



# Video Epitomes

VINCENT CHEUNG AND BRENDAN J. FREY

*Electrical and Computer Engineering, University of Toronto, Toronto, ON, M5S 3G4, Canada*

vincent@psi.toronto.edu

frey@psi.toronto.edu

NEBOJSA JOJIC

*Machine Learning and Applied Statistics, Microsoft Research, Redmond, WA 98052, USA*

jojic@microsoft.com

*Received March 8, 2006; Accepted November 8, 2006*

**Abstract.** Recently, “epitomes” were introduced as patch-based probability models that are learned by compiling together a large number of examples of patches from input images. In this paper, we describe how epitomes can be used to model video data and we describe significant computational speedups that can be incorporated into the epitome inference and learning algorithm. In the case of videos, epitomes are estimated so as to model most of the small space-time cubes from the input data. Then, the epitome can be used for various modeling and reconstruction tasks, of which we show results for video super-resolution, video interpolation, and object removal. Besides computational efficiency, an interesting advantage of the epitome as a representation is that it can be reliably estimated even from videos with large amounts of missing data. We illustrate this ability on the task of reconstructing the dropped frames in video broadcast using only the degraded video and also in denoising a severely corrupted video.

**Keywords:** epitome, video summarization, em algorithm, variational technique, super-resolution, inpainting, object removal, image restoration, missing data

## 1. Introduction

The technique of using small image patches to account for high-order statistics in image and video data continues to grow in popularity in the vision community. One of the first uses of patches was to find corresponding points in neighboring video frames for computing optical flow. Instead of simply matching patches, Jepson and Black introduced probabilistic patch models that account for outliers (Jepson and Black, 1993). Soon after that, Wang and Adelson showed that patches could be used for efficient video compression (Wang et al., 1994). In related work, “textons” use image patches within a structured mathematical framework, to account for texture using a patch-based representation (Zhu et al., 2002). In these cases, patches were used primarily for analyzing image data.

More recently, patches have been used successfully for synthesizing images and videos. Patches from one part of an image have been stitched together to synthesize new

images with similar texture, or to in-paint texture into an interior region (Efros and Freeman, 2001; Criminisi et al., 2003). This approach has also been used to fill in missing or occluded regions of video data (Wexler et al., 2004). Image denoising has been performed by averaging patches of similar intensity (Buades et al., 2005). Libraries of patches derived from high-resolution images have been used for super-resolution, both in static images and video sequences (Freeman et al., 2002; Bishop et al., 2003). Patch-based image models have been used for the purpose of “texture transfer,” also known as “unsupervised image analogies” (Rosales et al., 2003).

To jointly analyze and synthesize data, patch-based probability models are introduced in Jojic et al. (2003). These models, called “epitomes,” compile patches drawn from input images into a condensed image model that represents many of the high-order statistics in the input image. An advantage of learning a patch-based model from the input data is that the model can be used for a variety of inference tasks, such as image/texture classification,

in-painting, super-resolution, texture transfer, and image segmentation (Jojic et al., 2003). Epitomes have found uses in other application areas, including speech, motion capture data, and molecular biology.

In this paper, we address two problems: (1) How 2D image epitomes can be extended through time to form 3D space-time epitomes; (2) How epitomes can be compiled and applied in a computationally efficient manner and in particular, in a way that is significantly more efficient than using a library of patches. After describing a general framework for learning video epitomes, we demonstrate its use in several disparate applications.

## 2. Video Epitomes

Figure 1 outlines the procedure used to learn a video epitome. Viewing the input video as a 3D space-time volume, a large number of 3D training patches are drawn from the video. The learning algorithm is used to compile these patches into an “epitome”—a video that is smaller in both space and time, but contains many of the spatial and temporal patterns in the original input video. We derive the epitome learning algorithm by specifying a generative model, which explains how the input video can be generated from the epitome (in the opposite direction shown in the figure). The advantage of specifying a generative model is that the model is more adaptive than a non-generative technique and it can be used as a sub-component in other systems. Here, we introduce a 3D model of video similar to the 2D epitome model described in Jojic et al. (2003).

We treat a video sequence as a 3D array  $\mathbf{v}_{x,y,t}$  of real-valued pixel measurements (R, G, and B color channels in this paper), with  $x \in \{1, \dots, X_v\}$ ,  $y \in \{1, \dots, Y_v\}$ ,  $t \in \{1, \dots, T_v\}$ . The epitome  $e$  models the video using a set of probability distributions arranged on a grid of size  $X_e \times Y_e \times T_e$ , considerably smaller than the video, i.e.,  $X_e Y_e T_e < X_v Y_v T_v$ . We view the epitome  $e_{x,y,t}$  as a 3D array of probability distributions. A particular pixel value  $\mathbf{v}$

can be evaluated under any of the probability distributions in  $e$ . For example, for the epitome coordinates  $x_e, y_e, t_e$ , the probability density at the pixel value stored in the entry  $x_v, y_v, t_v$  of the video is  $e_{x_e, y_e, t_e}(\mathbf{v}_{x_v, y_v, t_v})$ . Since the pixel measurements are continuous in nature, it is necessary to parameterize each of the epitome distributions. In our experiments, we use a single parametric form, a three-dimensional Gaussian distribution parameterized by a different mean and diagonal covariance matrix for each entry,

$$e_{x,y,t}(\cdot) = \mathcal{N}(\cdot; \mu_{x,y,t}, \phi_{x,y,t}), \quad (1)$$

where  $\mu_{x,y,t}$  is the mean and  $\phi_{x,y,t}$  is the covariance matrix (e.g. for RGB values). The diagonal covariance matrix decouples color channel computations. Because of this, we will treat the measurements  $\mathbf{v}_{x,y,t}$  as scalar in the following derivations, which the reader can use to derive the full color model.

The epitome models the video by modeling 3D patches sampled from the video. These patches can have any shape, but to keep notation simple, we will assume each patch has linear, axis-aligned boundaries, and we think of each patch as a “cube”. Each patch is defined in terms of the ordered set of pixel coordinates in the patch,  $\mathcal{S}$ . For example, a  $10 \times 10 \times 5$  video patch starting at position 8, 9 in frame 7 of the video is described by the set  $\mathcal{S} = \{8, \dots, 17\} \times \{9, \dots, 18\} \times \{7, \dots, 11\}$ . We assume the coordinates in  $\mathcal{S}$  are ordered, so that  $\mathcal{S}(k)$  refers to the  $k$ th coordinate in  $\mathcal{S}$ , e.g.,  $\mathcal{S}(1) = (8, 9, 7)$  in the above example.

While  $\mathbf{v}$  denotes the observed pixel values at all coordinates in the video,  $\mathbf{v}_{\mathcal{S}}$  denotes the observed pixel values in a small video cube at coordinates  $\mathcal{S}$ , and  $\mathbf{c}_{\mathcal{S}}$  denotes the pixel values in the same cube, as predicted by the epitome. As described below, the goal of the learning algorithm is to make the predicted video cubes similar to the observed video cubes, i.e.,  $\mathbf{c}_{\mathcal{S}} \approx \mathbf{v}_{\mathcal{S}}$ . The cube  $\mathbf{c}_{\mathcal{S}}$  is modeled using a set of distributions in the epitome. We use  $e_{\mathcal{T}}$  to denote the set of distributions from the epitome  $e$  at coordinates  $\mathcal{T}$ .

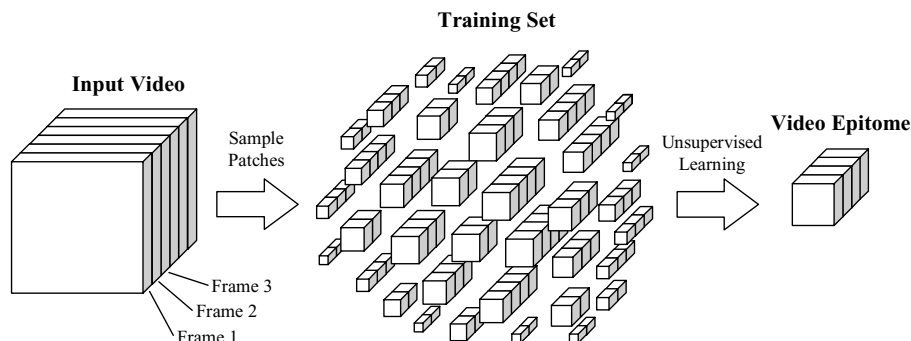


Figure 1. Learning the epitome of a video.

Assuming that the cube corresponding to  $\mathcal{T}$  is equal in size to the cube corresponding to  $\mathcal{S}$ , we can define a one-to-one coordinate correspondence so that the set ordering is preserved; then, the notation  $\mathcal{T}_{\mathcal{S}}$  is used. The probability density evaluated using distributions at coordinates  $\mathcal{T}$  and predicted values  $\mathbf{c}_{\mathcal{S}}$  is then  $e_{\mathcal{T}}(\mathbf{c}_{\mathcal{S}})$ :

$$p(\mathbf{c}_{\mathcal{S}}|\mathcal{T}_{\mathcal{S}}) = e_{\mathcal{T}_{\mathcal{S}}}(\mathbf{c}_{\mathcal{S}}) = \prod_{k=1}^{|\mathcal{T}|} e_{\mathcal{T}_{\mathcal{S}}(k)}(\mathbf{c}_{\mathcal{S}(k)}). \quad (2)$$

The above equation assumes independence of pixels given the responsible set of epitome distributions, with the idea of capturing correlations simply through mapping sets of measurements to the overlapping sets of distributions.

The above equation describes the probability model for an individual cube. To obtain a probability model for the entire input video,  $\mathbf{v}$ , we need to address the issue of how to account for overlapping cubes. If two sets of coordinates  $\mathcal{S}$  and  $\mathcal{S}'$  overlap, then the corresponding cubes  $\mathbf{c}_{\mathcal{S}}$  and  $\mathbf{c}_{\mathcal{S}'}$  are used to predict overlapping cubes in the input video  $\mathbf{v}$ , and so they should make similar predictions for overlapping pixels. The most obvious approach to account for overlap is to include a constraint that the predictions are identical in regions of overlap. However, this approach requires the introduction of a partition function, making learning significantly more difficult. We take a different, novel, approach (Jojic et al., 2003), where we treat the cubes  $\mathbf{c}_{\mathcal{S}}$  as independent even if they share coordinates, but enforce the agreement in the overlapping regions during inference, as described in the next section.

We now define a generative model of video sequences that is suitable for all applications described in the experimental section. The first step in the generative process consists of generating a predicted cube  $\mathbf{c}_{\mathcal{S}}$  for every possible coordinate set  $\mathcal{S}$  in the input video. This is accomplished by first randomly choosing a patch  $\mathcal{T}_{\mathcal{S}}$  from the epitome using a uniform distribution, and then generating  $\mathbf{c}_{\mathcal{S}}$  using the distribution  $e_{\mathcal{T}_{\mathcal{S}}}(\mathbf{c}_{\mathcal{S}})$  defined by (2). Then, for each pixel coordinate  $x, y, t$  in the video, all overlapping cubes  $\{\mathcal{S} : (x, y, t) \in \mathcal{S}\}$  are combined to make a single prediction for the observed pixel  $\mathbf{v}_{x,y,t}$  at that coordinate. This generative process is illustrated in Fig. 2.

During free energy minimization, we will force the predictions to agree, so the exact form of the combination function is not important. Here, we assume that to generate a consistent video pixel  $\mathbf{v}_{x,y,t}$ , the contributions from all overlapping video cubes  $\{\mathcal{S} : (x, y, t) \in \mathcal{S}\}$  are averaged, and Gaussian noise with variance  $\sigma_{x,y,t}^2$  is added to all three channels:

$$\begin{aligned} & p(\mathbf{v}_{x,y,t}|\{\mathbf{c}_{\mathcal{S}} : (x, y, t) \in \mathcal{S}\}) \\ &= \mathcal{N}\left(\mathbf{v}_{x,y,t}; \frac{\sum_{\mathcal{S}} \sum_k [\mathcal{S}(k) = (x, y, t)] \mathbf{c}_{\mathcal{S},k}}{\sum_{\mathcal{S}} \sum_k [\mathcal{S}(k) = (x, y, t)]}, \sigma_{x,y,t}^2\right) \\ & p(\mathbf{v}|\{\mathbf{c}_{\mathcal{S}}\}) = \prod_{x,y,t} p(\mathbf{v}_{x,y,t}|\{\mathbf{c}_{\mathcal{S}} : (x, y, t) \in \mathcal{S}\}), \quad (3) \end{aligned}$$




Figure 2. Generative process for a single video pixel,  $\mathbf{v}_{x,y,t}$ , which is represented by a dot in the input video. The overlapping video cubes in the input video containing the pixel  $x, y, t$  are  $\{\mathcal{S}_{i_1}, \dots, \mathcal{S}_{i_n}\}$ , two of which are outlined in gray in the input video. For each of these video cubes, there is a predicted cube. These predicted cubes,  $\{\mathbf{c}_{\mathcal{S}_{i_1}}, \dots, \mathbf{c}_{\mathcal{S}_{i_n}}\}$ , are generated from the epitome by randomly selecting locations in the epitome  $\{\mathcal{T}_{\mathcal{S}_{i_1}}, \dots, \mathcal{T}_{\mathcal{S}_{i_n}}\}$ , represented by dots in the video epitome, and then generating the pixel values for the predicted cubes under the distribution given by the epitome for these locations,  $e_{\mathcal{T}_{\mathcal{S}_{i_1}}}(\mathbf{c}_{\mathcal{S}_{i_1}}), \dots, e_{\mathcal{T}_{\mathcal{S}_{i_n}}}(\mathbf{c}_{\mathcal{S}_{i_n}})$ . Each of these cubes  $\mathbf{c}_{\mathcal{S}_i}$  make predictions for the pixel  $\mathbf{v}_{x,y,t}$  and these predictions are combined to generate a single pixel value.

where  $[ ]$  is Iverson's indicator function, i.e.,  $[true] = 1$ ,  $[false] = 0$ . We use the notation  $\mathbf{c}_{\mathcal{S},k}$  for the  $k$ -th pixel in  $\mathbf{c}_{\mathcal{S}}$  to emphasize that the video cubes  $\mathbf{c}_{\mathcal{S}}$  are treated as independent, and so a pixel in the video cube  $\mathbf{c}_{\mathcal{S}}$  is not uniquely defined by the global coordinate  $\mathcal{S}(k)$ , and is instead, potentially different from the pixels in other cubes that overlap  $\mathcal{S}(k)$ . However, as described by the above equations, for each coordinate  $(x, y, t)$ , a single final pixel  $\mathbf{v}_{x,y,t}$  is generated by adding noise to the average of all video cubes overlapping with coordinate  $(x, y, t)$ .

Usually, a subset of the input video patches provides accurate sufficient statistics for learning an epitome. In practice, we randomly sample the set  $\{\mathcal{S}\}$  of video cubes to be used for learning (e.g. from the set of all video cubes of a certain size) so that with high probability, or absolute certainty, we would expect to see every pixel from the input video in at least one video cube. The joint distribution over all variables can be written  $p(\mathbf{v}, \{\mathbf{c}_{\mathcal{S}}, \mathcal{T}_{\mathcal{S}}\}) = p(\mathbf{v}|\{\mathbf{c}_{\mathcal{S}}\}) \prod_{\mathcal{S}} p(\mathbf{c}_{\mathcal{S}}|\mathcal{T}_{\mathcal{S}}) p(\mathcal{T}_{\mathcal{S}})$ . We often assume  $p(\mathcal{T}_{\mathcal{S}})$  is uniform for the sake of simplicity.

### 3. Learning Video Epitomes

Learning under this model using the expectation-maximization algorithm is not possible because the exact

posterior is intractable. To overcome this problem, a variational technique is used to approximate the posterior. Learning can be viewed as minimizing a free energy cost function (Frey and Jojic, 2005):

$$F = \sum_S \sum_{\mathcal{T}_S} \int_{\mathbf{c}_S} q(\{\mathcal{T}_S, \mathbf{c}_S\}) \log \frac{q(\{\mathcal{T}_S, \mathbf{c}_S\})}{p(\mathbf{v}, \{\mathbf{c}_S, \mathcal{T}_S\})}, \quad (4)$$

where  $q(\{\mathcal{T}_S, \mathbf{c}_S\})$  is the auxiliary joint probability distribution over the set of epitome patches  $\{\mathcal{T}_S\}$  and the set of video cubes  $\{\mathbf{c}_S\}$  for all coordinate patches  $\{\mathcal{S}\}$ . It turns out that the free energy (Neal and Hinton, 1998) bounds the log-likelihood of the input video,  $F \geq -\log p(\mathbf{v})$ , where  $p(\mathbf{v}) = \sum_{\{\mathcal{T}_S\}} \int_{\{\mathbf{c}_S\}} p(\mathbf{v}, \{\mathbf{c}_S, \mathcal{T}_S\})$ , the proof of which can be found in the Appendix. So, by optimizing  $F$ , we can maximize  $p(\mathbf{v})$ . In fact, the closer  $q$  is to the true posterior  $p(\{\mathcal{T}_S, \mathbf{c}_S\}|\mathbf{v})$ , the tighter the bound, so  $q$  is an approximation to the posterior distribution. By choosing an appropriate form for  $q$ , we can achieve two goals: Obtain an efficient inference and learning algorithm by decoupling variables. Constrain inference and learning to the space where overlapping cubes agree, as proposed in Jojic et al. (2003). We choose the  $q$ -distribution as follows:

$$q(\{\mathcal{T}_S, \mathbf{c}_S\}) = \prod_S q(\mathcal{T}_S)q(\mathbf{c}_S). \quad (5)$$

$q(\mathcal{T}_S)$  is a discrete distribution over all possible patch locations in the epitome shaped as  $\mathcal{S}$ . To enforce the overlap constraint,  $q(\mathbf{c}_S)$  is a product of Dirac functions:

$$q(\mathbf{c}_S) = \prod_k \delta(\mathbf{c}_{S,k} - \nu_{S(k)}). \quad (6)$$

As before,  $\mathbf{c}_{S,k}$  denotes the pixel  $\mathcal{S}(k)$  in cube  $\mathbf{c}_S$ .  $\nu_{S(k)}$  denotes the variational parameter for the pixel  $\mathcal{S}(k) = (x, y, t)$ , which is *shared* by all patches  $\mathcal{S}$  that contain the coordinate  $(x, y, t)$ , thus constraining the patches to agree in overlapping pixels.

The free energy obtained using the above  $q$ -distribution leads to a tractable iterative learning algorithm. Setting to zero the derivatives of  $F$  w.r.t. the posterior cube colors  $\nu_{x,y,t}$ , the details of which are in the Appendix, we obtain the following update rule:

$$\nu_{x,y,t} \leftarrow \frac{\frac{\mathbf{v}_{x,y,t}}{\sigma_{x,y,t}^2} + \sum_{S,k:\mathcal{S}(k)=(x,y,t)} \sum_{\mathcal{T}_S} q(\mathcal{T}_S) \frac{\mu_{\mathcal{T}_S(k)}}{\phi_{\mathcal{T}_S(k)}}}{\frac{1}{\sigma_{x,y,t}^2} + \sum_{S,k:\mathcal{S}(k)=(x,y,t)} \sum_{\mathcal{T}_S} q(\mathcal{T}_S) \frac{1}{\phi_{\mathcal{T}_S(k)}}}. \quad (7)$$

This update sets the ‘‘hidden video’’  $\nu$  to a weighted combination of the input video and the top-down prediction for the video, as given by the current epitome. The weights are the inverse noise variances for the video and the epitome. Setting to zero the derivatives of  $F$  w.r.t.

the posterior epitome responsibilities  $q(\mathcal{T}_S)$ , we obtain

$$q(\mathcal{T}_S) \leftarrow \frac{p(\mathcal{T}_S)e_{\mathcal{T}_S}(\nu_S)}{\sum_{\mathcal{T}} p(\mathcal{T})e_{\mathcal{T}}(\nu_S)}. \quad (8)$$

This update is similar to computing the responsibilities of components in a mixture of Gaussians. For each cube  $\nu_S$ , the distribution over its position in the epitome is proportional to the probability it was generated from each position. To perform inference on video  $\mathbf{v}$  when the epitome is given, these two updates can be iterated until convergence.

If the epitome is to be learned from the video  $\mathbf{v}$ , then the two equations above are iterated in combination with the following two updates, obtained by setting to zero the derivatives of  $F$  w.r.t. the epitome parameters  $\mu_{x,y,t}$  and  $\phi_{x,y,t}$ :

$$\mu_{x_e, y_e, t_e} \leftarrow \frac{\sum_{\mathcal{T}, k:\mathcal{T}(k)=(x_e, y_e, t_e)} \sum_S q(\mathcal{T}_S = \mathcal{T}) \nu_{S(k)}}{\sum_{\mathcal{T}, k:\mathcal{T}(k)=(x_e, y_e, t_e)} \sum_S q(\mathcal{T}_S = \mathcal{T})} \quad (9)$$

$$\phi_{x_e, y_e, t_e} \leftarrow \frac{\sum_{\mathcal{T}, k:\mathcal{T}(k)=(x_e, y_e, t_e)} \sum_S q(\mathcal{T}_S = \mathcal{T}) (\nu_{S(k)} - \mu_{x_e, y_e, t_e})^2}{\sum_{\mathcal{T}, k:\mathcal{T}(k)=(x_e, y_e, t_e)} \sum_S q(\mathcal{T}_S = \mathcal{T})}. \quad (10)$$

The first update sets the mean pixel value in the epitome to the average value of all pixel values from video cubes  $\nu_S$ , weighted by the probability that the cube is aligned with the pixel in the epitome,  $q(\mathcal{T}_S = \mathcal{T})$ . The second update is similar, except it accounts for the variance, not the mean value. All inference algorithms in this paper are based on these four equations and the derivations for these equations are shown in the Appendix. In the next section, we address the issue of the relative dimensions of the epitome, the subsequent section explains how the above equations define different video processing applications, and in Section 6 we derive an efficient algorithm that implements these equations.

#### 4. Trading Off Space and Time

The video epitome can compress a video both spatially and temporally. The size of the epitome acts as a knob that can be tuned to adjust the amount of compression in both space and time. The balance between space and time has a profound effect on the resulting video epitome and how it models the video.

Figure 3(a) shows a few frames from a sample video where a toy car drives around a mat. In one extreme, the video epitome can emphasize spatial compression, as in Fig. 3(b). In this circumstance, individual motion patterns are isolated and several frames in the epitome are dedicated to each motion. Note that the epitome is taken over




Figure 3. Varying the dimensions of the epitome. (a) A few frames from a  $208 \times 288 \times 115$  video sequence of a toy car driving around a mat. (b) Several frames from its  $60 \times 60 \times 80$  video epitome, demonstrating an epitome that is spatially small. (c) A  $208 \times 288 \times 5$  video epitome, which is on the other extreme and significantly compresses time. (d) A  $208 \times 288 \times 5$  video epitome initialized in a more structured manner to encourage spatial features to align to those in the original sequence. All these epitomes contain approximately the same number of pixels and are each more than 20 times smaller than the original video. A variety of patch sizes were used in learning the epitome, varying from two to four along the time dimension and  $10 \times 20$  to  $15 \times 25$  and  $15 \times 10$  to  $20 \times 12$  spatially.

a torus, that is, the epitome circularly wraps at the edges in order to maximize the use of all the available pixels. Conversely, Fig. 3(c) shows a video epitome that greatly compresses the time dimension of the video. With just a few frames to work with, the video epitome models multiple motion patterns simultaneously within its frames.

Usually the mean of the epitome is initialized by sampling each pixel in the epitome independently from a gaussian distribution with the same mean and variance as that of the original video sequence. An inherent structure can be imposed upon the epitome by using a more structured initialization. For example, the epitome in Fig. 3(d) has the same dimensions and was learnt in the same manner as the epitome in Fig. 3(c), but the spatial structures are aligned with those in the original sequence. This effect was accomplished by initializing each pixel in the epitome to be the median value of the pixels across all the frames in that exact spatial location, plus a small amount of noise. The posterior epitome responsibility,  $q(\mathcal{I}_S)$ , from (8), then generally favors coordinates such that the spatial locations match the original video sequence because of the high agreement in pixel values.

These video epitomes contain approximately the same total number of pixels, but have much different appearances. However, in all of these epitomes, the essential structural and motion components are maintained. While the video epitome can itself be useful for visual purposes, its true power arises when used within a larger model for performing various video processing applications as

described in the next section. When used for a data processing application, the dimensions of the epitome should achieve a balance between space and time so that both large spatial and temporal features can be modeled.

## 5. Epitomic Video Processing

First, we note that a very high variance parameter  $\sigma_{x,y,t}^2$  effectively severs all the video cubes  $\mathbf{c}_S$  from the observation  $\mathbf{v}_{x,y,t}$ . According to the generative model Eq. (3), an excessive level of noise  $\sigma_{x,y,t}^2$  will make the generated value  $\mathbf{v}_{x,y,t}$  dominated by Gaussian noise and thus independent from the epitome generated video cubes. Because of this, the inferred color  $v_{x,y,t}$  shared for the coordinate  $(x, y, t)$  in all video cubes is dominated by the epitome prediction, as seen in (7). On the other hand, when  $\sigma_{x,y,t}^2$  is very low, (7) is dominated by the observation  $\mathbf{v}_{x,y,t}$ . This general observation leads to several video reconstruction applications of the inference algorithm described in the previous section, including filling-in missing data, obstruction removal, video interpolation, denoising, and super-resolution.

In a large area of high variance  $\sigma^2$ , iterating (7) and (8) will fill the inferred values  $\nu$  with the epitome generated cubes that tend to agree in the overlapping regions. To better understand this, consider a randomly initialized set of  $\nu$  values. The inference step (8) for the cube  $\mathbf{c}_S$  takes the current guess at  $\nu_S$  and evaluates how likely each

epitome cube  $e_T$  is to generate  $v_S$ . Since the posterior epitome responsibilities share the  $v$  values, the (probabilistically) chosen epitome cubes tend to have a higher level of agreement in the overlaps than if a set of random epitome cubes were selected as responsible for the video cubes. Now, applying (7) will replace the initialized values  $v$  with the average votes from the epitome cubes deemed likely to have been generated according to  $q(\mathcal{T}_S)$ . Since the likely video cubes have a moderate level of agreement, the generated  $v$  texture will now be more consistent with the epitome texture. Iterating these two steps will lead to the solution where  $\mathcal{T}_S$  are chosen so as to “quilt” a random, but consistent texture from the epitome. Note also, that in these steps, each new estimate for a  $v_{x,y,t}$  is a weighted average of *all* epitome means  $\mu_{x_e, y_e, t_e}$  weighted by both the inverse epitome variances  $\phi_{x_e, y_e, t_e}$  and the probabilities of all possible cube mappings that lead to mapping epitome entry  $(x_e, y_e, t_e)$  to video entry  $(x, y, t)$ . After several iterations, however, the  $q(\mathcal{T}_S)$  become fairly peaked and consistent.<sup>1</sup>

If the area of high variance  $\sigma^2$  is surrounded by the area of low variance, then the quilting steps described above will be automatically conditioned on satisfying the agreement of the  $v$  values in the areas of low variance. Thus, we can perform video repair operations by setting  $\sigma^2$  values to be high in the areas of video that are either missing, or considered bad, or that we simply want to replace with the texture from the epitome.

An interesting additional property of the algorithm in the previous section is that the epitome learning can be performed *jointly* with the repair operations, even when some of the input video pixels are known to be unreliable. If the unreliable (or unwanted) pixels are given high noise parameters  $\sigma^2$ , then iterating all four of the Eqs. (7)–(10) will construct the epitome only from the reliable pixels. (Only the reliable pixels propagate into  $v$  values, which in turn define the epitome statistics in each step (10)).

Next, we define several video reconstruction tasks by using a given  $X \times Y \times T$  video  $\mathbf{f}_{x,y,t}$  to form the input  $\mathbf{v}_{x,y,t}$ ,  $\sigma_{x,y,t}^2$  to the epitome inference algorithm, with all parameters initialized in an uninformative way. We consider the resulting set of parameters  $v$  a video reconstruction. In comparison to “texture quilting” techniques used in Efros and Freeman (2001) for example, the epitome offers better generalization and potentially greater computational efficiency. The cost of learning an epitome is proportional to the product of the data size and the epitome size, while the cost of the quilting operations is proportional to the product of the library size and the reconstruction size. If the epitome is much smaller than the training data, then the total cost of using the epitome as the library becomes lower than the cost of using the entire training video as the library. Furthermore, epitomes can be used as a set of pointers to the original training data, so that upon the computation of the posterior  $q(\mathcal{T}_S)$ ,

the training data cubes corresponding to the several best matching epitome cubes  $\mathcal{T}_S$  can be searched over to find potentially even better candidates for quilting. This offers the advantage of using all the original training cubes in quilting while still using the computational benefit and regularization of the epitome representation.

### 5.1. Denoising

If the noise level  $\sigma^2$  in the given video  $\mathbf{f}$  is known and constant, then we set  $\mathbf{v} = \mathbf{f}$  and  $\sigma_{x,y,t}^2 = \sigma^2$  and iterate (7)–(10), starting with non-informative initialization. This procedure effectively performs denoising by averaging similar pieces of the video. The video epitome serves as a nexus of self-similarity estimation, and iteration is necessary to estimate the epitome jointly with the denoising operation. The size of the epitome is critical in this operation. If the epitome is assumed to be very large, then, the denoising operation may still copy a lot of noisy pixels into the epitome. On the other hand, a small epitome will be forced to be more liberal in the definition of similarity, due to the lack of resources to capture the diversity. Averaging many video cubes into a small epitome may lead to excessive blurring.

If the noise level  $\sigma^2$  is not known, but is assumed to be *uniform* in the video, it can be estimated from the data by adding the following step to the algorithm

$$\sigma^2 = \frac{1}{XYT} (v_{x,y,t} - \mathbf{v}_{x,y,t})^2, \quad (11)$$

which follows from setting to zero the derivative of  $F$  with respect to  $\sigma^2$ .

### 5.2. Super-Resolution

Suppose that we want to increase the video resolution of the given video  $\mathbf{f}_{x,y,t}$  by a factor of  $n$ . To guide us in this task, let us assume that we have a high resolution video  $\mathbf{h}_{x,y,t}$  of a similar scene or scenes. We can then iterate (7)–(10) using  $\mathbf{v} = \mathbf{h}$  and small noise levels  $\sigma_{x,y,t}^2 = \epsilon$ . If  $\mathbf{h}$  is assumed to be somewhat noisy, we can increase  $\sigma_{x,y,t}^2$  or learn it as described in the previous subsection. The resulting high resolution epitome  $e$  can now be used to iterate only (7) and (8) on the input video defined as:

$$\begin{aligned} \mathbf{v}_{x,y,t} &= \mathbf{f}_{x/n, y/n, t} \\ \sigma_{x,y,t}^2 &= \epsilon + \sigma^2 [\text{mod}(x, n) > 0][\text{mod}(y, n) > 0], \end{aligned} \quad (12)$$

where  $\epsilon$  is a small number and  $\sigma^2$  is large, and  $[\ ]$  is an indicator function. In other words,  $\mathbf{v}$  is a  $nX \times nY \times T$  video created by nearest-neighbor resolution enhancement of  $\mathbf{f}$ , but the map  $\sigma_{x,y,t}^2$  labels as unreliable (noisy) the pixels which do not have coordinate divisible by  $n$ . Iterating (7)

and (8) will replace the unreliable pixels with the best matching texture of  $e$  learned from  $\mathbf{h}$ .

The reader can appreciate that a similar approach can be utilized to fill in the missing “in-between” pixels even if the missing pixels are not uniformly inserted as above.

### 5.3. Object Removal and General Missing Data Reconstruction

We now consider the general case of missing value reconstruction. Define the set of coordinates  $\mathcal{G}$  for which the measurements in  $\mathbf{v}$  are considered good, and the set of coordinates  $\mathcal{M}$  for which the measurements are missing, unreliable, or need to be replaced for any reason. When learning the epitome, the set of training video cubes  $\{\mathcal{S}\}$  is chosen so that none of the cubes overlap with missing data. Then, by setting variances  $\sigma_{\mathcal{G}}^2 = \epsilon$  to be small and  $\sigma_{\mathcal{M}}^2 = \sigma^2$  to be large, the algorithm given by iterating (7), (8), (9), and (10), will produce the reconstruction  $\mathbf{v}_{\mathcal{M}}$ , quilted from the epitome which captures the patterns in  $\mathbf{v}_{\mathcal{G}}$ . For example,  $\mathcal{M}$  could mark the spatio-temporal segment which has an unwanted object.

### 5.4. Video Interpolation

Video interpolation is a special case of the missing data problem and is also similar to that of super-resolution. Here, however, we describe a different version of the problem, one in which a similar video of higher temporal resolution is *not* given.

Formally, we can set up the problem as follows. For a sequence of frames  $\mathbf{f}_{x,y,u}$  we know that some frames are missing, so that the given sequence  $u = 1, \dots, U$  corresponds to the true frames  $t(u)$ . For instance, for a video interpolation task, we would have that  $t = \{1, 3, 5, 7, \dots\}$ . On the other hand, when a video is broadcast over the Internet, some clients will experience dropped frames in random shorter or longer bursts. Then,  $t$  follow a pattern like  $t = \{1, 2, 3, 4, 5, 15, 16, 17, 18, 19, 20, 21, 29, 30, 31, 32, \dots\}$ . For video interpolation tasks, we can define the input to the epitome algorithm as:

$$\begin{aligned} \mathbf{v}_{x,y,t(u)} &= \mathbf{f}_{x,y,u} \\ \sigma_{x,y,t(u)}^2 &= \epsilon, \quad \sigma_{x,y,t \neq t(u)}^2 = \sigma^2, \end{aligned} \quad (13)$$

with  $\epsilon$  and  $\sigma^2$  being a small and a large variance respectively. Then Eqs. (7)–(10) are iterated to jointly learn the epitome and fill in the missing frames in the reconstruction  $\mathbf{v}$ . When the pattern in  $t$  is uniform, as in the case of video interpolation, there is a danger that the epitome will be updated in a similar pattern, for example by decoupling the use of odd and even frames. It is thus gener-

ally useful to constrain the epitome learning so that this does not happen. A useful constraint is that the neighboring distributions  $e_{x,y,t}$  and  $e_{x,y,t+1}$  are similar. In our experiments we simply used potential functions  $\Psi_{x,y,t}$  connecting neighboring frames in the epitome as

$$\Psi_{x,y,t} = \max_{x_n, y_n, t_n \in N(x,y,t)} e^{-(\mu_{x,y,t} - \mu_{x_n, y_n, t_n})^2 / \psi}, \quad (14)$$

where  $N(x, y, t)$  denotes a neighborhood of  $x, y, t$  outside the frame  $t$ , e.g.,  $N(x, y, t) = \{x - \delta x, \dots, x + \delta x\} \times \{y - \delta y, \dots, y + \delta y\} \times \{t - 1\} \cup \{x - \delta x, \dots, x + \delta x\} \times \{y - \delta y, \dots, y + \delta y\} \times \{t + 1\}$ . These potential functions ensure that the neighboring frames are deformed versions of one another, and define the prior

$$p(e) = \frac{1}{Z} \prod_{x_e, y_e, t_e} \Psi_{x_e, y_e, t_e}. \quad (15)$$

By including this prior and re-evaluating the derivative of the resulting free energy<sup>2</sup> with respect to the mean  $\mu_{x,y,t}$  we can see that (9) changes to

$$\mu_{x_e, y_e, t_e} = \frac{\mu_{\hat{x}_{en}, \hat{y}_{en}, \hat{t}_{en}} / \psi + \hat{\mu}_{x_e, y_e, t_e}^v / \phi_{x_e, y_e, t_e}}{1/\psi + 1/\phi_{x_e, y_e, t_e}}, \quad (16)$$

where  $\hat{\mu}_{x_e, y_e, t_e}^v$  denotes the estimate from (9), and  $(\hat{x}_{en}, \hat{y}_{en}, \hat{t}_{en}) = \arg \max_{(x_{en}, y_{en}, t_{en}) \in N(x_e, y_e, t_e)} \Psi_{x,y,t}$ .

To perform video interpolation we iterate (7), (8), (16), and (10).

Temporal interpolation and super-resolution can be combined in a straight forward manner to derive a super-resolved smooth video from a low-resolution video if it has enough pseudo-repeating structure and the sensor has a short spatio-temporal averaging field.

## 6. The Shifted Cumulative Sum Method

Careful study of the epitome learning rules reveals that many operations are repeated multiple times. For example, a single pixel  $\mathbf{v}_{x,y,t}$  in the video gets evaluated under every epitome distribution  $e_{x_e, y_e, t_e}$  several times in different combinations  $\mathcal{S}, \mathcal{T}$  in which  $(x, y, t) \in \mathcal{S}$  and  $(x_e, y_e, t_e) \in \mathcal{T}$ . Since the pixels are treated as independent, then for each patch of coordinates  $\mathcal{S}$ , a product of individual distributions is evaluated, or in the log domain, a sum of elements  $\log e_{x_e, y_e, t_e}(\mathbf{v}_{x,y,t})$  is computed. The first observation to make here is that since the epitome algorithm requires that many different overlapping patches are evaluated in this way, it becomes computationally more efficient to compute  $\log e_{x_e, y_e, t_e}(\mathbf{v}_{x,y,t})$  for all combinations of  $(x_e, y_e, t_e)$  and  $(x, y, t)$  and then sum the appropriate elements for each  $\mathcal{S}, \mathcal{T}$  combination, than to apply (8) directly.



Figure 4. Video super resolution. (Left) A single frame from a low resolution  $135 \times 90 \times 32$  sequence. (Middle) The  $360 \times 240$  bicubic interpolation of the frame. (Right) The  $360 \times 240$  video epitome super-resolution result. A  $325 \times 225 \times 8$  video epitome was learnt from a  $360 \times 240 \times 17$  high resolution sequence captured when the camera physically zoomed into a similar scene at a different point in time. Super-resolution was performed by reconstructing the low resolution sequence with the epitome containing the features from the high resolution sequence. Video cubes of size  $75 \times 75 \times 4$ ,  $50 \times 50 \times 3$ , and  $25 \times 25 \times 2$  were used during learning and reconstruction. The video of this super-resolution result is available at <http://www.psi.toronto.edu/~vincent/videoepitome.html>.

Furthermore, overlapping video cubes even share many of the summations of the  $\log e_{x_e, y_e, t_e}(v_{x, y, t})$  terms. In fact, if the overlapping cubes  $\mathcal{S}$  used in epitome operations are *all* possible cube patches of a given size (or a set of sizes), then these shared summations can be exploited in a systematic way. In particular, consider a *shifted* cumulative sum matrix

$$C(x_s, y_s, t_s) = \sum_{x < x_s, y < y_s, t < t_s} \log e_{\hat{x}_e, \hat{y}_e, \hat{t}_e}(v_{x, y, t}),$$

where  $\hat{x}_e = (x_e + x_{e_s}) \bmod X_e$  and similarly for  $\hat{y}_e$  and  $\hat{t}_e$ , which is computed for a particular offset  $x_{e_s}, y_{e_s}, t_{e_s}$ . For the cube patch  $\mathcal{T}$  starting at  $\mathcal{T}(1) = (x_{e_s}, y_{e_s}, t_{e_s})$ , we can compute at once the epitome likelihoods for all equally shaped cubes  $\mathcal{S}$  as a linear combination of up to eight elements of the  $C$  matrix.<sup>3</sup> This means that for all  $XYT$  input cube patches  $\mathcal{S}$  of a certain size, and all  $X_e, Y_e, T_e$ , the computation of all epitome cube likelihoods is computed in  $o(XYT X_e Y_e T_e)$  time regardless of the cube size  $|\mathcal{S}|$ , thus allowing us to work with multiple size video cubes in our experiments, from very small cubes capturing fine details to very large ones that help strengthen the long-range correlations in the epitome without increasing the computational complexity.

## 7. Experiments

*Video Super Resolution.* An optical zoom in modern cameras allows the user to trade the field of view for the level of captured detail. The user often desires to capture both the large context of the scene and the detail in it, by first capturing a wide angle shot followed by a large zoom (as shown in the set-up video on the web page) and slow scene scanning at the constant zoom level (omitted in the set-up video for brevity). We can use the approach described in Section 5.2 to super resolve the original low resolution wide-angle shot  $\mathbf{f}$ , using the epitome learnt on

$\mathbf{h}$ , the high-resolution “scene scanning” video captured at the higher zoom level.

In our example, the wide shot captures a large plant moving in the wind. Later, the camera zooms in and scans the plant creating the shot used for training the high resolution epitome. As described in Section 5.2, the high resolution textural and shape features of the plant, as well as the motion patterns, are then used to compute a super-resolved version of the wide shot. Note that the content of the wide shot is similar but far from identical to the content of the high resolution training data. Figure 4 shows a single frame of the super-resolution result compared to bicubic interpolation, and the web page also shows a larger super-resolution sequence of the plant.

*Reconstructing Dropped Frames from a Video Broadcast.* Real-time streaming videos are becoming increasingly popular on the web. In streaming video, often video frames are dropped because of a lack of bandwidth. Client-side recovery of these missing frames using only the successfully received frames would greatly increase the experience with streaming videos.

Figure 5 shows a few frames from a video sequence where the dropped frames effect was simulated and the missing frames were then reconstructed using the video epitome as described in Section 5.4. The video epitome is able to consolidate the various disconnected frames and assemble a comprehensible set of motion patterns in the video sequence. In the video on the web page, we show the received frames on the left (with freezes during the frame drops) and on the right, the video reconstruction using only the video on the left as the epitome input. Note that due to the large number of missing frames, most original motion patterns have never been seen by the epitome algorithm in their uncorrupted form anywhere in the received broadcast. Exemplar-based approaches (cf. Wexler et al., 2004), while able to fill in some missing frames provided that similar motions appear elsewhere in

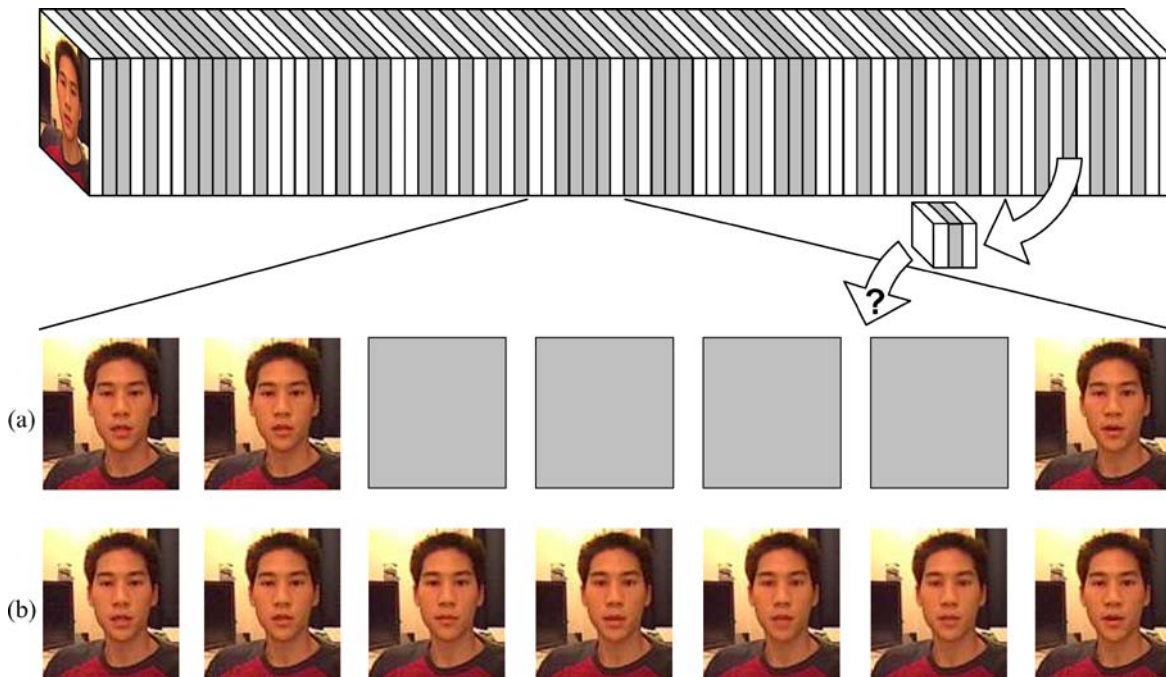


Figure 5. Dropped-frames experiment. This experiment deals with the problem of reconstructing missing frames when receiving streaming video over the Internet. The problem is illustrated by the diagram at the top, where gray indicates frames that were dropped during transmission. (a) Several frames from the sequence in which four of the frames are missing. Most video cubes from the sequence contain missing frames, so they cannot simply be stitched together to fill in the missing frames as with exemplar-based methods. (b) The reconstructed missing frames using the video epitome of the sequence with dropped frames. The arrival time of the frames in this  $90 \times 100 \times 39$  sequence was modeled as a discrete time Bernoulli process whereby the inter-arrival time of the frames was governed by a geometric distribution. Frames were expected to arrive in each time slot and any unfilled time slots between arrived frames were considered to be dropped frames. The mean number of missing frames between arrived frames was one. A  $90 \times 100 \times 13$  video epitome was learnt from only the non-dropped frames using video patches of size  $63 \times 70 \times 6$ , and  $36 \times 40 \times 5$ , and  $30 \times 30 \times 4$ . This video sequence is available at <http://www.psi.toronto.edu/~vincent/videoepitome.html>.




Figure 6. Video in-painting experiment using the video epitome. (a) Several frames from a  $115 \times 83 \times 53$  video. (b) Removal of the fire hydrant from the video by considering these pixels as missing. (c) In-painted frames using a  $30 \times 30 \times 15$  epitome with  $20 \times 20 \times 5$ ,  $15 \times 15 \times 4$ ,  $10 \times 10 \times 3$ , and  $5 \times 5 \times 2$  video cubes. Video available at <http://www.psi.toronto.edu/~vincent/videoepitome.html>.




Figure 7. Missing channels experiment. A  $115 \times 83 \times 24$  corrupted video sequence where each of the RGB color channels of each pixel was missing with 50% probability was recovered by learning its video epitome and filling in the missing values using the epitome. The  $40 \times 40 \times 12$  epitome was learnt from only the observed values and does not contain missing data because of the consolidation of the  $20 \times 20 \times 4$ ,  $15 \times 15 \times 3$ , and  $10 \times 10 \times 2$  video cubes from the video into the smaller video epitome during learning. Video available at <http://www.psi.toronto.edu/~vincent/videoepitome.html>.

the video sequence, cannot handle this situation, as nearly every video cube in the sequence contains missing data.

*Video in-painting.* Video epitomes contain the video’s basic structural and motion characteristics, which are useful for in-painting applications. The goal of video in-painting is to fill in missing portions of a video, which can arise with damaged films or occluding objects. Video in-painting can be performed with the video epitome as described in Section 5.3.

Figure 6 and its corresponding video on the web show the results of in-painting on a video of a girl walking. The results here are similar to those of Wexler et al. (2004). Part way through the video, the girl is occluded by a fire hydrant. Removing the fire hydrant can be formulated as a video reconstruction problem by considering these pixels as missing, learning the video epitome only on the observed pixel values, and performing epitome inference while setting the variances,  $\sigma^2$  to be high in the area occupied by the hydrant. The video epitome is able to compress the basic walking motion into several frames and transfers this motion pattern into the missing pixels.

*Denoising.* In Section 5.3, we discussed the general case of filling in arbitrary missing data in video given only the corrupted video. Figure 7 illustrates the potential power of the video epitome. A highly corrupted video was created in which each of the RGB color channels of each pixel was missing with 50% probability. The known bad channels for each pixel were marked by high noise variances  $\sigma^2$ . The video was then given to the epitome learning algorithm and then reconstructed by iterating (7)–(10). Despite high levels of corruption, the repetitive motion in the video helped the epitome learn a clean representation and reconstruct the video  $v$ . This video is available on the web page.

## 8. Conclusion

An epitome represents many of the high-order statistics in array data using a more compact array, and is amenable to faster searching and better generalization, compared to patch libraries. Here, we developed an efficient algorithm for learning video epitomes and performing various inference tasks. We illustrated some of the applications of video epitomes, including video enhancement and editing tasks. Videos demonstrating a variety of applications are available at <http://www.psi.toronto.edu/~vincent/videoepitome.html>.

In comparison to reconstruction techniques based on a library of known data patches, an advantage of epitomes is that they can be trained directly on corrupted or degraded data, as long as the data is repetitive. Video tends to be highly redundant and is thus quite well-suited

to analysis using epitomes. Because epitomes provide a representation that retains the natural flow of the input data, and offer significant computational and statistical advantages over patch libraries, we believe they will find many uses in data analysis and computer vision.

## Appendix: Derivation of the Video Epitome Learning Algorithm

Learning under the video epitome model is performed by using a variational expectation-maximization algorithm, which centers around minimizing the free energy cost function in (4). The free energy bounds the log-likelihood of the input video, so by optimizing  $F$ , the log-likelihood of the input video is maximized, the proof of which can be shown by using Jensen’s inequality:

$$\begin{aligned}
 F &= \sum_S \sum_{\mathcal{T}_S} \int_{\mathbf{c}_S} q(\{\mathcal{T}_S, \mathbf{c}_S\}) \log \frac{q(\{\mathcal{T}_S, \mathbf{c}_S\})}{p(\mathbf{v}, \{\mathbf{c}_S, \mathcal{T}_S\})} \\
 &= - \sum_S \sum_{\mathcal{T}_S} \int_{\mathbf{c}_S} q(\{\mathcal{T}_S, \mathbf{c}_S\}) \log \frac{p(\mathbf{v}, \{\mathbf{c}_S, \mathcal{T}_S\})}{q(\{\mathcal{T}_S, \mathbf{c}_S\})} \\
 &\geq - \log \sum_S \sum_{\mathcal{T}_S} \int_{\mathbf{c}_S} q(\{\mathcal{T}_S, \mathbf{c}_S\}) \frac{p(\mathbf{v}, \{\mathbf{c}_S, \mathcal{T}_S\})}{q(\{\mathcal{T}_S, \mathbf{c}_S\})} \\
 &= - \log \sum_S \sum_{\mathcal{T}_S} \int_{\mathbf{c}_S} p(\mathbf{v}, \{\mathbf{c}_S, \mathcal{T}_S\}) \\
 &= - \log p(\mathbf{v}).
 \end{aligned}$$

Learning progresses by iteratively updating the approximation to the true posterior,  $q(\{\mathcal{T}_S, \mathbf{c}_S\})$  and the variational parameters: the hidden video,  $v$ , and the mean and variance of the epitome,  $\mu$  and  $\phi$ . The update equations for the  $q$ -distribution and the parameters (7–10) are obtained by setting to zero the derivatives of the free energy cost function, the procedure of which is shown below.

The posterior epitome responsibilities  $q(\mathcal{T}_S)$  are obtained by using a Lagrange multiplier during optimization to constrain  $q$  to be a proper distribution:

$$\begin{aligned}
 \frac{\partial(\lambda(\sum_{\mathcal{T}} q(\mathcal{T})-1)+F)}{\partial q(\mathcal{T}_S)} &= \frac{\partial}{\partial q(\mathcal{T}_S)} \left( \lambda \left( \sum_{\mathcal{T}} q(\mathcal{T}) - 1 \right) \right. \\
 &\quad \left. + \sum_S \sum_{\mathcal{T}} \int_{\mathbf{c}_S} \left( \prod_S q(\mathcal{T}_S) q(\mathbf{c}_S) \right) \log \frac{\prod_S q(\mathcal{T}_S) q(\mathbf{c}_S)}{p(\mathbf{v}, \{\mathbf{c}_S, \mathcal{T}_S\})} \right) \\
 0 &= \frac{\partial}{\partial q(\mathcal{T}_S)} \left( \lambda q(\mathcal{T}_S) + q(\mathcal{T}_S) \log q(\mathcal{T}_S) \right. \\
 &\quad \left. - q(\mathcal{T}_S) \int_{\mathbf{c}_S} \left( \prod_S q(\mathbf{c}_S) \right) \log p(\mathbf{v}, \{\mathbf{c}_S, \mathcal{T}_S\}) \right)
 \end{aligned}$$

$$\begin{aligned}
0 &= \lambda + \log q(\mathcal{T}_S) + 1 - \log p(\mathbf{v}, \{v_S\}, \mathcal{T}_S) \\
q(\mathcal{T}_S) &= \exp(-\lambda - 1) p(\mathbf{v}, \{v_S\}, \mathcal{T}_S) \\
q(\mathcal{T}_S) &= \exp(-\lambda - 1) p(\mathbf{v}|\{v_S\}) p(\mathcal{T}_S) e_{\mathcal{T}_S}(v_S) \\
1 &= \sum_{\mathcal{T}} q(\mathcal{T}) = \exp(-\lambda - 1) p(\mathbf{v}|\{v_S\}) \sum_{\mathcal{T}} p(\mathcal{T}) e_{\mathcal{T}}(v_S) \\
\exp(-\lambda - 1) p(\mathbf{v}|\{v_S\}) &= \frac{1}{\sum_{\mathcal{T}} p(\mathcal{T}) e_{\mathcal{T}}(v_S)},
\end{aligned}$$

by substituting this last result into the equation for  $q(\mathcal{T}_S)$  above, the following update is obtained:

$$q(\mathcal{T}_S) \leftarrow \frac{p(\mathcal{T}_S) e_{\mathcal{T}_S}(v_S)}{\sum_{\mathcal{T}} p(\mathcal{T}) e_{\mathcal{T}}(v_S)}.$$

The update equation for the posterior cube colors  $v_i$ , where  $i = (x, y, t)$ , are derived by minimizing the free energy w.r.t.  $v_i$ , and as such, only terms involving  $v_i$  need to be considered,

$$\begin{aligned}
\frac{\partial F}{\partial v_i} &= -\frac{\partial}{\partial v_i} \sum_S \sum_{\mathcal{T}} \int_{\mathbf{c}_S} \left( \prod_S q(\mathcal{T}_S) q(\mathbf{c}_S) \right) \\
&\quad (\log p(\mathbf{v}_i|\{\mathbf{c}_S\}) + \log p(\mathbf{c}_S|\mathcal{T}_S)) \\
0 &= \frac{\partial}{\partial v_i} \sum_{S,k:S(k)=i} \sum_{\mathcal{T}_S} q(\mathcal{T}_S) (\log p(\mathbf{v}_i|\{v_S\}) \\
&\quad + \log e_{\mathcal{T}_S(k)}(v_{S(k)})) \\
0 &= \frac{\partial}{\partial v_i} \left( (v_i - v_i)^2 / 2\sigma_i^2 \right. \\
&\quad \left. + \sum_{S,k:S(k)=i} \sum_{\mathcal{T}_S} q(\mathcal{T}_S) (v_i - \mu_{\mathcal{T}_S(k)})^2 / 2\phi_{\mathcal{T}_S(k)} \right) \\
0 &= -(v_i - v_i) / \sigma_i^2 \\
&\quad + \sum_{S,k:S(k)=i} \sum_{\mathcal{T}_S} q(\mathcal{T}_S) (v_i - \mu_{\mathcal{T}_S(k)}) / \phi_{\mathcal{T}_S(k)},
\end{aligned}$$

gathering and isolating variables gives the following result:

$$v_{x,y,t} \leftarrow \frac{\frac{v_{x,y,t}}{\sigma_{x,y,t}^2} + \sum_{S,k:S(k)=(x,y,t)} \sum_{\mathcal{T}_S} q(\mathcal{T}_S) \frac{\mu_{\mathcal{T}_S(k)}}{\phi_{\mathcal{T}_S(k)}}}{\frac{1}{\sigma_{x,y,t}^2} + \sum_{S,k:S(k)=(x,y,t)} \sum_{\mathcal{T}_S} q(\mathcal{T}_S) \frac{1}{\phi_{\mathcal{T}_S(k)}}}.$$

Optimizing for the epitome parameter  $\mu_j$ ,  $j = (x_e, y_e, t_e)$ , is done in a similar manner:

$$\begin{aligned}
\frac{\partial F}{\partial \mu_j} &= -\frac{\partial}{\partial \mu_j} \sum_S \sum_{\mathcal{T}} \int_{\mathbf{c}_S} \left( \prod_S q(\mathcal{T}_S) q(\mathbf{c}_S) \right) \log p(\mathbf{c}_S|\mathcal{T}_S) \\
0 &= \frac{\partial}{\partial \mu_j} \sum_{\mathcal{T},k:\mathcal{T}(k)=j} \sum_S q(\mathcal{T}_S = \mathcal{T}) \log p(v_{S(k)}|\mathcal{T}_S = \mathcal{T}) \\
0 &= \frac{\partial}{\partial \mu_j} \sum_{\mathcal{T},k:\mathcal{T}(k)=j} \sum_S q(\mathcal{T}_S = \mathcal{T}) (v_{S(k)} - \mu_j)^2 / 2\phi_j
\end{aligned}$$

$$\begin{aligned}
0 &= \sum_{\mathcal{T},k:\mathcal{T}(k)=j} \sum_S q(\mathcal{T}_S = \mathcal{T}) (v_{S(k)} - \mu_j) / \phi_j \\
\mu_{x_e, y_e, t_e} &\leftarrow \frac{\sum_{\mathcal{T},k:\mathcal{T}(k)=(x_e, y_e, t_e)} \sum_S q(\mathcal{T}_S = \mathcal{T}) v_{S(k)}}{\sum_{\mathcal{T},k:\mathcal{T}(k)=(x_e, y_e, t_e)} \sum_S q(\mathcal{T}_S = \mathcal{T})}.
\end{aligned}$$

The steps in deriving the update equation for the epitome parameter  $\phi_j$ , where  $j = (x_e, y_e, t_e)$ , follows the same initial steps as above:

$$\begin{aligned}
\frac{\partial F}{\partial \phi_j} &= \frac{\partial}{\partial \phi_j} \sum_{\mathcal{T},k:\mathcal{T}(k)=j} \sum_S q(\mathcal{T}_S = \mathcal{T}) \\
&\quad \times \left( (v_{S(k)} - \mu_j)^2 / 2\phi_j + \frac{1}{2} \log \phi_j \right) \\
0 &= \sum_{\mathcal{T},k:\mathcal{T}(k)=j} \sum_S q(\mathcal{T}_S = \mathcal{T}) \\
&\quad \times \left( - (v_{S(k)} - \mu_j)^2 / 2\phi_j^2 + \frac{1}{2\phi_j} \right)
\end{aligned}$$

$$\begin{aligned}
\phi_{x_e, y_e, t_e} &\leftarrow \\
&\frac{\sum_{\mathcal{T},k:\mathcal{T}(k)=(x_e, y_e, t_e)} \sum_S q(\mathcal{T}_S = \mathcal{T}) (v_{S(k)} - \mu_{x_e, y_e, t_e})^2}{\sum_{\mathcal{T},k:\mathcal{T}(k)=(x_e, y_e, t_e)} \sum_S q(\mathcal{T}_S = \mathcal{T})}.
\end{aligned}$$

## Acknowledgments

We thank P. Anandan, A. Blake, B. Freeman, M. Irani, A. Kannan, J. Malik, E. Simoncelli, and S. C. Zhu for helpful discussions. This work was supported by a Discovery Grant from NSERC. V. Cheung was financially supported by an NSERC Canada Graduate Scholarship and B. Frey was financially supported by a Premier's Research Excellence Award. B. Frey is a Fellow of the Canadian Institute for Advanced Research.

## Notes

1. For example, if the responsible epitome cube for the video cube at  $\{1, \dots, 5\} \times \{1, \dots, 5\} \times \{1, \dots, 5\}$  is at epitome coordinates  $\{11, \dots, 15\} \times \{11, \dots, 15\} \times \{11, \dots, 15\}$ , then the video cube at  $\{2, \dots, 6\} \times \{3, \dots, 7\} \times \{4, \dots, 8\}$  will tend to be mapped to epitome coordinates  $\{12, \dots, 16\} \times \{13, \dots, 17\} \times \{14, \dots, 18\}$ , so that the for instance, the video coordinate  $(3, 4, 5)$  maps to the epitome coordinate  $(13, 14, 15)$  in both mappings. When the inference procedure does not result in such translational consistency, then it usually results in appearance consistency, i.e., each pixel tends to map to similar distributions  $e_{x,y,t}$  in all mappings.
2. We ignore the normalization constant  $Z$ .
3. For example, for  $\mathcal{T} = \{5, \dots, 10\} \times \{10, \dots, 20\} \times \{2\}$ , if we set  $(x_{es}, y_{es}, t_{es}) = (5, 10, 2)$ , the resulting shifted cumulative matrix  $C$  leads to  $\log e_{\mathcal{T}}(v_{[15,20]}) \log e_{\mathcal{T}}(v_{[15,20] \times [50,60] \times [17]}) = C(20, 60, 17) - C(15, 60, 17) - C(20, 50, 17) + C(15, 50, 17)$ .

## References

- Bishop, C., Blake, A., and Marthi, B. 2003. Super-resolution enhancement of video. In *Proc Artificial Intelligence and Statistics*.

- Buades, A., Coll, B., and Morel, J.-M. 2005. A non-local algorithm for image denoising. In *Proc. IEEE Computer Vision and Pattern Recognition*, pp. 60–65.
- Criminisi, A., Pérez, P., and Toyama, K. 2003. Object removal by exemplar-based inpainting. In *Proc. Conf. Computer Vision and Pattern Recognition*, pp. 721–728.
- Efros, A.A. and Freeman, W.T. 2001. Image quilting for texture synthesis and transfer. In *Proc. SIGGRAPH*, pp. 341–346.
- Freeman, W.T., Jones, T.R., and Pasztor, E.C. 2002. Example-based super-resolution. *IEEE Computer Graphics and Applications*, pp. 56–65.
- Frey, B.J. and Jojic, N. 2005. A comparison of algorithms for inference and learning in probabilistic graphical models. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 27(9):1392–1416.
- Jepson, A. and Black, M. 1993. Mixture models for optical flow computation. In *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, pp. 760–761.
- Jojic, N., Frey, B.J., and Kannan, A. 2003. Epitomic analysis of appearance and shape. In *Proc. IEEE Intern. Conf. Computer Vision*, pp. 34–41.
- Neal, R.M. and Hinton, G.E. 1998. A view of the em algorithm that justifies incremental, sparse, and other variants. In *Learning in Graphical Models*, Jordan, M.I. (Ed.), pp. 355–368.
- Rosales, R., Achan, K., and Frey, B.J. 2003. Unsupervised image translation. In *Proc. IEEE Intern. Conf. Computer Vision*, pp. 472–478.
- Wang, J.Y.A., Adelson, E.H., and Desai, U.Y. 1994. Applying mid-level vision techniques for video data compression and manipulation. In *Proc. SPIE on Digital Video Compression on Personal Computers: Algorithms and Technologies*, pp. 116–127.
- Wexler, Y., Shechtman, E., and Irani, M. 2004. Space-time video completion. In *Proc. IEEE Computer Vision and Pattern Recognition*, pp. 120–127.
- Zhu, S.C., Guo, C., Wu, Y.N., and Wang, Y. 2002. What are textons? In *Proc. 7th European Conf. on Computer Vision-Part IV*, pp. 793–807.