Modeling and Analyzing 3D Shapes using Clues from 2D Images

Minglun Gong

Dept. of CS, Memorial Univ.

- **Modeling Flowers from Images**
  - Reconstructing 3D flower petal shapes from a single image is difficult
  - Embedding priori knowledge into the 3D reconstruction process helps
- **Labeling 3D Shapes using Images**
  - Existing shape retrieval approaches do not give us examples with similar topology
  - Customized approach can better utilize available training data

Two Recent Projects

Part I

Modeling Flowers from Images

- **Flowers are ubiquitous but not easy to model**
  - **Complex structures**
  - **Complex geometries**

Flower Modeling

- **Automatic modeling w/ botanical knowledge**
  - Floral diagram & inflorescence
  - May not look real
- **Interactive modeling w/ traditional interface**
  - Only limited primitive shapes can be used
  - Tedious work
  - Requires expertise

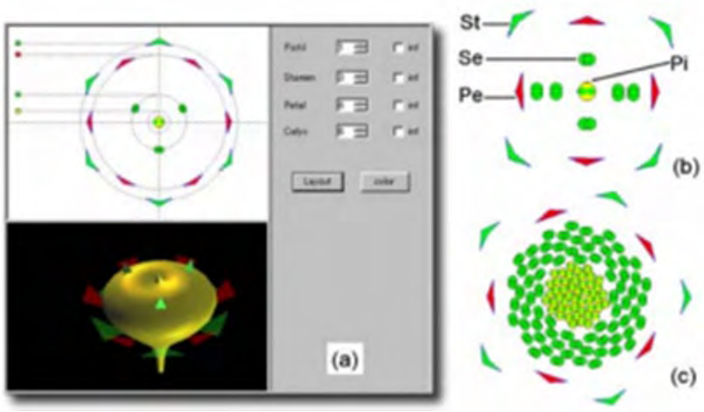

[Ijiri et al. 2005]

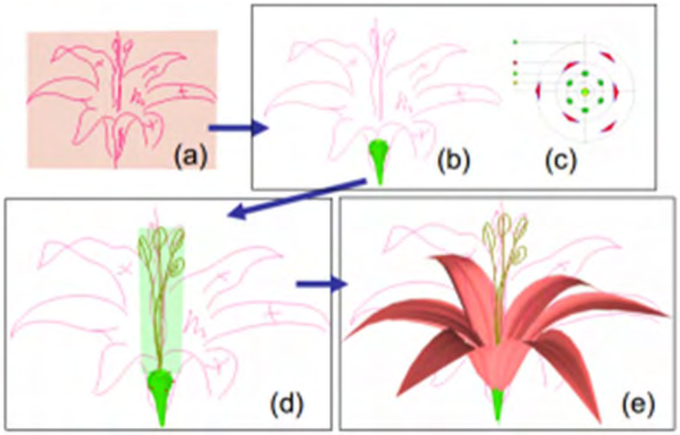

[Ijiri et al. 2006]

Existing Modeling Techniques

- **Modelers often refer to photos when modeling**
  - **Photos remain the cheap and easy means to capture the reality**
  - **Multiple photos of the same flower are not always available**
  - **Users have to estimate shape in a “modeling-and-checking” fashion**

Modeling from Images

Motivations

Allows users to directly interact with input photo to reconstruct the 3D flower shape

Input

User Interaction

Result

Key Observations

- **Petals from a single flower usually share a similar shape**
- **They repeat around the flower central axis , roughly forming a surface of revolution**
- **Their observed shapes in 2D images vary, due to differences in projecting directions**

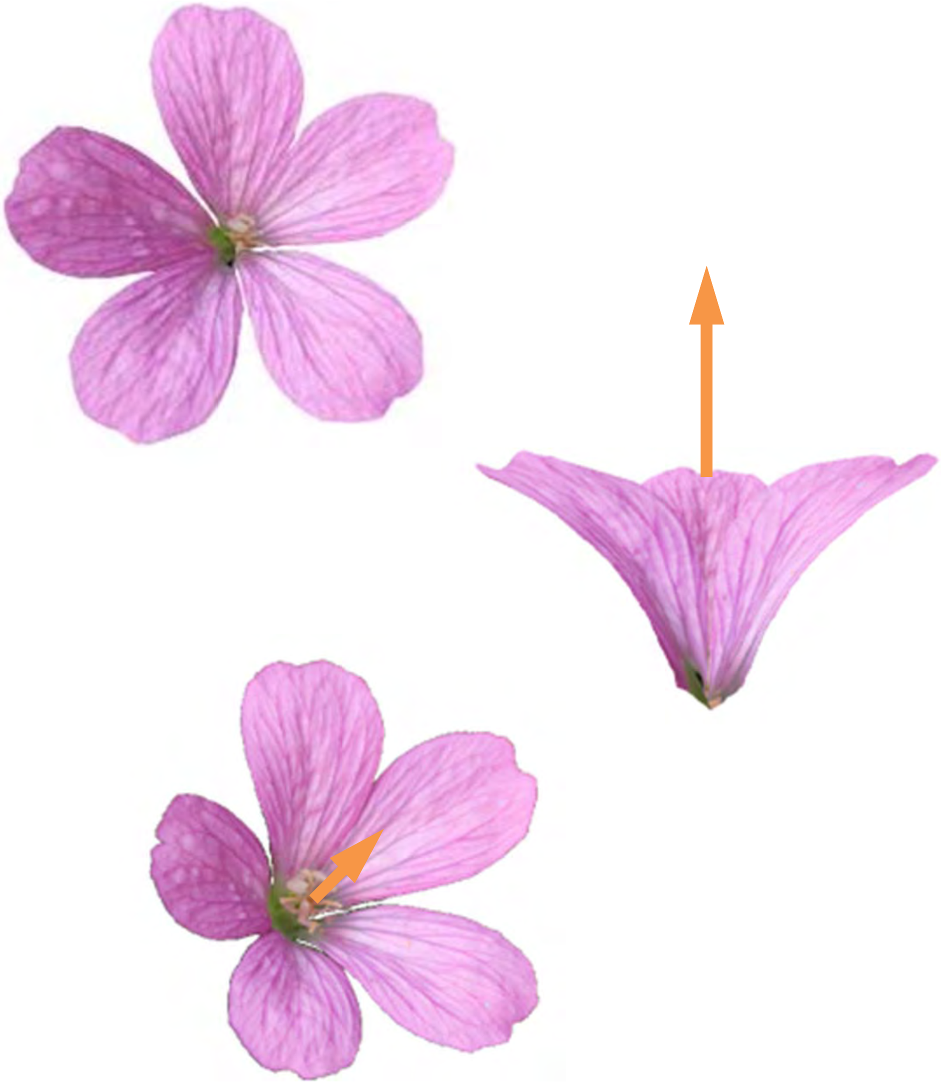

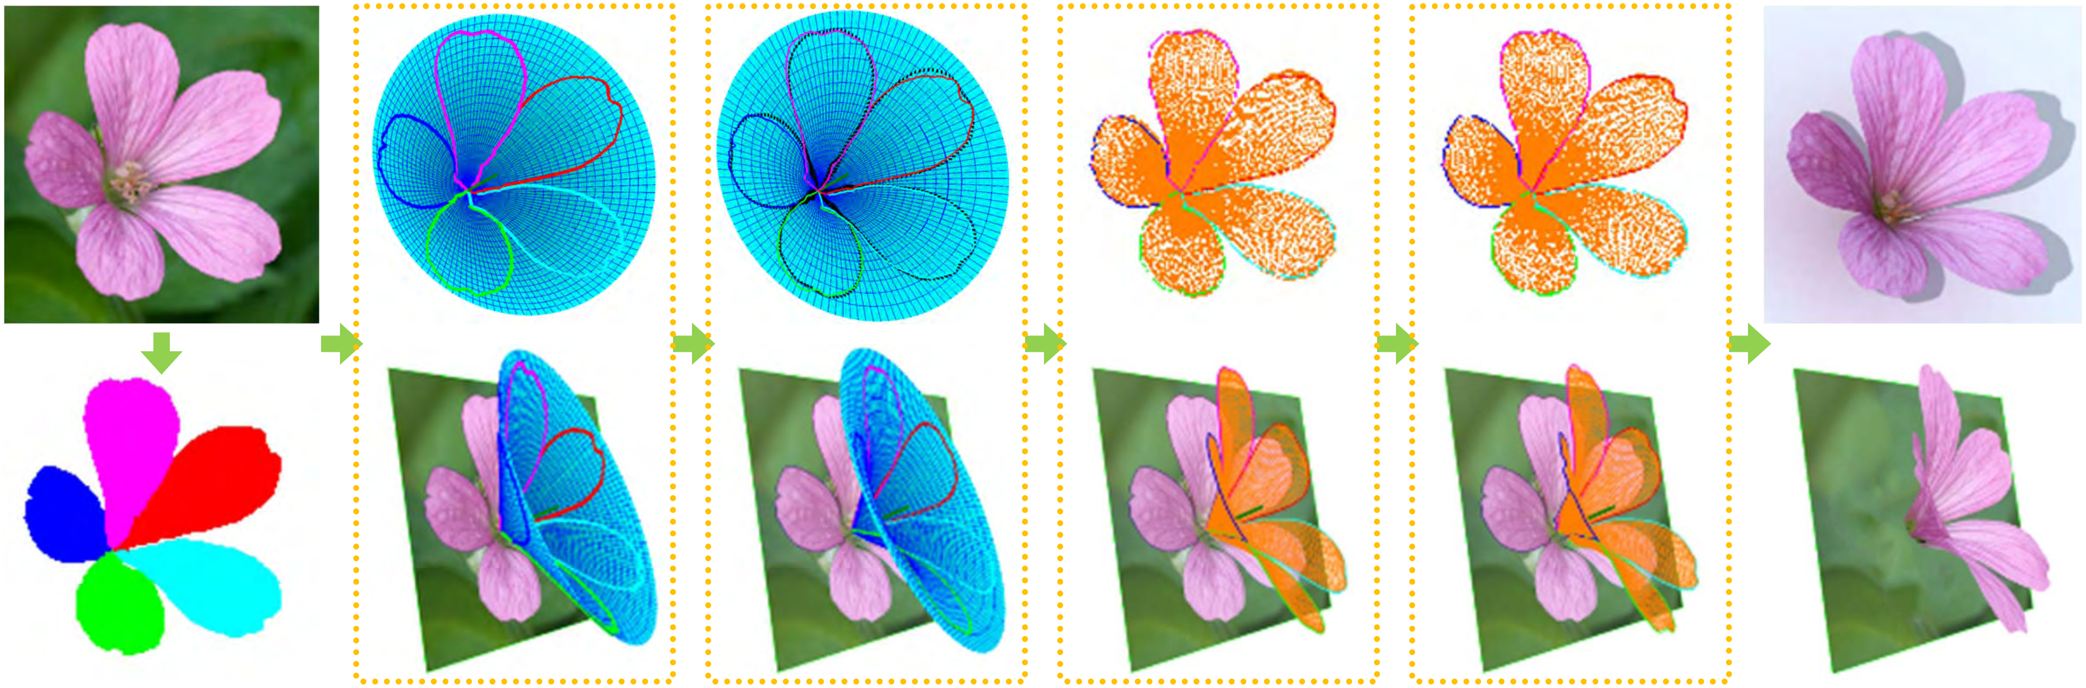

(a) pre-processing | (b) cone fitting | (c) surface fitting | (d) joint fitting | (e) individual fitting | (e) texture mapping

Algorithm Pipeline

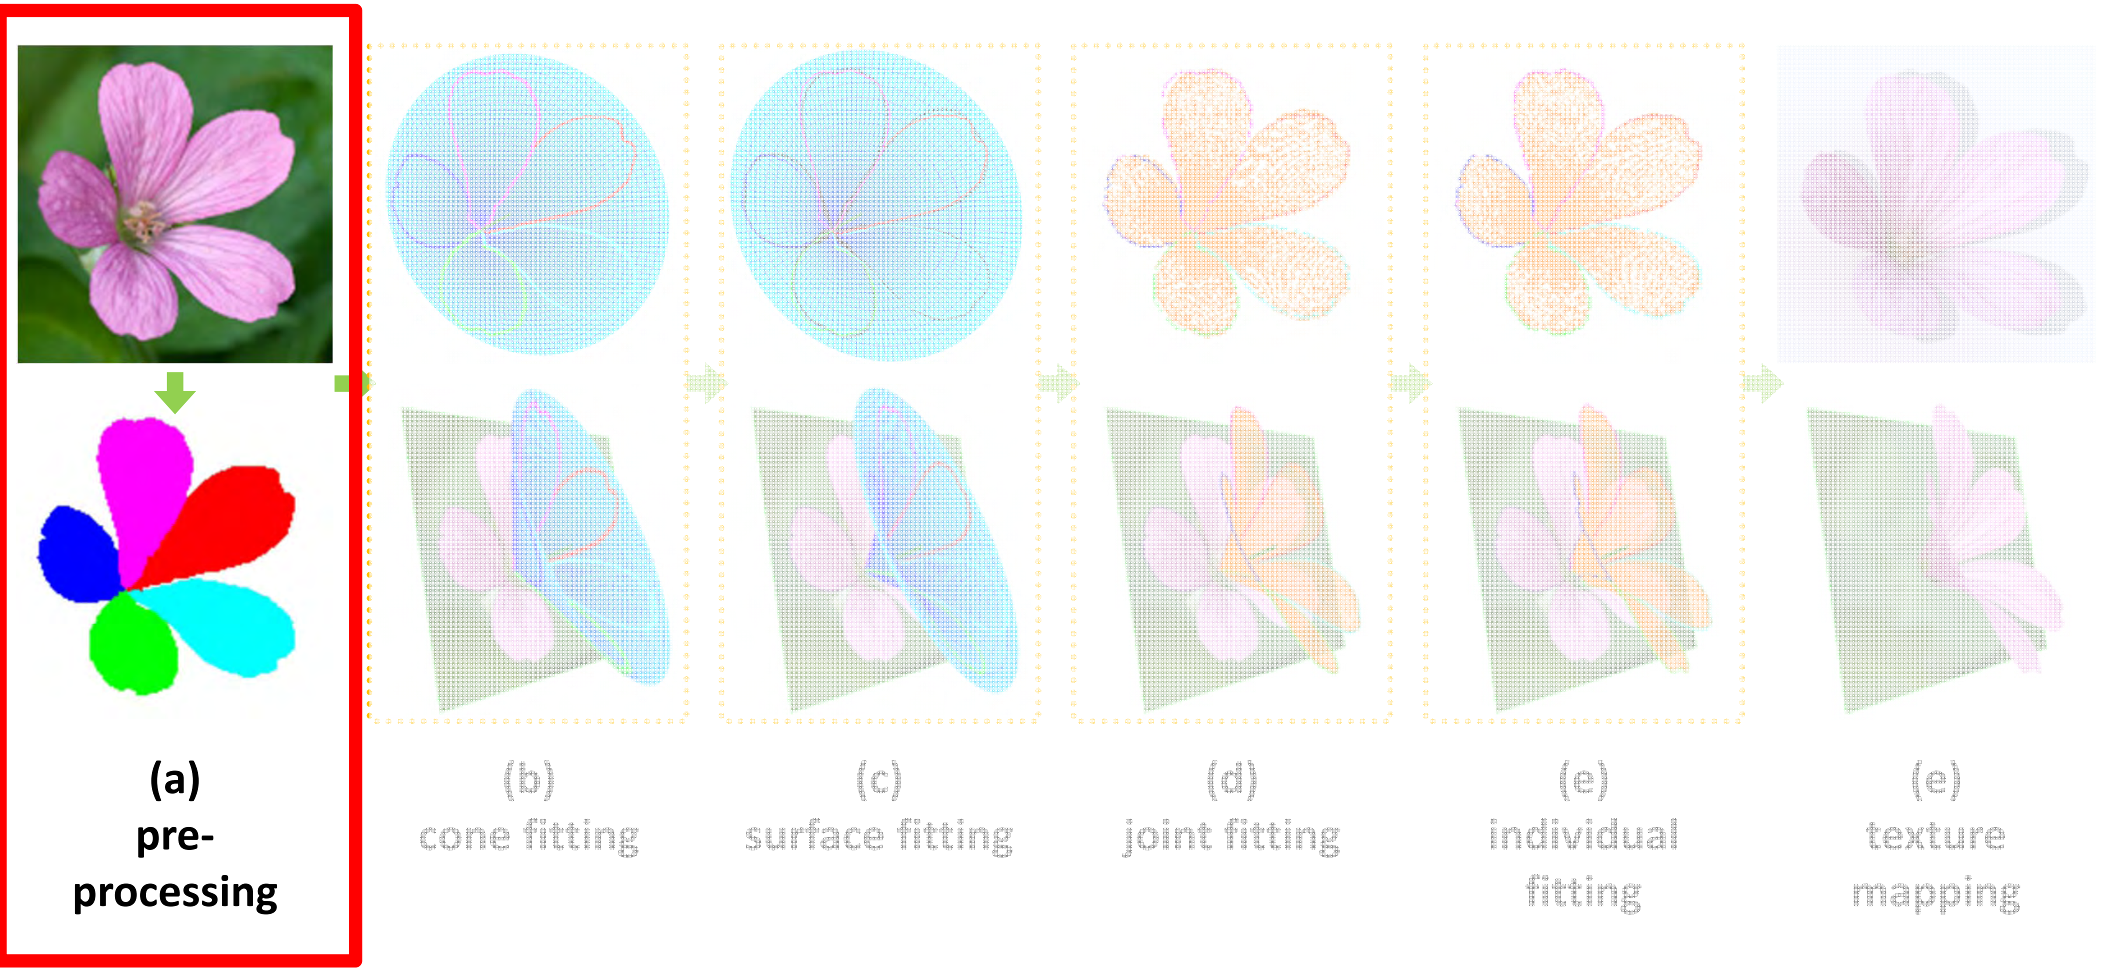

Preprocessing

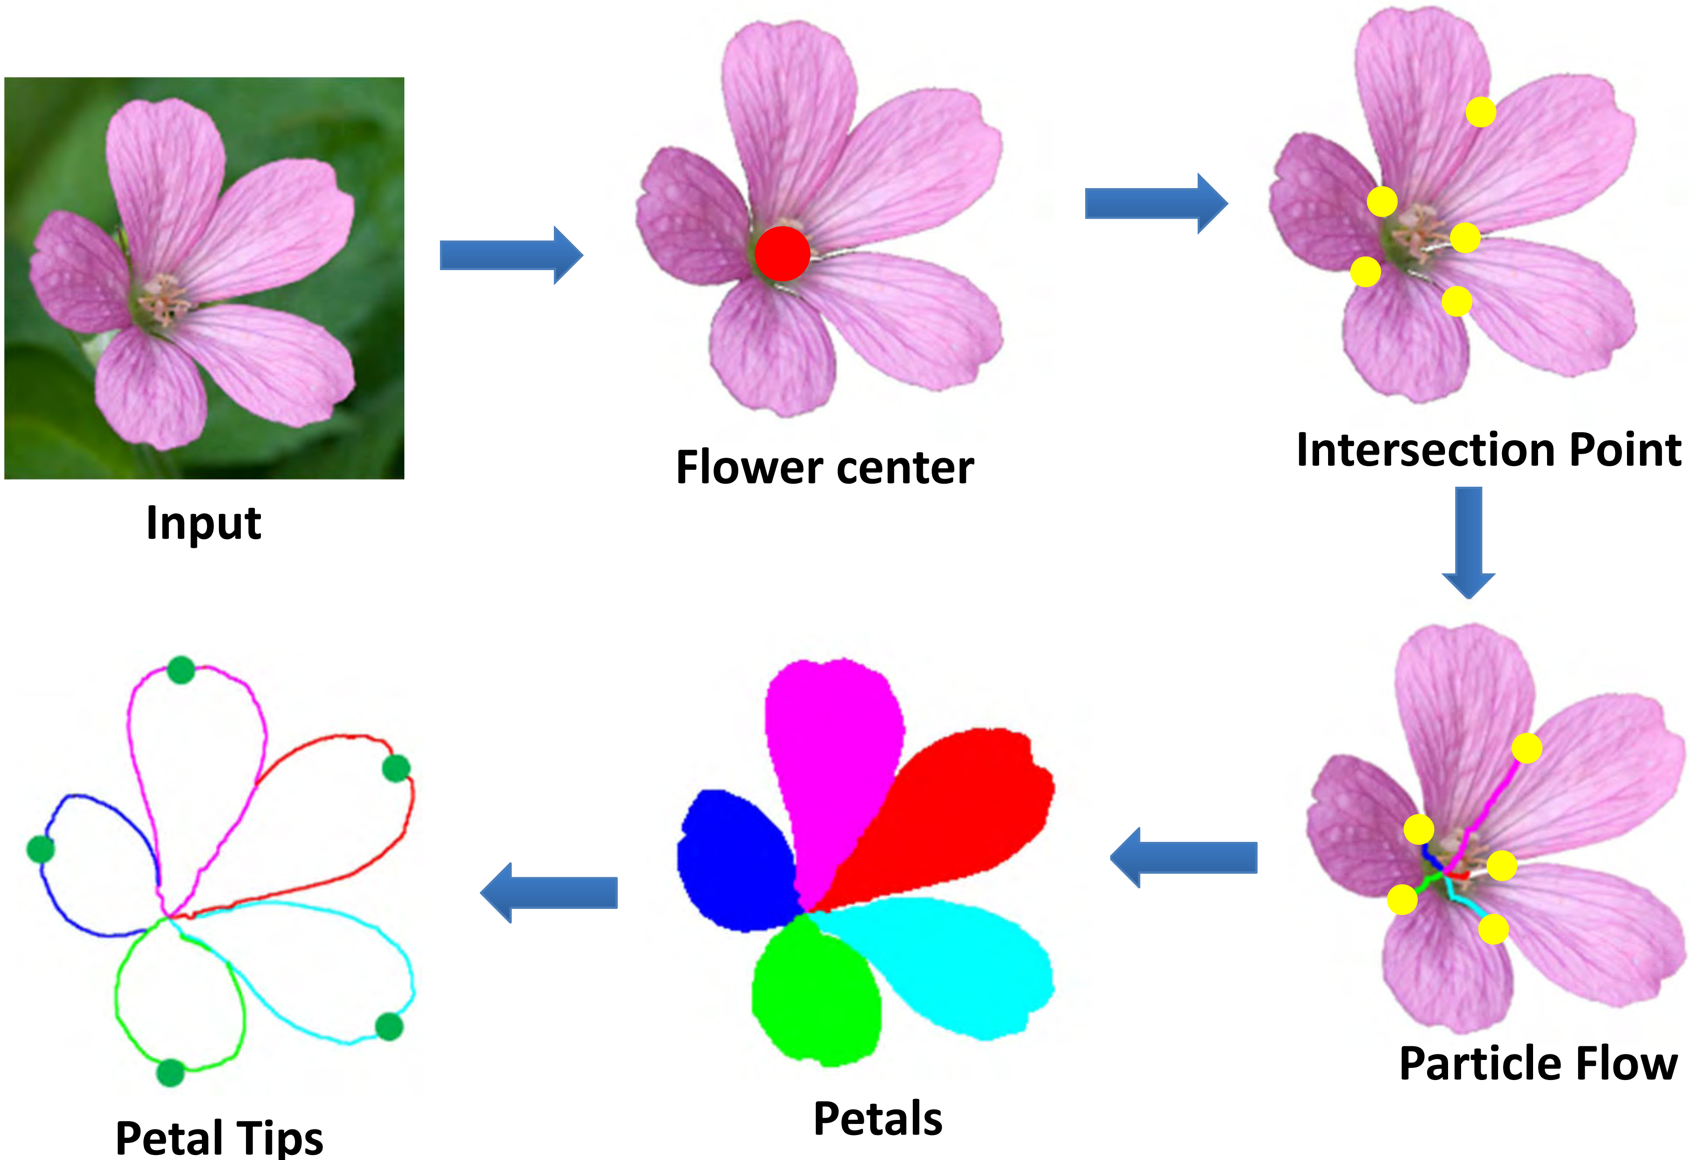

Locate Center & Segment Petals

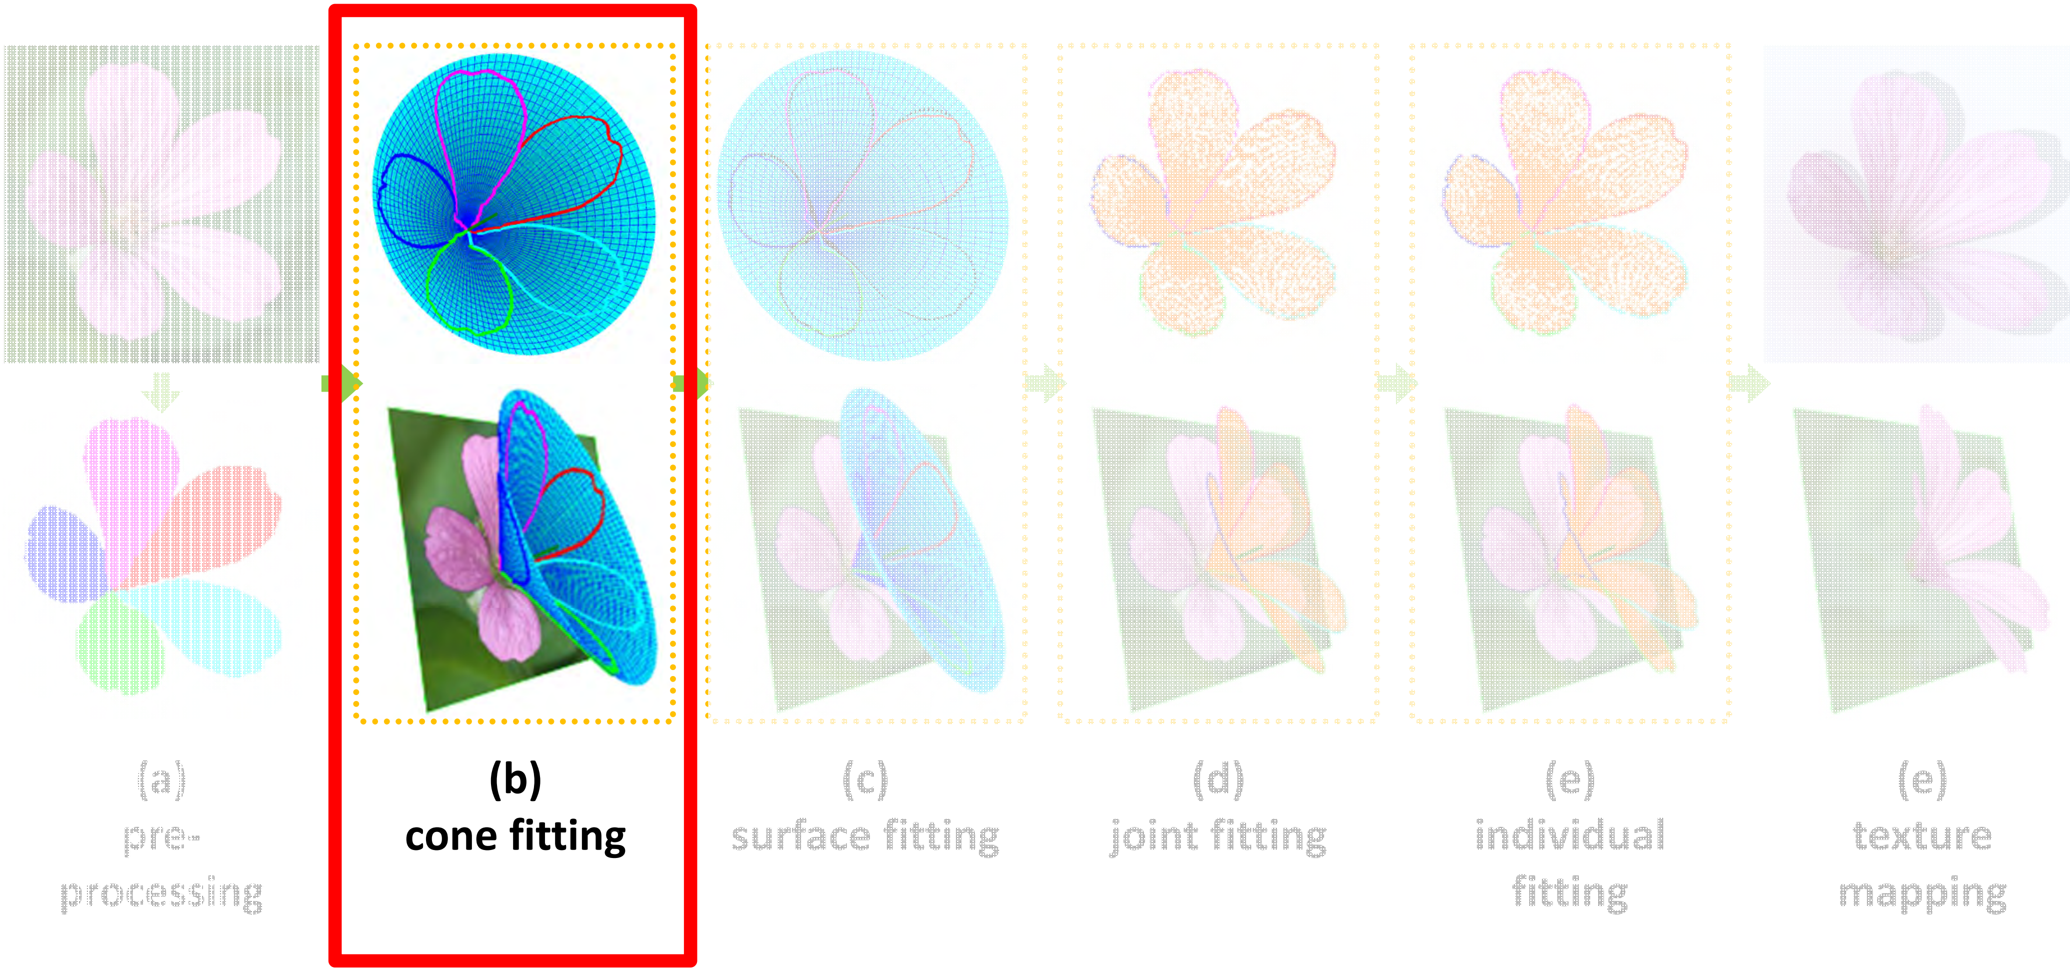

Flower Orientation Estimation

- **A flower can be roughly represented by a cone**
- **Input can be seen as the projection of a cone, under the parallel projection assumption**

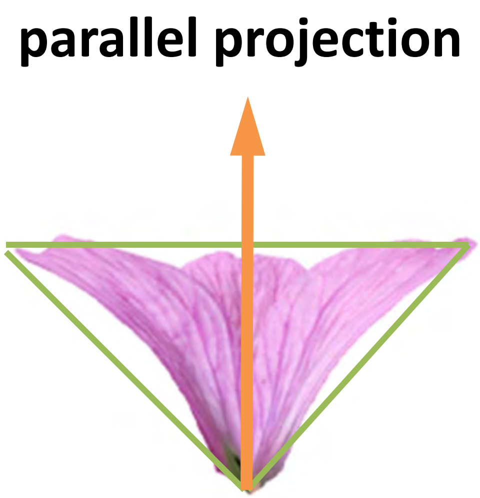

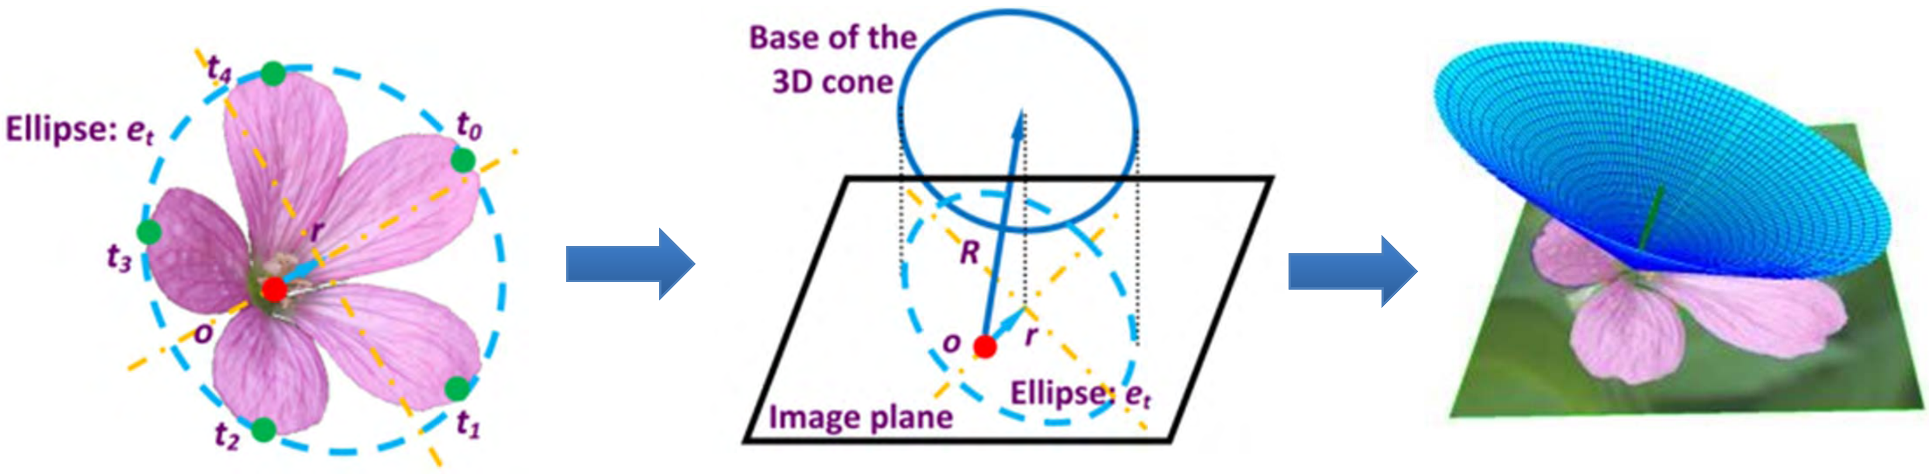

Cone Fitting

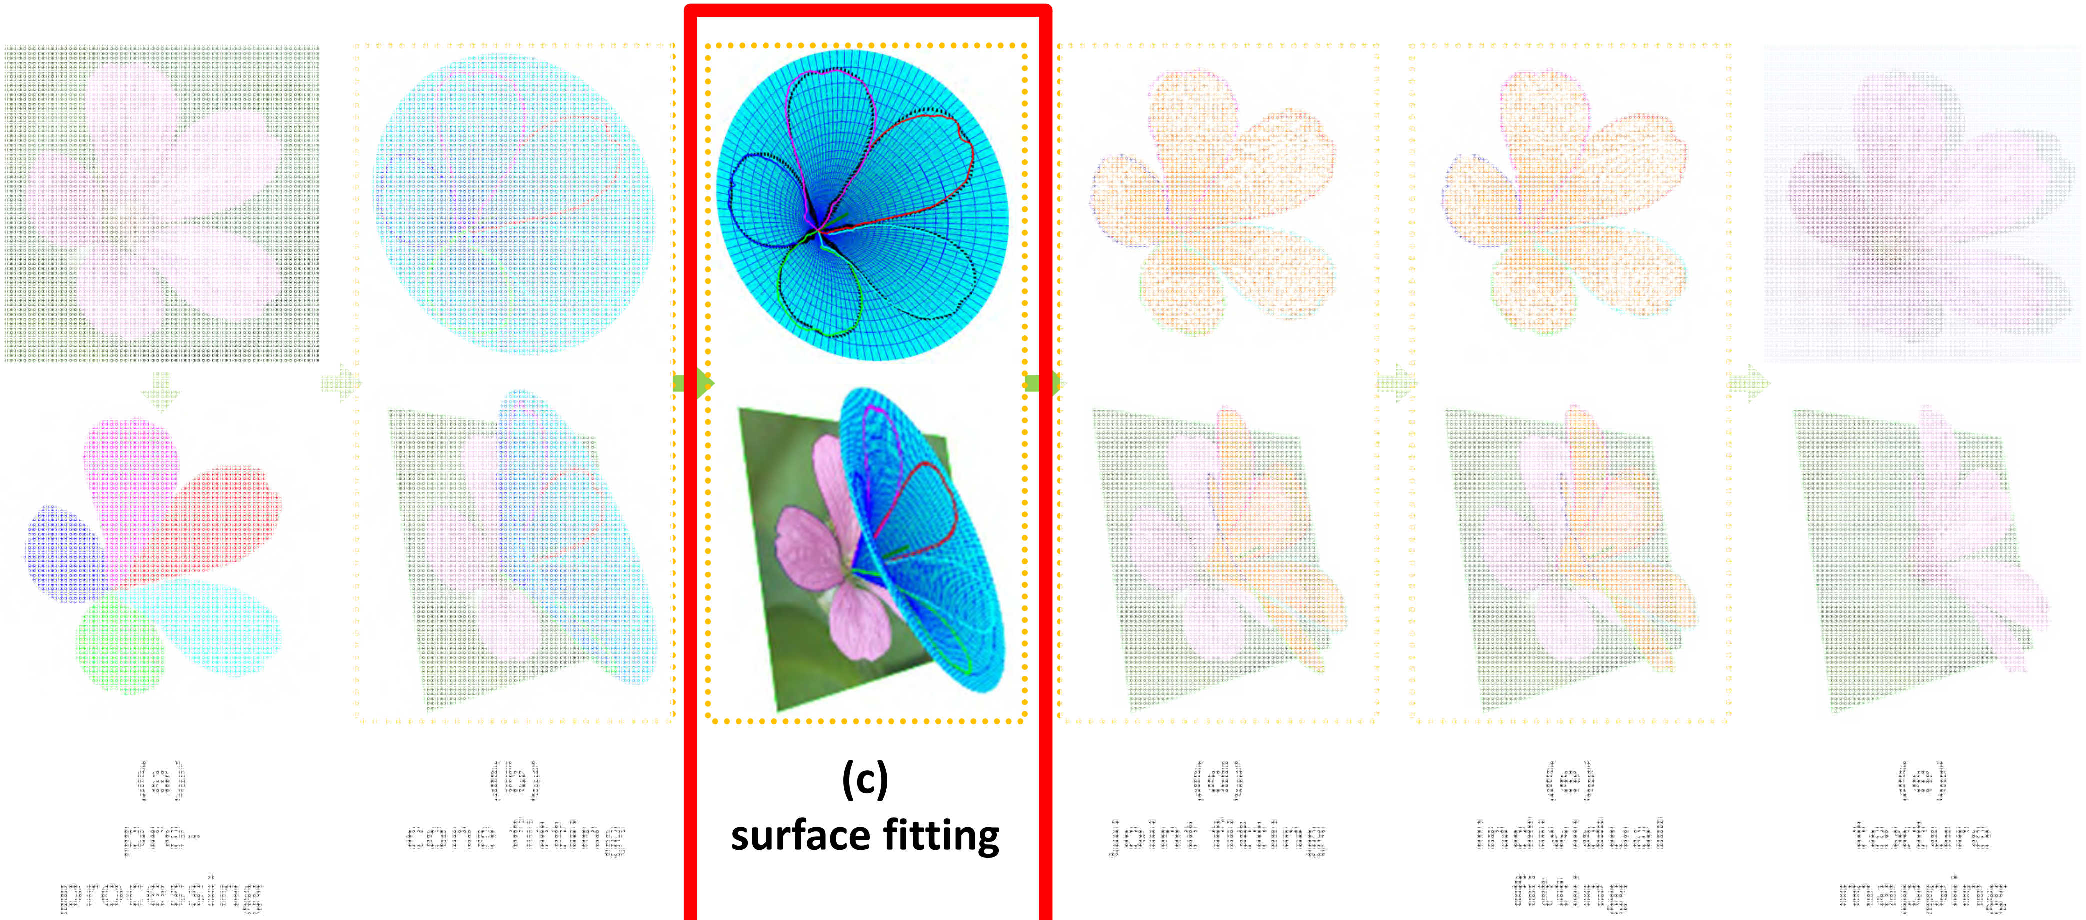

Flower Underlying Surface Fitting

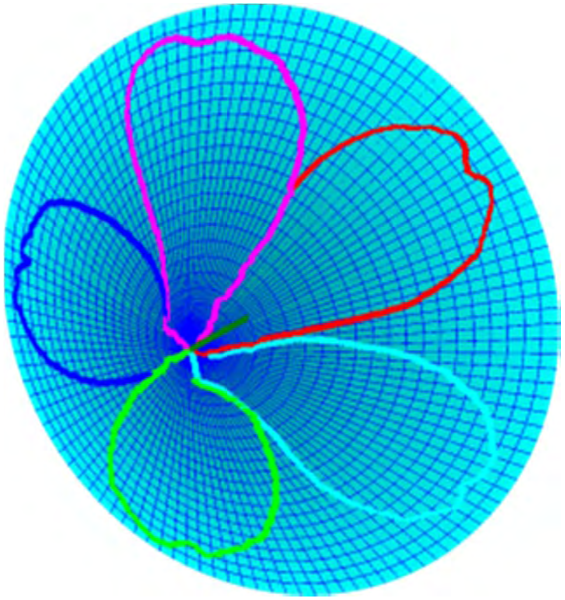

Generate petal contour template

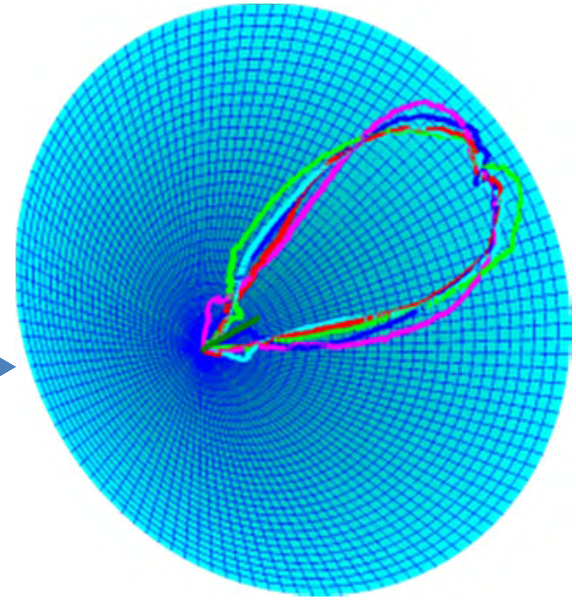

Color — observed contour in the photo

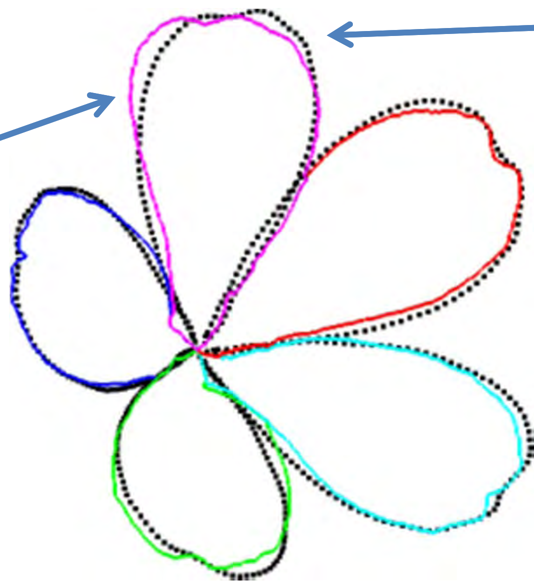

Black — projections of the contour template

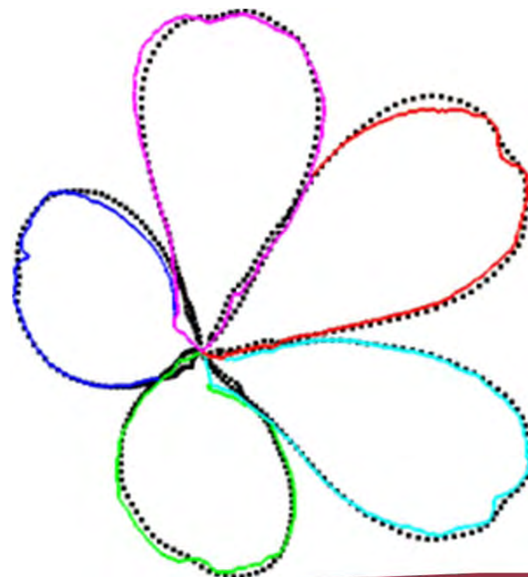

Fit a Surface of Revolution

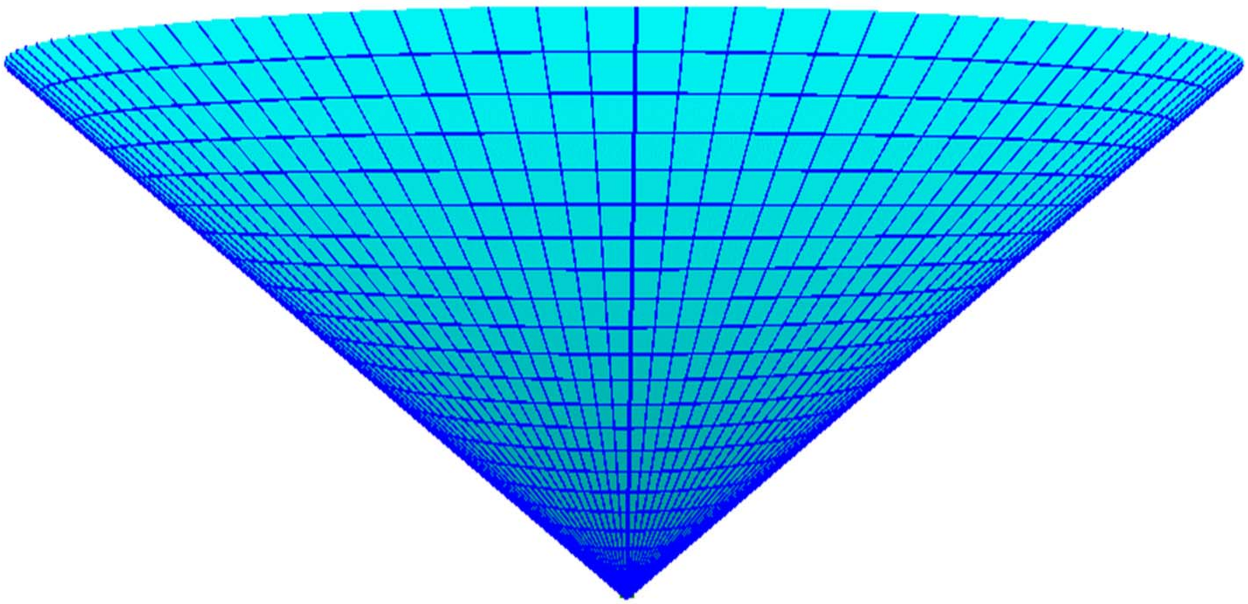

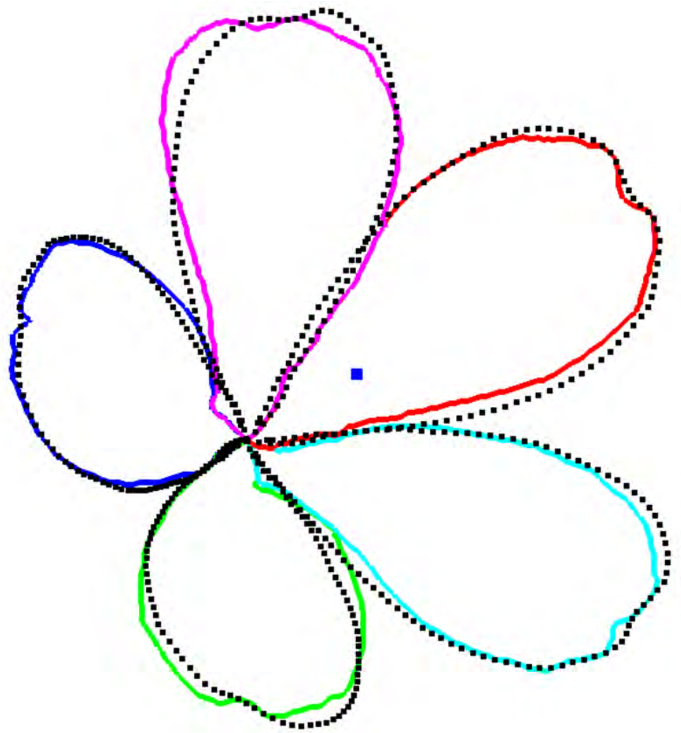

Surface Fitting Demo

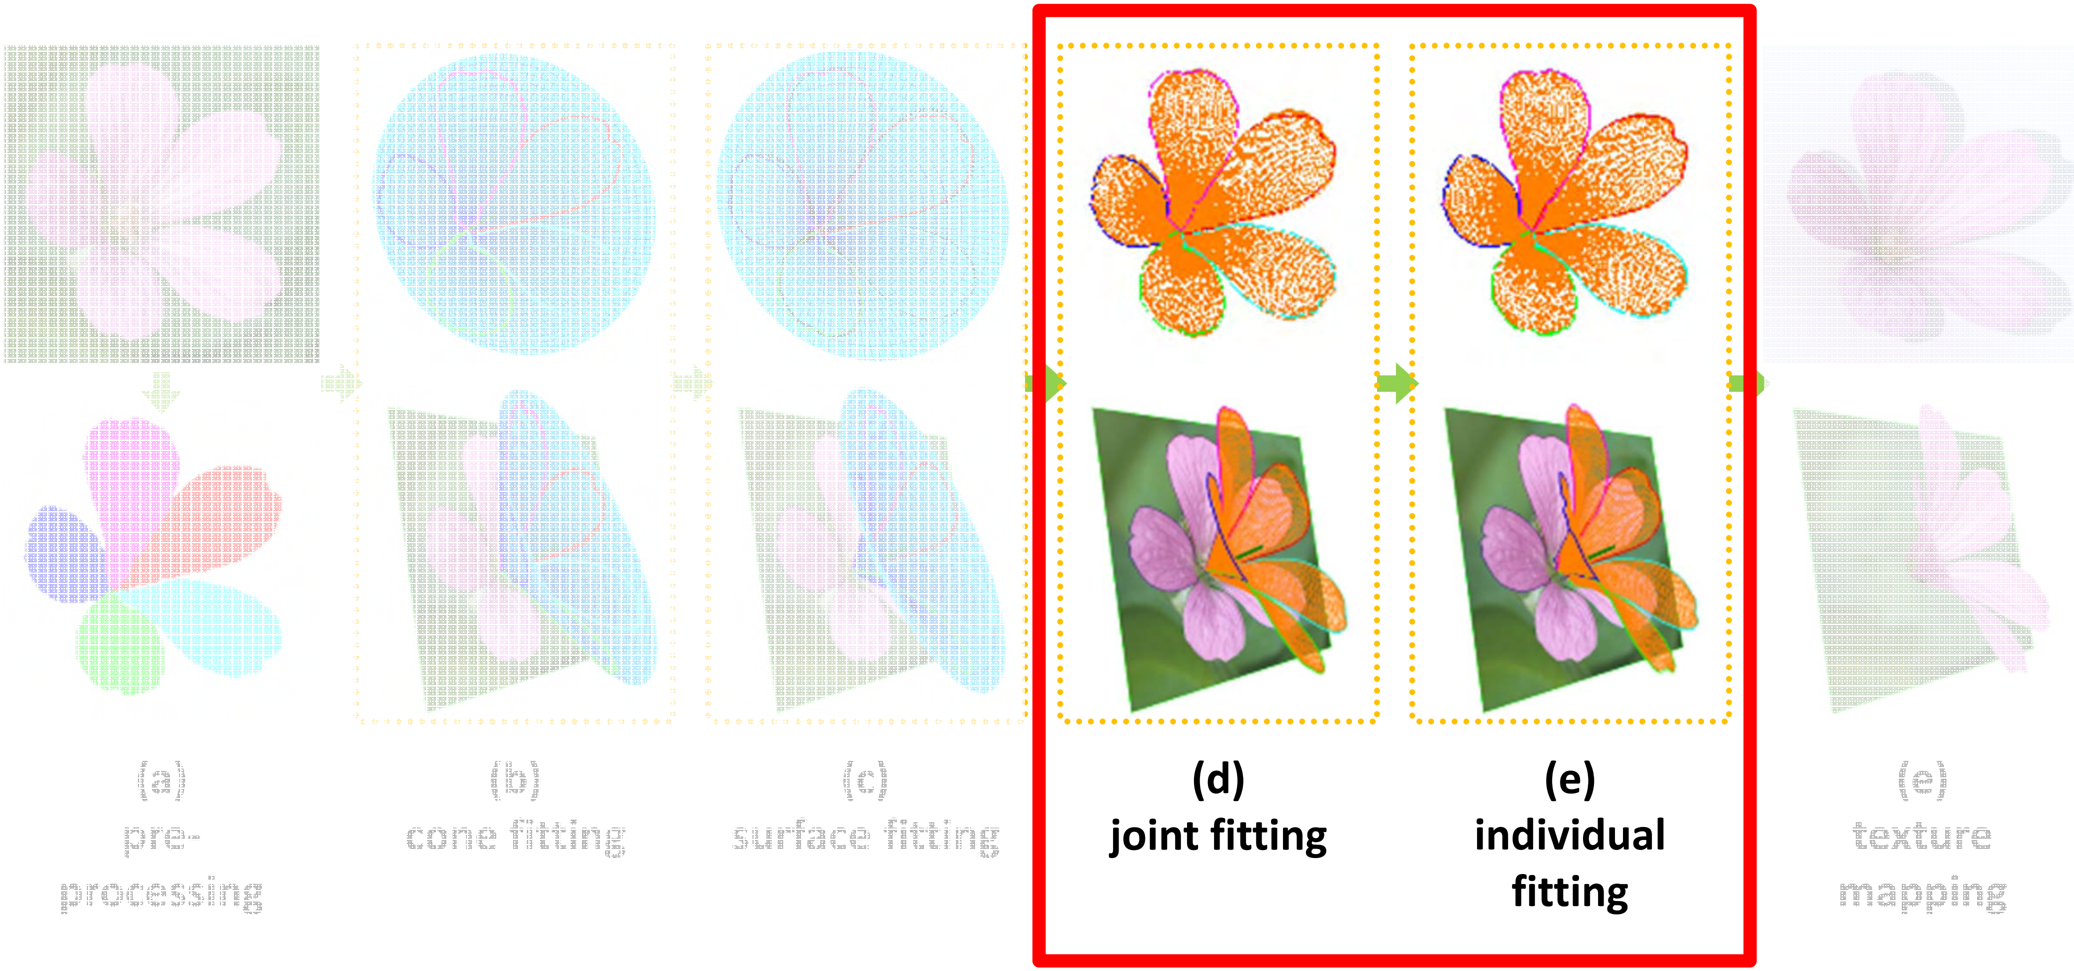

Petal Deformation

Petal Fitting: Joint and Individual

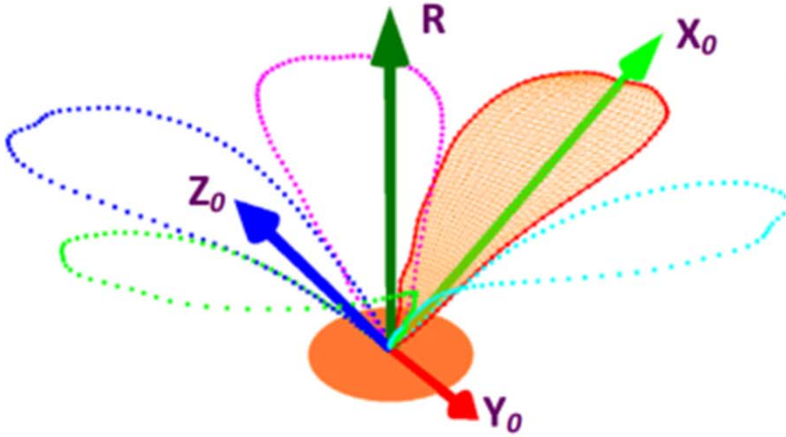

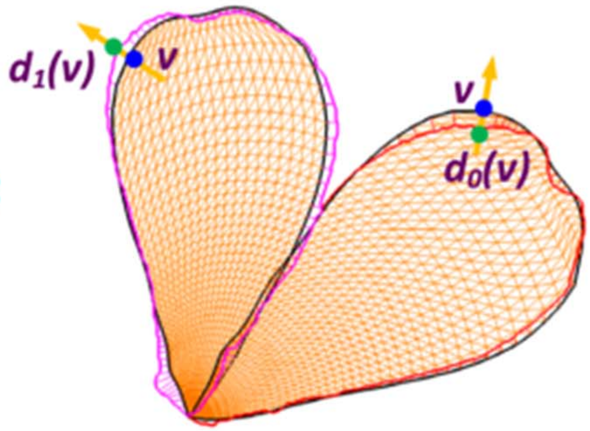

Move boundary vertices along their normal

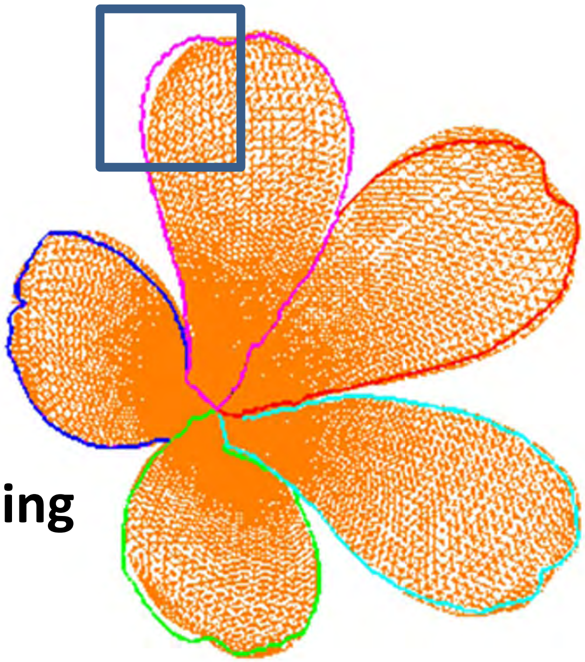

Joint fitting

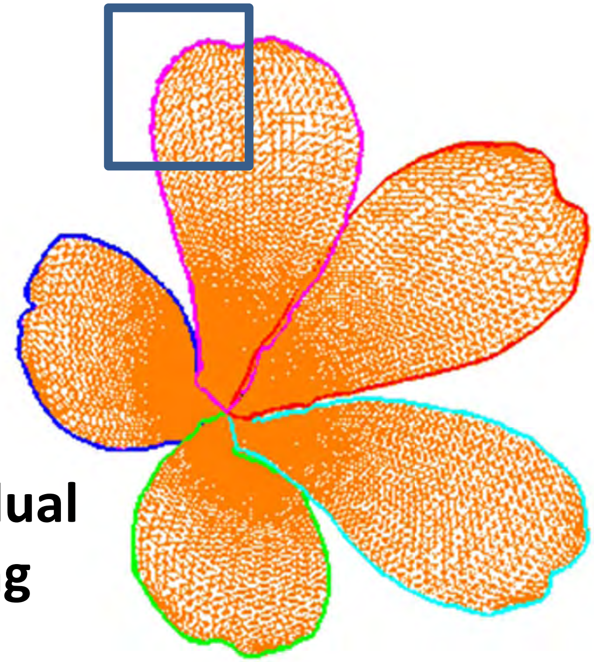

Individual fitting

- **Occlusion handling**

- **Flower of multiple layers**

Extension of the Pipeline

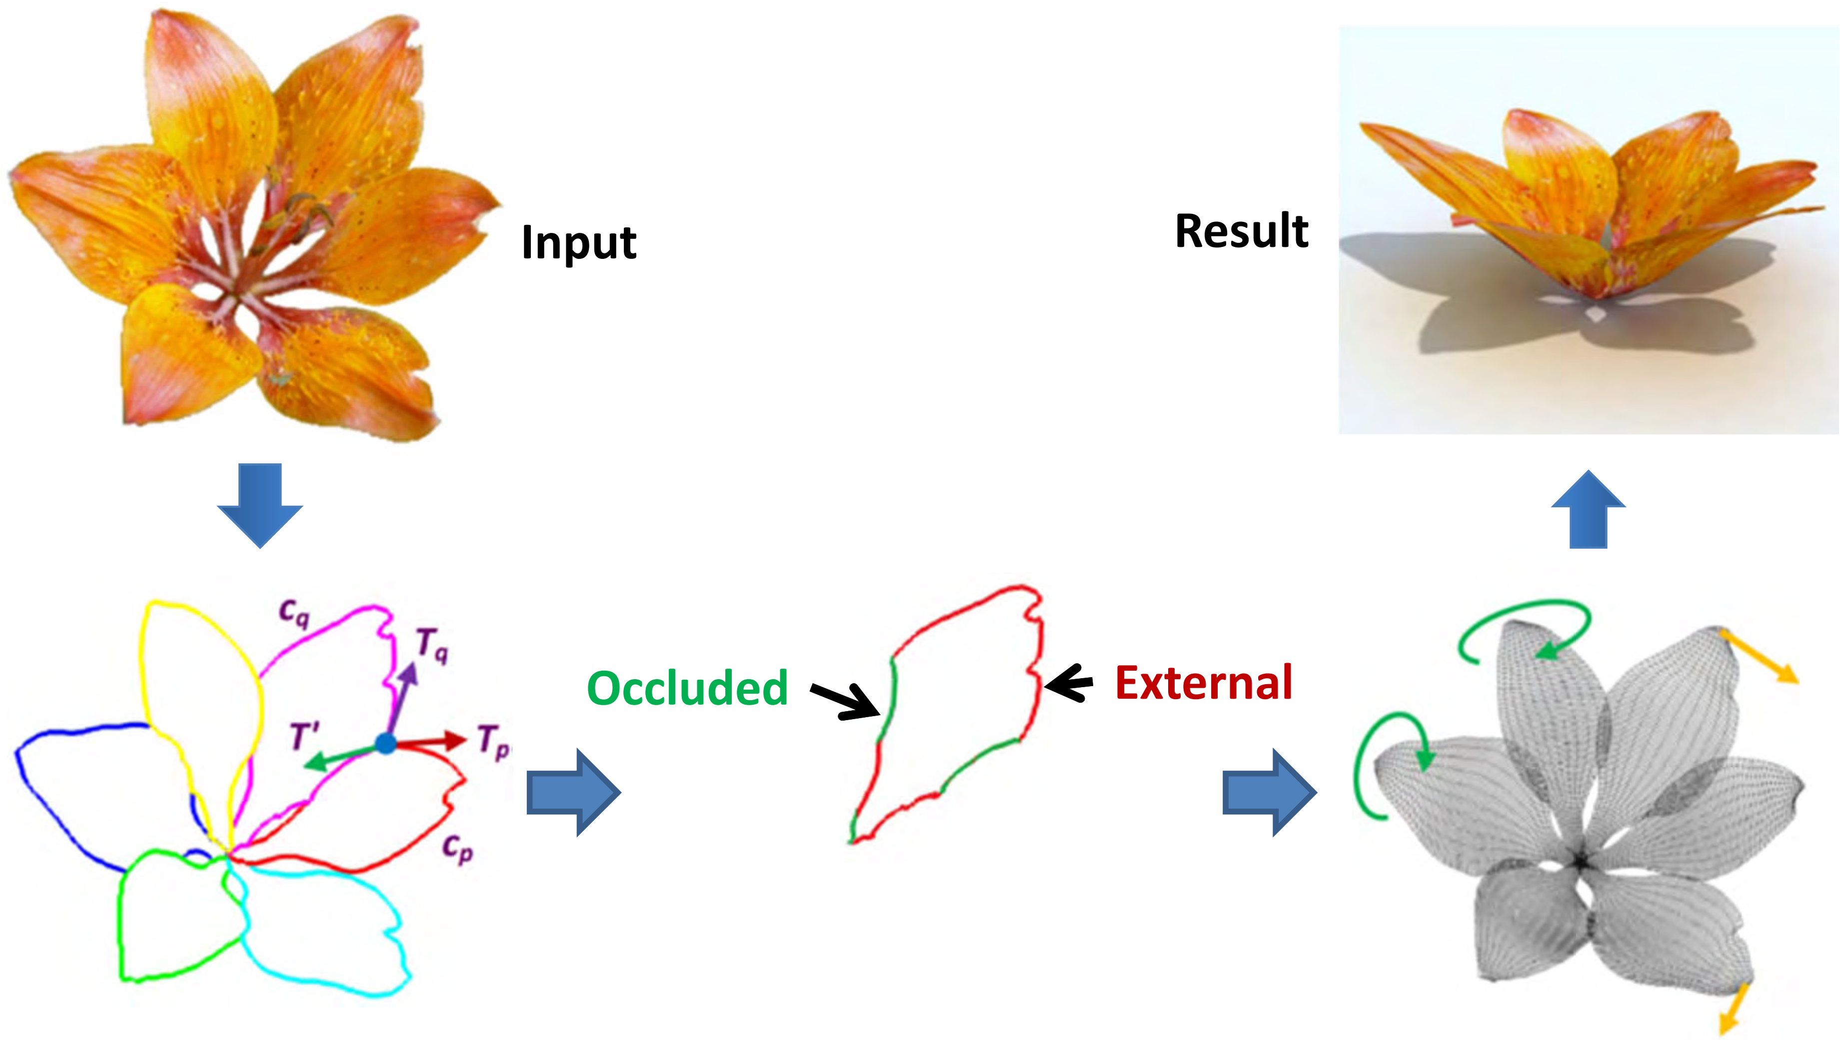

Occlusion Handling

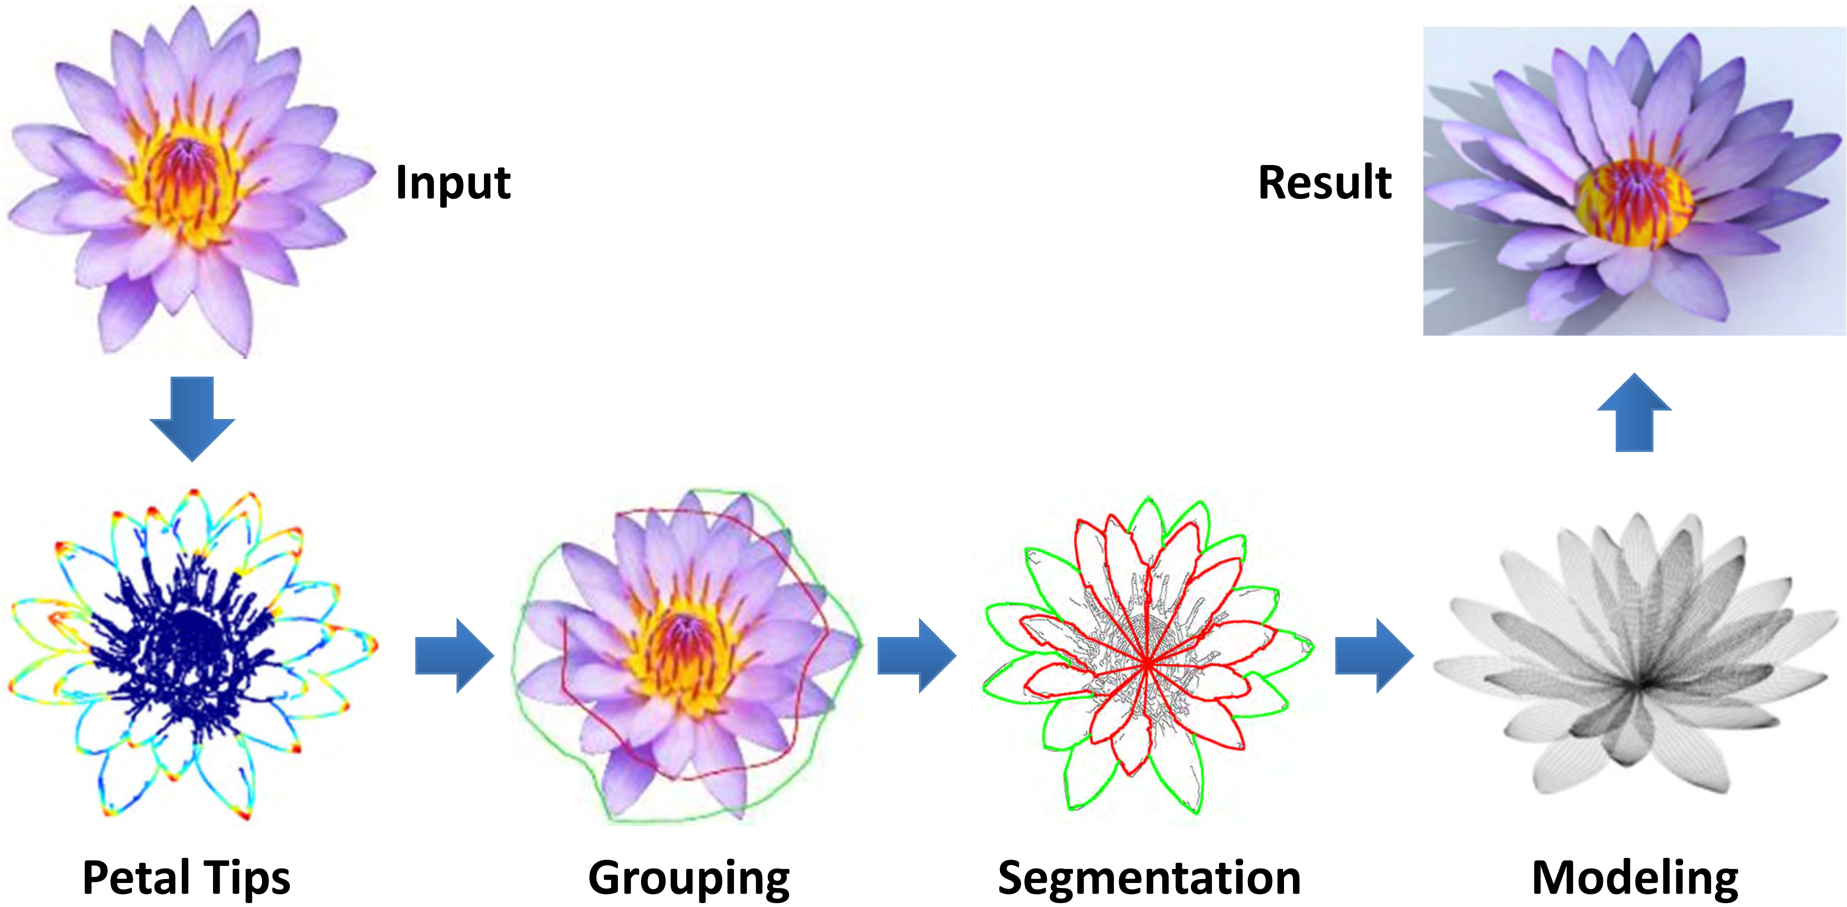

Multiple-Layer Flower

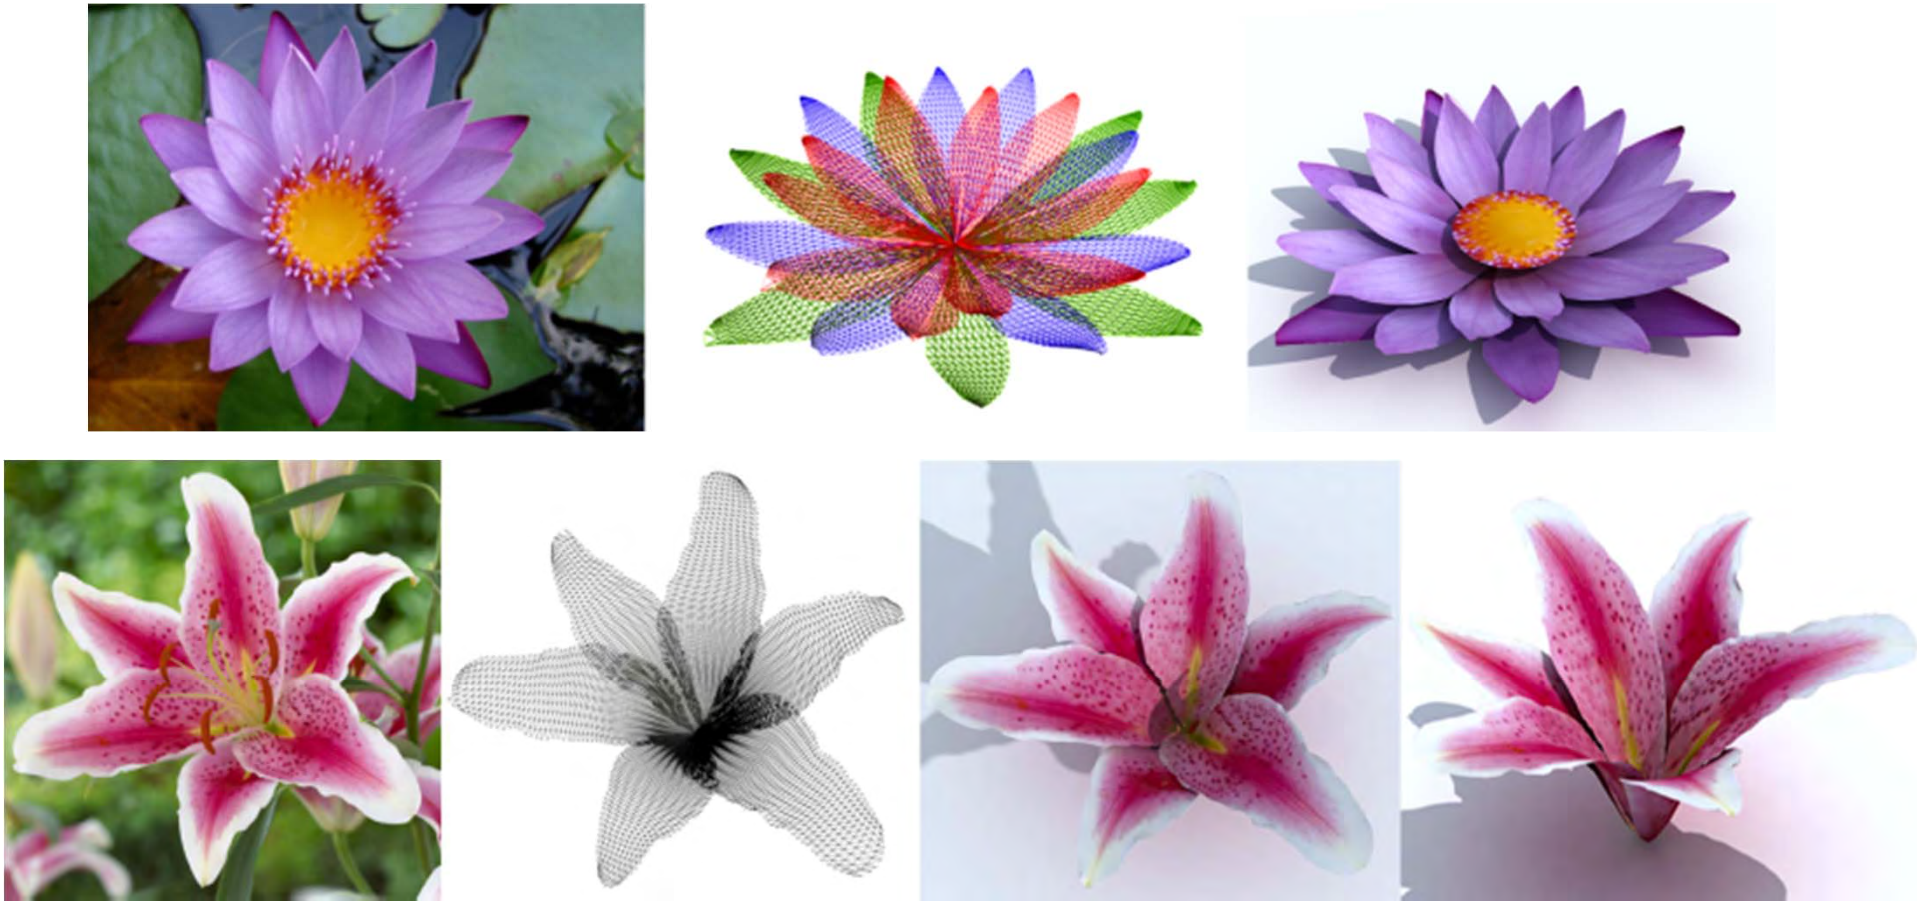

Results

More Results

- **Cannot capture geometric details on the petal surface**
- **Input photo cannot be taken:**
  - **from a highly oblique view direction**
  - **directly above the flower**

Limitations

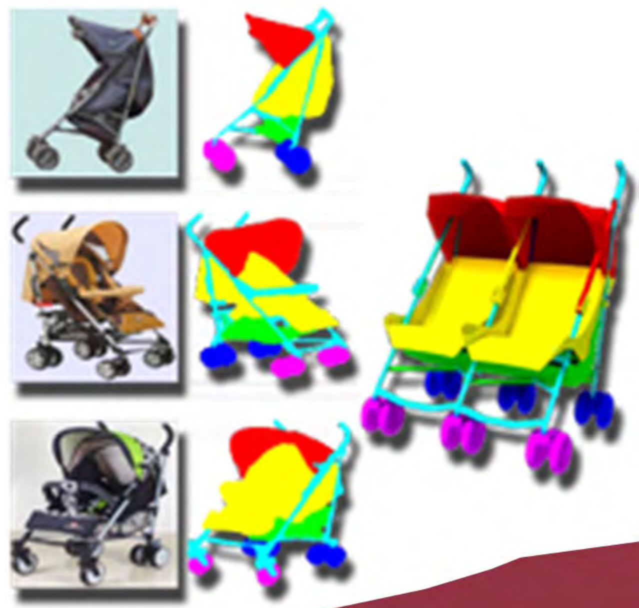

Part II

Labeling 3D Shapes using Images

- **Segmentation is one of the most fundamental tasks in shape analysis**
- **Low-level cues (minimal rule; convexity) alone are insufficient**

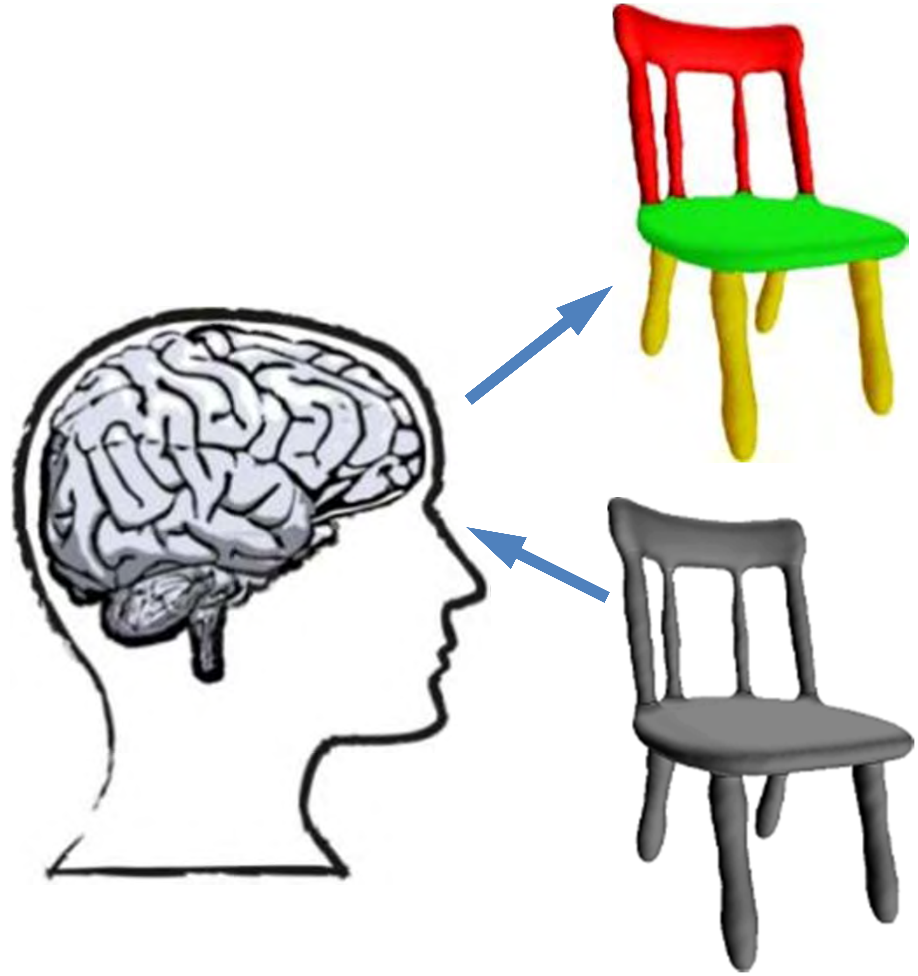

Segmentation of 3D Shapes

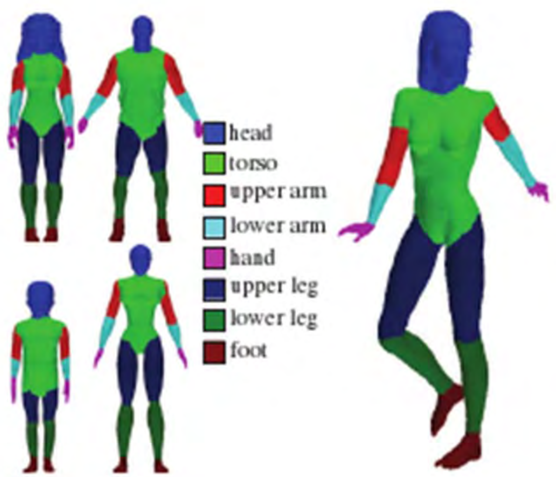

Learning segmentation
[Kalograkis et al. 10]

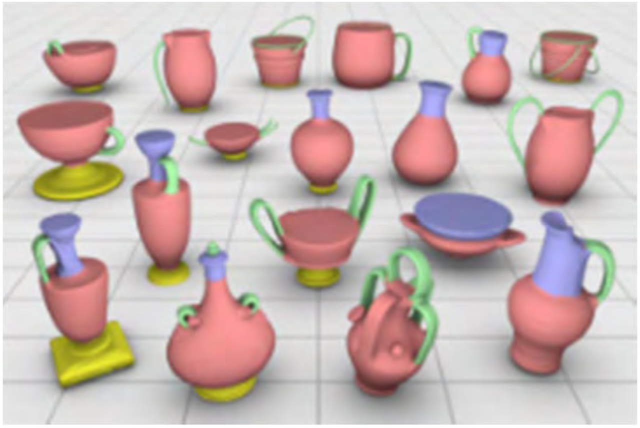

Unsupervised co-analysis
[Sidi et al. 2011]

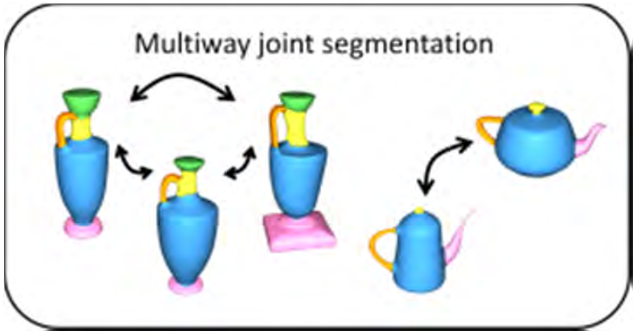

Joint segmentation
[Huang et al. 2011]

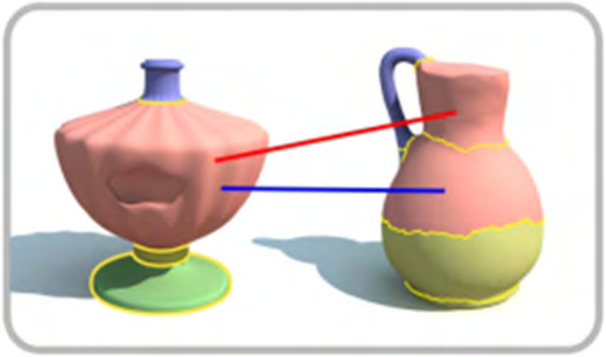

Active co-analysis
[Wang et al. 2012]

Existing Knowledge-driven Approaches

- How many 3D models of strollers, golf carts, gazebos, …?
- Not enough 3D models = insufficient knowledge
- Labeling 3D shapes is also a non-trivial task

380 labeled meshes over 19 object categories

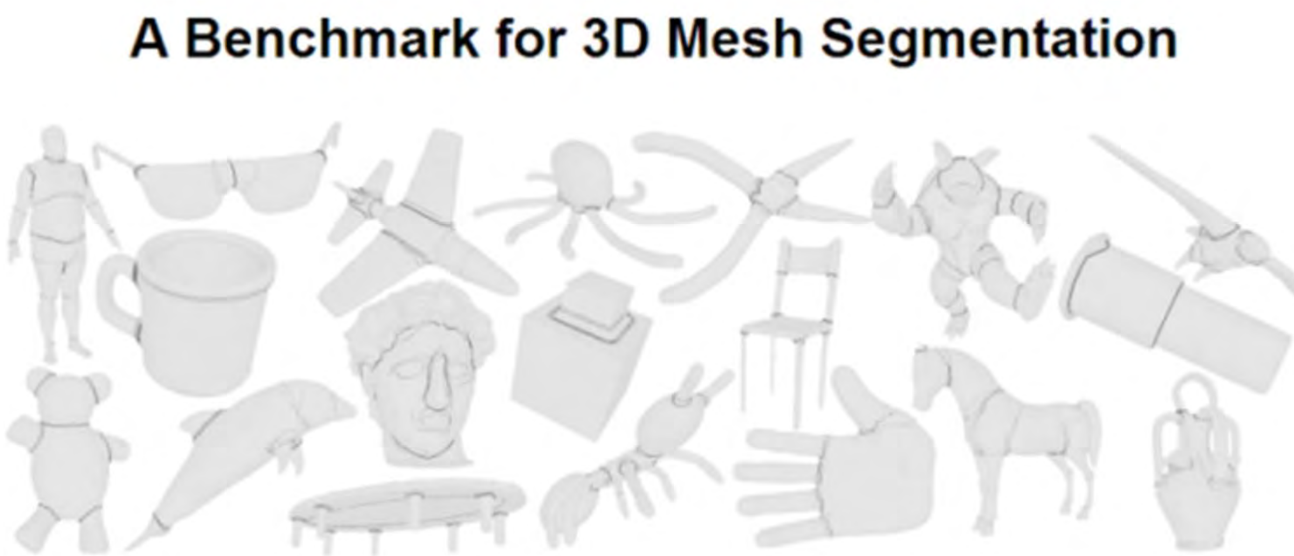

3D Data Challenge: Amount

Self-intersecting; non-manifold

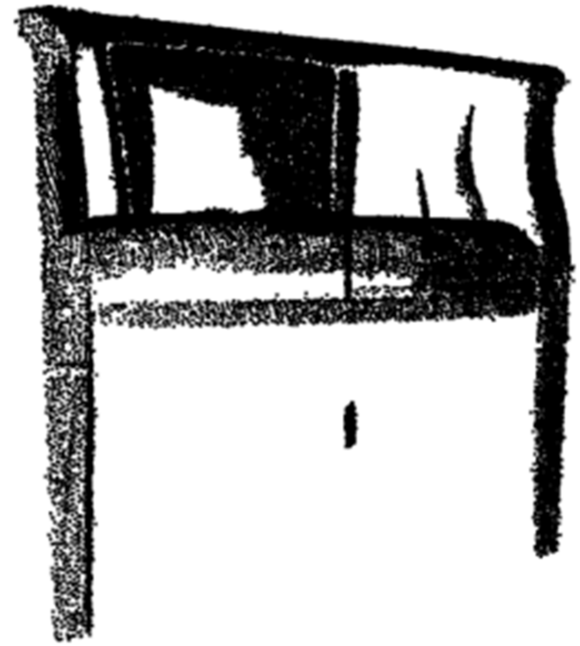

Incomplete

Real-world 3D models (e.g., those from Tremble Warehouse) are often imperfect

3D Data Challenge: Quality

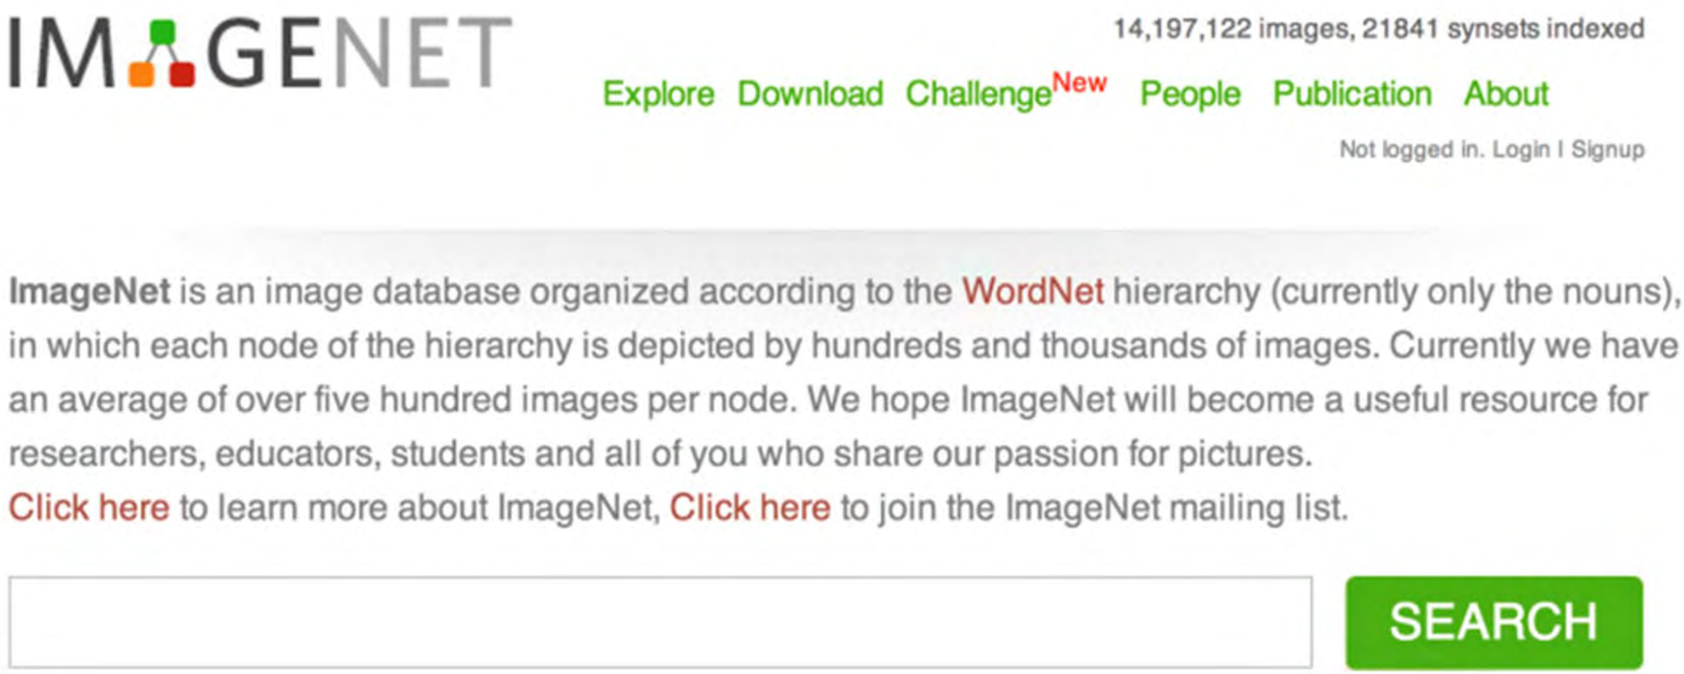

About 14 million images across almost 22,000 object categories

Labeling images is quite a bit easier than labeling 3D shapes

Many More Images Available

- Treat a 3D shape as a set of projected binary images
  - Label these projections by learning from vast amount of image data
- Then propagate the labels to the 3D shape
  - Allow us to analyze imperfect 3D shapes

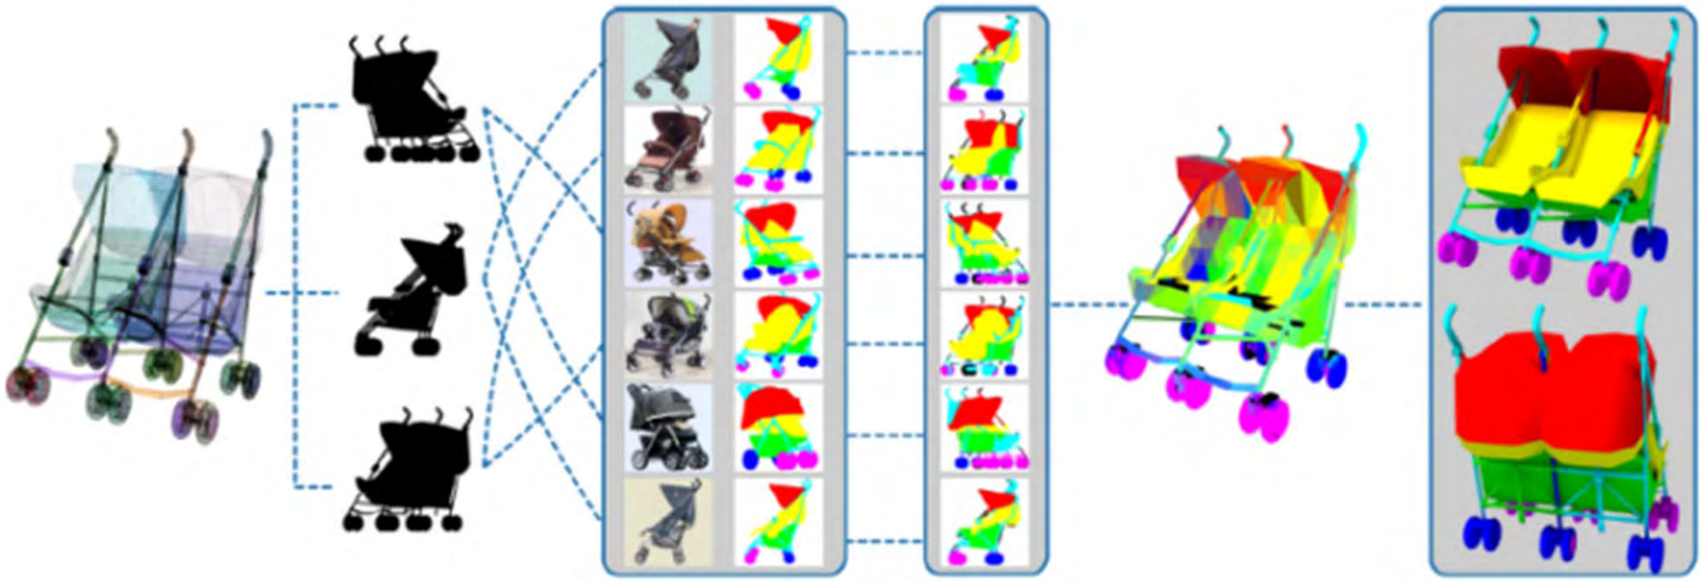

Motivations

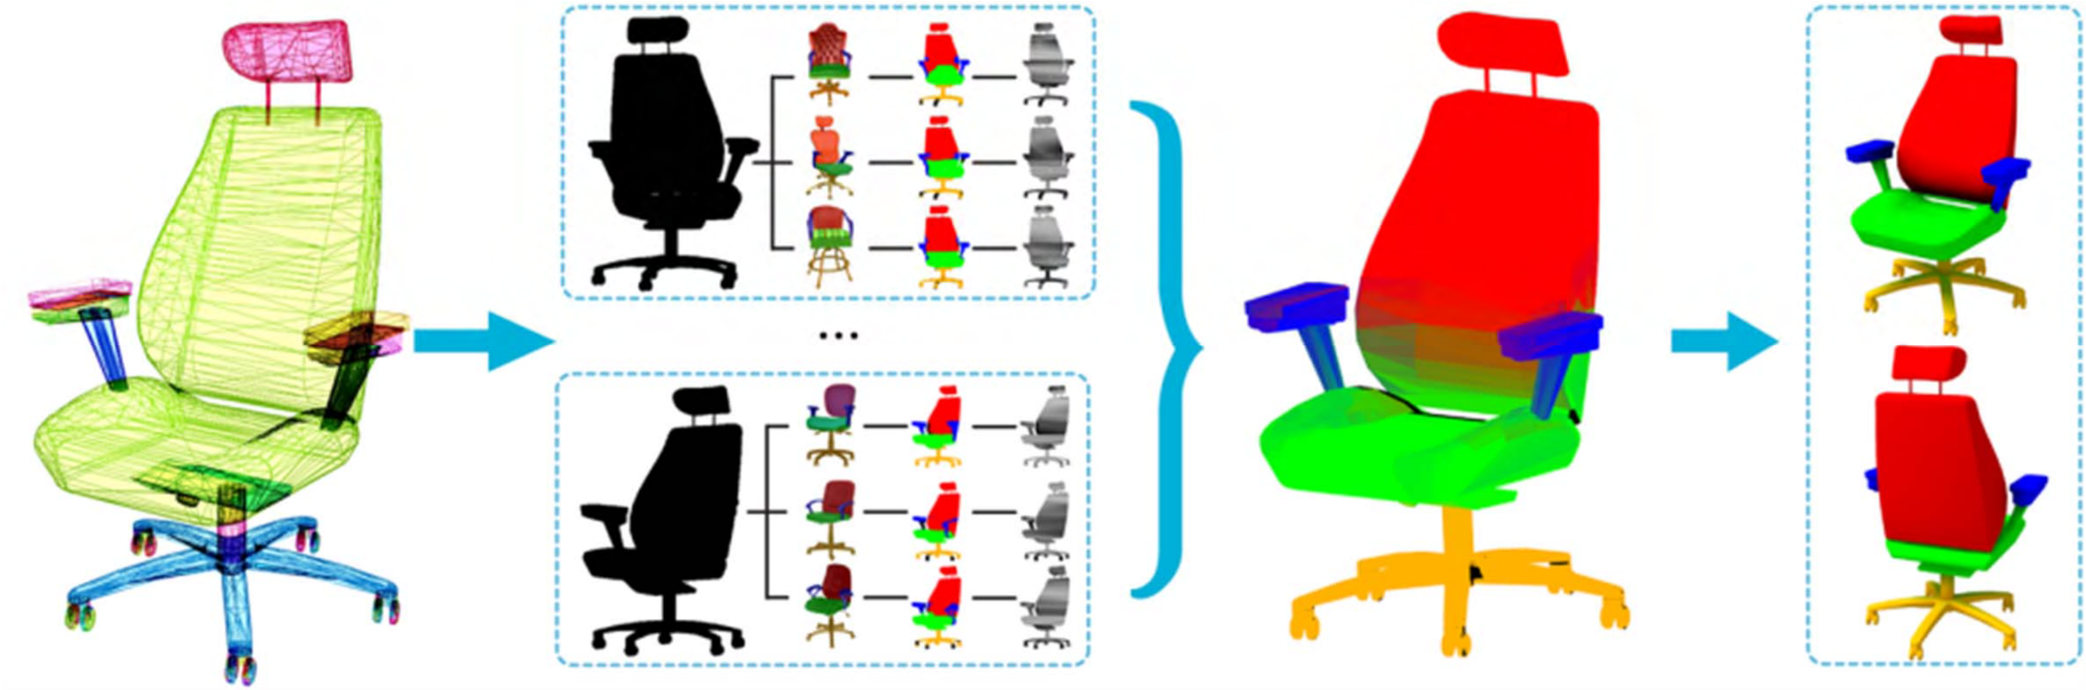

(a) project input shape

(b) image retrieval & label transfer

(c) back projection

(e) Graph cuts optimization

Algorithm Pipeline

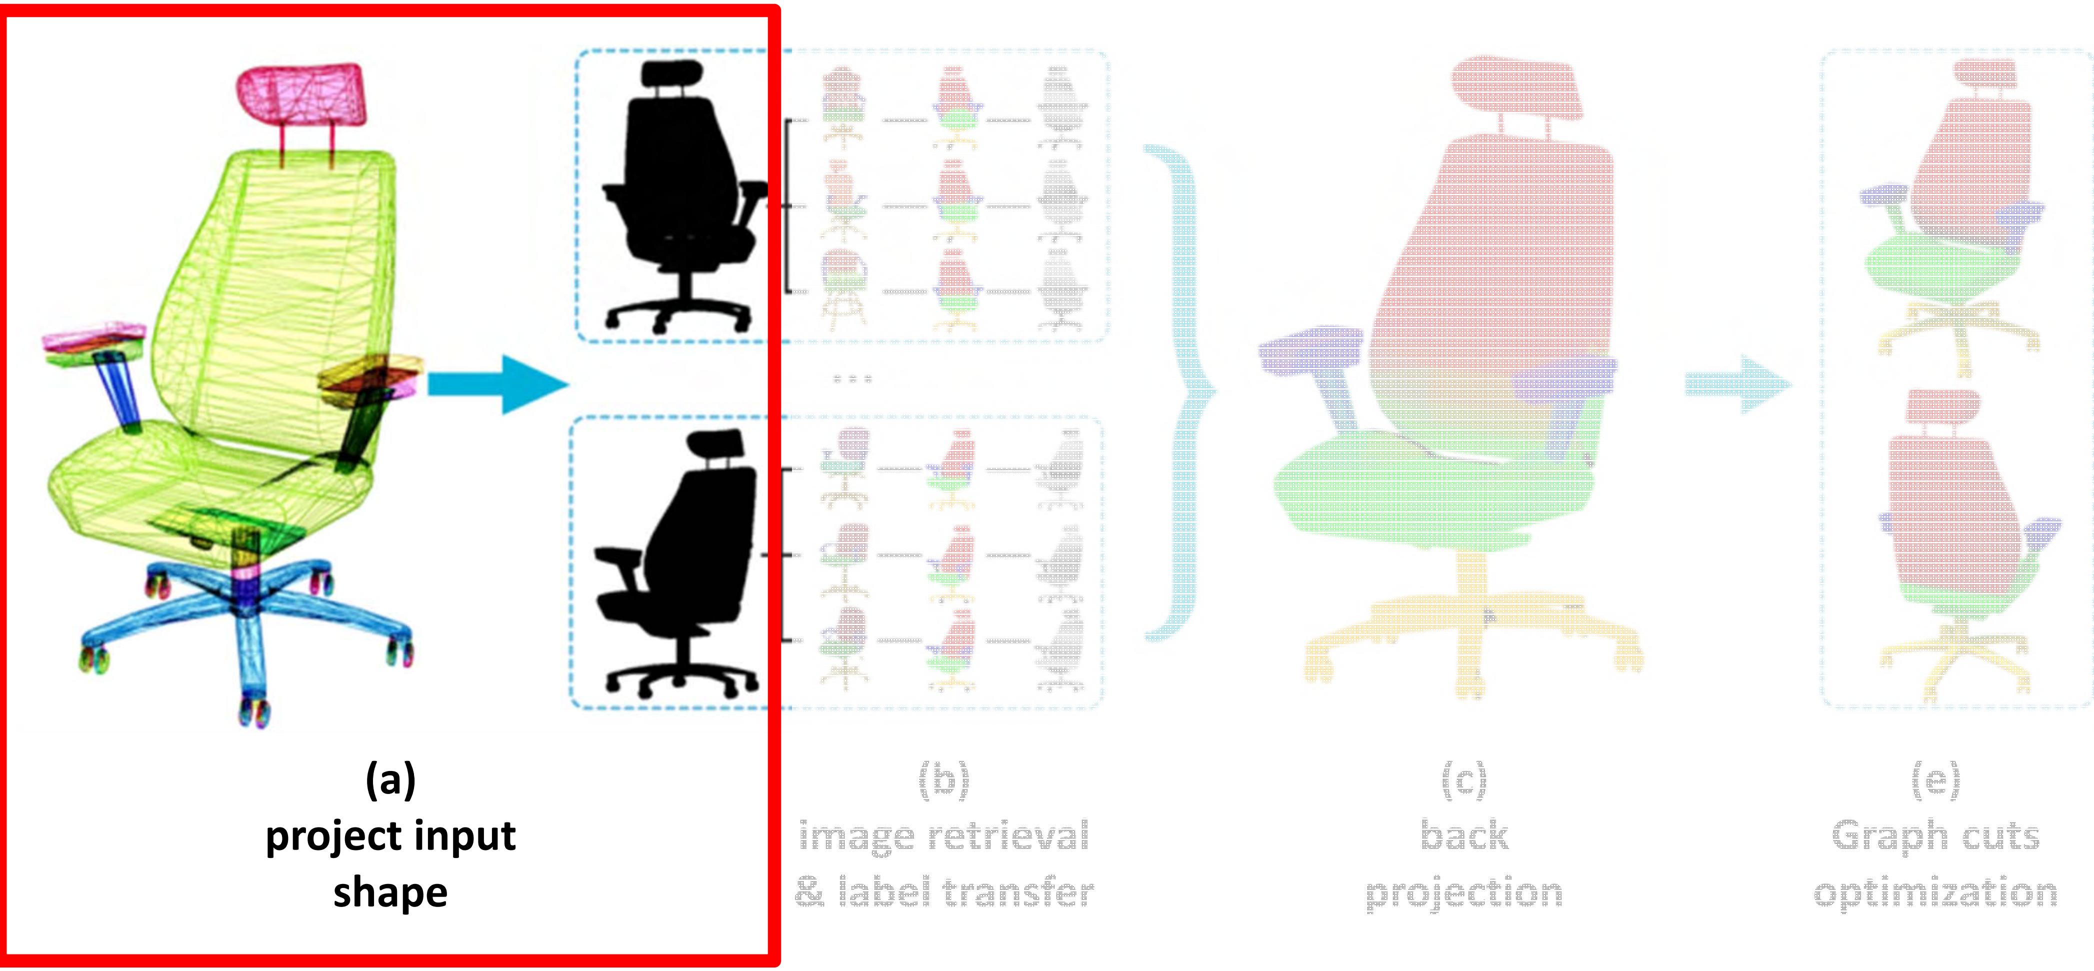

Project Input Shape

- **Assume objects are upright oriented**
  - **Most objects are!**
- **Project an input 3D shape from multiple pre-set viewpoints**

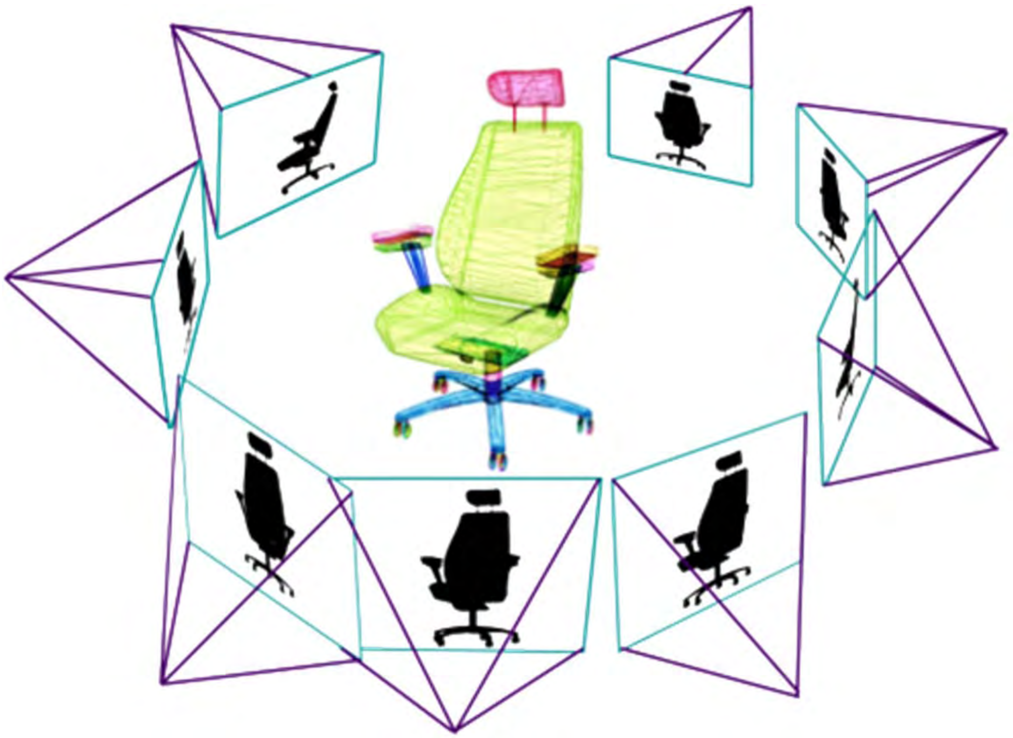

Generate 2D Projections for the Input

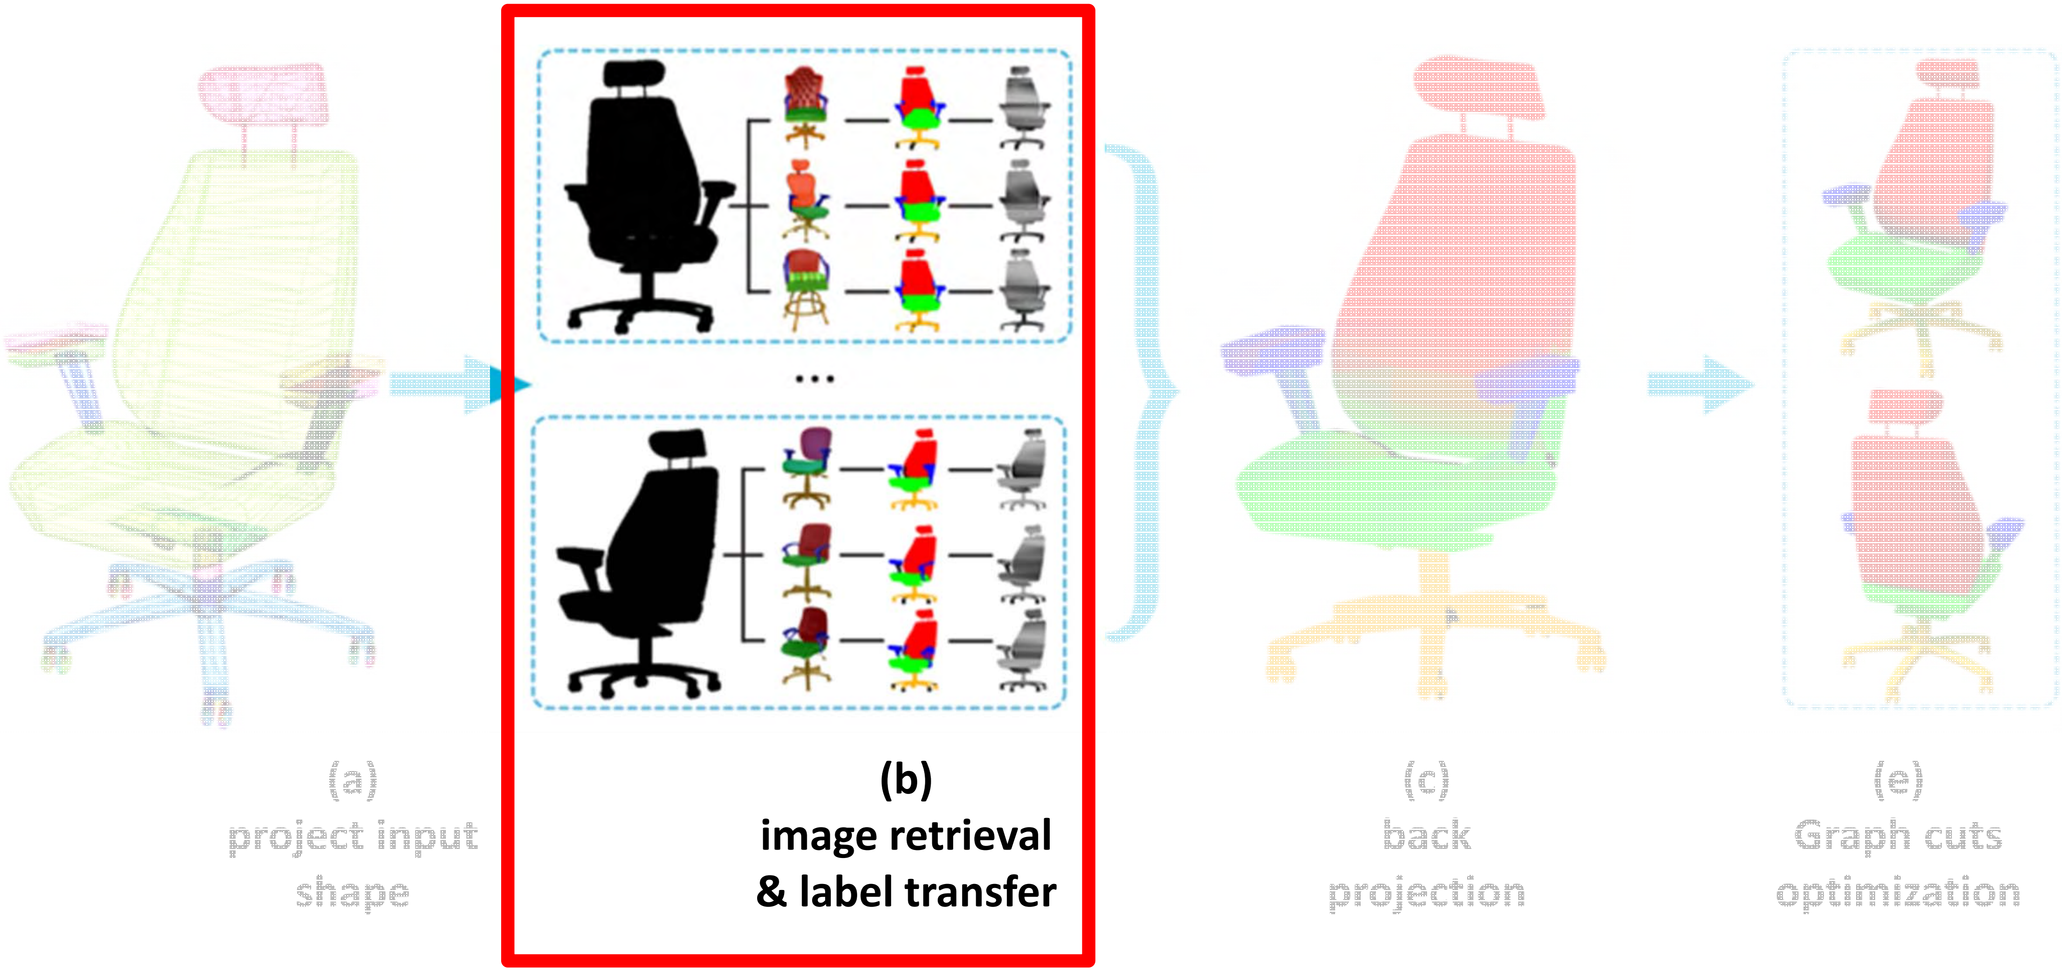

Label the Projections

- For each projection, retrieve top matches from the set of labelled images
- A novel 2D shape retrieval approach is proposed
  - Look for shapes with similar topology but ignore differences in part scales

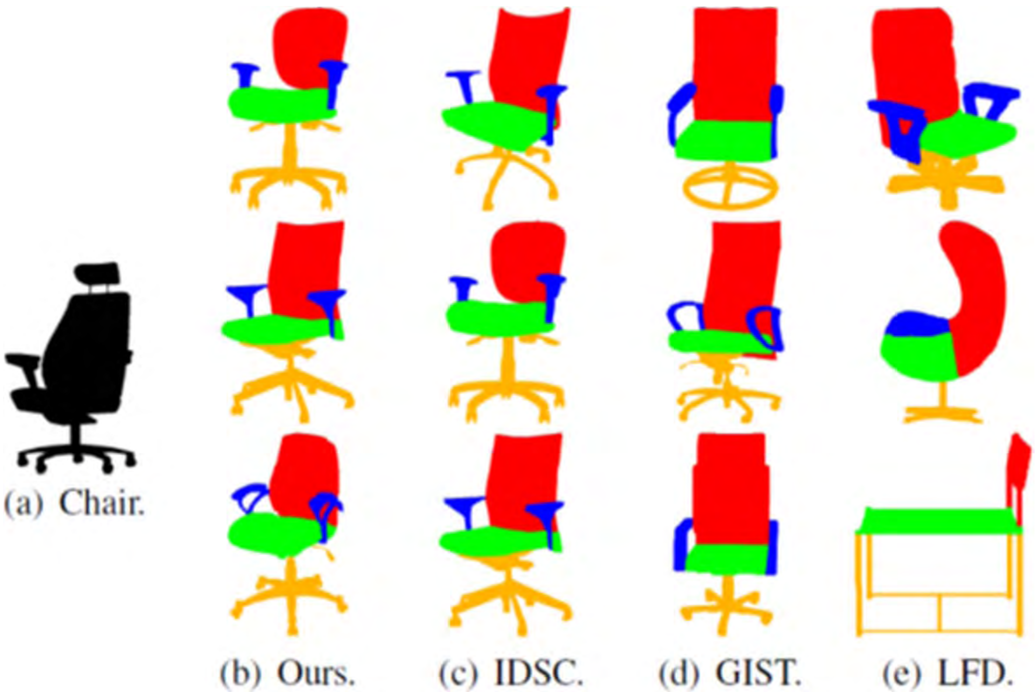

Retrieve Labeled Images

- **Split image into horizontal slabs**
  - Optimal corresponding slabs are found using a dynamic time warping (DTW) algorithm
  - Dissimilarity between 2 images is the sum of dissimilarities of corresponding slabs
- **Effectively warp images before compare**

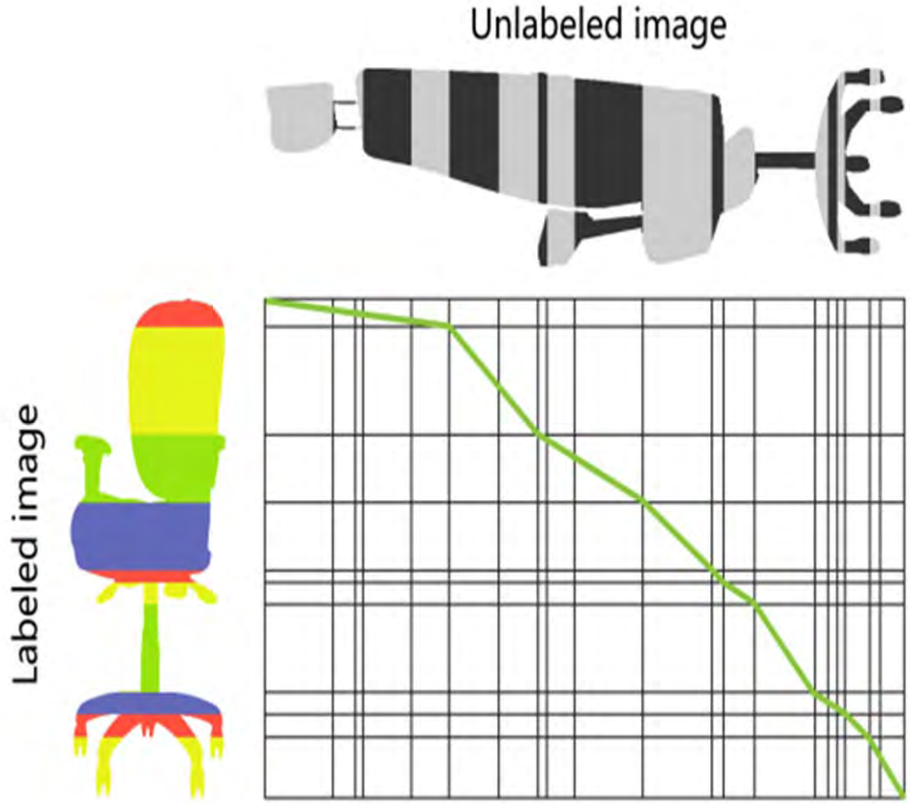

Image Comparison by Slab Matching

- **The top matching labeled images are used to transfer labels to the projection**
  - **Images are warped vertically to align with projection**
  - **Label transfer is done per corresponding horizontal slabs**

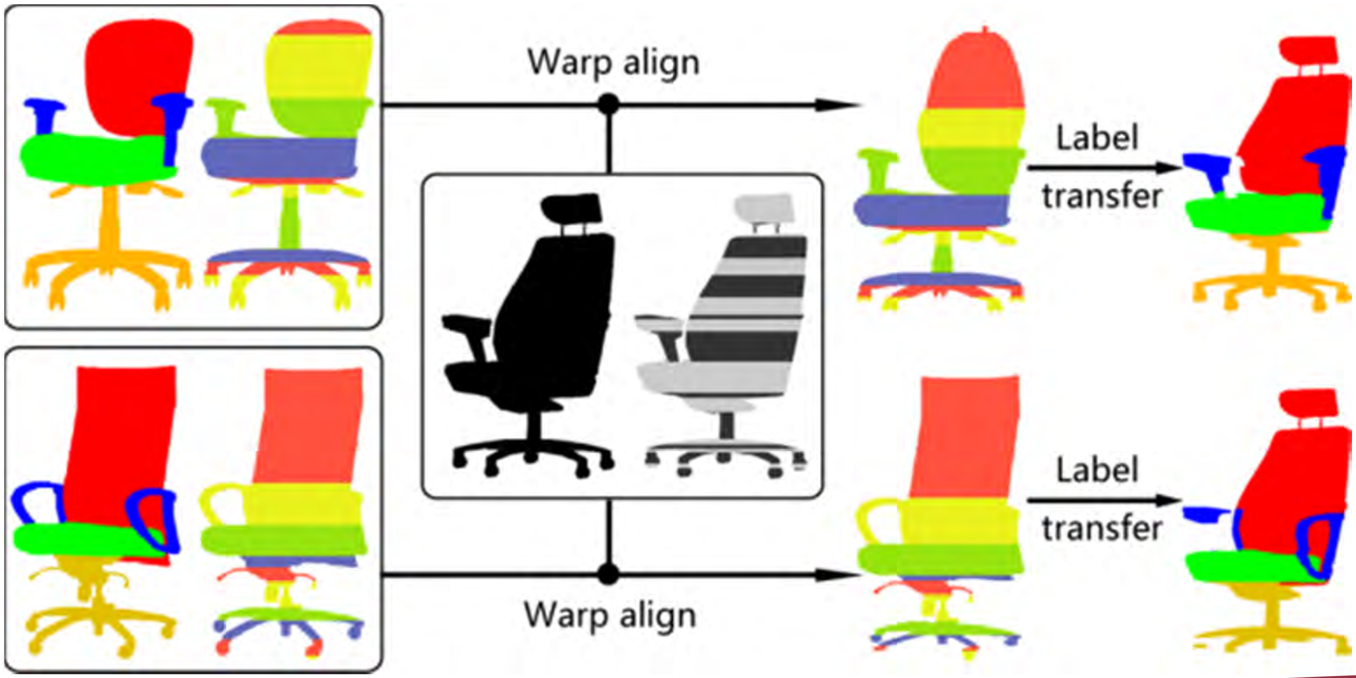

Label Transfer

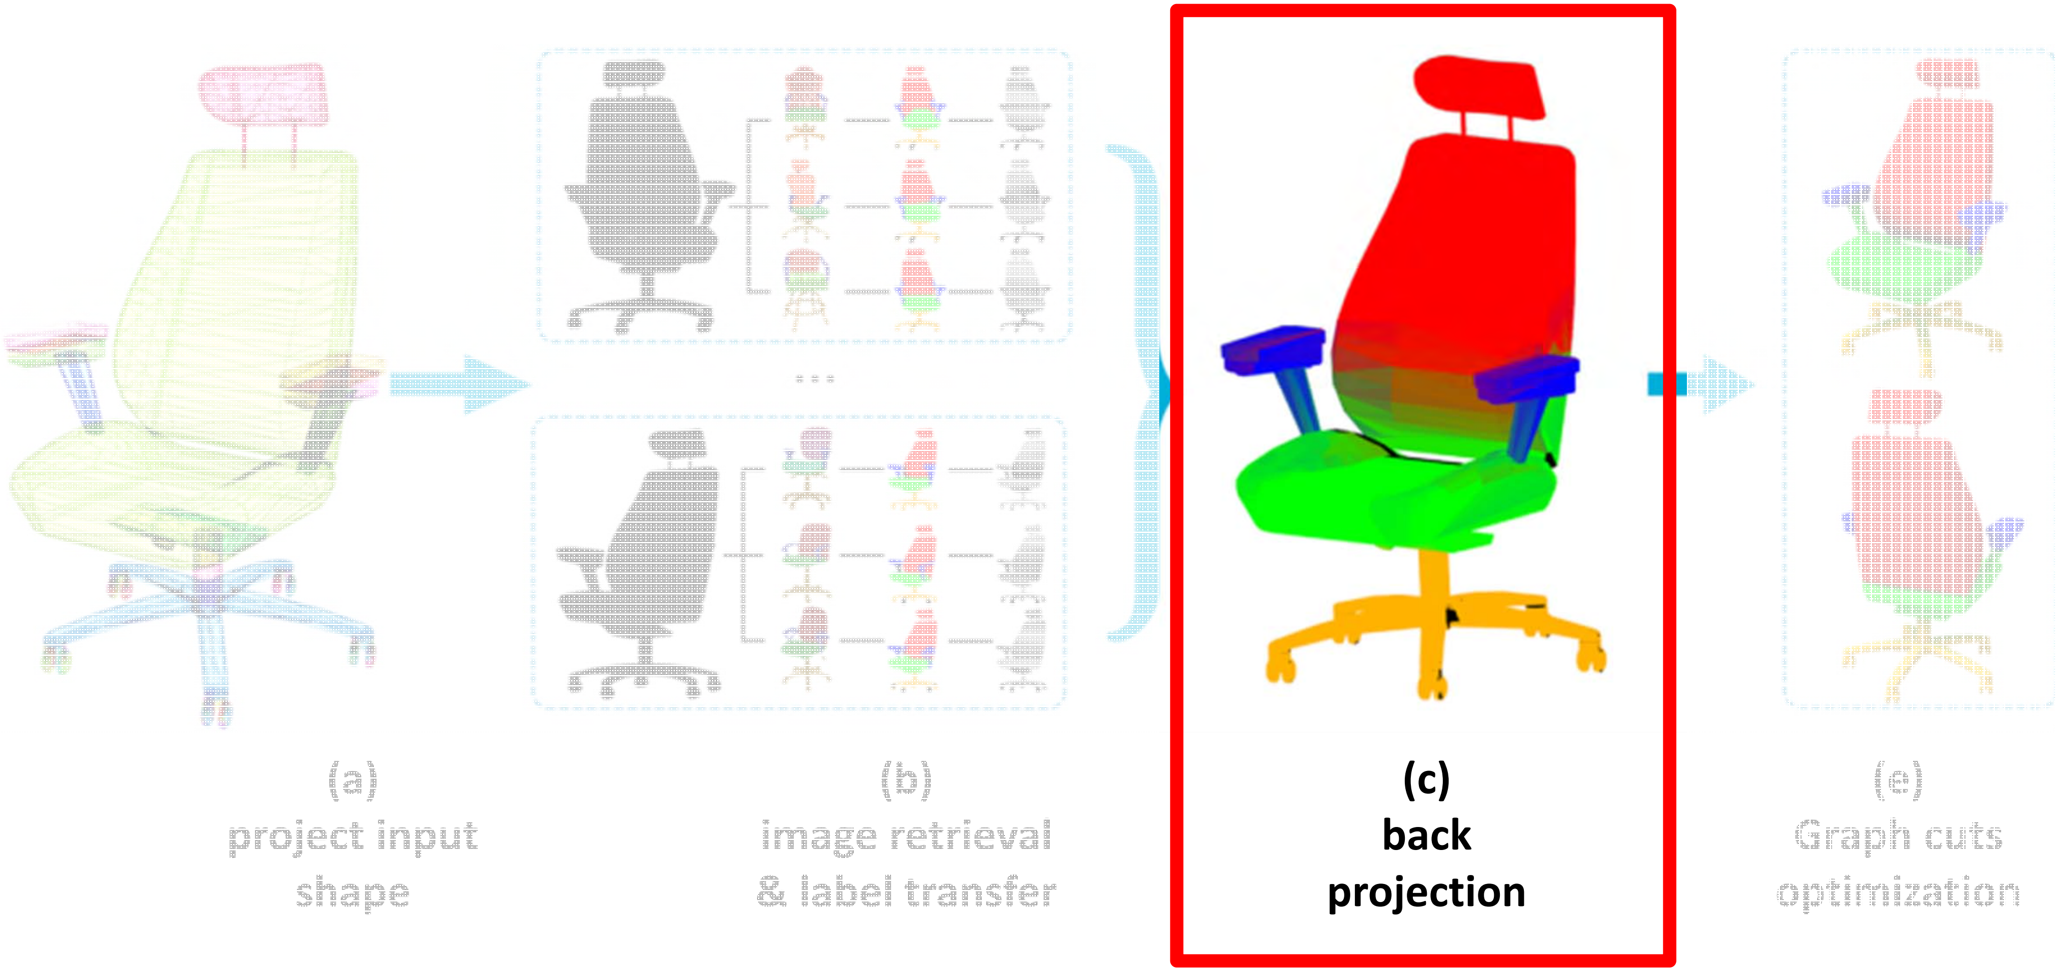

Back Projection

- **Compute a label probabilistic map over the input 3D model:**
  - Each 3D primitive covers multiple pixels in multiple projections
  - Each projection has multiple retrieved images
  - Per-pixel labels and confidences are gathered & integrated

Back Project Labels to 3D Surface

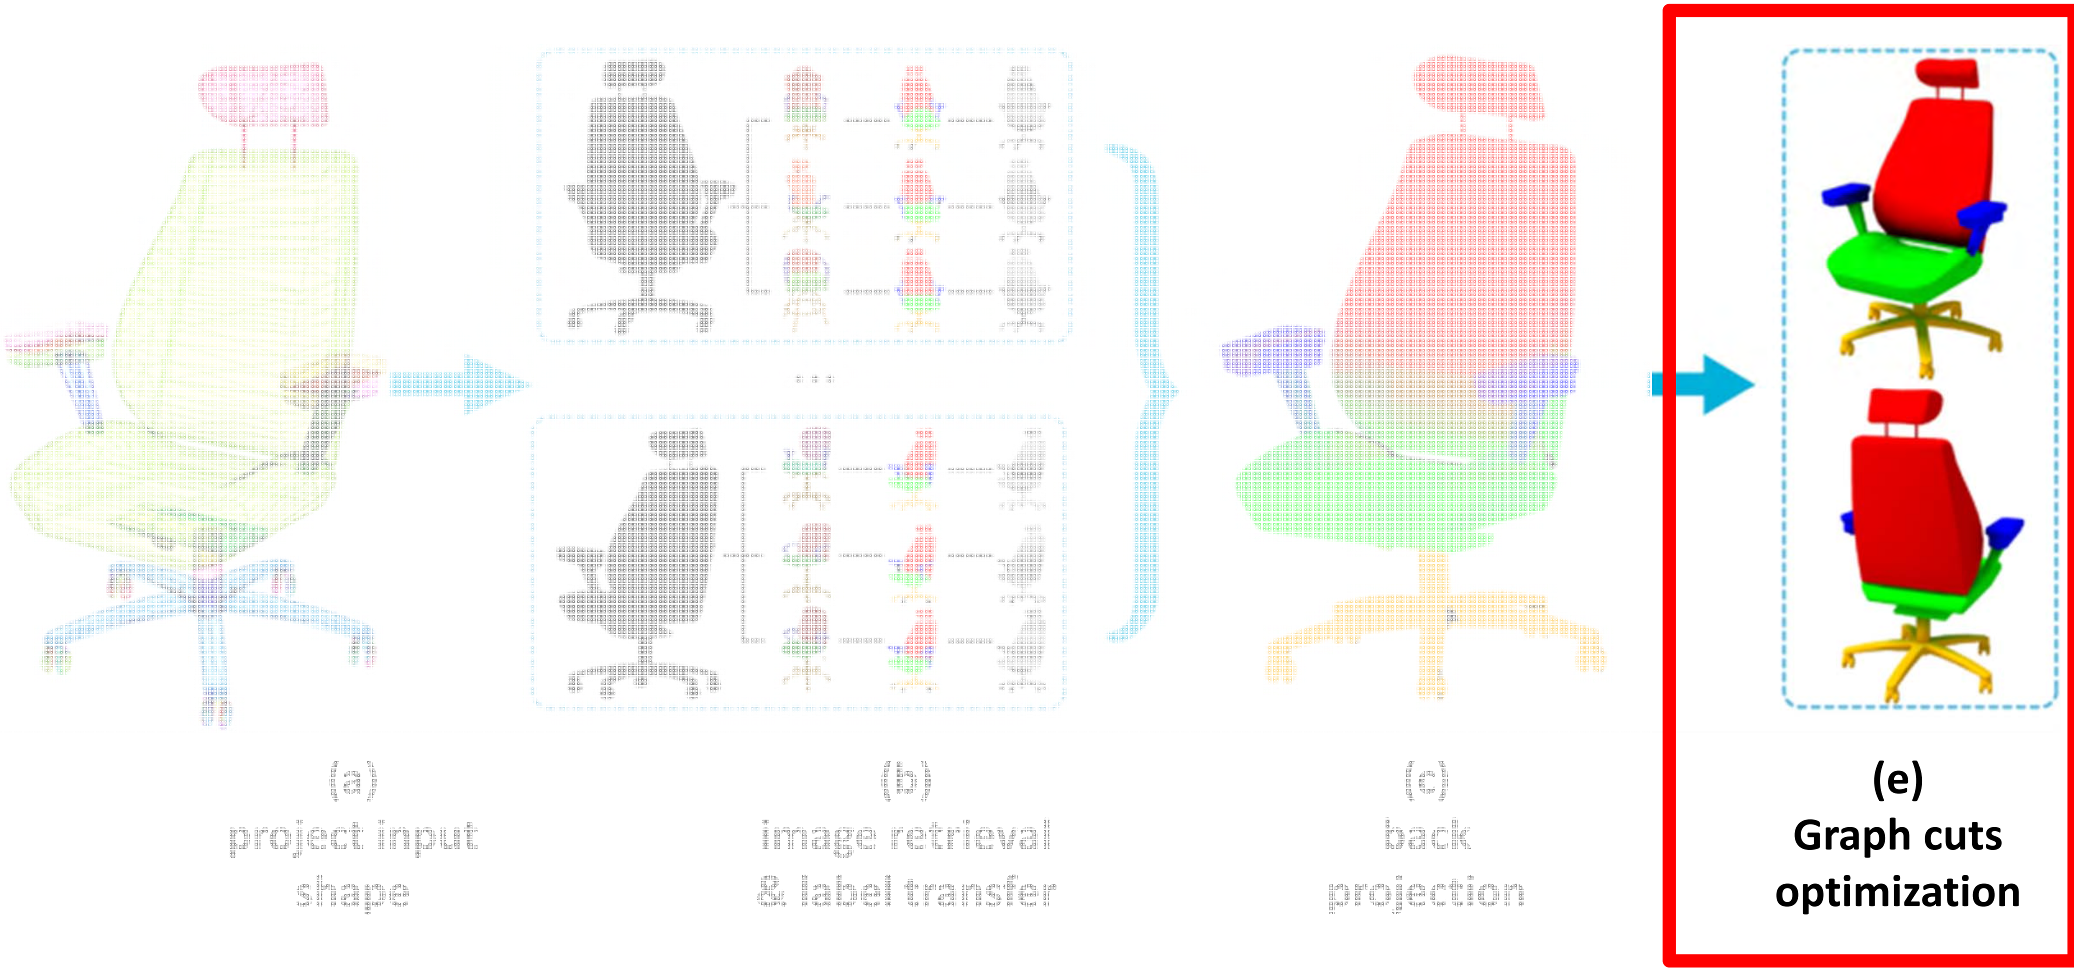

Graph Cuts Optimization

- **Use multi-label alpha expansion graph cuts algorithm**
  - **Data term is based on the probabilistic map**
    - **Each primitive shall be assigned to the label with highest probability**
  - **Smoothness term is based on the connectivity and proximity among the primitives**
    - **Adjacent and connected primitives shall have the same label**

Labeling with Graph Cuts

- **11 object categories:**
  - **About 2600 labeled images in total**
  - **3D shapes tested have self-intersections, as well as other data artifacts**

Category	# parts	# photos	Category	# parts	# photos
Chair	4	500	Stroller	6	400
Truck	3	400	Lamp	3	344
Vase	4	300	Table	2	250
Bike	5	181	Pavilion	3	60
Guitar	3	20	Fourleg	5	234
Robot	4	174			

Results: Real Image Training Dataset

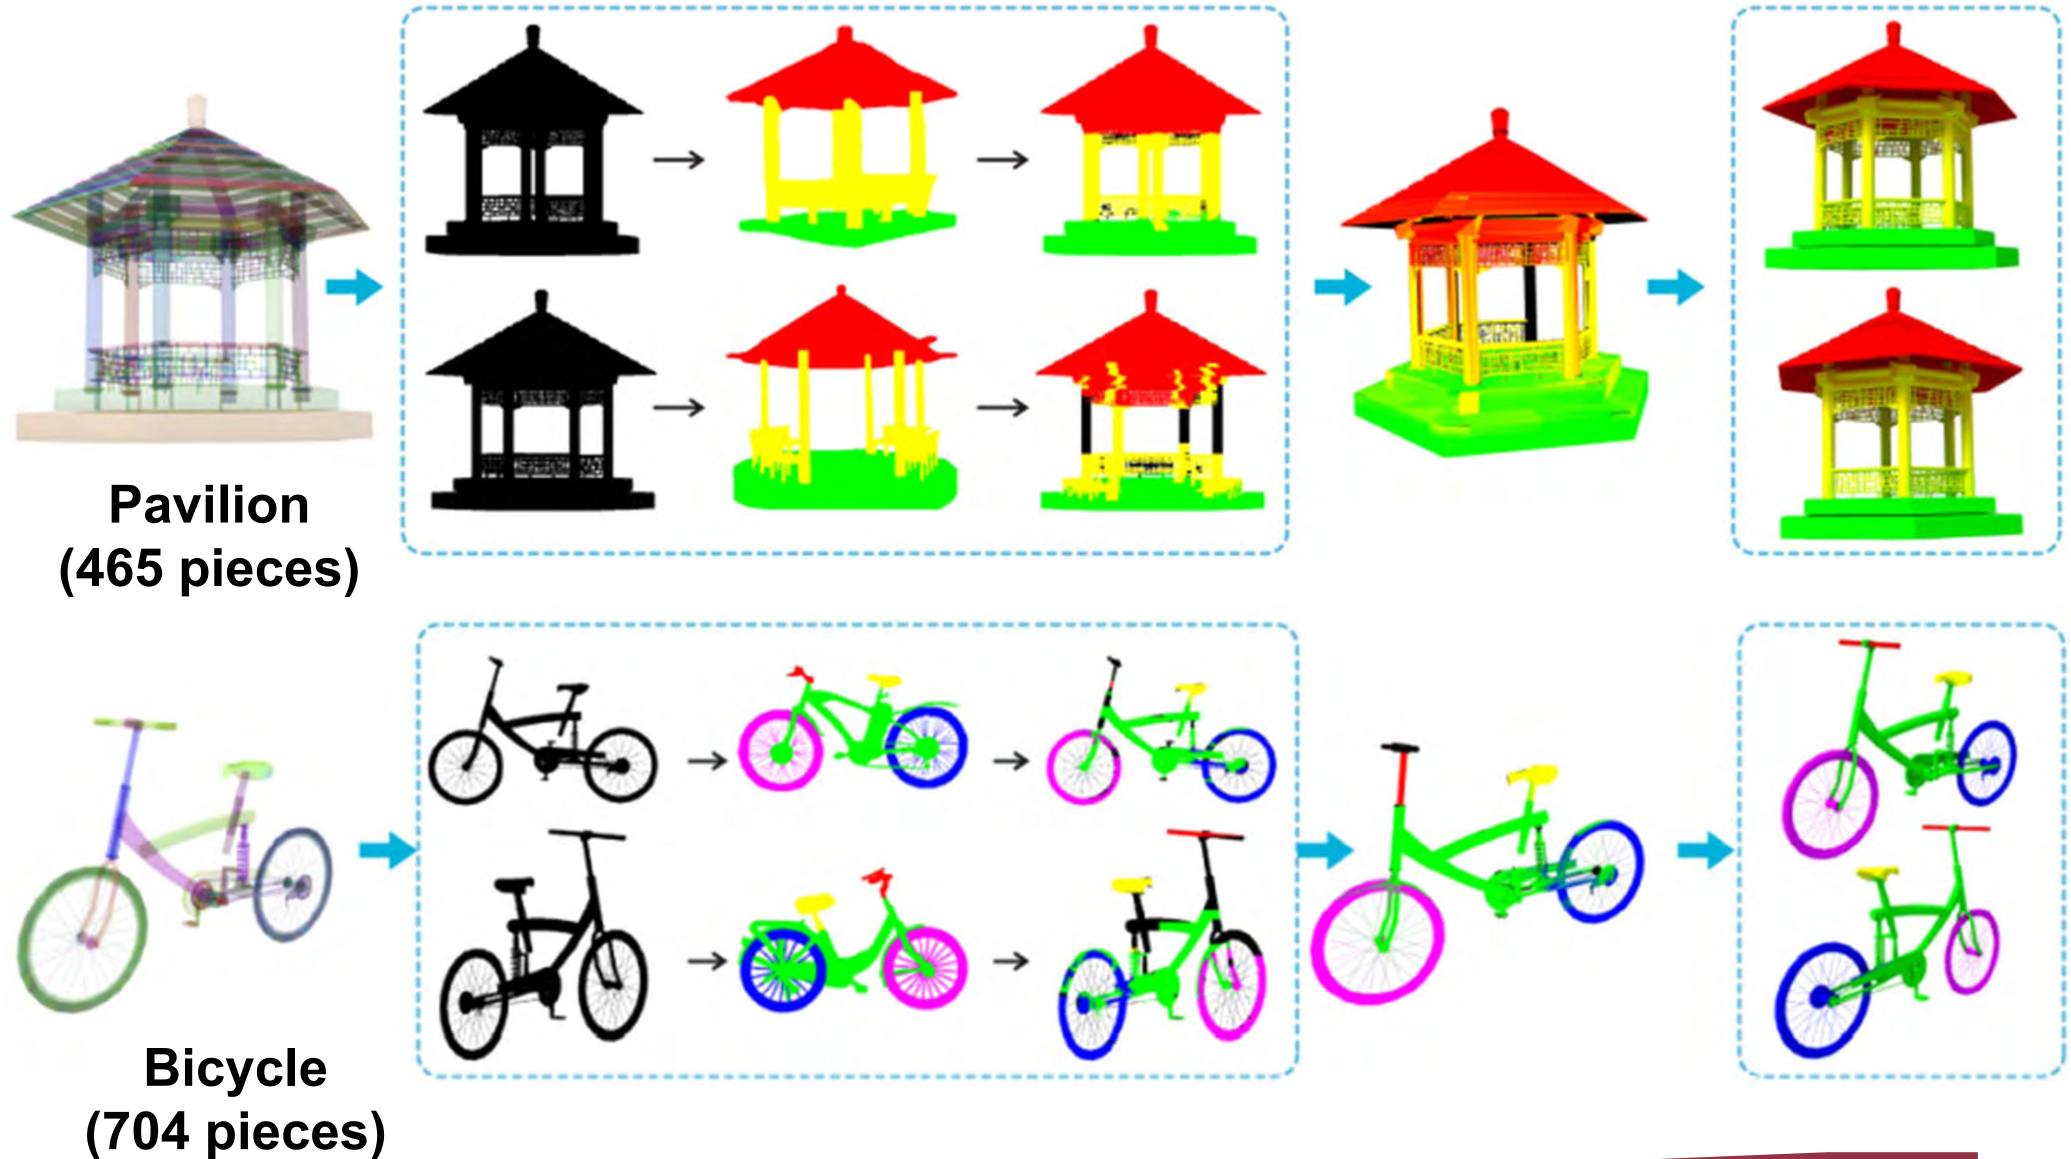

Complex Topology

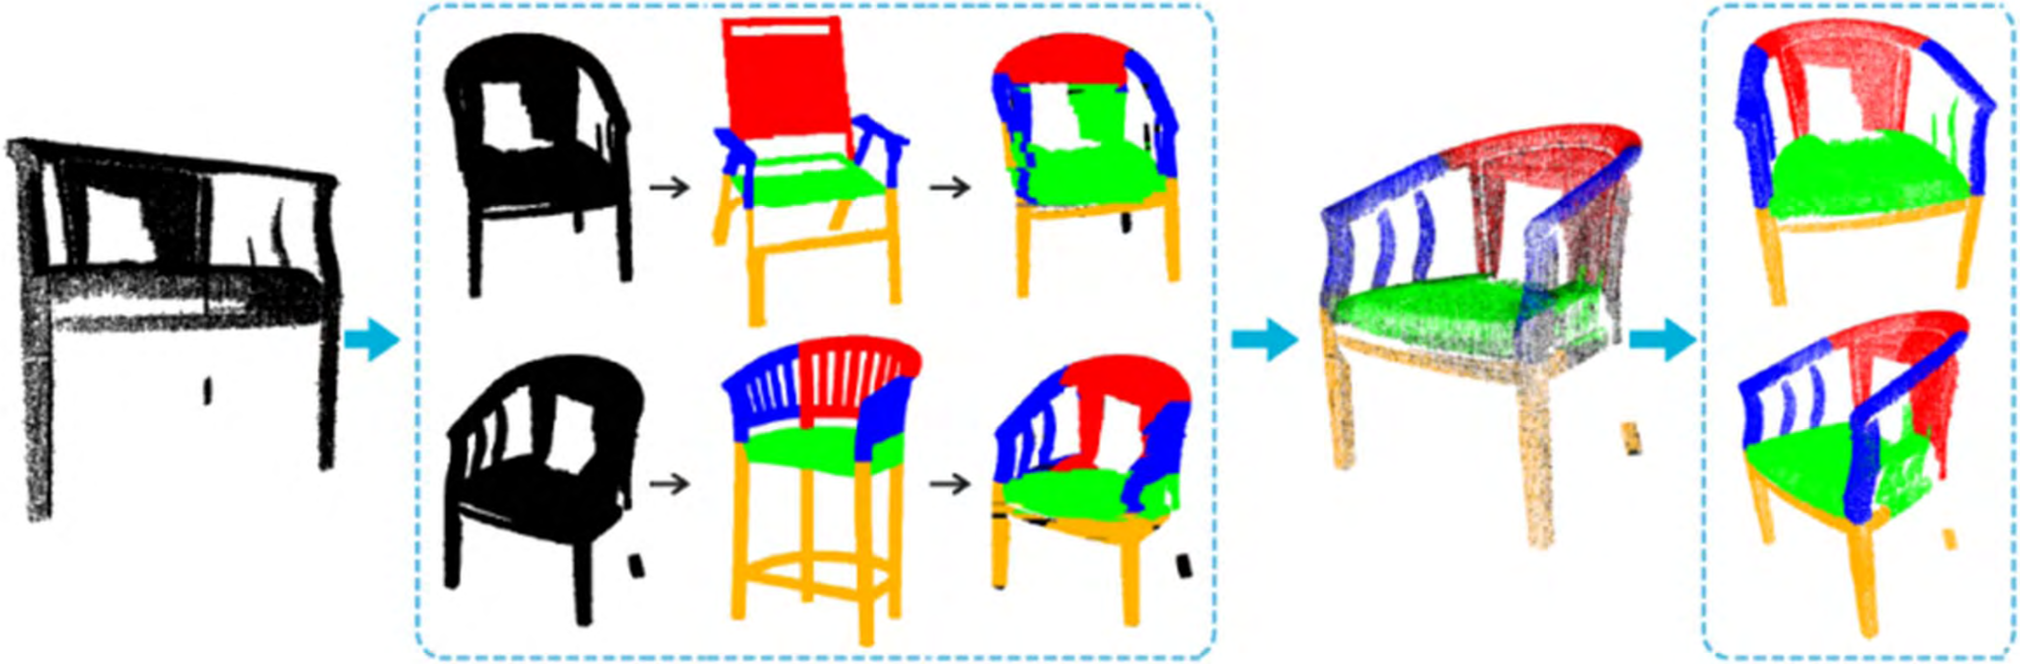

Incomplete Model: Point Clouds

- **Comparison with data-driven mesh segmentation**
  - **Models are manifold, complete, & no self-intersections**
  - **Training images are projections of labeled 3D models**
  - **Same inputs, training data, & experimental settings**

	Ours				Kalogerakis et. al [2010]			
Set	SB19	SB12	SB6	SB3	SB19	SB12	SB6	SB3
Chair	**99.2**	**99.6**	**97.9**	93.4	98.5	98.4	97.8	**97.1**
Table	**99.6**	**99.6**	**99.6**	**99.4**	99.4	99.3	99.1	99.0
Vase	**91.9**	**90.5**	**89.7**	**80.8**	87.2	85.8	77.0	74.3
Mech	94.6	**91.3**	**90.2**	**90.6**	94.6	90.5	88.9	82.4
Cup	99.1	**99.6**	97.5	94.4	**99.6**	96.0	**99.1**	**96.3**
Fourleg	67.9	54.3	59.1	58.6	**88.7**	**86.2**	**85.0**	**82.0**
Human	63.8	55.6	51.1	48.0	**93.6**	**93.2**	**89.4**	**83.2**

Quantitative Evaluation

- **Inherent from projective analysis:**
  - Do not fully capture 3D info
- **Inherent to data-driven approaches:**
  - Knowledge has to be in data
- **Upright assumption:**
  - Not designed for articulated shapes

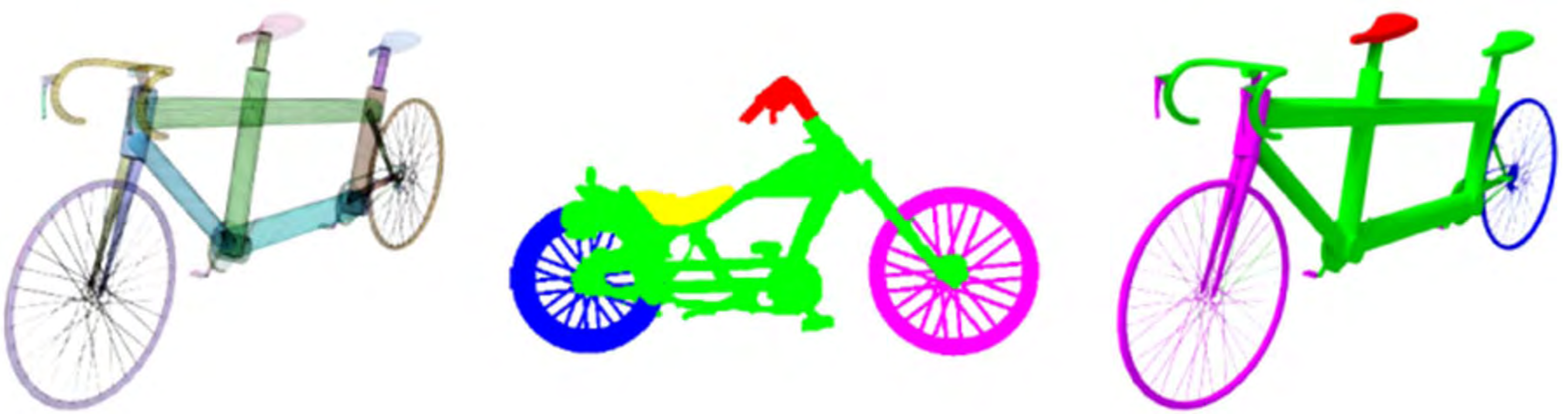

Limitations


Thank
you

- An image worth a thousand words. From images, we humans are able to infer the 3D shape of an object and to decompose the object into semantically meaningful parts. Now, is it possible to teach computers to do these tasks?
- Two recent research projects that work along this direction will be presented in this talk. The first one investigates how the 3D modeling of flower head can be facilitated using a single photo of the flower. The core idea is that flower head typically consists of petals of similar 3D geometries, yet their observed shapes on 2D images vary due to differences in projecting directions. Exploiting this variation allows us to reconstruct the 3D geometry of the petals from a single image.
- The second project studies how to segment 3D models into semantically meaningful parts based on knowledge learned from labeled 2D images. Here the input 3D model is treated as a collection of 2D projections, which are labeled using training images of similar objects. The 3D model is then segmented by summarizing the labeling for its projections. Here the key is, for each query projection, how to retrieve objects with similar semantic parts and transfer their labels over.

Abstract

- Feilong Yan, Minglun Gong, Daniel Cohen-Or, Oliver Deussen, & Baoquan Chen: Flower reconstruction from a single photo. Computer Graphics Forum 33(2 - Eurographics): 439-447. May 2014.

- Yunhai Wang, Minglun Gong, Tianhua Wang, Daniel Cohen-Or, Hao Zhang, & Baoquan Chen: Projective analysis for 3D shape segmentation. ACM Transactions on Graphics 32(6 - SIGGRAPH Asia). November 2013.

Related Publications