

ACTION-GUIDED ATTENTION FOR VIDEO ACTION ANTICIPATION

Anonymous authors

Paper under double-blind review

ABSTRACT

Anticipating future actions in videos is challenging, as the observed frames provide only evidence of past activities, requiring the inference of latent intentions to predict upcoming actions. Existing transformer-based approaches, which rely on dot-product attention over pixel representations, often lack the high-level semantics necessary to model video sequences for effective action anticipation. As a result, these methods tend to overfit to explicit visual cues present in the past frames, limiting their ability to capture underlying intentions and degrading generalization to unseen samples. To address this, we propose Action-Guided Attention (AGA), an attention mechanism that explicitly leverages predicted action sequences as queries and keys to guide sequence modeling. Our approach fosters the attention module to emphasize relevant moments from the past based on the upcoming activity and combine this information with the current frame embedding via a dedicated gating function. The design of AGA enables post-training analysis of the knowledge discovered from the training set. Experiments on the widely adopted EPIC-Kitchens-100 benchmark demonstrate that AGA generalizes well from validation to unseen test sets. Post-training analysis can further examine the action dependencies captured by the model and the counterfactual evidence it has internalized, offering transparent and interpretable insights into its anticipative predictions.

1 INTRODUCTION

Anticipating future actions in videos is challenging in computer vision, with broad implications for assistive systems, robotics, autonomous vehicles, and interactive entertainment. The core difficulty stems from the need to predict upcoming actions based solely on observed video frames, which provide only indirect evidence of subtle human intentions. Unlike action recognition, where its annotations are present and synchronized with observed frames, action anticipation, on the other hand, requires inferring future action labels from up-to-current observations. The observed frames often contain merely partially revealed and ambiguous visual cues, and the fact that the same past observations could lead to multiple possible future outcomes further makes the task inherently uncertain.

The emergence of the vision transformer has appeared as a dominant paradigm in video action anticipation modeling. However, existing vision transformer approaches typically rely on conventional self-attention over the token representation transformed from the video pixels. This design principle could have advantages in the action recognition task, as self-attention can effectively construct the visual patterns and match their synchronized annotations. Nonetheless, the inherent non-deterministic nature of the future makes action anticipation a more complex task than pattern recognition. Conventional self-attention can therefore be misled by visual clutter, potentially leading to overfitting.

To address this problem, we propose leveraging action predictions to sequentially guide the focus of attention. We call our method Action-Guided Attention (AGA), where both queries and keys are represented by action probabilities. The design builds upon the fundamentals of dot-product attention, which attends to values through correlations between query and key representations. Instead of relying on pixel-level features, AGA uses predicted actions as high-level semantic guidance. This formulation explicitly models the idea of predicting the next action conditioned on previously pre-

dicted actions, learning dependencies among actions to improve the anticipation of future events. The output of the AGA aggregates relevant past moments and further adaptively mixes the information derived from the current frame inputs through a dedicated gating function to balance between history context and current evidence.

On EPIC-Kitchens-100, AGA generalizes effectively from validation to unseen test sets, with a consistently narrow gap that suggests resistance to overfitting on partially observed video inputs. Its robustness is further validated on EPIC-Kitchens-55 and the sparsely annotated EGTEA Gaze+.

Moreover, the design of AGA also enables post-training analysis to uncover the knowledge learned during training. Through forward analysis and backward analysis, we can examine the action dependencies captured by the model and the counterfactual evidence it has internalized, offering transparent and interpretable insights into its anticipative predictions.

2 RELATED WORK

A broad spectrum of anticipation methods have been proposed over the past five years. Sequence modeling approaches, primarily based on recurrent networks (e.g., RNN, GRU, LSTM), were initially explored to map observed features to future actions. RULSTM (Furnari & Farinella, 2020) employs two LSTMs, the first summarizes the observed video sequence, while the second unrolls predictions according to the temporal distance to the start of the future action. By explicitly incorporating the anticipation interval and unrolling toward the end of the interval, this architecture established a strong baseline on large-scale video action anticipation benchmarks such as EPIC-Kitchens-55 (Damen et al., 2018) and EPIC-Kitchens-100 (Damen et al., 2022). However, sequential recurrent models are prone to error accumulation as time progresses. To address this, SRL (Qi et al., 2021) introduces a regulation mechanism that emphasizes novel information at each timestamp, contrasting it with previously observed content and modeling its correlation with past frames, thereby improving upon the RULSTM baseline. Building on this line, ImagineRNN (Wu et al., 2020) further enhances unrolling by learning to explicitly predict frame-wise differences within the anticipation interval in a contrastive manner.

More recently, transformer-based architectures have emerged, delivering significant improvements in action recognition and anticipation. The Anticipative Visual Transformer (AVT) (Girdhar & Grauman, 2021) introduced a causal attention decoder on top of the standard Vision Transformer framework for action anticipation, achieving state-of-the-art performance at its time of publication. MemViT (Wu et al., 2022) extended AVT by storing longer historical context within the attention keys and values through token compression techniques. Later, RaftFormer (Girase et al., 2023) addressed the high computational cost of attention by optimizing the model design for efficiency and faster inference.

Besides, several works have explored multi-modal integration to enhance prediction accuracy. S-GEAR (Diko et al., 2024) emphasized semantic representations by jointly supervising visual and language branches. AFFT (Zhong et al., 2023) adapted the GPT-2 architecture and proposed an efficient fusion mechanism capable of incorporating additional modalities such as audio and optical flow. InAViT (Roy et al., 2024) leveraged prior information in the form of hand masks to disentangle human interactions from environmental clutter. More recent directions investigate the use of vision-language models and large language models for anticipation (Zhang et al., 2023; Mittal et al., 2024; Wang et al., 2025). Other extensions include semantic action augmentation (Qiu & Rajan, 2025) and long-term anticipation (Zhao et al., 2023; Zatsarynna et al., 2025), further broadening the scope of multi-modal anticipation research.

Prior work also explores the semantic label space in anticipation, but unlike AGA they do not feed predictions back into the attention mechanism to reweight the visual representation. Abu Farha et al. (2018) use RNN/CNN predictors to decode future actions in the label space. Ke et al. (2019) refine predictions by injecting temporal features into an attention module. Zhao & Wilde (2020) condition the initial visual observation on an action-time sequence through a conditional GAN to generate diverse futures. However, in all these methods the predictions remain separate from feature extraction, they do not guide the attention over visual tokens. AGA is distinct in that it uses the model’s own evolving action predictions to dynamically focus attention on semantically meaningful visual cues.

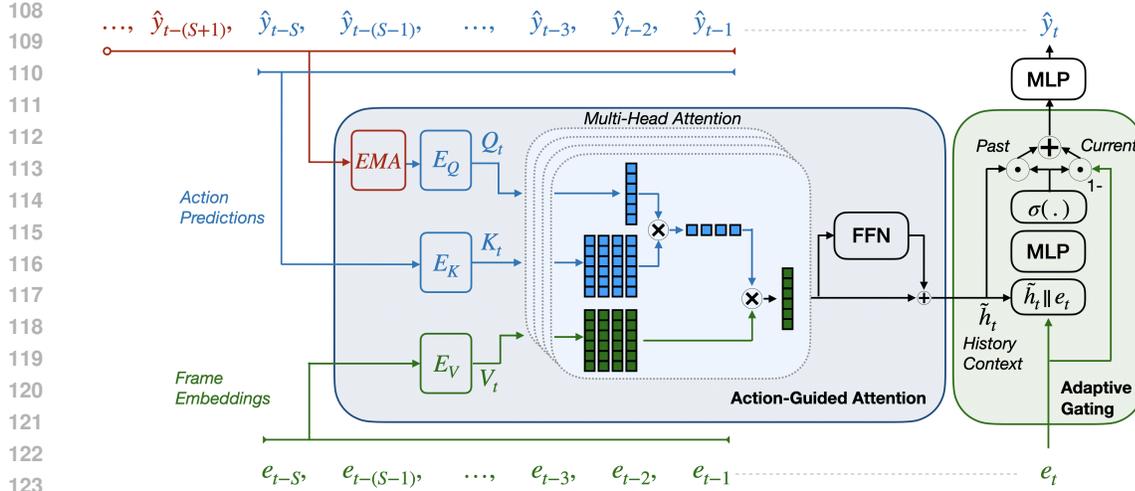


Figure 1: **Architecture overview.** The model consists of two modules. The *Action-Guided Attention* uses the most recent S action predictions as keys, the exponential moving average (EMA) of all predicted actions as the query, and S frame embeddings as values to generate a history context \tilde{h}_t . The *Adaptive Gating* then integrates this history context with the current frame embedding e_t to produce a fused representation, which is mapped to the new prediction \hat{y}_t .

In this work, we focus on a fundamental challenge in video action anticipation, which causes overfitting due to incomplete and uncertain observations associated with the future action. We propose an alternative attention design, AGA, that incorporates high-level, task-guided action representations as queries and keys, thereby mitigating attention to over-reliance on unreliable visual cues from partial video evidence. Our study considers explicitly anticipation methods that use RGB video frames as input, excluding those that integrate multi-modal or auxiliary knowledge sources such as text.

3 METHOD

Let $(x_t)_{t \geq 0}$ be a sequence of video frames sampled every $\Delta t > 0$ seconds, where $x_t \in \mathbb{R}^{C \times H \times W}$. For indices $r \leq s$, we write $x_{r:s}$ to denote the subsequence $(x_{\max(0,r)}, \dots, x_s)$, that is, values with $r < 0$ are clipped at zero; if s exceeds the total sequence length then set it to the last index, and $s \geq 0$. Here, t denotes the discrete frame index, corresponding to real time $t \cdot \Delta t$. Given the observed frames $x_{0:T}$, the goal is to predict the future action $y_{T+t_a} \in \mathbb{R}^{N_c}$, represented as a one-hot vector over N_c classes, which occurs $t_a \cdot \Delta t$ seconds after the last observed frame. For each index $0 < t < T$, the model outputs a probability distribution $\hat{y}_t \in [0, 1]^{N_c}$ as an estimate of the true future action y_{t+t_a} .

3.1 ACTION-GUIDED ATTENTION

We introduce a conditional attention mechanism that extends the standard formulation of Vaswani et al. (2017) by conditioning queries (Q) and keys (K) on semantic-level action predictions, while using frame-encoded video features as values (V). Ideally, the ground-truth action label y_t would provide the most reliable signal for guiding the estimation of future actions. However, because the ground-truth is unavailable at inference time, we instead rely on the self-predicted distribution \hat{y}_t as an approximate guide. In dot-product attention, correlations between query and key yield weights that select relevant values. Our design represents the features of a video frame by its corresponding action prediction, and then utilizes the dot-product between prediction embeddings to aggregate relevant frame features to forecast the anticipated action. The overall architecture is shown in Figure 1.

At each timestep t , the input video frame x_t is first processed by a frame-based backbone $f_{\text{backbone}}(\cdot)$, followed by an trainable encoder $f_x(\cdot)$ that extracts the frame features e_t :

$$e_t = f_x(f_{\text{backbone}}(x_t)).$$

The trainable function $f_x(\cdot)$ finetunes features from a frozen backbone, keeping the backbone modular and interchangeable for any image architecture.

These features are subsequently stored in a first-in, first-out (FIFO) queue of size S , a hyperparameter that determines the temporal window over which the model can reference past information. The number of the keys and values is thus dependent on S .

At each timestep, the frame features stored in the queue serve as the values, while the predictions from the same time, which can be viewed as semantic tags summarizing the information up to the current moment, are used to construct the keys. Specifically, instead of using the perceptual features e , we employ the past predictions \hat{y} to represent high-level semantics in the attention mechanism. Formally, we define:

$$K_t = E_K(\hat{y}_{t-S:t-1}) \quad V_t = E_V(e_{t-S:t-1}), \quad (1)$$

where $E_K(\cdot)$ and $E_V(\cdot)$ are implemented as multilayer perceptrons (MLP).

Correspondingly, the query also leverages the sequence of action predictions to match the semantic level with the keys. To effectively incorporate longer temporal dependencies, we apply an exponential moving average (EMA) over the past action predictions to form the query:

$$\bar{y}_t = \alpha \hat{y}_{t-1} + (1 - \alpha) \bar{y}_{t-1} \quad Q_t = E_Q(\bar{y}_t), \quad (2)$$

where \bar{y}_0 is initialized as the zero vector and α is empirically set to 0.8 (see Section 4.3 and Appendix B.1.2 for experiments on the choice of α). The function $E_Q(\cdot)$ consists of an MLP. Note that K_0 and V_0 are undefined and are never used for attention. **When the queue is empty, the first input frame, the attention is bypassed and its output is initialized to a zero vector. Starting from the second frame, the queue contains at least one timestep for both Q and K , allowing the attention mechanism to begin processing information even if the queue is not yet full.**

Then, with Q_t , K_t , and V_t defined, the multi-head dot-product attention is applied:

$$h_t = \text{MultiHead}(Q_t, K_t, V_t) = (\text{head}_1 \parallel \dots \parallel \text{head}_h) W^o,$$

where

$$\text{head}_i = \text{Softmax} \left(\frac{(Q_t W_i^Q) (K_t W_i^K)^\top}{\sqrt{d}} \right) (V_t W_i^V). \quad (3)$$

Here, h_0 is initially set to a zero-vector, W_i^Q , W_i^K , W_i^V are the trainable parameters for multi-heads, and W^o is the output weight for aggregating the multi-head attention outputs.

As usual, the output of the attention module is further transformed by a feedforward network (FFN), implemented as an MLP, with a residual connection. The FFN complements the attention mechanism by introducing non-linear transformations, while the learnable parameters within the attention are confined to the input and output embedding projections:

$$\tilde{h}_t = h_t + \text{FFN}(h_t).$$

For both multi-head attention and FFN the normalization is performed in PreNorm style (Xiong et al., 2020) using the RMSNorm layer (Zhang & Sennrich, 2019). The outputs of the multi-head attention and the FFN are then fused with e_t through the Adaptive Gating.

3.2 ADAPTIVE GATING

The relevance of the history context and the current visual evidence vary over time. Therefore, we componentwise fuse the history \tilde{h}_t and evidence e_t using a gating vector $g_t \in [0, 1]^d$ that has the same dimension as \tilde{h}_t and e_t , resulting in the output

$$o_t = g_t \odot \tilde{h}_t + (1 - g_t) \odot e_t,$$

where \odot denotes componentwise multiplication. The gating vector itself is computed as $g_t = \sigma(\text{MLP}(\tilde{h}_t \parallel e_t))$, where the MLP is composed of two linear layers and a ReLU activation function while $\sigma(\cdot)$ denotes the sigmoid function. Adaptive Gating is illustrated in Figure 1.

Entries of g_t close to 1 emphasize the history in those components, while entries near 0 favor the current visual evidence. This adaptive mixing lets the model rely on past or presence as needed.

216 Finally, the new prediction \hat{y}_t is obtained from the Adaptive Gating output o_t as

$$217 \hat{y}_t = \text{Softmax}(\text{MLP}(o_t)),$$

218 by applying an MLP, consisting of two linear layers with a ReLU activation to produce the logits,
219 followed by a Softmax.

222 3.3 FORWARD ANALYSIS

223 The design of AGA gives us insights into how a trained model works. Specifically, we can answer
224 questions of the kind: Given a currently observed action, which past actions does the model consider
225 as relevant for predicting the next action? We refer to this as forward analysis.

226 For example consider the actions $oc = \textit{open cupboard}$ and $cc = \textit{close cupboard}$. Suppose that in
227 the current frame the cupboard is opened and the history contains both, frames where the cupboard
228 is opened and where it is being closed. Which of those are more relevant? To find out we create
229 a query for the current observation, a key for each candidate action, and then compute the weights
230 that the attention mechanism gives to each candidate with regard to the query.

231 More precisely, for each action a , let y_a be the one-hot distribution with $y_a(a) = 1$ and $y_a(c) = 0$
232 for all action classes $c \neq a$. In our example, we create for oc the query embedding $Q_{oc} = E_Q(y_{oc})$
233 and for oc and cc together the key embedding $K_{oc,cc} = E_K(y_{oc} \| y_{cc})$. Using the Softmax expression
234 in Equation 3, we compute for each head i a vector $w^{(i)} = (w_{oc}^{(i)}, w_{cc}^{(i)})$ of attention weights and
235 then take the average $\bar{w} = 1/h \sum_{i=1}^h w^{(i)}$ of these vectors. The component \bar{w}_{oc} reflects the average
236 relevance given to frames where the cupboard is opened and \bar{w}_{cc} where it is closed. Figure 2 shows
237 actual weights for this example obtained from an AGA model trained on EPIC-Kitchens-100 dataset.

240 3.4 BACKWARD ANALYSIS

241 While forward analysis explains which past predictions the model *actually* relies on when producing
242 its output, backward analysis instead asks a counterfactual question: If the next action were a , what
243 changes in the past predictions would make the model assign a high probability to a ?

244 Let $Y = (\hat{y}_{t-S}, \dots, \hat{y}_{t-1})$ be the sequence of past predictions and let $\hat{y}_t = f_\theta(Y, X)$ denote the
245 predicted distribution at time t given the sequence of frames $X = X_{t-S:t}$. During training, each \hat{y}_i
246 is a probability distribution, but as a function f_θ is defined for arbitrary real vectors Y . Hence, for
247 fixed X , the following expressions are well defined for any Y .

248 Suppose the model has produced the prediction \hat{y}_t , while the correct next action is a , represented
249 by the one-hot distribution y_a . The discrepancy between prediction and target is measured by the
250 cross-entropy

$$251 H(y_a, \hat{y}_t) = - \sum_c y_a(c) \log \hat{y}_t(c) = - \log \hat{y}_t(a),$$

252 where the sum is taken over all class labels c . Since \hat{y}_t depends on the past predictions Y , we can
253 ask how Y can be changed so that the resulting prediction comes close to y_a . To capture this, we
254 define

$$255 L(Y) = H(y_a, f_\theta(Y, X)) = - \log (f_\theta(Y, X)(a)).$$

256 Minimizing $L(Y)$ amounts to finding past predictions that make the target action a more probable
257 for the given model. We can find an approximate local minimum by gradient descent, choosing a
258 step size η and starting from the original predictions $Y^{(0)}$, iterating through

$$259 Y^{(j+1)} = Y^{(j)} - \eta \nabla_Y L(Y^{(j)}),$$

260 until the loss function plateaus, that is,

$$261 |L(Y^{j+1}) - L(Y^j)| < \epsilon$$

262 for a threshold ϵ , or when the maximum number of iteration steps *iter* has been exceeded, that is,
263 $j \geq \textit{iter}$.

264 The resulting Y^* need not be itself a valid sequence of probability distributions, but it can still be
265 decomposed into vectors as long as distributions and we can compare it with the actual distributions

Methods	Modality	Overall Classes			Unseen Classes			Tail Classes		
		Action	Verb	Noun	Action	Verb	Noun	Action	Verb	Noun
RULSTM (Furnari & Farinella, 2020)	RGB, Flow, Obj	11.2	25.3	26.7	9.7	19.4	26.9	7.9	17.6	16.0
AVT+ (Girdhar & Grauman, 2021)	RGB, Obj	12.6	25.6	28.8	8.8	20.9	22.3	10.1	19.0	22.0
AVT++ (Girdhar & Grauman, 2021)	RGB, Flow, Obj	16.7	26.7	32.3	12.9	21.0	27.6	13.8	19.3	24.0
AFFT-TSN+ (Zhong et al., 2023)	RGB, Flow, Obj, Audio	13.4	19.4	28.3	9.9	14.0	24.2	10.9	12.0	19.5
AFFT-Swin+ (Zhong et al., 2023)	RGB, Flow, Obj, Audio	14.9	20.7	31.8	12.1	16.2	27.7	11.8	13.4	23.8
RAFTformer (Girase et al., 2023)	RGB	14.7	27.4	34.0	-	-	-	-	-	-
RAFTformer-2B (Girase et al., 2023)	RGB	15.4	30.1	34.1	-	-	-	-	-	-
S-GEAR (Diko et al., 2024)	RGB, Obj	14.7	32.0	25.9	-	-	-	-	-	-
S-GEAR-2B (Diko et al., 2024)	RGB, Obj	15.3	25.5	31.7	-	-	-	-	-	-
S-GEAR-4B (Diko et al., 2024)	RGB, Obj	15.5	26.6	32.6	-	-	-	-	-	-
AGA (Ours, Swin-B)	RGB	16.9	30.8	36.4	13.5	22.3	30.0	14.9	25.8	30.0

Table 1: **EK100 Test Results.** Evaluation on a unseen test set via the official challenge server.

in $Y^{(0)}$. In the difference $\Delta Y = Y^* - Y^{(0)}$ positive entries indicate supportive changes and negative entries indicate blocking changes for the target action a . This provides a counterfactual perspective on the dependencies captured by the model, complementary to forward analysis.

Qualitative results of both forward and backward analyses are presented in the experiments section, where they provide concrete insights into the model decision-making process and the knowledge that our model can explore in the training dataset.

4 EXPERIMENTAL RESULTS

4.1 EXPERIMENTAL SETUP

We evaluated our proposed method using the TSN (Wang et al., 2016) and Swin-B (Liu et al., 2021) backbones, with AGA applied on frozen features as a single attention layer with hidden size 2048 and 16 heads. **The queue size S is set to 16 (see Appendix B.1.1 for experiments on the choice of S).** Models were trained for 50 epochs with a batch size of 128, using AdamW with a learning rate of 2×10^{-4} , weight decay of 0.01, and a cosine schedule. All input frames were resized to 224×224 . Performance was measured using Mean Top-5 Recall (MT5R), which averages recall across classes to mitigate the dominance of frequent actions. **Following the training protocol of (Girdhar & Grauman, 2021), we compute class-reweighting factors with the inverted class frequency, and apply these weights when supervising the anticipation predictions with the ground-truth annotations.** Results were reported at the dataset-specific anticipation interval τ_a . **We converted τ_a (provided in seconds) into the corresponding anticipation index t_a (used by our method) as the number of time units Δt to reach τ_a , that is, $\tau_a = t_a \cdot \Delta t$.**

Experiments were conducted on three benchmarks. **EPIC-Kitchens-100 (EK100)** (Damen et al., 2022) contains 100 hours of video with 3,806 actions, including 67,217 training and 9,668 validation segments, **and was evaluated at $\tau_a = 1s$ and $\Delta t = 1s$.** **EPIC-Kitchens-55 (EK55)** (Damen et al., 2018) contains 55 hours of video with 2,513 actions, including 23,492 training and 4,979 validation segments, and was evaluated using Top-1 and Top-5 accuracy as well as MT5R **at $\tau_a = 1s$ and $\Delta t = 1s$.** **EGTEA Gaze+** (Li et al., 2018) comprises 28 hours of egocentric video with 106 actions, including 8,299 training and 2,022 validation clips, where we followed the protocol in (Girdhar & Grauman, 2021) and **reported Top-1 accuracy and mean Top-1 recall at $\tau_a = 0.5s$ and $\Delta t = 0.5s$.**

All the benchmark scores are evaluated on the time point $T \cdot \Delta t + \tau_a$, where T is the sequence length.

4.2 RESULTS

Tables 1 and 2 present the performance of the proposed AGA in comparison with prior methods on EPIC-Kitchens-100 for the unseen test set and the validation set, respectively. On the validation set, all methods are evaluated using the RGB modality only, and approaches that exploit additional modalities or external knowledge are excluded to ensure fair comparison. With the TSN backbone (highlighted in gray background colors), AGA achieves 17.5% accuracy, outperforming AFFT at 16.4%, S-GEAR at 14.9%, and other baselines. Replacing TSN with the stronger Swin-B backbone

Methods	Backbone	Overall Classes			Unseen Classes			Tail Classes		
		Action	Verb	Noun	Action	Verb	Noun	Action	Verb	Noun
RULSTM (Furnari & Farinella, 2020)	TSN	13.3	27.5	29.0	-	-	-	-	-	-
AVT-h (Girdhar & Grauman, 2021)	TSN	13.6	27.2	30.7	-	-	-	-	-	-
AVT-h (Girdhar & Grauman, 2021)	irCSN152	12.8	25.5	28.1	-	-	-	-	-	-
AVT-h (Girdhar & Grauman, 2021)	AVT-b	14.9	30.2	31.7	-	-	-	-	-	-
AFFT (Zhong et al., 2023)	TSN	16.4	21.3	32.7	13.6	24.1	25.5	14.3	13.2	25.8
AFFT (Zhong et al., 2023)	Swin-B	17.6	23.4	33.7	15.2	24.5	25.4	15.3	15.6	26.5
MeMViT 16x4 (Wu et al., 2022)	MViTv2-16	15.1	32.8	33.2	9.8	27.5	21.7	13.2	26.3	27.4
MeMViT 32x3 (Wu et al., 2022)	MViTv2-24	17.7	32.2	37.0	15.2	28.6	27.4	15.5	25.3	31.0
RAFTformer (Girase et al., 2023)	MViTv2-16	17.6	33.3	35.5	-	-	-	-	-	-
RAFTformer-2B (Girase et al., 2023)	MViTv2-16 + 24	19.1	33.8	37.9	-	-	-	-	-	-
S-GEAR (Diko et al., 2024)	TSN	14.9	25.8	29.8	-	-	-	-	-	-
S-GEAR (Diko et al., 2024)	ViT-B	18.3	31.1	37.3	-	-	-	-	-	-
S-GEAR-2B (Diko et al., 2024)	ViT-B x2	19.6	32.7	37.9	-	-	-	-	-	-
AGA (Ours)	TSN	17.5	32.2	35.7	11.9	31.4	24.9	16.8	26.9	31.8
AGA (Ours)	Swin-B	18.8	32.5	38.7	16.3	34.4	28.5	18.4	27.4	35.0

Table 2: **EK100 Validation Results.** Comparison of methods on the validation set using only RGB input. Methods highlighted in gray use the same TSN backbone weights.

Methods	Backbone	Top-1 Acc	Top-5 Acc	MT5R
RULSTM (Furnari & Farinella, 2020)	TSN	13.1	30.8	12.5
TempAgg (Sener et al., 2020)	TSN	12.3	28.5	13.1
ImagineRNN (Wu et al., 2020)	TSN	13.7	31.6	-
SRL (Qi et al., 2021)	TSN	-	31.7	13.2
AVT-h (Girdhar & Grauman, 2021)	TSN	13.1	28.1	13.5
AVT-h (Girdhar & Grauman, 2021)	AVT-b	12.5	30.1	13.6
AVT-h (Girdhar & Grauman, 2021)	irCSN152	14.4	31.7	13.2
DCR (Xu et al., 2022)	TSN	13.6	30.8	-
DCR (Xu et al., 2022)	irCSN152	15.1	34.0	-
DCR (Xu et al., 2022)	TSM	16.1	33.1	-
RAFTformer (Girase et al., 2023)	TSN	13.8	-	-
S-GEAR (Diko et al., 2024)	TSN	15.6	32.8	-
S-GEAR (Diko et al., 2024)	irCSN152	16.2	33.1	-
S-GEAR (Diko et al., 2024)	ViT-B	15.8	34.5	-
AGA (Ours)	TSN	13.5	32.1	14.3
AGA (Ours)	Swin-B	16.3	37.4	16.6

Table 3: **EK55 Validation Results.** Top-1/Top-5 action accuracy and MT5R at $\tau_a = 1s$. Methods highlighted in gray use the same TSN backbone weights.

further improves performance from 17.5% to 18.8%, with noticeable gains observed on unseen and tail classes. Methods that ensemble features from multiple backbones are highlighted in gray.

On the unseen test set, evaluated through the official challenge server, AGA maintains consistent performance, demonstrating strong generalization ability. Many methods listed in Table 1 are primarily competition submissions that rely on multi-modal inputs or ensemble strategies, whereas AGA achieves competitive scores without such enhancements. The gap between validation and test is consistently small, indicating that AGA generalizes effectively beyond the validation set compared with other methods. Leveraging action predictions to guide attention prevents the model from over-emphasizing the visual clutters, which is especially important in action anticipation where uncertainty is high and observations are less reliable than in other video recognition tasks.

Tables 3 and 4 compare the performance of AGA on two smaller-scale datasets, EK55 and EGTEA Gaze+. EGTEA Gaze+ provides sparse annotations only for the target action, without detailed labels for the preceding frames in the pre-action observation window. Accordingly, our model must rely entirely on its own predictions of prior activities without supervision. Nevertheless, empirical results show that AGA maintains strong performance, even under constrained annotation availability and sparse supervision. On EK55, S-GEAR reports higher accuracy than AGA under the same TSN backbone, likely due to its use of higher-resolution inputs (384×384 compared to 224×224). Moreover, EK55 evaluation is majorly reported in overall accuracy rather than in class-mean recall, which biases results toward common classes and obscures performance on rare actions.

378
379
380
381
382
383
384
385

Methods	Top-1 Acc			Mean Top-1 Recall		
	Action	Verb	Noun	Action	Verb	Noun
I3D-Res50 (Carreira & Zisserman, 2017)	34.8	48.0	42.1	23.2	31.3	30.0
FHOI (Liu et al., 2020)	36.6	49.0	45.5	32.5	32.7	25.3
TSN-AVT-h (Girdhar & Grauman, 2021)	39.8	51.7	50.3	28.3	41.2	41.4
AVT (Girdhar & Grauman, 2021)	43.0	54.9	52.2	35.2	49.9	48.3
TSN-AFFT (Zhong et al., 2023)	42.5	53.4	50.4	35.2	42.4	44.5
AGA (Ours, TSN)	43.5	54.3	52.2	35.5	43.8	46.6
AGA (Ours, Swin-B)	45.4	55.9	54.3	37.4	46.5	49.3

386 Table 4: **EGTEA Gaze+ Validation Results.** Top-1 accuracies and Mean Top-1 Recall at $\tau_a = 0.5s$.
387 Methods highlighted in gray use the same TSN backbone weights.
388

389
390
391
392
393
394
395

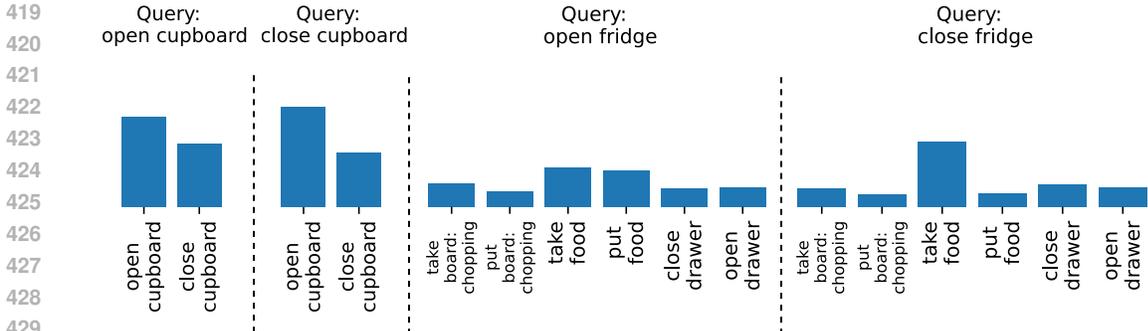
Configuration	MT5R	EMA α	MT5R
Reference (LSTM)	14.5	$\alpha = 0.0$	17.4
Baseline (Causal Attention)	15.9	$\alpha = 0.2$	18.7
Action-Guided Attention	18.2	$\alpha = 0.4$	18.4
Action-Guided Attention + Adaptive Gating	18.8	$\alpha = 0.6$	18.5
		$\alpha = \mathbf{0.8}$	18.8
		$\alpha = 1.0$	18.2

396 Table 5: Ablation study of proposed Action-Guided Attention. Table 6: Selection of EMA α for Q .
397

398
399
400 **4.3 ABLATION STUDY**

401 Table 5 presents the ablation study of the proposed model. The analysis begins with a baseline
402 model, causal attention (e.g., AVT). Replacing the standard dot-product attention with queries and
403 keys defined by action predictions yields a substantial improvement, raising accuracy from 15.9%
404 to 18.2%. In this setting, adaptive gating is disabled, so the attention output is directly added to
405 the current frame input without weighting past and current features. Incorporating adaptive gating
406 provides additional gains, further increasing accuracy from 18.2% to 18.8%. For reference, an
407 LSTM baseline trained under the same conditions is also reported.

408 Table 6 examines the impact of different α values in the exponential moving average used to form
409 the query. Larger values of α make the query depend more heavily on the most recent prediction.
410 When $\alpha = 0$, the query collapses to a constant (projected from the zero vector) across all timesteps,
411 leaving attention unconditioned and driven solely by the keys, which results in poor anticipation
412 accuracy. For $0.2 \leq \alpha \leq 0.8$, the query aggregates action dynamics over time, providing a coherent
413 action-guided signal and improving performance. When $\alpha = 1$, performance again drops because
414 the query depends only on the latest prediction. Based on empirical results, we adopt $\alpha = 0.8$ in our
415 model. All ablations reported here were conducted on the EK100 validation set and evaluated using
416 MT5R. [Appendix B.1.2 informs about studies on the other two datasets used in the experiments.](#)
417



430 Figure 2: **Forward Analysis** identifies which past actions the model attends to when predicting its
431 next action in response to a query. This analysis was conducted on the model trained with EK100.

	$t-3$	$t-2$	$t-1$	t
Video observation				
Original top-5 anticipation	dry ladle close cupboard eat squash move tray take tongs	dry ladle eat squash take pan close cupboard take tongs	close cupboard search container dry ladle take pan move tray	close cupboard dry ladle eat squash insert grater take pan
Backward analysis (Target to "take pan")	close cupboard dry ladle take pan move tray take tongs	take pan move sausage eat squash take tongs take spatula	take pan search container put pan move tray insert pan	take pan search container put pan put ladle take ladle

Figure 3: **Backward Analysis.** The figure compares the top-5 original predicted actions with the counterfactual supportive actions optimized toward the target action *take pan*. Each column represents a timestep; the final column shows the anticipated output. Suppressed actions are highlighted in red and promoted actions appear in green. This example is drawn from the EK100 validation set.

4.4 FORWARD ANALYSIS

Figure 2 illustrates the attention focus on past actions given different queries. Since AGA treats past actions (keys) as an unordered set, the order shown along the x-axis is not indicative. Four examples are analyzed using *open/close cupboard* and *open/close fridge* as queries to examine how the model allocates attention across past actions.

In the first two examples, where the queries are *open cupboard* and *close cupboard*, the model consistently attends more on earlier *open cupboard* rather than *close cupboard*. This behavior suggests that the model learns to revisit earlier events and search for evidence that indicates the duration of an *open cupboard* state, which may not be visible when the cupboard is closed. Moreover, the attention distribution conditioned on *open cupboard* is more uniform, implying that the next action following an *open* event is less certain and can be supported by multiple cues. By contrast, conditioning on *close cupboard* produces a relatively sharper distribution, reflecting more substantial confidence that the next action is linked to objects accessed during the preceding open period.

In the third and fourth examples, where the queries are *open fridge* and *close fridge*, the model shows a similar pattern. Conditioning on *open fridge* yields a more uniform attention across past actions, though food-related interactions receive higher weights. When conditioned on *close fridge*, the attention becomes more concentrated, focusing primarily on *take food*, indicating that the model leverages this evidence to predict subsequent actions involving the *food* that has been taken before. Overall, these examples suggest that AGA captured plausible dependencies and learnt meaningful associations between actions to focus on scenes in the past relevant for the immediate future.

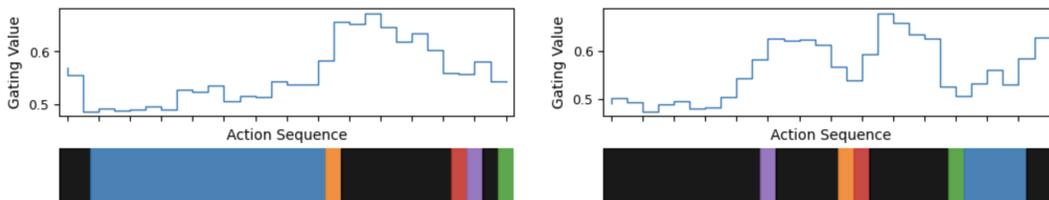


Figure 4: **Visualization of Adaptive Gating ratio.** The gating values, displayed alongside the action sequence, demonstrate context-aware behavior. In background regions (black), the gate retains historical context; in action regions (colored), it prioritizes current visual evidence.

4.5 BACKWARD ANALYSIS

Figure 3 presents the backward analysis of the model trained on the EK100 dataset.

The example is taken from the EK100 validation set. The original top-5 predicted actions are shown along with the counterfactual top-5 supportive actions via proposed backward analysis optimized toward the target action *take pan* (which originally ranked 5th). Each column corresponds to a timestep, with the final column (gray background) representing the anticipated output of the video clip. Actions marked in red indicate those suppressed by backward analysis, while those in green indicate promoted actions.

In the original top-5 predictions, the model prioritized actions such as *eat squash*, *dry ladle*, and *close cupboard*, with *take pan* largely absent. This pattern reveals a tendency to favor immediate, contextually dominant actions, as reflected in the diversity of verbs and nouns. Backward analysis conditioned on the counterfactual target *take pan* discovers an alternative reasoning path. Supportive actions involving containers and utensils, such as *take spatula*, *put pan*, *insert pan*, and *take ladle*, forming a semantically coherent trajectory toward the target. These results, obtained with model weights frozen, suggest that the model has already encoded multiple plausible futures. However, when faced with broad uncertainty, it originally distributed focus across diverse environmental cues. Backward analysis highlights a counterfactual path that converges on utensils, offering insights into both the model inference process and the latent knowledge it has acquired.

Backward analysis in the paper was conducted using a stopping criterion of $\epsilon = 1e-6$, a step size of $\eta = 1e2$, and a maximum of $iter = 5000$. Under this configuration, the optimization consistently identified counterfactual actions within the top-10 predictions as the new top-1 choice, validated across 30 validation samples drawn from EK100. When further varying the step size from $1e-1$ to $1e5$, we found that the same behavior remained stable for step sizes in the range $1e1$ to $1e4$, using the same stopping criterion.

4.6 ADAPTIVE GATING

Figure 4 shows two action sequences together with their average gating ratios, retrieved from the adaptive gating module during inference. Each sequence contains colored segments representing actions and black segments representing background or non-action content. The adaptive gate balances historical and current information: higher values emphasize past context, while lower values prioritize present input. Empirical observations reveal that fluctuations in the gating value often correlate with action transitions. For instance, the gate tends to rise during background segments to preserve past information, and then shifts attention to the current input when an action occurs. Importantly, the adaptive gate is not explicitly supervised by action boundaries; its context-aware behavior emerges naturally from training on the anticipation task.

5 CONCLUSION

We present AGA, Action-Guided Attention for video action anticipation, which leverages past predictions to create semantic representations that guide the model’s attention to parts of the history relevant for anticipating the immediate future. In addition, AGA dynamically mixes past and present visual information with its adaptive gating mechanism. Experiments show resistance to overfitting on top of overall competitive performance. Furthermore, AGA enables post-training forward and backward analysis, offering insight into the learned action dependencies and the reasoning process behind anticipation.

6 REPRODUCIBILITY STATEMENT

Implementation details are presented in Sections 3, 4.1, and Appendix B; links to the code implementation, video-processing scripts, and the preprocessed dataset will be released upon acceptance.

REFERENCES

- 540
541
542 Yazan Abu Farha, Alexander Richard, and Juergen Gall. When Will You Do What? — Anticipat-
543 ing Temporal Occurrences of Activities. In *IEEE Conference on Computer Vision and Pattern*
544 *Recognition*, 2018.
- 545 Joao Carreira and Andrew Zisserman. Quo vadis, action recognition? a new model and the kinetics
546 dataset. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2017.
- 547
548 Dima Damen, Hazel Doughty, Giovanni Maria Farinella, Sanja Fidler, Antonino Furnari, Evangelos
549 Kazakos, Davide Moltisanti, Jonathan Munro, Toby Perrett, Will Price, et al. Scaling egocentric
550 vision: The EPIC-Kitchens dataset. In *European Conference on Computer Vision*, 2018.
- 551 Dima Damen, Hazel Doughty, Giovanni Maria Farinella, Antonino Furnari, Evangelos Kazakos,
552 Jian Ma, Davide Moltisanti, Jonathan Munro, Toby Perrett, Will Price, et al. Rescaling egocentric
553 vision: Collection, pipeline and challenges for EPIC-Kitchens-100. *International Journal of*
554 *Computer Vision*, 130(1):33–55, 2022.
- 555 Anxhelo Diko, Danilo Avola, Bardh Prenkaj, Federico Fontana, and Luigi Cinque. Semantically
556 guided representation learning for action anticipation. In *European Conference on Computer*
557 *Vision*, 2024.
- 558
559 Antonino Furnari and Giovanni Maria Farinella. Rolling-unrolling LSTMs for action anticipation
560 from first-person video. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 43
561 (11):4021–4036, 2020.
- 562 Harshayu Girase, Nakul Agarwal, Chiho Choi, and Karttikeya Mangalam. Latency matters: Real-
563 time action forecasting transformer. In *IEEE/CVF Conference on Computer Vision and Pattern*
564 *Recognition*, 2023.
- 565 Rohit Girdhar and Kristen Grauman. Anticipative video transformer. In *IEEE/CVF International*
566 *Conference on Computer Vision*, 2021.
- 567
568 QiuHong Ke, Mario Fritz, and Bernt Schiele. Time-Conditioned Action Anticipation in One Shot.
569 In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019.
- 570
571 Yin Li, Miao Liu, and James M Rehg. In the eye of beholder: Joint learning of gaze and actions in
572 first person video. In *European Conference on Computer Vision*, 2018.
- 573 Miao Liu, Siyu Tang, Yin Li, and James M Rehg. Forecasting human-object interaction: joint pre-
574 diction of motor attention and actions in first person video. In *European Conference on Computer*
575 *Vision*, 2020.
- 576
577 Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo.
578 Swin transformer: Hierarchical vision transformer using shifted windows. In *IEEE/CVF Interna-*
579 *tional Conference on Computer Vision*, 2021.
- 580 Himangi Mittal, Nakul Agarwal, Shao-Yuan Lo, and Kwonjoon Lee. Can’t make an omelette with-
581 out breaking some eggs: Plausible action anticipation using large video-language models. In
582 *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024.
- 583
584 Zhaobo Qi, Shuhui Wang, Chi Su, Li Su, Qingming Huang, and Qi Tian. Self-regulated learning
585 for egocentric video activity anticipation. *IEEE Transactions on Pattern Analysis and Machine*
586 *Intelligence*, 45(6):6715–6730, 2021.
- 587 Yihui Qiu and Deepu Rajan. Action sequence augmentation for action anticipation. In *International*
588 *Conference on Learning Representations*, 2025.
- 589
590 Debaditya Roy, Ramanathan Rajendiran, and Basura Fernando. Interaction region visual transformer
591 for egocentric action anticipation. In *IEEE/CVF Winter Conference on Applications of Computer*
592 *Vision*, 2024.
- 593 Fadime Sener, Dipika Singhania, and Angela Yao. Temporal aggregate representations for long-
range video understanding. In *European Conference on Computer Vision*, 2020.

- 594 Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez,
595 Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in Neural Infor-*
596 *mation Processing Systems*, 2017.
- 597 Bingle Wang, Yao Tian, Shunzhou Wang, and Le Yang. Multimodal large models are effective
598 action anticipators. *IEEE Transactions on Multimedia*, 2025.
- 600 Limin Wang, Yuanjun Xiong, Zhe Wang, Yu Qiao, Dahua Lin, Xiaoou Tang, and Luc Van Gool.
601 Temporal segment networks: Towards good practices for deep action recognition. In *European*
602 *Conference on Computer Vision*, 2016.
- 603 Chao-Yuan Wu, Yanghao Li, Karttikeya Mangalam, Haoqi Fan, Bo Xiong, Jitendra Malik, and
604 Christoph Feichtenhofer. MeMViT: Memory-augmented multiscale vision transformer for effi-
605 cient long-term video recognition. In *IEEE/CVF Conference on Computer Vision and Pattern*
606 *Recognition*, 2022.
- 607 Yu Wu, Linchao Zhu, Xiaohan Wang, Yi Yang, and Fei Wu. Learning to anticipate egocentric
608 actions by imagination. *IEEE Transactions on Image Processing*, 30:1143–1152, 2020.
- 610 Ruibin Xiong, Yunchang Yang, Di He, Kai Zheng, Shuxin Zheng, Chen Xing, Huishuai Zhang,
611 Yanyan Lan, Liwei Wang, and Tiejian Liu. On layer normalization in the transformer architecture.
612 In *International Conference on Machine Learning*, 2020.
- 613 Xinyu Xu, Yong-Lu Li, and Cewu Lu. Learning to anticipate future with dynamic context removal.
614 In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022.
- 615 Olga Zatsarynna, Emad Bahrami, Yazan Abu Farha, Gianpiero Francesca, and Juergen Gall. Manta:
616 Diffusion mamba for efficient and effective stochastic long-term dense action anticipation. In
617 *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2025.
- 618 Biao Zhang and Rico Sennrich. Root mean square layer normalization. In *Advances in Neural*
619 *Information Processing Systems*, 2019.
- 620 Hang Zhang, Xin Li, and Lidong Bing. Video-LLaMa: An instruction-tuned audio-visual language
621 model for video understanding. *arXiv preprint arXiv:2306.02858*, 2023.
- 622 He Zhao and Richard P. Wilde. On Diverse Asynchronous Activity Anticipation. In *European*
623 *Conference on Computer Vision*, 2020.
- 624 Qi Zhao, Shijie Wang, Ce Zhang, Changcheng Fu, Minh Quan Do, Nakul Agarwal, Kwonjoon Lee,
625 and Chen Sun. AntGPT: Can large language models help long-term action anticipation from
626 videos? In *International Conference on Learning Representations*, 2023.
- 627 Zeyun Zhong, David Schneider, Michael Voit, Rainer Stiefelhagen, and Jürgen Beyerer. Anticipative
628 feature fusion transformer for multi-modal action anticipation. In *IEEE/CVF Winter Conference*
629 *on Applications of Computer Vision*, 2023.
- 630
631
632
633
634
635
636
637
638
639
640
641
642
643
644
645
646
647

f_x	E_Q, E_K	E_V	FFN	MLP in Adaptive Gating	MLP in classifier
LayerNorm	LayerNorm	LayerNorm	Linear (2048, 1024)	Linear (4096, 512)	LayerNorm
ReLU	ReLU	ReLU	GELU	ReLU	ReLU
Dropout	Dropout	Dropout	Dropout	Dropout	Dropout
Linear ($d_{backbone}$, 2048)	Linear (2048, 512)	Linear (2048, 2048)	Linear (1024, 2048)	Linear (512, 2048)	Linear(2048, N_c)
ScaleNorm					

Table 7: Modeling components in AGA.

Queue Size	AGA (MT5R)
S=4	18.4
S=16	18.8
S=30	18.0

Table 8: Ablation study with different queue sizes on EK100.

A IMPLEMENTATION DETAILS

For the TSN baseline, we utilized a BN-Inception backbone with pre-extracted features provided by the official EPIC-Kitchens action anticipation repository¹. For the Swin-B Transformer, we used open-source implementations with pre-trained weights from the *timm* library². Input frames for both architectures were resized to 224×224 with pixel values rescaled to the range $[-1, 1]$.

The EK100 dataset was sampled at 1 fps over 30-frame sequences, the EK55 dataset at 1 fps over 10-frame sequences, and the EGTEA Gaze+ dataset at 0.5 fps over 10-frame sequences. These sampling configurations align with the anticipation intervals described in the main manuscript.

Figure 7 details the architectures of f_x , E_Q , E_K , E_V , the FFN, and the MLPs. We applied a dropout rate of 0.6 on EK100 and 0.4 on both EK55 and EGTEA Gaze+ datasets.

B ADDITIONAL EXPERIMENTS

We carried out experiments to determine optimal hyperparameters for AGA and identify its computational cost. With another experiment, we investigated how errors in the predictions \hat{y} impact model accuracy. Finally, we examined how representing prediction uncertainty improves the accuracy of the AGA model.

B.1 SELECTION OF HYPERPARAMETERS

Two important hyperparameters for AGA are the length S of the queue from which the model can reference past information (Equation 1) and the smoothing factor α in the exponential moving average that determines the temporal window of the query (Equation 2).

B.1.1 QUEUE LENGTH

The experiments with EK100 reported in the paper were conducted with a queue length of $S = 16$. We also ran experiments for queue lengths 4 and 30. The mean top-5 recall for the different queue lengths is shown in Table 8.

B.1.2 SMOOTHING FACTOR

Intuitively, EMA injects temporal context into the query and smooths out model prediction jitter. The optimal smoothing constant α depends on factors such as the video sampling rate and the target action duration. As reported in Table 6 of the paper, we found in the EK100 experiment that the MT5R was between 18.4 and 18.8 for $0.2 \leq \alpha \leq 0.8$. This suggests that the accuracy is relatively stable for α in that range.

¹<https://github.com/epic-kitchens/C3-Action-Anticipation>

²<https://github.com/huggingface/pytorch-image-models,v0.5.4>

EMA α	Top-1	Top-5	MT5R
0.0	15.0	34.5	15.5
0.2	15.5	35.8	15.9
0.4	16.1	37.4	16.4
0.6	16.3	37.2	16.3
0.8	16.3	37.4	16.6
1.0	15.7	36.1	15.9

Table 9: Accuracy differences resulting from varying α on EK55.

EMA α	Top-1	Recall
0.0	41.8	34.7
0.2	43.9	36.2
0.4	44.9	37.3
0.6	45.2	37.1
0.8	45.4	37.4
1.0	43.8	36.0

Table 10: Accuracy differences resulting from varying α on EGTEA Gaze+.

Sequence Length	AVT (GFLOPs)	AGA (GFLOPs)
8	137.28	123.91
16	274.56	248.29
32	549.12	497.57

Table 11: Computational cost comparison.

Sequence Length	AVT (ms)	AGA (ms)
8	14.97	11.4737
16	20.57	23.4944
32	33.93	47.3397

Table 12: Inference time comparison.

We conducted the same study for the other two datasets used in the experiments. Tables 9 and 10 report the results of sweeping the smoothing factor α over $[0, 1]$. We observe that the optimal scores result from a value of $\alpha = 0.8$, which is consistent with the findings on EK100, while even across datasets a choice of α within the range $0.2 \leq \alpha \leq 0.8$ leads to relatively stable accuracy.

B.2 COMPUTATIONAL COST OF AGA

To evaluate the inference-time computational cost of AGA, we compared it against the AVT baseline. We created random tensors to mimic how sequences of frames are processed in a real inference setting. At each inference iteration we passed the random input to the model and let it do its processing: AVT in one go, due to its transformer architecture, AGA frame by frame, due to its sequential nature.

In Table 11, we show the total number of floating-point operations (FLOPs) computed for both models. AGA required slightly fewer FLOPs, despite using the heavier Swin-B backbone, versus AVT’s ViT-B

In Table 12, we show the inference times for AVT and AGA (with backbone Swin-B) measured on a single A100 GPU. The differences can be explained by the way in which the two models process the input. While AGA processes the sequence progressively, AVT handles the entire sequence in parallel.

B.3 PROPAGATION OF ERRORS

We wanted to understand how robust our model is with respect to errors in the predictions. We set up an experiment where AGA is subjected to inaccurate or noisy predictions that disrupt its action-guidance signal. As input, AGA received 30 frames from a 30-second clip from the EK100 validation set, with 1 second intervals in between. It had to predict the action in the last frame based on the preceding 29. For each sequence, we randomly created a permutation π of the 30 numbers $i = 0, \dots, 29$ and performed 30 runs. In run i , the predictions $\hat{y}_{\pi(j)}$, where $0 \leq j < i$, were forcibly reset to a uniform distribution, that is, each action was assigned the probability $\frac{1}{N_c}$. In this manner an increasing random subset of the predictions was made meaningless. In the table below, the first column contains the number of frames for which the the classifier output (i.e., \hat{y}) has been reset and the second column shows the resulting mean Top-5 recall (MT5R).

The experiment was run over the entire validation set of EK100. Figure 5 shows that the accuracy steadily decreases from 18.8% to 16.3%.

756
757
758
759
760
761
762
763
764
765
766
767
768
769
770
771
772
773
774
775
776
777
778
779
780
781
782
783
784
785
786
787
788
789
790
791
792
793
794
795
796
797
798
799
800
801
802
803
804
805
806
807
808
809

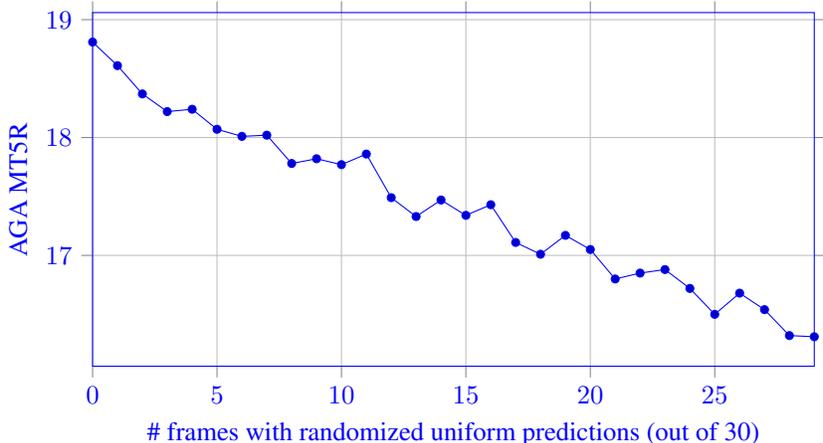


Figure 5: Robustness analysis of AGA against the error occurred in frame prediction.

Methods	Train On	Inference On	Overall Classes			Unseen Classes			Tail Classes		
			Action	Verb	Noun	Action	Verb	Noun	Action	Verb	Noun
GT	GT	GT	16.9	30.6	36.7	14.6	34.4	26.5	16.1	25.3	32.8
GT	Self-Pred	GT	17.1	31.9	36.0	14.1	32.9	26.3	16.6	26.8	32.2
Self-Pred (Top-1 One-Hot)	Self-Pred (Top-1 One-Hot)	Self-Pred (Top-1 One-Hot)	17.2	31.8	37.7	14.5	35.6	27.7	16.3	26.6	34.2
Self-Pred (Top-1 One-Hot)	Self-Pred	Self-Pred (Top-1 One-Hot)	17.5	33.4	36.4	16.3	32.6	27.7	16.5	28.4	32.5
AGA	Self-Pred	Self-Pred	18.8	32.5	38.7	16.3	34.4	28.5	18.4	27.4	35.0

Table 13: Comparison between different action-guidance signals for Q and K .

B.4 SELF-PREDICTION VS. GROUND TRUTH ACTIONS

AGA learns predictions \hat{y} that are probability distributions over the set of all action classes with a loss function that measures the difference to a one-hot distributions y_a for a specific action class a as the target.

Intuitively, during its training, AGA strives to generate predictions that are ever closer to the one-hot ground truth. It is therefore natural to ask whether and how the performance of a model changes if the predictions, which are generated to resemble the ground truth, are replaced with actual ground truth distributions. Whether or not this is the case should give us more insight into how AGA works.

We created an experiment where we compared the use of the ground truth as opposed to the self-predicted \hat{y} in query and keys. For frames with an action label a , we represented the ground truth with a one-hot y_a . For frames without label, occurring in background our transition segments, we used one-hot vectors representing a generic background class.

For additional insight, we included experiments where we simplified the self-predicted distributions to one-hot vectors $\hat{y}_{\top 1}$, carrying a 1 for the action with the highest probability in \hat{y} , thus indicating the Top-1 predicted action, and in this way aligning the representation to the ground truth format.

For both formats, ground truth and one-hot self-prediction, we ran two versions of the experiment: one where we used the revised format for both training and inference, and another one where we used it only for the inference.

We ran the experiments on the EK100 validation set. The results are shown in the Table 13. Here, “GT” stands for Ground Truth, that is, usage of the one-hot y_a , “Self-Pred” for Self-Predication, that is, usage of the \hat{y} , and “Self-Pred (Top-1 One-Hot)” for the usage of the vectors $\hat{y}_{\top 1}$. The figures refer to the accuracy on EK100 and the columns are identical to those in Tables 1 and 2 of the paper.

The results show that AGA, which leverages the full self-prediction distribution, achieves the highest accuracy. This can be intuitively explained by three factors:

- The full self-prediction carries richer intrinsic information than a one-hot ground truth label (as further evidenced by the accuracy drop when binarizing the self-prediction into one-hot);
- Self-prediction still provides meaningful signals when no action is annotated during the observation window (for instance, background or transition segments);
- Although GT provides the correct Top-1 future label, it lacks information about alternative plausible futures, whereas the (imperfect but semantically informative) self-prediction distribution better reflects the uncertainty and structure of the observations.

C MORE QUALITATIVE EXAMPLES

We provide additional qualitative examples from the EK100, EK55, and EGTEA Gaze+ datasets in Figures 6-11, where clips (a)–(d) illustrate successful predictions, with the ground-truth action revealed at the latest timestep, and clips (e)–(h) highlight failure cases in which the top-5 prediction of the model omit the target action. Figures are best viewed at full width and zoom in for full detail.

C.1 EPIC-KITCHENS-100 (EK100)

Figure 6 presents successful cases from the EK100 dataset. In video (a), the scene evolves slowly and the washing activity is clear. The noun *plate* is also unambiguous, which allows the model to converge on *wash plate*, while alternative verbs such as *hold*, *insert*, and *put* are gradually ruled out. In video (b), the target action *throw bin* occurs between the third and seventh frames and is correctly identified by the model. Video (c) shows a case with little movement, where the verb prediction narrows to a small set of consistent options over time, and the model ultimately selects the correct object. Video (d) is more challenging because the subject intent is unclear. The model nonetheless predicts a coherent sequence: *throw food* followed by *close bin*, while assigning lower probability to *open bin* once the hand interaction with the bin becomes visible.

Figure 7 illustrates failure cases. In video (e), the verb is predicted correctly but the noun is misclassified as *pizza*. In videos (f)–(h), errors result either from insufficient visual detail ((f)) or from ambiguity introduced by multiple plausible object candidates ((g) and (h)).

The EK100 examples demonstrate that the model can handle everyday kitchen activities with gradual scene changes, but performance degrades when object categories are visually similar or when frames lack sufficient detail.

C.2 EPIC-KITCHENS-55 (EK55)

Figure 8 highlights successful cases from the EK55 dataset, including challenging scenarios. In videos (a) and (b), the target objects (*container* and *board:cutting*, respectively) are not visible in the frames. Nevertheless, the model anticipates the actions *take container* and *take board:cutting* by leveraging contextual cues. Video (c) demonstrates the capacity to distinguish between similar actions, predicting *put plate* rather than *take plate*, based on the subject intention inferred from earlier frames. Accurate prediction requires retention of long-term evidence, which is supported by the adaptive gating mechanism and the AGA design. In video (d), the ongoing activity of cooking pasta is anticipated, with the model predicting *put-down spoon* as the next action in the final frame.

Figure 9 presents failure cases where the model encounters ambiguity or insufficient context. In video (e), the object *tofu* is occluded, leading to a misprediction as *container* because the opaque surface prevents visibility of the contents. Video (f) illustrates noun misclassification caused by visual similarity, with *olive* confused for *celery*. In video (g), the verb is misclassified because the hand interaction with *oil* is only partially visible. In video (h), the model confuses *chilli* with visually similar objects such as *tomato* and *pepper*.

The EK55 examples highlight the advantage of AGA and adaptive gating for leveraging long-term context, while also revealing difficulties when objects are occluded or visually confusable.

C.3 EGTEA GAZE+

Figure 10 provides qualitative examples from the EGTEA Gaze+ dataset, covering both successful and challenging scenarios. In video (a), the model captures subtle visual cues, such as the sponge disappearing from the bottom edge of the frame, which leads to the correct prediction of *put sponge*. In video (b), the model predicts *put container* and links the container to the *tomato* by recalling information from the initial frames, demonstrating the use of long-term memory to form object–action associations. Video (c) shows a case where observed movement suggests the verb *take* instead of the ground-truth verb *put* for the *bread:container*. Although direct evidence for returning the bread to the refrigerator is absent, the correct action still appears as the second-ranked prediction, reflecting the model strategy of coupling related verbs such as *put* and *take*. In video (d), the model addresses a sequential scenario involving *take cheese:container*, where one action is visible and the other remains hidden. Accurate recognition requires identifying the container holding the cheese, which can only be inferred from the presence of its cover.

Figure 11 presents failure cases caused by insufficient evidence or excessive ambiguity. In videos (e) and (f), the absence of salient details leads to incorrect predictions. Video (g) demonstrates noun ambiguity: the model predicts verbs confidently but assigns nearly equal probability to multiple visible objects. In video (h), the observation provides minimal context, leaving the model unable to determine the correct action.

The EGTEA Gaze+ examples highlight the model capacity to connect subtle visual cues with long-term memory, while also revealing the difficulty of disambiguating actions in situations with limited observations or an excessive number of potential object candidates.

918
919
920
921
922
923
924
925
926
927
928
929
930
931
932
933
934
935
936
937
938
939
940
941
942
943
944
945
946
947
948
949
950
951
952
953
954
955
956
957
958
959
960
961
962
963
964
965
966
967
968
969
970
971



Figure 6: Four video clips from the **EPIC-Kitchens-100** validation set, illustrating **correct** predictions. Each clip demonstrates the last eight frames along with the top-5 action anticipations at $\tau_a = 1s$, with the ground truth highlighted in bold green.



1018 Figure 7: Four video clips from the **EPIC-Kitchens-100** validation set, showcasing **incorrect** pre-
1019 dictions. Each clip demonstrates the last eight frames and the top-5 action anticipations at $\tau_a = 1s$,
1020 with the ground truth revealed in the final frame. The correct verb and noun are highlighted in bold
1021 green.
1022
1023
1024
1025

1026
1027
1028
1029
1030
1031
1032
1033
1034
1035
1036
1037
1038
1039
1040
1041
1042
1043
1044
1045
1046
1047
1048
1049
1050
1051
1052
1053
1054
1055
1056
1057
1058
1059
1060
1061
1062
1063
1064
1065
1066
1067
1068
1069
1070
1071
1072
1073
1074
1075
1076
1077
1078
1079



Figure 8: Four video clips from the **EPIC-Kitchens-55** validation set, illustrating **correct** predictions. Each clip displays the last eight frames and the top-5 action predictions. The ground truth is highlighted in bold green in the latest frame, occurring at $\tau_a = 1s$.

1134
1135
1136
1137
1138
1139
1140
1141
1142
1143
1144
1145
1146
1147
1148
1149
1150
1151
1152
1153
1154
1155
1156
1157
1158
1159
1160
1161
1162
1163
1164
1165
1166
1167
1168
1169
1170
1171
1172
1173
1174
1175
1176
1177
1178
1179
1180
1181
1182
1183
1184
1185
1186
1187



Figure 10: Four video clips from the **EGTEA Gaze+** validation set, highlighting **correct** predic-
tions. Each clip displays the last eight frames and the top-5 action predictions. The ground truth is
highlighted in bold green in the latest frame, occurring at $\tau_a = 0.5s$.



Figure 11: Four video clips from the EGTEA Gaze+ validation set, showcasing incorrect predictions. Each clip displays the last eight frames and the top-5 action predictions. The ground truth is highlighted in bold green in the latest frame, occurring at $\tau_a = 0.5s$.

1242 D LLMS USAGE
1243

1244 The authors utilized LLM powered AI tools (GPT-5 and Grammarly) to proofread sentences and
1245 identify grammatical errors.
1246

1247
1248
1249
1250
1251
1252
1253
1254
1255
1256
1257
1258
1259
1260
1261
1262
1263
1264
1265
1266
1267
1268
1269
1270
1271
1272
1273
1274
1275
1276
1277
1278
1279
1280
1281
1282
1283
1284
1285
1286
1287
1288
1289
1290
1291
1292
1293
1294
1295