Normative Multi-Agent Systems and Human-Robot Interaction

Stephen Cranefield[1] and Bastin Tony Roy Savarimuthu[1]

***Abstract*— This position paper provides an overview of the study of social norms in the *normative multi-agent systems* (NorMAS) community, and presents avenues for cross-fertilisation between the NorMAS and social robotics communities.**

I. INTRODUCTION

For the last three decades, researchers in the field of Normative Multi-Agent Systems (NorMAS) have studied how the concept of norms from human society can be adapted, modelled and incorporated into computational mechanisms to promote social order in societies of software agents. Initially, this endeavour was focused largely on open systems of autonomous software agents, but as human communication has become increasingly mediated by computers, the field has begin to consider how NorMAS reasoning mechanisms can be used to enable socially aware interaction within societies comprising both humans and software agents. However, the field has largely not considered the specific requirements of human-robot interaction.

This position paper reviews the concept of norms and norm-aware agents as conceptualised by NorMAS researchers, and considers some possible areas for cross-fertilisation between this field and human-robot interaction.

II. CONCEPTUALISATIONS OF NORMS

A range of models and representations of norms have been proposed in the NorMAS literature. Norm languages based on deontic logic are common [1], [2], allowing norms of obligation, prohibition and (sometimes) permission to be expressed logically, often with extra features such as conditions, deadlines and sanctions. Norm representations based on temporal logic [3], probabilistic logic progamming [4] and event sequences [5], [6] have also been proposed.

In contrast, simulation studies on the emergence of norms and the effects of sanctions on norm compliance often adopt game theory style models, where sets of numerical parameters represent strategies for specific social dilemmas [7], [8].

In recent years, multi-agent reinforcement learning approaches have also been adapted to enable the learning of socially beneficial rather than selfish behaviours [9], [10]. These are represented by *policies* mapping states to actions.

III. NORM-AWARE AGENTS AND SOCIETIES

Agents that are norm-aware should be able to identify existing norms, and to plan and choose their actions given knowledge of these norms. This includes understanding when their actions may fulfill or violate these norms. Note that as agents are usually considered to be automonous, an agent can choose to violate a norm and risk a sanction if it is better off to do so. In this section, we highlight a few of the research questions that have been addressed by researchers in the field of normative multi-agent systems.

[1]Stephen Cranefield and Tony Savarimuthu are in the Department of Information Science, University of Otago, Dunedin, New Zealand {stephen.cranefield,tony.savarimuthu}@otago.ac.nz

(a) ***How do agents come to know about norms?*** We consider three possible answers (that are not mutually exclusive):

(i) Norms may be created and published or broadcast by an informed and empowered designer (a human, an institution or a software agent [2], [11]–[13]). Human design is only feasible when norms are static. The field of *norm synthesis* [2] considers how software agents can monitor a society, detect undesirable interference between its members, and generate new norms or adapt existing ones to discourage these conflicts. However, it seems unlikely that human members of an agent society would automatically accept norms imposed on them, and such mechanisms would need to be combined with social choice mechanisms to recognise the humans' individual sense of agency.

(ii) Norms may be learned from observation and experience [14]. Work on learning symbolically represented norms has used a range of learning mechanisms, including frequent episode data mining [5], [6], plan recognition [15], probabilistic inference using Bayesian [3], [16] and Dempster Shafer [17] approaches, and probabilistic inductive logic programming [4] (we note that the last two works are from researchers in the fields of human-robot interaction and social robotics).
Evidence for the existence of norms may come from recognising *signalling actions* that indicate the application of a reward or sanction (these could be expressions of approval or disapproval or more overt reactions). For example, the frequent episode mining approach can identify prohibition norms that are the most frequent sequences of actions followed by a negative signalling action [6]. However, these are not the only possible forms of evidence. When agents' goals and their possible plans (at least for publicly observable behaviour) can be inferred, plans that are seldom followed can reinforce obligation and prohibition norm hypotheses that would explain the selection of alternative plans [15]. A Bayesian approach allows both forms of evidence to be combined [3], and could easily accommodate additional types of evidence such as advice about norms from other agents, suitably moderated by some measure of the advising agent's trustworthiness [18], [19]. However, we believe that evidence from observing

signalling actions has a special role in gaining confidence that an identified norm represents truly *normative* rather than merely *normal* behaviour [20].

(iii) Norms may be proposed by a *norm entrepreneur* and subsequently spread through a majority of the society. While this process has been studied at an abstract level by researchers in the field of international studies [21]–[23], there appears to be very little prior work on computational mechanisms for norm entrepreneurship [24].

(b) ***What is the lifecycle of dynamic social norms, and how can agents track their status?*** Several norm lifecycle models (with minor variations) have been proposed in the NorMAS community over the years, and an overview of such works can be found in the recent work of Morris-Martin et al. [25]. The lifecycle models, in general, describe how a norm is proposed, propagated (or spread), eventually adopted and then may possibly lose relevance in an agent society. The propagation step may involve a variety of mechanisms such as spreading of norms through explicit communication, applying rewards for compliance and/or sanctions for violations, or copying the observed behaviour of other agents, especially successful ones [5], [6]. A norm may become obsolete due to losing salience to current conditions or changes to the goals and/or membership of the society.

For example, researchers have proposed a norm-recommendation system [26] based on tracking the status of the norm in a community to recommend whether an agent (e.g., a robot) should follow or violate a norm based on factors such as the life-stage of a norm (e.g., emergent vs. mature), its uptake (a waxing or waning norm) and the severity of sanctions [27].

(c) ***How does knowledge of norms interact with other reasoning processes, such as goal creation and plan selection?*** In multi-agent systems, agents are often conceptualised in terms of the belief-desire-intention (BDI) practical reasoning architecture [28]. A BDI agent is considered to have *goals*, *plans* that are indexed by the goals they can achieve and the contexts they apply to, and *intentions*: the plan instantiations the agent is currently committed to (given that resources are finite, and focused effort is often needed to make progress towards a goal). Researchers have developed agent architectures such as n-BDI [29] and N-Jason [30] that consider norms as an important construct in the reasoning cycle along with beliefs, desires and intentions. A norm-aware BDI agent employs norm deliberation during goal creation and plan selection, i.e., it adopts goals and plans to satisfy obligations or avoid prohibited actions.

Knowledge of norms can also allow agents to adopt more efficient plans of action, under the assumption that some or all other agents will follow the norms. This assumption may be justified by monitoring the compliance of other agents [31], by the existence of robust and consistent sanctioning mechanisms, or by maintaining information about the trustworthiness of other agents [18], [19]. However, the connection between these mechanisms and plan choice in BDI agents has not gained much attention.

IV. CROSS-FERTILISATION WITH ROBOTICS

This section identifies five avenues for cross-fertilisation between NorMAS and social robotics.

First, most NorMAS research is simulation-based. Therefore, symbolic representations of the physical and social state of the world are easily obtained and there are no real-time demands on reasoning. In contrast, human-robot interaction involves creating knowledge from sensor data, and is likely to require both high level symbolic and sub-symbolic real-time reasoning for safe operation. Research on human-robot interaction will identify more computationally demanding use cases for normative reasoning that challenge the direct application of existing NorMAS techniques.

Second, to improve situated norm awareness of robots in human-robot teams, researchers can adopt or adapt normative architectures such as n-BDI and N-Jason that consider norms as top-level entities that influence agents' intentions and choice of plans, as outlined in Section III. While robots may have some planning requirements that differ from those of traditional BDI agents (e.g., path planning), addressing these by extending the existing norm-aware practical reasoning theories, architectures and software platforms should provide a faster path to developing norm-aware social robots with declarative goals and plans. These approaches would also facilitate communication with human partners in terms of these high-level cognitive concepts that fit well with human understanding of practical reasoning [32].

Third, norm conflict identification and resolution has seldom been addressed in human-robot collaborations. For example, a robot following a norm it acquired in one context may, in another, run into conflicts with humans or other robots. Works in NorMAS on these areas (e.g., [33]) hold promise to be applied in robotic systems.

Fourth, robots could be active partners in norm entrepreneurship within human-robot teams. Norm-capable robots could be norm entrepreneurs by proposing new or improved norms to their human partners. The techniques used in norm synthesis to avoid undesirable world states or agent interactions could be adapted for use in a peer-to-peer partnership model. Robots could also assist human norm entrepreneurs to propagate (suitably justified) norms by exemplifying them and explaining them to others. In both cases, new mechanisms would be needed to explain the purpose and benefits of newly proposed norms or modifications to old norms. For robot-generated norms and explanations to be effective, it may be necessary for the robots to explictly consider the humans' mental models of the task and robot capabilities [34].

Fifth, robots are likely to require fast non-symbolic reasoning when interacting physically. Thus, there is a tension between the representations needed for robot action learning and selection and those used in traditional NorMAS reasoning. This distinction is similar to the contrast between System 1 and System 2 thinking in humans (as studied in the work of Kahneman [35] and considered in the context of artificial intelligence by Booch et al. [36]). Robotics offers a promising avenue to explore the exchange of normative representations between these two types of reasoning. One challenge is to bridge the gap between the state-to-action mappings ("policies") learned via reinforcement learning (commonly applied in robotics) and the symbolic norm expressions used in NorMAS approaches, especially in the presence of norms involving temporal patterns of behaviour. While deep reinforcement learning using recurrent neural networks can model agent states that depend on past events [37], we are not aware of existing techniques to map between the resulting policies and symbolic norm expressions.

We believe the research avenues described above can aid towards the creation of norm-aware robotic systems.

References

[1] F. Dignum, "Autonomous agents with norms," *Artificial Intelligence and Law*, vol. 7, pp. 69–79, 1999.

[2] J. Morales, M. López-Sánchez, J. A. Rodríguez-Aguilar, M. J. Wooldridge, and W. W. Vasconcelos, "Automated synthesis of normative systems," in *Proceedings of the 12th International Conference on Autonomous Agents and Multi-Agent Systems*. IFAAMAS, 2013, pp. 483–490.

[3] S. Cranefield, F. Meneguzzi, N. Oren, and B. T. R. Savarimuthu, "A Bayesian approach to norm identification," in *22nd European Conference on Artificial Intelligence*. IOS Press, 2016, pp. 622–629.

[4] Z.-X. Tan, J. Brawer, and B. Scassellati, "That's mine! learning ownership relations and norms for robots," in *Proceedings of the Thirty-Third AAAI Conference on Artificial Intelligence*. AAAI Press, 2019, pp. 8058–8065.

[5] B. T. R. Savarimuthu, S. Cranefield, M. Purvis, and M. K. Purvis, "Obligation norm identification in agent societies," *Journal of Artificial Societies and Social Simulation*, vol. 13, no. 4, 2010.

[6] ——, "Identifying prohibition norms in agent societies," *Artificial Intelligence and Law*, vol. 21, no. 1, pp. 1–46, 2013.

[7] ——, "Role model based mechanism for norm emergence in artificial agent societies," in *Coordination, Organizations, Institutions, and Norms in Agent Systems III*, ser. Lecture Notes in Computer Science, vol. 4870. Springer, 2007, pp. 203–217.

[8] S. Sen and S. Airiau, "Emergence of norms through social learning," in *Proceedings of the 20th International Joint Conference on Artificial Intelligence*, 2007, pp. 1507–1512.

[9] Y. F. Chen, M. Everett, M. Liu, and J. P. How, "Socially aware motion planning with deep reinforcement learning," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2017, pp. 1343–1350.

[10] A. L. Bazzan, "Aligning individual and collective welfare in complex socio-technical systems by combining metaheuristics and reinforcement learning," *Engineering Applications of Artificial Intelligence*, vol. 79, pp. 23–33, 2019.

[11] J. Campos, M. López-Sánchez, and M. Esteva, "A case-based reasoning approach for norm adaptation," in *International Conference on Hybrid Artificial Intelligence Systems*. Springer, 2010, pp. 168–176.

[12] D. Corapi, A. Russo, M. D. Vos, J. A. Padget, and K. Satoh, "Normative design using inductive learning," *Theory and Practice of Logic Programming*, vol. 11, no. 4-5, pp. 783–799, 2011.

[13] J. Morales, M. López-Sánchez, J. A. Rodríguez-Aguilar, W. W. Vasconcelos, and M. J. Wooldridge, "Online automated synthesis of compact normative systems," *ACM Transactions on Autonomous and Adaptive Systems*, vol. 10, no. 1, pp. 2:1–2:33, 2015.

[14] B. T. R. Savarimuthu, R. Arulanandam, and M. Purvis, "Aspects of active norm learning and the effect of lying on norm emergence in agent societies," in *International Conference on Principles and Practice of Multi-Agent Systems*. Springer, 2011, pp. 36–50.

[15] N. Oren and F. Meneguzzi, "Norm identification through plan recognition," Presented at the 15th International Workshop on Coordination, Organizations, Institutions, and Norms in Agent Systems, arXiv:2010.02627, 2013.

[16] S. Cranefield and A. Dhiman, "Identifying norms from observation using MCMC sampling," in *Proc. of the 30th International Joint Conference on Artificial Intelligence*. International Joint Conferences on Artificial Intelligence, 2021, (in press).

[17] V. Sarathy, M. Scheutz, and B. F. Malle, "Learning behavioral norms in uncertain and changing contexts," in *Proceedings of the 8th IEEE International Conference on Cognitive Infocommunications*, 2017, pp. 301–306.

[18] R. Falcone, C. Castelfranchi, H. L. Cardoso, A. Jones, and E. Oliveira, *Norms and Trust*. Springer, 2013, pp. 221–231.

[19] J. F. Hübner, L. Vercouter, and O. Boissier, "Instrumenting multi-agent organisations with artifacts to support reputation processes," in *Coordination, Organizations, Institutions and Norms in Agent Systems IV*, ser. Lecture Notes in Computer Science, vol. 5428. Springer, 2008, pp. 96–110.

[20] R. Murali, S. Patnaik, and S. Cranefield, "Mining international political norms from the GDELT database," in *Coordination, Organizations, Institutions, Norms, and Ethics for Governance of Multi-Agent Systems XIII*, ser. Lecture Notes in Computer Science, vol. 12298. Springer, 2020, pp. 35–56.

[21] M. Finnemore and K. Sikkink, "International norm dynamics and political change," *International Organization*, vol. 52, no. 4, p. 887–917, 1998.

[22] E. A. Nadelmann, "Global prohibition regimes: the evolution of norms in international society," *International Organization*, vol. 44, no. 4, p. 479–526, 1990.

[23] A. M. Nah, "Networks and norm entrepreneurship amongst local civil society actors: advancing refugee protection in the Asia Pacific region," *The International Journal of Human Rights*, vol. 20, no. 2, pp. 223–240, 2016.

[24] M. J. Hoffmann, "Entrepreneurs and norm dynamics: An agent-based model of the norm life cycle," 2017. [Online]. Available: https://sckool.org/entrepreneurs-and-norm-dynamics-an-agent-based-model-of-the-no.html

[25] A. Morris-Martin, M. De Vos, and J. Padget, "Norm emergence in multiagent systems: a viewpoint paper," *Autonomous Agents and Multi-Agent Systems*, vol. 33, no. 6, pp. 706–749, 2019.

[26] B. T. R. Savarimuthu, J. Padget, and M. A. Purvis, "Social norm recommendation for virtual agent societies," in *International Conference on Principles and Practice of Multi-Agent Systems*. Springer, 2013, pp. 308–323.

[27] B. T. R. Savarimuthu and S. Cranefield, "Norm creation, spreading and emergence: A survey of simulation models of norms in multi-agent systems," *Multiagent and Grid Systems*, vol. 7, no. 1, p. 21–54, 2011.

[28] A. S. Rao and M. P. Georgeff, "BDI agents: From theory to practice," in *Proceedings of the First International Conference on Multiagent Systems*. The MIT Press, 1995, pp. 312–319.

[29] N. Criado, E. Argente, and V. Botti, "Normative deliberation in graded BDI agents," in *German Conference on Multiagent System Technologies*. Springer, 2010, pp. 52–63.

[30] J. Lee, J. Padget, B. Logan, D. Dybalova, and N. Alechina, "N-Jason: Run-time norm compliance in AgentSpeak(L)," in *International Workshop on Engineering Multi-Agent Systems*. Springer, 2014, pp. 367–387.

[31] S. Ranathunga, S. Cranefield, and M. K. Purvis, "Integrating expectation monitoring into BDI agents," in *Programming Multi-Agent Systems - 9th International Workshop, ProMAS 2011*, ser. Lecture Notes in Computer Science, vol. 7217. Springer, 2011, pp. 74–91.

[32] M. Bratman, *Intention, Plans, and Practical Reason*. Harvard University Press, 1987.

[33] W. W. Vasconcelos, M. J. Kollingbaum, and T. J. Norman, "Normative conflict resolution in multi-agent systems," *Autonomous Agents and Multi-Agent Systems*, vol. 19, no. 2, pp. 124–152, 2009.

[34] S. Kambhampati, "Synthesizing explainable behavior for human-AI collaboration," in *Proc. of the 18th International Conference on Autonomous Agents and MultiAgent Systems*. IFAAMAS, 2019, p. 1–2.

[35] D. Kahneman, *Thinking, Fast and Slow*. New York: Farrar, Straus and Giroux, 2011.

[36] G. Booch, F. Fabiano, L. Horesh, K. Kate, J. Lenchner, N. Linck, A. Loreggia, K. Murgesan, N. Mattei, F. Rossi, and B. Srivastava, "Thinking fast and slow in AI," *Proc. of the AAAI Conference on Artificial Intelligence*, vol. 35, no. 17, pp. 15 042–15 046, 2021.

[37] M. Hausknecht and P. Stone, "Deep recurrent Q-learning for partially observable MDPs," in *AAAI Fall Symposium on Sequential Decision Making for Intelligent Agents (AAAI-SDMIA15)*, 2015.