

Head and eye egocentric gesture recognition for human-robot interaction using eyewear cameras *

Javier Marina-Miranda¹ and V. Javier Traver²

¹Universitat Jaume I, E12071-Castellón, Spain

²Institute of New Imaging Technologies, Universitat Jaume I
E12071-Castellón, Spain
vtraver@uji.es

Abstract

Non-verbal communication plays a particularly important role in a wide range of scenarios in Human-Robot Interaction (HRI). Accordingly, this work addresses the problem of human gesture recognition. In particular, we focus on head and eye gestures, and adopt an egocentric (first-person) perspective using eyewear cameras. We argue that this egocentric view may offer a number of conceptual and technical benefits over scene- or robot-centric perspectives.

A motion-based recognition approach is proposed, which operates at two temporal granularities. Locally, frame-to-frame homographies are estimated with a convolutional neural network (CNN). The output of this CNN is input to a long short-term memory (LSTM) to capture longer-term temporal visual relationships, which are relevant to characterize gestures.

Regarding the configuration of the network architecture, one particularly interesting finding is that using the output of an internal layer of the homography CNN increases the recognition rate with respect to using the homography matrix itself. While this work focuses on action recognition, and no robot or user study has been conducted yet, the system has been designed to meet real-time constraints. The encouraging results suggest that the proposed egocentric perspective is viable, and this proof-of-concept work provides novel and useful contributions to the exciting area of HRI.

1 INTRODUCTION

With the advances in Social Robotics, Human-Robot Interaction (HRI) is an increasingly important area of research. Much work has been conducted on the *robot* side of the interaction so as to generate gaze and body motions which are expressive [2], natural, and effective for interaction [28] and collaborations [39, 41], and correspond well to human speech [34, 17]. However, given the importance of non-verbal cues in communication, a successful interaction requires the proper analysis of *human* action [20] as well. While head and eye-based gestures have been investigated in the broader context of human-computer interaction (HCI) [33, 12, 43, 42, 16], less efforts have

arguably been devoted in the HRI domain, for instance, to assist and improve human-robot task performance [5, 14].

In HCI, head gestures can be useful for interacting with mobile devices. For instance, touch input can be enriched with slight head motions for shortcut commands [16]. Combining gaze and head allows for more natural and accurate interaction [43]. Based on smart glasses, gaze and head can be combined following the paradigm of selection-by-looking followed by manipulation-by-moving [33]. Since small objects can be hard to select using only gaze, subtle head movements can provide additional cues for more accurate hands-free pointing [42]. As a further recent evidence, a synergistic use of gaze and head movements can provide users with more flexibility for gaze-based point-and-select tasks [37]. Although these concepts and techniques are useful and, to some extent, also applicable in HRI, most research in HCI focuses on interaction with *on-screen* objects. In contrast, gesture in HRI should consider wider and more natural interaction with a *physical* robot and, optionally, with *in-scene* objects. Notwithstanding this kind of differences, a common conceptual framework for gestures [7] in the more general field of human-machine interaction may increase cross-fertilization of otherwise separate research efforts.

In this work, the use of visual-based egocentric perspective for head and eye gesture recognition in the context of HRI is proposed. We argue that the first-person view (FPV) may offer benefits over a third-person view (TPV), as follows: (1) the hard problem of localizing and segmenting human body parts is naturally circumvented; (2) no external (scene or robot) cameras are required; (3) no need that cameras point in specific direction, so the robot might respond to user gestures even if it is not facing the person or is unaware of their presence; (4) the robot does not need to be nearby (it can be in another room in a building or even far away); (5) multiple robots can potentially be interacted with using a single eyewear device. Therefore, scenarios where the egocentric perspective can be helpful include initiating a dialogue when the robot is not yet ready; interacting with scene objects, where gaze cues are relevant to understand or disambiguate head gestures (e.g. “grasp *this*” command); in contexts where robot-centric or scene-centric vision are either not possible (e.g. the robot has no camera, or it is remotely operated, or in outdoors interaction), or affordable (e.g. surveillance-like cameras not available in every room). There are certainly also some drawbacks such as that the eyewear devices may be uncomfortable to wear, and they might

DOI: 10.1109/LRA.2022.3180442

be less cost-effective if multiple persons have to interact with a shared robot. Occlusions may affect both FPV and TPV, although differently, so a redundant system with both types of cameras can perform synergistically and overcome this limitation.

Despite the potential advantages of the egocentric-based head and eye gesture recognition, this problem has been scarcely addressed in HRI. Exceptions include a system to request and guide a robot to find some object [46], helping navigation to wheelchair users with limited hand mobility [26], or predicting user intention from gaze in grasping tasks [25] and hand-held robots [40]. Most of these works have in common that they seek to assist disabled people. While these are very important target users to consider, we believe the egocentric perspective can benefit a larger general population, in particular when it comes to interacting with robots in our daily-life environments. In fact, human gestures recognition can assist human workers in collaborative tasks in industrial settings [29], and play an important role for effective and natural human-robot communication in a wide range of applications in social robotics [4], such as education or entertainment. At least in some of those contexts, an egocentric vision perspective represents a valuable alternative or complement to the mainstream robot-centric and scene-centric visual sensors. For instance, when the human subject is holding some objects, smart glasses can capture useful head and eye cues for hands-free interaction [35]. Human-based egocentric sensors are also helpful in controlling remote robots such as aerial ones [3].

To address the problem of egocentric head and eye gesture recognition, we propose a motion-based, computer vision system which operates at two temporal levels of the incoming egocentric visual streams. First, frame-to-frame motions are analyzed with a convolutional neural network (CNN) for homography estimation. The output of this network is used as input to a long short-term memory (LSTM) so as to capture longer-range temporal dependencies required to characterize gestures. LSTMs have recently been proposed in the context of human action recognition, for example using 3D skeleton data encoding parts and joints of the human body [27, 38, 8], or in multimodal approaches [18].

The contributions of this work are: (1) An egocentric perspective for head and eye gesture recognition for HRI. This contrasts with the mainstream work which uses robot- or scene-centric vision, and hence further work in this area is promoted. (2) The exploration of an homography estimation network and its joint use with LSTMs. Although the combination of CNN and LSTMs is not new, the study of a recent homography estimation network, and the fusion of head and eye information are novel aspects. (3) Design decisions and experimental work that provide additional insights for interested practitioners to develop similar systems. The good computational and recognition performance are encouraging, and foster further investigation, for which some possibilities are suggested.

2 METHODOLOGY

To address the problem of head and eye gesture recognition, an eyewear device with a world camera (W) and one or two eye cameras (E) is assumed.

2.1 Gestures

Some studies propose generic gestures for deictic hand gestures [9], grasp types [13], or combining hand gestures and language in the context of HRI [31], but nothing was found regarding head and eye gestures, specifically for the context of HRI, that could guide our work.

When gestures-to-commands mappings are arbitrary, like geometric shapes [10], the gestures need to be learned and remembered. This issue can be addressed by using on-screen drawings that induce gaze patterns [30], or make user's gaze follow object contour [22]. While valuable, these approaches have some limitations such as that gestures may still be unnatural, or executing them may require some skill, and the need of a display hinders its applicability in some contexts, like many HRI scenarios. More generally speaking, there are a number of factors affecting the usability of gaze [19] and head gestures [26], which are important to consider in many practical settings.

In this work, the following criteria for head and eye gestures were considered: (1) oriented to social robots; (2) general (not application-dependent); (3) universal (easily understandable by users); (4) natural (avoiding gestures which are hard to perform or memorize); and (5) supportive (to express common communication needs).

With these criteria in mind, an initial catalog of interaction units (IUs) was proposed (Table 1). The IUs have been classified according to the taxonomy of illocutionary acts [36], which includes *representatives* (Rep) for assertions, *directives* (Dir) for requests or commands, *commissives* for promises, *expressives* (Exp) for psychological states, and *declarations*. Although this proposal is not definitive and this work only explores a subset of it (the boldfaced IUs, illustrated in Fig. 1 and in videos in the folder 'gestures' of Supplementary material), we believe it is a good starting point that can guide further work.

The five head gestures (nodding, turning left/right, and tilting left/right) considered in [43] resemble our general-purpose gestures. With respect to [46], our work shares the idea of using eyewear to communicate with a robot, and exploits the head movements, even though with different purposes and procedures. The human-robot interaction scenarios where the proposed system might be useful, share commonalities to the kind of non-verbal communication situations considered in [6], with the difference that that work is based on hand gestures rather than head-eyes. As an example, one specific usage case for the proposed gesture set might be one question-based educational or entertainment game with yes/no answers: after some conversation initiation (ComeHere), the robot may ask questions for which the human interactant may respond affirmatively (NoddingHead), negatively (ShakingHead), or show some hesitation (Maybe); the robot may also provide some information for which the user may react by displaying understanding (NoddingHead), disagreement (ShakingHead), or surprise (Surprise), which might be used by the robot to guide subsequent interaction.

2.2 Network architecture and training

For gesture recognition, we propose a motion-based approach which processes image sequences at two temporal levels

Table 1: Proposed IU catalog. The IUs in bold are followed by the identifier used in this paper. The type of IU and involved eyewear images are also given

Interaction unit	Type	Image(s)
Neutral (Neutral)	-	-
Come here (ComeHere)	Dir	W
Go away	Dir	TBD
Stop	Dir	TBD
Listen! Look! Hey!	Dir	TBD
Head nod (NoddingHead)	Rep	W, E
Partial nodding	Rep	W, E
Head shake (ShakingHead)	Rep	W, E
Positive surprise (Surprise)	Exp	W, E
Negative surprise	Exp	TBD
Maybe (Maybe)	Rep	W




Figure 1: Simplified description of interaction units in terms of rotations (R_i) around, or translations (T_i) along some axis $i \in \{X, Y, Z\}$. (a). Side-to-side motions (b) are involved in **ShakingHead** via R_Y , and **Maybe** via R_Z . Up-and-down motions (c) are involved in **NoddingHead** and **ComeHere** via R_X , with some differences (e.g. repetitions in **ComeHere** more time-spaced). **Surprise** may involve a quick gentle backwards motion (d) combining T_Z and R_X .

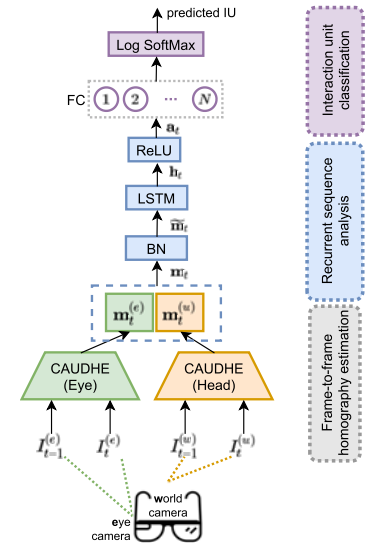


Figure 2: Proposed architecture for head and eye gesture recognition. If only head gestures are considered, only the head CAUDHE instance is required, and no concatenation of the world and eye output features is needed.

For training, the available pretrained CAUDHE's weights¹ were used, and the weights of the LSTM and the FC layer were supervisedly learned using the true IU labels, with the negative log likelihood as the loss function. Fig. 2 illustrates the complete proposed model, which includes two CAUDHE instances, for the head and the eye images, respectively. The outputs of these two CNNs are concatenated and input to the LSTM network, as reasoned before.

For the actual training of the rest of the model (the sequence analysis and IU classification parts in Fig. 2), the precomputed features (\mathbf{m}_t) from the training videos were split into snippets of length $S = 40$ and overlap $O = 30$. Batches of $M = 32$ snippets were used for training speedup, better performance of the BN layer, and faster convergence.

2.3 Practical considerations

Two important design factors for the system to perform efficiently and effectively are the frame rate and the image sizes, since they play important roles in different aspects, and a careful design and proper tradeoff is called for. Since CAUDHE assumes a small baseline (B) in the input frame pairs, the higher the frame rate (f), the more likely this constraint will be satisfied. On the other hand, the world camera of the eyewear device can operate at several spatial resolutions (from 1080p Full HD to QVGA), but the highest frame rates are only selectable

¹<https://github.com/JirongZhang/DeepHomography>

Table 2: Final configuration of acquisition (A) and working (W) resolutions after subsampling

Camera	Image size, $w \times h$ (px)		frame rate (fps)	
	A	W	A	W
World	640×480	192×144 (30%)	60	Up to 30
Eye	192×192	192×192 (100%)		




Figure 3: Pupil Core headset. Please refer to pupil-labs.com/products/core and [23] for technical details

with lower resolutions of these world images. Furthermore, higher quality images have a notable GPU memory footprint, often exceeding the available resources. Therefore, since frames will be later resized, it makes little sense to acquire frames at high spatial resolution.

When using low frame rates, we observed blurring artifacts in the world images when fast head movements were performed. These effects were alleviated with $f > 30$ fps. Thus, if images are high quality, homography estimation should improve. However, computational resources impose constraints on both the image size (to prevent memory overflow) and the frame rate (to keep under the limited computational power). Therefore, both spatial and temporal subsampling were required (Table 2), and their effects are discussed below (subsection 3.4).

2.4 Inference and online recognition

For efficiency at inference time, the input video frames are fed into CAUDHE in batches of K frame pairs for processing K homographies in parallel on the GPU. The K resulting features (\mathbf{m}_t) are reshaped and fed to the LSTM as a K -length sequence. Although predictions of IUs are performed frame-wise, for more robust classification, majority voting can be applied to the resulting K predicted IU labels. Here, $K = 10$ was chosen as a good tradeoff between computational efficiency and classification delay.

3 EXPERIMENTS

3.1 Configuration

For the purposes of this work, the images from the world camera and from the right-eye camera of the Pupil Core headset [23] (Fig. 3) were used.

To the best of our knowledge, no public egocentric gesture dataset exists that meets the required criteria (subsection 2.1).

Table 3: Summary of experimental conditions. Image (W=World, E=Eye), Features (H=Homography, D=Deep), IUs (A=all, A⁻=all but Surprise, B=two IUs)

Aspect ▼ \ Figure ►	4	5	6	7	8, 9, 10, 11
Images	W	{W, E, W+E}			W+E
Features	{H, D}	D			
IUs	A ⁻	A	B	A ⁻	

Table 4: Distribution of instances of IUs in the dataset

IU	Length (s)	%	# instances
Neutral	423.77	70.63	N/A
ShakingHead	37.38	6.23	25
NoddingHead	34.73	5.79	26
ComeHere	27.13	4.52	30
Maybe	46.74	7.79	30
Surprise	30.25	5.04	32
All IUs	600.00	100	—

Therefore, we collected a dataset with 10 minutes of video, half indoors (close-to-camera walls and objects) and half outdoors (distant background), with about 150 seconds (10 videos \times 15 seconds/video) from each of four participants². The participants, volunteer master students, were each briefly instructed about the gestures, without much detail, so that performance was as naturally as possible. Some IUs required additional guidance (see section 4). Prior to the actual recording, a calibration step (about 10–20 seconds long) per participant plus a few trial gestures were performed. Each of the 40 videos (10 videos per participant) contains 3–4 non-overlapping repetitions of a single non-neutral IU. The IU statistics of the dataset are summarized in Table 4. Video frames were manually annotated with the corresponding true IU for supervised learning. Since Surprise turned out to be a particularly subtle IU, it is considered in Fig. 6 in subsection 3.3 to separately study the effect of including the eye images, while the rest of the reported results consider the other five IUs. All tests include the world and eye images except in subsection 3.2, which focuses on the effect of deep features. Conditions are summarized in Table 3.

The Neutral IU is more represented in the dataset, which results in a significant class imbalance situation. To deal with it, the Neutral class was under-sampled (a ratio of training snippets from this class are discarded) so as to reduce the imbalance level. This was proven very effective and used in all reported experiments.

To account for the stochastic nature of training neural models, 30 repetitions were performed, over which performance statistics on the validation sets are reported. Experiments were run on modest low-cost hardware (PC, 8 GB, 1.8 GHz, GeForce MX250 with 2 GB GPU). Most software was developed using the Python ecosystem (NumPy, pandas, scikit-learn, PyTorch, etc.).

The model was trained for 30 epochs, with Adam optimizer, a learning rate of $5 \cdot 10^{-4}$, weight decay penalty of 10^{-2} ,

²This work has been approved by the deontological board of Universitat Jaume I (file code "CD/70/2021", date June 17th, 2021). Human subjects participated after their informed consent.




Figure 4: Maximum accuracy with different motion representations used as input to the LSTM

$\beta_1 = 0.9$, $\beta_2 = 0.999$, and PyTorch's *ReduceLROnPlateau* scheduler. A stratified, IU label-based split of 75% of training instances was used, except in subsection 3.5.

3.2 Homography vs. deep features

For gesture recognition, an obvious first choice is to use the 8 values of the estimated homography as input to the LSTM (\mathbf{m}_t in Fig. 2). However, we are not concerned on motion estimation itself, but on a qualitative representation of motion that can properly characterize the different interaction units. Therefore, we hypothesized that using the output of a hidden layer of CAUDHE could be beneficial. In particular, the output of the last 2D average pooling layer of CAUDHE (512 features) was used as deep features, and compared with the 8-dim homography values. As found by the distribution of the maximum accuracy (Fig. 4), using the deep features for \mathbf{m}_t turns out to be favorable, which indicates that this qualitative representation of motion is richer and more discriminative than the homography estimates. Similar results were obtained for other performance metrics (precision, recall, and F_1 score).

3.3 Using the world and/or eye images

We evaluated how the world ($I_t^{(w)}$) and eye ($I_t^{(e)}$) images contribute to the recognition performance. It can be observed (Fig. 5) that the performance is lower when only eye images are considered, which is not surprising given that the interaction units are heavily head-based. Although the joint consideration of the world and eye images does not improve performance, it is interesting that the recognition does not degrade even when eye images do not bring clearly distinctive information. This is a promising indication that the model might leverage on eye images for capturing commonalities and individual differences for a larger set of images from more participants and repetitions, or when more interaction units, some relying more on eye cues (blinks, winks, subtle motions, etc.), are considered. As an illustrative example, the recognition of the Surprise IU benefits significantly from including the eye images, with an F_1 score about twice that of using the world images only (Fig. 6).

As a controlled experiment, we considered a weighted concatenation $[\alpha_e \mathbf{m}_t^{(e)}, \alpha_w \mathbf{m}_t^{(w)}]$, and trained the network sepa-




Figure 5: Average accuracy with world and/or eye images




Figure 6: Average F_1 for Surprise with and without eye images

rately for NoddingHead vs Neutral and Surprise vs Neutral with $\alpha_e = 1$ and varying $\alpha_w \in [0, 2]$, so as to include $\alpha_e \ll \alpha_w$, $\alpha_e \simeq \alpha_w$ and $\alpha_e \gg \alpha_w$. Results (Fig. 7) show that the convergence is faster and the final performance better with bigger α_w for NoddingHead, whereas α_w has smaller relevance for the performance for Surprise. This is a sign that eye images provide higher discriminative cues than world images for some IUs such as Surprise.

3.4 Effect of subsampling

Regarding spatial subsampling, it was observed that 192×144 (30% of the original VGA resolution) worked well in most scenarios. For temporal subsampling, the analysis of the network convergence with different frame rates (Fig. 8) reveals that, although all get to approximately the same recognition accuracy, convergence is faster with higher frame rates. For instance, for $f = 10$ fps, more than 10 epochs are required to get an accuracy which is obtained with 5 or less epochs for the other frame rates. The likely reason for this behavior is that the baseline between two consecutive frames is smaller with higher frame rates, and this benefits the homography estimation network. Furthermore, the robustness against smaller frame rates can be attributed to the use of deep features. For the final version of the proposed system, $f = 20$ fps is chosen as a good tradeoff solution.




Figure 7: Performance for varying concatenation weight α_w . Left: NoddingHead vs Neutral. Right: Surprise vs Neutral




Figure 8: Average accuracy with different frame rates

3.5 Generalization ability

To evaluate the ability of the proposed model to recognize IUs performed by different subjects (authors) or against different backgrounds (places), we compare the stratified dataset split to the author-based splits (i.e. leaving-one-author out), and place-based splits (i.e. indoor and outdoor recordings are separately considered in the training and test splits). Results (Fig. 9) indicate that in the author-based case the maximum accuracy has larger variability than in the stratified case, since different subjects perform the gestures differently, but in most training repetitions the maximum accuracy was nonetheless high (about 84%). Regarding the place-based split, the performance metric has lower variability and similar or better average value than the stratified case, which is a good sign of the ability of CAUDHE to deal with scenes with backgrounds at different depths, and possibly the positive effect of using deep features. Furthermore, being motion-based, the recognition approach is not biased by the different backgrounds, which might be distracting with appearance-based approaches. Taken these results together, and given the limited data available, it can be argued that the model generalizes reasonably well to different scenes and users.




Figure 9: Distribution of the maximum accuracy on validation videos during training, for different dataset split methods

Predicted IU	True IU				
	Ne	Sh	Nh	Ch	Mb
Ne	2122 96%	31 16%	12 7%	34 25%	14 7%
Sh	11 0%	158 84%	0 0%	0 0%	0 0%
Nh	46 2%	0 0%	136 82%	30 22%	2 1%
Ch	5 0%	0 0%	9 5%	72 53%	0 0%
Mb	18 1%	0 0%	8 5%	0 0%	192 92%

Figure 10: Confusion matrix. Ne=Neutral, Sh=ShakingHead, Nh=NoddingHead, Ch=ComeHere, Mb=Maybe

3.6 Overall performance

To understand which IUs are generally better recognized, or with which other IUs they can be misrecognized, it is useful to analyze the confusion matrix (Fig. 10). It is revealed that Neutral is recognized well by our model. Maybe is recognized with a high accuracy, and is the best recognized “non-neutral” IU. ShakingHead is recognized fine as well, but is confused sometimes with Neutral. Finally, NoddingHead and ComeHere can also be mistaken, as they both imply vertical head movements, and their difference is somehow subtle and hard to distinguish. As the ComeHere movements are more subtle, it is more frequently confused with Neutral. Globally, this performance can be understood by the similarity among groups of the IUs (Fig. 1), and by other factors such as the variability in the performance by different participants, and the limited dataset size (section 4). As a representative performance summary, an overall accuracy of 90% resulted from averaging the peak validation accuracy over 40 training runs, for these five IUs, using deep features and both the world and eye images.

3.7 Online prediction

Finally, in addition to global performance metrics, it is insightful to see the predicted IUs over time. To that end, the network is fed with consecutive pairs of frames, and its output (the predicted IU) at each time step is reported as is, unfiltered. Results (Fig. 11) indicate that ShakingHead (Fig. 11a) and Maybe (Fig. 11d) are generally the best recognized interaction units, and the model tends to label incorrectly only the frames close to transitions. NoddingHead (Fig. 11b) is confused with ComeHere and even Maybe, but the recognition performance is still acceptable. In the case of ComeHere (Fig. 11c), although the overall accuracy is low (actions are often misrecognized as NoddingHead), the correct IU is detected on the onset of the IU, which is the most important part in practical situations. A video demonstration is given in folder ‘online’ in the Supplementary material, where majority voting (subsection 2.4) is applied for robustness.

4 DISCUSSION

Data aspects. The proposed approach for egocentric motion-based gesture recognition has been proven reliable in the collected dataset. Further work would be required to test the system with more interaction units and more data. Class imbalance was addressed by under-sampling the majority class. It could be explored whether additional mechanisms [21] bring some further benefit.

Gestures. At collection time, some IUs (NoddingHead and ShakingHead) simply required some prompts from the experimenter to ensure natural performance, but some other IUs required some hints (to avoid “shrugging shoulders” for Maybe, since that would not involve head or eyes), or suggestions for some exaggeration (Surprise was too subtle otherwise). In the future, elicitation studies might help waive such indications for more natural performance. The gesture catalog was proposed towards meeting criteria such as universality and usability. However, some gestures are subject to differences in context of use and cultural variations, with implications in their meaning [11], social acceptability [1, 15], etc. Therefore, adaptation of some gestures to different applications or world regions may be required.

Network design. Fusion techniques other than feature concatenation for combining world and eye images could be explored. Possibly, an alternative, local encoding of eye images (e.g. optic-flow based) would be preferable over the current global motion assumption. An open interesting question is whether performance can be improved if the homography estimation network is trained or fine-tuned with gesture-specific sequences instead of using the available pretrained weights.

Egocentric view. Since the eyewear device is light-weight, it can be reasonably comfortable to wear, mostly in comparison to virtual-reality devices. It is also preferable over head-mounted cameras because it allows the recognition of gestures where eye cues are important, including deictic gestures. In addition, having a view closer to the wearer’s eyes is generally desirable. Beyond gestures, gaze provides proxies for cognitive states (e.g. attention) which can be relevant for the interaction. Although the human-egocentric perspective is promising and potentially useful in some practical scenarios, this view

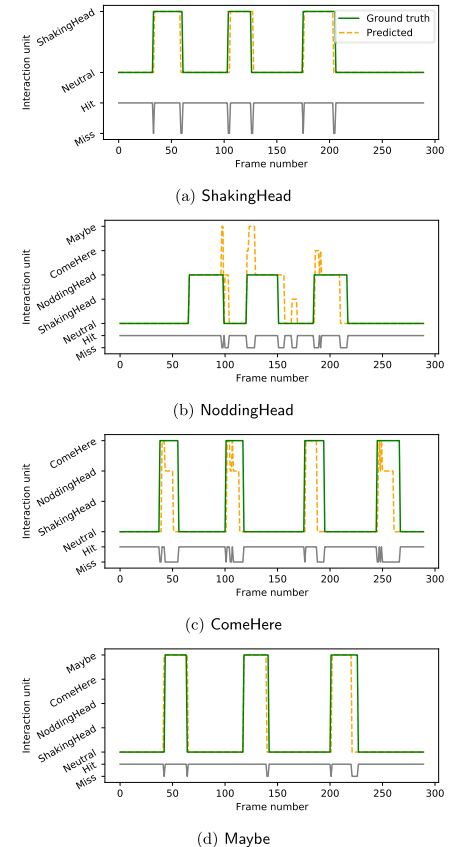


Figure 11: Time diagrams of some validation videos involving various IUs. The bottom gray line helps to appreciate correct recognitions (“hit”) and misrecognitions (“miss”)

can be complemented, where possible, with robot-egocentric and scene-centric visual sensors (e.g. for face, upper-body or full-body analysis), thus aiming at more robust recognition and wider applicability. While this work focuses on human gesture recognition, the interesting related problem of human-to-robot motion retargeting [44, 24] might be explored from an alternative, egocentric perspective, for generating robot (head) movements.

5 CONCLUSIONS

A motion-based head and eye gesture recognition framework has been proposed that leverages on egocentric visual data from an eyewear device. The approach relies on the combi-

nation of an homography estimation convolutional neural network for frame-to-frame motion characterization, and a long short-term memory for capturing longer-range visual dependencies. Five general-purpose interaction units have been shown to be recognized with reasonable high accuracy while the system can run in real time with commodity hardware. Moreover, the system performs equally well in a range of frame rates (e.g. 20-30 fps). These are very encouraging results for practical human-robot interaction requirements. One of the most interesting findings is that using the output of a hidden layer of the homography network increases the recognition performance over using the actual output of raw homography values.

ACKNOWLEDGMENT

We are grateful to the participants in the video capture sessions.

References

- [1] F. Alallah, A. Neshati, Y. Sakamoto, K. Hasan, E. Lank, A. Bunt, and P. Irani. Performer vs. observer: whose comfort level should we consider when examining the social acceptability of input modalities for head-worn display? In *Proc. of the ACM Symposium on Virtual Reality Software and Technology*, 2018.
- [2] M. Avdic, N. Marquardt, Y. Rogers, and J. Vermeulen. Machine body language: Expressing a smart speaker’s activity with intelligible physical motion. In *Designing Interactive Systems*, 2021.
- [3] W. Bentz, S. Dhanjal, and D. Panagou. Unsupervised learning of assistive camera views by an aerial co-robot in augmented reality multitasking environments. In *Intl. Conf. on Robotics and Automation (ICRA)*, 2019.
- [4] C. Breazeal, K. Dautenhahn, and T. Kanda. *Springer Handbook of Robotics*, chapter Social Robotics, pages 1935–1972. 2016.
- [5] C. Breazeal, C. Kidd, A. Thomaz, G. Hoffman, and M. Berlin. Effects of nonverbal communication on efficiency and robustness in human-robot teamwork. In *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, 2005.
- [6] H. Brock, S. Sabanovic, K. Nakamura, and R. Gomez. Robust real-time hand gestural recognition for non-verbal communication with tabletop robot Haru. In *IEEE Intl. Conf. on Robot and Human Interactive Communication (RO-MAN)*, pages 891–898, 2020.
- [7] A. Carfi and F. Mastrogiovanni. Gesture-based human-machine interaction: taxonomy, problem definition, and analysis. *IEEE Trans. on Cybernetics*, 2021.
- [8] F. Carrara, P. Elias, J. Sedmidubsky, and P. Zezula. LSTM-based real-time action detection and prediction in human motion streams. *Multimedia Tools and Applications*, 78(19), Oct. 2019.
- [9] H. Cochet and J. Vauclair. Deictic gestures and symbolic gestures produced by adults in an experimental context: Hand shapes and hand preferences. *Laterality*, 19(3), 2014.
- [10] T. L. Craig, C. A. Nelson, S. Li, and X. Zhang. Human gaze commands classification: A shape based approach to interfacing with robots. In *IEEE/ASME Intl. Conf. on Mechatronic and Embedded Systems and Applications (MESA)*, 2016.
- [11] H. C. David Matsumoto nad C. Hwang. Cultural similarities and differences in emblematic gestures. *Journal of Nonverbal Behavior*, 37, 2013.
- [12] H. Drewes, A. De Luca, and A. Schmidt. Eye-gaze interaction for mobile phones. In *Mobility*, 2007.
- [13] T. Feix, J. Romero, H.-B. Schmiedmayer, A. M. Dollar, and D. Kragic. The GRASP taxonomy of human grasp types. *IEEE Trans. on Human-Machine Systems*, 46(1), 2016.
- [14] K. Fujii, G. Gras, A. Salerno, and G.-Z. Yang. Gaze gesture based human robot interaction for laparoscopic surgery. *Medical Image Analysis*, 44, 2018.
- [15] T. Hossain, M. F. Islam, W. Delamare, F. Chowdhury, and K. Hasan. Exploring social acceptability and users’ preferences of head- and eye-based interaction with mobile devices. In *Intl. Conf. on Mobile and Ubiquitous Multimedia*, 2021.
- [16] S. Hueber, C. Cherek, P. Wacker, J. Borchers, and S. Voelker. Headbang: Using head gestures to trigger discrete actions on mobile devices. In *Intl. Conf. on Human-Computer Interaction with Mobile Devices and Services*, 2020.
- [17] C. T. Ishi, D. Machiyashiki, R. Mikata, and H. Ishiguro. A speech-driven hand gesture generation method and evaluation in Android robots. *IEEE Robotics and Automation Letters*, 3(4), 2018.
- [18] M. M. Islam and T. Iqbal. Multi-GAT: A graphical attention-based hierarchical multimodal representation learning approach for human activity recognition. *IEEE Robotics and Automation Letters*, 6(2), 2021.
- [19] H. Istance and A. I. Hyrskykari. Supporting making fixations and the effect on gaze gesture performance. In *Conf. on Human Factors in Computing Systems*, 2017.
- [20] Y. Ji, Y. Yang, F. Shen, H. T. Shen, and X. Li. A survey of human action analysis in HRI applications. *IEEE Trans. on Circuits and Systems for Video Technology*, 30(7), July 2020.
- [21] J. M. Johnson and T. M. Khoshgoftaar. Survey on deep learning with class imbalance. *Journal of Big Data*, 6:27, 2019.
- [22] F. Jungwirth, M. Haslgrübler, and A. Ferscha. Contour-guided gaze gestures: using object contours as visual guidance for triggering interactions. In *Proc. of the Symposium on Eye Tracking Research and Applications (ETRA)*. ACM, 2018.
- [23] M. Kassner, W. Patera, and A. Bulling. Pupil: An open source platform for pervasive eye tracking and mobile gaze-based interaction. In *Adjunct Proc. of the ACM Intl. Joint Conf. on Pervasive and Ubiquitous Computing*, 2014.
- [24] H. Khalil, E. Coronado, and G. Venture. Human motion retargeting to Pepper humanoid robot from uncalibrated videos using human pose estimation. In *IEEE Intl. Conf. on Robot and Human Interactive Communication (RO-MAN)*, 2021.
- [25] D. Kim, B. B. Kang, K. B. Kim, H. Choi, J. Ha, K.-J. Cho, and S. Jo. Eyes are faster than hands: A soft wearable robot learns user intention from the egocentric view. *Science Robotics*, 4(26), 2019.
- [26] M. Kutbi, X. Du, Y. Chang, B. Sun, N. Agadakos, H. Li, G. Hua, and P. Mordohai. Usability studies of an egocentric vision-based robotic wheelchair. *Journal of Human-Robot Interaction*, 10(1), July 2020.
- [27] Y. Li, C. Lan, J. Xing, W. Zeng, C. Yuan, and J. Liu. Online human action detection using joint classification-regression recurrent neural networks. In *European Conf. on Computer Vision (ECCV)*, 2016.
- [28] C. Liu, C. T. Ishi, H. Ishiguro, and N. Hagita. Generation of nodding, head tilting and eye gazing for human-robot dialogue interaction. In *ACM/IEEE Intl. Conf. on Human-Robot Interaction (HRI)*, 2012.
- [29] H. Liu and L. Wang. Gesture recognition for human-robot collaboration: A review. *International Journal of Industrial Ergonomics*, 68, 2018.
- [30] P. Majoranta, J. Laitinen, J. Kangas, and P. Isokoski. Inducing gaze gestures by static illustrations. In *Proc. of the Symposium on Eye Tracking Research and Applications (ETRA)*, 2019.
- [31] C. Matuszek, L. Bo, L. Zettlemoyer, and D. Fox. Learning from unscripted deictic gesture and language for human-robot interactions. *Proc. of the AAAI Conf. on Artificial Intell.*, 28(1), June 2014.
- [32] T. Nguyen, S. W. Chen, S. S. Shivakumar, C. J. Taylor, and V. Kumar. Unsupervised deep homography: A fast and robust homography estimation model. *IEEE Robotics and Automation Letters*, 3(3), 2018.
- [33] T. Nukarinen, J. Kangas, O. Špakov, P. Isokoski, D. Akkil, J. Rantala, and R. Raisamo. Evaluation of HeadTurn: An interaction technique using the gaze and head turns. In *NordCHI*, 2016.
- [34] J. Ondras, O. Celiktutan, P. Bremner, and H. Gunes. Audio-driven robot upper-body motion synthesis. *IEEE Trans. on Cybernetics*, 2020.
- [35] K.-B. Park, S. H. Choi, J. Y. Lee, Y. Ghasemi, M. Mohammed, and H. Jeong. Hands-free human-robot interaction using multimodal gestures and deep learning in wearable mixed reality. *IEEE Access*, 9, 2021.
- [36] J. R. Searle. A taxonomy of illocutionary acts. In *Language, Mind and Knowledge*. University of Minnesota Press, 1975.
- [37] L. Sidenmark and H. Gellersen. Eye&Head: synergetic eye and head Movement for gaze pointing and selection. In *Symposium on User Interface Software and Technology*, 2019.
- [38] S. Song, C. Lan, J. Xing, W. Zeng, and J. Liu. Spatio-temporal attention-based LSTM networks for 3D action recognition and detection. *IEEE Trans. on Image Processing*, 27(7), 2018.
- [39] C. Stahl, D. Anastasiou, and T. Latour. Social telepresence robots: The role of gesture for collaboration over a distance. In *Proc. of the Pervasive Technologies Related to Assistive Environments Conf. (PETRA)*, 2018.
- [40] J. Stolzenwald and W. W. Mayol-Cuevas. Rebellion and obedience: The effects of intention prediction in cooperative handheld robots. In *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, 2019.
- [41] Y. Terzioğlu, B. Mutlu, and E. Şahin. Designing social cues for collaborative robots: The role of gaze and breathing in human-robot collaboration. In *ACM/IEEE Intl. Conf. on Human-Robot Interaction (HRI)*, 2020.
- [42] O. Špakov, P. Isokoski, and P. Majoranta. Look and lean: Accurate head-assisted eye pointing. In *Proc. of the Symposium on Eye Tracking Research and Applications (ETRA)*, 2014.
- [43] O. Špakov and P. Majoranta. Enhanced gaze interaction using simple head gestures. In *ACM Conf. on Ubiquitous Computing (Ubicomp)*, 2012.
- [44] H. Zhang, W. Li, J. Liu, Z. Chen, Y. Cui, Y. Wang, and R. Xiong. Kinematic motion retargeting via neural latent optimization for learning sign language. *IEEE Robotics and Automation Letters*, 7(2), Apr. 2022.
- [45] J. Zhang, C. Wang, S. Liu, L. Jia, N. Ye, J. Wang, J. Zhou, and J. Sun. Content-aware unsupervised deep homography estimation. In *European Conf. on Computer Vision (ECCV)*, 2020.
- [46] J. Zhang, L. Zhuang, Y. Wang, Y. Zhou, Y. Meng, and G. Hua. An egocentric vision based assistive co-robot. In *IEEE Intl. Conf. on Rehabilitation Robotics (ICORR)*, 2013.