

General Model of Human Motivation and Goal Ranking

Bart Gajderowicz, Mark S. Fox, Michael Grüninger

Department of Mechanical and Industrial Engineering

University of Toronto

5 King's College Road

Toronto, ON M5S 3G8, Canada

Abstract

In this article, we describe high-fidelity human behaviour emulation model capable of ranking and re-ranking goals during plan execution based on changing emotional modes of an agent. Our model assumes the agent is rational but its reasoning is bounded. The agent's reasoning process incorporates emotions and basic human needs to emulate changes in human behaviour under cognitive limitations. The majority of cognitive systems that incorporate emotions rely on reactive models that elicit predetermined responses to emotional modes. Our model demonstrates how human emotions change during the execution of a plan independent of specific events that may elicit such responses. The initial goals of the agent are grounded in basic human needs outlined by Maslow's Hierarchy. Once a plan is generated under the cognitive limitations of the agent and execution begins, goals are re-ranked based on an emotional re-evaluation of the plan's progress. The result is a high-fidelity, domain-independent, general theory of motivation based on human needs and emotions. We demonstrate the algorithm with a use-case from the social service domain by emulating the behaviour of homeless clients in response to an intervention program.

Introduction

In this article, we show how emulation of human behaviour can be improved by extending artificial intelligence (AI) planning methods with behaviour psychology theories in novel ways. We describe a high-fidelity human behaviour emulation model capable of ranking and re-ranking goals based on changing emotional modes of an agent.

Unlike existing systems with specific goals defined *a priori*, our agent's goals are grounded by abstract representations of basic human needs defined by Maslow's Hierarchy (MH) (Maslow 1943). The hierarchy categorizes human needs into five levels: physiological (e.g. eat, sleep), security (e.g. shelter, safety), social (e.g. family, community), esteem (e.g. self-esteem, pride), and self-actualization (e.g. be a good parent, rewarding job). According to MH goals at lower levels must be satisfied before moving onto goals at higher levels. Empirical evidence suggests that for

some subpopulations the levels do not have a strict ordering (Henwood et al. 2015). Rather, goals are reevaluated and re-ranked in a manner more dynamic than assumed by the hierarchy. In recent work, a simplified version of MH has been incorporated into a cognitive architecture (CA) as a motivation model for simulating *artificial* agents (Shi et al. 2013). Our work relies on the complete version of MH used by social scientists for modeling *human* motivation and behaviour.

Our work also extends goal ranking and the planning process by incorporating human emotions. Most existing AI systems and CAs rely on appraisal theory and arousal theory to categorize emotions and emotional responses to events (Reisenzein et al. 2013; Lin, Spraragen, and Zyda 2012). In such systems, emotional states like anger, fear and happiness are associated with agent responses to events with explicit rules or probabilities. Our work takes a more general approach by generalizing the effects of emotions by simulating changes in emotional states as agents execute actions in any plan towards their goals. The Emotional Cycle of Change (ECOC) provides the psychological foundation for changing behaviour due to emotions (Gajderowicz, Fox, and Grüninger 2017b).

Finally, our model assumes the agent is rational but its reasoning is bounded. A rational agent maximizes their utility towards a goal, as defined by social science and AI literature (Gratch and Marsella 2004; Russell 1997; Zafirovski 2005). By bounds, we mean limitations defined by the theory of bounded rationality (BR), mainly cognitive limitations, missing information, and limited computation time (Simon 1996). However, AI and social science disciplines take different views on how best to understand human behaviour as rational or irrational within such limits. Our work merges these views by extending the rational planning reasoner STRIPS, with a plan utility function that incorporates MH and ECOC theories. The result is a high-fidelity, domain-independent general theory of motivation based on human needs and emotions.

To demonstrate the validity of our work, we summarize an experiment that compares the goal ranking of our model to the changing needs of homeless clients, a traditionally difficult population to plan client-centred services for (Henwood et al. 2015). We use data from a real social service intervention program called Housing First, administered by Calgary Homeless Foundation (CHF) (<http://calgaryhomeless.com>).

Background

Belief-desire-intention (BDI) is one of the most popular models used to represent goal-driven agents. Within BDI, desires represent the goals of an agent. Traditionally, goal creation, ranking, and re-ranking is either provided *a priori* or is calculated based on an agent's beliefs and intentions during the decision-making process. In this section, we describe how goals are ranked in existing AI planning and CA systems, and the limitations in applying such systems to emulation of human behaviour.

Basic Goal Ranking: Classical planning systems are based on the STRIPS architecture that solve planning problems (Fikes, Hart, and Nilsson 1972). A STRIPS planning problem has three main components: initial state of the world, a set of goals that should be true, and the action schema (*AS*) that provides actions that transition the world from one state to another. Basic goal ranking has been based on the order goals were initially specified in. Early on researchers recognized that it was not always feasible to search all possible plans to find the most optimal one (Simon 1996). Like humans, computers are also bounded by missing information and computation time. To find optimal solutions more efficiently, heuristic search was introduced. Heuristic search relies on various methods for estimating the optimal path and searching that path first. Most heuristic methods are based on probabilities and estimated distances. For example, the A* heuristic search calculates plan utility as shortest distance from start state to goal state (Hart, Nils, and Raphael 1968). Depending on the application, domain-specific problem characteristics could be used to prioritize goals and rank plans. In scheduling, heuristics based on feasibility of completing a plan within certain time windows and availability of resources saw great improvements in the efficiency of finding a plan (Sadeh and Fox 1995). This research demonstrated sizable improvements variable and value orderings have in the efficiency of a search problem.

Goal Categorization: As problem definitions become more abstract and search algorithms more sophisticated, it is possible to categorize goals to improve search efficiency further by excluding goal rankings known to be non-optimal (Hendler, Tate, and Drummond 1990). In AI and social sciences, the most fundamental goal categories are *achievement goals* and *maintenance goals* (Norman and Long 1995). Consider a typical application of AI agents; the planetary exploration rover. A teleological rover will have various predefined "*achievement goals*" provided *a priori*, such as performing experiments and taking photographs. These generally embody why an AI agent was created in the first place. To support achievement goals, "*maintenance goals*" must be continuously monitored, including power management and adaption to its terrain. Any required maintenance goals must first be satisfied before achievement goals are pursued.

AI planners and CAs have successfully incorporated goal categories to reduce search time. For example, Fikes et al. provide an extension to STRIPS that creates MACROPs, partial plans for specific goals (Fikes, Hart, and Nilsson 1972). Extensions to STRIPS like ABSTRIPS and NOAH provide a hierarchical search spaces (Fox 1983). Here, special "critical goals" represent highly ranked achievement

goals that must be included in the final plan. NOAH allows for the creation of goal-pair ordering used to resolve conflicts during the search. CAs can organize goals and partial plans in ways that mimic human memory management. For example, Soar and ACT-R store already created partial plans called "chunks" in memory for later use (Lin, Spraragen, and Zyda 2012). The original ICARUS has a special group of static "top-level goals" provided *a priori*, and can re-rank sub-goals during execution, while an extension allows the re-ranking of top-level goals as well (Choi 2010).

Goal Weights: Like people, machines must also reason with bounded rationality. When relying on incomplete information or insufficient resources, an assignment of weights to goals is often required (Hendler, Tate, and Drummond 1990). Many classical planners have incorporated probabilities, and detect changes in those probabilities during plan execution, triggering a re-planning process when needed (Blythe 1998; Little and Thiebaut 2007). Probabilistic planners like MAXPLAN extend work in operations research specifically to focus on finding optimal solutions (Majercik and Littman 1998). Systems like ACT-R assign utilities to production rules to indicate their usefulness (Anderson et al. 2013). Capturing weights as "preferences" is especially helpful when ranking goals provided by humans. Some planning algorithm incorporate constraints defined by human subjects to resolve conflicting preferences (Brafman and Chernyavsky 2005). In project requirements engineering, preferences are used to select most desired and realistic requirements and create a work schedule that prioritizes client needs (Liaskos et al. 2010).

Human Motivations and Behaviour: Once user preferences are captured, we can rely on algorithms to identify conflicting or unreasonable beliefs. Humans have many such beliefs, often leading to observed behaviour being categorized as biased, emotional, or irrational (Zafirovski 2005). Biased and irrational behaviour has been studied extensively in social sciences (Kahneman 2003). This article focuses on the emotional factors that influence goal ranking in AI planning and CAs. Majority of AI systems and CAs rely on either arousal theory or appraisal theory to incorporate emotions (Gratch and Marsella 2004; Lin, Spraragen, and Zyda 2012). Arousal theory is based on "drives" that motivate us to perform certain activities associated with positive or negative valance (Gratch and Marsella 2004). In appraisal theory, emotions act as categories for interpreting our perception of the world (Lin, Spraragen, and Zyda 2012). Emotions control how internal relationships form between an agent's beliefs, desires, and intentions.

Some systems associate emotional appraisal with events using explicit rules (Lin, Spraragen, and Zyda 2012). These include CAs like ALEC and MAMID, and AI Planners like ACRES/WILL, ActAddAct and EM-ONE. The use of appraisal theory is often supplemented with the cognition-based emotion model of Ortony, Clore, & Collins (OCC). OCC provides discrete emotions like fear, joy, and sadness that are associated with events. Valance can be assigned to emotions, with positive valance assigned to higher ranked goals. CAs that use appraisal theory with OCC include Soar-Emote and H-CogAff. AI planning systems include FA-

tiMA, EM, FLAME, Émile, and work by Gmytrasiewicz et al (Lin, Spraragen, and Zyda 2012). The arousal approach represents goal ranking as weighted drives that control emotional responses to events. CAs that utilize arousal include work by Ahn and a number of ACT-R extensions by Belkin et al, Fum et al., and Cochran et al. (Lin, Spraragen, and Zyda 2012). AI planners utilize arousal theory and weighted drives to rate the utility of plans during the search process (Gratch and Marsella 2004; Russell 1997).

Limitations of Existing Work: Unlike AI applications, social science relies on theories like Maslow’s Hierarchy that provide the source of human achievement goals and their ranking along basic human needs (Maslow 1943). However, key limitations exist for incorporating human-like goal ranking and emotions in AI and cognitive systems.

First, most systems focus on creating and executing plans in the most efficient way while minimizing or ignoring human limitations and bounded rationality (Kosinski and Zaczek-Chrzanowska 2003; Edelman 2015). Second, human-like cognitive impairments caused by faulty memory and cognitive biases are not easily captured (Hallion and Ruscio 2011). Some research into simulating certain types of impairments exists. For example, the Soar system’s episodic memory was evaluated to investigate how well it could function as more uncertainty is introduced (Derbinsky, Li, and Laird 2012) or different types of memory loss occur (Nuxoll et al. 2010). However, it would be infeasible to generalize such results to populations underrepresented in such studies. Also, associating positive and negative valence with specific actions ignores the complexity of human emotions, relying on social norms rather than unique biases of individuals. Hence, existing definitions of agent “drives” are too vague and inflexible to represent human emotions (Kleinginna and Kleinginna 1981).

Method

The methodology presented in this article begins with the mapping of Maslow’s needs to domain specific goals for an agent, as defined by (Gajderowicz, Fox, and Grüninger 2017a). Possible goals and actions are defined in the action schema (AS), which is domain-specific. The initial goal ranking is unique to the agent when the first plan is created. Maslow’s hierarchy determines goal ranking during execution when plan utility is recalculated. AS defines actions for reducing hunger, such as buying food or as in the homeless domain, going to a soup kitchen.

Once plan execution begins, a combination of MH ranking and ECOC stages is used to re-rank goals and change the utility of a plan. According to the ECOC theory, while individuals execute a plan, their emotions change in a way that is not directly related to their performance or a linear utility function. In our model, ECOC controls how an agent’s emotions impact their perception of success, and how that perception controls action selection and goal re-ranking during re-planning (Gajderowicz, Fox, and Grüninger 2017b).

How do AI systems and CAs address human bounded rationality? Economists and AI practitioner focus on understanding the internal processes of decision-making, which we call the “reasoning view”: an objective understanding




Figure 1: Action a with preconditions (pr_i^a) and postconditions (po_j^a).

of choices (Etzioni 1988; Russell 1997). Within psychology and sociology, rationality is a reference point, and researchers focus on interpreting observed behaviour, which we call the “behavioural view”: a subjective understanding of choices (Simon 1996; Etzioni 1988). This section discusses how these two views can be merged to provide a more general theory of motivation based on the relationship between Maslow’s Hierarchy, emotions, and decision making.

Human Behaviour as a Planning Problem

A STRIPS problem definition used in our work initializes the state of the world ($init(S_{init})$) and the agent’s goals ($goals(G)$) based on circumstances and goals expressed by an actual homeless client in the CHF dataset. Due to BR, the agent needs to re-rank goals and accommodate unplanned consequences of actions. AS contains domain specific actions. Each action a has preconditions (pr^a) that must be true before executing a and postconditions (po^a) that are true after execution, as per Figure 1. A state is explicitly defined as true by the statement $true(state)$. A state is explicitly defined as false by the statement $not(state)$. An action controls postconditions by either adding or deleting such states.

The STRIPS planning algorithm searches for a sequence of actions that change the initial state of the world from S_{init} to ΔS where $G \subseteq \Delta S$. The planning algorithm traverses a search tree of all possible states in the world. Nodes of the tree represent states, while edges represent actions that change one state to another.

According to BR, the search process traverses the search tree within BR bounds, defined as $BR(I, C, T)$. The *information bound* (I) causes certain branches in the search tree to be pruned prematurely as they do not match known information, where $I \in (0, 100)$ indicating percentage of information known. The *cognitive bound* (C) limits the complexity of reasoning a STRIPS planner can perform. In our model, C is the depth of tree traversal possible, where $C \in R$ ($R = \text{real numbers}$). Visited branches past a certain depth are pruned and not visited during the search. Finally, the *time bound* (T) limits the number of nodes the algorithm visits while traversing the search tree, where $T \in R$.

Plan Utility: Once a plan P is generated, the plan’s utility U_P is calculated. U_P is based on the utilities of actions U^a in P , incorporating their contribution towards the agent’s goals and changes in ECOC stages. To calculate U_P , each action’s utility is first calculated. For an action a , each precondition has a weight (pr_i^a) where $pr_i^a = 1$ if $pr_i^a \subseteq S$ and 0 otherwise, and $pr_i^a \in pr^a$. Each postcondition has a weight (po_j^a) where $po_j^a = 1$ if $po_j^a \subseteq G$ and 0 otherwise, and $po_j^a \in po^a$.

A special action weight ($pow_j^a(mh)$) considers an action’s postconditions that satisfy an outstanding goal for a particu-

lar MH level mh , as per Equation 1. This equation captures two characteristics of the relationship between MH levels. First, lower MH level goals are ranked higher than higher level goals, with physiological being highest ranked and self-actualization the lowest. Second, differences between levels are exponentially larger as one moves up the hierarchy from physiological to self-actualization.

$$pow_j^a(mh) = 1 - (po_j^{mh} - \min(MH))^{1/e} \quad (1)$$

Interim actions are ones that satisfy subgoals rather than MH goals directly. Interim actions are considered costs and have a constant negative weight of $pow_j^a(int) = -0.01$.

Each action also has an expectation of success associated with it based on previously satisfied goals at the same MH level or other interim actions. Rather than using a simple ratio of completed goals (G^S) to total goals ($|G|$), the ratio is adjusted using the function ECOC in Equation 2. It approximates the ECOC graph described in (Gajderowicz, Fox, and Grüniger 2017a). Note that according to the ECOC theory an agent that starts with zero achieved goals has a non-zero expectation of success. According to Equation 2 the estimated value is $ecoc(0) = 0.75$.

$$ecoc(X) = \begin{cases} 0.6 - \frac{\sin(8x - 1) + \cos(8x)}{x - 2}, & \text{if } X \leq 0.8; \\ X, & \text{otherwise.} \end{cases} \quad (2)$$

Each of the action weights described are then combined to calculate the utility of an action using its MH and interim goal weights, with U_{mh}^a defined by Equation 3 and U_{int}^a defined by Equation 4.

$$U_{mh}^a = \min(pr_i^a) \times \frac{\sum_j (po_j^a \times pow_j^a(mh))}{|po^a|} \times ecoc\left(\frac{|G_{mh}^S|}{|G_{mh}|}\right) \quad (3)$$

$$U_{int}^a = \min(pr_i^a) \times \frac{\sum_j (po_j^a \times pow_j^a(int))}{|po^a|} \times ecoc\left(\frac{|G_{int}^S|}{|G_{int}|}\right) \quad (4)$$

There is one factor that has not been captured yet: the contribution an interim action makes towards satisfying an MH goal. Until now, only U_{mh}^a have taken MH goals into account through postconditions, while U_{int}^a have no direct connections to MH goals. To rectify this, a final weight (aw_j^a) measures an action's overall contribution to a particular MH goal state in a plan. Each MH goal has a sub-plan P_{mh} required to satisfy that goal from initial state s^0 to goal state s_{mh}^g . We define the distance between states s^i and s^j in a plan with $dist(s^i, s^j)$. To calculate an action's contribution to P_{mh} , the action's distance from s^0 relative to $dist(s^0, s_{mh}^g)$ is calculated, as per Equation 5. Here, s^k is a state resulting from postconditions for action a that contribute to satisfying goals in sub-plan P_{mh} . Equation 6 combines U_a and all aw_k^a weights for action a to calculate its utility (U_P^a) for plan P .

$$aw_k^a = \frac{dist(s^0, s^k)}{dist(s_{mh}^0, s_{mh}^g)} \quad (5)$$

Table 1: Search Strategies

Strategy	Description
none	Select first plan found.
noneswap	Perform pairwise swapping of actions on first plan; choose plan with $max U_P$.
planutil	Search all plans; choose plan with $max U_P$.
planutil-swap	Search all plans; select plan with $max U_P$; perform pairwise swapping of actions; choose plan with $max U_P$.

$$U_P^a = U^a \times \sum_k aw_k^a \quad (6)$$

Finally, plan utility U_P can be calculated, as per Equation 7. An agent now has the ability to select a plan based on its utility. How the plan is chosen is based on one of the four search strategies in Table 1.

$$U_P = \frac{\sum U_P^a}{|P|} \quad (7)$$

Re-planning: Once plan P is generated and selected, the agent begins its execution. After the next available action is executed, U_P is recalculated. There are two main differences between the original U_P and U_P after execution. First, at least one goal has been achieved, and the expectation of success with $ecoc()$ is recalculated. Second, if a plan's utility falls below an "emotional threshold" execution stops, and certain goals are deferred until later. The re-planning process then begins. Re-planning is controlled by two thresholds that capture how an agent responds to the U_P during execution that were not considered during plan creation.

$ecoc^{Th}$ is the threshold that represents an emotional limit an agent can handle before pausing and reevaluating their plan. If $U_P < ecoc^{Th}$, plan execution halts and the plan's actions are considered for removal. $action^{Th}$ is the threshold that controls what actions are removed after a plan is halted. If $U^a < action^{Th}$, the action is removed along with any goals the action satisfies. Using the remaining goals, a new plan is created and execution begins again. Once a plan is successfully completed, meaning $U_P > ecoc^{Th}$ and all goals are satisfied, previously removed goals are added back, and a new plan is generated to achieve these outstanding goals. The cycle continues until all goals are satisfied.

If $U_P < ecoc^{Th}$ but either no goals can be removed or all goals are removed, a new plan is created and executed without considering $ecoc^{Th}$. This is called a "forced" plan. Any goals satisfied by a "forced" plan accumulate up to three times in the final goal count. If a "forced" plan cannot satisfy its goals it becomes a "failed" plan. Goals of a "failed" plan persist until the end of the simulation. After a "forced" or a "failed" plan, a new plan is created for the remaining goals. The new plan's S_{init} is the state that existed the last time $U_P > echo^{Th}$. The cycle continues until all goals are satisfied or the last plan fails.

Table 2: Dependent Variables

Variable	Description
ΔG_{act}	Actual goal ranking, grouped by 3-month periods.
ΔG_{sim}	Simulated goal ranking, grouped by re-planning cycles.
err_{traj}	Trajectory Error: ΔG_{act} versus ΔG_{sim} .
$learn_{rate}$	Learning Rate: required by agent to complete or fail plan.

Table 3: Environment Model

Variable	Description	Source
Services	Services satisfying MH goals.	AS
Actions	Actions satisfying any goals.	AS

Experiment: Homeless Client Emulation

The objective of this experiment is to find a model with a goal ranking and emotional thresholds that emulate how goals of homeless clients are re-ranked (ΔG) while participating in the Housing First program administered by CHF. Our hypothesis is that the use of MH needs and ECOC stages emulate client trajectory of changes in goals better than relying on STRIPS alone. A standard STRIPS planner with a breadth-first forward search is used to create a plan P that satisfies an agent’s goals G . The details of the experiment are given in (Gajderowicz, Fox, and Grüninger 2017a). Here, we summarize the metrics used and the results.

The score ($score_M$) for rating a model is the error between the final simulated goal ranking ΔG_{sim} and actual ranking ΔG_{act} for a participant in the CHF trial. A smaller score indicates a smaller error and a better model, where $score_M \in (0, 1)$. $score_M$ is the mean of err_{traj} , calculated as the mean square difference between ΔG_{act} and ΔG_{sim} goal ranking, and $learn_{rate}$, the ratio representing a simulated agent’s ability to go through the ECOC stages relative to the actual CHF client’s ability.


Experiment dependent variables are defined in Table 2. The environment and agent behaviour model variables are defined in Tables 3 and 4.

By executing tests for all combinations of search strategies $Search_S$, goal rankings MH and MH^R , and values for $ecoc^{Th}$ and $action^{Th}$ defined in Table 4, we have a total of 240 tests. The $score_M$ values for a subset of tests is shown in Figure 2. Here $score_M$ is calculated for all combinations of $ecoc^{Th}$ and $action^{Th}$ values in combination with MH^R goal ranking and “planutilswap” search strategy.

Analysis: The complete analysis is provided in (Gajderowicz, Fox, and Grüninger 2017a). Here we briefly discuss how to interpret the results and evaluate our hypothesis. For $score_M$ values for the subset of tests in Figure 2, we see that $ecoc^{Th} \geq 0.1$ and $action^{Th} = 0.3$ produce the best model to emulate the actual CHF participant. By comparing the results for all 240 model configuration combinations (not shown here), we can perform similar analysis and see which agent configuration is best. $score_M = 0.093$ in Figure 2 was the lowest $score_M$

Table 4: Agent Behaviour Model

Variable	Description	Value/Source
$BR(I, C, T)$	Bounded rationality	$I=100\%$, $C=40, T=2000$
SO	Search operator	Forward
$ecoc()$	Expectation function	Equation 2
$ecoc^{Th}$	ECOC threshold	0.0 – 0.6
$action^{Th}$	Action threshold	0.0 – 0.5
MH^G	Goals mapped to MH.	AS
$goals(G)$	Initial MH goal ranking (MH reversed).	$MH (MH^R)$
$init(S)$	Initial state (S) of world.	CHF data
$search_S$	Search strategy used.	none, nonswap, planutil, planutilswap

Figure 2: Model score ($score_M$) for tests using MH^R goal ranking and “planutilswap” strategy.

across all 240 tests. The “planutil” search strategy with the same goal ranking and thresholds produce the same score.

Hypothesis: Based on these results, we confirm our hypothesis that relying on both MH goal ranking and ECOC produces a better emulation of a person’s behaviour than without. The best models were those that did not ignore ECOC thresholds, where $ecoc^{Th} > 0$ and $action^{Th} > 0$. Also, MH^R goal ranking produced better models than MH for “planutil” and “planutilswap” search strategies.

Rational Behaviour: Since individuals cannot help but act on their emotions, any “unusual” behaviour in response to ECOC stages is not necessarily irrational. Also, as observed by (Henwood et al. 2015), goal ranking that differs from MH is not necessarily irrational. The preferred MH^R of the emulated agent may simply reflect their priorities at the time. The best indication of rationality is the search strategy. As per Table 1, we say that the “none” strategy is the least rational and “planutilswap” the most. The model that best emulated the CHF participant uses the “planutil” and “planutilswap” strategies, hence the agent is rational.

Conclusion

In this article, we presented a general model of human motivation and goal ranking. This model relies on abstract hu-

man needs and a novel way of incorporating emotions in an AI planner. Our model is of high-fidelity due to two key contributions: 1) low level abstraction of initial agent goals that replace traditionally predetermined goals; 2) model's independence from predefined emotional responses. The model demonstrated that clients act rationally but rank their goals in what may be perceived as unusual, but not irrational. The homeless client domain offers unique challenges that may not exist outside of this domain. The high-fidelity and generality of our model may benefit analysis of other populations that are also difficult to study.

Acknowledgments

The authors would like to thank the Calgary Homeless Foundation for sharing their dataset with us, and state the our findings do not reflect the views of the Foundation.

References

- Anderson, M. L.; Gomaa, W.; Grant, J.; and Perlis, D. 2013. An Approach to Human-Level Commonsense Reasoning. In Tanaka, K.; Berto, F.; Mares, E.; and Paoli, F., eds., *Paraconsistency: Logic and Applications*. Dordrecht: Springer Netherlands. chapter 12, 201–222.
- Blythe, J. 1998. *Planning under Uncertainty in Dynamic Domains*. Phd, Carnegie Mellon University.
- Brafman, R. I., and Chernyavsky, Y. 2005. Planning with Goal Preferences and Constraints. *Proceedings of the International Conference on Automated Planning and Scheduling (ICAPS)* 182–191.
- Choi, D. 2010. Nomination and Prioritization of Goals in a Cognitive Architecture. In *10th International Conference on Cognitive Modeling*, 25.
- Derbinsky, N.; Li, J.; and Laird, J. E. 2012. A Multi-Domain Evaluation of Scaling in a General Episodic Memory. *Aaai2012* 193–199.
- Edelman, S. 2015. The minority report: some common assumptions to reconsider in the modelling of the brain and behaviour. *Journal of Experimental & Theoretical Artificial Intelligence* 3079(September):1–26.
- Etzioni, A. 1988. Normative-affective factors: Toward a new decision-making model. *Journal of Economic Psychology* 9(2):125–150.
- Fikes, R. E.; Hart, P. E.; and Nilsson, N. J. 1972. Learning and executing generalized robot plans. *Artificial Intelligence* 3(1972):251–288.
- Fox, M. S. 1983. *Constraint-Directed Search: A Case Study of Job-Shop Scheduling*. Phd, Carnegie-Mellon University.
- Gajderowicz, B.; Fox, M. S.; and Grüninger, M. 2017a. Report on the General Model of Motivation and Goal Ranking. Working Paper. *Centre for Social Services Engineering, University of Toronto*. Revised: 21 August, 2017.
- Gajderowicz, B.; Fox, M. S.; and Grüninger, M. 2017b. Requirements for Emulating Homeless Client Behaviour. In *Proceedings of the AAAI Workshop on Artificial Intelligence for Operations Research and Social Good*, 7. San Francisco: AAAI Press.
- Gratch, J., and Marsella, S. C. 2004. A domain-independent framework for modeling emotion. *Cognitive Systems Research* 5(4):269–306.
- Hallion, L. S., and Ruscio, A. M. 2011. A Meta-Analysis of the Effect of Cognitive Bias Modification on Anxiety and Depression. *Psychological Bulletin* 137(6):940–958.
- Hart, P. E.; Nils, J. N.; and Raphael, B. 1968. Formal Basis for the Heuristic Determination of Minimum Cost Paths. (2):100–107.
- Hendler, J. A.; Tate, A.; and Drummond, M. 1990. AI Planning : Systems and Techniques. *AI Magazine* 11(2):61–77.
- Henwood, B. F.; Derejko, K.-S.; Couture, J.; and Padgett, D. K. 2015. Maslow and Mental Health Recovery: A Comparative Study of Homeless Programs for Adults with Serious Mental Illness. *Administration and Policy in Mental Health and Mental Health Services Research* 42(2):220–228.
- Kahneman, D. 2003. Maps of Bounded Rationality: Psychology for Behavioral Economics. *The American Economic Review* 93(5):1449–1475.
- Kleinginna, P. R., and Kleinginna, A. M. 1981. A categorized list of motivation definitions, with a suggestion for a consensual definition. *Motivation & Emotion* 5(3):263–291.
- Kosinski, W., and Zaczek-Chrzanowska, D. 2003. Pavlovian, Skinner, and Other Behaviourists' Contributions to AI. In Ellis, S. R.; Slater, M.; and Alexander, T., eds., *Workshop on Intelligent Motion and Interaction Within Virtual Environments*, 135–148. London, UK: NASA.
- Liaskos, S.; McIlraith, S. A.; Sohrabi, S.; and Mylopoulos, J. 2010. Integrating preferences into goal models for requirements engineering. *Proceedings of the 2010 18th IEEE International Requirements Engineering Conference, RE2010* 135–144.
- Lin, J.; Spraragen, M.; and Zyda, M. 2012. Computational Models of Emotion and Cognition. *Advances in Cognitive Systems* 2:59–76.
- Little, I., and Thiebaux, S. 2007. Probabilistic planning vs replanning. *ICAPS Workshop on Planning Competitions: Past, Present and Future*.
- Majercik, M., and Littman, M. L. 1998. MAXPLAN : A New Approach to Probabilistic Planning. In Simmons, R. G.; Veloso, M. M.; and Smith, S. F., eds., *Artificial Intelligence Planning Systems*, 86–93. Pittsburgh, USA: AAAI Press.
- Maslow, A. H. 1943. A theory of human motivation. *Psychological Review* 50(4):370–396.
- Norman, T., and Long, D. 1995. Goal creation in motivated agents. *Intelligent Agents* 277–290.
- Nuxoll, A.; Tecuci, D.; Ho, W. C.; and Ningxuan, W. 2010. Comparing Forgetting Algorithms for Artificial Episodic Memory Systems. *Remembering Who We Are Human Memory for Artificial Agents Symposium* (April):14–20.
- Reisenzein, R.; Hudlicka, E.; Dastani, M.; Gratch, J.; Hindriks, K. V.; Lorini, E.; and Ch. Meyer, J.-J. 2013. Computational Modeling of Emotion: Toward Improving the Inter- and Intradisciplinary Exchange. *IEEE Transactions on Affective Computing* 4(3):246–266.
- Russell, S. J. 1997. Rationality and intelligence. *Artificial Intelligence* 94(1-2):57–77.
- Sadeh, N. M., and Fox, M. S. 1995. Variable and value ordering heuristics for the job shop scheduling constraint satisfaction problem. *Artificial Intelligence Journal*.
- Shi, Z.; Zhang, J.; Yue, J.; and Qi, B. 2013. A Motivational System For Mind Model CAM. In *2013 AAAI Fall Symposium Series*, 79–86. Arlington, Virginia, USA: AAAI Press.
- Simon, H. A. 1996. *The Sciences Of The Artificial*. Cambridge, MA, USA: The MIT Press, third edit edition.
- Zafirovski, M. 2005. Is Sociology the Science of the Irrational? Conceptions of Rationality in Sociological Theory. *The American Sociologist* 36(1):85–110.