

Multi-agent preference reasoning: aggregation, elicitation, and manipulation

Francesca Rossi

Preferences are ubiquitous in everyday decision making. They should therefore be an essential ingredient in every reasoning tool. The talk will consider multi-agent systems where each agent expresses his preferences over the possible decisions, and the goal is to aggregate such preferences in order to choose a satisfying decision.

After defining some preference modelling frameworks, such as soft constraints and CP-nets, we will discuss their expressiveness and computational properties. We will then consider ways of aggregating preferences of multiple agents, also in the presence of various forms of uncertainty, such as missing or vague data. Uncertainty will also be studied in a context of preference elicitation, to show how eliciting preferences can help in removing enough uncertainty to take good decisions. Multi-agent preference aggregation will be presented also by exploiting classical notions and results from fields such as voting theory, computational social choice, and matching. Various properties of preference aggregation will be presented and studied, such as fairness and non-manipulability, which are very desirable for multi-agent decision support systems.