From *SIAM News*, Volume 37, Number 7, September 2004

Imaging Science

Point Clouds in Imaging Science

Point clouds, one of the most primitive and fundamental manifold representations, were the subject of a minisymposium in Salt Lake City. Guillermo Sapiro, acting chair of the SIAM Activity Group on Imaging Science, was the organizer; the speakers were Facundo Memoli (University of Minnesota), Dario Ringach (University of California, Los Angeles), Gunnar Carlsson (Stanford University), and Partha Niyogi (University of Chicago). Key images and ideas from the session are presented here.

Laser range scanners and other 3-D shape-acquisition devices, which have applications in many disciplines, are popular sources of point cloud data. These devices generally provide raw data in the form of noisy, unorganized point clouds representing surface samples. Researchers are looking for ways to work directly with point clouds and thus to avoid the intermediate step of fitting a surface to the data (a step that can both add to the computational complexity and introduce errors).

Point clouds are also encountered in the representation of high-dimensional manifolds by samples. Such high-dimensional and general co-dimension data appears in almost all disciplines, from computational biology to image analysis to finance. Because it is impossible to perform manifold reconstruction for extremely high-dimensional data, the task needs to be performed directly on the raw data—that is, on the point cloud.

■ ■ ■

Facundo Memoli described a geometric framework for comparing manifolds given by point clouds. The underlying theory is based on Gromov–Hausdorff distances, leading to isometric invariant and completely geometric comparisons. This theory, which includes rigid motion transformations as a particular case, is embedded in a probabilistic setting, as derived from a random sampling of manifolds, and then combined with results on matrices of pairwise geodesic distances for a computational implementation of the framework. Complementing the theoretical and computational results, experiments have been done with real three-dimensional shapes (see Figure 1).

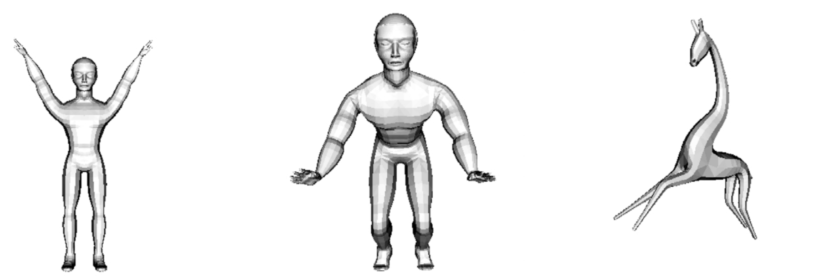

Figure 1. *Data used in the theoretical and computational work described by Facundo Memoli. With the geometric framework derived from this theory, viewers can automatically recognize that the left-hand and middle three-dimensional objects are just bendings of each other, and that they are different from the one at the right. (Data courtesy of R. Kimmel, whose work motivated in part the theoretical and computational results presented by Memoli.)*

Dario Ringach presented challenges in neuroscience for mathematical and computational scientists working on point cloud data. Analysis of neural activity, in which the processing of high-dimensional signals is important, is one example. The brain is composed of large numbers of individual cells (neurons) that communicate among themselves to solve rather sophisticated problems, such as parsing an image or speech and generating accurate motor movements. Understanding how populations of cells, which by themselves appear to be fairly elementary computing devices, work together to solve complex problems is one of the paramount problems in neuroscience.

Using new methods, investigators have been able to measure the activity of large populations of cells on millisecond time scales. In one such method, called voltage-sensitive dye imaging, the brain is perfused with a dye whose optical properties (fluorescence) change as a function of the voltage across the membrane of the neuron (which also determines the firing rate of action potentials). The data stream in such experiments consists of rapid (about 250 Hz) high-dimensional (512 × 512-pixel) images, which represent light distribution generated by the fluorophor across the cortical surface. The goal is to understand the dynamics of brain activity in this high-dimensional space.

Kenet et al. recently discovered that, using the dye perfusion method, direct optically recorded images of activity generated spontaneously in the visual

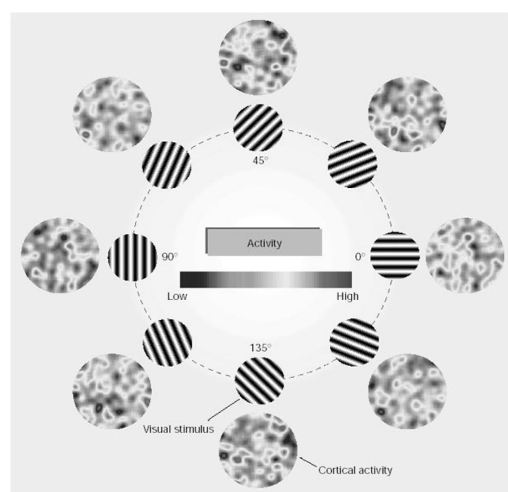

Figure 2. *Mapping orientation. The larger circles show activity across thousands of neurons in a patch of the feline visual cortex (the part of the brain that processes visual information) when the animals look at variously oriented stimuli. A particular "orientation map" is produced for each different visual orientation. Kenet et al. suggest that such orientation maps are intrinsic states (attractors) of the brain, created spontaneously even when a subject's eyes are closed.*

cortex, without visual stimulation, resemble images of activity evoked by oriented patterns. It is as if the brain, in the absence of stimulation, "hallucinates" the presence of a visual pattern (Figure 2). The initial discovery was based simply on the researchers' observation that some spontaneous states resembled states obtained with stimulation. They later used dimension-reduction techniques (a version of Kohonen's self-organizing maps) to analyze the data and verify that their intuition was correct. Further questions then arise: How can the manifold structure of cortical activity, and some of its qualitative topological features, be determined from basic principles? How can activity patterns be compared across different experimental conditions? It seems clear that methods being developed for the analysis of point cloud data hold promise for settling some of these issues.

Gunnar Carlsson discussed applications of algebraic topology to the analysis of high-dimensional data sets. Algebraic topology offers a way to make precise mathematical and computational sense of the intuitive notion of "connectivity information," i.e., the decomposition of spaces into connected components or the presence of "tunnels" or "voids" in the space. Such phenomena are detected by computing so-called Betti numbers, which are the ranks of certain vector spaces attached to the space. These vector spaces are called the homology groups.

Adapting these mathematical ideas to the point cloud domain requires several ideas. The first is the "Cech complex," which associates to a set of point cloud data a simplicial complex, using only information about the interpoint distances in the point cloud. The second is "persistent homology," developed by H. Edelsbrunner and colleagues, which makes it possible to study homology simultaneously at various scales and thus keeps the computations from becoming excessively expensive. The third, the method of "witness complexes" of de Silva and Carlsson, allows the study of complexes that, although much smaller than the Cech complex, nevertheless capture the relevant topological information.

In his talk, Carlsson discussed the application of these methods to a particular data set, consisting of 3×3 pixel patches (viewed as vectors in 9-dimensional Euclidean space) in natural images, from a database of images constructed by van Hateren et al. D. Mumford et al. extracted high-contrast patches and suitably normalized them so that they lie on the unit 7-sphere in 8-dimensional Euclidean space. The intuition of Mumford was that this data should be concentrated around a circle, corresponding to the angle made by a line through the pixel patch dividing it into light and dark regions and the x-axis.

Carlsson and de Silva studied this data set by looking at only those points in the whole data set that lie in the densest regions, but allowing themselves several different measures of density. Their topological study verified Mumford's intuition for a coarse measure of density but revealed interesting new structure at a finer measure of density. Specifically, for the finer notion of density, the data set was concentrated around a union of three circles, one main circle and two subsidiary circles, each of which intersects the main circle at two points and which do not intersect each other. The researchers interpret the main circle as corresponding to pixel patches in which the intensity function is linear in the plane variables x and y; the two subsidiary circles are seen as corresponding to quadratic functions in x and y separately. The method appears to have a great deal of potential for data sets that a priori sit in dimensions too high for visualization, and for which there are no natural projections to study.

Partha Niyogi and colleagues consider pattern recognition and machine learning problems from point cloud data. In many pattern classification tasks the data ostensibly lives in a very high-dimensional space, and classical methods face the so-called curse of dimensionality. The usual solution is to consider reduction to a lower-dimensional space and to perform pattern classification in the reduced-dimensional space.

Niyogi et al. have adopted a more geometrically motivated strategy. They assume that the data lies on or close to a low-dimensional manifold embedded in the ambient higher-dimensional space. If this manifold is known, harmonic analysis can lead to various families of functions that are defined naturally and intrinsically on it. From these families, it is possible to pick a good classifier/regressor by optimizing on the basis of the empirical data.

In the more interesting case in which the manifold is unknown, the researchers proceed by collecting a lot of unlabeled data. These data are then used to estimate a graph (more generally a simplicial complex) approximation to the manifold. Harmonic functions on the manifold are then estimated from the graph and optimized on the labeled training data. This leads to a class of geometrically motivated algorithms for learning from labeled and unlabeled data. Empirical performance evaluation for tasks in vision, speech, and text classification shows that the algorithms are able to utilize unlabeled data effectively to boost performance. Because the labeling of examples requires manual

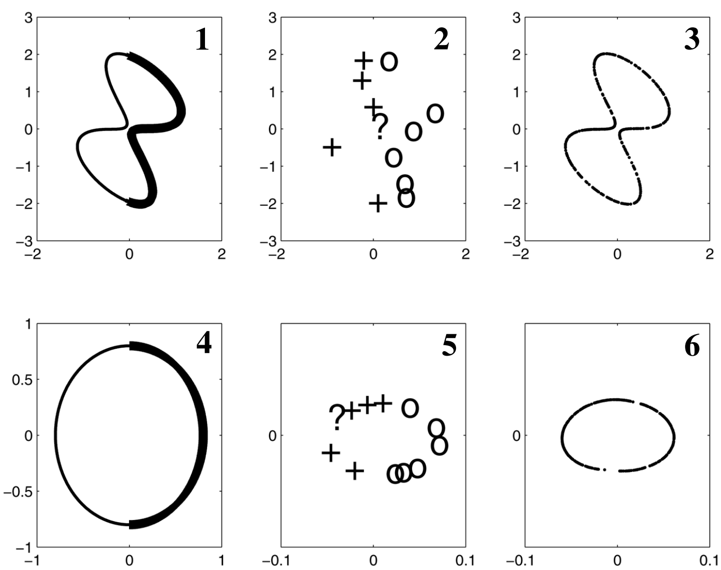

Figure 3. *The potential usefulness of unlabeled examples. The true distribution of patterns is on the manifold (a curve in this case), as shown in panel 1. Given only labeled examples (the circles and crosses of panel 2), it is hard to classify a novel example, such as that indicated by the question mark. If the manifold is known through unlabeled examples, however, it is seen that the data really lies on a closed curve (panel 3); classification becomes much easier with harmonic functions, which make it possible to re-embed the data points more naturally, as shown in panels 4 through 6. Panel 4 shows the ideal representation of the curve; panel 5, the embedding of the labeled examples; and panel 6, the embedding of all other examples.*

supervision and is expensive, the possibility of being able to learn with mostly unlabeled examples and only a small number of labeled examples is extremely attractive in many tasks (see Figure 3).

For Further Reading

T. Kenet, D. Bibitchkov, M. Tsodyks, A. Grinvald, and A. Arieli, *Spontaneously emerging cortical representations of visual attributes*, Nature 425:6961 (2003), 954–956.

F. Memoli and G. Sapiro, *Comparing Point Clouds*, IMA Technical Report, April 2004 (http://www.ima.umn.edu/preprints/apr2004/apr2004.html).