

SIMON FRASER
UNIVERSITY



CVPR is a contemporary art exhibition
-- Garbage is a source for impact --

3D Scene Understanding for Vision, Graphics,
and Robotics Workshop @ CVPR2020

Yasutaka Furukawa

We are scared of reviewers...



CVPR is a contemporary art exhibition
-- Garbage is a source for impact --

3D Scene Understanding for Vision, Graphics, and Robotics

CVPR 2020 Workshop, Virtual, June 15th, 2020

[Introduction](#) [Program](#) [Schedule](#) [Previous Workshop](#)



News

Due to the pandemic, our workshop will totally be virtual this year. We will host an online chat room for communication with the speakers and Q&A. The details are coming soon. Looking forward to meet you online!

Invited talks and oral presentations will be presented live or by recorded videos in the same Zoom room, all of the talks will have live Q&A session, please refer to the [Program](#) for recorded videos and more details.

Invited Speakers



Kristen Grauman (UT Austin)



Sergey Levine (UC Berkeley)



Andreas Geiger (University of Tübingen)



Yasutaka Furukawa (Simon Fraser University)



Daniel Ritchie (Brown University)



Jeannette Bohg (Stanford University)



Shuran Song (Columbia University)



Andrea Tagliasacchi (Google Brain)



Katerina Fragkiadaki (Carnegie Mellon University)

Opening Remark




David Forsyth (University of Illinois Urbana-Champaign)

ScanNet Indoor Scene Understanding Challenge

CVPR 2020 Workshop, Seattle, WA

June 19, 2020



Introduction

3D scene understanding for indoor environments is becoming an increasingly important area. Application domains such as augmented and virtual reality, computational photography, interior design, and autonomous mobile robots all require a deep understanding of 3D interior spaces, the semantics of objects that are present, and their relative configurations in 3D space.

We present the first comprehensive challenge for 3D scene understanding of entire rooms at the object instance-level with 5 tasks based on the ScanNet dataset. The ScanNet dataset is a large-scale semantically annotated dataset of 3D mesh reconstructions of interior spaces (approx. 1500 rooms and 2.5 million RGB-D frames). It is used by more than 480 research groups to develop and benchmark state-of-the-art approaches in semantic scene understanding. A key goal of this challenge is to compare state-of-the-art approaches operating on image data (including RGB-D) with approaches operating directly on 3D data (point cloud, or surface mesh representations). Additionally, we pose both object category label prediction (commonly referred to as semantic segmentation), and instance-level object recognition (object instance prediction and category label prediction). We propose five tasks that cover this space:


- **2D semantic label prediction:** prediction of object category labels from 2D image representation
- **2D semantic instance prediction:** prediction of object instance and category labels from 2D image representation
- **3D semantic label prediction:** prediction of object category labels from 3D representation
- **3D semantic instance prediction:** prediction of object instance and category labels from 3D representation
- **Scene type classification:** classification of entire 3D room into a scene type





For each task, challenge participants are provided with prepared training, validation, and test datasets, and automated evaluation scripts. In addition to the public train-val-test split, benchmarking is done on a hidden test set whose raw data can be downloaded without annotations; in order to participate in the benchmark, the predictions on the hidden test set are uploaded to the evaluation server, where they are evaluated. Submission is restricted to submissions every two weeks to avoid finetuning on the test dataset. See more details at http://kaldir.vc.in.tum.de/scannet_benchmark/documentation if you would like to participate in the challenge. The evaluation server leaderboard is live at http://kaldir.vc.in.tum.de/scannet_benchmark/.

CAD model
Structured geometry
Semantically segmented geometry

- Model editing
- texture/normal/bump/effects mapping
- Physics simulations
- Interactions (VR/AR)
- Architectural analysis
- ...

[Transformer]






Structured Geometry Reconstruction




CVPR 2009



CVPR 2014




ICCV 2017




ECCV 2018/ICCV 2019



ECCV 2012



ICCV 2015




ECCV 2018



CVPR 2020



CVPR 2018/2019



Ice-age


Revolution




CVPR 2009



CVPR 2014




ICCV 2017




ECCV 2018/ICCV 2019



ECCV 2012



ICCV 2015



ECCV 2018



CVPR 2020



CVPR 2018/CVPR 2019

Ice-age (no impact)


Revolution (impact)



CVPR 2009




CVPR 2014



ECCV 2012



ICCV 2015




ECCV 2018




CVPR 2018/CVPR 2019



CVPR 2020



ECCV 2018/ICCV 2019



ICCV 2017



Aaron Hertzmann

June 8 at 9:19 AM · 🌐



I've posted a new essay here which is the product of my recent attempts to better understand contemporary art. There are some pretty fundamental things here that I didn't get from any of the books that I read on this topic, nor from my undergraduate art degree. I'm really grateful to **Jason Salavon** for all his time spent discussing this with me.



AARONHERTZMANN.COM

What is contemporary art?

I recently had some private conversations with artists that changed how I understand contemp...

Like · Reply · 4d



Aaron Hertzmann replied · 1 Reply



Alyosha Efros

I like the comparison to scientific communities. Another commonality is that, just like 90% of CVPR papers will turn out to be junk (either immediately or after a few years), so is 90% of contemporary art. This is normal and we shouldn't be disturbed by it. Just as I am happy to find a few CVPR papers that I like, I should feel happy if, at any contemporary art exhibition, I find 1-2 pieces that speak to me.

Like · Reply · 4d



1



Aaron Hertzmann

https://en.wikipedia.org/wiki/Sturgeon%27s_law

...

EN.WIKIPEDIA.ORG

Sturgeon's law - Wikipedia

Like · Reply · 4d



2

Like · Reply · 4d



Aaron Hertzmann replied · 1 Reply



Alyosha Efros

I like the comparison to scientific communities. Another commonality is that, just like 90% of CVPR papers will turn out to be junk (either immediately or after a few years), so is 90% of contemporary art. This is normal and we shouldn't be disturbed by it. Just as I am happy to find a few CVPR papers that I like, I should feel happy if, at any contemporary art exhibition, I find 1-2 pieces that speak to me.

...



1

Like · Reply · 4d



Aaron Hertzmann

https://en.wikipedia.org/wiki/Sturgeon%27s_law

...

EN.WIKIPEDIA.ORG

Sturgeon's law - Wikipedia

Like · Reply · 4d



2

Like · Reply · 4d

→  Aaron Hertzmann replied · 1 Reply



Alyosha Efros

I like the comparison to scientific communities. Another commonality is that, just like 90% of CVPR papers will turn out to be junk (either immediately or after a few years), so is 90% of contemporary art. This is normal and we shouldn't be disturbed by it. Just as I am happy to find a few CVPR papers that I like, I should feel happy if, at any

...

CVPR is a contemporary art exhibition, I find 1-2 pieces that speak to me.



1

Like · Reply · 4d



Aaron Hertzmann

https://en.wikipedia.org/wiki/Sturgeon%27s_law

...

EN.WIKIPEDIA.ORG

Sturgeon's law - Wikipedia

Like · Reply · 4d



2



[Main page](#)
[Contents](#)
[Current events](#)
[Random article](#)
[About Wikipedia](#)
[Contact us](#)
[Donate](#)

[Contribute](#)
[Help](#)
[Community portal](#)
[Recent changes](#)
[Upload file](#)

[Tools](#)
[What links here](#)

Not logged in [Talk](#) [Contributions](#) [Create account](#) [Log in](#)

[Article](#)

[Talk](#)

[Read](#)

[Edit](#)

[View history](#)



Sturgeon's law

From Wikipedia, the free encyclopedia

Sturgeon's law (or **Sturgeon's revelation**), is an [adage](#) that states that "ninety percent of everything is crap."

The adage was coined by [Theodore Sturgeon](#), an American [science fiction author](#) and critic. The adage was inspired by Sturgeon's observation while [science fiction](#) was often derided for its low quality by critics, the majority of examples of works in other fields could equally be seen to be of low quality, and science fiction was thus no different in that regard from other [art](#).

Some make a distinction between the revelation ("ninety percent of everything is crap") and the law ("nothing is always absolutely so").

Contents [hide]

- [History](#)
- [The original Sturgeon's law](#)
- [See also](#)
- [Notes](#)
- [External links](#)

CVPR is a contemporary art exhibition

Let's stop pretending that all our
2000 CVPR papers convey good ideas

Garbage is a source for impact

We are scared of reviewers...



Ice-age (no impact)


Revolution (impact)



CVPR 2009




CVPR 2014



ECCV 2012



ICCV 2015




ECCV 2018




CVPR 2018/CVPR 2019



CVPR 2020



ECCV 2018/ICCV 2019




ICCV 2017

Ice-age (no impact)

Revolution (impact)




CVPR 2009



Metropolitan Museum of Art


ECCV 2012



ICCV 2015



CVPR 2014





ECCV 2018




CVPR 2018/CVPR 2019



CVPR 2020



ECCV 2018/ICCV 2019

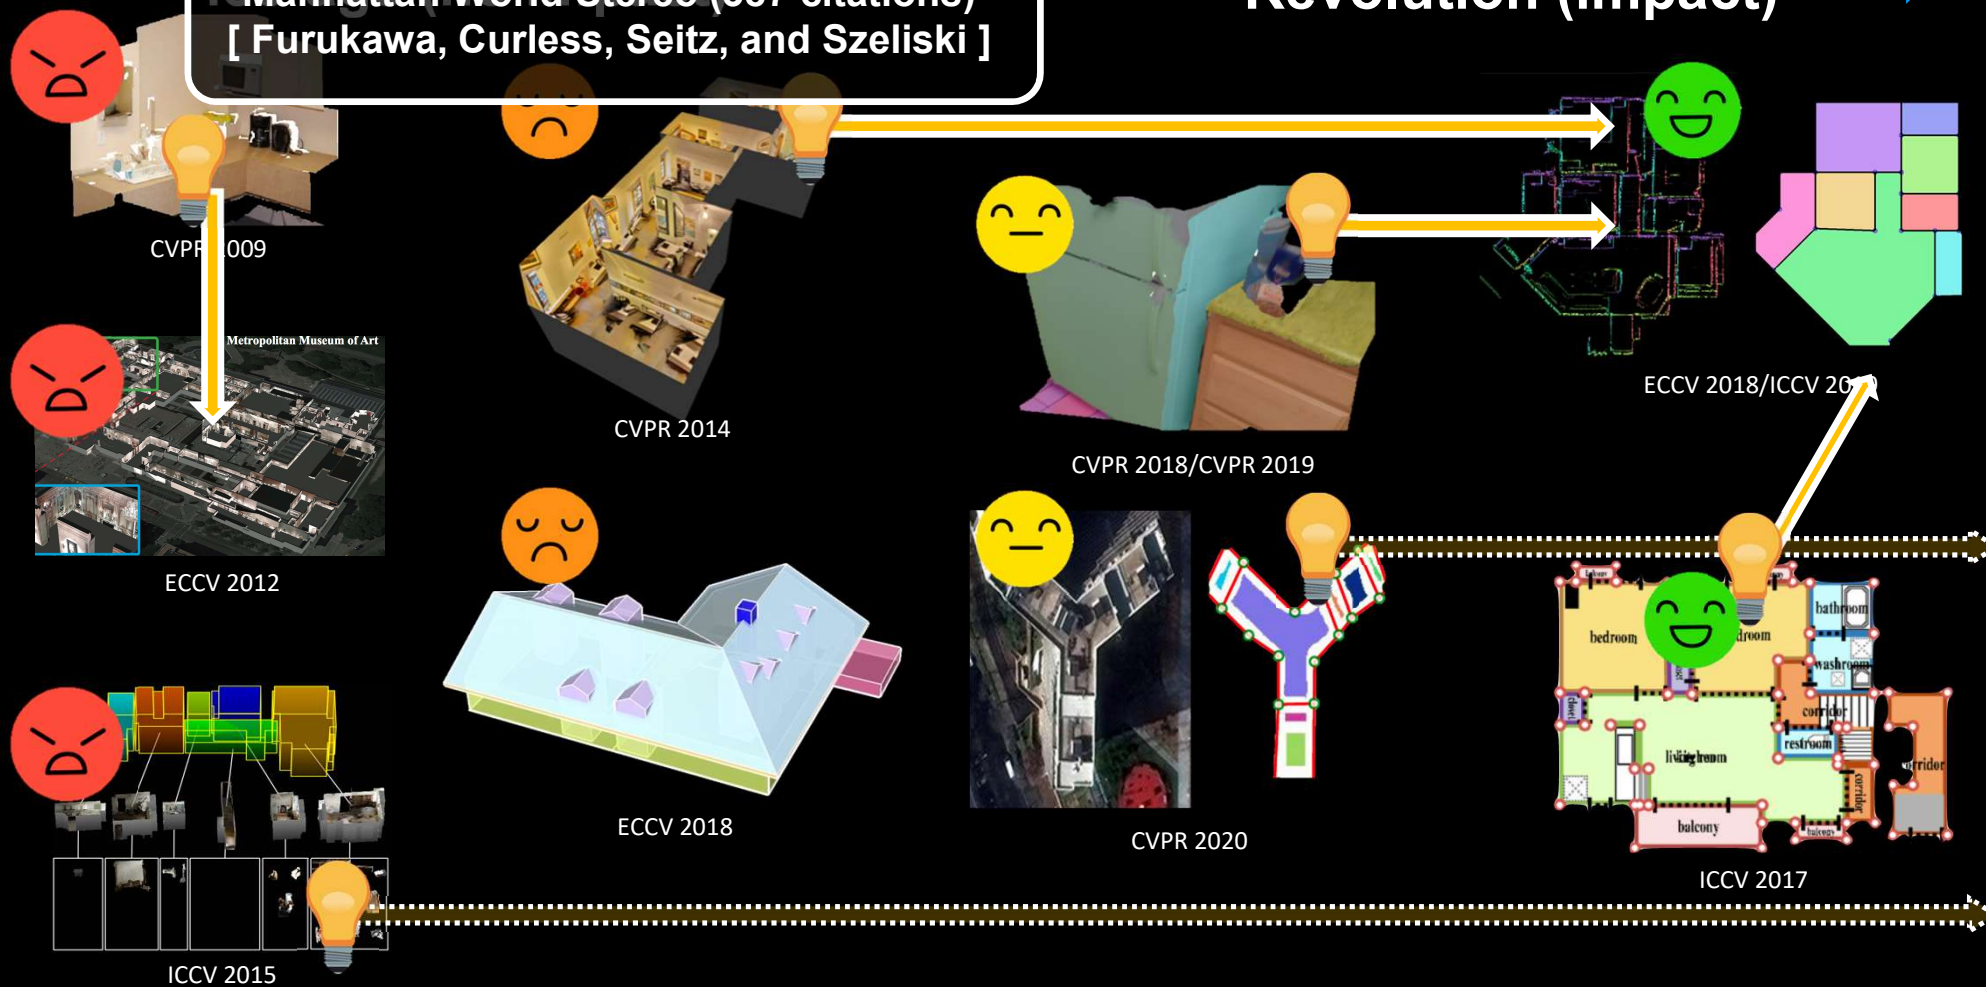


ICCV 2017



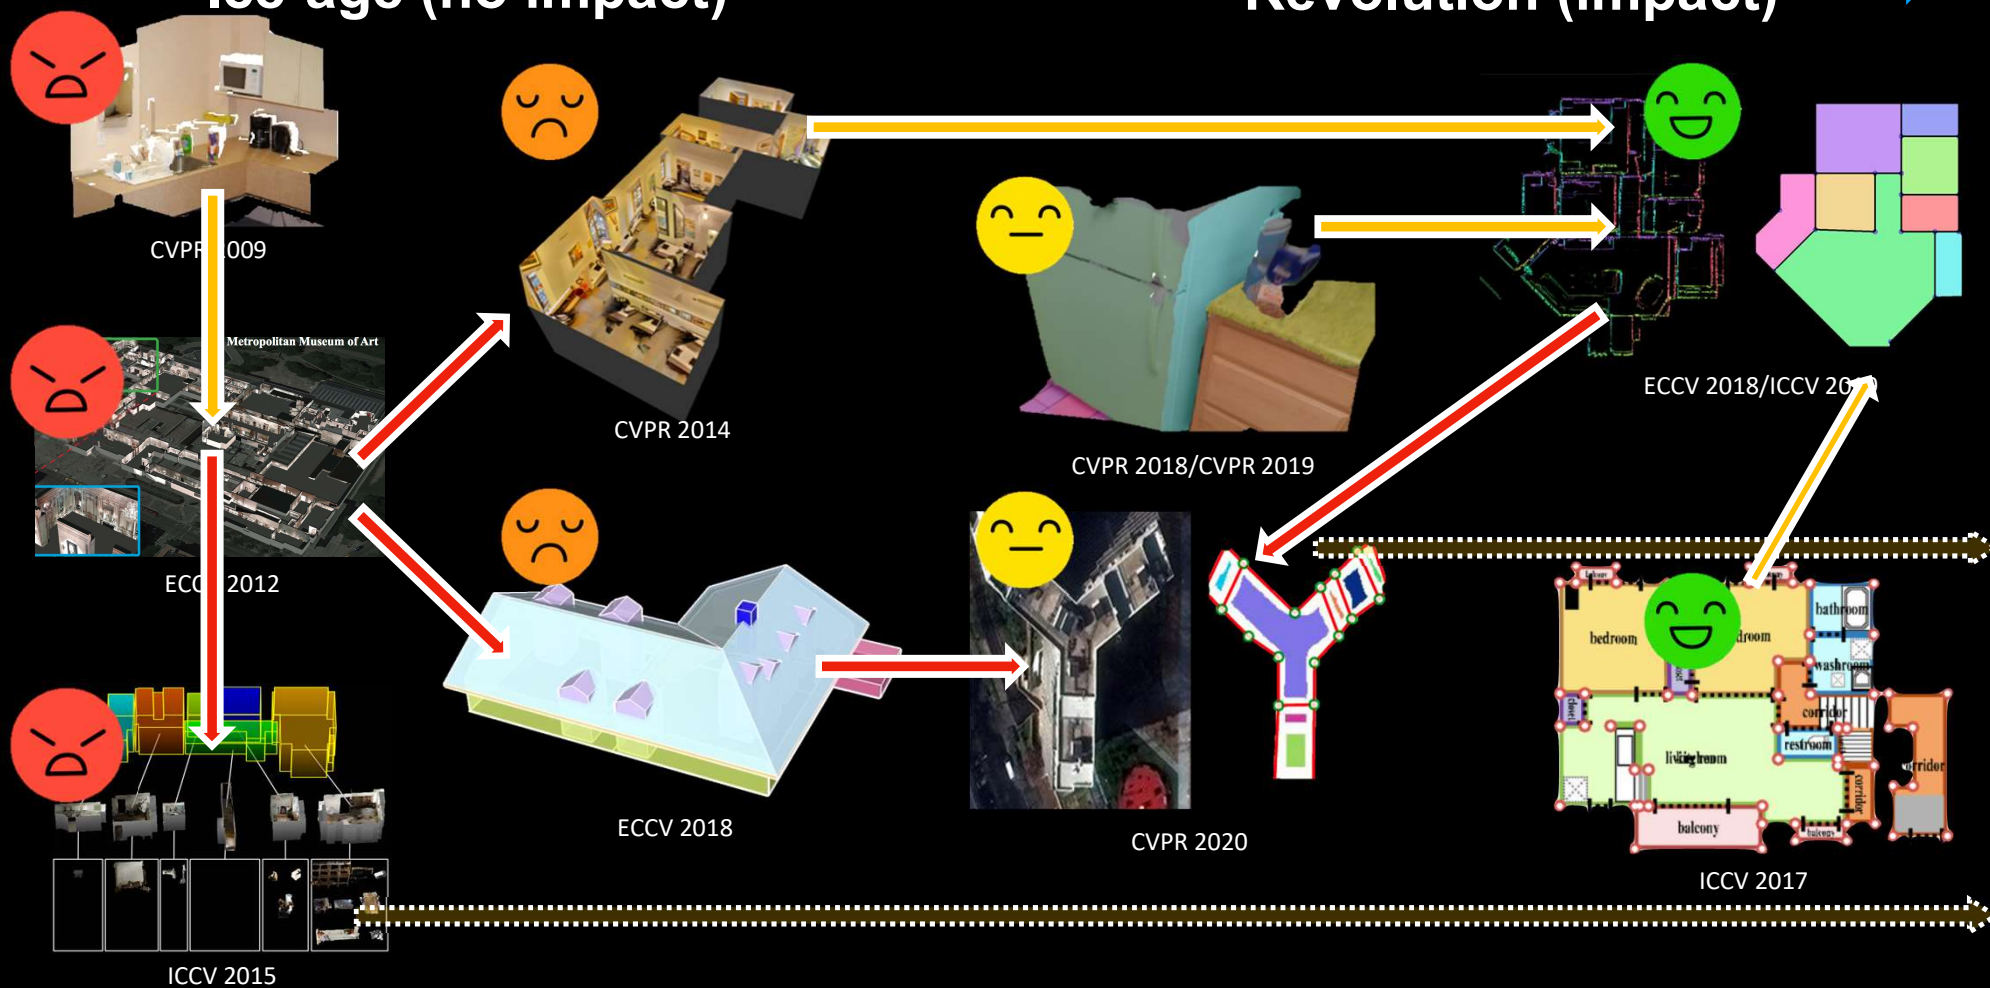
Manhattan World Stereo (387 citations)
[Furukawa, Curless, Seitz, and Szeliski]

Revolution (impact)



Ice-age (no impact)

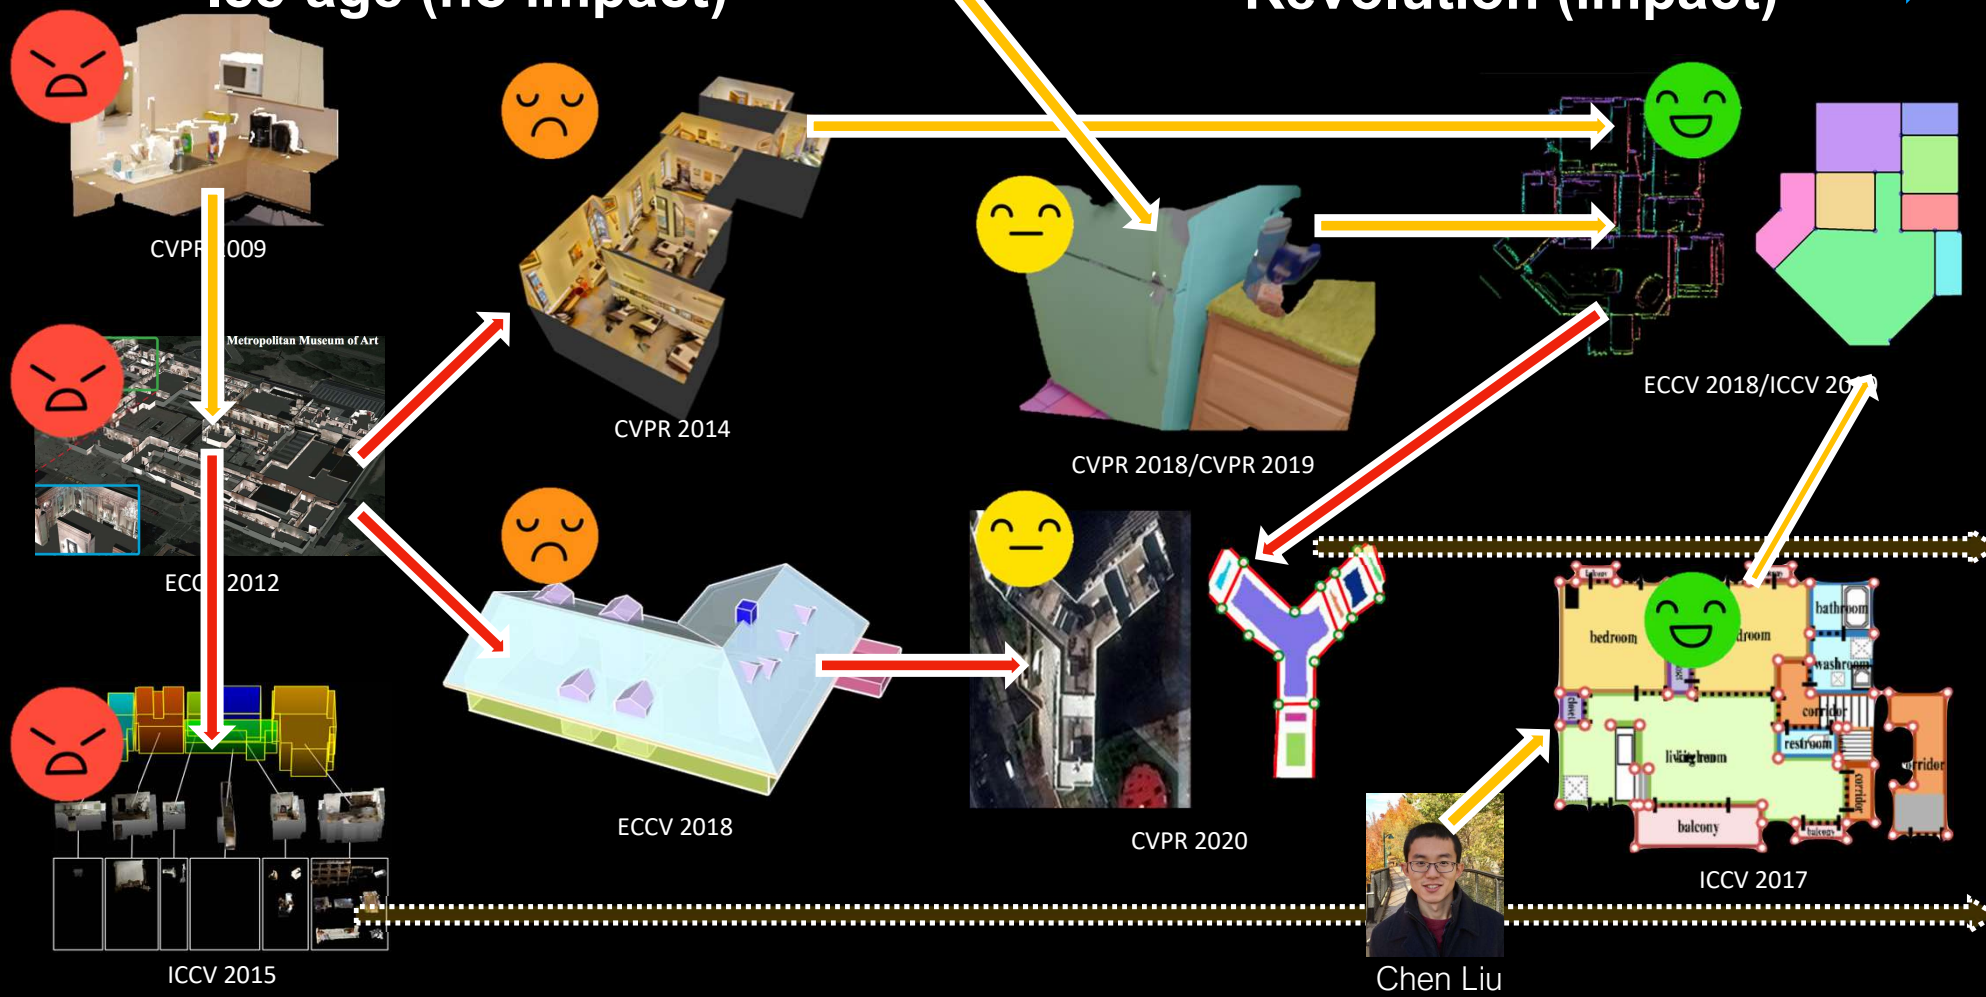
Revolution (impact)



Research idea-flow

Ice-age (no impact)

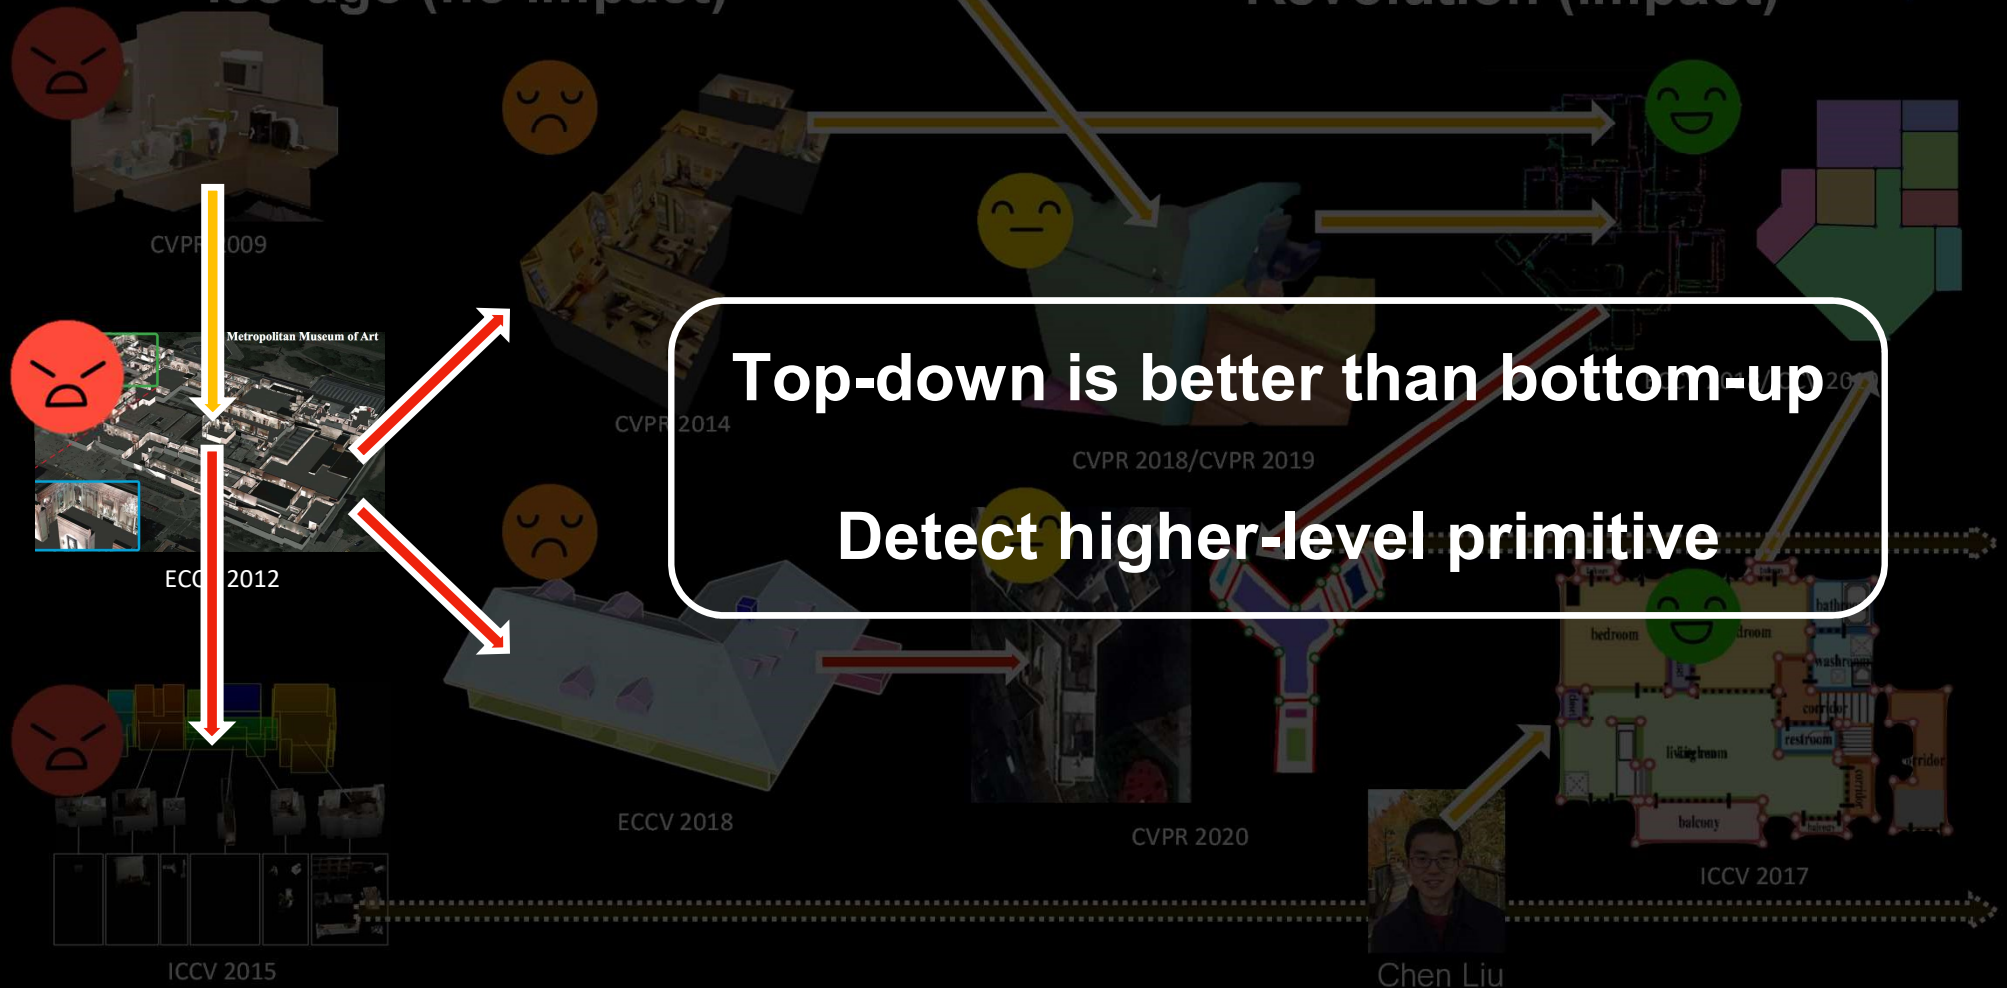
Revolution (impact)



Research idea-flow


Ice-age (no impact)

Revolution (impact)




Detect higher-level primitives

3D points




Floorplan




Bottom-up: Corner Detection

3D points





Floorplan





Bottom-up: Corner Detection

One missing corner





Bottom-up: Corner Detection

One missing corner






1. Cannot miss a single primitive.
2. Enumerate much more than necessary.
3. # of primitives to enumerate grow exponentially as the primitive DoF grows.

Top-down: Region Detection



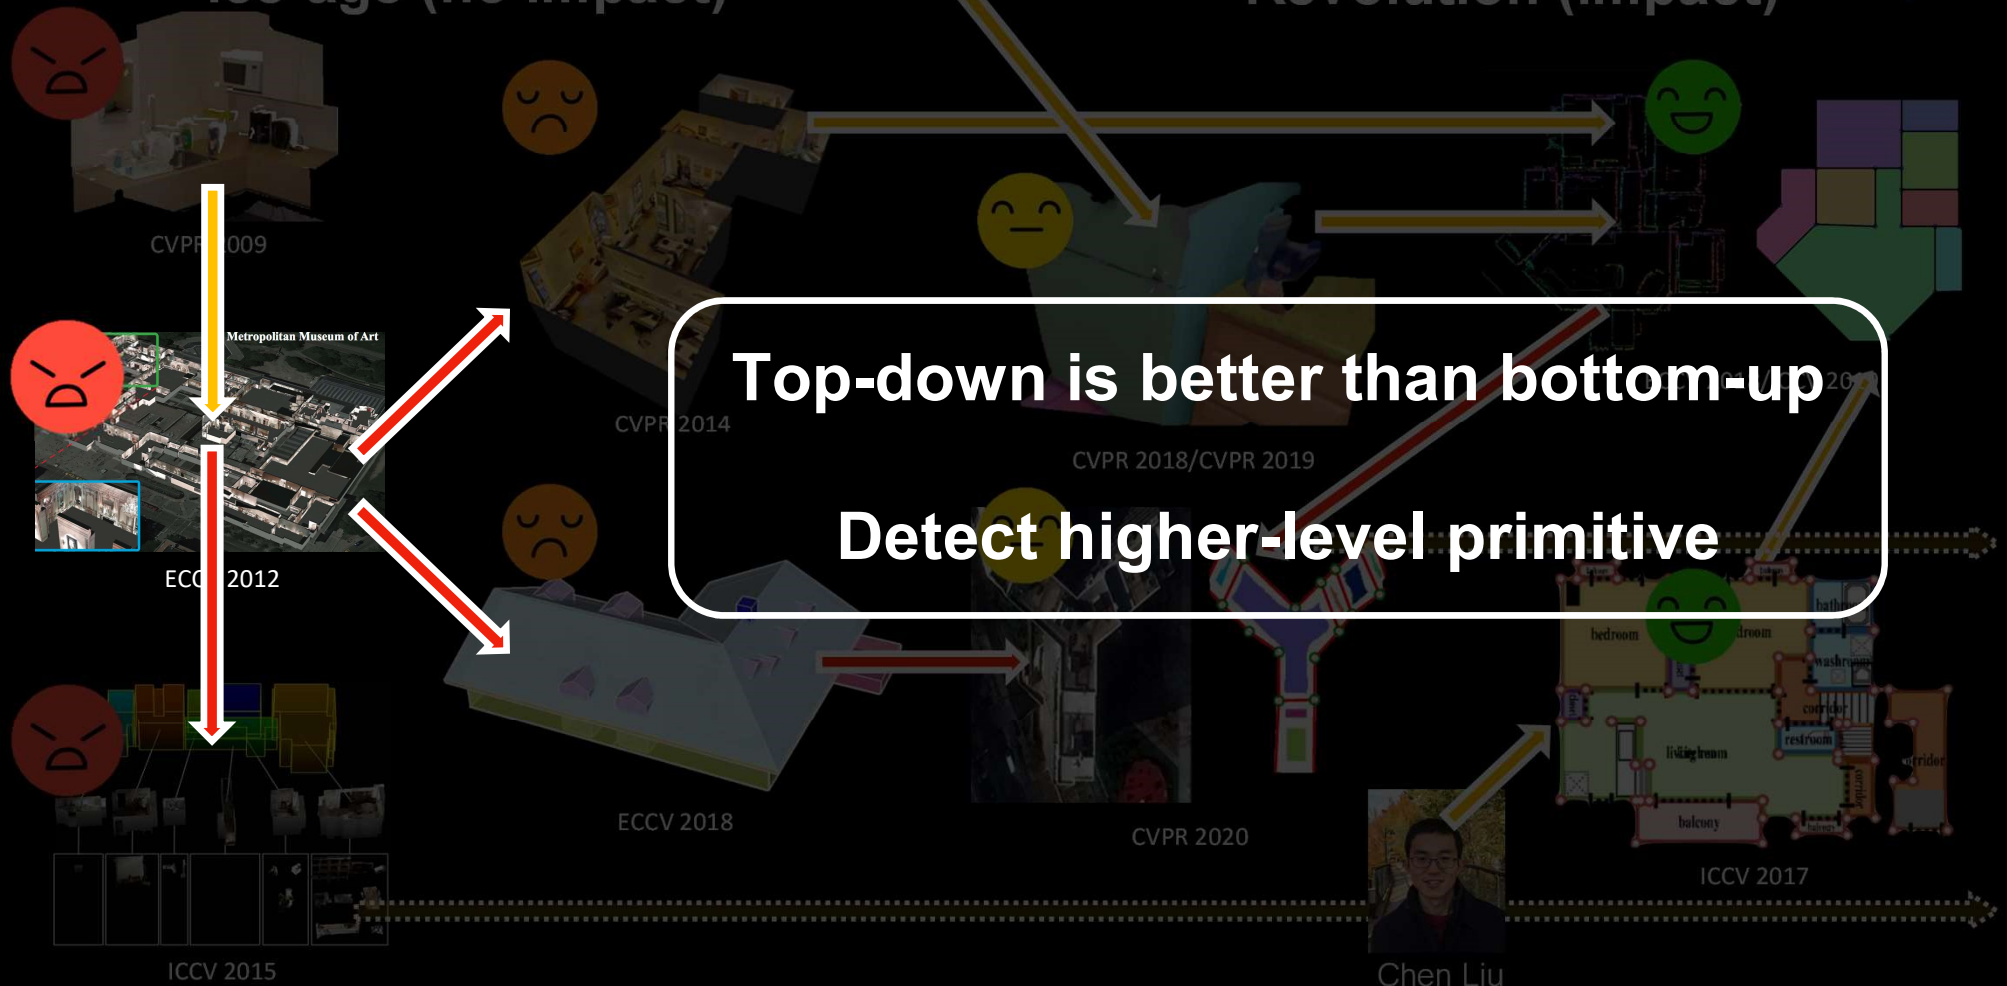
Top-down: Region Detection



Research idea-flow

Ice-age (no impact)

Revolution (impact)

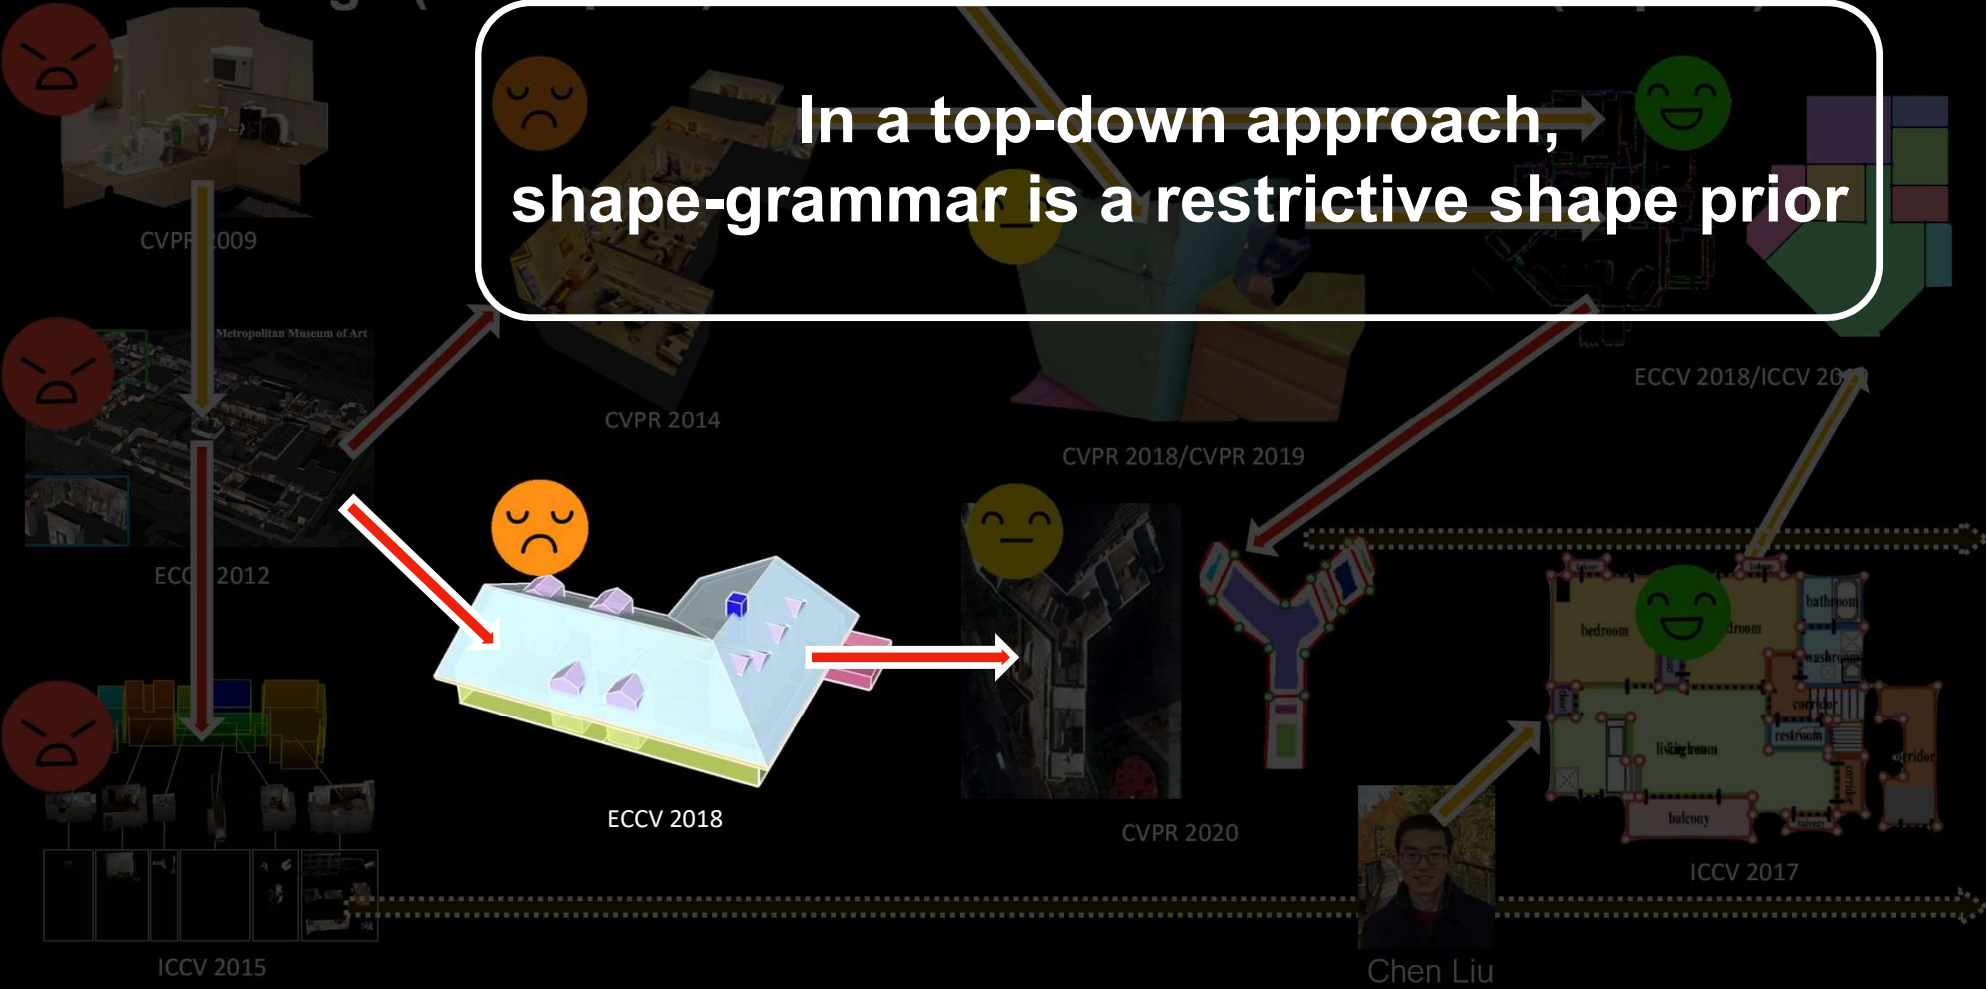


Research idea-flow

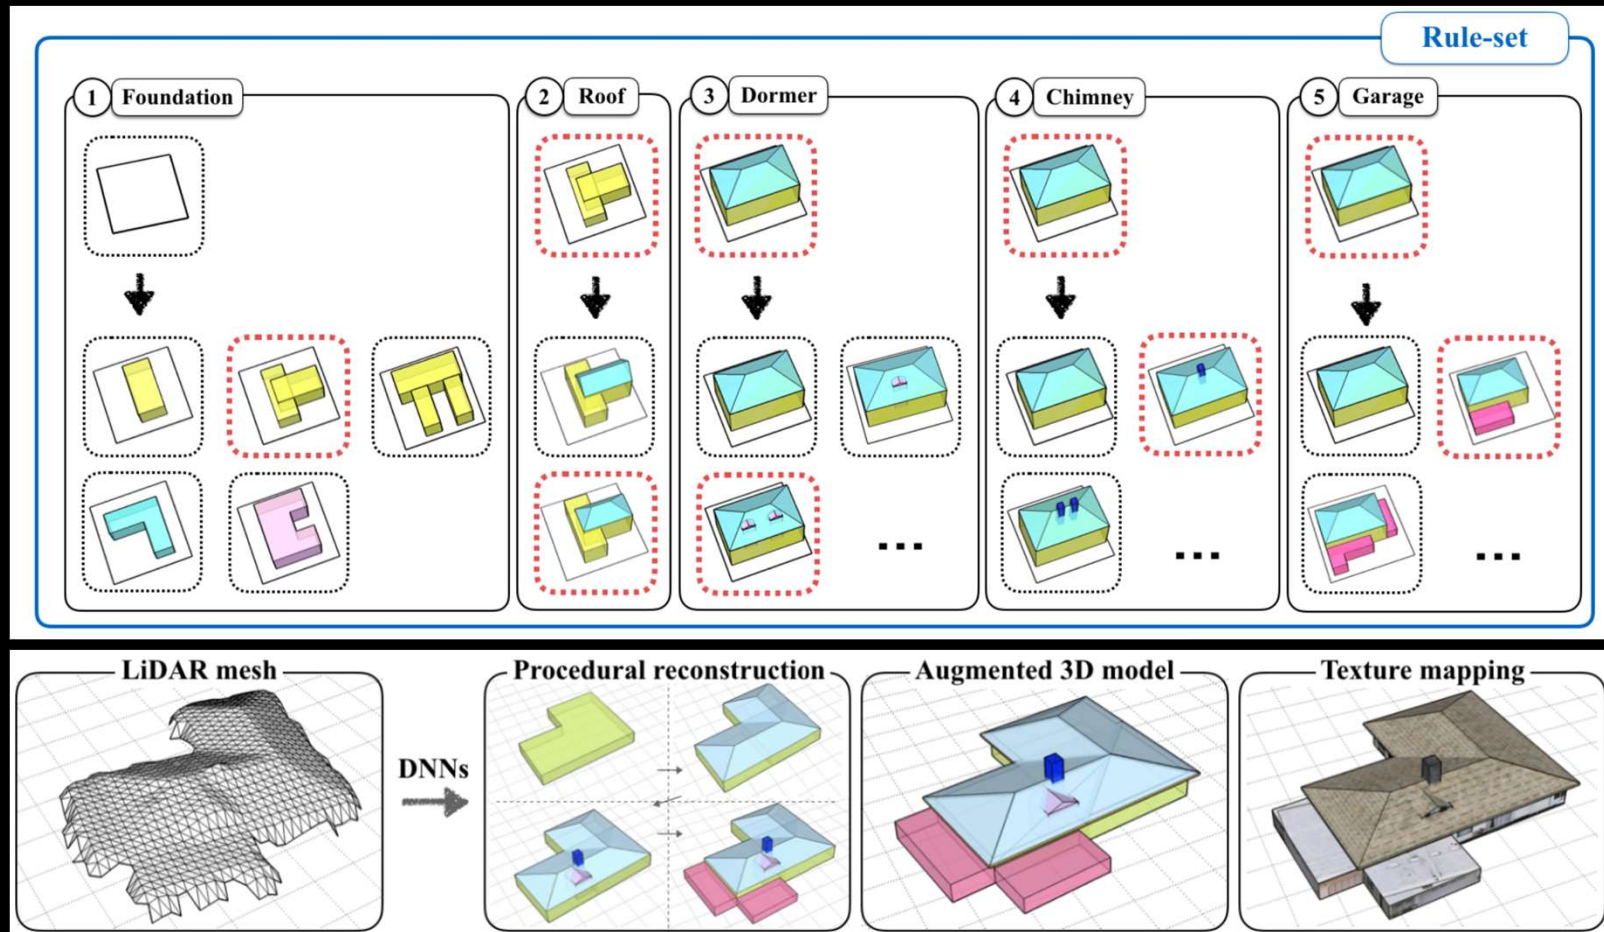
Ice-age (no impact)

Revolution (impact)

In a top-down approach, shape-grammar is a restrictive shape prior





Shape grammar as a shape prior





Grammar cannot model complex buildings



Manual design






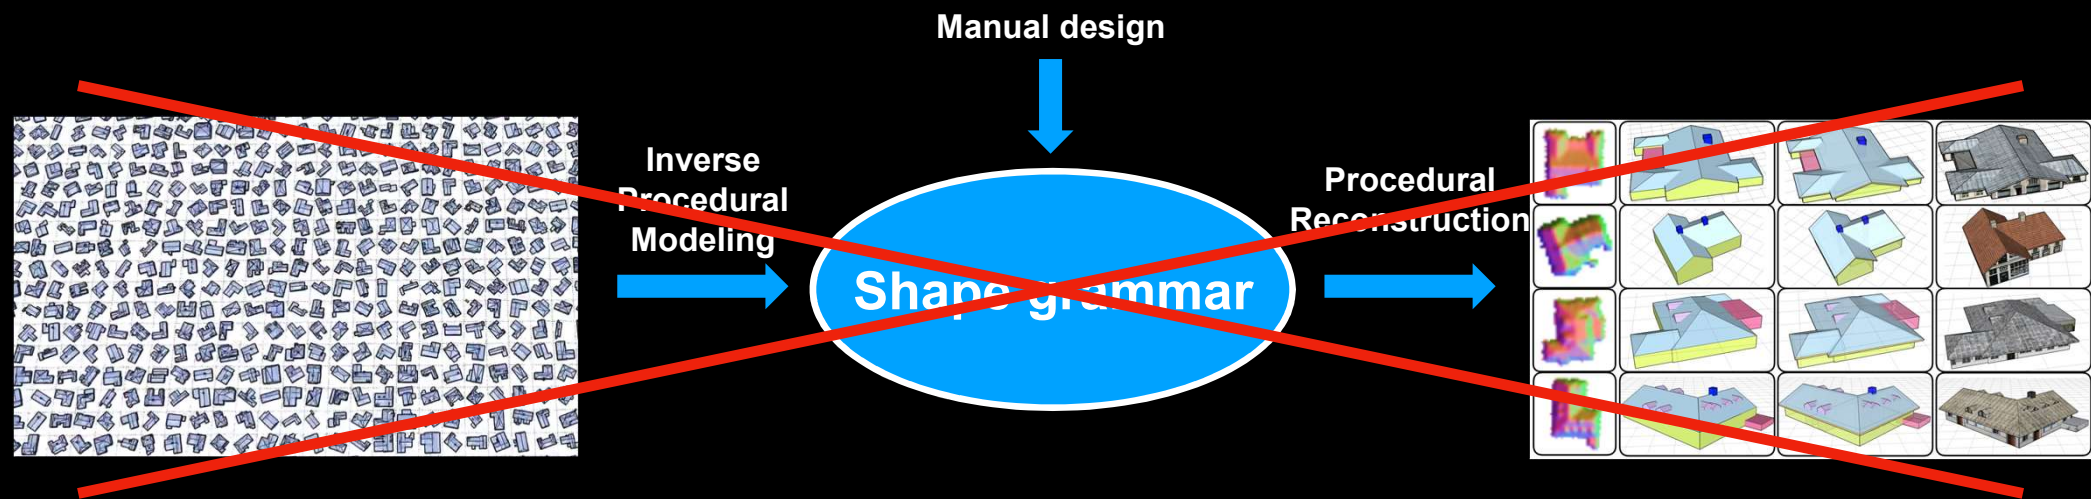
Inverse
Procedural
Modeling



Shape grammar

Procedural
Reconstruction





Research idea-flow

Ice-age (no impact)

Revolution (impact)

In a top-down approach, shape-grammar is a restrictive shape prior



CVPR 2009



Metropolitan Museum of Art

ECCV 2012



ICCV 2015



In a top-down approach, shape-grammar is a restrictive shape prior



ECCV 2018



CVPR 2020



CVPR 2018/CVPR 2019



ECCV 2018/ICCV 2019



ICCV 2017



Chen Liu

Research idea-flow

Ice-age (no impact)

Revolution (impact)

DNN enables the use of optimization instead of replacing

CVPR 2009

CVPR 2014



ECCV 2018/ICCV 2019

FCC 2012

ECCV 2018

CVPR 2020

ICCV 2017

ICCV 2015


Chen Liu



Outdoor Architecture Vectorization




Google Maps - satellite (2D)



CAD model / Architect drawing







**Optimization
(Integer programming)**

DNN enables optimization instead of replacing

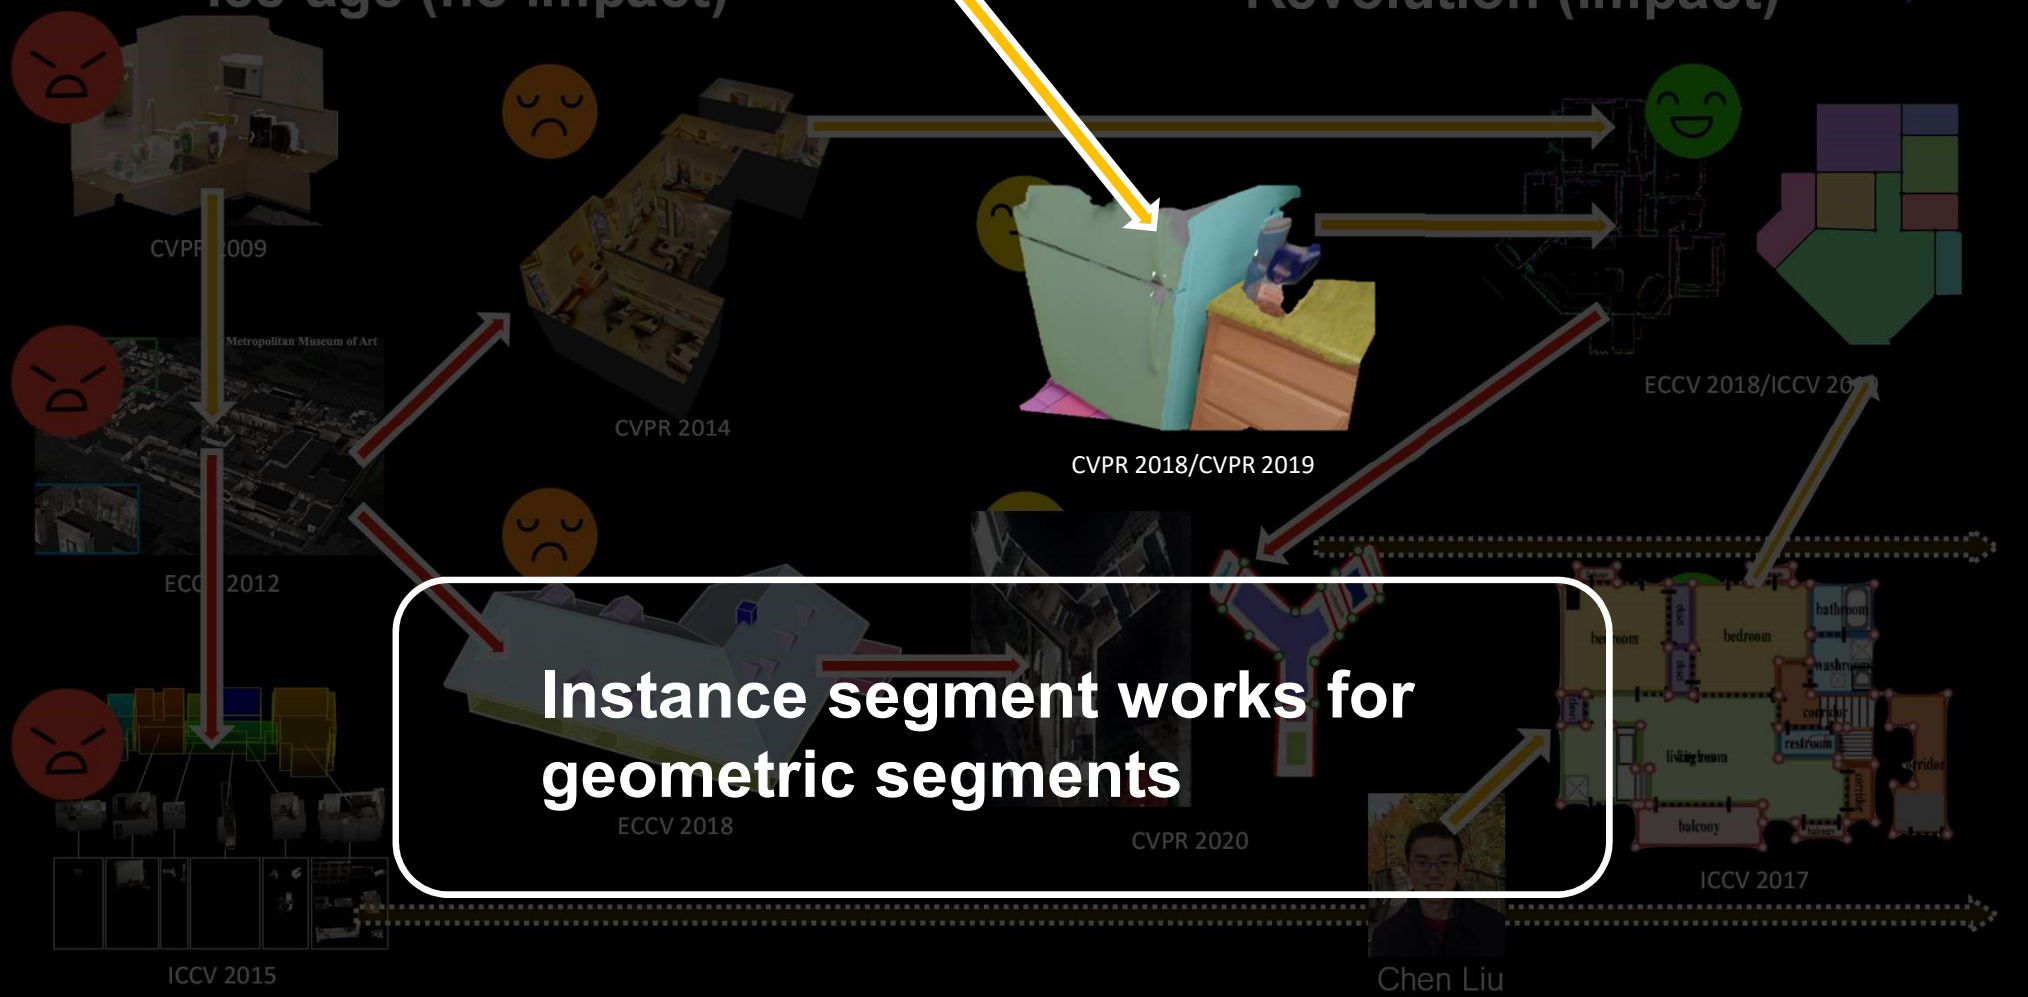
- DNN shrinks candidate pool
- DNN make energy landscape easy

	Scalability	Inference power
MRF	A few million vars	Pairwise Mostly submodular
IP	A few hundred vars	Higher-order Non-submodular

Research idea-flow

Ice-age (no impact)

Revolution (impact)




Research idea-flow


1. Neural topological inference
(avoid complex optimization)
2. Implicit grammar learning



CVPR 2018/CVPR 2019



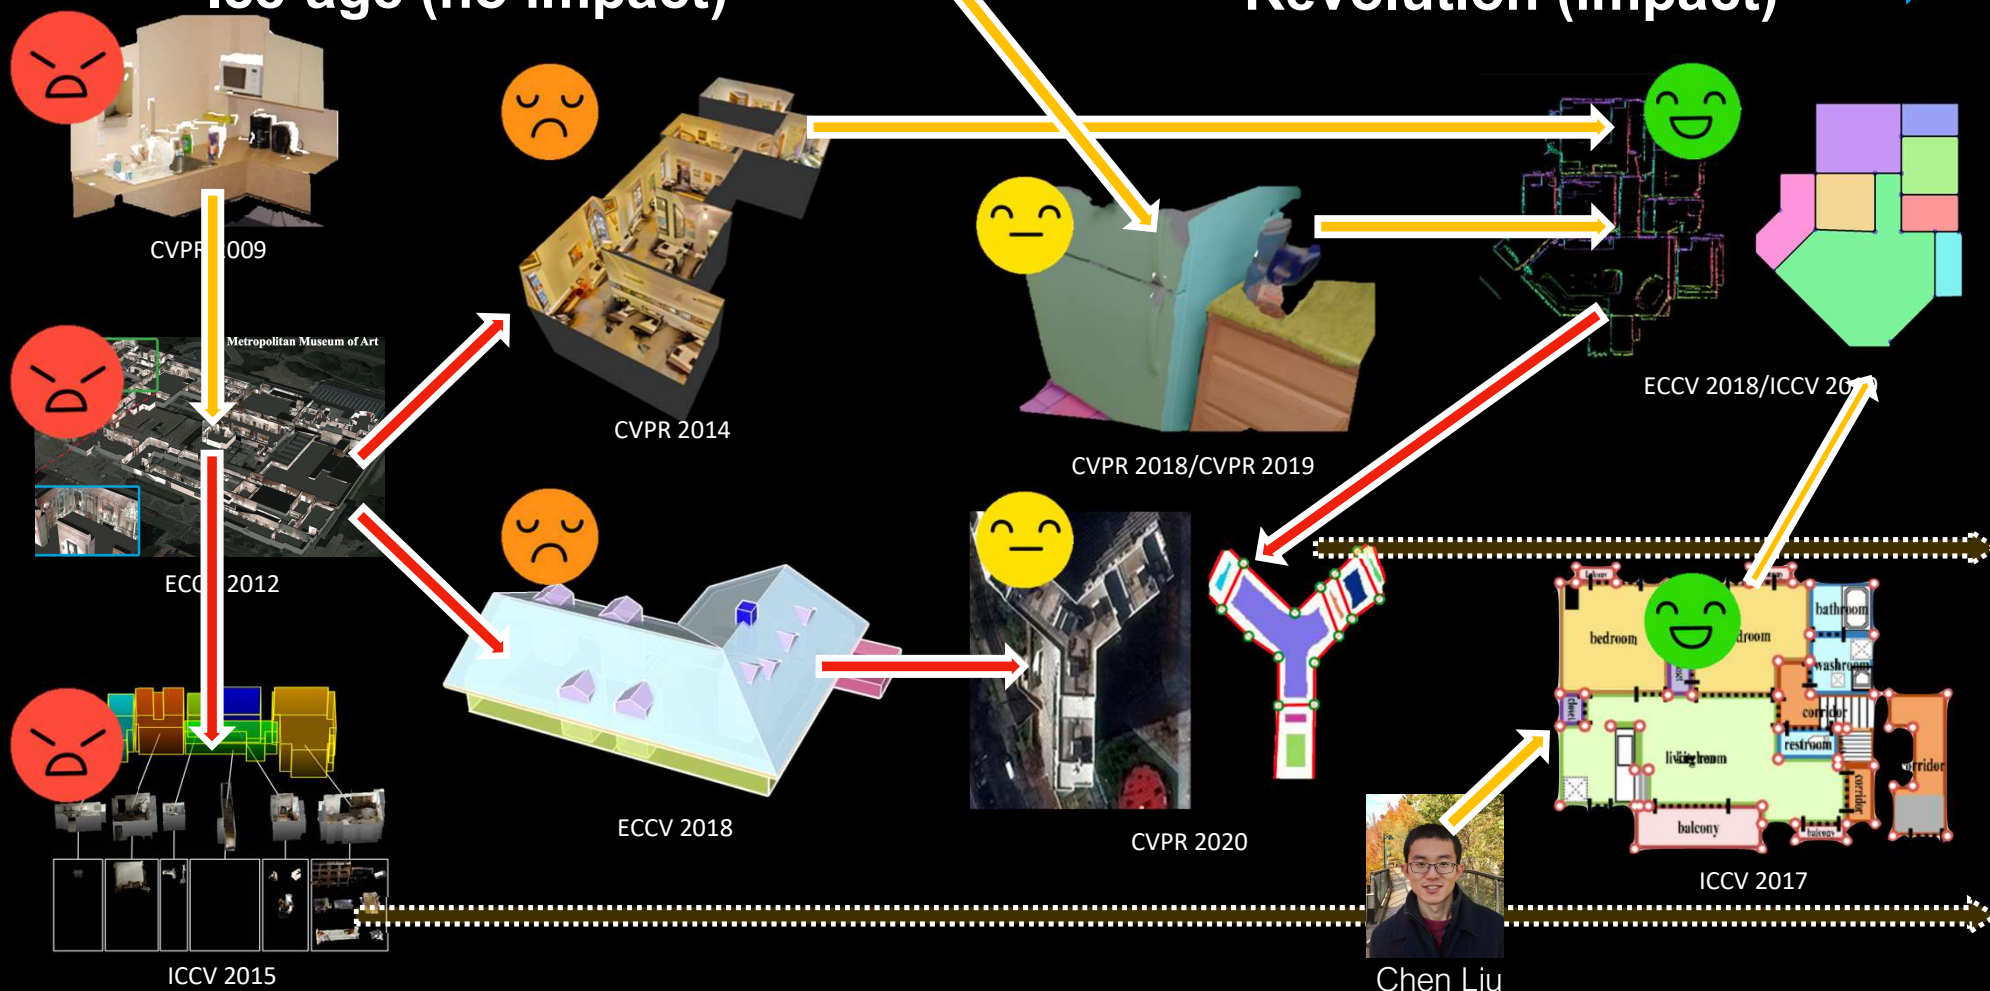
ECCV 2018/ICCV 2019



Research idea-flow

~~Ice-age (no impact)~~

Revolution (impact)



Ice-age (no impact)


Research idea-flow
Revolution (impact)



CVPR 2009




CVPR 2014



ECCV 2012



ICCV 2015





ECCV 2018



CVPR 2018/CVPR 2019



CVPR 2020



ECCV 2018/ICCV 2019



ICCV 2017

Ice-age (no impact)


Research idea-flow
Revolution (impact)



CVPR 2009




CVPR 2014



ECCV 2012



ICCV 2015





ECCV 2018



CVPR 2018/CVPR 2019



CVPR 2020



ECCV 2018/ICCV 2019



ICCV 2017

Garbage ideas lead to impact



Find your masterpiece at CVPR