SpiderMatch: 3D Shape Matching with Global Optimality and Geometric Consistency

Supplementary Material

Paul Roetzer[1] Florian Bernard[1] University of Bonn[1]

8. Geometric Consistency Evaluation

Previous works, e.g. [21], use conformal distortion errors to evaluate geometric consistency. Yet, conformal distortion only quantifies the *smoothness* of the resulting matching as it only considers deformation of individual triangles [22]. Consequently, smooth but non-geometrically consistent matchings yield low conformal distortion errors as can be seen in Fig. 10, in which we show various qualitative results with respective conformal distortion errors as well as Dirichlet energies.

Yet, we want to quantify *geometric consistency*. For that, we compute the Dirichlet energy [14, 24, 45] in the following manner. We consider the Laplacian $L^{\mathcal{X}}$ of shape $\mathcal{X}$ and the difference between vertex positions of matched vertices, i.e. deformation vectors $\Delta = V^{\mathcal{X}} - P_{\mathcal{Y}}^{\mathcal{X}} \bar{V}^{\mathcal{Y}}$ (where $\bar{V}^{\mathcal{Y}}$ are the to $V^{\mathcal{X}}$ rigidly aligned vertices $V^{\mathcal{Y}}$ of shape $\mathcal{Y}$ using ground truth matchings[3]) and $P_{\mathcal{Y}}^{\mathcal{X}} \in \{0,1\}^{|V^{\mathcal{X}}| \times |V^{\mathcal{Y}}|}$ is the point-wise matching matrix assigning vertices of $\mathcal{X}$ to vertices of $\mathcal{Y}$. With that, we compute the Dirichlet energy as

$$E_{\text{Dirichlet}} = ||\Delta||^2_{L^{\mathcal{X}}} = \operatorname{tr}\left(\Delta^T L^{\mathcal{X}} \Delta\right). \tag{4}$$

In Fig. 11, we show cumulative Dirichlet energy scores for different methods. We can see that our method produces low Dirichlet energy scores which validates that our method indeed produces geometrically consistent results.

[3]note that we only use the ground truth for evaluation purposes

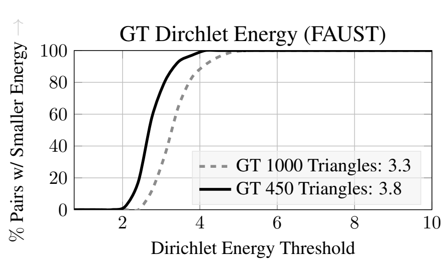

Figure 9. Comparison of **Dirichlet Energies for ground truth (GT) matchings** at different resolutions on FAUST dataset. Numbers in legends are mean Dirichlet Energies across all pairs. We can see that Dirichlet energies at resolution of 450 triangles and at resolution of 1000 are not fully comparable (even when using ground truth matchings as in this plot) due to different densities of deformation vectors Δ.

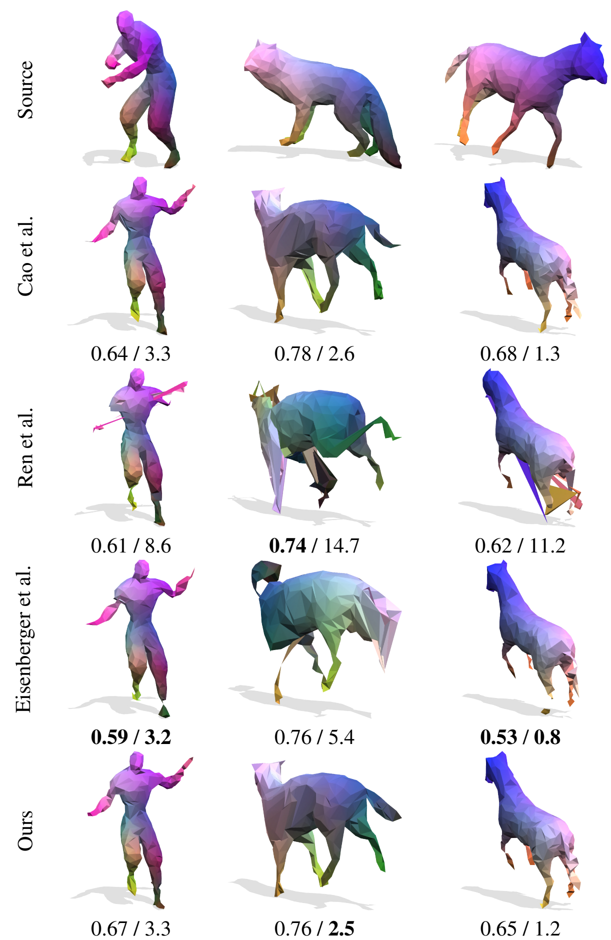

Figure 10. Qualitative results and comparison of values of **Conformal Distortion Errors / Dirichlet Energies** for shown pairs. Even though results for Ren et al. and Eisenberger et al. are not geometrically consistent (emphasised by triangulation transfer via the computed matching), certain conformal error values are lower than geometrically consistent results computed with our method. Further, we can see from shown examples that Dirichlet energy better quantifies geometric consistency due to values being lower for geometrically consistent solutions and vice versa.

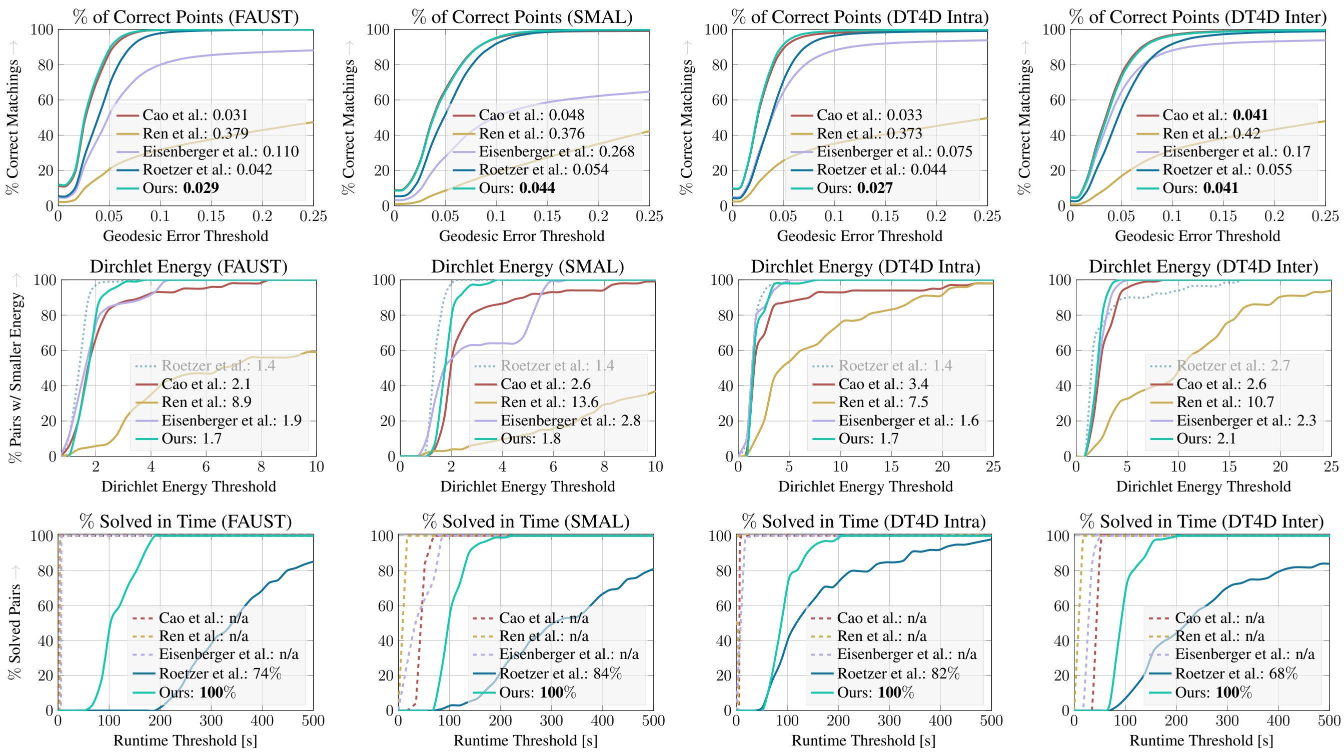

Figure 11. (**Top**) For convenience we repeat the the geodesic error plots on respective datasets from the main paper. Values in legends are mean geodesic errors. (**Middle**) Comparison of Dirichlet Energies on datasets FAUST, SMAL and DT4D. Numbers in legends are mean Dirichlet energies across all pairs. Note that Dirichlet energies on different shape resolution are not fully comparable due to different density of deformation vectors Δ (cf. also Fig. 9). Consequently, we only include Dirichlet energies of Roetzer et al. for reference as matchings for Roetzer et al. are computed at resolution of 450 triangles while matchings for all other methods are computed with resolution of 1000 triangles. (**Bottom**) Percentage of instances solved in certain time (in seconds). The values in legends are the percentage of instances solved to global optimality for each method. Note that only Roetzer et al. and our approach are able to certify global optimality (thus we put "n/a" values for Cao et al., Ren et al. and Eisenberger et al.). Our approach outperforms Roetzer et al. w.r.t. runtime even though shapes used by our method have *more than 2x more triangles* (1000 triangles for our method versus 450 triangles for Roetzer et al.). (**Conclusion**) Considering all metrics, our method finds the best balances between matching performance, geometric consistency guarantees and runtimes among all comptetitors and across all datasets.

9. Runtime Experiments

In Fig. 1 right, we show a runtime comparison of our method, Roetzer et al. and Cao et al. We sample five random pairs from the FAUST dataset and decimate the individual shapes to the respective amount of triangles and measure runtimes of the opimisation. The thick line is the mean runtime among the five pairs at each individual resolution.

Furthermore, we show the percentage of instances solved within a given runtime in Fig. 11 bottom. Even though our method considers shapes with more than twice the amount of triangles than the approach by Roetzer et al., we can see that our approach is much faster than Roetzer et al. (which is the only other method which guarantees geometric consistency).

10. Implementation Details

We avoid severe discretisation artefacts by adding additional edges to the edges of the triangle meshes of respective shapes $\mathcal{X}$ and $\mathcal{Y}$, see Fig. 12 for an illustration. This has the effect that the optimisation might find better shortest paths which would otherwise not be within the feasible region of problem 3 (i.e. due to different connectivity of neighbouring vertices).

We note that Roetzer et al. [55] optimises for an elastic energy consisting of a bending term (which compares curvatures at vertices of both shapes) and a membrane term (which penalises triangle deformations). This elastic energy does not perform as good as learned features and thus we consider it an unfair comparison to our method which utilises learned features. Hence, we replace the elastic energy originally proposed in [77] and adopted by [55] with the same learned features [14] as used by our method (note we compute feature differences for Roetzer et al. as shown in [21]). We emphasise that this improves the approach by Roetzer et al. w.r.t. runtime and matching performance.

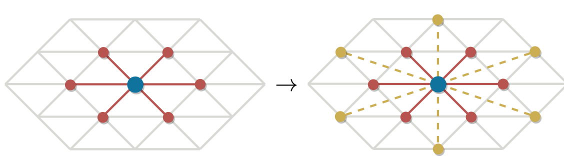

Figure 12. Inspired by edge flipping for intrinsic Delaunay triangulations (see [67]), we augment edges $E^{\mathcal{X}}$, $E^{\mathcal{X}}$ with **additional edges** (shown in yellow) that allow to directly reach the two-ring neighbourhood from the blue vertex. We only add yellow edges for pairs of triangles which are non-obtuse (i.e. all angles of both triangles are smaller than π).

11. Ablation: Resolution and Accuracy

In Fig. 13 (left), we qualitatively show a sphere with 50 triangles matched to a sphere with 1000 triangles, as well as a sphere with 100 triangles matched to a sphere with 1000 triangles, which showcases that our method can handle severe discretisation differences. In Fig. 13 (right), we show matching error using five instances from the FAUST dataset, where we fix the resolution of the target shape to 1000 triangles and match different resolutions of the source shape. We observe that our method can handle a reasonable amount of resolution differences.

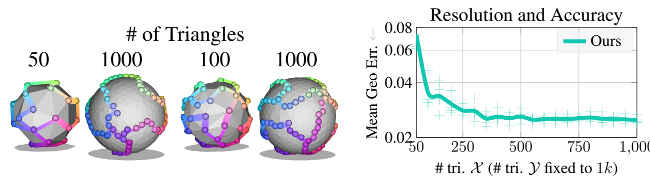

Figure 13. (**Left**) Qualitative results of matching a sphere with 50 triangles to a sphere with 1000 triangles, as well as matching a sphere with 100 triangles respectively to a sphere with 1000 triangles. (**Right**) Mean geodesic errors for our method when the resolution of the target shape is fixed to 1000 triangles and the resolution of the source shape is varied.

12. Additional Qualitative Results

We show qualitative results of Roetzer et al. in Fig. 14 which complete the qualitative comparison of Fig. 6 in the main paper.

Furthermore, in Fig. 15, we show failure cases of our method. Our method does not result in any globally wrong matchings (e.g. left-right flips happening for Eisenberger et al.). Yet, discretisation artefacts may lead to local distortions especially visible if triangulation is transfered from one shape to the other (cf. also Fig. 16).

Finally, in Fig. 16, we show qualitative results by transfering triangulation from one shape to the other via the computed matching. By doing so, geometric inconsistencies can be detected more easily. We note that only methods by Cao et al., Ren et al. and Eisenberger et al. can produce geometrically inconsistent solutions. Furthermore we note that our method does (by definition, see Problem 3) *not* lead to such inconsistencies and to instead leads to generally smoother matchings (cf. also Fig. 11). For a full picture of qualitative results, we include *all qualitative results for all methods on all datasets* with triangulation transfer in the supplementary material as html webpage.

Figure 14. **Qualitative comparison** for Roetzer et al. Pairs in individual columns are identical to pairs in columns of Fig. 6 in main paper. We note that we compute matchings for Roetzer et al. at 450 triangles (instead of 1000 triangles used for all other methods) to keep runtimes manageable (cf. also Fig. 1 right).

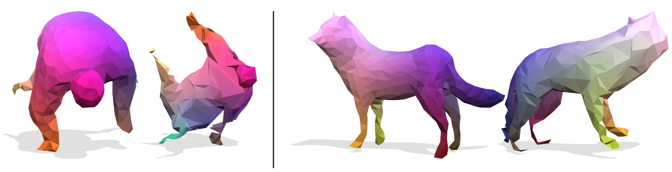

Figure 15. **Failure cases** of our method. Even though our constraints ensure geometric consistency, discretisation artefacts may lead to local triangle deformations (see arms for left shape pair or legs for right shape pair) when transferring triangulation from one shape to the other.

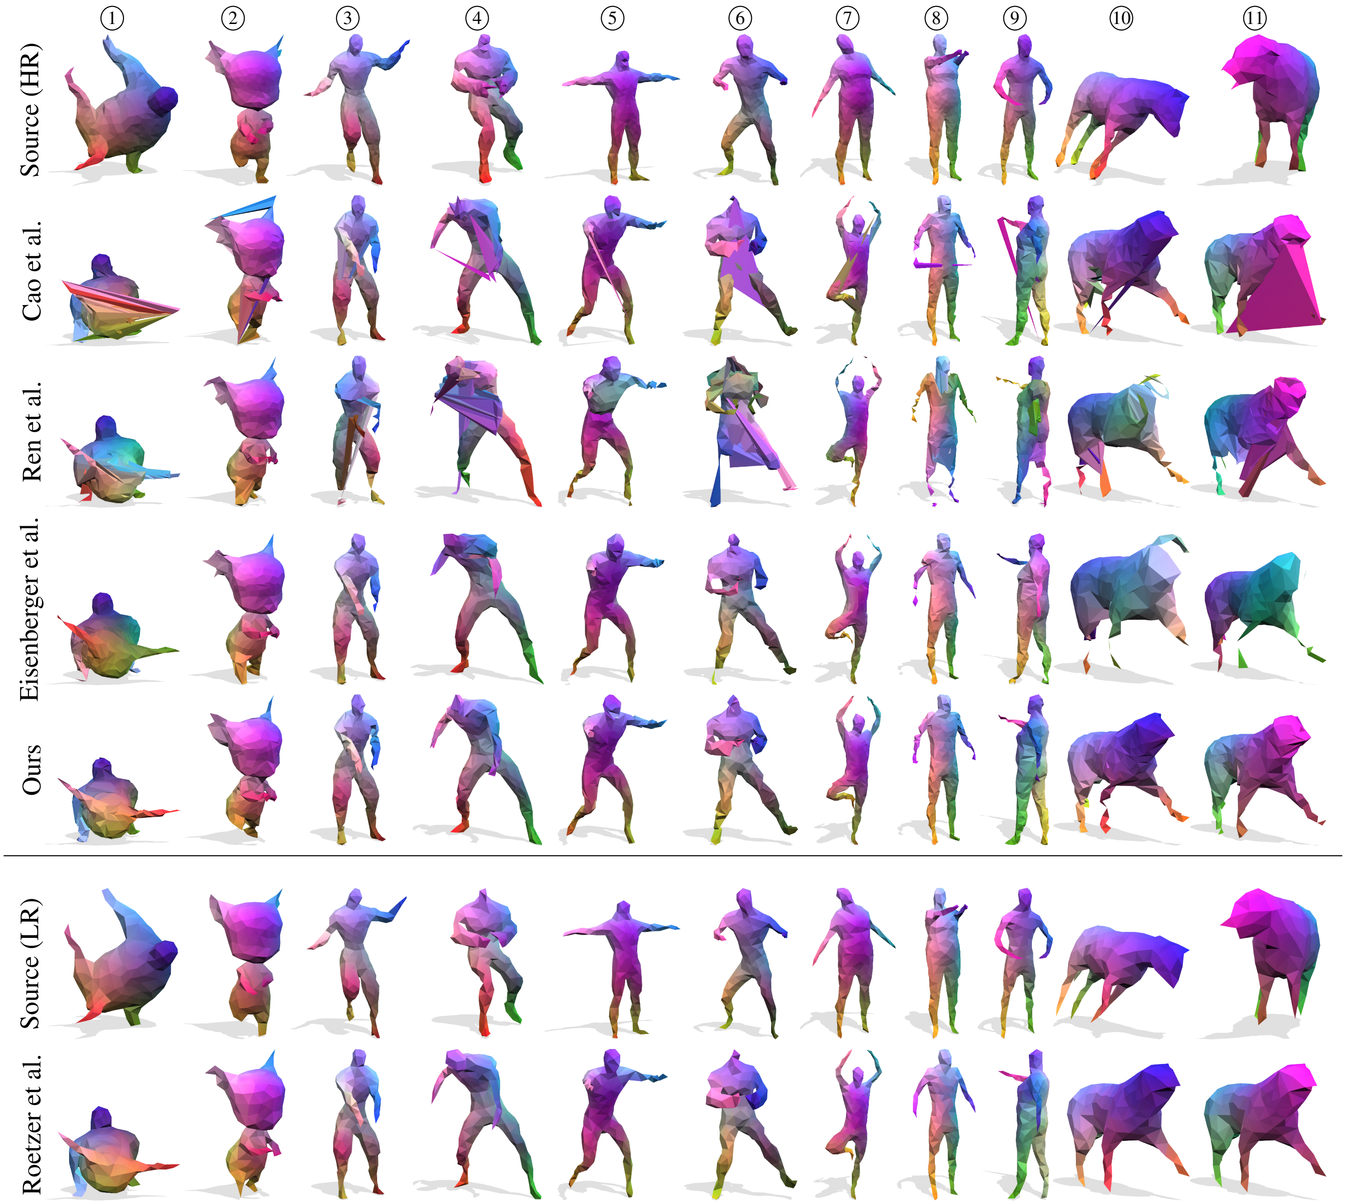

Figure 16. **Qualitative comparison** of shapes from datasets FAUST, SMAL and DT4D with triangulation transfer (triangulation is transfered via the computed matching). "Source (HR)" is the source shape for Cao et al., Ren et al., Eisenberger et al. and ours decimated to 1000 triangles while "Source (LR)" is the source shape for Roetzer et al. decimated to 450 triangles (target shapes have 1000 and 450 triangles respecitvely). By transfering triangulation from one shape to the other via the computed matching we can see geometric inconsistencies of solutions by large deformed triangles (cf. ①, ②, ⑪ of Cao et al., ⑤, ⑥ of Ren et al. and ⑩, ⑪ of Eisenberger et al.). Such deformations of triangles do not arise for Roetzer et al. and for our method since both methods can guarantee geometrically consistent matchings. Despite Roetzer et al. producing visually the best results (especially when transfering triangulation), we note that it is the slowest method among all competitors, even though shapes have less than 50% the amount of triangles compared to shapes used for all other methods (450 versus 1000 triangles).