

FUNDAMENTALS AND OVERVIEW

At the heart of my research interests is statistical machine learning (*e.g.*, automatic data analysis and modeling, computational approaches for learning from experience), which I also join with more specialized fields: HCI, computer vision/graphics, and data visualization/clustering, and more recently bioinformatics.

My central research interest is to create algorithmic and mathematical concepts and tools that would help computers exhibit forms of data processing skills akin to what has been regularly associated to human (or biological system) intelligence; hence my work in the above specialized areas. This is not intended to limit our scope, since in many cases it is possible to augment these skills. Examples include: characterizing intrinsic variability of observed phenomena (dimensionality reduction, compression, manifold learning), learning concepts from examples (visual recognition, language understanding), solving tasks from apparently insufficient information (realizing 3D scene attributes from a 2D picture, speech understanding from a faulty radio, protein folding from amino-acid sequences), providing remarkable quick and good answers to amazingly large (fairly unstructured) problems (game playing, recognition, search, system self-organization), and using high-order concepts (learning to learn). My research consists of employing and building mathematical and algorithmic concepts for approaching problems like these. Challenging open problems and fascinating new concepts, some yet to be formalized, make machine learning and related application fields an extraordinary area of research.

The field has already provided us with valuable applications in areas such as signal processing, coding and compression, recognition, data mining, and vision/graphics. I believe in the complementarity of theory and practice. Theory is vital to creating *the big picture* and to formally advancing our level of understanding. It compactly represents the essence of a vast number of observations and experiences. Theory possesses great academic importance of its own. Applications provide insights, define new problems, and motivate. More importantly, applications satisfy many human, social, and economical needs, and therefore constitute a primordial area of development, which have and will continue to attract industrial and governmental interest. This interest has even been accentuated with the rapid technological development observed in the past decade (processing and storage power) and data availability (high-throughput methods). Thus, data analysis in general, must be part of the future plans of any competitive and leading institution.

PAST EXPERIENCE AND ACHIEVEMENTS

I believe that statistics and probabilistic inference offer one of the most useful principled means to formalize and analyze the problems just described (*e.g.*, through graphical models). However, this setting by no means provides all the answers. In fact, probabilistic inference is computationally intractable in many interesting cases. As a result, many open problems have surfaced and connections to a diversity of fields have been drawn. My experience is based largely on this view. Yet, different perspectives (*e.g.*, non probabilistic) have proven important in our understanding; thus, they can be complementary and of interest. New paradigms are essential.

My interest in the field started during my undergraduate studies in operations research and optimization. I studied methods for queue analysis and performance enhancement in a computer network serving requests. During my master's years, with my advisor Stan Sclaroff, I got truly interested in computer vision. My master's thesis dealt with visual surveillance and human-computer interaction. I developed a method for localization and tracking of moving targets such as people and designed a way to recover their 3D trajectory based on 2D positions in an image sequence. This later allowed me to develop a motion recognition system that could, under certain conditions, differentiate between human actions independently of camera position. This work generated extensive industrial interest. General view-independent recognition remains an unsolved problem in vision.

Later, I became directly interested in machine learning and anticipated that probabilistic inference could be used for recovering 3D articulated structure from a single 2D image (a classic ill-posed problem). However, I found that exact inference was intractable for this problem. As part of my PhD., with Stan Sclaroff, I developed an original method to solve the 3D from 2D problem using approximate inference by creating a novel way of integrating generative and discriminative models. This was one of the first approaches that could perform this vision task in an automatic way, and also represented a departure from previous paradigms that tried to solve the problem by using purely camera geometry or tracking. This method has many important uses, including HCI, motion analysis, surveillance, etc. Although motivated by vision, it can be used more generally in intractable

inference problems where one has access to a relatively accurate generative model. In collaboration with Joni Alon, Vassilis Athitsos and Matheen Siddiqui we later extended the method to allow multiple cameras. During an internship at MERL, with Matt Brand, I helped in developing a method for synthesizing human motion from tracking data using HMM's and entropic priors. While working on my PhD., I also became involved in more general problems and concepts in machine learning. I joined Tommi Jaakkola's machine learning group meetings at MIT, who also provided excellent research advice and deep insights into the fundamental problems.

In 2001 I joined the artificial intelligence group at University of Toronto as a postdoctoral fellow where I approached several problems. With Kannan Achan and Brendan Frey we proposed a new general way to perform image processing in an unsupervised fashion. We proposed the idea of *learning image processing tasks* by simply providing example images with the desired statistics or properties that we wanted the degraded or input image to acquire. This method had immediate applications for reconstruction and restoration, edge detection, etc., and was also ideal for non-photorealistic rendering, texture transfer, and other graphics tasks. As an example, it allowed us to automatically re-render an image in the style of another image by simply providing the two images in question. Remarkably, using this idea, all of these tasks were unified by the same basic principle: approximate probabilistic inference in factor graphs where the new image was formed by latent patches, related to the input image by a finite set of unknown transformations. With Brendan Frey, I also developed a probabilistic interpretation and extensions of spectral clustering (SC). We showed how to cast SC into statistical inference and also provided an algorithm that learns the similarity function between data points and allows us to use class-dependent similarity measures, by looking at graph-encoding latent random variables. This method has potential connections to (mixture) of manifolds learning. Also, we have formulated a related method for learning to cluster by example. At a more abstract level, we are currently developing new methods for approximate inference in factor graphs with high-order factors (large clique size). We are studying a method that based on information theory measures, can guide the reduction of complexity during message passing.

THE FUTURE

I aim to focus on both theoretical progress and practical applications. Specifically, I will study methods for approximate inference, especially in models with high-order random dependencies, a problem that has not yet received the adequate attention. Intractability is the stone that hinders a broader use of probabilistic methods in complex real-world applications. I believe that the analysis of the information theoretic aspects of the problem is one of the best paths to progress. Brendan Frey and I have gained some insight into potential solutions. For inference, I am also concerned with new paradigms for efficiently exploring high posterior probability regions.

There is a captivating connection between inference algorithms and combinatorial optimization problems (e.g., belief propagation and K-SAT), and by further investigating these connections, years of theoretical findings could be exploited for joint progress. I plan to study connections to other fields. Other examples are: (1) convex analysis/optimization, which can be employed to create new learning concepts (e.g., for semi-supervised classification) or approach old, yet unsolved problems (e.g., approximate inference), and (2) information theory, which can be used, to guide information exchange applications (e.g., HCIs, data retrieval).

In other applications, by extending my earlier work, the concept of learning image processing by example can be further generalized to digital signal processing. This project approaches very common type of tasks, and shows a user-intuitive way to define them. It is very relevant for computer graphics, in particular image-based rendering and texture generation. The project is expected to attract considerable industrial attention (e.g., entertainment industry). I expect to keep exploring articulated 3D reconstruction, for this there are collaboration plans with Sclaroff's group. Substantial funding interest, especially on HCI also make this project attractive.

In data mining, partly based on my previous clustering work, I intend to *learn to cluster* using spatial, graph-based relationships among points in labeled data sets. This is relevant for understanding/visualizing data and is tied to (partly/fully) unsupervised classification. I have also gained an understanding of some fundamental bioinformatics problems, an area of major growth, relevant to data modeling. Some interesting problems are modeling protein folding, improving the accuracy of micro-array measurements, and analyzing DNA sequences.

Vast potential for support, development, and interaction with other fields along with challenging and motivating problems make these remarkably exciting areas of research. I pursue the development of ground-breaking algorithms and concepts in machine learning and data analysis in general, with emphasis in HCI, vision, data mining/visualization, and bioinformatics. I believe that my appeal for and experience with research conducting in academia will help create and maintain a favorable and productive research environment.

I enjoy witnessing excellence in teaching, an occupation that I like very much. Thus, I am particularly attentive to what sets excellent teaching apart from the standard.

After attending classes for more than 20 years, five of these in undergraduate and at least five more in graduate courses in Boston University, MIT, and University of Toronto, I have developed a valuable concept of what differentiates a well taught from an inadequately taught class, and an inspiring from a hollow one. My teaching, TA-ing, and mentoring experience have greatly contributed to the development of this concept.

My experience in teaching goes back to my undergraduate studies, where I taught the mentoring sessions in statistics and calculus for full-sized classes, as part of my undergraduate scholarship duties. Later, I taught diverse programming tutorials when working for the computer lab. During my graduate studies in Boston University I benefited from a research scholarship, however I chose to be the teaching assistant for several graduate classes. I was the TA for the programming languages class taught by Prof. Zlateva, the computer vision class, the introductory, and the advanced computer graphics classes, taught by my advisor Prof. Sclaroff. I was also the mentor for students involved in undergraduate research in vision. I gained extra experience during four semesters as the lab tutor, giving the lab lectures and also occasionally the class lectures for Profs. Veilleux and Zlateva, in the areas of data structures and OO C++ programming at introductory and advanced levels. In addition, I gave tutorials and invited lectures in vision and learning classes, and as part of my research work I give more frequent presentations to larger audiences in conferences and as invited speaker. As part of my postdoctoral teaching work at Toronto, I have been co-advisor of a Ph.D. student and given tutorial presentations on machine learning and vision topics. I have found my experience rewarding and also instructive to myself. I believe all of this experience positively contributed to the development of my teaching concept.

The aspects that I consider to be the key for first-rate science teaching are: (1) No class is complete if there is not attempt to connect the specific topic with other, more familiar concepts. I find this particularly useful when introducing complex material since it is a means to present the new concepts within the *big picture*. This is one of the most useful pieces of information for learning and it is also critical in intuition building. (2) It is imperative to give students the opportunity to get used and handle the rigorous and formal theoretical foundations and definitions of the field. At the adequate moment (not always at first), concepts have to be presented with the suitable mathematical rigor; otherwise the student may lose his/her chance to being able to communicate at the appropriate scientific level. (3) Teacher motivation and stimulus provides the basic driving force for the student enjoyment of the presented material and possibly his/her later research interests. This is of particular importance at the introductory level, such as undergraduate. (4) A remarkably important aspect that I stress is building critical skills by proper exposure to different viewpoints.

In order to achieve these goals, I favor both blackboard (for hands-on explanation) and slide presentations (for illustrations and concepts requiring longer design time). I encourage active class participation, reserving part of the grade to encourage this (a concept that I though very positive during my undergraduate). I consider written (paper) homework assignment and mathematical exercises, along with quizzes, necessary for consolidating what was learn in class and for building the student's confidence. Programming assignments are also very constructive; however, I believe that no time should be wasted in projects that require long programming hours with little or no substance. I favor introducing undergraduate students to more advanced projects by partitioning responsibilities. Guidance is critical here, but I expect the results to be more rewarding.

Since my field of expertise is multidisciplinary, it can be used to diversify the teaching areas. I am qualified to teach a variety of classes, depending of the department needs, such as computer vision, machine learning (artificial intelligence), probabilistic inference, information theory, computer graphics, data mining, Bayesian probability, signal processing, algebra, discrete math, and more basic data structures or programming classes. I consider it positive and appealing to have classes where different areas collaborate in non-traditional ways, *e.g.*, machine learning and graphics or learning and visualization. I would enjoy creating graduate research seminars for more specialized topics in my field of expertise. I am aware of the need for regular revision of class content, offerings, and planning in an academic environment. In case I am given the opportunity, I will be very much interested in collaborating at this level.

In general, I am confident that my strong commitment to excellence in teaching, along with my experience and knowledge of the field should help me provide a valuable contribution to any department I join.