From *SIAM News*, Volume 42, Number 6, July/August 2009

Abel Prize Science Lecture: Revolutionary Work in Geometry and Shape Analysis

By Guillermo Sapiro

On May 19, at a ceremony held in Oslo, Mikhail Gromov received the 2009 Abel Prize from His Majesty King Harald of Norway. Gromov, who was honored "for his revolutionary contributions to geometry," divides his time between the Institut des Hautes Études Scientifiques in France and the Courant Institute of Mathematical Sciences at New York University.

Stylish Celebration of Outstanding Mathematics

The award ceremony was the highlight of a weeklong celebration in Oslo, held in great style and marked by warm Norwegian hospitality. On May 18, the Norwegian Academy of Science and Letters hosted a dinner in honor of the Abel Laureate, following a wreath-laying at the Abel Monument (the Norwegian mathematician Niels Henrik Abel died in 1829, at the age of 26). Present at the May 19 award ceremony, in addition to Gromov and King Harald, were Her Majesty Queen Sonja and Tora Aasland, the minister of research and higher education. The ceremony included speeches by Kristian Seip (current chair of the Abel Committee), Øyvind Østerud (president of the Norwegian Academy of Science and Letters), and Gromov himself. Beautiful music by trumpeter Arve Henriksen and the musicians of Trio Mediaeval (which many of us hurried to download from iTunes) added to the elegance of the ceremony.

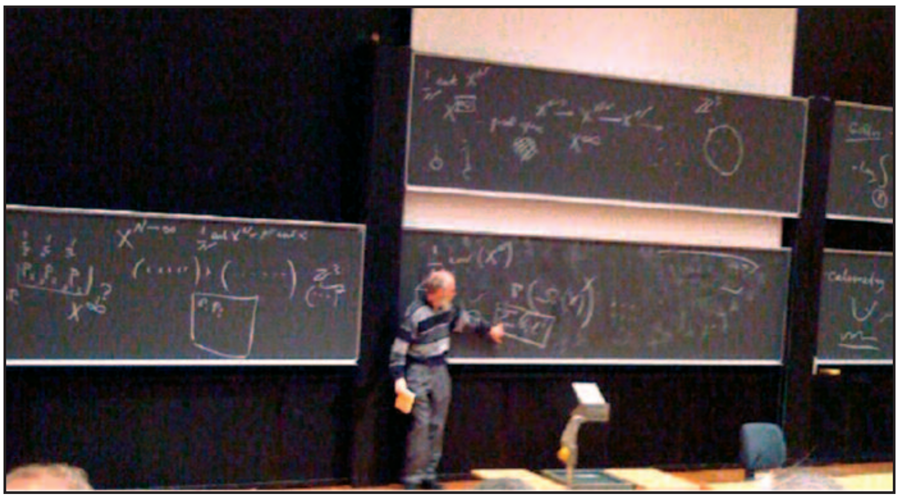

2009 Abel Laureate Mikhail Gromov. Photo by Guillermo Sapiro.

That evening, the Norwegian government gave a banquet in honor of the Abel Laureate at the Akershus Castle (and provided a unique chance to see a lot of mathematicians in tuxedos!). Ingrid Daubechies of Princeton University, a former member of the Abel Committee, gave a memorable speech in which she connected mathematics and science with Maurice Sendak, author of beloved children's books, including *Where the Wild Things Are*. The celebrations ended on May 20 with a reception at the Norwegian Academy of Science and Letters, where Enrique R. Pujals of IMPA (Brazil) was awarded the Ramanujan Prize for young mathematicians from developing countries.

The scientific component of the celebrations consisted of four Abel Lectures, given at the University of Oslo. Mikhail Gromov led off with an inspiring chalk-and-board lecture in which he showed, once again, how real problems from the physical sciences can motivate very deep and difficult mathematical challenges. Gromov's NYU colleague Jeff Cheeger followed with the second Abel Lecture, titled "How does he do it?" Along with some of Gromov's deep mathematical results, Cheeger recalled numerous interactions he has had over the years with the Abel Laureate. Through such anecdotes, he provided helpful insight into Gromov's extraordinary thinking and mathematical approach and philosophy. In the third Abel Lecture—titled "Geometry everywhere: Fiat lux!"—Martin Bridson of Oxford University focused on revolutionary ideas that Gromov has brought to discrete group theory.

Both Cheeger and Bridson made it clear that the important theorems proved by Gromov are just the beginning of his contributions; equally important are the actual proofs and methodologies, filled as they are with brilliant and unexpected innovations. The speakers provided numerous illustrations of Gromov's new, unpredictable, and very innovative approaches into different areas of mathematics. The lectures, in the words of the Abel Committee, are intended "to give a general audience a glimpse of the mathematics of the Abel Laureate and to convey to the general mathematician the importance and impact of his work." Cheeger and Bridson achieved this aim and more, leaving the audience with a clear understanding of the committee's decision to award the prestigious Abel Prize to Gromov.

The scientific component of the ceremonies ended with the Science Lecture, which "is intended for the broader scientific community and aims to highlight connections between the work of the Abel Laureate and other sciences." I had the honor of presenting this year's Science Lecture, which I titled "One small step for Gromov, one giant leap for shape analysis: A window into the 2009 Abel Laureate's contribution in computer vision and computer graphics."

The Gromov–Hausdorff Distance and Shape Analysis

One of the major challenges of computer vision and shape analysis is object recognition. This is an extremely difficult problem, in that we have to be able to distinguish thousands of object categories, each characterized by tremendous variability. Consider, for example, the many different types of chairs that have been fabricated, or the many different types of bikes in use over the years. The same object, moreover, can present itself in very different forms (Figure 1). The Science Lecture concentrated on the application of Gromov's contributions to perform three-dimensional object recognition that is invariant (or at least robust) to such deformations. At the heart of this work are Gromov's metric approach to shape and the Gromov–Hausdorff distance.

The Hausdorff distance, which is very well known in computer vision and shape analysis, is defined as follows:

$$(Z,d)\text{metric space ; } X,Y \subset Z$$
$$d_H(X,Y) :=$$
$$\max\left\{\sup_{x\in X}\inf_{y\in Y} d(x,y), \sup_{y\in Y}\inf_{x\in X} d(x,y)\right\}$$
$$d_H(X,Y) :=$$
$$\inf\{r>0 : X \subset U_r(Y) \text{ and } Y \subset U_r(X)\}.$$

Figure 1. *A subject can take on many different forms. From http://sarahbrumgart.com/images/main_yoga.jpg.*

Here $U_r(\cdot)$ stands for the r-dilation of the set. For shapes in 3D Euclidean space (in which case d is the Euclidean distance), it is often assumed in the shape analysis community that one of the shapes is allowed to optimally rotate and translate before the Hausdorff distance is computed and a shape comparison is obtained via

$$\inf_{\theta,\vec{t}} d_H\left(X, R(Y,\theta)+\vec{t}\right),$$

where R stands for a rotation. Computing the Hausdorff distance is impractical, and very good approximations have been developed (see, for example, [10] and references therein).

An important observation from the above definition is that Z is fixed. Even if rotations and translations are considered, or any d-metric-preserving transformation in the general case, the optimization does not include changing Z. The Gromov–Hausdorff distance (which Gromov, in [5], originally called simply *Hausdorff distance*), introduces flexibility in Z as well. More formally, Gromov considers shapes as metric spaces themselves, and distances between shapes accordingly become distances between metric spaces. For the case of shape analysis described here, the Gromov–Hausdorff distance is then defined as follows:

$$\text{Given 2 compact metric spaces:}$$
$$(X,d_X),(Y,d_Y)$$
$$d_{GH} := \inf_{Z,f,g} d_H^Z(f(X),g(Y))$$
$$f: X \to Z;\;\; g: Y \to Z$$
$$\text{isometric embeddings.}$$

Notice that the optimization includes Z. Isometric embedding in a predefined finite-dimensional space, such as Euclidean space, as commonly done in the vision and manifold learning communities, is in general suboptimal (an explicit example can be found in [5]). It is quite clear that d_{GH} is invariant to isometric transformations—that is, transformations of the object that preserve the intrinsic distances (often called *bends* in the computer vision literature). This is in part what motivated the introduction of the Gromov–Hausdorff distance into computer vision in [7,9].

The Gromov–Hausdorff distance d_{GH} has a number of important properties, one of which is that it is equivalent [3] to

$$d_{GH}(X,Y) := \frac{1}{2}\inf_{C} \sup_{(x,y),(x',y')\in C}$$
$$\left|d_X(x,x') - d_Y(y,y')\right|.$$

Here, C is a correspondence, meaning that every point in X has at least a corresponding point in Y, and vice versa (see Figure 2). In other words, the Gromov–Hausdorff distance d_{GH} measures the deformation of the intrinsic distances on the shapes (d_X and d_Y) when the shapes are matched. It is this interpretation that permits the development of computational techniques for approximating the Gromov–Hausdorff distance. Such techniques are based on explicit bounds [7,9], optimization approaches [1], or relaxations motivated by this distance [6,8].

Gromov's metric framework provides a lot of flexibility for the selection of the intrinsic pair-wise d_X and d_Y. In the work mentioned so far, geodesic distances have been used; recently, the use of diffusion distance [4] has been proposed [2]. This distance, which is related to and can be computed from the Laplace–Beltrami operator, is an average among all intrinsic paths (of certain lengths) connecting the corresponding points, and as such is significantly more robust to noise, topological changes, and missing data; see Figure 3 for examples of such scenarios. Any of these shape modifications can create drastic changes in the geodesic distance, but are in general not so critical for the diffusion distance. Therefore, when considering $d_X = d_Y$ = "diffusion distance" and the Gromov–Hausdorff distance d_{GH}, we obtain a bending invariant 3D shape recognition framework that is robust to real shape deformations like those depicted in Figure 3.

Figure 4 shows significant improvement in the optimal correspondence C with the diffusion distance in the presence of point-wise topological changes. In this figure, Voronoi regions around the corresponding points are shown in different colors, and we can see how the colors on the right (diffusion) match much more nicely than those in the middle (geodesic), indicating a better point-wise matching, as in the ground truth (left). This is reflected in a significantly improved 3D shape clustering and recognition [2]. Figure 5 shows the pairwise Gromov–Hausdorff diffusion distances for sets of shapes undergoing all the changes illustrated in Figure 3 (darker colors indicate lower d_{GH}); as expected, objects from the same class are closer.

Figure 2. *Shapes come with a metric, indicated here by pairwise distances and paths between points.*

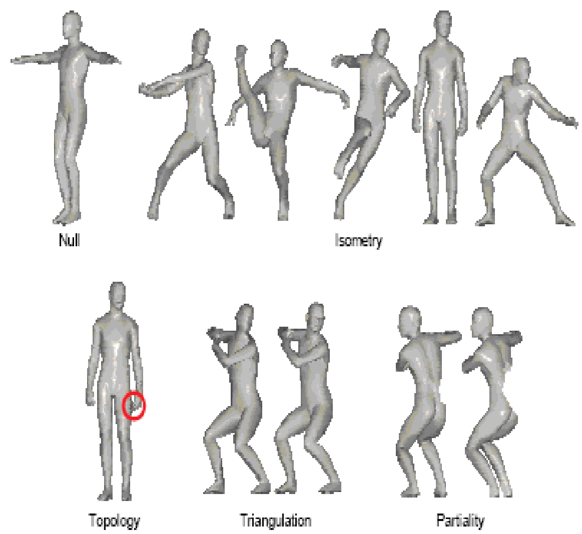

Figure 3. *Common shape deformations of various types, including bends, topological changes, triangulation differences, and missing data.*

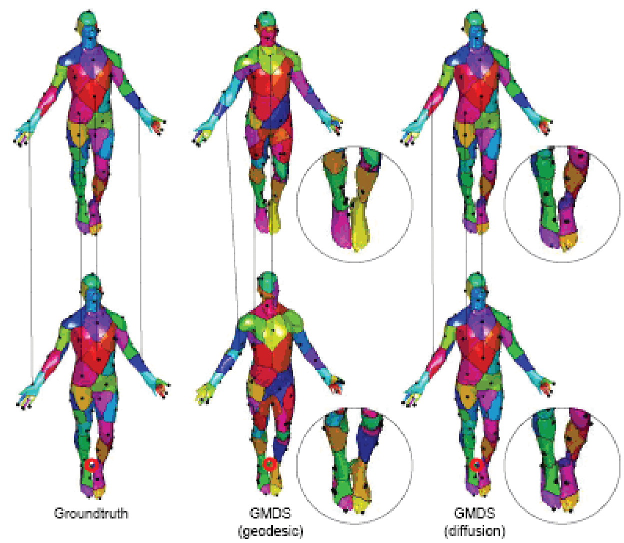

Figure 4. *Considering diffusion distances (right) in the Gromov metric framework for shape correspondence leads to better matching than with the geodesic distance (middle), as indicated by the corresponding matching colors for the same body parts (ground truth on the left).*

This short article only begins to introduce Gromov's work in computer vision. Without even covering all of chapter 3½ in his *Green Book* [5], it has already mentioned leading results in 3D shape analysis, recognition, and matching. Applications in 3D shape analysis also open a number of questions in the metric approach to shape analysis introduced by Gromov that have yet to be answered. Among them are the development of a computational theory with tight bounds, like the one developed for the Hausdorff distance by the computational geometry community; further analysis of the shape manifold resulting from the Gromov–Hausdorff distance; and development of further connections with the Cheeger isoperimetric constant when the diffusion distance is used. As indicated in the title of the Abel Science Lecture, this represents just a small step for Gromov, and a very small window into his revolutionary contributions.

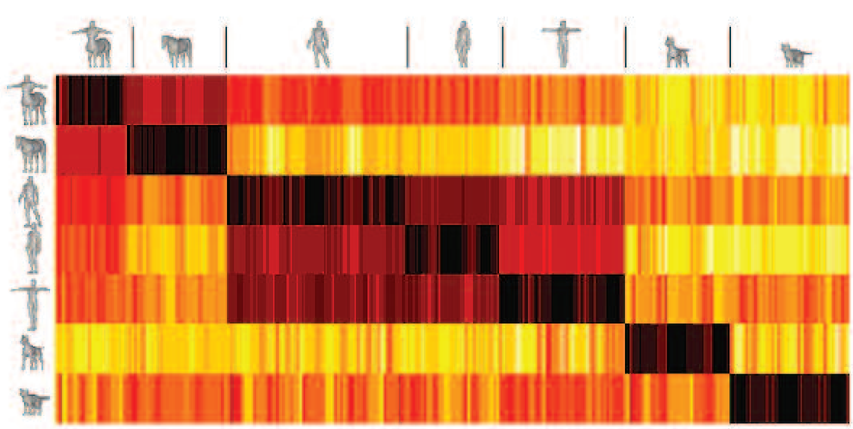

Figure 5. *Distances from null transformations of each shape class (rows) to all instances in the dataset (columns), which includes topological deformations, bends, missing data, and changes in the triangulation (see Figure 3). Brighter color represents larger distances (lower similarity).*

Acknowledgments

I thank Facundo Mémoli, with whom the travels into the Gromov world started during his PhD; my colleagues and co-authors of [2]; the Abel Committee for giving me the opportunity to present the Abel Science Lecture; and Doug Arnold for guidance. Special thanks of course to Professor Gromov for hosting me at IHES in June 2009 and for creating a structure as beautiful as the Gromov–Hausdorff distance.

References

[1] A.M. Bronstein, M.M. Bronstein, and R. Kimmel, *Numerical Geometry of Nonrigid Shapes*, Springer, New York, 2008.

[2] A.M. Bronstein, M.M. Bronstein, M. Mahmoudi, R. Kimmel, and G. Sapiro, *A Gromov–Hausdorff framework with diffusion geometry for topologically-robust non-rigid shape matching*, IMA Preprint, 2008.

[3] D. Burago, Y. Burago, and S. Ivanov, *A Course in Metric Geometry*, American Mathematical Society, Providence, RI, GSM 33, 2001.

[4] R.R. Coifman and S. Lafon, *Diffusion maps*, Appl. Computat. Harmonic Anal., 21 (2006).

[5] M. Gromov, *Metric Structures for Riemannian and non-Riemannian Spaces*, Birkhäuser, Boston, 1999.

[6] M. Mahmoudi and G. Sapiro, *Three-dimensional point cloud recognition via distributions of geometric distances*, Graphical Models, 71 (2009).

[7] F. Mémoli, *Estimation of Distance Functions and Geodesics and its Use for Shape Comparison and Alignment: Theoretical and Computational Results*, PhD Thesis, University of Minnesota, May 2005.

[8] F. Mémoli, *On the use of Gromov–Hausdorff distances for shape comparison*, Proceedings of the Symposium on Point Based Graphics, Prague, 2007.

[9] F. Mémoli and G. Sapiro, *A theoretical and computational framework for isometry invariant recognition of point cloud data*, Foundations of Computational Mathematics, 5 (2005).

[10] M. Tang, M. Lee, and Y.J. Kim, *Interactive Hausdorff distance computation for general polygonal models*, ACM Trans. Graphics (SIGGRAPH), 2009.

Guillermo Sapiro, editor-in-chief of SIAM Journal on Imaging Sciences*, is a professor in the Department of Electrical and Computer Engineering at the University of Minnesota.*