



A novel agent-based, evolutionary model for expressing the dynamics of creative open-problem solving in small groups

Alex Doboli¹ · Simona Doboli²

Accepted: 28 August 2020 / Published online: 27 October 2020
© Springer Science+Business Media, LLC, part of Springer Nature 2020

Abstract

Understanding the process of producing creative responses to open-ended problems solved in small groups is important for many modern domains, like health care, manufacturing, education, banking, and investment. Some of the main theoretical challenges include characterizing and measuring the dynamics of responses, relating social and individual aspects in group problem solving, incorporating soft skills (e.g., experience, social aspects, and emotions) to the theory of decision making in groups, and understanding the evolution of processes guided by soft utilities (hard-to-quantify utilities), e.g., social interactions and emotional rewards. This paper presents a novel theoretical model (TM) that describes the process of solving open-ended problems in small groups. It mathematically presents the connection between group member characteristics, interactions in a group, group knowledge evolution, and overall novelty of the responses created by a group as a whole. Each member is modeled as an agent with local knowledge, a way of interpreting the knowledge, resources, social skills, and emotional levels associated to problem goals and concepts. Five solving strategies can be employed by an agent to generate new knowledge. Group responses form a solution space, in which responses are grouped into categories based on their similarity and organized in abstraction levels. The solution space includes concrete features and samples, as well as the causal sequences that logically connect concepts with each other. The model was used to explain how member characteristics, e.g., the degree to which their knowledge is similar, relate to the solution novelty of the group. Model validation compared model simulations against results obtained through behavioral experiments with teams of human subjects, and suggests that TMs are a useful tool in improving the effectiveness of small teams.

Keywords Modeling · Groups · Problem solving · Knowledge evolution · Agent model

1 Introduction

Employing small groups to perform complex activities is critical to modern life. Due to the omnipresence of the Internet and other socially-connected environments, a growing number of applications address groups, teams, and even entire communities [1–6]. Related applications include group polling, decision making, problem solving,

engineering design, and many more [7–12]. A number of academic disciplines (mainly in social sciences) have traditionally studied group activities, more notably sociology, organizational design, psychology, learning sciences, and philosophy. Their main topics of study address group interaction, motivation, management, and leadership [13–18].

Moreover, significant transformations have recently emerged due to new computing technologies that change the way in which groups are formed, interact, and evolve. Transformations pertain to the nature of interactions between group members, the type of group memory, the context in which groups operate, the kind of cues that become available to members as a result of group participation, and group level strategies and procedures that emerge during the process.

The new ways in which humans interact with and through machines can be distinguished along five criteria. First, the nature of interactions among group members can be synchronous (e.g., members must participate

✉ Alex Doboli
alex.doboli@stonybrook.edu

Simona Doboli
simona.doboli@hofstra.edu

¹ Department of Electrical and Computer Engineering, Stony Brook University, Stony Brook, 117954-2350, NY, USA

² Department of Computer Science, Hofstra University, Hempstead, 11549, NY, USA

at the same time) or asynchronous (i.e. members can participate at different time moments), overlapping or sequential communication, through permanent or temporary media (such as depending on how responses are stored), through linearized interactions (if responses are only progressing with respect to current outputs) or cyclical (a reached consensus can be challenged by considering a different premise), absence or inclusion of personal preferences, beliefs, and emotions, and many others. Second, groups can be distinguished depending on the type of the implemented group memory, such as short term memory when participants access only a limited amount of information (like similar or preferred information, beliefs describing subjective or global preferences), and incubated knowledge (i.e. insight produced after a certain amount of times from the initial moment when the output was created). Third, the physical context in which groups operate influences if outcomes must immediately be produced (e.g., in real-time for participants in physical proximity), or if they interact through a medium that allows sampling and filtering of available outcomes. Fourth, new human - machine interaction solutions can generate a large variety of cues that become available to group members, like summarizations, trends, and content extraction through data analytics and Machine Learning (ML). Cues can be used to connect based on previous outcomes (convergence) or to expand towards previously unexplored situations (divergence). Fifth, various group level strategies and procedures can emerge depending on the group size, composition, and management. Groups do not have centralized control, however, different constraints can be set on the available resources or participants' characteristics.

Combining the options of the above five criteria produces a large variety of new group interaction mechanisms. Moreover, interactions between group members are strongly connected to group behavior and final results [4, 14, 18, 19]. Still, many open research issues remain. For example, explaining the apparent contradiction between selfish individuals and altruistic populations is under debate [19]. The causal impact of parameters on group behavior is often difficult to distinguish and characterize, especially for co-occurring parameters [19]. It is likely that multiple hypotheses can explain a certain observed situation, which demands for additional observation to further narrow down the possible hypotheses. Also, it has been explained that aggregated measures, like fitness, utility, and usefulness, which drive evolution selection, relate to basic parameters that are hard to observe, characterize, and model [19]. In general, developing more effective computing technologies for team-based applications requires understanding the connection between group interactions, group behavior, and evolution of group responses. New theoretical models are needed to capture the details of this connection.

This paper describes a novel theoretical model (TM) to describe solving of open-ended problems in a small-group setting, e.g., around five members in a group. In particular, the model refers to the problem of devising solutions to improving the American health care system, but other similar problems can be tackled too. The model mathematically describes the connection between group member characteristics, interactions inside a group, group knowledge evolution, and overall novelty of the responses created by the group as a whole. Every group member is characterized by his/her local (personal) knowledge, a way of interpreting the knowledge, beliefs, priorities and preferences, resources, motivation, and emotions associated to problem goals and concepts. The local (personal) knowledge of an agent is similar to a weak ontology [20], as it is expandable and specific to each agent. Five solving strategies can be employed by a participant to generate new knowledge that can be of different types, e.g., it refers to affordability, quality, education, and prevention, and includes different aspects of a solution, like what is the described solution (goal), how is the solution implemented to produce a desired outcome, who is performing the activity, for whom is the activity conducted, and why is the solution needed (reason). Group responses form the solution space, in which responses are grouped into categories based on their similarity and organized on abstraction levels. The solution space also includes concrete features and samples as well as the causal sequences that indicate how concepts are logically connected with each other. The nature of group interactions, like different degrees of competition or collaboration between members, emerges as a result of the strategies that are selected by members. Moreover, group knowledge evolution can be convergent, if response similarity increases, or divergent, if similarity decreases. The model was used to explain how group member characteristics relate to group solution creativity. The two main requirements of the model are having a way to decide the similarity of two knowledge items, as well as a method to describe the purpose of each knowledge token. The presented implementation of the model uses WordNet [21] for satisfying the two requirements.

The principal goal pursued by this work is to understand the dynamics of the responses to open-ended problems collectively tackled by small groups. Such situations are frequent in modern and emerging domains, including health care (e.g., teams of medical doctors must create holistic, customized treatment plans), industry, banking and investment [1, 2, 6, 10, 22, 23]. These areas include situations in which having global control based on global data is difficult. Several independent studies have recently argued that future careers will seek creative problem solving skills as well as capabilities to work effectively in teams. Studies suggest that team organization and management

will be increasingly important, including the management of softer individual skills, such as participant's expertise and experience, personal beliefs, motivation, and resources. Such skills can be both implicit or explicit, can change during problem solving (i.e. motivation might increase or decrease depending on the group dynamics), and can be partially observed and learned during the process. Another application refers to markets with minuscule or even without price [24], like open source software development. These markets pay less emphasis on monetary retribution, but stress motivational and emotional rewards to the participants. The paper argues that future markets will redraw the balance between rationality vs. emotions and motivation in decision making as hard-to-quantify incentives will be increasingly common.

This paper suggests that models describing group-level problem solving is a corner stone for a software environment to manage and optimize group behavior and human-machine interaction. The system would continuously observe knowledge evolution, group interactions, and participant's motivation, emotions, and resources. The system could improve group creativity by recognizing the opportunities unused by the group (opportunities that should be further explored), reorganizing and rebalancing the interactions between team members, and gaining new insight on how to solve a problem. Insight gaining include understanding the causality of trade-offs and limitations (bottlenecks) of present solutions, identifying solutions that could tackle the causality of bottlenecks, and how the interaction between group members influences their beliefs, motivation, emotions, and decision making (including how certain cues offered by others influences a participant's decisions). The computational model supports defining new metrics and data analytics that can be used understand how group level parameters can be managed to optimize solution effectiveness, or to predict expected outcomes for a group.

The paper has the following structure. Section 2 presents the tackled problem and summarizes related work. Section 3 discusses the proposed TM. Section 4 details knowledge self-organization during problem solving. Section 5 describes the software implementation, followed by experiments. Section 7 discusses the results. Conclusions end the paper.

2 Problem description and related work

The considered group collaborations have the following characteristics. Participants are required to solve an open-ended problem about which each of them is knowledgeable, e.g., a problem of general interest. Open-ended problems can have multiple solutions that differ by their characteristics, like novelty, usefulness, and quality [25]. Then the

goal is to maximize specific attributes of the group solutions (outcomes), e.g., solution novelty. All groups include a small number of participants, such as between four and six persons. Participants present their solutions through a shared environment that stores all solutions for an unlimited amount of time. Each participant can access all other group members' solutions. At every step, a member can propose a new solution or expand the solution offered by another group member. Participants interact asynchronously, meaning that they can present their solutions at any time without requiring the presence of others. There is no limit on the number of the solutions by a participant, and there is no time limit set for the time used to create an outcome. Groups do not interact with each other. Independent raters evaluates the characteristics of the found solutions.

Example . A typical case study represents the following experiment performed at University of Texas at Arlington [26]. Twenty groups participated to the following experiment. Each group had five members. Groups were asked to solve the following, general domain, open-ended problem: "You and your team members have been elected by the people of your country to serve as members of a committee. Your task as members is to generate as many ideas as possible that might help reform America's healthcare system". Participants were instructed by the authors of the experiment as follows: "Please list as many ideas as you can think of. Your ideas can be about anything related to the Healthcare system including finance, education, insurance, technology, mental health, long term care, pharmaceuticals, emergency room services, preventive care, health care providers etc. These are just some examples of different areas that you could think about. You can explore these areas along with any others that you can think of. Please do not restrict yourself to any one aspect of healthcare." Participants posted their solutions using an Internet-based environment that stored and displayed to the rest of the group all present solutions and their anonymized authors. Participants had the option to create a new response or to cite a previous answer when generating a new solution. The identity of a participant was unknown to the other group members. The novelty of the solutions was rated by a group of independent raters.

For this experiment, participants produced answers like the following responses [26]: "Have more research on mental health"; "Pharmaceuticals that can be used for something illegal should be carefully distributed"; or "Instead of having doctors volunteer at the facilities and clinics that accept no insurance. Why not pay doctors for their services so they will care more about the people that cannot get insurance". Participants could cite previous responses to propose a new solution. For example, the solution "Better technology provided in smaller clinics" was

cited when the following solution was created: “Having more smaller clinics being able to provide the same technology like bigger clinics”.

The presented work devised a theoretical model (TM) that declaratively connects the elements and parameters of a group to those of its participants, so that the group’s behavior and knowledge evolution can be explained and compared to other groups. TM parameters can be identified by minimizing the error to the observed experimental results. Compared to other models on group behavior, TM presents the evolutionary process through which new solutions were devised by a group. Knowledge evolution depends on group interactions, individual behavior, and cognitive processes grounded in knowledge semantics. The traditional concept of fitness was revised to express knowledge evolution.

2.1 Related work

Open-ended problem solving has been the focus of computational creativity [27–31]. Related systems implement different creative paradigms, including analogy, concept combinations, domain-specific rule-based reasoning systems, cognitive architectures, and transformational and evolutionary systems [27, 31, 32]. The work in [28] uses ontologies to deconstruct and reconstruct inputs into features and structures during conceptual recombination. Evolutionary algorithm, including Genetic Algorithms (GAs) have been suggested for creative applications, like composing music, generating stories, drawing, and engineering design. New electronic circuits were created using GAs [33, 34]. However, the creativity of such design is much lower than that of circuits invented by humans [35]. A recent semi-autonomous system suggests creating new recipes or menus by mimicking human creative activities, like finding a problem, acquiring a problem, gathering information, incubation, generating ideas, concept combination, selecting the best ideas, and externalizing solutions [31]. Solutions are described as schema that express ingredients and steps for cooking a dish. Our recent work proposed InnovA, a cognitive architecture for computational innovation through robust divergence [29]. Similar to [31], InnovA is inspired by human cognitive activities. However, it uses a different knowledge representation in which concepts are associated based on their similarity, and causally linked in sequences depending on their purposes in the final solution. Creating new solutions involves identifying new building blocks [36], symbolically understanding the similarities and differences between concepts [37], generalization, causal reasoning [38], and various types of concept combinations [39, 40].

In addition to the above creative activities, open-ended problem solving in teams, groups, or large communities

includes two other facets [41, 42]: (i) the emergence of insight by a participant or few participants possibly due to cues produced by the group, and (ii) the readiness of the group to understand, propagate, and embrace the new insight. Crosson and Grachter summarize existing theories that span from psychology, e.g., models of individual cognition, to models for economic behavior [41]. From a cognitive point of view, insight is a sudden restructuring of the problem representation. It is likely to represent a discontinuity (divergence) in the space of solutions uncovered by an individual or a group. Even though a person is convinced about the correctness of his/her insight, in a small group setting, it is possible that the group rejects an insight that does not resonate well with the group ideas and opinions, even though the individual might be passionate about the new solution. It has been explained by research in organization science that teams in which members share information about pursued goals and solutions are more efficient in reaching solutions than teams in which members do not interact with each other [42–44]. Creating a shared memory model between participants has been an important objective in team training, e.g., teams in commercial aviation [42, 43, 45]. Alternatively, the occurrence of discontinuities in large communities, like paradigm shifts, has been a main recent research topic in sociology and social networking with the main goal of understanding technological trends and which outcomes are more likely to produce a significant impact [6, 8, 16, 46–48]. There is a consensus among researchers that knowledge dynamics in a group or community includes a sequence of four stages: stagnation when little new information is created as current barriers are difficult to overcome, marginal and incremental change of the current paradigms, short periods of radical changes when new insight and solutions become available to the group, and swift adaptation of the new yet superior paradigm [49, 50].

As summarized later in Figs. 1, 4, and 5, from a theoretical point of view, models are at three abstraction levels: the bottom level uses agent-based models to study individual-related aspects (parameters), like the importance of diversity (knowledge, knowledge interpretation, strategy, and rewards) on team performance [2, 50–52]. The second level aggregates the individual outcomes to create cumulative, stochastic or differential algebraic models typically used to describe a group’s dynamics [14, 16, 53]. Finally, the most abstract models offer an interpretation of the group behavior as black box models that indicate how the parameters of the dynamic process influence the behavior (i.e. properties) of the group [42, 44, 54]. Often, such models do not indicate causal relationships, like the parameters that determine a certain group-level outcome. Agent-based models conduct an optimization process in which every agent explores locally, while global

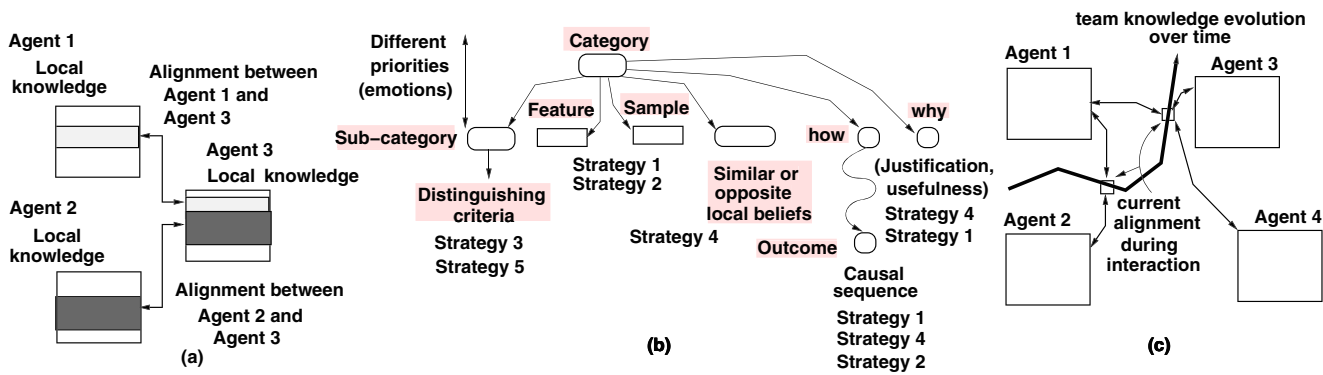


Fig. 1 (a) Interaction between individuals, (b) local knowledge structure, and (c) group knowledge evolution

search is achieved through interactions among agents. Social network-based statistical and topological analysis uses aggregated models to understand static and dynamic properties of a group, including which parameter values of the model decide optimized properties (behavior) of the entire population. A summary of the three categories is offered next.

1. Agent-based models study group behavior, including problem framing and solving as well as propagation of new knowledge, by setting up an agent model, in which an agent expresses a group member or a sub-group (i.e. a team) [55, 56]. An updated review of related work is offered in [57]. Agents perform with certain probabilities a set of activities, like sampling the problem space, sampling the solution space, experimenting new solutions, and interacting with other agents. Agents are rewarded depending on the activities of the other agents as well as the value (utility) of their outputs for the entire group. The expressed parameters of the agent models, as well as their level of details and connections to other parameters, are described depending on the purpose of the models. Early work by Lakatos and Feyerabend considers that paradigm shifts occur only after sufficient accumulation of complementary innovations [58, 59], even though the impact of individual innovations or the quality of the overall progress are hard to evaluate [58]. Understanding the nature of the problem space, e.g., technological landscape, is a difficult yet critical problem, as problem framing decides to a large degree the usefulness of the problem solving activity. Sickert et al. consider a simple description of the technological landscape as quadratic polynomials with local and global optima [50], but this assumption might be less effective for cross-disciplinary problems, for which such unifying functions are hard to identify. They propose an agent model to study the social dynamics of paradigm shifts. Agents learn the characteristics of physical-technological landscapes by attaching weights to their individual search vs. *social influence* by

other agents [50]. Computationally, social influence is a means to aggregate the impact of agents on other agents, thus represent a self-organization feature through which agents are influenced by the group [56, 60]. Each member assigns specific weights to each of the peers. The model is capable to produce dynamic situations that correspond to incremental changes of the solutions as well as paradigm shifts around disruptive ideas. Experiments indicate that strong social influence is a barrier in front of paradigm shifts. Hence, a lower social influence is needed to propagate knowledge from the surrounding peers, beyond the local technological landscape available to an agent, while enabling the propagation of disruptive ideas, like sudden insight. Also, large communities with strong social influence between their members are less likely to adopt innovative changes. Hong and Page show that agents with diverse problems solving perspectives and heuristics (selected randomly from a population) produce better solutions than a team of best performing agents or than the most effective agent [2]. In addition to the support offered by an agent-based model, the authors also prove mathematically the result. The diversity of the agents' local perspective enable reaching of better optimal solutions than if only high performing but similar agents are being considered. Recently, Gabora and Tseng studied the optimized ratio among imitators and creators, the degree to which social regulation of the creativity level of an individual helps the success of the entire population, and the usefulness of chaining simple outputs for solving open-ended problems [52]. Other agent models focus on the effect of complementarity among group members, populations sizes, minimum number of required pioneers, impact of the shapes of technological landscapes [61–63], individual adaptation strategies to landscape changes [50, 64], dynamic knowledge representation [65, 66], and conceptual reorganization of individuals' knowledge representations [67]. We previously studied the impact of different problem framing strategies and possible resource allocation schemes on the successful evolution of a community [51].

2. Social network based statistical analysis assumes that there is a connection between the topology of the network representing the social interactions of a group or community and the behavior or role of its members for the entire group [53]. Then, insight about the network static and dynamic behavior can be inferred from the metrics computed over the topological properties of the network. Seminal work by Granovetter suggests that weak ties in a group are important for increasing creativity as they serve as bridges between clusters of members with similar ideas [14, 68]. The strength of a tie depends on the amount of interaction, reciprocity, and emotional intensity. Perry-Smith and Shalley propose a social network model in which the peripheral nodes with many connections outside the network are more likely to acquire unique insight and generate more groundbreaking contributions [53]. However, as their contributions are accepted by the community, and thus peripheral nodes are more linked through new connections, their centrality within the group increases. The behavior of the peripheral nodes is called *spiraling model*. The authors explain that weak ties improve creativity as they enable access to diverse perspectives and approaches to problems, different outlooks, various interests, access to a wider group of people, and utilization of less redundant information [69]. Weak ties are also the most effective way of spreading rumors [14], but also help in validating alternatives as members identify less with the ideas and beliefs of a certain group [53]. In conclusion, the spiraling model suggests that many weak ties and few strong links help creativity, however, beyond a certain threshold, weak ties reduce focus and attention by introducing too much distraction [70].

While weak ties seem to be important for instrumental purposes when trust is not critical [71], strong ties are necessary for problems that involve complex and proprietary information, and for which trust and affect are needed between group members [72]. Strong ties encourage members to share information in competitive environments [71]. In recent work, Montjoye et al. show that strong ties are more correlated with the performance of a small team (e.g., four persons) than competencies and personality [16]. The studied teams were working on problems like high-level programming, data modeling, and simple ML algorithms. In this work, the strength of a tie is measured as a ratio between how often two persons were in each other's vicinity and the total number of participants. Strong ties are likely to be important because they enable the access to the needed piece of information available for the entire group. A limitation of statistical analysis of social networks relates to the capability to relate various topological metrics of the networks to group behavior. Traditionally, for small teams, this link has been realized by using surveys in which participants individually answered

to questions describing the components describing ties [14, 16]. However, these approach does not easily scale to very large groups (e.g., thousands of participants), and also it is not clear how group level self-structuring processes relate to the concept of tie strengths. Moreover, the causality of creating ties of various strengths is not fully understood.

Statistical models have been proposed to understand social networks [73]. A social network is expressed as a dynamical, connected system with unknown states. States are identified based on a set of observed outcomes (e.g., signals) over time [73]. For example, Pan et al. propose the influence model in which conditional influence distributions between a state and its neighbors are computed over time using hidden Markov models. Influences model the strengths of interactions between neighbors. The social network dynamics is described by the changes of influence distributions, and have been used to study the unknown social structure of a group, the functional roles of participants, and traffic patterns [73].

3. Models for the design of group mechanisms. This category of research has studied the way in which the design of a certain organization influences the behavior and performance of individuals and organizations. Classical research has shown that high performance teams use different adaptation strategies than low performance teams in the case of high demand situations [42]. Adaptation can include different aspects, like the team composition, coordination strategies, decision making methods, and interaction between team members. Pursued work is often experimental. For example, the shared situational mental model considers that the shared knowledge model describes a common model about the problems as well as a shared model about the environment, including the team members, their abilities and expected reactions [42]. An optimized interaction strategy senses and anticipates the needs and responses of other members due to the problem changes. Interaction and coordination mechanisms are both explicit and implicit, however, in the presence of increased stress and constraints, more effective teams show an increased level of implicit coordination through the shared mental models. Along the same lines, Wohl presents the stimulus - hypothesis - option - response model, in which teams alternate between implicit and explicit interaction, periodic explicit communication being required to update and reinforce the shared mental models of the participants [43, 74]. Similarly, Salas and Cannon-Bowers consider that an optimized organization has superior sensitivity and responsiveness to changes [44]. The rate of correct responses generated by a team increases with the upward-anticipation ratio, which is the ratio between the number of information transfers from subordinates to the leader and the number of requests by the leader. Shared

mental models reduce the number of errors committed by a team as each member offers a different views on the causal explanations of certain errors, which leads to better prediction of the effects of faults as well as the capacity of the group to solve the errors [45]. The model of *zone of proximity facilitation* considers that teams of experts (with superior background knowledge and interaction skills) perform better than expected for complex problems in aviation, but perform worse than expected for simple problems [75]. On explanation, is that experts interact less for simpler problems as they try to solve them individually.

Work has also studied the mechanisms through which certain patterns and self-organization behavior emerge due to the characteristics of the members as well as the global policies of the group. An important topic has been the emergence of information or decision cascades in a group [54]. Cascades describe situations in which individuals decide mainly based on others' previous decisions than based on their local information. Cascading produces superior decisions than individual decision making, but can also result in poor production and performance of a group [54]. An experimental study focused on the relationship between emergence of information cascades and the type of the institutions [54]. Three institution types were considered, individualistic institutions in which agents are rewarded based on the quality of their announced decisions, majority-rule institutions, in which agents are rewarded depending on whether the majority decisions were good or bad, and conformity-rewarding institutions in which agents are rewarded depending on their similarity with the majority. The study considers that individuals can make three types of decisions: Bayesian-Nash equilibrium behavior, private information revealer, public information seeker, and strategic player. Experiments show that cascades are more likely to emerge if rewards are made based on announced decisions. Also, cascades are more likely in individualistic organizations than in majority-rule organizations, and also more likely in conformity rewarding organizations [54]. The theoretical model supporting the work is based on General Decision Weight model, which is a weighted sum of the public announcements and private knowledge, a model which assumes a naive Bayesian interpretation of decision making.

3 Model description

This section presents the proposed theoretical model (TM) for describing knowledge evolution in small teams. Compared to other models, TM focuses on presenting bottom-up the connection between the evolutionary process through which new knowledge is created by a team, and the individual characteristics of team members (agents).

Figure 1 summarizes the three main components of TM. As shown in Fig. 1(a), the characteristics of team interactions depend on the alignment between agents, e.g., the similarity of their local knowledge (i.e. previous experiences) and exchanged cues (e.g., information, social and emotional cues). Each agent uses its local knowledge (Fig. 1(b)) to create new knowledge, i.e. solutions to the problem that is jointly solved by the team. The new knowledge is produced through activities, like top-down and bottom-up reasoning, implicit and explicit knowledge, and reusing previous solutions for new needs. The group knowledge evolves as a result of the new knowledge, as shown in Fig. 1(c).

TM uses a declarative description similar to [76] among others. Declarative equations express the connections between related parameters, but do not specifically indicate how relations are actually approximated by more-detailed sub-models. The goal of this work was to offer a comprehensive description of the aspects that influence team knowledge evolution. However, to explain its usefulness, Section 5 offers an implementation of TM, but other realizations of the model are possible too. TM predictions were related to team behavior observed in experiments.

The remaining part of the section offers a bottom-up presentation of the TM components: (i) solution space, (ii) individual agents, (iii) group behavior, and (iv) team knowledge evolution.

3.1 Solution space description

The description of the solutions created by a group for a problem must be grounded in the semantics of natural language, express the explicit and implicit nature of knowledge, and support identification of the cognitive activities pursued during problem solving, like top-down and bottom-up reasoning. Grounding in semantics is achieved for the proposed TM by indicating the role (meaning) served by different parts of a solution as characterized by WordNet [21], like the words expressing the solution and the relations among parts. WordNet is a database that stores the lexical components of a language (e.g., words), and their semantics expressed through attributes, i.e. typicality, membership, synonyms, polysemy, and so on. The nature of knowledge and pursued cognitive activities is presented through solution clustering by different criteria. Group behavior and solution evolution are then expressed based on the changes in solution clustering.

Frame descriptions (FDs) present the role (meaning) of the parts of a solution (e.g., a text response in the considered experiments). A separate FD is constructed for each response. A detailed presentation of FDs has been offered in [10].

- A frame description (FD) is illustrated in Fig. 2(a). Each FD includes one or more frames. Each frame has two nodes that are nouns in the response. They are linked by an arc labeled with the verb that relates the two nouns. The first noun might be linked through the verb to a list of nouns that are connected in the response through keywords “and” or “or”. A negation of the noun, expressed in a FD by node “not”, is defined by words with negating semantics, like “not”, “but”, “no”, and so on. Nouns can be characterized by a number of qualifiers, like adjectives.
- Frames can be hierarchically connected in a FD, like in Fig. 2(b). Hierarchical descriptions are used to provide additional details about nouns that are part of higher level frames.
- Figure 2(c) shows the different kinds of frames that can occur in a FD: (i) frames for goals express the goal of a solution, (ii) frames for outcome present the expected output (result) of a solution, (iii) frames for implementation indicate a way of creating an output, like a partial or complete solution to generate the output, (iv) frames for justification describe reasons supporting a solution, and (v) frames for reinforcement express reasons, like disfavoring conditions, in spite of which the solution still holds. The nature of frames is indicated

by specific words, like “so that” introduces a frame for outcome, “since”, “so”, and “because” indicate a frame for justification, and “matter” defines a frame for reinforcement. Not all frames must occur in a response.

Example . The response “Improving quality by using the latest technology to decrease the cost in most hospitals” includes the following frames: “improving quality” represents the goal of the solution. Part “using the latest technology” indicates the proposed implementation, and “decrease the cost” is the expected outcome (output) of the solution. Part “most hospitals” indicates the place of the outcome.

The frames for goals and outcomes (frames (i) and (ii)) are further divided into three kinds, as shown in Fig. 2(d). Frames of type *what* express what specifically the goal aims to achieve. For example, a frame would state that “Healthcare must be available to everyone”. The frame indicates a concept related to healthcare and a feature (e.g., becoming available) suggested for the concept. Second, frames of type *who* present the persons or institutions to create a certain feature of the healthcare-related feature. For example, a feature states that “Doctors must offer more compassionate treatment”. Third, frames of the type *for who* indicate the beneficiary of a certain improvement. For

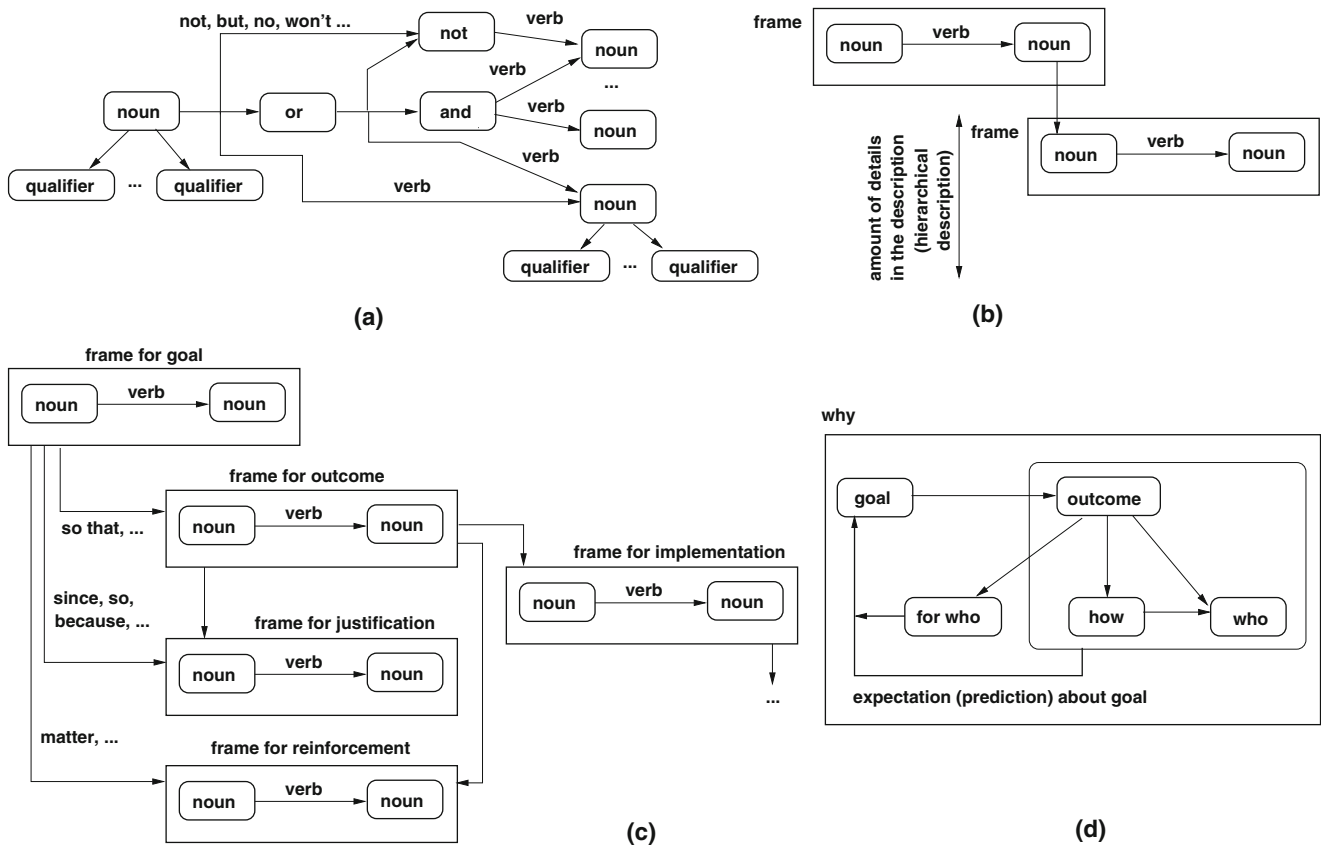


Fig. 2 Frame descriptions: (a) frame structure, (b) hierarchical descriptions, (c) frame types, and (d) spanned space

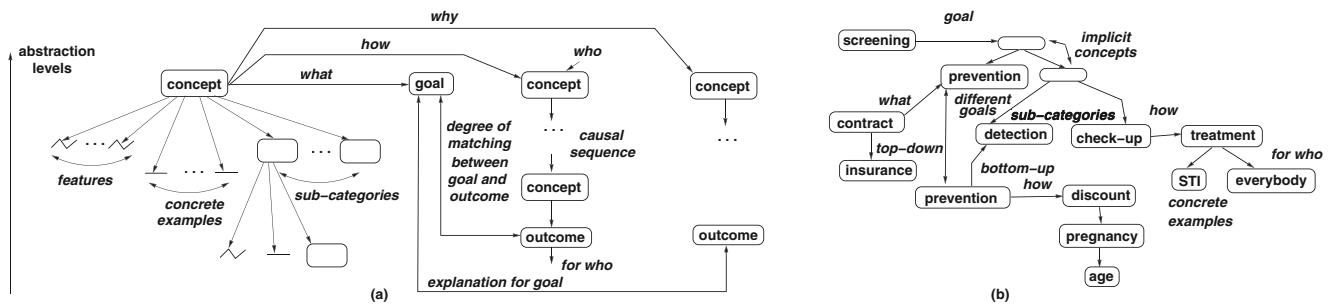


Fig. 3 (a) Solution space description along a hierarchy of abstraction levels and (b) an example

example, “Low income people will receive free healthcare”. Figure 2(d) depicts the three kinds, and their relationship to goals and outcomes. Each frame kind refers to a different dimension, but dimensions have the resulting *outcome* as a common reference.

The following equation relates the description of a solution outcome to the way in which the outcome is realized. It states that within the context of the solution (e.g., the reasons expressed by frames *why*), the *what* frame with its associated features and beneficiaries (i.e. as described by frames *for who*) is a declarative description (model) of the realization as presented by frames *how* composed with the related frames *who* (the agents carrying out the activity described by *how*). Operator \circ connects the different frames together.

$$\begin{aligned}
 & \text{what} \circ \text{outcome}^{(1)} \circ (\text{for who}^{(1)} \dots \circ \text{for who}^{(i)}) \dots \circ \text{outcome}^{(k)}(\dots) \models \\
 & \text{how}^{(1)} \circ (\text{who}^{(1)} \dots \circ \text{who}^{(m)}) \dots \circ \text{how}^{(n)}(\dots) \mid \text{why} \quad (1)
 \end{aligned}$$

Example : Let’s consider the following two responses that refer to a common goal. The first response states that “Elderly people should receive free healthcare from the federal government.” The second response states that “Elderly people should be able to see any doctor”. Both responses share a common *what* (to improve access to healthcare) as well as a common *for who* (elderly people). Thus, the two responses have the same left side of their (1). However, their right sides are different, as they refer to different *how* and *who*, e.g., two different ways of improving healthcare access.

Equation (1) is created using multiple responses that can be possibly offered by multiple members. Not all equation components are always present. For example, the responses offered by a group might include only fragments of the right side, i.e. how to improve healthcare, or fragments of the left side, e.g., for who are healthcare improvements meant. If only the right side is present (hence, information on *what*, outcomes, and *for who* is missing) then another group member can produce expectations about the elements of the left side, even if these predictions are not explicitly

stated. A certain implicit variety results as every member has different expectations. A similar implicit variety exists for expectations for the right side, if only the left side of the equation is present.

The five types of frames (*what*, *who*, *for who*, *how*, and *why*) are further classified into four types depending on the nature of the outcome. For the studied healthcare problem, the four kinds are as follows: (i) *Affordability* refers to improving access to healthcare (e.g., funds, costs, insurance, etc.). (ii) *Quality* presents responses about increasing the quality of healthcare (like the degree to which it satisfies the needs of the population). (iii) *Prevention* includes responses on preventing health issues (e.g., health life time, immunization, etc.), and (iv) *Education* has solutions on educating people on health related topics. The four categories are related to this application. In general, for any application, responses can pertain to the following four categories: (i) *providing the resources* needed for improvement, (ii) *nature of improvement*, such as what is the exact improvement, (iii) *avoiding* the need for the specific improvement, and (iv) *learning* about the specifics of the tackled problem from previous experience, so that future improvements are optimized.

The frames of the responses created by participants during problem solving form the evolving solution space. The solution space is represented as shown in Fig. 3(a). The representation is a network of concepts corresponding to a hierarchy of abstraction levels. Every concept might have associated features (attributes), concrete examples corresponding to the concept, and sub-categories that share the same characteristics as the overall concepts but also have their own distinguishing elements. A concept might be connected to a goal meant for the outcome (*what* frames), an implementation expressing a set of concepts connected through *how* frames (causal sequences), and a set of concepts in *why* frames (reasons for a goal or implementation). A mathematical description of the concept representation was presented in [10, 77, 78]. The novelty of a response is assessed using a weighted sum, in which goal and *for who* have the highest weights, features have middle weights, and the structure of the causal relations, their

nature, and kinds of concepts used in the implementations have the lowest weights [25, 79]. Predictions through weighted sum were calibrated using human ratings and they match well those ratings [25, 79]. Agent response novelty was also evaluated by comparing solution characteristics with the characteristics of solutions rated by human raters and produced by creative teams during experiments with human subjects, e.g., solution diversity, and effect of weak ties on similarity distances between the concepts in the solution space representation [53, 69].

The novelty of a response is assessed using a weighted sum, in which goal and *for who* have the highest weights, features have middle weights, and the structure of the causal relations, their nature, and kinds of concepts used in the implementations have the lowest weights [25, 79]. The predictions through the weighted sum matches well human ratings [25, 79].

Example . Figure 3(b) illustrates a simple example for the solution space description. Let's assume that the following responses are created in this sequence, R1: "screening for prevention and detection", R2: "prevention of pregnancy using age discounts", R3: "insurance contracts should include prevention". The frames of the four responses are gradually clustered depending on the similarity of their frames of the same kind. Top-down and bottom-up reasoning is performed as new responses are created. Besides, similar concepts representing similar frames in the responses represent alternatives of the same, more general concept, e.g., *detection* and *checkup*. Both are sub-categories that can be further detailed through top-down reasoning. Also, concrete examples are connected to more abstract concepts, like *STI* to *treatment*. Different *what frames* can correspond to the same goal, like *contract* and *screening* to prevention. Finally, implicit sub-categories can result when less similar concepts serve the same purpose (e.g., correspond to the same frame), like *prevention* and the sub-category corresponding to *checkup* and *detection*. The proposed TM describes in Section 3.4 the group behavior by indicating how the solution space description is changed by subsequent responses. Section 6 illustrates the solution space created during behavioral experiments with human subjects.

We have previously used the presented solution space description for expressing the solution space for different open-ended and design problems, such as circuit design [10, 77, 78, 80]. Our solution space representation includes elements similar to prototype schema [81], but emphasize the connection between concepts based on their roles in solutions. The representation includes causal association sequences, which show how frames with different roles produce a solution. Sequences describe the reasoning process of a team during problem solving, and are important

to highlight solutions that are less similar to the previous cases. Other related but less similar representations include a metadata representation for case-based reasoning [82], a network representation to highlight the activities performed during problem solving [83], patterns used in concept generation [84], ontologies [85, 86], semantic web [87], and various rule-based systems [88]. Frame descriptions (FDs) are relational descriptions between the elements of a response, similar to other relational notations, like relational schema [89]. Descriptive attributes and slots in relational schema focus mainly on the roles (purposes) of the elements. FDs also present how roles are achieved, such as suggested implementations.

3.2 Individual behavior

Figure 4(a) illustrates the model of an individual's (agent's) behavior in a team. Its software implementation was shown in Fig. 6. The model describes an agent's capability to produce a new solution with its current resources. The capability is a trade-off between the expected benefits by the agent and the required resources, including its motivation. Among other models [55, 57], the presented agents are most similar to the model of Budaev et al. [90]. Their rules are selected by an agent using an individual genome. Instead of a genome, our agents have local knowledge, which represents an individual hierarchy of concepts linked in a causal sequence, like in Fig. 3(a). The local knowledge models the previous experience of an agent, i.e. its previous responses. Maximizing the expected benefit or minimizing the cost of a new response balances between divergent and convergent responses, e.g., responses that are less or more similar to previous solutions.

The model is as follows. The local knowledge of an agent is his/her experience, beliefs (e.g., the way of interpreting certain situations), and priorities. Cognitively, the local knowledge corresponds to an individual's memory [91]. The local knowledge is a collection of interconnected networks, each network including concepts, goals, outcomes, expected impact on the group, and associated emotions. It has been shown that some concepts suggest a certain outcome and create enough motivation to produce a response [25, 79]. Figure 1(b) shows the structure of the local knowledge of a group member. It includes concepts in responses grouped in categories similar to the models discussed in [10, 77, 78, 80]. Categories can include sub-categories, which inherit the features of the more abstract category, but also have distinguishing parameters and constraints that separate it from other sub-categories pertaining to the same category. Each category can have multiple features (attributes) and several specific samples (exemplars) belonging to the category. Categories might also be related to other semantically similar categories (synonyms) or semantically opposite

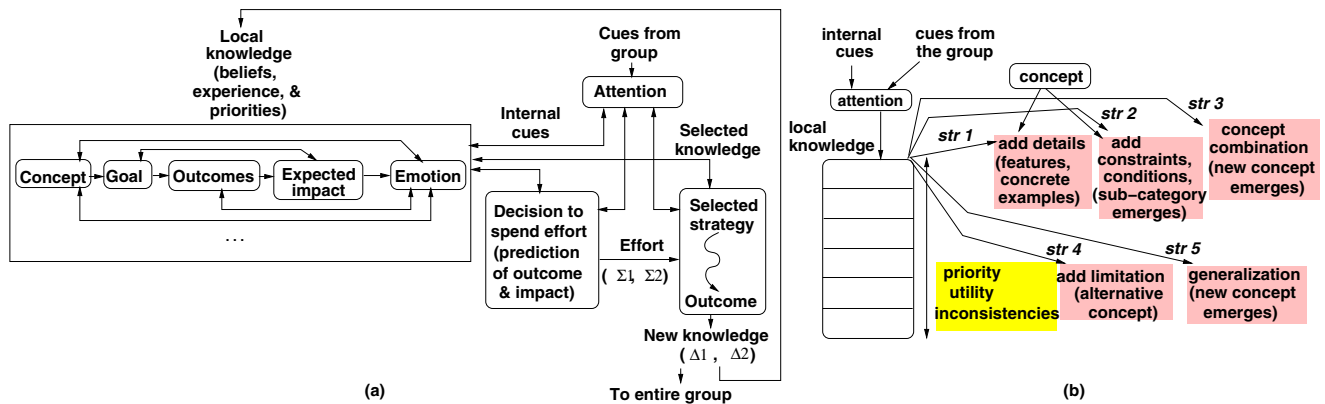


Fig. 4 (a) Model for individual behavior and (b) extension of the solution space by an individual (agent)

categories (antonyms). Every category can have a *how* component that present the way in which the outcome is created by a causal sequence. Categories might also have a *why* element which justifies the reason for certain features, outcome, or *how* component. Justification relates to a category's usefulness. Categories are organized into the overall local knowledge based on their interpretation, priority, and attached emotions. The software implementation in Section 5 uses data tables `cognition_table` and `emotion_table` to store an agent's local knowledge and attached interpretation and emotional content, respectively.

An agent receives two kind of cues: internal cues based on the local knowledge of the individual, and external cues based on the responses from the group. In the model, cues are frames from a response. Some cues capture the attention of the individual. Certain knowledge is selected from the agent's local knowledge by a cue. Next, the individual estimates the expected output and its impact on the group (hence, the rewards to be received by the individual from the group), and decides if a certain amount of effort should be utilized to produce the envisioned response. The selected knowledge, external cues, and selected strategy generates a response that is new knowledge for the group. Responses have two components: component Δ_1 is common knowledge (already existing), and component Δ_2 is the new knowledge created by the solving process. Component Δ_2 receives a certain attention (i.e. degree of surprise), which connects to a certain emotion and motivation to consider the cue for creating a new response. The new knowledge is added to an agent's local knowledge and also propagated through the group through mechanisms, like dissemination, visibility, and prestige. The software implementation depicted in Fig. 6 and discussed in Section 5 uses module `Observe differences to experience` to compute components $\Delta_{1,2}$ in the model. Module `Model attention to selected experience and specific words` implements the attention received by cues.

As shown in Fig. 4(b), an individual utilizes five strategies to solve a problem (labeled as str_i) [39, 80]: Strategy 1 creates an incremental modification of a response by adding new features to an existing response. It adds more details (e.g., features and concrete examples) to an existing concept, implementation (*how* frame), or reason (*why* frame). The causal sequence and nature of the response does not change (e.g., its sub-category, type, and kind). Strategy 2 combines an existing response with concrete features from another kind of response without changing the sub-category, type, and kind of the first response. It creates a sub-category of an existing concept by adding new constraints or conditions to an existing concept. Strategy 3 combines a more abstract response from a category with more abstract response from another category, thus creating a new sub-category. Strategy 4 rejects an existing response arguing for another approach. It introduces an alternative concept in response to a limitation of an existing concept. Strategy 5 generalizes a single or a set of responses, hence a new sub-category is created starting from a more concrete outcome.

Section 5 and Fig. 6 discuss the using of the five strategies in the model implementation.

The application order of strategies influences the degree to which attention, emotion, and motivation is aroused. For example, a repeated application of Strategy 1 might add a sequence of new features to an existing response without significantly changing its characteristics. As successive responses are similar, it is likely to reduce the motivation of a participant to further apply the strategy. Applying Strategy 3 can identify a new sub-category that is emotionally loaded for an agent. It can trigger enough motivation to create new responses. Similarly, if a high emotion is associated to a certain concept, Strategy 4 increases an agent's motivation to participate to a contradictory argument, even though little new knowledge is being generated in the process. Module `Estimate emotion` in Fig. 6 models agent emotion.

Agents select the next strategy based on a trade-off between expected benefits and effort. The trade-off is characterized by two variables: (i) similarity of the local knowledge of two group members, like the distance between the used concepts, and (ii) similarity in their interpretation of the knowledge, including priority, motivation, and emotion. The first variable also relates to usefulness of the concepts, the gap in local knowledge, the amount of overlaps between local knowledge, the similarity of goals, and the similarity with previous work. The second variable is also linked to prestige, visibility, and amount of effort already put into work. Modules `Estimate effort to create outcome` and `Estimate reward of outcome` in Fig. 6 implement the estimation of effort and benefits, respectively.

Based on the two variables, there are four types of groups: The first kind includes groups with members having similar local knowledge and similar interpretation. Such groups are likely to have strong interaction among members, and repeated application of Strategy 1 through which new details are added to the similar responses. The second kind is groups in which members have similar local knowledge, but which are interpreted differently. These members might still interact a lot, but mainly through Strategy 4, as their views are opposite. The third kind contains groups with more dissimilar local knowledge, but similar interpretations of the fundamental concepts. Such groups interact less unless they bridge their less similar local knowledge. Members are likely to apply Strategy 3 but also Strategy 1. Finally, the fourth kind refers to groups with dissimilar local knowledge and dissimilar interpretations. Such groups have little interaction. Members are expected to perform Strategy 3, Strategy 2, and Strategy 4. The two variables are likely to be aggregated as a single parameter, like in one-reason decision making [92].

3.3 Interactions between individuals

During problem solving, group members interact with each other by presenting responses to the group and then being offered responses by others. Every member can decide to ignore responses, to read them, and to consider them in a response that further elaborates the message. In addition, a member must decide which of the responses from those available it will consider for elaboration. For example, in Fig. 1(a), Member 3 must decide if he/she will interact with Member 1 or with Member 2. This is a complex decision. For example, Member 3 might decide to elaborate on the response of Member 2, as it seems to be more interesting and/or easier to extend. However, an elaboration of the response by Member 1 might extend a response that has a greater flexibility (potential) to lead to novel responses. Member interactions continuously tackle the

trade-off between local benefits of individuals and global benefits of the group.

Group interaction is modeled as a continuous process in which mini-groups are being *ad-hoc* formed between the members that use each other's ideas. There is no other definition of a mini-group than the fact that one person uses the responses posted by another member. Figure 1(c) illustrates this process. The existence of the interaction indicates a certain *alignment* of the local knowledges of the involved members meaning that there is a certain similarity between their knowledge items (concepts, goals, features, outcomes, and causal sequences) while they might have similar or opposing interpretations that, however, create enough motivation to a group member to select and elaborate a certain response. Still, the creation of a specific alignment prevents the production of another alignment, thus a certain direction of evolution is not pursued by the group. For example, two members that interact through repeated Strategy 4 create a sequence of repeated contradictions of each other's responses, even though another alignment would steer knowledge evolution along a more productive direction. Then, group problem solving is a process representing the *sequence of interactions* between some participants (called mini-groups) during which their local knowledge aligns, and responses are produced that can serve as cues to illuminate new knowledge items (which otherwise might be hard to access). Optimizing group problem solving requires determining the best sequence of interaction among members (or the best sequence of mini-groups) to create responses of the highest novelty. Alternatively, the trajectory of group knowledge evolution must be found, so that the beliefs of members are sufficiently closely aligned for them to interact which exchanging ideas.

During interaction, certain concepts (e.g., words) of the responses raise the attention of other group members and carry enough emotion to motivate them to produce an answer (module `Model attention to selected experience and specific words` in Fig. 6). A member might drop his/her current flow of responses. Hence, a possibility to improve group interactions is by dynamically identifying which concepts (words) are more likely to produce a response from others due to their associated emotion and motivation. Also, during interaction, certain causal relations between concepts and outcomes might be pursued by a participant, even though the causal relation does not capture the entire trade-off between the parameters involved in the response. Interaction with another member during which Strategy 4 is repeatedly applied might expose 'hidden' trade-offs.

Group behavior is described by the responses generated by member interaction as described by the following

equation:

$$\Delta(\text{how}, \text{who}) \approx \Delta\text{how} + \Delta\text{who} + \Delta\text{how} \Delta\text{who} \quad (2)$$

The equation states that the overall changes when concepts *how* and *who* change depend on the modifications of the causal sequence *how* (Δhow), concept *who* (Δwho), and the relation (combination) of the two changes. Component Δhow depends on using new building blocks (BBs), new ways of connecting the BBs to the outcome, and new features that are introduced. Component Δwho suggests a new population to carry out the process. Component $\Delta\text{how} \Delta\text{who}$ presents new features, BBs, and connections that result in the context of the new population performing the activity. Module *Observe differences to experience* in Fig. 6 computes the various kinds of Δ s of the TM equations.

Example : Agent 1 proposes a solution in which employees receive health insurance from the employer. Next, Agent 2 modifies the solution by indicating that employees receive health insurance from the federal government. $\Delta\text{how} = 0$ as both solutions refer to offering health insurance. $\Delta\text{who} = 1 - \text{Similarity}(\text{employer}, \text{federal government})$, which is the dissimilarity between concepts *employer* and *federal government*.

New responses are also created using the following equation:

$$\Delta(\text{what}, \text{for who}) \approx \Delta\text{what} + \Delta\text{for who} + \Delta\text{what} \Delta\text{for who} \quad (3)$$

The equation presents that the overall changes when concepts *what* and *for who* change depend on the modifications of the goal *what* (Δwhat), concept *for who* ($\Delta\text{for who}$), and the relation (combination) of the two changes.

Example : Agent 1 suggests full coverage for everyone, then Agent 2 proposes basic coverage for elderly. In this case, $\Delta\text{what} = 1 - \text{Similarity}(\text{full coverage}, \text{basic coverage})$, and $\Delta\text{for who} = 1 - \text{Similarity}(\text{everyone}, \text{elderly})$. $\Delta\text{what} \Delta\text{for who} = (1 - \text{Similarity}(\text{full coverage}, \text{basic coverage})) \circ (1 - \text{Similarity}(\text{everyone}, \text{elderly}))$, which is the combination of the two dissimilarities.

Equations (2) and (3) describe the extension of the group and local knowledge during group knowledge evolution. It considers member interaction. The changes of the components are used by the five strategies.

When responses from two distinct response evolution traces are combined, e.g., using strategies 2 or 3, (1) of the new response is described by the following equations:

$$\text{what}^{(1)} \circ \text{what}^{(2)} \circ \text{outcome}^{(1)} \circ \text{outcome}^{(2)} (\text{for who}^{(1)} \circ \text{for who}^{(2)}) \models \text{how}^{(1)} \circ \text{how}^{(2)} (\text{who}^{(1)} \circ \text{who}^{(2)}) | \text{why}^{(1)} \circ \text{why}^{(2)} \quad (4)$$

Note that not all the components of the equation might be used during a combination. The combination bridges the semantic gap between the two categories (and their traces), which can be addressed using similar or dissimilar interpretations of the combined elements. A property-based combination results if the two *what* frames are combined (left side of the equation), and a relation-based combination originates when the two *how* frames are combined (right side of the equation).

Example : Agent 1 suggests that free insurance should be given to elderly people. Agent 2 indicates that any doctor should treat any elderly person. $\text{what}^{(1)} = \text{what}^{(2)}$, which is access to health care. Similarly, $\text{for who}^{(1)} = \text{for who}^{(2)}$, which are elderly persons. Finally, $\text{outcome}^{(1)} = \text{free insurance}$ and $\text{outcome}^{(2)} = \text{should be treated by any doctor}$. Hence, $\text{outcome}^{(1)} \circ \text{outcome}^{(2)}$ suggests free insurance for treatment by any doctor.

Equation (4) might be instantiated during knowledge evolution as simpler situations. For example:

$$\text{what}^{(1)} \circ \text{what}^{(2)} \models \quad (5)$$

The two combined responses include only *what* frames, hence the right side of (4) is missing. There is an opportunity to identify the missing *for who*, *how* and *who* frames. Or, as in the following equation:

$$\text{what}^{(1)} \models \text{how}^{(2)} \quad (6)$$

which combines the *how* frame of a response (implementation) with the *what* frame of another response. Opportunities includes identifying possible *for who* and *who* frames for the combination. Another possibility is as follows:

$$\text{what}^{(1)} \models \text{who}^{(2)} \quad (7)$$

The *who* frame of a response is combined with the *what* frame of another response. There are opportunities to include alternative *for who* and *how* frames for the combination. Other kinds of combinations result depending on the missing parts of (4).

Combinations in which certain frames are missing, similar to the previous three equations, produce opportunities to identify new knowledge beyond the boundaries of the current concepts in the solution space. Such opportunities lead to divergence from the current concepts, but the divergence is controlled by (e.g., within the boundaries) the combined

frames. Some of the opportunities might be exploited by subsequent decisions or might be missed.

Figure 5(a) mathematically models the likelihood of producing change Δ_k by member k . The figure shows two kinds of changes: moving from $\Delta_k^{(i)}$ to $\Delta_k^{(j)}$, or repeating the same kind of change like the present change, e.g., $\Delta_k^{(i)}$ to $\Delta_k^{(i)}$. For example, such a transition could describe a change from $\Delta(\text{how, who})$ to a change $\Delta(\text{what, for who})$. Each arc is labeled with the strategy that can produce change $\Delta_k^{(j)}$ from the current change $\Delta_k^{(i)}$:

$$str_k^{(i,j)} = decision_k(cue, Similarity_k, impact_k, motivation_k) \tag{8}$$

Figure 1(b) illustrates how different strategies produce a certain kind of change. The change might or might not be transmitted to the entire group depending on its propagation infrastructure, like dissemination procedure, visibility, importance of prestige, and so on. For example, Strategy 1 or Strategy 2 can add a new feature or a new specific sample to the current knowledge. $decision_k$ models member k 's way of selecting the problem solving strategy for a given cue that attracted his/her attention, the similarity with the local knowledge, the associated importance and emotion, and the energy required to select and use the strategy and any other related knowledge. The bottom module of Fig. 6 illustrates the using of the strategies in the software implementation of the model.

At each time moment t , a group member j makes a decision between generating a response pertaining to an existing category (sub-category), or creating a solution that belongs to a new category. The first situation indicates response convergence, and the second case indicate response divergence:

$$str_j(t) = p_j^{(conv)} str^{(conv)} + (1 - p_j^{(conv)}) str^{(div)} \tag{9}$$

where parameter $p^{(conv)}$ is the likelihood for member j selecting a converging strategy, and $(1 - p^{(conv)})$ is the likelihood of selecting a diverging strategy. Converging strategies are strategies 1, 4, and 5 as they do not change the status-quo of the group knowledge. Strategies 2 and 3 are diverging as they create solutions that use new concept

categories (sub-categories). Moreover, if motivation is also considered, (9) can be extended as follows:

$$str_j(t) = motivation_j^{str_k^{(conv)}}(t) p_j^{(conv)} str^{(conv)} + motivation_j^{str_p^{(div)}}(t) (1 - p_j^{(conv)}) str^{(div)} \tag{10}$$

Parameter $motiv_j^{str_k^{(conv)}}(t)$ represents member j 's motivation in using the converging strategy $str_k^{(conv)}$, and parameter $motiv_j^{str_p^{(div)}}(t)$ is the member's motivation in using the diverging strategy $str_p^{(div)}$.

Example : Figure 5(c) summarizes the link between two successive responses by two interacting agents. The first response cued the local knowledge of the second agent. The second agent produced a new solution in response. The first response had a certain semantics within the local knowledge of the second agent. The semantics is described by a distribution of similarities between the concepts in the local knowledge and the concepts in the response (overlapping region).

The new solution that results by applying a strategy has features that pertain to the overlapping region. The distance between the new solution and the first response (which served as a cue) depends on the amount of overlap, e.g., smaller overlaps create smaller distances (e.g., dissimilarity) between the two responses, while larger overlaps originate greater dissimilarity.

The following equation results after aggregating (10) at the group level:

$$str_{Group} \equiv \cup_{j \in Group} str_j(t) \tag{11}$$

It represents the ratio in which the different solving strategies are used by the group members at time t . Ideally, a group would attempt to maximize the current opportunities for convergence, while providing sufficient divergence to continue in the future (e.g., continuation of the group evolution process).

$$\max \cup_{j \in Group} motivation_j^{str_k^{(conv)}}(t) p_j^{(conv)} str^{(conv)} | \cup_{k \in Group} motivation_k^{str_p^{(div)}}(t) (1 - p_k^{(conv)}) str^{(div)} \tag{12}$$

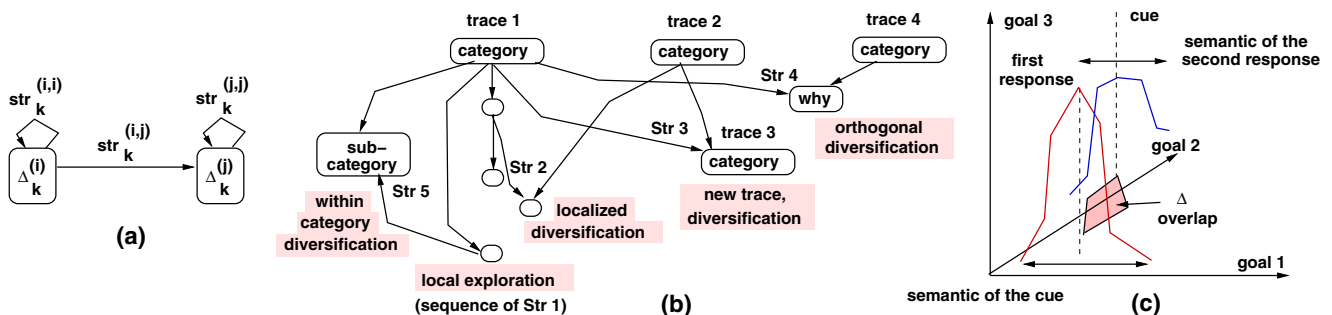


Fig. 5 (a) Modeling of knowledge evolution, (b) characterization of knowledge evolution, and (c) description of the changes in responses

Maximizing convergence implies that the group pragmatically receives all benefits that can be obtained by expanding the unexplored categories. Providing sufficient divergence indicates that less related categories are also identified, and/or that previously unrelated categories are being connected through new responses. Hence, the distance between these categories decreases. The second component represents some degree of *altruism* [19], as there is no guarantee that the dissimilar responses will have an impact on the group's future responses.

3.4 Group knowledge evolution

Group behavior is an evolutionary process in which the group's solution space (e.g., knowledge related to the problem) is extended through the solutions proposed by the group members. Traditional computational models for evolutionary processes are guided by well-defined fitness functions, which express the degree to which solutions meet the problem requirements. In this case, however, there is no explicit fitness function, as the criteria according to which the group raters will assess the solution's quality is unknown to the participants when devising new responses. Instead, every member produces a new outcome based on its own (local) knowledge and beliefs, selected responses by others, and the decision to spend a certain amount of effort on creating the outcome, i.e. motivation.

3.4.1 Characterization of knowledge evolution

Figure 5(b) illustrates the characterization of the knowledge evolution process, e.g., what kind of changes are introduced by each solving strategy to the solution space. The figure describes the relation between the different kinds of interactions in a group. The following five situations are possible depending on the used strategy.

1. A sequence of Strategy 1 (*str* 1) produces *localized extension* of a category as new details (features and samples) are added to the category and to its causal sequence. A set of possible features and concrete samples might be described for the category. All responses remain within the general characteristics of the overall category. Moreover, the response sequence over time creates a *trace* of responses that share the same category, hence a set of defining attributes. According to (2) and (3), the distance of a new response to the previous concepts ($response_{prev}$, which might be explicitly cited during group-level solving) is as follows:

$$Dist(response, response_{prev}) \approx \Delta_{what}, \Delta_{who}, \Delta_{for\ who}, \Delta_{how} \quad (13)$$

The distance of the response to the previous one depends linearly on the added changes (Δ). A high response similarity suggests that agents have highly overlapping local knowledge (with respect to this category) and similar interpretations.

2. *Localized diversification* occurs if responses pertaining to a category, i.e. those of trace 1, are combined with concrete features and samples pertaining to another category, e.g., like those of trace 2. Strategy 2 (*str* 2) is used in the combination process. The concrete features and samples are used because of their perceived utility within the responses of the original trace (i.e. trace 1). As the causal sequence of the new solution remains the same, a relation-based combination [79] is used to connect the features of the response of trace 2 to the new response. Similar to (13), the distance of the new response as compared to that of trace 1 ($response_{trace\ 1}$) is additive with the modifications introduced by the concrete elements from the response of trace 2:

$$Dist(response, response_{trace\ 1}) \approx \Delta_{what_{trace\ 2}}, \Delta_{who_{trace\ 2}}, \Delta_{for\ who_{trace\ 2}}, \Delta_{how_{trace\ 2}} \quad (14)$$

The case indicates a partial overlapping of the two participants' knowledge structures (including some more detailed features) together with a similar interpretation.

3. Strategy 3 (*str* 3) introduces a new category by combining the categories of two separate traces, like trace 1 and trace 2 in the figure. It introduces a *stronger diversification* than that through Strategy 2, as the new category defines a new trace distinct from the previous ones. Strategy 3 implements property-based combination [79], meaning that the characteristics of the new combination are more than the changes from the individual traces. The distance from the responses of the two original traces is characterized by the following equation:

$$Dist(response, response_{trace\ 1}, response_{trace\ 2}) \approx \Delta_{trace\ 2}, \Delta_{trace\ 1} \quad (15)$$

where $\Delta_{trace\ 1}$ and $\Delta_{trace\ 2}$ describe a Δ_{what} , Δ_{who} , $\Delta_{for\ who}$, or Δ_{how} pertaining to the two traces. The new category has different distinguishing parameters and constraints than the two original constraints. They might or might not be explicitly identified by the participants [25, 79]. Experiments suggest that their explicit identification and labeling can help creating more new responses pertaining to the new category [25, 79]. This case shows a partial overlapping of the knowledge structures, but involving less concrete items as compared to the situations for strategies 1 and 2. The interpretation is similar for the participating members.

4. Strategy 4 (*str* 4) introduces an *orthogonal diversification*, meaning that the category and responses pertaining

to this trace have an opposite (conflicting) meaning to those from the trace they refer to. Strategy 4 refers to *why* frames. It indicates an overlapping of the member's knowledge structures while their interpretation is opposite.

5. Strategy 5 (*str* 5) creates a new sub-category by generalizing the characteristics of more concrete responses. For example, such a generalization can occur after introducing a feature of a new kind (either through strategy 1 or through strategy 2) from another trace. Strategy 5 introduces a new sub-trace within the initial category, hence its *local diversification* can be stronger than that produced through strategy 2.

Because the combination is at a concrete level, the likelihood of creating a correct combination between the responses of two traces is higher for Strategy 2 than for Strategy 3. It is more probable that inconsistencies and errors in the new response are detected, if more precise details are considered than if more abstract reasoning is performed, e.g., using Strategy 3. Therefore, it is likely that producing more Strategy 2 combinations along similar traces followed by using Strategy 5 to abstract the responses as a more general sub-category is likely to be more robust in terms of solution correctness than by directly using Strategy 3. There can be different strategies for robust diversification, so that the new traces are also feasible (possible) and correct.

3.4.2 Continuation of the evolution process

Instead of fitness, the group evaluation process utilizes the concept of *continuation*. It is the groups capability to continue to produce new knowledge under a certain set of resource constraints set for the members (e.g., time, energy, personnel, motivation). Continuation at time t is mathematically defined as follows:

$$\text{Continuation}(t) \equiv \text{New knowledge}(\text{Knowledge}(t), \text{Resources}(t)) \quad (16)$$

The success of a group's evolution is defined by its maximum continuation,

$$\max_t \text{Continuation}(t) \quad (17)$$

Group member j has the capability to produce a new response, if the resources required to create the answer is less than a threshold value Thresh_j that depends on his/her motivation (motiv_j):

$$\text{Resources}_j(t, \text{Dist}_j(\text{local_know}_j, \text{beliefs}_j, \text{selected}_j(\text{Knowledge}(t))), \text{str}_j(t)) \leq \text{Thresh}_j(\text{motiv}_j) \quad (18)$$

The resources required for member j depend on the distance Dist_j (i.e. similarity) between his/her own knowledge (*local.know_j*) and the selected ideas from the groups knowledge at time t (*Knowledge(t)*), his/her beliefs (*beliefs_j*) on the meaning of the knowledge concepts and features (including the attached emotional information), and

the solving strategy (*str_j*) used to produce the response. The selection process of ideas from the group knowledge is represented by *selected_j*, which incorporates elements like capacity to interact with the other members (visibility of the members), and their prestige. If the resource constraint is met then the new outcome created by a member is as follows:

$$\text{Response}_j(t) = \text{str}_j(t, \text{local_know}_j, \text{beliefs}_j, \text{selected}_j(\text{Knowledge}(t))) \quad (19)$$

The agent implementation in Fig. 6 shows module `Estimate effort to create outcome` used to predict the required effort to produce a solution, and module `Estimate reward of outcome` to predict the expected rewards of a solution. The cost function used by the evolutionary process is also shown in the figure.

The variety of the group's knowledge includes all concept categories, sub-categories, features, and examples produced by the members as responses to the problem. Similarly, the similarity of a group's knowledge represents the set of distances among the concepts in the responses. A category (sub-category) is *unexplored* if members are still motivated to produce responses pertaining to the category (sub-category), e.g., they decide it is worth spending the effort to generate a new response in that category. Or that the benefit according to the criteria of interest exceeds the required effort. The *unexplored variety* are the unexplored categories of a group's variety. A *group's opportunities* at a certain moment represent the unexplored variety of the group. A group's opportunities are its flexibility in continuing the evaluation process. A group continues its existence as long it has sufficient flexibility. A straightforward observation is that a group *converges* in its evolution, hence eventually dies out, if an increasing amount of resources is needed to produce new knowledge, or there is an increasing threshold value. A group has enough *divergence*, thus has the potential of remaining alive, if it generates new categories and sub-categories that attract enough motivation. Hence, a group's evolutionary continues as long as there is enough unexplored variety within its knowledge, and this variety can attract enough resources from the members.

4 Self-organization during knowledge evolution

The dynamics of applying the five strategies in Fig. 5(b) can create a certain self-organization of the responses depending on the participants' characteristics and interactions, like knowledge similarity, interpretation, motivation, resources, and existing infrastructure for knowledge dissemination. Self-organization emerges from a certain set of common

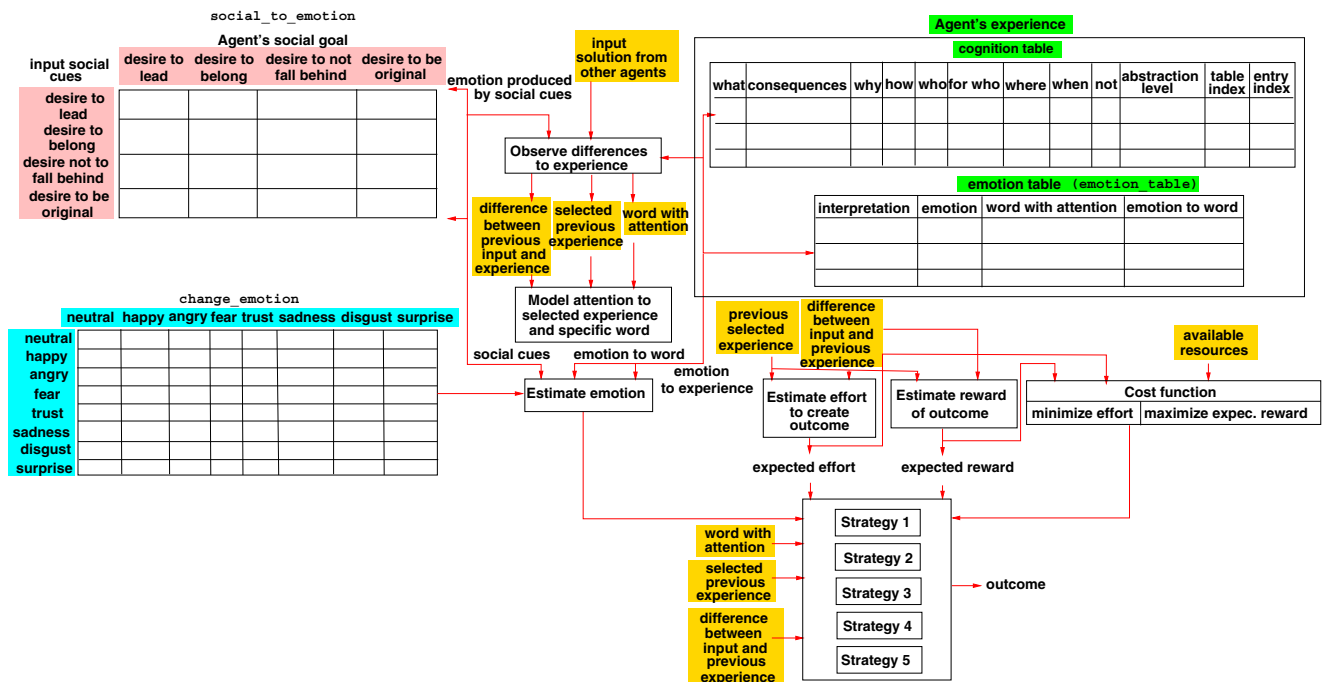


Fig. 6 Structure of an agent's implementation

parameters of a group members. Once self-organization emerges, it perpetuates (conserves) itself to some degree.

The model defines self-organization starting from the likelihood of concept set S being used in a new response ($likelihood(S)$).

The self-organization patterns in Fig. 5(c) result depending on the nature of the responses in set S and how they relate to the responses that use them.

1. *Modular self-organization* results if responses in set S and those that use them (denoted as set R) pertain to the same category. This case describes *convergent* evolution, in which subsequent responses are similar to the original category and the responses in set S . The causal sequences (if they exist) of the responses are also similar. The reasons why responses are suggested are also similar. Responses in set R are alternatives to the common category. Every new response details a certain aspect of the category, which has not been improved by the previous responses in set S . This increases the quality of the solutions in the category and their usefulness. A special self-organization case exists if nodes in set R pertain to the same trace and set S includes a single node, which is the *anchor* (originator) of the modular structure, like in Fig. 5(c).

The following metrics are defined in the model:

1. *Distance distributions*: The metric refers to the distribution of the distances (e.g., path lengths) for the concepts in sets S and R . Longer distances suggest that responses

in set S were at higher abstraction level and had enough flexibility to be further detailed by more responses. They also indicate that agents had enough resources and motivation to continue to refine the responses in set S , thus these responses have been considered to be important.

2. *Degree of alternatives*: The metric describes the number of alternatives of the concepts in a modular structure.
3. *Degree of similarity*: The metric expresses the amount of similarity among concepts and their features and concrete samples. This arguably represents the level of knowledge redundancy present in the solution space.

2. *Spreading-out self-organization* occurs if set S has a single node, and the nodes in set R belong to different traces. The node in set S is a *hub*. Figure 5(c) shows the situation.

3. *Merging self-organization* is shown in Fig. change(c) in which set R has a single node, the *node of convergence*, and the nodes in set S belong to different traces. Such self-organization structures are created through strategies 2 and 3. Strategy 2 adds a new feature or sample, which are likely to extend the flexibility of the response, thus the opportunity of more subsequent steps. Strategy 3 produces a new category or sub-category or a new sample that can be generalized through Strategy 5 as a new category or sub-category. The new category is different from the converging traces, but the likelihood of it being feasible is higher as it partially relies on existing categories.

5 Software implementation of the computational model

Figure 6 presents our implementation of the proposed model. The implementation encodes the data and relations of the declarative TM. Other implementations are possible too, if different modeling solutions are used for the equations. For example, different heuristics can be utilized to characterize the similarities and differences between responses (Δ s in the TM equations) between responses, knowledge representations, agent motivation, attention, emotion, and social cues.

Every agent stores its local knowledge (experience) in two separate tables. Section 3.2 presents an agent's local knowledge. Table `cognition_table` saves the frames (clauses) that are known to the agent, such as the frames it possessed at the beginning of problem solving and the outcomes that were generated and learned during the process. The table corresponds to Fig. 4(b) in TM. It stores an agent's knowledge representation, like in Figs. 3(a) and (b). Equations (18) and (19) refer to the local knowledge. Every entry of the table stores an outcome, including its related frames, like the *what*, *goals (consequences)*, *why*, *how*, *who*, *for who*, *where*, *when*, and *not* frames. In addition, the entry stores the abstraction level of the outcome, as well as two indexes that point to the outcome for which the present row is an instance. Index `index_table` is a pointer to the table entry including the more abstract outcome, and `entry_index` points to the clause of the abstract outcome that is detailed by the current row. The abstraction level is incremented every time a new outcome is added to detail one of the existing clauses. Every entry also stores the pair of pointers `synonym_table` and `synonym_entry` pointing to another table entry that includes a similar outcome but with different goal. This shows situations in which the same outcome has other meanings in different contexts. Another pair of pointers, `similar_table` and `similar_entry`, point to the output and frame that instantiate the same abstract outcome and frame while being different than that at the current index. The table also stores for each frame the variety of its detailing outcomes.

The second table, `emotion_table` stores aspects related to the attention and emotion associated to each outcome. Figure 4(b) and (8), (10), and (12) describe the using of emotions to model the agent's behavior. It includes the interpretation of the outcomes, such as its gain which is the expected rewards minus the effort needed to produce the outcome. Entries `associated_attention` and `emotion` store the degree of attention and the emotion associated to the entire outcome. Entry `word_with_attention` presents the clause that is perceived to be the most significant in

setting the *what* clause and the consequence of the outcome. This clause represents the main parameter of the causal relation of the outcome. Entries `attention_to_word` and `emotion_to_word` describe the attention and emotion associated to the most significant clause of the outcome.

The agent model in Fig. 6 corresponds to Fig. 4. It first observes the differences of the solutions from other agents and the agent's experience in table `cognition_table`. Similarity is found using the metrics provided by WordNet [21]. The module identifies the previous outcomes that are most similar to the current input (`selected_previous_experience`), the frames that differ in the previous experience and input, and the word in the difference that captures the highest attention of the agent. Note that the module finds the differences between input and most similar experience while receiving emotion information that corresponds to social cues from others. The emotion information is used to select outcomes that are emotionally consonant with the social cues, like outcomes for which the agent's interpretation and attention are high, and also their associated emotion does not conflict with the social cues.

The next component models the attention of the selected outcome experience and the most dominant word to decide if they exceed the attention threshold of the agent. A threshold constant (`Thresh_Attention`) is used to model different kind of personalities. The outcomes and dominant words that exceed the threshold barrier are further processed by the agent. Module `Estimate_emotion` predicts the emotion produced by aggregating the received social cues, the agent's current emotion, and the emotions associated to the selected outcomes and dominant word (stored in Table `emotion_table`). Then, module `Estimate_effort_to_create_outcome` uses the Bass model to predict the effort required to create an output [51], and module `Estimate_reward_of_outcome` predicts the expected reward of the new outcome using the similarity with previous, similar outcomes. The estimations for effort and reward are used in the implementations of (8), (10), (12), (16), and (18).

Next, the agent creates and then learns a new outcome. The outcome uses the selected outcome stored in the experience table as well as the observed differences, the predicted rewards and efforts, the modeled emotions, and the agent's objectives as represented by the continuation function in (16) and (17). The created outcome implements the five strategies (Sections 3.1 and 3.4.1) corresponding to the four kinds of variations and the contradiction (blocking) of a previous outcome, situation in which an outcome for the same *what* clause uses different other clauses in the solution.

Emotions are modeled as an FSM with its current state corresponding to an agent's present emotional state. The FSM states correspond to basic emotions, i.e. neutral, happy, angry, fear, trust, sadness, disgust, and surprised.

The FSM has two different states for each kind of emotion, one for the medium level and one for the acute level. Table `change_emotion` encodes the transitions between states of emotions due to the emotional cues received by an agent. Figure 7(a) presents agent emotion modeling. Module `estimate_emotion` in Fig. 6 uses the Finite State Machine (FSM) description to express changes in an agent’s emotions. The current emotional state is an FSM state. Changes in the current emotion correspond to the FSM transitions depending on the emotions associated to the selected entry in the local knowledge (field `emotion` in table `emotion_table` in Fig. 6), the emotion associated to the main concept of the response (field `emotion_to_word` in table `emotion_table`), and the emotion associated to social cues.

Different approaches have been proposed to model emotions, personality traits, and social interactions [93–96]. Modeling has been challenging due to many correlated, subjective parameters [94]. The software implementation used a model similar to [97, 98]. The model in [97] uses nine primary and three secondary emotions of two valences. Emotional states correspond to homeostatic states, which change due to internal and external cues. Even though methods are not mature yet, the implementation includes emotion and social interaction modeling to capture their effect on the selected knowledge, decisions, group interactions, and knowledge evolution. Future work will continue to improve the used models. Moreover, future work will also attempt to address new challenges that are likely to occur. For example, the current lockdown due to Covid-19 pandemic forced many teams to work remotely. Due to the lack of physical proximity, remote interactions changed the way social and emotional cues are perceived and interpreted by others, thus team behavior was modified too.


Social interactions are modeled as Finite State Machines (FSMs), in which every state indicates the current dominant social attitude of the agent towards the other agents in the group. Every agent has a social profile that is defined with respect to its own social goals, like the desire to lead, the desire to belong to a group, the desire to not fall behind

with respect to the produced outcomes, and the desire to be original as compared to the other agents. The social cues an agent receives from the others are also of the four kinds. Table `social_to_emotion` encodes the emotions that result from the social interactions, e.g., happiness, angry, trust, and so on. The emotions associated to social cues were modeled as shown in Fig. 7(b). The depicted model corresponds to table `social_to_emotion` in Fig. 6. Table rows correspond to an agent, and columns to the agent it is interacting with. Table entries indicate the emotional cue triggered by social interactions. Rows and columns are labeled with interaction profiles, like *DTL* describes agents that would like to lead, *DTNF* denotes agents that want to remain clustered with the other agents, *DTDO* presents agents which like to imitate others, and *DTB* are agents that like to contradict anyone.

Example : Two different kinds of conflicts can exist between team members: task and relationship conflicts [99]. Task conflicts relate to task-specific aspects, like goals and solutions, while relationship conflicts indicate interpersonal incompatibilities between team members. First, we assume a team of four medical doctors that must devise a holistic, customized treatment plan. A task conflict emerges between them due to their different perspectives about the pros and cons of a treatment plan. Increasing the likelihood of finding a novel treatment requires first finding a set of broader criteria on which all doctors agree on, e.g., there is an alignment of their local knowledge (Fig. 1(a)).

Next, let’s assume there is such an incompatibility between two members, i.e. both want to be team leaders, which corresponds to rows and columns labeled *DTL* in Fig. 7. This social interaction produces *angry* emotions, which encourages the two agents to continuously select Strategy 4. This strategy describes a situation in which the two agents keep contradicting each other, because they have opposite understandings of their responses, e.g., field `not` in table `cognition_table` (Fig. 6) is set.

Figure 8 illustrates the learning hierarchy implemented for each agent. The hierarchy has the following six levels:



	DTL		DTNF		DTDO		DTB	
	mild	high	mild	high	mild	high	mild	high
DTL	mildly angry	very angry	mildly happy	very happy	mildly trusting	very trusting	mildly angry	very angry
	high	very angry	very happy	very happy	very trusting	very trusting	very angry	very angry
DTNF	mild	neutral	mildly happy
	high	mildly happy	mildly happy
DTDO								
...								

Fig. 7 (a) Emotion modeling and (b) social interaction modeling

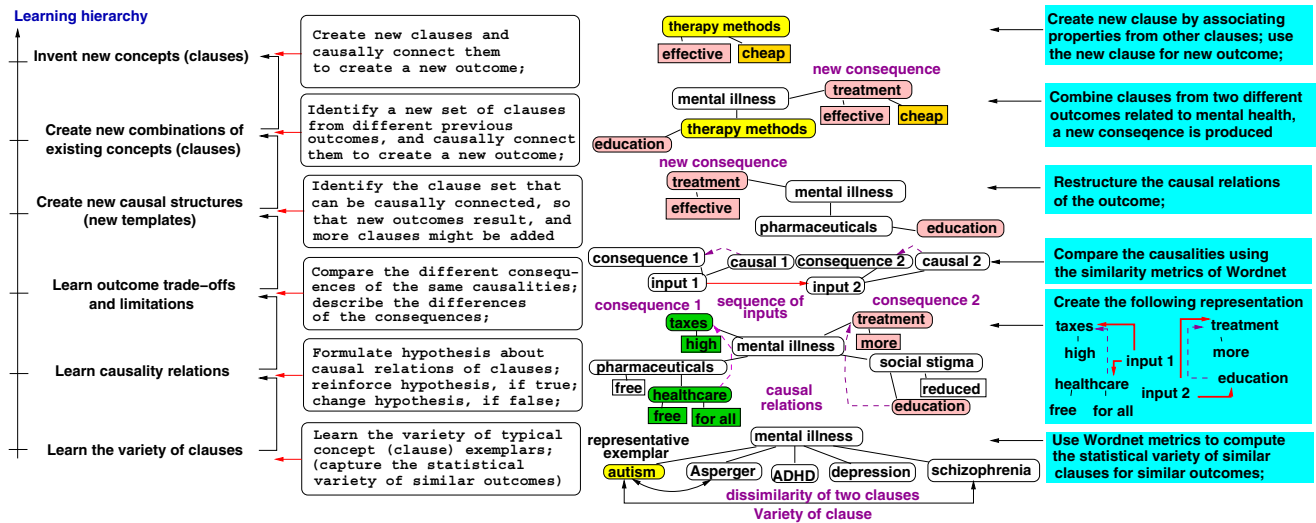


Fig. 8 Learning hierarchy

1. The bottom level learns the variety of clauses, such as the statistical variety of the clauses used in outcomes. Each clause has a typical exemplar and the variety of the related alternatives. The implementation uses the similarity metrics provided by WordNet [21]. For example, the concept “mental illness” has “autism” as a typical exemplars, and “Asperger”, “ADHD”, “depression”, and “schizophrenia” as related instances. The variety associated to “mental illness” includes the similarities of the four latter instances and the typical exemplar. This level mainly implements associative learning, which identifies the invariant attributes of the concepts [37, 78].
2. The next bottom level learns the main causal relations of the outcomes. The algorithm produces a relation (shown as a purple dashed line in the figure) between the most dominant clause of the outcome and the outcome consequence. Two causal relations are shown in the figure, one used green boxes and the other pink boxes. The first causal relation indicates that “free healthcare for all”, one of the *how* clauses used in an outcome, acts as a causal relation to the consequence that “taxes will be high”. The second causal relation describes that “education” produces “more treatment”. Note that the same outcome, “mental illness care” has two causal relations defined by two distinct outcomes. The learning algorithm generates a representation as shown by the bottom second, light blue box.
3. The third level learns the trade-offs and limitations of certain outcomes. It pairwise compares all causal relations associated to the same outcome. The similarity of the consequences is found using the metrics offered by WordNet.
4. The fourth level creates new causal structures (templates) through restructuring of the causal relations of the current templates. For example, “education about pharmaceuticals” is used to produce the new consequence that “effective treatment” is produced for “mental illness”.
5. The fifth level creates new combinations using clauses from different outcomes for the same *what* clause, e.g., “mental illness”. For example, “education” and “therapy methods” are combined together in a new solution. Both were previously *how* clauses, but after combination, one became the new *what* clause, e.g., “therapy methods”.
6. The top-most learning method involves creating a new frame (clause), and then using it to produce a new outcome. It is achieved by adding new properties to an existing frame, i.e. “effective” and “cheap” to “therapy methods”.

Figure 8 suggests that the local knowledge structure of an agent continuously expands along the six levels by adding new (i) typical exemplars and alternatives, (ii)-(iii) causal relations and trade-offs of the outcomes, (iv) causal templates, (v) clause combinations, and (vi) frames and outcomes. Thus, the expandable, local knowledge structures are weak ontologies. As defined in [20], as opposed to an ontology which is a shared conceptualization, a weak ontology allows insertion of new details to an ontology, and is not shared by all agents. We argue that knowledge representations similar to weak ontologies are more appropriate for describing team activities, as team members rarely share precisely the same knowledge and its interpretation. An intriguing future work topic is

understanding the nature of the logic that can be used for reasoning based on the local knowledge representation of each agent.

6 Experimental results

This section presents support for the proposed theoretical model (TM) using experimental data obtained from the study performed by Dr. Paulus and Dr. Kenworthy at University of Texas at Arlington [26] as well as data from simulating the TM implementation. The section discusses four kinds of results: (i) the nature of the solution space representations (as shown in Fig. 3) built for the responses created by twenty groups of participants that solved the problem discussed in Section 2. (ii) The way in which members of the same group reacted to each other's responses, as a way to validate the model expressing interactions among members, and depicted in Figs. 1, 4, and 5. (iii) The relation between the novelty of the solution space, as it was characterized by human raters, and the characteristics of the participants. (iv) The insight obtained from simulating the TM implementation.

The analysis used the responses about the affordability aspect of improving healthcare. The selection was justified by the fact that we felt that TM validation would be reasonable only if responses of the same kind were considered. The representations were separately constructed for the responses generated by every group, and then used in the analysis.

6.1 Solution space representations

First, this subsection offers a detailed presentation of the solution space descriptions, followed by the depiction of the solution space descriptions of the groups with the highest, median, and lowest solution novelty. These solutions were created during experiments with human subjects [26]. The shown descriptions are representative for all groups. A more detailed presentation of the nature of the categories that emerged in different groups was also discussed in this subsection (Table 1).

Figure 9, and Tables 2 and 3 present the characteristics of the solutions space description for the twenty groups that participated to the experiment. Figure 9(a) summarizes the solution space produced by the most creative group, Group 1. It shows the description template that resulted for different frames of the solution.

As also shown in Table 2, the produced solutions pertain to a smaller number of different kinds of concepts as compared to the less creative groups. However, some concepts are well explored, as shown in the figure, including the set of related features, the set of real examples that

are identified as important for the concepts, and the related sub-categories. Concept features were further developed by subsequent responses that indicated the related sub-categories, like *what* or *how* frames presenting what features actually mean or how they are realized. Similarly, subsequent responses add extra features to the identified real examples or to the selected sub-categories. The solution space representation is reasonably deep in terms of its number of abstraction levels. Also, it includes only very few instances in which Strategy 4 and Strategy 5 were utilized by the group members, suggesting that they did not attempt to create competing (opposite) alternatives for the same concept, or to generalize concrete aspects as more abstract concepts. There was a small number of reinforcements to show that one member fully agrees with another solution. This observations suggests that group members, while accepting each other's solutions, were mainly preoccupied by proposing responses that had a certain amount of modifications (Δ).

Partially similar solution spaces were created by the other groups with more novel responses, like groups 2, 3, 4, and 5 in Table 2. Their space representations include a similar number of abstraction levels, even though not all components (features, real examples, and sub-categories) are always present for the most developed concepts. Also, there is a larger number of real examples that were offered in the responses. There is an increased number of times when group members proposed alternative solutions through Strategy 4, and a larger number of reinforced solutions, showing that certain members fully agreed with the solutions of others. These results suggest an increased number of more detailed responses and more responses that disagree with each other.

Figure 9(b) and (c) depict the solution space representations for groups 8 and 9, which created solutions of median novelty (according to human raters). The two corresponding rows in Table 2 summarize the structural features of their representations. The most discussed concepts are less elaborated as compared to the previous groups, like they include a set of sub-categories and a set of features or a set of real examples, but not both. The graph structure is in general less deep than the previous representations, while concepts with more levels of abstraction indicate situations, in which group members disagreed with each other, therefore repeatedly used Strategy 4 to argument their responses (*why* frames). The real examples described concrete situations, constraints, or specific cases that presented *for who* and *why* frames.

Figure 9(d) and the corresponding row in Table 3 show the characteristics of the solution space representation created by Group 19. The group produced answers of a very low novelty. Similar to Group 1, responses referred to a small number of broad concepts. While similarly to

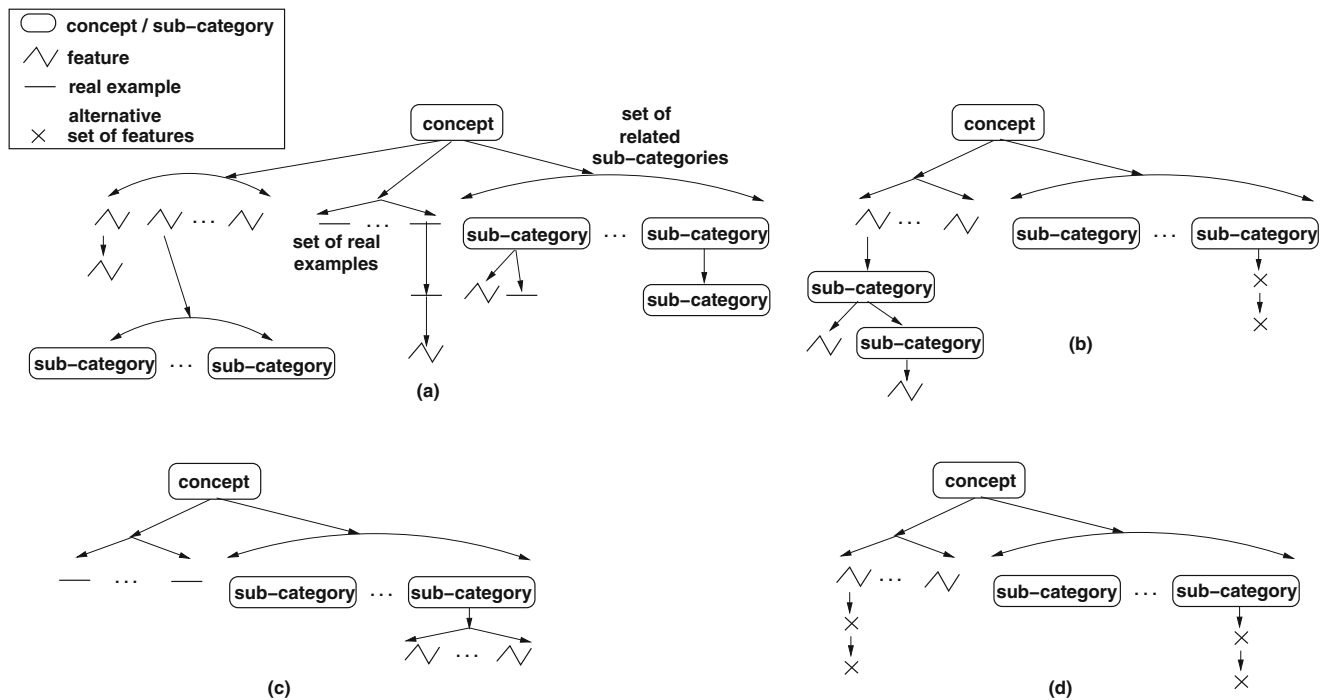


Fig. 9 Different solution space descriptions for group problem solving

groups with responses of median novelty, the representation includes a set of features and a set of sub-categories, there is an increased number of alternative responses, which are created through Strategy 4, hence were meant to replace each other. Strategy 4 was applied to frames *what* and *who* suggesting that there were disagreements on problem framing, such as the significance of the needs to be addressed to improve healthcare.

Due to the increased fragmentation (lower clustering) of the generated solution space, the average path lengths between concepts is larger for less creative groups than for more creative groups.

6.2 Group interactions

Group interactions were distinguished depending on the amount of overlapping of the group members' responses and how these responses were interpreted with respect to the problem goal. The following interaction types were observed for the experiments with human subjects [26]:

1. *Focused belief overlaps.* Some group members had similar responses about a certain aspect, e.g., three members had similar ideas about offering health insurance at work. The broad idea was further detailed by one participant, who added the constraint on giving full coverage by the employer in case of accidents. The new response was produced using Strategy 3. Another member distinguished accidents that were avoidable (Strategy 4), and changed the response by adding a detailed feature related to this constraint. Starting

from a common, shared abstract idea, the three members pursued a set of steps that added more details to the abstract concept. The final response is likely to be acceptable to all members.

2. *Mediated solution overlaps.* A set of members shared ideas that while not entirely identical, they defined a semantic region that was acceptable to them. For example, one group member suggested free health insurance, while another one proposed low cost insurance followed by free insurance for persons with low income (Strategy 2). While the three responses are not semantically equivalent, the second members mediated the two responses by adding the constraint referring to low income persons. The final response was acceptable to the first member too. A similar situation occurred also in another group between the responses of having free insurance and insurance discounts.

3. *Abstract idea overlaps.* A set of members agreed on an abstract idea, but then had different interpretations of the idea. For example, Fig. 10(a) shows the situation in which all group members agreed that healthcare must be affordable. Eight different interpretations were offered on what affordable healthcare actually means. The responses span a broad range, including opposing responses that healthcare should be free or not free. This case is a situation in which the group (Group 1) produced a large variety of responses. Even though some responses conflicted with each other, members did not engage in using Strategy 4 to support their opposing views. A more constrained version of this situation is if a new response was a special case of two

Table 1 Experimental results for different agent models

Aver. path lengths (2-3)	Clause variety (<#what:#consequence:#how:#who:#for who:#w here:#when>) (4-5)		Hierarchy levels (#level 1, #level 2, etc.) (#Str.1.;#Str.2.;#Str.3.;#Str.4.;#Str.5) (6-7)		Used strategies (#Str.1.;#Str.2.;#Str.3.;#Str.4.;#Str.5) (8-9)		Continuation (#min.cost:#max.know.) (10-11)		Resources (max:final) (12-13)			
	Agent 1	Agent 2	Agent 1	Agent 2	Agent 1	Agent 2	Agent 1	Agent 2	Agent 1	Agent 2		
0%	0.133	0.04	<15;1;7;2;2;12;11;2>	<16;2;0;0;1;3;3;1>	25(25)	46(38;7;1)	<2;4;0;0;0>	<4;12;1;0;3>	6;0	8;12	8;0;0.5	13.59;1.62
25%	0.037	0.049	<4;1;1;1;3;6;0>	<5;2;2;2;0;6;4;3>	48(37;9;2)	42(28;12;2)	<15;11;2;1;0>	<7;12;2;0;1>	19;10	10;12	17.34;8.3	13.7;4.5
50%	0.04	0.14	<14;9;1;2;12;13;3;1>	<15;8;0;10;8;9;3;1>	47(31;13;3)24(21;3)	47(31;13;3)24(21;3)	<1;13;2;0;1>	<1;3;0;0;1>	8;22	4;2	16.22;4.42	8.0;0.91
100%	0.064	0.065	<7;3;0;4;3;8;1;1>	<4;2;1;4;5;4;0>	37(28;9)	37(26;10;1)	<4;10;0;1;6>	<1;9;1;0;12>	7;14	1;22	22.0;11.5	22.5;10.8
			<11;3;5;2;0;10;8;8>	<4;1;2;2;0;3;2;1>								
			<1;2;4;1;0;4;4;5>	<2;2;1;0;0;2;0;1>								
			<11;10;3;3;8;8;4;0>	<10;14;6;6;12;4;0>								
			<15;4;7;3;5;3;0;0>	<13;4;5;1;3;6;1;0>								

more general responses produced by other two members. While one of the members did not adopt the new response, the other member elaborated on it by adding more details.

4. *Repeated instantiations.* A broad concept accepted by the group was further extended, as members produced real examples describing the concept. Arguably, the concrete examples acted as cues that illuminated other related examples by the same person or by another group member. For example, Fig. 10(b) shows a case in which the cost aspect of health services was instantiated for concrete cases, like ER, braces, X rays, doctor visits, and so on. Another example considered the broad concept of offering cheap pharmaceuticals, which then was instantiated as concrete examples, like using natural remedies, or offering cheap vaccination and condoms.

5. *Consolidated idea.* Group members had conflicting ideas about a response. Figure 10(c) illustrates this case. Responses mainly followed Strategy 4. New responses were mostly *why* frames and added few new features or sub-categories to the more abstract concept. Another case is the one in which the members fully agreed with each other, hence reinforced their beliefs without adding any new details to the responses.

The following observations were made about the ideas shared by group members. All members of the group with highest novelty (Group 1) were concerned about insurance, as all provided a response on this topic. However, there were fewer detailed ideas. Three members agreed about having insurance at the job, out of which two agreed on having job-related insurance with the exception of avoidable accidents. Two members agreed on giving free insurance to low income persons. Two members agreed that every doctor must accept any insurance. Even though four members were concerned about pharmaceuticals, there was only a reduced idea overlapping, as one member indicated that pharmaceuticals should be affordable, a response that was cited by a second member when suggesting that their cost must be reduced. Similarly, for the group with the next highest novelty, four group members were concerned about the need to offer affordable insurance, however, there were only pairwise agreement on the more detailed issues, like different pairs agreed that insurance should be offered at work and be provided in case of lay-offs too, taxing marijuana, free or affordable medicine, and low interest rates for medical expenses. As group response novelty decreases, there was higher incidence of cases in which members had opposing ideas. For example, two members of Group 5 had distinct responses on how to prevent abusing the healthcare system. There was disagreement between them on how the verification process should be implemented. Also, there was a higher incidence about detailed agreements between group members. For example, three members of Group 6 agreed not only on healthcare

Table 2 Parameter ranges and correlation coefficients

Group	concept (category)	# abstr. levels	#concepts / # features	# real examples	# str. 4	# str. 5	# common abstr. / # common detail.	# reinforc.
Group 1	pharmaceuticals	4	5 / 1	1	0	0	7 / 0	1
	insurance	4	23 / 5	4	1	0	14 / 0	0
	employers	4	6 / 0	0	0	0	2 / 0	0
Group 2	insurance	4	8 / 2	1	0	0	12 / 2	0
	medicine	2	2 / 1	0	0	0	2 / 0	0
Group 3	pharmaceuticals	2	3 / 3	1	0	0	5 / 0	0
	insurance	3	17 / 2	6	1	0	15 / 2	0
	emergency	2	1 / 2	0	0	0	2 / 0	0
	SF	2	4 / 0	4	1	0	0 / 0	0
	treatment	3	8 / 1	4	0	1	4 / 0	0
	dentist	2	1 / 1	0	0	0	0 / 0	0
Group 4	pharmaceuticals	3	4 / 4	2	1	0	2 / 0	1
	insurance	3	4 / 3	1	2	0	17 / 11	2
	cost	3	3 / 1	0	0	0	0 / 0	1
	healthcare	4	19 / 10	3	0	1	14 / 21	4
	doctors	2	3 / 0	0	1	0	6 / 0	2
	people	2	3 / 3	0	0	0	0 / 0	1
	senior citizens	2	3 / 0	0	0	0	0 / 0	1
	payment plan	4	9 / 1	1	2	0	11 / 0	5
Group 5	insurance	4	23 / 4	3	6	0	15 / 0	0
	treatment	2	7 / 3	5	1	0	0 / 2	0
	health professionals	4	11 / 0	2	2	0	6 / 0	1
	people help	2	2 / 0	2	1	1	0 / 0	0
	people	2	7 / 0	3	3	0	8 / 0	0
	like Canada	2	3 / 0	1	1	0	0 / 2	0
	government	2	2 / 0	0	0	0	2 / 2	0
	health education	3	2 / 1	0	2	0	0 / 0	0
Group 7	medication	3	3 / 3	0	0	0	3 / 0	0
	treatment	4	13 / 4	3	1	0	4 / 2	2
	health coverage	6	20 / 13	0	3	0	24 / 0	1
	people	4	5 / 1	0	0	0	0 / 0	0
	not connected	3	4 / 1	1	0	0	0 / 0	0
	doctors	3	5 / 0	0	0	0	4 / 0	0
	medical school	3	3 / 2	0	0	0	4 / 0	0
Group 8	medicine	4	4 / 6	1	0	1	5 / 0	1
	surgeries	5	5 / 1	0	0	0	2 / 0	2
	mentally ill	3	3 / 1	0	0	0	0 / 0	0
	ER	3	4 / 1	3	1	1	3 / 0	1
	ambulance	3	2 / 1	1	0	0	0 / 0	0
	hospitals	3	2 / 1	3	0	0	0 / 0	0
	insurance	3	2 / 1	0	0	0	0 / 0	0
	insurance company	2	3 / 0	0	1	0	2 / 0	0
	people	3	6 / 0	0	0	0	4 / 0	2
	low income	3	4 / 0	1	0	0	2 / 0	1
	students	2	2 / 0	0	0	0	0 / 0	0
	company	2	2 / 0	0	0	0	0 / 0	0
	healthcare	4	11 / 11	3	3	1	13 / 0	4
medical students	2	2 / 1	0	1	1	0 / 0	1	

Table 2 (continued)

Group	concept (category)	# abstr. levels	#concepts / # features	# real examples	# str. 4	# str. 5	# common abstr. / # common detail.	# reinforc.
Group 9	insurance	2	9 / 1	1	3	0	2 / 0	0
	cost	7	34 / 10	9	11	0	23 / 8	6
	elderly, impaired	2	2 / 0	1	0	0	0 / 0	0
	schools	3	3 / 1	1	1	0	0 / 0	0
	government	3	7 / 1	1	0	0	2 / 0	0
	healthcare	5	9 / 1	4	8	0	6 / 0	2
Group 10	insurance	4	13 / 1	3	1	0	9 / 2	1
	cost	3	5 / 4	3	0	0	2 / 2	0
	people	4	4 / 1	1	0	0	4 / 0	0
	based on income	3	7 / 2	3	0	0	0 / 0	0
	students	3	4 / 1	1	0	0	2 / 2	1
	pharmaceutical companies	3	3 / 0	0	0	0	0 / 0	0
	health care	4	16 / 3	2	0	0	10 / 2	3

having a lower price, but also on offering a bare minimum of services to everyone. Still, similar to groups with high novelty, there were instances when members agreed on the more abstract concepts too, while they generated an increased variety of more detailed responses for abstract concepts. Members in groups with low novelty often disagreed about topics with high emotional content, like vaccination, replicating medical services in other western countries, globalization, universal insurance, or health insurance for illegal immigrants.

6.3 Response evolution depending on group characteristics

Figure 11 presents the evolution of the solution space that was rated by human raters as having the highest novelty [26]. It corresponds to Group 1. Responses were relatively equally distributed among the four frame categories, *what*, *who*, *for who*, and *how*. Responses transitioned among the four categories, usually by using Strategy 3 to create a new outcome. Strategy 2 was applied to create a sub-category by adding constraints and conditions to a more abstract concept, like concept “health insurance given by employer” was further refined into a

new sub-category by adding the constraint “in case of work accidents”. Strategy 2 was rarely utilized. The longer sequences of responses pertaining to the same category were mostly produced using Strategy 1, when multiple related real examples were enumerated by a member, like “long term care”, “vaccination”, “drugs”, and “natural remedies”. Strategy 1 was also used for creating responses that described a concrete, but well-known situation (example), like offering health insurance to college students. Longer sequences of responses in the same frame category were created by repeatedly using Strategy 4. Strategy 5 was seldom used, but when used, the general concept was not further developed by the group. Thus, Strategy 5 had little impact on the solution space.

Figure 12 illustrates the evolution of the solution space for the group with the least novel responses (Group 19). In addition to responses that sampled the four categories, there was also a strong modularity of the outcomes indicating that responses had a higher similarity with each other, hence performing a localized sampling (search) of a certain aspect. There were much fewer responses describing *what* frames as compared to the previous group. Most of the responses were *how* frames with some indication of elements referring to *what* and *for who*, but with very few




Fig. 10 Types of interactions between group members

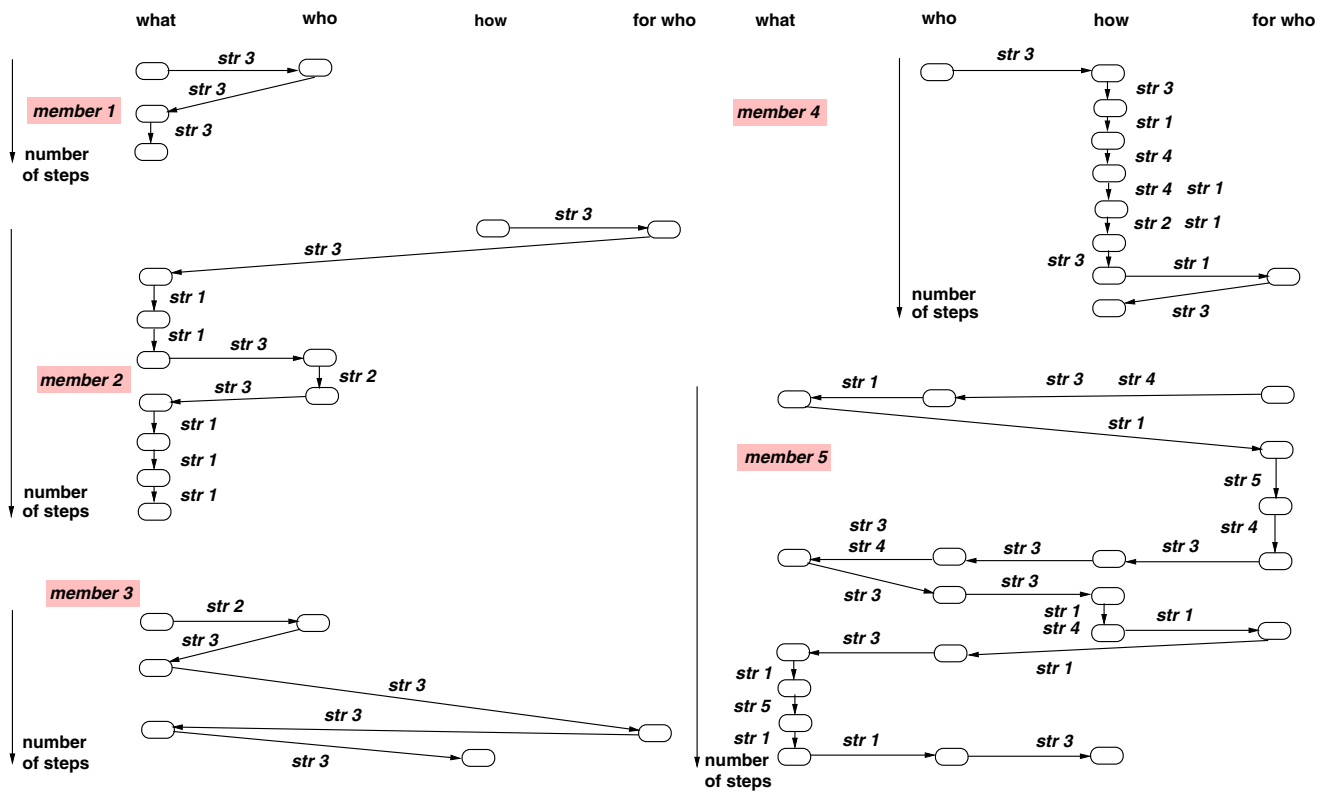


Fig. 11 Solution space evolution for the most creative group

mentioning *who* elements. The proposed solutions were mostly a subset of the solutions found by the previous group, such as the idea that healthcare for students should be supported by parents' insurance until students graduate from college.

Each member produced a much smaller number of responses as compared to the previous group. The transition from one category of responses to another category was produced mostly through responses produced by using Strategy 3. Hence, responses described more abstract

outcomes that could have served as categories or sub-categories of the solution space. However, these responses were usually not followed-up by further elaborations to increase the diversity of their category (sub-category). Strategy 1 was used to add details to a given situation, including responses that were suggested through Strategy 4 as alternatives to another response. Responses describing *how* frames referred to concepts with strong emotional content (and likely to trigger a certain kind of response from the others), e.g., taxation, involvement of government, and

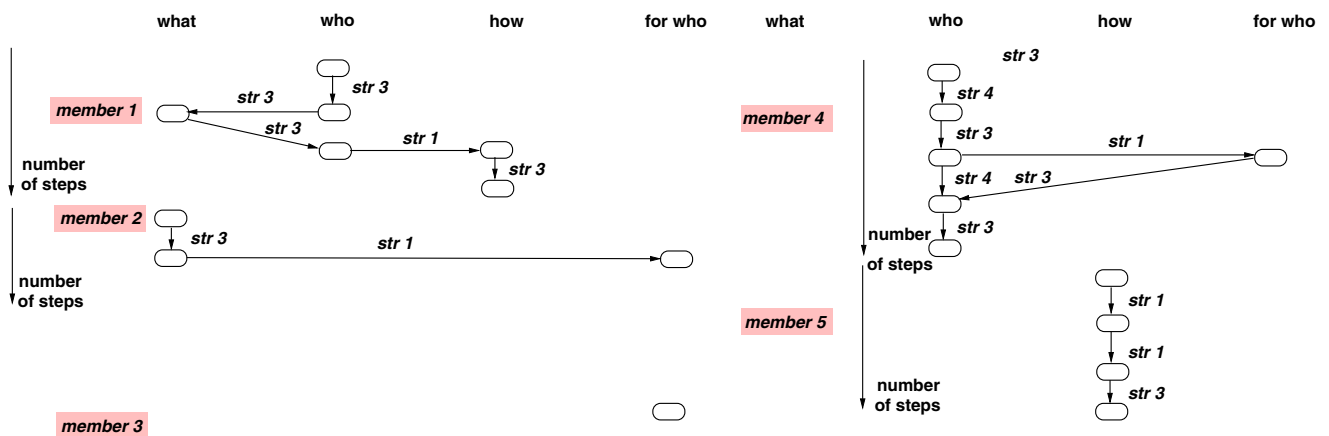


Fig. 12 Solution space evolution for the least creative group

Table 3 Parameter ranges and correlation coefficients

Group	concept (category)	# abstr. levels	#concepts / # features	# real examples	# str. 4	# str. 5	# common abstr. / # common detail.	# reinforc.
Group 11	pharmaceuticals	3	3 / 2	0	0	0	2 / 0	0
	insurance	4	19 / 6	3	6	0	17 / 4	3
	people	4	9 / 2	1	1	1	6 / 0	2
	illegals	2	2 / 0	0	0	0	0 / 0	0
	cost	3	3 / 3	2	0	0	2 / 0	0
	government	2	2 / 0	0	1	0	0 / 0	0
	borders	2	1 / 1	0	0	0	0 / 0	0
Group 12	pharmaceuticals	3	3 / 1	0	0	0	2 / 0	0
	insurance	5	12 / 6	1	4	0	11 / 0	2
	finance	4	5 / 4	1	0	0	0 / 0	0
	everyone	2	2 / 0	1	0	0	0 / 0	0
	people	4	10 / 1	1	1	0	5 / 0	1
	most Americans	3	3 / 0	0	1	0	2 / 0	0
	preexisting conditions	3	2 / 1	1	1	0	0 / 0	0
	government help	3	2 / 1	0	1	0	0 / 0	0
Group 13	healthcare	4	19 / 1	4	2	1	6 / 2	3
	prescriptions	3	2 / 1	0	0	0	0 / 0	0
	cost	5	10 / 1	3	0	0	2 / 0	1
	those	4	10 / 0	4	4	0	10 / 0	1
	primary care	3	3 / 1	1	1	0	2 / 0	0
	government	3	4 / 0	0	0	1	0 / 0	3
	like Canada	2	2 / 0	1	0	0	0 / 0	0
Group 15	healthcare	5	21 / 5	3	1	1	13 / 2	3
	insurance	3	6 / 5	1	0	0	2 / 0	0
	insurance companies	3	4 / 0	1	0	0	0 / 0	0
	everyone	6	14 / 0	0	8	0	9 / 0	1
	cost	6	15 / 5	2	7	1	13 / 1	3
	doctors	4	5 / 0	0	1	0	4 / 0	0
	healthcare	5	31 / 5	2	9	1	18 / 0	3
Group 16	medicine	5	6 / 1	0	2	0	2 / 0	0
	insurance	4	4 / 1	1	1	0	2 / 2	0
	machines	2	1 / 1	1	0	0	0 / 0	0
	healthcare	3	2 / 1	3	0	0	0 / 0	0
	taxes	3	5 / 0	0	1	0	0 / 0	0
	doctors	4	3 / 2	1	1	0	2 / 0	0
	healthcare	9	43 / 9	6	9	1	20 / 3	8
Group 17	drugs	3	4 / 0	0	0	0	6 / 0	1
	insurance	2	4 / 0	0	0	0	3 / 0	0
	HIV/AIDS	4	5 / 1	1	0	0	0 / 0	0
	government	4	7 / 0	0	0	0	0 / 0	0
	healthcare	4	14 / 5	2	0	0	3 / 0	1
	healthcare	5	40 / 7	6	5	0	20 / 9	14
Group 18	insurance	4	5 / 2	1	1	0	5 / 0	1
	America	5	4 / 1	0	0	0	0 / 0	0
	finances	5	7 / 1	0	4	0	9 / 0	1
	healthcare	4	/ 10	2	5	0	14 / 0	3

Table 3 (continued)

Group	concept (category)	# abstr. levels	#concepts / # features	# real examples	# str. 4	# str. 5	# common abstr. / # common detail.	# reinforc.
Group 19	pharmaceuticals	4	7 / 2	0	2	0	5 / 0	0
	companies	2	2 / 0	0	0	0	0 / 0	1
	healthcare	5	27 / 6	2	11	0	15 / 2	1
Group 20	insurance companies	5	9 / 2	3	5	0	9 / 2	0
	like other countries	10	15 / 1	2	8	0	16 / 2	4
	healthcare	6	27 / 7	2	11	0	20 / 0	1

military budget. Such concepts might quickly trigger the using of Strategy 4, if other members did not agree with a certain response, as well as detailed arguments (i.e. specific cases or examples) in order to support the alternative.

6.4 Agent model simulation

This subsection summarizes the experimental results using the agent model implementation presented in Section VI. The model was calibrated using the responses of the experiment with human subjects in [26]. The frames used in responses were used to set-up the database utilized to create the agents' local memory (experience). Solution quality was estimated by computing the shortest paths for the similarities of the frames used in the solutions created by the agents. Shorter paths among frames suggests that responses span a greater area of the solution space without creating separate solution clusters with few concepts to bridge the gaps between clusters. Clustered solutions have less variety, hence lower novelty.

Experiments were run with model parameters that mimic the experiment with human subjects, as well as some cases that were not covered by the experiment. The model was run for a number of parameters, including the number of agents in a small group (i.e. 2, 3, and 4 agents), the number of responses created by each agent (e.g., 30, 60, and 100 responses), the size of an agent's initial experience (local knowledge) (i.e. 20, 30, and 40 ideas), and the amount of shared ideas among agents in a group, 0.0%, 25%, 50%, and 100% of shared ideas. Other considered parameters were the threshold between effort and available resources, which is used to select between minimizing effort (cost) or maximizing knowledge (considered ratios were 0.15, 0.375, and 0.6), the upper similarity degree perceived by an agent for distinct concepts (used values were 0.2, 0.5, and 0.7, the higher value suggesting only that words with a high WordNet similarity are considered to be similar in the model), and different models for spending resources and receiving rewards for the created solutions.

Table 1 summarizes the experimental results obtained for agent pairs with 0.0%, 25%, 50%, and 100% shared ideas.

The case with 0% shared ideas describes independent agents that interact with each other independently of the meaning of their knowledge. This model is similar to [90]. The case with 100% shared information represents the situation in which all agents access the same knowledge, which is the traditional agent model [55, 57]. Columns 2 and 3 show the final average paths lengths of the two agents. Shorter path lengths indicate higher novelty [53, 69]. Columns 4 and 5 show for each agent the number of clauses in all their responses. The values are an indicator of the response varieties. Columns 6 and 7 present the number of concepts on the hierarchy levels of agent local knowledge, e.g., number of concepts on the lowest level, number of concepts on the next level of abstraction, and so on. Columns 8 and 9 present how many times each of the five strategies was used during problem solving. Columns 10 and 11 indicate the used continuation criteria for knowledge evolution, e.g., the number of times cost was minimized and the number of times knowledge was maximized. Finally, columns 12 and 13 show the agent's maximum and final resource amounts.

We summarized next the main observations of the agent model simulations. Creative teams are preferred. Results show that reducing the final average paths lengths of the agents (columns 2-3) correlates to their capability of accessing resources (columns 12-13). The path length is high if an agent runs out of resources, e.g., Agent 1 for case 0% (Column 12) or Agent 2 for case 50% (Column 13). More resources not only increase the number of responses of an agent, but also increase the likelihood of having strategies 2, 3, and 5. Strategies 2 and 3 bridge the gaps between clusters of similar responses, and Strategy 5 increase the height of the knowledge hierarchy. Columns 6 and 7 also illustrate this aspect. Bridging gaps between different clusters reduces the fragmentation (increases the clustering) of the responses. Finally, having more strategies 2 and 3 increases the total number of clauses (columns 4 and 5), as they combine the clauses in different responses. Results show that case 25% produces the best access to resources of the two agents, which finally leads to the shortest final average paths lengths. The best results were obtained when knowledge maximization was selected

more often than cost minimization. Agents performed an increased number of concept combinations through Strategy 3 and generalizations through Strategy 5.

For the results in Table 1, Fig. 13 illustrates the distribution of frames for less (case (a)) and more creative pairs (case (b)). The numbering of the frames is as follows: (1) - *what* frame, (2) - *consequence*, (3) - *why*, (4) - *how*, (5) - *who*, (6) - *for who*, (7) - *where*, and (8) - *when*. Black rectangles are frames that existed in the agent’s local memory, and pink rectangles are new frames that were combined to create responses. In case (a), Agent 1 produced less novel solutions. The average shortest paths of frame similarities was about twice larger than the average shortest paths for Agent 2. Similar to the results in Figs. 11 and 12 for the experiment with human subjects, Fig. 13 shows that more novel solutions include a broader frame variety. These are created by more “zig-zagging” between different frame types through consecutive responses. In Fig. 13(b), the two agents cover well each frame kind, except frames *why*. These frames however are important to justify a solution, and less to decide its novelty. In contrast, Fig. 13(a) shows fewer frames *how*, *who*, and *for who*.

We detailed next the differences in decision making between Agent 1 (low novelty) and Agent 2 (high novelty) (Fig. 13(a)). Agent 1 runs fairly fast out of resources (by iteration 14), even though it selects the conservative policy to minimize the amount of resources used to create an output. The first steps include incremental refinement through Strategy 1 and combining detailed frames received from agent 2 with its own detailed frames (in the local memory) through Strategy 2. The ratio of received rewards and spent effort is low. In contrast, Agent 2 has enough resources for all iterations, because it received large rewards during the first iterations. Agent 2 adopts the strategy of maximizing knowledge, which even though has a

higher cost produces higher returns. The agent performs a moderate number of generalizations (Strategy 5) and few abstract concept combinations (Strategy 3). Agent 2 mostly adopts the idea provided by Agent 1, but does not incrementally refine its own ideas (Strategy 1). Also, it has been noticed that the bridging between different idea clusters is realized at the beginning of the evolutionary process when enough resources are available to the agents.

Due to the nature of the used models for cost and rewards (e.g., Bass model and exponential model, respectively) [51], agents cannot increase their resources beyond iteration 15. The main reason is that there are no more new solutions being generated for higher abstraction levels. These solutions receive higher rewards, because they are more flexible in spawning incremental refinements. Simulations show that high rewards resulted if an agent adopted the other agent’s idea, which it then combined with its own, original clauses. These solutions scored high on the rewards scale while requiring less resources due to their similarity with the previous solutions of the agent. Incremental changes (Strategy 1) created rewards only during the initial steps, while concept combinations through Strategy 2 or 3 offered rewards for longer sequences of steps. Simulations for longer number of iterations (i.e. 60 and 100 steps) support this observation.

Similar results were obtained in the case in which agents shared an increased number of starting ideas in their local memory (e.g., 25% and 50%). Having more common frames increased (in a solution) the number of frames associated to the *what* frame, however the average shortest path between frames did not decrease, if one agent run out of resources. Further increasing the percentage of shared ideas reduced the diversity of the frames associated to a *what* frame. The upper similarity that is perceived by an agent for distinct concepts had similar effects.

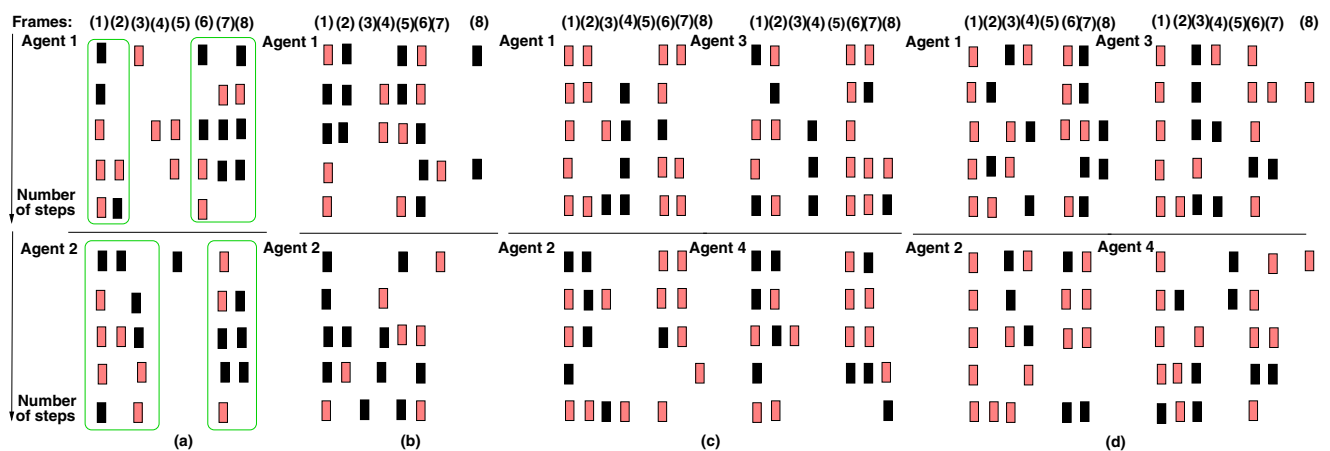


Fig. 13 Solution space evolution for (a) less creative and (b) more creative pairs, and for (c) less creative and (d) more creative groups of four members (black rectangles are frames from local memory and pink rectangles are added frames)

A similar situation exists for groups of four members (Fig. 13(c) and (d)). The frames of groups with less novel solutions are more clustered, as shown in Fig. 13(c). Frames are more uniformly distributed for groups with solutions of higher novelty, like in Fig. 13(d). As groups run faster out of resources than pairs, it is less often that all agents in a group achieve small shortest paths between their solution frames. There are no concept combinations through Strategy 3 and no generalizations using Strategy 5. Solutions are mostly at lower abstraction levels, involving concrete frames. Members who have initially more different frames are less likely to create solutions with frames connected to the solutions of the other agents. Hence, their shortest paths are longer than those of the latter agents.

7 Discussion

Responses of groups with high novelty continuously shifted between different frame types (*what*, *who*, *for who*, and *how* frames). Shifts from one category to another were mostly achieved through Strategy 3 applied to more abstract concepts. This observation suggests that the cognitive effort was smaller when combined concepts were more abstract. In contrast, responses of groups with low novelty stayed mostly in a single category, like *how* frame. Hence, creative groups sampled more the meaning of a concept, features, or sub-categories in terms of the goal facets, but without considering existing trade-offs. Group interactions included rare cases when a member used within his/her own solution context a certain change suggested by another member, or the case in which a change previously used by the member was applied again but to the response of another member. The two cases describe situations in which knowledge was moved across the contexts of different responses. Members did not explore the connections between solution facets. Participants mainly considered direct connections between goals, concepts, and outcomes, and neglected situations that can emerge due to the interactions between two or multiple changes that were made to a response.

The previous observations were also supported by simulations of the TM implementation. Agents obtained smaller average shortest paths between the frames of their solutions, hence a better coverage of the solution space, if they balanced well the effort (cost) to create responses and their available resources. Agents that focused only on incremental extensions ran out of resources, even though the cost of such extensions was small. The best approach was to gather resources during the initial steps by maximizing the produced knowledge through strategies that combined concepts, like strategies 2 and 3. All agents achieved small average shortest paths if they were able to implicitly distribute their frames, so that each of them

maintained a number of unique frames while sharing some frames with the other agents. This was more likely, if agents shared a moderate number of frames of their local knowledge, or if a lower threshold was adopted when two frames were considered similar. Increasing the number of simulation steps did not produce significant differences. Also, increasing the number of agents in a group reduced the likelihood that all achieved a good coverage of the solution space, as providing enough resources to all of them was hard.

There was some similarity at more abstract levels between the responses of different groups, but the similarity decreased as more details were added. The novelty of the responses was higher, if the group identified a larger variety that instantiated the abstract concept. A higher variety was more likely to emerge, if members had local knowledge that partially overlapped, hence bordered a solution space region in which the differences in responses could be addressed without triggering strong opposing responses, like those for Strategy 4. Such a case was more likely, if the proposed differences were due to different trade-offs considered by participants, while sufficient flexibility was maintained for the responses of others. The response novelty was lower when participants enumerated concrete features and real examples, even if the number of such instances was higher. Even for a certain overlapping between the meaning (semantics) of real examples, there was less triggering of new responses that bridge the meaning of the real examples. It can be argued that the real examples were concrete, previous experiences. Participants with similar experiences were more likely to enumerate or to agree on the same concrete features. Solution spaces with average novelty had a higher fragmentation than those for group responses with high or low novelty. However, there is less elaboration of the variants (e.g., the concept branches in the graph representation), which suggests that the local knowledges of members overlapped less with each other. Also, there was a higher semantic distance between the enumerated variants. Finally, group responses had low novelty when members had opposing valences on an issue, hence they pursued a sequence of contradicting steps (Strategy 4) or when they fully agreed with each other, hence creating a large number of reinforcements. Novelty was also reduced, if members used many synonyms to denote the same concept or feature.

Group members were rarely aware of the trade-offs involved in a response. For example, the solution to have free health care conflicted with the economic costs of the solution, the quality of the solution, and its scalability for a large population. Participants were usually concerned only about one main goal while ignoring the other problem facets. Tackling a broad concept from the different perspectives of a trade-off was a way of creating a larger variety of similar responses using strategies 2

and 3. However, in other cases, the existence of a trade-off had a negative impact, if the competing aspects had a high priority for two participants. Not being aware of the intrinsic trade-off of the solution was likely to trigger using Strategy 4, through which incompatible solutions were proposed. This suggests that the existence of trade-offs for a response frames it as an ill-defined problem. Hence, it encourages using solving specific strategies for such problems. Moreover, the same participants might have made inconsistent decisions, i.e. they argued for a certain solution considering a facet of the trade-off, or argued against the same solution, if another facet was analyzed. For example, a participant suggested free healthcare as a global solution to improve healthcare, but then in a different context, proposed that users should be held accountable if they misuse the insurance.

There can be different inconsistencies between the beliefs of each participant, or the beliefs of a participant and the group. Beliefs tend to stay unchanged during the entire experiment showing that the received responses, including using frames *why* and Strategy 4, did not change the beliefs of members. Inconsistencies occurred when more details were added to a general concept, or if a new goal facet was considered by a member.

Groups produced solution spaces with different degrees of fragmentation (clustering) with respect to the nature and similarity of the concepts (frames) in their solutions. Groups that originated responses with the highest and lowest novelty had less fragmentation (e.g., fewer number of abstract concepts to which responses belong), while median groups created responses of higher fragmentation. During interaction, group members aligned their responses with responses of other members. Conditions, constraints, and concrete situations under which an alignment was acceptable were explicitly or implicitly identified. The analysis also pointed out elements that members associated to certain responses, including attention and priority (which are indicators of the associated emotions and expected utility), inconsistencies in using beliefs during reasoning, and flexibility in accepting different beliefs. Some responses showed the emergence of stable, invariant relations between certain concepts, e.g., health insurance provided by the employer. The meaning of a concept related with the meaning of other concepts in a fixed way, which was independent on the response variations. Hence, it is reasonable to conclude that a response set can be characterized by a pair of components: the first component indicates the invariant part, and the second component the variations (e.g., disjoint concepts) of the responses.

It can be argued from the point of view of mathematical logic that responses with multiple changes cannot produce superior results than those with a single change, as those with multiple changes are a special case of those with a

single change [76]. However, this observation is valid only if the multiple changes are not correlated with each other through the problem definition. For example, providing cheaper ER for everyone might be unsustainable due to large costs. The high cost is a bottleneck of the solution. Reducing the ER cost involves reducing the cost of regular treatment that is likely to lower its quality. Reducing cost without reducing quality can be achieved only if the number of ER visits is reduced, such as by better prevention, or if new technology is created that can lower cost, like by increasing automation. The relation between cost and technology or prevention explains responses with higher novelty than responses with a single change.

8 Conclusions

This paper proposes a novel theoretical model (TM) to describe the problem solving process of open-ended problems in small groups. TM presents declaratively the connection between the characteristics of the participants, the interactions in a group, the evolution of the group's knowledge, and the overall novelty of the responses created by the group. The model presents each participant as an agent with the following components: (i) local knowledge, (ii) a way of interpreting the knowledge, (iii) priorities and preferences, (iv) resources, (v) motivation, and (vi) emotions associated to goals and concepts. Agents select among five solving strategies to generate new knowledge and responses. Group responses form a solution space, in which responses are clustered into categories based on response similarity and organized on abstraction levels. The solution space includes concrete features and samples, as well as causal sequences that indicate how concepts are logically connected with each other. Understanding the process of producing responses to open-ended problems solved in small groups is important for many modern domains, like health care, manufacturing, banking, investment, and other.

The model was used to explain how member characteristics, e.g., the degree to which their beliefs are similar, relate to the solution novelty of the group. Experiments highlighted a few interesting observations. Groups that originated responses with the highest and lowest novelty identified fewer abstract concepts, while the median groups created responses pertaining to more concepts. The response novelty was higher, if group members had personal knowledge that partially overlapped without triggering opposing responses. The response novelty was lower when participants enumerated concrete examples, even when more such examples were offered. Participants with similar experiences were more likely to agree with each other about similar concrete features. Solution spaces with average novelty had less elaboration of the variants, which suggests that

members had less overlapping knowledge. Finally, group responses had low novelty when members had opposing beliefs on an issue (hence they pursued a sequence of contradictions) or when they fully agreed with each other. Responses of groups with high novelty continuously shifted between different frame types (*what, who, for who* and *how* frames). In contrast, responses of groups with low novelty stayed mostly in a single category, such as *how* frame. This observation suggests that more creative groups sampled to a higher degree the meaning of a concept, features, or subcategory in terms of the goal facets, even though without considering the specifics of the trade-offs that exist. This insight was supported also by the simulation of the TM implementation, as discussed in Section 7.

Members were rarely aware of the trade-offs involved in responses. The connections between facets were not explored. Participants were concerned only about one main goal while ignoring the other problem facets. Tackling an abstract concept from the various perspectives of a trade-off can be a way to produce a larger variety of similar responses. The observation suggests that participants mainly considered direct connections between goals, concepts, and outcomes, and neglected any situations that can emerge due to the interactions between two or multiple changes that were made to a response. Also, different kinds of inconsistencies can exist between the beliefs of a participant and those of the group. Beliefs tend to stay unchanged during the experiment.

TM agents use five strategies to extend the solution space of an open-ended problem. The limitations of transformational systems, e.g., systems that produce new solutions by applying a set of rules, have been discussed in [27, 81]. It is however important to note that transformational concepts have been a cornerstone in high-level synthesis of digital electronic designs, in which behavioral descriptions are automatically converted into circuit designs [100–102]. It has arguably been the most successful application of transformation-based design in engineering. However, transformation-based design has been less successful for applications which are less algorithmic, like analog and mixed-signal design [103–105]. Previous work has suggested using Genetic Algorithms [33, 34], but produced solutions are less creative than those generated by human designers [35]. Our recent approach proposed a cognitive architecture that mimicks the cognitive activities of designers [38]. It uses an evolutionary process that uses five transformation rules similar to those presented in this paper. In our experience, we don't think that the rules are a significant barrier in limiting the nature of the created solutions. Rules allow moving top-down and bottom-up through the solution space representations, abstractions, instances, and concept combinations. Instead, we think that the main limitation in creating more novel solutions is in that the system cannot

produce novel concepts (i.e. building blocks), beyond the concepts that already occurred in designs. For the discussed TM, it means that the solutions will not include other words than those that appear in previous solutions. We performed preliminary work on a system that would automatically create new concepts (building blocks) [36], but our future work will address this issue in more detail.

References

- Cummings J, Kiesler S, Zadeh R, Balakrishnan A (2013) Group heterogeneity increases the risks of large group size: a longitudinal study of productivity in research groups. *Psychol Sci* 24(6):880–890
- Hong L, Page S (2004) Groups of diverse problem solvers can outperform groups of high-ability problem solvers. *Proc Natl Acad Sci* 101(46):16385–16389
- Malone T, Laubacher R, Dellarocas C (2010) The collective intelligence genome. *MIT Sloan Manag Rev* 51(3):1–6
- Nijstad B, Stroebe W (2006) How the group affects the mind: a cognitive model of idea generation in groups. *Personal Soc Psychol Rev* 10(3):186–213
- Paulus PB, Coskun H (2012) *Group creativity* Levine JM (ed), Elsevier, Amsterdam
- Uzzi B, Mukherjee S, Stringer M, Jones B (2013) Atypical combinations and scientific impact. *Science*, 342(6157)
- Daniels B, Ellison C, Krakauer D, Flack J (2016) Quantifying collectivity. *Curr Opin Neurobiol* 37(3):106–113
- Jones B, Wuchty S, Uzzi B (2008) Multi-university research teams: Shifting impact, geography, and stratification in science. *Science* 322:1259–1262
- Kluge A, Sauer J, Schuler K, Burkolter D (2009) Designing training for process control simulators: a review of empirical findings and current practices, theoretical issues in ergonomics. *Science* 10:489–509
- Liu X, Li H, Doboli A, Doboli S (2019) Modeling and extraction of insight information on the evolution of community-level, creative problem solving. In: Paulus P, Nijstad B (eds) *Oxford handbook of group creativity and innovation*. Oxford University Press, Oxford
- Paulus PB, Korde R (2013) How to get the most creativity and innovation out of groups and teams. In: Chan J, Thomas K (eds) *Handbook of research on creativity*. Edward Elgar Publishing, Cheltenham, pp 493–507
- Thomas L, Lleras A (2009) Covert shifts of attention function as an implicit aid to insight. *Cognition* 111(2):168–174
- De Dreu CKW, Nijstad BA, Bechtoldt MN, Baas M (2011) Group creativity and innovation: a motivated information processing perspective. *Psychol Aesthet Creat Arts* 5(1):81–89
- Granovetter M (1973) The strength of weak ties. *Am J Sociol* 78(6):1360–1380
- Kozlowski SWJ, Ilgen DR (2006) Enhancing the effectiveness of work groups and teams. *Psychol Sci Public Int* 7:77–124
- Montjoye Y, Stopczynski A, Shmueli E, Pentland A, Lehmann S (2014) The strength of the strongest ties in collaborative problem solving. *Nature* 4:1–6
- Mullen B, Johnson C, Salas E (1991) Productivity loss in brainstorming groups: a meta-analytic integration. *Basic Appl Soc Psychol* 12:3–23
- Nakui T, Paulus PB, van der Zee KI (2011) The role of attitudes in reactions to diversity in work groups. *J Appl Soc Psychol* 41:2327–2351

19. Sober E (1984) *The nature of selection: evolutionary theory in philosophical focus*. University of Chicago Press, Chicago
20. Portmann E, Meier A (2010) A fuzzy grassroots ontology for improving weblog extraction. *J Digit Inf Manag* 8(4):276–284
21. Fellbaum C (1998) *Wordnet. an electronic lexical database*. MIT Press, Cambridge
22. Jackson SE (1996) The consequences of diversity in multidisciplinary work teams. In: West MA (ed) *Handbook of work group psychology*. Wiley, Chichester, pp 53–75
23. West MA, Richter AW (2008) *Climates and cultures for innovation and creativity at work*. In: Zhou JJ, Shalley CE (eds) *Handbook of organizational creativity*. Taylor & Francis Group, New York, pp 211–236
24. Rifkin J (2014) *The zero marginal cost society*. MacMillan Publishers, New York
25. Doboli A, Umbarkar A (2014) The role of precedents in increasing creativity during iterative design of electronic embedded systems. *Des Stud* 35(3):298–326
26. Coursey M, Williams B, Kenworthy J, Paulus P, Doboli S (2018) Divergent and convergent group creativity in an asynchronous online environment. *J Creat Behav*, <https://doi.org/10.1002/jocb.363>:1–14
27. Boden M (1998) Creativity and artificial intelligence. *Artif Intell* 103(1):347–356. Artificial intelligence 40 years later
28. Hung E, Choy C (2013) Conceptual recombination: a method for producing exploratory and transformational creativity in creative works. *Knowl-Based Syst* 53:1–12
29. Li H, Liu X, Jiao F, Doboli A, Doboli A (2018) Innova: a cognitive architecture for computational innovation through robust divergence and its application for analog circuit design. *IEEE Trans CADICS* 37(10):1943–1956
30. Varshney L (2019) Mathematical limit theorems for computational creativity. *IBM J Res Develop* 63(1):2:1–2:12
31. Varshney L, Pintel F, Varshney K, Bhattacharjya D, Schorgendorfer A, Chee Y (2019) A big data approach to computational creativity: the curious case of chef watson. *IBM J Res Develop* 63(1):7:1–7:18
32. Wang K, Nickerson J (2017) A literature review on individual creativity support systems. *Comput Hum Behav* 74:139–151
33. Kruskamp W, Darwin DL (1995) Cmos opamp synthesis by means of genetic algorithm. In: *Proceeding design automation conference* pp 433–438
34. McConaghy T, Palmers P, Gao P, Steyaert M, Gielen G (2009) *Variation-aware analog structural synthesis*. Springer, New York
35. Ferent C, Doboli A (2011) Measuring the uniqueness and variety of analog circuit design features. *Integr VLSI J* 44(1):39–50
36. Li H, Jiao F, Doboli A (2016) Analog circuit topological feature extraction with unsupervised learning of new substructures. In: *Proceeding design automation and test in europe conference*, pp 1509–1512
37. Ferent C, Doboli A (2013) Symbolic matching and constraint generation for systematic comparison of analog circuits. *IEEE Transactions on CADICS* 32(4):616–629
38. Jiao F, Li H, Doboli A (2018) Modeling and extraction of causal information in analog circuits. *IEEE Trans CADICS* 37(10):1915–1928
39. Jiao F, Montano S, Doboli A (2015) Knowledge-intensive, causal reasoning for analog circuit topology synthesis in emergent and innovative applications. In: *Design, automation and test in europe conference*
40. Jiao F, Doboli A (2016) Causal reasoning mining approach to analog circuit verification. *Integr VLSI J* 55:376–383
41. Croson R, Gachter S (2010) The science of experimental economics. *J Econ Behav Org* 73(1):122–131
42. Serfaty D, Entin E, Johnston J (1998) *Team coordination training Making Decisions under stress: Implications for Individual and Team Training*
43. Orasanu JM, Fischer U (1992) Team cognition in the cockpit: linguistic control of shared problem solving. In: *Proceedings of the 14th annual conference of the cognitive science society*. Hillsdale, N J Erlbaum
44. Salas E, Cannon-Bowers JA (1997) Methods, tools, and strategies for team training. In: Quinones M, Ehrenstein A (eds) *Training for a rapidly changing workplace: applications of psychological research*. American Psychological Association, Massachusetts, pp 249–279
45. Rasmussen J, Nixon P, Warner F (1990) Human error and the problem of causality in analysis of accidents [and discussion]. *Philos Trans Royal Soc London. Series B, Biol Sci* 327(1241):449–462
46. Garriga H, von Krogh G, Spaeth S (2013) How constraints and knowledge impact open innovation. *Strat Manag J* 34(9):1134–1144
47. Hagemann V, Kluge A (2017) Complex problem solving in teams: the impact of collective orientation on team process demands. *Front Psychol* 8
48. Shi X, Leskovec J, McFarland DA (2010) Citing for high impact. In: *Proceedings of the 10th annual joint conference on digital libraries, JCDL '10*, New York, NY, USA. ACM, pp 49–58
49. Kuhn T (1962) *The structure of scientific revolutions*. University of Chicago Press, Chicago
50. Rodriguez-Sickert C, Cosmelli D, Claro F, Fuentes M (2015) The underlying social dynamics of paradigm shifts. *PLOS ONE* 10(9):1–12 09
51. Doboli A, Doboli S (2014) Iflows: A novel simulation model for predicting the effectiveness of a research community. In: *Proceedings of the 2014 IEEE symposium series on computational intelligence (SSCI)*, Orlando, FL
52. Gabora L, Tseng S (2017) The social benefits of balancing creativity and imitation: evidence from an agent-based model. *Psychology of Aesthetics, Creativity, and the Arts*, 04
53. Perry-Smith J, Shalley C (2003) The social side of creativity: a static and dynamic social network perspective. *Acad Manag Rev* 28(1):89–106
54. Hung A, Plott C (2001) Information cascades: replication and an extension to majority rule and conformity-rewarding institutions. *Am Econ Rev* 91(5):1508–1520
55. Bonabeau E (2002) Agent-based modeling: methods and techniques for simulating human systems. *Proc Nat Academy Sci* 99(suppl 3):7280–7287
56. Sun R (2001) Cognitive science meets multi-agent systems: a prolegomenon. *Philos Psychol* 14(1):5–28
57. DeAngelis D, Diaz S (2019) Decision-making in agent-based modeling: a current review and future prospectus. *Front Ecol Evol* 6:237
58. Lakatos I (1978) *The methodology of scientific research programmes: philosophical papers, vol I*. Cambridge University Press, Cambridge
59. Feyerabend P (1978) *Science in a free society* nlb
60. Nowak A, Szamrej J, Latane B (1990) From private attitude to public opinion: a dynamic theory of social impact. *Psychol Rev* 9(3):362–376
61. Bhattacharyya S, Ohlsson S (2010) Social creativity as a function of agent cognition and network properties: a computer model. burlington, MA: ALPHATECH, *Social Networks*(32) pp 263–278
62. Evoc LG (2008) A computer model of the evolution of culture
63. Spencer GM (2012) Creative economies of scale: an agent-based model of creativity and agglomeration. *J Econ Geogr* 12(1):247–271

64. Rodriguez A, Kim B, Lee J-M, Coh B-Y, Jeong M (2015) Graph kernel based measure for evaluating the influence of patents in a patent citation network. *Expert Syst Appl* 42(3):1479–1486
65. Le Ber F, Chouvet M-P (1999) An agent-based model for domain knowledge representation. *Data Knowl Eng* 29(2):147–161
66. An G (2015) Introduction of a framework for dynamic knowledge representation of the control structure of transplant immunology: employing the power of abstraction with a solid organ transplant agent-based model. *Front Immunol* 6:561
67. Henrich J, Boyd R (2002) On modeling cognition and culture: why replicators are not necessary for cultural evolution. *J Cogn Cult* 2:87–112
68. Granovetter M (1983) The strength of weak ties: a network theory revisited. *Sociol Theor* 1:201–233
69. Burt RS (1997) The contingent value of social capital. *Adm Sci Q* 42:339–365
70. Scott J (1991) *Social network analysis: a handbook* london: Sage
71. Bouty I (2000) Interpersonal and interaction influences on informal resource exchanges between r&d researchers across organizational boundaries. *Acad Manag J* 43(1):50–65
72. Hansen MT (1999) The search-transfer problem: weak ties in sharing knowledge across organization subunits. *Adm Sci Q* 44:82–111
73. Pan W, Dong W, Cebrian M, Kim T, Fowler J, Pentland A (2012) Modeling dynamical influence in human interaction: using data to make better inferences about influence within social systems. *IEEE Signal Proc Mag* 29(2):77–86
74. Wohl J (1981) Force management decision requirements for air force tactical command and control. *IEEE Trans Syst Man Cybern* 11(9):618–639
75. Nokes-Malach T, Meade M, Morrow D (2012) The effect of expertise on collaborative problem solving. *Thinking Reasoning* 18(1):32–58
76. Stenning K, Van Lambalgen M (2008) *Human reasoning and cognitive science*. MIT Press, Cambridge
77. Ferent C, Doboli A (2013) An axiomatic model for concept structure description and its application to circuit design. *Knowl-based Syst* 45:114–133
78. Doboli A, Umbarkar A, Doboli S, Betz J (2015) Modeling semantic knowledge structures for creative problem solving: studies on expressing concepts, categories, associations, goals and context. *Knowl-based Syst* 78:34–50
79. Umbarkar A, Subramanian V, Doboli A, Doboli S (2014) Two experimental studies on creative concept combinations in modular design of electronic embedded systems. *Des Stud* 35(1):80–109
80. Jiao F, Montano S, Ferent C, Doboli A, Doboli S (2015) Analog circuit design knowledge mining: discovering topological similarities and uncovering design reasoning strategies. *IEEE Trans CADICS* 34(7):1045–1059
81. Gero JS (1990) Design prototypes: a knowledge representation schema for design. *AI Mag* 4(11):26–36
82. Manenti L, Sartori F (2011) Using visual information analysis to explore complex patterns in the activity of designers. *Metadata Support Retriev Revise Solut Case-based Reason* 6(3/4):185–194
83. Cash P, Stankovic T, Storga M (2014) Using visual information analysis to explore complex patterns in the activity of designers. *Des Stud* 35(1):1–28
84. Taura T, Yamamoto E, Fasiha M, Goka M, Mukai F, Nagai Y, Nakashima H (2012) Constructive simulation of creative concept generation process in design: a research method for difficult-to-observe design-thinking processes. *J Eng Des* 23(4):297–321
85. Calegari S, Loregian M (2006) Using dynamic fuzzy ontologies to understand creative environments. In: Proceedings of the international conference on flexible query answering systems, pp 404–415
86. Jurasica I, Mysopoulos J, Yu E (2004) Ontologies for knowledge management: an information systems perspective. *Knowl Inf Syst* 6:380–401
87. Shadbolt N, Hall W, Berners-Lee T (2006) The semantic web revisited. *IEEE Intelligent Systems*
88. Gupta A, Forgy C, Newell A, Wedig R (1986) Parallel algorithms and architectures for rule-based systems. In: Proceedings of the 13th annual international symposium on computer architecture, pp 28–37
89. Getoor L, Friedman N, Koller D, Pfeffer A (2001) Learning probabilistic relational models. In: Dzeroski S, Lavrac N (eds) *Relational data mining*. Springer, Berlin, pp 18–34
90. Budaev S, Giske J, Eliassen S (2018) Aha: a general cognitive architecture for darwinian agents. *Biol Inspired Cognit Architect* 25:51–57
91. Anderson J (1995) *Learning and memory. an integrated approach*. Wiley, New York
92. Gigerenzer G, Todd P (1999) *The ABC research group. simple heuristics that make us smart*. Oxford University Press, Oxford
93. Elliott C (1992) The affective reasoner: a process model of emotions in a multi-agent system. The institute for the learning sciences technical report #32, Ph.D. Thesis, Northwestern University
94. Laajaj R, Macours K, Pinzon Hernandez D, Arias O, Gosling S, Potter J, Rubio-Codina M, Vakis R (2019) Challenges to capture the big five personality traits in non-weird populations. *Science Advances*, 5(7)
95. Marinier R, Laird J (2004) Toward a comprehensive computational model of emotions and feelings. *ICCM*
96. Osuna E, Rodriguez L-F, Gutierrez-Garcia OJ, Castro L (2020) Development of computational models of emotions: a software engineering perspective. *Cogn Syst Res* 60:1–19
97. Becker-Asano C, Wachsmuth I (2010) Affective computing with primary and secondary emotions in a virtual human. *Autonomous Agents and Multi-Agent Systems*, 20(32)
98. Gadanho S (2003) Learning behavior-selection by emotions and cognition in a multi-goal robot task. *J Mach Learn Res* 4:385–412 01
99. Humphrey S, Aime F, Cushenbery L, Hill A, Fairchild J (2017) Team conflict dynamics: implications of a dyadic view of conflict for team performance. *Organ Behav Hum Decis Process* 142:58–70
100. Gupta R (1995) *Co-Synthesis Of hardware and software for digital embedded systems*. Kluwer
101. Vanhoof J, Van Rompaey K, Bolsens I, Goossens G, De Man H (1993) *High-level synthesis for real-time digital signal processing*. Springer, New York
102. Doboli A, Vemuri R (2001) Integrated high-level synthesis and power-net routing for digital design under switching noise constraints. In: Proceeding of the 38th design automation conference
103. Doboli A, Vemuri R (2003) Exploration-based high-level synthesis of linear analog systems operating at low/medium frequencies. *IEEE Trans CADICS* 22
104. Doboli A, Vemuri R (2003) Behavioral modeling for high-level synthesis of analog and mixed-signal systems from vhdl-ams. *IEEE Trans CADICS* 22
105. Doboli A, Tang H, Zhang H (2003) Towards high-level synthesis of analog and mixed-signal systems from vhdl-ams specifications - a case study for a sigma delta analog to digital converter. In: Forum on specification and design languages (FDL)

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.