

Qualitative Spatiotemporal Representations of Episodic Memory for Strategic Reasoning

Will Hancock and Kenneth D. Forbus

Northwestern University

wwhancock@u.northwestern.edu, forbus@northwestern.edu

Abstract

Episodic memory is crucial to an intelligent agent’s ability to cope with a diverse array of tasks. Following Hayes’ notion of histories, we hypothesize that qualitative spatiotemporal representations play an important role in organizing episodic memories about continuous phenomena. This paper describes a novel representation (**QualitativeSpatioTemporalEpisodicMemory**) that carves up continuous properties into qualitative regions of space and time. These representations ground experience in local, qualitative, structured episodes. We evaluate this work by showing that our agent can effectively learn strategy game decisions from expert replays using these representations.

1 Introduction

Episodic memory is a persistent contextualized store of specific events (Tulving, 1983), often involving both spatial and temporal aspects. It is a powerful cognitive mechanism, since it supports learning by experience, e.g., using analogy to understand one’s options when making a decision based on prior experiences.

In qualitative reasoning research, Hayes proposed the notion of *histories* (Hayes, 1978;1989) as a general framework for representing change in continuous domains. Whereas the situation calculus describes changes in terms of discrete events and provides no spatial constraints (leading to the infamous Frame Problem), histories represent change in terms of pieces of space-time, based on the objects involved in the changes.

Cognitive systems research has focused on representations of episodic memory that are either completely task independent (Laird & Derbinsky, 2009; Menager & Choi, 2016), or task specific (Brom et al., 2016). By contrast, we propose to utilize the notion of histories to ground episodic memory in local, qualitative, structured episodes. The idea is to use the spatial aspects of the entities involved in an event or situation, combined with their temporal duration, to provide a representation for changes over time. Histories have been heavily used in qualitative reasoning (Forbus, 2019); thus we hypothesize that they will prove invaluable for organizing episodic memories for continuous aspects of domains.

Previously, we showed that histories can provide a useful representation for strategic reasoning (Hancock et al., 2020). Our current hypothesis is that histories can provide a useful representation for episodic memories and therefore support learning from experience. This paper describes our novel representation (QSTEM) and an investigation of this hypothesis using the strategy game Freeciv as a testbed.

2 Background and Related Work

We begin with a brief overview of our test domain (Freeciv), our visual processing system (CogSketch), and our learning mechanism (analogy).

2.1 Freeciv

Freeciv is an open-source strategy game based on Sid Meier’s Civilization II. The player must establish a multi-city civilization and manage its growth, economic vitality, and scientific progress, while simultaneously developing a military for defense and offense. The player wins by either sending a colony ship to Alpha Centauri or conquering the world.

Freeciv is an excellent domain for AI research due to its complexity. A typical game board consists of 4,000 tiles with varying terrain. Games typically last for hundreds of turns, and each turn involves many decisions. Some decisions are global across the entire civilization, such as setting the tax rate, determining the next technology to research, and engaging in diplomacy. Workers must be kept busy modifying terrain. Military units must defend cities and conduct attacks on opponents when at war. By contrast, Go is played on a 19x19 grid with uniform, immutable spatial properties which are always visible from the start. In addition, each turn in Go only involves a decision to place one piece.

Freeciv is especially useful for exploring histories because important behaviors happen at multiple grain sizes. For example, there is typically an expansion phase, where a player builds out multiple cities, to stake out desirable territory and deny it to competitors. Wars can cause the expansion or contraction of a player’s civilization, depending on their success, making decomposing time based on the set of cities a useful distinction. While the expansion phase enables reasoning at the level of an entire civilization, reasoning must also occur at the level of individual units and cities. Deciding what re-

source should be built next in a city should benefit from analyzing experiences built on histories similar spatially to its local area as well as temporally local to its needs.

2.2 CogSketch

CogSketch (Forbus et al., 2011) is a sketch understanding system that provides a model of high-level visual processing. It provides multiple, hierarchical levels of visual representation, including decomposing digital ink into edges, combining edges into entities, and grouping based on gestalt principles.

The basic level of organization for ink in CogSketch is the *glyph*, one or more ink strokes that are taken to represent some entity (abstract or concrete) in the sketch. People can generate glyphs using a pen or mouse to produce digital ink. Glyphs can also be automatically produced via visual analysis of images (Chen et al. 2019).

CogSketch can compute various spatial relationships between glyphs, including adjacency, relative position, topological relationships, and relative size. Properties of individual glyphs can also be computed, such as shape attributes (roundness) as well as glyph decompositions, such as shape skeletons and Voronoi diagrams. Moreover, CogSketch can decompose glyphs into edges, and group visual entities based on gestalt properties. These capabilities have enabled it to model a variety of visual problem-solving tasks, including Ravens' Progressive Matrices (Lovett & Forbus, 2017), making it a useful platform for human-like visual reasoning.

Here we make use of topological relations automatically calculated between spatial regions. We rely mainly on the region connection calculus (RCC8) (Cohn et al., 1997) and describe our spatial representations further in section 3.5.

Prior Work on Blobs

(McLure et al., 2012) formalized an encoding schema for geographical regions within Freeciv. They introduced three distinct types: trafficability, terrain, and continent blobs. In this paper, we use continent blobs. Continent blobs correspond to the geographical concept of continents. There are two distinct types: land and water regions. Each blob consists of a contiguous set of Freeciv tiles of the same land or water type.

Footprint Regions

In addition to continent, terrain, and trafficability blobs, (Hancock et al., 2020) introduced a new type of blob, the *footprint* of a compound entity. A footprint is the region defined by a convex hull of a group of similar objects with spatial extent. That previous study used civilization footprints. This paper introduces two new types of footprints:

1. City footprints: a city site and collection of surrounding regions under a city's control
2. Unit footprints: a spatially local group of units under a single player's control

Information about each footprint is encoded along with its local spatial relations. Different footprint types exhibit different behaviors over time. The civilization footprint grows and shrinks over time as cities are founded and destroyed. City

footprints are static; they exist as long as the city exists and do not change in size. Unit footprints can grow and shrink, but also change location, as units are mobile. Trajectory information about enemy unit blobs is encoded relative to our player's city and civilization footprints. Currently, trajectory is represented as a Boolean; either a unit blob is approaching some other region, or it is not. This information is only encoded when an enemy is approaching. We define this condition as a monotonically decreasing distance with a minimum history of some number of turns, in this study, three turns. Given this condition, we record the fact that an enemy unit blob is approaching the civilization footprint, as well as the closest city footprint.

2.3 Analogical Learning

One goal of this research is to demonstrate how analogy can support episodic memory retrieval. To do this, we use the Structure Mapping Engine (SME) (Forbus et al. 2017) for comparison, the MAC/FAC retrieval system (Forbus et al., 1995), and the SAGE generalization system (McLure et al. 2015). These analogical mechanisms are tightly integrated in the underlying reasoning engine and provide the mechanisms for retrieval, comparison, and transfer. SME not only assesses the similarity of a precedent to the current situation, but also projects the previous solution into the new case by translating entities to their mapped equivalents, in the form of candidate inferences. Thus, analogy provides adaptation automatically. The unit of comparison and retrieval is a case. In this approach, cases are not entire games but instead facts temporally and spatially local to some entity of interest. For example, cases can capture a decision about what improvements to build, what tiles to work, and what technologies to research. This work focuses on the task of deciding what improvements to build (i.e. production decisions), (see section 3.7).

When a production decision is made in the game, a spatio-temporal snapshot is constructed and stored in the game context. Case snapshots collected from expert replays form the corpus of relevant decisions. That is, each production decision type (city walls, warriors, etc.) has a SAGE generalization pool which maintains a set of automatically constructed generalizations and outliers. A generalization pool is also a MAC/FAC case library, so adding a new case causes the most similar item to be retrieved. If the degree of similarity is over the assimilation threshold, then the new case is assimilated into the retrieved item. This assimilation process produces a new generalization if the retrieved item is an outlier, or if the retrieved item is a generalization, updates it with the contents of the new case. Every statement in a generalization has a probability based on the fraction of cases in which it was true, and non-identical entities are replaced by more abstract entities. If the similarity of the retrieved item is low, the new case is added to the pool as a new outlier.

When our agent is playing autonomously, it makes production decisions by considering what the expert chose to do in a similar context. First, a case is constructed in the same manner as it was constructed for the expert's action when it watched them play. This case is then used as a probe, and

similar experiences (or generalizations built from experiences) are retrieved (see section 4).




Figure 1: A gpool representing the decision to build settlers. Each white circle corresponds to a generalization, and each black dot a case in that generalization.

3 Qualitative Spatiotemporal Representations

This section describes our history representation. We start by describing our guiding principles for encoding. We then describe quantities, which form the basis for reasoning about continuous change. Next, we discuss how continuous quantities are incorporated into a temporal graph, giving a relational representation of a region in time. Finally, we discuss the incorporation of spatial information, as well as other domain information that is also relevant to strategic reasoning.

3.1 Encoding Principles

We seek to construct representations that are concise, sparse, and local. *Concise* means that the representations respect the relevance principle (Forbus, 2019) by not making unnecessary distinctions. Regions where a relevant property is constant should be one entity, subdivided only when required due to some other important constraint. *Sparsity* is an aid to learning. There is always a tradeoff in the amount of information encoded. Too little, and there isn't enough signal to learn the appropriate distinctions. Too much, and learning is made more difficult because the space of hypotheses is larger. More entities means more potential relations between them, which means more work to analogically match descriptions. All else being equal, keeping the number of entities small is preferable. This is one reason that qualitative representations are better than, for example, using tiles or other highly granular quantitative representations as commonly done in reinforcement learning research. *Locality* means that networks of relationships should be computed between entities that are spatially and temporally local to one another. This makes the set of relationships computed for two descriptions more likely to

match when they are similar. This locality heuristic is already built into several of CogSketch's algorithms, e.g. positional relationships are only computed between adjacent glyphs. The same concept has been extended to computing relations between temporal intervals.

3.3 Quantities

Freeciv players must reason about a wide variety of continuous properties and try to optimize many different aspects of their civilization at the same time. For instance, resource management involves maintaining gold supplies, production rates, and scientific output, while combat requires reasoning about relative unit strengths, horde sizes, and terrain types suitable for combat, etc.

These quantities have instantaneous values; the current gold in the treasury, the number of cities controlled by a player, and many more are directly available to the player. To effectively reason about the game, a player must also monitor the ways that these quantities are changing. The amount of gold may be the same in two different scenarios, but if it is rapidly increasing in one and diminishing in the other, these situations should be represented as two qualitatively distinct states.

Temporal bounds can be determined for a given quantity and its change over time. We first describe levels of quantity encodings and then describe how the spatial and temporal extent is determined for quantity encodings. Currently, we define three distinct levels of quantity encodings:

1. Sign of the derivative: can be in one of three states, constant, monotonically increasing, or monotonically decreasing
2. Value sign: this quantity's state is defined by its value being zero, positive, or negative.
3. Value: this quantity's state is determined solely by its value.

We use one of these encoding levels for every Freeciv quantity used in cases. For example, the civilization footprint size is simply the area of the civilization footprint. We use sign of derivative for this quantity, i.e. its current growth is marked as constant, positive, or negative. Cities are usually founded on the first turn immediately making growth of the footprint area positive. Since many turns see no change in the number of cities, we treat periods of monotonic nondecreasing as episodes. Hence growth occurs until a player loses a city; this may be attributed to enemy conquest, city starvation, or other reasons. The first temporal interval for this quantity is thus defined by the bound of the initial growth period. When a city is lost, a qualitative state is introduced corresponding to the period in which the size of the national footprint is monotonically nonincreasing.

Certain game quantities are strictly monotonically nondecreasing or nonincreasing, such as the cardinality of scientific discoveries. When encoded into a QSTEM representation, quantity encoding (3) provides useful temporal information for these types of quantities. For example, a temporally local

qualitative interval for scientific discovery will encode the time period since the last discovery until the present.

3.4 Temporal Representations

Our representation of time draws inspiration from Allen’s work on temporal graphs (Allen, 1983). Specifically, our representation takes temporal intervals as primitives, with relations between intervals encoded as graph edges. We use Allen’s Interval Algebra (AIA) to denote these relations. Allen notes: “This representation is designed explicitly to deal with the problem that much of our temporal knowledge is relative, and hence cannot be described by a date”.

Intervals are defined by properties of continuous quantities. They correspond to concepts such as a specific period of growth or a period of invasion for a city. AIA encodes the relationship between intervals, e.g. the invasion began during a period of growth.

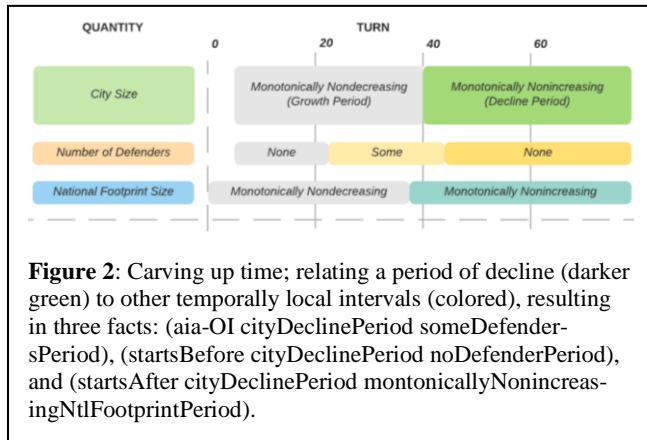


Figure 2: Carving up time; relating a period of decline (darker green) to other temporally local intervals (colored), resulting in three facts: (aia-OI cityDeclinePeriod someDefenderPeriod), (startsBefore cityDeclinePeriod noDefenderPeriod), and (startsAfter cityDeclinePeriod monotonicallyNonincreasingNtlFootprintPeriod).

The graph denoted by a collection of such intervals may contain many edges, thus violating the conciseness principle. We take a simple approach to limiting complexity, relating intervals only when adjacent or overlapping. Figure 2 illustrates this, where the nonincreasing city size interval is related to other quantity intervals. Temporal relations between the number of defenders and national footprint size are not reified. This choice for encoding follows our general guidelines; representations should be concise.

The nature of experience is that we do not know when a present temporal state will end. AIA relations between two intervals depend on knowing the endpoint of at least one of them. If one interval in question is complete, it may be possible to assign an AIA relation with an interval that is incomplete. If not, or if both are incomplete, we reduce the set of possible relations to (startsBefore, startsAfter, and startsAt-SameTime).

3.5 Representations of Space

The spatial component of an episode is determined by the previously described blob regions and spatial relations between them. Our encoding focuses mainly on relations defined by the region-connection calculus (RCC8). Briefly, RCC8 encodes connectivity between regions, including containment, partial overlap, non-connectedness, etc.

How might region relations guide strategic reasoning? We take the relationship between a city footprint and a national footprint as an example. A city footprint may lie on the border or interior region of the national footprint. RCC8 is able to distinguish these two situations (rcc8-NonTangentialProperPart vs rcc8-TangentialProperPart). For the task of making a production decision, this distinction may mean life or death for a city, since cities on the border of a footprint are much more likely to be attacked than those on the interior.

A city footprint may also overlap an enemy region or may share a DMZ with an enemy. These concepts are also encoded into our representation (see section 3.7).

In addition to RCC8 relations, we also utilize the concept of bisection. A region A bisects a region B if the difference of B – A results in more than one contiguous region. This relation is computed between all glyphs and reified when it holds.

3.6 Connecting Time and Space

So far we have discussed both representation of time as well as space, but how should our representation connect them? To follow the locality principle, only quantity changes that are active in some spatial context should be considered. To support analogical learning, these relations should be reified to support higher-order structure.

Locality is straightforward for space; all RCC8 relations except for rcc8-Disconnected constitute a local relation. Similarly, locality is well-defined for time. The combination of space and time is less straightforward. Abstract quantities such as gold reserves (and its corresponding temporal intervals of growth and decay) do not have a clear spatial extent. We use a simple heuristic: as gold reserves are an attribute of a civilization, their spatial extent is confined to the corresponding civilization footprint. Similarly, attributes of individual cities are confined to the footprints of their respective cities.

As mentioned previously, time is grounded by quantities and their values. Space is grounded by some region. For production decisions, this means the national and city footprints corresponding to the city making a decision. In our representation, time and space meet where there is some target entity that has both a spatial and temporal extent. This is indeed the case for cities, and the spatial city footprint is associated with the temporal duration relating to the city’s size.

For example, consider the representation for the Freeciv city of Chicago. The temporal qualitative state of Chicago is designated by the growth or decline of its city size. At the start of a game, this implies a period of growth. Since Chicago also has a spatial extent (its footprint region), we associate these two quantities: (spatialExtentOf citySizeInterval FootprintOfChicago). Temporal intervals specific to Chicago (e.g. number of defenders) are then related to citySizeInterval using the methods described in 3.4.

3.7 Case Construction

The episodes that we are constructing are part of a much larger spatiotemporal web. Limiting what facts should be included is crucial to building concise representations.

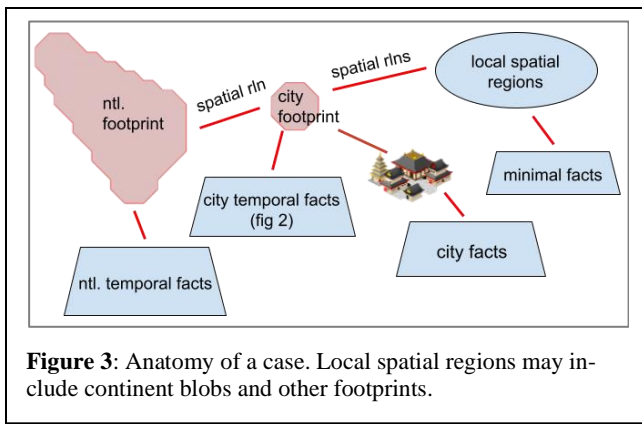


Figure 3: Anatomy of a case. Local spatial regions may include continent blobs and other footprints.

Construction starts by taking a spatially bound entity as an argument (here, a Freeciv city). Facts about this city are included, as defined by an existing case constructor that outputs any information known about the city.

Next, we query CogSketch for objects that share some relation to the city’s footprint. For each of these objects, we reify the spatial relation, along with a list of minimal case facts for that object (figure 3). For an enemy national footprint region, this is the set of isa statements about that region (isa region HostileTerritory), for example.

Parent/child relations between regions of the same type are handled separately. A city footprint region will always have some relation to its parent national footprint. In this case, the child’s temporal relation to the parent’s is reified, along with spatial and temporal facts relevant to the parent footprint.

Then, temporal intervals are reified for the city footprint. Its spatial change (growth, decline, constant) is related to quantity intervals relevant to that city (figure 2).

4 Experiment

Recalling past experience can assist an agent in deciding what actions to take. For this experiment, our agent learns production decisions from expert replays.

Production decisions in Freeciv are crucial to performing well in the game. A city can make many different types of resources. Combat units and city walls are needed to adequately defend a city. New settlers are needed to found new cities. There is an inherent tradeoff between growth and defense early in the game. Too much growth leads to underdefended cities. Too little growth means that a civilization is unable to keep technological pace with its opponents.

We hypothesize that using QSTEM to learn production decision from an expert can improve the performance of our agent. Learning from experts can be complicated, however. A novice that watches a grandmaster chess game may make no improvement because they are unable to understand why the experts are making certain decisions. Reasoning in Freeciv certainly requires high level strategies, but simple spatial awareness is pertinent. Similarly, time matters as well. Different decisions are made when enemies are at the gate vs. when one is left alone. Our representation takes these types of factors into account, and thus we hypothesize that episodes

represented this way should allow our agent to effectively learn from an expert player.

To generate the expert scenarios, ten Freeciv games were played until turn 100, or until the player lost. Each time a production decision was made, an episodic case was recorded. Overall, 460 cases were produced. The ten Freeciv maps used were partially revealed so that the entire starting continent was visible to the player.

To learn from these episodes, our agent makes production decisions based on those made in similar spatiotemporal contexts to the expert replays. A probe case is first constructed for the city making the decision, and a list of generalizations is returned, ordered by their structural similarity to the probe case.

In Freeciv, some resources are unavailable for production until a certain technology has been discovered. In the case where the most similar resource is unavailable, the next structurally similar resource is considered. If the three most similar resources are unavailable, the system defaults to a decision that aligns with its current goals. This is the same decision-making process as our baseline agent.

Our baseline agent builds upon existing work in qualitative reasoning (Forbus and Hinrichs, 2019). Fundamentally, it is a goal-based agent that leverages a learned qualitative model of Freeciv for decision making. The agent makes investment decisions about continuous quantities (e.g. should I increase the number of defenders in this city) and solely considers type-level goal activation as its decision criterion. That means that it only distinguishes alternatives on the basis of intrinsic properties i.e. properties shared by all members of their class. Extrinsic properties such as location, spatial configuration, and move points are therefore not considered. By contrast, our agent explicitly represents and learns these types of continuous extrinsic properties.

5 Results

Each agent type (baseline and episodic) was tested in ten new scenarios and statistics for the games were recorded. The maps in these scenarios were partially revealed in the same manner as the expert domains; the entire continent on which

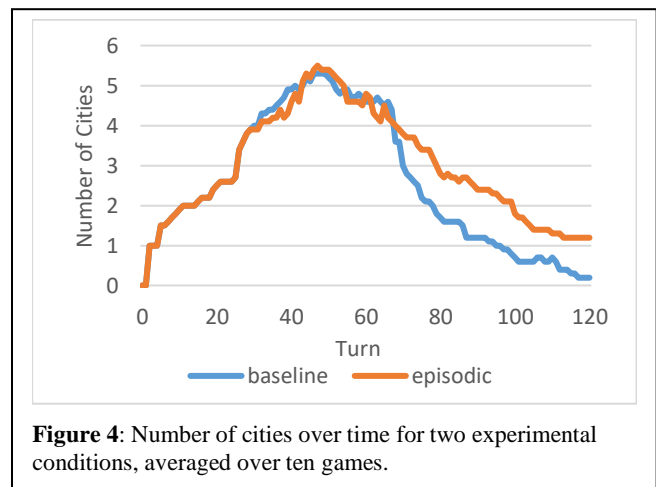


Figure 4: Number of cities over time for two experimental conditions, averaged over ten games.

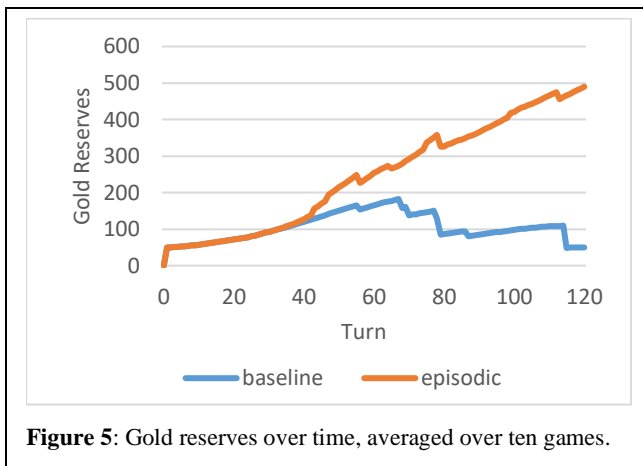


Figure 5: Gold reserves over time, averaged over ten games.

the agent started was unhidden. Success in Freeciv is complicated; there is no single metric to gauge performance. Civilization size, average city size, gold, science achievement, and many other factors go into determining an agent’s performance. Overall, surviving later in the game corresponds to better performance. However, it may be possible to survive by focusing only on defense while ignoring other aspects of one’s civilization. For this reason, we look at the number of cities that have survived as well as the amount of gold accumulated at turn 120. Both results are significant at $p < .05$.

While our results show an improvement, Freeciv is still a very difficult game. Evidence from watching our episodic agent play shows that it is making contextually relevant decisions. Border cities adjacent to enemies tend to specialize in defense, whereas interior cities are more likely to produce settlers and gold. In figure 4, both conditions are increasing their territory size at turn 40. The episodic condition begins to specialize certain cities towards gold production around the same time (figure 5) and is still able to better defend its cities and survive to later turns than the baseline.

6 Discussion and Future Work

In this paper, we described a history-based representation for qualitative spatiotemporal episodic memories. Furthermore, we showed that this representation is useful for learning production decisions from expert replays in the game of Freeciv.

The three quantity encodings outlined in this paper are a good start to reason with minimal domain knowledge about the complexities of a strategy game such as Freeciv. The learned episodes should be able to assist in further refining causal models. As the agent gains experience, it can learn limit points; the number of defenders sufficient to protect a city, the optimal number of cities for a particular spatiotemporal state, etc.

Another potential use for QSTEM in Freeciv is future state prediction. An episodic case encodes what is presently true in some circumstance; by retrieving similar cases from past experience tied with their eventual outcomes, analogy could be levered to predict spatial and temporal relations that might hold in the future. If my civilization is being invaded, looking at similar situation in the past might provide insight into whether my city will survive. This might be encoded with

varying temporal relations, e.g. a candidate inference indicating that there is a future growth phase, or a city size decline phase outlasting an invasion phase. Knowing what is likely to hold in the future could help with decision making in the present.

We hypothesize that QSTEM may be useful outside of strategy games due to its flexibility in representing continuous quantities. These representations may allow reframing an agent’s viewpoint; if larger pieces of space and longer temporal durations are not proving to reliably explain phenomenon, finer grained distinctions can be made by re-encoding.

Another possibility for extending this work is to support more general learning of continuous conceptual knowledge. We hypothesize that QSTEM will be useful for learning event types in interactive settings. Humans can identify behaviors such as rolling, bouncing, throwing, etc. Reasoning about these types of phenomena requires understanding continuous change in spatiotemporal contexts.

QSTEM provides both a framework for describing such changes, as well as a set of default encodings that make assumptions about the types of limit points one expects to see in the world. These assumptions cover phenomena such as an object being in motion or not, but not the concept of an object being near to another one. We would like to investigate more in depth the kinds of representations that are necessary for both learning a broad array of concepts, as well as allowing an agent to demonstrate that it understands these concepts by acting in the world.

Acknowledgments

This material is based upon work supported by the Air Force Office of Scientific Research under Award No. FA9550-16-1-0138.

References

- Allen, J. F. (1983). Maintaining knowledge about temporal intervals. *Communications of the ACM*, 26(11), 832–843.
- Brom, C., Pešková, K., & Lukavský, J. (2007). What Does Your Actor Remember? Towards Characters with a Full Episodic Memory. In M. Cavazza & S. Donikian (Eds.), *Virtual Storytelling. Using Virtual Reality Technologies for Storytelling* (Vol. 4871, pp. 89–101). Springer Berlin Heidelberg.
- Chen, K., Rabkina, I., McLure, M. D., & Forbus, K. D. (2019). Human-Like Sketch Object Recognition via Analogical Learning. *Proceedings of the AAAI Conference on Artificial Intelligence*, 33(01), 1336–1343.
- Cohn, A. G., Bennett, B., Gooday, J., & Gotts, N. M. (1997). Qualitative Spatial Representation and Reasoning with the Region Connection Calculus. *GeoInformatica*, 1(3), 275–316.
- Forbus, K. (2019). *Qualitative Representations: How People Reason and Learn about the Continuous World*. MIT Press.

- Forbus, K., Ferguson, R. W., Lovett, A., & Gentner, D. (2017). Extending SME to Handle Large-Scale Cognitive Modeling. *Cognitive Science*, *41*(5), 1152–1201.
- Forbus, K., Gentner, D., & Law, K. (1995). MAC/FAC: A model of similarity-based retrieval. *Cognitive Science*, *19*(2), 141–205.
- Forbus, K., & Hinrichs, T. (2019). *Qualitative Reasoning about Investment Decisions*. 7.
- Forbus, K., Usher, J., Lovett, A., Lockwood, K., & Wetzel, J. (2011). CogSketch: Sketch Understanding for Cognitive Science Research and for Education: Topics in Cognitive Science. *Topics in Cognitive Science*, *3*(4), 648–666.
- Hancock, W., Forbus, K. D., & Hinrichs, T. R. (2020). *Towards Qualitative Spatiotemporal Representations for Episodic Memory*. 33rd International Workshop on Qualitative Reasoning.
- Hayes, P. J. (1978). *The Naive Physics Manifesto*. Institut pour les etudes semantiques et cognitives.
- Hayes, P. J. (1989). Naive Physics I: Ontology for Liquids. In *Readings in qualitative reasoning about physical systems* (pp. 484–502). Morgan Kaufmann Publishers Inc.
- Laird, J. E., & Derbinsky, N. (2009). *A Year of Episodic Memory*. Workshop on Grand Challenges for Reasoning from Experiences.
- Lawrence, Z., & Peterson, D. J. (2016). *Mentally walking through doorways causes forgetting: The location updating effect and imagination*.
- Lovett, A., & Forbus, K. (2017). Modeling visual problem solving as analogical reasoning. *Psychological Review*, *124*(1), 60–90.
- McLure, M. D., & Forbus, K. D. (2012). *Encoding Strategies for Learning Geographical Concepts via Analogy*. 9.
- McLure, M. D., Friedman, S. E., & Forbus, K. D. (2015, February 10). Extending Analogical Generalization with Near-Misses. *Twenty-Ninth AAAI Conference on Artificial Intelligence*. Twenty-Ninth AAAI Conference on Artificial Intelligence.
- Menager, D. H., & Choi, D. (2016). *A Robust Implementation of Episodic Memory for a Cognitive Architecture*. Cognitive Science.
- Tulving, E. (1983). *Elements of Episodic Memory*. Clarendon Press.