

Exploiting moral values to choose the right norms

Maite Lopez-Sanchez
Math & Comp Science dept.
Universitat de Barcelona
maite_lopez@ub.edu

Marc Serramia, Juan A. Rodriguez-Aguilar
Artificial Intelligence Research Institute (IIIA-CSIC)
Campus UAB. Bellaterra, Spain
{marcserr, jar}@iiia.csic.es

Javier Morales, Michael Wooldridge
Dept. of Computer Science
University of Oxford, UK
{javier.morales, mjw}@cs.ox.ac.uk

Carlos Ansotegui
Universitat de Lleida, Spain
carlos@diei.udl.cat

Abstract

Norms constitute regulative mechanisms extensively enacted in groups, organisations, and societies. However, ‘choosing the right norms to establish’ constitutes an open problem that requires the consideration of a number of constraints (such as norm relations) and preference criteria (e.g. over involved moral values). This paper advances the state of the art in the Normative Multiagent Systems literature by formally defining this problem and by proposing its encoding as a linear program so that it can be automatically solved.

Introduction

Norms are coordination mechanisms that have been extensively used to regulate the interactions between individuals within groups, organisations, and societies. In traditional societies, most norms have been long established. Nevertheless, new-born or highly dynamic social groups require an active decision on which norms should be enacted, and this decision is by no means straightforward. The reasons are twofold. First, the norms to choose from may not be independent (i.e. they can be related to each other). Second, different preference criteria (such as, for example, deployment costs or involved moral values) may be applied when choosing the right norms to enact.

In the context of social dynamism –and its corresponding effects on regulations–, e-participation and e-governance ICT systems (Weerakkody and Reddick 2012; DeTar 2013; Jackson and Kuehn 2016) are gaining relevance. Accordingly, electronic participation initiatives are being launched in a large variety of settings. These can range from top public governmental institutions at country level (such as France (Par 2017) or New Zealand (New 2016a; New 2016b)), down to public local institutions (such as Reykjavik (Rey 2017), Madrid (Mad 2017), or Barcelona (BCN 2017) municipalities), or small social groups (Jackson and Kuehn 2016). Overall, these initiatives open to their participants the possibility of providing their opinion on individual proposals but, to the best of our knowledge, none of them tackle the problem of choosing the “best” set of proposals. Moreover, they mainly focus on citizen support as their reference criterion. One may argue that citizen support encompasses other

criteria, but it does not explicit them, and this has a detrimental effect over transparency in the consultation process.

Within the Artificial Intelligence research area of Normative Multiagent Systems (Andrighetto et al. 2013), this paper advances the state of the art by modeling a series of decision-making problems that social communities confront when choosing the set of policies to establish. In order to do so, we first identify three different norm relationships – namely, generalisation, exclusivity, and substitutability – and we then consider norm representation power, cost, and associated moral values as alternative explicit preference criteria. Thereafter, we show that these decision-making problems can be encoded as (optimisation) linear programs, and hence solved with the aid of state-of-the-art solvers. The consideration of moral values associated to norms as an additional criterion also constitutes a relevant contribution since, although values have been studied in argumentation (Bench-Capon and Atkinson 2009; Modgil 2006) and in multi-agent institutions (Kohler et al. 2014), we are not aware of other authors considering the values that norms support.

The paper is structured as follows. First, next section introduces norms and related concepts. Then, subsequent sections characterise our different optimisation problems. Particularly, fourth section introduces moral values into our multi-objective optimisation problem. Finally, last section is devoted to conclusions and future work.

Norms and their relations

Within the area of normative multiagent systems (Andrighetto et al. 2013), problems such as norm synthesis (Shoham and Tennenholtz 1995; Agotnes and Wooldridge 2010), norm emergence (Griffiths and Luck 2010; Villatoro, Sabater-Mir, and Sen 2011), or norm learning (Savarimuthu et al. 2013; Campos et al. 2013; Riveret et al. 2014) have been widely studied. Although all these approaches consider norms as coordination mechanisms, they differ in their formal definition. Our notion of norm is based on a simplification of the one in (López, Luck, and d’Inverno 2002). Thus, here we formally consider a norm n_i as a pair $\theta(\rho, ac)$, where: θ is a deontic operator (prohibition, permission, or obligation); ρ is a description of the addressee entity, namely, the agent required to comply with the norm; and ac is an action –from a set of actions– that entities can perform in a specific domain. Example 1 illustrates this definition.

As for norm relations, they have also been previously studied in the literature. Thus, for example, Grossi and Dignum (Grossi and Dignum 2005) study the relation between abstract and concrete norms, whereas Kollingbaum, Vasconcelos et al. (Kollingbaum et al. 2007; Vasconcelos, Kollingbaum, and Norman 2009) focus on norm conflicts—and solve them based on first-order unification and constraint solving techniques. In this paper, we borrow two of the relations from (Morales et al. 2015) and define three norm relationships (namely, *exclusivity*, *substitutability*, and *generalisation*). Informally, it is considered that: two norms are mutually exclusive when they are incompatible; two norms are substitutable if they are interchangeable; and a norm is more general than another one when it subsumes its regulation (i.e., its regulation scope is wider).

Specifically, considering N as a non-empty set of norms, the *exclusivity* relation is a binary relation $R_x \subseteq N \times N$. If $(n_i, n_j) \in R_x$ we say that n_i, n_j are incompatible or mutually exclusive. R_x is an irreflexive, symmetric, and intransitive relation.

The *substitutability* relation is a binary relation $R_s \subseteq N \times N$. If $(n_i, n_j) \in R_s$, we say that norms n_i, n_j are interchangeable or substitutable. Based on substitutability relationships, we introduce the notion of *substitution chain* as follows. Given two norms, $n_i, n_k \in N$, we say that n_k is *connected by substitutabilities* to n_i if there is a non-empty subset of norms $\{n_1, \dots, n_p\} \subseteq N$ such that $(n_1, n_2), \dots, (n_{p-1}, n_p) \in R_s$, $n_1 = n_i$, and $n_p = n_k$. Henceforth, a new relationship $\mathcal{S} \subseteq N \times N$ will contain the pairs of norms that are connected by substitutabilities. In particular, notice that if $(n_i, n_j) \in R_s$, then $(n_i, n_j) \in \mathcal{S}$. R_s is an irreflexive, symmetric, and transitive relation.

Finally, the *direct generalisation* relation is a binary relation $R_g \subseteq N \times N$. If $(n_i, n_j) \in R_g$, we say that n_i is more general than n_j , or directly, it generalises n_j . Notice that R_g is irreflexive, anti-symmetric, and intransitive (this is so because if $(n_i, n_j) \in R_g$, $\nexists n_k \in N$ s.t. $(n_i, n_k), (n_k, n_j) \in R_g$). Transitivity is captured through the notion of indirect generalisation and the so-called *ancestors*. Given two norms, $n_k, n_i \in N$, we say that n_k is an ancestor of n_i if there is a subset of norms $\{n_1, \dots, n_p\} \subseteq N$ such that $(n_1, n_2), \dots, (n_{p-1}, n_p) \in R_g$, $n_1 = n_k$, and $n_p = n_i$. Henceforth, given a norm $n_i \in N$, we will note its ancestors as $\mathcal{A}(n_i)$.

Previous concepts allows us to define the so-called norm net.

Def. 1 A norm net is a pair $NN = \langle N, R \rangle$, where N stands for a set of norms and $R = \{R_g, R_x, R_s\}$ contains generalisation, exclusivity and substitutability relationships over the norms in N . The relationships in R are mutually exclusive, namely $R_g \cap R_x = \emptyset$, $R_g \cap R_s = \emptyset$, and $R_x \cap R_s = \emptyset$.

Although relations are mutually exclusive (for instance, two norms cannot be exclusive and substitutable simultaneously), it is worth noticing that exclusivity relations propagate along generalisation relationships. Thus, for example, if we consider four different norms $n_i, n_j, n_k, n_x \in N$, and we have that n_i is more general than n_j , which in turn is more general than n_k (i.e., $(n_i, n_j), (n_j, n_k) \in R_g$), then, if

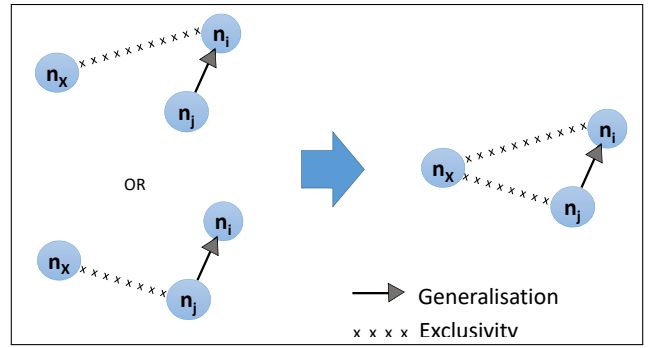


Figure 1: Exclusivity propagation examples.

one of these three norms are exclusive with norm n_x (that is, if, for instance $(n_x, n_j) \in R_x$), then all three norms along their generalisation relationships are also exclusive with n_x (i.e., $(n_x, n_i), (n_x, n_k) \in R_x$). Top part of Figure 1 illustrates how an exclusivity relation between norms n_x and n_i ($(n_x, n_i) \in R_x$, depicted with an “x dotted line”) is propagated downwards to n_j , the norm n_i generalises ($(n_i, n_j) \in R_g$), so that $(n_x, n_j) \in R_x$. Notice that Figure 1 draws generalisation relationships with an arrow line with the arrowhead pointing towards the general norm n_i . Similarly, as bottom part of Figure 1 shows, if it is the case that norm n_x is exclusive with n_j , then this exclusivity relation $(n_x, n_j) \in R_x$ is also propagated upwards in the generalisation relationship $(n_x, n_i) \in R_x$.

Furthermore, given a norm net $NN = \langle N, R \rangle$, we will refer to any subset of the norms in N as a *norm system*. The challenge then lies in selecting the right norm system out of a norm net. In general, we will be interested in norm systems incorporating as many norms as possible but excluding overlapping nor conflicting norms.

Considering that exclusivity relationships capture conflicts between norms whereas substitutability and generalisation relationships capture *redundancy* or overlap (or in the case of generalisation, subsumption), the following characterisation of norm systems naturally follows.

Def. 2 Given a norm net $NN = \langle N, R \rangle$, we say that a norm system $\Omega \subseteq N$ is *conflict-free* iff for each $n_i, n_j \in \Omega$, $(n_i, n_j) \notin R_x$.

Def. 3 Given a norm net $NN = \langle N, R \rangle$, we say that a norm system $\Omega \subseteq N$ is *non-redundant* iff for each $n_i, n_j \in \Omega$: (i) $(n_i, n_j) \notin R_g$ and $n_j \notin \mathcal{A}(n_i)$; and (ii) $(n_i, n_j) \notin \mathcal{S}$

Def. 4 Given a norm net $NN = \langle N, R \rangle$, we say that a norm system $\Omega \subseteq N$ is *sound* iff it is both conflict-free and non-redundant.

Thus, we aim at finding sound norm systems that satisfy certain criteria. Next section is devoted to further elaborate on that. Before that, though, it may be helpful to illustrate a norm net with an example.

Example 1 Figure 2 illustrates an example of a norm net that includes some norms (rules) of border control at an international airport. Norms are depicted as circles labeled

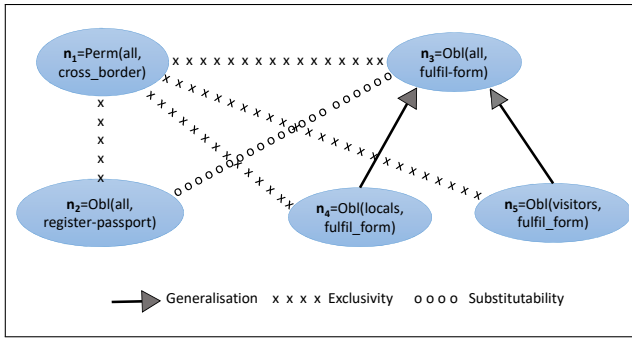


Figure 2: Norm net example: rules of border control at an international airport.

as n_1, \dots, n_5 respectively. In particular, they are defined as follows:

- n_1 : *Permission*(all_passengers, cross_border)
- n_2 : *Obligation*(all_passengers, register_passport)
- n_3 : *Obligation*(all_passengers, fulfil_form)
- n_4 : *Obligation*(locals, fulfil_form)
- n_5 : *Obligation*(visitors, fulfil_form)

Norm n_1 rules free movement of passengers, allowing all passengers to cross the border without any additional action. On the other hand, norm n_2 requires all passengers to register their passport, and there is still a third rule n_3 that requires them to fulfil a form asking for passport information such as passport number, holder's name or address.

Regarding exclusivity relationships, first and second norms are exclusive ($(n_1, n_2) \in R_x$) because it is not possible to ask passengers to perform an action (in this case, register their passport) and allow them to go ahead and simply cross the border. The same reasoning applies for first and third norms, and thus, n_1 and n_3 are also exclusive ($(n_1, n_3) \in R_x$).

Additionally, there is a substitutability relationship between second and third norms ($(n_2, n_3) \in R_x$), since it is possible to control which passengers are actually crossing the border by registering their password or by asking them to fulfil a form. Passengers may even have to abide by both norms, since they could actually do both things despite of its redundancy. Figure 2 shows this substitutability relationship with an “o dotted” line.

Having a closer look to norm n_3 , we can see that asking all passengers to fulfil a form is a generalisation of two other norms: n_4 , which requires local passengers to fulfil a form; and n_5 , which requires foreign passengers (visitors) to do the same. Formally, we have $(n_3, n_4) \in R_g$ and $(n_3, n_5) \in R_g$.

Finally, exclusive relationship between n_1 and n_3 ($(n_1, n_3) \in R_x$) propagates down to norms n_4 and n_5 , so that their exclusivity with n_1 is also taken into account ($(n_1, n_4), (n_1, n_5) \in R_x$).

Norm system multi-objective optimisation

As previously mentioned, this paper tackles the problem of choosing the right set of norms by considering certain

constraints and preference criteria. Constraints are set to guarantee sound norm systems whereas preference criteria are specified as functions to optimise. This section details how soundness constraints and two specific criteria –namely, norm representation and norm deployment cost– can be encoded as a linear program that can in turn be solved with state-of-the-art solvers such as CPLEX or Gurobi.

On the one hand, norm representation is a preference criteria that encapsulates the purpose of incorporating as many norms as possible out of those proposed in a norm net, since they are all considered to be acceptable. Thus, informally, we will aim at the norm system that *represents* the largest number of norms in the norm net. On the other hand, we cannot ignore the fact that norm deployment has associated costs and that the expenses derived from imposing norms should be bounded by the available budget. Norm costs may represent monetary expenses derived from regulatory processes –such as norm establishment or norm enforcement– as well as non-monetary aspects –such as social implications or political correctness– as long as they can be somehow quantified. In any case, we intend to minimise incurred costs while maximizing norm representation.

Formally, we consider a representation power function, $r : N \rightarrow \mathbb{R}$ to be a linear function that computes a norm's representation power: a real value that encapsulates the fact that a norm cannot only represent itself but also all the norms it generalises. Besides linearity, the only condition that we impose on r is that $r(n_i) \leq r(n_j)$ for each $n_j \in \mathcal{A}(n_i)$. Hence, the representation power of a normative system Ω can be readily obtained by adding the representation power of its norms, namely $\rho(\Omega) = \sum_{n_i \in \Omega} r(n_i)$.

Similarly, we assume that the cost of a norm system can be obtained by adding the individual costs of its norms, namely $cost(\Omega) = \sum_{n_i \in \Omega} c(n_i)$, where $c(n_i)$ stands for the cost of norm n_i . Furthermore, we make the (reasonable) assumption that $cost(\Omega)$ is bounded by a maximum budget b that is available to cover the expenses of imposing those norms in the resulting norm system.

From that, we can cast the decision problem as the following multi-objective optimisation problem.

Problem 1 Given a norm net $NN = \langle N, R \rangle$ ¹, a representation power function r , and a fixed budget b , the maximum norm system problem with limited budget (MNSPLB) is the problem of finding a sound norm system $\Omega \subseteq N$ with maximum representation power (i.e., there is no other norm system $\Omega' \subseteq N$ such that $\rho(\Omega') > \rho(\Omega)$) and minimum cost (i.e., $\nexists \Omega' \subseteq N$ s.t. $cost(\Omega') < cost(\Omega)$) limited by some non-negative budget b .

Lemma 1 The complexity of the maximum norm system problem is at least NP-Hard.

Proof. 1 The proof goes trivially by reduction of the maximum independent set problem, which is known to be an NP-Hard optimisation problem (Karp 1972), to the maximum norm system problem. Consider that we want to find the maximum independent set of a graph $G = (V, E)$. Now

¹Notice that we assume knowledge about candidate norms to enact and the relationships between such norms.

say that each vertex in V stands for a norm and each edge in E stands for an exclusivity relationship in R_x . From this follows that finding the maximum independent set of G amounts to solving the maximum norm set problem on the norm net $\langle V, \{R_x\} \rangle$, where the representation power and cost functions are defined as $r(v) = 1$, $cost(v) = 1$ for each $v \in V$ and $b = |V|$.

Next, we show how to solve the MNSPLB by encoding the optimisation problem as a linear program. Thus, consider a norm net $NN = \langle N, R \rangle$, and a set of binary decision variables $\{x_1, \dots, x_{|N|}\}$, where each x_i encodes the decision of whether norm n_i is selected (taking value 1) for a norm system or not (taking value 0). Thus, solving the MNSPLB amounts to solving the following:

$$\begin{aligned} \max \sum_{i=1}^{|N|} x_i \cdot r(n_i) \\ \min \sum_{i=1}^{|N|} x_i \cdot c(n_i) \end{aligned} \quad (1)$$

subject to a number of constraints (including soundness restrictions).

First, *exclusivity constraints* prevent that two mutually exclusive (incompatible) norms are jointly selected to be part of a norm system. Thus, the following constraints must hold:

$$x_i + x_j \leq 1 \quad \text{for each } (n_i, n_j) \in R_x \quad (2)$$

Second, *substitutability constraints* avoid that interchangeable norms are simultaneously selected. This amounts to enforcing that any pair of norms that are connected by substitutabilities cannot be simultaneously selected, namely:

$$x_i + x_j \leq 1 \quad \text{for each } (n_i, n_j) \in \mathcal{S} \quad (3)$$

Third, *generalisation constraints* avoid redundancy by imposing that a norm cannot be selected together with any of the norms that it directly generalises. Given a norm n_i , the set of norms generalised by n_i is defined as $Children(n_i) = \{n_j | (n_i, n_j) \in R_g\}$. Then, formally, the following constraints must hold:

$$x_i + x_j \leq 1 \quad n_j \in Children(n_i) \quad 1 \leq i \leq |N| \quad (4)$$

Moreover, all the children of a norm cannot be simultaneously selected. Formally:

$$\sum_{n_j \in Children(n_i)} x_j < |Children(n_i)| \quad 1 \leq i \leq |N| \quad (5)$$

Additionally, a norm cannot be simultaneously selected together with any of its ancestors, namely:

$$x_i + x_k \leq 1 \quad n_k \in \mathcal{A}(n_i) \quad 1 \leq i \leq |N| \quad (6)$$

Fourth, we must also consider the binary constraints corresponding to the norm decision variables, namely:

$$x_i \in \{0, 1\} \quad 1 \leq i \leq |N| \quad (7)$$

Finally, a further constraint ensures that the cost of the norm system does not exceed the limited budget $b \geq 0$:

$$\sum_{i=1}^{|N|} c(n_i) \cdot x_i \leq b, \quad (8)$$

The linear program encoding the MNSPLB requires $|N|$ binary decision variables, $2 \cdot |R_g| + |R_x| + |\mathcal{S}|$ pairwise constraints (equations 2,3,4, and 6), and $|P(R_g)|$ inequality constraints (equation 5), where $P(R_g) = \{n_i | (n_i, n_j) \in R_g\}$. Hence, the number of constraints is linear with the number of norm relationships in a norm net.

Moral values in norm decision making

Choosing the right set of norms may also imply the explicit consideration of a third preference criteria encapsulating ethical aspects that allow to discern ‘right’ from ‘wrong’. In this manner, if we assume that the society has overall preferences over moral values, and if we are able to assess the moral values promoted by the norms in a given norm network, then we can aim at maximising moral value support when computing the optimal norm system. The principle that we adhere to is: *the more preferred the values supported by a norm system, the more preferred that norm system is*. Thus, this section is devoted to extend the multi-objective decision-making problem introduced in previous section to account for the *moral values* supported by norms.

To do so, we must formally define some concepts. We start by introducing moral values. Inspired on Ethics, where moral values are often considered to be moral standards to distinguish between right and wrong (good or evil), we adapt some value-related definitions by Bench-Capon et al. (Bench-Capon and Atkinson 2009) and consider V as a non-empty set of moral values in a society. We will assume that moral values are unrelated objects and that there is a total order (no ties) \succ over the moral values in V that reflects the global social preferences over them. Without loss of generality, we can assume that $v_1 \succ v_2 \succ \dots \succ v_{|V|}$. From this, we obtain the moral value(s) supported by each norm by means of a function $val : \mathcal{N} \rightarrow \mathcal{P}(V)$. Thus, $val(n_i)$ stands for the set of moral values promoted by norm n_i .

Now we introduce a utility function that will allow us to capture the total ordering over moral values. The utility of a moral value can be calculated as:

$$u(v_i) = 1 + \sum_{k=i+1}^{|V|} u(v_k) \quad (9)$$

where $1 \leq i \leq |V|$. From this definition, it is clear that $u(v_i) > u(v_j) \Leftrightarrow v_i \succ v_j$ and we can readily calculate the *value support* of a norm n_i by adding the utility of the moral values supported by the norm as follows:

$$u_n(n_i) = \sum_{v \in val(n_i)} u(v) \quad (10)$$

This allows us to compute the *value support* for a given norm system $\Omega \subseteq N$ by adding the utility of the moral values supported by each one of its norms as:

$$u_N(\Omega) = \sum_{n \in \Omega} u_n(n) \quad (11)$$

Notice that utility function u_N allows us to lift the preferences defined as a linear order over single moral values to a preference relation over sets of norms. Thus, we will say that $\Omega \succ \Omega' \Leftrightarrow u_N(\Omega) > u_N(\Omega')$. Interestingly, the lifting of preferences provided by the u_N utility function satisfies two desirable properties: (i) *responsiveness* (Barbera, Hammond, and Seidl 2004); and (ii) *monotonicity*. Informally, responsiveness (also called pairwise-dominance), states that if in a norm system $\{n_2, n_3\}$, norm n_3 is replaced by a *better* (supporting more preferred moral values) norm, e.g. n_1 , then $\{n_2, n_1\}$ makes a better norm system. Monotonicity states that if $\Omega \supset \Omega'$, then $\Omega \succ \Omega'$. These observations are formally captured in the following lemma.

Lemma 2 *The utility function u_N guarantees responsiveness and monotonicity.*

Proof. 2 *To prove responsiveness it suffices to show that given a norm system Ω such that $n_i \in \Omega$, $n_j \notin \Omega$, and $n_j \succ n_i$, then $\Omega \setminus \{n_i\} \cup \{n_j\} \succ \Omega$. Let us note $\Omega_{-i} = \Omega \setminus \{n_i\}$. Since $u_N(\Omega) = u_N(\Omega_{-i}) + u_n(n_i) < u_N(\Omega_{-i}) + u_n(n_j) = u_N(\Omega_{-i} \cup \{n_j\})$, then $\Omega \setminus \{n_i\} \cup \{n_j\} \succ \Omega$ holds. As for monotonicity, this immediately follows from the definition of utility of a moral value in Eq. 9. Given two norm systems such that $\Omega \supset \Omega'$, it is clear that $u_N(\Omega) > u_N(\Omega')$ since the value support for each norm in $\Omega \setminus \Omega'$ is greater or equal than 1, and therefore $\Omega \succ \Omega'$.*

At this point, we can quantitatively compare norm systems based on the moral values that they support. Hence, we are ready to define a new multi-objective optimisation problem involving moral values as an extension of problem 1.

Problem 2 *Given a norm net $NN = \langle N, R \rangle$, a representation power function r , a fixed budget b , a set of values V and a total order \succ over its values, the value-based maximum norm system problem with limited budget (VMNSPLB) is the problem of finding a sound norm system $\Omega \subseteq N$ with maximum representation power, minimum cost limited by some non-negative budget b , and maximum value support.*

Coding the VMNSPLB problem amounts to extend the multi-objective problem in Eq. 1 with this additional objective of value support maximisation:

$$\max \sum_{i=1}^{|N|} x_i \cdot u_n(n_i) \quad (12)$$

The resulting multi-objective optimisation problem can be formulated as a single-objective problem by aggregating the three objectives by means of scalarisation (Hwang and Masud 2012). This is achieved by: (i) *normalising* the representation, cost, and value support values; and (ii) *prioritising* these three different preference criteria by means of weights.

On the one hand, normalisation is performed by considering maximum values. In this manner, we can readily normalise representation values by considering a maximum representation power computed as $\mathcal{R}_{max} = \sum_{n_j \in G_N} r(n_j)$, where G_N stands for the set of norms that are not directly generalised by any other norm. As for cost values, we can normalise by means of the maximum budget b . Regarding

value support, we can safely normalise moral values by considering $\mathcal{V}_{max} = \sum_{i=1}^{|N|} u_n(n_i)$.

On the other hand, prioritisation can be performed by employing three different weights w_r, w_c, w_v that respectively measure the relative importance of maximising norm representation power, minimising norm costs, and maximising moral value support.

Putting all this together, we can cast the VMNSPLB problem as a single-objective optimisation problem that can be solved by the following linear problem:

$$\begin{aligned} \max \left[\frac{w_r}{\mathcal{R}_{max}} \cdot \sum_{i=1}^{|N|} x_i \cdot r(n_i) + w_c \cdot \left(y - \frac{1}{b} \sum_{i=1}^{|N|} x_i \cdot c(n_i) \right) + \right. \\ \left. + \frac{w_v}{\mathcal{V}_{max}} \cdot \sum_{i=1}^{|N|} x_i \cdot u_n(n_i) \right] \end{aligned} \quad (13)$$

subject to constraints from equations 2 to 8 together with the following constraint related with prioritisation weights:

$$w_r + w_c + w_v = 1 \quad w_r, w_c, w_v \in [0, 1] \quad (14)$$

and two additional constraints related with y : a binary indicator variable that allows us to turn the minimisation problem in equation 1 into a maximisation problem, since finding the norm system with minimum (normalised) cost ($\frac{1}{b} \sum_{i=1}^{|N|} x_i \cdot c(n_i)$) amounts to maximising expression $y - \frac{1}{b} \sum_{i=1}^{|N|} x_i \cdot c(n_i)$. Hence, y must satisfy that:

$$y \in \{0, 1\} \quad (15)$$

$$y \leq \sum_{i=1}^{|N|} x_i \leq M \cdot y \quad (16)$$

where M is a very large number.² Furthermore, it is worth mentioning that this indicator variable guarantees that no cost is added to the objective function if no norm is chosen.

Finally, notice that the specification above corresponds to a maximization problem whose constraints are all inequalities. Hence, it is in standard form and it can be solved with state-of-the-art linear program solvers such as CPLEX or Gurobi.

Example 2 *Adding moral values to previous example (see Fig. 2), Figure 3 considers that n_1 supports a “free movement of persons” moral value (v_1), whereas n_2, \dots, n_5 support a “safety” moral value (v_2). Let us consider that: i) “free movement” is preferable than “safety” (namely, $v_1 \succ v_2$), then, by Eq. 9, $u(v_1) = 2$ and $u(v_2) = 1$; and that ii) moral values are the only criterion to consider (i.e., $w_r = w_c = 0$ and $w_v = 1$). Therefore, our problem amounts to finding the sound norm system that has maximum value support. Then, if we encode the problem, a linear program solver results in $\Omega = \{n_1\}$ as a value-optimal sound norm system³.*

²In our problem M can be defined to be strictly larger than $|N|$.

³In fact, there are three alternative solutions $\Omega = \{n_1\}$, $\Omega = \{n_2, n_4\}$, $\Omega = \{n_2, n_5\}$ that constitute three different value-optimal sound norm systems.

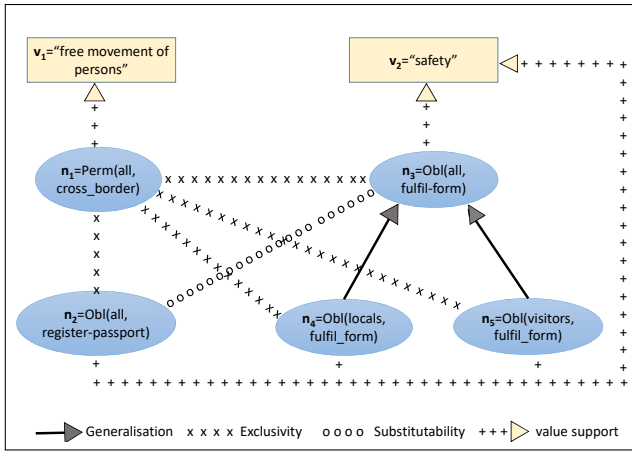


Figure 3: Example of rules of border control (n_1, \dots, n_5) together with the values they support (v_1, v_2).

Conversely, if we consider norm representation and value support to have the same priority (i.e., $w_r = w_v = 0.5$, $w_c = 0$) and we consider the following recursive computation of the representation power function

$$r(n_i) = 1 + \sum_{n_j \in \text{Children}(n_i)} r(n_j) \quad (17)$$

then, a linear program solver results in $\Omega = \{n_3\}$ as the only value-optimal sound norm system.

Discussion: dealing with norms in force

So far we have considered a normative process that starts from scratch. This applies for new-born communities willing to establish their regulations. However, other scenarios –such as dynamic organisations or evolving societies– may already count on a collection of norms currently in force. If this is the case, the decision making process must consider such norms as part of the norm reasoning together with the norm net containing a new collection of candidate norms. In what follows, we discuss how this situation can be readily accommodated in our optimisation-based framework.

Henceforth, we shall consider that N_0 stands for a set of norms in force and $NN = \langle N, R \rangle$ stands for a norm net containing candidate norms to reason about together with their relationships. Consequently, we build an *extended norm net* that merges the norms in force with the norm net. This will result in a new norm net $NN' = \langle N_0 \cup N, R \cup R_0 \rangle$, where R_0 stands for generalisation, exclusivity and substitutability relationships holding between the norms in force (in N_0) and the candidate norms in the norm net, namely in N .

Example 3 Figure 4 shows an extension of previous example (for the sake of simplicity, it does not show moral values). The norms in force (i.e., in N_0) are the following:
 n_6 : Prohibition(all_passengers, unattend_luggage)
 n_7 : Obligation(all_passengers, passport_control)

Norm n_6 prohibits all passengers to leave their luggage unattended and norm n_7 forces all passengers to go

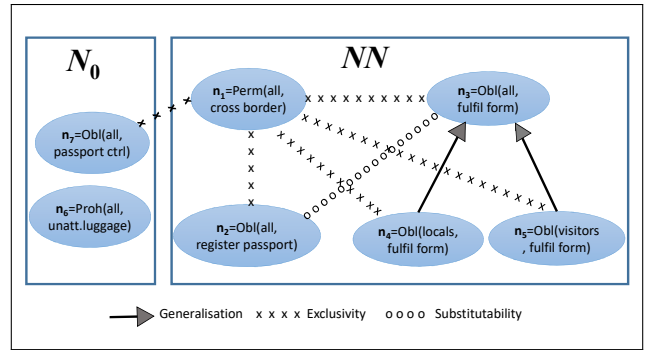


Figure 4: Example of N_0 , a norm net NN , and an exclusivity relation between norms in N_0 and N .

through passport control. Notice that there is a new exclusivity relation between n_7 and n_1 , so that $(n_7, n_1) \in R_0$ will be added to the new norm net NN' .

Considering the resulting norm net NN' suffices for encoding the VMNSPLB problem so to determine the new optimal sound norm system. However, the new norm system may not preserve those norms in N_0 . Therefore, if norms in force must be preserved (or in other words, if current norm system can only be extended), then we need to add further constraints to our optimisation problem to guarantee norms in N_0 are always selected. Namely, $x_n = 1$ for all $n \in N_0$.

Conclusions and future work

This paper contributes to the state of the art in the Normative Multi-Agent Systems research area by modeling the problem of choosing the right set of norms to enact in new-born (and dynamic) social groups. The proposed formalisation includes multiple constraints and objectives into optimisation problems which are encoded as linear programs which can then be solved with state-of-the-art solvers.

On the one hand, constraints guarantee, among other things, to obtain sound (i.e., conflict free and non-redundant) norm systems. On the other hand, optimisation objectives guide the search for the optimal norm system and make explicit social preference criteria –such as costs and supported moral values– that are typically implicit in decision making processes. In fact, including moral values in the decision making constitutes a specially relevant contribution of the paper, since they encapsulate the socially agreed notion of *rightness*.

We are currently working on using this approach to enhance both transparency and the decision process in participatory tools (Mad 2017; BCN 2017) where participants express their opinions over individual proposals without considering their combination. However, our work opens many paths to future research. First, it is worth studying the hardness of different norm decision scenarios depending on e.g. the density of norm relationships. Second, reasoning about norms and moral values could be taken a step further by considering that norms support values positively or negatively. Moreover, we could deal with a partial order over moral values or perform automated discovery of norm relationships.

References

- [Agotnes and Wooldridge 2010] Agotnes, T., and Wooldridge, M. 2010. Optimal Social Laws. In *Proceedings of the International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, 667–674.
- [Andrighetto et al. 2013] Andrighetto, G.; Governatori, G.; Noriega, P.; and van der Torre, L. W. 2013. *Normative multi-agent systems*, volume 4. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik.
- [Barbera, Hammond, and Seidl 2004] Barbera, S.; Hammond, P.; and Seidl, C., eds. 2004. *Handbook of utility theory*. Boston, Mass. [u.a.]: Kluwer.
- [BCN 2017] 2017. City of Barcelona participation portal. <https://decidim.barcelona>.
- [Bench-Capon and Atkinson 2009] Bench-Capon, T. J. M., and Atkinson, K. 2009. Abstract argumentation and values. In *Argumentation in Artificial Intelligence*. 45–64.
- [Campos et al. 2013] Campos, J.; Lopez-Sanchez, M.; Salamó, M.; Avila, P.; and Rodríguez-Aguilar, J. A. 2013. Robust regulation adaptation in multi-agent systems. *ACM Transactions on Autonomous and Adaptive Systems* 8:1–27. IF 1.286, Q1 Computer Science, Theory & Methods cat.
- [DeTar 2013] DeTar, C. 2013. *InterTwinkles: Online Tools for Non-Hierarchical, Consensus-Oriented Decision Making*. Ph.D. Dissertation, Media Arts and Sciences at the Massachusetts Institute of Technology.
- [Griffiths and Luck 2010] Griffiths, N., and Luck, M. 2010. Norm Emergence in Tag-Based Cooperation. In *Proceedings of COIN*.
- [Grossi and Dignum 2005] Grossi, D., and Dignum, F. 2005. From abstract to concrete norms in agent institutions. In *Proceedings of the Third international conference on Formal Approaches to Agent-Based Systems, FAABS'04*, 12–29. Berlin, Heidelberg: Springer-Verlag.
- [Hwang and Masud 2012] Hwang, C.-L., and Masud, A. S. M. 2012. *Multiple objective decision making methods and applications: a state-of-the-art survey*, volume 164. Springer Science & Business Media.
- [Jackson and Kuehn 2016] Jackson, S. K., and Kuehn, K. M. 2016. Open source, social activism and” necessary trade-offs” in the digital enclosure: A case study of platform cooperative, loomio. org. *tripleC: Communication, Capitalism & Critique. Open Access Journal for a Global Sustainable Information Society* 14(2):413–427.
- [Karp 1972] Karp, R. M. 1972. Reducibility among combinatorial problems. In *Complexity of computer computations*. Springer. 85–103.
- [Kohler et al. 2014] Kohler, T.; Steghoefer, J.-P.; Busquets, D.; and Pitt, J. 2014. The value of fairness: Trade-offs in repeated dynamic resource allocation. 1–10. IEEE.
- [Kollingbaum et al. 2007] Kollingbaum, M. J.; Norman, T. J.; Preece, A.; and Sleeman, D. 2007. *Norm Conflicts and Inconsistencies in Virtual Organisations*. Berlin, Heidelberg: Springer Berlin Heidelberg. 245–258.
- [López, Luck, and d’Inverno 2002] López, F. L.; Luck, M.; and d’Inverno, M. 2002. Constraining autonomy through norms. In *AAMAS*, 674–681. ACM.
- [Mad 2017] 2017. City of Madrid participation portal. <https://decide.madrid.es/>.
- [Modgil 2006] Modgil, S. 2006. Value based argumentation in hierarchical argumentation frameworks. In *Proceedings of the 2006 Conference on Computational Models of Argument: Proceedings of COMMA 2006*, 297–308. Amsterdam, The Netherlands, The Netherlands: IOS Press.
- [Morales et al. 2015] Morales, J.; Lopez-Sanchez, M.; Rodríguez-Aguilar, J. A.; Wooldridge, M.; and Vasconcelos, W. 2015. Synthesising liberal normative systems. In *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems, AAMAS '15*, 433–441. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems.
- [New 2016a] 2016a. Citizen space in new zealand’s ministry of justice: Better family violence law. <https://consultations.justice.govt.nz/policy/family-violence-law/>.
- [New 2016b] 2016b. Pensions policy institute in new zealand: Citizen’s pension: Lessons from new zealand and state pension reform: The consultation response. <http://www.pensionspolicyinstitute.org.uk/press/press-releases/citizens-pension-lessons-from-new-zealand-and-state-pension-reform-the-consultation-response>.
- [Par 2017] 2017. Parlement citoyens: institutional participation portal for the discussion of french laws. <https://parlement-et-citoyens.fr/>.
- [Rey 2017] 2017. City of Reykjavík participation portal. <http://reykjavik.is/en/participation>.
- [Riveret et al. 2014] Riveret, R.; Nepomuceno, E.; Pitt, J.; and Artikis, A. 2014. Self-governance by transfiguration: From learning to prescription changes. 70–79. IEEE.
- [Savarimuthu et al. 2013] Savarimuthu, B.; Cranefield, S.; Purvis, M.; and Purvis, M. 2013. Identifying prohibition norms in agent societies. *Artificial Intelligence and Law* 21(1):1–46.
- [Shoham and Tennenholtz 1995] Shoham, Y., and Tennenholtz, M. 1995. On social laws for artificial agent societies: off-line design. *Artificial Intelligence* 73(1-2):231–252.
- [Vasconcelos, Kollingbaum, and Norman 2009] Vasconcelos, W. W.; Kollingbaum, M. J.; and Norman, T. J. 2009. Normative conflict resolution in multi-agent systems. *Autonomous Agents and Multi-Agent Systems* 19(2):124–152.
- [Villatoro, Sabater-Mir, and Sen 2011] Villatoro, D.; Sabater-Mir, J.; and Sen, S. 2011. Social instruments for robust convention emergence. In *IJCAI*, 420–425.
- [Weerakkody and Reddick 2012] Weerakkody, V., and Reddick, C. G. 2012. *Public sector transformation through e-government: experiences from Europe and North America*. Routledge.