



HAL
open science

COMEDIAN: Self-supervised learning and knowledge distillation for action spotting using transformers

Julien Denize, Mykola Liashuha, Jaonary Rabarisoa, Astrid Orcesi, Romain Hérault

► **To cite this version:**

Julien Denize, Mykola Liashuha, Jaonary Rabarisoa, Astrid Orcesi, Romain Hérault. COMEDIAN: Self-supervised learning and knowledge distillation for action spotting using transformers. WACV 2024 - IEEE/CVF Winter Conference on Applications of Computer Vision, Jan 2024, Waikoloa, United States. 10.48550/arXiv.2309.01270 . hal-04462159

HAL Id: hal-04462159

<https://hal.science/hal-04462159>

Submitted on 16 Feb 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

COMEDIAN: Self-Supervised Learning and Knowledge Distillation for Action Spotting using Transformers

Julien Denize^{*†} Mykola Liashuha^{*} Jaonary Rabarisoa^{*} Astrid Orcesi^{*}
Romain Hérault[‡]

^{*} Université Paris-Saclay, CEA, List, F-91120, Palaiseau, France

`firstname.lastname@cea.fr`

[†] Normandie Univ., INSA Rouen, LITIS, 76801, Saint Etienne du Rouvray, France

`firstname.lastname@insa-rouen.fr`

[‡] Normandie Univ., UNICAEN, ENSICAEN, CNRS, GREYC, 14000 Caen, France

`firstname.lastname@unicaen.fr`

Abstract

We present COMEDIAN, a novel pipeline to initialize spatiotemporal transformers for action spotting, which involves self-supervised learning and knowledge distillation. Action spotting is a timestamp-level temporal action detection task. Our pipeline consists of three steps, with two initialization stages. First, we perform self-supervised initialization of a spatial transformer using short videos as input. Additionally, we initialize a temporal transformer that enhances the spatial transformer’s outputs with global context through knowledge distillation from a pre-computed feature bank aligned with each short video segment. In the final step, we fine-tune the transformers to the action spotting task. The experiments, conducted on the SoccerNet-v2 dataset, demonstrate state-of-the-art performance and validate the effectiveness of COMEDIAN’s pretraining paradigm. Our results highlight several advantages of our pretraining pipeline, including improved performance and faster convergence compared to non-pretrained models. Source code is available here: <https://github.com/juliendenize/eztorch>.

1. Introduction

The field of computer vision has witnessed remarkable progress in recent years, and in particular in video analysis since the Deep Learning era. To make the best use of Deep Learning models, vast amounts of data and high computational power are needed. Temporal Action Detection (TAD) has been gaining attention as it has different applications in our lives ranging from home automation to sports analysis. TAD involves the identification of when specific actions occur within a video, enabling comprehensive understanding

and meaningful insights into the dynamics of a given scenario. Action Spotting [23] is a specific TAD task whose goal is to predict actions at a precise timestamp and therefore requires a temporally precise prediction. Modeling actions in videos faces several issues such as the sparsity of actions and the intricate relationships between them.

While a lot of video data is readily available, annotating actions presents significant challenges. It suffers from the inherent subjectivity of annotators to interpret when an action starts or ends. Moreover, the process of manual annotation requires considerable time and resources, limiting the scalability of annotating large datasets.

Recent advancements in vision Transformers [17,35] for video analysis [1,36,47] showed them surpassing the traditionally used Convolutional Neural Networks (CNNs) but required dedicated architecture design to reduce computational cost [1,3,39,41,58–60]. By capturing long-range dependencies and leveraging global context, Transformers analyses better complex sequences. However, their effectiveness crucially depends on proper initialization and access to ample labeled training data [1,36].

The use of pretraining methods has shown tremendous progress in enhancing the capabilities of Transformers [1,7,36]. Pretraining can be achieved through two approaches: Supervised Learning (SL) using labeled data that requires extensive annotation efforts, and Self-Supervised Learning (SSL) that leverages unlabeled data. SSL for video representation learning [15,21,52,57] has shown promising results, demonstrating significant improvements in generalization and narrowing the gap with supervised learning approaches. Knowledge Distillation (KD) [28] also serves as a powerful tool to initialize a network by transferring knowledge from another network or a collection of models. De-

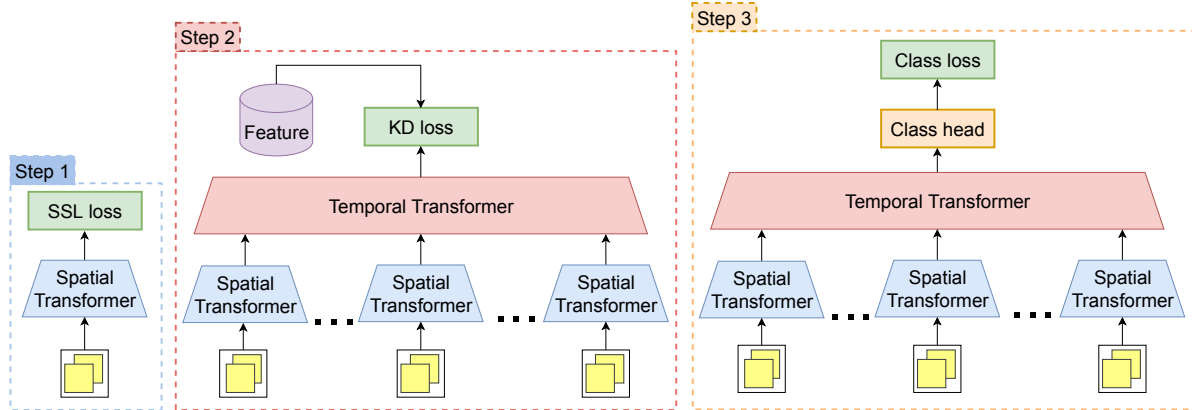


Figure 1. Overview of COMEDIAN training pipeline. Step 1: Pretraining of the spatial transformer. Step 2: Pretraining of the spatial and temporal transformers. Step 3: fine-tuning to the action spotting task.

pending on how the networks are obtained and the losses used to distill, KD can be considered as SL [53] or SSL [22].

In this work, we propose the COMEDIAN approach, which combines self-supervised learning and knowledge distillation to initialize a spatio-temporal transformer for action spotting. Our method involves two transformers: a spatial transformer, which learns short context information from frames extracted from small videos, and a temporal transformer, which enriches the local context with global information. The initialization process consists of two stages: the first stage focuses on the spatial transformer via SSL, while the second stage initializes both spatial and temporal transformers via KD. KD is employed from a pre-computed bank of representations aligned with each output temporal token. Importantly, COMEDIAN leverages unlabeled video data for initialization, effectively addressing the aforementioned challenges associated with transformers.

To evaluate our approach, we conducted experiments on the action spotting task on the SoccerNet-v2 [14] dataset, which contains soccer matches with 17 distinct actions varying in semantics and occurrence. Our contributions can be summarized as follows:

- We propose COMEDIAN a combined self-supervised learning and knowledge distillation pipeline illustrated in Fig. 1 to initialize transformers for action spotting.
- We demonstrate that COMEDIAN achieves state-of-the-art performance for action spotting on the SoccerNet-v2 dataset, showcasing the effectiveness of our self-supervised and knowledge distillation pipeline.
- We provide a comprehensive analysis of the benefits of pretraining with knowledge distillation, including improved performance and faster convergence compared to non-pretrained models.

2. Related Work

Video Transformers. Vision Transformers [17] (ViT) capture long-term dependencies better than recurrent models or convolutional networks. It relies on a tokenizer, that embeds patches of the input, and self-attention [55]. Standard self-attention is computationally heavy and several other attention mechanisms have been proposed such as Swin [35] or DeiT [53]. Transformers can be applied to videos by adapting the tokenizer [1, 3, 36]. However as videos increase the number of tokens, video transformers are computationally heavy, and various strategies have been proposed to reduce their cost. VTN [41] adds a temporal encoder on top of a ViT while ViViT [1] and TimeFormer [3] propose several factorizations of space-time attention. ViViT as VTN found a spatio-temporal hierarchical model offers the best trade-off between performance and cost which led to several methods [4, 31]. Previously mentioned transformers focused on short videos, e.g. ≤ 5 seconds and some architectures have been developed to capture long-range dependencies on longer videos via a sliding window that keeps relevant information from the past with a memory [58, 59] or a recurrence [39, 60] mechanism.

In our work, we consider a spatio-temporal hierarchical model without changing specifically the architecture to capture long-term dependencies as we also want to capture bidirectional short-term dependencies.

Pretraining. Pretraining has been crucial to unleashing image [17] and video [1] transformers either via Supervised Learning (SL) on a large dataset [17] or Self-Supervised Learning (SSL) [7, 10]. Contrastive learning [54] is a kind of state-of-the-art SSL on images [6, 9, 16, 18, 25, 27] that has been adapted to video [13, 15, 21, 37, 45, 46]. It pulls representations of positive views based on the input while pushing a large number of other representations. Recently, SCE [15] has shown that leveraging estimated inter-instance relations with contrastive learning improves performance.

Masked Modeling approaches emerged with transformers and showed better performance than contrastive learning to learn local features [2,26,42,64]. It masks a part of the input and reconstructs the signal either at pixel-level like MAE [20,26,52], at features-level [22,57], or by predicting visual tokens [50,56]. However, these methods assume lots of redundancies are present in the video which is true for short videos with few view variations but does not hold for complex videos such as soccer matches. Finally, Knowledge distillation (KD) [28] is another pretraining approach that distills information from a teacher or a collection of teacher models to a student and has been successfully applied for Supervised Learning [53] as well as Self-Supervised Learning [19,22,34,43]. Multiple approaches emerged to learn spatial and temporal features separately or decoupled to learn a global representation [30,44]. Notably, [63] proposes a multi-step pretraining method that decouples spatial and temporal information through two different networks recoupled via a self-distilled network. [62] performs spatial and temporal contrastive learning at multiple hierarchies in the model to separate spatial and temporal features.

In our work, we pretrain our hierarchical model with a contrastive SSL initialization of the spatial transformer. Then, our global model is pretrained with a KD loss from an extracted bank of features aligned with all the output tokens. This loss leverages temporal masking and soft contrastive learning to maintain local-temporal information enriched in a global context. Therefore, our goal is to learn multiple local-temporal representations, not one global.

Action Spotting. Action Spotting is a timestamp-level Temporal Action Detection (TAD) first introduced for the dataset SoccerNet [23]. This dataset has been extended to more videos and more actions in SoccerNet-v2 [14] and the tight-Average mean Average Precision (t-AmAP) has been introduced to evaluate precise detection within thresholds of 1 to 5 seconds. Several approaches have been proposed to tackle this task that can be divided into two categories. First, most approaches build a temporal architecture on top of a feature extractor [5,8,11,12,24,38,40,48,49,51,65], and second, few others train an end-to-end network [29,66]. The first kind of approach reduces the computational cost of experiments however makes them rely on a feature extractor for generalization. Notably, Baidu [65] proposed the extraction features of five 3D CNN models pretrained on Kinetics [32] and finetuned on SoccerNet-v2. The features are then plugged into a temporal action detector. Spivak [49] used these Baidu features coupled with features from a pretrained ResNet-152 to train an anchor-based approach that first classifies actions falling in a few-second temporal radius and shifts the predictions using a temporal regressor. For end-to-end approaches, E2E-spot [29] proposed to train a CNN spatial encoder on top of a simple recurrent model that performed competitively with the previous approach.

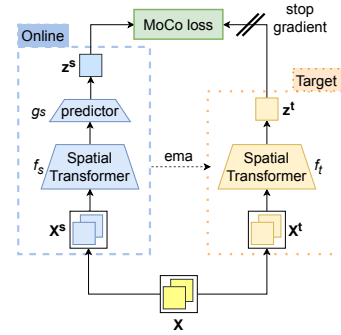


Figure 2. Pretraining of the spatial transformer.

As for other domains, the attention mechanism has been studied to improve performance [40,48,49,65,66].

In our work, we propose an end-to-end transformer-based action spotting approach that assigns actions that fall in frames in a small temporal radius.

3. Method

3.1. Overview

Our approach seeks to learn locally precise temporal features enriched with a larger context for action spotting. Therefore our model is composed of three different parts:

- A spatial transformer embeds the information of a small temporal window and outputs one token embedding.
- A temporal transformer that takes as input the token outputs of the spatial transformer on consecutive windows and outputs the same number of tokens. The temporal transformer enriches the representation of local windows with the knowledge of a larger context.
- A linear head is applied on each temporal output token to perform the classification of the action classes.

To train the model, we perform three different training steps described in the latter subsections: (1) pretraining of the spatial transformer on small windows in Sec. 3.2, (2) pretraining of the spatial and temporal transformers on large windows in Sec. 3.3, (3) fine-tuning of the model on the action spotting task in Sec. 3.4.

3.2. Spatial pretraining

To train the spatial transformer, we follow the Self-Supervised Contrastive method MoCo [27] illustrated in Fig. 2. We use a Siamese architecture containing an online and a target branch. For the online branch, the model contains a transformer f_s and a predictor g_s . The target branch contains a copy of the transformer updated by the exponential moving average, or ema, of the online transformer.

Each video $\mathbf{X} \in \mathbb{R}^{T_s \times H \times W \times C}$ of T_s frames, width W , height H , and C channels from the dataset is augmented by two different distributions of data augmentations A^1 and A^2

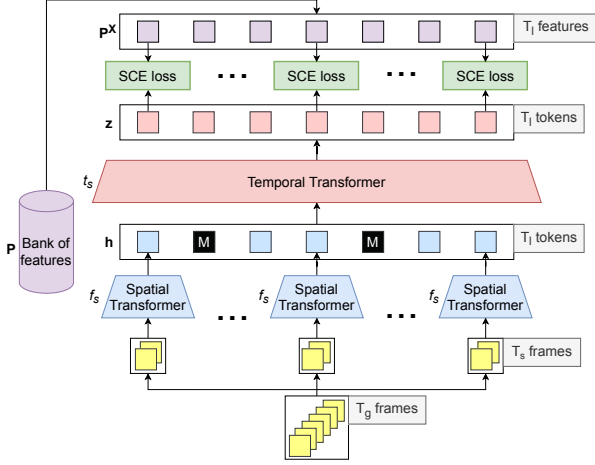


Figure 3. Pretraining of the spatial and temporal transformers via knowledge distillation of a bank of features with the SCE loss. Some spatial output tokens are masked.

to form positive views $\mathbf{X}^1 = a^1(\mathbf{X})$, $\mathbf{X}^2 = a^2(\mathbf{X})$ with $a^1 \sim A^1$ and $a^2 \sim A^2$. We pass both views in both transformers to compute the representations $\mathbf{z}^{1s} = g_s(f_s(\mathbf{X}^1))$, $\mathbf{z}^{2s} = g_s(f_s(\mathbf{X}^2))$, $\mathbf{z}^{1t} = f_t(\mathbf{X}^1)$ and $\mathbf{z}^{2t} = f_t(\mathbf{X}^2)$. A momentum memory buffer \mathbf{Q} of size $M \gg N$ is maintained on the target representations to provide negatives.

We apply the MoCo loss on each representation as follows:

$$\mathcal{L}_{MoCo} = -\frac{1}{2} (l_M(\mathbf{z}^{1s}, \mathbf{z}^{2t}) + l_M(\mathbf{z}^{2s}, \mathbf{z}^{1t})), \quad (1)$$

$$l_M(\mathbf{z}, \mathbf{k}) = \log \left(\frac{\exp(\mathbf{z} \cdot \mathbf{k} / \tau)}{\exp(\mathbf{z} \cdot \mathbf{k} / \tau) + \sum_{j=1}^n \exp(\mathbf{z} \cdot \mathbf{Q}_j / \tau)} \right). \quad (2)$$

3.3. Spatio-temporal pretraining

To pretrain the spatial transformer f_s and the temporal transformer t_s , we adapt the soft contrastive self-supervised loss SCE to perform knowledge distillation as illustrated in Fig. 3. by considering the following inputs:

- A large video $\mathbf{X} \in \mathbb{R}^{T_g, C, H, W}$ of T_g frames that is divisible by T_s with $T_l = T_g / T_s$ sampled from a dataset.
- A bank of spatio-temporal features $\mathbf{P} \in \mathbb{R}^{M_P, D_t}$ of size M_P and dimension D_t that can be aligned temporally with small window of T_s frames within any sampled large video. More specifically, each small window is associated with the temporally closest feature of its middle frame. Therefore for each sample \mathbf{X} a set of features $\mathbf{P}^{\mathbf{X}} \in \mathbb{R}^{T_l, D_t}$ is selected from \mathbf{P} . As a preprocessing stage, this bank of features is extracted from a pretrained model.

The video is augmented with $a^3 \sim T^3$ such as $\mathbf{X}^3 = a^3(\mathbf{X})$. Each global window is split into $T_l = T_g / T_s$ smaller windows and the input is reshaped to

(T_l, T_s, C, H, W) . It passes through the spatial transformer to output $\mathbf{h} = f_s(\mathbf{X}^3)$ with $\mathbf{h} \in \mathbb{R}^{T_l \times D}$ and D the output dimension of the spatial transformer. A masking ratio α_1 is applied to replace $\alpha_1 * T_l$ tokens with a learned mask token. The temporal transformer adds a positional embedding to each token and computes $\mathbf{z} = t_s(\mathbf{h})$ with $\mathbf{z} \in \mathbb{R}^{T_l, D_t}$ and D_t the dimension of the temporal output tokens.

The SCE loss is applied on each token of \mathbf{z} with the associated set of features $\mathbf{P}^{\mathbf{X}}$ as follows. First, a target relation distribution s^2 is computed on each of the features in $\mathbf{P}^{\mathbf{X}}$ with the complete bank. A one-hot label is mixed with this distribution with a coefficient λ to form the target distribution \mathbf{w}^2 . Then, each token in \mathbf{h} predicts this target distribution by computing its similarity distribution s^1 with the complete bank of features. Finally, the loss is applied:

$$s_{ik}^1 = \frac{\exp(\mathbf{z}_i^1 \cdot \mathbf{P}_k / \tau)}{\sum_{j=1}^{M_P} \exp(\mathbf{z}_i^1 \cdot \mathbf{P}_j / \tau)}, \quad (3)$$

$$s_{ik}^2 = \frac{\mathbb{1}_{\mathbf{P}_i^{\mathbf{X}} \neq \mathbf{P}_k} \cdot \exp(\mathbf{P}_i^{\mathbf{X}} \cdot \mathbf{P}_k / \tau_m)}{\sum_{j=1}^{M_P} \mathbb{1}_{\mathbf{P}_i^{\mathbf{X}} \neq \mathbf{P}_j} \cdot \exp(\mathbf{P}_i^{\mathbf{X}} \cdot \mathbf{P}_j / \tau_m)}, \quad (4)$$

$$w_{ik}^2 = \lambda \mathbb{1}_{\mathbf{P}_i^{\mathbf{X}} = \mathbf{P}_k} + (1 - \lambda) s_{ik}^2, \quad (5)$$

$$\mathcal{L}_{SCE} = -\frac{1}{T_l} \sum_{i=1}^{T_l} \mathbf{w}_i^2 \log(s_i^1). \quad (6)$$

The KD enforces that each temporal token contains the information of its corresponding smaller window while allowing contextual information from the larger window thanks to the temporal transformer. The SCE loss enables the spatio-temporal token representation to model the relations among spatio-temporal small windows that the bank of features contains. Depending on how the features are extracted, the KD can be considered as supervised or self-supervised as discussed in Sec. 4.2.

3.4. Fine-tuning

The input video \mathbf{X} considered in this section has the same shape as for spatio-temporal pretraining and passes through the spatial and temporal transformer that outputs $\mathbf{z} = t_s(f_s(a^4(\mathbf{X})))$ with $a^4 \sim A^4$, a data augmentation.

To train the model to the action spotting task with C classes, a classification head c_s is placed upon the temporal transformer and is applied on each of the temporal output tokens to predict $\hat{\mathbf{y}} = c_s(\mathbf{z})$ with $\hat{\mathbf{y}} \in \mathbb{R}^{T_l, C}$ as illustrated in Fig. 4. Each token is associated with the average timestamp to the corresponding small window they represent. The masking strategy from the spatiotemporal pretraining is maintained during training by randomly masking $\alpha_2 * T_l$ spatial output tokens.

For supervision, each ground truth action that falls into the T_g sampled window is associated with the input frame




Figure 4. Fine-tuning to the action spotting task. Some spatial output tokens are masked.

timestamps that fall into a temporal radius displacement ϵ with the action. For each T_s smaller window, if at least one of its frames is associated with an action, the label associated with its temporal output token is 1 and otherwise 0 to form the vector of label $\mathbf{y} \in \mathbb{R}^{T_l, C}$.

The action classes are considered independently, therefore we apply a Sigmoid activation function to the classifier. During training, the Binary Cross Entropy (BCE) loss is computed for each class at each timestamp as follows:

$$\mathcal{L}_{BCE} = -\frac{1}{T_l \times C} \sum_{t=1}^{T_l} \sum_{c=1}^C y_{tc} \log(\hat{y}_{tc}) \quad (7)$$

During inference, the predictions are assigned to the output timestamps of their temporal token. A sliding window with overlap is performed on the videos. For overlapped predictions, a strategy is applied to only keep one prediction per class per timestamp, such as keeping the maximum or the average of predictions.

4. Empirical study

In this section, we will first review the implementation details of each step, then perform an ablation study on the different parts of our pipeline, and finally compare ourselves with the state of the art.

4.1. Implementation details

All experiments are averaged on three seeds.

Dataset. We performed our study on the SoccerNet-v2 [14] action spotting dataset. It contains soccer matches divided into two halves of about 45 minutes. It has three annotated splits of 17 classes with 300 matches for training, 100 for validation, and 100 for testing. There is also a challenge

split that contains 50 videos for which annotations are not given. The metric used is the tight-Average mean Average Precision, t-AmAP for short, which evaluates predictions that fall on average between 1-5 seconds. We extracted the video frames at resolution 398×224 at 2 Frames Per Second (FPS). The dataset provides pre-computed Baidu [65] features at 1 FPS for all splits. We performed a PCA on these features to reduce the dimension to 512 for distillation.

Spatial and Temporal transformer architectures. For the spatial transformer, two different architectures are used: ViT [17] and Swin [35]. The temporal transformer is a stack of 4 attention layers from a ViT architecture. For ViT, the global architecture corresponds to the ViViT model 2 [1]. We keep this name and refer to the Swin-based architecture as ViSwin. More details can be found in the supplementary material. For optimization, we used the ADAMW optimizer with a weight decay of 0.05. The initial learning rate depends on the step as well as the backbone and is detailed in the supplementary material.

Spatial pretraining. To perform the pretraining of the spatial transformer via MoCo [27], we use SCE’s hyperparameters [15]. More specifically we use a projector for the online and target branches and a predictor for the online branch. For data sampling, all sub-videos of 1 second are used. Details can be found in the supplementary material.

Spatio-temporal pretraining. To perform the pretraining of the spatial and temporal transformers, we follow SCE [15]. We apply a projector on top of each output temporal token and we distill information from the reduced Baidu features. Details can be found in the supplementary material. For data sampling, we randomly extract 150 videos of 32 seconds, or 64 frames, per match at each epoch.

Fine-tuning. To perform fine-tuning, common data augmentations are applied as well as mixup [61]. 100 videos per match per epoch are sampled uniformly. The classifier is first initialized and then the whole backbone is fine-tuned. Details can be found in the supplementary material.

Inference. During inference, a sliding window with half overlap is applied on all videos. For multiple timestamp classifications, the maximum of predictions per action is kept. No data augmentation is applied. A hard Non-Maximum Suppression (NMS) of a 5-second window is applied. The 6 first and last seconds of each window are ignored to keep predictions with past or future context.

4.2. Ablation study

In this subsection, we will make various ablations to highlight the advantages of our 3-step approach, our masking strategy, and how the model and data sampling affect performance. The majority of the ablation study is performed on ViViT Tiny to reduce the cost of training. The models in this section are trained on the training split and evaluated on the validation split.

Model	Params (M)	GFLOPs	t-AmAP (%)
ViViT T	7.5	41.2	64.7
ViViT S	29.1	149.5	65.9
ViSwin T	55.9	145.6	66.1

Table 1. Influence of the model architecture on the t-AmAP.

Depth	Params (M)	GFLOPs	t-AmAP (%)
4	7.5	41.2	64.7
6	8.3	41.3	65.4
8	9.2	41.3	65.3

Table 2. Influence of temporal depth on ViViT-T for the t-AmAP.

Architectures. We test two architectures for the spatial part, ViT [17] and Swin [35]. Because the output embedding dimension of the spatial transformer is the input dimension of the temporal transformer, the number of parameters increases quadratically with the spatial dimension. As going deeper with Swin increases the token dimension, its output dimension token is large, leading to a larger number of parameters for ViSwin’s temporal encoder than ViViT’s. We compare the performance of ViViT and ViSwin in Tab. 1. ViSwin Tiny shows an improvement over ViViT Tiny with +1.4 percentage points (p.p). However, this improvement comes with a price of about 7.5 times more parameters. Going from Tiny to Small for ViViT improved by +1.2 p.p for about 4 times parameter. But, its GFLOPs are slightly higher than ViSwin Small suggesting the Swin spatial transformer is more efficient for larger networks.

We also test a deeper temporal transformer as the majority of computation comes from the spatial transformer by design [1]. Indeed for ViT, it sees for a global window 6,272 tokens whereas the temporal transformer only has 32. Therefore, besides increasing the number of parameters, the cost of making a deeper temporal transformer is computationally negligible in comparison with a deeper spatial transformer. The baseline is a depth of 4 blocks of attention and we increase it to 6 and 8. The results are reported in Tab. 2 and show that increasing the temporal depth to 6 increases the t-AmAP by 0.7 p.p and going deeper decreases performance. In contrast with ViViT applied to action classification [1] we increase performance with a deeper temporal transformer probably because action spotting requires modeling more complex temporal dependencies.

Size of context. The size of the temporal context influences how our model perceives actions. We study this influence in Tab. 3 by increasing 2 and 4 times the temporal context. To keep the computational cost the same between different sizes, we adapt the batch size adequately. Increasing it 2 times improved the results by +1.3 p.p and a larger con-

Window duration (s)	t-AmAP (%)
32	64.7
64	66.0
128	65.5

Table 3. Influence of temporal length on ViViT-T on the t-AmAP.

α_1	α_2	t-AmAP (%)
<i>None</i>		
0.00	0.00	64.7
<i>Only fine-tuning</i>		
0.00	0.50	55.5
<i>Only pretraining</i>		
0.25	0.00	64.9
0.50	0.00	65.1
0.75	0.00	64.8
<i>Both</i>		
0.25	0.25	65.2
0.50	0.50	65.0
0.75	0.75	63.5

Table 4. Influence of masking ratio during spatio-temporal pre-training (α_1) and fine-tuning (α_2) on ViViT-T on the t-AmAP.

text shows a slight decrease. This shows that a large temporal context leads to a better understanding of soccer actions.

Masking. Steps 2 and 3 of our training pipeline incorporate a temporal masking strategy. This masking has two goals: limit model overfitting and make the temporal transformer focus on contextual information instead of just aligning its output with its input. We show the advantage of this masking strategy in Tab. 4 by masking only during pretraining, only during fine-tuning, or both. First, masking during only fine-tuning drastically decreases performance by -9.2 p.p. Masking during pretraining increases results by up to 0.4 p.p for 50% tokens masked and masking during both steps increases up to 0.5 p.p for 25% tokens masked. Performance decreases with further masking. These results suggest it is necessary to initialize the mask token during pretraining and the percentage of tokens to mask is different during pretraining and fine-tuning for optimal performance.

Steps. Our training pipeline consists of three steps, each adding complexity. Here, we validate the usefulness of each step. We evaluate the quality of our learned representation in step 1 by comparing its performance with a supervised pretrained ViT Tiny model on ImageNet 21k that contains 14 million labeled diverse images. To ensure a fair comparison with spatiotemporal pretrained backbones, when step 2 is not performed, we perform a longer fine-tuning.

We report results in Tab. 5. Each step consistently improves performance. Indeed, training from scratch reaches

Step 1	Step 2	Step 3 epochs	t-AmAP (%)
x	x	100F	48.1
SN	x	100F	54.7
IN	x	100F	57.7
SN	✓	50F	62.2
x	✓	30C + 20F	60.0
SN	✓	30C + 20F	64.7
IN	✓	30C + 20F	65.0

Table 5. Influence of the pretraining steps and the number of fine-tuning epochs on ViViT-T on the t-AmAP. *SN* stands for SoccerNet-v2 MoCo self-supervised pretraining, and *IN* for ImageNet21k supervised pretraining. *C* stands for training the classifier and *F* for fine-tuning the whole model.

Depth	Sequence	Masking	t-AmAP (%)
x	x	x	64.7
✓	x	x	65.4
x	✓	x	66.0
x	x	✓	65.2
✓	✓	x	66.1
x	✓	✓	66.2
✓	✓	✓	66.6

Table 6. Influence of best parameters for temporal depth and length, and the masking strategy on ViViT-T on the t-AmAP.

48.1% t-AmAP. Adding step 1 increases up to 54.7% for SSL pretraining and 57.7% for ImageNet pretraining. This suggests that the spatial transformer takes advantage of initialization from a large diversity of data and that our SSL pretraining can be improved. Step 2 further improves results, even with random spatial initialization which reached 60.0% t-AmAP. With SoccerNet weights, it increases to 64.7%, and with ImageNet weights, it reaches 65.0%. The gap between SoccerNet and Imagenet pretraining in step 1 is almost closed in step 2. As our SSL approach was trained on videos, while ImageNet weights were obtained on images, we argue that our second pretraining stage benefits from having initial spatiotemporal features. Initializing the temporal transformer accelerates convergence and improves results compared to training from scratch or step 1. This reduces the cost of pretraining, allowing future work to perform fast experiments in the fine-tuning phase.

All together. In Tab. 6, we test adding together the different best hyperparameters for a deeper temporal transformer, larger temporal context, and temporal masking. Previously, we showed that the larger improvement came from increasing the temporal context so we add other components to it. Increasing temporal depth adds 0.1% whilst using the masking strategy adds 0.2% which makes them marginal in comparison with previous improvements. How-

Features	Pretraining		t-AmAP (%)
	Dataset	Fine-tuned	
SCE [15]	K400		63.6
SCE [15]	K400	✓	65.7
Baidu [65]	K400	✓	66.6

Table 7. Influence of features to perform KD on ViViT-T on t-AmAP. Supervised features provide best results and self-supervised features of SCE achieve good performance.

NMS	Ignore (s)	Window (s)	t-AmAP (%)
<i>Default inference</i>			
hard	6	5	66.6
<i>Best inference for soft and hard NMS</i>			
hard	12	3	67.1
soft	12	10	68.0

Table 8. Best inference parameters on ViViT-T on the t-AmAP.

ever, combining the three improves 0.6% to attain our best result of 66.6%. This confirms that a large temporal context is the most determining component to improve performance and that the masking strategy scales with the number of parameters and ensures new information is learned.

Bank of features. We change the bank of features used from Baidu [65], which necessitates fine-tuning of 5 models pretraining on Kinetics400 [33] to obtain, with two options: extracted features from SCE [15] pretrained R3D50 on Kinetics400, and its fine-tuned version to the action spotting task. The clips used for the R3D50 last 4 seconds and fine-tuning is performed on the middle frame. We report results in Tab. 7. Baidu features achieve best performance thanks to its 5 aggregated models, but SCE fine-tuned, which is 1 model, is enough to achieve competitive performance. Also, using the self-supervised model loses $-2.1p.p$ but opens an interesting perspective toward pretraining a self-supervised model on a closer domain for feature extraction.

Inference pipeline. The inference has also a huge impact on performance. There are 3 parameters that we take into account: whether to use *hard* or *strong* NMS, the number of seconds to ignore at the beginning and end of each window prediction, and the size of the NMS window. In Tab. 8, we provide the results of the best parameters that we found for hard and soft NMS which are detailed in the supplementary material and we empirically observe a better performance for soft NMS.

4.3. Comparison with the State of the Art

Implementation details. For comparison with the state of the art on the test and challenge splits, we use the best settings found in Sec. 4.2. The results labeled *ens.* means we use the average predictions of 3 seeds. For the test, resp.

Method	Input	t-AmAP (%)
NetVLAD++ [24]	F	11.5
AImageLab RMSNet [51]	F	28.8
Baidu [65]	F	47.1
Faster-TAD [8]	F	54.1
SpotFormer [5]	F	60.9
E2E-Spot [29]	I	61.8
Spivak [49]	F	65.1
COMEDIAN (ViViT-T)	I	70.7
COMEDIAN (ViSwin-T)	I	71.6
COMEDIAN (ViViT-T - ens.)	I	72.0
COMEDIAN (ViSwin-T - ens.)	I	73.1

Table 9. Comparison with the state of the art on the test split of SoccerNet-v2. *F* stands for methods using a feature extractor, *I* for methods end-to-end with image inputs.

Method	Input	t-AmAP (%)
<i>Challenge 2022 leaderboard</i>		
Baidu [65]	F	49.56
Transformer-AS [66]	I	52.04
Faster-TAD [8]	F	64.88
E2E-Spot [29]	I	66.73
Spivak [49]	F	67.81
<i>Challenge 2023 submission</i>		
Spivak* [49]	F	68.33
COMEDIAN (ViViT-T - ens.)	I	68.38

Table 10. Comparison with the state of the art on the challenge split of SoccerNet-v2. *F* stands for methods using a feature extractor, *I* for methods end-to-end with image inputs.

challenge, split we use the training and validation, resp. all annotated, splits for steps 2 and 3.

Comparison on test split. We report our results in Tab. 9. We compare ourselves with methods that use a sequence of images as input or a feature extractor. We observe that COMEDIAN with ViViT Tiny provides a significant improvement over the state of the art by +5.6 p.p on t-AmAP. ViSwin Tiny increases performance by +0.9 p.p but at a high cost in terms of computational usage. Finally using an ensemble of our 3 seed, we achieve 72.0% t-AmAP for ViViT Tiny and 73.1 % for ViSwin Tiny. These results empirically prove that our approach even with a small network produces state-of-the-art results by using our initializing pipeline. It is worth noting that we perform a simple fine-tuning stage. Previous approaches only focused on the fine-tuning part and because the two are not mutually exclusive, it opens interesting perspectives for future work to build better fine-tuning upon our approach.

Comparison on challenge split. We report our results in Tab. 10. We compare with some participants from the Chal-

lenge 2022. We also report the result of the baseline of the Challenge 2023 [49] in which we participated, which is an improved version of the winner of 2022. Our proposed COMEDIAN achieves 68.38% on global t-AmAP for ViViT Tiny. Contrary to the test split we do not have a gap with previous state-of-the-art methods and achieve +0.56 p.p in comparison with 2022 best result and +0.05 p.p in comparison with 2023’s. Because of the opacity of the challenge split’s labels, it is difficult to investigate the discrepancy between the test and the challenge. Compared with the best end-to-end methods E2E-Spot [29], our approach achieves a more significant improvement of +1.64 p.p on t-AmAP.

5. Discussion

We discuss potential improvements and future directions for enhancing COMEDIAN. For spatial pretraining, we identify three directions: increasing the amount of training data, the number of frames, and designing a SSL task for semantically sparse images. For spatio-temporal pretraining, we suggest exploring temporal models better suited to action spotting and improving KD by using SSL-learned features from a large close domain. In the fine-tuning stage, we suggest improving the labelization and designing losses that make use of relationships between different action classes and are capable of dealing with challenging actions such as card-related ones. Finally, our study focused on soccer action spotting and studying the generalization to other datasets and general TAD is an interesting perspective.

6. Conclusion

In this chapter, we introduce COMEDIAN a novel approach for Action Spotting that leverages self-supervised learning and knowledge distillation to initialize a spatio-temporal transformer. It achieves state-of-the-art results on the SoccerNet-v2 action spotting task, demonstrating the effectiveness of the proposed pipeline. By utilizing unlabeled video data for pretraining, we address the subjective and resource-intensive manual labeling processes for action spotting. The pretraining cost is leveraged by a faster and better convergence during fine-tuning. While our approach shows promising results, there are areas for improvement in the pretraining and fine-tuning steps and we hope that our approach will open the path to new methods to increase performance on action spotting and temporal action detection with spatio-temporal transformer models.

Acknowledgement

This publication was made possible by the use of the Factory-AI supercomputer, financially supported by the Ile-de-France Regional Council, and the HPC resources of IDRIS under the allocation 2023-AD011014382 made by GENCI.

References

- [1] Anurag Arnab, Mostafa Dehghani, Georg Heigold, Chen Sun, Mario Lucic, and Cordelia Schmid. Vivit: A video vision transformer. In *International Conference on Computer Vision*, pages 6816–6826, 2021. 1, 2, 5, 6
- [2] Hangbo Bao, Li Dong, Songhao Piao, and Furu Wei. Beit: BERT pre-training of image transformers. In *International Conference on Learning Representations*, 2022. 3
- [3] Gedas Bertasius, Heng Wang, and Lorenzo Torresani. Is space-time attention all you need for video understanding? In Marina Meila and Tong Zhang, editors, *38th International Conference on Machine Learning*, pages 813–824, 2021. 1, 2
- [4] Adrian Bulat, Juan-Manuel Pérez-Rúa, Swathikiran Sudhakaran, Brais Martínez, and Georgios Tzimiropoulos. Space-time mixing attention for video transformer. In *Advances in Neural Information Processing Systems 34*, pages 19594–19607, 2021. 2
- [5] Mengqi Cao, Min Yang, Guozhen Zhang, Xiaotian Li, Yilu Wu, Gangshan Wu, and Limin Wang. Spotformer: A transformer-based framework for precise soccer action spotting. In *24th International Workshop on Multimedia Signal Processing*, pages 1–6. IEEE, 2022. 3, 8
- [6] Mathilde Caron, Ishan Misra, Julien Mairal, Priya Goyal, Piotr Bojanowski, and Armand Joulin. Unsupervised learning of visual features by contrasting cluster assignments. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems*, 2020. 2
- [7] 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 International Conference on Computer Vision*, pages 6706–6716, 2021. 1, 2
- [8] Shimin Chen, Chen Chen, Wei Li, Xunqiang Tao, and Yandong Guo. Faster-tad: Towards temporal action detection with proposal generation and classification in a unified network. *arXiv*, abs/2204.02674, 2022. 3, 8
- [9] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey E. Hinton. A simple framework for contrastive learning of visual representations. In *Proceedings of the 37th International Conference on Machine Learning*, pages 1597–1607, 2020. 2
- [10] Xinlei Chen, Saining Xie, and Kaiming He. An empirical study of training self-supervised vision transformers. In *International Conference on Computer Vision*, pages 9620–9629, 2021. 2
- [11] Anthony Cioppa, Adrien Delière, Silvio Giancola, Bernard Ghanem, Marc Van Droogenbroeck, Rikke Gade, and Thomas B. Moeslund. A context-aware loss function for action spotting in soccer videos. In *Conference on Computer Vision and Pattern Recognition*, pages 13123–13133, 2020. 3
- [12] Abdulrahman Darwish and Tallal El Shabrawy. STE: spatiotemporal encoder for action spotting in soccer videos. In Rainer Lienhart, Thomas B. Moeslund, and Hideo Saito, editors, *5th International ACM Workshop on Multimedia Content Analysis in Sports*, pages 87–92. ACM, 2022. 3
- [13] Ishan R. Dave, Rohit Gupta, Mamshad Nayeem Rizve, and Mubarak Shah. TCLR: temporal contrastive learning for video representation. *Computer Vision and Image Understanding*, page 103406, 2022. 2
- [14] Adrien Delière, Anthony Cioppa, Silvio Giancola, Meisam Jamshidi Seikavandi, Jacob V. Dueholm, Kamal Nasrollahi, Bernard Ghanem, Thomas B. Moeslund, and Marc Van Droogenbroeck. Soccernet-v2: A dataset and benchmarks for holistic understanding of broadcast soccer videos. In *Conference on Computer Vision and Pattern Recognition Workshops*, pages 4508–4519, 2021. 2, 3, 5
- [15] Julien Denize, Jaonary Rabarisoa, Astrid Orcesi, and Romain Hérault. Similarity contrastive estimation for image and video soft contrastive self-supervised learning. *Machine Vision and Applications*, 34(6), 2023. 1, 2, 5, 7
- [16] Julien Denize, Jaonary Rabarisoa, Astrid Orcesi, Romain Hérault, and Stéphane Canu. Similarity contrastive estimation for self-supervised soft contrastive learning. In *Winter Conference on Applications of Computer Vision*, pages 2705–2715, 2023. 2
- [17] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations*, 2021. 1, 2, 5, 6
- [18] Debidatta Dwibedi, Yusuf Aytar, Jonathan Tompson, Pierre Sermanet, and Andrew Zisserman. With a little help from my friends: Nearest-neighbor contrastive learning of visual representations. In *2021 International Conference on Computer Vision*, pages 9568–9577, 2021. 2
- [19] Zhiyuan Fang, Jianfeng Wang, Lijuan Wang, Lei Zhang, Yezhou Yang, and Zicheng Liu. SEED: self-supervised distillation for visual representation. In *International Conference on Learning Representations*, 2021. 3
- [20] Christoph Feichtenhofer, Haoqi Fan, Yanghao Li, and Kaiming He. Masked autoencoders as spatiotemporal learners. In *Advances in Neural Information Processing Systems*, 2022. 3
- [21] Christoph Feichtenhofer, Haoqi Fan, Bo Xiong, Ross B. Girshick, and Kaiming He. A large-scale study on unsupervised spatiotemporal representation learning. In *Conference on Computer Vision and Pattern Recognition*, pages 3299–3309, 2021. 1, 2
- [22] Yuting Gao, Jia-Xin Zhuang, Shaohui Lin, Hao Cheng, Xing Sun, Ke Li, and Chunhua Shen. Disco: Remedying self-supervised learning on lightweight models with distilled contrastive learning. In *17th European Conference on Computer Vision*, pages 237–253, 2022. 2, 3
- [23] Silvio Giancola, Mohieddine Amine, Tarek Dghaily, and Bernard Ghanem. Soccernet: A scalable dataset for action spotting in soccer videos. In *Conference on Computer Vision and Pattern Recognition*, pages 1711–1721, 2018. 1, 3
- [24] Silvio Giancola and Bernard Ghanem. Temporally-aware feature pooling for action spotting in soccer broadcasts. In *Conference on Computer Vision and Pattern Recognition Workshops*, pages 4490–4499, 2021. 3, 8

- [25] Jean-Bastien Grill, Florian Strub, Florent Althé, Corentin Tallec, Pierre H. Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Ávila Pires, Zhaohan Guo, Mohammad Gheshlaghi Azar, Bilal Piot, Koray Kavukcuoglu, Rémi Munos, and Michal Valko. Bootstrap your own latent - A new approach to self-supervised learning. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems*, 2020. 2
- [26] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross B. Girshick. Masked autoencoders are scalable vision learners. In *Conference on Computer Vision and Pattern Recognition*, pages 15979–15988, 2022. 3
- [27] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross B. Girshick. Momentum contrast for unsupervised visual representation learning. In *Conference on Computer Vision and Pattern Recognition*, pages 9726–9735, 2020. 2, 3, 5
- [28] Geoffrey E. Hinton, Oriol Vinyals, and Jeffrey Dean. Distilling the knowledge in a neural network. *arXiv*, abs/1503.02531, 2015. 1, 3
- [29] James Hong, Haotian Zhang, Michaël Gharbi, Matthew Fisher, and Kayvon Fatahalian. Spotting temporally precise, fine-grained events in video. In *17th European Conference on Computer Vision*, pages 33–51, 2022. 3, 8
- [30] Lianghua Huang, Yu Liu, Bin Wang, Pan Pan, Yinghui Xu, and Rong Jin. Self-supervised video representation learning by context and motion decoupling. In *Conference on Computer Vision and Pattern Recognition*, pages 13886–13895, 2021. 3
- [31] Sukjun Hwang, Miran Heo, Seoung Wug Oh, and Seon Joo Kim. Video instance segmentation using inter-frame communication transformers. In *Advances in Neural Information Processing Systems*, pages 13352–13363, 2021. 2
- [32] Ioannis Kakogeorgiou, Spyros Gidaris, Bill Psomas, Yan-nis Avrithis, Andrei Bursuc, Konstantinos Karantzalos, and Nikos Komodakis. What to hide from your students: Attention-guided masked image modeling. In Shai Avidan, Gabriel J. Brostow, Moustapha Cissé, Giovanni Maria Farinella, and Tal Hassner, editors, *17th European Conference on Computer Vision*, pages 300–318, 2022. 3
- [33] Will Kay, João Carreira, Karen Simonyan, Brian Zhang, Chloe Hillier, Sudheendra Vijayanarasimhan, Fabio Viola, Tim Green, Trevor Back, Paul Natsev, Mustafa Suleyman, and Andrew Zisserman. The kinetics human action video dataset. *arXiv*, abs/1705.06950, 2017. 7
- [34] Soroush Abbasi Koohpayegani, Ajinkya Tejankar, and Hamed Pirsiavash. Compress: Self-supervised learning by compressing representations. In *Advances in Neural Information Processing Systems*, 2020. 3
- [35] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *International Conference on Computer Vision*, pages 9992–10002, 2021. 1, 2, 5, 6
- [36] Ze Liu, Jia Ning, Yue Cao, Yixuan Wei, Zheng Zhang, Stephen Lin, and Han Hu. Video swin transformer. In *Conference on Computer Vision and Pattern Recognition*, pages 3192–3201, 2022. 1, 2
- [37] Guillaume Lorre, Jaonary Rabarisoa, Astrid Orcesi, Samia Ainouz, and Stéphane Canu. Temporal contrastive pretraining for video action recognition. In *Winter Conference on Applications of Computer Vision*, pages 651–659, 2020. 2
- [38] Behzad Mahaseni, Erma Rahayu Mohd Faizal Abdullah, and Ram Gopal Raj. Spotting football events using two-stream convolutional neural network and dilated recurrent neural network. *IEEE Access*, pages 61929–61942, 2021. 3
- [39] Tim Meinhardt, Alexander Kirillov, Laura Leal-Taixé, and Christoph Feichtenhofer. Trackformer: Multi-object tracking with transformers. In *Conference on Computer Vision and Pattern Recognition*, pages 8834–8844, 2022. 1, 2
- [40] Hiroaki Minoura, Tsubasa Hirakawa, Takayoshi Yamashita, Hironobu Fujiyoshi, Mitsuru Nakazawa, Yeongnam Chae, and Björn Stenger. Action spotting and temporal attention analysis in soccer videos. In *17th International Conference on Machine Vision and Applications I*, pages 1–6, 2021. 3
- [41] Daniel Neimark, Omri Bar, Maya Zohar, and Dotan Asselmann. Video transformer network. In *International Conference on Computer Vision Workshops*, pages 3156–3165, 2021. 1, 2
- [42] Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, Mahmoud Assran, Nicolas Ballas, Wojciech Galuba, Russell Howes, Po-Yao Huang, Shang-Wen Li, Ishan Misra, Michael G. Rabbat, Vasu Sharma, Gabriel Synnaeve, Hu Xu, Hervé Jégou, Julien Mairal, Patrick Labatut, Armand Joulin, and Piotr Bojanowski. Dinov2: Learning robust visual features without supervision. *arXiv*, abs/2304.07193, 2023. 3
- [43] Wonpyo Park, Dongju Kim, Yan Lu, and Minsu Cho. Relational knowledge distillation. In *Conference on Computer Vision and Pattern Recognition*, pages 3967–3976, 2019. 3
- [44] Rui Qian, Shuangrui Ding, Xian Liu, and Dahua Lin. Static and dynamic concepts for self-supervised video representation learning. In *European Conference on Computer Vision*, pages 145–164. Springer, 2022. 3
- [45] Rui Qian, Tianjian Meng, Boqing Gong, Ming-Hsuan Yang, Huiheng Wang, Serge J. Belongie, and Yin Cui. Spatiotemporal contrastive video representation learning. In *Conference on Computer Vision and Pattern Recognition*, pages 6964–6974, 2021. 2
- [46] Adrià Recasens, Pauline Luc, Jean-Baptiste Alayrac, Luyu Wang, Florian Strub, Corentin Tallec, Mateusz Malinowski, Viorica Patraucean, Florent Althé, Michal Valko, Jean-Bastien Grill, Aäron van den Oord, and Andrew Zisserman. Broaden your views for self-supervised video learning. In *International Conference on Computer Vision*, pages 1235–1245, 2021. 2
- [47] Javier Selva, Anders S. Johansen, Sergio Escalera, Kamal Nasrollahi, Thomas B. Moeslund, and Albert Clapés. Video transformers: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pages 1–20, 2023. 1
- [48] Yuzhi Shi, Hiroaki Minoura, Takayoshi Yamashita, Tsubasa Hirakawa, Hironobu Fujiyoshi, Mitsuru Nakazawa, Yeongnam Chae, and Björn Stenger. Action spotting in soccer videos using multiple scene encoders. In *26th International*

- Conference on Pattern Recognition*, pages 3183–3189, 2022. 3
- [49] João V. B. Soares, Avijit Shah, and Topojoy Biswas. Temporally precise action spotting in soccer videos using dense detection anchors. In *International Conference on Image Processing*, pages 2796–2800, 2022. 3, 8
- [50] Chen Sun, Austin Myers, Carl Vondrick, Kevin Murphy, and Cordelia Schmid. Videobert: A joint model for video and language representation learning. In *International Conference on Computer Vision*, pages 7463–7472, 2019. 3
- [51] Matteo Tomei, Lorenzo Baraldi, Simone Calderara, Simone Bronzin, and Rita Cucchiara. Rms-net: Regression and masking for soccer event spotting. In *25th International Conference on Pattern Recognition*, pages 7699–7706, 2020. 3, 8
- [52] Zhan Tong, Yibing Song, Jue Wang, and Limin Wang. VideoMAE: Masked autoencoders are data-efficient learners for self-supervised video pre-training. In *Advances in Neural Information Processing Systems*, 2022. 1, 3
- [53] Hugo Touvron, Matthieu Cord, Matthijs Douze, Francisco Massa, Alexandre Sablayrolles, and Hervé Jégou. Training data-efficient image transformers & distillation through attention. In *38th International Conference on Machine Learning*, pages 10347–10357, 2021. 2, 3
- [54] Aäron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv*, abs/1807.03748, 2018. 2
- [55] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems*, pages 5998–6008, 2017. 2
- [56] Rui Wang, Dongdong Chen, Zuxuan Wu, Yinpeng Chen, Xiyang Dai, Mengchen Liu, Yu-Gang Jiang, Luowei Zhou, and Lu Yuan. BEVT: BERT pretraining of video transformers. In *Conference on Computer Vision and Pattern Recognition*, pages 14713–14723, 2022. 3
- [57] Chen Wei, Haoqi Fan, Saining Xie, Chao-Yuan Wu, Alan L. Yuille, and Christoph Feichtenhofer. Masked feature prediction for self-supervised visual pre-training. In *Conference on Computer Vision and Pattern Recognition*, pages 14648–14658, 2022. 1, 3
- [58] Chao-Yuan Wu, Yanghao Li, Karttikeya Mangalam, Haoqi Fan, Bo Xiong, Jitendra Malik, and Christoph Feichtenhofer. Memvit: Memory-augmented multiscale vision transformer for efficient long-term video recognition. In *Conference on Computer Vision and Pattern Recognition*, pages 13577–13587, 2022. 1, 2
- [59] Mingze Xu, Yuanjun Xiong, Hao Chen, Xinyu Li, Wei Xia, Zhuowen Tu, and Stefano Soatto. Long short-term transformer for online action detection. In *Advances in Neural Information Processing Systems*, pages 1086–1099, 2021. 1, 2
- [60] Jiewen Yang, Xingbo Dong, Liujuan Liu, Chao Zhang, Jiajun Shen, and Dahai Yu. Recurring the transformer for video action recognition. In *Conference on Computer Vision and Pattern Recognition*, pages 14043–14053, 2022. 1, 2
- [61] Hongyi Zhang, Moustapha Cissé, Yann N. Dauphin, and David Lopez-Paz. mixup: Beyond empirical risk minimization. In *6th International Conference on Learning Representations*, 2018. 5
- [62] Zehua Zhang and David Crandall. Hierarchically decoupled spatial-temporal contrast for self-supervised video representation learning. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 3235–3245, 2022. 3
- [63] Benjia Zhou, Pichao Wang, Jun Wan, Yanyan Liang, Fan Wang, Du Zhang, Zhen Lei, Hao Li, and Rong Jin. Decoupling and recoupling spatiotemporal representation for rgb-d-based motion recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20154–20163, 2022. 3
- [64] Jinghao Zhou, Chen Wei, Huiyu Wang, Wei Shen, Cihang Xie, Alan L. Yuille, and Tao Kong. ibot: Image BERT pre-training with online tokenizer. *arXiv*, abs/2111.07832, 2021. 3
- [65] Xin Zhou, Le Kang, Zhiyu Cheng, Bo He, and Jingyu Xin. Feature combination meets attention: Baidu soccer embeddings and transformer based temporal detection. *arXiv*, abs/2106.14447, 2021. 3, 5, 7, 8
- [66] He Zhu, Junwei Liang, Chengzhi Lin, Jun Zhang, and Jianming Hu. A transformer-based system for action spotting in soccer videos. In *5th International ACM Workshop on Multimedia Content Analysis in Sports*, pages 103–109, 2022. 3, 8