



Minerva Access is the Institutional Repository of The University of Melbourne

Author/s:

Sonenberg, E;Miller, T;Pearce, AR;Felli, P;Muisse, CJ;Dignum, F

Title:

Social planning for social HRI

Date:

2016

Citation:

Sonenberg, E., Miller, T., Pearce, A. R., Felli, P., Muise, C. J. & Dignum, F. (2016). Social planning for social HRI. Baxter, P (Ed.) Trafton, G (Ed.) Lemaignan, S (Ed.) Cognitive Architectures for Social Human-Robot Interaction, arxiv. <https://doi.org/10.48550/arxiv.1602.06483>.

Persistent Link:

<https://hdl.handle.net/11343/297381>

Social planning for social HRI

Liz Sonenberg
Tim Miller
Adrian Pearce

Department of Computing and
Information Systems
University of Melbourne
Australia

Contact: l.sonenberg@unimelb.edu.au

Paolo Felli
University of Nottingham
UK

Christian Muise
MIT, Boston
USA

Frank Dignum
University of Utrecht
The Netherlands

Index Terms—Nested belief reasoning; Theory of Mind reasoning; Multi-agent planning; Epistemic planning.

I. RESEARCH CONTEXT

Making a computational agent ‘social’ has implications for how it perceives itself and the environment in which it is situated, including the ability to recognise the behaviours of others at various levels – simple actions, goals and intentions. Hence fundamental elements of an architecture for social agents must allow for management of social motivations - i.e. to reach social goals, not only practical goals - and must model and account for actions having both practical and social effects. Further it has been argued that to build social agents it is not sufficient to just add a few ‘social modules’ to existing architectures: while multilayer computational cognitive models have been studied for some time, c.f. [1], a new layered deliberation architecture is required that at the higher level(s) naturally accommodates analysis of decision choices that take into account both rich context and future projections of possible consequences, yet does not rely on computational expensive deep reasoning capability [2], [3], [4].

In the work reported here, we do not attempt to address the ‘large’ questions associated with the design of a fully integrated computational cognitive architecture; rather we adopt a relatively narrow focus on exploiting and extending epistemic planning mechanisms to achieve run-time generation of plans in rich multi-actor contexts, i.e. we seek to construct *social plans* in settings where the social context is relevant in the assessment of the beliefs and capabilities of others, and in making appropriate choices of what to do next.

Our approach has been informed by our experience with the BDI model of agency [5] and several associated agent architectures - architectures that were introduced to support a balance of deliberative and reactive behaviours, and that in their instantiation are reliant on domain-specific expert knowledge acquisition to provide a *knowledge level* view [6], c.f. [7], [8]. We are also supporters of the position that logic-based techniques are well suited to represent social reasoning and through which to engineer effective mechanisms, c.f. [9], [10], [11].

Fundamental concepts we build on include: reasoning about the beliefs of others, including their beliefs about others; establishing common ground; and the use of stereotypes. So a few words about each.

Exploiting mutual awareness to enable a participant engaged in collaborative activity with others to select an appropriate action typically involves *Theory of Mind* (ToM) reasoning [12], [13], i.e., reasoning about the knowledge, beliefs, perspectives and reasoning of other participants. Agent-based computational models have been used to investigate higher-order ToM in varied scenarios, including alignment with human performance data in some cases e.g., [14], [15], [16], [17], [18], [19].

A specific element of ToM reasoning is *grounding*, or establishing common ground, i.e. an important mechanism by which participants engaged in joint activity coordinate their respective understandings of matters at hand. This construct arises from a model of conversation developed by Herbert Clark [20] and since studied widely in many fields, including social psychology, e.g. [21], HCI, e.g. [22], philosophy, e.g. [23]. Finding computationally amenable representations and mechanisms that allow agents interacting with humans to keep track of the activity, and their understanding of other participants in the same activity, remains a challenge, c.f. [4], [24]. Exploring alternative definitions of grounding, allowing for subtle and important variations in the notions of knowledge, belief and acceptance, is one aspect we have investigated [25], [26].

To efficiently take action in settings without the forms of full information needed for ToM reasoning, humans often reason in terms of the (reference) groups to which they and others belong, and the role structures and stereotypical behaviours associated with those reference groups. Steps have been taken towards equipping agents with similar computational capabilities, e.g., [27], [28], [29].

Now to *social planning*. Planning research has for some time yielded highly efficient mechanisms for plan synthesis suiting single-agent scenarios. Input to a planner includes descriptions of the world and effects of available actions, the initial state(s) that the world might be in before the plan-execution agent performs any actions, and the desired

objectives, such as achieving a goal or performing a specific task. The output typically consists of either a plan (a sequence of actions for the agent to perform) or a policy (an action to perform for each state). However, such descriptions are often insufficient for agents operating in multi-agent environments. In such environments, a planning agent must consider that other agents have their own actions and mental states, and that these actions and mental states can affect the outcomes and interpretation of its own actions. Thus, such reasoning is inherently a social task.

In environments where an agent needs to plan its interactions with others, computational complexity increases: the actions of the other agents can induce a combinatorial explosion in the number of contingencies to be considered, making both the search space and the solution size exponentially larger, hence demanding novel methods. A recent advance is the development of *epistemic planning* [30]: planning according to the knowledge or belief (and iterated knowledge or belief) of other agents, allowing the specification of a more complex class of planning domains than those mostly concerned with simple facts about the world.

Building on this work and on recent advances in non-deterministic planning, we have made progress towards the challenge of efficient reasoning both *with* and *about* the incomplete, higher-order, and possibly incorrect beliefs of other individuals as part of the planning process, and how we can plan considering the actions of others. Our work involves descriptions and demonstrations-in-use of novel mechanisms for stereotypical and empathetic reasoning, explorations of how this can be used as a theory of mind, and planning while considering the actions of others [27], [31], [32], [33], [34], [35].

II. CHALLENGE SCENARIOS

We offer three scenarios that provide challenging settings for social planning.

SCENARIO 1

This scenario illustrates the need for complex reasoning with others, allowing for possibly limited or faulty perceptions by others of their environment.

Consider a self-driving car and a pedestrian each approaching an intersection. A safe plan for each is to wait for the other to go, resulting in a stalemate. With human participants, such encounters are generally resolved with social cues: e.g. one signalling to the other using a nod of the head or hand signal. In such cases, cues such as establishing eye contact generate a common belief that each party understands who will go first, and each party understands that each understands this, etc. For a self-driving car to achieve similar interactions with a pedestrian, it will need both sophisticated sensing technology (to accurately recognise the nod or hand-signal) and also rich internal computational mechanisms to interpret the signal. However, even physical signals often require social context for their correct interpretation. For example, a young child's inability to correctly assess the belief of others, and therefore, the common belief between themselves and

a driver, mean that the driver must consider this when planning its action, and may behave more cautiously.

SCENARIO 2

This scenario is inspired by the Wumpus Hunt and demands agents engage in strategic and social reasoning. It has been used to demonstrate the power of theory of mind reasoning [27], [31].

The lord of a castle is informed by a peasant that a Wumpus is dwelling in a dungeon nearby. It is known that the Wumpus can be killed by one hunter alone only if asleep; if awake, two hunters are required. The lord then tasks the peasant to go to fetch the White Knight, his loyal champion, and hunt down the beast together. The White Knight is known for being irreprehensible, trustworthy and brave; however, the peasant does not know any knight, and neither their looks. While looking for the White Knight, he runs into the Black Knight and, believing him the White Knight, tells him about the quest.

There is some additional information that needs to be taken into account: on one hand, the knight knows how a Wumpus can be killed by two hunters, but he is aware that a simple peasant may get scared by the thought of confronting an awake Wumpus. Also, the peasant can not hunt and is unable to see whether the Wumpus is awake (he can not approach unnoticed), but the knight can. Therefore it is not clear to him whether the peasant can be of any help to the quest. On the other hand, the knight is aware of the misunderstanding: he knows that the peasant attributes to him all the good qualities of the White Knight, so the peasant is confident that the knight won't put him in danger whenever possible. While on the road, they agree on a protocol: they will enter the dungeon from two sides, and the Knight will use a whistle to signal whether the Wumpus is awake, then they will attack.

SCENARIO 3

A more difficult challenge problem can be found with the multi-player board game of deception and bluff, *Hattari* [36].

Hattari involves a crime scene, three suspects, one victim, and clues. The task is to guess who is the culprit, to accuse him or to deceive the other players! Each player receives a "suspect profile" and 5 accusation markers. Three suspect profiles are placed upright in the center of the table, and one profile is placed face down, next to the other three. That is the victim of the crime. The goal is to unmask the the culprit among the three standing suspects. The rules of the game involve selective sharing of information, but also manipulation of incomplete information among the players, through passing around of pieces as players take turns.



Although we have incorporated some of our research on epistemic planning into a (limited) implementation of Hattari [37], creation of an artificial player that could participate meaningfully in a game where humans exploit and interpret body language as they navigate the possibilities of bluffing and deception seems far beyond current technologies.

III. WORKSHOP DISCUSSION QUESTIONS

- 1) *Why should you use cognitive architectures - how would they benefit your research as a theoretical framework, a tool and/or a methodology?*

Our interest is directly in the design of cognitive architectures as the basis for executable strategic collaboration and teamwork involving hybrid human-agent teams.

- 2) *Should cognitive architectures for social interaction be inspired and/or limited by models of human cognition?*

Cognitive architectures should be inspired by models of human cognition. Modelling the cognitive architecture after concepts of human cognition seems to allow us to better prepare agents for human-agent interaction. Further, while explorations with computational models cannot directly shed light on human cognition, c.f. [38], experiments with computational cognitive models can contribute to analyses of potential building blocks for mechanisms involved in coordination in joint action, whether it be in purely human, or human-robot interaction contexts.

- 3) *What are the functional requirements for a cognitive architecture to support social interaction?*

Too many to enumerate here... But, as mentioned above, a cognitive architecture should at least have components modeling the (social) identities, social context and social triggers and effects of actions. In short, representations of the social reality of the partners in the interaction are required.

- 4) *How the requirements for social interaction would inform your choice of the fundamental computational structures of the architecture (e.g. symbolic, sub-symbolic, hybrid, ...)?*

Computational structures should be hybrid. For low level interactions and time constrained feedback loops, some very efficient and robust mechanisms are needed that seem best to be represented sub-symbolically. However, for longer term social actions, it is necessary to have symbolic representations in order to deliberate, on the run, about the (social) effects of actions.

- 5) *What is the primary outstanding challenge in developing and/or applying cognitive architectures to social HRI systems?*

Outstanding challenges include: identifying and exploiting 'sweet spots' in the expressivity-efficiency tradeoff in the engineering of computational artefacts; finding an effective (domain specific) balance between design-time knowledge engineering and run-time learning; signalling of state (in

both directions) between human and artificial participants in joint activity; integration of the diverse perceptual, cognitive and social aspects in a plausibly effective system; establishment of metrics and evaluation methods that allow terms such as "plausibly effective" to be precisely defined and formally demonstrated.

- 6) *Devise a social interaction scenario that current cognitive architectures would likely fail, and why.*

The beginnings of candidate scenarios are offered above. To provoke failure, what is needed are scenarios exhibiting *social brittleness* - i.e. where the normal course of interaction fails due to different expectations or assumptions as a result of different social understandings, and a repair has to be found.

IV. FINAL REMARKS

Even though our focus is on cognitive mechanisms as essential components of an integrated cognitive architecture for effective social robots, and we have some exploratory work on human communication patterns [39], we recognise there are many topics important in such architectures that we do not attempt to address - spatial reasoning [40], dialogue actions [7], multimodal inputs [41], action signalling [24], the link between perception and action [42], [43], and comparisons between logic-based reasoning and other approaches such as game theory [44] and probabilistic reasoning [45] ... to name but a few!!

ACKNOWLEDGEMENTS

Much of the work reported here was carried out while two of the authors (Felli & Muise) were employed by the University of Melbourne with the financial support of the Australian Research Council Discovery Projects Grant DP130102825 *Foundations of Human-Agent Collaboration: Situation-Relevant Information Sharing*. Additional information about the project can be found at <http://agentlab.cis.unimelb.edu.au/project-hac.html>

REFERENCES

- [1] P. Thagard, "Cognitive architectures," *The Cambridge handbook of cognitive science*, pp. 50–70, 2012.
- [2] F. Dignum, R. Prada, and G. J. Hofstede, "From autistic to social agents," in *Proceedings of the 2014 International Conference on Autonomous Agents and Multi-agent Systems*, ser. AAMAS '14, 2014, pp. 1161–1164.
- [3] G. A. Kaminka, "Curing robot autism: A challenge," in *Proceedings of the 2013 International Conference on Autonomous Agents and Multi-agent Systems*, ser. AAMAS '13, 2013, pp. 801–804.
- [4] E. Pacherie, "Intentional joint agency: shared intention lite," *Synthese*, vol. 190, no. 10, pp. 1817–1839, 2013. [Online]. Available: <http://dx.doi.org/10.1007/s11229-013-0263-7>
- [5] M. Georgeff, B. Pell, M. Pollack, M. Tambe, and M. Wooldridge, "The Belief-Desire-Intention model of agency," in *Intelligent Agents V: Agents Theories, Architectures, and Languages*, ser. Lecture Notes in Computer Science, J. Miller, A. Rao, and M. Singh, Eds. Springer, 1999, vol. 1555, pp. 1–10.
- [6] A. Newell, "The Knowledge Level," *Artificial Intelligence*, vol. 18, no. 1, pp. 87–127, 1982.
- [7] S. Lemaignan and R. Alami, "Explicit knowledge and the deliberative layer: Lessons learned," in *Intelligent Robots and Systems (IROS), 2013 IEEE/RSJ International Conference on*, 2013, pp. 5700–5707.

- [8] E. Norling, "What should the agent know?: the challenge of capturing human knowledge," in *Proceedings of the 7th International Joint Conference on Autonomous Agents and Multiagent Systems*, L. Padgham and D. Parkes, Eds., 2008, pp. 1225–1228.
- [9] F. Dignum and L. Sonenberg, "A dialogical argument for the usefulness of logic in MAS," in *Journal of Artificial Societies and Social Simulation*, vol. 7, no. 4, 2004, <http://jasss.soc.surrey.ac.uk/7/4/8/logic-in-abss/context.html>.
- [10] B. Edmonds, "Comments on a dialogical argument for the usefulness of logic in MAS," in *Journal of Artificial Societies and Social Simulation*, vol. 7, no. 4, 2004, <http://jasss.soc.surrey.ac.uk/7/4/8/logic-in-abss/context.html>.
- [11] W. Reich, "Reasoning about other agents: a plea for logic-based methods," in *Journal of Artificial Societies and Social Simulation*, vol. 7, no. 4, 2004, <http://jasss.soc.surrey.ac.uk/7/4/4.html>.
- [12] A. I. Goldman, "Theory of mind," in *The Oxford Handbook of Philosophy of Cognitive Science*, E. Margolis, R. Samuels, and S. P. Stich, Eds. Oxford University Press, 2012, ch. 17.
- [13] D. Premack and G. Woodruff, "Does the chimpanzee have a theory of mind?" *Behavioral and Brain Sciences*, vol. 1, pp. 515–526, 12 1978.
- [14] H. de Weerd, R. Verbrugge, and B. Verheij, "How much does it help to know what she knows you know? an agent-based simulation study," *Artif. Intell.*, vol. 199, pp. 67–92, 2013.
- [15] —, "Negotiating with other minds: the role of recursive theory of mind in negotiation with incomplete information," *Autonomous Agents and Multi-Agent Systems*, pp. 1–38, 2015.
- [16] S. G. Ficci and A. Pfeffer, "Modeling how humans reason about others with partial information," in *Proceedings of the 2008 International Conference on Autonomous Agents and Multi-agent Systems*, 2008, pp. 315–322.
- [17] L. M. Hiatt and J. G. Trafton, "Understanding second-order theory of mind," in *Proceedings of the Tenth Annual ACM/IEEE International Conference on Human-Robot Interaction, March 2015*. ACM, 2015, pp. 167–168.
- [18] B. F. Malle, "Social robots and the tree of social cognition," in *HRI Workshop Proceedings – Cognition: A Bridge between Robotics and Interaction*, 2015, pp. 13–14.
- [19] S. Thill and T. Ziemke, "Interaction as a bridge between cognition and robotics," in *HRI Workshop Proceedings – Cognition: A Bridge between Robotics and Interaction*, 2015, pp. 25–31.
- [20] D. Wilkes-Gibbs and H. H. Clark, "Coordinating beliefs in conversation," *Journal of Memory and Language*, vol. 31, no. 2, pp. 183 – 194, 1992.
- [21] Y. Kashima, O. Klein, and A. E. Clark, "Grounding: Sharing information in social interaction," in *Social Communication*, K. Fiedler, Ed. Psychology Press, 2007, pp. 27–77.
- [22] D. Brock and J. Trafton, "Cognitive representation of common ground in user interfaces," in *UM99 User Modeling*, ser. CISM International Centre for Mechanical Sciences, J. Kay, Ed. Springer Vienna, 1999, vol. 407, pp. 287–289.
- [23] K. Allan, "What is Common Ground?" in *Perspectives on Linguistic Pragmatics*, A. Capone, F. Lo Piparo, and M. Carapezza, Eds. Springer, 2013, ch. 11, pp. 285–310.
- [24] C. Vesper, "How to support action prediction: Evidence from human coordination tasks," in *The 23rd IEEE International Symposium on Robot and Human Interactive Communication, RO-MAN*, 2014, pp. 655–659.
- [25] T. Miller, J. Pfau, L. Sonenberg, and Y. Kashima, "Logics of common ground," 2016, under review. Preliminary version available at <http://people.eng.unimelb.edu.au/tmiller/pubs/logics-of-common-ground.pdf>.
- [26] J. Pfau, T. Miller, and L. Sonenberg, "Modelling and using common ground in human-agent collaboration during spacecraft operations," in *Proceedings of SpaceOps 2014 Conference*. American Institute of Aeronautics and Astronautics, 2014, pp. 1–15.
- [27] P. Felli, T. Miller, C. Muise, A. R. Pearce, and L. Sonenberg, "Computing social behaviours using agent models," in *International Joint Conference on Artificial Intelligence, IJCAI*, 2015.
- [28] J. Pfau, Y. Kashima, and L. Sonenberg, "Towards agent-based models of cultural dynamics: A case of stereotypes," in *Perspectives on Culture and Agent-based Simulations*. P, 2014, pp. 129–147.
- [29] B. G. Silverman, D. Pietrocola, B. Nye, N. Weyer, O. Osin, D. Johnson, and R. Weaver, "Rich socio-cognitive agents for immersive training environments: Case of NonKin Village," *Autonomous Agents and Multi-Agent Systems*, vol. 24, no. 2, pp. 312–343, 2012.
- [30] T. Bolander and M. B. Andersen, "Epistemic planning for single- and multi-agent systems," *Journal of Applied Non-Classical Logics*, vol. 21, no. 1, pp. 9–34, 2011.
- [31] P. Felli, T. Miller, C. Muise, A. Pearce, and L. Sonenberg, "Artificial social reasoning: Computational mechanisms for reasoning about others," in *Social Robotics - 6th International Conference, ICSR 2014. Proceedings*, ser. Lecture Notes in Computer Science, M. Beetz, B. Johnston, and M. Williams, Eds., vol. 8755. Springer, 2014, pp. 146–155.
- [32] T. Miller, C. Muise, P. Felli, A. R. Pearce, and L. Sonenberg, "Knowing whether" in proper epistemic knowledge bases," in *The 30th AAAI Conference on Artificial Intelligence*, 2016.
- [33] C. Muise, F. Dignum, P. Felli, T. Miller, A. R. Pearce, and L. Sonenberg, "Towards team formation via automated planning," in *International Workshop on Coordination, Organisation, Institutions and Norms in Multi-Agent Systems*, 2015.
- [34] C. Muise, V. Belle, P. Felli, S. McIlraith, T. Miller, A. R. Pearce, and L. Sonenberg, "Planning over multi-agent epistemic states: A classical planning approach," in *Proceedings of 29th AAAI Conference on Artificial Intelligence (AAAI)*, B. Bonet and S. Koenig, Eds., 2015.
- [35] C. Muise, P. Felli, T. Miller, A. R. Pearce, and L. Sonenberg, "Leveraging FOND planning technology to solve multi-agent planning problems," *Distributed and Multi-Agent Planning (DMAP-15)*, pp. 83–90.
- [36] MissMerc007, "How to play Hattari," 2013, accessed Feb. 2016. [Online]. Available: <https://youtu.be/CbyMYCiQ79I>
- [37] C. Muise, "Hattari demo," 2015, accessed Feb. 2016. [Online]. Available: <http://hattari.haz.ca/>
- [38] R. Sun, "Theoretical status of computational cognitive modeling," *Cognitive Systems Research*, vol. 10, no. 2, pp. 124–140, 2009.
- [39] A. Butchibabu, J. Shah, and L. Sonenberg, "Implicit coordination strategies for effective team communication," *Human Factors*, 2016, in press - accepted Feb 12, 2016.
- [40] M. Warnier, J. Guitton, S. Lemaignan, and R. Alami, "When the robot puts itself in your shoes: Managing and exploiting human and robot beliefs," in *RO-MAN*. IEEE, 2012, pp. 948–954.
- [41] M. Sridharan, "Integrating visual learning and hierarchical planning for autonomy in human-robot collaboration," in *AAAI Spring Symposium on Designing Intelligent Robots: Reintegrating AI II*, 2013.
- [42] P. E. Baxter, J. de Greeff, and T. Belpaeme, "Cognitive architecture for human-robot interaction: towards behavioural alignment," *Biologically Inspired Cognitive Architectures*, vol. 6, pp. 30–39, 2013.
- [43] J. Pfau, Y. Kashima, and L. Sonenberg, "A two-level computational architecture for modeling human joint action," in *Proceedings of 13th International Conference on Cognitive Modelling (ICCM)*, N. Taatgen, M. van Vugt, J. Borst, and K. Mehlhorn, Eds., 2015, pp. 1–6.
- [44] A. S. Goodie, P. Doshi, and D. L. Young, "Levels of theory-of-mind reasoning in competitive games," *Journal of Behavioral Decision Making*, vol. 25, no. 1, pp. 95–108, 2012.
- [45] A. Stuhlmüller and N. Goodman, "Reasoning about reasoning by nested conditioning: Modeling theory of mind with probabilistic programs," *Cognitive Systems Research*, vol. 28, pp. 80–99, 2014.