Iterative Cage-based Registration for Dynamic Shape Capture

Yann Savoye*
INRIA Bordeaux University, France.

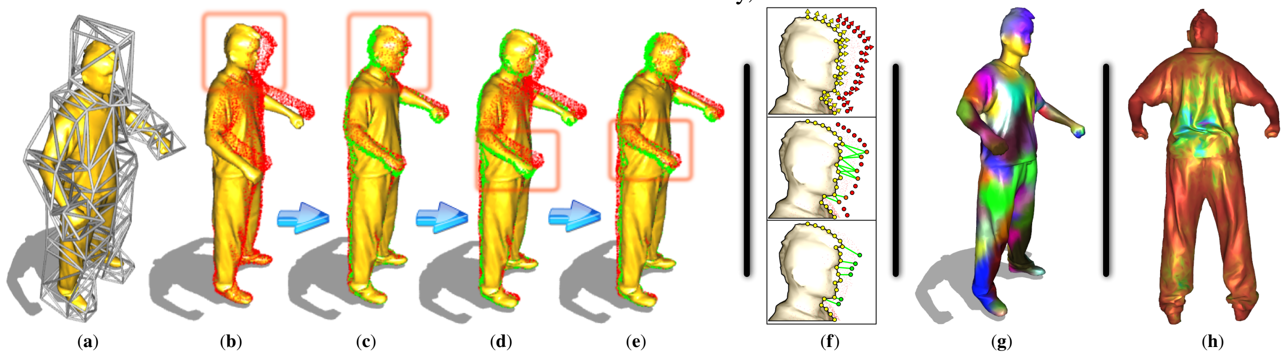

Figure 1: ***Life-like Dynamic Registration of Clothed Organic Shape.*** *A generic humanoid-type cage is enveloping a laser-scanned template mesh roughly, to perform scalable registration of whole-body scanned models (**a**). We aim at deforming the source template (in yellow color) toward a target point cloud (in red color) (**b**). The skin-detached registration procedure pulls the cage-based geometry toward corresponding target locations (**c**). After a sufficient number of decade iterations, the next target point cloud is reconstructed (**d**), and the overall iterative procedure is repeated (**e**). Normal-guided pairwise correspondences (in green color) are pruned, satisfying a criteria based on smoothed normals difference (**f**). Consequently, handle-aware overlapping harmonic-rigidities (**g**) are well-suited to register the non-rigid edge-length deviation (**h**) of the surface with controllable flexibility.*

1 Introduction

Recent advances in low-cost dynamic scanning turn the cross-parametrization of non-rigid animatable surface into a vision-oriented ill-posed problem. In contrast with [Li et al. 2012], we propose a novel detail-preserving registration approach with resolution-independent control. Furthermore, our skin-detached surface registration avoids patch-based segmentation or affine fitting to maintain the local plasticity, as required in [Budd and Hilton 2010]. In particular, we leverage the problem of highly non-rigid spacetime registration by employing an elasto-plastic coarse cage. Thus, we perform scalable handle-aware harmonic shape registration, relying on the high-level of shape abstraction offered by the space-based paradigm. To the best of our knowledge, our technique is the first to investigate handle-aware elastic overlapping-rigidities for registering life-like dynamic shapes in full-body clothing.

2 Handle-Aware Detached Registration

Non-Rigid Registration Setup. We propose to evolve the fixed connectivity $\mathcal{F}$, offered by the given dense template mesh $\mathcal{M} = \{\mathcal{V}, \mathcal{F}\}$, by registering roughly-and-temporally a given set $\mathcal{P} = \{\mathcal{P}_0, \cdots, \mathcal{P}_l\}$ of l unstructured time-varying point clouds. We assume no prior knowledge about the temporal matching. The geometry of the template mesh is written by $\mathcal{V}$. Let us denote $\Omega \subset \mathbb{R}^3$ the bounded domain included by m control cage-handles enveloping the static template mesh. We designate by $\mathbf{c}_j$ the current location of the j^{th} cage handle in the global coordinates system. This cage polytope structure is augmented by the Laplace-Beltrami Operator $\mathcal{L}(\cdot)$ with non-uniform cotangent weights, allowing scalable template mesh registration. Associated differential cage coordinates δ encode each cage-handle relatively to its neighborhood in the cage connectivity. Finally, the geodesic-aware relationship between the volumetric subspace and the static template is encapsulated by a bi-harmonic rigging process, computed once at the default pose.

Normal-Guided Pairwise Correspondences. We adopt a similar fuzzy-yet-robust geometric strategy than [Budd and Hilton 2010] to infer a minimal set of compatible feature correspondences $\mathcal{S} = \{s_k : (k, \mathbf{q}_k, \gamma_k)\}$, updated at each intra-frame iteration. A unique target location $\mathbf{q}_k \in \mathbb{R}^3$ is obtained for the k^{th} current template vertex by averaging candidates in the current point cloud of $\mathcal{P}$. After the outliers pruning, each correspondence is weighted by γ_k defined as the dot product of smoothed pairwise normals.

*ysavoye@siggraph.org

Iterative Elasto-Plastic Optimization. For each point cloud, our cage-handle curve registration process alternates successively between deformation optimization and correspondences. At each intra-frame iteration $t+$, we initialize the cage geometry with location $\mathbf{c}^t$ obtained at the previous iteration t. Then, we update $\mathcal{L}(\cdot)$, the corresponding δ and we infer $\mathcal{S}$. Finally, driven by the correspondences propagation in the subspace, a new cage pose $\mathbf{c}^{t+}$ is estimated by solving the following variational objective function:

$$\underset{\{\mathbf{c}_1^{t+},\cdots,\mathbf{c}_m^{t+}\}}{\operatorname{argmin}} \left(\alpha \cdot \sum_{j=1}^{m} \left\| \mathcal{L}^t\left(\mathbf{c}_j^{t+}\right) - \delta_j^t \right\|_2^2 + \beta \cdot \sum_{s_k \in \mathcal{S}^t} \gamma_k \cdot \left\| \mathbf{q}_k - \sum_{j=1}^{m} w_{kj} \cdot \mathbf{c}_j^{t+} \right\|_2^2 \right)$$

where $w_{kj} : \Omega \to \mathbb{R}$ is the biharmonic weight for a given cage handle j with respect to the k^{th} template vertex, as proposed in [Jacobson et al. 2011]. Consequently, the registered template geometry is generated by a cage-based warping field with low-distortion.

Weight-Control Update Rules. The data-term weight-control is initialized at $\beta = 0.01$ and increases along iterations by following an exponential growth rule to promote the constraint-guided bending energy. The weight-control α enforcing the shape-prior is set-up to 1 and slightly decreases to relax the deformation stiffness prior.

3 Conclusions

Our new approach is a first step toward the automatic template-based registration of highly non-rigid dynamic shape using low-dimensional space-based encoding. We train the effectiveness of our algorithm by aligning several real-world datasets of [Vlasic et al. 2008]. The main advantage of our iterative optimization remains in the simultaneous cross-reconstruction of dynamic shape, and skin-detached registration of reusable temporal curves expressing the clothed-body deformations. In brief, we proposed a new system that registers shape variations while preserves the life-likeness of captured data, and acquires reusable consistent surface parameters. We expect to pursue our on-going efforts to perform better fine-tuned non-rigid alignment for large organic motion.

References

Budd, C., and Hilton, A. 2010. Temporal alignment of 3d video sequences using shape and appearance. In *CVMP'10: 9th Conference on Visual Media Production*.

Jacobson, A., Baran, I., Popović, J., and Sorkine, O. 2011. Bounded biharmonic weights for real-time deformation. *ACM Trans. Graph. 30*.

Li, H., Luo, L., Vlasic, D., Peers, P., Popović, J., Pauly, M., and Rusinkiewicz, S. 2012. Temporally coherent completion of dynamic shapes. *ACM Trans. Graph. 31*.

Vlasic, D., Baran, I., Matusik, W., and Popović, J. 2008. Articulated mesh animation from multi-view silhouettes. *ACM Trans. Graph. 27* (August), 97:1–97:9.