
Abstract

Smoothing (diffusion) is a fundamental task in low-level vision and image processing. In the context of natural images, where edges (sharp discontinuities) play an important psychovisual role, the smoothing process needs to respect two contrasting requirements at the same time – enhanced diffusion in homogeneous regions and suppressed diffusion around edges. Clearly, it is impossible to meet these requirements using a fixed Gaussian (heat) kernel. The seminal work of Perona and Malik showed that this can be achieved using a non-linear heat equation, where an image-dependent kernel is used to regulate the diffusion at a point. Following this work, several transform or filtering-based approaches have been proposed for non-linear smoothing. The focus of this thesis is on two popular smoothers, the bilateral and the non-local means filters, and their application to high-dimensional data. By high-dimensional data, we mean color and hyperspectral images, videos, patch-based representations, flow-fields, etc. where the dimension of the spatio-range space is large.

Brute-force computation of high-dimensional filters in real-time is a challenging task. To address this issue, researchers have come up with approximation algorithms which can accelerate the computation without significantly compromising the visual quality. Unfortunately, most of these algorithms are tailored to work with gray images and it is difficult to extend them to color images. The situation is more challenging for hyperspectral images where the range dimension is much larger than color images.

The thesis has two important contributions. The first of these is a fast algorithm for bilateral filtering of gray images. In particular, we show that by approximating the kernel using a truncated Fourier series, we can express the non-linear filtering using fast convolutions. However, since the number of terms in the approximation grows exponentially with the dimension, it is difficult to scale the algorithm for high-dimensional filtering. We show that this can be overcome by approximating both the underlying data and the kernel. The algorithm emerging from the approximation essentially involves clustering and fast convolutions. Importantly, the algorithm is conceptually simple, easy to implement, and comes with a guarantee on the approximation error which is not enjoyed by existing algorithms. We present filtering results on gray, color and hyperspectral images, patch-based representations, videos, and flow fields to demonstrate the effectiveness of the proposed algorithms.