
Learning transport operators for image manifolds: supplement

Benjamin J. Culpepper
Department of EECS
Computer Science Division
University of California, Berkeley
Berkeley, CA 94720
bjc@cs.berkeley.edu

Bruno A. Olshausen
Helen Wills Neuroscience Institute
& School of Optometry
University of California, Berkeley
Berkeley, CA 94720
baolshausen@berkeley.edu

In cases where topology cannot be captured by a matrix exponential operator, we would like the operators to capture local structure. One manifold of interest to us is the Klein bottle, since studies have shown that it is the topology of natural image patches [1] and is reflected in the organization of cortical orientation maps [2]. However, there are many pairs of points on the Klein bottle that cannot be related through an operator whose trajectory exactly tracks the manifold. To see this, recall that the surface of a Klein bottle can be viewed as a 2-dimensional sheet with a periodic horizontal boundary, and a reflected vertical boundary, as depicted in figure 1(a). Two points related by non-zero horizontal and vertical displacements are connected through a diagonal line on this sheet. Locations A and B on the sheet correspond to the same point, as do C and D, due to the side identifications. The line intersects itself at point E, and at the intersection, the directions of travel are perpendicular. Since the effect of the Lie group operators is entirely determined by \mathbf{x} , there can be no operator \mathbf{A} that exhibits this global behavior. However, when given two nearby points, our inference procedure should deliver an \mathbf{A} that can be used to locally interpolate between points along the surface; as the separation distance increases, the interpolation path will begin to slip off. For the case of the Klein bottle, we investigate the extent to which this slippage occurs by attempting to interpolate over sufficiently long segments.

Related pairs of points on a Klein bottle are generated by choosing two angles θ_0, ϕ_0 uniformly at random from $[0, 2\pi]$; two related angles θ_1, ϕ_1 are produced by sampling from two von Mises distributions with means θ_0 and ϕ_0 , and concentration $\kappa = 5$. We use the 4-dimensional embedding of [3]: $\mathbf{x}_t = (\cos \theta_t \cos \phi_t, \sin \theta_t \cos \phi_t, 2 \cos \frac{\theta_t}{2} \sin \phi_t, 2 \sin \frac{\theta_t}{2} \sin \phi_t)$. Our interpolation along the path shown in figure 1(c) is three times better than linear interpolation (as measured by MSE), and when both techniques have twice as many segments, ours is superior by an order of magnitude. This illustrates the robustness of our algorithm in a case where the underlying data cannot be captured by a linear ODE: a piecewise model, where each piece is approximated by a linear ODE using an \mathbf{A} generated from a learned basis Ψ , gives a sensible answer.

References

- [1] Carlsson, G., Ishkhanov, T., de Silva, V., & Zomorodian, A. (2007) On the Local Behavior of Spaces of Natural Images. *Int. J. Comput. Vision*, 76(1): 1-12.
- [2] Swindale, N.V. (1996) Visual cortex: Looking into a Klein bottle. *Current Biology* 6(7): 776-779.
- [3] Tompkins, C. (1941) A flat Klein bottle isometrically embedded in Euclidean 4-space. *Bull. Am. Math. Soc.*, 47(6): 508.

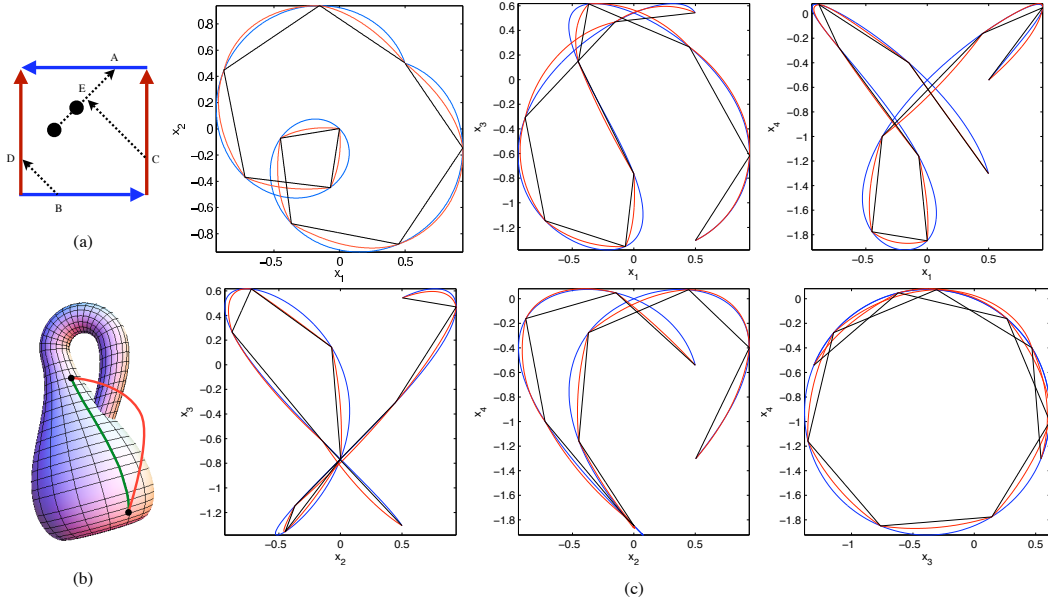


Figure 1: **Learning the Klein bottle manifold.** (a) The line connecting the two points intersects itself at E, and at this intersection the line is perpendicular to itself, making clear why a single \mathbf{A} operator which interpolates between two such points cannot globally stay on the Klein bottle. (b) The green path connects two points on the manifold through a path along it. Although both endpoints lie on the manifold, the red path shows an example (for purposes of illustration) of an interpolated path between them does not. (c) Real piecewise approximations made by our algorithm, compared to a true trajectory along the manifold (in the embedded 4-space). The two underlying angles are initialized to $\theta_0 = \pi/4$, $\phi_0 = -\pi/4$ then scanned linearly to $\theta_1 = -3\pi$, $\phi_1 = 3/2\pi$ in 10 equally space segments. The endpoints of the resulting 4-space curves (blue) are then used to estimate an operator \mathbf{A} that interpolates between them (red). This interpolation is compared to linear interpolation in the 4-space (black).