Isometric Multi-Shape Matching —Supplementary Material—

Maolin Gao[1] Zorah Lähner[1,3] Johan Thunberg[2] Daniel Cremers[1] Florian Bernard[1]

[1] Technical University of Munich [2] Halmstad University [3] University of Siegen

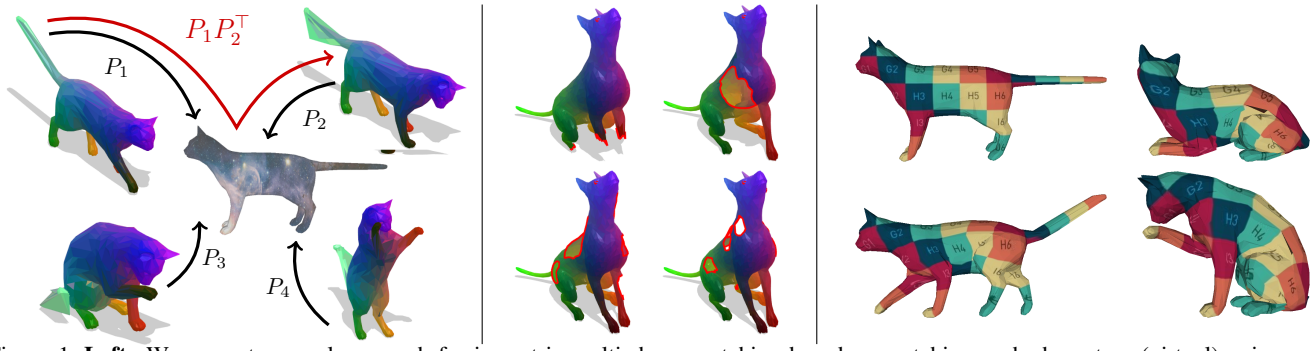

Figure 1. **Left:** We present a novel approach for isometric multi-shape matching based on matching each shape to a (virtual) universe shape (shown semi-transparent). Our formulation represents point-to-point correspondences between shapes i and j as the composition of the shape-to-universe permutation matrix P_i and the universe-to-shape permutation matrix P_j^T. By doing so, the pairwise matchings $P_{ij} = P_iP_j^\top$ are by construction cycle-consistent. **Middle:** Our formulation successfully solves isometric multi-matching of partial shapes. **Right:** Due to the cycle-consistency we can use our correspondences to faithfully transfer textures across a shape collection.

1. Theoretical Analysis (with Proofs)

Lemma 1 $\langle U_t^\top \mathbf{\Phi} Q_t, U_{t+1}^\top \mathbf{\Phi} Q_t\rangle \geq \langle U_t^\top \mathbf{\Phi} Q_t, U_t^\top \mathbf{\Phi} Q_t\rangle$ *holds for any* t.

Proof: According to (13), the function $\langle U_t^\top \mathbf{\Phi} Q_t, U \mathbf{\Phi} Q_t\rangle$ is maximised w.r.t. U over $\mathbb{P}$ for the choice $U = U_{t+1}$. Our claim follows immediately from this. □

Proposition 2 (Monotonicity of U-update) *The objective values cannot decrease through the U-update step (13), and* $\langle U_{t+1}^\top \mathbf{\Phi} Q_t, U_{t+1}^\top \mathbf{\Phi} Q_t\rangle \geq \langle U_t^\top \mathbf{\Phi} Q_t, U_t^\top \mathbf{\Phi} Q_t\rangle$ *holds.*

Proof: We prove the proposition by using Lemma 1. Recalling that $Z = \mathbf{\Phi} Q_t Q_t^\top \mathbf{\Phi}^\top$, we can see that

$$
\begin{aligned}
0 &\leq \|U_{t+1}^\top \mathbf{\Phi} Q_t - U_t^\top \mathbf{\Phi} Q_t\|_F^2 && \text{(S1)}\\
&= \langle U_{t+1}^\top \mathbf{\Phi} Q_t, U_{t+1}^\top \mathbf{\Phi} Q_t\rangle - 2\langle U_{t+1}^\top \mathbf{\Phi} Q_t, U_t^\top \mathbf{\Phi} Q_t\rangle \\
&\quad + \langle U_t^\top \mathbf{\Phi} Q_t, U_t^\top \mathbf{\Phi} Q_t\rangle && \text{(S2)}\\
&= \langle U_{t+1}^\top, U_{t+1}^\top Z\rangle - 2\langle U_{t+1}^\top, U_t^\top Z\rangle + \langle U_t^\top, U_t Z\rangle\,. && \text{(S3)}
\end{aligned}
$$

From Lemma 1 and using the symmetry of Z, we know that $\langle U_{t+1}^\top, U_t^\top Z\rangle \geq \langle U_t^\top, U_t^\top Z\rangle$. By transitivity this leads to

$$0 \leq \langle U_{t+1}^\top, U_{t+1}^\top Z\rangle - 2\langle U_{t+1}^\top, U_t^\top Z\rangle + \langle U_{t+1}^\top, U_t^\top Z\rangle\,,$$

so that

$$\langle U_{t+1}^\top, U_t^\top Z\rangle \leq \langle U_{t+1}^\top, U_{t+1}^\top Z\rangle\,.$$

□

Lemma 3 *In each iteration* t, $\langle U_{t+1}^\top \mathbf{\Phi} Q_t, U_{t+1}^\top \mathbf{\Phi} Q_{t+1}\rangle \geq \langle U_{t+1}^\top \mathbf{\Phi} Q_t, U_{t+1}^\top \mathbf{\Phi} Q_t\rangle$ *holds.*

Proof: Analogously to the proof of Lemma 1, and according to (15), the choice $Q = Q_{t+1}$ is the element maximising the expression $\langle U_{t+1}^\top \mathbf{\Phi} Q_t, U_{t+1}^\top \mathbf{\Phi} Q\rangle$ w.r.t. Q over $\mathbb{O}$. □

Proposition 4 (Monotonicity of Q-update) *The objective values cannot decrease through the Q-update (15), and* $\langle U_{t+1}^\top \mathbf{\Phi} Q_{t+1}, U_{t+1}^\top \mathbf{\Phi} Q_{t+1}\rangle \geq \langle U_{t+1}^\top \mathbf{\Phi} Q_t, U_{t+1}^\top \mathbf{\Phi} Q_t\rangle$ *holds.*

Proof: The proof is analogous to Prop. 2. For $\overline{Z} = \mathbf{\Phi}^\top U_{t+1} U_{t+1}^\top \mathbf{\Phi}$ we observe that

$$0 \leq \|U_{t+1}^\top \mathbf{\Phi} Q_t - U_{t+1}^\top \mathbf{\Phi} Q_{t+1}\|_F^2 \tag{S4}$$

$$= \langle Q_t, \overline{Z} Q_t \rangle - 2\langle Q_t, \overline{Z} Q_{t+1} \rangle + \langle Q_{t+1}, \overline{Z} Q_{t+1} \rangle . \tag{S5}$$

From Lemma 3 we have $\langle Q_t^\top, \overline{Z} Q_{t+1} \rangle \geq \langle Q_t^\top, \overline{Z} Q_t \rangle$, so that our claim follows by transitivity. □

By combining these properties we obtain the following immediate result regarding Algorithm 1:

Theorem 5 (Convergence)
The sequence $(f(U_t, Q_t))_{t=1,2,\dots}$ is convergent and Algorithm 1 terminates in finite time.

Proof: For any t we have $U_t \in \mathbb{P}$ and $Q_t \in \mathbb{O}$. Hence, the value of $f(U_t, Q_t)$ is bounded from above (both $\mathbb{P}$ and $\mathbb{O}$ are compact sets). Combined with the monotonicity of the U-update (Prop. 2) and Q-update (Prop. 4), this shows that the sequence $(f(U_t, Q_t))_{t=1,2\dots}$ converges. □

2. Details on Experimental Setup

In this section we clarify further details regarding the experimental setup.

Synchronisation of ZoomOut results. Running ZoomOut produces the pairwise correspondences $\{P_{ij}\}$, and the pairwise functional maps $\{\mathcal{C}_{ij}\}$ between all pairs of shapes. In general, the pairwise correspondences and the pairwise functional maps are not cycle-consistent. In order to obtain cycle-consistent shape-to-universe representations we apply synchronisation, as we explain next.

For isometric shapes, the spectra of the Laplace-Beltrami operator are the same for all shapes. Moreover, the pairwise functional maps have a band-diagonal structure, where the band-width depends on the largest multiplicity of the spectra, see [6] for details. Hence, we first set all elements of $\mathcal{C}_{ij}$ to 0 that are outside the diagonal band of radius $r = 6$, i.e. $\mathcal{C}_{ij}(s,t) = 0$ whenever $|s-t| > r$. Subsequently, we project the "band-filtered" $\mathcal{C}_{ij}$ onto $\mathbb{O}_b$ using singular value decomposition (SVD), since isometric shapes must lead to orthogonal functional maps. Eventually, we use orthogonal transformation synchronisation [11, 2] in order to obtain the shape-to-universe functional maps $\{\mathcal{C}_i\}$, which we stack into the block-matrix Q.

In order to obtain U, we first represent all LBO eigenfunctions Φ_i in terms of the universe, i.e. $\Phi_i \mathcal{C}_i$, and then stack them all into the matrix

$$\mathbf{\Psi} = \begin{bmatrix} \Phi_1 \mathcal{C}_1 \\ \vdots \\ \Phi_k \mathcal{C}_k \end{bmatrix} . \tag{S6}$$

Eventually, we obtain the shape-to-universe matching matrix $U \in \mathbb{P}$ by performing a constrained clustering, where the features used for clustering are the inner products between the eigenfunctions in the universe representation. This means that the rows of the matrix $\mathbf{\Psi}\mathbf{\Psi}^\top$ are used as features for clustering. This is motivated by the (constrained) clustering interpretation of partial permutation synchronisation, see e.g. [13, 3]. For performing the clustering, we first apply the Successive Block Rotation Algorithm (SBRA) [3], followed by projecting the result onto the set $\mathbb{P}$. Further details can be found in [3].

Symmetries. Bringing symmetric shapes into correspondence is well-known to be a challenging problem [12]. To avoid symmetric flips it is common practice to incorporate an additional symmetry descriptor into shape matching formulations, as for example done in [5]. We follow this path, and make use of a symmetry descriptor for finding the ZoomOut initialisation. We emphasise that the symmetry descriptor is not used after the multi-shape matching methods have been initialised.

Parameters. For the experiments that consider full shapes (on the TOSCA, FAUST and SCAPE datasets), there exists a bijection between all shapes within a category, hence $m_i = m_j$ for all i, j. Thus, we set the universe size d to the number of vertices present in each shape, i.e $d = m_i$. In all experiments, we fix the relative objective improvement to machine precision, i.e. $\epsilon \approx 2.2 \cdot 10^{-16}$.

3. Multi-Matching of Partial Shapes

This section will provide more details on the experiments of Section 5.2 in the main paper. Strictly speaking, partial shapes do not fulfil the isometry assumption due to missing parts that affect geodesic distances. However, in the case of finding a matching between a full shape and shape with holes, both of the same class, there is a close relationship (see Fig. 1). [9] discusses how spectral properties change in this case, and the necessary adjustment of our pipeline is based on this theory. Finding correspondences for partial-to-partial cases is a much more challenging and open problem, and due to a lack of robust initialisation methods, as a proof-of-concept we show results on small datasets with only minor deformations. See Fig. 2 for qualitative results. Note that the cycle error of our approach is guaranteed to be zero even in the partial multi-matching case, and the runtime is similar as in the full matching case.

Problem formulation. Partiality can be handled naturally in our approach due to the universe formulation. Since each P_i maps the points of $\mathcal{X}_i$ to a *subset* of the d universe points, this case boils down to choosing the correct universe points. Assuming that all given partial shapes represent parts of the same full shape, the optimal universe would model exactly the full geometry.

The functional maps $\mathcal{C}_i$ need to be adjusted slightly for

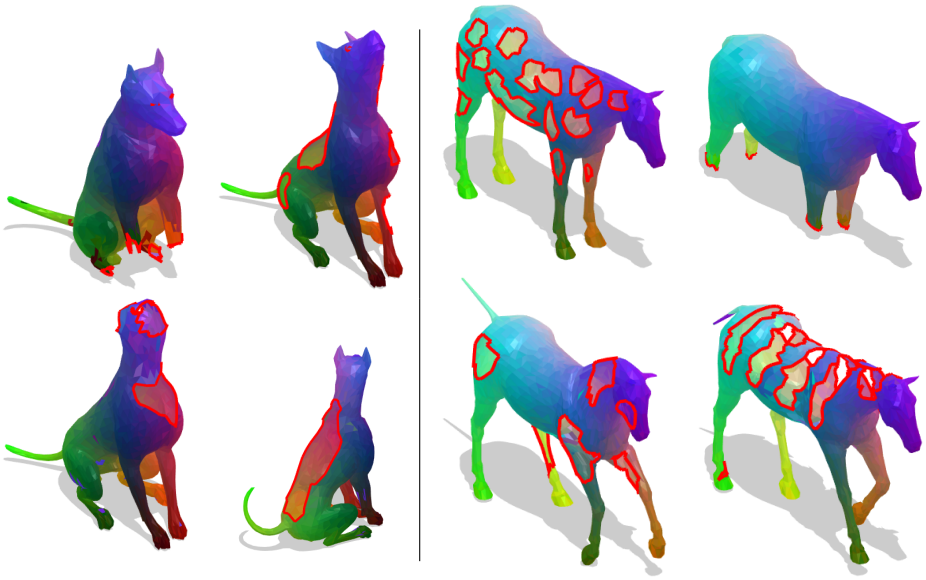

Figure 2. Our results on deformed partial shapes of the TOSCA dog and horse classes. Although partial-to-partial matching is a very challenging setting, our method produces high quality results. These results serve as a proof-of-concept that our method is applicable to such partial settings.

this setting. As explained in Section 3.2 in the main paper, square orthogonal $\mathcal{C}$s model area-preservation. This is meaningful for isometries, but, since partial shapes literally miss some areas, it does not hold in this case. Instead, we use the theory about partial functional maps provided in [9]. According to [9], functional maps for the partial case have *slanted* diagonals and the area preservation only holds in one direction. Additionally, some LBO eigenfunctions of the full shape do not appear on the partial shapes, such that each $\mathcal{C}_i$ needs to map to a higher dimensional space, and only choose the corresponding eigenfunctions there. Therefore, instead of being square, the matrices are rectangular, and we adjust the definition of the orthogonality constraint as

$$\mathbb{O}_b^P = \left\{ \mathcal{C} \in \mathbb{R}^{b \times b'} : \mathcal{C}\mathcal{C}^\top = \mathbf{I}_b \right\}, \tag{S7}$$

where $b' > b$, and we chose $b' = 1.2b$ in all our partial experiments. Note that it does not require any modification in our optimisation pipeline and our problem formulation is capable of handling this more challenging case.

Initialisation. For the full multi-shape matching pipeline, we used functional maps [8] and ZoomOut [7] to get an estimation for each $\mathcal{C}_{ij}$ and P_{ij}. However, they are not well-suited for directly performing partial-to-partial matching. Instead, we directly compute $\{P_i\}$ between each partial and the full shape using a combination of SHOT [10], Heat Kernel Signature [4], Wave Kernel Signature [1] and symmetry descriptors, which are subsequently refined using a partiality-adjusted version of ZoomOut to obtain the shape-to-universe initialisation for IsoMush. To initialise U and Q in our partial IsoMush, we adjusted ZoomOut to use rectangular functional maps (S7) by always involving a larger (fixed) number of eigenfunctions on the full (universe) shape than on the partial shapes in every ZoomOut iteration (according to the theory from [9]). We fix this ratio to be $0.2b$ across all our partial experiments.

4. Additional Qualitative Results

We show qualitative results on FAUST in Fig. 3, the complete results on SCAPE in Fig. 4, as well as additional qualitative results of different TOSCA classes in Fig. 5 and Fig. 6. Fig. 3 shows the main source of errors for our method on FAUST, which are front-back flips. This is due to the intrinsic front-back near-symmetry of humans and descriptors that do not discriminate well between these. All ZoomOut variants (including the initialisation for our method) suffer from this problem. Note that even though the correspondence is flipped, our results are still cycle-consistent.

References

[1] Mathieu Aubry, Ulrich Schlickewei, and Daniel Cremers. The wave kernel signature: A quantum mechanical approach to shape analysis. In *ICCV Workshops*, 2011. 3

[2] Florian Bernard, Johan Thunberg, Peter Gemmar, Frank Hertel, Andreas Husch, and Jorge Goncalves. A Solution for Multi-Alignment by Transformation Synchronisation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2161–2169, 2015. 2

Figure 3. Qualitative examples of correspondences on FAUST registrations. Black indicates no matching due to non-bijectivity. Our results contain the least noise and are cycle-consistent. ‡CONSISTENTZOOMOUT obtains cycle-consistent C_{ij}, but not P_{ij}. (Best viewed magnified on screen)

[3] Florian Bernard, Johan Thunberg, Jorge Goncalves, and Christian Theobalt. Synchronisation of partial multi-matchings via non-negative factorisations. *Pattern Recognition*, 92:146–155, 2019. 2

[4] Michael M Bronstein and Iasonas Kokkinos. Scale-invariant heat kernel signatures for non-rigid shape recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2010. 3

[5] Luca Cosmo, Emanuele Rodolà, Andrea Albarelli, Facundo Mémoli, and Daniel Cremers. Consistent partial matching of shape collections via sparse modeling. In *Computer Graphics Forum*, volume 36, pages 209–221, 2017. 2

[6] Haggai Maron, Nadav Dym, Itay Kezurer, Shahar Kovalsky, and Yaron Lipman. Point registration via efficient convex relaxation. *ACM Transactions on Graphics (TOG)*, 35(4):1–12, 2016. 2

[7] Simone Melzi, Jing Ren, Emanuele Rodolà, Abhishek Sharma, Peter Wonka, and Makis Ovsjanikov. Zoomout: Spectral upsampling for efficient shape correspondence. *ACM Transactions on Graphics (Proc. SIGGRAPH Asia)*,

Figure 4. Complete qualitative results of correspondences on SCAPE. Black indicates no matching due to non-bijectivity. Our results contain the least noise and are cycle-consistent, although there is one outlier shape where neither HIPPI nor our method could recover from a bad initialisation. ‡CONSISTENTZOOMOUT obtains cycle-consistent $\mathcal{C}_{ij}$, but not P_{ij}. (Best viewed magnified on screen)

2019. 3

[8] Maks Ovsjanikov, Mirela Ben-Chen, Justin Solomon, Adrian Butscher, and Leonidas Guibas. Functional maps: a flexible representation of maps between shapes. *ACM Transactions on Graphics (TOG)*, 31(4):1–11, 2012. 3

[9] Emanuele Rodolà, Luca Cosmo, Michael Bronstein, Andrea Torsello, and Daniel Cremers. Partial functional correspondence. *Computer Graphics Forum (CGF)*, 2016. 2, 3

[10] Samuele Salti, Federico Tombari, and Luigi Di Stefano. Shot: Unique signatures of histograms for surface and texture description. *Computer Vision and Image Understanding*, 125:251–264, 2014. 3

[11] A Singer and Y Shkolnisky. Three-Dimensional Structure Determination from Common Lines in Cryo-EM by Eigenvectors and Semidefinite Programming. *SIAM J Imaging Sciences*, 4(2):543–572, June 2011. 2

[12] Yifan Sun, Zhenxiao Liang, Xiangru Huang, and Qixing Huang. Joint map and symmetry synchronization. In *European Conference on Computer Vision (ECCV)*, pages 251–264, 2018. 2

[13] Roberto Tron, Xiaowei Zhou, Carloes Esteves, and Kostas Daniilidis. Fast Multi-Image Matching via Density-Based Clustering. In *International Conference on Computer Vision (ICCV)*, 2017. 2

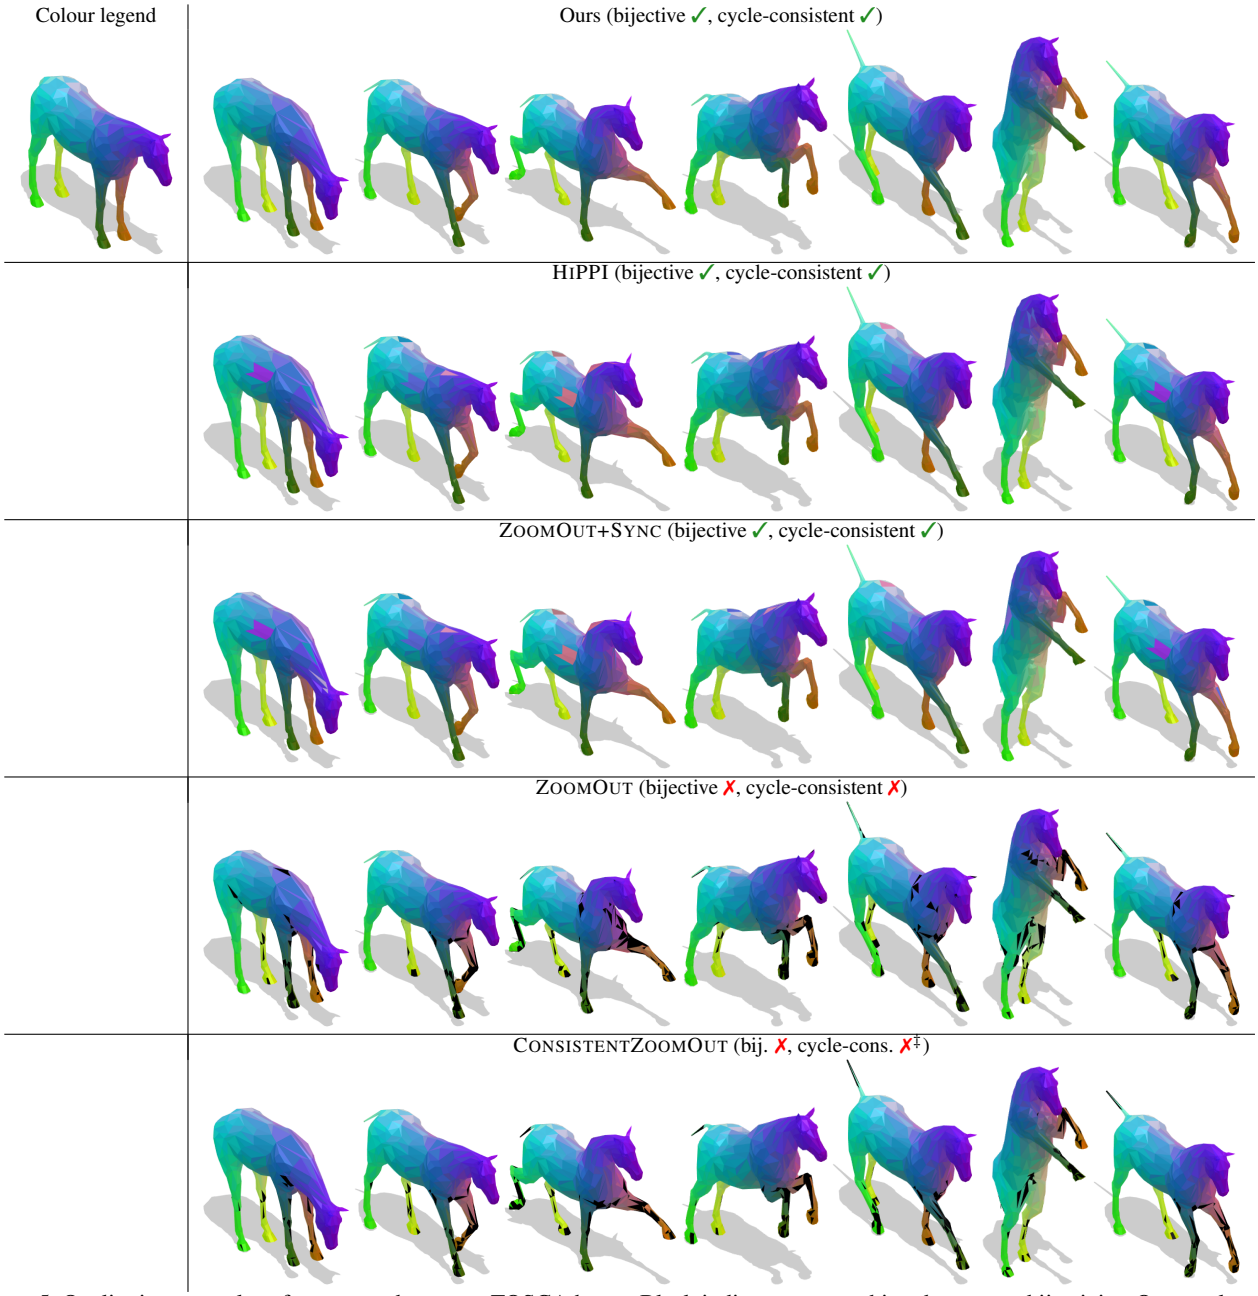

Figure 5. Qualitative examples of correspondences on TOSCA horse. Black indicates no matching due to non-bijectivity. Our results contain the best results and are cycle-consistent. ‡CONSISTENTZOOMOUT obtains cycle-consistent $\mathcal{C}_{ij}$, but not P_{ij}. (Best viewed magnified on screen)

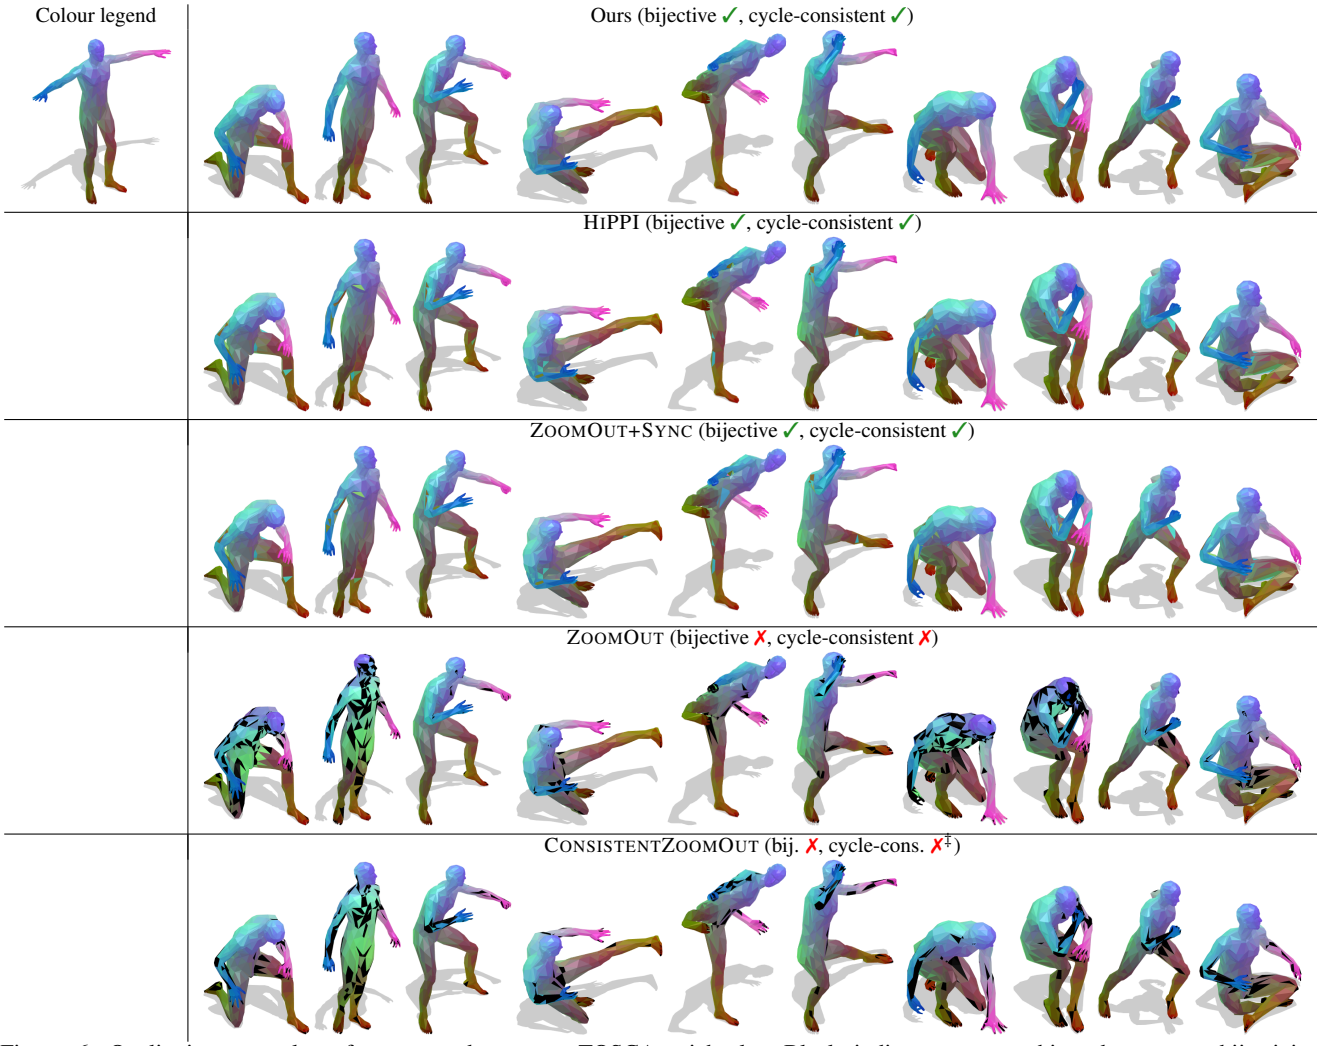

Figure 6. Qualitative examples of correspondences on TOSCA michael. Black indicates no matching due to non-bijectivity. ‡CONSISTENTZOOMOUT obtains cycle-consistent $\mathcal{C}_{ij}$, but not P_{ij}. (Best viewed magnified on screen)