

See discussions, stats, and author profiles for this publication at: <https://www.researchgate.net/publication/313556478>

Requirements for Emulating Homeless Client Behaviour

Conference Paper · February 2017

CITATION

1

READS

496

3 authors, including:



Bart Gajderowicz

University of Toronto

8 PUBLICATIONS 13 CITATIONS

SEE PROFILE



Mark Stephen Fox

University of Toronto

236 PUBLICATIONS 8,718 CITATIONS

SEE PROFILE

Some of the authors of this publication are also working on these related projects:



Callisto: An Intelligent Project Management System [View project](#)



PolisGnosis Project: Representing and Analysing City Indicators [View project](#)

Requirements for Emulating Homeless Client Behaviour

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 the field of social services, practitioners work with homeless clients to help them exit homelessness. Due to the uniqueness of each client and the dynamic nature of their environment, practitioners must choose from a variety of theories and techniques to perform the required tasks. A major obstacle in automating these tasks is the lack of high-fidelity models that emulate client behaviour. The purpose of this paper is to identify the requirements for emulating a homeless client and how artificial intelligence (AI) can fulfill these requirements. This paper also highlights how such an emulation can contribute to the areas of social service policy evaluation, understanding client behaviour, and education.

Introduction

In the field of social services, practitioners, program administrators and researchers are tasked with helping clients sustainably exit homelessness. These groups work together to complete any one of the following three steps: 1) analyze a client's unique needs, goals, abilities, and circumstances, 2) predict what specific actions and subgoals will help them exit homelessness, and finally 3) change the client's behaviour by motivating them to perform the required actions.

Towards achieving these tasks, emulation of client behaviour serves three purposes: to evaluate intervention policies, to understand a specific client's behaviour, and to educate practitioners, communities, and clients. Policy evaluation relies on low-fidelity simulation models that can scale to a large number of clients. Understanding client behaviour requires a high-fidelity model that represents both the client and their particular circumstances. Finally, education also requires a high-fidelity client emulation model, but one that a human user can interact with.

When creating systems that assist with these tasks, practitioners and researchers at all levels are faced with three major obstacles. First, the information a client provides to a practitioner about their past is often incomplete or inaccurate. The acquisition of reliable information requires a strong working relationship between the practitioner and the client that is built on trust, a long and often unsuccessful process (Chen and Ogden 2012). Second, even if required

information is available, the high level of uncertainty impacting a client's behaviour and environment makes predictive analysis difficult. Third, homeless clients fail to follow through on commitments they make to practitioners or themselves (Kulkarni, Bell, and Wylie 2010). Persistent family and health problems like anxiety and depression, as well as constant safety risks cause clients to focus less on long-term goals and more on immediate short-term goals and specific departures from a plan (Gajderowicz, Fox, and Grüninger 2014).

Policy Evaluation

In the 1970's, it was argued that due to technological limitations it was unrealistic to emulate a client's decision-making process (Duncan and Curnow 1978). Instead research focused on the information available for assessing the professionals working with clients by associating outcomes with specific intervention programs. Large-scale simulation systems have used this data to predict how certain socioeconomic groups respond to changes in policy.

Most notable successes have been in simulating service providers within healthcare systems. Much of this work is based on explicit objectives with well-defined utility functions that optimally maximize resources while minimizing cost (Early 1999). The main drawbacks to this line of research are the assumptions that patients will behave like participants in a study: be available and prepared when appointments are scheduled and correctly follow-through on instructions. A recent survey of research analyzing healthcare services in the UK found that out of 142 papers reviewed, 51 used some form of simulation of client characteristics, but only one focused on real-time simulation of human behaviour (Fakhimi 2013).

The only categorical exception to healthcare research that treats patient behaviour as the main stochastic variable is the work done on emergency services like ambulance schedules and disaster response units where situational uncertainty is to be expected (Li et al. 2012). Models that require complex interaction between multiple agents have not gained much traction. For those that have, large-scale real-time simulations have been successful for models that rely on simple patient needs and patterns of behaviour, such as movement of people amongst their relationship networks during a virus outbreak (Beeler, Aleman, and Carter 2012).

Understanding Client Behaviour

Due to the uniqueness of each client and the dynamic and unpredictable nature of their environment, practitioners rely on “whatever works” to assess and modify client behaviour (Bricker and Tollison 2011). A variety of evidence-based assessment methods have been created to help understand how a client becomes homeless, a process referred to as their “pathways to homelessness” (Kim et al. 2010). To assess a client’s current state, questionnaires such as the “Service Prioritization Decision Assistance Tool” (SPDAT) measure their “vulnerability index” based on past and current circumstances. Once a client’s needs have been identified, a variety of assessment tools help to match those needs to available resources and services.

To address the last obstacle, changing a client’s behaviour, a variety of methods are used. For example, motivational interviewing (MI) is a technique for facilitating change in a client by identifying their own intrinsic motivations and life experiences to satisfy those motivations (Bricker and Tollison 2011). MI is used in conjunction with other techniques like Acceptance and Commitment Therapy to promote normative values and reduce destructive behavior.

Education and Assessment

Since the 1980’s, advancements in modeling techniques have allowed practitioners to more closely study client behaviour (Doolin 1986; Wolch and Rowe 1992). This work categorized clients along multiple dimensions, focusing on needs, constraints, and decision-making for different client demographics. A new picture of social clients began to emerge, moving away from negative stereotypes towards the real impact constrained resources and conflicting needs have on their lives.

One category of research into the lives of clients has focused on role-playing and role reversal for the purpose of education and assessment. These techniques involve simulating an interaction between a client, community, or service provider where at least one party, human or computer, is acting as someone else. The interaction is based on real client scenarios and behaviour patterns compiled into scripts and educational material.

The most complete and interactive client emulations rely on human actors and are applied to the education and evaluation of students across a number of social science disciplines (Logie and Bogo 2013). With technological advancements in speech understanding, virtual human interaction has also advanced. By “virtual” we mean an interaction where at least one party is a computer-based system simulating a human. A number of “virtual patient” projects are underway (Talbot et al. 2012; Smokowski and Hartung 2003). In each, virtual responses are programmed by researchers based on known responses by actual clients to automate the “virtual patient’s” interaction with human users.

Projects that educate the client and their community have also benefited from role-playing and role reversal. A special workshop consisted of an interactive theatrical play to help students understand and relate to the experiences of fellow students who experienced homelessness or were refugees

(Day 2002). A number of tools have been developed over the years to teach clients and youths at risk of becoming homeless how to better manage their actions and interpret their environment. For example, the Smart Talk software package is a tool for social workers to teach clients new strategies for resolving conflict without violence (Bosworth, Espelage, and DuBay 1998). It incorporates computer games, simulations, cartoons, animation, and interactive interviews. Focusing their attention on measures that deescalate conflicts, users learn from a repertoire of scripted non-violent responses to various scenarios.

Client emulation for the purpose of education suffer from several major drawbacks when compared to their real-life counterparts. First, an automated approach requires lower fidelity models that reduce expressiveness of client behaviour. This in turn reduces the emotional consequences a human user feels in reaction to their emulated partner (Bogo et al. 2014). Second, actors may produce inconsistent results reducing the reliability of assessments. Finally, the use of actors is most effective when the individuals interacting with the actors can pretend the scenario is real. This is especially difficult when an adult actor is playing a young child. In short, current emulation methods lack the type of engagement required during real interaction with clients.

The purpose of this paper is to identify the requirements for emulating a homeless client and the role AI plays towards achieving this goal. The next section provides challenges faced by the three key applications within social services mentioned above. We then provide a number of ways AI addresses these challenges, and finish off the chapter with a summary of requirements for emulating homeless clients using AI methods.

Challenges

There are separate but related challenges that must be addressed in unison.

Data Collection Data on homeless clients is primarily collected through interviews. It provides the empirical foundation for simulating this population. Point-in-time counts of homeless populations provide a snapshot of the count and demographics of the people living on the streets on a particular night, as well as their geographical and seasonal distribution (Meghan Henry et al. 2014). More detailed data is collected by interviewing participants in the context of a particular study that is evaluating an intervention program.

Each data collection method, however, has its limitations. Analysis performed on point-in-time data is limited to descriptive analysis of trends over time. For individual studies that capture data continuously, data is often limited to the participants’ level of satisfaction with the program being evaluated with limited follow-up after participants complete the program. Due to the difficult nature of tracking and interviewing clients, some of this information is “administrative data” obtained by the same organizations being evaluated (Wyrwich et al. 2007). The acquisition of reliable information directly from participants also has its challenges. It requires a strong working relationship between the interviewer and the client that is built on trust. Unfortunately

building trust is a long and often unsuccessful process and information provided may be unreliable (Chen and Ogden 2012). Finally, there are ethical dilemmas on the part of the program administrator. Special arrangements must often be made for individuals in a way that jeopardizes the integrity of the study (Falvo 2009).

These fundamental limitations of social service studies are a direct consequence of the participants' irregular lives and must be considered when evaluating the quality of the data collected.

Participant Characteristic Selection and Retention

Studies must limit their conclusions about impact to the remaining participants. Hence the final analysis is not simply based on initial hypothesis, but on what can be inferred from the changes exhibited by a statistically significant subset of remaining participants. The conclusion is then limited to the characteristics and demographics of those participants. For example, if a study focused on the effects of "housing first" intervention program and only one elderly male was left in the study, no conclusion could be made about the effects of "housing first" program on this demographic.

Following positive conclusions, studies are repeated with the objective to understand an intervention programs's impact on specific client populations (Stergiopoulos and Herrmann 2003). Once outcomes are known, specific factors are investigated. For a follow-up study to be effective, participant characteristics must match those of the original study. However, followup studies suffer from the same participation retention issues as the original study.

Macro-Predictions and Micro-Emulation At the macro-level, the quality of a simulation model depends on its ability to make predictions about the impact of a policy on a group of clients. At the micro-level, the quality of the emulation model depends on its ability to respond to specific situations comparably to a client with similar characteristics.

At the macro-level, it is difficult to translate successful programs into successful policy. For a program to be implemented as a policy, its success must be proven by a study or pilot, and be cost-effective by saving or making money. Evidence-based policy goes a step further and relies on specific data-driven studies to determine what policy has the highest chance of success within a specific political, economic, and social context (Stanhope and Dunn 2011). Given the uniqueness of each client, a high fidelity metric would be required to evaluate a policy on the entire homeless population and the various demographics it represents. In the absence of such models, specifically chosen categories of clients are identified from that population. This reduction in client representation allows a policy to focus on a single issue. The clear draw back of this approach is that a real policy is never implemented in isolation. There are other factors that change its effectiveness that are not considered during the study.

At the micro-level, the dynamic nature of a client's environment reduces the probability they will follow through on instructions provided by a practitioner. Practitioners have found that "motivating" rather than instructing their clients is more successful (Bricker and Tollison 2011). This is in

stark contrast to the regularity with which AI systems follow instructions provided. AI systems have well-defined goals and rules for satisfying the system's objectives. Identifying motivational factors and operationalizing motivation is still an open question in the social sciences (Kleinginna and Kleinginna 1981).

Rationality, Reasoning, and Behaviour Different sub-fields within social science address human behaviour in different ways. Economists assume rational behaviour based on a reasoning process that maximizes available utility towards a goal (Russell 1997). Psychologists rely on associating past experience with observed behaviour in a way that explains a client's reasoning process (Simon 1967).

Given observed behaviour, how do you distinguish what is rational and irrational? Is the observer completely rational or are they missing information that would explain a subject's behaviour as rational? We give the following definitions to make answering these questions easier.

Reasoning is the process of making a decision based on inference, also called "reasoning process."

Intuition is the process of making a decision without reasoning. Also called "intuitive process."

Behaviour is performed or planned actions and mannerisms in response to some stimuli, based on reasoning and intuition.

Next, a clear distinction needs to be made between what it means to be rational and irrational.

Rational reasoning maximizes available utility towards a goal.

Irrational reasoning is not rational to some degree.

Rational behaviour is based on rational reasoning.

Irrational behaviour is based on irrational reasoning.

Given these definitions, it would be impossible to emulate a rational client's behaviour by understanding their reasoning alone since it is impossible to capture all relevant factors (Kahneman 2003; Simon 1967). We must instead begin with observations of client behaviour, followed by client emulation using a rational reasoning process, a task well-suited for AI reasoners.

AI Requirements for Emulating Clients

Many AI systems simulate human tasks at different cognitive levels. High-level tasks include scheduling and planning with applications in various fields from robotics to industrial engineering (Sadeh and Fox 1990). Agent-based simulation allows an artificial agent to simulate social behaviours by following well-defined interaction rules (Gajderowicz, Fox, and Grüninger 2014). While an AI planner may be able to derive a plan for a client to follow, current planners are not designed with a homeless client's limitations in mind. First, rather than emulating human cognition, AI planners focus on optimizing performance for specialized tasks (Rintanen and Homann 2001). Second, goals are provided externally along with well-defined rules for identifying intermediate actions and subgoals.

This section outlines requirements that need to be met if client emulation by AI systems is to be achieved.

Observing and Capturing Behaviour

Collecting data from various sensors in a system is a reliable form of building and verifying models. Increasingly, metrics for city services are being developed as part of a move towards “smart cities” (Fox 2015). Except for clients in extreme cases or requiring immediate emergency services, most information known about individual clients is qualitative rather than quantitative, obtained through interviews or questionnaires as part of a treatment or study (Bricker and Tollison 2011; Nousiainen 2015; Rabinovitch 2015). This information captures the client’s life experiences as a sequence of events, choices, perceived factors, and emotional states. AI still lacks metrics that successfully monitor behaviour of clients using these social services.

Cognitive impairments caused by faulty memory and cognitive biases are a major issue with homeless clients (Beck 2014; Hallion and Ruscio 2011). Cognitive systems like ACT-R, PRODIGY, ICARUS, and Soar provide the architectural foundation for representing agent states using memory modules that can be made artificially impaired (Langley, Laird, and Rogers 2009). For example, the Soar system’s episodic memory was evaluated to see how well it could function as more uncertainty was introduced through prolonged interaction with the environment (Derbinsky, Li, and Laird 2012). If external regularity is assumed however, internal uncertainty is reduced by adjusting to the external “reality.”

In a separate study, the Soar system was modified to exhibit different types of memory loss, a requirement for cognitive systems that have a finite amount of storage (Nuxoll et al. 2010). As was expected, the modified system’s performance decreased at different rates for different types of memory loss. Unfortunately, it would be impossible to perform such detailed studies on the memories of homeless clients. Even if that was not so, a client still could not explain reasoning for many actions, especially when that reasoning was based on suppressed memories following a tragic event.

Reasoning From Observations

In addition to memories, emulating behaviour requires a suitable reasoning engine that can recreate a client’s past behaviour given their existing memories, then predict what choices will be made in the future given their current constraints. Different cognitive and environmental constraints must be taken into account. For example, clients employ different reasoning strategies that depend on different sources of stress. Given that a high percentage of clients suffer from anxiety, many of their choices depend on instinct rather than long and careful deliberation (Hwang et al. 2012; Beck and Clark 1997). Clients under high levels of duress caused by negative past experiences often incorporate negative and positive coping strategies that impact their decision-making (Wolch and Rowe 1992).

Goals Because humans behaviour is goal-driven, client emulation must replicate a client’s goal management strate-

gies. In AI and social sciences, goals fall into two separate categories: *achievement goals* that are provided *a priori* and *maintenance goals* that are added to a plan once certain requirements are not met (Grant and Dweck 2003; Hindriks and Van Riemsdijk 2008). In AI, achievement goals are provided externally by the system’s designers. Social science relies on theories like Maslow’s Hierarchy that provide the source of goals and their ranking (Maslow 1943). A social service practitioner will often depend on multiple theories and techniques to understand a client’s goals and goal ordering.

Search What kind of search are clients performing when looking for ways to achieve their goals? In most cases, it can be assumed that not all information is known, and some type of heuristics search is being performed. Constrained-directed heuristic search applies when a client knows or encounters the constraints they face while working towards their goals (Fox, Sadeh, and Baykan 1989; Beck and Fox 1998). In some situations they must work backwards from a goal, employing a backward search strategy, or work forward from their current state employing a forward search strategy. Very often when goals and means are not known clients will perform explorative search. For humans, explorative search is driven by unsatisfied basic needs or “meaning making,” an innate need to reduce uncertainty about the world around them (Engel et al. 2013). Once some information about required states is obtained, an island search strategy can be deployed from and to the states known to be on the path towards a goal. Generally, one can assume that a combination of each of these plays a role, and some variation of an explorative constrained-directed heuristic search strategy is employed.

Bounded rationality As mentioned previously, it is impossible to understand a client’s reasoning process by an observer. We can however consider factors that impact a client’s seemingly irrational behaviour. Bounded rationality provides the framework for incorporating cognitive and environmental constraints that may help understand a client’s behaviour (Simon 1955).

There are four main types of bounds influencing an individual.

Memory (short-term) : Humans are able to store and retrieve a finite amount of information into short-term memory (Simon 1972).

Cognition : Gaps in memory prevent complete knowledge about what strategy to choose when selecting an action and how to evaluate alternative actions (Simon 1972). Psychological disorders like anxiety cause incorrect cognitive associations between events and responses (Simon 1967).

Experiences : The aspiration and rates of success for past goals impact what information is stored and how it is structured (Simon 1955; 1967).

Biology : Real-time needs interrupt goal-driven behaviour caused by uncertain environmental events and physiological needs (Simon 1967).




Figure 1: Emotional Cycle of Change showing a client’s emotional “mood” regarding goals and “expectation” of success.

Different reasoners have incorporated parts of bounded rationality either directly or indirectly. For example early planners based on the STRIPS system provide ways of overcoming physical computing and memory limitations by grouping information into manageable chunks or a network of beliefs (Fikes, Hart, and Nilsson 1972; Sacerdoti 1974; Georgeff and Lansky 1987). If plan execution monitoring is possible, techniques like re-planning and contingency planning are used to adjust plans during execution created with missing information (Alfasi and Portugali 2004).

Emotions How do emotions influence our behaviour? There are many theories of emotions that contribute to behaviour. The vast majority, like BDI (belief-desire-intention) and OCC (Ortony, Clore and Collins) rely on “drives” that form a direct connection between a stimuli and a response (Reisenzein et al. 2013). In this work, emotional states control behaviour using a “triggering” threshold with positive or negative valence of emotions, where each emotional trigger is associated with specific actions (Gratch and Marsella 2004). There are several limitations that make this approach unsuitable for emulating homeless clients. First, it is difficult to identify positive and negative valence of emotions in specific situations described by the client. Second, “drives” are a vague representation of emotions while most systems rely on a predetermined and static assignment of valence to stimuli and a response (Kleinginna and Kleinginna 1981).

Creating an abstract “drive” is also problematic. First, abstract drives are represented as “motivations” but are difficult to operationalize as “motivation” lacks a formal definition (Kleinginna and Kleinginna 1981). Second, levels of abstraction are not modelled after human methods of abstraction but on methods that optimize a particular algorithm.

Instead of relying on specific connections between a stimuli and response, AI researchers may have more luck in mimicking the methods practitioners use to overcome the limitations introduced by “drives.” Practitioners observe clients as their behaviour changes from one situation to another. For example, when moving from destructive to constructive behaviour, a client transitions through multiple *stages of change*. A practitioner must recognize these stages and guide the client accordingly. The “Emotional Cycle of Change” is one way practitioners recognize stages of change, as presented in Figure 1 (Kelly and Connor 1979).

If a client is overly optimistic about a goal they may be in

the Uninformed Optimism (UO) stage. Here the practitioner recognizes that unknown prerequisites and unforeseen consequences will prevent the client from achieving their goals. Once faced with this reality, a client may lose confidence in the Informed Pessimism (IP) stage before succumbing to hopelessness in the Valley of Despair (VOD) stage. A system that relies on a graph of mood changes similar to ECOC does not need to know what stimuli triggered the action, only that overtime behaviour changes according to a pattern found in the ECOC graph.

Learning From Observations

In behaviour psychology, there are two main methods for learning information: *classical conditioning* where a “stimulus acquires the ability to produce a response” and *operant conditioning* where the “consequences that follow some behaviour increase or decrease the likelihood of that behaviour” in the future (Kosinski and Zaczek-Chrzanowska 2003). AI has been successful in mimicking both types of learning through reinforcement when the relationship between stimuli and response is known.

Quite often, however, a client may not know why they behaved in a certain way and can’t predict how they will behave in the future with any real certainty. Such choices can be based on psychological issues or suppressed negative memories from the past. When the direct connection between stimuli and response is not known, *latent learning* is required.

Latent learning has been applied most successfully to applications that require exploration of the entire search space before making a decision. The Procedural Reasoning System is one of the original examples, and was used to create a spatial model for a robot’s surroundings before performing tasks like “leave room” or “go to room X” (Georgeff and Lansky 1987). Such systems are self-organizing and rely on mental models of their environment (Kosinski and Zaczek-Chrzanowska 2003). Latent learning is a seldom used method for configuring a model due to its impracticality. Latent learners do not offer convexity of a training problem, making it impossible to know how close to a solution an algorithm is (Aslan et al. 2013). It expands the dimensionality of the problem and the search tree without clues as to the solution. Clients are however not able to create an entire map of their environment. They must instead rely on heuristics for exploration and chunking for retention of partial information.

Summary of Requirements

For a system to fully emulate homeless clients it must first replicate, in a general way, the flexibility and adaptability humans exhibit when faced with limited cognition and memory management. Table 1 summarizes the requirements for a high-fidelity client emulation model. A portion of this work is already achieved by cognitive systems and planners that reason under bounded rationality in one way or another. However, more work is required to base this reasoning on reliable information gathered from clients through interviews and observations.

Table 1: Client Emulation Requirements

Requirement	Descriptions
Abstraction	Exhibit abstract thinking that mimics a human process of abstraction.
Generalize	Able to generalize perception of events and memories.
Intuition	Rely on intuition rules to provide guidance when immediate response is needed, using direct stimuli-response connections.
Memory	Rely on memory of past experience to provide guidance in pursuing goals.
Feedback	Rely on monitoring and feedback to learn stimuli-response associations through reinforcement and latent learning.
Explore	Create plans by exploring the environment and reasoning about knowledge.
Explore and Learn	Create goals by exploring the environment and reasoning about learned knowledge.
Monitor	Monitor current behaviour and progress towards goal, and modify or abandon plans.
Non-linear Reasoning	Capable of non-linear reasoning and consider multiple scenarios at once.
Multiple Strategies	Apply different behaviour strategies that incorporate the situation and memories affecting a client's mental and emotional state.
Problem Relaxation	Relax or strengthen components of problem definitions based on bounded rationality.

Conclusion

Since the 1970's, a wealth of data regarding homeless clients has become available, with more recent work providing data that gives a new perspective on the difficulties clients face and the impact on their behaviour. This paper highlights key applications within social services where AI research can make considerable contributions to current state-of-the-art data collection and analysis methods. Requirements for a computer-based high-fidelity client emulation model are provided that extend these methods using AI techniques.

References

Alfasi, N., and Portugali, J. 2004. Planning Just-in-Time versus planning Just-in-Case. *Cities* 21(1):29–39.

Aslan, Ö.; Cheng, H.; Schuurmans, D.; and Zhang, X. 2013. Convex two-layer modeling. *Advances in Neural Information Processing Systems* 29:2985–2993.

Beck, A. T., and Clark, D. A. 1997. An information processing model of anxiety: Automatic and strategic processes. *Behaviour Research and Therapy* 35(1):49–58.

Beck, J. C., and Fox, M. S. 1998. A Generic Framework for Constraint-Directed Search and Scheduling. *AI Magazine* 19(4):101–130.

Beck, J. C. 2014. MIE1619S 2014 : Readings. 8126:1–7.

Beeler, M. F.; Aleman, D. M.; and Carter, M. W. 2012. A large simulation experiment to test influenza pandemic behavior. In *Simulation Conference (WSC)*, 1–7.

Bogo, M.; Rawlings, M.; Katz, E.; and Logie, C. H. 2014. Using Simulation in Assessment and Teaching:

OSCE Adapted for Social Work (Objective Structured Clinical Examination). Technical report, University of Toronto, Toronto, Ontario.

Bosworth, K.; Espelage, D.; and DuBay, T. 1998. A computer-based violence prevention intervention for young adolescents: Pilot study. *Adolescence* 33(132):785–196.

Bricker, J. B., and Tollison, S. J. 2011. Comparison of Motivational Interviewing with Acceptance and Commitment Therapy: A conceptual and clinical review. *Behavioural and cognitive psychotherapy* 39(5):541–559.

Chen, F. P., and Ogden, L. 2012. A Working Relationship Model That Reduces Homelessness Among People With Mental Illness. *Qualitative Health Research* 22(3):373–383.

Day, L. 2002. 'Putting Yourself in Other People's Shoes': The Use of Forum Theatre to Explore Refugee and Homeless issues in Schools. *Journal of Moral Education* 31(1):21–34.

Derbinsky, N.; Li, J.; and Laird, J. E. 2012. A Multi-Domain Evaluation of Scaling in a General Episodic Memory. *Aaai2012* 193–199.

Doolin, J. 1986. Planning for the special needs of the homeless elderly. *The Gerontologist* 26(3):229–31.

Duncan, I., and Curnow, R. 1978. Operational Research in the Health and Social Services. *Journal of the Royal Statistical Society. Series A* 141(2):153–194.

Early, D. W. 1999. A microeconomic analysis of homelessness: An empirical investigation using choice-based sampling. *Journal of Housing Economics* 8:312–327.

Engel, A. K.; Maye, A.; Kurthen, M.; and König, P. 2013. Where's the action? The pragmatic turn in cognitive science. *Trends in Cognitive Sciences* 17(5):202–209.

Fakhimi, M. 2013. Operations research within UK health-care: a review. *Journal of Enterprise Information Management* 26(1):21–49.

Falvo, N. 2009. Homelessness, Program Responses, and an Assessment of Toronto's Streets to Homes Program. Technical Report February.

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.; Sadeh, N. M.; and Baykan, C. 1989. Constraint Heuristic Search. volume 1, 309–315.

Fox, M. S. 2015. PolisGnosis Project: Representing and Analysing City Indicators.

Gajderowicz, B.; Fox, M. S.; and Grüniger, M. 2014. Requirements for an Ontological Foundation for Modelling Social Service Chains. In Guan, Y., and Liao, J., eds., *Proceedings of the 2014 Industrial and Systems Engineering Research Conference*.

Georgeff, M. P., and Lansky, A. L. 1987. Reactive reasoning and planning. *Proceedings of the sixth national conference on artificial intelligence (AAAI-87)* 677–682.

Grant, H., and Dweck, C. S. 2003. Clarifying achievement goals and their impact. *Journal of personality and social psychology* 85(3):541–553.

- 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.
- Hindriks, K. V., and Van Riemsdijk, M. B. 2008. Satisfying maintenance goals. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)* 4897 LNAI:86–103.
- Hwang, S. W.; Stergiopoulos, V.; O'Campo, P.; and Gozdzik, A. 2012. Ending homelessness among people with mental illness: the At Home/chez Soi randomized trial of a Housing First intervention in Toronto. *BMC Public Health* 12:787.
- Kahneman, D. 2003. Maps of Bounded Rationality: Psychology for Behavioral Economics. *The American Economic Review* 93(5):1449–1475.
- Kelly, D., and Connor, D. R. 1979. The emotional cycle of change. *The 1979 Annual Handbook for Group Facilitators*.
- Kim, M. M.; Ford, J. D.; Howard, D. L.; and Bradford, D. W. 2010. Assessing trauma, substance abuse, and mental health in a sample of homeless men. *Health & social work* 35(1):39–48.
- 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.
- Kulkarni, S.; Bell, H.; and Wylie, L. 2010. Why Don't They Follow Through? Intimate partner survivors' challenges in accessing health and social services. *Family & Community Health* 33(2):94–105.
- Langley, P.; Laird, J. E.; and Rogers, S. 2009. Cognitive architectures: Research issues and challenges. *Cognitive Systems Research* 10(2):141–160.
- Li, A. C.; Nozick, L.; Xu, N.; and Davidson, R. 2012. Shelter location and transportation planning under hurricane conditions. *Transportation Research Part E: Logistics and Transportation Review* 48(4):715–729.
- Logie, C. H., and Bogo, M. 2013. A Critical Appraisal of the Use of Standardized Client Simulations in Social Work Education. *Journal of Social Work Education* 49(1):37–41.
- Maslow, A. H. 1943. A theory of human motivation. *Psychological Review* 50(4):370–396.
- Meghan Henry; Cortes, A.; Shivji, A.; and Buck, K. 2014. The 2014 Annual Homelessness Assessment Report to Congress, Part 1: Point-in-time estimates of homelessness.
- Nousiainen, K. 2015. Reflecting narrative interview context as performance: interviews with former homeless persons with intoxication and mental health problems. *Nordic Social Work Research* 5(2):129–142.
- 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.
- Rabinovitch, H. R. 2015. *Snapshot of an object in motion: quantifying homelessness*. Master, Simon Fraser University.
- 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.
- Rintanen, J., and Homann, J. 2001. An overview of recent algorithms for AI planning. *Künstliche Intelligenz* 01(2):5–11.
- Russell, S. J. 1997. Rationality and intelligence. *Artificial Intelligence* 94(1-2):57–77.
- Sacerdoti, E. D. 1974. Planning in a Hierarchy of Abstract Spaces. *Artificial Intelligence* 5(1974):115–135.
- Sadeh, N. M., and Fox, M. S. 1990. Variable and Value Ordering Heuristics for Activity-Based Job-Shop Scheduling. In *Fourth International Conference on Expert Systems and Operation Managements*, 134–144.
- Simon, H. A. 1955. A Behavioral Model of Rational Choice. *The quarterly journal of economics* 69(1):99–118.
- Simon, H. A. 1967. Motivational and emotional controls of cognition. *Psychological review* 74(1):29–39.
- Simon, H. A. 1972. Theories of Bounded Rationality.
- Smokowski, P. R., and Hartung, K. 2003. Computer Simulation and Virtual Reality : Enhancing the Practice of School Social Work. *Journal of Technology in Human Services Computer Simulation and Virtual Reality* 21(1/2):37–41.
- Stanhope, V., and Dunn, K. 2011. The curious case of Housing First: the limits of evidence based policy. *International journal of law and psychiatry* 34(4):275–82.
- Stergiopoulos, V., and Herrmann, N. 2003. Old and homeless: a review and survey of older adults who use shelters in an urban setting. *Canadian journal of psychiatry. Revue canadienne de psychiatrie* 48(6):374–80.
- Talbot, T. B.; Sagae, K.; Bruce, J.; and Rizzo, A. 2012. Sorting Out the Virtual Patient: How to Exploit Artificial Intelligence, Game Technology and Sound Education Practices to Create Engaging Role-Playing Simulations. *International Journal of Gaming and Computer-Mediated Simulations* 4:1–19.
- Wolch, J. R., and Rowe, S. 1992. On the Streets: Mobility Paths of the Urban Homeless. *City & Society* 6(2):115–140.
- Wyrwich, K. W.; Metz, S. M.; Kroenke, K.; Tierney, W. M.; Babu, A. N.; and Wolinsky, F. D. 2007. Measuring patient and clinician perspectives to evaluate change in health-related quality of life among patients with chronic obstructive pulmonary disease. *Journal of general internal medicine* 22(2):161–70.