THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Towards a Taxonomy of Task-Oriented Domains of Dialogue

Citation for published version:

Marques, T 2015, Towards a Taxonomy of Task-Oriented Domains of Dialogue. in Q Chen, P Torroni, S Villata, J Hsu & A Omicini (eds), *PRIMA 2015: Principles and Practice of Multi-Agent Systems: 18th International Conference, Bertinoro, Italy, October 26–30, 2015, Proceedings*. Lecture Notes in Computer Science, vol. 9387, Springer International Publishing, pp. 510-518. https://doi.org/10.1007/978-3-319-25524-8_33

Digital Object Identifier (DOI):

[10.1007/978-3-319-25524-8_33](https://doi.org/10.1007/978-3-319-25524-8_33)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Peer reviewed version

Published In:

PRIMA 2015: Principles and Practice of Multi-Agent Systems

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

Towards a Taxonomy of Task-Oriented Domains of Dialogue

Tânia Marques

School of Informatics
University of Edinburgh
Edinburgh EH8 9AB, United Kingdom
`tmarques@inf.ed.ac.uk`

Abstract. To deal with a broad spectrum of domains, intelligent agents have to generate their own task-oriented dialogue that stems from the need to interact with another agent when solving their own individual task. Most work created to date has either been focused on the task or on the dialogue, but not on both. A taxonomy that describes how the characteristics of a domain determine the types of dialogue needed would be useful, both for understanding how to create agents that are more adaptable to different domains, and also to facilitate reusing previous work. In this paper, we present a number of dimensions that could be included in such a taxonomy, and illustrate how they could be used to determine the nature of dialogue needed in a particular type of domain.

Keywords: Agent Communication, Taxonomy, Task-Oriented Dialogue

1 Introduction

Creating agents who are able to automatically determine how to function in different domains — in other words, agents who adapt their interactions to the domain-level problem they are facing — would be a useful addition to the state of the art models of interaction, which typically are tuned manually to the specific domain of application. This is only possible if the agents communicate with other agents when the task demands it. To date most of the work in the agents community has focused on creating agents that are very good at solving a given task when communication is not needed [1, 2], very good at communicating when recipes to solve the task are given [3, 4], or very good only in very specific settings [5, 6]. It is necessary to bridge the gap between these trends of research to create more flexible agents, capable of dealing with a broader spectrum of domains.

A taxonomy showing how the characteristics of a particular type of domain influence the dialogue needed can be a useful tool. It would help understand what is required of dialogue when building task-oriented domain-independent agents, adaptable to several tasks and situations. It would make explicit the connection between the task and the potential dialogue, while helping grounding the communication to the physical task-domain. Forging the link between task

and the discourse can also be of value for reusing work that was previously done in the literature by understanding how it correlate with the domain task. For instance, if the agents need to perform a joint plan, then work related to argumentation might be relevant for them to exchange arguments about the adequacy of a certain plan.

The purpose of this paper is to discuss the need of a taxonomy that forges a link between the task and dialogue for agent interaction by showing a first attempt to create it, and explaining how this taxonomy could be used. To that end, we describe a number of dimensions that define task-oriented domains in terms of characteristics that might influence the types of dialogue needed in such domains. It should be noted that this is an initial attempt where the simple case of one on one interaction between agents, and a common language is assumed. The dimensions described should not be seen as an exhaustive list of all independent dimensions needed to describe a task-oriented domain. Nor should the taxonomy be seen as the only one that should be used.

This paper is structured as follows: related work is described in Section 2. Section 3 is divided in two parts: we start by presenting a number of dimensions used to characterize task-oriented domains, and then we describe how a taxonomy could be created from the dimensions presented. Some examples illustrating its use can be found in section 4. Section 5 concludes by presenting some limitations of the taxonomy presented here and future work.

2 Related Work

In the literature, several taxonomies were proposed for the classification of speech acts [7, 8]. These were intended for describing human language and are unnecessarily expressive for multi-agent systems. The Foundation for Intelligent Physical Agents (FIPA) [9] uses a small subset of those speech acts in the ACL language such as *request* and *inform*. Their focus was on defining a minimal subset of generic models of communication semantics, rather than on how the speech acts can actually be used based on the task being performed. Our taxonomy does not aim to classify speech acts or the minimal set needed to communicate like the previous works do. Instead, it attempts to use the existing classifications and the task being performed to help identify the types of speech acts that might be needed by a particular agent.

In the area of argumentation, Walton and Krabbe [10] have provided a taxonomy to determine which type of dialogue should be used by identifying the goal behind it (see Figure 1). Their dialogue typology is composed of six formal models:

- **Information-Seeking:** agents ask for information from the other party;
- **Inquiry:** the two parties attempt to answer a question whose answer is unknown by both, but may be answered with their joint knowledge;
- **Persuasion:** one party attempts to change the other party’s beliefs;
- **Negotiation:** the participants bargain over the allocation of resources;

- **Deliberation:** the parties collaborate to know which course of action should be done in which situation;
- **Eristic:** the participants quarrel verbally to vent grievances.

We are interested in understanding how a similar taxonomy can be created for task-oriented agents where communication is not defined by the goal of the dialogue, but it is instead determined by the characteristics of the domain and the task to be performed.

Fig. 1. A categorization to determine the type dialogue, according to Walton and Krabbe pp. 81 [10].

3 Taxonomy

Creating a taxonomy that tells us which type of dialogue is going to be needed by a particular agent can only be achieved if we understand the task and domain where the agent operates. In this section, we look at some dimensions that can be found in the literature, and attempt to understand how they could be used to develop a taxonomy of task-oriented domains of dialogue.

3.1 Finding Dimensions of Task-Oriented Domains

In this section, we present a number of dimensions that characterize task-oriented domains. This allows us to identify how domains determine the communication that is performed by an agent, and consequently, it also tell us something about the implementation requirements behind such an agent. This list is not exhaustive and there is not enough evidence to say that these dimensions are totally independent. In spite of that, we believe that it is representative of the type of dimensions that must be found to create a taxonomy of task-oriented domains of dialogue.

Observability: There is a great incentive to use communication in domains with partial observability and/or incomplete information. Planning for decentralized agents with partial information of the world is significantly more complex than planning for a single agent [11], but if information is shared between agents, the number of possible belief-states is reduced and the problem becomes equivalent to single-agent planning. Cooperative map acquisition by agents is another example where sharing information has been shown to be beneficial [12] to reduce the time needed for acquiring a plausible map. Lack of information is also an issue when the goal of the agents is finding an equilibrium strategy, where no agent could gain from changing strategy, because it might not be possible to find any equilibrium with the known information. In such a scenario, the agents' best choice would be to announce private information [13].

The communication mechanisms in partial observable domains can usually be seen as information-seeking. The questions, however, will depend on the nature of the domain. If the agent needs to acquire knowledge concerning the current state, then this can be seen as a transmission of actions and/or observations using speech acts such as *inform* or *assert*. When the agent seeks to know the likely outcome of the action, then its questions will possibly be requests of the information known by the other agent regarding that action. In the particular case where the questions stem from not knowing the opponent, then the agents will most likely ask about the other's intentions. If it is not possible to confirm the information given by the other agents, then information-seeking will only make sense if the agents are sincere, and sincerity is a rational principle only when the preference functions of the agents align [14].

Types of Actions: The actions that an agent can perform in a determined setting will influence whether dialogue is needed. If there is a finite amount of resources, the agent might want to ask for some of the resources from the other agent (e.g. [15]). This exchange requested is an action that the agent is unable to perform by itself. In a similar fashion, there may also exist actions that the other agent is able to prevent this agent from executing. For example, if the other agent blocks this agent's path. These actions which are naturally dependent of the interaction of the agents are called public in multi-agent planning. They include joint actions, where agents need to synchronize to perform them simultaneously. In contrast, private actions cannot be influenced by other agents and are thus independent of the interaction [16].

Succinctly, we can consider three types of actions: actions that need to be performed by the other agent — usually actions that have to be requested to be executed or avoided — to allow this agent to perform a certain action, private actions that the agent can perform by itself, and joint actions where both agents need to perform it simultaneously. Depending on the action's type, the agents may decide to jointly plan their activity by using a dialogue of deliberation, or may decide to be more competitive and use a process of negotiation [17].

Preferences over costs: Communication is not necessarily advantageous in all domains, and it may even be better to avoid it, especially when the cost incurred is too high or the available bandwidth is very restricted. The solution may be to be selective in the type of messages that are exchanged, as for example in [18], where the agents only send the messages that are the most valuable for team performance, or it could even be finding solutions that do not involve communication such as the one proposed in [19], where agents use deduction based on sensory information. However, there are also cases where the cost of the task exceeds the communication cost. For example, in domains presented in [21] where the agents can reduce the walked distance if they exchanged their individual tasks amongst themselves. Depending on the domain, the agent will have a stronger preference over minimizing the cost of the communication or the cost of the task.

Available Interactions: Dialogue is also influenced by the interactions that can be performed between agents in a certain domain. In [20] Tan discusses three types of interaction in cooperative settings: sharing observations, sharing actions that happened, and sharing learned policies or plans. In other settings, exchange of resources (e.g. [15]), or exchange of tasks might also be available (e.g. [21]). The availability of resources is usually scarce and limited, thusly the exchange of resources is more likely to be competitive and require negotiation. On the other hand, the sharing of observations or plans if under an assumption of sincerity is less likely to be competitive, leading to deliberative dialogues.

3.2 Creating a Taxonomy

In section 3.1, we identified four dimensions of task-oriented domains that seem to influence the existence of dialogue and its characteristics: observability, types of actions, preferences over costs, and available interactions. Now, we need to understand how they relate to specific types of dialogue. The categorization of dialogues proposed by Walton and Krabbe [10] mainly focuses on argumentative settings, but can also be used for a broader range of domains due to the heterogeneity of the categories. Most of the speech acts proposed by FIPA (Foundation for Intelligent Physical Agents) in the ACL language [9], the standard language for agent communication, can also be mapped into these categories. For instance some of its speech acts such as *propose* correspond to a dialogue of negotiation, and *query* can be seen in information-seeking dialogue. Therefore, it seems to be justifiable to use this categorization to create an initial model of how a taxonomy for task-oriented domains of dialogue could look like.

From the literature presented, we can see a prevalence of: (a) information-seeking, corresponding to information sharing; (b) negotiation in competitive settings; and (c) deliberation when performing joint tasks. This does not mean that other types of dialogue are not relevant. Negotiation benefits from persuasion and inquiry is a type of information seeking. However, for now, we will focus on the information-seeking, negotiation and deliberation. These categories of dialogue are presented in table 1, along with the a set of speech acts from

Table 1. Types of dialogue used in the taxonomy

Type of Dialogue	FIPA Speech Acts	Benefit	Fields of Interest
Negotiation	propose; accept-proposal; reject-proposal; call-proposals	Reach agreement; Get the best deal to oneself	Automated negotiation; Argumentation; Preferences handling
Deliberation	request; agree; refuse	Reach agreement; Build a joint plan	Argumentation; Shared plans generation; Automated planning
Information-Seeking	inform; query; confirm; disconfirm	Share information; Common understanding of the world	Knowledge representation; Automated reasoning; Belief revision
No Dialogue	—	Avoid cost of dialogue	Centralized multi-agent planning; Offline optimization

FIPA associated with each category, the benefit for the agent of using it, and some research fields that might be relevant for creating multi-agent system with that type of dialogue.

Considering the types of dialogue and the dimensions presented, we inferred the relations illustrated in Figure 2. Our reasoning is that an agent might not need to communicate when it has full observability and there are no actions that it cannot perform by itself or that are influenced by other agents. Even so, the agent might decide to communicate if the cost of the task could be decreased by doing so. This is usually possible when an agent may exchange tasks or resources, otherwise the agent will not benefit from the interaction. There are two distinct cases to consider when communication is needed: the case with partial observability and the one with full observability. In the former, the agent’s priority will usually be to exchange information if that is possible (e.g. observations, outcomes, intentions). In both cases, the agents would benefit from interacting when performing a joint task. If they are able to exchange plans, then the agents might recur to the use of deliberation, or they might as well simply negotiate tasks and resources if it is not possible to reason cooperatively about the plans to be performed. Even when the agents are not performing a joint task, there might be actions or exchanges that are needed or may affect the agent’s goal. This corresponds to the case where the agent needs to communicate, but it is not able to exchange information and it is not performing a joint action. Negotiation of resources or tasks might be suitable in such cases to bargain with the other agent in order to obtain what is needed to reach its goals.

The categories used in the taxonomy are not mutually exclusive and it is possible for agents to require more than one type of dialogue in a certain domain. As stated before, this is only an initial attempt to create a taxonomy for task-oriented domains of dialogue, and more work is needed to fully understand how communication is influenced by the task and the domain. Yet, it is possible to imagine some examples where this taxonomy could be useful. We present two examples of this in the following section.

Fig. 2. A categorization to determine the type dialogue that the agents need to perform taking in consideration the domain they inhabit.

4 Examples

In this section, we present two examples of how the taxonomy presented could be useful for creating multi-agent systems. In the first example, we use the taxonomy to identify the type of dialogue for a given domain. In the second example, we explain how the taxonomy could potentially help in the creation of more flexible agents able to deal with different domains.

Example 1: The most obvious application of the taxonomy is to identify the most adequate type of dialogue for a particular domain. We will exemplify this with two agents whose task is to vacuum a number of rooms separated by walls that are not soundproof. For efficiency purposes, the agents will avoid repeated work. Looking at the taxonomy, it seems that an information-seeking dialogue fits this problem. Communication is needed, because there is partial information (the agents do not see which rooms have been vacuumed), and information exchange is possible (the walls are not soundproof). After knowing the type of dialogue, the concrete speech acts that might be useful can be found in table 1. In this particular case, the speech act inform might be enough for this system. The different types of dialogues are also closely related to certain fields of research. Thusly, it can give us an idea that the work done in knowledge representation and belief revision might be relevant for modeling this multi-agent system.

Example 2: Another potential application of such taxonomy is the creation of more flexible agents that can deal with several domains. Imagine an agent or a set of agents able to generate the different types of dialogue presented in the taxonomy. If the agent was able to identify the dialogue needs from the domain description given, then it could use a subset of algorithms or sub-agents that can deal with such a dialogue. In other words, given a description of a domain and task, the agent would look at the task, and decide if it could solve the task without communication, and if it discovers that it can not, then it would proceed to see if there is partial information in the world, and so on. For this, the agent must have some algorithm to obtain the input needed. For example, if the agent identifies the type of dialogue as negotiation, then it knows that it needs to create a set of deals with the resources available, and generate a dialogue of negotiation. This is a theoretical abstract use of the taxonomy, but it shows that it might be possible to create more flexible communicative agents, if we identify how the domain correlates with the types of dialogue needed in it.

5 Conclusions and Further Work

In this paper, we presented an initial study of how to create a taxonomy where domains are categorized according to a number of dimensions which determine the dialogue required by agents inhabiting those domains. We believe that such a taxonomy can help identify the requirements for creating more flexible agents able to deal with different domains, and it can give insights of how work created in different areas of research correlate with each other. The taxonomy presented is simple, and there is still a lot of work to be done. For example, it is not easy to understand how preferences over the communication can influence the dialogue. In the future, we plan to look at the literature to identify the different techniques used to reduce or avoid communication overhead and how they influence the types of dialogue. We also intend to drop the assumption regarding the agents' honesty. We imagine that these directions will lead to a more accurate taxonomy, where there might be very little or no dialogue when the communication cost is very high or when trust between agents is not possible. Another possible direction consists of analyzing how dialogue is influenced by the absence of a shared language amongst the agents. We predict that this might increase the negotiation required even in domains where this type of dialogue is not common, due to the need of agreeing on a symbol for a particular object when its definition differs amongst agents. It would also be interesting to explore how the social needs of the other agents may also affect the agent's actions in a certain domain.

Acknowledgments. The research presented in this paper has been funded by the European Community's Seventh Framework Programme (FP7/2007-2013) under grant agreement no. 567652 *ESSENCE: Evolution of Shared Semantics in Computational Environments* (<http://www.essence-network.com/>).

References

1. Ghallab, M., Nau, D., Traverso, P.: *Automated Planning: Theory and Practice*. Morgan Kaufmann, San Francisco (2004)
2. Meneguzzi, F., de Silva, L.: Planning in BDI agents: a survey of the integration of planning algorithms and agent reasoning. *The Knowledge Engineering Review*, 30, pp. 1–44 (2015)
3. van Benthem, Johan.: *Argumentation in artificial intelligence*. Eds. Iyad Rahwan, and Guillermo R. Simari. 47, Springer (2009)
4. Jennings, N. R., Faratin, P. et al: Automated negotiation: prospects, methods and challenges. *Group Decision and Negotiation*, 10(2), pp. 199–215 (2001)
5. Necoletta, F., Viganó, F., Colombetti, M.: Agent communication and artificial institutions. *Autonomous Agents and Multi-Agent Systems* 14(2), pp. 121–142 (2007)
6. Goldman, C. V., Zilberstein, S.: Optimizing information exchange in cooperative multi-agent systems. In: the 2nd international joint conference on Autonomous agents and multiagent system. pp. 137–144, ACM (2003)
7. Austin, J. L.: *How To Do Things With Words*. Eds.. J.O. Urmson and Marina Sbisa). Oxford: Oxford University Press (1975).
8. Searle, J. R.: A taxonomy of illocutionary acts. Eds. Gunderson. *Language in Society*, 5(1-23), pp.344–369 (1975)
9. O'Brian, P.D., Nicol, R.C.: FIPA: Towards a Standard for Software Agents. *BT Technology Journal*. 16(3), 51–59 (1998)
10. Walton, D. N., Krabbe, E. C. W.: *Commitment in Dialogue: Basic Concepts of Interpersonal Reasoning*. SUNY Series in Logic and Language. State University of New York Press (1995)
11. Bernstein, D. S., Givan, R. et al: The complexity of decentralized control of Markov decision processes, *Mathematics of Operations Research*, 27(4) pp. 819-840 (2002)
12. López-Sánchez, M., Esteva, F. et al: Map generation by cooperative low-cost robots in structured unknown environments. *Autonomous agents*. Springer (1998)
13. Jackson, M. O., Simon, L. K. et al: Communication and equilibrium in discontinuous games of incomplete information. *Econometrica*, pp. 1711–1740 (2002)
14. Asher, N., Lascarides, A.: *Strategic Conversation, Semantics and Pragmatics*, 6(2), pp. 1–62 (2013)
15. Gal, Y., Grosz, B. et al: Colored Trails: a Formalism for Investigating Decision-Making in Strategic Environments. *Workshop on Reasoning, Representation and Learning in Computer Games*. pp. 25–30. AAAI Press, Menlo Park (2005)
16. Brafman, R. I., Domshlak, C.: From One to Many: Planning for Loosely Coupled Multi-Agent Systems. In: the 8th International Conference on Automated Planning and Scheduling, pp. 28–35 (2008)
17. Doran, J. E., Franklin, S. R. J. N. et al: On cooperation in multi-agent systems. *The Knowledge Engineering Review*, 12(3), pp. 309–314 (1997)
18. Roth, M., Simmons, R., Veloso, M.: What to communicate? Execution-time decision in multi-agent POMDPs. In: *Distributed Autonomous Robotic Systems* 7, pp. 177–186, Springer (2006)
19. Genesereth, M. R., Ginsberg, M. L., Rosenschein, J. S.: *Cooperation without communication, Heuristic Programming Project*, Computer Science Department, pp. 51–57, Stanford University (1984)
20. Tan, M.: Multi-agent reinforcement learning: Independent vs. cooperative agents. In: the 10th international conference on machine learning, pp. 330–337 (1993)
21. Rosenschein, J. S., Zlotkin, G.: *Rules of Encounter: Designing Conventions for Automated Negotiation Among Computers*, MIT Press, Cambridge, MA, USA (1994)