_s comprises J 3D skeleton joints, alongside a text prompt t_j representing the action class label. Initially, the skeleton encoder g_s is pretrained for the task of skeleton action recognition. This pretraining enables SkeletonCLIP to learn a joint embedding with language using an already pretrained CLIP (Radford et al. 2021) text encoder g_t .

The skeleton encoder yields a skeleton representation $z_i^s = \frac{1}{T_s} \sum g_s(S_i)$ for an input sequence S_i from its penultimate layer, whereas the text encoder g_t yields a feature $z_j^t = g_t(t_j)$ for each action class text prompt t_j . To learn a joint skeleton-language embedding, the pretrained skeleton and text encoders are trained jointly while keeping the text encoder frozen. This joint embedding is optimized by maximizing the cosine similarity between skeleton and text

embeddings, quantified by the cross-entropy loss:

$$\mathcal{L}_{CE}(z_i^s, z_i^t) = - \sum_i \log \frac{\exp(\text{sim}(z_i^s, z_i^t)/\tau)}{\sum_j \exp(\text{sim}(z_j^s, z_j^t)/\tau)} \quad (1)$$

where τ is a temperature parameter and $\text{sim}(x, y)$ denotes the cosine similarity between x and y . This joint training enforces the skeleton encoder to align with the language embeddings, enabling it to perform zero-shot skeleton action recognition. This language grounding is essential for incorporating skeleton features into VLMs and LVLMs, which is detailed next.

3.2 Skeleton Induced VLM (SKI-VLM)

Consider a video $V_i \in \mathbb{R}^{T_v \times 3 \times H \times W}$ comprising T_v frames at a spatial resolution of $H \times W$ and a text prompt t_j encapsulating the action categories within a predefined template. To obtain their representations within the VLM embedding space, the input videos and text prompts are processed through the CLIP (Radford et al. 2021) image encoder f_v and text f_t encoder respectively as

$$z_i^v = \frac{1}{T_v} \sum f_v(V_i); \quad z_j^t = f_t(t_j) \quad (2)$$

The video-level representation z_i^v , derived from the CLIP image encoder, is obtained by averaging the feature representations across frames. Typically, the joint embedding $(z_i^v \cdot z_j^t)$ is trained using various finetuning techniques for

learning discriminative video representation. These include fully finetuning (Rasheed et al. 2023) the CLIP embeddings using $\mathcal{L}_{CE}(z_i^v, z_j^t)$, partially finetuning (Ni et al. 2022) using adaptors, or fully finetuning in addition to knowledge distillation from CLIP (Huang et al. 2024). The trained dual encoder model is referred to as a VideoCLIP model, capable of performing zero-shot action recognition in RGB videos.

To integrate 3D skeleton information into the joint video-text embeddings ($z_v \cdot z_t$), we leverage an online knowledge distillation strategy, *SkeletonCLIP Distillation* (SCD). It is performed across the Skeleton-Language (SkeletonCLIP) and Vision-Language (VideoCLIP) models, both of which are pretrained on the training distribution. This pretraining ensures the alignment of the modality-specific features (z_i^v or z_i^s) with the textual features (z_i^t). Since the textual features differ across modalities (videos and skeletons), hereon we will denote the video-based and skeleton-based textual representations as z_i^{tv} and z_i^{ts} , respectively.

SkeletonCLIP Distillation (SCD). For SCD, the model with two dual encoders takes as input video sequences V_i , corresponding 3D skeleton sequences S_i , and text prompts t_j describing action labels, as illustrated in Figure 2 (middle). The VideoCLIP processes the video sequences and text prompts using the encoders f_v and f_t , respectively. Simultaneously, the SkeletonCLIP processes the 3D skeleton sequences and text prompts using the encoders g_s and g_t , respectively. SCD is then applied to integrate visual and skeleton information within a language-contextualized space. This space is constructed by minimizing the similarity between modality-specific features (z_i^v or z_i^s) and text features (z_i^{tv} or z_i^{ts}). The resulting language-contextualized features in the visual and skeleton domains, F_{LV} and F_{LS} , are defined as $F_{LV} = z_i^v \cdot (z_j^{tv})^T$ and $F_{LS} = z_i^s \cdot (z_j^{ts})^T$.

We adopt an online knowledge distillation approach where both the SkeletonCLIP and VideoCLIP models are trainable. This distillation, conducted within the language-contextualized space, aims to minimize the similarity between the respective modalities and the text features, thereby enabling VideoCLIP to infer skeleton-text correlations when applicable. To ensure the relevance of 3D skeleton information for action recognition, we jointly train the SCD with action recognition losses. The total loss (\mathcal{L}) during this training phase is represented as the sum of the cross-entropy losses for video-text and skeleton-text embeddings, combined with a distillation loss \mathcal{L}_D (mean squared error):

$$\mathcal{L} = \mathcal{L}_{CE}(z_i^v, z_i^{tv}) + \mathcal{L}_{CE}(z_i^s, z_i^{ts}) + \alpha \mathcal{L}_D(F_{LV}, F_{LS}) \quad (3)$$

where α is the weight allocated to the distillation loss to balance it with the cross-entropy losses. Through SCD, SkeletonCLIP enriches VideoCLIP with temporal information, resulting in the Skeleton-Induced VideoCLIP (SKI-VLM).

During **inference**, only the VideoCLIP (SKI-VideoCLIP) is utilized, eliminating the need for skeleton data and avoiding additional computational overhead. We implement a teacher-student framework to learn the SKI-VLM representation through SCD, which can be integrated with various student VLMs (Rasheed et al. 2023; Ni et al. 2022; Huang et al. 2024).

3.3 Skeleton Induced LVLM (SKI-LVLM)

3D skeleton features can also be employed to guide visual features, thereby enhancing the semantic reasoning and generation capabilities of Large Vision Language Models (LVLMs). These enhancements are expected to improve the quality of descriptions generated by LVLMs, as they can focus on human joints and their motion while generating textual content. To integrate skeleton information within LVLMs, the features z_i^s for an input skeleton sequence are first extracted from SkeletonCLIP. Given the effectiveness of SCD, where collaborative learning of modality-specific dual encoders proved highly effective, we adopt a similar design approach for integrating skeletons within LVLMs.

In this approach, both visual (z_i^v) and skeleton (z_i^s) features extracted from the RGB+skeleton input sequence are fed into their respective projection layers (\mathcal{T}_v and \mathcal{T}_s), which are then input to the LLM:

$$Q_v = \mathcal{T}_v(z_i^v) \in \mathbb{R}^{F_v \times K}, \quad Q_s = \mathcal{T}_s(z_i^s) \in \mathbb{R}^{F_s \times K} \quad (4)$$

where F_v and F_s represent the dimensionality of the video and skeleton features, respectively, and K is the shared dimensionality to which these features are projected. Similar to the visual projector \mathcal{T}_v , the skeleton projector \mathcal{T}_s learns the mapping of skeleton tokens from the skeleton-language space to the input space of the LLM. Notably, the visual features here are extracted from the CLIP encoder without finetuning on the training distribution, to maintain the generalizability of the RGB information fed to the LLM.

Thus, the input to the LLM consists of Q_v , Q_s , and the tokenized text query $Q_t \in \mathbb{R}^{F_t \times K}$ (where F_t is the dimensionality of the text features), arranged in the following template: **[USER: $\langle Q_t \rangle \langle Q_v \rangle \langle Q_s \rangle$, Assistant:].** SKI-LVLM is trained collaboratively on video-skeleton-text triplets using the autoregressive training objective from (Maaz et al. 2024). The weights of the encoders and the LLM are frozen, and only the projection layers are trained, as illustrated in Figure 2 (right). During **inference**, only the visual input is utilized, eliminating the need for skeleton data and its projector \mathcal{T}_s , making it practical for real-world applications.

4 Experimental Results

Datasets. For *zero-shot* (ZS) *action recognition*, we evaluate our SKI-VLMs on the large-scale NTU-RGB+D-60 (NTU60) (Shahroudy et al. 2016) and NTU-RGB+D-120 (NTU120) (Liu et al. 2019) datasets. NTU120 contains approximately 114K video-pose pairs across 120 action classes, while NTU60 includes around 57K pairs for 60 classes. We adopt the evaluation splits from (Gupta, Sharma, and Sarvadevabhatla 2021), using 55/5 and 48/12 splits for NTU60, and 110/10 and 96/24 splits for NTU120.

For *dense video captioning*, we train our SKI-LVLM on NTU120 video-instruction pairs and evaluate on the Charades dataset (Sigurdsson et al. 2016) following (Reilly et al. 2024). NTU120 video-instruction pairs are generated by captioning single frames (mid-video) with COGVLM (Wang et al. 2023) and creating 100K question-answer pairs using GPT 3.5 turbo. Details on prompts are in the Appendix D. We will release these video-instruction

pairs to the community. We evaluate video captioning performance using Llama 3.1 (Meta 2024) on the five VideoChatGPT (Maaz et al. 2024) metrics: *Correctness of Information*, *Detail Orientation*, *Contextual Understanding*, *Temporal Understanding* and *Consistency*.

Implementation Details. For **SKI-VLMs**, we use XCLIP (Ni et al. 2022), ViFiCLIP (Rasheed et al. 2023), and FROSTER (Huang et al. 2024) as the student VLMs (VideoCLIP). SkeletonCLIP is trained on seen classes of NTU60 (55, 48) and NTU120 (110, 96) for action recognition. The skeleton encoder is pretrained for 140 epochs, followed by alignment with a CLIP Text encoder for 100 epochs. SCD requires 2 epochs. We use a learning rate of 2.25×10^{-5} with cosine decay and set α to 0.01 for NTU60 and 10.0 for NTU120. SKI-XCLIP, SKI-ViFiCLIP, and SKI-FROSTER denote SKI-VLM with respective backbones. For **SKI-LVLM**, embedding dimensions are $F_v = 1024$, $F_s = 216$, $K = 4096$. The input to the projection layers (\mathcal{T}_v & \mathcal{T}_s) preceding the LLM are 356 visual and 256 skeleton tokens. We train SKI-LVLM and its baselines for 3 epochs with batch size 32, learning rate $2e^{-5}$ on 8 A6000 48GB GPUs.

Method	Modality			NTU60 48/12	NTU120 110/10
	V	T	S		
Tri-modal Align. (Guzhov et al. 2021)	✗	✓	✓	8.3	11.9
Cross-projection Align. (Wu et al. 2021)	✗	✓	✓	8.4	25.5
XCLIP (Ni et al. 2022)	✓	✓	✗	38.9	57.8
FROSTER (Huang et al. 2024)	✓	✓	✗	43.9	65.2
ViFiCLIP (Rasheed et al. 2023)	✓	✓	✗	48.2	70.0
SkeletonCLIP	✗	✓	✓	35.1	63.0
ViFiCLIP + SkeletonCLIP	✓	✓	✓	17.5	48.8
SKI-XCLIP	✓	✓	○	42.2	66.1
SKI-FROSTER	✓	✓	○	44.4	68.5
SKI-ViFiCLIP	✓	✓	○	52.0	77.5

Table 1: Performance comparison of methods for aligning skeleton features with video-text embeddings on NTU60 and NTU120 datasets, demonstrating the effectiveness of integrating language-contextualized 3D skeleton features into the CLIP embedding space. *V*, *S* and *T* denote the video, skeleton, and text modalities respectively. (○ indicates that skeleton features were used only during training)

4.1 ZS Action Recognition using SKI-VLM

For zero-shot action recognition, we evaluate models on unseen action classes within each split. We first demonstrate the challenge of learning a common embedding space for video, text, and skeletons. We then justify the superiority of VideoCLIP architectures over SkeletonCLIP and highlight the robustness of SKI-VLMs. Finally, we present state-of-the-art zero-shot action recognition results on NTU datasets.

How can we introduce 3D skeleton features in the CLIP embedding space? In the audio domain, methods such as AudioCLIP (Guzhov et al. 2021) and Wav2CLIP (Wu et al. 2021) employ contrastive learning and cross-projection techniques to align audio with the CLIP embedding space. Inspired by these approaches, we explore aligning skeleton features with the video-text embedding space. However, as shown in Table 1, neither tri-modal alignment (implemented

following AudioCLIP) nor cross-projection alignment (implemented following Wav2CLIP) effectively aligns skeleton and text representations. This limitation stems from the lack of large-scale symmetric video-skeleton-text datasets necessary for effective contrastive learning. The results of SKI-VLMs highlights the effectiveness of integrating language-contextualized 3D skeleton features into the vision-language embedding space, enhancing the learning of generalized action representations.

SkeletonCLIP vs VideoCLIP. SkeletonCLIP grounds skeleton features within a language-contextualized space. As shown in Table 1, SkeletonCLIP achieves performance comparable to VideoCLIP models like XCLIP, FROSTER, and ViFiCLIP. To further enhance zero-shot recognition performance, we combined the strengths of VideoCLIP and SkeletonCLIP by directly fusing language-contextualized skeleton features with VLM features. However, this approach led to poorer action recognition performance, likely due to conflicting gradients from the different modalities. Thus, VLMs that incorporate only video and text modalities demonstrate more promising performance in zero-shot action recognition compared to skeleton-only models.

Robustness of SKI-VLMs. Our SKI-VLMs can be implemented using any existing dual encoders. In Table 1, we demonstrate the implementation of SKI-VLMs with XCLIP, ViFiCLIP, and FROSTER, where they outperform their respective baselines by up to **+14.3%**, **+10.7%**, and **+5.1%**. This highlights the robustness of our SKI-VLMs when integrated with any student VLM.

Method	Modality			NTU60		NTU120	
	V	T	S	55/5	48/12	110/10	96/24
SynSE (Gupta, Sharma, and Sarvadevabhatla 2021)	✗	✓	✓	75.8	33.3	62.7	38.7
SMIE (Zhou et al. 2023)	✗	✓	✓	78.0	40.2	65.7	45.3
STAR (Chen et al. 2024)	✗	✓	✓	81.4	45.1	63.3	44.3
CLIP (Radford et al. 2021)	✓	✓	✗	54.5	20.2	35.0	15.2
XCLIP (Ni et al. 2022)	✓	✓	✗	76.0	38.9	57.8	49.7
ViFiCLIP (Rasheed et al. 2023)	✓	✓	✗	79.9	48.2	70.0	56.6
FROSTER (Huang et al. 2024)	✓	✓	✗	79.1	43.9	65.2	34.0
SKI-ViFiCLIP	✓	✓	○	82.2	52.0	77.5	59.3

Table 2: Comparison of Zero-Shot Action Recognition Accuracy on NTU60 and NTU120 datasets. *V*, *S* and *T* denote the video, skeleton, and text modalities respectively. (○ indicates that skeleton features were used only during training)

State-of-the-art comparison. In Table 2, we benchmark SKI-VLMs against leading skeleton-text and video-text zero-shot action recognition models. The inferior performance of skeleton-text models like SynSE (Gupta, Sharma, and Sarvadevabhatla 2021) and SMIE (Zhou et al. 2023) can be attributed to their lack of appearance information. Furthermore, adapting these models to different data distributions is challenging due to variations in skeleton configurations arising from disparate depth sensors (Zhang 2012) or skeleton extraction techniques (Rogez, Weinzaepfel, and Schmid 2019). We also note that FROSTER (Huang et al. 2024), despite outperforming ViFiCLIP (Rasheed et al. 2023) and XCLIP (Ni et al. 2022) in web video datasets such as Kinetics (Kay et al. 2017), UCF-101 (Soomro, Zamir, and

Metric	VCGPT	NTU-VCGPT	SK-VCGPT	SKI-LVLM
Correctness	24.8	25.5	40.6	40.6
Detail Orientation	43.2	21.6	55.0	54.8
Contextual Understanding	33.8	29.0	52.4	52.4
Temporal Understanding	20.0	16.9	30.6	31.6
Consistency	31.6	34.1	38.4	38.2
Average	30.7	25.4	43.4	43.5

Table 3: Dense Caption Generation Performance of SkeletonCLIP induced LVLM on Charades Dataset (VCGPT = VideoChatGPT)

Shah 2012), and HMDB (Kuehne et al. 2011), falls short in ADL datasets, showing the unique challenges posed by ADL compared to web videos. FROSTER’s reliance on image-based CLIP knowledge transfer is less effective on NTU due to the lack of scene-specific contextual information in NTU.

SKI-ViFiCLIP surpasses all other zero-shot action recognition models. Notably, applying SCD to a video-based CLIP model (ViFiCLIP) significantly enhances its zero-shot action recognition performance on all NTU60 and NTU120 splits by up to +5.4% and +7.8%, respectively. This demonstrates the enhanced “generalizability” of SKI-VLM when incorporating skeletons into VLMs.

4.2 Dense Video Captioning using SKI-LVLM

In Table 3, we assess the video captioning performance of SKI-LVLM using the five metrics proposed in (Maaz et al. 2024), scaled to a range of 1-100 following (Reilly et al. 2024). SKI-LVLM is benchmarked against VideoChatGPT (VCGPT) and NTU trained VideoChatGPT (NTU-VCGPT). VCGPT is trained on a large-scale Instruction Tuning dataset derived from ActivityNet (Caba Heilbron et al. 2015), while NTU-VCGPT denotes VideoChatGPT trained exclusively on video-instruction pairs from trimmed NTU120 videos (see Ablation D for details). Additionally, we include a variant of SKI-LVLM, denoted SK-VCGPT, which is trained on NTU120 instruction pairs and corresponding 3D skeleton features obtained from a pre-trained skeleton backbone (Hyperformer). This allows us to investigate the value of language-contextualized skeleton features in SKI-LVLM, as SK-VCGPT incorporates raw 3D skeleton features without language contextualization.

Surprisingly, we find that SK-VCGPT performs comparably with SKI-LVLM, contrasting with our findings in traditional VLMs where language contextualization is essential. This suggests that feature integration mechanisms may differ in LVLMs, indicating that language contextualization in LVLMs plays a less critical role. However, we observe that the introduction of 3D skeletons greatly improves the performance of the LVLMs, highlighting the importance of integrating skeleton features into LVLM training.

4.3 Ablation Study

All our ablation studies have been conducted with ViFiCLIP as the dual encoder (VideoCLIP).

Variants of Knowledge Distillation (KD). KD between SkeletonCLIP and VideoCLIP can be implemented using different strategies. In Table 4, we compare SCD with

Method	NTU60 48/12	NTU120 110/10	Harmonic Mean
Feature-level KD <i>w/o Projection</i>	52.9	65.3	58.4
Feature-level KD <i>w Projection</i>	53.2	64.7	58.4
Offline KD	53.3	70.8	60.8
Online KD (ours)	52.0	77.5	62.2

Table 4: Performance on different Distillation techniques

Loss Function	NTU60 48/12	NTU120 110/10
Contrastive Loss	46.9	70.0
KL Divergence Loss	48.5	73.1
Mean Squared Error (MSE) Loss	52.0	77.5

Table 5: Effect of choice of distillation loss functions

feature-level distillation and offline KD. Feature-level distillation, used in video-skeleton action recognition to enhance RGB-based encoders (Das et al. 2021; Reilly and Das 2024), is sub-optimal for large-scale datasets in SKI-VLMs. Feature space mismatches between the teacher (SkeletonCLIP) and student (VideoCLIP) are magnified in large-scale settings, leading to inefficiencies in knowledge transfer. Online distillation at the logit level, directly aligning network outputs during training, offers a more robust solution for handling large-scale data. Additionally, using a projection layer to map features into a common space for distillation, as explored in (Huang et al. 2024), improves performance on NTU60 but doesn’t consistently boost accuracy across larger datasets. Unlike online KD in SCD, where both networks are trainable, offline KD freezes the teacher and trains only the student. Our findings suggest that while offline and feature-level KD perform better in low-data settings, online KD consistently outperforms them in large-scale datasets. This highlights the benefit of collaborative training between SkeletonCLIP and VideoCLIP, allowing the latter to acquire meaningful representations and achieve higher accuracy.

Loss Configuration. In Table 5, we perform an ablation study on the choice of distillation loss for the implementation of SCD. We find that employing MSE for SCD results in notably improved performance compared to using Contrastive Loss or KL Divergence Loss.

5 Conclusion

In this paper, we introduced Skeleton-Induced (SKI) models, including SKI-VLMs and SKI-LVLMs, to integrate skeleton information into vision-language embeddings. This enables the models to focus on human keypoints for better action modeling, crucial for learning discriminative video representations in ADL. Our experiments show that SKI-VLMs achieve significant gains in zero-shot action recognition without needing skeletons during inference, while SKI-LVLM improves LVLMs’ understanding of subtle actions, enhancing semantic reasoning and dense caption generation.

This work is a first step towards building multimodal foundation models that incorporate the human skeleton modality. Future research will explore replacing SkeletonCLIP with other effective skeleton-language models.

Acknowledgements

We thank the members of the Genius Lab at UNC Charlotte for helpful discussions. This work is supported in part by the National Science Foundation (IIS-2245652). Additionally, this material is based on research partially supported by the Chateaubriand Fellowship of the Office for Science and Technology of the Embassy of France in the United States.

References

- Achiam, J.; Adler, S.; Agarwal, S.; Ahmad, L.; Akkaya, I.; Aleman, F. L.; Almeida, D.; Altschmidt, J.; Altman, S.; Anadkat, S.; et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Aytar, Y.; Vondrick, C.; and Torralba, A. 2016. SoundNet: Learning Sound Representations from Unlabeled Video. In *Proceedings of the 30th International Conference on Neural Information Processing Systems, NIPS'16*, 892–900. Red Hook, NY, USA: Curran Associates Inc. ISBN 9781510838819.
- Caba Heilbron, F.; Escorcia, V.; Ghanem, B.; and Carlos Niebles, J. 2015. Activitynet: A large-scale video benchmark for human activity understanding. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 961–970.
- Chen, Y.; Guo, J.; He, T.; and Wang, L. 2024. Fine-Grained Side Information Guided Dual-Prompts for Zero-Shot Skeleton Action Recognition.
- Cheng, K.; Zhang, Y.; He, X.; Chen, W.; Cheng, J.; and Lu, H. 2020. Skeleton-Based Action Recognition with Shift Graph Convolutional Network. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Chiang, W.-L.; Li, Z.; Lin, Z.; Sheng, Y.; Wu, Z.; Zhang, H.; Zheng, L.; Zhuang, S.; Zhuang, Y.; Gonzalez, J. E.; Stoica, I.; and Xing, E. P. 2023. Vicuna: An Open-Source Chatbot Impressing GPT-4 with 90%* ChatGPT Quality.
- Craato, N.; Weinzaepfel, P.; Alahari, K.; and Schmid, C. 2019. MARS: Motion-Augmented RGB Stream for Action Recognition. In *CVPR*.
- Darcet, T.; Oquab, M.; Mairal, J.; and Bojanowski, P. 2023. Vision Transformers Need Registers. *ArXiv*, abs/2309.16588.
- Das, S.; Dai, R.; Koperski, M.; Minciullo, L.; Garattoni, L.; Bremond, F.; and Francesca, G. 2019. Toyota Smarthome: Real-World Activities of Daily Living. In *Int. Conf. Comput. Vis.*
- Das, S.; Dai, R.; Yang, D.; and Bremond, F. 2021. VPN++: Rethinking Video-Pose embeddings for understanding Activities of Daily Living. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 1–1.
- Das, S.; Sharma, S.; Dai, R.; Bremond, F.; and Thonnat, M. 2020. Vpn: Learning video-pose embedding for activities of daily living. In *European Conference on Computer Vision*, 72–90. Springer.
- Duan, H.; Zhao, Y.; Chen, K.; Lin, D.; and Dai, B. 2022. Revisiting Skeleton-based Action Recognition. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2959–2968.
- Feng, Y.; Lin, J.; Dwivedi, S. K.; Sun, Y.; Patel, P.; and Black, M. J. 2024. ChatPose: Chatting about 3D Human Pose. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Girdhar, R.; El-Nouby, A.; Liu, Z.; Singh, M.; Alwala, K. V.; Joulin, A.; and Misra, I. 2023. Imagebind: One embedding space to bind them all. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 15180–15190.
- Gupta, P.; Sharma, D.; and Sarvadevabhatla, R. K. 2021. Syntactically Guided Generative Embeddings for Zero-Shot Skeleton Action Recognition. *2021 IEEE International Conference on Image Processing (ICIP)*, 439–443.
- Gupta, S.; Hoffman, J.; and Malik, J. 2016. Cross Modal Distillation for Supervision Transfer.
- Guzhov, A.; Raue, F.; Hees, J.; and Dengel, A. R. 2021. Audioclip: Extending Clip to Image, Text and Audio. *ICASSP 2022 - 2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 976–980.
- Hinton, G.; Vinyals, O.; and Dean, J. 2015. Distilling the Knowledge in a Neural Network. *arXiv:1503.02531*.
- Huang, X.; Zhou, H.; Yao, K.; and Han, K. 2024. FROSTER: Frozen CLIP is a Strong Teacher for Open-Vocabulary Action Recognition. In *International Conference on Learning Representations*.
- Kay, W.; Carreira, J.; Simonyan, K.; Zhang, B.; Hillier, C.; Vijayanarasimhan, S.; Viola, F.; Green, T.; Back, T.; Natsev, P.; et al. 2017. The kinetics human action video dataset. *arXiv preprint arXiv:1705.06950*.
- Kuehne, H.; Jhuang, H.; Garrote, E.; Poggio, T.; and Serre, T. 2011. HMDB: a large video database for human motion recognition. In *2011 International Conference on Computer Vision*, 2556–2563. IEEE.
- Lin, B.; Zhu, B.; Ye, Y.; Ning, M.; Jin, P.; and Yuan, L. 2023. Video-LLaVA: Learning United Visual Representation by Alignment Before Projection. *arXiv preprint arXiv:2311.10122*.
- Liu, J.; Shahroudy, A.; Perez, M.; Wang, G.; Duan, L.-Y.; and Kot, A. C. 2019. NTU RGB+D 120: A Large-Scale Benchmark for 3D Human Activity Understanding. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Lu, J.; Clark, C.; Zellers, R.; Mottaghi, R.; and Kembhavi, A. 2022. Unified-IO: A Unified Model for Vision, Language, and Multi-Modal Tasks. *arXiv preprint arXiv:2206.08916*.
- Maaz, M.; Rasheed, H.; Khan, S.; and Khan, F. S. 2024. Video-ChatGPT: Towards Detailed Video Understanding via Large Vision and Language Models. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (ACL 2024)*.
- Meta. 2024. The Llama 3 Herd of Models. *arXiv:2407.21783*.

- Ni, B.; Peng, H.; Chen, M.; Zhang, S.; Meng, G.; Fu, J.; Xiang, S.; and Ling, H. 2022. Expanding Language-Image Pretrained Models for General Video Recognition.
- Pan, J.; Lin, Z.; Zhu, X.; Shao, J.; and Li, H. 2022. ST-Adapter: Parameter-Efficient Image-to-Video Transfer Learning. In Koyejo, S.; Mohamed, S.; Agarwal, A.; Belgrave, D.; Cho, K.; and Oh, A., eds., *Advances in Neural Information Processing Systems*, volume 35, 26462–26477. Curran Associates, Inc.
- Qian, R.; Li, Y.; Xu, Z.; Yang, M.-H.; Belongie, S.; and Cui, Y. 2022. Multimodal open-vocabulary video classification via pre-trained vision and language models. *arXiv preprint arXiv:2207.07646*.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; Krueger, G.; and Sutskever, I. 2021. Learning Transferable Visual Models From Natural Language Supervision. In *International Conference on Machine Learning*.
- Rasheed, H.; Khattak, M. U.; Maaz, M.; Khan, S.; and Khan, F. S. 2023. Finetuned CLIP models are efficient video learners. In *The IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- Reilly, D.; Chakraborty, R.; Sinha, A.; Govind, M. K.; Wang, P.; Bremond, F.; Xue, L.; and Das, S. 2024. LLAVIDAL: A Large LAnGuage ViSion Model for Daily Activities of Living. *arXiv*, 2406.09390.
- Reilly, D.; and Das, S. 2024. Just Add π ! Pose Induced Video Transformers for Understanding Activities of Daily Living. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Rogez, G.; Weinzaepfel, P.; and Schmid, C. 2019. LCR-Net++: Multi-person 2D and 3D Pose Detection in Natural Images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Ruan, L.; Hu, A.; Song, Y.; Zhang, L.; Zheng, S.; and Jin, Q. 2023. Accommodating Audio Modality in CLIP for Multimodal Processing. In *AAAI Conference on Artificial Intelligence*.
- Shahroudy, A.; Liu, J.; Ng, T.-T.; and Wang, G. 2016. NTU RGB+D: A Large Scale Dataset for 3D Human Activity Analysis. In *IEEE Conf. Comput. Vis. Pattern Recog.*
- Shi, L.; Zhang, Y.; Cheng, J.; and Lu, H. 2019. Two-stream adaptive graph convolutional networks for skeleton-based action recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 12026–12035.
- Sigurdsson, G. A.; Varol, G.; Wang, X.; Farhadi, A.; Laptev, I.; and Gupta, A. 2016. Hollywood in Homes: Crowdsourcing Data Collection for Activity Understanding. In *European Conference on Computer Vision (ECCV)*.
- Simonyan, K.; and Zisserman, A. 2014. Two-stream convolutional networks for action recognition in videos. In *Advances in neural information processing systems*, 568–576.
- Soomro, K.; Zamir, A. R.; and Shah, M. 2012. UCF101: A Dataset of 101 Human Actions Classes From Videos in The Wild. *CoRR*, abs/1212.0402.
- Su, Y.; Lan, T.; Li, H.; Xu, J.; Wang, Y.; and Cai, D. 2023. PandaGPT: One Model To Instruction-Follow Them All. *arXiv preprint arXiv:2305.16355*.
- Sun, Q.; Fang, Y.; Wu, L.; Wang, X.; and Cao, Y. 2023. Eva-clip: Improved training techniques for clip at scale. *arXiv preprint arXiv:2303.15389*.
- Touvron, H.; Lavril, T.; Izacard, G.; Martinet, X.; Lachaux, M.-A.; Lacroix, T.; Rozière, B.; Goyal, N.; Hambro, E.; Azhar, F.; Rodriguez, A.; Joulin, A.; Grave, E.; and Lample, G. 2023. LLaMA: Open and Efficient Foundation Language Models. *ArXiv*, abs/2302.13971.
- Wang, M.; Xing, J.; and Liu, Y. 2021. ActionCLIP: A New Paradigm for Video Action Recognition. *ArXiv*, abs/2109.08472.
- Wang, W.; Lv, Q.; Yu, W.; Hong, W.; Qi, J.; Wang, Y.; Ji, J.; Yang, Z.; Zhao, L.; Song, X.; Xu, J.; Xu, B.; Li, J.; Dong, Y.; Ding, M.; and Tang, J. 2023. CogVLM: Visual Expert for Pretrained Language Models. *ArXiv*, abs/2311.03079.
- Wu, H.-H.; Seetharaman, P.; Kumar, K.; and Bello, J. P. 2021. Wav2CLIP: Learning Robust Audio Representations from Clip. *ICASSP 2022 - 2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 4563–4567.
- Yan, S.; Xiong, Y.; and Lin, D. 2018. Spatial temporal graph convolutional networks for skeleton-based action recognition. In *Thirty-second AAAI conference on artificial intelligence*.
- Yan, T.; Zeng, W.; Xiao, Y.; Tong, X.; Tan, B.; Fang, Z.; Cao, Z.; and Zhou, J. T. 2024. CrossGLG: LLM Guides One-shot Skeleton-based 3D Action Recognition in a Cross-level Manner. *arXiv preprint arXiv:2403.10082*.
- Yang, T.; Zhu, Y.; Xie, Y.; Zhang, A.; Chen, C.; and Li, M. 2023. AIM: Adapting Image Models for Efficient Video Understanding. In *International Conference on Learning Representations*.
- Zhang, D.; Li, S.; Zhang, X.; Zhan, J.; Wang, P.; Zhou, Y.; and Qiu, X. 2023. SpeechGPT: Empowering Large Language Models with Intrinsic Cross-Modal Conversational Abilities. *arXiv:2305.11000*.
- Zhang, H.; Li, X.; and Bing, L. 2023. Video-LLaMA: An Instruction-tuned Audio-Visual Language Model for Video Understanding. *ArXiv*, abs/2306.02858.
- Zhang, Z. 2012. Microsoft Kinect Sensor and Its Effect. In *IEEE MultiMedia*, volume 19.
- Zhou, Y.; Cheng, Z.-Q.; Li, C.; Geng, Y.; Xie, X.; and Keuper, M. 2022. Hypergraph Transformer for Skeleton-based Action Recognition. *arXiv preprint arXiv:2211.09590*.
- Zhou, Y.; Qiang, W.; Rao, A.; Lin, N.; Su, B.; and Wang, J. 2023. Zero-shot Skeleton-based Action Recognition via Mutual Information Estimation and Maximization. In *Proceedings of the 31st ACM International Conference on Multimedia*, 5302–5310.
- Zhu, B.; Lin, B.; Ning, M.; Yan, Y.; Cui, J.; Wang, H.; Pang, Y.; Jiang, W.; Zhang, J.; Li, Z.; Zhang, W.; Li, Z.; Liu, W.; and Yuan, L. 2023. LanguageBind: Extending Video-Language Pretraining to N-modality by Language-based Semantic Alignment. *ArXiv*, abs/2310.01852.

Appendix

A Overview

The Supplementary material is organized as follows:

- Section B: Effectiveness of SKI-VLM for ADL
- Section C: Ablations
 - Section C.1: SkeletonCLIP Text Encoder
 - Section C.2: Pretraining Strategy
- Section D: LVLM Video Instruction Data Curation
- Section E: Integrating SkeletonCLIP in VLMs
- Section F: Distillation Loss Weightage




Figure 3: **Attention Map Visualization:** Comparison between ViFiCLIP and SKI-VLM. While ViFiCLIP struggles to identify the critical areas responsible for actions, SKI-VLM accurately focuses on the relevant joints, such as hands and face, for actions like *Sneeze/Cough*.

B Effectiveness of SKI-VLM for ADL

In Figure 3, we present the attention map visualizations¹ of ViFiCLIP and SKI-VLM on a sample from the NTU dataset. These heatmaps highlight the important regions of the video frames that contributed most to the model’s prediction. In the figure, the second row consists of the attention maps of the ViFiCLIP model, where the heatmaps are dispersed across large areas, indicating a broad focus. In contrast, the final row shows SKI-VLMs GradCAM visualization, which precisely emphasizes the relevant human joints involved in the action. This demonstrates that incorporating skeleton information through SKI-VLM enables a more targeted focus on the critical regions of the video compared to ViFiCLIP.

C Ablations

C.1 SkeletonCLIP Text Encoder

In this section, we evaluate the configuration of the SkeletonCLIP Text Encoder, specifically examining whether the

¹The visualizations tend to emphasize static elements of the background due to absence of registers during training, as discussed in (Darcet et al. 2023). These high background values are disregarded in our analysis.

Table 6: Effect of SkeletonCLIP Text Encoder on SKI-ViFiCLIP

Method	NTU60 48/12	NTU120 110/10	Harmonic Mean
Trainable Text Encoder	52.4	72.2	60.7
Frozen Text Encoder (ours)	52.0	77.5	62.2

text encoder should be trainable or frozen. As shown in Table 6, SKI-ViFiCLIP achieves superior performance on average with a frozen Text Encoder compared to a trainable one. This outcome highlights the effectiveness of a frozen Text Encoder in aligning the skeleton distribution more closely with the pre-aligned text and video distributions, enhancing skeleton-text-video alignment during the SkeletonClip Distillation.

C.2 Pretraining Strategy

Table 7: Pretraining Strategy of Components of SKI-VLM

Method	Pretraining	NTU60 48/12	NTU120 110/10	Harmonic Mean
Online Distillation		50.0	10.0	16.7
	+ SkeletonCLIP	49.2	38.7	43.3
	+ ViFiCLIP	53.0	69.6	60.2
	+ SkeletonCLIP + ViFiCLIP	52.0	77.5	62.2

To validate our model’s pretraining strategy, we conducted experiments across various scenarios: no pretraining, SkeletonCLIP pretraining, ViFiCLIP pretraining, and combined SkeletonCLIP + ViFiCLIP pretraining. As shown in Table 7, the model demonstrates superior performance when both SkeletonCLIP and ViFiCLIP are pretrained. This result supports our hypothesis that pretraining and aligning SkeletonCLIP are essential for it to effectively serve as a teacher model during distillation. Pretraining the skeleton encoder is critical for aligning skeleton and text representations within SkeletonCLIP. Likewise, a pretrained ViFiCLIP plays a key role in extracting discriminative video-text representations from the input videos and text prompts.

D LVLM Video Instruction Data Curation

To generate the the video instruction data for training the LVLM, for each video we first crop out the person(s) performing the action. This helps in eliminating unnecessary background information in the videos. Subsequently a single frame is selected from the video and CogVLM (Wang et al. 2023) is used to generate its caption. CogVLM employs Vicuna v1.5 7B (Chiang et al. 2023) as its primary language model and EVA2-CLIP-E (Sun et al. 2023) as the visual transformer encoder, with input images set at 224 × 224 pixels. The prompt used for captioning is - "Give a detailed description of the actions happening and describe the image, include motions and the objects interacted by the person."




Figure 4: Illustration of how SkeletonCLIP can be easily integrated with VLMs like ViFiCLIP, FROSTER and XCLIP (left to right)

E Integrating SkeletonCLIP in VLMs

Figure 4 demonstrates SKI-VLM implemented using SkeletonCLIP as the teacher and XCLIP, ViFiCLIP, and FROSTER as the student models for performing SCD.

F Distillation Loss Weightage




Figure 5: Impact of α in SKI-VLM for NTU48 and NTU110

In Figure 5, we report the accuracy of SKI-VLM on NTU48 and NTU110 for various values of α . We observe that for NTU48, the optimal α is 0.01, while for NTU110, an α value of 10.0 yields the best accuracy. This indicates that in the low data regime (NTU48), placing greater emphasis on the classification loss (cross entropy) is crucial to ensure the model effectively learns from the limited data. Conversely, in the large-scale dataset (NTU110), a higher α value allows for more effective leveraging of the SkeletonCLIP through the distillation loss, helping the model to better generalize by aligning with the teacher model's guidance.