

# Social Cue Detection and Analysis Using Transfer Entropy

Haoyang Jiang  
Monash University  
Melbourne, Victoria, Australia  
haoyang.jiang@monash.edu

Elizabeth A. Croft  
University of Victoria  
Victoria, British Columbia, Canada  
Monash University  
Melbourne, Victoria, Australia  
ecroft@uvic.ca

Michael G. Burke  
Monash University  
Melbourne, Victoria, Australia  
University of Edinburgh  
Edinburgh, Scotland, United Kingdom  
michael.g.burke@monash.edu

## ABSTRACT

Robots that work close to humans need to understand and use social cues to act in a socially acceptable manner. Social cues are a form of communication (i.e., information flow) between people. In this paper<sup>1</sup>, a framework is introduced to detect and analyse a class of perceptible social cues that are nonverbal and episodic, and the related information transfer using an information-theoretic measure, namely, transfer entropy. We use a group-joining setting to demonstrate the practicality of transfer entropy for analysing communications between humans. Then we demonstrate the framework in two settings involving social interactions between humans: object-handover and person-following. Our results show that transfer entropy can identify information flows between agents and when and where they occur. Potential applications of the framework include information flow or social cue analysis for interactive robot design and socially-aware robot planning.

## CCS CONCEPTS

• **Human-centered computing** → *Collaborative and social computing*.

## KEYWORDS

Social cues, Transfer entropy, Nonverbal communication, Socially-aware robots

## 1 INTRODUCTION

Social robots are a category of robots that work physically close to humans and are designed to interact with non-expert users. User acceptance requires that these robots operate in a socially acceptable manner and are able to react to, or exchange, social cues that play an important role during interaction. While people frequently exchange social cues using verbal and gestural signals, they also exchange rich information through their gait, posture and walking patterns [34]. For instance, in a handover task, a giver typically reaches forward to indicate their intent to give an item and a receiver observes this cue and reaches forward indicating readiness to accept the item. Many nonverbal, subtle cues like this reaching cue might not be immediately apparent to a casual observer, but play an important role when people collaborate or interact with each other [44].

Socially-aware behaviour is a multi-lateral process, and to be useful in robots it needs to be predictable, adaptable and easily understood by humans [34]. This requires that social robots sense and react to social information conveyed by humans while simultaneously conveying useful social information to humans. However, due

to hardware limitations, environmental constraints and the robot's primary task, not all robots can communicate through dedicated audio or visual display channels. Therefore, being able to capture and use measurable social cues embedded in kinematic measurements (e.g., pose and motion) is an important step for social robots to achieve more acceptable and anthropomorphic behaviours while working with humans. Detecting and reacting to subtle social cues is crucial to move beyond explicit and often contrived or overt human-robot interaction. Unfortunately, this is very challenging because humans or agents can express themselves in a variety of ways, which means that the cues themselves are often difficult to elucidate. This work seeks to address this challenge by defining a class of measurable social cues and proposing a framework that allows these cues to be detected automatically.

In signal detection theory, Green and Swets [15] introduced the concept of *response bias*, which refers to how much evidence an observer requires before responding to a given signal. Our cue detection framework builds on this theoretical notion. In our framework, the given signal refers to social cues. If a cue is below the response bias of the observer, it is a silent cue, meaning that one cannot tell if a cue exists by solely observing the observer. For the purposes of this work, we focus on those cues that are above the response bias. We define a *perceptible social cue* as *an event generated by an agent that influences the behaviour of the other agent observing the cue*. This paper aims to formulate a method to detect perceptible cues of this form.

The proposed perceptible social cue analysis framework accepts raw information captured by sensors on robots, and seeks to primarily provide answers to the following questions:

- When or where is a perceptible social cue activated?
- What is the direction of the exchanged cue?

A core contribution of the proposed framework is to model implicit communication between agents, and thereby the exchange of perceptible social cues, as an *exchange of information*. This allows us to use information-theoretic approaches to measure continuous information flows between agents, and threshold these to identify perceptible social cues. Herein, cue transfer is analysed using an information-theoretic measure, Transfer Entropy (TE), a statistical measure of the amount of directed transfer of information between two systems [35]. Other contributions include:

- a general framework that is able to detect arbitrary cues from raw data of social interactions without pre-designing or predefining a set of cues.
- a method of computing TE locally with multi-dimensional features using neural networks.

<sup>1</sup>This paper has been accepted by HRI'24.

We validate the framework in three unique settings: group joining, handover and person-following, to showcase the broad applicability of perceptible social cue analysis.

## 2 BACKGROUND

### 2.1 Social Cues

It is broadly recognised that social cues play an important role in communication. However, for social cue detection, researchers usually predefine a fixed set of social cues composed of known postures and gestures such as nodding or head shaking [43][5][6]. NovA [2] is a well-developed system for nonverbal signal detection, which detects postures and gestures using event-based gesture analysis [23], and classifies them into a predefined set of cues. This system also provides a movement expressivity measurement. However, these measurements are not used for cue detection.

Social cues are widely used in robotics, especially for Human-Robot Interaction (HRI). Tomari et al. proposed a socially-aware navigation planner for wheelchair robots, tracking the head orientations of the participants to estimate personal space assuming humans are more protective of the space in front of them [41]. Hansen et al. developed an adaptive system for natural interaction between mobile robots and humans based on the person’s pose and position, which estimates the interaction intention of the user and then uses it as a basis for socially-aware navigation based on a person’s social space [42]. Escobedo et al. used the commonly visited destinations of a wheelchair robot’s user to estimate the probable intended destination of the user and then accept the user’s face and voice commands for navigation [11]. These works are dependent on the proxemics theory proposed by Hall, who defined general social zones of humans [16]. However, the concept of social zones relies on averaged heuristics. Hall validated his study only for US citizens [34], which means these metrics may not translate well to other cultural contexts. Many human-following robot designs use social cues, e.g., in [28] where body orientation is used to estimate the intended turning direction. In [30], the relative position between a human and a robot is used as a feature to anticipate the human. Hu et al. [20] use human orientation as an input for anticipatory robot behaviours. In general, most robots are designed to act on heuristic or manually selected social cues, and it is unclear whether these social cues are reliable or repeatable. Consequently, most current research focuses on detecting these cues to predict intent or behavior.

Cue detection usually requires tuning and is not necessarily repeatable or cross-cultural. There are no examples of research where interaction cues are automatically sourced. We tackle these issues in this work. The perceptible social cue detection framework proposed here identifies social cues both spatially and in time from raw data, requiring no predefined cue sets and comparatively less manual specification. Since it is data-driven, the framework is agnostic to social norms or heuristics and can be applied to any context where motion data can be tracked.

### 2.2 Communicating intent

The study of social cues is related to intent communication, which refers to behaviours that allow an observer to quickly and correctly infer the intention of the agent generating the behaviour [7][10].

The research in [26] shows that motions that communicate intent increase perceived safety during virtual human-robot path-crossing tasks. In [7], the authors consider robustness, efficiency and energy as universal costs in a reinforcement learning scheme, with results showing an increase in the ability of a human to interpret a robot’s intention. Research studying projected visual legibility cues [19], has shown that projected arrows are generally more interpretable than flashing lights in a navigation setting. Importantly, the ways people communicate and understand intent are different, and there is no universal method to measure the communication of intent. However, if communicating intent is interpreted as the transfer of information between agents, this concept can be captured by the proposed framework, which could be used as a standard for measurement and to guide the selection of robot motions or actions.

### 2.3 Transfer Entropy

TE is a measure that allows the analysis of the information transfer and potential causal relationships between two simultaneous time series. In the economics literature, Baek et al. use TE to analyse the market influence of companies in the U.S. stock market [1]. He and Shang [17] compare different TE methods for analysing the relationship between 9 stock indices from the U.S., Europe and China. TE has been used to analyse animal-animal or animal-robot interactions such as [36][33]. It has also been adopted to analyse joint attention [38] and model pedestrian evacuation [45]. Berger et al. apply TE in robotics to detect human-to-robot perturbations using low-cost sensors [3]. The above studies succeed in quantizing the information transfer or the relationship between their targets using TE, but tend to focus on specific features or aspects of interest. Our work seeks to provide a more general framework for detecting and analysing perceptible social cues, by looking for changes in information transfer over time and in space.

Mutual information (MI) is an information measure that captures the shared information between random variables. Klyubin et al. proposed the concept of empowerment for intrinsic motivation for reinforcement learning, which measures the maximum MI (channel capacity) between the agent’s actuation and their sensors [24]. This is expanded by Mohamed and Rezende with a lower complexity maximisation approach to MI [29]. Jaques et al. use MI as a social influence reward for multi-agent deep reinforcement learning in Sequential Social Dilemmas (SSDs) [25] to encourage collaborations between agents [21]. While MI measures the shared information, TE measures the time-asymmetric information transfer. Unlike MI, the asymmetric property of TE allows us to analyse the directionality of information flow, which is beneficial for analysing the exchange of social cues. By applying TE, we aim to quantize social cue transfer to better understand and use social cues for robots. In addition, most TE-related research reduces multi-dimensional features down to a single dimension to compute TE, while in practice, most features are multi-dimensional. We tackle this issue by using neural networks to build the probability distributions needed to estimate TE.

## 3 METHODOLOGY

We model a perceptible cue as an exchange of information between two agents, using information-theoretic approaches to measure the levels of information exchanged. While it is difficult to measure

extremely subtle information transfer, given sensor capability limits and the varied time horizons over which cues can occur, it is possible to identify a restricted subset of cues within the information stream. To enable this, we make the following assumptions:

- To be detected, a perceptible cue needs to exceed some significance level or base threshold of information transfer between agents.
- To be detected, a perceptible cue needs to occur within a finite time period.

Our framework is illustrated in Fig. 1 and described below.

### 3.1 Transfer Entropy

Transfer entropy is closely related to the concept of *Wiener-Granger causality*<sup>2</sup> [14]. TE,  $T_{Y \rightarrow X}$ , can be defined as the conditional mutual information between two variables  $X_t$  and  $Y_t$  [4], which is formulated as follows.

$$T_{Y \rightarrow X}^{(k,l)}(t) = I(X_t : Y_t^{(l)} | X_t^{(k)}) \\ = H(X_t | X_t^{(k)}) - H(X_t | X_t^{(k)}, Y_t^{(l)}) \quad (1)$$

Here,  $I$  denotes the mutual information and  $H$  Shannon's entropy [37]. This equation defines the information transfer from  $Y$  to  $X$ , where  $t$  is the time, and  $k, l$  are the history length of  $X_t$  and  $Y_t$ . In some literature,  $X$  is termed the target and  $Y$  the source. TE can be intuitively interpreted as the reduction of uncertainty in a state  $X$  predicted solely based on its own history when an additional information source  $Y$  is introduced [4]. The key idea of the proposed framework is to use TE to measure the information transfer of perceptible social cues.

### 3.2 Overall framework workflow

**1. Preparation.** We start by identifying a target  $X$  and source  $Y$ . In social cue analysis scenarios, the source can be a feature that transmits cues of interest (e.g., a pedestrian's head orientation), and the target features that would be influenced by the cues (e.g., another pedestrian's pose). Next, time series data is collected for both the target and source. Prior to modelling, we use Takens' delay embedding [39] to create higher-dimensional embeddings for the time-history series of features,

$$X_t^{(d)} = (x(t - \delta), x(t - 2\delta), \dots, x(t - (d - 1)\delta)) \\ Y_t^{(d)} = (y(t - \delta), y(t - 2\delta), \dots, y(t - (d - 1)\delta)), \quad (2)$$

where  $d$  is the history length and  $\delta$  is the unit time step. The history length for the target and source do not necessarily need to be the same, one can also denote them separately such as  $k, l$  in Eq. (1). In general, the history needs to be long enough to capture the temporal relationship between cue cause and effect. In our experiments, we match this with estimates of human response times.

**2. Model the Target.** We next establish a baseline model, conditioned on only the history features of the target. The framework allows the application of different probabilistic prediction models. For instance, we can use a simple linear vector autoregressive model (VAR) to model the target  $X_t$ .

$$X_t = c + \alpha X_t^{(d)} + \epsilon \quad (3)$$

where  $c$  is the constant intercept of the model,  $\alpha$  is the time-invariant matrix that matches the dimension of  $X_t^{(d)}$  and  $\epsilon$  an error term. We can also model the target using other modelling techniques such as a multilayer perceptron (MLP), a Gaussian process or potentially more complex neural network architectures. Regardless of the underlying model, we intentionally form the target output distribution as a conditional multivariate Gaussian with mean  $f(X_t^{(d)})$  and covariance  $f^\sigma(X_t^{(d)})$ ,

$$P(X_t | X_t^{(d)}) \sim \mathcal{N}_{X_t}(f(X_t^{(d)}), f^\sigma(X_t^{(d)})), \quad (4)$$

This simplifies later TE calculations. After obtaining the base model conditioned only on the target history, the source is included to build a second model of the target. For example, assuming a vector autoregressive model, we obtain

$$X_t = c + \alpha X_t^{(d)} + \beta Y_t^{(d)} + \epsilon \\ P(X_t | X_t^{(d)}, Y_t^{(d)}) \sim \mathcal{N}_{X_t}(f(X_t^{(d)}, Y_t^{(d)}), f^\sigma(X_t^{(d)}, Y_t^{(d)})). \quad (5)$$

Here, we illustrate the process using a vector autoregressive model, but models that capture non-linear effects could also be used. We use neural networks to model behaviours for the experiments below.

**3. Compute and Analyse Transfer Entropy.** Modelling targets using conditional multivariate Gaussian random variables allows Shannon's differential entropy to be calculated as follows,

$$H(\mathbf{x}) = - \int p(\mathbf{x}) \log p(\mathbf{x}) d\mathbf{x} \\ = \frac{D}{2} (1 + \log(2\pi)) + \frac{1}{2} \log |\sigma| \quad (6)$$

Here,  $D$  denotes the dimension of the target variable. Using (6) to calculate the entropy for each model, one can calculate the TE from the source to target as

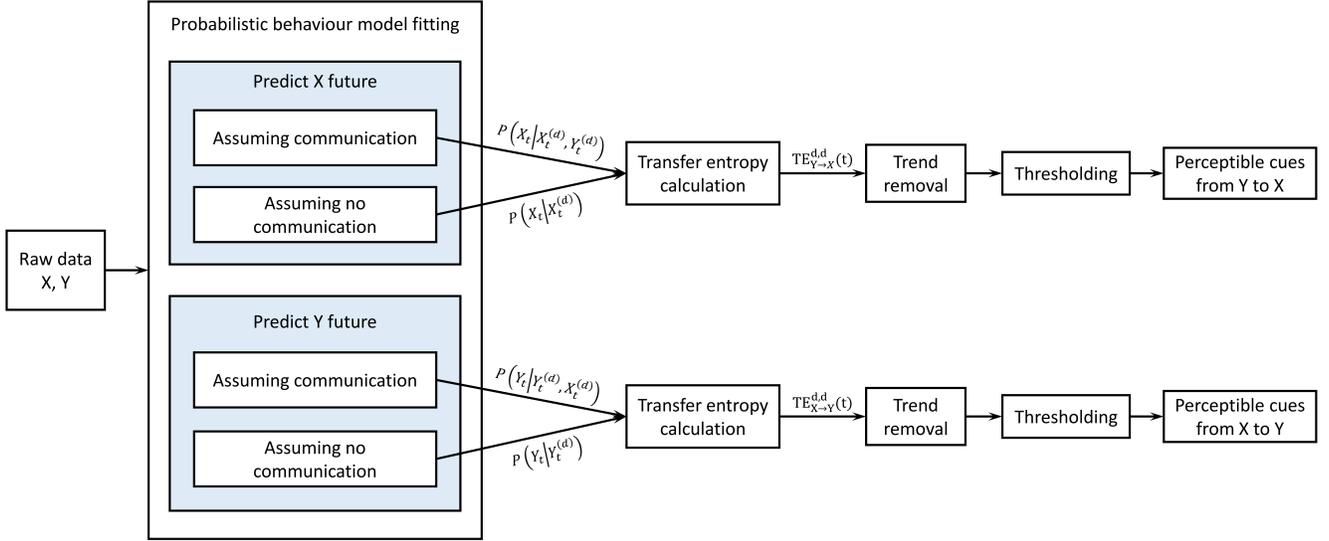
$$T_{Y \rightarrow X}^{(d,d)}(t) = I(X_t : Y_t^{(d)} | X_t^{(d)}) \\ = H(X_t | X_t^{(d)}) - H(X_t | X_t^{(d)}, Y_t^{(d)}) \quad (7)$$

Finally, we can analyse the information transfer from the source to the target using the measured TE.

### 3.3 Computing Transfer Entropy

The signals received from agents are usually in continuous form. Researchers often discretise these continuous signals in order to statistically compute TE using histograms or other frequentist approaches [3][31]. Kernel density estimation has also been recommended in [35] and frequently used in the information analysis toolbox [27], a popular tool for modelling TE in economics. However, these methods require a large amount of data to build sufficiently dense probability distributions, especially if the feature dimension is high [31]. Therefore, they are not helpful for local time series analysis with a limited number of samples, which means these methods are not suitable for computing local TE in order to analyse information transfer spatially or temporally. Even if we have sufficient data, building high-dimensional density distributions is not computationally efficient, which means these approaches are unsuitable for online algorithms. To tackle this issue, we use neural networks in our framework to estimate distributions, and then calculate entropy directly from the estimated conditional probabilities.

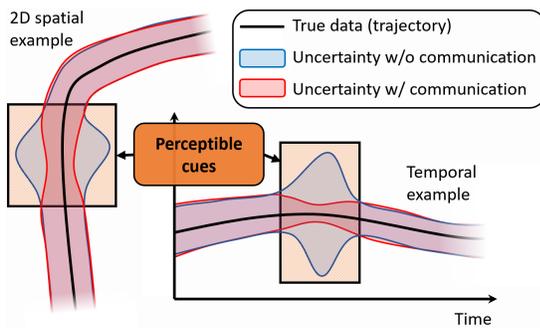
<sup>2</sup>These are equivalent when linear Gaussian models are used.



**Figure 1: Illustration of the proposed framework.**  $X, Y$  represent two agents in a social interaction scenario. Raw data is used to fit probabilistic behaviour model with different conditions, the output conditional probabilities are then used to compute transfer entropy (TE) in both information transfer directions. Double exponential smoothing (DES) is used to remove the general TE trend followed by the thresholding in order to find perceptible cues using the TE.

This allows us to estimate TE locally and continuously, which is also computationally more efficient and achieves spatial and temporal information transfer analysis. It should be noted that since our data is continuous and entropy is computed using Shannon’s differential entropy, the transfer entropy can be negative indicating the absence of information being communicated.

### 3.4 Cue detection



**Figure 2: An illustration of the perceptible cue detection framework.**

A graphical illustration of the intuition underlying perceptible cue detection is given in Fig. 2. We use the TE to detect perceptible social cues.<sup>3</sup> Positive TE means the information transfers from the source to the target. However, there can be continuous information

<sup>3</sup>We provide an example at <https://github.com/jhy9968/TEscd.git>

transfer during social interactions. We are most interested in the set of special events that trigger observer responses. Therefore, to meet the criteria of a perceptible social cue in our framework, we require TE measures to meet the requirements that: 1) the TE has to be larger than zero indicating general information transfer; and 2) the TE has to be larger than a threshold to be identified as a perceptible cue. To determine the threshold, we use double exponential smoothing (DES) [13] to find the moving mean and standard deviation of the TE. Let us denote  $\mu_t$ ,  $\sigma_t$  and  $T_t$  as the moving mean, moving standard deviation of TE and current TE at time  $t$ . The DES works as follows:

$$\begin{aligned} \mu_t &= \alpha T_t + (1 + \alpha)(\mu_{t-1} + b_{t-1}) \\ \sigma_t &= (1 - \alpha)(\sigma_{t-1} + \alpha(T_t - \mu_{t-1} - b_{t-1})(T_t - \mu_{t-1})) \\ b_t &= \beta(T_t - x_{t-1}) + (1 - \beta)b_{t-1} \end{aligned} \quad (8)$$

Here,  $b_t$  is the smoothed trend at time  $t$ ,  $\alpha$  is a smoothing factor and  $\beta$  a trend smoothing factor. Then the perceptible cue threshold is computed as

$$threshold_t = \mu_{t-1} + \gamma \sigma_{t-1}, \quad (9)$$

where  $\gamma$  is a tunable parameter. We use  $\gamma = 3$  in the work. The thresholding process filters out TE trends, which represent the accumulated information during continuous information exchanges. TE regions that meets the requirement  $TE_t > 0$  and  $TE_t > threshold_t$  indicate perceptible social cues. In order to further remove the effect of the trend, we apply a 1st order high-pass filter to the TE before cue detection. Since human reaction time to visual cues approximately ranges from 0.18-0.9s without distractions [12][18], the critical frequency of the high-pass filter is set to 1Hz in this

work. In the current framework, the selection of  $\alpha$  and  $\beta$  is related to the time constant of DES, which depends on the expected cue period. The time constants for smoothing  $\tau_\alpha$  and trend smoothing  $\tau_\beta$  are calculated as follows:

$$\tau_\alpha = -\frac{\Delta T}{\ln(1-\alpha)} \quad \tau_\beta = -\frac{\Delta T}{\ln(1-\beta)}, \quad (10)$$

where  $\Delta T$  is the sampling time interval of the data. A basic principle is that  $\tau_\alpha$  should be roughly the length of one cue in the setting, and  $\tau_\beta$  should be generally shorter than  $\tau_\alpha$ .

## 4 EXPERIMENTS

In this section, we first show that TE can be used to quantify information transfer in social interactions. Then we validate the framework in two settings to demonstrate its ability to detect perceptible social cues both temporally and spatially.

### 4.1 TE for social interactions - Group-Joining

To appraise the ability of TE to study social interactions, we first study a group-joining activity. The CongreG8 dataset [46] contains 380 full-body motion trials of free-standing conversational groups of three humans and a newcomer who approaches the groups with the intent of joining them. Four participants play a game called *Who's the Spy*. The game contains three group members and one adjudicator (the newcomer). Three group members gather at the centre of the arena, two are given cards with the same item, the third is given a card with a different item. Group members take turns describing the item on their card. Meanwhile, the adjudicator walks around the arena listening to the group members' discussion. Once the adjudicator has identified a spy, they join the group and point out the spy. The interactions are grouped into two categories. When the adjudicator approaches the group and the group members accommodate the adjudicator (e.g., a group member moves to make space), the trial is labelled as *Welcome*. Alternatively, if the group members stand still and ignore the adjudicator, the sample is labelled as *Ignorance*. Due to the high level of freedom within this experiment, participants' behaviours are largely different between each trial resulting in large variance in the sampled data.

We apply the general workflow (3.2) to the *Welcome* and *Ignorance* scenarios separately to compute the TE for each trial for two directions (from the adjudicator to the group members, and the reverse) and study the peak values of TE. We set the position of adjudicator or group members as target features depending on the direction of transfer being investigated. We consider the group members as a whole, so a single target/source is defined for all the group members. A re-sampling rate of 10Hz and history length of 10 with a unit time step of 0.1s for the embedding are applied. Since this is a relatively more complex dataset, a Gaussian Emission Variational Autoencoder (VAE) [22] with a two-hidden-layer encoder, a two-hidden-layer decoder and a latent space with 8 nodes is used to model the distributions over targets.

In this experiment, the adjudicator leads the group-joining event. Depending on the scenario, the responses from the group members are different. For the *Welcome* scenarios, group members should respond more to the perceptible social cues from the adjudicator when compared to the *Ignorance* scenarios. Therefore, we expect to see a difference in peak TE from the adjudicator to the group

members because we assume that the peak TE occurs during the joining process when the majority of information transfer from the adjudicator to the group happens. Therefore, we hypothesize that: The peak values of TE from the adjudicator to the group members should have large differences for the *Welcome* and *Ignorance* scenarios, when compared with peak TE in the other direction (group members to adjudicator).

To test this hypothesis, we calculate the peak TE values of each trial for both scenarios and both information transfer directions. Then we conduct a two-sided t-test between the two scenarios for both directions. This test (Table 1) shows a significant difference between *Ignorance* and *Welcome* settings in the adjudicator to group member direction, but not from group members to the adjudicator, supporting our hypothesis. This also supports our proposal that TE can correctly quantify information transfer in social settings.

**Table 1: Peak transfer entropy analysis t-test results.**

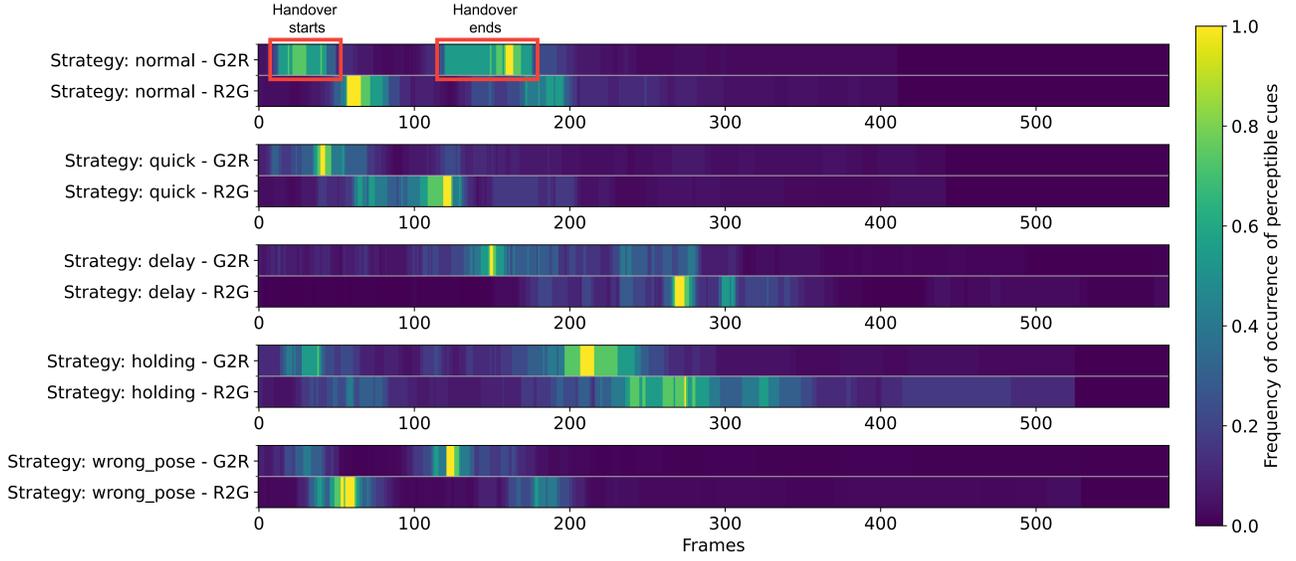
<i>Information transfer direction</i>	<i>p-value</i>
Adjudicator to Group	<b>0.0008</b>
Group to Adjudicator	0.9246

### 4.2 Temporal cue detection - Human-Human Handover

We test the framework's ability to identify *when* perceptible social cues occur in a human-human handover scenario. We use a dataset [8] comprising over 1000 recordings collected from 18 right-handed volunteers performing human-human handovers. 6-axis inertial data was acquired from two smartwatches on participants' wrists. During their single-blind experiments, a volunteer (receiver) and an experimenter (giver) form a pair. The two participants start from the diagonal corners of a square experimental area and walk towards the centre of the square to perform handovers. Several strategies are applied during the handovers:

- *Normal*: experimenter gives the object in a normal fashion.
- *Quick*: experimenter moves their arm faster.
- *Delay*: once the volunteer initiates the transferring gesture, the experimenter keeps their arm still ( $\approx 2s$ ) before reaching towards the volunteer to give the object.
- *Holding*: experimenter holds the object in place ( $\approx 2s$ ) after both persons have touched it, i.e., they do not release the object once the volunteer has grasped it.
- *Wrong pose*: as the volunteer initiates the transfer gesture, the experimenter unexpectedly moves their arm towards the volunteer's left shoulder, an unnatural pose for right-handed volunteers.

The original sampling rate of the dataset is approximately 7Hz, which we interpolated to a frame rate of 115Hz to increase the number of window samples to analyse. We used wrist angular velocity data from experiments where the same object, a ball, is used. In this experiment, the magnitude of the triaxial angular velocity is the only feature analysed for perceptible social cues. We use a history length of 4 with a unit time step of 0.14s for the Taken's delay embeddings. These numbers were selected empirically but could be chosen based on model predictive quality. The average human reaction time to detect visual stimuli is approximately 0.18-0.20s



**Figure 3: Frequency of occurrence of perceptible cues in the handover task (G - Giver; R - Receiver). The colour indicates the frequency of the occurrence of perceptible cues at each time frame. Cues from the giver always occur before the cues from the receiver. Two high cue count regions are usually identified in each graph indicating the start and end cues of handover respectively (An example is highlighted for the normal scenario). The gap between high cue count regions shows the delay between cues, and the concentration of high cue count regions shows the relative variance in time of the cue. Each graph matches with their corresponding handover strategy.**

[40], so 0.4s is a reasonable window to capture human reactions. For cue detection, we set the  $\alpha = 0.005$  and  $\beta = 0.01$ , which gives us  $\tau_\alpha = 1.7s$  and  $\tau_\beta = 0.9s$ . We expect that perceptible cues occur when the handover starts and ends with the giver initiating the handover and the receiver reaching in response to their cues [32].

To demonstrate the ability of the framework, we detect the perceptible cue regions for all the trials using the proposed framework and generate plots that visualise the frequency of occurrence of perceptible cues at each frame for different strategies and information transfer directions. The results are shown in Fig. 3. We trim the data so the frame consistently starts when the distance between the two participants is 1m. In a regular handover scenario, two main social cues should occur at the start and end of a handover respectively [32]. We also expect the cues from a giver to be followed by the cues from the receiver. This is visible in the *normal* strategy result, with two high cue count regions in both directions. The first and second regions correspond to the start and end of the handover respectively. We can see that high cue count regions in the receiver-to-giver (R2G) direction follow the giver-to-receiver (G2R) direction. Using the *normal* strategy as a baseline for comparison, we can also analyse the other scenarios. In the *quick* scenario, due to the faster movements, the cues occur and end faster when compared to the *normal* strategy; i.e, the spacing between cue counts is comparatively shorter. The fast movement of the giver could also confuse the receiver, resulting in a wider spread of the receiver’s reaction cue. In the *delay* scenario, the participant was not given instructions on how long to delay during

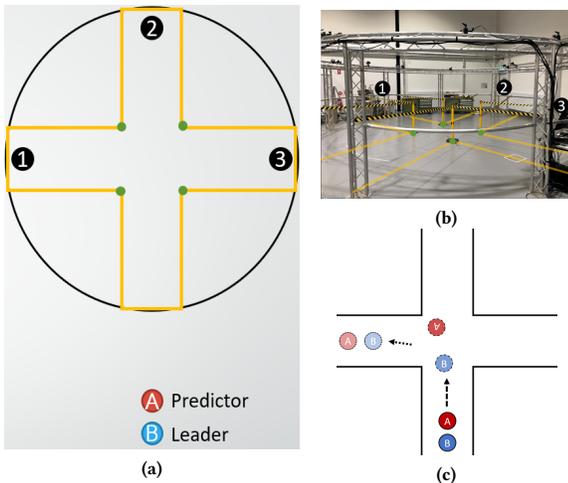
the experiment, thus the time of the starting cue of the handover is widely distributed. We observe a large delay before the first high cue count region appears, and the cues are more widely distributed. For the *holding* strategy, the time gap between the starting and ending cue should be longer compared to *normal*. This larger gap is observed between the first and second high cue count regions in both directions. For the *wrong pose* scenario, the wrong pose of the giver could confuse the receiver, resulting in weaker responses from the receiver. Therefore, the plots appear to be similar to the *normal* scenario. However, we observe the first high cue count region in the G2R direction has a lower cue occurrence frequency compared to *normal*. This analysis demonstrates that the proposed framework can identify *when* perceptible social cues occur. Additionally, interesting behaviour patterns can be extracted from the TE analysis, which can be useful for downstream anthropomorphic and personalized robot interaction design.

### 4.3 Spatial cue detection - Person-Following

We next apply the proposed framework to *spatially* analyse perceptible social cue transfer in a socially-aware navigation scenario requiring substantial non-verbal communication.

To maximise the exchange of perceptible social cues, we designed a leader-predictor front-following task simulating a person-following robot scenario (Fig. 4c) in a simulated intersection assembled using retractable barriers in a Vicon motion capture arena (Fig. 4b). The arena has a diameter of approximately 5.0m, and the

junction size is a 1.6m×1.6m square to permit comfortable side-by-side walking [9]. The scenario involves two human participants: the predictor A and the leader B. At the beginning of each trial, B is asked to secretly select one of the destinations shown in Fig. 4a. A is asked to actively stay in front of B and attempt to reach the unknown destination in advance of B. The starting position and formation are illustrated in Fig. 4a. This design avoids potential bias where the initial pose of participants might influence their preference after entering the intersection. When each trial begins, B walks naturally towards the selected destination, treating A as an unknown pedestrian (i.e., keeping a comfortable distance, no direct communication, but still following general social norms such as adjusting speed to avoid collision). After entering the intersection (the black circle in Fig. 4a), A can move freely, provided they do not interfere with B’s progress. Direct communication channels such as speaking or hand signals are forbidden, forcing participants to communicate via more subtle motion cues. This ‘front-following’ task involves high levels of observation, prediction and collaboration, providing an interesting setting to analyse the use of social cues. Seven pairs (14 individuals aged 19 to 31, 3 females, 11 males) were recruited from the Monash University Clayton campus. Participants provided consent at the start of the experiment <sup>4</sup>.



**Figure 4: Person-following task design. (a) An illustration of the starting position and formation. (b) A side view of the experimental arena. (c) An example leader-predictor front-following task (turning left scenario).**

The Vicon system provides a default frame rate of around 200 Hz. Vicon markers on a cap and a belt are used to track A and B’s head and hip pose, with the centre position and the head and hip orientation recorded for each participant.<sup>5</sup> Head and hip velocity are post-calculated using the positional information. Orientation and velocity vectors are projected onto the *xy* plane and then normalised. To apply the proposed framework, we take position, velocity, head and hip orientation as potential sources of cues and set each participant’s position as the target. A variational autoencoder

model is fit to the target using embeddings with a history length of 4 and a unit time step of 0.1s. For cue detection, we set the  $\alpha = 0.01$  and  $\beta = 0.05$ , which gives us time constants  $\tau_\alpha = 0.5s$  and  $\tau_\beta = 0.1s$ . Compared to the handover task, the cues in this experiment can be quicker, subtler and noisier. Our hypothesis is that perceptible cues appear in locations where a decision needs to be made or communicated about travel direction.

Similar to the handover experiment, we detect perceptible cue regions for all the trials using the proposed framework and generate plots that visualise the perceptible cue frequency at each spatial location in a 2-dimensional grid space representing the experiment arena. The results are shown in Fig. 5. We observe that for the leader to predictor (B2A) direction, most high-frequency cue regions appear in the region leading to the intersection. For the predictor to leader (A2B) direction, high-frequency cue regions often appear inside or slightly before the intersection. This outcome is consistent with the leader’s role in the front-following task, who first transfers cues to the predictor (B2A) prior to entering the intersection. This communication is usually followed by cues confirming the travel direction sent by the predictor (A2B) before leaving the intersection. This pattern of cue communication in physical space is clearly shown in Fig. 5. Although in theory, we would expect the perceptible cue regions to be similar for tuning left and right scenarios, the results show that they have noticeable differences. Most predictors (A) preferred to stay in the top left corner of the intersection while waiting and observing the leader (B) during the experiment <sup>6</sup>. This matches with post-experiment survey results that showed 42.9% of the predictors usually prefer to walk on the left side of the road or a corridor (28.6% right and 28.6% no preference). Therefore, physical space occupancy could be an influential factor for information transfer. This highlights the potential of the framework to capture information about cultural navigation biases.

The analysis above shows the proposed framework can identify *where* perceptible social cues occur.

## 5 DISCUSSION

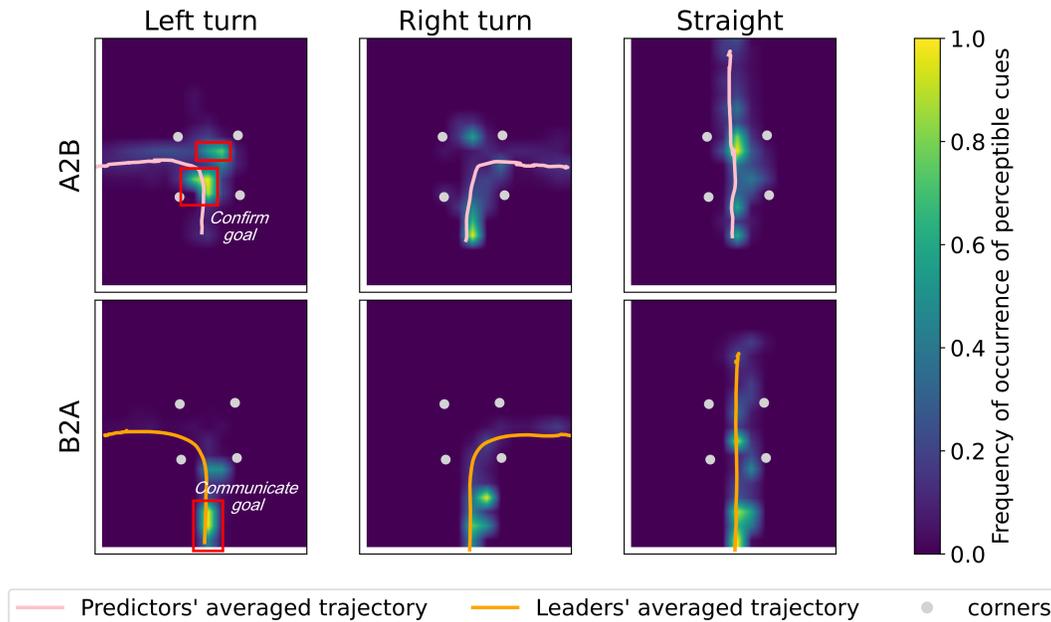
Our experiments have demonstrated the capability of the proposed framework. In addition to detecting perceptible social cues, the proposed framework can be potentially used to conduct further analysis. For instance, it has good scalability in the sense that it allows us to analyse detailed cues related to individual motion features, or combined cues if we consider all motions together. This scalability can be used to zero in on the source of cues embedded in motion. In addition, as shown in 4.1, the magnitude of TE potentially represents the strength of information transfer. This means it is possible to compare the strength of each cue in a setting.

While our experiments have focused on human interactions, we believe the proposed framework is broadly applicable in the field of HRI. Studying human-human cues can help robot designers identify suitable sources of information or cues for more natural human-robot interaction. The proposed framework could also be used to analyse social cues or information transfer during collaboration between humans and robots, for example in handover between humans and manipulators. In the context of socially-aware navigation,

<sup>4</sup>Approved by the Monash University ethics committee. Project ID: 33090

<sup>5</sup>The data is publicly available at <https://doi.org/10.26180/24719034.v1>

<sup>6</sup>This experiment is done in Australia where the norm is to keep left on sidewalks.



**Figure 5: Frequency of occurrence (cue count) of perceptible cues in the front-following task (A - Predictor; B - Leader). The colour indicates the frequency of the occurrence of perceptible cues at a physical location. Cues sent from the leader to the predictor to communicate the goal usually occur further away from the intersection, while cues sent from the predictor to the leader to confirm direction usually occur closer to or inside the intersection. High cue count regions are annotated in the left turn scenario as an example. This pattern in physical space is consistent for the three scenarios.**

we could design navigation algorithms to control the strength of information transfer. For example, we could intentionally reduce the influence on pedestrians during socially-aware navigation by minimising TE, thus minimising the influence of a robot on pedestrians. A robot could also take actions to increase information transfer to convey more information and potentially influence pedestrian motion if needed. This approach is potentially more flexible, adaptive and inclusive than current methods [42][28][20], which are designed based on Hall’s proxemics theory and the concept of social zones [16] that has been criticised for not considering the diverse range of social norms demonstrated by humans [34].

## 6 CONCLUSIONS AND FUTURE WORK

In this paper, we propose a framework for analysing perceptible social cue information transfer using transfer entropy. We have used a group-joining experiment to show the ability of TE to analyse information exchange during social interactions. We have applied the proposed framework to two unique settings, namely: object-handover and person-front-following, and demonstrated its capability of identifying the temporal and spatial occurrence of perceptible social cues. The proposed framework can be used for analysing cue information transfer of human-human interactions to identify cues of potential interest to robot designers, but could also be applied to analyse social cue transfer during human-robot collaborations.

Extending the proposed framework to real-time human-robot collaboration is a particularly exciting area of future work. This would require technical advancements and modelling choices to allow for distributions to be learned online in new settings, but this

could be simplified if pre-trained distributions are obtained from previously observed settings. Further user studies are needed to determine if similar social cues can be elicited or generated when a robot is a participant in an interaction, and these studies are an exciting next step for us. We plan to assess the capability of the proposed framework in less controlled settings and compare it with other existing cue detection methods.

## ACKNOWLEDGMENTS

We are grateful to Dr. Wesley Chan for his assistance at the early stage of this research and Mr. Brandon Johns for technical support with the Vicon system.

## REFERENCES

- [1] Seung Ki Baek, Woo-Sung Jung, Okyu Kwon, and Hie-Tae Moon. 2005. Transfer Entropy Analysis of the Stock Market. In *arXiv*. arXiv, arXiv. <https://doi.org/10.48550/arXiv.physics/0509014> arXiv:physics/0509014
- [2] Tobias Baur, Ionut Damian, Florian Lingensfelder, Johannes Wagner, and Elisabeth André. 2013. NovA: Automated Analysis of Nonverbal Signals in Social Interactions. In *Human Behavior Understanding*, Albert Ali Salah, Hayley Hung, Oya Aran, and Hatice Gunes (Eds.). Springer International Publishing, Cham, 160–171.
- [3] Erik Berger, David Müller, David Vogt, Bernhard Jung, and Heni Ben Amor. 2014. Transfer entropy for feature extraction in physical human-robot interaction: Detecting perturbations from low-cost sensors. In *2014 IEEE-RAS International Conference on Humanoid Robots* (Madrid, Spain). IEEE, New York, NY, USA, 829–834. <https://doi.org/10.1109/HUMANOIDS.2014.7041459>
- [4] Terry Bossomaier, Lionel Barnett, Michael Harré, and Joseph T. Lizier. 2016. Transfer Entropy. In *An Introduction to Transfer Entropy: Information Flow in Complex Systems*. Springer International Publishing, Cham, 65–95. [https://doi.org/10.1007/978-3-319-43222-9\\_4](https://doi.org/10.1007/978-3-319-43222-9_4)
- [5] Konstantinos Bousmalis, Marc Mehu, and Maja Pantic. 2013-02-01. Towards the automatic detection of spontaneous agreement and disagreement based on

- nonverbal behaviour: A survey of related cues, databases, and tools. *Image and Vision Computing* 31, 2 (2013-02-01), 203–221. <https://doi.org/10.1016/j.imavis.2012.07.003>
- [6] Alexandra Bremers, Alexandria Pabst, Maria Teresa Parreira, and Wendy Ju. 2023. Using Social Cues to Recognize Task Failures for HRI: A Review of Current Research and Future Directions. In *arXiv*. arXiv, arXiv. <https://doi.org/10.48550/arXiv.2301.11972> [cs]
- [7] Baptiste Busch, Jonathan Grizou, Manuel Lopes, and Freek Stulp. 2017. Learning Legible Motion from Human–Robot Interactions. *Int J of Soc Robotics* 9, 5 (2017), 765–779. <https://doi.org/10.1007/s12369-017-0400-4>
- [8] Alessandro Carfi, Francesco Foglino, Barbara Bruno, and Fulvio Mastrogiovanni. 2019. A multi-sensor dataset of human-human handover. *Data in Brief* 22 (2019), 109–117. <https://doi.org/10.1016/j.dib.2018.11.110>
- [9] Minnie Doss. 2023. Two Lane Corridor Dimensions & Drawings | Dimensions.com.
- [10] Anca D. Dragan, Kenton C.T. Lee, and Siddhartha S. Srinivasa. 2013. Legibility and predictability of robot motion. In *2013 8th ACM/IEEE International Conference on Human-Robot Interaction (HRI)* (Tokyo, Japan). IEEE, New York, NY, USA, 301–308. <https://doi.org/10.1109/HRI.2013.6483603>
- [11] Arturo Escobedo, Anne Spalanzani, and Christian Laugier. 2014. Using social cues to estimate possible destinations when driving a robotic wheelchair. In *2014 IEEE/RSJ International Conference on Intelligent Robots and Systems* (Chicago, IL, USA). IEEE, New York, NY, USA, 3299–3304. <https://doi.org/10.1109/IROS.2014.6943021>
- [12] Thomas Fugger, Bryan Randles, Anthony Stein, William Whiting, and Brian Gallagher. 2000-01-01. Analysis of Pedestrian Gait and Perception-Reaction at Signal-Controlled Crosswalk Intersections. *Transportation Research Record* 1705 (2000-01-01), 20–25. <https://doi.org/10.3141/1705-04>
- [13] Everette S. Gardner Jr. 1985. Exponential smoothing: The state of the art. *Journal of Forecasting* 4, 1 (1985), 1–28. <https://doi.org/10.1002/for.3980040103>
- [14] C. W. J. Granger. 1969. Investigating Causal Relations by Econometric Models and Cross-spectral Methods. *Econometrica* 37, 3 (1969), 424–438. <https://doi.org/10.2307/1912791>
- [15] David M. Green and John A. Swets. 1966. *Signal detection theory and psychophysics*. John Wiley, Oxford, England.
- [16] Edward T. Hall. 1966. *The hidden dimension: man's use of space in public and private*. The Bodley Head Ltd, London.
- [17] Jiayi He and Pengjian Shang. 2017. Comparison of transfer entropy methods for financial time series. *Physica A: Statistical Mechanics and its Applications* 482 (2017), 772–785. <https://doi.org/10.1016/j.physa.2017.04.089>
- [18] Frouke Hermens, Markus Bindemann, and A. Mike Burton. 2017-01-01. Responding to social and symbolic extrafoveal cues: cue shape trumps biological relevance. *Psychological Research* 81, 1 (2017-01-01), 24–42. <https://doi.org/10.1007/s00426-015-0733-2>
- [19] Nicholas J. Hetherington, Elizabeth A. Croft, and H. F. Machiel Van der Loos. 2021. Hey Robot, Which Way Are You Going? Nonverbal Motion Legibility Cues for Human-Robot Spatial Interaction. *IEEE Robot. Autom. Lett.* 6, 3 (2021), 5010–5015. <https://doi.org/10.1109/LRA.2021.3068708> arXiv:2104.02275
- [20] Jwu-Sheng Hu, Jyun-Ji Wang, and Daniel Minare Ho. 2014-04. Design of Sensing System and Anticipative Behavior for Human Following of Mobile Robots. *IEEE Trans. on Industrial Electronics* 61, 4 (2014-04), 1916–1927. <https://doi.org/10.1109/TIE.2013.2262758>
- [21] Natasha Jaques, Angeliki Lazaridou, Edward Hughes, Caglar Gulcehre, Pedro Ortega, Dj Strouse, Joel Z. Leibo, and Nando De Freitas. 2019. Social Influence as Intrinsic Motivation for Multi-Agent Deep Reinforcement Learning. In *Proceedings of the 36th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 97)*, Kamalika Chaudhuri and Ruslan Salakhutdinov (Eds.). PMLR, CA, USA, 3040–3049.
- [22] Diederik P. Kingma and Max Welling. 2022. Auto-Encoding Variational Bayes. In *arXiv*. arXiv, arXiv. <https://doi.org/10.48550/arXiv.1312.6114> arXiv:1312.6114 [cs, stat]
- [23] Felix Kistler, Birgit Endrass, Ionut Damian, Chi Tai Dang, and Elisabeth André. 2012-07-01. Natural interaction with culturally adaptive virtual characters. *Journal on Multimodal User Interfaces* 6, 1 (2012-07-01), 39–47. <https://doi.org/10.1007/s12193-011-0087-z>
- [24] A.S. Klyubin, D. Polani, and C.L. Nehaniv. 2005. Empowerment: a universal agent-centric measure of control. In *2005 IEEE Congress on Evolutionary Computation* (Edinburgh, UK), Vol. 1. IEEE, New York, NY, USA, 128–135 Vol.1. <https://doi.org/10.1109/CEC.2005.1554676>
- [25] Joel Z. Leibo, Vinicius Zambaldi, Marc Lanctot, Janusz Marecki, and Thore Graepel. 2017. Multi-Agent Reinforcement Learning in Sequential Social Dilemmas. In *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems* (São Paulo, Brazil) (AAMAS '17). International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 464–473.
- [26] Christina Lichtenthaler, Tamara Lorenzy, and Alexandra Kirsch. 2012. Influence of legibility on perceived safety in a virtual human-robot path crossing task. In *2012 IEEE RO-MAN: The 21st IEEE International Symposium on Robot and Human Interactive Communication* (Paris, France). IEEE, New York, NY, USA, 676–681. <https://doi.org/10.1109/ROMAN.2012.6343829>
- [27] Joseph T. Lizier. 2014. JIDT: An Information-Theoretic Toolkit for Studying the Dynamics of Complex Systems. *Frontiers in Robotics and AI* 1 (2014). <https://doi.org/10.3389/frobt.2014.00011> arXiv:1408.3270 [nlin, physics:physics]
- [28] Weiming Mi, Xiaozhe Wang, Ping Ren, and Chenyu Hou. 2016. A System for an Anticipative Front Human Following Robot. In *Proceedings of the International Conference on Artificial Intelligence and Robotics and the International Conference on Automation, Control and Robotics Engineering* (Kitakyushu, Japan) (ICAIR-CACRE '16). Association for Computing Machinery, New York, NY, USA, Article 4, 6 pages. <https://doi.org/10.1145/2952744.2952748>
- [29] Shakir Mohamed and Danilo J. Rezende. 2015. Variational Information Maximisation for Intrinsically Motivated Reinforcement Learning. In *Proceedings of the 28th International Conference on Neural Information Processing Systems - Volume 2* (Montreal, Canada) (NIPS'15). MIT Press, Cambridge, MA, USA, 2125–2133.
- [30] George P. Moustris and Costas S. Tzafestas. 2016. Intention-based front-following control for an intelligent robotic rollator in indoor environments. In *2016 IEEE Symposium Series on Computational Intelligence (SSCI)* (Athens, Greece). IEEE, New York, NY, USA, 1–7. <https://doi.org/10.1109/SSCI.2016.7850067>
- [31] N. Orange and N. Abaid. 2015-12-01. A transfer entropy analysis of leader-follower interactions in flying bats. *The European Physical Journal Special Topics* 224, 17 (2015-12-01), 3279–3293. <https://doi.org/10.1140/epjst/e2015-50235-9>
- [32] Valerio Ortenzi, Akansel Cosgun, Tommaso Pardi, Wesley P. Chan, Elizabeth Croft, and Dana Kulić. 2021-12. Object Handovers: A Review for Robotics. *IEEE Transactions on Robotics* 37, 6 (2021-12), 1855–1873. <https://doi.org/10.1109/TRO.2021.3075365> Conference Name: IEEE Transactions on Robotics.
- [33] Maurizio Porfiri. 2018. Inferring causal relationships in zebrafish-robot interactions through transfer entropy: a small lure to catch a big fish. *AB&C* 5, 4 (2018), 341–367. <https://doi.org/10.26451/abc.05.04.03.2018>
- [34] J. Rios-Martinez, A. Spalanzani, and C. Laugier. 2015. From Proxemics Theory to Socially-Aware Navigation: A Survey. *Int J of Soc Robotics* 7, 2 (2015), 137–153. <https://doi.org/10.1007/s12369-014-0251-1>
- [35] Thomas Schreiber. 2000. Measuring Information Transfer. *Phys. Rev. Lett.* 85, 2 (2000), 461–464. <https://doi.org/10.1103/PhysRevLett.85.461>
- [36] Irena Shaffer and Nicole Abaid. 2020. Transfer Entropy Analysis of Interactions between Bats Using Position and Echolocation Data. *Entropy* 22, 10 (2020), 1176. <https://doi.org/10.3390/e22101176>
- [37] C. E. Shannon. 1948. A Mathematical Theory of Communication. *Bell System Technical Jour.* 27, 3 (1948), 379–423. <https://doi.org/10.1002/j.1538-7305.1948.tb01338.x>
- [38] Hidenobu Sumioka, Yuichiro Yoshikawa, and Minoru Asada. 2007. Causality detected by transfer entropy leads acquisition of joint attention. In *2007 IEEE 6th Int. Conf. on Development and Learning* (London, UK). IEEE, New York, NY, USA, 264–269. <https://doi.org/10.1109/DEVLRN.2007.4354069>
- [39] Floris Takens. 1981. Detecting strange attractors in turbulence. In *Dynamical Systems and Turbulence, Warwick 1980*, David Rand and Lai-Sang Yung (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 366–381.
- [40] P. D. Thompson, J. G. Colebatch, P. Brown, J. C. Rothwell, B. L. Day, J. A. Obeso, and C. D. Marsden. 1992. Voluntary stimulus-sensitive jerks and jumps mimicking myoclonus or pathological startle syndromes. *Movement Disorders* 7, 3 (1992), 257–262. <https://doi.org/10.1002/mds.870070312>
- [41] Razali Tomari, Yoshinori Kobayashi, and Yoshinori Kuno. 2014. Analysis of Socially Acceptable Smart Wheelchair Navigation Based on Head Cue Information. *Procedia Computer Science* 42 (2014), 198–205. <https://doi.org/10.1016/j.procs.2014.11.052>
- [42] Soren Tranberg Hansen, Mikael Svenstrup, Hans Jorgen Andersen, and Thomas Bak. 2009. Adaptive human aware navigation based on motion pattern analysis. In *RO-MAN 2009* (Toyama, Japan). IEEE, New York, NY, USA, 927–932. <https://doi.org/10.1109/ROMAN.2009.5326212>
- [43] Jacqueline Urakami and Katie Seaborn. 2023-03-15. Nonverbal Cues in Human–Robot Interaction: A Communication Studies Perspective. *ACM Transactions on Human-Robot Interaction* 12, 2 (2023-03-15), 22:1–22:21. <https://doi.org/10.1145/3570169>
- [44] Alessandro Vinciarelli. 2009. Social Computers for the Social Animal: State-of-the-Art and Future Perspectives of Social Signal Processing. In *User Modeling, Adaptation, and Personalization*, Geert-Jan Houben, Gord McCalla, Fabio Pianesi, and Massimo Zancanaro (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 1–1. [https://doi.org/10.1007/978-3-642-02247-0\\_1](https://doi.org/10.1007/978-3-642-02247-0_1)
- [45] Wei Xie, Dongli Gao, and Eric Waiming Lee. 2022. Detecting Undeclared-Leader-Follower Structure in Pedestrian Evacuation Using Transfer Entropy. *IEEE Trans. on Intelligent Transportation Systems* 23, 10 (2022), 17644–17653. <https://doi.org/10.1109/TITS.2022.3161813>
- [46] Fangkai Yang, Yuan Gao, Ruiyang Ma, Sahba Zojaji, Ginevra Castellano, and Christopher Peters. 2021. A dataset of human and robot approach behaviors into small free-standing conversational groups. *PLOS ONE* 16, 2 (2021), e0247364. <https://doi.org/10.1371/journal.pone.0247364>