SIAM Activity Group on Imaging Science:

Forward-Looking Panel Discussion

Tuesday April 13, 2010

Fadil Santosa led the discussion and had invited a panel of four speakers:

1. Michael Bronstein (Technion)
2. Garrett Funka-Lea (Siemans)
3. Yali Amit (Chicago)
4. Josh Nolting (GeoEye)

Between 20 and 30 people were in attendance at the beginning of the panel. Each panelist spoke for 10-15 minutes, following which there was a short discussion.

Michael Bronstein (Technion)

Title: Computational metric geometry: an old new tool in image sciences

Dr Bronstein discussed metric geometry and its applications in imaging science. He mentioned that using metric geometry may have applications in image retrieval, categorization, tracking, detection/recognition as well as restoration and alignment. He made the point that the choice of metric is key as the metrics define the invariance of the system. Since the idea is to map the similarity of images to some metric space, the metric depends on the sought transition from data space to the embedding space. He mentioned briefly that these techniques could be used for detecting video piracy.

Dr Bronstein identified the following challenges:

- Theory: Approximate symmetry notion: group-like structure, comparing data from different spaces.
- Computational: Efficient solution of minimum distortion correspondence problems (Gromov-Hausdorff distance), efficient algorithms for embedding into interesting metric spaces.
- Applications: problems that can be formulated in terms of metric geometry.

Garrett Funka-Lea (Siemans)

Dr. Funka-Lea discussed primarily the roadblocks and successes in getting imaging technology into the hospitals. He began by discussing how image analysis was initially sold or designed as a way of replacing doctors; he brought up the basic problem that people are not willing to trust the computer to make decisions about their health and well-being. He emphasized the need to develop a confidence measure through the use of large-scale clinical trials and through improvements in the statistical understanding of the algorithms and their likelihood of incorrect diagnosis. He advocated doing things that would fit into the work-flow of how doctors already do things. Specifically, he mentioned that reconstruction algorithms are generally trusted; still improvements that can be made (e.g. faster algorithms), but doctors can't design them themselves and they trust the images.

The need to understand and merge multiple data sets, joint imaging with multiple modalities, and the management of the sorts of large data sets generated are also important problems of current interest. Where he sees the most room for improvement is in visualization; e.g. being able to segment and isolate the heart, make it easier for the doctors to see what they want to see anyway, rather than diagnosis by computer. There was some discussion of the ability (or lack of) to train computers to do conceptual learning and inductive reasoning to help correlate multiple data sets. He also mentioned that the most effective approaches seem to perform the tasks that the doctors do not like, for example, the doctor identifies a tumor on a 2D slice of an image and the computer finds its edges in 3D.

Yali Amit, University of Chicago

Dr Amit discussed current issues in computer vision, and the goal of having computers interpret images in ways similar to how our visual system does. Current techniques are quite successful at detecting objects in images, but are not able to make semantic interpretations of the resulting images. He emphasized the need to develop approaches that can work with less training data than the thousands of examples and billions of non-examples that is used in current techniques. He thinks the future is along these lines, developing algorithms that can adapt themselves and learn in a less primative way than the billion examples used now. Specifically he mentioned that variability of the objects are not generally modeled and that for these sorts of techniques to scale to larger problems, training must become possible with smaller data sets. He sees promising directions in including statistical tools and in creating generative models for the data at the local or object scales.

Josh Nolting (GeoEye)

Dr Nolting discussed several specific aspects of image processing: feature extraction, registration, and change detection. He mentioned that in general these tools do not work in sufficient generality to be of use without a great deal of supervision, as a result, most of these algorithms are not used in practice. He emphasized the cost of fixing errors, that if a person has to go through and move or redraw the segmentation, for example, it takes more time than it would take the person to segment the image manually to begin with. He emphasized the need for accuracy over a reasonable range of problems from a single algorithm; he used the example of finding the middle of a road, not the general vicinity of the road and that this must be robust to slightly different scenarios. For change detection, the current ratio of change detection to false alarms is not large enough, there is a lot of change from view angle and lighting and we need algorithms that are capable of distinguishing this from actual changes in the scene. He sees a great deal of potential in utilizing threading and GPUs to process multiple images simultaneously as well as to speed up processing algorithms.

Discussion

Following the talks of the panelists, there was a short discussion with audience members. A few key points of that discussion are listed here.

- The question was asked whether there are completely new ways of imaging, needed or in the pipeline, for something you really want to detect or measure, beyond MRI, CT etc? This was followed briefly by Dr Funka-Lea mentioning that MRI is something of a black art, with a large parameter space, from one hospital to the next it may be different, there is lots of progress and processing; this also happens in different modalities but is most notable for MRI.

- The question was also raised as to whether doctors can do a really good analysis with no prior info besides the image. The answer was that for obvious problems yes. For newer imaging modalities they see a lot of things that they're not sure whether they're normal or not, which is problematic with full-body scans for example.

- A point was raised about measurements vs judgments, perhaps we don't want to have the computers making decisions; the counter-point that the goal is to reduce the time required to go from imaging to diagnosis was also discussed.

- There was some discussion of the nature of human vision and its relation to how we train computers to identify objects. The question was raised as to whether our visual system is just a big database? The point was made that the brain is not a nearest neighbor classifier, these don't actually do so well in reality, there is a lot of structure in the visual-neural system – it is not necessarily the case that computer vision should mimic the human visual system.

- There was some discussion as to whether we do have 'data retrieval systems' in our brains that allow us to access quickly and accurately past experience, e.g. champion chess players, and whether that might be a good model for computer learning as well.

- The point was made that while we shouldn't solve problems that are more difficult than we need, we do need to always keep in mind the next step. Simple training does not go there.