

Qualitative Modeling of Social Relationships

Kenneth D. Forbus

Qualitative Reasoning Group, Northwestern University
2233 Tech Drive, Evanston, IL, 60208 USA
forbus@northwestern.edu

Abstract

Social reasoning is a key capability in human cognition. Formalizing social reasoning can both improve our understanding of human cognition and support building AI systems that can perform it. The advantages of qualitative representations, such as abstraction of numerical values and compositional causal relationships between quantities, become especially important in domains where basic properties to formulate mathematical models are missing. However, social reasoning provides new challenges for qualitative reasoning, since, like many everyday reasoning problems, it involves fluently moving between discrete representations of events/actions and continuous causal models. This paper explores the hypothesis that the continuous aspects of social reasoning can be effectively modeled in qualitative process theory plus two extensions. These extensions are (1) incorporating discrete changes in the language of influences and (2) modeling aspects of episodic memory via sets of cases representing experience. We illustrate these ideas by formalizing aspects of social relationships involved in friendship.

Introduction

Social life is complicated. So much so that sociality is increasingly viewed as a driver of the evolution of intelligence (e.g. Tomasello, 2001). Thus understanding how social reasoning works can better help understand human cognition, as well as providing part of the foundation for creating AI systems that can understand our social world and even someday participate effectively in it.

We seem to think about some aspects of social life in continuous terms. For example, we can talk about events bringing us closer to someone, and gauge which of our friends is more likely to be relied on in a tough situation. However, as with many everyday phenomena, the level of precision in available information is a mismatch with the requirements of traditional mathematical modeling methods. To gather numerical data requires having some notion of units, for example – when we speak of two people being close, how would we quantify that as a number? There have been attempts to model social relationships mathematically, but there is little quantitative data upon which to base such models, nor constraints on their internal parameters. Thus formalisms developed in the qualitative reasoning community offer a way to build models that are closer to what kinds of information are actually available about the phenomena.

However, social reasoning raises interesting challenges for qualitative reasoning. It requires drawing conclusions from experiences in the everyday world, which means incorporating rich representations of events, their participants, and relationships among them. It requires shifting between continuous and discrete models of actions, as when one takes a series of discrete interactions in the aggregate to approximate derivatives over intervals. And it requires estimating properties of interactions over experiences, hence addressing properties of episodic memory.

The rest of this paper describes an initial attempt to model the continuous aspects of social relationships qualitatively, using qualitative process theory (Forbus, 1984;2019) with two extensions. We begin by providing some background both about social psychology and some ontological assumptions we build upon for handling events, finite symbol quantity values, sets, and cases. Then we describe the two extensions to QP theory. The first is to re-purpose a discrete notion of influence due to Kim (1993) to handle causal reasoning about the occurrence of events on continuous parameters. The second is a very simple formalization of episodic memory and how to connect changes in events to changes in qualitative values. Next, we describe a basic encoding of social relationships, focusing on continuous properties. An extended example shows how these ideas can be used together to do at least one aspect of social reasoning. We close with conclusions and future work.

Background

We begin with the aspects of social psychology we are drawing upon, then the ontological assumptions in addition to QP theory that we need to make.

Some Social Psychology

Given the complexity of human social life, it is perhaps unsurprising that social psychology has not settled on a single, universally agreed upon theory of social relationships, let alone a formal version of such a theory. There are multiple frameworks with varying degrees of adoption. For example, Fisk (1992, 2004) argues that social relationships can be broken down into four categories. The first, communal sharing, focuses on what people have in common. This includes being members of the same family, workplace, club, and so on. The second, authority ranking, describes interactions in terms of ordered differences, such as seniority, age, gender,

or caste. The third, equality matching, uses balances of contributions, such as turn-taking and exchanges of favors. The fourth, market pricing, uses money or some implicit continuous parameter to evaluate interactions. This includes rents, dividends, interest rates, and evaluating relative benefits of a relationship to those involved. Each of these categories is described as modes of interactions, “mods” in Fisk (2004), which can be combined with culturally specific prototypes and patterns, “preos” in Fisk (2004), to describe the practices of a group. In compositional modeling terms, mods and preos are analogous to model fragments, with preos modifying mods, such that a situation model composed of such model fragment instances would be a model of how that cultural group operates. While examining this analogy more closely could be productive, we focus here on the effects of interactions on individuals and their relationships, rather than constraining what interactions people will have, which is more in the territory of social norms (Malle et al. 2019; Olson & Forbus, 2021).

While Fisk’s account focuses on providing a mechanism, other social psychology models such as Kelley et al. (2003) focus more on cataloging the phenomena. Kelley et al. argue that the construct of situation is central in social psychology, because the situations that social beings find themselves in are major factors in determining what they do. Thus classifying types of situations serves the purpose of carving the phenomena up into units amenable for analysis. Kelly et al. (2003) argues that recognizing such situations is an important force in our evolution. Formalizing these situations, again, would be an interesting enterprise, but would take us far beyond what a QR focus can provide.

To provide an initial focus for modeling, we build on the account of friendship due to Rawlins (1992). Rawlins’ account is informal and descriptive, not mechanistic. He argues that friendship is a kind of social relationship that is (1) voluntary, (2) negotiated by both parties involved, (3) provide some sense of equality for the parties involved, (4) require mutual involvement, and (5) have an affective component¹. The last three characteristics are promising for qualitative modeling because they seem to involve continuous factors. For example, one factor in equality is that the needs and desires of both friends are important, making large or long-duration imbalances something important to detect. Similarly, mutual involvement is measured by both parties being willing to spend time together at an appropriate amount and frequency. Finally, pairs of friends are closer to each other than people who are not friends. They know each other’s histories and build up considerable shared history. The notion of closeness is analyzed below, since it seems central to friendship and social relations.

Rawlins also proposed a seven-stage model of the trajectory of friendship, encompassing both its growth and decline. The first stage are *role-delimited interactions*, e.g. the relationships you have with other people when you are shopping, driving, mentoring, etc. Within such roles, the next step involves what Rawlins calls *friendly relations*, where there is more mutual disclosure beyond what the roles require. Then comes the initiation of interactions outside the roles, a stage Rawlins calls *moves-toward-friendship*. Then comes *nascent friendship*, where interactions are no longer following role stereotypes at all, and norms for what and

how to communicate are established for the relationship. Some topics might be declared out of bounds, (e.g. religion or politics), and continued compliance with those mutually agreed-upon norms is part of the process of deepening trust. Interactions with others start to take the friendship into account as well. At some point nascent friendships become *stabilized friendships*, where interaction patterns and norms are stable and mutually agreeable. Rawlins observed that stabilized friendships fall into three types: *active friendships* involve regular mutual interactions currently, versus *dormant friendships* where mutual interactions have tapered off, although it could be quickly restarted by interacting again. Finally, *commemorative friendships* are those where the bulk of the interactions were in the past, with only minimal current interactions. The last two stages that Rawlins identified are involved in the dissolution of friendship. A *waning friendship* is one that starts to decline in its importance in our lives. The sources of waning might be a reduction in closeness, due to one or both parties not investing enough in events that sustain friendship, or negative events like betrayals. Finally, in the post-friendship phase, the friendship continues to provide memories that influence future relationships, e.g. great activities to have done, experiences best avoided.

Being informal, there are many open questions in Rawlins’ account. For example, how can one tell a post-friendship from a dormant stabilized friendship? While model fragments might be used to encode some aspects of these various stages, the criteria that should be used to define limit points to transition between them is far from clear. Nevertheless, building up formal qualitative models of such theories might both provide better social reasoning for AI systems, as well as perhaps helping produce more formal but still qualitative social psychology theories. Such formal qualitative models would enable the generation of testable predictions, while at the same time being better suited to the kinds of evidence and data available, compared to traditional mathematical models. This is not unprecedented, as the work of Bredeweg et al. (2008) in ecology and de Jong (2008) in genetic regulatory networks illustrate. Here we will start small and build up a simple model of social relationships that expresses some continuous aspects of friendship.

Ontological Assumptions

For the aspects of this model that lie outside QP theory, we freely draw upon the OpenCyc ontology, as used in NextKB². OpenCyc is a subset of Cycorp’s Cyc KB (Lenat et al. 1990) contents that is freely available³. It provides a broad commonsense ontology which can be linked to QP theory constructs and is grounded in natural language (Forbus, 2023), making it useful for describing the open-ended nature of events and relationships in the world. For example, there are 16,715 subcategories of Event, and 2,643 distinct role relationships that express their properties. This broad vocabulary is useful given the nature of human social life.

The NextKB ontology provides support for Hayes’ (1985) notion of histories, where change is represented in terms of bounded pieces of space/time whose properties

¹ This summary draws upon Wrench et al. (2023)’s account of Rawlins’ work.

² <https://www.qrg.northwestern.edu/nextkb/index.html>

³ Creative Commons Attribution license.

vary along those axes. In OpenCyc, the kinds of things that can have histories are instances of `SpatialThing-Localized`, which inherits from `TemporalThing` and `SpatialThing`. The predicate `holdsIn` provides a modal operator that specifies that, over a given temporal extent denoted by a `TemporalThing`, a given proposition holds. This provides a means of specifying what properties are true during an event, for example. The logical function `AtFn` is used to denote the spatiotemporal slice of an entity during a subset of a history, e.g. the speed of a falling object is higher at a point later in its trajectory than at its start.

The Cyc ontology supports microtheories (Guha 1992), a form of context. All reasoning is performed with respect to some microtheory and those microtheories it inherits from (via the `genlMt` relation). Qualitative states and models of the contents of the minds of others are both implemented via microtheories, for instance.

NextKB supports the traditional QR notion of a quantity as a fluent, which takes different values at different times. In addition to ordinal and signs, NextKB inherits a well worked out ontology of symbolic values and numerical properties from OpenCyc. Properties such as `Happiness` are ontologized so that instances of them are values, either symbolic or numerical. The symbolic values include the kind of finite symbol algebra commonly found in QR. Logical functions (e.g. `HighAmountFn`) are used to provide a general approach to specifying such values. OpenCyc also provides a rich collection of units and conversions between them. Units are represented by logical functions, so that non-atomic terms like `(HoursDuration 3)` bundle units with values.

The preferences of an agent can be expressed via `prefers`, a ternary predicate taking an agent and two sentences, meaning that the agent prefers situations in which the first sentence is true over those in which the second sentence is true.

NextKB supports both intensional and extensional representations of sets. An operator is provided for gathering the bindings satisfied by a conjunction of statements (`the-ClosedRetrievalSetOf`) evaluated over an extensionally specified set. We will use this operation in modeling operations over episodic memory below.

Extensions to QP Theory for Social Reasoning

QR has mostly focused on continuous processes, but there have been interesting exceptions. Simmons (1983) described a notion of discrete process to handle reasoning about geological processes, which happen too slowly to be directly observed (earthquakes and volcanos excepted), but whose occurrences over historical time are important to understand. His representation of change used operations on a diagrammatic representation of layers under the Earth. The representations below are typically also considered discrete events, and while a more physically grounded event representation would include a spatial/diagrammatic component, we do not do this here. Integration of QP theory with discrete planning has been done via STRIPS operators in environments (Forbus, 1989), compiling processes into operators for a temporal planner (Hogge, 1987), and more tightly integrated with plan operators (Drabble, 1993). The extension here to handle discrete effects in events is closest to Hogge (1987), but the representations are used differently in

reasoning. The second extension concerns modeling episodic memory. Some prior work has explored the use of QR in the representation of episodic memories (Hancock & Forbus, 2021), but not in constructing a formal model of episodic memory per se.

Handling Discrete Effects in Events

Consider a pleasant outing undertaken by two friends, Pat and Kit. They walked through the woods, picnicked in a clearing, and swam in a lake. Each of these constituent events of the outing could be decomposed further into sub-events as needed. Some of those sub-events can, in turn, be viewed as including occurrences of continuous processes, such as walking and swimming. In other cases, the sub-events may best be viewed as discrete events. For example, the picnic consists of particular events of eating, drinking, and conversing, bookended by setting up the picnic and cleaning up afterwards. But some of these events can be construed in terms of continuous processes. Thus the event of drinking a glass of wine can be decomposed if needed into a set of movements of the glass/liquid combination, the pouring of the liquid, and so on. To infer the causal import of events, we need to combine continuous and discrete models of effects across multiple levels of events. For example, given the wine consumption at the picnic, are either Pat or Kit in shape to go swimming? Answering this question does not require fine-grained decomposition of constituent events, only knowing that wine was consumed, and some means of estimating how much. This is an example of a key problem in commonsense reasoning: determining how to compute effects of events and processes across levels of description without getting bogged down in irrelevant details.

For continuous processes, we use the qualitative mathematics of influences from QP theory. That is, qualitative proportionalities (`qprop+`, `qprop-`) provide representations of partial information about algebraic causal connections (e.g. (`qprop+ (level (ContainedLiquid wine glass)) (mass (ContainedLiquid wine glass))`)). Direct influences (`i+`, `i-`) represent partial information about integral causality (e.g. (`i- (energy Kit) (rateFn (Walking Kit))`)). For discrete changes, we adopt H. Kim's (1993) extensions originally developed to encode abrupt changes:

`(increase <qty>)` indicates that `<qty>` increases over the interval of interest

`(decrease <qty>)` indicates that `<qty>` decreases over the interval of interest

`(increaseBy <qty> <amt>)` indicates that `<qty>` increases by `<amt>` over the interval of interest

`(decreaseBy <qty> <amt>)` indicates that `<qty>` decreases by `<amt>` over the interval of interest.

Unlike `i+`, `i-`, these relationships make no specification as to the derivative of `<qty>` at any particular sub-interval for the interval of instance, they only concern the net effect across the interval in question. In Kim (1993) this was used, for instance, to model the effects of combustion in a four-cycle engine, which is for some purposes modeled as an impulse. The same ambiguity regarding when within an interval that a change happens is used here for intervals covering substantial intervals of time, e.g. a picnic.

Returning to our picnic example, `(increaseBy (Wine-ConsumedFn Kit) (GlassesFn 3))` states that, in the event for which this statement appears, Kit consumed three glasses of wine. If `P1` denotes the picnic, then

```
(holdsIn P1
  (increaseBy (WineConsumedFn Kit)
    (GlassesFn 3)))
```

Such a conclusion might be reached by combining the amount of wine Kit drank across the entire picnic, i.e.

```
(evaluate ?n-drinks
  (TheClosedRetrievalSetOf ?drinks
    (and (occursDuring ?sub-e P1)
      (isa ?sub-e DrinkingEvent)
      (doneBy ?sub-e Kit)
      (substanceConsumed ?sub-e Wine)
      (amountConsumed ?sub-e (GlassesFn ?n))
      (unifies ?drinks (GlassesFn ?n)))))
```

The extraction of the total number of glasses consumed is straightforward.

So far, we have focused on the physical aspects of Kit and Pat's outing. But what are the social effects? Again, we only have finite qualitative symbol systems and ordinal relationships to express preference information. This means results will often be ambiguous, but that is the nature of qualitative reasoning. A social reasoner must evaluate as best it can the effects on the participants in terms of whether or not it is a positive, negative, or mixed experience. This requires taking their preferences into account. We do not assume perfect information. Each social reasoner must do the best that it can with the information it has.

There are four representations of preferences to consider between two people: Their own beliefs about what each of them prefers, and their own beliefs about what the other prefers⁴. These can be kept distinct via microtheories, e.g. (PrefersBeliefsOfFn Kit Pat) is the set of preferences that Pat believes that Kit has. There are multiple dimensions that could be compared for evaluative purposes, e.g. fiscal requirements, physical stamina, etc. and more likely this set is simply a subset of the larger set of beliefs that Pat has about Kit. We stick with preferences here for simplicity.

Reasoning about Change in Sets of Events

Our intuitions about friendship, as well as Rawlin's model, tell us that our feelings are determined in part by our recollections of shared experiences with our friend. Thus we need to have a way to model the relevant types of memories and how they change over time. We are focused here on episodic memories, specifically memories of events that an agent has participated in with other people, making them grist for building/evaluating that social agent's relationship with that other (or others). We will represent each memory via a case (implemented as a microtheory) whose contents are one or more occurrences of events. For example, two friends having a dinner out can be described in terms of such a microtheory.

People's memories are personal, subjective and noisy. The same event (or network of interlocking events, without loss of generality) might be remembered in very different ways by the participants⁵. We use the relationship (episodicMemoryOf <mt> <Person>) to indicate that the microtheory <mt> of events is part of <Person>'s episodic memories.

For each Person, there is a microtheory denoted by (MemoriesOfFn <Person>) whose contents are that person's memories. Thus episodicMemoryOf statements in that person's MemoriesOfFn microtheory indicate that that person does indeed have that memory. episodicMemoryOf statements in someone else's MemoriesOfFn microtheory represent someone else's belief as to what that person remembers about an experience. We assume that there are memories beyond episodic memories in MemoriesOfFn. For our purposes, we include other narratives (e.g. stories that a person has understood from conversation, reading, and watching), but not semantic memory or skill memory, because these are not used in similar ways in social reasoning. The episodic memories of an agent (i.e. the extension of episodicMemoryOf statements) will be denoted using EpisodicMemoryFn.

As people gain experience, and as their knowledge about their experiences changes (learning more about a party, forgetting a slight that you perceived), this can lead to changes in social relationships. Such changes are not tied to the underlying parameter changes as directly as the normal qualitative mathematics of influences, nor are they continuous integrations anymore. The notion of derivative across intervals of time representing occurrences of discrete events is still useful even though it is more granular. We define the usual signs of derivatives in terms of ordinal relationships over a quantity being tracked that is affected by an event. This means we need to extend ordinals across sets of events, and extend changes to include both changes in parameters (being updated about an event, or forgetting aspects of an event).

We define the extension of a quantity Q over a set of episodic memories M as (quantityAspect Q M). Thus we can say

```
(> (quantityAspect Enthusiasm
  (subsetOfType Snorkeling
    (EpisodicMemoriesOf K)))
  (quantityAspect Enthusiasm
    (subsetOfType DentalWork
      (EpisodicMemoriesOf K))))
```

Evaluating such ordinal relationships is easy when there are numerical values, as can sometimes be done with physical domains. That is not an option here, but fortunately QR research has provided some useful ways of eking out conclusions from partial information. For example, bipartite graph partitioning of opposing signs used in influence resolution can sometimes generate answers when enough ordinal relationships between specific events are known. Similarly, symbolic algebras can often provide ordinal information across a broader range of quantities. For instance, if (VeryHighAmountFn Enthusiasm) applied to Snorkeling and (VeryLowAmountFn Enthusiasm) applied to DentalWork, the ordinal above would follow from these.

A particularly subtle effect is when a value has increased, but not so much as to be detectable via ordinals computed over aspects. Enough such small increases can lead to a derivative increase if the comparison were across

⁴ In the early stages of a relationship, preferences can be misstated in order to increase closeness, e.g. expressing interest in board games early on but then refusing to play once married <https://www.ny-times.com/2024/05/31/magazine/judge-john-hodgman-on-compulsory-game-nights.html>

⁵ See Kurosawa's Rashomon for an extreme example. <https://en.wikipedia.org/wiki/Rashomon>

a broader range of memories, as when values “break over” in order of magnitude representations (Dauge, 1993). A similar mechanism might be useful here.

A Simple QP Model of Social Relationships

Now we have enough representational machinery built up to construct a simple model of social relationships. We start by formalizing the concept of a social relationship between two people. We formalize social relationships as conceptual entities, represented by model fragments. Each involves a pair of instances of `Person`⁶.

```
(defModelFragment SocialReIn
  :participants ((?me :type Person)
                 (?other :type Person))
  :conditions ((knowEachOther ?me ?other))
  :consequences ((hasQuantity
                  (ClosenessFn ?self))
                 <...>)
```

The variable `?self` is a meta-linguistic convention: `?self` is always bound to the model fragment instance, so that its properties can be specified in the definition of the model fragment type. More consequences of this model fragment are enumerated below.

Recall that in model fragments, the variables for participants define role relations. So if Pat knows Kit, then

```
(isa SR1 SocialReIn)
(me SR1 Pat)
(other SR1 Kit)
```

Notice that this relationship is unidirectional – if `knowEachOther` is a symmetric predicate, then the first implies a social relationship in the other direction:

```
(isa SR2 SocialReIn)
(me SR2 Kit)
(other SR2 Pat)
```

This allows for asymmetric relationships, e.g. Pat might feel more close to Kit than the other way around. By building in a perspective `SocialReIn` also better supports reasoning from that perspective, e.g. an agent (`me`) reasoning about whether someone else (`other`) might be called upon to help or participate in some other mutual activity.

We will take the closeness that a person feels for another (i.e. (`ClosenessFn SR1`) as how close Pat feels to Kit) as a quantity that, when sufficiently high, causes them to believe that the other person is a friend. In other words, there is a limit point at which this transition happens, but we will not specify it except to denote it as (`FriendLimitPointFn SR1`).

Closeness appears to depend on multiple factors. One aspect is shared background, e.g. if two people are routinely engaged in the same kinds of activities and have overlapping social networks, they have a built-in basis of common ground. Another aspect is shared experiences, which builds

up a shared history together. Notice that this shared history does not always have to be enjoyable: A slogan in the US military is “Shared pain leads to unit cohesion.” This assumes the source of the shared pain is outside the dyad, e.g. an unexpected thunderstorm ruining a picnic might bring people closer together, especially if they worked together to ameliorate its negative effects. On the other hand, if one of the planners was feckless and ignored a weather forecast predicting a serious thunderstorm, that would most likely decrease, not increase, closeness.

To capture shared experiences, we model episodic memory as a set of cases. Recall that

(`EpisodicMemoryFn <agent>`) denotes the set of cases that constitute the episodic memory of agent `<agent>`. We define a subset of episodic memories relevant to a social relationship by those which mention another `Person` as

(`InteractionEpisodicMemoriesFn <agent> <other>`) and those episodic memories relevant to a social relationship via

(`SocialReInMemoriesFn <SocialReIn>`) consisting of the `InteractionEpisodicMemoriesFn` for the `me` and other of that relationship.

We split positive and negative aspects of events because people seem to track them separately. For example, we can distinguish between an event whose net impact is small because it only had a small positive impact, or because there are large impacts of opposite signs, the latter being a more fraught situation. The positive and negative effects of events involving the people in a social relationship will be represented by two quantities, (`PosExperienceFn ?self`) and (`NegExperienceFn ?self`). These are accumulations over the set of episodic memories of the `me` of the relationship for memories where the other is involved. Thus when a new event is experienced, its positive and negative impacts will be considered in estimating these quantities. Recall that there are no numerical values associated with these quantities, by assumption. Instead, ordinal values are updated based on local information. Suppose a new event `E'` is added to episodic memory, and it was more positive than negative. Then, treating closeness as an extensive parameter, the new value for closeness, whatever it is, is higher than the value before this event. The two people have become closer. Should it have been more negative than positive, the new value for closeness would be recorded as less than the prior value, i.e. $D_s = -1$.

This method of tracking causal changes due to differences caused by adding events⁷ relies on local changes and the accumulation of ordinals across time. This detailed record-keeping may or may not be psychologically plausible, and it does not provide easy comparison across people, e.g. are you closer to one friend than another? There might be a summarization mechanism that tracks accumulation of ordinal changes through changes in a parallel symbolic algebra representation, e.g. `HighAmountFn` transitioning to `VeryHighAmountFn`. We discuss possible quantitative extensions in future work.

⁶ In OpenCyc, `Person` inherits from `IntelligentAgent` and `SocialBeing`, but does not include organizations, hence apt for this purpose

⁷ Exactly how much forgetting (as opposed to failing to consciously retrieve) occurs in episodic memory is still an open question. If episodes are

forgotten, does that process somehow update the quantities that it was involved in changing? This seems unlikely.

We introduce the following additional quantities for `SocialReln`, with the understanding that this set is likely incomplete:

- `TrustLevelFn`: How much you trust the other person with regard to information-sharing.
- `ReliabilityFn`: How likely will they do what they say they will do.
- `HelpfulnessFn`: How likely they are to be willing to help do something.
- `InterestsOverlapFn`: How many of your interests do they share?
- `FriendsOverlapFn`: How much do your social networks overlap?

What do these parameters depend upon? Let us start with trust. The common business metaphor “trust is built in drops and lost in buckets” suggests an accumulation, albeit asymmetric in flow rates, which would make it an extensive parameter. For each event `E` in `InteractionEpisodicMemoriesFn` for a social relationship `S`, for every information sharing norm `N` in `S`, either `E` is agnostic with respect to `N`, or represents compliance with `N`, or represents a violation of `N`. Depending on the relationship between `E` and `N`, `E` either doesn’t contribute to either experience parameters, contributes to (`PosExperienceFn S`), or contributes to (`NegExperienceFn S`). In other words, adhering versus violating norms on information sharing in the relationship should increase/decrease closeness as well as trust. A separate `PosTrustRateFn` and `NegTrustRateFn` are introduced to represent the effects of adherence/violation to trust, with the magnitude of `NegTrustRateFn` being much larger than `PosTrustRateFn`, to model the gradual accumulation of trust and sharp dissolution of it upon betrayal. (An order of magnitude relationship might not be amiss here.) Note that, in addition to asymmetry, keeping these rates separate from their effect on closeness should better enable modeling that someone might get closer again to another while no longer trusting them with regard to keeping secrets.

Reliability can be modeled as a ratio of commitments honored to commitments made, e.g. the consequences of `SocialReln` should include

```
(qprop+ (ReliabilityFn ?self)
  (NCommitmentsHonoredFn ?self)
(qprop- (ReliabilityFn ?self)
  (NCommitmentsMadeFn ?self))
```

Where `NCommitmentsMadeFn` and `NCommitmentsHonoredFn` are the cardinalities of sets whose members consist of the set of commitments made by the other and the number of those which were honored, respectively. This illustrates the importance of reputation: An agent must assess this information either directly, from experiences with the other, or indirectly, from what yet other agents say about the other agent in the relationship. That said, there are commitments with varying levels of importance, e.g. keeping secret a surprise birthday party versus keeping secret that someone plans to leave their job. So cardinality is unlikely to be sufficient to capture the causal relationship here. This is not an isolated case, so we should develop a general representation that can be specialized via reasoning appropriately. Let us define `ImportanceScoreFn` as a binary function whose first argument is a continuous quantity and whose second argument is a social agent. Its range is in turn a unary function

whose domain is a set of events and whose range is a set of quantity values. For example,

```
((ImportanceScoreFn eventValue <me>)
  (EventsWithCommitments <other> <me>))
```

provides a set of quantity values for the set of events,

where `(EventsWithCommitments <other> <me>)` expands to

```
(TheClosedRetrievalSetOf ?e
  (and (isa ?e Event)
    (commitmentInEvent ?e <other>)
    (eventValue ?e ?v)))
```

and `(EventsCommitmentsHonored <other> <me>)` is

```
(TheClosedRetrievalSetOf ?e
  (and (isa ?e Event)
    (commitmentInEvent ?e <other>)
    (commitmentHonoredInEvent ?e <other>)
    (eventValue ?e ?v)))
```

We are assuming finite symbol values here for simplicity, accumulating ordinal information about particular properties of events is also doable but somewhat more complex.

Helpfulness can be defined analogously, i.e. a positive influence based on the number of times the other helped the me agent by some action, including joint activity, plus a negative influence based on the perceived number of opportunities to be helpful that the other agent had but did not take. Again this could be based simply on cardinality, the number of times they were helpful, or each contribution could be scaled based on utility, e.g. giving someone \$5 versus giving someone your kidney for a transplant.

The final two quantities that seem relevant are the overlaps in interests and in friends. Again, these might simply be cardinality in set intersections or scaled based on significance. For example, in the US it was once not uncommon for spouses to belong to different political parties, whereas now political orientation is typically a gating factor on long-term involvement. Similarly, if one shares close friends with someone else, that is likely to have more impact on closeness than sharing assorted random acquaintances.

This formulation of sets and importance measures to define quantity values is quite different from the traditional notion of directly influenced parameters, where the derivative of a quantity is specified continuously over time. In traditional continuous change models, effects accumulate continuously. For these parameters, the values change discretely, as the members of particular sets of events, interests, or people change over time (and perhaps change in evaluation as well). Nonetheless, the compositional causal relationships do seem to capture the intended effects of changes in the sets and in the evaluations of members of those sets.

Example: Planning an Outing

Pleasant outings often involve planning, which should take into account the preferences of the people involved. Suppose Kit did the planning for the outing described above and has already decided to propose a picnic and a swim. There

are two paths to the clearing, one a pleasant amble and another requiring climbing equipment. Suppose further that

```
(PrefersBeliefsOfFn Kit Kit):
(attitudeTowardsType Kit Walking
  (HighAmountFn Enthusiasm))
(attitudeTowardsType Kit RockClimbing
  (VeryHighAmountFn Dislike))
(PrefersBeliefsOfFn Kit Pat):
(attitudeTowardsType Pat Walking
  (MediumAmountFn Enthusiasm))
(attitudeTowardsType Pat RockClimbing
  (HighAmountFn Enthusiasm))
(attitudeTowardsType Pat Swimming
  (HighAmountFn Enthusiasm))
```

We assume that these `attitudeTowardsType` statements are generated via computations over episodic memories, both of things that they have done together but also Kit's understanding of Pat's self-reports or third-party stories about Pat.

Given straightforward reasoning about relative magnitudes and the negative relationship between enthusiasm and dislike, we get a conflict in preferences:

```
(prefers Pat (activityInPlan P1 RockClimbing)
  (activityInPlan P1 Walking))
(prefers Kit (activityInPlan P1 Walking)
  (activityInPlan P1 RockClimbing))
```

where `activityInPlan` means that the plan denoted by the first argument includes one or more instances of the concept denoted by the second argument. So while Kit believes that Pat would prefer rock climbing, Kit does not want to be miserable, which would make the outing less fun for them both, and so proposes walking, which both should find acceptable.

Differences in beliefs can lead to surprises:

```
(PrefersBeliefsOfFn Pat Pat):
(attitudeTowardsType Pat Walking
  (MediumAmountFn Enthusiasm))
(attitudeTowardsType Pat RockClimbing
  (HighAmountFn Enthusiasm))
(attitudeTowardsType Pat Snorkeling
  (HighAmountFn Enthusiasm))
(attitudeTowardsType Pat Swimming
  (LowAmountFn Enthusiasm))
```

Kit may have inferred their belief that Pat is very enthusiastic about swimming from hearing that Pat is a snorkeling enthusiast. But if Pat's love of snorkeling comes from seeing coral reefs, swimming in a lake just isn't the same thing, hence the low enthusiasm for swimming per se. Nonetheless, it would still be a net positive experience for Pat, and hence his acceptance of the proposal.

Conclusions and Future Work

Like many other aspects of commonsense, social reasoning seems to have continuous aspects, and this paper argues that QP theory with two extensions may be able to formalize those aspects. These extensions bridge from qualitative

modeling to the more discrete world of events and the accumulation of these events into episodic memories that are analyzed to track relationship parameters over time. This account relies on higher-order representations, e.g. microtheories for representing cases and states of belief and states of affairs in the world.

There are two next steps. The first is to expand the formalization to handle more phenomena. For instance, if someone is betrayed, how does that impact their interpretation of their episodic memories of prior interactions? Are the continuous aspects of people's models stored in a distributed fashion (e.g. Friedman et al. 2018), so that different models for another person are retrieved under different classes of situations? The second is to implement the non-QP aspects of the reasoning described here, to test these ideas at reasonable scale. We plan to explore whether or not this account can be extended to support story understanding, e.g. to predict changes in social relationships between characters as the events of a story unfold.

There is a looming open question: How far can a purely qualitative account go, especially as the size of episodic memory grows? Is a quantitative substrate inevitable, to facilitate cross-person comparisons? Models of emotion, for example, compute appraisal variables (Gratch & Marsella, 2004; Wilson et al. 2013), and similar computations could be used for the quantities used here. It could be that some internal quantities are used to track the impact of experiences, but that qualitative representations are used to facilitate planning and prediction. This is a question worth exploring.

Acknowledgments

This research was sponsored by the US Air Force Office of Scientific Research under award number FA95550-20-1-0091.

References

- Bredeweg, B., Salles, P., Bouwer, A., Liem, J., Nuttle, T., ... Zitek A. (2008). Towards a structured approach to building qualitative reasoning models and simulations. *Ecological Informatics*, 3, 1-12.
- Dauge, P. (1993). Symbolic reasoning with relative orders of magnitude. *Proceedings of the Thirteenth IJCAI*, pp. 1509-1515. Chambéry, France.
- de Jong, H. (2008). Qualitative modeling and simulation of bacterial regulatory networks. In A. Uhrmacher & M. Heiner (Eds.), *Computational Methods in Systems Biology*. Springer-Verlag, Berlin, Germany.
- Drabble, B. (1993). EXCALIBUR: A program for planning and reasoning with processes. *Artificial Intelligence*, 62, 91-140.
- Fisk, A. (1992). The four elementary forms of sociality: Framework for a unified theory of social relations. *Psychological Review*, 99, 689-723.
- Fisk, A. (2004). Relational Models Theory 2.0. In Nick Haslam (Ed.) *Relational Models Theory: A Contemporary Overview*. LEA Associates, New Jersey.
- Forbus, K. (1984). Qualitative process theory. *Artificial Intelligence*, 24, 85-168.

Forbus, K. (1989). Introducing actions into qualitative simulation. *Proceedings of the 11th IJCAI*, pp. 1273-1279, Detroit, MI.

Forbus, K. (2019). *Qualitative Representations: How People Reason and Learn about the Continuous World*. MIT Press, Cambridge, MA.

Forbus, K. (2023) Domain Theories for Commonsense Reasoning from Language-Grounded Ontologies. *Proceedings of the 36th International Workshop on Qualitative Reasoning* Krakow, Poland.

Friedman, S., Forbus, K., & Sherin, B. (2018). Representing, Running, and Revising Mental Models: A Computational Model. *Cognitive Science*, 1110-1145. DOI:10.1111/cogs.12574.

Gratch, J. & Marsella, S. (2004). A domain-independent framework for modeling emotion. *Journal of Cognitive Systems Research*, 5(4):269-306.

Guha, R. (1992) *Contexts: a formalization and some applications*. PhD. Thesis, Stanford University.

Hancock, W. & Forbus, K. (2021). Qualitative Spatiotemporal Representations of Episodic Memory for Strategic Reasoning. *Proceedings of QR-2021*.

Hayes, P. (1985). Naïve Physics I: ontology for liquids. In J.R. Hobbs & R. Moore (Eds.), *Formal Theories of the Commonsense World*. Ablex: Norwood, NJ.

Hogge, J. (1987). Compiling plan operators from domains expressed in qualitative process theory. *Proceedings of AAAI 1987*, AAAI Press, Menlo Park, CA.

Kelley, H., Holmes, J., Kerr, N., Reis, H., Rusbult, C., & Van Lange, P. (2003). *An Atlas of Interpersonal Situations*. Cambridge University Press, Cambridge, England.

Kim, H. (1993). *Qualitative reasoning about fluids and mechanics*. Ph.D. dissertation and ILS Technical Report, Northwestern University.

Lenat, Douglas B.; Guha, R. V.; Pittman, Karen; Pratt, Dexter; Shepherd, Mary (August 1990). "Cyc: Toward Programs with Common Sense". *Communications ACM*. 33 (8): 30–49. doi:10.1145/79173.79176

Malle, B.F.; Bello, P.; and Scheutz, M. (2019). Requirements for an Artificial Agent with Norm Competence. *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society* (pp. 21-27).

Olson, T. & Forbus, K. (2021). Learning Norms via Natural Language Teachings. *Proceedings of the 9th Annual Conference on Advances in Cognitive Systems 2021*.

Rawlins, W. (1992). *Friendship matters: Communication, dialectics, and the life course*. Routledge, 2017.

Simmons, R. (1983). *Representing and reasoning about change in geologic interpretation*. Ph.D. dissertation, MIT, Cambridge, MA.

Tomasello, M. (2001). *The Cultural Origins of Human Cognition*. Harvard University Press, Cambridge, MA.

Wilson, J., Forbus, K., and McClure, M. (2013). Am I Really Scared? A Multi-phase Computational Model of Emotions. *Proceedings of the 2nd Conference on Advances in Cognitive Systems*.

Wrench, J., Punyanunt-Carter, N., & Thweatt, K. (2023) Chapter 10: Friendship Relationships, in *Interpersonal Communication: A Mindful Approach to Relationships*, LibreTexts. <https://milnepublishing.geneseo.edu/interpersonalcommunication/chapter/10/>