



1st StruCo3D Workshop

Structural and Compositional Learning on 3D Data

Questions/Feedbacks: [kaichun \[at\] cs.stanford.edu](mailto:kaichun@cs.stanford.edu)

geometry.stanford.edu/struco3d




Organizers



Kaichun Mo
Stanford



Paul Guerrero
Adobe



Kai Wang
Brown



Antoni Rosinol
MIT



Fei Xia
Google



Danfei Xu
Stanford



Songfang Han
UCSD



Minhyuk Sung
KAIST



Shubham Tulsiani
CMU



Dieter Fox
University of Washington & NVIDIA




Niloy Mitra
UCL & Adobe



Siddhartha Chaudhuri
Adobe & IIT Bombay


3D Data and Applications




Robotics



Augmented Reality

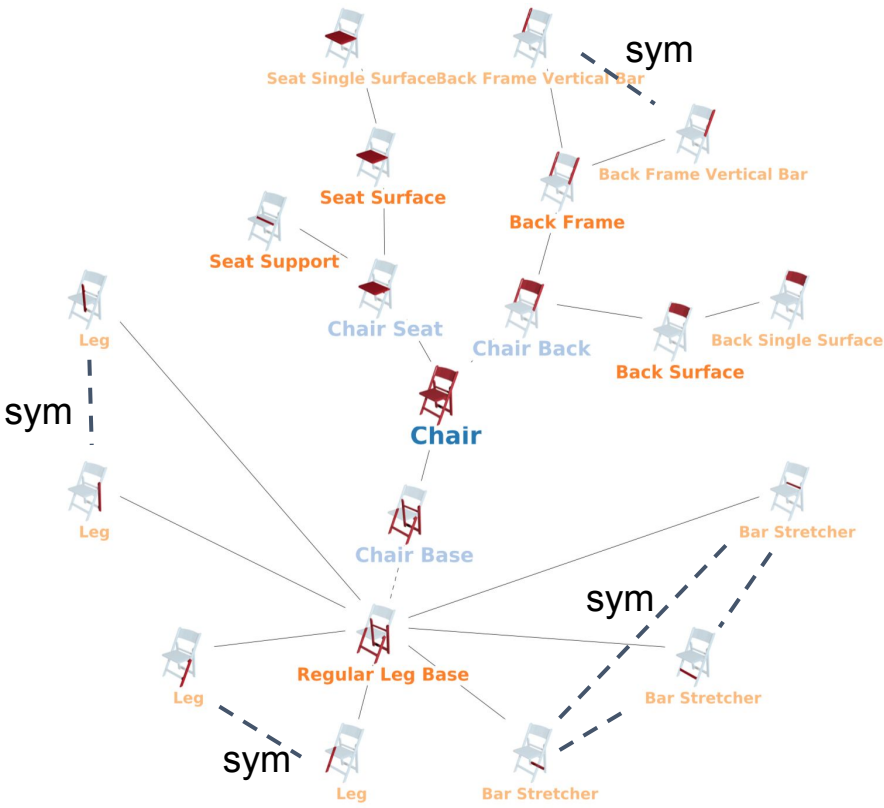


Autonomous driving




Medical Image Processing

3D Data Structures




```


for(i<4)
  sphere(pos=(0,i,0),
         color=8)
for(i<4)
  cylinder(pos=(1,2,i),
          color=5+i/2)
for(i<3)
  for(j<2+i)
    cube(pos=(2+i,1+j,0),
         color=(6-j))
  
```



RGB Image



Geometric Scene Graph



Symbolic Scene Graph



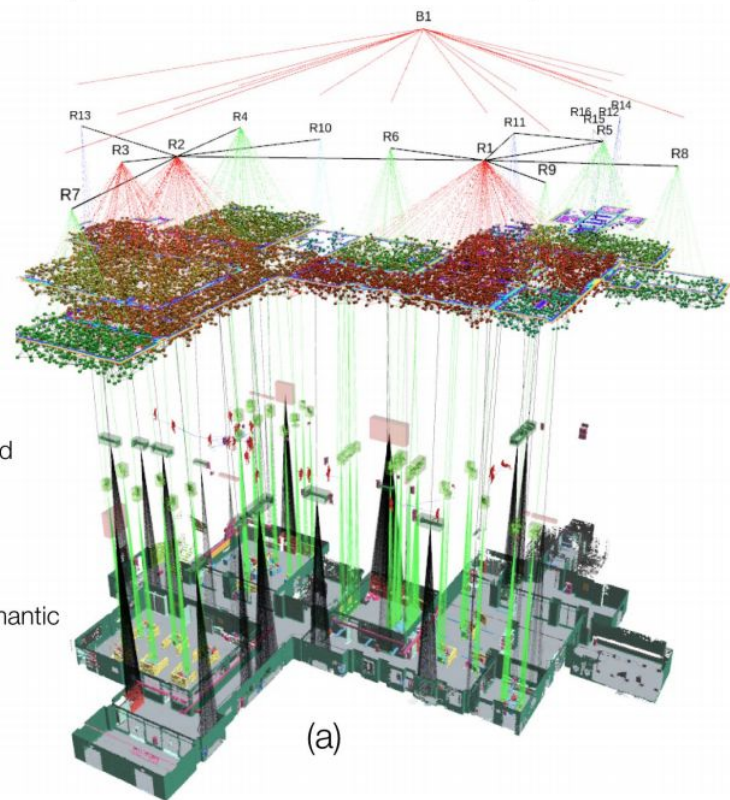
Layer 5: Buildings

Layer 4: Rooms

Layer 3: Places and Structures

Layer 2: Objects and Agents

Layer 1: Metric-Semantic Mesh



(a)

Image credits: Kaichun Mo, Fenggen Yu, Yifeng Zhu, Yunchao Liu, Antoni Rosinol, Kai Wang

3D Compositional Learning


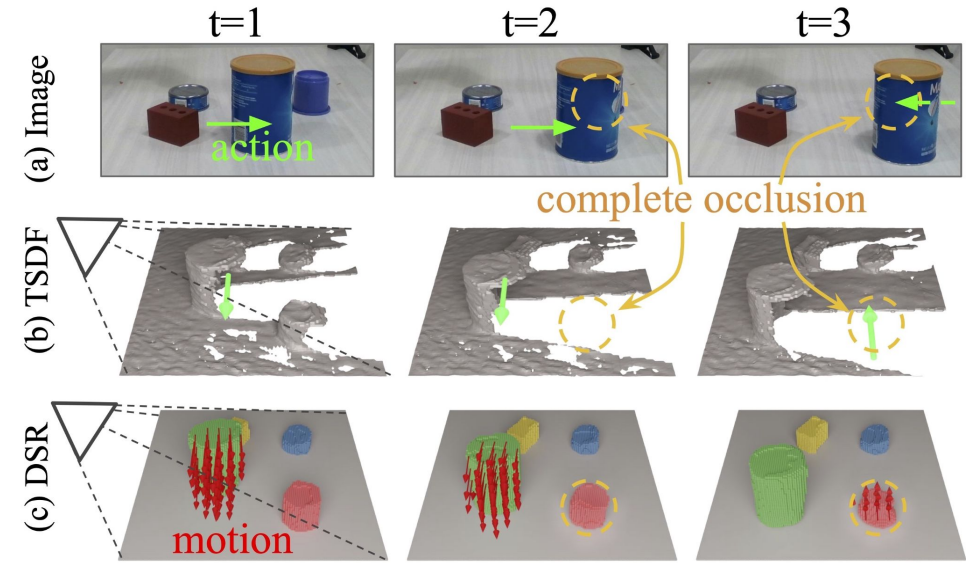
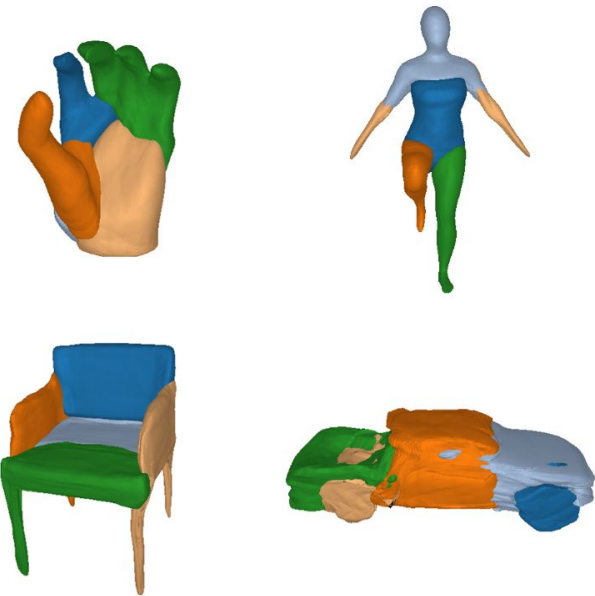


Image credits: Kyle Genova, Despoina Paschalidou, Songfang Han, Jason Zhang, Zhenjia Xu, Luca Weihs

Topics Covered

- How to design, discover, and represent structures for various 3D data?
- What structure to use for different task and applications?
- How people from different fields design and use 3D structure?
- What are the state-of-the-art algorithms over 3D data structure?
- How to discover implicit 3D data structures and use them for tasks?
- How should we mix structural and non-structural approaches?
-

Workshop Schedule

geometry.stanford.edu/struco3d/schedule.html



| Time (EDT) | Title | Speakers |
|--------------------|--|--|
| 7:00 – 7:10 AM | Welcome & Introductions | Kaichun Mo |
| 7:10 – 9:05 AM | Session 1: Structural Representations for 3D Scenes [ZOOM] | Helisa Dhamo, Siyuan Huang, Daniel Ritchie, Luca Carlone <i>(Hosts: Kai Wang, Antoni Rosinol)</i> |
| 9:10 – 11:05 AM | Session 2: Compositional Structures in Robotics [ZOOM] | Shuran Song, Krishna Murthy, Yifeng Zhu, Roozbeh Mottaghi <i>(Hosts: Fei Xia, Danfei Xu)</i> |
| 11:10 – 11:40 AM | Accepted Paper Poster Session 1 [Gatherly] | |
| 12:30 AM – 2:25 PM | Session 3: Structural Representations for 3D Shapes [ZOOM] | Leonidas J. Guibas, Hao (Richard) Zhang, Gopal Sharma, Jason Zhang <i>(Hosts: Minhyuk Sung, Kaichun Mo)</i> |
| 2:30 – 4:25 PM | Session 4: 3D Structural and Compositional Discovery [ZOOM] | Thomas Funkhouser, Hao Su, Kyle Genova, Despoina Paschalidou <i>(Hosts: Paul Guerrero, Songfang Han)</i> |
| 4:30 – 5:25 PM | Panel Discussion [ZOOM] | Shuran Song, Roozbeh Mottaghi, Hao Su, Hao (Richard) Zhang, Daniel Ritchie, Thomas Funkhouser <i>(Hosts: Shubham Tulsiani, Niloy Mitra)</i> |
| 5:30 – 6:00 PM | Accepted Paper Poster Session 2 [Gatherly] | |

Four Talk Sessions

geometry.stanford.edu/struco3d/schedule.html



| Time (EDT) | Title | Speakers |
|--------------------|--|---|
| 7:00 – 7:10 AM | Welcome & Introductions | Kaichun Mo |
| 7:10 – 9:05 AM | Session 1: Structural Representations for 3D Scenes | Helisa Dharmo, Siyuan Huang, Daniel Ritchie, Luca Carlone (Hosts: Kai Wang, Antoni Rosinol) |
| 9:10 – 11:05 AM | Session 2: Compositional Structures in Robotics | Shuran Song, Krishna Murthy, Yifeng Zhu, Roozbeh Mottaghi (Hosts: Fei Xia, Danfei Xu) |
| 11:10 – 11:40 AM | Accepted Paper Poster Session 1 | |
| 12:30 AM – 2:25 PM | Session 3: Structural Representations for 3D Shapes | Leonidas J. Guibas, Hao (Richard) Zhang, Gopal Sharma, Jason Zhang (Hosts: Minhyuk Sung, Kaichun Mo) |
| 2:30 – 4:25 PM | Session 4: 3D Structural and Compositional Discovery | Thomas Funkhouser, Hao Su, Kyle Genova, Despoina Paschalidou (Hosts: Paul Guerrero, Songfang Han) |
| 4:30 – 5:25 PM | Panel Discussion | Shuran Song, Roozbeh Mottaghi, Hao Su, Hao (Richard) Zhang, Daniel Ritchie, Thomas Funkhouser (Hosts: Shubham Tulsiani, Niloy Mitra) |
| 5:30 – 6:00 PM | Accepted Paper Poster Session 2 | |



Four Talk Sessions

geometry.stanford.edu/struco3d/schedule.html



| Time (EDT) | Title | Speakers |
|--------------------|--|--|
| 7:00 – 7:10 AM | Welcome & Introductions | Kaichun Mo |
| 7:10 – 9:05 AM | Session 1: Structural Representations for 3D Scenes | Helisa Dharmo, Siyuan Huang, Daniel Ritchie, Luca Carlone <i>(Hosts: Kai Wang, Antoni Rosinol)</i> |
| 9:10 – 11:05 AM | Session 2: Compositional Structures in Robotics | Shuran Song, Krishna Murthy, Yifeng Zhu, Roozbeh Mottaghi <i>(Hosts: Fei Xia, Danfei Xu)</i> |
| 11:10 – 11:40 AM | Accepted Paper Poster Session 1 | |
| 12:30 AM – 2:25 PM | Session 3: Structural Representations for 3D Shapes | Leonidas J. Guibas, Hao (Richard) Zhang, Gopal Sharma, Jason Zhang <i>(Hosts: Minhyuk Sung, Kaichun Mo)</i> |
| 2:30 – 4:25 PM | Session 4: 3D Structural and Compositional Discovery | Thomas Funkhouser, Hao Su, Kyle Genova, Despoina Paschalidou <i>(Hosts: Paul Guerrero, Songfang Han)</i> |
| 4:30 – 5:25 PM | Panel Discussion | Shuran Song, Roozbeh Mottaghi, Hao Su, Hao (Richard) Zhang, Daniel Ritchie, Thomas Funkhouser <i>(Hosts: Shubham Tulsiani, Niloy Mitra)</i> |
| 5:30 – 6:00 PM | Accepted Paper Poster Session 2 | |


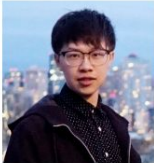




Four Talk Sessions

Each Session: 2 Hours + 4 talks

geometry.stanford.edu/struco3d/schedule.html



| Oct 16, 07:10 AM | Oct 16, 09:05 AM | Session 1: Structural Representations for 3D Scenes | |
|------------------|------------------|---|---|
| Oct 16, 07:10 AM | Oct 16, 07:15 AM | Opening Remarks for Session 1 Workshop Organizers (Kai Wang, Antoni Rosinol) | |
| Oct 16, 07:15 AM | Oct 16, 07:30 AM |  | [Spotlight] Describing 3D Indoor Scenes via Semantic Scene Graphs Helisa Dhamo (TUM) |
| Oct 16, 07:30 AM | Oct 16, 07:45 AM |  | [Spotlight] Compositional Structure in 3D Vision and Language Siyuan Huang (UCLA) |
| Oct 16, 07:45 AM | Oct 16, 08:25 AM |  | [Keynote] Learning to Infer and Generate Programs for 3D Shapes and Scenes Daniel Ritchie (Brown) |
| Oct 16, 08:25 AM | Oct 16, 09:05 AM |  | [Keynote] 3D Dynamic Scene Graphs: High-level Scene Understanding for Embodied Intelligence Luca Carlone (MIT) |



2 Spotlight Talks (10-min + Q&A)
from Students / Junior Researchers

2 Keynote Talks (30-min + Q&A)
from Professors / Senior Researchers

Four Talk Sessions

geometry.stanford.edu/struco3d/schedule.html



Invited Keynote Speakers



Leonidas J. Guibas
Stanford



Hao (Richard) Zhang
SFU




Daniel Ritchie
Brown




Luca Carlone
MIT



Thomas Funkhouser
Princeton & Google



Hao Su
UCSD



Shuran Song
Columbia University



Roozbeh Mottaghi
AI2 & University of Washington

Four Talk Sessions

geometry.stanford.edu/struco3d/schedule.html



Invited Spotlight Speakers




Gopal Sharma
UMass Amherst



Jason Zhang
CMU



Helisa Dharmo
TUM



Siyuan Huang
UCLA



Kyle Genova
Google Research



Despoina Paschalidou
Max Planck ETH



Krishna Murthy
Mila



Yifeng Zhu
UT Austin

Four Talk Sessions

geometry.stanford.edu/struco3d/schedule.html



Please Do Not Hesitate to Ask Questions!

- Live Q&As after Each Talk (type in Zoom chat box; ask for unmute)
- Leave questions in the asynchronous chats in ICCV Q&As
- Panel Discussion

Poster Presentations

workshopdaytwoposters.event.gatherly.io



| Time (EDT) | Title | Speakers |
|--------------------|--|---|
| 7:00 – 7:10 AM | Welcome & Introductions | Kaichun Mo |
| 7:10 – 9:05 AM | Session 1: Structural Representations for 3D Scenes | Helisa Dharmo, Siyuan Huang, Daniel Ritchie, Luca Carlone (Hosts: Kai Wang, Antoni Rosinol) |
| 9:10 – 11:05 AM | Session 2: Compositional Structures in Robotics | Shuran Song, Krishna Murthy, Yifeng Zhu, Roozbeh Mottaghi (Hosts: Fei Xia, Danfei Xu) |
| 11:10 – 11:40 AM | Accepted Paper Poster Session 1 | |
| 12:30 AM – 2:25 PM | Session 3: Structural Representations for 3D Shapes | Leonidas J. Guibas, Hao (Richard) Zhang, Gopal Sharma, Jason Zhang (Hosts: Minhyuk Sung, Kaichun Mo) |
| 2:30 – 4:25 PM | Session 4: 3D Structural and Compositional Discovery | Thomas Funkhouser, Hao Su, Kyle Genova, Despoina Paschalidou (Hosts: Paul Guerrero, Songfang Han) |
| 4:30 – 5:25 PM | Panel Discussion | Shuran Song, Roozbeh Mottaghi, Hao Su, Hao (Richard) Zhang, Daniel Ritchie, Thomas Funkhouser (Hosts: Shubham Tulsiani, Niloy Mitra) |
| 5:30 – 6:00 PM | Accepted Paper Poster Session 2 | |



Poster Presentations

workshopdaytwoposters.event.gatherly.io



Archival Workshop Publications [[CVF Official Archive](#)]

| | | |
|---|---|---|
| 1 | <p>Pinak Paliwal, Vikas Paliwal</p> <p>14D Systems Inc. and UC Berkeley</p> <p>[PDF] [Video]</p> <p>Poster Session 1</p> | <p>3D Scene Angles using UL Decomposition of Planar Homography</p> <p>Show Abstract</p> |
| 2 | <p>Rinon Gal, Amit H Bermano, Hao Zhang, Danny Cohen-Or</p> <p>Tel Aviv University, Simon Fraser University</p> <p>[PDF] [Supp] [Video]</p> <p>Poster Session 1 & 2</p> | <p>MRGAN: Multi-Rooted 3D Shape Representation Learning with Unsupervised Part Disentanglement</p> <p>Show Abstract</p> |
| 3 | <p>Siddharth Katageri, Shashidhar V Kudari, Akshaykumar Gunari, Ramesh Tabib, Uma Mudenagudi</p> <p>KLE Technological University</p> <p>[PDF] [Video]</p> <p>Poster Session 1 & 2</p> | <p>ABD-Net: Attention Based Decomposition Network for 3D Point Cloud Decomposition</p> <p>Show Abstract</p> |

Poster Presentations



workshopdaytwoposters.event.gatherly.io

Non-archival Short Presentations

| | | |
|---|---|--|
| 4 | Jianglong Ye, Yuntao Chen, Naiyan Wang, Xiaolong Wang TuSimple, Chinese Academy of Sciences, UCSD [PDF] [Video] Poster Session 2 | Online Adaptation for Implicit Object Tracking and Shape Reconstruction in the Wild Show Abstract |
| 5 | Rishabh Baghel, Abhishek Trivedi, Tejas Ravichandran, Ravi Kiran Sarvadevabhatla IIIT Hyderabad [PDF] [Video] Poster Session 1 & 2 | MeronymNet: A Unified Framework for Part and Category Controllable Generation of Objects Show Abstract |
| 6 | Jiashun Wang, Huazhe Xu, Medhini Narasimhan, Xiaolong Wang UCSD, UC Berkeley [PDF] [Video] Poster Session 2 | Multi-Person 3D Motion Prediction with Multi-Range Transformers Show Abstract |
| 7 | Yuzhe Qin, Yueh-Hua Wu, Shaowei Liu, Hanwen Jiang, Ruihan Yang, Yang Fu, Xiaolong Wang UCSD, National Taiwan University / Academia Sinica, UIUC [PDF] [Video] Poster Session 2 | DexMV: Imitation Learning for Dexterous Manipulation from Human Videos Show Abstract |
| 8 | Hanwen Jiang, Shaowei Liu, Jiashun Wang, Xiaolong Wang UCSD [PDF] [Video] Poster Session 1 & 2 | Hand-Object Contact Consistency Reasoning for Human Grasps Generation Show Abstract |
| 9 | Jiteng Mu, Weichao Qiu, Adam Kortylewski, Alan Yuille, Nuno Vasconcelos, Xiaolong Wang UCSD, Johns Hopkins University [PDF] [Video] Poster Session 1 & 2 | A-SDF: Learning Disentangled Signed Distance Functions for Articulated Shape Representation Show Abstract |

| | | |
|----|--|--|
| 10 | Alexey Bokhovkin, Vladislav Ishimtsev, Emil Bogomolov, Denis Zorin, Alexey Artemov, Evgeny Burnaev, Angela Dai TUM, Skoltech, NYU [PDF] [Video] Poster Session 1 | Towards Part-Based Understanding of RGB-D Scans Show Abstract |
| 11 | Yining Hong, Qing Li, Song-Chun Zhu, Siyuan Huang UCLA, Beijing Institute for General Artificial Intelligence, Tsinghua University, Peking University [PDF] [Video] Poster Session 1 & 2 | VLGrammar: Grounded Grammar Induction of Vision and Language Show Abstract |
| 12 | Mikaela Angelina Uy, Vladimir G. Kim, Minhyuk Sung, Noam Aigerman, Siddhartha Chaudhuri, Leonidas Guibas Stanford University, Adobe Research, KAIST, IIT Bombay [PDF] [Supp] [Video] Poster Session 2 | Joint Learning of 3D Shape Retrieval and Deformation Show Abstract |
| 13 | Minghua Liu, Minhyuk Sung, Radomir Mech, Hao Su UCSD, KAIST, Adobe Research [PDF] [Video] Poster Session 1 & 2 | DeepMetaHandles: Learning Deformation Meta-Handles of 3D Meshes with Biharmonic Coordinates Show Abstract |
| 14 | Yikai Li, Jiayuan Mao, Xiuming Zhang, William T. Freeman, Joshua B. Tenenbaum, Noah Snavely, Jiajun Wu MIT CSAIL, Shanghai Jiao Tong University, Google Research, Stanford University [PDF] [Supp] [Video] Poster Session 1 & 2 | Multi-Plane Program Induction with 3D Box Priors Show Abstract |

Poster Presentations

workshopdaytwoposters.event.gatherly.io



Program Committee

- Mika Uy (Stanford)
- Ian Huang (Stanford)
- Li Yi (Tsinghua)
- Zhenyu Jiang (UT Austin)
- Minghua Liu (UCSD)
- Jiayuan Gu (UCSD)
- Tiange Luo (UCSD)
- Jie Yang (Chinese Academy of Sciences)
- Kenny Jones (Brown)
- Xianghao Xu (Brown)
- Wamiq Para (KAUST)
- Nilesh Kulkarni (University of Michigan)
- Yufei Ye (CMU)



Poster Presentations

workshopdaytwoposters.event.gatherly.io



Archival Workshop Publications [[CVF Official Archive](#)]

| | | |
|---|--|---|
| 1 | <p>Pinak Paliwal, Vikas Paliwal</p> <p>14D Systems Inc. and UC Berkeley</p> <p>[PDF] [Video]</p> <p>Poster Session 1</p> | <p>3D Scene Angles using UL Decomposition of Planar Homography</p> <p>Show Abstract</p> |
| 2 | <p>Rinon Gal, Amit H Bermano, Hao Zhang, Danny Cohen-Or</p> <p>Tel Aviv University, Simon Fraser University</p> <p>[PDF] [Supp] [Video]</p> <p>Poster Session 1 & 2</p> | <p>MRGAN: Multi-Rooted 3D Shape Representation Learning with Unsupervised Part Disentanglement</p> <p>Show Abstract</p> |
| 3 | <p>Siddharth Katageri, Shashidhar V Kudari, Akshaykumar Gunari, Ramesh Tabib, Uma Mudenagudi</p> <p>KLE Technological University</p> <p>[PDF] [Video]</p> <p>Poster Session 1 & 2</p> | <p>ABD-Net: Attention Based Decomposition Network for 3D Point Cloud Decomposition</p> <p>Show Abstract</p> |

Poster Presentations

2 x 30-min Gatherly Sessions

workshopdaytwoposters.event.gatherly.io

2021 ICCV OCTOBER 11-17 VIRTUAL Get Gatherly!

Welcome to 10/16 Workshops Poster Session, Kaichun Mo!

- As people join, they'll appear in the **Map** and the **People** tab.
- To **start a video chat**, click someone's avatar or click a green huddle to join an ongoing conversation.

Kaichun Mo (me)

Audio & Video settings Help Network connection: FAST

Structural and Com... 1 person Brochure

1 Elevators (Select a floor) Kaichun Mo Exit event

Tip: Click an empty space on the map to exit the elevator.

Search for a floor

No people

Self-supervised Learning for Next-Generation Industry-level Autonomous Driving No people

2 Structural and Compositional Learning on 3D Data 1 person

1st Workshop, Benchmark and Challenge on Human Trajectory and Pose Dynamics Forecasting in the



Panel Discussion

geometry.stanford.edu/struco3d/schedule.html



| Time (EDT) | Title | Speakers |
|--------------------|--|---|
| 7:00 – 7:10 AM | Welcome & Introductions | Kaichun Mo |
| 7:10 – 9:05 AM | Session 1: Structural Representations for 3D Scenes | Helisa Dharmo, Siyuan Huang, Daniel Ritchie, Luca Carlone (Hosts: Kai Wang, Antoni Rosinol) |
| 9:10 – 11:05 AM | Session 2: Compositional Structures in Robotics | Shuran Song, Krishna Murthy, Yifeng Zhu, Roozbeh Mottaghi (Hosts: Fei Xia, Danfei Xu) |
| 11:10 – 11:40 AM | Accepted Paper Poster Session 1 | |
| 12:30 AM – 2:25 PM | Session 3: Structural Representations for 3D Shapes | Leonidas J. Guibas, Hao (Richard) Zhang, Gopal Sharma, Jason Zhang (Hosts: Minhyuk Sung, Kaichun Mo) |
| 2:30 – 4:25 PM | Session 4: 3D Structural and Compositional Discovery | Thomas Funkhouser, Hao Su, Kyle Genova, Despoina Paschalidou (Hosts: Paul Guerrero, Songfang Han) |
| 4:30 – 5:25 PM | Panel Discussion | Shuran Song, Roozbeh Mottaghi, Hao Su, Hao (Richard) Zhang, Daniel Ritchie, Thomas Funkhouser (Hosts: Shubham Tulsiani, Niloy Mitra) |
| 5:30 – 6:00 PM | Accepted Paper Poster Session 2 | |



Panel Discussion


geometry.stanford.edu/struco3d/schedule.html




Hosts



Leonidas J. Guibas
Stanford



Hao (Richard) Zhang
SFU




Daniel Ritchie
Brown



Luca Carlone
MIT



Thomas Funkhouser
Princeton & Google



Hao Su
UCSD



Shuran Song
Columbia University



Roozbeh Mottaghi
AI2 & University of Washington

Panelists

Panel Discussion

geometry.stanford.edu/struco3d/schedule.html



People from different fields or backgrounds use different structural and compositional representations of their 3D data for different applications. We bring them together in this workshop to have an explicit discussion of the advantages and disadvantages of different representations and approaches, as well as to share, discuss and debate the diverse opinions regarding the following questions:

- Which types of structure should we use for different tasks and applications in graphics, vision and robotics?
- How should we factorize a given problem into sparse concepts that make up the structure?
- How should we factorize different types of 3D data into sparse sets of components, relationships, or operators?
- Which algorithms are best suited for a given type of structure?
- How should we mix structural and non-structural approaches?
- Which parts of a problem are suited for structural approaches, and which ones are better handled without structure?

+ Your Questions!

2021 **ICCV** OCTOBER 11-17
VIRTUAL



1st StruCo3D Workshop

Structural and Compositional Learning on 3D Data

Enjoy the Talks, Papers, and Panel!

geometry.stanford.edu/struco3d

