Towards Robust Perception Systems in Real World Environments

Feb 10, 2022
B Exam

Hubert Lin

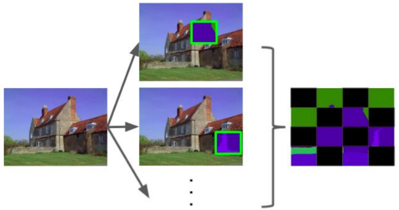

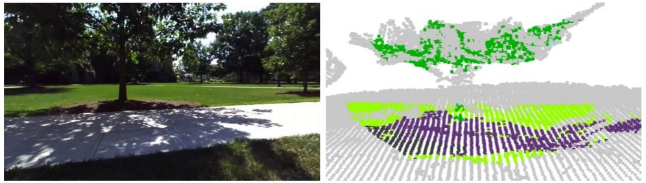

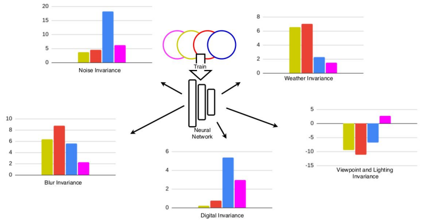

Introduction

Visual perception systems useful for many applications:

- Robotics
- Self-driving
- Visual discovery
- Medical diagnostics
- …

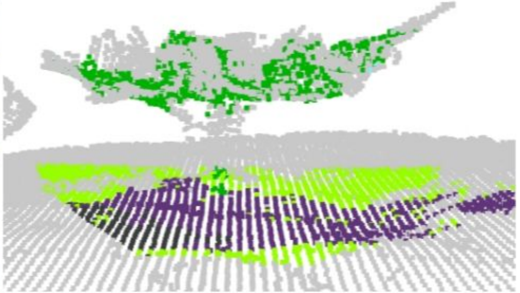

Many modern systems are based on neural networks.

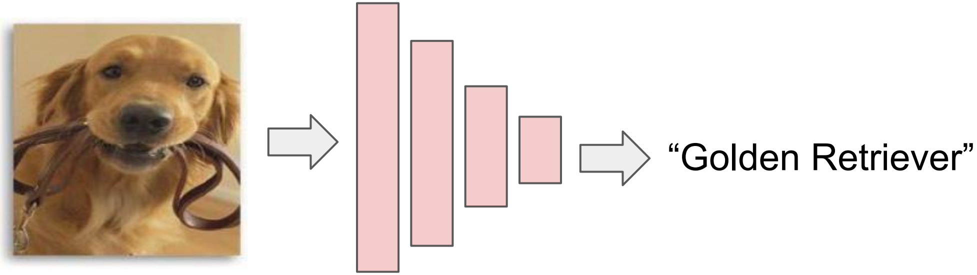

Many imperfections in real world images…

Noise

Camera
Blur

Digital
Manipulation

Unconventional
Viewpoint

Rightmost photo: instagram @*goldenlifeof_duke*

Effect of noisy images

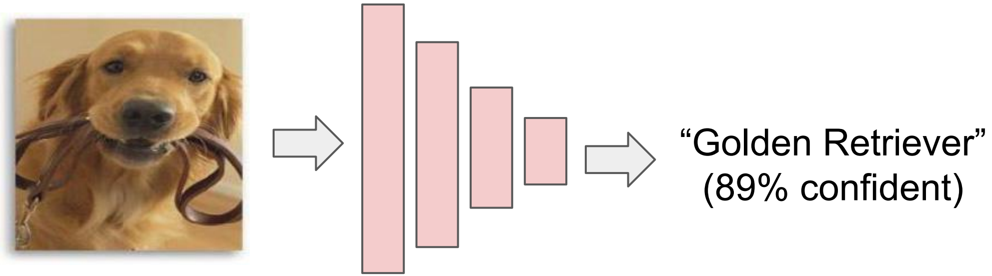

- Humans can recognize objects[1] and materials[2] in non-ideal images.
- Neural networks may struggle to perform well.

[1]Geirhos et al 2018, Generalisation in Humans and DNNs
[2]Sharan et al 2014, Accuracy and speed of material categorization in real-world images

Effect of noisy images

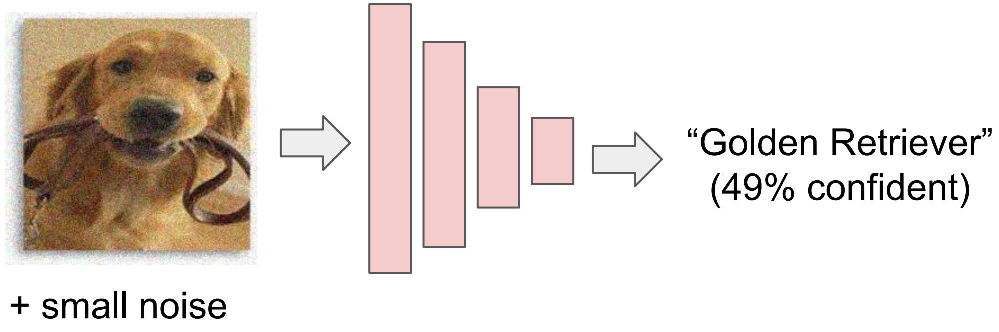

- Humans can recognize objects[1] and materials[2] in non-ideal images.
- Neural networks may struggle to perform well.

[1]Geirhos et al 2018, Generalisation in Humans and DNNs
[2]Sharan et al 2014, Accuracy and speed of material categorization in real-world images

Effect of noisy images

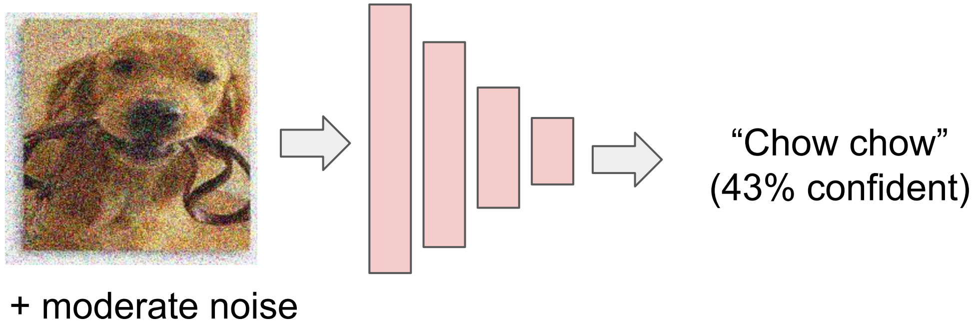

- Humans can recognize objects[1] and materials[2] in non-ideal images.
- Neural networks may struggle to perform well.

[1]Geirhos et al 2018, Generalisation in Humans and DNNs
[2]Sharan et al 2014, Accuracy and speed of material categorization in real-world images

Effect of noisy images

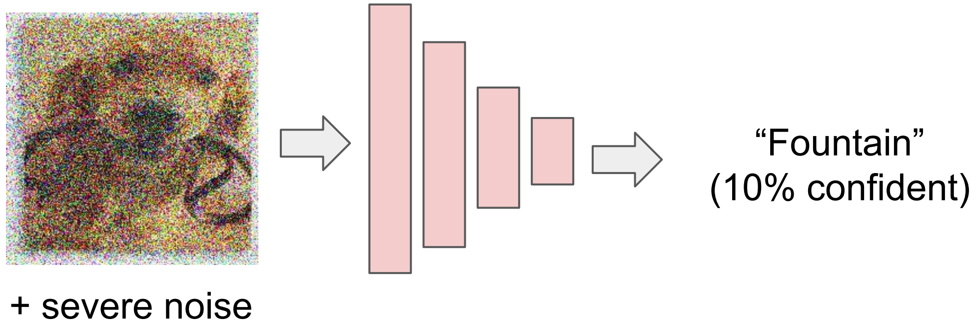

- Humans can recognize objects[1] and materials[2] in non-ideal images.
- Neural networks may struggle to perform well.

[1]Geirhos et al 2018, Generalisation in Humans and DNNs
[2]Sharan et al 2014, Accuracy and speed of material categorization in real-world images

Talk Outline

Many challenges in improving perception systems in real world.

1. **Better annotation tools.**
   [ICCV 2019]

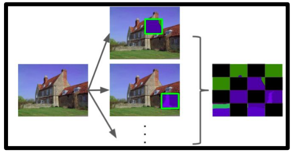

2. **Learning robust visual invariances.**
   [PLOS One 2021, FAPER ICPR 2020, CVPR 2021]

3. **Reasoning about perception uncertainties.**
   [ICRA 2020]

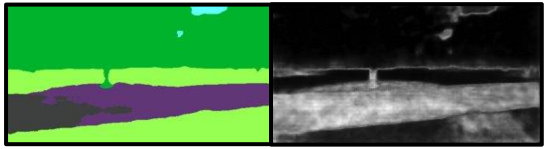

4. **Summary.**

This talk will primarily focus on **(2)**.

Talk Outline

Many challenges in improving perception systems in real world.

1. **Better annotation tools.**
   [ICCV 2019]

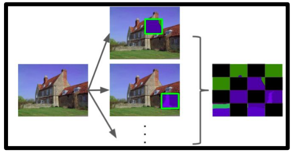

2. **Learning robust visual invariances.**
   [PLOS One 2021, FAPER ICPR 2020, CVPR 2021]

3. **Reasoning about perception uncertainties.**
   [ICRA 2020]

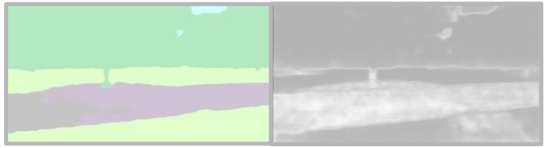

4. **Summary.**

Illustrative example

Giraffes face left in training set.

Unseen image: Giraffe?

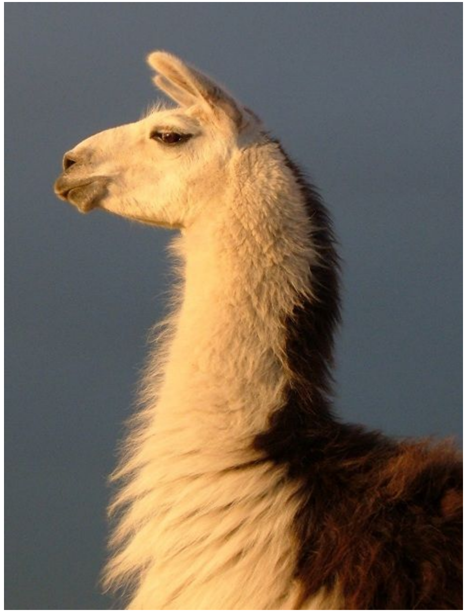

Need more data

One possible solution: train with more images of giraffes in different poses.

Original Dataset

New Labeled Data

What kind of labels are useful?

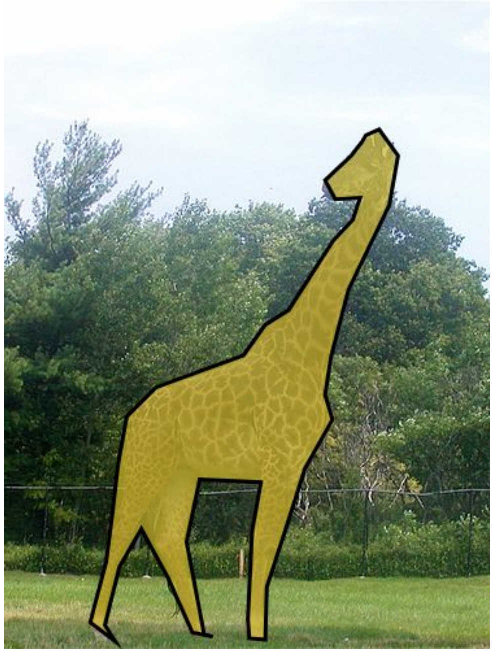

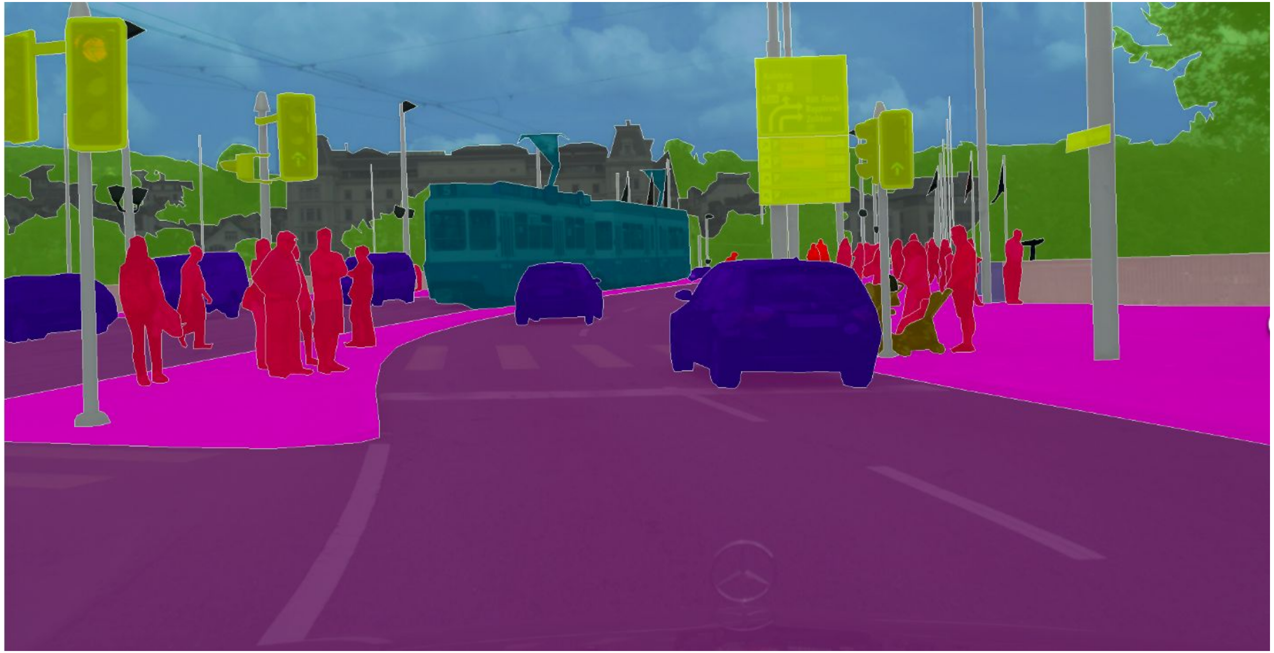

Lin et al 2014; Cordts et al 2016

Cheap but coarse alternatives

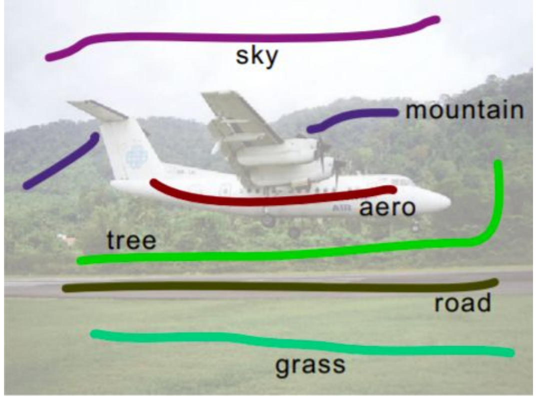

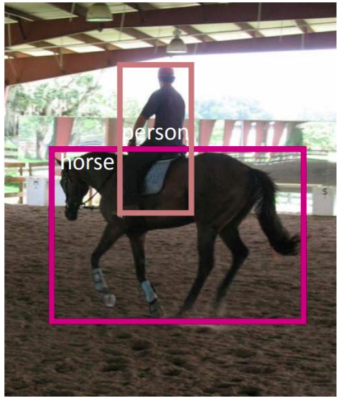

Cordts et al 2016; Bearman et al 2015; Lin et al 2016; Dai et al 2015

Block Annotation

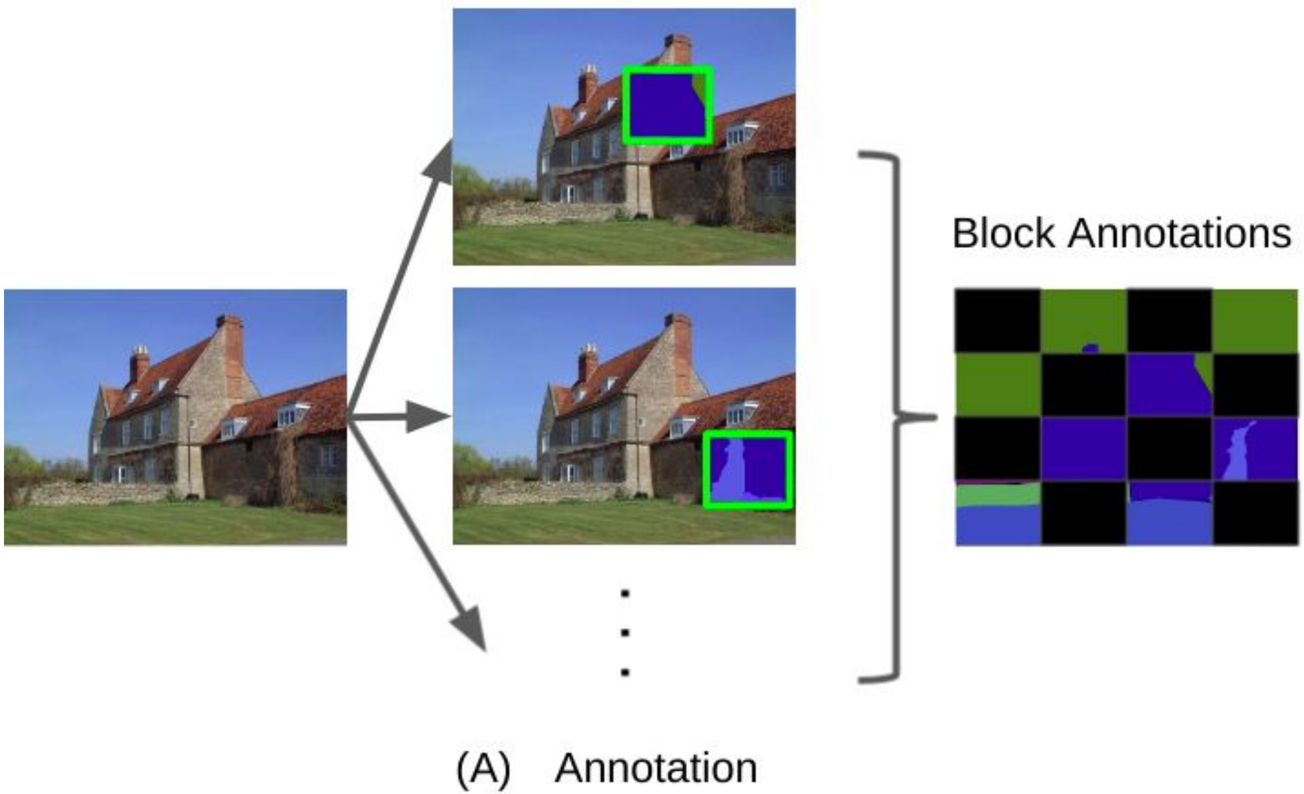

Lin, Upchurch, Bala, ***"Block Annotation: Better Image Annotation with Sub-Image Decomposition"****, ICCV 2019*

Block Annotation

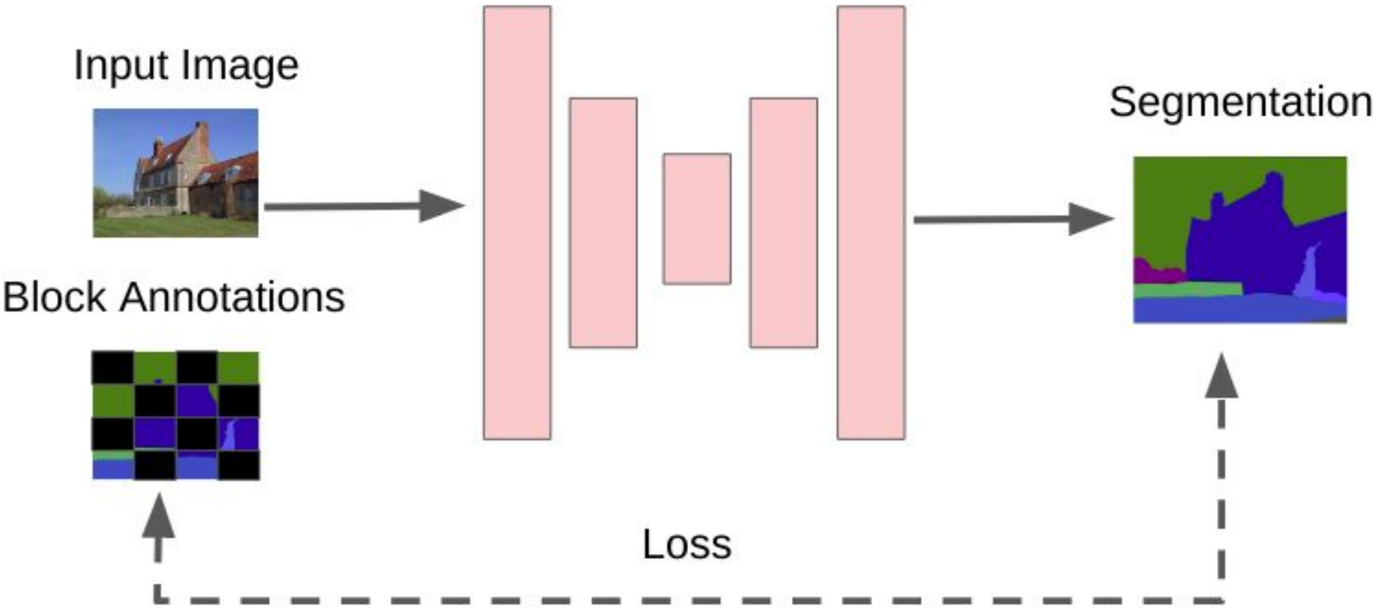

(B) Segmentation

Lin, Upchurch, Bala, ***"Block Annotation: Better Image Annotation with Sub-Image Decomposition"***, *ICCV 2019*

Key Findings: Block Annotation

Crowdworkers produce high quality annotations, and more cheaply than conventional methods.

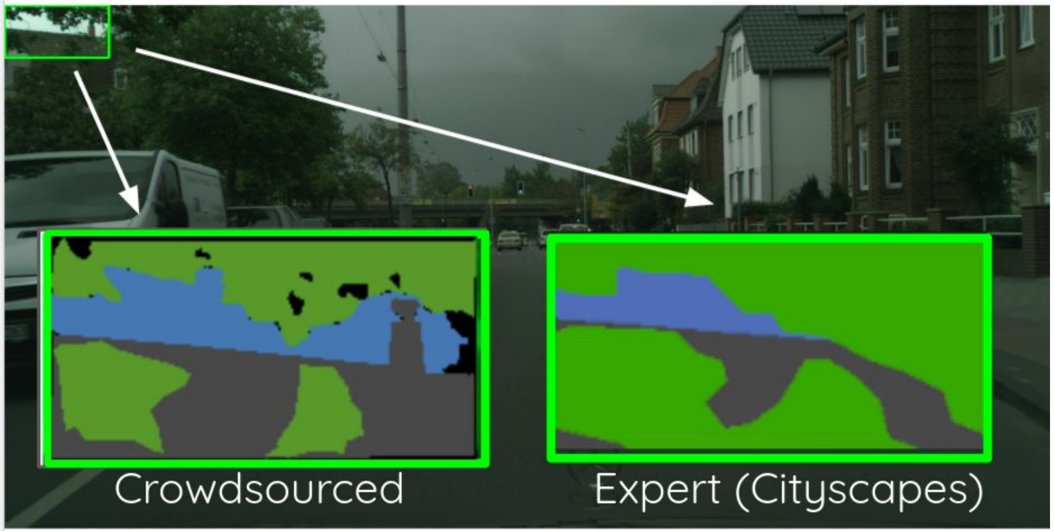

Key Findings: Block Annotation

High performing semantic segmentation models learned – up to 97% of full supervision performance with 1/10th annotation time.

Cityscapes	**Ours:** Block (7 min)	Coarse (7 min [14])	Full Supervision (90 min [14])
mIOU (%)	**72.1**	68.8	77.7

Pascal	**Ours:** Block (25 sec)	Scribbles (25 sec [36])	Full Supervision (4 min [41])
mIOU (%)	**67.2**	63.1 [36]	69.6

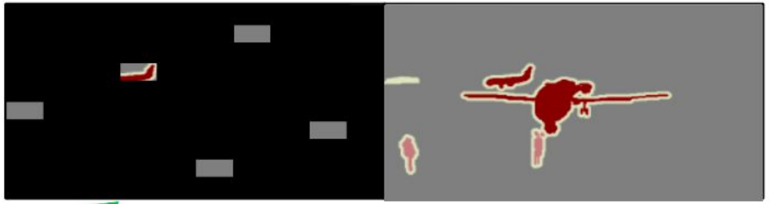

Talk Outline

Many challenges in improving perception systems in real world.

1. **Better annotation tools.**
   [ICCV 2019]

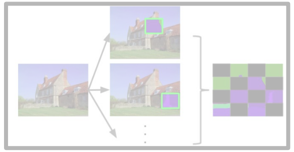

2. **Learning robust visual invariances.**
   [PLOS One 2021, FAPER ICPR 2020, CVPR 2021]

3. **Reasoning about perception uncertainties.**
   [ICRA 2020]

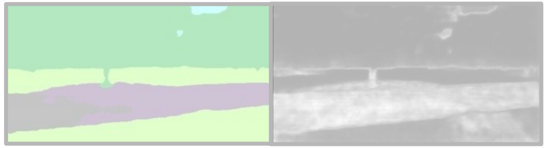

4. **Summary.**

What about synthetic data?

Problem: Giraffes always face towards left in original dataset.

- Label more data – expensive.
- Alternative? Synthetically create images+labels by applying a left-right reflection to the existing set of images.

Original Dataset

+

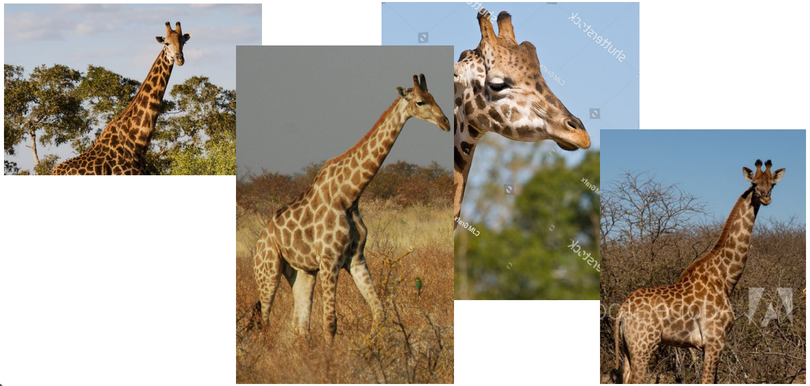

Data Augmentation

Image transformations encourage the network to ignore some signals in the data.

- Reflected image pairs: Network will not rely on left-right orientation when classifying an animal.

Data Augmentation

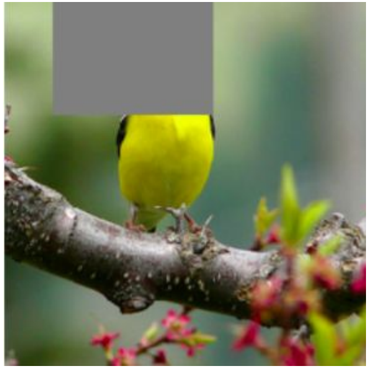

CutOut

AutoAugment

RandAugment

AugMix

DeepAugment

Common goals:
(a) Preserve semantics.
(b) Manipulate non-robust features.

DeVries and Taylor 2017; Zhang et al 2018; Yun et al 2019; Cubuk et al 2018,2019; Hendrycks et al 2020

Artwork: Fra Carnevale
“The Birth of the Virgin” 1467

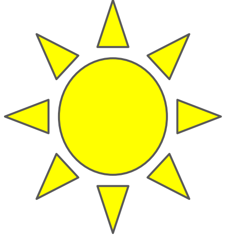

Cavanagh 2005, *“The Artist as Neuroscientist”*

Paintings as Implicit Data Augmentation

Artworks implicitly encode human visual invariances by omitting or altering unimportant details for perception.

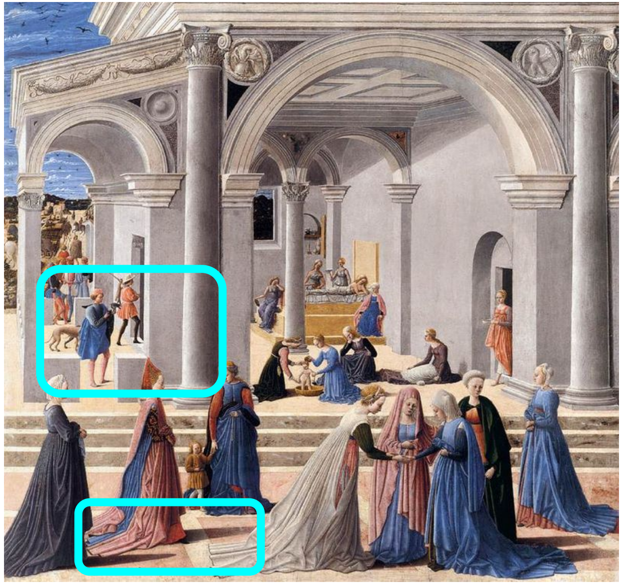

See e.g.: Cavanagh 2005; Mamassian 2008; Gombrich 1960

Materials In Paintings

https://materialsinpaintings.tudelft.nl/

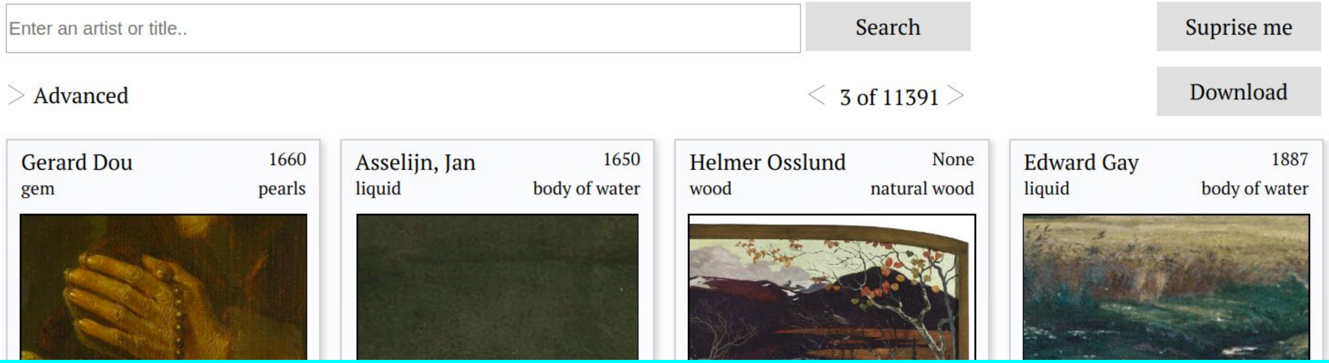

Van Zuijlen, Lin, Bala, Pont, Wijntjes, ***"Materials In Paintings (MIP): An interdisciplinary dataset for perception, art history, and computer vision"***, *PLOS One 2021*

Painterly Biases

Highlights on glass cups:

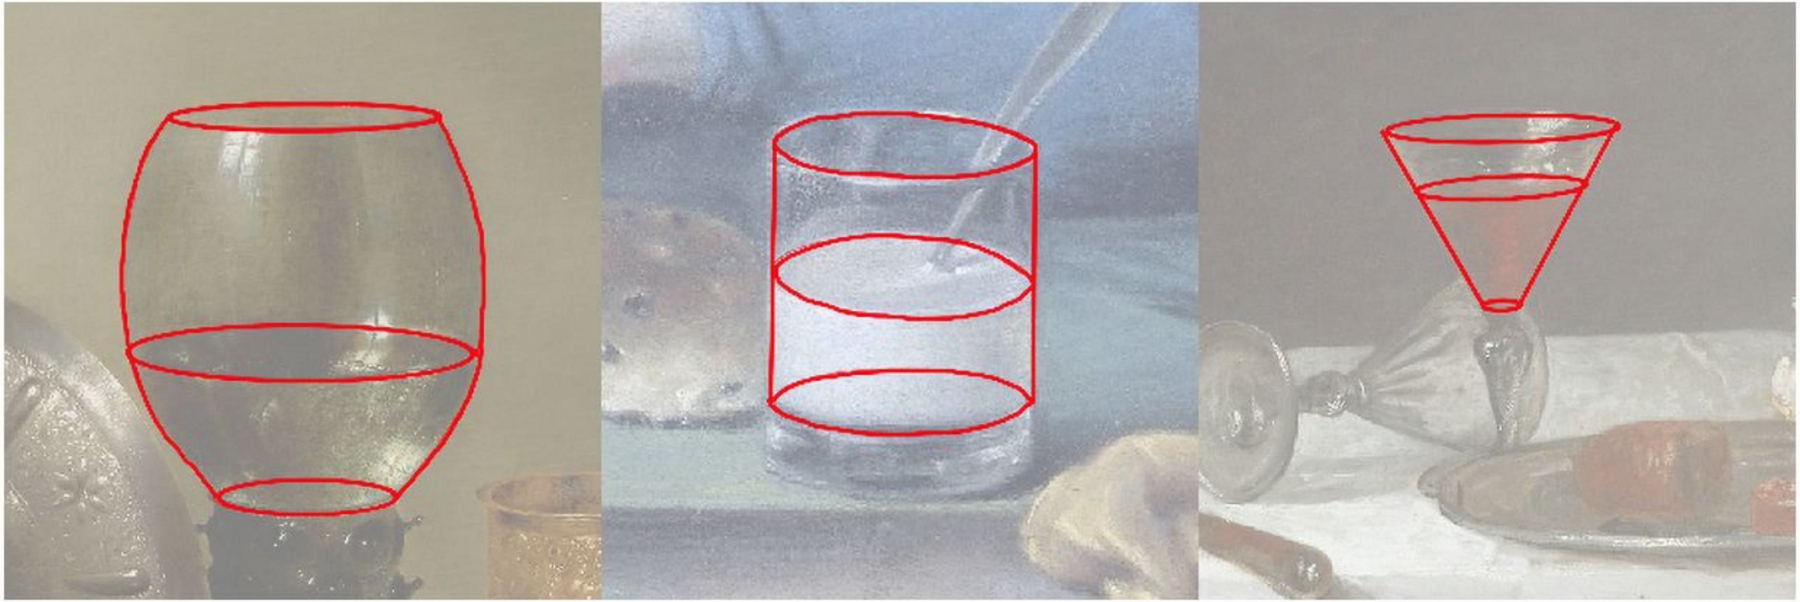

How are glass highlights depicted?

1. Paintings have more localized highlights.

2. Painting highlights are less ambiguous.
   - 50% higher agreement (recall) between participants.

Highlight Heatmaps

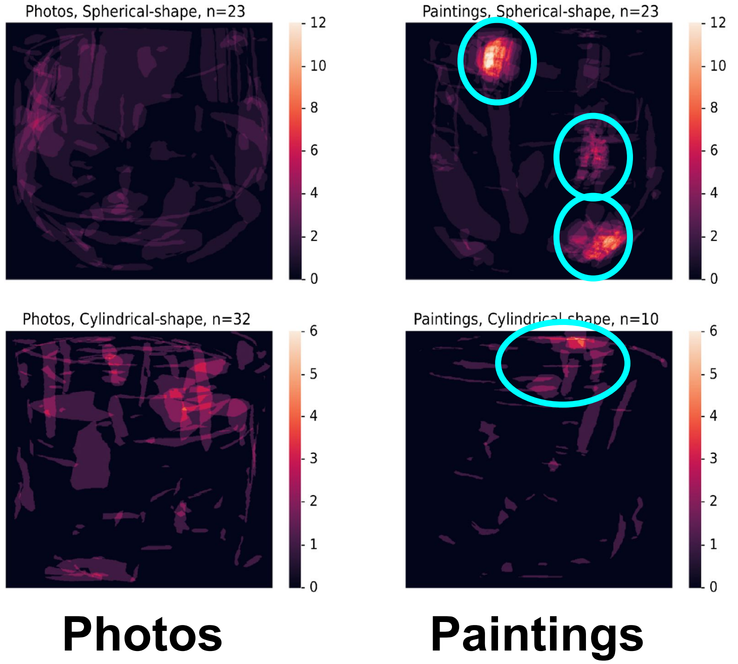

Learning From Painterly Biases

Local cues like highlights on the satin or silk fabrics are emphasized by artists.



Lin, Van Zuijlen, Wijntjes, Pont, Bala, ***"Insights from a Large-Scale Database of Material Depictions in Paintings"***, *FAPER ICPR 2020*

Learning From Painterly Biases

Compare models trained on paintings or photos to distinguish satin from cotton.

- Assess human preference for cues used by models.

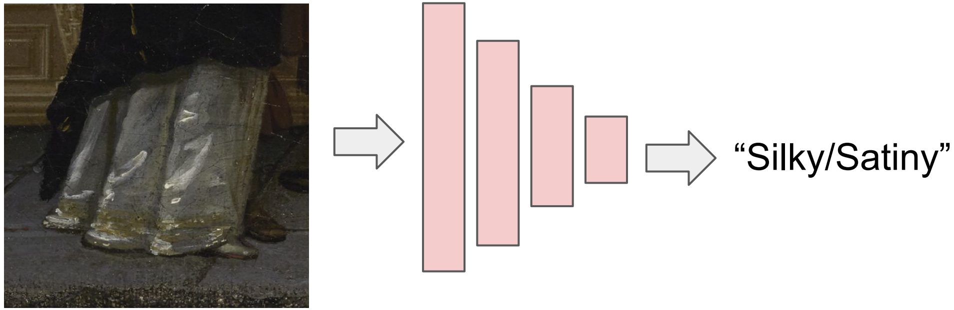

Human Preferences for Cues

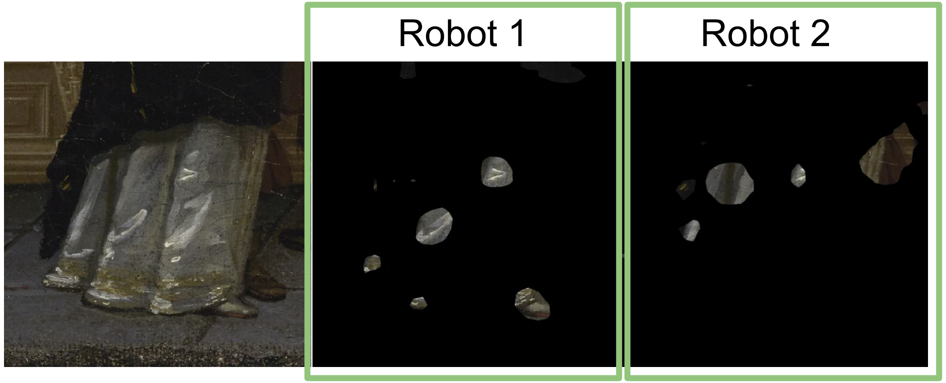

"Two different robots think these regions in the image look like silk/satin. Which robot do you agree with more?"

Humans are shown cues used by each classifier and prompted to select which set of cues they prefer.

Which cues do humans prefer?

	MEAN	...	Silk/Satin Photos
Photo Classif. Preferred	$44.9 \pm 1.9\%$		$48.9 \pm 3.1\%$
Painting Classif. Preferred	$55.1 \pm 1.9\%$		$51.1 \pm 3.1\%$

- Overall, cues learned from paintings preferred over cues learned from photos.
- For photos of silk/satin, cues learned from paintings equally preferred to cues learned from photos despite domain shifts.

"Fake" Paintings via Style Transfer

Style transfer: methods for creating painting-like images from photos

Giraffe photo

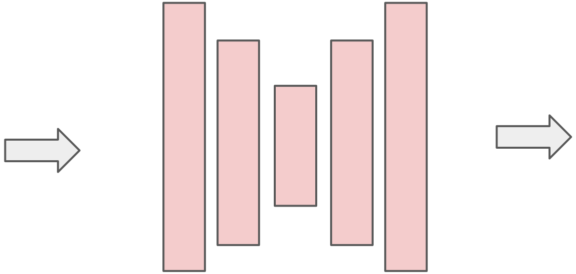

Giraffe in the style of a Monet painting

Sanakoyeu et al 2018

“Fake” Paintings via Style Transfer

Style transfer: methods for creating painting-like images from photos

Huang and Belongie 2017; Sanakoyeu et al 2018

Learning from Paintings vs Stylized Images

Do models learn similar behaviors from paintings and stylized images?

- Does style transfer allow us to replace paintings altogether?

Painting

Stylized Photo

Lin, Van Zuijlen, Wijntjes, Pont, Bala, ***"What Can Style Transfer and Paintings Do For Model Robustness?"***, *CVPR 2021*

Datasets

Materials:

- Photographs of materials from existing datasets (MINC[1], COCO[2])
- Paintings of materials from Materials in Paintings (MIP[3])

Objects:

- Existing dataset of photos, paintings, cartoons, and sketches (PACS[4]).

[1] Bell et al 2015; [2] Lin et al 2014; [3] Van Zuijlen et al 2021; [4] Li et al 2017

Evaluating Model Behavior

Interested in model behavior in real-world settings with imperfect images.

- High accuracy on these images = model is more "robust"

Noise

Camera Blur

Digital Manipulation

Unconventional Viewpoint

Evaluating Model Behavior

Accuracy with respect to **common image corruptions**:

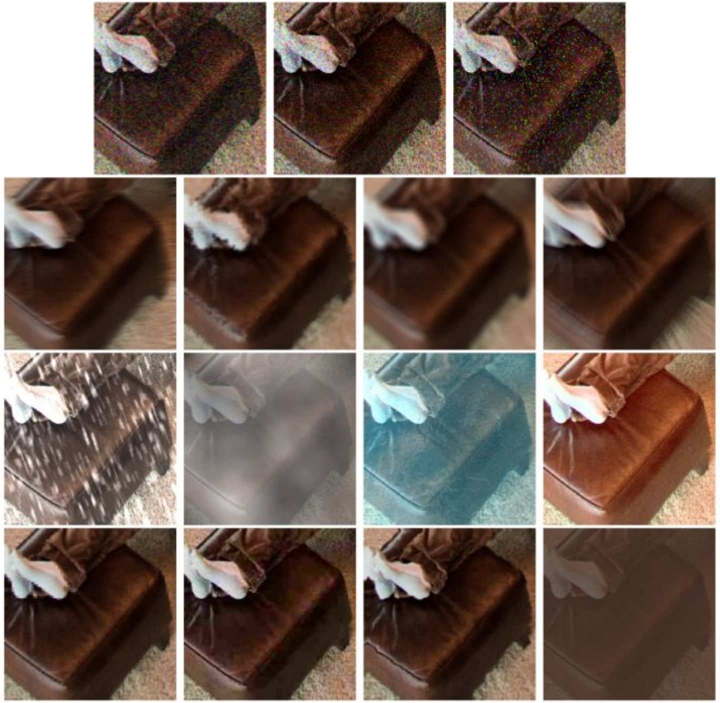

Noise

Blur

Weather

Digital

Hendrycks and Dietterich 2019

Evaluating Model Behavior

Accuracy with respect to **out-of-distribution photos** (different viewpoint, lighting):

Materials → FMD[1]

PACS → Subset of YFCC100M[2]

[1] Sharan et al 2014; [2] Thomee et al 2015

Research Questions

1. Does learning from paintings improve model robustness?
2. Does learning from stylized images improve model robustness?
3. How do models trained on paintings differ from models trained on stylized images?

Lin, Van Zuijlen, Wijntjes, Pont, Bala, ***"What Can Style Transfer and Paintings Do For Model Robustness?"****, CVPR 2021*

Experiment Setup

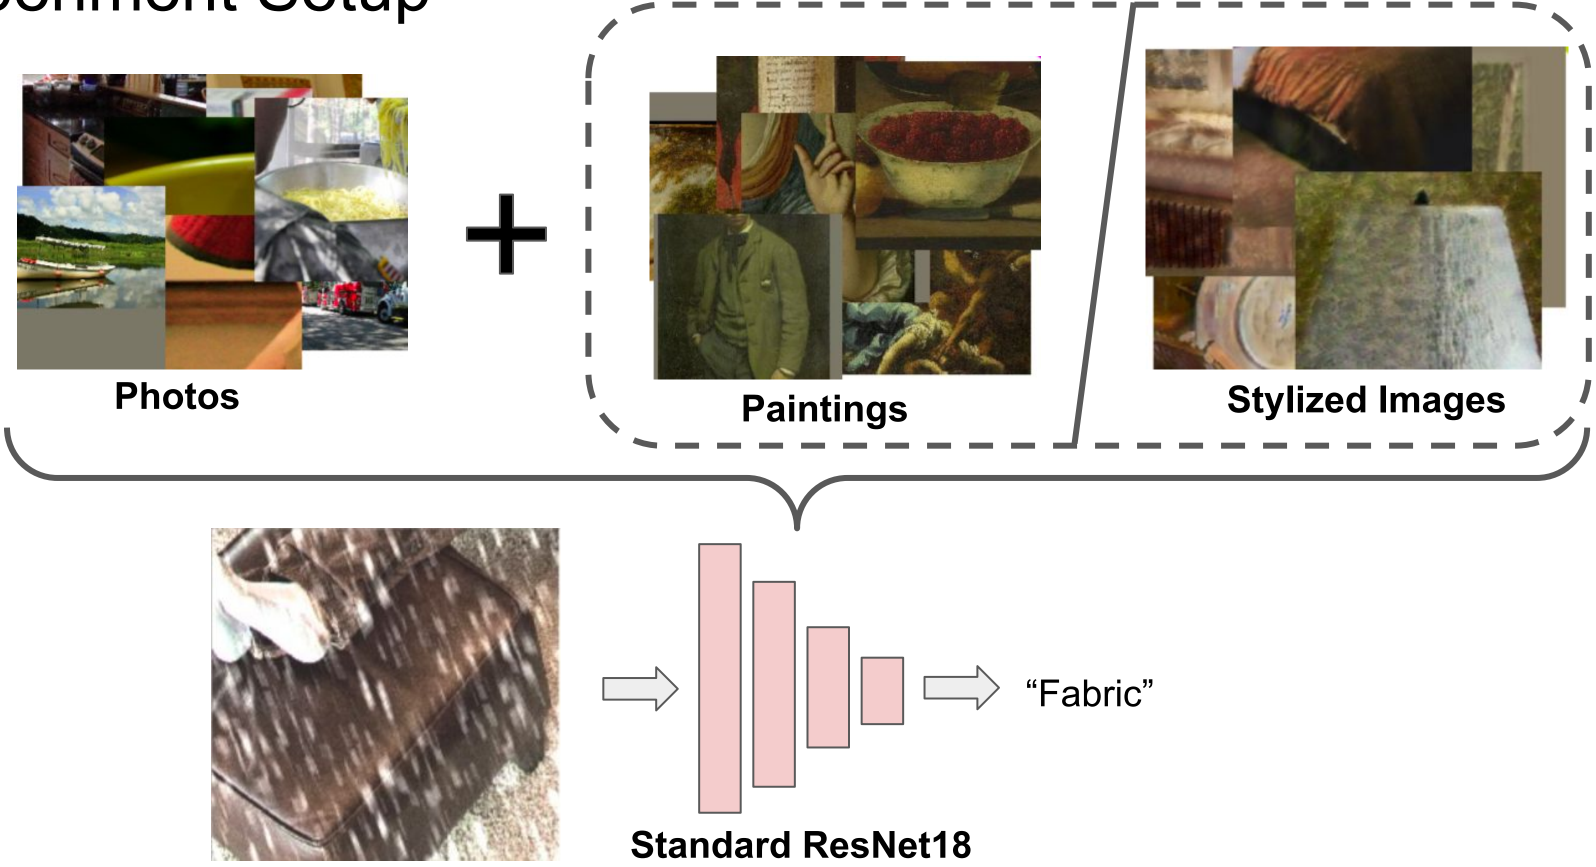

Research Questions

1. Does learning from paintings improve model robustness?

Does learning from paintings improve robustness?

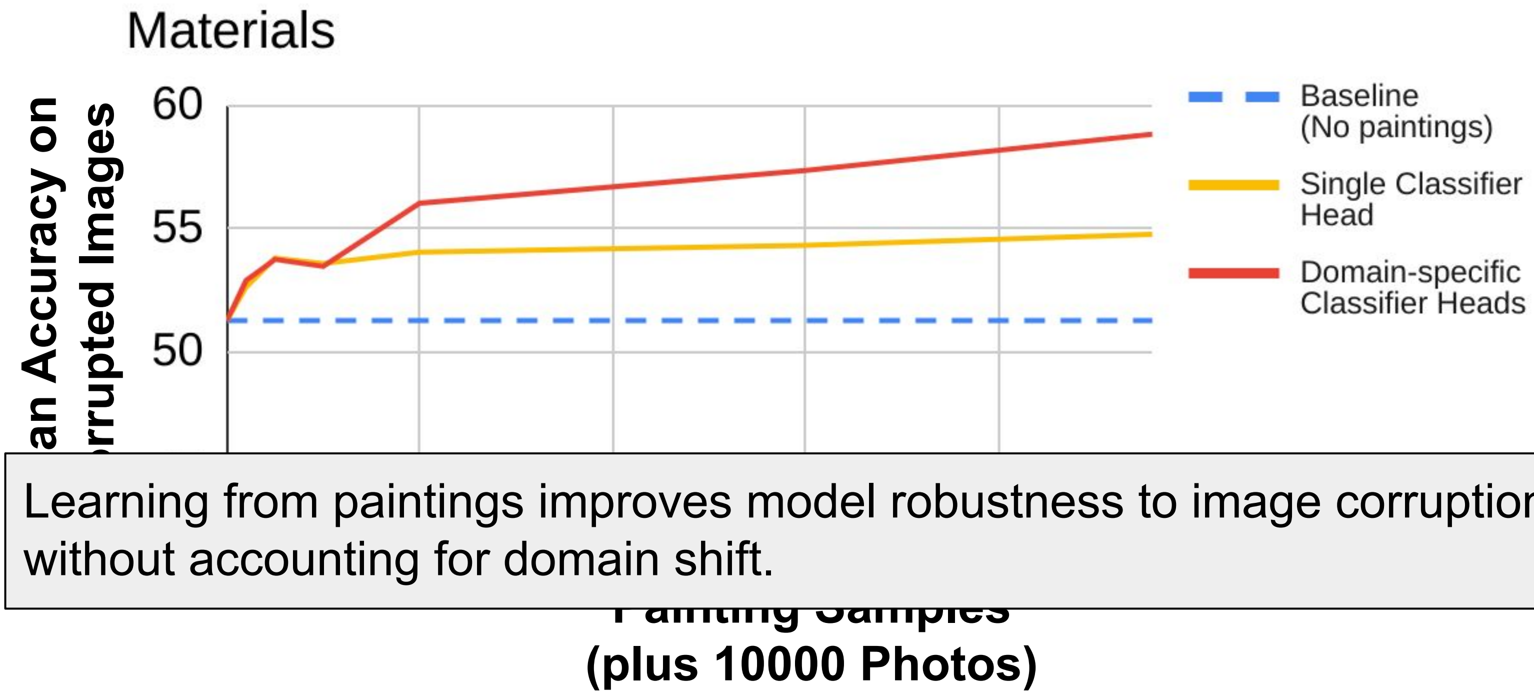

Learning from paintings improves model robustness to image corruptions, even without accounting for domain shift.

Is learning from paintings data-efficient?

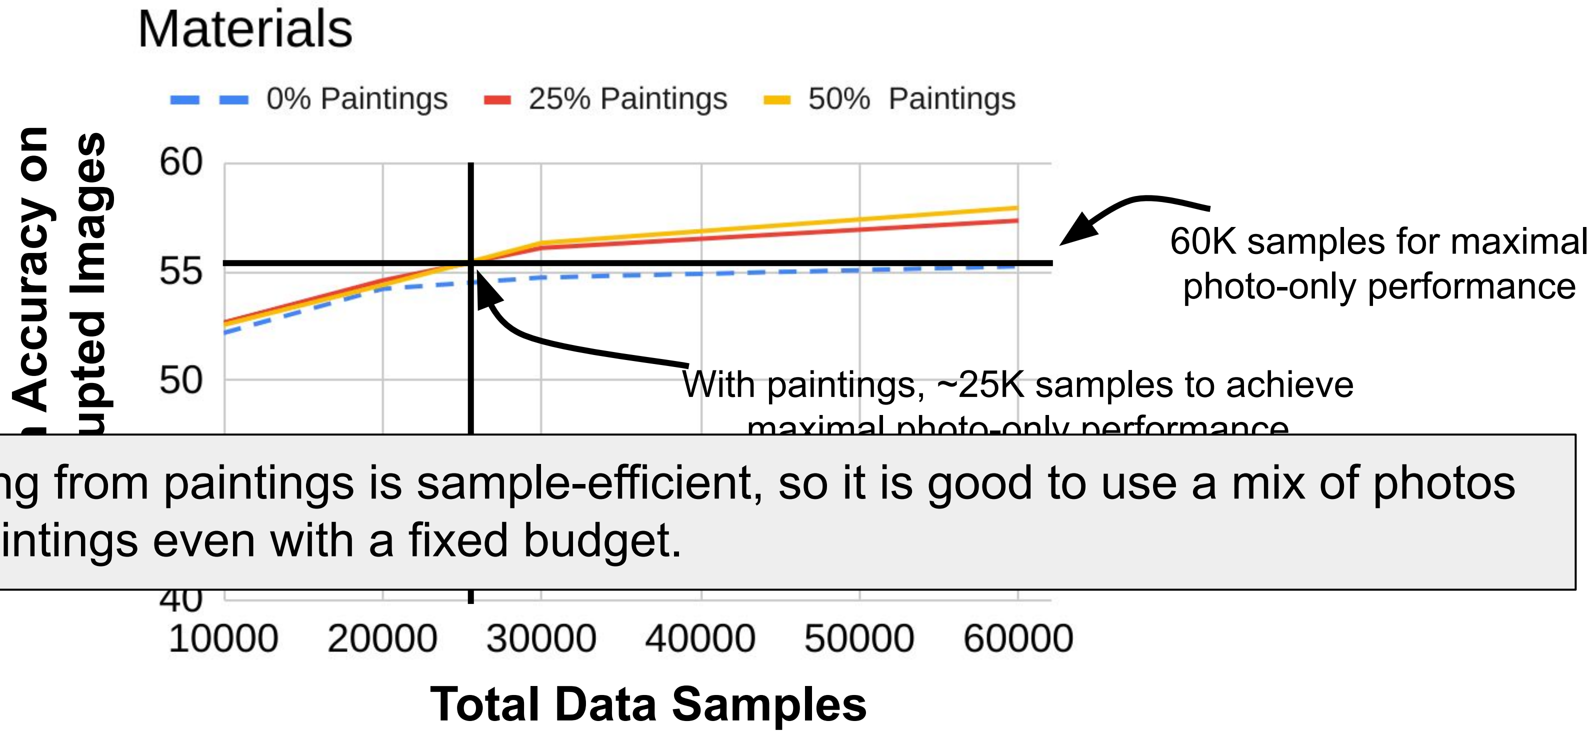

Learning from paintings is sample-efficient, so it is good to use a mix of photos and paintings even with a fixed budget.

Can sketches and cartoons work too?

Other artwork like sketches and cartoons are also perceptually meaningful.



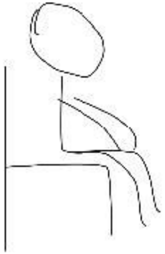


Can sketches and cartoons work too?

Training Data (# Samples)	Mean Corruption Acc (%)
Materials	
Photo (30K)	54.73±0.25
Photo + **Painting** (15K + 15K)	**56.31**±0.27 (+)
PACS	
Photo (1500)	76.16±0.34
Photo + **Painting** (750 + 750)	**79.41**±0.55 (+)
Photo + Cartoon (750 + 750)	75.38±0.36 (−)
Photo + Sketch (750 + 750)	73.85±0.39 (−)
DomainNet [28]	
Photo (120K)	36.59±0.12
Photo + **Painting** (90K + 30K)	**39.00**±0.14 (+)
Photo + Sketch (90K + 30K)	37.57±0.22 (+)
[illegible]	[illegible]
[illegible]	[illegible]
VisDA [29]	
Photo (30K)	**65.97**±0.33
Photo + Rendering (15K + 15K)	63.90±0.21 (−)

Paintings are uniquely useful due to their balance of realism and abstraction.

Research Questions

1. Does learning from paintings improve model robustness?
    - YES – improves robustness to image corruptions.
    - Cost-effective compared to only photos.
    - More abstract art forms do not enable such improvements.

Research Questions

1. Does learning from paintings improve model robustness?
    - YES – improves robustness to image corruptions.
    - Cost-effective compared to only photos.
    - More abstract art forms do not enable such improvements.
2. Does learning from stylized images improve model robustness?

Does learning from stylized images improve robustness?

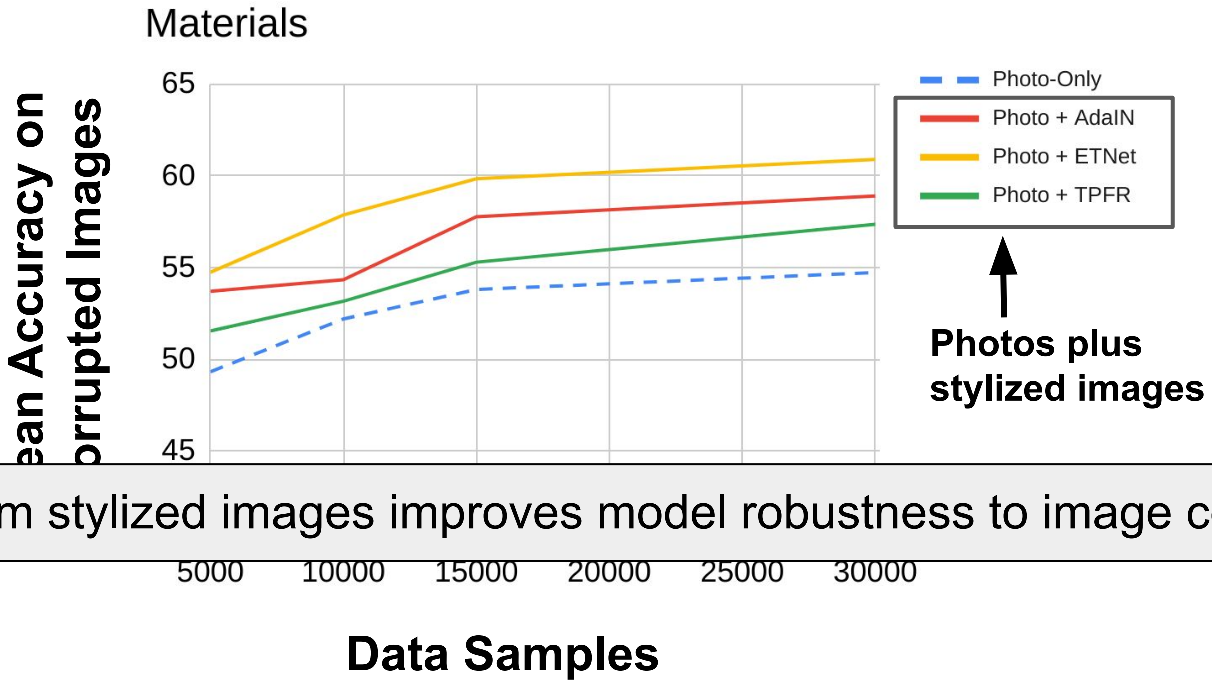

Learning from stylized images improves model robustness to image corruptions.

AdaIN: Huang and Belongie 2017; ETNet: Song et al 2019; TPFR: Svoboda et al 2020

Do stylized images need painting styles?

Arbitrary style transfer applies style from a source image to a target image.

- Do style source images need to be paintings?

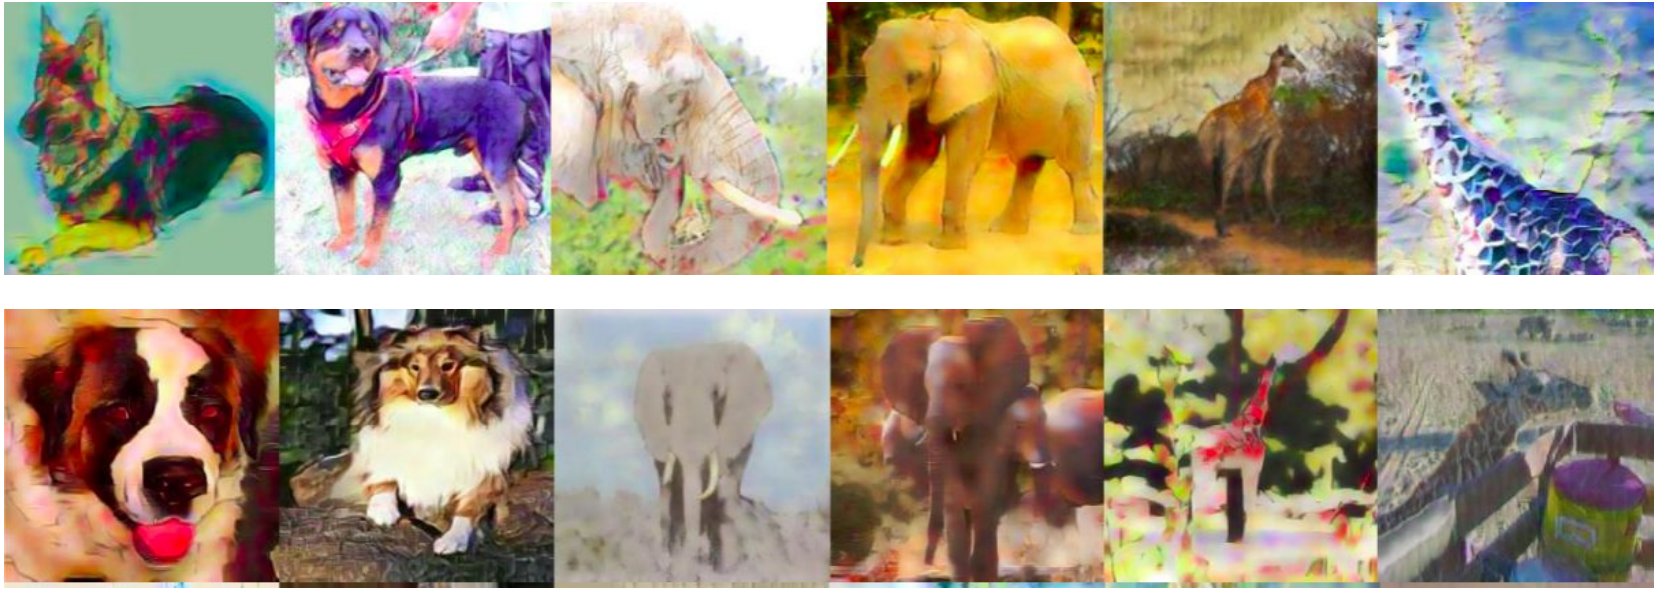

TOP: Photos stylized by paintings. **BOTTOM**: Photos stylized by photos.

Do stylized images need painting styles?

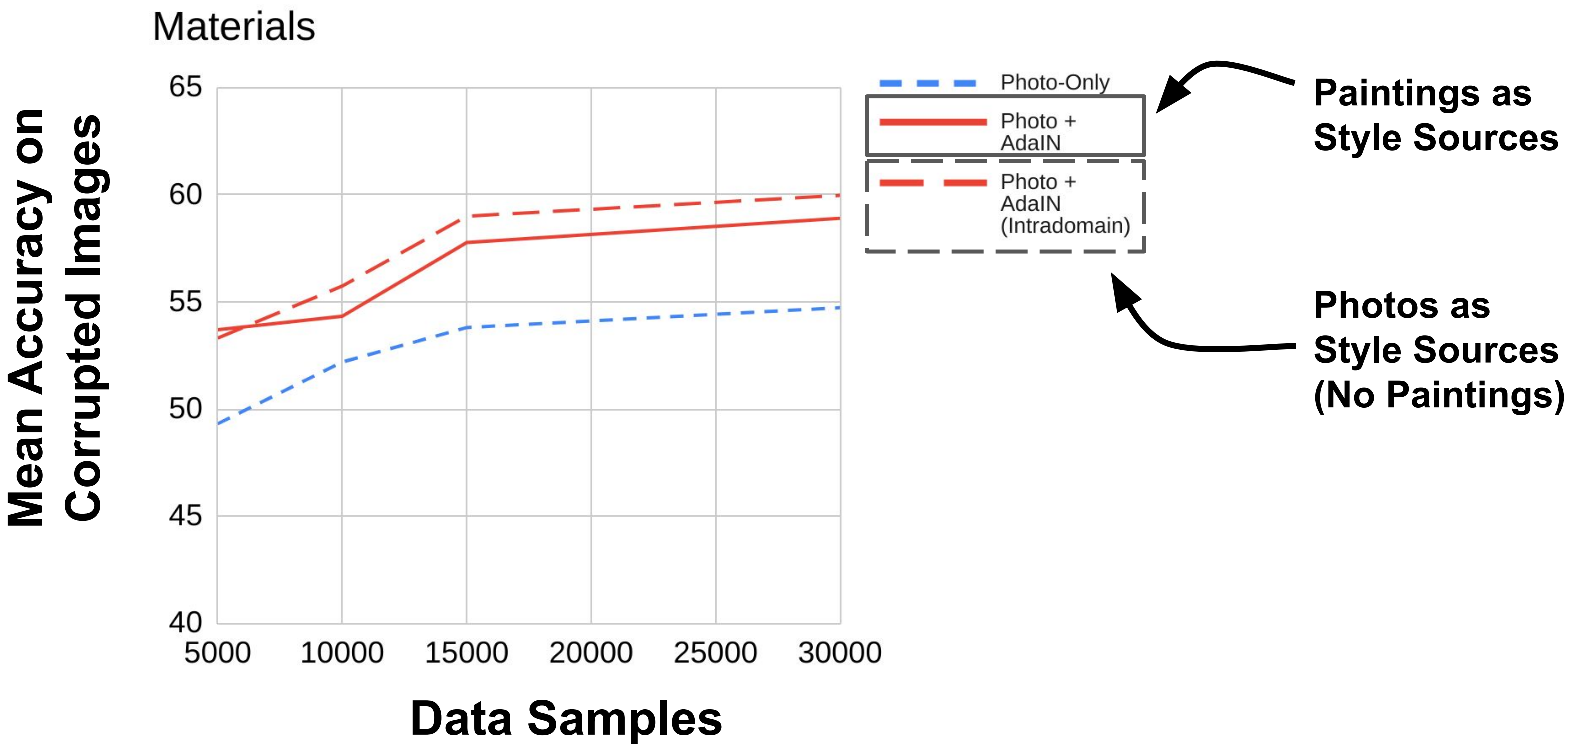

Do stylized images need painting styles?

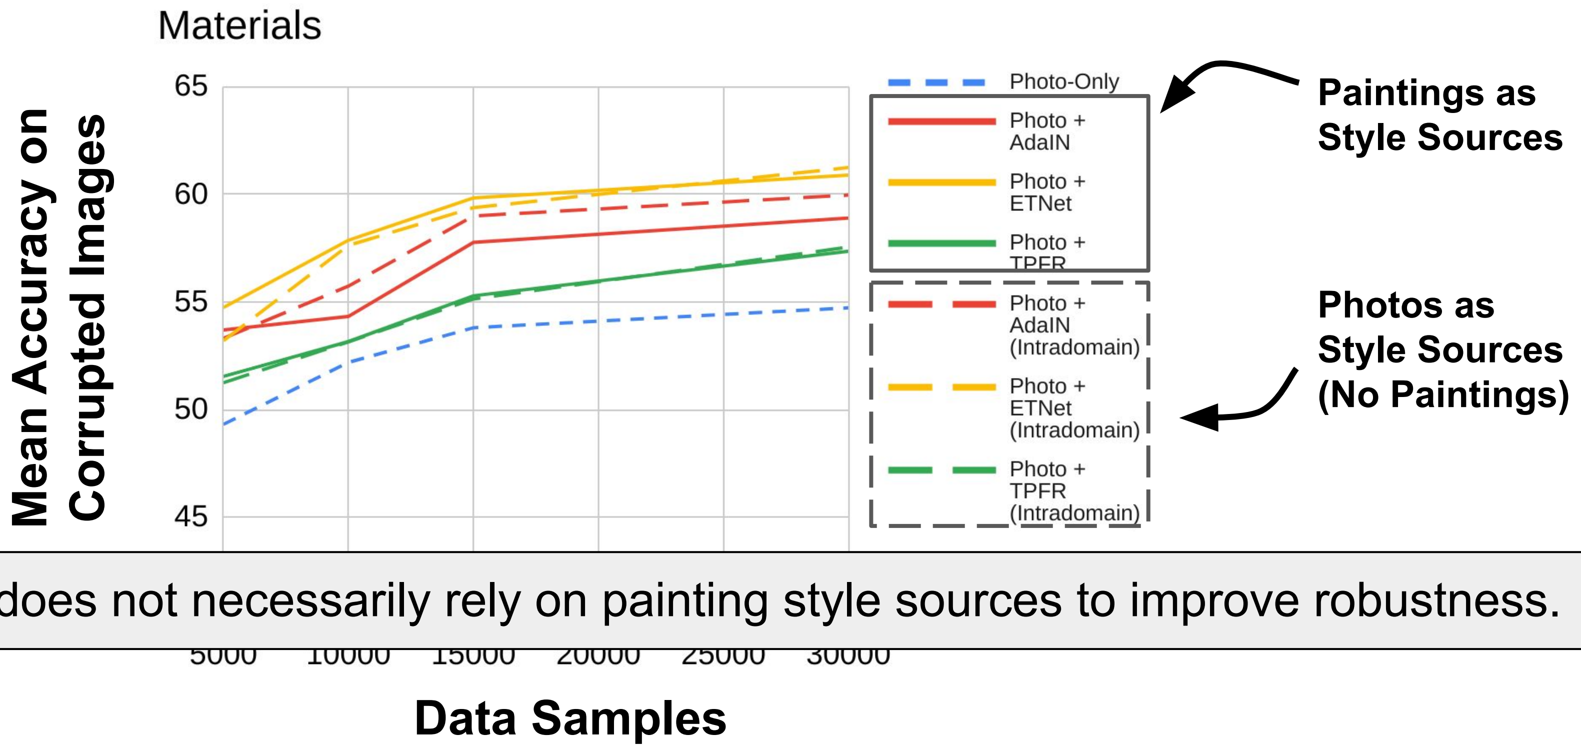

Stylization does not necessarily rely on painting style sources to improve robustness.

Research Questions

1. Does learning from paintings improve model robustness?
    - YES – improves robustness to image corruptions.
    - Cost-effective compared to only photos.
    - More abstract art forms do not enable such improvements.
2. Does learning from stylized images improve model robustness?
    - YES – improves robustness to image corruptions.
    - Does not necessarily require painting styles.

Research Questions

1. Does learning from paintings improve model robustness?
   - YES – improves robustness to image corruptions.
   - Cost-effective compared to only photos.
   - More abstract art forms do not enable such improvements.
2. Does learning from stylized images improve model robustness?
   - YES – improves robustness to image corruptions.
   - Does not necessarily require painting styles.
3. How do models trained on paintings differ from models trained on stylized images?

How do paintings and stylized images differ?

Method	**Accuracy (Image Corruptions)**	**Accuracy (OOD Photos)**
Materials (30K samples / domain)		
Photos-only	54.73±0.25	41.33±0.62
Photos + Stylization		
Photos + Paintings		
PACS (1.5K samples / domain)		
Photos-only	76.16±0.34	82.57±0.00
Photos + Stylization		
Photos + Paintings		

How do paintings and stylized images differ?

Method	**Accuracy (Image Corruptions)**	**Accuracy (OOD Photos)**
Materials (30K samples / domain)		
Photos-only	54.73±0.25	41.33±0.62
Photos + Stylization	**62.67**±0.03	34.54±0.91
Photos + Paintings		
PACS (1.5K samples / domain)		
Photos + Stylization	**87.27**±0.10	77.43±0.84
Photos + Paintings		

Stylization improves robustness to image corruptions, but hurts view generalization.

Stylized images have diverse textures, but same background contexts and views

Diverse textures: helps against image corruptions
Same background and views: hurts against new views

How do paintings and stylized images differ?

Method	Accuracy (Image Corruptions)	Accuracy (OOD Photos)
Materials (30K samples / domain)		
Photos-only	54.73±0.25	41.33±0.62
Photos + Stylization	**62.67**±0.03	34.54±0.91
Photos + Paintings	57.92±0.09	**43.92**±0.47
PACS (1.5K samples / domain)		
Photos + Stylization	**87.27**±0.10	77.43±0.84
Photos + Paintings	79.65±0.49	**85.43**±0.70

Paintings improve robustness to both image corruptions and novel views.

Paintings have diverse textures, and have ambiguous backgrounds

Diverse textures: helps against image corruptions.
Ambiguous background, focused on foreground: helps against new views.

Why stylization > paintings against image noise?

Similar textures, but stylization much better.

Reasons?

1. Corrupted images are share similar background and views to training; model uses these features.

2. Invisible high frequency textures?

Method	Accuracy (Image Corruptions)
Materials (30K samples / domain)	
Photos-only	54.73±0.25
Photos + Stylization	**62.67**±0.03
Photos + Paintings	57.92±0.09
PACS (1.5K samples / domain)	
Photos-only	76.16±0.34
Photos + Stylization	**87.27**±0.10
Photos + Paintings	79.65±0.49

Why stylization > paintings against image noise?

Image Power Spectrum

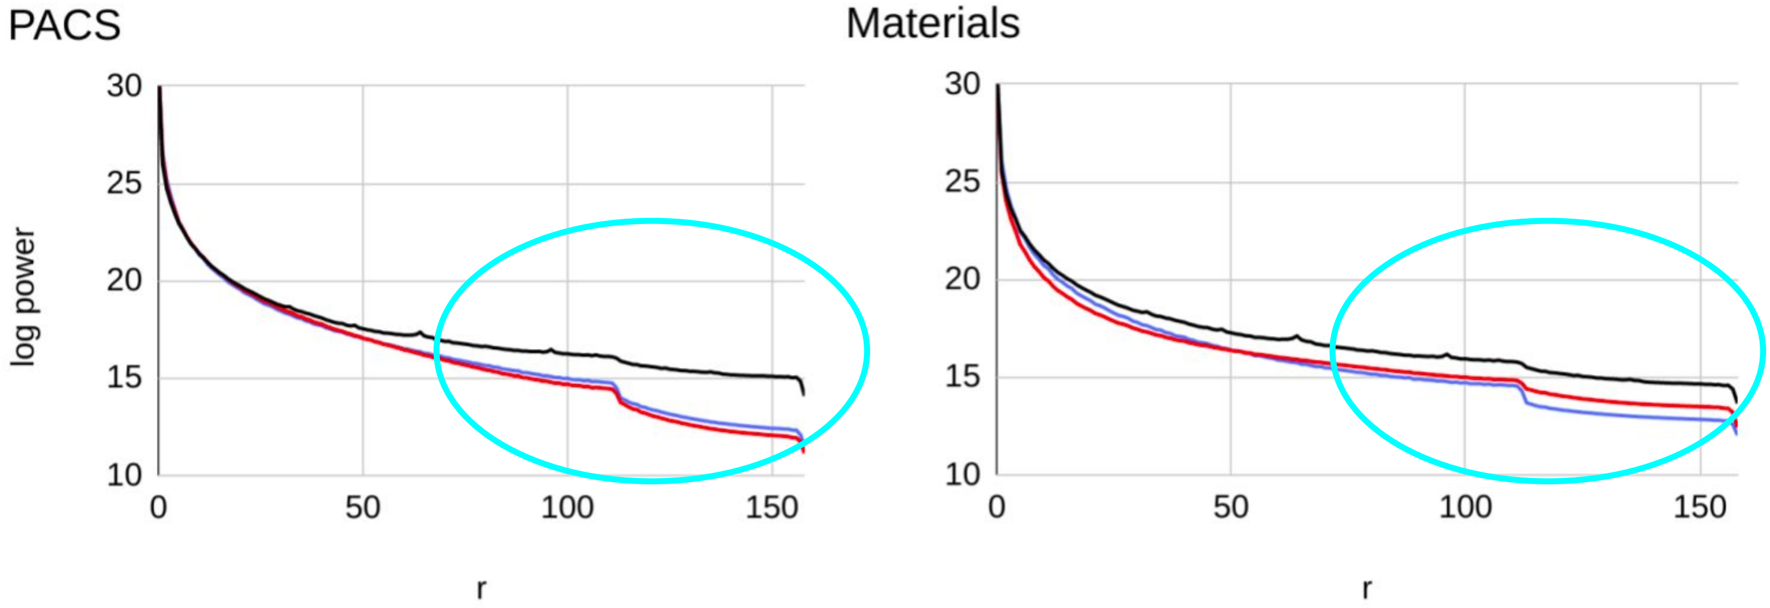

BLACK: Stylized Images. **RED:** Paintings. **BLUE:** Photos.

Stylized images contain larger magnitude high frequency components.

Why stylization > paintings against image noise?

Original Image

Low Frequency Only

Why stylization > paintings against image noise?

Method	**Accuracy (Images with Noise)**	
	Materials (30K samples / domain)	*PACS (1.5K samples / domain)*
Photos + Stylization	61.87±0.16	85.98±0.56
Photos + Stylization (Low Freq. Images)	45.82±1.36	77.55±2.60
Photos + Paintings	49.82±0.56	68.83±0.83
Photos + Paintings (Low Freq. Images)	44.95±0.66	71.16±1.31

Stylized images contain imperceptible high-frequency signals that greatly improve noise robustness.

Are paintings and stylized images complementary?

Method	**Accuracy (Image Corruptions + OOD Photos)**	
	Materials (30K samples / domain)	*PACS (1.5K samples / domain)*
Photos-only	48.03±0.21	79.37±0.17
Photos + Stylization		
Photos + Paintings		
Photos + Stylization + Paintings		

Are paintings and stylized images complementary?

Method	**Accuracy (Image Corruptions + OOD Photos)**	
	Materials (30K samples / domain)	*PACS (1.5K samples / domain)*
Photos-only	48.03±0.21	79.37±0.17
Photos + Stylization	48.56±0.45	82.35±0.37
Photos + Paintings	50.92±0.22	82.54±0.59
Photos + Stylization + Paintings		

Are paintings and stylized images complementary?

Method	**Accuracy (Image Corruptions + OOD Photos)**	
	Materials (30K samples / domain)	*PACS (1.5K samples / domain)*
Photos-only	48.03±0.21	79.37±0.17
Photos + Stylization	48.56±0.45	82.35±0.37
Photos + Paintings	50.92±0.22	82.54±0.59
Photos + Stylization + Paintings	**51.49**±0.69	**85.42**±0.18

Models learn complementary invariances from paintings and stylization.

Research Questions

1. Does learning from stylized images improve model robustness?
   - YES – improves robustness to image corruptions.
   - Does not necessarily require painting styles.
2. Does learning from paintings improve model robustness?
   - YES – improves robustness to image corruptions.
   - Cost-effective compared to only photos.
   - More abstract art forms do not enable such improvements.
3. How do models trained on paintings differ from models trained on stylized images?
   - Stylized images greatly improve robustness to corruptions, but hurts generalization to new views. Paintings improve robustness to both.
   - Stylized images contain imperceptible noise that improve robustness.

Key Findings: Learning Robust Invariances from Paintings

How can paintings help our perception models?

1. Artists emphasize cues like highlights to help viewers understand scenes.
2. Models trained on paintings may learn to use more interpretable cues.
3. Models trained on paintings are robust to image corruptions and novel views.
4. “Fake” paintings produced by style transfer greatly strengthen model robustness to noise while harming generalization to novel views.
5. Learning from both paintings and stylized images allow models to learn useful complementary invariances that boost robustness overall.

Talk Outline

Many challenges in improving perception systems in real world.

1. **Better annotation tools.**
   [ICCV 2019]

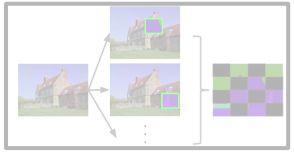

2. **Learning robust visual invariances.**
   [PLOS One 2021, FAPER ICPR 2020, CVPR 2021]

3. **Reasoning about perception uncertainties.**
   [ICRA 2020]

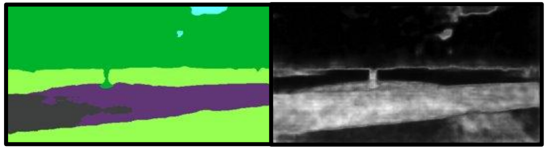

4. **Summary.**

Toy problem: Solved?

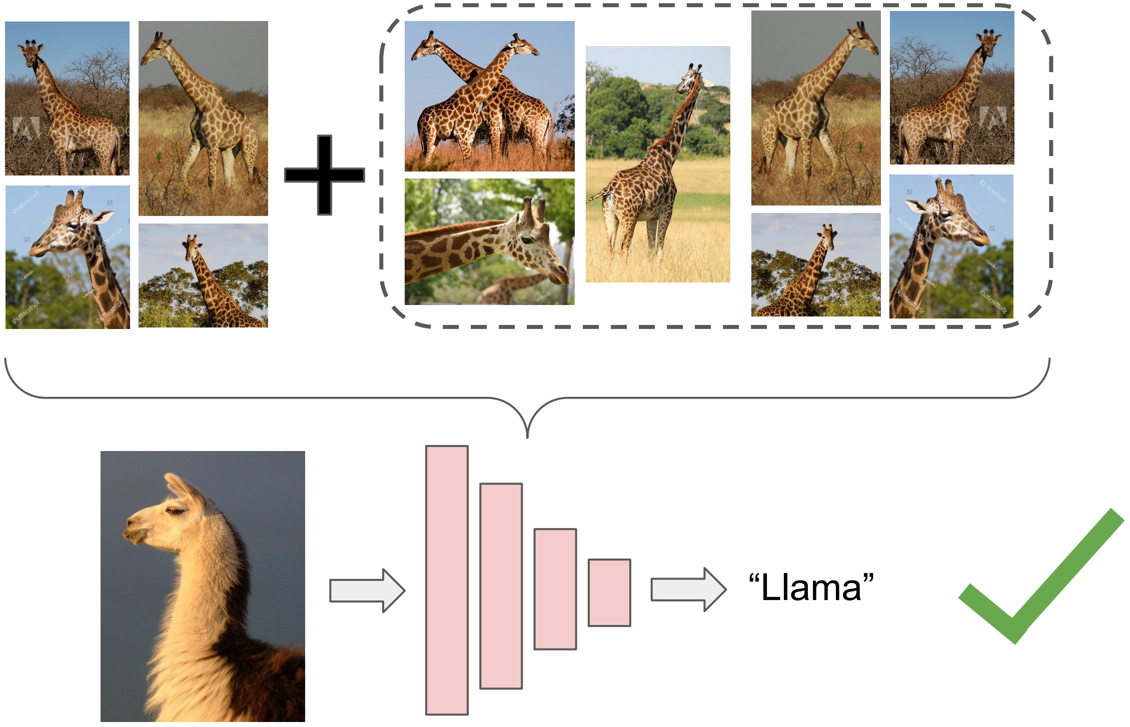

Toy problem: Solved?

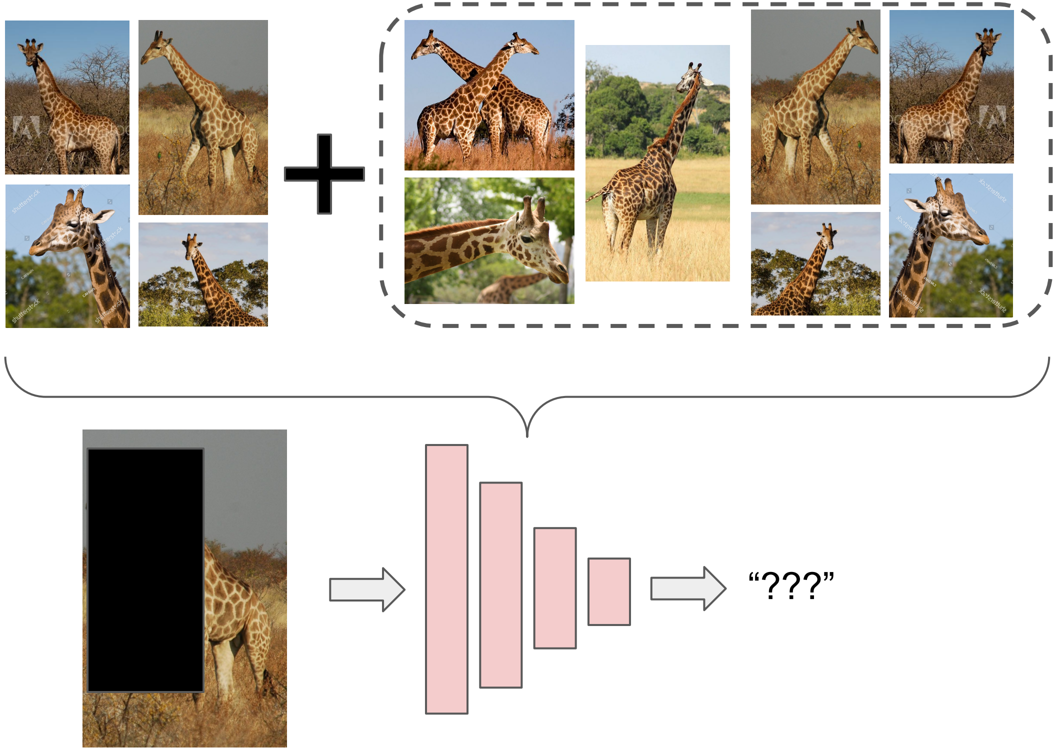

Toy problem: Solved?

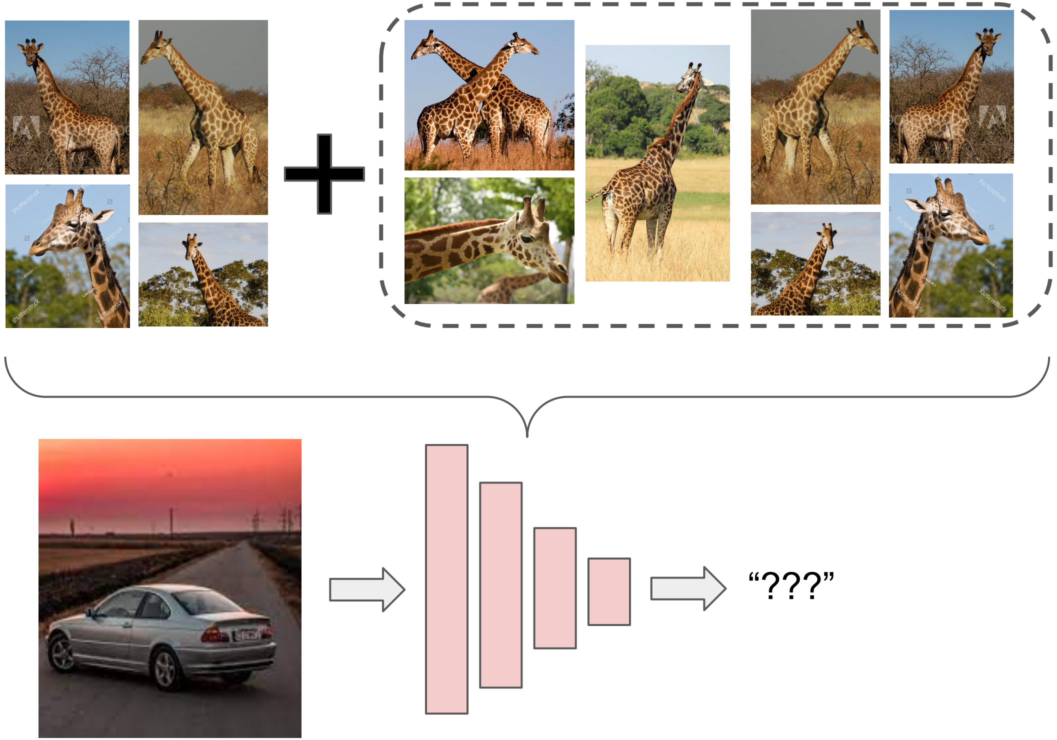

Modeling Uncertainty

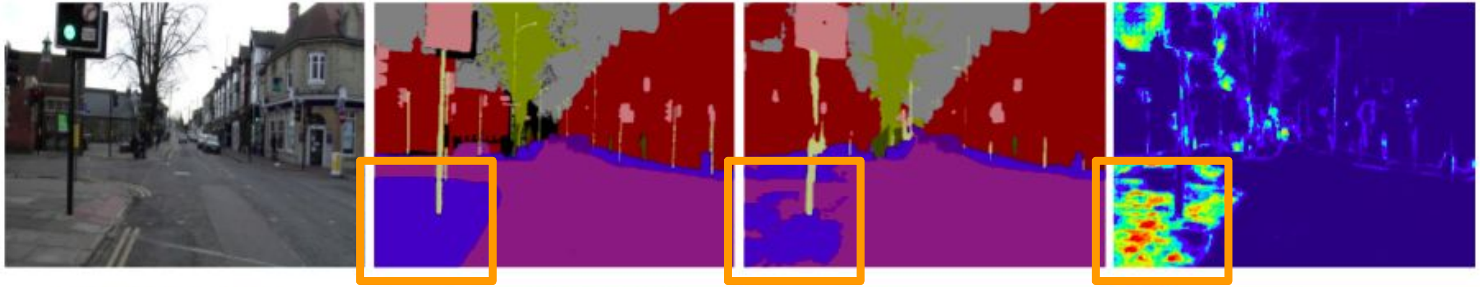

(a) Input Image (b) Ground Truth (c) Semantic Segmentation (e) Epistemic Uncertainty

Incorrect predictions in regions with high model uncertainty

Kendall and Gal 2017

Unstructured Real World Navigation

Han*, <u>Lin*</u>**,** Banfi*, Bala, Campbell, ***"DeepSemanticHPPC: Hypothesis-based Planning over Uncertain Semantic Point Clouds"****, ICRA 2020*

Environment Semantics and Uncertainty

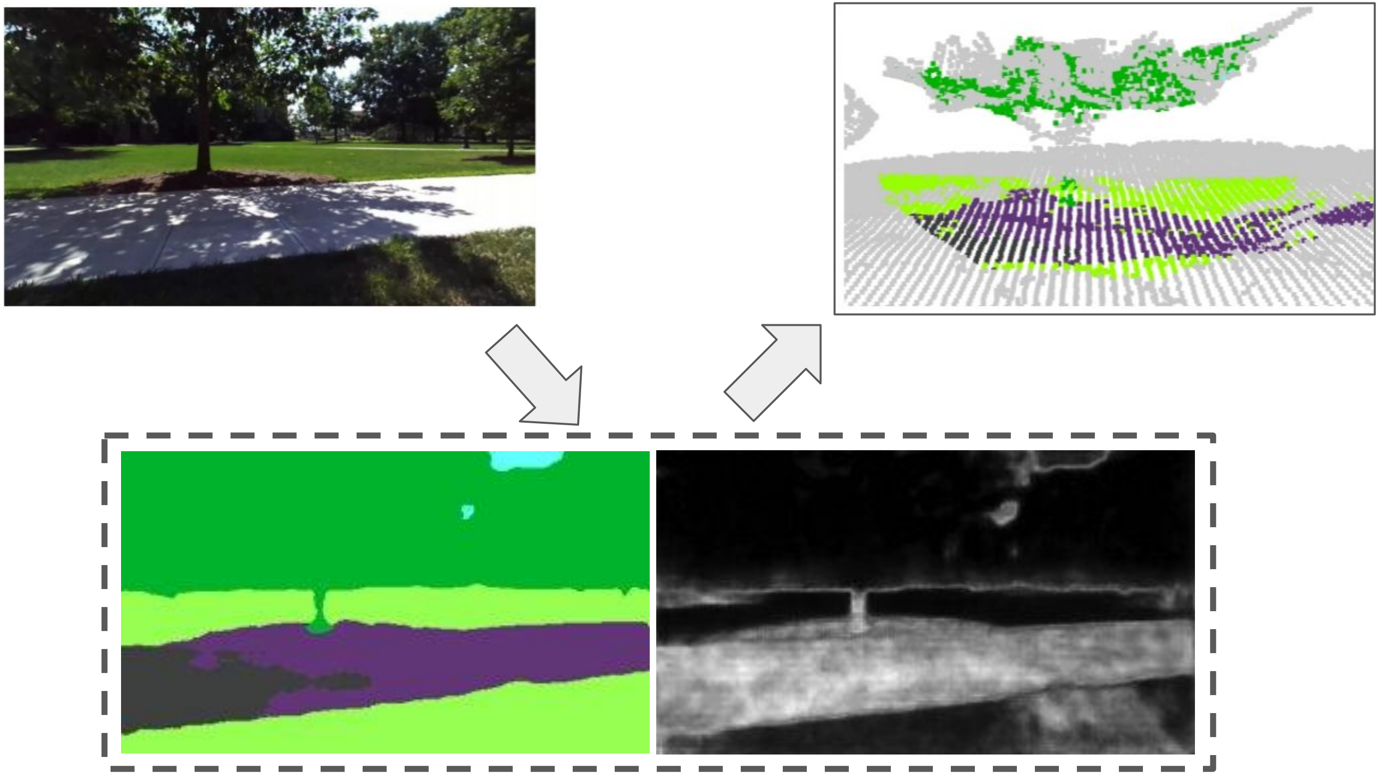

Navigation Task

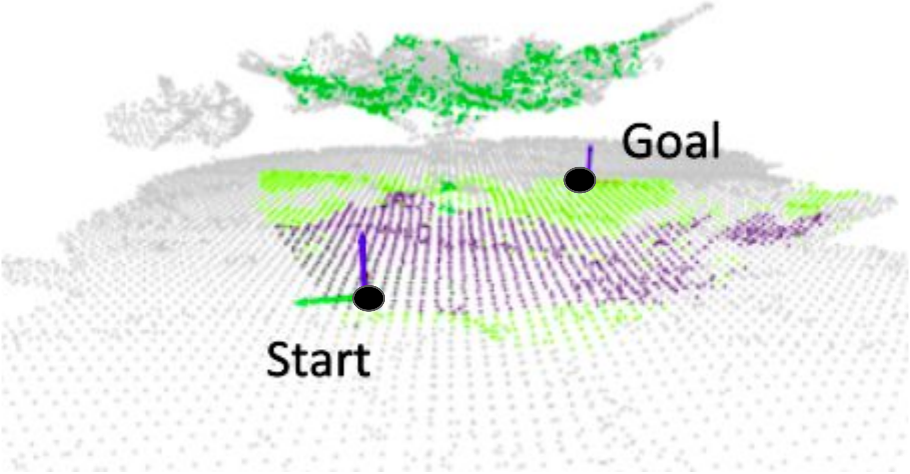

Planned Paths

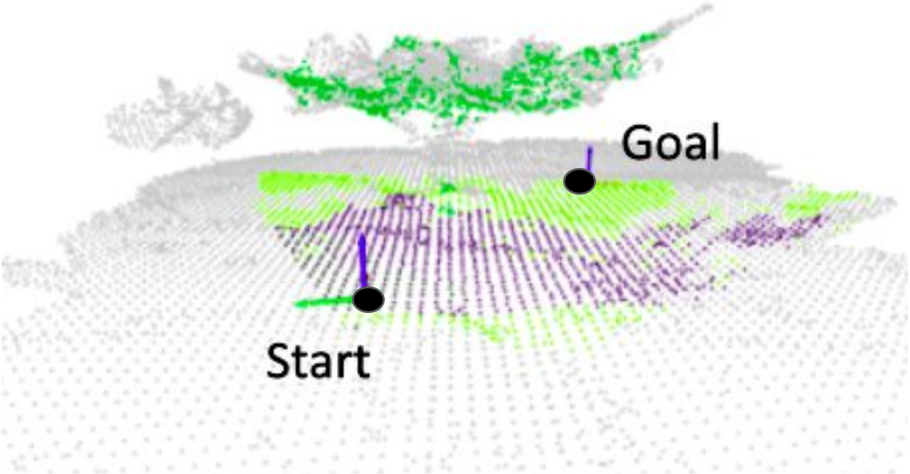

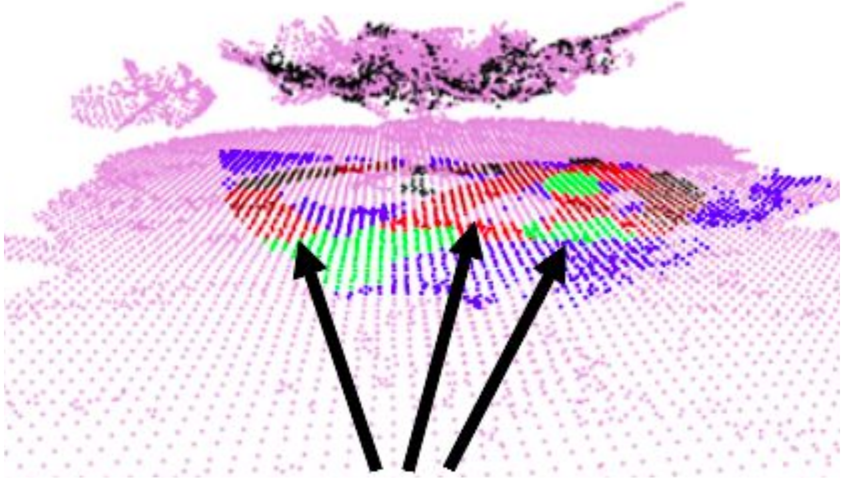

Uncertain or unsafe regions in potential paths

Uncertainty Reduction

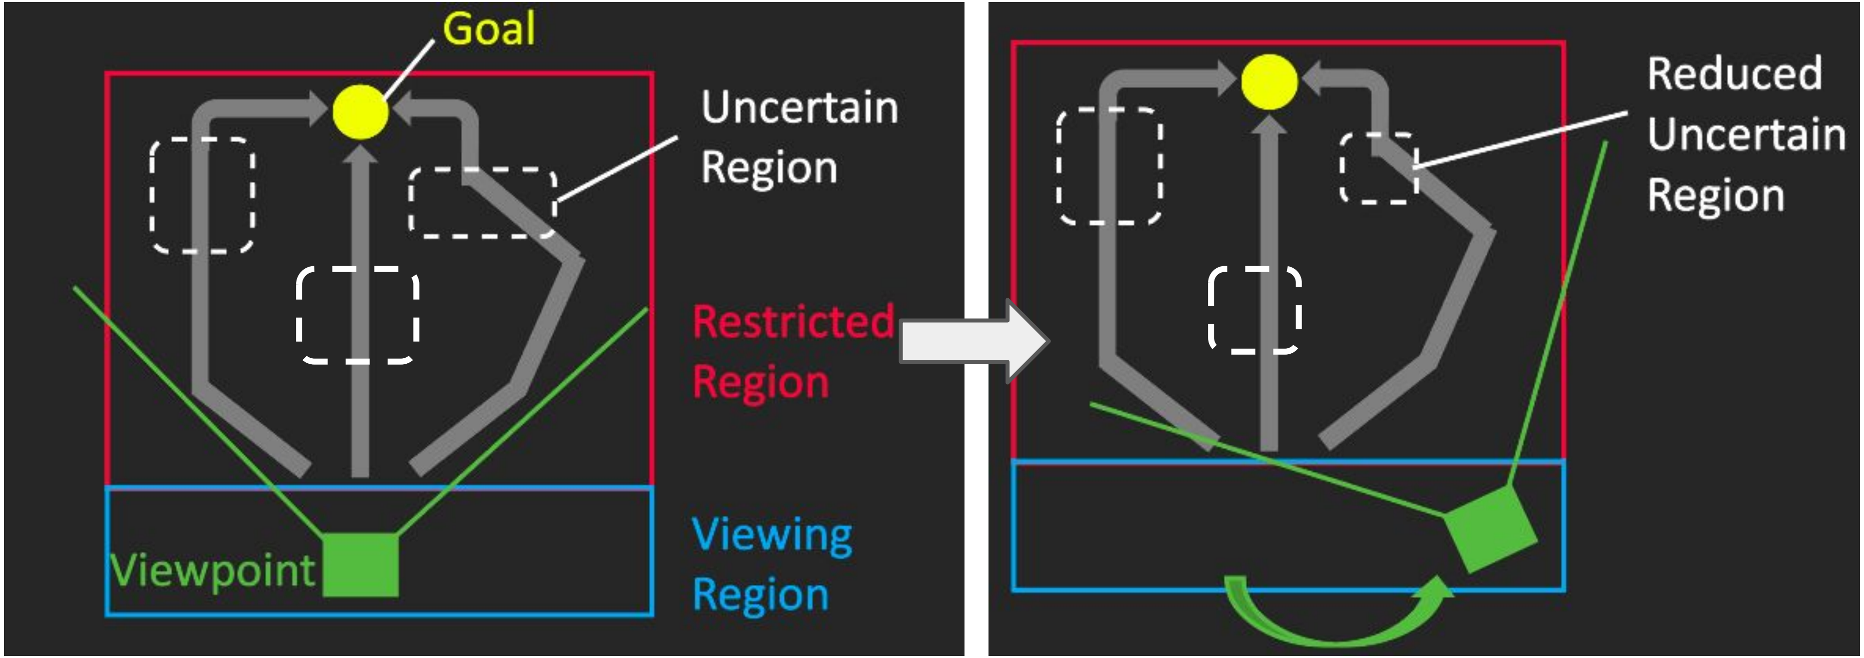

Current Measurement

Updated Measurement

Real World Test Environments

Key Findings: Unstructured Real World Navigation

Reasoning about **semantics with uncertainty** allows higher path safety than (B1) only geometry and (B2) semantics without uncertainty reduction.

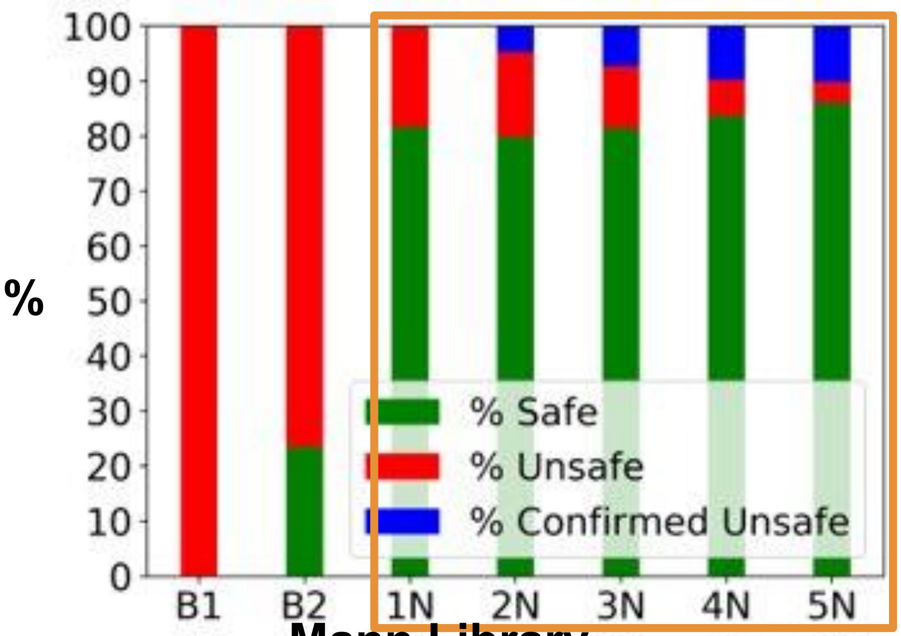

Mann Library

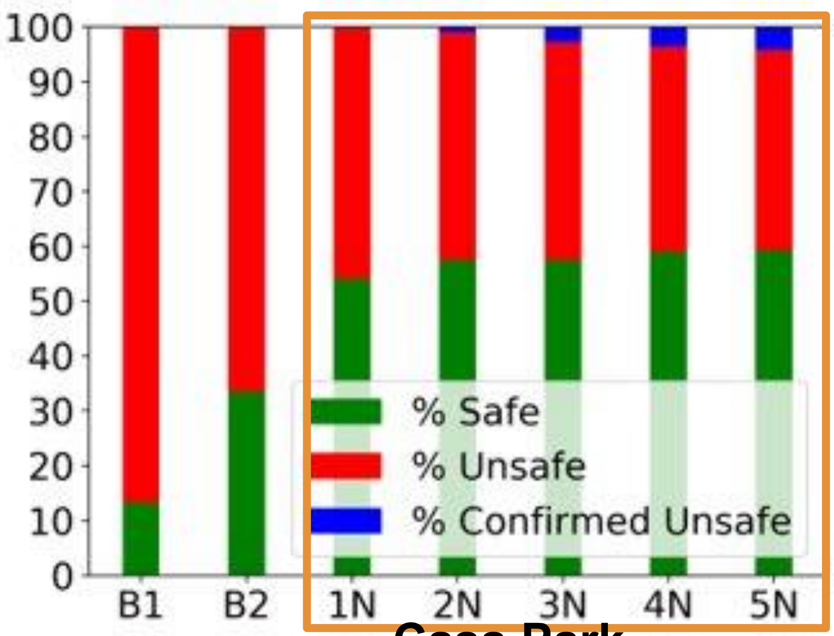

Cass Park

Talk Outline

Many challenges in improving perception systems in real world.

1. **Better annotation tools.**
   [ICCV 2019]

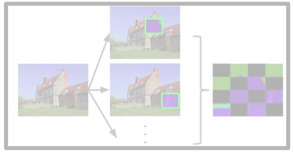

2. **Learning robust visual invariances.**
   [PLOS One 2021, FAPER ICPR 2020, CVPR 2021]

3. **Reasoning about perception uncertainties.**
   [ICRA 2020]

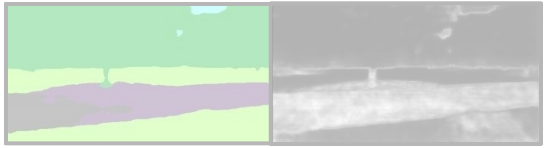

4. **Summary.**

Summary

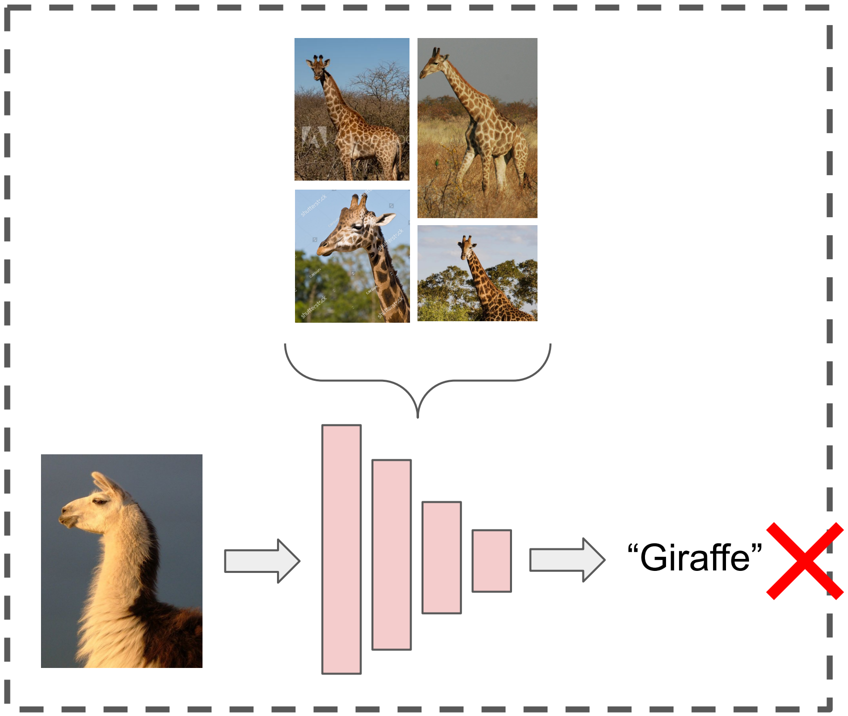

Training data and learned models are imperfect

Efficiently annotate more data

Encourage model to learn robust invariances

Reason about perception uncertainties when using model in real world

Future Directions

1. Better annotation tools.
    a. Which images or image regions to label?

2. Learning robust invariances from paintings.
    a. Improved style transfer algorithms.
    b. Implications for synthetic data in computer vision – physical realism goal?
    c. Better methods for learning from paintings – domain generalization methods fail.

3. Reasoning about perception failures.
    a. Combine with online adaptation and continual learning.

Acknowledgements

Research group and collaborators from 2016 to 2022


Kavita Bala

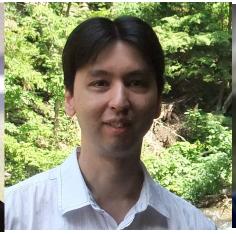
Paul Upchurch


Scott Wehrwein


Balazs Kovacs

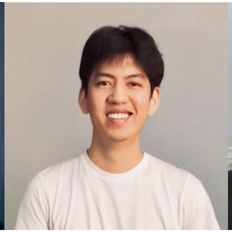
Fujun Luan


Utkarsh Mall


Hadi AlZayer

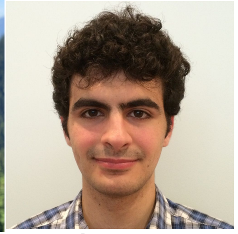
Aaron Gokaslan


Mitchell Van Zuijlen


Maarten Wijntjes


Sylvia Pont

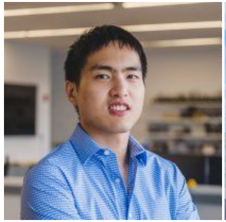
Yutao Han


Jacopo Banfi


Mark Campbell

Funding: NSERC, NSF, ONR

Thank you!

References

Lin, Upchurch, Bala, ***"Block Annotation: Better Image Annotation with Sub-Image Decomposition"****, ICCV 2019*

Van Zuijlen, Lin, Bala, Pont, Wijntjes, ***"Materials In Paintings (MIP): An interdisciplinary dataset for perception, art history, and computer vision"****, PLOS One 2021*

Lin, Van Zuijlen, Wijntjes, Pont, Bala, ***"Insights from a Large-Scale Database of Material Depictions in Paintings"****, FAPER ICPR 2020*

Lin, Van Zuijlen, Wijntjes, Pont, Bala, ***"What Can Style Transfer and Paintings Do For Model Robustness?"****, CVPR 2021*

Han*, Lin*, Banfi*, Bala, Campbell, ***"DeepSemanticHPPC: Hypothesis-based Planning over Uncertain Semantic Point Clouds"****, ICRA 2020*

Removed + Backup Slides…

Painting + Style Robustness Backup Slides

Learning from Stylized Images

What do models learn from stylized images ("fake" paintings)?

- More shape-based decisions, similar to humans.

Cat shape with
elephant texture

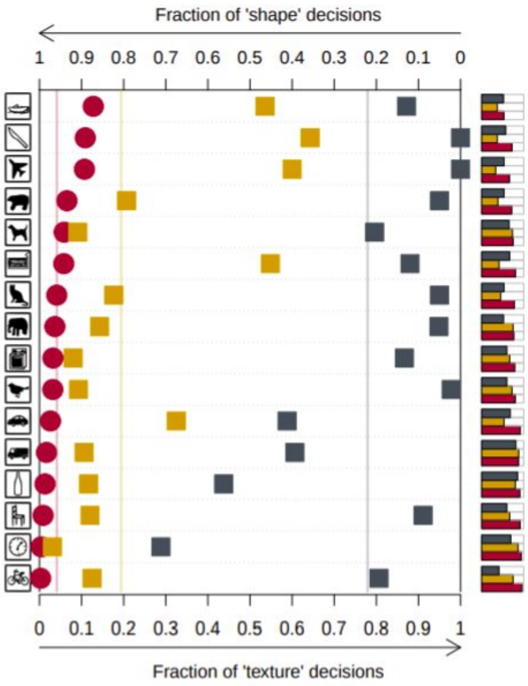

Geirhos et al 2019

Style Semantic Diversity vs Robustness

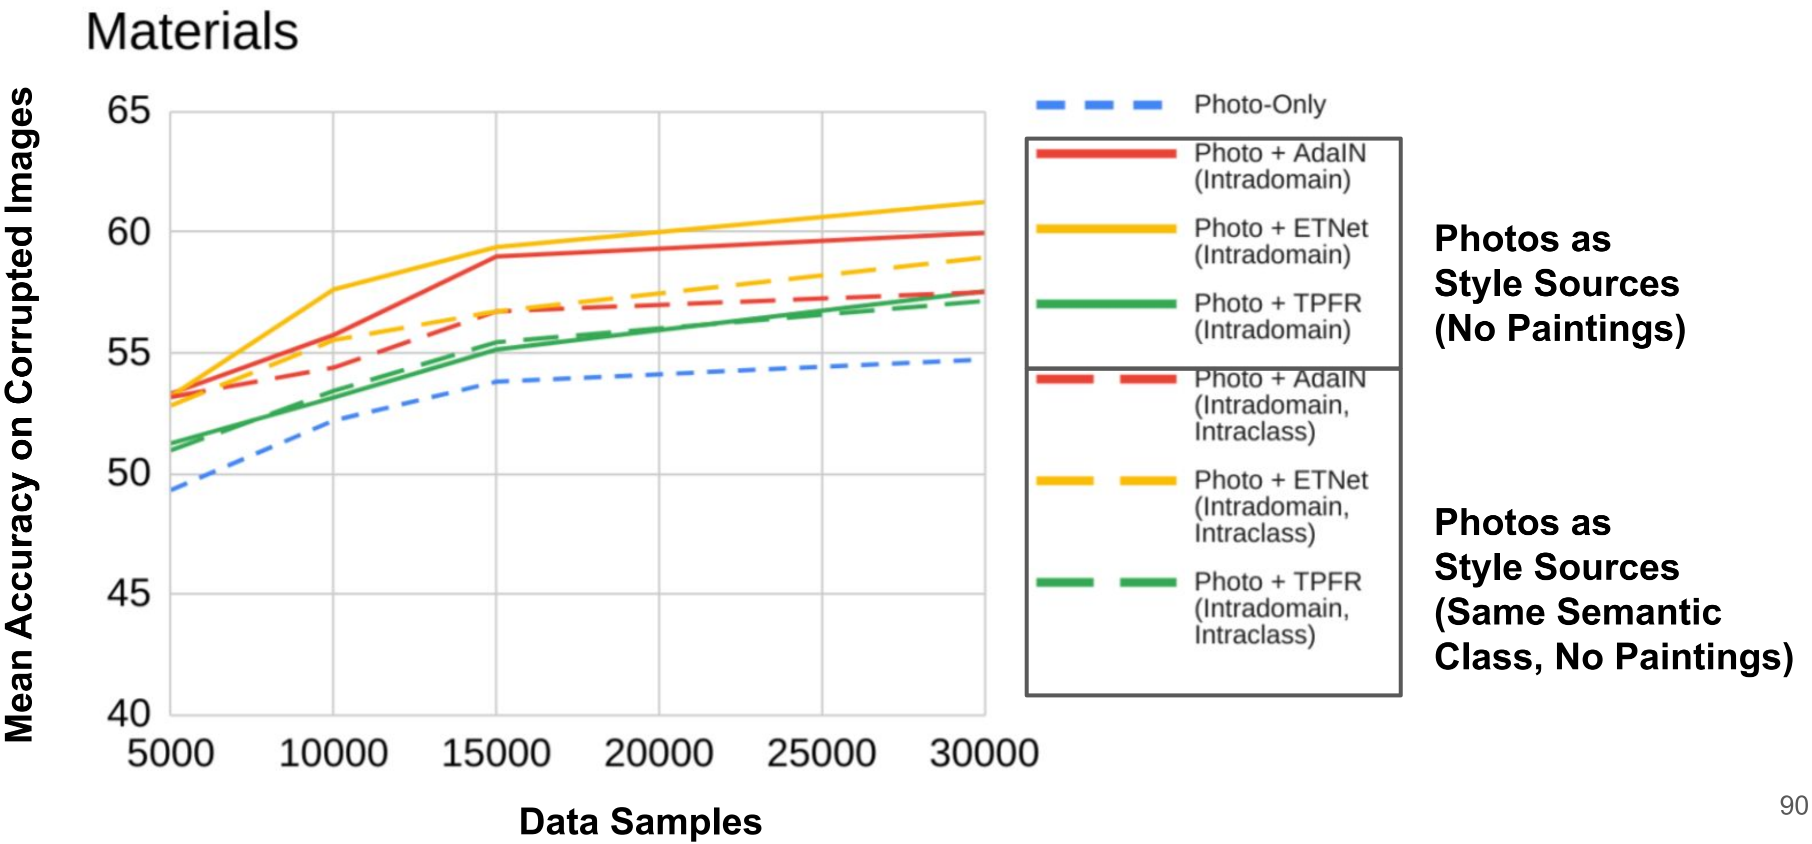

Style Strength vs Robustness

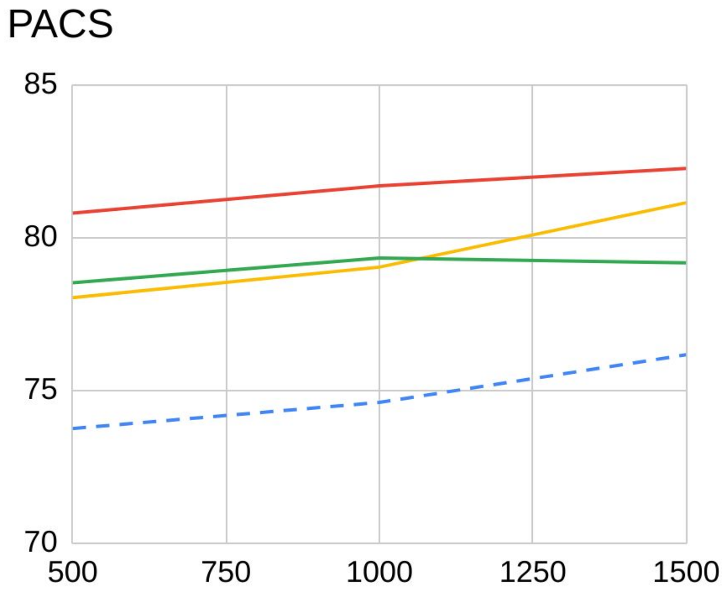

Method	Painting	Intradomain	Intradomain (Intraclass)
AdaIN	1.58±0.93	1.28±0.79	1.16±0.85
ETNet	2.33±1.09	2.13±1.04	1.81±1.03
TPFR	1.52±0.90	1.38±0.87	1.27±0.91

Table 6: **Style (Gram Matrix) Distance.** Gram matrices computed from ImageNet pretrained ResNet18 features on PACS. Mean distance between `(image, stylized image)` pairs is reported. ↑ distance implies ↑ style difference. ± denotes standard deviation across 1.5K pairs.

Stylization vs Paintings: Per-Corruption Accuracy

Method	Noise	Blur	Weather	Digital
Materials (30K Samples/Domain)				
Photo-Only	43.70±0.65	58.76±0.14	55.25±0.33	61.20±0.69
Photo + SACL	61.87±0.16	64.36±0.20	57.49±0.24	66.55±0.17
Photo + Painting	49.82±0.56	61.03±0.13	56.69±0.10	64.15±0.14
Photo+SACL (LF)	45.82±1.36	64.24±0.39	57.06±0.13	66.37±0.29
Photo+Painting (LF)	44.95±0.66	60.87±0.29	56.82±0.23	63.69±0.46
PACS (1.5K Samples/Domain)				
Photo-Only	62.64±1.48	72.75±0.04	83.24±0.22	86.33±0.14
Photo + SACL	85.98±0.56	84.61±0.15	89.73±0.33	88.74±0.48
Photo + Painting	68.83±0.83	75.80±0.95	86.88±0.66	87.07±0.14
Photo+SACL (LF)	77.55±2.60	85.4±0.11	88.93±0.22	88.53±0.15
Photo+Painting (LF)	71.16±1.31	75.97±0.71	86.82±0.37	87.35±0.36

Path Planning Backup Slides

Semantic Segmentation with Uncertainty

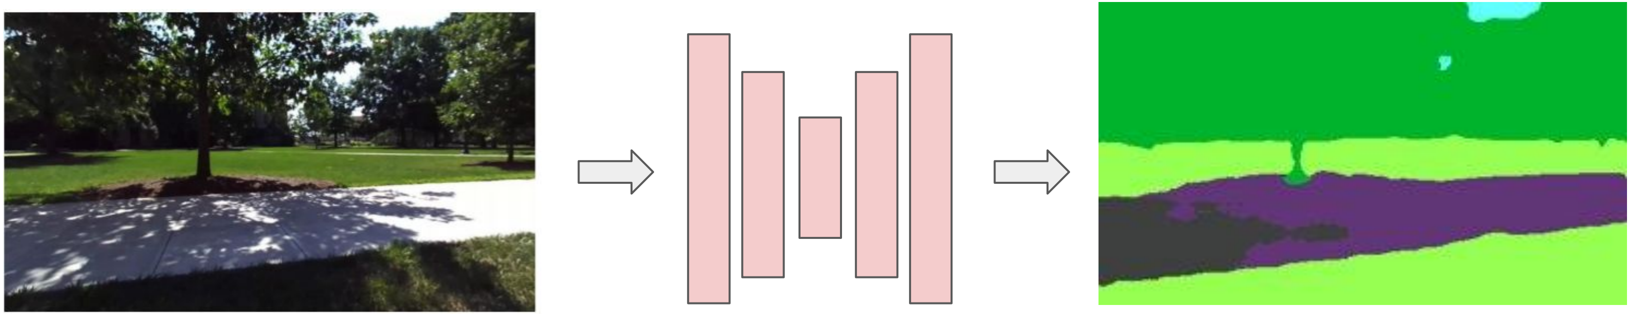

Deeplabv3+[1]

[1]Chen et al 2018

Semantic Segmentation with Uncertainty

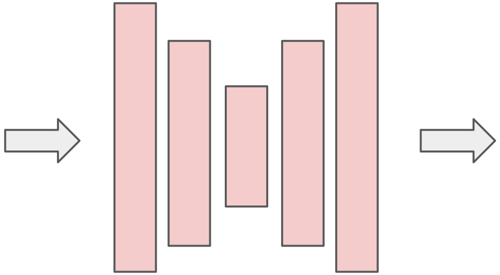

Deeplabv3+
w/ dropout

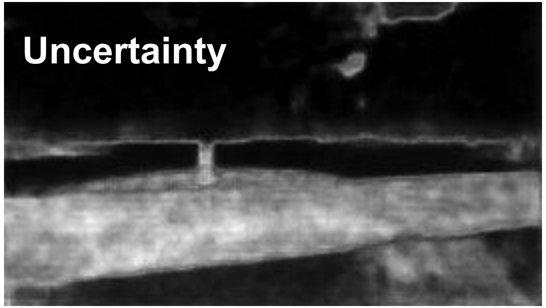

$$\mathbf{p}^{(i,j)x} = \frac{1}{T}\sum_{t=1}^{T} \mathbf{s}_t^{(i,j)x}(y|X)$$

$$\boldsymbol{\sigma}^{(i,j)x} = \sqrt{\frac{\sum_{t=1}^{T}\left(\mathbf{s}_t^{(i,j)x}(y|X) - \mathbf{p}^{(i,j)x}\right)^2}{T-1}}$$

Kendall and Gal 2017; Gal and Ghahramani 2015

Environment Semantics and Uncertainty

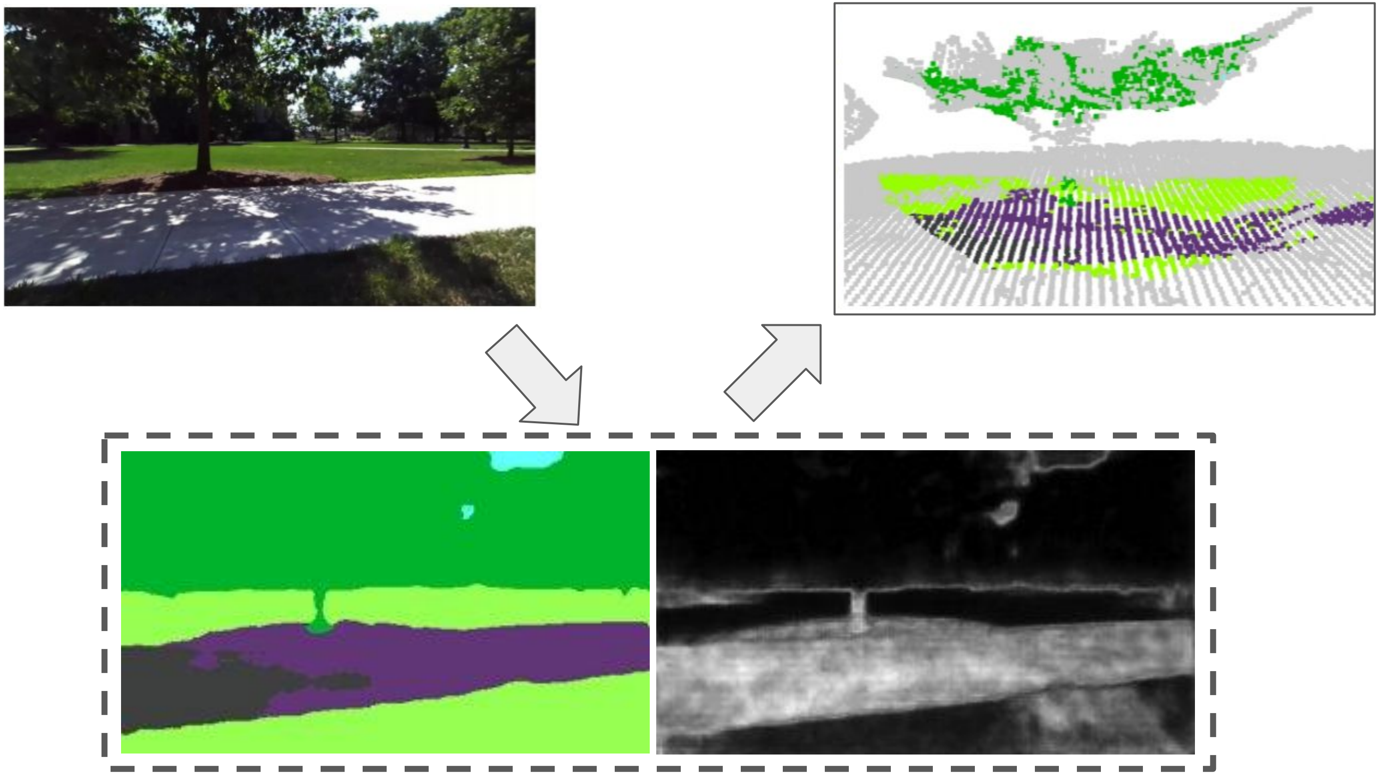

Navigation Task

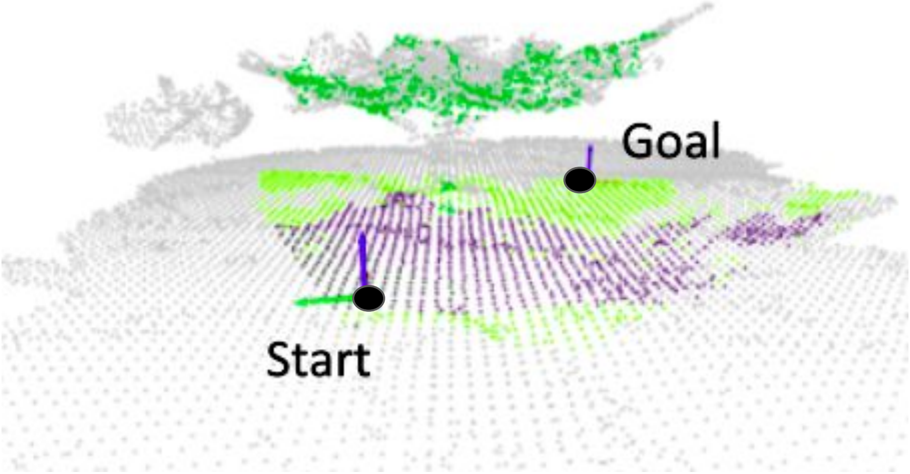

Multihypothesis Path Planner

Planner:

- RRT with sampling biased towards goal.
- Multiple paths: remove large regions after planning, and re-plan a new path.

Feasibility constraints: [Krusi et al 2017]

- Contact with the terrain surface.
- Static traversability (bounded roll and pitch angles).
- Kinematic constraints (motion primitives + bounded continuous curvature).

Multihypothesis Path Planner

Basic RRT:

RRT Visualization

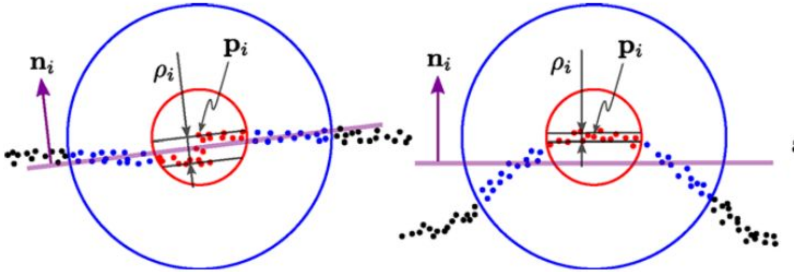

Terrain Smoothness

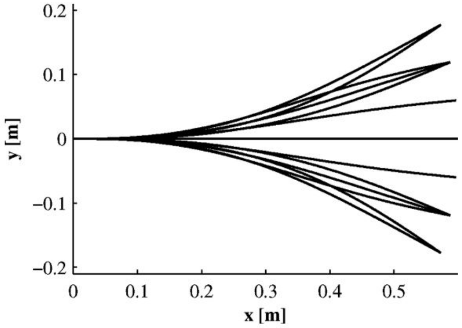

Motion Primitives

LaValle 1998 (Figure: Wikipedia); Krusi et al 2017

Planned Paths

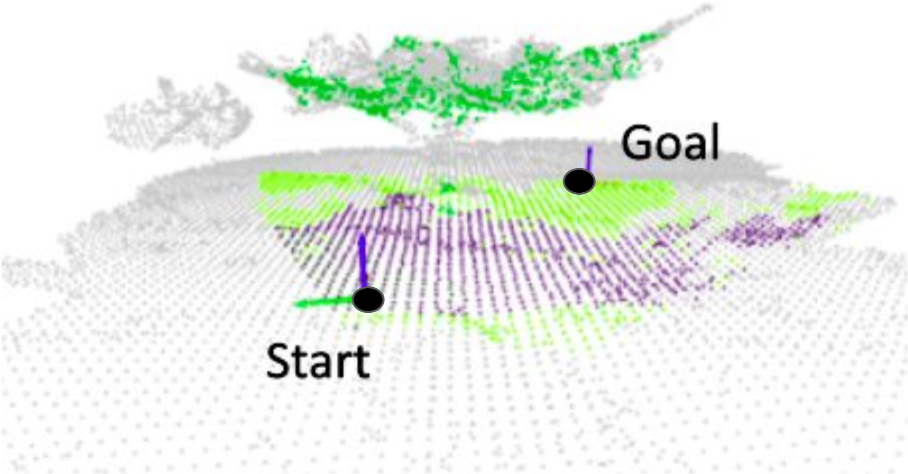

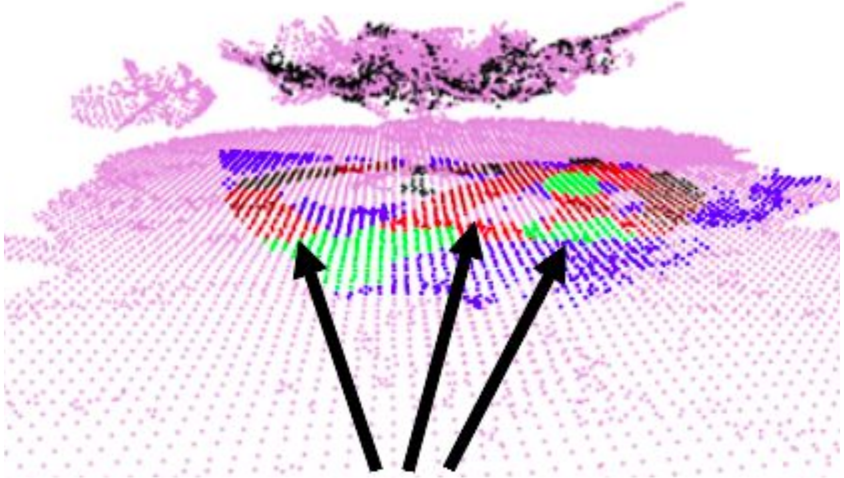

Uncertain or unsafe regions in potential paths

$$\begin{cases} \textbf{safe} \text{ if } p_S^i - w_\sigma \sigma^i \geq \theta_s; \\ \textbf{unsafe} \text{ if } p_U^i - w_\sigma \sigma^i \geq \theta_u; \\ \textbf{unclear} \text{ otherwise.} \end{cases}$$

Uncertainty Reduction

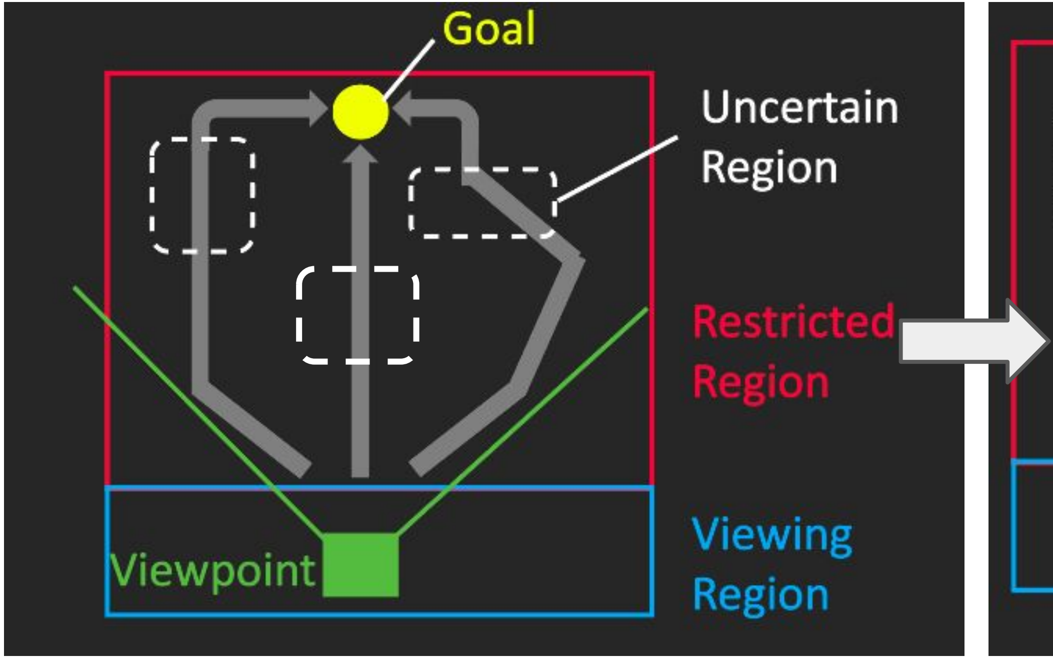

Current Measurement

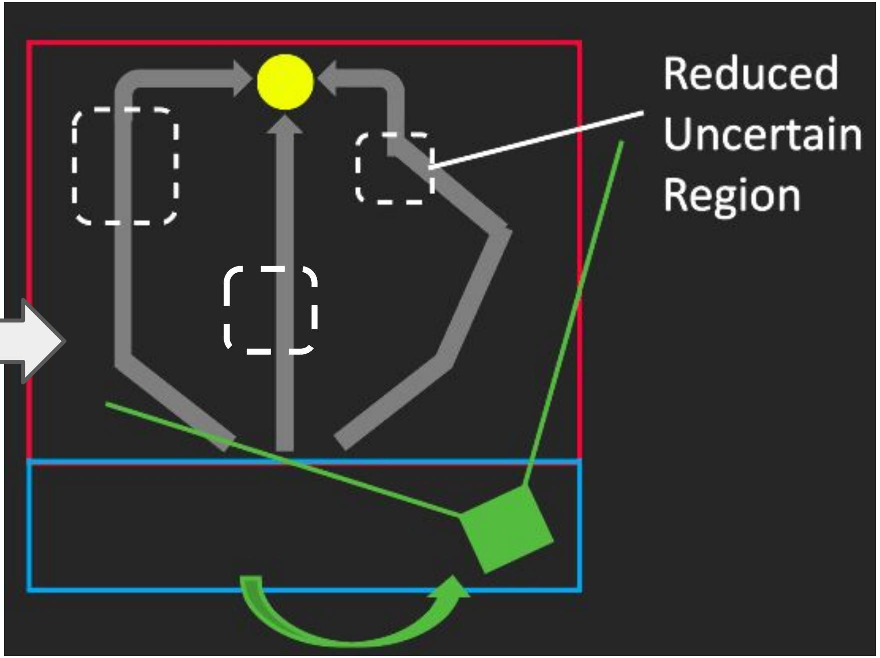

Updated Measurement

Uncertainty Reduction

How to select new viewpoint?

Camera pose heuristics:

- ↓ Distance to visible path nodes.
- ↑ Viewing angle (vs initial viewing orientation).
- ↑ Number of path nodes seen from view.

Uncertainty reduction heuristics:

- ↑ Uncertainty of visible path nodes.
- ↑ Pixel coverage visible path nodes projected onto view.

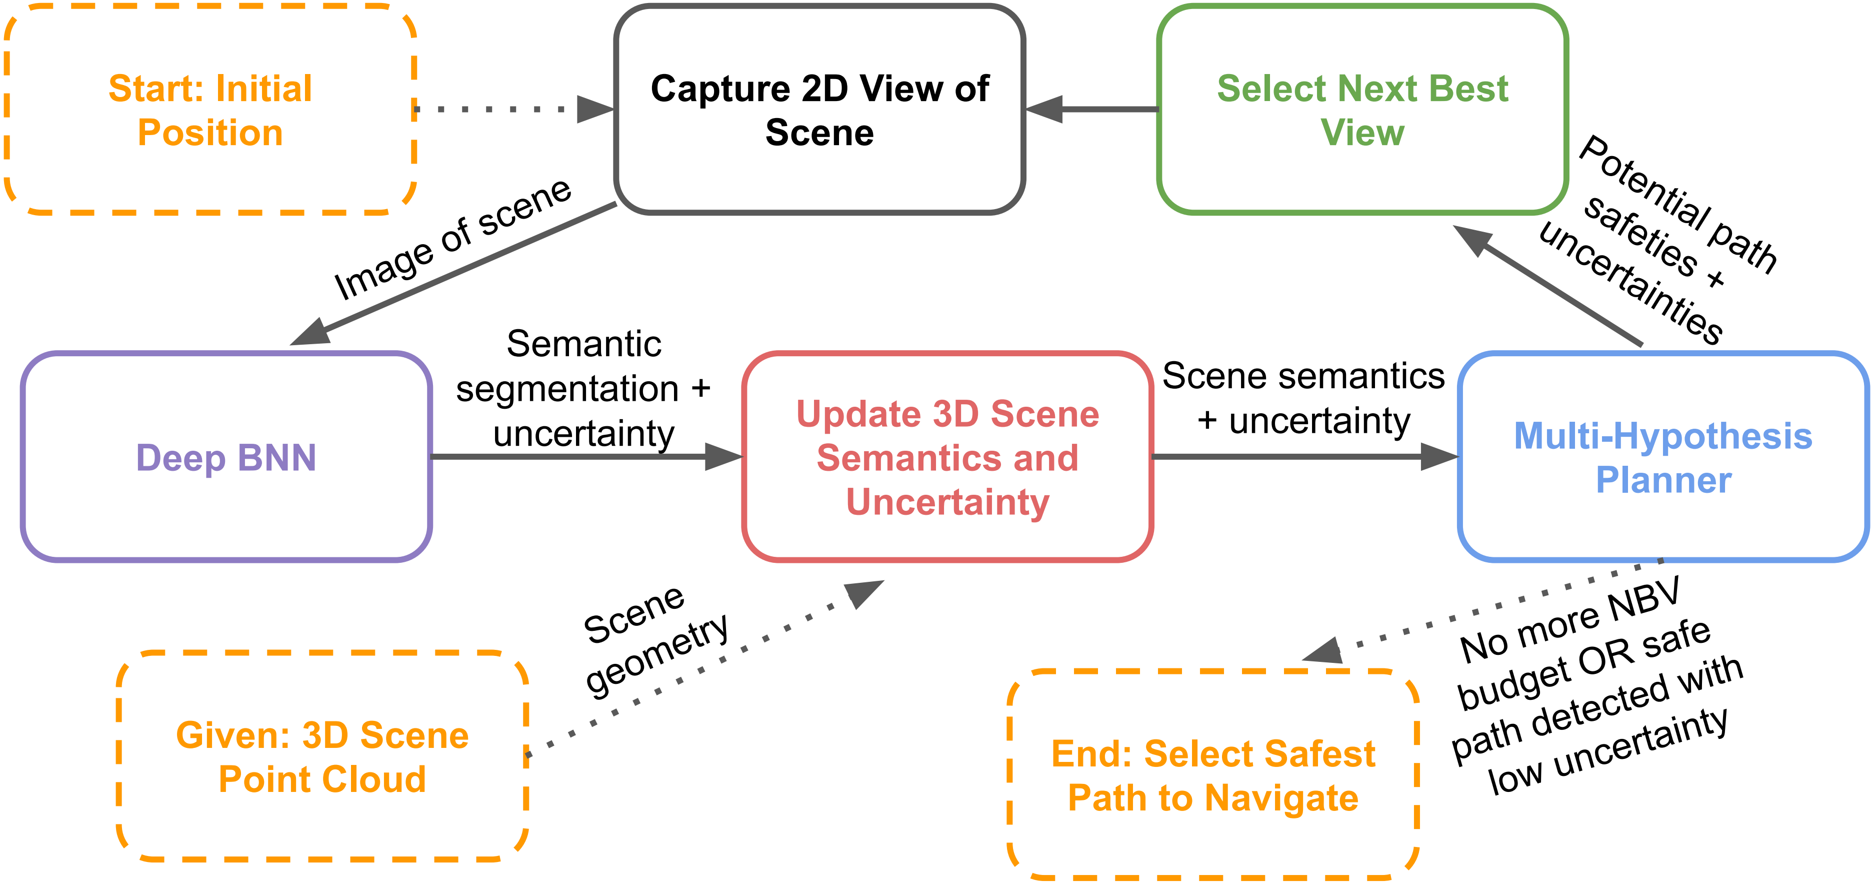

Han*, <u>Lin*</u>, Banfi*, Bala, Campbell, ***"DeepSemanticHPPC: Hypothesis-based Planning over Uncertain Semantic Point Clouds"****, ICRA 2020*

Key Findings: Unstructured Real World Navigation

Reasoning about **semantics with uncertainty** allows higher path safety than (B1) only geometry and (B2) semantics without uncertainty reduction.

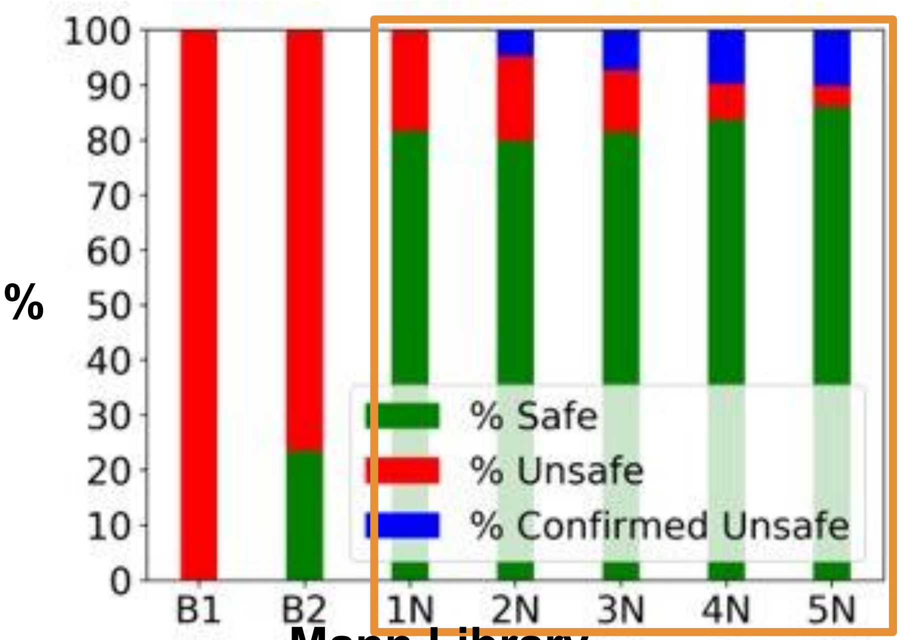

Mann Library

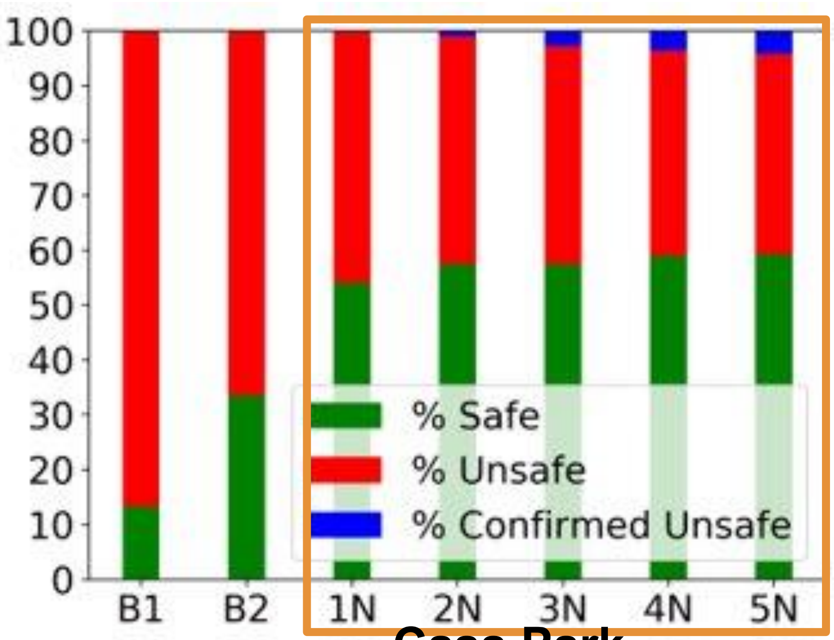

Cass Park

Key Findings: Unstructured Real World Navigation

Accounting for **viewing angle+distance (geometry)** and **uncertainty of viewable path nodes** is important for selecting good measurements.

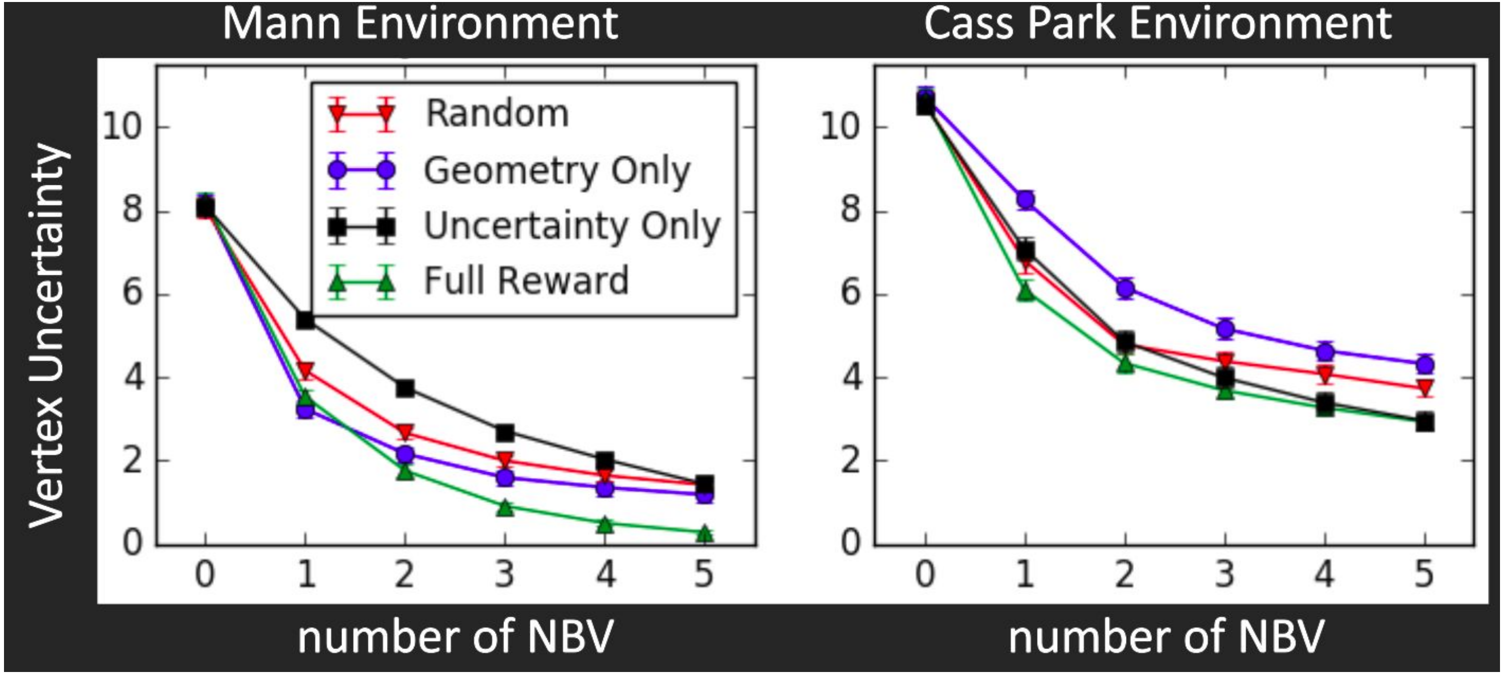

Uncertainty Reduction

Projected / Estimated View
(given point cloud)

True Captured View of Environment

$$p_S^i = \sum_{j \in S} p_j^i,\ p_U^i = 1 - p_S^i = \sum_{j \in U} p_j^i$$

$$\sigma^i = \min\left(\sqrt{\sum_{j \in S} \sigma_j^{i2}}, \sqrt{\sum_{j \in U} \sigma_j^{i2}}\right)$$