Int J Comput Vis (2013) 103:178–189
DOI 10.1007/s11263-012-0579-7

Error-Tolerant Image Compositing

Michael W. Tao · Micah K. Johnson · Sylvain Paris

Received: 26 January 2011 / Accepted: 20 September 2012 / Published online: 13 October 2012

Abstract Gradient-domain compositing is an essential tool in computer vision and its applications, e.g., seamless cloning, panorama stitching, shadow removal, scene completion and reshuffling. While easy to implement, these gradient-domain techniques often generate bleeding artifacts where the composited image regions do not match. One option is to modify the region boundary to minimize such mismatches. However, this option may not always be sufficient or applicable, e.g., the user or algorithm may not allow the selection to be altered. We propose a new approach to gradient-domain compositing that is robust to inaccuracies and prevents color bleeding without changing the boundary location. Our approach improves standard gradient-domain compositing in two ways. First, we define the boundary gradients such that the produced gradient field is nearly integrable. Second, we control the integration process to concentrate residuals where they are less conspicuous. We show that our approach can be formulated as a standard least-squares problem that can be solved with a sparse linear system akin to the classical Poisson equation. We demonstrate results on a variety of scenes. The visual quality and runtime complexity compares favorably to other approaches.

Keywords Gradient-domain compositing · Visual masking

1 Introduction

Gradient-domain compositing is an essential technique at the core of many computer vision applications such as seamless cloning (Prez et al. 2003; Agarwala et al. 2004; Georgiev 2006; Jia et al. 2006), panorama stitching (Levin et al. 2006; Agarwala 2007; Sivic et al. 2008), inpainting (Whyte et al. 2009), shadow removal (Finlayson et al. 2009), scene completion (Hays and Efros 2007), and reshuffling (Cho et al. 2010). These methods first delineate the composited regions, then compute a target gradient field and boundary conditions from these regions, and finally solve the Poisson equation to reconstruct an image. A major issue with gradient-domain compositing is that the combined gradient field may not be integrable; that is, an image with gradients that match the target field as well as the specified boundary conditions may not exist. Existing work mitigates this issue by moving the boundary to more carefully combine the merged regions. However, when the combined images are widely different, this strategy may not be sufficient. Or, if the user has specified the boundary by hand, he or she may not want it to be altered. For instance in Fig. 1, the selection cannot be modified because the tree trunks have to abut the pyramids. Even with boundary refinement, the target gradient fields may be far from integrable, yielding color leaks and halos typical of Poisson-based methods.

In this paper, we present an algorithm for minimizing artifacts in gradient-domain image compositing. We characterize the origin of typical bleeding artifacts and analyze

Electronic supplementary material The online version of this article (doi:10.1007/s11263-012-0579-7) contains supplementary material, which is available to authorized users.

M.W. Tao (✉)
University of California, Berkeley, CA, USA
e-mail: mtao@berkeley.edu

M.K. Johnson
Massachusetts Institute of Technology, Cambridge, MA, USA
e-mail: kimo@csail.mit.edu

S. Paris
Adobe Systems, Inc., Cambridge, MA, USA
e-mail: sparis@adobe.com