



QR Code:
(github)

Mentor: Rayan Saab
rsaab@ucsd.edu

Mentor: Alex Cloninger
acloninger@ucsd.edu

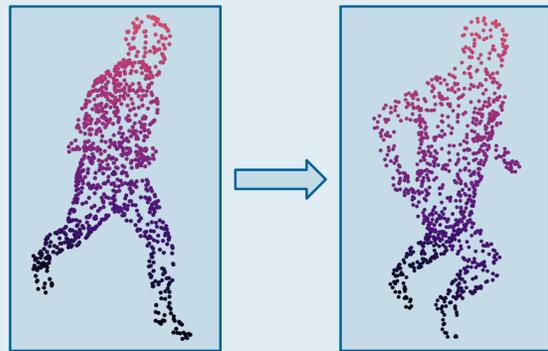
Matt Tokunaga
m2tokunaga@ucsd.edu

Quy-Dzu Do
qndo@ucsd.edu

Four Dimensional Embedding Increases Accuracy in Body Region Matching

Background

- High spatiotemporal resolution point clouds common in applications
- Matching similar regions across point clouds for action recognition, skeleton construction
- Optimal transport for one-to-one point cloud mapping; struggles with human bodies
- **Key failure: confusion between right vs. left legs**



Data

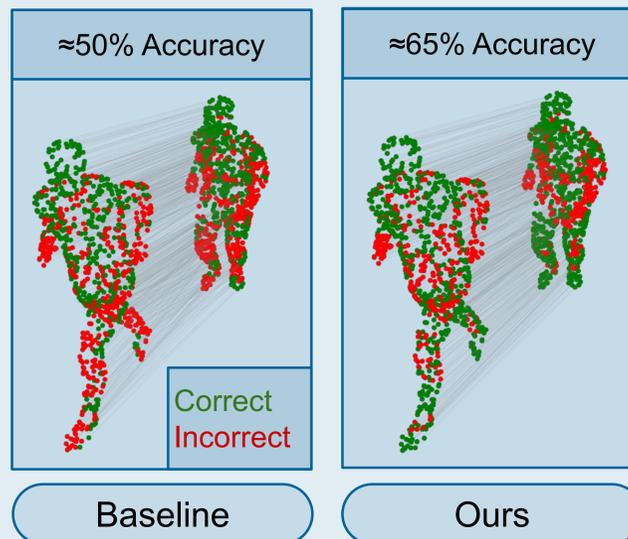
- Point clouds sampled from AMASS human body meshes
- 1,000 surface points per frame
- Walking and running sequences (≈ 120 fps)
- Mesh-derived body regions as ground truth
- 16 manually defined regions



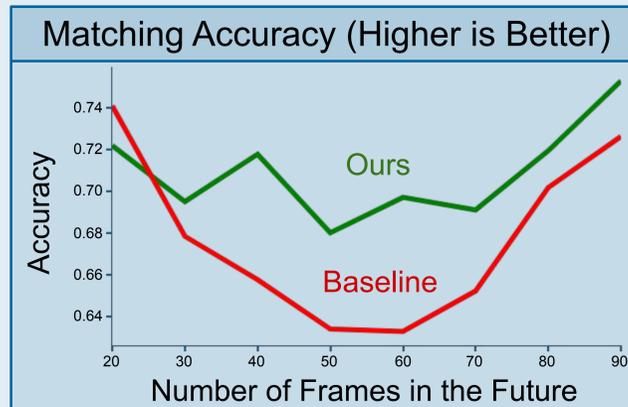
Research Question

How can we improve an optimal transport-based point matching for body regions?

Results

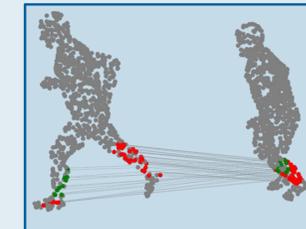


Huge lower body improvement



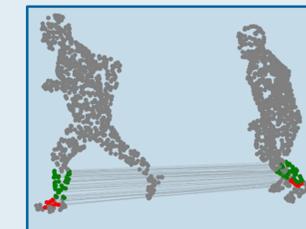
Better accuracy for further frames

Naive algorithm maps left shin points to both legs

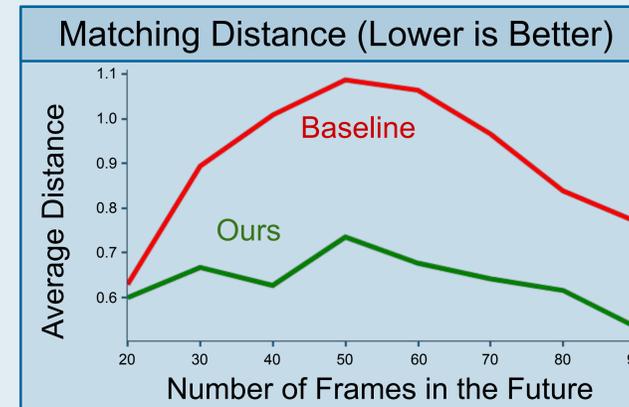


Baseline

Our algorithm keeps left shin points on the left side

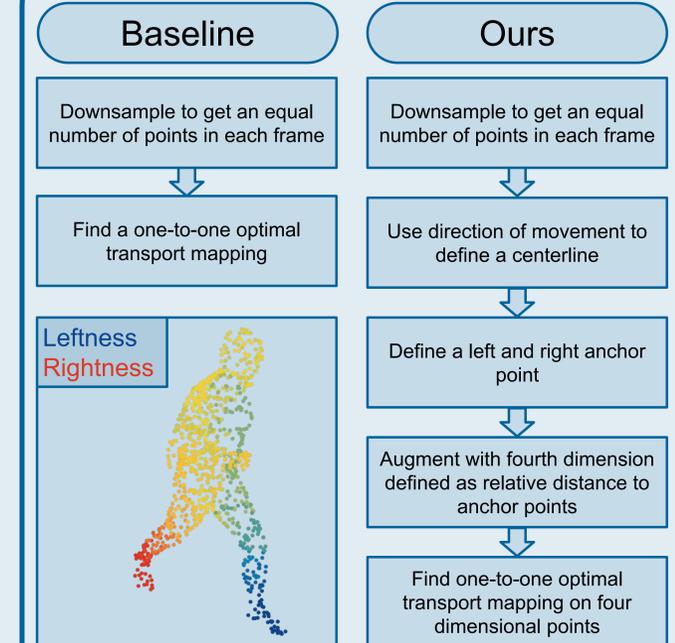


Ours

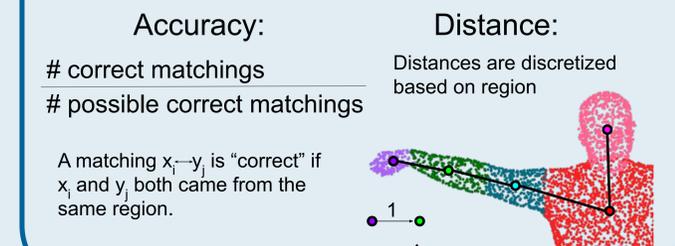


Much closer on average, even on mismatches

Methods



Evaluation



Conclusions

- Disambiguating left vs. right extremities improves region-preserving point cloud mapping
- Higher accuracy overall
- Less severe errors when mismatches occur