

Computational Creativity in Naturalistic Decision-Making

Magnus Jändel

Swedish Defence Research Agency
Stockholm, SE-16490, Sweden
magnus.jandel@foi.se

Abstract

Creativity can be of great importance in decision-making and applying computational creativity to decision support is comparatively feasible since novelty and value often can be evaluated by a reasonable human effort or by simulation. A prominent model for how humans make real-life decisions is reviewed and we identify and discuss six opportunities for enhancing the process with computational creativity. It is found that computational creativity can be employed for suggesting courses of action, unnoticed situation features, plan improvements, unseen anomalies, situation reassessments and information to explore. For each such enhancement opportunity tentative computational creativity methods are examined. Relevant trends in decision support research are related to the resulting framework and we speculate on how computational creativity methods such as story generation could be used for decision support.

Introduction

Creativity and decision-making Before the battle of Austerlitz 1805, Napoleon deceptively maneuvered to create the impression of weakness and indecision in the French forces. The opposing Russo-Austrian army took the bait, attacked and fell into a carefully prepared trap resulting in a crushing defeat. Detecting deception requires an act of creativity where the reality of the situation is discerned behind a screen of trickery. European history could have taken a different turn with more creative leadership on the Russian and Austrian side. Leaders of today are likewise challenged to be more creative. Given the progress of computational creativity in other fields, it is therefore interesting to pursue its application to decision-making.

Computational creativity for decision support The key problem in computational creativity is how to automatically assess the novelty and creative value of an idea, concept or artifact that has been generated by computational means (Boden, 2009). This is a very difficult problem in art where novelty is judged by comparing to extensive traditions and evaluation would engender implementation of computational esthetic taste. Decision support is fundamentally less challenging. Novelty is often judged against a reasonably short list of options that are known to the decision makers and value is evaluated by analyzing how the idea works in the situation at hand. Computer simulations are increasing-

ly employed for assisting decision makers and it is often quite feasible to use simulations for automatic evaluation of suggested ideas. Given the comparative straightforwardness of applying computational creativity to decision support, it appears that there are surprisingly few applications. Some of these are discussed and put in context after that we have introduced the framework that is the main result of this paper.

Decision-making models Applying computational creativity to any given area of decision-making requires substantial domain knowledge and it is often difficult to see how methods generalize to other domains. Our strategy is therefore to identify generic approaches by analyzing how formal decision-making models can be extended to include computational creativity techniques.

Somewhat simplified, decision-making models can be partitioned into two general classes: rational models and naturalistic models. The former prescribes how decisions ought to be made while the latter describes how people really make decisions. Many naturalistic models surpass, however, their purely descriptive origins and offer some suggestions on how intuitive decision-making can be improved.

In the following two sections we analyze how computational creativity tools can extend rational and naturalistic decision-making models respectively with a strong focus on a particularly prominent naturalistic model.

Rational decision-making models

Rational decision-making models provide methods for how to optimally select an action from a set of alternative actions (Towler 2010). In utility-based decision-making it is for example assumed that each action leads to a set of outcomes and that the probability of each outcome is known or can be estimated. Furthermore, each outcome has a utility which is a real-valued variable and the task of the decision-maker is to select the action that is most likely to optimize the utility of the outcome. Other rational schemes extend this approach to cases with multiple objectives and multiple constraints (Triantaphyllou, 2002). The main US army Military Decision-Making Process (MDMP) is for example essentially a rational process where it is required

that at least three different courses of action should be compared.

Rational decision-making models characteristically give little guidance on how to generate the set of action alternatives although it is tacitly assumed that more alternatives make for better decisions. Since the 1980s, it has been claimed that decision makers typically don't employ rational models (Kahneman, Slovic, and Tversky 1982) and it appears further that leaders don't find rational models to be efficient (Yates, Veinott and Patalano 2003) and that generating more alternatives actually can be detrimental for decision quality (Johnson and Raab 2003).

Computational creativity could assist rational decision-making by suggesting criteria, enriching the set of action alternatives, envisioning possible outcomes of actions and suggesting factors that should be considered in mental or computer simulations. The methods that could be employed for this are often quite similar to corresponding methods in naturalistic decision-making which is the focus of this paper.

Extended naturalistic decision-making model

Naturalistic decision-making models are inspired by research in how decisions are made in domains such as business, firefighting and in military areas. Investigations indicate that experienced and effective leaders evaluate the nature of the situation intuitively and rarely consider more than one course of action (Klein 2003).

Figure 1 summarizes a leading naturalistic decision-making model: the recognition-primed model (RPD). For the moment, please ignore the symbols **CC1**, **CC2**, ... **CC6**. This paragraph briefly reviews work of Klein and coworkers on RPD (Klein, Calderwood and Clinton-Cirocco 1986; Ross, Klein, Thunholm, Scmitt, and Baxter, 2004; Klein 2008). The experienced decision maker evaluates the state of affairs and will normally recognize a familiar type of situation. Recognition means that the relevant cues or indicators are pinpointed; expectancies on how the situation will appear and unfold are identified; what kind of goals that are reasonable to pursue are recognized and a typically short list or courses of actions are found. In the following we use *course of action* or *action* to designate the conceptual level of a top-level plan that if implemented will consist of a chain of component actions or plan elements. As a reality-check the expectancies are analyzed and compared to available information. Any anomalies found trigger an iteration of the recognition process where more information may be sought and the situation is reassessed, sometimes leading to a major shift in how the situation is perceived. Eventually, the decision maker arrives to a satisfactory anomaly-free situation-recognition and selects the most promising course of action for scrutiny. The consequences of performing the selected course of action is simulated either mentally or by computer. This may lead to that the course of action is rejected and another option is selected for a new round of simulation. Frequently it is found that the course of action is promising but that it has some unwanted consequences.

Rather than rejecting the course of action, decision makers will try to repair the plan by modifying the chain of plan elements that implement the course of action. It is implicit in Figure 1 that modified courses of actions are re-simulated. Eventually the decision-maker will find a satisfactory course of action which will be implemented. Note that the RPD process does not include a search for the optimal course of action, the optimal implementation or quantitative utility criteria. If the plan satisfies the recognized goals it is deemed to be ready for implementation.

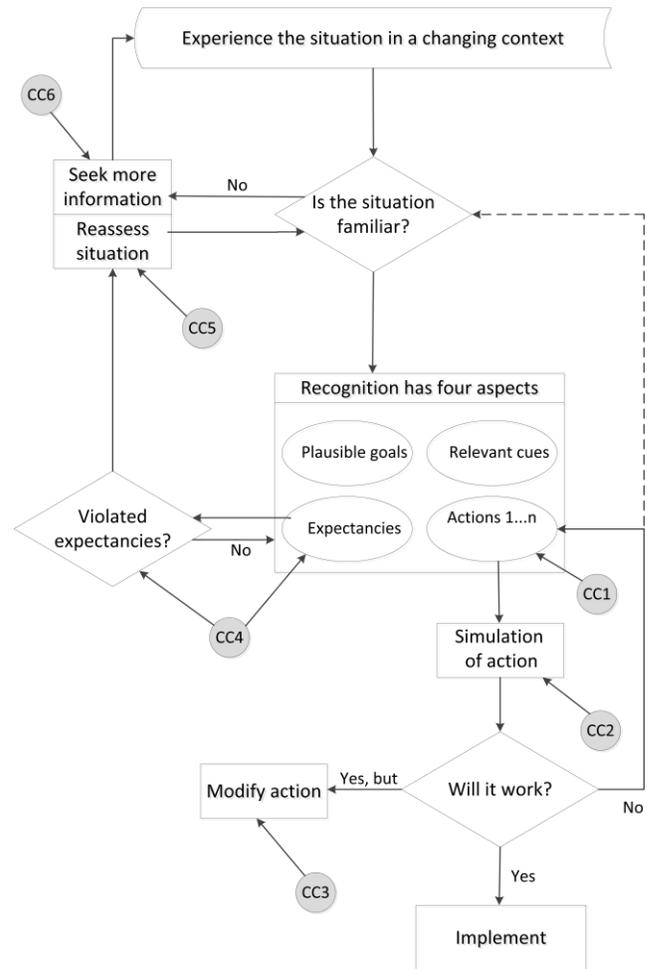


Figure 1. Computational creativity extensions to the recognition-primed model. Filled circles denote computational creativity agents. Everything else in the figure is quoted from Klein (2008).

How can decision makers that use RPD or some similar naturalistic decision-making model take advantage of computational creativity? In Figure 1, we mark six slots where a computational creativity agent could be plugged into the RPD process. The computational creativity agents are called **CC1**, **CC2**, ... **CC6** and these symbols are used in the following to highlight where the different computational creativity extensions are mentioned. For each agent we provide a mnemonic tag, discuss in which way it could improve decision-making and provide a sketch of at least

one creativity technology that could be applicable. Finally we provide a speculative example illustrating why creative input could be of great value in the current decision-making phase.

CC1 (proposing actions): Computational creativity could be used for suggesting a broader range of courses of action in a recognized situation. The **CC1** agent would work under the assumption that the situation and the relevant goals are correctly identified and that the creative task is to find unrecognized course of action alternatives that lead towards the “plausible goals” in Figure 1. The decision maker has identified the nature of the situation which will suggest a well-defined search space of actions. Some of the actions in the search space are explicitly known to the decision maker and would hence be found in the list of actions indicated in Figure 1. The RPD process evaluates listed actions. Creative suggestions should hence point to feasible actions that are significantly different from already listed actions. The **CC1** algorithm must define a similarity metric in action space and the list of actions that are known to the decision-maker should be available to the **CC1** agent so that it can avoid searching too close to known actions. Ideally, the **CC1** agent uses a simulation engine for confirming the approximate validity of courses of action but it might also be possible to fall back on human adjudication. Candidate courses of action that are far from known courses of action according to the metric and pass the simulation test are suggested to the decision maker and added to the known list. A government wanting to integrate an island population to the mainland society may for example consider courses of actions such as building a bridge, airport or ferry terminal. The **CC1** agent, realizing that known courses of action all relate to physical connectivity, may suggest courses of actions such as investing in telepresence or locating a new university to the island.

CC2 (proposing features): Simulations are never completely realistic but will always model some aspects of the situation at hand with higher fidelity than others and also ignore many other aspects. The **CC2** agent could suggest features that should be included in computer or mental simulations. Such ideas might be crucial for success since the acuity of the simulation is essential for the quality of the plan that implements the selected course of action. Consider for example a decision maker trying to control flooding caused by a burst dam. The core simulation would be concerned with modeling how actions influence the flow of water. A **CC2** agent searching historical records of floods could come up with the suggestion that modeling the spread of cholera might be important. The **CC2** agent could for example grade the importance of candidate features by measuring how often they are mentioned in news stories on flood-related events.

CC3 (reparing plans): Computational creativity could be used for suggesting how a promising but somewhat flawed plan can be repaired. Assume that simulation has exposed

at least one problem with the current course of action and that the decision makers have the mental or computational means for re-planning but are out of ideas. The task of the **CC3** agent is to provide an idea for how the problem can be solved. The main planning process can then use the idea for driving the next iteration of re-planning. Consider a case in which the main planning process is a planning algorithm (Ghallab, Nau, and Traverso 2004) that works by searching for a chain of plan elements that implements the course of action. Each plan element has a set of prerequisites and a set of consequences. The planner searches for chains of plan elements where all prerequisites are satisfied, the consequences at the end of the chains match the goals and the general direction of the plans is consistent with the currently considered course of action. If the planning algorithm fails to find a problem-free plan, the **CC3** agent could suggest a new plan element. This creative output is validated if the planner solves or alleviates the problem by using the suggested action element in a modified plan. The task of the **CC3** agent could be construed as search in the space of possible plan elements where the identified problem may be used for heuristic direction of the search. Note that the **CC3** agent should not be another planner that by explicit planning guarantees that the suggested plan element solves the problem. It is sufficient that suggested plan elements have a high probability of contributing to the solution. The **CC3** agent should obviously be aware of the present set of plan elements that are used by the main planner and avoid suggesting elements that are identical or very similar to currently known plan elements. A government may for example have selected reduction of the national carbon footprint as the chief course of action for environmental protection but fails to find a plan that reaches the target. A **CC3** agent could then suggest enhanced weathering, where crushed rock absorbs carbon dioxide, as a new plan element.

CC4 (identifying anomalies): Computational creativity could be used for identifying anomalous expectancies in the current perspective on the situation. When a decision maker has identified the situation as familiar it is often difficult to notice aspects of the situation that do not fit into the familiar context. It is crucial to find any anomalies since this might trigger a radical reassessment of the situation. The **CC4** agent is best applied when the decision-maker has exhausted the manual search for anomalous expectancies and is ready to proceed with action evaluation. A simple version of the combinatorial approach could explore the space of situation features searching for pairs of features that in combination stand out as anomalous. This could be done by investigating second-order attributes of the features and noting how combinations of attributes interact. The obscure features method (McCaffrey and Spector 2011) might be adapted for this purpose. Simulation methods that are used for evaluating actions could also be applied to examining anomaly candidates with validated anomalies escalated for human consideration. The Russian and Austrian leaders at Austerlitz would have benefited

from a **CC4** agent suggesting that Napoleon's uncharacteristic eagerness to negotiate and seemingly panicky abandonment of important positions were anomalies deserving serious attention.

CC5 (situation assessment): Supporting reassessment of the situation is the most challenging creative task. Imagine that the decision maker has noted a number of anomalies indicating that the present situation recognition is flawed but no viable alternative interpretations pop up in human minds. People are often locked into habitual trains of thought and this behavior is frequently aggravated by time pressure, fear and group-think. Computational creativity is free from such human frailties and might be able to suggest new ways of looking at the situation. A single idea might be enough for providing the Aha! experience that releases the intuitive power of the decision maker. A simple implementation of a **CC5** agent could use a library of case histories enshrining human expert findings in a broad range of circumstances. A sufficiently small volume of decision-making experience could, as noted by M. Boden (personal communication), advantageously be codified as a check list. **CC5** agents would be needed only in contexts in which the total span of assessment possibilities is large and inscrutable. A police officer leading the investigation of suspected arson in an in-door food market could for example benefit from the suggestion that spontaneous combustion of pistachio nuts might be an alternative perspective on the evidence (see Hill (2010) for further information on spontaneous combustion). The **CC5** agent would in this case use encyclopedic knowledge that pistachio nuts are a kind of food and that pistachio nuts are subject to spontaneous combustion combined with records of historical cases in which suspected arson has been found to be explained by spontaneous combustion.

CC6 (recommending information): Computational creativity could be used for suggesting what kind of information that could support reassessment or resolution of apparently anomalous expectancies. Decision makers often have access to vast archives and abundant streams of news and reports. Selecting what deserves attention is a difficult and sometimes creative task. Decision makers would be biased by their present understanding of the situation so the **CC6** agent might be able to provide a fresh perspective. The task of the **CC6** agent is quite similar to that of the **CC2** agent; it must explore the space of information sources and information aspects for the purpose of identifying novel and valuable pieces of information. A doctor confronted with anomalous symptoms could for example get suggestions from a **CC6** agent regarding which medical tests to apply.

Discussion and conclusions

In this section we will first discuss some current applications of computational creativity to decision support in relation to the framework described in the previous section

and then speculate on how selected approaches to computational creativity could be applied to decision support.

Examples of computational creativity in decision support Computer chess is probably the most advanced current application of computational creativity in decision support. Grand masters learn creativity in chess by studying how computers play (Bushinsky 2012). The main reason for this success is that chess is a very complex but deterministic game that readily can be simulated. The complexity of the game makes it possible for a computer to discover solutions that escape the attention of humans and simulation combined with heuristic assessment of positions enable automatic evaluation of computer generated solutions. The main components of creativity - novelty and value - are therefore attainable by chess programs. Referring to Fig. 1, we note that chess players use computational creativity mainly for suggesting courses of action (**CC1**) and repairing plans (**CC3**).

Tan and Kwok (2009) demonstrate how Conceptual Blending Theory (Fauconnier and Turner 2002) can be used for scenario generation intended for defense against maritime terrorism. The scenarios are examples of **CC5** agent output since they assist decision makers in assessing situations that may look peaceful and familiar at first sight but in which creative insights may reveal an insidious attack pattern.

Deep Green is a DARPA research program that aims for a new type of decision support for military leaders (Surdu and Kittka 2008) According to the Deep Green vision commanders should sketch a course of action using an advanced graphical interface while AI assistants work out the consequences and suggests how the plan can be implemented. The AI assistants are guided by thousands of simulations that explore how the situation could evolve and what factors that are important. According to our analysis in the previous section Deep Green seems to include development of **CC2** and **CC3** computational creativity agents although computational creativity is not explicitly mentioned in the Deep Green program.

Emerging tools Creative story generation could be turned into tools for decision support. Story generation techniques that spin a yarn connecting a well-defined initial state with a given final state (Riedl and Sugandh, 2008) could be used by **CC2** agents for suggesting improvements in simulations forecasting the outcome of plans. The **CC2** agent could generate stories that starts with the present situation, implements the course of action under consideration and ends with failure. Analysis of the generated story could give decision makers insights into aspects and circumstances that should be simulated carefully. **CC3** agents could also use the stories for suggesting countermeasures. With a comprehensive domain-related supply of vignettes, story generation might even be used for situation assessment by **CC5** agents. It is interesting to note that the techniques of vignette-based story generation are similar to those of planning algorithms and simulation engines. The difference is in purpose rather than in methodology. Story

generators aim for novelty, planners for optimality and simulations for sufficiently realistic modeling of some relevant aspects of reality.

It can be difficult for decision makers to fully understand the ramifications of goals that have been adopted by opponents or partners. This may cause errors in situation recognition and in identifying relevant expectancies in the current situation assessment. Agent-based story generation where an open-ended story evolves driven by conflicting goals (Meehan 1976) could be useful for both **CC4** and **CC5** decision support agents. Such stories could give a fresh perspective from a different point of view and help identifying anomalies and possibly inspire reassessment of the situation.

Consider a **CC3** agent that, as discussed in the previous section, is tasked with coming up with new plan elements for the purpose of repairing a failed plan. Li et al. (2012) extends conceptual blending (Fauconnier and Turner, 2002) to incorporate goals with application to algorithms for generating hypothetical gadgets engineered to fulfill the goals. This methodology could be applied to algorithms for **CC3** agents in which the goals are derived from the needs of the jammed planning process and generated “gadgets” would be plan elements with prerequisites that can be fulfilled in the context of the problem-ridden plan and consequences designed to be instrumental for unjamming the planning process.

Jändel (2013) describes information fusion systems extended with computational creativity agents of type **CC5**. The agents aid in uncovering deceit by comparing generic deception strategies to the present situation and guide the fusion process to explore alternative situation assessments.

Future applications There are many research opportunities in the confluence of computational creativity and naturalistic decision-making both with respect to algorithms for the six types of agents indicated in this paper and for research into the effect and efficiency of computational creativity in various domains of decision-making.

Pioneering areas of application will probably be in high-stake strategic decision-making where time and resources are at hand and leaders are willing to go to great lengths in order to minimize risks and ensure decision quality. Bridgehead applications will therefore most likely be in fields such as defense strategy, major economic and environmental decisions and strategic business planning. As methods and tools evolve and the level of automation increases computational creativity will increasingly be applied also to operative and tactical decision-making.

Acknowledgments

This research is financed by the R&D programme of the Swedish Armed Forces.

References

Boden, M. 2009. Computer models of creativity. *AI Magazine* 30(3):23–34.

Bushinsky, S. 2009. Deus ex machina a higher creative species in the game of chess. *AI Magazine* 30(3):63–69.

Fauconnier, G., and Turner, M. 2002. *The Way We Think: Conceptual Blending and the Mind's Hidden Complexities*. Basic Books.

Ghallab, M.; Nau, D.S.; and Traverso, P. 2004. *Automated Planning: Theory and Practice*. Morgan Kaufmann.

Hill, L. G. 2010. *ShockWave Science and Technology Reference Library, V 5, Non-Shock Initiation of Explosives*. Springer.

Johnson, J., and Raab, M. 2003. Take the first: Option-generation and resulting choices. *Organizational Behavior and Human Decision Processes* 91:215229.

Jändel, M. 2013. Computational creativity for counterdeception in information fusion. Unpublished, submitted to 16th Int. Conf. on Information Fusion.

Kahneman, D.; Slovic, P.; and Tversky, A. 1982. *Judgement under uncertainty: Heuristics and biases*. Cambridge University Press.

Klein, G.; Calderwood, R.; and Clinton-Cirocco, A. 1986. Rapid decision-making on the fireground. In *Proceedings of the Human Factors and Ergonomics Society*, 576–580.

Klein, G. 2003. *Intuition at work*. New York: Doubleday.

Klein, G. 2008. Naturalistic decision making. *Human Factors* 50:456–460.

Li, B.; Zook, A.; Davis, N.; and Riedl, M. . 2012. Goal driven conceptual blending: A computational approach for creativity. In *Proceedings of the 2012 International Conference on Computational Creativity*.

McCaffrey, T., and Spector, L. 2011. How the obscure features hypothesis leads to innovation assistant software. In *Proceedings of the Second International Conference on Computational Creativity*.

Meehan, J. 1976. *The Metanovel: Writing stories by computer*. Ph.D. Dissertation, Yale.

Riedl, M. O., and Sugandh, N. 2008. Story planning with vignettes: Toward overcoming the content production bottleneck. In *Proceedings of the 1st Joint International Conference on Interactive Digital Storytelling: Interactive Storytelling, ICIDS '08*, 168–179. Berlin, Heidelberg: Springer.

Ross, K.; Klein, G.; Thunholm, P.; Scmitt, J.; and Baxter, H. 2004. The recognition-primed decision model. *Military Review* July-August, 6–10.

Surdu, J. R., and Kittka, K. 2008. The deep green concept. In *Proceedings of the 2008 Spring simulation multiconference, SpringSim '08*, 623–631. San Diego, CA, USA: Society for Computer Simulation International.

Tan, K.-M. T., and Kwok, K. 2009. Scenario generation using double-scope blending. In *AAAI Fall Symposium*.

Towler, M. 2010. *Rational decision making: An introduction*. Wiley.

Triantaphyllou, E. 2002. *Multi-Criteria Decision Making Methods: A Comparative Study*. Kluwer.

Yates, J.; Veinott, E.; and Patalano, A. 2003. Hard decisions, bad decisions: On decision quality and decision aiding. In Schneider, S., and Shanteau., eds., *Emerging perspectives on judgment and decision research*. Cambridge University Press. 13–63.