

# Requirements for an Artificial Agent with Norm Competence

[Authors' Disciplines: Psychology, Cognitive Science, Philosophy, Computer Science, Robotics]

Affiliations and Addresses Go Here Using the Affiliation and Address Style (Required)  
publications19@aaai.org

## Abstract

Human behavior is frequently guided by social and moral norms, and no human community can exist without norms. Robots that enter human societies must therefore behave in norm-conforming ways as well. However, currently there is no solid cognitive or computational model available of how human norms are represented, activated, and learned. We provide a conceptual and psychological analysis of key properties of human norms and identify the demands these properties put on any artificial agent that incorporates norms—demands on the format of norm representations, their structured organization, and their learning algorithms.

No human community can exist without norms (Hechter and Opp 2001; Ullmann-Margalit 1977), and many past human communities have gone extinct with suboptimal systems of norms (Wilson 2002). It stands to reason that communities that include both humans and machines as partners will also not succeed without norms. If this is true then we need to understand and formalize what norms are in the human mind—how people represent, learn, activate, update, and deploy norms to guide their behavior—so that we can effectively design artificial agents with appropriate capacities to represent and obey norms.

If an artificial agent is to acquire human norms, its formalisms and algorithms must be informed by the properties of human norms—how humans represent norms, learn them, and use them to guide behavior. We introduce here core properties of human norms and define the demands these properties put on any artificial agent that incorporates norms. These demands range from the format of norm representations to their structured organization, from learning algorithms to communication skills.

In sociology and experimental economics, the importance of norms has long been recognized (Parsons 1951; Schelling 1960). These literatures try to explain human cooperation despite the individual's rational self-interest, and norms are an external force that constrains human action. But it is not known how such external forces can operate cognitively and computationally. A person complying with norms must have something in their mind that allows their action to conform to the norm, and that something may be called a *norm*

*representation*. A sizable set of studies has examined the automatic activation of such norm representations by situation cues—for example, garbage on the floor triggers the “don't litter” norm (Cialdini, Reno, and Kallgren 1990), or the sight of a library triggers the “be quiet” norm (Aarts and Dijksterhuis 2003). But no cognitive model has been offered that specifies at least some of the key properties of norm representations. This is what we attempt to do here.

## Properties of Human Norms

### Working Definition of Norm

We adopt the following working definition of a norm ([citation concealed]; cf. Bicchieri 2006; Brennan et al. 2013):

A norm is an instruction to (not) perform a type of action in a given context, provided that a sufficient number of individuals in a community (i) indeed follow this instruction and (ii) demand of each other to follow the instruction.

This definition captures both the “external” aspect of norms (that they are obeyed and enforced by communities) and their “internal” aspect (that they guide actions). Accordingly, “instruction” can sometimes refer to one person or group literally instructing another person, sometimes to a self-instruction that motivates a person to perform a certain action. The term *action* covers a broad class—including physical (observable) or mental acts, omissions, as well as acts of bringing about a certain outcome.

The separation into conditions (i) and (ii) follows a long tradition of considering both “descriptive” and “injunctive” elements of norms (Bicchieri 2006; Gibbs 1965; Cialdini, Reno, and Kallgren 1990).

We can call condition (i) the *prevalence component* of a norm—that members of a community do in fact follow a norm (with a certain degree of consistency); and we can call condition (ii) the *normative demand component* of a norm—the degree to which community members demand of one another that each follow the norm. Because of this normative

demand, violations of norms often lead to sanctions (e.g., criticism, isolation, reform, punishment). However, the existence of sanctions need not be part of the definition of a norm, as some have argued (e.g., Bendor and Swistak 2001). It is entirely conceivable that for some norm in some communities, no sanctions have been necessary to uphold a norm. That would obviously not make it any less of a norm.

For something to be a *norm* requires that people in a relevant community meet conditions (i) and (ii). For someone to have a *norm representation* requires that the agent knows conditions (i) and (ii); and for someone to show *norm compliance*, the agent both knows conditions (i) and (ii) and tries to follow the instruction *because* of conditions (i) and (ii).

*Related concepts.* Norms differ from other action guides, such as preferences, goals, and collective habits. The normative force condition marks this difference. A lot of people put milk in their coffee, but they do not demand it of each other, so this action is not a norm but a wide-spread preference. By contrast, getting in line to order coffee is a norm, because that is what people would expect of each other. Norms also differ from values, and the notion of being instructions to act in a particular context marks this difference. Values (such as fairness, freedom, dignity) typically govern a larger class of possible actions/outcomes across a wider range of contexts.

### Implied and Suggested Properties

We discuss now six properties of norms, either implied by or further elaborated from the working definition, and develop demands that these properties put on artificial agents' norm representations.

#### 1. Multiple norm types

A first property is that norms can be of multiple deontic types: at least prescriptions, prohibitions, and permissions. Thus, the representation of a norm-guided action must signal which of the types governs the particular action.

#### 2. Context sensitivity

A second critical property of norms is that they are context-specific and must somehow be activated by characteristic features of a given context. In initial research we asked people to state the prescription norms that apply for a variety of everyday scenes (e.g., board room, jogging path), and 95% of these norms were specific to the context [cit. concealed].

#### 3. Community prevalence

A third property is that genuine norms have sufficient prevalence—that is, most community members comply with a given norm and are aware of that collective compliance. In initial research we found that the top 8 prescription or prohibition norms that people stated for various everyday scenes showed prevalence rates between 94% (most prevalent) and 21% (8<sup>th</sup>-most prevalent). Beyond the actual agreement rates, people also have a *belief* about the norm's

prevalence in the community, which we can call the *prevalence parameter*. Though this belief is not necessarily accurate (Prentice & Miller, 1996), it will often be based on reasonably representative behavioral data.

#### 4. Graded normative demand

A rarely noted property is that (at least) prescriptions and prohibitions come in degrees of community demand [cit. concealed]. At least in English, terms of prescription can capture low demand (“it is suggested to A”) to high demand (“it is required to A”), with further gradations in between. Likewise, terms of prohibition can capture low demand (“it is frowned upon to A”) to high demand (“it is forbidden to A”). In preliminary research we found that people show high consensus in ordering these terms along a dimension of normative demand, and Figure 1 shows such orderings for all three norm types [cit. concealed]. Thus, human norm representations include a graded normative demand parameter.

A norm's degree of normative demand is likely to be closely related to its prevalence, but the two are not reducible to each other. In general, the more strongly people demand of others to conform to norm  $N_i$ , the more people will obey it. But when community norms change, strong demand sometimes lingers even though prevalence is declining (Mack 2018); and for some norms of only modest normative demand, prevalence may be high (e.g., if the benefits of norm conformity are considerable). Even though the exact relationship between prevalence and normative demand is unknown, as a first approximation we can assume that demand is a linear function of prevalence and other factors such as severity of consequences.

#### 5. Resolving norm conflict by normative demand

Sometimes norms stand in conflict with one another such that, in the given context, every action violates at least one norm (e.g., in moral dilemmas). Because normative demand comes in degrees, violating some norms will be more costly than violating others, so norm conflict resolution will have to take graded normative demand into account.

#### 6. Structural representation of norms

The norms relevant to a particular context must be organized in some way, and one plausible organization is a network structure that has some core (strong or consensual) norms at the center and other (weaker or less consensual) norms in

required		frowned upon
expected	optional	objectionable
called for	discretionary	unacceptable
encouraged	tolerated	barred
suggested		forbidden

Figure 1. Linguistic terms referring to degrees of normative demand for prescriptions (left), prohibitions (right), and permissions (middle).

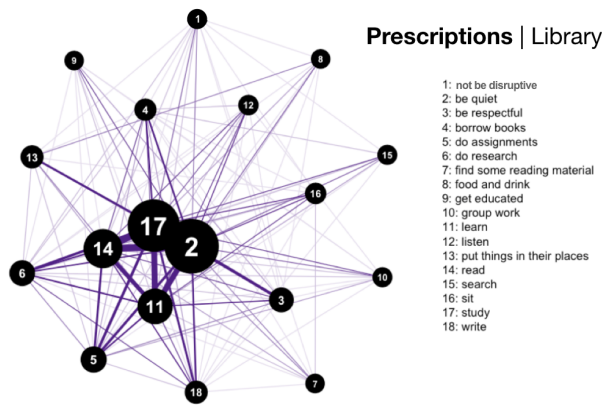


Figure 2. Network representation of prescriptions people generated when exposed to the picture of a library. Larger circles indicate norms that were mentioned by more people.

the periphery. In our preliminary research we found that norms generated by a sample of participants for specific contexts clustered together in such networks (see Figure 2 as an illustration; [cit. concealed]). We also found that, when exposed to a particular context, norms mentioned earlier were more prevalent—thus, they were likely at the core of the network.

## Formal Representation

Most formal approaches to norm representations have used logical formalisms (Bringsjord, Arkoudas, and Bello 2006; Pereira and Saptawijaya 2007). One way of formalizing context-specific norms is to introduce a deontic operator  $\mathbb{D}$  (instantiating a prescription, prohibition, or permission) for an action  $A$ , thus  $N_i := C_j \rightarrow \mathbb{D}(A_k)$ .  $C_j$  may be defined extensionally as a set of preconditions in the world under which  $A_k$  is prescribed/prohibited/permitted. It is implausible that these preconditions are linked together as a long conjunction; more plausible is an overall likelihood estimation that aggregates the presence of features into something akin to a sufficient statistic. Given a sample of features from a population, it should be more likely that  $C_j$  holds than that any other  $C$  holds, even if the specific sample may slightly differ from instance to instance.

However, the material implication  $C_j \rightarrow \mathbb{D}(A_k)$  has a number of unattractive properties for deontic reasoning (e.g., Chisholm 1963). An alternative formulation would be  $N_i := \mathbb{D}(A_k, C_j)$ , in which the deontic operator establishes a relation between actions and contexts. To capture the parameter of *graded normative demand*, however, the classical interpretation of  $\mathbb{D}$  must be expanded to take on a value between 0 and 1, where 0 = prohibited, 0.5 = permitted (optional), and 1.0 = prescribed. A recent proposal further integrated the *prevalence* parameter by defining prevalence as the uncertainty over an estimated deontic value, formally a

confidence interval around the normative demand estimate [cit. concealed].

## Requirements for Norm-Competent Agents

Summarizing the above properties of norm representations and the suggested formal representation we can now identify a number of requirements for an artificial agent to appropriately represent and learn norms.

### 1. Multiple Norm Types and Action Planning

The three deontic categories (prescriptions, prohibitions, or permissions) must be properly mapped onto the agent’s planning and action modules. *Prescriptions* generate goals, with a goal priority value dictated by the prescriptions’ normative demand values. *Permissions* affirm current goals that are already set or pursued by the system. *Prohibitions* stop the pursuit of goals, if during this pursuit any actions or consequences fall under prohibited actions or outcomes. In principle, one might wish that during planning *any* considered action path be compared to the agent’s norm network to avoid violating prohibitions, but this strategy could quickly lead to computational explosion. The number of norms against which the candidate action is compared must be restricted. Limiting the contexts the agent may find itself in would help. In addition, the agent’s planning could be primarily guided by context-specific *prescriptions* because, if the right ones are activated, they are safe to pursue, and their numbers will generally be manageable. If the agent is about to pursue a goal, then it could first be compared against this more limited prescription set for the given context, and if there is no match, a question of clarification may be in order (e.g., “This goal is not among my duties...”).

### 2. Context Sensitivity and Context Recognition

The property of context sensitivity requires that the agent recognize what context it is in and activate the context-appropriate norms. It is currently unknown exactly how humans recognize contexts and how context activates the relevant norms. But it is clear that many elements go into context: space (e.g., the room one enters), time (e.g., morning vs. evening), event type (e.g., party, debate), who is present in what role (e.g., friends, authorities), agent’s own role (e.g., assistant, guest), and more. In order to activate the right set of norms, the norms “preconditions” must be tested. at least two approaches are available [cit. concealed]. One, the agent may collect a sufficient number of features from the observed environment and from its knowledge base that make being in a particular context  $C_j$  likely, which then activates the norm set that is relevant for  $C_j$ . Second, the agent may activate specific norms relevant to specific features in the environment (e.g., a chair, someone else’s phone, someone giving a speech) without necessarily making an overall categorical determination which “context” it is in; instead, it would rely on the world to naturally make features co-occur

such that the set of feature-activated norms turns out to be the right set for the particular context.

Whichever model is correct, both context (feature) recognition and norm activation put enormous demands on the agent in terms of perception, attention, and categorization capacities. In general, sophisticated scene understanding is not possible for today's artificial agents. Incredible progress has been made in object classification, but recognizing relations that define contexts (e.g., *when this person in this role performs this action in this space*) is currently out of reach. What is challenging about relational information is that many relations that uniquely pick out contexts need to be recognized over different temporal scales, and all sorts of relations could exist (in principle) between the various relations (i.e. objects, events, etc.) in a scene. Considering all possible relations and testing whether they indicate that one is in a certain context is computationally inefficient.

Memory for the objects and events that typically appear in contexts might provide initial hypotheses for which context one might be in, but verifying and disconfirming these hypotheses requires further action. One candidate to supplement memory-based relational processing is *attention* (Kim, Ricci, and Serre 2018), which offers several advantages. First, top-down attentional guidance will narrow the space of relations to be considered, and attention plausibly realizes so-called "visual routines" (Ullman 1984), which can be run to verify whether or not a relation holds (Bello et al. 2018; Yuan, Uttal, and Franconeri 2016).

Artificial agents might also do better when constraints are in place on the contexts to be recognized. First, the agent could be deployed in a limited domain that has a limited number of contexts (e.g., a nursing home robot, a cafeteria bussing robot). The agent could then be equipped with a knowledge base of reasonably reliable indicators of the likely (sub)contexts it could find itself. Searching for such tell-tale indicators (rather than scanning all possible scene features) would greatly speed up processing and increase accuracy of context recognition. Second, hypotheses about which context the agent is in and hypotheses about which norms apply could mutually constrain one another. That is, some initially activated norms (e.g., activated by a few salient objects or persons in the scene) might serve to narrow down just what context an agent is in. As norms are highly context specific, the co-activation of even just a small subset of initially activated norms could make it very likely that one is in context  $C_j$ , rather than in any other candidate context. This hypothesis of being in  $C_j$  would then initiate selective attention to additional features that tend to uniquely characterize  $C_j$ , further (dis)confirming this context hypothesis. By extension, the context hypothesis makes predictions about additional applicable norms that can be tested by observing other agents' behaviors or by calculating their joint likelihood with the subset of initially activated norms.

### 3. Prevalence Parameter

An agent's prevalence parameter would reflect an estimate of the proportion of people in the community who actually obey the norm. Mere observation will typically not suffice to achieve such an estimate (see our discussion of the limits of observation below), so the designer will have to provide a starting value that can be updated, through both observing community members and querying them about their norm perceptions and practices (just like humans in foreign countries query locals about their norms).

### 4. Normative Demand Parameter

Including a graded normative demand parameter will require collating evidence from a number of sources in the relevant community: prevalence of norm compliance (because stronger norms tend to be followed more consistently); intensity of demand expressions (e.g., verbal exhortations, warnings); and intensity of sanctions upon violation (e.g., yelling at a person for committing a violation).

### 5. Norm conflict resolution algorithms

An agent facing conflicts among norms will have to rely on the conflicting norms' demand parameters to minimize violation costs. For example, Kasenberg and Scheutz (2018) proposed an algorithm within a Markov Decision Process framework where norms are represented in linear temporal logic as temporal expressions that the agent intends to make true. In doing so, agents attempt to obey as many norms as possible, closely monitor the relative normative demand (importance), and minimize aggregated violation costs (whereby greater costs accrue for stronger norms and for temporally more extended states of violation).

### 6. Structured Organization

Building norm representations that have a structured organization (e.g., network structure) will be challenging because little is known about how human norm networks are organized. A plausible hypothesis is that such networks have a core and a periphery, with core nodes being more prevalent and/or more important. Nodes would represent action instructions (with a normative demand parameter), and edges would stand for the probability of co-activation in the same context. It can be expected that human norm networks have "small-world properties": Nodes in the network are connected to only a small subset of all nodes and have high degrees of clustering, but the clusters can be traversed with relatively short path lengths (i.e., there are many context-specific subnetworks that are nonetheless connected by linking nodes). These are characteristics we find in semantic networks (Steyvers and Tenenbaum 2005). Norm networks, of course, are not just another version of semantic networks. For example, norms and their high degree of context specificity would likely show even more clustering and sparseness; moreover, the parameters associated with norms (demand, prevalence, etc.) differ considerably from the properties of word meanings.

## Implications for Norm Learning

We now turn to the greatest challenge of designing norm-competent artificial agents: how they could acquire norms. We begin with the case of human norm learning and then develop implications for norm learning in artificial agents.

### Human Norm Learning

Most generally, norm learning refers to the process of extracting  $\mathbb{D}(A_k | C)$  relations from evidence in the world, estimating prevalence and normative demand. What is this evidence in the world that reveals norms? We categorize this evidence into four types and briefly discuss each type.

#### Explicit instructions

The most direct evidence for a norm is its declaration, in symbols (e.g., signs) or verbal utterances. However, signs are rare (consider how few laws are publicly displayed), and verbal norm instructions may be even rarer (Wright and Bartsch 2008). When provided, however, such explicit instructions can both provide information about prevalence (e.g., “Everybody here is making a donation...”) and also scale community demand by choosing the graded linguistic expressions of normative demand, as shown in Figure 1.

#### Behavior patterns

Observing other community members’ behavior is a powerful second type of evidence for the presence of a norm (e.g., Milgram, Bickman, and Berkowitz 1969). For example, by looking at others in a cafeteria we deduce whether we are expected to bus our own dishes or someone else does it. Already from ages 2 to 3 on, children readily infer norms from other people’s behavior (Rakoczy, Warneken, and Tomasello 2008). Mere behavioral observation, however, provides only limited information. First, whereas trends of behavior reveal the prevalence of prescriptions, they are sparse with respect to prohibitions, because when people comply with prohibitions there is typically no behavior to be observed. Second, prevalent behaviors performed by a number of people can also be *desired* by those individuals, rather than reflecting a norm. For example, on a warm summer day at the beach many people eat ice-cream, and many people stand in line at the ice-cream stand; only the latter behavior is norm-guided. To differentiate norm-guided from desired behavior additional information is needed. On the behavior side itself, a group of people all performing a certain behavior in highly similar ways increases the likelihood of a norm operating. Patiently waiting one’s turn in line is quite orderly and uniform, whereas eating the ice cream afterwards shows more variability. Beyond the behavior itself, critical evidence to distinguish norms from desires lies in consequences of the observed behavior, the third type of evidence.

#### Behavior Consequences

Relevant consequences include at least two kinds: costs for the agent and benefits for other people. If an agent’s

foregone alternative behaviors would be individually more attractive, then the observed behavior is *costly* for the agent and suggests the presence of a norm (Henrich, 2009). Everybody would prefer to order ice-cream right when they arrive at the stand, so waiting in line is costly and likely norm-guided. In addition, if an agent’s observed behavior causes *benefits* to others, then this provides evidence for a norm, as with tipping, table manners, holding doors open, etc. Conversely, rare behaviors with visible negative impact on others (taking another person’s ice-cream instead of purchasing one) suggest a violated norm of prohibition.

#### Social (dis)approval

Community members’ expressions of approval or disapproval of a performed behavior constitute the fourth type of evidence. Disapproval, such as chiding someone who cuts in front of the line, clearly reveals a violated prohibition, both to the person who violated the norm and to an observer. Disapproval often comes in degrees through varying facial, verbal, and bodily signals, so normative demand can be inferred; indeed, Brauer and Chaurand (2010) found that a violation’s perceived degree of “deviance” (i.e., a proxy for normative demand) is a strong predictor of likelihood of expressed social disapproval.

Expressions of approval for a performed behavior may suggest a prescriptive norm that has been met, but such approvals are fairly rare and increase primarily when the behavior exceeds, rather than just meets, the relevant norm. Nobody gets praised for standing in line or treating others with respect, precisely because the norm has made compliance literally “normal.” Approval for omissions could indicate prohibitions that were upheld, but such praise is even rarer (“good job for not cheating on the test”).

### Machine Norm Learning

Artificial agents, just as humans, should be able to learn from instruction, observation of behavior and consequences, and from social (dis)approval. Learning norms from instruction (written or spoken) is challenging for many reasons, not the least is that the preconditions (context) and the action will always be underspecified. Nonetheless, some success has been reported in robots learning recipes (analogous to cooking norms) from written data (Nyga and Beetz 2012) and learning new action norms from spoken commands (Scheutz et al. 2017). In the broadest sense, programming a robot with a set of a priori norms is a form of teaching by instruction as well, though still challenging because a robot that knows *If C then A* needs to identify the instances in which *C* holds and instantiate the right kind of *A*.

Learning norms from observation may be enabled by Inverse Reinforcement Learning (IRL; e.g., Arai and Suzuki 2014), which has been proposed to ensure that agents “align with human values” (Russell, Dewey, and Tegmark 2016). There, the agent observes other agents’ behavior as well as

the rewards and punishments those agents receive, and it derives a value function that encodes what the proper behaviors are. IRL algorithms can grasp behavior patterns in specific contexts and may be able to code for degrees of norm demand. But without further enrichment, this approach cannot distinguish between actions that benefit the individual agent and actions that benefit the community (Arnold, Kasenberg, and Scheutz 2017), or between norms that hold for some people but not for others (e.g., observing students and teachers in the classroom, the agent would infer that the teacher violates norms).

Reinforcement learning (RL) approaches to norms (Abel, MacGlashan, and Littman 2016; Li et al. 2015) are responsive to rewards and punishment, which could come in the form of social (dis)approval. Thus, if trustworthy community representatives give the agent feedback, it could learn appropriate actions for specific contexts. However, we would not want learning agents to “experiment” in social environments and learn from trial and error what appropriate actions it should take. Such training could occur in virtual and game-like worlds, but creating those worlds may be as difficult as building norms into the agent from the start.

Recent work by Savarimuthu et al. (2013) identified norms from sanctioning behavior in multi-agent simulations. This approach takes advantage of the diagnostic evidence of social (dis)approval, but in natural environments it faces the problem that praise is infrequent (hence learning prescriptions becomes difficult), and it would learn prohibitions only when it commits or observes a violation (which is a costly form of learning). In addition, disapproval is less frequently expressed in societies with high norm compliance—which is exactly where one would want to “raise” an artificial learning agent).

Most generally, none of the foregoing approaches learn norm *representations*; they learn only how to *behave* in accordance with observed human patterns or human sanctions. As a result, they cannot express what they have learned, only that the learned action is the “best” in this context. But the standard of what is best may not be a norm but could be a desire or habit, or even a physically convenient movement (e.g., going upright is not a norm but anatomically convenient for humans). Moreover, such agents cannot represent norm conflicts, because without having representations of norms they cannot diagnose norms as standing in conflict. Yet, representing norm conflicts, and explaining how they should be resolved, will be critical in human-robot interaction, when action recommendations differ or when the human is surprised by the robot’s behavior. Trying to address these limitations, some authors have developed algorithms that learn explicit norm representations from observed behavior and are therefore able to recognize norm conflicts and attempt to resolve those conflicts (Kasenberg and Scheutz 2018, 2017).

None of the current algorithms, however, can exploit the rich information contained in (foregone and actual) consequences of observed behaviors. To that end, agents would have to infer an observed person’s goals, assess costs and benefits for the person and for other individuals, and compare actual to counterfactual actions (e.g., the agent’s foregone benefits as an indicator of prescription norms). Thus, agents would need to have social-cognitive capacities to ground their norm competence.

Given the constraints of learning from observation (e.g., no observed data on prohibitions, no sanctioning data on prescriptions) and the current limitations of algorithms for such learning, observation alone will not generate norm-competent agents. There would be too few data points to grasp the complex context specificity of human norms, too little knowledge about important distinctions in the behavior stream, and no sense of which observed behaviors represent norms, rather than desires or habits.

Currently, and for the foreseeable future, the safest and most effective way of designing norm-competent agents would therefore be a “hybrid” approach (Allen, Smit, and Wallach 2005) in which *a priori* legal, moral, and social norms combine with abilities to learn new norms and update existing norms. Science would have to identify the relevant *a priori* norms (e.g., for a robot in a particular role in a particular community) and implement them in ways that replicate key properties of human norms, such as context specificity, graded normative demand, and network organization. Successful artificial agents would then update this starting package by learning from instruction, observing behaviors, consequences, and social (dis)approval, and requesting advice when necessary. Continuous teaching by instruction and observation is attractive in part because robots deployed in social contexts will be surrounded by teachers. Not all community members are equally good teachers (and as we know from the cases of Tay and Chappie, some will actively try to corrupt the agent). But these problems arise for human children learning norms as much as for robots learning norms. We must trust human communities to make up for the failings of some teachers and find ways to correct individual agents’ missteps.

## Conclusion

We have drawn a map of human norm representations that can guide the development of comparable norm representations in artificial agents. Creating such norm-competent agents is a significant challenge, but it is vital for a future society in which human and artificial agents will co-exist, both guided by the social and moral norms of their shared community.

## References

- Aarts, H., and Dijksterhuis, A. 2003. The Silence of the Library: Environment, Situational Norm, and Social Behavior. *Journal of Personality and Social Psychology* 84 (1): 18–28.
- Abel, D., MacGlashan, J., and Littman, M. L. 2016. Reinforcement Learning as a Framework for Ethical Decision Making. In *AAAI Workshop: AI, Ethics, and Society, Volume WS-16-02 of 13th AAAI Workshops*. AAAI Press.
- Allen, C., Smit, I., and Wallach, W. 2005. Artificial Morality: Top-down, Bottom-up, and Hybrid Approaches. *Ethics and Information Technology* 7 (3): 149–155.
- Arai, S., and Suzuki, K. 2014. Encouragement of Right Social Norms by Inverse Reinforcement Learning. *Journal of Information Processing* 22 (2): 299–306.
- Arnold, T., Kasenberg, D., and Scheutz, M. 2017. Value Alignment or Misalignment – What Will Keep Systems Accountable? In *The Workshops of the Thirty-First AAAI Conference on Artificial Intelligence: Technical Reports, WS-17-02: AI, Ethics, and Society*, 81–88. Palo Alto, CA: The AAAI Press.
- Bello, P., Lovett, A., Briggs, G., and O’Neill, K. 2018. An Attention-Driven Computational Model of Human Causal Reasoning. In *Proceedings of the 40th Annual Meeting of the Cognitive Science Society*, 1353–1358. Austin, TX: Cognitive Science Society.
- Bendor, J., and Swistak, P. 2001. The Evolution of Norms. *American Journal of Sociology* 106 (6): 1493–1545.
- Bicchieri, C. 2006. *The Grammar of Society: The Nature and Dynamics of Social Norms*. New York, NY: Cambridge University Press.
- Brauer, M., and Chaurand, N. 2010. Descriptive Norms, Prescriptive Norms, and Social Control: An Intercultural Comparison of People’s Reactions to Uncivil Behaviors. *European Journal of Social Psychology* 40 (3): 490–499.
- Brennan, G., Eriksson, L., Goodin, R. E., and Southwood, N. 2013. *Explaining Norms*. New York, NY: Oxford University Press.
- Bringsjord, S., Arkoudas, K., and Bello, P. 2006. Toward a General Logicist Methodology for Engineering Ethically Correct Robots. *Intelligent Systems, IEEE* 21 (4): 38–44.
- Chisholm, R. M. 1963. Contrary-to-Duty Imperatives and Deontic Logic. *Analysis* 24 (2): 33–36.
- Cialdini, R. B., Reno, R. R., and Kallgren, C. A. 1990. A Focus Theory of Normative Conduct: Recycling the Concept of Norms to Reduce Littering in Public Places. *Journal of Personality and Social Psychology* 58 (6): 1015–1026.
- Gibbs, J. P. 1965. Norms: The Problem of Definition and Classification. *American Journal of Sociology* 70 (5): 586–594.
- Hechter, M., and Opp, K.-D. 2001. *Social Norms*. Russell Sage Foundation.
- Kasenberg, D., and Scheutz, M. 2017. Interpretable Apprenticeship Learning with Temporal Logic Specifications. In *Proceedings of the 56th IEEE Conference on Decision and Control (CDC 2017)*, 4914–4921. Piscataway, NJ: IEEE Press.
- . 2018. Norm Conflict Resolution in Stochastic Domains. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence*.
- Kim, J., Ricci, M., and Serre, T. 2018. Not-So-CLEVR: Learning Same-Different Relations Strains Feedforward Neural Networks. *Interface Focus* 8 (4): 20180011.
- Li, J., Meneguzzi, F., Fagundes, M., and Logan, B. 2015. Reinforcement Learning of Normative Monitoring Intensities. In *Proceedings of the International Workshop on Coordination, Organization, Institutions and Norms in Multi-Agent Systems*.
- Mack, A., ed. 2018. Changing Social Norms. *Social Research: An International Quarterly* 85 (1): 1–271.
- Milgram, S., Bickman, L., and Berkowitz, L. 1969. Note on the Drawing Power of Crowds of Different Size. *Journal of Personality and Social Psychology* 13 (2): 79–82.
- Nyga, D., and Beetz, M. 2012. Everything Robots Always Wanted to Know about Housework (but Were Afraid to Ask). In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 243–250. Piscataway, NJ: IEEE Press.
- Parsons, T. 1951. *The Social System*. Glencoe, IL: Free Press.
- Pereira, L. M., and Saptawijaya, A. 2007. Modelling Morality with Prospective Logic. In *Progress in Artificial Intelligence*, edited by J. Neves, M. F. Santos, and J. M. Machado, 99–111. Lecture Notes in Computer Science. Springer Berlin Heidelberg.
- Rakoczy, H., Warneken, F., and Tomasello, M. 2008. The Sources of Normativity: Young Children’s Awareness of the Normative Structure of Games. *Developmental Psychology* 44 (3): 875–881.
- Russell, S., Dewey, D., and Tegmark, M. 2016. Research Priorities for Robust and Beneficial Artificial Intelligence. *ArXiv Preprint ArXiv:1602.03506*.
- Savarimuthu, B. T. R., Cranefield, S., Purvis, M. A., and Purvis, M. K. 2013. Identifying Prohibition Norms in Agent Societies. *Artificial Intelligence and Law* 21 (1): 1–46.
- Schelling, T. C. 1960. *The Strategy of Conflict*. Cambridge, MA: Harvard University Press.
- Scheutz, M., Krause, E., Oosterveld, B., Frasca, T., and Platt, R. 2017. Spoken Instruction-Based One-Shot Object and Action Learning in a Cognitive Robotic Architecture. In *Proceedings of the 16th Conference on Autonomous Agents and Multi-Agent Systems*, 1378–1386. AAMAS ’17. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems.
- Steyvers, M., and Tenenbaum, J. B. 2005. The Large-Scale Structure of Semantic Networks: Statistical Analyses and a Model of Semantic Growth. *Cognitive Science* 29 (1): 41–78.
- Ullman, S. 1984. Visual Routines. *Cognition* 18 (1–3): 97–159.
- Ullmann-Margalit, E. 1977. *The Emergence of Norms*. Clarendon Library of Logic and Philosophy. Oxford: Clarendon Press.
- Wilson, D. S. 2002. *Darwin’s Cathedral: Evolution, Religion, and the Nature of Society*. Chicago: University of Chicago Press.
- Wright, J. C., and Bartsch, K. 2008. Portraits of Early Moral Sensibility in Two Children’s Everyday Conversations. *Merrill-Palmer Quarterly* 54 (1): 56–85.
- Yuan, L., Uttal, D., and Franconeri, S. 2016. Are Categorical Spatial Relations Encoded by Shifting Visual Attention between Objects? *PLOS ONE* 11 (10): e0163141.