



# Predicting Gaze in Egocentric Video by Learning Task-Dependent Attention Transition

Yifei Huang<sup>1</sup> , Minjie Cai<sup>1,2</sup> , Zhenqiang Li<sup>1</sup>, and Yoichi Sato<sup>1</sup> 

<sup>1</sup> The University of Tokyo, Tokyo, Japan

{hyf, cai-mj, lzq, ysato}@iis.u-tokyo.ac.jp

<sup>2</sup> Hunan University, Changsha, China

**Abstract.** We present a new computational model for gaze prediction in egocentric videos by exploring patterns in temporal shift of gaze fixations (attention transition) that are dependent on egocentric manipulation tasks. Our assumption is that the high-level context of how a task is completed in a certain way has a strong influence on attention transition and should be modeled for gaze prediction in natural dynamic scenes. Specifically, we propose a hybrid model based on deep neural networks which integrates task-dependent attention transition with bottom-up saliency prediction. In particular, the task-dependent attention transition is learned with a recurrent neural network to exploit the temporal context of gaze fixations, *e.g.* looking at a cup after moving gaze away from a grasped bottle. Experiments on public egocentric activity datasets show that our model significantly outperforms state-of-the-art gaze prediction methods and is able to learn meaningful transition of human attention.

**Keywords:** Gaze prediction · Egocentric video · Attention transition

## 1 Introduction

With the increasing popularity of wearable or action cameras in recording our life experience, egocentric vision [1], which aims at automatic analysis of videos captured from a first-person perspective [4, 6, 21], has become an emerging field in computer vision. In particular, as the camera wearer’s point-of-gaze in egocentric video contains important information about interacted objects and the camera wearer’s intent [17], gaze prediction can be used to infer important regions in images and videos to reduce the amount of computation needed in learning and inference of various analysis tasks [5, 7, 11, 36].

This paper aims to develop a computational model for predicting the camera wearer’s point-of-gaze from an egocentric video. Most previous methods have formulated gaze prediction as the problem of saliency detection, and computational models of visual saliency have been studied to find image regions that are likely to attract human attention. The saliency-based paradigm is reasonable because it is known that highly salient regions are strongly correlated with actual

gaze locations [27]. However, the saliency model-based gaze prediction becomes much more difficult in natural dynamic scenes, *e.g.* cooking in a kitchen, where high-level knowledge of the task has a strong influence on human attention.

In a natural dynamic scene, a person perceives the surrounding environment with a series of gaze fixations which point to the objects/regions related to the person’s interactions with the environment. It has been observed that the attention transition is deeply related to the task carried out by the person. Especially in object manipulation tasks, the high-level knowledge of an undergoing task determines a stream of objects or places to be attended successively and thus influences the transition of human attention. For example, to pour water from a bottle to a cup, a person always first looks at the bottle before grasping it and then change the fixation onto the cup during the action of pouring. Therefore, we argue that it is necessary to explore the task-dependent patterns in attention transition in order to achieve accurate gaze prediction.

In this paper, we propose a hybrid gaze prediction model that combines bottom-up visual saliency with task-dependent attention transition learned from successively attended image regions in training data. The proposed model is mainly composed of three modules. The first module generates saliency maps directly from video frames. It is based on a two-stream Convolutional Neural Network (CNN) which is similar to traditional bottom-up saliency prediction models. The second module is based on a recurrent neural network and a fixation state predictor which generates an attention map for each frame based on previously fixated regions and head motion. It is built based on two assumptions. Firstly, a person’s gaze tends to be located on the same object during each fixation, and a large gaze shift almost always occurs along with large head motion [23]. Secondly, patterns in the temporal shift between regions of attention are dependent on the performed task and can be learned from data. The last module is based on a fully convolutional network which fuses the saliency map and the attention map from the first two modules and generates a final gaze map, from which the final prediction of 2D gaze position is made.

Main contributions of this work are summarized as follows:

- We propose a new hybrid model for gaze prediction that leverages both bottom-up visual saliency and task-dependent attention transition.
- We propose a novel model for task-dependent attention transition that explores the patterns in the temporal shift of gaze fixations and can be used to predict the region of attention based on previous fixations.
- The proposed approach achieves state-of-the-art gaze prediction performance on public egocentric activity datasets.

## 2 Related Works

**Visual Saliency Prediction.** Visual saliency is a way to measure image regions that are likely to attract human attention and thus gaze fixation [2]. Traditional saliency models are based on the feature integration theory [35] telling that an image region with high saliency contains distinct visual features such as color,

intensity and contrast compared to other regions. After Itti *et al.*'s primary work [19] on a computational saliency model, various bottom-up computational models of visual saliency have been proposed such as a graph-based model [13] and a spectral clustering-based model [15]. Recent saliency models [16, 25, 26] leveraged a deep Convolutional Neural Network (CNN) to improve their performance. More recently, high-level context has been considered in deep learning-based saliency models. In [8, 31], class labels were used to compute the partial derivatives of CNN response with respect to input image regions to obtain a class-specific saliency map. In [40], a salient object is detected by combining global context of the whole image and local context of each image superpixel. In [29], region-to-word mapping in a neural saliency model was learned by using image captions as high-level input.

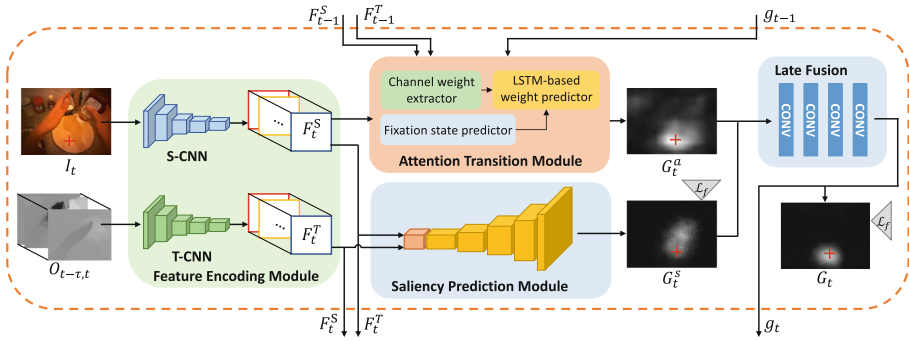
However, none of the previous methods explored the patterns in the transition of human attention inherent in a complex task. In this work, we propose to learn the task-dependent attention transition on how gaze shifts between different objects/regions to better model human attention in natural dynamic scenes.

**Egocentric Gaze Prediction.** Egocentric vision is an emerging research domain in computer vision which focuses on automatic analysis of egocentric videos recorded with wearable cameras. Egocentric gaze is a key component in egocentric vision which benefits various egocentric applications such as action recognition [11] and video summarization [36]. Although there is correlation between visually salient image regions and gaze fixation locations [27], it has been found that traditional bottom-up models for visual saliency is insufficient to model and predict human gaze in egocentric video [37]. Yamada *et al.* [38] presented a gaze prediction model by exploring the correlation between gaze and head motion. In their model, bottom-up saliency map is integrated with an attention map obtained based on camera rotation and translation to infer final egocentric gaze position. Li *et al.* [24] explored different egocentric cues like global camera motion, hand motion and hand positions to model egocentric gaze in hand manipulation activities. They built a graphical model and further combined the dynamic behaviour of gaze as latent variables to improve the gaze prediction. However, their model is dependent on predefined egocentric cues and may not generalize well to other activities where hands are not always involved. Recently, Zhang *et al.* [39] proposed the gaze anticipation problem in egocentric videos. In their work, a Generative Adversarial Network (GAN) based model is proposed to generate future frames from a current video frame, and gaze positions are predicted on the generated future frames based on a 3D-CNN based saliency prediction model.

In this paper, we propose a new hybrid model to predict gaze in egocentric videos, which combines bottom-up visual saliency with task-dependent attention transition. To the best of our knowledge, this is the first work to explore the patterns in attention transition for egocentric gaze prediction.

### 3 Gaze Prediction Model

In this section, we first give overview of the network architecture of the proposed gaze prediction model, and then explain the details of each component. The details of training the model are provided in the end.



**Fig. 1.** The architecture of our proposed gaze prediction model. The red crosses in the figure indicate ground truth gaze positions. (Color figure online)

#### 3.1 Model Architecture

Given consecutive video frames as input, we aim to predict a gaze position in each frame. To leverage both bottom-up visual saliency and task-dependent attention transition, we propose a hybrid model that (1) predicts a saliency map from each video frame, (2) predicts an attention map by exploiting temporal context of gaze fixations, and (3) fuses the saliency map and the attention map to output a final gaze map.

The model architecture is shown in Fig. 1. The feature encoding module is composed by a spatial Convolutional Neural Network (S-CNN) and a temporal Convolutional Neural Network (T-CNN), which extract latent representations from a single RGB image and stacked optical flow images respectively. The saliency prediction module generates a saliency map based on the extracted latent representation. The attention transition module generates an attention map based on previous gaze fixations and head motion. The late fusion module combines the results of saliency prediction and attention transition to generate a final gaze map. The details of each module will be given in the following part.

#### 3.2 Feature Encoding

At time  $t$ , the current video frame  $I_t$  and stacked optical flow  $O_{t-\tau,t}$  are fed into S-CNN and T-CNN to extract latent representations  $F_t^S = h^S(I_t)$  from the current RGB frame, and  $F_t^T = h^T(O_{t-\tau,t})$  from the stacked optical flow images for later use. Here  $\tau$  is fixed as 10 following [32].

The feature encoding network of S-CNN and T-CNN follows the base architecture of the first five convolutional blocks in Two Stream CNN [32], while omitting the final max pooling layer. We choose to use the output feature map of the last convolution layer from the 5-th convolutional group, i.e., *conv5\_3*. Further analysis of different choices of deep feature maps from other layers is described in Sect. 4.4.

### 3.3 Saliency Prediction Module

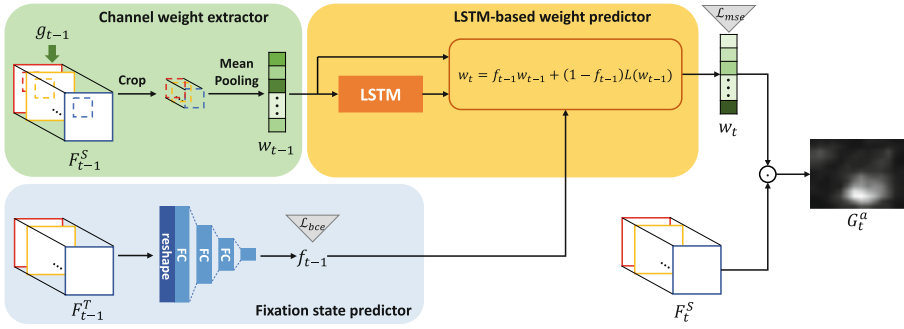
Biologically, human tends to gaze at an image region with high saliency, i.e., a region containing unique and distinctive visual features [34]. In the saliency prediction module of our gaze prediction model, we learn to generate a visual saliency map which reflects image regions that are likely to attract human gaze. We fuse the latent representations  $F_t^S$  and  $F_t^T$  as an input to a saliency prediction decoder (denoted as  $S$ ) to obtain the initial gaze prediction map  $G_t^s$  (Eq. 1). We use the “3dconv + pooling” method of [12] to fuse the two input feature streams. Since our task is different from [12], we modify the kernel sizes of the fusion part, which can be seen in detail in Sect. 3.7. The decoder outputs a visual saliency map with each pixel value within the range of  $[0, 1]$ . Details of the architecture of the decoder is described in Sect. 3.7. The equation for generating the visual saliency map is:

$$G_t^s = S(F_t^S, F_t^T) \quad (1)$$

However, a saliency map alone does not predict accurately where people actually look [37], especially in egocentric videos of natural dynamic scenes where the knowledge of a task has a strong influence on human gaze. To achieve better gaze prediction, high-level knowledge about a task, such as which object is to be looked at and manipulated next, has to be considered.

### 3.4 Attention Transition Module

During the procedure of performing a task, the task knowledge strongly influences the temporal transition of human gaze fixations on a series of objects. Therefore, given previous gaze fixations, it is possible to anticipate the image region where next attention occurs. However, direct modeling the object transition explicitly such as using object categories is problematic since a reliable and generic object detector is needed. Motivated by the fact that different channels of a feature map in top convolutional layers correspond well to spatial responses of different high-level semantics such as different object categories [9, 41], we represent the region that is likely to attract human attention by weighting each channel of the feature map differently. We train a Long Short Term Memory (LSTM) model [14] to predict a vector of channel weights which is used to predict the region of attention at next fixation. Figure 2 depicts the framework of the proposed attention transition module. The module is composed of a channel weight extractor (C), a fixation state predictor (P), and a LSTM-based weight predictor (L).



**Fig. 2.** The architecture of the attention transition module.

The channel weight extractor takes as input the latent representation  $F_{t-1}^S$  and the predicted gaze point  $g_{t-1}$  from the previous frame.  $F_{t-1}^S$  is in fact a stack of feature maps with spatial resolution  $14 \times 14$  and 512 channels. From each channel, we project the predicted gaze position  $g_{t-1}$  onto the  $14 \times 14$  feature map, and crop a fixed size area with height  $H_c$  and width  $W_c$  centered at the projected gaze position. We then average the value of the cropped feature map at each channel, obtaining a 512-dimensional vector of channel weight  $w_{t-1}$ :

$$w_{t-1} = C(F_{t-1}^S, g_{t-1}) \tag{2}$$

where  $C(\cdot)$  indicates the cropping and averaging operation,  $w_{t-1}$  is used as feature representation of the region of attention around the gaze point at frame  $t - 1$ .

The fixation state predictor takes the latent representation of  $F_{t-1}^T$  as input and outputs a probabilistic score of fixation state  $f_{t-1}^p = P(F_{t-1}^T) \in [0, 1]$ . Basically, the score tells how likely fixation is occurring in the frame  $t - 1$ . The fixation state predictor is composed by three fully connected layers followed by a final softmax layer to output a probabilistic score for gaze fixation state.

We use a LSTM to learn the attention transition by learning the transition of channel weights. The LSTM is trained based on a sequence of channel weight vectors extracted from images *at the boundaries of* all gaze fixation periods with ground-truth gaze points, *i.e.* we only extract one channel weight vector for each fixation to learn its transition between fixations. During testing, given a channel weight vector  $w_{t-1}$ , the trained LSTM outputs a channel weight vector  $L(w_{t-1})$  that represents the region of attention at next gaze fixation. We also consider the dynamic behavior of gaze and its influence on attention transition. Intuitively speaking, during a period of fixation, the region of attention tends to remain unchanged, and the attended region changes only when saccade happens.

Therefore, we compute the region of attention at current frame  $w_t$  as a linear combination of previous region of attention  $w_{t-1}$  and the anticipated region of attention at next fixation  $L(w_{t-1})$ , weighted by the predicted fixation probability  $f_{t-1}^p$ :

$$w_t = f_{t-1}^p \cdot w_{t-1} + (1 - f_{t-1}^p) \cdot L(w_{t-1}) \quad (3)$$

Finally, an attention map  $G_t^a$  is computed as the weighted sum of the latent representation  $F_t^S$  at frame  $t$  by using the resulting channel weight vector  $w_t$ :

$$G_t^a = \sum_{c=1}^n w_t[c] \cdot F_t^S[c] \quad (4)$$

where  $[c]$  denotes the  $c$ -th dimension/channel of  $w_t/F_t^S$  respectively.

### 3.5 Late Fusion

We build the late fusion module (LF) on top of the saliency prediction module and the attention transition module, which takes  $G_t^s$  and  $G_t^a$  as input and outputs the predicted gaze map  $G_t$ .

$$G_t = LF(G_t^s, G_t^a) \quad (5)$$

Finally, a predicted 2D gaze position  $g_t$  is given as the spatial coordinate of maximum value of  $G_t$ .

### 3.6 Training

For training gaze prediction in saliency prediction module and late fusion module, the ground truth gaze map  $\hat{G}$  is given by convolving an isotropic Gaussian over the measured gaze position in the image. Previous work used either Binary Cross-Entropy loss [22], or KL divergence loss [39] between the predicted gaze map and the ground truth gaze map for training neural networks. However, these loss functions do not work well with noisy gaze measurement. A measured gaze position is not static but continuously quivers in a small spatial range, even during fixation, and conventional loss functions are sensitive to small fluctuations of gaze. This observation motivates us to propose a new loss function, where the loss of pixels within small distance from the measured gaze position is down-weighted. More concretely, we modify the Binary Cross-Entropy loss function ( $\mathcal{L}_{bce}$ ) across all the  $N$  pixels with the weighting term  $1 + d_i$  as:

$$\mathcal{L}_f(G, \hat{G}) = -\frac{1}{N} \sum_{i=1}^N (1 + d_i) \{ \hat{G}[i] \cdot \log(G[i]) + (1 - \hat{G}[i]) \cdot \log(1 - G[i]) \} \quad (6)$$

where  $d_i$  is the euclidean distance between ground truth gaze position and the pixel  $i$ , normalized by the image width.

For training the fixation state predictor in the attention transition module, we treat the fixation prediction of each frame as a binary classification problem.

Thus, we use the Binary Cross-Entropy loss function for training the fixation state predictor. For training the LSTM-based weight predictor in the attention transition module, we use the mean squared error loss function across all the  $n$  channels:

$$\mathcal{L}_{mse}(w_t, \hat{w}_t) = \frac{1}{n} \sum_{i=1}^n (w_t[i] - \hat{w}_t[i])^2 \quad (7)$$

where  $w_t[i]$  denotes the  $i$ -th element of  $w_t$ .

### 3.7 Implementation Details

We describe the network structure and training details in this section. Our implementation is based on the PyTorch [28] library. The feature encoding module follows the base architecture of the first five convolutional blocks (*conv1* ~ *conv5*) of VGG16 [33] network. We remove the last max-pooling layer in the 5-th convolutional block. We initialize these convolutional layers using pre-trained weights on ImageNet [10]. Following [32], since the input channels of T-CNN is changed to 20, we average the weights of the first convolution layer of T-CNN part. The saliency prediction module is a set of 5 convolution layer groups following the inverse order of VGG16 while changing all max pooling layers into upsampling layers. We change the last layer to output 1 channel and add sigmoid activation on top. Since the input of the saliency prediction module contains latent representations from both S-CNN and T-CNN, we use a 3d convolution layer (with a kernel size of  $1 \times 3 \times 3$ ) and a 3d pooling layer (with a kernel size of  $2 \times 1 \times 1$ ) to fuse the inputs. Thus, the input and output sizes are all  $224 \times 224$ . The fixation state predictor is a set of fully connected (FC) layers, whose output sizes are 4096,1024,2 sequentially. The LSTM is a 3-layer LSTM whose input and output sizes are both 512. The late fusion module consists of 4 convolution layers followed by sigmoid activation. The first three layers have a kernel size of  $3 \times 3$ , 1 zero padding, and output channels 32,32,8 respectively, and the last convolution layer has a kernel size of 1 with a single output channel. We empirically set both the height  $H_c$  and width  $W_c$  for cropping the latent representations to be 3.

The whole model is trained using Adam optimizer [20] with its default settings. We fix the learning rate as  $1e-7$  and first train the saliency prediction module for 5 epochs for the module to converge. We then fix the saliency prediction module and train the LSTM-based weight predictor and the fixation state predictor in the attention transition module. Learning rates for other modules in our framework are all fixed as  $1e-4$ . After training the attention transition module, we fix the saliency prediction and the attention transition module to train the late fusion module in the end.

## 4 Experiments

We first evaluate our gaze prediction model on two public egocentric activity datasets (**GTEA Gaze** and **GTEA Gaze Plus**). We compare the proposed



model with other state-of-the-art methods and provide detailed analysis of our model through ablation study and visualization of outputs of different modules. Furthermore, to examine our model’s ability in learning attention transition, we visualize output of the attention transition module on a newly collected test set from GTEA Gaze Plus dataset (denoted as **GTEA-sub**).

#### 4.1 Datasets

We introduce the datasets used for gaze prediction and attention transition.

**GTEA Gaze** contains 17 video sequences of kitchen tasks performed by 14 subjects. Each video clip lasts for about 4 min with the frame rate of 15 fps and an image resolution of  $480 \times 640$ . We use videos 1, 4, 6–22 as a training set and the rest as a test set as in Yin *et al.* [24].

**GTEA Gaze Plus** contains 37 videos with the frame rate of 24 fps and an image resolution of  $960 \times 1280$ . In this dataset each of the 5 subjects performs 7 meal preparation activities in a more natural environment. Each video clip is 10 to 15 min long on average. Similarly to [24], gaze prediction accuracy is evaluated with 5-fold cross validation across all 5 subjects.

**GTEA-sub** contains 227 video frames selected from the sampled frames of GTEA Gaze Plus dataset. Each selected frame is not only under a gaze fixation, but also contains the object (or region) that is to be attended at the next fixation. We manually draw bounding boxes on those regions by inspecting future frames. The dataset is used to examine whether or not our model trained on GTEA Gaze Plus (excluding GTEA-sub) has successfully learned the task-dependent attention transition.

#### 4.2 Evaluation Metrics

We use two standard evaluation metrics for gaze prediction in egocentric videos: Area Under the Curve (AUC) [3] and Average Angular Error (AAE) [30]. **AUC** is the area under a curve of true positive rate versus false positive rate for different thresholds on the predicted gaze map. It is a commonly used evaluation metric in saliency prediction. **AAE** is the average angular distance between the predicted and the ground truth gaze positions.

#### 4.3 Results on Gaze Prediction

**Baselines.** We use the following baselines for gaze prediction:

- *Saliency prediction algorithms:* We compare our method with several representative saliency prediction methods. More specifically, we used Itti’s model [18], Graph Based Visual Saliency (GBVS [13]), and a deep neural network based saliency model as the current state of the art (SALICON [16]).
- *Center bias:* Since egocentric gaze data is observed to have a strong center bias, we use the image center as the predicted gaze position as in [24].

- *Gaze prediction algorithms*: We also compare our method with two state-of-the-art gaze prediction methods: the egocentric cue-based method (Yin *et al.* [24]), and the GAN-based method (DFG [39]). Note that although the goal of [39] is gaze anticipation in future frames, it also reported gaze prediction in the current frame.

**Performance Comparison.** The quantitative results of different methods on two datasets are given in Table 1. Our method significantly outperforms all baselines on both datasets, particularly on the AAE score. Although there is only a small improvement on the AUC score, it can be seen that previous method of DFG [39] has already achieved quite high score and the space of improvement is limited. Besides, we have observed from experiments that high AUC score does not necessarily mean high performance of gaze prediction. The overall performance on GTEA Gaze is lower than that on GTEA Gaze Plus. The reason might be that the number of training samples in GTEA Gaze is smaller and over 25% of ground truth gaze measurements are missing. It is also interesting to see that the center bias outperforms all saliency-based methods and works only slightly worse than Yin *et al.* [24] on GTEA Gaze Plus, which demonstrates the strong spatial bias of gaze in egocentric videos.

**Table 1.** Performance comparison of different methods for gaze prediction on two public datasets. Higher AUC (or lower AAE) means higher performance.

Metrics	GTEA Gaze Plus		GTEA Gaze	
	AAE (deg)	AUC	AAE (deg)	AUC
Itti <i>et al.</i> [18]	19.9	0.753	18.4	0.747
GBVS [13]	14.7	0.803	15.3	0.769
SALICON [16]	15.6	0.818	16.5	0.761
Center bias	8.6	0.819	10.2	0.789
Yin <i>et al.</i> [24]	7.9	0.867	8.4	0.878
DFG [39]	6.6	0.952	10.5	0.883
Our full model	<b>4.0</b>	<b>0.957</b>	<b>7.6</b>	<b>0.898</b>


**Ablation Study.** To study the effect of each module of our model, and the effectiveness of our modified binary cross entropy loss (Eq. 6), we conduct an ablation study and test each component on both GTEA Gaze Plus and GTEA Gaze datasets. Our baselines include: (1) single-stream saliency prediction with binary cross entropy loss (**S-CNN bce** and **T-CNN bce**), (2) single-stream saliency prediction with our modified bce loss (**S-CNN** and **T-CNN**), (3) two-stream saliency prediction with bce loss (**SP bce**), (4) two-stream input saliency prediction with our modified bce loss (**SP**), (5) the attention transition module (**AT**), and our full model.

Table 2 shows the results of the ablation study. The comparison of the same framework with different loss functions shows that our modified bce loss function is more suitable for the training of gaze prediction in egocentric video. The SP module performs better than either of the single-stream saliency prediction (S-CNN and T-CNN), indicating that both spatial and temporal information are needed for accurate gaze prediction. It is important to see that the AT module performs competitively or better than the SP module. This validates our claim that learning task-dependent attention transition is important in egocentric gaze prediction. More importantly, our full model outperforms all separate components by a large margin, which confirms that the bottom-up visual saliency and high-level task-dependent attention are complementary cues to each other and should be considered together in modeling human attention.

**Table 2.** Results of ablation study


Metrics	GTEA Gaze plus		GTEA Gaze	
	AAE (deg)	AUC	AAE (deg)	AUC
S-CNN (bce)	5.61	0.893	9.90	0.854
T-CNN (bce)	6.15	0.906	10.08	0.854
S-CNN	5.57	0.905	9.72	0.857
T-CNN	6.07	0.906	9.6	0.859
SP (bce)	5.63	0.918	9.53	0.860
SP	5.52	0.928	9.43	0.861
AT	5.02	0.940	9.51	0.857
Our full model	<b>4.05</b>	<b>0.957</b>	<b>7.58</b>	<b>0.898</b>

**Visualization.** Figure 3 shows qualitative results of our model. Group (1a, 1b) shows a typical gaze shift: the camera wearer shifts his attention to the pan after turning on the oven. SP fails to find the correct gaze position in (1b) only from visual features of the current frame. Since AT exploits the high-level temporal context of gaze fixations, it successfully predicts the region to be on the pan. Group (2a, 2b) demonstrates a “put” action: the camera wearer first looks at the target location, then puts the can to that location. It is interesting that AT has learned the camera wearer’s intention, and predicts the region at the target location rather than the more salient hand region in (2a). In group (3a, 3b), the camera wearer searches for a spatula after looking at the pan. Again, AT has learned this context which leads to more accurate gaze prediction than SP. Finally, group (4a, 4b) shows that SP and AT are complementary to each other. While AT performs better in (4a), and SP performs better in (4b), the full model combines the merits of both AT and SP to make better prediction. Overall, these results demonstrate that the attention transition plays an important role in improving gaze prediction accuracy.



**Fig. 3.** Visualization of predicted gaze maps from our model. Each group contains two images from two consecutive fixations, where a happens before b. We show the output heatmap from the saliency prediction module (**SP**) and the attention transition module (**AT**) as well as our full model. The ground truth gaze map (the rightmost column) is obtained by convolving an isotropic Gaussian on the measured gaze point.

**Cross Task Validation.** To examine how the task-dependent attention transition learned in our model can generalize to different tasks under same (kitchen) scene, we perform a cross validation across the 7 different meal preparation tasks on GTEA Gaze Plus dataset. We consider the following experiment settings:



**Fig. 4.** AUC and AAE scores of cross task validation. Five different experiment settings (explained in the text below) are compared to study the differences of attention transition in different tasks.

- **SP:** The saliency prediction module is treated as a generic component and trained on a separate subset of the dataset. We also use it as a baseline for studying the performance variation of different settings.
- **AT\_d:** The attention transition module is trained and validated under different tasks. Average performance of 7-fold cross validation is reported.
- **AT\_s:** The attention transition module is trained and validated on two splits of the same task. Average performance of 7 tasks is reported.
- **SP+AT\_d:** The late fusion on top of **SP** and **AT\_d**.
- **SP+AT\_s:** The late fusion on top of **SP** and **AT\_s**.


Quantitative results of different settings are shown in Fig. 4. Both AUC and AAE scores show the same performance trend with different settings. **AT\_d** works worse than **SP**, while **AT\_s** outperforms **SP**. This is probably due to the differences of gaze behavior contained in different tasks. However, **SP+AT\_d** with the late fusion module can still improve the performance compared with **SP** and **AT\_s**, even with the context learned from different tasks.

#### 4.4 Examination of the Attention Transition Module

We further demonstrate that our attention transition module is able to learn meaningful transition between adjacent gaze fixations. This ability has important applications in computer-aided AR system, such as implying a person where to look next in performing a complex task. We conduct a new experiment on the GTEA-sub dataset (as introduced in Sect. 4.1) to test the attention transition module of our model. Since here we focus on the module’s ability of attention transition, we omit the fixation state predictor in the module and assume the output of the fixation state predictor as  $f_t = 0$  in the test frame. The module

takes  $w_t$  calculated from the region of current fixation as input and outputs an attention map on the same frame which represents the predicted region of the next fixation. We extract a 2D position from the maximum value of the predicted heatmap and calculate its rate of falling within the annotated bounding box as the transition accuracy.

We conduct experiments based on different latent representations extracted from the convolutional layer: *conv5\_1*, *conv5\_2*, and *conv5\_3* of S-CNN. The accuracy based on the above three convolutional layers are 71.7%, 83.0%, and 86.8% respectively, while the accuracy based on random position is 10.7%. We also tried using random channel weight as the output of channel weight predictor to compute attention map based on the latent representation of *conv5\_3*, and the accuracy is 9.4%. This verifies that our model can learn meaningful attention transition of the performed task. Figure 5 shows some qualitative results of the attention transition module learned based on layer *conv5\_3*. It can be seen that the attention transition module can successfully predict the image region of next fixation.



**Fig. 5.** Qualitative results of attention transition. We visualize the predicted heatmap on the current frame, together with the current gaze position (red cross) and ground truth bounding box of the object/region of the next fixation (yellow box). (Color figure online)

## 5 Conclusion and Future Work

This paper presents a hybrid model for gaze prediction in egocentric videos. Task-dependent attention transition is learned to predict human attention from previous fixations by exploiting the temporal context of gaze fixations. The task-dependent attention transition is further integrated with a CNN-based saliency model to leverage the cues from both bottom-up visual saliency and high-level attention transition. The proposed model achieves state-of-the-art performance in two public egocentric datasets.

As for our future work, we plan to explore the task-dependent gaze behavior in a broader scale, *i.e.* tasks in an office or in a manufacturing factory, and to study the generalizability of our model in different task domains.

**Acknowledgments.** This work was supported by JST CREST Grant Number JPMJCR14E1, Japan.

## References

1. Betancourt, A., Morerio, P., Regazzoni, C.S., Rauterberg, M.: The evolution of first person vision methods: a survey. *IEEE Trans. Circ. Syst. Video Technol.* **25**(5), 744–760 (2015)
2. Borji, A., Itti, L.: State-of-the-art in visual attention modeling. *IEEE Trans. Pattern Anal. Mach. Intell.* **35**(1), 185–207 (2013)
3. Borji, A., Tavakoli, H.R., Sihite, D.N., Itti, L.: Analysis of scores, datasets, and models in visual saliency prediction. In: *ICCV* (2013)
4. Cai, M., Kitani, K.M., Sato, Y.: A scalable approach for understanding the visual structures of hand grasps. In: *ICRA* (2015)
5. Cai, M., Kitani, K.M., Sato, Y.: Understanding hand-object manipulation with grasp types and object attributes. In: *Robotics: Science and Systems* (2016)
6. Cai, M., Kitani, K.M., Sato, Y.: An ego-vision system for hand grasp analysis. *IEEE Trans. Hum. Mach. Syst.* **47**(4), 524–535 (2017)
7. Cai, M., Lu, F., Gao, Y.: Desktop action recognition from first-person point-of-view. *IEEE Trans. Cybern.* **PP**(99), 1–13 (2018)
8. Cao, C., et al.: Look and think twice: capturing top-down visual attention with feedback convolutional neural networks. In: *ICCV* (2015)
9. Chen, L., et al.: SCA-CNN: spatial and channel-wise attention in convolutional networks for image captioning. In: *CVPR* (2017)
10. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: a large-scale hierarchical image database. In: *CVPR* (2009)
11. Fathi, A., Li, Y., Rehg, J.M.: Learning to recognize daily actions using gaze. In: Fitzgibbon, A., Lazebnik, S., Perona, P., Sato, Y., Schmid, C. (eds.) *ECCV 2012*. LNCS, vol. 7572, pp. 314–327. Springer, Heidelberg (2012). [https://doi.org/10.1007/978-3-642-33718-5\\_23](https://doi.org/10.1007/978-3-642-33718-5_23)
12. Feichtenhofer, C., Pinz, A., Zisserman, A.: Convolutional two-stream network fusion for video action recognition. In: *CVPR* (2016)
13. Harel, J., Koch, C., Perona, P.: Graph-based visual saliency. In: *NIPS* (2007)
14. Hochreiter, S., Schmidhuber, J.: Long short-term memory. *Neural Comput.* **9**(8), 1735–1780 (1997)
15. Hou, X., Harel, J., Koch, C.: Image signature: highlighting sparse salient regions. *IEEE Trans. Pattern Anal. Mach. Intell.* **34**(1), 194–201 (2012)
16. Huang, X., Shen, C., Boix, X., Zhao, Q.: SALICON: reducing the semantic gap in saliency prediction by adapting deep neural networks. In: *ICCV* (2015)
17. Huang, Y., Cai, M., Kera, H., Yonetani, R., Higuchi, K., Sato, Y.: Temporal localization and spatial segmentation of joint attention in multiple first-person videos. In: *ICCV Workshop* (2017)
18. Itti, L., Koch, C.: A saliency-based search mechanism for overt and covert shifts of visual attention. *Vis. Res.* **40**(10–12), 1489–1506 (2000)
19. Itti, L., Koch, C., Niebur, E.: A model of saliency-based visual attention for rapid scene analysis. *IEEE Trans. Pattern Anal. Mach. Intell.* **20**(11), 1254–1259 (1998)
20. Kingma, D.P., Ba, J.: Adam: a method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014)
21. Kitani, K.M., Okabe, T., Sato, Y., Sugimoto, A.: Fast unsupervised ego-action learning for first-person sports videos. In: *CVPR* (2011)

22. Kuen, J., Wang, Z., Wang, G.: Recurrent attentional networks for saliency detection. In: CVPR (2016)
23. Land, M.F.: The coordination of rotations of the eyes, head and trunk in saccadic turns produced in natural situations. *Exp. Brain Res.* **159**(2), 151–160 (2004)
24. Li, Y., Fathi, A., Rehg, J.M.: Learning to predict gaze in egocentric video. In: ICCV (2013)
25. Lin, Y., Kong, S., Wang, D., Zhuang, Y.: Saliency detection within a deep convolutional architecture. In: AAAI Workshops (2014)
26. Pan, J., Sayrol, E., Giro-i Nieto, X., McGuinness, K., O'Connor, N.E.: Shallow and deep convolutional networks for saliency prediction. In: CVPR (2016)
27. Parkhurst, D., Law, K., Niebur, E.: Modeling the role of salience in the allocation of overt visual attention. *Vis. Res.* **42**(1), 107–123 (2002)
28. Paszke, A., et al.: Automatic differentiation in PyTorch (2017)
29. Ramanishka, V., Das, A., Zhang, J., Saenko, K.: Top-down visual saliency guided by captions. In: CVPR (2017)
30. Riche, N., Duvinage, M., Mancas, M., Gosselin, B., Dutoit, T.: Saliency and human fixations: state-of-the-art and study of comparison metrics. In: ICCV (2013)
31. Simonyan, K., Vedaldi, A., Zisserman, A.: Deep inside convolutional networks: visualising image classification models and saliency maps. arXiv preprint [arXiv:1312.6034](https://arxiv.org/abs/1312.6034) (2013)
32. Simonyan, K., Zisserman, A.: Two-stream convolutional networks for action recognition in videos. In: NIPS (2014)
33. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. arXiv preprint [arXiv:1409.1556](https://arxiv.org/abs/1409.1556) (2014)
34. Sugano, Y., Matsushita, Y., Sato, Y.: Appearance-based gaze estimation using visual saliency. *IEEE Trans. Pattern Anal. Mach. Intell.* **35**(2), 329–341 (2013)
35. Treisman, A.M., Gelade, G.: A feature-integration theory of attention. *Cogn. Psychol.* **12**(1), 97–136 (1980)
36. Xu, J., Mukherjee, L., Li, Y., Warner, J., Rehg, J.M., Singh, V.: Gaze-enabled egocentric video summarization via constrained submodular maximization. In: CVPR (2015)
37. Yamada, K., Sugano, Y., Okabe, T., Sato, Y., Sugimoto, A., Hiraki, K.: Can saliency map models predict human egocentric visual attention? In: Koch, R., Huang, F. (eds.) ACCV 2010. LNCS, vol. 6468, pp. 420–429. Springer, Heidelberg (2011). [https://doi.org/10.1007/978-3-642-22822-3\\_42](https://doi.org/10.1007/978-3-642-22822-3_42)
38. Yamada, K., Sugano, Y., Okabe, T., Sato, Y., Sugimoto, A., Hiraki, K.: Attention prediction in egocentric video using motion and visual saliency. In: Ho, Y.-S. (ed.) PSIVT 2011. LNCS, vol. 7087, pp. 277–288. Springer, Heidelberg (2011). [https://doi.org/10.1007/978-3-642-25367-6\\_25](https://doi.org/10.1007/978-3-642-25367-6_25)
39. Zhang, M., Teck Ma, K., Hwee Lim, J., Zhao, Q., Feng, J.: Deep future gaze: gaze anticipation on egocentric videos using adversarial networks. In: CVPR (2017)
40. Zhao, R., Ouyang, W., Li, H., Wang, X.: Saliency detection by multi-context deep learning. In: CVPR (2015)
41. Zhou, B., Khosla, A., Lapedriza, A., Oliva, A., Torralba, A.: Learning deep features for discriminative localization. In: CVPR (2016)