Algebraic vision

organized by
Sameer Agarwal, Max Lieblich, and Rekha Thomas

Workshop Summary

This Workshop brought together a diverse range of researchers from throughout algebraic geometry, differential geometry, optimization, and computer vision, to explore a new area we call Algebraic Vision.

Computer vision is a fundamental part of the modern world, used everywhere from the Defense Department to highway tolls to mobile phone panoramic images to augmented and virtual reality. At its core lies a hard kernel of classical projective and Euclidean geometry. While statistical methods are often used to solve large problems in computer vision, the fundamental mathematical aspects of the underlying problems play an absolutely essential role in algorithms enabling computations that are fast enough to be practical. As an example, cutting-edge computational methods use hard-coded Groebner basis calculations, arising from the solution of a system of polynomial equations describing the projective geometry of pairs of camera images, to perform camera position and orientation estimates in a single microsecond. This is essential because thousands of such estimates need to be performed to overcome measurement noise, and if one is to track (for example) video data, these thousands of estimates must all be done 30 times per second.

At the same time, there are numerous beautiful and fundamental mathematical questions about the geometry of computer vision. For instance, the process of 3d reconstruction from sequences of 2d images can be understood using certain incidence correspondences in products of projective spaces. These incidence correspondences themselves have interesting moduli, and in general (for example) they form a dense subset of a single component of the Hilbert scheme. The real algebraic geometry of linear sections of these incidence correspondences is the key object of study for the reconstruction problem; while the complex geometry is not too difficult to understand, the real points of the resulting varieties are still only vaguely understood.

Some topics common to algebraic geometry and computer vision have developed over time along separate threads that have been only loosely woven together. One beautiful example of this that arose at the workshop has to do with the edges of 2d images of 3d objects. Computer vision practitioners have studied the images of algebraic surfaces; in 1996, Forsyth famously showed that (roughly speaking) the edge of the shadow determines the original surface, up to projective transformation. However, it turns out that this problem, in algebraic geometry, dates back to Chisini in 1944, and it was solved by Moishezon in 1981 in the case of a generic projection (which is the case considered by Forsyth). Even though Forsyth consulted with algebraic geometers in preparing his paper, the Chisini problem and Moishezons resolution of it for generic projections is sufficiently obscure within the subject to have evaded detection.

Organizing a workshop on algebraic vision is complex. Computer vision uses mathematical language to describe phenomena that are not entirely mathematical. Mathematicians, generally speaking, dont know how computationally feasible or efficient their theorems are. The techniques that are common knowledge on one side (for example, RANSAC for computer vision or the functorial approach for algebraic geometry) are mysterious on the other. As such, rather than focusing from the start on two or three problems, much of our time at the workshop was devoted to bringing these groups together. Each day had a theme. In the morning there were two or three talks about the theme given by experts in the area.

Here are the daily themes.

Monday: Introduction to vision (with two broad overview talks by Hartley and Ponce)

Tuesday: Multiview varieties and tensors (with talks by Sturmfels, Ottaviani, and Lieblich)

Wednesday: Minimal problems and computational algebraic geometry (with talks by Pajdla, Hauenstein, and Safey El Din)

Thursday: Shapes, silhouettes, and contours (with talks by Trager, Olver, and Kogan)

Friday: Carlsson-Weinshall duality (with a talk by Kahl) and a morning problem session

The first two afternoons were devoted to cross-fertilization in ask the experts sessions, where participants could learn more about particular topics from other participants with expertise in those topics. The second afternoon also featured a problem session. Other afternoons were devoted to working on problems that arose throughout the workshop.

The workshop stimulated a tremendous amount of discussion and thinking, and led to the creation of several new collaborations. We expect this activity to continue and expand into the future.