
Limitations of Human-Centric Decision Making: An Observer's Perspective

Bart Gajderowicz

BARTG@MIE.UTORONTO.CA

Mark S. Fox

MSF@MIE.UTORONTO.CA

Michael Grüninger

GRUNINGER@MIE.UTORONTO.CA

Department of Mechanical and Industrial Engineering, University of Toronto, 5 King's College Road, Toronto, M5S 3G8, CA

Abstract

Traditional decision theories based on economic theories of rationality make assumptions about the rationality of the human-like subjects being observed. Over the years, such assumptions have come into question with a number of extensions created to address such assumptions. These include sequential, subjective, and dynamic theories of behaviour, as well as factors impacting the subject, including bounded rationality and bias. However, the observer remains omniscient, not impacted by bounded rationality or biases. At the same time, the observer plays a major role in the development and configuration of models in artificial intelligence and cognitive architectures, which require details understanding of the subject's limitations and reasoning processes. This paper proposes that the observer is not omniscient but, like the subject, is also bounded. We evaluate decision theories to highlight which axioms make assumptions about the observer's omniscience and must be modified, removed, or can be retained. The goal is to identify a method towards a human-centric decision theory for bounded observers.

1. Introduction

This paper evaluates limitations of existing decision theories and artificial intelligence architectures for emulating decision-making of human-like subjects. We review key axioms of existing decision theories, mainly single decision theory, sequential decision theory, and dynamic choice theory. These theories are evaluated to understand their use by current artificial intelligence systems and cognitive architectures to interpret observed behaviour of a subject.

Modern theories recognize and categorize biases and aim to use them for modelling seemingly "irrational" decision-making of human-like subjects. The observer gathers information about subjects within a system, categorizing their biases, constraints, goal preferences, and cognitive impairments to create autonomous agents that emulate the behaviour of observed subjects. However, the observer is assumed to be omniscient, aware of all factors impacting the subject's decision-making. Four key components of human-like decision-making are missing in representing the observer. First, a sufficient representation of bounded rationality for the observer are not provided, who is assumed to be omniscient. Second, the observer is assumed to know the subject's goal preferences which are stable over time without contradictions. Third, emotional evaluation of the subject's own state

and probability of desired outcomes is dependent on a static assignment of emotion-based utility to specific events. Fourth, triggers for replanning during execution are assumed to always improve the plan over time.

Current work in artificial intelligence (AI) focuses on creating reasoning mechanisms that equal or exceed human performance. However, decision-making of human-like agents in response to a dynamic environment are insufficiently human-like. Many planners and cognitive architectures rely on optimal memory management solutions to identify the most efficient way to make and change decisions. With some recent exceptions, many methods have focused on making algorithms more efficient in finding optimal solutions and scale to a large number of human-like agents (Kosinski & Zaczek-Chrzanowska, 2003; Edelman, 2015; Stolzmann & Butz, 2000). Optimality is often achieved by removing human-like errors and biases while mimicking human-like decision-making with respect to social and structural factors. As with classical decision theories, the underlying reasoning process and belief system is assumed to be sufficiently understood and implemented using AI methods. This includes the subject's preferences and goal reasoning methods. Also, emotions are not sufficiently dynamic, but rather rely on emotion models that are statically associated with specific events *a priori*. This does not capture the dynamic nature of emotions and their influence on our decision-making (Seo & Feldman Barrett, 2007).

2. Theories of Decision Making

In this section, we highlight the main theories of decision-making. We discuss their key axioms from the perspective of the observer. We also discuss how emotions and biases are represented.

2.1 Single Decision Theory

The ability to emulate behaviour is generally assumed to be done by an observer with an outside perspective. The subject, then, is the decision maker being observed for the purpose of emulation. According to classical decision theory, the observer is assumed to be omniscient. They know the state of the subject, their preferences, their constraints, and the choices available to them. The subject is also assumed to be rational, meaning they maximize the utility of the means available to them towards satisfying their goals according to their preferences. von Neumann and Morgenstern (1944) proposed calculations that make it possible to derive the utility of a decision, leading to the development of expected utility (EU) and classical decision theory. This work is based on four key axioms, prefixed with VNM.

Axiom VNM-1 Completeness: A preference is assigned to all pairs of choices.

Axiom VNM-2 Transitivity: Order of preferred choices is maintained across all choices.

Axiom VNM-3 Continuity: No outcome is so bad that it is not worth a gamble with a sufficiently high probability of success.

Axiom VNM-4 Independence: Utility of a plan is independent from the probability of the plan's outcome.

The axioms make assumptions about how stable and invariant the subject's preferences and objectivity are when calculating utility of actions. Axiom VNM-1 states that utility is assigned to all

possible choices a subject can make at any given time. Axiom VNM-2 states that the order of preferred goals is maintained across all combinations of choices, where one order does not contradict another. Axiom VNM-3 states that any risk is worth taking if the benefit is sufficiently high. Axiom VNM-4 states that the utility of a plan is not dependent on the likelihood of the outcome being successful.

2.2 Role of Bounded Rationality

Decision theory assumes the observer and subject are omniscient in their knowledge, and rational in how that knowledge is used to make decisions. Bounded rationality provides a framework for incorporating cognitive and environmental constraints that may help interpret a subject's seemingly "irrational" behaviour (Simon, 1955). According to bounded rationality, there are three main types of bounds influencing a subject's decisions (Simon, 1955, 1967, 1972). First, the information required to find an optimal solution is either incomplete or incorrect. Second, cognitive limitations prevent a subject from considering all possible factors, limiting them to only the most trivial problems. Third, the subject simply does not have enough computational time to evaluate all options required to solve a complex problem. Within bounded rationality, subjects compensate for lack of information, cognitive limitations or time by adopting a complex network of normative factors. For example, social norms and individual biases allow one to determine their action based on the actions of others (Etzioni, 1993; Kahneman, 2003; Zafirovski, 2005). Many factors that lead to seemingly irrational behaviour, have been identified, including emotional beliefs and desires, biases, irrational reasoning processes, as well as normative social factors (Etzioni, 1988; Zafirovski, 2005; Ladouceur et al., 1988).

2.3 Subjective Decision Theory

To understand a subject's choices, the observer must have a sufficient understanding of the subject's goals used to make choices that satisfy them. Once identified, goal preferences can be represented with a utility weight, calculated using an extension of decision theory called subjective decision theory (Savage, 1954; Jeffrey, 1990). The order of preferred goals can be expressed as nominal and cardinal preferences. Nominal preferences provide the order of preferred goals (Peterson, 2009). Cardinal preferences indicate the degree to which one goal is preferred over another (Von Neumann et al., 1944; Wold et al., 1952). Subjective decision theory provides several methods for calculating the utility of goal states and choices.

Jeffrey's (1990) theory of subjective expectation utility makes two key contributions. First, all probabilities are strictly based on Bayesianism. Second, utility and probability are associated with predicates that can be true or false in a particular world. This allows a subject to consider multiple worlds when calculating the utility of actions, an important aspect of human behaviour (Cooper, 1999). Preferences are dependent on the outcome that may be different in different worlds, meaning one can't assume the same choice, with the same probability, will have the same outcome in two different worlds. Consider states s_i and s_j , a sequence of actions A^x and A^y , and a utility function $U(A)$ for sequence utility. Jeffrey's axioms, prefixed with J, state that:

Axiom J-1. Averaging: A proposition cannot be better or worse than any of its realizations (worlds). It follows then that if, for some set of preference ordering \prec , s_i and s_j are mutually incompatible then combining them has no impact on the overall preferences, where:

$$s_i \prec s_j \iff s_i \prec (s_i \cup s_j) \prec s_j$$

Axiom J-2. Impartiality: Given two sequences A^x and A^y , if $U(A^x) = U(A^y)$ and they have no actions in common, then a new state s_i changes $U(A^x)$ and $U(A^y)$ in the same way.

Savage’s (1954) theory introduced the idea of thinking in terms of a process with individual states and an outcome rather than simply as individual choices. Outcomes are the situations and events a subject cares about and has control over, i.e. their goals. States are scenarios that a subject has no control over. States are also the source of uncertainty in the decision-making process due to the subject’s limitation of seeing the entire state of their world. A subject’s preferences for outcomes are based on beliefs about the probability of success. Savages theory requires six axioms to be true, prefixed with S:

Axiom S-1. The relation \prec between two goal states is complete and transitive, combining the completeness and transitivity axioms VNM-1 and VNM-2.

Axiom S-2. Sure Thing Principle: Choice preferences can be evaluated independently if the probability of one choice is not impacted by the outcome of another. This is an extension to independence axiom VNM-4.

Axiom S-3. State Neutrality: The preference of an outcome is independent of that current state, and has no impact on $U(A^x)$, an extension to independence axiom VNM-4.

Axiom S-4. Preference of an outcome is independent of the outcome’s utility, an extension to the independence axiom VNM-4.

Axiom S-5. A subject must not be indifferent to sequences, and there must be some difference in utility between one sequence and another.

Axiom S-6. Non-Atomicity: If a sequence A^x is already preferred to sequence A^y , where $A^x \prec A^y$, then a change to some state s_k made true by A^x and A^y will not affect the preference unless the probability of s_k is sufficiently high.

2.4 Dynamic Choice Theory

Given that the subject is bounded, the plan they select will inevitably have mistakes that become apparent only during the execution phase. Hence, the subject’s plan will inevitably be un-executable in its original form. During this phase, the subject must monitor plan execution, and respond to feedback when expected outcomes do not match actual outcomes. Given the feedback, the subject can reevaluate their preferences to focus on what is required or most probable, rather than preferred. Following such adjustments, the subject can create a new plan that uses a more realistic goal ordering. When a plan is no longer executable, the subject must stop and reevaluate their goals, current state, and construct a new plan. This replanning process occurs every time a new plan is worth pursuing over the current plan.

Dynamic choice theory assumes the observer knows what reasonable sequences a subject might follow based on a strategy for calculating plan utility (Bermúdez, 2009). Hammond (1976) has

identified three such strategies, mainly *myopic*, *sophisticated* and the *resolute*, with a comparative analysis by McClennen (1990). The *myopic* strategy assumes choices are independent of any future or past choices. The *sophisticated* strategy considers the utility of all sequences starting from the current time, and chooses the sequence with the lowest risk, recalculated after each action is executed. Finally, the *resolute* strategy also calculates the utility of all sequences but only at the beginning of the decision process. Once a sequence is selected, utility is not recalculated and the originally chosen sequence is followed until the end.

2.5 Roles of Emotions and Bias

While subjective decision theory allows an observer to calculate utility of goals and plans from observed behaviour, some subjectivity is based on bias and emotions. For example, many models in economics have attempted to incorporate bounded rationality with varying degrees of success (Rubinstein, 1998). *Behavioural economics* is a field within economics that attempts to address some of the shortcomings of the rational agent theory by categorizing different types of biases exhibited by individuals, incorporating sociology, psychology, biology along with certain notions in economics (Wilkinson & Klaes, 2012; Bolton & Ockenfels, 2012). *Game theory* has been used by economists, psychologists and political scientists to simulate behaviour for non-economic problems and incorporate a variety of observable biases (Başar & Olsder, 1995).

Emotions play a central role in perceiving events around us, and provide an opportunity to model the subjectivity of individual choices (Izard, 2007; Rodrigues et al., 2014; Ivanović et al., 2015). This includes perception of goals and consequences of our actions towards satisfying those goals. There are many theories of emotions and architectures that contribute to human-like behaviour. The vast majority, like BDI (belief-desire-intention) and OCC (Ortony, Clore and Collins) rely on “drives” that form a direct connection between some stimuli and a response (Reisenzein et al., 2013; Gratch & Marsella, 2004; Zelazo & Cunningham, 2009). The cognitive theory of emotions called “appraisal theory” captures how events are appraised before a response is triggered (Scherer et al., 2001). The use of appraisal theory is often supplemented with the OCC model that provides responses and valence defined *a priori* with discrete emotions like fear, joy, and sadness (Ortony et al., 2014). As a rational subject that is also emotional, they are assumed to improve their choices over time.

3. Limitations for Human-Centric AI Reasoning

The axioms presented in the previous section are preserved only if we assume the observer modeling the subject’s behaviour within an AI architecture is also omniscient. However, assuming that like the subject the observer is also bounded, this section evaluates the axioms to determine which are preserved and which must be abandoned or modified.

3.1 Bounded Observer and Subject

Due to bounded rationality, the assumption that either the observer or the subject are omniscient are not always applicable. Outside of a controlled lab setting, it is near impossible for the observer to

control for all factors that contribute to a subject’s decision-making. An observer’s evaluation could produce a reasonable plan a subject could follow. The same plan, however, may not be executable in real life. Due to gaps in the observer’s knowledge, the subject would be triggered to replan often throughout the execution phase.

Axiom VNM-1 (completeness) is a problematic requirement in decision theory (Hansson, 2005). First, the knowledge bound and limited memory prevents a subject from having knowledge about all actions and possible outcomes, hence all states cannot be assigned a preference. Second, the cognitive and time bounds prevent a planning algorithm from generating a search tree and visiting all states to assign them a preference. Also, if an observer knows that a subject is either risk taking or risk averse from past observations, they could reasonably infer whether the subject thinks a gamble is worth taking for some gain, and Axiom VNM-3 (continuity) is preserved.

3.2 Goal Preferences

It is difficult to define a realistic goal ordering and preferences for a subject using traditional representations. The required axioms in decision theory and expected utility for governing relationships between preferences used in utility-maximization do not always hold for human desires and preferences (Scott, 2000). When the relationships do hold, preferences become contextualized in a situation and the type of decision being made (Arrow, 1963). To simplify the modeling effort, most important neoclassic economists treat preferences as stable or given (Ladouceur et al., 1988).

Axiom VNM-2 (transitivity) is a problematic requirement and generally assumed to be a weak relation between preferred states rather than a strong relation (Hansson, 2005). For example, say a subject’s goal preferences are $b \prec a$, $c \prec b$, and $a \prec c$. Any sequence with all three goal states must break one of the preferences. For example, either:

$$b \prec a \wedge c \prec b \implies a \not\prec c \quad \text{or} \quad b \prec a \wedge a \prec c \implies c \not\prec b$$

As Arrow (1963) points out, a strong preference relation becomes contextualized in a particular scenario. Hence, for two plans that have different scenarios goal transitivity may not hold, and Axiom VNM-2 is not preserved. Coincidentally, Savage’s Axiom S-1 is not preserved as it relies on transitivity of goal preferences. Axiom S-1 also relies on completeness Axiom VNM-1 which is not preserved, as per section 3.1.

Changes in perception of preferences from one time step to another during execution prevent some axioms from being preserved. Axiom VNM-4 (independence) is not preserved since during execution the newly discovered outcomes of actions reduce the utility of a plan. Hence, plan utility is dependent, in part at least, on the perceived outcome of each action. Axiom S-2 requires that two events are evaluated independently. However, since the independence axiom VNM-4 is not preserved, S-2 is also not preserved. Axiom S-3 (state neutrality) is not preserved since actions with highest probability of success from the current state are preferred.

Jeffrey’s axioms J-1 and J-2 for averaging and impartiality are not preserved due to the order imposed by action preconditions. First, Jeffrey’s *averaging* axiom J-1 is not preserved due to the change in a goal order has on its utility in relation to other goals. Since transitivity is not preserved, there is no guarantee that adding a new goal will not change the scenario in a way that changes goal preference and the utility of the entire plan. Consider the union of two plans, $P^x \cup P^y$. Certain ac-

tions in plan P^y may undo actions in plan P^x , requiring action reordering and potentially changing the plan's utility. Next consider Jeffrey's *impartiality* axiom J-2. If two plans P^x and P^y finish with utility, say $U(P^x) = U(P^y)$, it is not guaranteed that a new action or goal state will change their utility in the same way. If a new action a is added to each plan, a 's preconditions may be true in plan P^x but not P^y . Changing the order of actions to accommodate the preconditions may change utility of the plan.

4. Consequences for Artificial Intelligence

By assuming the observer, as the builder of AI models, is bounded we must accept that the axioms discussed in the previous chapter are not preserved as they assume an omniscient observer. This has an impact on how behaviour of the subject can be emulated. We now evaluate how the preserved axioms impact subject emulation by a bounded observer.

4.1 What is Preserved?

There are two characteristics shared by the three preserved axioms VNM-3, S-5 and S-6. First, the utility of a sequence is used to decide between sequences, not the utility of single action. Second, once a subject has committed to one sequence over another, a sufficiently large change in utility is required for a subject to select another sequence. In the context of AI planning and cognitive architectures, this means that the subject will select a plan with maximum utility, then execute and monitor that plan. During execution, the subject will trigger a replanning process when either: 1) a plan with a sufficiently higher utility is found or 2) the current plan's utility falls sufficiently low relative to its original utility during the planning phase.

Recall that Axiom VNM-3 (continuity) states that "no outcome is so bad that it is not worth a gamble with a sufficiently high probability of success." Axiom VNM-3 is preserved because it is possible that a sufficiently large change in plan utility after replanning will prompt a subject to change their course of action and follow the new plan. Axiom S-5 (sequence indifference) states that "a subject must not be indifferent to sequences, and there must be some difference in utility between one sequence and another." Axiom S-5 is preserved because the subject is not indifferent to plan utility. If two plans have different utilities, the rational subject will prefer the plan with maximum utility. If two plans have the same utility, only then is the subject indifferent to the plan. Finally, Axiom S-6 (sequence non-atomicity) is preserved because a change to an already chosen sequence will convince a subject to choose the new plan only if that change significantly changes each sequence's utility.

$$P_{i,j}^x \cup P_{j,k}^y \cup \dots \cup P_{v,w}^z = P^O \quad (1)$$

From the perspective of the observer, rather than emulating one continuous plan, they need to emulate a set of partial plans, as per Equation 1. Given rational plans P^x , P^y to P^z , and partial plans $P_{i,j}^x, P_{j,k}^y$ to $P_{v,w}^z$ chosen and executed by the subject from time steps i to w , the observed plan P^O is the union of partial plans executed by the subject. For a given plan P^x at time step j the subject triggered the replanning process. The new plan P^y starting at time step j needed to be

replanned at time step k . The final replanning was triggered at time step v with the partial plan $P_{v,w}^z$ being executed. Hence a rational subject is not myopic, sophisticated or resolute. Instead, a subject is resolute up to a point when replanning is required and utility recalculated. Since the partial plans are combined continuously, the utility function required to model such a process must be continuous.

4.2 Continuous Goal Preference Utility

The subject is assumed to have a set of goals provided *a priori* as so-called “achievement goals” (Grant & Dweck, 2003). Known goal preferences are often provided by individuals for specific scenarios (Ladouceur et al., 1988; Baier & McIlraith, 2008; Liaskos et al., 2010). To allow for a continuous evaluation of goal preferences, goal ranking must be provided for a sequence of actions across multiple scenarios. This differs from the ranking between independent pairs of goals, as discussed in section 3.2. This sequence would include not just the achievement goals but also the “maintenance goals” (Hindriks & Van Riemsdijk, 2008). Any observed behaviour that does not satisfy known achievement goals is assumed to be satisfying maintenance goals required to satisfy achievement goals.

A number of continuous planning systems have been developed (Myers, 1999). Hierarchical planners have provided a representation of planning at different levels of abstraction, allowing for continuous revisions of plans. The FPE architecture (Bai et al., 2015) monitors a number of meta-properties including effects, perceptual inspection, conditions, as well as intention enaction and selection. Such a representation leads well to a continuous representation of goals and plans that are continuously generated and ranked. Langley et al. (2017) extended the PUG architecture with PUG/X which has the ability to execute and monitor a selected plan based on goal utilities in continuous domains, and replan when needed.

Unknown human-centric achievement and maintenance goals can be grounded in theories of behaviour from psychology like Maslow’s (1943) hierarchy. For example, being social is a goal everyone is assumed to have at some point in their life. Eating food and feeling secure can be considered maintenance goals that help the subject satisfy their achievement goals. Once goals are identified, their preference can be represented with a utility weight, calculated using a subjective decision theory (Gajderowicz et al., 2018). Gajderowicz et al. (2018) express goal utility as nominal and cardinal preferences.

4.3 Continuous Bias and Emotion-Based Utility

Bias and emotion-based utility evaluation provide an opportunity to act as a trigger for replanning. Neoclassical rational agent models have been extended with biased decisions by behavioural economics. Theories based on behavioural economics have retained many of the basic principles of rational choice theory, and mostly extend them with cognitive limitations related to specific scenarios and biases found in economic literature (Kahneman, 2003). However, some work has extended behavioural economic into non-economic domains, such as addiction. Bretteville-Jensen (2003) points out that rational subjects take into account past, present, and future consumption. During

consumption, utility of each dose is calculated as the “consumption capital” over continuous time. For each consumer, an addiction threshold changes their reason for use from casual to full addiction.

Game theory also allows for modelling of biased decisions over the extent of a game (Başar & Olsder, 1995). Traditionally, game theory assumes subjects are not bounded. Specifically, it assumes the rules of the game are well defined, all subjects are aware of these rules, know their utility function and act to maximize it. In *cooperative game theory*, it is also assumed that every subject trusts that other agents are rational and maximize their own utility towards their respective goals (Peleg & Sudhölter, 2007). Recent work has focused on continuous games, where subjects are aware of past experiences and change their choices asynchronously in continuous time (Leng et al., 2018).

Rather than always optimizing as assumed by classical expected utility, a human-like subject will instead reevaluate a plan when their perception of the plan changes. A subject may believe they will act optimally during the planning phase, when in fact their mood may change during the execution phase. Changes in emotional perception of events during execution are a common response to true but unsuspected consequences of action, but have not been included in existing models of human-like decision-making (Samsonovich, 2012; Dias & Paiva, 2011; Lin et al., 2012). Over time, perception of an event may switch from positive valence to negative, and vice versa, as defined by the Emotional Cycle of Change (ECOC) (Kelley & Connor, 1979; Dyson & Brown, 2006; Lancaster & Gray, 1982; Mashazi, 2002). Relying on optimistic and pessimistic stages of ECOC to calculate a subject’s emotional mood can be used to trigger replanning when mood falls below some predetermined threshold. Gajderowicz et al. (2018) incorporate ECOC as a continuous mood utility function adopted for the homeless population.

5. Conclusion

This paper identifies the limitation of existing decision theories for emulating and interpreting behaviour of human-like agents. Assumptions about the omniscience of the observer and rationality of the subject are challenged and evaluated as they pertain to AI planning and cognitive architectures. This analysis highlights which axioms in existing theories are not preserved when modelling behaviour of human-like subjects. Recognizing the observer is omniscient, extension to existing theories should focus on work in psychology and sociology to evaluate factors impacting an agent from observations. Such work would benefit the evaluation of subjects during the planning, monitoring and execution phases of a subject’s behaviour. It provides a more complete picture of factors impacting behaviour that are not observable directly, but only as patterns over time. Such continuous functions used to calculate plan utility can be applied to replanning, goal reasoning, and emulation of behaviour.

References

- Arrow, K. J. (1963). *Social choice and individual values*. London: Yale University Press.
- Başar, T., & Olsder, G. J. (1995). *Dynamic noncooperative game theory*. Cambridge, MA: Academic Press.

- Bai, Y., Pearce, C., Langley, P. W., Barley, M., & Worsfold, C. (2015). An architecture for flexibly interleaving planning and execution. *Proceedings of the Third Annual Conference on Advances in Cognitive Systems*. Atlanta, GA: Cognitive Systems Foundation.
- Baier, J. A., & McIlraith, S. A. (2008). Planning with preferences. *AI Magazine*, 29, 25–36.
- Bermúdez, J. L. (2009). *Decision theory and rationality*. Oxford University Press.
- Bolton, G. E., & Ockenfels, A. (2012). Behavioral economic engineering. *Journal of Economic Psychology*, 33, 665–676.
- Bretteville-Jensen, A. L. (2003). Rational addiction and injection of heroin. In R. E. Vuchinich & N. Heather (Eds.), *Choice, behavioural economics and addiction*, 281–308. Elsevier.
- Cooper, C. R. (1999). Multiple selves, multiple worlds: Cultural perspectives on individuality and connectedness in adolescent development. *Proceedings of the 1999 Minnesota Symposia on Child Psychology* (p. 25). New York and London: Psychology Press.
- Dias, J., & Paiva, A. (2011). Agents with emotional intelligence for storytelling. *Proceedings of the Fourth International Conference on Affective Computing and Intelligent Interaction* (pp. 77–86). Memphis, TN: Springer.
- Dyson, S., & Brown, B. (2006). The implications of strategy and method. In A. Bryman (Ed.), *Social theory and applied health research*, 130–147. New York: McGraw-Hill.
- Edelman, S. (2015). The minority report: Some common assumptions to reconsider in the modelling of the brain and behaviour. *Journal of Experimental and Theoretical Artificial Intelligence*, 3079, 1–26.
- Etzioni, A. (1988). Normative-affective factors: Toward a new decision-making model. *Journal of Economic Psychology*, 9, 125–150.
- Etzioni, A. (1993). Normative-affective choices. *Human Relations*, 46, 1053–1069.
- Gajderowicz, B., Fox, M. S., & Grüninger, M. (in press). The role of goal ranking and emotional utility in dynamic replanning strategies. *Advances in Cognitive Systems*. Cognitive Systems Foundation.
- Grant, H., & Dweck, C. S. (2003). Clarifying achievement goals and their impact. *Journal of personality and social psychology*, 85, 541–553.
- Gratch, J., & Marsella, S. C. (2004). A domain-independent framework for modeling emotion. *Cognitive Systems Research*, 5, 269–306.
- Hammond, P. J. (1976). Changing tastes and choice coherent dynamic. *The Review of Economic Studies*, 1, 159–173.
- Hansson, S. O. (2005). *Decision theory: A brief introduction* (Technical Report). Department of Philosophy and the History of Technology, Royal Institute of Technology, Stockholm.
- Hindriks, K. V., & Van Riemsdijk, M. B. (2008). Satisfying maintenance goals. *Lecture Notes in Computer Science*, 4897, 86–103.
- Ivanović, M., Budimac, Z., Radovanović, M., Kurbalija, V., Dai, W., Badica, C., Colhon, M., Ninković, S., & Mitrović, D. (2015). Emotional agents - state of the art and applications. *Com-*

- puter Science and Information Systems*, 12, 1121–1148.
- Izard, C. E. (2007). Basic emotions, natural kinds, emotion schemas, and a new paradigm. *Perspectives on Psychological Science*, 2, 260–280.
- Jeffrey, R. C. (1990). *The logic of decision*. University of Chicago Press.
- Kahneman, D. (2003). Maps of bounded rationality: Psychology for behavioral economics. *The American Economic Review*, 93, 1449–1475.
- Kelley, D., & Connor, D. R. (1979). The emotional cycle of change. *The 1979 Annual Handbook for Group Facilitators*.
- Kosinski, W., & Zaczek-Chrzanowska, D. (2003). Pavlovian, Skinner, and other behaviourists' contributions to AI. *Proceedings of the 2003 Workshop on Intelligent Motion and Interaction Within Virtual Environments* (pp. 135–148). London, UK: NASA.
- Ladouceur, R., Gaboury, A., Dumont, M., & Rochette, P. (1988). Gambling: Relationship between the frequency of wins and irrational thinking. *Journal of Psychology*, 122, 409–414.
- Lancaster, J., & Gray, C. J. (1982). Change agents as leaders in nursing.
- Langley, P. W., Choi, D., Barley, M., Meadows, B., & Katz, E. (2017). Generating, executing, and monitoring plans with goal-based utilities in continuous domains. *Proceedings of the Fifth Annual Conference on Advances in Cognitive Systems* (pp. 1–12). Troy, NY: Cognitive Systems Foundation.
- Leng, A., Friesen, L., Kalayci, K., & Man, P. (2018). A minimum effort coordination game experiment in continuous time. *Experimental Economics*, 21, 549–572.
- Liaskos, S., McIlraith, S. A., Sohrabi, S., & Mylopoulos, J. (2010). Integrating preferences into goal models for requirements engineering. *Proceedings of the Eighteenth IEEE International Requirements Engineering Conference* (pp. 135–144). Sydney: IEEE Computer Society.
- Lin, J., Spraragen, M., & Zyda, M. (2012). Computational models of emotion and cognition. *Advances in Cognitive Systems*, 2, 59–76.
- Mashazi, M. I. (2002). *A model for the integration of provincial and local authority nurses rendering primary health care services in a district*. Doctoral dissertation, University of South Afrika, Pretoria, South Africa.
- Maslow, A. H. (1943). A theory of human motivation. *Psychological Review*, 50, 370–396.
- McClennen, E. F. (1990). *Rationality and dynamic choice: Foundational explorations*. Cambridge, MA: Cambridge University Press.
- Myers, K. L. (1999). CPEF: A continuous planning and execution framework. *AI Magazine*, 20, 63.
- Ortony, A., Clore, G. L., & Collins, A. (1989). *The cognitive structure of emotions*. New York: Cambridge University Press.
- Peleg, B., & Sudhölter, P. (2007). *Introduction to the theory of cooperative games*. Berlin: Springer-Verlag.
- Peterson, M. (2009). *An introduction to decision theory*. New York: Cambridge University Press.

- Reisenzein, R., Hudlicka, E., Dastani, M., Gratch, J., Hindriks, K. V., Lorini, E., & Meyer, J.-J. C. (2013). Computational modeling of emotion: Toward improving the inter- and intradisciplinary exchange. *IEEE Transactions on Affective Computing*, 4, 246–266.
- Rodrigues, S. H., Mascarenhas, S., Dias, J., & Paiva, A. (2014). A process model of empathy for virtual agents. *Interacting with Computers*, 27, 371–391.
- Rubinstein, A. (1998). *Modeling bounded rationality*. Cambridge, MA: MIT Press.
- Samsonovich, A. V. (2012). *Modeling social emotions in intelligent agents based on the mental state formalism* (Technical report FS-12-02). AAI.
- Savage, L. J. (1954). *The foundations of statistics*. Wiley.
- Scherer, K. R., Schorr, A., & Johnstone, T. (2001). *Appraisal processes in emotion: Theory, methods, research*. Oxford University Press.
- Scott, J. (2000). Rational choice theory. *Understanding contemporary society: Theories of the present*, 129, 671–85.
- Seo, M., & Feldman Barrett, L. (2007). Being emotional during decision making - Good or bad? An empirical investigation. *Academy of Management Journal*, 50, 923–940.
- Simon, H. A. (1955). A behavioral model of rational choice. *The quarterly journal of economics*, 69, 99–118.
- Simon, H. A. (1967). Motivational and emotional controls of cognition. *Psychological Review*, 74, 29–39.
- Simon, H. A. (1972). Theories of bounded rationality. In C. B. McGuire & R. Radner (Eds.), *Decision and organization*, chapter 8. Amsterdam: North-Holland Publishing.
- Stolzmann, W., & Butz, M. (1999). Latent learning and action-planning in robots with anticipatory classifier systems. *Proceedings of the Second International Workshop on Learning Classifier Systems* (pp. 301–317). Orlando, FL: Springer.
- Von Neumann, J., Morgenstern, O., & Others (1944). *Theory of games and economic behavior*.
- Wilkinson, N., & Klaes, M. (2012). *An introduction to behavioral economics*. Palgrave.
- Wold, H., Shackle, L. S., & Savage, L. J. (1952). Ordinal preferences or cardinal utility? *Econometrica*, 20, 661–664.
- Zafirovski, M. (2005). Is sociology the science of the irrational? Conceptions of rationality in sociological theory. *The American Sociologist*, 36, 85–110.
- Zelazo, P. D., & Cunningham, W. A. (2009). Executive function: Mechanism underlying emotion regulation. In J. Gross (Ed.), *Handbook of emotion regulation*, chapter 7. Guilford Publications.