The Creeping Virtuality of Place

Inderjeet Mani

inderjeet period mani at gmail period com
sites.google.com/site/inderjeetmani

Places are inherently dynamic. The boundaries of administrative regions (e.g., provinces and countries) and natural regions (e.g., coastlines, ecosystems, mountains, ice sheets, tectonic plates) can alter over time, as can their names. Cities have a daily influx and egress of people and vehicles, buildings go up and down, streets sometimes get repaired, added, etc. Even when a place persists, its geographical constituents may change. Some of this information is clearly more volatile than the rest. Places also mediate between entities and events of significance to us, and space. As a result, places have a subjective quality. They reflect a network of associations, involving landmarks deemed salient for various reasons. These are all properties assigned to a place by a speaker, and may or may not correspond to the properties assigned to a place by any other speaker. The psychological maps of Milgram (1976) reflect this subjectivity, as do contemporary social-network based maps, e.g., TagMaps[1].

These properties of dynamicity and subjectivity present interesting challenges when producing mashups that align different data sources. Freeware tools such as MIPLACE (Mani et al., 2010), and commercial tools such as MetaCarta, PinPoint, etc., aim at interpreting absolute and relative references to places in free-form text and integrating them with maps. These tools can generate interesting mashups when combining their output, e.g., in GML[2], KML[3], etc., with other content about places. Consider the following extract from a travel blog[4], with times, events, and places highlighted based on the TimeML (Pustejovsky et al. 2005) and Spatial ML[5] standards.

> The following morning, as I was planning to ride farther that day, I left at dawn while Gregg and Brooks were still asleep. I biked 30 miles to the town of *Agua Azul* where I played for 4 hours in *waterfalls* and clean cool *pools*.

Figure 1 shows a snapshot from a Google Earth tour generated from the blog, from Plotnick (2008).

---

[1] TagMaps are derived from user-tagged Flickr photos that contain geo-locations. See tagmaps.research.yahoo.com .
[2] www.opengis.net/gml
[3] code.google.com/apis/kml/documentation
[4] www.rideforclimate.com/journals/?cat=3
[5] sourceforge.net/projects/spatialml

Dagstuhl Seminar Proceedings 10131
Spatial Representation and Reasoning in Language : Ontologies and Logics of Space
http://drops.dagstuhl.de/opus/volltexte/2011/2732

FIGURE 1: Google Earth Tour from Travel Blog

The question is how precise such an alignment of textual narrative and maps needs to be. Consider dynamicity. No historian would use a contemporary map of Asia to accompany Xenophon's *Anabasis* (where the time-frame is ca. 490 B.C.). While most gazetteers of place names used as a resource by language processing tools reflect only a current snapshot, it is undesirable for a system to miss references to Paphlagonia or Phrygia, or to assign an incorrect footprint for Pakistan and Palestine.

Now consider subjectivity. A particular narrator's view of New York from an airliner queued at 10,000 feet needs to be distinguished from that of a character he writes about who is standing facing west from the 10^{th} floor of the Empire State Building.

In order to address these twin problems, we need a richer level of indexicality for places, that we can use to reason with when aligning information. Let us define entities to be objects with identifiers and other properties, including, following Hornsby and Egehofer (2000), a history, namely a sequence of non-overlapping timestamps, representing different times when the entity is predicated to exist. Each timestamp will be represented as a time interval based on the interval calculus (Allen 1984), with possible metric constraints on the interval. Places are entities with spatial properties. The spatial properties will include its topological relationships to other places, represented in terms of RCC-8 (Randell et al. 1992, Cohn et al. 1997), or the 9-intersection calculus of Egenhofer and Herring (1990), as well as distance and orientation relations.

Entities such as people or vehicles can be treated as places with humanoid properties. They are located via spatial relations with other places. A particular speaker e1 located at p1 can assert at time t2 that a vehicle e2 was parked at time t1 < t2 at a place p2 west of place p3. We then have the following scenario in Figure 2, assuming particular lifetimes for each object. Let us allow p2 and p3 to cease to exist and then be reincarnated. Notice that we have left out the outermost embedding assertion -- the extradiegetic level, to use the narratological terminology of Genette (1980). Notice also that times are absolute, whereas places are asserted to exist for particular periods of time.

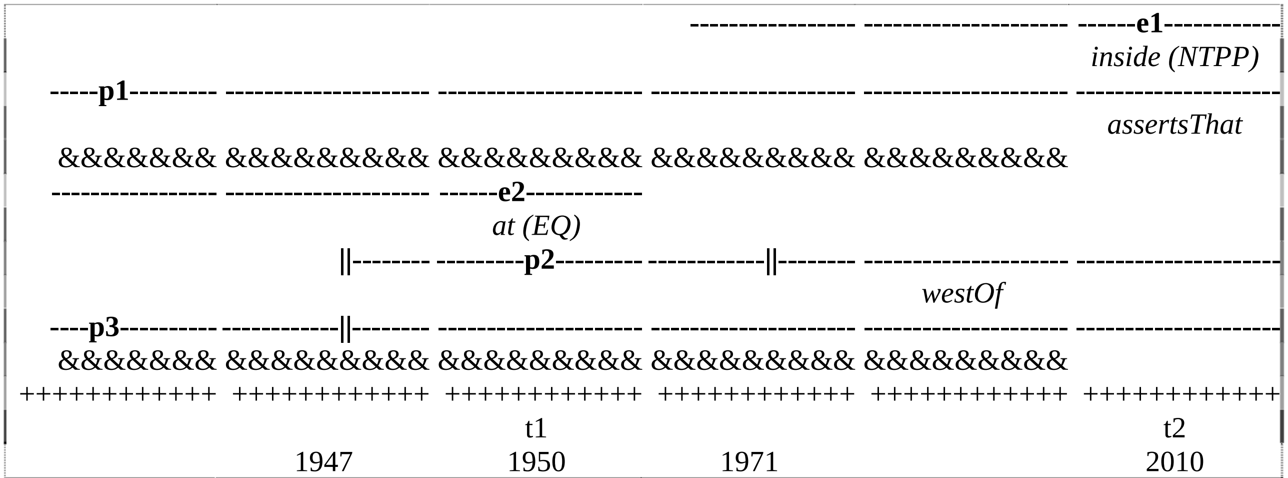

FIGURE 2 Lifetimes of Entities and Places At Different Diegetic Levels

To make the story of Figure 2 more realistic: sitting in a seminar room in Dagstuhl in 2010, I can assert that my grandfather who died before I was born traveled by ship in 1950 to Pakistan, which was created in 1947 after he was born, and was reconstituted (in 1971, minus East Pakistan which became Bangladesh) after he died. I can further assert that the reconstituted Pakistan is (wholly) west of India. I can also assert that India was reincarnated before I was born, while Dagstuhl existed only once.

The main differences between this scenario and reasoning purely within the interval and spatial calculi are the introduction of histories and the relationship of narration. Statements about histories, e.g., the lifetime and death of entities are accommodated via a relation e *existsAt* t. The relation of narration, that e *assertsThat* P, is more problematic. Here P is being treated as if it were a fictional proposition. Classical theories such as Lewis (1973) treat a fictional proposition P as a counterfactual conditional, true if and only if there is some possible world where (i) both the facts of the narrative as a whole are true and P is true, and (ii) that is closer to the actual world than every possible world where those facts are true and P is false. However, a vague appeal to similarity between worlds offers little help to a modal reasoner. While there has been encouraging progress in theorem-proving using the modal theory of abstract objects (including fictional ones) by Fitelson and Zalta (2007), far more research is needed to integrate this with spatio-temporal reasoning. It may be possible to consider *assertsThat* as introducing a branch in time, allowing us to exploit Computational Tree Logic, e.g., (Huth and Ryan 2004). Thus, one could state that for some path in time, there exists a past set of states where the proposition P is true.

References

Allen, James. 1984. Towards a General Theory of Action and Time. *Artificial Intelligence,* 23, 123-154.

Cohn, A. G., Bennett, B., Gooday, J. and Gotts, N. M. 1997. Qualitative Spatial Representation and Reasoning with the Region Connection Calculus. *GeoInformatica*, 1, 275–316.

Egenhofer, M. and Herring, J. (1990). Categorizing Binary Topological Relations Betwe Regions, Lines, and Points in Geographic Databases/ Technical Report, Department of Surveying Engineering, University of Maine, 1990.

Fitelson, Brian and Zalta, Edward. 2007. Steps Toward a Computational Metaphysics. *Journal of Philosophical Logic*, 36/2, 227–247.

Genette, Gérard. 1980. *Narrative Discourse*, trans. Jane Lewin. Cornell University Press, Ithaca, New York.

Hornsby, Kathleen and Egenhofer, Max. 2000. Identity-Based Change: A Foundation for Spatio-Temporal Knowledge Representation. *International Journal of Geographical Information Science* 14 (3): 207-224.

Huth, Michael and Ryan, Mark. 2004. *Logic in Computer Science* (Second Edition). Cambridge University Press.

Lewis, David. 1973. *Counterfactuals*. Harvard University Press. Cambridge, MA.

Mani, I., Doran, C., Harris, D., Hitzeman, J., Quimby, R., Richer, J., Wellner, B., Mardis, S., and Clancy, S. 2010. SpatialML: Annotation Scheme, Resources, and Evaluation. *Language Resources and Evaluation*, 44, 3, 263-280.

Mani, I., Wellner, B., Verhagen, M., Lee, C. M., and Pustejovsky, J. 2006. Machine Learning of Temporal Relations. *Proceedings of the 44th Annual Meeting of the Association for Computational Linguistics (COLING-ACL)*, Sydney, Australia, 753-760. ‘

Milgram, Stanley. 1976. Psychological Maps of Paris. In, Milgram, S. *Environmental Psychology: People and Their Physical Settings*, 2nd ed. Holt, Rinehart and Winston, New York, pp. 104–124.

Plotnick, A. 2008. Space-Time Annotation. Workshop on Methodologies and Resources for Processing Spatial Language at the Sixth international conference on Language Resources and Evaluation (LREC’2008), Marrakech, 31 May, 2008. http://www.sfbtr8.unibremen.de/SpatialLREC/

Pustejovsky, J., Ingria, B., Sauri, R., Castano, J., Littman, J., Gaizauskas, R., Setzer, A., Katz, G., and Mani, I. 2005. The Specification Language TimeML. In Mani, I., Pustejovsky, J., and Gaizauskas, R (eds.). *The Language of Time: A Reader*. Oxford University Press, New York, 549-562.

Randell, D. A., Cui, Z. and Cohn, A. G. 1992. A Spatial Logic Based on Regions and Connection, Proc. 3rd Int. Conf. on Knowledge Representation and Reasoning, Morgan Kaufmann, San Mateo, pp. 165–176.