

Qualitative Reasoning about Investment Decisions

Kenneth D. Forbus, Thomas Hinrichs

Qualitative Reasoning Group, Northwestern University

2233 Tech Drive, Evanston, IL, 60208, USA

forbus@northwestern.edu

Abstract

Qualitative decision-making focuses on framing problems, to identify simple solutions when possible, and identify what additional information is needed in more complex situations. This paper explores using qualitative models to enable autonomous systems to make investment decisions, using an open-source strategy game (Freeciv) as an experimental platform. Two kinds of investments, increasing capacity and increasing capabilities, are identified, and tied to the idea of functional subsystems within a system being constructed and/or managed. Drawing on ideas from prior research in qualitative preferences, we describe algorithms for making investment decisions that have been tested in Freeciv.

1 Introduction

Traditional theories of decision-making are quantitative, focusing on making optimal decisions. Such techniques have had many practical applications. But the framing of decision problems themselves remains informal, with key parts of the expertise residing only in the minds of the analysts. Just as qualitative reasoning for science and engineering focuses on framing problems and providing simple solutions with low data requirements, qualitative decision-making [Forbus & Hinrichs, 2018] involves formulating decision problems in ways that either allow them to be solved with little data, or identify what additional information is needed. Thus, like the use of qualitative reasoning in science and engineering, qualitative decision-making has the potential to lead to improved tools to support human decision-makers and to more autonomous AI systems whose decisions are understandable to their human partners.

Complex systems, such as companies, governments, and even entire civilizations, give rise to many kinds of decisions. Some of the most difficult are investment decisions, where resources are expended to provide future, rather than immediate, benefits. This paper describes a simple qualitative model of investment decisions, using the open-source strategy game Freeciv as a domain. Freeciv is useful because playing it involves many kinds of decisions, including investment decisions. We start by describing the background we draw upon and related work. Then we describe

our approach, including the ideas of functional subsystems, types of investments, and algorithms for using these ideas in Freeciv. An experiment is described that evaluates the potential for these algorithms to improve autonomous decision-making.

2 Background and Related Work

We start by examining prior work on qualitative decision-making, then we examine AI research that uses Freeciv.

2.1 Prior work on qualitative decision-making

Qualitative reasoning has been successfully used to capture aspects of decision-making previously, including investment decisions. The most recent is [Rovira et al. 2018], who use fuzzy linguistic term sets to map from language to values, and the Technique for Order Preference by Similarity to the Ideal Solution [Hwang & Yoon, 1991], which then ranks alternatives with respect to positive and negative ideal solutions. These techniques enabled them to assess the degree of consensus among a set of professional investors about investment opportunities. Our problem is different in that we are providing an autonomous system with the ability to use qualitative information in making its decisions in a simulated world.

Another line of work that has inspired us is research on preferences [Rossi et al. 2011]. Some traditional economic theorists postulate that people are capable of producing, for any object, a numerical estimate of its utility.¹ That people have trouble comparing the utility of different types of objects, and that there can be multiple dimensions of utility, suggests that instead it can be worth formulating methods to deal with the kind of ordinal (i.e. qualitative) distinctions that are often far easier to obtain from people. In this paper we draw upon ideas for reasoning about qualitative preferences from work summarized by [Santhanam et al. 2016]. Specifically, they point out that many problems can be broken down into a set of *dimensions*, within which various attributes can be measured, at least ordinally, with preferences expressed across dimensions. For example, in designing a smartphone, battery size is viewed as a positive from

¹ Some even postulate a unit for such measures, the *util*.
https://en.wikipedia.org/wiki/Cardinal_utility

the perspective of battery life, but a negative from the perspective of weight. The Tradeoff-enhanced Conditional Preference nets (TCP-Net) formalism [Brafman et al. 2006] is the closest to the scheme we use here. CP-Nets enable preferences to be expressed between values of variables based on the values of other variables (hence, conditional preferences). TCP-Nets go further in providing ways to express preferences among dimensions, e.g. weight is more important than battery life. Priorities across functional sub-systems, as explained below, are a means of influencing decision-making based on properties of the world.



Figure 1: A Freeciv civilization

2.2 Prior AI research using Freeciv

Freeciv² is an excellent domain for AI research because it provides a *constructive dynamic world*. While the simulated world is expressed as a 2D map, it is far more complex than Go: There are by default 4,000 tiles, each of which has different types of terrain, randomly generated at the start of the game, and which can be improved by players to some degree. By contrast, Go is played on a 19x19 grid, whose spatial properties are uniform, cannot be changed, and are always visible from the start. In Freeciv, players must explore, found cities, build a transportation network, build up economies, and defend themselves against aggressors (or start wars themselves). There are classic guns/butter tradeoffs in managing cities, i.e. providing food and luxuries versus investing in economic production and scientific research. A deep tree of technologies, whose fruits are uncovered via research, provide an expanding set of capabilities and opportunities. For example, there are over thirty improvements that can be built in cities, and fifty different types of units that can be constructed. There can also be many players, making the modeling of allies and opponents, and when they switch loyalties, quite complex.

Consequently, we and many others have used Freeciv for research purposes. For example, [Goel & Rugaber, 2015] have developed a software environment for creating knowledge-based agents to tackle sub-components of the game, such as managing a single city. A combination of

text analysis and Monte Carlo simulation was used by [Brannan et al. 2012] to do reasonably well at simplified versions of Freeciv, i.e. much smaller boards and many fewer turns, so that most of the complexity of the game was factored out. We note that most constructive dynamic worlds do not have existing simulators where eight-way lookahead Monte Carlo search can be used, making their approach of limited utility.

Our own work uses the Companion cognitive architecture [Forbus & Hinrichs, 2017] to explore autonomous decision-making. This has included using qualitative models to represent strategic thinking in the game [Hinrichs & Forbus, 2015], learning qualitative models from reading advice about the game [McFate et al. 2014], and learning about the game via demonstration [Hinrichs & Forbus, 2012a]. The nature of constructive dynamics worlds is what drove us to develop *type-level qualitative representations* [Hinrichs & Forbus, 2012b], which we use in this work, as well as the learned knowledge from the prior efforts in the Companion’s qualitative model of the game.

3 Our Approach

We treat decisions as having two basic forms. *Continuous value decisions* involve choosing a numerical value for a parameter of a system, such as a tax rate or a percentage of one’s budget to give to charity. *Enumerable decision* involve making a selection from a finite set of choices, as when deciding which movie to go see, or which bank to open an account with. This paper focuses on enumerable decisions, but see [Hinrichs & Forbus, 2019] for an approach to learning to make continuous value decisions.

Evaluating choices naturally requires exploring the tradeoffs between costs and benefits. Being able to construct relevant descriptions of costs and benefits as part of model formulation for decision-making is important. Costs are defined in terms of a type of resource, e.g. time or money. We formalize costs abstractly via `COSTFn`, a unary function whose domain is the ontological category `ResourceType` and whose range consists of unary functions whose domain is open (i.e., `Thing` in the OpenCyc ontology) and whose range consists of continuous quantities. Thus

```
(valueOf ((CostFn MonetaryValue) Ticket2)
 (USDollarFn 9))
```

indicates that the cost of a particular movie ticket is \$9. This use of a higher-order function, by contrast with the two-argument version in [Forbus & Hinrichs 2018], supports type-level qualitative representations, which are used below. In Freeciv, the types of resources include gold points (i.e. money), food points and production points that are generated by cities.

Costs can be immediate (e.g. the price of a car) or enduring (e.g. the cost of maintaining that car). Maintenance costs can be defined as higher-order functions in terms of resources as well. For example, units in Freeciv may have maintenance costs in terms of gold, food, and/or production points. In the real world, owning a car subjects one to peri-

² <http://www.freeciv.org/>

odic additional costs, such as fuel, repairs, and license fees. These costs are enduring, in that they last through the period of owning the car, but are not continuous during that period. A qualitative model of car ownership is responsible for identifying the kinds of costs that could be incurred, but not for specifying a quantitative level of detail that might be used, for example, in a spreadsheet trying to estimate total cost of ownership for various vehicles.

Benefits are somewhat more subtle than costs, because they are often intangible. Private ownership of cars is currently popular in many countries because of perceived convenience relative to costs. Innovations such as congestion taxes in crowded cities can be viewed as methods to push societal costs back onto those who benefit most directly from car ownership, as does demand pricing for high-speed lanes for some roads in Texas. Just as with costs, some benefits are immediate and others are enduring. An example of an immediate benefit in real life is buying resources, such as food or water, since then one can use them. An example of an enduring benefit is taking a class, since presumably the knowledge and skills gained persist long after the class is over. In Freeciv, cities can produce things that have immediate benefits, such as military units and coinage. Other city products have enduring benefits, e.g. building a library leads to higher science output in that city. Research in Freeciv is another example of enduring benefits, since (in the game at least) technologies once discovered are never forgotten.

The distinction between immediate and enduring benefits leads us to make a distinction about types of investments. Some investments increase *capacity*, i.e. provide new resources that can be used in many ways. Military units can be used to defend cities, or attack current enemies. Stockpiling gold enables construction to be sped up. Other investments increase *capabilities*, i.e. provide new types of resources that provide additional freedom of action. Doing research in Freeciv, for example, leads to the ability to create new kinds of buildings and units, which must then be produced by cities in order to add them to that civilization’s capacity.

So far our discussion of capacities and capabilities has been very abstract. In dealing with a complex, growing dynamical system such as a Freeciv civilization, it is useful to consider it in terms of *functional subsystems*. We tacitly do this when we describe one country as an economic powerhouse and another as a military juggernaut – we are focusing on one subsystem of what is an extremely complex system. The importance of decomposing complex artifacts into different perspectives has long been known in science and engineering practice, with one formalization provided by compositional modeling [Falkenhainer & Forbus, 1991]. For example, a smartphone design team must consider power, thermal properties, radio performance, and software properties. These multiple perspectives can and do trade off against each other, e.g. the Samsung Note 7 recall, where optimizing for battery life without adequate care in terms of thermal properties was responsible. While engineering design can involve awesome complexity, with the possible

exception of software, it does not have to deal with adversaries attacking it while it is growing, which is a common occurrence in strategy games.

For decision-making in strategic problems, functional subsystems provide a means of expressing priorities. A civilization in Freeciv needs to expand as rapidly as possible, producing new settlers to found new cities when feasible and building other units, such as workers, to provide economic benefits. On the other hand, if war has broken out, then beefing up the military, and choosing technologies that could lead to military benefits, becomes higher priority. Table 1 shows the functional subsystems our qualitative model of Freeciv uses.

System	Unit/Building	Technology
Growth	Settler	Pottery
Economic	Marketplace	Currency
Research	Library	Writing
Awareness	Diplomat	Writing
Military	Warrior	Wheel
Transportation	Worker	Bridge Building

Table 1: Functional subsystems in a qualitative model of Freeciv, with an example unit or building and an example technology that enhances them. The order indicates the Companion’s default priorities.

The qualitative model includes assertions that indicate which technologies benefit what systems, e.g. the assertion `(techAddsSystemCapability FC-Tech-Pottery GrowthSystem)`

expresses the connection between these two concepts in Table 1. It should be noted that technologies can impact multiple systems: for example, Writing is a prerequisite for building both Libraries and Diplomats, and hence its discovery adds capabilities to both Research and Awareness systems.

Treating functional subsystems as priorities provides a means of stratification of options during decision-making. If money is tight, then taking the bus, while slower, beats driving a car and paying for parking, whereas if time is viewed as a more important constraint than cost, the reverse will be true. Functional subsystems serve as dimensions in decision-making, following [Santhanam et al. 2016], with specific properties of the outputs of a choice serving as the within-dimension properties over which there are preferences.

Algorithm 1: Making Research Decisions

Inputs: A list of topics that can currently be researched, a priority list of functional systems

Output: The next topic to research

1. For each system S in the priority list,
 - a. If any of the research topics provides a capability for it, choose one of them at random and return it as the output.

Algorithm 1 shows our method for making research decisions. Since the priority list is complete and every technology contributes to at least one system, there will always be a research topic chosen. Notice that this method does not exploit any lookahead. This is deliberate, since the technology tree contains richly connected dependencies. This means that players will end up researching most of the technologies in any case, to reach advanced ones needed to win the game, and so the main question is how to make local choices that provide immediate relevant benefits. A more sophisticated algorithm might regress the properties of what could be built backwards, e.g. if the military situation required ranged units then researching the Warrior Code before Horseback Riding could be advantageous, because it enables Archers to be built, whereas if mobility were more important, researching them in the reverse order would provide Horsemen sooner.

Decisions as to what to build in a city take into account the benefits that a constructed item has relative to particular systems. These are determined by rules constraining this predicate:

```
(benefitOfProduction <product> <system>
  <value>)
```

where *<product>* is a unit, building, or coinage, *<system>* is a functional subsystem, and *<value>* is an integer used for expressing preference. These rules break down into three categories. The first uses attributes of the product, e.g. for military units, attack strength and defense strength are numerical values that are used as values, since a unit with a larger defense strength provides more benefit than a unit of lesser defense strength. At the level of functional subsystem, we do not distinguish between attack and defense strength, since any preferences as to those should depend on the plans that the unit is being built to serve. City defenders should typically have high defense strength, for example. But if one of the defense plans is to have a defender attack a unit heading towards a city, in that case ranged weapons and a larger attack strength might be preferred. For simplicity, we postpone such considerations to future work.

The second category of rules determining the benefit of a product choice involves using the qualitative model to calculate benefit relative to a system. For example, the fact that building a Library in a city increases its science output is expressed in the type-level qualitative model as

```
(positivelyDependsOn-TypeType
  (MeasurableQuantityFn cityScienceTotal)
  FreeCiv-City FC-Building-Library
  cityHasImprovement)
```

This results in a value of 5, chosen so that values derived from the qualitative model are higher than hand-generated values (e.g. Coinage has a benefit of 1), so that creating buildings such as Marketplace and Bank will be more attractive, since their relevance to the economic system is similarly derived from the qualitative model. We view learning to adjust such priorities by experimentation as an important avenue for future work. Only seven rules suffice for benefit computations, since most of the information is being extracted from the type-level qualitative model.

Algorithm 2 describes a method for using these concepts of benefits and functional subsystems to make production decisions. Since everything that can be built benefits at least one kind of system, and the priority list includes all systems, there is always some system that will show a benefit, and hence this algorithm will always return. We note that this algorithm might not make the best decision – there could be two equal-value options for a system *S*, but one of which will also benefit a lower-priority system, so it could turn out to be a better choice. Notice that we are also not taking maintenance costs into account – that requires the Companion to have a better model of its current economic state than it has. A more subtle problem is the sorcerer’s apprentice scenario, i.e. what keeps the system from producing the same thing over and over again? For buildings, the choices provided by the game engine prevent producing multiple copies of the same building, e.g. building two Libraries in the same city is impossible. For military units, an additional criterion in the simulation API provides a maximum number of units that any city should support. For coinage, which can be produced indefinitely, the simulation API automatically produces an event each turn to force the Companion to re-examine any coinage production decisions.

Algorithm 2: Making Production Decisions

Inputs: A list of possible products for a city and a priority list of functional systems.

Output: A decision about what to build next in that city.

1. For each system *S* in the priority list,
 - a. For each product *P* in the possible product list for that city, calculate the value of its benefit for *S*, if any.
 - b. If there are benefits for *S*, choose the *P* with the highest value.

Currently Companion capabilities are not capable of exploiting the full complexities of the simulation. Production can be changed in progress, for example, and gold spent to rapidly finish building something to satisfy an urgent need (e.g. a defender for a city about to be attacked). Our system currently does not do either of these things. Moreover, it currently does not plan to build Wonders, nor does it establish trade routes. Nevertheless, its current capabilities are sufficient to provide some evidence for the utility of these ideas, as the experiment described next indicates.

4 Experiments

To evaluate these ideas, we focus on research decisions (i.e. investment in capability) and production decisions (i.e. investment in capacity) in Freeciv. Since we are focused on enumerable decisions, we leave the rates of investment in production and research fixed at their default settings across all conditions. We look at the algorithms above contrasted with baseline algorithms.

For research decisions, we contrast two methods:

1. Select the lowest-depth technology in the technology tree that can be researched. If there are several

with the same lowest depth, choose based on alphabetical order.

2. Use priorities over functional subsystems to select which of the technologies that can be currently researched to work on (Algorithm 1 above).

The closest technology heuristic is a reasonable baseline because dependencies among the technologies means that most early technologies must be explored sooner or later. Most of them also provide valuable capabilities for various systems (e.g. better defenders, buildings that enhance economic growth), so they are worth having.

For production decisions, we contrast two methods:

1. Random choice. Of the things that a city can build, one of them is chosen randomly to build next. (Wonders are excluded, since in the early game they take too long to build.)
2. Use priorities over functional subsystems and properties of what is built to select what to build (Algorithm 2 above).

We use random choice as a baseline because some kinds of entities (i.e. units) can be built multiple times, so any fixed ordering would not cover the range of possibilities.

We examine behavior across two kinds of scenarios:

- **Empty maps:** These test the ability of a strategy to grow and expand, factoring out warfare and competition for territory.
- **Combat maps:** These worlds have four additional civilizations, to test the ability to trade off between warfare and economics.

In graphs, we use “E” as a suffix to indicate results restricted to empty maps, and “T” as a suffix to indicate results

restricted to combat maps. We use small maps, only 2,000 tiles instead of the default 4,000, because we want frequent contact between civilizations in the combat maps. The fraction of land was set to be 85%, compared to the default of 30%, so that all civilizations will start on the same continent and they will have the most room to grow. Huts were not allowed, because they randomly provide technologies, units, cities or kill units, thereby muddying the signal. Ten empty maps and ten combat maps were generated using these parameters, with the same twenty maps used in each condition.

Our experiment has a two-by-two design,

- Research decision strategy
 - CR: Closest research
 - QR: Qualitative decision
- Production decision strategy
 - RP: Random production
 - QP: Qualitative production

We use two measures for evaluation: (1) number of cities that are built, and (b) amount of gold accumulated. Number of cities matters because a powerful economy is needed to win the game, and the strength of a civilization’s economy depends on the number and strength of its cities. The amount of gold accumulated is another measure of economic strength, and since it can be used to speed up production, it provides more freedom of action.

We predicted that qualitative decision-making would outperform the baselines, i.e. QRQP > CRRP. This was borne out for both number of cities and amount of gold ($p < 0.05$). Figure 2 shows the averages across all maps for the two conditions.

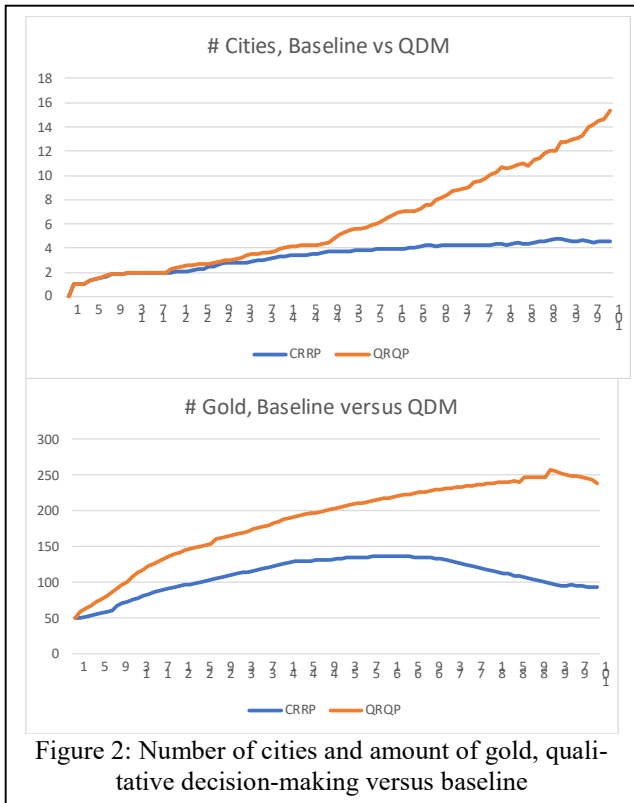


Figure 2: Number of cities and amount of gold, qualitative decision-making versus baseline




Figure 3: City production across conditions by map types

Notice that the amount of gold begins to drop for both conditions, albeit more severely for the baseline. There are two reasons for this. First, the maintenance costs of units and buildings were not considered in any of these strategies, and as they expand, these costs start becoming significant. Second, in 18% of the combat games, the Companion's civilization was eradicated by its opponents, which sets the gold to zero at the end of those games. (The current grasp of tactics in the Companion player is rudimentary at best, since this is a topic we plan to have it learn in future work.) We could find no statistically significant difference across conditions for predicting survival, although random production games accounted for only one of the losses, with qualitative decision making on production used in six of the losses. Close inspection as to why indicates that, since there is little else to build in early stages of the game, random production games led to large numbers of military units being produced (even in empty maps, which in that case were a needless drain on the civilization). This made the very small number of cities that were produced extremely well defended. Moreover, the high expansion rate with qualitative production choices led to larger-footprint civilizations and a larger selection of targets for their neighbors.

The productivity can be seen by comparing city production across conditions, as illustrated in Figure 3. By prioritizing growth when possible, qualitative production strategies do far better in terms of creating a large civilization than the baseline.

What about gold production? To our surprise, the combination of baseline research and qualitative production decision-making significantly out-performed the use of qualitative decision-making strategies for both types of maps, as shown in Figure 4. Comparing the research paths chosen by

the two methods in the empty maps (where each path will always be the same) illuminates why. The first technology that enables a new building which provides economic capabilities is Currency, which enables the construction of Marketplaces. In the baseline research path, this is the ninth technology discovered (which requires many turns). Because the qualitative research strategy has, after Growth, strengthening the Economic functional subsystem as its priority, it reaches Currency much earlier, as the fourth technology it discovers. Once it discovers Currency, building Marketplaces becomes a priority, if Settlers (which fuel growth) cannot currently be built in a city due to population constraints. Absent this alternative for production, coinage is produced when economics is a priority. Thus the baseline research path provides a much longer period where coinage is the preferred option. In fact, across the empty maps, coinage was the choice for production 957 times for baseline research games versus only 339 times for qualitative research games, with 118 Marketplaces being built in baseline research games versus 141 in qualitative research games. While gold is useful, Marketplaces provide more long-term benefits, demonstrating a limitation in using gold production as a measure of investment prowess.

5 Summary and Future Work

This paper illustrates how qualitative representations can be used to capture aspects of decision-making about investments in constructive dynamic worlds. We built upon prior research in QR, especially work on qualitative preferences, and showed how simple ideas of costs, benefits, and functional subsystems can be used to help autonomous systems make better investments.

There are still many things to do. First, the current production decision-making process only uses a crude form of situational awareness (i.e. at war or not). Detecting threats in advance requires better models of what happens during the game, which we suspect can be acquired through expanding the notion of episodic memory presented in (Forbus & Hinrichs 2018) to broader spans of time. Having finer-grained models of time estimates (e.g. Hancock et al. 2018) could help better determine alternate courses of action, e.g. building a defender versus moving one in from elsewhere when the barbarians are at the gates. Second, we would like to examine if one of the algorithms for TCP-nets could be used to provide better decisions, compared to our simple priority-based scheme. Third, we would like to explore other domains beyond Freeciv, to see how general these ideas are.

Acknowledgments

This research was supported by the US Air Force Office of Scientific Research.

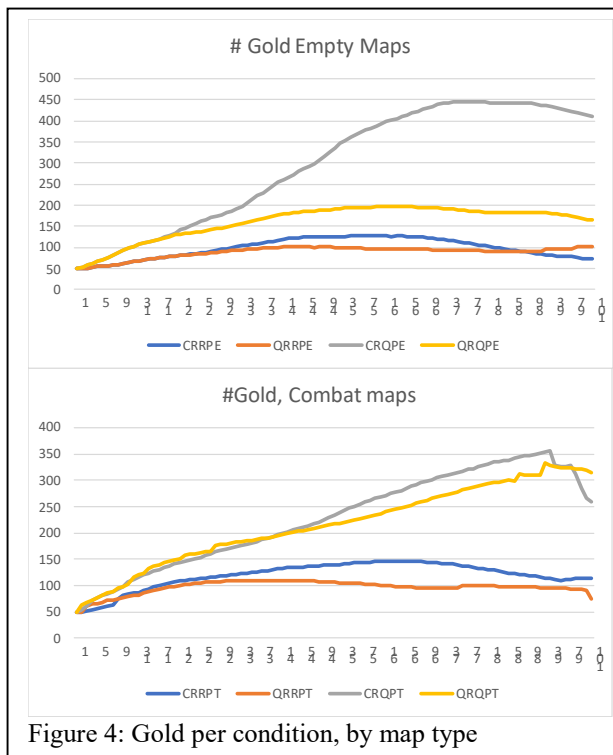


Figure 4: Gold per condition, by map type

References

- [Brafman et al. 2006] Brafman, I., Domshlack, C., & Shimony, S. (2006). On graphical modeling of preferences and importance. *Journal of Artificial Intelligence Research*, 25:389-424.
- [Branavan et al. 2012] Branavan, S., Silver, D., & Barzilay, R. (2012). Learning to Win by Reading Manuals in a Monte-Carlo Framework. *Journal of AI Research*, 43:661-704.
- [Falkenhauer & Forbus, 1991] Falkenhauer, B. and Forbus, K. (1991). Compositional modeling: Finding the right model for the job. *Artificial Intelligence*, 51, 95-143.
- [Forbus & Hinrichs, 2017] Forbus, K.D. & Hinrichs, T. (2017) Analogy and Relational Representations in the Companion Cognitive Architecture. *AI Magazine*.
- [Forbus & Hinrichs, 2018] Forbus, K. & Hinrichs, T. (2018) Qualitative Reasoning for Decision-Making: A Preliminary Report. *Proceedings of QR 2018*, Stockholm.
- [Goel & Rugaber, 2015] Goel, A., & Rugaber, S. (2015). Interactive Meta-Reasoning: Towards a CAD-Like Environment for Designing Game-Playing Agents. In Besold, T., Schorlemmer, M., & Smaill, A. (Eds.) *Computational Creativity Research: Towards Creative Machines*. Atlantis Thinking Machines, Vol. 7, Atlantis Press, Paris.
- [Hancock et al. 2018] Hancock, W., Forbus, K., & Hinrichs, T. (2018). Optimizing Strategy Game Planning via Combining Analogical Learning with Function Learning In *Proceedings of the 6th Annual Conference on Advances in Cognitive Systems*. Stanford.
- [Hinrichs & Forbus, 2012a] Hinrichs, T. and Forbus, K. (2012). Learning Qualitative Models by Demonstration. *Proceedings of the 26th AAAI Conference on Artificial Intelligence*, 207-213, July.
- [Hinrichs & Forbus, 2012b] Hinrichs, T. and Forbus, K. (2012). Toward Higher-Order Qualitative Representations. *Proceedings of the Twenty-sixth International Workshop on Qualitative Reasoning*. Los Angeles CA, July.
- [Hinrichs & Forbus, 2015] Hinrichs, T., and Forbus, K. (2015). Qualitative models for strategic planning. *Proceedings of the 3rd Annual Conference on Advances in Cognitive Systems*.
- [Hinrichs & Forbus, 2019] Hinrichs, T., & Forbus, K. (2019). How Qualitative Models and Improve Learning by Experimentation. *Proceedings of QR 2019*.
- [Hwang & Yoon, 1981] Hwang, C. & Yoon, K. (1981). Multiple criteria decision-making. *Lecture Notes in Economics and Mathematical Systems*, 186:58-191.
- [McFate et al. 2014] McFate, C.J., Forbus, K. and Hinrichs, T. (2014). Using Narrative Function to Extract Qualitative Information from Natural Language Texts. *Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence*, Québec City, Québec, Canada.
- [Rossi et al. 2011] Rossi, F., Venable, K., & Walsh, T. (2011). *A Short Introduction to Preferences: Between Artificial Intelligence and Social Choice*. Synthesis Lectures on Artificial Intelligence and Machine Learning, Morgan & Claypool.
- [Rovira et al. 2018] Rovira, X., Riu, D., Giones, F., Agell, N., Casabayo, M. (2018). Helping Investors to Assess New Ventures Proposals by a Linguistic Multi-Criteria Decision Tool. *Proceedings of QR 2018*, Stockholm.
- [Santhanam et al. 2016] Santhanam, G., Basu, S., & Honavar, V. (2016). *Representing and Reasoning with Qualitative Preferences*, Synthesis Lectures on Artificial Intelligence and Machine Learning, Morgan & Claypool.