Spatial Aggregation for Qualitative Assessment of Scientific Computations

Chris Bailey-Kellogg
Purdue Computer Science Dept.
250 N. Univ. St.
West Lafayette, IN 47907
cbk@cs.purdue.edu

Naren Ramakrishnan
2160L Torgersen Hall
Department of Computer Science
Virginia Tech, VA 24061
naren@cs.vt.edu

Abstract

Qualitative assessment of scientific computations is an emerging application area that applies a data-driven approach to characterize, at a high level, phenomena including conditioning of matrices, sensitivity to various types of error propagation, and algorithmic convergence behavior. This paper develops a spatial aggregation approach that formalizes such analysis in terms of model selection utilizing spatial structures extracted from matrix perturbation datasets. We focus in particular on the characterization of matrix eigenstructure, both analyzing sensitivity of computations with spectral portraits and determining eigenvalue multiplicity with Jordan portraits. Our approach employs spatial reasoning to overcome noise and sparsity by detecting mutually reinforcing interpretations, and to guide subsequent data sampling. It enables quantitative evaluation of properties of a scientific computation in terms of confidence in a model, explainable in terms of the sampled data and domain knowledge about the underlying mathematical structure. Not only is our methodology more rigorous than the common approach of visual inspection, but it also is often substantially more efficient, due to well-defined stopping criteria. Results show that the mechanism efficiently samples perturbation space and successfully uncovers high-level properties of matrices.

Introduction

A significant trend has recently developed in the numerical and scientific computing community, with researchers employing *data-driven* methodologies to assess broad, explicative features characterizing scientific computations. A typical application is in the analysis and "explaining away" of computations that may not be producing correct results. Such an approach is in stark contrast to the traditional focus on developing ever-stable numerical methods, to avoid the dangers of finite-precision computations. The philosophical shift is significant, since as Chaitin-Chatelin and Frayssé point out, "the aim is no longer to control the computing error, but rather to extract [high-level] meaning from results." (Chaitin-Chatelin & Frayssé 1996) Data-driven methodologies have been successfully applied in studying phenomena including conditioning of matrices, sensitivity to various types of error propagation, and algorithmic convergence behavior.

This approach to analysis of scientific computations has been labeled by one scientific computing group as *qualitative computing* (Chaitin-Chatelin & Traviesas 2002), and has many parallels to its namesake in the artificial intelligence community, particularly qualitative spatial reasoning in the style advocated by (Abelson *et al.* 1989). (1) It utilizes *imagistic reasoning*, expressing data in a suitable spatial context, and extracting and analyzing higher-level spatial structures in order to characterize, explain, or optimize a phenomenon of interest (Yip & Zhao 1996). Just as dynamical systems are qualitatively characterized with phase portraits (Yip 1991) and kinematic mechanisms in configuration spaces (Joskowicz & Sacks 1991), scientific computations are now characterized in perturbation and precision spaces. (2) It must harness *domain-specific knowledge*, here encoding the underlying mathematical structure, e.g. semantics arising from perturbation theory and computational invariants. Such information is important in order to discard spurious models of data, and to define rigorous acceptability criteria for models. (3) It should *integrate data collection and analysis*, focusing sample collection to gather information most useful for the key purpose of extracting qualitative insight. Intelligently closing-the-loop by interleaving sample selection and model assessment carefully exercises the degrees of freedom in collecting data in order to extract reinforcing interpretations that overcome noise and sparsity.

These facets present novel research issues when encountered in the context of assessing scientific computations. Our goal in this paper is to cast the assessment problem in a formal setting emphasizing qualitative model selection and iterative data collection. We do this within a spatial aggregation mechanism that identifies structures in the data, according to domain knowledge about the mathematical structure. To the best of our knowledge, this work presents the *first* systematic algorithms for performing complete imagistic analyses on matrix perturbation datasets (as opposed to relying on human visual inspection (Chaitin-Chatelin & Frayssé 1996)), and which interleave data collection and model evaluation until a high-confidence model is obtained. We first introduce the underlying problem domain and briefly survey background work in spatial aggregation, before describing our specific qualitative analysis mechanism. Results on important classes of matrix problems are then presented.

Qualitative Assessment of Matrix Eigenstructure

Many tasks in scientific computing involve assessing the eigenstructure of a given matrix, i.e., determining the locations of eigenvalues, their multiplicities, and their sensitivities to perturbations. Eigenstructure helps characterize the stability, sensitivity, and accuracy of numerical methods as well as the fundamental tractability of problems. For instance, the eigenstructure of a matrix underlying the behavior of a harmonic oscillator gives valuable information about the vibration characteristics of the oscillator, as a function of forced inputs.

The *spectral portrait* has emerged as a popular tool for graphically visualizing eigenstructure in the complex plane; it displays how the eigenvalues of a matrix change as perturbations are introduced. The spectral portrait of a matrix $\mathcal{A}$ is defined as the map in the complex plane:

$$\Phi(z) = \log_{10} \|\mathcal{A}\|_2 \; \|(\mathcal{A} - zI)^{-1}\|_2, \tag{1}$$

where I is the identity matrix. The singularities of Φ are located at the eigenvalues of $\mathcal{A}$, and the analysis determines the sensitivity of computation to numerical imprecision by analyzing how the map decreases (from ∞) moving away from the eigenvalues. To see why, notice that when z is an eigenvalue, the expression $(\mathcal{A} - zI)^{-1}$ is undefined because the determinant of $(\mathcal{A} - zI)$ is zero, and in such a case $\|(\mathcal{A} - zI)^{-1}\|_2$ is taken to be ∞. When z is near an eigenvalue, $\|(\mathcal{A} - zI)^{-1}\|_2$ will be a large but finite number and can be expressed in terms of $\|\mathcal{A}\|_2$ as $\frac{1}{10^{-k}\|\mathcal{A}\|_2}$, for some k. Here, $\Phi(z)$ is simply given by k.

Why is the spectral portrait useful? Notice that the condition $\|(\mathcal{A} - zI)^{-1}\|_2 \geq \frac{1}{10^{-k}\|\mathcal{A}\|_2}$ describes the region enclosed by a level curve labeled k in the spectral portrait. It can be proved (e.g., see (Chaitin-Chatelin & Frayssé 1996)) that such a region also denotes the eigenvalues of perturbed matrices $\mathcal{A} + E$ where $\|E\|_2 \leq 10^{-k}\|\mathcal{A}\|_2$. Hence the level curves in a spectral portrait correspond to perturbation magnitudes, and the region enclosed by a level curve contains all possible eigenvalues that are equivalent with respect to perturbations of the given magnitude. The above discussion will lead us to believe that the spectral portrait should consist of small circles around the eigenvalues, with radii given by $\Phi(z) = k$. For many matrices this is indeed true. Unfortunately, for other classes of matrices (known as 'non-normal' (Golub & Van Loan 1996)) this ideal behavior does not emerge, and this is where the graphical depiction of the spectral portrait becomes an invaluable tool.

Fig. 1(a) shows the spectral portrait of a non-normal matrix with eigenvalues at 1, 2, 3, and 4. The qualitative interpretations from such a portrait are several-fold. First, if the matrix is represented only to a certain (normwise) perturbation level, and we try to compute its eigenvalues, then we might obtain any point inside the region enclosed by the corresponding curve! The level curves summarize the extent to which we can afford loss of precision. Observe that at small perturbation levels, the spectral portrait is disconnected (which is desirable), but it gets connected at higher

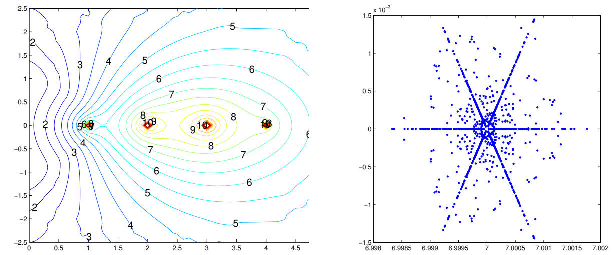

(a) (b)

Figure 1: Qualitative assessment of matrix eigenstructure by analysis of portraits in the complex plane. (a) Spectral portrait. Perturbation levels increase going outward from singularities at the (unperturbed) eigenvalues, and a curve surrounding multiple eigenvalues indicates a level of precision (labeled with negative logarithm base 10) beyond which those eigenvalues cannot be distinguished. Eigenvalues at 2 and 3 are the most sensitive to perturbation, but that at 4 becomes indistinguishable under a larger perturbation, and 1 merges under a very large perturbation. (b) Jordan portrait. Multiple positive and negative perturbations cause the computed eigenvalues to lie on the vertices of two slightly shifted polygons, yielding a star with degree twice the eigenvalue multiplicity. The six-fold symmetry of the star reveals a multiplicity of three for the eigenvalue at 1.

levels. Second, notice that there is a "drift" toward the eigenvalue at 1 as perturbations are increased. This gives important information about the "defectiveness" of the eigenvalues and the non-normality of the matrix. Finally, the spectral portrait suggests how iterative computations involving the matrix will converge; the defectiveness of the eigenvalues indicates that there might be periods of instability and/or periods of slow convergence.

Having determined the relationship between the eigenvalues at given perturbation levels, we typically desire to ascertain the (geometric) multiplicities of the eigenvalues, for further insight into matrix stability. This information is summarized in the so-called Jordan decomposition, whose direct computation is often difficult or impossible (Golub & Van Loan 1996). Once again, we can employ a data-driven approach, here determining the multiplicities by qualitatively assesing a different portrait in the complex plane, which we call a *Jordan portrait*. The Jordan portrait uses the property that, due to finite precision arithmetic, the *computed eigenvalues* will actually be the *true eigenvalues* of a suitably perturbed matrix. Specifically, the computed eigenvalues corresponding to an actual eigenvalue λ_i under perturbation δ are given by:

$$\lambda_i + |\delta|^{\frac{1}{\rho_i}} e^{\frac{i\phi}{\rho_i}}, \tag{2}$$

where λ_i is of multiplicity ρ_i, and the phase ϕ of the perturbation δ ranges over $\{2\pi, 4\pi, \ldots, 2\rho_i\pi\}$ if δ is positive and over $\{3\pi, 5\pi, \ldots, 2(\rho_i + 1)\pi\}$ if δ is negative. These phase variations imply that the computed eigenvalues lie on the vertices of a regular polygon with $2\rho_i$ sides, centered on λ_i, and with diameter influenced by $|\delta|$. Therefore, qualitative analysis of the *symmetry* in a Jordan portrait that graph-

Samples Evaluate perturbation level (Eq. 1) for locations in complex plane, initially on a coarse uniform grid.
Spatial aggregates Aggregate interpolated samples with Delaunay triangulation. Construct iso-contour curves from connected points at same perturbation level.
Correspondence Establish analogy for connected points on different curves; abstract into curve merge events.
Qualitative model merge tree; constructed by tracking iso-contour correspondence outward from eigenvalues. For each pair of eigenvalues, evaluate model by separation and surroundedness.
Sampling policy Sample until absolute separation and absolute surroundedness are above limits, and relative values are converging.

Table 1: Spatial aggregation for spectral portrait analysis.

Samples Perform random normwise perturbation (Eq. 2) at desired level.
Spatial aggregates Aggregate samples into triangles. Identify congruent triangles by geometric hashing.
Correspondence Establish analogy between congruent vertices; abstract as rotation.
Qualitative model polygon isomorphism; obtained by identifying rotations applicable to many points. Evaluate model by likelihood calculation.
Sampling policy Collect additional samples at same or subsequent perturbation level, until posterior distribution of models is peaked enough.

Table 2: Spatial aggregation for Jordan portrait analysis.

ically *superimposes* numerous such perturbed calculations reveals the multiplicity ρ_i. For example, the Jordan portrait in Fig. 1(b) depicts a 6-gon around the eigenvalue at 1, hence indicating that the eigenvalue multiplicity is three.[1]

Spatial Aggregation Mechanism

As the previous section illustrates, matrix eigenstructure can be characterized by imagistic analysis of responses to matrix perturbations. Our goal is to automatically and confidently extract high-level matrix properties through application of spatial reasoning. We develop here an integrated mechanism that abstracts perturbation data into spatial structures, quantitatively evaluates and compares models of matrix properties, and collects additional perturbation samples as needed. For spectral portraits, the analysis can be viewed as yielding a model we call a *merge tree* (Fig. 2, right), indicating for each pair of eigenvalues the perturbation level at which the iso-contours surrounding them merge. For Jordan portraits, our model is a *polygon isomorphism* (Fig. 3) capturing the symmetry of perturbations around an eigenvalue, and thereby revealing the multiplicity. The key components and knowledge for the spatial reasoning mechanism are summarized in Tabs. 1 and 2; detailed in the following subsections; and illustrated in Figs. 2 and 3.

[1]The 'noise' around the star is an artifact of the given matrix; while this 'hallucination' effect is not modeled in this paper, it can pose problems for automatic imagistic analysis.

Extracting Spatial Structures

We ground our analysis mechanism in the Spatial Aggregation Language (SAL) (Bailey-Kellogg, Zhao, & Yip 1996; Yip & Zhao 1996; Zhao, Bailey-Kellogg, & Fromherz 2003), which defines a set of uniform operators and data types that exploit domain knowledge to identify structures in spatial data. SAL was developed in order to make explicit the reasoning performed by a number of successful AI applications (e.g., dynamical systems analysis (Yip 1991), kinematic mechanism analysis (Joskowicz & Sacks 1991)), and to provide a suitable vocabulary and mechanism for new applications. Subsequent reasoners utilized SAL in a number of different areas, including decentralized control design for thermal regulation (Bailey-Kellogg & Zhao 2001), object manipulation (Zhao, Bailey-Kellogg, & Fromherz 2003), weather data analysis (Huang & Zhao 1999), analysis of diffusion-reaction morphogenesis (Ordóñez & Zhao 2000), and identification of pockets underlying gradient fields (Bailey-Kellogg & Ramakrishnan 2001).

Two key aspects of SAL are particularly relevant here: use of a *spatial aggregate* representation identifying *relationships* among spatial objects in data, and exploitation of *domain knowledge* in uncovering these aggregates. The spatial aggregates in matrix eigenstructure analysis identify relationships among perturbations in portraits, and the domain knowledge coveys the underlying mathematical properties. We note that the matrices being studied are defined by different applications, so additional, specific knowledge about those domains could provide even more guidance; we do not pursue that here but focus on generic portrait analysis.

Our spectral portrait analysis algorithm groups perturbation samples into iso-contours and "tracks" these contours outward from eigenvalues through containment and merging events. The construction of iso-contours follows a stadard spatial aggregation paradigm; the tracking requires an extension discussed below. Samples are taken initially on a uniform grid, and isocontours extracted by interpolating locations at which desired perturbation levels are achieved. Spatial aggregates are computed by (1) localizing computation by aggregating interpolated points into a Delaunay triangulation (relating nearby objects); (2) grouping connected points in a perturbation level (relating similar objects); and (3) constructing curve objects for the groups of points (composing lower-level objects into higher-level objects).

Our Jordan protrait analysis algorithm groups perturbation samples into polygons and recognizes symmetry by a rotation aligning polygons. (The eigenvalue multiplicity is then apparent from the polygon degree.) As with spectral portraits, the first step follows a standard spatial aggregation approach, and the second requires extensions. Samples are taken by random normwise perturbation, and so appear "scattered" around the complex plane. The underlying mathematics indicates, however, that there is a structure in the portrait, such that, except for noise, points lie on vertices of regular polygons. In order to identify the polygon, we first group triples of points and compose them into triangles whose congruence will allow a subsequent analysis step to identify the symmetry defining the polygon.

In both cases, then, spatial aggregation groups samples into some higher-level structure (iso-contours, triangles). The abstraction step allows further analyses to treat the objects "holistically," with more global properties (e.g. curvature and congruence) that aren't defined on mere sets of lower-level objects. In particular, the key to the applications studied here is *correspondence* relationships. Correspondence among level curves in spectral portraits reveals the merge events indicating when eigenvalues become indistinguishable. Similarly, correspondence among triangles of Jordan samples reveals the underlying symmetry. As an imagistic analysis aid, correspondence has been well studied in computer vision and pattern recognition. Our purpose here is to distill essential principles into the traditional SAL mechanism, exploiting the hierarchical composition of spatial aggregates, and demonstrating the effective use of such a correspondence mechanism in overcoming noise and sparisty in interpreting portraits.

In the context of SAL computations, it is easy to see that correspondence is intimately connected with locality and continuity, since close, similar parts of one object typically correspond to close, similar parts of another. Further, identification of sites at which this isn't true can be interesting events about which to reason (e.g. in the spectral portrait application, this distinguishes straight containment of iso-contours from merge events). Thus our correspondence mechanism leverages the fact that a SAL hierarchy composes lower-level objects into higher-level objects based on locality and continuity. Our mechanism has two key steps: (2) establish *analogy* as a relation among lower-level constituents of higher-level objects; (2) establish *correspondence* between higher-level objects as an *abstraction* of the analogy between their constituents.

Definition 1 (Analogy) *Given a set L of lower-level SAL objects that have been composed into a set H of higher-level objects, $l_i, l_j \in L$ is said to be an* analogous pair *with respect to $h_k, h_l \in H$ if $l_i \in h_k, l_j \in h_l$. An* analogy *is a relation A on L with respect to H containing a subset of the possible analogous pairs.*

Definition 2 (Correspondence) *Given sets L and H and an analogy A as in Definition 1, a* correspondence *on H abstracting A is a relation $M \subseteq H \times H$ defined for each pair of objects with analogous components in A, so that the objects can be transformed (approximately) onto each other.*

The analogy between constituents is well-defined only because of the context provided by the higher-level objects; higher-level correspondence then captures a more global view of the local matches. In spectral portrait analysis, analogy is defined for the interpolated samples on separate curves, connected by Delaunay edges. Correspondence then identifies pairs of curves with a strong set of analogies. In Jordan portrait analysis, analogy is defined by identifying (e.g. with geometric hashing) pairs of congruent triangles and mapping the vertices between each pair. Correspondence abstracts the analogy into a rotation such that analogous vertices map onto each other.

Representing and Evaluating Models

Our goal is to identify and evaluate models explaining the perturbation data in terms of high-level matrix properties. The qualitative spatial models considered here capture spatial objects and their relationships, but do not otherwise impose any restrictions on how models are posited; in particular, a model could be derived by means other than SAL, and could be represented using a rich vocabulary of relationship types (e.g. (Cui, Cohn, & Randell 1992)). Given a model, derived either by the aggregation / correspondence mechanism or by some other means, we desire to decide how well the data support it, that is, to quantify confidence in the posited relationships. We employ a confidence function $f(M, d)$ to assess a qualitative spatial model M in terms of a dataset d.

In a spectral portrait, a model encapsulates level curves and their merging relationships in the merge tree structure (can also be a forest), with leaves for eigenvalues, internal nodes for curves, and a parent function such that a parent's perturbation level is one larger than its child's. Such a model indicates, for each pair of eigenvalues, the magnitude of the smallest perturbation level at which they are equivalent. The dataset d is a sampled grid of points in the complex plane, and we would like to ensure that the eigenvalues are definitely separate before their lowest common ancestor, and that the level curve for that ancestor definitely surrounds the eigenvalues. For each pair of eigenvalues i and j, let v be the level of the ancestor, and define $R_{i,v-1}$, $R_{j,v-1}$, and $R_{ij,v}$ to be the regions (sets of samples) containing respectively i, j, and both, as computed by region growing with samples of at most the given perturbation level. We measure the separation by the smallest number of samples on any grid walk between regions $R_{i,v-1}$ and $R_{j,v-1}$. We measure the surroundedness by the maximum flow on the grid for $R_{ij,v}$, from sources $R_{i,v-1}$ to sinks $R_{j,v-1}$ (i.e. strong connection between the regions).

In a Jordan portrait, the polygon isomorphism model relates pairs of vertices of two polygons. Such a model captures the symmetry relation of the sample points by the induced rotation θ around the eigenvalue λ, minimizing distance between related pairs. We can solve a straightforward least-squares problem for θ and λ, and calculate multiplicity ρ as $180°/\theta$. The dataset d is a random set of sample points in the complex plane, each of which can be viewed as an i.i.d. estimate of the eigenvalue. Thus we define the score $f(M, d)$ to be the likelihood $P(d|M)$. This likelihood can be simply given as $\Pi_k P(d_k|\lambda, \rho, \delta)$ where λ and ρ are the eigenvalue location and multiplicity as revealed by the model M, and δ is the perturbation level. The likelihood of each data element $d_k = a + ib$ is in turn assessed using Eq. 2, distributing probability mass of 1 among all vertices of the implied polygon mapping for the eigenvalue and perturbation level corresponding to point d_k. For instance, each vertex of the triangle isomorphism in Fig. 1(b) at a given perturbation level receives probability of $\frac{1}{6}$. This particular formulation of f can be viewed as assessing the posterior of a model under a uniform prior over all models.

Intelligently Collecting Additional Samples

Interleaving model generation, model assessment, and additional data collection steps can be viewed as causing a progressive peaking of the posterior distribution of models (under a probabilistic interpretation of model assessment). We treat data analysis in the simultaneous updating sense, although the Jordan portrait application lends itself to sequential updating (owing to the i.i.d. samples). In other words, with each additional data point, the analysis is performed over the entire dataset accumulated so far; this means that some models hastily ruled out in earlier stages could become relevant later. At some point, enough confidence is gained so that sampling can be stopped. Until that point, decisions must be made as to which additional samples to collect, under the degrees of freedom provided. A *sampling policy* defines the stopping criteria and sample selection strategy. It could be designed by optimization or based on heuristic considerations.

In the spectral portrait, data collection is driven by systematic global subsampling, where we begin with an uniform grid and the sampling strategy has the choice of expanding the grid (to ensure a merge tree rather than forest) or recursively subsampling (to gain more confidence). The design of the stopping criterion for the spectral portrait borrows from discretization refinement in numerical quadrature (Krommer & Ueberhuber 1998), where *pairs of formulas* with different levels of approximation are used so that one of the formulas is *expected* to yield more accuracy than the other, for most classes of problems. Between such successive applications, if the estimated answer doesn't change significantly, we declare to the model to be a sufficient approximation of the given function. Similarly, to assess the quality of the spectral portrait model, we evaluate the stability of merge and separation events, under the metrics provided by model assessment, with respect to additional subsampling.

In the Jordan portrait, data collection is driven by increasing the number of independent samples from which the multiplicity is estimated. The sampling strategy here must decide whether to sample (i) at the same perturbation level, (ii) at a higher perturbation level, or (iii) at the same perturbation level unless the number of posited models increased (thereby avoiding hallucination). Stopping criteria for the Jordan portrait assess the peakedness of the posterior distribution for M and whether the log-likelihood score f crosses a certain threshold (guaranteeing adequacy of a dataset before generalization).

Results

We have applied the our qualitative analysis mechanism on the testbank of matrices described in (Chaitin-Chatelin & Frayssé 1996). These matrices are drawn from a variety of application domains such as hydrodynamics, stress analysis, control theory, fluid flow, and membrane transport phenomena. For want of space, we illustrate our results on only two matrices—*Rose* and *Brunet*—from this collection, highlighting the intuition behind the mechanism and performance results.

Fig. 2 illustrates the spectral portrait analysis for the 10×10 *Rose* matrix, motivated by the problem of determining the roots of the underlying characteristic equation. *A priori*, a double-factorial number of binary merge trees are possible, but the approach presented here eliminates almost all of them without even considering them. Instead, it considers only one plausible tree for a given number of samples, and decides whether or not the merge events captured in the tree can be clarified. Depending on model assessment, the mechanism then subsamples or expands, or halts. While some amount of correspondence is found even with the coarse initial grid (top of Fig. 2), the model is a forest and the model assessment yields low confidence (few separating points). As a result, our mechanism computes additional samples on a finer, larger grid (bottom of Fig. 2), yielding high confidence in the (correct) curve merge tree shown. The mechanism converges from this point when additional data are collected.

We have applied this approach to a variety of matrices (e.g., polynomial companion matrices, matrices from finite discretization of continuous models), with different numbers and spacings of roots. In each case, the correspondence mechanism correctly identifies the correct model with high confidence, essentially declaring that any other model that would be proposed would be highly inconsistent with the data, after 1–3 subsamples and 1–3 grid expansions.

Fig. 3 demonstrates the results of applying our Jordan portrait analysis with the 10×10 *Brunet* matrix. The top part uses a small set of sample points, while the bottom two parts use a larger set and illustrate a good vs. bad correspondence. As the number of samples increases, so does the risk of model "hallucination" — finding some subset of points that by chance happen to correspond, as in bottom of Fig. 3. This illustrates the importance of monitoring relative model confidence and controlling the sampling to avoid over-sampling.

We tested 10 matrices across 4–10 perturbation levels. To study the effect of sampling strategy, we organized data collection into rounds of 6–8 samples each and experimented with the three sample collection policies mentioned earlier. We varied a tolerance parameter for triangle congruence from 0.1 to 0.5 (effectively increasing the number of models posited) and determined the number of rounds needed to determine the Jordan form, for each of the sampling policies described earlier. Policy 1 required an average of 1 round at a tolerance of 0.1, up to 2.7 rounds at 0.5. Even with a large number of models proposed, additional data quickly weeded out bad models. Policy 2 fared better only for cases where policy 1 was focused on lower perturbation levels, and policy 3 was preferable only for the Brunet-type matrices. In other words, there is no real advantage to moving across perturbation levels! In retrospect, this is not surprising since our Jordan form computation treats multiple perturbations (irresp. of level) as independent estimates of eigenstructure.

The effectiveness of our spatial aggregation mechanism lies as much in the targeted use of domain knowledge as in the design of a suitable sampling policy, as revealed by the following count of feasible models:

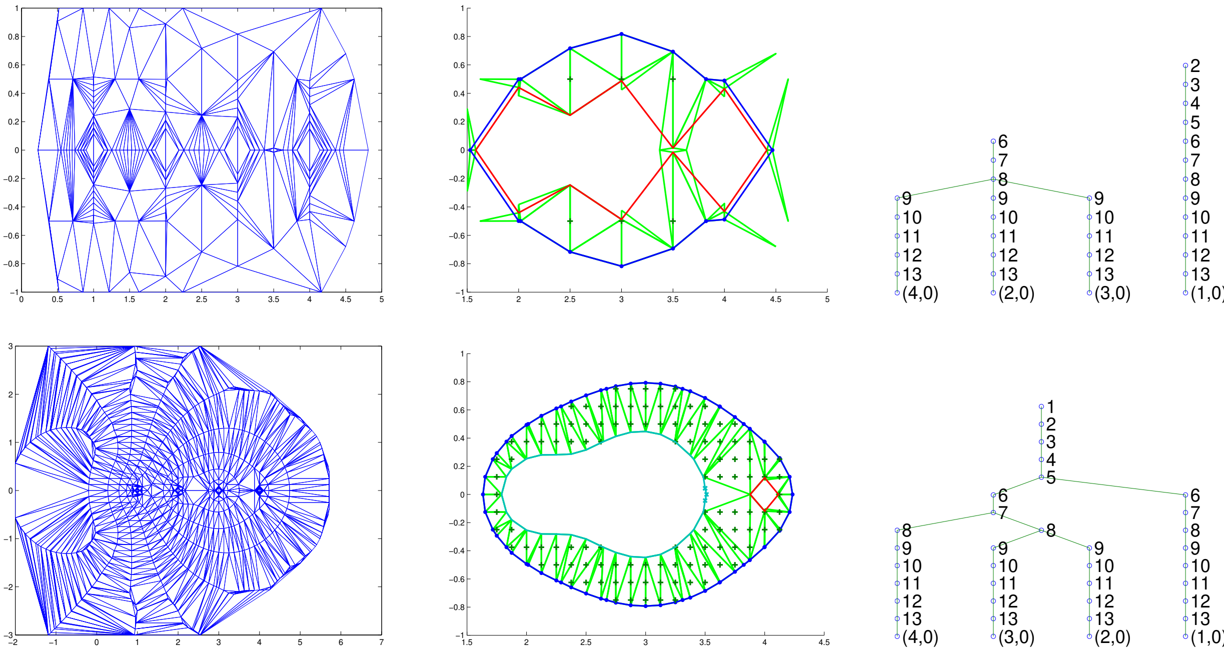

Figure 2: Example of *Rose* matrix spectral portrait analysis for (top) a small, coarse grid, and (bottom) an extended, subsampled grid. (left) Delaunay triangulation analogy for interpolated points comprising contours. (middle) Analogy at an example merge event; separating samples marked with “+”. (right) Curve merge tree: eigenvalues at bottom; node for curves labeled with perturbation level. Merge events indicate at what perturbation level descendant eigenvalues are indistinguishable.

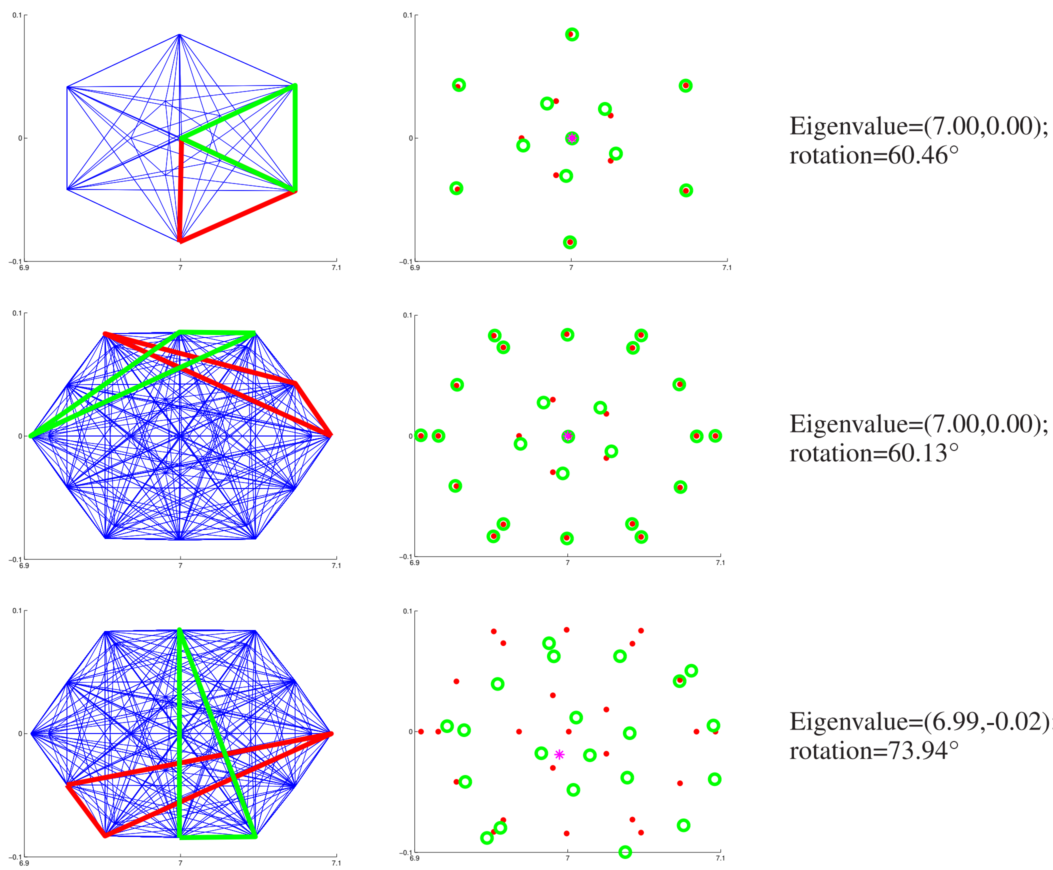

Figure 3: Example of *Brunet* matrix Jordan portrait analysis, for (top) small sample set; (middle) larger sample set; (bottom) larger sample set but lower-scoring model. (left) Approximately-congruent triangles. (middle) Evaluation of correspondence in terms of match between original (red dots) and rotated (green circles) samples. (right) Associated eigenvalue and rotation symmetry.

application	initial	round 1	round 2	round 3
spectral (Rose)	15	1 (forest)	1	1
Jordan (Brunet)	1000	1	1	2

Discussion

Automated imagistic reasoning, of the form presented here, is an important aid to reasoning about the quality, efficiency, and robustness of numerical computations involving matrices. Since matrix computations underly many areas of science and engineering (especially where linear models are employed), qualitative analysis of spectral portraits and Jordan portraits will play a central role in these domains.

The analysis of correspondence using our spatial aggregation mechanism is similar in spirit to that of (Huang & Zhao 1999) for weather data interpretation, and can be seen as a significant generalization, formalization, and application of techniques studied there for finding patterns in meteorological data. Similarly, our methodology captures and generalizes the computations employed in object recognition, allowing the body of research developed there to be applied to a broader class of applications, such as in scientific computing. This work is also very much philosophically aligned with work on structure-mapping (e.g. (Falkenhainer, Forbus, & Gentner 1989)), which seeks to develop a cognitively plausible theory and computationally efficient engine for analogical reasoning. The cited paper points out that "analogies are about relations rather than simple features." As with recent work from that same group (and others) on sketching (e.g. (Forbus, Tomai, & Usher 2003)), our mechanism applies this key insight to spatial reasoning; an important aspect of this paper is the integration of analogical relationships within the context of spatial aggregation relationships.

Similarly to compositional modeling (Falkenhainer & Forbus 1991), we advocate targeted use of domain knowledge, and as with qualitative/quantitative model selection (e.g. (Capelo, Ironi, & Tentoni 1998)), we seek to determine high level models for empirical data. Our focus is on problems requiring integrated qualitative model analysis and sampling to overcome sparsity and noise in spatial datasets. A possible direction of future work is to explore if the inclusion-exclusion methodology popular in grid algorithms (Bekas *et al.* 2001) can be fruitfully harnessed in a SAL-based framework. Our long-term goal is to study data collection policies and their relationships to qualitative model determination. The decomposable nature of SAL computations promises to both support the design of efficient, hierarchical algorithms for model estimation as well as provide a deeper understanding of the recurring roles that domain knowledge plays in spatial data analysis.

Acknowledgements

We received very helpful comments for a preliminary version of this work (Bailey-Kellogg & Ramakrishnan 2003) at the 2003 Qualitative Reasoning Workshop, and would like to thank in particular Ron Ferguson, Ken Forbus, Liliana Ironi, Stefania Tentoni, and Feng Zhao. We also acknowledge the encouraging suggestions from Francoise Chaitin-Chatelin. This work was funded in part by grants from the National Science Foundation to CBK (IIS-0237654) and NR (EIA-9984317).

References

Abelson, H.; Eisenberg, M.; Halfant, M.; Katzenelson, J.; Sacks, E.; Sussman, G.; Wisdom, J.; and Yip, K. 1989. Intelligence in scientific computing. *CACM* 32(5):546–562.

Bailey-Kellogg, C., and Ramakrishnan, N. 2001. Ambiguity-Directed Sampling for Qualitative Analysis of Sparse Data from Spatially Distributed Physical Systems. In *Proc. IJCAI*, 43–50.

Bailey-Kellogg, C., and Ramakrishnan, N. 2003. Active Data Mining of Correspondence for Qualitative Assessment of Scientific Computations. In *Proc. QR*, 23–30.

Bailey-Kellogg, C., and Zhao, F. 2001. Influence-Based Model Decomposition for Reasoning about Spatially Distributed Physical Systems. *AIJ* Vol. 130(2):pages 125–166.

Bailey-Kellogg, C.; Zhao, F.; and Yip, K. 1996. Spatial Aggregation: Language and Applications. In *Proc. AAAI*, 517–522.

Bekas, C.; Kokiopulou, E.; Koutis, I.; and Gallopoulos, E. 2001. Towards the Effective Parallel Computation of Matrix Pseudospectra. In *Proceedings of the International Conference on Supercomputing (ICS)*, 260–269.

Capelo, A.; Ironi, L.; and Tentoni, S. 1998. Automated Mathematical Modeling from Experimental Data: An Application to Material Science. *IEEE Transactions on Systems, Man, and Cybernetics* Vol. 28(3):pages 356–370.

Chaitin-Chatelin, F., and Frayssé, V. 1996. *Lectures on Finite Precision Computations*. SIAM Monographs.

Chaitin-Chatelin, F., and Traviesas, E. 2002. Qualitative Computing. In *Handbook of Numerical Computation*. SIAM.

Cui, Z.; Cohn, A.; and Randell, D. 1992. Qualitative simulation based on a logical formalism of space and time. In *Proc. AAAI*, 679–684.

Falkenhainer, B., and Forbus, K. 1991. Compositional Modeling: Finding the Right Model for the Job. *AIJ* Vol. 51(1-3):pages 95–143.

Falkenhainer, B.; Forbus, K.; and Gentner, D. 1989. The structure mapping engine: Algorithm and examples. *AIJ* 41:1–63.

Forbus, K.; Tomai, E.; and Usher, J. 2003. Qualitative spatial reasoning for visual grouping in sketches. In *Proc. QR*, 79–86.

Golub, G., and Van Loan, C. 1996. *Matrix Computations*. Johns Hopkins University Press. Third Edition.

Huang, X., and Zhao, F. 1999. Relation-Based Aggregation: Finding Objects in Large Spatial Datasets. In *Proc. IDA*.

Joskowicz, L., and Sacks, E. 1991. Computational kinematics. *AIJ* 51:381–416.

Krommer, A., and Ueberhuber, C. 1998. *Computational Integration*. SIAM.

Ordóñez, I., and Zhao, F. 2000. STA: Spatio-Temporal Aggregation with Applications to Analysis of Diffusion-Reaction Phenomena. In *Proc. AAAI*, 517–523.

Yip, K., and Zhao, F. 1996. Spatial Aggregation: Theory and Applications. *JAIR* Vol. 5:pages 1–26.

Yip, K. 1991. *KAM: A system for intelligently guiding numerical experimentation by computer*. MIT Press.

Zhao, F.; Bailey-Kellogg, C.; and Fromherz, M. 2003. Physics-based encapsulation in embedded software for distributed sensing and control applications. *Proceedings of the IEEE* 91:40–63.