Towards a generalized spatio-temporal understanding for the semantic web

Michael J Pan, Rob Raskin
NASA Jet Propulsion Laboratory
California Institute of Technology
{ michael.j.pan | rob.raskin } @ jpl.nasa.gov

Abstract

While there have been decades of research on spatio-temporal reasoning (STR) logics, the obstacle preventing a complete understanding of our world is the lack of spatio-temporal structures in existing ontologies. This is particularly true for semantic web efforts: for how do we begin to create a semantic web when we cannot define the semantics of the entities with which we which we wish to populate this web? In this paper we will outline some of the particular issues involved, focusing in particular on spatio-temporal knowledge representation. We begin with a discussion of various application requirements, provide a fundamental overview of ontological and spatio-temporal theories, and propose a more general approach for representing and reasoning spatio-temporal relations.

Introduction

The problem of a complete spatio-temporal understanding, despite attracting much research, remains open. For particular applications, however, a complete understanding is not necessary. Object-oriented programming languages is one such example. The classes created within these languages follow an inheritance hierarchy, and are generally assumed to exist solely within the running time and space of the program, and thus devoid of spatio-temporal properties. Anthropological taxonomies such as those of human evolution typically have associated with each class an approximate region and interval for which the class lived. The necessary representation can be quite simple, requiring simple values for the spatio-temporal fields. For GIS and earth science observational systems, various methods for spatial representation are available, and each offers benefits and carries associated costs. Particularly for earth science phenomena, events and their spatial locations are parameterized with respect to their temporal values, requiring a understanding of spatial-temporal interactions, especially with regard to change and motion. One last application that comes to mind is that of version management or source control (7). In addition to the distinctions made between transaction and valid times (i.e. when the changes are made and when the same changes become effective), it is possible to treat reconstruction of document versions as either a script or set-theoretic disjunction. Here time is an imperil factor, while spatial considerations can be disregarded. As we build a semantic web, we would like to formalize a generalized spatio-temporal understanding that is adaptable to any application.

Properties of time and space

As with all properties, the values we assign to specific instances of space and time are relative. They are relative to the reference point and the unit of measurement, and not absolute to themselves.

In this paper we assume an arbitrary reference point and unit of measurement, and also assume that differences in time and space derive from differences in values of these measures and their references. Those with different points of references and different units need an interface layer to provide a translation into each other, or into a third encompassing set.

Time

There are many considerations to temporal reasoning. The following is a quick overview of some of the fundamentals of temporal theories.

Temporal conceptualizations There are at least six senses for time as it is used in language (4). Time is

1. a dimension in that it is a property whose values are orthogonal to all other properties,
2. a plenum in that it contains a set of temporally possible worlds or timelines,
3. an interval in that it is a segment of a particular timeline,
4. a point in that it divides an entity into two parts and belongs to neither,
5. a duration in that it is the length of an interval in the time dimension, and
6. a position in that it is a point on a specific timeline.

Temporal models The various temporal models enable reasoning with respect to the situation.

1. Time is linear. Movement along this timeline is unidirectional and constant. This is the most commonly utilized model.

2. Time can be past branching. This is useful in abduction- in determining the causes that has led to the current world state.
3. Time can be future branching. This is useful in speculation about future world states given current actions.
4. Time can be circular. Production cycles and seasonal cycles are examples of these.

Temporal topology Temporal topologies defined how time is represented within an ontology.

1. Time is a set of discrete points. Each pair of adjacent points define an indivisible temporal interval.
2. Time is a dense set of points. Between each pair of points lies an infinite number of points.
3. Time is a continuum. Each subdivision of time is an interval.

Temporal point logic A point divides the time line into two halves, but is a part of neither. It is simply the intersection of a temporal plane (which spans across the branches) with a particular timeline, and it has no duration. Given two points P_1 and P_2

1. P_1 and P_2 denote the same point $\rightarrow P_1 = P_2 \rightarrow$ equal(P_1, P_2),
2. P_1 occurs after $P_2 \rightarrow P_1 > P_2 \rightarrow$ greater(P_1, P_2),
3. P_1 occurs before $P_2 \rightarrow P_1 < P_2 \rightarrow$ less(P_1, P_2).

Temporal interval logic It seems logical, and perhaps even more efficient computationally, to define a point as an interval with two identical endpoints. One soon realizes that such a definition becomes infinitely recursive. On the other hand, a single representation within a knowledge base does have its appeal.

Our approach holds that there are two fundamental types of time duration- points and intervals. An interval is the duration of time between two specific points on the time line. We additionally distinguish open versus closed intervals by stating whether an interval includes the point itself. A closed interval does, but an open one does not.

Given two intervals F and G such that

(1) { has(start-time, X_f), has(end-time, Y_f) } $\subseteq$ F ,
{ has(start-time, X_g), has(end-time, Y_g) } $\subseteq$ G ,
$X_f < Y_f$,
$X_g < Y_g$.

and assuming closed intervals,

1. equals(F,G) $\leftarrow X_f = X_g, Y_f = Y_g$.
2. precedes(F,G) $\leftarrow Y_f < X_g$.
3. meets(F,G) $\leftarrow Y_f = X_g$.
4. overlaps(F,G) $\leftarrow X_f$ ¡ Y_g, $Y_f > X_g$, $X_f < X_g$.
5. starts(F,G) $\leftarrow X_f = X_g, Y_f < Y_g$.
6. finishes(F,G) $\leftarrow X_f > X_g, Y_f = Y_g$.
7. during(F,G) $\leftarrow X_f > X_g, Y_f < Y_g$.
8. We obtain the 6 other axioms, which themselves are inverses of axioms 2-7, simply by exchanging all X_f and X_g, and Y_f and Y_g.

Space

The previous axioms generally hold for spatial reasoning as well. When we reason spatially, we usually are only interested whether X is in Y, i.e. whether there is an overlap relation.

Certain issues not considered in temporal reasoning arise within spatial reasoning. While humans have a stronger conceptual grasp of space than they do of time, reasoning in all three spatial dimensions simultaneously dramatically increases the complexity. Additionally, while humans inherently apply a linear scale to spatial locations, most locations with which we reason are based on Earth, which is roughly spherical in nature. Time is also somewhat easier to deal with as we conceive of it as a unidirectional dimension, whereas it is possible to move bidirectionally along spatial directions. While the temporal interval predicates apply generally to spatial reasoning as well, because we consider the three spatial dimensions simultaneously, the set of possible predicates expands exponentially. To our relief, we do not make use of all of them. The “overlaps” and “borders”, “distant” seem to be the three relations used in spatial reasoning, corresponding to the “overlaps”, “meets”, and “precedes” predicates respectively.

Earth science observational systems intrinsically require spatio-temporal understandings that are typically supported by a GIS system. The problem of locating entities is in the coordinate system with which we use. The choice of a stable reference point, especially since all the entities in the universe are moving away from each other seems impossible. Limiting a coordinate system to a single planet (in our case, Earth) makes the problem slightly easier, but not by much. One commonly used system on Earth is the latitude/longitude/surface elevation system. Numerous problems arise with this system. Distances and area calculations vary according to the latitude, and land surface elevation is not constant, while sea level, though more uniform, changes with the tides and temperature and other climate conditions.

That is the motivation behind the numerous geo-spatial coordinate systems, with one of the most commonly used being the latitude/longitude (polar) coordinate system. Nevertheless, it barely begins to resolve even just the issues mentioned above. These issues necessitate a difficult task in practical solutions for spatial storage and reasoning.

Even as we assume that locations are generally flat, given only a 2 1/2 dimensional representation (where we only consider a subset of the elevation), we still have the problem of representing objects (and locations) spatially. A rectangular approach is the simplest method but most things are not rectangular. Similarly, representation in a circular areal method has the same problem. Objects and locations are polygonal in nature. Reasoning with polygons incites complex computational geomentry issues which we leave for discussion elsewhere.

Spatio-temporal Perspectives

Sowa (6) differentiates between continuants and occurrents, though he admits that such distinctions are based upon the temporal range in context. Similarly, in his critique of Allen's theories, Galton (3) proposed a difference between objects in motion and objects at rest- that entities are either exclusively in states of motion or states of position.

Representation of such requires two differing schema, but it should not matter whether an entity has a temporal component or is in a particular state for all entities. All physical entities and even some abstract entities have a temporal component. The important thing is the ability to represent such once we have knowledge of it. Additionally, the same entity should not switch categories simply based upon the perspective taken.

The four perspectives with which we conceptualize entities and reason with respect to time and space are:

1. objective,
2. temporally relative,
3. spatially relative, and
4. spatio-temporally relative.

These perspectives are defined relative to the reference point from which the perspective is taken. The "objective" perspective is one in which the reference point lies independent of space and time. One such example is "the Mayan civilization." Some examples of the "temporally relative" perspective are "last decade", "tomorrow", and "until now." The "spatially relative" perspective is one in which the reference point is spatially based. The compass directions- "north", "east", "west", "south"- are all examples of such. The last perspective is the "spatio-temporally relative" perspective, where the reference point is both temporally and spatially positioned.

We would like a schema capable of all perspectives, yet it would be redundant to represent the same situation in all of the perspectives. Additionally, changes in reference points such as the progression of current time would require recalculation of the latter three. We thus propose associating all entities within the knowledge base using one (preferably, but not necessarily the objective) perspective, and computing the other perspectives only when necessary.

We briefly discuss representation of such information in ontologies later in this paper.

Parameterizing time

We come to the problem of how to represent time and space within the ontology, and how to associate entities within time and space, and determine spatial location at any point in time, or vice versa.

Time and space are simply properties with which entities can be associated, and all physical entities are associated with a spatio-temporal region. That spatio-temporal region may be undefined or unknown. Certain entities that do not exist in space also have temporal ranges. The year "1999 AD" is seemingly a strictly temporal property, but then again, "1999 AD" is only relevant to Earth, and not to Mars.

While we agree with Galton about the entailment between motion and position, we disagree that any object is ever in a state of position. We make the simple observation that all objects are actually in a state of motion. Objects perceived to be at rest are merely those who have a zero velocity relative to another object. All objects on Earth are in constant revolution, and all objects in the universe are in constant expansion; states of position only occur with relation to a reference point. An attempt to represent entities as either in motion or position is thus self defeating if it does not accurately represent the state of the world as it truly is.

Similarly, the distinction between continuants and occurrents is trivial, especially the part where such distinctions depend upon the time range in consideration.

The approach necessary to store spatio-temporal information, given a 3 dimensional Cartesian space is to treat time as a parameter for a set of spatial functions, since time is easily conceived as a one dimentional unidirectional line. Thus for

(2) time t, and functions i(t), j(t) and k(t)
such that i,j,k respectively calculate the values of 3D space for any t, i.e.
$i(t) = x$,
$j(t) = y$,
$k(t) = z$.

These functions allow us to calculate the movement of the entity over a period of time, potentially over the entity's lifetime. Also, while we use the example of Cartesian space to illustrate the task of parameterization, it should be noted that the use of other reference systems is entirely possible. The latitude/longitude (latlong) system used in GIS systems is an example of an Earth-based polar coordinate system.

The above is an idealistic solution, and that it may not necessarily be possible for us to be able to determine these three functions for all (or even any of the) entities. Sometimes the only data available may be a sequence of values over a discrete set of time positions. If such is the case, one would then compile a list of observations at certain temporal instances and the corresponding values at those instances. When the observation frequency becomes high enough, we then have an accurate representative sample. For simple entities such as sound waves, achieving the Nyquist limit is enough. For more complex behaviors, the necessary frequency remains to be determined. The best we can do then, is to record all the information we have on any particular entity, and interpolate between these observances. For scientific observational systems such as climate monitoring, recording the extent of events at frequent, regular, periodic intervals over its lifetime probably allows us to chart its progress to an acceptable level of accuracy. Parameterization is general enough to accomodate for both continuous and discrete spatio-temporal models and datasets.

Ontology

We make use of an ontology to store the information of the objects with which we reason. Following (5), we define our ontology in a manner supportive of dynamic semantics, where entities are defined

(3) entity(attribute, value).

This is in contrast to the traditional relational definition

(4) entity$(a_1, a_2, ..., a_n)$ where $\forall i = 1..n, a_i$ is an attribute.

Table 1 illustrates both the mapping of relational structures to the mereological representation as well as the differences between the proposed mapping and traditional conceptualization within a relational schema.

Relations	Traditional	Proposed
Tables	Entity sets	Entities
Rows	Entities	Parts
Columns	Entity attributes	Meta-properties
R × C	Attribute values	<attribute, value> pairs

Table 1: Decomposition of relational schema- how the proposed differs from the traditional approach

We can add or remove attributes without changing our relational definition nor the structure of the queries by which we compute our semantics, allowing us to dynamically manipulate our ontology. We are also not bound to specify the categories to which each entity belongs, whether it be in motion or at rest, or a continuant or occurrent. Nor are we constrained to store the information in any specific spatio-temporal perspective.

One issue yet to be resolved is determination of the level at which we should associate the spatio-temporal values in the ontology. One way to support for entities to remain the same entity within our ontology despite changes to its parts (e.g. classical example of the cat that loses its tail is still the same cat) is to associate temporal ranges to each of its parts. A redefinition of our ontology to

(5) entity(attribute, value, valid_start_time, valid_end_time)

easily accomplishes this goal. For applications such as observations of continuous events at discrete points in time however or for cyclical events, this is not efficient. There is no need to redefine the ontology and add two additional pieces of information for each entity part, when all parts (and the entity iteself) have the same temporal range. Insertion of tuples

(6) <valid_start_time, T_{start}>
<valid_end_time, T_{end}>

would do just as well, but not accomodate entities with changing parts. Further work in this area is necessary.

Discussion

There are three other types of reasoning closely intertwined with STR as we have described it. They are qualitative reasoning, contextual reasoning, and inferential reasoning. Qualitative reasoning is another way to formulate the STR problem. Our products are directed towards the scientific community, and thus requires the rigidity of the spatio-temporal logics as given. Without contextual reasoning, STR would be tremendously difficult if not impossible. Inferential reasoning derives from STR abilities.

Qualitative reasoning

While we have only discussed this paper within the context of quantative reasoning, the proposed ontological structure is flexible and supports the storing of spatio-temporal values as either qualitative or quantitative values, or both. To enable a qualitative framework rather than the quantitative one we have described (which uses interval logic to determine how two locations relate to each other), we would store how locations relate to each other qualitatively, as opposed to how they both relate to a pre-selected point of reference. While the ontology is capable of storing these values, we would need additional knowledge to compute transitivity of relations, i.e. discover XR_3Z given

(7) XR_1Y and YR_2Z

Whether or not a system uses reasons qualitatively or quantitatively will depend upon the requirements of the system. General knowledge systems will more likely use the former, while scientific systems require the latter.

Contextual reasoning

There is a rich literature in contextual reasoning as well, so we shall keep our discussion of it brief. Contextual reasoning requires a theory of mind- an understanding by an agent that another agent discussing another topic than the one it itself is discussing. With regards to STR, that means an ability to differentiate between its current (spatio-temporal) location, the location being discussed by itself, the location being discussed by the other agent, and the location the other agent is located, i.e.

(8) AGENT-1-LOCATION ↔ TOPIC-1-LOCATION ↔ TOPIC-2-LOCATION ↔ AGENT-2-LOCATION.

Even if we initially that agent location and agent topic location are the same, we still need contextual reasoning, as we would have this model where there is still an incompatibility between the two agents.

(9) AGENT-1/TOPIC-1-LOCATION ↔ AGENT-2/TOPIC-2-LOCATION.

The location under discussion may be different than the location of either or both agents; we would not need contextual reasoning if it were otherwise. Both agents need to work towards the state such that each agent is in its own location, and the location in discussion can be any arbitrary location, including one of the two locations where the two agents are located, i.e.

(10) AGENT-1-LOCATION ↔ TOPIC-LOCATION ↔ AGENT-2-LOCATION.

Because we would like the ability to switch between perspectives, and because we are only using one perspective for ST entity representation, contextual reasoning provides with the ability to determine first that there is a difference in the locations and how one location relates to another.

Complexity We view the taxonomy of locations as a forest of trees, where children nodes have a finer granularity than the parent node. One path on such a tree may be

(11) Earth → North America → United States → California → California → Los Angeles → Union Square.

Once an agent determines that the other agent is in another agent (and for the sake of simplicity, presumes that each agent is in the context of its own location), the first agent will traverse the tree of locations where the leaves are coordinates. It traverses up the tree until it finds the two locations lowest on the tree having the same parent, each a distinct ancestor of their respective locations, and uses that as the basis for inferring relative perspective. The second agent also needs to realize that the first is using contextual reasoning, and thus should search for siblings of its ancestors.

For the generating agent, the first depth at which the context location differs between two agents is an O(log n) search with respect to the number of locations. For the second agent who needs to compute the relation between two locations given the context difference information from the first agent, it searches for all siblings of its ancestors. Expected running time should be O(ln n) to search to traverse up and down the tree. The worst running time should be O(n) if there are only n entities.

Inference resolution

Inferential abilities provide humans with the power to derive previously unknown facts from a particular knowledge base. Given the mapping

(12) $\{$ has(start-time, '[X]-12-31-23-59-50'), has(end-time, '[X+1]-1-1-0-0-0') $\} \subseteq$ NEW-YEAR-COUNTDOWN,

we associate it with an entity such as the New Year celebration

(13) $\{$ has(event, NEW-YEAR-COUNTDOWN) $\} \subseteq$ NEW-YEAR-CELEBRATION.

An inference engine, equipped with a rule such that

(14) $\{$ has(event, Y) $\} \subseteq$ X → overlaps(X, Y),

where the overlaps relation is defined as a temporal containment (1), immediately realizes that New Year's celebration lasts for at least 10 seconds, if not longer.

Conclusion

In this paper, we have reviewed the fundamentals to, identified the problem of perspectives and representation within, and proposed the use of a mereological approach to representation along with time parameterization as necessary components to STR. Additionally we have related STR to other types of reasoning- qualitative, contextual and inferential. The theoretical foundations laid out in this paper were made with the realization of the semantic web in mind, and are by no means complete. It is unquestionable that more work needs to be accomplished.

References

[1] J Allen. (1984) "Maintaining knowledge about temporal intervals", *Communications of the ACM*, 26(11), pp. 832-843.

[2] Juan Carlos Augusto. (2001) "The logical approach to temporal reasoning", *Artificial Intelligence Review*, Vol 16, No 4, pp 301-303.

[3] Antony Galton. (1990) "A critical examination of allen's theory of action and time", *Artificial Intelligence*, Amsterdam, v. 42 pp. 159-188.

[4] Patrick Hayes. (1996) "A catalog of temporal theories", *Technical Report UIUC-BI-AI-96-01*, The Beckman Institute, University of Illinois.

[5] Michael Pan, Rob Raskin. "Towards a set-theoretic approach for formalizing ontologies", *Proceedings of the 5th International Workshop on Computation Semantics*, Tilburg, Netherlands.

[6] John F Sowa. (2000) *Knowledge representation: logical, philosophical, and computational foundations*, Brooks/Cole, Pacific Grove, CA.

[7] Shu-Yao Chien, Vassilis J Tsotras, Carlo Zaniolo. (2001) "XML document versioning", *SIGMOD Records*, 30(3).