

# A satisficing agreements model

Guido Boella

Università degli Studi di Torino  
guido@di.unito.it

Gabriella Pigozzi

Université Paris Dauphine  
gabriella.pigozzi@dauphine.fr

Marija Slavkovic

University of Luxembourg  
marija.slavkovic@uni.lu

Leendert van der Torre

University of Luxembourg  
leon.vandertorre@uni.lu

**Abstract**—Satisficing, the concept proposed by Herbert Simon, as an approach to reaching agreements is little explored. We propose a model for satisficing agreement reaching for an adaptive collaborative group of agents. The group consists of one human agent familiar with the problem and arbitrarily many artificial agents. Our model raises to the team level the recognition-primed decision model constructed in the field of cognitive decision-making by using social choice for reaching group opinions.

**Index Terms**—agreement reaching; satisficing; judgment aggregation;

## I. INTRODUCTION

Agreements are essential to the problem in agent coordination. Negotiation, in its many forms as argumentation, auctions, bargaining etc., is seen as an essential technology for reaching agreements [1]. However, negotiation is not suited for cooperation in fast changing environment. Negotiation protocols require several rounds of exchanges between the agents before an agreement is reached. Consequently the situation can change while the agents are still negotiating about the situation. It is in dangerous environments that we would like to replace the on-site human teams with robots and drones.

Agreement-reaching problems can be addressed through traditional decision-making; see for example Chapter 1 of [2]. Decision-making is driven by the concept of rationality associated with the decision-maker. A rational agent acts in his own best interest, i.e., chooses, given his knowledge about the world, those options that are optimal in the sense that they maximize the agent's expected utility. Optimizing is difficult when the agents' resources are limited.

People are not good rationalizers, (pg.ix [3]), but they are able to coordinate successfully under time pressure even when all adequate information is not available, when their goals are unclear and the procedures they have to follow are poorly defined. Can these skills of high adaptiveness be advanced to artificial agents? The answer to this question begets another question: how do people make decisions under time pressure, in dynamic conditions in uncertain environments?

Gary A. Klein and his associates, studied how firefighter commanders make decisions under extreme time pressure, [4]. They found that, when a commander has prior experience with a problem, which is normally the case, he acts according to the *recognition-primed decision* (RPD) model summarized in Figure 1.

Klein *et al.* found that the RPD model exemplifies Herbert Simons's [5] notion of satisficing since the observed comman-




Fig. 1. The recognition-primed decision model (pg. 203, [4]).

ders were looking for the first workable course of action rather than trying to find the best possible option.

Although firefighters operate as teams, the RPD model is a model of a single agent - that of the commander who coordinates the team. The commander is the one who assesses the situation, generates possible courses of action, evaluates them, and gives orders for actions that should be implemented by the rest of the team. The coordination through the commander is applicable when the commander is on the ground together with the team that executes his orders. Otherwise, the commander is not able to assess the situation, nor make evaluations by himself.

To apply the RPD model to teams, *e.g.*, when the agents on the ground are robots or drones that coordinate remotely with the commander, we need to raise it from an individual raised to team level. This is the question that we address here, *how to raise the recognition-primed decision model to a team level for use in multi-agent systems?* Particularly, how can we do so without relying on negotiation?

We consider a mixed human-robot team in which there is one human, called *initiator*, which has a role similar to that of the firefighter's commander. The rest of the agents are artificial, called *executors*. Unlike the commander, the initiator is not on the ground and has to fully rely on the executors for the following processes: situation assessment, verifying expectancies, and evaluating the potential course of action.

The situation assessment is the process in which the initiator matches the problem with a possible solution (a goal). The goal is good enough if and only if a certain combination of *cues* can be identified as (not) present. Verifying expectancies is the process in which the agreements on cue's presence or absence and goals are verified as the agents proceed with pursuing the goal. While the commander of the RPD model can identify the cues himself, the initiator needs to obtain an agreement on them, by considering the opinions of the executors. The initiator also needs the input from the executors for verifying expectancies.

The reasoning according to the RPD model is very fast;

Klein *et al.* estimated that the fireground commanders make around 80 percent of their decisions in less than a minute (pg.4, [6]). Inevitably, the making of decisions can be expected to be longer when there is agreement involved. The executor needs to be able to get all of the information from the agents at once and deduce the courses of action from it. We advance that social choice, in particular judgment aggregation [7], should be used as an agreement reaching technology in this case.

This paper is structured as follows. In Section II we present the RPD model of Klein. In Section III we propose a model of team coordination obtained by raising the RPD model to the team level and in Section IV we address the issue of reconsidering agreements. In Section V we address the social choice aspects of the model and in Section VI we make our conclusions. We do not consider issues such as how exactly the initiator produces the plan, or how the communication between the agents is implemented.

## II. THE (SINGLE AGENT) RPD MODEL

The process of reaching satisficing agreements consists of these steps: identifying a possible goal and a cue pattern that establish that goal as satisficing, agreeing on whether the goal is good enough by agreeing on whether the cue pattern is present, establishing a plan for the goal and verifying that the agreements are in accordance with the current state of the world. The agreement on whether certain cues are present or not implicitly is an agreement on what the state of the world is.

Consider a group of ten cleaner agents in charge of cleaning a shopping mall, all using the same supply closet in their daily tasks. Their remote human supervisor is Mrs. J. The agents execute their individual tasks independently and autonomously. At one point three agents report that there is a large box in front of the closet. If most of the agents cannot access the supply closet, they cannot conduct their tasks. Mrs. J contacts the agents to inquire if the box is a problem. Seven of the physically bigger agents are unable to access the supply closet. Mrs. J needs to find a way to resolve the problem. There are many courses of action which she can consider: leave the box, abort their tasks and wait for the situation to change; attempt to access another supply closet; attempt to move the box; etc. Each of these courses of actions is potentially risky: waiting may be too long; another supply closet is used by another group and may not contain enough supplies, moving the box may cause damage, etc. The difficulty in comparing options is that Mrs. J would have to consult other supervisors to ensure that she: a) thought of every possibility; b) knows the exact plans each of the agents is pursuing; and c) knows their exact state and capabilities.

Klein *et al.* found [4] that a fire ground commander faced with a problem first would try to match the problem to a problem-solution pattern based on his experience. Once he recognizes a potential course of action, the clues that indicate that this solution would work and what actions should be taken, the commander runs a mental simulation of the course of action. If he finds a reason why the course of action may

be inapplicable, the commander modifies the course of action to accommodate for that reason.

Mrs. J can employ the same model up to the level of matching the problem-solution pattern based on her experience. She cannot fully assess the situation and run the mental model by herself since she is not in front of the supply closet. She needs to get the relevant clues from the agents on the ground.

Mrs. J knows from experience that the box may not be removed for a long time, hence she starts by considering the option to use another supply closet ( $s_1$ ). There are two closets per floor. Mrs. J identifies the cues which establish  $s_1$  as a good solution: proceeding with  $s_1$  is possible if the agents know where it is ( $c_1$ ), and if they can either find the closet key ( $c_2$ ) or an agent that uses that closet ( $c_3$ ). Instead of running a mental simulation of  $s_1$ , Mrs. J elicits opinions from the agents on the cues and their conclusions regarding  $s_1$ . It is sufficient that at least one agent reports that  $c_1$ ,  $c_2$  or  $c_3$  is present. However,  $s_1$  is not adopted since, see Table I, no one is aware of the locations of the key. Mrs. J now considers

	$c_1$	$c_2$	$c_3$	$s_1 \leftrightarrow (c_1 \wedge (c_2 \vee c_3))$
$A_1 - A_4$	yes	no	no	no
$A_5$	no	no	no	no
$A_6, A_7$	yes	not know	no	not know
hline Group	yes	no	no	no

TABLE I  
THE GROUP'S OPINIONS REGARDING SOLUTION  $s_1$ .

another option: having the agents move the box ( $s_2$ ). She does not know the exact state of the agents, but has a general knowledge of their capacities and configuration. She also does not know exactly how big or how heavy the problematic box is. She does know that if there is one agent who can lift the box long enough for rollers to be inserted under it ( $c_4$ ), or if the agents can push an object ( $c_5$ ), and attempt lifting or pushing without damaging themselves ( $c_6$ ), then  $s_2$  is a solution they can pursue. In addition to eliciting the opinions of the agents, Mrs. J also consults an expert opinion, the manufacturer of the agents, to see if pushing is something their motors can sustain. Since there are five agents, see Table II, that can attempt  $s_2$

	$c_4$	$c_5$	$c_6$	$s_2 \leftrightarrow ((c_3 \vee c_4) \wedge c_5)$
$A_1 - A_3$	yes	no	yes	yes
$A_4$	no	no	no	no
$A_5$	no	yes	no	no
$A_6, A_7$	no	not know	yes	yes
Expert	-	-	yes	-
Group	yes	no	yes	yes

TABLE II  
THE GROUP'S OPINIONS REGARDING SOLUTION  $s_2$ .

without self-damage and three that are of the opinion that they can lift the box, Mrs. J decides to proceed with  $s_2$  as a goal and devises a set of actions, a plan, for the agents  $S = \{A_1, A_2, A_3, A_6, A_7\}$ . Her plan  $p_1$  is to have any agent find a roller, then  $A_1$  lifts the box,  $A_6$  and  $A_7$  push the roller under the box and then push the box away from the closet. She proposes this plan to the agents in  $S$ . The agents can either approve the plan or report a constraint as to why they cannot execute it. Agent  $A_1$  attempted to lift the box and failed so he reports that he is unable. Mrs. J changes the plan and now

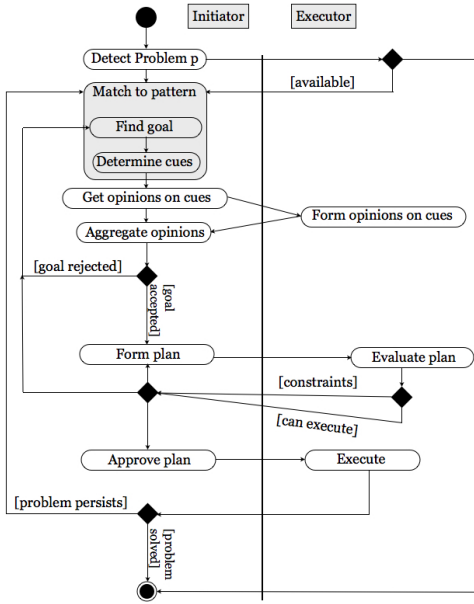


Fig. 2. The satisficing agreement model.

this task of lifting the box is assigned to  $A_2$  and  $A_3$  together. This new plan is approved by all the agents and they proceed to execution.

### III. SATISFICING AGREEMENT MODEL

We generalize the RPD model to a group of agents, by externalizing the processes of: situational assessment, verification of expectancies and plan evaluation. One agent, the human, is in charge of matching the present problem with a possible solution. This agent is called an initiator, in our example this is Mrs. J. The rest of the agents are called executors, in the example these are all ten cleaners. Figure 2 depicts the agreement-reaching process for both the initiator (left hand side) and an executor (right hand side). Once the initiator recognizes a problem, she contacts the agents to determine who is available for solving it. She proceeds with finding a possible solution and the cues under which this solution is good enough to pursue. Whether the cues are satisfied is determined by aggregating the opinions of the executors. If no solution is approved, as long as the problem persists, the initiator waits for a possibility to arise.

Once a solution is deemed acceptable, the initiator proposes a tentative plan. The initiator has some, but not full, knowledge of the capabilities of the executors. Hence, it is up to the executors to evaluate the roles assigned to them by the plan. The executors approve the plan or report the constraints which make the plan unfeasible. The initiator includes these constraints when forming another plan. If the constraints are such that the initiator cannot form a plan, the goal is dropped and another solution is proposed. The initiator may propose a plan and the agents execute it, but the problem can persist. In this case the initiator proposes another solution.

### IV. AGREEMENT RECONSIDERATION

“Members of adaptive teams utilize their pooled resources (i.e., knowledge gained from learning) to adjust their actions

according to situational requirements” (pg. 1190, [8]). The adaptation in our model is executed by the process of verifying expectancies, but also by updating the agreements as soon as unscheduled new information becomes available. Regardless of whether the update is scheduled, e.g., after a task is executed, or caused by a new observation, there are two types of information that can be cause for update or revision: a new constraint on an agreed plan, or a cue value is determined.

At any point after a plan is adopted and execution starts, an agent may report a constraint regarding the plan. The initiator adapts the plan accordingly and informs all the agents of the change. A cause for a plan revision is also the revision in the situational assessment.

The agreements reached on the cues are an estimate of what the state of the world is. An information contradicting the agreed value of a cue may lead to the change in the decision to pursue a goal.

Consider for example that, after the execution of the plan  $p_2$  had commenced, the agreed value “no” of  $c_2$  is shown to be inconsistent with the new knowledge, the initiator learning about the location of the key to the other supply closet. The initiator can replace the aggregated value for  $c_2$  and deduce a new conclusion for  $s_1$  considering the earlier opinions of the agents on  $c_1$  and  $c_3$ . As Table III shows, the agent’s values in column  $c_2$  are ignored, the group value in the same column is the new information and all values in the rightmost column are obtained by re-aggregating the new information with the old.

However, it is not always possible to adapt the situational assessment by revising only the agreement. When the new information concerns a cue that is logically related to other cues, the revision of the agreement may lead the agents to make the wrong situational assessment. Consider for example a goal  $g$  which is adopted if and only if cues  $c'$  and either  $c''$  or  $c'''$  are present, namely  $g \leftrightarrow (c' \wedge (c'' \vee c'''))$ . Furthermore, the nature of the cues is such that  $c'' \rightarrow c'''$ . Let the agreement be that  $g$  is adopted and the agents were unanimous regarding the presence of  $c'$ ,  $c''$  and  $c'''$ . After the agreement is made, the agents observe that  $c''$  is not present. If the initiator revises the only agreement, he can conclude that the goal is still good since  $c' \vee c'''$  holds even when  $\neg c''$  is the case. However, he cannot be sure if the agents considered  $c'''$  present in its own right or as a consequence of  $c'' \rightarrow c'''$ . In this case the initiator has to ask the agents to individually reconsider their own opinions and then aggregate these to reach the new agreement.

Whether reconsideration should be applied on an agreement or not depends also on the status of the group activities. For example, if the group had already commenced executing plan  $p_2$  towards goal  $s_2$ , a new agreement to pursue  $s_1$  and consequently drop  $s_2$  means wasting the resources already invested towards  $s_2$ . The initiator should thus only monitor the cues for the rejected goals only until a goal is adopted and a plan for it is approved for execution.

	$c_1$	$c_2$	$c_3$	$s_1 \leftrightarrow (c_1 \wedge (c_2 \vee c_3))$
$A_1 - A_4$	yes	no	no	not know
$A_5$	no	no	no	no
$A_6, A_7$	yes	not know	no	not know
Group	yes	yes	no	yes

TABLE III

RECONSIDERING THE GROUP'S OPINIONS REGARDING SOLUTION  $s_1$ .

## V. SOCIAL CHOICE TO AGREEMENT

Reaching an agreement on whether a pattern of cues is present, and whether the entailed conclusion to pursue a goal  $s$  is “yes” or “no”, is the situation in which a group of individuals is faced with a set of binary decisions. The decisions are interconnected, *e.g.*, the alternatives chosen for some of the decisions constrain the alternatives that can consistently be given to others. The problem is to aggregate the binary decisions of each agent on each criterion, into a collective set of binary decisions. This aggregation problem is a problem studied in *judgment aggregation* [7]. Unlike with negotiation, the agreement using judgment aggregation is obtained by aggregating the opinions from all the agents in one step.

Since each agent may have a different view or different information about the world, the input from all agents should be considered, possibly assigning different weights to the inputs of different agents. For example, the opinion of the consulted expert, regarding cue  $c_6$  in the case of  $s_2$ , should be more relevant than the opinions of the agents themselves.

## VI. CONCLUSION

We presented a model of reaching agreements for coordination that is constructed by externalizing parts of the recognition-primed decision (RPD) model of Klein *et al.*[9], [4]. Our aim is to explore whether satisficing can be used as an approach to reaching agreements. We chose to construct our model by raising the RPD model to the team level for three reasons. First, because the RPD model is itself recognized to be an instance of the satisficing concept. Second, the RPD model describes the behavior of the commanders of human teams. Although this is a single agent model its context is that of coordination. Third, the RPD describes the decision-making done by experienced commanders in dynamic settings. Agreements in dynamic, uncertain environments are ones we want to consider.

The main characteristic of the firefighters studied by Klein *et al.* is that they constitute highly adaptive teams. Burke *et al.* [8] define team adaptation as a change in team performance in response to salient cues that leads to a functional outcome for the team. The adaptive cycle of the team adaptation model presented in [8] is characterized by four core constructs: (1) situation assessment; (2) plan formulation; (3) plan execution, via adaptive interaction processes; and (4) team learning. The adaptive cycle is further characterized by *emerging cognitive states*, “which serve as both proximal outcomes and inputs to the this cycle”, (pg. 1192, [8]). The model we propose is in accordance with this team adaptation model. The emerging states in the case of our model are the agreements regarding

goals, value of cues and the adopted/refuted plans. The verification of expectancies and the valuation of plans are the way in which the emerging states are reconsidered. Learning is the process we do not explicitly considered in our model. We focus on giving a conceptual model of agreement reaching and not of adaptation, although adaptation happens concurrently. Nevertheless, we acknowledge that from the aspect of improving team performance, learning is an important process both for the initiator and the executors. The initiator can improve his accuracy in matching a problem with a goal and cues, while the executors can learn to improve their plan evaluation and cue observation accuracy.

Simon’s concept of satisficing does not offer a specific design principle and it has been initially identified with heuristic search. The agreement model we present is a different “implementation” of satisficing, since no heuristics is employed to find a solution to a problem. The model heavily depends on the initiator’s ability to match a goal to a problem and identify the relevant cues. It is the cues that play the role of satisficing criterion that establishes the goal as “good enough”.

The experience of the human initiator consists of cases, which are the product of learning. A (human) commander gathers cases from personal experience, but also by exchanging experiences with colleagues. A set of cases can be supplied to an (artificial) initiator agent. However, the power of the (human) commander is in the ability to swiftly recognize similar cases by association. For an artificial agent to be an initiator he necessarily has to be able to recognize and recall patterns in cases by association.

The use of judgment aggregation and the reconsideration of situational assessments are issues that deserve more consideration than what we had been able to assign to them here. An extended version of this paper is intended to mend these shortcomings.

## REFERENCES

- [1] S. Ossowski, “Coordination and agreement in multi-agent systems,” in *Cooperative Information Agents XII*, ser. Lecture Notes in Computer Science, M. Klusch, M. Pechoucek, and A. Polleres, Eds. Springer Berlin Heidelberg, 2008, vol. 5180, pp. 16–23.
- [2] M. Peterson, *An Introduction to Decision Theory*, 1st ed. Cambridge University Press, May 2009.
- [3] B. Hardy-Vallée, Ed., *Cognitive Decision-Making: Empirical and Foundational Issues*. Cambridge: Cambridge Scholars Press, January 2007.
- [4] G. Klein, R. Calderwood, and A. Clinton-Cirocco, “Rapid decision making on the fire ground: The original study plus a postscript,” *Journal of Cognitive Engineering and Decision Making*, vol. 4, pp. 186–209, 2010.
- [5] H. A. Simon, “A behavioral model of rational choice,” *The Quarterly Journal of Economics*, vol. 69, no. 1, pp. 99–118, 1955.
- [6] G. Klein, *Sources of Power: How People Make Decisions*. The MIT Press, February 1999.
- [7] L. List and B. Polak, “Introduction to judgment aggregation,” *Journal of Economic Theory*, vol. 145, no. 2, pp. 441 – 466, 2010.
- [8] C. S. Burke, K. C. Stagl, E. Salas, and D. Pierce, L.; Kendall, “Understanding team adaptation: A conceptual analysis and model,” *Journal of Applied Psychology*, vol. 91, no. 6, pp. 1189–1207, November 2006.
- [9] G. Klein, “Naturalistic decision making,” *Human Factors: The Journal of the Human Factors and Ergonomics Society*, vol. 50, no. 3, pp. 456–460, June 2008.