

Discrete Optimization for Shape Matching Supplementary Materials

Jing Ren¹, Simone Melzi², Peter Wonka¹, Maks Ovsjanikov³

¹KAUST ²Sapienza University of Rome ³LIX, École Polytechnique, CNRS

In the supplementary materials, we focus on the analysis of the Effective Functional Map Refinement (EFMR) method that we proposed in Sec. 5.3. in the main paper.

Overview We first give the full details of the EFMR algorithm in Sec. 1. To analyze the new energy and the EFMR method, we perform tests on some synthetic data in Sec. 2. Specifically, we would like to study three problems: (1) how well different energies characterize the desirable maps? E.g., how many local minima there are and are these local minima all corresponding to desirable maps? (2) For a map refinement method, where does it converge to starting from different initializations? E.g., what are the termination points of different refinement methods. (3) Does the termination point of a refinement method happen to be a local minima of the optimized energy? We investigated ICP [OBCS*12], ZoomOut [MRR*19], and our EFMR regarding these three problems. In Sec. 3 we further discuss the relationship between EFMR and other methods. We also perform an ablation study in Sec. 4 to verify that each component of the new energy matters.

1. Effective Functional Map Refinement

In this section, we discuss how to apply our discrete solver to the new energy that combines the orthogonality, Laplacian commutativity, and bijectivity on functional maps from both directions of a shape pair, which is defined as:

$$E_{\text{new}}(C_{12}, C_{21}) = w_1 E_{\text{bi}}(C_{12}, C_{21}) + w_2 E_{\text{ortho}}(C_{12}) + w_2 E_{\text{ortho}}(C_{21}) + w_3 E_{\text{lap}}(C_{12}) + w_3 E_{\text{lap}}(C_{21}) \quad (1)$$

Our goal is to find a pair of *proper* functional maps from both directions that minimize the above energy:

$$\min_{C_{12} \in \mathcal{P}_{12}, C_{21} \in \mathcal{P}_{21}} E_{\text{new}}(C_{12}, C_{21}) \quad (2)$$

ALGORITHM 1: Effective Functional Map Refinement (EFMR)

Goal : $\min_{C_{12} \in \mathcal{P}_{12}, C_{21} \in \mathcal{P}_{21}} E_{\text{new}}(C_{12}, C_{21})$
Output: Proper functional maps C_{12}, C_{21} ; Pointwise maps T_{12}, T_{21}
Parameters: weights W_i and α, N , some stopping criterion
while *Stopping Criterion Not Met* **do**
 $\Phi_1 = \Phi_{S_1}^{(k)}, \Phi_2 = \Phi_{S_2}^{(k)}, \Delta_1 = \Delta_{S_1}^{(k)}, \Delta_2 = \Delta_{S_2}^{(k)}$
for $iter = 1:N$ **do**
 Construct X_1, X_2, Y_1, Y_2 :
 $X_1 = \Phi_1 [w_1 C_{21}, (w_2 + \alpha) C_{12}^T, w_3 \Delta_1]$
 $X_2 = [w_1 \Phi_2, (w_2 + \alpha) \Phi_2, w_3 \Phi_2 \Delta_2 C_{12}]$
 $Y_1 = [w_1 \Phi_1, (w_2 + \alpha) \Phi_1, w_3 \Phi_1 \Delta_1 C_{21}]$
 $Y_2 = \Phi_2 [w_1 C_{12}, (w_2 + \alpha) C_{21}^T, w_3 \Delta_2]$
 Compute pointwise maps:
 $\Pi_{12} = \text{knn}(Y_2, Y_1)$
 $\Pi_{21} = \text{knn}(X_1, X_2)$
 Compute proper functional maps:
 $C_{12} = \Phi_{S_2}^\dagger \Pi_{21}^* \Phi_{S_1}$
 $C_{21} = \Phi_{S_1}^\dagger \Pi_{12}^* \Phi_{S_2}$
end
 $k \leftarrow k + 1$
end

We can construct the relaxed version of our new energy in the following (as discussed in Sec. 4 and Sec. 5):

$$E_{\text{new}}^{\text{relax}} = w_1 \|\Phi_{S_2}^\dagger \Pi_{21} \Phi_{S_1} C_{21} - Id\|_F^2 + w_1 \|\Phi_{S_1}^\dagger \Pi_{12} \Phi_{S_2} C_{12} - Id\|_F^2 + w_2 \|\Phi_{S_2}^\dagger \Pi_{21} \Phi_{S_1} C_{12}^T - Id\|_F^2 + w_2 \|\Phi_{S_1}^\dagger \Pi_{12} \Phi_{S_2} C_{21}^T - Id\|_F^2 + w_3 \|\Phi_{S_2}^\dagger \Pi_{21} \Phi_{S_1} \Delta_{S_1} - \Delta_{S_2} C_{12}\|_F^2 + w_3 \|\Phi_{S_1}^\dagger \Pi_{12} \Phi_{S_2} \Delta_{S_2} - \Delta_{S_1} C_{21}\|_F^2 + \alpha \|\Phi_{S_2}^\dagger \Pi_{21} \Phi_{S_1} C_{12}^T - Id\|_F^2 + \alpha \|\Phi_{S_1}^\dagger \Pi_{12} \Phi_{S_2} C_{21}^T - Id\|_F^2 \quad (3)$$

where we plugged the hard constraints into the original energy and add two soft regularizers with weight α . We then have the updating rules for our discrete solver: (1) update the pointwise maps by $\Pi_{12} = \arg \min_{\Pi_{12}} E_{\text{new}}^{\text{relax}}(\Pi_{12} | \Pi_{21}, C_{12}, C_{21})$ and $\Pi_{21} = \arg \min_{\Pi_{21}} E_{\text{new}}^{\text{relax}}(\Pi_{21} | \Pi_{12}, C_{12}, C_{21})$; (2) compute the proper functional maps by $C_{12} = \Phi_{S_2}^\dagger \Pi_{21} \Phi_{S_1}$ and $C_{21} = \Phi_{S_1}^\dagger \Pi_{12} \Phi_{S_2}$.

See Algorithm 1 for full details of applying the discrete solver to



Figure 1: Dense correspondences on SHREC’19 using our method from random initialization.

optimize the new energy we proposed with the progressive upsampling technique. In Sec. 6.3 of the main paper, we quantitatively show that our EFMR method achieves the state-of-the-art accuracy on SHREC’19 benchmark. Fig. 1 and Fig. 2 show two qualitative examples of the strong convergence of our EFMR method that can handle trivial and even random initializations.

In the following, we demonstrate the advantages of our new energy and our EFMR algorithm on a synthetic dataset, where we can explore the complete proper functional map search space to find the local minima of different energies, and to find the converged/terminating points of different refinement methods, including ICP, ZoomOut and ours.

2. Quantitative Justification on Synthetic Data

To illustrate the advantages of our new energy Eq. (1) and our algorithm EFMR, we consider shapes with a small number of vertices $n = 3, 4, 5$, for which it is possible to explore the *full* search space of pointwise maps by enumerating all possible pointwise maps in both directions. We then study the effectiveness of different energies in distinguishing good maps within this space and study the convergence power of different refinement methods.

Specifically, Fig. 3 shows an example of a pair of triangle shapes. In this case we know that there are $3^3 = 27$ possible pointwise maps across the two shapes in each direction, and the *product space* of two directions contains $27 \times 27 = 792$ total elements.

To evaluate how well different energies characterize the search space, we define the topology on the map space as illustrated in Fig. 3. Namely two maps are connected if they differ by one entry. We then extend this to the product space by connecting two map pairs if either of the maps in the pair are connected. This allows us to define the notion of a *local minimum* in the map search space w.r.t. a given energy: i.e., a map pair is a local minimum if and only if its energy value is smaller than that of its neighbors. In this test, we compare our new energy to the ICP energy and the ZOOMOUT energy on the complete map product space. The number of local minima of different energies is reported in Table. 1. Note that our new energy has significantly fewer local minima than the other two energies.

We also study the convergence of different methods starting from different initial maps. Specifically, for a shape pair with $n(n \leq 10)$ vertices, we run ICP at dimension n (i.e., the maximum dimension), and run ZOOMOUT and our method from dimension 2 to n on the tested initializations in the product space. When $n > 10$, we run the algorithms at (or up to) 30 dimensions. We then measure the convergence of different methods as the number of unique maps




Figure 2: Refining a trivial map with all vertices from the source mapped to a single vertex on the target. We compare our approach to ICP at a fixed dimension and to Zoomout [MRR*19], showing the initial and the refined functional and pointwise maps (via color transfer).

after refinement (see Table 1). We can see that our method has a strong convergence power so that out of more than 500 random initialization we obtain 2-4 unique maps, while as a comparison, ZOOMOUT converged to 25 - 2K unique maps.

For the shapes with more than five vertices, the size of the map product space quickly become intractable. Therefore, we only test on a relatively small set of initial maps and check the convergence of different methods. For example, Fig. 4 shows an example of two heart-shape meshes with 6 vertices. Our method successfully converged to the two global optima of the geodesic distortion, while ZOOMOUT fails to return a good map. Table 1 also includes the results of human shape pairs with 10, 100, 1000, 5000 vertices, where we tested different refinement on 500 randomly generated initial maps. Again, our method converged to significantly fewer maps with better quality than the other two methods, ICP and ZOOMOUT.

3. Relation to previous approaches

Our EFMR method is related to several previous approaches. Most importantly the Iterative Closest Point (ICP) refinement, commonly used in functional maps literature, e.g., [OBBS*12, PBB*13]. This approach is equivalent to the orthogonality-based map update scheme in fixed dimension, with an additional singular value projection step. Our EFMR method can also be regarded as an extension to the recent ZoomOut [MRR*19] and the bidirectional ZoomOut [RMOW20] method, that combines the orthogonality, bijectivity, and Laplacian commutativity using the progressive upsampling technique. We also note that our use of variable splitting is related to the work of [PBB*13], where the authors proposed to estimate a *permutation of input descriptors*, while promoting functional map diagonality. In contrast, we directly recover dense pointwise correspondences without any input descriptors. Finally, our proper functional map optimization is related to promoting pointwise maps, as in [NO17]. However, rather than promoting this property in the functional domain, we link pointwise and functional map estimation and, again, avoid the use of descriptor functions. In summary our method can be thought of as a strong generalization of both ICP and ZoomOut. Our generalization plays a crucial role in practice, where, as we demonstrate below, it is able to recover good maps without any initialization, which is not possible for either ICP and ZoomOut.

At a high level, our method is also related to other recent refinement techniques, such as PMF [VLR*17, VLB*17], BCICP [RPWO18], and Reversible Harmonic Maps [ESBC19]. Unlike these methods, however, we show how well-established functional

Table 1: Synthetic evaluation. For the shapes with small number of vertices ($n = 3, 4, 5$), we can enumerate all possible pointwise maps to study how well different energies can characterize the complete search space and how well different methods can refine the pointwise maps from arbitrary initialization. We report the number of local minima of different energies and the number of converged points/maps in the search space of different methods. For the case ($n > 5$) where the complete search space is intractable, we test different methods on 500 randomly selected initializations. As a result, our new energy has much less number of local minima than the ICP and ZOOMOUT energy. At the same time, our new method can converge to much less number of maps. Specifically, for $n \leq 6$, the converged maps obtained by our method are indeed the global optima of the geodesic distortion measure, i.e., the direct map and the symmetric map.

# vertex n	search space size n^n	# tested Ini in the product space	# local minima of the energy			# converged maps of the method			# converged maps are local minima		
			ICP	ZOOMOUT	Ours	ICP	ZOOMOUT	Ours	ICP	ZOOMOUT	Ours
3	27	$27^2 = 729$	117	118	29	520	25	2	96	12	2
4	256	$256^2 \approx 65K$	9156	7537	853	12240	70	2	2153	7	2
5	3125	$3125^2 \approx 9.7M$	810K	1235K	29K	425K	2160	3	62K	1005	3
6	46656	46656	-	-	-	7251	25	2	-	-	-
10	10^{10}	500	-	-	-	500	77	2	-	-	-
100	100^{100}	500	-	-	-	500	404	2	-	-	-
1000	1000^{1000}	500	-	-	-	500	497	4	-	-	-
5000	5000^{5000}	500	-	-	-	500	498	4	-	-	-

Table 2: Ablation Study. We tested 50 FAUST shape pairs. For each shape pair, we randomly generate 100 initial pointwise maps, we then apply our method with different parameter choices (w_1, w_2, w_3, N), and measure how many unique maps are obtained after refinement. We also measure the accuracy and bijectivity of the converged maps and report the average over all tested shape pairs.

Metrics \ Methods	Rand Ini	ICP	ZM	N	Ours																					
					w_1	w_2	w_3	1	1	1	1	1	1	1	1	1	1	10	10	10	10	10	10	10	10	10
# Converged Maps	100	100	100	-	76	66	66	66	67	12	12	32	25	26	28	9	7	6	7	7	2	2	3	3	3	3
Accuracy ($\times 10^{-3}$)	771	789	648	-	637	275	225	238	236	647	638	195	121	118	133	445	113	115	112	105	552	505	58.7	50.3	50.7	50.6
Bijectivity ($\times 10^{-3}$)	890	884	863	-	863	670	629	628	616	46.1	45.6	36.0	34.2	34.4	34.6	716	475	502	517	500	43.0	42.4	34.3	32.5	33.0	33.0

map energies can be used to optimize dense pointwise correspondences directly. As a result, our method is significantly more robust, is more scalable (as it does not require maintaining large dense matrices).

4. Ablation Study

Here we verify by an ablation study that every component of our new energy in Eq. (1) is effective. We have four main parameters in the energy and the algorithm, the number of inner loops N , and the weights w_1, w_2, w_3 for different terms, namely the bijectivity, orthogonality, and Laplacian Commutativity. To avoid arbitrary selections of these four parameters, we conduct our ablation study in a more principled way. First of all, we always include the orthogonality term and set $w_2 = 1$, given that [MRR*19] has already demonstrated the effectiveness of the orthogonality term. Second, we only consider $w_1 = 0$ or $w_1 = 1$. The reason is that the orthogonality and the bijectivity terms both impose some sort of orthonormal constraint on the functional maps. Without any prior knowledge, it is difficult to set different weights for these two terms. Finally, for a fixed choice of w_1, w_2 , we test different choices of w_3 , the weight of the Laplacian Commutativity term. Moreover, for a fixed choice of the weights, we also test $N = 1$ or $N = 10$ to verify if the inner loop can help improve the robustness of our algorithm. Note that,

ZOOMOUT can be regarded as a special case of our EFMR with the parameter setting of $N = 1, w_1 = w_3 = 0, w_2 = 0$.

We then test our algorithm with different parameter choices on 50 pairs of FAUST human shapes [BRLB14]. We randomly generated 100 different initializations for each shape pair. We report the number of converged maps and the quality including the map accuracy and the bijectivity of the converged maps averaged across all the tested pairs in Table 2. We can see that we need to combine all the three terms in our energy to achieve the best performance. For any selection of parameters, our method outperforms the competitors highlighting that the novel discrete optimization injected in the iterative process improves the quality of the maps independently from the parameters. We stress that the discrete optimization was not considered at all in the previous method and has to be considered as the main theoretical and practical contribution of the proposed method.

References

[BRLB14] BOGO F., ROMERO J., LOPER M., BLACK M. J.: FAUST: Dataset and evaluation for 3D mesh registration. In *Proc. CVPR* (Columbus, Ohio, 2014), IEEE, pp. 3794–3801. 3

[ESBC19] EZUZ D., SOLOMON J., BEN-CHEN M.: Reversible harmonic maps between discrete surfaces. *ACM Trans. Graph.* 38, 2 (2019), 15:1–15:12. 2

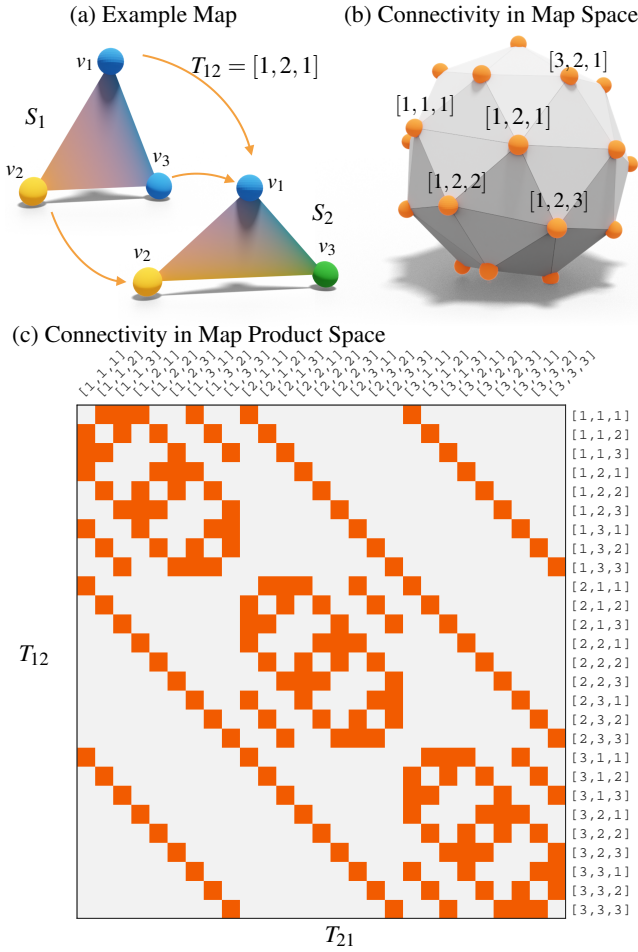


Figure 3: Illustration of the Map Product Space. (a) Shows a pair of triangles and an example map T_{12} , where the first and the third vertex on S_1 are both mapped to the first vertex on S_2 , and the second vertex on S_1 is mapped to the second vertex on S_2 . The color coding on the vertices visualizes the corresponding map. (b) We define the connectivity between the pointwise maps in the following way: T^i is connected to T^j if and only if T^i can become T^j if we only change one entry of T^j . The connectivity of the maps between these two triangles can be visualized as a shape in (b), where each node represents a single map and is connected to 6 other nodes/maps. (c) When we are dealing with the map product space (T_{12}, T_{21}) , we can similarly define the connectivity between two pairs of bidirectional maps in this map product space: the map pair (T_{12}^i, T_{21}^j) is connected to the map pair (T_{12}^p, T_{21}^q) if and only if T_{12}^i is connected to T_{12}^p or T_{21}^j is connected to T_{21}^q . The connectivity of the map produce space on this pair of triangle can be visualized as matrix, where each entry of the matrix represents a map pair. Each entry is connected to the other red entries in the same row or column.

[MRR*19] MELZI S., REN J., RODOLÀ E., SHARMA A., WONKA P., OVSJANIKOV M.: Zoomout: Spectral upsampling for efficient shape cor-

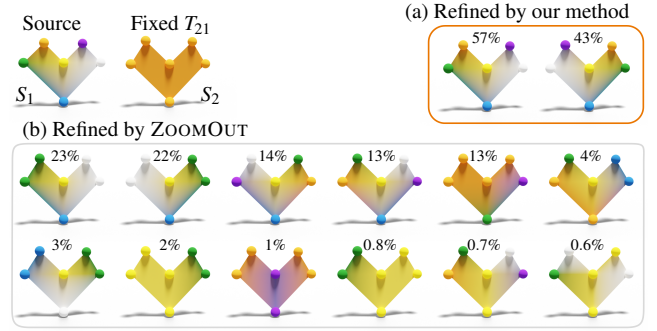


Figure 4: Test on two heart-shaped meshes with 6 vertices. For the point-wise map T_{12} from shape S_1 to S_2 , there are 6^6 different maps. Therefore, we enumerate all of them and apply ZOOMOUT to refine these maps, where the refined maps are converged to 25 different maps. Here we show the first 12 maps with the highest occurrence in (b). At the same time, we also use our method to refine those T_{12} s as well while keeping the the map from the other direction T_{21} as a trivial map, as shown above. As a comparison, our method converged to two maps, including the correct direct map and the symmetric map as shown in (a). Also notice that, the ZoomOut-refined maps never converged to the correct maps shown in (a).

respondence. *ACM Transactions on Graphics (TOG)* 38, 6 (Nov. 2019), 155:1–155:14. 1, 2, 3

[NO17] NOGNENG D., OVSJANIKOV M.: Informative descriptor preservation via commutativity for shape matching. *Computer Graphics Forum* 36, 2 (2017), 259–267. 2

[OBCS*12] OVSJANIKOV M., BEN-CHEN M., SOLOMON J., BUTSCHER A., GUIBAS L.: Functional maps: a flexible representation of maps between shapes. *ACM Transactions on Graphics (TOG)* 31, 4 (2012), 30:1–30:11. 1, 2

[PBB*13] POKRASS J., BRONSTEIN A. M., BRONSTEIN M. M., SPRECHMANN P., SAPIRO G.: Sparse modeling of intrinsic correspondences. *Computer Graphics Forum* 32, 24 (2013), 459–468. 2

[RMOW20] REN J., MELZI S., OVSJANIKOV M., WONKA P.: Maptree: Recovering multiple solutions in the space of maps. *ACM Trans. Graph.* 39, 6 (Nov. 2020). 2

[RPWO18] REN J., POULENARD A., WONKA P., OVSJANIKOV M.: Continuous and orientation-preserving correspondences via functional maps. *ACM Transactions on Graphics (TOG)* 37, 6 (2018). 2

[VLB*17] VESTNER M., LÄHNER Z., BOYARSKI A., LITANY O., SLOSSBERG R., REMEZ T., RODOLÀ E., BRONSTEIN A., BRONSTEIN M., KIMMEL R.: Efficient deformable shape correspondence via kernel matching. In *3D Vision (3DV), 2017 International Conference on* (2017), IEEE, pp. 517–526. 2

[VLR*17] VESTNER M., LITMAN R., RODOLÀ E., BRONSTEIN A., CREMERS D.: Product manifold filter: Non-rigid shape correspondence via kernel density estimation in the product space. In *Proc. CVPR* (2017), pp. 6681–6690. 2