## Data-driven

## Geometry Processing

Vladimir (Vova) Kim
Adobe Research


## Modeling 3D Assets



## Modeling Garments

## Direct 3D Edit



Physics-driven Pattern Adjustment for Direct 3D Garment Editing. Aric Bartle, Alla Sheffer, Vladimir G. Kim, Danny Kaufman, Nicholas Vining, Floraine Berthouzoz SIGGRAPH 2016

## Modeling Body Parts



## Modeling Body Parts



Customized Software to Optimize Circumferential Pharyngoesophageal Free Flap Reconstruction. Oleksandr Butskiy, Vladimir G. Kim, Brent Chang, Donald Anderson, and Eitan Prisman. Laryngoscope, 2017

Modeling is Difficult


## Why Data-Driven?

## Find desired model in existing dataset



Exploring Collections of 3D Models using Fuzzy Correspondences. V. Kim, W. Li, N. Mitra, Š. Chaudhuri, S. DiVerdi, T. Funkhouser, SIGG'RAPH 2012

## Why Data-Driven?

Create new model from existing parts


## Why Data-Driven?

## Understand context and variations



ShapeSynth: Parameterizing Model Collections for Coupled Shape Exploration and Synthesis. M. Averkiou, V. Kim, Y. Zheng, and N. Mitra, Eurographics 2014

## What is the right "Data"

3D shapes
Categories
Parts
Part Labels

## Machine Learning Pipeline

## Training

Get Data

## Label Data

## Train Model

-     -         -             -                 -                     -                         -                             -                                 -                                     -                                         -                                             -                                                 -                                                     -                                                         -                                                             -                                                                 -                                                                     -                                                                         - 


## Testing

- Label input

Apply Model

- Sample from model -


## Example: Shape Completion

## Use trained model to complete a partial scan



Data-Driven Structural Priors for Shape Completion. Minhyuk Sung, Vladimir G. Kim, Roland Angst, and Leonidas Guibas SIGGRAPH Asia 2015

## Example: Shape Completion

Use trained model to complete a partial scan


Data-Driven Structural Priors for Shape Completion. Minhyuk Sung, Vladimir G. Kim, Roland Angst, and Leonidas Guibas SIGGRAPH Asia 2015

## Example: Shape Completion

Use trained model to complete a partial scan


Labeled Input
Conditional Sample
Data-Driven Structural Priors for Shape Completion. Minhyuk Sung, Vladimir G. Kim, Roland Angst, and Leonidas Guibas SIGGRAPH Asia 2015

## Machine Learning Pipeline

## Training

## Get Data

## Label Data

## Train Mode

Testing

## Label input

Apply Model
Sample from mode

## Part Labeling

## Somebody provided data for you

## Part Labeling

## Somebody provided data for you

Human labeling

- Experts (aka "graduate students")
- Crowd workers


## Part Labeling

## Somebody provided data for you

Human labeling

- Experts (aka "graduate students")
- Crowd workers

ML algorithms

## Part Labeling

## Somebody provided data for you

Human labeling

- Experts (aka "graduate students")
- Crowd workers

ML algorithms

Some combination of above

## Training Data

Where do we get segmented shapes?

- Active learning for labeling [Wang'12]
- Dataset: 1090 meshes, 2-5 parts per mesh, watertight manifold



## Training Data

Where do we get segmented shapes?

- Active learning for labeling [Wang'12]
- Dataset: 1090 meshes, 2-5 parts per mesh, watertight manifold

by a human
does not directly optimize efficiency (active learning objective: train the best classifier)


## Training Data

Where do we get segmented shapes?

- Active learning for labeling [Wang'12]
- Dataset: 1090 meshes, 2-5 parts per mesh, watertight manifold
- Manual annotation [Chen'09]
- Dataset: 380 meshes, 19 parts per mesh, watertight manifold



## Training Data

Where do we get segmented shapes?

- Active learning for labeling [W_
- Dataset: 1090 meshes, 2-5 parts
- Manual annotation [Chen’09]
- Dataset: 380 meshes, 19 parts pe



## Training Data

Where do we get segmented shapes?

- Active learning for labeling [Wang'12]
- Dataset: 1090 meshes, 2-5 parts per mesh, watertight manifold
- Manual annotation [Chen'09]
- Dataset: 380 meshes, 19 parts per mesh, watertight manifold

Verify all


## Training Data

## Where do we get segmented shapes?

- Active learning for labeling [Wang'12]
- Dataset: 1090 mesh
- Manual annotation
- Dataset: 380 meshe


## Verify all

 results!

## Training Data

Where do we get segmented shapes?

- Active learning for labeling [Wang'12]
- Dataset: 1090 meshes, 2-5 parts per mesh, watertight manifold
- Manual annotation [Chen'09]
- Dataset: 380 meshes, 19 parts per mesh, watertight manifold



## Our Approach

## Verify algorithmic predictions



## Label Propagate

## Verify

A Scalable Active Framework for Region Annotation in 3D Shape Collections Li Yi, Vladimir G. Kim, Duygu Ceylan, I-Chao Shen, Mengyan Yan, Hao Su, Cewu Lu, Qixing Huang, Alla Sheffer, and Leonidas Guibas SIGGRAPH Asia 2016

## Our Approach

Verify algorithmic predictions


Which models to label \& which labels to verify?

## Utility Function

several shapes



## Utility Function

several shapes
annotation set $\mathcal{A}$

## Utility Function

label propagation
annotation set $\mathcal{A}$

## Utility Function


annotation set $\mathcal{A}$

## Utility Function


annotation set $\mathcal{A}$

## Utility Function



## Utility Function



## Utility Function



## Utility Function



## Utility Function



## Utility Function



## Utility Function



## Utility Function

$T=t_{\text {annotation }}(\mathcal{A})+t_{\text {verification }}(\mathcal{V},\{q\})$

## Utility Function

## $T=t_{\text {annotation }}(\mathcal{A})+t_{\text {verification }}(\mathcal{V},\{q\})$


$\tau_{\text {annotation }} \approx 30 s /$ label

## Utility Function

## $T=t_{\text {annotation }}(\mathcal{A})+t_{\text {verification }}(\mathcal{V},\{q\})$


$\tau_{\text {annotation }} \approx 30 s /$ label

$\tau_{\text {identification }} \approx 0.3 s /$ label
$\tau_{\text {click }} \approx 1.1 s /$ label

## Utility Function

$N_{\text {good }}$
$T$

## Utility Function

$N_{\text {good }}$

## $T$

$\{q\}$ is not known before selecting $\mathcal{A}$ and $\mathcal{V}$

## Utility Function

$N_{\text {good }}$
$T$
$\{q\}$ is not known before selecting $\mathcal{A}$ and $\mathcal{V}$
Estimate probability and use expectation

## Utility Function

$N_{\text {good }}$

## $T$

$\{q\}$ is not known before selecting $\mathcal{A}$ and $\mathcal{V}$
Estimate probability and use expectation

$$
\frac{\mathbb{E}\left[N_{\text {good }}\right]}{\mathbb{E}[T]}
$$

## Utility Function

$N_{\text {good }}$

## $T$

$\{q\}$ is not known before selecting $\mathcal{A}$ and $\mathcal{V}$
Estimate probability and use expectation

$$
\frac{\mathbb{E}\left[N_{\text {good }}\right]}{\mathbb{E}[T]}
$$

## Utility Function

$$
\underset{\mathcal{A}, \mathcal{V}}{\operatorname{argmin}} \frac{\mathbb{E}\left[N_{\text {good }}\right]}{\mathbb{E}[T]}
$$

## Utility Function

$$
\underset{\mathcal{A}, \mathcal{V}}{\operatorname{argmin}} \frac{\mathbb{E}\left[N_{\text {good }}\right]}{\mathbb{E}[T]}
$$

## Select Annotation Set: Beam Search

Select Verification Set: Greedy Search

## Results

Dataset

- 30,000 shapes
-90,000 parts
car (7496)



## Results

## Dataset

- 30,000 shapes
- 90,000 parts

motorbikes



## Time Saving

Comparison on manifold shapes

- Chen'09 (manual) : about x12 more expensive
- Wang'12 (active) : about x2 more expensive


## Time Saving

Comparison on manifold shapes

- Chen'09 (manual) : about x12 more expensive
-Wang'12 (active) : about x2 more expensive
Comparison on data "in the wild"



## Impact of Data

## Training

## Get Data

## ---------------------

## Testing



## Impact of Data

[1] PointNet: Deep Learning on Point Sets for 3D Classification and Segmentation. Hao Su, Charles Qi, Kaichun Mo, Leonidas Guibas. CVPR 2017
[2] SyncSpecCNN: Synchronized Spectral CNN for 3D Shape Segmentation. Li Yi, Hao Su, Xingwen Guo, Leonidas Guibas. CVPR 2017
[3] 3D Shape Segmentation with Projective Convolutional Networks. Evangelos Kalogerakis, Melinos Averkiou, Subhransu Maji and Siddhartha Chaudhuri. CVPR 2017

## Are we done with labeling?

Smaller than analogues image datasets $\approx 30,000$ vs $\approx 300,000$ segmentations

Geometry varies in resolution (we only capture


## Are we done with labeling?

Smaller than analogues image datasets $\approx 30,000$ vs $\approx 300,000$ segmentations

## Geometry varies in resolution (we only capture



Efficiency

## Part Labeling

## Somebody provided data for you

Human labeling

- Experts (aka "graduate students")
- Crowd workers

ML algorithms

Some combination of above

## Somebody provided data for you


v 图 2008+Bugatti+Veyron+16
v : < skp118>
: <base_b>
:: <body_a>
: <brakelig10>
: < brakelig12>
: < brakelig13>
: <bumper_f05>
: <bumper_f18>
: <bumper_f21>
:: <bumper_r26>
: \lldoor_lef10>
: \lldoor_rig14>
: <engine_a>
:: <engine_b>
: <exhaust_08>
: <exhaust_10>
: <fender_f01>
: <fender_f05>
: <hood_c01>
: <sidemirr01>
: \llsidemirrO2>
: <sidemirrO3>
: <sidemirr04>
:: <sidemirr07>
: <sidemirror>
: <skirt_le22>
: <skirt_ri22>
: <trunk_a>
-1 surindan f 0 ?

## Machine Learning Pipeline

## Training

Get Data

## Label Data



Label input
Sample from model

## Challenge

Input scene graphs are messy and inconsistent


## Scene Graphs

Input data is messy and inconsistent


## Our Approach

## Training (robust to inconsistencies)

## Get Data



Train Model (Part-Based Analysis, Mesh
Segmentaiton)

## Testing

## Apply Model

(Mesh
Segmentation)

## Our Approach

## Training (robust to inconsistencies)

## Train Model <br> (Part-Based Analysis, Mesh Segmentaiton)

Learning Hierarchical Shape Segmentation and Labeling from Online Repositories. Li Yi, Leonidas Guibas, Aaron Hertzmann, Vladimir G. Kim, Hao Su, and Ersin Yumer SIGGRAPH 2017

## Objective Function

$$
E(\theta, p, c, M)=\lambda_{c} E_{c}+\lambda_{s} E_{s}+\lambda_{d} E_{d}+\lambda_{m} E_{m}-H
$$



## Objective Function

$E(\theta, p, c, M)=\lambda_{c} E_{c}+\lambda_{s} E_{s}+\lambda_{d} E_{d}+\lambda_{m} E_{m}-H$
Embedding of every node in a scene graph


## Objective Function

$$
E(\theta, p, c, M)=\lambda_{c} E_{c}+\lambda_{s} E_{s}+\lambda_{d} E_{d}+\lambda_{m} E_{m}-H
$$

Clustering: labels


## Objective Function

$$
E(\theta, \underline{p, c}, M)=\lambda_{c} E_{c}+\lambda_{s} E_{s}+\lambda_{d} E_{d}+\lambda_{m} E_{m}-H
$$

Clustering: labels, centroids


## Objective Function

$E(\theta, p, c, M)=\lambda_{c} E_{c}+\lambda_{s} E_{s}+\lambda_{d} E_{d}+\lambda_{m} E_{m}-H$

## Objective Function

$$
\begin{gathered}
E(\theta, p, c, M)=\lambda_{c} E_{c}+\lambda_{s} E_{s}+\lambda_{d} E_{d}+\lambda_{m} E_{m}-H \\
E_{c}=\sum_{\text {parts centroids }} \sum_{\text {part,centroid } \| f(\text { part })-\text { centroid } \|}
\end{gathered}
$$



## Objective Function



## Objective Function

$$
\begin{aligned}
& E(\theta, p, c, M)=\lambda_{c} E_{c}+\lambda_{s} E_{s}+\lambda_{d} E_{d}+\lambda_{m} E_{m}-H \\
& E_{s}=\sum_{\text {part } 1 \text { part 2 }} \sum_{i} \| f(\text { part 1) }-f(\text { part 2)\| iff almost identical } \\
& \text { or tags are the same }
\end{aligned}
$$



## Objective Function

$$
\begin{aligned}
& E(\theta, p, c, M)=\lambda_{c} E_{c}+\lambda_{s} E_{s}+\lambda_{d} E_{d}+\lambda_{m} E_{m}-H \\
& E_{s}=\sum_{\text {part }} \sum_{\text {part 2 }} \frac{\| f(\text { part 1)-f(part 2)\| iff almost identical }}{} \text { or tags are the same }
\end{aligned}
$$



## Objective Function

$$
\begin{aligned}
& E(\theta, p, c, M)=\lambda_{c} E_{c}+\lambda_{s} E_{s}+\lambda_{d} E_{d}+\lambda_{m} E_{m}-H \\
& E_{d}=\sum_{\text {part 1 part 2 }} \max (0, \sigma-\| f(\text { part 1)-f(part 2)\|} \\
& \text { iff on the same shape } \\
& \text { but tags are different }
\end{aligned}
$$

## Objective Function



## Objective Function

$$
E(\theta, p, c, M)=\lambda_{c} E_{c}+\lambda_{s} E_{s}+\lambda_{d} E_{d}+\lambda_{m} E_{m}-H
$$

Parent-child relationships \& regularization

EM optimization:
E step: optimize for $p$.
M step: optimize for $\theta, c, M$

## Part-based Analysis Results



## Hierarchical Mesh Segmentation

What if part was not in a separate node?

- NN to predict per-face labels (use part-base analysis tags)
- Preserve parent-child relationships
- MRF to get smooth boundaries
- Use connected components if available


## Results



## Results



## Evaluation

Clustering (Normalized Mutual Information $\leq 1$ )

- Chance $=0.034$
- Features $=0.348$
- Ours $=0.573$

Part Tagging (accuracy $\leq 1$ )

- Chance $=0.139$
- Features $=0.823$
- Ours $=0.910$


## Crowd vs Ambient Data



## Crowd vs Ambient Data



## Crowd vs Ambient Data

Not all ambient data is equally useful


## Use Case

## Fe

## Project Felix

## Summary

Data is as important as algorithms

Make data collection easier

- Use all available meta-data
- Crowd-source as necessary
- Automatically propagate and then manually verify labels


## Future Work

## Training

Get Data

## Label Data

## Train Model

## T -4



## Deep Learning for Geometry

Geometry processing as neural network layers


## Deep Learning for Geometry

## Geometry processing as neural network layers



Convolutional Neural Networks on Surfaces via Seamless Toric Covers. Haggai Maron, Meirav Galun, Noam Aigerman, Miri Trope, Nadav Dym, Ersin Yumer, Vladimir G. Kim, and Yaron Lipman. SIGGRAPH 2017

## Deep Learning for Geometry

Geometry processing as neural network layers

- Make input more regular:


## Parameterization

(translation-invariant, seamless)
unstructure geometry $\left\langle x_{1}, y_{1}, z_{1}\right\rangle$


Tutte


GIM


Seamless


Ours

Convolutional iveural ive works on surraces via seanmess ionic Lovers. naggai iviaron, Meirav Galun, Noam Aigerman, Miri Trope, Nadav Dym, Ersin Yumer, Vladimir G. Kim, and

## Deep Learning for Geometry

Geometry processing as neural network layers

- Make input more regular: Parameterization
- Make input more consistent:

Correspondence

- Learn how to parameterize and
optimize for consistency
Parametric Layer
Regular input

Convolutional Neural Networks on Surfaces via Seamless Toric Covers. Haggai Maron, Meirav Galun, Noam Aigerman, Miri Trope, Nadav Dym, Ersin Yumer, Vladimir G. Kim, and Yaron Lipman. SIGGRAPH 2017

## Deep Learning for Geometry

## Geometry processing as neural network layers

- Generative Models Geometric Modeling
- Loss Function

Shape Metric Spaces


Regular input


NN


Re-synthesis
Rendering

## Deep Learning for Geometry

Geometry processing as neural network layers


# Acknowledgements 

Two main papers:


Li Yi

Applications:


Future Work:


Other contributors (alphabetical)

Noam Aigerman
Roland Angst
Floraine Berthouzoz
Duygu Ceylan
Stephen DiVerdi
Nadav Dym
Thomas Funkhouser
Leonidas Guibas

Meirav Galun
Aaron Hertzmann
Qixing Huang
Danny Kaufman
Yaron Lipman
Wilmot Li
Cewu Lu
Niloy J. Mitra

Hao Su
Alla Sheffer
I-Chao Shen
Miri Trope
Nicholas Vining
Mengyan Yan

Ersin Yumer



# Acknowledgements 

Two main papers:


Applications:


Future Work:


Other contributors (alphabetical) Interns @ Adobe

Noam Aigerman
Roland Angst
Floraine Berthouzoz
Duygu Ceylan
Stephen DiVerdi
Nadav Dym
Thomas Funkhouser
Leonidas Guibas

Meirav Galun
Aaron Hertzmann
Qixing Huang
Danny Kaufman
Yaron Lipman
Wilmot Li
Cewu Lu
Niloy J. Mitra

Hao Su
Alla Sheffer
I-Chao Shen
Miri Trope
Nicholas Vining
Mengyan Yan


Ersin Yumer
Researchers @ Adobe

