

Instrumental Properties of Social Testbeds

Javier Insa-Cabrera

José Hernández-Orallo

Dep. de Sistemes Informàtics i Computació,
Universitat Politècnica de València

jinsa@dsic.upv.es

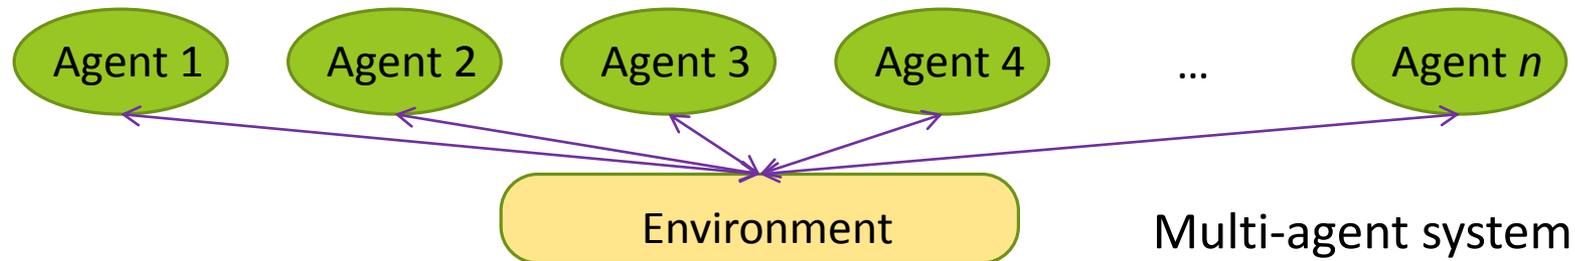
jorallo@dsic.upv.es

Outline

- **Social intelligence testbeds**
- **Social Properties**
 - **Interactivity, non-neutrality, competitive and cooperative anticipation**
- **Instrumental Properties**
 - **Discrimination, grading, boundedness, team symmetry, reliability, efficiency**
- **Univocal Properties**
 - **Validity**
- **Examples of application to some environments**
- **Conclusions**

Social intelligence testbeds

- Issues about the evaluation of social intelligence



- What makes a MAS social? The agents or the environment?
 - Some have studied this focussing on the agents:
 - Hibbard's adversarial matching pennies (Hibbard 2008-2011).
 - Darwin-Wallace distribution (Hernandez-Orallo et al 2011).
- What makes *social* and *general* intelligence different?
- How can the influence of the other agents be regulated?

Can we set some criteria that a multi-agent environment should meet in order to evaluate social intelligence?

Social intelligence testbeds

- We first analyse multi-agent systems in terms of:
 - Usual MAS with actions, observations and rewards.
 - Simultaneous for every agent.
 - Agent slots and line-ups
 - The same environment can be instantiated with different sets of agents, leading to very different behaviours.
 - Teams
 - In practice, it is unlikely that alliances and coalitions appear spontaneously.
 - We consider the existence of previously defined teams
 - Rewards are the same for all members in a team.

Social intelligence testbeds

- We use a customary definition:

$$\Upsilon(\Pi, w_L, M, w_M, w_S) \triangleq \sum_{\mu \in M} w_M(\mu) \sum_{i=1}^{N(\mu)} w_S(i, \mu) \sum_{l \in L^{N(\mu)}(\Pi)} w_L(l) R_i(\mu[l])$$

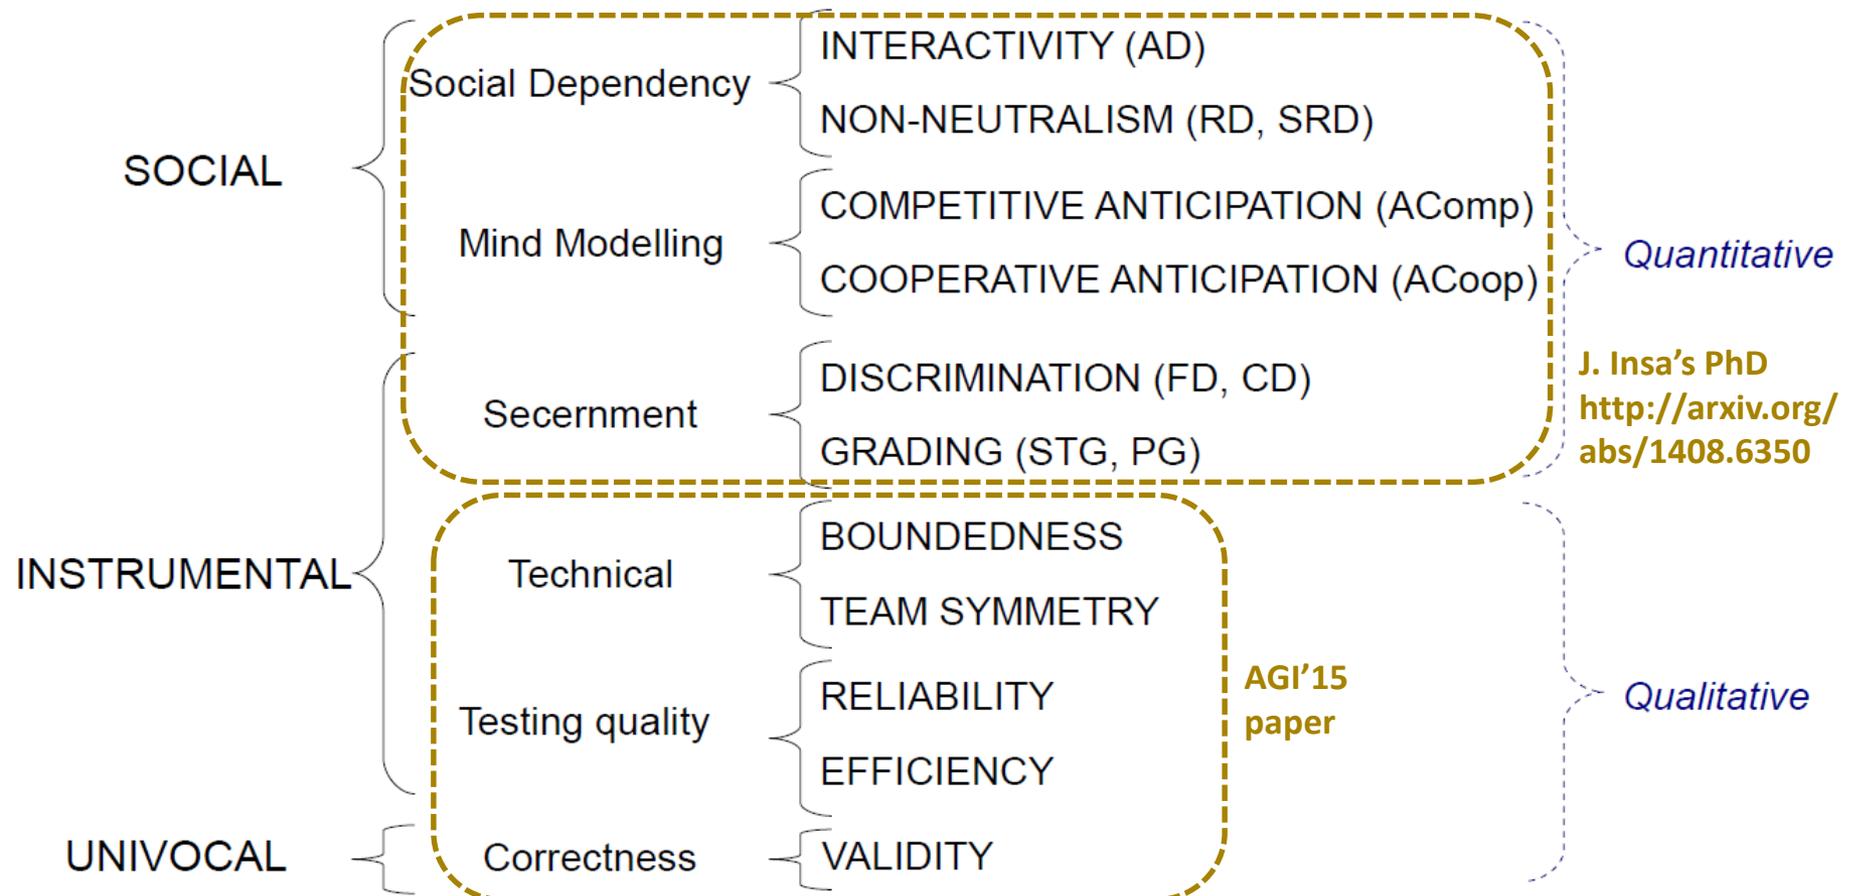
- Set of agents (e.g., robocup players)
- Distribution of line-ups (e.g., Pr teammates and opponents)
- Set of environments (e.g., several game configurations)
- Distribution of environments (e.g., Pr configurations)
- Distribution of slots (e.g., positions of the evaluated agent)

Social intelligence testbeds

- In a multi-agent environment:
 - A rich configuration may lack any social interaction if other agents have no effect on the reward of the evaluated agent.
 - The ability of the opponents is key, especially for competitive social intelligence.
 - The ability of the teammates is also key, especially for cooperative social intelligence.
 - The way in which we sample the distributions is also important.

Social intelligence testbeds

- We have introduced a series of formal properties to analyse the suitability of a multi-agent environment to evaluate social intelligence:



Social properties

- Interactivity (Action dependency):
 - Action sensitivity to other agents.
 - Whether the inclusion of different agents in the multi-agent environment has an effect on what the evaluated agent does.
- Non-neutrality (Reward / slot result dependency)
 - Effect of other agents on the evaluated agent's rewards.
 - From the six forms of symbiosis in ecology:
 - Neutralism (0,0), amensalism (0,-), commensalism (+,0), competition (-,-), mutualism (+,+), and predation/parasitism (+,-).
 - We can simplify this to neutralism, cooperation (including commensalism and mutualism) and competition (including the rest).
 - Non-neutrality measure: 0 (neutralism) > 0 (cooperation), <0 (competition)

Social properties

- Competitive anticipation
 - The evaluated agent can perform better if their opponents/competitors can be well anticipated.
 - It is measured relative to the results against random agents.
- Cooperative anticipation
 - The evaluated agent can perform better if their teammates/cooperators can be well anticipated.
 - It is measured relative to the results with random agents.

Instrumental properties

- Discrimination
 - Given a set of agents, we want the testbed to give significantly different values to the agents so that their social abilities can be discriminated.
- Grading (strict total grading or partial grading)
 - Measures how much the metrics resemble a total order or, more precisely, how frequent is that for three agents (a,b,c) if $a \leq b$, $b \leq c$ then $a \leq c$, when placed in different slots.
 - This can be calculated for a strict total order or for a partial order

Instrumental properties

- Boundedness
 - Weights for environments, agents and line-up being bounded (or being probability measures).
 - Zero-sum teams (in the limit). Given several teams, the sum of rewards of all teams sum up to 0.
- Team symmetry
 - If we make the environment team-symmetric, in terms of positions inside the team (intra-team) and between teams (inter-team), we do not need the slot distribution.
 - Many games are not team-symmetric:
 - Prey-predator
 - Football (goalkeepers very different from other players)

Instrumental properties

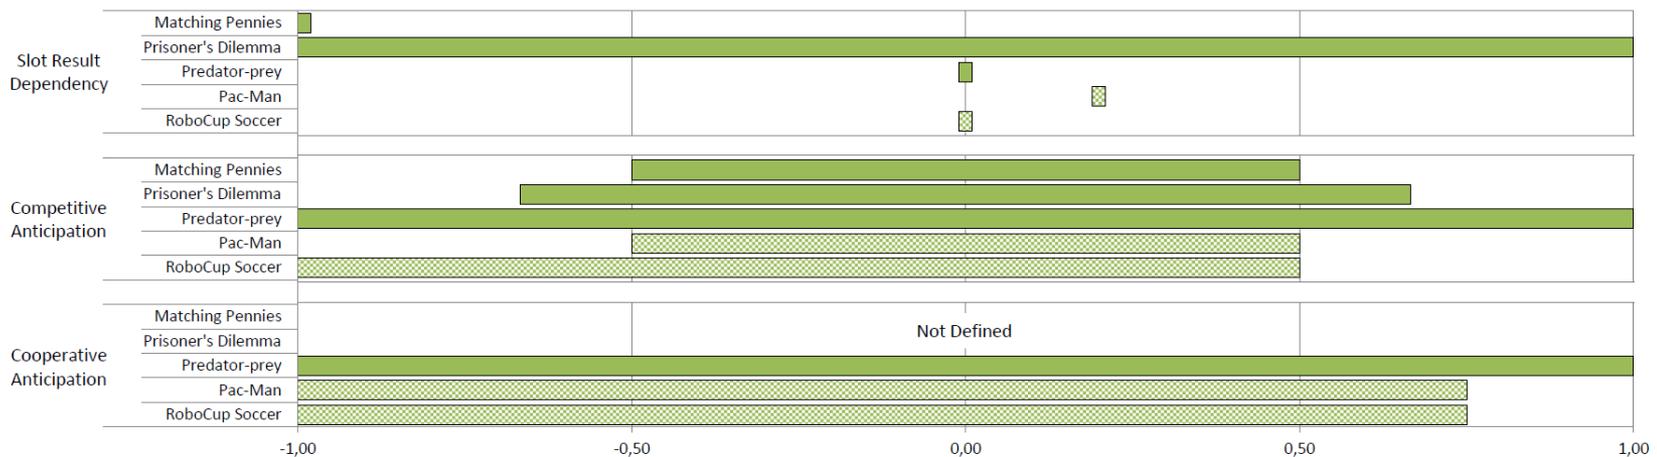
- Reliability:
 - How close the measured value is to the actual value given by the definition.
 - Tests sample over the distributions of environments, slots and agents, and have to limit trial duration.
- Efficiency
 - How much reliability can be achieved in terms of the time devoted to testing.
 - It depends on how representative and effective the sampling over the distributions is.

Univocal properties

- Validity:
 - Main testbed pitfalls may originate from two reasons.
 - If the testbed allows for good performance without social intelligence.
 - Social characteristics are not very relevant and general intelligence must suffice.
 - If social intelligent agents do not get good performance in the testbed.
 - The test may measure some other abilities that are not social intelligence.

Application to some environments

- We have applied the properties to several MAS:
 - Five MAS environments/games have been analysed:
 - Matching pennies (any slot)
 - Prisoner's dilemma (any slot)
 - Predator-prey (3 predators, 1 prey, evaluatee acts in predator slot)
 - Pac-man (any slot)
 - RoboCup Soccer (any slot)
 - Using *all possible agents*.



Application to some environments

- The ranges are wide if all possible agents are considered.



- The analysis changes radically when using families of agents instead of all.

Application to some environments

- For the instrumental properties there is more diversity.

	Boundedness	Symmetry	Validity	Reliability	Efficiency
Matching Pennies	✓	✗	✗	✗	✓
Prisoner's Dilemma	✓	✓	✗	✓	✓
Predator-prey	✓	✗	✗	✗	✓
Pac-Man	✗	✗	✗	✗	✗
RoboCup Soccer	✓	✗	✗	✗	✗

- Validity problems originate because many other abilities are more relevant than social intelligence for these environments.
 - Also, the first two lack cooperation.
- Reliability problems, as many environments are stochastic.
 - Even with same line-up and slots, results can be very different.
 - With several repetitions, the average can converge fast for some of them (efficiency).

Conclusions

- We have derived a series of formal, effective properties to characterise multi-agent systems in terms of how necessary and sufficient social intelligent is for them.
- The properties are more fine-grained and allow for a more informative characterisation of a testbed.
 - Go well (but controversially) beyond game theory equilibria and other properties.
- Using five environments as examples, we have seen that the set of agents that is considered is crucial.
 - Considering all possible agents leads to virtually any possibility in any game.
- Main questions for future work.
 - Define reasonable subsets of agents, using agent description languages and see how the ranges for the properties change for these subsets.
 - How many different games/environments are necessary so that the particularities of the games/environments are finally irrelevant for the aggregate measure?
 - Communication and language have been left out.