Reconstruction and Recognition

ICVSS 2008

Sicily ~ 14-19 July

international computer vision summer school

HOW TO CONDUCT A LITERATURE REVIEW
AND DISCOVER THE CONTEXT OF AN IDEA
(WITH PRIZE!)

Stefano Soatto
UCLA, USA

The prize was assigned to:
Pashmina Bendale, University of Cambridge, UK

READING GROUP ABSTRACT

Let's say you attending a conference and find a particular talk intriguing, or a particular idea catchy. Maybe you are considering working in the same area, or expanding on the idea as part of your research. The first thing to do, your advisor would tell you, is to read up on the matter. If the authors have done their homework, you will go to the list of references, and that would give you the context in which their idea originated. Unfortunately, it is often difficult to discover all the work that is related to a given idea, and a recent experiment indicates that often the authors are not all that knowledgeable about the work they actually cite (M. Simkin and V.P.
Roychowdhury, "Read before You Cite!", Complex Systems, 2003, 14, 269-274, indicate that authors only read about 20% of the papers they cite). So you do not really want to rely entirely on the list of references provided at the end of the paper, and you want to cast your net a little wider.
But where to start? Google, you say. Ok. But unfortunately most search engines do not index work prior that, say, 1980. And even if they did, how do you choose, among the thousands of hit, which ones to read? The most cited ones, you say. Fine, but that may get you trapped in a local minimum in which - after a few people start citing a given work, everyone else follows, including you, and you miss out to all the work leading to that idea, that may contain seeds of more interesting ideas or development.

It is important to trace back an idea to its root not only to give proper credit to their originators, but also to avoid re-discovering the same concept over and over, which slows down progress of the academic community as a whole (a small degree of rediscovery of the wheel is unavoidable, and sometime healthy).

As part of this discussion, we will do a scholarly exercise. Let's pretend that the talk that got your interest in the first place used three ideas: Ada-boost, Mean-shift, and SIFT. Prior to coming to the school, your homework is to research the literature and trace the roots of these three ideas, as far back as you can (ok, Aristotle may be a bit of a stretch, but don't stop at 1999 either). During the session, which will be interactive, each student will explain to the class the process by which he or she arrived at the roots of an idea (not necessarily the first paper that introduced "SIFT" , "mean-shift" or "ada-boost", which would be trivial, but the papers that introduced the concepts that are at the core of these ideas).

The outcome of this exercise will be a tree, having the paper of interest as root, and going back as far as possible. To add some spice to the exercise, I will offer a prize to the student that can link the highest number of papers tracking back the idea, that do not appear in standard search engines (say Google Scholar). 100$ for each paper that the student can argue is meaningfully related to the idea in question, with a 10% premium if the authors are deceased, up to a maximum of 1,000$.

Homework/Syllabus

Students are required to read up on Ada-boost, Mean-shift and SIFT, and follow the guidelines above to arrive at the school with a report, in the form of a brief document, describing the dependency tree discovered in the exercise.