Int J Comput Vis (2012) 98:123–145
DOI 10.1007/s11263-011-0501-8

Estimating the Natural Illumination Conditions from a Single Outdoor Image

Jean-François Lalonde · Alexei A. Efros · Srinivasa G. Narasimhan

Received: 24 January 2011 / Accepted: 3 October 2011 / Published online: 27 October 2011

Abstract Given a single outdoor image, we present a method for estimating the likely illumination conditions of the scene. In particular, we compute the probability distribution over the sun position and visibility. The method relies on a combination of weak cues that can be extracted from different portions of the image: the sky, the vertical surfaces, the ground, and the convex objects in the image. While no single cue can reliably estimate illumination by itself, each one can reinforce the others to yield a more robust estimate. This is combined with a data-driven prior computed over a dataset of 6 million photos. We present quantitative results on a webcam dataset with annotated sun positions, as well as quantitative and qualitative results on consumer-grade photographs downloaded from Internet. Based on the estimated illumination, we show how to realistically insert synthetic 3-D objects into the scene, and how to transfer appearance across images while keeping the illumination consistent.

Keywords Illumination estimation · Data-driven methods · Shadow detection · Scene understanding · Image synthesis

1 Introduction

The appearance of a scene is determined to a great extent by the prevailing illumination conditions. Is it sunny or overcast, morning or noon, clear or hazy? Claude Monet, a fastidious student of light, observed: "A landscape does not exist in its own right... but the surrounding atmosphere brings it to life... For me, it is only the surrounding atmosphere which gives subjects their true value." Within the Grand Vision Problem, illumination is one of the key variables that must be untangled in order to get from pixels to image understanding.

But while a lot of work has been done on modeling and using illumination in a laboratory setting, relatively little is known about it "in the wild", i.e. in a typical outdoor scene. In fact, most vision applications treat illumination more as a nuisance—something that one strives to be invariant to—rather than a source of signal. Examples include illumination adaptation in tracking and surveillance (e.g. Stauffer 1999), or contrast normalization schemes in popular object detectors (e.g. Dalal and Triggs 2005). Alas, the search for the ultimate illumination invariant might be in vain (Chen et al. 2000). Instead, we believe there is much to be gained by embracing illumination, even in the challenging, uncontrolled world of consumer photographs.

In this paper, we propose a method for estimating natural illumination (sun position and visibility) from a single outdoor image. To be sure, this is an extremely difficult task, even for humans (Cavanagh 2005). In fact, the problem is severely underconstrained in the general case—while some images might have enough information for a reasonably precise estimate, others will be completely uninformative. Therefore, we will take a probabilistic approach, estimating illumination parameters using as much information as may be available in a given image and producing the maximum likelihood solution (see Fig. 1).

So what information about illumination is available in a single image? Unfortunately, there is no simple answer. When we humans perform this task, we look at different parts of the image for clues. The appearance of the sky can

J.-F. Lalonde (✉) · A.A. Efros · S.G. Narasimhan
School of Computer Science, Carnegie Mellon University, Pittsburgh, PA, USA
e-mail: jlalonde@cs.cmu.edu

A.A. Efros
e-mail: efros@cs.cmu.edu

S.G. Narasimhan
e-mail: srinivas@cs.cmu.edu