

Didn't see that coming: a survey on non-verbal social human behavior forecasting

German Barquero

Johnny Núñez

Sergio Escalera

Universitat de Barcelona and Computer Vision Center, Spain

GBARQUGA9@ALUMNES.UB.EDU

JNUNEZCA11@ALUMNES.UB.EDU

SERGIO@MAIA.UB.ES

Zhen Xu

Wei-Wei Tu

4Paradigm, Beijing, China

XUZHEN@4PARADIGM.COM

TUWEIWEI@4PARADIGM.COM

Isabelle Guyon

LISN (CNRS/INRIA) Université Paris-Saclay, France, and ChaLearn, USA

GUYON@CHALEARN.ORG

Cristina Palmero

Universitat de Barcelona and Computer Vision Center, Spain

CRPALMEC7@ALUMNES.UB.EDU

Abstract

Non-verbal social human behavior forecasting has increasingly attracted the interest of the research community in recent years. Its direct applications to human-robot interaction and socially-aware human motion generation make it a very attractive field. In this survey, we define the behavior forecasting problem for multiple interactive agents in a generic way that aims at unifying the fields of social signals prediction and human motion forecasting, traditionally separated. We hold that both problem formulations refer to the same conceptual problem, and identify many shared fundamental challenges: future stochasticity, context awareness, history exploitation, etc. We also propose a taxonomy that comprises methods published in the last 5 years in a very informative way and describes the current main concerns of the community with regard to this problem. In order to promote further research on this field, we also provide a summarized and friendly overview of audiovisual datasets featuring non-acted social interactions. Finally, we describe the most common metrics used in this task and their particular issues.

Keywords: Behavior forecasting, Human motion prediction, Social signal prediction, Social robots, Socially interactive agents, Dyadic interactions, Triadic interactions, Multi-party interactions, Backchanneling, Engagement

1. Introduction

Communication among humans is extremely complex. It involves an exchange of a continuous stream of social signals among interactants, to which we adapt and respond back. These social signals are manifested as non-verbal behavioral cues like facial expressions, body poses, hands gestures, or vocal feedback. We, as humans, have the innate capability of identifying, understanding, and processing social cues and signals, which is the core of our social intelligence (Vinciarelli et al., 2009). Similarly, we are also inherently capable of anticipating, up to some extent, these social signals. For instance, we do not need the speaker to actually end their speech before we know that a turn-taking event is close (Ondáš and Pleva, 2019). We are prepared in advance. In a similar way, we can anticipate a social

action like a handshake by the correct observation and interpretation of simultaneously occurring visual cues from the other interlocutor, like a verbal greeting while their hand is approaching. In fact, recent works in neuroscience believe that such anticipation is the motor for cognition. In particular, this current, called *predictive processing*, supports the idea that the brain is constantly generating and updating a mental model of the environment by comparing predicted behaviors to actual observations (Walsh et al., 2020). Very interestingly, some works have successfully observed interpersonal predictive processing signals during social interactions (Thornton et al., 2019; Okruszek et al., 2019). This suggests the importance that behavior forecasting may have as a pathway to the ultimate behavioral model.

If successfully modeled, such forecasting capabilities can enhance human-robot interaction in many applications. For instance, wherever turn-taking events are frequent and very dynamic (e.g., multi-party conversations), any degree of anticipation is extremely beneficial. Being able to anticipate the next speaker or when a listener will disengage are key to efficiently handle such situations. Also, forecasting has direct applications to robot behavior synthesis for social interactions in two directions. First, being able to anticipate the interactants’ behavior can help the agent to behave in a more socially-aware way. And second, being able to predict one’s own behavior, even by only few milliseconds, can save a valuable computation time. In fact, the longer the future that we are able to predict, the better the robot can prepare for the execution of a movement.

Unfortunately, providing robots or virtual agents with social forecasting capabilities is extremely difficult due to the numerous particularities of the problem. First, the large amount of variables driving a social interaction makes it a highly dimensional problem. For instance, in the previous handshake example, even if the agent detects the approach of the hand as a visual cue and anticipates a handshake, it could miss in the case where the hand is grabbing an object. On top of that, predicting the future always poses problems related to its stochasticity. The plausibility of several equally probable future scenarios complicates the development and evaluation of forecasting models.

In the past years, research on non-verbal social behavior forecasting has followed distinct paths for social signal prediction and computer vision fields, although they share most of their fundamental concerns. For example, the human motion forecasting field does not usually refer to any social signal forecasting work (Mourot et al., 2021), even though some of them predict visual social cues, or action labels. And vice versa. This survey wants to unify non-verbal social behavior forecasting for both fields, describe its main challenges, and analyze how they are currently being addressed. To do so, we establish a taxonomy which comprises all methodologies applied to multi-agent (human-human, or human-robot/virtual agent) scenarios and presented in the most recent years (2017-2021). In particular, we focus on works that exploit at least one visual cue. We also engage in a discussion where we foresee some methodological gaps that might become future trends in this field. Besides, we summarize the publicly available datasets of social interactions into a comprehensive and friendly survey. Finally, we present and discuss on the usual evaluation metrics for non-verbal social behavior forecasting.

This survey is organized as follows. First, in Section 2, we formulate the non-verbal social human behavior forecasting problem (Figure 1), and introduce a taxonomy for it. We start by identifying and discussing the main challenges associated to the task, and describe

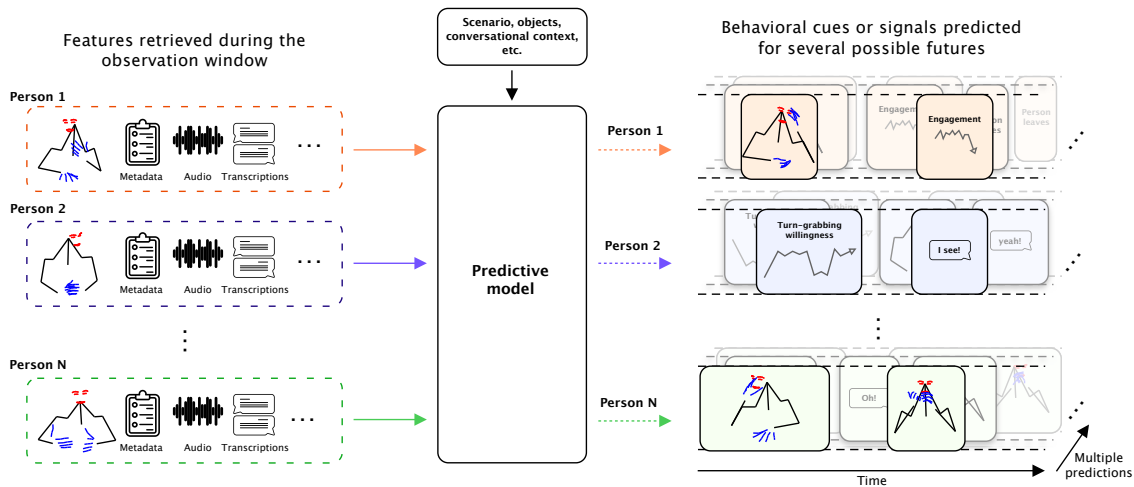


Figure 1: Visual representation of the generic social behavior forecasting problem reviewed in this work. Given a set of features of N multiple interactants observed from the past and the contextual information about the interaction, M future sequences composed of behavioral social cues/signals are predicted.

how past works have addressed them (Sections 2.1 to 2.5). Then, we thoroughly review the state-of-the-art methodologies proposed for non-verbal social behavior forecasting. In particular, we split the survey into methods predicting low-level (e.g., landmarks, facial action units) and high-level (social cues and signals) representations of non-verbal behavior, in Sections 2.6.1 and 2.6.2, respectively. Then, in Section 3, an extensive collection of datasets featuring audiovisual social interactions is presented. Datasets are classified according to the interaction scenario (dyadic, triadic, group, or several groups), task (conversational, collaborative, or competitive), and setting (participants are standing or seated). We also provide a summary table that allows the reader to easily compare the low- and high-level behavioral annotations provided as part of each dataset. Section 4 presents the most popular metrics for assessing the accuracy, diversity, and realism of the predicted behavior. In Section 5, we provide a discussion on general trends and current challenges, as well as possible future research directions for non-verbal social behavior forecasting. Finally, in Section 6, we review the ethical concerns regarding non-verbal social behavior forecasting and its real-world applications.

2. Taxonomy

Our taxonomy, see Figure 2, includes all approaches that predict non-verbal human behavior in socially interactive scenarios. These scenarios include at least two subjects socially interacting together, typically referred to as *focused interactions*. Also, we understand forecasting in the strictest sense of the word. Namely, only information observed *before the prediction starts* is used. Such constraints leave co-speech generative methods (Liu et al.,

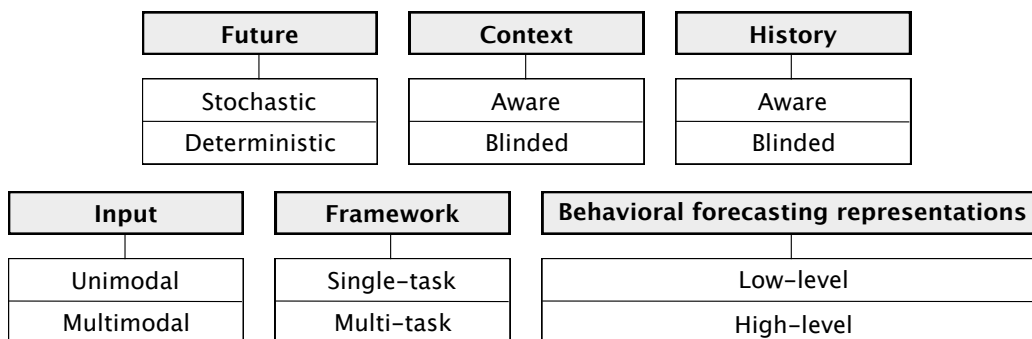


Figure 2: The proposed taxonomy for non-verbal social human behavior forecasting in socially interactive scenarios.

2021; Kucherenko et al., 2021) or pedestrian trajectories forecasting out of our scope. The former leverages the future speech and the latter does not usually feature a focused interaction. Additionally, the approaches needed to exploit *at least one visual cue* (e.g., landmarks, image, visually annotated behavioral labels). We acknowledge though many works that use exclusively lexical or audio features to predict non-verbal social cues such as backchannel opportunities or turn-taking events (Ortega et al., 2020; Jang et al., 2021). On the other hand, the taxonomy is very flexible with regards to the typology of predicted human behavior. Therefore, we include works that predict low-level behavioral representations such as landmarks, head pose or image (Table 1), but also high-level ones like social cues and social signals (Table 2) (Vinciarelli et al., 2009). We encourage the reader to accompany the survey lecture with those tables, as they provide a synthesized view of the methodologies and a comparison among them. In addition, we refer the reader to Figure 1 for an illustration of our definition of behavior forecasting.

Next, we present and describe the main challenges related to non-verbal social behavior forecasting that are currently being actively addressed by the community. They are enclosed within each dimension of our proposed taxonomy: future perspective (Section 2.1), context exploitation (Section 2.2), history awareness (Section 2.3), input modalities (Section 2.4), and framework (Section 2.5). Finally (Section 2.6), we detail and compare related works according to their behavioral forecasting representation (e.g., landmarks, action labels).

2.1. Future perspective

A common approach in single-person behavior forecasting consists in embracing the future uncertainty and exploiting it by predicting multiple futures (*multimodal*, or *stochastic*) (Aliakbarian et al., 2021; Hassan et al., 2021; Mao et al., 2021a). However, this research line has not been fully exploited for social scenarios yet. Instead, most works propose methods which assume that the observed future is unique (*deterministic*) (Adeli et al., 2021; Guo et al., 2021; Wang et al., 2021b; Barquero et al., 2022), thus ignoring multiple future behaviors that may co-exist and be equally plausible. This hypothesis removes some challenges associated to stochastic approaches such as the sampling choice among several generated futures, or the assessment of the realism and plausibility of all predicted futures.

However, this simplification comes at a high cost: the predictive model is penalized for generating plausible and realistic behaviors which do not match the ones observed in the dataset. To alleviate this, many works reduce the dimensionality of the forecasting objective (e.g., action labels) (Sanghvi et al., 2020; Airale et al., 2021), or provide extensive contextual information in order to narrow the future space and therefore reduce its stochasticity (Corona et al., 2020b; Adeli et al., 2020, 2021). Still, some works forecasting low-level behavioral representations (e.g., landmarks) report a strong regression-to-the-mean effect in the predictions (Barquero et al., 2022). Some works tried to tackle this problem in several contexts. For example, Feng et al. (2017) proposed building a specific high-frequencies predictor which made the generated facial expressions more realistic. In the context of social signals forecasting, Raman et al. (2021) reasoned that such effect was linked to the availability of similar future signals triggered at different future points. To mitigate it, they proposed to inject the time offset at which social signals triggered into the past encoding. In general though, deterministic works complement the quantitative evaluation with qualitative visualizations that help assess the realism and smoothness of the predictions.

2.2. Context exploitation

Social interaction among humans is dynamic and strongly influenced by many external factors. The human behavioral model that drives a conversation between a professor with a student is drastically different from that of a conversation among friends (Reis et al., 2000). Even with the same interactants setting, the place where the interaction happens (e.g., in a bar, at a conference, at home) may change the whole dynamic of their behavior. In a similar way, a handshake might become a handover if the approaching hand is holding an object (Shu et al., 2016; Corona et al., 2020b). Although considering all external factors that might influence the behavior is still impractical, some works consider using some contextual information. Accordingly, we split between *context-aware* methods, which were introduced by Corona et al. (2020b) and consider at least one modality of contextual information, and *context-blinded* methods, which focus on the target person only (Wang et al., 2021a). Most context-aware methods reviewed in Section 2.6 leveraged the partners’ behavioral information. Additionally, other works introduced approaches that also considered the presence and trajectory of objects (Shu et al., 2016; Corona et al., 2020b; Adeli et al., 2021), or even the whole visual scene (Adeli et al., 2020). These methods prove particularly useful in contexts where the behavior is strongly driven by the interaction with the scene.

2.3. History awareness

By definition, social interactions evolve over time, generating multiple long-range temporal dependencies. An event at the beginning of an interaction may impact and alter the rest of it. Furthermore, forecasting may benefit from observing very long sequences (e.g., >10 seconds) of interactions in order to tune a generic behavioral model to work with the interactants and the specific conversational context.

Although few works attempt to exploit the history in the single-person motion forecasting domain (Mao et al., 2020, 2021b), there are few social *history-aware* works. Although they do not detail how long the history can be, Chu et al. (2018) encoded a history of past text sequences and facial expressions with variable length to improve their forecasting

capabilities. [Guo et al. \(2021\)](#) and [Katircioglu et al. \(2021\)](#) incorporated motion attention to propagate observed motions to the future, theoretically even when the motion has not been seen in training time. However, both considered small histories of 2 seconds, which only favors the propagation of short repetitive motion. We are not aware of methods that consider much longer historical data, or that learn in an online and adaptive fashion the unique characteristics of each person’s behavior.

2.4. Input modalities

In addition to the interaction context, the speech, voice tone or other information related to the person of interest or the others may influence their behavior. Naturally, such multimodal data needs to be exploited in a specific way in order to fully profit from it. Therefore, we distinguish between *unimodal* and *multimodal* methods, which combine the visual modality with at least another modality as input to make their predictions.

Most common multimodal settings combined landmarks, body/head pose, or visual cues with past utterance transcriptions ([Chu et al., 2018](#); [Hua et al., 2019](#); [Ueno et al., 2020](#); [Barquero et al., 2022](#)), acoustic features ([Türker et al., 2018](#); [Ahuja et al., 2019](#); [Ueno et al., 2020](#); [Goswami et al., 2020](#); [Woo et al., 2021](#); [Jain and Leekha, 2021](#); [Murray et al., 2021](#); [Ben-Youssef et al., 2021](#)), speaker’s metadata ([Raman et al., 2021](#); [Barquero et al., 2022](#)), or with combinations of the previous modalities ([Ishii et al., 2020](#); [Huang et al., 2020](#); [Blache et al., 2020](#); [Ishii et al., 2021](#); [Boudin et al., 2021](#)). The most common way to exploit different modalities together consists in simply concatenating their embedded representations. Although this has proven to work for several applications, extracting relevant information from multiple modalities is not always straightforward ([Barquero et al., 2022](#)).

2.5. Framework

During the course of an interaction, humans exchange multiple social signals. Turn taking, agreement, politeness, empathy, disengagement, etc, are some examples. In some cases, such signals might be inferred from the same set of social cues, which is the perfect environment for *multi-task* learning. This paradigm, which consists in learning several tasks with the same model, has already helped to improve the results of single task models in many other fields ([Zhang and Yang, 2021](#)). In our context, few works have explored multi-task frameworks. [Ishii et al. \(2020, 2021\)](#) explored the benefits from predicting several social signals and cues at the same time (turn taking, turn-grabbing willingness, and backchannel responses). [Chu et al. \(2018\)](#) proposed a method to predict the next facial Action Units (AUs) by also predicting the future speech content.

2.6. Behavioral forecasting representations

2.6.1. LOW-LEVEL

When interacting with humans, virtual or robotic agents need to be able to reciprocate non-verbal cues at all dimensions. Given the relevance of this problem, many works aim at forecasting low-level representations of non-verbal social cues, see [Table 1](#). Such representations can be non-semantic representations such as raw image or audio, or semantic ones such as landmarks, head pose, or gaze directions.

Authors	Method highlights	Multimodal input	Context [†]	→	→	Prediction	Scenario
Face							
Huang and Khan 2017	Two GANs. Face image generation from the partner’s past expressions.	✗	Partner	5s	1f	FL*	Remote dyadic conv.
Feng et al. 2017	VAEs. Parallel low-frequency and high-frequency models favor realism.	✗	Partner	3s	Ⓒ ₅	FL	Remote dyadic conv.
Chu et al. 2018	(Bi-) LSTMs. Multi-task (text+face) setting trained with RL.	Transcripts	Partner	Variable	IPU	AU	Movies dyadic conv.
Chen et al. 2019	LSTM+GAN. Face image generation from the partner’s past expressions.	✗	Partner	0.4s	Ⓒ ₁₀	AU+HP*	TV dyadic conv.
Ueno et al. 2020	Bi-GRU. Attention among sequential embeddings of input modalities.	Transcripts +Audio	Partner	Variable	1f	AU	Triadic conv.
Woo et al. 2021	LSTM. Simultaneous prediction for both participants.	Audio	Partner	0.8s	Ⓒ ₁	AU+HP	Dyadic conv.
Pose (upper or full body)							
Shu et al. 2016	MCMC. Joints functional grouping and sub-events learning.	✗	Partner +Object	0.4s	Ⓒ ₅	BL	Constrained social actions
Ahuja et al. 2019	LSTM/TCN. Dynamically attends to monadic and dyadic behavior models.	Audio	Partner	Variable	1f	BL	Dyadic conv.
Hua et al. 2019	LSTM. Distinct models while speaking (co-speech) and listening (forecasting).	Transcripts	Partner	2.8s	2.8s	UBL	Dyadic conv.
Honda et al. 2020	LSTM/GRU. Joint encoding and decoding for both interactants.	✗	Partner	0.5s	1s	BL	Competitive fencing
Corona et al. 2020b	GATs+RNNs. Interactions among objects and subjects modeled.	✗	Partners +Object	1s	2s	BL	Human-object interactions
Adeli et al. 2020	GRU. Scene understanding and multi-person encoding (social pooling).	Raw image	Partners +Scene	0.6s/ 1.3s	0.6s/ 0.4s	BL	In-the-wild interactions
Adeli et al. 2021	GATs+RNNs. Interactions among objects, subjects and scene modeled.	Raw image	Partner +Scene +Objects	0.6s/1s	0.6s/1s	BL	In-the-wild interactions
Raman et al. 2021	GRU/MLP. Social processes definition and prediction offset injection.	Speaking status	Partners’ features	10f	10f	HP+Body Pose	Triadic conv.
Yasar and Iqbal 2021	GRU+Attention. Interpretable latent space. Cross-agent attention.	✗	Partners	-	0.6/1.6s	BL	Diverse interactions
Wang et al. 2021b	Transformer+DCT. Local- and global-range transformers.	✗	Partners	1s	3s	BL	Groups of interactions
Guo et al. 2021	Transformer+GCN+DCT. Cross-interaction motion attention (early-fusion).	✗	Partner	2s	Ⓒ ₄	BL	Dancing interactions
Katircioglu et al. 2021	Transformer+GCN+DCT. Pairwise motion attention (late-fusion).	✗	Partners	2s	1s	BL	Dancing interactions
Wang et al. 2021a	GCN+DCT. Strong and simple baseline with training tricks.	✗	✗	0.6s	0.6s	BL	In-the-wild interactions
Whole body (face+pose+hands)							
Barquero et al. (2022)	LSTM/GRU, TCN, Transformers, and GCN. Weakly supervised with noisy labels.	Audio/ Transcripts /Metadata	Partner	4s	Ⓒ ₁₀ /2s	FL+UBL +HL	Dyadic conv.

Table 1: Summary of papers forecasting low-level representations of non-verbal behavior. All works are history-blinded, with deterministic future, and use at least one visual input modality. Abbreviations: →|, observation window length; |→, prediction window length; Ⓒ_X, Future autoregressively predicted in steps of X frames (or seconds, when specified); 1f, only the immediate next frame is predicted; AU, action units; HP, head pose; FL, face landmarks; (U)BL, (upper) body landmarks; HL, hands landmarks; LSTM, long short-term memory; Bi, bidirectional; GRU, gated recurrent unit; VAE, variational autoencoder; GAN, generative adversarial network; DCT, discrete cosine transform; TCN, temporal convolutional network; MCMC, Markov chain Monte Carlo; RL, reinforcement learning. *: Incorporates image generation. †: partners’ information used matches the *Prediction* column.

Image and audio. In single human motion forecasting, we find methods that leverage pose motion prediction and generative methods to infer the future image frames (Walker et al., 2017; Zhao and Dou, 2020). There are few works proposing similar two-step approaches for image-based future social behavior forecasting. In an interview setting, Huang and Khan (2017) proposed a method to generate contextually valid facial images of the interviewer from the past interviewee’s facial expressions. To do so, they trained two Generative Adversarial Networks (GAN). The first one, conditioned on the interviewee’s recent facial expressions, produced the interviewer expression. The second one was trained to transform the generated expression into a real face image. Chen et al. (2019) proposed a face-to-face conversation system that also generated real-looking faces in two steps. They designed different models for forecasting behavior during speaking and listening phases. While the former was a co-speech generative method, the latter predicted the future AUs and head pose with a recurrent Long-Short Term Memory (LSTM) unit that only leveraged the past facial gestures of the speaker. Finally, a GAN conditioned on the predicted AUs and head pose generated the face image. Regarding the prediction of future raw audio output, there are no works that propose such architecture, to the best of our knowledge. A common path is to predict verbal behavior like textual content and include a Text-To-Speech model to generate the speech (Saeki et al., 2021).

Face. Most methods either focus on lower-dimensional representations of the face such as AUs or explicitly learnt representations. Regarding the latter, and aiming at replicating realistic facial gestures, Feng et al. (2017) proposed a Variational Auto-Encoder (VAE) that was trained to explicitly learn a lower-dimensional space for representing facial expressions. This bottleneck helped to reduce the dimensionality of the problem. Then, in order to promote the generation of subtle social cues (e.g., blinking, or eyebrow raising), the encoded past facial expressions of the user and an interactive embodied agent were processed by two specialized predictors, each focusing on either high or low frequencies. Very interestingly, instead of treating it as a regression problem, they clustered the learnt facial latent space and predicted the future expressions in the resulting discrete space. As a result, the regression-to-the-mean effect was mitigated. Chu et al. (2018) also proposed to detach both low- and high-frequency movements generation. Their multimodal model encoded a history representation of past text and facial gestures together (AUs) with the last observed text and facial expression. Then, the future text and the coarse and subtle face expressions were independently predicted in a multi-task setting that was trained with Reinforcement Learning (RL). Finally, they incorporated an adversarial discriminator that promoted the generation of diverse and realistic conversational behavior. Ueno et al. (2020) presented a multimodal approach that embedded text, visual, and audio sequence with bi-directional two-layered GRUs. Then, they were fused by an attention-weighted average layer to predict the face expression of the partner during the immediate feedback response. This visual response was then used to generate the textual feedback, resembling a multi-task setting. Unfortunately, this method cannot be applied to iteratively predict the evolution of the facial expression response in a pure forecasting fashion, as it uses modalities from the immediate previous step as input. Very recently, Woo et al. (2021) described an ongoing research that aims at leveraging audio and the context to predict the future facial expressions and head motion forecasting.

Pose. The very first attempt to forecast non-verbal body behavior in social interactions was carried out for robot learning of *social affordance* (Shu et al., 2016). In their work, Shu et al. presented a Markov Chain Monte Carlo (MCMC) based algorithm that iteratively discovered latent subevents, important joints, and their functional grouping. Their method also considered past trajectories of objects to successfully predict the agent’s behavior while performing handshakes, high-fives, or object handovers. Favored by the appearance of new and bigger datasets featuring social interactions (von Marcard et al., 2018; Andriluka et al., 2018; Joo et al., 2019a), Recurrent Neural Networks (RNNs) quickly became the standard in human motion forecasting (Martinez et al., 2017; Hua et al., 2019; Honda et al., 2020). However, Honda et al. (2020) observed that recurrent models used for single human motion forecasting are not suitable for highly interactive situations like fencing. In their work, they presented a general framework that provided single human motion forecasting methods with the ability to model interpersonal dynamics. To do so, both encoder and decoder LSTMs received as input the previous skeleton (either observed or predicted) concatenated with the hidden state of the opponent in the previous timestep. As a result, the simultaneous behavior forecasting of both players encoded the interpersonal dynamics of the interaction, making the predicted movements more accurate and coherent in the context of competitive fencing. While previous approaches focused on scenarios strongly driven by interpersonal dynamics, Ahuja et al. (2019) emphasized the imbalance between intrapersonal and interpersonal dynamics in dyadic conversations, with considerably less instances from the later. They warned that, in such scenarios, interpersonal dynamics could end up being ignored. To mitigate this issue, they proposed a dyadic residual-attention model (DRAM) that smoothly transitioned between monadic- and dyadic-driven behavior generation. Results showed that their model successfully identified non-verbal social cues like head nod mirroring or torso pose switching and generated proper reactions. Hua et al. (2019) also supported the use of the partner’s cues but restricted it to the modeling of the listener’s behavior. In their approach, they presented a human-robot body gesture interaction system built similarly to the system of Chen et al. (2019) for facial gestures synthesis. Similarly, they also leveraged two specialized methods for the speaking (co-speech generator) and listening (behavior forecasting) phases of the interaction. In contrast to Chen et al. (2019) though, they incorporated the speaker’s speech transcription as an extra predictive feature for the listener’s behavior.

Corona et al. (2020b) adverted to the fact that, on top of interactions with other humans, human motion is also inherently driven by interactions with objects. To model such interactions, they proposed a method which learnt a semantic graph of human-object interactions during the past observations. Then, the interactions graph was recurrently injected to a RNN in order to generate a context encoding. Both context vector and observed body poses were jointly decoded by a fully connected layer to predict the residuals (motion) of the next pose. As a result, their learnt behavioral model recognized and adapted to the particular dynamics of the scene. Additionally, they proposed to use the context vector to also predict the future motion of the scene objects, and update the context vector accordingly. They reported that their method was state of the art in what scene and human activities understanding refers. With a similar concept in mind, Adeli et al. (2020) proposed an action-agnostic context-aware method. The main difference is that they used spatio-temporal visual features directly extracted from the scene image, so-called context

features. Additionally, they introduced a social pooling module that merged the interactants’ behavior embeddings in a socially invariant feature vector. Then, the concatenated individual, social, and context features were decoded by a GRU module for each person. Differently from [Corona et al. \(2020b\)](#), the decoding stage did not take into account the interactants’ future behavior. In a newer work, [Adeli et al. \(2021\)](#) replaced the social pooling module by a graph attention network (GATs) that modeled interactions among individuals and objects. First, the historic of each person represented as joints-wise attention graph was fed to an RNN to get rid of the temporal dimension. Then, all RNNs outputs were used to build a human-human and a human-object graph attention network, which underwent an iterative message passing algorithm whose flow alternated between both of them. The respective social and context-aware encodings were concatenated to the spatio-temporal visual features and used as the initial hidden state of the RNN-based decoder. In contrast to their prior work, at each person’s decoding step, the hidden state was refined by the human-to-human attention graph in order to decode the future motion in a socially aware manner. Similarly to [Corona et al. \(2020b\)](#), they also observed that the socially aware decoding of the predictions improved the overall accuracy.

Very recently, several approaches introduced Transformer-like architectures ([Vaswani et al., 2017](#)) which outperformed previous RNN-based ones ([Yasar and Iqbal, 2021](#); [Wang et al., 2021b](#); [Guo et al., 2021](#); [Katircioglu et al., 2021](#)). [Yasar and Iqbal \(2021\)](#) proposed to encode individually the multiple agents’ joints positions, velocities, and accelerations. Then, cross-agent attention was applied among the latent space to generate socially aware representations, which followed two subsequent paths. First, these representations went through a two-streams adversarial discriminator that sampled discrete and continuous latent variables. The authors reported that such configuration favored the latent space interpretability. Their analysis on such variables showed that their method effectively captured the underlying dynamics of human motion. Finally, the socially aware latent representations underwent individual recurrent decoders that autoregressively predicted the future sequence of poses. The independent generation of poses represented its main limitation, as generated poses might not be socially coherent. [Wang et al. \(2021b\)](#) proposed to encode local- and global-range dependencies (intra- and inter-personal dependencies, respectively) with two specialized transformer encoders. The past motion of the person of interest was transformed by means of a Discrete Cosine Transform (DCT) ([Ahmed et al., 1974](#)), which was then fed to the local-range transformer performing self-attention. At the same time, the global-range transformer encoder applied self-attention across different subjects and different time steps. A spatial positional encoding was added to the global encodings to help the network cluster different individuals in different social interaction groups. Finally, the transformer decoder leveraged the last observed pose as the query, and the local- and global-range encodings as both keys and values in order to generate the whole predicted sequence at once, which was then fed to a linear and an Inverse DCT layer. Additionally, an adversarial loss was used to ensure the realism of the generated behavior. The authors argued that, by predicting the whole motion sequence at once, they prevented generating freezing motion. They reported state-of-the-art and qualitative impressive results in various datasets with several prediction window lengths (up to 3 seconds) and synthetically generated crowded scenarios (up to 15 people). [Guo et al. \(2021\)](#) provided the motion attention concept originally proposed by [Mao et al. \(2020\)](#) for single human motion prediction with a

mechanism to exploit the dyadic dynamics. To do so, they refined the keys and the values of both individuals by applying attention with those of the interactant (cross-interaction attention). The main benefit of motion attention is driven by its capacity of repeating historical patterns even for longer observed windows than the ones used for training. In a highly interactive scenario like dancing, they showed quantitative and qualitative improvements over the naive adaptation of their base method to interactive scenarios (Mao et al. 2020’s method with concatenation of inputs). Very similarly, Katircioglu et al. (2021) recently presented an analogous adaptation of Mao et al. (2020). Instead of refining each interactants’ keys and values with the others’, Katircioglu et al. (2021) suggested having two branches to exploit the single and multi-person dynamics through self-attention and pairwise attention, respectively, and merge them after the decoding stage. Leveraging the interactant’s motion relative to the person of interest’s coordinates helped to model the interaction. Similarly to the cross-interaction attention, the pairwise attention also outperformed the concatenation-based base method and provided much more interactive predictions. As its main limitation, they raised the point that the fact that each subject has their own dancing style might sometimes cause unsatisfactory results. Curiously, the three state-of-the-art transformer-based methods integrated the DCT to predict a whole motion sequence at once in a non-recurrent manner to avoid freezing motions. This already devises a future trend in low-level behavioral representations forecasting.

In contrast to the previous highly complex approaches, Wang et al. (2021a) recently proposed an unimodal and context-blinded method which beat its multimodal and context-aware competitors in a multi-person motion prediction benchmark (Adeli et al., 2020, 2021). They used the work of Mao et al. (2019) as backbone, which consisted of cascaded Graph Convolutional Networks (GCNs) applied to the DCT of the joints. They proved that using several training tricks such as interpolation of invisible/missing joints, data augmentation, boundary filtering, or curriculum learning, among many others, may be more effective than leveraging more complex networks.

Hands. Even though anticipating the hands’ motion and gestures might be useful for social behavior modeling, we did not find any work within the scope of our survey. Most related work on hands focus on human-object affordance (Lee et al., 2018; Corona et al., 2020a), or hands motion prediction in non-social contexts (Luo and Mai, 2019).

Whole body. Few works have attempted to jointly model the behavior of body and face (Grafsgaard et al., 2018; Joo et al., 2019a). However, they do not fall within the scope of this work as all of them used future information of either another modality (e.g., text, speech) or the interactant. Very recently, a behavior forecasting competition leveraging whole-body landmarks was held within the ChaLearn LAP DYAD@ICCV’21 workshop (Palmero et al., 2022). The common trend observed during the competition coincides with the classic path for body pose forecasting: recurrent encoder-decoder architectures with adversarial losses that ensure realism. Although none of the teams beat the competition baseline, the organizers identified some of the main challenges. The usage of noisy labels, the highly stochastic nature of the hands, or the mostly static nature of the dataset (seated dyadic conversations) are some examples. Motivated by this workshop’s benchmark, Barquero et al. (2022) proposed several state-of-the-art methodologies that outperformed the competition’s baseline. Consistently to the recent findings in body pose forecasting, they also found that Transformer-like architectures provided the best results in whole-body

behavior forecasting. Interestingly, their best results were obtained by only leveraging the prediction of one part of the body at a time (face, pose, and hands). They hypothesized that it could be due to the significant behavioral differences among parts of the body. They also underlined the need of larger datasets to model such high dimensional problems.

2.6.2. HIGH-LEVEL

The ability to understand social signals or behaviors lies in the correct detection of their several distinctive associated social cues (Vinciarelli et al., 2009). Therefore, their *early-detection* or *anticipation* is of utmost importance in many social applications (Ondáš and Pleva, 2019). Both social cues and signals are comprised within our definition of high-level representations of non-verbal behavior, see Table 2. Note that our survey includes works aiming at predicting such behavioral representations at *any* time in the future. This also includes works aiming at the *immediate future* (e.g., decision making, behavior generation). This distinction is noticeable by looking at the future length column in Table 2.

Social cues. Backchannel responses are among the most explored social cues¹ in human-robot interaction scenarios. Such cues can be vocal (e.g., 'Mmmh!', 'Well...!') or visual (e.g., head nodding), and are of utmost importance in order to keep the interacting user engaged (Krauss et al., 1977). Earlier classic approaches built handcrafted sets of rules that triggered generic backchannel responses (Al Moubayed et al., 2009; Poppe et al., 2010). In fact, forecasting backchannel subtypes (generic, agreement, disagreement, surprise, fear, etc.) had traditionally required different levels of semantic processing. Blache et al. (2020) proposed a novel single-route backchannel predictive model that revisited the rule-based classic paradigm and predicted backchannels in real time at a fine-grained level. Their method used prosodic, discourse, semantic, syntax, and gesture features. In contrast with previous approaches that used an observation window as long as the last utterance, they proposed to extract features from bigger semantic units by means of discourse markers. More recently, the collection, annotation and release of bigger datasets favored the appearance of data-driven automated multimodal methods for backchannel prediction. For example, Boudin et al. (2021) used a logistic classifier that was trained on visual cues, prosodic and lexico-syntactic features in order to predict not only the backchannel opportunity but also their subtype associated (generic, positive, or expected). The choice of such a simple classifier was driven by the small dataset available. They showed the superior performance of the multimodal combinations in both tasks.

Other works also focused on the visual dimension of backchannel responses. For example, in a human-robot interaction, Huang et al. (2020) proposed a multimodal Support Vector Machine (SVM) that fused prosodic, verbal (word-based), and visual (head motion and gaze attention) features from only the human interlocutors to generate behavior based on nods and gaze attention switches. They argued that different behaviors needed to be modeled for the three possible situations: speaking, listening, and idling (while no one speaks). Even though their results show a fair prediction capability, the model was only tested for an immediate predicted reaction (next frame). We expect the model to struggle with

1. The definitions of social cues and signals used in this work are borrowed from the Social Signal Processing domain (Vinciarelli et al., 2009; Raman et al., 2021). Accordingly, a *social signal* refers to the relational attitudes displayed by people. Such signals are high-level constructs resulting from the cues perception.

Authors	Method highlights	Multimodal input	Context	→	→	Prediction	Scenario
Social cues							
Sanghvi et al. 2020	GRU+Attention. Social signals are gated with social attention.	✗	Partners' features	15 actions	Imm.	Behavioral actions	Simulated group interactions
Huang et al. 2020	SVM. Multimodal generation of nodding and attention behavior from listeners' behavior only.	Acoustic +Transcripts (words)	Listeners' features	1-10s	Imm	Nodding + attention	Multi-party meeting (HRI)
Blache et al. 2020	Rule-based. It leverages multimodal features at a discourse level.	Acoustic +Transcripts (discourse)	Listener's features	Discourse length	Imm.	Backchannel opportunity +subtypes	Doctor - patient dialogs
Jain and Leekha 2021	LSTM. Semi-supervised annotation method.	Acoustic	Listener's features	3s	3s	Verbal/visual backchannel opportunity	Dyadic conv.
Ishii et al. 2021	MLP. Multimodal multi-task framework for backchannel opportunity prediction.	Acoustic +Transcripts (utterance)	Listener's features	IPU	Imm.	Backchannel opportunity	Dyadic conv.
Boudin et al. 2021	Logistic regression. Multimodal prediction of feedback subtypes.	Acoustic +Transcripts (word)	Listener's features	IPU	Imm.	Backchannel opportunity +subtypes	Dyadic conv.
Murray et al. 2021	LSTM. Useful data augmentation techniques.	Acoustic	Partner's features	2s	Imm.	Nodding	Remote dyadic conv.
Airale et al. 2021	LSTM+GAN. Pooling module and dual-stream discriminator.	✗	Partners' actions	3s	3s	Behavioral actions	Cocktail party
Raman et al. 2021	GRU/MLP. Social processes definition and prediction offset injection.	Speaking status	Partners' features	10f	10f	Speaking status	Triadic conv.
Social signals							
Ishii et al. 2017	SVM. Speaker's and listeners' head motion and synchrony leveraged.	✗	Listeners' head pose	3s	0-4s	Next-utterance timing	Multi-party meeting
Türker et al. 2018	SVM+LSTM. Fusion of classifiers. Many prediction window lengths tested.	Acoustic	Partner's features	3s/2s	0-1s	Nodding/Turn-taking	Dyadic conv.
van Doorn 2018	AdaBoost. Forecasting when an engagement breakdown will occur.	✗	Partners' features	10s	0-30s	Engagement breakdown	Cocktail party
Ishii et al. 2019	SVM+SVR. Three-step model that combines mouth and gaze visual cues.	✗	Listeners' mouth + gaze cues	1.2s	2s	Next speaker + utterance interval	Multi-party meeting
Ishii et al. 2020	MLP. Multimodal multi-task framework for turn-changing anticipation.	Acoustic +Transcripts (utterance)	Listener's features	IPU	+/- 0.6s	Turn-taking	Dyadic conv.
Goswami et al. 2020	Random Forests / ResNet. Visual cues leveraged for the first time for this task.	Acoustic	Speaker's acoustic features	3s	3s	Disengagement	Storytelling to children
Müller et al. 2020	LSTM. Multimodal encoding predicts gaze aversion probability.	Acoustic	Partner's features	6.4s	0.2-5s	Gaze aversion	Remote dyadic conv.
Ben-Youssef et al. 2021	Logistic regression. Combinations of multiple modalities explored.	Acoustic	Partner's features	10s	5s	Engagement breakdown	Dyadic interaction (HRI)

Table 2: Summary of papers forecasting high-level representations of non-verbal behavior. Abbreviations: IPU, inter-pausal unit; Imm., immediate future; conv., conversation; HRI, human-robot interaction; GRU, gated recurrent unit; MLP, multi-layer perceptron; LSTM, long short-term memory; SVM: support vector machine; GAN, generative adversarial network; SVR, support vector regression.

an earlier anticipation as no intentional or behavioral information was encoded. In fact, the authors claimed that the generation of fully autonomous behavior in groupal human-robot interactions is still beyond their capabilities. One of the main reasons behind such pessimistic point of view lays on the numerous particularities of behavior forecasting. For example, backchanneling periods in a conversation are short and infrequent, which leads to a huge imbalanced problem. [Murray et al. \(2021\)](#) proposed a data augmentation method

that tried to mitigate this issue when predicting head nodding. The data augmentation focused on the frequency acoustic features and consisted in warping them over time, masking blocks of utterances, and masking blocks of consecutive frequency channels. A similar technique was explored for the head pose, which was also warped in space and time, and masked over time. The experiments showed important improvements when forecasting head nodding with the combination of these strategies and an LSTM. Another big challenge is the extremely time-consuming annotation of social datasets. In the backchannel scenario, the highly multimodal nature of an interaction requires the annotator to pay attention to the audio, speech, and visual content before making a decision. This process is tedious and prone to errors. Very recently, [Jain and Leekha \(2021\)](#) proposed a semi-supervised method for identification of listener backchannels that was able to detect up to the 90% of the backchannels with only a small subset of labeled data (25%). More importantly, it identified the type of the signals associated around 85% of the times. The authors showed that models trained on such noisy labels were able to keep a 93% and a 96% of the performance with respect to those trained with the cleaned annotations for the tasks of backchannel opportunity prediction and signal classification, respectively. Their general methodology can be adapted for other conversational datasets. Although its validation in other datasets is still pending, it represents an important first step to speed-up the annotation processes and reduce the workforce that they require.

Clearly, the prediction of backchannel responses has attracted a lot of attention. However, they are not the only type of non-verbal behavioral social cues. In a more general framework, few works have tried to predict the future development of an interaction leveraging low-level action labels (e.g., speaking, idling, laughing), motivated by the recent release of annotated audiovisual datasets ([Alameda-Pineda et al., 2016](#); [Cabrera-Quiros et al., 2018](#)). [Sanghvi et al. \(2020\)](#) proposed a model that used the visual features (e.g., location, gaze orientation) from all interactants to predict the target’s future actions (e.g., speak, listen, leave). To do so, a GRU encoded the features of each interactant concatenated with those from the target person. Then, similarly to the attention-based methods reviewed in [Section 2.6.1](#), their method applied attention across all social encodings, which were then fed to a pooling layer so that an arbitrary number of individuals could be handled. Finally, two dense layers converted the output of the pooling layer to the probability distribution defining the next conversational action. They reported that, thanks to the social attention, group annotations were not required. [Airale et al. \(2021\)](#) also defined the non-verbal behavior forecasting as a discrete multi-sequence generation problem. Their methodology was radically different though: a GAN conditioned on the observed interaction, which was encoded with an LSTM. In the generative stage, a socially aware hidden state was computed at each timestep. The strategy consisted in using a pooling module to update the hidden states of each person’s LSTM decoder and convert them to new socially aware LSTM hidden states for the next decoding step. As a result, the decoded actions were obtained in a coherent way with respect to the actions generated for all surrounding subjects. Additionally, they presented a two-streams novel adversarial discriminator. The first branch corresponded to the classic one, so it favored the realism of individual action sequences disregarding any contextual information. The second one combined the predicted actions sequences with a pooling module of all individuals in the scene to ensure that the generated interactions as a

whole were realistic. As a result, the consistency across the predicted action sequences for all participants of the interaction was preserved.

Social signals. The ability to recognize human social signals and behaviors like turn-taking, disengagement, or agreement, and act accordingly are the keys to develop socially intelligent agents (Vinciarelli et al., 2009). There are many other human capabilities which are carried out unconsciously, like anticipation. For example, a person starts building their response to the speaker’s speech before their turn is over, thanks to anticipating the turn taking event (Ondáš and Pleva, 2019). Consequently, research has focused on trying to find the most important social cues when it comes to anticipating the appearance of social signals of interest. For example, Ishii et al. (2017) found that the amount of head movement of the current speaker, next speaker, and listeners had fairly good prediction capabilities with regards to the next utterance timing. They proposed a light SVM method that could be deployed to any agent equipped with a camera or a depth sensor. Similarly, in multi-party conversations, Ishii et al. (2019) discovered that the speaker’s and listeners’ mouth-opening transition patterns could be used to predict the next speaker and the time interval between the current utterance ends and the next utterance begins. To prove it, they developed a three-step system. First, an SVM model predicted whether a turn-changing would be produced next. If the answer was yes, then another SVM predicted the next speaker. Finally, independent SVR models for turn-changing and turn-keeping events predicted the utterance interval. The good results of the predictive model suggested the importance of visual cues in forecasting the conversational flow. Its exploitation could potentially help conversational agents to raise the participants’ engagement before the start of the next utterance.

Actually, turn-taking modeling has always attracted a lot of attention. An appropriate turn management is very much needed for a smooth and fluent conversation, which is determining for a pleasant human-robot interaction. In this line, Türker et al. (2018) made one of the first attempts to exploit multiple modalities (acoustic features and visual cues) to predict turn-taking. Among their contributions, they presented an approach that summarized each acoustic feature into a set of statistical measures across the temporal axis (e.g., mean, deviation, skewness). As a result, thanks to the removal of the temporal axis, a simple SVM could be used as classifier. Unfortunately, their tests with turn-taking and head-nodding behavior forecasting showed a superior performance of the recurrent LSTM alternative, thus proving the relevance of the temporal dependencies for such tasks. They also showed that, while the unimodal (acoustic features) forecasting results were close to random, the multimodal performance was promising. Many subsequent works have presented successful multimodal methodologies for forecasting other social signals. For example, Ishii et al. (2020) analyzed the relationship between turn-holding and grabbing willingness and the actual turn-change. Although they found discrepancies between willingness and actual turn-changing behavior, building a multimodal multi-task model to simultaneously predict both turn-willingness and turn-changing behaviors improved the results for both tasks. In a posterior work, Ishii et al. (2021) expanded their multimodal multi-task framework with the backchannel opportunity prediction task. They showed that, while backchannel prediction benefited from the multi-task learning, no improvement came from adding the turn-changing prediction. Among their conclusions, they stated that, in both cases, simultaneously employing features from the speaker and the listener helped to improve the predictions. One of their

main limitations though was the limited exploration of fusion strategies for the modalities used, as their choice was simply concatenating the three feature vectors of audio, text, and video, before a dense layer.

Disengagement is another very important social signal to be considered when designing interactive agents. If identified with enough anticipation, the speaker can prevent it from happening with backchannel responses or by making an engaging hands gesture, for example. [Van Doorn \(2018\)](#) made an attempt to anticipate whether a user would leave the conversation and when. Unfortunately, in the first task, they only slightly outperformed the random baseline with one of the many models trained (AdaBoost), and were not successful at the second. Similarly in a human-robot interaction scenario, [Ben-Youssef et al. \(2021\)](#) aimed at anticipating a premature ending of an interaction. As part of their experiments, a logistic regression classifier was trained with all the possible multimodal combinations. They found that the best results were achieved with the combination of the distance to the robot, the gaze and the head motion, as well as facial expressions and acoustic features. Surprisingly, the choice of the classifier had little effect over the final results (logistic regression, random forest, multi-layer perceptron, or linear-discriminant analysis). The small influence of the classifier choice has been consistent along all works reviewed in this section. We hypothesize that this is due to the little need of further processing for such simple and already high-level representations. As a matter of fact, [Goswami et al. \(2020\)](#) found few performance differences between a random forest and a ResNet when predicting disengagement in the context of storytelling with children, who are prone to disengage very easily. In this work, they also predicted whether a low or a high degree of backchanneling was needed to keep the listener engaged. They assessed the prediction capabilities of many visual cues never used before for this task such as pupil dilation, blink rate, head movements, or facial AUs. Interestingly, they found that the gaze features and the speech pitch were among the most important features to be considered for disengagement prediction. These findings are consistent with those from the work of [Müller et al. \(2020\)](#), who found eye contact and speaker turns to be the most informative modalities when it comes to anticipating averted gaze during dyadic interactions. In their work, the authors also tested other less powerful modalities including face- and gaze-related attributes, expressions and speaker information.

As observed in the course of the survey, there is an important heterogeneity with respect to methodologies used for social signal/cues forecasting. [Joo et al. \(2019a\)](#) tried to establish a generic definition which homogenized all methodologies: a *social signal prediction* (SSP) model. The SSP model defined a framework to model the dynamics of social signals exchanged among interacting individuals in a data-driven way. It consisted in using the target’s past behavior and their interactants’ to predict its future behavior. However, their definition implied that a separate function was learnt for every person. [Raman et al. \(2021\)](#) solved this issue in their formulation of *social processes* (SP), in which a simultaneous prediction of the behavior of all the individuals was considered.

3. Datasets

A recurrent problem observed in our survey and one of the main challenges of non-verbal social behavior forecasting is the lack of large annotated datasets. In order to provide the reader with an overview of the currently available datasets, we briefly go through them in

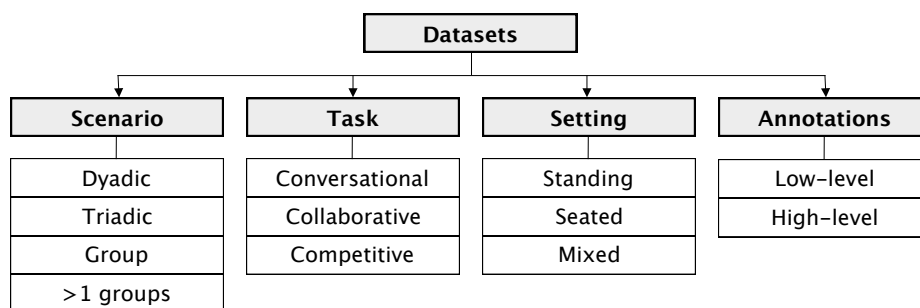


Figure 3: Classification of audiovisual datasets featuring non-acted social interactions.




Figure 4: Samples representative of the types of dataset scenarios included in this survey.

this section. Note that we restrict the survey to *publicly available* datasets that feature *audiovisual non-acted social interactions*. The taxonomy that we present, see [Figure 3](#), groups them into scenarios (dyadic, triadic, group, and >1 groups), tasks (conversational, collaborative, and competitive), and settings (standing, or seating) that elicit different behavioral patterns. [Figure 4](#) shows illustrative examples of the scenarios considered in this part of the survey.

We summarize and compare all datasets reviewed in [Table 3](#). They appear classified into first-person (egocentric), third-person (mid-distance camera), and computer-mediated

	Dataset	Scenario	Task	Setting	Content	#Subjects	Size	Behavioral Annotations	
								Low-level	High-level
Third-person view									
McCowan et al. 2005	AMI	Group	Conv.	Seated	A,T	?	100h	✗	Turn Taking, Gestures, Emotions, Game Decisions
Douglas-Cowie et al. 2007	HUMaine	Multiple	Multiple	Seated	A,P,T	309	>26h	Face Expression	Face Gestures, Emotions
Van Son et al. 2008	IFADV	Dyadic	Conv.	Seated	A	34	5h	Gaze	Emotions, Turn talking, Feedback Responses
Bertrand et al. 2008	CID	Dyadic	Conv.	Seated	A	16	8h	✗	Phonetics, Prosody, Morphology, Syntax, Discourse, Face Gestures
Edlund et al. 2010	Spontal	Dyadic	Conv.	Seated	A	?	60h	Motion Capture	✗
Hung and Chittaranjan 2010	IDIAP Wolf	Group	Comp.	Seated	A	36	7h	✗	Speaking Segments, Deceptive/Non-Deceptive Roles, Game Decisions
Lücking et al. 2012	SaGA	Dyadic	Conv.	Seated	A,T	50	4.7h	✗	Gestures
Soomro et al. 2012	UCF101	Multiple	Multiple	Mixed	A	?	27h	✗	Action labels
Sanchez-Cortes et al. 2012	ELEA	Triadic, Group	Collab.	Seated	A,Q	102	10h	✗	Power, Dominance, Leadership, Perceived Leadership, Competence, Likeness
Reh et al. 2013	MMDB	Dyadic	Conv.+ Collab.	Seated	A,P	121	~10h	Face Expressions, Gaze	Vocalizations, Verbalizations, Vocal Affect, Gestures
Vella and Paggio 2013	MAMCO	Dyadic	Conv.+ Collab.	Standing	A,P,T	12	~1h	✗	Turn Overlap
Bilakhia et al. 2015	MAHNOB	Dyadic	Conv.	Seated	A	60	11.6h	Face Expressions, Head, Body, and Hands Motion, Postural Shifts	✗
Vandeventer et al. 2015	4D CCDB	Dyadic	Conv.	Seated	A,T	4	~0.5h	Face Expressions, Head Motion, Gaze	Back/Front Channelling, (dis)Agreement, Happiness, Surprise, Thinking, Confusion, Head Nodding/Shaking/Tilting
Salter et al. 2015	The Tower Game	Dyadic	Comp.	Standing	A	39	9.5h	Face Landmark, Gaze, Person Tracking	✗
Naim et al. 2015	MIT Interview	Dyadic	Conv.	Seated	A,T	69	10.5h	Face Expressions	Friendliness, Presence, Engagement, Excitement, Focused, Calm, Authentic
Shukla et al. 2016	MuDERI	Dyadic	Conv.	Seated	A,B,P	12	~7h	✗	Valence, Arousal
Alameda-Pineda et al. 2016	SALSA	>1 groups	Conv.	Standing	A,P	18	1h	Position, Head, Body Orientation	F-formations
Edwards et al. 2016	CONVERSE	Dyadic	Conv.+ Collab.	Standing	A	16	8h	Body Landmarks, Gaze, Face Expressions	✗
Beyan et al. 2016	Leadership Corpus	Group	Collab.	Seated	A,P,Q	64	~7h	✗	Leadership
Chou et al. 2017	NNIME	Dyadic	Conv.	Seated	A,B,T	44	11h	✗	Emotions

Table 3: Datasets that feature audiovisual non-acted social interactions and are publicly available. They are presented grouped by recording setup (third-person, egocentric, and computer-mediated). Abbreviations: Conv. conversational; Collab., collaborative, Comp.; competitive, A, audiovisual; P, psychological; B, biosignals; T, transcriptions; Q, questionnaires; IMU, Inertial measurement unit; ?, value not found; *, robot interaction.

NON-VERBAL SOCIAL BEHAVIOR FORECASTING

	Dataset	Scenario	Task	Setting	Content	#Subjects	Size	Behavioral Annotations	
								Low-level	High-level
Third-person view									
Georgakis et al. 2017	CONFER	Multiple	Comp.	Seated	A	54	2.4h	Face Landmarks, Person Tracking	Conflict Intensity
Paggio and Navarretta 2017	NOMCO	Dyadic	Conv.	Standing	A,T	12	~1h	Head Motion, Face Expressions, Body Landmarks	Emotions, Gestures
Bozkurt et al. 2017	JESTKOD	Dyadic	Conv.	Standing	A,P	10	4.3h	Body Landmarks, Body Motion	Activation, Valence and Dominance
Andriluka et al. 2018	PoseTrack	Multiple	Multiple	Mixed	A	>300	~45h	Body Landmarks	✗
von Marcard et al. 2018	3DPW	Multiple	Multiple	Mixed	A	?	0.5h	3D Body Landmarks	✗
Mehta et al. 2018	MuPoTS-3D	Dyadic, Triadic	Conv.	Seated	A	8	0.05h	3D Body Landmarks	✗
Lemaignan et al. 2018	PInSoRo*	Dyadic	Multiple	Mixed	A	120	45h	Face, Body Landmarks	Task Engagement, Social Engagement, Social Attitude
Cabrera-Quiros et al. 2018	MatchNMingle	>1 groups	Conv.	Seated	A,P	92	2h	Acceleration, Proximity	Social Actions, Cues and Signals, F-Formations
Celiktutan et al. 2019	MHHRI*	Dyadic, Triadic	Conv.	Seated	A,B	18	4.2h	Wrist Acceleration	Self-Reported Engagement
Joo et al. 2019b	CMU Panoptic	Triadic, Group	Multiple	Mixed	A,P,T	?	5.5h	3D Body, Face and Hands Landmarks	Speaking Status, Social Formation
Carreira et al. 2019	Kinetics-700	Multiple	Multiple	Mixed	A	?	1805h	✗	Action Labels
Lee et al. 2019	Talking With Hands 16.2M	Dyadic	Conv.	Standing	A	50	50h	Body, Hands Landmarks	✗
Zhao et al. 2019	HACS	Multiple	Multiple	Mixed	A	?	861h	✗	Action Labels
Monfort et al. 2020	Moments in time	Multiple	Multiple	Mixed	A	?	833h	✗	Action Labels
Chen et al. 2020	DAMI-P2C	Dyadic	Conv.	Seated	A,P,T	68	21.6h	✗	Parent Perception, Engagement, Affect
Maman et al. 2020	GAME-ON	Triadic	Multiple	Standing	A,Q	51	11.5h	Motion Capture	Cohesion, Leadership, Warmth, Competence, Competitvity, Emotions, Motivation
Khan et al. 2020	Vyaktiv	Dyadic	Conv.	Seated	A,P,T	38	~6.7h	✗	Lexical Annotations
Schiphorst et al. 2020	Video2Report	Dyadic	Conv. (Medical Visit)	Mixed	A	4	~7.1h	Body Landmarks	Action Labels (Medical)
Park et al. 2020	K-EmoCon	Dyadic	Conv.	Seated	A,B	32	2.8h	Accelerometer	✗
Yang et al. 2021	CongreG8	Triadic, Group	Comp.	Standing	A,P,Q	38	~28h	Motion Capture	✗
Martín-Martín et al. 2021	JRDB-Act	Multiple	Multiple	Mixed	A	?	1h	Body Point Cloud, Person Detection and Tracking	Atomic Action Labels, Social Formations
Doyran et al. 2021	MUMBAI	Dyadic	Collab +Comp.	Seated	A,P,Q	58	46h	Body, Face Landmarks	Emotions
Palmero et al. 2022	UDIVA v0.5	Dyadic	Conv.+ Collab. +Comp.	Seated	A,Q,T	134	80h	Face, Body, Hands Landmarks, Gaze	✗

Table 3: (Continuation) Datasets that feature audiovisual non-acted social interactions and are publicly available. They are presented grouped by recording setup (third-person, egocentric, and computer-mediated). Abbreviations: Conv. conversational; Collab., collaborative, Comp.; competitive, A, audiovisual; P, psychological; B, biosignals; T, transcriptions; Q, questionnaires; IMU, Inertial measurement unit; ?, value not found; *, robot interaction.

	Dataset	Scenario	Task	Setting	Content	#Subjects	Size	Behavioral Annotations	
								Low-level	High-level
First-person view (Egocentric)									
Bambach et al. 2015	EgoHands	Dyadic	Collab +Comp.	Seated	A	4	1.2h	Hands Segmentation	✗
Yonetani et al. 2016	PEV	Dyadic	Conv.	Seated	A	6	~0.5h	✗	Actions and Reactions Labels
Silva et al. 2018	DoMSEV	Multiple	Multiple	Mixed	A	?	80h	IMU, GPS	Actions Label
Abebe et al. 2018	FPV-O	Multiple	Conv. (Office)	Mixed	A	12	3h	✗	Actions Label
Grauman et al. 2021	Ego4D	Multiple	Multiple (Daily Life)	Mixed	A	923	3670h	Gaze	✗
Computer-Mediated (Online Interactions)									
McKeown et al. 2010	SEMAINE	Dyadic	Conv.	Seated	A,T	20	~6.5h	✗	Emotions, Epistemic States, Interaction Actions, Engagement
Ringeval et al. 2013	RECOLA	Dyadic	Collab.	Seated	A,B,P	46	3.8h	✗	Valence, Arousal, Agreement, Dominance, Performance, Rapport, Engagement, Utterance
Cafaro et al. 2017	NoXi	Dyadic	Conv.	Standing	A,T	87	25h	Body and Face Landmarks, Smiling, Head and Hands Gestures	Arousal, Engagement, Turn Talking
Feng et al. 2017	Learn2Smile	Dyadic	Conv.	Seated	A	500	~30h	Face Landmarks	✗
Kossaifi et al. 2019	SEWA DB	Dyadic	Conv.	Seated	A,T	398	44h	Face Landmarks and Action Units, Head and Hands Gestures	Valence, Arousal, Liking/Disliking, Agreement, Mimicry, Backchannel, Laughs

Table 3: (Continuation) Datasets that feature audiovisual non-acted social interactions and are publicly available. They are presented grouped by recording setup (third-person, egocentric, and computer-mediated). Abbreviations: Conv. conversational; Collab., collaborative, Comp.; competitive, A, audiovisual; P, psychological; B, biosignals; T, transcriptions; Q, questionnaires; IMU, Inertial measurement unit; ?, value not found; *, robot interaction.

(e.g., video-conferences) recording setups due to their significant differences regarding their possible applications. Datasets recorded from very distant third-person views, or egocentric views where the camera is carried by a non-interacting agent were discarded due to their poor social behavior content. Usually, third-view datasets consist of structured interactions where participants need to follow basic directives which favor spontaneous and fluent interactions. Despite the fact that conversations are the most common interaction structure, there are datasets which aim at fostering specific social signals like leadership, competitiveness, empathy, or affect, and therefore engage the participants in competitive/cooperative scenarios ([Hung and Chittaranjan, 2010](#); [Sanchez-Cortes et al., 2012](#); [Rehg et al., 2013](#); [Ringeval et al., 2013](#); [Vella and Paggio, 2013](#); [Bambach et al., 2015](#); [Salter et al., 2015](#); [Edwards et al., 2016](#); [Beyan et al., 2016](#); [Georgakis et al., 2017](#); [Yang et al., 2021](#); [Doyran et al., 2021](#); [Palmero et al., 2022](#)). Other datasets, instead, record in-the-wild interactions during the so-called cocktail parties ([Alameda-Pineda et al., 2016](#); [Cabrera-Quiros et al., 2018](#)) and represent very interesting benchmarks to study group dynamics. Thanks to the camera portability during the collection, egocentric datasets can record social behavior in

less constrained environments. Very recently, [Grauman et al. \(2021\)](#) released more than 3000 hours of in-the-wild egocentric recordings of human actions, which also include social interactions. Finally, the computer-mediated recording setup elicits a very particular behavior due to the idiosyncrasies of the communication channel ([McKeown et al., 2010](#); [Ringeval et al., 2013](#); [Cafaro et al., 2017](#); [Feng et al., 2017](#); [Kossaifi et al., 2019](#)). For example, the latency or the limited field of view might affect the way social cues and signals are transmitted and observed.

With regards to the setting, participants might interact while standing or while seated. Some datasets may include videos with both configurations (e.g., several independent interactive groups). The most frequent scenario consists in dyadic interactions due to their special interest for human-robot interaction and human behavior understanding, and their lower behavioral complexity when compared to bigger social groups. Triadic ([Sanchez-Cortes et al., 2012](#); [Mehta et al., 2018](#); [Celiktutan et al., 2019](#); [Joo et al., 2019b](#); [Yang et al., 2021](#)) and bigger social gatherings ([McCowan et al., 2005](#); [Hung and Chittaranjan, 2010](#); [Beyan et al., 2016](#); [Joo et al., 2019b](#); [Yang et al., 2021](#)), which are referred to as *groups*, are less commonly showcased scenarios. Datasets featuring several simultaneous groups of interactions, as long as they showed *focused interactions*, are also included ([Alameda-Pineda et al., 2016](#); [Cabrera-Quiros et al., 2018](#)).

Regarding the content released, although originally available, some of them did not release the audio of the videos showcased due to privacy concerns. This is especially frequent for egocentric videos, as the unconstrained recording of the interactions obstructs the collection of consent forms. Other common content typologies consist of psychological data (e.g., personality questionnaires), biosignals monitoring (e.g., heart rate, electrocardiogram, electroencephalograms) and transcriptions. The latter is considerably less frequent due to its tedious manual annotation process ([McCowan et al., 2005](#); [Douglas-Cowie et al., 2007](#); [McKeown et al., 2010](#); [Lücking et al., 2012](#); [Vella and Paggio, 2013](#); [Vandeventer et al., 2015](#); [Naim et al., 2015](#); [Chou et al., 2017](#); [Paggio and Navarretta, 2017](#); [Cafaro et al., 2017](#); [Joo et al., 2019b](#); [Kossaifi et al., 2019](#); [Chen et al., 2020](#); [Khan et al., 2020](#); [Palmero et al., 2022](#)). The most frequent low-level annotations that the datasets provide are the participants' body poses and facial expressions ([Douglas-Cowie et al., 2007](#); [Rehg et al., 2013](#); [Bilakhia et al., 2015](#); [Vandeventer et al., 2015](#); [Naim et al., 2015](#); [Edwards et al., 2016](#); [Cafaro et al., 2017](#); [Feng et al., 2017](#); [Georgakis et al., 2017](#); [Paggio and Navarretta, 2017](#); [Bozkurt et al., 2017](#); [Andriluka et al., 2018](#); [von Marcard et al., 2018](#); [Mehta et al., 2018](#); [Lemaignan et al., 2018](#); [Joo et al., 2019b](#); [Kossaifi et al., 2019](#); [Schiphorst et al., 2020](#); [Doyran et al., 2021](#); [Palmero et al., 2022](#)). Given their annotation complexity, they are usually automatically retrieved with tools like OpenPose ([Cao et al., 2019](#)), and manually fixed or discarded. Others use more complex retrieval systems like motion capture, or mocap ([Edlund et al., 2010](#); [Maman et al., 2020](#); [Yang et al., 2021](#)). However, the characteristics of the mocap recording setup and the special suit that participants wear could unintentionally bias the elicited behaviors during the interaction. Finally, the annotation of high-level social signals is often led by the needs of the study for which the dataset was collected. Indeed, some of the datasets have been complementary annotated and added in posterior studies. As a result, most common high-level labels consist of elicited emotions ([McCowan et al., 2005](#); [Douglas-Cowie et al., 2007](#); [van Son et al., 2008](#); [McKeown et al., 2010](#); [Naim et al., 2015](#); [Vandeventer et al., 2015](#); [Chou et al., 2017](#); [Paggio and Navarretta, 2017](#); [Maman et al., 2020](#); [Doyran et al.,](#)

2021), action labels (Soomro et al., 2012; Yonetani et al., 2016; Silva et al., 2018; Abebe et al., 2018; Carreira et al., 2019; Zhao et al., 2019; Schiphorst et al., 2020; Monfort et al., 2020; Martín-Martín et al., 2021), and social cues/signals (Hung and Chittaranjan, 2010; Sanchez-Cortes et al., 2012; Ringeval et al., 2013; Vandeventer et al., 2015; Shukla et al., 2016; Bozkurt et al., 2017; Cafaro et al., 2017; Feng et al., 2017; Lemaignan et al., 2018; Cabrera-Quiros et al., 2018; Celiktutan et al., 2019; Chen et al., 2020; Maman et al., 2020).

4. Metrics

The metrics used in behavior forecasting can be clustered in three well-defined branches by the property of the behavior that they assess, see [Figure 5](#). This classification will be used in this section to present and group the metrics.

On one side, works that predict high-level behavioral representations such as social cues or signals often need to compare single or multiple discrete predictions to a single observed ground truth (*one-to-one*). For univariate classification, the Area Under the Curve (AUC) of the Receiver Operating Curve (ROC) is a common choice, as it very well describes the overall performance of the model. The classic accuracy, precision, recall, and F1-score are valid alternatives for both single and multi-class classification. On the other side, low-level behavioral representations are often continuous and therefore conceived as regression tasks. For those, the L1 and L2 distances have traditionally been the golden standard. Variants of L2-based metrics have been used depending on the field of application. For example, in human motion forecasting with joints, the Mean-Per-Joint-Position Error (MPJPE), the Percentage of Correct Keypoints (PCK), or the cosine similarity are very frequent options. In this context, [Adeli et al. \(2021\)](#) raised the common problem of missing data (e.g., occluded or out-of-view joints). To address this, they proposed metrics that evaluated the models performance under several visibility scenarios: the Visibility-Ignored Metric (VIM), the Visibility-Aware Metric (VAM) and the Visibility Score Metric (VSM). Differently, in behavior forecasting with raw image or audio, quality metrics like the Mean-Squared Error (MSE), the Structural Similarity Index Measure (SSIM) ([Wang et al., 2004](#)), or the Peak Signal to Noise Ratio (PSNR) are better suited. Indeed, this task shares many similarities with video prediction, so any metric from that field can be leveraged for ours ([Oprea et al., 2020](#)).

A common argument used against accuracy-based metrics is supported on the idea that there are multiple valid and equally plausible futures. For example, using a distance-based metric such as mean-squared error in hands gestures forecasting could end up penalizing models which forecast a gesture over another which could be also suitable to that situation. Similarly, a method could forecast the correct high-frequency gesture but under a small delay and yield a very low accuracy score. To account for such future stochasticity, works embracing this paradigm try to predict all possible future sequences ([Section 2.1](#)). In order to make sure that a representative spectrum of all future possibilities is predicted, stochastic approaches need to measure both the *accuracy* and *diversity* of their predictions ([Yuan and Kitani, 2020](#)). Most of these works predict low-level representations that are continuous in the future (e.g., trajectories, poses). As a result, stochastic metrics herein presented may be biased towards such representations. However, most of them are directly applicable to high-level representations forecasting.




Figure 5: Classification of metrics commonly used in non-verbal social behavior forecasting.

The accuracy of methods under the stochastic assumption can be quantified under two different paradigms. The first consists in assessing that at least one of the predicted futures is accurate and matches the ground truth (*many-to-one*). To do so, first, the predicted sample most similar to the ground truth is selected. Then, its accuracy can be simply computed with any of the deterministic metrics previously presented (e.g., Average Displacement Error, ADE, Yuan and Kitani 2020). The second, instead of assuming that at least one predicted sample is certain, consists in generating a hypothetical set of multiple ground truths (*many-to-many*). They usually do so by grouping similar past sequences. The future sequences from those grouped observations are considered their *multimodal* ground truth. In this scenario, multimodal adaptations of the deterministic metrics are usually leveraged. Examples include the Multimodal ADE (MMADE), and the Multimodal FDE (MMFDE) (Mao et al., 2021a). Nonetheless, any one-to-one metric can be computed across all possible futures and then averaged to get a multimodal score. Regarding the quantification of the diversity across the multiple predicted futures, the analysis is usually restricted to either the *predictions space* or an *intermediate latent space*. For direct comparison in the predictions space, some works use the Average Pairwise Distance (APD), which is calculated as the average L2 distance between all pairs of predictions (Yuan and Kitani, 2020; Mao et al., 2021a; Aliakbarian et al., 2021). The higher this metric is, the higher the variability among predictions. Similarly, the Average Self Distance (ASD) and the Final Self Distance (FSD) were proposed (Yuan and Kitani, 2019). They both measure, for each future sample from the predictions set Y , its minimum distance to the closest future from Y (in terms of L2 distance, for example). All timesteps are averaged for the ASD, and only the last timestep is considered for FSD. These metrics penalize methods sampling repeated futures. Regarding the diversity-based metrics computed in the latent space, the usual choices are the popular Frechet Inception Distance (FID) and the Inception Score (IS), which are distribution-based metrics that measure the generation fidelity (Yang et al., 2018; Aliakbarian et al., 2020; Cai et al., 2021). Additionally, Cai et al. (2021) presented a *Diversity* score that estimates the feature-based standard deviation of the multiple generated outputs from the same past observation.

Finally, the realism, naturalness, smoothness and plausibility of behavior forecasting is often visually assessed by human raters. To do so, some works prepare questionnaires asking questions like “*Is the purpose of the interaction successfully achieved?*”, “*Is the synthesized agent behaving naturally?*”, “*Does the synthesized agent look like human rather than a robot?*” (Shu et al., 2016). The answers are usually scales of discrete values (e.g., between

0 and 5). During the evaluation process, those questions are presented while showing a behavior sampled either from the ground truth or from the predictive model (Kucherenko et al., 2021). This prevents the introduction of biases in the humans’ ratings. The posterior usage of statistical hypothesis tests helps to conclude whether the predicted behavior is humanlike (Feng et al., 2017; Chu et al., 2018; Woo et al., 2021). Unfortunately, the amount of unique questions used in the literature is very large, with very few of them being repeated across studies (Fitriani et al., 2019). There have been few attempts to create a unified measurement instrument that helps reduce such heterogeneity (Fitriani et al., 2020, 2021). Alternatively, other works use pairwise evaluations. In those, the human rater selects the most human-like behavior among a pair of samples. A recent study observed the superior inter-rater reliability of pairwise evaluations, which favored it over questionnaires for small-scale studies (Wolfert et al., 2021a). Overall, most subjective evaluations lack a systematic approach that ensures high methodological quality. As a result, the extraction of systematic conclusions is often difficult (Wolfert et al., 2021b). Additionally, one must be aware though of the possible biases induced by this type of assessments. First, the sampling of the subjects that participate in the qualitative evaluation may include biases towards certain subgroups. For example, the participants selection of a qualitative analysis performed by a graduate student may contain inherent biases towards people from academia. On the other side, the bias could be originated in the choice of the samples to evaluate, the way they are displayed, or even the order in which they are presented. Although most of them are not completely avoidable, they must be taken into consideration and minimized if possible.

5. Discussion

The forecasting of low-level representations like landmarks or facial action units has been recently tackled with deep learning methods such as recurrent neural networks, graph neural networks, and transformers. The usage of such deep and data-hungry models has been encouraged by the recent availability of large-scale multi-view datasets, often annotated in a semi-automatic way (Joo et al., 2019b; Palmero et al., 2022). The increasing accuracy of monocular and multi-view automated methods for face, pose, and hands estimation has contributed in reducing the annotation effort. Still, the largest available datasets that provide thousands of hours of audiovisual material and feature the widest spectrum of behaviors do not provide such annotations (Carreira et al., 2019; Zhao et al., 2019; Monfort et al., 2020; Grauman et al., 2021). In contrast, the automated methods for high-level representations recognition such as feedback responses or atomic action labels are not accurate enough to significantly help in their annotation procedures. Consequently, such annotations are scarce, and are only available for small datasets, as shown in our survey. Accordingly, recent works have opted for classic methods such as SVM, AdaBoost, and simple recurrent neural networks, which have traditionally worked fairly well with small datasets. We expect future work on high-level behavior forecasting to also explore semi- and weakly-supervised approaches (Jain and Leekha, 2021).

Latest works that focus on forecasting low-level representations have proposed methods that successfully exploit interpersonal dynamics in very specific scenarios (e.g., dancing) by using cross-interaction attention (Katircioglu et al., 2021; Guo et al., 2021). In other tasks where these dynamics may not be so strong and frequent (e.g., conversational), or simply

not sufficiently captured by the visual cue chosen (e.g., landmarks), the adequate inclusion of such features still has further room for improvement (Barquero et al., 2022). In fact, the influence of the scenario and the input representation on performance is also observed in works that focus on forecasting high-level representations. For instance, using head and body pose as input and output, each represented by a selected 3D landmark coordinate and a normal, Raman et al. (2021) observed that the addition of features from the other interlocutors harmed performance for the three predicted categories (including speaking status) when applied to an in-the-wild mingling scenario. In contrast, in a structured triadic scenario, head and body location were better predicted when adding such information (MSE of up to 15.84 cm vs. 18.20 cm), whereas orientation was better predicted without. When using input features other than landmarks, most works tend to benefit from the addition of the partner’s features, even in conversational scenarios. For instance, Ishii et al. (2021) showed consistently superior performance when using speaker and listener features compared to using just one of them for backchannel (F1 score of 85.2% vs. up to 74.2%) and turn-changing (F1 score of 61.7% vs. up to 59.2%) prediction in a seated dyadic scenario.

With respect to the methodological trends of low-level representation forecasting, we foresee a bright future for methods that use representations in the frequency space. Very recent works have reported promising results with such architectures, especially by helping to alleviate the very limiting freezing motion commonly observed in deterministic approaches. Related to the future perspective, we also expect future works to explore the stochastic point of view. Although many works from the single human behavior forecasting field have already found benefits from assuming the future stochasticity and thus predicting several futures, their translation to social settings remains unexplored. We think that the implementation of socially aware adversarial losses, like the dual-stream one presented by Airale et al. (2021) for behavioral action forecasting, could help to build systems capable of generating diverse, plausible, and socially coherent behavior. There are other research lines that have provided many benefits in other fields that also remain uninvestigated. First, the exploration of multimodal approaches has been timidly addressed in the past only to generate immediate next motions (Hua et al., 2019; Ahuja et al., 2019), or in very preliminary and naive ways (Barquero et al., 2022). Future works should also test self-supervised learning techniques, which have shown their power in conceptually close applications like video prediction (Oprea et al., 2020), and could be similarly applied to this field. Furthermore, models that update the learnt behavioral model according to each person’s individual behavioral patterns via meta-learning are also very promising (Moon and Seo, 2021).

A popular question when reading works that tackle forecasting problems like ours is whether the results could be transferred to real applications. We all like to see state-of-the-art methods on top of benchmarks and directly choose them for our target application. Instead, we should stop right before making a choice and ask ourselves: *do those numeric accuracy values suit real life scenarios?* Obviously, the answer is not straightforward, and greatly depends to a large degree on the amount of error assumable in each application. For example, terminating an interaction during a life-threatening human-robot assistance (e.g., assisted surgery, rescues) does not have room for errors, while doing it during a shopping assistance does. In order to assess this adequacy of the model performance to the real target task, works on social cues/signals forecasting (e.g., backchannel opportunity prediction, interaction breakdown) often perform objective and subjective studies by means of robotic

or virtual agents (Huang et al., 2020; Murray et al., 2021). For example, these tests have been used to prove the great capacities of models that generate backchannel responses when it comes to successfully keeping the user engaged during human-robot interactions (Murray et al., 2021). For low-level representations though, there is still a lack of extensive studies that assess the transferability of the results to the target scenario. This is mainly due to the extra constraints posed by these low-level representations. In fact, some works highlight the current possibilities and limitations of behavior forecasting with such representations. For example, future behavior in competitive interactions like fencing strongly depends on the player’s decisions in answer to the competitor’s. Within this context, Honda et al. (2020) showed inferior performance for the rapid and highly stochastic motion of the dominant arm (PCKs of 71.8% and 66.4% for dominant hand and elbow, respectively) than for the other parts of the body (average PCK of 77.1%). This has been consistent in the literature for other less interactive scenarios like face-to-face conversations. Barquero et al. (2022) showed superior accuracy for behavior forecasting for face and upper body torso (errors of 12.70 and 5.75 pixels in average for a prediction of 2 seconds) than for hands (25.15 pixels). This represents an important bottleneck for hands forecasting, where research is almost nonexistent despite of their importance in human communication. The authors also showed superior performance (error of 5.34 pixels) than a naive but strong baseline (6.00 pixels) for the short term (<400ms). This opens new possibilities for providing human-robot interactive agents with fairly accurate anticipation capabilities. For example, the proper activation of the actuators of a robot may benefit from any extra milliseconds of anticipation. In general though, we think that landmarks-based behavior forecasting is still immature, and will strongly benefit from further research efforts. Another concerning issue related to this topic lies on the typology of the data leveraged for forecasting. Only models that make their predictions solely leveraging automatically retrieved data can be successfully applied to real life scenarios. Actually, and similarly to the low-level scenario, we expect the forecasting of high-level behavioral representations to greatly benefit from the development of new accurate methods to automatically retrieve social cues/signals from raw image/audio data.

Finally, we want to raise awareness of what we consider one of the main bottlenecks of behavior forecasting: the evaluation metrics. An evaluation metric must always illustrate how well a method does for the target task. While this sentence may seem trivial when thinking of classic classification or regression tasks, it is an important source of controversy in the behavior forecasting field. For instance, the distance between the generated and the ground truth futures does not describe the coherence of the pose in all future steps, neither the realism of the movements. In fact, it does not even guarantee that a method with low error performs poorly, as the predictions may simply not match the ground truth, which is a sample of the multiple and equally plausible set of futures. Although one would conclude that a proper evaluation always contains a qualitative analysis, multiple behavioral dimensions may escape from human raters and therefore bias it. For instance, it is not trivial to build a qualitative analysis that also assesses the coherence of the predicted behavior with respect to the behavioral patterns specific to the subject, the context, or even the events from the mid- to long-term past. We hope that the recent appearance of behavior forecasting benchmarks and specific datasets will encourage the community to find better-suited metrics and evaluation protocols that will boost the research progress in this field.

6. Ethics

We have discussed many applications for good where non-verbal social behavior forecasting might be valuable. Personalized pedagogical agents (Davis, 2018) that maximize learner’s attention and learning, empathetic assistive robots for hospital patients or dependent people (Andrist et al., 2015; Esterwood and Robert, 2021), and collaborative robots for industrial processes or even surgeries (Sexton et al., 2018) are few examples. However, each new technology comes with its own pitfalls and limitations. In fact, these algorithms may unintentionally hold important biases that lead to unfairness in the task being performed. For example, the implementation of behavior forecasting algorithms in security borders or migration controls might lead to undesired outcomes (McKendrick, 2019) interfering with human rights (Akhmetova and Harris, 2021). Furthermore, the interacting user should always be aware of the presence of such forecasting systems, the possible manipulations or persuasion techniques attached, and their ultimate goal. Unfortunately, providing the user with these descriptions is not always easy because most of the times such systems are neither transparent nor explainable. Therefore, the incorporation of specific techniques to promote such interpretability is of utmost importance in order to build trust with the user. On the other side, it is also important to consider the potential vulnerabilities that such systems may have and how users might exploit them driven by unethical purposes. This is especially important for assistive or collaborative robots, which often involve very sensitive scenarios. Finally, although data protection regulations vary across countries (Guzzo et al., 2015), data privacy and data protection must ensure informational self-determination and consensual use of the information that can be extracted with the methods presented herein. In this sense, frameworks such as the EU General Data Protection Regulation (GDPR²) provide excellent safeguards for establishing ethical borders that should not be crossed.

7. Conclusion

In this survey, we provided an overview of the recent approaches proposed for non-verbal social behavior forecasting. We formulated a taxonomy that comprises and unifies recent (since 2017) attempts of forecasting low- or high-level representations of non-verbal social behavior up to date. By means of this taxonomy, we identified and described the main challenges of the problem, and analyzed how the recent literature has addressed them from both the sociological and the computer vision perspectives. We also presented all audiovisual datasets related to social behavior publicly released up to date in a summarized, structured, and friendly way. Finally, we described the most commonly used metrics, and the controversy that they often raise. We hope this survey can help bring the human motion prediction and the social signal forecasting worlds together in order to jointly tackle the main challenges of this field.

2. <https://gdpr.eu/>.

Acknowledgments

Isabelle Guyon was supported by ANR Chair of Artificial Intelligence HUMANIA ANR-19-CHIA-0022. This work has been partially supported by the Spanish project PID2019-105093GB-I00 and by ICREA under the ICREA Academia programme.

References

- Girmaw Abebe, Andreu Català, and Andrea Cavallaro. A first-person vision dataset of office activities. In *IAPR Workshop on Multimodal Pattern Recognition of Social Signals in Human-Computer Interaction*, pages 27–37, 2018.
- Vida Adeli, Ehsan Adeli, Ian Reid, Juan Carlos Nibbles, and Hamid Rezatofghi. Socially and contextually aware human motion and pose forecasting. *IEEE Robotics and Automation Letters*, 5:6033–6040, 10 2020.
- Vida Adeli, Mahsa Ehsanpour, Ian Reid, Juan Carlos Nibbles, Silvio Savarese, Ehsan Adeli, and Hamid Rezatofghi. Tripod: Human trajectory and pose dynamics forecasting in the wild. *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021.
- Nasir Ahmed, T. Natarajan, and Kamisetty R Rao. Discrete cosine transform. *IEEE transactions on Computers*, 100(1):90–93, 1974.
- Chaitanya Ahuja, Louis Philippe Morency, Yaser Sheikh, and Shugao Ma. To react or not to react: End-to-end visual pose forecasting for personalized avatar during dyadic conversations. *2019 International Conference on Multimodal Interaction*, pages 74–84, 10 2019.
- Louis Airale, Dominique Vaufreydaz, and Xavier Alameda-Pineda. Socialinteractiongan: Multi-person interaction sequence generation. *arXiv preprint arXiv:2103.05916*, 2021.
- Roxana Akhmetova and Erin Harris. Politics of technology: the use of artificial intelligence by us and canadian immigration agencies and their impacts on human rights. In *Digital Identity, Virtual Borders and Social Media*. Edward Elgar Publishing, 2021.
- Sames Al Moubayed, Malek Baklouti, Mohamed Chetouani, Thierry Dutoit, Ammar Mahdhaoui, J-C Martin, Stanislav Ondas, Catherine Pelachaud, Jérôme Urbain, and Mehmet Yilmaz. Generating robot/agent backchannels during a storytelling experiment. In *2009 IEEE International Conference on Robotics and Automation*, pages 3749–3754. IEEE, 2009.
- Xavier Alameda-Pineda, Jacopo Staiano, Subramanian Ramanathan, Ligia Maria Batrinca, Elisa Ricci, B. Lepri, Oswald Lanz, and N. Sebe. Salsa: A novel dataset for multimodal group behavior analysis. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 38:1707–1720, 2016.

- Sadegh Aliakbarian, Fatemeh Sadat Saleh, Mathieu Salzmann, Lars Petersson, and Stephen Gould. A stochastic conditioning scheme for diverse human motion prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5223–5232, 2020.
- Sadegh Aliakbarian, Fatemeh Saleh, Lars Petersson, Stephen Gould, and Mathieu Salzmann. Contextually plausible and diverse 3d human motion prediction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 11333–11342, 2021.
- Mykhaylo Andriluka, Umar Iqbal, Eldar Insafutdinov, Leonid Pishchulin, Anton Milan, Juergen Gall, and Bernt Schiele. Posetrack: A benchmark for human pose estimation and tracking. *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5167–5176, 2018.
- Sean Andrist, Bilge Mutlu, and Adriana Tapus. Look like me: matching robot personality via gaze to increase motivation. In *Proceedings of the 33rd annual ACM conference on human factors in computing systems*, pages 3603–3612, 2015.
- Sven Bambach, Stefan Lee, David J. Crandall, and Chen Yu. Lending a hand: Detecting hands and recognizing activities in complex egocentric interactions. *2015 IEEE International Conference on Computer Vision (ICCV)*, pages 1949–1957, 2015.
- German Barquero, Johnny Núñez, Zhen Xu, Sergio Escalera, Wei-Wei Tu, Isabelle Guyon, and Cristina Palmero. Comparison of spatio-temporal models for human motion and pose forecasting in face-to-face interaction scenarios. In *Understanding Social Behavior in Dyadic and Small Group Interactions*, volume 173 of *Proceedings of Machine Learning Research*, pages 107–138, 2022.
- Atef Ben-Youssef, Chloé Clavel, and Slim Essid. Early detection of user engagement breakdown in spontaneous human-humanoid interaction. *IEEE Transactions on Affective Computing*, 12:776–787, 7 2021.
- Roxane Bertrand, Philippe Blache, Robert Espesser, Gaëlle Ferré, Christine Meunier, Béatrice Priego-Valverde, and Stéphane Rauzy. Le cid-corpus of interactional data-annotation et exploitation multimodale de parole conversationnelle. *Traitement automatique des langues*, 49(3):pp–105, 2008.
- Cigdem Beyan, Nicoló Carissimi, Francesca Capozzi, Sebastiano Vascon, Matteo Bustreo, Antonio Pierro, Cristina Becchio, and Vittorio Murino. Detecting emergent leader in a meeting environment using nonverbal visual features only. *Proceedings of the 18th ACM International Conference on Multimodal Interaction*, 2016.
- Sanjay Bilakhia, Stavros Petridis, Anton Nijholt, and Maja Pantic. The mahnob mimicry database: A database of naturalistic human interactions. *Pattern Recognition Letters*, 66:52–61, 2015.
- Philippe Blache, Massina Abderrahmane, Stéphane Rauzy, and Roxane Bertrand. An integrated model for predicting backchannel feedbacks. *Proceedings of the 20th ACM International Conference on Intelligent Virtual Agents, IVA 2020*, 10 2020.

- Auriane Boudin, Roxane Bertrand, Stéphane Rauzy, Magalie Ochs, and Philippe Blache. A multimodal model for predicting conversational feedbacks. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 12848 LNAI:537–549, 2021.
- Elif Bozkurt, Hossein Khaki, Sinan Keçeci, Bekir Berker Turker, Yücel Yemez, and Engin Erzincin. The jestkod database: an affective multimodal database of dyadic interactions. *Language Resources and Evaluation*, 51:857–872, 2017.
- Laura Cabrera-Quiros, Andrew Demetriou, Ekin Gedik, Leander van der Meij, and Hayley Hung. The matchmingle dataset: a novel multi-sensor resource for the analysis of social interactions and group dynamics in-the-wild during free-standing conversations and speed dates. *IEEE Transactions on Affective Computing*, 12(1):113–130, 2018.
- Angelo Cafaro, Johannes Wagner, Tobias Baur, Soumia Dermouche, Mercedes Torres Torres, Catherine Pelachaud, Elisabeth André, and Michel F. Valstar. The noxi database: multimodal recordings of mediated novice-expert interactions. *Proceedings of the 19th ACM International Conference on Multimodal Interaction*, 2017.
- Yujun Cai, Yiwei Wang, Yiheng Zhu, Tat-Jen Cham, Jianfei Cai, Junsong Yuan, Jun Liu, Chuanxia Zheng, Sijie Yan, Henghui Ding, Xiaohui Shen, Ding Liu, and Nadia Magnenat-Thalmann. A unified 3d human motion synthesis model via conditional variational auto-encoder. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 11645–11655, 2021.
- Zhe Cao, Gines Hidalgo, Tomas Simon, Shih-En Wei, and Yaser Sheikh. Openpose: realtime multi-person 2d pose estimation using part affinity fields. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 43(1):172–186, 2019.
- João Carreira, Eric Noland, Chloe Hillier, and Andrew Zisserman. A short note on the kinetics-700 human action dataset. *arXiv preprint arXiv:1907.06987*, 2019.
- Oya Celiktutan, Efstratios Skordos, and Hatice Gunes. Multimodal human-human-robot interactions (mhhri) dataset for studying personality and engagement. *IEEE Transactions on Affective Computing*, 10:484–497, 2019.
- Huili Chen, Yue Zhang, Felix Weninger, Rosalind Picard, Cynthia Breazeal, and Hae Won Park. Dyadic speech-based affect recognition using dami-p2c parent-child multimodal interaction dataset. *Proceedings of the 2020 International Conference on Multimodal Interaction*, 2020.
- Zezhou Chen, Zhaoxiang Liu, Huan Hu, Jinqiang Bai, Shiguo Lian, Fuyuan Shi, and Kai Wang. A realistic face-to-face conversation system based on deep neural networks. *Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops*, 2019.
- Huang-Cheng Chou, Wei-Cheng Lin, Lien-Chiang Chang, Chyi-Chang Li, Hsi-Pin Ma, and Chi-Chun Lee. Nnime: The nthu-ntua chinese interactive multimodal emotion corpus.

- 2017 Seventh International Conference on Affective Computing and Intelligent Interaction (ACII)*, pages 292–298, 2017.
- Hang Chu, Daiqing Li, and Sanja Fidler. A face-to-face neural conversation model. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018.
- Enric Corona, Albert Pumarola, Guillem Alenya, Francesc Moreno-Noguer, and Grégory Rogez. Ganhand: Predicting human grasp affordances in multi-object scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5031–5041, 2020a.
- Enric Corona, Albert Pumarola, and Guillem Alenyà. Context-aware human motion prediction. *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6992–7001, 2020b.
- Robert O. Davis. The impact of pedagogical agent gesturing in multimedia learning environments: A meta-analysis. *Educational Research Review*, 24:193–219, 2018.
- Ellen Douglas-Cowie, Roddy Cowie, Ian Sneddon, Cate Cox, Orla Lowry, Margaret McRorie, Jean-Claude Martin, Laurence Devillers, Sarkis Abrilian, Anton Batliner, Noam Amir, and Kostas Karpouzis. The humane database: Addressing the collection and annotation of naturalistic and induced emotional data. In *International conference on affective computing and intelligent interaction*, pages 488–500, 2007.
- Metehan Doyran, Arjan Schimmel, Pinar Baki, Kübra Ergin, Batikan Türkmen, Almila Akdag Salah, Sander Bakkes, Heysem Kaya, Ronald Poppe, and A. A. Salah. Mumbai: multi-person, multimodal board game affect and interaction analysis dataset. *Journal on Multimodal User Interfaces*, 15(4):373–391, 2021.
- Jens Edlund, Jonas Beskow, Kjell Elenius, Kahl Hellmer, Sofia Strömbergsson, and David House. Spontal: A swedish spontaneous dialogue corpus of audio, video and motion capture. In *Proceedings of the Seventh International Conference on Language Resources and Evaluation*, 2010.
- Michael Edwards, Jingjing Deng, and Xianghua Xie. From pose to activity: Surveying datasets and introducing converse. *Computer Vision and Image Understanding*, 144: 73–105, 2016.
- Connor Esterwood and Lionel P Robert. A systematic review of human and robot personality in health care human-robot interaction. *Frontiers in Robotics and AI*, page 306, 2021.
- Will Feng, Anitha Kannan, Georgia Gkioxari, and C. Lawrence Zitnick. Learn2smile: Learning non-verbal interaction through observation. *IEEE International Conference on Intelligent Robots and Systems*, 2017-September:4131–4138, 12 2017.
- Siska Fitrianie, Merijn Bruijnes, Deborah Richards, Amal Abdulrahman, and Willem-Paul Brinkman. What are we measuring anyway? -a literature survey of questionnaires used in studies reported in the intelligent virtual agent conferences. In *Proceedings of the 19th ACM International Conference on Intelligent Virtual Agents*, pages 159–161, 2019.

- Siska Fitrianie, Merijn Bruijnes, Deborah Richards, Andrea Bönsch, and Willem-Paul Brinkman. The 19 unifying questionnaire constructs of artificial social agents: An iva community analysis. In *Proceedings of the 20th ACM International Conference on Intelligent Virtual Agents*, pages 1–8, 2020.
- Siska Fitrianie, Merijn Bruijnes, Fengxiang Li, and Willem-Paul Brinkman. Questionnaire items for evaluating artificial social agents-expert generated, content validated and reliability analysed. In *Proceedings of the 21st ACM International Conference on Intelligent Virtual Agents*, pages 84–86, 2021.
- Christos Georgakis, Yannis Panagakis, Stefanos Zafeiriou, and Maja Pantic. The conflict escalation resolution (confer) database. *Image and Vision Computing*, 65:37–48, 2017.
- Mononito Goswami, Minkush Manuja, and Maitree Leekha. Towards social & engaging peer learning: Predicting backchanneling and disengagement in children. *arXiv preprint arXiv:2007.11346*, 7 2020.
- Joseph Grafsgaard, Nicholas Duran, Ashley Randall, Chun Tao, and Sidney D’Mello. Generative multimodal models of nonverbal synchrony in close relationships. *Proceedings - 13th IEEE International Conference on Automatic Face and Gesture Recognition, FG 2018*, pages 195–202, 6 2018.
- Kristen Grauman, Andrew Westbury, Eugene Byrne, Zachary Q. Chavis, Antonino Furnari, Rohit Girdhar, Jackson Hamburger, Hao Jiang, Miao Liu, Xingyu Liu, Miguel Martin, Tushar Nagarajan, Ilija Radosavovic, Santhosh K. Ramakrishnan, F. Ryan, Jayant Sharma, Michael Wray, Mengmeng Xu, Eric Z. Xu, Chen Zhao, Siddhant Bansal, Dhruv Batra, Vincent Cartillier, Sean Crane, Tien Do, Morrie Doulaty, Akshay Erapalli, Christoph Feichtenhofer, Adriano Fragomeni, Qichen Fu, Christian Fuegen, Abraham Gebreselasie, Cristina Gonzalez, James M. Hillis, Xuhua Huang, Yifei Huang, Wenqi Jia, Weslie Khoo, Jáchym Kolár, Satwik Kottur, Anurag Kumar, Federico Landini, Chao Li, Yanghao Li, Zhenqiang Li, Karttikeya Mangalam, Raghava Modhugu, Jonathan Munro, Tullie Murrell, Takumi Nishiyasu, Will Price, Paola Ruiz Puentes, Merey Ramazanova, Leda Sari, Kiran K. Somasundaram, Audrey Southerland, Yusuke Sugano, Ruijie Tao, Minh Phuoc Vo, Yuchen Wang, Xindi Wu, Takuma Yagi, Yunyi Zhu, Pablo Arbeláez, David J. Crandall, Dima Damen, Giovanni Maria Farinella, Bernard Ghanem, Vamsi K. Ithapu, C. V. Jawahar, Hanbyul Joo, Kris Kitani, Haizhou Li, Richard A. Newcombe, Aude Oliva, Hyun Soo Park, James M. Rehg, Yoichi Sato, Jianbo Shi, Mike Zheng Shou, Antonio Torralba, Lorenzo Torresani, Mingfei Yan, and Jitendra Malik. Ego4d: Around the world in 3, 000 hours of egocentric video. *arXiv preprint arXiv:2110.07058*, 2021.
- Wen Guo, Xiaoyu Bie, Xavier Alameda-Pineda, and Francesc Moreno-Noguer. Multi-person extreme motion prediction with cross-interaction attention. *arXiv preprint arXiv:2105.08825*, 5 2021.
- Richard A Guzzo, Alexis A Fink, Eden King, Scott Tonidandel, and Ronald S Landis. Big data recommendations for industrial–organizational psychology. *Industrial and Organizational Psychology*, 8(4):491–508, 2015.

- Mohamed Hassan, Duygu Ceylan, Ruben Villegas, Jun Saito, Jimei Yang, Yi Zhou, and Michael J Black. Stochastic scene-aware motion prediction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 11374–11384, 2021.
- Yutaro Honda, Rei Kawakami, and Takeshi Naemura. Rnn-based motion prediction in competitive fencing considering interaction between players. *The British Machine Vision Conference*, 2020.
- Minjie Hua, Fuyuan Shi, Yibing Nan, Kai Wang, Hao Chen, and Shiguo Lian. Towards more realistic human-robot conversation: A seq2seq-based body gesture interaction system. *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 1393–1400, 2019.
- Hung Hsuan Huang, Seiya Kimura, Kazuhiro Kuwabara, and Toyoaki Nishida. Generation of head movements of a robot using multimodal features of peer participants in group discussion conversation. *Multimodal Technologies and Interaction 2020, Vol. 4, Page 15*, 4:15, 4 2020.
- Yuchi Huang and Saad M Khan. Dyadgan: Generating facial expressions in dyadic interactions. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, 2017.
- Hayley Hung and Gokul Chittaranjan. The idiap wolf corpus: exploring group behaviour in a competitive role-playing game. In *Proceedings of the 18th ACM international conference on Multimedia*, pages 879–882, 2010.
- Ryo Ishii, Shiro Kumano, and Kazuhiro Otsuka. Prediction of next-utterance timing using head movement in multi-party meetings. *Proceedings of the 5th International Conference on Human Agent Interaction*, 2017.
- Ryo Ishii, Kazuhiro Otsuka, Shiro Kumano, Ryuichiro Higashinaka, and Junji Tomita. Prediction of who will be next speaker and when using mouth-opening pattern in multi-party conversation. *Multimodal Technologies and Interaction 2019, Vol. 3, Page 70*, 3: 70, 10 2019.
- Ryo Ishii, Xutong Ren, Michal Muszynski, and Louis-Philippe Morency. Can prediction of turn-management willingness improve turn-changing modeling. *Proceedings of the 20th ACM International Conference on Intelligent Virtual Agents*, 2020.
- Ryo Ishii, Xutong Ren, Michal Muszynski, and Louis Philippe Morency. Multimodal and multitask approach to listener’s backchannel prediction: Can prediction of turn-changing and turn-management willingness improve backchannel modeling? *Proceedings of the 21st ACM International Conference on Intelligent Virtual Agents, IVA 2021*, 21:131–138, 9 2021.
- Vidit Jain and Maitree Leekha. Exploring semi-supervised learning for predicting listener backchannels. *Conference on Human Factors in Computing Systems - Proceedings*, 5 2021.

- Jin Yea Jang, San Kim, Minyoung Jung, Saim Shin, and Gahgene Gweon. Bpm_mt: Enhanced backchannel prediction model using multi-task learning. *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 3447–3452, 2021.
- Hanbyul Joo, Tomas Simon, Mina Cikara, and Yaser Sheikh. Towards social artificial intelligence: Nonverbal social signal prediction in a triadic interaction. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10873–10883, 2019a.
- Hanbyul Joo, Tomas Simon, Xulong Li, Hao Liu, Lei Tan, Lin Gui, Sean Banerjee, Timothy Godisart, Bart C. Nabbe, I. Matthews, Takeo Kanade, Shohei Nobuhara, and Yaser Sheikh. Panoptic studio: A massively multiview system for social interaction capture. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 41:190–204, 2019b.
- Isinsu Katircioglu, Costa Georgantas, Mathieu Salzmann, and Pascal Fua. Dyadic human motion prediction. *arXiv preprint arXiv:2112.00396*, 12 2021.
- Shahid Nawaz Khan, Maitree Leekha, Jainendra Shukla, and Rajiv Ratn Shah. Vyaktiv: A multimodal peer-to-peer hindi conversations based dataset for personality assessment. *2020 IEEE Sixth International Conference on Multimedia Big Data (BigMM)*, pages 103–111, 2020.
- Jean Kossaifi, Robert Walecki, Yannis Panagakis, Jie Shen, Maximilian Schmitt, Fabien Ringeval, Jing Han, Vedhas Pandit, Antoine Toisoul, Björn Schuller, et al. Sewa db: A rich database for audio-visual emotion and sentiment research in the wild. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 43(3):1022–1040, 2019.
- Robert M Krauss, Connie M Garlock, Peter D Bricker, and Lee E McMahon. The role of audible and visible back-channel responses in interpersonal communication. *Journal of personality and social psychology*, 35(7):523, 1977.
- Taras Kucherenko, Patrik Jonell, Youngwoo Yoon, Pieter Wolfert, and Gustav Eje Henter. A large, crowdsourced evaluation of gesture generation systems on common data: The genea challenge 2020. In *26th International Conference on Intelligent User Interfaces*, pages 11–21, 2021.
- Gilwoo Lee, Zhiwei Deng, Shugao Ma, Takaaki Shiratori, Siddhartha S. Srinivasa, and Yaser Sheikh. Talking with hands 16.2m: A large-scale dataset of synchronized body-finger motion and audio for conversational motion analysis and synthesis. *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 763–772, 2019.
- Jangwon Lee, Haodan Tan, David Crandall, and Selma Šabanović. Forecasting hand gestures for human-drone interaction. In *Companion of the 2018 ACM/IEEE International Conference on Human-Robot Interaction*, pages 167–168, 2018.
- Séverin Lemaignan, Charlotte Edmunds, Emmanuel Senft, and Tony Belpaeme. The pinsoro dataset: Supporting the data-driven study of child-child and child-robot social dynamics. *PLoS ONE*, 13, 2018.

- Yu Liu, Gelareh Mohammadi, Yang Song, and Wafa Johal. Speech-based gesture generation for robots and embodied agents: A scoping review. In *Proceedings of the 9th International Conference on Human-Agent Interaction*, pages 31–38, 2021.
- Andy Lücking, Kirsten Bergmann, Florian Hahn, Stefan Kopp, and Hannes Rieser. Data-based analysis of speech and gesture: the bielefeld speech and gesture alignment corpus (saga) and its applications. *Journal on Multimodal User Interfaces*, 7:5–18, 2012.
- Ren C Luo and Licong Mai. Human intention inference and on-line human hand motion prediction for human-robot collaboration. In *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 5958–5964. IEEE, 2019.
- Lucien Maman, Eleonora Ceccaldi, Nale Lehmann-Willenbrock, Laurence Likforman-Sulem, Mohamed Chetouani, Gualtiero Volpe, and Giovanna Varni. Game-on: A multimodal dataset for cohesion and group analysis. *IEEE Access*, 8:124185–124203, 2020.
- Wei Mao, Miaomiao Liu, Mathieu Salzmann, and Hongdong Li. Learning trajectory dependencies for human motion prediction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9489–9497, 2019.
- Wei Mao, Miaomiao Liu, and Mathieu Salzmann. History repeats itself: Human motion prediction via motion attention. *European Conference on Computer Vision*, pages 474–489, 2020.
- Wei Mao, Miaomiao Liu, and Mathieu Salzmann. Generating smooth pose sequences for diverse human motion prediction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 13309–13318, 2021a.
- Wei Mao, Miaomiao Liu, Mathieu Salzmann, and Hongdong Li. Multi-level motion attention for human motion prediction. *International Journal of Computer Vision*, pages 1–23, 2021b.
- Roberto Martín-Martín, Mihir Patel, Hamid Rezaatofghi, Abhijeet Shenoi, JunYoung Gwak, Eric Frankel, Amir Sadeghian, and Silvio Savarese. Jrdb: A dataset and benchmark of egocentric robot visual perception of humans in built environments. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PP, 2021.
- Julieta Martinez, Michael J. Black, and Javier Romero. On human motion prediction using recurrent neural networks. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2891–2900, 2017.
- Iain McCowan, Jean Carletta, Wessel Kraaij, Simone Ashby, S Bourban, M Flynn, M Guillemot, Thomas Hain, J Kadlec, Vasilis Karaiskos, et al. The ami meeting corpus. In *Proceedings of the 5th international conference on methods and techniques in behavioral research*, volume 88, page 100, 2005.
- K. McKendrick. *Artificial intelligence prediction and counterterrorism*. London: UK, 2019.

- Gary McKeown, Michel F. Valstar, Roddy Cowie, and Maja Pantic. The semaine corpus of emotionally coloured character interactions. *2010 IEEE International Conference on Multimedia and Expo*, pages 1079–1084, 2010.
- Dushyant Mehta, Oleksandr Sotnychenko, Franziska Mueller, Weipeng Xu, Srinath Sridhar, Gerard Pons-Moll, and Christian Theobalt. Single-shot multi-person 3d pose estimation from monocular rgb. *Proceedings - 2018 International Conference on 3D Vision, 3DV 2018*, pages 120–130, 10 2018.
- Mathew Monfort, Bolei Zhou, Sarah Adel Bargal, Alex Andonian, Tom Yan, Kandan Ramakrishnan, Lisa M. Brown, Quanfu Fan, Dan Gutfreund, Carl Vondrick, and Aude Oliva. Moments in time dataset: One million videos for event understanding. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 42:502–508, 2020.
- Hee-Seung Moon and Jiwon Seo. Fast user adaptation for human motion prediction in physical human–robot interaction. *IEEE Robotics and Automation Letters*, 7(1):120–127, 2021.
- Lucas Mourot, Ludovic Hoyet, François Le Clerc, François Schnitzler, and Pierre Hellier. A survey on deep learning for skeleton-based human animation. In *Computer Graphics Forum*. Wiley Online Library, 2021.
- Michael Murray, Nick Walker, Amal Nanavati, Patricia Alves-Oliveira, Nikita Filippov, Allison Sauppe, Bilge Mutlu, and Maya Cakmak. Learning backchanneling behaviors for a social robot via data augmentation from human-human conversations. *5th Annual Conference on Robot Learning*, 6 2021.
- Philipp Müller, Ekta Sood, and Andreas Bulling. Anticipating averted gaze in dyadic interactions. *ACM Symposium on Eye Tracking Research and Applications*, 2020.
- Iftexhar Naim, M Iftexhar Tanveer, Daniel Gildea, and Mohammed Ehsan Hoque. Automated prediction and analysis of job interview performance: The role of what you say and how you say it. In *2015 11th IEEE international conference and workshops on automatic face and gesture recognition (FG)*, volume 1, pages 1–6. IEEE, 2015.
- Lukasz Okruszek, Aleksandra Piejka, Adam Wysokiński, Ewa Szczepocka, and Valeria Manera. The second agent effect: interpersonal predictive coding in people with schizophrenia. *Social Neuroscience*, 14(2):208–213, 2019.
- Stanislav Ondáš and Matúš Pleva. Anticipation and its applications in human-machine interaction. In *Proceedings of the 19th Conference Information Technologies-Applications and Theory (ITAT 2019)*, pages 152–156, 2019.
- Sergiu Oprea, Pablo Martinez-Gonzalez, Alberto Garcia-Garcia, John Alejandro Castro-Vargas, Sergio Orts-Escolano, Jose Garcia-Rodriguez, and Antonis Argyros. A review on deep learning techniques for video prediction. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2020.

- Daniel Ortega, Chia Yu Li, and Ngoc Thang Vu. Oh, jeez! or uh-huh? a listener-aware backchannel predictor on asr transcriptions. *ICASSP, IEEE International Conference on Acoustics, Speech and Signal Processing - Proceedings*, 2020-May:8064–8068, 5 2020.
- Patrizia Paggio and Costanza Navarretta. The danish nomco corpus: multimodal interaction in first acquaintance conversations. *Language Resources and Evaluation*, 51:463–494, 2017.
- Cristina Palmero, German Barquero, Julio C. S. Jacques Junior, Albert Clapés, Johnny Núñez, David Curto, Sorina Smeureanu, Javier Selva, Zejian Zhang, David Saeteros, David Gallardo-Pujol, Georgina Guilera, David Leiva, Feng Han, Xiaoxue Feng, Jennifer He, Wei-Wei Tu, Thomas B. Moeslund, Isabelle Guyon, and Sergio Escalera. Chalearn LAP challenges on self-reported personality recognition and non-verbal behavior forecasting during social dyadic interactions: Dataset, design, and results. In *Understanding Social Behavior in Dyadic and Small Group Interactions*, volume 173 of *Proceedings of Machine Learning Research (PMLR)*, pages 4–52, 2022.
- Cheul Young Park, Narae Cha, Soowon Kang, Auk Kim, Ahsan H. Khandoker, Leontios J. Hadjileontiadis, Alice H. Oh, Yong Jeong, and Uichin Lee. K-emocon, a multimodal sensor dataset for continuous emotion recognition in naturalistic conversations. *Scientific Data*, 7, 2020.
- Ronald Poppe, Khiet P Truong, Dennis Reidsma, and Dirk Heylen. Backchannel strategies for artificial listeners. In *International Conference on Intelligent Virtual Agents*, pages 146–158. Springer, 2010.
- Chirag Raman, Hayley Hung, and Marco Loog. Social processes: Self-supervised forecasting of nonverbal cues in social conversations. *arXiv preprint arXiv:2107.13576*, 7 2021.
- James M. Rehg, Gregory D. Abowd, Agata Rozga, M. Romero, Mark A. Clements, Stan Sclaroff, Irfan Essa, Opal Y. Ousley, Yin Li, Chanh Kim, Hrishikesh Rao, Jonathan C. Kim, Liliana Lo Presti, Jianming Zhang, Denis Lantsman, Jonathan Bidwell, and Zhefan Ye. Decoding children’s social behavior. *2013 IEEE Conference on Computer Vision and Pattern Recognition*, pages 3414–3421, 2013.
- Harry T Reis, W Andrew Collins, and Ellen Berscheid. The relationship context of human behavior and development. *Psychological bulletin*, 126(6):844, 2000.
- Fabien Ringeval, Andreas Sonderegger, Jürgen S. Sauer, and Denis Lalanne. Introducing the recola multimodal corpus of remote collaborative and affective interactions. *2013 10th IEEE International Conference and Workshops on Automatic Face and Gesture Recognition (FG)*, pages 1–8, 2013.
- Takaaki Saeki, Shinnosuke Takamichi, and Hiroshi Saruwatari. Incremental text-to-speech synthesis using pseudo lookahead with large pretrained language model. *IEEE Signal Processing Letters*, 28:857–861, 2021.
- David A. Salter, Amir Tamrakar, Behjat Siddiquie, Mohamed R. Amer, Ajay Divakaran, Brian Lande, and Darius Mehri. The tower game dataset: A multimodal dataset for

- analyzing social interaction predicates. *2015 International Conference on Affective Computing and Intelligent Interaction (ACII)*, pages 656–662, 2015.
- Dairazalia Sanchez-Cortes, Oya Aran, Marianne Schmid Mast, and Daniel Gática-Pérez. A nonverbal behavior approach to identify emergent leaders in small groups. *IEEE Transactions on Multimedia*, 14:816–832, 2012.
- Navyata Sanghvi, Ryo Yonetani, and Kris Kitani. Mgpi: A computational model of multi-agent group perception and interaction. In *Proceedings of the 19th International Conference on Autonomous Agents and MultiAgent Systems*, pages 1196–1205, 2020.
- Laura Schiphorst, Metehan Doyran, Sabine Molenaar, A. A. Salah, and Sjaak Brinkkemper. Video2report: A video database for automatic reporting of medical consultancy sessions. *2020 15th IEEE International Conference on Automatic Face and Gesture Recognition (FG 2020)*, pages 552–556, 2020.
- Kevin Sexton, Amanda Johnson, Amanda Gotsch, Ahmed A Hussein, Lora Cavuoto, and Khurshid A Guru. Anticipation, teamwork and cognitive load: chasing efficiency during robot-assisted surgery. *BMJ quality & safety*, 27(2):148–154, 2018.
- Tianmin Shu, M. S. Ryoo, and Song-Chun Zhu. Learning social affordance for human-robot interaction. *IJCAI International Joint Conference on Artificial Intelligence*, 2016-January:3454–3461, 4 2016.
- Jainendra Shukla, Miguel Barreda-Ángeles, Joan Oliver, and Domenec Puig. Muder: Multimodal database for emotion recognition among intellectually disabled individuals. In *The Eight International Conference on Social Robotics*, 2016.
- Michel Silva, Washington Ramos, Joao Ferreira, Felipe Chamone, Mario Campos, and Erickson R Nascimento. A weighted sparse sampling and smoothing frame transition approach for semantic fast-forward first-person videos. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2383–2392, 2018.
- Khurram Soomro, Amir Roshan Zamir, and Mubarak Shah. Ucf101: A dataset of 101 human actions classes from videos in the wild. *arXiv preprint arXiv:1212.0402*, 2012.
- Mark A Thornton, Miriam E Weaverdyck, and Diana I Tamir. The social brain automatically predicts others’ future mental states. *Journal of Neuroscience*, 39(1):140–148, 2019.
- Bekir Berker Türker, Engin Erzin, Yücel Yemez, and Metin Sezgin. Audio-visual prediction of head-nod and turn-taking events in dyadic interactions. *Interspeech*, pages 1741–1745, 2018.
- Ryosuke Ueno, Yukiko I. Nakano, Jie Zeng, and Fumio Nihei. Estimating the intensity of facial expressions accompanying feedback responses in multiparty video-mediated communication. *ICMI 2020 - Proceedings of the 2020 International Conference on Multimodal Interaction*, 20:144–152, 10 2020.

- Felix van Doorn. Rituals of leaving: Predictive modelling of leaving behaviour in conversation, 2018.
- Rob J.J.H. van Son, Wieneke Wesseling, Eric Sanders, and Henk van den Heuvel. The ifadv corpus: a free dialog video corpus. In *Proceedings of the Sixth International Conference on Language Resources and Evaluation (LREC'08)*, 2008.
- Jason Vandeventer, Andrew J. Aubrey, Paul L. Rosin, and A. David Marshall. 4d cardiff conversation database (4d ccdb): a 4d database of natural, dyadic conversations. *Auditory-Visual Speech Processing*, 2015.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008, 2017.
- Alexandra Vella and Patrizia Paggio. Overlaps in maltese: a comparison between map task dialogues and multimodal conversational data. In *4th Nordic Symposium on Multimodal Communication*, pages 21–29, 2013.
- Alessandro Vinciarelli, Maja Pantic, and Hervé Bourlard. Social signal processing: Survey of an emerging domain. *Image and vision computing*, 27(12):1743–1759, 2009.
- Timo von Marcard, Roberto Henschel, Michael J. Black, Bodo Rosenhahn, and Gerard Pons-Moll. Recovering accurate 3d human pose in the wild using imus and a moving camera. *European Conference on Computer Vision (ECCV)*, pages 601–617, 2018.
- Jacob Walker, Kenneth Marino, Abhinav Gupta, and Martial Hebert. The pose knows: Video forecasting by generating pose futures. In *Proceedings of the IEEE international conference on computer vision*, pages 3332–3341, 2017.
- Kevin S Walsh, David P McGovern, Andy Clark, and Redmond G O’Connell. Evaluating the neurophysiological evidence for predictive processing as a model of perception. *Annals of the new York Academy of Sciences*, 1464(1):242–268, 2020.
- Chenxi Wang, Yunfeng Wang, Zixuan Huang, and Zhiwen Chen. Simple baseline for single human motion forecasting. *ICCV SoMoF Workshop*, 2021a.
- Jiashun Wang, Huazhe Xu, Medhini Narasimhan, Xiaolong Wang, and Uc San Diego. Multi-person 3d motion prediction with multi-range transformers. *Thirty-Fifth Conference on Neural Information Processing Systems*, 2021b.
- Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612, 2004.
- Pieter Wolfert, Jeffrey M Girard, Taras Kucherenko, and Tony Belpaeme. To rate or not to rate: Investigating evaluation methods for generated co-speech gestures. In *Proceedings of the 2021 International Conference on Multimodal Interaction*, pages 494–502, 2021a.

- Pieter Wolfert, Nicole Robinson, and Tony Belpaeme. A review of evaluation practices of gesture generation in embodied conversational agents. *arXiv preprint arXiv:2101.03769*, 2021b.
- Jieyeon Woo, Catherine Pelachaud, and Catherine Achard. Creating an interactive human/agent loop using multimodal recurrent neural networks. *WACAI*, 2021.
- Dingdong Yang, Seunghoon Hong, Yunseok Jang, Tianchen Zhao, and Honglak Lee. Diversity-sensitive conditional generative adversarial networks. In *International Conference on Learning Representations*, 2018.
- Fangkai Yang, Yuan Gao, Ruiyang Ma, Sahba Zojaji, Ginevra Castellano, and Christopher E. Peters. A dataset of human and robot approach behaviors into small free-standing conversational groups. *PLoS ONE*, 16, 2021.
- Mohammad Samin Yasar and Tariq Iqbal. A scalable approach to predict multi-agent motion for human-robot collaboration. *IEEE Robotics and Automation Letters*, 6(2): 1686–1693, 2021.
- Ryo Yonetani, Kris M. Kitani, and Yoichi Sato. Recognizing micro-actions and reactions from paired egocentric videos. *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2629–2638, 2016.
- Ye Yuan and Kris Kitani. Dlow: Diversifying latent flows for diverse human motion prediction. In *European Conference on Computer Vision*, pages 346–364. Springer, 2020.
- Ye Yuan and Kris M Kitani. Diverse trajectory forecasting with determinantal point processes. In *International Conference on Learning Representations*, 2019.
- Yu Zhang and Qiang Yang. A survey on multi-task learning. *IEEE Transactions on Knowledge and Data Engineering*, 2021.
- Hang Zhao, Zhicheng Yan, Heng Wang, and Lorenzo Torresani. Hacs: Human action clips and segments dataset for recognition and temporal localization. *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 8667–8677, 2019.
- Yang Zhao and Yong Dou. Pose-forecasting aided human video prediction with graph convolutional networks. *IEEE Access*, 8:147256–147264, 2020.