

ECE544 Lecture Notes
Fall 2006

Statistical Image and Video Processing

Prof. Pierre Moulin
University of Illinois at Urbana–Champaign

©2006 by Pierre Moulin. All rights reserved.

Chapter I

Introduction

The fields of digital image and video processing have grown tremendously over the past 30 years and may be viewed as quite mature in many respects. Yet new applications have emerged in disciplines as varied as imaging, networking, and security. In these new applications as well as in the more conventional ones, the need for high performance processing and the availability of powerful computers have fueled the development of sophisticated algorithms.

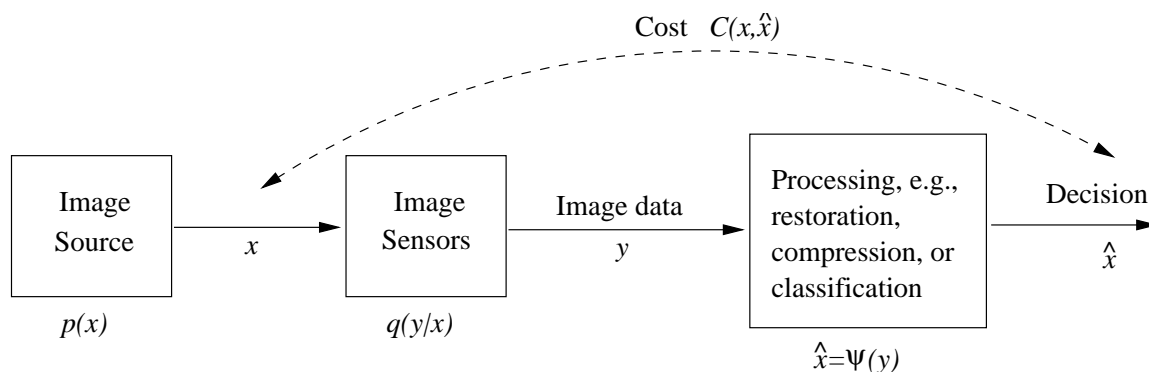
A newcomer to this field is struck by the huge variety of methods used in different applications. For instance, image restoration makes widespread use of Wiener filters, which assume stationary random models for images. On the other hand, image compression makes use of completely different techniques, as the flowchart of a JPEG compression algorithm would indicate. Despite the bewildering variety of existing image processing techniques, there are a few powerful, basic principles that can be used to guide the design of image processing algorithms. The goal of this course is to introduce these concepts and investigate their applicability to image and video processing. Specifically, we will

- develop statistical models for images and sensors;
- study optimal processing techniques for applications such as restoration and compression;
- learn how to recognize the implicit assumptions that lie behind suboptimal techniques.

Our analytical framework is as follows. Whether the application at hand involves restoration, compression or classification, uncertainty about the image x to be recovered (partially or fully) or transmitted to a receiver may be modeled by a probability distribution $p(x)$ defined over a set \mathcal{X} of images of interest. Additional uncertainty is introduced by the image measurement process which may introduce thermal noise, photon noise, limited spatial resolution, nonlinearities, saturation effects, and other degradations. The data (measurements) y take values in a space \mathcal{Y} . We use a stochastic model to describe this uncertainty, and more specifically a conditional distribution $q(y|x)$ for the data.

We design an algorithm that produces an output $\hat{x} = \psi(y)$, where ψ is viewed as a decision function. The nature of this decision is application-dependent. In image restoration, ψ is an estimator, producing a restored image \hat{x} ; in image compression, ψ is a quantization mapping, producing a bistring; in pattern recognition and image hashing applications, ψ produces a classification index. Associated with the decision function ψ is a cost $C(x, \hat{x})$ for making decision \hat{x} when the underlying image is x .

- In image restoration, $C(x, \hat{x})$ would be a distortion function based on a perceptual model or on some tractable metric such as Euclidean norm.
- In image compression, $C(x, \hat{x})$ would be the weighted sum of a distortion term and a codelength term.
- In pattern recognition, $C(x, \hat{x})$ would represent the cost of making an incorrect decision.



Once $p(x)$, $q(y|x)$ and $C(x, \hat{x})$ are specified, one may seek the decision function ψ that minimizes the average cost $\mathbb{E}C(X, \psi(Y))$ associated with ψ . This is a natural framework for developing optimal (in a perceptual sense) image restoration algorithms, optimal (in a rate-distortion sense) image compression algorithms, and optimal (in a classification performance sense) classifiers.

The primary difficulty consists in identifying realistic models for $p(x)$, $q(y|x)$ and $C(x, \hat{x})$, and in making the optimization problems tractable (this introduces practical constraints on the possible models p , q , C). Modeling of $p(x)$ is a difficult problem. The physics of the sensors determine $q(y|x)$. The choice of a suitable $C(x, \hat{x})$ is application-dependent and often requires some knowledge of the Human Visual System.

When the solution above is hard to find, an appropriate substitute can often be identified, e.g., ψ that minimizes an approximate cost function.

A graduate-level course in random processes is needed to formulate, understand and analyze the statistical models used in this course. Some familiarity with Image Processing is also necessary to develop the necessary intuition about the design of statistical models.