



Exploring Temporal Dependencies in Multimodal Referring Expressions with Mixed Reality

Elena Sibirtseva^{1(✉)}, Ali Ghadirzadeh^{1,2}, Iolanda Leite¹, Mårten Björkman¹,
and Danica Kragic¹

¹ Division of Robotics, Perception and Learning,
EECS at KTH Royal Institute of Technology, Stockholm, Sweden
{elenasi,algh,iolanda,celle,dani}@kth.se

² Intelligent Robotics Research Group, Aalto University, Espoo, Finland

Abstract. In collaborative tasks, people rely both on verbal and non-verbal cues simultaneously to communicate with each other. For human-robot interaction to run smoothly and naturally, a robot should be equipped with the ability to robustly disambiguate referring expressions. In this work, we propose a model that can disambiguate multimodal fetching requests using modalities such as head movements, hand gestures, and speech. We analysed the acquired data from mixed reality experiments and formulated a hypothesis that modelling temporal dependencies of events in these three modalities increases the model's predictive power. We evaluated our model on a Bayesian framework to interpret referring expressions with and without exploiting the temporal prior.

Keywords: Multimodal interaction · Human-robot interaction · Referring expressions · Mixed reality

1 Introduction

In most industrial applications, robots typically work in isolation from humans in repetitive tasks with or without very little interaction with humans. There has been however, the need for developing collaborative robots that can communicate their intent to humans, but also understand human communicative behaviours [21]. This means that we need to go beyond designing classical pre-scripted robots for industrial settings, and more towards assistive robot co-workers with interaction capabilities that empower human workers. For humans to establish successful communication with robotic agents, robots need to use multisensory approaches to perceive human multimodal data and *interpret social cues that communicate humans' intent*.

One of the challenges in understanding human intent is the multisensory fusion problem [11, 23]. The goal of multisensory fusion is to get data from different sensors, combine it in some fashion and, ultimately, a come up with a model

of the user’s intentions. The main problem is interpretation of each modality in combination with each other. Our focus in this research is currently on the following modalities: head movements, hand gestures and speech. The interaction scenario we are interested in is a fetching task, where a human participant explains which object he/she needs from the shared workspace and the robot has to interpret the request from the multimodal sensor observations (Fig. 1).



Fig. 1. Interaction scenario. The human agent is requesting a robot to give them Lego blocks. Any modality can be used in a natural way, no restriction to the interaction is applied. The only tracked modalities are head movements, hand gestures and speech. The participant wears a mixed reality headset to see which block to request and what is the robot’s estimation of their command.

Recent studies focused on intent recognition by combining different features from speech with gaze fixations [1], head movements [25], and gestures [6]. However, in non-guided natural human-robot interaction this approach has its own limitations. Our previous human study [10] showed that participants often look at each other more than at the target object or spend more time looking at the next object in the sequence while still describing the previous one. This led us to look for more high level behaviour patterns that consist of events happening in all the modalities. Our hypothesis is that it is important to look at *when* a certain event happened (e.g. head fixation, pointing gesture) given events in other modalities and *not for how long*. This way we assume that individual events in modalities can be combined in higher level behaviour patterns based on temporal dependencies.

In this paper we present our findings on the answers to the following questions:

- **Q1:** Do common temporal patterns emerge in participants’ behaviour during the fetching request task?
- **Q2:** If we encode these patterns as temporal priors, will they be helpful in inferring the intended object from multimodal referring expressions?
- **Q3:** Are temporal patterns common across most participants or are they person-dependent?

We discuss what we learned from the analysis of a human study and how we see the future development of efficient and natural human-robot interaction in shared workspaces.

2 Related Work

Disambiguation of referring expressions is a well-researched topic in the human-robot interaction community. While written text understanding can be performed in batches, real-time interaction requires continuous reference resolution.

Many works focused exclusively on the language part of the request through incremental reference resolution [4, 7]. However incremental reference resolution is sometimes not enough to completely disambiguate a verbal request. Additional information can be inferred from other modalities, since studies show that people convey considerable amount of information through non-verbal cues [21]. For instance, through gaze [1, 13], head movements [25], and gestures [6]. Our focus is on combining three modalities - speech, head movements, and pointing gestures.

While originally relationships between modalities were encoded through a heuristic approach [3], currently probabilistic graphical models [17, 20] and deep learning [24, 27] are the most common ways to handle the multimodal representation. We are interested in investigating multimodal behaviour patterns and modelling them explicitly in a probabilistic manner. Thus we implemented multimodal fusion as a Bayesian filter, which already showed promising results for reference resolution [26].

More specifically, Whitney et al. [26] developed a Bayesian filter to calculate the belief of an object being the target given observed person’s gestures and speech. In this approach, the longer a person is pointing at an object, the more probable it is to be the target. Basing prediction on the longest fixation in continuous modalities such as pointing and gaze [9] are a common way to model them. However, as was shown in our previous study [10], when the complexity of the task is increased, nosiness of these modalities increases accordingly to such an extent that it becomes nearly impossible to make a prediction solely from the longest fixation.

Behaviour studies [2] showed that gestures are temporally coordinated with speech, when people are retelling a scene from their favourite movie. Based on these findings, we hypothesise that by incorporating timing of gestures and head fixations with relation to speech in our model, we can filter out unrelated to

the request non-verbal behaviour. In our work, we expand Whitney’s framework [26] by learning a temporal prior and adding it to the observations update. Our focus is on the temporal relationships between events in modalities and whether this prior can increase filter’s accuracy.

3 Methodology

3.1 Bayesian Filter

Having three continuous modalities, we want to fuse them together in order to get a probability distribution over all objects (\mathcal{X}), given observed speech, head, and pointing inputs (\mathcal{Z}) at each time step (t). We apply probabilistic inference based on Bayesian filtering [22]. The hidden state, $x_t \in \mathcal{X}$, is the target object in the scene that the person is currently referencing. The robot observes the user’s non-verbal actions and speech, \mathcal{Z} , and at each time step estimates a distribution over the current state, x_t (Fig. 2):

$$p(x_t | z_{0:t}). \quad (1)$$

First, a prediction about current state is made based only on previous observations and then two types of update are made: time update without any contextual information and observation update, as proposed in [26].

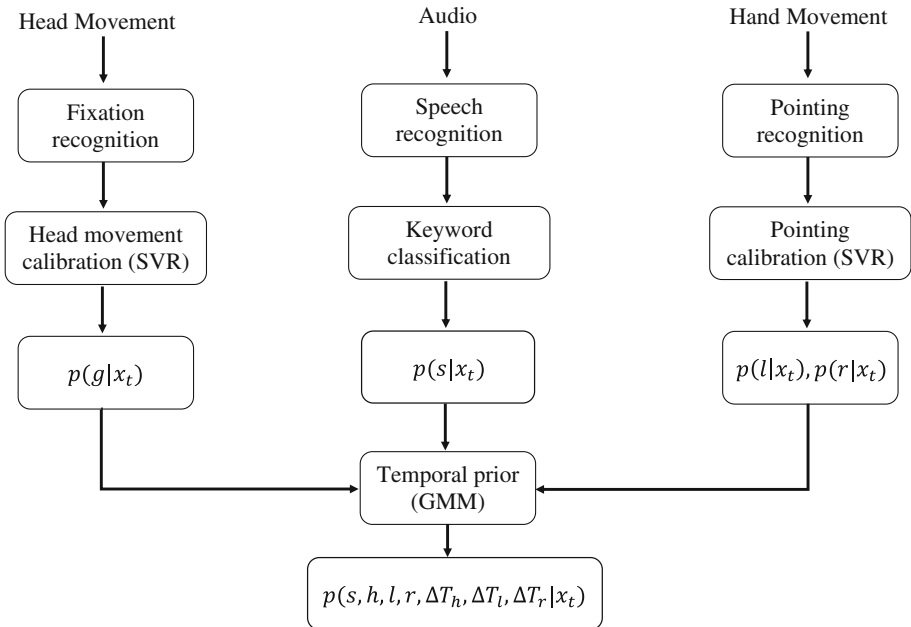


Fig. 2. Observation update with a temporal prior

3.2 Observation Update

The posterior distribution of x_t given a history of observations, $p(x_t|z_{0:t})$, also known as the belief $\mathcal{B}(x_t)$, is obtained using the Bayesian rule:

$$\mathcal{B}(x_t) = p(x_t|z_{0:t}) = \frac{p(z_t|x_t) \times p(x_t|z_{0:t-1})}{p(z_t|z_{0:t-1})} \propto p(z_t|x_t)p(x_t|z_{0:t-1}). \quad (2)$$

By substituting $p(x_t|z_{0:t-1}) = \sum_{x_{t-1} \in \mathcal{X}} p(x_t|x_{t-1})p(x_{t-1}|z_{0:t-1})$ in the above equation (considering Markovian properties), the Bayes filter algorithm can be used to obtain a recursive update rule:

$$\mathcal{B}(x_t) = p(z_t|x_t) \sum_{x_{t-1} \in \mathcal{X}} p(x_t|x_{t-1})\mathcal{B}(x_{t-1}), \quad (3)$$

where, $p(x_t|x_{t-1})$ is the transition probability found similarly as in [26],

$$p(x_t|x_{t-1}) = \begin{cases} c, & \text{if } x_t = x_{t-1}. \\ \frac{1-c}{|\mathcal{X}|-1}, & \text{otherwise.} \end{cases}, \quad (4)$$

where c is a constant value.

The observation model calculates the probability of the observation given the state. Each observation is a set of the user's head movement, hands pointing, and speech, $\langle h, l, r, s \rangle$ where:

- h represents a 3D vector of roll, pitch and yaw angles of the head.
- l represents a 3D vector as the direction of the participant's left index finger.
- r represents a 3D vector as the direction of the participant's right index finger.
- s represents a list of words uttered by the participant.

More formally, the observation model looks as follows:

$$p(z_t|x_t) = p(h, l, r, s|x_t). \quad (5)$$

We factor the expression by assuming that each observation is conditionally independent of the others given the target object. In other words, if we know the intended target object, knowledge about e.g., right hand pointing does not provide any further information about the head movements. This results in the following factorization:

$$p(z_t|x_t) = p(h_t|x_t) \times p(l_t|x_t) \times p(r_t|x_t) \times p(s_t|x_t). \quad (6)$$

In the following, we discuss how the above likelihoods can be modelled in our proposed approach.

Head Movement. We first learn a model $p \leftarrow f_h(h)$ that maps an angular position of the participant head (h) into a 2D position on the table (p) where the participant is looking at. Following the guidelines of the device [8], we calibrated it as an eye-tracker by training a Support Vector Regression (SVR) [19] with a RBF kernel ($C=10$, $\text{gamma}=5$) on 14 known points on the table. Participants were asked to look at each point for a duration of 1950 ms out of which the first 700 ms period was ignored. This calibration process results in ± 4.85 cm gaze positioning error.

Similar to the earlier study [26], we assign distributions over different head angular positions according to the distance between the corresponding gaze location and the target object location, i.e.,

$$p(h_t|x_t = i) \propto \exp[-(f_h(h_t) - p_i)^T \Sigma_h (f_h(h_t) - p_{x_t})], \quad (7)$$

where, p_i is the position of the i_{th} object on the table, and Σ_h is a diagonal co-variance matrix with trainable parameters.

Hand Gestures. Similarly, two separate SVM models are trained to map the directions of the left ($p \leftarrow f_l(l)$) and right ($p \leftarrow f_r(r)$) hands to the corresponding 2D positions on the table. Pointing detection is made with the help of LeapMotion device. The same expression as in Eq. 7 is used to assign probability distributions over left l_t and right r_t hand pointing directions conditioned on the target object x_t .

Speech. First, we use speech recognition to convert audio to text and then perform keywords dictionary-based classification to identify the following speech events:

- attribute - adjectives that describes size, shape or colour of a Lego block from the workspace, e.i. *red, large, cylinder, etc.*
- deictic - i.e. *here, there, this, that, etc.*
- other - any other word that is not included into the previous two categories

As an extra speech event, the beginning of a verbal request is detected from the audio directly. These specific classes were inferred from the initial data collection. The highest correlation was shown between them and events in other modalities.

After detecting speech events, we represent it with a unigram model. Namely, we take each word w in a given speech input s and calculate the probability that, given state x_t , that word would have been spoken.

$$p(s|x) = \prod_{w \in s} p(w|x) \quad (8)$$

3.3 Temporal Priors

The main hypothesis of this paper, is whether incorporating the knowledge of temporal correlations between high-level events in the input modalities can help the robot to better understand the intentions of a human while the person is referring to something. In order to validate this hypothesis, we propose to use temporal conditional probabilities to represent the observation likelihood introduced in Eq. 2. This yields,

$$\begin{aligned} p(s_t, h_t, \Delta T_h, l_t, \Delta T_l, r_t, \Delta T_r | x_t) = \\ = p(s_t | x_t) \times p(h_t | x_t) \times p(l_t | x_t) \times p(r_t | x_t) \times p(\Delta T_h, \Delta T_l, \Delta T_r | x_t), \end{aligned} \quad (9)$$

where $\Delta T_h = T_s - T_h$, i.e., the time difference between the speech and head movement events. Similarly, $\Delta T_l = T_s - T_l$ and $\Delta T_r = T_s - T_r$. We used T_s as the time reference, since it is less affected by noises compared to the other modalities. Furthermore, we assumed independence between, e.g., the current value of the head position and its event time. However, the time differences between the events are highly correlated.

In this paper, multivariate Gaussian Mixture Models (GMMs) are used to represent the PDF of $p(\Delta T_h, \Delta T_l, \Delta T_r | x_t)$. We assume that modalities are temporally dependent; thus, the co-variance matrix is learned alongside kernels' means. We train GMMs with an Expectation Maximisation (EM) algorithm. Online adaptation of the model is performed via Maximum A Posteriori (MAP) estimation approach, as in Reynolds' work [16], due to a limited number of samples that we are able to collect during the interaction.

4 Data Collection

In order to test our hypothesis of existing temporal dependencies in behaviour patterns, we collected data of people interacting with a robot controlled by a wizard of Oz.

Our previous experiments showed that human behaviour varies dramatically in human-human interaction versus human-robot interaction. For instance, if the person has a human partner, they were more prone to use gestures and look at their partner. On the other hand, with a robot partner, participants favoured other modalities, like speech and exaggerated head movements. Moreover, robots can interpret some modalities with more precision than humans do. For humans, it is easier to understand where the person is pointing at than where they are looking. For robots using head movement tracking sensors, this modality becomes much more precise and easy to interpret than hand gestures. Our goal is to recreate the data collection scenario as close as possible to the target settings of the real human-robot interaction we plan our robot to operate in.




Fig. 3. First person view of the mixed reality interface. The participant is instructed to request an Lego block, that has an augmented orange circle around it. Robot’s guess is indicated as a white cylinder. The virtual robot shows the future trajectory of a picking up motion before the participant confirms whether the object was disambiguated correctly. (Color figure online)

4.1 Scenario

In our study, participants are instructed through a projection in a mixed reality headset to request Lego blocks (see Fig. 3) from the robot in an ambiguous environment, i.e., there are several blocks of the same colour and shape. Thus, it is impossible to disambiguate a human request only from speech and the interpretation of other modalities is necessary. Mixed reality was chosen as the way to convey robot’s current belief and augment additional information on the shared workspace, based on the results of our previous human study [18]. A human wizard interprets the human requests by looking at what a robot would be able to sense and tries to infer the intended object. When the human participant acknowledges that the robot understood which object the participant meant, the mixed reality headset suggests the next object for the participant to describe to the wizard. The wizard’s interface contains data from all the sensors. Namely,

- 3D position and rotation of the participant’s head tracked by the headset¹ and updated at frequency 60 Hz;
- Projection from the centre of their head on the table;
- 3D coordinates of both hands, a projection from the index finger on the table when pointing occurs. The original frame rate of the sensor is 120 Hz but to align data streams of head and hand tracking, we record only each second

¹ <https://www.microsoft.com/en-us/hololens>.

- frame, resulting in 60 Hz frequency. Tracking is performed by the Leap Motion sensor²;
- Speech recognition represented as text, acquired from using Microsoft Speech Platform³. Speech data is recorded every time the dictation hypothesis is updated. We don't wait for the utterance to be completed, and instead employ a riskier but also faster approach (Fig. 4).



Fig. 4. Wizard interface with wizard's guess visualised as a white cylinder around an object. Multimodal input is represent as (a) text from speech recognition; (b) white rectangle being participant's head position and rotation, while a blue circle on the table surface is a projected vector from the centre of the head; (c) skeleton of a hand and a projected position from its finger on the table plane as a green circle. (Color figure online)

4.2 Participants

Subjects were recruited using mailing lists and flyers on the university campus. A total of 30 subjects (16 female, 14 male), ages between 23 and 34 ($M = 27.7$), participated in the data collection. All participants had to meet the following requirements: to be fluent in English, not require glasses to see objects 1.5-2 away from them (due to the mixed reality head-mounted display) and not have any colour vision deficiencies. In general, participants indicated their experience with digital technology as $M = 1.3$ on a scale from 5 to 1 (where 1 denotes "very

² <https://www.leapmotion.com/>.

³ [https://msdn.microsoft.com/en-us/library/hh361572\(v=office.14\).aspx](https://msdn.microsoft.com/en-us/library/hh361572(v=office.14).aspx).

highly”). Moreover, 57% had some experience with virtual reality and only 10% tried augmented/mixed reality head-mounted displays before.

4.3 Dataset

Each participant made a total of 20 fetching requests. The time of each request was not fixed; the start time of the request is considered to be the moment a participant was shown a new object in the Mixed Reality interface, and the request was considered resolved when a participant verbally confirmed the robot’s guess. Thus, enabling the data collection to proceed to the next randomly selected object and marking the current timestamp as the end of the request. Overall, 600 requests were collected with a total duration of 429 min of uninterrupted recording. Each request consisted of multiple datapoints with the following fields: a timestamp, a 3D vector of *head position*, a boolean variable representing whether the current head movement is a *fixation*, a 3D vector of each hand index *finger positions*, a boolean variable indicating whether the current gesture was *pointing*, the *text of the verbal request* so far from speech recognition, the current *target object*, and current *wizard’s guess* of the target object. The final dataset contains $N = 30705$ datapoints.

5 Results

This section presents our findings on the questions **Q1 - Q3**.

Q1: Do common temporal patterns emerge in participants’ behaviour during the fetching request task?

A pre-processing step is performed before training temporal priors encoded as GMMs. All fixations from both head movements and gestures are labelled as intentional or accidental. By intentional, we imply a fixation on the target object. All the other fixations are considered accidental, or noise. In a request, fixations on the minimum distance from the ground truth target object are labelled as intended. Moreover, time intervals of head fixations and pointing gestures are computed relative to the events in the speech modality. As a result, the training data set consists of time intervals and labels of head fixations, gestures and types of the corresponding event in the speech modality. Finally, through leave-one-participant-out cross validation, we train GMMs temporal priors on the training dataset. Analysis of the GMMs densest regions discovers three most common temporal patterns in participants’ behaviour, namely:

- **P1** Head fixation + beginning of the verbal request
- **P2** Head fixation + deictic keyword + pointing gesture
- **P3** Object attribute keyword + head fixation

Let’s observe how these patterns appear in the human-robot interaction during a fetching request. On the Fig. 5, we visualise a timeline of events in each modality during one of the participant’s request. This request contains all

three of the common patterns (highlighted with rectangles). In the beginning of the request, the participant, firstly, fixates his/her head on an unrelated to the request object. Then, he starts to verbalise his request with words “please gimme...”. Just before the beginning of the utterance, the first common pattern occurs (**P1**) - the participant fixates on the target object. After that several unintentional head fixations are detected alongside with the left hand pointing gesture in an uninformative direction. At the timestamp 800 ms the second pattern (**P2**) is detected - deictic keyword “this” with an intended head fixation and the left hand pointing at the target object. Later on, at timestamps 1300 ms and 1800 ms we can observe **P3**. The participant clarifies his request by saying an attribute key-word “small” and “blue” simultaneously with head fixations on the target object.




Fig. 5. A timeline of events in each modality during request 3 from participant 27. The top line is speech recognition of participant’s request. Black rectangles represent fixations on the target object, white on any other. Grey rectangles indicate common behaviour patterns. (Color figure online)

Q2: If we encode these patterns as temporal priors, will they be helpful in inferring the intended object from multimodal referring expressions?

To answer this question, we compare performance of Bayesian filter with (**BF+TP**) and without (**BF**) temporal priors through leave-one-participant-out cross validation. Bayesian filter without temporal priors (**BF**) is considered as a control condition. Our evaluation consists of two measures - **accuracy** (%) and decision making **time** (sec). Accuracy is measured as a ratio of correctly disambiguated target objects to the total number of requests. Decision making time is computed from the moment a person started speaking and until the robot makes a final decision which object is the target.

There are two possible scenarios of how a model can make a decision.

- *Voluntary decision.* We define a decision making line as 85%, a commonly used value for such scenarios [26]. This means that if probability of an object being the target reaches 85% or higher, given all previous multimodal observations, then the model selects it as a target object.
- *Forced decision.* If during a request no object’s probability to be the target one crossed the decision line and there are no datapoints left in this request, then the object with the highest probability is chosen the target.

Models have to make forced decisions due to the way we collected the dataset. A human wizard was imitating a robot during the data collection process; therefore, a request was considered resolved when the wizard guessed the correct target object. However, our models have a limited understanding of multimodal human behaviour and are not as sophisticated as a human wizard reasoning.

To find out what causes the main issues for the models, compared to the human wizard, we analysed only cases where models had to make forced decisions. We manually checked video recordings of the interactions and discovered that the majority of such cases contained utterances such as “to the left”, “in the corner”, “he same as the previous one” A human is able to infer much more information from such phrases, while our speech system is only corpus-based. This bottleneck can be addressed by implementing a more sophisticated natural language recognition system, for instance, BERT [5].

For models’ evaluation we take into account both voluntary and forced decisions.

Two types of evaluation is performed. In the first case, only the first attempt of the model to make a decision is considered. In the second, though, we employ the same way of interaction as during the data collection phase. The model can make several attempts to guess the target object while there is still data left in the request. After each guess it gets feedback whether the guess is correct or not. If it is incorrect, the model excludes the previous object from the possible objects set and proceeds disambiguation participant’s request as before. For the multiple attempts case, the decision making time is measured from the beginning and until the guess is either correct or the request is over and the model is forced to make the final guess.

We performed a repeated measurement one-tailed t-test to test significance of our results on the 95% interval. According to the Table 1, **BF+TP** dramatically ($p < 0.00$) decreases decision making time from 24.99 ± 7.94 sec to 15.32 ± 3.08 sec. Accuracy of the Bayesian Filter with Temporal Priors ($M = 68.45$, $SE = 5.73$) is also significantly ($p < 0.00$) higher than without ($M = 55.83$, $SE = 12.01$).

In the multiple attempts case (Table 2), the tendency of **BF+TP** ($M_{time} = 18.85$, $SE_{time} = 3.73$, $M_{acc} = 86.22$, $SE_{acc} = 4.34$) outperforming BF on both measurements is even more evident. Time and accuracy of **BF** does not significantly ($p > 0.05$) change from the first attempt. Number of attempts per request gives us insight into why this is the case. For BF, number of attempts is nearly one per request ($M = 1.14$), while **BF+TP** can make 2.58 on average. An explanation to this can be drawn from the Table 2, specifically the decision making time. **BF** takes approximately 1.6 times more to make the first decision and it does not have enough time left of the request to make an accurate guess on the second attempt. In a multiple attempt scenario **BF+TP** model can potentially make more guesses on the same amount of data, while being more accurate than the control condition.

Therefore, we can conclude that temporal priors have a significant positive influence on both decision making time and accuracy for the both evaluation scenarios.

Table 1. Models evaluation on the first attempt at disambiguating a multimodal referring expression

Model	Time (sec)	Accuracy
BF	24.99 ± 7.94	55.83% ± 12.01%
BF+TP	15.32 ± 3.08	68.45% ± 5.73%
BF+TP+OA	13.41 ± 2.93	76.58% ± 5.65%

Table 2. Models evaluation on multiple number of attempts at disambiguating a multimodal referring expression

Model	Time (sec)	Accuracy	# Attempts
BF	25.50 ± 8.11	57.38% ± 12.24%	1.14
BF+TP	18.85 ± 3.73	86.22% ± 4.34%	2.58
BF+TP+OA	18.92 ± 3.70	89.16% ± 4.28%	2.63

Q3: Are temporal patterns common across most participants or person-dependent?

Our approach to **Q3** is to evaluate what is the effect of online adaptation (**BF+TP+OA**) on the decision making time and accuracy versus no participant-based adaptation (**BF+TP**). The model with online adaptation of temporal priors is performed in the same fashion as in the previous section, through leave-one-participant-out cross validation. However, we iteratively update the GMMs temporal priors by feeding them datapoints from the previously resolved request. The following request disambiguation is made with the refitted with all the previous requests model. This implies that with time, the GMMs become more fitted to the temporal patterns of this particular participant. The first requests of each participant in both models with and without adaptation are based on the same temporal priors GMMs.

Our results show that while **BF+TP+OA** accuracy ($M = 76.58, SE = 5.65$) increases significantly ($p = 0.04$) in comparison to the temporal priors without online adaption during the first attempt (Table 1). Decision making time ($M = 13.41, SE = 2.93$), even though has a decreasing trend, is not statistically significant ($p = 0.13$).

For the multiple attempts case (Table 2), adaptation results also show the best accuracy ($M = 89.16, SE = 4.28$) out of the evaluated models. In regards of decision making time ($M = 18.92, SE = 4.28$) and number of attempts, there

was no statistically significant difference found between models **BF+TP** and **BF+TP+OA**.

We can reason that adaptation has a positive effect on model’s accuracy, slightly adjusting temporal priors for each participant. The structure of the common patterns stays mostly the same between participants, while the Gaussian Mixtures shift to accommodate to the different timing of each participant individually.

6 Conclusion

In this paper, we explored temporal dependencies in multimodal human-robot interaction and developed a Bayesian-based model to evaluate our hypothesis. We developed a system in Mixed Reality to efficiently collect data of humans interacting with a robot in a fetching scenario. As our results showed, taking temporal dependencies between high-level events in all input modalities (i.e. fixations in head movements, key words in speech, etc.) increases the model’s speed and accuracy of the person’s intention predictions. Moreover, we tested how online adaptation influences results of the prediction and found out that, while both speed and accuracy increase, the change is not as dramatic as between using a Bayesian filter with or without temporal priors. Thus, we came to a conclusion that common temporal patterns exist in human behaviour during referencing objects and have a significant impact on the intention prediction. We encoded temporal priors as a Gaussian Mixture Model and used it with the Bayesian filter to compute probabilities of objects being the target ones.

The next step for this project is to test how scalable our approach is to more complex tasks. Our initial motivation to explore temporal dependencies comes from our previous work [10], where participants were building furniture together. The main challenge there came from the nosiness of input modalities. And the more complex the interaction, the nosier participant’s behaviour is. In other words, participants are less distracted and more focused on the task during simple interactions, such as fetching requests. We see the potential benefits of employing temporal priors to tackle nosiness in the more complex interactions.

Another direction will be to add more modalities and explore how they can be represented as high-level events and encoded into temporal behaviour patterns. For instance, body posture and gaze tracking. A more nuanced, not key-words based, approach to natural language understanding can also enrich the possibilities for diverse interactions.

And finally, in the future we would like to focus more on the deep reinforcement learning approaches to multimodal human-interaction. So far our study was necessary for gaining a deeper understanding of human behaviour and multimodal data. However, we want to move away from feature engineering and formulate our human-robot interaction scenario as a deep reinforcement learning problem. Recent studies in HRI showed impressive results in employing deep reinforcement learning for various applications [12, 14, 15]. The main challenge for deep learning approaches is the lack of training data from human studies

but we plan to tackle this problem using our current Bayesian-based model to simulate human behaviour data as a prior for the deep reinforcement learning model.

Acknowledgements. This work is supported by the SSF (Swedish Foundation for Strategic Research) projects COIN.

References

1. Admoni, H., Srinivasa, S.: Predicting user intent through eye gaze for shared autonomy. In: Proceedings of the AAAI Fall Symposium Series: Shared Autonomy in Research and Practice (AAAI Fall Symposium), pp. 298–303 (2016)
2. Bavelas, J., Gerwing, J., Healing, S.: Hand and facial gestures in conversational interaction. In: Holtgraves, T.M. (ed.) The Oxford Handbook of Language and Social Psychology, pp. 111–130. Oxford University Press, Oxford (2014)
3. Bolt, R.A.: “Put-that-there”: voice and gesture at the graphics interface, vol. 14. ACM (1980)
4. Chai, J.Y., et al.: Collaborative effort towards common ground in situated human-robot dialogue. In: Proceedings of the 2014 ACM/IEEE International Conference on Human-Robot Interaction, pp. 33–40. ACM (2014)
5. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: Bert: pre-training of deep bidirectional transformers for language understanding. arXiv preprint [arXiv:1810.04805](https://arxiv.org/abs/1810.04805) (2018)
6. Duarte, N., Tasevski, J., Coco, M., Raković, M., Santos-Victor, J.: Action anticipation: reading the intentions of humans and robots. arXiv preprint [arXiv:1802.02788](https://arxiv.org/abs/1802.02788) (2018)
7. Funakoshi, K., Nakano, M., Tokunaga, T., Iida, R.: A unified probabilistic approach to referring expressions. In: Proceedings of the 13th Annual Meeting of the Special Interest Group on Discourse and Dialogue, pp. 237–246. Association for Computational Linguistics (2012)
8. Harezlak, K., Kasprowski, P., Stasch, M.: Towards accurate eye tracker calibration-methods and procedures. Proc. Comput. Sci. **35**, 1073–1081 (2014)
9. Huang, C.M., Andrist, S., Sauppé, A., Mutlu, B.: Using gaze patterns to predict task intent in collaboration. Front. Psychol. **6**, 1049 (2015)
10. Kontogiorgos, D., Sibirtseva, E., Pereira, A., Skantze, G., Gustafson, J.: Multimodal reference resolution in collaborative assembly tasks. In: Proceedings of the 4th International Workshop on Multimodal Analyses Enabling Artificial Agents in Human-Machine Interaction, pp. 38–42. ACM (2018)
11. Lalanne, D., Nigay, L., Robinson, P., Vanderdonckt, J., Ladry, J.F., et al.: Fusion engines for multimodal input: a survey. In: Proceedings of the 2009 international conference on Multimodal interfaces, pp. 153–160. ACM (2009)
12. Lathuilière, S., Massé, B., Mesejo, P., Horaud, R.: Neural network based reinforcement learning for audio-visual gaze control in human-robot interaction. Pattern Recogn. Lett. **118**, 61–71 (2018)
13. Mehlmann, G., Häring, M., Janowski, K., Baur, T., Gebhard, P., André, E.: Exploring a model of gaze for grounding in multimodal HRI. In: Proceedings of the 16th International Conference on Multimodal Interaction, pp. 247–254. ACM (2014)
14. Minh, V., et al.: Human-level control through deep reinforcement learning. Nature **518**(7540), 529 (2015)

15. Qureshi, A.H., Nakamura, Y., Yoshikawa, Y., Ishiguro, H.: Robot gains social intelligence through multimodal deep reinforcement learning. In: 2016 IEEE-RAS 16th International Conference on Humanoid Robots (Humanoids), pp. 745–751. IEEE (2016)
16. Reynolds, D.A., Quatieri, T.F., Dunn, R.B.: Speaker verification using adapted gaussian mixture models. *Digit. Sig. Process.* **10**(1–3), 19–41 (2000)
17. Savran, A., Cao, H., Nenkova, A., Verma, R.: Temporal bayesian fusion for affect sensing: combining video, audio, and lexical modalities. *IEEE Trans. Cybern.* **45**(9), 1927–1941 (2015)
18. Sibirtseva, E., et al.: A comparison of visualisation methods for disambiguating verbal requests in human-robot interaction. arXiv preprint [arXiv:1801.08760](https://arxiv.org/abs/1801.08760) (2018)
19. Smola, A.J., Schölkopf, B.: A tutorial on support vector regression. *Stat. Comput.* **14**(3), 199–222 (2004)
20. Srivastava, N., Salakhutdinov, R.: Learning representations for multimodal data with deep belief nets. In: International Conference on Machine Learning Workshop, vol. 79 (2012)
21. Thomaz, A., Hoffman, G., Cakmak, M., et al.: Computational human-robot interaction. *Found. Trends® Robot.* **4**(2–3), 105–223 (2016)
22. Thrun, S., Burgard, W., Fox, D.: Probabilistic Robotics. MIT Press, Cambridge (2005)
23. Turk, M.: Multimodal interaction: a review. *Pattern Recogn. Lett.* **36**, 189–195 (2014)
24. Venugopalan, S., Xu, H., Donahue, J., Rohrbach, M., Mooney, R., Saenko, K.: Translating videos to natural language using deep recurrent neural networks. arXiv preprint [arXiv:1412.4729](https://arxiv.org/abs/1412.4729) (2014)
25. Veronese, A., Racca, M., Pieters, R.S., Kyrki, V.: Probabilistic mapping of human visual attention from head pose estimation. *Front. Robot. AI* **4**, 53 (2017)
26. Whitney, D., Eldon, M., Oberlin, J., Tellex, S.: Interpreting multimodal referring expressions in real time. In: 2016 IEEE International Conference on Robotics and Automation (ICRA), pp. 3331–3338. IEEE (2016)
27. Yao, L., et al.: Describing videos by exploiting temporal structure. In: Proceedings of the IEEE International Conference on Computer Vision, pp. 4507–4515 (2015)