

Intention-Conditioned Long-Term Human Egocentric Action Anticipation

Esteve Valls Mascaró¹Hyemin Ahn²Dongheui Lee^{1,3}

Abstract

To anticipate how a person would act in the future, it is essential to understand the human intention since it guides the subject towards a certain action. In this paper, we propose a hierarchical architecture which assumes a sequence of human action (low-level) can be driven from the human intention (high-level). Based on this, we deal with long-term action anticipation task in egocentric videos. Our framework first extracts this low- and high-level human information over the observed human actions in a video through a Hierarchical Multi-task Multi-Layer Perceptrons Mixer (H3M). Then, we constrain the uncertainty of the future through an Intention-Conditioned Variational Auto-Encoder (I-CVAE) that generates multiple stable predictions of the next actions that the observed human might perform. By leveraging human intention as high-level information, we claim that our model is able to anticipate more time-consistent actions in the long-term, thus improving the results over the baseline in Ego4D dataset. This work results in the state-of-the-art for Long-Term Anticipation (LTA) task in Ego4D by providing more plausible anticipated sequences, improving the anticipation scores of nouns and actions. Our work ranked first in both CVPR@2022 and ECCV@2022 Ego4D LTA Challenge.

1. Introduction

In our everyday life, reasoning about next actions is essential before executing a certain complex task. Humans can project oneself into the future through a constructive imagination system, that allows them to anticipate the future actions of themselves as well as others. Based on this, vari-

ous applications such as task and motion planning (TAMP), or human-robot-collaboration can be realized. For instance, after understanding that someone is about to make a salad, we can first anticipate his/her next actions to later assist this person to prepare ingredients. This field of research in computer vision, known as Long-Term Anticipation (LTA) of human actions, aims at forecasting the actions that a human will perform most likely based on the past observations.

The fundamental challenge of LTA of human actions is the inherent uncertainty of the future. The uniqueness of the human being results in high variability of how each of us executes a certain task. Moreover, this behaviour may vary for the same individual at different moments and also depending on the environment. However, despite the theoretical high variability of possible predictions, the future often has only a limited number of plausible outcomes. Inspired by this, we hypothesize that the arbitrariness of future events can be narrowed down through conditioning on past observations, which would imply the context of the whole task. For instance, if we observe a human cutting a tomato and we know that the human's intention is to make a salad, the variety of ingredients that the human might interact with can be reduced. This human intention, defined as a high-level hypothesis, conditions the human behavior and reduces the variability and uncertainty of the future.

Therefore, we develop a methodology that aims to constrain the variability of future actions based on the human intention estimated from past observations. We predict a hierarchical structure from a sequence of videos, each depicting a particular human action. From this given video clip sequence, we define two different levels of abstraction, as visualized in Fig. 1. First, we explore the human intention as the highest concept that defines the ultimate objective pursued by the human while performing a task. Second, we make use of the sequence of low-level actions to perform a certain task. We aim to mimic a human's reasoning process when predicting which actions to perform given a high-level instruction: (i) determine the current status, or context, of the task by observing the steps already completed (past low-level actions), (ii) plan the next steps based on the ultimate goal (intention) of the task. Therefore, based on the historical sequence of low-level actions performed by a human, we propose to exploit the high-level human intention as a guidance that conditions the next actions that will be

¹Esteve Valls Mascaró and Dongheui Lee are with Autonomous Systems, Technische Universität Wien (TU Wien), Vienna, Austria (e-mail: {estev.valls.mascaró, dongheui.lee}@tuwien.ac.at).

²Hyemin Ahn is with Artificial Intelligence Graduate School (AIGS), Ulsan National Institute of Science and Technology (UNIST), Ulsan, Korea (e-mail: hyemin.ahn@unist.ac.kr).

³Dongheui Lee is also with the Institute of Robotics and Mechatronics, German Aerospace Center, Wessling, Germany.

This work is funded by Marie Skłodowska-Curie Action Horizon 2020 (Grant agreement No. 955778) for project 'Personalized Robotics as Service Oriented Applications' (PERSEO).

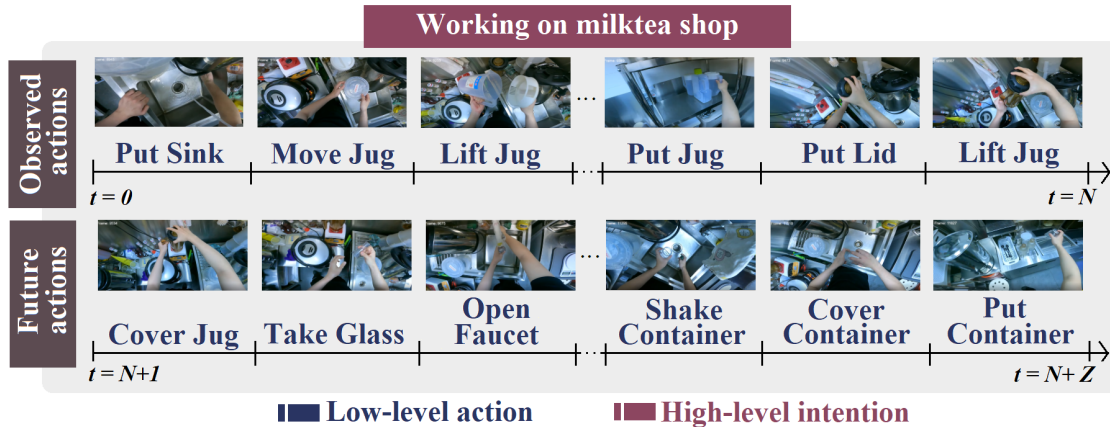


Figure 1: **An example of the hierarchy structure of a human task.** Egocentric sequence of videos of a human ‘Working on milktea shop’ (in purple, describing the high-level human intention) from Ego4d [13]. In blue, a sequence of low-level actions labels performed by the camera wearer is shown. This paper proposes a methodology that understands a human task based on this hierarchical structure. Our model extracts a high-level human intention information and N action labels from the observed sequence of N clips (first row) to facilitate the anticipation of low-level Z actions in its future (second row).

executed in the future.

Following the division of our methodology, we define our framework as a two-step. First, we propose a Hierarchical Multitask Multi-Layer Perceptrons (MLP) Mixer (H3M), to classify each observed video to an action label, as well as to extract the overall intention of the human. The MLP Mixer-based architecture [32] has been empirically shown to be an optimal model that uses the repeated MLP layers through temporal and spatial channels. Our H3M is designed as a multitask network [6] to exploit dependencies between low-level actions and high-level intentions while making the network more efficient. Secondly, we design an Intention-Conditioned Variational Autoencoder (I-CVAE), to anticipate the user’s future actions conditioned by the human intention and the observed past actions. Variational Autoencoder (VAE) [17] has been shown in [20, 21] to effectively model the human action sequence distribution. However, to narrow the uncertainty of the future, our model is based on Conditional VAE (CVAE), inspired by [24], leveraging the inferred human intention as a latent condition that provides a guideline to the model to anticipate the upcoming action sequence.

To demonstrate the effectiveness of our approach, we make use of the most diverse dataset of human videos currently available, Ego4D [13], and in particular we evaluate our results in the LTA benchmark. Ego4D provides first-person videos of humans experiencing everyday activities around the world. In the case of the LTA task, it proposes to predict the future sequence of actions of the camera user from the untrimmed video of its past. Low-level actions have been annotated in videos, which are classified depending on a scenario, that we understand as human

intention, as can be seen in Fig. 1

Finally, we report quantitative results that lead our approach to win both CVPR@2022 and ECCV@2022 Ego4D Long-Term Action Anticipation (LTA) challenges. We extend the results with a detailed discussion based on our ablation study. To sum up, the contributions can be summarized as follows:

1. It aims to extract the sequence of low-level actions as well as the human intention from a video through a multi-task hierarchical structure.
2. It promotes the use of high-level intention as a condition to anticipate the future sequence of actions.
3. It provides detailed analysis of long-term human action anticipation task, based on ablation studies, that aims to point out new research questions.

2. Related work

Long-Term Anticipation (LTA) has been a fundamental challenge in the computer vision research community. In the following section, we discuss the most relevant research in that field. Then, we review several works for hierarchical extraction and generative models.

2.1. Long-Term Anticipation

Predicting human’s future events has been covered in computer vision based on different tasks, such as generating skeleton motion [24, 29, 31], predicting future human trajectory [11, 23, 34] or generating a sequence of future action labels [20, 21, 28]. Due to the broadness of the literature, we will focus on works related to the action sequence anticipation task.

Existing literature [1, 2, 20, 21, 22, 26, 28] can be classified into two categories in terms of how to deal with the future. On the one hand, researchers aimed at simplifying the problem by modelling future as a deterministic function based on the observed video, regardless of future uncertainty [2, 22]. To overcome this simplification, two-step approaches were also taken into account in [2]: Recurrent neural network (RNN) [5] was used to first infer the action label sequence from the observed video, and then feed these classified actions into another RNN to forecast the future action sequence. However, [22] argued that anticipating unseen sequence of future action labels directly from observed videos provides more contextual cues.

On the other hand, inspired by the probabilistic nature of the future, [1] attempted to model the uncertainty through recursively sampling a future action from a softmax distribution obtained through a deterministic RNN based on previously predicted actions. To adapt the probabilistic essence in the architecture, Variational Autoencoder (VAE) [17] model, combined with point process models, was proposed in [21] to learn a latent distribution conditioned on the observed action sequence. Based on that, future actions are generated. [20] reused a VAE-approach but modelled a Multi-Head attention-based (MHA) Variational RNN to encode the latent distribution, which is concatenated in the RNN’s hidden state to decode the next actions recursively.

Our work is also inspired by both generative approaches [20, 21] but attempts at minimizing the future uncertainty by adding human intention as a condition learned from the latent distribution. While experimental results [28] demonstrated the key role of recent actions for immediate future predictions, we claim that knowing the intention guides the intelligent agent through a better long-term anticipation.

More recent work such as [26] demonstrated the use of latent goals, obtained from the observed actions, as a feature representation used to anticipate the next action. However, this latent representation is not explainable as it does not consist of language-based labels. Moreover, [26] only attempted to select one next action based on RNN proposed candidates, then dealing with short-term anticipation. On the contrary, we propose a single high-level explainable label, the human intention, as a guidance for the anticipation of long-term sequence of low-level actions. This attempt to hierarchically model the future is inspired by [30], which coped with the uncertainty by abstracting the level of prediction when the model confidence score was low. Our goal differs in nature, as we aim to always predict low-level actions, but encourage confidence to the model based on the high-level intention.

To the best of our knowledge, our work is the first attempt to decompose the future in a two-level explainable hierarchical structure. This design allows to deal with time-uncertainty by top-down approaches: high-level intention is used for robust anticipation of the low-level actions.

2.2. Generative Models

Understanding and modelling the data distribution and generalizing it to unseen scenarios is key for a deep generative model. When generative models succeed in synthesizing plausible data, it is assumed that the model has learned the data distribution properly. Deep generative models can be applied to a wide range of domains, from the generation of images [3, 25, 27] to text [4, 15], skeleton movement [24, 29, 31] and the synthesis of forthcoming actions [20, 21].

Despite the enormous variety of works regarding generative models, a recent convention is to classify deep generative models into three different directions [12, 14, 17]. First, Generative Adversarial Networks (GANs) [3, 12, 29] benefit from adversarial training of two networks that contest to maximize its own objective function in opposed tasks, thus encouraging the opponent network to improve its performance for generating data. Second, Variational Autoencoders (VAEs) [15, 17, 20, 21, 24] learn to encode large quantity of data into low-dimensional latent space, and then reconstruct the original data based on that latent space representation. Finally, Diffusion Models (DMs) [14] have become popular recently due to their great performance in image synthesis tasks [25, 27]. DMs work by destroying the data iteratively through Gaussian noise addition, and learning how to reverse the noise injection process by gradually denoising the sample.

This work is inspired by [20, 21] that aim to generate future human actions based on observed actions by VAEs, but refers to [24] more to condition the prediction result based on the estimated human intention. Our architecture is based on [24], where a Transformer Encoder-Decoder based architecture [33] is designed to synthesize human motions conditioned on a categorical action through the use of distribution parameter tokens. These parameters encode the context information (past observations and intention) and serve as conditioners for the transformer decoder.

3. Methodology

Long-Term Anticipation of human actions needs to exploit temporal dependencies among the observed actions to generate plausible human action sequences in the future. Our two-step approach first aims at understanding the observed actions through a Hierarchical MultiTask MLP Mixer (H3M), described in Section 3.2 in a bottom-up approach. Then, a Transformer-based Encoder-Decoder structure [33] is proposed for the Intention-Conditioned Variational Auto-Encoder (I-CVAE) in Section 3.3, inspired by [24]. Fig. 2 shows the overview of the proposed framework, that represents how the structure is extracted from the observed past and used to condition the future actions to be generated.

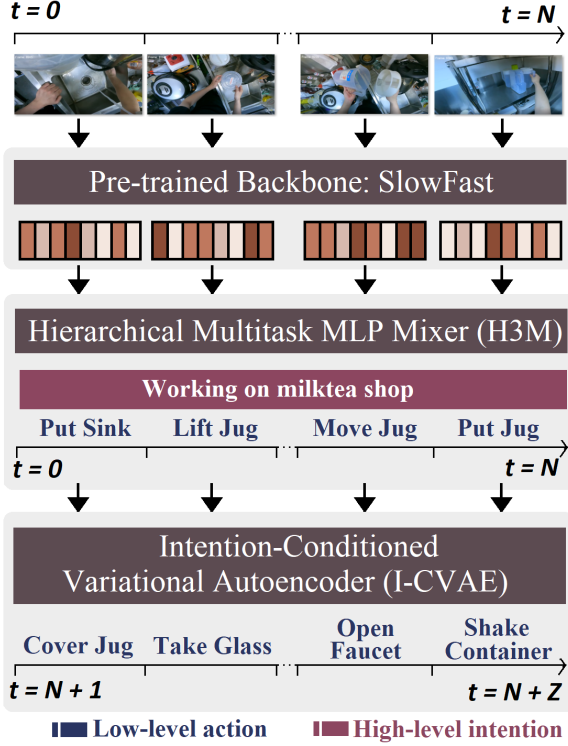


Figure 2: **Overall proposed framework.** Provided pre-extracted features for $N = 4$ observed videos are fed to our Hierarchical Multitask MLP Mixer model (H3M) to obtain low-level action labels and high-level intention. Results are fed into our Intention-Conditioned Variational AutoEncoder (I-CVAE) that anticipates subsequent $Z = 20$ actions.

3.1. Problem Formulation

Let $\mathbf{a}_t = (v_t, n_t)$ denote an action label at time t , which consists of verb label v_t and noun label n_t . Then, past N observed action of the camera wearer from a given untrimmed video can be represented as $\mathbf{A}_{obs} = [\mathbf{a}_1, \dots, \mathbf{a}_N]$. Our aim is to predict the future sequence of Z actions such as $\mathbf{A}_{pred} = [\mathbf{a}_{N+1}, \dots, \mathbf{a}_{N+Z}]$, by generating K possible sequences \mathbf{A}_{pred} to account for variations.

For a single video clip $\mathbf{V}_t = \{\mathbf{I}_t, \dots, \mathbf{I}_{t+T}\}$ which consists of image frames \mathbf{I} from time t to $t + T$, Ego4d benchmark provides a set of feature vectors $\mathbf{F} = \{\mathbf{f}_1, \dots, \mathbf{f}_T\}$, where $\mathbf{f}_t \in \mathbb{R}^{2304}$ denotes a feature vector obtained from image frames in one-second video. Each f is obtained from SlowFast architecture [10] pre-trained with Kinetics dataset [16]. We apply zero-padding in \mathbf{F} to ensure the same T duration for all video clips.

Then, from the observed N video clips, we obtain a N -sequence of pre-extracted visual features $\mathbf{F}_{obs} = [\mathbf{F}_1, \dots, \mathbf{F}_N]$, which we use as an input for our H3M model. Obtained two-level outputs composes the input for I-CVAE: (i) a high-level intention prediction \hat{I} is denoted as

a unique label that classifies the overall goal of the camera wearer for the whole task; and (ii) a low-level action sequence prediction $\hat{\mathbf{A}}_{obs}$. Finally, both \hat{I} and $\hat{\mathbf{A}}_{obs}$ are used for generating K possible future action sequences $\hat{\mathbf{A}}_{pred}$.

Next, we describe our proposed two independent architectures that aim to (i) obtain the hierarchical structure of the human task and (ii) anticipate the low-level actions conditioned by the human intention.

3.2. Hierarchical Multitask MLP Mixer (H3M)

As shown in Fig. 3, the pre-extracted features $\mathbf{F}_{obs} = [\mathbf{F}_1, \dots, \mathbf{F}_N]$ that describe the visual information of the untrimmed egocentric video are fed into a Hierarchical Multitask MLP Mixer (H3M) architecture. The Action Mixer model processes \mathbf{F}_t in parallel to encode the visual information that define a given $\mathbf{a}_t = (v_t, n_t)$. Let $\mathbf{F}_t \in \mathbb{R}^{T \times 2304}$, composed by T feature patches f_t . Mixer layer is defined in [32] and is composed of 2 MLP blocks encapsulated between two matrix transposition operations that capture the global context of f_t . The first MLP block identifies strong temporal dependencies and mixes the data among each tokenized patched, while the second MLP block leverages spatial same-patch features. Each MLP block consists of two fully-connected layers with a GELU activation function in-between. Finally, global average pooling is applied in the T dimension to project $\mathbb{R}^{T \times 2304}$ to \mathbb{R}^{2304} , encoding the action representation \mathbf{x}_t for an observed action. The action head of the H3M projects and classifies each \mathbf{x}_t as \hat{v}_t and \hat{n}_t through a Fully Connected layer, obtaining N $\hat{\mathbf{a}}_t = (\hat{v}_t, \hat{n}_t)$ pairs to conform $\hat{\mathbf{A}}_{obs}$. Finally, the intention head applies a second MLP Mixer to leverage the global context of the sequence of action representation features $\mathbf{X} = [x_1, \dots, x_N]$ and classifies the human intention \hat{I} in the observed video. The overall hierarchical multitask classifier obtains the top-low level information from past observation.

3.3. Intention-Conditioned Variational Autoencoder (I-CVAE)

Fig. 4 shows the Intention-Conditioned Variational Autoencoder (I-CVAE) architecture. Encoder-Decoder Transformer structure is used to leverage temporal dependencies between past actions to anticipate future sequence of human verb-noun pairs. Due to the nature of VAE, in the training scenario the model learns the conditional probabilistic distribution in the encoder through exploiting both observed actions \mathbf{A}_{obs} and future actions \mathbf{A}_{pred} . In its application, only the decoder block is used without any provided information of its future.

First, each action pair $\mathbf{a}_t = (v_t, n_t)$ of \mathbf{A}_{obs} is projected independently through learnable embedding matrices into $\mathbf{e}_t = (e_t^v, e_t^n)$, where e_t^v and e_t^n consist of d -dimensional embedding vector for verb and noun, respectively. Therefore, the verb embeddings are represented




Figure 3: **Detailed structure of H3M architecture.** First, pre-extracted padded features are fed into an Action MLP Mixer [32] to obtain clip level features (as green circles). These features are used (i) to obtain verb-noun pair through a fully-connected pair (action head); (ii) to obtain a video representation through the Intention MLP Mixer which is classified as an intention class. Definition of Mixer Layer is inherited from [32].

as $\mathbf{E}_{obs}^v = [e_1^v, \dots, e_N^v]$ for the observed sequence and $\mathbf{E}_{pred}^v = [e_N^v, \dots, e_{N+Z}^v]$ for the future sequence (same procedure for nouns, with \mathbf{E}_{obs}^n and \mathbf{E}_{pred}^n). To strengthen the dependencies between verb and noun as a common action, \mathbf{E}^v and \mathbf{E}^n are concatenated to $\mathbf{E} \in \mathbb{R}^{N \times 2d}$ in the Action Embedder. Intuitively, then, our Transformer Encoder-Decoder aims to reconstruct the \mathbf{E}_{pred} in the decoder given only the \mathbf{E}_{obs} and the intention label I as conditions.

Inspired by [24], a condition label I is projected into extra-learnable distribution parameters μ and Σ to inherit the action sequence representation after the Transformer Encoder. These tokens are prepended to the sequence of embedded representations of both the observed and forecasting actions to pool the time dimension. This whole embedded sequence are summed to a sinusoidal Positional Encoder (PE) and fed into the Transformer encoder, as can be seen in the top diagram of Fig. 4. The obtained parameters $\hat{\mu}$ and $\hat{\Sigma}$ are then used to define a latent distribution based on the reparametrization trick from VAE. We sample $z \in \mathbb{R}^{2d}$

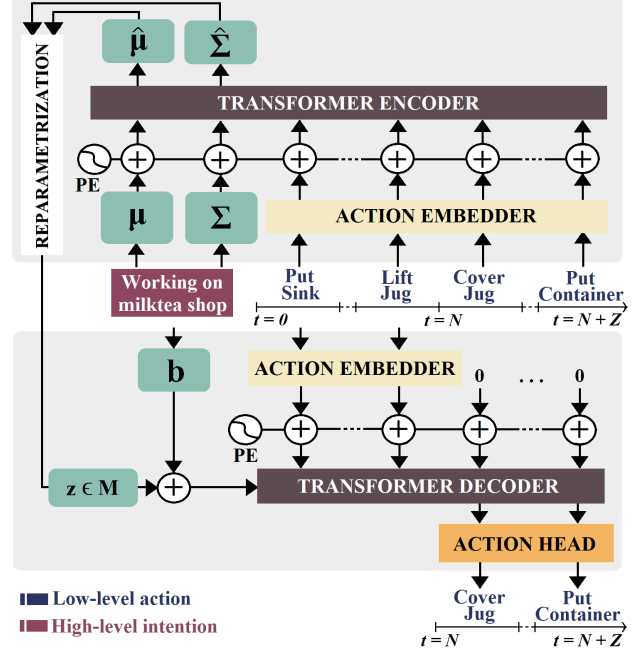


Figure 4: Detailed structure of I-CVAE architecture, illustrating the encoder (top) and decoder (bottom) of our Transformer-based CVAE model. Given a sequence of $N + Z$ actions and an Intention label, the encoder outputs distribution parameters ($\hat{\mu}$ and $\hat{\Sigma}$) that encode all sequence information. Inspired by [24], extra learnable parameters per intention are used (μ and Σ) to obtain $\hat{\mu}$ and $\hat{\Sigma}$ and sample the latent future action representation $z \in \mathbb{R}^M$, where M is the latent dimension of the Transformer. The decoder takes a latent vector z , the N observed actions and the intention I to output the representation sequence of Z actions to anticipate. I is used to determine the learnable b . Positional Encoder (PE) gives the time-component knowledge to the decoder. Finally, an Action Head compound by two fully-connected layers projects each action representation into a verb-noun pair.

from this latent distribution in the decoder.

During the decoding phase, as illustrated in the bottom diagram of Fig. 4, only the N observed actions \mathbf{A}_{obs} and intention label I are used to condition the future generation of next Z actions \mathbf{A}_{pred} . A learnable bias parameter b is used to shift the latent representation vector z to an intention-dependent space. The observed actions embeddings \mathbf{E}_{obs} are appended to a zero-vector for each Z actions to predict and summed with PE to form the input pattern to the Transformer Decoder. Moreover, the shift latent space is also fed in the transformer to condition the anticipation of $\hat{\mathbf{E}}_{pred}$ based on I . Finally, for each $t = [N+1, \dots, N+Z]$, the predicted future actions representations \hat{e}_t are fed into the Action Head, which applies a fully-connected layer to

classify \hat{e}_t into $\hat{\mathbf{a}}_t = (\hat{v}_t, \hat{n}_t)$.

To guide I-CVAE to decode plausible actions representations, an L2 reconstruction loss is applied between \mathbf{E}_{pred} and $\hat{\mathbf{E}}_{pred}$. Finally, to ensure the correct classification of the actions $\hat{\mathbf{a}}_t$, a weighted cross-entropy loss is used between \mathbf{A}_{pred} and $\hat{\mathbf{A}}_{pred}$.

4. Experiments

In this chapter, we present the Ego4D Dataset [13] and its proposed baselines and report and compare our quantitative results. Finally, we demonstrate the effectiveness of our approach through an ablation study.

4.1. Dataset - Ego4D

To appropriately anticipate the future, it is necessary to understand in detail the observed actions. Human Action Recognition (HAR) from video is itself a large computer vision research field, with increasing interest over egocentric view datasets [7, 13]. Ego4D [13] is the most extensive daily-life egocentric video dataset, currently available to research. It is notable that the Ego4D authors also provide pre-extracted features for each second of video. These features are obtained through SlowFast 8x8 model [10]. Due to the recent publication of the Ego4D dataset, only baselines results are provided for comparison. The Forecasting Benchmark from Ego4D (which includes LTA) consists of 120 hours of annotated videos from 53 different scenarios. The annotations provided contain 478 noun types and 115 verb types, with a total amount of 4756 action classes among training and validation set. Ego4D has a long-tailed distribution both for nouns and verbs categories, resulting in a high imbalanced dataset.

4.2. Evaluation Metrics and Baselines

Metrics. Following the proposed evaluation protocols in Ego4D LTA benchmark [13], we report the Edit Distance (ED) metric, as shown in Equation 1, computed as the Damerau-Levenshtein distance [8, 18] over sequences of predictions of verbs, nouns and actions. This metric accounts for small variations in the action sequence as predicting over long time horizons is subject to uncertainty. The proposed evaluation selects the best of K generated sequences according to the smallest metric. The lower the ED, the more similar the anticipated sequences to the reality.

$$\Delta E((\hat{n}_{z,k}^{(j)}, \hat{v}_{z,k}^{(j)})_{z=1}^Z, (n_z^{(j)}, v_z^{(j)})_{z=1}^Z) \quad (1)$$

Baselines. We compare our proposed framework with the baseline model proposed in Ego4D [13], where a trimmed video is used to predict K different plausible sequence of future actions. The baseline consists of (i) an encoder backbone for obtaining N video-level features based on SlowFast [10]; (ii) a transformer-based aggregation module that combines the previously extracted clip-level fea-

	VERB	NOUN	ACTION
Baseline [13]	0.739	0.780	0.943
Video+CLIP [9]	<u>0.74</u>	<u>0.77</u>	<u>0.94</u>
Ours	0.741	0.739	0.930

Table 1: Edit Distance (ED) comparison of long-term human action anticipation in Ego4D dataset. Scores are obtained directly from their reported results. Here, bold fonts denote the best result and underlines denote the second best result among all approaches.

tures through self-attention mechanisms; (iii) a Multi-Head decoder network with Z heads (one head per future time step) that generate sequence of future actions. K future action possible sequences are generated by sampling the predicted future action distribution K times. Moreover, we also compare our results with a CLIP-based model [9] that leverages multi-modality to encode the visual information of the observed actions. To the best of our knowledge, no other existing work has been reported to tackle the anticipation of very long-term actions (20 actions in advance).

4.3. Quantitative Results

We report our results for the LTA task in Table 1 based on the test set of the Ego4D LTA dataset. In this experiment, our framework predicted the $N = 6$ observed actions and the overall intention from the past, to anticipate $Z = 20$ future actions by generating $K = 5$ sequences. Our framework performs similar or better in all the LTA metrics defined: ED for verbs, nouns and overall actions. We claim that conditioning the generative model through the intention highly improves the ED for nouns, thus performing better in the overall action anticipation.

Due to our limited computational resources, training was performed independently for each module. Next we will describe the quantitative evaluation first for the H3M module and then for I-CVAE as stand-alone models. As LTA-Ego4D Forecasting benchmark is private, the ground truth from the testing set is not provided. Therefore, to validate our hypothesis, we perform an ablation study which is based on the results obtained from validation set.

4.3.1 H3M

Our model is able to recognize verb-noun pairs with similar performance as the baseline, as reported in Table 2. However, our model is trained on pre-extracted features \mathbf{F} for each clip, and not the image-based video clip \mathbf{V} . Due to the lower dimensionality of these features ($\mathbf{F} \in \mathbb{R}^{T \times 2304}$, with $T = 14$) several techniques were applied to avoid overfitting. We applied Gaussian noise injection (defined as N in Table 2) in the pre-extracted input features to improve robustness of the classifier. By adding multi-task approach and sharing inner layers among the tasks (M), we inherited

Accuracy	INTENTION		VERB		NOUN	
	TOP1	TOP5	TOP1	TOP5	TOP1	TOP5
Ours (M+I+N)	78.50	93.27	20.44	55.05	<u>19.32</u>	39.65
Ours (I+N)	<u>76.33</u>	94.06	20.12	55.14	18.64	39.95
Ours (N)	74.92	94.20	20.18	<u>56.02</u>	19.16	<u>40.02</u>
Baseline	-	-	22.06	56.90	20.92	41.40

Table 2: Performance of H3M with different training strategies, compared to the baseline using the accuracy metric. M: multitask surrogate loss (sharing weights). I: focal loss to solve class imbalance. N: noise injection. Here, bold fonts denote the best result and underlines denote the second best result among all approaches.

INTENTION	VERB		NOUN	
	TOP1	TOP5	TOP1	TOP5
Correct	20.13	54.32	19.48	41.58
Error	21.48	57.46	18.79	33.28

Table 3: Intention influence on action classification. Scores based on the accuracy metric.

implicit data augmentation: as there was a need of modelling a representation of three tasks (classifying intention, verb and noun), the model was forced to extract a better pattern from the inner shared layers to optimize results for each task. Finally, focal loss [19] modulation factor ($\beta = 0.99$) for the cross-entropy loss was applied to address the class imbalance (I) of intentions, verbs and nouns. The best results were obtained when an initial training was done to recognize the intention using multi-task approach, and then fine-tuning the action layers for the verb-noun recognition.

In Table 2, it is shown that applying all techniques (M+I+N) obtains the best Top-1 accuracy results of our framework, but our model slightly decreases its performance regarding Top-5 predictions. We claim that by dealing with imbalance of the Ego4D dataset through Focal Loss, the model takes more risk by attempting to predict less-frequent actions, which impacts the accuracy. Finally, our model is limited by the low-dimensionality of inputs, which are pre-extracted visual features as mentioned above, but is still able to approximate to the baseline.

We also compare the influence of predicting intention correctly to the accuracy of action classification, illustrated in Table 3. The results show that there is a significant and direct relationship between noun and intention. By conditioning the action-level prediction framework through the intention, it is shown that the performance in terms of noun prediction is improved. As our intention label is obtained, in this work, as the ‘scenario’ ground-truth of Ego4D, we argue that while verbs are more related to the human motion, nouns are coupled with the scenario, which refer to the environment of the human. Therefore, conditioning based on the environment has direct effect on the accuracy of noun-classification.

Moreover, we further evaluate the influence of the intention as the context for our H3M model. For a given intention label, only a few verbs and nouns are observed. Then, we define an out-of-context error as predicting a verb or a noun which is unseen in a given intention. For instance, if the model predicts the action ‘drive bike’ in a video where the human intention is ‘washing a dog’, we claim that the model has an out-of-context error. To determine the observed classes for a given intention, we create a bag of nouns and verbs used for each. If we predict a noun that cannot be found for the current intention bag, we determine this as an out-of-context error. We can observe from our experiments that no verb is out of context, while 14.56% of nouns are. These results strengthen the observation that (i) nouns and intentions have a significant relationship, (ii) verbs are less conditioned by intention and they are more related to the history of actions performed, as will be empirically shown in the next section.

4.3.2 I-CVAE

We evaluate the performance of our standalone I-CVAE model based on the ground-truth action and intention labels provided by Ego4d dataset. We report Edit Distance (ED) for the time horizon $Z = 20$ ($ED@Z = 20$) as our evaluation metric for both nouns and verbs.

First, we study the performance of I-CVAE under the variation of the number of observed actions N . Fig. 5 shows that the best results are obtained when leveraging $N = 4$. Observing longer past causes the model to focus on less-relevant cues, thus collapsing in less plausible action sequences. In addition, we observe that the noun performance is inversely affected by N , anticipating with less confident if N is lower, but not verbs. We claim that these results are caused by the different behaviours of verbs and nouns during a sequence of actions. On the one hand, verbs tend to change more frequently, but they usually appear in repeated patterns. For example, a sequence of ‘take-move-open-wash-close-put’ is typical when ‘washing’ an object. On the other hand, nouns are less variant as a human usually interacts with a certain object for longer duration. Reasoning that the next object that the human will interact with next depends on the environment as well as previous verbs.

This phenomenon is also shown in Fig. 6, where the ED for different N is represented at each of the $Z = 20$ time-horizons to forecast. Observing longer sequences ($N = 8$) increases the variance of observed actions, which leads the model to worsen the short-term anticipation. This stresses the need of exploiting temporal dependencies using our Transformer-based model. However, only accounting the last action performed ($N = 1$) caused implausible anticipations, mostly of nouns, due to the lack of context. Moreover, Fig. 6 also illustrates how the longer time horizon estimation affects each action component differently: nouns per-




Figure 5: Evaluation of I-CVAE trained based on different number of observed actions N .

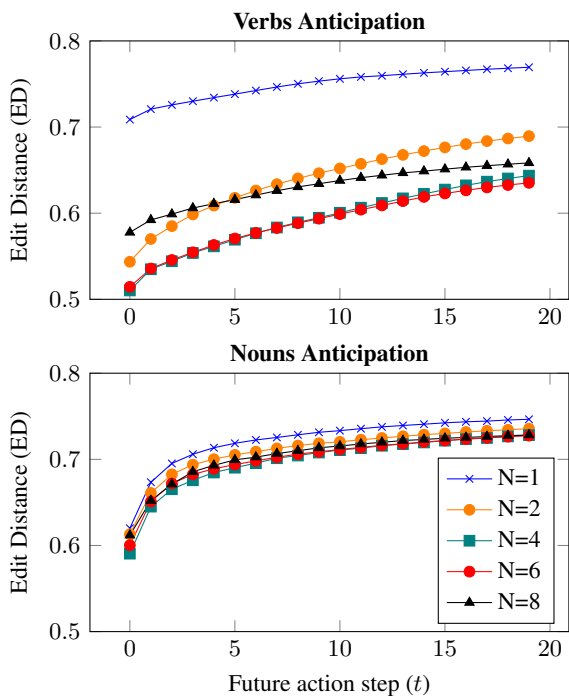


Figure 6: Edit Distance (ED) at each time step t depending on the number of observed actions N .

formance is reduced linearly with each timestep compared to the logarithmic reduction of verbs performance. To sum up, we have shown the two long-standing challenges in human long-term action anticipation: (i) learning the correct pattern of human behaviours, which is highly influenced by understanding the repeating patterns of verbs in a task; and (ii) exploring future tasks to perform in a given environment, based on leveraging visual cues to provide subsequent nouns that a human is going to interact with.

N	ED@20 verbs		ED@20 nouns		ED@20 Action	
	I✓	I✗	I✓	I✗	I✓	I✗
2	0.743	0.762	0.747	0.786	0.932	0.953
4	0.742	0.751	0.736	0.765	0.932	0.943
6	0.741	0.748	0.740	0.753	0.930	0.938
8	0.747	0.755	0.747	0.758	0.935	0.939

Table 4: Comparison of ED@20 for verbs, nouns and actions in the test set depending on the use of the intention (I✓) or not (I✗) in the end-to-end approach: I-CVAE uses the H3M predictions from the N observed clips.

4.3.3 H3M + I-CVAE

Finally, we investigate the performance of our whole framework based on the end-to-end evaluation. First, H3M classifies the actions and the intention from the observed clips. Then, based on these predictions, our I-CVAE model anticipates the $Z = 20$ actions in the future. In Table 4 we evaluate the LTA task in the test set for different N and under the influence of the intention as a condition. Results strengthen the importance of using the intention as a guideline to make more realistic anticipations. Therefore, we are able to confirm our hypothesis of using the intention as a high-level task knowledge to narrow the arbitrariness of the future. These results are also aligned with the different behaviours of verbs and nouns previously discussed: using $N = 4$ and $N = 6$ provides a sufficient context for the model to understand the task at hand. Observing higher number of actions causes the model to overfit.

5. Conclusion

In this work, we propose a two-module framework, which consists of a Hierarchical Multitask MLP Mixer (H3M) and Intention-Conditioned Variational Autoencoder (I-CVAE), to efficiently exploit human intention as a condition to increase the confidence of the model for long-term human action sequence anticipation. Our H3M module leverages hierarchy structure to classify the observed human actions and intentions. Then, by conditioning the future through our I-CVAE, our framework anticipates better the long-term action sequence. These results reinforce the importance of defining different level of abstractions for task contextualization. Experiments demonstrate that conditioning the model through the human intention has a direct effect when improving realistic anticipation. Finally, we conduct an extensive ablation study to investigate the differences between anticipating nouns and verbs. We describe these behaviors with different pattern observed in the Ego4D dataset, which are essential to tackle the Long-Term Anticipation (LTA). Then, our work outperforms the state-of-the-art in the Ego4D LTA task, mainly when forecasting nouns and actions. Our work ranked first in both CVPR@2022 and ECCV@2022 Ego4D LTA Challenge.

References

- [1] Yazan Abu Farha and Juergen Gall. Uncertainty-aware anticipation of activities. In *Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops*, pages 0–0, 2019.
- [2] Yazan Abu Farha, Alexander Richard, and Juergen Gall. When will you do what?-anticipating temporal occurrences of activities. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5343–5352, 2018.
- [3] Andrew Brock, Jeff Donahue, and Karen Simonyan. Large scale GAN training for high fidelity natural image synthesis. In *International Conference on Learning Representations*, 2019.
- [4] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020.
- [5] Jerome T Connor, R Douglas Martin, and Les E Atlas. Recurrent neural networks and robust time series prediction. *IEEE transactions on neural networks*, 5(2):240–254, 1994.
- [6] Michael Crawshaw. Multi-task learning with deep neural networks: A survey. *arXiv preprint arXiv:2009.09796*, 2020.
- [7] Dima Damen, Hazel Doughty, Giovanni Maria Farinella, Sanja Fidler, Antonino Furnari, Evangelos Kazakos, Davide Moltisanti, Jonathan Munro, Toby Perrett, Will Price, et al. The epic-kitchens dataset: Collection, challenges and baselines. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 43(11):4125–4141, 2020.
- [8] Fred J Damerau. A technique for computer detection and correction of spelling errors. *Communications of the ACM*, 7(3):171–176, 1964.
- [9] Srijan Das and Michael S. Ryoo. Video + clip baseline for ego4d long-term action anticipation, 2022.
- [10] Christoph Feichtenhofer, Haoqi Fan, Jitendra Malik, and Kaiming He. Slowfast networks for video recognition. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 6202–6211, 2019.
- [11] Harshayu Girase, Haiming Gang, Srikanth Malla, Jiachen Li, Akira Kanehara, Karttikeya Mangalam, and Chiho Choi. Loki: Long term and key intentions for trajectory prediction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9803–9812, 2021.
- [12] Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks, 2014.
- [13] Kristen Grauman, Andrew Westbury, Eugene Byrne, Zachary Chavis, Antonino Furnari, Rohit Girdhar, Jackson Hamburger, Hao Jiang, Miao Liu, Xingyu Liu, 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*, pages 18995–19012, 2022.
- [14] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems*, 33:6840–6851, 2020.
- [15] Zhiting Hu, Zichao Yang, Xiaodan Liang, Ruslan Salakhutdinov, and Eric P Xing. Toward controlled generation of text. In *International conference on machine learning*, pages 1587–1596. PMLR, 2017.
- [16] Will Kay, Joao Carreira, Karen Simonyan, Brian Zhang, Chloe Hillier, Sudheendra Vijayanarasimhan, Fabio Viola, Tim Green, Trevor Back, Paul Natsev, et al. The kinetics human action video dataset. *arXiv preprint arXiv:1705.06950*, 2017.
- [17] Diederik P Kingma, Max Welling, et al. An introduction to variational autoencoders. *Foundations and Trends in Machine Learning*, 12(4):307–392, 2019.
- [18] Vladimir I Levenshtein et al. Binary codes capable of correcting deletions, insertions, and reversals. In *Soviet physics doklady*, volume 10, pages 707–710. Soviet Union, 1966.
- [19] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 2980–2988, 2017.
- [20] Siyuan Brandon Loh, Debaditya Roy, and Basura Fernando. Long-term action forecasting using multi-headed attention-based variational recurrent neural networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2419–2427, 2022.
- [21] Nazanin Mehrasa, Akash Abdu Jyothi, Thibaut Durand, Jiawei He, Leonid Sigal, and Greg Mori. A variational auto-encoder model for stochastic point processes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3165–3174, 2019.
- [22] Yan Bin Ng and Basura Fernando. Forecasting future action sequences with attention: a new approach to weakly supervised action forecasting. *IEEE Transactions on Image Processing*, 29:8880–8891, 2020.
- [23] Hyun Soo Park, Jyh-Jing Hwang, Yedong Niu, and Jianbo Shi. Egocentric future localization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4697–4705, 2016.
- [24] Mathis Petrovich, Michael J Black, and Gül Varol. Action-conditioned 3d human motion synthesis with transformer vae. In *Proceedings of the IEEE/CVF International Conference on Computer Vision and Pattern Recognition*, pages 10985–10995, 2021.
- [25] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*, 2022.
- [26] Debaditya Roy and Basura Fernando. Action anticipation using latent goal learning. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 2745–2753, 2022.
- [27] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily Denton, Seyed Kamyar Seyed Ghasemipour, Burcu Karagol Ayan, S Sara Mahdavi, Rapha Gontijo Lopes, et al. Photorealistic text-to-image diffusion models with deep language understanding. *arXiv preprint arXiv:2205.11487*, 2022.
- [28] Fadime Sener, Dipika Singhania, and Angela Yao. Temporal aggregate representations for long-range video understand-

- ing. In *European Conference on Computer Vision*, pages 154–171. Springer, 2020.
- [29] Ziyang Song, Dongliang Wang, Nan Jiang, Zhicheng Fang, Chenjing Ding, Weihao Gan, and Wei Wu. Actformer: A gan transformer framework towards general action-conditioned 3d human motion generation. *arXiv preprint arXiv:2203.07706*, 2022.
- [30] Dídac Surís, Ruoshi Liu, and Carl Vondrick. Learning the predictability of the future. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12607–12617, 2021.
- [31] Guy Tevet, Sigal Raab, Brian Gordon, Yonatan Shafir, Daniel Cohen-Or, and Amit H Bermano. Human motion diffusion model. *arXiv preprint arXiv:2209.14916*, 2022.
- [32] Ilya O Tolstikhin, Neil Houlsby, Alexander Kolesnikov, Lucas Beyer, Xiaohua Zhai, Thomas Unterthiner, Jessica Yung, Andreas Steiner, Daniel Keysers, Jakob Uszkoreit, et al. Mlp-mixer: An all-mlp architecture for vision. *Advances in Neural Information Processing Systems*, 34:24261–24272, 2021.
- [33] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.
- [34] Takuma Yagi, Karttikeya Mangalam, Ryo Yonetani, and Yoichi Sato. Future person localization in first-person videos. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7593–7602, 2018.