

Data Hiding and the Statistics of Images

Ingemar J. Cox

Department of Computer Science
University College London
Gower Street
London WC1E 6BT
UK

Abstract. The fields of digital watermarking, steganography and steganalysis, and content forensics are closely related. In all cases, there is a class of images that is considered “natural”, i.e. images that do not contain watermarks, images that do not contain covert messages, or images that have not been tampered with. And, conversely, there is a class of images that is considered to be “unnatural”, i.e. images that contain watermarks, images that contain covert messages, or images that have been tampered with.

Thus, at the simplest level, watermarking, steganalysis and content forensics reduce to a two-class classification problem. Specifically, the recognition of natural and unnatural images. A fundamental question is whether all natural images share some common statistical properties. And are these distinct from the statistical properties of unnatural images? These questions are key to understanding the limitations of data hiding technologies with respect to false alarm rates. In this paper we review work pertaining to these questions.