IEEE 2017 Conference on Computer Vision and Pattern Recognition

Discriminative Optimization: Theory and Applications to Point Cloud Registration

Jayakorn Vongkulbhisal[1,2], Fernando De la Torre[1], and João Paulo Costeira[2]
[1]Carnegie Mellon University, USA [2]IST, Universidade de Lisboa, Portugal
http://humansensing.cs.cmu.edu/DO

What is a good cost function for point cloud registration?

Example: Three 2D point cloud registration cases.

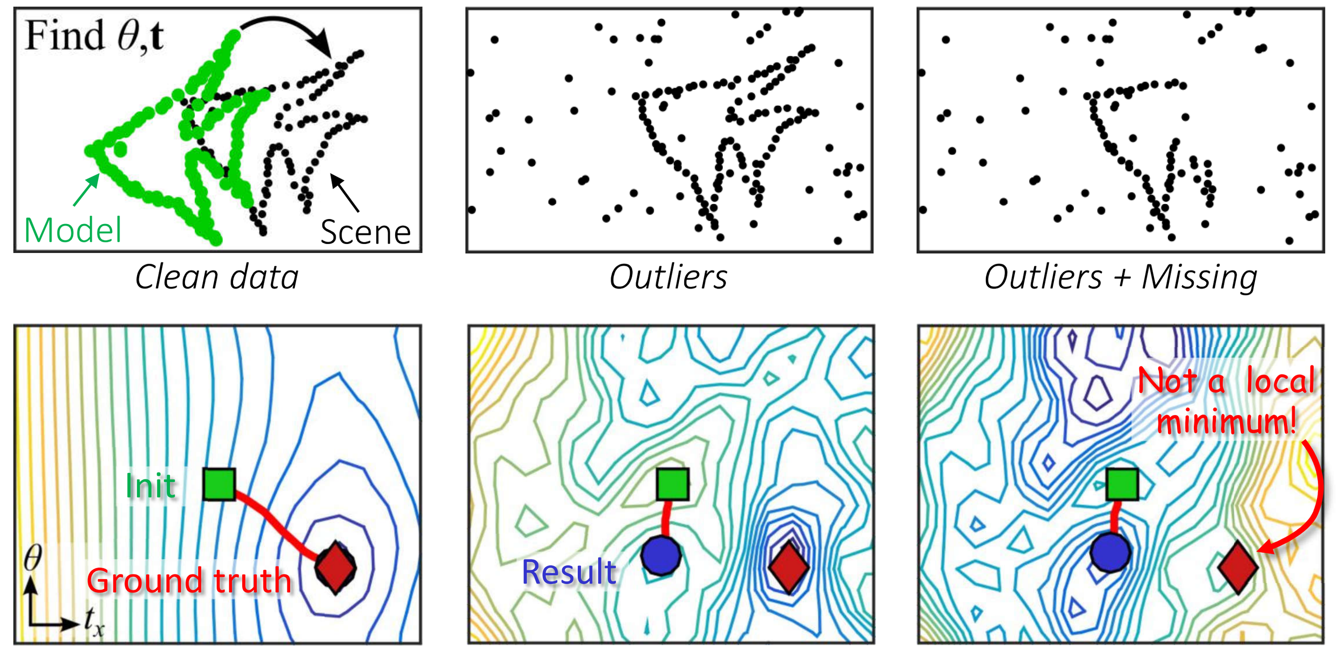

Clean data | *Outliers* | *Outliers + Missing*

'Clean' surface | *Multiple local minima* | *Ground truth is not a local minimum!*

ICP [1] solves the following optimization problem:

$$\min_{\substack{p_{ij}\in\{0,1\},\\ \mathbf{R}\in SO(2),\mathbf{t}\in\mathbb{R}^2}} \sum_{i,j} p_{ij}\|\mathbf{R}\mathbf{m}_i+\mathbf{t}-\mathbf{s}_j\|_2^2$$

Problem: Noise, outliers, missing data, etc., can deform the cost surfaces, resulting in incorrect solutions.

Three problems with the existing optimization-based approaches:
- What is a good cost function for a given problem?
- What is an efficient way to find a solution?
- How can we avoid local minima and bad solutions?

Discriminative Optimization (DO)

Idea: Follow the vector fields from the initializations to the solutions.

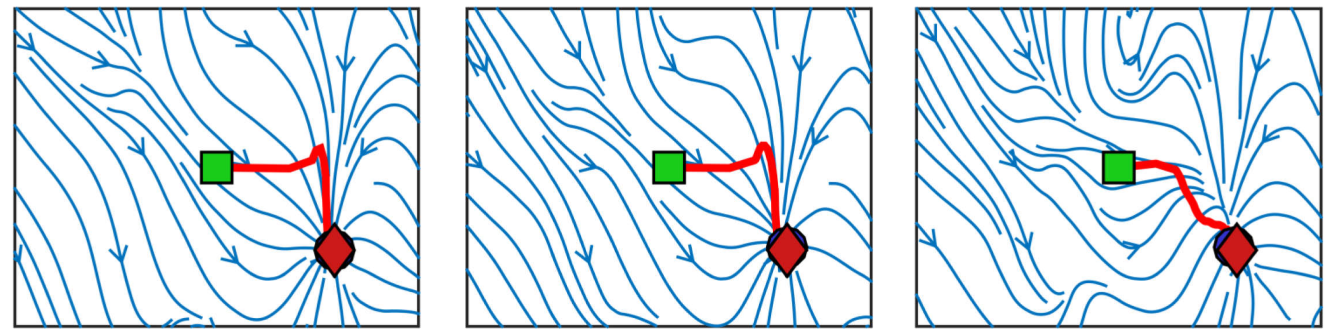

Update directions of DO points to the 'correct' solution.

Bonus: DO also expands the region of convergence.

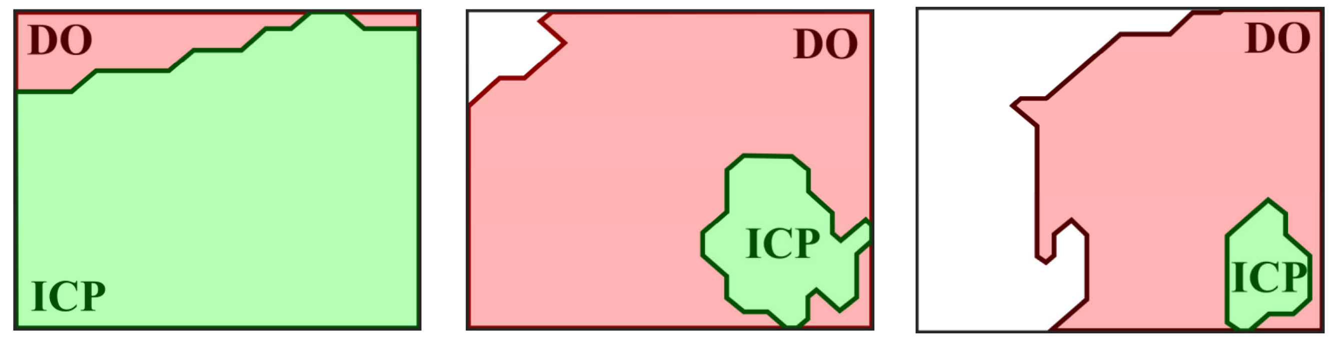

How does DO work?

Update rule

$$\mathbf{x}_{t+1} = \overbrace{\mathbf{x}_t}^{\text{Estimate at time } t} - \underbrace{\mathbf{D}_{t+1}}_{\substack{\text{A linear map}\\\text{(to be learned)}}}\underbrace{\mathbf{h}(\mathbf{x}_t)}_{\text{Feature function}}$$

Update vector: $\mathbf{D}_{t+1}\mathbf{h}(\mathbf{x}_t)$

Related works: Cascaded regression [2], Supervised Descent (SDM) [3], etc.

Theoretical result

A function $\mathbf{f}:\mathbb{R}^p\to\mathbb{R}^p$ is *strictly* monotone at $\mathbf{x}_*\in\mathbb{R}^p$ if

$$(\mathbf{x}-\mathbf{x}_*)^\top\mathbf{f}(\mathbf{x})>0, \forall\mathbf{x}\in\mathbb{R}^p\setminus\{\mathbf{x}_*\}.$$

If there exists $\hat{\mathbf{D}}$ such that $\hat{\mathbf{D}}\mathbf{h}^{(i)}$ is strictly monotone at $\mathbf{x}_*^{(i)}$ for all i, then the training error of DO strictly decreases in every iteration:

$$\sum_{i=1}^n\|\mathbf{x}_{t+1}^{(i)}-\mathbf{x}_*^{(i)}\|_2^2<\sum_{i=1}^n\|\mathbf{x}_t^{(i)}-\mathbf{x}_*^{(i)}\|_2^2.$$

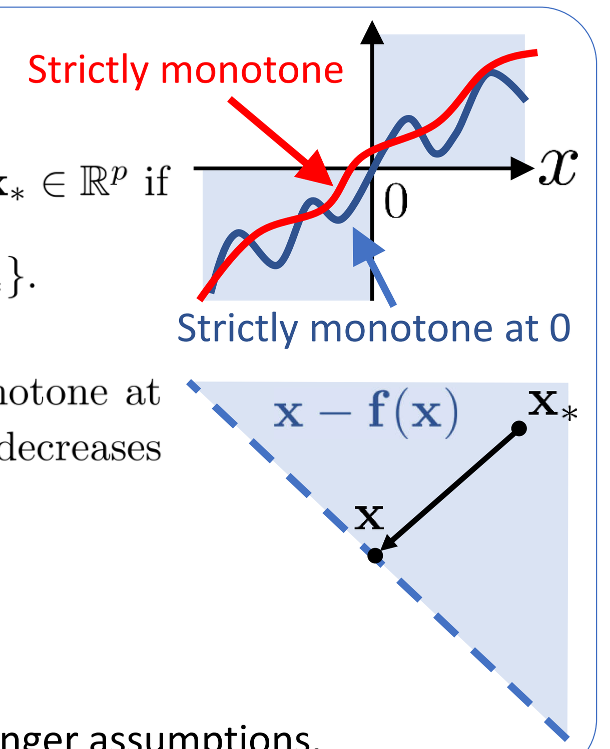

Bonus: Training error converges to zero under stronger assumptions.

Training

Goal: Learn vector fields from different initializations to the solutions

<u>Training data</u>

$\{(\mathbf{x}_0^{(i)},\mathbf{x}_*^{(i)},\mathbf{h}^{(i)})\}_{i=1}^n$ — Initialization: $\mathbf{x}_0^{(i)}$; Ground truth: $\mathbf{x}_*^{(i)}$; Feature function: $\mathbf{h}^{(i)}$

$t:=0$

<u>Learn the maps</u>

$$\mathbf{D}_{t+1}=\arg\min_{\tilde{\mathbf{D}}}\sum_{i=1}^n\|\mathbf{x}_*^{(i)}-\mathbf{x}_t^{(i)}+\tilde{\mathbf{D}}\mathbf{h}^{(i)}(\mathbf{x}_t^{(i)})\|_2^2$$

<u>Update the data</u>

$$\mathbf{x}_{t+1}^{(i)}=\mathbf{x}_t^{(i)}-\mathbf{D}_{t+1}\mathbf{h}^{(i)}(\mathbf{x}_t^{(i)})$$

$t:=t+1$

Terminate? (max #iter / small updates) — No: back to Learn the maps; Yes: end

Solving new problem instance

<u>Input</u>: $\mathbf{x}_0,\mathbf{h},\{\mathbf{D}_t\}_{t=1}^T$

$t:=1$

<u>Update</u>: $\mathbf{x}=\mathbf{x}-\mathbf{D}_t\mathbf{h}(\mathbf{x})$

$t\geq T$? — No: $t:=t+1$, back to Update; Yes: Terminate? (max #iter / small updates) — No: $t:=T$, back to Update; Yes: end

Advantages of DO

- No explicit cost function
- Robust against noise and outliers
- Empirically shown to avoid bad solutions with larger region of convergence
- The $\mathbf{h}$ function does not need to be differentiable/continuous
- Computationally efficient

DO for Point Cloud Registration

Parametrization

- Parametrize $\mathbf{x}$ as the Lie algebra of the rigid transformation
- Function $\mathbf{h}$ as a histogram of local density of scene points in each side of each model point

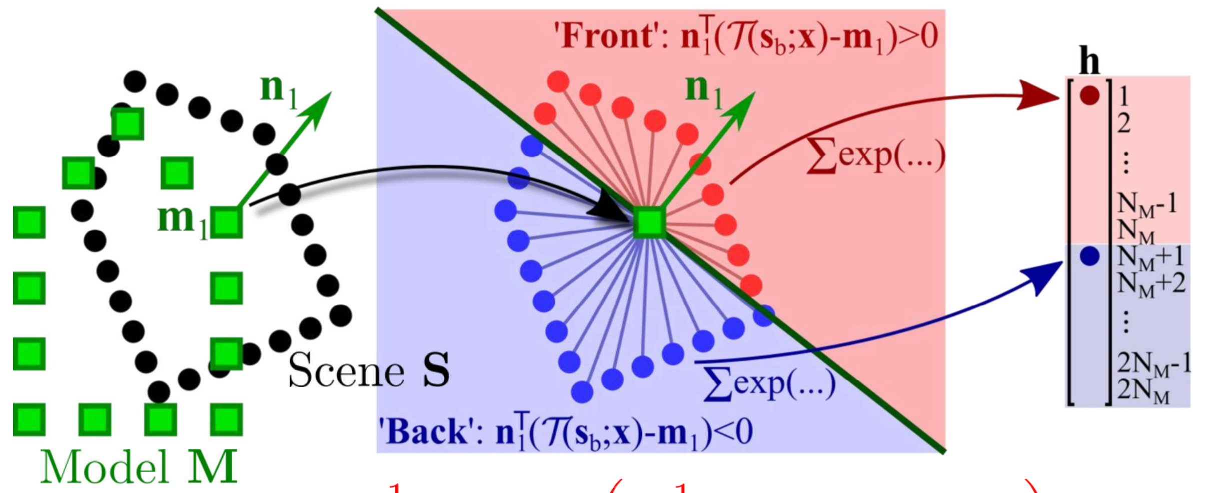

$$[\mathbf{h}(\mathbf{x};\mathbf{S})]_a=\frac{1}{z}\sum_{\mathbf{s}_b\in S_a^+}\exp\left(-\frac{1}{\sigma^2}\|\mathcal{T}(\mathbf{s}_b;\mathbf{x})-\mathbf{m}_a\|^2\right)$$

$$[\mathbf{h}(\mathbf{x};\mathbf{S})]_{a+N_M}=\frac{1}{z}\sum_{\mathbf{s}_b\in S_a^-}\exp\left(-\frac{1}{\sigma^2}\|\mathcal{T}(\mathbf{s}_b;\mathbf{x})-\mathbf{m}_a\|^2\right)$$

Training the maps

Generated synthetically perturbed data from the model and trained 30 maps with $\mathbf{x}_0^{(i)}=\mathbf{0}_6$, and $\mathbf{x}_*^{(i)}=$ ground truth transformation parameters.

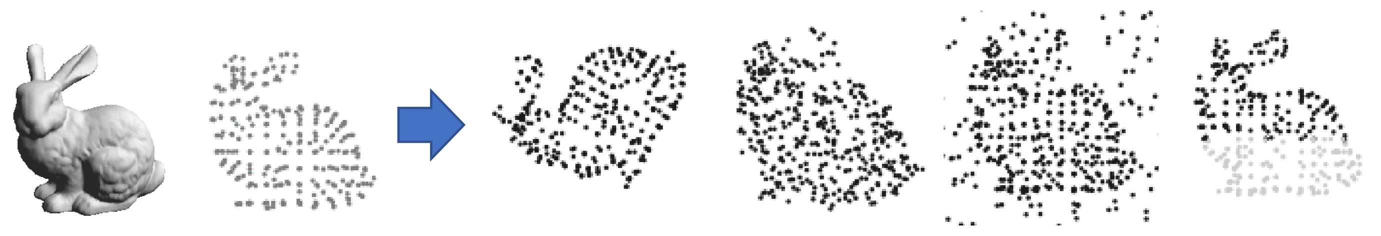

Model | Rotate/translate | Noise | Outliers | Incompl. data

Acknowledgment This research was supported in part by Fundação para a Ciência e a Tecnologia (project FCT [UID/EEA/50009/2013] and a PhD grant from the Carnegie Mellon-Portugal program), the National Science Foundation under the grants RI-1617953, and the EU-Horizon 2020 project #731667 (MULTIDRONE). The content is solely the responsibility of the authors and does not necessarily represent the official views of the funding agencies.

Experiments

Synthetic results

- Stanford bunny, 500 rounds / setting
- With 5 perturbations in the data

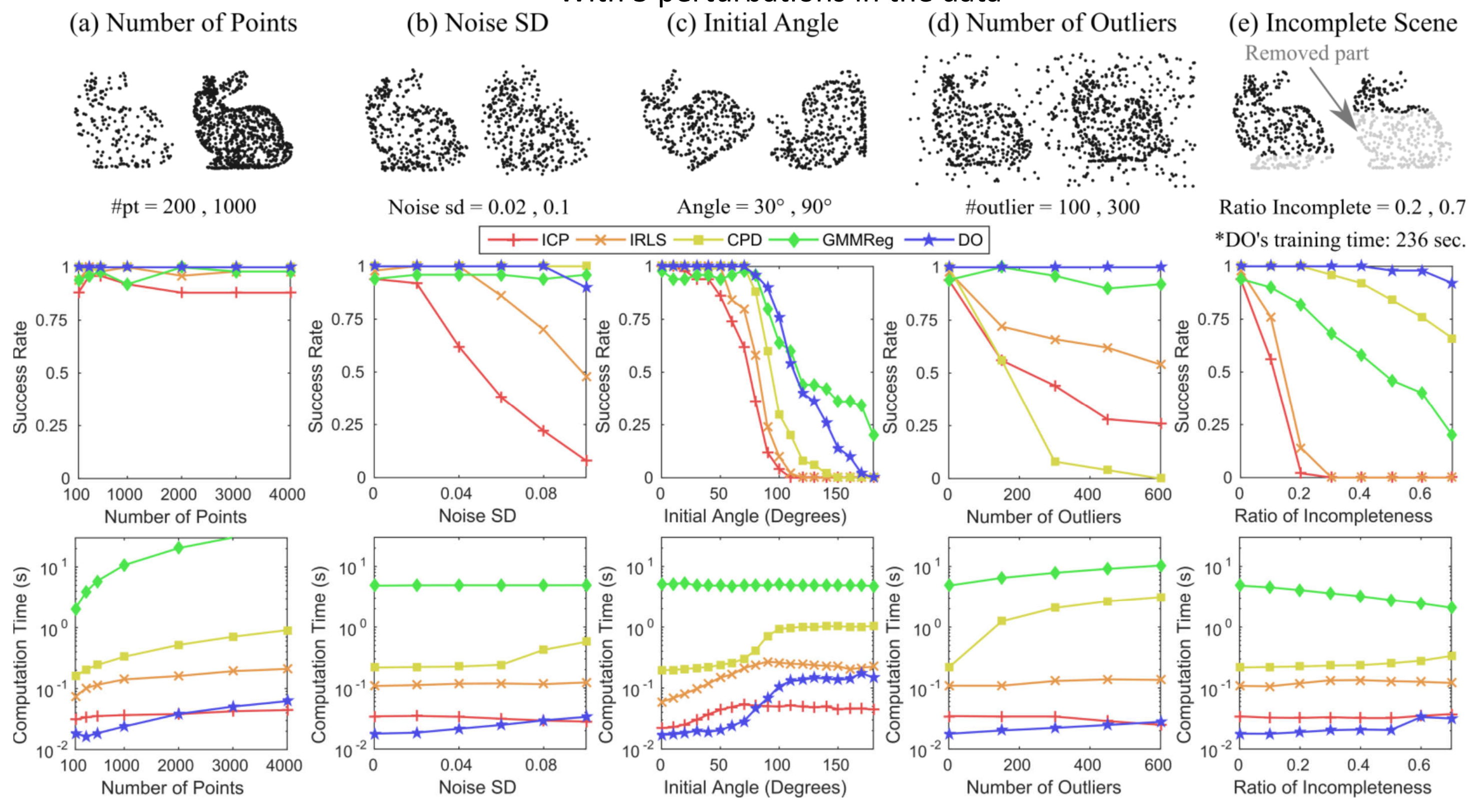

Range-scan results

- UWA [4] dataset, 4 shapes, 50 scenes, 50 rounds / setting
- Structured outliers

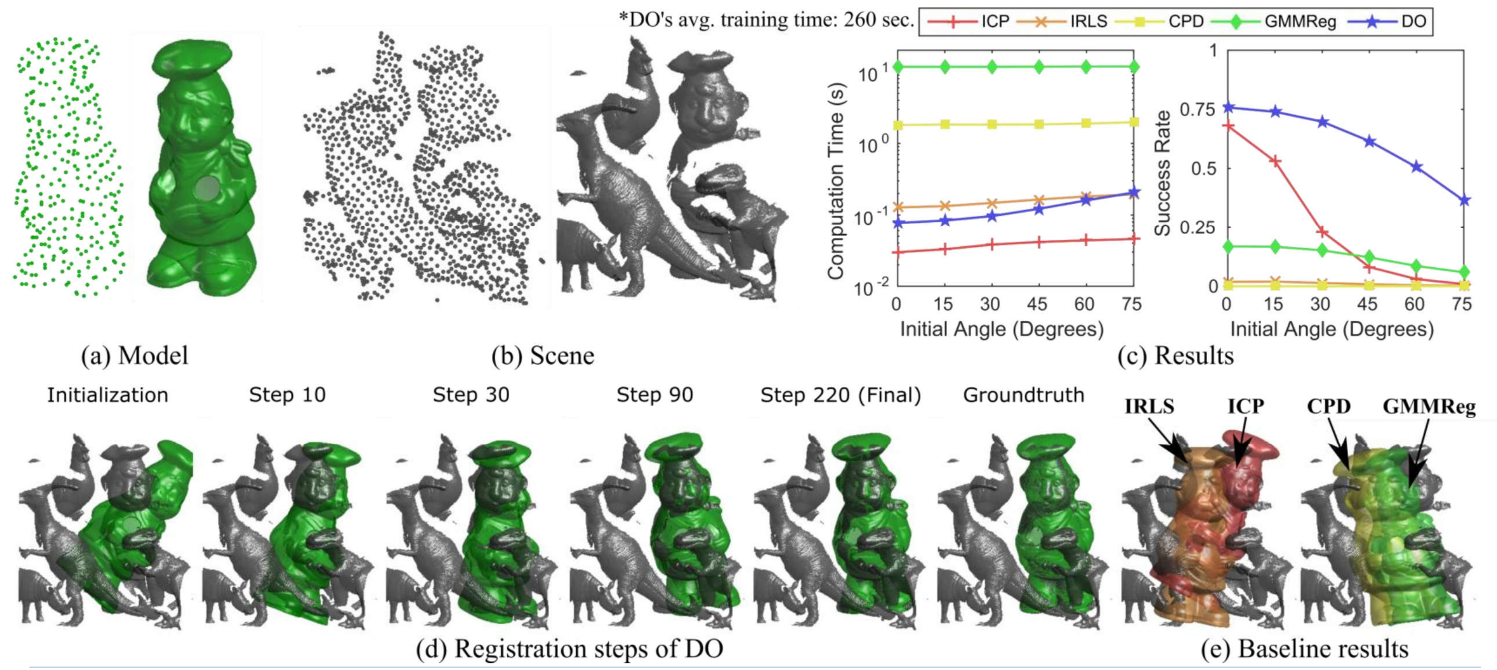

Application: 3D Object tracking

Tracked two shapes on 3D video from Kinect, DO runs at ~40ms/frame

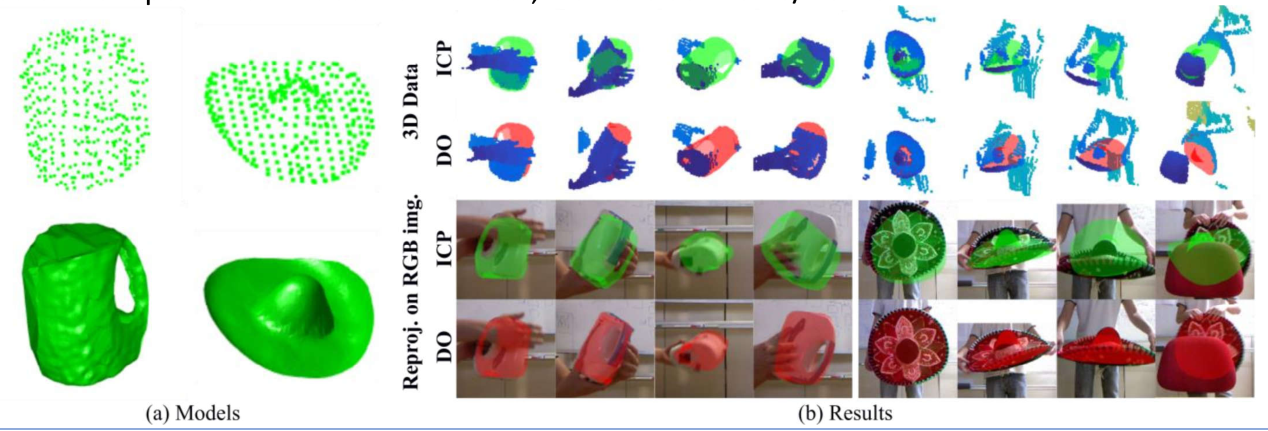

References
[1] P. J. Besl and H. D. McKay, "A method for registration of 3-D shapes," *IEEE TPAMI*, 14(2):239–256, 1992.
[2] P. Dollar, P. Welinder, and P. Perona. "Cascaded pose regression." In *CVPR*, 2010.
[3] X. Xiong and F. De la Torre. "Supervised descent method and its application to face alignment." In *CVPR*, 2013.
[4] A. Mian, M. Bennamoun, and R. Owens. "On the repeatability and quality of keypoints for local feature-based 3D object retrieval from cluttered scenes." *IJCV*, 89(2):348–361, 2010.
[5] B. Jian and B. C. Vemuri. "Robust point set registration using gaussian mixture models." *IEEE TPAMI*, 33(8):1633–1645, 2011 .
[6] A. Myronenko and X. Song. "Point set registration: Coherent point drift." *IEEE TPAMI*, 32(12):2262–2275, 2010.
[7] P. Bergstrom and O. Edlund. "Robust registration of point sets using iteratively reweighted least squares." *Computational Optimization and Applications*, 58(3):543–561, 2014.