

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing this collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. **PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.**

1. REPORT DATE (DD-MM-YYYY) 03-03-2008		2. REPORT TYPE Final Report		3. DATES COVERED (From - To) 03/07-03/08	
4. TITLE AND SUBTITLE New Directions in Complex Data Analysis for Emerging Applications				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER FA9550-07-1-0438	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S) Hamid Krim				5d. PROJECT NUM	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) North Carolina State University, Raleigh, NC 27695				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) AFOSR: 875 North Randolph Street Suite 325 Room 3112 Arlington, Virginia 22203				10. SPONSOR/MONITOR'S ACRONYM(S) Dr. Jon Sjogren	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION / AVAILABILITY STATEMENT					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT The workshop was organized into four thematic sessions including interdisciplinary trends, algebraic geometry, differential geometry and statistics, and topological and geometric features of data. The invited speakers were directed to consider the workshop as a visionary forum. The presentations were all limited to ten minutes with an emphasis on new directions, open problems, and provocative speculation. These were followed by breakout sessions where the presented material was further discussed in small groups. The charge to the participants was to provide concrete and coherent recommendations for new research directions.					
15. SUBJECT TERMS Algebraic Geometry, Topological modeling, Differential geometry, Dimension Reduction and data representation, Graphical representation					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON Dr. Jon Sjogren
a. REPORT	b. ABSTRACT	c. THIS PAGE			19b. TELEPHONE NUMBER (include area code) (703) 696 6564

**Workshop Final Report:
NEW DIRECTIONS IN COMPLEX DATA
ANALYSIS FOR EMERGING APPLICATIONS**

**Beckenridge, Colorado
March 28, 2008**

Hamid Krim

June 13, 2008

Contents

1	Workshop Description	2
2	Workshop Organization	3
3	Problems in Complex Data Analysis	5
4	New Directions in Geometry	7
5	Open Research Problems	10
5.1	Set-to-set Pattern Recognition	10
5.2	Statistical Signal Processing (SSP)	12
5.3	Algebraic Geometry and Control	12
5.4	Geometry and Shape Theory	14
5.5	Algebro-Geometric Tools for Shape Theory	15
5.6	Detection, Comparison and Analysis of Sampled Manifolds . .	17
5.7	Algebraic Geometry for Image Processing	18
5.8	Representation and Reconstruction	20
5.9	Low-Dimensional Embeddings	20

1 Workshop Description

The workshop venue was Breckenridge Resort, Breckenridge, Colorado. This was an ideal environment in which to exchange ideas, debate new and important research trends, and discuss avenues of attack on problems of fundamental importance to the analysis of complex data sets. For two and a half days, the participants presented overviews of compelling problems and speculated on future directions. Each of the four main sessions consisted of a thematic set of lectures followed by a small group breakout session for focussed discussion resulting in reports to the main group. These reports summarized connections between the presentations and opportunities for future work. Many interesting ideas arose from these discussions and they are summarized in this report. A common theme that emerged was the significance of geometry (algebraic, computational and differential) for aiding knowledge discovery from data sets that are both massive and high dimensional. Two main thought procedures formed the basis for meeting the challenge of the highly complex and data-intensive applied problems researchers now face. One line of thought advocates the collection of all data followed by a reduction step and then a processing step while the other prefers to combine the pruning with the collection stage (compressed sensing, ensure the sensor measurement only relays reduced data, so that no redundant information is taken into account further up the line). The mathematics driving these approaches are, at times, distinct in fundamental ways. One can contrast these ideas with approaches that seek massive non-compressed data collection with the hope that simultaneous and parallel processing of such data will afford the extraction of yet more discriminating information. The participants as a whole spoke to the increasing importance of mathematical theory for making progress in these data processing problems and the importance of collaborations involving practitioners and theoreticians. Specific opportunities for such collaborations are detailed in this report.

The importance of algebraic, computational and differential geometry and the potential for these areas to contribute new data processing algorithms was evident at the workshop. Curves for instance, over the last couple decades, have been of crucial importance in investigating localization and tracking issues for tumors in medical imagery, or for detecting and tracing the spatial evolution of holes in the ozone layer in environmental sciences. Statistical or

other characteristics shared by a fraction of the data which translate into a coherent geometrical entity afford a simpler and perhaps more direct means to "divide and conquer" and to interpret otherwise complex and daunting data. While such a curve based approaches have proven to be very successful, their application to higher dimensional data with more complex geometry is very limited. Such difficulty arises in many applications where the curse of dimensionality very quickly becomes an issue, starting with 3D bodies to other processes lying in higher dimensions but yet associated with common characteristics which may be used to some advantage. One such example, relevant in security applications, is the characteristic space where human face data lie. Such data measurably lie in thousands of dimensions, while the information-bearing submanifold dimension is in reality a small fraction of that. Such embedding clearly has a significant impact on the subsequent computational load as may be rigorously understood through the Whitney Embedding theorem in Algebraic geometry. Such a theoretical framework could be used to guide the construction and optimal implementation of practical algorithms to address practical problems. Other connections which arise in the course of further exploring high dimensional data may arise in connection with various group invariants such as curvature (the fundamental Euclidean invariant for curves). Other invariants, including topological invariants, may be constructed for manifolds. These may in turn play a crucial role in characterizing high dimensional data and in determining the most appropriate data structures with which to represent and compare them. Algebraic structure-preserving morphing, formalized in category theory, could lead to additional insights into how to best compare data which are typically subjected to transformations.

In what follows we describe the organization of the workshop, give an overview of some problems in complex data analysis, and provide a summary of new directions and opportunities.

2 Workshop Organization

The workshop was organized into four thematic sessions including interdisciplinary trends, algebraic geometry, differential geometry and statistics, and topological and geometric features of data. The invited speakers were di-

rected to consider the workshop as a visionary forum. The presentations were all limited to ten minutes with an emphasis on new directions, open problems, and provocative speculation. These were followed by breakout sessions where the presented material was further discussed in small groups. The charge to the participants was to provide concrete and coherent recommendations for new research directions. Below is a list of the speakers organized according to the thematic areas.

- Inter-disciplinary Trends (Monday afternoon, 1:00-4:30)
 - Doug Cochran (session chair)
 - Davi Geiger
 - Louis Scharf
 - Yoshihisa Shinagawa
 - Peter Schroeder
 - Rina Tannenbaum
- Algebraic geometry (Monday evening, 6:30-9:30)
 - Chris Peterson
 - Jayant Shah
 - Jon Sjogren (session chair)
 - Andrew Sommese
 - Peter Stiller
 - Allen Tannenbaum
- Differential Geometry and Statistics (Tuesday Morning, 8:30-11:45)
 - Y. Baryshnikov
 - Huiling Le
 - Stacey Levine
 - Peter Olver
 - Tony Yezzi (session chair)

- Laurent Younes
- Topological Geometric Features of Data (Tuesday Evening, 6:30-9:30)
 - Emmanuel Candes
 - David Dreisigmeyer
 - Peter Giblin
 - Michael Kirby (session chair)
 - Hamid Krim

The workshop had the following organizational committee:

Co-General Chairs:

Hamid Krim, Professor, Department of Electrical Engineering, North Carolina State University, Raleigh, NC

Michael Kirby, Professor, Dept. of Mathematics, Colorado State University, Fort Collins, CO

Anthony Yezzi, Professor, School of Electrical and Computer Engineering, Georgia Tech, Atlanta, GA

Local Arrangements:

Michael Kirby, Professor, Mathematics Department, Colorado State University, Fort Collins, CO

All speakers were invited by the organizers. In addition to presenting talks all speakers participated in small group and large group discussions and also provided material for this workshop report.

3 Problems in Complex Data Analysis

Complex data come in a variety of forms: it may be the surface of an object, or a field of sensor measurements, or a temporal sequence of images. Given

such data, gleaned embedded information and exploiting it remain the principal goals of the data analysis. Such analysis entails the discovery of the intrinsic structure of the data which would in turn constitute its characterization and hence its parsimonious representation. An example, is that of a 2D surface of an object being captured by a simple graphical model whose weights represent detailed geometric information. While preserving the information (topological as well as geometric) of a given object, this *simpler* representation clearly yields a *reduced form* with significant computational advantages, and a very useful statistical framework for classification and recognition applications. Dimension reduction of a data set is often based on information that is characteristic but does not necessarily require the whole space to be expressed or summarized. An example of that is the representation of a surface by a set of sampled curves on the surface. This may be accomplished by defining an intrinsic characteristic function, referred to as a Morse function, which by its denseness nature, provides a good measure of the surface. Sampling such a function, in effect is tantamount to sampling a surface along a "curve" dimension. Modeling these samples will effectively yield a significant reduction in data while preserving the intrinsic geometric information.

Such problems arise not only in 3D data in Data-Base archival/retrieval and homeland security, but are predominant in video data in surveillance and entertainment industry, remote sensing/atmospheric data in weather forecasting, medical and genomic research where MASSIVE AMOUNTS OF DATA are collected, analyzed and exploited. In security applications, daunting amounts of video data, chemical sensor data, temperature data and audio may all be simultaneously collected, which if not properly processed and mined for the important information, would be a waste and a security risk.

The staggering and increasing number of IED's used in military and civil wars is one of the most deadly threats faced by our national forces and careful and persistent surveillance of likely neighborhoods for such events are among the crucial applications we may explore. Human silhouettes, for example lie in a very reduced dimensional space (e.g. human silhouettes under different postures) which needs to be discovered and used as a guideline in understanding the environment.

In medicine and biology the exploration of processes spans orders of mag-

nitude in scale (meters to nanometers) starting with tissue to proteins and molecules. At each level a great deal of geometrical and topological structure appears and gives rise to the wonders of biological life as we know it.

4 New Directions in Geometry

In the meeting, there was a clear sense that differential geometry, algebraic geometry, topology, group theory, statistics and functional analysis are each of fundamental importance in the analysis of large data sets. It is too much to ask for one person to be an expert in all of these fields. At the same time, there were a number of themes and ideas that cut across disciplinary boundaries. This solidified the feeling of many that teams of pure mathematicians, applied mathematicians, statisticians and engineers must work together to understand the strengths (and language) of each of the others in the group, to understand how to phrase problems in the setting of others in the group and to understand how the collection of individual expertise within the group can be combined to exceed the results obtainable by individuals. Within the meeting it felt apparent that a team approach will provide the best chances to solve the fundamental challenges we now face in synthesizing, comprehending, extracting the information contained in massive, high dimensional data sets for these problems are, by their very nature, interdisciplinary in scope. It is exciting for an algebraic geometer to see that Schubert varieties, Stiefel manifolds, Lie groups, Flag varieties, families of metrics, etc., have an important role to play in the analysis of data. One can only imagine that the topologists, functional analysts and differential geometers feel much the same with the application of Morse theory, shape spaces, energy minimization techniques, compressed sensing and homological methods all yielding new results and insights. Parameter spaces arose many times in different contexts. Sometimes, a non-obvious metric on the parameter space or a novel application of a dimension reduction procedure yielded surprisingly strong results. For instance, we saw new functionals applied to shape spaces, novel metrics on Grassmann varieties and projections of data at the level of data collection all leading to advances. Some ideas from differential and algebraic geometry which the meeting attendees strongly suspected would play a useful role in the near future include:

- (1) Maps of parameter spaces into vector spaces: Place an object at the center of a sphere. Consider the set of all digital pictures that can be obtained by pointing the camera at the object while requiring the lens of the camera to lie on the sphere. This is equivalent to the set of all digital images that can be obtained by fixing a camera so that it points at a defined center of the object then allowing the object itself to rotate in any possible manner about the center. The set of all possible digital images of the object (obtained in this manner) corresponds to the image of a map of $SO(3)$ into the vector space generated by all possible digital images. This image is sometimes referred to as a pose manifold (for a fixed illumination condition).
- (2) Maps of vector spaces into parameter spaces: Fix a camera and an object then consider the collection of all digital images obtainable by varying the illumination conditions of the object. It is not difficult to see that the weighted average of two different images is again an obtainable image thus the collection forms a convex set. It has been shown that the vast majority of the energy of such a data set can be captured by a relatively low dimensional linear space. By fixing the dimension of this linear space (call this dimension L) we can associate a vector space to each object, pose pair. This linear space corresponds to a point on the Grassmann variety of L dimensional subspaces of the vector space generated by all possible images.
- (3) Vector bundles: Now consider the collection of all possible digital images obtainable by allowing variations in both pose and illumination conditions. For each fixed pose we have an L dimensional vector space. We can think of the entire data set as a vector bundle over the pose manifold. Fixing an illumination condition corresponds to taking a section of this bundle. Varying over the entire data set yields a map of $SO(3)$ into the Grassmann variety.
- (4) Fiber bundles: When data is collected over two variations of state one can consider the sub data obtained by fixing one state and varying the other state. As in the case of vector bundles, this yields a map of one state space into the moduli space of possible fibers with each fiber corresponding to the other state space.

- (5) Schubert varieties: For many parameter spaces there are fundamental sub-objects (playing a role similar to that of subspaces of a vector space). For instance, consider the Grassmannian of all rank two subspaces of a four dimensional vector space. Let L be a fixed one dimensional vector space. The subvariety of the Grassmannian consisting of all the two dimensional spaces which contain L is an example of a Schubert variety. The Algebro-Geometric tools that have been developed in the context of Schubert varieties are likely to be useful in data problems.
- (6) Riemannian Manifolds for Continuous Curves and Surfaces: While families of curves with a fixed and finite number of landmark points have been extensively studied in finite dimensional vector spaces, only recently has attention begun to branch toward infinite dimensional Riemannian manifolds for continuous curves where the metric on the manifold is formulated to be independent of the parametric or implicit representation of the curves. This type of formulation is relevant, for example, if one wishes to compute an optimal morphing between two curves or the average of two curves in cases where there are no easy ways to sample the curves or extract a finite number of prominent geometric feature points. While valid Riemannian metric spaces for continuous curves has received modest levels of attention recently, nothing has been done yet for surfaces and higher dimensional geometric datasets (3-folds, 4-folds, etc...). Such study is important as a precursor to the dimension reduction step. When considering infinite dimensional geometric entities, it is crucial to understand the larger manifold where these objects live in order to properly understand the finite dimensional submanifolds which may help us analyze the geometric data with computationally reasonable algorithms.

The general themes in the six points listed above are Parameter Spaces, Maps, Fibers of Maps, Incidences, Relations between Incidences, Riemannian metrics for infinite dimensional shape manifolds. It is expected that in the more distant future, further tools from algebraic and differential geometry will come into play but for the near future these seem like sure bets.

5 Open Research Problems

5.1 Set-to-set Pattern Recognition

High dimensional data contains patterns that hold valuable information about a physical process, human activity or a battlefield situation. We view a set of such patterns, such as sets of images generated by video, as a family if it has a common feature or set of features. Families of patterns have best bases, i.e., they live in spaces of reduced dimension. There are additional subspaces of reduced dimension that intersect these spaces known as Schubert Varieties. Algorithmically exploiting the idea of a family of patterns and the induced geometry may lead to new algorithms. In particular, we note that it appears very promising to compare sets of images to sets of images rather than to simple compare still images to prototypes of interest.

Families of patterns may live in subspaces, submanifolds or subsets in high ambient dimensions. We need fast(er) algorithms to characterize these distinctions. Understanding the way data sits in space is important for selecting algorithms for data understanding. For example, digital images of faces do not form a subspace as it is not closed under addition. Can this observation be used to guide our data processing?

As evidenced by these examples and others in this conference report, the themes of differential and algebraic geometry as applied to the characterization of information in large data sets clearly emerged over the course of the workshop. Animated discussions indicated that this is an exciting new area with many open questions concerning such topics as sampling theory, geometric invariants and issues as fundamental as correct measures of distance or similarity in this geometric framework. We observe that characterizing things through geometry is not simply doing the same things with a new vocabulary. For example, as will be described below, certain questions in data processing find their natural language in geometry and outside this setting are seemingly intractable. Though while we can create a Rosetta stone to compare working in flat and curved spaces the questions that can be answered in the latter domain extend what we can conclude about the origin and nature of large data sets.

As an example of the power of geometry, consider the recognition of ob-

jects over a variation in state. For instance, the face recognition problem with variations in illumination may be viewed in this way. This is widely considered to be one of the most challenging aspects of the face recognition problem. It is natural to remove, or normalize away such complicating variations and to reduce the problem to normalized prototype comparison. However, this approach fails to exploit an intrinsic feature of the problem, namely, the way that illumination varies over an object actually contains discriminating information that should be retained at all costs. Further, rather than remove this variation, one should seek to collect such information when it is available.

Geometrically such a variation can be quantified as a mapping a of set of images associated with a face under different lighting conditions to a point on a nonlinear parameter space, e.g., the Grassmannian. The power of this encoding is that now a point in a parameter space (really representing megabytes of pixel data) can be compared with other points (images of other subjects) in a natural way using one of the many metrics that are widely known in differential geometry.

This framework can be referred to as the image set-to-set paradigm which can be summarized as follows:

- An instance of a representation of a pattern is a set of observations.
- The characterization of a single class is a collection of such sets.
- Our objective is to match an unlabeled set of images with a class.

We may now pose many fundamental open problems.

How separated can P k dimensional planes be in n -dimensions? How does this separation depend on the resolution n , population size P and representation dimension k ? This is a data packing problem for points generated by real data on the Grassmannian.

Given a cloud of data how does one identify the independent variables? Random projection works in linear setting but is not optimal for nonlinear setting or if the data is not sparse.

What geometric characteristics of the data can be quantified (e.g., symmetry, curvature, dimension) and used as a guide to algorithm selection?

What geometric invariants do objects possess that are invisible in standard vector space representations in high dimensions but appear like flashing red lights in the correct parameter spaces?

5.2 Statistical Signal Processing (SSP)

Geometry and SSP may offer interesting opportunities for future work. A challenging problem in SSP remains robust methods for solving multi-modal optimization problems in radar, sonar, and communication, and perhaps geometry has something to say. Another tough problem is to solve large inverse problems for equalization, inverse imaging, and so on, at very high rates for very nonstationary problems. One may speculate that SSP could be re-worked along geometric lines, rather than subspace lines, to produce a theory more general and more powerful than what we have. A cautionary note is that the SVD seems to be a powerful bridge between geometry and linear algebra. Perhaps we could understand more clearly why it is so powerful.

Are there avenues open to integrate existing approaches? It seems that many of the intuitions from SSP and communications might be integrated into geometry to get at the question of complexity, modelling, compression, and processing. Information theory seems to be missing, even though we are talking about measurements. (In fairness it is missing in much of SSP as well.) We would argue that even in 3 years a more profound understanding of the limits of subspace modelling might be gained.

Are there opportunities to advance basic mathematics by considering applications in this regime? Mathematicians in geometry and topology should collaborate with mathematical engineers to explore a new set of ideas related to concepts like bandwidth, power, capacity, rate-distortion, rank reduction, and so on.

5.3 Algebraic Geometry and Control

Families of dynamical systems present key problems which arise in trying describe their algebro-geometric properties. Families of dynamical systems

appear in all aspects of systems and control theory. Indeed, the essential need for feedback in control systems is the fact that the plant model is only an approximation, and so we must in reality design for a whole family of plants. Of course, the appropriate notion of family depends upon the type of problem in which we are interested. For example, in robust control, families are modelled by certain natural norm bounded perturbations of a given nominal plant. This is a local analytic point of view.

In the early 1970's, R.E. Kalman undertook a global algebraic approach to the problem of system parametrization when he constructed a *universal parameter space* of linear time invariant systems of fixed state space and input/output dimensions. In doing this, he initiated a powerful algebro-geometric framework for studying families of linear time invariant dynamical systems. This approach has had major ramifications in algebraic systems theory and basically opened up a new branch of study. Indeed, a whole conference was dedicated at Harvard in 1979 just to consider this research area. Even today more than 25 years later, prominent researchers are continuing along this research stream.

Besides the introduction of algebraic geometry into control, Kalman's work illuminated the deep connection between invariant theory and a number of control problems. Given the introduction of invariant theory and algebraic geometry into control, it was only a small step to bring *geometric invariant theory* into the picture. Indeed, geometric invariant theory may be regarded as an algebro-geometric manifestation of classical invariant theory. It was devised by David Mumford precisely in the context of universal families (or *moduli spaces*) of algebraic varieties.

The purpose of our briefing was to give a geometric-invariant theoretic construction of the Kalman space and using this construction to derive some of its key geometric properties and to describe possible new research directions in systems and signal processing using these type of techniques. Such methods appear in many applications including image processing and the statistical analysis of data (e.g., GPCA) all of which impacts the information sciences.

Where is the field going? It has been now more than thirty-five years since, Kalman initiated the geometric approach to families of systems outlined above. Of course, even today the concept of family remains fundamental

in systems and control, and is really the underlying object of study in both adaptive and robust control. Especially relevant is adaptive control with its emphasis on the notion of identification, since much of the interest in the the topology and geometry of the moduli spaces of systems was precisely for identification theoretic reasons. However it is important to note that despite the uses of very high powered techniques from topology, complex analysis, Lie groups, algebras, differential and algebraic geometry, the precise global structure of these universal families and parameter spaces still remains an open problem to this day, and one of active research interest.

There has also been an extensive program of research using a local analytic approach in both adaptive and robust control. It turned out to be highly profitable both from the theoretical and practical standpoints to consider families defined in weighted H^∞ balls using techniques from operator and interpolation theory. There is also much work being carried out on the melding of the robust and adaptive control approaches to system uncertainty and families of systems. In short, the study of families of systems whether from the algebraic or analytic, local or global point of view lies at the heart of feedback control theory. Certainly, the Kalman construction of the moduli space of dynamical systems is one of the major achievements in this area. Algebraic geometry allows one to investigate the internal parameters on families of systems in a completely principled way which makes it a powerful tool in the information sciences.

5.4 Geometry and Shape Theory

We can sometimes characterize manifolds of invariants of general objects such as Kendall shape spaces for Euclidean shapes, Grassmannian manifolds as a model for affine shapes and the infinite dimensional manifold of curves in 2D. However, invariants of a set of sample data, e.g. the curves, shapes or affine shapes of surfaces of a collection of aeroplanes or human faces, usually lie in very low dimensional submanifolds of such large manifolds. One possible way to represent the common feature of the invariants of the data is by using the exponential map and a modified PCA method, described as follows, to determine the submanifolds in which the invariants of the data lie. If we can express the exponential map ‘exp’ at points of the manifolds,

we can use the inverse of the map ‘exp’ to map the invariants of the data to a relevant tangent space, for example, that at the ‘mean’ of the data, and perform modified PCA, to take into account the metric structure of the tangent space, there to determine a lower dimensional subspace of the tangent space which then gives an appropriate submanifold. When the data set is rich enough, such submanifolds will give a very good approximation of those in which the invariants of, say, aeroplanes or human faces lie. This can certainly be achieved for Kendall shape spaces and Grassmannian manifolds. The Grassmannian manifold of affine shapes may be an attractive object to explore.

Kendall shape space and infinite-dimensional manifolds of curves in 2D. To use an infinite-dimensional manifold of curves in 2D in practice, one needs to use finite-dimensional approximations. Hence it is important to understand its finite-dimensional equivalents. One possibility is to investigate its relationship with Kendall shape spaces and to explore the limit of some kind of subspaces of Kendall shape spaces as the number of vertices of configurations tends to infinity. There are many different possible embedding structures for Kendall shape spaces when the number of vertices increases, which make it possible to consider various possible limit schemes when the number of vertices tends to infinity. However, the resulting limits are very likely to be too big for practical applications and so one might consider restricting to the limits of sequences of submanifolds of Kendall shape spaces.

5.5 Algebro-Geometric Tools for Shape Theory

A very interesting topic for future research would be to explore the role that algebraic geometry might play to reconcile two different approaches to shape analysis, i.e., differential Riemannian geometry and the calculus of variations versus the Kendall approach to shape analysis using landmarks.

Severe pathologies associated with the metric structure that has been implicitly utilized in gradient descent methods for geometric active contours have been observed. By geometric active contours, we refer to those active contour models which are not dependent upon the representation of the curve (either its parameterization or its implicit formulation). The underlying norm assumed by these geometric gradient techniques is a geometric version of L2

in which the arclength measure is used when integrating along the curve. A surprising result is that the associated inner product does not produce a well-behaved Riemannian metric on the manifold of curves. The resulting space is incomplete, the resulting energy on homotopies is not lower semi-continuous, and the relaxation of this energy yields a distance of zero between any two curves in the space. PDE flows to reduce this energy are ill-posed and the simplest attempts to repair this problem through the use of conformal factors in the metric still do not guarantee the existence of geodesics between any two given curves in the space. This has been the primary reason behind the exploration of Sobolev metrics for shape analysis and for the design of new active contour models.

The Kendall theory does not suffer from any of the above problems. This, however, is not surprising since the resulting spaces are finite dimensional. The price one pays for using the Kendall theory is the need to select semantically meaningful landmark points on the shape and to fix the number of landmarks used when comparing multiple shapes. In the limit, as the number of landmarks goes to infinity, one obtains a specific parameterization of the shape. One could consider a sort of geometric limit by considering equally spaced landmarks according to arclength, and then let the number increase to infinity. In this case, the only remaining degree of freedom in this sampled geometric representation would be a cyclic permutation of the landmarks. For any finite number of equally spaced landmarks, the resulting quotient space to represent the curve would consist of n -dimensional Euclidean space with equivalence classes given by both the similarity group and cyclic permutations.

The tools available in algebraic geometry used to consider such spaces may provide transformative insights to this problem. Since we know that in the infinite continuous limit, the resulting L^2 metric structure breaks down, algebraic geometry may shed more light and lead to a better understanding of this phenomenon by studying the behavior of these finite dimensional spaces modulo the group of cyclic permutations and as the dimension increases. Not only could this lead to a better understanding of the continuous limit, which is where much recent research into shape spaces is concerned, but it could also lead to a better understanding of the finite dimensional Kendall approach to shape and how one should go about choosing an optimal number of landmarks given the constraints and parameters of their particular application.

5.6 Detection, Comparison and Analysis of Sampled Manifolds

Of all the problems research directions mentioned in the Breckenridge meeting maybe the most challenging relate to the detection, comparison and analysis of sampled manifolds. Suppose one has a point cloud that comes from a noisily sampled submanifold in a high dimensional space. There is a significant need, motivated by a broad range of applications, to develop robust algorithms for reconstructing basic geometrical properties — dimension of the submanifold, topology, curvatures and other invariants, metrics and geodesic distances, etc. — from the point cloud samples. Methods requiring a preliminary triangulation appear to be inadequate and indirect to handle densely and noisily sampled objects. Invariant signature recognition, particularly noise reduced signatures based on joint invariants, integral invariants, and semi-differential invariants, requires new statistical sampling methods and comparison of the resulting invariant signature submanifolds. For example, how often does one need to sample two submanifolds (using some probabilistic distribution) to be 99% sure they are the same or different? Can techniques from compressed sampling be applied, i.e. how can one formulate a theory of compressed sampling of submanifolds? How can one effectively apply learning algorithms to objects with non-flat intrinsic geometry?

A range of shape and submanifold metrics have already been proposed. However, despite much work in this area, many key issues, both intrinsic and extrinsic, remain underdeveloped and properly tested. An even more basic question is whether metrics are the correct mathematical construct required to compare shapes and submanifolds. Further analysis of the pros and cons of metric geometry versus more general geometries is required.

Classification and detection of symmetries extends the domain of interest to "currents", representing multiply parameterized submanifolds. For example, the number of discrete symmetries of an object can be found by determining the index of the signature — how many times the signature is covered by the original. Further development of distance and other joint invariant histograms appears promising, but needs testing and comparison with other approaches.

A better understanding of image and signature statistics would be of

importance in comparing, classifying and analyzing scenes and objects in images. For example, can one develop natural curvature (or torsion, or ..) statistics to enable classification of objects? Recognition and reconstruction of scenes from stereo, video, etc. requires understanding how the differential invariants and other invariant quantities behave under projection to the screen, with only preliminary results available to date. Applications include object & target recognition, tracking, motion, scene reconstruction, etc. Extensions to more general transformation groups, including infinite-dimensional pseudo-groups, should be pursued.

Invariant numerical algorithms are just beginning to be developed and applied to systems arising in applications — image processing, fluid mechanics, invariant flows, and so on. This fall under the general area now known as "Geometric Integration", which has received much attention and development in other parts of the world, but where the US seems lagging at present. Combining methods from the discrete variational calculus and moving frames seems a very promising way to develop symmetry-preserving codes with potential benefits. Use of the underlying geometry, e.g. circles in the case of conformal geometry or conic sections in affine geometry, is promising, but requires a more extensive development of techniques and testing on real world problems.

5.7 Algebraic Geometry for Image Processing

Efficiently recognizing three dimensional arrangements of features on an object from a single two dimensional view requires an approach that is view and pose invariant. Existing methods often rely on computationally expensive template matching. Those methods use comparisons against templates created for all possible views; with the infinite number of possibilities being approximated by some finite number of views. To carry out an invariant approach to target recognition, we need to exploit properties and relationships that are geometrically intrinsic to the objects and/or images being compared. Our approach to view and pose independence (as well as coordinate independence) starts with a characterization of a configuration of features by its geometric invariants. The specific group to which things should be invariant is a function of the sensor type. We then derive a fundamental

set of equations that express, in an invariant way, the relationship between the 3D geometry and its "residual" in a 2D (or 1D) image. These equations completely and invariantly describe the mutual 3D/2D constraints. Once derived, they can be exploited in a number of ways. For example, from a given 2D configuration, we are able to determine a set of nonlinear constraints on the geometric invariants of the 3D configurations capable of producing that given 2D configuration, and thereby arrive at a test for determining the object being viewed. Conversely, given a 3D geometric configuration (features on an object), we are able to find a set of equations that constrain the invariants of the images of that object; helping to determine if that object appears in selected images. With these results in hand, future work includes three major problems:

- object/image metrics on shape spaces to provide a distance (difference) between two object configurations, two image configurations, or an object and an image pair in pose invariant, coordinate free terms,
- reconstruction of an object's 3D shape from 2D sensed information, either from multiple sensors or multiple images of a moving object,
- statistical issues surrounding random shapes, distributions of shapes, and noise in object recognition.

Dealing with data on certain manifolds, most notably Grassmann manifolds appears to be a fruitful new direction in the analysis of complex data. Appropriate metrics and also procedures for fitting subvarieties to such data need to be developed. The general question of invariant features of high dimensional data under projections to lower dimensions is also an interesting one. It appears that some aspects of our techniques could be applied to such high dimensional problems. Finally, problems in signal processing may have nice geometric formulations in terms of secant varieties of rational normal curves, where the same sort of metrics on Grassmannians play a role in finding the optimal answer.

5.8 Representation and Reconstruction

Representation: One of the most challenging problems in visual inference is that of representing objects, scenes, categories etc. in ways that trade off invariance/insensitivity to nuisances of image-formation (viewpoint, illumination, occlusion), and at the same time retain discriminative power. In particular, it can be easily shown that any viewpoint invariant statistic is not shape discriminant. However, that is true for "worst-case" invariants, that is image statistics that are invariant to any possible viewpoint, for objects of any possible scene. Because scenes are not generic (the shapes of objects are highly non-generic), we must find ways to embed natural "scene" statistics (typical shapes, typical illumination, typical camera motions) into the design of local feature descriptors that can support decision tasks such as classification or recognition.

3-D reconstruction: Reconstructing the 3-D structure (shape) and appearance (reflectance) of complex surfaces hinges on assumptions about illumination and reflectance properties of the scene. The most common assumptions (Lambertian reflection, diffuse illumination) have worked well so far in laboratory environments, but have failed the test of real scenes, such as outdoors, or complex objects such as vegetation, human skin, shiny indoor materials such as polished surfaces. Formalizing the reconstruction problem in ways that takes into account complex reflectance models requires devising models that have generative power, that is models that can synthesize images that exhibit the non-Lambertian phenomena that we want to capture. Because both shape and reflectance are unknown, reconstruction typically boils down to solving infinite-dimensional optimization problem, and it is important to devise multi-grid, multi-resolution methods that can produce results in useful computational time (order of minutes, not hours or days), to impact applications in cartography, navigation, surveillance etc.

5.9 Low-Dimensional Embeddings

One of the central ideas when dealing with high-dimensional data is the concept of a manifold. Specifically, one often assumes that the data is sampled from some underlying manifold that one wants to process in some way. This

concept allows for the use of various analytic theorems, e.g., Whitney's embedding theorem, for reducing the dimensionality and/or complexity of the data.

When dealing with a high-ambient dimensional embedding, a common theme is to define a function over the underlying manifold that is in some sense optimized for a specific characteristic that one finds desirable. Examples of this are the Karcher mean of a set of points on a manifold and, the optimal Whitney projection direction for manifold-valued data. Now one has to deal with high-ambient dimensional optimization in order to improve the representation of the raw data.

This interplay between working with high-dimensional data and dealing with large-scale optimization problems is a two-way street. One can employ techniques for dimensionality reduction in order to make the optimization problem practical. Alternately, large-scale optimization problems are ubiquitous in dimensionality reduction routines.

It so happens that a significant portion of what goes by the name 'non-linear programming' can be viewed as applied Morse theory. So now one is examining the level sets of some function defined over a manifold. Here, however, the manifold corresponds to the constraints that are present in the problem formulation. One can, in principle, reconstruct the topology of the underlying manifold by solving the optimization problem.

This shift in viewpoint is significant because again we see that the concept of a manifold reappears. So learning how to deal with (potentially low-dimensional) manifolds embedded in a high ambient dimension has important implications beyond the obvious applications to dimensionality reduction. We have the curious situation where the techniques used to solve a problem themselves can be improved by the solution to the problem itself