Abstract: Computer vision systems seek to recover properties of the physical world from measurements of reflected light. To do so, they must solve ill-posed estimation problems by leveraging the statistical structure present in natural scenes. In the first part of the talk, I will introduce a new inference framework that can efficiently reason with different notions of spatial structure, at different scales and over different regions across the visual field. This allows the accurate recovery of continuous-valued maps of scene properties---depth, surface orientation, reflectance, motion, etc.---from image data. Specifically, I will describe a method that uses this framework to estimate scene depth from a single image, by training a neural network to produce dense probabilistic estimates of different elements of local geometric structure, and harmonizing these estimates to produce consistent depth maps. Beyond better inference algorithms, optimized optical sensors and measurement strategies also contribute to the success and accuracy of vision systems. In the second part of the talk, I will describe an automatic data-driven approach to the co-design of optical measurement and computational inference. Using color imaging as an example task, I will introduce a framework where the camera's measurement process is encoded as a neural-network "layer", whose learnable weights parameterize the possible measurement choices for the sensor. This layer is trained together with a network that carries out inference on the corresponding measured intensities, with the common objective of maximizing the quality of the final output. I will show that this approach is able to automatically discover a measurement strategy that, when used with the jointly learned inference network, significantly outperforms traditional sensor designs.

Bio: Ayan Chakrabarti is currently a Research Assistant Professor at the Toyota Technological Institute at Chicago. He completed his PhD. in Engineering Sciences f