Journal of Knowledge and Information Systems, (11)5, 2007

Visual Transformation for Interactive Spatio-Temporal Data Mining

Yang Cai*, Richard Stumpf**, Timothy Wynne**, Michelle Tomlinson**,

Sai Ho Chung*, Xavier Boutonnier*, Matthias Ihmig*,

Rafael Franco* and Nathaniel Bauernfeind*

*Carnegie Mellon University, Ambient Intelligent Lab, 4720 Forbes Ave., Pittsburgh, PA 15213, USA, ycai@cmu.edu
** NOAA, 1305 East-West Highway, N/SCI1, Silver Spring, MD 20910, USA, richard.stumpf@noaa.gov

Abstract.

Analytical models intend to reveal inner structure, dynamics or relationship of things. However, they are not necessarily intuitive to humans. Conventional scientific visualization methods are intuitive, but limited by depth, dimensions, and resolution. The purpose of this study is to bridge the gap with transformation algorithms for mapping the data from an abstract space to an intuitive one, which includes shape correlation, periodicity, multi-physics, and spatial Bayesian. We tested this approach with the oceanographic case study. We found that the interactive visualization increases robustness in object tracking and positive detection accuracy in object prediction. We also found that the interactive method enables the user to process the image data at less than 1 minute per image versus 30 minutes per image manually. As a result, our test system can handle at least 10 times more data sets than traditional manual analysis. The results also suggest that minimal human interactions with appropriate computational transforms or cues may significantly increase the overall productivity.

keywords: vision, visualization, interaction, data mining

Received Oct 3, 2005
Revised Nov 15, 2006
Accepted Dec 6, 2006

1. Introduction

According to Constructionism, people not only discover things, but also construct them and interact with them. For over half a century, people have been developing computing systems for discovering lawful patterns in letters, numbers, words and images. The research has expanded into the computational study of the process of scientific discovery, producing such well-known classic AI programs as BACON (Langley, 1977) and DENDRAL (Linsay, Buchanan and Feigenbaum, 1993). However, *autonomous* discovery systems have rarely been used in the real world. While many factors have contributed to this, the most chronic difficulties seem always to fall into two categories: the representation of the prior knowledge that people bring to their tasks, and the awareness of new context. Many difficult scientific discovery tasks can only be solved in *interactive* ways, by combining intelligent computing techniques with intuitive and adaptive user interfaces. It is inevitable that human intelligence is used in scientific discovery systems. For example, the human eyes can capture complex patterns and detect anomalies in a data set. The human brain can easily manipulate perceptions (shape, color, balance, time, distance, direction, speed, force, similarity, likelihood, intent and well-being) to make decisions. The process consists of *perception* and *interaction* and it is often ubiquitous and autonomous. We refer to this kind of intelligence as *Ambient Intelligence (AmI)* (Aarts and Marzano, 2003; Cai, 2005).

Ambient Intelligence is originally referred to smart electronic environments that are sensitive and responsive to the presence of people. Since its introduction in the late 1990s, the field has matured and expanded into cognitive aspects of interactions. Ambient Intelligence is about human interaction with information in a way that permits humans to spot interesting signs in massive data sources - building tools that capitalize on human strengths and compensate for human weaknesses in order to enhance and extend discovery capabilities. For example, people are much better than machines at detecting patterns in a visual scene, while machines are better than people at manipulating streams of numbers.

1.1. Studies in interactive visualization

The human-computer interactive discovery systems have risen from many professional applications. Zudilova and Sloot investigate the practical deployment of virtual reality systems in the medical environment (Zudilova and Sloot, 2005). They explore the multi-modal interaction of virtual reality and desktop computers in Virtual Radiology Explorer. Pryke and Beale present their interactive data mining system that helps users gain insight from the dynamically created virtual data space (Pryke and Beale, 2005). Cowell et al. present the architecture of a next-generation analytical environment for scientific discovery within continuous, time-varying data streams (Devaney, et al, 2005). Devaney et al. develop a high-end CAVE-based virtual laboratory for scientific discovery. By minimizing the time between the generation of a scientific hypothesis and the test of that idea, It enables scientific investigations at the speed of thought. Cai et al. (Cai, et al, 2005) develop a game-based interactive scientific visualization system for interdisciplinary biological discovery and education.

Understanding how people sense, understand, and use images and words in everyday work and life can eventually help us design more effective discovery

systems. Hubona and Shirah investigate how various spatial depth cues, such as motion, stereoscopic vision, shadows, and scene background, are utilized to promote the perceptual discovery and interpretation of the presented imagery in 3D scientific visualization (Hubona and Shirah, 2005).

1.2. Studies in spatiotemporal data mining

Spatiotemporal data are widely used in the remote sensing community. The imagery data usually contain multi-spectrum information such as visible, infrared, UV, or microwave signals. Scientists need to find out the physical properties behind the data flow and predict the future trends such as global warming, flood, harmful algae blooms or river plumes. Spatiotemporal dynamics is not only about frames but also multidimensional interaction and evolution. Spatiotemporal data mining is often a computationally hard problem. For example, a co-location mining problem for over 10,000 locations would take over several hours to solve by a computer. In many cases, the inverse physics process is complicated. Heuristics has to be used for approximate estimations.

Contemporary data mining methods intend to discover patterns without domain knowledge. Many models are based on statistics and geospatial applications. For example, the Bayesian model is used to geographically profile serial killers; neural networks are used to model urban development dynamics, where the spatiotemporal patterns and space-time transitions are extracted. Spatial auto-regression is developed for spatial data mining. Furthermore, Variogram is used for spatial variation analyses. However, spatial data are often sparse and noisy. For example, normally over 80 percent of optical remote sensing satellite images over Florida contain clouds so that the objects beneath remain unseen. There are many artifacts in the images due to the noise in data acquisition and processing. Furthermore, our knowledge about the spatial dynamics is often limited. Therefore, a fully autonomous data mining is nearly impractical. Human intervention is necessary. It is desirable to use computers' numerical power and humans' pattern recognition capabilities.

Spatiotemporal data mining is also an induction process that derives conclusions from massive data. Models, such as regression, neural network, and statistic models, can be extended to spatiotemporal domains. These models are quantitative and programmable. However, it assumes that the future is what it used to be. As a result, they only work in a well-controlled environment with 'clean data'. They are better for short-term prediction, because they need less knowledge or parameters.

Physics models are developed for augmenting dynamic objects in images. For example, deformable finite meshed templates are used to model faces (Sonka, et al, 1998). Particle filters (Zhou, Chellappa and Moghaddam, 2005) are used to describe the spatiotemporal feature distributions as particle dynamics. Particle filters assume the particle distribution follows a statistical pattern and a linear or non-linear transformation trajectory. Unfortunately, many spatiotemporal data include overwhelming noises and missing data that prevent using statistical assumptions properly. Object tracking sometimes fails due to the discontinuity of the patterns.

Field theories are developed to represent cognition dynamics with physical analogy. For example, Leyton discovers that shape is energy (Leyton, 1992). His shape process theory reveals the rules for recovering original shapes based on

symmetry and other properties. Lewin (Bandini and Pavesi, 2002) uses metaphors from well-known systems (e.g. physics), in order to understand and explain a rather unknown system (e.g. psychology). In explaining phenomena around the concept of behavior, Lewin uses the metaphor of a 'force field' such as psychological tension, psychological force, and fluidity.

Studies of Cellular Automata (CA) suggest the possibility of modeling artificial life forms (Hardy, et al, 1976; Toffoli and Margolus, 1987; Wolfram, 2002; Wolfram, 1984*a*; Wolfram, 1983*a*; Wolfram, 1983*b*; Wolfram, 1985; Wolfram, 1984*b*), as well as at the large scale, such as urban and economic dynamics (Gardner, 1980). Furthermore, CA has been employed in simulating and verifying various types of physical interactions, such as the expansion patterns of different lattice gases, surface tension of liquids, and diffusion-limited aggregation (Essam, 1980). Although CA's rules are simple and local, it can simulate complex spatiotemporal interactions that are challenging to other computational methods.

1.3. Objectives of this study

Analytical models intend to reveal inner structure, dynamics or relationship of things. However, they are not necessarily intuitive to humans. Conventional scientific visualization methods are intuitive, but limited by depth, dimension, and resolution. To bridge the gap, transformation algorithms are needed to map the data from an abstract space to an intuitive one. For example, a spectrogram maps an invisible sound to a visible frequency-intensity-time space. The convergence of scientific visualization and data mining creates a new domain for visual data mining. Sensing and understanding together enables humans to gain insight from a large amount of data. This approach integrates a human's Ambient Intelligence and analytical computation to form a coherent knowledge discovery environment.

The goal of this study is to develop a set of visual transformation algorithms that incorporate computer vision, multi-physics simulation, and machine learning models to solve the tracking and prediction problems across multiple databases. Embedding vision into the conventional data mining algorithms enable us to automate the data pre-preprocessing and increase accuracy in the overall process. In contrast to the conventional data mining algorithm that is normally context-free and independent from the physical constraints, we introduced a visual multiple physics model such as Cellular Automata to simulate the physical, biological, and chemical phenomena. This provides a mechanism to explore the deep knowledge about spatiotemporal patterns while the accurate computational models are not available or impractical. For example, the current data assimilation for weather models takes several hours so that it prevents real-time weather forecasting. Machine learning has been a key feature in many data mining algorithms. In this project, we extend learning algorithms to spatial and temporal space and integrate them with vision and simulation models. As learning models can be data mining tools alone, they also support vision and simulation models for calibration or optimization of parameters. Moreover, human-computer interaction and scientific visualization methods are introduced as a comprehensive interface, which increases the robustness of the process.

The spatiotemporal data mining process involves multi-physics phenomena and domain-oriented knowledge discovery. No single visualization method or sin-

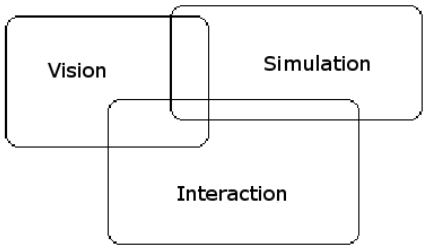

Fig. 1. Illustration of the system architecture that consists of computer vision, multi-physics simulation and user interaction

gle interaction can cover the whole process. Instead, multiple visualization methods and human-computer interactions are embedded inside the complex process. To bridge the analytical models and human visual perception, we have developed a platform of the *visual transformation* models that enable the two-way visual interactions, which are subtle but profound.

In this paper, we present general interactive cellular automata for tracking and predicting the spatiotemporal dynamics in an imagery flow. The model consists of a set of parameters and rules that control spatiotemporal dynamics in the form of a shape, such as diffusion, transport, and collision. In this model, cells not only interact with their neighbors *locally*, but also contain the *holistic* shape of a cluster of cells as an object or a *lump*. Figure 2 shows an illustration of the system architecture where the cellular automata model consists of Bayesian inference, artificial life rules, and multi-physics rules such as multi-body collision, etc. The visualization functions include fuzzy blobs, isosurfaces, correlation values, periodicity transform, animation and polarized glasses. On the other side, the micro interactions include shape adjustment, parametric control, rotation, pan, zoom, cut, and slide windows.

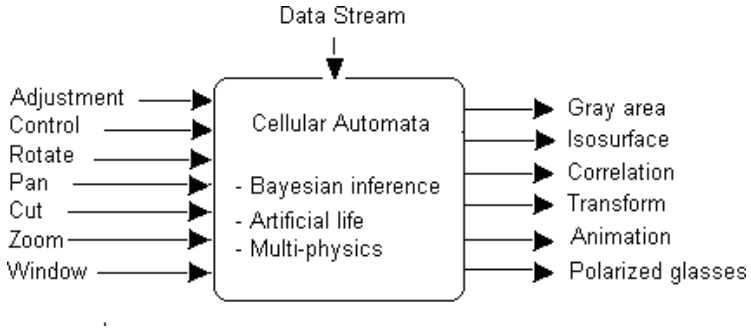

Fig. 2. Visual transformation for interactive data mining

In this paper, we aim to develop computational algorithms for tracking and predicting the movement of objects from images and numbers. For tracking problems, given an object in an image sequence (t=1,..,n), we try to find the object at t = n+1 and beyond. For prediction problems, given databases of historical data and current physical and biochemical conditions, we try to predict the

occurrence of the object of interest at a particular time and location. Figure 3 illustrates a generalized spatiotemporal data mining problem.

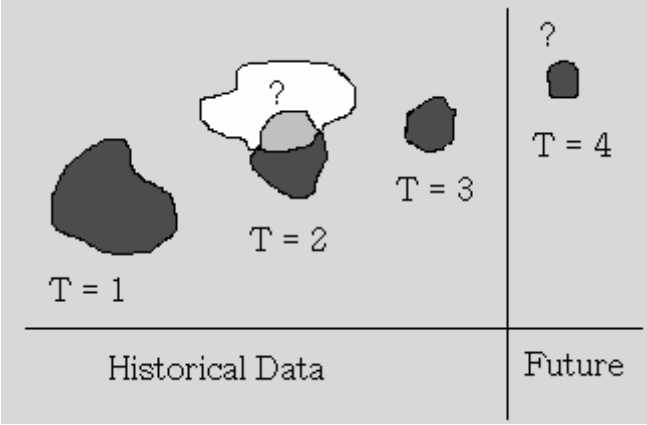

Fig. 3. Illustration of the spatial tracking and prediction problems.

To validate the methods developed in this project, we conducted one case study with the real world databases: tracking and prediction of harmful algal blooms in the Gulf of Mexico area.

2. Shape Correlation Transform

Given a selected object in a very large imagery database, tracking the movement of the object over time with multiple sensors is a challenging task for both humans and computers. For example, it takes about 30 minutes for scientists to analyze one frame of the multi-spectrum satellite image. On the other hand, computer vision-based tracking is faster, but they are not very reliable in many cases. The goal of the interactive object-tracking is to combine the machine vision with human-computer interaction so that the computer takes care of most of the time-consuming iterative task and humans contribute the minimal initiation and evaluation processes.

The interaction protocol is developed for human-computer interaction during the tracking task. In the protocol, the user identifies the tracking target. The computer takes over the tracking frame by frame and superimposes the tracking results on top of the original images. At a certain point, the computer stops the tracking and sends the signal to the user for more information. After the re-initiation, the computer moves on till the next break point. The key element in this protocol is the context-awareness algorithm. In this study, we use the correlation filter to measure the similarity of the object in the current frame and the object in the previous frame. The correlation filter is built by using Fast Fourier Transform and Correlation Theorem:

$$C = IFFT\left(FFT(a) * conj\left(FFT(b)\right)\right) \tag{1}$$

where a is the test image while b is the reference object in the previous image

to be tracked. $X = FFT(x)$ and $x = IFFT(X)$ represent Discrete Fast Fourier Transform and Inverse Discrete Fast Fourier Transform respectively.

$$X(k) = \sum_{j=1}^{N} x(j)\omega_N^{(j-1)(k-1)} \tag{2}$$

$$x(j) = \sum_{j=1}^{N} X(k)\omega_N^{-(j-1)(k-1)} \tag{3}$$

where $\omega_N = e^{(-2\pi i)fN}$, N denotes the length of the vector and symbol. The symbol $*$ denotes array multiplication and $conj$ returns the complex conjugate of the argument. The pixel that gives the highest correlation value has the highest confidence that the reference object is at that pixel location. The correlation filter alone cannot reliably track the object over a long period of time because it faces a dilemma when the object breaks into pieces. The computer would select one of the pieces and track it. The scenario probably is not what we want for object tracking due to lack of control of the process. Figure 4 shows the scenario where the object breaks into two pieces and the computer is attracted to the small piece at the fourth frame.

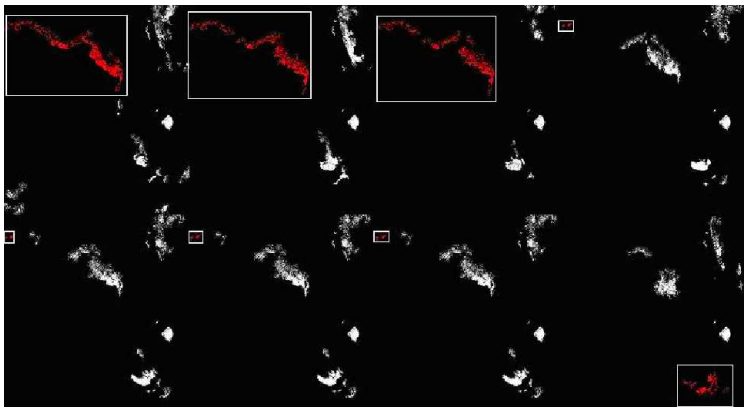

Fig. 4. Automatic tracking of a target which has split into 2 pieces.

In the interaction protocol, object tracking using correlation filter is based on the similarity of the objects between images. If there is a large difference in the values given by the correlation filter between two consecutive images, we can conclude that the objects have changed significantly during this time period. However, we cannot determine whether the object in the previous frame has disappeared in the subsequent frame or if the object has just deformed significantly. Scientists' experience and intuition would help determine if the object still appears in the next image in this case. We performed an experiment to track an object in a series of 84 images. In frame 48, the correlation dropped by more than 50% from the previous frame. Originally, the algorithm would stop at this point because it cannot identify the object being tracked. However, instead of terminating the tracking, the computer prompted the user to select a particular object in the new image so that the tracking could go on. In our experiment, the user selected an object that was most alike the previous object

and the tracking progress continued until the end of series. There is also another application that human interaction may help. When the object breaks into more than one piece, the original algorithm would only try to find the single best match of the previous tracking object. With the interaction method, the user can select more than one object and multiple targets will be under the tracking process.

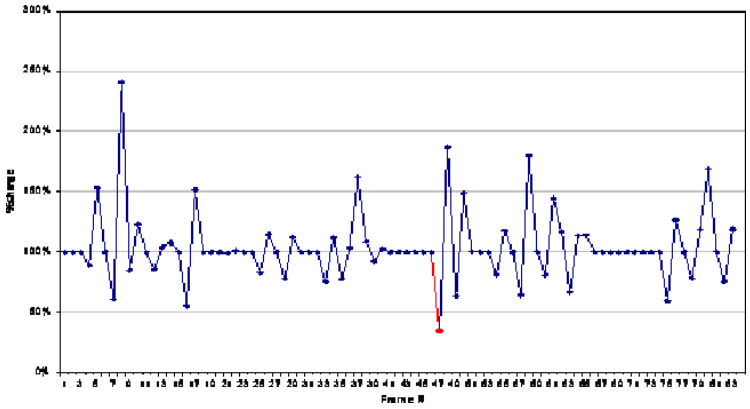

Fig. 5. Correlation Value indicates the 'breaking' point at the 50% level.

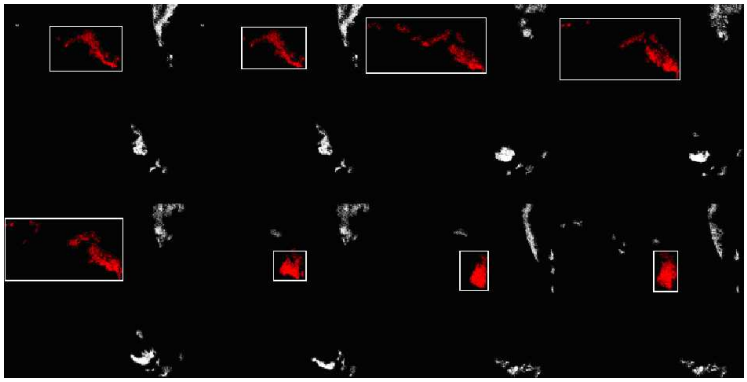

Fig. 6. Interactive tracking results.

Table 1. Automated versus Interaction.

Case	Automatic	Interactive
acceptable % for the coherent target	73/79 = 92%	73/79 = 92%
acceptable % for the target split to two pieces	48/79 = 60%	71/79 = 90%

3. Spatial Periodicity Transform

Although humans are capable of identifying spatial temporal and patterns, such as the repetitions in textures, and the rhythms in music, human perceptions

often miss the hidden patterns if the data set is complex. Computational transforms are needed for semi-automatically identifying underlying patterns. Perhaps the best known is the Fourier Transform, which attempts to explain a data set as a weighted sum of sinusoidal basis elements. However, none of the available methods directly search for periodicities, repetitions, or regularities in the data. Sethares builds on a classical technique, called the Buys-Ballot table, for the determination of hidden periodicities, the Periodicity Transform (Sethares and Staley, 1999). In his method, the algorithm searches for the best periodic characterization of the length N signal x. The underlying technique is to project x onto a periodic subspace.

In this study, we expand the one-dimensional Periodicity Transform into the two-dimensional space of the cellular automata. We construct spatial and temporal windows to allow the user to select the region of interest (in the red box of longitude and latitude coordinates), on a particular time frame. The algorithm starts with cropping the data within the ROI (Region Of Interest) box and sorting the data by time. Then it takes the Fast Fourier Transform (FFT) and extracts the highest value. For each point, we take the average value of the corresponding time. We then simply calculate the remaining signal and reprocess it. The outputs include the detected periods and their energy. This method enables user to interact with spatial periodicity transform in an interactive mode, which is very important to many applications such as oceanographic studies.

From the Fourier Transform, we select the frequency of the component with the highest energy. Then we turn this frequency (f) into a period (p) with the formula:

$$p = \frac{1}{f} \qquad \text{for all points } x_i \tag{4}$$

The equation of the power (energy) is the following:

$$Power = \|x\|^2 = \frac{1}{p^2} \sum_{i=0}^{p^2-1} x(i)^2 \tag{5}$$

There are many ways to render the spatial periodicities. First, we use two-dimensional colored cells to represent the power (energy) for a particular period component, e.g. 365.25 days. In this case, a user needs to select the period component as an input. Fig.7 illustrates an example of the two-dimensional power ratios in a database of *Karenia Brevis* cell counts in the Gulf of Mexico for 50 years. The brighter the cell, the higher the power ratio, which means that the probability of finding the target every year at the same time is higher. Due to the heterogeneity of the data set, pre-processing is necessary. That consists of computing a space mean of the number of cells and the estimation of missing data spots where we do not have any measurement. Based on the periodicity transform and human-computer interaction, we can see on Figure 9c that for each set of data we have a one-year period component.

Another two-dimensional rendering is to tile the individual period spectrum diagram into a two dimensional matrix. Each cell has its periodicities along with their powers. This matrix is called cellular periodigram and is represented in Figure 8. The advantage of this method is its all-in-one solution. However, its scalability is limited by the number of cells per screen.

Finally, we created the visual interaction interface for the spatial periodicity transform. We assign an interactive window for two-dimensional data points. In

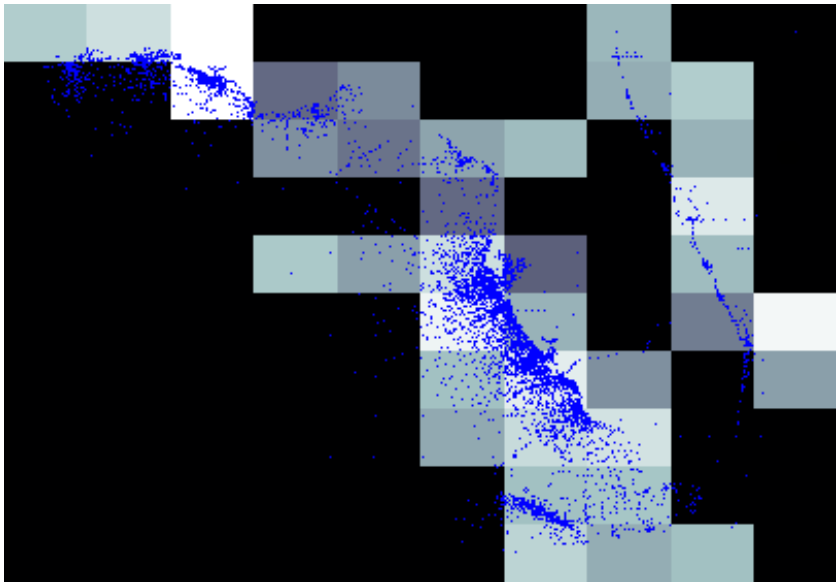

Fig. 7. Map of the power ratios with measurement points overlayed. The lighter the background, the more powerful it is in the 365.25 day period.

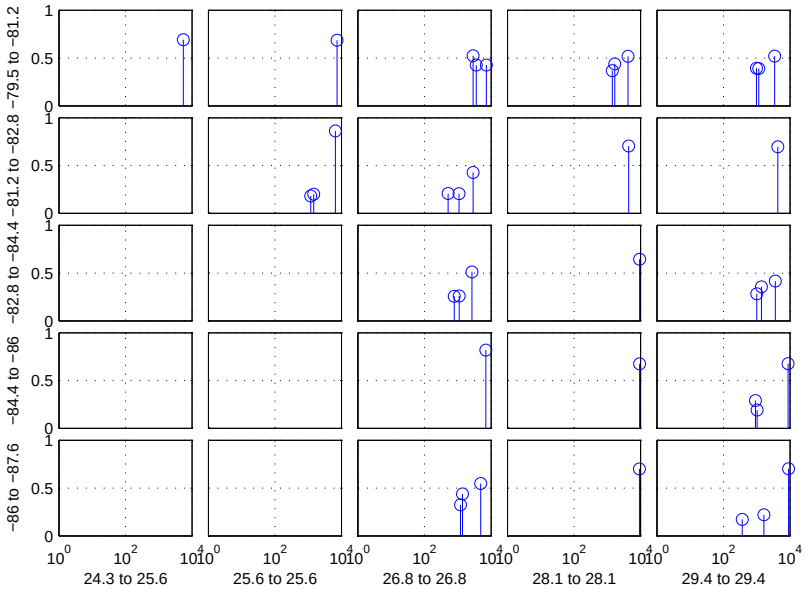

Fig. 8. Spatial based Frequency Analysis

another window, we display the associated results of the periodicity transform, where the x-coordinate is the period and the y-coordinate is the power ratio of the component. Figures 9a through 9c are examples of the interactive data mining results. The data is drawn from 50 years of the *Karena Brevis* cell counts from Florida coast. Each piece of data contains the geographical coordinates and level of concentration.

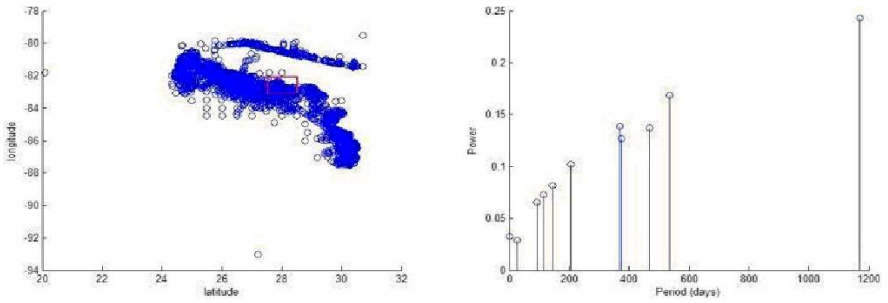

(a) All samples with narrow time window and the periodicity detection.

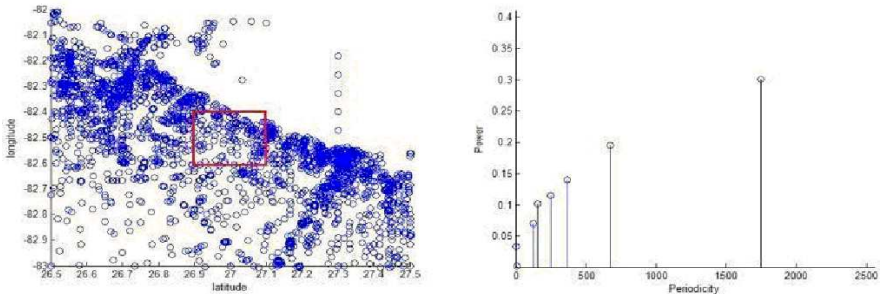

(b) Median window (ROI) selection and the periodicity detection .

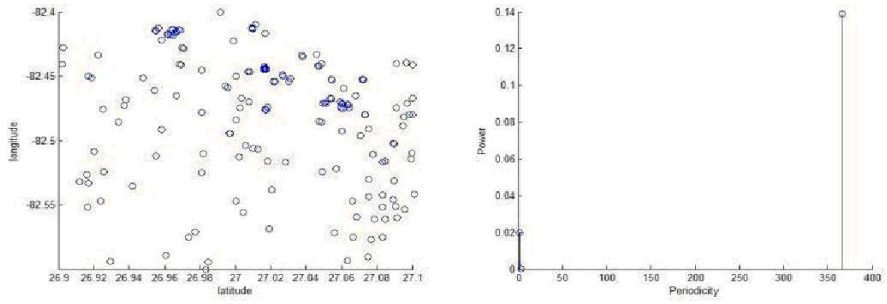

(c) Small window (ROI) selection and the periodicity detection.

Fig. 9. Different zooming windows for the analysis of periodicity

4. Multi-Physics Transform

Conventional data mining models do not involve physical parameters. On the other hand, traditional physical models only deal with particular phenomena

with required calibration process. In this study, we try to merge the two with a simple multi-physics model and an interaction process that overlays the results from the simulated model on top of the collected data.

We build the simulation model for predicting the spatiotemporal dynamics. Given the results obtained from the previous sections, we create scenarios of object movements by modeling the artificial life, with diffusion, growth, shrink, transport and collision. The growth diffusion function changes the object shape by a predefined structuring element. The dilation or erosion morphology is based on the assumption of a glass-like surface where no friction exists. They can only deform the outline of a shape *uniformly* based on the population. It does not reflect the intrinsic, external non-linear, or random factors that contribute to the shape process. In order to simulate the non-uniform shape growth, we investigate percolation clustering (Essam, 1980). Figure 10 shows an example of the percolation process started from a given shape.

In the real-world, it is rarely the case that cells are allowed to grow or move uninhibitedly in an infinite amount of space; rather, their growths and movements are often affected by environmental factors such as external forces: like winds and currents, and the presence of obstacles. A collision causes deformation. The extent of a deformation depends on the elasticity of an object. For a rigid object, it will stop at the collision contact point and keep the same shape. We classify objects as rigid, elastic, and flexible. The two extreme cases are straightforward. Rigid objects don't deform after the impact. Flexible objects completely deform along the outline of the collided hard objects. For the case of an elastic object and a rigid object, we use the following rules to model the deformation:

1. *Detect the collision points*
2. *Move a portion of the collided cells to sideways, the ratio is proportion to the elasticity of the object. The bigger the elasticity, the more the cells move sideways.*
3. *Continue to deform until the motion stops.*

Figure 11 shows a simple dynamics of growth and collision of the artificial life, which starts at the left frame with a very small region. Through some internal factors, the life grows larger. But once the growth hits the wall, the growth spreads to the side starting from the third frame. The frames above were generated using a Cellular Automata simulation with a relatively low elasticity.

To simulate the dynamics of artificial life, diffusion using Cellular Automata is implemented for the growth and shrinking of cells, collision of cells onto a boundary, and wind translations. The simulation rules are listed in the Appendix.

The surface current can also be combined with the Bayesian prediction model to analyze the dynamics of artificial life. So far, Bayesian prediction model is a better prediction than most other models because of its simplicity.

The CA simulation model can then use the predicted output and simulate future shape deformations, collisions, and wind translations using surface current conditions, such as temperature which would affect how the shape deforms, grows, and shrinks. The dynamics are detailed by simulating predictions and comparing them to the actual movement of artificial life. The error of simulation compared to actual movements will allow humans to better understand the dynamics of the internal and external factors of artificial life. With better understanding, the CA model can be improved to incorporate more rules dealing

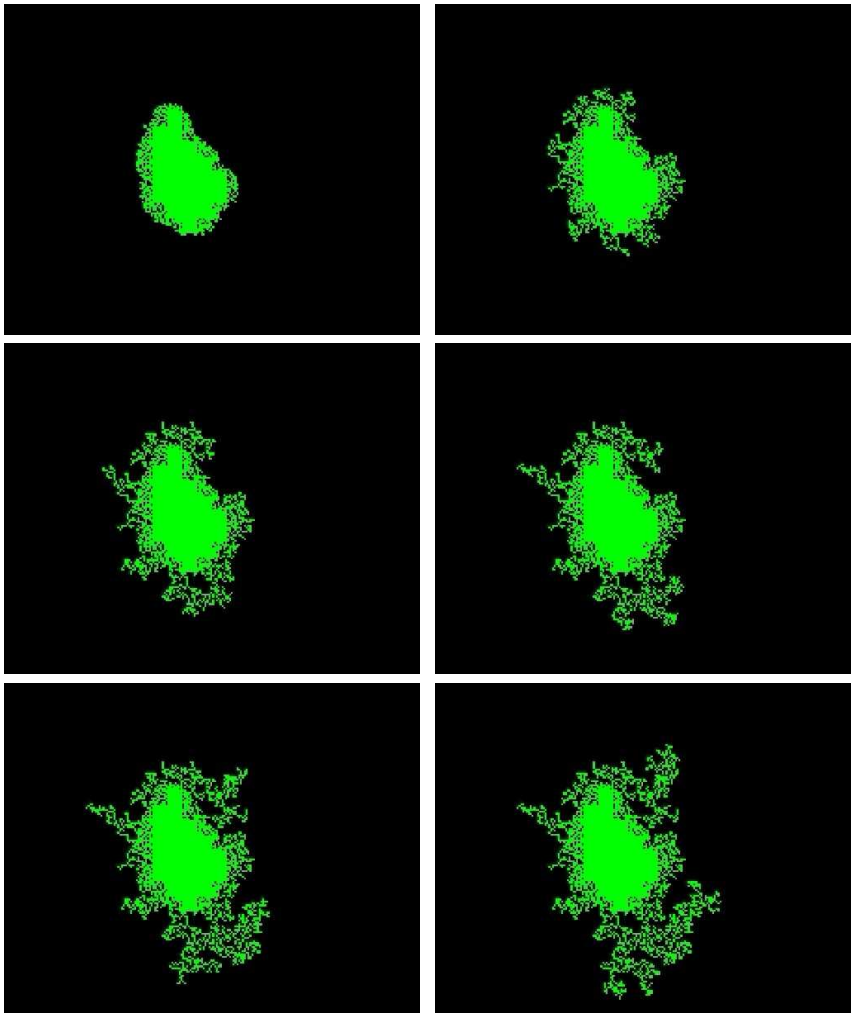

Fig. 10. A simulated result of diffusion with resistance. The image is projected onto the dynamic data and calibrated by visual interactions

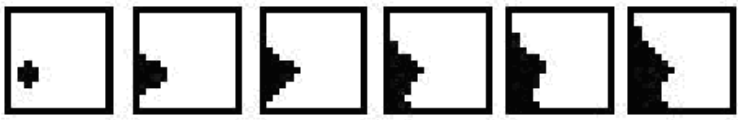

Fig. 11. A simulated result of the cellular growth and collision. From left to right, the cells grow and deform at the left border.

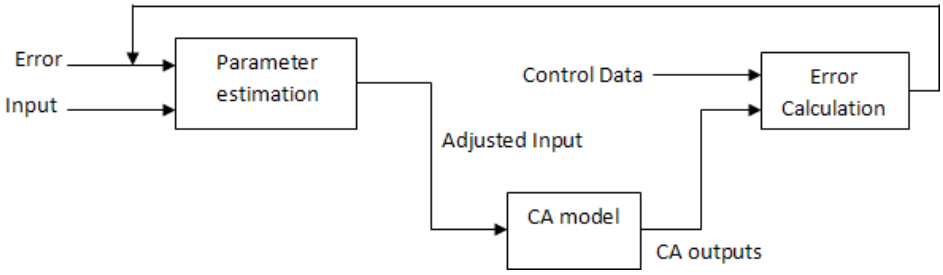

Fig. 12. Interactive parameter estimation with CA visualization and manual adjustment.

Fig. 13. The dual-channel polarized projectors for overlaying the simulated results on the observation data.

with artificial life. This cycle can be repeated for the computer and the human to learn together.

Visualization of the artificial life is implemented with the dual-channel display system, where we display the simulated results on the left channel and the ground truth data on the right channel. The polarized display two datasets simultaneously. Furthermore, the multi-modal display system is scalable from a desktop monitor to a large projected screen.

5. Spatial Bayesian Transform

Here we define a Bayesian cellular automata model to represent the spatial stochastic formation of an artificial life 'alga.' The alga lives in a two-dimensional space divided by a grid of cells. In this model, we consider the location of an alga (x, y), time (t), and other physical stressors. For historical images, c denotes the presence of an alga. Given the location in (x, y), and time (t), the following

equation can be used to find the probability that location is inhabited by alga at that time.

$$P\left(c|x,y,t\right) = \frac{P(c)P(x|c)P(y|c)P(t|c)}{P(x,y,t)} \tag{6}$$

Sometimes when additional evidence e such as the surface temperature is present, it is necessary to use the recursive property of Bayesian to get a more accurate value of the probability:

$$P(c|x,y,t,e_1,e_2,\ldots,e_k) = \frac{P(c)P(x|c)P(y|c)P(t|c)\prod_{i=1}^{k}P(e_i|c)}{P(x,y,t,e_1,e_2,\ldots,e_k)} \tag{7}$$

Calculating each of the probabilities is simple:

$$P(c) = \frac{N_c}{N} \tag{8}$$

$$P(x|c) = \frac{N_{x,c}}{N} \tag{9}$$

N_c is the number of pixel occurrences of c in all images, N is the total number of pixels in all images, and $N_{x,c}$ is the number of pixel occurrences of c at location x in all images. Similar calculations to $P(x|c)$ are made for $P(y|c)$, $P(t|c)$, and $P(e|c)$.

Just calculating the probability of algae present is not good enough. To visualize this prediction, we need a threshold α to determine the pixel of the output image. The output pixel $O(x,y)$ is a Boolean value decided by the equation below:

$$O(x,y) = P(c|x,y,t,e_1,e_2,\ldots,e_k) \geq \alpha \tag{10}$$

For images outputting the probability, the value of each grid represents the probability the grid is inhabited by artificial life. The higher the probability, the darker the region is on the probability image. Probability images are used to predict the location of life given the location, time, and any additional evidence.

5.1. Gray area rendering and interaction

The probabilities under those certain conditions/evidence, such as time and temperature, can be calculated. Figure 14 shows examples of the estimated target locations.

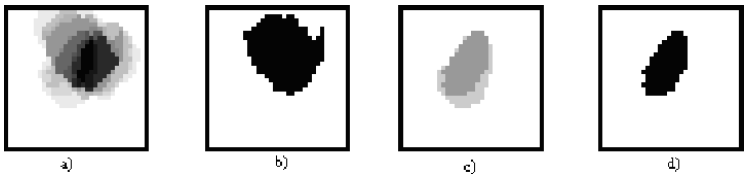

Fig. 14. a) Probability of location being inhabited by algae using a single input temperature = 15℃. b) output image of using 15℃ and , c) probability of location being inhabited by algae using a single input time = January, d) output image of January.

Prediction using one input becomes much more inaccurate during months

with cold temperature because during certain times of the year, no alga is detected when the temperature falls below 15℃. In Figure 14, a) and c) are poor probability models because of the lack of evidence for the Bayesian model to make an accurate prediction; this is indicated by the large number of cells in each image with very low probability (light colored cells). When the model is fed with both temperature and time: 15℃ and January, it becomes much more accurate, as shown in Figure 15.

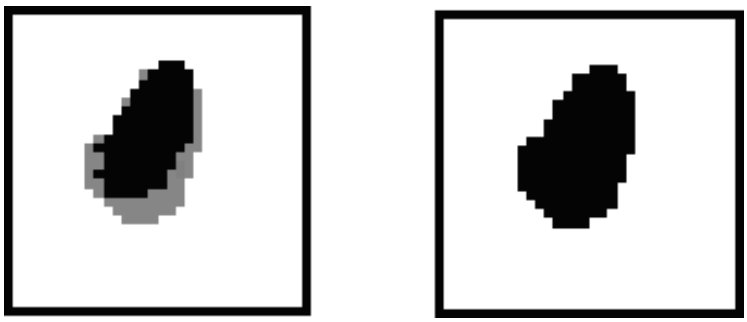

Fig. 15. Left: probability of location being inhabited by algae. (Note the reduced amount of light cells). Right: predicted location after applying the interaction (threshold setting).

5.2. Isosurface rendering and interaction

Isosurface is an interactive tool for volumetric visualization. The isosurface plots in Figure 16 have latitude, longitude, and time. Each slice of the blue isosurface with respect to the time axis represents the predicted image at time $t = z$ with some constant α where α is the probability threshold constant or confidence level. Humans can interact with the iso-surface plot and be able to make more accurate predictions by adjusting the value of α. This visualization tool can be used by the user to determine what confidence level, or α to use.

In Figure 16, the horizontal axes are the latitude and longitude axes of predicted image. Each time slice of the blue iso-surface represents the predicted location of life. For the figure above, $\alpha = 0.4$ was used for the prediction.

Isosurface images can be viewed from any angle and the user has the choice of choosing the time interval to view the plot, and the α threshold. Through visualizing different slices of the plot with different α values, the user can interact with the Bayesian probability model to produce more accurate images by adjusting the α value for each time period.

We verify the process with a case study predicting the movement of harmful algal bloom (HAB) of *Karena Brevis*. In this case, the HABs are lumps, rather than individual cells. The input data are satellite images and cell counts. The prediction model uses the satellite images as additional evidence. A Spatiotemporal Bayesian prediction model is implemented using a modified version of the equations above. The inputs for the prediction are the location in latitude and longitude, the time, and the patch of the image for the region of interest. For the prediction, the probability of HAB present is compared to the probability of not present. The higher probability will determine the classification. Because the

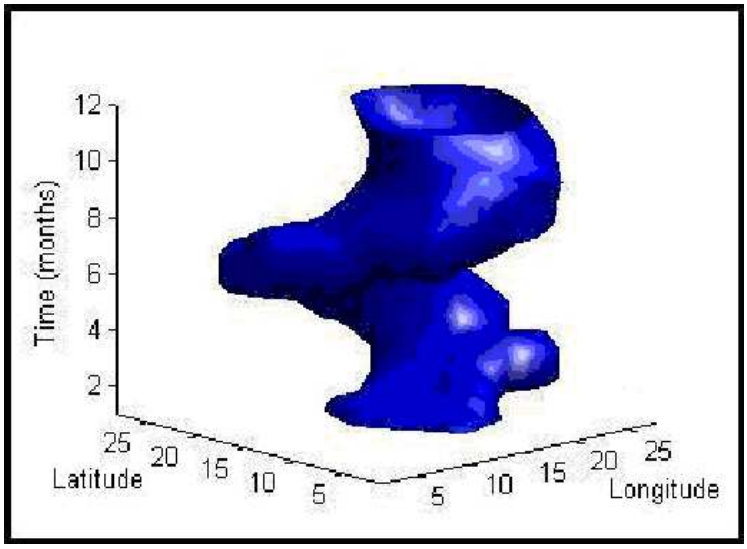

Fig. 16. Iso-surface plot of predicted image $\alpha = 0.8$.

denominator for calculating both of the probabilities is the same, the prediction equation summarizes to:

$$SB(x, y, t, I) = \underset{c \in C}{argmax}\, P(c)P(x|c)P(y|c)P(t|c)P(I_{x,y}|c) \tag{11}$$

I is input image, and $I_{x,y}$ is the $c \in C$ patch for the region of interest. See the next section "Usability Study" for more details.

6. Usability Study

We investigate the visual transformation methods in the context of our real-world scientific research projects in oceanographical studies. Therefore, our approach has to make sense to scientists who conduct day-to-day analyses and discovery tasks. Table 4 summarizes the advantages and disadvantages of the visualization modes and interaction modes in a matrix form.

We tested our interactive spatiotemporal data mining system with the field data of cell counts from the Gulf of Mexico area and SeaWiFS satellite images from NASA. We quantified our results, for example, positive accuracy, which describes the percentage of predictions which were correct when Harmful Algal Bloom cell count, was greater than or equal to 5,000. The measurements are defined as following:

- Positive accuracy is the percent of the cases in which target is present and the model predicted correctly.
- Positive accuracy = confirmed positive / (confirmed positive + false negative).
- Positive detection is the percent of all predictions that are correct.
- Positive detection = (confirmed positive + confirmed negative) / (sum).

Table 2. Interactive Spatio-Temporal results

Method	False positive	Confirmed positive	False negative	Confirmed negative	Sum	Positive detection	Positive accuracy	Negative accuracy
Image only	44	17	6	306	373	86.60%	73.91%	87.43%
$SB^1w/oSDT^2$	161	423	142	1658	2384	87.29%	74.87%	91.15%
SBw/SDT	176	445	120	1643	2384	87.58%	78.76%	90.32%
$SBw/SDT\&Int^3$	166	441	124	1653	2384	87.84%	78.05%	90.87%
$SBw/SDT\&Or^4$	173	445	120	1646	2384	87.71%	78.76%	90.49%

[1]- Spatiotemporal Bayesian Model [2]- Sparse Data Treatment
[3]-Interpolated images [4]- Original Images

Table 3. Manual-only results (Tomlinson, et al, 2004)

Method	False positive	Confirmed positive	False negative	Confirmed negative	Sum	Positive detection	Positive accuracy	Negative accuracy
Reference Results	5	36	23	124	188	85.10%	61.02%	96.12%

7. Conclusion

Analytical models intend to reveal inner structure, dynamics, or relationship of things. However, they are not necessarily intuitive to humans. Conventional scientific visualization methods are intuitive but limited by depth, dimensions and resolutions. The purpose of this study is to bridge the gap with transformation algorithms for mapping the data from an abstract space to an intuitive one.

The visual transformation algorithms incorporate computer vision, multi-physics simulation, and machine learning models to solve the tracking and prediction problems across multiple databases. In particular, the visual transformers include shape correlation for interactive tracking, periodicity for periodical patterns, multi-physics for simulating the complex surface dynamics, and spatial Bayesian for spatio-temporal prediction.

Embedding computer vision into the conventional data mining algorithms enable us to automate the data pre-precessing and increase accuracy in the overall process.

In contrast to the conventional data mining algorithm, which is normally context-free and independent from the physical constraints, we introduced a visual multiple physics model Cellular Automata to simulate the physical, biological and chemical phenomena. This provides a mechanism to explore the deep knowledge of spatiotemporal patterns while the accurate computational models are not available or impractical.

In this study, the human-computer interactions are embedded into the data mining process with multi-modal interfaces such as desktop computers and stereo projectors. The interactive visualization is not only a display but also a learning and problem solving process.

We have tested this approach with the oceanographic data mining from NASA's 8-year SeaWiFS satellite database along with NOAA's 50-year in-situ cell count databases. We found that the interactive visualization increases robustness in object tracking and positive detection accuracy in object prediction. We found that the interactive method enables the user to process the image data at less than 1 minute per image versus 30 minutes per image manually. As a result, our test system can handle significantly more data sets (e.g. 5,000 samples) than traditional manual analysis (e.g. 188 sample).

The empirical benchmarks and usability studies suggest that the visual in-

Table 4. Comparisons of the visual transformation methods.

Visualization	Interaction	Advantage	Disadvantage
Spatial Bayesian Transform with Gray-scale rendering	Adjusting probability threshold	Showing the possible algae formation location intuitively	One frame per time slice, lack of confidence indication for none measurement spots
Spatial Bayesian Transform with iso-surface	Rotation, pan, zoom, and cutting planes	Combining spatiotemporal data in one object	Obscure and over-confidence for long-term prediction
Shape Correlation Transform	Human intervention at the breaking points from the shape correlation indicator	Improvement of the tracking continuity.	The correlation indicator is just one visual cue. It may not be always working for complex scenes
Periodicity Transform with zooming windows	Sliding windows for spatiotemporal data selection	The mathematical transform enables user to discover the periodicity easily and the slide window enables flexible data selection	How to select the suitable windows is sensitive to the detection results. Also pre-processing of the raw data is necessary to recover the missing spots.
Multi-Physics Transform on a desktop	Adjusting the input parameters	Intuitive and fewer input parameters	The adjustment process could be time-consuming.
Multi-Physics Transform with the stereo projectors and polarized glasses	Adjusting the balance of the two channels	Displaying the ground truth and simulated data simultaneously	Additional equipment for users. Possible fatigue after starring for a while

teractions can improve the data mining quality. However, where, when and how to engage the visual interaction are key factors in the process. The results also suggest that minimal human interactions with appropriate computational transforms or cues may significantly increase the overall productivity.

Acknowledgements. The authors would like to thank the editor of this special issue Dr. Elena Zudilova-Seinstra for her comments and support. The authors appreciate the

comments and suggestions from Karen Meo, Kai-Dee Chu, Steven Smith, Gene Carl Feldman and James Acker from NASA. Also, many thanks to Christos Faloulus and his students from Carnegie Mellon University for their inputs.This study is supported by NASA ESTO grant AIST-QRS-04-3031, Army Research Office, Cylab, Cylab-Korea, KISA and General Motors.

Appendix: CA Simulation Rules

The growth rule is as followed given the growth amount N:

1. Pick an arbitrary cell neighboring the region of life
2. If the sum of neighboring cells ≥ 2 and the arbitrary cell is '0' then make the cell '1'
3. Repeat steps 1 and 2 until the desired N has been reached

The shrink rule is as followed given the shrink amount M:

1. Pick an arbitrary cell on the edges of the region of life
2. If the sum of neighboring cells ≤ 2 and the arbitrary cell is '1' then make the cell '0'
3. Repeat steps 1 and 2 until the desired N has been reached, or a maximum iteration threshold has been reached. If the threshold is reached, it means the region of life has shrunk to a very condensed region and will not easily reduce in size.

The collision rule is as followed given elasticity E:

1. If any cells cross the boundary, proceed to the next step.
2. Treat the rows and columns outside of the boundary as '1'
3. If the sum of neighboring cells around any cell is ≥ 4 then make the cell '1'. This phase will add the cells onto the boundary row, or in the case that the region of life has already collided with a boundary, add cells close to the previously collided cells.
4. Repeat this step E times since the more elastic the collision, the more the cells will spread.

The wind translation rules are as followed given wind speed W and direction:

1. Using a constant for speed of translation with relationship to wind speed. (e.g. $.15 * W$)
2. Move the shape at the speed of one iteration at a time in the x and y direction.
3. If there is a collision with a boundary, spread along the boundary, but continue moving.

References

Aarts, E. and Marzano, S. (2003), *The new everyday = views on ambient intelligent*, ambient intelligent for everyday life, lecture notes in artificial intelligence, lnai 3864 edn, Springer, 2006.

Bandini, S. and Pavesi, G. (2002), 'Controlled generation of two-dimensional patterns based on stochastic cellular automata', *Future Generation Computer Systems* **18**, 973–981.

Cai (2005), *Ambient Intelligence for Scientific Discovery*, Vol. 3345, Springer.

Cai, et al (2005), *Ambient Diagnostics, Ambient Intelligence for Scientific Discovery*, Vol. 3345, Springer.

Cai, Y. and Abscal, J. (2006), *Ambient Intelligent for Everyday Life, Lecture Notes in Artificial Intelligence, LNAI 3864*, Springer.

Cowell, A. J., Havre, S., May, R. and Sanfilippo, A. (2005), *Scientific Discovery Within Data Streams, Ambient Intelligence for Scientific Discovery*, Vol. 3345, Springer.

Devaney, et al (2005), *Science at the Speed of Thought, Ambient Intelligence for Scientific Discovery*, Vol. 3345, Springer.

Essam, J. W. (1980), 'Percolation theory', *Rep. Prog. Phys.* **43**, 833.

Gardner, M. (1980), 'Mathematical games', *Scientific American* **243**, 18–28.

Hardy, et al (1976), 'Molecular dynamics of a classical lattice gas: Transport properties and time correlation functions', *Phys. Rev. A* **13**(5), 1949–1961.

Hubona, G. S. and Shirah, G. W. (2005), *Spatial Cues in 3D Visualization, Ambient Intelligence for Scientific Discovery*, Vol. 3345, Springer.

Karl Fu, Y. C. (2006), Spatiotemporal data mining with cellular automata, Technical report, Cylab.

Langley, P. (1977), 'Bacon: A production system that discovers empirical laws', *IJCAI 1977* pp. 344–344.

Lewin, K. (1975), *Spatial Cues in 3D Visualization, Ambient Intelligence for Scientific Discovery*, Greenwood Press Westport.

Leyton, M. (1992), *Symmetry, Causality, Mind*, MIT Press.

Linsay, R., Buchanan, B. and Feigenbaum, E. (1993), 'Dendral: a case study of the first expert system for scientific hypothesis formation', *Artificial Intelligence* **61**(2), 209–261.

Pryke, A. and Beale, R. (2005), *Interactive Comprehensible Data Mining, Ambient Intelligence for Scientific Discovery*, Vol. 3345, Springer.

Sethares, W. A. and Staley, T. W. (1999), 'Periodicity transforms', *IEEE Transactions on Signal Processing* **47**(11), 2953–2962.

sethares, W. A. and Staley, T. W. (2001), 'Meter and periodicity in musical performance', *Journal of New Music Research* **22**(5), 1–11.

Sonka, et al (1998), 'Image processing', *Analysis and Machine Vision* .

Stauffer, D. (1994), *Introduction to Percolation Theory*, Taylor and Francis.

Toffoli, T. and Margolus, N. (1987), *Cellular automata machines: a new environment for modeling*, MIT Press.

Tomlinson, et al (2004), 'Evaluation of the use of seawifs imagery for detecting karenia brevis harmful algal blooms in the eastern gulf of mexico', *Remote Sensing of Environment* **91**, 293–303.

Vichniac, G. Y. (1984), 'Simulating physics with cellular automata', *Physica D Nonlinear Phenomena* **10**, 96–116.

von Neumann, J. (1966), *The Theory of Self-reproducing Automata*, Univ. of Illinois Press, Urbana, IL.

Wolfram, S. (1983*a*), 'Statistical mechanics of cellular automata', *Rev. Mod. Phys.* p. 55:601.

Wolfram, S. (1983*b*), 'Statistical mechanics of cellular automata', *Rev. Mod. Phys.* **55**, 601–644.

Wolfram, S. (1984*a*), 'Cellular automata as models for complexity', *Nature* p. 311:419.

Wolfram, S. (1984*b*), 'Universality and complexity in cellular automata', *Physica D 10* pp. 1–35.

Wolfram, S. (1985), 'Twenty problems in the theory of cellular automata', *Physica Scripta T9* pp. 170–183.

Wolfram, S. (2002), 'A new kind of science', *Champaign, IL: Wolfram Media* pp. 60–70 and 886.

Yeh, A. and Li, X. (2002), 'Urban simulation using neural networks and cellular automata for land use planning', *Symposium on Geospatial Theory, Processing and Applications* .

Zhou, S., Chellappa, R. and Moghaddam, B. (2005), 'tracking and recognition using apprearance-adaptive models in particle filters', *IEEE Trans. On Image Processing* **13**(11), 1491–1506.

Zudilova, E. V. and Sloot, P. M. (2005), *Multi-modal Interaction in Biomedicine, Ambient Intelligence for Scientific Discovery*, Vol. 3345, Springer.

Author Biographies

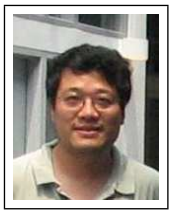

Yang Cai is Director of Ambient Intelligence Lab and faculty of Cylab and Institute of Complex Engineered Systems (ICES), Carnegie Mellon University, and Professor of Industrial Design at Modern Industrial Design Institute, Zhejiang University, P.R. China. He was Systems Scientist at Human-Computer Interaction Institute, Senior Scientist in CMRI at CMU, and Senior Designer for Daimler Chrysler. Cai's interests include pattern recognition, visualization and Ambient Intelligence. He co-chaired international workshops in Ambient Intelligence for Scientific Discovery, Vienna, 2004 and AmI for Everyday Life, Spain, 2005 and Digital Human Modeling, UK, 2006. He is editor of the Lecture Notes in Artificial Intelligence, LNAI 3345 and LNAI 3864, published by Springer. He was NASA Faculty Fellow in 2003 and 2004.

Richard Stumpf is a senior oceanographer and team leader of Remote Sensing National Oceanic and Atmospheric Administration (NOAA), Center for Coastal Monitoring and Assessment, Silver Spring MD, where he leads 6 to 10 team members developing remote sensing capabilities for NOAA. Dr. Stumpf has extensively published papers on remote sensing for monitoring and forecasting harmful algal blooms and river plumes. He received his Ph.D. in Oceanography.

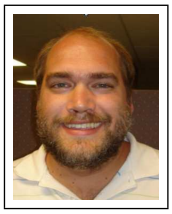

Timothy Wynne is an oceanographer with I.M. systems Group and NOAA. Primarily his work at NOAA has involved ocean color imagery with an emphasis on algal bloom detection. He has also used remotely sensed data to quantify resuspension events. He has a M.S. in Oceanography from Old Dominion University and a B.S. in Marine Science from the Richard Stockton College of New Jersey.

Michelle Tomlinson has been an oceanographer with the Center for Coastal Monitoring and Assessment, National Ocean Service, NOAA since 2002. Her current research focuses on the application of satellite derived ocean color sensors (SeaWiFS, MODIS, MERIS) to detect, monitor, and forecast the occurrence of harmful algal blooms. This work has led to the development of an operational forecast system for harmful Karenia brevis blooms in the Gulf of Mexico. She received her B.S. in Marine Science Biology from Southampton College of Long Island University, and a M.S. in Oceanography from Old Dominion University.

Mattias Ihmig is pursuing his Ph.D. in the area of software defined radio at Munich Technical University, while he is working at BMW, Germany. He was an intern graduate student at Carnegie Mellon University, USA. His interests include stereo vision, wireless networks and intelligent systems.

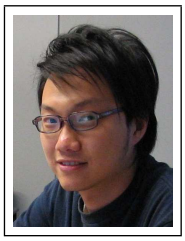

Daniel Sai Ho (Daniel) Chung is a MS degree student at Institute of Networked Information, Carnegie Mellon University. He has been a research assistant in the Ambient Intelligence Lab since 2004, where he developed data mining and wireless video streaming systems for NASA and TRB-sponsored projects.

Xavier Boutonnier is a Research Assistant at Carnegie Mellon University, CYLAB - Ambient Intelligence Lab, Pittsburgh, PA, USA. He is a Master of Science degree student at the National Superior School of Electronics of Toulouse (ENSEEIHT) in France. He specialized in Signal, Image, Acoustic and optimization. He has been working with Dr. Yang Cai on the NASA-sponsored data mining project. His favorite fields of application are Video, Image, acoustic, and other signal processing.

Rafael Franco is a Master student in Electronics and Telecommunications for the Engineering School of ENSEEIHT in France. Currently, he is an intern at Cylab working on assignments related to information visualization and wireless user positioning.

Nathaniel Bauernfeind is a Research Assistant at Carnegie Mellon University, CYLAB - Ambient Intelligence Lab, Pittsburgh, PA, USA. He is an undergraduate Computer Science and Math student in the School of Computer Science at CMU. His research interests include computational algorithms, 3D graphics programming, and artificial intelligence. He is working with Dr. Yang Cai on a driving simulator that focuses on algorithmic automation for General Motors.