

MemX: An Attention-Aware Smart Eyewear System for Personalized Moment Auto-capture

YUHU CHANG, School of Computer Science, Fudan University, China and Shanghai Key Laboratory of Data Science, Fudan University, China

YINGYING ZHAO*, School of Computer Science, Fudan University, China and Shanghai Key Laboratory of Data Science, Fudan University, China

MINGZHI DONG, School of Computer Science, Fudan University, China and Shanghai Key Laboratory of Data Science, Fudan University, China

YUJIANG WANG, Department of Computing, Imperial College London, United Kingdom

YUTIAN LU, School of Computer Science, Fudan University, China and Shanghai Key Laboratory of Data Science, Fudan University, China

QIN LV, University of Colorado Boulder, United States

ROBERT P. DICK, University of Michigan, United States

TUN LU, School of Computer Science, Fudan University, China and Shanghai Key Laboratory of Data Science, Fudan University, China

NING GU, School of Computer Science, Fudan University, China and Shanghai Key Laboratory of Data Science, Fudan University, China

LI SHANG*, School of Computer Science, Fudan University, China and Shanghai Key Laboratory of Data Science, Fudan University, China

*Corresponding authors

Authors' addresses: Yuhu Chang, yhchang14@fudan.edu.cn, School of Computer Science, Fudan University, Shanghai, China, 200438, Shanghai Key Laboratory of Data Science, Fudan University, Shanghai, China, 200438; Yingying Zhao, yingyingzhao@fudan.edu.cn, School of Computer Science, Fudan University, Shanghai, China, 200438, Shanghai Key Laboratory of Data Science, Fudan University, Shanghai, China, 200438; Mingzhi Dong, mingzhidong@gmail.com, School of Computer Science, Fudan University, Shanghai, China, 200438, Shanghai Key Laboratory of Data Science, Fudan University, Shanghai, China, 200438; Yujiang Wang, yujiang.wang14@imperial.ac.uk, Department of Computing, Imperial College London, London, United Kingdom; Yutian Lu, 20210240098@fudan.edu.cn, School of Computer Science, Fudan University, Shanghai, China, 200438, Shanghai Key Laboratory of Data Science, Fudan University, Shanghai, China, 200438; Qin Lv, qin.lv@colorado.edu, University of Colorado Boulder, Boulder, Colorado, United States; Robert P. Dick, dickrp@umich.edu, University of Michigan, Ann Arbor, Michigan, United States; Tun Lu, lutun@fudan.edu.cn, School of Computer Science, Fudan University, Shanghai, China, 200438, Shanghai Key Laboratory of Data Science, Fudan University, Shanghai, China, 200438; Ning Gu, ninggu@fudan.edu.cn, School of Computer Science, Fudan University, Shanghai, China, 200438, Shanghai Key Laboratory of Data Science, Fudan University, Shanghai, China, 200438; Li Shang, lishang@fudan.edu.cn, School of Computer Science, Fudan University, Shanghai, China, 200438, Shanghai Key Laboratory of Data Science, Fudan University, Shanghai, China, 200438.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2021 Association for Computing Machinery.

2474-9567/2021/6-ART56 \$15.00

<https://doi.org/10.1145/3463509>

This work presents *MemX*: a biologically-inspired attention-aware eyewear system developed with the goal of pursuing the long-awaited vision of a personalized visual Memex. *MemX* captures human visual attention on the fly, analyzes the salient visual content, and records moments of personal interest in the form of compact video snippets. Accurate attentive scene detection and analysis on resource-constrained platforms is challenging because these tasks are computation and energy intensive. We propose a new temporal visual attention network that unifies human visual attention tracking and salient visual content analysis. Attention tracking focuses computation-intensive video analysis on salient regions, while video analysis makes human attention detection and tracking more accurate. Using the YouTube-VIS dataset and 30 participants, we experimentally show that *MemX* significantly improves the attention tracking accuracy over the eye-tracking-alone method, while maintaining high system energy efficiency. We have also conducted 11 in-field pilot studies across a range of daily usage scenarios, which demonstrate the feasibility and potential benefits of *MemX*.

CCS Concepts: • **Human-centered computing** → **Mobile devices**.

Additional Key Words and Phrases: Smart glasses, attention-aware, energy-efficient, eye tracking, video instance segmentation

ACM Reference Format:

Yuhu Chang, Yingying Zhao, Mingzhi Dong, Yujiang Wang, Yutian Lu, Qin Lv, Robert P. Dick, Tun Lu, Ning Gu, and Li Shang. 2021. MemX: An Attention-Aware Smart Eyewear System for Personalized Moment Auto-capture. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 5, 2, Article 56 (June 2021), 23 pages. <https://doi.org/10.1145/3463509>

1 INTRODUCTION

First presented by Vannevar Bush then restated by Jim Grey, Personal Memex paints the vision of the capability to capture and record personal experiences in daily life [6]. Rapid improvement in mobile camera and signal processing technology has put us on the verge of realizing this vision. In particular, cameras in wearable forms, such as GoPro, Google Clips, and Spectacles from Snapchat, enable people to take photos and record videos anywhere, anytime.

However, challenges remain. Personalization is the core value of Personal Memex. Videos and photos captured by devices must match people’s personal interests to remain relevant and valuable. A fundamental limitation of existing wearable cameras is the inability to automatically detect and selectively record the visual content of personal interest. Most existing wearable cameras, such as GoPro and Spectacles, require manual intervention for content recording. Empowered by machine vision technologies, recently released Google Clips performs auto-capture if interesting scenes are detected. However, scene interest is estimated solely by the built-in machine vision algorithm: instead of being personalized to the user, it provides a one-size-fits-all solution. As noted by a reputable 3rd-party reviewer, quoted here “Clips’ standout feature – artificially intelligent auto-capture – is too unpredictable, keeping it from being truly satisfying” [26].

Recent advances in eye-tracking technologies offer the capability to measure human visual attention [3, 9]. Based on physiology and psychology research [5], human attention serves as an information filter and prioritization strategy [5, 28], which selectively determines the allocation of cognitive resources. It operates as a gate that connects human inner consciousness with the outer world [37]. Via attention gate, human selects the prioritized information regarding the most interesting areas or objects from a huge amount of unstructured perceived information [28]. Physiologically, humans direct their gaze to objects of interest [35]. Human eye movement consists of several dynamic patterns, including saccade, smooth pursuit, vergence, and fixation [42]; saccade and smooth pursuit account for most eye movements. Saccades are conjugate eye motions, where gaze is rapidly shifted from one location or object to another. They can be voluntarily directed or involuntarily triggered by visual stimuli. Smooth pursuit is smooth eye motion allowing the gaze to continuously follow an object of interest. Saccades and smooth pursuit may alternate [4]. They can be measured and used to produce visual attention time series, as shown in Fig. 1. Eyewear systems may include, or be retrofitted with, eye-tracking technologies, e.g., inward-facing eye cameras, to measure human regional gaze fixation within the field of view, and infer the focused region of human visual attention.

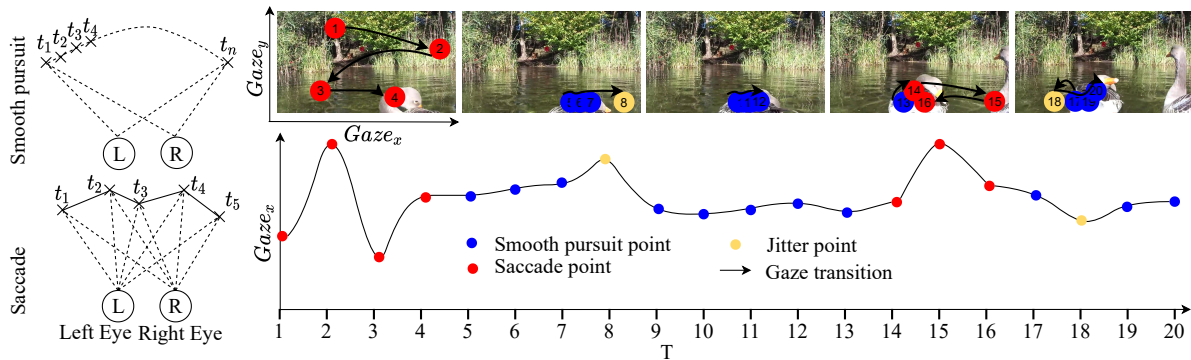


Fig. 1. (Left column) shows two eye movement types: saccade (top) and smooth pursuit (down). (Right column) An example of human gaze trajectory profile, which is a dynamic and noisy process. $Gaze_x$ and $Gaze_y$ denote the gaze position along the x axis and y axis, respectively.

Despite its promising potential, eye tracking suffers from limitations that complicate practical use. First, eye tracking does not directly detect the object of human visual interest. Instead, it estimates human visual regional focus by measuring the angular direction of eye gaze relative to the field of view of the eyewear. Measurement noise, e.g., angular offset and head motion, can easily produce incorrect visual attention measurements. In addition, eye movement is noisy. Human gaze frequently experiences transient jitter: gaze may drift away from the object of interest due to other visual stimuli or noise (As shown in Fig. 1). Therefore, eye tracking alone is inadequate for reliable attention tracking. A potential remedy to this problem is to augment eye tracking with visual scene analysis using video analytics techniques, e.g., identifying the salient object of human visual focus using instance segmentation [49] and analyzing the spatial visual context using object detection and recognition [41]. This motivates us to combine eye-tracking with video analytics to enable accurate human visual attention detection and tracking. However, machine vision techniques are computationally demanding and energy-intensive, limiting adoption in resource-constrained wearable settings. These challenges must be addressed in order to accomplish our goal of accurate and efficient human visual attention tracking.

This paper presents *MemX*, a biologically-inspired attention-aware eyewear system to enable automated, highly personalized capture of interesting visual content, which is recorded in compact video snippets. We retrofit off-the-shelf glasses with an inward-facing eye camera and a forward-facing world camera. The eye camera performs eye tracking to determine human gaze direction. The world camera samples and analyzes the field of view. We propose a new temporal visual attention (TVA) network, which unifies eye tracking and video analysis, to enable accurate and computationally efficient human visual attention tracking and salient visual content analysis. Specifically, the eye camera performs eye tracking to determine human visual regional focus, which confines computation-intensive video analytics tasks solely to regions of personal interest, thereby minimizing computing time and energy costs. Working in tandem, the video analytics component analyzes the visual scene content and identifies the salient object of human visual focus as well as the spatial visual context, thereby improving the accuracy and stability of human visual attention tracking. In addition, the high-resolution scene frame is used as supplementary input to the TVA network for more accurate attention detection. Such high-resolution scene frames are captured only when likely attention is detected, which significantly reduces energy consumption by decreasing the average sampling frequency. Furthermore, eye movement types, i.e., saccades or smooth pursuit, act as an “intermediate” supervision to be incorporated into the learning process.

This work makes the following contributions.

- (1) We propose a new temporal visual attention (TVA) network that unifies eye-tracking and video analytics, to enable accurate and computation-efficient human visual attention tracking.
- (2) We develop *MemX*, a new biologically inspired, attention-aware eyewear system. Equipped with the proposed TVA network, *MemX* automatically detects human visual attention and uses it to analyze and record visual content of personal interest.
- (3) *MemX* is evaluated using the YouTube-VIS dataset [49] and 30 participants. Experimental results show that compared to the eye-tracking-alone method, *MemX* significantly improves the attention tracking accuracy, while maintaining high system energy efficiency.
- (4) *MemX* can be applied in several personal visual content capture scenarios, such as sightseeing, lifelogging, academic events (e.g., meetings), shopping, and sports logging. We conduct 11 in-field pilot studies with different potential usage scenarios to evaluate *MemX* and illustrate some of its potential uses and benefits. According to the questionnaire survey, 96.05% of the moments of interest are successfully captured by *MemX*. In addition, the pilot studies demonstrate significant energy savings as compared with the record-everything case, with 86.36% energy savings on average.

The rest of the paper is organized as follows. Section 2 surveys related work. Section 3 presents the proposed *MemX* system and algorithm. Section 4 details *MemX* system design and implementation. Section 5 presents and discusses the experimental results. Section 6 presents the in-field pilot studies. We conclude the work in Section 7.

2 RELATED WORK

This section summarizes the most relevant work in the areas of eye tracking, eye movement classification, visual attention, human interest, and computation-efficient video analytics.

2.1 Eye Tracking

There has been extensive work on eye tracking and estimation. Two comprehensive reviews of existing approaches are provided by Hansen et al. [11] and Cazzato et al. [8]. These existing methods can be categorized into two classes: model-based approaches [44, 48] and appearance-based approaches [50, 52]. For model-based methods, high-quality images and appropriate lighting conditions are critical, while appearance-based methods leverage more facial features, such as head orientation, from large-scale training data. Compared with model-based methods, appearance-based methods are less sensitive to image quality. Recent work has shown that appearance-based methods can offer good generalization capability even without user-specific data [51]. Thanks to the emerging large-scale data sets and deep learning techniques, the performance of learning-based eye-tracking models has been steadily improving [22, 43, 47]. Our system unifies eye tracking and video analysis to track human visual attention, allowing video analysis to focus on salient visual content, thereby reducing computation and energy costs.

2.2 Eye Movement Classification

As a fundamental step towards eye tracking [21], eye movement classification has been extensively studied and a wide range of methods have been proposed, including classical algorithms [32], Bayesian Mixed Models (BMM) [38], and deep learning [30]. Most classical approaches identify saccade and fixation movements using velocity-threshold identification, dispersion-threshold identification, and hidden Markov model identification [10]. Later, Tafaj et al. [39] and Santini et al. [34] added eye movement classification of smooth pursuit in BMM. Recently, deep learning has been used to classify eye movements. For example, Zemblys et al. proposed a CNN-based architecture for sequence-to-sequence eye movement classification [30]. Startsev et al. combined 1D-CNN with BLSTM to classify eye movements as fixation, saccade, and smooth pursuit in windows of up to 1 s [36]. Our

system uses eye movement types as a supplement “intermediate” supervision to drive system-level learning in *MemX*, and also identifies the instants of visual attention.

2.3 Visual Attention and Eye Tracking

Due to the limited processing capabilities of biological vision system, attention (and high spatial resolution capture by the fovea) is directed to a small portion of a scene [7], thereby reducing the computational demands on the visual cortex. Leveraging attention, humans can flexibly control limited computational resources [25] to process the most relevant and preferential sensory information [3]. Recently, extensive studies have combined visual attention and eye-tracking technologies. For instance, Hwang et al. proposed to use eye-tracking to quantify human visual attention in an online shopping environment [15]. Katz et al. used eye-tracking to identify regional gaze fixation on a respiratory function monitor for medical diagnosis [19]. Wang et al. applied visual attention in unsupervised video object segmentation [45]. These works promote the idea of combining eye-tracking and video analytics methods to enable accurate and efficient human attention detection and tracking.

2.4 Visual Attention and Human Interest

Previous works in the field of physiology and psychology research demonstrate that eye movements reveal human attention and gaze positions indicate interesting and important viewed content. Jain et al. pointed out that eye movements can reveal attention’s change or shift [16]. For example, saccade means quick visual attention shift, fixation means the visual attention rests at the same location, and smooth pursuit suggests visual attention follows a moving target smoothly [16]. Tag et al. [40] described a smart eyewear platform capable of attention tracking by measuring the cognitive engagement of people in different situations. Later, Abdelrahman et al. [1] pointed out that eye data can be indicative for human cognition [1, 23]. For instance, gaze positions can reveal the attention locus [1]. Similarly, Kranthi et al. [23] suggested that gaze can be used as an indicator of importance when humans watch a video. That is, gaze positions indicate important or interesting viewed content [23].

2.5 Computation-Efficient Machine Vision

Deep learning has achieved great success in machine vision applications, such as object detection, object classification, and semantic segmentation. Deep models, such as ResNet [12], DenseNet [14], MobileNetV2 [33], and MobileNetV3 [13] are widely used in machine vision applications. However, deep models incur high computation costs, limiting adoption in resource-constrained mobile environments. Light-weight computation-efficient network architecture is therefore an important research branch of deep learning communities. Take for example the architecture of Google’s MobileNetV2 [33] which is a representative light-weight network designed for mobile device. Linear bottlenecks and inverted residual blocks were used as the basic structure in MobileNetV2, aiming to solve the issue of feature degradation during training. Also, as pointed out by Lubana et al., energy consumption significantly depends on the transferred resolutions in imaging systems [27]. Therefore, they proposed that energy consumption can be dramatically reduced if only the task-related information is input to deep models. Inspired by the above work, we propose a new lightweight temporal visual attention network, which fuses two visual input sources, i.e., temporal eye movements and field-of-view scene video, to offer higher computation efficiency while still guarantee task performance.

3 TEMPORAL VISUAL ATTENTION NETWORK

3.1 Design Motivations

The vision of the personalized Memex builds on the key observation that content capture must match a user’s personal interest in order to be relevant and valuable. As visual attention directly reflects human attentive interest, the goal of this work is to design an accurate and efficient method for human visual attention tracking and salient

content analysis. In this section, we first describe human visual attention, and then discuss the capabilities and limitations of eye-tracking and video analysis techniques, which motivate the proposed temporal visual analytical network, a unification of these two techniques.

Humans direct their foveas, the high-resolution regions of their retinas, toward locations attracting their attention: foveal orientations reveal locations of visual attention. Since real-world scenes often contain multiple stationary and moving objects, attention is a temporally and spatially selective process, driven by top-down influences and bottom-up sensing mechanisms. It starts with steering the eye gaze towards an object attracting attention, then continuously follows the object of interest, until new visual stimuli or top-down goals cause the fovea to drift away.

From the eye-tracking perspective, as stated in Section 1 and illustrated in Fig. 1, human eye movement patterns can be mainly classified into two distinct and alternating phases, namely saccade and smooth pursuit [24]. Saccade, the rapid conjugate eye shift, is triggered either voluntarily or by an external visual stimulus, indicating the change or a new moment of human visual attention. Smooth pursuit, on the other hand, refers to slow and smooth eye movements following the same object of interest. Together, a human visual attention epoch starts with a saccade phase, followed by a smooth pursuit phase. Therefore, eye movement analysis, e.g., saccade-smooth pursuit transition detection, can be used to detect and classify the moments of human visual attention. In addition, as described in Section 2, recent work on eye tracking provides relatively accurate estimation of human visual regional focus by measuring the angular direction of eye gaze relative to the eyewear's field of view. Furthermore, eye tracking is computationally efficient, as it requires low-resolution eye images and light-weight algorithms. However, as stated in Section 1, eye tracking does not directly detect the object of interest, and cannot independently provide accurate and stable indications of human visual interest, due to limited tracking resolution and inherently noisy eye movement patterns (as shown in Fig. 1).

A key motivation of this work is that video analysis techniques, e.g., object detection and classification, offer object or semantic level information, which is complementary to regional-level eye tracking. It is thus possible to augment eye tracking with video analysis techniques to perform accurate and stable human attention detection. However, object and semantic level video analysis tasks are data- and computation-intensive. Consider recent work on video instance segmentation (VIS) [49]; VIS is one of the most broadly-used video analysis tasks, consisting of simultaneous detection, segmentation, and tracking of object instances. The latest work on VIS adopts ResNet-50-FPN or ResNet-100-FPN [17] as a feature extractor and uses regional proposal network [31] to identify possible regions of objects. However, the high energy demand of the VIS task limits its adoption in computation and energy-constrained wearable settings.

We observe that eye tracking can help improve the efficiency of video analysis tasks. First, eye tracking efficiently captures human visual regional focus, focusing video analysis on regions of interest and eliminating analysis of other, superfluous regions, thereby improving computational efficiency. Second, eye tracking can detect human visual attention, enabling video analysis to be triggered only when human attention is likely. Furthermore, eye movement types, i.e., saccades or smooth pursuit, can serve as an “intermediate” supervision that can be incorporated into the learning process to regularize the learning problem, leading to more accurate decisions.

In summary, working in tandem, eye tracking and video analysis can potentially deliver accurate and efficient human visual attention tracking and serve as the foundation to support personalized moment auto-capture.

3.2 Pipeline of the TVA Network

Here, we describe the overall pipeline of the proposed temporal visual attention (TVA) network that we deploy in the *MemX* system to detect the user's visual attention.

Fig. 2 shows the proposed TVA framework that unifies eye-tracking and video analysis to detect the user’s visual attention. Two video streams serve as the inputs of this framework, which include (1) the scene stream $\mathbf{I} = \{\mathbf{I}_1, \mathbf{I}_2, \dots, \mathbf{I}_m\}$ captured by the forward-facing world camera; and (2) the video sequence $\mathbf{E} = \{\mathbf{E}_1, \mathbf{E}_2, \dots, \mathbf{E}_m\}$ captured by the inward-facing eye camera, where m is the total number of time steps. For a given time step t , the TVA network predicts whether the user is attentive to a newly-appeared instance or tracks a previously-attentive instance for the current frame \mathbf{I}_t with a binary result $a_t \in \{0, 1\}$. If true ($a_t = 1$), *MemX* invokes high-resolution video recording. Otherwise ($a_t = 0$), the current frame \mathbf{I}_t is discarded to save energy.

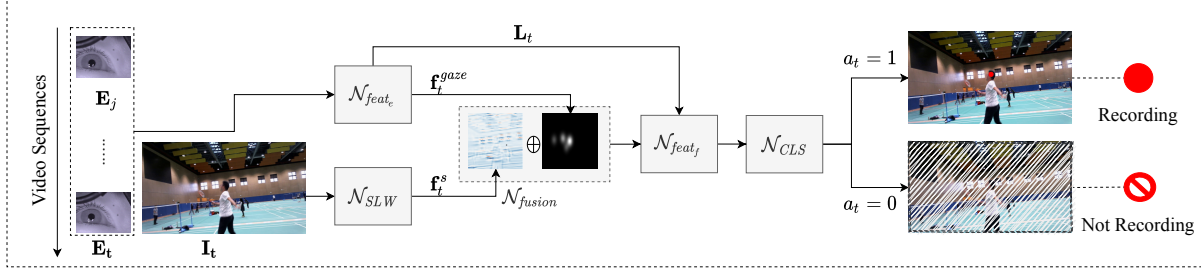


Fig. 2. Overall pipeline of the proposed Temporal Visual Attention (TVA) network. The TVA network distills historical eye movement phases likelihood \mathbf{L}_t and gaze features \mathbf{f}_t^{gaze} from the low-resolution eye camera, along with object and semantic features \mathbf{f}_t^s from the high-resolution world camera, and then performs feature fusion and classification for human attention detection and content auto-capture.

This work develops *MemX* to enable automated, personalized capture of interesting visual content, with the energy-constrained wearable form factor. To this end, the proposed TVA network aims to deliver accurate and energy-efficient visual attention detection by unifying eye-tracking and video analysis. To generate temporally-consistent predictions, the TVA network continuously senses the inward-facing low-resolution eye video and extracts two essential gaze representations, which are the likelihood \mathbf{L}_t of historical eye movement types and the gaze features \mathbf{f}_t^{gaze} extracted from the historical gaze positions, as shown in Fig. 2. Those two gaze representations are fused together to make an initial prediction that a potential attention epoch is occurring. In particular, the likelihood \mathbf{L}_t of historical eye movement types are used to capture gaze saccade-smooth pursuit transition, which is a clear indicator of the beginning of an attention epoch. However, the gaze transition likelihood \mathbf{L}_t alone may not be reliable since the eye movements are subtle and difficult to identify. Therefore, we also use the historical gaze positions \mathbf{P}_t to improve detection robustness. Intuitively, if the majority of the recent gaze positions fall into a small circular region, we can predict with higher confidence that the user’s eye movement has indeed entered a smooth pursuit phase. Moreover, as pointed out by prior works such as pupil detection and tracking [18], which is suitable for energy-constrained wearable scenarios, eye tracking is computationally efficient. Eye tracking may therefore be used for every frame, which can effectively minimize the invocation of computation-intensive video analysis processing and energy-intensive high-resolution video recording.

As mentioned earlier, eye tracking alone is not sufficiently accurate for attention detection. To boost the accuracy of attention detection, the forward-facing high-resolution world camera is also used only when human attention is likely. The TVA network then extracts low-level scene features \mathbf{f}_t^s from the high-resolution frame \mathbf{I}_t , and fuse it with the gaze features \mathbf{f}_t^{gaze} extracted from historical gaze positions \mathbf{P}_t . Compared with the size of the original scene image (224×224), the spatial resolution of the gaze features \mathbf{f}_t^{gaze} is more compact (56×56), and it can also guide our TVA network to focus on analyzing the attentive local region, which further reduces the computational cost.

The obtained scene features \mathbf{f}_t^s are then fused with the likelihood \mathbf{L}_t of historical eye movement types and gaze features \mathbf{f}_t^{gaze} in the TVA network to better predict the temporal attention of the user.

Instead of treating eye-tracking and video analysis as two separate tasks, we design a light-weight network to perform feature fusion. Specifically, we first apply a shallow network \mathcal{N}_{SLW} to obtain features \mathbf{f}_t^s for frame \mathbf{I}_t , i.e., $\mathbf{f}_t^s = \mathcal{N}_{SLW}(\mathbf{I}_t)$. We then map the historical eye tracking results to gaze position related features \mathbf{f}_t^{gaze} and the likelihood of eye movement phases \mathbf{L}_t . The mapping function is denoted as \mathcal{N}_{feat_e} , i.e., $\mathbf{L}_t, \mathbf{f}_t^{gaze} = \mathcal{N}_{feat_e}(\mathbf{E}_j, \dots, \mathbf{E}_t)$. After that, we fuse \mathbf{f}_t^{gaze} and \mathbf{f}_t^s using a predefined operation \mathcal{N}_{fusion} that makes those features complementary to each other so as to augment the detectability of salient region focus. In addition, \mathbf{L}_t is leveraged as an “intermediate” supervision to drive the learning process. The intuition is that if attention is drawn to an object, historical eye movements may follow a detectable pattern, which can be used as prior knowledge to supervise the learning network. For example, an eye-movement sequence: *saccade*, *smooth pursuit*, \dots , *smooth pursuit*, can suggest the occurrence of attention with high confidence. Finally, the above output goes through a classification model \mathcal{N}_{CLS} to generate the final decision of a_t (attention or not). Next, we provide the detailed implementation of the proposed TVA network.

Given an incoming eye frame \mathbf{E}_t , we want to detect user attention, i.e., whether the user is potentially gazing at an object within the field of view. We experimentally define the occurrence of potential attention as the majority of the gaze positions falling within a region with area A or smaller during a time period T . The impacts of A and T on attention detection performance are evaluated in Section 5.

If possible attention is detected at time t , we then trigger the world camera and sense scene frame \mathbf{I}_t to further boost the confidence level that the user is indeed paying attention to an object. Specifically, let $\mathbf{P}_t = \{\mathbf{p}_{t-N+1}, \dots, \mathbf{p}_{t-1}, \mathbf{p}_t\}$ be the predicted historical gaze positions during time steps $(t - N, t]$ where $\mathbf{p}_k \in \mathbb{Z}^2$, and let $\mathbf{L}_t = \{\mathbf{l}_{t-N+1}, \dots, \mathbf{l}_{t-1}, \mathbf{l}_t\}$ be the predicted historical eye movement types likelihood where each $\mathbf{l}_i \in \mathbb{R} \cap [0, 1]$. The goal of the TVA network is to determine from \mathbf{I}_t , \mathbf{P}_t and \mathbf{L}_t whether the user is actually gazing at an object within the scene frame, i.e., $a_t = \mathcal{N}_{TVA}(\mathbf{P}_t, \mathbf{I}_t, \mathbf{L}_t)$, and \mathcal{N}_{TVA} represents the proposed TVA network.

After that, a shallow network \mathcal{N}_{SLW} is applied to the scene image \mathbf{I}_t to generate the scene features \mathbf{f}_t^s with size $56 \times 56 \times 24$, i.e., $\mathbf{f}_t^s = \mathcal{N}_{SLW}(\mathbf{I}_t)$. Since the historical gaze positions \mathbf{P}_t are discrete and not differentiable, we relax each historical gaze position into a continuous heatmap of size 56×56 based on Gaussian distributions. Specifically, for the k^{th} gaze position $\mathbf{p}_k \in \mathbb{Z}^2$, the value $h_{\mathbf{p}_i}$ for a location $\mathbf{p}_i \in \mathbb{Z}^2$ in the heatmap \mathbf{H}_k can be computed as

$$h_{\mathbf{p}_i} = \frac{w_k}{\sigma\sqrt{2\pi}} \exp\left(-\frac{D^2(\mathbf{p}_i, \mathbf{p}_k)}{2\sigma^2}\right), \quad (1)$$

where $D(\mathbf{p}_i, \mathbf{p}_k)$ represents the Euclidean distance between \mathbf{p}_i and \mathbf{p}_k , and w_k is a pre-defined weight related to the time step k . Applying the Gaussian Relaxation in Eq. 1 to each gaze position in \mathbf{P}_t will generate N heatmaps $\{\mathbf{H}_{t-N+1}, \dots, \mathbf{H}_{t-1}, \mathbf{H}_t\}$ of resolution 56×56 . We denote those heatmaps as \mathbf{f}_t^{gaze} .

Finally, The gaze heatmaps \mathbf{f}_t^{gaze} , the scene features \mathbf{f}_t^s and their element-wise product $\mathbf{f}_t^{gaze} \circ \mathbf{f}_t^s$ are concatenated channel-wise to fuse information from gaze analysis and scene images. The resulting features, which are of size $56 \times 56 \times (24 + 8)$, are then fed into one convolution layer Conv0 with 56 output channels. A following fully-connected layer FC0 gradually reduces the channels into two representing the probability of the binary action a_t . In addition, inspired by the work of Wang et al. [46], we append the likelihood \mathbf{L}_t of historical eye movement types with the input tensor of FC0 layer to supervise the learning during training and help with the inference during testing.

4 MEMX: AN ATTENTION-AWARE EYEWEAR SYSTEM

This section describes *MemX*'s hardware design and TVA network integration to support the operation of personalized moment auto-capture. Since energy efficiency is a key focus of *MemX*, this section also describes the energy model of *MemX*.

4.1 *MemX* Hardware Design

MemX is prototyped using smart eyewear with an inward-facing eye camera and a forward-facing world camera. The eye camera performs gaze tracking. The world camera samples and analyzes the field of view. Fig. 3 illustrates the hardware prototype of *MemX*. The first generation of *MemX* uses a Logitech B525 with 1,280×720 resolution as the front-view camera installed on the upper side of the eyewear frame, and a Sony IMX 291 with 320×240 resolution as the eye camera, secured by a standalone arm. We are in the process of developing the second-generation hardware prototype. A key modification is to remove the standalone arm to improve the usability and user-friendliness.

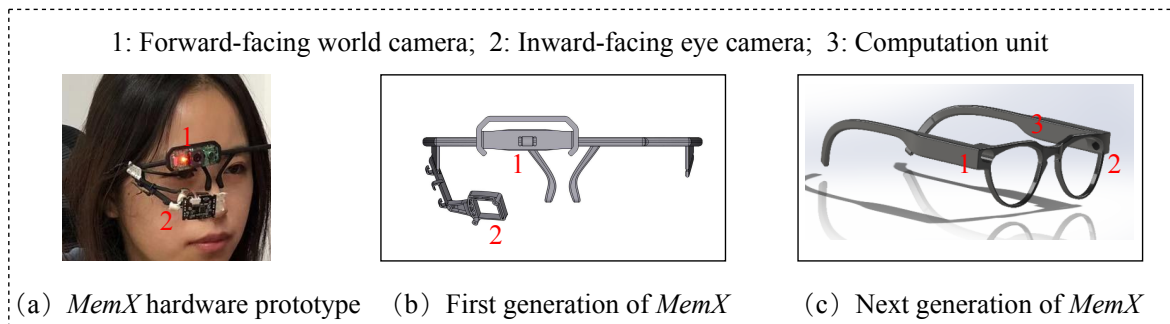


Fig. 3. *MemX* smart eyewear hardware prototype.

The first-generation prototype uses the Jetson Xavier NX mobile computing platform. We are in the process of migrating the *MemX* design to the Ambarella CV28 low-power vision platform directly integrated into one of the eyewear legs, with a battery in the other leg, similar to the hardware design of Spectacles from Snapchat. The CV28 platform is equipped with built-in advanced image processing, high-resolution video encoding, and CVflow computer vision processing capabilities [2]. The typical power consumption of CV28 is in the range of 500 mW, which is suitable for battery-powered wearable design.

4.2 *MemX* Software Operation

MemX is equipped with the proposed TVA network, which uses attention tracking to achieve personalized moment auto-capture. The TVA network continuously tracks eye gaze to detect potential attention events. If a possible attention event is detected, the TVA network will turn on the world camera, and eye gaze and scene sequences are both fed into the TVA network to make the final attention decision.

As described in Section 3, attention is a direct indicator of salient visual content. Furthermore, the attention level, or the level of interest, can be quantified by T , the duration of the smooth pursuit phase, and longer duration implies stronger personal interest. Therefore, T controls the selectivity of moment recording. Furthermore, the proposed TVA network uses computation-efficient eye-tracking to detect potential attention events. A larger T makes eye-tracking more selective. Therefore, the fusion stage of the TVA network, which is more computation and energy intensive, will be triggered less frequently, thereby improving system energy efficiency. The details of

how T affects the performance and energy consumption of *MemX* are discussed in Section 5. One other design parameter of *MemX* is the duration of each video snippet. In the current implementation, *MemX* continuously records a detected moment till a pre-defined duration threshold is reached.

4.3 Energy Model

The energy consumption of *MemX* is mainly contributed by the world camera, eye camera, the TVA network, and also high-resolution video recording. Next, we characterize the energy consumption of these key components.

4.3.1 Energy Consumption of Imaging Pipelines. The operation of an imaging pipeline starts with sensing incoming light and converting it into electrical signals. Then, an image signal processor (ISP) receives the electrical signals and encodes them into a compressed format. The energy consumption of imaging pipeline is contributed by the following three components [27]: E_{sensor} (image sensor), E_{ISP} (ISP), and E_{comm} (communication), as follows:

$$E_{imaging} = E_{sensor} + E_{ISP} + E_{comm}. \quad (2)$$

Next, we discuss the energy consumption of the aforementioned components.

(1) *Energy consumption of image sensing.* The operation of an image sensor consists of three states, i.e., idle, active, and standby. The power consumption of the standby state is negligible (typically in the range of 0.5-1.5 mW), thus ignored from the energy model. Eq. 3 defines sensor energy, as follows:

$$E_{sensor} = P_{sensor, idle} \cdot T_{exp} + P_{sensor, active} \cdot T_{sensor, active}, \quad (3)$$

where $P_{sensor, idle}$ and $P_{sensor, active}$ are the average power consumption when the sensor is in the idle state and active state, respectively. T_{exp} is the exposure time, and the image sensor is idle during the exposure phase. T_{active} is the time duration when the image sensor is active, which is determined by the ratio of transferred frame resolution R_{frame} to the external clock frequency f , i.e., R_{frame}/f . Here, $P_{sensor, idle}$ and f can be viewed as sensor-specific constants, and $P_{sensor, active}$ is a linear function of sensor resolution R ($R \geq R_{frame}$).

(2) *Energy consumption of ISP.* The ISP operates in two states: idle and active. It is active during image processing (T_{ISP}) and idle during image sensing. The time for image sensing is the sum of exposure time T_{exp} and the transferring time of frame (in pixels), i.e., R_{frame}/f . The energy consumption of ISP is then determined as follows.

$$E_{ISP} = P_{ISP, active} \cdot T_{ISP} + P_{ISP, idle} \cdot (T_{exp} + R_{frame}/f), \quad (4)$$

where $P_{ISP, active}$ and $P_{ISP, idle}$ are the average power consumption of the ISP in the active and idle state, respectively.

(3) *Energy consumption of communication interface.* The energy consumption of the communication interface E_{comm} is a linear function of the number of transferred frame pixels R_{frame} [27], as follows:

$$E_{comm} = k \cdot R_{frame}, \quad (5)$$

where k is a design-specific constant determined by the communication interface.

As can be seen from Eq. 2, 3, 4, and 5, the energy consumption of cameras highly depends on the sensor resolution. For *MemX*, the energy consumption of the high-resolution world camera is significantly higher than that of the eye camera. Therefore, the TVA network uses eye-tracking alone to first detect potential attention events, and only then turns on the world camera for accurate attention detection and salient content recording, thereby effectively minimizing the energy cost from the world camera.

4.3.2 Energy Consumption of MemX. *MemX* is equipped with the proposed TVA network. The operation of the TVA network consists of two stages. First, the TVA network continuously performs eye-tracking to detect potential attention events. This stage has high energy efficiency, thanks to the low data rate of the eye camera and computationally-efficient eye tracking. When a possible attention event is detected, the TVA network invokes the second stage, using a light-weight network to perform eye-scene feature fusion in order to finalize the attention decision. We denote the average power of the two stages as $P_{eye\ tracking}$ and P_{fusion} , respectively.

The energy consumption of *MemX* is then formulated as follows:

$$E_{MemX} = T_{always-on} \times (P_{eye\ camera} + P_{eye\ tracking}) + T_{fusion} \times (P_{world\ camera} + P_{fusion}) + T_{auto-captured} \times (P_{world\ camera} + P_{encoding-storing}), \quad (6)$$

where $T_{always-on}$ is the operation time of *MemX*, $P_{eye\ camera}$ and $P_{world\ camera}$ are the power consumption of the inward-facing eye camera and the forward-facing world camera, respectively, T_{fusion} is the operation time of eye-scene feature fusion, $T_{auto-captured}$ is the operation time of high-resolution video recording when human visual attention events are detected, $P_{encoding-storing}$ is the power consumption of the host processor during high-resolution video recording, mostly contributed by video encoding and storage. Compared with the eye-tracking alone stage, the second stage is more data and computation intensive, but is only triggered when potential attention events are detected, and stays off most of the time, thus effectively reducing energy cost. In addition, the TVA network consists of a new light-weight network architecture, which is significantly more efficient than the existing VIS-based design, e.g., MaskTrack R-CNN architecture [49]. As shown in the experimental studies, the proposed TVA network significantly outperforms the VIS-based method in terms of system energy efficiency.

5 EVALUATIONS

This section evaluates *MemX*, the proposed attention-aware smart eyewear system for personalized moment auto-capture. We first detail the experimental settings, including user studies, data collections, and evaluation metrics. Then we present the evaluation results to quantify the performance and efficiency of *MemX*. The studies described in this section are conducted in a controlled lab environment. In-field pilot studies will be presented in the next section.

5.1 Evaluation Methodology

5.1.1 Experimental Setup. The controlled lab setting is shown in Fig. 4. A group of participants wear *MemX*, sit in front of a computer, and watching a sequence of video clips. During each study, video clips are recorded through the field of view of *MemX*, which simultaneously tracks the participant’s visual attention. In total, 30 participants are recruited. A summary of the participants are provided below.

- **Gender:** 10 (33.3%) female and 20 (66.7%) male,
- **Age:** between 21 and 35,
- **Vision:** normal vision (7, 23.3%), nearsighted ≤ 3.0 diopters (8, 26.7%) and ≥ 3.0 diopters (15, 50.0%).

5.1.2 Video Data Preparation. In this work, we adopt the YouTube-VIS dataset¹ [49] as “micro-benchmark in a controlled setting” to quantitatively evaluate such technical capability of *MemX*. In addition, we use pilot studies to further evaluate *MemX* in real-life scenarios.

Specifically, the YouTube-VIS dataset covers a wide range of complex real-world scenarios. For instance, the dataset consists of temporally dynamic scenes, such as sports, as well as spatially diverse scenes, such as multiple objects within the same scene. Therefore, using the YouTube-VIS dataset, we can design a wide range of interesting testing cases to evaluate the accuracy and efficiency of *MemX* for eye tracking and attention detection. For

¹<https://youtube-vos.org/dataset/vis/>



Fig. 4. *MemX* experimental study: The controlled lab setting.

example, the participant is initially attracted by a white duck within a scene consisting of multiple animals, and then his or her attention shifts towards a different animal. In addition, as a widely used dataset, the YouTube-VIS dataset includes necessary object annotation information regarding object location, classification, and instance segmentation. Such annotations help us to accurately set up the ground truth in terms of the objects of interest during the experiments.

The YouTube-VIS dataset consists of a 40-category label set and 2,238 videos with released annotations. Each video snippet lasts 3 to 5 seconds with a frame rate of 30 fps, and every 5th frame is annotated for each video snippet. We first divide the 2,238 videos into four classes based on their content, including *Animal*, *People*, *Vehicle*, and *Others*. As a result, the four classes contain 1487, 437, 215, and 99 videos, respectively. We then randomly select videos in each class, resulting in 1,000 videos total (670, 197, 97, and 36), which are used in this study. Fig. 5 visualizes cases of the 4-class video frames in YouTube-VIS dataset.

Since each video snippet lasts 3 to 5 seconds in the YouTube-VIS dataset, we concatenate multiple video snippets to form videos with a duration of about 7-15 minutes (approximately 100 video snippets). The concatenated video is then provided to the participants to watch. We recruited 30 volunteers to collect their eye data when they watch the concatenated videos. To alleviate the potentially person-dependent bias, we ensure each video snippet is watched by several participants (three in this work). Given the 30-recruited participants, we select 1,000 video snippets from the entire YouTube-VIS datasets.

5.1.3 Eye Data Collection. *MemX* uses the inward-facing eye camera to capture eye video data. We calibrate *MemX* to correlate the world camera with the eye camera. The calibration method follows Pupil Capture² [18]. Each participant watches approximately 100 videos randomly selected from the four classes in our benchmark video data. That ensures each video in our benchmark can be watched by three participants. We pre-select the target attentive object in each video snippet as users' visual interest, and we guide participants to gaze at the pre-assigned object and track their motion. We then obtain the users' eye video dataset that is synchronized with the video dataset.

In the following experiments, the video dataset and the eye dataset are randomly divided into training set, test set, and validation set with a 70%:10%:20% ratio.

²<https://docs.pupil-labs.com/core/software/pupil-capture/#calibration>




Fig. 5. YouTube-VIS dataset illustration. The bounding boxes mark the pre-selected objects of interest.

5.1.4 Evaluation Metrics. We use precision and recall to evaluate the accuracy of human visual attention tracking of *MemX*, defined as follows:

$$precision = \frac{TP}{TP + FP} \text{ and} \quad (7)$$

$$recall = \frac{TP}{TP + FN}, \quad (8)$$

where *TP* (true positive) denotes when *MemX* correctly identifies the attentive object of interest, *FN* (false negative) denotes when *MemX* fails to identify the attentive object of interest, and *FP* (false positive) denotes when *MemX* incorrectly identifies an object of interest that the user actually did not pay attention to. In summary, the higher the precision and recall, the better the accuracy of *MemX*. In addition we also consider average precision (AP) which jointly considers precision and recall measures, as follows.

$$AP = \sum_n (recall_n - recall_{n-1}) \times precision_n, \quad (9)$$

where $precision_n$ and $recall_n$ are the precision and recall at the n th threshold.

To evaluate the energy efficiency of *MemX*, we consider the following hardware settings. The *MemX* prototype is equipped with a Sony IMX 291 based eye camera and a Logitech B525 based world camera. The Logitech B525 camera supports a maximum resolution of 2.07 M-pixels. In addition, we target the Ambarella CV28 low-power computer vision SoC platform, which is equipped with advanced image processing, high-resolution video encoding, and CVflow computer vision processing capabilities. The power estimation is based on the model described in Section 4.

5.1.5 Training TVA Network. This work uses Adam [20] as the optimizer, with an initial learning rate of 0.01, which is decreased by 10% after every 30 epochs until 100 epochs have occurred. For the shallow network \mathcal{N}_{SLM} in TVA, we adopt the first two blocks of MobileNetV2 to extract the scene frame features as MobileNetV2 is

efficient enough for mobile devices [33]. We use pre-trained weights provided in prior work [33] as a starting point for training TVA. During training, we freeze \mathcal{N}_{SLM} during the first 30 epochs.

5.1.6 Baselines. To our best knowledge, there is no prior work that targets the same research challenge and is directly comparable to this work. For comparison, we consider the following two baselines.

(1) *Eye-tracking-alone.* This approach uses eye tracking alone to capture potential attentive visual content. Specifically, we first use eye tracking to detect the saccade-smooth pursuit transition, an indication of potential visual attention shift. Then, we measure the gaze regional focus relative to the field of viewing during a smooth pursuit phase or a fixation phase for a time period. We experimentally define the occurrence of actual attention as when 90% gazes are located in a close region with area $A = (0.05 \times W) \times (0.05 \times H)$ for a time period t , where W and H denote the width and height of viewing scene frame, respectively, and 0.05 is the rescaling ratio. Considering that the Logitech B525 camera has a 69° diagonal field-of-view³ and angular error is approximately with median value of 3.45° [8], we estimate the rescaling ratio as $3.45/69$ (i.e., 0.05) in this work. The effectiveness of how T affects the performance of attention detection is evaluated in Section 5.2.2. The following experiments use the eye-tracking-alone method to establish the baseline accuracy of the proposed work.

(2) *VIS-based method.* This method uses eye tracking and VIS-based object detection to jointly capture potential attentive visual content, and the VIS-based object detection task adopts the MaskTrack R-CNN architecture [49]. Different from the proposed TVA network, in the VIS-based method, eye tracking and VIS-based object detection are two independent, parallel tasks. As a result, VIS-based object detection is always on, which may introduce significant computation and energy overhead. As will be shown in the experimental results, the processing speed of the VIS-based method is only 0.6 frames per second, preventing it from practical adoption into wearables. The purpose of including this method in the experiments is to establish the baseline energy efficiency of the proposed work.

(3) *Saliency-map-based method.* The saliency detection method [29] aims to predict the salient regions that has potentially attracted the use's attention, which is somehow similar to the target of the proposed TVA network. For a fair and complete comparison, we adopt and evaluate the saliency detection method in [29] as an additional baseline. We will refer it as the saliency-map-based method for simplicity.

5.2 Results

5.2.1 Overall Performance. Table 1 shows the accuracy and energy efficiency comparison between *MemX* and the three baseline methods. The precision, recall, and average precision AP of each method are shown in columns 2, 3 and 4. The energy reduction of the proposed work and the eye-tracking-alone based baseline method relative to that of the VIS-based baseline method is shown in column 5.

As the table shows, the eye-tracking-alone method achieves low precision and recall: it cannot be used for accurate and stable attention tracking. In contrast, the proposed TVA-based method significantly improves the attention tracking precision and recall, with only a slight increase in energy consumption. In addition, the proposed TVA network achieves higher average precision than the two baseline methods.

Among these four methods, the VIS-based method achieves the highest precision and recall, which however is associated with significant energy and computation penalties. Specifically, compared with the proposed TVA-based method, the energy consumption of the VIS-based method is $2.7\times$ higher. Furthermore, using the Jetson Xavier NX platform, the processing speed of the VIS-based method is 0.6 frames per second, which is too slow for practical use. In contrast, the processing speed of the proposed TVA-based method exceeds 30 frames per second, offering at least a $50\times$ speedup.

³www.logitech.com/assets/64667/b525-datasheet.ENG.pdf

Compared with the VIS-based method, the saliency-map-based method is with significant accuracy degradation. This is due to the fact that the saliency-map-based method can only provide region-level information, instead of object-level information. An object often contains multiple regions, and human's visual attention on a specific object may temporally shift across the multiple regions of the same object. As a result, the saliency-map-based method fails to accurately detect whether a person gazes at the same object. Compared with the saliency-map-based method, the VIS-based method offers object-level information, thus offering better visual attention analysis accuracy. Also, the proposed TVA-based method outperforms the saliency-map-based method. This is because the TVA network is designed to provide object-level information in a light-weight architecture complemented with the regional-level eye-tracking task. Experimental results demonstrate that, compared against the saliency-map-based method, the proposed TVA-based method achieves 58.48% accuracy improvement (average precision) with 45.64% energy savings gain.

In summary, the three baseline methods suffer from either serious accuracy or energy efficiency limitations. In comparison, the proposed TVA network strikes a good balance between accuracy and efficiency for attention tracking.

Table 1. Accuracy and energy efficiency comparison.

Method	Precision	Recall	Average Precision	Energy Savings
TVA-Based Method (Proposed)	87.50%	86.40%	93.65%	70.61%
Eye-Tracking-Alone Method	40.52%	12.70%	35.93%	76.97%
VIS-Based Method	98.08%	90.26%	93.40%	0.00%
Saliency-Map-Based Method	40.08%	37.94%	35.17%	24.97%

5.2.2 Usability Discussion. *MemX* uses attention tracking to control personalized moment auto-capture. As described in Section 4, the duration of the smooth pursuit phase indicates the level of interest. In *MemX*, the proposed TVA network uses the parameter T to determine how long a potential smooth pursuit phase has been detected using eye tracking before triggering high-precision attention tracking and salient content recording. By adjusting the value of T , *MemX* controls the selectivity of moment recording, as well as the energy efficiency of the system.

Fig. 6a shows the impact of T on the trigger rate α of *MemX*. As shown in this figure, the trigger rate α decreases as T increases. Fig. 6b shows the impact of T on the energy consumption of attention detection. As shown in this figure, the energy saving increases as T increases. This is due to the fact that, with the increase of T , the eye-tracking component of the TVA network becomes more selective, and the data and computation intensive fusion stage is less likely to be triggered, thereby reducing the energy cost.

For the in-field pilot studies (described in the next section), T is empirically set to 1 s. Based on the feedback from the pilot studies, this T value strikes a good balance between selectivity of salient content recording and system energy efficiency.

5.2.3 Case Study. Next, we present two cases to offer further insights regarding why the proposed TVA-based method outperforms the eye-tracking-alone method for attention detection.

(1) *A False Negative Case.* As described in Section 1, eye-tracking-alone cannot provide accurate attention tracking, due to limited tracking resolution and inherently noisy eye movement patterns. Fig. 7 shows such a case. As we can see, the user's attention focuses on a fast-moving vehicle, and gaze trajectory is jittering and noisy. Therefore, it is challenging for eye tracking to distinguish between saccade and smooth pursuit. In contrast,




Fig. 6. Effects of T on trigger rate α and energy savings.

the TVA network leverages both regional information provided by eye tracking and object and semantic level information by video analytics, thus is able to accurately detect the fast-moving vehicle as the object of interest.



Fig. 7. A false negative case illustration.

(2) *A False Positive Case.* Consider a frequent inattentive event in which a person is not paying attention to the surrounding environment and gaze is fixed on one position. Eye tracking alone will mistakenly classify this case as a qualified attention event. In contrast, the TVA network is able to detect such a false positive case. Specifically, as shown in Fig. 8, the TVA network recognizes that the regional visual focus along the gaze trajectory does not match a specific object. Therefore, there is no salient visual object that draws the user's attention.




Fig. 8. A false positive case illustration.

6 PILOT STUDY

MemX can be applied to a wide range of potential application scenarios. We have conducted multi-round user interviews to explore potential usage scenarios. Suggested popular scenarios include daily lifelogging, traveling and sightseeing, sports logging, public event summary, etc. After examining these scenarios, we observe that

different scenarios exhibit distinct characteristics. In particular, we have identified two key attributes, namely scene complexity and human attention dynamics to categorize individual scenarios as follows.

- (1) **Scene complexity** depends on the number of objects in each scene as well as their motion patterns. A simple scene might contain few, mostly stationary, objects while complex scene might contain many fast-moving objects.
- (2) **Human attention dynamics** depends on the frequency that human attention switches among objects. For example, in a basketball game, human attention is highly dynamic. Book reading, on the other hand, is mostly stationary.

6.1 Data Collection and Description

MemX aims to enable automated, personalized capture of interesting visual content, with energy-constrained wearable form factor. To evaluate the accuracy regarding personalized moment auto-capture and energy efficiency of *MemX*, we select a set of potential usage scenarios, covering the four combinations of the two characteristics. We then conduct in-field pilot studies targeting these scenarios, including 11 pilot studies in total (5 female, 6 male, aged between 21 and 40), covering a diverse range of real-life scenarios. During each study, the participants wear *MemX* to auto-capture visual moments of interest in the form of short video clips. For comparison purposes, the complete visual experience of each study is also recorded by the world camera of *MemX* (baseline video). In total, equipped with *MemX*, 11 participants recorded approximately 245 minutes of baseline video, with about 37 minutes of video clips auto-captured as moments of interest. Fig. 9 shows some exemplary moment snapshots captured by *MemX*.

6.2 Accuracy of *MemX* in Pilot Study

6.2.1 Overall Accuracy. To evaluate the accuracy of *MemX*, after each user study, we asked the participant to review the baseline video and manually mark the moments reflecting his or her true interest during recording (ground truth), which are then compared against the video clips auto-captured by *MemX*. Table 2 summarizes the accuracy of *MemX*. As shown in Table 2, *MemX* can accurately detect and automatically capture 96.05% of visual moments of interest across the 11 pilot studies. In other words, the video clips auto-captured by *MemX* accurately reflect the participants' moments of interest.

Table 2. Summary of the in-field pilot study.

No.	Scenario	Record (minute)	Auto-Captured (minute)	Duration Reduction	Number of True Interest	Number of Missed Interest	Energy Savings
1	Lifelogging	1.76	0.13	92.61%	3	0	90.68%
2	Academic events	4.72	0.67	85.74%	5	0	83.80%
3	Grocery shopping	3.71	0.47	87.41%	3	0	85.48%
4	Sports logging	5.12	0.29	94.37%	3	0	92.43%
5	Going for a walk	9.01	1.14	87.30%	5	0	85.37 %
6	Cooking	30.78	6.28	79.61%	6	0	77.67 %
7	Playing game	22.74	0.44	98.08%	5	2	96.15 %
8	Driving #1	32.49	4.69	85.58%	5	1	83.64%
9	Shopping mall	71.24	18.94	73.42%	22	0	71.49%
10	On campus	31.82	0.96	96.99%	10	0	95.06%
11	Driving #2	31.15	3.07	90.15%	6	0	88.22%



Fig. 9. Exemplary moment snapshots captured by *MemX* in the 11 pilot studies. The first column shows the 11 volunteers wearing *MemX*. The rest are the captured moment snapshots. The red circle indicates the measured visual attention by *MemX*. The 11 studies, from top to bottom, are (1) lifelogging, (2) academic events, (3) grocery shopping, (4) sports logging, (5) going for a walk, (6) cooking, (7) playing game, (8) driving #1, (9) shopping mall, (10) on campus, and (11) driving #2.

MemX analyzes the relationship between human attention and interest from the following three aspects simultaneously: (1) the temporal transition from saccade to smooth pursuit, which suggests potential visual attention shift; (2) the gaze duration of following a moving target or fixating on a stationary target, which qualitatively measures the current interest level. In general, the longer the duration, the more interested the user might be; and (3) scene analysis and understanding, which helps to detect cognitively whether there are potentially interesting objects within the region of gaze points. By jointly considering the aforementioned three aspects, *MemX* is able to tackle special corner cases such as mind-wandering and driving scenarios. For example, the user may gaze at a position for a long time unconsciously when mind wandering. In this case, *MemX* leverages scene understanding to help decide whether a potential target of interest exists or not. If not, *MemX* filters out those moments. In another case, the user's attention may temporally drift away and then quickly shift back if no interesting objects are detected. Such attention shifts can be captured and then discarded (due to short duration) by *MemX*. We have included pilot studies for these corner cases. The pilot studies demonstrate that *MemX* can successfully filter out these corner cases and capture the true moments of interest.

6.2.2 Two Exemplary Cases. The following two exemplary cases provide further intuition. Fig. 10 and Fig. 12 show true moments of interest auto-captured by *MemX* and otherwise discarded by *MemX*, respectively. Specifically, Fig. 10 and Fig. 12 show successive video frame sequence recorded by *MemX* with marked gaze positions (red circle). Fig. 11 and Fig. 13 illustrate the time-series normalized gaze distance between two successive frames. For the true moments of interest shown in Fig. 10 and Fig. 11, we can observe that: (1) there is a saccade-smooth pursuit transition at approximately 2.40 seconds; (2) after that, most of the eye movements are smooth pursuit or fixation; and (3) the attention is located on a girl who is playing guitar. In a driving case, as shown in Fig. 12, the driver gazes at a fixed location in his field-of-view without any obvious target or object. As shown in Fig. 13, we can see some gaze shifts, such as at time 17 seconds. However, those gaze shifts do not qualify for possible attention because we cannot find one stable object that the user continuously focuses on. Thus, those moments are discarded by *MemX*.

6.3 Energy Efficiency of *MemX* in Pilot Study

Energy efficiency is essential to wearable devices. In *MemX*, high-resolution video capture and content recording pipeline through the world camera is energy demanding. *MemX* significantly reduces such use only when potential visual attention and moments of interest are detected. As shown in Table 2, the pilot studies demonstrate that the duration that *MemX* triggers the world camera and records moments of interest accounts for a small percentage of the total usage time. Furthermore, in *MemX*, even though the eye tracking process is always on, this stage has high energy efficiency, thanks to the low data rate of the eye camera and energy-efficient TVA network architecture design. The always-on eye-tracking process is approximately 51.98x more energy efficient than the high-resolution video capture and recording pipeline. In addition, the light-weight fusion network is approximately 43.75x more energy efficient than the VIS pipeline [49, 53]. Overall, the pilot study demonstrates that, compared with the record-everything baseline, *MemX* effectively improves system energy efficiency by 86.36% on average. Based on the pilot studies, we estimate that, equipped with 0.36 Wh battery (similar to Spectacles v2), *MemX* is able to support 8 hours of continuous operation after fully charged, which can meet typical daily usage requirement without frequent charging.

6.4 Discussion

After the pilot studies, we have conducted a questionnaire involving the 11 users who participated in the pilot studies to explore potential personal usage scenarios. All 11 participants agree that human visual attention may serve as an easy-to-use information filter and event detection mechanism for visual content gathering. Based on their feedback, the top-3 popular potential usage scenarios include sightseeing, lifelogging, and sports logging.




Fig. 10. Examples of video frames with marked gaze points (red circle) when a participant is enjoying a busker show scenario.

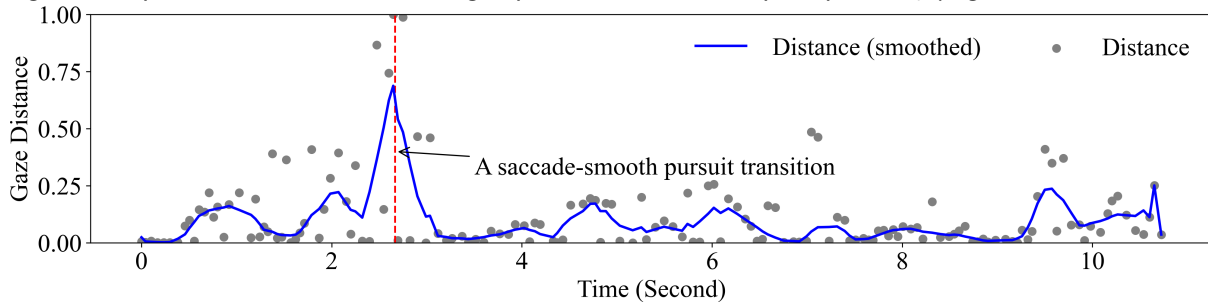


Fig. 11. An example of time-series gaze distance variation when a participant is enjoying a busker show scenario.



Fig. 12. Examples of video frames with marked gaze points (red circle) when a participant is driving.

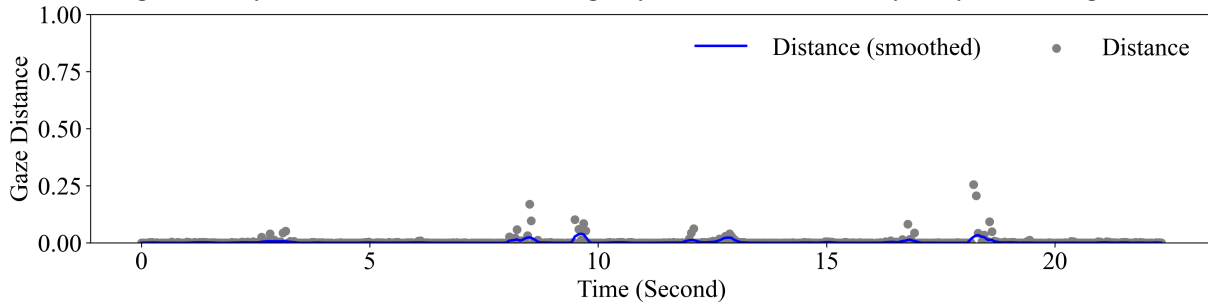


Fig. 13. An example of time-series gaze distance variation when a participant is driving.

Some of their comments are quoted below. “*MemX is pretty cool when I do sightseeing. With MemX, I can record wonderful scenery and interesting encounters effortlessly.*” “*I can hand-freely record memorable moments, e.g., social gathering, during my daily life.*” “*MemX is perfect for sports, such as cycling, probably a better choice than GoPro.*”

Besides personal usage cases, we would also like to explore other domain segments, e.g., industry, education, and gaming, to support on-site visual information gathering and remote communication and interaction. In particular, *MemX*, as a convenient human-computer interface method, can be integrated with the fast-growing AR and VR technologies. Furthermore, we are in the process of developing a video-editing software to automatically create high-quality video journals, e.g., vlogs, using the visual moments captured by *MemX*. Our goal is to enable a complete personal interest aware visual moment auto-capture and content creation framework, with the end goal of fulfilling the long-awaited vision of a personalized visual Memex.

However, we have also identified several limitations of the current version of *MemX* in terms of video quality. In particular, motion compensation is a must-have feature for many eyewear usage scenarios. Furthermore, attention-aware smart glasses may introduce privacy concerns. First, using *MemX*, users can conduct scene recording in a more discrete fashion. Second, content captured by *MemX* discloses the user’s personal interest. Our future work will focus on further improving the video quality, and more importantly, address the privacy concerns introduced by attention-aware personalized moment auto-capture devices.

7 CONCLUSIONS

This work aims to realize the decades-long vision of the personalized visual Memex, emphasizing the importance of content capture which must reflect personal interest to stay relevant and valuable. We have developed *MemX*, a biologically-inspired attention-aware eyewear system to enable auto-capture of personalized attentive visual content, and record moments of personal interest in the form of compact video snippets. *MemX* is equipped with a new temporal visual attention (TVA) network, which unifies eye-tracking and video analysis to enable accurate and computation-efficient human visual attention tracking and salient visual content analysis. *MemX* is evaluated using the YouTube-VIS dataset and 30 participants. Our results show that, compared with the eye-tracking-alone method, *MemX* significantly improves the attention tracking accuracy, while maintaining high system energy efficiency. In addition, we have conducted 11 in-field pilot studies with different potential usage scenarios, which demonstrate the feasibility and potential benefits of *MemX*. We envision that *MemX* can potentially benefit a wide range of personal visual content capture scenarios, such as sightseeing, lifelogging, travel experience recording, and event abstraction.

ACKNOWLEDGMENTS

This work was supported in part by the National Natural Science Foundation of China under Grant No. 62090025 and 61932007 and in part by the National Science Foundation of the United States under grant CNS-2008151.

REFERENCES

- [1] Yomna Abdelrahman, Anam Ahmad Khan, Joshua Newn, Eduardo Velloso, Sherine Ashraf Safwat, James Bailey, Andreas Bulling, Frank Vetere, and Albrecht Schmidt. 2019. Classifying attention types with thermal imaging and eye tracking. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 3, 3 (2019), 1–27.
- [2] Ambarella. 2020. Ambarella introduces CV28M SoC with CVflow to enable new categories of intelligent sensing devices. <https://www.ambarella.com>
- [3] Katharina Anton-Erxleben and Marisa Carrasco. 2013. Attentional enhancement of spatial resolution: linking behavioural and neurophysiological evidence. *Nature Reviews Neuroscience* 14, 3 (2013), 188–200.
- [4] Laxmidhar Behera, Indrani Kar, and Avshalom C Elitzur. 2005. A recurrent quantum neural network model to describe eye tracking of moving targets. *Foundations of Physics Letters* 18, 4 (2005), 357–370.
- [5] Broadbent and E. D. 1957. A mechanical model for human attention and immediate memory. *Psychological Review* 64, 3 (1957), 205–215.
- [6] Vannevar Bush et al. 1945. As we may think. *The atlantic monthly* 176, 1 (1945), 101–108.

- [7] Marisa Carrasco. 2011. Visual attention: The past 25 years. *Vision research* 51, 13 (2011), 1484–1525.
- [8] Dario Cazzato, Marco Leo, Cosimo Distanto, and Holger Voos. 2020. When I Look into Your Eyes: A Survey on Computer Vision Contributions for Human Gaze Estimation and Tracking. *Sensors* 20, 13 (2020), 3739.
- [9] E. Chong, Y. Wang, N. Ruiz, and J. M. Rehg. 2020. Detecting Attended Visual Targets in Video. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 5395–5405. <https://doi.org/10.1109/CVPR42600.2020.00544>
- [10] Shagen Djanian. 2019. *Eye Movement Classification Using Deep Learning*. Ph.D. Dissertation. Master Thesis. Department of Electronic Systems, Aalborg University, Fredrik
- [11] Dan Witzner Hansen and Qiang Ji. 2009. In the eye of the beholder: A survey of models for eyes and gaze. *IEEE transactions on pattern analysis and machine intelligence* 32, 3 (2009), 478–500.
- [12] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 770–778.
- [13] Andrew Howard, Mark Sandler, Grace Chu, Liang-Chieh Chen, Bo Chen, Mingxing Tan, Weijun Wang, Yukun Zhu, Ruoming Pang, Vijay Vasudevan, et al. 2019. Searching for mobilenetv3. In *Proceedings of the IEEE International Conference on Computer Vision*. 1314–1324.
- [14] Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q Weinberger. 2017. Densely connected convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 4700–4708.
- [15] Yoon Min Hwang and Kun Chang Lee. 2018. Using an eye-tracking approach to explore gender differences in visual attention and shopping attitudes in an online shopping environment. *International Journal of Human-Computer Interaction* 34, 1 (2018), 15–24.
- [16] Eakta Jain, Yaser Sheikh, Ariel Shamir, and Jessica Hodgins. 2015. Gaze-driven video re-editing. *ACM Transactions on Graphics (TOG)* 34, 2 (2015), 1–12.
- [17] Jeremiah W Johnson. 2018. Adapting mask-rcnn for automatic nucleus segmentation. *arXiv preprint arXiv:1805.00500* (2018).
- [18] Moritz Kassner, William Patera, and Andreas Bulling. 2014. Pupil: an open source platform for pervasive eye tracking and mobile gaze-based interaction. In *Proceedings of the 2014 ACM international joint conference on pervasive and ubiquitous computing: Adjunct publication*. 1151–1160.
- [19] Trixie A Katz, Danielle D Weinberg, Claire E Fishman, Vinay Nadkarni, Patrice Tremoulet, Arjan B Te Pas, Aleksandra Sarcevic, and Elizabeth E Foglia. 2019. Visual attention on a respiratory function monitor during simulated neonatal resuscitation: an eye-tracking study. *Archives of Disease in Childhood-Fetal and Neonatal Edition* 104, 3 (2019), F259–F264.
- [20] Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014).
- [21] Oleg V Komogortsev, Sampath Jayarathna, Do Hyong Koh, and Sandeep Munikrishne Gowda. 2010. Qualitative and quantitative scoring and evaluation of the eye movement classification algorithms. In *Proceedings of the 2010 Symposium on Eye-Tracking Research & Applications*. ACM, 65–68.
- [22] Kyle Krafka, Aditya Khosla, Petr Kellnhofer, Harini Kannan, Suchendra Bhandarkar, Wojciech Matusik, and Antonio Torralba. 2016. Eye tracking for everyone. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 2176–2184.
- [23] Rachavarapu Kranthi Kumar, Kumar Moneish, Gandhi Vineet, and Subramanian Ramanathan. 2018. Watch to Edit: Video Retargeting using Gaze. 205–215.
- [24] R John Leigh and David S Zee. 2015. *The neurology of eye movements*. OUP USA.
- [25] Grace W Lindsay. 2020. Attention in Psychology, Neuroscience, and Machine Learning. *Frontiers in Computational Neuroscience* 14 (2020), 29.
- [26] Cheryl Lynn Low. 2018. Google Clips review: A smart, but unpredictable camera. <https://www.engadget.com/2018-02-27-google-clips-ai-camera-review.html>
- [27] Ekdeep Singh Lubana and Robert P. Dick. 2018. Digital Foveation: An Energy-Aware Machine Vision Framework. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* 37, 11 (Nov. 2018), 2371–2380. <https://doi.org/10.1109/TCAD.2018.2858340>
- [28] Matei Mancas, Vincent P Ferrera, Nicolas Riche, and John G Taylor. 2016. *From Human Attention to Computational Attention*. Vol. 2. Springer.
- [29] Sophie Marat, Tien HoPhuoc, Lionel Granjon, Nathalie Guyader, Denis Pellerin, and Anne Guérin-Dugué. 2009. Modelling spatio-temporal saliency to predict gaze direction for short videos. *International Journal of Computer Vision* 82, 3 (2009), 231.
- [30] Raimondas Zemblyns, Diederick, C, Niehorster, Kenneth, and Holmqvist. 2018. gazeNet: End-to-end eye-movement event detection with deep neural networks. *Behavior Research Methods* (2018).
- [31] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. 2015. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Advances in neural information processing systems*. 91–99.
- [32] Dario D Salvucci and Joseph H Goldberg. 2000. Identifying fixations and saccades in eye-tracking protocols. In *Proceedings of the 2000 symposium on Eye tracking research & applications*. ACM, 71–78.
- [33] Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. 2018. Mobilenetv2: Inverted residuals and linear bottlenecks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 4510–4520.
- [34] Thiago Santini, Wolfgang Fuhl, Thomas Kübler, and Enkelejda Kasneci. 2016. Bayesian identification of fixations, saccades, and smooth pursuits. In *the Ninth Biennial ACM Symposium*.

- [35] Nelson Silva, Tobias Schreck, Eduardo Veas, Vedran Sabol, Eva Eggeling, and Dieter W. Fellner. 2018. Leveraging Eye-Gaze and Time-Series Features to Predict User Interests and Build a Recommendation Model for Visual Analysis. In *Proceedings of the 2018 ACM Symposium on Eye Tracking Research & Applications (ETRA '18)*. Association for Computing Machinery, New York, NY, USA, Article 13, 9 pages. <https://doi.org/10.1145/3204493.3204546>
- [36] Mikhail Startsev, Ioannis Agtzidis, and Michael Dorr. 2019. 1D CNN with BLSTM for automated classification of fixations, saccades, and smooth pursuits. *Behavior Research Methods* 51, 2 (2019), 556–572.
- [37] Elizabeth Styles and Elizabeth Styles. 2006. *The Psychology of Attention, 2nd Edition*.
- [38] Enkelejda Tafaj, Gjergji Kasneci, Wolfgang Rosenstiel, and Martin Bogdan. 2012. Bayesian Online Clustering of Eye Movement Data. In *Proceedings of the Symposium on Eye Tracking Research and Applications (ETRA '12)*. Association for Computing Machinery, New York, NY, USA, 285–288. <https://doi.org/10.1145/2168556.2168617>
- [39] Enkelejda Tafaj, Thomas C. Kübler, Gjergji Kasneci, Wolfgang Rosenstiel, and Martin Bogdan. 2013. Online Classification of Eye Tracking Data for Automated Analysis of Traffic Hazard Perception. In *Artificial Neural Networks and Machine Learning – ICANN 2013*, Valeri Mladenov, Petia Koprinkova-Hristova, Günther Palm, Alessandro E. P. Villa, Bruno Appollini, and Nikola Kasabov (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 442–450.
- [40] Benjamin Tag, Ryan Mannschreck, Kazunori Sugiura, George Chernyshov, Naohisa Ohta, and Kai Kunze. 2017. *Facial Thermography for Attention Tracking on Smart Eyewear: An Initial Study*. Association for Computing Machinery, New York, NY, USA, 2959–2966. <https://doi.org/10.1145/3027063.3053243>
- [41] A. Torralba and P. Sinha. 2001. Statistical context priming for object detection. In *Proceedings Eighth IEEE International Conference on Computer Vision. ICCV 2001*, Vol. 1. 763–770 vol.1. <https://doi.org/10.1109/ICCV.2001.937604>
- [42] Kang Wang, Hui Su, and Qiang Ji. 2019. Neuro-inspired Eye Tracking with Eye Movement Dynamics. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 9831–9840.
- [43] Kang Wang, Rui Zhao, Hui Su, and Qiang Ji. 2019. Generalizing eye tracking with bayesian adversarial learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 11907–11916.
- [44] Shuo Wang, Ming Jiang, Xavier Morin Duchesne, Elizabeth A Laugeson, Daniel P Kennedy, Ralph Adolphs, and Qi Zhao. 2015. Atypical visual saliency in autism spectrum disorder quantified through model-based eye tracking. *Neuron* 88, 3 (2015), 604–616.
- [45] Wenguan Wang, Hongmei Song, Shuyang Zhao, Jianbing Shen, Sanyuan Zhao, Steven CH Hoi, and Haibin Ling. 2019. Learning unsupervised video object segmentation through visual attention. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 3064–3074.
- [46] Yujiang Wang, Mingzhi Dong, Jie Shen, Yang Wu, Shiyang Cheng, and Maja Pantic. 2020. Dynamic Face Video Segmentation via Reinforcement Learning. In *Proc. IEEE/CVF Conf. on Computer Vision and Pattern Recognition*. 6959–6969.
- [47] En Teng Wong, Seanglidet Yean, Qingyao Hu, Bu Sung Lee, Jigang Liu, and Rajan Deepu. 2019. Gaze Estimation Using Residual Neural Network. In *2019 IEEE International Conference on Pervasive Computing and Communications Workshops (PerCom Workshops)*. IEEE, 411–414.
- [48] Erroll Wood and Andreas Bulling. 2014. Eytat: Model-based gaze estimation on unmodified tablet computers. In *Proceedings of the Symposium on Eye Tracking Research and Applications*. 207–210.
- [49] Linjie Yang, Yuchen Fan, and Ning Xu. 2019. Video instance segmentation. In *Proceedings of the IEEE International Conference on Computer Vision*. 5188–5197.
- [50] Xucong Zhang, Yusuke Sugano, and Andreas Bulling. 2018. Revisiting Data Normalization for Appearance-Based Gaze Estimation. In *Proc. International Symposium on Eye Tracking Research and Applications (ETRA) (2018-03-28)*. 12:1–12:9. <https://doi.org/10.1145/3204493.3204548>
- [51] Xucong Zhang, Yusuke Sugano, and Andreas Bulling. 2019. Evaluation of Appearance-Based Methods and Implications for Gaze-Based Applications. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. ACM, 416.
- [52] Xucong Zhang, Yusuke Sugano, Mario Fritz, and Andreas Bulling. 2017. Mpiigaze: Real-world dataset and deep appearance-based gaze estimation. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 41, 1 (2017), 162–175.
- [53] Yingying Zhao, Mingzhi Dong, Yujiang Wang, Da Feng, Qin Lv, Robert Dick, Dongsheng Li, Tun Lu, Ning Gu, and Li Shang. 2021. A Reinforcement-Learning-Based Energy-Efficient Framework for Multi-Task Video Analytics Pipeline. [arXiv:2104.04443](https://arxiv.org/abs/2104.04443)