Resource Models in Qualitative Decision-Making: A Proposal

Kenneth D. Forbus & Thomas R. Hinrichs
Qualitative Reasoning Group, Northwestern University
forbus@northwestern.edu

Abstract

Qualitative decision-making aims to make the formulation of decision problems involving continuous aspects of the world automatic. This paper proposes the idea of *qualitative resource models*, which can be used to rule out potential options and provide ranking information about alternatives from the perspectives of the resources they require. Two kinds of scenarios are used to illustrate these ideas, defense in a strategy game and figuring out what to do about dinner.

1 Introduction

Decision-making is one of the key problems that intelligent agents face, whether operating autonomously or in a supporting role with human collaborators. In most decision-making research, the focus is on how to make decisions, often focusing on optimality. In qualitative decision-making (Forbus & Hinrichs 2018; 2019), the goal is to use ideas from qualitative reasoning to provide representations that enable systems to formulate their own decision problems, when there are continuous aspects to be considered. Arguably any AI system capable of generating multiple plans to achieve a goal can be said to make decisions, but how should the alternatives in those decisions be evaluated? Often they involve factors that can be considered as continuous, such as costs in time, money, or other material. Moreover, the situations that agents are faced with are often only partially specified, especially when planning for the future. In addition, exact models and data are rarely available in everyday circumstances. These reasons suggest that qualitative representations of continuous aspects of decision-making can potentially provide a valuable service in reasoning.

This paper extends our prior work on qualitative decision-making with the idea of *qualitative resource models*. Every alternative considered in a decision involves some form of costs, in terms of resources. Qualitative resource models make such costs explicit and represent causal relationships among the parameters of the decision problem, in order to support reasoning about alternatives with minimal information. We begin by describing qualitative resource models, including how they provide a layer over events, actions, and other conceptual structures used to describe alternatives. This includes both continuous resources and discrete resources. The role of investments, i.e. activities undertaken to improve the set of available resources, is also discussed. Next we illustrate resource models with two extended examples. The first consists of deciding how to defend a city under attack in a strategy game. The second consists of deciding what to do about dinner, an everyday decision which has become more fraught during the pandemic.

2 Qualitative Resource Models

A resource is something that is used to carry out an action or event A. We follow the standard definitions of a consumable resource being something that is used up and a durable resource being something needed in an event or action but persists after its use. In baking, for example, the flour that went into a recipe is a consumable resource and the oven used is a durable resource. The costs of A in consumable resources consist of a set of quantities, one per consumable resource. We use the conventions of (Forbus & Hinrichs 2019) for representing costs. For example, a representation of the typical costs of dinner at a fancy restaurant can be described as

```
(valueOf-Type (CostFn Money)
              DinnerNiceRestaurant
              (Dollar 50 100))
(valueOf-Type (CostFn Time)
              DinnerNiceRestaurant
              (Hour 2 6))
```

That is, events that are instances of the concept of having dinner at a nice restaurant cost between $50-$100 and they last between 2-6 hours. (We use the OpenCyc conventions[1] that units, here `Dollar` and `Hour`, are represented via logical functions, which with two arguments indicates an interval and with one argument indicates a specific value.)

[1] The NextKB knowledge base uses the OpenCyc ontology, and its documentation provides an introduction to its conventions and tools for using it: https://www.qrg.northwestern.edu/nextkb/index.html

Type-level qualitative models (Hinrichs & Forbus, 2012) are useful for reasoning abstractly, and enable the expression of general causal models, e.g.

```
(qprop+TypeType (CostFn Money)
                (CostFn Time)
   DinnerNiceRestaurant DinnerNiceRestaurant
   same)
```

That is, the longer a dinner at a nice restaurant takes, the more expensive it is likely to be.

Type-level models of costs can be derived from statistics gleaned from particular instances of events (Hancock et al. 2018). The instance-level relation connecting a specific fluent for a quantity of an event to its value is `valueOf`. For example, a specific dinner that has already occurred, D1, might lead to recorded costs such as

```
(valueOf ((CostFn Money) D1)(Dollar 75))
(valueOf ((CostFn Time) D1)(Hour 3))
```

As usual in qualitative reasoning, we often will not know specific values. There are several notions of value, in addition to intervals as introduced so far, that are useful in qualitative decision-making. The first are ordinal relations. We might not know the specific cost of dinner D2, but know only

```
(greaterThan ((CostFn Money) D2)
             ((CostFn Money) D1))
```

The second kind of qualitative value are stratified values. Order of magnitude representations (e.g. Agell et al. 2006) are one example of stratified values. Another are common named ranges. Some of these are defined in terms of specific intervals, e.g. `TensOfMinuteDuration`, `SeveralHoursDuration` from OpenCyc. Others provide relative values. For example, there is no specific interval associated with either of the OpenCyc values of `Inexpensive` and `Expensive`, but their relative ordering is known.

Durable resources, as noted above, are resources used by an alternative but not consumed during that usage. Space is an example of a durable resource. For cooking a meal, for example, the burners on a stove are a commonly used durable resource whose availability shapes events. The AI planning literature addresses the assignment of durable resources. They are of interest here mainly because investment decisions often involve adding durable resources, e.g. buying a new piece of kitchen equipment.

Resources are a bridge between constructs used to design behavior and those actions and events used to implement said behaviors. Possible behaviors, which form alternatives to decide among, are constructed by planning. For example, one might consider eating dinner at a food truck versus a nice restaurant. For any alternative A, we assume a set of events {E} which constitute A and whose costs define the cost of A. Unsurprisingly, money costs are additive, e.g.

```
(c+TypeType (CostFn Money) (CostFn Money)
   DecisionAlternative Event
   subEvent)
```

(`c+TypeType` is the type-level version of `c+`, which in QP theory provides a term for the sum that constrains a quantity.) Time costs can be additive as well, if the subevents are strictly sequential. We will assume this to be the case here for simplicity.

2.1 Decisions and alternatives

We consider a decision to consist of a set of alternatives, one of which needs to be chosen in order to make that decision. Decisions can have associated constraints indicating requirements for success. Deadlines are a common form of constraint, e.g. money must arrive in a bank account before it can be spent. Alternatives for a decision that violate a deadline may be ruled out, e.g. if an expense will be billed tomorrow, then mailing a check to a bank in another country is unlikely to be an effective way to deal with the situation. Stratified values can often rule out alternatives, e.g. something mailed typically takes several days to arrive, which is longer than one day.

Traditional decision-making often involves optimization, seeking the best alternative. The use of qualitative representations for costs means that relative costs can sometimes, but not always, be determined. This degree of ambiguity means that qualitative decision-making will not always be able to identify an optimal solution, but it does provide the factors that need to be considered if a more quantitative model needs to be constructed.

2.2 Investment Decisions

Investment decisions involve actions or events aimed at improving an agent's resources, so that its other activities are more efficient and/or effective. In strategy games, for example, building installations or units is a form of investment, as is researching new technologies, which in turn expand the range of capabilities an agent has. We organize the set of capabilities around the idea of *functional subsystems* (Forbus & Hinrichs, 2019), e.g. a household might have a living space system consisting of the rooms that the participants live in, a financial system consisting of the finances of its participants, and a cooking system, the equipment available for food preparation. Deciding whether or not to buy a microwave oven, for example, may involve examining the typical activities conducted in the household to estimate how much that investment might improve them.

Shared resources means that decisions can interact. The decision to purchase a microwave oven, for example, means those funds are not available to repair the bathroom. Resource models provide a common perspective that identifies such potential interactions between decisions, due to shared influences on cost quantities.

2.3 Modeling an Agent's Activities

Understanding what might make a good investment requires knowing about the intended mix of future activities. How can the mix of activities of an agent or an organization to be modeled? First, we need conventions for representing activities. AI architectures have used several constructs for this. The most common are plan representations, e.g. Hierarchical Task Network (HTN) tasks (Georgievski & Aiello, 2015). These ground out in primitive actions, each of which has one or more associated type of event associated with its execution. Thus the resource model for an instance of an HTN task network reduces to the resource model for the constituent events associated with their primitive actions. For example, in NextKB, the Freeciv primitive action `do-Move` is understood to be representable in terms of instances of the `Movement-TranslationEvent` concept. High-level procedural representations, e.g. (Morley & Meyers, 2004) ground out similarly. We assume an agent records instances of task/procedure executions, including in those records the events that they caused.

We propose that the SAGE model of analogical generalization (McLure et al. 2015) can be used to learn models of the mix of an agent's activities from experience. SAGE associates a *generalization pool* with each concept to be learned. These generalization pools incrementally accumulate examples, constructing and maintaining a set of generalizations and outliers representing its model of that concept. The accumulation process uses analogical retrieval (Forbus et al. 1995) to find the closest prior example or generalization when a new example is added. If the new example matches something sufficiently similar (as measured via the Structure-Mapping Engine (SME; Forbus et al. 2017)), they are assimilated. That is, if the retrieved item is an example, the structural alignment SME computes is used to initialize the generalization, with aligned facts having probability 1.0 and non-aligned facts having probability 0.5. If the retrieved item is a generalization, the probabilities associated with the aligned facts are updated accordingly. Thus, each generalization is a probabilistic relational schema, where the probability of each statement is the frequency with which something matching it appeared in that very similar set of examples. A generalization pool can have multiple generalizations, thereby providing a means of modeling disjunctive concepts. Examples not added to any generalization constitute outliers, at least relative to the system's experience at that point.

Activities involve functional subsystems, e.g. finding dinner might be solved via the activity of cooking a meal, which uses the kitchen system. Similarly, defending a city might be solved by adding military units to it, which uses the military system. Thus the sets of activities an agent is involved in provides information as to how its systems are used. SAGE generalizations track the number of instances that have been assimilated into them, and by analyzing these statistics across a generalization pool, the relative usage of systems can be estimated. This provides information that can be used to prioritize investments, since frequently used systems are more likely to be worth improving. This is especially true if activities involving them have been less than successful – again, something that can be tracked via SAGE's probabilities, assuming the success and failure of particular activities is recorded.

To explore how these ideas might play out in decision-making, we next examine two extended examples, one from a strategy game, and one from everyday life.

[2] We are ignoring the maintenance costs of units here, which consume production points from their home city.

3 Defending a City in Freeciv

Strategy games like Freeciv, an open-source version of Civilization 2, provide useful testbeds for exploring decision-making for several reasons. First, they involve multiple interacting factors, thereby forcing players to consider tradeoffs, e.g. the classic guns/butter tradeoff in economic improvement versus military strength. Second, they are constructive, requiring players to build cities, improve terrain, and do research to improve capabilities. Thus they provide a good testbed for investment decisions as well.

Here we examine a problem that comes up routinely in such games, namely defending a city from enemy attacks. The tactic for handling this problem is called `DefendingA-Position` in our system. Invoking this tactic introduces a new decision, namely which of several possible alternative sub-tactics should be used to address it. For `DefendingA-Position`, these include:

1. `DefendByReinforcement`: Move a military unit into the city.
2. `DefendByBuilding`: Use the city's production capacity to create a new military unit. This can take a few turns or many, depending on what is being built.
3. `DefendByBuying`: By paying gold, production can be sped up to a single turn.

Each of these sub-tactics are implemented in turn via taking actions, which if taken will give rise to events, each of which should have statements contributing to qualitative resource models. Let us examine what these models should be. Since `DefendByReinforcement` is reassigning an existing unit, its only cost is movement time[2]:

```
(qprop+TypeType (CostFn Time) (DistanceFn path)
        Movement Movement same)
```

That is, the time cost depends on the distance for the path involved in moving the unit into the city. Different units move at different rates, so a more complete model must take this factor into account as well. For `DefendByBuilding`, the time cost is the number of turns required to build the object, a parameter that can be measured directly from the underlying simulation:

```
(qprop+TypeType (CostFn Time)
      (MeasurableQuantityFn numTurnsToBuild)
      CreationEvent FreeCiv-City doneby)
```

The `DefendByBuying` tactic trades money for time:

```
(qprop+TypeType (CostFn Money)
      (MeasurableQuantityFn numTurnsToBuild)
      Buying FreeCiv-City doneBy)
```

That is, the further an improvement is from completion, the more money must be spent to finish it.

The circumstances in which `DefendingAPosition` is invoked can lead to additional constraints. Some strategies pre-position defending units in cities even in peacetime. This represents both an opportunity cost, i.e. not building economic improvements, and depending on the type of government the civilization is under, a decrease in citizen happiness. Other strategies only add defenders when a war breaks out, or when an attacker is detected on their way to a city. For tactics where the cost is time, that time must be smaller than the arrival time of the attackers. For `DefendByBuying`, the danger is bankruptcy, where improvements and military units are sold off to pay debts. Thus if multiple cities are under attack, the tradeoffs between these strategies must be considered carefully.

4 Investments in Cooking

Households in modern societies have several categories of activities they can use to have dinner. Abstractly, these include eating at a restaurant, getting takeaway food, and cooking. Each of these broad solutions have variations, such as dining at a nice restaurant versus a food truck, or shopping for high-quality ingredients and preparing multiple dishes, versus microwaving pre-packaged food. The relative frequency of these activities will vary from household to household, of course, and within the same household over time. We assume that, in the household participants' experiences, they have robust generalizations for both the broad categories of activities as well as the more specific subcategories (e.g. picking up BBQ at Hecky's). A high-income household that prioritizes work activities might have a heavily refined set of models for restaurant dining and takeaway, with relatively sparse models for cooking. On the other hand, a household that prioritizes saving money probably will have robust models for different solutions involving cooking, and relatively fewer restaurant and takeaway models.

The Covid-19 pandemic provides examples of how people's routines were upended, forcing them to come up with other ways to, among other things, find dinner. During lockdowns, restaurant dining was no longer an option, and some extreme lockdowns in some countries even forbade takeaway service. Whatever preferences households had before, with lockdowns, cooking (and in many places, takeaway) service suddenly gained higher priority. This change in priority, we assume, provides a signal to consider investments that could improve the efficiency and/or effectiveness of the newly prioritized activities. Thus households that ignored or neglected their kitchen systems previously found themselves considering how they might improve them. There is evidence that this is the case. In the US, for example, sales of bread makers rose by over 600%, electric pasta maker sales rose by 462%, soda machines sales rose by 283%, and deep freezers sales rose by 45% in 2020[3]. This was, of course, for high-income households, which also had the option of simply relying on far more takeaway, leading to massive increases in orders through food delivery services.

One factor in these investment decisions is space. Kitchen appliances must be stored somewhere when not in use, and take up counter space in use. Deep freezers are typically relegated to basements in houses, and there is rarely space for them in apartments. The surge in home-buying during the US last year is commonly attributed to the perceived need to radically expand the spatial resources available for household systems, due to working from home.

Building out resource models to capture investment decisions such as these will need more refined models of preferences, in order to evaluate benefits as well as costs. How much is it worth spending to make a regularly cooked dish produced more efficiently, or expand the range of what can be prepared beyond current routines? Prior work on qualitative models in representing preferences and decision-making, outlined next, might be productively linked to qualitative resource models to tackle this.

5 Related Work

Several previous lines of research in qualitative reasoning on decision-making have inspired this research. Rovira et al. (2018)'s use of linguistic term sets to map from language to values provides a promising approach for more flexible expressions of costs and other decision factors. Rossi et al.'s (2011) and Santhanam et al.'s (2016) work on preferences shows that qualitative values can indeed be valuable in decision-making applications, including Benaroch & Dhar's (1995) work on investment decisions.

Freeciv has been used by several researchers previously. Branavan et al. (2012) used reinforcement learning guided by NLP over a manual to learn to play Freeciv, but required the game engine to be used for lookahead, something which is not available in most decision-making situations. Goel & Rugaber (2015) examined how meta-reasoning could be used in designing game agents, whereas our concern is how to build autonomous agents that make their own decisions.

6 Conclusions and Future Work

This paper proposes qualitative resource models to provide a new kind of linkage between qualitative representations and other forms of knowledge. That is, qualitative resource models reflect continuous properties of events, in terms of dimensions of costs and the factors that drive them. Resource models appear to be promising as part of the process of formulating decision problems. In some cases, resource models should suffice to prune infeasible alternatives and ascertain which alternatives are best, based on partial information. In other cases, more quantitative data might be needed, in which case the qualitative resource model helps identify what data is needed. The use of analogical generalization to construct models of an agent's activities also seems promising.

We see two next steps in future work. The first is to finish implementing qualitative resource models in the Companion cognitive architecture, and test them for decision-making in Freeciv. The second is to formalize benefits, in terms of both

[3] Stockpiling Germaphobs Ignite Unlikely Boom: Appliances. Bloomberg News, May 13 2020, [Deep Freezers, Bread Makers Sell Out in Coronavirus Spending Boom - Bloomberg]

reduction of costs and in terms of positive evaluative quantities (e.g. improving the strength of a system), to be able to reason better about opportunity costs and investment decisions.

Acknowledgments

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

References

[Agell et al. 2006] Agell, N., Rovira, X., Sanchez, M., and Prats, F. (2006). Using Orders of Magnitude in Multi-Attribute Decision-Making. *Proceedings of QR 2006*.

[Benaroch & Dhar, 1995] Benaroch, M., & Dhar, V. (1995). Controlling the complexity of investment decisions using qualitative reasoning techniques. *Decision Support Systems*, 15(2):115-131.

[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.

[Forbus, 1984] Forbus, K. (1984). Qualitative process theory. *Artificial Intelligence*, 24, 85-168.

[Forbus & Hinrichs, 2017] Forbus, K. & 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.

[Forbus & Hinrichs, 2019] Forbus, K., & Hinrichs, T. (2019). Qualitative Reasoning about Investment Decisions. Proceedings of the 32nd International Workshop on Qualitative Reasoning (QR 2019). Macao, China.

[Forbus et al. 1995.] Forbus, K., Gentner, D. & Law, K. (1995). MAC/FAC: A model of similarity-based retrieval. *Cognitive Science*, 19:141-205.

[Forbus et al. 2017] Forbus, K., Ferguson, R., Lovett, A., and Gentner, D. (2017). Extending SME to handle large-scale cognitive modeling. Cognitive Science, DOI: 10.1111/cogs.12377, pp 1-50

[Georgievski & Aiello, 2015] Georgievski, I., & Aiello, M. (2015). HTN Planning: Overview, comparison, and beyond. *Artificial Intelligence*, 222:124-156.

[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, 2012] 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.*

[McLure et al. 2015] McLure, M., Friedman, S., & Forbus, K. (2015) Extending Analogical Generalization with Near-Misses. *Proc. AAAI 2015*.

[Morley & Meyers, 2004] Morley, D. & Meyers, K. (2004). The SPARK agent framework. *Proc. AAMAS'04*, pp. 714-721.

[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.