Neural Parts: Learning Expressive 3D Shape Abstractions with Invertible Neural Networks

Despoina Paschalidou[1,5,6] Angelos Katharopoulos[3,4] Andreas Geiger[1,2,5] Sanja Fidler[6,7,8]

[1]MPI for Intelligent Systems Tübingen [2]University of Tübingen [3]Idiap [4]EPFL
[5]Max Planck ETH Center for Learning Systems [6]NVIDIA [7]University of Toronto [8]Vector Institute

Motivation

There exists a **trade-off** between **reconstruction quality** and **interpretability** in primitive-based methods.

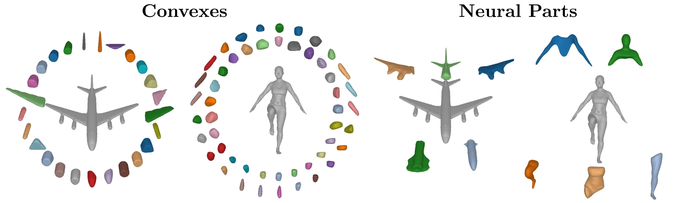

Primitive-based representations rely on **simple shapes** for decomposing complex 3D objects into parts.

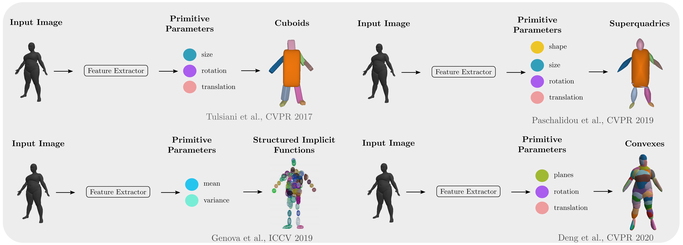

We define primitives via a **deformation between shapes** parametrized as a **learned homeomorphism**.

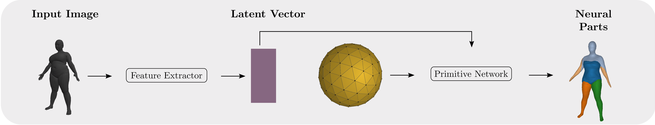

Neural Parts **decouple reconstruction quality from the number of primitives**.

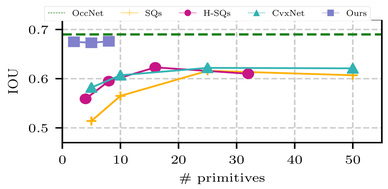

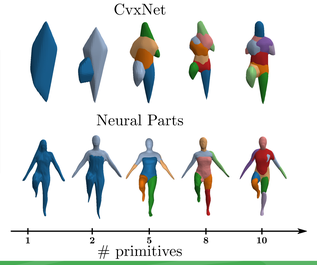

Our Method

A **homeomorphism** $\phi_{\boldsymbol{\theta}} : \mathbb{R}^3 \rightarrow \mathbb{R}^3$ is $\mathbf{x} = \phi_{\boldsymbol{\theta}}(\mathbf{y})$ and $\mathbf{y} = \phi_{\boldsymbol{\theta}}^{-1}(\mathbf{x})$,where $\mathbf{x}$ and $\mathbf{y}$ are 3D points in two topological spaces and $\phi_{\boldsymbol{\theta}} : Y \rightarrow X$, $\phi_{\boldsymbol{\theta}}^{-1} : X \rightarrow Y$ are **continuous bijections**.

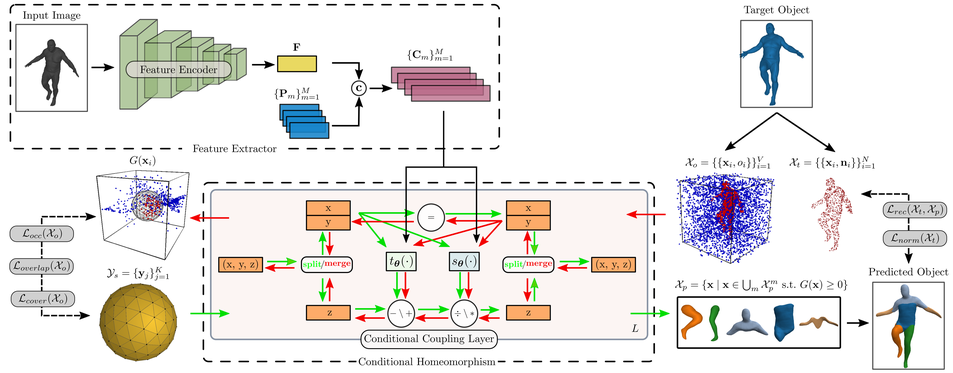

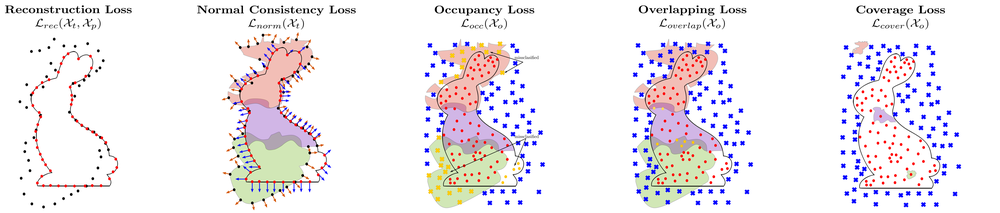

The surface of the predicted and the target surface should match	The normals of the predicted and the target shape should match	The volume of the predicted and the target shape should match	Prevent overlapping primitives	Prevent degenerate primitive arrangements

Is the INN really necessary?

The use of the INN ensures that we always learn **valid homeomorphisms** and allows to impose **constraints on the predicted primitives**.

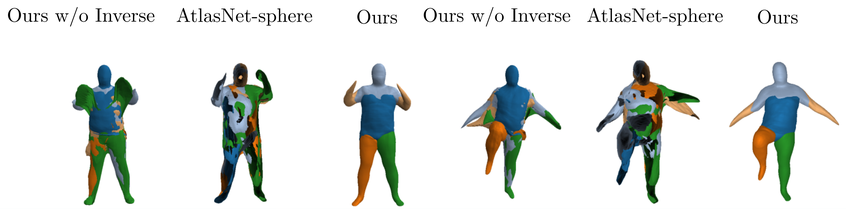

Experiments

Single Image 3D Reconstruction on ShapeNet:

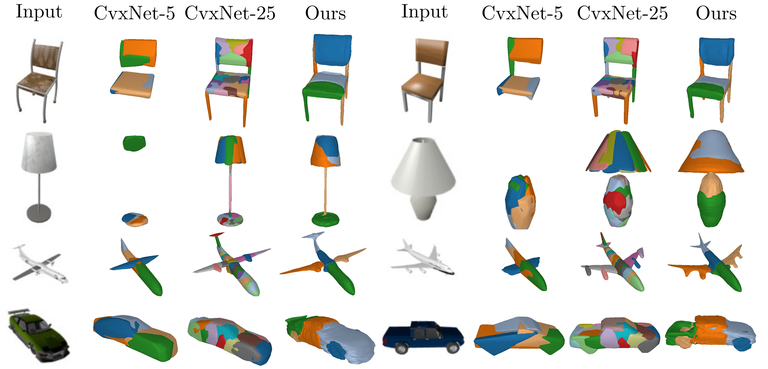

Single Image 3D Reconstruction on D-FAUST and FreiHAND:

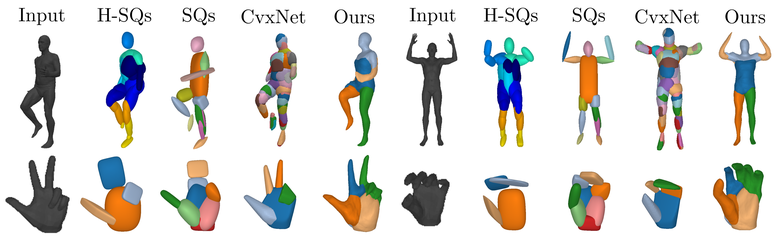

Semantic Consistency:

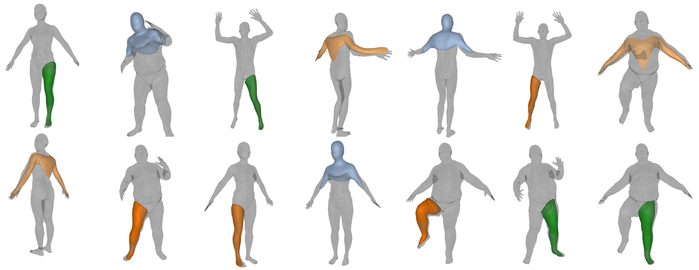

https://paschalidoud.github.io/neural_parts