ECCV 2018 (Munich)
Trip report

Kobus Barnard

Interdisciplinary Visual Intelligence Lab
http://ivilab.org

DARPA

NSF

Trip reporting

- Expect to give a presentation if you attend a conference, workshop, extended tutorial, etc.

Conferences in CS and AI in particular

- Top conferences are very competitive
- Top vision conferences are: ICCV, CVPR, ECCV
- Top machine learning conferences are NIPS and ICML
- Additional AI conferences are AAAI, ICHCAI, and ICLR

- Posters versus orals
  - Acceptance is usually done first, then orals are chosen
  - Orals give you some bragging rights. They are some of the better papers, but also are chosen to be of more general interest, and fit into a larger balanced schedule.

ECCV 2018 (the explosion continues)

- My first big conference in a while

- We had a paper in ECCV 2016, but I did not go because I waited too long to decide to go (Kyle went).

- ECCV 2018 was capped at 3200 (and was sold out for the last month)

- CVPR 2017 was of the order of 7000

- NIPS sold out in 11 minutes!

Munich

What do you notice walking about?

- People smoke —a lot!
  - Lots of outdoor seating at restaurants
- Commuting by bike (without helmets) is very popular
  - Pedestrians are prey

Venue: Gasteig, which is relatively famous philharmonic hall


Gasteig


Questions?
Poster No. P-1B-02
Code & Data: github.com/12dmodel/deep_motion_mag

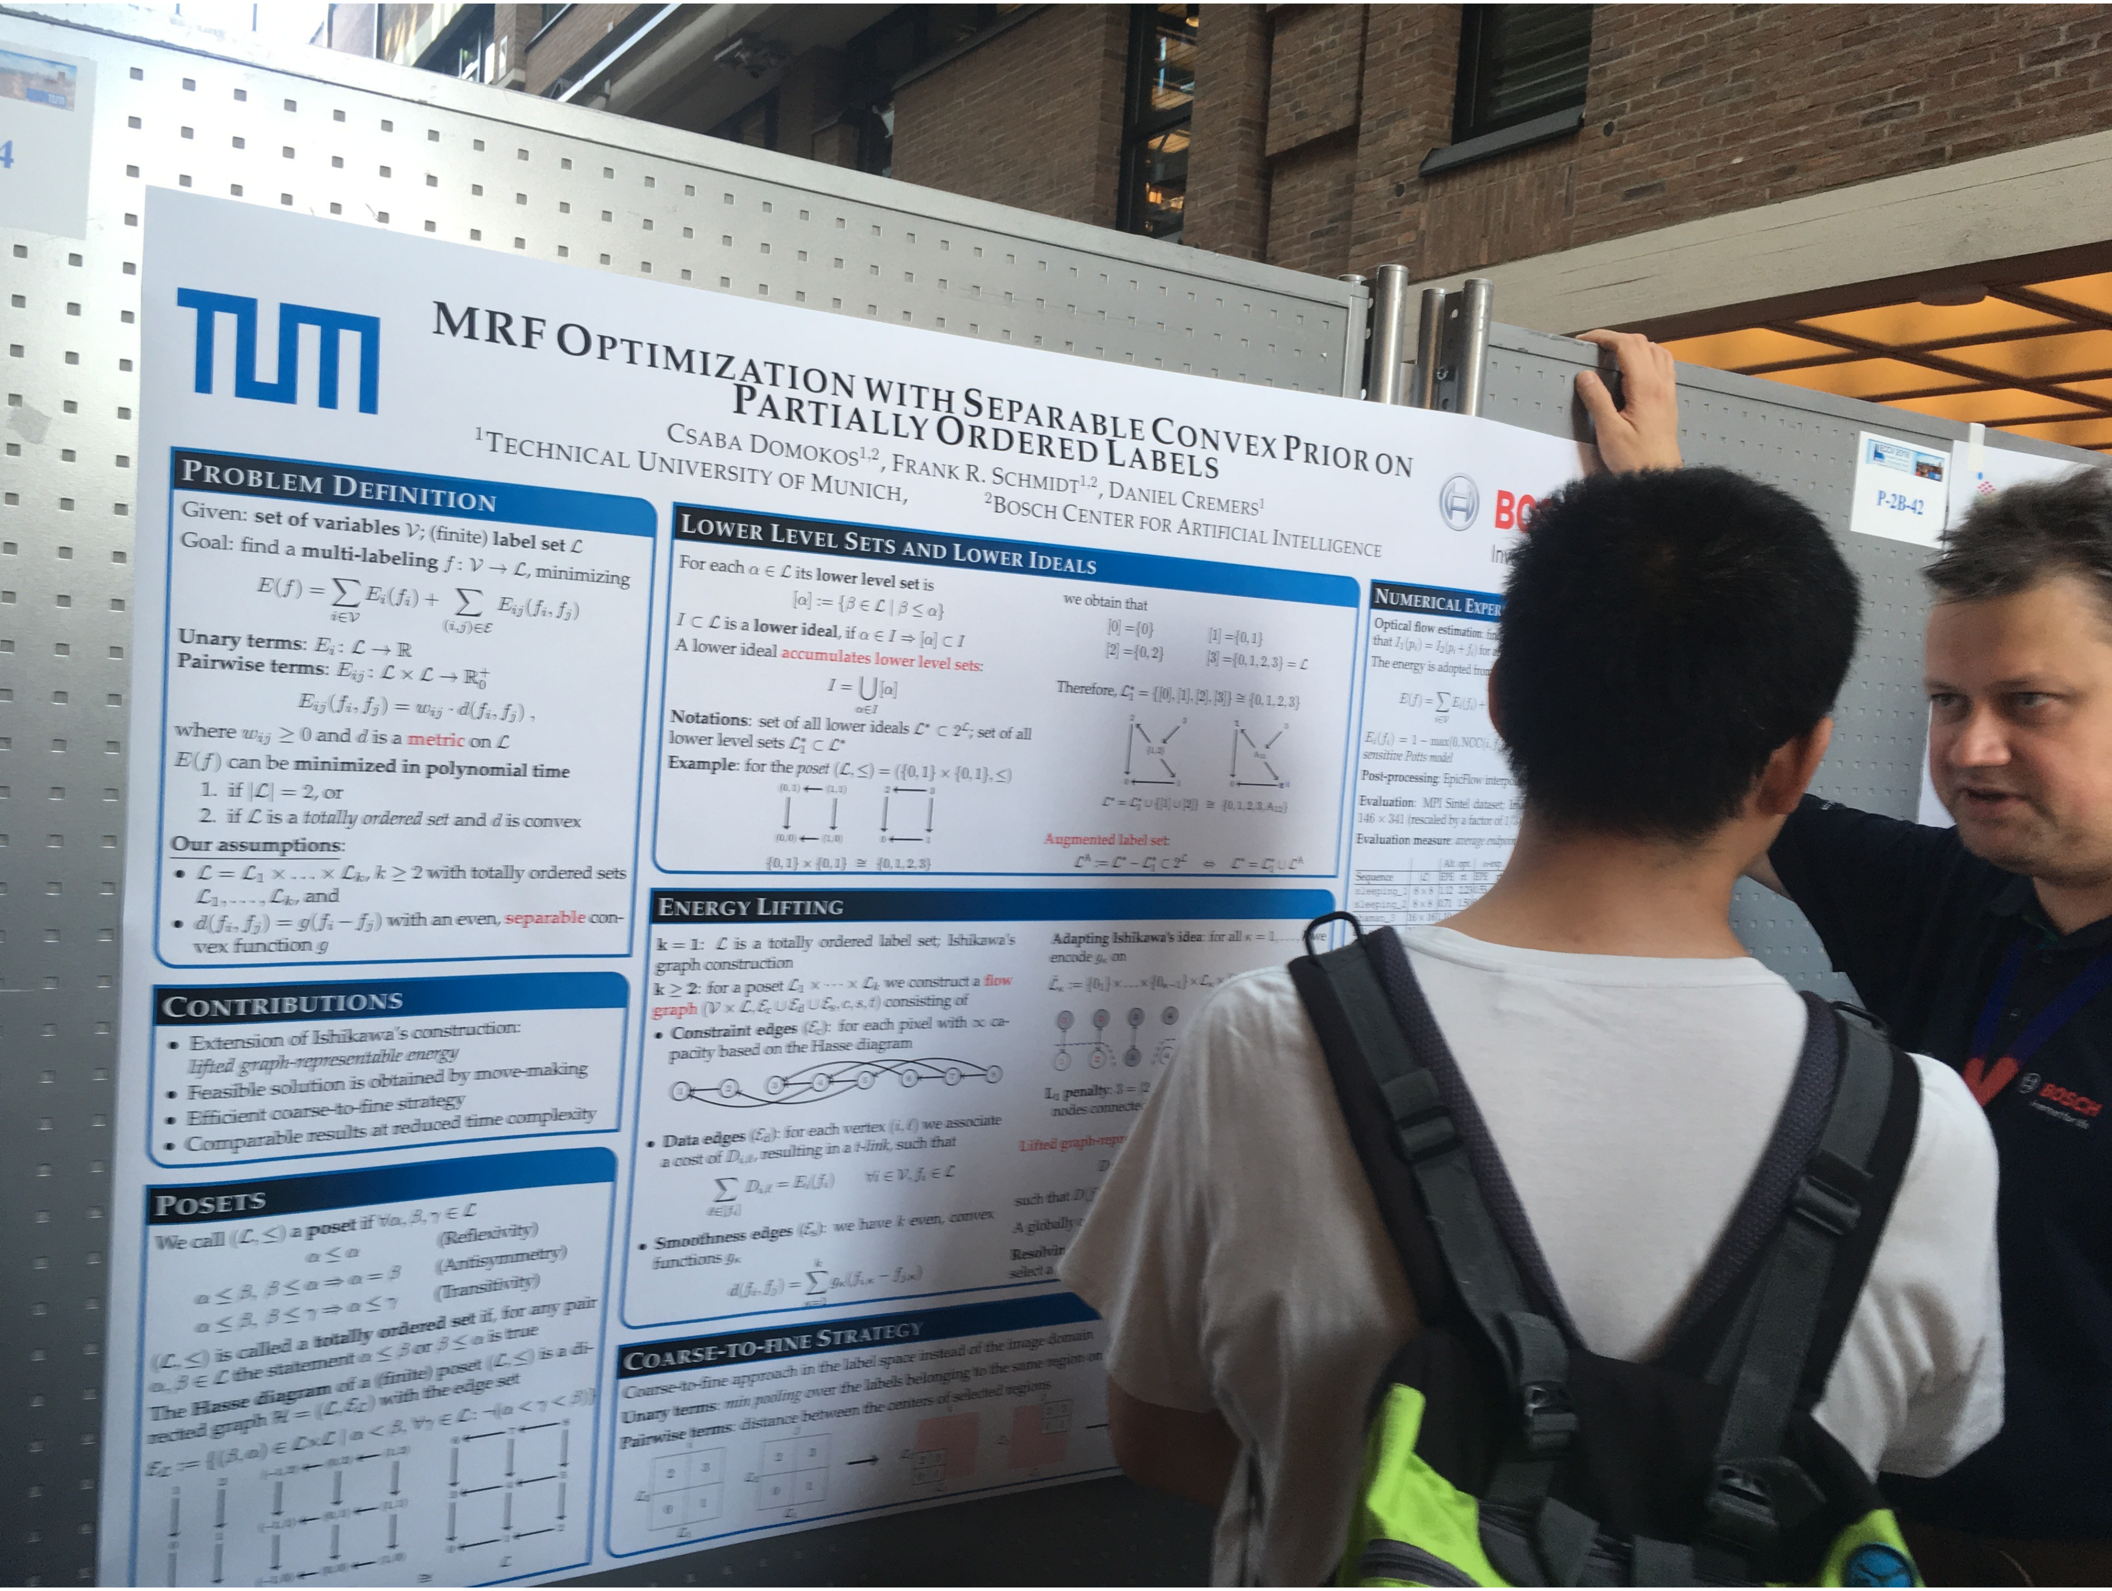
MRF Optimization with Separable Convex Prior on Partially Ordered Labels
Csaba Domokos1,2, Frank R. Schmidt1,2, Daniel Cremers1
1Technical University of Munich, 2Bosch Center for Artificial Intelligence
Problem Definition
Given: set of variables $\mathcal{V}$; (finite) label set $\mathcal{L}$
Goal: find a multi-labeling $f: \mathcal{V} \to \mathcal{L}$, minimizing
$E(f) = \sum_{i\in\mathcal{V}} E_i(f_i) + \sum_{(i,j)\in\mathcal{E}} E_{ij}(f_i, f_j)$
Unary terms: $E_i: \mathcal{L} \to \mathbb{R}$
Pairwise terms: $E_{ij}: \mathcal{L} \times \mathcal{L} \to \mathbb{R}_0^+$
$E_{ij}(f_i, f_j) = w_{ij} \cdot d(f_i, f_j)$,
where $w_{ij} \geq 0$ and d is a metric on $\mathcal{L}$
$E(f)$ can be minimized in polynomial time
1. if $|\mathcal{L}| = 2$, or
2. if $\mathcal{L}$ is a totally ordered set and d is convex
Our assumptions:
• $\mathcal{L} = \mathcal{L}_1 \times \ldots \times \mathcal{L}_k$, $k \geq 2$ with totally ordered sets $\mathcal{L}_1, \ldots, \mathcal{L}_k$, and
• $d(f_i, f_j) = g(f_i - f_j)$ with an even, separable convex function g
Contributions
• Extension of Ishikawa's construction: lifted graph-representable energy
• Feasible solution is obtained by move-making
• Efficient coarse-to-fine strategy
• Comparable results at reduced time complexity
Posets
We call $(\mathcal{L}, \leq)$ a poset if $\forall \alpha, \beta, \gamma \in \mathcal{L}$
$\alpha \leq \alpha$ (Reflexivity)
$\alpha \leq \beta, \beta \leq \alpha \Rightarrow \alpha = \beta$ (Antisymmetry)
$\alpha \leq \beta, \beta \leq \gamma \Rightarrow \alpha \leq \gamma$ (Transitivity)
$(\mathcal{L}, \leq)$ is called a totally ordered set if, for any pair $\alpha, \beta \in \mathcal{L}$ the statement $\alpha \leq \beta$ or $\beta \leq \alpha$ is true
The Hasse diagram of a (finite) poset $(\mathcal{L}, \leq)$ is a directed graph $\mathcal{H} = (\mathcal{L}, \mathcal{E}_\mathcal{L})$ with the edge set
Lower Level Sets and Lower Ideals
For each $\alpha \in \mathcal{L}$ its lower level set is
$[\alpha] := \{\beta \in \mathcal{L} \mid \beta \leq \alpha\}$
$I \subset \mathcal{L}$ is a lower ideal, if $\alpha \in I \Rightarrow [\alpha] \subset I$
A lower ideal accumulates lower level sets:
$I = \bigcup_{\alpha\in I} [\alpha]$
Notations: set of all lower ideals $\mathcal{L}^* \subset 2^\mathcal{L}$; set of all lower level sets $\mathcal{L}_1^* \subset \mathcal{L}^*$
Example: for the poset $(\mathcal{L}, \leq) = (\{0,1\} \times \{0,1\}, \leq)$
$\{0,1\} \times \{0,1\} \cong \{0,1,2,3\}$
we obtain that
$[0] = \{0\}$ $[1] = \{0,1\}$
$[2] = \{0,2\}$ $[3] = \{0,1,2,3\} = \mathcal{L}$
Therefore, $\mathcal{L}_1^* = \{[0],[1],[2],[3]\} \cong \{0,1,2,3\}$
Augmented label set:
Energy Lifting
k = 1: $\mathcal{L}$ is a totally ordered label set; Ishikawa's graph construction
k ≥ 2: for a poset $\mathcal{L}_1 \times \cdots \times \mathcal{L}_k$ we construct a flow graph consisting of
• Constraint edges: for each pixel with ∞ capacity based on the Hasse diagram
• Data edges: for each vertex we associate a cost of $D_{i,\ell}$, resulting in a t-link, such that
• Smoothness edges: we have k even, convex functions g_κ
Adapting Ishikawa's idea:
Lifted graph-repr
Coarse-to-fine Strategy
Coarse-to-fine approach in the label space instead of the image domain
Unary terms: min pooling over the labels belonging to the same region
Pairwise terms: distance between the centers of selected regions
Numerical Exper
Optical flow estimation
The energy is adopted from
Post-processing: EpicFlow
Evaluation: MPI Sintel dataset
Evaluation measure:
BOSCH
P-2B-42

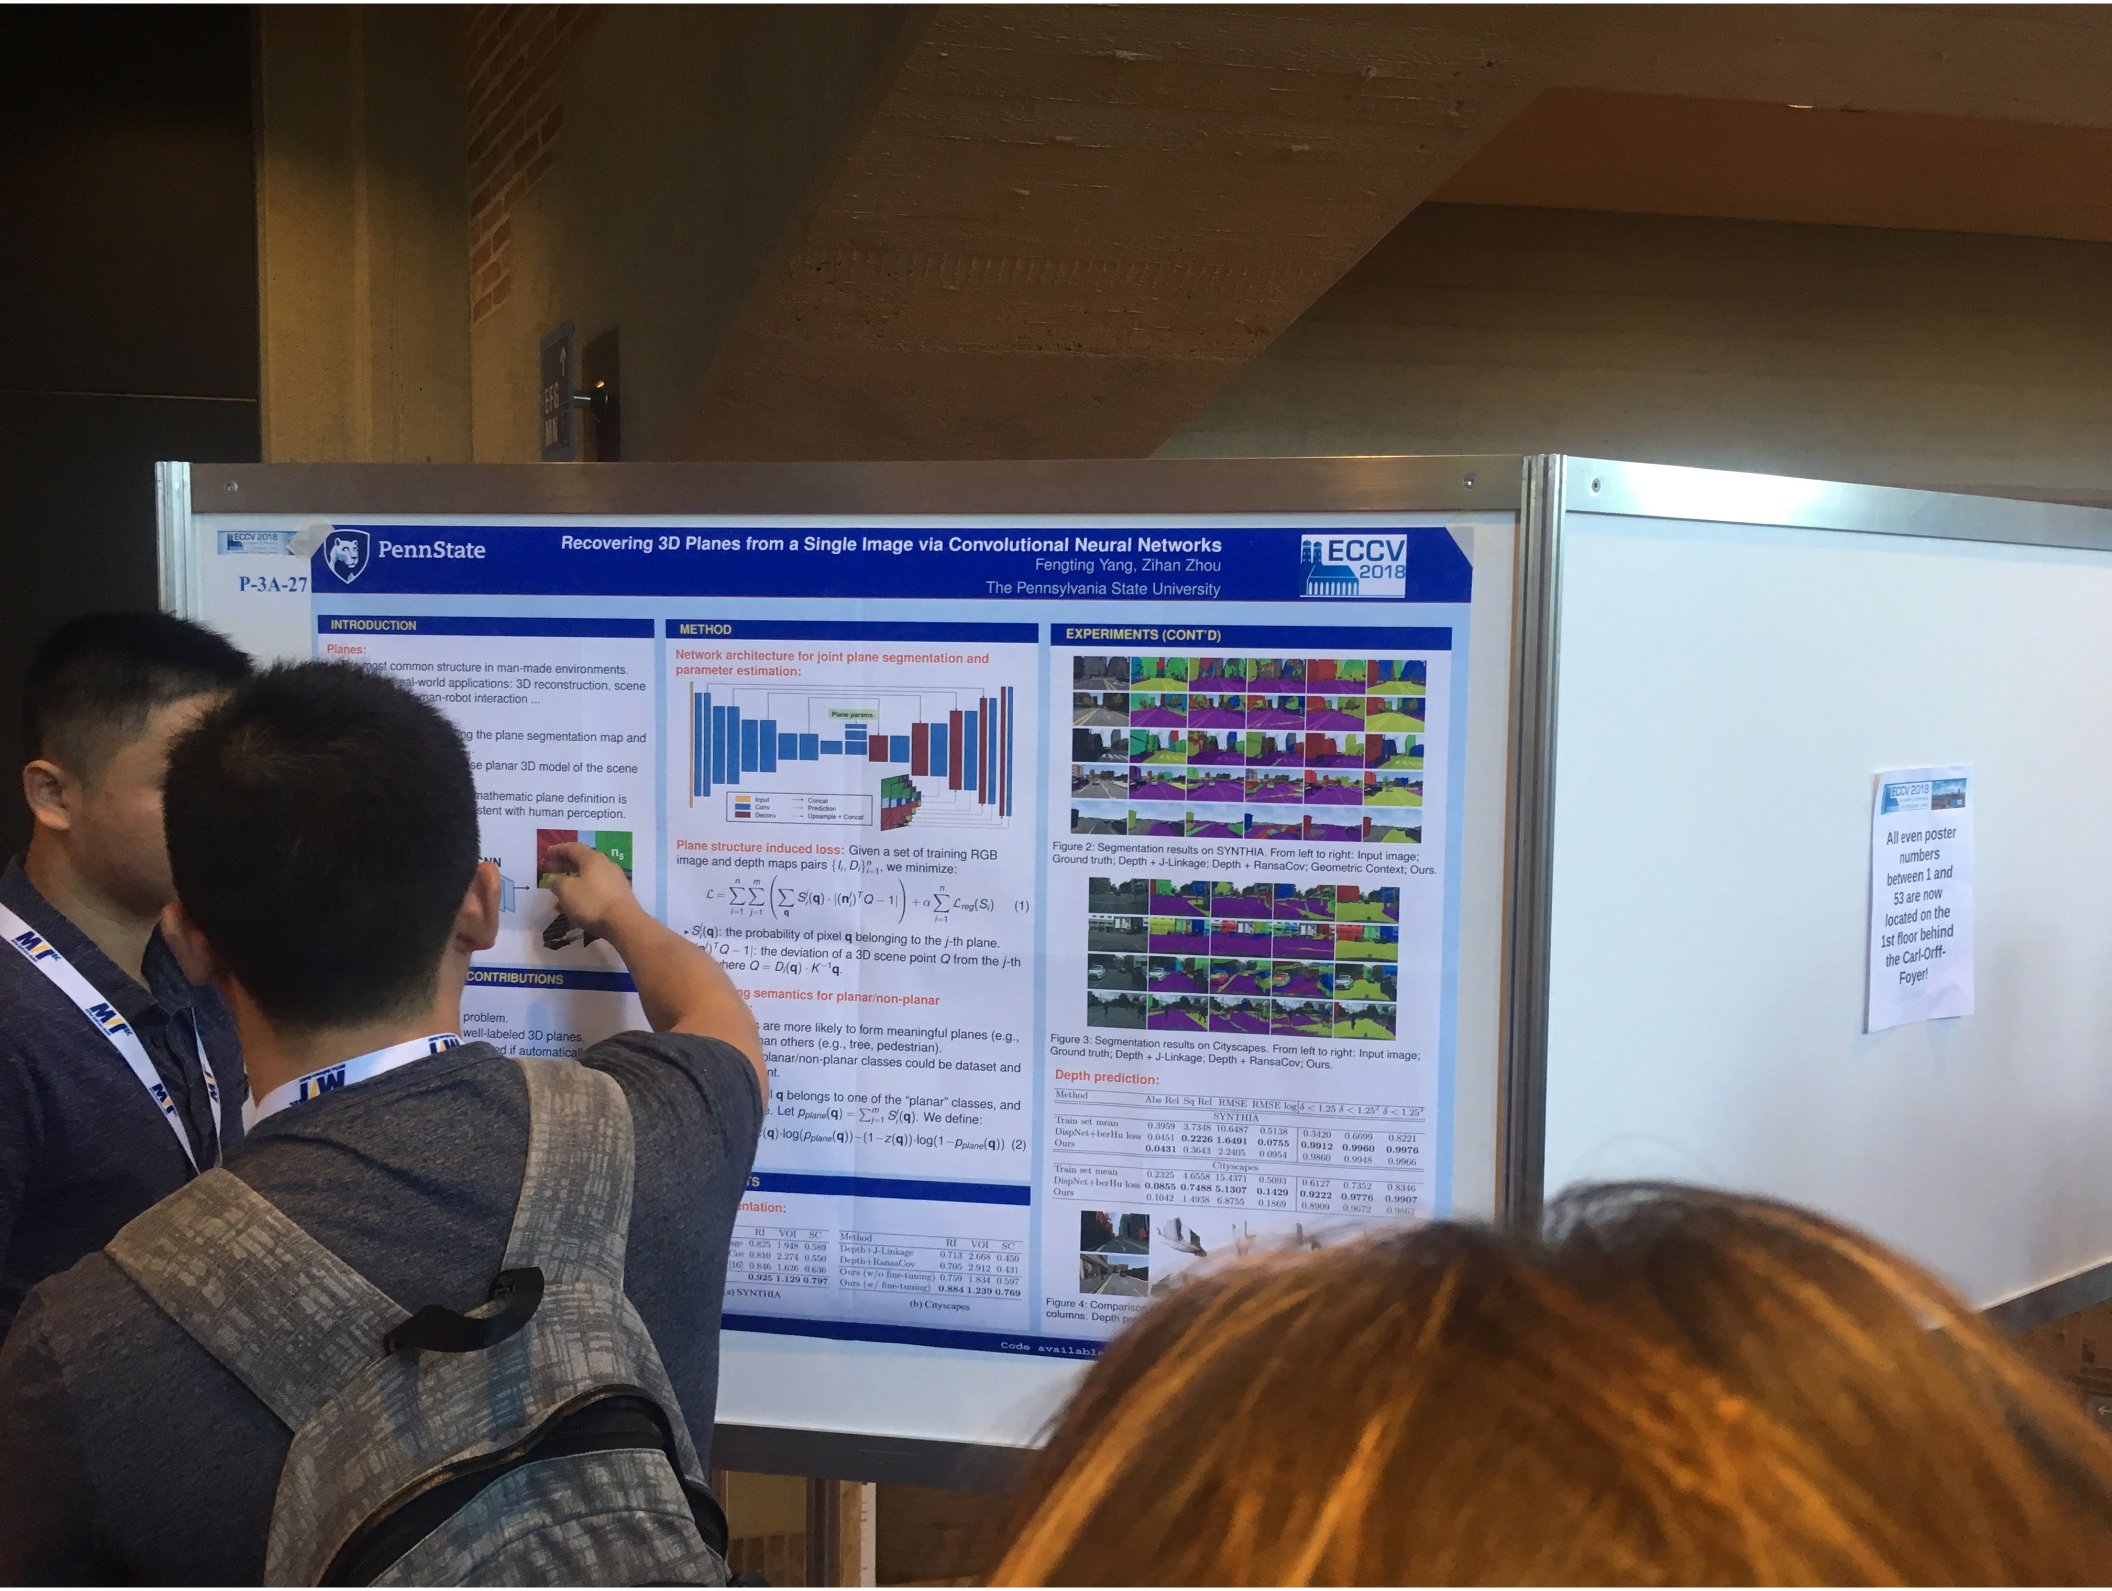

P-3A-27
PennState
Recovering 3D Planes from a Single Image via Convolutional Neural Networks
Fengting Yang, Zihan Zhou
The Pennsylvania State University
ECCV 2018
INTRODUCTION
Planes:
most common structure in man-made environments.
real-world applications: 3D reconstruction, scene
man-robot interaction ...
ng the plane segmentation map and
se planar 3D model of the scene
mathematical plane definition is
stent with human perception.
CONTRIBUTIONS
problem.
well-labeled 3D planes.
METHOD
Network architecture for joint plane segmentation and parameter estimation:
Plane structure induced loss: Given a set of training RGB image and depth maps pairs, we minimize:
$S_j(\mathbf{q})$: the probability of pixel $\mathbf{q}$ belonging to the j-th plane.
the deviation of a 3D scene point Q from the j-th
here $Q = D_i(\mathbf{q}) \cdot K^{-1}\mathbf{q}$.
g semantics for planar/non-planar
are more likely to form meaningful planes (e.g.,
an others (e.g., tree, pedestrian).
planar/non-planar classes could be dataset and
belongs to one of the "planar" classes, and
We define:
entation:
(a) SYNTHIA
(b) Cityscapes
EXPERIMENTS (CONT'D)
Figure 2: Segmentation results on SYNTHIA. From left to right: Input image; Ground truth; Depth + J-Linkage; Depth + RansaCov; Geometric Context; Ours.
Figure 3: Segmentation results on Cityscapes. From left to right: Input image; Ground truth; Depth + J-Linkage; Depth + RansaCov; Ours.
Depth prediction:
SYNTHIA
Cityscapes
Train set mean
DispNet+berHu loss
Ours
Figure 4: Compariso
Code available
All even poster numbers between 1 and 53 are now located on the 1st floor behind the Carl-Orff-Foyer!

What is HOT?

Everything!

Social Trends

- Industry is consuming academia!
  - New affiliations that I just learned about
    - Sven Dickinson (doing a 2 year leave at Samsung from UT), Raquel Urtasun (now at Uber), Noah Snavely (50% at Google), …
- Self-driving cars
  - Bottomless well of money
  - Self-flying personal helicopters and delivery drones
- Industry is starting to dominate computer vision and ML conferences
  - Lots of papers out of the big labs (Google, FAIR, …)
  - Huge number of nice “booths”

Deep learning for the sake of humanity.
Deep learning for the sake of humanity.
Vision
Tracking
Autonomy
Explainable AI
Lifelong Learning
Context Awareness
Panasonic β
SAP

Industry topics that come to mind

- Self-driving cars
- Sports
- Augmented reality
- Mapping
- Data annotation!


ACTIVITY RECOGNITION
AND AUTOMATED
BROADCAST OVERLAY
AI-POWERED
SPORTS
ANALYTICS
SPORTLOGiQ
iQ
LEARN MORE AT SPORTLOGIQ.COM

EMBEDDED
COMPUTER VISION
AND
DEEP LEARNING
TECHNOLOGY
Mighty Ai
Training Data as a Service

Clova
NAVER | LINE
A HUMAN-ORIENTED ARTIFICIAL INTELLIGENCE PLATFORM
JOINTLY DEVELOPED BY NAVER AND LINE

lyft
Level
5
Self-Driving
Division


Map data at scale
from images
contributed by anyone
TAXI

You Cannot Serve Two Masters: The Harms of Dual Affiliation

Ben Recht, David A. Forsyth, and Alexei Efros • Aug 9, 2018

Facebook would like to have computer science faculty in AI committed to work 80% of their time in industrial jobs and 20% of their time at their university. They call this scheme "co-employment" or "dual affiliation." This model assumes people can slice their time and attention like a computer, but people can't do this. Universities and companies are communities, each with their particular missions and values. The values of these communities are often at odds, and researchers must choose where their main commitment lies. By committing researchers to a particular company's interests, this new model of employment will harm our colleagues, our discipline, and everyone's future. Like many harms, it comes with benefits for some. But the harm in this proposal outweighs the benefits. If industry wants to support and grow academic computer science, there are much better ways to achieve this.

The proposal will harm our discipline, because it will distract established talent from the special roles of academics: curiosity driven research. Academic scholarship has an excellent record of pursuing ideas into places that are exciting and productive, even if they don't result in immediate, tangible benefits and especially if they ruffle the feathers of established, powerful institutions. You can't do that if 80% of your time is spent not annoying a big company. Though big companies belabor promises of complete intellectual freedom to faculty, that can't and won't happen because the purpose of companies is to make money for shareholders.

The proposal harms our students directly. Our faculty at their best secure everyone's future by teaching talented students how to understand the challenges facing the broader world. Such mentorship is enriched by the courage, independence, security, and trained judgement of senior scholars to guide students' perspectives on what is worth doing, what is likely irrelevant, and what is wrong. Engaging with a student body requires an all-in commitment, both in teaching and advising roles. Faculty primarily working elsewhere means cancelled classes. Faculty wedded to a company means advice that's colored by the interest of the company.

The proposal harms our future because it will stifle innovation. University researchers have a great historical record of disruptive entrepreneurism — for example, Google dates back to a paper from the Stanford digital library project. Smooth transitions from academic research to industrial practice are widely encouraged: most universities allow faculty to consult at 20% time, do year-long sabbaticals in industry, or take leave to start companies in order to promote such transitions. But there's a big difference between an industrial leave and a long-term commitment. You can't do disruptive entrepreneurism if 80% of what you do is owned by a big company. Part of the point of being a big company is to control your environment by crushing, containing, or co-opting inconvenient innovations. Faculty who sign on are subject to a huge gravitational force and are hard pressed not to annoy the big company they work for.

Like many really dangerous bargains, the harms are diffuse, and the benefits are focused. One kind of benefit is for faculty who sign on: in addition to the higher industrial salaries, working at a big company provides a chance to lead a team of research engineers to execute large-scale projects that may be used by millions. But another, more alarming, benefit is for big companies: all those potentially disruptive or potentially annoying ideas are now owned or controlled by the big company. Perhaps that's ~~the point of~~ why management supports the proposal.

If industry really wants to help scale and advance computer science research, it's easy to do. Do what many companies are already doing, but do much more of it. Give fellowships to graduate students and scholarships to undergraduate students. Employ students as interns. Pay for named chairs and new buildings. Give lots of faculty small amounts of research money. Make and publish open datasets. Give us easy access to industrial scale computing resources. But don't raid our faculty and tell us it's good for us.

Hot fringe topics —swapping clothes

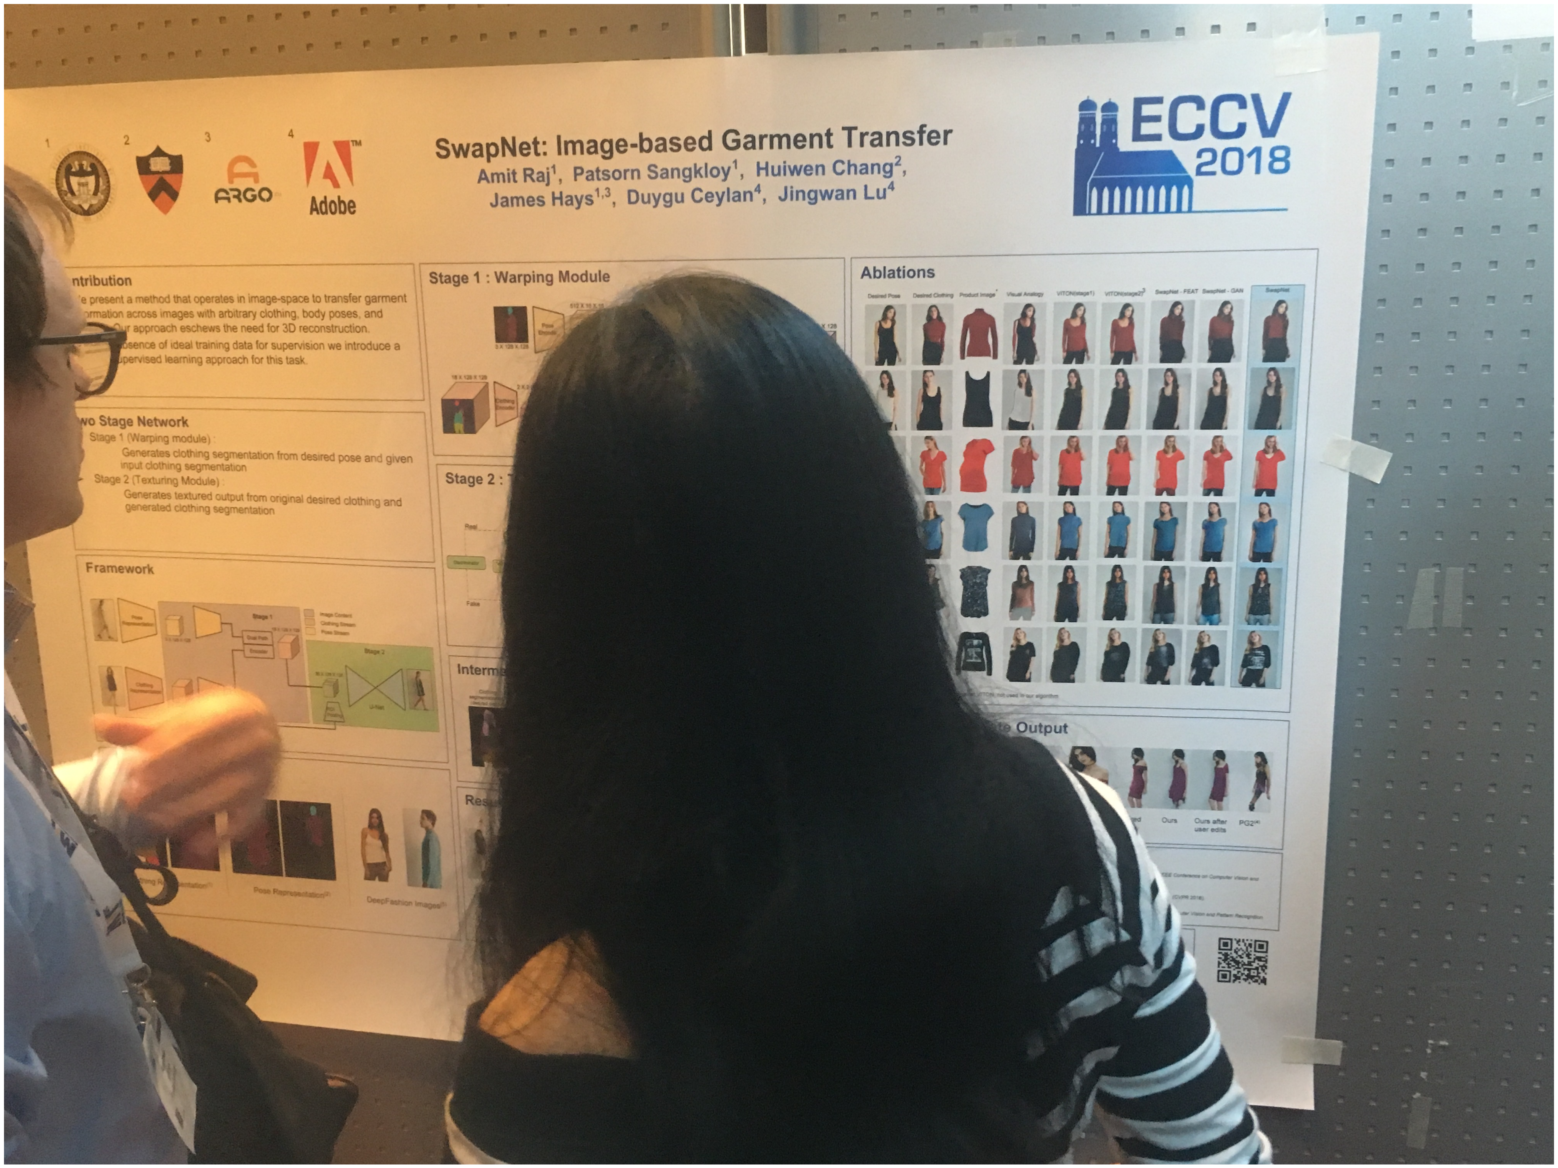

Hot fringe topics —gaze

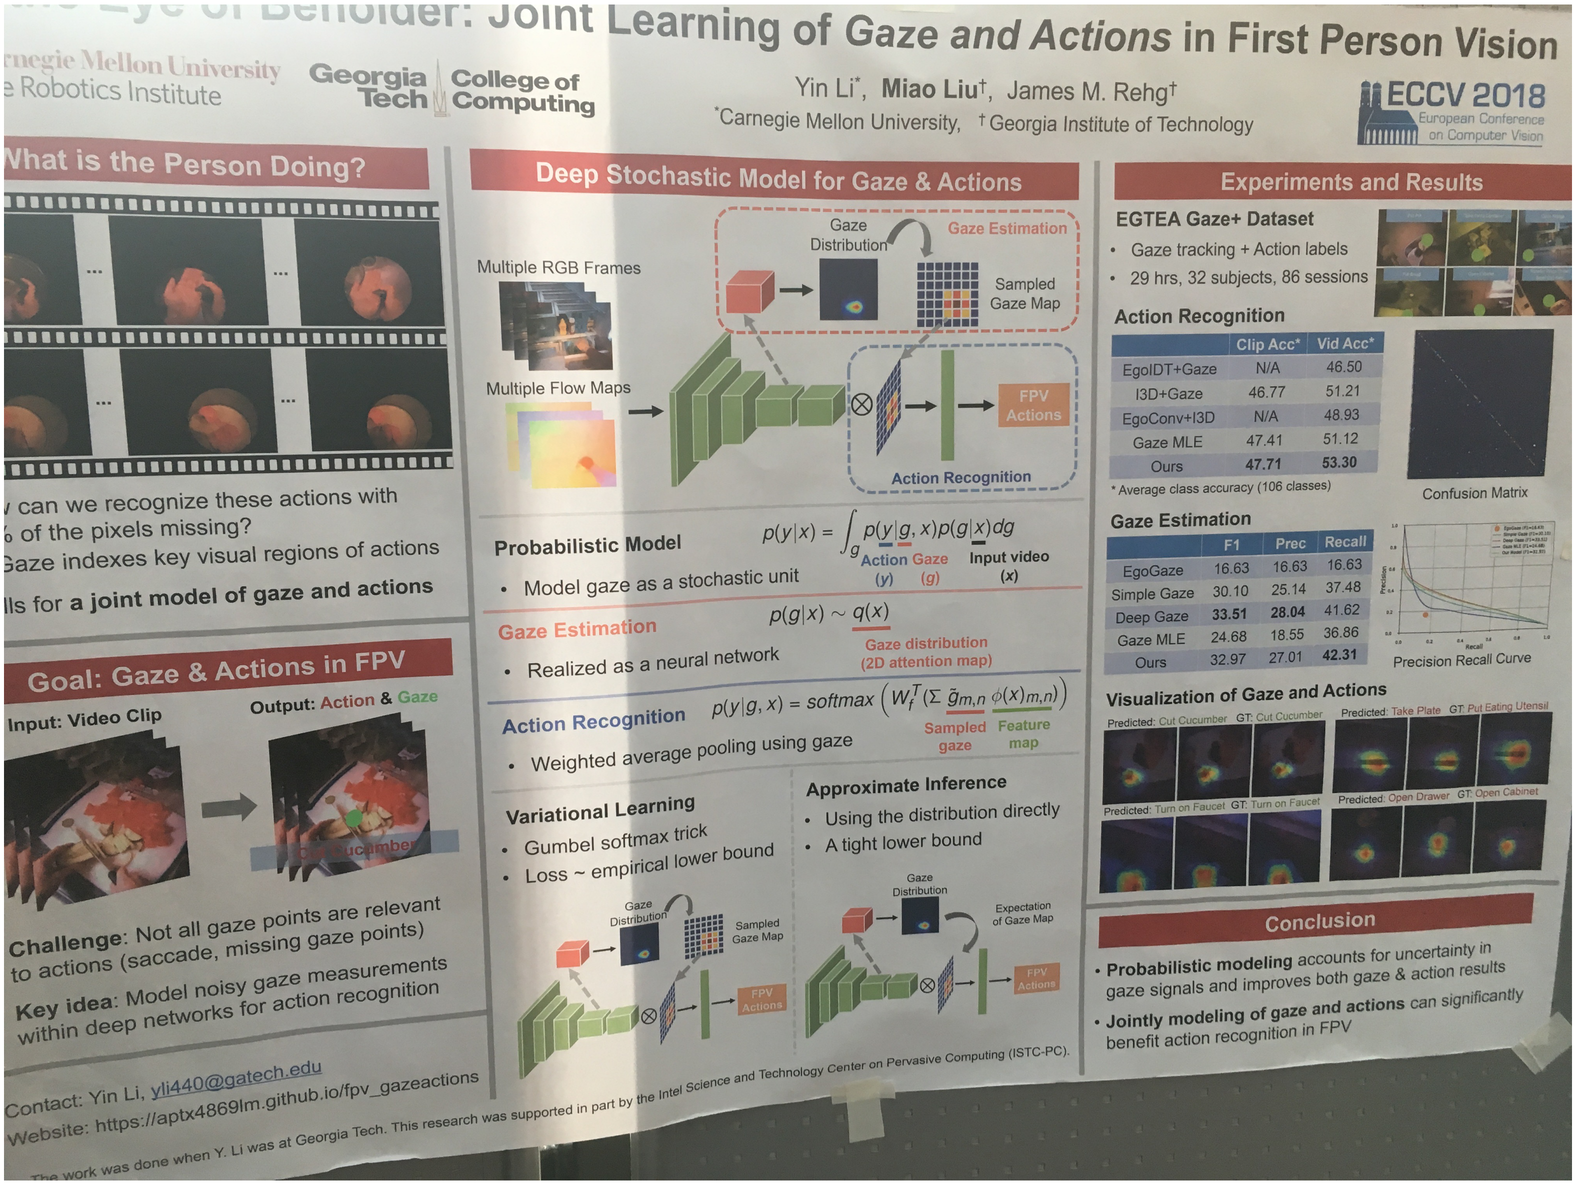

Hot fringe topics

- Swapping cloths
- Gaze
- Augmented reality (e.g., Microsoft HoloLens)
- Any subtopic related to self driving cars
  - EG, finding lanes
- RANSAC
- Explainable deep nets
- Sound and images

Hot main ideas (I)

- Deep Learning
  - Your thing **must** be called "…Net", e.g., SwapNet, ModelNet, GazeNet, …

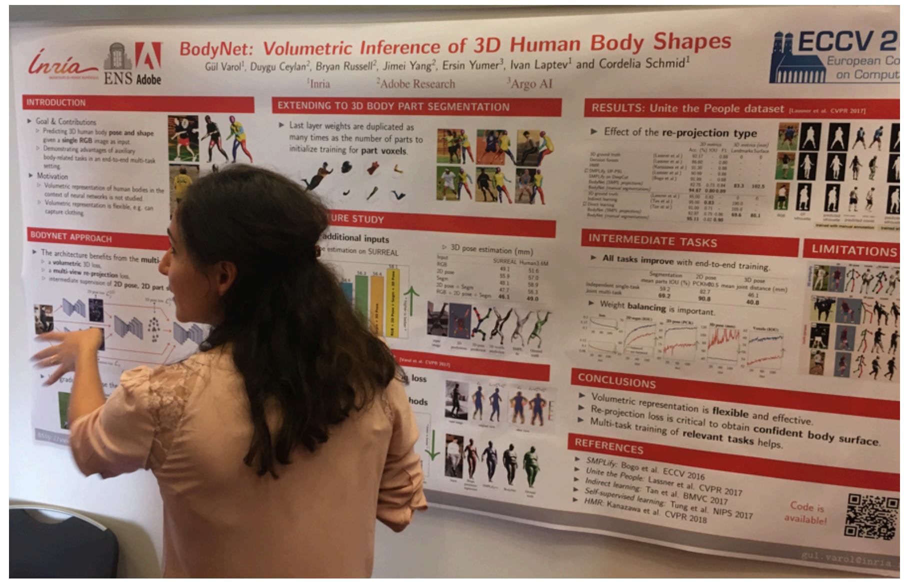

- 3D representation
  - Space filling voxels
  - Meshes

- Deep nets are being built more intelligently, with more vision knowledge built in
  - New phrase: A problem can be "resistant to deepification"

Hot main ideas (II)

- Modular end-to-end
  - Context is self-driving cars
    - Explainability
    - Technical challenge is making the modules differentiable (for training)

- Embodied vision

- Simulation versus reality

Party time


ECCV 2018
Welcome
ECCV 2018
Welcome


ECCV 2018
Welcome


ECCV 2018
Welcome
ECCV 2018
Welcome


Welcome

For next week

- We will start going through some of the papers

- I have read the title and author list of all them
  – There are only about 750 of them (;-))

- I have selected promising ones on topics like
  – 3D representation (biased towards mesh)
  – State of the art ideas in deep learning
  – Gaze
  – Emotion (e.g., FACS) recognition
  – Optimization

- tracking
  - Mir_Rayat_Imtiaz_Hossain_Exploiting_temporal_information_ECCV_2018_paper.pdf
  - Efstratios_Gavves_Long-term_Tracking_in_ECCV_2018_paper.pdf
  - Chanho_Kim_Multi-object_Tracking_with_ECCV_2018_paper.pdf
- sound
  - Ruohan_Gao_Learning_to_Separate_ECCV_2018_paper.pdf
- slic
  - Hiroaki_Santo_Light_Structure_from_ECCV_2018_paper.pdf
- shrimp
  - Yiming_Qian_Simultaneous_3D_Reconstruction_ECCV_2018_paper.pdf
- shape
  - Zorah_Laehner_DeepWrinkles_Accurate_and_ECCV_2018_paper.pdf
  - Thomas_Probst_Model-free_Consensus_Maximization_ECCV_2018_paper.pdf
  - Thomas_Probst_Incremental_Non-Rigid_Structure-from-Motion_ECCV_2018_paper.pdf
  - Nanyang_Wang_Pixel2Mesh_Generating_3D_ECCV_2018_paper.pdf
  - Ladicky_From_Point_Clouds_ICCV_2017_paper.pdf
  - Jiajun_Wu_Learning_3D_Shape_ECCV_2018_paper.pdf
  - Gul_Varol_BodyNet_Volumetric_Inference_ECCV_2018_paper.pdf
  - Angjoo_Kanazawa_Learning_Category-Specific_Mesh_ECCV_2018_paper.pdf
- scene-geometry-and-physics
  - Zhijian_Liu_Physical_Primitive_Decomposition_ECCV_2018_paper.pdf
  - Tianfan_Xue_Seeing_Tree_Structure_ECCV_2018_paper.pdf
  - Tian_Ye_Interpretable_Intuitive_Physics_ECCV_2018_paper.pdf
- optimzation
  - Tat-Jun_Chin_Robust_fitting_in_ECCV_2018_paper.pdf
  - Siddharth_Tourani_MPLP_Fast_Parallel_ECCV_2018_paper.pdf
  - Csaba_Domokos_MRF_Optimization_with_ECCV_2018_paper.pdf
- illumination
  - Zhengqi_Li_CGIntrinsics_Better_Intrinsic_ECCV_2018_paper.pdf
  - Wei-Chiu_Single_Image_Intrinsic_ECCV_2018_paper.pdf
- human-object-interaction
  - Keizo_Kato_Compositional_Learning_of_ECCV_2018_paper.pdf

- gaze
  - Tobias_Fischer_RT-GENE_Real-Time_Eye_ECCV_2018_paper.pdf
  - Seonwook_Park_Deep_Pictorial_Gaze_ECCV_2018_paper.pdf
  - Heinonen_Eemeli.pdf
  - Eunji_Chong_Connecting_Gaze_Scene_ECCV_2018_paper.pdf
  - egocentric
- fusion
  - Ming_Liang_Deep_Continuous_Fusion_ECCV_2018_paper.pdf
- face-detailed
  - Weixuan_Chen_DeepPhys_Video-Based_Physiological_ECCV_2018_paper.pdf
  - Ciprian_Corneanu_Deep_Structure_Inference_ECCV_2018_paper.pdf
- explainable-AI
- emotion
  - Guosheng_Hu_Deep_Multi-Task_Learning_ECCV_2018_paper.pdf
  - Benitez-Quiroz_EmotioNet_An_Accurate_CVPR_2016_paper.pdf
  - Albert_Pumarola_Anatomically_Coherent_Facial_ECCV_2018_paper.pdf
- ECCV-18-notes.txt.org
- deep-learning
  - Yuxin_Wu_Group_Normalization_ECCV_2018_paper.pdf
  - Xin_Wang_SkipNet_Learning_Dynamic_ECCV_2018_paper.pdf
  - Viorica_Patraucean_Massively_Parallel_Video_ECCV_2018_paper.pdf
  - Tae-Hyun_Oh_Learning-based_Video_Motion_ECCV_2018_paper.pdf
  - Safa_Cicek_SaaS_Speed_as_ECCV_2018_paper.pdf
  - Rene_Ranftl_Deep_Fundamental_Matrix_ECCV_2018_paper.pdf
  - Ramprasaath_Ramasamy_Selvaraju_Choose_Your_Neuron_ECCV_2018_paper.pdf
  - Pei_Wang_Towards_Realistic_Predictors_ECCV_2018_paper.pdf
  - Navaneeth_Bodla_Semi-supervised_FusedGAN_for_ECCV_2018_paper.pdf
  - Michael_Moeller_Lifting_Layers_Analysis_ECCV_2018_paper.pdf
  - Liang_Mi_Variational_Wasserstein_Clustering_ECCV_2018_paper.pdf
  - Konstantin_Shmelkov_How_good_is_ECCV_2018_paper.pdf
  - Jiqing_Wu_Wasserstein_Divergence_For_ECCV_2018_paper.pdf
  - Gedas_Bertasius_Object_Detection_in_ECCV_2018_paper.pdf
  - Dong_Su_Is_Robustness_the_ECCV_2018_paper.pdf
  - Chenxi_Liu_Progressive_Neural_Architecture_ECCV_2018_paper.pdf
  - Antonio_Torralba_Interpretable_Basis_Decomposition_ECCV_2018_paper.pdf
- bibtex.bib
- amimal-trap
  - Beery_Recognition_in_Terra_ECCV_2018_paper.pdf

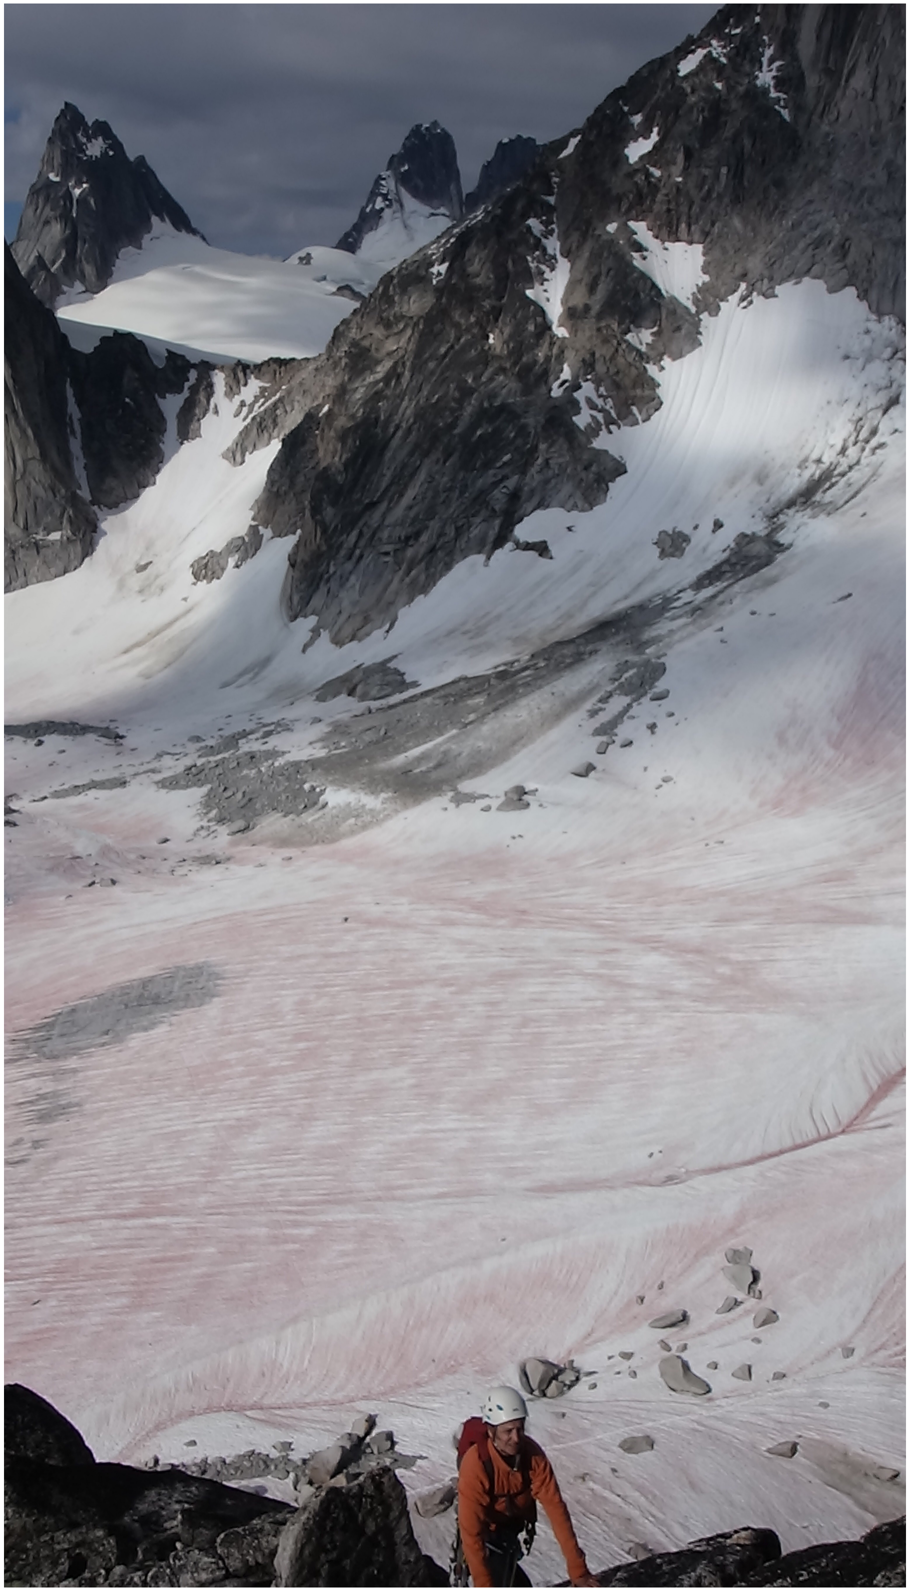



Interdisciplinary Visual Intelligence Lab

http://ivilab.org