

THE FUNDAMENTALS

At the heart of my research interests is machine learning (roughly described as data analysis and modeling), which I join with more specialized fields: computer vision, graphics, data mining and recently bioinformatics.

My central research interest is to create algorithmic and mathematical concepts and tools that would help computers exhibit some form of data processing skills akin to what has been regularly associated to human (biological system) intelligence; hence my work in the above specialized areas. Examples of these processing skills are: characterizing intrinsic variability of observed phenomena (dimensionality reduction, compression, manifold learning), learning concepts from examples (visual recognition, language/speech understanding, gene finding), 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 (sometimes fairly unstructured) problems (game playing, motion planning, recognition, 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 this field an extraordinary area of research.

This 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 create the big picture of knowledge, to formally advance our level of understanding. Theory compactly represents the essence of a vast number of observations and experiences. These areas possess great academic and theoretical importance of their 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

There are many ways to tackle these problems. I believe that statistics and probabilistic inference offer one of the most useful principled means to formalize and analyze them (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. They can be complementary and of interest also.

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 is still an unsolved problem in vision.

Later, I became more interested in machine learning and believed that probabilistic inference could be used for recovering 3D articulated structure from a single 2D image (a classic ill-posed problem). However, I noticed that exact inference was intractable for this problem. For my PhD., with Stan Sclaroff, I developed an original method to solve the 3D from 2D problem using probabilistic 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. In my PhD I also became involved in more general problems and concepts in machine learning when 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, restoration, edge detection and segmentation and was also ideal for non-photorealistic rendering, texture transfer, and other graphics effects. 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 and was related to the input image by a finite set of unknown transformations. With Brendan Frey, I also developed probabilistic interpretations and extensions to spectral clustering (SC). We showed how to cast SC into statistical inference and apart from this, provided an algorithm that learns the similarity function between data points and allows us to use class dependent similarity measures by looking at full graphs as a latent random variable. This method has connections to (mixture) of manifolds learning.

At a more abstract level, we are currently developing new methods for approximate inference. Specifically we approached the problem of inference in factor graphs with high-order factors (large clique size). We have proposed a method that based on information theory, guides inference in these complicated models. Other projects include: learning to cluster by example and approximations to statistical properties of random graphs.

THE FUTURE

I aim to focus on both theoretical progress and real 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 plausible solutions.

There is a captivating connection between intractable inference and NP-complete problems (e.g., belief propagation and K-SAT), by further investigating these connections one can exploit years of theoretical findings to advance both fields. Thus, I plan to study connections to other fields. As another example, convex analysis can be employed as a means to formulate new learning concepts and also approach old, yet unsolved problems.

At the applications side, by extending my earlier work, the concept of learning image processing by example can be further generalized to digital signal processing. The 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. As a whole, it is expected to attract considerable industrial attention (including from entertainment). In vision, I will also continue research in articulated 3D reconstruction in collaboration with Stan 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 graph-based relationships among points in labeled data sets. This is tied to unsupervised and partially labeled classification. I became involved in bioinformatics, an area of major growth that can be further advanced with statistical methods. I plan to focus on three problems: the design of high-throughput experiments using information theory, the discovery of regulatory networks from micro-array data, and the improvement of accuracy in micro-array measurements. There exist excellent data availability and potential for interaction with biologists.

Vast potential for support, development, and interaction along with challenging and motivating problems make the field a remarkably exciting area of research. I pursue the development of ground-breaking algorithms and concepts in machine learning and data analysis in general, with emphasis in vision (my experience bias) and graphics, but also focused on data mining and bioinformatics. I feel that my experience of research conducting and funding in academia should 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, at least 10 of those in undergraduate and 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 studies in Venezuela, 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 gained extra experience during four semesters as the laboratory tutor, giving the lab lectures and 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 occasional tutorial 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. 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.

Since my field of expertise is multidisciplinary, I am qualified to **teach a variety of classes**, depending of the department needs, such as computer vision, machine learning, information theory, computer graphics, data mining, Bayesian probability, signal processing, algebra, discrete math, and more basic data structures or programming classes. 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.

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.