thrive SIGGRAPH2019
LOS ANGELES • 28 JULY – 1 AUGUST

DEEP VIEW SYNTHESIS FROM SPARSE PHOTOMETRIC IMAGES

Zexiang Xu[1], Sai Bi[1], Kalyan Sunkavalli[2], Sunil Hadap[3], Hao Su[1], Ravi Ramamoorthi[1]

[1]University of California, San Diego
[2]Adobe Research
[3]Lab 126, Amazon

Render real scenes

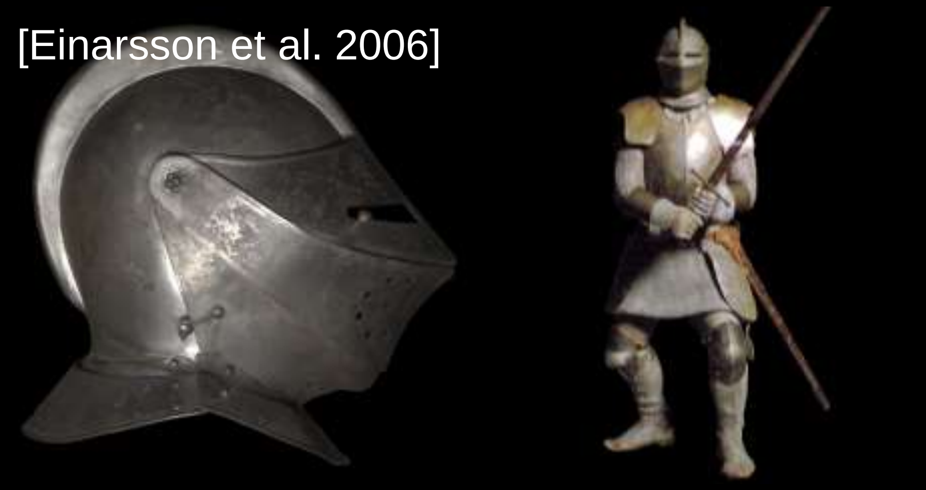
[Einarsson et al. 2006]

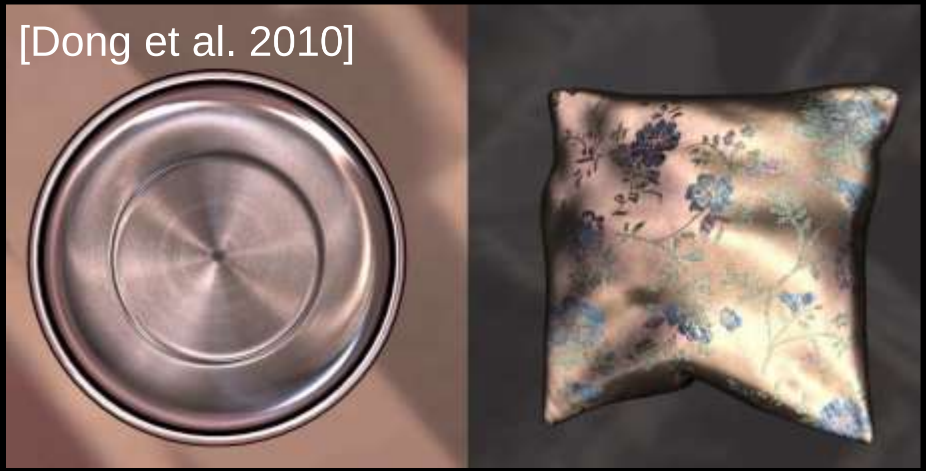
[Dong et al. 2010]


[Schwartz et al. 2011]


[Zickler et al. 2005]

Appearance of a scene

- Geometry

[Furukawa and Ponce 2008]

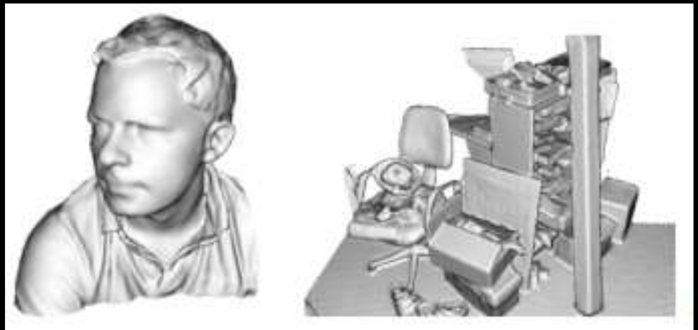

[Newcombe et al. 2011]

- Materials

[Xu et. al 2016]

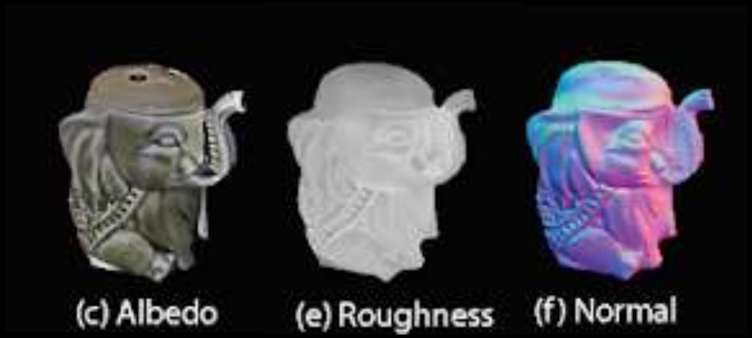

[Li et. al 2018]

Appearance of a scene

- Geometry

- Materials

- Realistic rendering

Appearance of a scene

- Geometry

- Materials

- Realistic rendering

Appearance of a scene

Light
Transport
Function

- Realistic rendering

[Einarsson et al. 2006]

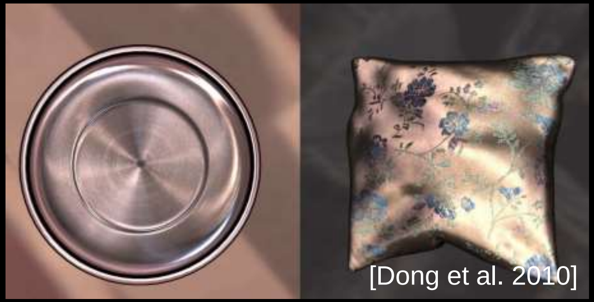

[Dong et al. 2010]

[Schwartz et al. 2011]

Light transport acquisition

Light
Transport
Function

[Matusik et al. 2002]

Image-based relighting

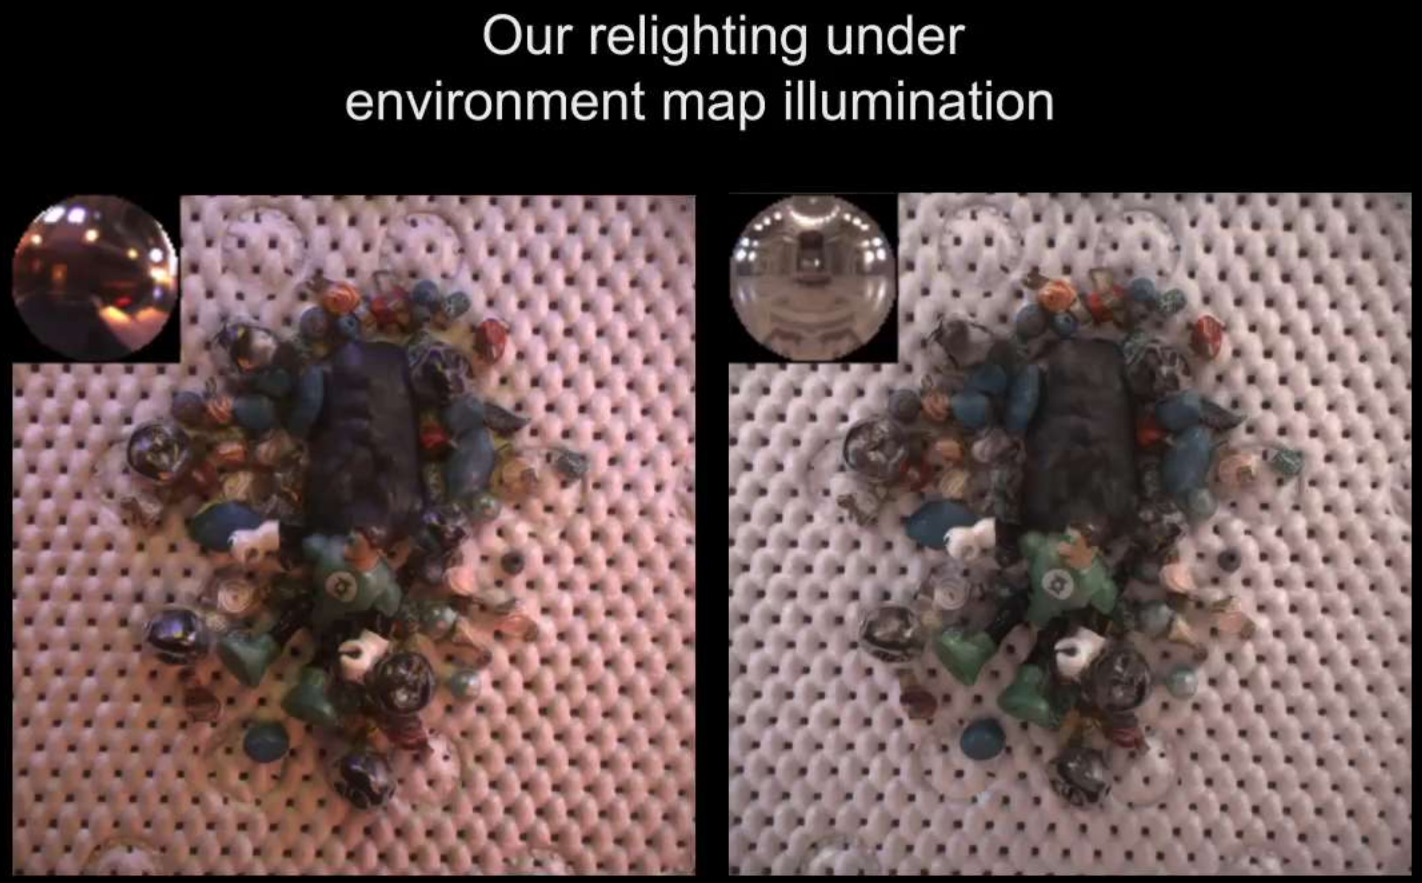

[Xu et al. 2018]

Light transport acquisition for changing view

Sparse input views

Novel view appearance

Novel view synthesis

[Chen and Williams 1993]


[Flynn et al. 2016]


[Kalantari et al. 2016]


[Levoy and Hanrahan 1996]


[Penner and Zhang 2017]

[Zhou et al. 2018]

- Unstructured views
- Small baseline
- Natural illumination

Sparse sampling for light transport acquisition

- Large baseline
- Controlled lighting

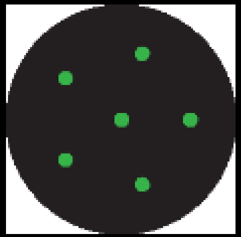

Preview

- Large baseline
- Controlled lighting

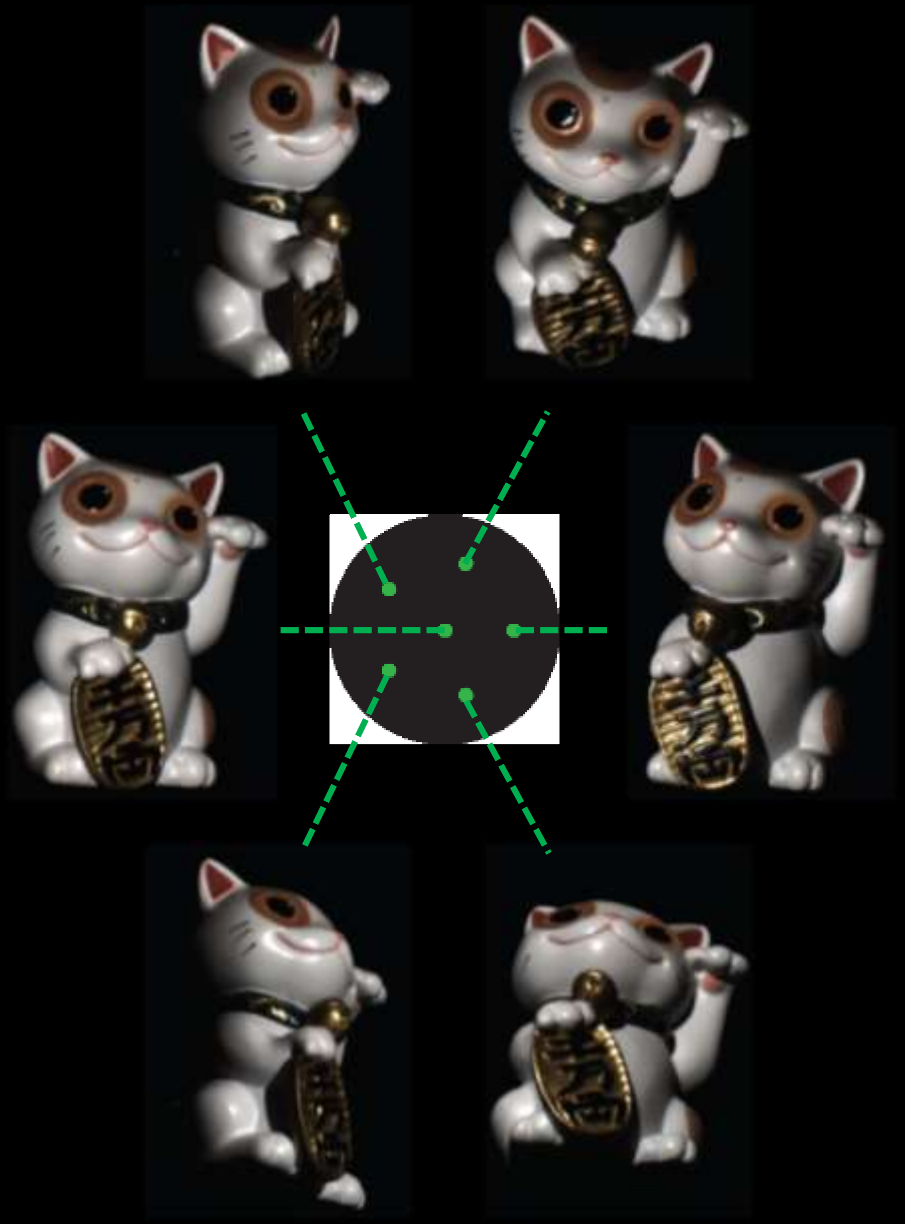

Preview

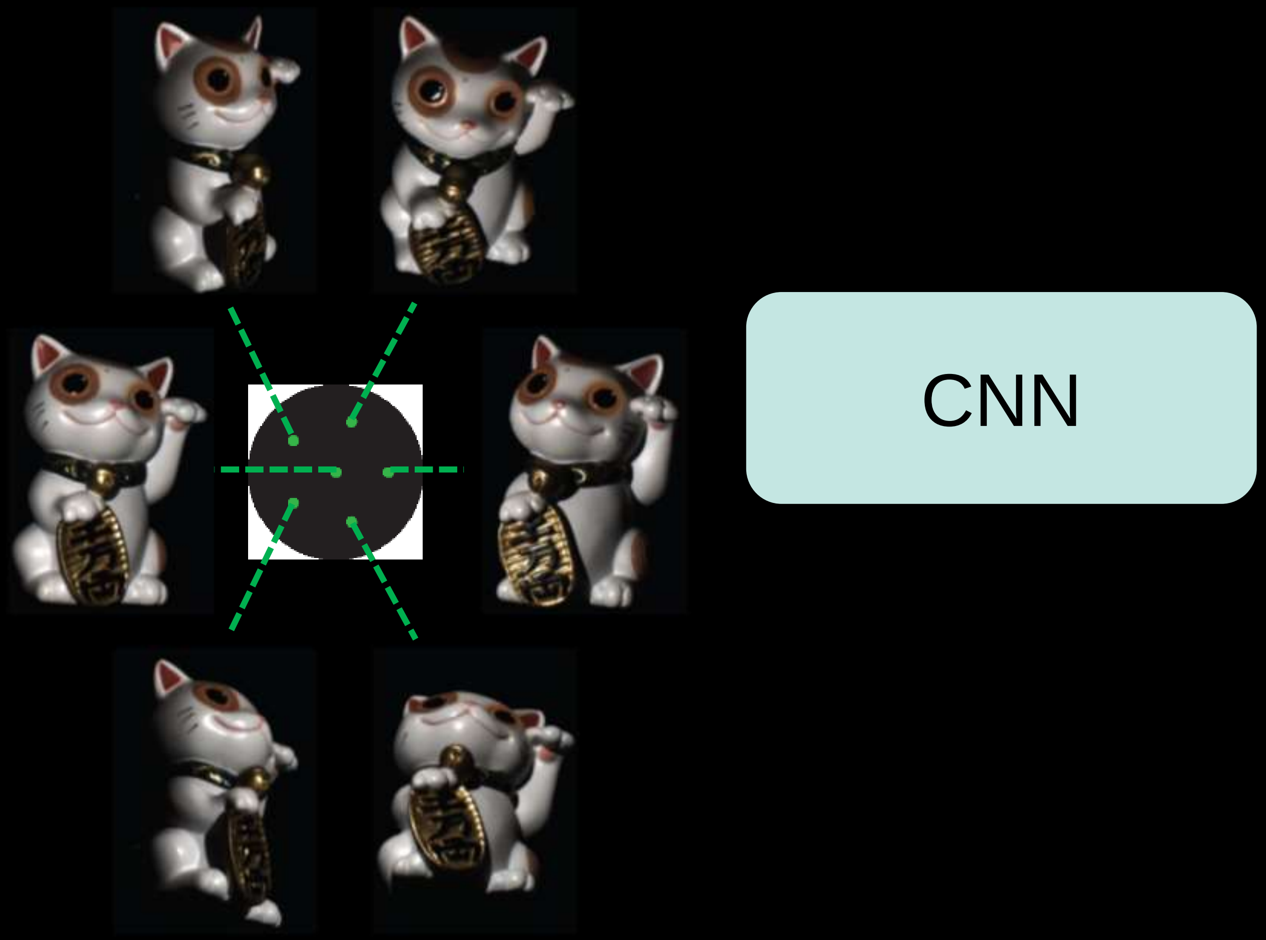

Preview

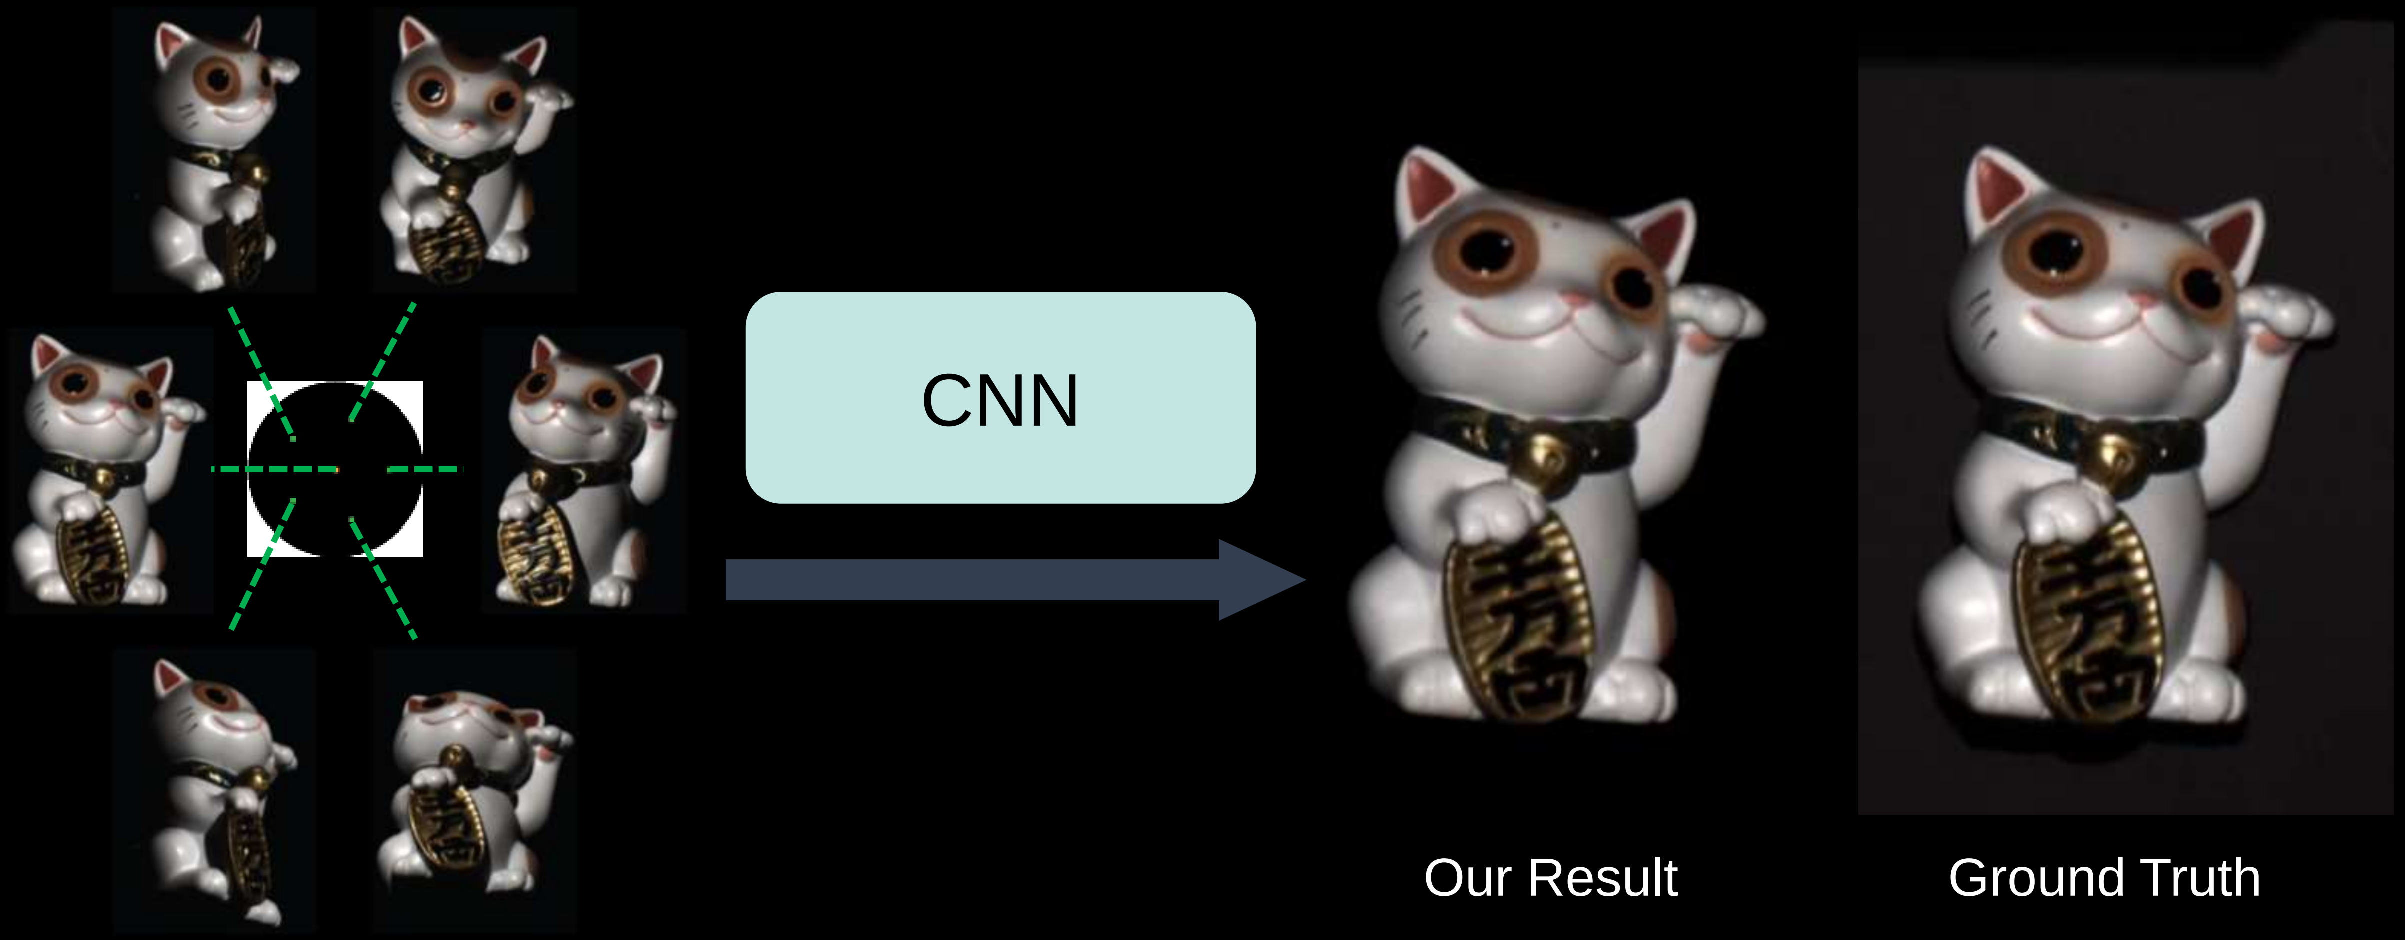

Acquisition configuration

- Sparse
- Good coverage

Acquisition configuration

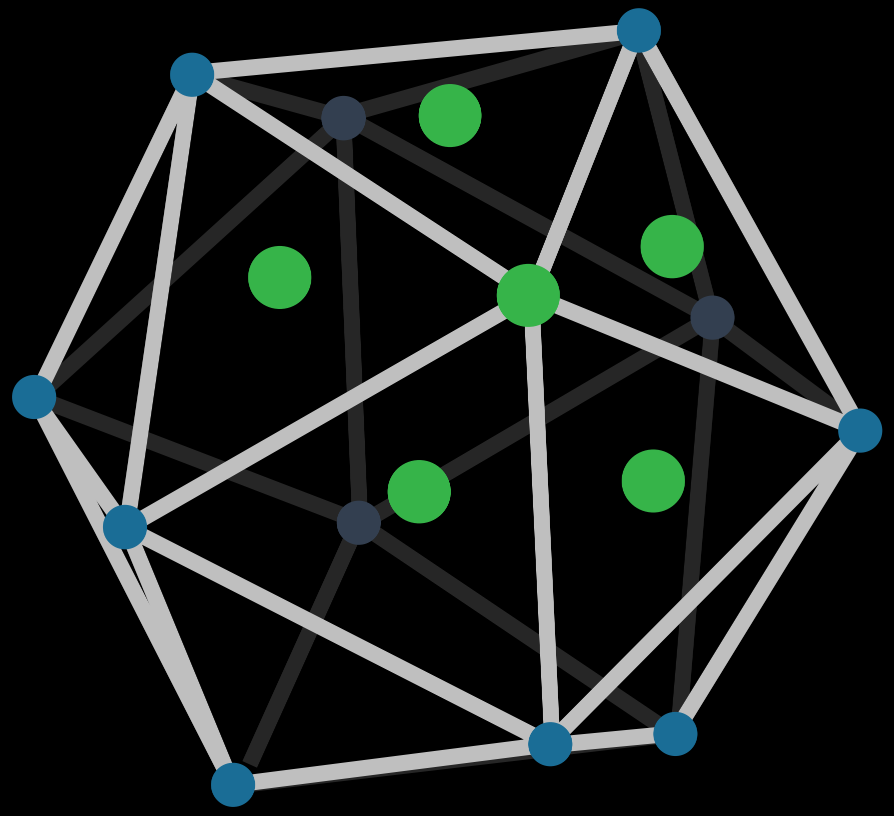

Icosahedron

- 12 vertices
- 20 faces
- Symmetric

Acquisition configuration

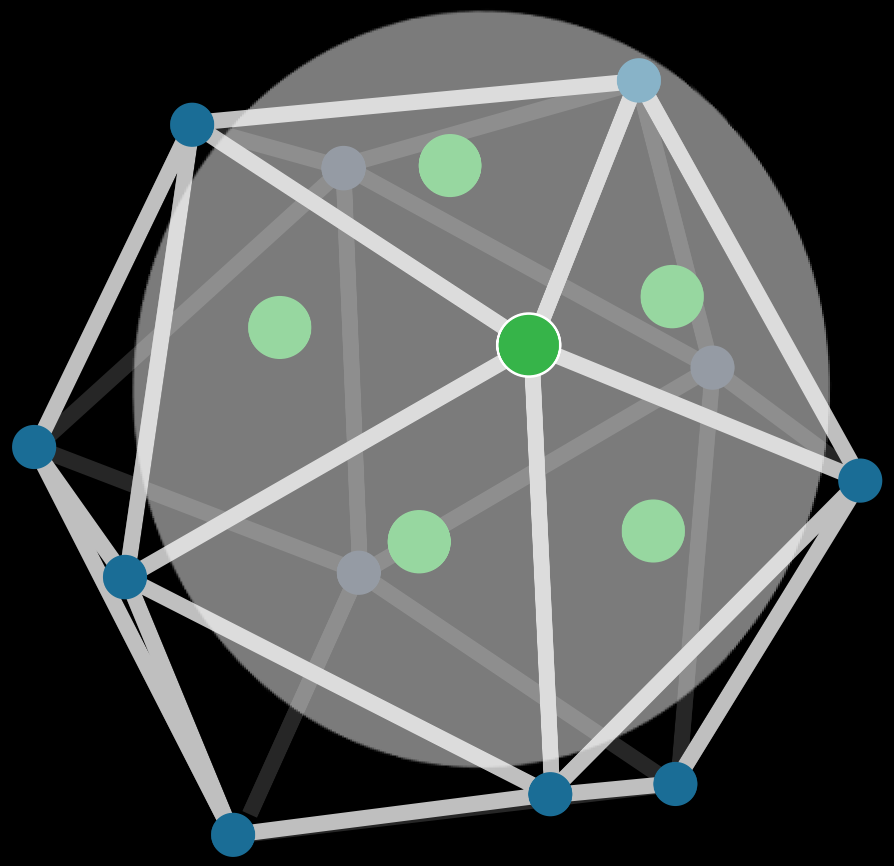

Icosahedron

Acquisition configuration

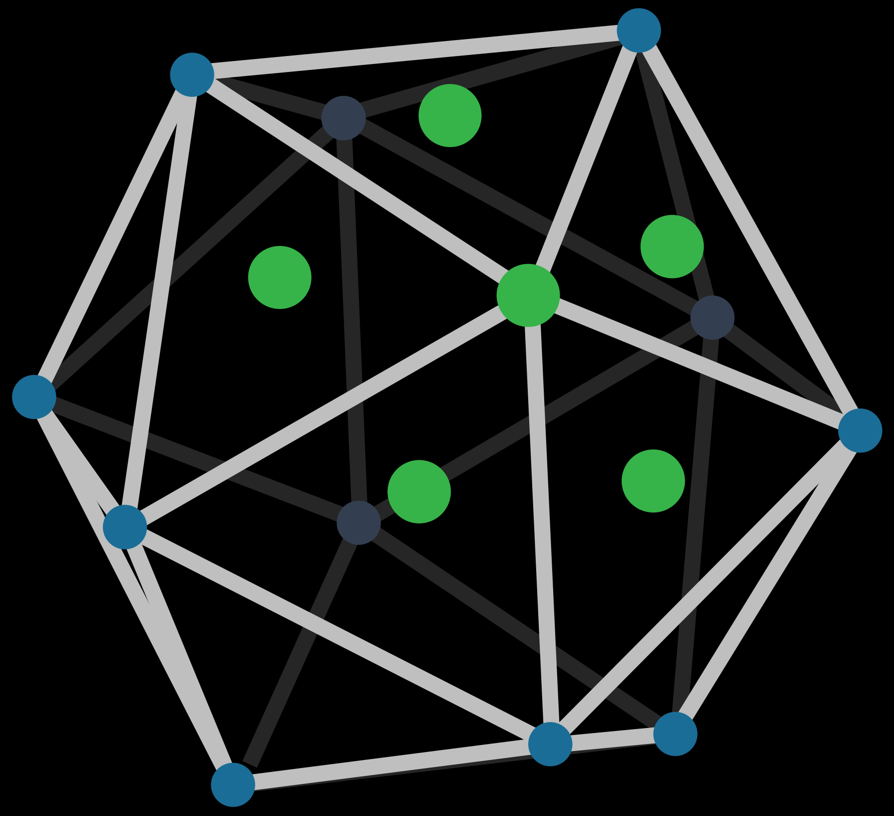

Icosahedron

Acquisition configuration

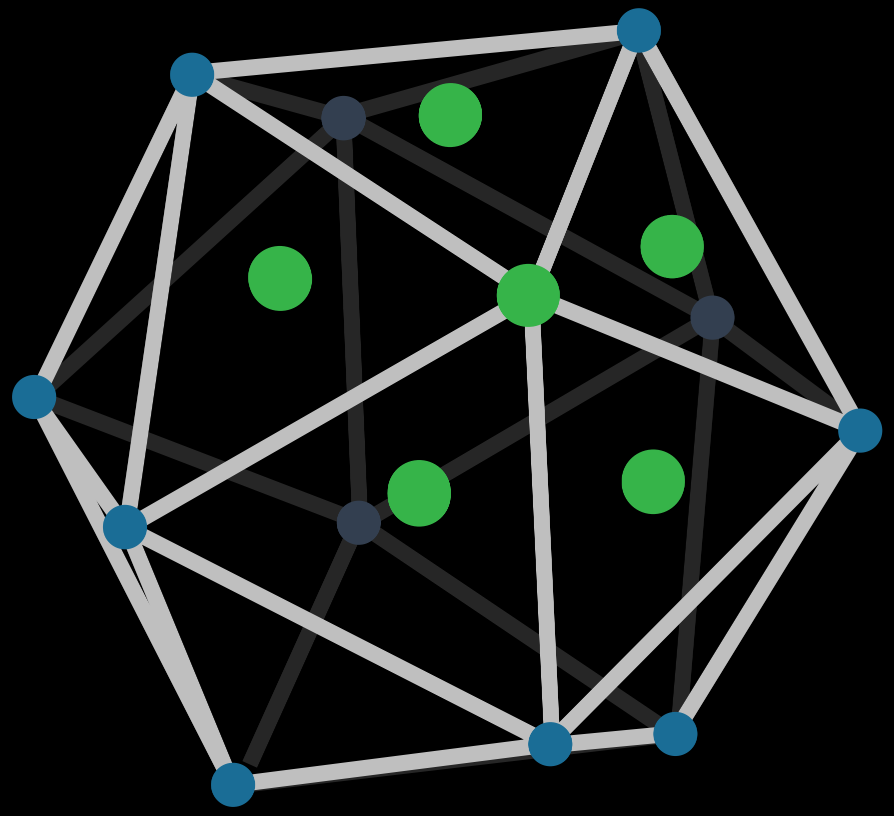

Icosahedron

Acquisition configuration

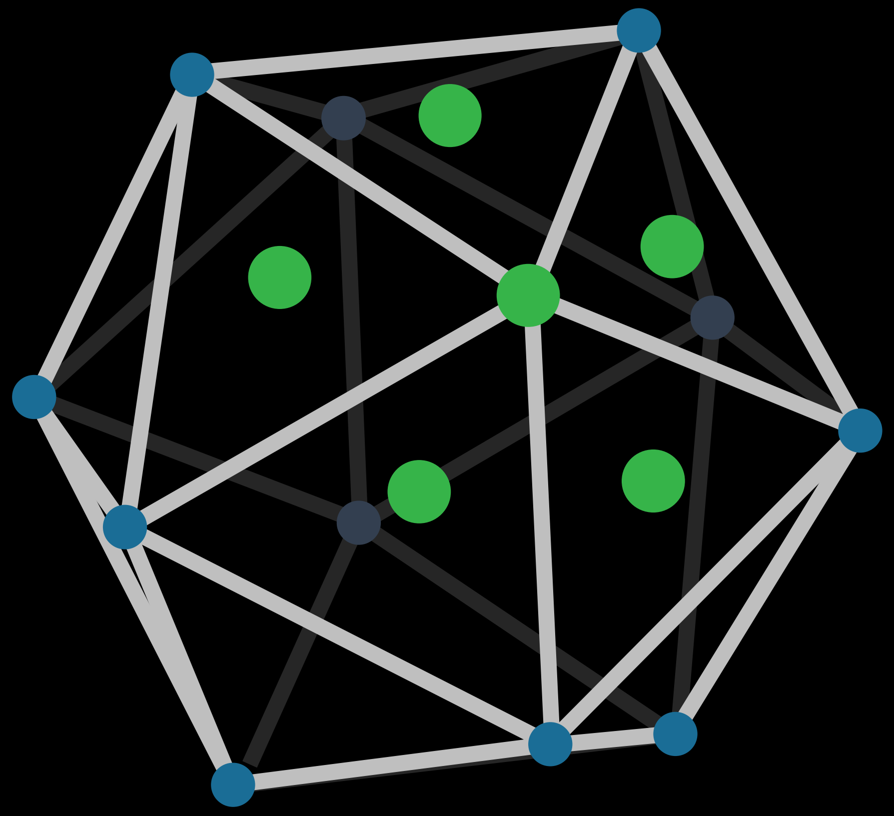

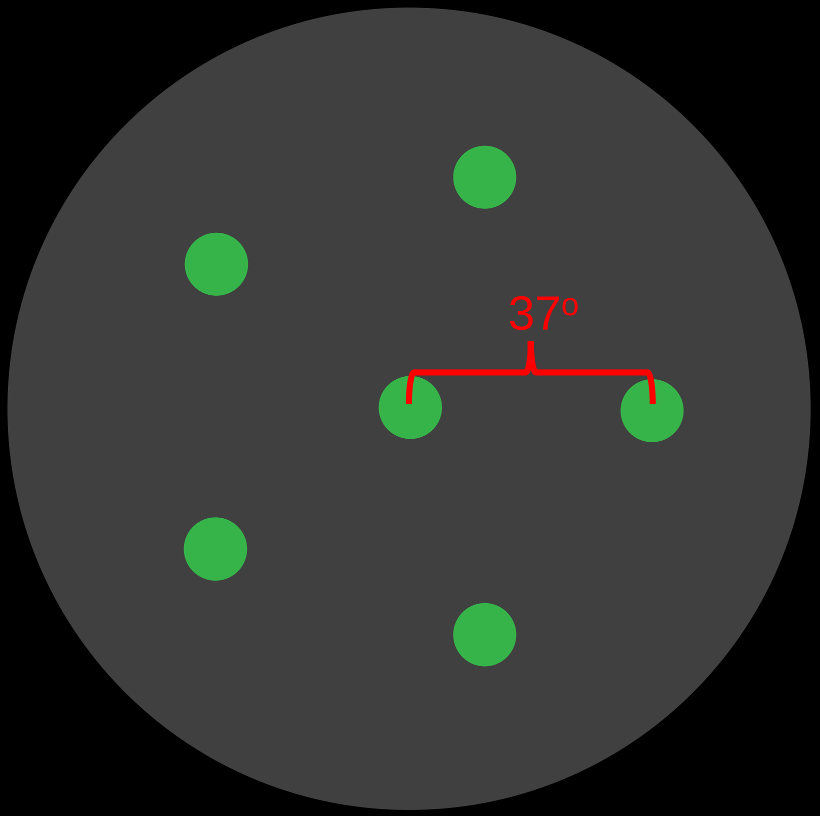

Acquisition configuration

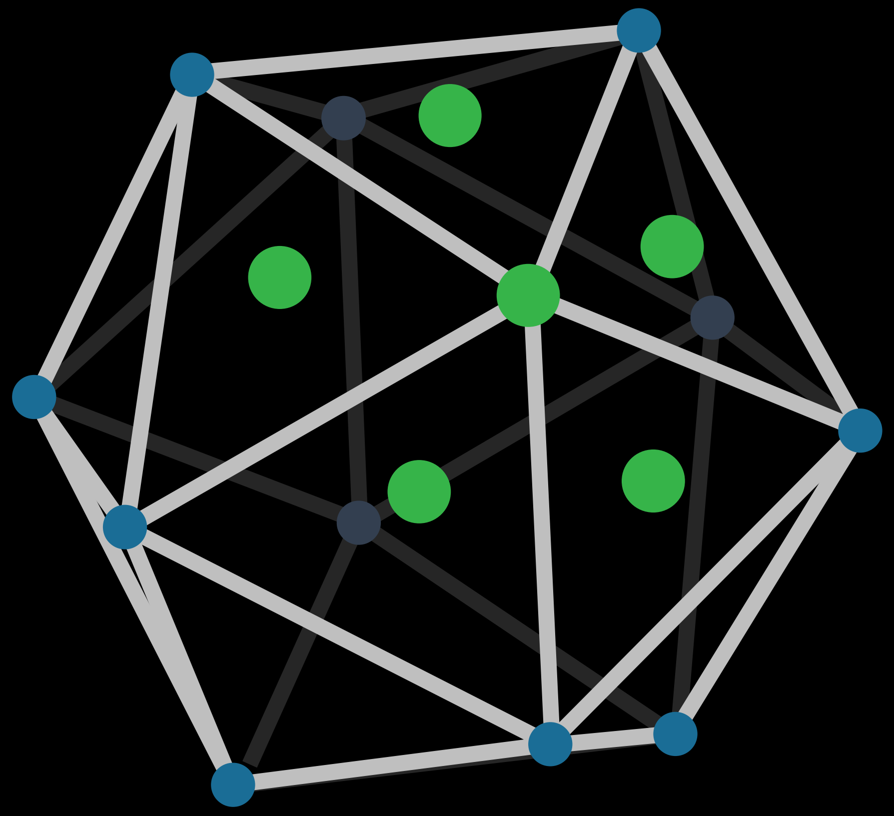

Icosahedron

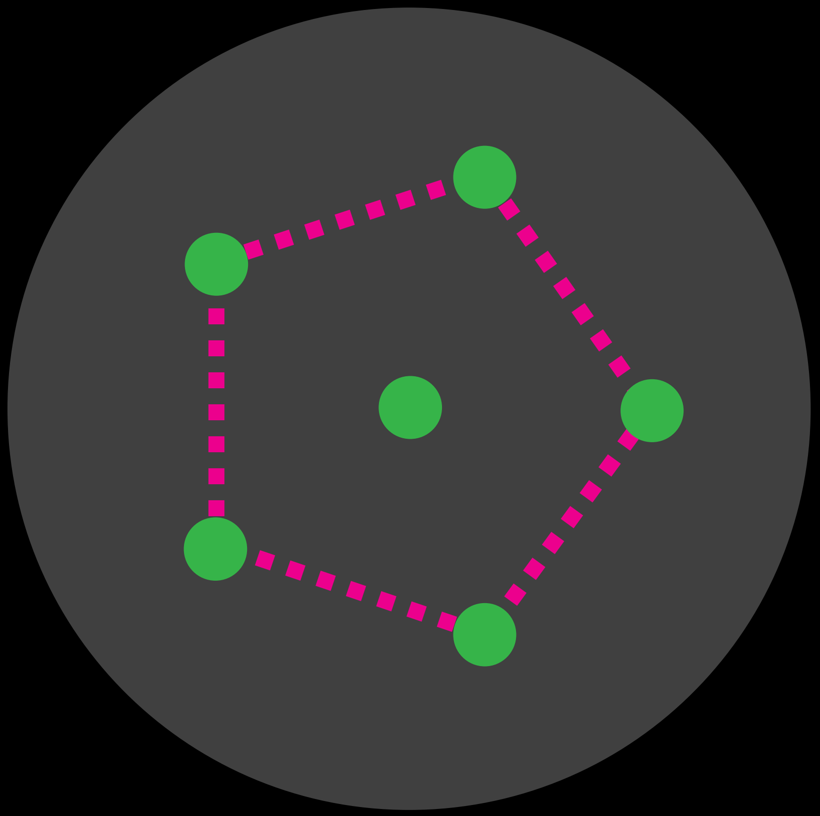

Synthetic scenes

Geometry:

Procedurally Generated Objects

Material images courtesy: Allegorithmic and Adobe Stock

Reflectance:

Adobe Stock Material

Synthetic scenes

Geometry:

Procedurally Generated Objects

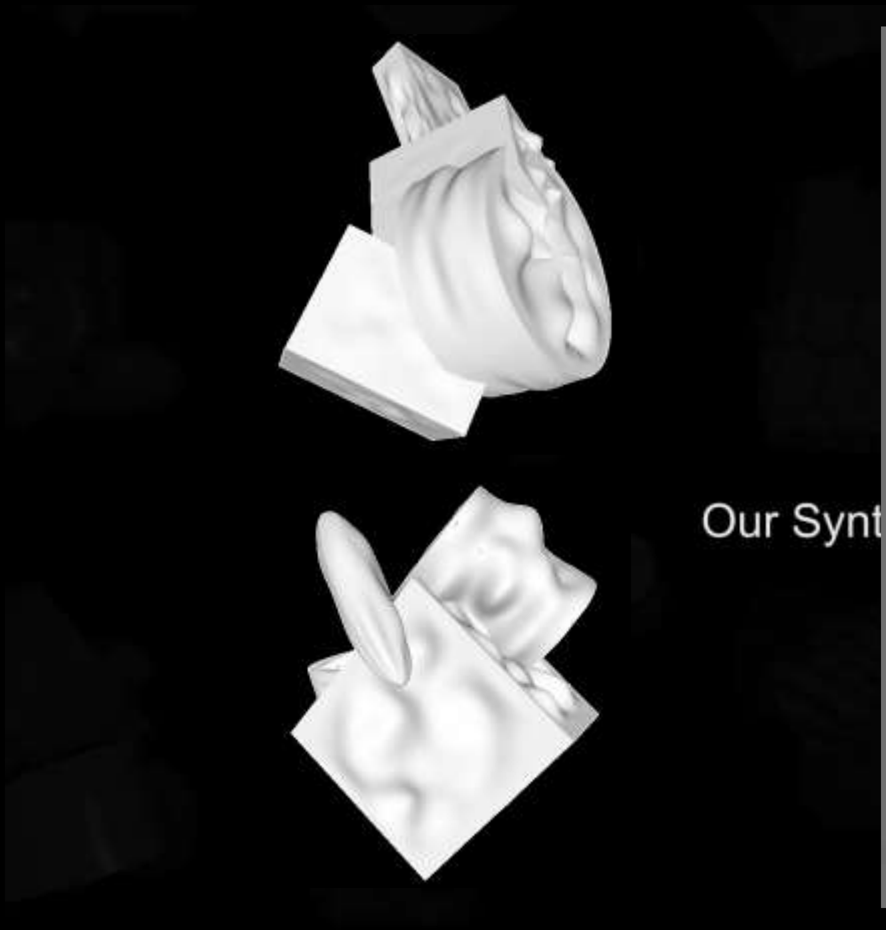

Reflectance:

Adobe Stock Material

Overview

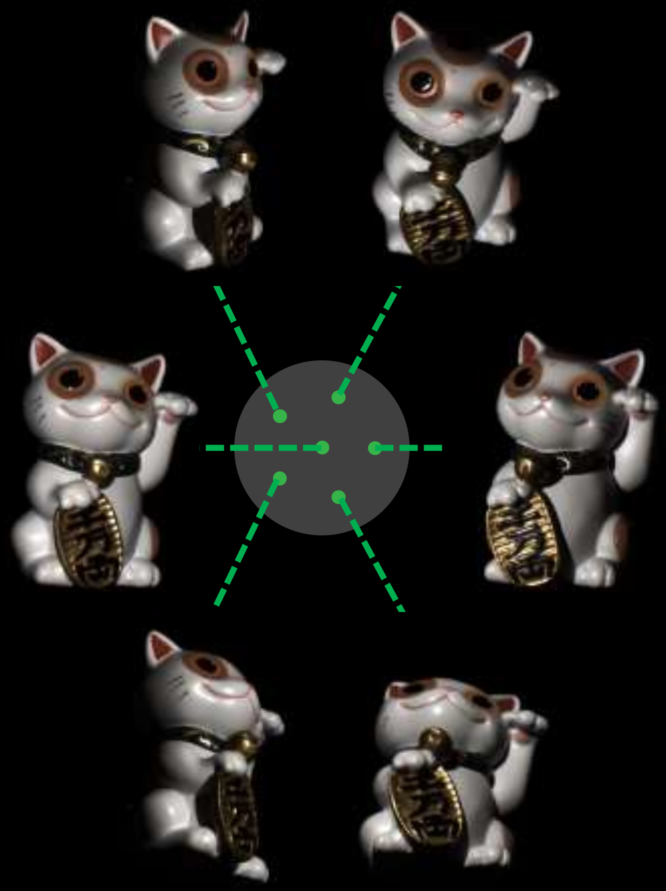

Input views

CNN

Overview

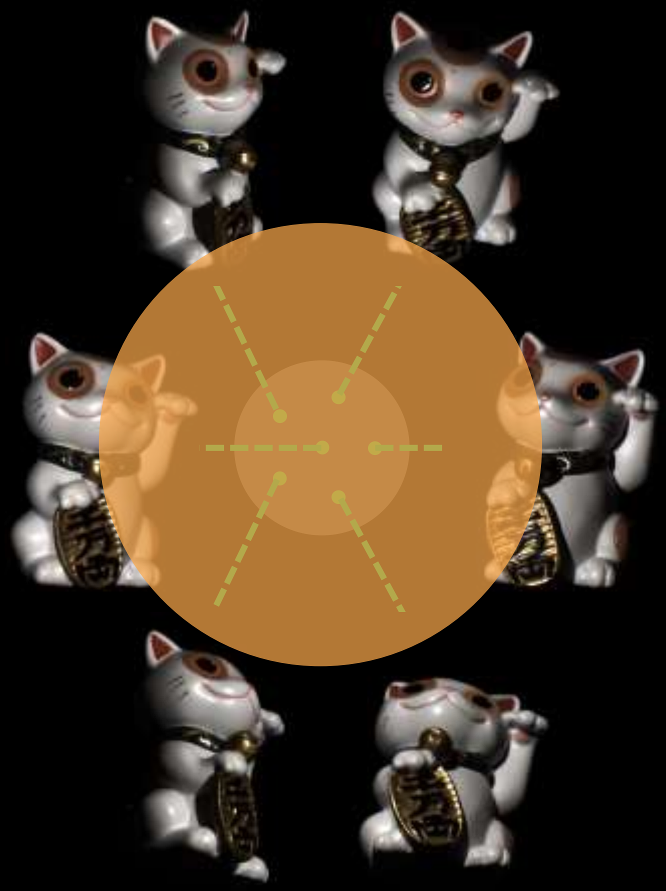

Input views

CNN

Overview

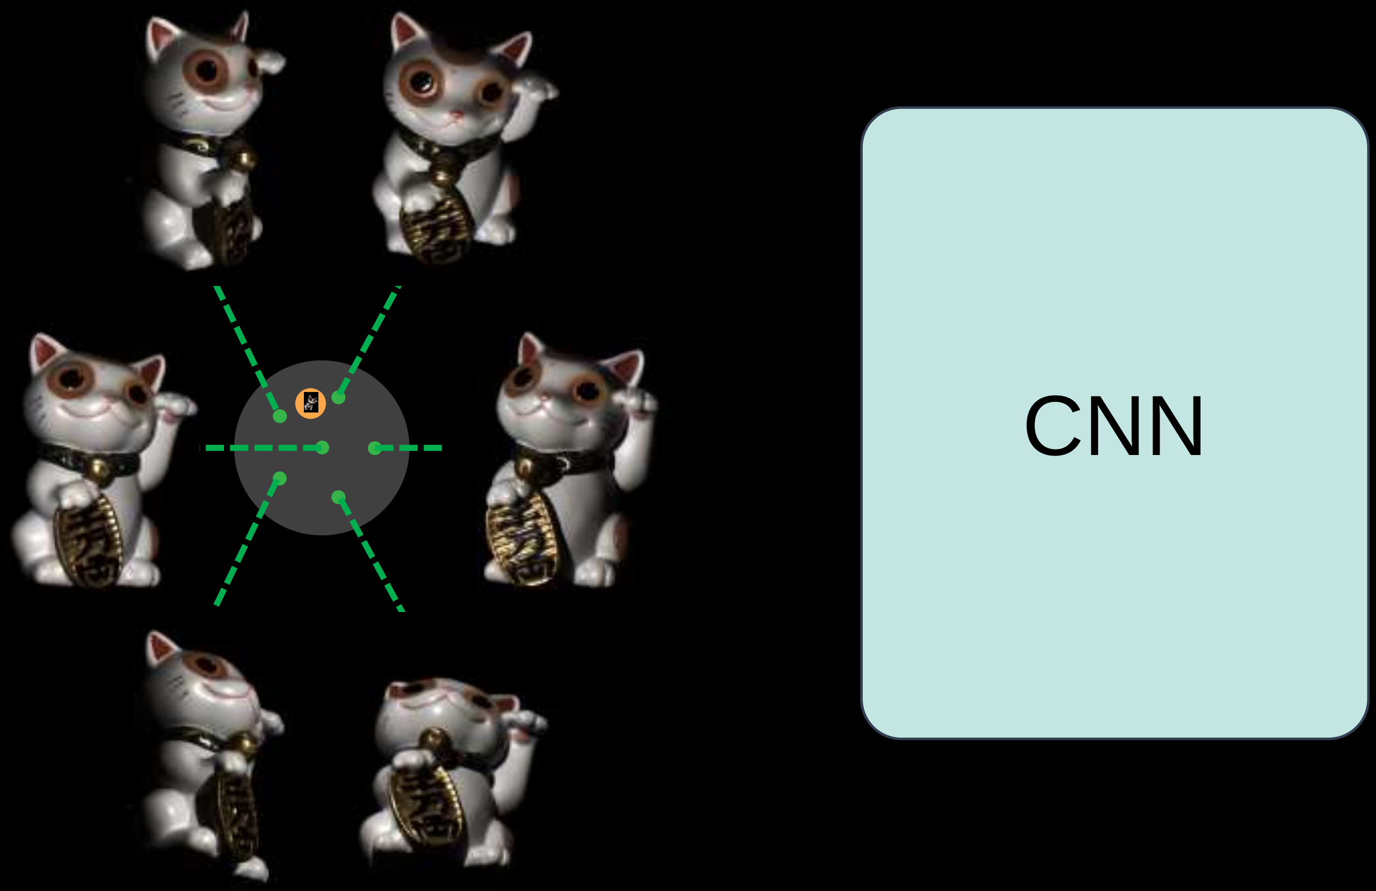

Input views

Overview

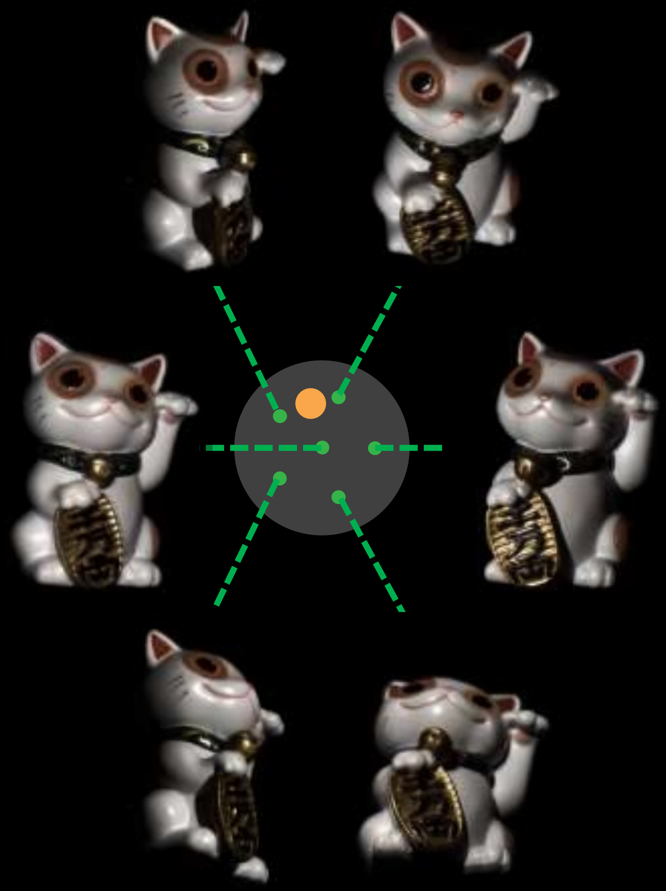

Input views

CNN

Novel view

Overview

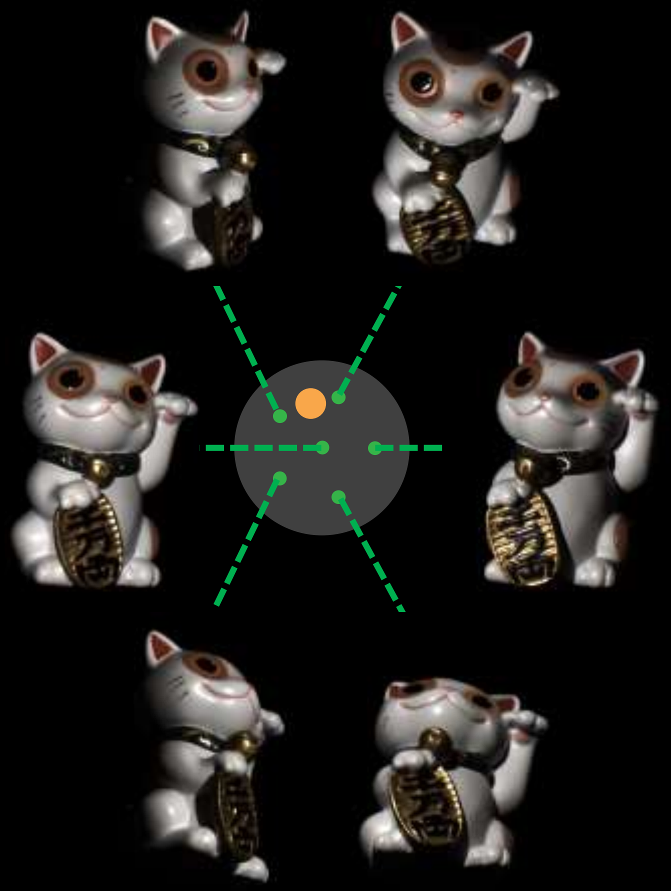

Input views

Novel view

Plane sweep volume

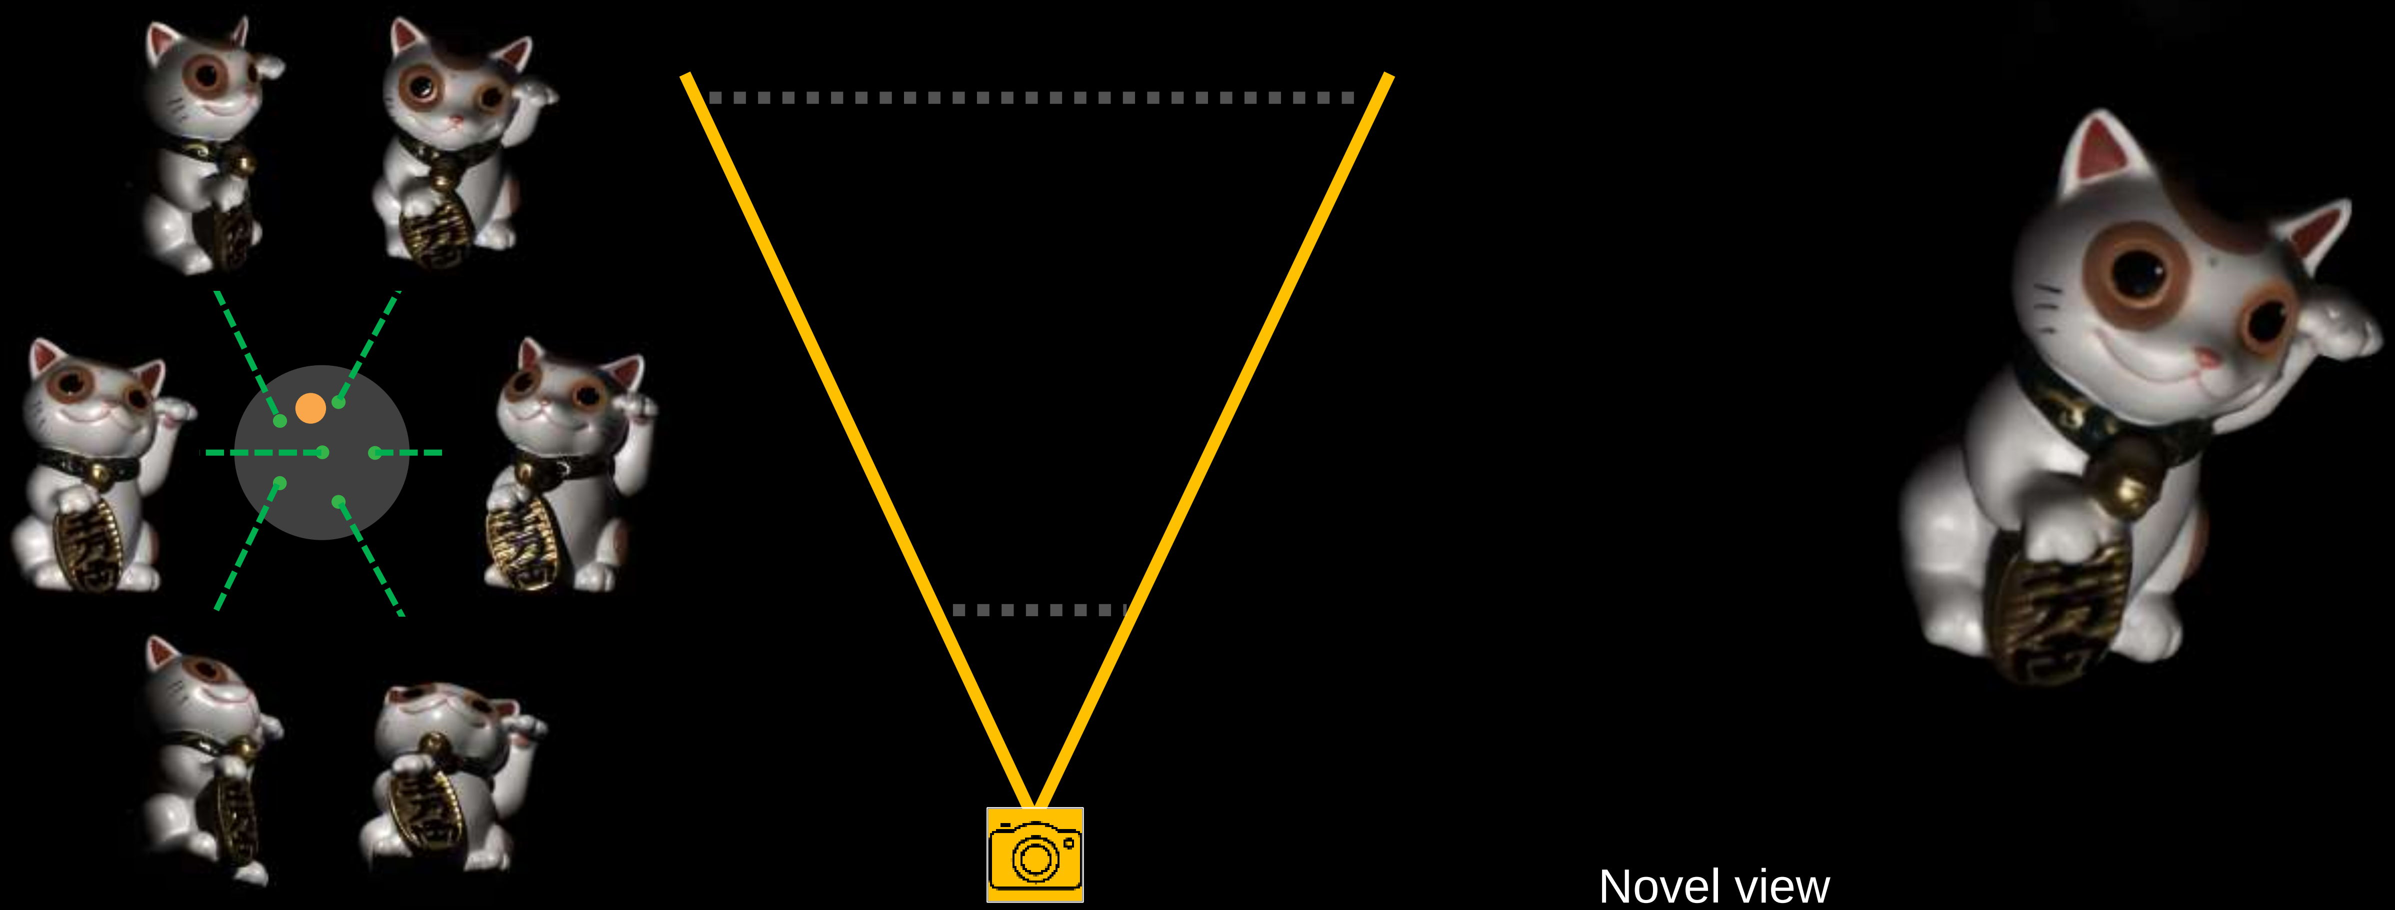

Input views

Novel view

Plane sweep volume

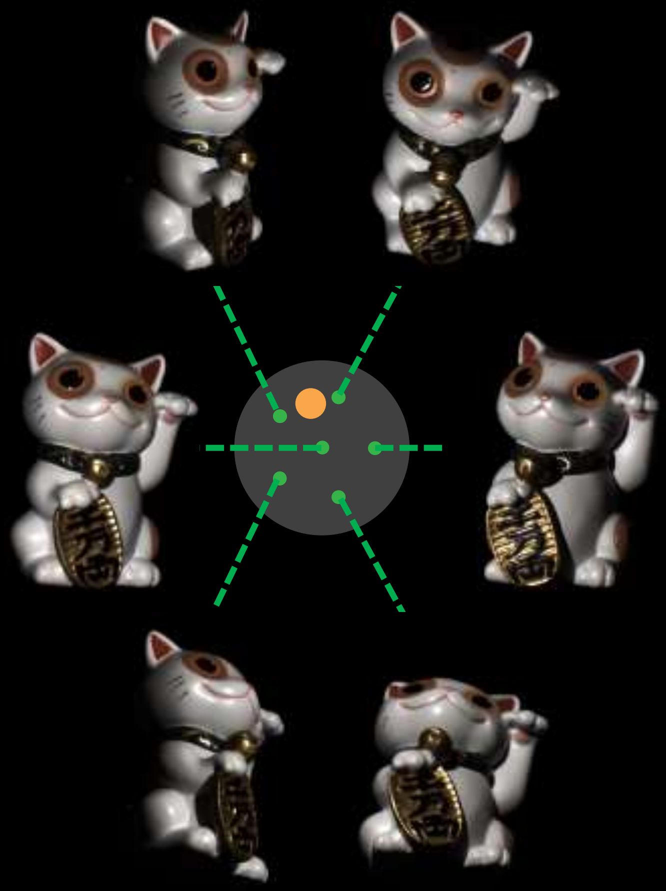

Input views

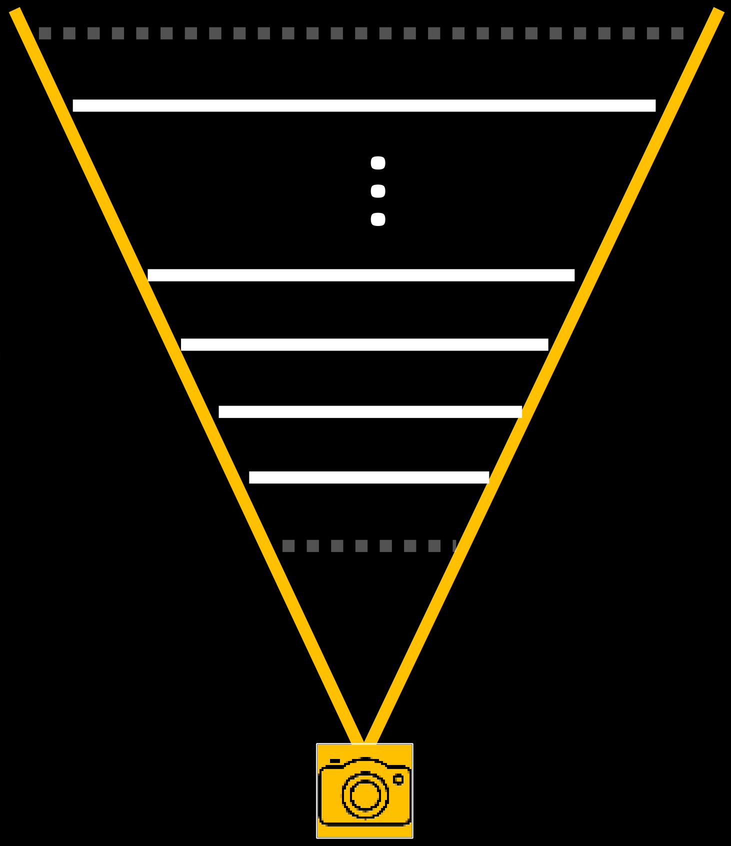

Novel view

Plane sweep volume

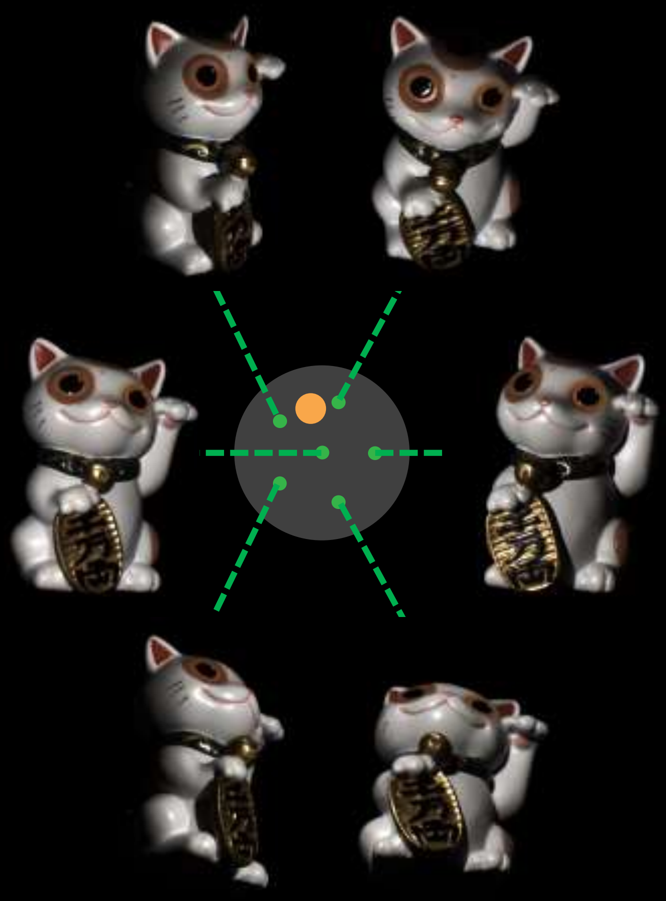

Input views

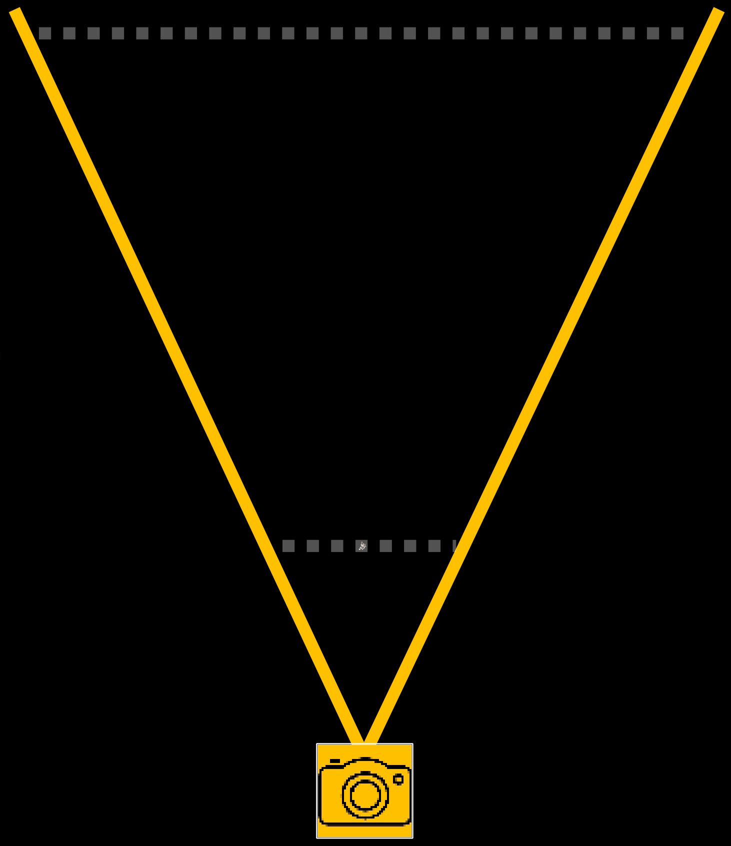

Novel view

Plane sweep volume

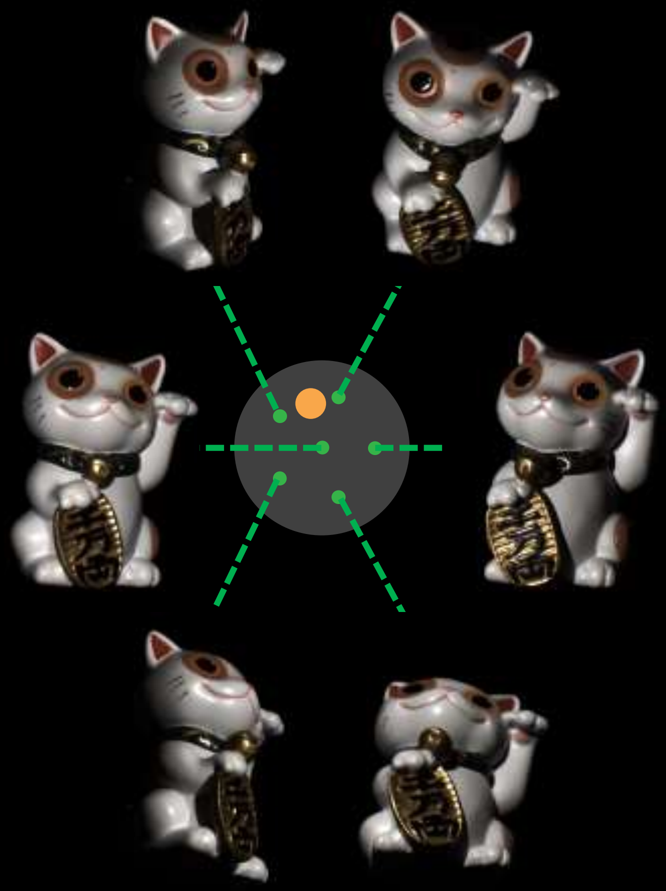

Input views

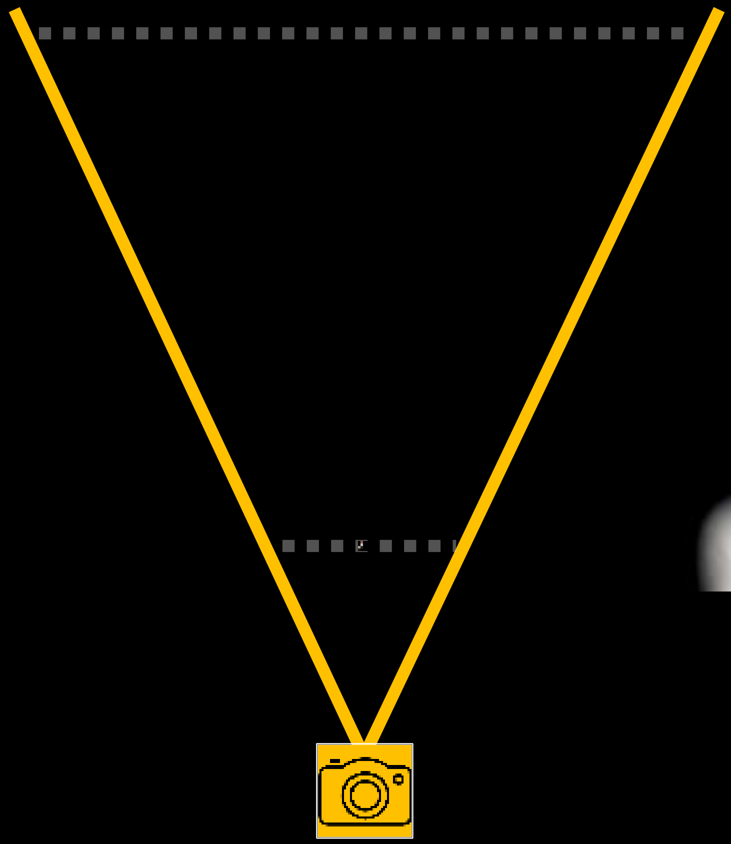

Novel view

Plane sweep volume

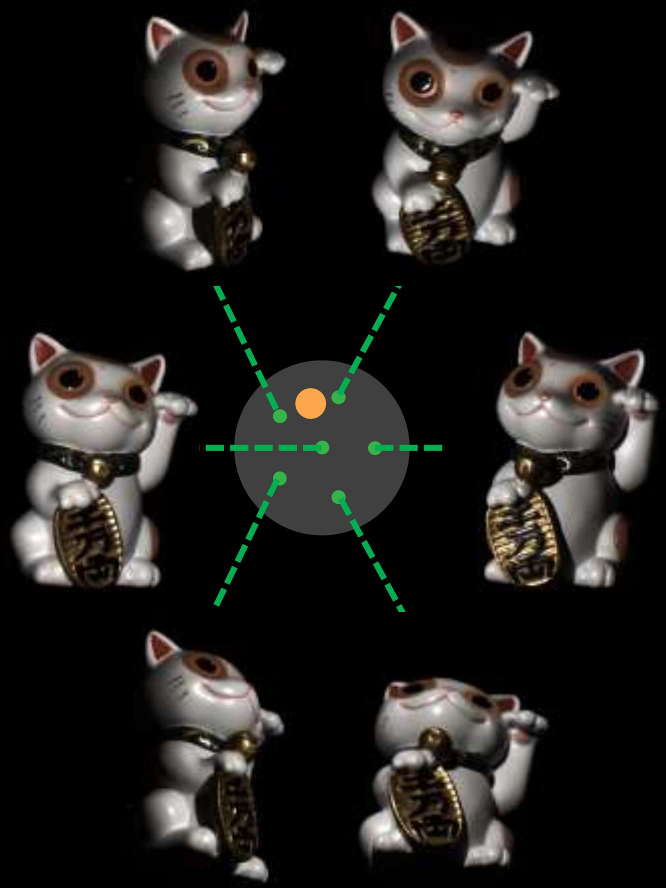

Input views

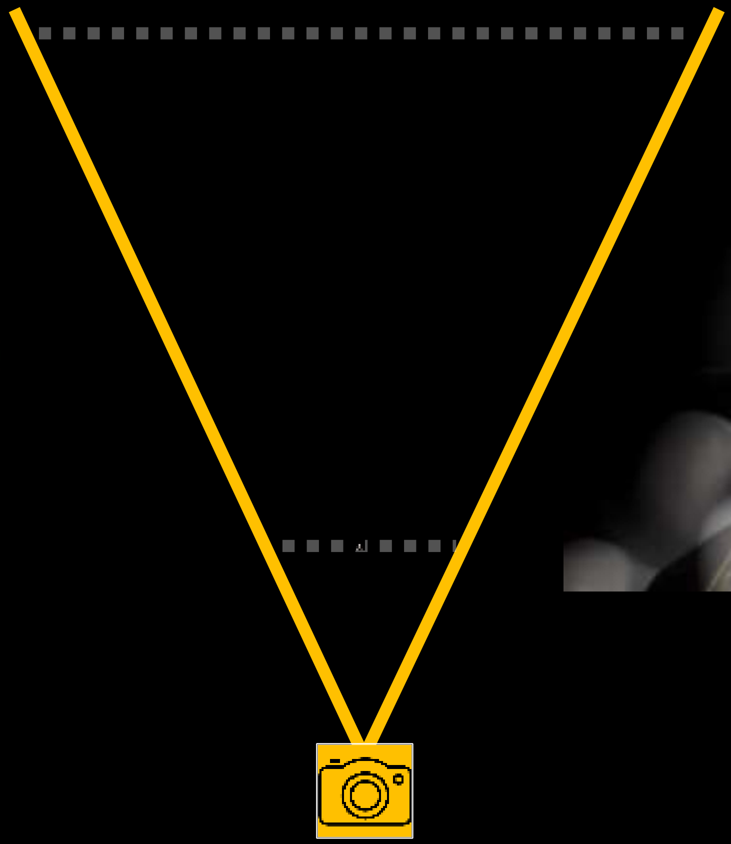

Novel view

Plane sweep volume

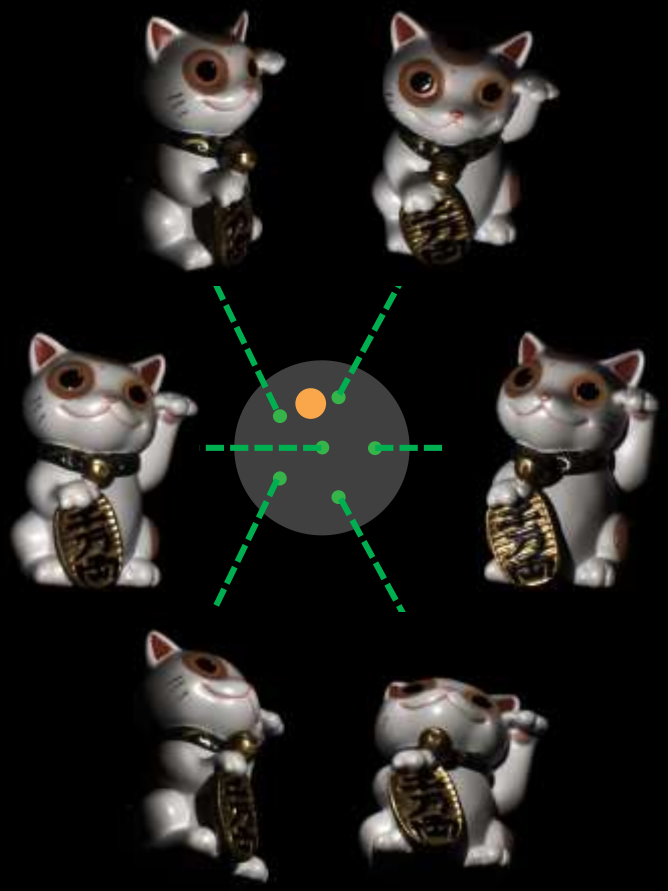

Input views

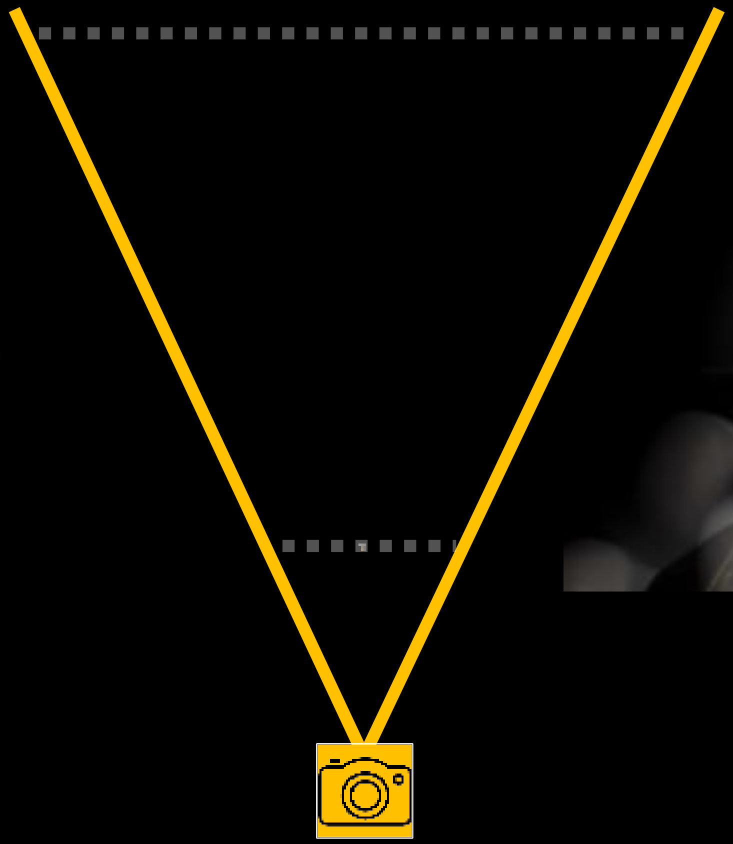

Novel view

Plane sweep volume

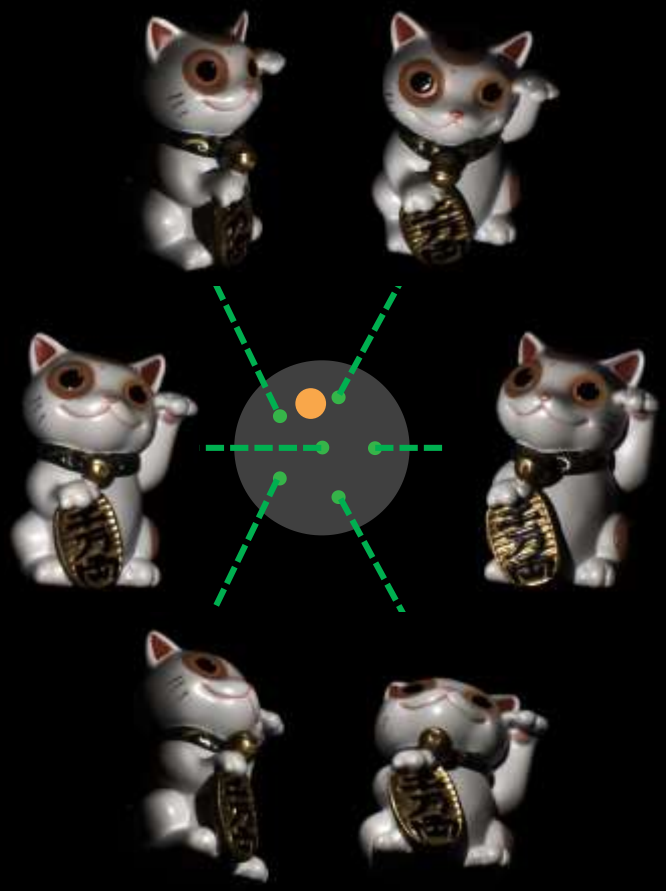

Input views

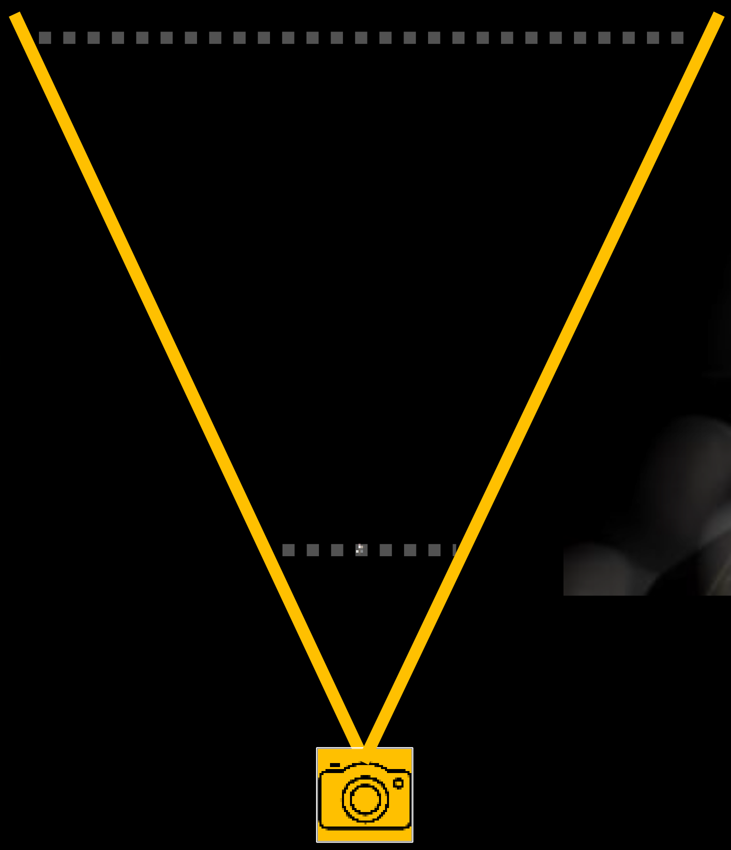

Novel view

Plane sweep volume

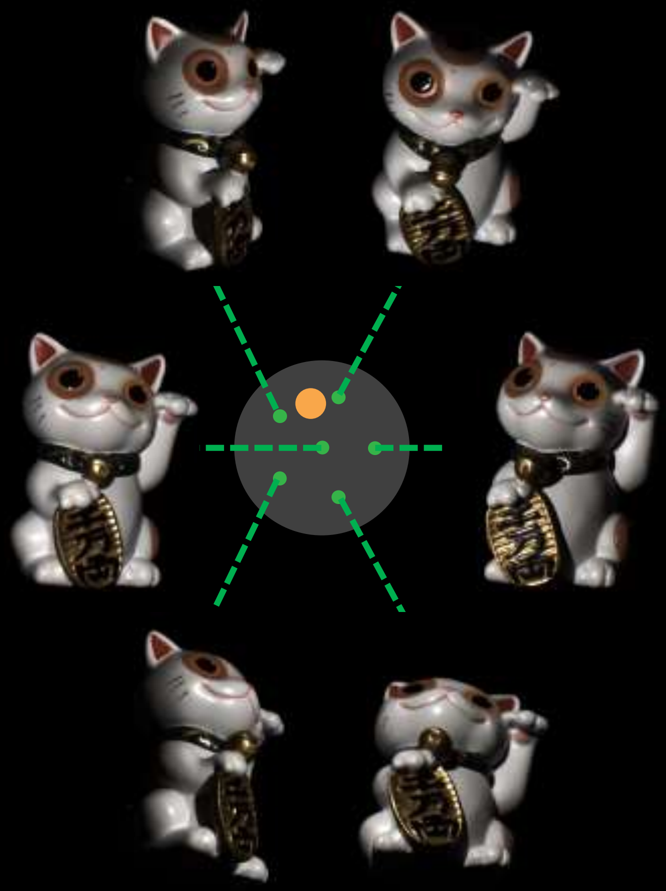

Input views

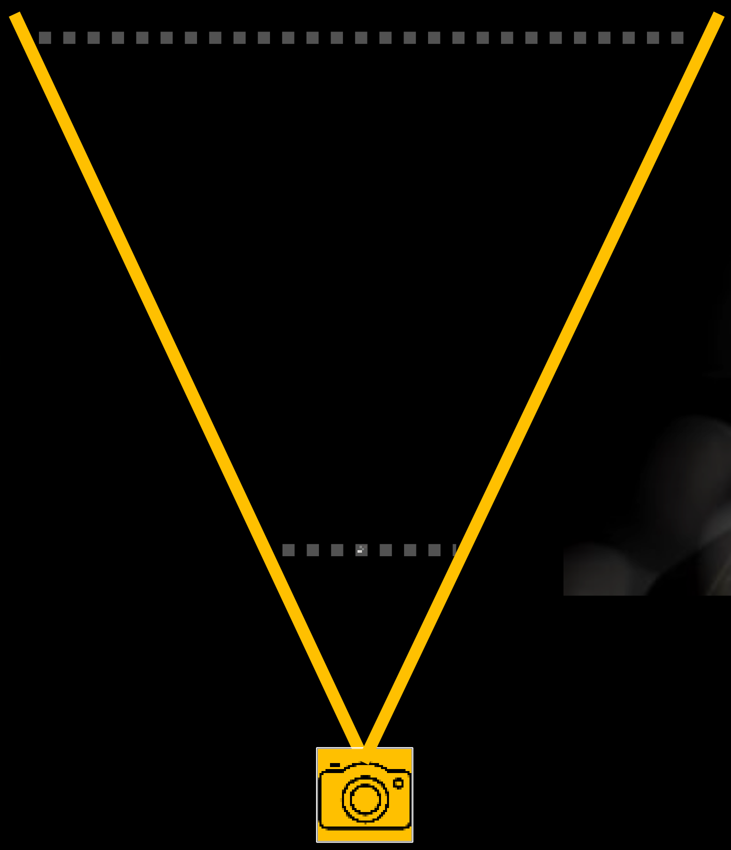

Novel view

Plane sweep volume

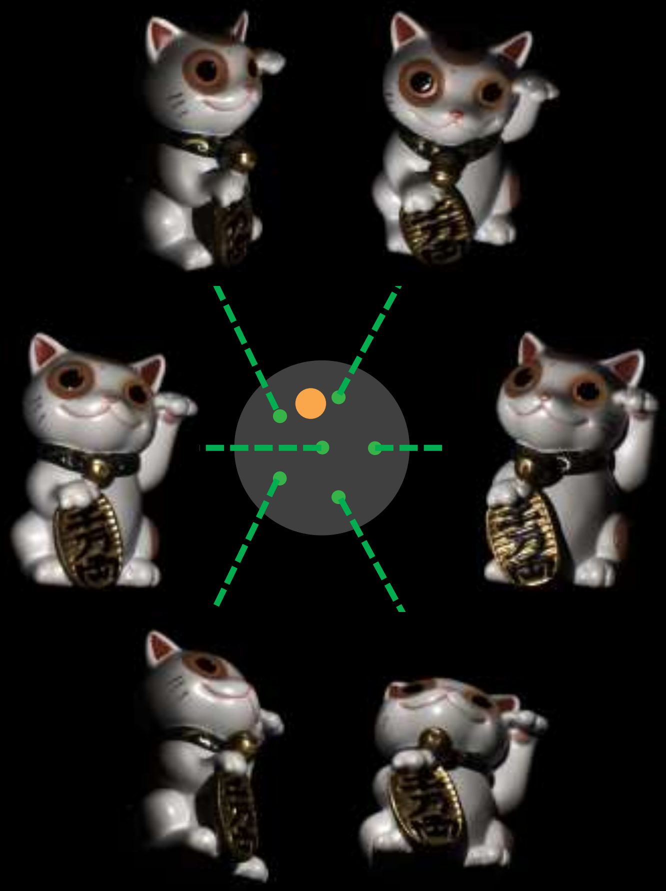

Input views

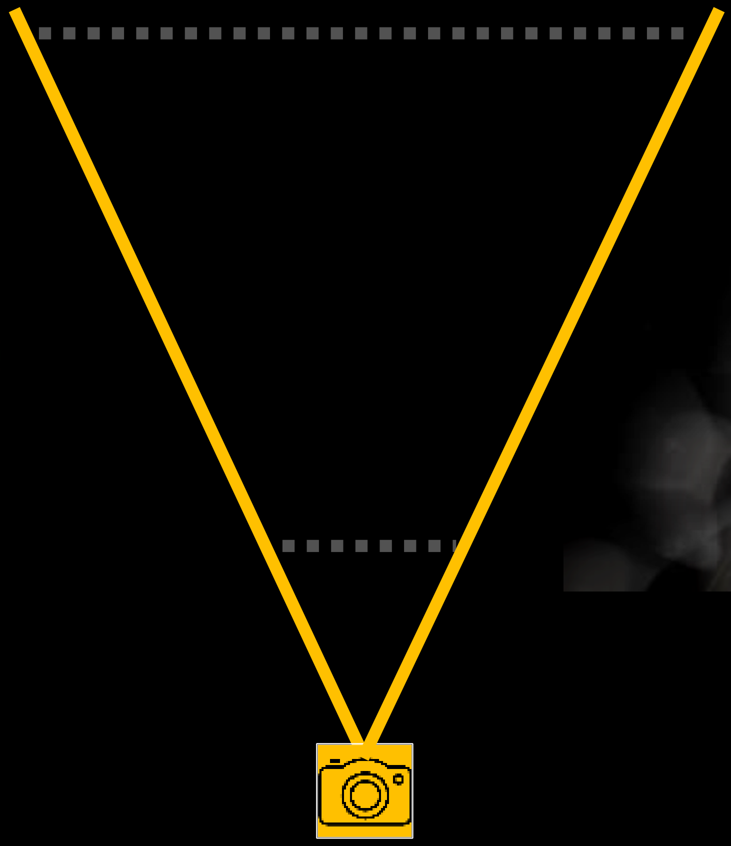

Novel view

Plane sweep volume

Plane sweep volume

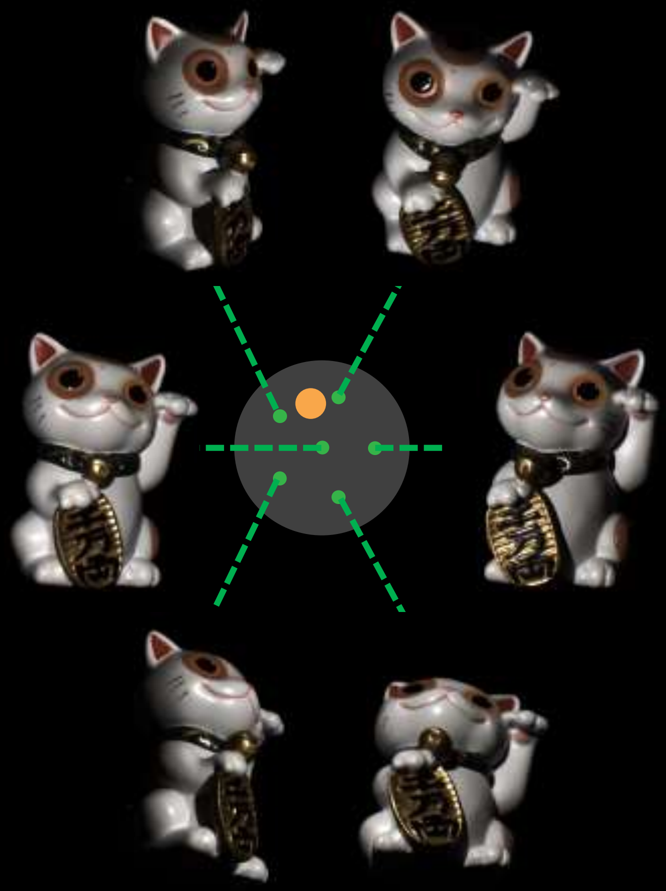

Input views

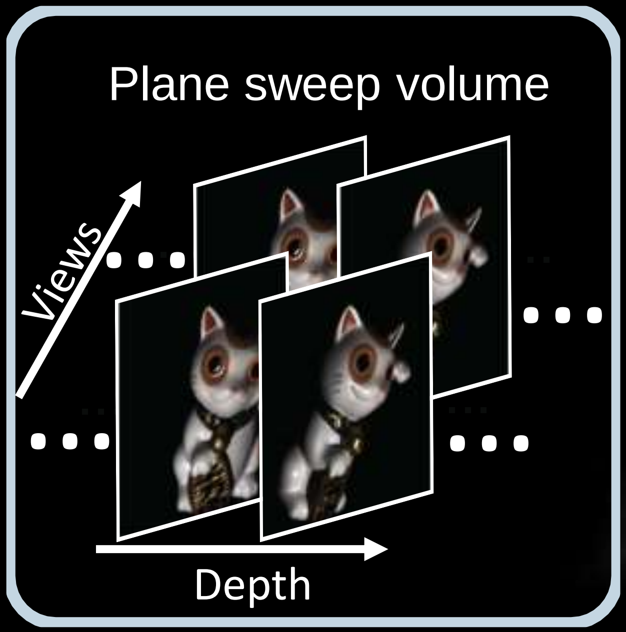

Novel view

Attention maps

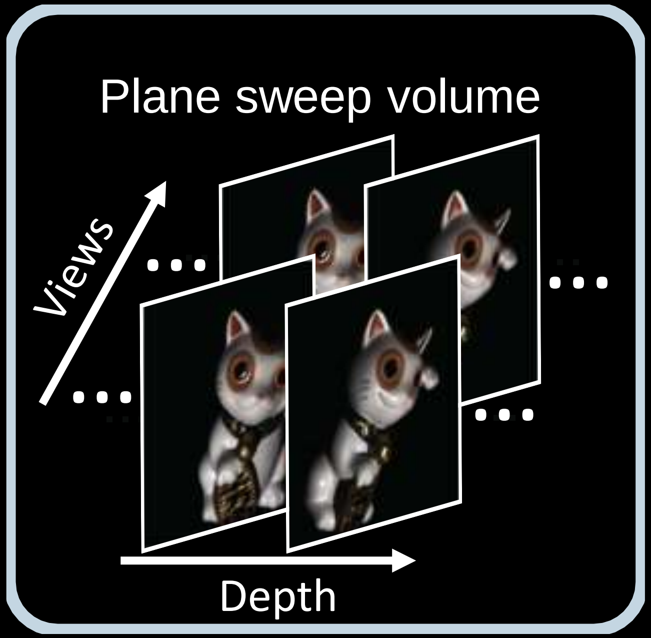

Attention maps

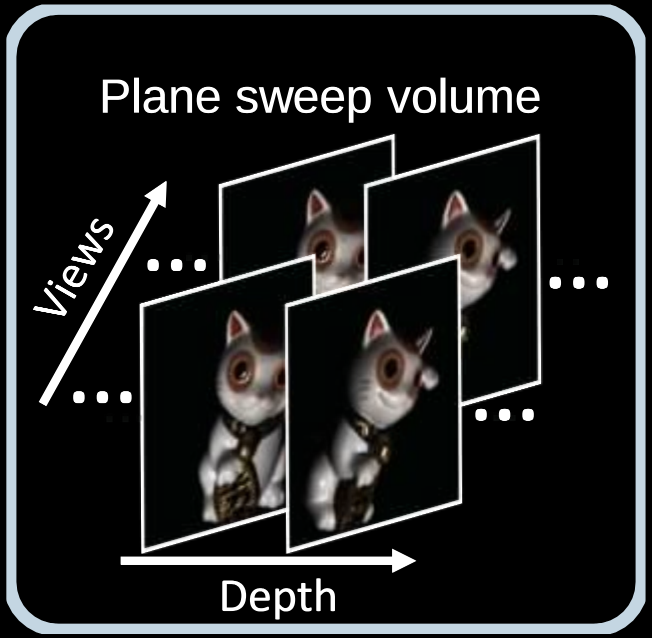

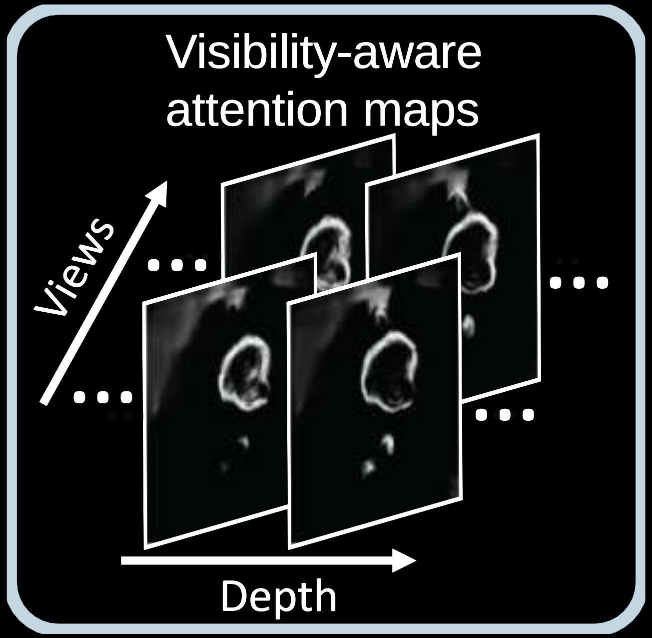

Attention maps

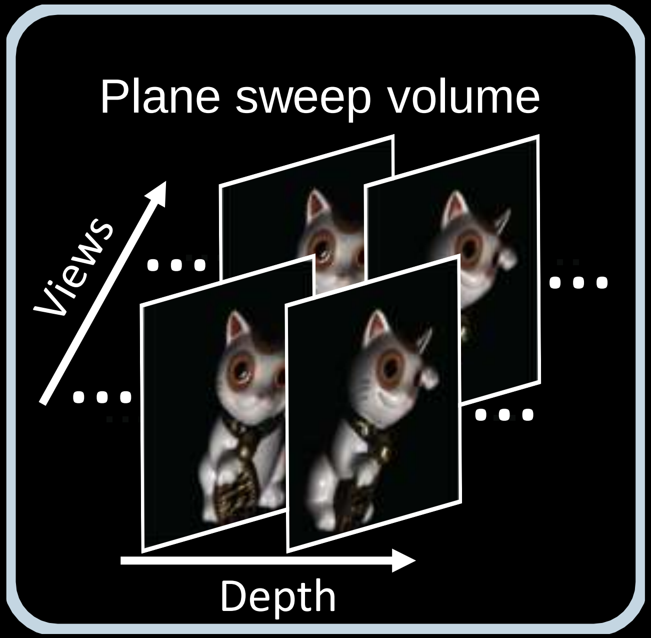

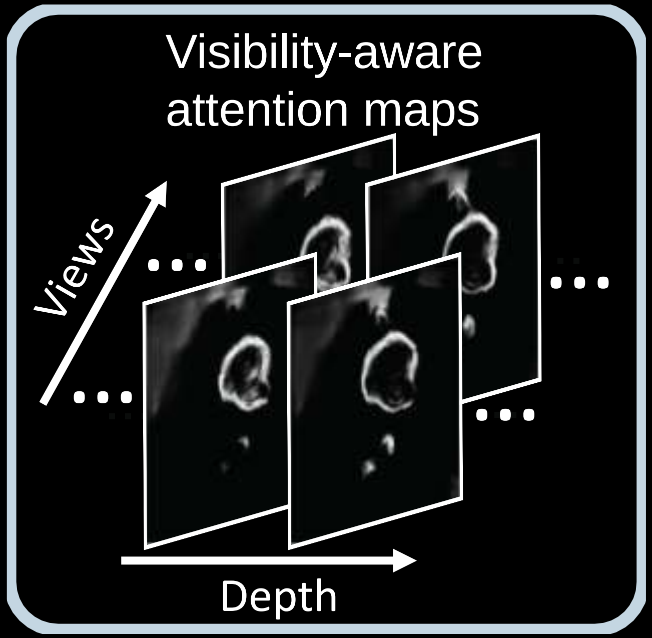

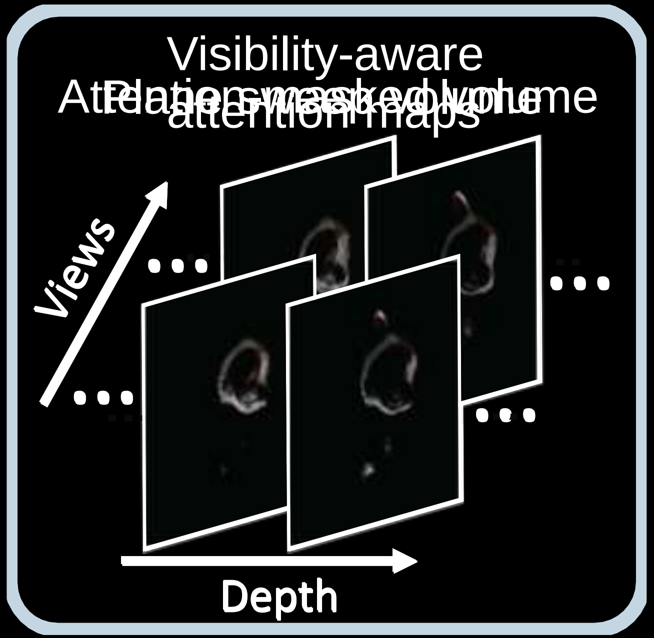

Attention maps

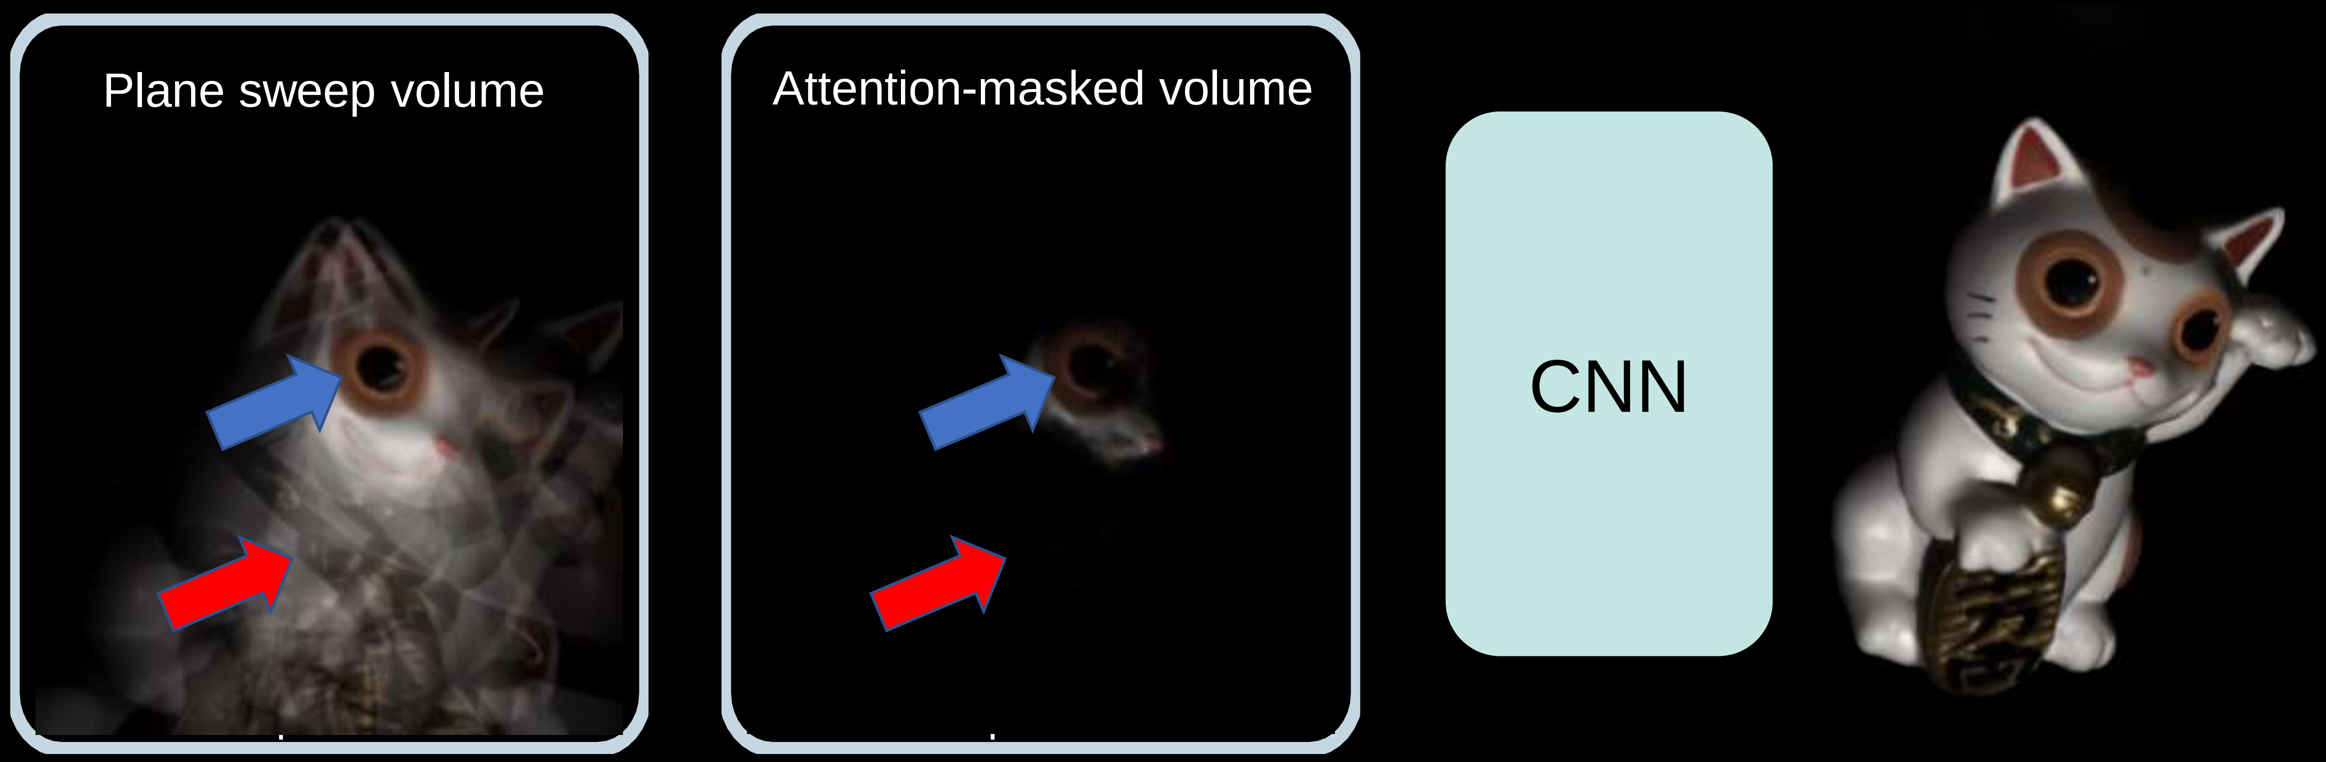

Our network

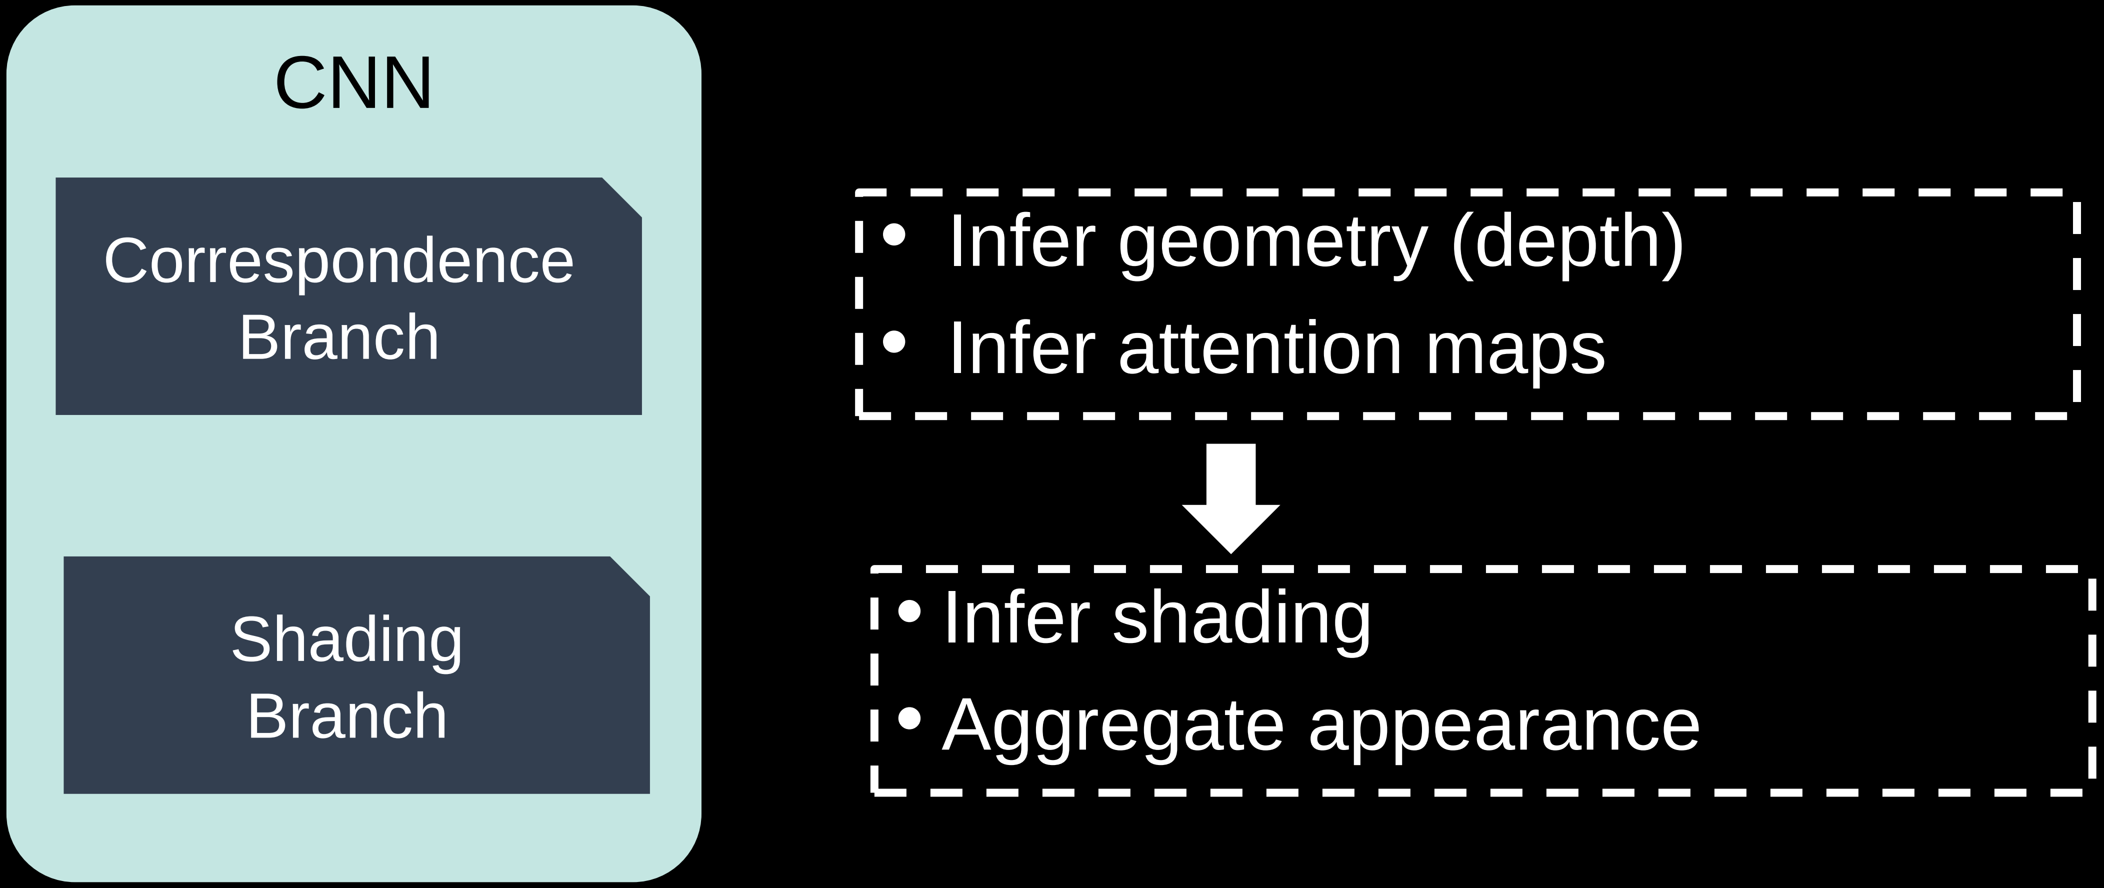

CNN
Correspondence Branch
Shading Branch
• Infer geometry (depth)
• Infer attention maps
• Infer shading
• Aggregate appearance

Correspondence branch

- Infer geometry (depth)
- Infer attention maps

Correspondence branch

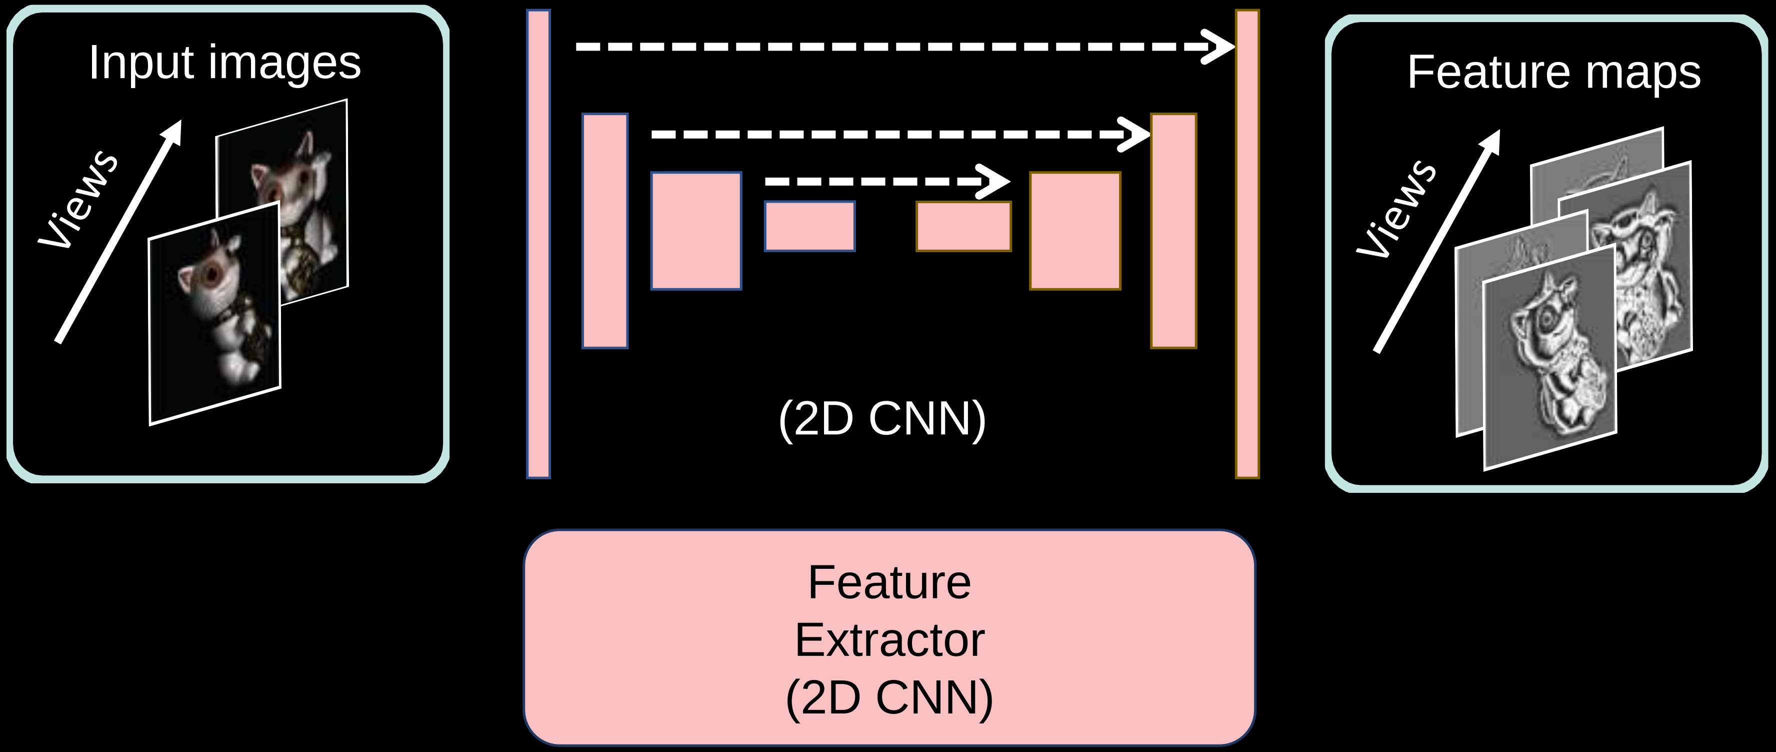

Correspondence branch

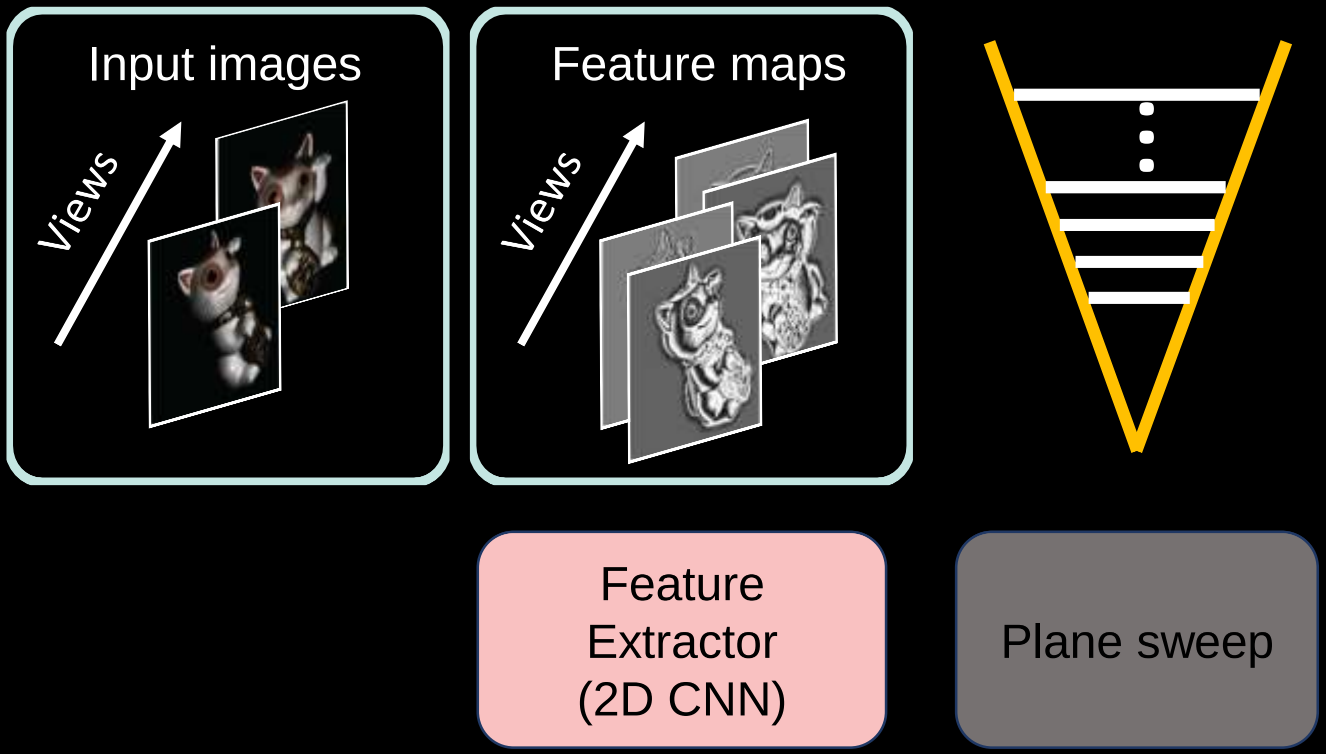

Correspondence branch

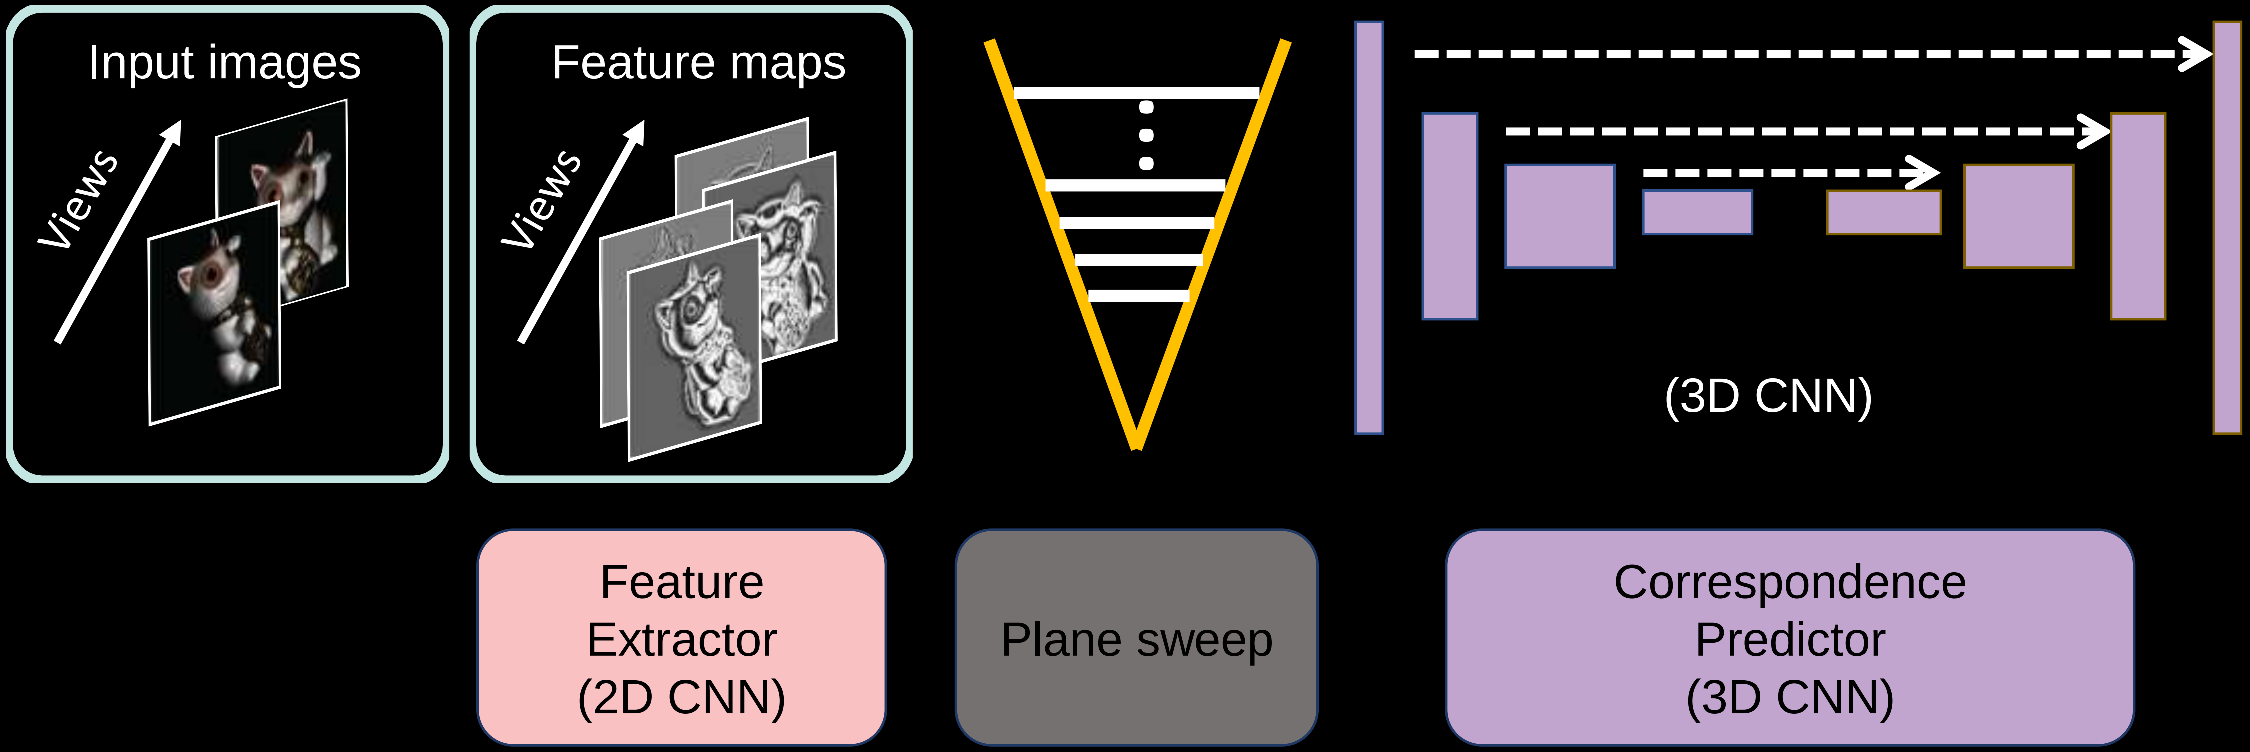

Correspondence branch

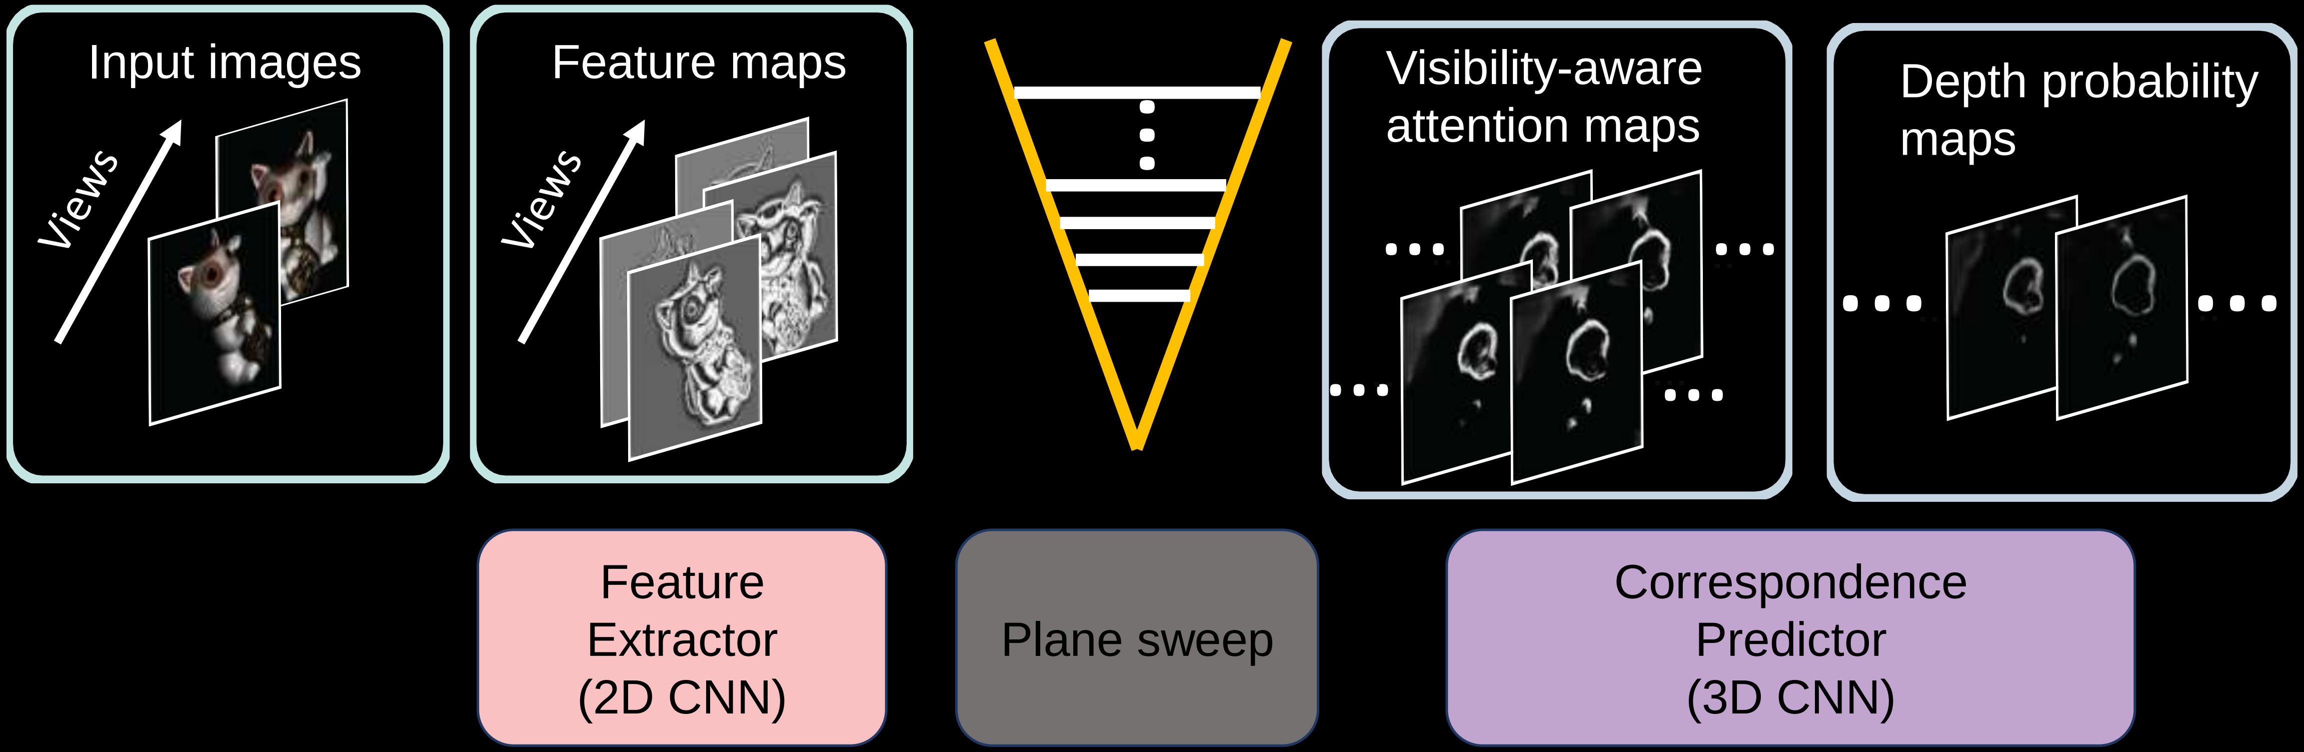

Correspondence branch

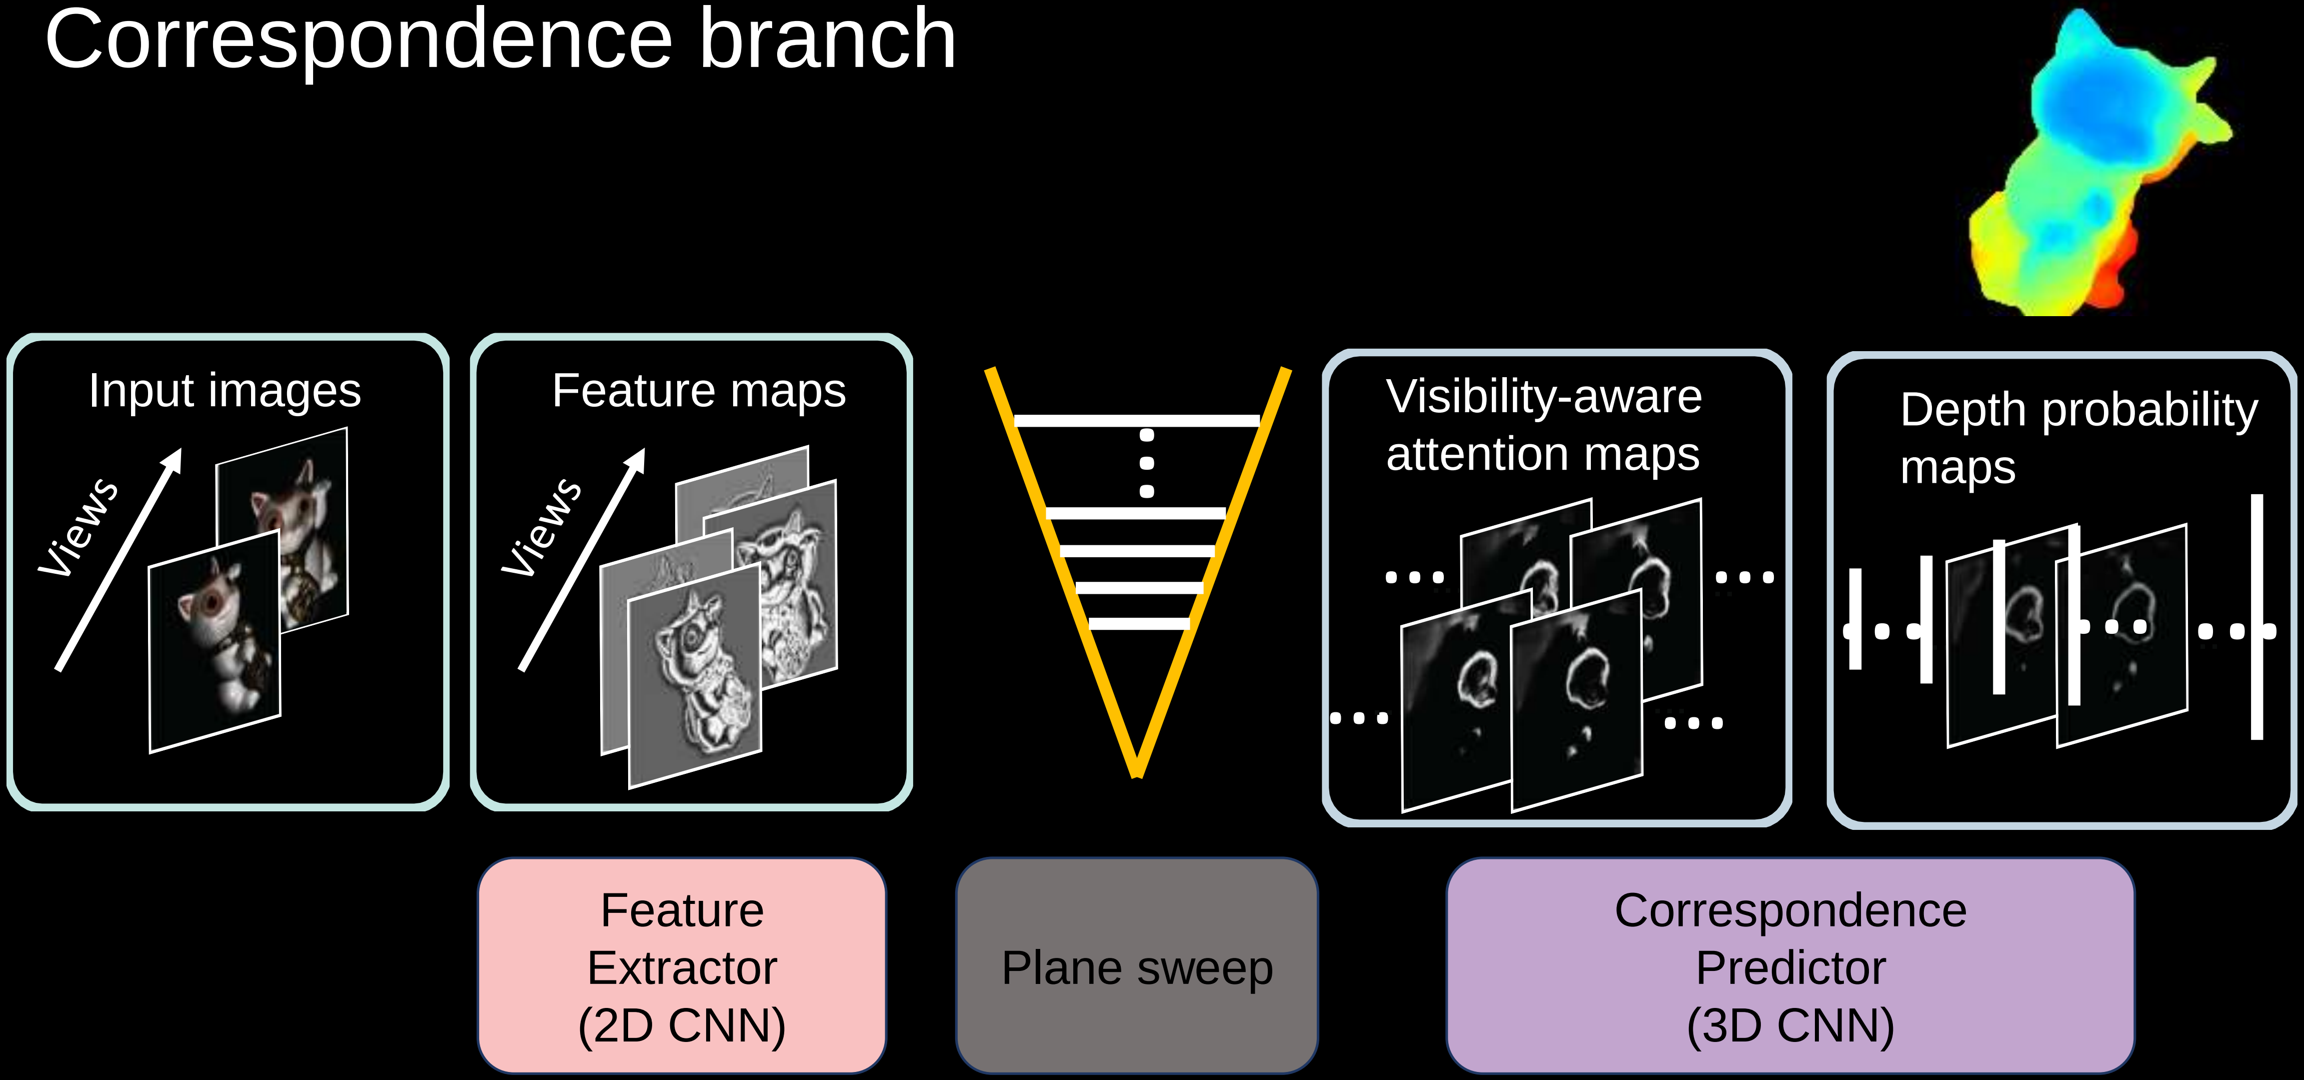

Shading branch

Visibility-aware attention maps

Depth probability maps

Input images

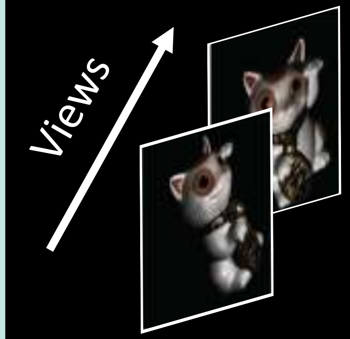

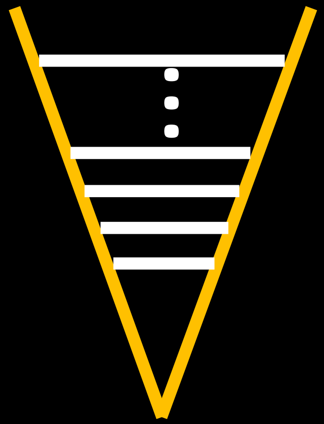

Shading branch

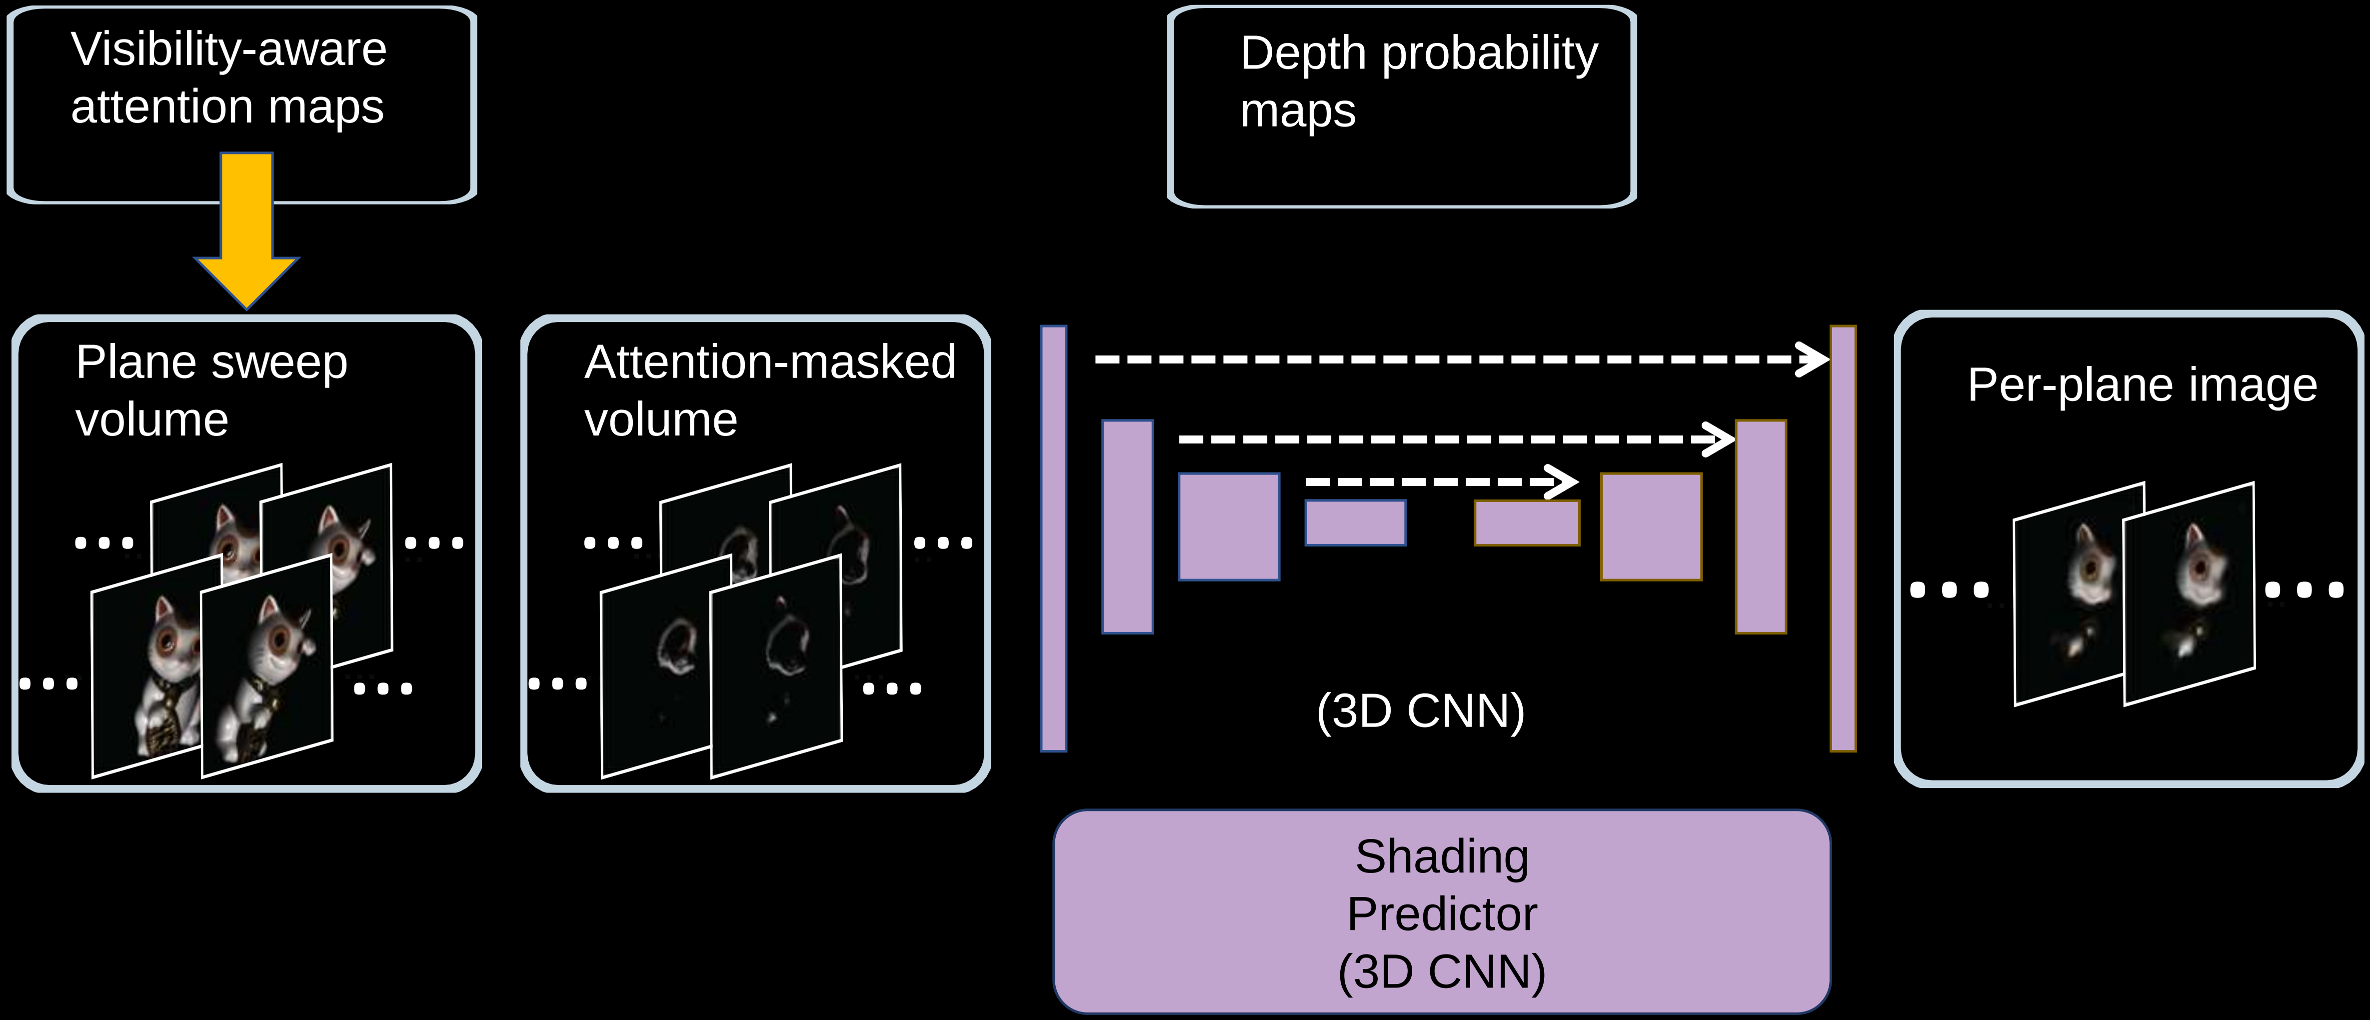

Shading branch

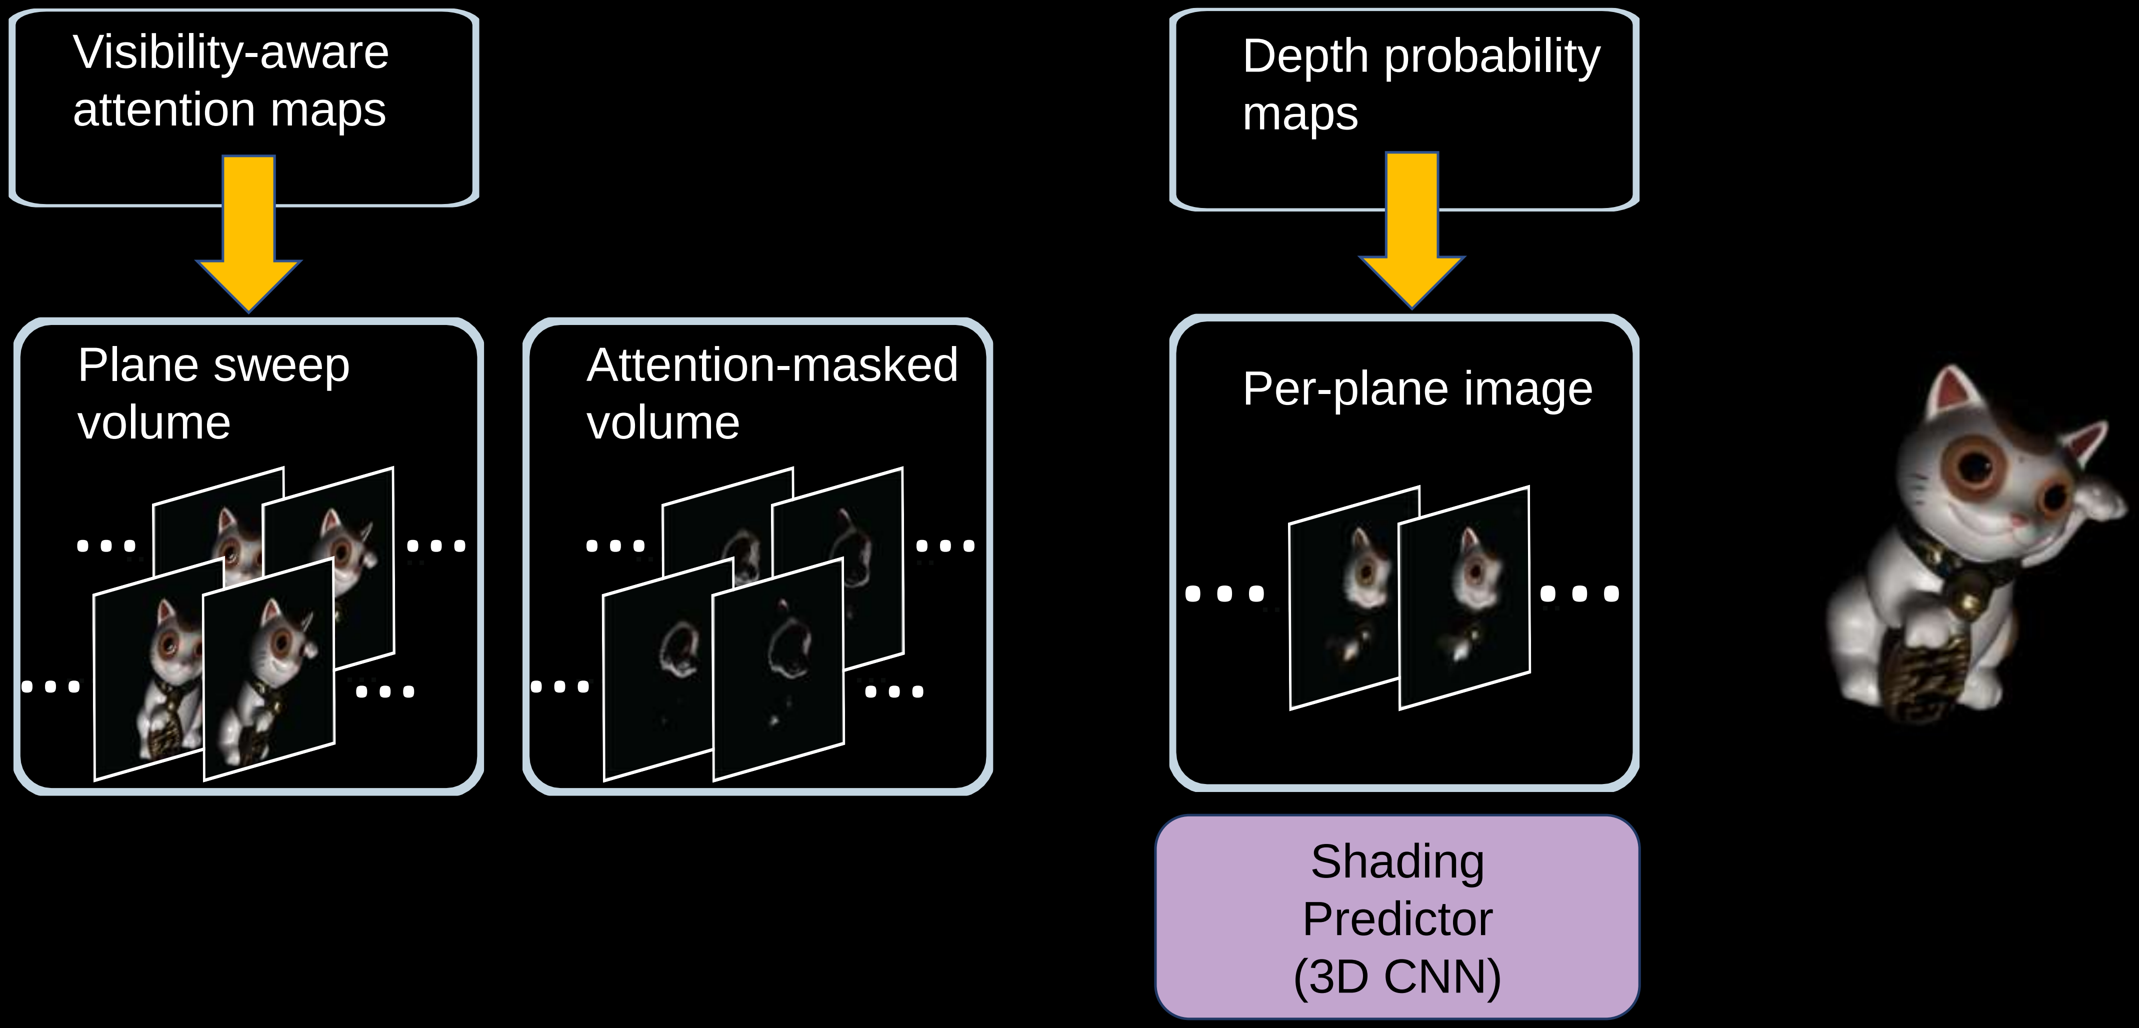

Corr-Branch + Shade-Branch

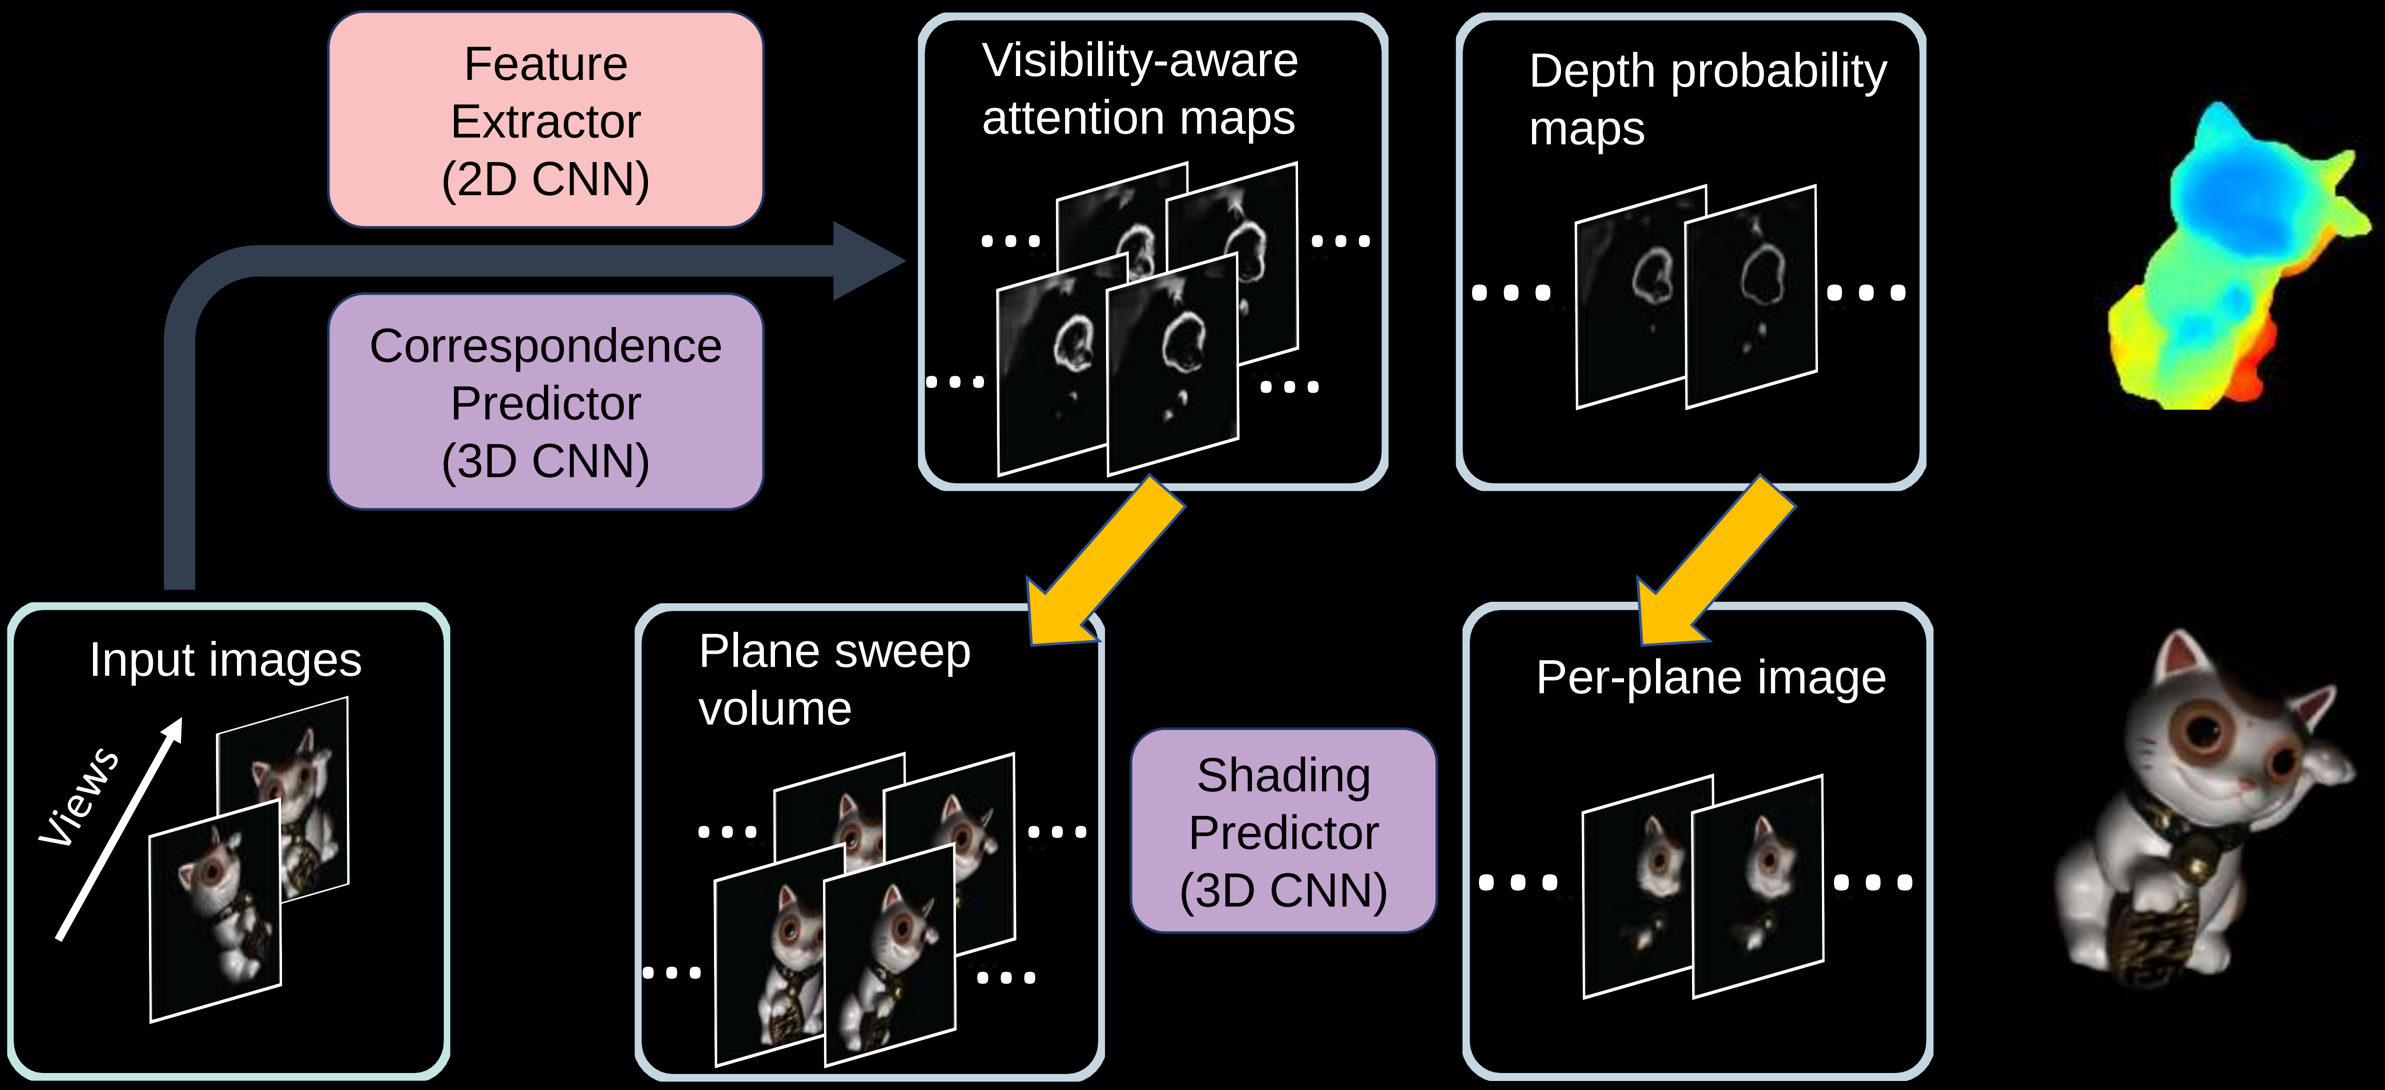

Real Data Results

Data #1: input images and corresponding views

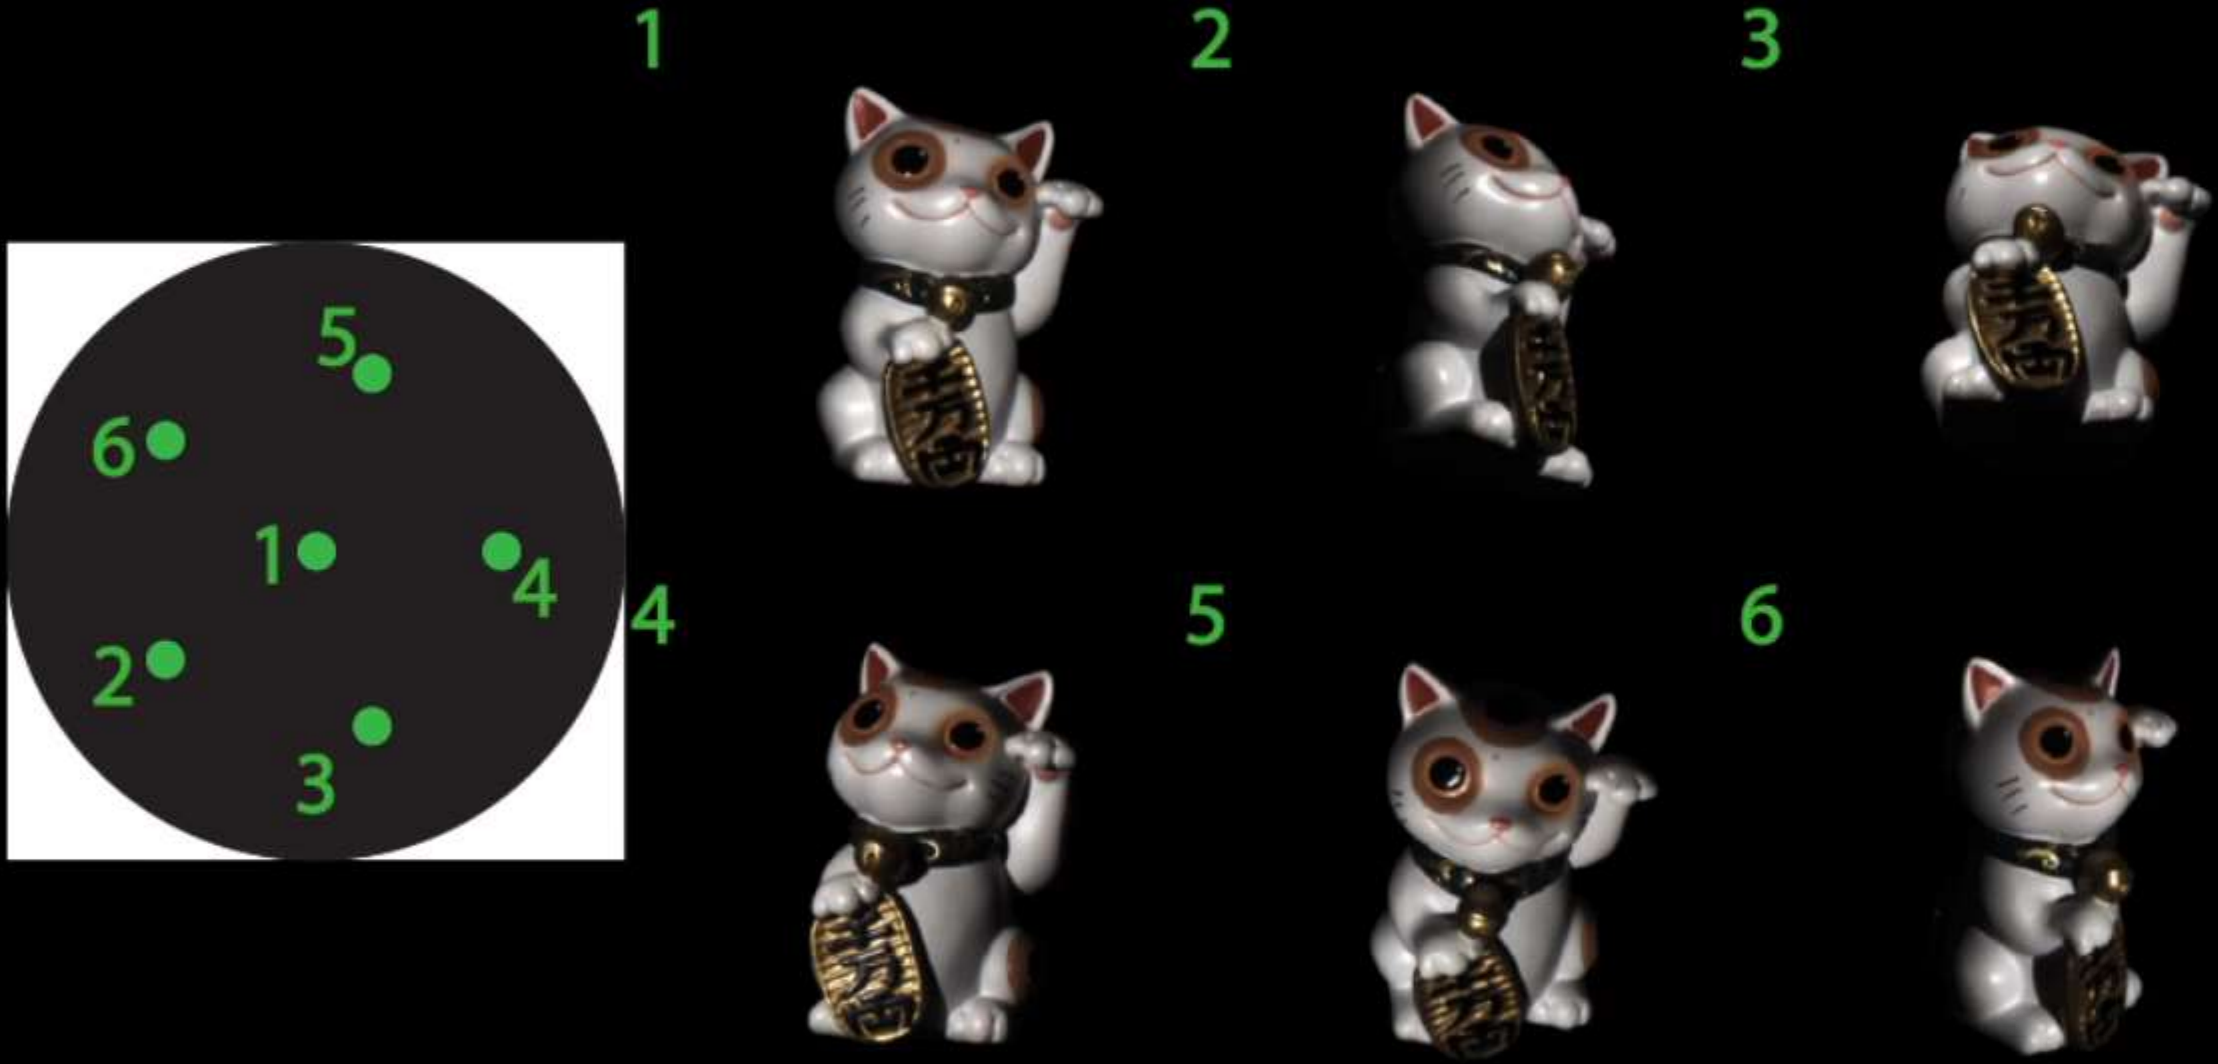

Data #1: compare with [Penner and Zhang 2017]
using the same inputs

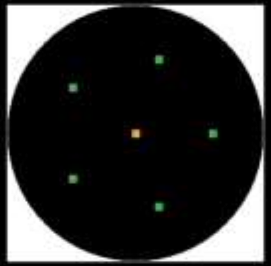

Penner and
Zhang 2017

Our results

Ground truth

(Some views are occluded)

Data #1: compare with [Sun et al. 2018] using the same inputs

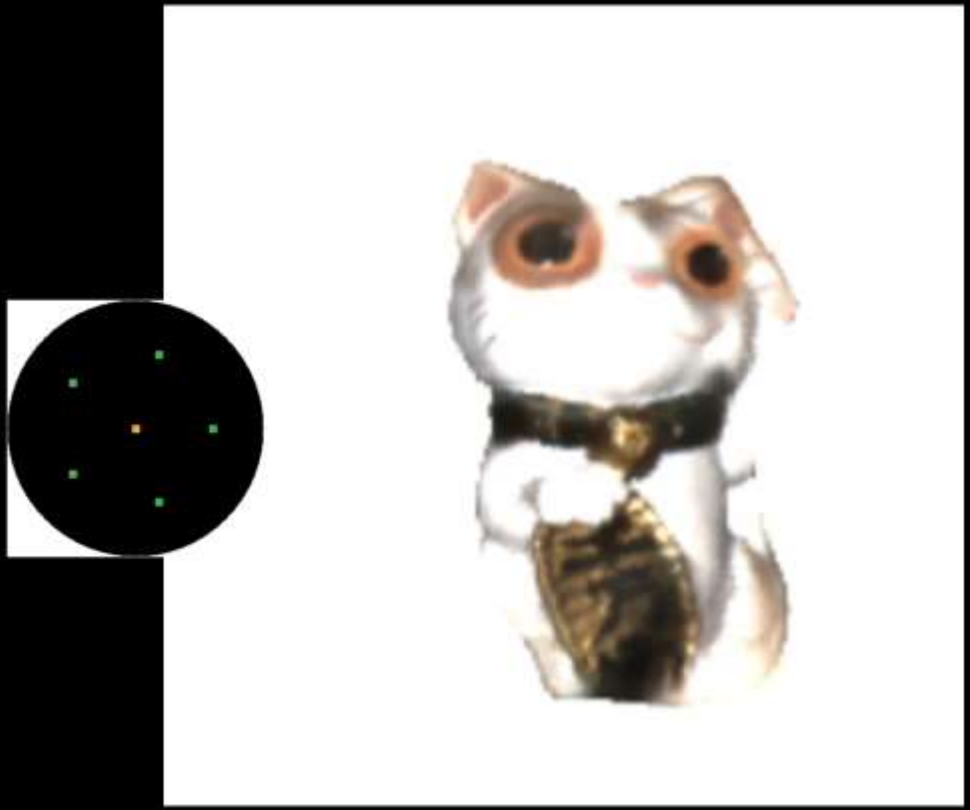

Sun et al. 2018

Our results

Ground truth

(Some views are occluded)

Data #2: input images and corresponding views

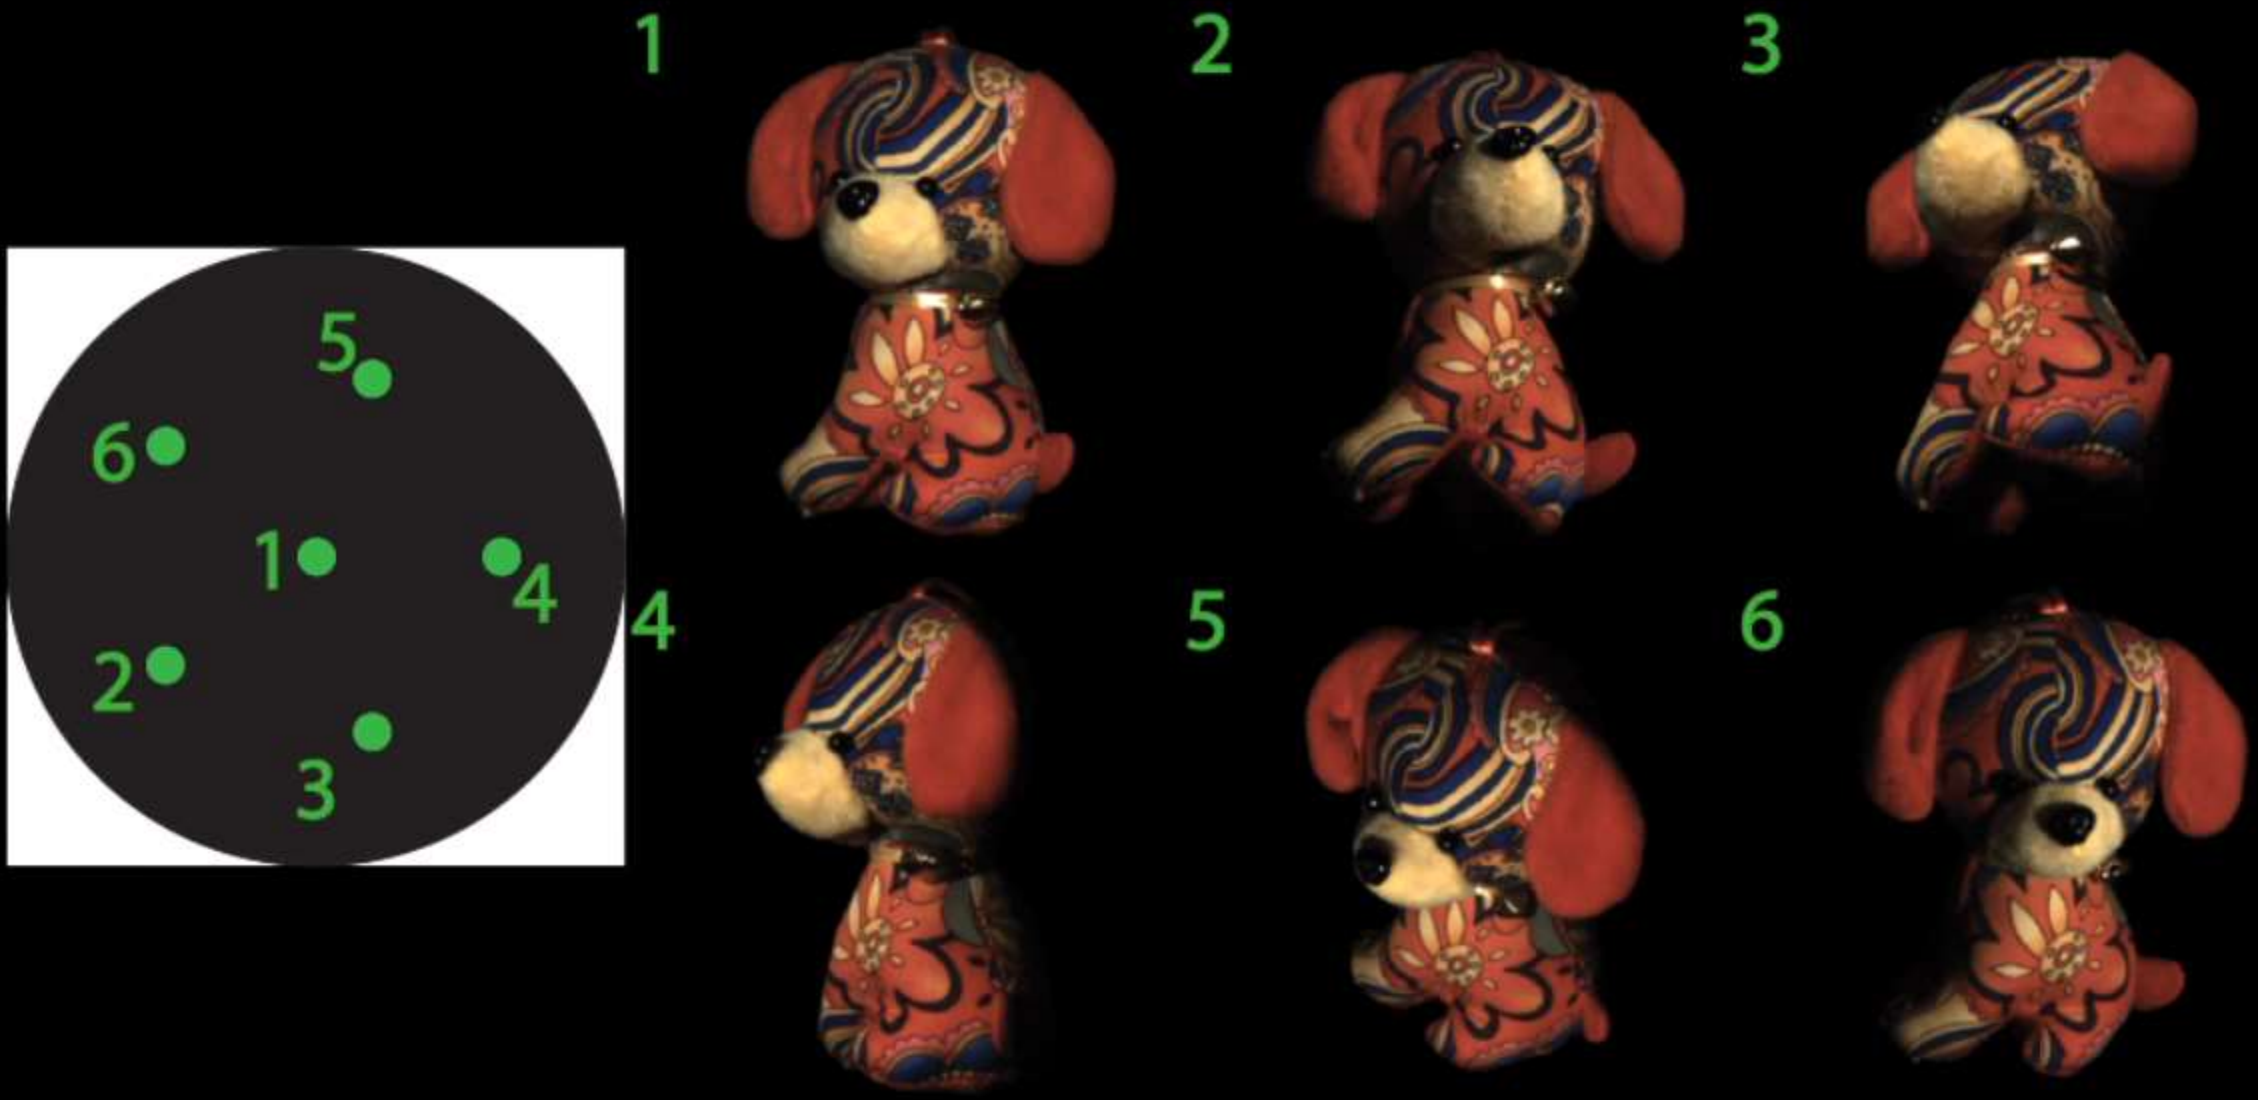

Data #2: our results and ground truth

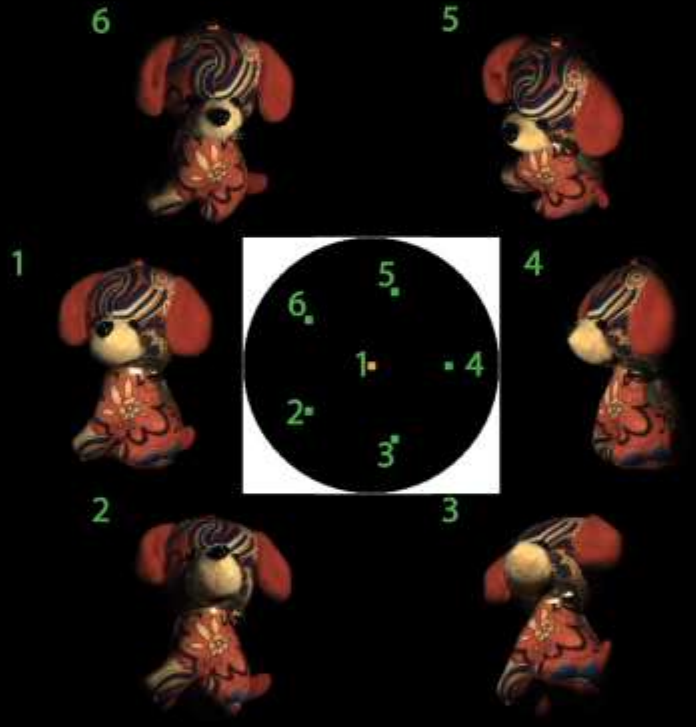

Inputs and viewing directions

Our results

Ground truth

(Some views are occluded)

Data #3: input images and corresponding views

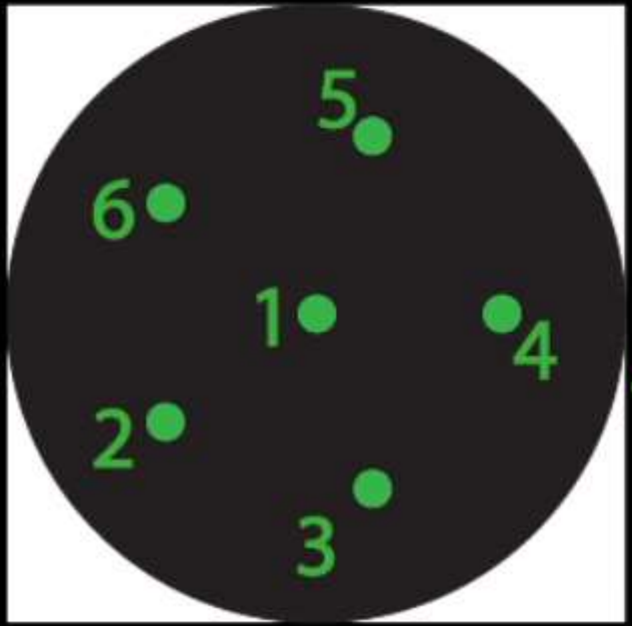

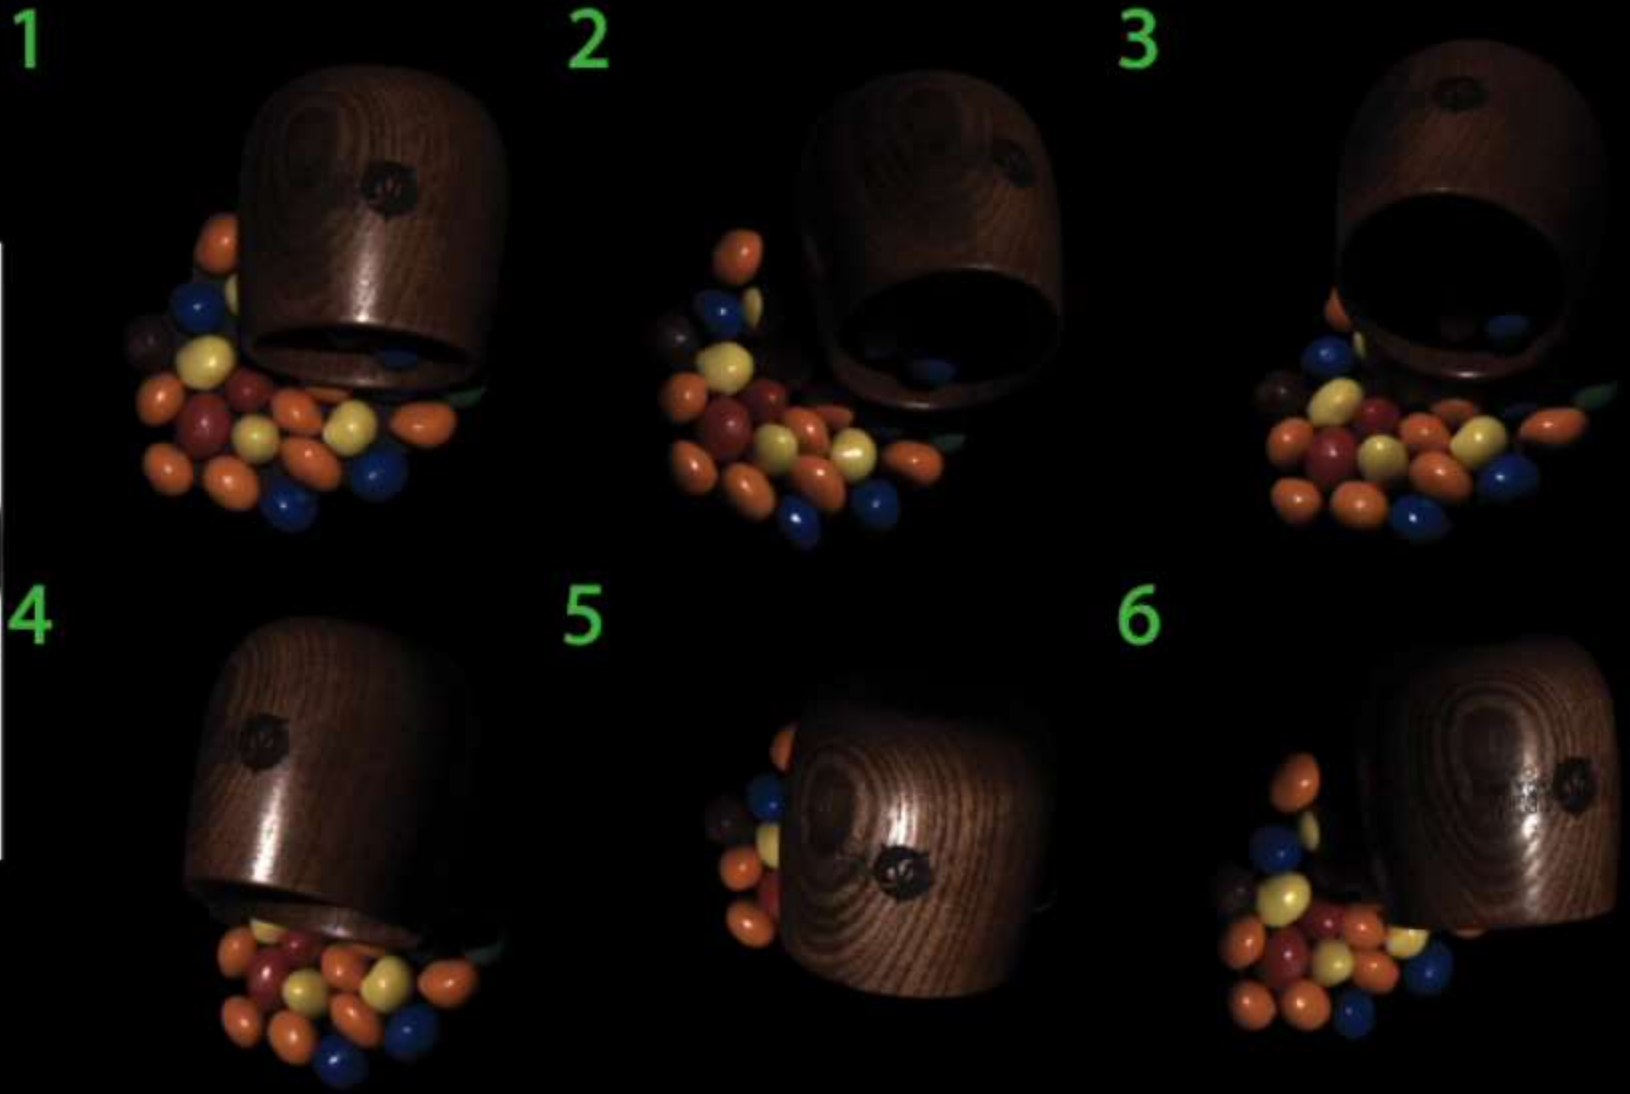

Data #3: our results and ground truth

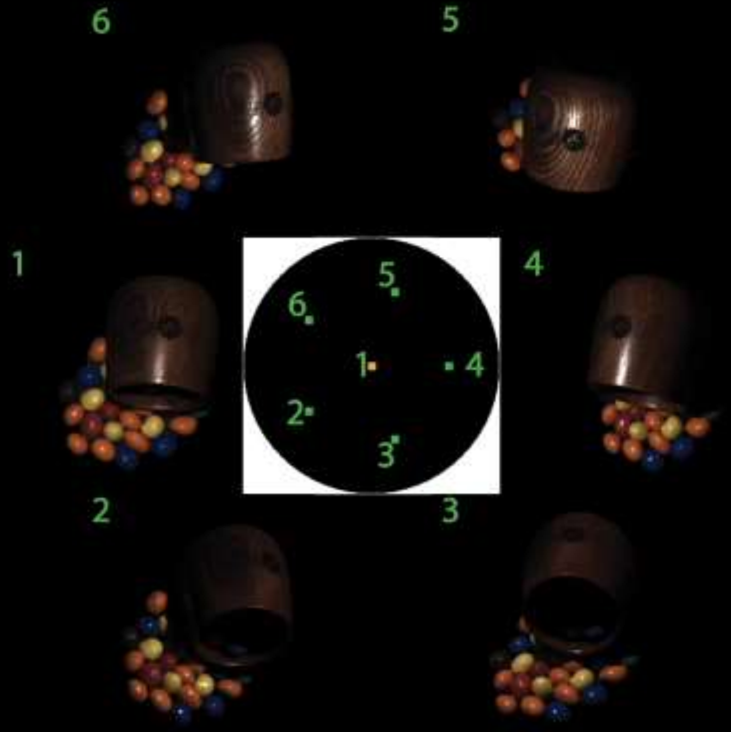

Inputs and viewing directions

Our results

Ground truth

(Some views are occluded)

Data #4: input images and corresponding views

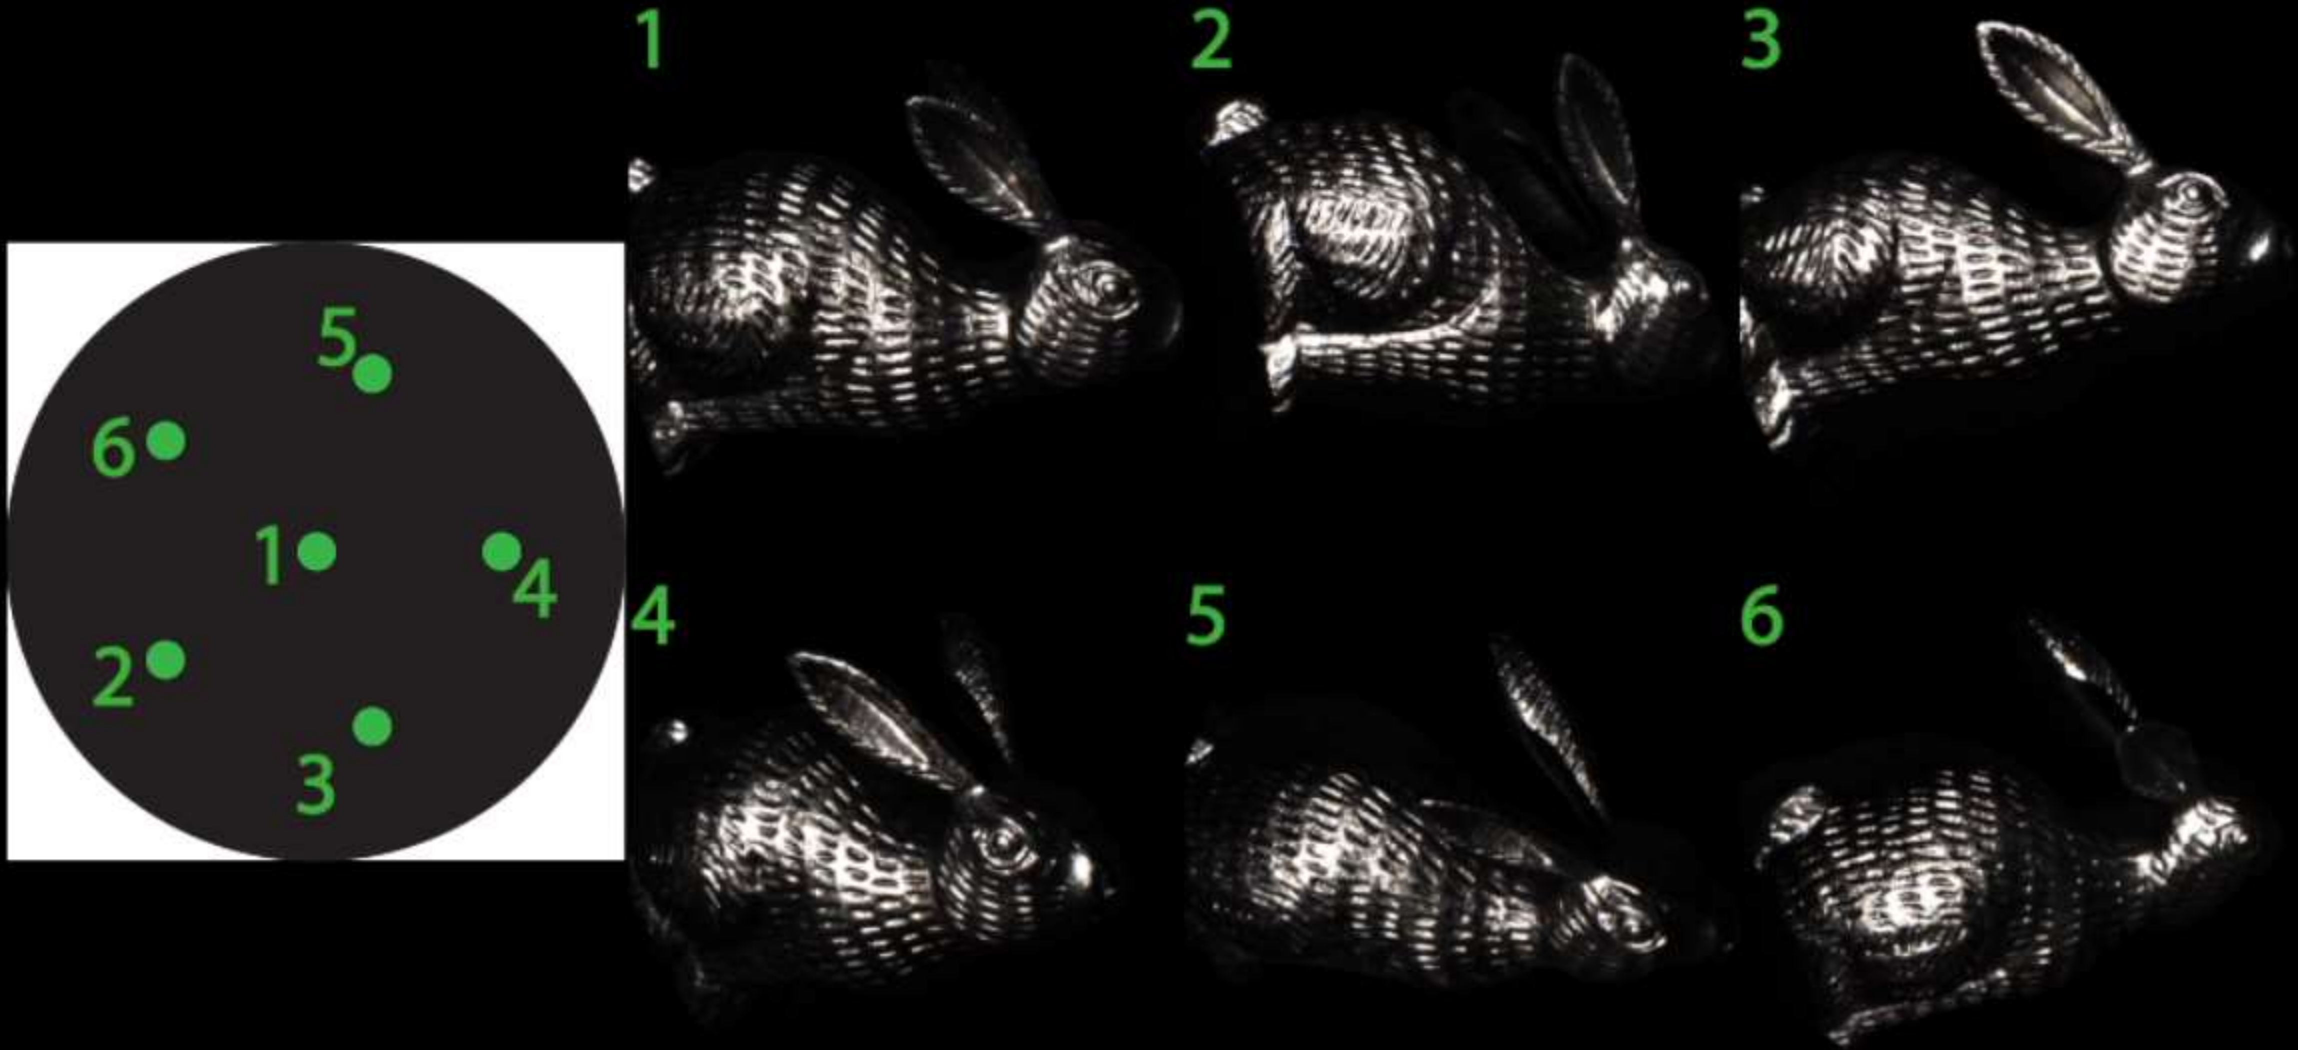

Data #4: our results and ground truth

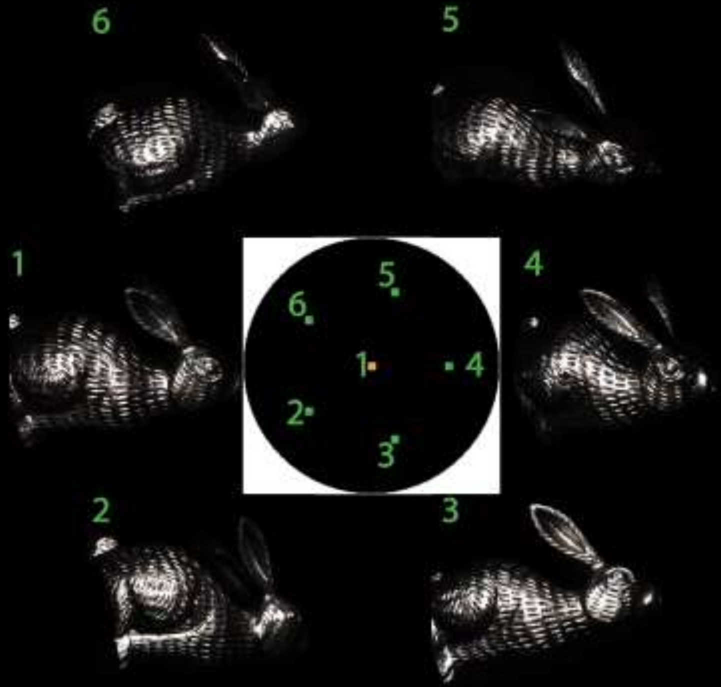

Inputs and viewing directions

Our results

Ground truth

(Some views are occluded)

Novel view relighting

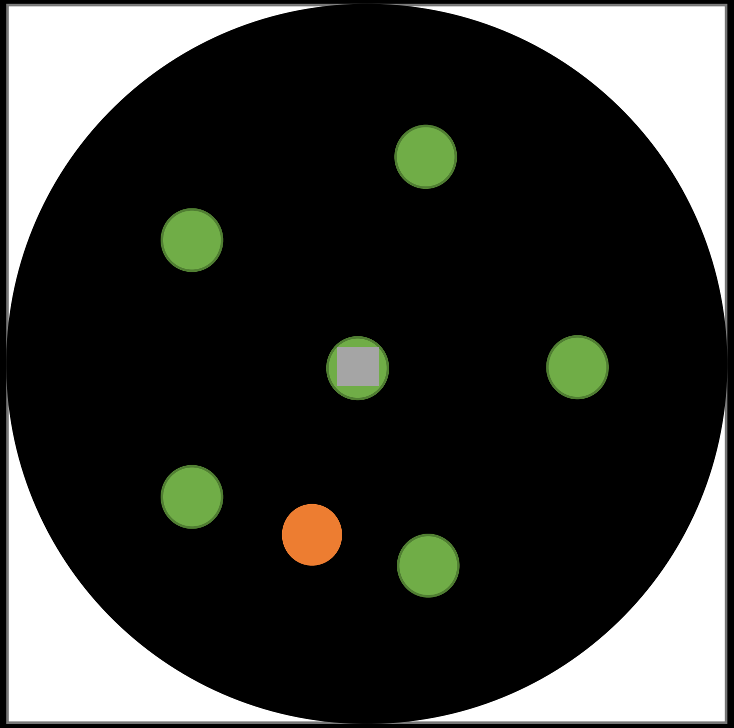

Novel view relighting

Data #4: our novel view relighting results

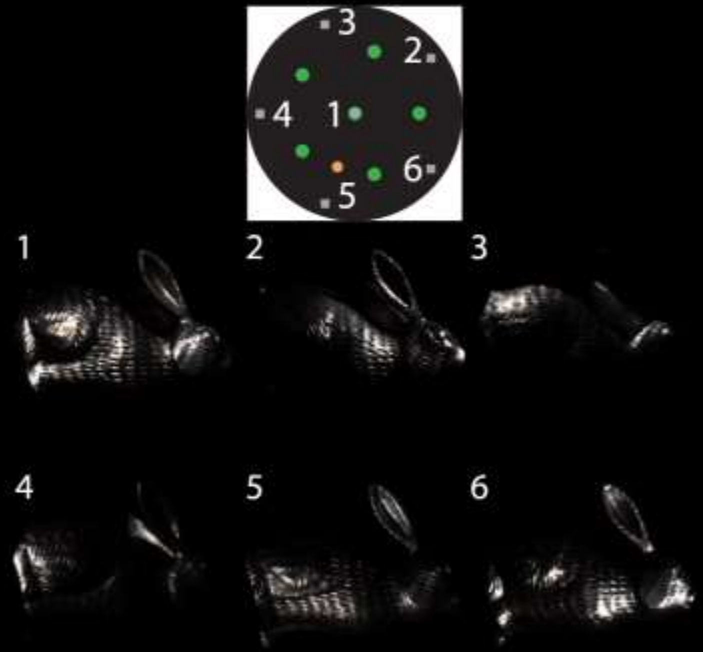

Our synthesized images

Relighting

Environment map 1

Multi-view stereo

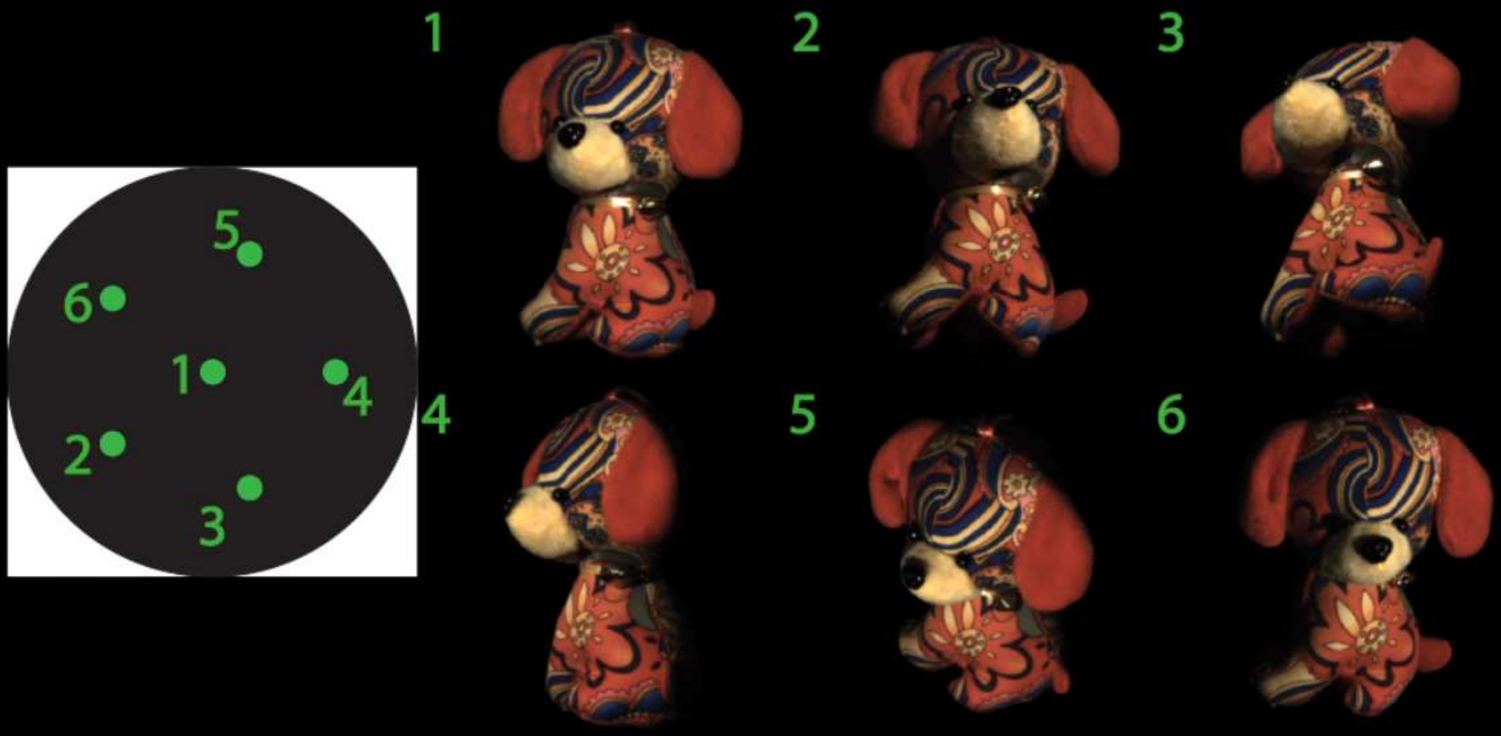

Input images

Reconstruction

Data #2: Multi-view stereo from synthesized images

Reconstruction from 56 captured images

Reconstruction from 56 synthesized images using our method

Reconstruction from 56 synthesized images using our method (rendered with color)

Limitations

- Highly specular objects
  - 64 x 64 image crops for training
  - Limited receptive field

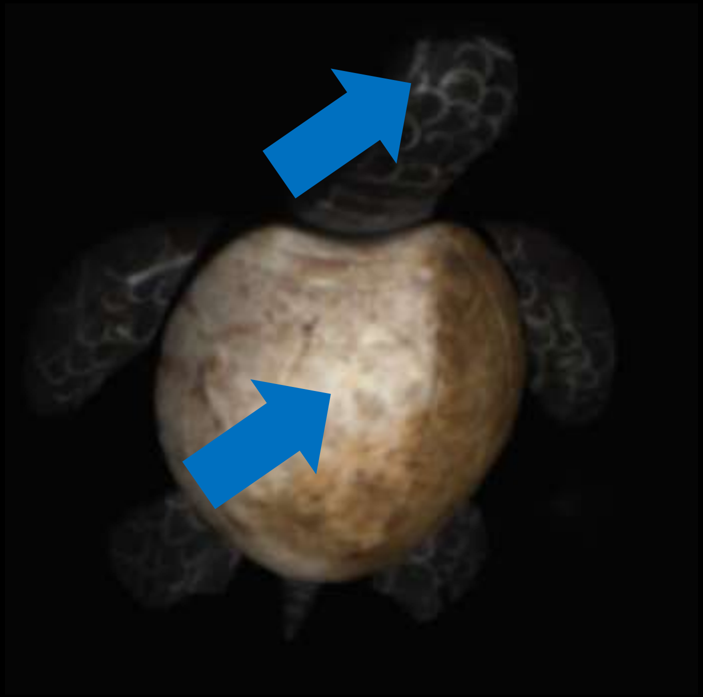

Our result

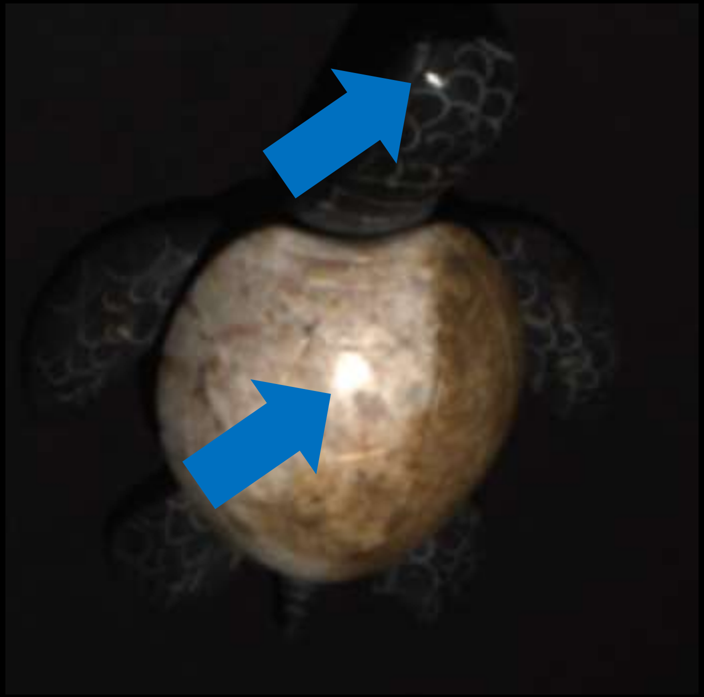

Ground truth

Limitations

- Highly specular objects
  - 64 x 64 image crops for training
  - Limited receptive field
- Highly non-convex shape
  - Visible from 1 or 2 views

Our result

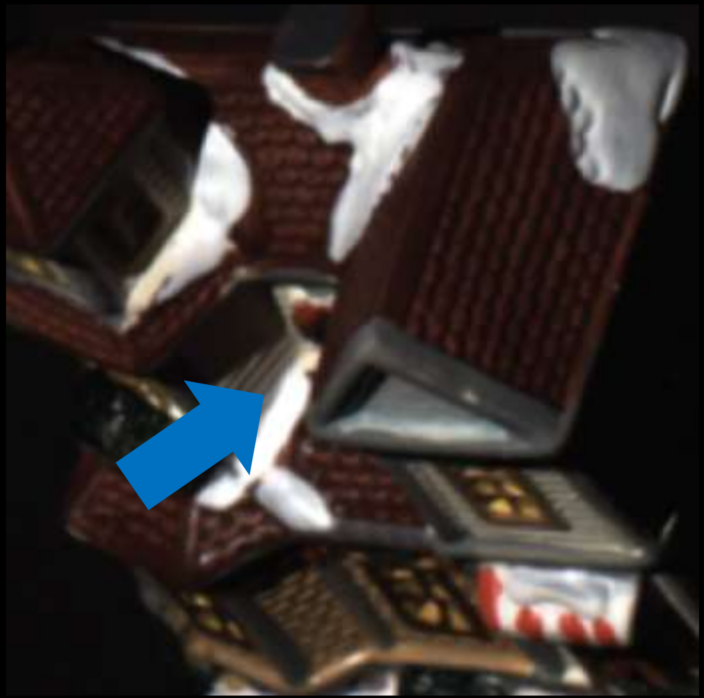

Ground truth

Conclusion

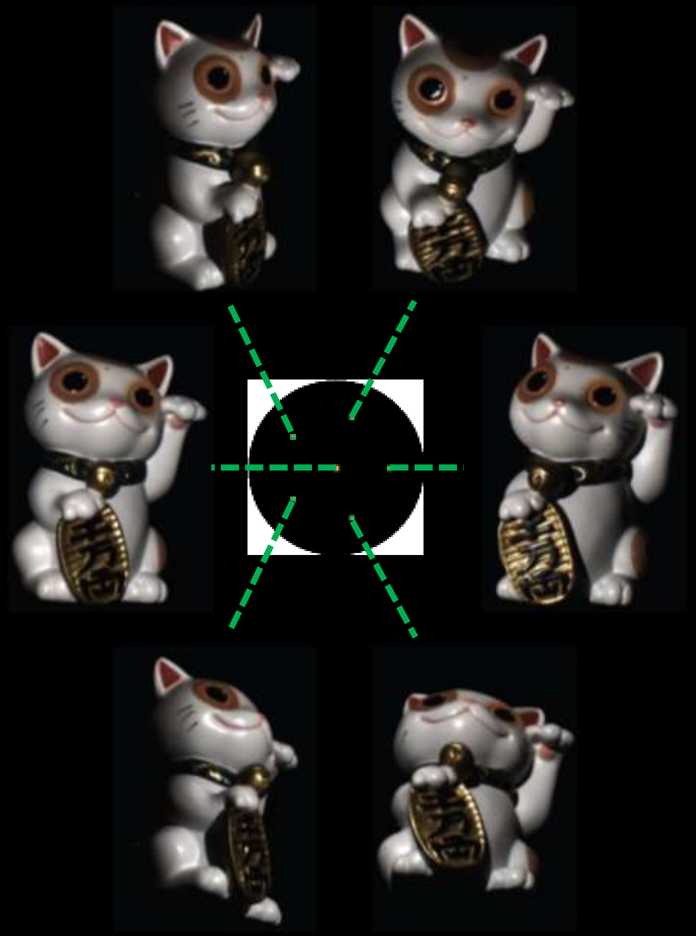

Our Result

Ground Truth

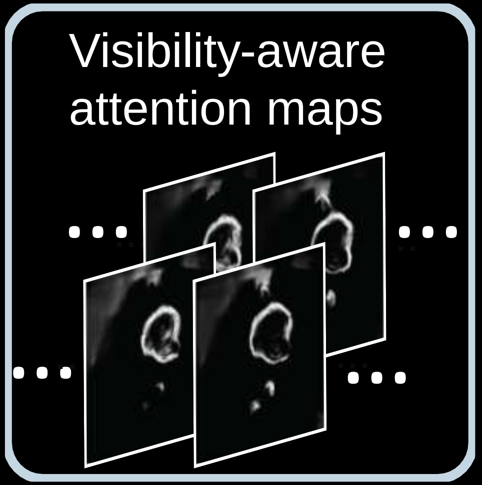

Conclusion

Our Result

Ground Truth

Novel view relighting

Multi-view stereo

Acknowledgements

- Pratul Srinivasan and Zhengqin Li
- NSF grants 1617234, 1703957
- ONR grant N000141712687
- Adobe
- Adobe Research Fellowship
- Powell-Bundle Fellowship
- Ronald L. Graham Chair
- UC San Diego Center for Visual Computing

THANK YOU!