

Who’s in Charge? Roles and Responsibilities of Decision-Making Components in Conversational Robots

PIERRE LISON, University of Oslo, Norway

CASEY KENNINGTON, Boise State University, USA

Software architectures for conversational robots typically consist of multiple modules, each designed for a particular processing task or functionality. Some of these modules are developed for the purpose of making decisions about the next action that the robot ought to perform in the current context. Those actions may relate to physical movements, such as driving forward or grasping an object, but may also correspond to communicative acts, such as asking a question to the human user. In this position paper, we reflect on the organization of those decision modules in human-robot interaction platforms. We discuss the relative benefits and limitations of modular vs. end-to-end architectures, and argue that, despite the increasing popularity of end-to-end approaches, modular architectures remain preferable when developing conversational robots designed to execute complex tasks in collaboration with human users. We also show that most practical HRI architectures tend to be either robot-centric or dialogue-centric, depending on where developers wish to place the “command center” of their system. While those design choices may be justified in some application domains, they also limit the robot’s ability to flexibly interleave physical movements and conversational behaviours. We contend that architectures placing “action managers” and “interaction managers” on an equal footing may provide the best path forward for future human-robot interaction systems.

Additional Key Words and Phrases: software architectures, dialogue management, robot planning, human-robot interaction

1 INTRODUCTION

Conversational interaction allows humans to communicate with robots using a medium that is both intuitive and highly expressive (i.e., speech), but endowing robots with the ability to comprehend, represent, and produce spoken responses and integrate a speech system into the robots’ own states and plans has been a subject of research for many years with limited success. The mere integration of an off-the-shelf automatic speech recognizer into the robotic platform is far from sufficient, as system developers need to address a host of additional challenges in communication, including how to handle turn-taking, dealing with noise, uncertainties and misunderstandings, and building common ground [37]. Those phenomena, which extend well beyond speech processing, are precisely what the spoken dialogue systems (SDS) community has been addressing for many years. If humans are to naturally interact with robots, then robots need SDS—particularly, multimodal SDS, given the intrinsically multimodal nature of face-to-face interactions [8].

Both robotic platforms and multimodal SDS are complex, multifaceted technical systems, but their integration in a single architecture raise a number of additional challenges that the community needs to face as robots that converse with humans using speech are expected to become more commonplace. In this position paper, we focus on one central question for autonomous systems: *decision making*. The problem of selecting the next action to take given the current state (as perceived by the robot sensors) has been widely studied in robotics [46] and has led to a wide range of specialized decision systems, from motion and path planning [18, 39] to object manipulation [5] and task planning [21]. Those decisions may either rely on reactive mechanisms [11], planning algorithms [26] or, increasingly, data-driven strategies optimised via imitation or reinforcement learning [12, 47].

SDS also need to make decisions about what to say next, given the current state of the conversation (as perceived/understood by the agent). Those decisions are typically made by a dedicated module called the *dialogue manager* [35, 60], sometimes also called the *interaction manager* in human–robot interaction [1]. The dialogue manager is responsible for deciding what to say and when to say it, and is itself typically divided in two sub-tasks: dialogue

state tracking and action selection. While the goal of dialogue state tracking is to maintain a representation of the current dialogue state given the stream of speech inputs and other observations, action selection relies on this state representation to select the most appropriate system response.

A key question for action selection in both robotics and in multimodal SDS is how to operate in large state-action spaces. This question becomes particularly acute when both types of systems are combined into a single platform. In this position paper, we review some of the design choices related to the “chain of command” underpinning software architectures for conversational robots. In particular, we discuss the relative merits and limitations of modular architectures in light of the growing popularity of end-to-end approaches. We also observe that many HRI systems tend to adopt either a robot-centric or dialogue-centric view of the decision chain, often influenced by the particular research background of their developers. We argue that such design biases, while justifiable in some application domains, may also limit the system ability to handle complex situations including both physical and conversational tasks. The view advocated in this paper is that task/motion managers (responsible for the physical actions executed by the robot) and interaction managers (responsible for the robot’s conversational behaviour) should ideally be placed side-by-side as part of a larger, modular architecture.

2 BACKGROUND

In this section, we review some seminal and recent research literature on decision making in SDS and robotics. The reviews are by no means comprehensive, but should give the reader enough background to understand the architectural issues discussed in Section 3.

2.1 Decision-Making in SDS

The *Information State* approach to dialogue management [31, 51] has been highly influential in SDS research and continues to inspire modern dialogue systems. The approach revolves around a symbolic representation of the current dialogue state, which includes all types of information relevant for managing the system’s conversational behaviour, such as questions under discussion, commitments, beliefs, and intentions from dialogue participants. Two sets of rules operate on this dialogue state:

- Update rules modify/extend the dialogue state on the basis of observations, such as new user inputs.
- Decision rules then select the most appropriate dialogue move(s) on the basis of the dialogue state.

Traum and Larsson [51] already note that there is no universal agreement on what should constitute a decision maker because different systems might carve up system modularity in different ways, some centralizing all decision-making in a single module while others rely on a more de-centralized distribution of labor (for example, to a natural language generation module that takes high-level information from the decision maker and generates an actual utterance).

Early implementations of the Information State framework used symbolic representations of the dialogue state (and task/domain knowledge) along with handcrafted rules for both state update and action selection. Those handcrafted rules are now often replaced by data-driven models, as detailed in the large literature on dialogue state tracking and dialogue policy optimization. Dialogue state tracking [17, 23, 57] seeks to predict the current dialogue state (typically represented as slot-value pairs) based on the stream of user inputs. Based on the current state, the next action is then selected by a dialogue policy optimized through supervised or reinforcement learning [41, 48, 62].

Also seminal is the early work of [9] on plan-based approaches to dialogue management, building on earlier work by [3]. Planning, the authors argue, is necessary if an agent is to be *intentional*, defined as having choice with

commitment. Although plan-based frameworks are still used in a number of application domains where the task completion necessitates long-term reasoning [45], most current approaches rely on dialogue policies specifying a (learned) direct mapping between states and actions¹. If the dialogue state is represented as probability distribution over possible slot-value pairs, this mapping can be formalized as a Partially Observable Markov Decision Process (POMDP), as shown by [59, 61]. POMDPs have also been employed in robotics [56], although they are – as in SDS research – increasingly replaced by neural methods.

End-to-end approaches that directly produce system responses from the user inputs (or their transcription in the case of speech-enabled interfaces) using encoder-decoder neural models are also increasingly popular [16, 20, 63, 64]. Those approaches do not compute any explicit representation of the dialogue state, but encode the user inputs and the dialogue history into latent vector representations that are then employed to generate the system response token-by-token. Those end-to-end models are typically trained from dialogue corpora, either directly or on the basis of existing large neural language models.

2.2 Decision-Making in Robotics

Decision-making and planning constitute a central problem in robotics, encompassing various tasks from high-level task planning [21, 25] to path planning [18, 22] and to object manipulation/grasping and motion planning [5, 38, 39]. Given a particular robot morphology and degrees of freedom, the planning algorithm seek to derive a path to an expected target configuration or location of the robot. This path is then employed to determine which actuators or motors to activate, and when. Robot planning can either rely on model-based approaches with formal planning languages such as PDDL and its variants [6, 7, 14], or employ data-driven methods, often based on (deep) reinforcement learning [4, 33].

[15] recently reviewed the literature of planning and control of mobile robots that can navigate a space autonomously using one or more of various sensors (e.g., laser or vision-based guidance). Interestingly, in recent years autonomous mobile robots are seeing an increase in a decentralized decision making module (particularly for navigation). Planning is still central, but smaller decisions can be made by other modules more specific to a particular aspect of the robot, a claw for example. This is similar to work done in SDS, where many systems only use the decision making module for high-level decisions, though the issue of division of labor is by no means settled.

Similar to, though distinct from, robotics is autonomous vehicle control, which also requires (fast, high-stakes) real-time decision making in novel situations. [52]’s review of the autonomous vehicle control literature identified a need to fill in a gap between integrating continuous sensor information and making decisions based on more abstract logic-based reasoning. Interestingly, they mention beliefs, goals, and intentions which have been important in planning and decision making literature in several fields, including SDS [9], as explained above.

2.3 Decision-Making in Human–Robot Interaction

As a general trend, the choice of decision-making framework for human–interaction frameworks seems to be strongly influenced by the particular background and research interests of their developers. In other words, SDS researchers will often view the interaction manager as constituting the central decision module of their platform [28, 29, 36]. Under this view, task planning is either subordinated to this interaction manager or encompassed into it. Similarly, robotics

¹This direct mapping from state to actions can be seen as representing precomputed plans for every possible state. This “precomputation” is often achieved through reinforcement learning, i.e. by traversing large numbers of possible dialogue trajectories and using the resulting outcomes to estimate the expected cumulative reward $Q(s, a)$ of an action a in a given state s . Those Q -values are now typically expressed as a neural network. By contrast, online planning algorithms determine the desirability of each action at runtime, using e.g. forward planning [44].

researchers will typically place task/motion planners at the center of their architecture and will often relegate dialogue management to a second-order component limited to handling speech inputs/outputs [10, 13].

Much of the literature that reports on the integration of (multimodal) SDS in robotic platforms seems to be focused on either SDS or robotics, with few examples of practical systems integrating conversational abilities to robots executing complex physical tasks. Conversational robots are typically developed and evaluated in the context of social tasks that either do not require much movement, as in [27]) or where the movements are subordinated to communicative functions, such as gestures [30, 53] or facial expressions [2].

3 OPEN QUESTIONS

In this section, we review two key questions related to the design of software architectures for conversational robots:

- where to adopt modular vs. end-to-end architectures ;
- how the decision modules should interact with one another.

3.1 Modular vs. End-to-end SDS

End-to-end SDS rely on a single model – typically a neural network – that takes an (often text-based) input and directly produce an output, such as a question-answering system that provides an answer to a question, or a social chatbot that produces responses given text input. End-to-end systems often focus on the ability to produce a plausible response to any input, based on an encoder-decoder model trained or fine-tuned from dialogue data. The “decisions” made by the model correspond in this case to the successive selection of tokens in the response. In other words, end-to-end conversational models do not have any explicit concept of task or plan – they are only optimized to produce plausible responses. As such, their use for task-oriented systems remains challenging, although see [55, 58] for a few attempts.

In contrast, modular systems are made up of modules that have well-defined roles in the system, and which can be made to communicate with each other. For example, a prototypical SDS is often made up five modules including automatic speech recognition (ASR) that transcribes speech to a text representation of the human utterances, natural language understanding (NLU) that takes the text and yields a semantic abstraction, dialogue management (including dialogue state tracking) that makes the high-level decision about the next action to take given the semantic abstraction, natural language generation (NLG) that takes the dialogue manager’s decision and determines which words to use and in what order, and text-to-speech (TTS) which actually utters the response.

Both types of approaches have their advantages and limitations. End-to-end systems are conceptually simpler (even though their neural architectures may be large and composed of many layers), and do not suffer from cascading failures as in modular systems. As they are fully differentiable, they may also be optimized directly from the observed inputs and outputs, without needing to rely on annotations or semantic abstractions.

However, end-to-end approaches also have a number of shortcomings. Most importantly, from a designer perspective, their conversational behaviour is much more difficult to control, adjust, and interpret. System designers can never be sure of what the system will end up producing as response, and the tendency of generative language models to e.g. hallucinate facts is well documented [24]. As they are primarily text-based, their integration with non-linguistic knowledge sources may also be challenging, although there exists a number of approaches integrating neural models with knowledge bases [17, 20] or image/videos [32, 49]. Finally, even though few-shots learning approaches may be employed to mitigate problems of data scarcity [42, 54], their portability to scenarios with no or limited data remains difficult.




Fig. 1. Architectural choices for decision-making components in conversational human-robot interaction platforms. The left side shows a SDS-centric view in which the interaction manager is the top decision-making module, and the right side a robot-centric view where the robot planners and controllers take precedence. At the center, a schematic architecture in which both decision-making modules are set on an equal footing and operate on the basis of a shared memory and common resources.

In our view, end-to-end models may be useful for speech-enabled robots whose primary purpose is to serve as social companion. However, modular approaches remain preferable for conversational robots set to execute physical tasks in the real world, either alone or in collaboration with humans. This position is motivated by several considerations. Most importantly, task-oriented robots are themselves developed around modular architectures such as ROS [43], and a modular SDS is much easier to integrate in those platforms than a monolithic, end-to-end conversational model. This is particularly true when the interaction management needs to be tightly interfaced with other decision-making modules for task, path or motion planning, which need to rely on dedicated representations and algorithms and cannot be reduced to a token-by-token generation task. In addition modular approaches also provide better control and interpretability, which is a decisive factor when developing complex robotic platforms operating in the real world.

3.2 Interactions between Decision-Making Modules

Which modules should be responsible for making decisions about the (linguistic and non-linguistic) actions that can be performed by the robot? As noted in Section 2.3, most current approaches tend to adopt either a SDS-centric view or a robot-centric view of the decision chain. In the SDS-centric view, the main center of command is the interaction manager, which is typically retooled to include decisions for robot actions. In contrast, a robot-centric view places the task/motion planning as the center stage, only including the SDS as a secondary subsystem limited to receiving speech inputs and/or communicating information back to the user.

Both design choices may certainly be appropriate in certain application domains. For instance, SDS-centric architectures may work well for social robots designed to provide stimulating social interactions with their human interlocutors. In those cases, managing the conversational behaviour of the robot is indeed the main objective to which all modules are subordinated. Similarly, robot-centric architectures are well-suited for robots that are primarily developed for the execution of a particular physical task, and only use SDS as a way to convey information to its human users.

However, we contend that such architectures are insufficient for robotic systems that aim to achieve complex, collaborative tasks including mixed-initiative dialogues with human users. Indeed, such scenarios will often interleave task execution and spoken interactions in complex, often unpredictable ways. This interplay of behaviours is difficult to address in a single decision-making module, and is better handled through an architecture where the interaction manager operates in parallel to the robot planners and action controllers. Conceptual visualizations for some of these different decision making architectures can be seen in Figure 1.

The interaction manager and the other robot-related components need, however, to interact closely, as their actions may impact one another. For instance, the robot may not be able to make gestures while in motion. In addition, the interaction manager may need to obtain information about the physical state of the robot and/or status of the task to complete in order to answer questions uttered by the user. As shown in Figure 1, one way to allow for an efficient exchange of information is to rely on a shared memory that can be read/written by those decision-making modules, and can also be used to resolve potential conflicts between the actions selected by each component.

Recent work on the EVA platform shows an impressive system that is equal parts SDS- and “robot-” (i.e., virtual agent-) centric in that both the decisions made about system speech acts and virtual agent articulation are both treated as first-class citizens and important to the unfolding interaction.² The underlying model is plan-based and, from our understanding, despite being fairly modular is quite centralized in how decisions are made. Modeling such a system of course is challenging as state and actions spaces can potentially be prohibitively large. Other work has used a centralized module for dialogue and robot action decision making, though the set of actions were fairly limited [34, 50]

4 DISCUSSION & CONCLUSION

Decision making in SDS and robotics is a very important part of a workable system in both fields. In this paper, we looked at how the fields differ, and offered some possible paths forward on how to model decision making when SDS and robots are brought together. In order for the community to share resources such as models and platforms, it is prudent to find a principled solution to what often seems to be ad-hoc modeling decisions. We have by no means comprehensively solved the problem of determining how decision-making components should interact in complex human-robot interaction systems. Moreover, we did not discuss how different state representations (i.e., how inputs and actions are represented in a system) can affect modeling decisions, nor did we touch on the importance of time horizons in decision making (i.e., some decisions are near-term, while others are longer-term like in chess where there is always a next move, but there is also an overall strategy that includes planning moves beyond just the next one). There may be a way to systematically test different methods of decision making to determine important trade offs using a task that is equal parts dialogue and robot control, such as [40], a virtual robot recipe task, which has been annotated with dialogue acts [19], which we leave for future work.

REFERENCES

- [1] Bahareh Abbasi, Natawut Monaikul, Zhanibek Rysbek, Barbara Di Eugenio, and Miloš Žefran. 2019. A multimodal human-robot interaction manager for assistive robots. In *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 6756–6762.
- [2] Samer Al Moubayed, Jonas Beskow, Gabriel Skantze, and Björn Granström. 2012. Furhat: A Back-Projected Human-Like Robot Head for Multiparty Human-Machine Interaction. In *Cognitive Behavioural Systems*. Springer Berlin Heidelberg, 114–130.
- [3] James Frederick Allen. 1979. *A plan-based approach to speech act recognition*. Ph. D. Dissertation.
- [4] Szilárd Aradi. 2020. Survey of deep reinforcement learning for motion planning of autonomous vehicles. *IEEE Transactions on Intelligent Transportation Systems* 23, 2 (2020), 740–759.
- [5] Aude Billard and Danica Kragic. 2019. Trends and challenges in robot manipulation. *Science* 364, 6446 (2019), eaat8414.
- [6] Gerard Canal, Michael Cashmore, Senka Krivić, Guillem Alenyà, Daniele Magazzeni, and Carme Torras. 2019. Probabilistic planning for robotics with ROSPlan. In *Towards Autonomous Robotic Systems: 20th Annual Conference, TAROS 2019, London, UK, July 3–5, 2019, Proceedings, Part I 20*. Springer, 236–250.
- [7] Michael Cashmore, Maria Fox, Derek Long, Daniele Magazzeni, Bram Ridder, Arnau Carrera, Narcis Palomer, Natalia Hurtos, and Marc Carreras. 2015. Rosplan: Planning in the robot operating system. In *Proceedings of the international conference on automated planning and scheduling*, Vol. 25. 333–341.
- [8] Oya Celiktutan, Efstratios Skordos, and Hatice Gunes. 2017. Multimodal human-human-robot interactions (mhhri) dataset for studying personality and engagement. *IEEE Transactions on Affective Computing* 10, 4 (2017), 484–497.

²<https://openstream.ai/>

- [9] Philip R Cohen and Hector J Levesque. 1990. Intention Is Choice with Commitment*. *Artif. Intell.* (1990).
- [10] Ricardo Cruz, Luís Garrote, Ana Lopes, and Urbano J Nunes. 2018. Modular software architecture for human-robot interaction applied to the InterBot mobile robot. In *2018 IEEE International Conference on Autonomous Robot Systems and Competitions (ICARSC)*. IEEE, 17–23.
- [11] L De Silva and H Ekanayake. 2008. Behavior-based robotics and the reactive paradigm a survey. In *2008 11th International Conference on Computer and Information Technology*. IEEE, 36–43.
- [12] Bin Fang, Shidong Jia, Di Guo, Muhua Xu, Shuhuan Wen, and Fuchun Sun. 2019. Survey of imitation learning for robotic manipulation. *International Journal of Intelligent Robotics and Applications* 3 (2019), 362–369.
- [13] Tobias Fischer, Jordi-Ysard Puigbò, Daniel Camilleri, Phuong DH Nguyen, Clément Moulin-Frier, Stéphane Lallée, Giorgio Metta, Tony J Prescott, Yiannis Demiris, and Paul FMJ Verschure. 2018. iCub-HRI: a software framework for complex human–robot interaction scenarios on the iCub humanoid robot. *Frontiers in Robotics and AI* 5 (2018), 22.
- [14] Maria Fox and Derek Long. 2003. PDDL2. 1: An extension to PDDL for expressing temporal planning domains. *Journal of artificial intelligence research* 20 (2003), 61–124.
- [15] Giuseppe Fragapane, René de Koster, Fabio Sgarbossa, and Jan Ola Strandhagen. 2021. Planning and control of autonomous mobile robots for intralogistics: Literature review and research agenda. *Eur. J. Oper. Res.* 294, 2 (Oct. 2021), 405–426.
- [16] Jianfeng Gao, Michel Galley, and Lihong Li. 2019. *Neural approaches to conversational AI: Question answering, task-oriented dialogues and social chatbots*. Now Foundations and Trends.
- [17] Shuyang Gao, Abhishek Sethi, Sanchit Agarwal, Tagyoung Chung, Dilek Hakkani-Tur, and Amazon Alexa AI. 2019. Dialog State Tracking: A Neural Reading Comprehension Approach. In *20th Annual Meeting of the Special Interest Group on Discourse and Dialogue*. 264.
- [18] Alessandro Gasparetto, Paolo Boscariol, Albano Lanzutti, and Renato Vidoni. 2015. Path planning and trajectory planning algorithms: A general overview. *Motion and Operation Planning of Robotic Systems: Background and Practical Approaches* (2015), 3–27.
- [19] Spandana Gella, Aishwarya Padmakumar, Patrick Lange, and Dilek Hakkani-Tur. 2022. Dialog acts for task-driven embodied agents. (Sept. 2022). arXiv:2209.12953 [cs.CL]
- [20] Marjan Ghazvininejad, Chris Brockett, Ming-Wei Chang, Bill Dolan, Jianfeng Gao, Wen-tau Yih, and Michel Galley. 2018. A knowledge-grounded neural conversation model. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 32.
- [21] Marc Hanheide, Moritz Göbelbecker, Graham S Horn, Andrzej Pronobis, Kristoffer Sjö, Alper Aydemir, Patric Jensfelt, Charles Gretton, Richard Dearden, Miroslav Janicek, et al. 2017. Robot task planning and explanation in open and uncertain worlds. *Artificial Intelligence* 247 (2017), 119–150.
- [22] Abdelfetah Hentout, Abderraouf Maoudj, and Mustapha Aouache. 2022. A review of the literature on fuzzy-logic approaches for collision-free path planning of manipulator robots. *Artificial Intelligence Review* (Sept. 2022).
- [23] Léo Jacqmin, Lina M Rojas Barahona, and Benoit Favre. 2022. "Do you follow me?": A Survey of Recent Approaches in Dialogue State Tracking. In *Proceedings of the 23rd Annual Meeting of the Special Interest Group on Discourse and Dialogue*. 336–350.
- [24] Ziwei Ji, Nayeon Lee, Rita Frieske, Tiezheng Yu, Dan Su, Yan Xu, Etsuko Ishii, Yejin Bang, Andrea Madotto, and Pascale Fung. 2022. Survey of hallucination in natural language generation. *Comput. Surveys* (2022).
- [25] Erez Karpas and Daniele Magazzeni. 2020. Automated planning for robotics. *Annual Review of Control, Robotics, and Autonomous Systems* 3 (2020), 417–439.
- [26] Karthik Karur, Nitin Sharma, Chinmay Dharmatti, and Joshua E Siegel. 2021. A survey of path planning algorithms for mobile robots. *Vehicles* 3, 3 (2021), 448–468.
- [27] Casey Kennington, Kotaro Funakoshi, Yuki Takahashi, and Mikio Nakano. 2014. Probabilistic multiparty dialogue management for a game master robot. In *Proceedings of the 2014 ACM/IEEE international conference on Human-robot interaction - HRI '14*. ACM, Bielefeld, Germany, 200–201.
- [28] Casey Kennington, Daniele Moro, Lucas Marchand, Jake Carns, and David McNeill. 2020. rrSDS: Towards a robot-ready spoken dialogue system. In *Proceedings of the 21th annual meeting of the special interest group on discourse and dialogue*. 132–135.
- [29] Bernd Kiefer and Christian Willms. 2020. Implementing diverse robotic interactive systems using VONDA. In *International Joint Conference on Artificial Intelligence-Pacific Rim International Conference on Artificial Intelligence*.
- [30] Stefan Kopp, Kirsten Bergmann, and Ipke Wachsmuth. 2008. Multimodal communication from multimodal thinking—towards an integrated model of speech and gesture production. *International Journal of Semantic Computing* 2, 01 (2008), 115–136.
- [31] Staffan Larsson and David R Traum. 2000. Information state and dialogue management in the TRINDI dialogue move engine toolkit. *Natural language engineering* 6, 3-4 (2000), 323–340.
- [32] Hung Le, Doyen Sahoo, Nancy Chen, and Steven Hoi. 2019. Multimodal Transformer Networks for End-to-End Video-Grounded Dialogue Systems. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*. 5612–5623.
- [33] Xiaoyun Lei, Zhian Zhang, and Peifang Dong. 2018. Dynamic path planning of unknown environment based on deep reinforcement learning. *Journal of Robotics* 2018 (2018).
- [34] Pierre Lison. 2012. Probabilistic Dialogue Models with Prior Domain Knowledge. In *SigDial2012*. Association for Computational Linguistics, Seoul, South Korea, 179–188.
- [35] Pierre Lison. 2015. A hybrid approach to dialogue management based on probabilistic rules. *Comput. Speech Lang.* 34, 1 (2015), 232–255.
- [36] Pierre Lison and Casey Kennington. 2016. OpenDial: A Toolkit for Developing Spoken Dialogue Systems with Probabilistic Rules. In *Proceedings of ACL-2016 System Demonstrations*. Association for Computational Linguistics, Berlin, Germany, 67–72. <https://doi.org/10.18653/v1/P16-4012>

- [37] Matthew Marge, Carol Espy-Wilson, Nigel G Ward, Abeer Alwan, Yoav Artzi, Mohit Bansal, Gil Blankenship, Joyce Chai, Hal Daumé III, Debadeepta Dey, et al. 2022. Spoken language interaction with robots: Recommendations for future research. *Computer Speech & Language* 71 (2022), 101255.
- [38] Naresh Marturi, Marek Kopicki, Alireza Rastegarpanah, Vijaykumar Rajasekaran, Maxime Adjigble, Rustam Stolkin, Aleš Leonardis, and Yasemin Bekiroglu. 2019. Dynamic grasp and trajectory planning for moving objects. *Autonomous Robots* 43 (2019), 1241–1256.
- [39] MG Mohanan and Ambuja Salgoankar. 2018. A survey of robotic motion planning in dynamic environments. *Robotics and Autonomous Systems* 100 (2018), 171–185.
- [40] Aishwarya Padmakumar, Jesse Thomason, Ayush Shrivastava, Patrick Lange, Anjali Narayan-Chen, Spandana Gella, Robinson Piramithu, Gokhan Tur, and Dilek Hakkani-Tur. 2021. TEACH: Task-driven Embodied Agents that Chat. (Oct. 2021). arXiv:2110.00534 [cs.CV]
- [41] Baolin Peng, Xiujun Li, Jianfeng Gao, Jingjing Liu, and Kam-Fai Wong. 2018. Deep Dyna-Q: Integrating Planning for Task-Completion Dialogue Policy Learning. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. 2182–2192.
- [42] Baolin Peng, Chenguang Zhu, Chunyuan Li, Xiujun Li, Jinchao Li, Michael Zeng, and Jianfeng Gao. 2020. Few-shot Natural Language Generation for Task-Oriented Dialog. In *Findings of the Association for Computational Linguistics: EMNLP 2020*. 172–182.
- [43] Morgan Quigley, Ken Conley, Brian P. Gerkey, Josh Faust, Tully Foote, Jeremy Leibs, Rob Wheeler, and Andrew Y. Ng. 2009. ROS: An open-source robot operating system. In *Workshops at the IEEE International Conference on Robotics and Automation*.
- [44] Stéphane Ross, Joelle Pineau, Sébastien Paquet, and Brahim Chaib-Draa. 2008. Online planning algorithms for POMDPs. *Journal of Artificial Intelligence Research* 32 (2008), 663–704.
- [45] Milene Santos Teixeira and Mauro Dragoni. 2022. A review of plan-based approaches for dialogue management. *Cognitive Computation* 14, 3 (2022), 1019–1038.
- [46] Edward Schmerling, Karen Leung, Wolf Vollprecht, and Marco Pavone. 2018. Multimodal probabilistic model-based planning for human-robot interaction. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 3399–3406.
- [47] Bharat Singh, Rajesh Kumar, and Vinay Pratap Singh. 2022. Reinforcement learning in robotic applications: a comprehensive survey. *Artificial Intelligence Review* (2022), 1–46.
- [48] Pei-Hao Su, Milica Gašić, and Steve Young. 2018. Reward estimation for dialogue policy optimisation. *Computer Speech & Language* 51 (2018), 24–43.
- [49] Anirudh Sundar and Larry Heck. 2022. Multimodal Conversational AI: A Survey of Datasets and Approaches. In *Proceedings of the 4th Workshop on NLP for Conversational AI*. Association for Computational Linguistics, Dublin, Ireland, 131–147. <https://doi.org/10.18653/v1/2022.nlp4convai-1.12>
- [50] Josue Torres-Foncesca, Catherine Henry, and Casey Kennington. 2022. Symbol and Communicative Grounding through Object Permanence with a Mobile Robot. In *Proceedings of the 23rd Annual Meeting of the Special Interest Group on Discourse and Dialogue*. Association for Computational Linguistics, Edinburgh, UK, 124–134.
- [51] David Traum and Staffan Larsson. 2003. The Information State Approach to Dialogue Management. In *Current and New Directions in Discourse and Dialogue*. Springer, 325–353.
- [52] S M Veres, L Molnar, N K Lincoln, and C P Morice. 2011. Autonomous vehicle control systems — a review of decision making. *Proc Inst Mech Eng Part I J Syst Control Eng* 225, 2 (March 2011), 155–195.
- [53] Petra Wagner, Zofia Malisz, and Stefan Kopp. 2014. Gesture and speech in interaction: An overview. , 209–232 pages.
- [54] Yaqing Wang, Quanming Yao, James T. Kwok, and Lionel M. Ni. 2020. Generalizing from a Few Examples: A Survey on Few-Shot Learning. *Comput. Surveys* 53, 3, Article 63 (jun 2020), 34 pages. <https://doi.org/10.1145/3386252>
- [55] Tsung-Hsien Wen, David Vandyke, Nikola Mrkšić, Milica Gasic, Lina M Rojas Barahona, Pei-Hao Su, Stefan Ultes, and Steve Young. 2017. A Network-based End-to-End Trainable Task-oriented Dialogue System. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 1, Long Papers*. 438–449.
- [56] David Whitney, Eric Rosen, James Macglashan, Lawson L S Wong, and Stefanie Tellex. 2017. Reducing errors in object-fetching interactions through social feedback. In *Proceedings - IEEE International Conference on Robotics and Automation*.
- [57] Jason D Williams, Antoine Raux, and Matthew Henderson. 2016. The dialog state tracking challenge series: A review. *Dialogue & Discourse* 7, 3 (2016), 4–33.
- [58] Yunyi Yang, Yunhao Li, and Xiaojun Quan. 2021. Ubar: Towards fully end-to-end task-oriented dialog system with gpt-2. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 35. 14230–14238.
- [59] Steve Young, Milica Gašić, Simon Keizer, François Mairesse, Jost Schatzmann, Blaise Thomson, and Kai Yu. 2010. The Hidden Information State model: A practical framework for POMDP-based spoken dialogue management. *Comput. Speech Lang.* 24, 2 (2010), 150–174.
- [60] Steve Young, Milica Gašić, Blaise Thomson, and Jason D Williams. 2013. POMDP-based statistical spoken dialog systems: A review. *Proc. IEEE* 101, 5 (2013), 1160–1179.
- [61] Steve Young, Milica Gašić, Blaise Thomson, and Jason D Williams. 2013. POMDP-based statistical spoken dialog systems: A review. *Proc. IEEE* 101, 5 (2013), 1160–1179.
- [62] Haodi Zhang, Zhichao Zeng, Keting Lu, Kaishun Wu, and Shiqi Zhang. 2022. Efficient Dialog Policy Learning by Reasoning with Contextual Knowledge. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 36. 11667–11675.
- [63] Weinan Zhang, Yiming Cui, Yifa Wang, Qingfu Zhu, Lingzhi Li, Lianqiang Zhou, and Ting Liu. 2018. Context-sensitive generation of open-domain conversational responses. In *Proceedings of the 27th international conference on computational linguistics*. 2437–2447.

- [64] Yizhe Zhang, Siqi Sun, Michel Galley, Yen-Chun Chen, Chris Brockett, Xiang Gao, Jianfeng Gao, Jingjing Liu, and William B Dolan. 2020. DIALOGPT: Large-Scale Generative Pre-training for Conversational Response Generation. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*. 270–278.