Symmetry-Robust 3D Orientation Estimation

Chris Scarvelis
On the job market soon!

Massachusetts Institute of Technology

Backflip AI

David Benhaim
Paul Zhang

The problem.

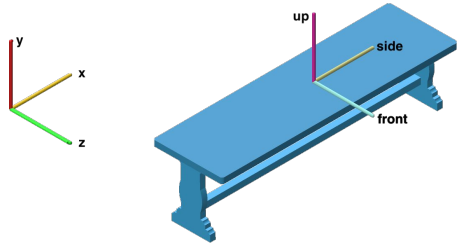

Estimate a shape's *orienter:*
its side, up, and front axes.

Learn an *orienter function:*

$$f : \mathcal{S} \to SO(3)$$

Maps shapes to predicted orientations $\hat{\Omega}_S$

Key challenge:
If a shape has rotational symmetries, then its orienter map is one-to-many – not a function!

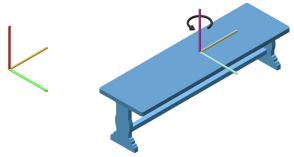

Figure 4: The solution $f^*(S)$ to Problem 2 evaluated at the bench shape S may be any rotation about the y-axis.

One symmetry → infinitely many bad solutions to naive L_2 regression!

Our solution.

$$\min_{f:\mathcal{S}\to SO(3)} \mathop{\mathbb{E}}_{\substack{R\sim U(SO(3)) \\ (S,\Omega_S)\in\mathcal{D}}} \left[\min_{Q\in\hat{\mathcal{R}}_S} \|f(RS) - RQ\Omega_S\|_F^2 \right]$$

<u>First stage: *Quotient regression*</u>
Sends $f(RS) \approx \Omega_{RS}$
up to a symmetry $Q \in \hat{\mathcal{R}}_S$
We quotient by *cube flips*.

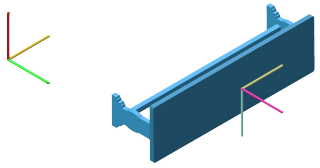

Correctly oriented outputs – *up to a cube flip* $Q^* \in \hat{\mathcal{R}}_S$

$$f^*(RS) = RQ^*\Omega_S$$

<u>Second stage</u>: Find the cube flip by solving a *<u>discrete classification</u>* problem.

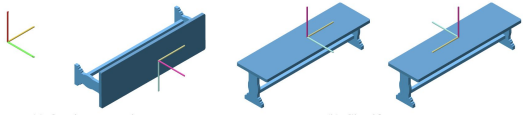

(a) Quotient regression output (b) Classifier outputs

Recovers the correct orientation up to a symmetry of the shape.

Results.

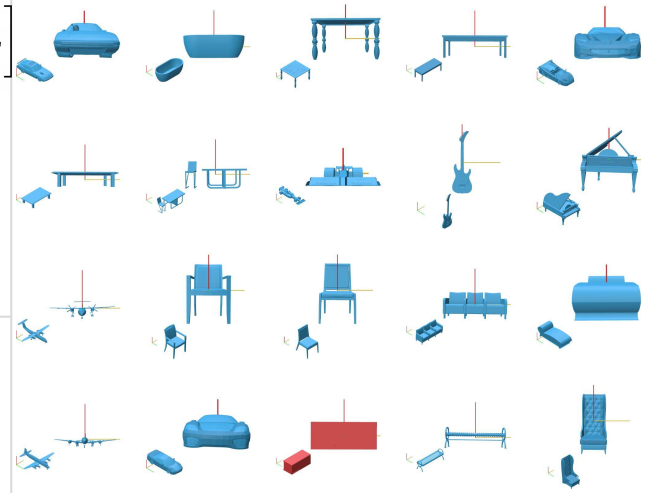

- **Recovers full orientations – upright *and* front-facing.**
- **64.6% lower error rate vs. prior SOTA on up-axis prediction.**
- **Works on *all* of Shapenet – not just select classes.**

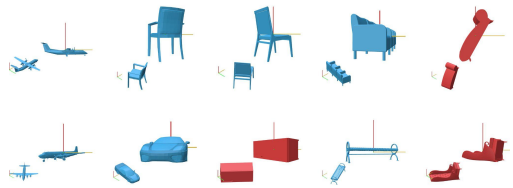

Prior SOTA: Upright-Net (Pang et al. 2022)