




Contents lists available at ScienceDirect

Robotics and Autonomous Systems

journal homepage: www.elsevier.com/locate/robot

Self-supervised prediction of the intention to interact with a service robot

Gabriele Abbate^{a,*}, Alessandro Giusti^a, Viktor Schmuck^b, Oya Celiktutan^b, Antonio Paolillo^a

^a Dalle Molle Institute for Artificial Intelligence (IDSIA), USI-SUPSI, Lugano, Switzerland

^b Department of Engineering, King's College London, United Kingdom

ARTICLE INFO

Keywords:

Self-supervised learning
Human-robot interaction
Social robotics

ABSTRACT

A service robot can provide a smoother interaction experience if it has the ability to *proactively* detect whether a nearby user intends to interact, in order to adapt its behavior e.g. by explicitly showing that it is available to provide a service. In this work, we propose a learning-based approach to predict the probability that a human user will interact with a robot before the interaction actually begins; the approach is *self-supervised* because after each encounter with a human, the robot can automatically label it depending on whether it resulted in an interaction or not. We explore different classification approaches, using different sets of features considering the pose and the motion of the user. We validate and deploy the approach in three scenarios. The first collects 3442 natural sequences (both interacting and non-interacting) representing employees in an office break area: a real-world, challenging setting, where we consider a coffee machine in place of a service robot. The other two scenarios represent researchers interacting with service robots (200 and 72 sequences, respectively). Results show that, even in challenging real-world settings, our approach can learn without external supervision, and can achieve accurate classification (i.e. AUROC greater than 0.9) of the user's intention to interact with an advance of more than 3 s before the interaction actually occurs.

1. Introduction

Many emerging applications of robots have the potential of assisting humans in everyday life tasks or automating jobs in the future [1]. Examples include social robots offering assistance at receptions [2], in hospitality sectors [3] or at home [4]; navigation guidance in public spaces [5] or personal care [6]; and object delivery [7]. In such situations, robots should automatically understand the human intention to interact well before the interaction starts to be more proactive and offer relevant services.

The very initial phase of these interactions plays an important role in establishing an effective Human-Robot Interaction (HRI), in which the user first sees the robot and decides to approach and engage with it. When users are unfamiliar with the situation, e.g., because they enter a new environment and are unsure about the right action to take, the robot's behavior is crucial to determine if this approach phase yields a successful interaction and a good user experience [8].

Consider the everyday-life scenario of a skilled human receptionist operating in a bustling lobby. They can anticipate the arrival of a client detecting cues in the client's movement and body language well before they reach the reception desk. In these circumstances, the receptionist welcomes the client without being distracted by other nearby people who are not interested in interacting. This behavior makes it clear to

the client that the receptionist is indeed available for interaction and is the right person to approach for assistance. Albeit it might just appear as a pleasant but superfluous detail for a client who already knows what to do, this subtle behavior of the receptionist can reassure novel users who might be intimidated or confused in an unfamiliar situation.

To enable the widespread deployment and social acceptance of robots in everyday life scenarios (see Fig. 1), they must develop similar skills, namely, anticipate and adapt to human intentions. Indeed, an effective service robot should have the following skills: (i) keeping track of nearby people; (ii) predicting when an approaching person intends to interact with it; and (iii) reacting accordingly. In this paper, we use off-the-shelf tools to solve the first point and focus our contribution on the second skill, which is general and mostly independent of the type of robot. Once the intention of the user has been detected, reaction strategies can be designed according to the specific robot hardware and sensory equipment.

Our primary contribution is a learning-based method that enables the robot to classify whether each tracked person intends to interact with it or not. As input, we use body motion cues that are provided by off-the-shelf video or RGB-D sensing subsystems. The probability that the person will interact is updated in real-time and can trigger a reaction of the robot when exceeding a threshold. Eventually, each tracked

* Corresponding author.

E-mail address: gabriele.abbate@supsi.ch (G. Abbate).

<https://doi.org/10.1016/j.robot.2023.104568>

Received 28 July 2023; Received in revised form 11 September 2023; Accepted 23 October 2023

Available online 31 October 2023

0921-8890/© 2023 The Author(s). Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).



Fig. 1. For a robot providing information in a corridor of a public building (top) or serving a chocolate treat to a passerby (bottom), it is crucial to proactively detect the human intention to interact even when the user is still at a distance, in order to adopt behaviors perceived as friendly, demonstrate availability to interact, and more efficiently offer the relevant services.

person either interacts with the robot, or leaves without interacting; the corresponding data sequences, considered with hindsight, provide additional training data that the robot collects without the need for manual labeling, or any other form of external supervision. Variations of this concept have been known and applied in various fields of robotics research since the mid-2000s [9–20], denoted with the term *self-supervised learning*, which highlights that the robot autonomously generates labeled data for the task of interest.

The remainder of the paper is organized as follows. After reviewing related work (Section 2), we describe our approach (Section 3) and its implementation (Section 4); experimental results are presented in Section 5. We finally derive our conclusions in Section 6, discussing future work directions.

2. Related work

Nonverbal communication cues [21], such as body motion and language, play a central role in HRI, from both users’ and robots’ perspective [22,23]. However, the perception of social nonverbal behaviors is a challenging task to solve in HRI [24], especially for the first phases of the interactions [8]. Nonetheless, it is important to be able to predict the intention to interact with the robot so that an effective reaction strategy can be well accommodated to the users’ needs. For example, human intention navigation is inferred using motion features [25]. In the context of collaborative tasks, the human intention is estimated from gaze and motion features in virtual reality [26] or analyzing the motion performed in front of a humanoid robot [27]. In these cited works, the intention of the human is intended to be related to the next action to take in the context of an ongoing activity. Similarly, other systems based on body motion cues are used to classify the social behavior of humans standing in the robot’s proximity [28,

29]. In our work, instead, we aim at predicting the human intent well before the interaction actually starts. The intention to interact based on gaze and body motion has also been proposed as a tool to evaluate the engagement of a user standing in front of a system at a fixed distance [30]. Our work focuses on a more general scenario since we want to predict the intention of any users free to move into social spaces. It is worth mentioning that a significant body of work bases the intention recognition only on gaze cues, as can be found in a recent review [31]. However, our work aims at predicting human intention from far distances, where the performances of gaze trackers are expected to decrease. Our work is more similar to approaches using multi-modal features, including body motion, to train a binary classifier predicting users’ intention to interact [32–34] or to assess the intensity of human engagement intention [35]. These works rely on hand-labeled datasets collected in controlled environments, which are expensive, and sometimes unfeasible, to acquire for each deployment scenario. In contrast, our approach is self-supervised, as discussed below.

In a standard supervised paradigm, one would need to collect large training datasets composed of a large number of tracks representing a given human in the robot’s vicinity, and manually provide labels assigning a class to every track depending on whether the human interacts with the robot or not. In contrast, our work relies on the robot’s ability to reconsider its experience in hindsight, and automatically assign a label to each recorded track, depending on whether it eventually resulted in an interaction with the robot or not. This is a form of *self-supervised robot learning*, that derives labels from data available to the robot only *after* the sample was observed; robots capable of self-supervised learning rely on data collected in previous experiences by their own sensors in order to self-generate meaningful supervision, a paradigm initially adopted in robotics for segmentation of traversable terrain [9–11], then applied to other tasks such as grasping [12–14] and long-range sensing for navigation [15–20]. It is worth noting that in the recent deep learning literature, the term “self-supervised” has a different meaning: it denotes the practice of using pretext tasks [36–38] for learning useful data representations [39] from large amounts of unlabeled data.

One of the advantages of Self-Supervised Learning (SSL) approaches is that they allow the system to continuously update its models with new training data acquired on the spot. This is especially valuable in our scenario, as the robot can use these data to learn human behavior cues that are specific to its deployment environment. A related but different field of research is *continual learning* [40], which provides methods to efficiently adapt models as new training data becomes available, without having to store the entire training dataset and avoiding the problem of catastrophic forgetting. In our work, we adopt a simpler approach: we store the entire dataset and retrain the model from scratch, without resorting to continual learning techniques.

3. Approach

3.1. Problem formulation

We consider a robot standing in an environment shared with humans, some of which might approach the robot in order to engage with it. The robot is equipped with sensors capable to detect and track people at least within a distance of 4 m, i.e. the robot’s *social space* [24,41], but possibly beyond. During normal operation, people routinely pass nearby the robot, entering and exiting the robot’s social space; occasionally, some users engage with the robot.

We define \mathcal{F}_r as a fixed frame centered on the robot. For each tracked person, the robot is capable to estimate the pose of their torso (\mathcal{F}_t) and head (\mathcal{F}_h) frames. In particular, we denote as $p_t \in \mathbb{R}^2$ and θ_t the planar position and orientation of \mathcal{F}_t w.r.t. \mathcal{F}_r , respectively. The distance of the person from the robot is $d = \|p_t\|$. Similarly, θ_h indicates the orientation of \mathcal{F}_h in \mathcal{F}_r around the vertical axis. Finally, the variable $v_t \in \mathbb{R}^2$ indicates the person’s linear velocity. Note that

the position and orientation of a person's torso w.r.t \mathcal{F}_r and its velocity are informative of their proxemics [23] and are also useful to determine which proxemic zone [24,41] they occupy. The head orientation is also indicative of the user's gaze and is expected to be informative of their intention.

We tackle the problem of predicting the intention of a person to interact with the robot, as soon as possible before the interaction begins. To this end, we make use of information captured about possible interacting people, elaborated by different classifier architectures, as described in the following.

3.2. Sensing and features

In our study, we make use of the proxemics, i.e. the analysis of motion cues of interacting users. More into detail, proxemics analyses the way a user uses or occupies the social space in order to infer useful information for the interaction [23]. Proxemics concepts are particularly suitable to our scope as (i) they are very representative of the intention to interact and (ii) the quantities that define them can be conveniently measured with state-of-the-art robotic sensors. In particular, the RGB-D sensor used for data collection is the Azure Kinect [42]. The SDK of this sensor provides the detection and tracking of the human skeletons appearing in its field of view. More into detail, each detected skeleton is given an ID and defined as a tree of frames along the kinematic structures of the user. From the spatial information of the skeleton, the motion of the user can be easily extracted and used for our intention prediction module. Such data is saved in an anonymous way, i.e., no RGB-D images are stored: only the metric information required by the classifier is logged.

For our analysis, we take into account different sets of features. First of all, we consider the distance or the orientation of the person's torso:

$$f_1 = d, \quad \text{and} \quad f_2 = (\sin \theta_t, \cos \theta_t)^T \in \mathbb{R}^2. \quad (1)$$

where, according to machine learning best practices [43], we encode the torso orientation into its sin and cos functions, to account for the fact that the feature is cyclical, and thus the representation of angle 0° should be close to 359° . The third set that we consider contains the torso position:

$$f_3 = p_t \in \mathbb{R}^2. \quad (2)$$

The fourth set of features gathers torso position and orientation together:

$$f_4 = (p_t^T, \sin \theta_t, \cos \theta_t)^T \in \mathbb{R}^4 \quad (3)$$

In the fifth set, we also include the orientation of the head:

$$f_5 = (p_t^T, \sin \theta_t, \cos \theta_t, \sin \theta_h, \cos \theta_h)^T \in \mathbb{R}^6 \quad (4)$$

and in the last one, we add the velocity of the torso as well:

$$f_6 = (p_t^T, \sin \theta_t, \cos \theta_t, \sin \theta_h, \cos \theta_h, v_t^T)^T \in \mathbb{R}^8. \quad (5)$$

The sets of features include the notions of proxemics at different levels. Their comparison allows for analyzing the contribution of the different proxemics elements in the prediction of the interaction.

3.3. Classification approach

To solve the problem, we train a binary classifier that takes as input a feature vector describing a tracked person at a given time and outputs the probability that the person will interact with the robot.

The classifier is trained on a dataset D composed of several *sequences*. A sequence represents a person tracked by the robot over time and is composed of multiple *samples* (one per timestep). The sequence begins when the person enters the robot's social space and is first seen by the sensor; it ends when the person either begins their interaction




Fig. 2. To collect data, the motion of people walking in a break area is monitored to predict their intention to interact with a coffee machine.

with the robot or exits the social space of the robot without interacting. The dataset is denoted as

$$D = \left\{ f_{i,j}, y_{i,j} \right\}_{i=1,j=1}^{N_j, S} \quad (6)$$

where f is the feature vector, and y the label; the subscripts i and j indicate the i th sample of the j th sequence, respectively; S is the number of sequences, whereas N_j is the number of samples contained in the j th sequence. For a given sequence j , all labels $y_{i,j}$ are the same: 0 if the person did not interact with the robot; 1 if they did.

Assuming that the robot has the ability to detect that the user has engaged in an interaction, the true label of a sequence becomes available as soon as the sequence ends. This enables the robot to grow its training dataset without external supervision and to iteratively improve the classifier performance in a self-supervised way.

From an implementation point of view, we investigate different classifiers: Logistic Classifier (LC), Random Forest (RF), and Multilayer Perceptron (MLP) using their implementation provided by the scikit-learn library [44]; and Long Short-Term Memory (LSTM) which is implemented using PyTorch [45]. The MLP is composed of 2 hidden layers with 30 neurons each, using sigmoid activations; the LSTM is composed of 2 long short-term memory [46] cells with a 10-dimensional hidden state each; both models have approximately 1500 trainable parameters.

It is worth noting that the input of the classifier is limited to the information related to a single subject, i.e., the one whose intention to interact is being classified, whereas it does not include information about other people. However, during both training and inference, the presence of multiple people is easily handled by our approach since each person is tracked and processed independently. Since the classifier is computationally light (i.e. on a standard laptop, it runs at 30 FPS, which is the sensor's maximum frame rate) and can be instantiated in parallel for several users, we actually handle multi-user prediction by tracking and classifying all the people appearing in the field of view of the sensor (see Section 5.2).

4. Experimental scenarios

We test our approach in different scenarios, presented in Table 1 and described in detail in the following.

4.1. Real-world interactions at a coffee break area

We collect a real-world, challenging dataset of human-machine interactions in which humans behave naturally. In particular, we consider a coffee machine placed in a break area neighboring a corridor of an

Table 1
Scenarios considered in our analysis.

| Scenario | Setting | User | Agent | Self labeling | Sequences number | Mode | Agent behavior |
|--------------|-------------|---------|----------------|----------------|------------------|--------------|----------------|
| Coffee break | In-the-wild | Unaware | Coffee machine | Distance based | High (3422) | Train & Test | Passive |
| Waiter robot | Controlled | Actor | Robo-master | Vision based | Medium (200) | Train & Test | Reactive |
| Info robot | Controlled | Actor | HSR-B robot | Touch based | Low (72) | Test | Passive |

office building (Fig. 2). During the day, many people pass through the corridor, some of them stop in the break area, and some others approach the machine to have a coffee. This scenario is interesting and convenient for our analysis as we can observe the spontaneous behavior of the users who plan to interact with the coffee machine, in a natural context with many challenging complications and distractors: other users hanging around chatting; users approaching the general area to reach a nearby tap or fridge; users queuing up to use the machine. In this scenario, we have collected 3422 unique sequences of tracked skeletons, accounting for more than 12 h of recorded data. Recorded users come from a heterogeneous sample of people, mainly employees and guests who have access to the break area. The users are informed about the presence of the sensor above the coffee machine. However, they are unaware of the scope of the data collection. In this way, we ensure that their behavior is as natural as possible. Non-sensitive data (i.e. only the skeletons of the users) are recorded.

In this specific scenario, sequences should be ideally labeled by considering when a user operates the machine, e.g. by pressing a button on it; similarly one might expect a service robot to easily determine when a user engages with it. However, in our case, we do not have access to the machine firmware and we cannot read its internal state. Therefore, we rely on the sensor used for data collection to automatically generate labels. To do so, we use the following distance-based heuristic: interaction is detected when a person stays very close (i.e. within a distance of 1 m) to the coffee machine for an uninterrupted period of 5 seconds; we assume that the interaction takes place at the end of this period; all samples, coming from the same sequence, in the preceding 10 seconds are labeled $y = 1$. We empirically verified that such criterion is very effective as a proxy to detect actual interactions, and we use it to automatically generate labels in this scenario.

4.2. Chocolate handover by a waiter robot

In the second scenario, we use a wheeled omnidirectional robot (DJI Robomaster EP [47]), placed on a table in the vicinity of the Azure Kinect [42] sensor (Fig. 1, bottom). The robot behaves as a waiter who serves chocolate treats to people passing by. During the data collection, the robot does not perform any motion. Data can be self-labeled using a simple vision-based approach based on image-based detections taken with the robot's onboard camera with which we can automatically detect whether users take the chocolate or not. The recorded data consists of 200 sequences of a single user performing the same number of interacting and non-interacting actions.

In the deployment phase, instead, we provide the robot with reactive behavior. If an interaction is predicted, the robot enacts a reaction by turning its LEDs on and orienting itself towards the user yielding the highest probability. At the same time, the robot extends its arm handing out a chocolate treat to the user: this acknowledges that the robot has seen the user and is available to interact. When no interaction is predicted, the robot gets back to its initial orientation, turns the LEDs off, and retracts its arm. Such behavior has been tested with users aware of the interactions in a controlled environment. The users of these tests were informed about the purpose of the experiments and gave their consent to participate in the data collection.

4.3. Information service robot

Finally, we propose a controlled evaluation setup with the Toyota Human Support Robot series B (HSR-B) [48] robot placed in a U-shaped corridor (see Fig. 1, top). The robot is equipped with the Azure Kinect sensor on its head, oriented horizontally w.r.t. the floor, at a height of about 1.3 m. We have collected evaluation data from the behavior of 12 participants. The participants who walk through the corridor, initially cannot see the robot, after the first curve notice it, and adjust their behavior according to their intention to interact. In this specific data collection setup, participants act as actors, i.e., they are informed of the presence of the robot. Furthermore, in half of the cases, they are told to pretend that they do not wish to or do not have time to interact with the robot. In this way, they provide non-interaction sequences for our dataset. In the other half of the cases, participants are instructed to walk to the robot when they see it and touch its head, which we considered the interaction trigger for this evaluation scenario. Each test participant produced 3 samples of not interacting with the robot, and 3 that recorded interaction, resulting in a total of 72 samples. The data collection protocol was approved by the Ethical Committee of King's College London, United Kingdom (Review reference: LRS/DP-22/23-35586).

5. Results

We report the experimental analysis carried out in each scenario described in Section 4. First, we perform offline experiments on the large and challenging coffee break dataset presented in Section 4.1, comparing different feature sets and classification approaches. Based on these experiment results, we then select the most promising combination of features and classifier for the experimental validation within the other two scenarios that involve actual robots. The presented results can be further qualitatively evaluated in the video accompanying the paper.

5.1. Offline experiment in the coffee break scenario

5.1.1. Sample-level performance

We compare different feature sets (Section 3.2) and classification approaches (Section 3.3) using the dataset collected in the coffee break scenario. We partition the set of the recorded sequences into 5 evenly-sized non-overlapping groups. Then, for each combination of feature set and classifier, we use a 5-fold cross-validation approach to compute predictions for all the samples in all the sequences. In particular, the samples in all the sequences of a given group are classified by a model trained on all the sequences belonging to the 4 remaining groups.

We then consider all samples from all sequences to compute performance metrics. In particular, we report the Area Under the ROC Curve (AUROC): a robust binary classification metric that does not depend on a choice of threshold, and ranges between 0.5 (for a non-informative classifier, e.g. one always reporting the majority class) and 1.0 (an ideal classifier). It can be interpreted as the probability that, taking a random sample from a person who did not interact, and a random sample from a person who eventually interacted, the classifier assigns to the former a lower score than the latter. When computed on all testing samples pooled together, all models score very high when using feature sets that include the distance-based information (see Table 2). The reason is that the person's distance from the device is a very strong cue of whether the person ends up interacting with it.

However, we aim to evaluate the ability of our approach to classify a person's intention to interact *independently* on their distance from the device. A more informative metric in our context is therefore the AUROC computed among samples that all lie approximately at the same distance; within this group of samples, the distance feature alone loses its discriminative ability. Thus, we partition all our testing samples in seven distance bins, determined in such a way to have an approximately

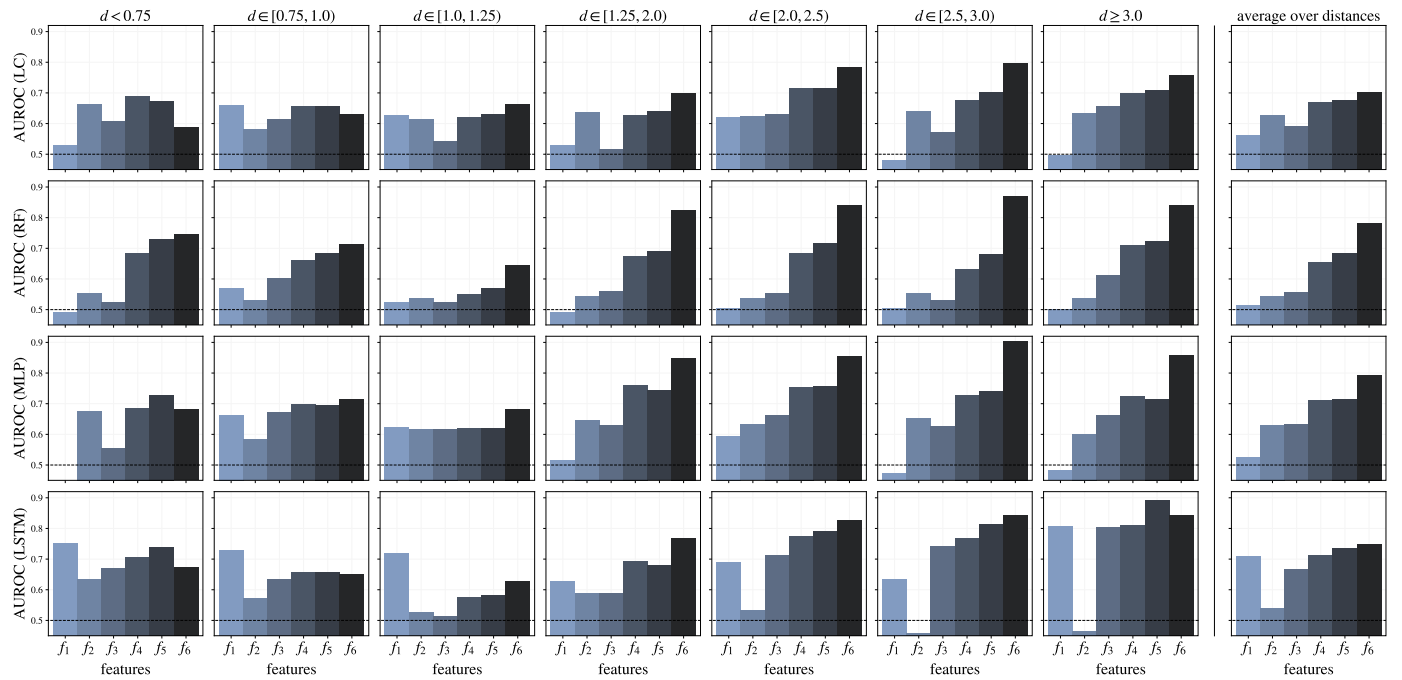


Fig. 3. Coffee break scenario: performance of the classifiers according to the AUROC metric for the different models (from top to bottom: LC, RF, MLP, and LSTM); tested in different ranges of distance (from left to right, ranging from below 0.75 m to above 3.5 m) and on average over all the distance ranges (last column); and using different sets of features (from f_1 to f_6 for each column of the histograms from left to right). The horizontal dotted line denotes the performance of a noninformative classifier (AUROC = 0.5).

Table 2

AUROC for different classifiers (rows) and feature sets (columns) tested on all samples pooled together (i.e., without distance-based binning).

| | f_1 | f_2 | f_3 | f_4 | f_5 | f_6 |
|------|-------|-------|-------|-------|-------|-------|
| LC | 0.909 | 0.663 | 0.897 | 0.901 | 0.901 | 0.906 |
| RF | 0.838 | 0.559 | 0.872 | 0.896 | 0.905 | 0.931 |
| MLP | 0.908 | 0.666 | 0.914 | 0.925 | 0.921 | 0.940 |
| LSTM | 0.919 | 0.659 | 0.894 | 0.906 | 0.895 | 0.913 |

uniform amount of samples per bin: $d < 0.75$ m, $d \in [0.75, 1)$ m, $d \in [1, 1.25)$ m, $d \in [1.25, 2)$ m, $d \in [2, 2.5)$ m, $d \in [2.5, 3)$ m, and $d \geq 3$ m; this yields 7 AUROC values for each model, each representing its performance on people in a given distance bin; we then average these values together to get an overall metric describing how good a model is to determine user's intention, independently of their distance from the machine. Fig. 3 reports this metric, separately for each distance bin, and averaged over all the distances. We observe that:

- As expected, in non-recurrent models (LC, RF, MLP), f_1 alone is not informative according to the chosen metric (the AUROC is always close to 0.5).
- Consistently over all the models and distances, richer features yield better results.
- The LSTM model does not benefit when provided with explicit velocity information, since this can be already captured by the model itself, which operates on sequential data. For the same reason, the LSTM model performs significantly better than chance, even when given only the distance feature as input, since it can capture and exploit distance variations over time.
- When the models are provided with rich features, predicting performance at short distances is harder (lower AUROC) than at long distances. This can be explained considering the characteristics of our dataset: people in the vicinity of the device often mingle around it for a long time, chatting with others or being busy with other tasks, even if they do not end up interacting with the machine; people that are approaching from afar, in contrast, exhibit clearer intention in their body language and gaze; this

also explains why, for people that lie far from the device, providing orientation and velocity information is very beneficial to performance, whereas the same does not hold for people nearby.

5.1.2. Sequence-level performance

While sample-level performance is a relevant metric to robustly compare different classification approaches, in a real deployment we care about the ability of the approach to correctly classify the intent of a nearby person, as early as possible after the person is first detected. Therefore, we now limit our analysis to the LSTM approach using the f_5 set, which shows the most promising performances for higher distances with no need to explicitly encode velocity features. We report sequence-level metrics, computed as follows.

We consider each sequence in the testing set separately; we evaluate every sample in the sequence and simulate taking an irreversible decision (e.g. to acknowledge the person's presence and demonstrate availability to interact) as soon as the probability returned by the classifier exceeds a given threshold. A sequence for which such probability never exceeds the threshold is a *true negative* if the user does not interact with the robot, or a *false negative* if it eventually does. A sequence for which such probability exceeds the threshold for at least one sample is a *true positive* if the person eventually interacts, or a *false positive* otherwise. Then, we can compute the *true positive rate* (i.e., the *recall*), *false positive rate*, and *precision*. For true positives, we also track the *advance detection time*: the period (in seconds) between the first time the probability exceeds the threshold and the moment in which the interaction actually occurs.

Fig. 4 reports how these metrics change as a function of the threshold. We observe that the resulting AUROC is well above 0.5, indicating a good ability of the approach to discriminate sequences that eventually interact from those that do not; as the threshold increases, the advance detection time (averaged over true positives) decreases, as the system takes a decision later in the sequence, i.e., when the person is closer to the robot.

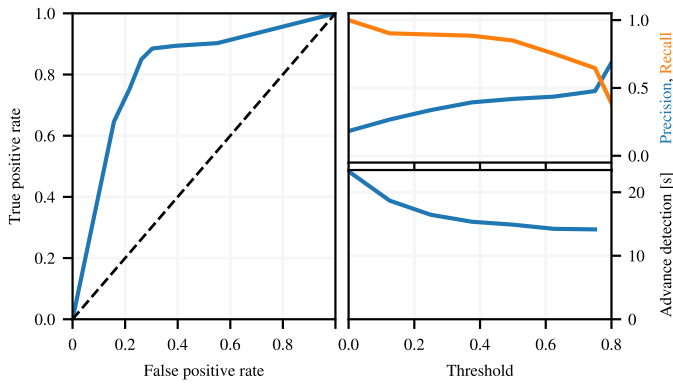


Fig. 4. Coffee break scenario: ROC curve for sequence-level performance (left); Precision, Recall, and Advance detection time w.r.t the threshold of the classifier (right).

5.1.3. Self-supervised learning

We test the ability of the system to improve its performance as new training data is collected in a self-supervised way [14,40]. In particular, we split the collected sequences into 10 disjoint, equally-sized, temporally-contiguous groups. Each group has about 340 sequences and we refer to it in the following as a “day” of data, assuming that the robot is placed in an area with limited visitors. A crowded hall might see the same number of sequences in one hour or less. We then consider a setting in which the robot is deployed with no training at day 0: the robot collects data for one day, then trains a new model using all collected data up to that day, which will be used and evaluated in the following day; the process is repeated for a total of 10 days.

Fig. 5 reports the improvements in the performance measured over the considered period; statistics are reported over 20 runs of the experiment, obtained by randomly shuffling the order of the days. For each run, we take the average AUROC computed over each distance bin independently as explained in 5.1.1. We observe that median performance steeply increases in the first 4 days (about 1500 sequences); additional training further improves AUROC, with reduced returns. Note that, although we did not test this in our current experiment, this approach would be able to automatically adapt to domain shift over time, i.e. caused by changing user demographics, or changing the spatial layout of the environment.

5.2. Robot validation experiments

5.2.1. Self-supervised learning on the waiter robot

We leverage the self-supervised nature of the dataset collected as described in Section 4.2 to implement the behavior of the waiter robot. Similarly to the coffee break scenario, we split the available data in 3 “days” and assume that each one of them is incrementally added to the dataset as time goes by. At day 0 the robot starts without actual training and passively collects data. Each day, new data is recorded and a new model is trained using k -fold cross-validation, where k equals the number of available days of data. We then compute AUROCs for each model and we observe the performance incrementally increasing from 0.500 on day 0, to 0.871 on day 1 and 0.927 on day 2. On day 3, the robot becomes very confident about its prediction as the model yields an AUROC of 0.944. At this point, the robot can start enacting the reaction strategy described in Section 4.2.

The qualitative performance of the model can be observed in Fig. 6 and the supplementary video. Our perception module correctly detects whether someone is approaching the robot to take the chocolate, or is simply passing nearby the robot. This behavior would not be possible to realize if the model used distance-based features only. Furthermore, the model works with multiple users at the same time. In this case, the robot shows availability to interact with the closest user whose intention to interact is predicted.




Fig. 5. Coffee break scenario: AUROC of the model on each day of the self-supervised learning experiment (see text). Boxplots report statistics over 20 runs of the experiment.

5.2.2. Performance in the information robot scenario

Finally, we consider the classifier learned in the waiter robot scenario, and test it in a new setting with different users. In particular, we compute the performance of such a classifier on the dataset collected in the information robot scenario. Fig. 7 reports solid sequence-based metrics: AUROC is approximately equal to 0.99; also, for a threshold of about 0.86 we get a recall of 1, and a precision of about 0.90, maintaining an average advance detection time of more than 3 s, which is a reasonable prediction time for the considered scenario.

A qualitative evaluation of the performance is shown in Fig. 8 and the supplementary video.

5.3. Discussion

The results obtained in the three scenarios illustrate that the proposed approach works well to detect users’ intention to interact before the interaction actually happens. The robot validation experiments exhibit better performance than the experiments in the coffee break scenario: in fact, the former is a controlled environment with users that were specifically tasked to interact with the robot; the latter relies on a more challenging dataset collected in the wild. Nevertheless, the AUROC value computed for the sequence-based analysis in the coffee break scenario confirms the reliability of the approach.

The advance detection time, which is always >3 s, demonstrates that the approach works well in practice. In fact, considering an average human walking speed of 1.35 m/s [49], we can argue that we are able to predict the intention of an approaching person and proactively anticipate them to successfully start an interaction. This intuition is extensively verified in the attached video, where the classifier evaluated in Sections 5.2.1 and 5.2.2 is successfully deployed. In the video, we try to challenge the classifier, recording difficult sequences in which people approach the robot just to pass by it without interacting. In these circumstances, the use of only distance-based features would not be sufficient to correctly classify the user’s intention to interact. Indeed in the recorded sequences, our classifier that makes use of a rich feature set (f_5 in the presented experiments) manages to successfully detect the user’s intentions.

Furthermore, the SSL experiments advocate for the practical use of the proposed approach. Starting from the realistic assumption that a robot can easily determine whether a person has interacted with it (e.g. by pressing a button, or starting a conversation), the results shown in Sections 5.1.3 and 5.2.1 demonstrate how a robot could be deployed in an unknown environment. Most importantly, during the deployment, the robot can autonomously collect new data, improve its predictions, and start to proactively engage people in interactions.

Moreover, the choice of using only spatial-based features extracted from the skeleton of the user, avoiding RGB-D data, proved to be




Fig. 6. Waiter robot scenario: deployment of the classifier. If no interaction is predicted (snapshots on the left) the robot does not react at all. If a user is classified as intending to interact (center), the robot orients its body towards them, turns its LEDs on, and extends its arm to hand out a chocolate treat. When multiple users are detected (right) the robot orients itself towards the closest person that is predicted as intending to interact.




Fig. 7. Information robot scenario: ROC curve for sequence-level performance (left); Precision, Recall, and Advance detection time w.r.t the threshold of the classifier (right).

ideal to make the approach more robust and general. Indeed, using skeleton-derived data as features allows us to be independent of the users' appearance and more robust w.r.t. to the scene background. In fact, we have obtained strong performances with the classifier trained in the waiter robot scenario even when deployed in a new scenario, such as the information robot one, without retraining. Both the video and Fig. 8 showcase this important aspect, displaying sequences from the robot information scenario and the predictions returned by the waiter robot classifier. Also, avoiding image based information makes data collection and processing easier, both in term of computation and privacy concerns.

Finally, both Fig. 6 and the video show how the approach can be deployed to handle multiple users. The input to the classifier is limited to the information related to one single subject (the one whose intention is actually classified) and does not include information related to other neighboring persons that might be influencing the subject's behavior. However, both during training and inference, the presence of multiple people is easily handled as each person is tracked and processed independently. Indeed, in the waiter robot scenario, we demonstrate that we can handle multiple users and interact with the closest person who is predicted to interact. The video qualitatively shows that the robot can proactively behave even when multiple users are present at the same time.

6. Conclusions

We have presented a self-supervised learning approach to predict the user's intention to interact with a robot. To this end, we have collected three datasets in different interaction contexts and settings, with different sizes, containing hundreds of body-tracked users interacting with agents, even within real everyday-life scenarios. We have tested the system with various classification approaches to assess the relevance of the features containing information on the user's pose and motion. We have also simulated the deployment of our strategy in a self-supervised learning fashion and tested it at both sample and sequence levels. Furthermore, we have validated our approach in real human-robot interaction experiments, and involving two different robot platforms. Finally, we have also shown a strategy to proactively react to the user's intention. The presented results are also reported in the supplementary video.

In the future, we will investigate different robot reaction strategies and the way they affect the interaction from the users' perspective. Similarly, we will analyze how the presence of multiple people influences the user's intention to interact. To this end, we plan to augment the feature set of the classifier with information about the people neighboring the tracked user. Furthermore, we will also consider the role played by the robot's appearance and test our framework with different robot platforms. Finally, we plan to conduct an extensive data collection session in public environments and in different social contexts.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

Data will be made available on request.

Acknowledgments

This work was supported by the European Union through the project SERMAS, by the Swiss State Secretariat for Education, Research and Innovation (SERI) under contract number 22.00247, and by the Swiss National Science Foundation grant n. 213074.




Fig. 8. Information robot scenario: snapshots taken from the robot's sensor during two sequences. The bounding boxes that are superimposed on the acquired image report the output of the classifier: red boxes mean a low probability of interaction, whereas green boxes indicate a higher probability of interaction. Top: a user walks through the corridor without interacting with the robot; the system does not predict any interaction. Bottom: another user approaches the robot and the system correctly predicts the intention to interact in advance.

Appendix A. Supplementary data

Supplementary material related to this article can be found online at <https://doi.org/10.1016/j.robot.2023.104568>.

References

- [1] A. Paolillo, F. Colella, N. Nosenko, F. Schiano, W. Stewart, D. Zambrano, I. Chappuis, R. Lalive, D. Floreano, How to compete with robots by assessing job automation risks and resilient alternatives, *Science Robotics* 7 (65) (2022) eabg5561.
- [2] M.K. Lee, S. Kiesler, J. Forlizzi, Receptionist or information kiosk: how do people talk with a robot? in: *ACM Conference on Computer Supported Cooperative Work*, 2010, pp. 31–40.
- [3] A. Tuomi, I.P. Tussyadiah, J. Stienmetz, Applications and implications of service robots in hospitality, *Cornell Hosp. Q.* 62 (2) (2021) 232–247.
- [4] G.A. Zachiotis, G. Andrikopoulos, R. Gornez, K. Nakamura, G. Nikolakopoulos, A survey on the application trends of home service robotics, in: *IEEE Int. Conf. on Robotics and Biomimetics*, 2018, pp. 1999–2006.
- [5] L. Palopoli, A. Argyros, J. Birchbauer, A. Colombo, D. Fontanelli, A. Legay, A. Garulli, A. Giannitrapani, D. Macii, F. Moro, et al., Navigation assistance and guidance of older adults across complex public spaces: the DALi approach, *Intell. Serv. Robot.* 8 (2015) 77–92.
- [6] J. Mišeikius, P. Caroni, P. Duchamp, A. Gasser, R. Marko, N. Mišeikienė, F. Zwilling, C. de Castelbajac, L. Eicher, M. Früh, H. Früh, Lio-a personal robot assistant for human-robot interaction and care applications, *IEEE Robot. Autom. Lett.* 5 (4) (2020) 5339–5346.
- [7] D. Lee, G. Kang, B. Kim, D.H. Shim, Assistive delivery robot application for real-world postal services, *IEEE Access* 9 (2021) 141981–141998.
- [8] J. Avelino, L. Garcia-Marques, R. Ventura, A. Bernardino, Break the ice: a survey on socially aware engagement for human-robot first encounters, *Int. J. Soc. Robot.* 13 (8) (2021) 1851–1877.
- [9] H. Dahlkamp, A. Kaehler, D. Stavens, S. Thrun, G.R. Bradski, Self-supervised monocular road detection in desert terrain, in: *Robotics: Science and Systems*, 2006.
- [10] D. Stavens, S. Thrun, A self-supervised terrain roughness estimator for off-road autonomous driving, in: *Proceedings of the Twenty-Second Conference on Uncertainty in Artificial Intelligence*, AUAI Press, 2006, pp. 469–476.
- [11] A. Lookingbill, J. Rogers, D. Lieb, J. Curry, S. Thrun, Reverse optical flow for self-supervised adaptive autonomous robot navigation, *Int. J. Comput. Vis.* 74 (2006) 287–302.
- [12] T. Mar, V. Tikhonoff, G. Metta, L. Natale, Self-supervised learning of grasp dependent tool affordances on the iCub Humanoid robot, in: *IEEE Int. Conf. on Robotics and Automation*, 2015, pp. 3200–3206.
- [13] S. Levine, P. Pastor, A. Krizhevsky, J. Ibarz, D. Quillen, Learning hand-eye coordination for robotic grasping with deep learning and large-scale data collection, *Int. J. Robot. Res.* 37 (4–5) (2018) 421–436.
- [14] N. Churamani, S. Kalkan, H. Gunes, Continual learning for affective robotics: Why, what and how? in: *Int. Symp. on Robot and Human Interactive Communication*, 2020, pp. 425–431.
- [15] M. Nava, J. Guzzi, R.O. Chavez-Garcia, L.M. Gambardella, A. Giusti, Learning long-range perception using self-supervision from short-range sensors and odometry, *IEEE Robot. Autom. Lett.* 4 (2) (2019) 1279–1286.
- [16] D. Gandhi, L. Pinto, A. Gupta, Learning to fly by crashing, in: *IEEE/RSJ Int. Conf. on Intelligent Robots and Systems*, 2017, pp. 3948–3955.
- [17] M.A. Bekhti, Y. Kobayashi, K. Matsumura, Terrain traversability analysis using multi-sensor data correlation by a mobile robot, in: *IEEE/SICE Int. Symp. on System Integration*, 2014, pp. 615–620.
- [18] C.A. Brooks, K. Iagnemma, Self-supervised terrain classification for planetary surface exploration rovers, *J. Field Robot.* 29 (3) (2012) 445–468.
- [19] R. Hadsell, P. Sermanet, J. Ben, A. Erkan, M. Scoffier, K. Kavukcuoglu, U. Muller, Y. LeCun, Learning long-range vision for autonomous off-road driving, *J. Field Robot.* 26 (2) (2009) 120–144.
- [20] M. Nava, A. Paolillo, J. Guzzi, L.M. Gambardella, A. Giusti, Uncertainty-aware self-supervised learning of spatial perception tasks, *IEEE Robot. Autom. Lett.* 6 (4) (2021) 6693–6700.
- [21] J. Urakami, K. Seaborn, Nonverbal cues in human-robot interaction: A communication studies perspective, *ACM Trans. Hum.-Robot Interact.* 12 (2) (2023) 1–21.
- [22] N. Gasteiger, M. Hellou, H.S. Ahn, Factors for personalization and localization to optimize human-robot interaction: A literature review, *Int. J. Soc. Robot.* (2021) 1–13.
- [23] S. Saunderson, G. Nejat, How robots influence humans: A survey of nonverbal communication in social human-robot interaction, *Int. J. Soc. Robot.* 11 (2019) 575–608.
- [24] J. Rios-Martinez, A. Spalanzani, C. Laugier, From proxemics theory to socially-aware navigation: A survey, *Int. J. Soc. Robot.* 7 (2015) 137–153.
- [25] P. Agand, M. Taherahmadi, A. Lim, M. Chen, Human navigational intent inference with probabilistic and optimal approaches, in: *IEEE Int. Conf. on Robotics and Automation*, 2022, pp. 8562–8568.
- [26] A. Belardinelli, A.R. Kondapally, D. Ruiken, D. Tanneberg, T. Watabe, Intention estimation from gaze and motion features for human-robot shared-control object manipulation, in: *IEEE/RSJ Int. Conf. on Intelligent Robots and Systems*, 2022, pp. 9806–9813.
- [27] S. Vinanzi, C. Goerick, A. Cangelosi, Mindreading for robots: Predicting intentions via dynamical clustering of human postures, in: *Joint IEEE 9th International Conference on Development and Learning and Epigenetic Robotics*, 2019, pp. 272–277.
- [28] A. Zaraki, M. Giuliani, M.B. Dehkordi, D. Mazzei, A. D'ursi, D. De Rossi, An RGB-D based social behavior interpretation system for a humanoid social robot, in: *RSI/ISM International Conference on Robotics and Mechatronics*, 2014, pp. 185–190.

- [29] A. Gaschler, S. Jentzsch, M. Giuliani, K. Huth, J. de Ruiter, A. Knoll, Social behavior recognition using body posture and head pose for human-robot interaction, in: *IEEE/RJS Int. Conf. on Intelligent Robots and Systems*, 2012, pp. 2128–2133.
- [30] J. Schwarz, C.C. Marais, T. Leyvand, S.E. Hudson, J. Mankoff, Combining body pose, gaze, and gesture to determine intention to interact in vision-based interfaces, in: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 2014, pp. 3443–3452.
- [31] A. Belardinelli, Gaze-based intention estimation: principles, methodologies, and applications in *HRI*, 2023, arXiv:2302.04530 [cs].
- [32] M. Brenner, H. Brock, A. Stiegler, R. Gomez, Developing an engagement-aware system for the detection of unfocused interaction, in: *Int. Symp. on Robot and Human Interactive Communication*, 2021, pp. 798–805.
- [33] D. Vaufraydaz, W. Johal, C. Combe, Starting engagement detection towards a companion robot using multimodal features, *Robot. Auton. Syst.* 75 (2016) 4–16.
- [34] Y. Kato, T. Kanda, H. Ishiguro, May I help you? - Design of human-like polite approaching behavior-, in: *10th ACM/IEEE International Conference on Human-Robot Interaction*, 2015, pp. 35–42.
- [35] J. Bi, F.-c. Hu, Y.-j. Wang, M.-n. Luo, M. He, A method based on interpretable machine learning for recognizing the intensity of human engagement intention, *Sci. Rep.* 13 (1) (2023) 2537.
- [36] L. Jing, Y. Tian, Self-supervised visual feature learning with deep neural networks: A survey, *IEEE Trans. Pattern Anal. Mach. Intell.* (2020).
- [37] C. Doersch, A. Zisserman, Multi-task self-supervised visual learning, in: *IEEE International Conference on Computer Vision*, 2017, pp. 2051–2060.
- [38] M. Nava, A. Paolillo, J. Guzzi, L.M. Gambardella, A. Giusti, Learning visual localization of a quadrotor using its noise as self-supervision, *IEEE Robot. Autom. Lett.* 7 (2) (2022) 2218–2225.
- [39] Y. Bengio, A. Courville, P. Vincent, Representation learning: A review and new perspectives, *IEEE Trans. Pattern Anal. Mach. Intell.* 35 (8) (2013) 1798–1828.
- [40] T. Lesort, V. Lomonaco, A. Stoin, D. Maltoni, D. Filliat, N. Díaz-Rodríguez, Continual learning for robotics: Definition, framework, learning strategies, opportunities and challenges, *Inf. Fusion* 58 (2020) 52–68.
- [41] N. Marquardt, S. Greenberg, Informing the design of proxemic interactions, *IEEE Pervasive Comput.* 11 (2) (2012) 14–23.
- [42] Microsoft, Azure Kinect sensor SDK system requirements, 2023, <https://learn.microsoft.com/en-us/azure/kinect-dk/system-requirements>, Accessed: 2023.
- [43] T. Mahajan, G. Singh, G. Bruns, G. Bruns, T. Mahajan, G. Singh, An experimental assessment of treatments for cyclical data, in: *Computer Science Conference for CSU Undergraduates*, Vol. 6, 2021, p. 22.
- [44] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, E. Duchesnay, Scikit-learn: Machine learning in Python, *J. Mach. Learn. Res.* 12 (2011) 2825–2830.
- [45] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Kopf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, S. Chintala, PyTorch: An imperative style, high-performance deep learning library, in: H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché Buc, E. Fox, R. Garnett (Eds.), *Advances in Neural Information Processing Systems 32*, Curran Associates, Inc., 2019, pp. 8024–8035.
- [46] S. Hochreiter, J. Schmidhuber, Long short-term memory, *Neural Comput.* 9 (8) (1997) 1735–1780.
- [47] DJI, Robomaster EP core, 2023, <https://www.dji.com/ch/robomaster-ep-core>, Accessed: 2023.
- [48] T. Yamamoto, K. Terada, A. Ochiai, F. Saito, Y. Asahara, K. Murase, Development of human support robot as the research platform of a domestic mobile manipulator, *ROBOMECH J.* 6 (1) (2019) 1–15.
- [49] R.W. Bohannon, A. Williams Andrews, Normal walking speed: a descriptive meta-analysis, *Physiotherapy* 97 (3) (2011) 182–189, <http://dx.doi.org/10.1016/j.physio.2010.12.004>, URL: <https://www.sciencedirect.com/science/article/pii/S0031940611000307>.



Gabriele Abbate is a researcher and software engineer at the Dalle Molle Institute for Artificial Intelligence (IDSIA, USI-SUPSI) in Lugano since 2019. He received a double Master's degree in Informatics from Università della Svizzera Italiana (Switzerland) and Università degli Studi di Milano-Bicocca (Italy) in 2018. At IDSIA, he leads technical development in AI and machine learning for various projects, including image classification, segmentation, large-scale 3D urban point clouds, and gesture detection. He's also active in human–robot interaction, with a focus on robotics simulation, virtual reality tools, software integration, and rehabilitation systems.



Alessandro Giusti is a Professor of AI for Autonomous Robotics at the Dalle Molle Institute for Artificial Intelligence (IDSIA, USI-SUPSI) in Lugano, Switzerland. He leads a group of 10 researchers and PhD students working on Self-Supervised robot learning, real-world Deep Learning applications, and Human–Robot Interaction. He teaches university courses at the Swiss Master of Science in Engineering (Zurich), USI (Lugano), SUPSI (Lugano), and Politecnico di Milano on Robotics, Computer Vision, Deep Learning and Data Science topics.



Viktor Schmuck received his M.Sc. degree in Medialogy from Aalborg University, Aalborg, Denmark, in 2018. He spent one year as a Researcher with Aalborg University, Aalborg, DK. From 2019 to 2023, he has been a Ph.D. student with the Center for Robotics Research, Department of Engineering, King's College London, as a member of the Social AI & Robotics Laboratory. His research focuses on human–robot interaction, machine learning within the scope of human and group behavior understanding, and socially-aware navigation of robots. He is currently a Postdoctoral Research Associate at the Social AI & Robotics Laboratory at KCL, primarily working on socially-aware robotic applications focused on Human–Robot Interaction.



Oya Celiktutan is an Associate Professor (Senior Lecturer) in the Center for Robotics Research, Department of Engineering, King's College London, United Kingdom, where she is the Head of Social AI & Robotics Laboratory, <https://sairlab.github.io/>. Her research focuses on two overarching questions: how to model human behavior from multimodal data and how to transfer these models to robots for learning, action, and interaction. Through an interdisciplinary lens, she explores the intersection of machine learning and human–robot interaction.



Antonio Paolillo is a researcher at the Dalle Molle Institute for Artificial Intelligence (IDSIA, USI-SUPSI) in Lugano. Previously, he was a post-doc at EPFL, Lausanne, Switzerland (2019–20); Idiap Research Institute, Martigny, Switzerland (2018–19); CNRS-University of Montpellier, France (2015–17). He received his Ph.D. and M.Sc. from Sapienza University of Rome, Italy, in 2015 and 2011, respectively. He visited the CNRS-AIST Joint Robotics Laboratory, Tsukuba, Japan (2015); CNRS-University of Montpellier, France (2014); and Örebro University, Sweden (2010). He is an Associate Editor for RA-L. His research interests include robotic control, machine learning and AI for robotics, human–robot interaction, and rehabilitation robotics.