Data-Driven Revision of Conditional Norms in Multi-Agent Systems

Davide Dell'Anna, Natasha Alechina, Fabiano Dalpiaz, Mehdi Dastani, Brian Logan

Department of Information and Computing Sciences, Utrecht University, The Netherlands

{d.dellanna, n.alechina, f.dalpiaz, m.dastani, b.logan}@uu.nl

Universiteit Utrecht

Norm Revision in Normative Multi-Agent Systems

Normative Multi-Agent Systems are MASs where **norms** are enforced to ensure **system-level objectives**.
Example. A smart traffic system where traffic norms ensure smooth traffic flow and low CO_2 emissions.

Problem. When the system-level objectives change (e.g., new levels of CO_2 to reduce air pollution), the norms need to change too.

Research Question. How to automatically revise norms to align them with the new systems objectives?

Solution. DDNR (Data-Driven Norm Revision), a data-driven approach to the automatic synthesis of norms from a data set of traces describing the behavior of the agents in the system.

Download the full paper and source code of DDNR!

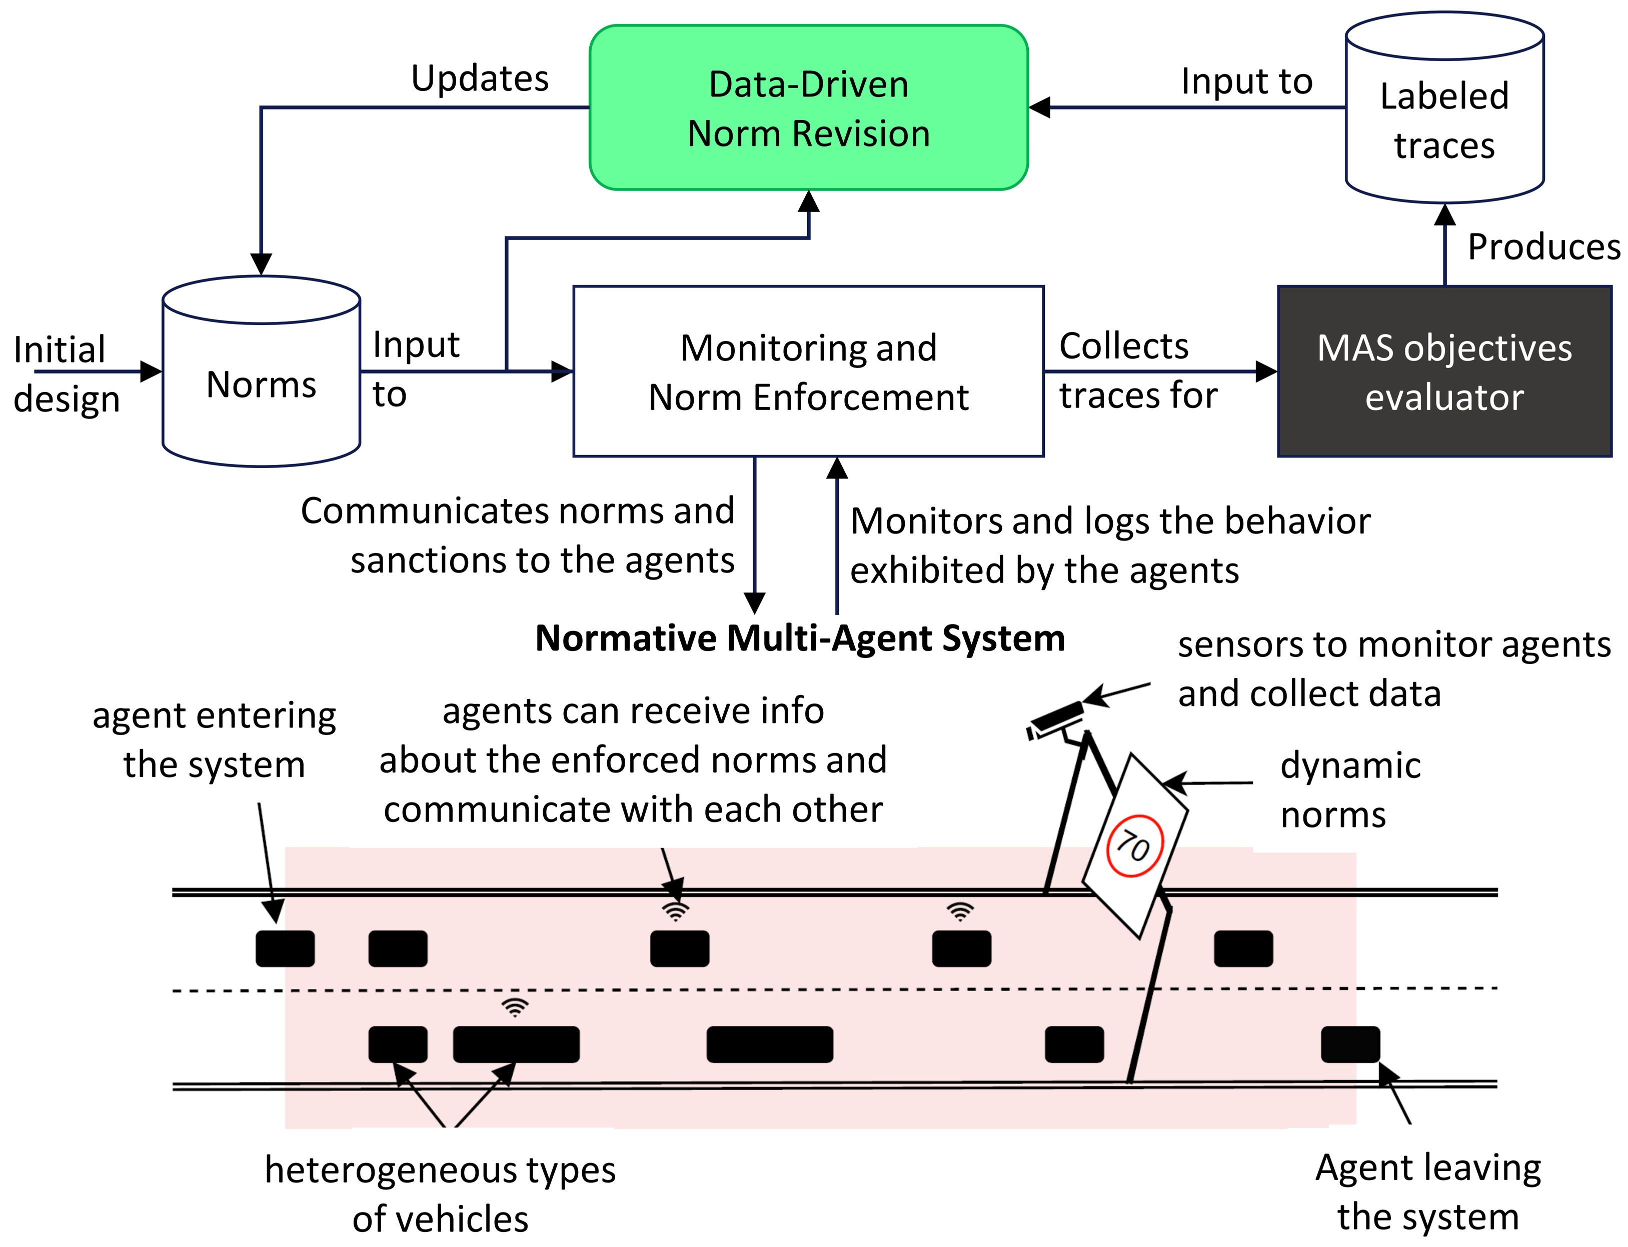

DDNR: Data-Driven Norm Revision

Data Set of Finite Execution Traces (Agent Behaviors)

Example. A trace represents a vehicle's journey through the highway and it is labeled as *positive* if the travel time and the CO_2 emitted are adequate.

$(\{km_1, sp_{30}, car\}, \{km_2, sp_{22}, car\}, \ \ldots, \{km_9, sp_{18}, car\}, \{km_{10}, sp_{14}, car\})$

Conditional Prohibitions with Deadlines (ϕ_C, ϕ_P, ϕ_D)

Example. $(km_2 \wedge car, sp_{70}, km_7)$ represents a norm "If a car enters the 2^{nd} km of the highway, it is prohibited from driving faster than 70 km/h until it reaches the 7^{th} km of the highway".

Data set Γ					Norm n	Examples of weakening of n: $n_1 = (a \wedge b, c, e)$ $n_2 = (a, c \wedge d, e)$ $n_3 = (a, c, e \vee a)$			Examples of strengthening of n: $n_4 = (a \vee c, c, e)$ $n_5 = (a, c \vee i, e)$ $n_6 = (a, c, e \wedge f)$			Examples of alteration of n: $n_7 = (a \vee c, c \vee j, e \vee b)$
Trace	s_1	s_2	s_3	Label	(a, c, e)	n_1	n_2	n_3	n_4	n_5	n_6	n_7
γ_1	a,b	c,d	e,f	*negative*	✓	✓	✓		✓	✓	✓	
γ_2	a	c,d	k,l	*negative*	✓		✓		✓	✓	✓	✓
γ_3	a,b	i,j	e,f	*positive*						✓		
γ_4	a,b	i.j	k,l	*positive*						✓		
γ_5	g,h	c,d	e,f	*negative*					✓			✓
γ_6	g,h	c,d	k,l	*negative*					✓			✓
γ_7	g,h	i,j	e,f	*positive*								
γ_8	g,h	i.j	k,l	*negative*								
Accuracy of the norm w.r.t. Label					62.5%	50%	62.5%	37.5%	**87.5%**	37.5%	62.5%	75%

✓in cell (i, j) indicates that trace i *violates* norm j. Empty cell (i, j) indicates that trace i is compliant with j.
Coloured cells indicate different types of states, used for the Synthesis Step (details in the paper).

DDNR: Synthesis Step

Synthesises a set $\mathcal{R}(n)$ of candidate revisions of n based on Γ. 3 types of revisions:
Alteration prohibit different behaviors.
Weakening prohibit fewer behaviors.
Strengthening prohibit more behaviors.

DDNR: Selection Step

Chooses from $\mathcal{R}(n)$ the revised norm n^* with **highest accuracy**, i.e., the norm most in line with the labeling of traces w.r.t. the system-level objectives.

Empirical Evaluation via Traffic Simulations of the Highway Section

Experiments. Traffic simulations via SUMO traffic simulator.
Agents. 50% cars, 50% trucks. 75% always norm-compliant, 25% ignores the norms (may or may not violate them, e.g., due to traffic jams).
Norms. 100 different initial speed limit norms.
Traces. 100 data sets (1 per norm), each with 1500 traces (1 per vehicle).
Objectives. $CO_2 \leq 100g/s$ and Travel Time $\leq 450s$.

Results. The revised norms are significantly more accurate than the initial ones (ca. +30% with alteration), i.e., significantly better aligned with the system-level objectives.

Additional experiments (in the paper).
Revision of multiple norms: average accuracy change ca. +25%.
Generalization to unseen traces: average accuracy change ca. +13%.

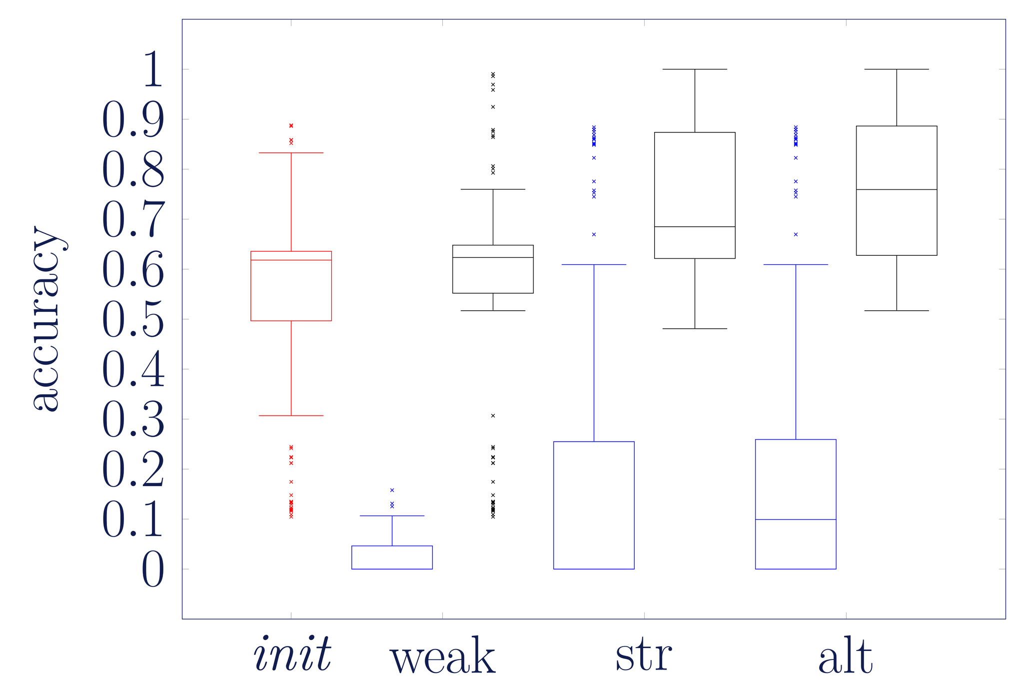

Accuracy of the initial 100 norms, accuracy change, and accuracy of the 100 revised norms.