STA: Spatio-Temporal Aggregation with Applications to Analysis of Diffusion-Reaction Phenomena*†

Iván Ordóñez
Dept. of Computer and Information Science
The Ohio State University
2015 Neil Avenue
Columbus, OH 43210
iordonez@cis.ohio-state.edu

Feng Zhao
Xerox Palo Alto Research Center
3333 Coyote Hill Road
Palo Alto, CA 94304
zhao@parc.xerox.com

Abstract

Spatio-temporal data sets arise when time-varying physical fields are discretized for simulation or analysis. The study of these data sets is essential for generating qualitative interpretations for human understanding. This paper presents Spatio-Temporal Aggregation (STA), a system for recognizing and tracking qualitative structures in spatio-temporal data sets. STA algorithms record and maintain temporal events and compile event sequences into concise history descriptions. This is carried out at several levels of description, from the bottom up, until a high level description of the system's temporal evolution is obtained. STA has been demonstrated on a class of diffusion-reaction systems in two dimensions and has successfully generated high-level symbolic descriptions of systems similar to those produced by scientists through carefully hand-tuned computational experiments.

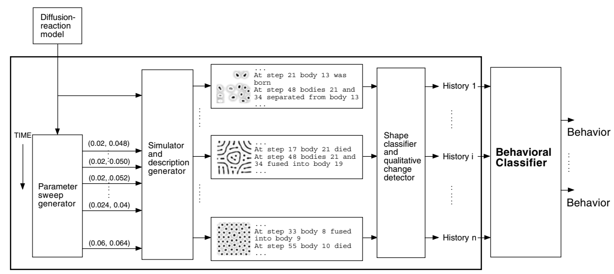

Figure 1: STA catalogs qualitatively distinct behavioral classes, represented as spatio-temporal patterns, for a diffusion-reaction system. The simulator generates multiple system evolutions, each corresponding to a different set of system parameter values and initial condition. Each evolution is compiled into an event history. The classifier identifies behavioral classes from the set of event histories.

Introduction

Spatio-temporal data sets arise when time-varying physical fields are discretized for the purpose of simulation or analysis. The study of these data sets is essential in scientific visualization, or generating qualitative interpretations.

A qualitative description of a physical field recognizes several events: the existence of coherent objects, their persistence through time, and their abrupt change. The classification of qualitative events based on topological and geometric characteristics of the involved objects and the nature of the transformations they undergo yields insight into the aggregated behavior of the system.

This paper describes Spatio-Temporal Aggregation, or STA, a temporal extension to Spatial Aggregation. This extension addresses systems that vary over time by recognizing and tracking structures in spatio-temporal data sets. STA is applied to a class of diffusion-reaction systems in two dimensions and it successfully generates high-level symbolic descriptions about the systems. In addition, by comparing multiple system histories, STA classifies systems with different parameterizations into equivalence classes, each of which contains members that exhibit qualitatively similar behaviors. This method is applied to the Gray-Scott (GS) model of glycolysis. It carries out an automated series of observations of temporal evolutions of this model, extracting a set of behavior-based classes of temporal evolutions. The approach has proved useful in that the classification scheme it generates is similar to one previously obtained by a scientist through carefully hand-tuned computational experiments and qualitative assessment by human observers (Pearson 1993). The operation of this application is sketched in Figure 1.

Other researchers have addressed the problem of generating high-level descriptions of physical systems. For instance, Williams and Millar (1996) develop a method for large-scale modeling and apply it to the thermal modeling of a smart building. STA is similar to their work in that it models complex systems through decomposition, but differs in that STA models more complex spatio-temporal dynamics, and produces symbolic descriptions. Crawford, Farquhar and Kuipers (1990) automatically generate qualitative differential equations from physical models. Their work considers temporal change, but not spatially distributed systems. Hornsby and Egenhofer (1997) study qualitative represen-

*This work was conducted as part of the research in the Intelligent Simulation Group at Ohio State. Support for the Group is provided in part by ONR Young Investigator award N00014-97-1-0599, NSF Young Investigator award CCR-9457802, a Sloan Research Fellowship, NSF grant CCR-9308639, and a Xerox Foundation grant.

†A longer version of the paper appears in AAAI'00.

tations of change, such as an object's continuation, separation and fusion, and construct hierarchies of change, but they do not attempt to apply these objects to continuous fields. Forbus, Nielsen and Faltings (1991) developed the CLOCK project, which uses qualitative spatial reasoning to automatically analyze and qualitatively predict the behavior of fixed-axis mechanisms, such as mechanical clocks. Their approach is suitable for mechanical systems of rigid parts, while ours is best suited for continuous fields that exhibit high-level properties such as quasi-uniform regions.

The main contribution of this paper is a computational system that analyzes very large sets of unstructured data to produce descriptions of qualitatively distinct aggregate objects and events.

Spatio-Temporal Aggregation

STA significantly extends the functionality of Spatial Aggregation (SA) in the temporal dimension. SA provides a uniform vocabulary and mechanism for representing and reasoning about spatial fields. For a full description of the SA field ontology and operators see Yip and Zhao (1996) and Bailey-Kellogg (1999).

Existing applications of SA abstract over domains such as phase spaces and configuration spaces, in which time is only implicitly represented. Others deal with physical spaces in a fixed, steady state. In all these cases the field, as an ontology, and all the conceptual layers built on top of it, are static. Problems that use time are not necessarily outside the domain of Spatial Aggregation. For example, KAM (Yip 1989) is used to study Hamiltonian systems, which describe frictionless motion. These systems are studied in phase space, where temporal variation is implicitly represented. More in general, SA could be used to study time-varying systems as simple static systems where time has been represented as an extra spatial dimension. On the other hand, STA offers, beyond such approaches, the ability to reason about time-varying systems without having to compute and store the entire space-time volume beforehand.

Sophisticated techniques have been developed to address the problem of temporal tracking in fields (Silver and Wang 1997). It seems natural to find whether there is a generalization of these tracking approaches, which would let them deal with not just one, but multiple abstraction layers.

The main addition made to the SA standard vocabulary by STA is the *update* operator, which takes a field or an object space and applies a set of transformations corresponding to the passage of a time interval.

- Updates on Neighborhood Graphs: For a set of objects S, a neighborhood graph is a relation $R \in S \times S$. When objects in space come into existence, cease to exist or change positions, their adjacencies may be modified.
- Updates on Object Classes: Changes in adjacencies may cause objects to cease to belong to a certain class or to start belonging to a new class. Classes are connected sets of objects; therefore, changes in R may affect classes. Also, changes in the intrinsic properties of the objects may also affect the way they are classified.

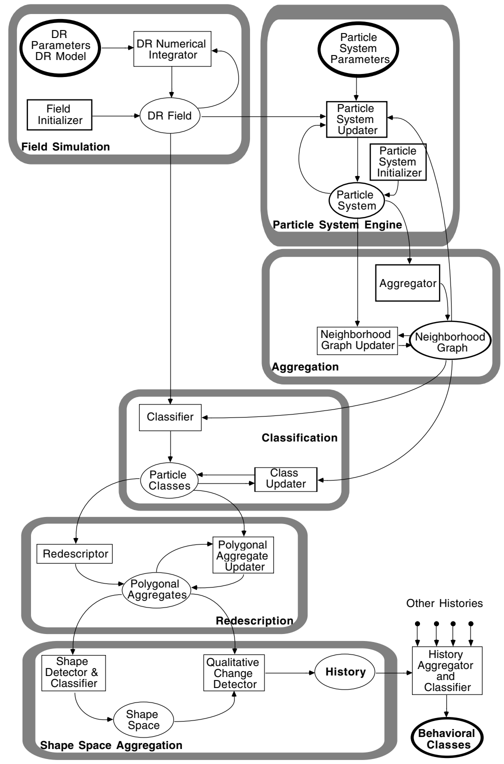

Figure 2: STA application to the analysis of diffusion-reaction systems. A field simulator generates system evolutions, which are tracked by the particle system. A chain of aggregation, classification and re-description is maintained to track high-level objects. Qualitative changes are registered into event histories. The history aggregator and classifier take multiple histories and identify behavioral classes.

- Updates on Re-described Objects: Changes in classes of objects may affect the way higher-level objects are re-described, depending on what features are kept in the re-description process and which are abstracted away.

Kinetic Data Structures: Reasoning about Change Detection

STA employs ideas from Kinetic Data Structures (KDS) to maintain the consistency of neighborhood graphs, object classes and re-described objects. KDS have been developed in robotics to maintain a set of geometric relations among distributed data (Basch, Guibas and Hershberg 1997). The problem KDS address consists of determining under which conditions the structure of certain geometric constructs is altered given that the elements are subject to particular motion laws.

Update Mechanisms

We enhance the static SA to include certificate-violation based update mechanisms adapted from KDS. This is done first at the neighborhood graph level, by associating the graph (namely, its vertices and its adjacencies) with a set of certificates that establish how much deformation the graph can take without undergoing a structural change. The classifier operator now does not only map objects to classes via the neighborhood graph, but it also maps graph changes due to certificate violations to class changes.

Application to Diffusion-Reaction Systems

We present a structure-identification algorithm for describing and classifying instances of diffusion-reaction systems that exhibit highly organized spatio-temporal structure. Figure 2 illustrates the operation of the algorithm.

Tracking High-Level Structures

The existence of coherent structures in a field implies that there are regions of approximately uniform characteristics. Once those regions are identified, characteristics such as topology and temporal behavior can be studied. The Field Simulation module (see Figure 2) generates the field and its changes, but is unaware of the existence of high-level structures.

Diffusion-reaction fields are sampled by the STA algorithm using particle systems (see corresponding block in Figure 2). Particles have the advantage of being persistent: they have discrete identities and hence whatever happens to them can be tracked in time with ease.

We use a simple algorithm that allows the particle system to adapt itself to changes in the field, always maintaining an adequate sampling. The algorithm is a modification of a method introduced by Witkin and Heckbert (1994). It allows particles to move across the field, repelling each other, thereby occupying space uniformly. Moreover, they modify their distribution and density to compensate for under or over-sampling.

The sampling particles are used to construct a spatial subdivision. The subdivision is computed by dividing the space into simplices whose vertices are the particles, and whose edges constitute a neighborhood relation for the particles. The simplices need to be small and non-sharp, so a Delaunay triangulation is used.

As the field varies in time, so does the position of the particles. This, in turn, causes the spatial subdivision to change: some edges cease to exist and some new ones arise at every time step. However, given the assumption that the underlying field changes slowly, the vast majority of edges and triangles are preserved through successive time steps, even though their shape is slightly changed.

The static construction of a neighborhood graph constitutes the *aggregation* operator in SA. The corresponding block in Figure 2 represents the enhanced STA aggregate operation, which maintains the neighborhood graph as the particle system changes.

Cluster boundaries are associated with field regions of high gradient. Those regions are identified using iso-lines, continuous zones of uniform or near-uniform field value. The particle placement algorithm previously described is used to approximate iso-line contours of uniform regions. This algorithm requires the ability to determine class equivalence between particles (the *classification* block in Figure 2).

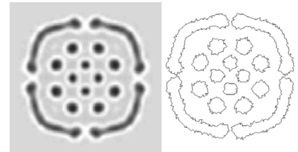

Figure 3: Subdivision generated from a particle system that samples a diffusion-reaction system

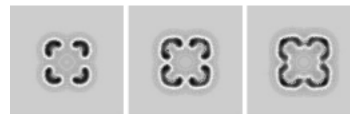

Figure 4: Successive snapshots of the evolution of a Gray-Scott diffusion-reaction system

The extraction of structures from the spatial subdivision is analogous to a pixel-based region growing algorithm, with the difference that the element of aggregation is not the pixel, but the sampling particle. The block that does this is labeled *redescription*. In Figure 3 the result of carrying out this process is exemplified.

STA records not only catastrophic events (such as object collisions), but also events that involve a single object modifying its shape. We use a shape-recognition and classification method called the Multiple Curvature Segmentation Algorithm, introduced by Dudek and Tsotsos (1997). Objects are placed in a shape space, and they are clustered by similarity. See the *Shape Space Aggregation* block in Figure 2.

Extracting Behavioral Descriptions

The STA algorithmic components we have described so far take as input a time-evolving diffusion-reaction system and produce the following descriptions:

- A detailed history of qualitatively significant events, including births, deaths, collisions and fusions of objects, and their changes in shape, and
- A summary of significant events that have taken place in the history.

The last two blocks of Figure 2 indicate the final summarization process of the STA application: multiple histories as

```
At step 88 body 3 was born
At step 88 body 2 was born
At step 88 body 1 was born
At step 88 body 0 was born
At step 229 bodies 0 (born 88), 3 (born 88) fused into body 1
At step 237 body 2 (born 88) fused into body 1
```

Table 1: A segment of a history: each entry is a time-stamped event. Notice that two fusion events are recorded. In them, the larger object preserves its identity, and the smaller ones are said to have fused to it.

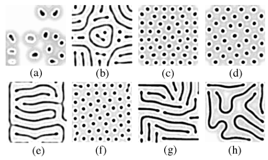

Figure 5: Snapshots for DR system evolutions

Cluster 1:	History (h)
Cluster 2:	Histories (e) and (g)
Cluster 3:	History (b)
Cluster 4:	Histories (c), (d) and (f)
Cluster 5:	History (a)

Table 2: Behavioral classes discovered by STA

generated above are compared, and then classified according to behavioral similarity.

A Sample Session

We now present a short run of the history-generation part of the program. It records the events that take place in an evolving diffusion-reaction field. For instance, when a system such as that shown in Figure 4 evolves, the program can generate a history file such as that of Table 1.

The program can also compare several histories and group them into classes of similar behavior. For the systems on Figure 5, the groups in Table 2 were discovered. Compare these with the classes discovered by Pearson (1993), shown in Figure 6: cluster 4 corresponds to pattern (b); cluster 2 to (c) and cluster (5) to (a).

Conclusions

This paper has described a novel computational system, STA, for reasoning about time-varying fields such as diffusion-reaction systems. STA extends Spatial Aggregation to make explicit the representation of time and temporal change.

STA has been demonstrated on a complex dynamical system that exhibits multiple, qualitatively different behaviors. This demonstration accounts for approximately 83% of the observations meticulously carried out by Pearson as documented in his 1993 paper. What this research contributes that had not been done before is the automatic differentiation of pattern classes by behavior.

STA makes use of various techniques, namely, operations of abstraction of change, kinetic data structures and geometric shape classification. How well would these techniques do if applied outside of this domain? We expect that a straightforward application of STA to problems that require extensive contextual and non-geometric knowledge would not work as well. For example, tracking objects for computer vision requires solving problems such as that of object occlusion and representation from incomplete information, not to mention the existence of multiple perspectives, different levels of illumination and reflectance, etc. In order to address those problems, STA needs to integrate additional domain specific techniques from computer vision. Similarly, the problem of examining weather patterns also requires extensive domain knowledge. While this problem seems more amenable to treatment from a STA perspective, it would still require integrating specific techniques such as those developed by Huang and Zhao (2000) with the STA tracking mechanism.

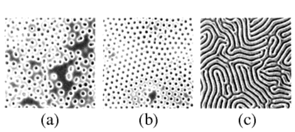

Figure 6: Patterns discovered by Pearson (1993) on the Gray-Scott system

References

Bailey-Kellogg, C. 1999. The Spatial Aggregation Language for Modeling and Controlling Distributed Physical Systems. Ph.D. diss., Dept. of Computer and Information Science, The Ohio State Univ.

Basch, J.; Guibas, L.; and Hershberg, J. 1997. Data Structures for Mobile Data. In *Proceedings of the 8th ACM-SIAM Symposium on Discrete Algorithms* 747–756.

Crawford, J; Farquhar, A.; and Kuipers, B. 1990. QPC: A Compiler from Physical Models into Qualitative Differential Equations. In *Proceedings of the National Conference on Artificial Intelligence (AAAI-90),* AAAI/MIT Press.

Dudek, G. and Tsotsos, J. 1997. Shape Representation and Recognition from Multiscale Curvature. *Computer Vision and Image Understanding* 68(2): 170–189.

Forbus, K.; Nielsen, P.; and Faltings, B. 1991. Qualitative Spatial Reasoning: The Clock Project. *Artificial Intelligence* 51(1–3): 417–471.

Hornsby, K. and Egenhofer, M. 1997. Qualitative Representation of Change. In *Proceedings of the International Conference on Spatial Information Theory (COSIT-97)*, Lecture Notes in Computer Science, 1329:15–33.

Huang, X. and Zhao, F. 2000. Relation-Based Aggregation: Finding Objects in Large Spatial Datasets. *Inter. J. of Intelligent Data Analysis.* To appear.

Pearson, J. 1993. Complex Patterns in a Simple System. *Science*, 261.

Silver, D. and Wang, X. 1997. Tracking and Visualizing Turbulent 3D Features. *IEEE Transactions on Visualization and Computer Graphics* 3:129–141.

Williams, B. C. and Millar, B 1996. Automated Decomposition of Model-based Learning Problems. In *Proceedings of QR-96.*

Witkin, A. and Heckbert, P. 1994. Using Particles to Sample and Control Implicit Surfaces. In *Computer Graphics, SIGGRAPH 94 Proceedings.*

Yip, K. 1989. KAM: Automatic Planning and Interpretation of Numerical Experiments using Geometrical Methods. Ph.D. diss., Massachusetts Institute of Technology.

Yip, K. and Zhao, F. 1996. Spatial Aggregation: Theory and Applications. *Journal of Artificial Intelligence Research*, 5.