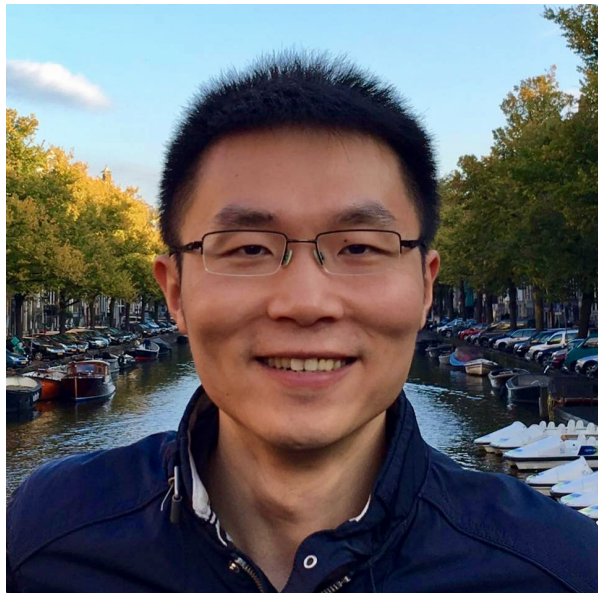IARPA Video LINCS Proposers' Day Lightning Talk

Jia Deng

Princeton University

Feb 7, 2024

PI: Jia Deng

- Associate Professor
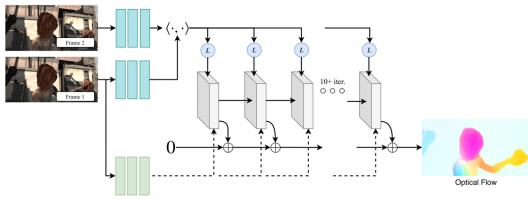- Department of Computer Science
- Princeton University
- jiadeng@princeton.edu
- Lab: https://pvl.cs.princeton.edu

Capability 1: 3D Reconstruction

- Optimization-inspired neural networks (neural + classical)
  - Recurrent neural networks + differentiable optimization layers
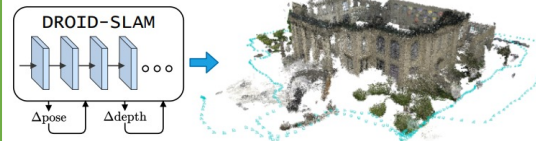 - Large improvements over prior SOTA on many 3D vision tasks

Optical Flow

RAFT
- ECCV 2020
- 30-40% less error

ECCV Best Paper

SLAM
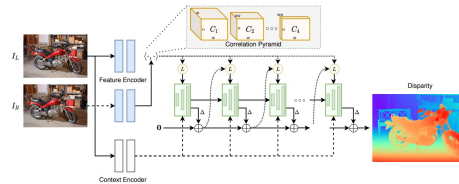
DROID-SLAM
- NeurIPS 2021
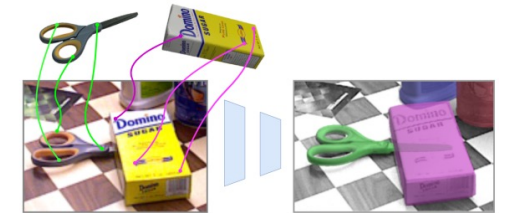- 60-80% less error

NeurIPS Oral

Stereo Depth

RAFT-Stereo
- 3DV 2021
- 30% less error

3DV Best Student Paper

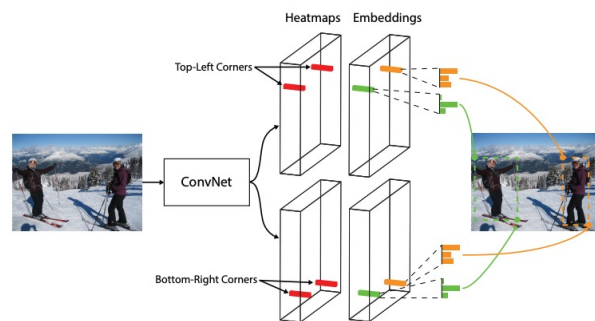**6DoF Object Pose**

Coupled Iterative Refinement
- ECCV 2022

Co-Winner, BOP Challenge

Capability 2: Object Recognition

- Extensive experience in object recognition
- PI was first author of ImageNet

Object Detection
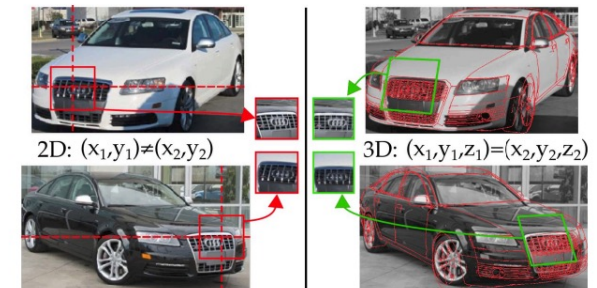
CornerNet
- ECCV 2018

Citations: 3978

Human Pose

Hourglass Networks
- ECCV 2016

Citations: 5818

Vehicle Recognition

Fine-grained vehicle recognition
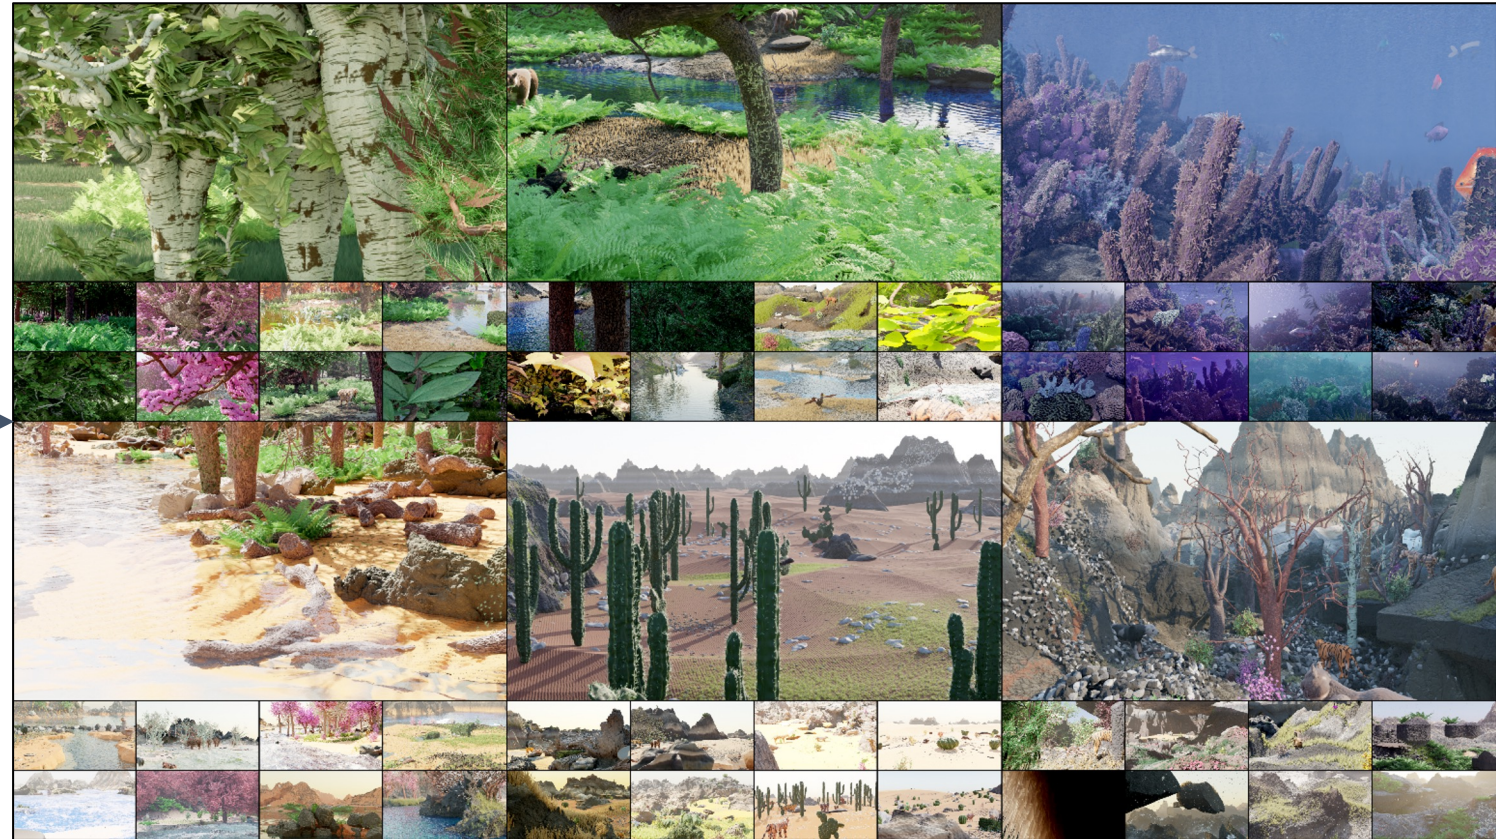- CVPR 2013 Workshop

Citations: 3333

Capability 3: Procedural Synthetic Data

- Infinigen: a procedural *generator* of infinite synthetic data (graphics-based)
  - CVPR 2023
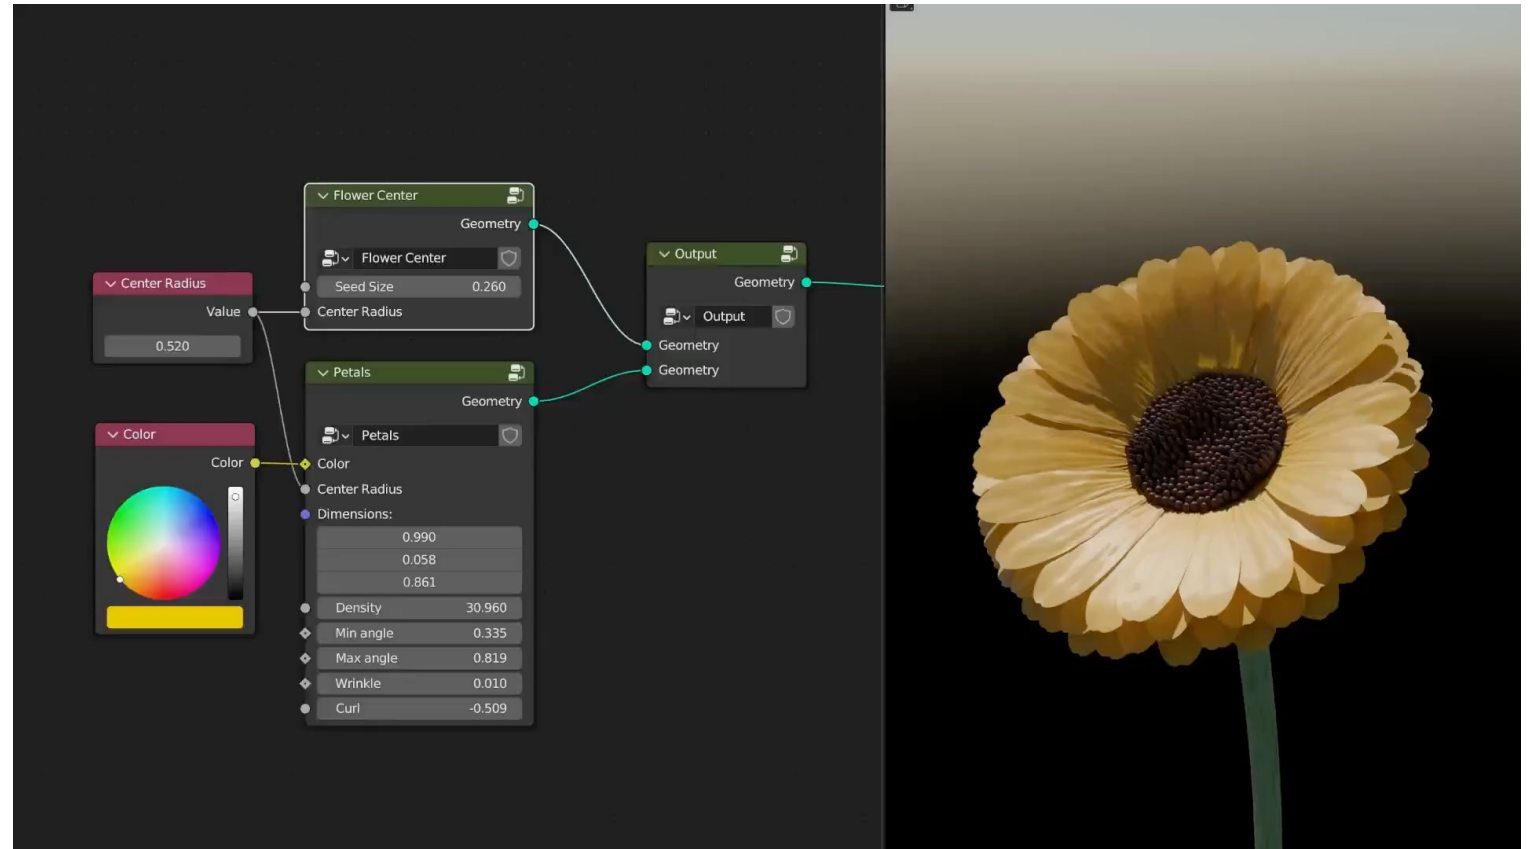 - infinigen.org
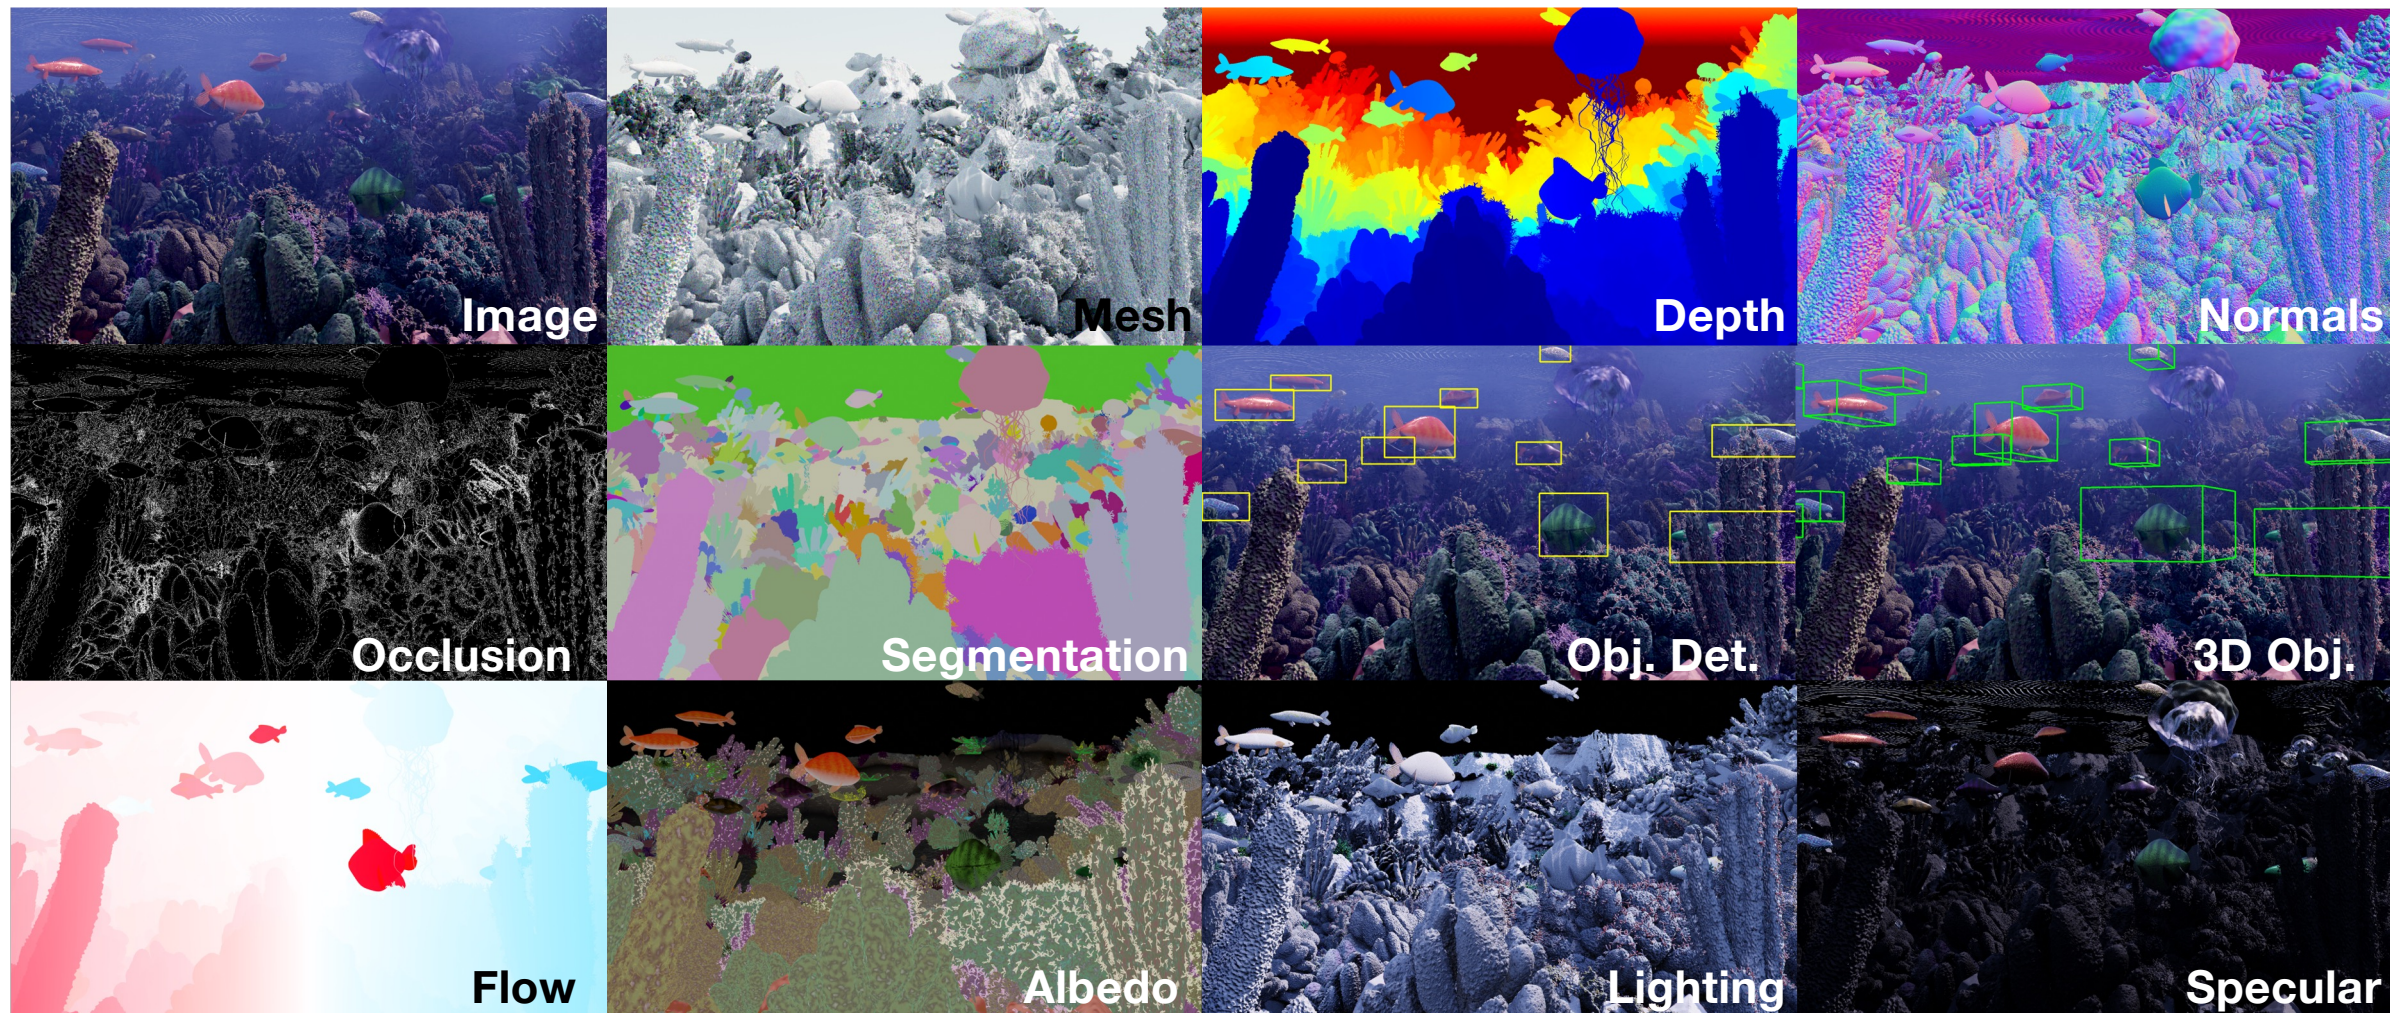

Random seeds +
Optional user controls

(Random sample of generated scenes)

Infinigen is 100% Procedural

- Generated from scratch using randomized mathematical rules
- Everything (shape, texture, material, lighting, arrangement, animation) is procedural
- No external assets
- No AI

Infinigen creates automatic labels

Infinigen is free and open-source (BSD)

https://github.com/princeton-vl/infinigen

Customizing Infinigen for Video LINCS

- Data is often the bottleneck

- Infinite tailored synthetic data:

  - Custom procedural objects (e.g. vehicles) placed in diverse procedural scenes