UNIVERSITY OF CALIFORNIA, BERKELEY

BERKELEY · DAVIS · IRVINE · LOS ANGELES · RIVERSIDE · SAN DIEGO · SAN FRANCISCO

SANTA BARBARA · SANTA CRUZ

DEPARTMENT OF ELECTRICAL ENGINEERING
AND COMPUTER SCIENCES

Deva Ramanan
Computer Science Division
547 Soda Hall, Mail Code 1776
Berkeley, California 94720-1776

Phone: (510) 642-6193
E-Mail: ramanan@eecs.berkeley.edu

December 23, 2004

To the Faculty Search Committee:

I am writing to apply for a tenure-track assistant professor position. My background is computer vision and machine learning. I expect to file my dissertation no later than June 2005.

My graduate work has focused on the analysis of human movement in video. This is a hard problem because body parts can be small and can move fast & unpredictably. The vision community has struggled with this challenge for almost 50 years. This technology, when mature, has many applications; video summary/mining, Human-Computer Interaction (HCI), motion capture (for computer graphics), human behavior analysis, and security surveillance (a topic of great interest nowadays).

The main insight we bring is a new model of uncertainty in video data that links tracking and object detection. Combining ideas from both research areas, we develop the first demonstrably reliable system for large-scale people tracking. We introduce a novel evaluation methodology for tracking and applied it on a massive scale, evaluating our tracker on *hundreds of thousands* of frames (several orders of magnitude more than previous systems).

Our system works by automatically learning models of appearance for the object being tracked. These appearance models are a sophisticated class of object models known as pictorial structures (which have existed for thirty years in the vision community). Our algorithm is the first and only known method for learning pictorial structures. Using these models to detect objects in other images, we introduce the paradigm of using video to learn models for object detection. Other research groups have now adopted this approach.

Our model reveals fundamental flaws in basic algorithms (such as Kalman and particle filtering) commonly used for tracking. These insights apply when tracking any object, whether it be a face, a hand, or a car. Our observations have implications for other tasks that I would like to explore, such as object detection and video compression/encoding. Given our video model, we show our tracking algorithm to be an instance of approximate inference on that model. I am interested in looking at other approximate algorithms from the machine learning community, both in general and in the context of vision problems.

Our work has almost turned people-tracking into a usable technology. We are at the point where we can begin addressing the exciting applications above. We have used our tracker to identify activities by taking an "analysis as synthesis" approach for activity recognition. We synthesize a novel motion by cutting and pasting existing clips of annotated motion such that the new motion looks like the tracked frames. The synthesized annotations serve as the activity descriptions. By linking tracking (a vision task) with motion synthesis (a graphics task), we again benefit from technology transfer between different fields.

Given that I had a fellowship for all my graduate years, my formal teaching experience is limited to TAing one semester for the undergraduate *Artificial Intelligence* course at Berkeley. (Informally, our lab has a great tradition of older students mentoring the newer ones). I did acquire practical knowledge of how to manage a class; I held weekly discussion sections, office hours, and helped write and grade the exams. In general I quite enjoy speaking in front of groups, whether it be a classroom lecture or a conference presentation. As such I had a wonderful time TAing, and I look forward leading my own class in the future.

As requested, I have enclosed my CV (including references), research statement, and teaching statement. The easiest way to contact me is through e-mail, ramanan@eecs.berkeley.edu. I look forward to hearing from you.

Sincerely,

Deva Ramanan
EECS, Computer Science Division
University of California, Berkeley.