

Figure 1: STA automatically catalogs qualitatively distinct behavioral classes, represented as spatio-temporal patterns, for a diffusion-reaction system. The simulator generates multiple system evolutions, each of which corresponds to a different set of system parameter values and initial condition. Each evolution, described as a sequence of qualitative events such as birth, death, or separation of objects, in conjunction with object shape transitions, is compiled into an event history. The classifier identifies behavioral classes from the set of event histories.

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 representations 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. Many other spatio-temporal reasoning systems cannot address such large systems because the sheer size of the data sets causes them problems such as combinatorial explosions. STA avoids such problems through intelligent decomposition and aggregation.




Figure 2: Three snapshots of a time-varying Gray-Scott diffusion-reaction system

A Case Study: Diffusion-Reaction Systems

An interesting instance of time-varying nonlinear dynamical systems is the set of phenomena known as Diffusion-Reaction. These phenomena are of great scientific importance, because they are associated with the problem of Morphogenesis, first addressed by Turing (1952). Particularly interesting instances, where noticeable patterns emerge and vary in seemingly unpredictable ways, will be examined.

The Gray Scott Model of Glycolysis

The phenomenon of glycolysis is found in virtually all living organisms. The Gray-Scott model of glycolysis is a diffusion-reaction system, characterized by the following equations:

$$\begin{aligned} \frac{\partial u}{\partial t} &= D_u \nabla^2 u - uv^2 + F(1 - u) \\ \frac{\partial v}{\partial t} &= D_v \nabla^2 v + uv^2 - (F + k)v, \end{aligned} \quad (1)$$

where u and v are concentrations of two reactants, D_u and D_v are their diffusion rates, and F and k are reaction parameters. This system is of interest not only as a model of glycolysis, but also because it exhibits a variety of behaviors unlike anything observed before in theoretical or numerical studies. Pearson (1993) first observed the strikingly varied patterns exhibited by the GS system, such as the one seen in Figure 2. Pearson et al. (1994) have argued that since glycolysis occurs inside the cell, it is possible that patterns such

as the above could form within it. Furthermore, they observe that the process of mitosis, through which cells divide, requires the formation of a bipolar structure known as the mitotic spindle, which is likely governed by simple physical processes such as chemical reactions and diffusion, rather than by complex genetic mechanisms.

As the parameters vary in the GS system, it undergoes qualitative transformations in its behavior. We apply STA to develop a program that can observe various system evolutions for different parameter values, and from this observation generate descriptions of the qualitative events that took place for each case. These descriptions are later used by the system to classify the instances into groups of similar behavior.

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. It builds a multi-layer, increasingly more abstract representation of a spatial field. Objects of each layer are formed as aggregates of lower-level objects. A neighborhood graph is constructed on the set of objects within each layer, and the objects are partitioned into equivalence classes with respect to their features, e.g., color, temperature, or pressure, as well their spatial adjacency. Each class is then re-described as a single object at the next higher level. The same process of aggregation, classification, and re-description repeats with more abstract relations at the next level. For a full description of the SA field ontology and operators see Yip and Zhao (1996) and Bailey-Kellogg (1999).

Temporal Changes

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. STA allows for the observation and representation of events as they happen, a feature that might be useful for real-time systems. For instance, our diffusion-reaction application, as sketched in Figure 1, records events such as the birth and death of spatial clusters in a diffusion-reaction field.

Aggregation and Persistence

Sophisticated techniques have been developed to address the problem of temporal tracking in fields (Silver and Wang

1997). It would seem 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, in the SA style.

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. This operation allows for changes in an object's features, position and existence, and it affects all levels of conceptual entities: objects, neighborhood graphs, equivalence classes and inter-layer mappings. The notion of update implies the premise that these conceptual entities are persistent. Thus, a neighborhood graph on a particular abstraction layer at time $t + 1$ should be conceived as a revision of the graph on that layer at time t , rather than as a new construct built from scratch. For instance, the dark areas in the fields seen in Figure 2 are objects, which may change in shape or position, while preserving their identities.

- Updates on Neighborhood Graphs: For a set of objects S , a neighborhood graph is a relation $R \in S \times S$ that does not contain any elements of the identity relation. When the objects in space come into existence, cease to exist or change positions, their adjacencies may be modified (thus changing R by removing elements from or adding elements to it). The changes in the neighborhood graph due to a change in a single object may remain localized in space, or may propagate everywhere, depending on the nature of the graph.
- Updates on Object Classes: Adjacency is a fundamental criterion to establish object equivalence in STA. Therefore, 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 (for any two elements in a class, there is a path between them made of elements of R); therefore, changes in R may affect classes. On the other hand, even if the adjacencies are not altered, changes in the intrinsic properties of the objects may also affect the way they are classified. Changes in classification are annotated as sets of objects added or removed from each class, as well as classes that are newly formed or newly deceased.
- 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. For example, if clusters of objects in space constitute classes and they are re-described as convex hulls, internal changes in the clusters do not affect the higher level objects as long as they do not involve the hull. Therefore it is necessary to have mechanisms that detect lower-level changes that affect the structure of higher-level re-described objects.

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.

Structural failure in a KDS is detected by maintaining a set of validity certificates, predicates that determine the conditions under which the current conditions of the system are valid. When a certificate is violated, an event is said to occur. The event is then processed and the certificates are updated to reflect the new conditions.

The existing corpus of research on finding and maintaining good certificates for various data structures is rich and varied. We are much less concerned with the particulars of each data structure and its maintenance algorithm than with the fact that such algorithms exist, and that they all fit within a single model of change as a violation of a certificate. Because STA addresses various levels of description, it is necessary to add conceptual mechanisms to determine the relevance of each certificate for the structure at the next abstraction level.

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.

The certificate-violation mechanism is extended to include the detection of non-geometrical change, namely, change in intrinsic object properties and existential change. The former kind of certificates exist at the classifier level, but not at the neighborhood graph level; typically it will consist of simple inequalities that test whether certain object features are within certain ranges. The latter exists at all levels. Also, detection at the re-description level requires being able to determine what low-level objects are relevant to the structure of higher-level objects. Because we know which objects are involved in which certificates, all certificates that contain relevant lower-level objects are needed for re-description. Such a filtering scheme is general, and allows for a unifying method to reason about abstraction of change.

Tracking Change in Time

We have developed a unifying reasoning scheme to deal with the propagation of change through an aggregation chain. We now focus on how to provide support for mechanisms that seek to interpret this change.

Keeping track of change in a system may be useful in many applications. For example, when studying transitional phases in self-organizing systems (such as the formation of

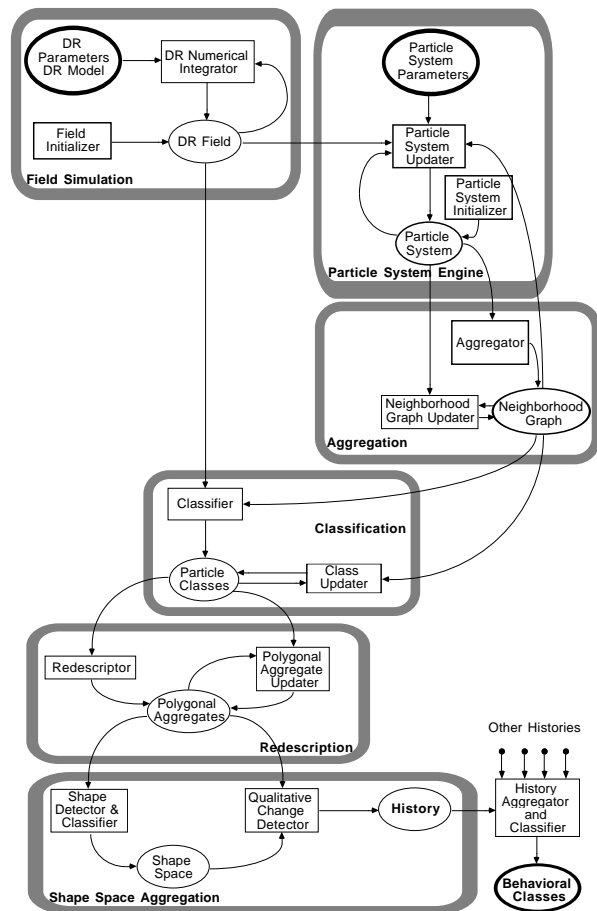


Figure 3: The flow diagram of STA application to the analysis of diffusion-reaction systems. The field simulator generates system evolutions, which are sampled and tracked by the particle system engine. A chain of aggregation, classification and re-description is maintained on top of the particle system, to identify and track high-level objects. Qualitative changes are later detected, generating event histories. The history aggregator and classifier take multiple histories and identify behavioral classes.

convection rolls in boiling water), researchers need to determine which kinds of events precede such transitions. We need, thus, to have a generic methodology to represent a history. More specifically, we require the ability to register a sequence of events that take place at various aggregation levels, namely changes in spatial objects, neighborhood graphs and object classes.

A history should register relevant change. The ease with which this can be done depends on how well the update mechanisms at various levels work. On one extreme, there is no attempt at any updates, and all structures are reconstructed from scratch at fixed intervals. In such conditions, finding relevant change is very difficult, since there is no knowledge to start with to draw correspondences. On the other extreme, there is a good update mechanism that operates on structures with a high degree of locality and that ex-

explicitly generates all change events. Such approach requires virtually no extra work from a history-tracking mechanism, which only needs to log these changes with their respective time stamps.

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. The algorithm is based on the central idea that qualitative structures of a spatial field can be constructed from an adaptive spatial subdivision rather than directly from a regularly discretized field. This adaptive subdivision changes as the field changes, but the identity of its structural components is persistent. The persistence of these components simplifies the correspondence between successive temporal snapshots. Figure 3 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. In the GS model, each region clearly belongs to one of two classes, low or high pH. The fewer the number of classes observed, and the larger the regions of uniform attributes are, the higher the organization perceived by an observer. These two global attributes of a field may vary with relative independence of each other, and each contributes significantly to the perception of coherence. For these reasons, local uniformity is one of the main features to look for when studying patterns. Once regions of uniformity are identified, characteristics such as topology and temporal behavior can be studied. The Field Simulation module (see Figure 3) generates the field and its changes, but is unaware of the existence of high-level structures.

Sampling Through a Particle System

Diffusion-reaction fields are sampled by the STA algorithm using particle systems (see corresponding block in Figure 3). Particles have the advantage of being persistent: they have discrete identities and hence whatever happens to them can be tracked in time with ease. Furthermore, any structures constructed by aggregating particles can also be tracked, because the identities of such constructs can be established recursively through a simple heuristic from the identities of its components. For example, one such simple heuristic is the following: if constructs A and B, existing at different time instants, share a majority of their components, they can be said to be identical. The particles must behave in such a way that they sample the field accurately. Therefore they must exist in large densities wherever the field gradient is large, and in low densities where it is small.

We consider 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. For this purpose, a Gaussian energy function is used. For any two particles i

and j , their mutual energy is

$$E_{ij} = \alpha e^{-\frac{|r_{ij}|^2}{2\sigma_i^2}}, \quad (2)$$

where α is a global constant and r_{ij} is their distance. The energy for each particle is given by

$$E_i = \sum_{j=1}^n E_{ij} + E_{ji}. \quad (3)$$

Particles are assigned a velocity that is negatively proportional to the gradient of energy, such that their local energy (for particle i , the part of its energy that does not depend on σ_j) is minimized. Moreover, they modify their distribution and density to compensate for under or over-sampling, by adaptively changing each σ_i to maintain the local energy of each particle constant, and splitting or dying when this parameter falls outside of a pre-defined range.

Aggregating a Particle System

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. It offers the added advantage that it can be computed efficiently in two dimensions. Also, this triangulation is a superset of the closest-neighbor graph, and therefore it captures the notion of spatial locality: local variations in a particle's position cause changes in the triangulation that do not propagate beyond its immediate vicinity.

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. Because of the local nature of the Delaunay triangulation, these updates do not propagate far.

The static construction of a neighborhood graph constitutes the *aggregation* operator in SA. The corresponding block in Figure 3 represents the enhanced STA aggregate operation, which maintains the neighborhood graph as the particle system changes.

Description through Iso-Lines

Cluster boundaries are associated with field regions of high gradient. Those regions can be identified using iso-lines, continuous zones of uniform or near-uniform field value. The ratio of field value change to the distance between iso-lines gives an estimation of the gradient. Therefore, a field that is characterized by near-uniform regions that vary smoothly is well described by iso-lines that sample evenly spaced field values. Temporal variations in fields will be studied through the examination of geometric and topological change in iso-lines.

The particle placement algorithm previously described is used to approximate iso-line contours of uniform regions.




Figure 4: Subdivision generated from a particle system that samples a diffusion-reaction system

This algorithm requires the ability to do two things: to determine class equivalence between particles (the *classification* block in Figure 3), and to evaluate distances between particles that take the field into consideration. Since this case study is of two-dimensional diffusion-reaction systems, the distances between particles can be computed by using the field values as additional dimensions, with an appropriate scaling constant. Class equivalence for adjacent particles is computed by thresholding the distance between the particles in feature space, that is, by considering only the values of their properties. 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 in Figure 3 is labeled *redescription*. In Figure 4 the result of carrying out this process is exemplified.

Keeping Track of Shape Changes

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. When an object moves from one shape cluster to another, a qualitative shape transformation is said to have taken place (see the *Shape Space Aggregation* block in Figure 3).

Putting it All Together: 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, specified as transitions from one shape cluster to another, and
- A summary of significant events that have taken place in the history, including records of the most common shapes and the most common events.

The last two blocks of Figure 3 indicate the final summarization process of the STA application: multiple histories as generated above are compared, and then classified according to behavioral similarity.




Figure 5: Successive snapshots of the evolution of a Gray-Scott diffusion-reaction system

```

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.

A Sample Session: Classifying Patterns According to Behavior

We now present a short run of the history-generation part of the program.

The program records the events that take place in an evolving diffusion-reaction field. For instance, when a system such as that shown in Figure 5 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 6, the groups in Table 2 were discovered. Compare these with the classes discovered by Pearson (1993), shown in Figure 7: 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 tem-




Figure 6: Snapshots for DR system evolutions. Histories were later classified into groups of similarity.

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 the STA application




Figure 7: Patterns discovered by Pearson (1993) on the Gray-Scott system. His classification is similar to that produced by the STA application presented in this paper.

poral change. For this purpose, various abstract operations were introduced to represent the notions of persistence and change. Common qualitative events such as birth, death, collision, separation, acquisition or loss of components or properties were identified for objects in spatio-temporal domains.

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. Our STA application classified multiple instances of this system using not only the appearance of static snapshots, but also the full extension of the behavior of the systems through a relatively long time interval. 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.

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.
- Beers, H. and Berkow, R. eds. 1999. *The Merck Manual*. Merck Research Laboratories, seventeenth edition.
- 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.
- Ordóñez, I. 1999. STA: Spatio-Temporal Aggregation of Physical Fields with Applications to Analysis of Diffusion-Reaction Phenomena. Ph.D. diss., Dept. of Computer and Information Science, The Ohio State Univ.
- Pearson, J; Doolen, G.; Hasslacher, B. Reynolds, B.; Horthemke, W.; Lee, K.; McCormick, W.; Swinney, H.; and Ponce-Dawson, S. 1994. Chemical Pattern Formation. *Parallel Computing Research*, 2(3).
- 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.
- Turing, A. 1952. The Chemical Basis of Morphogenesis. *Philosophical Transactions of the Royal Society (B)* 237:37–72.
- 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.