

Exploring the Manifold of Image Patches

Yevgen Matviychuk, Shannon M. Hughes
 Department of Electrical, Computer, and Energy Engineering
 University of Colorado at Boulder

yevgen.matviychuk@colorado.edu, shannon.hughes@colorado.edu

Abstract

In image processing, the idea of considering an entire image as a collection of its small overlapping regions has recently attained significant popularity. Algorithms that work with image patches, instead of an image as a whole, produce state-of-the-art results in solving reconstruction problems. Meanwhile, it has been shown that manifold models may provide a good approximation for the set of natural image patches. In this paper we consider one such popular model. We will show some interesting and aesthetically appealing geometric structures arising from an attempt to visualize an embedding of the patch manifold in three-dimensional space.

Introduction

Consider a photograph of a natural scene taken with a digital camera and imagine zooming into very small parts of this image, for example, only 5×5 or 10×10 pixels. You may observe that such image regions—often called patches—if sufficiently small look very similar across the photograph. Regardless of the part of the original photograph you were zooming into, you likely arrived to either a patch of almost uniform intensity (e.g., the sky), a patch containing a sharp contrast edge (e.g., horizon line), or possibly a part of some pattern or texture (e.g., woven fabric). In either case the resulting patch can be very accurately described with just a few parameters, such as the color of the uniform fill or the position of the edge.

As an illustrative example, consider an image of a zebra shown in Fig. 1. All patches on the zebra's hide can be well approximated with simple black-and-white wedges and therefore described with just two parameters: for example, the distance d from the black and white boundary to the center of the patch and the angle α of the normal vector \vec{n} to the boundary pointing from black to white. Therefore, even though each patch may contain dozens of pixels, we need only two numbers to specify it. Potentially this offers a concise and robust description of the entire image as a collection of its (possibly overlapping) patches.



Figure 1: An example of a natural image. Sufficiently small patches of the image can be well approximated with just a pair of parameters.

The phenomenon of similarity between different image patches was recently observed and used in numerous successful patch-based image processing algorithms. Methods based on jointly filtering similar image patches show state-of-the-art performance in removing noise [3], increasing image resolution, or reconstructing images from their very few linear measurements [4]. Furthermore, patch-based methods produce exceptional results in solving higher-level image processing problems. For example, content aware editing tools in Adobe Photoshop allow users to fill in missing parts of images using patches from the rest of the image or even seamlessly rearrange image parts [1].

Another important observation about the set of image patches comes from viewing them as points residing in a high-dimensional space. Indeed, each $n \times n$ -pixel patch is just a vector in $\mathbb{R}^{n \cdot n}$. However, due

to their inherent structure, valid image patches occupy only a small subset of this large space. Interestingly, this subset often exhibits a manifold geometry [5, 6]. Returning to our example of high-contrast wedge patches, it can be shown that they lie on a two-dimensional manifold (a surface parametrized by continuous parameters α and d) in the high-dimensional ambient space.

Next, we will make an attempt to explore and visualize the underlying manifold of image patches using a simple approximating model. We will generate interesting 3D structures by restricting the ambient space of patches $\mathbb{R}^{n \times n}$ to only three dimensions.

Model of Gradient Image Patches

We consider a simple yet popular analytical model of gradient $n \times n$ patches. It was used in [5] and [2] to analyze the statistical and topological properties of the set of natural image patches. In particular, it was shown in [2] that the underlying patch manifold has the topology of a Klein bottle.

To set up the model, let $\mathbf{p}, \mathbf{q} \in \mathbb{R}^n$ be two vectors of n numbers that are equally spaced on the interval $[-1, 1]$. All pairs $(\mathbf{p}_i, \mathbf{q}_j)$, for $i, j = 1, \dots, n$, form a coordinate grid, which we will use to build our patches. We define the intensity of each pixel (i, j) as a polynomial function of its coordinates on this grid, $P(\mathbf{p}_i, \mathbf{q}_j) = c(a\mathbf{p}_i + b\mathbf{q}_j)^2 + d(a\mathbf{p}_i + b\mathbf{q}_j)$, where $a, b, c, d \in \mathbb{R}$ are some constant parameters. To bound the energy of the patches, we additionally require each pair (a, b) and (c, d) to lie on the unit circle, i.e. $a^2 + b^2 = 1$ and $c^2 + d^2 = 1$. Finally, we define the polynomial $P(p, q)$ as:

$$P(\mathbf{p}_i, \mathbf{q}_j) = \sin \beta (\cos \alpha \cdot \mathbf{p}_i + \sin \alpha \cdot \mathbf{q}_j)^2 + \cos \beta (\cos \alpha \cdot \mathbf{p}_i + \sin \alpha \cdot \mathbf{q}_j), \quad (1)$$

where the parameters $\alpha, \beta \in [0, \dots, 2\pi]$ are introduced to take into account the constraints on a, b, c, d . Examples of patches generated with this method for different values of α and β are shown in Fig. 2. We note that α has a meaning of an angular parameter setting the direction of the gradient, and β regulates the amount of second order gradients (ridges) in the patches.

Alternatively, we can think of the polynomial P evaluated on a fixed coordinate pixel grid as a mapping from the space of parameters $(\alpha, \beta) \in \mathbb{R}^2$ to the high-dimensional space of patches $P: \mathbb{R}^2 \rightarrow \mathbb{R}^{n \times n}$, where it defines a two-dimensional smooth manifold \mathcal{M} . We will attempt to visualize an embedding of this surface in 3D. For this, we restrict our attention to only three corner pixels of each patch, the ones with (p, q) coordinates equal to $\begin{bmatrix} -1 \\ 1 \end{bmatrix}$, $\begin{bmatrix} 1 \\ -1 \end{bmatrix}$, and $\begin{bmatrix} 1 \\ 1 \end{bmatrix}$. We evaluate the polynomial P at these positions and treat the resulting triplets as Cartesian (x, y, z) -coordinates of points in three-dimensional space:

$$\begin{bmatrix} x \\ y \\ z \end{bmatrix} = \begin{bmatrix} \sin \beta (\sin \alpha - \cos \alpha)^2 + \cos \beta (\sin \alpha - \cos \alpha) \\ \sin \beta (\cos \alpha - \sin \alpha)^2 + \cos \beta (\cos \alpha - \sin \alpha) \\ \sin \beta (\cos \alpha + \sin \alpha)^2 + \cos \beta (\cos \alpha + \sin \alpha) \end{bmatrix}. \quad (2)$$




Figure 2: Examples of patches generated using the Eq. 1 for different values of α and β .

Now changing α and β allows us to generate any point on the embedding of the approximation to the patch manifold \mathcal{M} in 3D and thus visualize this surface.

Visualization of the Patch Manifold

In our experiments, instead of simply generating points on \mathcal{M} for different pairs of α and β we fix one of these parameters and gradually change the other in the interval $[0, 2\pi]$. This method produces closed curves (loops) on the surface of the manifold. Some plots of these curves in 3D (viewed from several different directions) are shown in Fig. 3 and Fig. 4.




Figure 3: Different views of the loops on the (same) patch manifold generated with Eq. 2. Each loop corresponds to a fixed value of parameter α , while β is changing from 0 to 2π .




Figure 4: Different views of the loops on the (same) patch manifold generated with Eq. 2. Each loop corresponds to a fixed value of parameter β , while α is changing from 0 to 2π .

To be able to create 3D objects to visualize the considered patch manifold model we add some width to these curves, effectively making them wires. We then combine both sets of loops generated for fixed α and β , which connects otherwise disjoint curves. This effectively results in a mesh on the surface of the manifold, which can be printed on a 3D printer. Figure 5 shows a visualization of this object (viewed from different perspectives) and a photograph of the printed figurine.

Finally, Fig. 6 shows yet another possible figurine created based on the described patch manifold model. Here we use only the curves generated for several constant values of β (while gradually changing α) and include only a half of the manifold surface.

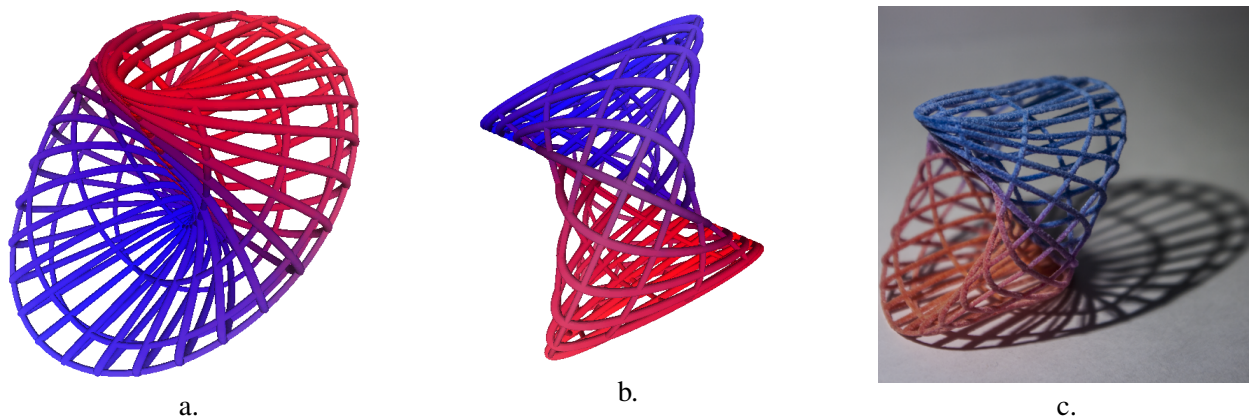


Figure 5: (a.-b.): Visualization of the mesh on the surface of the patch manifold in 3D viewed from two different viewpoints. (c.) Photograph of the printed 3D model. Notice the self intersection of the surface caused by restricting the embedding only to three dimensions.




Figure 6: Another possible variation of a sculpture based on the discussed patch manifold model (viewed from two different positions).

Conclusion

In this paper we considered a simple but popular analytical approximation to the manifold of natural image patches. Visualization of the arising surface not only allows one to gain deeper understanding of the underlying patch manifold but can also be used to create aesthetically pleasing pictures and 3D objects.

References

- [1] C. Barnes, E. Shechtman, A. Finkelstein, and D. B. Goldman. Patchmatch: a randomized correspondence algorithm for structural image editing. *ACM Trans. Graph.*, 28(3):24:1–24:11, 2009. (document)
- [2] G. Carlsson, T. Ishkhanov, V. Silva, and A. Zomorodian. On the local behavior of spaces of natural images. *Int. J. Comput. Vision*, 76(1):1–12, January 2008. (document)
- [3] K. Dabov, A. Foi, V. Katkovnik, and K. Egiazarian. Image denoising by sparse 3-D transform-domain collaborative filtering. *IEEE Transactions on Image Processing*, 16(8):2080–2095, 2007. (document)
- [4] A. Danielyan, A. Foi, V. Katkovnik, and K. Egiazarian. Spatially adaptive filtering as regularization in inverse imaging: compressive sensing, upsampling, and super-resolution. 2010. (document)
- [5] A. B. Lee, K. S. Pedersen, and D. Mumford. The nonlinear statistics of high-contrast patches in natural images. *Int. J. Comput. Vision*, 54(1-3):83–103, August 2003. (document)
- [6] G. Peyré. Manifold models for signals and images. *Computer Vision and Image Understanding*, 113(2):249–260, February 2009. (document)