מכון ויצמן למדע

THE WEIZMANN INSTITUTE OF SCIENCE

FACULTY OF MATHEMATICS AND COMPUTER SCIENCE

Vision and Robotics Seminar

Room 1 ,Ziskind Building
on Thursday, Jan 14, 2016
at 12:15

Oren Friefeld
MIT

From representation to inference: respecting and exploiting mathematical structures in computer vision and machine learning

Abstract:

Stochastic analysis of real-world signals consists of 3 main parts: mathematical representation; probabilistic modeling; statistical inference. For it to be effective, we need mathematically-principled and practical computational tools that take into consideration not only each of these components by itself but also their interplay. This is especially true for a large class of computer-vision and machine-learning problems that involve certain mathematical structures; the latter may be a property of the data or encoded in the representation/model to ensure mathematically-desired properties and computational tractability. For concreteness, this talk will center on structures that are geometric, hierarchical, or topological.

Structures present challenges. For example, on nonlinear spaces, most statistical tools are not directly applicable, and, moreover, computations can be expensive. As another example, in mixture models, topological constraints break statistical independence. Once we overcome the difficulties, however, structures offer many benefits. For example, respecting and exploiting the structure of Riemannian manifolds and/or Lie groups yield better probabilistic models that also support consistent synthesis. The latter is crucial for the employment of analysis-by-synthesis inference methods used within, e.g., a generative Bayesian framework. Likewise, imposing a certain structure on velocity fields yields highly-expressive diffeomorphisms that are also simple and computationally tractable; particularly, this facilitates MCMC inference, traditionally viewed as too expensive in this context.

Time permitting, throughout the talk I will also briefly touch upon related applications such as statistical shape models, transfer learning on manifolds, image warping/registration, time warping, superpixels, 3D-scene analysis, nonparametric Bayesian clustering of spherical data, multi-metric learning, and new machine-learning applications of diffeomorphisms. Lastly, we also applied the (largely model-based) ideas above to propose the first learned data augmentation scheme; as it turns

out, when compared with the state-of-the-art schemes, this improves the performance of classifiers of the deep-net variety.