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]Max Planck Institute for Intelligent Systems Tübingen [2]University of Tübingen
[3]Idiap Research Institute, Switzerland [4]École Polytechique Fédérale de Lausanne (EPFL)
[5]Max Planck ETH Center for Learning Systems [6]NVIDIA [7]University of Toronto [8]Vector Institute

https://paschalidoud.github.io/neural_parts

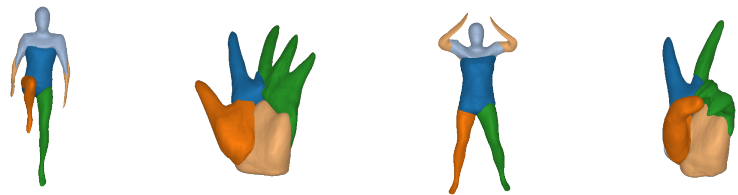

Primitive-based Representations

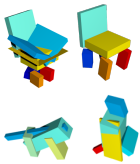

Cuboids
Tulsiani et al. 2017

Primitive-based Representations

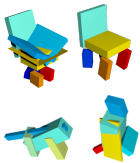

Cuboids
Tulsiani et al. 2017

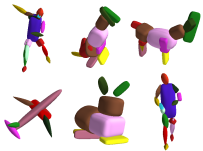

Superquadrics
Paschalidou et al. 2019

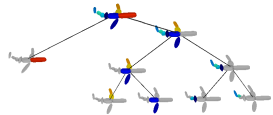

Hierarhical Superquadrics
Paschalidou et al. 2020

Primitive-based Representations

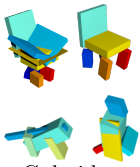

Cuboids
Tulsiani et al. 2017

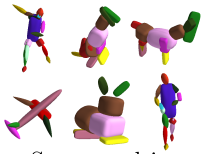

Superquadrics
Paschalidou et al. 2019

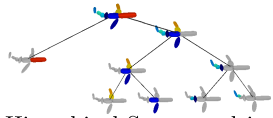

Hierarhical Superquadrics
Paschalidou et al. 2020

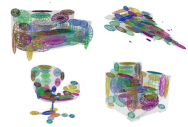

Gaussians
Genova et al. 2019

Primitive-based Representations

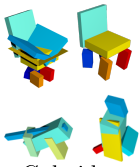

Cuboids
Tulsiani et al. 2017

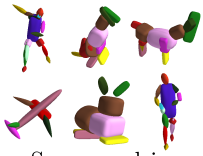

Superquadrics
Paschalidou et al. 2019

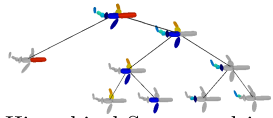

Hierarhical Superquadrics
Paschalidou et al. 2020

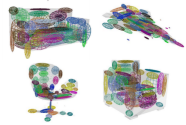

Gaussians
Genova et al. 2019

Convexes
Deng et al. 2020

Primitive-based Representations

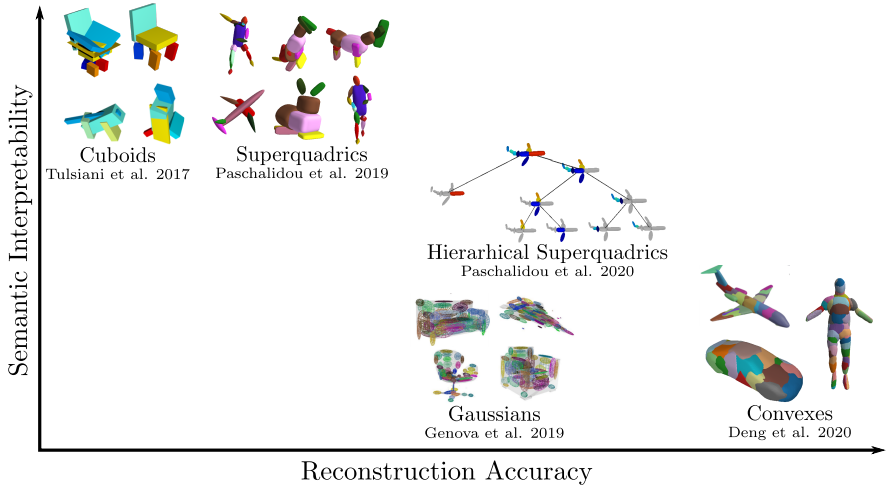

Primitive-based thar rely on simple shapes require a large number of parts for accurate reconstructions.

Primitive-based thar rely on simple shapes require a large number of parts for accurate reconstructions.

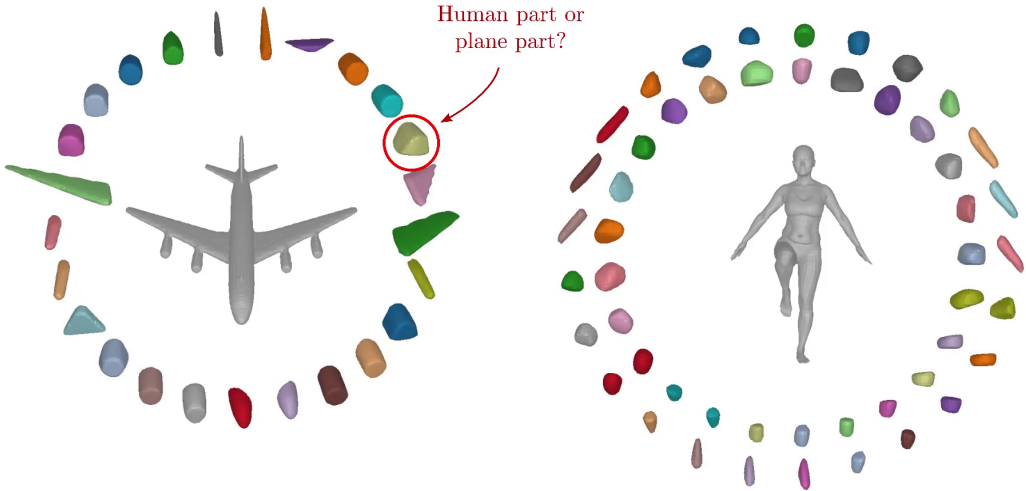

Neural Parts yield accurate and semantic reconstructions using an order of magnitude less parts.

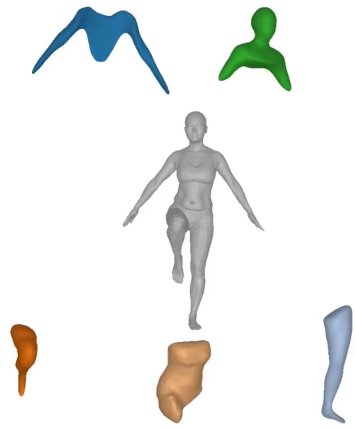

Primitive-based Learning

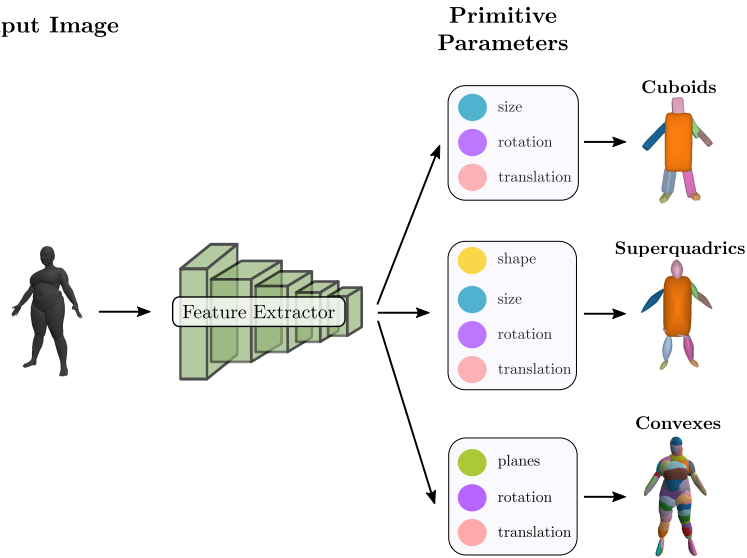

Primitive-based Learning

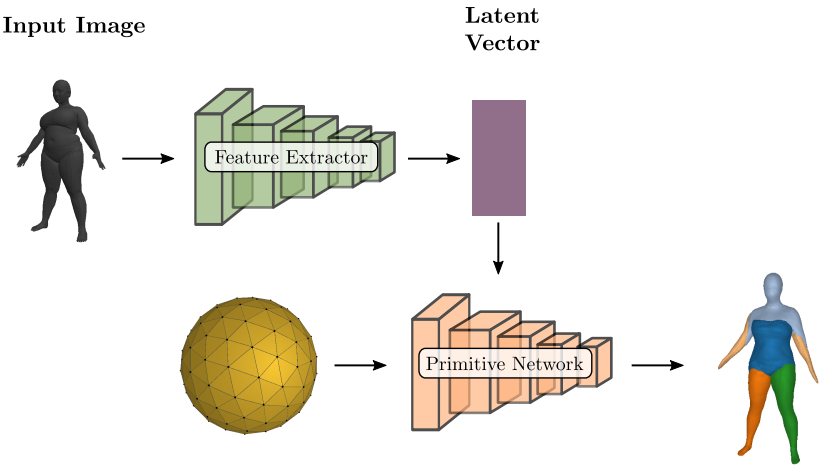

Primitives as Homeomorphic Mappings

A **homeomorphism is a continuous map** between two topological spaces Y and X that preserves all topological properties. In our setup, a homeomorphism $\phi_{\boldsymbol{\theta}} : \mathbb{R}^3 \to \mathbb{R}^3$ is

$$\mathbf{x} = \phi_{\boldsymbol{\theta}}(\mathbf{y}) \text{ and } \mathbf{y} = \phi_{\boldsymbol{\theta}}^{-1}(\mathbf{x})$$

where $\mathbf{x}$ and $\mathbf{y}$ are 3D points in X and Y and $\phi_{\boldsymbol{\theta}} : Y \to X$, $\phi_{\boldsymbol{\theta}}^{-1} : X \to Y$ are continuous bijections.

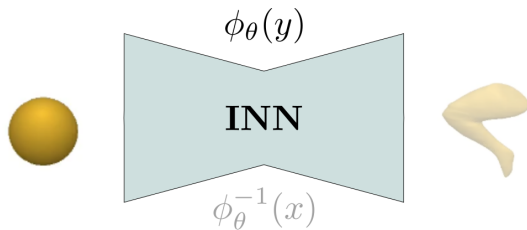

Primitives as Homeomorphic Mappings

A **homeomorphism is a continuous map** between two topological spaces Y and X that preserves all topological properties. In our setup, a homeomorphism $\phi_{\boldsymbol{\theta}} : \mathbb{R}^3 \to \mathbb{R}^3$ is

$$\mathbf{x} = \phi_{\boldsymbol{\theta}}(\mathbf{y}) \text{ and } \mathbf{y} = \phi_{\boldsymbol{\theta}}^{-1}(\mathbf{x})$$

where $\mathbf{x}$ and $\mathbf{y}$ are 3D points in X and Y and $\phi_{\boldsymbol{\theta}} : Y \to X$, $\phi_{\boldsymbol{\theta}}^{-1} : X \to Y$ are continuous bijections.

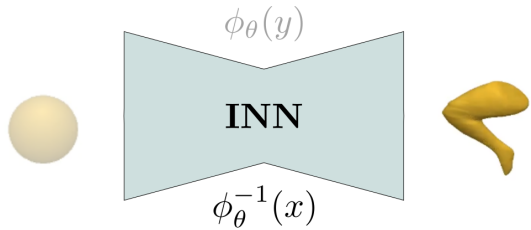

Training Neural Parts

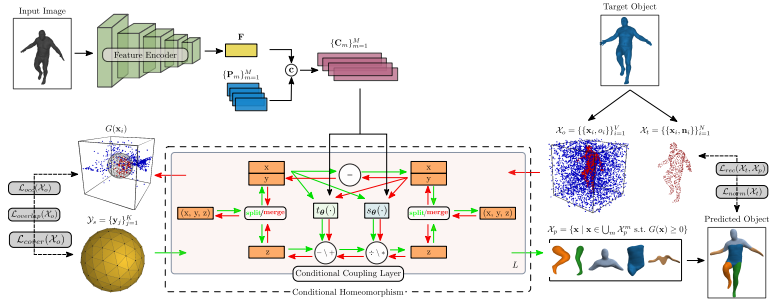

- **Reconstruction Loss**: The surface of the target and the predicted shape should match.
- **Occupancy Loss**: The volume of the target and the predicted shape should match.
- **Normals Consistency Loss**: The normals of the target and the predicted shape should match.
- **Overlapping Loss**: Prevents overlapping primitives.
- **Coverage Loss**: Prevents degenerate primitive arrangements.

Representation Power of Primitive-based Representations

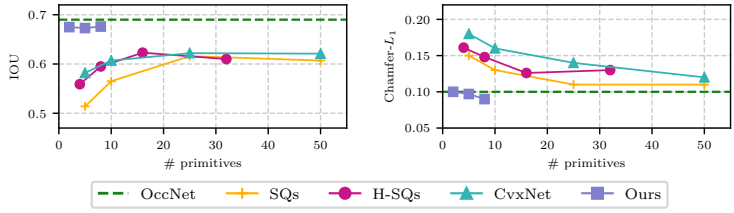

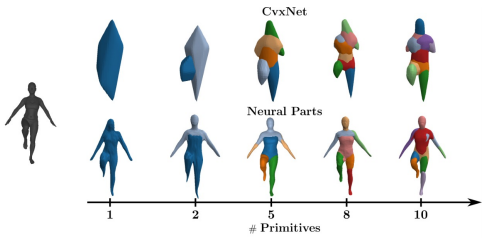

Single-view 3D Reconstruction on D-FAUST

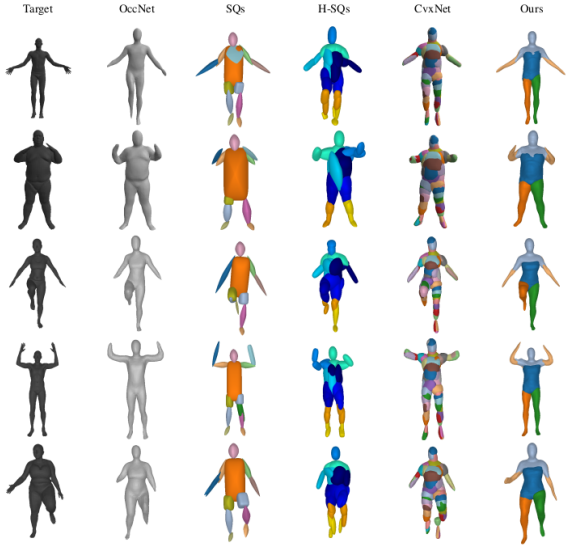

Semantic Consistency of Humans in Motion

Single-view 3D Reconstruction on FreiHAND

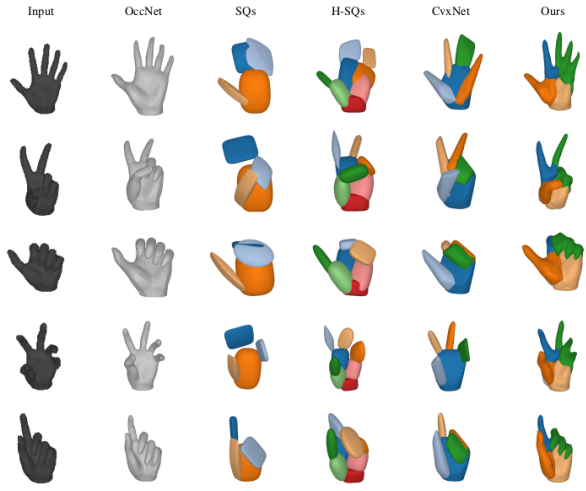

Single-view 3D Reconstruction on ShapeNet

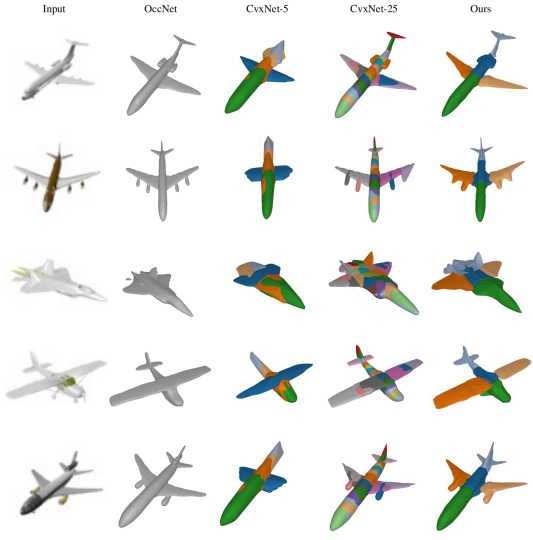

Single-view 3D Reconstruction on ShapeNet

Single-view 3D Reconstruction on ShapeNet

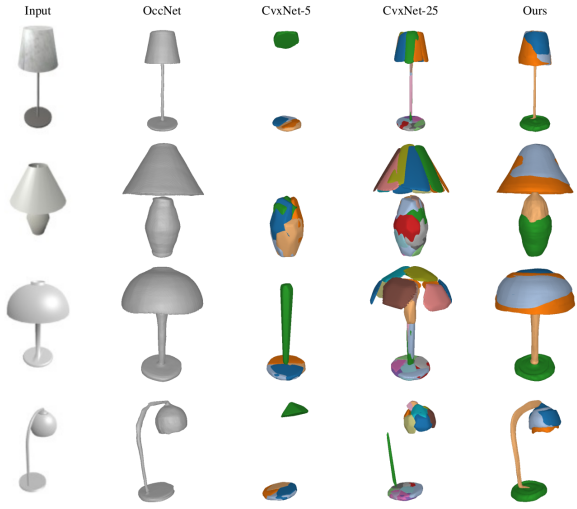

Check our project page for code and additional results!

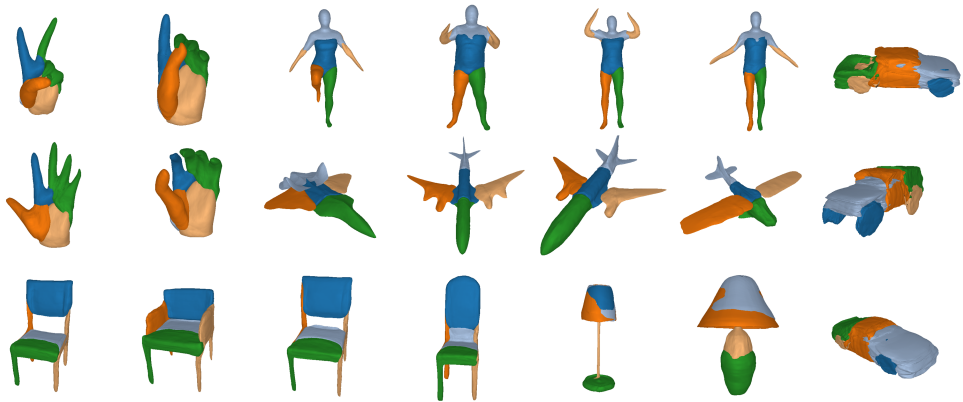

https://paschalidoud.github.io/neural_parts