Computational microscopy

George Barbastathis
Department of Mechanical Engineering, and
Singapore-MIT Alliance for Research and Technology (SMART) Centre
Massachusetts Institute of Technology

A computational imaging system consists of a physical part and an algorithmic part. The physical part manipulates photons so as to produce an intermediate representation of the object, which we call the "raw image," to be captured by the camera. Necessarily, the raw image is the intensity of a coherent or partially coherent optical field. The algorithmic part then processes the raw image to produce a final representation, which we call the "image," that is acceptable to the user. The purpose of such division of tasks, then, is to reduce the expense and burden on the optical components, transferring instead some of the design degrees of freedom to the computation, which nowadays is ubiquitously available and cheap. Indeed, this mode of imaging has existed for decades in modalities where direct imaging isn't available, such as magnetic resonance imaging, computed tomography, and particle imaging.

The basis of most computational imaging algorithms is what we will refer to as the Tikhonov-Wiener approach. Let f denote the unknown object which we wish to image, g the (intensity) raw image captured by the camera or detector, and H the "forward operator," which expresses the electromagnetic relationship between f and g. The object estimate $\hat{f}$ is then obtained as

$$\hat{f} = \mathrm{argmin}_f \left\{ ||Hf - g||^2 + \alpha\Phi(f) \right\}, \tag{1}$$

where $\Phi(f)$ is the "regularizer" expressing prior knowledge about the class of objects being imaged and α is a weighting parameter. The purpose of Φ is to penalize estimates $\hat{f}$ that do not meet certain criteria that we know *a priori* that the objects satisfy. Examples are: positivity, e.g. for fluorescent objects; and sparsity, which is almost universally applicable as long as an appropriate "sparsifying" transform can be found. Common sparsifying transforms are wavelet decompositions and, more generally, dictionaries which learn the sparse representation from examples of objects within the class of interest.

In this tutorial I will discuss in detail the historical origins of the Tikhonov-Wiener approach and some of its early successes; its application to microscopy by considering several modalities: bright and dark field, confocal, fluorescence, and quantitative phase; the intuitive explanation of regularizers and how they are chosen, sometimes but not always wisely; algorithmic implementations of the optimization problem expressed in (1); and, lastly, the recent emergence of a class of techniques that replace (1) with a machine learning algorithm that learns how to approximate Φ and, sometimes, H as well from examples.