

Sentiment-Based Engagement Strategies for Intuitive Human-Robot Interaction

Thorsten Hempel^a, Laslo Dinges^b and Ayoub Al-Hamadi^c

*Neuro-Information Technology, Faculty of Electrical Engineering and Information Technology,
Otto von Guericke University, Magdeburg, Germany*

Keywords: Human-Robot Interaction, Approaching Strategy, Sentiment Estimation, Emotion Detection, Anticipating Human Behaviors, Approaching People, Productive Teaming.

Abstract: Emotion expressions serve as important communicative signals and are crucial cues in intuitive interactions between humans. Hence, it is essential to include these fundamentals in robotic behavior strategies when interacting with humans to promote mutual understanding and to reduce misjudgements. We tackle this challenge by detecting and using the emotional state and attention for a sentiment analysis of potential human interaction partners to select well-adjusted engagement strategies. This way, we pave the way for more intuitive human-robot interactions, as the robot's action conforms to the person's mood and expectation. We propose four different engagement strategies with implicit and explicit communication techniques that we implement on a mobile robot platform for initial experiments.

1 INTRODUCTION

Once utopian, robots are increasingly moving from industrial and laboratory settings to the real-world in order to assist humans in everyday tasks. In this process, human-robot interaction became a central point of interest in robotic research that investigates the manifold challenges in performing interactive and collaborative tasks. As a general principle, the robot needs to be acquired with the necessary skills to enable intuitive interactions with arbitrary human interaction partners, regardless of the corresponding task, the human's intention and communication behavior. To this end, a fundamental objective is the anticipation of appropriate strategies to proactively approach or evade people based on the situation-specific context, such as the person's mood, attitude (Müller-Abdelrazeq et al., 2019; Elprama et al., 2016) and resentments towards robots (Naneva et al., 2020; Miller et al., 2021). First, it requires the determination of people's interest in interacting at all (Satake et al., 2009; Finke et al., 2005). Thereupon, the robot has to find an appropriate strategy for either engaging people to establish an interaction or to avoid a confrontation if desired. In order to reach these under-

lying decisions, the robot has to cope with a row of sub-tasks: estimating people's focus of attention (Dini et al., 2017), predicting their mental state, and carrying out a (dis-)engagement strategy (Avelino et al., 2021).

Especially, the estimation of the interaction willingness is a very challenging task, as this mental state tends to provide only vague social signals. As an alternative approach, many methods (Oertel et al., 2020) go after the detection of engagement in order to determine if a person already approached the robot and waits for the robot's reaction. This state expresses itself more clearly, e.g., in voice commands (Foster et al., 2017), gestures and proxemics, but it undermines the proactive approach and neglects situations where the human partner is in need of help, but too uncertain or not aware of the robot.

In this work, we close the gap between the lack of proactive behavior and advanced mental state predictions by introducing a sentiment-based engagement strategy.

By combining the focus of attention and emotions, we derive a sentiment analysis that allows us to select fine-grained behavior patterns that exceed current binary engagement and disengagement strategies. We implement our approach on a mobile robot platform and execute our strategies using explicit and implicit embodied communication. Our models are trained on minimum sized networks in order to perform our

^a <https://orcid.org/0000-0002-3621-7194>

^b <https://orcid.org/0000-0003-0364-8067>

^c <https://orcid.org/0000-0002-3632-2402>

method on mobile systems with limited computational resources.

2 RELATED WORK

In recent years, proxemic rules have been a popular tool to draw conclusions about the different states of human-robot interactions (Mumm and Mutlu, 2011) and how spatial zones can be leveraged to improve them (Syrdal et al., 2007). (Repiso et al., 2018) predict appropriate encounter points to achieve a natural engagement with a group of people. Similarly, (Satake et al., 2009) proactively approach detected people to offer help. (Fischer et al., 2021) verbally greet people and dynamically adapt the voice volume based on the distance to the target person. Following (Kendon, 1990), (Carvalho et al., 2021) apply the Kendon’s greeting model to approach people while tracking the human mental states of the interaction using multiple features, such as gaze, facial expressions, and gestures. Yet, all of these methods neglect if the human counterpart is actually interested and ready to initiate an interaction. This is problematic, as in the case of negative attitude, the robot’s interaction efforts can be perceived as rude and annoying (Bröhl et al., 2019). An initial attempt to address this was presented by (Kato et al., 2015), who only approach and offer help if they detect signs of attention towards the robot. But this still doesn’t incorporate the actual mental attitude of the human counterpart towards the interaction itself.

The recent advances in the area of deep-learning opened up new possibilities for the visual analysis of humans, starting from general detection tasks to the estimation of facial micro expressions. This allows the aggregation of additional relevant information, such as emotions (Chuah and Yu, 2021; Spezialetti et al., 2020), to better understand the behavior and intentions of the human interaction partner. To the best of our knowledge, we are the first to leverage this advancement by predicting sentiment states based on visual emotion estimation to improve interaction strategy for robots. We summarize our main contributions as follows:

- We propose an image-based sentiment analysis to gauge the current interaction preference of potential human interaction partners.
- We propose four different behavior strategies (Engage, Attract, Ignore, Avoid) that comprises a number of different explicit and implicit communication modalities.



Figure 1: The TIAGo robot platform with the key operations used in our method.

- We design ultra lightweight models to implement our method and integrate it on a mobile robot platform.
- We conduct initial experiments in laboratory settings for an early evaluation.

3 METHOD

This section explains each step of our proposed method for sentiment-based interaction behavior.

The sentiment analysis is based on two main features: head pose and emotion estimation. The head pose estimation is a leading indicator to gauge the current visual focus of attention of the person. The emotion estimation predicts the current emotional state of the person. Together, it allows the assumption about a person that *a*) is seeking for interaction, *b*) is indecisive about it or *c*) doesn’t want to engage with the robot at all. Determining each sentiment state enables to perform a dedicated robot behavior that corresponds with the expected reaction from the person and, thus, improves an intuitive interaction. Figure 2 illustrates an overview of our proposed system. In the following subsection, we will give details about each component of the system and its interplay.

3.1 Robot Platform

We use the TIAGo robot from PAL Robotics, a mobile service robot for indoor environments. It is equipped with an RGB-D camera mounted inside its head, that can be yawed and pitched to dynamically perceive the environment. For grasping and moving objects, the

robot has a manipulator arm with a 5 finger gripper. The torso can be lifted to adjust the robot's height, and integrated speakers allow the output of voice commands. These abilities allow the combination of explicit and implicit communication modalities that we use in our method.

- *Head Following*: The robot's head aims at and follows a moving person as an implicit signal of attentiveness without actively approaching.
- *Body Following*: If a person exceeds the range of the robot's moving head, the entire robot base rotates in order to keep track of the person. This behavior communicates an attention commitment and is therefore an even stronger signal as the *head following*.
- *Torso Lifting*: Lifting the torso is a rather subtle communication cue, as it is perceived similarly to the social convention of standing up while greet someone.
- *Speech*: Approaching people with speech represents an explicit communication strategy, that stronger commits them to a reaction. Hence, speech should be only used if the approached person is highly expected to welcome an interaction.

3.2 Visual Attention

Typical indicators for gauging the human visual attention are the gaze and, more coarsely, the head pose. As the gaze depends on the eyes, which takes only a small part in face images, it is prone to errors. Therefore, we estimate the visual attention of the surrounding persons based on a head pose prediction algorithm. At first, we locate faces in the image stream using an ultra light SSD face detector. Then, the face crops are further processed by the 6DRepNet (Hempel et al., 2022), that uses a rotation representation to directly regress yaw, pitch, and roll angle of the faces. To assure real-time processing capabilities, we replace the original 6DRepNet backbone with the efficient MobileNetv3-Small (Howard et al., 2019). The new head pose prediction model is able to maintain 90% of its accuracy compared to the original model, but is downsized to only one tenth of its parameters count.

3.3 Emotion Estimation

In order to assess the attitude towards robots as well as differentiated subjective states such as dislike or willingness to cooperate, basic emotions can be used as indicating features. Naturally, these are frequently

communicated by facial expressions. Hence, automatic facial expression recognition is an essential method to generate features for Human-Robot-Collaboration and attention prediction.

Previous facial expression recognition methods adhere to the established pipeline of face detection, landmark extraction, and action unit (AU). Then, emotions are classified using these AUs as feature vectors (Wegrzyn et al., 2017; Vinkemeier et al., 2018; Werner et al., 2017). However, end-to-end learning on a single but more comprehensive, and better generalizing database, often outperforms such traditional approaches. The AffectNet database (Mollahosseini et al., 2017) for example contains about 500k in-the-wild samples for neutral and the basic emotions happy, disgust, fear, surprise, anger, and sadness. Furthermore, it is the only database, that also includes ground truth of valence and arousal (VA), which are, in contrast to the basic emotion classes, continuous labels $\in [-1, 1]$ which can be used for regression tasks. VA is less intuitively interpretable by humans, but it allows differentiation between facial expressions of varying degrees. This is particularly useful for capturing modest but long-term changes in expressed emotion. In addition, a multitask network (simultaneous training of classification and regression) also improved the accuracy of the classification. Furthermore, Valence summarizes essential information about several emotions in one parameter (negative, neutral, positive), which may facilitate the development of heuristic behavior rules for some scenarios.

In our former work, we proposed a multitask network based on YOLOv3 (Dinges et al., 2021). For the current paper, we compared EfficientNetV2-S, ResNet18, and MobileNetV3-Large. Despite the poorer converging training loss, MobileNet, which is optimized for limited hardware, achieved almost the same accuracy on the test set as EfficientNet (deviation $< 1\%$). In addition, we used a weighted random sampler to handle the imbalanced dataset and re-train the network on the training and test set to generate an application model with increased robustness to unseen data.

3.4 Sentiment Analysis and Approaching Strategy

Detecting emotional expressions combined with the visual focus of attention allows perceiving important communicative signals. As in human-human interaction, you may be more likely to greet a stranger who smiles at you but avoid eye contact if that stranger is scowling. We transfer these principles to deter-

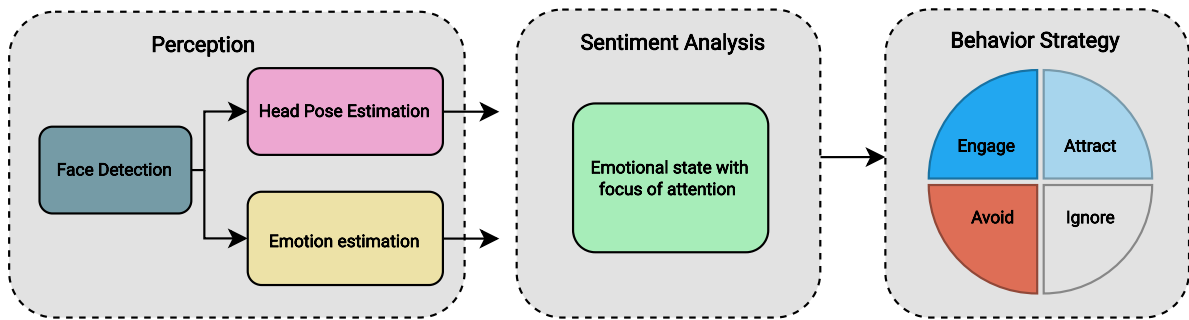


Figure 2: Method overview of our proposed system. Detected faces are used to estimate the head pose and the corresponding emotions based on facial expressions. Both information is used to determine the emotional state with the corresponding visual focus of the person's attention that indicates the sentiment and allows a fine-tuned behavior strategy.

	Facing	Averted
Happy	ENGAGE: Lift torso, follow with head and base, and approach by speech	ATTRACT: Lift torso and follow with head
Surprise		
Neutral		
Uncertain		
Sad	AVOID: Look down and avoid attraction	IGNORE: Follow with head
Fear		
Disgust		
Anger		

Figure 3: Behavior strategies based on the emotional and attentive state.

mine states of sentiment based on the combination of the different emotions and visual attention. Is a person showing visual signs of fear or anger, while he is attentive towards the robot, it indicates a negative attitude towards the robot. In this case, a passive robot behavior without hectic movements can help to defuse the situation and to calm the person. If a person is feared, but not attentive towards the robot, the reason for the fear might not be the robot. Executing a more active robot behavior to show awareness could be therefore a more suitable strategy. Hence, our sentiment analysis allows us to compartmentalize the common hard binary engagement strategy into soft behavior. Instead of choosing between engaging and not engaging, we choose between four different behaviors: Engage, Attract, Avoid, Ignore.

- *Engage*: Engaging will extend the torso of the robot to signal the awareness of the person's presence. Additionally, voice commands will greet the people to actively initiate an interaction. This

strategy is called, when the attention is directed at the robot with clear positive emotional attitude.

- *Attract*: Positive emotional attitude without attention towards the robot leads to more passive behavior. The goal is to signal the offer for interaction and help without actively committing to an interaction. This is done by extending the torso to show awareness, facing the person, but without verbally greeting. This way, a non-verbal offer for support (especially in case of uncertainty and sadness) is subtly given, that can be easily accepted or ignored.
- *Avoid*: The avoiding strategy is executed when strong negative emotions (Fear, Disgust, Anger) are facing the robot. In this case, the robot avoids eye contact (keeps the face outside the camera center) and takes a passive part. Especially in case of fear, the goal is to increase the feeling of safety by avoiding unexpected motions.
- *Ignore*: If the negative emotions are not directed at the robot, the robot mostly ignores the person, only following it with its head to capture changes in emotional or attention states.

Figure 3 shows all combinations of emotional and attention states with corresponding behavior strategy.

4 IMPLEMENTATION

We used the ROS¹ middleware to implement our proposed solution on our robot platform, as this also offers a simplified integration into other ROS-based robot platforms. Face detection, head pose and emotion estimation are executed in separate nodes, that forward the necessary information to perform the sentiment analysis and the action controls. While the latter is performed on the robot itself, the inference of

¹<https://www.ros.org/>




Figure 4: Sentiment analysis in a laboratory setting. The person is tracked using a face detector, followed by a head pose and emotion estimation. As the person shows attentiveness and happy, a positive sentiment towards the robot is expected and a proactive initiation strategy selected.

the neural networks is outsourced to a separate notebook that is mounted on the shoulder of the robot. This ensures the complete mobility of the robot, while exploiting additional computational resources. The notebook has a Quadro RTX 4000 processing all models in parallel with the following inference times:

- Face Detection: 6.7 ms
- Head Pose Estimation: 1.4 ms
- Emotion Estimation: 6.3 ms

The visual perception can be therefore performed in real time and is able to continuously track people in the surroundings.

5 EXPERIMENT

We conducted an initial experiment in a laboratory setting to test our implemented method on the TIAGo robot platform. The objective of the experiment was to gain experience about the model's performances in real application scenarios. Figure 4 shows a snapshot from one of our tests. The robot platform detects a face in the environment and predicts head pose and emotional state. The focus of attention is on the robot (indicated by the green arrow) in conjunction with expressing happiness, so the robot initiates an interaction by following the engage strategy.

While the models provide very robust predictions, the implementation is currently entirely based on single shot detections. Yet, over the duration of face tracking, facial expressions change intermediately and short distractions make the head pose occasionally oscillating. This leads to abrupt changes of

states within short time periods. Hence, to improve the robustness of the sentiment analysis, conclusions about the current state should incorporate the predictions from multiple time frames.

6 CONCLUSION

In this paper, we present a sentiment-based methodology to improve the intuitiveness and reasoning in human-robot interactions. Our method focuses on the visual perception of mental states and the focus of attention to derive a sentiment analysis for suitable interaction initiation strategies. We take also into account the case where people are not interested in an engagement at all, which has been mainly neglected in the recent literature. To implement our method, we trained lightweight perception models that we integrate on a mobile robot platform. Finally, we conducted initial experiments using a mobile robot platform to analyze the performance of our models and the overall system.

In the future, we will extend our work by embedding our sentiment analysis in a temporal-probabilistic framework to improve the robustness in real-world scenarios. Here, we will also incorporate the valence and arousal predictions to evaluate the emotion's intensity and transitions between emotional states. Further enhancements offer the inclusion of the robot's mobility to actively engage and disengage people or to add additional interaction modalities, such as following people by command, getting out of the way of passing people, or approaching a place that offers a better overview of the surroundings.

Finally, we will conduct a study with test subjects to evaluate our approach in terms of intuitiveness and sympathy.

ACKNOWLEDGEMENTS

This work is funded and supported by the Federal Ministry of Education and Research of Germany (BMBF) (AutoKoWAT-3DMA under grant Nr. 13N16336) and German Research Foundation (DFG) under grants AI 638/13-1, AI 638/14-1 and AI 638/15-1.

REFERENCES

- Avelino, J., Garcia-Marques, L., Ventura, R., and Bernardino, A. (2021). Break the ice: a survey on

- socially aware engagement for human–robot first encounters. *International Journal of Social Robotics*, 13:1851 – 1877.
- Bröhl, C., Nelles, J., Brandl, C., Mertens, A., and Nitsch, V. (2019). Human-robot collaboration acceptance model: Development and comparison for germany, japan, china and the usa. *I. J. Social Robotics*, 11:709–726.
- Carvalho, M., Avelino, J., Bernardino, A., Ventura, R. M. M., and Moreno, P. (2021). Human-robot greeting: tracking human greeting mental states and acting accordingly. *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 1935–1941.
- Chuah, S. H.-W. and Yu, J. (2021). The future of service: The power of emotion in human-robot interaction. *Journal of Retailing and Consumer Services*, 61:102551.
- Dinges, L., Al-Hamadi, A., Hempel, T., and Al Aghbari, Z. (2021). Using facial action recognition to evaluate user perception in aggravated hrc scenarios. In *2021 12th International Symposium on Image and Signal Processing and Analysis (ISPA)*, pages 195–199.
- Dini, A., Murko, C., Yahyanejad, S., Augsdörfer, U., Hofbaur, M., and Paletta, L. (2017). Measurement and prediction of situation awareness in human-robot interaction based on a framework of probabilistic attention. In *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 4354–4361.
- Elprama, S., El Makrini, I., Vanderborgh, B., and Jacobs, A. (2016). Acceptance of collaborative robots by factory workers: a pilot study on the role of social cues of anthropomorphic robots.
- Finke, M., Koay, K. L., Dautenhahn, K., Nehaniv, C., Walters, M., and Saunders, J. (2005). Hey, i’m over here - how can a robot attract people’s attention? In *RO-MAN 2005. IEEE International Workshop on Robot and Human Interactive Communication, 2005.*, pages 7–12.
- Fischer, K., Naik, L., Langedijk, R. M., Baumann, T., Jelínek, M., and Palinko, O. (2021). Initiating human-robot interactions using incremental speech adaptation. In *Companion of the 2021 ACM/IEEE International Conference on Human-Robot Interaction, HRI ’21 Companion*, page 421–425, New York, NY, USA. Association for Computing Machinery.
- Foster, M. E., Gaschler, A., and Giuliani, M. (2017). Automatically classifying user engagement for dynamic multi-party human–robot interaction. *International Journal of Social Robotics*, 9.
- Hempel, T., Abdelrahman, A. A., and Al-Hamadi, A. (2022). 6d rotation representation for unconstrained head pose estimation. In *2022 IEEE International Conference on Image Processing (ICIP)*, pages 2496–2500.
- Howard, A., Sandler, M., Chu, G., Chen, L.-C., Chen, B., Tan, M., Wang, W., Zhu, Y., Pang, R., Vasudevan, V., Le, Q. V., and Adam, H. (2019). Searching for mobilenetv3. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*.
- Kato, Y., Kanda, T., and Ishiguro, H. (2015). May i help you? - design of human-like polite approaching behavior-. In *2015 10th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pages 35–42.
- Kendon, A. (1990). *Conducting Interaction: Patterns of Behavior in Focused Encounters*. Cambridge University Press, Cambridge, U.K.
- Miller, L., Kraus, J., Babel, F., and Baumann, M. (2021). More than a feeling—interrelation of trust layers in human-robot interaction and the role of user dispositions and state anxiety. *Frontiers in Psychology*, 12.
- Mollahosseini, A., Hasani, B., and Mahoor, M. H. (2017). Affectnet: A database for facial expression, valence, and arousal computing in the wild. *IEEE Transactions on Affective Computing*, 10:18–31.
- Müller-Abdelrazeq, S. L., Schönefeld, K., Haberstroh, M., and Hees, F. (2019). *Interacting with Collaborative Robots—A Study on Attitudes and Acceptance in Industrial Contexts*, pages 101–117. Springer International Publishing, Cham.
- Mumm, J. and Mutlu, B. (2011). Human-robot proxemics: Physical and psychological distancing in human-robot interaction. In *2011 6th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pages 331–338.
- Naneva, S., Gou, M. S., Webb, T. L., and Prescott, T. J. (2020). A systematic review of attitudes, anxiety, acceptance, and trust towards social robots. *International Journal of Social Robotics*, pages 1–23.
- Oertel, C., Castellano, G., Chetouani, M., Nasir, J., Obaid, M., Pelachaud, C., and Peters, C. (2020). Engagement in human-agent interaction: An overview. *Frontiers in Robotics and AI*, 7.
- Repiso, E., Garrell, A., and Sanfeliu, A. (2018). Robot approaching and engaging people in a human-robot companion framework. In *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 8200–8205.
- Satake, S., Kanda, T., Glas, D. F., Imai, M., Ishiguro, H., and Hagita, N. (2009). How to approach humans?- strategies for social robots to initiate interaction. In *2009 4th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pages 109–116.
- Spezialetti, M., Placidi, G., and Rossi, S. (2020). Emotion recognition for human-robot interaction: Recent advances and future perspectives. *Frontiers in Robotics and AI*, 7.
- Syrdal, D. S., Lee Koay, K., Walters, M. L., and Dautenhahn, K. (2007). A personalized robot companion? - the role of individual differences on spatial preferences in hri scenarios. In *RO-MAN 2007 - The 16th IEEE International Symposium on Robot and Human Interactive Communication*, pages 1143–1148.
- Vinkemeier, D., Valstar, M. F., and Gratch, J. (2018). Predicting folds in poker using action unit detectors and decision trees. In *FG*, pages 504–511.
- Wegrzyn, M., Vogt, M., Kireclioglu, B., Schneider, J., and Kissler, J. (2017). Mapping the emotional face. How

individual face parts contribute to successful emotion recognition. *PLOS ONE*, 12(5):e0177239.

Werner, P., Handrich, S., and Al-Hamadi, A. (2017). Facial action unit intensity estimation and feature relevance visualization with random regression forests. In *2017 Seventh International Conference on Affective Computing and Intelligent Interaction (ACII)*, volume 2018-Janua, pages 401–406. IEEE.

