



Norm emergence in multiagent systems: a viewpoint paper

Andreas Morris-Martin¹ · Marina De Vos¹ · Julian Padget¹

Published online: 30 September 2019
© The Author(s) 2019

Abstract

Norms are utilised in agent societies to encourage acceptable behaviour by the participating agents. They can be established or revised from the top-down (authority) or from the bottom-up (populace). The study of norm creation from the bottom-up—or norm emergence/convergence—shows evidence of increasing activity. In consequence, we seek to analyse and categorize the approaches proposed in the literature for facilitating norm emergence. This paper makes three contributions to the study of norm emergence. Firstly, we present the different perspectives of norms and their impact on the norm emergence process, with the aim of comparing their similarities and differences in implementing the norm life cycle. Secondly, we identify the characteristics that support norm emergence that are observed in the emergence literature. Finally, we identify and propose future topics for study for the community, through a discussion of the challenges and opportunities in norm emergence.

Keywords Norms · Norm emergence · Norm convergence · Norm life cycle

1 Introduction

Norms are defined as either rules of expected behaviour in a society or as behaviour that is common in a society. For the purpose of this paper, we adopt the latter norm definition, behaviour that is common in a society. Savarimuthu et al. [89] cite Tuomela [99]’s norm categorization as follows: r-norms (rule norms), s-norms (social norms), m-norms (moral norms) and p-norms (prudential norms). Rule norms are those norms that are imposed by an authority e.g. a law. Social norms are those norms that apply to large groups such as a whole society and often based on customs and conventions. Moral norms are intended to appeal to an individual’s conscience, and prudential norms are based on rationality [89]. In agent societies, two types of norms are generally distinguished: norms and law [98], depending on

✉ Andreas Morris-Martin
a.l.morris.martin@bath.ac.uk

Marina De Vos
m.d.vos@bath.ac.uk

Julian Padget
j.a.padget@bath.ac.uk

¹ University of Bath, Bath, UK

who established the norm and how they are enforced. In the literature, the term norm often refers to s-norms, while laws are r-norms.

Norms are generally aimed at maintaining behaviour patterns that are acceptable to the majority of the population. Notwithstanding, there are situations where a norm that is observed within a population may not be to everyone's benefit. For example, Székely et al. [98] study the prevalence of extortion rackets as a norm in locations that are Mafia-controlled. Though both human and agent societies use norms as a way to regulate behaviours, the way in which the norms are used in the decision-making process differs greatly. In the decision making process of people, ethics and social norms are entrenched in the action selection process [34]. Agent societies however hardly ever make decisions in this way. Instead agents typically use some utility, resource or time-optimizing method/algorithm predetermined by the programmer in their decision making [34]. It is therefore necessary to design mechanisms for the inclusion of norms in the decision making algorithms of agents.

Societies do not remain the same over time: as they evolve, so inevitably will the norms that govern behaviours within them. As a consequence, societies need mechanisms to facilitate these changes of norms throughout their lifetime. This is especially relevant to agent societies, as designers can allow for this to happen. Norms can be designed into agents from inception, or forced upon agents, or they can emerge from the agents themselves [21]. Frequently in agent societies, the norms governing the behaviour of agents are specified at design time and remain the same for the lifetime of the society. This can work if the society is static and fully understood, but causes problems when the society is dynamic and circumstances can change over its lifetime or when the state space is too extensive to model entirely. For these dynamic societies, when norms are specified at design time only, if situations change and new norms need to be considered then all the agents in the society will need to be redesigned. However, there may be institutions/organisations that should not be halted for this to occur. A solution would be to have the norms originate or emerge from the agents. For a large number of agent societies, norm emergence is observed when a large percentage of the population adopts the norm or chooses the same action.

Having norms emerge from the bottom up is not just easier for the agent but less expensive to implement [90,106]. According to Yu et al. [106] “Due to the expense and inefficiency of having a centralized policing enforcer to formulate and specify social norms in a prescriptive manner, it is more desirable to enable social norms to evolve and emerge on their own without relying on any centralized authority.” Additionally, Axelrod [12] suggests that “what is needed is a theory that accounts not only for the norms existing at any point in time but also how norms change over time”. It is no surprise then that several pieces of work have focused on investigating how norms emerge in agent societies.

There are already several works examining research in norm emergence in multiagent systems. A survey by Savarimuthu and Cranefield [86] analyses the norm life cycle and presents an expanded five stage cycle. It focuses on the categorisation of the simulation models into one of these five stages, based on their area of focus. Our survey is different because we examine the literature specifically on norm emergence and identify characteristics that support or can prevent the emergence of norms as seen across the literature. Additionally we observe that some of the recommendations for simulation models identified in Savarimuthu and Cranefield [86] have been implemented in some of the literature presented in this survey, allowing us the opportunity to discuss the effectiveness of those recommendations as they contribute to the norm emergence process.

Another survey by Haynes et al. [54] discusses the engineering of emergent norms and examines how the approaches presented in the literature are useful in implementing the three steps they identify in the process of engineering emergent norms. In addition, they examine

how the approaches presented are suitable for different types of multiagent systems. The concept of how emergent norms can be engineered is not the focus of this paper but is similar to an idea we present as an opportunity for future research in Sect. 4.

The aim of this viewpoint article is to examine how concepts studied to date in norm emergence can be applied to research in *normative* multiagent systems and more specifically how prescriptive norms—norms that are explicitly represented within the system—can emerge through monitoring the behaviour of the agents within the system. We present the notion of agents in the norm emergence literature that indirectly and unintentionally utilise norms and show how this coincides with explicit norms in normative multiagent systems. Moreover we discuss how these implicit norms/behaviours can be encoded and explicitly represented as norms in normative multiagent systems. We postulate that the mechanisms in place allowing individual agents to suggest norms in norm emergence can be used as a source for norm synthesis in normative multiagent systems. And finally we posit that concepts such as agents proposing new behaviours and agents learning these behaviours from each other can be implemented in autonomous systems in addressing the concept of self governance.

This paper is written with the assumption that the reader is familiar with norms and is interested in norm emergence from either a theoretical or practical perspective. We present a comprehensive analysis of the key works on the process of norm emergence and agent-based simulation models that study norm emergence. The aim is not to be exhaustive but to provide the reader with an overview to support understanding the approaches used in the research presented, and to select the most appropriate technique(s) as we categorise them.

The paper is structured as follows. Section 2 introduces the reader to background concepts on norms and norm emergence, while providing insights to aid in understanding the rest of the viewpoint paper. Section 3 provides a detailed overview of the state of the art, together with a classification of reviewed papers based on a number of selected and argued characteristics, which forms the main body of our paper. Based on this classification of features, the paper concludes in Sect. 4 highlighting and discussing future challenges and opportunities for norm emergence as a whole and its various sub-processes.

2 Norms and norm emergence

This section briefly presents a range of topics whose purpose is to enable the reader to understand the fundamental concepts in the study of norms and norm emergence in multiagent systems and agent-based simulation models.

We begin in Sect. 2.1 discussing how norm compliance or adherence can be modelled in regimented systems where prohibited actions are unavailable, or where agents always comply with norms. The remainder of Sect. 2.1 considers that agents are free to decide whether to adopt norms or not, and examines the reasons why agents with freedom of choice might still adopt norms. We discuss the two main reasons why agents may adopt norms: increased payoff and the avoidance of punishment in Sects. 2.1.2 and 2.1.3, respectively. Then in Sect. 2.2 we address why agents might deliberately choose to violate norms.

We continue in Sect. 2.3 where we discuss alternative perspectives on norms, as the literature seems to be divided along these lines. We also discuss how norms are represented. Following this, in Sect. 2.4 we discuss the concept of norm emergence and briefly outline some potential sources of the norms that emerges in Sect. 2.5. Section 2.6, discusses the norm life cycle and how the implementation differs depending on the different perspectives. Finally Sect. 2.7 summarises the section.

2.1 Norm adherence

Norms are beneficial to multiagent societies because they reduce the computation requirement for an agent to make a decision [90], ensure society's goals are met and function as a conflict resolution strategy [101].¹ For example, Székely et al. [98] found that social norms can help to resolve disputes without expensive legal representation. Though the work by Székely et al. [98] refers to a human society, there is potential for similar benefits in an agent society. When there are several possible actions to take, an agent needs to perform some (possibly computationally expensive) deliberation before selecting an action to execute, when a norm prevails, the agent can just select the normative action.

However some agents will not adopt norms just because they exist, there must also be some benefit to the agent or some reason why an agent ought to adopt a norm, since a software agent's decision making is normally inherently rational. Many social scientists have studied why norms are adhered to. Some of the reasons include: fear of authority; rational appeal of the norms; feelings such as shame, embarrassment and guilt arising from non-adherence [89]. Criado et al. [36] incorporating Castelfranchi [27]'s research, proposes the following strategies for norm compliance: (i) unconditional—always fulfil norms, (ii) instrumental—fulfil when they benefit oneself, (iii) cooperative—fulfil when they benefit the whole society, and (iv) benevolent—fulfil when it benefits an agent whom they want to favour.

2.1.1 Regimented versus non-regimented systems

Norms are adhered to in some systems because agents have no other choice but to do so, as prohibited actions are restricted or unavailable. Alechina et al. [4] for example utilise a guard function during runtime that can restrict the transition to states that would result in a norm violation. In other systems, agents do not have a choice whether to adhere to or violate a norm because they are programmed always to fulfil norms. This is referred to as terminal adoption [5]. Bench-Capon and Modgil [18] argue that this type of regimented system should not be considered as operating through norms. Similarly, the regimenting of social laws and conventions as argued by Villatoro et al. [101] limit the autonomy of an agent and is not a desirable solution in most situations. The modelling of agents to make decisions about whether to adopt a norm or not contributes to maintaining/preserving agent autonomy.

2.1.2 Norm adherence for increased payoff

In some systems, agents receive feedback on their actions in the form of payoffs, where payoffs denote the outcome or result derived from a previous interaction. A positive payoff is seen as a reward and a negative payoff a punishment. Agents often choose to adhere to a norm in order to increase their (positive) payoff in an interaction. A norm in this context is referred to as the preferred action for the majority of agents among the other actions for a situation, which we describe as the emergence perspective, and discuss in more detail in Sect. 2.3.

In the literature, most agents are programmed as inherently selfish: they only cooperate when it is the only or easiest way to achieve a higher payoff, making this one of the agent's

¹ The following concepts where (i) norms promote coherent behaviour, (ii) norms help to reduce disagreement, and (iii) norms reduce computational requirement in decision making, derive from the literature in social sciences and economics, specifically the seminal works of Ostrom [79] and North [77]. Some researchers also credit the work of Axelrod [12] and Shoham and Tennenholtz [96].

main reasons for adopting a norm. Savarimuthu et al. [92] demonstrate how a selfish agent can change their behaviour to become a sharer to increase their own payoff or to accomplish their goal. The observation of this behaviour change is interpreted as the agent adopting a sharing norm. In principle, an agent programmed to maximise payoff will choose the action with the highest payoff, without recognising that it is a norm and without any inclination to adopt or participate in the emergence of a particular norm(s). This is in contrast to norms from the prescriptive perspective, see Sect. 2.3, where agents are deliberate in their action selection with regards to adopting norms.

Sen and Airiau [93] investigate whether norms would emerge in the instance where several combinations of actions give rise to the same payoff. This is important, since most studies reflect the norm that emerges is also the action that yields the highest payoffs. They found that all actions with equal payoff emerged with the same frequency [93]. Not surprisingly, it takes longer for norms to evolve for larger action sets, where several actions can result in the same payoff [93]. Additionally Hu and Leung [57] investigate the use of several actions with equal payoff and found that there was a diverse set of actions that emerged within small groups/communities but no global emergence observed. These results lend support to the argument that agents choose normative actions only to increase their payoff, and if payoffs were constant for any action, it would be difficult for agents to converge to a single one.

2.1.3 Norm adherence to avoid punishment

Norm adherence/compliance is also linked to punishments existing for non-conformance. Norms exist in a given social setting: individuals usually act in a certain way and are often “punished” when seen not to be acting in some socially accepted way [12]. When societies have agents that enforce punishment of violating agents, the punishment for violation is a deterrent against violation in the future. This punishment can be given either by a central authority or distributed enforcer agents. Savarimuthu et al. [90] propose that it is useful to have distributed agents act as enforcers of norms by applying punishment to defaulters, since it would be difficult and expensive to have a central authority monitor and punish defaulting agents. However having distributed enforcers is not an automatic fix, as a society needs to have sufficient number of enforcers to convert non-conforming agents to conforming ones [90]. Similarly having too many distributed enforcers is not without its drawbacks. Balke et al. [14] hypothesise that there comes a point when the cost of enforcement outweighs the benefits.

Boella and van der Torre [20] and Savarimuthu et al. [90] present two distinct approaches to the implementation of distributed enforcers. On the one hand, Savarimuthu et al. [90] allow punisher agents the option of deciding whether to punish or not, as punishing can have effects on the punisher agent itself. So it is normal for a punisher agent to observe a violation and ignore the defaulting agent, when it is not willing to bear the cost of punishment. On the other hand, Boella and van der Torre [20] utilises defender agents on whom obligations are imposed, that require them to monitor violations of norms within the society and apply sanctions. The framework in Boella and van der Torre [20] is based on the premise that if agents can observe the effect of their action or non-action on the system, it will impact their decision either to comply with or resist a norm.

Vanhee et al. [100] pose the question, but do not conclude either way, of whether norms should be enforced via violations and sanctions, or whether they should be set up so that no violations occur. This question raises another pertinent one, if there were no punishments for violation, would agents still adopt norms? The answer to the latter, though not apparent, would be significant, as it would determine if agents only adopt norms to avoid punishment,

and would violate them if punishments did not exist, which is similar to the position of [12]. However, if and when punishment for violations exist, agents must be made aware of the consequences of their violations before they make decisions, as they may choose to avoid the punishment.

Another important concept with norms and punishment is that there may be a misconception that suggests that, when one observes a certain type of action, then it must mean that there is no norm, legal or social, regulating that particular action in a normative system. However this might not always be true. It is very possible that a norm does exist but no punishment is being carried out [98]. When no punishment is carried out and the norms are not beneficial to the agents or in line with their goals, agents may not adopt norms.

2.2 Deliberate norm violation

Violations or sub-ideal behaviour of agents is inevitable [45]. Though it is usually beneficial for agents to comply with norms as described above, there are some situations where it might be better for the agent to violate norms. Agents may choose to violate norms when violations provide a better way to achieve goals [95]; when personal goals conflict with organisational goals e.g. in common pool resources [46,48]; or when they need to violate one set of norms to comply with another set of norms [1,45].

In Shams et al. [95], the agent chooses a plan that maximises overall utility, as opposed to the norm compliant one. Utility is calculated by deducting penalty costs for violating norms from the value gained from goals satisfied. Bench-Capon and Modgil [18] use value-based reasoning to tell an agent when and how to violate a norm in a complex system. They define an acceptable ordering of values for the agent to base their reasoning on since for some complex systems the reasons to adopt a norm in one situation may lead to violating it in another. However, they warn that trying to impose a specific value ordering on agents can lead to a conformist society and limit their autonomy.

Ajmeri et al. [1] and Gasparini et al. [45] demonstrate strategies for making decisions in situations where violating one set of norms in order to adhere to another set of norms is also the rational decision. Ajmeri et al. [1] refer to conflicting commitments and the use of some importance ordering to make a rational decision when these occur. In this work, the resolution of conflicts uses dominance relations, where the satisfaction of a more dominant instance overrides the violation of a less dominant one, and the agent is still seen as compliant even though one or more instances are violated. Santos et al. [85] survey several conflict resolution mechanisms proposed in the literature, some of which involve the violation of some norms to comply with others.

Gasparini et al. [45] consider that norms and the violations of norms vary in severity. The concept of determining the severity of norm violations is important since there may be some norms that it is reasonable to violate and others whose violation can have catastrophic effects. Gasparini et al. [45] propose a mechanism for determining the severity of norms and use of contrary-to-duty obligations² to repair failure states. In their research, they show that agents may choose to violate a less severe norm now, even when no conflict is involved, in order to avoid violating a more severe norm in the future. The mechanism defines a preference relation to determine from a particular state, which of the next possible states is more or less preferred based on the severity of norm violations.

² Contrary-to-duty obligations [39,82] are common concepts in deontic logic to recover from norm violations. Contrary-to-duty obligations are obligations that are instituted only when some other obligation has been violated.

2.3 Perspectives and representations of norms

Norms can be looked at from two distinct perspectives, norms as deontic concepts, and norms as a preference behaviour, respectively. Norms as deontic concepts views norms as permissions/prohibitions and obligations, or commitments. Conte and Castelfranchi [32] refer to this perspective as “norms as prescriptions”, while Savarimuthu et al. [92] refer to it as the prescriptive approach. This deontic logic perspective on norms is widely studied in the literature. Research here includes but is not limited to: (i) frameworks to describe and model norms as deontic logic in institutions³ for example: InstAL[31], JaCaMo [22], Operetta [2]; (ii) the development of agents that reason about norms in their decision making for example: BDI (Belief–Desire–Intention) and BOID (Belief–Obligation–Intention–Desire) [19,24,35,38], Normative KGP (Knowledge–Goals–Plans) agents [84], NBDI(Norm–Belief–Desire–Intention) [41] (iii) approaches to synthesising normative systems for example: IRON and SENSE [70,71,73], AOCMAS [26], Guard functions in [4]. (iv) and more recently values and norms, for example [18,34,94].

The following provide a more extensive collection of research on prescriptive norms: the NorMAS 2012 Handbook [8], the Handbook of Deontic Logic and Normative Systems [44], Handbook of Research on Multi-Agent Systems: Semantics and Dynamics of Organizational Models [40] and more recently Social Coordination Frameworks for Social Technical Systems [3] for example.

The other perspective looks at norms as a predetermined or computed preference behaviour to execute from among a set of behaviours in a given situation. Here norms that emerge are seen as a particular behaviour, strategy, policy or action from a set of possible actions in a similar situation. Conte and Castelfranchi [32] refer to this perspective as norms as conventions, while Savarimuthu et al. [92] refer to it as the emergence approach. The majority of societies, in the literature on norm emergence, adopt the notion of norm from this latter perspective as is clear from Table 1.

Conte and Castelfranchi [32] posit that the norm literature presents them as a dichotomy, where at one end there is the rational view, in which conventions are considered and emergence is studied, while at the opposing end lies the prescriptive view, characterised by deliberate issuing of norms or laws. The authors believe that there is a need for a unified view of norms and attempts to fill the gap between the views, by putting forward a bridge between conventions and prescriptions that considers distributed goals and a complementary perspective of prescriptions, as opposed to explicit issuance. They describe the notion of a social norm as the belief that a given behaviour is prescribed within the population, and that behaviour is followed once it is believed to be obliged. This conformance results in prescribing of that belief, which in turn results in its spreading.

Another useful distinction in the literature is how norms are represented in agent societies: explicit versus implicit norm representation. Agent societies with an explicit representation commonly have an internal representation of norms in the beliefs of the agent, such as with BDI agents, examples of which include Beheshti [15], Morales et al. [73], Conte and Castelfranchi [32]. There are other cases where there is an external representation of the norm in a distinct location, which represents the normative system implemented as a common knowledge source that can be referenced [31,58,61]. Additionally, there are agent societies where there is both an internal and an external representation of the norms for example, Boella and van der Torre [19] and dos Santos Neto et al. [41].

³ An institution is a collection of norms outside the agent(s) that monitors their behaviour in the context of these norms

Agent societies with implicit norm representations do not use the term “norm” explicitly. Rather, agents have a predetermined strategy for deciding on a particular behaviour in a given situation. They are able to learn or modify their strategies either by mimicking another agent’s strategy, utilising a machine learning mechanism, advice learning or some data mining mechanism [91]. When a majority of the agents in a society utilise the same strategy for every instance of the situation over time, a norm is deemed to have emerged. Notably for most of the norm emergence mechanisms in the literature, an implicit representation of norms is used, see Table 1. When an implicit norm representation is used, there is no need for agents to have an internal representation of norms as in the case of some BDI agents and consequently no stored representation of a norm. To our knowledge, the works of Savarimuthu et al. are the only research utilising an implicit representation of norms where there is an explicit reference to a norm. They refer to the stored policies as “proposal norm” and “acceptance norm” [87–89].

When there is no explicit norm representation, what essentially happens is that agents learn best response or rational action based on interactions. Then, those observing the society may make the inference that a norm has emerged. This questions the validity of saying that “a norm has emerged” since in essence the agents do not have the concept of a norm, only that this action seems to be the most popular or beneficial one in this situation. Can we then safely say that a norm has emerged? Beheshti [15] argues that an agent can observe behaviours during interactions and even learn these new norms, but until the agent modifies its internal BDI (belief, desires and intentions) memory, then it is unlikely that those learned behaviours will affect their actions. Beheshti [15] explores an approach to internalise norms by proposing the Cognitive Social Learner (CSL). CSLs implement norms through an iterative process, where normative behaviour is developed incrementally within the cognitive model of the agent [15]. Subsequently the norm emerges as observed behaviour in the agent.

Explicit norm representation is usually the characteristic of societies with norms defined as laws, legal norms, and where the adherence to them are likely monitored and enforced: the prescriptive perspective. However social norms can also have an explicit norm representation, and those norms may or may not be monitored or punished as in Andrighetto et al. [6]. Additionally there may be a combination of both social and legal norms as in Frantz et al. [43]. Implicit norm representation is characteristic of norms referred to as social norms which are generally not punishable, however there may be cases where punishment is used as an approach for propagating norms and promoting norm emergence, for example Savarimuthu et al. [90], Balaraman and Singh [13], Lotzmann et al. [66]. Implicit norm representation is observed in societies with an emergence perspective of norms, see Table 1.

2.4 Norm emergence

When agents interact with other agents, a pattern of interaction is observed and agents are able to learn the appropriate action for future interactions. It can be said that this pattern “emerges” from those interactions. The exposition of Mintz-Woo [69] on the work of Axelrod [12] describes the study of norm emergence as an evolutionary approach, where norms arise through iterated games. These iterated games facilitate the repeated interaction among agents for the norms to be learnt. But during an interaction, there are several approaches that agents can use to learn norms: imitation, machine learning techniques, advice learning or data mining mechanisms [91]. The mechanism used to learn the norm and its impact on the emergence of norms in a society is discussed in Sect. 3.

Norm emergence or convergence is widely accepted to have happened, or a norm is said to have emerged, when a predetermined percentage of the population observes the

Table 1 Summary of different perspectives and representation of norms used in the literature surveyed

	Norm perspective		Norm representation		Norm emergence studied
	Emergence	Prescriptive	Implicit	Explicit	
Alechina et al. [4]		✓		✓	
Aldewereld and Dignum [2]		✓		✓	
Balaraman and Singh [13]		✓		✓	✓
Beheshti and Sukthankar [16]	✓		✓		✓
Beheshti et al. [17]		✓		✓	✓
Bench-Capon and Modgil [18]		✓		✓	
Boissier et al. [22]		✓		✓	
Boella and van der Torre [20]		✓		✓	
Boella and van der Torre [19]		✓		✓	
Boyd and Richerson [23]		✓		✓	
Broersen et al. [24]		✓		✓	
Brooks et al. [25]	✓		✓		✓
Campos et al. [26]		✓		✓	
Chakrabarti and Basu [28]	✓		✓		✓
Chao Yu et al. [29]	✓		✓		✓
Cliffe et al. [31]		✓		✓	
Conte and Castelfranchi [32]		✓		✓	
Criado et al. [35]		✓		✓	
Dascalu et al. [37]	✓		✓		✓
Dastani and van der Torre [38]		✓		✓	
dos Santos Neto et al. [41]		✓		✓	
Franks et al. [42]	✓		✓		✓
Frantz et al. [43]		✓		✓	✓
Ghorbani et al. [48]	✓		✓		✓
Ghorbani and Bravo [46]	✓		✓		✓
Hoffmann [55]	✓		✓		✓
Hofmann et al. [56]	✓		✓		✓
Hassani-Mahmooei and Parris [53]	✓		✓		✓
Hao and Leung [49]	✓		✓		✓
Hao et al. [50]	✓		✓		✓
Hao et al. [51]	✓		✓		✓
Hao et al. [52]	✓		✓		✓
Hu and Leung [57]	✓		✓		✓
Hübner et al. [58]		✓		✓	
Lopez [65]	✓		✓		✓
Lotzmann et al. [66]		✓		✓	✓
Lee et al. [61]		✓		✓	
Mahmoud et al. [67]	✓		✓		✓
Mashayekhi et al. [68]		✓		✓	

Table 1 continued

	Norm perspective		Norm representation		Norm emergence studied
	Emergence	Prescriptive	Implicit	Explicit	
Mukherjee et al. [74]	✓		✓		✓
Mukherjee et al. [75]	✓		✓		✓
Mungovan et al. [76]	✓		✓		✓
Morales et al. [70]		✓		✓	
Morales et al. [72]		✓		✓	
Morales et al. [71]		✓		✓	
Morales et al. [73]		✓		✓	
Ohtsuki et al. [78]	✓		✓		✓
Peleteiro et al. [80]	✓		✓		✓
Peleteiro et al. [80]	✓		✓		✓
Riveret et al. [83]		✓		✓	✓
Sadri et al. [84]		✓		✓	
Savarimuthu et al. [87]	✓		✓		✓
Savarimuthu et al. [88]	✓		✓		✓
Savarimuthu et al. [89]	✓		✓		✓
Savarimuthu et al. [90]	✓		✓		✓
Savarimuthu et al. [91]	✓		✓		✓
Savarimuthu et al. [92]	✓		✓		✓
Sen and Airiau [93]	✓		✓		✓
Cranefield et al. [34]		✓		✓	
Serramia et al. [94]		✓		✓	
Swarup et al. [97]	✓		✓		✓
Villatoro et al. [101]	✓		✓		✓
Villatoro et al. [103]	✓		✓		✓
Villatoro et al. [102]	✓		✓		✓
Vouros [105]	✓		✓		✓
Yu et al. [106]	✓		✓		✓

norm or chooses the same action. The terms emergence and convergence appear to be used interchangeably in the literature. Of the norm emergence mechanisms investigated, only a few considered convergence when 100% of the population chooses the same action, for example Villatoro et al. [101], Mungovan et al. [76] and Savarimuthu et al. [88]. Realistically, 100% convergence is improbable, and is not straightforward to achieve, consequently the acceptable convergence rate is 90% [60] for the rest of the literature cited in this paper except for one other, namely, Savarimuthu et al. [89]. Savarimuthu et al. [89] observes norm convergence when more agents choose that norm over any other competing norms, bringing the percentage down considerably depending on the number of competing norms. This view can potentially result in an undesirable situation of unstable or oscillatory norm emergence, when there is a large action set, or there are numerous competing norms. In their 2011 survey, Savarimuthu and Cranefield observe from the simulation models reviewed that percentages range from 35 to 100.

Sub-conventions [102,103] are an obstacle to 100% convergence. They can emerge in agent societies, where a smaller set of agents converge to a norm that is different from the norm of the majority of the population. These agents are usually located in isolated regions where they are not likely to interact with the larger population [93,102]. Villatoro et al. [102] refers to regions of a network that can maintain a sub-convention as Self Reinforcing Substructures (SRS). Unfortunately, these sub-conventions can persist throughout the simulation run if no mechanism is in place to resolve them. Villatoro et al. [102] advance that agents must be employed with the ability to recognise when they are in a SRS, and employ social instruments to resolve such a dilemma.

Hu and Leung [57] study local norm emergence, similar to the concept of sub-conventions, versus global norm emergence. They study a community network structure where there are dense connections in small communities and sparse connections between communities within a population. Results show that from a set of norms with equal payoff, diverse norms emerge in the system within individual communities: local norm emergence, with no global norm emergence observed.

In the literature, there is little or no mention of how long it is acceptable for this percentage to persist to really be able to say the norm has emerged. The observation is that of instant emergence, where norm emergence is observed the moment it happens. The premise is that the norm once learnt will persist over time. However flip flopping between “norms” may be possible, therefore norm stability should be a consideration in norm emergence and is furthered discussed in Sect. 4.6. Though the stability of the norm that emerged is not a major consideration in the emergence literature, the stability of norms is investigated in the norm synthesis literature for example in Morales et al. [73].

2.5 Source of the emerging norm

As in human society, acceptable behaviour of the occupants of the society are sometimes observed or directed by individuals of influence or individuals within some social circle. Some societies like the ones in Savarimuthu et al. [88] introduce the concept of role models and a norm advisor, which mimic the concept of influential people in a human society. Agents seek advice from these types of agents—role models and a norm advisor—thereby reducing the computation requirement for having to determine the best action among a set of possible actions. Agents however still have the ability to accept or reject the advice of a role model [88].

Agent societies can use these ideas to inject agents that can spread the agenda of the designers of a system, as these mechanisms are successful in ensuring that a norm propagates within a society. In 2007, Savarimuthu et al. proposes a variation to the role model mechanism employed by [88]. In their framework agents do not whole-heartedly adopt the norm of the role model when they choose to accept their advice. Instead agents modify their Personal Norm (a numeric value) to move closer to their role model’s norm. At the end of each iteration, agents choose their role model by sending a request to their best performing neighbour. Balaraman and Singh [13] investigate the influence a role model or team leader can have on other agents. They observe that teams in an organisation develop unique norms from constant interaction with each other [13]. Results from Balaraman and Singh [13] show that a team leader serves as a role model for the group and the behaviour of the team leader influences the members. When a team leader violates a norm, it weakens the norm and encourages team members also to violate, and similarly for compliance [13]. This shows the effect a role model adopting or violating a norm can have on other members of the group [13].

Similarly, norm entrepreneurs [55] are inserted into the system for the sole purpose of suggesting a new rule or norm to all agents, who then replace their worst performing rule with that one. Results suggest that the norm which emerges in the system is the norm suggested by the norm entrepreneur. This is evident even if a different norm was prevailing beforehand, and as long as the reach of the norm entrepreneurs is a subset of the population, of around 30%. It is noteworthy that though norm entrepreneurs are responsible for which norm emerges, their existence has no effect on the rate of emergence [55]. Norm entrepreneurs are able to influence the norm which emerges because they act as fixed agents utilising a single strategy every time when interacting with other agents. This helps to reinforce the action for learning agents. A drawback however, is that norm entrepreneurs can potentially be used to undermine a society.

Norm emergence from the bottom up, though useful, is not suitable for every situation. There are situations that exist where norm emergence can be difficult or expensive [42] and centralised distribution is also undesirable. Those cases require the injection of specialised agents to influence convention emergence. These agents may be planted in the society sometimes at pivotal positions [42] to act as role models.

2.6 Norm life cycle

The study of norm emergence cannot be accomplished without an understanding of the norm life cycle. The widely accepted norm life cycle based on simulation studies include the following three stage process: norm formation/creation, norm propagation and norm emergence [86]. Andrighetto et al. [7] however express the norm life cycle as a cyclic process encompassing the following four stages: (i) norm generation, (ii) spreading, (iii) norm stability, (iv) norm evolution. The process is perceived to begin with the generation stage and restarts after the evolution stage. In evolution, some norms may decay as new ones arise, they may evolve and others may become codified into law.

Norm formation or creation here refers to the introduction of a norm into the system. Norm propagation refers to spreading or distribution of norms to agents within the system and norm emergence is where a percentage of the population is observed to be adopting the norm. Savarimuthu and Cranefield [86] saw the need for, and proposed an expanded five stage norm life cycle which they posit “broadly captures the processes associated with the norm life-cycle”: (i) norm creation: norms are created by some mechanism, (ii) identification: agent becoming aware of the norm, (iii) spreading: distribution of the norm, (iv) enforcement: sanctioning mechanism to discourage violators, and (v) emergence.

We believe however that it is remiss to define norm emergence as the point of reaching the threshold, without considering the process which led to that stage. We cannot observe a percentage of the population following a norm, if the norm had not first existed in the population through the action(s) of an agent or a group of agents. Additionally, it is necessary for the existing norm to be spread to other agents through observation, communication or interaction before the emergence of the norm is observed. Consequently in order to fully situate the research on norm emergence, we put forward a refined definition of norm emergence. We posit that norm emergence should be defined as the “process” whereby a population of agents reach a predefined threshold of agents following the same norm. This means the norm emergence process will then include the creation and spreading of the norm and finally culminates with the observation of a percentage of agents following the norm, making the norm emergence process what is now widely accepted as the norm life-cycle. The norm life

cycle for societies with implicit or explicit norm representations generally follow the same stages but the implementation of those stages are vastly different.

2.6.1 Norm life cycle with implicit norms

The norm life cycle in a society with implicit norms usually exhibits the following stages: (i) creation, (ii) propagation, (iii) emergence. Norm creation refers to the initial or predetermined strategy to play from a set of available actions. This may be a fixed strategy, randomly selected or selected based on a predetermined probability. This is followed by norm propagation, where the norm is learnt or spread by and to other agents. Norm propagation in societies with implicit norms is not usually deliberate. In some cases agents are not aware that their behaviour is being observed and mimicked, and in other cases they share this behaviour willingly with agents who request this information.

Norm propagation is not achieved by the norm being communicated directly to agents, such as “there exists a norm to do this”, but by the observation of an agent’s action or the payoff an agent receives from taking an action. Alternatively norm propagation can take the form of the appropriate strategy being communicated via information sharing. One effective way of propagating a norm is using a leader to inform the followers, which is known as oblique transmission [23]. However this is not applicable in all scenarios. According to Boyd and Richerson [23], there are three ways by which a social norm can be propagated from one member of the society to another. They are vertical transmission (from parents to offspring), oblique transmission (from a leader of a society to the followers) and horizontal transmission (from peer to peer interactions) [23].

The norm propagation phase is a combination of the identification and spreading sub-processes identified by Savarimuthu and Cranefield [86] which is also observed in societies with explicit norm representation. Agents observe an action in a given situation and choose to do the same when faced with that situation in the future, without the explicit concept of a norm and norm adoption. Then finally, emergence where a percentage of agents execute this action for every instance of the triggering situation. This cycle is implemented in a large body of literature, from which we select Villatoro et al. [101], Mukherjee et al. [74], Savarimuthu et al. [89] and Vouros [105].

2.6.2 Norm life cycle with explicit norms

The norm life cycle in a society with explicit norm representation exhibits the following stages: (i) creation, (ii) identification, (iii) adoption, (iv) propagation, (v) emergence. The focus is on an agent recognising a norm and making the decision to adopt the norm. Norm creation is when a norm is introduced into the society. Usually the source of the norm is not mentioned, as the focus is on the agents recognising and adopting the norm. This is in contrast with the literature on norm emergence with implicit representation however, as it is usually clear where the new norm originates, for example a normative advisor [87], a role model [13,87,88], a norm entrepreneur [55] or another participating agent.

Creation is followed by the norm identification phase, sometimes also referred to as norm recognition. Since the norm is explicitly represented, agents need to become aware of it, which can be achieved through observation and inference, or by communication of the norm. At this stage, agents might place an internal representation of the norm within their beliefs [6]. However as in the case of Lee et al. [61], agents do not make an internal representation of the norm but refer to an external representation during their decision making. The internal

and external representations of explicitly stated norms are discussed in more detail in Section 2.3.

After the identification phase, agents are required to make a decision about the adoption of the norm. The assumption is that agents would only internalise a norm that they have decided to adopt. However internalisation is usually where the agent recognises that there is a norm existing, whether by communication or observation. The agents have to store the norm internally—have it as part of their beliefs—in order to be able to reason about its adoption. This adoption phase, or reasoning about adoption, results in the agent taking a position on whether to adopt or ignore the norm. To do so an agent needs to consider if the new norm conflicts with existing norms and goals. Conflict resolution is a normal process during this stage. A review of conflict resolution mechanisms is discussed in detail in Santos et al. [85].

The decision for or against norm adoption is normally captured in agent beliefs and some agents drop or remove the norms that they have decided not to adopt [41]. Andrighetto et al. [6] believe that agents adopt norms unless they have good reason to do otherwise. Some societies, when an agent adopts a norm, will apply the norm in every relevant situation in the future, as with the normative agents in Criado et al. [35], while others will reason about adopting the norm on a case by case basis as with the graded normative agents in Criado et al. [35] and agents in dos Santos Neto et al. [41] and Bench-Capon and Modgil [18].

Norm propagation, when used, is usually deliberate, where agents communicate the existence of a norm to other agents. In some situations agents may decide on the existence of norms through their observations of other agents' actions or punishments. The receiving agent must now go through the process of identification and adoption.

Research on explicit norms typically ends after the agent has adopted the norm. The emergence of a norm could be inferred, but this is not made explicit or even mentioned. We suggest that emergence can be observed when a predetermined threshold of agents have adopted the norm. This cycle, excluding the emergence stage, is implemented in, for example Boella and van der Torre [19], dos Santos Neto et al. [41], Andrighetto and Conte [5].

2.7 Discussion/summary

In this section we discussed the reasons why agents may adopt or violate norms. We then offer an interpretation of norms by describing the two main perspectives of norms: the prescriptive approach, where norms are described using deontic logic and the emergence approach, where norms are seen as a preference action from a set of actions. The literature on norm emergence is based primarily on the emergence perspective and implicit norm representation as demonstrated in Table 1, where the majority of the literature discussed are categorised accordingly. Further, we looked at what is norm emergence and how it is achieved in societies. Though the literature of norms examines them from different perspectives and considers the norm life cycle differently, it is useful to understand each perspective. The benefit to researchers from either side are twofold, firstly to highlight the commonalities within the two perspectives and secondly to identify the differences towards the unification of the approaches [32] and facilitating norm emergence in a society with both approaches to norms.

We summarise that the literature with implicit representations of social norms aligns with the emergence perspective of norms, where norms are understood as a preference behaviour, while the literature with explicit representations of social or legal norms aligns with the prescriptive perspective of norms (Table 1). The emergence perspective of norms is commonly associated with the study of norm emergence. The prescriptive perspective of norms is commonly associated with the study of the norm life cycle, an agent's internal process

of norm adoption, the specification of appropriate normative systems and norm synthesis. Social norms or the norms of the emergence perspective are usually referred to as conventions in the literature on the prescriptive approach. In the next section we take a closer look at the features of the various norm emergence approaches.

3 Characteristics of norm emergence mechanisms

Each simulation model or mechanism to support norm emergence is different, however there are some characteristics, whether of the society or the agents within the society, that have been studied across all the mechanisms in the literature. At one end of the spectrum there are those characteristics that help to facilitate or speed up the emergence, while at the other, there are characteristics that can delay the emergence of a norm, or more seriously, impede that emergence altogether. In analysing the commonalities among the mechanisms in the literature, we have identified the following characteristics which support norm emergence in various ways: (i) social topology, (ii) agents' cognitive abilities or decision making mechanisms, (iii) propagation mechanism, (iv) agents' use of learning algorithms, (v) observation capabilities, (vi) offline methods or online methods.

3.1 Social topology

The social/interaction or network topology specifies how agents are connected in the society, which ultimately determines how they interact. The social topology of the agents has a direct impact on the emergence of norms. Agents usually interact with agents they are connected to, therefore more connections typically result in more interactions. This is especially true for the cases where norms are learnt based on the interactions between agents, as agents adopt the strategies of their neighbours. Yu et al. [106] posit that agent interactions should not be ad-hoc but should mimic the real world where interactions are structured and occur based on relationships and social networks. Yu et al. [106] present scale free and small world networks as those that most closely mimic the connections in real world interactions. Similarly, Savarimuthu et al. [87] investigate norm emergence over three types of networks: fully connected, random networks and scale free. They find that scale free is better suited to norm propagation largely corroborating the position of Yu et al. [106], though not considering small world networks.

Random networks, though not representative of human real world interactions, are still a valid type of network for study in the agent literature, because they may be seen as representative of a network of agents that utilise an application without having knowledge of who the other users are or their whereabouts e.g. a file sharing application. Though they do not represent typical human societies, they are still useful in simulating some environments.

One can imagine that some network topologies have a greater likelihood to support norm emergence because it ensures that an ample number of agents are connected to other agents and that those connections are sufficiently distributed. The connectivity of a network influences the rate at which information can be shared among the network since information sharing is achieved via these connections. Subsequently the density of the connections, and neighbourhood size, determines the number of agents that one agent can reach, which in turn affects the rate at which agents can share information. For example consider the structure of the scale free and small world networks where agents are connected in a set of large and

small hubs, that also connect to outliers or leaves, versus networks such as rings that have agents connected only to two other agents, or fully-connected networks.

Research has corroborated that network topology or structure has a significant impact on the interactions that can occur between agents as the number of interactions, and whether two agents actually interact, etc. are all determined by the network structure. Mahmoud et al. [67] demonstrate their model's effectiveness for achieving norm emergence on lattices and some configurations of small world network topologies, but their model is ineffective in scale-free networks. This is due to the nature of connections in the scale-free networks, as the meta-norm model relies on observations of neighbours to work. There exists one contrasting view in the work of Hao et al. [52], where underlying topology had no effect on the emergence of norms using the mechanism presented. These results speak to the effectiveness of the strategy used as this is not observed in any of the other works investigated. Hao et al. [51]'s study of individual action learners (IALs) and joint action learners⁴ (JALs) also show the JALs being able to converge to the optimal policy on all topologies for all types of games investigated, but with varying rates and successes. The success and rate of convergence for the IALs however was significantly affected by the underlying topology.

Neighbourhood size refers to the number of agents connected in smaller groups within a given network topology. Mukherjee et al. [75] show the effect of varying neighbourhood size for agents within a given distance in a grid world, and conclude that smaller neighbourhood sizes result in faster norm emergence. Additionally, they find that there are several other characteristics of networks that need to be considered namely, average path length, degree distribution, clustering coefficient, diameter, etc.. Highly-connected agents play more games, meaning they interact more and become aware of the best norm before poorly-connected agents [87].

Neighbourhood size also determines how connected agents are in the wider populations where small neighbourhoods mean tightly-connected smaller groups. In Savarimuthu et al. [87] results show that when the average degree of connectivity increases in a network, so does the rate of convergence in random and scale free networks. So more connected networks result in faster convergence rates and connectivity here speaks to both neighbourhood size and topology.

Neighbourhood size and topology can have a direct impact on convergence of a norm within a society because of the existence of sub-conventions. Sub-conventions can emerge in agent societies where smaller sets of agents converge to a norm that is different from the majority population [103]. These sub-conventions can persist and affect full convergence of the preferred norm. The topology of the network can mean some agents are poorly connected, and can have adopted a different norm within the small group. In addition, it could mean that agents within these small groups have little or no interaction with the rest of the network so they cannot be influenced by their actions, as observed in Hu and Leung [57].

3.2 Agents' cognitive ability or decision making mechanism

Cognitive ability or the decision making mechanism refers to how the agent makes a decision. An agent can be programmed to follow a single action every time or randomly follow an action selected from a set of actions. An agent might also be programmed to determine its action based on the situation or, more sophisticated, internalise the state of the environment and intelligently select the most appropriate action. In order to understand the research landscape we put forward a categorisation of agents' cognitive capabilities into low, medium and high

⁴ Joint Action Learners (JALs) were introduced by Claus and Boutilier [30].

as depicted in Table 2. Low-cognitive reasoning ability connotes an agent with a basic action selection strategy, such as fixed action selection or random choice among several actions. Medium connotes an agent with the ability to choose from a set of choices based on the given context. We will also classify agents that use a learning algorithm to learn the best response action based on previous history as medium (Table 2). And finally, high is for an agent that is able to reason about the environment, norms and the actions available to it and choose an appropriate action based on these.

The cognitive ability of an agent can have an effect on both the rate of emergence and emergence in general. Some works demonstrate though that even low-cognitive agents can converge to a norm in a society through interaction and simple mimicking or through reaching a consensus on norms [46,48]. Agents in Ghorbani et al. [48] and Ghorbani and Bravo [46] have low-cognitive abilities that either imitate the strategy of their best performing neighbour, or simply try different combinations of predefined values, but doing so only when their current strategy is performing poorly.

Some systems utilise low-cognitive reasoning ability agents with a fixed strategy. If all the agents in a system utilise the same action, it may be concluded that a norm has emerged, when in fact those agents were all pre-programmed with a fixed action, this “norm” was designed into the agent rather than emerged naturally. Alternatively, if a society of fixed strategy agents is composed of sets of agents with different strategies by set, that society can never have a dominant strategy [29], unless emergence is seen from the perspective of Savarimuthu et al. [89], where once a majority follows one norm, then emergence is achieved. Interestingly, studies by Hao and Leung [49], Mukherjee et al. [75], Sen and Airiau [93] demonstrate how the introduction of a small number of fixed strategy agents in a society with learning agents can influence which norm emerges.

Villatoro et al. [103] utilise low-cognitive agents with random choice among several actions, where at first the agents randomly choose an action from a set of n actions, to interact with another agent and similarly Savarimuthu et al. [90] model agents who can choose either to litter or not to litter in a given interaction. Brooks et al. [25] model agents who choose a strategy based on a given probability towards that particular strategy. The outcome of the interaction determines if agents increase the likelihood or probability of choosing that action in the future [25]. Savarimuthu et al. [87] present a similar concept where agents have two strategies to choose from: a group norm and a personal norm. Agents use an autonomy value that determines how often they utilise which of the strategies. The majority of norm emergence systems investigate agents with low-cognitive abilities as Table 2 demonstrates.

Vouros [105] models agents that may be considered of medium-cognitive ability, in which an agent can fulfil a combination of different roles, which sometimes have conflicting or incompatible goals. In Vouros’s scenario, agents must coordinate to carry out a joint task, or utilise a joint resource, by choosing a mutually convenient time to do so. Agents have a specific strategy for scheduling, based on their role, and must decide which strategy to use to schedule the resource. They decide on a strategy that is either more suited to the team leader or to the coordinator role [105].

High-cognitive ability agents are not commonly used in work that investigates norm emergence, as shown in Table 2, where only low or medium cognitive abilities are inferred. This situation is probably premised on the general view that emergence must involve simple interactions. However the definition of emergence is limited to the observation of a particular phenomena for a predetermined percentage of actors. There is no premise that the observation must be preceded by simple interactions. We postulate that high-cognitive agents who deliberately decide on an action can, over time, demonstrate emergence when a predeter-

Table 2 Summary of characteristics of the emergence mechanisms studied in the literature surveyed

	Social topology	Cognitive abilities	Propagation mechanism	Learning algorithms	Learning method			Observation capabilities			Offline methods
					Pairwise	Collective	Global	None	Private	Local	
Balaraman and Singh [13]		Low	Normative advisor & punishment								✓
Beheshti et al. [17]		Medium	Learning from interaction & punishment	✓	✓			✓			
Beheshti and Sukthankar [16]	✓	Low	Punishment		✓				✓		
Brooks et al. [25]		Low	Learning from interaction					✓			
Chakrabarti and Basu [28]		Low	Learning from interaction	✓	✓						
Chao Yu et al. [29]	✓	Medium	Learning from interaction	✓	✓			✓			
Franks et al. [42]	✓	Low	Learning from interaction	✓	✓				✓		
Frantz et al. [43]		Medium	Learning from interaction	✓	✓			✓			
Dascalu et al. [37]		Low	Learning from interaction	✓	✓			✓			
Ghorbani et al. [48]		Low	Learning from interaction		✓						✓ *
Ghorbani and Bravo [46]	✓	Low	Learning from interaction	✓	✓				✓		
Hao and Leung [49]		Medium	Learning from interaction	✓	✓				✓		
Hao et al. [50]		Medium	Learning from interaction	✓	✓				✓		
Hao et al. [51]	✓	Medium	Learning from interaction	✓	✓				✓		
Hao et al. [52]	✓	Medium	Learning from interaction	✓	✓			✓			

Table 2 continued

	Social topology	Cognitive abilities	Propagation mechanism	Learning algorithms	Learning method			Observation capabilities			Offline methods	
					Pairwise	Collective	Global	None	Private	Local		Global
Hassani-Mahmooei and Parris [53]		Low	Learning from interaction & Punishment	✓							✓	
Hofmann et al. [56]	✓	Medium	Learning from interaction	✓							✓	
Hoffmann [55]		Low	Normative advisor, learning from interaction & punishment	✓							✓	
Hu and Leung [57]	✓	Medium	Learning from interaction	✓				✓				
Lopez [65]	✓	Low	Punishment								✓	
Lotzmann et al. [66]		Medium	Punishment								✓	
Mahmoud et al. [67]	✓	Low	Punishment								✓	
Mukherjee et al. [74]	✓	Medium	Learning from interaction	✓							✓	
Mukherjee et al. [75]	✓	Medium	Learning from interaction	✓							✓	
Mungovan et al. [76]	✓	Medium	Learning from interaction	✓							✓	
Ohtsuki et al. [78]	✓	Low	Learning from interaction	✓							✓	
Peleteiro et al. [80]	✓	Low	Learning from interaction	✓							✓	

Table 2 continued

	Social topology	Cognitive abilities	Propagation mechanism	Learning algorithms	Learning method			Observation capabilities			Offline methods
					Pairwise	Collective	Global	None	Private	Local	
Riveret et al. [83]		Medium	Punishment								✓
Savarimuthu et al. [87]	✓	Low	Normative advisor & role models								✓
Savarimuthu et al. [88]	✓	Low	Role models								✓
Savarimuthu et al. [89]	✓	Low	Role model								✓
Savarimuthu et al. [90]		Low	Punishment								✓
Savarimuthu et al. [91]		Low	Learning from interaction & Role model	✓							✓
Savarimuthu et al. [92]		Low	Learning from interaction					✓			
Sen and Airiau [93]	✓	Low & medium	Learning from interaction	✓							✓
Swarup et al. [97]	✓	Low	Learning from interaction	✓							✓
Villatoro et al. [101]	✓	Medium	Learning from interaction	✓					✓		
Villatoro et al. [103]	✓	Low	Learning from interaction	✓					✓		
Villatoro et al. [102]	✓	Low	Learning from interaction	✓							✓
Vouras [105]	✓	Medium	Learning from interaction	✓							✓
Yu et al. [106]	✓	Medium	Learning from interaction	✓				✓			✓

Text in cells with black text represents characteristics that are explicitly stated/investigated. Text in cells in bold represents characteristics that have been inferred from the analysis of the paper. The ✓ * in the single case of Franks et al. [42] denotes that an offline method was implemented but its results were investigated in an online method

mined percentage adopt a particular action. We believe this opens an area for future research which we explain further in Sect. 4.

3.3 Propagation mechanism

A propagation mechanism is the way in which everyone in the society is made aware of the existence of the norm. Mechanisms described in the literature include: (i) normative advisor, (ii) role model, (iii) learning from interaction, and (iv) enforcement or punishment, each of which we now discuss in more detail.

3.3.1 The normative advisor

The normative advisor is a leader agent who has global knowledge and access to broadcast information or can change global parameters. This concept is similar to a role model (see below), however only one normative advisor exists who proposes strategies to the population. This method of propagation is classified as oblique transmission. In some cases, once the normative advisor proposes a new strategy it is adopted, while in others, agents have the autonomy to decide whether to adopt the proposed strategy. Savarimuthu et al. [87] considers a norm advisor mechanism, where the advisor updates the group norm value for all agents by computing an average strategy based on the successes and failures of all the agents within the society. Agents whose strategy was unsuccessful will change their personal norm closer to the group norm. Similarly, Balaraman and Singh [13] utilises a team leader, though the team leader does not explicitly inform agents of their intent to adopt or violate a norm: when team members observe either action, it influences their decision to do the same, which coincides with the concept of a normative advisor. Norm entrepreneurs can also be considered within this category, as they are inserted into the system to suggest a norm to the agents, which is usually the norm that is observed to have emerged within the system [55]. If a single norm entrepreneur is used, we can classify it as a normative advisor, but if multiple norm entrepreneurs are used, then they can be appropriately classified as role models.

3.3.2 The role model

Role models are agents distributed throughout the population that can advise other agents on strategies. Agents select their role model and decide whether to accept their advice or not. The role model mechanism is a useful one, as it mimics how people's actions are influenced by each other and consequently how the actions of an agent can affect the actions of other agents in the future. Moreover, the actions of highly influential agents are likely to be reproduced by the rest of the population. This method of propagation is classified as horizontal transmission but can also represent a form of oblique transmission if we think of role models as being multiple leaders. Studies by Savarimuthu et al. [87,88,91] each demonstrate success with employing role model agents within the society, where agents can choose a role model among another agents, based on their success within the society. An agent generally copies the action of the role model. A slightly different approach by Savarimuthu et al. [87] has agents modifying rather than copying the strategy of their role model, to move closer to that of the role model. This is effective, since the strategy here is a numeric value which can be adjusted in this way, and which serves to illustrate the principle, but it would be more complicated with a less abstract representation of strategy.

3.3.3 Learning from interaction

This is when agents use some learning mechanism and adapt their strategies based on either the outcome of an interaction(s), the actions of another agent, or information about the success of another agent's strategy. The type or nature of interaction, whether agents interact with random agents or only their neighbours, is an important issue in norm emergence. This method of norm propagation is widely studied in the literature, see Table 2, and is classified as horizontal transmission.

From the literature, we see that research where agents learn from interaction consider how the agents in a population interact: either with random agents or nearby agents. While random interactions might seem unrealistic at first sight, there are applications of multiagent systems where random interactions are more reflective of what actually takes place e.g. file sharing or flight reservations. Investigating norm emergence using random interactions is quite popular [25,28,49,50,74,89,93,101–103]. In random interactions, at any time step, two agents are randomly chosen to interact with each other. Results show that norms do emerge in such cases and the rate of emergence is deemed acceptable in the context of the experiments reported.

Different from researches on random interactions in its true sense, Mungovan et al. [76] propose Weighted Random Interactions, in which they argue that random interactions in the real world are not truly random, but that people are more likely to come into contact with particular people, for example, friends of friends, people in the same bus, at the same party or in the neighbourhood, rather than a random individual. Therefore, they investigate the effects of weighted random interactions on emergence, by allowing random interactions with people based on social distance in a network. Results show that increasing the frequency and weighting of random interactions leads to faster convergence to a norm, given two alternatives, and also higher levels of norm conversion [76]. This result is consistent with results from the literature on purely random interactions.

On the other hand, the literature also includes many investigations of interactions with nearby agents or close neighbours [53,57,80,87,92,97,106]. Researchers argue that this method is preferred as it reflects human interaction. However it is not without its challenges as it can give rise to sub-conventions [103].

3.3.4 Enforcement or punishment

This characterises when agents can become aware of the existence of a norm governing a particular action because a punishment is received for executing that action. Savarimuthu et al. [91] refers to experiential learning, or learning by doing, where they suggest that an agent can learn about norms by doing an action and being punished for doing so. This method of propagation is classified as horizontal transmission when distributed enforcers/punishers are used, and oblique transmission when a central authority punisher is used. Savarimuthu et al. [90] investigate how norms can be established by using punishment. They use agents that act as punishers when they observe other agents violating norms. They also investigate the use of a common knowledge source that informs agents of the state of the environment; they find that as agents become aware of the unacceptable state of the environment, they begin to punish the violators which results in behaviour change [90]. Similarly, Lotzmann et al. [66] utilise norm invocations, which they define as “utterances by an actor when he or she takes offence or is satisfied by an action of another actor”. There are negative and positive norm invocations, where negative norm invocations act as punishments. Agents issue negative norm invocations to other agents when they observe them behaving in a way that the observing

agent finds unacceptable. Agents who receive a negative norm invocation for an action will reduce the likelihood of choosing that action in the future [66].

Enforcement or punishment as a norm propagation method uses distributed enforcer agents to punish non-conforming agents. The intention is to convert non-conforming agents to conforming ones, which can be achieved since rational agents usually aim to increase their payoff, but punishments provide either negative payoff, no payoff or small payoffs, compared to acceptable actions. Agents will then change their behaviour in order to avoid these punishments. Notwithstanding, there are instances when agents are allowed to continue violating because their violations are not observed [67], or enforcer agents choose not to punish [90]. Therefore for enforcement/punishment as a norm propagation method to be successful, there needs to be an adequate number of enforcer agents [90] who are able to observe violating agents [67] which is dependent on their social topology. It is important to note though that having too many enforcer agents is also ill-advised [14].

3.4 Agents' use of a learning algorithm

Savarimuthu et al. [91] investigates the effects of active learning, which can be either experiential, observational or communication based, on the learning of norms. They posit that a mechanism that combines the three would be useful in negating the deficiencies of each approach in isolation, for example lying in communication based methods [91]. Results show that the combination results in faster norm convergence, when compared to the experiential approach alone, and experiential with observational. Unexpectedly though, using the combination approach with the presence of only 2 liars in the system, can destabilise any norm that has emerged much faster, than when an observational with communication based approach is used.

Chao Yu et al. [29] investigates the effects of three learning strategies on norm emergence namely Q-Learning, WoLF-PHC and Fictitious Play (FP). Results show that Q-learning evolves fastest to a norm, followed by WoLF-PHC, then Fictitious play. Similarly, Mukherjee et al. [75] and Vouros [105] as part of their research study the effects of the same group of algorithms with identical results. However Vouros [105], in addition to the above mentioned three, also utilises Highest Cumulative Reward-based models. Given that Q-Learning performs the best of the four algorithms, for those researches where the learning algorithm remains constant, it could potentially be the one of choice—see Table 3—as it is for Savarimuthu et al. [91] and Beheshti et al. [17]. A significant proportion of the literature analysed investigate an agent's use of a learning algorithm, see Table 2.

Chakrabarti and Basu [28] utilise a cellular automata technique for selecting the learning algorithm used by agents. This means that an agent observes the actions of their neighbours to the left, right, top and bottom, they observe the learning algorithm used, and based on the most popular one, the agent adjusts its own learning algorithm. Emergence of norms using the cellular automata technique is more rapid than without, and the payoff is better. However with a coordination game, the convergence rate is better but the payoff is poor.

In addition to studying the effects of different learning algorithms, Chao Yu et al. [29] also investigate the effect of having agents with no learning ability in the society, which they refer to as non-learning agents. In a society of only non-learners with a fixed action, a norm can never emerge as would be expected, while a society with non-learners and a fixed observation strategy can benefit from a small number of learners and achieve emergence [29]. The introduction of a small proportion of fast learning agents can facilitate norm convergence in a large scale agent society and the results corroborate this, as with the introduction of a

small proportion (e.g., 10%) of Q-learning agents in the population, the convergence time is steeply reduced and further increasing the proportion of Q-learning agents steadily decreases the convergence time [29]. However a norm emerges very slowly in a society of non-learners with conflicting strategies. Based on the location of agents within the society, one may find that the society emerges to two conflicting norms based on the location of learner agents within the society. In the experiment done, after 500 runs the society had not yet converged to one norm [29].

A related factor for the learning algorithm is whether the method/strategy used for learning is pairwise or collective. Yu et al. [106] and Hao et al. [52] study how collective learning compares to pairwise learning. Pairwise learning refers to an agent learning or adopting an action based on a single interaction with another agent in a given round, while collective learning refers to the agent's ability to learn from interactions with several agents in a given round. With collective learning, an agent's action at any one time is determined by an aggregation of all its best response actions into a single action that will be played with each neighbour in the following run. Results show that a social norm will evolve in societies using both types of learning [52,106]. Nearly all agents reach the consensus using collective learning and the rewards or payoff are greater for collective learning [106]. Agents using collective learning are able to converge quickly on a norm, so the emergence rate is also quicker. Hao et al. [52] found that a norm emerged each time using their collective learning strategy and that varying network topology had no effect.

The majority of the literature on norm emergence utilises a pairwise learning method as shown in Table 2, though evidence suggests collective learning achieves faster emergence with a higher adoption rate. Collective learning is considered to mimic human decision making as is multiagent reinforcement learning. The nature of multiagent reinforcement learning's use of interaction history in the decision process is what aligns with human decision making. Multiagent reinforcement learning algorithms refer to an algorithm in which an agent chooses an action based on past interactions with multiple agents. Multiagent reinforcement learning algorithms can be applied on interaction data gathered from both pairwise and/or collective learning methods. In the literature, multiagent reinforcement learning is widely used for example Chakrabarti and Basu [28], Jianye et al. [59], Villatoro et al. [101], Chao Yu et al. [29], Hao and Leung [49], Hao et al. [50]. The drawback of both collective learning and multiagent reinforcement learning methods is that agents must have sufficient memory to keep track of past interactions, however in some situations this is not readily available. Villatoro et al. [101] shows how the size of memory affects how much history is recorded, and how it also affects the speed of decision making and ultimately norm convergence. Alternatively, some studies investigate agents choosing an action based only on the current interaction and no information about past actions is required, namely Mukherjee et al. [74,75], Villatoro et al. [103], Sen and Airiau [93].

Agents can make decisions either by using local exploration, which refers to determining the best response action based on the agent's own history, or global exploration, which takes the agent's chosen best action and checks it against best actions in the world, making a selection of the best action among all. The agent updates its action with the best-response action and plays the same action with all agents in the next time step. This activity is repeated for each time step. If agents are going to be designed to learn norms based on past interactions with other agents, they need to be equipped with a learning algorithm to facilitate this. The rate of emergence achievable is purportedly the determining factor in this decision, however the learning algorithm cannot be considered in isolation. The method used for learning, whether pairwise or collective, must be identified as it directly impacts the rate of emergence.

3.5 Observation capabilities

An important factor to consider when instituting mechanisms for the emergence of norms is whether agents have access to information about other agents. Some studies assume that agents do, but the opposite can also be true and consequently the emergence of norms via private interactions is a subject for study. In some cases, agents have access to only a subset of information or a limited amount of information. Most of the studies presented reflect agents having access to some information about another agent, see Table 2. Yu et al. [106] however did not require agents to know about their neighbour's payoff or actions, and this did not hinder emergence. Mukherjee et al. [75] also conclude that private information is sufficient to facilitate norm emergence in a society of learning agents.

There are different levels of access to information in the studies analysed. For example, Jianye et al. [59] builds on readily available information as the agent chooses a best action among its own best actions and those in the world. This means that the agent has access to the information about the best responses of all other agents. Perfect visibility refers to an agent having access to all the information about all other agents [13]. This is also referred to as global observation [103].

In the real world, global observation is not always practical, as it might not be possible or feasible to have agents know everything about all others. Alternatively, there is local observation where agents have access to information about a subset of agents. Local observation is predominantly studied in the literature as shown in Table 2. Hassani-Mahmooui and Parris [53] demonstrate the usefulness of having some information available, because agents can inform other agents of a defaulting agent by placing them on a blacklist, which must be made visible to all agents. This is essentially a compromise, by making small bits of information accessible globally, similar to the common knowledge source in Savarimuthu et al. [90].

Hao et al. [51]'s results show that joint action learners are able to converge to an optimal policy faster than individual action learners and with greater success over different topologies. The authors attribute this feat to the fact that joint action learners have access to more information than individual action learners. Joint action learners not only have information about their own history and history of their neighbours, like individual action learners, but also to that of their neighbours' neighbours as well. Consequently, we can conclude that the more information an agent has access to in their decision making, the faster the emergence of a norm. Collective learning strategies allow agents to interact with all their neighbours in a round, providing more experience to learn from at the end of a round. Studies investigating collective learning versus pairwise learning show that collective learning converges to a norm at a faster rate than pairwise learning [28,29,52,106]. Hao et al. [51] demonstrate the benefit of observation, as they are able to achieve the same effect as collective learning but without the need to interact with all the agents. They instead observe the history of interactions of their neighbours and their neighbours' neighbours.

Mahmoud et al. [67] discuss the use of meta-norms as a way for agents to observe the actions of the other agents and to punish agents when violating. Axelrod [12] describe meta-norms as the establishment of a norm that says an agent must punish those who do not punish a defection of another norm. This also highlights the issue of observability, as results show that in some topologies, where agents are limited to observing only those agents to which they are connected, some agents are able to continue violating norms, because their actions are less likely to be observed. The meta-norm model is consequently ineffective when agents are unable to observe the interactions of other agents. The effectiveness of the emergence strategy or learning algorithm is tied, in most cases, to the amount of information an agent has access to about another agent. The number of agents operating in the society, the nature

of the interactions and the sensitivity of the information will factor in to a decision of how much data can be made available to the agent, see Table 3.

3.6 Offline or online methods

Online methods refer to a mechanism in place during the runtime of the society being modelled. Norm emergence, whereby agents converge on a norm, is a form of online synthesis using a decentralised approach [90]. Most societies that simulate norm emergence using implicit norm representations can be considered to be adopting an online method—see Table 3—as the norm emerges in the society during the live interactions with agents. This can be clearly observed in Table 2 where the offline column is almost empty. The only exception is Franks et al. [42] who utilise an offline method to determine beforehand the optimal positions in a society where an agent has the most influence on the society. Then during runtime, agents can be injected at those positions determined earlier. An online method to this approach is also mentioned where agents are injected at optimal positions determined while the simulation is live. The process of synthesising norms, as discussed in Sect. 2.6, is crucial for the smooth running of a society by avoiding conflicting states. Societies that utilise explicit norm representation often must synthesise a set of norms, as opposed to having norms emerge.

Norm synthesis can be done both offline and online. Online synthesis is useful for determining the set of norms that avoid conflicting states while the system is running [4,26,68,70–72]. Offline synthesis, though useful, is difficult to do when all possible states of the system are not known during design time or can change based on interactions of actors within the system. In those cases, online synthesis is more appropriate to synthesising norms during runtime, allowing new norms to be synthesised as circumstances change. Morales et al. [73] present an offline approach to synthesising norms for complex coordination situations of a system by simulating simple interdependent games of potential conflict situations, where agents can fully perceive each other and coordinate. The norms determined by the simulation run can then be applied to the system. An alternative approach is illustrated in the scenarios describing weak (permission–prohibition) and strong (prohibition–obligation) norm conflict in Li et al. [63,64], while Corapi et al. [33], Athakravi et al. [9], Li [62] chronicle the development of a framework built on inductive logic programming, starting from an existing (possibly empty) norm set, observation traces and normative conditions over final (normative) states, to synthesize new norms or revise existing ones.

3.7 Discussion/summary

There are several characteristics of agent societies that affect norm emergence (convergence). The ones discussed above are limited to the characteristics analysed from the literature presented. We summarise the findings of this section in Table 3.

The social/interaction topology of the agents plays a key role in norm emergence as the connectivity of the agents determines how many other agents an agent can interact with and influence, and has a direct impact on the rate at which information can be shared, ultimately leading to emergence. Interestingly though Hao et al. [52] show that their mechanism for norm emergence was not affected by the underlying technology. Additionally, the network topology needs to be considered based on what type of society needs to be modelled, as in the case of modelling a natural human phenomena, the scale free or small world network is more appropriate (Table 3).

Table 3 Overview of the key findings of the characteristics of norm emergence based on our literature survey as discussed in Sect. 3

Characteristic	Types investigated	Findings
Social topology	Fully connected, grid, lattices, small world, scale free, random	human like—scale free and small world. Other—random, fully connected, grid, lattices. More connectivity results in more influence/reach which results in faster rate of emergence
Agents' cognitive abilities	Categories: low, medium, high	Only low and medium used. Even the most simple learning strategy, mimicking, can result in norm emergence
Propagation mechanism	Normative advisor, role model, learning from interaction, punishment	centralised—normative advisor, punishment vs. decentralised—role model, learning from interaction, punishment
Agents' use of learning algorithms	Q-Learning, WoLF-PHC, fictitious play (FP), highest cumulative reward based models used. Collective vs. pairwise learning method.	All result in norm emergence. Q-Learning outperforms all resulting in faster norm emergence. Useful when agent has no information about other agents, can learn based on previous history. Collective learning method outperforms pairwise learning but requires more computation and memory store
Observation capabilities	None, private—own history of previous actions and rewards, local—access to some or all information about neighbouring agents, global—access to some or all information about entire population	The level of observation available affects the rate of emergence where more information available to the agent results in faster norm emergence. However the context of the system determines which can be used, sensitive information or too many agents may prove global observation risky or impractical. Emergence can be observed in either method. The choice depends on the propagation mechanism used as well
Offline or online methods	Offline or online	The concept of emergence is undoubtedly an online method

The cognitive abilities or decision making mechanisms of agents is a useful characteristic, as even low-cognitive agents are capable of achieving emergence. The literature surveyed only considers agents with low or medium cognitive abilities, with the majority being low-cognitive ability agents (see Table 2). There is potential to investigate high cognitive agents and how they might affect the emergence of norms in a society and also to compare the rate of emergence based on the three types of agents discussed (low, medium and high cognitive ability).

Understanding the different propagation or emergence mechanism used is significant, as depending on the type of agent system required, one method will be more suitable than the other. For example, the role model and normative advisor mechanisms are in line with simulations requiring either distributed or central authority, respectively.

Mechanisms that require agents to have the ability to learn their strategy from their interaction history will have to deploy a learning algorithm. The Q-Learning algorithm consistently outperforms the WoLP-PHC and the Fictitious Play algorithm, and also Highest Cumulative Reward-based models in the literature and we propose should be the algorithm of choice for faster emergence rates (see Table 3). The method or strategy used for learning will depend on the constraints of computational power and memory capacity. Collective learning, though more effective than pairwise learning, does require an agent to interact with all of its neighbours, which will be computationally demanding, as will keeping track of all the interactions, which affects memory requirements.

The observation capability of agents is also important, as some approaches assume agents have global knowledge, while in other approaches agents do not keep track of any information, and also have no knowledge about their neighbours. The more information available to the agent, the faster a norm will emerge. The majority of the literature cited has agents that have access to some information about another agent (Table 2). The amount of information that an agent has access to will determine what mechanism can be employed for propagation, and also influence the learning algorithm and learning approach used. For example, agents cannot utilise a collective learning strategy if they have no or limited memory of past interactions.

Finally the emergence method employed can be applied either offline or online. Offline methods are employed at design time, before the simulation is executed, and online methods, where norms are allowed to emerge while the simulation is in execution (Table 3).

An understanding of how each characteristic affects norm emergence is vital in developing approaches to facilitate norm emergence. The choice of characteristics to include in any mechanism will depend on the requirements of the domain to be implemented and how it will be modelled. Section 4 provides additional details on how these characteristics can be used as a check-list for designing agent systems.

4 Future challenges and opportunities

Our analysis of the literature enables us to identify the following challenges and opportunities that exist in the study of norms and norm emergence. Table 4 highlights those identified and the key findings which we discuss in more detail below.

4.1 Design of multiagent systems

It is reasonable to utilise the characteristics that are identified to support norm emergence as design issues for new systems, where we may find that all characteristics may or may not

Table 4 Overview of future challenges and opportunities

Future challenge/opportunity	Findings/recommendations
Design of multiagent systems	All characteristics may not be applicable in all contexts. Different combinations of characteristics will be suitable when modelling different types of multiagent systems. The context of the domain to be modelled will aid in deciding which characteristics to utilise
Additional characteristics	There is an opportunity to study the effect of the emergence mechanism across varying types of games and for agents with a learning algorithm utilising different exploration strategies. The results of Hao et al. [51,52] with success across varying topologies necessitate further research from existing mechanisms to enable successful emergence irrespective of underlying topology
Online norm synthesis	There is need for additional mechanisms that enable centralised online norm synthesis and the introduction of mechanisms that enable decentralised online norm synthesis for prescriptive norm systems
High-cognitive ability agents demonstrating norm emergence	It is worthwhile to explore whether a population of high cognitive ability agents utilising varying non-simplistic action selection strategies would demonstrate the emergence of a particular action, when faced with a set of actions to execute in a given situation. We believe the true test of the emergence of a norm is observed, when agents faced with several actions of equal payoff would all (or most) independently choose the same action
Generic or adaptive framework to support norm emergence	There is potential to investigate a generic or adaptive mechanism that is capable of facilitating the emergence of norms in normative agent societies, irrespective of the characteristics of the society
Stable emergence	Norms that emerge but persist for a short time can potentially result in unstable systems. Research on the emergence perspective should consider norm stability, which features more in the prescriptive perspective literature
Prescribing both social and legal norms	Societies with prescriptive norms already have legal norms explicitly represented. Social norms existing as conventions are assumed to exist, but it would be difficult to determine if or when social norms emerge. If social norms are to be explicitly represented in addition to legal norms, they will allow us to speak confidently on the emergence of norms
Converting social norms to legal norms	Social norms can become legal norms as part of the norm life cycle presented by Andrighetto et al. [7]. The behaviour that emerges within a society, as social norms, can be prescribed as the legal norms for the society, allowing agents to participate in determining the norms that will govern them. Haynes et al. [54] present the concept of the engineering of emergent norms allowing the encouragement of beneficial behaviour and discouraging non-beneficial behaviours

not be applicable (Table 4). The characteristics can be used as a check-list for the design of new agent societies to support norm emergence. A designer can prioritise the characteristics based on the needs of the environment and select the appropriate method of implementation of each for their intended domain. For example, a designer will need to consider the interaction topology of the agents when deciding how best to have information shared among the agents. Small world and scale-free topologies are appropriate for modelling natural human interaction and should be the options selected if the society is intended to simulate a human phenomena.

However the designer must take into account that the small world and scale-free topologies can pose problems, where smaller fractions of the network can adopt a norm different from the entire society. The designer will then need to ensure that the propagation method chosen considers the effects of the smaller fractions and potentially consider employing rewiring and an observation mechanism as in Villatoro et al. [102,103].

If a particular propagation mechanism is priority, it will influence the choice of interaction topology. A norm advisor is capable of broadcasting to all agents within the network irrespective of their location, thereby negating the effects of the interaction topology. Conversely the success of the remaining propagation methods: learning by interaction, punishment or role models will depend greatly on the interaction topology of a network as agents use a distributed approach for propagation in each. A complementary concept to the role model mechanism would be to determine the most influential agent and their position in order to maximise the spread of information or behaviour within the population. This is the nexus of Franks et al. [42]’s work to investigate how to measure empirically an agent’s influence, or the point in a network with high(est) influence. After the influential points are determined, an agent is injected at those particular positions with a goal to spread a particular norm that is not evident in the society. Results show that the norm propagated by the injected agent successfully emerges within the society [42].

The cognitive abilities of the agent appear to be a less significant factor in the design of agent systems, since even low-cognitive agents can converge to a norm. This means designers can model the abilities of the agent based on the needs of the system and it is very probable that norms will emerge. The case of modelling high-cognitive agents to investigate emergence has to our knowledge not been studied and presents an opportunity for research as will be discussed below. The learning algorithm used seems a straightforward decision to make, when the rate of convergence is important and, based on existing literature, the options commonly tested are the Q-Learning, WoLF-PHC, Fictitious Play (FP) algorithms and Highest Cumulative Reward-based models. The Q-Learning algorithm consistently outperforms WoLF-PHC, Fictitious Play (FP) and Highest Cumulative Reward-based models in all of the literature presented, thus making it the clear choice for the learning algorithm, when agents are to be modelled with the ability to learn from past interactions (Table 3). A future research topic can involve studying whether other machine learning techniques could perform better. In order to decide on the appropriate observation capabilities of the agent, the designer will need to consider the learning algorithm to be used and the cognitive abilities of the agent. The observation capabilities of the agent become a significant consideration when the agents are to be modelled to rely on previous knowledge or knowledge about other agents.

4.2 Additional characteristics

There remain a few characteristics that have not (yet) been extensively investigated in the literature and should be explored in the future: the types of game played when modelling agent interaction in a game theoretic manner; and the type of exploration used by the agents to find new norms. Additionally, we believe some existing norm emergence mechanisms could be improved to allow for greater success across varying topologies.

The learning mechanism used in Hao et al. [52] was able to successfully observe the emergence of norms irrespective of the underlying topology. This is distinct from what occurs in other norm emergence mechanisms. Similarly, the joint action learners in Hao et al. [51] were also able to emerge to a norm over all the networks, though with varying rates of emergence. This we believe deserves more wide-ranging investigation (Table 4). We note

that the learning mechanisms used in both Hao et al. [52] and Hao et al. [51] involved the agents having access to information about their own interactions and the interactions of their neighbours. In Hao et al. [51], joint action learner agents have access to information about their neighbour's neighbours as well. The insight here is that, if agents learn from not just their interactions but the interactions of other agents, then they can converge faster and more successfully to a norm. This strategy is similar to that of collective learning agents, where interacting with more agents to have more data to learn from is used. However this strategy reduces the computational requirement of having to actually interact with the agents instead simply being able to observe the interactions of those other agents.

The majority of research that studies norm emergence from a game theoretic point of view demonstrates a given mechanism's ability to simulate the emergence of norms for a single type of game. The only exceptions are Hao et al. [52] and Hao et al. [51]. The mechanism used in Hao et al. [52] is able to successfully demonstrate emergence in all the game types investigated: coordination, anti-coordination, coordination game with high penalty and fully stochastic coordination game with high penalty. Both the joint action learner (JAL) and individual action learner (IAL) action selection mechanisms in Hao et al. [51] are investigated over two deterministic cooperative games and two stochastic cooperative games, for both single state and general cooperative games. JALs are successful across all topologies in all single state deterministic and stochastic cooperative games as well as the general deterministic cooperative games. They however fail for some of the general stochastic cooperative games. IALs are successful in all single state deterministic and stochastic cooperative games for all topologies, except the random topology, however with slower emergence rates than the JALs. In general deterministic cooperative games, IALs are successful on all networks except random networks but consistently fail for a certain percentage of games across all four networks tested in general stochastic cooperative games. The results above motivate us to suggest that other successful mechanisms for norm emergence, that have been applied to one type of game, should be evaluated across varying types of games to determine their overall success (see Table 4). We believe that a mechanism to support norm emergence must be applicable across varying types of games or must be capable of adapting to varying types of games and other characteristics of a system, an opinion we further discuss in Sect. 4.5.

Mechanisms for norm emergence that utilise a learning algorithm have predominantly done so using an ϵ -greedy exploration strategy. In our opinion, these works just state that the agent uses this ϵ -greedy exploration strategy, without providing any justification. This is likely because ϵ -greedy exploration is one of the most widely used exploration strategies, but not the only one, for example, another commonly used strategy is Boltzmann exploration. Exploration is important in reinforcement learning as some norms cannot emerge in the absence of exploration. Consequently, we believe it is prudent for future research on learning agents in norm emergence to investigate the effect of various exploration strategies on the rate of norm emergence (Table 4), especially given the importance of exploration strategies in reinforcement learning to find the appropriate balance with exploration and exploitation.

We highlight one attempt at investigating exploration in Hao et al. [52], where they investigate two types of collective learning strategies: local (EV-l) and global (EV-g). In EV-l, the exploration is made before determining the overall best-response action, whereas in EV-g, the exploration is done after the determination of the best-response action. Both strategies utilise the ϵ -greedy exploration but at different times in the process. Results show that agents are able to converge to norms faster using local exploration [52].

4.3 Online norm synthesis

The literature on the prescriptive approach to norms, with specific reference to the synthesis of norms, has potential for further study on the employment of more online implementations. Norm synthesis is defined as the process of determining the set of norms that avoid conflicting states [70,96]. Usually in the literature, a new norm simply appears from nowhere. The study is then on how, during runtime, the agents go through the recognition, internalisation and adoption processes to determine whether the new norm will be adopted. There is however little or no mention of the source of the norm. A typical example is Serramia et al. [94] who successfully implement an approach for selecting the “right” subset of norms for the normative system from a large set of norms. The focus of the study is not on the origin of the set of norms used but instead on the method used in the selection process. They employ linear programming to select norms that cannot be classified as generalisable, pairwise-compatible, or substitutable while at the same time promoting moral values as specified by the designers [94].

There are a few exceptions however, for example Morales et al. [70–72], Alechina et al. [4], Mashayekhi et al. [68] and Campos et al. [26] that investigate mechanisms for the online synthesis of norms. This in contrast to Li [62], who advances a synthesis mechanism for offline use aimed at designers, although it could in principle be delivered as an online service.

Mashayekhi et al. [68] present Silk, which consists of a centralised mechanism, an agent, referred to as a generator that monitors member agents interactions and recommends norms to resolve conflicts. Silk also includes individual agents that use reinforcement learning to guide their decision making. When agents leave, they share their knowledge with incoming agents, who decide whether to adopt or violate a norm based on this new knowledge.

We propose that there is a need for additional mechanisms for the online synthesis of norms using a centralised approach, and that mechanisms using a distributed approach should be explored, under the constraint of potentially limited or just local knowledge (Table 4). In their offline method for synthesising an evolutionary stable normative system, Morales et al. [73] present SENSE—“System for Evolutionary Norm SynthEsis”—which is similar to the centralised method in Morales et al. [70,71]. They suggest that a distributed method of their approach is possible [73].

Similarly Campos et al. [26] present AOCMAS which enable an organisation with adaptive capabilities. At runtime, assistant agents propose the regulations, rules or legal norms, for the system after partially observing the state of the organisation. They utilise case-based reasoning (CBR) to check for regulations from existing stored solutions or propose new solutions. Each proposed set of regulations is voted on by all assistants, and the majority decision becomes the new set of regulations. The system includes a feedback loop, where they evaluate the effectiveness of the proposed regulations and update the stored solution or remove it.

The approach by Alechina et al. [4] presents Guards, which are functions that assess a fixed history of a run and identify which states can be termed safe, meaning not leading to a violation. The remaining states are restricted from being accomplished. One benefit of this approach is that there is no need to know all the potential unsafe state transitions at the beginning of a run, or at design time, as this can be determined at intervals during runtime.

In principle agents from Morales et al. [70–73] can violate the norms synthesised, but rarely do so, since it is the best strategy to adopt the norms, whereas agents in Alechina et al. [4] do not have the option to do so, as the result of the synthesis process is a regimented system where identified states are no longer accessible. Agents in Mashayekhi et al. [68]

are free to violate laws but they normally respect them. Agents also utilise the information passed on to them. On the other hand Campos et al. [26] boasts regulations without monitoring or enforcement mechanisms for norm compliance, instead in the presence of several non-compliant agents, regulations become more restrictive prompting changed behaviour to encourage less restrictive regulations.

4.4 High-cognitive ability agents demonstrating norm emergence

To our knowledge, no research has been done to determine if agents with high cognitive abilities can demonstrate the emergence of norms using their action selection strategies. This is a gap that we have established in the research, as shown on Table 2, where in the column “agents’ cognitive abilities”, none of the papers surveyed have been classified as high. High cognitive agents, often practical reasoning agents, are assumed to be able to reason about the environment, norms and the actions available to them, and past performance, in order to choose an appropriate action. Usually these agents go through this action selection process for every decision, which means when posed with similar situations, they can potentially choose a different action in each instance. Currently in the emergence literature, agents use somewhat simplistic methods for deciding if a particular action from a set of possible ones will be collectively utilised. We believe the real test of the emergence of a norm would be demonstrated if high cognitive agents using non-simplistic reasoning and without knowledge of other agents’ actions will still display this sort of preferred action selection in similar situations. At this point there would be no confusion as to whether a particular norm has emerged.

In order to achieve this, a study will need to consider actions where each results in equal payoff. The use of actions with equal payoff is important as rational agents might choose the action with the highest payoff based on the fact that it is the highest payoff only. Therefore if several actions yield equal payoff, then we are able to determine if different agents with different reasoning will still adopt the same action. Existing research show that when all actions have equal payoff agents take long to converge to a norm [93] or never do [57]. Consequently, it is worthwhile to investigate if the emergence of a norm, where one action is the preferred action from a set of actions, will be observed and sustained over time, in a society consisting of these types of agents, and with relevant actions resulting in equal payoff (Table 4). We nonetheless recognise that, in order to influence high-cognitive ability agents to choose the same action from a set of actions with equal payoff, there must exist some alternative criterion or mechanism (e.g. effect on values held) that differentiates the actions for these agents.

We now present some literature utilising agents we believe demonstrates high-cognitive abilities and briefly explain their non-simplistic action selection process. The following only represent a small sample of the extensive literature that exists in this area of research. Visser et al. [104] show how BDI agents can select plans based on preferences. They describe preferences as soft constraints, meaning they do not have to be satisfied for the goal to be satisfied. Plans are annotated with the preferences of the users. In plan selection, each plan is considered based on how well it meets user preferences, and the most preferred, based on the computed value, is attempted first [104].

Bench-Capon and Modgil [18] consider explicit ethical agents, that have the capacity to reason about the actions they will perform, and the norms governing those actions. They propose a value-based reasoning approach, which utilises argumentation, allowing autonomous agents to operate in unforeseen situations. The sort of reasoning proposed is intended to be

closer to the reasoning in humans than previous algorithmic methods utilised in intelligent agents. Agents operating at this cognitive level are capable of deciding when to violate or adhere to a norm in a given context, based on the ordering of their values. This recent study by Bench-Capon and Modgil [18] builds on concepts formulated in Atkinson and Bench-Capon [10] and Atkinson and Bench-Capon [11].

A similar use of value-based reasoning for the purpose of plan selection is demonstrated in Cranfield et al. [34], where plans are filtered, not just for their applicability to the situation, but also based on the effect it will have on values held by the agent. Additionally, Petruzzi et al. [81] utilise social capital as an incentive for agents to participate in and choose actions that benefit a group rather than themselves. Social capital represents attributes of individuals for example, trustworthiness, social network and institutions, that help them to choose appropriate actions in collective situations [81]. They present a social capital framework that supplies the data store for the attributes and also defines the processes that assist agents in deciding how to act in appropriate situations—updating, evaluating and decision-making with social capital.

The use of high-cognitive ability agents will undoubtedly necessitate the exploration of ethical issues in multiagent systems, an area that is recently being actively researched.

4.5 Generic or adaptive framework to support norm emergence

It would be useful if a generic or adaptive framework existed that can facilitate norm emergence across various types of agent societies. The identification of the characteristics of a normative system that supports norm emergence can be used as a starting point in the development of a framework that combines the most influential characteristics into a single framework. The assumption is that such a framework could potentially facilitate norm emergence in various types of normative multiagent societies. Recent research has been successful in promoting norm emergence, but we believe their main drawback is that they have singled out a characteristic or in a few cases, two characteristics and have utilised those to support norm emergence. There is potential to investigate if a generic mechanism that combines the characteristics will be capable of facilitating the emergence of norms in normative agent societies, irrespective of the properties of the society. It is possible that the resulting framework may be an adaptive one that utilises a fixed combination of characteristics, for a given type of agent society, or a single optimal one that can effectively promote norm emergence in any situation.

The work by Mahmoud et al. [67] and Hao et al. [52] demonstrate strategies that are successful across varying characteristics of societies. Mahmoud et al. [67] presents a mechanism of metanorms over different topologies. They found that utilising a new dynamic policy adaptation approach results in the establishment of norms in scale-free networks. The basic approach was unable to facilitate the establishment of norms, even though it was successful in lattices and small world networks [67]. Hao et al. [52] investigate two learning strategies for norm emergence, collective learning EV-l—local exploration—and collective learning EV-g—global exploration—across various types of coordination and anti-coordination games and across different topologies. Results demonstrate that when compared with pairwise learning and social learning, the strategies outperform them across a wider variety of games and topologies [52]. We believe the results of the preceding studies show that there is viability in designing a generic or an adaptive strategy for all types of agent societies (Table 4).

4.6 Stable emergence

Norm emergence in the literature cited studies the emergence of norms as isolated instances without much consideration for the stability of the norms that emerged within the society. It is likely that norms that have emerged can become oscillatory if observed over time when the dynamics of the environment changes and agents continue to modify their strategies selfishly. Though the literature of emergence is normally only concerned with reaching the prescribed threshold, we believe the duration of the norm that emerged within the society should also be investigated. A norm that emerges but persists only for a short time to be replaced by a different norm would pose problems for the stability of a system (see Table 4). It would be futile to design systems on the premise of norm emergence only if norms that emerge can potentially become oscillatory.

Savarimuthu et al. [91] observe cycles of emergence, where the emergence of a particular norm is not sustained, because agents are lying. Their research demonstrates that once agents lie, a norm that emerges can be replaced by a different norm over time and this cycle can continue indefinitely. The literature on norm synthesis is more concerned with the stability of norms. In Morales et al. [73] they present the development of an evolutionary stable strategy (ESS), which is one that is followed by all agents on the assumption that all the other agents are following same, and where the agent cannot do better or benefit from deviating from this strategy. SENSE [73] keeps exploring new norms whenever a conflict situation arises. Over time, a set of norms remains stable but even then a disturbance can result in a norm from this set being removed and a different one added. If eventually the set of norms returns to the initial stable set then it is classified as an ESS. Therefore, they posit that a set of strategies is an ESS, if the norms contained within the normative system changes because of a disturbance, but is restored to the original set after some time [73].

Though this speaks to identifying the appropriate set of norms that should be contained in the normative system, there is potential for a study on the stability of the norms that emerge in a society by investigating the effects of the presence of liars or norm entrepreneurs who are injected into the simulation to try to change the prevailing norm(s). Additionally other situations that can potentially affect the stability of a norm that has emerged should also be investigated. An example could be norms that are good for individuals, but are ultimately bad for the collective (e.g. common pool resources).

4.7 Prescribing both social and legal norms

Conte and Castelfranchi [32] speak about the unification of the perspectives of norms: prescriptive and emergence. Their work presents a way that both types of norms—social and legal—can be prescribed using the same formalism in one system [32]. There has not been other research, to our knowledge, that investigates providing a single formalism to represent both social and legal norms. The research by Frantz et al. [43] is the closest to this concept, where they present dynamic deontics utilising Interval Type-2 Fuzzy Sets. They suggest that the use of dynamic deontics allows for a wider spectrum of understanding of norms where it is not limited to must—obligations—may—permissions—or may not—prohibitions—but to a larger set consisting of additional terms such as must not, should not, may not, may, should and must [43].

It is common for a human society to be governed by both social and legal norms and one could expect that agent societies might benefit from the same situation. The legal norms would represent the non-negotiable behaviour that cannot be left to chance for emergence, or

for which emergence is unable to properly regulate. The social norms would represent those behaviours that can emerge but do not require a legal standing, though it can be good for agents to follow them to keep the society in a mutually acceptable state. Moreover, Morales et al. [73] posit that there are situations, where numerous interdependent conflict situations exist, where norm emergence cannot properly synthesise norms for multiagent systems. Consequently, a combination of both social and legal norms in one society is potentially good for the society. Societies that utilise legal norms will likely have same, explicitly represented, based on findings in the literature, see Sect. 2.3 and Table 1. It would be useful for these societies to be able to represent social norms explicitly as well. The explicit representation of social norms will allow agents to be able to reason about adopting them in the same way as they do legal norms. This would eliminate the questions surrounding whether the agents are actually adopting norms or just acting selfishly. We believe explicitly representing both social and legal norms would help us to properly understand if norms do emerge based on the interaction of agents that are aware of the existence of these norms (Table 3).

In a study of extortion rackets in the Sicilian and Georgian Mafia, Székely et al. [98] discuss how important it is for social norms to be established to strengthen the legal norms that exist. They discuss how the difficulties in enforcement and monitoring can affect the effectiveness of laws and also how a social norm can support legal norms when they are in harmony, but conversely how social norms can undermine legal norms when they are in conflict [98]. When working in harmony, a social norm may be more likely to be reported when violated by members, which would increase enforcement of the law, or the social punishment may be enough to deter violators without the need for legal punishment [98]. They observe that when societies with similar circumstances sought to resolve the case of extortion rackets using legal norms, it was only in the case of Georgia that success was immediate, because it had accompanying social norms that helped to strengthen or reinforce the legal norms [98]. The benefits of the social and legal norms within a human society working in harmony for the overall benefit of the society is an opportunity for investigation in agent societies.

4.8 Converting social norms to legal norms

Brooks et al. [25] states that “norms typically become codified into law, but do not start out that way”. The view that norms can become laws is shared by others in the literature [87,90]. The concept of converting social norms to legal norms is in line with the norm life cycle as expressed by Andrighetto et al. [7] where in the norm evolution stage some norms may become codified into law, others may evolve and the remaining decay. A common example is the perspective of not smoking in public places. This began as an acceptable behaviour, a social norm, where the violation might have only been met by disgruntled onlookers, but over time has turned into a legal norm in the society, that is monitored and enforced centrally.

Normally, norm emergence is characterised as a percentage of the observed agents adopting the same action in a given situation. An external observer infers that a particular norm has emerged and the research ends. There is an opportunity for research where the preference action that emerges within the society can then be prescribed as a law in the society (see Table 4). Therefore, the action that has casually emerged within the society becomes one of the laws that govern the society, thereby turning the norm into a law that must be followed. This can be potentially implemented as an obligation for that action in a given situation. Ghorbani et al. [47,48] investigate a similar concept, where agents propose their strategy to be accepted as the norm to be followed by all the agents. The set of norms that are prescribed to govern a society is referred to as the normative system, sometimes referred

to as an institution. Ghorbani et al. [47,48] present an institution which consist of the most popular proposal or common norm from the agents. The nature of this can be considered to be the norm followed by the majority of the agents in the society, meaning the norm that emerged becomes the norm to govern the behaviour of the agents in the society.

Though the concept of turning social norms into legal norms has been mentioned in the literature [25,87,90], the only research into implementing the concept appears to be Ghorbani et al. [47,48]. The opportunity here cannot be missed as the implications of turning a social norm into a legal norm are significant. It entrenches a convention and associates civil and possibly criminal penalties to what had previously only been convention-breaking, as well as making something statute and immutable (without revising legislation) that was previously fluid.

A survey by Haynes et al. [54] examines the viability of this concept. They present the concept of the engineering of emergent norms which involves three main steps: (i) identification or detection of the possible emergent norm, (ii) evaluation of the benefit of the possible norm to individual agents or the system as a whole, and (iii) encouraging or discouraging the spread of the norm [54]. Haynes et al. [54] posit that some emergent behaviours in a system can be beneficial to the system and should be encouraged and spread, while some emergent behaviour should be discouraged once determined not to be beneficial. We propose that the encouragement of beneficial emergent behaviours will give rise to the emergence of obligation norms, potentially permission norms and the introduction of prohibition norms for substitute behaviour. Likewise, the discouragement of non-beneficial emergent behaviours could give rise to the emergence of prohibition norms, potentially obligation norms to avoid certain states and revocation of permission norms.

5 Conclusions

In this viewpoint paper we present an in-depth analysis of the literature on norm emergence, in order to identify the concepts that are applicable to the study of normative multiagent systems, with an emphasis on prescriptive norms and those concepts in normative multiagent systems that impact norm emergence. We commence with a discussion of norms and the reasons why agents may adopt or violate them. An examination of the literature on norm emergence also reveals that there is a specific perspective of norms, namely the emergence approach, that is different from that usually encountered in normative multiagent systems. Consequently, we present these different perspectives briefly, the prescriptive approach and emergence approach, and show how these approaches affect how norms are represented in agent societies (see Table 1). The prescriptive approach is usually represented using deontic notions, while the emergence approach appears as a stored or computed strategy for action selection.

Further, we propose a new definition of the norm emergence process which better represents the broader activities that lead to the point when emergence is observed. Norm emergence should be defined as the process whereby a population reaches a predefined threshold of agents following the same norm. It includes the creation and spreading of the norm, and culminates with the observation of a percentage of agents following the norm. Based on the above notion of the process of norm emergence, we subsequently present the norm life cycle, building on the earlier survey by Savarimuthu and Cranefield [86]. We find that the stages of the norm life cycle can be studied using either the emergence approach or the

prescriptive approach and we present the similarities and differences in the implementation of the norm life cycle that arise in consequence.

In our discussion of the implementation of the norm life cycle, we show how agents from the emergence perspective indirectly and unintentionally utilise norms coinciding with explicit norms in the prescriptive approach. This prompts us to suggest that emergent behaviours or conventions, as they are referred to in the prescriptive perspective, can be prescribed as laws to the benefit of multiagent systems.

We examine simulation models of norm emergence to identify any characteristics that have an effect on the emergence of a norm. We determine the following characteristics: (i) social topology, (ii) agents' cognitive abilities or decision making mechanisms, (iii) propagation mechanism, (iv) agents' use of learning algorithms, (v) observation capabilities, (vi) offline methods or online methods. These are identified and discussed, while Table 2 demonstrates which of the characteristics of norm emergence are investigated in the various simulation models cited. Table 3 summarises our findings.

We conclude the survey with a discussion of challenges and/or opportunities that exist in the study of norms with our findings summarised in Table 4. Based on our survey, we believe the following areas offer the most research potential: (i) design of agent systems, (ii) additional characteristics, (iii) online norm synthesis, (iv) high cognitive agents demonstrating norm emergence, (v) generic or adaptive framework to support norm emergence, (vi) stable emergence, (vii) prescribing both social and legal norms, and (viii) converting social norms to legal norms.

In our discussions in Sect. 4, we advance the notion for the encoding of implicit norms/behaviours as explicit representations for normative multiagent systems, thereby allowing individual agents to participate in the synthesis of norms contributing to self-governance. The literature on norm emergence cited provides several insights applicable to the study and improvement of normative multiagent systems in the future. Similarly, future research in norm emergence can attempt to fill the gaps identified and incorporate some of the concepts from normative multiagent systems as discussed throughout this paper.

Acknowledgements We wish to thank the editors and reviewers whose valuable comments and additional insights resulted in a much improved paper.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

References

1. Ajmeri, N., Jiang, J., Chirkova, R., Doyle, J., & Singh, M. P. (2016). Coco: Runtime reasoning about conflicting commitments. In S. Kambhampati (Ed.), *Proceedings of the twenty-fifth international joint conference on artificial intelligence* (pp. 17–23), IJCAI 2016, New York, NY, USA, 9–15 July 2016. IJCAI/AAAI Press. <http://www.ijcai.org/Abstract/16/010>.
2. Aldewereld, H., & Dignum, V. (2010). OperettA: Organization-oriented development environment. In M. Dastani, A. E. Fallah-Seghrouchni, J. Hübner, & J. Leite (Eds.), *Languages, methodologies, and development tools for multi-agent systems—third international workshop* (Vol. 6822, pp. 1–18), LADS 2010, Lyon, France, August 30–September 1, 2010, Revised Selected Papers, Springer, Lecture Notes in Computer Science. https://doi.org/10.1007/978-3-642-22723-3_1.
3. Aldewereld, H., Boissier, O., Dignum, V., Noriega, P., & Padget, J. (2016). *Social coordination frameworks for social technical systems* (1st ed.). Berlin: Springer.

4. Alechina, N., Bulling, N., Dastani, M., & Logan, B. (2015). Practical run-time norm enforcement with bounded lookahead. In Weiss, G., Yolum, P., Bordini, R. H., & Elkind, E. (Eds.), *Proceedings of the 2015 international conference on Autonomous Agents and Multiagent Systems, AAMAS 2015, Istanbul, Turkey, May 4–8*, (pp. 433–441), ACM. <http://dl.acm.org/citation.cfm?id=2772937>.
5. Andrighetto, G., & Conte, R. (2012). Cognitive dynamics of norm compliance: From norm adoption to flexible automated conformity. *Artificial Intelligence and Law*, 20(4), 359–381. <https://doi.org/10.1007/s10506-012-9135-6>.
6. Andrighetto, G., Campenni M, Conte, R., & Paolucci, M. (2007). On the immergence of norms: A normative agent architecture. In G. P. Trajkovski, & S. G. Collins (Eds.), *Emergent agents and socialities: Social and organizational aspects of intelligence, Papers from the 2007 AAAI fall symposium* (Vol. FS-07-04, pp. 11–18), Arlington, Virginia, USA, November 9–11, 2007. AAAI Press. AAAI Technical Report. <https://www.aaai.org/Library/Symposia/Fall/2007/fs07-04-003.php>.
7. Andrighetto, G., Castelfranchi, C., Mayor, E., McBreen, J., Lopez-Sanchez, M., & Parsons, S. (2013a). (Social) Norm dynamics. *Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik GmbH, Wadern/Saarbruecken, Germany*, <https://doi.org/10.4230/dfu.vol4.12111.135>, <http://drops.dagstuhl.de/opus/volltexte/2013/4002/>.
8. Andrighetto, G., Governatori, G., Noriega, P., & van der Torre, L. (2013b). *Normative multi-agent systems...*. <https://doi.org/10.4230/DFU.Vol4.12111.i>, <http://drops.dagstuhl.de/opus/volltexte/2013/4005/>.
9. Athakravi, D., Corapi, D., Russo, A., Vos, M. D., Padget, J. A., & Satoh, K. (2012). Handling change in normative specifications. In M. Baldoni, L. A. Dennis, V. Mascardi, & W. V. Vasconcelos (Eds.), *Declarative agent languages and technologies X—10th international workshop* (Vol. 7784, pp. 1–19), DALIT 2012, Valencia, Spain, June 4, 2012. Revised Selected Papers, Springer, Lecture Notes in Computer Science. https://doi.org/10.1007/978-3-642-37890-4_1.
10. Atkinson, K., & Bench-Capon, T. (2007). Practical reasoning as presumptive argumentation using action based alternating transition systems. *Artificial Intelligence*, 171(10–15), 855–874. <https://doi.org/10.1016/j.artint.2007.04.009>.
11. Atkinson, K., & Bench-Capon, T. (2016). States, goals and values: Revisiting practical reasoning. *Argument & Computation*, 7(2–3), 135–154. <https://doi.org/10.3233/AAC-160011>.
12. Axelrod, R. (1986). An evolutionary approach to norms. *American Political Science Review*, 80(04), 1095–1111. <https://doi.org/10.2307/1960858>.
13. Balaraman, V., & Singh, M. (2014). Exploring norm establishment in organizations using an extended Axelrod model and with two new metanorms. In *Proceedings of the 2014 summer simulation multicongference, society for computer simulation international, society for computer simulation international, SummerSim '14* (pp. 1–9).
14. Balke, T., De Vos, M., & Padget, J. (2013). Evaluating the cost of enforcement by agent-based simulation: A wireless mobile grid example. In G. Boella, E. Elkind, B. T. R. Savarimuthu, F. Dignum, & M. K. Purvis (Eds.), *PRIMA 2013: Principles and practice of multi-agent systems* (Vol. 8291, pp. 21–36), Springer Berlin Heidelberg, Berlin, Heidelberg. https://doi.org/10.1007/978-3-642-44927-7_3.
15. Beheshti, R. (2014). Normative agents for real-world scenarios. In *Proceedings of the 2014 international conference on Autonomous agents and multi-agent systems, international foundation for autonomous agents and multiagent systems* (pp. 1749–1750).
16. Beheshti, R., & Sukthankar, G. (2015). Modeling tipping point theory using normative multi-agent systems. In *Proceedings of the 2015 international conference on autonomous agents and multiagent systems, international foundation for autonomous agents and multiagent systems* (pp. 1731–1732).
17. Beheshti, R., Ali, A. M., & Sukthankar, G. R. (2015). Cognitive social learners: An architecture for modeling normative behavior. In *AAAI* (pp. 2017–2023).
18. Bench-Capon, T., & Modgil, S. (2017). Norms and value based reasoning: Justifying compliance and violation. *Artificial Intelligence and Law*, 25(1), 29–64. <https://doi.org/10.1007/s10506-017-9194-9>.
19. Boella, G., & van der Torre, L. (2003a). BDI and BOD argumentation. In *Proceedings of IJCAI workshop on computational models of natural argument* (Vol. 3).
20. Boella, G., & van der Torre, L. W. N. (2003b). Norm governed multiagent systems: The delegation of control to autonomous agents. In *2003 IEEE/WIC international conference on intelligent agent technology (IAT 2003)* (pp. 329–335), 13–17 October 2003. Halifax, Canada, IEEE Computer Society. <https://doi.org/10.1109/IAT.2003.1241092>.
21. Boella, G., van der Torre, L., & Verhagen, H. (2006). Introduction to normative multiagent systems. *Computational & Mathematical Organization Theory*, 12(2–3), 71–79. <https://doi.org/10.1007/s10588-006-9537-7>.
22. Boissier, O., Hübner, J. F., & Ricci, A. (2016). The JaCaMo framework. In H. Aldewereld, O. Boissier, V. Dignum, P. Noriega, & J. Padget (Eds.), *Social coordination frameworks for social technical systems* (pp. 125–151), Springer International Publishing, Cham. https://doi.org/10.1007/978-3-319-33570-4_7.

23. Boyd, R., & Richerson, P. (1988). Culture and the evolutionary process. Biology, Anthropology, Sociology, University of Chicago Press, <https://books.google.co.uk/books?id=MBg4oBsCKU8C>.
24. Broersen, J., Dastani, M., Hulstijn, J., Huang, Z., & van der Torre, L. (2001). The boid architecture—conflicts between beliefs, obligations, intentions and desires. In *In proceedings of the fifth international conference on autonomous agents* (pp. 9–16), ACM Press .
25. Brooks, L., Iba, W., & Sen, S. (2011). Modeling the emergence and convergence of norms. In T. Walsh (Ed.), *IJCAI 2011, proceedings of the 22nd international joint conference on artificial intelligence* (pp. 97–102), Barcelona, Catalonia, Spain, July 16–22, 2011, IJCAI/AAAI.
26. Campos, J., Lopez-Sanchez, M., Salamó, M., Avila, P., & Rodríguez-Aguilar, J. A. (2013). Robust regulation adaptation in multi-agent systems. *ACM Transactions on Autonomous and Adaptive Systems*, 8(3), 1–27. <https://doi.org/10.1145/2517328>.
27. Castelfranchi, C. (1999). Prescribed mental attitudes in goal-adoption and norm-adoption. *Artificial Intelligence and Law*, 7(1), 37–50. <https://doi.org/10.1023/A:1008363413485>.
28. Chakrabarti, P., & Basu, J. (2010). Emergence of norms in a society of heterogeneous agents influenced by the rules of cellular automata techniques. *International Journal on Computer Science and Engineering (IJCSSE)*, 2(3), 481–486.
29. Chao, Yu., Zhang, Minjie, & Ren, Fenghui. (2014). Collective learning for the emergence of social norms in networked multiagent systems. *IEEE Transactions on Cybernetics*, 44(12), 2342–2355. <https://doi.org/10.1109/TCYB.2014.2306919>.
30. Claus, C., & Boutilier, C. (1998). The dynamics of reinforcement learning in cooperative multiagent systems. In J. Mostow, & C. Rich (Eds.), *Proceedings of the fifteenth national conference on artificial intelligence and tenth innovative applications of artificial intelligence conference* (pp. 746–752), AAAI 98, IAAI 98, July 26–30, 1998, Madison, Wisconsin, USA. AAAI Press / The MIT Press. <http://www.aaai.org/Library/AAAI/1998/aaai98-106.php>.
31. Cliffe, O., Vos, M. D., & Padget, J. A. (2007). Embedding landmarks and scenes in a computational model of institutions. In J. S. Sichman, J. A. Padget, S. Ossowski, & P. Noriega (Eds.), *Coordination, organizations, institutions, and norms in agent systems III, COIN 2007, international workshops COIN@AAMAS 2007* (Vol. 4870, pp. 41–57), Honolulu, HI, USA, May 14, 2007, COIN@MALLOW 2007, Durham, UK, September 3–4, 2007. Revised Selected Papers, Springer, Lecture Notes in Computer Science. https://doi.org/10.1007/978-3-540-79003-7_4.
32. Conte, R., & Castelfranchi, C. (1999). From conventions to prescriptions: Towards an integrated view of norms. *Artificial Intelligence and Law*, 7(4), 323–340.
33. Corapi, D., Russo, A., De Vos, M., Padget, J., & Satoh, K. (2011). Normative design using inductive learning. *Theory and Practice of Logic Programming*, 11(4–5), 783–799. <https://doi.org/10.1017/s1471068411000305>.
34. Craneffeld, S., Winikoff, M., Dignum, V., & Dignum, F. (2017). No Pizza for You: Value-based plan selection in BDI agents. In *Proceedings of the twenty-sixth international joint conference on artificial intelligence* (pp. 178–184), IJCAI-17. <https://doi.org/10.24963/ijcai.2017/26>.
35. Criado, N., Argente, E., & Botti, V. J. (2010a). Normative deliberation in graded BDI agents. In J. Dix, & C. Witteveen (Eds.), *Multiagent system technologies, 8th German conference* (Vol. 6251, pp. 52–63), MATES 2010, Leipzig, Germany, September 27–29, 2010. Proceedings, Springer, Lecture Notes in Computer Science. https://doi.org/10.1007/978-3-642-16178-0_7.
36. Criado, N., Argente, E., & Botti, V. J. (2010b). Rational strategies for norm compliance in the n-bdi proposal. In M. D. Vos, N. Fornara, J. V. Pitt, & G. A. Vouros (Eds.), *Coordination, organizations, institutions, and norms in agent systems VI—COIN 2010 international workshops* (Vol. 6541, pp. 1–20), COIN@AAMAS 2010, Toronto, Canada, May 2010, COIN@MALLOW 2010, Lyon, France, August 2010. Revised Selected Papers, Springer, Lecture Notes in Computer Science. https://doi.org/10.1007/978-3-642-21268-0_1.
37. Dascalu, S., Scurtu, T., Urzica, A., Trascau, M., & Florea, A. M. (2013). Using norm emergence in addressing the tragedy of the commons. In *International conference on computational collective intelligence* (pp. 165–174), Springer. http://link.springer.com/chapter/10.1007/978-3-642-40495-5_17.
38. Dastani, M., & van der Torre, L. (2002). A classification of cognitive agents. In *Proceedings of the cognitive science society* (Vol. 24).
39. Dastani, M., Meyer, J. C., & Grossi, D. (2013). A logic for normative multi-agent programs. *Journal of Logic and Computation*, 23(2), 335–354. <https://doi.org/10.1093/logcom/exr026>.
40. Dignum, V. (2009). Handbook of research on multi-agent systems: Semantics and dynamics of organizational models: Semantics and dynamics of organizational models. information science reference, <https://books.google.co.uk/books?id=-UMAEGsMvfYC>.

41. dos Santos Neto, B. F., da Silva, V. T., & de Lucena, C. J. (2011). Developing goal-oriented normative agents: The NBDI architecture. In *International conference on agents and artificial intelligence* (pp. 176–191), Springer.
42. Franks, H., Griffiths, N., & Anand, S. S. (2013). Learning influence in complex social networks. In *Proceedings of the 2013 international conference on autonomous agents and multi-agent systems, international foundation for autonomous agents and multiagent systems* (pp. 447–454), Richland, SC, AAMAS '13. <http://dl.acm.org/citation.cfm?id=2484920.2484992>.
43. Frantz, C., Purvis, M. K., Nowostawski, M., & Savarimuthu, B. T. R. (2014). Analysing the dynamics of norm evolution using interval type-2 fuzzy sets, (pp. 230–237), IEEE <https://doi.org/10.1109/WI-IAT.2014.172>, <http://ieeexplore.ieee.org/document/6928190/>.
44. Gabbay, D., Horty, J., Parent, X., van der Torre, L., & van der Meyden, R. (2013). Handbook of deontic logic and normative systems. College Publications, <https://books.google.co.uk/books?id=iPchnwEACAAJ>.
45. Gasparini, L., Norman, T. J., & Kollingbaum, M. J. (2018). Severity-sensitive norm-governed multi-agent planning. *Autonomous Agents and Multi-Agent Systems*, 32(1), 26–58. <https://doi.org/10.1007/s10458-017-9372-x>.
46. Ghorbani, A., & Bravo, G. (2016). Managing the commons: A simple model of the emergence of institutions through collective action. *International Journal of the Commons*, 10(1), 200–219. <https://doi.org/10.18352/ijc.606>.
47. Ghorbani, A., Bots, P., Dignum, V., & Dijkema, G. (2013). MAIA: A framework for developing agent-based social simulations. *Journal of Artificial Societies and Social Simulation*, 16(2), 9. <https://doi.org/10.18564/jass.2166>.
48. Ghorbani, A., Bravo, G., Frey, U., & Theesfeld, I. (2017). Self-organization in the commons: An empirically-tested model. *Environmental Modelling & Software*, 96, 30–45. <https://doi.org/10.1016/j.envsoft.2017.06.039>.
49. Hao, J., & Leung, H. F. (2013). Achieving socially optimal outcomes in multiagent systems with reinforcement social learning. *ACM Transactions on Autonomous and Adaptive Systems*, 8(3), 1–23. <https://doi.org/10.1145/2517329>.
50. Hao, J., Leung, H. F., & Ming, Z. (2014). Multiagent reinforcement social learning toward coordination in cooperative multiagent systems. *ACM Transactions on Autonomous and Adaptive Systems*, 9(4), 1–20. <https://doi.org/10.1145/2644819>.
51. Hao, J., Huang, D., Cai, Y., & Hf, Leung. (2017a). The dynamics of reinforcement social learning in networked cooperative multiagent systems. *Engineering Applications of Artificial Intelligence*, 58, 111–122. <https://doi.org/10.1016/j.engappai.2016.11.008>.
52. Hao, J., Sun, J., Chen, G., Wang, Z., Yu, C., & Ming, Z. (2017b). Efficient and robust emergence of norms through heuristic collective learning. *ACM Transactions on Autonomous and Adaptive Systems*, 12(4), 23. <https://doi.org/10.1145/3127498>.
53. Hassani-Mahmooui, B., & Parris, B. W. (2014). Dynamics of effort allocation and evolution of trust: An agent-based model. *Computational and Mathematical Organization Theory*, 20(2), 133–154. <https://doi.org/10.1007/s10588-013-9164-z>.
54. Haynes, C., Luck, M., McBurney, P., Mahmoud, S., Ví tek, T., & Miles, S. (2017). Engineering the emergence of norms: A review. *The Knowledge Engineering Review*, 32, e18. <https://doi.org/10.1017/S0269888917000169>.
55. Hoffmann, M. J. (2003). *Entrepreneurs and norm dynamics: An agent-based model of the norm life cycle*. Tech. rep., Delaware, USA.
56. Hofmann, L. M., Chakraborty, N., & Sycara, K. (2011). The evolution of cooperation in self-interested agent societies: a critical study. In *The 10th international conference on autonomous agents and multi-agent systems-Volume 2, international foundation for autonomous agents and multiagent systems* (pp. 685–692).
57. Hu, S., & Leung, H. (2017). Achieving coordination in multi-agent systems by stable local conventions under community networks. In C. Sierra (Ed.), *Proceedings of the twenty-sixth international joint conference on artificial intelligence* (pp. 4731–4737), IJCAI 2017. Melbourne, Australia, August 19–25, 2017, ijcai.org. <https://doi.org/10.24963/ijcai.2017/659>.
58. Hübner, J. F., Boissier, O., & Bordini, R. H. (2010). A normative organisation programming language for organisation management infrastructures. In J. Padget, A. Artikis, W. Vasconcelos, K. Stathis, V. T. da Silva, E. Matson, & A. Polleres (Eds.), *Coordination, Organizations, Institutions and Norms in Agent Systems V* (pp. 114–129). Berlin: Springer.
59. Jianye, H., Sun, J., Huang, D., Cai, Y., & Yu, C. (2015). Heuristic collective learning for efficient and robust emergence of social norms. In *Proceedings of the 2015 international conference on autonomous*

- agents and multiagent systems, international foundation for autonomous agents and multiagent systems* (pp. 1647–1648), Richland, SC, AAMAS '15. <http://dl.acm.org/citation.cfm?id=2772879.2773366>.
60. Kittock, J. (1995). Emergent conventions and the structure of multi-agent systems. In *Lectures in complex systems: The proceedings of the 1993 complex systems summer school, Santa Fe Institute Studies in the Sciences of Complexity Lecture Volume VI* (pp. 507–521), Santa Fe Institute, Addison-Wesley.
 61. Lee, J., Padget, J., Logan, B., Dybalova, D., & Alechina, N. (2014). N-Jason: Run-time norm compliance in AgentSpeak(L). In F. Dalpiaz, J. Dix, & M. B. van Riemsdijk (Eds.), *Engineering multi-agent systems—second international workshop* (Vol. 8758, pp. 367–387), EMAS 2014, Paris, France, May 5–6, 2014. Revised Selected Papers, Springer, Lecture Notes in Computer Science. https://doi.org/10.1007/978-3-319-14484-9_19.
 62. Li, T. (2014). Normative conflict detection and resolution in cooperating institutions. PhD thesis, University of Bath, https://researchportal.bath.ac.uk/files/123040954/Tingting_thesis_final.pdf.
 63. Li, T., Balke, T., De Vos M., Padget, J. A., & Satoh, K. (2013a). A model-based approach to the automatic revision of secondary legislation. In E. Francesconi, & B. Verheij (Eds.), *International conference on artificial intelligence and law* (pp. 202–206), ACM. <http://doi.acm.org/10.1145/2514601.2514627>.
 64. Li, T., Balke, T., Vos, M. D., Padget, J., & Satoh, K. (2013b). Legal conflict detection in interacting legal systems. In K. D. Ashley (Ed.), *JURIX, IOS Press, frontiers in artificial intelligence and applications* (Vol. 259, pp. 107–116). <https://doi.org/10.3233/978-1-61499-359-9-107>.
 65. Lopez, F. (2015). Modeling emergence of norms in multi-agent systems by applying tipping points ideas. [arXiv:1508.04531](https://arxiv.org/abs/1508.04531) [cs].
 66. Lotzmann, U., Möhring, M., & Troitzsch, K. G. (2013). Simulating the emergence of norms in different scenarios. *Artificial Intelligence and Law*, 21(1), 109–138. <https://doi.org/10.1007/s10506-012-9130-y>.
 67. Mahmoud, S., Griffiths, N., Keppens, J., & Luck, M. (2017). Establishing norms with metanorms over interaction topologies. *Autonomous Agents and Multi-Agent Systems*, 31(6), 1344–1376. <https://doi.org/10.1007/s10458-017-9364-x>.
 68. Mashayekhi, M., Du, H., List, G. F., & Singh, M. P. (2016). Silk: A simulation study of regulating open normative multiagent systems. In *Proceedings of the twenty-fifth international joint conference on artificial intelligence* (pp. 373–379), AAAI Press, IJCAI'16. <http://dl.acm.org/citation.cfm?id=3060621.3060674>.
 69. Mintz-Woo, K. (2010). R. Axelrod—an evolutionary approach to norms. Retrieved 19 June 2017 from <http://www.lamsade.dauphine.fr/~airiau/Teaching/norms/Axelrod-Evolutionary.pdf>.
 70. Morales, J., López-Sánchez, M., Rodríguez-Aguilar, J. A., Wooldridge, M., & Vasconcelos, W. (2013). Automated synthesis of normative systems. In *Proceedings of the 2013 international conference on autonomous agents and multi-agent systems, international foundation for autonomous agents and multi-agent systems* (pp. 483–490).
 71. Morales, J., López-Sánchez, M., Rodríguez-Aguilar, J. A., Vasconcelos, W., & Wooldridge, M. (2015a). Online automated synthesis of compact normative systems. *ACM Transactions on Autonomous and Adaptive Systems*, 10(1), 1–33. <https://doi.org/10.1145/2720024>.
 72. Morales, J., López-Sánchez, M., Rodríguez-Aguilar, J. A., Wooldridge, M. J., & Vasconcelos, W. W. (2015b). Synthesising liberal normative systems. In Weiss, G., Yolum, P., Bordini, R. H., & Elkind, E. (Eds.), *Proceedings of the 2015 international conference on Autonomous Agents and Multiagent Systems, AAMAS 2015, Istanbul, Turkey, May 4–8, (pp. 433–441)*, ACM. <http://dl.acm.org/citation.cfm?id=2772936>.
 73. Morales, J., Wooldridge, M., Rodríguez-Aguilar, J. A., & López-Sánchez, M. (2018). Off-line synthesis of evolutionarily stable normative systems. *Autonomous Agents and Multi-Agent Systems*, <https://doi.org/10.1007/s10458-018-9390-3>.
 74. Mukherjee, P., Sen, S., & Airiau, S. (2007). Emergence of norms with biased interactions in heterogeneous agent societies. In *ACM international conferences on web intelligence and intelligent agent technology workshops* (pp. 512–515), 2007 IEEE/WIC/ IEEE. <https://doi.org/10.1109/WI-IATW.2007.115>, <http://ieeexplore.ieee.org/document/4427640/>.
 75. Mukherjee, P., Sen, S., & Airiau, S. (2008). Norm emergence under constrained interactions in diverse societies. In *Proceedings of the 7th international joint conference on Autonomous agents and multiagent systems-Volume 2, International Foundation for Autonomous Agents and Multiagent Systems* (pp. 779–786), AAMAS '08.
 76. Mungovan, D., Howley, E., & Duggan, J. (2011). The influence of random interactions and decision heuristics on norm evolution in social networks. *Computational and Mathematical Organization Theory*, 17(2), 152–178. <https://doi.org/10.1007/s10588-011-9085-7>.

77. North, D. (1990). Institutions, institutional change and economic performance. Institutions, Institutional Change, and Economic Performance, Cambridge University Press, <https://books.google.co.uk/books?id=oFnWbTqgNPYC>.
78. Ohtsuki, H., Hauert, C., Lieberman, E., & Nowak, M. A. (2006). A simple rule for the evolution of cooperation on graphs and social networks. *Nature*, *441*(7092), 502–505.
79. Ostrom, E. (2015). Governing the commons: The evolution of institutions for collective action. Canto Classics, Cambridge University Press, <https://books.google.co.uk/books?id=hHGgCgAAQBAJ>.
80. Peleteiro, A., Burguillo, J. C., Arcos, J. L., & Rodríguez-Aguilar, J. A. (2014). Fostering cooperation through dynamic coalition formation and partner switching. *ACM Transactions on Autonomous and Adaptive Systems*, *9*(1), 1–31. <https://doi.org/10.1145/2567928>.
81. Petrucci, P. E., Busquets, D., & Pitt, J. V. (2015). A generic social capital framework for optimising self-organised collective action. In *2015 IEEE 9th international conference on self-adaptive and self-organizing systems* (pp. 21–30), Cambridge, MA, USA, September 21–25, 2015. IEEE Computer Society. <https://doi.org/10.1109/SASO.2015.10>.
82. Prakken, H., & Sergot, M. J. (1996). Contrary-to-duty obligations. *Studia Logica*, *57*(1), 91–115. <https://doi.org/10.1007/BF00370671>.
83. Riveret, R., Rotolo, A., & Sartor, G. (2012). Probabilistic rule-based argumentation for norm-governed learning agents. *Artificial Intelligence and Law*, *20*(4), 383–420. <https://doi.org/10.1007/s10506-012-9134-7>.
84. Sadri, F., Stathis, K., & Toni, F. (2006). Normative KGP agents. *Computational & Mathematical Organization Theory*, *12*(2–3), 101–126. <https://doi.org/10.1007/s10588-006-9539-5>.
85. Santos, J. S., Zahn, J. O., Silvestre, E. A., Silva, V. T., & Vasconcelos, W. W. (2017). Detection and resolution of normative conflicts in multi-agent systems: A literature survey. *Autonomous Agents and Multi-Agent Systems*, *31*(6), 1236–1282. <https://doi.org/10.1007/s10458-017-9362-z>.
86. Savarimuthu, B. T. R., & Cranefield, S. (2011). Norm creation, spreading and emergence: A survey of simulation models of norms in multi-agent systems. *Multiagent and Grid Systems*, *7*(1), 21–54. <https://doi.org/10.3233/MGS-2011-0167>.
87. Savarimuthu, B. T. R., Purvis, M., Cranefield, S., & Purvis, M. K. (2007). Mechanisms for norm emergence in multiagent societies. In *6th international joint conference on autonomous agents and multiagent systems* (p. 173) (AAMAS 2007), Honolulu, Hawaii, USA, May 14–18, 2007. <https://doi.org/10.1145/1329125.1329335>.
88. Savarimuthu, B. T. R., Cranefield, S., Purvis, M., & Purvis, M. (2008). Role model based mechanism for norm emergence in artificial agent societies. In *Coordination, organizations, institutions, and norms in agent systems III* (pp. 203–217), Springer.
89. Savarimuthu, B. T. R., Cranefield, S., Purvis, M. K., & Purvis, M. A. (2009a). Norm emergence in agent societies formed by dynamically changing networks. *Web Intelligence and Agent Systems: An International Journal*, *7*(3), 223–232. <https://doi.org/10.3233/WIA-2009-0164>.
90. Savarimuthu, B. T. R., Purvis, M., Purvis, M., & Cranefield, S. (2009b). Social norm emergence in virtual agent societies. In M. Baldoni, T. C. Son, M. B. van Riemsdijk, & M. Winikoff (Eds.), *International Workshop on Declarative Agent Languages and Technologies* (pp. 18–28). Berlin: Springer.
91. Savarimuthu, B. T. R., Arulanandam, R., & Purvis, M. (2011a). Aspects of active norm learning and the effect of lying on norm emergence in agent societies. In *Proceedings of the 14th international conference on agents in principle, agents in practice* (pp. 36–50), Springer-Verlag, Berlin, Heidelberg, PRIMA'11. https://doi.org/10.1007/978-3-642-25044-6_6.
92. Savarimuthu, B. T. R., Arulanandam, R., & Savarimuthu, S. (2011b). Emergence of a sharing norm in a simulated hunter-gatherer society. In *Proceedings of the 2011 IEEE/WIC/ACM international conference on intelligent agent technology* (pp. 34–37), IAT 2011, Campus Scientifique de la Doua, Lyon, France, August 22–27, 2011. <https://doi.org/10.1109/WI-IAT.2011.141>.
93. Sen, S., & Airiau, S. (2007). Emergence of norms through social learning. In *Proceedings of the 20th international joint conference on artificial intelligence* (pp. 1507–1512), Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, IJCAI'07. <http://dl.acm.org/citation.cfm?id=1625275.1625519>.
94. Serramia, M., López-Sánchez, M., Rodríguez-Aguilar, J. A., Rodríguez, M., Wooldridge, M. J., Morales, J., & Ansótegui, C. (2018). Moral values in norm decision making. In E. André, S. Koenig, M. Dastani, G. Sukthankar (Eds.), *Proceedings of the 17th international conference on autonomous agents and multiagent systems* (pp. 1294–1302), AAMAS 2018, Stockholm, Sweden, July 10–15, 2018. International Foundation for Autonomous Agents and Multiagent Systems Richland, SC, USA / ACM. <http://dl.acm.org/citation.cfm?id=3237891>.
95. Shams, Z., De Vos, M., Padget, J., & Vasconcelos, W. W. (2017). Practical reasoning with norms for autonomous software agents. *Engineering Applications of Artificial Intelligence*, *65*, 388–399. <https://doi.org/10.1016/j.engappai.2017.07.021>.

96. Shoham, Y., & Tennenholtz, M. (1995). On social laws for artificial agent societies: Off-line design. *Artif Intell*, 73(1–2), 231–252. [https://doi.org/10.1016/0004-3702\(94\)00007-N](https://doi.org/10.1016/0004-3702(94)00007-N).
97. Swarup, S., Apolloni, A., & Fagyal, Z. (2011). A model of norm emergence and innovation in language change. In *The 10th international conference on autonomous agents and multiagent systems-Volume 2, international foundation for autonomous agents and multiagent systems* (pp. 693–700), AAMAS '11.
98. Székely, Á., Andrighetto, G., & Nardin, L. G. (2016). *Social norms and extortion rackets* (1st ed.). New York: Springer.
99. Tuomela, R. (1995). The importance of us: A philosophical study of basic social notions. Research in Applied Linguistic, Stanford University Press, <https://books.google.co.uk/books?id=Vj6sAAAAIAAJ>.
100. Vanhee, L., Aldewereld, H., & Dignum, F. (2011). Implementing norms? In *Proceedings of the 2011 IEEE/WIC/ACM international conferences on web intelligence and intelligent agent technology—Volume 03, IEEE Computer Society* (pp. 13–16), Washington, DC, USA. WI-IAT '11. <https://doi.org/10.1109/WI-IAT.2011.184>, <http://ieeexplore.ieee.org/document/6040693/>.
101. Villatoro, D., Sen, S., & Sabater-Mir, J. (2009). Topology and memory effect on convention emergence (pp. 233–240). IEEE. <https://doi.org/10.1109/WI-IAT.2009.155>, <http://ieeexplore.ieee.org/document/5284837/>.
102. Villatoro, D., Sabater-Mir, J., & Sen, S. (2011). Social instruments for robust convention emergence. In T. Walsh (Ed.), *IJCAI 2011, proceedings of the 22nd international joint conference on artificial intelligence* (pp. 420–425), Barcelona, Catalonia, Spain, July 16–22, 2011. IJCAI/AAAI. <https://doi.org/10.5591/978-1-57735-516-8/IJCAI11-078>.
103. Villatoro, D., Sabater-Mir, J., & Sen, S. (2013). Robust convention emergence in social networks through self-reinforcing structures dissolution. *ACM Transactions on Autonomous and Adaptive Systems*, 8(1), 1–21. <https://doi.org/10.1145/2451248.2451250>.
104. Visser, S., Thangarajah, J., Harland, J., & Dignum, F. (2016). Preference-based reasoning in BDI agent systems. *Autonomous Agents and Multi-Agent Systems*, 30(2), 291–330. <https://doi.org/10.1007/s10458-015-9288-2>.
105. Vouros, G. A. (2015). The emergence of norms via contextual agreements in open societies. In F. Koch, C. Guttmann, & D. Busquets (Eds.), *Advances in Social Computing and Multiagent Systems* (pp. 185–201). Cham: Springer.
106. Yu, C., Zhang, M., Ren, F., & Luo, X. (2013). Emergence of social norms through collective learning in networked agent societies. In *Proceedings of the 2013 international conference on Autonomous agents and multi-agent systems, international foundation for autonomous agents and multiagent systems* (pp. 475–482).

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.