Manifold and group action disentanglement

Yeari Vigder   Amit Moscovich

Motivation

Disentangled representation learning is a field in machine learning that aims to separate the underlying factors or variables that generate data into **distinct and interpretable components**. The goal is to **disentangle the essential features of the data,** such as object identity, pose, and lighting, so that they can be analyzed independently.

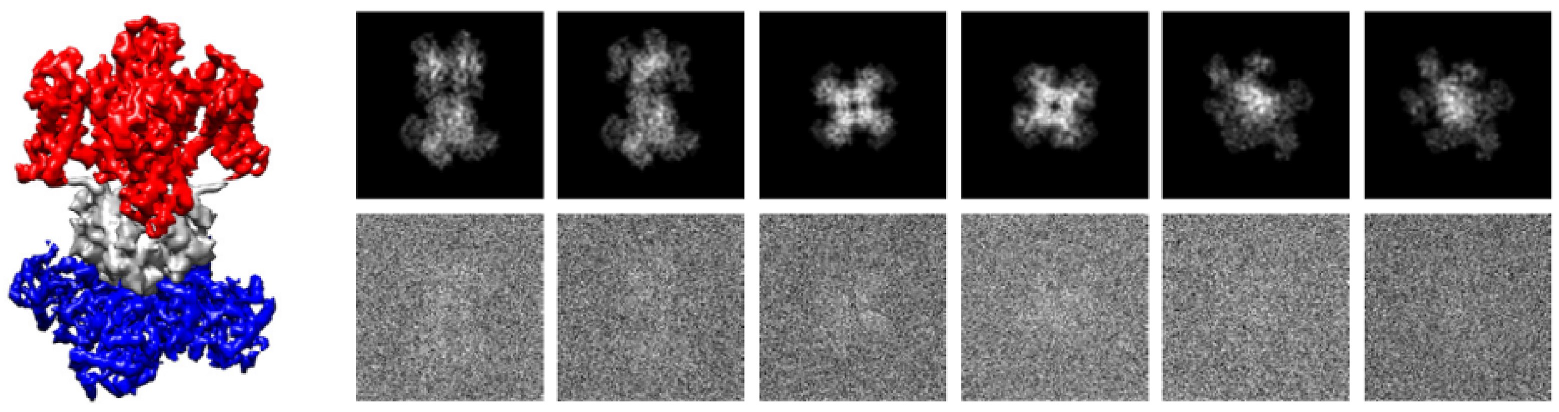

Figure 1. (left) 3D model of a macromolecule with two independent motions of the top and bottom. (right) 2D projection images of the same macromolecule captured from different orientations. The orientations correspond to the elements of the rotation group $SO(3)$ acting on the molecule. In addition, the intrinsic motion can be seen in the 2D images.

Graph Laplacian eigendecomposition for dim. reduction

Given a sample of N points $x_1, \dots, x_N$:

- Compute $W \in M_N$ the matrix of edge weights. Gaussian affinities are often used for the weights:
  $$W_{ij} = \exp^{\frac{\|x_i - x_j\|^2}{2\varepsilon}}$$
  Another common approach is to define W to be the Adjacency matrix of the k nn of each point.
- Compute D, the diagonal degree matrix that satisfies $D_{ii} = \sum_{j=1}^{N} W_{ij}$
- The RW Laplacian defined by $L_N = I - D^{-1}W$, and it is positive semi-definite matrix.
- Denote the eigenvectors of L_N that correspond to the smallest eigenvalues by $\phi_1, \dots, \phi_l \in \mathbf{R}^N$.
- Laplacian eigenvectors can be used for dimensionality reduction by the map,
  $$x_i \mapsto [\phi_2(x_i), \dots, \phi_l(x_i)].$$
  This approach was pioneered by [1] under the name Laplacian eigenmaps and later analyzed and extended by [3].
- When $N \to \infty$ The limit operator of L_N is the Laplace-Beltrami operator $\Delta_{\mathcal{M}}$, and the eigenvectors and the eigenvalues of L_N converges to eigenfunctions and the eigenvalues of $\Delta_{\mathcal{M}}$, respectively.

The Laplace-Beltrami operator on product manifolds

Theorem 1. Let $\mathcal{M}$ be a compact Riemannian manifold, without boundary or with Neumann boundary conditions. The solutions to the Helmholtz equation satisfy the following:
(i) The eigenvalues are real, non-negative, and tend to infinity,
$$0 = \lambda_1 \le \lambda_2 \le \lambda_3 \le \dots \to \infty.$$
(ii) There exists a complete orthonormal basis for $L^2(\mathcal{M})$ of real eigenfunctions $\{f_k\}$

Theorem 2. Let $\mathcal{M}_1, \mathcal{M}_2$ be compact Riemannian manifolds and let $\mathcal{M} = \mathcal{M}_1 \times \mathcal{M}_2$. The eigenfunctions of $\mathcal{M}_i$ are $\{f_k^{(i)}\}$ and the corresponding eigenvalues are $\{\lambda_k^{(i)}\}$. Let $\pi^{(i)} : \mathcal{M} \to \mathcal{M}_i$ denote the canonical projection of $\mathcal{M}$ to $\mathcal{M}_i$. Then the eigenfunctions of $\mathcal{M}$ are given by the pointwise products:
$$f_{m,n} = (f_m^{(1)} \circ \pi^{(1)})(f_n^{(2)} \circ \pi^{(2)})$$
The corresponding eigenvalues are the sums:
$$\lambda_{m,n} = \lambda_m^{(1)} + \lambda_n^{(2)}$$

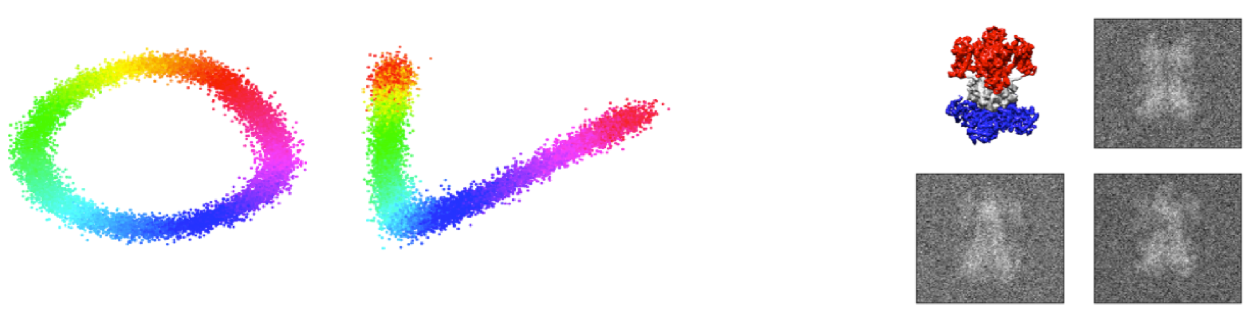

Figure 2. **Visualizing manifold Disentanglement**: (right) Images of a simulated molecule taken from a single orientation. This molecule exhibits a rotary motion of the top and a linear stretching motion of the bottom. (left) The disentangled representation of the two independent motion.

Laplacian eigenfunctions

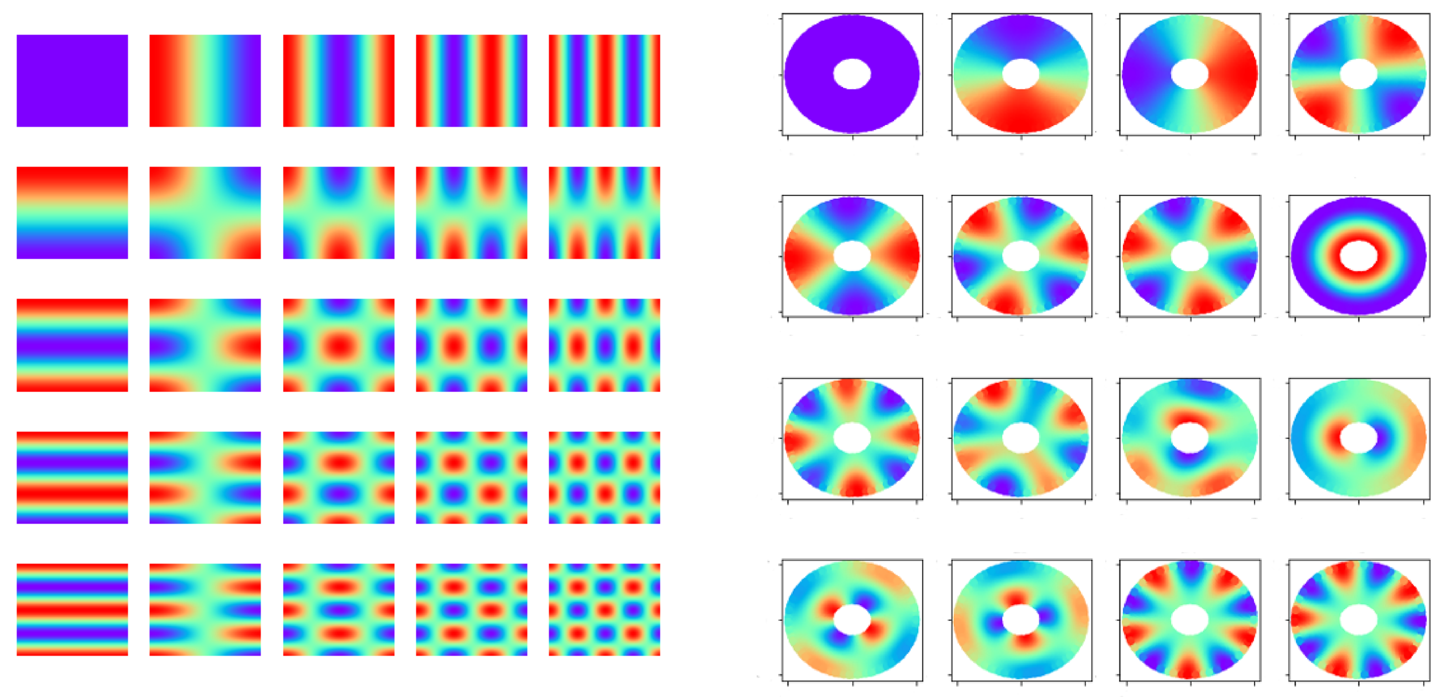

Figure 3. Eigenvectors of the Laplace operator on a rectangular domain and on an annulus.

Disentanglement of product manifold

```
Algorithm 1: Identification of individual factors
Data: Non-trivial eigenvectors {φ_1, ..., φ_N} of
      L_rw, sorted increasingly by their
      log-transformed eigenvalues λ_1, ..., λ_N.
Result: List of triplets (i, j, k) where φ_k ≈ φ_i φ_j
        and the corresponding similarity scores.
for k ← 1...N do
    maxS ← 0;
    for i, j < k do
        if |λ_i + λ_j − λ_k| < δ and
        S(φ_k, φ_i φ_j) > maxS then
            maxS ← S(φ_k, φ_i φ_j);
            i_max ← i;
            j_max ← j;
        end
    end
    if maxS > γ then
        add (i_max, j_max, k) to triplets;
    end
end
```

Figure 4. The algorithm presented in [4]

Current work

Today, our main focus is to extend [4], to disentangle group actions from other motion. Specifically, separating the rotation groups $SO(2)$ and $SO(3)$ from the intrinsic motion of the body.
More formally, Given a dataset $\{x_1, \dots, x_n\} \subseteq \mathbf{R}^D$ of a volume/image with an intrinsic motion that the group G acts on, and the Diffusion map $\Phi : \mathbf{R}^D \to \mathbf{R}^d$, Our goal is to find two maps $\Psi_1 : \mathbf{R}^D \to \mathbf{R}^{l_1}$, $\Psi_2 : \mathbf{R}^D \to \mathbf{R}^{l_2}$ s.t. Ψ_1 captures the structure of G and Ψ_1 captures the structure of $\mathcal{M}/G$ - the intrinsic motion component that is perpendicular to G.

Challenges

- Even in the simplest example of the annulus ($SO(2)$ acting on I), the Laplacian eigenfunctions do not include all individual factors of $\mathcal{M}_1$ and $\mathcal{M}_2$. So the triplet-finding algorithm in Figure 4 is insufficient.
- In non-commutative groups, such as $SO(3)$, we can't use classical Fourier theory.

References

[1] Mikhail Belkin and Partha Niyogi. Laplacian eigenmaps for dimensionality reduction and data representation. *Neural computation*, 2003.

[2] Ali Dashti et. al. Trajectories of the ribosome as a brownian nanomachine. *Proceedings of the National Academy of Sciences*, 2014.

[3] Boaz Nadler, Stephane Lafon, Ronald R. Coifman, and Ioannis G. Kevrekidis. Diffusion maps, spectral clustering and reaction coordinates of dynamical systems. *Applied and Computational Harmonic Analysis*, 2005.

[4] Sharon Zhang, Amit Moscovich, and Amit Singer. Product manifold learning. *AISTATS*, 2020.

הקרן הלאומית למדע
المؤسسة الإسرائيلية للعلوم
Israel Science Foundation