

Towards Dialogue-based Navigation with Multivariate Adaptation driven by Intention and Politeness for Social Robots

Chandrakant Bothe¹, Fernando Garcia², Arturo Cruz Maya²,
Amit Kumar Pandey², and Stefan Wermter¹

¹ Knowledge Technology, Department of Informatics
University of Hamburg, Hamburg, Germany
{bothe, wermter}@informatik.uni-hamburg.de

² SoftBank Robotics Europe, Paris, France
{ferran.garcia, arturo.cruzmaya, akpandey}@softbankrobotics.com

Abstract. Service robots need to show appropriate social behaviour in order to be deployed in social environments such as healthcare, education, retail, etc. Some of the main capabilities that robots should have are navigation and conversational skills. If the person is impatient, the person might want a robot to navigate faster and vice versa. Linguistic features that indicate politeness can provide social cues about a person's patient and impatient behaviour. The novelty presented in this paper is to dynamically incorporate politeness in robotic dialogue systems for navigation. Understanding the politeness in users' speech can be used to modulate the robot behaviour and responses. Therefore, we developed a dialogue system to navigate in an indoor environment, which produces different robot behaviours and responses based on users' intention and degree of politeness. We deploy and test our system with the Pepper robot that adapts to the changes in user's politeness. ³

Keywords: Social Robots · Dialogue System · Effect of Politeness · Natural Language Understanding · Human-Robot Interaction.

1 Introduction

The perception of politeness of a user can be a reflection of their patience during interaction. In addition to other factors such as the robot appearance, robot behaviour is a crucial aspect for their acceptance. Hence, politeness cues are intimately related to the dynamics of behavior and interaction [4, 7, 12, 23]. It is useful for adapting to the dynamic tension [20] that occurs as a user tries to maintain a sufficient degree of politeness while interacting with the robot. For example, sentence-initial *you* or an action directive verb can be impolite “*You*

³ This is the pre-print of the article published at ICSR 2018 and final authenticated version is available online at: https://link.springer.com/chapter/10.1007/978-3-030-05204-1_23

need to show...” or *“Show me the...”*, whereas sentence-medial *you* or sentence-initial *could* or *would* often indicates the politeness like in these sentences *“Could you show me...”* or *“Would you take me to...”*.

Multivariate adaptive and affective dialogue systems based on linguistic features have been subject to previous research [1, 8, 22]. The effect of politeness on the conversation is prominent and it has been researched in the sociolinguist community [7, 12, 20]. The effect of such a feature on human-robot interaction (HRI) has been a subject of study with various aspects: an impolite vs. a polite robot playing a game [5], in determining social robot acceptance with multicultural background people [21], making robots sociable and to achieve safe HRI [8]. Hence, a robot that can recognize the intention of the user during interaction should also adapt to the human’s linguistic behavioural changes. For example, different sociolinguistic features such as politeness, emotion, sentiment, etc. represent the user’s behavioural dynamics. If the user’s utterance is impolite, then he/she might be in a hurry and vice versa. In such cases, the robot might need to change its behaviour or even alter some actions or speed up or down the movements. We develop a modular dialogue system (DS) that can process such features and make a robot to adapt accordingly. The natural language understanding part of the DS uses recurrent neural networks (RNNs) [25, 26] and the Snips library [6] for extracting the structured information from the user input. The politeness detection is learned from the human-annotated corpus and fine-tuned for the scenario-specific data. The navigation of the Pepper robot is achieved by using the NAOqi framework. The robot behaviour and responses are driven by dialogue flow module using the intention and politeness. The main contributions of the present work towards bridging the gap between sociolinguistic research and HRI community are:

- developing a dialogue-based navigation for incorporating politeness, and
- incorporating sociolinguistic features for robotic behavioural modelling.

To the best of our knowledge, our system is the first dialogue-based navigation that incorporates politeness as an important social cue to drive the robot behaviour and responses.

2 Approach

We propose a dialogue system which takes into account the degree of politeness as a factor that affects the conversation flow and the robot behaviour. The dialogue system takes the intention and other information, for example, slot-value pairs (see details in Section 3.1) into account to understand the input utterance [25]. However, in the proposed model, the dialogue system processes sociolinguistic features to make inferences regarding the input utterance. The overall architecture is shown in Figure 1. The proposed system is customizable to any extent as the robot is controlled using a client-server architecture. The state and motion managers are wrapped into an application programming interface (API) as a server [10] and communicated via the dialogue flow module. The dialogue system can be accessed also if the robot is not connected to the server.

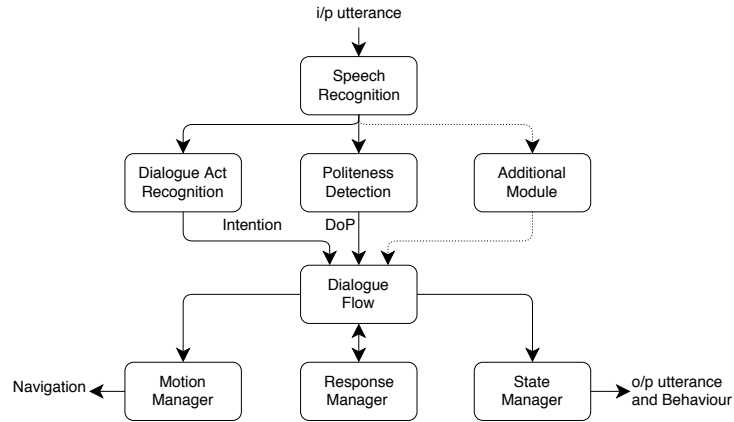


Fig. 1. The overall architecture of the dialogue system. DoP: degree of politeness.

3 Dialogue System

3.1 Natural Language Understanding

The input speech from a user is converted into text using the speech recognition module (from NAOqi). The language understanding module takes the converted input utterance forms a symbolic representation and provides the degree of politeness. The dialogue act recognition module is used to extract the symbolic representation and politeness detection module to detect the degree of politeness of that input utterance.

Dialogue Act Recognition Module The dialogue act (DA) recognition is a crucial process in the dialogue system. The task is to decode the input utterance and form the symbolic representation, such as dialogue acts and slot-value pairs. For example, the utterance “*could you please show me the retail department*” can be decoded as $\{da : TakeToPlace, room : retail\}$ where *da* represents the dialogue act or intention, *room* is a slot and *retail* its value. We created a dataset for the given scenario to be able to drive the conversation (some examples are given in Table 1). The following methods are used in conjunction for robustness by validating one another based on heuristics of their confidence values:

(1) Dialogue act recognition using RNNs: The architecture is shown in Figure 2, where RNNs are used in a hierarchical fashion to learn the dialogue acts and slot-value pairs [3, 14, 26]. The RNNs are better at encoding the contextual and sequential information in the utterance [14]. The dialogue acts are classified with the first layer of the RNN, preserving the utterance representation utt_rep1 . utt_rep1 is then used to recognize slots on the next layer of the RNN, producing a new utterance representation utt_rep2 . The values of slots are learned with next layer of the RNN using utt_rep2 and the detected slot as a switch (\in) for




Fig. 2. Dialogue acts and slot-value pairs recognition using RNNs.

the belonging values learned in this layer (see the output in Figure 5 for better understanding). We fit the model to the data and use the trained model for inference [3].

(2) Snips Natural Language Understanding (NLU) Engine: Snips NLU Engine⁴ is an open source Python library that uses two approaches: (a) a deterministic parser and (b) a probabilistic parser [6]. The deterministic parser is basically a pattern matching mechanism which uses regular expressions to parse the input utterance. The probabilistic parser uses logistic regression for intent classification and conditional random fields (CRFs) for slot filling. For the given input utterance, the engine provides the intention and slot-value pairs.

Politeness Detection Module The politeness detection module takes the input utterance as an input and computes its degree of politeness. An RNN is used to learn the degree of politeness from Stanford Politeness Corpus⁵[7]. We fine-tune the trained model for the dataset (mentioned in the previous section) that is created for the particular scenario to minimize uncertainty in prediction. The degree of politeness (DoP) varies from 1 to -1 (very polite to very impolite). For the sake of conceptual and computational simplicity, we discretized them into categories: polite (1), neutral (0) and impolite (-1); see the examples below:

DoP	Class	Utterance
1	polite	Could you please show me the education department?
0	neutral	Can you show me the education department?
-1	impolite	Show me the education department.

⁴ <https://snips-nlu.readthedocs.io>

⁵ <https://www.cs.cornell.edu/~cristian/Politeness.html>

Table 1. Examples of dialogue act and slot-value pairs

Dialogue acts	Examples	Slots	Values
Greeting	<i>Hello.</i> <i>Hi, how are you?</i>	no_slot	no_value
Thanking	<i>Thank you.</i> <i>Thank you very much.</i>	no_slot	no_value
TakeToPlace	<i>Could you show me the education department?</i> <i>Take me to the retail section.</i> <i>Can you take me to tourism department?</i>	room	retail education tourism
MoveRobot	<i>Please go ahead.</i> <i>Could you move ahead?</i> <i>Go back please.</i>	direction	forward backward right left
TurnRobot	<i>Can you turn right?</i> <i>Could you turn left.</i>		
Accept	<i>Yes, I would like to visit.</i>	no_slot	no_value
AbortRobot	<i>stop, wait, be careful...</i>	no_slot	no_value

Additional Module This module is open to adding additional sociolinguistic features such as sentiment, emotion, etc. Adding more features can increase the complexity of the dialogue system. However, it could be useful in some cases to incorporate multiple features and modalities to produce the required behaviour.

3.2 Dialogue Flow

The dialogue flow is a central engine of the system which communicates with most of the modules. It is implemented as a main function to drive the DS. A rule-based and probabilistic belief tracking or dialogue state tracking model could be used to maintain the dialogue flow [25]. We used a rule-based model where the dialogue flow module keeps track of the input dialogue acts and DoP and send them to the response manager to fetch responses. The complete state loop has a queue to store the context information of the preceding utterances. It is helpful to trigger new dialogue acts based on the context information. For example, if the last dialogue act is *TakeToPlace*, it triggers a new dialogue act called *FinishedOne* to inform the system that the last action was finished and asks the user if he/she wishes to visit the next place. Another loop keeps track of whether the user accepts or rejects the proposal using *Accept* and *Reject* dialogue acts. If one of the dialogue acts appears, the robot takes the user to the next location until either the list of locations is finished or the user rejects to visit the next place.

3.3 Response Manager

The response manager is responsible for picking up the right response for the given intention and degree of politeness. Pre-defined response templates are stored in a data file that is accessed continuously during the interaction.




Fig. 3. The behavioural model used to create the verbal and non-verbal responses based on the cumulative sum of the DoP. The Pepper robot shown in the right is in the position of the vertical orange line in the plot during the interaction.

4 Robot Control and Navigation

4.1 Robot Platform

Pepper is a 1.2 meter tall omnidirectional wheeled humanoid robot platform capable of exhibiting body language, perceiving and interacting with its surroundings, and move autonomously. Due to its 17 joints and 20 degrees of freedom (DoF) kinematic configuration and edgeless design, the system is suitable for safe HRI [18]. The platform is equipped with a large variety of sensors and actuators that ensure safe navigation and a high degree of expressiveness: LED's are distributed across the head (eyes and ears) and torso (shoulders) to support non-verbal communication by modifying colour and intensity. The microphones and speakers allow verbal interaction as well as environmental awareness. Sensing components include three laser sensors, two sonars and two infrared sensors located in the robot's base, as well as two cameras and a three-dimensional camera located in the head. Finally, the platform is powered by an Atom processor with a 1.91 GHz quad-core unit that allows the NAOqi SDK to orchestrate the different hardware elements as well as their access from other APIs.

4.2 State Manager

In order to produce the physical and verbal responses in accordance with the degree of politeness exhibited during the interaction, a behavioural model inspired by the valence and arousal model [2] has been designed. The model is given the discrete DoP computed from the last utterance being [1, 0, -1] and maps the cumulative sum of the politeness of the previous and current utterances to different actuators. In this way, a variability is provided to every single social cue that can vary in order to fit the interaction needs.




Fig. 4. The environment map created with the Pepper robot and gmapping from ROS.

The actuators used to externalize the robot’s change of state are the LED’s color [17], head pitch orientation [15], voice pitch [13] and navigation speed, and are mapped following the intuition (shown in Figure 3). For example, a user repetitively polite during the whole interaction will experience a decrement in the navigation speed of the robot, a head position oriented towards the user, green coloured eyes and a slightly higher voice pitch.

4.3 Motion Manager

The motion manager is responsible for navigation and can be operated in the following modes:

Tele-operation In this mode, the Pepper robot could be teleoperated with the help of the NAOqi framework using the *moveToward* function from ALMotion service and the keys on the keyboard are used for moving or stopping the robot.

Scripted Navigation The scripted navigation is achieved by commanding a robot to move to the specific positions/places with the known distances in the environment. This is also achieved with NAOqi framework using *moveTo* command from ALMotion service. We specify how far the robot has to move (in meters) and the orientations (in radians) it has to take during motion.

Navigation: Mapping and Planning This module requires the use of the Robot Operating System (ROS), an open source middle-ware framework. To fit our need for navigation, we have adopted the following approach for generating and post-processing the map. The current readings of the Pepper’s depth image are converted into virtual laser data, using the package *depthimage_to_laserscan* [19, 24]. An offline map (shown in Figure 4, and post-processed for testing purposes) can be acquired using *gmapping* (laser-based SLAM) [11]. Then, the localization is performed using Adaptive Monte Carlo Localization (*acml*) [9]. Finally, the navigation uses a global planner with a map with inflated obstacles (costmap) and a local costmap with observations from the virtual laser data. The Dialog Flow requests a location from the API server (on the robot using a virtual machine) using an ID and this one sends the coordinates to the ROS navigation stack to execute the path.




Fig. 5. Output of the DA recognition module.

5 Experiments and Results: A Real-World Scenario

The task is to navigate in the given environment to show a user the different departments. Our tour scenario in the lab consists of four departments: retail, education, tourism and healthcare as shown in the map in Figure 4. The robot is a guide which takes the user to the particular department using verbal interaction as mentioned in Section 3.2. When the user asks the robot, the input utterance gets processed by the DA recognition module which produces the result as shown in Figure 5. The politeness detection module provides the DoP of that utterance. The dialogue flow communicates this information with all the managers. The robot adapts its behaviour such as speeding up or down while navigating to the locations and changing the pitch of speech, changing the pitch angle of the head.

We tested our system on the Pepper robot with different users, expressing different levels of politeness. The behavioural changes and adaptation to speed change based on a change in DoP are shown in Figure 6. The robot behavioural adapts to the human being polite; the robot slows down and spends more time with the user. When the user is impolite, the robot speeds up and executes motion faster. The proposed behaviour of the robot for different situations shown in the figure is mainly to demonstrate the developed system and the efficacy of




Fig. 6. Robot internal state for polite (a) and impolite (b) interactions.

the proposed framework. The results indicate that the system is able to consider the linguistic features to modulate the navigation behaviour of the robot in a coherent theoretical and functional framework. As aforementioned, to the best of our knowledge such a framework and implementation in a practical situation is one of the first attempts of its kind. However, it is important to mention that the validation of the hypotheses about the most appropriate behaviours of the robot is not within the scope of this paper and it will require further investigation and user studies. As mentioned in the conclusion such studies are one of the next steps to utilize the framework for different situations. The demonstration video and dialogue logs of the generated graphs in Figure 6 are available at the SECURE EU Project website: <https://secure-robots.eu/fellows/bothe/secondment-project>.

6 Conclusions and Future Work

We developed a dialogue-based navigation system for integrating intention and politeness features for multivariate adaptation of the robot. We successfully deployed and tested our system on the robot with different levels of politeness. Currently, our work does not elicit the causal explanation for the behaviour and the multivariate adaptation of the robot. However, our experimental framework opens up a new challenge for the study of the effect of politeness in human-robot social interaction. We strongly believe that our work will be helpful in bridging the gap between sociolinguistic research and the HRI community. This research shall also be helpful in targeting the deployment of social-service robots with adaptation to sociolinguistic features such as politeness. In this work, the behaviours are based on previous research [8, 16, 21, 23]. The validation of the system is crucial and it will be addressed in future work through user studies.

Acknowledgements

This project has received funding from the European Union’s Horizon 2020 framework programme for research and innovation under the Marie Skłodowska-Curie Grant Agreement No. 642667 (SECURE), the Industrial Leadership Agreement (ICT) No. 779942 (CROWDBOT), and No. 688147 (MuMMER).

References

1. Adam, C., Johal, W., Pellier, D., Fiorino, H., Pesty, S.: Social Human-Robot Interaction: A New Cognitive and Affective Interaction-Oriented Architecture. In: International Conference on Social Robotics. pp. 253–263. Springer (2016)
2. Beck, A., Cañamero, L., Bard, K.A.: Towards an Affect Space for robots to display emotional body language . In: 19th International Symposium in Robot and Human Interactive Communication. pp. 464–469 (2010)
3. Bothe, C., Magg, S., Weber, C., Wermter, S.: Discourse-Wizard: Discovering Deep Discourse Structure in your Conversation with RNNs. arXiv:1806.11420 (2018)
4. Brown, P., Levinson, S.C.: Politeness: Some Universals in Language Usage, vol. 4. Cambridge University Press (1987)
5. Castro-González, Á., Castillo, J.C., Alonso-Martín, F., Olortegui-Ortega, O.V., González-Pacheco, V., Malfaz, M., Salichs, M.A.: The Effects of an Impolite vs. a Polite Robot Playing Rock-Paper-Scissors. In: International Conference on Social Robotics. pp. 306–316. Springer (2016)

6. Coucke, A., Saade, A., Ball, A., Bluche, T., Caulier, A., Leroy, D., Doumouro, C., Gisselbrecht, T., Caltagirone, F., Lavril, T., et al.: Snips Voice Platform: an embedded Spoken Language Understanding system for private-by-design voice interfaces. arXiv preprint arXiv:1805.10190 (2018)
7. Danescu-Niculescu-Mizil, C., Sudhof, M., Jurafsky, D., Leskovec, J., Potts, C.: A computational approach to politeness with application to social factors. In: Proceedings of ACL 2013 (Volume 1: Long Papers). pp. 250–259 (2013)
8. Fong, T., Nourbakhsh, I., Dautenhahn, K.: A survey of socially interactive robots. *Robotics and Autonomous Systems* **42**(3-4), 143–166 (2003)
9. Fox, D.: Adapting the Sample Size in Particle Filters Through KLD-Sampling. *The International Journal of Robotics Research* **22**(12), 985–1003 (2003)
10. Grinberg, M.: *Flask Web Development: Developing Web Applications with Python*. O’Reilly Media, Inc. (2018)
11. Grisetti, G., Stachniss, C., Burgard, W.: Improving Grid-based SLAM with Rao-Blackwellized Particle Filters by Adaptive Proposals and Selective Resampling. In: International Conference on Robotics and Automation. pp. 2432–2437 (2005)
12. Holmes, J., Stubbe, M.: *Power and Politeness in the Workplace: A Sociolinguistic Analysis of Talk at Work*. Routledge (2015)
13. Hubbard, D.J., Faso, D.J., Assmann, P.F., Sasson, N.J.: Production and perception of emotional prosody by adults with autism spectrum disorder. *Autism Research* **10**(12), 1991–2001 (2017)
14. Kumar, H., Agarwal, A., Dasgupta, R., Joshi, S., Kumar, A.: Dialogue Act Sequence Labeling using Hierarchical encoder with CRF. *AAAI Conference on Artificial Intelligence* pp. 3440–3447 (2018)
15. Lemaignan, S., Garcia, F., Jacq, A., Dillenbourg, P.: From Real-time Attention Assessment to “With-me-ness” in Human-Robot Interaction. In: International Conference on Human Robot Interaction. pp. 157–164 (2016)
16. Manav, B.: Color-emotion associations and color preferences: A case study for residences. *Color Research & Application* **32**(2), 144–150 (2007)
17. Nijdam, N.A.: *Mapping Emotion to Color*. Citeseer (2009)
18. Pandey, A., Gelin, R.: A Mass-Produced Sociable Humanoid Robot: Pepper: The First Machine of Its Kind. *IEEE Robotics Automation Magazine* pp. 40–48 (2018)
19. Perera, V., Pereira, T., Connell, J., Veloso, M.: Setting Up Pepper For Autonomous Navigation And Personalized Interaction With Users. arXiv:1704.04797 (2017)
20. Rogers, P.S., Lee-Wong, S.M.: Reconceptualizing Politeness to Accommodate Dynamic Tensions in Subordinate-to-Superior Reporting. *Journal of Business and Technical Communication* **17**(4), 379–412 (2003)
21. Salem, M., Ziadee, M., Sakr, M.: Marhaba, how may I help you?: Effects of Politeness and Culture on Robot Acceptance and Anthropomorphization. In: International Conference on Human-robot Interaction. pp. 74–81 (2014)
22. Shi, W., Yu, Z.: Sentiment Adaptive End-to-End Dialog Systems. In: Proceedings of ACL 2018. pp. 1509–1519 (2018)
23. Srinivasan, V., Takayama, L.: Help Me Please: Robot Politeness Strategies for Soliciting Help From People. In: Proceedings of the Conference on Human Factors in Computing Systems. pp. 4945–4955. ACM (2016)
24. Suddrey, G., Jacobson, A., Ward, B.: Enabling a Pepper Robot to provide Automated and Interactive Tours of a Robotics Laboratory. arXiv:1804.03288 (2018)
25. Ultes, S., Rojas Barahona, L.M., Su, P.H., Vandyke, D., Kim, D., Casanueva, I., Budzianowski, P., Mrkšić, N., Wen, T.H., Gasic, M., Young, S.: PyDial: A Multi-domain Statistical Dialogue System Toolkit. In: Proceedings of ACL 2017, System Demonstrations. pp. 73–78 (2017)

26. Yang, X., Chen, Y.N., Hakkani-Tür, D., Crook, P., Li, X., Gao, J., Deng, L.: End-to-End Joint Learning of Natural Language Understanding and Dialogue Manager. In: Proc. of IEEE ICASSP 2017. pp. 5690–5694 (2017)