Transfer Denoising with Hierarchical Dirichlet Process Hidden Markov Trees

Jyri J. Kivinen KIVINEN@CS.BERKELEY.EDU
Erik B. Sudderth SUDDERTH@CS.BERKELEY.EDU
Michael I. Jordan JORDAN@CS.BERKELEY.EDU

We develop hierarchical, nonparametric Bayesian models for wavelet representations of natural images. Individual wavelet coefficients at each scale are marginally distributed as Dirichlet process (DP) mixtures of Gaussians. The hidden assignments of observations to clusters at neighboring nodes are then linked by a multiscale, tree–structured graphical model. When learning, the transition distributions associated with different parent states are coupled via a hierarchical Dirichlet process (HDP) (Teh et al., 2006). The multiscale stochastic process defined by this hierarchical Dirichlet process hidden Markov tree (HDP-HMT) automatically adapts to the varying complexity of different datasets, and captures global, highly non-Gaussian statistical properties of natural images.

The HDP-HMT framework was initially proposed by Kivinen et al. (2007a), extending prior work on hidden Markov trees, local Gaussian scale mixtures, and the hierarchical Dirichlet process hidden Markov model (HDP-HMM) (Teh et al., 2006). Using a collapsed Gibbs sampler for learning HDP-HMT models from images, an effective empirical Bayesian image denoising algorithm was developed. Subsequently, truncated representations of the HDP were used to develop more efficient blocked sampling algorithms, allowing learning from large datasets (Kivinen et al., 2007b). These truncations also provide a mechanism for balancing computational efficiency and representational accuracy, while maintaining a nonparametric model.

In this abstract, we extend the HDP-HMT framework to allow statistics from a large database of natural images to be *transferred* to the restoration of noisy images of novel scenes. Figure 1 shows a graphical model for a single image which has been contaminated by additive, zero–mean Gaussian noise of known variance. It augments the model presented in Kivinen et al. (2007a) by a set of hidden, clean wavelet coefficients, enabling direct separation of the statistics of noise from that of undistorted images. Our denoising algorithm extends the blocked Gibbs sampler of Kivinen et al. (2007b) to also resample noisy coefficients.

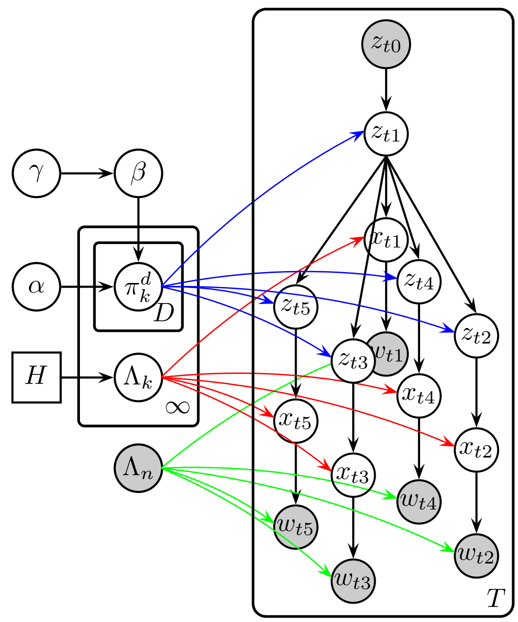

Figure 1. Two levels of an HDP-HMT in which hidden discrete states z_{ti} generate noise–free features x_{ti}. Observations w_{ti} are corrupted by additive Gaussian noise with covariance Λ_n. States at neighboring locations and scales are linked by direction dependent transition distributions π_k^d. A global measure β couples these transitions when learning, encouraging reuse of hidden states.

Most existing image denoising algorithms estimate unknown parameters directly from the noisy image at hand. While effective in some cases, at high noise levels there can be insufficient information, and flexible models may lead to significantly distorted reconstructions. Using an empirical Bayesian denoising algorithm, our results at low and moderate noise levels are comparable to BLS-GSM, a leading wavelet-based denoising method. However, at higher noise levels, increasing high-frequency artifacts start to reduce restoration quality (see figures 2&3).

By learning the statistics of a set of 100 clean natural images from the Berkeley segmentation dataset,

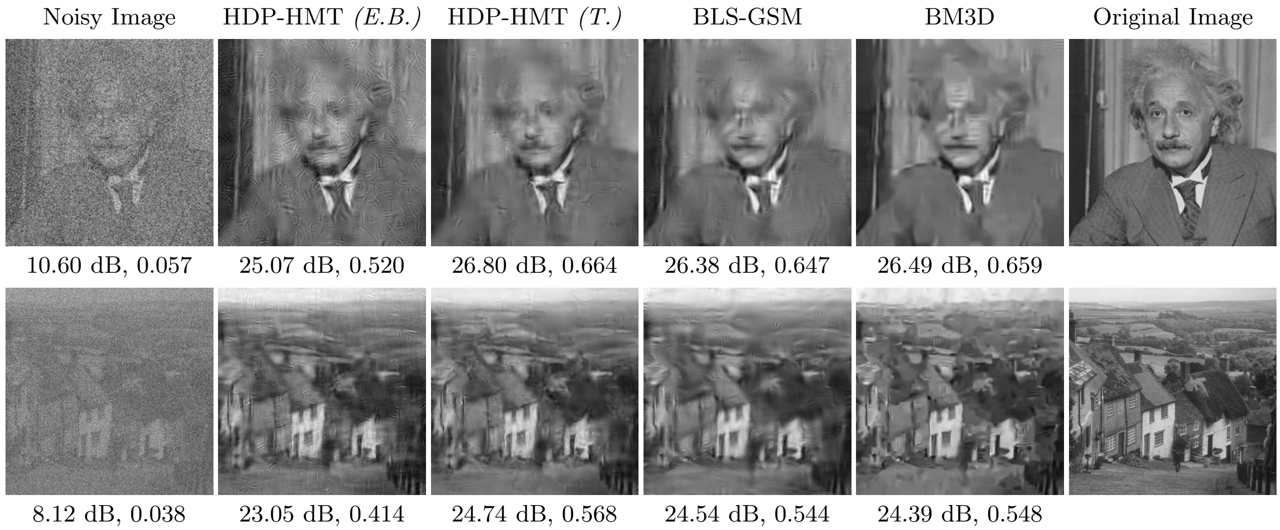

Figure 2. Denoising results for *einstein* and *hill* images contaminated by additive white Gaussian noise of standard deviation $\sigma = 75$ (top) and $\sigma = 100$ (bottom) pixels, respectively. We also report PSNR and perceptual mean structural similarity index values for the proposed HDP-HMT models (empirical Bayesian (E.B.) and using statistics transfer (T.)), a leading wavelet-based BLS-GSM (Portilla et al., 2003), and block-based method BM3D (Dabov et al., 2007).

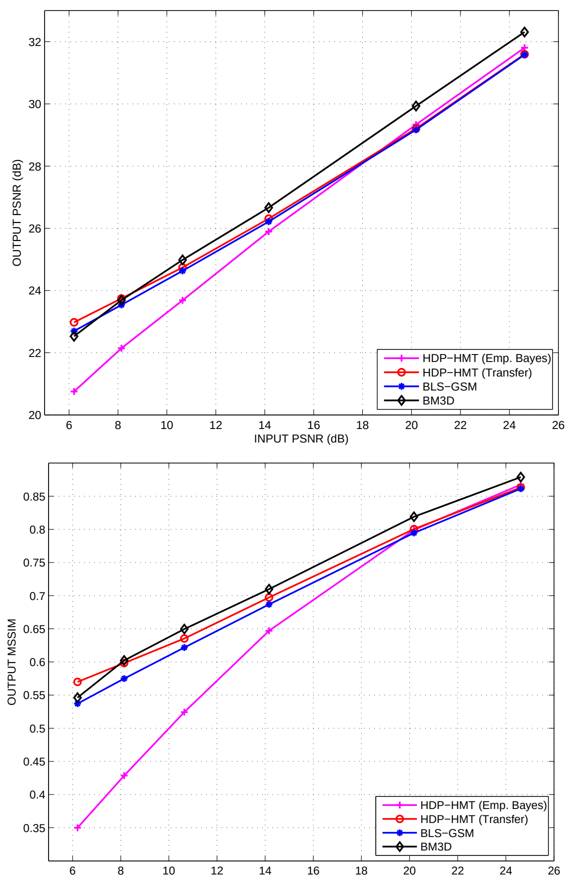

Figure 3. Average peak signal-to-noise ratio (PSNR) and mean structural similarity index (MSSIM) values as a function input PSNR, computed over the denoising results of an ensemble of standard denoising images (*cameraman, einstein, house, mandrill, peppers).*

the HDP-HMT learns that images typically contain many smooth or homogeneously textured regions, separated by sharp edges. The denoising algorithm transfers this prior knowledge by reusing multiscale hidden state patterns, resulting in better reconstruction of distorted textures at higher noise levels, especially with respect to the perceptual MSSIM criterion.

In these higher noise regimes, transfer denoising with the HDP-HMT also surpasses the performance of the BLS-GSM. At extreme noise levels, the results are comparable to or better than those of BM3D, a state-of-the-art algorithm which averages similar blocks of pixels. We expect that transfer of natural image statistics will prove useful for correcting other forms of image distortion, such as significant motion blur.

References

Dabov, K., Foi, A., Katkovnik, V., & Egiazarian, K. (2007). Image denoising by sparse 3d transform-domain collaborative filtering. *IEEE Trans. Image Proc.*, *16*.

Kivinen, J., Sudderth, E., & Jordan, M. (2007a). Image denoising with nonparametric hidden Markov trees. *ICIP*.

Kivinen, J., Sudderth, E., & Jordan, M. (2007b). Learning multiscale representations of natural scenes using Dirichlet processes. *ICCV*.

Portilla, J., Strela, V., Wainwright, M., & Simoncelli, E. (2003). Image denoising using scale mixtures of Gaussians in the wavelet domain. *IEEE Trans. Image Proc.*, *12*.

Teh, Y. W., Jordan, M. I., Beal, M. J., & Blei, D. M. (2006). Hierarchical Dirichlet processes. *JASA*, *101*.