Learning Behavioral Norms in Uncertain and Changing Contexts

Vasanth Sarathy and Matthias Scheutz
Dept. of Computer Science
Tufts University
Medford, MA
Email: {vasanth.sarathy,matthias.scheutz}@tufts.edu

Bertram F. Malle
Dept. of Cogn., Ling., and Psych. Sciences
Brown University
Providence, RI
Email: bfmalle@brown.edu

***Abstract*—Social and moral norms guide human behavior. Robots that enter human societies must therefore behave in norm-conforming ways as well to increase coordination, predictability, and safety in human-robot interactions. However, human norms are context-specific and laced with uncertainty, making the representation, learning, and communication of norms challenging. We propose a formal norm representation using deontic logic, Dempster-Shafer Theory, and a machine learning algorithm that allows an artificial agent to learn norms under uncertainty from human data. We demonstrate a novel cognitive capability with which an agent can dynamically learn norms while being exposed to distinct contexts, recognizing the unique identity of each context and the norms that apply in it.**

I. Introduction

Human social communities function more effectively a shared set of norms guides their actions [1], [2]. A system of norms has many advantages [3]: Norms streamline the selection of actions and standardize behaviors across situations and time; actions consistent with norms are more predictable and understandable; and norm-consistent actions increase coordination and cooperation in the community.

Some norms are explicitly laid down in laws and regulations, whereas social and moral norms reveal themselves more implicitly in the actions of community members. When such actions are observable, humans can learn what they should and should not do by observing others. When visiting a place of worship, for example, people can learn implicitly that talking is impermissible; when socializing at a bar, they realize that talking is permissible, perhaps even obligatory.

As robots are increasingly deployed in various roles in society—ranging from household robots (e.g., vacuum cleaners, lawn movers) to socially assistive robots (e.g., for elder care or physical therapy)—they must be prepared to follow the many social and moral norms that humans adhere to every day. They must be capable of going beyond pre-programmed norms and, like humans, learn the extant norms by carefully observing other members of society.

We aim at developing computational methods for automatically learning norms from observations. This requires accounting for how norms are cognitively represented, how they are swiftly activated in relevant contexts, and how they can be learned. Learning and applying context-specific norms is difficult. Different learning agents may possess different background knowledge and sensory capabilities, so they may differ in their norm inferences. For example, while one learner might conclude, after several observations, that talking is prohibited in a library setting, another might conclude it is permitted because a group of people were talking at the checkout counter. How can we account for these differences among norm learners and allow distinct agents to communicate and resolve their differences?

In this paper, we extend a recent theoretical proposal [3] as well as experimental and preliminary computational work [4] to introduce novel contributions particularly relevant to the problem of disparate agents learning norms from observation. Specifically, we propose a norm representation scheme (shared by all agents) that introduces a novel deontic modal operator, which equips deontic logic with *context-specificity* and *uncertainty* by relying on a formal framework called Dempster-Shafer theory [5]. We also provide an algorithm that learns norms and honors several critical properties of human norms. Crucially, we suggest that the proposed approach provides the necessary computational infrastructure needed for agents to learn a set of shared norms, as well as to communicate about and address individual differences.

We first provide background on Dempster-Shafer theory, an uncertainty processing framework supporting the novel norm representation. We then present the proposed representation, initial experimental evidence, and the algorithm for learning from observation. Finally, we provide a demonstration of the algorithmic agent's cognitive capability to dynamically learn norms while being exposed to distinct contexts.

II. Dempster-Shafer Theory Background

DS-Theory is a belief-theoretic mathematical framework that allows for combining pieces of uncertain evidence about various events to produce degrees of belief about the events. It has been extensively used in sensor fusion networks, object tracking, and network security. In DS-Theory, a set of elementary events of interest is called *Frame of Discernment* (FoD). The FoD is a finite set of mutually exclusive events $\Theta = \{\theta_1, ..., \theta_N\}$. The power set of Θ is denoted by $2^\Theta = \{A : A \subseteq \Theta\}$. Each set $A \subseteq \Theta$ has a certain weight, or *mass* associated with it. A *Basic Belief Assignment* (BBA) is a mapping $m_\Theta(\cdot) : 2^\Theta \to [0, 1]$ such that $\sum_{A \subseteq \Theta} m_\Theta(A) = 1$

and $m_\Theta(\emptyset) = 0$. The BBA measures the support assigned to the propositions $A \subseteq \Theta$ only. The subsets of A with non-zero mass are referred to as *focal elements* and comprise the set $\mathcal{F}_\Theta$. The triple $\mathcal{E} = \langle \Theta, \mathcal{F}_\Theta, m_\Theta(\cdot) \rangle$ is called the *Body of Evidence* (BoE). For ease of reading, we sometimes omit $\mathcal{F}_\Theta$ when referencing the BoE. Given a BoE $\langle \Theta, \mathcal{F}_\Theta, m_\Theta(\cdot) \rangle$, the *belief* for a set of hypotheses A is $Bel(A) = \sum_{B \subseteq A} m_\Theta(B)$. This belief function captures the total support that can be committed to A without also committing it to the complement A^c of A. The *plausibility* of A is $Pl(A) = 1 - Bel(A^c)$. Thus, $Pl(A)$ corresponds to the total belief that does not contradict A. The $uncertainty$ interval of A is $[Bel(A), Pl(A)]$, which contains the true probability $P(A)$. In the limit case with no uncertainty, we get $Pl(A) = Bel(A) = P(A)$.

DS-Theory extends Bayesian theory in several ways. First, it allows for assigning probabilistic measures to sets of these hypotheses allowing it to consider ignorant and ambiguous information. Second, DS-theory does not require assuming any prior distributions, which is useful when priors are difficult to justify. Third, DS-theoretic uncertainty generally refers to epistemic uncertainty and corresponds to beliefs held by agents about the world.

One recent development in DS-theory is an evidence filtering strategy. It has upgraded Dempster's original rule of evidence combination to accommodate the inertia of available evidence and address some challenges with respect to conflicting evidence [6]. In particular, consider the BoEs $\mathcal{E}_1 = \langle \Theta, \mathcal{F}_1, m_1(\cdot) \rangle$ and $\mathcal{E}_2 = \langle \Theta, \mathcal{F}_2, m_2(\cdot) \rangle$, and a given $A \in \mathcal{F}_2$. The updated belief (from iteration t to $t+1$) $Bel_{t+1} : 2^\Theta \rightarrow [0,1]$ and the updated plausibility $Pl_{t+1} : 2^\Theta \rightarrow [0,1]$ of an arbitrary proposition $B \subseteq \Theta$ are [1]:

$$Bel(B)^{\mathcal{E}_1}_{t+1} = p_1 \cdot Bel(B)^{\mathcal{E}_1}_t + p_2 \cdot Bel(B|A)^{\mathcal{E}_2}_t \qquad (1)$$

$$Pl(B)^{\mathcal{E}_1}_{t+1} = p_1 \cdot Pl(B)^{\mathcal{E}_1}_t + p_2 \cdot Pl(B|A)^{\mathcal{E}_2}_t \qquad (2)$$

where $p_1, p_2 \geq 0, p_1 + p_2 = 1$. The conditionals in the above equations are defined by Fagin-Halpern conditionals, which extend Bayesian conditional notions [7]. That is, for a BoE $\mathcal{E} = \langle \Theta, \mathcal{F}, m(\cdot) \rangle, A \subseteq \Theta$ and an arbitrary $B \subseteq \Theta$, the conditional beliefs and plausibility are given by [2]:

$$Bel(B|A)^{\mathcal{E}} = Bel(A \cap B)^{\mathcal{E}} / \left[Bel(A \cap B)^{\mathcal{E}} + Pl(A \setminus B)^{\mathcal{E}}\right] \qquad (3)$$

$$Pl(B|A)^{\mathcal{E}} = Pl(A \cap B)^{\mathcal{E}} / \left[Pl(A \cap B)^{\mathcal{E}} + Bel(A \setminus B)^{\mathcal{E}}\right] \qquad (4)$$

We build on this development to provide a unified probabilistic norm learning framework, grounded in a belief-theoretic approach.

III. Context-Specific, Belief-Theoretic Norm Representation

The cognitive science literature offers few investigations into the structure and representations of norms. Logical and deontological approaches have introduced formal representations of norm systems [8], [9], [10]. Though such formalizations suggest possible representations, they do not necessarily capture the properties that characterize actual human norm representation and learning.

A. Mathematical Formulation of a Norm System

We define the logical form of norms as follows:

Definition 1: (Context-Specific Norm). A context-specific norm $\mathcal{N}$ is an expression of the form:

$$\mathcal{N} \stackrel{\text{def}}{=} \mathbb{D}_C A \qquad (5)$$

for a formal language $\mathcal{L}$ together with deontic modal operators for obligatory ($\mathbb{O}$), forbidden ($\mathbb{F}$) and permissible ($\mathbb{P}$), respectively (collectively, denoted by $\mathbb{D}$). $C \in \mathcal{L}$ represents a context condition, and $A \in \mathcal{L}$ represents an action or state. The norm expression states that in context C the action or state A is either obligatory, forbidden, or permissible, or not obligatory, forbidden or permissible.

This norm definition build on an approach to normative reasoning and norm representation that some of us have taken previously [3], [8], [9], [4]. But we extend this approach by explicitly accounting for uncertainty about a norm representation as follows:

Definition 2: (Belief-Theoretic Norm). A context-specific belief-theoretic norm is an expression of the form:

$$\mathcal{N} \stackrel{\text{def}}{=} \mathbb{D}_C^{[\alpha,\beta]} A \qquad (6)$$

where $\mathbb{D}_C^{[\alpha,\beta]}$ is an uncertain context-specific deontic operator with $[\alpha, \beta]$ representing a Dempster-Shafer uncertainty interval for the operator, and $0 \leq \alpha \leq \beta \leq 1$. An uncertain deontic operator reduces to a standard deontic operator when there is no uncertainty, i.e., when $[\alpha, \beta] = [1, 1]$

Example 1: Consider an agent deliberating about actions it may or may not perform in a library. This situation can be represented as a Belief-Theoretic Norm System, $\mathcal{T}$, as follows:

$$\begin{aligned}
\mathcal{N}_1 &\stackrel{\text{def}}{=} \mathbb{O}_{Lib}^{[0.9,1]}\ quiet \\
\mathcal{N}_2 &\stackrel{\text{def}}{=} \mathbb{P}_{Lib}^{[0.8,0.95]}\ reading \\
\mathcal{N}_3 &\stackrel{\text{def}}{=} \mathbb{F}_{Lib}^{[0.9,1]}\ yelling \\
\mathcal{N}_4 &\stackrel{\text{def}}{=} \mathbb{O}_{Lib}^{[0,0.3]}\ talking \\
\mathcal{N}_5 &\stackrel{\text{def}}{=} \mathbb{F}_{Lib}^{[0.3,0.6]}\ talking
\end{aligned} \qquad (7)$$

The norms in this example have intuitive semantics. They specify that in the library (i.e., Lib), the agent is obligated ($\mathbb{O}$) to be in a certain state ($quiet$) or is prohibited ($\mathbb{F}$) from performing a certain action ($talking$), each with associated uncertainty intervals. The center of the interval denotes the point estimate of subjective certainty for the believed applicability of the deontic operator, and the width of the interval denotes the level of evidence for that belief. Norm representations $\mathcal{N}_1$, $\mathcal{N}_2$, and $\mathcal{N}_3$ have narrow uncertainty intervals close to 1, indicating strong support for an agent's confident belief (e.g., in the obligation to be quiet in the library). Norm $\mathcal{N}_4$ also has a narrow uncertainty interval but a center close to zero,

[1]We specify the BoE superscript for $Bel()$ and $Pl()$ as needed to be precise, especially when we are combining two distinct BoEs.

[2]The forward slash ("/") represents division, and the backward slash ("\") represents set difference.

indicating strong support for the belief that the operator does not apply. Finally, $\mathcal{N}_5$ has a wider interval and is centered near 0.5, indicating little evidence for either the belief in the norm's applicability or the belief in its inapplicability.

A belief-theoretic norm system, as proposed, enables the clean separation of norms from the evidence supporting them. The evidence may arrive in disparate forms from different sources and through a variety of sensors. The norm system then represents the agent's current belief about the set of extant norms, in light of the available evidence.

An agent may have a vast set of norm representations, but in any given situation, the agent is unlikely to reason with every one of those norms. The agent may instead consider only a subset of the entire system, such as those norms that apply to the specific context in which it finds itself. This idea is captured by a *norm frame*, defined below.

Definition 3: (Norm Frame). A norm frame $\mathcal{N}_k^\Theta$ is a set of k norms in which every norm has the same set of context conditions and the same deontic operator. Thus, in Example 1, norms $\mathcal{N}_3$ and $\mathcal{N}_5$ would constitute a norm frame.

This definition of a norm frame enables us to model an agent's reasoning about behavior in a context-specific manner. Such context specificity constrains and simplifies computation and better captures properties of norm representations in humans, as introduced next.

B. Human Data Supporting the Proposed Representation

We conducted two experiments (labeled "generation" and "detection") on human norm representations. They illustrate some of the cognitive properties of norms and provide support for the proposed formal representation, particularly the context-sensitivity of these norms.

1) Methodology: In the *generation* experiment [11], participants (*n* = 100 recruited from Amazon Mechanical Turk, AMT) viewed four pictures, one after another, that each displayed an ordinary scene (e.g., library, jogging path, board room). While inspecting each picture, participants typed actions that one is either "allowed" to perform in this scene (Permissions), or is "not allowed" to perform (Prohibitions), or is "supposed" to perform (Prescriptions). The resulting verbal responses were analyzed for agreement (i.e., how many people mentioned a given action for a given scene) and context specificity (i.e., whether an action mentioned as permitted in one scene was also mentioned as permitted in another scene). The researchers identified, for each scene, the seven actions most often mentioned as being permitted, and likewise the seven most often mentioned as being prescribed and as being prohibited. These top seven actions within each norm type and each scene were critical elements in the detection experiment.

In the *detection* experiment, participants (*n* = 360 recruited from AMT) viewed the same pictures as in the generation experiment. Their task was to consider each scene and judge 14 actions, one after another, for whether they were permitted (or prescribed, or prohibited) in the particular scene. Of the 14 actions assigned to a given scene and a given norm type (e.g., Library/permitted), 7 were the previously selected top seven actions for that scene and norm type; the other 7 were drawn from top seven actions mentioned in *different* scenes, but under the same norm type. Thus, the latter 7 were just as prevalent as the former 7 but were generated in different contexts.

2) Results: A key result from these studies was that the norms applicable to the various scenes showed remarkable context differentiation. Among the top seven prescribed actions for a given scene, aggregated over eight scenes), 95% were uniquely mentioned in one scene; for permitted actions, this rate was 84%.

IV. Norm Learning

A. Human Norm Learning

When learning social and moral norms, people have to master multiple norm types (permissions, prescriptions, prohibitions), use different learning mechanisms (e.g., observation, instruction), and take input from numerous sources. In this paper, we consider the process of learning permission norms from observation. We put our proposed computational framework to the test using responses from the sampling of community members in the previously described detection experiment.

Consider a person unfamiliar with libraries. Upon entering one for the first time, this person observes people reading and studying. Some are whispering to each other and others are talking loudly, even though there is a sign that says "No talking in the library." The person also notices people at the checkout desk checking out books and subsequently exiting the library, passing another sign that says "Remember to check out." The person observes, briefly, a child running alongside the stacks but then sitting down next to an adult.

How fast and confidently our protagonist learns the norms of a library will depend on the number of people performing each behavior, their appearance, age, and inferred expertise, perhaps evaluative frowns from others, and the salience and meaning of various physical objects. Moreover, there may be many overlapping (and even conflicting) contexts in the library, from the general library setting to the checkout desk to the restrooms and conference rooms. Norms may not apply with equal deontic force in these different settings. While learning the library norms, our learner must dynamically adjust to these contextual variations.

Below we offer a data representation format that incorporates these and other properties of the norm learning process.

B. Data Representation Format of Norm Learning

Consider a set $S = \{s_1, \ldots, s_n\}$ of n evidence sources. For example, s_i could be a patron in the library, a library staff person, or a sign near the entrance. Now, consider a norm frame $\mathcal{N}_k^\Theta$ comprising k norms (out of a larger possible set) that all share the same deontic type (here, permissions) and the same general context precondition (here, library).

Let an endorsement $e_{i,j}$ be the i^{th} data source's endorsement of the j^{th} norm, where $e \in \{0, 1, \epsilon\}$. The value $e_{i,j}$ asigns a truth value, indicating whether the source endorses the norm to be true (1), false (0) or unknown (ϵ). For example, an observation that a person is quiet in the library can be interpreted as showing that this person endorses the norm $\mathcal{N}_1$ to be true in this library context, hence $e_{i,\mathcal{N}_1} = 1$. The set Φ_{s_i} is a source's finite set of endorsements within a given norm frame, such that $|\Phi_{s_i}| = k$.

Informally, for a set of norms in a given context and for a particular source, we can learn about that source's endorsement of each norm; if we also assign a weight (e.g., reliability, expertise) to the source, we form a *data instance.* Multiple data instances (i.e., evidence from multiple sources) form a data set. More formally:

Definition 4: (Data Instance). A data instance $d = \langle \mathcal{N}_k^\Theta, s_i, \Phi_{s_i}, m_{s_i} \rangle$ is a tuple comprising a norm frame $\mathcal{N}_k^\Theta$, a specific source s_i, a set of endorsements Φ_{s_i} provided by that source, and a mass assignment m_{s_i} corresponding to the amount of consideration or reliability placed on source s_i for this instance, including its reliability in detecting context.

Definition 5: (Dataset). A dataset $\mathcal{D}$ is a finite set of n data instances $\{d_1, \ldots, d_n\}$.

With this data representation we can support (a) a range of source types (e.g., signs and symbols, observed behavior, and natural language), (b) different source reliability measures, (c) order effects (updates can be tuned, if necessary, to the order of received data), and (d) missing and imprecise data (we can use ϵ to represent ignorance). The data format also supports cases where the prior probability distributions are unknown or are not justifiable (it does not require any priors), and cases when there are norm dependencies (e.g., correlations between the prohibition to yell and the prohibition to talk). The proposed representation can also be extended to other learning mechanisms, such as trial-and-error and verbal instruction.

C. Automatically Learning Norms from Experimental Data

Next, we will represent the detection data, introduced in Section IIIB, using the data representation format defined in Section IVB. To recap, the detection experiment featured, for each scene, 14 potentially permissible actions, which can be represented as a norm frame $\mathcal{N}_k^\Theta$ with k = 14. Among these 14 actions, 7 had previously been declared as specifically permitted in the given scene, whereas the other 7 had been declared as permitted in different scenes. Each participant (or source s_i) responded with yes (1) or no (0) whether each of the 14 actions was permitted in the given scene. These 0 or 1 responses (and potential missing values, ϵ) formed a set of endorsements Φ_{s_i}, with $|\Phi| = 14$. All participants (sources) were considered equally reliable (i.e., assigned identical m_{s_i} weights), although this is by no means required in other cases.

Now that we have introduced these representations, we can turn to the task of formally defining the *norm learning problem* within our framework. This will then allow us to introduce an algorithm to analyze human norm data and derive a set of norms for a given context. As a reminder, according to Definition 2, any norm has a DS uncertainty interval $[\alpha, \beta]$ associated with it, which reflects the quality and consistency of the evidence for a given norm. Our learning problem becomes a parameter learning problem for estimating the values of the uncertainty interval for each norm in a norm frame:

Definition 6: (Norm Learning Problem). For a norm frame $\mathcal{N}_k^\Theta$ and dataset $\mathcal{D}$, compute the parameters $\alpha_1, \ldots, \alpha_k, \beta_1, \ldots, \beta_k$ of that norm frame.

In the human data set, each participant produces a vector of yes (1) or no (0) responses across the 14 actions for a given norm. Each of these vectors is a data instance d that represents a potential arrangement of true and false values for the given norm in a norm frame. Setting aside the possibility that $e_{i,j} = \epsilon$, each data instance thus provides a k-length string of 1s and 0s. The norm learning algorithm represents each string as a hypothesis in a set of hypotheses (termed Frame of Discernment in Dempster-Shafer theory) and assigns uncertainty parameters to each norm, updating those values as it considers each new data instance (i.e., each participant who contributes to the sample of norm endorsements in the community). **Algorithm 1**, displayed below, implements this form of norm learning from a human dataset.

Algorithm 1 getParameters($\mathcal{D}, \mathcal{N}_k^\Theta$)

1: $\mathcal{D} = \{d_1, \ldots, d_n\}$: Dataset containing n data instances
2: $\mathcal{N}_k^\Theta$: An unspecified norm frame containing k norms $\mathcal{N}$
3: Initialize DS Frame $\Theta = \{\theta_1, \ldots, \theta_{2^k}\}$
4: $m(\Theta) = 1$
5: **for all** $d \in \mathcal{D}$ **do**
6: **for all** $\mathcal{N} \in \mathcal{N}_k^\Theta$ **do**
7: Set learning parameters p_1 and p_2
8: $Bel(\mathcal{N}|d) = \frac{Bel(\mathcal{N} \cap d)}{Bel(\mathcal{N} \cap d) + Pl(d \setminus \mathcal{N})}$
9: $Pl(\mathcal{N}|d) = \frac{Pl(\mathcal{N} \cap d)}{Pl(\mathcal{N} \cap d) + Bel(d \setminus \mathcal{N})}$
10: $Bel(\mathcal{N})_{new} = p_1 \cdot Bel(\mathcal{N})_{prev} + p_2 \cdot Bel(\mathcal{N}|d)$
11: $Pl(\mathcal{N})_{new} = p_1 \cdot Pl(\mathcal{N})_{prev} + p_2 \cdot Pl(\mathcal{N}|d)$
12: **end for**
13: Set frame Θ with $Bel(\mathcal{N})_{new}$ and $Pl(\mathcal{N})_{new}$
14: **end for**
15: **for all** $\mathcal{N} \in \mathcal{N}_k^\Theta$ **do**
16: $\alpha_{\mathcal{N}} \leftarrow Bel(\mathcal{N})$, $\beta_{\mathcal{N}} \leftarrow Pl(\mathcal{N})$
17: **end for**
18: **return** $\alpha_1, \ldots, \alpha_k, \beta_1, \ldots, \beta_k$

The algorithm receives as input a dataset containing n data instances. It iterates through each data instance (line 5), and for each instance, it iterates through each norm in the norm frame (line 6). During each iteration, the algorithm sets hyper-parameters p_1 and p_2 (line 7). The hyper-parameters specify the amount of weight the algorithm will place on previous learned knowledge (p_1) and on the new data instance under consideration in this iteration (p_2). These parameters are then used to compute the Fagin-Halpern conditional beliefs and plausibilities for the norm given the current data instance (lines 8,9). The conditional beliefs and plausibilities are then used to compute the updated beliefs and plausibilities for each norm (lines 10,11). Finally, the algorithm outputs these updated beliefs and plausibilities for each of the norms (lines 15-18).

The result is a set of norms with uncertainty intervals. The width of the uncertainty interval represents the amount of evidence for the inferred norm (varying, for example, as a function of the number of respondents in the human data set);

the center of the interval corresponds to the estimated norm endorsement by the human respondents (approaching $0 \rightarrow$ not permitted; approaching $1 \rightarrow$ permitted).

D. Evaluation: Dynamic Context Shifting

We selected two smaller norm frames to evaluate the algorithm: six permission norms each for the contexts of library and boardroom. We constructed norm frames such that four permitted actions (reading, talking, walking, and listening) were the same in each context but with different endorsement rates (e.g., walking was hardly permissible in the boardroom but permissible in the library). If the algorithm captures the context-specificity of norms it would have to track the norm endorsements for any given action not in general but conditional on the specific context. We also included two permitted actions for which we had data for one context but not for the other: using computers in the library and drinking in the boardroom. This allowed us to evaluate the algorithm's ability to handle ignorance and incomplete information.

The algorithm also had to track how these norms are learned when observations are made in dynamic situations involving changing contexts and general contextual uncertainty. In an ideal learning scenario, the agent would obtain a dataset containing a collection of observations from a single context. The generation experiment described earlier provided such an ideal dataset. However, learning in the real world is far less perfect; data are often obtained in a streaming, unfolding manner through a series of observations made during a certain time window. Observations are made in context, the identity of which might be uncertain and may even change over time.

Consider the example of a norm-learning agent moving through a library between the reception, stacks, and through various conference rooms. At any given moment, the agent may not be entirely certain if it is in one context or another. As it approaches a boardroom, near the threshold, the agent may not be sure if it is within the confines of the boardroom context or within the confines of the general library context. Normative behavior is often learned in this messy manner through observations in changing and uncertain contexts.

Moreover, this contextual uncertainty can influence the agent's normative beliefs themselves, which in turn can result in deviations in normative behavior. Two identical agents who observe the same sequence of actions but differ in their reliability of identifying the context they are in will acquire a different set of norms. The algorithm must be able to track these variations in natural learning and should account for differences in normative beliefs between different agents.

For our experiment, we consider two agents (Agent 1 and Agent 2) that learn norms by observing actions in a library and a boardroom. For simplicity, we stipulate that at any given moment an agent can either be in the general library context or in a boardroom context. Also, we stipulate that the agents are moving from the general library area into a boardroom. Thus, they first observe actions in a library and then actions in a boardroom. However, the agents differ in their ability to accurately detect the context they are in.

Agent 1 has a reliable context detector and can, with complete certainty, identify its current context. Agent 2 has a more unreliable context detector, at least at the threshold between the general library area and a boardroom. Thus, Agent 2 is initially certain that it is in the library, but as it moves toward the boardroom it becomes uncertain about what context it occupies. Once it is completely in the boardroom it is again certain of its context.

We expect that the learning algorithm can capture not only the context dependency of the norms but also the agents' different normative beliefs (as a result of their different context detection abilities). We further expect that an agent that is more unsure of its context is also less certain about the applicability (truth or falsity) of a given norm. Moreover, we expect that because the agent is unsure of its context at the library-boardroom threshold, it will attribute observations during this time to both contexts (albeit to different degrees), thereby generally acquiring more data instances for norms in each context. This would then have the effect of strengthening the agent's belief in the norm that it learns, allowing for a narrowing of the uncertainty interval.

Figure 1 illustrates the success of the learning algorithm in both capturing the context sensitivity of norms and the differential effect of learning norms in different situations. We display four plots each showing single runs of the algorithm for whether the action of "talking" is permissible. The results show a wider uncertainty interval at first, but narrowing with accumulated data instances, which is consistent with a hypothetical agent roaming the library and then entering a boardroom. We also performed these single runs for the remaining 5 actions over both contexts for both agents. The results at the end of these runs are shown in Table I. Again, the uncertainty intervals $[\alpha, \beta]$ represent the degree of support for the rule in the observations (α) and the total belief that does not contradict the rule (β).

As predicted, Agent 1's learning of the library norm (Figure 1, top left) proceeds by converging to an optimal estimate of the community's norm endorsement but then holds steady once the agent has left the library context and enters the boardroom context. Conversely, this agent's learning of the boardroom norm (bottom left) begins to converge only once the agent has entered the boardroom. The algorithm approaches the descriptive statistics from the experimental data (which it was not given) but maintains a level of uncertainty that reflects the imperfect agreement in the human data.

In the case of Agent 2 learning the library norm (top right), because it is uncertain about the context in the middle of the run (positive sloping part of the red-dotted line), the agent continues to adapt its learning, finally settling on an interval towards the end of the run, which is a later convergence than that of Agent 1. Conversely, the agent's learning of the boardroom norm (bottom right) begins converging earlier than Agent 1. The result of the uncertainty in context is at the end of the run, where Agent 2 settles on an uncertainty interval that does not include the descriptive mean in the data. As predicted, this deviation suggests that Agent 2's learning at the

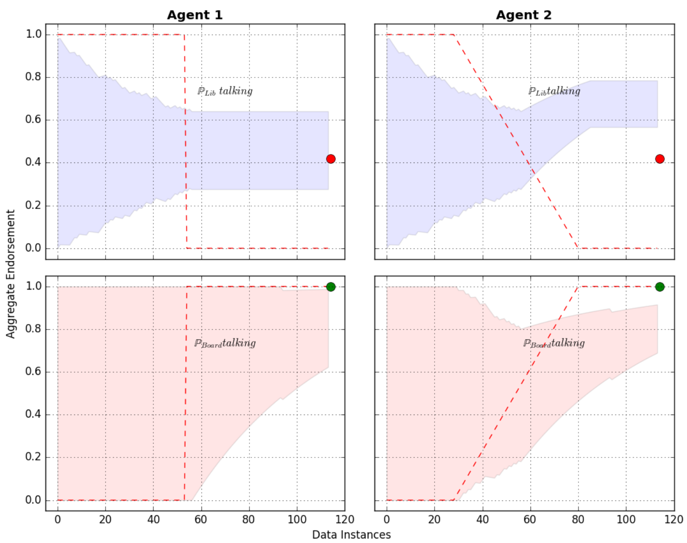

Fig. 1. Single runs of two Agents moving from a general library context towards a boardroom context. The red-dotted line shows the agent's certainty about its context as it moves. The narrowing shaded regions indicate converging uncertainty intervals as new data instances are processed. Colored dots represent the mean norm endorsements by experimental participants —the proportion of participants who answered yes to the question: "Is this action permitted here?" The algorithm displays flexibility in learning normative rules from observation for various agents with different detection capabilities.

library-boardroom threshold was influenced by both contexts, thereby increasing both the truth of the norm in the library and the falsity of the norm in the boardroom. Moreover, because threshold data is used in both contexts, the number of data points considered are increased, providing a tighter interval.

Agents equipped with the proposed learning methodology can not only dynamically learn from observations in evolving environments but can begin addressing mutual differences in their background knowledge and sensory capabilities. For example, an agent (such as Agent 2) can compare its learned uncertainty interval with human consensus statistics; that way it can either correct its beliefs or refine its context detection accuracy by attending to aspects of the environment it previously overlooked. Agents can also directly compare and contrast uncertainty intervals. For example, agents can agree if their confidence intervals overlap, or if the center of their intervals are close to each other. That said, being able to perform this sort of introspection would additionally require the agents to be aware of their differences, a challenge that is beyond the scope of this paper but the subject of future work.

TABLE I
UNCERTAINTY INTERVALS FOR AGENTS 1 AND 2.

Norm	Consensus	Agent 1	Agent 2
$\mathbb{P}_{Lib}^{[\alpha,\beta]}$ *reading*	1.0	[0.63, 1.0]	[0.75, 0.97]
$\mathbb{P}_{Board}^{[\alpha,\beta]}$ *reading*	0.93	[0.59, 0.95]	[0.73, 0.95]
$\mathbb{P}_{Lib}^{[\alpha,\beta]}$ *listening*	0.96	[0.61, 0.98]	[0.77, 0.99]
$\mathbb{P}_{Board}^{[\alpha,\beta]}$ *listening*	1.0	[0.63, 1.0]	[0.76, 0.99]
$\mathbb{P}_{Lib}^{[\alpha,\beta]}$ *talking*	0.42	[0.27, 0.64]	[0.57, 0.78]
$\mathbb{P}_{Board}^{[\alpha,\beta]}$ *talking*	1.0	[0.62, 1.0]	[0.69, 0.91]
$\mathbb{P}_{Lib}^{[\alpha,\beta]}$ *walking*	0.91	[0.57, 0.94]	[0.45, 0.67]
$\mathbb{P}_{Board}^{[\alpha,\beta]}$ *walking*	0.32	[0.21, 0.57]	[0.33, 0.55]
$\mathbb{P}_{Lib}^{[\alpha,\beta]}$ *usingComputers*	0.95	[0.60, 0.96]	[0.36, 0.98]
$\mathbb{P}_{Board}^{[\alpha,\beta]}$ *usingComputers*	-	[0, 1]	[0.13, 0.99]
$\mathbb{P}_{Lib}^{[\alpha,\beta]}$ *drinking*	-	[0, 1]	[0.29, 0.89]
$\mathbb{P}_{Board}^{[\alpha,\beta]}$ *drinking*	0.72	[0.44, 0.81]	[0.44, 0.81]

V. CONCLUSION

This paper introduces a formal representation of norms and a learning algorithm, with uncertain deontic operators within the framework of Dempster-Shafer theory. The representation captures the context specificity of norms which has been demonstrated as a central property of human norm representations. Our generalizable data format allows artificial agents to learn norms from different sources in varying contexts using disparate sensors. The proposed approach can also be highly useful to cognitive infocommunication [12] between humans and artificial agents. In particular, the shared norm representation format will make interactions with machines more intuitive for humans, and they may avail themselves of the machines' sensory channels, which could enhance human norm perception and learning [13].

VI. ACKNOWLEDGMENT

This research was funded in part by grants from the Office of Naval Research (ONR), N00014-16-1-2278, and from the Defense Advanced Research Projects Agency (DARPA), SIMPLEX 14-46-FP-097. The opinions expressed here are our own and do not necessarily reflect the views of ONR or DARPA.

REFERENCES

[1] C. Bicchieri, *The grammar of society: The nature and dynamics of social norms*. New York, NY: Cambridge University Press, 2006.
[2] E. Ullmann-Margalit, *The emergence of norms*. Oxford: Clarendon Press, 1977.
[3] B. Malle, M. Scheutz, and J. Austerweil, "Networks of social and moral norms in human and robot agents," in *A World with Robots*. Springer, 2017, pp. 3–17.
[4] V. Sarathy, M. Scheutz, J. Austerweil, Y. Kenett, M. Allaham, and B. Malle, "Mental representations and computational modeling of context-specific human norm systems," in *Proceedings of Cognitive Science*, 2017.
[5] G. Shafer, *A Mathematical Theory of Evidence*. Princeton University Press, 1976.
[6] D. A. Dewasurendra, P. H. Bauer, and K. Premaratne, "Evidence filtering," *IEEE Transactions on Signal Processing*, vol. 55, no. 12, pp. 5796–5805, 2007.
[7] R. Fagin and J. Y. Halpern, "A new approach to updating beliefs," *arXiv preprint arXiv:1304.1119*, 2013.
[8] S. Bringsjord, K. Arkoudas, and P. Bello, "Toward a general logicist methodology for engineering ethically correct robots," *IEEE Intelligent Systems*, vol. 21, no. 4, pp. 38–44, 2006.
[9] M. Scheutz and B. F. Malle, ""Think and do the right thing"—A plea for morally competent autonomous robots," in *IEEE International Symposium on Ethics in Science, Technology and Engineering*. IEEE, 2014, pp. 1–4.
[10] L. M. Pereira and A. Saptawijaya, "Modelling morality with prospective logic," *International Journal of Reasoning-based Intelligent Systems*, vol. 1, no. 3-4, pp. 209–221, 2009.
[11] Y. N. Kenett, M. M. Allaham, J. L. Austerweil, and B. F. Malle, "The norm fluency task: Unveiling the properties of norm representation," in *Poster presented at the 57th Annual Meeting of the Psychonomic Society, Boston, MA*, Nov. 2016.
[12] P. Baranyi and A. Csapo, "Cognitive infocommunications: coginfocom," in *11th International Symposium on Computational Intelligence and Informatics (CINTI)*. IEEE, 2010, pp. 141–146.
[13] ——, "Definition and synergies of cognitive infocommunications," *Acta Polytechnica Hungarica*, vol. 9, no. 1, pp. 67–83, 2012.