Tutorial

MAP inference in Discrete Models.

Pushmeet Kohli

Microsoft Research, UK

Many problems in Computer Vision are formulated in form of a random filed of discrete variables. Examples range from low-level vision such as image segmentation, optical flow and stereo reconstruction, to high-level vision such as object recognition. The goal is typically to infer the most probable values of the random variables, known as Maximum a Posteriori (MAP) estimation. This has been widely studied in several areas of Computer Science (e.g. Computer Vision, Machine Learning, Theory), and the resulting algorithms have greatly helped in obtaining accurate and reliable solutions to many problems. These algorithms are extremely efficient and can find the globally (or strong locally) optimal solutions for an important class of models in polynomial time. Hence, they have led to a significant increase in the use of random field models in computer vision and information engineering in general. This tutorial is aimed at researchers who wish to use and understand these algorithms for solving new problems in computer vision and information engineering. No prior knowledge of probabilistic models or discrete optimization will be assumed. The tutorial will answer the following questions: (a) How to formalize and solve some known vision problems using MAP inference of a random field? (b) What are the different genres of MAP inference algorithms? (c) How do they work? (d) What are the recent developments and open questions in this field?

Pushmeet Kohli is a research scientist in the Machine Learning and Perception group at Microsoft Research Cambridge, an associate of the Psychometric Centre and Trinity Hall, University of Cambridge. Pushmeet was awarded his PhD from Oxford Brookes in 2007 and was the first of Phil Torr's students to graduate from that group.

Pushmeet's research revolves around Intelligent Systems and Computational Sciences, and he publishes in the fields of Machine Learning, Computer Vision, Information Retrieval, and Game Theory. His current research interests include 'human behaviour analysis' and the 'prediction of user preferences'. Pushmeet is interested in designing autonomous and intelligent computer vision, bargaining and trading systems which learn by observing and interacting with users on social media sites such as Facebook. He is also investigating the use of new sensors such as KINECT for the problems of pose estimation, scene understanding and robotics.

Pushmeet has won a number of awards and prizes for his research. His PhD thesis, titled "Minimizing Dynamic and Higher Order Energy Functions using Graph Cuts", was the winner of the British Machine Vision Association's 'Sullivan Doctoral Thesis Award', and was a runner-up for the British Computer Society's 'Distinguished Dissertation Award'. Pushmeet's papers have appeared in Computer Vision (ICCV, CVPR, ECCV, PAMI, IJCV, CVIU, BMVC, DAGM), Machine Learning, Robotics and AI (NIPS, ICML, AISTATS, AAAI, AAMAS, UAI, ISMAR), Computer Graphics (SIGGRAPH, Eurographics), and HCI (CHI, UIST) conferences. They have won best paper awards in ICVGIP 2006, 2010, ECCV 2010 and ISMAR 2011. His research has also been the subject of a number of articles in popular media outlets such as Forbes, The Economic Times, New Scientist and MIT Technology Review. Pushmeet is a part of the Association for Computing Machinery's (ACM) Distinguished Speaker Program.

http://dx.doi.org/10.5244/C.26.2