

AFF-ttention! Affordances and Attention models for Short-Term Object Interaction Anticipation

Lorenzo Mur-Labadia¹, Ruben Martinez-Cantin¹, Jose J. Guerrero¹,
Giovanni Maria Farinella², and Antonino Furnari²

¹ University of Zaragoza, Spain

{lmur, rmcantin, jguerrer}@unizar.es

² University of Catania, Italy

{giovanni.farinella, antonino.furnari}@unict.it

Abstract. Short-Term object-interaction Anticipation (STA) consists of detecting the location of the next-active objects, the noun and verb categories of the interaction, and the time to contact from the observation of egocentric video. This ability is fundamental for wearable assistants or human-robot interaction to understand the user’s goals, but there is still room for improvement to perform STA in a precise and reliable way. In this work, we improve the performance of STA predictions with two contributions: 1) We propose STAformer, a novel attention-based architecture integrating frame-guided temporal pooling, dual image-video attention, and multiscale feature fusion to support STA predictions from an image-input video pair; 2) We introduce two novel modules to ground STA predictions on human behavior by modeling affordances.

First, we integrate an environment affordance model which acts as a persistent memory of interactions that can take place in a given physical scene. Second, we predict interaction hotspots from the observation of hands and object trajectories, increasing confidence in STA predictions localized around the hotspot. Our results show significant relative Overall Top-5 mAP improvements of up to +45% on Ego4D and +42% on a novel set of curated EPIC-Kitchens STA labels. We will release the code, annotations, and pre-extracted affordances on Ego4D and EPIC-Kitchens to encourage future research in this area.

Keywords: Short-term forecasting · Affordances · Egocentric video understanding

1 Introduction

Anticipating the future is a fundamental ability for assistive egocentric devices and to support human-robot interaction. For example, a smart wearable device could alert an electrical operator before they short-circuit a switchboard, or a home robot can support the human by turning on appliances or moving objects according to their forecasted long-term goal. Predicting the future state of the scene from egocentric visual observations is a growing research area [39, 45],

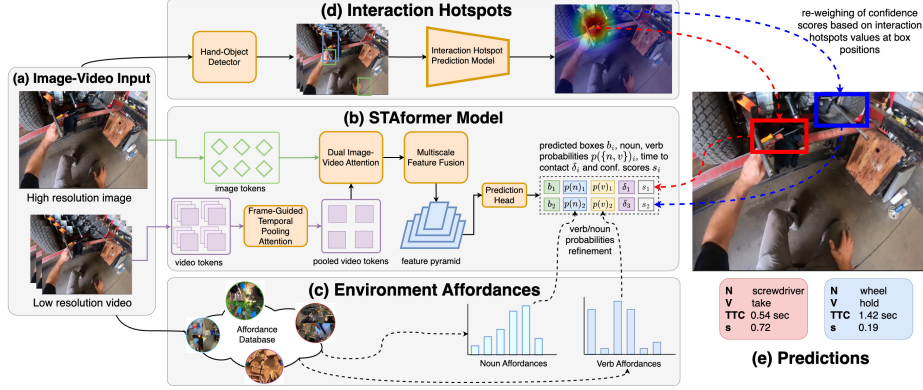


Fig. 1: (a) Our approach takes as input an image-video pair. (b) The input is processed by the proposed STAformer model which predicts object bounding boxes, the associated verb/noun probabilities, time-to-contact estimates and confidence scores. (c) Environment affordances are inferred from video and used to refine the predicted noun/verb probabilities. (d) Our model observes detected hand-object interactions in the video and predicts an interaction hotspot probability map, which is used to re-weight confidence scores based on box locations, leading to (e) our final predictions.

with works tackling action anticipation [5, 12, 14, 34, 46, 54, 55], locomotion prediction [3, 19, 21, 27, 37], hands trajectory forecasting [1, 24, 25], and next-active object detection [7, 11, 18, 43]. Recently, Grauman et al. [17] defined the Short-Term Object Interaction Anticipation (STA) task as the simultaneous prediction of the action and object category, the object’s bounding box, and the time to contact, and introduced an international challenge within the forecasting benchmark of the Ego4D dataset. Inspired by this challenge, the community proposed different approaches [4, 38, 42, 49–52]. Despite the progress in the area, our results show a large advantage over previous results, which highlights the room for improvement in accuracy and robustness.

Our aim with this work is to advance research in STA with two main contributions. First, we propose STAformer, a principled architecture unifying the computation of image and video inputs with attention-based components (Figure 1(a)-(b)). Differently from previous approaches [17, 38, 50], we explicitly designed STAformer to operate on an image-video input pair, which is specific to the considered STA task. Our architecture is a significant departure from convolutional baselines [17, 42] and aims to offer the convenience and state-of-the-art performance of attention-based feature extractors [2, 36] and components [53]. Second, to tackle the challenges associated with relating past visual observations to future events from video, we propose two effective modules designed to ground predictions into human behavior by modeling affordances. As highlighted in recent studies [40], human activities exhibit consistency in similar environments. Hence, we first leverage environment affordances [32], estimated by matching the input observation to a learned affordance database, to predict probability

distributions over nouns and verbs, which are used to refine verb and noun probabilities predicted by STAformer (Figure 1(c)). Our intuition is that linking a zone across similar environments captures a description of the feasible interactions, grounding predictions into previously observed human behavior. The second affordance module aims to relate STA predictions to a spatial prior of where an interaction may take place in the current frame. This is done by predicting an interaction hotspot [25], which is used to re-weight confidence scores of STA predictions depending on the object’s locations (Figure 1(e)).

Experiments on Ego4D [17] and a novel set of curated STA annotations on the EPIC-Kitchens dataset [6] highlight the effectiveness of the proposed approach, which obtains significant relative improvements of +45% on the validation set of Ego4D v1, +42.1% on the validation set of Ego4D v2, +30.3% on the private test set of Ego4D, and +42% on EPIC-Kitchens, when measured with the official overall Top-5 mAP evaluation measure. The proposed approach currently scores first on the Ego4D Short-Term object-interaction Anticipation leaderboard.³ Experiments also highlight the individual contributions of STAformer and the proposed modules to exploit affordances for STA prediction. In sum, the contributions of our work are as follows: 1) We introduce STAformer, a novel attention-based architecture specifically designed to process an input image-video pair, which achieves state-of-the-art performance on the two challenging Ego4D and EPIC-Kitchens benchmarks. 2) We propose two modules to ground STA predictions to human behavior by modeling environment affordances and interaction hotspots. The two modules are shown to be effective when coupled with STAformer, as well as previous architectures, which highlights the general usefulness of the approach. 3) We contribute a novel set of STA annotations, curated from public EPIC-Kitchens labels. This effectively provides the research community with a second large-scale and challenging benchmark for the STA task, besides the popular Ego4D. We will publicly release the open-source implementation of STAformer and the affordance modules, the proposed affordance databases pre-computed on Ego4D and EPIC-Kitchens, and the novel EPIC-Kitchens STA annotations.

2 Related works

Short-term Object Interaction Anticipation: Among seminal works, Furnari et al. [11] initially introduced the concept of Next-Active Objects (NAO), proposing to detect future interacted objects by analyzing their trajectories as observed from the first-person point of view. Differently from action anticipation [6], the NAO detection task is designed to provide grounded predictions in the form of bounding boxes, which can be particularly informative for wearable AI assistants or embodied robotic agents. Unlike traditional object detection [15], NAO prediction requires the ability to model the dynamics of the scene and anticipate the user’s intention. Jiang et al. [18] developed a method to predict the next-active

³ See the supplementary material for more details.

object location in the form of a Gaussian heatmap from a single RGB image, combining visual attention with probabilistic maps of hand locations. Ego-OMG [7] segments the NAO and predicts the interaction time using a contact anticipation map that captures scene dynamics. While previous works considered different task formulations and evaluation approaches, Grauman et al. [17] formalized NAO prediction by introducing the STA task and an associated challenge on the EGO4D dataset [17]. The initial baseline is composed of a Faster R-CNN branch to detect objects [15] and a SlowFast 3D CNN [10] for video processing. Subsequent research introduced architectural enhancements and alternative approaches. Chen et al. [4] employed pre-computed object detections using a DETR model and substituted SlowFast with a VideoMAE pre-trained ViT [52]. Pasca et al. [38] proposed TransFusion, which employs a language encoder for action context summary, performing multi-modal fusion with visual features. While previous works leveraged pre-extracted object detections for 2D image understanding, Ragusa et al. [42] introduced StillFast, an end-to-end framework unifying the processing of 2D images and video in a combined backbone. Thakur et al. [49] proposed GANO, an end-to-end model based on a transformer architecture including a novel guided attention mechanism. Guided attention was integrated within a StillFast architecture in [50], achieving state-of-the-art results. Thakur et al. [51] introduced NAOGAT, a multi-modal transformer that attends detected objects and includes a motion decoder to track object trajectories. Despite the progress, previous works show incremental results over the baselines, indicating significant potential improvement.

Affordances for Anticipation: Defined by Gibson [13], affordances are *the potential actions that the environment offers to the agent*. The computational perception of affordances has been investigated in different forms. A line of works predicts affordance labels of object parts, requiring strong supervision in the form of manually annotated masks [8, 30, 31, 35]. However, these methods are not “grounded” in human behavior as the annotator declares interaction regions outside of any interaction context [32]. Other works considered the problem of grounding affordance regions in images by leveraging videos depicting human-object interactions in a weakly supervised way, where only the action label is used as supervision without spatial annotations [16, 22, 26, 32]. Nagarajan et al. [32] introduced the concept of “interaction hotspots” as the potential spatial regions where the action can occur. Mur-Labadia et al. [29] create a 3D multi-label mapping of affordances extracted from egocentric video. Another line of work infers interaction hotspots from video by forecasting future hand movements to select candidate regions for future interactions [16, 18, 24, 25]. Few works studied scene affordances to predict a list of likely actions that can be performed in a given scene [33, 44]. In particular, Nagarajan et al. [33] proposed EGO-TOPO, a procedure to decompose a set of egocentric videos into a topological map encoding scene affordances. Despite the interest in affordances, only a few works investigated how to exploit them for future predictions. Montesano et al. [28] predicted affordance effects for human-robot interaction. Koppula et al. [20] used object affordances to anticipate human behavior in the form of motion trajectories of

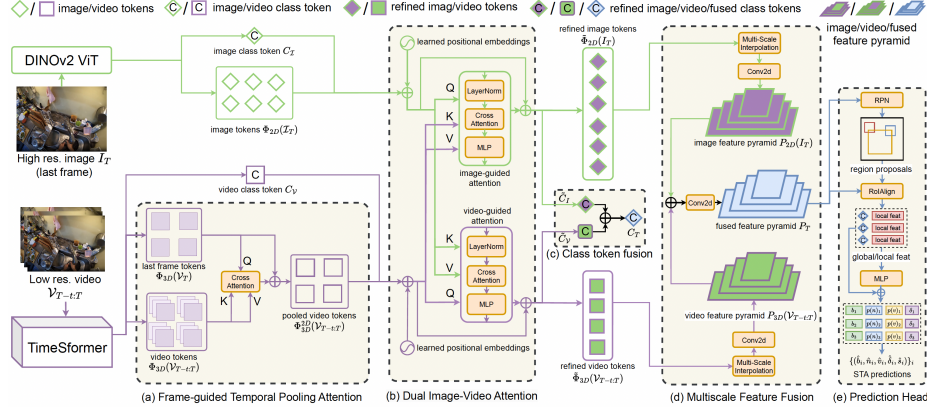


Fig. 2: STAformer architecture. DINO-v2 and TimeSformer extract 2D and 3D features from the image-video input. (a) Frame-guided temporal pooling attention spatially aligns video to image features. (b) Dual image-video attention enriches 2D features with temporal dynamics and 3D features with fine-grained image details. Image and video representations are joined to obtain a global class token (c) and a feature pyramid (d), from which we obtain the STA predictions (e).

objects and humans. Nagarajan et al. [33] showed how scene affordances learned from egocentric video can improve long-term action anticipation. Liu et al. [24] tackled action anticipation by jointly predicting egocentric hand motion, interaction hotspots, and future actions. Liu et al. [25] highlighted how interaction hotspots predicted by forecasting hand motion can support action anticipation.

3 STAformer Architecture for Short-Term Anticipation

As defined in [17], the goal of Short-Term object interaction Anticipation (STA) is to detect the Next-Active Object (NAO) from the observation of a given input video $\mathcal{V}_{:T}$ up to timestamp T . The model prediction are a set of detections, defined by the tuple $(b_i, n_i, v_i, \delta_i, s_i)$, denoting future interacted objects in the last observed frame I_T . Each bounding box b_i is associated with an object category label n_i (noun), a verb label indicating the interaction mode v_i , a time-to-contact δ_i indicating that the interaction will take place at time $T + \delta_i$, and a confidence score s_i . We propose STAformer, a novel architecture that leverages pre-trained transformer models for image and video feature extraction [36, 41] and introduces novel attention-based components for image-video representation fusion. We illustrate the architecture in Figure 2 and discuss it in the following.

Feature Extraction: We follow previous work [17, 42] and process a high resolution image $I_T \in \mathbb{R}^{h_s \times w_s \times 3}$ sampled from the input video $\mathcal{V}_{:T}$ at time T and a sequence of low-resolution frames $\mathcal{V}_{T-t:T} \in \mathbb{R}^{t \times h_f \times w_f \times 3}$ taken t time-steps before time T . First, we extract high-resolution 2D features from the I_T input with a DINOv2 model [36], obtaining a set of 2D image tokens $\Phi_{2D}(I_T)$ and a class

token C_I offering a global representation of the image. We also extract spatio-temporal 3D features from the $\mathcal{V}_{T-t:T}$ input with a TimeSformer model [2] in the form of video tokens $\Phi_{3D}(\mathcal{V}_T)$ and a class token C_V giving a global representation of the input clip.

Frame-guided Temporal Pooling Attention (Figure 2(a)): While the overall video tokens provide a spatio-temporal representation of the input video, STA predictions need to be aligned to the spatial location of the last video frame. The frame-guided temporal pooling attention maps video tokens to the spatial reference system of the last video frame, compressing the 3D representation obtained by the TimeSformer to a 2D one. The 3D video tokens $\Phi_{3D}(\mathcal{V}_{T-t:T})$ are mapped to 2D pooled video tokens denoted as $\Phi_{3D}^{2D}(\mathcal{V}_{T-t:T})$ adopting a residual cross-attention mechanism. Specifically, we compute query vectors from last-frame video tokens $\Phi_{3D}(\mathcal{V}_T)$ with a linear projection W_Q , while key and value vectors are computed from the overall video tokens $\Phi_{3D}(\mathcal{V}_{T-t:T})$ using the W_K and W_V linear projection layers. We obtain pooled video tokens with a residual multi-head attention (A) layer as follows:

$$\Phi_{3D}^{2D}(\mathcal{V}_{T-t:T}) = \Phi_{3D}(\mathcal{V}_T) + A\left(\underbrace{\Phi_{3D}(\mathcal{V}_T)W_Q}_{\text{queries}}, \underbrace{\Phi_{3D}(\mathcal{V}_{T-t:T})W_K}_{\text{keys}}, \underbrace{\Phi_{3D}(\mathcal{V}_{T-t:T})W_V}_{\text{values}}\right) \quad (1)$$

Used as queries, last-frame tokens guide an adaptive temporal pooling that summarizes the spatio-temporal feature map computed by the TimeSformer model and maps it to the 2D reference space of the last observed frame. The residual connection facilitates learning and lets the attention mechanism focus on enriching last-frame tokens with video tokens.

Dual Image-Video Attention fusion (Figure 2(b)): Image tokens $\Phi_{2D}(I_T)$ and pooled video tokens $\Phi_{3D}^{2D}(\mathcal{V}_{T-t:T})$ are spatially aligned, but carry different information, with image tokens encoding fine-grained visual features and video tokens encoding scene dynamics. This module adopts a residual dual cross-attention that aims to enrich image tokens with scene dynamics information coming from video tokens through image-guided cross-attention and, vice versa, video tokens with fine-grained visual information coming from image tokens through video-guided cross-attention. Prior to forwarding image and video tokens to the multi-head cross-attention modules, these are summed with learnable positional embeddings to capture insightful spatial relationships and normalized through a Layer Norm. The residual image-guided cross-attention is as follows:

$$\begin{aligned} [\tilde{\Phi}_{2D}(I_T), \tilde{C}_I] &= [\Phi_{2D}(I_T), C_I] + \\ &A\left(\underbrace{[\Phi_{2D}(I_T), C_I]W_Q}_{\text{queries}}, \underbrace{[\Phi_{3D}^{2D}(\mathcal{V}_{T-t:T}), C_V]W_K}_{\text{keys}}, \underbrace{[\Phi_{2D}^{3D}(\mathcal{V}_{T-t:T}), C_V]W_V}_{\text{values}}\right) \end{aligned} \quad (2)$$

where $[\cdot, \cdot]$ denotes concatenation along batch dimension, and W_Q , W_K , and W_V are linear projection layers. After the multi-head attention layer, the refined image representation $[\tilde{\Phi}_{2D}(I_T), \tilde{C}_I]$ is passed through a residual MLP. The video-guided cross-attention works in a similar way to compute refined video tokens

$\tilde{\Phi}_{3D}(\mathcal{V}_{T-t:T})$ and video class tokens \tilde{C}_V , but queries are computed from video tokens while keys and values are computed from image tokens.

Feature Fusion and prediction head (Figure 2(c)-(e)): Refined image and video class tokens are summed to obtain the overall class token $C_T = \tilde{C}_I + \tilde{C}_V$, a global representation of the input image-video pair (Figure 2(c)). Refined image tokens $\tilde{\Phi}_{2D}(I_T)$ are mapped to a multi-scale feature pyramid [23] $P_{2D}(I_T)$ by rescaling $\tilde{\Phi}_{2D}(I_T)$ to multiple resolutions using bilinear interpolation⁴, followed by a 3×3 convolution to compensate for interpolation artifacts. Refined video tokens $\tilde{\Phi}_{3D}(\mathcal{V}_{T-t:T})$ are mapped to a feature pyramid $P_{3D}(\mathcal{V}_{T-t:T})$ in the same way. The two feature pyramids are summed and passed through a 2D 3×3 convolution to obtain the fused feature pyramid P_T (Figure 2(d)). We adopt the prediction head⁴ proposed in [42] to obtain the final predictions $(\hat{b}_i, \hat{n}_i, \hat{v}_i, \hat{\delta}_i, \hat{s}_i)$. It is a modified version of the Faster-RCNN [15] that integrates specialized components for STA prediction. Note that while [42] uses global average pooling to obtain a global representation of the scene, we naturally use the class token C_T learned from the input image-video pair.

4 Leveraging affordances for human behavior grounding

While end-to-end STA architectures predict future human-object interactions from labeled data, in this section we show that it is beneficial to explicitly incorporate environment affordances based on linking functionally similar regions and interaction hotspots obtained from hand trajectories.

4.1 Leveraging environment affordances

Environment affordances [33] refer to all potential interactions that can be performed in a given physical zone. Our intuition is that a robust representation of environment affordances, learned from the observation of human activities in egocentric video, encapsulates the interaction that the user is going to perform next. We first build an affordance database grouping the training videos according to their visual similarity in activity-centric zones. At inference time, we match a novel video \mathcal{V}' to the most functionally similar zones in the affordance database, estimating the distribution of the affordable interactions in the new video. We use the nouns and verbs affordance distributions to refine the respective nouns and verbs probabilities predicted by STAformer.

Building the affordance database: We start extracting activity-centric zones from the training set following [33]. We build positive and negative frame pairs labels by counting homography estimation inliers, evaluating temporal coherence, and computing visual similarity with a pre-trained ResNet-152. A Siamese network \mathbb{L} is then trained on these pairs and used to predict the probability $\mathbb{L}(I, I')$ that two frames I and I' belong to the same zone. We then process all frames in a video sequence with \mathbb{L} to group video frames according to their

⁴ See the supplementary material for more details.

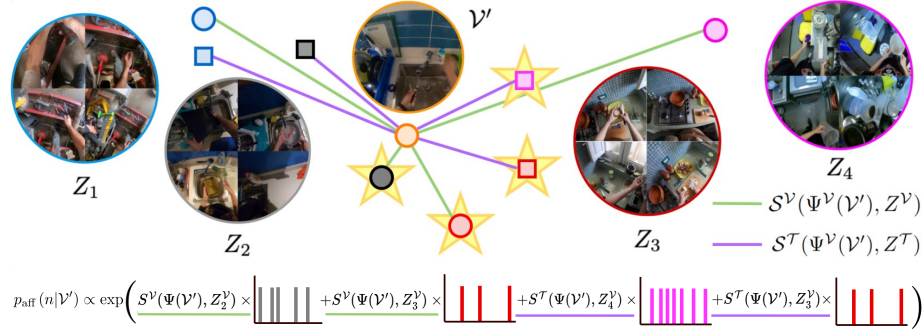


Fig. 3: Cross-environment inference of affordances: The input video \mathcal{V}' is matched to the affordance database comparing its visual representation $\Psi^{\mathcal{V}}(\mathcal{V}')$ to the visual $Z^{\mathcal{V}}$ (○) and text $Z^{\mathcal{T}}$ (□) zone descriptors. The affordance noun probability $p_{\text{aff}}(n|\mathcal{V}')$ is obtained by weighting the counts of nouns present in the top-2K nearest zones (★) according to the respective similarity \mathcal{S} . Example for $K=2$.

visual similarity in different zones.⁴ Each zone Z represents an activity-centric region composed of the group of visually similar images I_i^Z , their corresponding videos \mathcal{V}_i^Z , the associated narrations \mathcal{T}_i^Z , sets of nouns \mathcal{N}_i^Z and action verbs \mathcal{A}_i^Z appearing at least once in the STA annotations of all images I_i^Z . This represents a sort of *persistent memory* on how humans behave in each different environment. We obtain a visual descriptor $Z^{\mathcal{V}}$ and a text descriptor $Z^{\mathcal{T}}$ for each zone Z computing the average descriptors of videos within each zone: $Z^{\mathcal{V}} = \sum_{i=1}^{|Z|} \Psi_{\mathcal{V}}(\mathcal{V}_i^Z)/|Z|$, $Z^{\mathcal{T}} = \sum_{i=1}^{|Z|} \Psi_{\mathcal{T}}(\mathcal{T}_i^Z)/|Z|$. We adopt the dual encoder of EgoVLP-v2 [41] to extract video $\Psi^{\mathcal{V}}(\mathcal{V}_i^Z)$ and text $\Psi^{\mathcal{T}}(\mathcal{T}_i^Z)$ descriptors.

Inferring environment affordances: While [33] links similar activity-centric zones and trains a neural network to predict affordances directly from video, we found this approach suboptimal in our settings as we discuss in the results section. Instead, at inference time, we predict the nouns and verbs affordance distribution by matching a novel video \mathcal{V}' to zones related to functionally similar environments in the affordance database. Since we can only extract a visual descriptor from the novel video, $\Psi^{\mathcal{V}}(\mathcal{V}')$, we compute the visual cosine similarity $\mathcal{S}^{\mathcal{V}}(\Psi^{\mathcal{V}}(\mathcal{V}'), Z^{\mathcal{V}})$ and the video-text cross cosine similarity $\mathcal{S}^{\mathcal{T}}(\Psi^{\mathcal{V}}(\mathcal{V}'), Z^{\mathcal{T}})$ between the clip and each zone Z in the database. Despite being visually dissimilar, the cross distance relates different locations with similar functionality that affords the same interaction (i.e, painting a wall in India or painting a canvas with watercolor in Spain both afford to dip the brush in the paint).

As illustrated in Figure 3, we employ the K-Nearest Neighbour algorithm to identify the most similar zones to the given input \mathcal{V}' . We define the top-K visual zones $\mathcal{K}^{\mathcal{V}} = \{(Z_1^{\mathcal{V}}, S_1^{\mathcal{V}}), \dots, (Z_K^{\mathcal{V}}, S_K^{\mathcal{V}})\}$, where $S_k^{\mathcal{V}}$ is a shorthand notation for $\mathcal{S}^{\mathcal{V}}(\Psi^{\mathcal{V}}(\mathcal{V}'), Z_k^{\mathcal{V}})$, and the top-K narrative zones $\mathcal{K}^{\mathcal{T}} = \{(Z_1^{\mathcal{T}}, S_1^{\mathcal{T}}), \dots, (Z_K^{\mathcal{T}}, S_K^{\mathcal{T}})\}$. Combining both sets, $\mathcal{K} = \mathcal{K}^{\mathcal{V}} \cup \mathcal{K}^{\mathcal{T}} = \{(Z_i, S_i)\}_{i=1}^{2K}$ yields a total of $2K$ zones and their respective similarity scores, which we assume to share affordances with

\mathcal{V}' . We then define the probability of each noun $p_{\text{aff}}(n|\mathcal{V}')$ as an exponential distribution by weighting the noun appearance in each neighbouring zone \mathcal{N}^{Z_i} according to the respective similarity S_i :

$$p_{\text{aff}}(n|\mathcal{V}') \propto \exp\left(\sum_{(Z_i, S_i) \in \mathcal{K}} S_i \cdot \mathbb{1}_{n \in \mathcal{N}^{Z_i}}\right) \quad (3)$$

We apply the same procedure to predict the verb distribution $p_{\text{aff}}(v|\mathcal{V}')$.

STA predictions and environment affordances data fusion: Based on the environment affordances, we can predict probability distributions over *possible* nouns $p_{\text{aff}}(n|\mathcal{V})$ or verbs $p_{\text{aff}}(v|\mathcal{V}')$ given *past interactions in functionally similar zones*. Differently, the STA model will predict probability distributions of given nouns and verbs *being the next interactions* $p_{\text{sta}}(n|\mathcal{V}', I')$ and $p_{\text{sta}}(v|\mathcal{V}', I')$ directly from the input image-video pair, without explicitly considering the set of possible actions. We assume independence between the two predictions⁵ and perform data fusion by computing the unnormalized joint likelihoods:

$$\begin{aligned} p_{\text{fus}}(n|I', \mathcal{V}') &\propto p_{\text{aff}}(n|\mathcal{V}') \cdot p_{\text{sta}}(n|\mathcal{V}', I') \\ p_{\text{fus}}(v|I', \mathcal{V}') &\propto p_{\text{aff}}(v|\mathcal{V}') \cdot p_{\text{sta}}(v|\mathcal{V}', I') \end{aligned} \quad (4)$$

4.2 Leveraging interaction hotspots:

While our affordance database gives us information on which objects (nouns) and interaction modes (verbs) are likely to appear in the current scene, it does not give us any information on *where* the interaction will take place in the observed images. As noted in previous works [24, 25], observing how hands move in egocentric videos can allow us to predict the interaction hotspot [25, 33], a distribution over image regions indicating possible future interactions locations. We exploit this concept and include a module to predict an interaction hotspot by observing frames, hands, and objects. As Figure 4 illustrates, we hence re-weight the confidence scores s_i of STA predictions according to the location of the respective bounding box centers in the predicted interaction hotspot, to reduce the influence of false positive detections falling in areas of unlikely interaction.

Inferring interaction hotspots We base our interaction hotspot module on the work presented in [25] with some improvements. First, we fine-tune the hand object detector presented in [47] on EGO4D-SCOD [17] annotations, rather than using it out-of-the-box. Second, we extract stronger egocentric-aware frame features with the video part of the dual-encoder version of EgoVLP [41] pre-trained on Ego4D [41], instead of using a ConvNet as in [25].⁶ The model takes as inputs the features of the observed frames, besides the coordinates and features of both hands and pre-detected objects, and is trained to forecast the hand trajectory,

⁵ In practice, we build the two models with different architectures and training objectives to make the dependence weak.

⁶ See supp. for more information on the interaction hotspot prediction module.




Fig. 4: Refinement of confidence scores based on the interaction hotspots. The interaction hotspot model observes frames, hands, and objects and forecasts a map encoding the probability of the interaction in each pixel. STA confidence scores are re-weighted based on the probability values at the bounding box coordinate centers, reducing confidence in false positive predictions falling far from the interaction hotspot.

from which it predicts a distribution over plausible future contact points. Given the observed image-video pair $(I_T, \mathcal{V}_{T-t:T})$, the output of the model is a probability distribution over the spatial locations of I_T indicating the probability of interaction of each pixel denoted as $p_{ih}(x, y | I_T, \mathcal{V}_{T-t:T})$.

Fusing STA predictions with interaction hotspots: We exploit the interaction hotspots to refine the predictions of the STA model, assuming that regions close to the predicted interaction hotspots are more likely to contain the next active objects. Given a predicted box \hat{b}_i , we re-weight its related confidence score \hat{s}_i according to the location of the bounding box center $(\hat{c}_i^x, \hat{c}_i^y)$ in the interaction hotspot as following: $\hat{s}_i \cdot p_{ih}(\hat{c}_i^x, \hat{c}_i^y | I_T, \mathcal{V}_{T-t:T})$.

5 Results

We evaluate our approach on two challenging benchmarks: Ego4D [17] and EPIC-Kitchens [6]. Since no STA labels are available for EPIC-Kitchens, we extend and publicly release annotations to contribute a new benchmark for STA. We compare our model against different STA methods that either provide open-source implementations or report results in their papers [4, 17, 38, 42, 48, 50]. We adopt standard Noun (N), Noun+Verb (N+V), Noun+time-to-contact (N+ δ) and Noun+Verb+time-to-contact (All) Top-5 mean Average Precision (mAP).

5.1 Comparison with the state-of-the-art

Tables 1-2 report the results on the validation sets of Ego4D v1 and v2. Our method outperforms all previous approaches by wide margins, showing relative gains⁷ of up to +45.0% and +42.1% on v1 and v2 respectively when considering the mAP All measure. The significant improvements both in semantic, spatial, and temporal reasoning confirm the benefits of our two main contributions:

⁷ We compute the relative gain% of x concerning y as $100 \cdot (\frac{x-y}{y})$.

Table 1: Results in mAP on the validation split of Ego4D-STA v1. **Best results** in bold. Relative gain is with respect to second best

Model	N	N + V	N + δ	All
FRCNN+SF [17]	17.55	5.19	5.37	2.07
FRCNN+Feat. [48]	<u>22.01</u>	5.52	5.54	1.78
StillFast [42]	16.21	7.47	4.94	2.48
Transfusion [38]	20.19	<u>7.55</u>	<u>6.17</u>	2.60
STAformer	21.71	10.75	7.24	3.53
STAformer + AFF	24.36	12.00	7.66	3.77
Gain (rel %)	<u>+10.6</u>	<u>+58.9</u>	<u>+24.2</u>	<u>+45.0</u>

Table 2: Results in mAP on the validation split of Ego4D-STA v2. **Best results** in bold. Relative gain is with respect to second best.

Model	N	N + V	N + δ	All
FRCNN+SF [17]	21.00	7.45	7.07	2.98
InternVideo [4]	19.45	8.00	6.97	3.25
StillFast [42]	20.26	10.37	7.26	3.96
GANO v2 [50]	<u>20.52</u>	<u>10.42</u>	<u>7.28</u>	3.99
STAformer	24.85	13.45	7.41	4.90
STAformer+AFF	27.03	14.36	8.72	5.04
STAformer+MH	27.51	14.68	9.63	5.50
STAformer+MH+AFF	29.39	15.38	9.94	5.67
Gain (rel %)	<u>+43.3</u>	<u>+47.6</u>	<u>+36.5</u>	<u>+42.1</u>

Table 3: Results in mAP on the test split of Ego4D-STA. T denotes training data.

Model	T	N	N + V	N + δ	All
FRCNN+SF. [17]	v1	20.45	6.78	6.17	2.45
FRCNN+Feat. [48]	v1	20.45	4.81	4.40	1.31
InternVideo [4]	v1	24.60	9.18	<u>7.64</u>	3.40
Transfusion [38]	v1	<u>24.69</u>	<u>9.97</u>	7.33	3.44
StillFast [42]	v1	19.51	9.95	6.45	<u>3.49</u>
STAformer	v1	24.39	12.49	7.54	4.03
STAformer + AFF	v1	26.52	13.15	7.78	4.06
Gain (rel %)	v1	<u>+7.4</u>	<u>+32.1</u>	<u>+1.8</u>	<u>+13.1</u>
StillFast [42]	v2	25.06	13.29	9.14	5.12
GANO v2 [50]	v2	25.67	<u>13.60</u>	9.02	5.16
Language NAO	v2	<u>30.43</u>	13.45	<u>10.38</u>	5.18
STAformer	v2	30.61	16.67	10.06	5.62
STAformer+AFF	v2	32.39	17.38	10.26	5.70
STAformer+MH	v2	31.99	16.79	11.62	6.72
STAformer+MH+AFF	v2	33.50	17.25	11.77	6.75
Gain (rel %)	v2	<u>+10.1</u>	<u>+28.3</u>	<u>+13.4</u>	<u>+30.3</u>

Table 4: Results in mAP on the validation split of EPIC-Kitchens. **Best results** in bold. Relative gain is with respect to second best.

Model	N	N + V	N + δ	All
StillFast [42]	<u>21.24</u>	<u>12.41</u>	6.22	<u>3.28</u>
STAformer	24.16	15.55	7.08	4.31
STAformer + AFF	26.19	16.49	7.18	4.69
Gain (rel %)	<u>+23.3</u>	<u>+32.8</u>	<u>+15.4</u>	<u>+42.9</u>

STAformer and the integration of affordances. The joined semantic generalization capacity of environment affordances and the spatial refinement of interaction hotspots make STAformer + AFF excel in the $N + V$ mAP, obtaining a gain of +58.9% and 47.6% on v1 and v2, respectively.

Table 3 reports the results on the test split of Ego4D. Note that the test set is private, so we are only able to compare approaches showing test results in their papers. For fair comparisons, we report two settings with methods trained on v1 or v2 (a larger set also includes v1 annotations). Our method achieves consistent gains with respect to trained methods on v1, for instance, obtaining a +13.1% in mAP All and a +32.1% in $N + V$ mAP. Smaller but consistent gains of +7.4% and +1.8% are obtained for N and $N + \delta$ mAP, respectively. We observe similar gains when training on v2, with a +30.3% mAP All, +28.3% $N + V$ mAP, +13.4% mAP $N + \delta$ and +10.1% N mAP. It is worth noting that our approach also benefits from training on larger sets of data. Indeed, performance is improved in Table 3 when training on v2, with respect to our model trained on v1, increasing from 4.06 to 6.75 mAP All and jumping from 26.52 to 33.50 N mAP due to the

Table 5: Ablation study of the different components of STAformer on the v1-val split of Ego4D. ❄Encoder frozen 🛠Encoder last-blocks finetuned ⚙Full encoder trained.

Exp.	Image Encoder	Video Encoder	Temporal pooling	2D-3D Fusion	N	N + V	N + δ	All
[42]	R50 ⚙	X3D ⚙	Mean	Sum	16.21	7.52	4.94	2.48
A1	DINOv2 ❄	-	-	-	17.48	8.64	5.20	2.52
A2	DINOv2 ❄	DINOv2 ❄	Mean	Sum	15.82	7.65	4.11	2.19
A3	DINOv2 ❄	X3D ⚙	Mean	Sum	18.84	8.84	5.56	2.57
B1	DINOv2 ❄	TimeSformer 🛠	Mean	Sum	16.67	8.38	5.16	2.63
B2	DINOv2 ❄	TimeSformer 🛠	Conv	Sum	17.36	8.75	6.05	2.94
B3	DINOv2 ❄	TimeSformer 🛠	Frame-guided	Sum	19.78	10.04	6.35	3.39
C1	DINOv2 ❄	TimeSformer 🛠	Frame-guided	Dual $I \leftrightarrow V$ attn	20.08	10.21	6.51	3.47
C2	DINOv2 🛠	TimeSformer 🛠	Frame-guided	Dual $I \leftrightarrow V$ attn	21.71	10.75	7.24	3.53
C3	DINOv2 🛠	TimeSformer 🛠	Frame-guided	$I \rightarrow V$ c.attn	20.01	10.04	5.80	3.01
C4	DINOv2 🛠	TimeSformer 🛠	Frame-guided	$V \rightarrow I$ c.attn	20.12	10.31	6.30	3.35
C5	DINOv2 🛠	TimeSformer 🛠	MH.Frame-guided	MH.Dual $I \leftrightarrow V$ attn	23.02	11.57	7.86	3.85

joining effect of the affordances and the multi-head attention. Table 4 finally reports the results on EPIC-Kitchens. Since this benchmark is new, we train the official implementation of StillFast [42] on EPIC-Kitchens. Also in this case, our method achieves significant performance gains ranging from +15.5% in the case of N+ δ mAP to +42.9% in mAP All.

5.2 Ablation study

STAformer architecture: Table 5 compares the performance effects of the main components involved in the STAformer architecture. In experiment A1, we encode the image input with a pre-trained DINOv2 model [36] and discard the video, obtaining small gains with respect to the baseline [42]. While [42] fully trains both image-video encoders, the A1 version trains solely the STA prediction head and reflects the modelling capacity of DINOv2. However, simply extracting per-frame DINOv2 features and performing mean temporal pooling (Exp. A2), decreases the performance and indicates the limits of DINOv2 in modeling video dynamics. Using X3D [9], a convolutional 3D CNN, as the video encoder in Exp. A3 leads to improvements with respect to A1 (e.g., 18.84 vs 17.48 N mAP and 5.56 vs 5.20 N+ δ), indicating the advantage of appropriately encoding video dynamics.


We compare different versions of temporal pooling in experiments B1-B3 of Table 5 using a fine-tuned TimeSformer as video model.⁸ Computing the mean along the temporal dimension of the video features (exp. B1), leads to non-systematic gains compared to the image-only transformer baseline A1. Using a convolutional module for temporal pooling (exp. B2) helps modeling temporal cues, improving N+ δ mAP up to 6.05. However, the proposed frame-guided attention (exp. B3) achieves a joint spatio-temporal understanding of the video improving from 8.75 to 10.04 N+V mAP and from 2.94 to 3.39 All mAP.

Next, experiments C1-C5 of Table 5 assess the contribution of the proposed Dual Image-Video Attention module for 2D-3D feature fusion. Comparing ex-

⁸ We finetune the last three blocks of the model.

Table 6: Ablation of the effect of environment affordances and interaction hotspots on StillFast and STAformer. Results on the Ego4D val v1 split in Top-5 mAP.

Ours							Ours						
STA model	E.AFF	I.H	N	N+V	N+ δ	All	STA model	E.AFF	I.H	N	N+V	N+ δ	All
StillFast	✗	✗	16.20	7.47	4.94	2.48	STAformer	✗	✗	21.71	10.75	7.24	3.53
StillFast	✓	✗	18.44	8.46	5.47	2.85	STAformer	✓	✗	23.55	11.75	7.55	3.74
StillFast	✗	✓	17.82	7.62	5.05	2.53	STAformer	✗	✓	23.63	11.38	7.51	3.66
StillFast	✓	✓	19.34	8.58	5.55	2.95	STAformer	✓	✓	24.36	12.00	7.66	3.77
EGO-TOPO [33]							EGO-TOPO [33]						
STA model	E.AFF	I.H	N	N+V	N+ δ	All	STA model	E.AFF	I.H	N	N+V	N+ δ	All
StillFast	✓	✗	14.92	6.45	4.01	2.14	STAformer	✓	✗	17.21	8.45	5.32	2.64

**Fig. 5: Predicted environment affordances:** Linking across functionally similar environments (\mathcal{K}^V , \mathcal{K}^T) creates a robust affordance representation which captures the STA interaction. We show in orange the STA ground-truth label.

periments C1 vs. B3 shows small but consistent gains when dual image-video attention is used for fusion, as compared to simple sum fusion (20.08 vs. 19.78 N, 10.21 vs. 10.04 N + V, 6.51 vs. 6.35 N + δ and 3.47 vs. 3.39 All mAP), suggesting that it is beneficial to enrich image tokens with video tokens and vice versa for 2D-3D fusion. The effect is more significant when we finetune the last 3 blocks of the image encoder (exp C2), showing the benefits of adapting the generalistic feature space of DINOv2 to the egocentric perspective. Using standard cross-attention layers only with image tokens ($I \rightarrow V$ - C3) or video tokens ($V \rightarrow I$ - C4) as queries, while still allowing to outperform simple sum fusion (B3), performs worse than the proposed dual image-video attention (C2), suggesting again the need to incorporate the refinement of both modalities. Finally, incorporating multi-head attention on the temporal pooling and on the 2D-3D fusion (C5) produces a consistent improvement in all the metrics due to its ability to capture diverse patterns from multiple representation simultaneously.

Leveraging affordances for STA: Table 6 details the influence of environment affordances (E.AFF) and interaction hotspots (I.H), when integrated, separately and jointly, on the StillFast [42] baseline and the proposed STAformer model, showing in both cases consistent improvements. The E.AFF module refines the nouns and verbs probabilities, obtaining significant gains in N + V Top-5 mAP (8.46 vs. 7.47 in Stillfast and 11.75 vs. 10.75 in STAformer). However, training a NN classifier as in [33] does not produce a useful distribution of the affordances for later fusion with the STA probabilities. Our intuition is that the NN overfits




Fig. 6: Qualitative results. We show the predictions of our model in different scenarios.

to the interactions in the scene which are more obvious, losing the generalist quality of our predictions across environments. By re-weighting confidence scores based on the spatial prior provided by the interaction hotspots, the I.-H. module produces a general improvement in all the metrics (e.g., N mAP of 17.82 vs 16.20 in StillFast and 23.63 vs 21.71 in STAformer - mAP All of 2.53 vs 2.48 in StillFast and 3.66 vs 2.53 in STAformer). Combining environment affordances and hotspots brings significant improvements in both StillFast and STAformer. For instance, the proposed approach improves its N mAP from 21.72 to 24.36 and its All mAP from 3.53 to 3.77.

5.3 Qualitative results

Figure 5 visualizes the nouns and verbs affordance distribution for two query videos, along with the closest zones with respect to visual appearance and narrations. Though the ground truth STA class is not the top predicted class, it is present in both the predicted verb and noun affordances, validating our hypothesis that scenes with close descriptive and visual similarity afford the same interaction.⁹ Figure 6 reports the final top predictions of our model, showing in some cases multiple next-active objects due to the inherent uncertainty in forecasting tasks.⁹

6 Conclusions

In this paper, we addressed the problem of Short-Term object-interaction Anticipation (STA). Our main contributions are STAformer, a novel attention-based architecture for STA, and the integration of affordances to ground STA predictions into human behavior. Our results showcase the improvements given by the proposed architecture and affordance modules, which scores first on all splits of the challenging Ego4D and EPIC-Kitchens benchmarks. We also detailed the contribution of each individual component through ablations and showed that the integration of affordances is beneficial also to other STA architecture besides the proposed one. We will release the code and all the material, hoping that our work will enable future research in the area.

⁹ We report more qualitative results in the supplementary material.

Acknowledgements

Research at University of Catania has been supported by the project Future Artificial Intelligence Research (FAIR) – PNRR MUR Cod. PE0000013 - CUP: E63C22001940006. Research at the University of Zaragoza was supported by projects PID2021-125209OB-I00 and TED2021-129410B-I00 (MCIN/ AEI/10.13039/501100011033, FEDER/UE and NextGenerationEU/PRTR).

References

1. Bao, W., Chen, L., Zeng, L., Li, Z., Xu, Y., Yuan, J., Kong, Y.: Uncertainty-aware state space transformer for egocentric 3d hand trajectory forecasting. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 13702–13711 (2023)
2. Bertasius, G., Wang, H., Torresani, L.: Is space-time attention all you need for video understanding? In: ICML. vol. 2, p. 4 (2021)
3. Bi, H., Zhang, R., Mao, T., Deng, Z., Wang, Z.: How can i see my future? fvtraj: Using first-person view for pedestrian trajectory prediction. In: Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part VII 16. pp. 576–593. Springer (2020)
4. Chen, G., Xing, S., Chen, Z., Wang, Y., Li, K., Li, Y., Liu, Y., Wang, J., Zheng, Y.D., Huang, B., et al.: Internvideo-ego4d: A pack of champion solutions to ego4d challenges. arXiv preprint arXiv:2211.09529 (2022)
5. Chi, H.g., Lee, K., Agarwal, N., Xu, Y., Ramani, K., Choi, C.: Adamsformer for spatial action localization in the future. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 17885–17895 (2023)
6. Damen, D., Doughty, H., Farinella, G.M., Fidler, S., Furnari, A., Kazakos, E., Moltisanti, D., Munro, J., Perrett, T., Price, W., et al.: Scaling egocentric vision: The epic-kitchens dataset. In: Proceedings of the European conference on computer vision (ECCV). pp. 720–736 (2018)
7. Dessalene, E., Devaraj, C., Maynard, M., Fermuller, C., Aloimonos, Y.: Forecasting action through contact representations from first person video. IEEE Transactions on Pattern Analysis and Machine Intelligence (2021)
8. Do, T.T., Nguyen, A., Reid, I.: Affordancenet: An end-to-end deep learning approach for object affordance detection. In: 2018 IEEE international conference on robotics and automation (ICRA). pp. 5882–5889. IEEE (2018)
9. Feichtenhofer, C.: X3d: Expanding architectures for efficient video recognition. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 203–213 (2020)
10. Feichtenhofer, C., Fan, H., Malik, J., He, K.: Slowfast networks for video recognition. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 6202–6211 (2019)
11. Furnari, A., Battiato, S., Grauman, K., Farinella, G.M.: Next-active-object prediction from egocentric videos. Journal of Visual Communication and Image Representation **49**, 401–411 (2017)
12. Furnari, A., Farinella, G.M.: Rolling-unrolling lstms for action anticipation from first-person video. IEEE transactions on pattern analysis and machine intelligence **43**(11), 4021–4036 (2020)
13. Gibson, J.J.: The theory of affordances. Hilldale, USA **1**(2), 67–82 (1977)

14. Girdhar, R., Grauman, K.: Anticipative video transformer. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 13505–13515 (2021)
15. Girshick, R.: Fast r-cnn. In: Proceedings of the IEEE international conference on computer vision. pp. 1440–1448 (2015)
16. Goyal, M., Modi, S., Goyal, R., Gupta, S.: Human hands as probes for interactive object understanding. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 3293–3303 (2022)
17. Grauman, K., Westbury, A., Byrne, E., Chavis, Z., Furnari, A., Girdhar, R., Hamburger, J., Jiang, H., Liu, M., Liu, X., et al.: Ego4d: Around the world in 3,000 hours of egocentric video. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 18995–19012 (2022)
18. Jiang, J., Nan, Z., Chen, H., Chen, S., Zheng, N.: Predicting short-term next-active-object through visual attention and hand position. *Neurocomputing* **433**, 212–222 (2021)
19. Kitani, K.M., Ziebart, B.D., Bagnell, J.A., Hebert, M.: Activity forecasting. In: Computer Vision–ECCV 2012: 12th European Conference on Computer Vision, Florence, Italy, October 7–13, 2012, Proceedings, Part IV 12. pp. 201–214. Springer (2012)
20. Koppula, H.S., Saxena, A.: Anticipating human activities using object affordances for reactive robotic response. *IEEE transactions on pattern analysis and machine intelligence* **38**(1), 14–29 (2015)
21. Lee, Y.J., Ghosh, J., Grauman, K.: Discovering important people and objects for egocentric video summarization. In: 2012 IEEE conference on computer vision and pattern recognition. pp. 1346–1353. IEEE (2012)
22. Li, G., Jampani, V., Sun, D., Sevilla-Lara, L.: Locate: Localize and transfer object parts for weakly supervised affordance grounding. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 10922–10931 (2023)
23. Lin, T.Y., Dollár, P., Girshick, R., He, K., Hariharan, B., Belongie, S.: Feature pyramid networks for object detection. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 2117–2125 (2017)
24. Liu, M., Tang, S., Li, Y., Rehg, J.M.: Forecasting human-object interaction: joint prediction of motor attention and actions in first person video. In: Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part I 16. pp. 704–721. Springer (2020)
25. Liu, S., Tripathi, S., Majumdar, S., Wang, X.: Joint hand motion and interaction hotspots prediction from egocentric videos. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 3282–3292 (2022)
26. Luo, H., Zhai, W., Zhang, J., Cao, Y., Tao, D.: Learning visual affordance grounding from demonstration videos. *IEEE Transactions on Neural Networks and Learning Systems* (2023)
27. Marchetti, F., Becattini, F., Seidenari, L., Del Bimbo, A.: Multiple trajectory prediction of moving agents with memory augmented networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2020)
28. Montesano, L., Lopes, M., Bernardino, A., Santos-Victor, J.: Learning object affordances: from sensory–motor coordination to imitation. *IEEE Transactions on Robotics* **24**(1), 15–26 (2008)
29. Mur-Labadia, L., Guerrero, J.J., Martinez-Cantin, R.: Multi-label affordance mapping from egocentric vision. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 5238–5249 (2023)

30. Mur-Labadia, L., Martinez-Cantin, R., Guerrero, J.J.: Bayesian deep learning for affordance segmentation in images. In: 2023 IEEE international conference on robotics and automation (ICRA). IEEE (2023)
31. Myers, A., Teo, C.L., Fermüller, C., Aloimonos, Y.: Affordance detection of tool parts from geometric features. In: 2015 IEEE International Conference on Robotics and Automation (ICRA). pp. 1374–1381. IEEE (2015)
32. Nagarajan, T., Feichtenhofer, C., Grauman, K.: Grounded human-object interaction hotspots from video. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 8688–8697 (2019)
33. Nagarajan, T., Li, Y., Feichtenhofer, C., Grauman, K.: Ego-topo: Environment affordances from egocentric video. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 163–172 (2020)
34. Nawhal, M., Jyothi, A.A., Mori, G.: Rethinking learning approaches for long-term action anticipation. In: European Conference on Computer Vision. pp. 558–576. Springer (2022)
35. Nguyen, A., Kanoulas, D., Caldwell, D.G., Tsagarakis, N.G.: Object-based affordances detection with convolutional neural networks and dense conditional random fields. In: 2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). pp. 5908–5915. IEEE (2017)
36. Oquab, M., Darcet, T., Moutakanni, T., Vo, H., Szafraniec, M., Khalidov, V., Fernandez, P., Haziza, D., Massa, F., El-Nouby, A., et al.: Dinov2: Learning robust visual features without supervision. Transactions on Machine Learning Research (2024)
37. Park, H.S., Hwang, J.J., Niu, Y., Shi, J.: Egocentric future localization. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 4697–4705 (2016)
38. Pasca, R.G., Gavryushin, A., Kuo, Y.L., Hilliges, O., Wang, X.: Summarize the past to predict the future: Natural language descriptions of context boost multimodal object interaction. arXiv preprint arXiv:2301.09209 (2023)
39. Plizzari, C., Goletto, G., Furnari, A., Bansal, S., Ragusa, F., Farinella, G.M., Damen, D., Tommasi, T.: An outlook into the future of egocentric vision. arXiv preprint arXiv:2308.07123 (2023)
40. Plizzari, C., Perrett, T., Caputo, B., Damen, D.: What can a cook in Italy teach a mechanic in India? action recognition generalisation over scenarios and locations. In: ICCV2023 (2023)
41. Pramanick, S., Song, Y., Nag, S., Lin, K.Q., Shah, H., Shou, M.Z., Chellappa, R., Zhang, P.: EgoVLPv2: Egocentric video-language pre-training with fusion in the backbone. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 5285–5297 (2023)
42. Ragusa, F., Farinella, G.M., Furnari, A.: Stillfast: An end-to-end approach for short-term object interaction anticipation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 3635–3644 (2023)
43. Ragusa, F., Furnari, A., Livatino, S., Farinella, G.M.: The meccano dataset: Understanding human-object interactions from egocentric videos in an industrial-like domain. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 1569–1578 (2021)
44. Rhinehart, N., Kitani, K.M.: Learning action maps of large environments via first-person vision. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 580–588 (2016)

45. Rodin, I., Furnari, A., Mavroedis, D., Farinella, G.M.: Predicting the future from first person (egocentric) vision: A survey. *Computer Vision and Image Understanding* **211**, 103252 (2021)
46. Roy, D., Rajendiran, R., Fernando, B.: Interaction region visual transformer for egocentric action anticipation. In: *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*. pp. 6740–6750 (2024)
47. Shan, D., Geng, J., Shu, M., Fouhey, D.F.: Understanding human hands in contact at internet scale. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 9869–9878 (2020)
48. Team, E.: Short-Term object-interaction Anticipation quickstart. https://colab.research.google.com/drive/10k_6F106K8kX1S4sEnU62Ho0Bw_CPngR?usp=sharing (2023), [Online; accessed 03-March-2024]
49. Thakur, S., Beyan, C., Morerio, P., Murino, V., Del Bue, A.: Enhancing next active object-based egocentric action anticipation with guided attention. In *International Conference on Image Processing* (2023)
50. Thakur, S., Beyan, C., Morerio, P., Murino, V., Del Bue, A.: Guided attention for next active object@ ego4d sta challenge. *CVPR23 EGO4D Workshop STA Challenge* (2023)
51. Thakur, S., Beyan, C., Morerio, P., Murino, V., Del Bue, A.: Leveraging next-active objects for context-aware anticipation in egocentric videos. In: *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*. pp. 8657–8666 (2024)
52. Tong, Z., Song, Y., Wang, J., Wang, L.: Videomae: Masked autoencoders are data-efficient learners for self-supervised video pre-training. *Advances in neural information processing systems* **35**, 10078–10093 (2022)
53. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., Polosukhin, I.: Attention is all you need. *Advances in neural information processing systems* **30** (2017)
54. Zatsarynna, O., Abu Farha, Y., Gall, J.: Multi-modal temporal convolutional network for anticipating actions in egocentric videos. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 2249–2258 (2021)
55. Zhong, Z., Schneider, D., Voit, M., Stiefelhagen, R., Beyerer, J.: Anticipative feature fusion transformer for multi-modal action anticipation. In: *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*. pp. 6068–6077 (2023)