Functional correspondence by matrix completion

A. Kovnatsky[1], M. M. Bronstein[1], Xavier Bresson[2], and Pierre Vandergheynst[2]
[1]Faculty of Informatics, Università della Svizzera italiana (USI). [2]Signal Processing Lab, École Polytechnique Fédérale de Lausanne (EPFL).

We deal with the problem of finding dense intrinsic correspondence between manifolds in the functional formulation ([4], Fig. 1). We propose treating functional correspondence as *geometric matrix completion*. We show that our method compares favorably to state-of-the-art methods for non-rigid shape correspondence on the challenging Princeton benchmark [2]. The advantage of our method is especially pronounced in the *scarce data* settings. The proposed method is completely generic and can be applied to any manifolds and high-dimensional geometric data.

Functional maps Let X and Y denote two manifolds sampled at n and m points, respectively. Ovsjanikov et al. [4] propose to model the *functional correspondence* between the spaces $L^2(X)$ and $L^2(Y)$ as the $m \times n$ matrix $\mathbf{T}_0$, which maps a function $\mathbf{f} \in L^2(X)$ into $\mathbf{T}_0\mathbf{f} = \mathbf{g} \in L^2(Y)$. Traditional point-wise correspondence is a particular case of this model wherein $\mathbf{T}_0$ maps delta functions into delta functions.

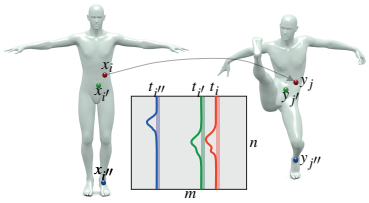

Figure 1: Illustration of functional correspondence as matrix completion. Column $\mathbf{t}_i = \mathbf{T}\delta_i$ of matrix $\mathbf{T}$ is the functional map of a delta at $\mathbf{x}_i$. Geometric structure is imposed on the columns to ensure smoothness, i.e. $\mathbf{t}_i$, $\mathbf{t}_{i'}$ corresponding to spatially close points $\mathbf{x}_i$, $\mathbf{x}_{i'}$ on X are similar.

Functional maps as matrix completion Kalofolias et al. [1] studied the problem of *matrix completion on graphs*, where the rows and columns of the matrix representing the correspondence have underlying geometric structure. They show that adding geometric structure to the standard matrix completion problem improves recovery results.

We use the same philosophy to formulate the problem of finding a functional map as matrix completion, whose rows and columns are interpreted as functions on the respective manifolds X and Y. For this purpose, we consider the matrix $\mathbf{T}$ as a collection of columns $\mathbf{T} = (\mathbf{t}_1, \dots, \mathbf{t}_n)$ or rows $\mathbf{T} = (\mathbf{t}^{1\top}, \dots, \mathbf{t}^{m\top})^\top$, where $\mathbf{t}_i$ and $\mathbf{t}^j$ denote the ith column and jth row of $\mathbf{T}$, respectively. The column $\mathbf{t}_i = \mathbf{T}\delta_i$ is the function on Y corresponding to a delta located at point $\mathbf{x}_i$ on X. Similarly, the row $\mathbf{t}^j = \delta_j^\top \mathbf{T}$ is the function on X corresponding to a delta located at point $\mathbf{y}_j$ on Y.

As in the classical matrix completion problem, we aim at recovering the unknown correspondence matrix $\mathbf{T}_0$ from a few observations of the form $\mathbf{T}_0\mathbf{F} = \mathbf{G}$, where $\mathbf{F} = (\mathbf{f}_1, \dots, \mathbf{f}_q)$ and $\mathbf{G} = (\mathbf{g}_1, \dots, \mathbf{g}_q)$ are q corresponding functions on X, Y, respectively. Furthermore, matrix $\mathbf{T} \approx \mathbf{T}_0$ should explain the data in a "simple" way, in the sense discussed in the following. Our problem comprises the following terms:

Data term The correspondence should respect the data, which is achieved by minimizing $\|\mathbf{TF} - \mathbf{G}\|_\mathrm{F}$.

Smoothness The correspondence must be *smooth*, in the sense that if $x_i, x_{i'}$ are two close points on X, then the respective corresponding functions are similar, i.e., $\mathbf{t}_i \approx \mathbf{t}_{i'}$ (see Figure 1). Similarly, for close points $y_j, y_{j'}$ on Y, the rows $\mathbf{t}^j \approx \mathbf{t}^{j'}$. Smoothness is achieved by minimizing the *row* and *column Dirichlet energy* $\mathrm{tr}(\mathbf{T}^\top \mathbf{L}_Y \mathbf{T}) + \mathrm{tr}(\mathbf{T}\mathbf{L}_X\mathbf{T}^\top)$.

Localization The correspondence is localized using the L_1-penalty $\|\mathbf{T}\|_1$, which, in combination with smoothness, results in a few non-zero elements that are close in space.

"Simplicity" By simplicity, we mean here that the correspondence matrix is 'explained' using a small number of degrees of freedom. The following models are commonly used in matrix completion and recommendation systems literature.

Low rank. A popular model is to find the matrix with smallest rank($\mathbf{T}$). However, this minimization is known to be NP-hard, and the nuclear norm $\|\mathbf{T}\|_*$ is typically used as a convex proxy, leading to

$$\min_{\mathbf{T}} \|\mathbf{TF} - \mathbf{G}\|_\mathrm{F}^2 + \mu_1 \mathrm{tr}(\mathbf{T}\mathbf{L}_X\mathbf{T}^\top) + \mu_2 \mathrm{tr}(\mathbf{T}^\top \mathbf{L}_Y \mathbf{T}) + \mu_3 \|\mathbf{T}\|_1 + \mu_4 \|\mathbf{T}\|_*. \quad (1)$$

where $\mu_1, \mu_2, \mu_3 > 0$ are parameters determining the tradeoff between smoothness and localization.

Low norm. Srebro et al. [6] rewrite problem (1) using the decomposition $\mathbf{T} = \mathbf{U}\mathbf{V}^\top$ with $\mathbf{U}$ and $\mathbf{V}$ of size $m \times k$ and $n \times k$, respectively. Note that k can be arbitrarily large. The nuclear norm is written as $\|\mathbf{T}\|_* = \frac{1}{2}(\|\mathbf{U}\|_\mathrm{F}^2 + \|\mathbf{V}\|_\mathrm{F}^2)$. Unlike (1), this problem is non-convex w.r.t. both $\mathbf{U}$ and $\mathbf{V}$, however behaves well for large k [6].

Subspace parametrization Let $\mathbf{L}_X$ and $\mathbf{L}_Y$ be the Laplacians of X and Y, and let $\Phi_k = (\phi_1, \dots \phi_k)$ and $\Psi_k = (\psi_1, \dots \psi_k)$ be the respective truncated Laplacian eigenbases. The number of variables in the problem (1) depends on the number of samples m, n, which may result in scalability issues for large ($m, n \sim 10^6$) manifolds. To overcome this issue, we approximate the $n \times k$ and $m \times k$ factors $\mathbf{U}, \mathbf{V}$ in the truncated Laplacian eigenbases of X and Y using $k' \geq k$ first expansion terms, $\mathbf{U} \approx \Psi_{k'}\mathbf{A}$ and $\mathbf{V} \approx \Phi_{k'}\mathbf{B}$, where matrices $\mathbf{A}, \mathbf{B}$ of the expansion coefficients are of size $k' \times k$. This leads to a subspace version of our problem,

$$\min_{\mathbf{A},\mathbf{B}} \quad \|\mathbf{A}\mathbf{B}^\top \Phi_{k'}^\top \mathbf{F} - \Psi_{k'}^\top \mathbf{G}\|_\mathrm{F}^2 + \mu_1 \mathrm{tr}(\mathbf{A}\mathbf{B}^\top \Lambda_{X,k'} \mathbf{B}\mathbf{A}^\top) + \quad (2)$$
$$\mu_2 \mathrm{tr}(\mathbf{B}\mathbf{A}^\top \Lambda_{Y,k'} \mathbf{A}\mathbf{B}^\top) + \mu_3 \|\Psi_{k'}\mathbf{A}\mathbf{B}^\top \Phi_{k'}^\top\|_1 + \frac{\mu_4}{2}(\|\Psi_{k'}\mathbf{A}\|_\mathrm{F}^2 + \|\Phi_{k'}\mathbf{B}\|_\mathrm{F}^2)$$

where we used the invariance of the Frobenius norm to orthogonal transformations and the fact that $\Phi_{k'}^\top \mathbf{L}_X \Phi_{k'} = \Lambda_{X,k'}$, $\Psi_{k'}^\top \mathbf{L}_Y \Psi_{k'} = \Lambda_{Y,k'}$ to simplify the expressions ($\Lambda_{\cdot,k}$ denotes a diagonal matrix of k respective eigenvalues).

Note that now the number of variables $2kk'$ is *independent of the number of samples*. We emphasize that k', k can be arbitrarily large and are dictated only by complexity considerations and not by the amount of data. This is one of the major advantages of our approach compared to [3, 4, 5].

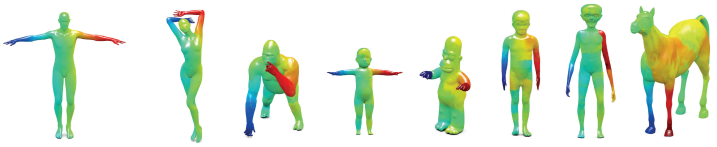

Figure 2: Examples of correspondence between non-isometric shapes. Leftmost shape is used as reference. Similar colors encode corresponding points.

[1] Vassilis Kalofolias, Xavier Bresson, Michael Bronstein, and Pierre Vandergheynst. Matrix completion on graphs. *arXiv:1408.1717*, 2014.

[2] Vladimir G Kim, Yaron Lipman, and Thomas Funkhouser. Blended intrinsic maps. *TOG*, 30(4):79, 2011.

[3] Artiom Kovnatsky, Michael M Bronstein, Alexander M Bronstein, Klaus Glashoff, and Ron Kimmel. Coupled quasi-harmonic bases. *Computer Graphics Forum*, 32:439–448, 2013.

[4] M. Ovsjanikov, M. Ben-Chen, J. Solomon, A. Butscher, and L. Guibas. Functional maps: A flexible representation of maps between shapes. *TOG*, 31(4), 2012.

[5] Jonathan Pokrass et al. Sparse modeling of intrinsic correspondences. *Computer Graphics Forum*, 32:459–468, 2013.

[6] Nathan Srebro, Jason Rennie, and Tommi S Jaakkola. Maximum-margin matrix factorization. In *Proc. NIPS*, 2004.

This is an extended abstract. The full paper is available at the Computer Vision Foundation webpage.