Learned Visual Sensors and Probabilistic Measurements

Ayan Chakrabarti, Washington University in St. Louis

As vision algorithms algorithms become increasingly successful ""in the wild"", many modern machines and devices come equipped with cameras for visual sensing. Our goal is to ensure that these cameras make optimal measurements that are most useful to these algorithms. At the same time, the versatility of visual information brings up the possibility that images captured by these on-device cameras could also be used by an adversary, that obtains these images, to recover sensitive information that we do not want disclosed. In this scenario, we want our measurements to be optimal in inhibiting, rather than promoting, inference. However, state-of-the-art inference algorithms for many tasks are based on the use of deep neural networks, and thus opaque to expert analysis on what these optimal measurement might be. In this talk, I will describe ongoing work on jointly learning camera designs and image encodings together with inference methods to attain these goals.

Moreover, as we push the limit on extracting as much scene information with as few measurements as possible, it is inevitable that there will be some uncertainty in the outputs---shape predictions, depth maps, motion estimates, etc.---of our computational camera systems. For these systems to be useful across various applications, it is useful that their outputs characterize this uncertainty instead of simply being the best guess of scene property values. To this end, I will describe our work on inference algorithms that produce probabilistic representations of the scene properties that they are tasked with measuring, so as to allow flexibility and use by downstream algorithms for a diverse range of applications.