LEARNING SHADING AND LIGHTING WITHOUT GROUND TRUTH

DAVID FORSYTH

Professor
Department of Computer Engineering
Bogazici University

ABSTRACT

Computer vision research has been revolutionized by a relatively straightforward recipe: obtain annotated data, and apply modern classification or regression techniques, as appropriate. This recipe has solved commercially valuable problems and built fame and fortune for many. But it has also stimulated an arms race -- it is now expensive and hard to use the recipe without large engiineering teams and a lot of money. Furthermore, there are many complicated but essentially simple details that need working out.

I contend that academic computer vision research, rather than working out these details, should look beyond this recipe. Do we really believe that animals have vision because the proprietor issued some early owners of an eyeball with a gold standard dataset? What do we do if we don't have, or can't get, appropriately labelled data? Two natural strategies -- fake the data, or find mathematical structure — seem particularly promising.

Physically based data faking strategies are now quite well established, and I will review strategies used for optic flow problems and for image defogging. The trick extends to other problems of a somewhat local character, and I will demonstrate how data faking leads to really strong methods for recovering bright images from very dark scenes.

Another problem of somewhat local character is intrinsic image decomposition. I will review the classical notion of intrinsic images, and various instatiations of this notion. The simplest version of the idea is a decomposition into albedo and shading. There are now good evaluation procedures for that problem. Suprisingly, early intrinsic image algorithms do quite well in modern competitions, likely because the underlying insights were right. It is quite straightforward to build fake data for this problem, and I will show how fake data can be used to achieve very good results.

DAVID FORSYTH

Professor
Department of Computer Engineering
Bogazici University

Simple faked data produces notably better results than either CGI data (hard and unreliable fakery) or real data (hard to get).

Equivariance properties -- for example, the requirement that cropping or scaling an image should crop or scale the albedo, but produce no other form of error – have been undervalued as constraints. I will demonstrate that aggressively imposing these properties on albedo estimates or brightening estimates leads to major improvements in performance.

Finally, I will use scene relighting — take a picture of a scene, and make it look as though the light were different — to illustrate the value of mathematical structure. This is a novel problem that can draw on a long mathematical tradition, which I will sketch.