A Roadmap for Self-Evolving Communities

Nardine Osman and Carles Sierra

Artificial Intelligence Research Institute (IIIA-CSIC), Bellaterra, Barcelona, Spain

Abstract. Self-organisation and self-evolution is evident in physics, chemistry, biology, and human societies. Despite the existing literature on the topic, we believe self-organisation and self-evolution is still missing from the IT tools (whether online or offline) we are building and using. In the last decade, human interactions have been moving more and more towards social media. The time we spend interacting with others in virtual communities and networks is tremendous. Yet, the tools supporting those interactions remain rigid. This position paper argues the need for self-evolving software-enabled communities, and proposes a roadmap for achieving this required self-evolution. The proposal is based on building normative-based communities, where community interactions are regulated by norms and community members are free to discuss and modify their community's norms. The evolution of communities is then dictated by the evolution of its norms.

Keywords: norms, interaction protocols, self-evolution

1 Motivation

The concept of social self-organisation has deep roots in political theory. In the Leviathan, Hobbes expresses the idea that the social structure, the State or *res publica*, is the result of a *social contract* or *pact*. The social contract is an agreement among free individuals in order to surrender power to a central authority in order to get a guarantee of peace. In essence, a social contract is the implicit acceptance of self-limiting individual freedom in order to guarantee self-preservation and a better life. The concept of social contract implies that individuals reflect on the community governance in order to decide on the means to reach commonwealth.

This capacity of individuals to reflect and decide on community governance is missing in the IT tools we are currently building. How individuals interact is decided by groups of engineers that design tools, determine the individuals' privacy levels, hinder some from performing certain tasks, and all this without the individuals having a say. We argue that it is time for us to take control of the evolution of our own virtual communities, which we refer to as software-enabled communities — examples of such communities are those based on social networks or any other community whose management relies on software, such as customer relationship management systems (CRMs) or conference management systems. To address this issue, we propose normative-based communities, where community interactions are regulated by norms and community members are free to discuss and modify their community's norms (or social contracts), while respecting their free will. The roadmap for self-evolution is then dictated by the lifecycle of norms, and the research required for supporting this lifecycle is varied and multidisciplinary — from social sciences and legal studies, to learning mechanisms, agreement technologies, natural language processing, formal logics, norm enforcement and regimentation, human computer interaction, and other software engineering techniques.

It is true that the notion of self-organisation is not new. Self-organisation, in general, refers to how a system reorganises to adapt to changes to its goals or environment, where the reorganisation is driven from within the system as opposed to some external control. In many cases, self-organisation is associated with emergence, where a structure appears at a higher level without it being represented at a lower level. Self-organisation has been studied in a number of areas, such as biology, chemistry, geology, sociology, as well as information technology [26]. In AI, self-organisation has mostly been inspired by naturally existing self-organisation models, although some new mechanisms have been designed specifically for software applications [27]. For example, self-organisation may be found in bio-inspired robot teams, where insect-based technology is used to aid robots self-organise to accomplish a task [44]. Another example is intrusion detection in network security, which is inspired from the natural immune system [20]. Swarm intelligence may be used to mimic insect foraging behaviour on the coordination and control of data network traffic [48]. With the notions of self-adaptability and long-lasting evolvability in mind, it has also been proposed that natural ecosystems can inspire building ecosystems of services [47].

In this paper, however, we do not propose to follow traditional self-organisation AI techniques, but to open a new direction in AI research with a novel method for self-evolving software-enables communities. The proposed research is original in the sense that we do not simply talk about self-organising software that usually imitates existing natural systems, but we study how the human users' evolving needs consciously direct evolution. We essentially propose to provide the human users with learning mechanisms that help them learn the best evolution path, along with agreement technologies that help them discuss, argue, and agree on their preferred path of evolution. The system then takes care of the formalisation, operationalisation, and enforcement of the agreed upon norms. This is a fundamentally different approach from existing self-organising systems.

2 Roadmap for Self-Evolving Communities

In this paper, we argue that it is time for us to take control of the evolution of our own virtual communities. To address this issue, we propose normative-based communities, where community interactions are regulated by norms and community members are free to discuss and modify their community's norms.

Norms are the rules that govern behaviour in groups and societies [8]. They motivate and influence individual actions by dictating what values, beliefs, attitudes, and behaviours are deemed appropriate or not. Social norms have been extensively studied by anthropologists, sociologists, philosophers, and economists in the hope of understanding how they motivate individual actions, influence market behaviour, and so on. In multiagent systems, the study of norms gained tremendous attention due to the critical issue of coordinating agent behaviour and actions. Although, unlike other social sciences, the distinction between social and legal norms has not been concrete in the field of computer science.

We adopt the idea of using norms to control, or mediate, community behaviour. This is motivated by the fact that software is usually engineered based on some notion of norms in mind; furthermore, the use of norms allows human users to discuss their interactions without the need for any technical knowledge about the software actually mediating their interactions. We note that by specifying a community through its norms, the evolution of the community becomes dictated by the evolution of its norms. As such, the evolution of software-enabled communities may be depicted by the lifecycle of norms, as illustrated by Figure 1.

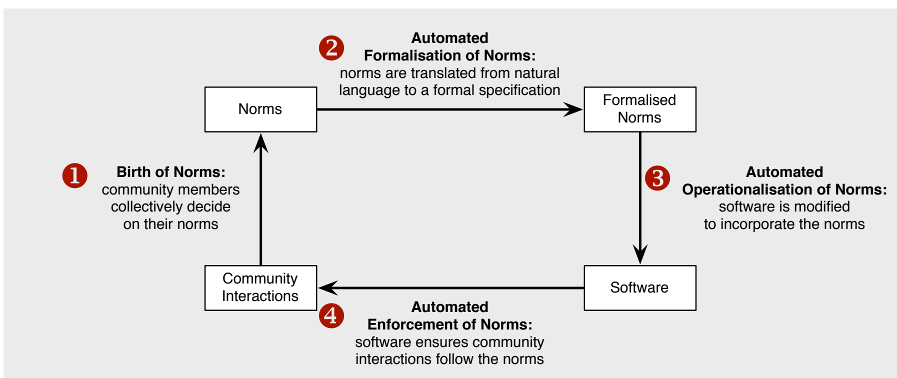

Fig. 1. The lifecycle of norms lays out the roadmap for self-evolving communities

We divide the lifecycle of norms in a software-enabled community into four different stages. In the first stage, community members discuss and agree on their community's norms. If community members do agree on a new set of norms, then the norms' lifecycle enters its second stage, where the agreed upon norms are automatically translated into some formal specification. Given the formal specification of norms, the lifecycle then enters the third stage, where the software-enabled community interactions are modified to operationalise the new norms. Finally, as the software runs to mediate community interactions, norms are enforced by ensuring that community interactions follow the agreed upon norms. The details of these four stages are presented next.

2.1 Birth of Norms

In this stage, community members propose norm modifications (including the abandoning and adoption of norms), discuss the different proposals amongst themselves, and collectively agree on their community's new (or modified) set of norms.

An interdisciplinary research is highly recommended at this stage. For example, **social sciences** can help understand the emergence of norms, their social function, their impact on behviour, as well as understand the possible deviance from norms. **Legal studies** can help understand the legal aspects of norms, such as the enforcement of norms, the effectivity and validity of norms, and the resolve of conflicting norms. **Learning mechanisms** can help learn from a community's past experience, or a similar community's past experience, and suggest norm modifications accordingly. Finally, to aid the collective agreement on community norms, **agreement technologies** will be needed.

Concerning learning mechanisms, learning from data has been one of the main subfields of artificial intelligence [2], yet the literature lacks the in-depth study of learning the best *norms* for multiagent interactions. In multiagent systems, some work has been carried out on norm synthesis [40]. Most norm synthesis (online) approaches, such as [38, 46, 37, 23], are based on norm emergence techniques that require agents to collaborate to synthesise their own norms. [29] does not require the active participation of agents in norm synthesis as norms are synthesised from the observation of agents' interactions. The norm synthesis of [30] is based on the capability of generalising norms, which allows characterising necessary conditions for coordination, avoiding over-regulation.

Concerning agreement technologies, we note that this field is one of the vibrant and active fields in multiagent systems. Agreement technologies aim at helping individuals collaboratively reach a decision, which is crucial when individuals are intelligent and autonomous, having their own goals and agendas to fulfil. The field is based on a number of models and mechanisms, such as argumentation and negotiation mechanisms, social choice theory and voting mechanisms, and trust and reputation models.

Argumentation theory uses logical reasoning to illustrate how conclusions may be drawn in cases of uncertainty and conflict. Argumentation frameworks in multiagent systems [32] have mostly been built on Dung's theory of abstract argumentation [17], where a directed graph is used to represent arguments (the nodes of the graph) and the attack relations between them (the directed links of the graph). A 'calculus of opposition' is then applied to help determine the winning argument. Dung's framework has been extended in a number of directions, such as adding preferences [3], allowing attacks on attacks [7], and adding weights to attacks [18]. One interesting extension to Dungs framework is one that incorporates social voting and is used in social networks [25, 11].

Social choice theory is concerned with the representation and aggregation of individual preferences. Voting mechanisms and ranking systems are considered to fall under the social choice category. Social choice theory studies a variety of aspects, such as the manipulation of a voting rule [10], the use of combinatorial domains where voters may specify a list of preferences [9], and the efficiency and fairness of voting rules [19].

A common issue that arises in building open distributed systems is the issue of trust and reputation. Trust has become an increasingly important concept in computer science [5, 22]. For example, there have been studies on the development of trust in e-commerce [34], on mechanisms to determine which sources to trust when faced with multiple conflicting information sources [54], and mechanisms for identifying which individuals to trust based on their past activity [1]. Trust is an especially important issue from the perspective of autonomous agents and multiagent systems [42]. In multiagent interactions, agents will have to reason about how much should they trust other entities (in various contexts). For reviews on trust in multiagent systems, we refer the reader to [33, 36].

Roadmap: *In the context of self-evolving communities, research in the information and computer Science field should focus on* **two main issues**. *(1) We should aim at building a system that is capable of automatically suggesting norms that can help improve the quality of interactions in a community, based on learning from similar past experiences. Improving the quality of interactions may refer to a number of issues, such as improving the individual user satisfaction, or improving the efficiency of the interactions by helping community members achieve their goals faster and easier. We note that learning mechanisms may also be combined with other techniques, such as analogical reasoning [45] or coherence theory [43]. (2) We should investigate the combination of different agreement technologies to help community members agree on how should their community evolve. For example, it would be interesting to study the combination of trust and argumentation mechanisms, for instance, by modifying the reasoning process to consider reputation/trust based measures of the strengths of arguments. It would also be interesting to study voting algorithms mediated by trust, with votes potentially weighted by the trustworthiness or degree of involvement in the community. The* **outcome** *of this research work would be to aid community members in agreeing on their community norms.*

2.2 Automated Formalisation of Norms

Norms agreed upon by community members need to be formalised in a language that can be understood by the system mediating community interactions. We expect community members to be non-technical people who discuss and specify norms in natural language. An automatic translator is then needed to translate the norms from their natural language specification to some formal specification that is comprehensible by the system.

The research that should support this stage should focus on **natural language processing**, which will help perform the automatic translation, and **formal logics**, which will define the formal language used to specify norms.

Concerning natural language processing, machine translation [49] is a sub-field that translates one natural language into another. This sub-field has been well studied and it has been even used in commercial applications, such as Google Translate or Yahoo's Babel Fish. The translation from a natural language to some logic, on the other hand, has attracted much less attention in the natural language processing field. Nevertheless, important research has been carried in that direction. [50, 52] illustrates and discusses the use of existing state-of-the-art in the automatic translation of regulatory rules in natural language into a machine readable formal representation. [39] translates a complete set of pediatric guideline recommendations into a controlled language (Attempto Controlled English – ACE). [51] adopts and applies a controlled natural language to constrain the domain of discourse in a sample discussion from an on-line discussion forums for e-government policy-making. The controlled natural language helps eliminate ambiguity and unclarity, and allows a logical representation of statements. Each of the policy statements is then automatically translated into first-order logic. [53] presents a linguistically-oriented, rule-based approach, for extracting conditional and deontic rules from regulations specified in natural language. Finally, [4, 6] present approaches for the logical representation of regulations.

Concerning the formal specification of norms, existing approaches may be divided into two main categories: *declarative* approaches and *procedural* ones. We note that in this paper, we refer to the declarative approaches as norms and the procedural ones as interaction models. Declarative formalisms are concerned with the expressiveness of norms, the formal semantics, and how to resolve conflicts arising from an inconsistent set of norms. Declarative approaches are usually based on deontic logic, which is the logic of duties. They deal with concepts like permissions, prohibitions, and obligations, which help specify who can do what, under what conditions, and so on.

Deontic-based policy languages have been used widely in hardware systems and networks for security reasons, trust negotiation, access control, authorisation, and so on. [41] defines policies to be "one aspect of information which influences the behaviour of objects within the system". [13] categorises policies into two types. The first type covers obligation policies for managing actions, which are usually event triggered condition-action rules. The basic concept is that specific events trigger specific actions, and the actions may only be executed if a predefined set of conditions is satisfied. The second type covers authorisation policies, which are usually used for access control.

In multiagent systems, several deontic based formal logics have been proposed for defining a normative specification of agents interaction, such as [31] and [14]. In [31], it is proposed that a community should be defined whenever a group of agents have some common goals and they need to act as a whole within the society in order to fulfil those goals. They propose to define the roles within the community, the relationships among them, which actions each role can do, and how the obligations are distributed among the roles. Each role has associated deontic notions that describe the role obligations, permissions and prohibitions. [14] extends the BDI model of agents to include goals, obligations, and norms; the proposal is based on providing a formal definitions of norms by means of some variation of deontic logic that includes conditional and temporal aspects.

Roadmap: *In the context of self-evolving communities, research should focus on* **two main issues**. *(1) We should aim at building an automatic translator that* *can translate norms specified in a natural language into some formal logic. (2) We should select, or design, a logic that fulfills the following three requirements: i) it should be expressive enough to capture the community members' needs; ii) it should be compatible with the natural language processing technique that needs to translate norms specified in natural language into the logic; and iii) it should be compatible with the software system that needs to operationalise those norms. The* **outcome** *of this research work would be to have the agreed upon norms translated into some formal logic that is ready to be operationalised by the system.*

2.3 Automated Operationalisation of Norms

Given a formal specification of norms, the ultimate goal is to enforce those norms. One approach to achieve this is through what is known as **norm regimentation**. In other words, the system mediating community interactions needs to be modified in order to operationalise those norms.

The literature provides a variety of solutions that deal with specifying and regulating interactions in multiagent systems based on the concept of following social norms [40], such as having contracts and commitments [15], organisational approaches [24], electronic institutions [16], distributed dialogues [35], and so on.

Two specifically interesting approaches are electronic institutions [16] and the lightweight coordination calculus [35]. In [16], it is argued that open multiagent systems can be effectively designed and implemented as agent mediated electronic institutions where heterogeneous (human and software) agents can participate, playing different roles and interacting by means of speech acts. An institution is defined by a set of roles that agents participating in the institution will play, a common language to allow heterogenous agents to exchange knowledge, the valid interactions that agents may have structured in conversations, and the consequences of agents' actions within an institution, captured by obligations that agents acquire and fulfil.

The lightweight coordination calculus (LCC) [35] is a process calculus, based on logic programming, that provides means of achieving coordination in distributed systems by enforcing social norms. The process calculus specifies what actions agents can perform, when they can perform such actions, under what conditions these actions may be carried out, and so on. However, unlike electronic institutions, there are no 'governors' that ensure that agents abide by norms. Of course, like all the approaches above, these rules are associated with roles rather than physical agents; and agents can play more than one role in more than one interaction. This provides an abstraction for the interaction model from the individual agents that might engage in such an interaction.

In the context of self-evolution, it is important to build a system whose interaction models can operationalise norms requested by its uers. We note, however, that not all norms are capable of being operationalised. For example, in an e-commerece context, the system can operationalise the norm that states that a buyer cannot rate the seller more than once, as the system can actually prevent the buyer from doing so. However, the system cannot operationalise a norm that states that only people with sufficient credit can bid, as credit is private information that cannot be accessed and assessed by the system. To deal with norms that cannot be operationalised, alternative **norm enforcement** methods will need to be applied, such as applying sanctions (punishments and rewards).[1 2] In this case, the system will need to ensure the enforcement of un-operationalised norms by checking the members' abidance to those norms and enforcing the appropriate sanctions accordingly. To achieve this, it is important to design a lightweight computational norm enforcement model that allows for the detection of norm violations and the application of remedial actions. This work may build upon existing work, as the literature is rich with norm engines that may be used for resolving conflicting norms and applying appropriate sanctions [21, 28, 12].

Roadmap: *In summary, in the context of self-evolution, research should focus on* **two main issues***: (1) building a mechanism that allows the system to operationalise the norms agreed upon by the community members; and (2) building a norm enforcement model that can check members' abidance to un-operationalised norms and enforce appropriate sanctions accordingly. The* **outcome** *of this research work would be to have a system that is capable of enforcing the norms agreed upon by the community members.*

2.4 Automated Enforcement of Norms

Finally, running the software that mediates community interactions will allow members of the community to interact, while ensuring the enforcement of the norms. The research that will support this stage will need to focus on **human computer interaction** techniques that are required to ensure an intuitive, user friendly interaction for non-technical community members. Research will also need to focus on **software engineering techniques** to ensure efficient community interactions. However, the context of self-evolution imposes an additional and critical challenge: how can the system be modified at runtime when different members may still be interacting together? Indeed, it is very unrealistic to assume that all members will need to complete their current interactions successfully and pause any future interactions for the system to get modified (as future interactions will need to follow the new norms). As such, research will need to address the issue of seamlessly adapting the system at runtime, and it may take inspiration from existing software compatibility models that allow different versions of an application to continue to communicate and collaborate.

Roadmap: *In summary, in the context of self-evolution, research should focus on* **two main issues***: (1) building a user friendly interface that would allow*

[1] The literature usually refers to the operationalisation of norms as 'norm regimentation', and alternative norm enforcement methods (such as the application of sanctions) as 'norm enforcement'. In this paper, however, we overload the word 'enforce' as we use it to describe both approaches.

[2] It may be argued that punishment and reward is not always the right approach for motivating the abidance to norms. Furthermore, what may be considered a punishment for one may be viewed as a reward for another. In this paper, although we talk about punishments and rewards, we concur that sanctions may be labeled more generally as the post-conditions of abiding with or breaking a norm.

community member to interact, following the norms, without any knowledge of their interactions and their complexity, and in cases when users need to discuss norms, norms and norm violations will need to be visualised in a user friendly and comprehensible manner; and (2) engineering the system in such a way that ensures efficient interactions and permits evolution at runtime in a seamless manner that does not interrupt community interactions. The **outcome** *of this research work would be to have the system run and manage evolving normative-based community interactions in an efficient, userfriendly, and seamless manner.*

3 Conclusion

We argue that just like human communities, e-communities (or virtual communities) also need to self-evolve. Instead of creating numerous rigid systems, which is very common in online social systems, what we should aim at instead is providing tools for creating self-evolving systems that adapt to the community's needs. We believe different communities should be governed by different rules. These rules should be an ever evolving set resulting from the aspirations of its members, and not a rigid set defined by a corporation (or a system designer) that does not take the community members' interests at heart. Furthermore, for the community's rules to be effective, they need to be tailored to the specific community, such as considering the character traits of the community's members.

We believe that the notion of self-evolving communities will revolutionise the way how software is build, as well as how we interact with software. We say that in the context of social interactions, software does not need to be rigid, nor customised for the individual, but adaptable to the group. Software engineers will need to design intelligent software that is capable of evolving according to the needs of the community as a whole, and users (or community members) should become more active in shaping their community by discussing and agreeing on their community rules.

In summary, the proposal is based on the traditional notion of mediating social interactions via norms. The roadmap for self-evolution is then dictaded by the lifecycle of norms, which we divide into four different stages (see Figure 1): (1) the birth of norms, (2) the automated formalisation of norms, (3) the automated operationalisation of norms, and (4) the automated enforcement of norms. Varied, and even interdisciplinary, research will be required for supporting these four stages, from social sciences and legal studies, to learning mechanisms, agreement technologies, natural language processing, formal logics, norm enforcement and regimentation, human computer interaction, and other software engineering techniques.

References

1. Adler, B.T., de Alfaro, L.: A content-driven reputation system for the Wikipedia. In: Proceedings of the 16th International World Wide Web Conference. Banff, Alberta, Canada (May 2007)

2. Alpaydin, E.: Introduction to machine learning. MIT press (2004)
3. Amgoud, L., Cayrol, C.: A reasoning model based on the production of acceptable arguments. Annals of Mathematics and Artificial Intelligence 34(1-3), 197–215 (Mar 2002)
4. Antoniou, G., Antoniou, G., Billington, D., Governatori, G., Maher, M.: On the modelling and analysis of regulations. In: In Proceedings of the Australian Conference on Information Systems (1999)
5. Artz, D., Gil, Y.: A survey of trust in computer science and the semantic web. J. of Web Semantics 5(2), 58–71 (2007)
6. Athan, T., Boley, H., Governatori, G., Palmirani, M., Paschke, A., Wyner, A.: OASIS LegalRuleML. In: Francesconi, E., Verheij, B. (eds.) International Conference on Artificial Intelligence and Law, ICAIL '13. pp. 3–12. ACM (2013)
7. Baroni, P., Cerutti, F., Giacomin, M., Guida, G.: AFRA: Argumentation framework with recursive attacks. International Journal of Approximate Reasoning 52(1), 19–37 (Jan 2011)
8. Bicchieri, C., Muldoon, R.: Social norms. In: Zalta, E.N. (ed.) The Stanford Encyclopedia of Philosophy. The Metaphysics Research Lab, Stanford University, Stanford, CA, USA, spring 2011 edn. (2011)
9. Brams, S., Kilgour, D., Zwicker, W.: The paradox of multiple elections. Social Choice & Welfare 15(2), 211–236 (1998)
10. Conitzer, V., Lang, J., Sandholm, T.: How many candidates are needed to make elections hard to manipulate? In: Proceedings of the 9th Conference on Theoretical Aspects of Rationality and Knowledge. pp. 201–214. ACM (2003)
11. Correia, M., Cruz, J., Leite, J.: On the efficient implementation of social abstract argumentation. In: Schaub, T., Friedrich, G., O'Sullivan, B. (eds.) 21st European Conference on Artificial Intelligence, ECAI 2014. Frontiers in Artificial Intelligence and Applications, vol. 263, pp. 225–230. IOS Press (2014)
12. Criado, N., Argente, E., Noriega, P., Botti, V.J.: MaNEA: A distributed architecture for enforcing norms in open MAS. Eng. Appl. of AI 26(1), 76–95 (2013)
13. Damianou, N.C., Bandara, A.K., Sloman, M.S., Lupu, E.C.: A survey of policy specification approaches. Tech. rep., Department of Computing, Imperial College of Science Technology and Medicine, London, UK (2002)
14. Dignum, F., Kinny, D., Sonenberg, L.: From desires, obligations and norms to goals. Cogn. Sci. Quarterly 2, 407–430 (2002)
15. Dignum, V., Meyer, J.J., Weigand, H.: Towards an organizational model for agent societies using contracts. In: Proceedings of AAMAS 2002. pp. 694–695. ACM, New York, NY, USA (2002)
16. d'Inverno, M., Luck, M., Noriega, P., Rodriguez-Aguilar, J.A., Sierra, C.: Communicating open systems. Artificial Intelligence 186, 38–94 (Jul 2012)
17. Dung, P.M.: On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. Artificial Intelligence 77(2), 321–357 (Sep 1995)
18. Dunne, P.E., Hunter, A., McBurney, P., Parsons, S., Wooldridge, M.: Weighted argument systems: Basic definitions, algorithms, and complexity results. Artificial Intelligence 175(2), 457–486 (Feb 2011)
19. Endriss, U., Maudet, N., Sadri, F., Toni, F.: Negotiating socially optimal allocations of resources. Journal of Artificial Intelligence Research 25(1), 315–348 (Mar 2006)
20. Foukia, N.: IDReAM: Intrusion detection and response executed with agent mobility architecture and implementation. In: Proceedings of AAMAS 2005. pp. 264–270. ACM (2005)

21. Gaertner, D., Garcia-Camino, A., Noriega, P., Rodriguez-Aguilar, J.A., Vasconcelos, W.: Distributed norm management in regulated multiagent systems. In: Proceedings of the 6th International Joint Conference on Autonomous Agents and Multiagent Systems. pp. 90:1–90:8. AAMAS '07, ACM, New York, NY, USA (2007)
22. Grandison, T., Sloman, M.: A survey of trust in internet applications. IEEE Commun. Surveys Tuts. 4(4), 2–16 (2000)
23. Griffiths, N., Luck, M.: Norm Emergence in Tag-Based Cooperation. In: Proceedings of COIN (2010)
24. Horling, B., Lesser, V.: A survey of multi-agent organizational paradigms. Knowl. Eng. Rev. 19, 281–316 (2004)
25. Leite, J.a., Martins, J.a.: Social abstract argumentation. In: Proceedings of the Twenty-Second International Joint Conference on Artificial Intelligence. pp. 2287–2292. IJCAI'11, AAAI Press (2011)
26. Marzo Serugendo, G., Gleizes, M.P., Karageorgos, A.: Self-organising systems. In: Di Marzo Serugendo, G., Gleizes, M.P., Karageorgos, A. (eds.) Self-organising Software, pp. 7–32. Natural Computing Series, Springer (2011)
27. Marzo Serugendo, G., Karageorgos, A., Rana, O., Zambonelli, F.: Engineering Self-Organising Systems, Lecture Notes in Computer Science, vol. 2977. Springer Berlin Heidelberg (2004)
28. Modgil, S., Faci, N., Meneguzzi, F., Oren, N., Miles, S., Luck, M.: A framework for monitoring agent-based normative systems. In: Proceedings of The 8th International Conference on Autonomous Agents and Multiagent Systems - Volume 1. pp. 153–160. AAMAS '09, International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC (2009)
29. Morales, J., López-Sánchez, M., Esteva, M.: Using Experience to Generate New Regulations. In: IJCAI. pp. 307–312 (2011)
30. Morales, J., Lopez-Sanchez, M., Rodriguez-Aguilar, J.A., Wooldridge, M., Vasconcelos, W.: Automated synthesis of normative systems. In: AAMAS 2013. pp. 483–490 (2013)
31. Pacheco, O., Carmo, J.: A role based model for the normative specification of organized collective agency and agents interaction. Autonomous Agents and Multi-Agent Systems 6, 145–184 (March 2003)
32. Rahwan, I., Simari, G.R.: Argumentation in Artificial Intelligence. Springer, 1st edn. (2009)
33. Ramchurn, S.D., Huynh, D., Jennings, N.R.: Trust in multi-agent systems. Knowledge Engineering Review 19(1), 1–25 (2004)
34. Resnick, P., Zeckhauser, R.: Trust among strangers in Internet transactions: Empirical analysis of eBay's reputation system. In: Baye, M.R. (ed.) The Economics of the Internet and E-Commerce, pp. 127–157. Elsevier Science (2002)
35. Robertson, D.: A lightweight method for coordination of agent oriented web services. In: Proceedings of AAAI Spring Symposium on Semantic Web Services. Stanford (2004)
36. Sabater, J., Sierra, C.: Review on computational trust and reputation models. AI Review 23(1), 33–60 (September 2005)
37. Sen, O., Sen, S.: Effects of social network topology and options on norm emergence. In: Proceedings of COIN. pp. 211–222 (2010)
38. Sen, S., Airiau, S.: Emergence of norms through social learning. In: IJCAI. pp. 1507–1512 (2007)
39. Shiffman, R.N., Michel, G., Krauthammer, M., Fuchs, N.E., Kaljurand, K., Kuhn, T.: Writing clinical practice guidelines in controlled natural language. In: Proceed-

ings of the 2009 Conference on Controlled Natural Language. pp. 265–280. CNL'09, Springer-Verlag, Berlin, Heidelberg (2010)
40. Shoham, Y., Tennenholtz, M.: On social laws for artificial agent societies: off-line design. Artif. Intell. 73, 231–252 (1995)
41. Sloman, M.: Policy driven management for distributed systems. Journal of Network and Systems Management 2, 333–360 (1994)
42. Teacy, W., Chalkiadakis, G., Rogers, A., Jennings, N.R.: Sequential decision making with untrustworthy service providers. In: Proceedings of the 7th International Conference on Autonomous Agents and Multiagent Systems (2008)
43. Thagard, P.: Coherence in thought and action. MIT Press (2002)
44. Thakoor, S., Morookian, J.M., Chahl, J., Hine, B., Zornetzer, S.: BEES: Exploring mars with bioinspired technologies. Computer 37(9), 38–47 (2004)
45. Veloso, M.M.: Planning and learning by analogical reasoning, vol. 886. Springer (1994)
46. Villatoro, D., Sabater-Mir, J., Sen, S.: Social instruments for robust convention emergence. In: IJCAI. pp. 420–425 (2011)
47. Viroli, M., Zambonelli, F.: A biochemical approach to adaptive service ecosystems. Inf. Sci. 180(10), 1876–1892 (2010)
48. Wedde, H.F., Farooq, M.: A comprehensive review of nature inspired routing algorithms for fixed telecommunication networks. Journal of Systems Architecture 52(8), 461–484 (Aug 2006)
49. Wilks, Y.: An artificial intelligence approach to machine translation. Springer (2009)
50. Wyner, A., Bos, J., Basile, V., Quaresma, P.: An empirical approach to the semantic representation of laws. In: Schäfer, B. (ed.) Legal Knowledge and Information Systems - JURIX 2012: The Twenty-Fifth Annual Conference. Frontiers in Artificial Intelligence and Applications, vol. 250, pp. 177–180. IOS Press (2012)
51. Wyner, A., van Engers, T., Bahreini, K.: From policy-making statements to first-order logic. In: Proceedings of the First International Conference on Electronic Government and the Information Systems Perspective. pp. 47–61. EGOVIS'10, Springer-Verlag, Berlin, Heidelberg (2010)
52. Wyner, A., Governatori, G.: A study on translating regulatory rules from natural language to defeasible logic. In: RuleML 2013. pp. 16.1–16.8. Seatle, USA (July 2013)
53. Wyner, A., Peters, W.: On rule extraction from regulations. In: Atkinson, K. (ed.) Legal Knowledge and Information Systems - JURIX 2011: The Twenty-Fourth Annual Conference. Frontiers in Artificial Intelligence and Applications, vol. 235, pp. 113–122. IOS Press (2011)
54. Yin, X., Han, J., Yu, P.S.: Truth discovery with multiple conflicting information providers on the Web. In: Proceedings of the Conference on Knowledge and Data Discovery (2007)