



UvA-DARE (Digital Academic Repository)

Object priors for classifying and localizing unseen actions

Mettes, P.; Thong, W.; Snoek, C.G.M.

DOI

[10.1007/s11263-021-01454-y](https://doi.org/10.1007/s11263-021-01454-y)

Publication date

2021

Document Version

Final published version

Published in

International Journal of Computer Vision

License

CC BY

[Link to publication](#)

Citation for published version (APA):

Mettes, P., Thong, W., & Snoek, C. G. M. (2021). Object priors for classifying and localizing unseen actions. *International Journal of Computer Vision*, 129(6), 1954–1971. <https://doi.org/10.1007/s11263-021-01454-y>

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <https://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, P.O. Box 19185, 1000 GD Amsterdam, The Netherlands. You will be contacted as soon as possible.



Object Priors for Classifying and Localizing Unseen Actions

Pascal Mettes¹ · William Thong¹ · Cees G. M. Snoek¹

Received: 19 May 2020 / Accepted: 5 March 2021 / Published online: 19 April 2021
© The Author(s) 2021

Abstract

This work strives for the classification and localization of human actions in videos, without the need for any labeled video training examples. Where existing work relies on transferring global attribute or object information from seen to unseen action videos, we seek to classify and spatio-temporally localize unseen actions in videos from image-based object information only. We propose three spatial object priors, which encode local person and object detectors along with their spatial relations. On top we introduce three semantic object priors, which extend semantic matching through word embeddings with three simple functions that tackle semantic ambiguity, object discrimination, and object naming. A video embedding combines the spatial and semantic object priors. It enables us to introduce a new video retrieval task that retrieves action tubes in video collections based on user-specified objects, spatial relations, and object size. Experimental evaluation on five action datasets shows the importance of spatial and semantic object priors for unseen actions. We find that persons and objects have preferred spatial relations that benefit unseen action localization, while using multiple languages and simple object filtering directly improves semantic matching, leading to state-of-the-art results for both unseen action classification and localization.

1 Introduction

The goal of this paper is to classify and localize human actions in video, such as *shooting a bow*, *doing a pull-up*, and *cycling*. Human action recognition has a long tradition in computer vision, with initial success stemming from spatio-temporal interest points (Chakraborty et al. 2012; Laptev 2005), dense trajectories (Wang et al. 2013; Jain et al. 2013), and cuboids (Kläser et al. 2010; Liu et al. 2008). Progress has recently been accelerated by deep learning, with the introduction of video networks exploiting two-streams (Feichtenhofer et al. 2016; Simonyan and Zisserman 2014) and 3D convolutions (Carreira and Zisserman 2017; Tran et al. 2019; Zhao et al. 2018; Feichtenhofer et al. 2019). Building on such networks, current action localizers have shown the ability to detect actions precisely in both space and time, e.g., (Gkioxari and Malik 2015; Hou et al. 2017;

Kalogeiton et al. (2017a); Zhao and Snoek 2019). Common amongst action classification and localization approaches is the need for a substantial amount of annotated training videos. Obtaining training videos with spatio-temporal annotations (Chéron et al. 2018; Mettes and Snoek 2019) is expensive and error-prone, limiting the ability to generalize to any action. We aim for action classification and localization without the need for *any* video examples during training.

In action recognition, many have explored the role of semantic action structures, from uncovering the grammar of an action (Kuehne et al. 2014) to enabling question answering in videos (Zhu et al. 2017). Language also plays a central role in zero-shot action recognition. Pioneering approaches transfer knowledge from attribute adjectives (Liu et al. 2011; Gan et al. (2016b); Zhang et al. 2015), object nouns (Jain et al. (2015a)), or combinations thereof (Wu et al. 2014). The supervised action recognition literature has already revealed the strong link between actions and objects for recognition (Gupta and Davis 2007; Jain et al. (2015b); Wu et al. 2007). Especially when object classification scores are obtained from large-scale image datasets (Deng et al. 2009; Lin et al. 2014) and matched with any action through word embeddings (Grave et al. 2018). We follow this object-based perspective for unseen actions. We add a generalization to spatio-temporal localization, by including local object detection scores and prior knowledge about prepositions, and we

Communicated by Vittorio Ferrari.

✉ Pascal Mettes
P.S.M.Mettes@uva.nl
William Thong
W.E.Thong@uva.nl
Cees G. M. Snoek
cgmsnoek@uva.nl

¹ Universiteit van Amsterdam, Amsterdam, The Netherlands

examine the linguistic relations between actions and objects to improve their semantic matching.

Our first contribution are three spatial object priors that encode local object and actor detections, as well as their spatial relations. We are inspired by the supervised action classification literature, where the spatial link with objects is well established, *e.g.*, (Gupta and Davis 2007; Kalogeiton et al. (2017b); Moore et al. 1999; Wu et al. 2007; Yao et al. 2011). To incorporate information about spatial prepositions without action video examples, we start from existing object detection image datasets and models. Box annotations in object datasets allow us to assess how people and objects are commonly related spatially. From discovered spatial relations, we propose a score function that combines person detections, object detections, and their spatial match for unseen action classification and localization. The spatial priors were previously introduced in the conference version (Mettes and Snoek 2017) preceding this paper.

Our second contribution, not addressed in (Mettes and Snoek 2017), are three semantic object priors. Common in unseen action recognition using objects is to estimate relations using word embeddings (Chang et al. 2016; Jain et al. (2015a); Li et al. 2019; Wu et al. 2016). They provide dense representations on which similarity functions are performed to estimate semantic relations (Mikolov et al. 2013). Similarities from word embeddings have several linguistic limitations relevant for unseen actions. Our semantic priors address three limitations with simple functions on top of word embedding similarities. First, we leverage word embeddings across languages to reduce semantic ambiguity in the action-object matching. Second, we show how to filter out non-discriminative objects directly from similarities between all objects and actions. Third, we show how to focus on basic-level names in object datasets to improve relevant matching. We combine the spatial and semantic object priors into a video embedding.

Experiments on five action datasets demonstrates the effectiveness of our six object priors. We find that the use of prepositions in our spatial-aware embedding enables effective unseen action localization using only a few localized objects. Our semantic object priors improve both unseen action classification and localization, with multi-lingual word embeddings, object discrimination functions, and a bias towards basic-level objects for selection. We also introduce a new task, action tube retrieval, where users can search for action tubes by specifying desired objects, sizes, and prepositions. Our object prior embedding obtains state-of-the-art zero-shot results for both unseen action classification and localization, highlighting its effectiveness and more generally, emphasizing the strong link between actions and objects.

The rest of the paper is organized as follows. Section 2 discusses related work. Sections 3 and 4 detail our spatial and

semantic object priors. Sections 5 and 6 discuss the experimental setup and results. The paper is concluded in Sect. 7.

2 Related Work

2.1 Unseen Action Classification

For unseen action classification, a common approach is to generalize from seen to unseen actions by mapping videos to a shared attribute space (Gan et al. (2016b); Liu et al. 2011; Zhang et al. 2015), akin to attribute-based approaches in images (Lampert et al 2013). Attribute classifiers are trained on seen actions and applied to test videos. The obtained attribute classifications are in turn compared to a priori defined attribute annotations. With the use of attributes, actions not seen during training can still be recognized. The attribute-based approach has been extended by using knowledge about test video distributions in transductive settings (Fu et al. 2015; Xu et al. 2017) and by incorporating domain adaptation (Kodirov et al. 2015; Xu et al. 2016). While enabling zero-shot recognition, attributes require prior expert knowledge for every action, which does not generalize to arbitrary queries. Hence we refrain from employing attributes.

Several works have investigated skipping the intermediate mapping to attributes by directly mapping unseen actions to seen actions. Li et al. (2016) and Tian et al. (2018) map features from videos to a semantic space shared by seen and unseen actions, while Gan et al. ((2016c)) train a classifier for unseen actions by performing several levels of relatedness to seen actions. Other works propose to synthesize features for unseen actions (Mishra et al. 2018, 2020), learn a universal representation of actions (Zhu et al. 2018), or differentiate seen from unseen actions through out-of-distribution detection (Mandal et al. 2019). All these works eliminate the need for attributes for unseen action classification. We also do not require attributes for our action classification, yet with the same model, we also enable action localization.

Several works have considered object classification scores for their zero-shot action, or event, classification by performing a semantic matching through word vectors (An et al. 2019; Bishay et al. 2019; Chang et al. 2016; Inoue and Shinoda 2016; Li et al. 2019; Jain et al. (2015a); Wu et al. 2016) or auxiliary textual descriptions (Gan et al. (2016a); Habibi et al. 2017). Objects provide an effective common space for unseen actions, as object scores are easily obtained by pre-training on existing large-scale datasets, such as ImageNet (Deng et al. 2009). Objects furthermore allow for a generalization to arbitrary unseen actions, since relevant objects for new actions can be obtained on-the-fly through word embedding matching with object names. In this work, we follow this line of work and generalize to spatio-temporal localization by modeling the spatial relations between actors

and objects. This allows us to perform action classification and localization within the same approach. Different from the common setup for zero-shot actions (Junior et al. 2019), we do not assume access to any training videos of seen actions. We seek to recognize actions in video without ever having seen a video before, solely by relying on prior knowledge about objects in images and their relation to actions.

To improve semantic matching, Alexiou et al (2016) correct class names to increase unseen action discrimination. Similar in spirit are approaches that employ query expansion (Dalton et al. 2013; de Boer et al. 2016) or textual action descriptions (Gan et al. (2016c); Habibian et al. 2017; Wang and Chen 2017) to make the action inputs more expressive. In contrast, we focus on improving the semantic matching itself to deal with semantic ambiguity, non-discriminative objects, and object naming.

2.2 Unseen Action Localization

Spatio-temporal localization of actions without examples is hardly investigated in the current literature. Jain et al. ((2015a)) split each test video into spatio-temporal proposals (Jain et al. 2017). Then for each proposal, boxes are sampled and individually fed to a pre-trained object classification network to obtain object scores. The object scores of each proposal are semantically matched to the action and the best matched proposal is selected as the location of interest. In this paper, we employ local object detectors and embed spatial relations between humans and objects. Where Jain et al. ((2015a)) implicitly assume that the spatial location of objects and the humans performing actions is identical, our spatial object priors explicitly model how humans and objects are spatially related, whether objects are above, to the left, or on the human. Moreover, we go beyond standard word embedding similarities for semantic matching between actions and objects to improve both unseen action classification and localization. Soomro and Shah (2017) investigate action localization in an unsupervised setting, which discriminatively clusters similar action tubes but does not specify action labels. In contrast, we seek to discover both action locations and action labels without training examples or manual action annotations.

Several works have investigated unseen action localization in the temporal domain. (Zhang et al. 2020) perform zero-shot temporal action localization by transferring knowledge from temporally annotated seen actions to unseen actions. Jain et al. (2020) learn an action localization model from seen actions in trimmed videos, enabling zero-shot temporal action localization by a semantic knowledge transfer of unseen actions. Sener and Yao (2018) learn to temporally segment actions in long videos in an unsupervised manner. Different from these works, we perform unseen action localization in space and time simultaneously.

2.3 Self-supervised Video Learning

Recently, a number of works have proposed approaches for representation learning for unlabeled videos through self-supervision. The general pipeline is to train a pre-text task on unlabeled data and transfer the knowledge to a supervised downstream task (Jing and Tian 2020) or by clustering video datasets without manual supervision (Asano et al. 2020). Pre-text tasks include dense predictive coding (Han et al. 2020), shuffling frames (Fernando et al. 2017; Xu et al. 2019), exploiting spatial and/or temporal order (Jenni et al. 2020; Tschannen et al. 2020; Wang et al. 2019), or by matching frames with other modalities (Afouras et al. 2020; Alayrac et al. 2020; Owens and Efros 2018; Patrick et al. 2020). Self-supervised approaches utilize unlabeled train videos to learn representations without semantic class labels. In contrast, we do not use any training videos and instead classify and localize actions using object classes and bounding boxes from images. Since we do not assume any video knowledge, common losses and notions from the zero-shot and self-supervised literature can not be leveraged. It is the object priors that still allow us to classify and spatio-temporally localize unseen actions in videos.

3 Spatial Object Priors

In unseen action localization, the aim is to discover a set of spatio-temporal action tubes from test videos for each action in the set of all actions $\mathcal{A} = \{A_1, \dots, A_C\}$, with C the total number of actions. Furthermore, unseen action classification is concerned with predicting the label of each test video from \mathcal{A} . For each action, nothing is known except its name. The evaluation is performed on a set of N unlabeled and unseen test videos denoted as \mathcal{V} . In this section, we outline how to obtain such a localization and classification with spatial priors from local objects using prior knowledge.

3.1 Priors from Persons, Objects, and Prepositions

For a test video $v \in \mathcal{V}$ and unseen action $a \in \mathcal{A}$, the first step of our approach is to score local boxes in the video with respect to a . For a bounding box b in video frame F , we define a score function $s(\cdot)$ for action class a . The score function is proportional to three priors.

Object prior I (person prior) The likelihood of any action in b is proportional to the likelihood of a person present in b . The first prior follows directly from our human action recognition task. The first condition is independent of the specific action class, as it must hold for any action. The score function therefore adheres to the following:

$$s(b, F, a) \propto Pr(\text{person}|b). \quad (1)$$

Object prior II (object location prior) The likelihood of action a in box b is proportional to the likelihood of detected objects that are (i) semantically close to action class a and (ii) the detection is sufficiently close to b .

The second prior states that the presence of an action in a box b also depends on the presence of relevant objects in the vicinity of b . We formalize this as:

$$s(b, F, a) \propto \sum_{o \in \mathcal{L}} \Psi(o, a) \cdot \max_{b' \in o_D(F, b)} Pr(o|b'), \tag{2}$$

where \mathcal{L} denotes the set of pre-trained object detections and $o_D(F, b)$ denotes the set of all object detections of object o in frame F that are near to box b . Empirically, the second object prior is robust to the pixel distance to determine the neighbourhood set $o_D(F, b)$ for box b , as long as it is a non-negative number smaller than the frame size. We use a value of 25 throughout. Function $\Psi(o, a)$ denotes the semantic similarity between object o and a and is defined as the word embedding similarity:

$$\Psi(o, a) = \cos(\phi(o), \phi(a)), \tag{3}$$

with $\phi(\cdot) \in \mathbb{R}^{300}$ the word embedding representation. The word embeddings are given by a pre-trained word embedding model, such as word2vec (Mikolov et al. 2013), FastText (Grave et al. 2018), or GloVe (Pennington et al. 2014).

Object prior III (spatial relation prior) The likelihood of action a in b given an object o with box detection d that abides object prior II, is proportional to the match between the spatial awareness of b and d with the prior spatial awareness of a and o .

The third prior incorporates spatial awareness between actions and objects. We exploit the observation that people interact with objects in preferred spatial relations. We do this by gathering statistics from the same image dataset used to pre-train the object detectors. By reusing the same dataset, we keep the amount of knowledge sources contained to a dataset for object detectors and a semantic word embedding. For the spatial relations, we examine the bounding box annotations for the person class and all object classes. We gather all instances where an object and person box annotation co-occur. We quantize the gathered instances into representations that describe coarse spatial prepositions between people and objects.

The spatial relation between an object box relative to a person box is quantized into a 9-dimensional grid. This grid represents how the object box is spatially distributed to the person box with respect to the following prepositions: {*above left, above, above right, left, on, right, below left, below, below right*}. Since no video examples are given in our setting, prepositions can only be obtained from prior image




Fig. 1 Intuition behind spatial object priors. The spatial relations (end of green arrows) of the two persons (red boxes) have different spatial relations with the detected skateboard (blue box). The spatial relations for the person on the left are a better match with the spatial relations obtained from prior knowledge. This match enforces the likelihood that the person on the left is involved in a skateboarding activity

sources and we therefore exclude relations such as *in front of* and *behind of*. Let $d_1(b, d) \in \mathbb{R}^9$ denote the spatial distribution of object box d relative to person box b . Furthermore, let $d_2(\text{person}, o)$ denote the gathered distribution of object o with respect to a person from the image dataset. We define the spatial relation function as:

$$\Phi(b, d, o) = 1 - \text{JSD}_2(d_1(b, d) || d_2(\text{person}, o)), \tag{4}$$

where $\text{JSD}_2(\cdot || \cdot) \in [0, 1]$ denotes the Jensen-Shannon Divergence with base 2 logarithm. Intuitively, this function determines the extent to which the 9-dimensional distributions match, as visualized in Fig. 1. The more similar the distributions, the lower the divergence, and the higher the score according to Equation 4.

Combined spatial priors Our final box score combines the priors of persons, objects, and spatial prepositions. We combine the three priors into the following score function for a box b with respect to action a :

$$s(b, F, a) = Pr(\text{person}|b) + \sum_{o \in \mathcal{O}} \Psi(o, a) \cdot \max_{b' \in o_D(F, b)} \left(Pr(o|b') \cdot \Phi(b, b', o) \right). \tag{5}$$

3.2 Linking Action Tubes

Given scored boxes in individual frames, we link boxes into tubes to arrive at a spatio-temporal action localization. We link boxes that have high scores from our object embeddings and have a high spatial overlap. Given an action a and boxes b_1 and b_2 in consecutive frames F_1 and F_2 , the link score is given as:

$$w(b_1, b_2, a) = s(b_1, F_1, a) + s(b_2, F_2, a) + \text{iou}(b_1, b_2), \tag{6}$$

where $\text{iou}(\cdot, \cdot)$ states the spatial intersection-over-union score. We solve the problem of linking boxes into tubes with the Viterbi algorithm (Gkioxari and Malik 2015). For a video V , we apply the Viterbi algorithm on the link scores to obtain spatio-temporal action tubes. In each tube, we continue linking as long as there is at least one box in the next frame with an overlap higher than 0.1 and with a combined action score of at least 1.0. Otherwise we stop linking. Incorporating the stopping criterion allows us to localize actions in time also, akin to (Gkioxari and Malik 2015). We reiterate this process until we obtain T tubes. The action score for a of an action tube t is defined as the average score of the boxes in the tube:

$$\ell_{\text{tube}}(t, a) = \frac{1}{|t|} \sum_{i=1}^{|t|} s(b_{t_i}, F_{t_i}, a), \quad (7)$$

where b_{t_i} and F_{t_i} denote respectively the box and frame of the i^{th} element in t .

Unseen action localization and classification For unseen action localization, we gather tubes across all test videos and rank the tubes using the scores provided by Equation 7. We can also perform unseen action classification using the spatial priors by simply disregarding the tube locations. For each video, we predict the action class label as the action with the highest tube score within the video.

3.3 Action Tube Retrieval

The use of objects with spatial priors extends beyond unseen action classification and localization. We can also perform a new task, dubbed action tube retrieval. This task resembles localization, as the goal is to rank the most relevant tubes the highest. Different from localization, we now have the opportunity to specify which objects are of interest and which spatial relations are desirable for a detailed result. Furthermore, inspired by the effectiveness of size in actor-object relations (Escorcia and Niebles 2013), we extend the retrieval setting by allowing users to specify a desired relative size between actors and objects. The ability to specify the object, spatial relations, and size allows for different localizations of the same action. To enable such a retrieval, we extend the box score function of Equation 5 as follows:

$$s(b, F, o, r, s) = Pr(\text{person}|b) + \max_{b' \in o_D(F, b)} \left(Pr(o|b') \cdot \Phi_r(b, b', r) \cdot \left(1 - \left|\frac{\text{size}(b')}{\text{size}(b)} - s\right|\right) \right), \quad (8)$$

where o denotes the user-specified object, $r \in \mathbb{R}^9$ the specified spatial relations, and s the specified relative size. The spatial relation function is modified to directly match box

relations to specified relations:

$$\Phi_r(b, d, r) = 1 - \text{JSD}_2(d_1(b, d)||r). \quad (9)$$

With the three user-specified objectives, we again score individual boxes first and link them over time. The tube score is used to rank the tubes across a video collection to obtain the final retrieval result.

4 Semantic Object Priors

Spatial object priors relying on local objects enables a spatio-temporal localization of unseen actions. However, local objects do not tell the whole story. When a person performs an action, this is typically happens in a suitable context. Think about someone *playing tennis*. While the tennis racket provides a relevant cue about the action and its location, surrounding objects from context, such as tennis court and tennis net, further enforce the action likelihood. Here, we add three additional object priors to integrate knowledge from global objects for unseen action classification and localization. We start from the common word embedding setup for semantic matching, which we extend with three simple priors that make for effective unseen action matching with global objects. Lastly, we outline how to integrate the semantic and spatial object priors for unseen actions. Figure 2 illustrates our proposal.

4.1 Matching and Scoring with Word Embeddings

To obtain action scores for a video $v \in \mathcal{V}$, the common setup is to directly use the object likelihoods from a set of global objects \mathcal{G} and their semantic similarity. Since \mathcal{G} typically contains many objects, the usage is restricted to the objects with the highest semantic similarity to action a :

$$\Psi(g, a) = \cos(\phi(g), \phi(a)) \text{ such that } g \in \mathcal{G}_a, \quad (10)$$

where \mathcal{G}_a the set of k most similar objects with respect to a . The video score function is defined as:

$$\ell_{\text{video}}(v, a) = \sum_{g \in \mathcal{G}_a} \Psi(g, a) \cdot Pr(g|v), \quad (11)$$

where $Pr(g|v)$ denotes the likelihood of g in v , as given by the softmax outputs of a pre-trained object classification network. Such an approach has shown to be effective for unseen action classification (Jain et al. (2015a)). Here, we identify three additional semantic priors to improve both unseen action classification and localization.




Fig. 2 Intuition behind our three semantic object priors. The red and orange distributions denote the word embeddings of the action kicking in English and Dutch. The closer to the center an object is, the higher the semantic similarity to the action. In (a), the object football is enforced, because its semantic similarity is high across languages, reducing

semantic ambiguity. In (b), the importance of grass is decreased, as it is also relevant for another action, while the opposite happens for goal post. In (c), the importance of football is increased and of adjudicator decreased, as football follows basic-level object naming, in contrast to adjudicator (a referee)

4.2 Priors for Ambiguity, Discrimination, and Naming

Similar to the common word embedding setup, for a video $v \in \mathcal{V}$, we seek to obtain a score for action $a \in \mathcal{A}$ using a set of global objects \mathcal{G} . Global objects generally come from deep networks (Mettes et al. 2020) pre-trained on large-scale object datasets (Deng et al. 2009). We build upon current semantic matching approaches by providing three simple priors that deal with semantic ambiguity, non-discriminative objects, and object naming.

Object prior IV (semantic ambiguity prior) A zero-shot likelihood estimation of action a in video v benefits from minimal semantic ambiguity between a and global objects \mathcal{G} .

The score of a target action depends on the semantic relations to source objects. However, semantic relations can be ambiguous, since words can have multiple meanings depending on the context. For example for the action *kicking*, an object such as *tie* is deemed highly relevant, because one of its meanings is a draw in a football match (Mettes and Snoek 2017). However, a *tie* can also denote an entirely different object, namely a *necktie*. Such semantic ambiguity may lead to the selection of irrelevant objects for an action.

To combat semantic ambiguity in the selection of objects, we consider two properties of object coherence across languages (Malt 1995). First, most object categories are common across different languages. Second, the formation of some categories can nevertheless differ among languages. We leverage these two properties of object coherence across languages by introducing a multi-lingual semantic similarity. For computing multi-lingual semantic representations of words at a large-scale, we are empowered by recent advances in the word embedding literature, where embedding models have been trained and made publicly available for many languages (Grave et al. 2018). In a multi-lingual setting, let L denote the total number of languages to use. Furthermore, let $\tau_l(g)$ denote the translator for language $l \in L$ applied to object g . Multi-lingual unseen action classification can then

be done by simply updating the semantic matching function to:

$$\psi_L(g, a) = \frac{1}{L} \sum_{l=1}^L \cos(\phi_l(\tau_l(g)), \phi_l(\tau_l(a))), \tag{12}$$

where ϕ_l denotes the semantic word embedding of language l . The multi-lingual semantic similarity states that for a high semantic match between object and action, the pair should be of a high similarity across languages. In this manner, accidental high similarity due to semantic ambiguity can be addressed, as this phenomenon is factored out over languages.

Object prior V (object discrimination prior) A zero-shot likelihood estimation of action a in video v benefits from knowledge about which objects in \mathcal{G} are suitable for action discrimination.

The second semantic prior is centered around finding discriminative objects. Only using semantic similarity to select objects ignores the fact that an object can be non-discriminative, despite being semantically similar. For example, for the action *diving*, the objects *person* and *diving board* might both correctly be considered as semantically relevant. The object *person* is however not a strong indicator for the action *diving*, as this object is present in many actions. The object *diving board* on the other hand is a distinguishing indicator, as it is not shared by many other actions.

To incorporate an object discrimination prior, we take inspiration from object taxonomies. When organizing such taxonomies, care must be taken to convey the most important and discriminant information (Murphy 2004). Here, we are searching for the most unique objects for actions, *i.e.*, objects with low inclusivity. It is desirable to select indicative objects, rather than focus on objects that are shared among many actions. To do so, we propose a formulation to predict the relevance of every object for unseen actions. We extend

the action-object matching function as follows:

$$\Psi_r(g, a) = \Psi(g, a) + r(g, \cdot, a), \quad (13)$$

where $r(g, \cdot, a)$ denotes a function that estimates the relevance of object g for the action a . We propose two score functions. The first penalizes objects that are not unique for an action a :

$$r_a(g, A, a) = \Psi(g, a) - \max_{c \in A \setminus a} \Psi(g, c). \quad (14)$$

An object g scores high if it is relevant for action a and for no other action. If either of these conditions are not met, the score decreases, which negatively affects the updated matching function.

The second score function solely uses inter-object relations for discrimination and is given as:

$$r_o(g, \mathcal{G}, a) = \Psi(g, a) - \frac{1}{|\mathcal{G}|} \sum_{g' \in \mathcal{G} \setminus g} \Psi(g, g')^{\frac{1}{2}}. \quad (15)$$

Intuitively, this score function promotes objects that have an intrinsically high uniqueness across the set of objects, regardless of their match to actions. The square root normalization is applied to reduce the skewness of the object set distribution.

Object prior VI (object naming prior) A zero-shot likelihood estimation of action a in video v benefits from a bias towards basic-level object names.

The third semantic prior concerns object naming. The matching function between actions and objects relies on the object categories in the set \mathcal{G} . The way objects are named and categorized has an influence on their matching score with an action. For example for the action *walking with a dog*, it would be more relevant to simply name the object present in the video as a *dog* rather than a *domesticated animal*, or an *Australian terrier*. Indeed, the *dog* naming yields a higher matching score with the action *walking with a dog* than the too generic *domesticated animal* or too specific *Australian terrier* namings.

As is well known, there exists a preferred entry-level of abstraction in linguistics, for naming objects (Jolicoeur et al. 1984; Rosch et al. 1976). The basic-level naming (Rosch et al. 1976; Rosch 1988) is a trade-off between superordinates and subordinates. Superordinates concern broad category sets, while subordinates concern very fine-grained categories. Hence, basic-level categories are preferred because they convey the most relevant information and are discriminative from one another (Rosch et al. 1976). It would then be valuable to emphasize basic-level objects rather than objects from other levels of abstraction. Here, we enforce such an emphasis by using the relative WordNet depth of the objects in \mathcal{G} to weight

each object. Intuitively, the deeper an object is in the WordNet hierarchy, the more specific the object is and vice versa. To perform the weighting, we start from the beta distribution:

$$\begin{aligned} \text{Beta}(d|\alpha, \beta) &= \frac{d^{\alpha-1} \cdot (1-d)^{\beta-1}}{B(\alpha, \beta)}, \\ B(\alpha, \beta) &= \frac{\Gamma(\alpha) \cdot \Gamma(\beta)}{\Gamma(\alpha+\beta)}, \end{aligned} \quad (16)$$

where d denotes the relative depth of an object and $\Gamma(\cdot)$ denotes the gamma function. Different values for α and β determine which levels to focus on. For a focus on basic-level we want to weight objects of intermediate level higher and the most specific and generic objects lower. We can do so by setting $\alpha = \beta = 2$. Setting $\alpha = \beta = 1$ results in the common setup where all objects are equally weighted. We incorporate the objects weights by adjusting the semantic similarity function between objects and actions.

Combined semantic priors We combine the three semantic object priors into the following function of global objects for unseen actions:

$$\ell_{\text{video}}(v, a) = \sum_{g \in \mathcal{G}_a} ((\Psi_L(g, a) + \Delta(o, \cdot, a)) \cdot \text{Beta}(d_g|\alpha, \beta)) \cdot Pr(g|v), \quad (17)$$

where d_g denotes the depth of object g , [0,1] normalized based on the minimum WordNet depth (2) and maximum WordNet depth (18) over all objects in \mathcal{G} . In this formulation, the proposed embedding is more robust to semantic ambiguity, non-discriminative objects, and non-basic level objects compared to Equation 10.

4.3 Object Prior Embedding

Unseen action localization and classification benefit from both a spatial and semantic priors. For unseen action localization, we obtain an object prior embedding by simply adding the tube score (Equation 7) and the score of the corresponding video (Equation 17). For unseen action classification we add the highest score of the tubes in the video with the video score.

5 Experimental Setup

5.1 Datasets

We experiment on UCF Sports (Rodriguez et al. 2008), J-HMDB (Jhuang et al. 2013), UCF-101 (Soomro et al. 2012), Kinetics (Carreira and Zisserman 2017), and AVA (Gu et al. 2018). Due to the lack of training examples, all these datasets still form open challenges in unseen action literature, even though high scores can be achieved with supervised

approaches on *e.g.*, UCF-101 (Carreira and Zisserman 2017; Zhao and Snoek 2019).

UCF Sports contains 150 videos from 10 actions such as *running* and *horse riding* (Rodriguez et al. 2008). The videos are from sports broadcasts. We employ the test split provided by Lan et al. (2011).

J-HMDB contains 928 videos from 21 actions such as *brushing hair* and *catching* (Jhuang et al. 2013), from HMDB (Kuehne et al. 2011). The videos focus on daily human activities. We employ the test split provided by Jhuang et al. (2013).

UCF-101 contains 13,320 videos from 101 actions such as *skiing* and *playing basketball* (Soomro et al. 2012). The videos are taken from both sports and daily activities. We employ the test split provided by Soomro et al. (2012).

Kinetics-400 contains 104,000 videos from 400 actions such as *playing monopoly* and *zumba* Carreira and Zisserman (2017) from Youtube videos. We use all videos as test for unseen action classification.

AVAv2.2 contains 437 15-minutes clips from movies covering 80 atomic actions such as *listening* and *writing* Gu et al. (2018). For 61 out of 64 validation videos, the YouTube links are still available and we use these as test videos for unseen action localization.

Note that for all datasets, we exclude the use of *any* information from the training videos. We employ the action labels and ground truth box annotations from the test videos to evaluate the zero-shot action classification and localization performance.

5.2 Object Priors Sources

Object scores and detections To obtain person and local object box detections in individual frames, we employ Faster R-CNN (Ren et al. 2015), pre-trained on MS-COCO (Lin et al. 2014). The pre-trained network includes the person class and 79 objects, such as *car*, *chair*, and *tv*. For the global object scores over whole videos, we apply a GoogLeNet (Szegedy et al. 2015), pre-trained on 12,988 ImageNet categories (Mettes et al. 2020). The object probability distributions are averaged over the sampled frames to obtain the global object scores. On all datasets except AVA, frames are sampled at a fixed rate of 2 frames per second. On AVA, we use the annotated keyframes as frames. All frames have an input size of 224x224 (Table 1).

Spatial priors sources For the spatial relations, we reuse the bounding box annotations of the training set of MS-COCO, as also used to pre-train the detection model, to obtain the prior prepositional knowledge between persons and objects.

Semantic priors sources For the semantic priors, we rely on FastText, pre-trained on 157 languages (Grave et al. 2018). This collection of word embeddings enables us to investi-

gate multi-lingual semantic matching between actions and objects. For the multi-lingual experiments, we employ five languages: English, French, Dutch, Italian, and Afrikaans. We obtain action and object translations first from Open Multilingual WordNet (Bond and Foster 2013). For the remaining objects and all actions, we use Google Translate with manual verification.

Code is available at <https://github.com/psmmettes/object-priors-unseen-actions>.

5.3 Evaluation Protocol

We follow the zero-shot action evaluation protocol of (Jain et al. (2015a); Mettes and Snoek 2017; Zhu et al. 2018), where no training is performed on a separate set of actions; the set of test actions are directly evaluated. For each dataset, we evaluate on the videos in the test set. For classification experiments where the number of test actions is lower than the total number of actions in the dataset, we perform five random selections and report the mean accuracy and standard deviation.

For unseen action localization, we compute the spatio-temporal (st) overlap between action tube a and ground truth b from the same video as:

$$\text{st-iou}(a, b) = \frac{1}{|\Omega|} \sum_{f \in \Omega} \text{iou}_f(a, b), \quad (18)$$

where Ω states the union of frames in a and b . The function $\text{iou}_f(a, b)$ is 0 if either one of the tubes is not present in frame f . For overlap threshold τ , an action tube is positive if the tube is from a positive video, the overlap with a ground truth instances is at least τ , and the ground truth instance has not been detected before. For unseen action localization, we report the AUC and video mAP metrics on UCF Sports and J-HMDB, following Mettes and Snoek (2017). On AVA, we report frame mAP, following Gu et al. (2018). Unless specified otherwise, the overlap threshold is 0.5. For unseen action classification, we evaluate using multi-class classification accuracy.

6 Results

6.1 Spatial Object Priors Ablation

In the first experiment, we evaluate the importance of spatial relations between persons and local object detections for unseen action classification and localization. We use the 80 local objects pre-trained on MS-COCO for this ablation study. We investigate the desired number of local objects to select per action and the effect of modelling spatial relations.

Results are shown in Table 1. When relying on only the first prior, person detections, we unsurprisingly obtain ran-

Table 1 Effect of spatial object priors for unseen action classification (acc, %) and localization (mAP@0.5, %), on UCF Sports

		Classification					Localization				
		Number of object detections					Number of object detections				
		0	1	2	5	10	0	1	2	5	10
Object prior I	Person	8.5	–	–	–	–	10.1	–	–	–	–
Object prior I+II	+ Objects	–	21.3	19.2	27.7	27.7	–	22.8	22.8	24.4	23.6
Object prior I+II+III	+ Spatial relations	–	12.8	25.5	29.8	29.8	–	26.0	22.4	27.0	22.8

Bold indicates best performance for each setting (relevant for the corresponding experiment)

We investigate three spatial prior settings; only person detections (I), person and object detections (I+II), and with the additional spatial prepositions between people and objects (I+II+III). For both unseen classification and localization, using the top five objects with spatial relations obtains the highest scores

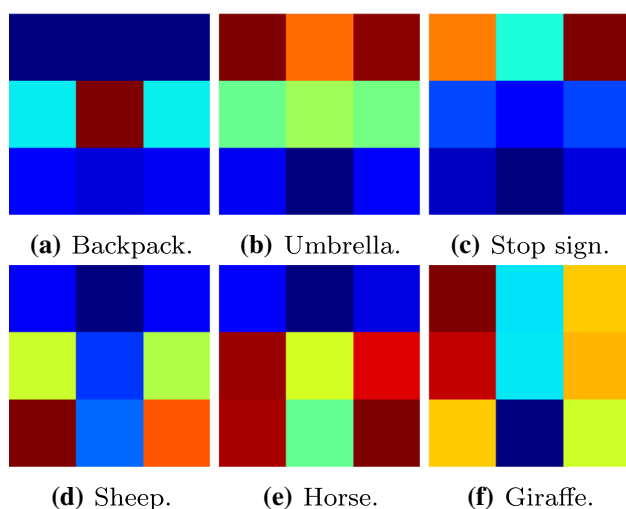


Fig. 3 Spatial preposition priors for six local objects. Different objects have different spatial preferences relative to persons. These prepositional preferences align with our intuitions of the objects, e.g., an umbrella tends to be above a person, while a backpack tends to be on a person

dom classification and localization scores, since there is no direct manner to differentiate actions. Naturally, the first object prior is still vital, since it determines which boxes to consider in video frames. When adding the second prior, we find that the scores improve drastically for both classification and localization. Objects are indicative for unseen actions, whether actions need to be classified or localized.

Lastly, we include the spatial preposition prior. This provides a further boost in the results, showing that persons and objects have preferred spatial relations that can be exploited. In Fig. 3, we provide six discovered spatial relations from prior knowledge that are used in our action localization.

The results of Table 1 show that for unseen action classification, more local objects improve accuracy as they provide a richer source for action discrimination. For action localization, having many local objects may hurt, as the local box scoring becomes noisier, resulting in action tubes with lower overlap to the ground truth. Based on the scores obtained in

Table 2 Object prior IV (semantic ambiguity prior) ablation

	UCF-101		
	Number of test classes		
	25	50	101
<i>Single language</i>			
English	52.6 ± 4.7	43.3 ± 2.1	33.3
Dutch	49.7 ± 3.5	40.4 ± 3.3	30.2
Portuguese	44.4 ± 4.6	37.8 ± 4.0	28.5
Afrikaans	43.6 ± 3.4	36.3 ± 4.0	27.7
French	44.7 ± 2.5	36.2 ± 3.4	27.5
German	38.0 ± 2.0	31.2 ± 2.6	26.0
<i>Multi-lingual</i>			
English and Dutch	54.3 ± 4.5	45.8 ± 2.4	35.7
All languages	51.8 ± 4.8	43.2 ± 2.7	32.8

Bold indicates best performance for each setting (relevant for the corresponding experiment)

Unseen action classification accuracies (%) on UCF-101 for multiple languages in the semantic matching for 25, 50, and 101 test classes. Combining two languages improves results. For this dataset, a combination of English and Dutch is best

this experiment, we recommend the use of spatial prepositions and five local object detections per action.

6.2 Semantic Object Priors Ablations

In the second experiment, we perform ablation studies on the three semantic object priors for semantic matching between unseen actions and objects. We evaluate unseen action classification on UCF-101. Throughout this experiment, we focus on global object classification scores from the 12,988 ImageNet concepts applied and averaged over sampled video frames.

Object prior IV (semantic ambiguity prior) We first investigate the importance of multi-lingual semantic similarity to deal with semantic ambiguity. We evaluate on three settings of UCF-101 for 25, 50, and 101 test classes. We perform this evaluation on all five individual languages, as well as their combination. We select the top-100 objects per action, following Mettes and Snoek (2017).




Fig. 4 Pairwise multilingual evaluation of all six languages on UCF-101 with all 101 test actions. The better the performance of the individual language, the more that language benefits others. For English, only adding the second best performing language (Dutch) is beneficial. When not taking English into account, we find that combining languages is mutually effective for seven out of the ten combinations

The results are shown in Table 2. We first observe that individually English performs better than the other four languages. Dutch performs roughly three percent point lower, while the other three languages perform five to nine percent lower. A likely explanation for the lower results of the other languages is that the starting language of the objects and actions is English. The object and action names of the other languages are translated from English. Translation imperfections and breaking up compound nouns into multiple terms result in less effective word representations. As a result, there is a gap between English and the other languages.

In Fig. 4, we show the relative accuracy scores for all language pairs on UCF-101 with all 101 test actions. We find that combining languages always boosts the least effective language of the pair. For the most effective English language, only the addition of Dutch results in a higher accuracy. For all other language pairs, the combined language performance is higher than the best individual language, except for German-Portuguese, German-Afrikaans, and Dutch-Portuguese. These are likely a result of poor individual performance (German) or low lexical similarity to other languages (Portuguese). Overall, multi-lingual similarity with English and Dutch results in an improvement of 1.7% 2.5% and 2.4% for 25, 50 and 101 classes. Further improvements are expected with better translations.

To investigate why multiple languages aid unseen action classification, we have performed a qualitative analysis for the action *field hockey penalty* in UCF-101. We consider the most similar objects when using English only and when using English and Dutch combined. Figure 5 shows that for English only, several of the top ranked objects are not correct due to semantic ambiguity. These objects include penal institution,




Fig. 5 Object prior IV (semantic ambiguity prior) analysis with multiple languages for unseen recognition of the action *field hockey penalty*. When relying on English, several irrelevant objects rank high due to semantic ambiguity (red boxes). When Dutch is added, ambiguous objects are downgraded, resulting in better recognition

Table 3 Object prior V (object discrimination prior) ablation

	UCF-101		
	Number of test classes		
	25	50	101
Standard setup	52.6 ± 4.7	43.3 ± 2.1	33.3
+ action-based discrimination	53.2 ± 4.3	44.3 ± 1.9	34.3
+ object-based discrimination	54.0 ± 3.6	44.7 ± 2.1	34.0

Bold indicates best performance for each setting (relevant for the corresponding experiment)

Unseen action classification accuracy (%) on UCF-101 with the proposed object discrimination functions for the English language. Both action-based and object-based discrimination aid recognition, especially when using fewer actions

field artillery, and field wormwood. Evidently, such objects were selected because of their similarity to the English words field and penalty, but they are not related to the action of interest. When adding Dutch to the matching, such objects are ranked lower, because the ambiguity of these objects do not translate to Dutch. Hence, more relevant objects are ranked higher, which is also reflected in the results, where the accuracy increases from 0.07 to 0.27 for the action.

We conclude that using multiple languages for semantic matching between actions and objects reduces semantic ambiguity, resulting in improved unseen action classification accuracy.

Object prior V (object discrimination prior) For the object discrimination prior ablation, we investigate both the proposed object-based and action-based prior variants. We again report on UCF-101 with 25, 50, and 101 test actions, with the top 100 objects selected per action.

The results in Table 3 show consistent improvements are obtained by both the action-based and the object-based vari-

Table 4 Object prior V (object discrimination prior) analysis for two UCF-101 actions

Apply eye makeup		Pizza tossing	
<i>Most discriminative objects</i>			
Makeup	0.63	0.88	Pizza
Eyeliner	0.57	0.85	Pepperoni pizza
Eye eyebrow pencil	0.52	0.82	Sausage pizza
Eyeshadow	0.51	0.79	Cheese pizza
Mascara	0.50	0.78	Anchovy pizza
<i>Least discriminative actions</i>			
Edam (cheese)	− 0.07	− 0.06	Argali (sheep)
Hokan (people)	− 0.07	− 0.07	Lhasa (terrier dog)
Lincoln (sheep)	− 0.07	− 0.07	Yautia (food)
Dicot (plant)	− 0.08	− 0.08	Caddo (people)
Loranthaceae (plant)	− 0.09	− 0.08	Filovirus (virus)

We show objects deemed most and least discriminative for *apply eye makeup* and *pizza tossing*, along with their scores. By finding out which objects are uniquely discriminative for an action in comparison to all other actions, we are able to highlight relevant objects and in turn improve unseen action classification

Table 5 Object prior VI (object naming prior) ablation

Weighting preference	α	β	accuracy
Uniform	1	1	43.3 ± 2.1
Specific only	5	1	7.6 ± 0.7
Generic only	1	5	30.2 ± 1.1
Basic-level	2	2	43.9 ± 2.0

Bold indicates best performance for each setting (relevant for the corresponding experiment)

Unseen action classification accuracy (%) on UCF-101 for 50 test classes using English. Only a small gain is feasible with a focus on basic-level objects compared to uniform weighting

ants. While the object-based taxonomy is preferred when recognizing 25 or 50 actions, the action-based taxonomy is preferred when recognizing 101 activities. In all three cases, incorporating a selection of the most discriminative objects yields better results. To highlight what kind of objects are boosted and subdued, we show the most and least discriminative objects of two actions in Table 4.

Object prior VI (object naming prior) For the third semantic object prior, we evaluate the effect of weighting objects based on their WordNet depth to understand whether a bias towards basic-level objects is desirable in unseen action classification. This experiment is performed on UCF-101 for 50 test actions.

Table 5 shows the results for the basic-level weighting preference compared to three baselines, *i.e.*, uniform (no preference), specific only, and generic only. We find that focusing on only the most specific or generic objects is not desirable and both result in a large drop in classification accuracy. The weighting preference for basic-level objects has a

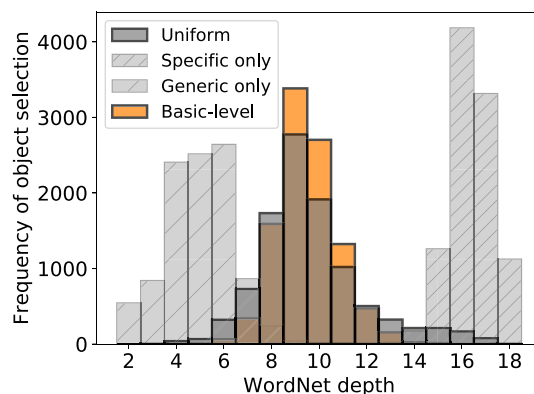


Fig. 6 Object prior VI analysis on UCF-101. Akin to our basic-level object prior does the uniform weighting result in a distribution that favors basic-level objects. This explains the competitive performance of uniform weights versus the basic-level object prior; a bias towards basic-level objects is inherent in large-scale object sources. An explicit basic-level prior provides marginal gains

slight increase in accuracy compared to uniform, although the difference is small. This results shows that a prior for basic-level objects is not as effective as the semantic ambiguity and object discrimination priors.

To better understand our results, we have analysed the WordNet depth distribution of the top 100 selected objects for all actions in UCF-101. The distributions are visualized in Fig. 6. The two extreme preference weightings select objects from expected depth distributions and focus on the leftmost or rightmost side of the depth spectrum. Similarly for the basic-level weighting, objects from intermediate depth are selected. The uniform weighting however behaves unexpectedly and does not result in a uniform object depth distribution. In fact, this function also favors basic-level objects. The reason for this behaviour is found in the depth distribution of all 12,988 objects. For large-scale object collections, the WordNet depth distribution favors basic-level objects, following a normal distribution. As a result, the depth distribution of the selected objects follows a similar distribution, hence creating an inherent emphasis on basic-level objects. The basic-level object prior puts an additional emphasis on these kinds of objects and ignores specific and generic objects altogether.

We conclude that a prior on basic-level objects *is* important for unseen actions. Such a bias is inherently incorporated in large-scale object sources and no additional weighting is required to assist the object selection, although a small increase is feasible (Fig. 7).

Combining semantic priors In Table 6, we report the unseen action classification performance on Kinetics-400 using the semantic priors. Our approach does not require any class labels and videos during training, enabling a 400-way unseen action classification. When performing 400-way classification, the semantic ambiguity (IV) and object naming (VI) priors are most decisive, resulting in an accuracy of 6.4%,

Table 6 Unseen action classification on Kinetics-400 for the three semantic priors

Object prior			Test actions		
Semantic ambiguity (IV)	Object discrimination (V)	Object naming (VI)	25	100	400
	Random		4.0	1.0	0.3
	English only		21.8±3.5	10.8±1.0	6.0
✓			20.9±4.1	10.8±1.0	6.3
✓	✓		21.2±3.9	10.7±1.0	6.1
✓		✓	22.0±3.7	11.2±1.0	6.4
✓	✓	✓	21.9±3.8	11.1±0.8	6.4

Even with hundreds of unseen actions, the object priors help to assign action labels to videos. Across the three action sizes, semantic ambiguity and object naming work best, especially when having more unseen actions to choose from

Table 7 The top-10 and bottom-10 performing actions (acc, %) on UCF-101 and Kinetics using an English-Dutch vocabulary

UCF-101		Kinetics	
Bowling	98.1	65.0	Playing poker
Ice dancing	96.8	54.2	Strumming guitar
Sumo Wrestling	92.2	51.5	Using segway
Horse riding	91.5	48.5	Golf chipping
Playing piano	91.4	48.1	Bowling
Playing sitar	90.4	43.6	Playing bass guitar
Rowing	89.8	41.5	Playing cymbals
Biking	89.6	40.8	Playing squash
Golf swing	89.2	39.2	Playing badminton
Playing violin	88.0	39.2	Playing cello
Yo yo	00.0	00.0	Zumba
Hammer throw	00.0	00.0	Skiing
Jump rope	00.0	00.0	Egg hunting
Front crawl	00.0	00.0	Exercising arm
Frisbee catch	00.0	00.0	Exercise with ball
Floor gymnastics	00.0	00.0	Crosscountry skiing
Jumping jack	00.0	00.0	Faceplanting
Writing on board	00.0	00.0	Feeding fish
Lunges	00.0	00.0	Eating chips
Pizza tossing	00.0	00.0	Situp

Across both datasets, our approach is effective for actions with clear object interactions (e.g., *bowling*, *playing instruments*, *horse riding*, and *biking*). Actions can not be recognized when they are without direct object interactions (e.g., fitness actions such as *jumping jacks*, *zumba*, and *exercising arm*) or when they use objects for which we have no detector or classifier (e.g., *yo yo* and *exercising with exercise ball*)

compared to 0.25% for random performance. For the Kinetics experiment, we evaluate unseen action classification as a function of the number of actions. For each size of the action vocabulary, we perform a random selection of the actions and perform 5 runs. We report both the mean and standard deviation.

For which actions are semantic priors effective? In Table 7, we show respectively the top and bottom performing actions

on UCF-101 and Kinetics when using our priors. On Kinetics, high accuracies can be achieved for actions such as *playing poker* (65.0%) and *strumming guitar* (54.2%), the accuracy is hampered by actions that can not be recognized, such as *zumba* and *situp*, likely due to the lack of relevant objects. Figure 8 divides the UCF-101 actions into three classes; person-object, person-person, and person-only, to analyse when semantic priors are effective and when not.

6.3 Combining Spatial and Semantic Priors

Based on the positive effect of the six individual spatial and semantic priors, we evaluate the impact on combining all priors for classification and localization. The results on UCF Sports are shown in Table 8. Naturally, spatial objects priors are leading for unseen action localization, since this is impossible with semantic priors only. The reverse holds for action classification, where semantic priors on global objects are leading. We do find that for both tasks, using a combination of all priors is best. We recommend to use a combination of the six object priors to best deal with unseen actions.

We show success and failure cases for unseen actions in Fig. 7. Adding the semantic priors on top of the spatial priors is especially beneficial when actions do not directly depend on an interacting object, see e.g. Fig. 7b. Since there is no relevant interacting object for the *diving* action, the corresponding tube relies solely on the person detection, resulting in a high overlap but with a low AP since the score is akin to non-*diving* tubes. Adding the scores from the semantic priors however, results in the highest diving score for the shown action tube over all other test tubes. Interestingly, the global objects from the semantic priors are ambiguous for the action, e.g., *diving suit*, but they still help for *diving*, as it is the only aquatic action.

6.4 Action Tube Retrieval

In the fourth experiment, we qualitative show the potential of our new task action tube retrieval. In this setting, users

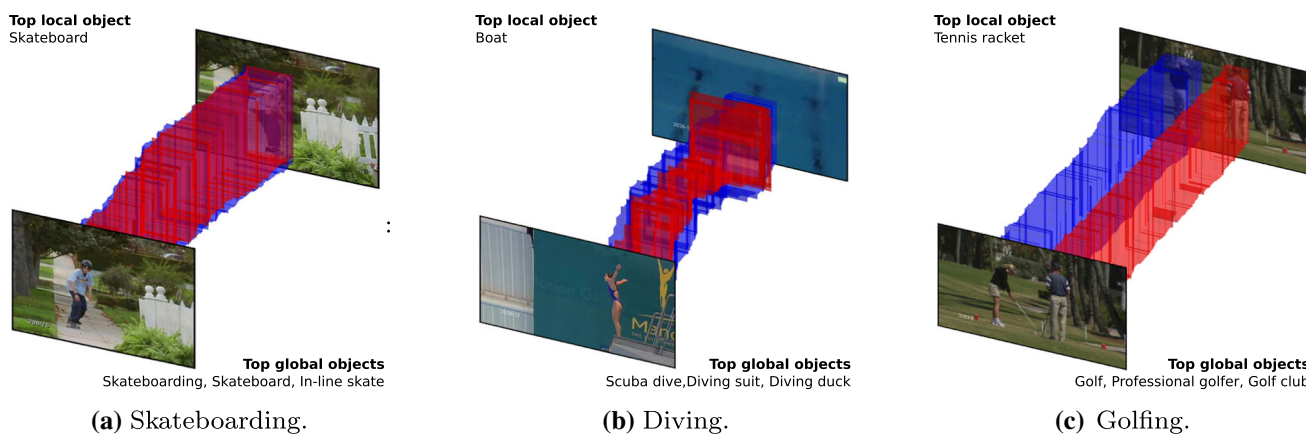


Fig. 7 Qualitative analysis on UCF Sports. For the video example of *skateboarding* we obtain a correct localization due to a clear match with relevant objects. The example of *golfing* obtains an incorrect localization. While the global objects are correct and relevant, the local object is

incorrect. Upon inspection, we found that this error was due to the limited vocabulary of the local objects; no golf-based objects are present in MS-COCO




Fig. 8 UCF-101 accuracies aggregated into three categories; person-object, person-person, and person-only. As expected our approach favors person-object interactions due to our object priors. Person-only actions, such as gymnastics and fitness actions, obtain lower scores, highlighting the importance of having relevant objects to recognize actions in our approach

query for desired objects, spatial prepositions, and optionally relative object size. In Fig. 9, we show three example queries along with top retrieved action locations.

6.5 Comparative Evaluation

In the fifth experiment, we compare our proposed approach to other works in action classification and localization without examples. For the classification comparison, we report on the UCF-101 dataset, since it is most used for this setting. For the localization comparison, we report on the other two datasets. For all comparisons, we use both spatial and semantic object priors.

Unseen action classification In Table 9, we show the unseen classification accuracies on UCF101 for three common

Table 8 Effect of combining spatial and semantic priors on the unseen action classification and localization results on UCF Sports

Object priors		Classification accuracy (%)	Localization mAP@0.5 (%)
spatial	semantic		
✓		29.8	27.0
	✓	59.6	–
✓	✓	68.1	34.9

Bold indicates best performance for each setting (relevant for the corresponding experiment)

For both tasks, combining semantic matching for global objects with spatial matching for local objects is beneficial

dataset splits using 101, 50, and 20 test classes. We first note the difference in scores with our conference version (Mettes and Snoek 2017), which are due to the three new semantic object priors. In the unseen setting, where no training actions are used, we are state-of-the-art. Moreover, we are competitive with zero-shot approaches that require extensive training on large-scale action datasets, such as Zhu et al. (2018) and Brattoli et al. (2020). Each approach employs different prior knowledge, making a direct comparison difficult. The comparison serves to highlight the overall effectiveness of our approach.

Unseen action localization In Table 10, we show the results for unseen action localization on UCF Sports and J-HMDB. The comparison is made to the only two previous papers with unseen localization results (Jain et al. (2015a); Mettes and Snoek 2017). On UCF Sports, we obtain an AUC score of 33.1%, compared to 7.2% for Jain et al. ((2015a)). We also outperform our previous work (Mettes and Snoek 2017), using spatial object priors only, by 2%, reiterating the empirical effect of semantic object priors. We furthermore provide mAP scores on both UCF Sports and J-HMDB. The larger

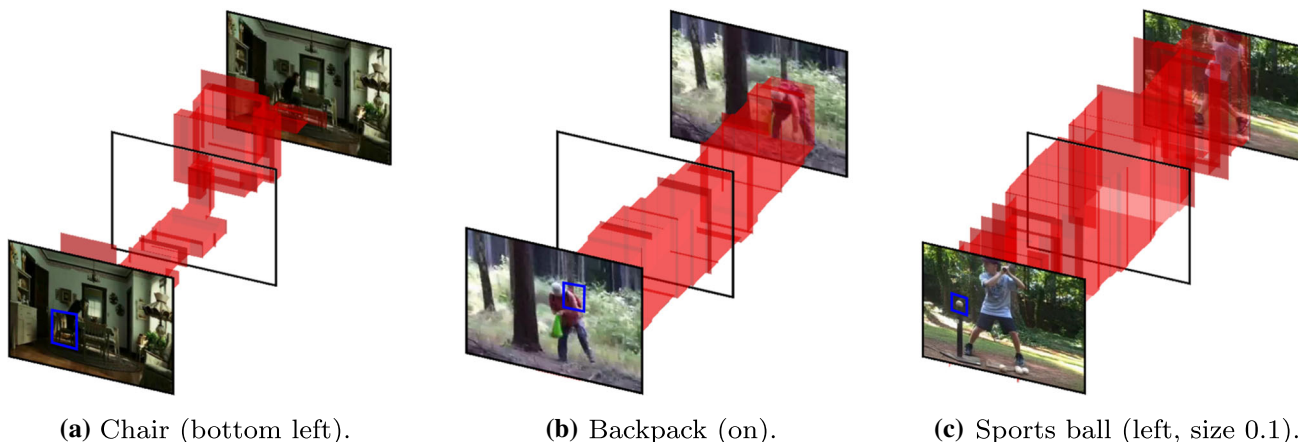


Fig. 9 Qualitative results for action tube retrieval on J-HMDB. The examples for *chair* and *backpack* show that our object embedding is capable of retrieving relevant action locations from user queries on the

fly. The example for *sports ball* shows that we can additionally request a preferred object size. In this example, a localization with a baseball is retrieved, since a small ball size was queried

Table 9 Comparison for unseen action classification accuracy (%) on UCF-101 for multiple numbers of test classes

	UCF-101		Accuracy
	Number of actions Train	Test	
Jain et al. ((2015a))	–	101	30.3
Mettes and Snoek (2017)	–	101	32.8
<i>This paper</i>	–	101	36.3
Zhu et al. (2018)	200	101	34.2
Brattoli et al. (2020)	664	101	37.6
Mettes and Snoek (2017)	–	50	40.4
<i>This paper</i>	–	50	47.3
An et al. (2019)	51	50	17.3
Bishay et al. (2019)	51	50	23.2
Mishra et al. (2020)	51	50	23.9
Mandal et al. (2019)	51	50	38.3
Zhu et al. (2018)	200	50	42.5
Brattoli et al. (2020)	664	50	48.0
Mettes and Snoek (2017)	–	20	51.2
<i>This paper</i>	–	20	61.1
Gan et al. ((2016b))	81	20	31.1
Bishay et al. (2019)	81	20	42.7
Zhu et al. (2018)	200	20	53.8

The train and test columns denote the number of action used for training and testing. Our approach is state-of-the-art in the unseen setting, where no training actions are used, and competitive to Zhu et al. (2018) and Brattoli et al. (2020), who require extensive training on ActivityNet and Kinetics respectively

gap in scores compared to the AUC metric on UCF Sports shows that we are now better at ranking correct action localizations at the top of the list for actions. Similarly for J-HMDB, we find consistent improvements across all overlap thresholds, highlighting our effectiveness for unseen action localization. We conclude that object priors matter for unseen

action classification and localization, resulting in state-of-the-art scores on both tasks.

Next to unseen action localization experiments on UCF Sports and J-HMDB, we also provide, for the first time, unseen localization on AVA. In Fig. 10, we show the frame AP for all 80 actions. We obtain a mean AP of 3.7%, compared to 0.7% for random scores with the same detected objects and persons. This result shows that large-scale multi-person action localization without training videos is feasible. Our zero-shot approach can identify contextual actions such as *play musical instrument* and *sail boat*, while it struggles with fine-grained actions that focus on person dynamics instead of object interaction, such as *crawl* and *fall down*.

The quantitative results on AVA show that large-scale unseen action localization is feasible, but multiple open challenges remain. In Fig. 11, we highlight three open challenges to improve localization performance. Most notably, it is unknown in the zero-shot setting how many actions occur at each timestep, while person-centric actions are often missed due to the lack of informative objects and context. Fine-grained actions (e.g., *listen to* versus *playing music*) are also difficult in dense scenes. Addressing these challenges require priors that go beyond objects, including but not limited to action priors and person skeleton priors.

7 Conclusions

This work advocates the importance of using priors obtained from objects to enable unseen action classification and localization. We propose three spatial object priors, allowing for spatio-temporal localization without examples. Additionally, we propose three semantic object priors to deal with semantic ambiguity, object discrimination, and object naming in the semantic matching. Even though no video examples are

Table 10 Unseen action localization comparisons on UCF Sports and J-HMDB using AUC and mAP across 5 overlap thresholds

	AUC					mAP				
	0.1	0.2	0.3	0.4	0.5	0.1	0.2	0.3	0.4	0.5
<i>UCF Sports</i>										
Jain et al. ((2015a))	38.8	23.2	16.2	9.9	7.2	–	–	–	–	–
Mettes and Snoek (2017)	43.5	39.3	37.1	35.7	31.1	47.4	43.5	42.1	32.0	23.2
<i>This paper</i>	47.3	43.0	40.7	37.9	33.1	61.2	54.2	54.0	41.5	34.9
<i>J-HMDB</i>										
Mettes and Snoek (2017)	34.6	33.3	30.5	26.8	23.0	27.5	27.0	23.2	19.2	15.1
<i>This paper</i>	37.3	37.1	33.9	31.0	26.7	32.1	31.5	27.2	22.6	17.6

Bold indicates best performance for each setting (relevant for the corresponding experiment) Across all settings, we obtain improved results, indicating the effectiveness of our approach




Fig. 10 Quantitative results of our approach with all six object priors on AVA. We show the frame AP over all classes on the validation videos. The mean AP over all classes is 3.7%, with notable high-performing actions that either involve clear interacting objects (*answer phone*, *player musical instrument*, and *sail boat*) or involve multiple people

that stand next to each other, in line with the spatial priors (*listen to* and *talk to a person*), highlighting that we can deal with multiple persons performing actions at the same time. Our approach struggles for single person actions, without any object interactions, such as *crawl* and *fall down*




Fig. 11 Challenges for unseen action localization with object priors in the wild on AVA keyframes (Gu et al. 2018). For each keyframe, we show the top three highest scoring actions (below frame) for the detected persons (red boxes), compared to the ground truth actions (above frame and blue boxes). In all three keyframes, at least one ground truth action is in our top actions due to relevant objects, resp. a phone in (a), a book

in (b), and an instrument in (c). The keyframes also show open challenges, e.g.,: it is unknown how many actions are relevant in a frame (a–c), person-centric actions are often missed (talk to in b and sit in c), and fine-grained actions can not be distinguished (listed to music versus playing instrument in c)

available during training, the object priors provide strong indications what actions happen where in videos. Due to the generic setup of our priors, we also introduce a new task, action tube retrieval, where users specify object type, spatial relations, and object size to obtain spatio-temporal locations on-the-fly. The use of spatial and semantic object priors results in state-of-the-art scores for unseen action classification and localization. We conclude that objects make sense for unseen actions when the set of actions is heterogeneous, as is the case in common action datasets. When actions become more fine-grained, *e.g.*, throwing versus catching a ball, spatial and semantic priors alone might not be sufficient, urging the need for causal temporal priors about objects and persons. For zero-shot interactions between persons, a fruitful source of priors to explore relate to knowledge about body pose.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Afouras, T., Owens, A., Chung, J. S., & Zisserman, A. (2020). Self-supervised learning of audio-visual objects from video. In *ECCV*.
- Alayrac, J. B., Recasens, A., Schneider, R., Arandjelović, R., Ramapuram, J., De Fauw, J., Smaira, L., Dieleman, S., & Zisserman, A. (2020). Self-supervised multimodal versatile networks. In *NeurIPS*.
- Alexiou, I., Xiang, T., & Gong, S. (2016). Exploring synonyms as context in zero-shot action recognition. In *ICIP*.
- An, R., Miao, Z., Li, Q., Xu, W., & Zhang, Q. (2019). Spatiotemporal visual-semantic embedding network for zero-shot action recognition. *Journal of Electronic Imaging*, 28(2), 93.
- Asano, Y. M., Patrick, M., Rupprecht, C., & Vedaldi, A. (2020). Labelling unlabelled videos from scratch with multi-modal self-supervision. In *NeurIPS*.
- Bishay, M., Zoumpourlis, G., & Patras, I. (2019). Tarn: Temporal attentive relation network for few-shot and zero-shot action recognition. In *BMVC*.
- de Boer, M., Schutte, K., & Kraaij, W. (2016). Knowledge based query expansion in complex multimedia event detection. *MTA*, 75(15), 9025–9043.
- Bond, F., & Foster, R. (2013). Linking and extending an open multilingual wordnet. In *Annual meeting of the Association for computational linguistics*.
- Brattoli, B., Tighe, J., Zhdanov, F., Perona, P., & Chalupka, K. (2020). Rethinking zero-shot video classification: End-to-end training for realistic applications. In *CVPR*.
- Carreira, J., & Zisserman, A. (2017). Quo vadis, action recognition? A new model and the kinetics dataset. In *CVPR*.
- Chakraborty, B., Holte, M. B., Moeslund, T. B., & González, J. (2012). Selective spatio-temporal interest points. *CVIU*, 116(3), 396–410.
- Chang, X., Yang, Y., Long, G., Zhang, C., & Hauptmann, A. G. (2016). Dynamic concept composition for zero-example event detection. In *AAAI*.
- Chéron, G., Alayrac, J. B., Laptev, I., & Schmid, C. (2018). A flexible model for training action localization with varying levels of supervision. In *NeurIPS*.
- Dalton, J., Allan, J., & Mirajkar, P. (2013). Zero-shot video retrieval using content and concepts. In *ICIKM*.
- Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., & Fei-Fei, L. (2009). Imagenet: A large-scale hierarchical image database. In *CVPR*.
- Escorcía, V., & Niebles, J. C. (2013). Spatio-temporal human-object interactions for action recognition in videos. In *ICCV workshops*.
- Feichtenhofer, C., Pinz, A., & Zisserman, A. (2016). Convolutional two-stream network fusion for video action recognition. In *CVPR*.
- Feichtenhofer, C., Fan, H., Malik, J., & He, K. (2019). Slowfast networks for video recognition. In *ICCV*.
- Fernando, B., Bilen, H., Gavves, E., & Gould, S. (2017). Self-supervised video representation learning with odd-one-out networks. In *CVPR*.
- Fu, Y., Hospedales, T. M., Xiang, T., & Gong, S. (2015). Transductive multi-view zero-shot learning. *IEEE TPAMI*, 37(11), 59.
- Gan, C., Lin, M., Yang, Y., de Melo, G., & Hauptmann, A. G. (2016a). Concepts not alone: Exploring pairwise relationships for zero-shot video activity recognition. In *AAAI*.
- Gan, C., Yang, T., & Gong, B. (2016b). Learning attributes equals multi-source domain generalization. In *CVPR*.
- Gan, C., Yang, Y., Zhu, L., Zhao, D., & Zhuang, Y. (2016c). Recognizing an action using its name: A knowledge-based approach. In *IJCV*.
- Gkioxari, G., & Malik, J. (2015). Finding action tubes. In *CVPR*.
- Grave, E., Bojanowski, P., Gupta, P., Joulin, A., & Mikolov, T. (2018). Learning word vectors for 157 languages. In *LREC*.
- Gu, C., Sun, C., Ross, D. A., Vondrick, C., Pantofaru, C., Li, Y., Vijayanarasimhan, S., Toderici, G., Ricco, S., Sukthankar, R., et al. (2018). Ava: A video dataset of spatio-temporally localized atomic visual actions. In *CVPR*.
- Gupta, A., & Davis, L. S. (2007). Objects in action: An approach for combining action understanding and object perception. In *CVPR*.
- Habibian, A., Mensink, T., & Snoek, C. G. M. (2017). Video2vec embeddings recognize events when examples are scarce. *TPAMI*, 39(10), 2089–2103.
- Han, T., Xie, W., & Zisserman, A. (2020). Memory-augmented dense predictive coding for video representation learning. In *ECCV*.
- Hou, R., Chen, C., & Shah, M. (2017). Tube convolutional neural network (t-cnn) for action detection in videos. In *ICCV*.
- Inoue, N., & Shinoda, K. (2016). Adaptation of word vectors using tree structure for visual semantics. In *MM*.
- Jain, M., Jegou, H., Bouthemy, & P. (2013). Better exploiting motion for better action recognition. In *CVPR*.
- Jain, M., van Gemert, J. C., Mensink, T., & Snoek, C. G. M. (2015a). Objects2action: Classifying and localizing actions without any video example. In *ICCV*.
- Jain, M., van Gemert, J. C., & Snoek, C. G. M. (2015b). What do 15,000 object categories tell us about classifying and localizing actions? In *CVPR*.
- Jain, M., van Gemert, J., Jégou, H., Bouthemy, P., & Snoek, C. G. M. (2017). Tubelets: Unsupervised action proposals from spatiotemporal super-voxels. *IJCV*, 124(3), 287–311.
- Jain, M., Ghodrati, A., & Snoek, C. G. M. (2020). Actionbytes: Learning from trimmed videos to localize actions. In *CVPR*.
- Jenni, S., Meishvili, G., & Favaro, P. (2020). Video representation learning by recognizing temporal transformations. In *ECCV*.
- Jhuang, H., Gall, J., Zuffi, S., Schmid, C., & Black, M. J. (2013). Towards understanding action recognition. In *ICCV*.

- Jing, L., & Tian, Y. (2020). Self-supervised visual feature learning with deep neural networks: A survey. In *TPAMI*.
- Jolicoeur, P., Gluck, M. A., & Kosslyn, S. M. (1984). Pictures and names: Making the connection. *Cognitive Psychology*, *16*(2), 243–275.
- Junior, V. L. E., Pedrini, H., & Menotti, D. (2019). Zero-shot action recognition in videos: A survey. In *CoRR*.
- Kalogeiton, V., Weinzaepfel, P., Ferrari, V., & Schmid, C. (2017a). Action tubelet detector for spatio-temporal action localization. In *ICCV*.
- Kalogeiton, V., Weinzaepfel, P., Ferrari, V., & Schmid, C. (2017b). Joint learning of object and action detectors. In *ICCV*.
- Kläser, A., Marszałek, M., Schmid, C., & Zisserman, A. (2010). Human focused action localization in video. In *ECCV*.
- Kodirov, E., Xiang, T., Fu, Z., & Gong, S. (2015). Unsupervised domain adaptation for zero-shot learning. In *ICCV*.
- Kuehne, H., Jhuang, H., Garrote, E., Poggio, T., & Serre, T. (2011). Hmdb: A large video database for human motion recognition. In *ICCV*.
- Kuehne, H., Arslan, A., & Serre, T. (2014). The language of actions: Recovering the syntax and semantics of goal-directed human activities. In *CVPR*.
- Lampert, C. H., Nickisch, H., & Harmeling, S. (2013). Attribute-based classification for zero-shot visual object categorization. *TPAMI*, *36*(3), 453–465.
- Lan, T., Wang, Y., & Mori, G. (2011). Discriminative figure-centric models for joint action localization and recognition. In *ICCV*.
- Laptev, I. (2005). On space-time interest points. *IJCV*, *64*(2–3), 107–123.
- Li, Y., Hu, Sh., & Li, B. (2016). Recognizing unseen actions in a domain-adapted embedding space. In *ICIP*.
- Li, Z., Yao, L., Chang, X., Zhan, K., Sun, J., & Zhang, H. (2019). Zero-shot event detection via event-adaptive concept relevance mining. *Pattern Recognition*, *3*, 91.
- Lin, T. Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., & Zitnick, C. L. (2014). Microsoft coco: Common objects in context. In *ECCV*.
- Liu, J., Ali, S., & Shah, M. (2008). Recognizing human actions using multiple features. In *CVPR*.
- Liu, J., Kuipers, B., & Savarese, S. (2011). Recognizing human actions by attributes. In *CVPR*.
- Malt, B. C. (1995). Category coherence in cross-cultural perspective. *Cognitive Psychology*, *29*(2), 85–148.
- Mandal, D., Narayan, S., Dwivedi, S. K., Gupta, V., Ahmed, S., Khan, F. S., & Shao, L. (2019). Out-of-distribution detection for generalized zero-shot action recognition. In *CVPR*.
- Mettes, P., & Snoek, C. G. M. (2017). Spatial-aware object embeddings for zero-shot localization and classification of actions. In *ICCV*.
- Mettes, P., & Snoek, C. G. M. (2019). Pointly-supervised action localization. *IJCV*, *127*(3), 263–281.
- Mettes, P., Koelma, D. C., & Snoek, C. G. M. (2020). Shuffled imagenet pre-training for video event detection and search. In *TOMM*.
- Mikolov, T., Sutskever, I., Chen, K., Corrado, G. S., & Dean, J. (2013). Distributed representations of words and phrases and their compositionality. In *NeurIPS*.
- Mishra, A., Verma, V. K., Reddy, M. S. K., Arulkumar, S., Rai, P., & Mittal, A. (2018). A generative approach to zero-shot and few-shot action recognition. In *WACV*.
- Mishra, A., Pandey, A., & Murthy, H. A. (2020). Zero-shot learning for action recognition using synthesized features. *Neurocomputing*, *2*, 13.
- Moore, D. J., Essa, I. A., & Hayes, M. H. (1999). Exploiting human actions and object context for recognition tasks. In *ICCV*.
- Murphy, G. (2004). *The big book of concepts*. London: MIT Press.
- Owens, A., & Efros, A. A. (2018). Audio-visual scene analysis with self-supervised multisensory features. In *ECCV*.
- Patrick, M., Asano, Y. M., Fong, R., Henriques, J. F., Zweig, G., & Vedaldi, A. (2020). Multi-modal self-supervision from generalized data transformations. arXiv.
- Pennington, J., Socher, R., & Manning, C. (2014). Glove: Global vectors for word representation. In *EMLNP*.
- Ren, S., He, K., Girshick, R., & Sun, J. (2015). Faster r-cnn: Towards real-time object detection with region proposal networks. In *NeurIPS*.
- Rodriguez, M. D., Ahmed, J., & Shah, M. (2008). Action MACH: A spatio-temporal maximum average correlation height filter for action recognition. In *CVPR*.
- Rosch, E. (1988). Principles of categorization. In E. E. Smith (Ed.), *Collins A* (pp. 312–322). Morgan Kaufmann: Readings in Cognitive Science.
- Rosch, E., Mervis, C. B., Gray, W. D., Johnson, D. M., & Boyes-Braem, P. (1976). Basic objects in natural categories. *Cognitive Psychology*, *8*(3), 382–439.
- Sener, F., & Yao, A. (2018). Unsupervised learning and segmentation of complex activities from video. In *CVPR*.
- Simonyan, K., & Zisserman, A. (2014). Two-stream convolutional networks for action recognition in videos. In *NeurIPS*.
- Soomro, K., & Shah, M. (2017). Unsupervised action discovery and localization in videos. In *ICCV*.
- Soomro, K., Zamir, A. R., & Shah, M. (2012). Ucf101: A dataset of 101 human actions classes from videos in the wild. arXiv:1212.0402.
- Szegedy, C., Liu, W., Jia, Y., Sermanet, P., Reed, S., Anguelov, D., Erhan, D., Vanhoucke, V., & Rabinovich, A. (2015). Going deeper with convolutions. In *CVPR*.
- Tian, Y., Ruan, Q., & An, G. (2018). Zero-shot action recognition via empirical maximum mean discrepancy. In *ICSP*.
- Tran, D., Wang, H., Torresani, L., & Feiszli, M. (2019). Video classification with channel-separated convolutional networks. In *ICCV*.
- Tschannen, M., Djolonga, J., Ritter, M., Mahendran, A., Houlsby, N., Gelly, S., & Lucic, M. (2020). Self-supervised learning of video-induced visual invariances. In *CVPR*.
- Wang, H., Kläser, A., Schmid, C., & Liu, C. L. (2013). Dense trajectories and motion boundary descriptors for action recognition. *IJCV*, *103*(1), 60–79.
- Wang, J., Jiao, J., Bao, L., He, S., Liu, Y., & Liu, W. (2019). Self-supervised spatio-temporal representation learning for videos by predicting motion and appearance statistics. In *CVPR*.
- Wang, Q., & Chen, K. (2017). Alternative semantic representations for zero-shot human action recognition. In *ECML*.
- Wu, J., Osuntogun, A., Choudhury, T., Philipose, M., & Rehg, J. M. (2007). A scalable approach to activity recognition based on object use. In *ICCV*.
- Wu, S., Bondugula, S., Luisier, F., Zhuang, X., & Natarajan, P. (2014). Zero-shot event detection using multi-modal fusion of weakly supervised concepts. In *CVPR*.
- Wu, Z., Fu, Y., Jiang, Y. G., & Sigal, L. (2016). Harnessing object and scene semantics for large-scale video understanding. In *CVPR*.
- Xu, D., Xiao, J., Zhao, Z., Shao, J., Xie, D., & Zhuang, Y. (2019). Self-supervised spatiotemporal learning via video clip order prediction. In *CVPR*.
- Xu, X., Hospedales, T., & Gong, S. (2016). Multi-task zero-shot action recognition with prioritised data augmentation. In *ECCV*.
- Xu, X., Hospedales, T., & Gong, S. (2017). Transductive zero-shot action recognition by word-vector embedding. *IJCV*, *123*(3), 309–333.
- Yao, B., Khosla, A., & Fei-Fei, L. (2011). Classifying actions and measuring action similarity by modeling the mutual context of objects and human poses. In *ICML*.
- Zhang, L., Chang, X., Liu, J., Luo, M., Wang, S., Ge, Z., & Hauptmann, A. (2020). Zstad: Zero-shot temporal activity detection. In *CVPR*.

- Zhang, Z., Wang, C., Xiao, B., Zhou, W., & Liu, S. (2015). Robust relative attributes for human action recognition. *PAA*, *18*(1), 157–171.
- Zhao, J., & Snoek, C.G.M. (2019). Dance with flow: Two-in-one stream action detection. In *CVPR*.
- Zhao, Y., Xiong, Y., & Lin, D. (2018). Trajectory convolution for action recognition. In *NeurIPS*.
- Zhu, L., Xu, Z., Yang, Y., & Hauptmann, A. G. (2017). Uncovering the temporal context for video question answering. *IJCV*, *124*(3), 409–421.
- Zhu, Y., Long, Y., Guan, Y., Newsam, S., & Shao, L. (2018). Towards universal representation for unseen action recognition. In *CVPR*.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.