




1950


FUTURE VISION



2017 MWF 1PM 368

COMPUTER VISION

Project 6 mark distribution



Paper 3: ScanNet (CVPR 2017)

Has everything!

Stereo, geometry reconstruction, dataset generation, mturk, classification, CNNs, dataset retrieval – oh my.



How does a depth camera work?

Stereo in infrared.




Stereo correspondence

- Let x be a point in left image, x' in right image
- Epipolar relation
 - x maps to epipolar line l'
 - x' maps to epipolar line l




How does a depth camera work?

Stereo in infrared.



Time of Flight





ScanNet: Richly-annotated 3D Reconstructions of Indoor Scenes

Angela Dai Angel X. Chang Manolis Savva Maciej Halber

Thomas Funkhouser Matthias Nießner

Stanford University
Princeton University
Technical University of Munich

CVPR 2017 (Spotlight)

BundleFusion: Real-time Globally Consistent 3D Reconstruction using Online Surface Re-integration

*Angela Dai¹ Matthias Nießner¹
Michael Zollhöfer² Shahram Izadi³
Christian Theobalt²*

¹Stanford University

²Max Planck Institute for Informatics

³Microsoft Research

(contains audio)

The final illusion...

...is a transformation.

January 2017...

You (actually me)



Computer vision –
it's pixels and stuff.

We learned about image formation.

Light, color, frequency, and filtering.

May 2017...

If I feed my CNN images with
more than 3 spectral
responses, will it still learn a
compressed 3-dimensional
basis for color?

You



We learned about camera geometry.

Camera calibration, stereo, multi-view geometry.

May 2017...

Is human perceptual shape
understanding more like
photogrammetry or more
like database learning?

You



We learned about features and matching.

Gradients, local interest points,
motion models, model fitting.

May 2017...

My brain can compute
correspondence so quickly,
but it's a pretty rough
approximate – maybe we can
simulate this to speed up
feature matching?

You



We learned about recognition.

Supervised and unsupervised learning,
detection, classification, validation.

May 2017...

Maybe I can use
dimensionality reduction
as a pre-process to
factor out the
illumination variance in
my small dataset...

You



We learned about ConvNets.

Neural networks, regularization, inception.

May 2017...

If we can fool a CNN easily by tweaking gradients, but not a human, then what visual information am I not yet modeling?

You



Modeling visual inference problems.

Your MATLAB-fu is also better.

What's next?

Spring 2018: CSCI 2951-I

- Computer Vision for Graphics and Interaction
- Read state of the art papers + present + discuss
- Learn *how* to read a paper
- Learn *how* to present a paper
- Open-ended project

Inferring physically-meaningful models.
How does augmented reality work?



Please give us feedback

- I will put up an anon Google form; please let me & the TAs know what you liked / what we can improve.
- The Critical Review