JUNE 18-22, 2023 CVPR VANCOUVER, CANADA

Australian National University

Autonomous Vehicle Localization By Leveraging Off-the-shelf Satellite Images

Presenter: Yujiao Shi

The Australian National University

JUNE 18-22, 2023 CVPR VANCOUVER, CANADA

Australian National University

Joint works with

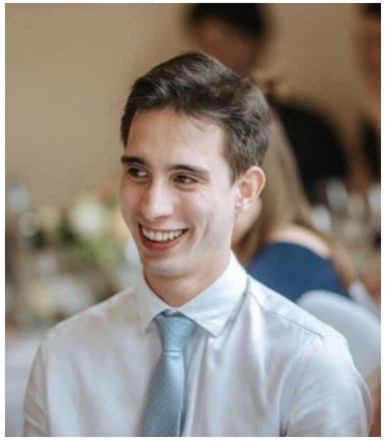

- Prof. Hongdong Li
- Dr. Xin Yu
- Dr. Dylan Campbell
- Dr. Liu Liu

Task Definition

Where am I?

Geo-tagged Satellite Image

Content Outline

(1) Coarse-level City-scale Localization

NeurIPS 2019
AAAI 2020

Localizing Orientation-Aligned Panoramas

CVPR 2020

Localizing Images with Unknown Orientation and limited FoV

(2) Increasing the Localization Accuracy

TPAMI 2022

Exhaustive Search

CVPR 2022

Pose Optimization

ECCV2022

Localization with Uncertainty Estimates

CVPR2023

Geometry-guided SliceMatch

Coarse-level City-scale Localization

- Part-1: Localizing Orientation Aligned Panoramas (AAAI 2020, NeurIPS 2019)
- Part-2: Localizing Images with Unknown Orientation and Limited FoV (CVPR 2020)

Image Retrieval

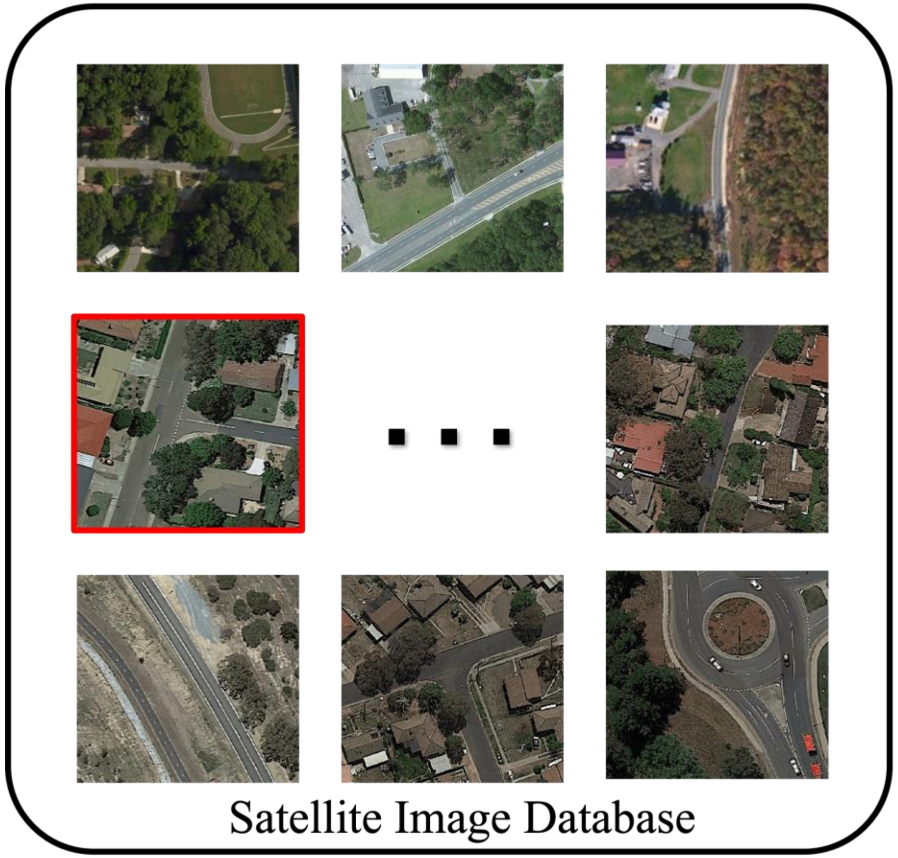

Constructing database:

1. Sample grids
2. Split to small patches

Image Retrieval

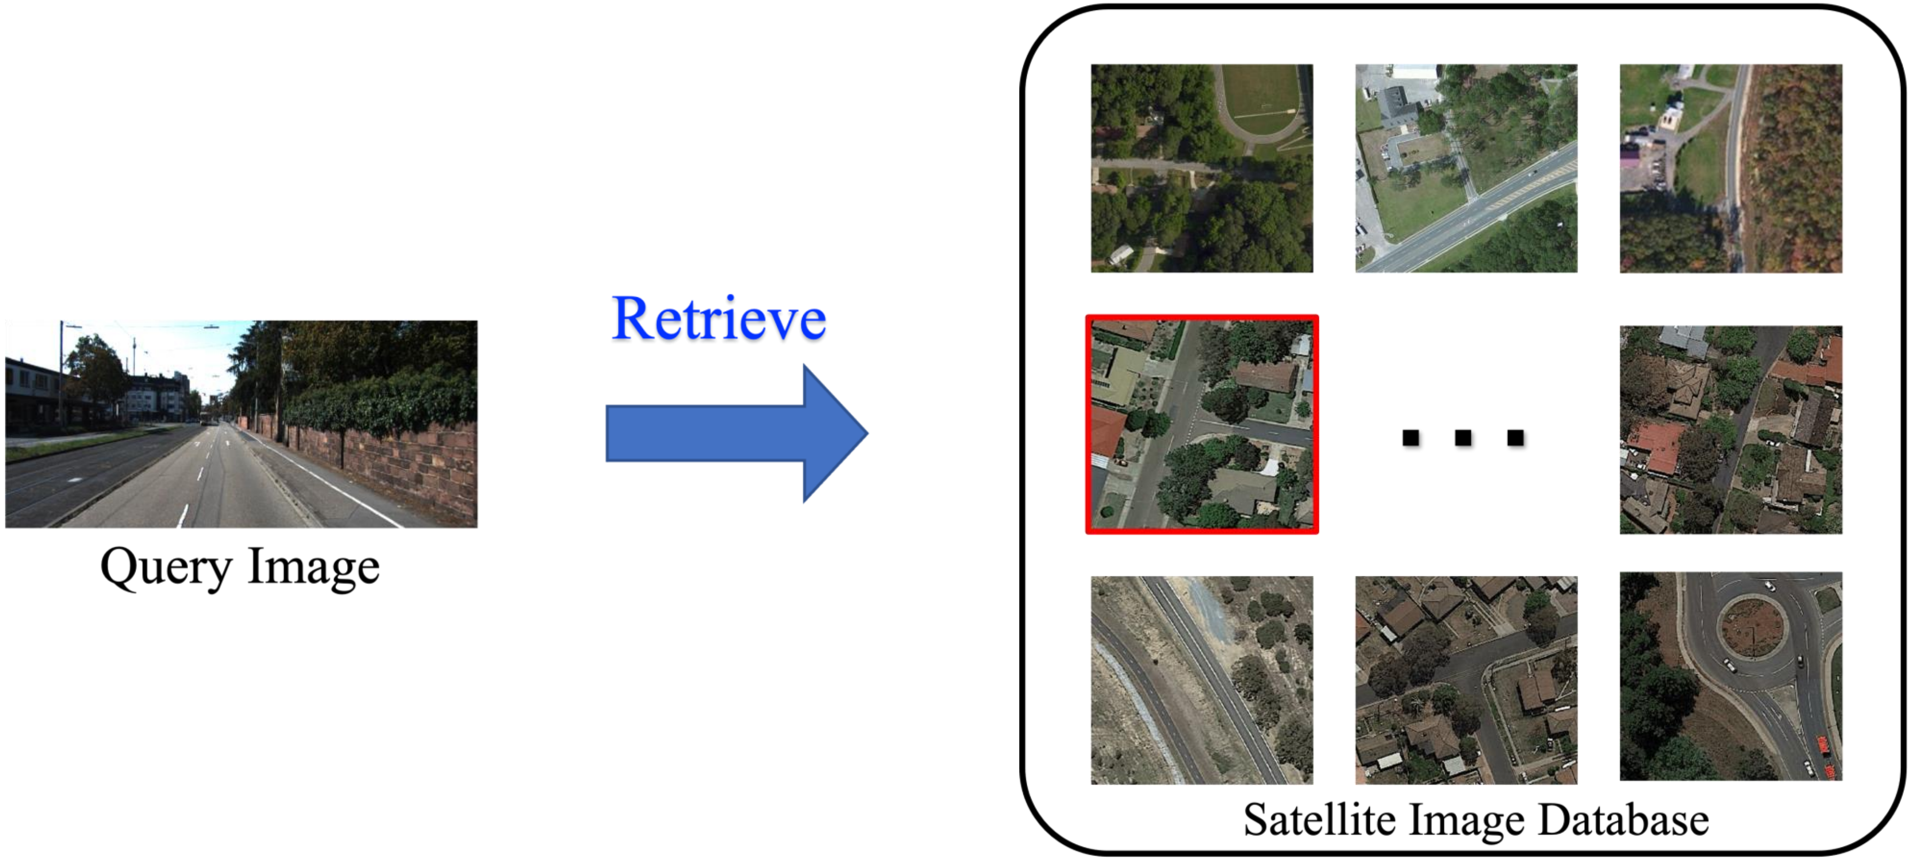

The GPS of retrieved satellite image is assigned as current location.

A Typical Solution

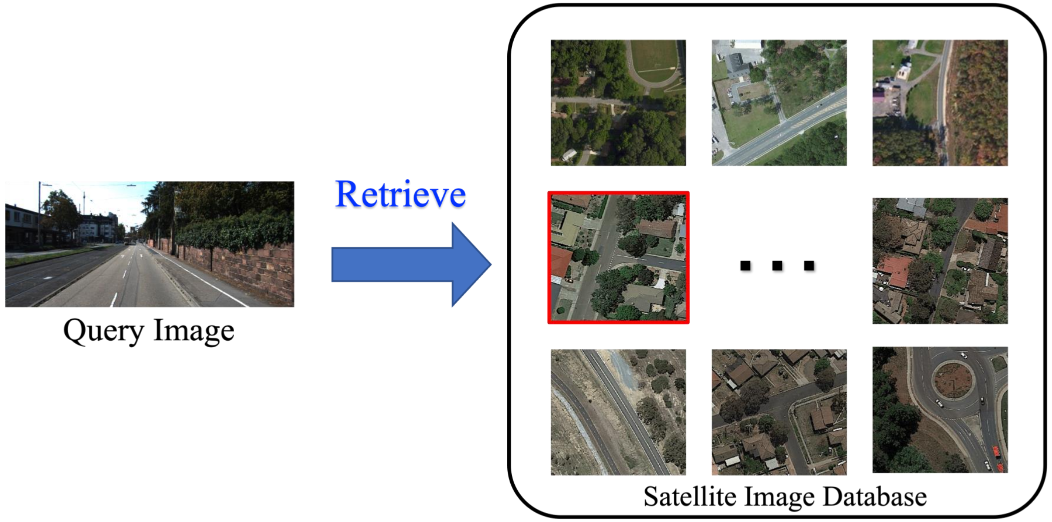

Task Formulation
-- Image Retrieval

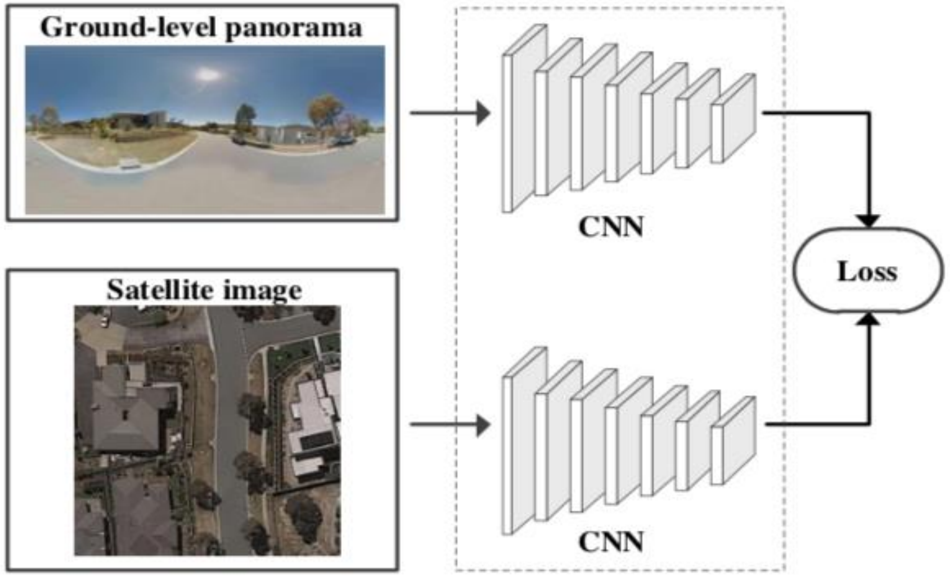

Solution
-- Deep Metric Learning

Challenges

1. Significant domain differences

2. Unknown orientation and limited FoV

3. Limited localization accuracy

Our Solutions

Prof. Hongdong Li

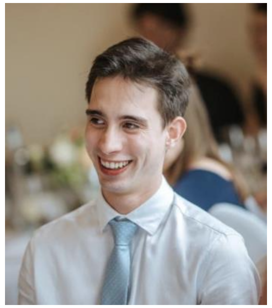

Dr. Dylan Campbell

Spatial-aware feature aggregation for image based cross-view geo-localization Y Shi, L Liu, X Yu, H Li NeurIPS 32	110	2019
Where am i looking at? joint location and orientation estimation by cross-view matching Y Shi, X Yu, D Campbell, H Li CVPR, 4064-4072	87	2020
Optimal feature transport for cross-view image geo-localization Y Shi, X Yu, L Liu, T Zhang, H Li AAAI oral 34 (07), 11990-11997	80	2020
Beyond Cross-view Image Retrieval: Highly Accurate Vehicle Localization Using Satellite Image Y Shi, H Li CVPR, 17010-17020	14	2022
Geometry-guided street-view panorama synthesis from satellite imagery Y Shi, DJ Campbell, X Yu, H Li TPAMI	11	2022
Accurate 3-DoF Camera Geo-Localization via Ground-to-Satellite Image Matching Y Shi, X Yu, L Liu, D Campbell, P Koniusz, H Li TPAMI	8	2022
CVLNet: Cross-View Semantic Correspondence Learning for Video-based Camera Localization Y Shi, X Yu, S Wang, H Li ACCV 2022 oral	2	2022

Dr. Xin Yu

Dr. Liu Liu

Github

Part-1: Localizing Orientation Aligned Panoramas

AAAI 2020 & NeurIPS 2019

Geometry Correspondences

JUNE 18-22, 2023 CVPR VANCOUVER, CANADA

Australian National University

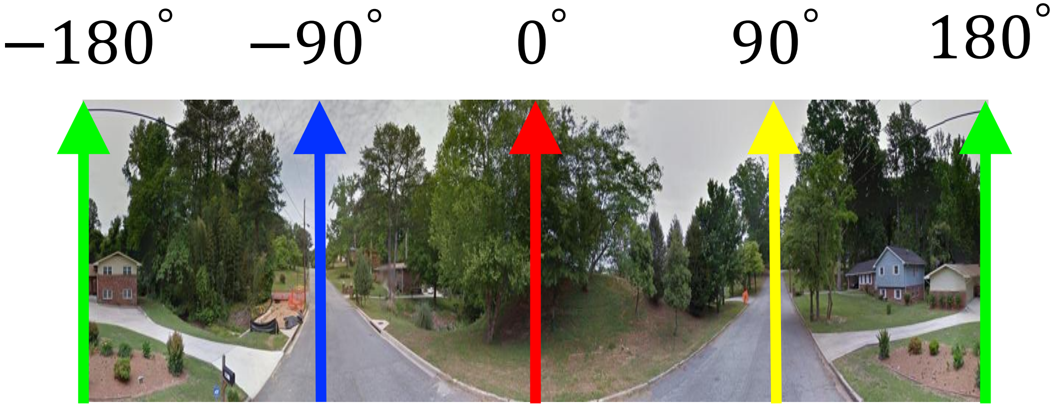

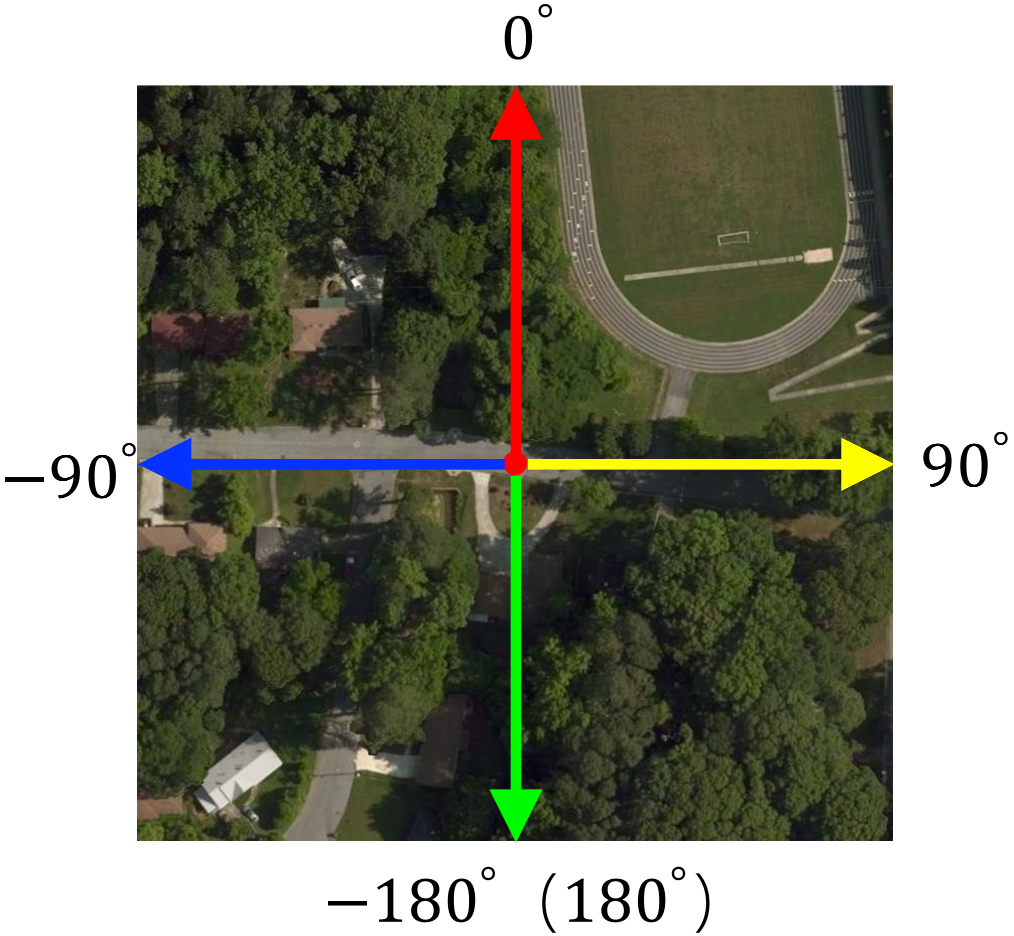

Optimal Feature Transport
(AAAI 2020)

1. Spatial layout information

2. Domain differences

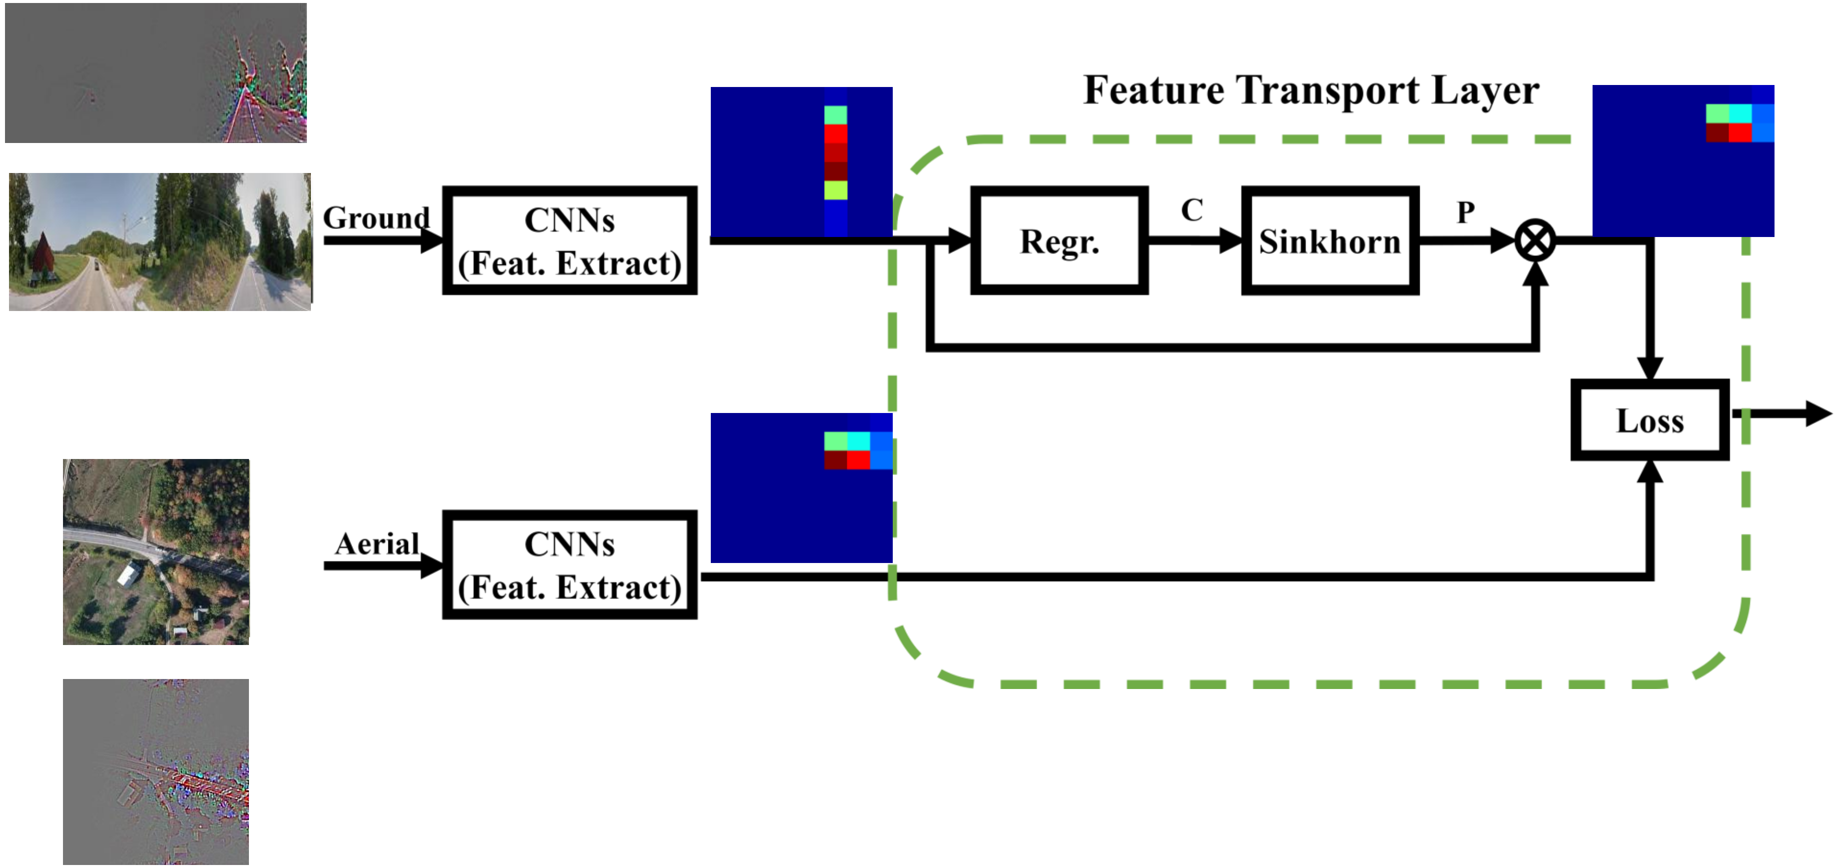

Shi, Yujiao, et al. "Optimal feature transport for cross-view image geo-localization." AAAI 2020.

Theory

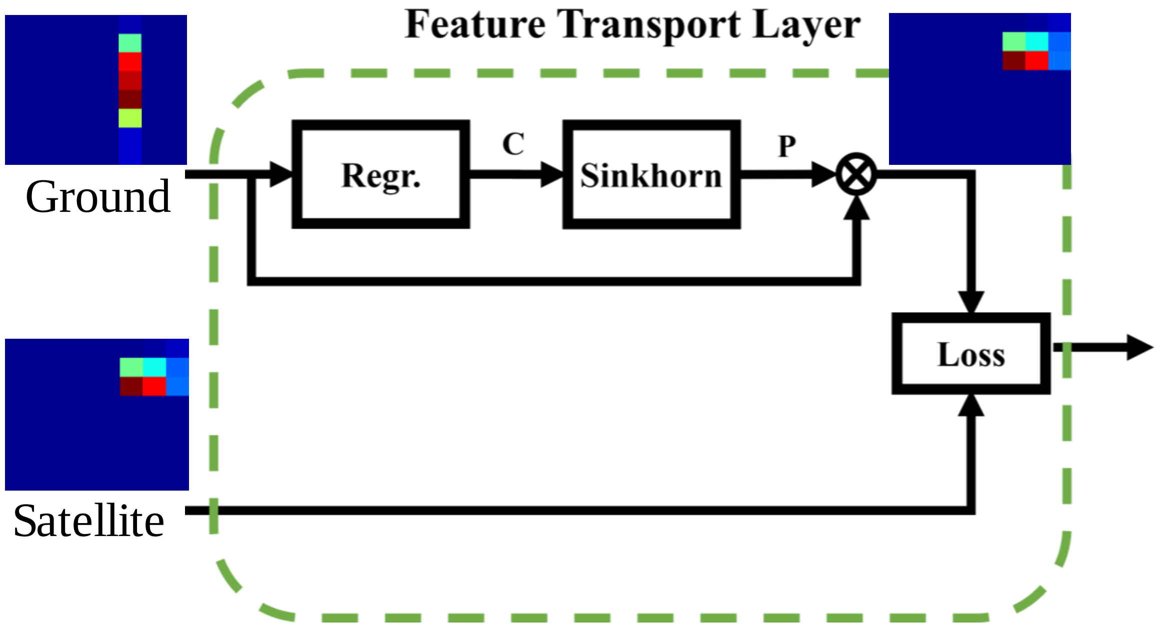

$$\mathbf{P}^* = \underset{\mathbf{P}\in\mathcal{P}}{\arg\min} \langle \mathbf{P}, \mathbf{C} \rangle_F - \lambda h(\mathbf{P})$$

Row normalization:

$$\mathcal{N}^r_{i,j}(\mathbf{C}') = \frac{c'_{i,j}}{\sum_{k=1}^{N} c'_{i,k}}$$

Column normalization:

$$\mathcal{N}^c_{i,j}(\mathbf{C}') = \frac{c'_{i,j}}{\sum_{k=1}^{N} c'_{k,j}}$$

Polar Transform
(NeurIPS 2019)

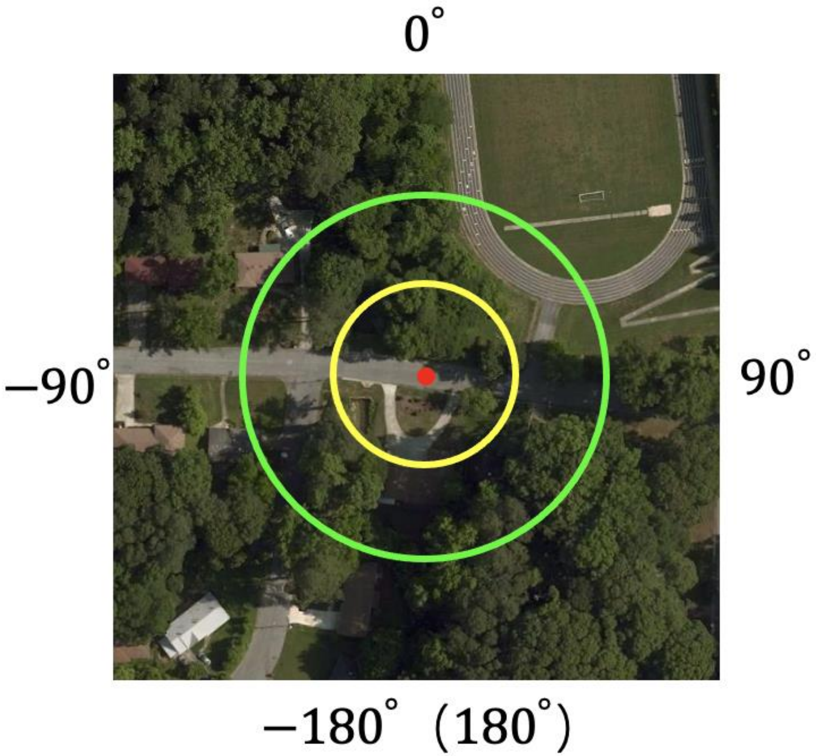

Satellite Image

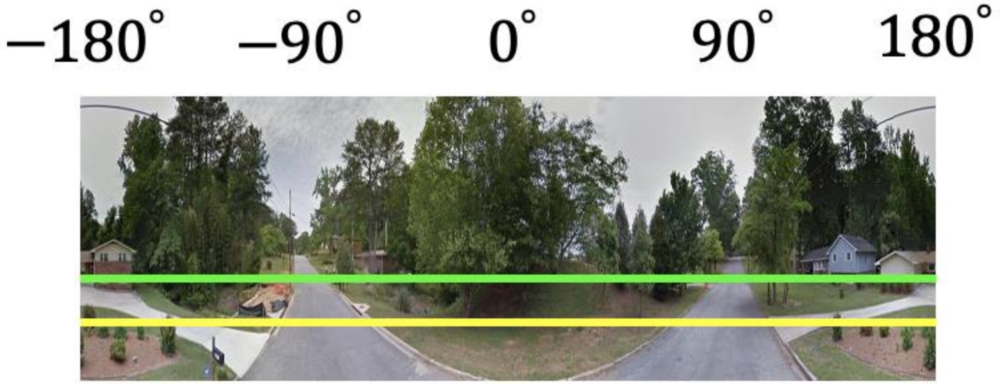

Ground Panorama


Polar-transformed Satellite Image

Framework

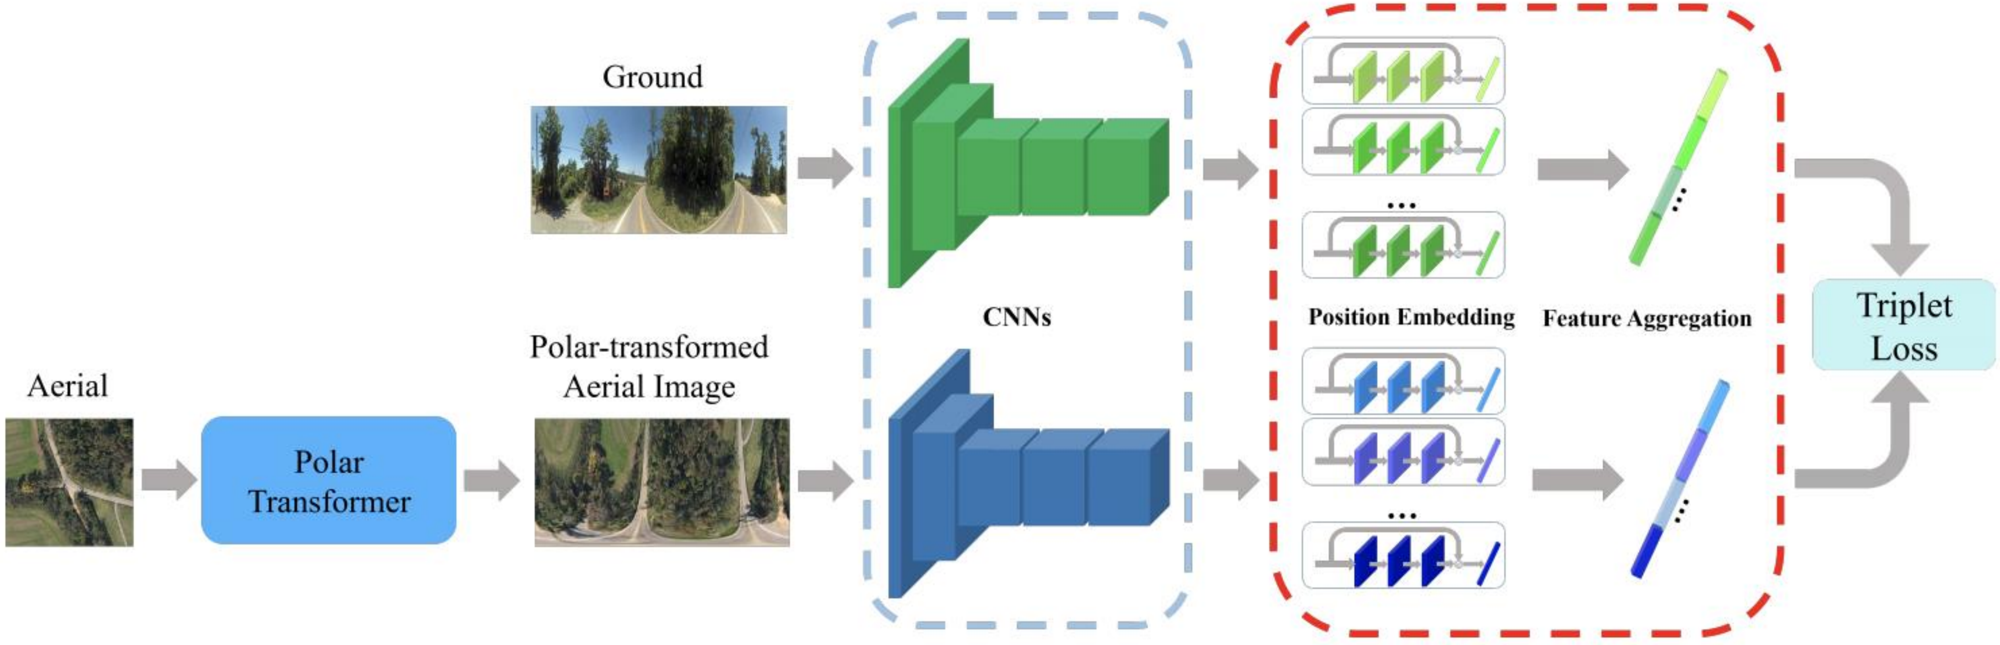

Spatial-aware feature aggregation (SAFA) module

Shi, Yujiao, et al. "Spatial-aware feature aggregation for image based cross-view geo-localization." NeurIPS 2019.

Spatial Attention

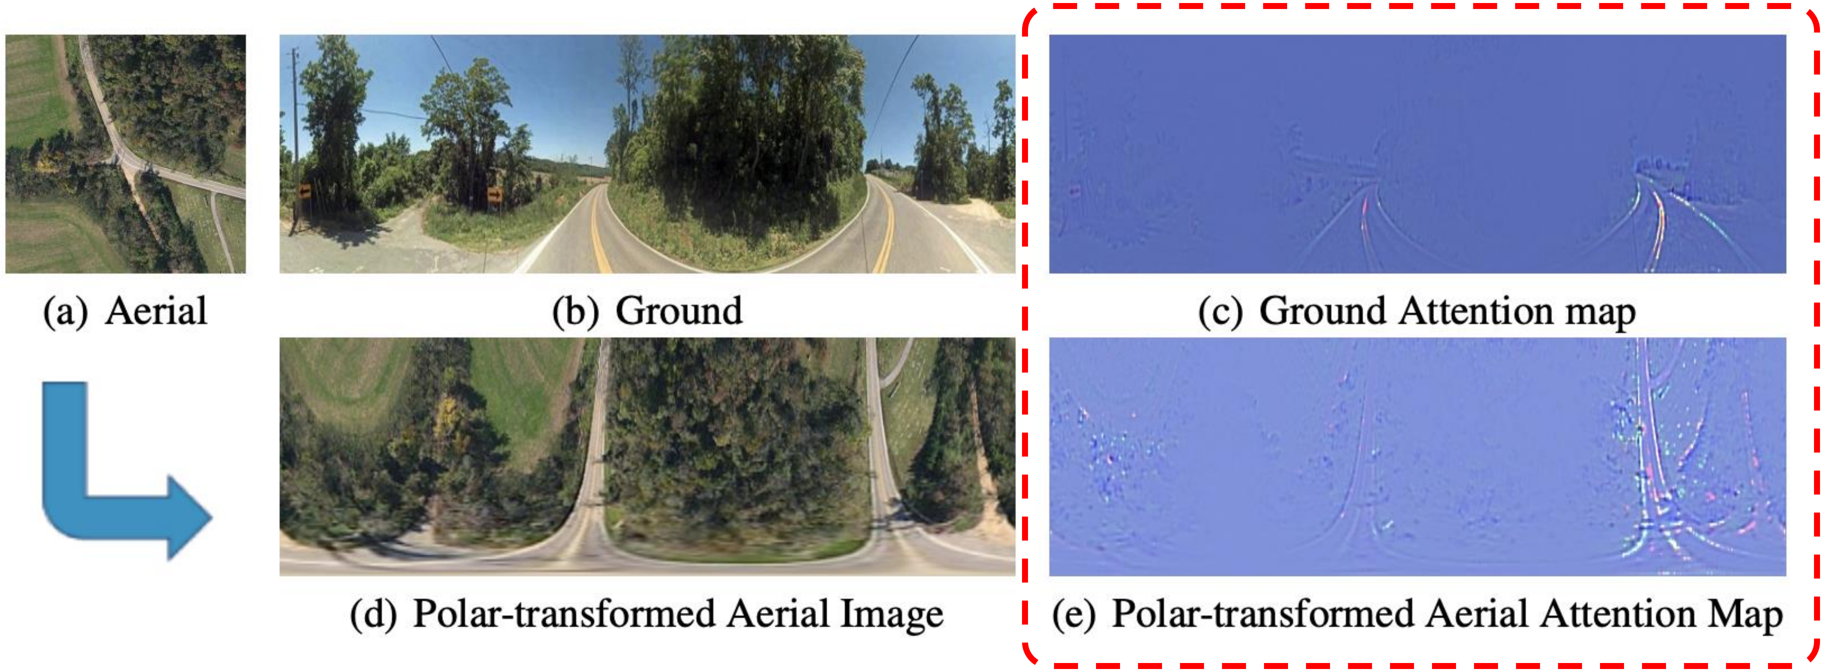

Localization Visualization

Orientation-aligned panoramas

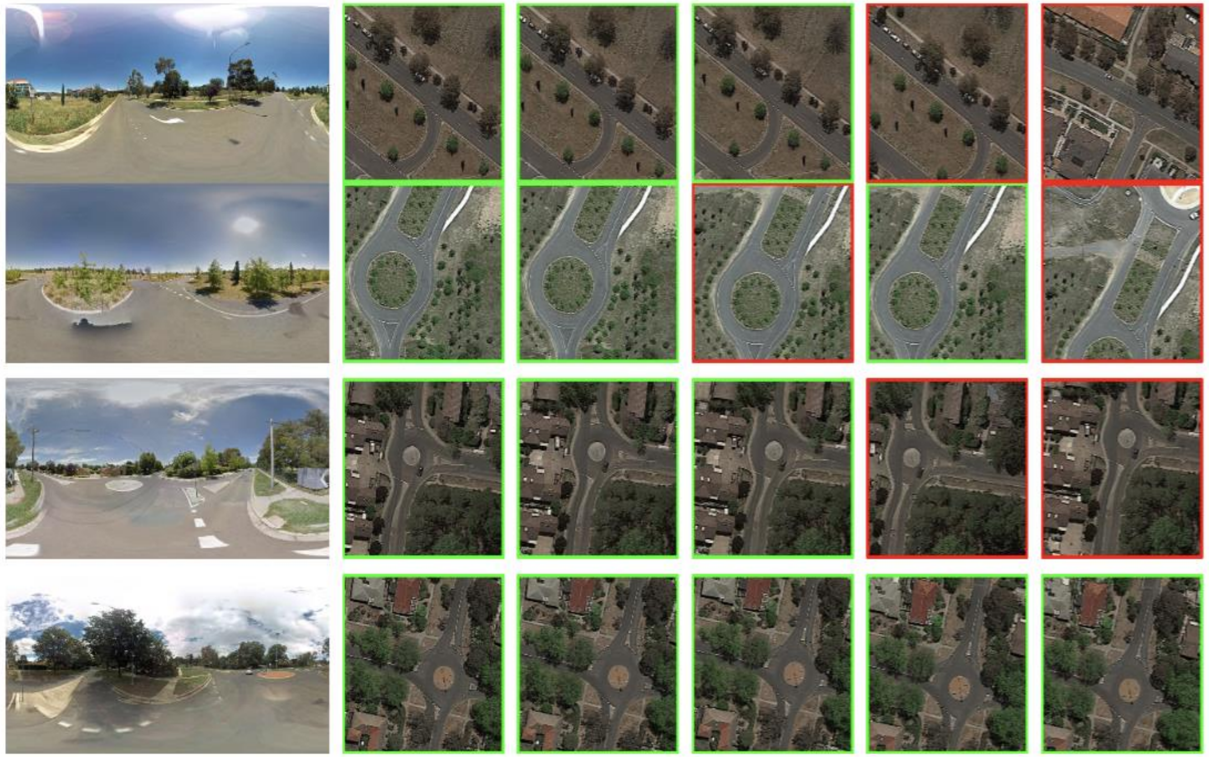

Query Image Top-1 Top-2 Top-3 Top-4 Top-5

Part-2: Localizing Images with Unknown Orientation and a limited FoV

CVPR 2020

Challenges

Satellite Image

Ground images captured at the same location but different orientations

Ground images with a limited FoV captured at the the ***same location but different orientation*** can be totally **different**.

Orientation Estimation

Dynamic Similarity Matching:

Estimate orientation alignment between two-view images.

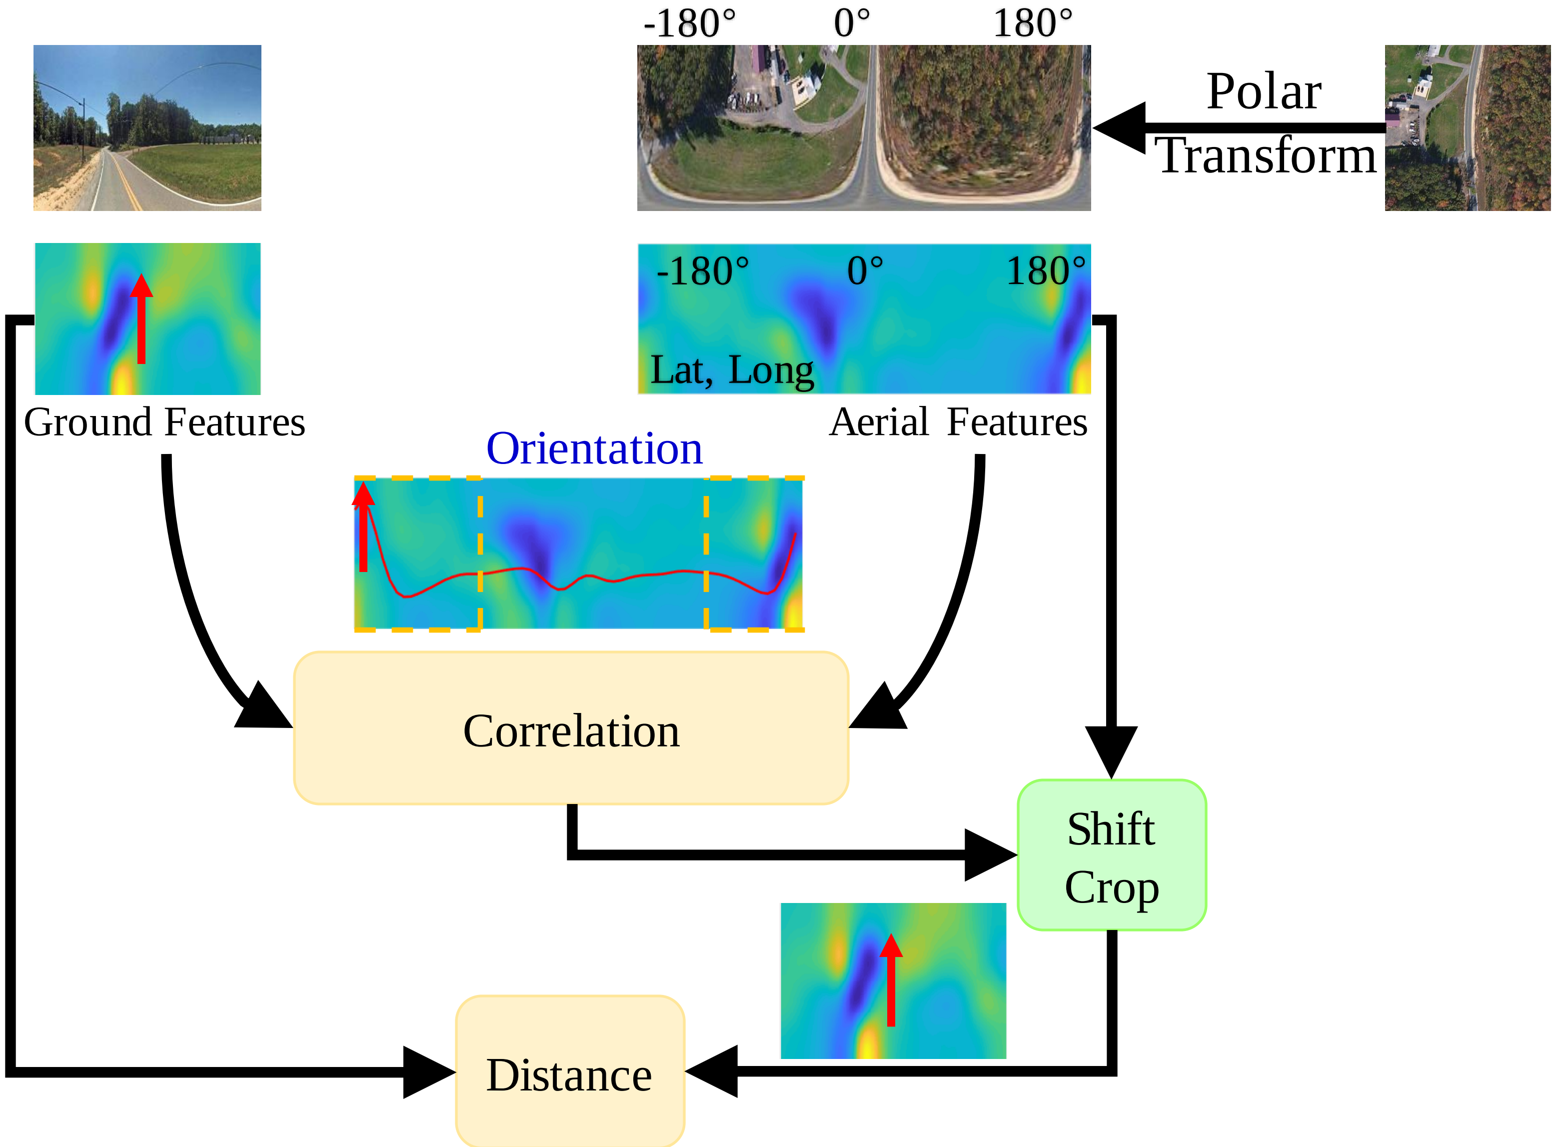

Framework

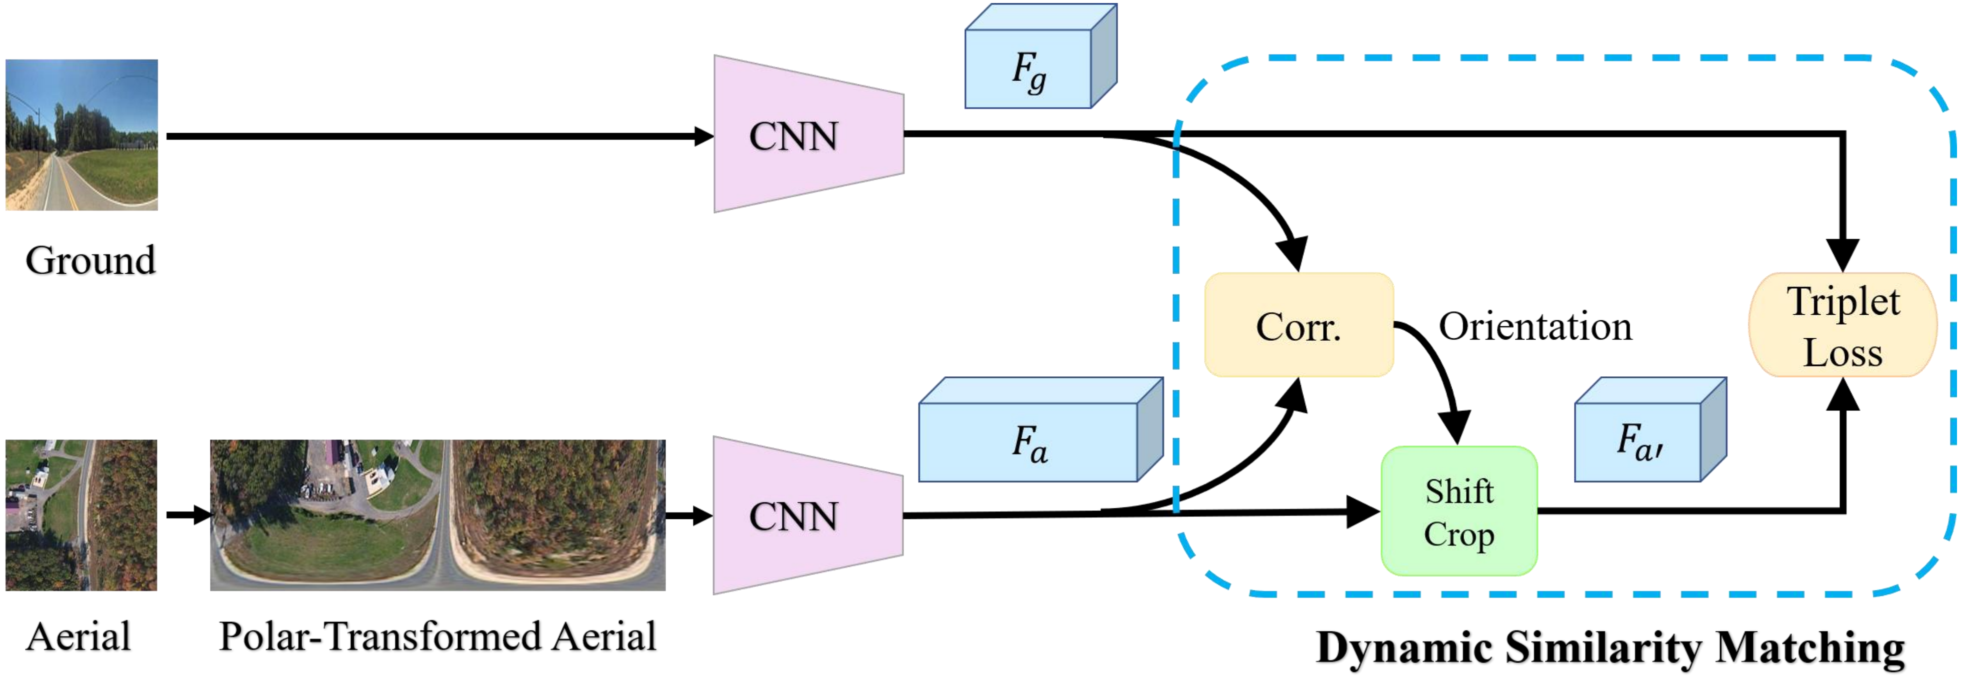

- Joint orientation and location estimation
- End-to-end

Shi, Yujiao, et al. "Where am i looking at? joint location and orientation estimation by cross-view matching." CVPR 2020.

Orientation Estimation

(a) FoV=360°

(b) FoV=180°

Table 4. Orientation prediction performance on correctly localized ground images.

Dataset	CVUSA				CVACT_val			
FoV	360°	180°	90°	70°	360°	180°	90°	70°
orien_acc	99.41	98.54	76.15	61.67	99.84	99.10	74.51	55.18
median_error	2.38	2.38	4.50	4.88	1.97	2.89	5.21	6.22

Limitation

Fig. 11: Examples of symmetric scenes (satellite images). At these locations, it is hard to determine the orientation (azimuth angle) of a ground image.

Localization Visualization

Unknown orientation and limited FoV

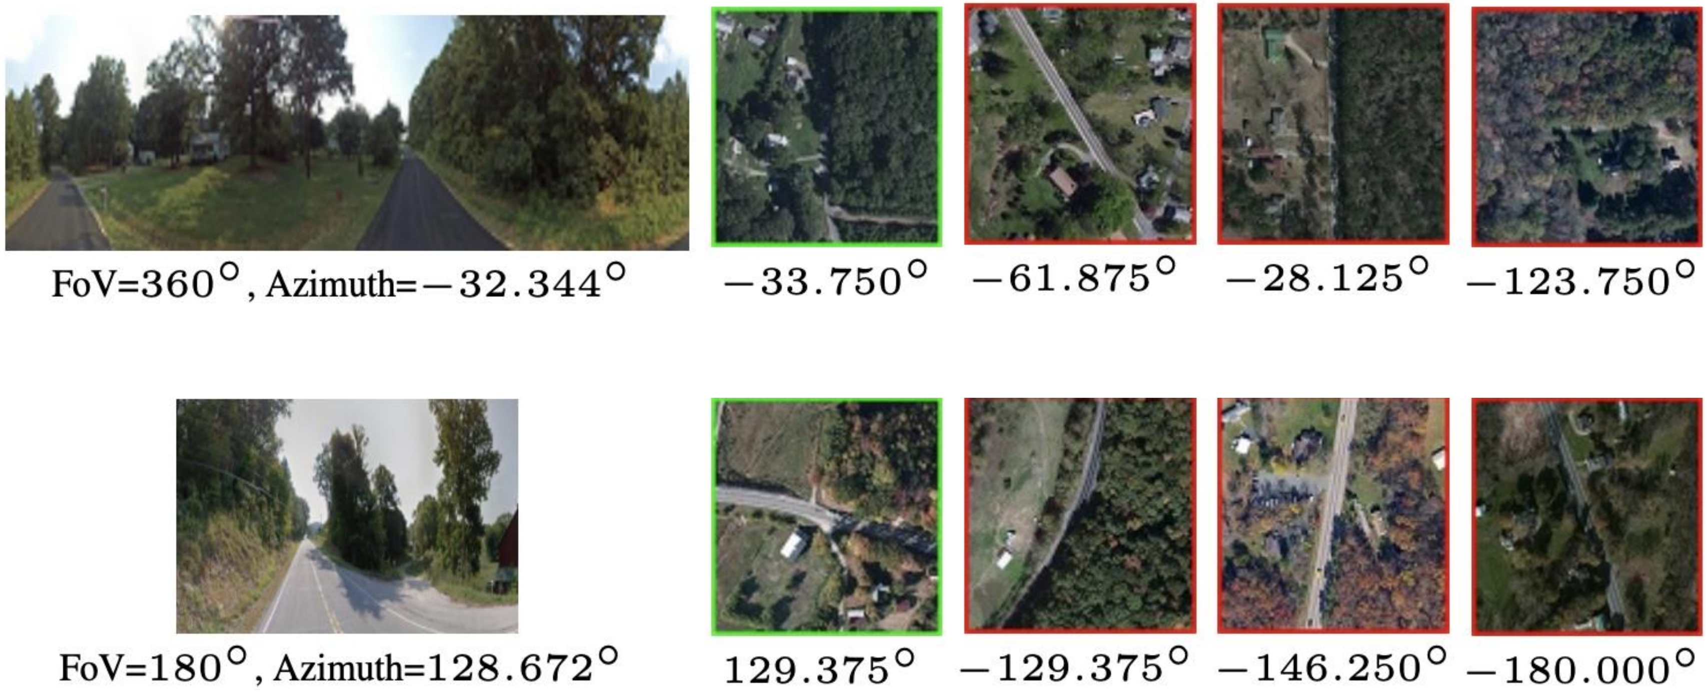

Query Image Top-1 Top-2 Top-3 Top-4

Drawbacks of Retrieval

- Poor Localization accuracy
  -- limited to the density of sampled grids.

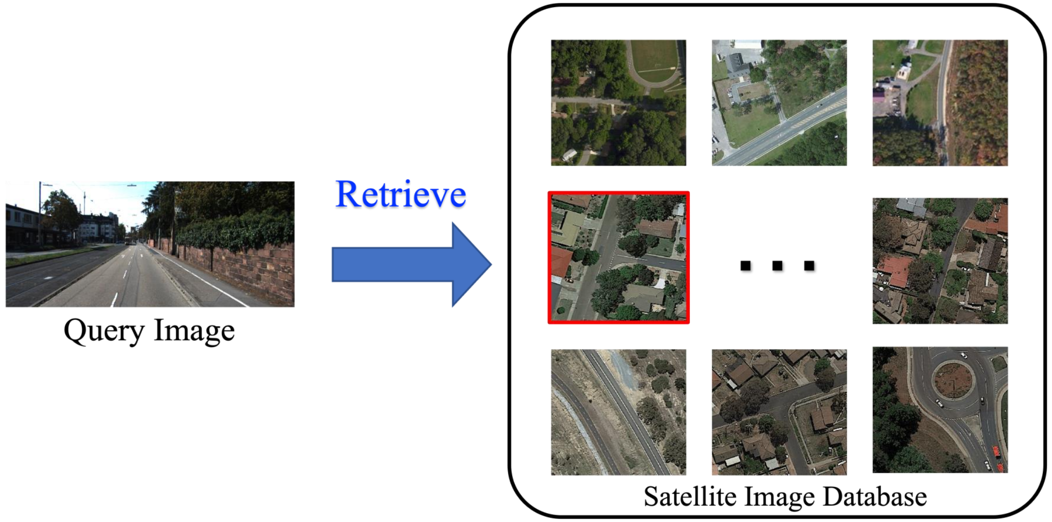

Increasing the Localization Accuracy

- Part-1: Exhaustive Search (TPAMI 2022)
- Part-2: Pose Optimization (CVPR 2021)
- Part-3: Localization with Uncertainty Estimates (ECCV 2022)
- Part-4: Geometry-guided SliceMatch (CVPR 2023)

Part-1

Exhaustive Search

TPAMI 2022

Exhaustive Search

Query Image

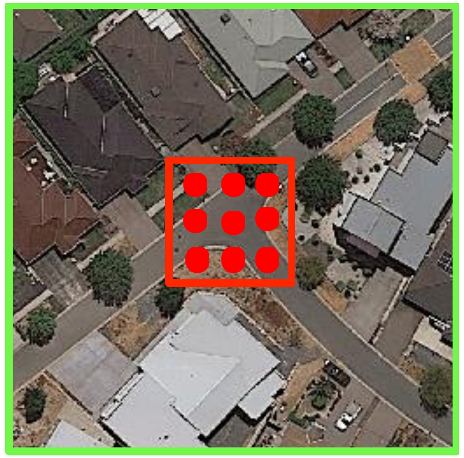

Top-1 Retrieved Satellite Image

Search every location per unit distance in the search region.

Localization Process

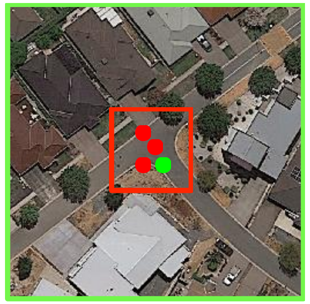

SSIM as similarity measure

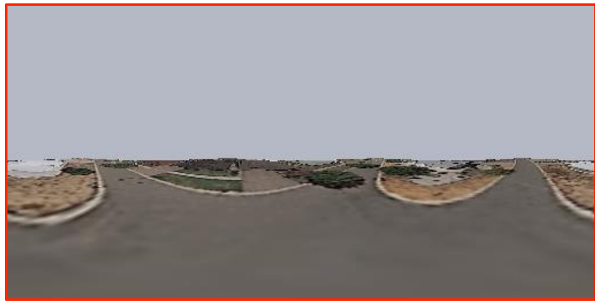

Projected at (0, 0)

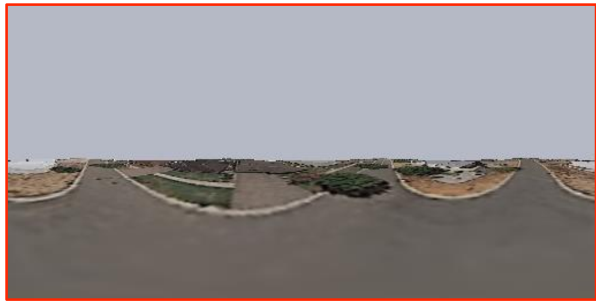

Projected at (-3, -8)

Projected at (4, -5)

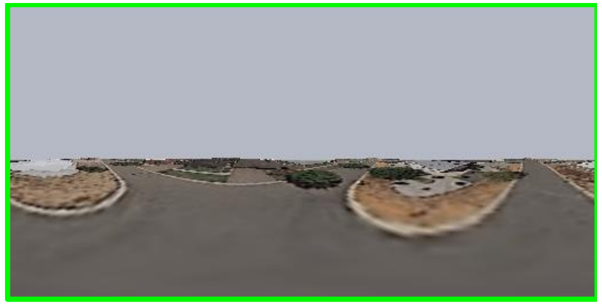

Projected at (9, 2)

Shi, Yujiao, et al. "Accurate 3-DoF Camera Geo-Localization via Ground-to-Satellite Image Matching." TPAMI 2022.

Unknown orientation and limited FoV

Unknown orientation and limited FoV

Change feature maps to original images.

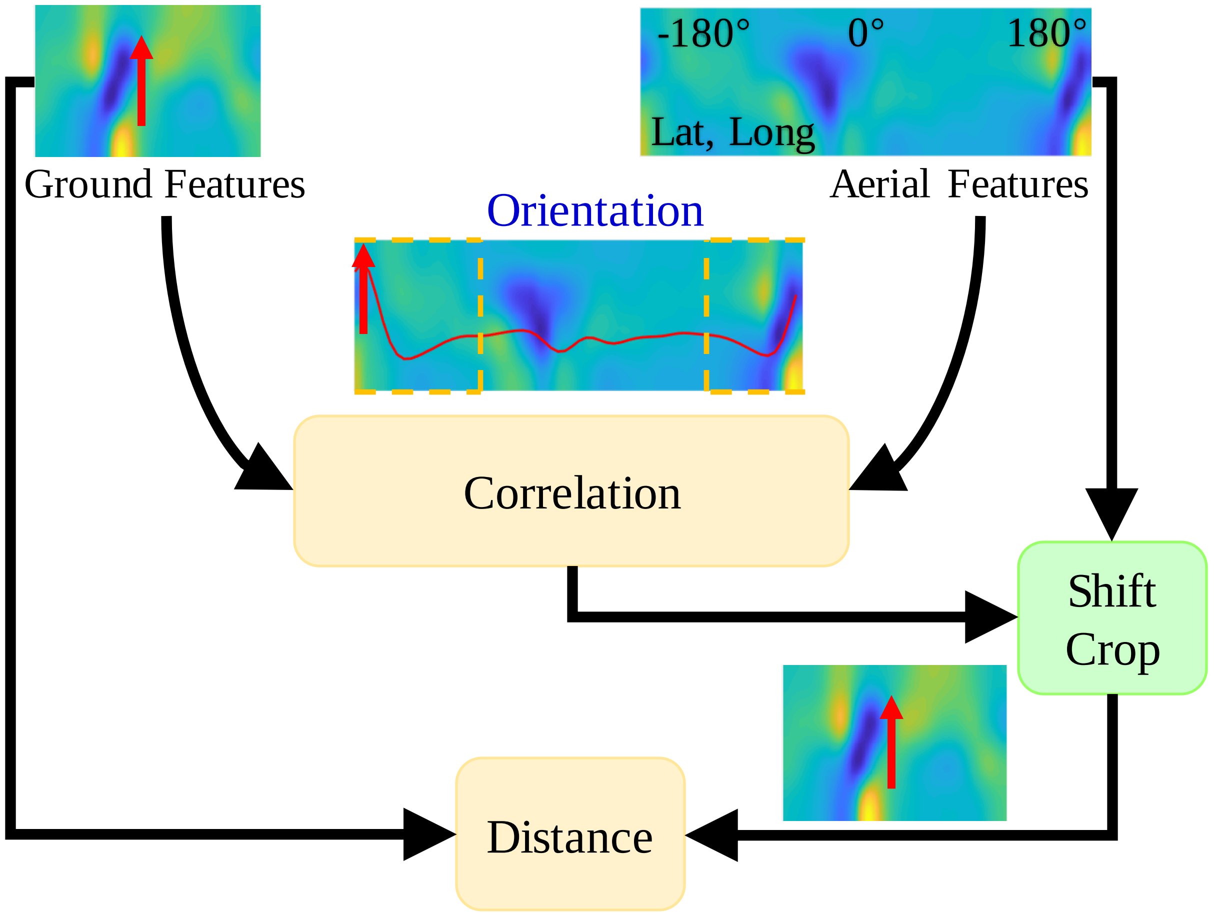

Fine-grained Localization

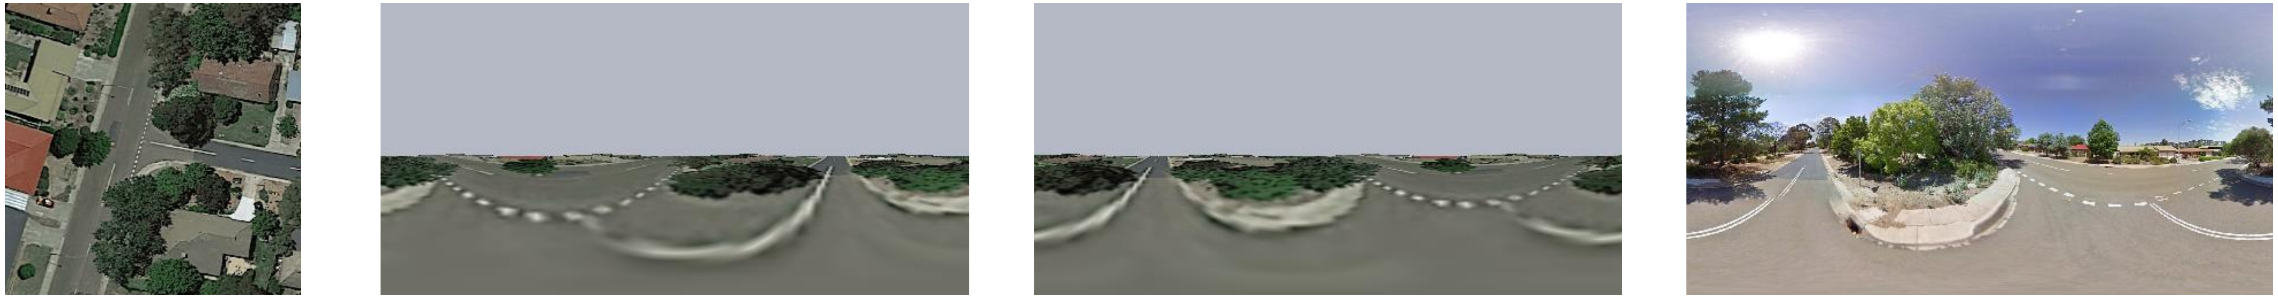

Top-1 retrieved image | Projected at satellite image center | Projected at (3, 3), pred_orien=-159° | gt_orien=-159°

Top-1 retrieved image | Projected at satellite image center | Projected at (19, -20), pred_orien=-129° | gt_orien=-130°

Localizing orientation-unknown panoramas

Fine-grained Localization

Top-1 retrieved image

Projected at satellite image center

Projected at (6, -3), pred_orien=154°

gt_orien=153°

Top-1 retrieved image

Projected at satellite image center

Projected at (15, -9), pred_orien=67°

gt_orien=68°

Localizing images with unknown-orientation and limited FoV

Projective Transform

- Project satellite image to a ground-view according to a relative location
- Assume ground plane homography

$$\begin{cases} u_i^s = u_0 + s z_2 \tan(\pi v_i^t / H_g) \cos(2\pi u_i^t / W_g) \\ v_i^s = v_0 - s z_2 \tan(\pi v_i^t / H_g) \sin(2\pi u_i^t / W_g), \end{cases}$$

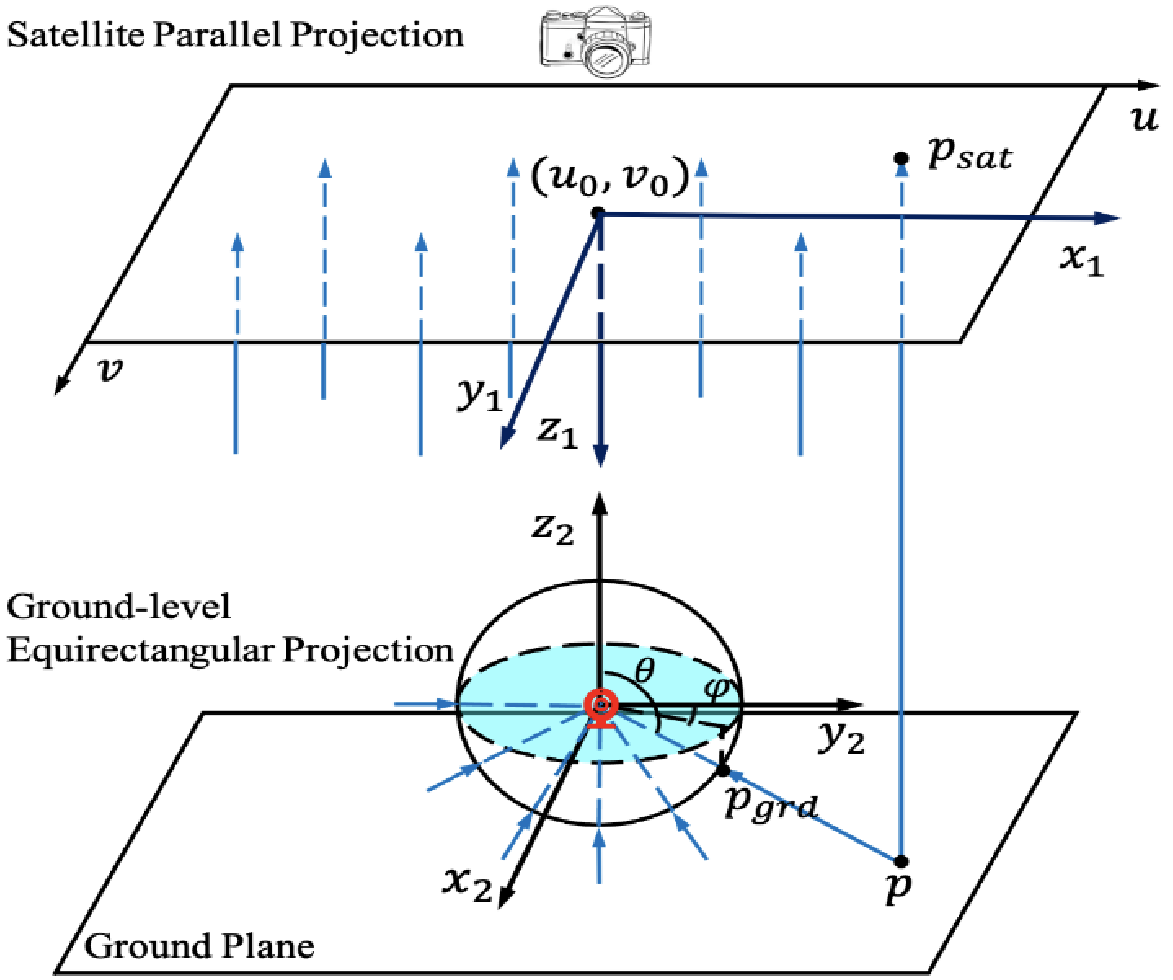

Quantitative Evaluation

User study

Matching Satellite Image

Projected Satellite Image at Ground GPS location

Projected Satellite Image at our estimated location

Query Image

Which one is correct? GPS location, or our estimated location?

(1) GPS location
(2) Our estimated location
(3) Both
(4) Neither

Quantitative Evaluation

JUNE 18-22, 2023 CVPR VANCOUVER, CANADA

Australian National University

User study

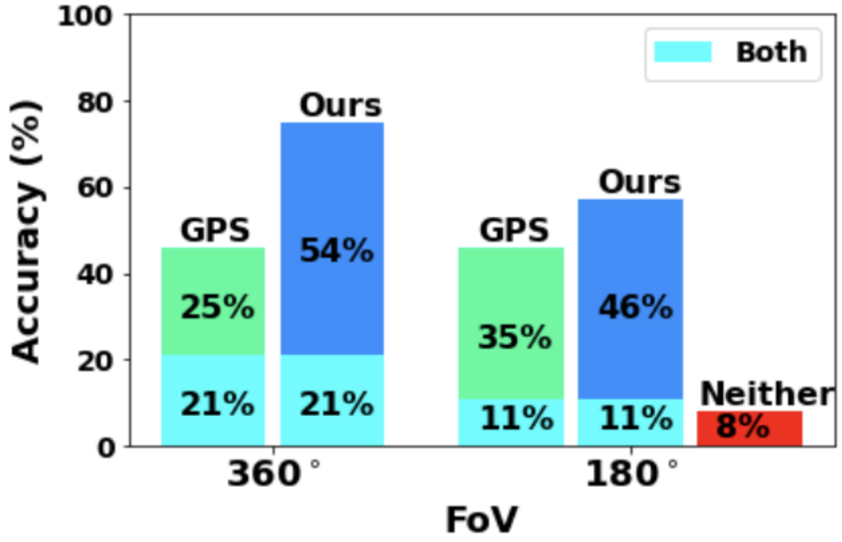

Fig. 14: User study results for fine-grained camera localization (orientation aligned). In this evaluation, users are asked to determine whether a location is correct or not. The color cyan indicates the portion of data where both GPS and our estimated locations are correct. The red bar "Neither" indicates the portion of data that cannot be localized when the FoV decreases from 360° to 180°.

Method Limitation

Matching Satellite Image

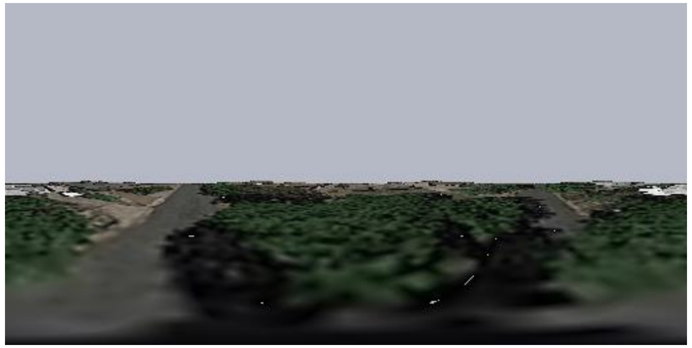

Projected Satellite Image at Ground GPS location

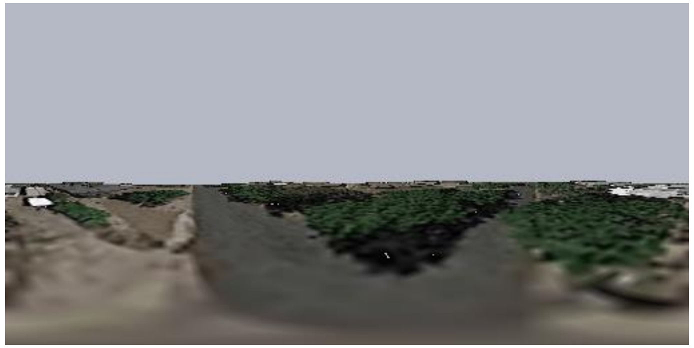

Projected Satellite Image at our estimated location

Query Image

Further Exploration

- **SSIM is not very effective**

  -- Better similarity metric? Feature Similarity

- **Exhaustive search is computationally expensive**

  -- Sophisticated search strategy LM pose optimization

Part-2

Pose Optimization

CVPR 2022

Pose Optimization

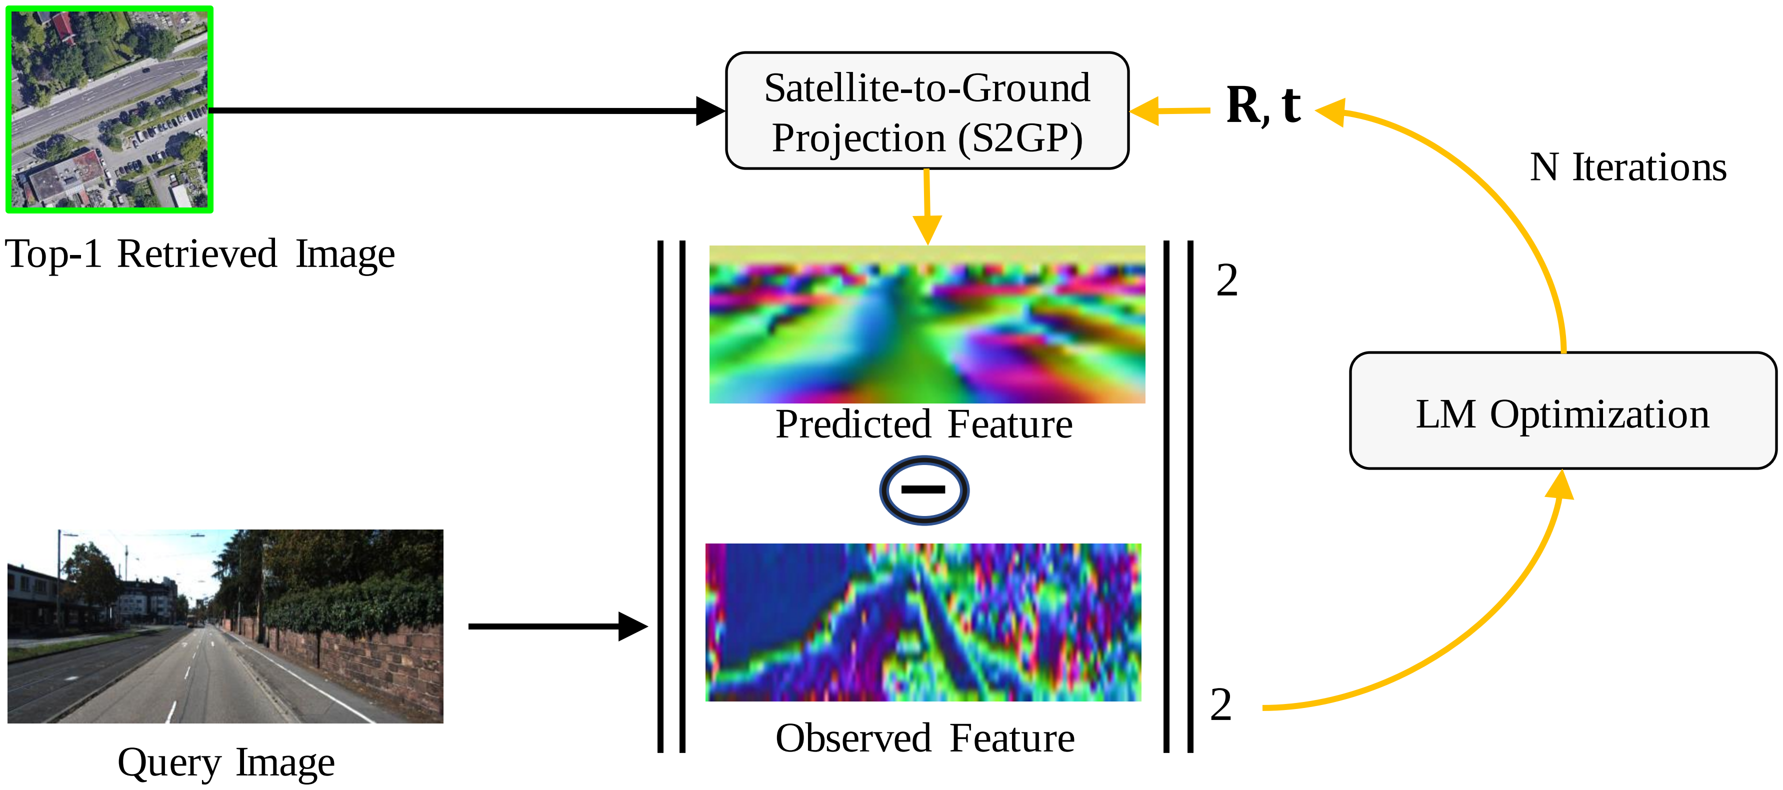

Shi, Yujiao, and Hongdong Li. "Beyond cross-view image retrieval: Highly accurate vehicle localization using satellite image." CVPR 2022.

LM Optimization

LM Objective:

$$\hat{\xi} = \arg\min_{\xi} \|\mathbf{e}^l\|_2^2,$$

$$\mathbf{e}^l = \mathbf{F}^l_{s2g} - \mathbf{F}^l_g.$$

$\mathbf{F}^l_{s2g}$ – projected ground features from satellite domain at the feature level l;
$\mathbf{F}^l_g$ – observed ground features at the feature level l;

Pose is updated by:

$$\xi_{t+1} = \xi_t + \widetilde{\mathbf{H}}^{-1}\mathbf{J}^T\mathbf{e},$$

$$\mathbf{J} = \frac{\partial \mathbf{F}_{s2g}}{\partial \xi} = \frac{\partial \mathbf{F}_{s2g}}{\partial \mathbf{p}_s}\frac{\partial \mathbf{p}_s}{\partial \xi}, \quad \text{and} \quad \mathbf{H} = \mathbf{J}^T\mathbf{J},$$

where t index iterations.

Framework

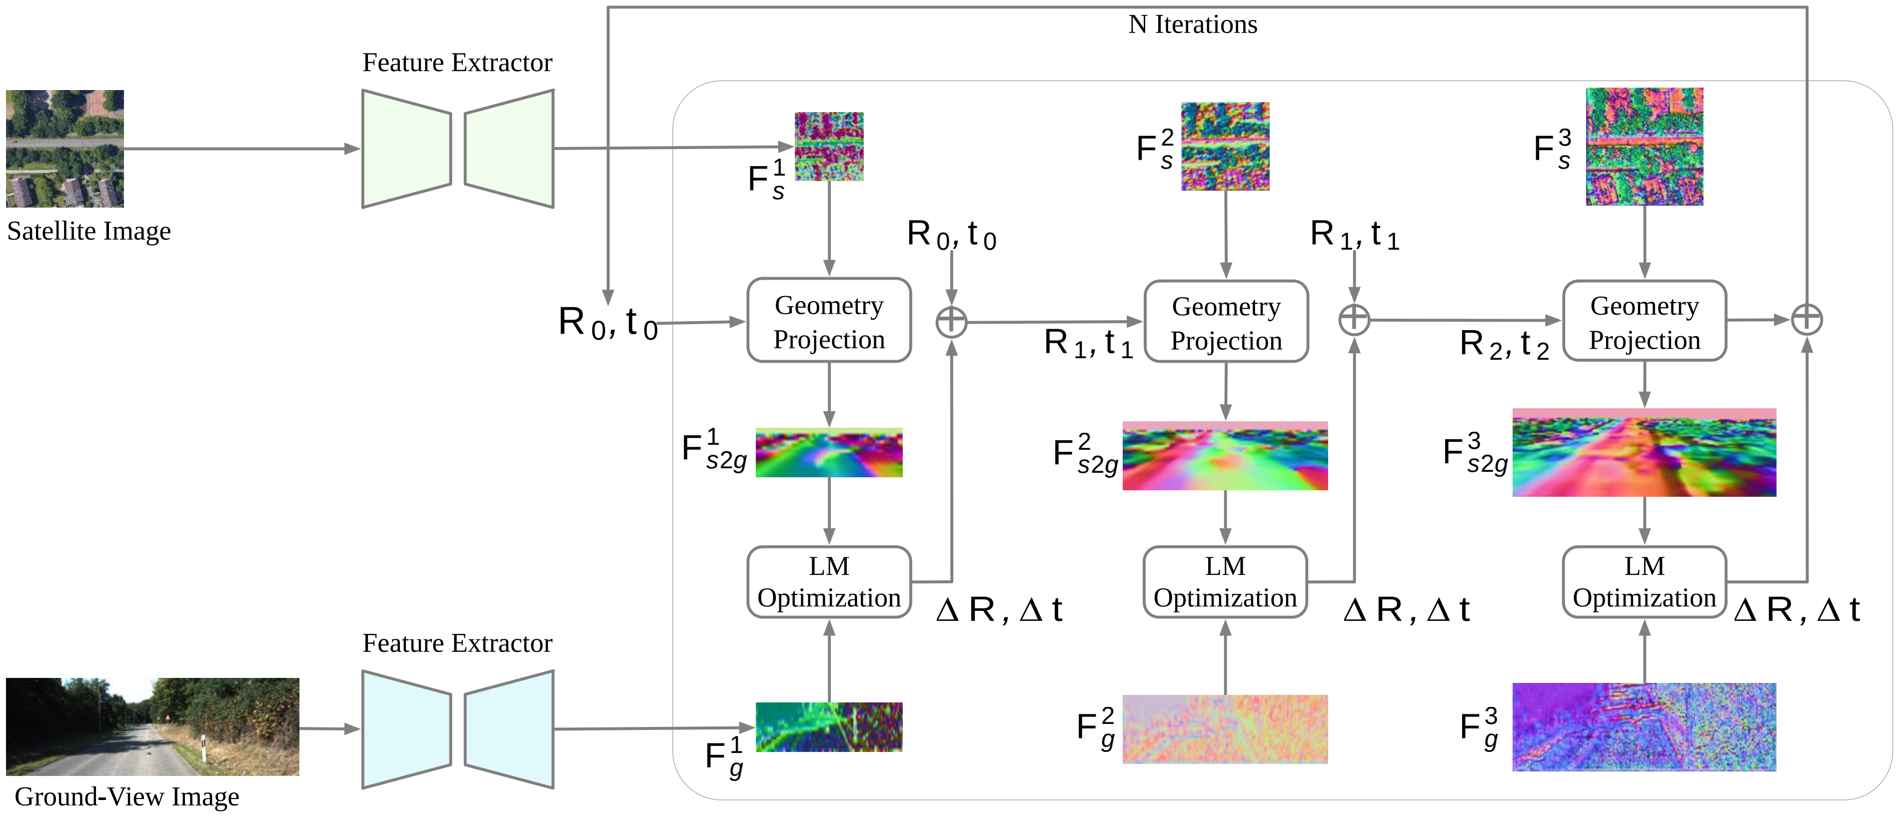

Multi-scale Coarse-to-fine Optimization

Satellite-to-ground Projection

Satellite parallel projection:

$$[u_s, v_s]^T = [\frac{z}{\alpha} + u_s^0, \frac{x}{\alpha} + v_s^0]^T, \qquad (1)$$

α – per-pixel real-word distance.

Ground (pin-hole) camera projection:

$$w[u^g, v^g, 1]^T = \mathbf{K}[x_c, y_c, z_c]^T, \qquad (2)$$

From ground camera to world camera:

$$[x, y, z]^T = \mathbf{R}([x_c, y_c, z_c]^T + \mathbf{t}), \qquad (3)$$

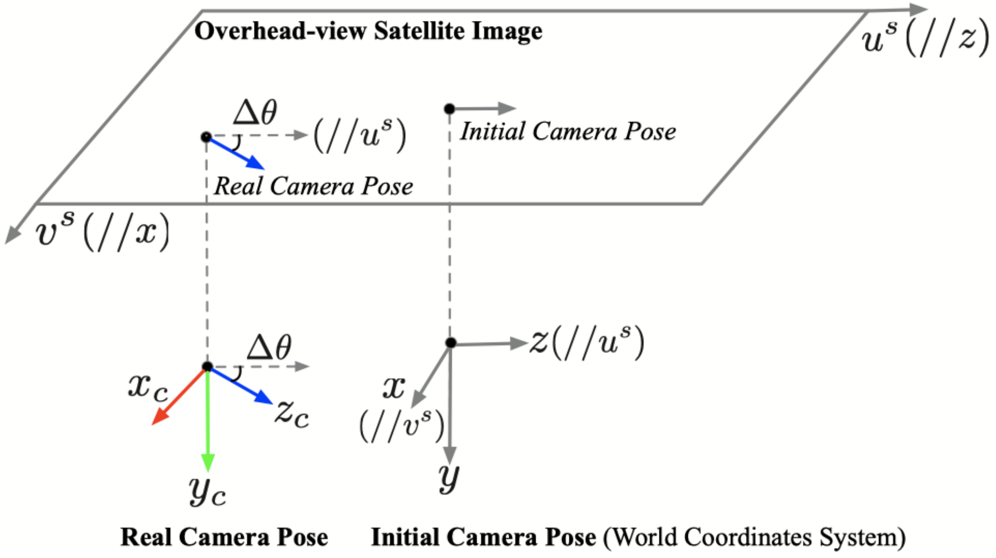

Fig. 2 Coordinates illustration.

From ground pixel to satellite pixel:

$$\begin{bmatrix} u_s \\ v_s \\ z \end{bmatrix} = \begin{bmatrix} \frac{1}{\alpha} & 0 & 0 \\ 0 & \frac{1}{\alpha} & 0 \\ 0 & 0 & 1 \end{bmatrix} \left(w\mathbf{R}\mathbf{K}^{-1} \begin{bmatrix} u_g \\ v_g \\ 1 \end{bmatrix} + \mathbf{R}\mathbf{t} \right) + \begin{bmatrix} u_s^0 \\ v_s^0 \\ 0 \end{bmatrix}. \qquad (4)$$

Results
– KITTI

Table 1. Performance comparison between our method and state-of-the-art methods on KITTI dataset.

	Test1									Test2								
	Lateral			Longitudinal			Azimuth			Lateral			Longitudinal			Azimuth		
	$d=1$	$d=3$	$d=5$	$d=1$	$d=3$	$d=5$	$\theta=1$	$\theta=3$	$\theta=5$	$d=1$	$d=3$	$d=5$	$d=1$	$d=3$	$d=5$	$\theta=1$	$\theta=3$	$\theta=5$
CVM-NET [17]	5.83	17.41	28.78	3.47	11.18	18.42	-	-	-	6.96	21.55	35.24	3.58	10.45	17.53	-	-	-
CVFT [49]	7.71	22.37	36.28	3.82	11.48	18.63	-	-	-	7.20	22.05	36.21	3.63	11.11	18.46	-	-	-
SAFA [46]	9.49	29.31	46.44	4.35	12.46	21.10	-	-	-	9.15	27.83	44.27	4.22	11.93	19.65	-	-	-
Polar-SAFA [46]	9.57	30.08	45.83	4.56	13.01	21.12	-	-	-	10.02	29.09	46.19	3.82	11.87	19.84	-	-	-
DSM [47]	10.12	30.67	48.24	4.08	12.01	20.14	3.58	13.81	24.44	10.77	31.37	48.24	3.87	11.73	19.50	3.53	14.09	23.95
VIGOR [70]	18.61	49.06	69.79	4.29	13.01	21.47	-	-	-	17.38	48.20	70.79	4.07	12.52	20.14	-	-	-
Ours	**35.54**	**70.77**	**80.36**	**5.22**	**15.88**	**26.13**	**19.64**	**51.76**	**71.72**	**27.82**	**59.79**	**72.89**	**5.75**	**16.36**	**26.48**	**18.42**	**49.72**	**71.00**

Results

(1) SGD Vs. (2) ADAM. Vs. (3) LM

Table 4. Performance comparison by using different optimizers on the KITTI dataset.

	Test1									Test2								
	Lateral			Longitudinal			Azimuth			Lateral			Longitudinal			Azimuth		
	$d=1$	$d=3$	$d=5$	$d=1$	$d=3$	$d=5$	$\theta=1$	$\theta=3$	$\theta=5$	$d=1$	$d=3$	$d=5$	$d=1$	$d=3$	$d=5$	$\theta=1$	$\theta=3$	$\theta=5$
SGD	16.86	39.60	51.15	4.72	15.29	25.39	10.05	30.37	49.80	16.06	38.41	50.29	5.00	15.34	25.70	9.98	30.03	50.13
ADAM	7.13	21.15	32.97	4.96	15.13	25.63	10.36	30.32	50.49	7.33	21.36	33.52	5.64	15.38	26.00	10.28	30.81	50.91
LM (**Ours**)	**35.54**	**70.77**	**80.36**	**5.22**	**15.88**	**26.13**	19.64	51.76	71.72	**27.82**	**59.79**	**72.89**	**5.75**	**16.36**	**26.48**	18.42	49.72	71.00

Search region: 40 m x 40 m, 20-degree

Results

(1) Net. Vs. (2) LM

Table 4. Performance comparison by using different optimizers on the KITTI dataset.

	Test1									Test2								
	Lateral			Longitudinal			Azimuth			Lateral			Longitudinal			Azimuth		
	$d=1$	$d=3$	$d=5$	$d=1$	$d=3$	$d=5$	$\theta=1$	$\theta=3$	$\theta=5$	$d=1$	$d=3$	$d=5$	$d=1$	$d=3$	$d=5$	$\theta=1$	$\theta=3$	$\theta=5$
Net	27.14	58.28	71.91	4.53	15.19	25.36	**45.56**	**93.19**	**99.76**	20.26	53.94	67.42	5.40	15.82	25.58	**42.03**	**92.32**	**99.81**
LM (**Ours**)	**35.54**	**70.77**	**80.36**	**5.22**	**15.88**	**26.13**	19.64	51.76	71.72	**27.82**	**59.79**	**72.89**	**5.75**	**16.36**	**26.48**	18.42	49.72	71.00

Search region: 40 m x 40 m, 20-degree

Pose Update Visualization

Weaknesses

Query Reference Query Reference

- Red arrow: initial pose
- Blue arrow: GT pose
- Green arrow: predicted pose
- Purpose dots: intermediate pose

Part-3 : ECCV2022

Visual Cross-View Metric Localization with Dense Uncertainty Estimates

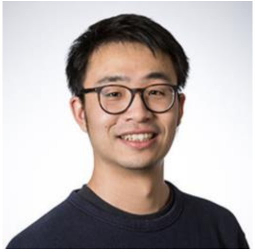
Zimin Xia


Olaf Booij


Marco Manfredi


Julian F. P. Kooij

Intelligent Vehicles Group, Technical University Delft, The Netherlands

TomTom, Amsterdam, The Netherlands

Code

Paper

Probabilistic Estimation

G

S

$p(X|G,S)$

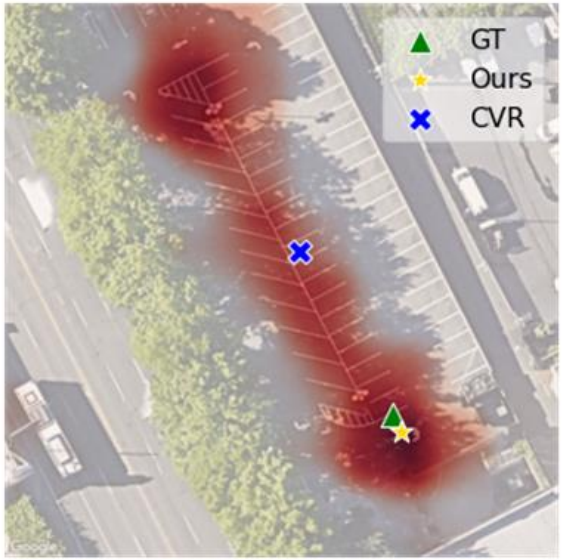

Framework

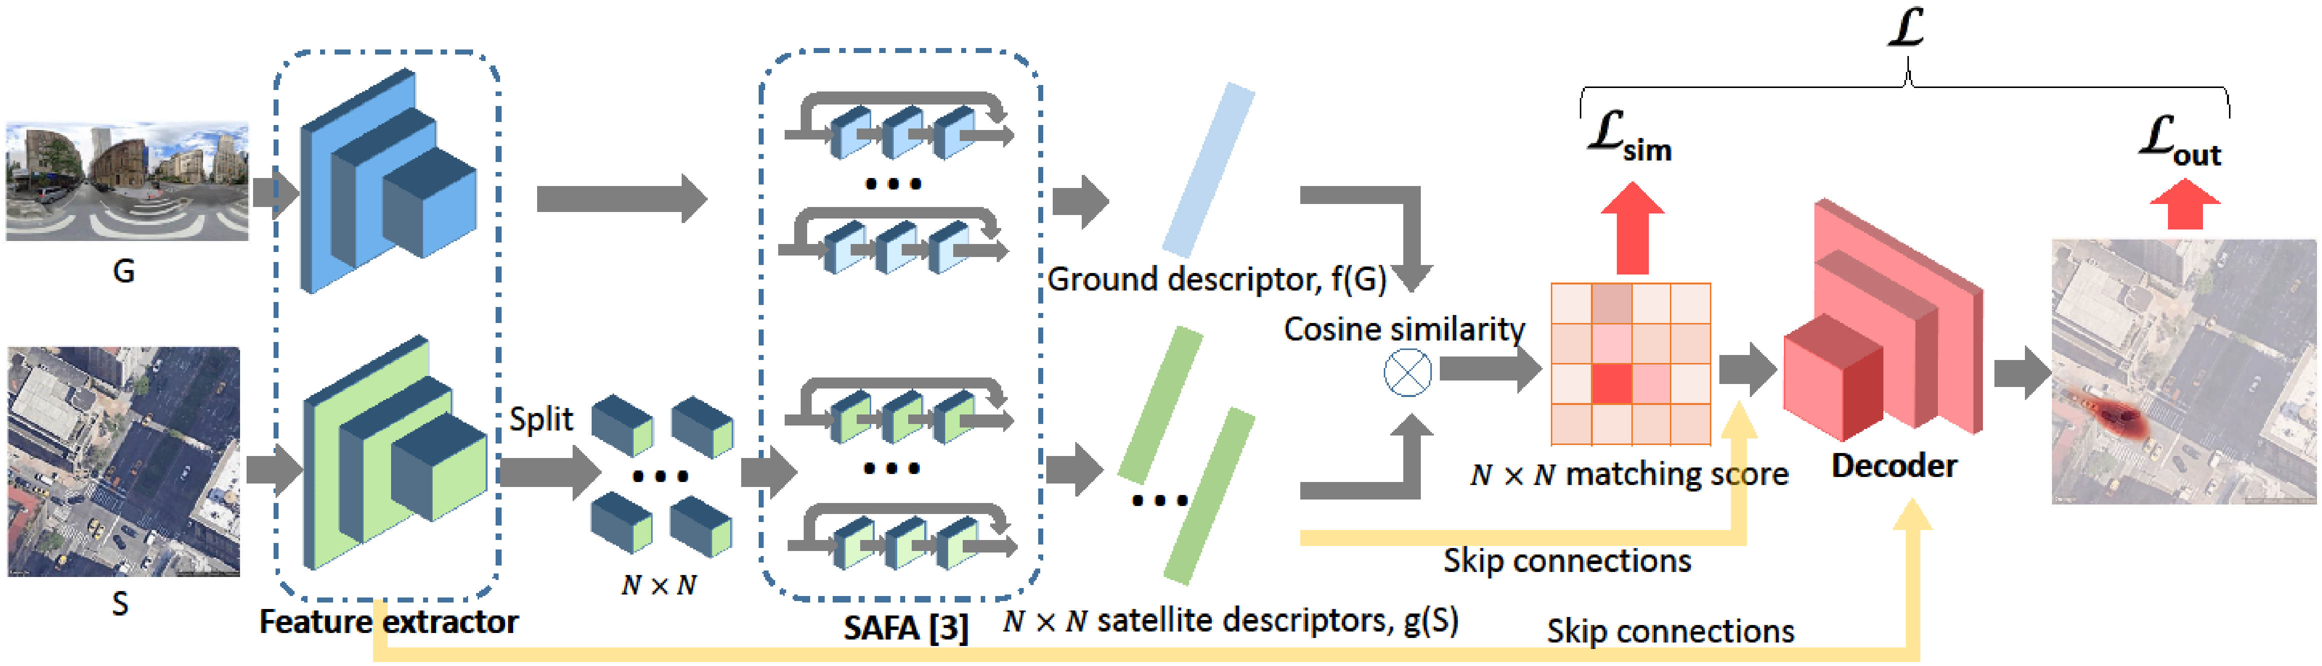

G
S
Feature extractor
Split
$N \times N$
SAFA [3]
Ground descriptor, f(G)
$N \times N$ satellite descriptors, g(S)
Cosine similarity
$\mathcal{L}$
$\mathcal{L}_{sim}$
$\mathcal{L}_{out}$
$N \times N$ matching score
Decoder
Skip connections
Skip connections

Visualization

JUNE 18-22, 2023 CVPR VANCOUVER, CANADA

Australian National University

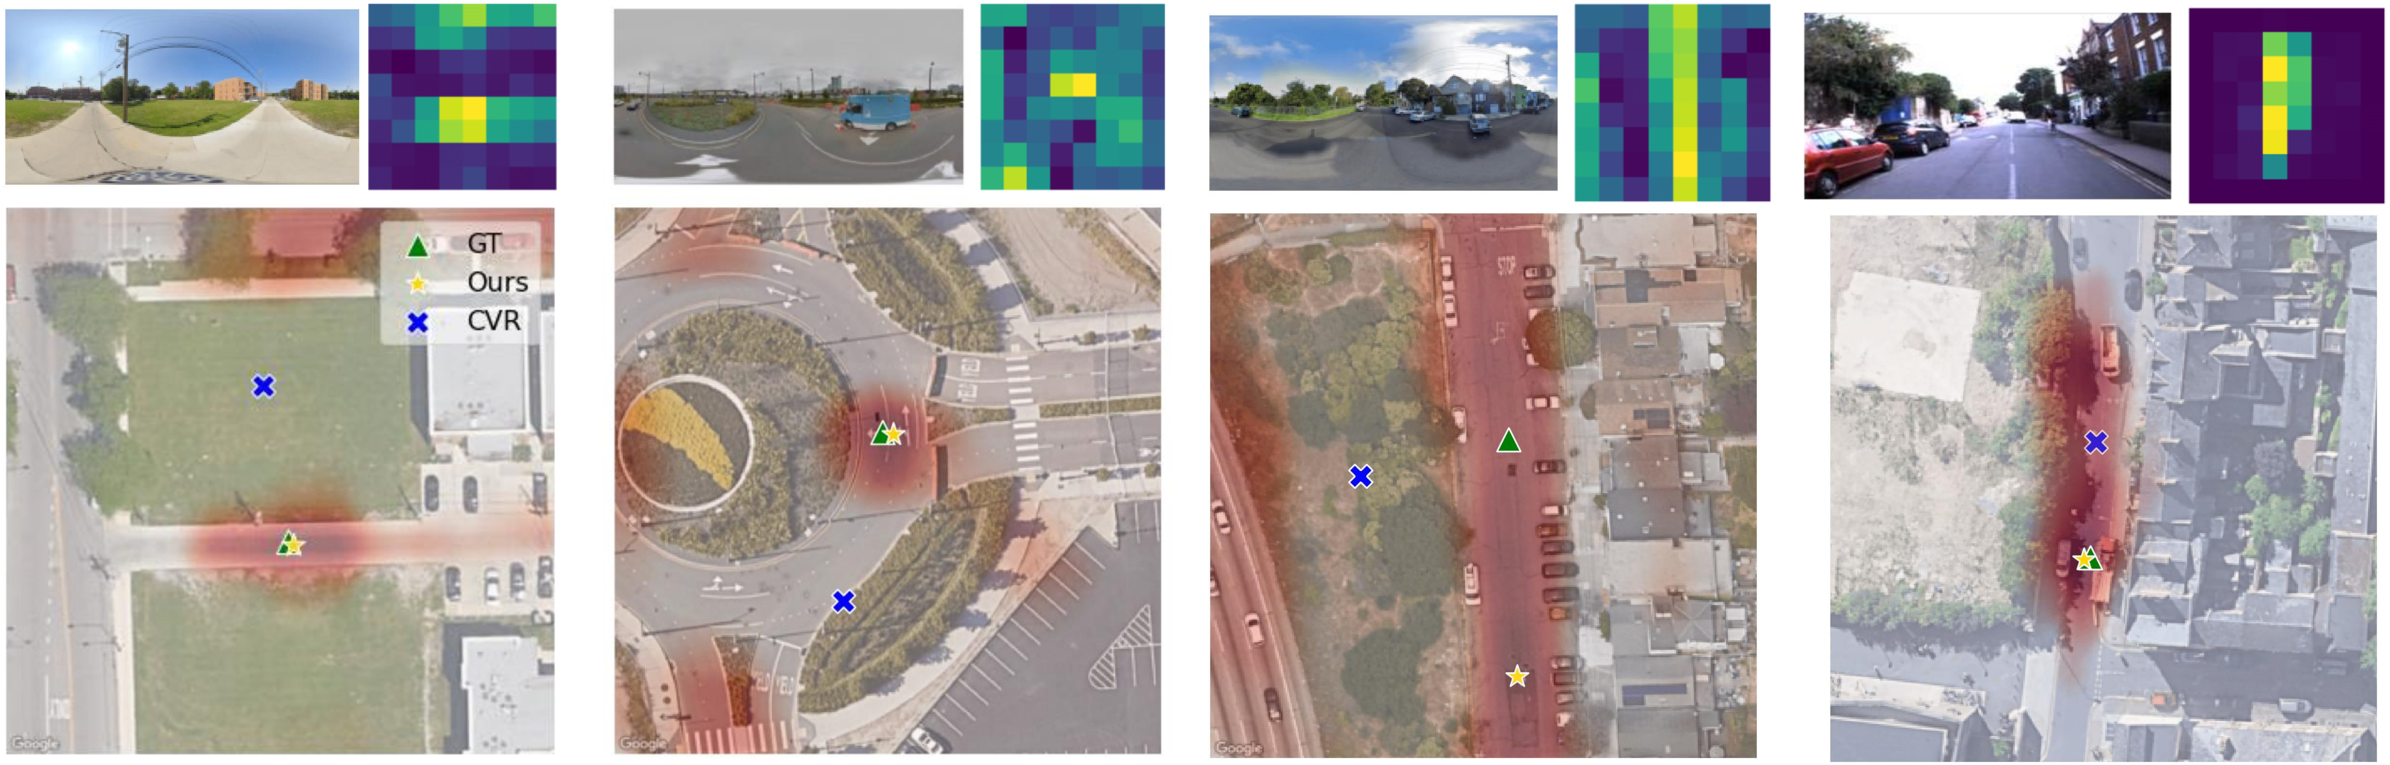

Part-4: CVPR 2023

SliceMatch: Geometry-guided Aggregation for Cross-View Pose Estimation


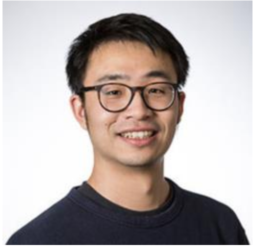

Ted Lentsch* | Zimin Xia* | Holger Caesar | Julian F. P. Kooij

Intelligent Vehicles Group, Delft University of Technology, The Netherlands

*Equal contribution

Poster ID: THU-AM-071

Code

Paper

Insight

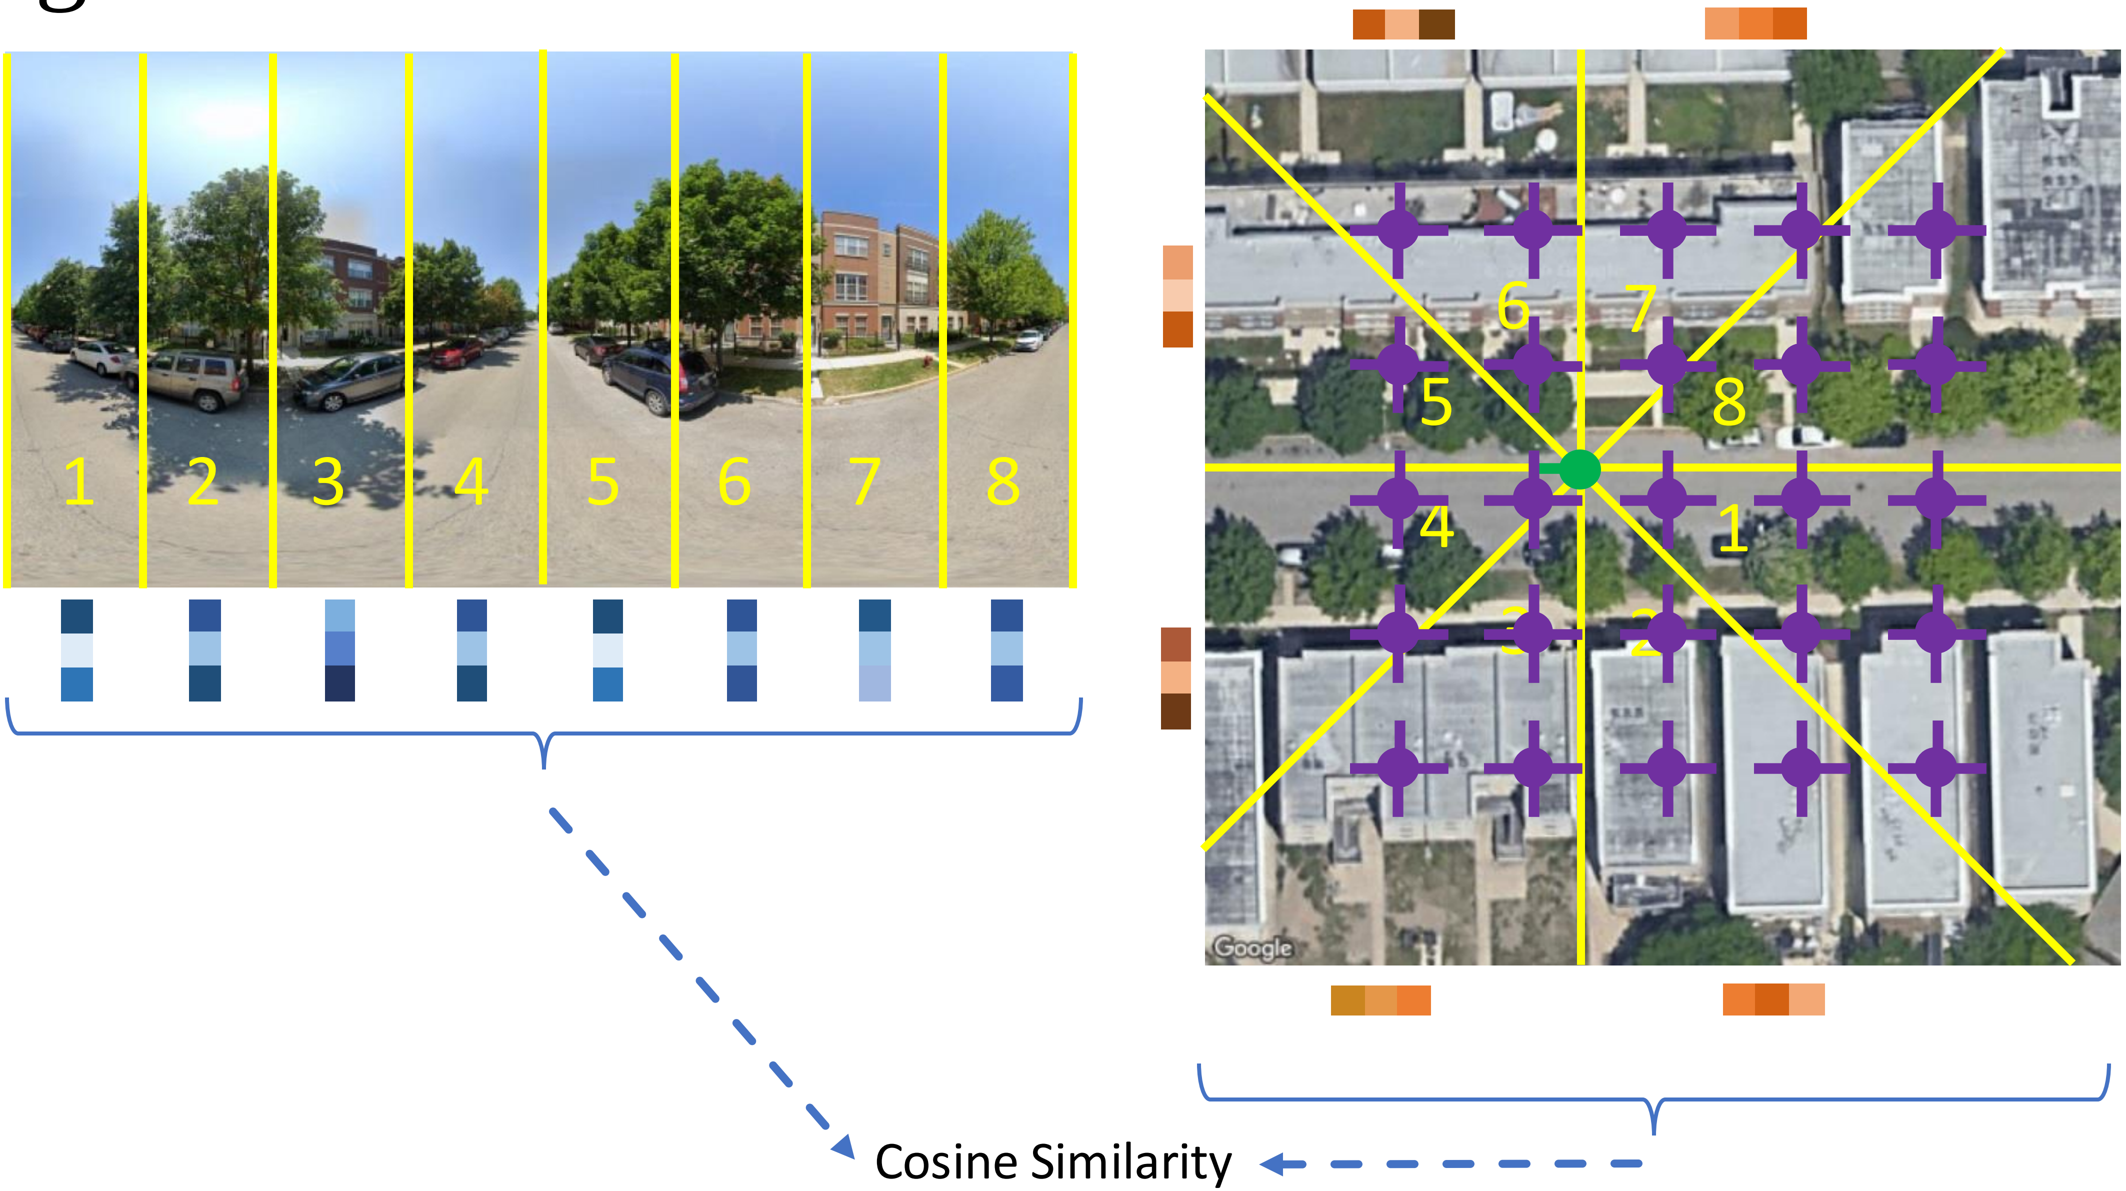

Framework

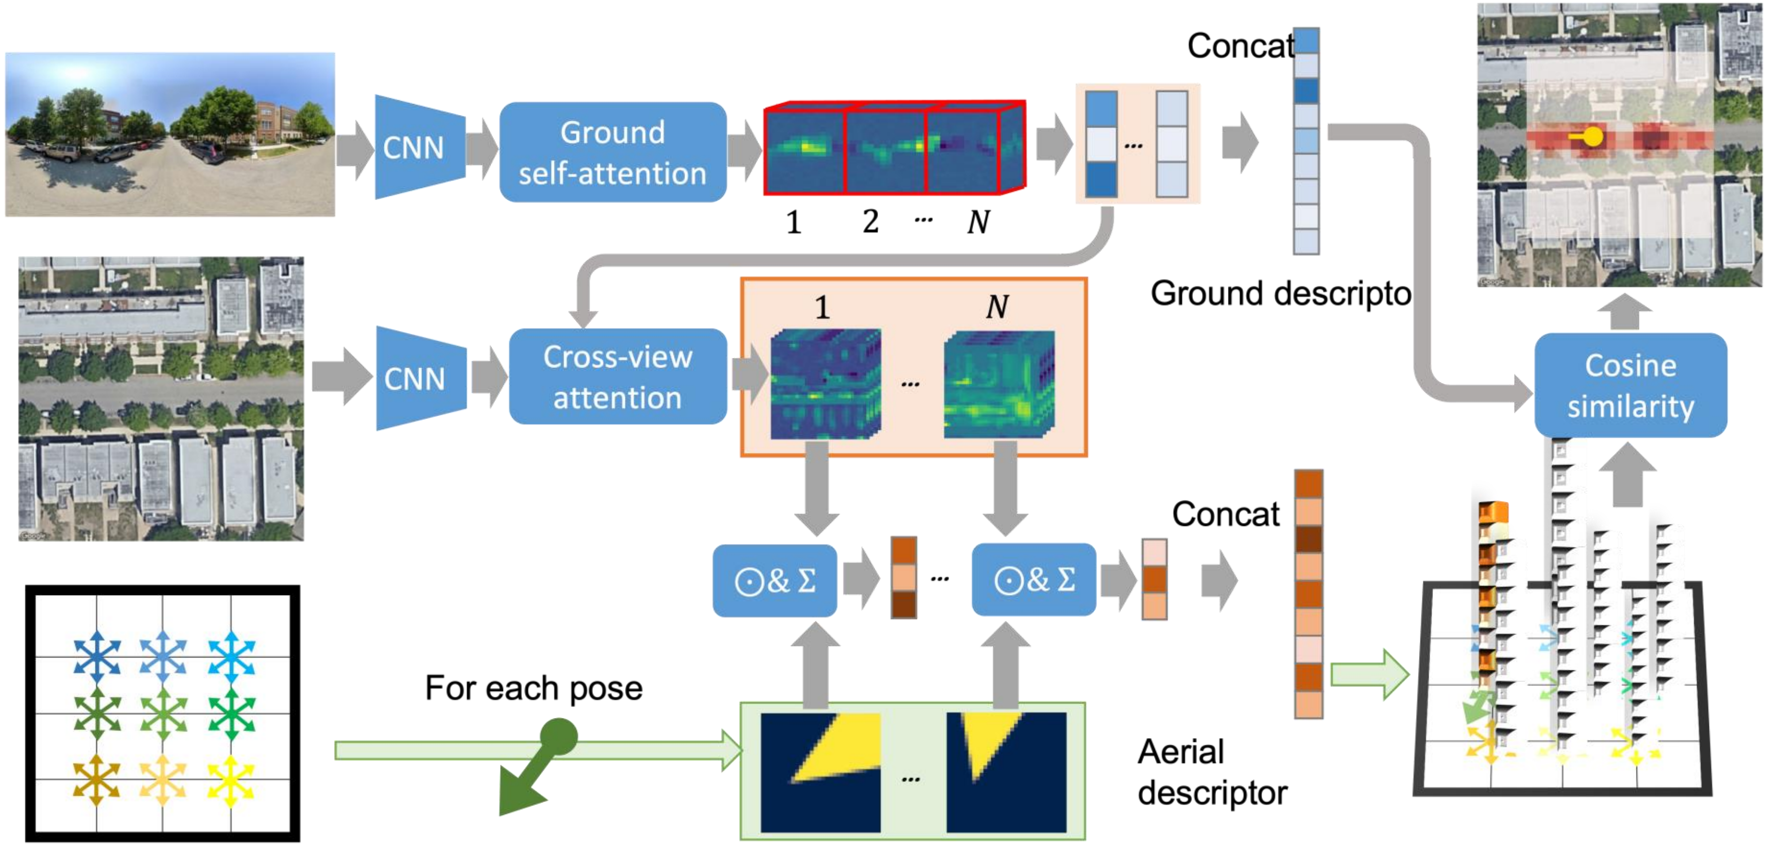

CNN
Ground self-attention
1
2
...
N
Concat
Ground descripto
CNN
Cross-view attention
1
N
...
⊙& Σ
⊙& Σ
Concat
For each pose
Aerial descriptor
Cosine similarity

JUNE 18-22, 2023 CVPR VANCOUVER, CANADA

Australian National University

More details and results can be found in the paper

Poster ID: THU-AM-071

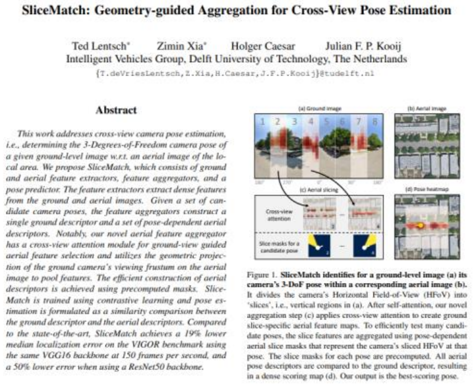

SliceMatch: Geometry-guided Aggregation for Cross-View Pose Estimation

Ted Lentsch* Zimin Xia* Holger Caesar Julian F. P. Kooij
Intelligent Vehicles Group, Delft University of Technology, The Netherlands
{T.deVriesLentsch,Z.Xia,H.Caesar,J.F.P.Kooij}@tudelft.nl

Abstract

This work addresses cross-view camera pose estimation, i.e., determining the 3-Degrees-of-Freedom camera pose of a given ground-level image w.r.t. an aerial image of the local area. We propose SliceMatch, which consists of ground and aerial feature extractors, feature aggregators, and a pose predictor. The feature extractors extract dense features from the ground and aerial images. Given a set of candidate camera poses, the feature aggregators construct a single ground descriptor and a set of pose-dependent aerial descriptors. Notably, our novel aerial feature aggregator has a cross-view attention module for ground-view guided aerial feature selection and utilizes the geometric projection of the ground camera's viewing frustum on the aerial image to pool features. The efficient construction of aerial descriptors is achieved using precomputed masks. SliceMatch is trained using contrastive learning and pose estimation is formulated as a similarity comparison between the ground descriptor and the aerial descriptors. Compared to the state-of-the-art, SliceMatch achieves a 19% lower median localization error on the VIGOR benchmark using the same VGG16 backbone at 150 frames per second, and a 50% lower error when using a ResNet50 backbone.

Figure 1. **SliceMatch identifies for a ground-level image (a) its camera's 3-DoF pose within a corresponding aerial image (b).** It divides the camera's Horizontal Field-of-View (HFoV) into 'slices', i.e., vertical regions in (a). After self-attention, our novel aggregation step (c) applies cross-view attention to create ground slice-specific aerial feature maps. To efficiently test many candidate poses, the slice features are aggregated using pose-dependent aerial slice masks that represent the camera's sliced HFoV at that pose. The slice masks for each pose are precomputed. All aerial pose descriptors are compared to the ground descriptor, resulting in a dense scoring map (d). Our output is the best-scoring pose.

https://github.com/tudelft-iv/SliceMatch

https://arxiv.org/abs/2211.14651

JUNE 18-22, 2023 CVPR VANCOUVER, CANADA

Australian National University

Satellite to Street-view Panorama Synthesis

TPAMI 2022

Task Description

- Given a satellite image and a precise location

- *What does it look like down there, as if one was standing right there?*

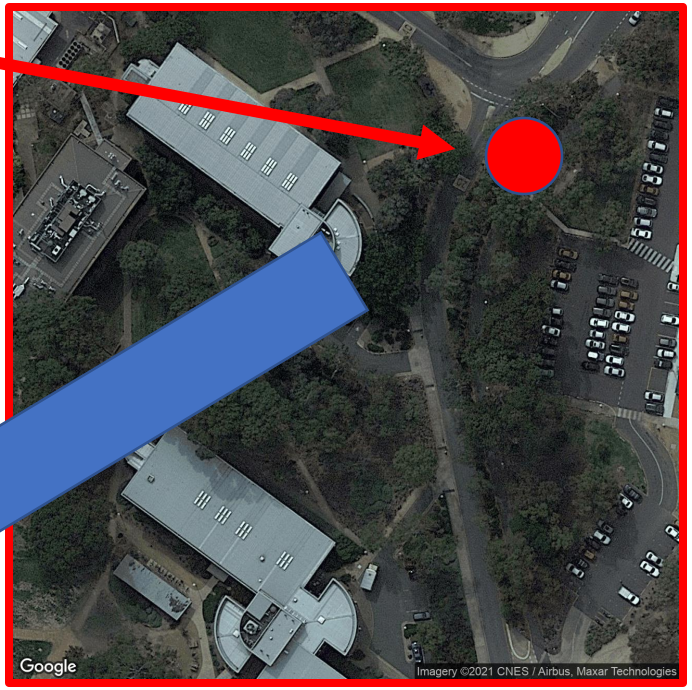

Satellite to Street-view Panorama Synthesis

Input: Satellite Image

Desired output: Street-View Panorama

Challenges:

(1) severe occlusions;
(2) different imaging modalities;
(3) seasonal / weather differences, et al.

Shi, Yujiao, et al. "Geometry-guided street-view panorama synthesis from satellite imagery." TPAMI 2022.

Geometry Correspondences

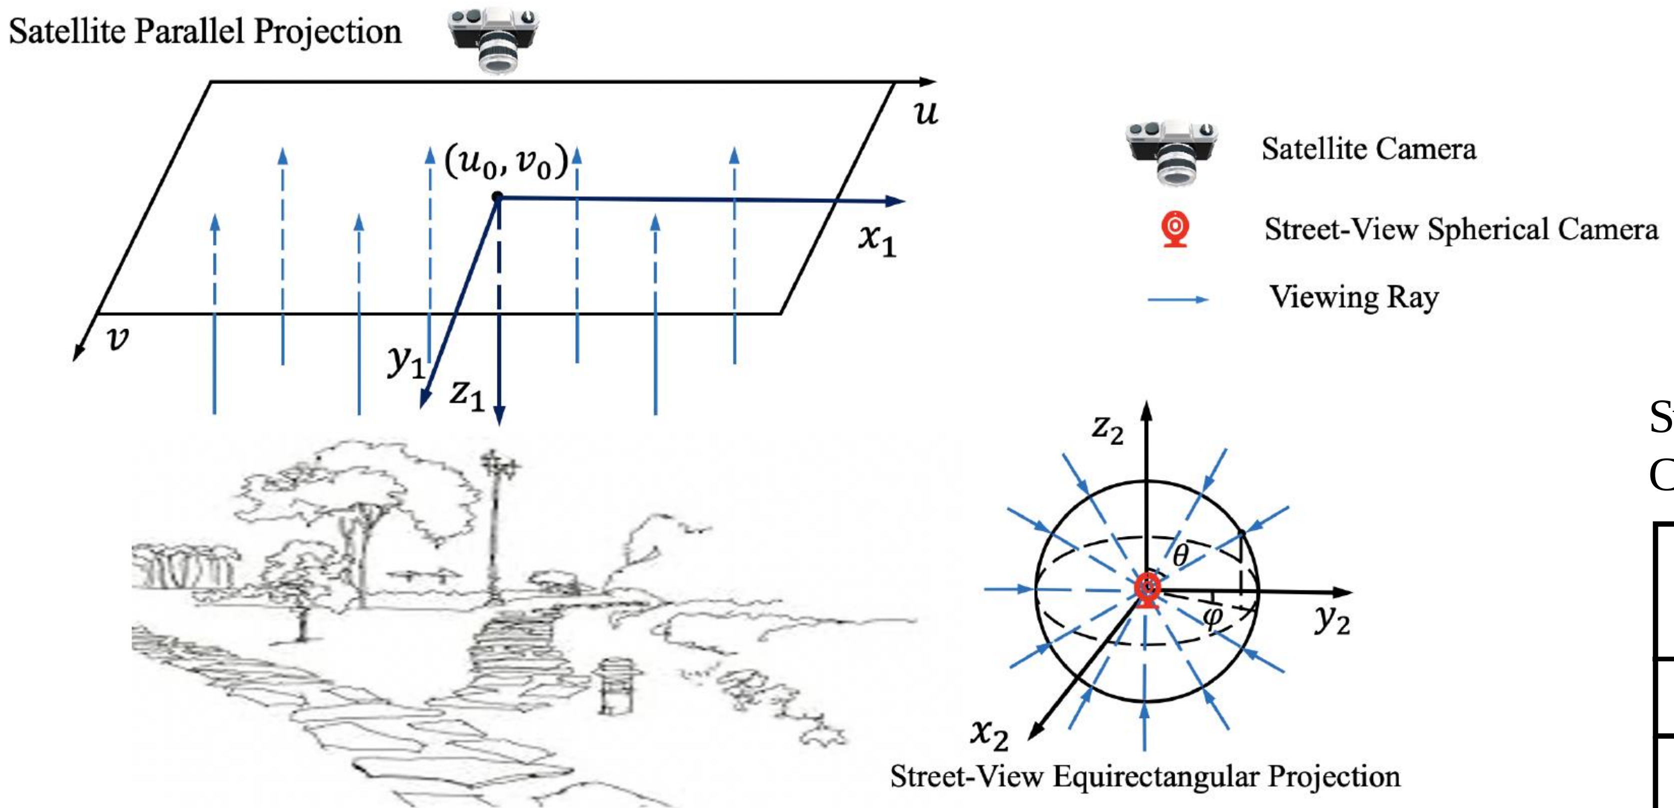

Street-view Panorama
Coordinate System

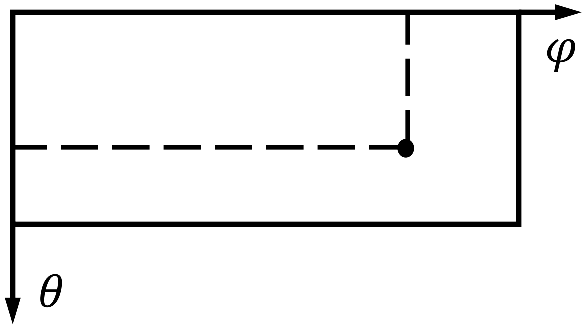

There is a deterministic mapping between the satellite to street-view cameras, ***determined by the heights of scene objects.***

Multi (height) Plane Image (MPI)

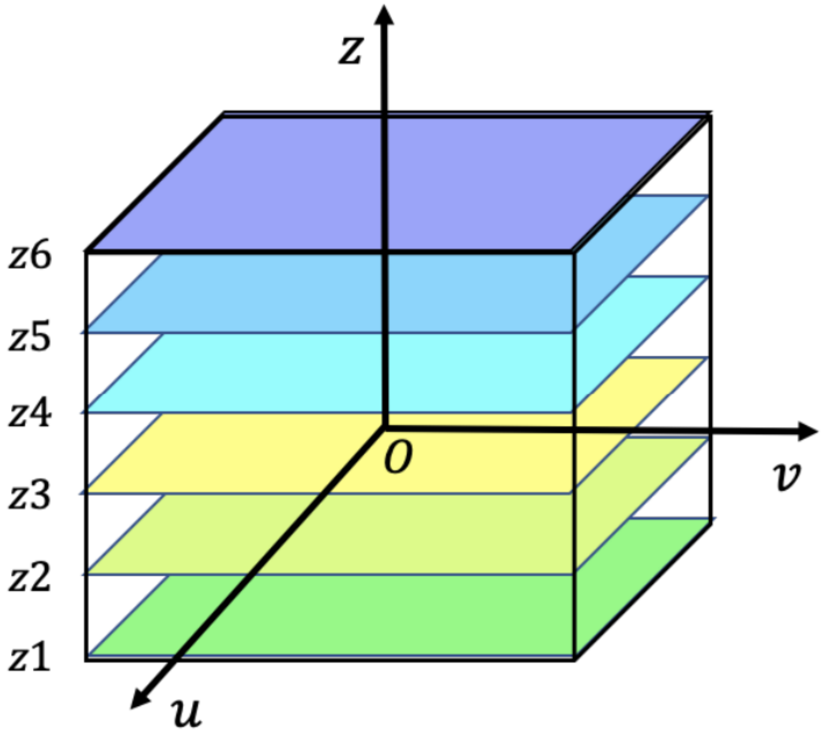

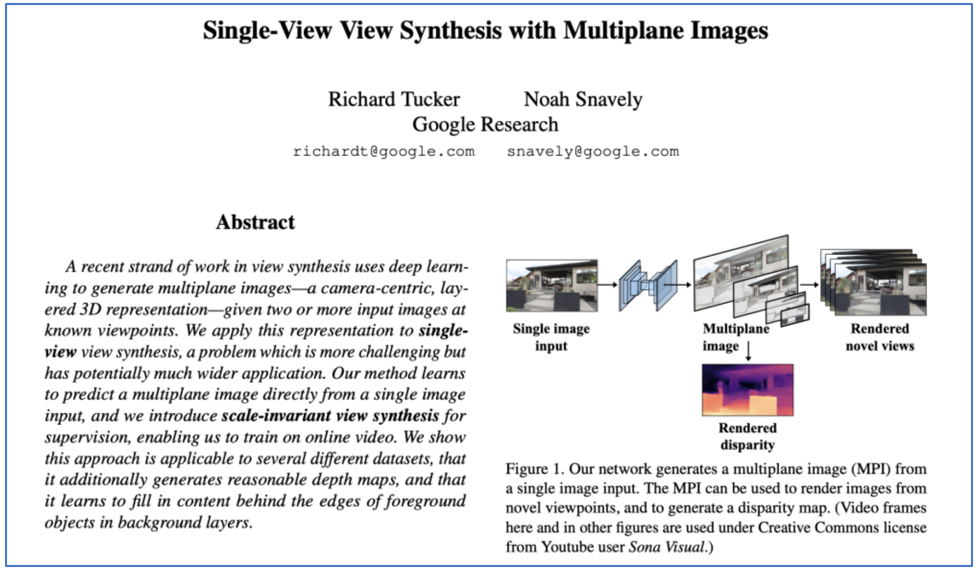

Single-View View Synthesis with Multiplane Images

Richard Tucker Noah Snavely
Google Research
richardt@google.com snavely@google.com

Abstract

*A recent strand of work in view synthesis uses deep learning to generate multiplane images—a camera-centric, layered 3D representation—given two or more input images at known viewpoints. We apply this representation to **single-view** view synthesis, a problem which is more challenging but has potentially much wider application. Our method learns to predict a multiplane image directly from a single image input, and we introduce **scale-invariant view synthesis** for supervision, enabling us to train on online video. We show this approach is applicable to several different datasets, that it additionally generates reasonable depth maps, and that it learns to fill in content behind the edges of foreground objects in background layers.*

Figure 1. Our network generates a multiplane image (MPI) from a single image input. The MPI can be used to render images from novel viewpoints, and to generate a disparity map. (Video frames here and in other figures are used under Creative Commons license from Youtube user *Sona Visual*.)

Satellite to Street-view Projection

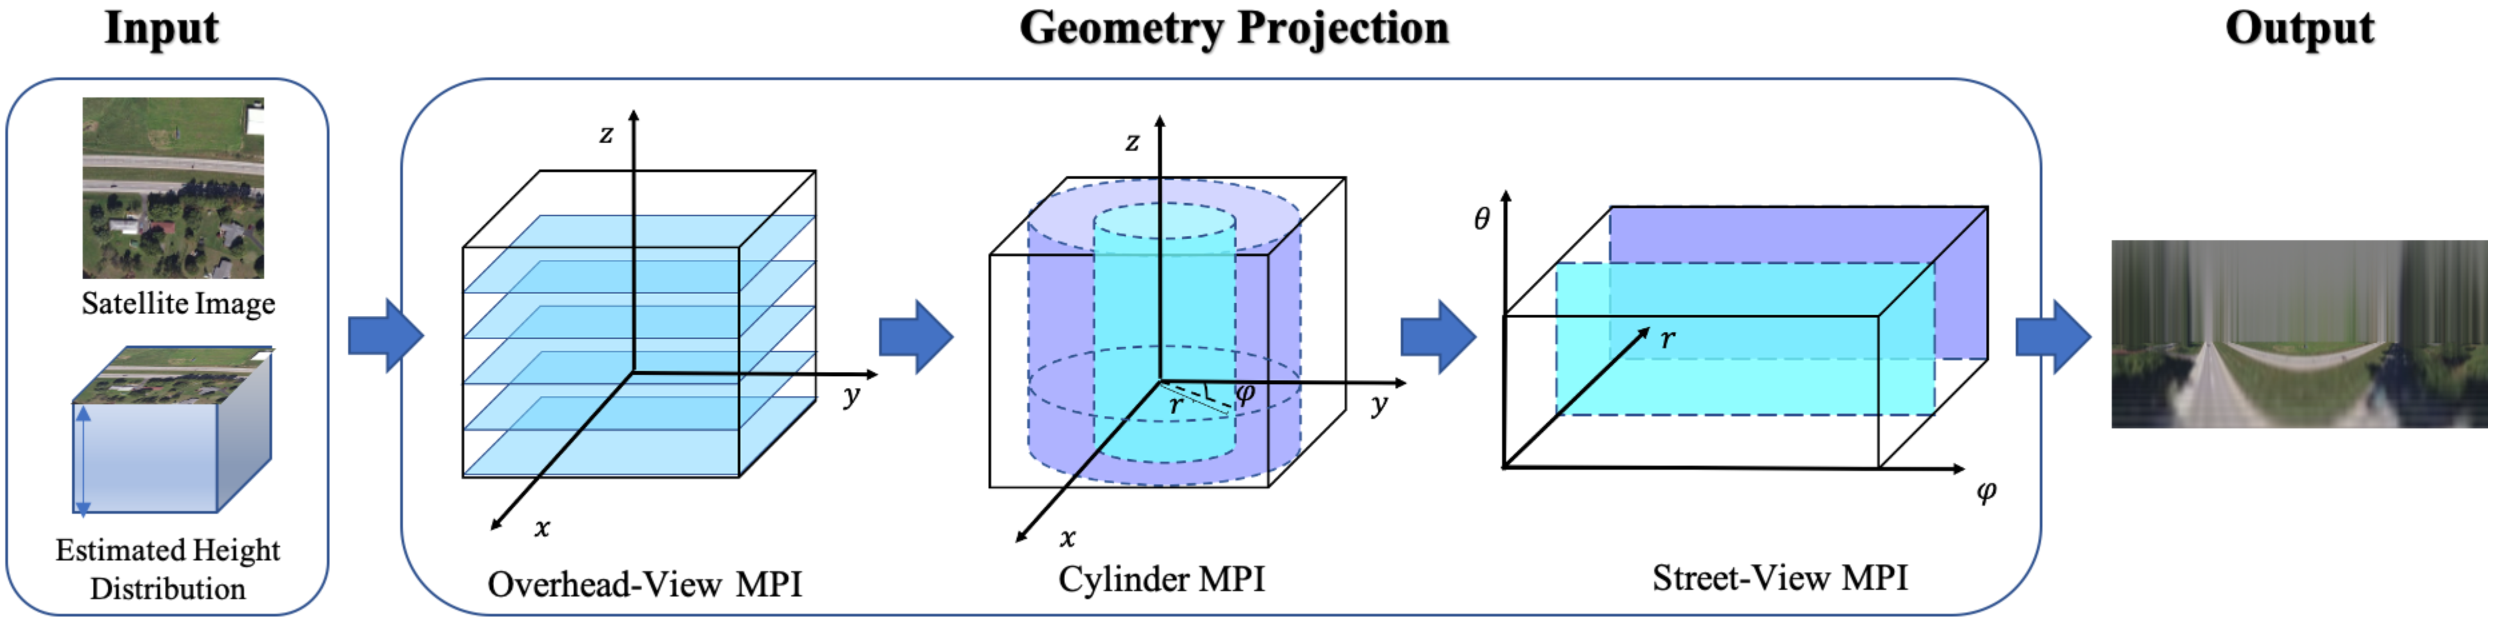

Framework

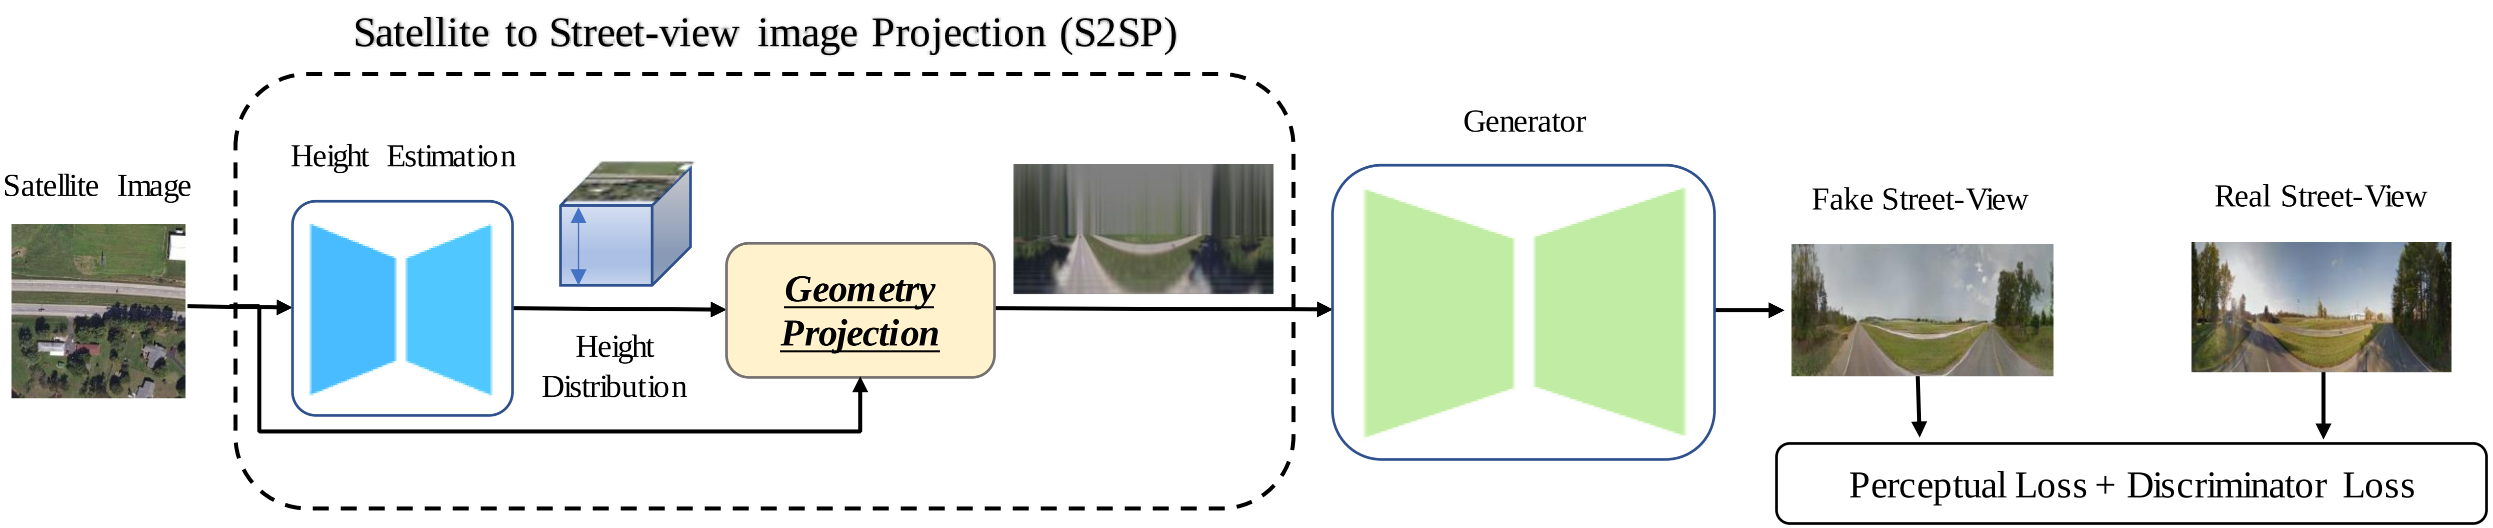
Satellite to Street-view image Projection (S2SP)
Satellite Image
Height Estimation
Height Distribution
Geometry Projection
Generator
Fake Street-View
Real Street-View
Perceptual Loss + Discriminator Loss

Intermediate Visualization

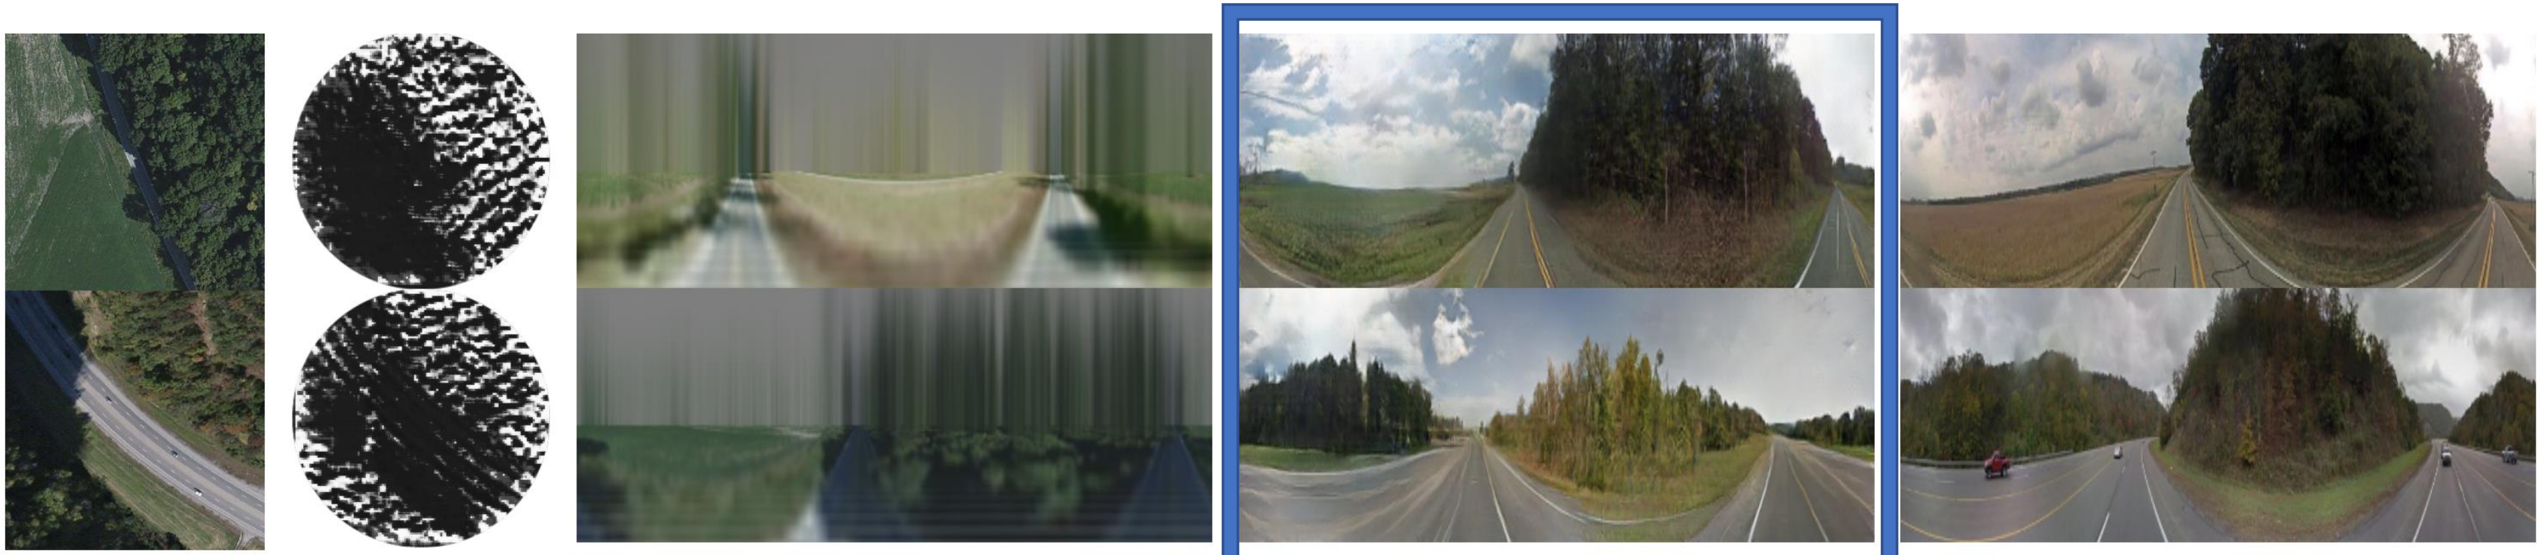

(a) Satellite Image

(b) Height Map

(c) Projected Satellite Image (Intermediate)

(d) Generated Street-View Panorama (Final)

(e) Real Street-View Panorama (Ground Truth)

Height Map Visualization

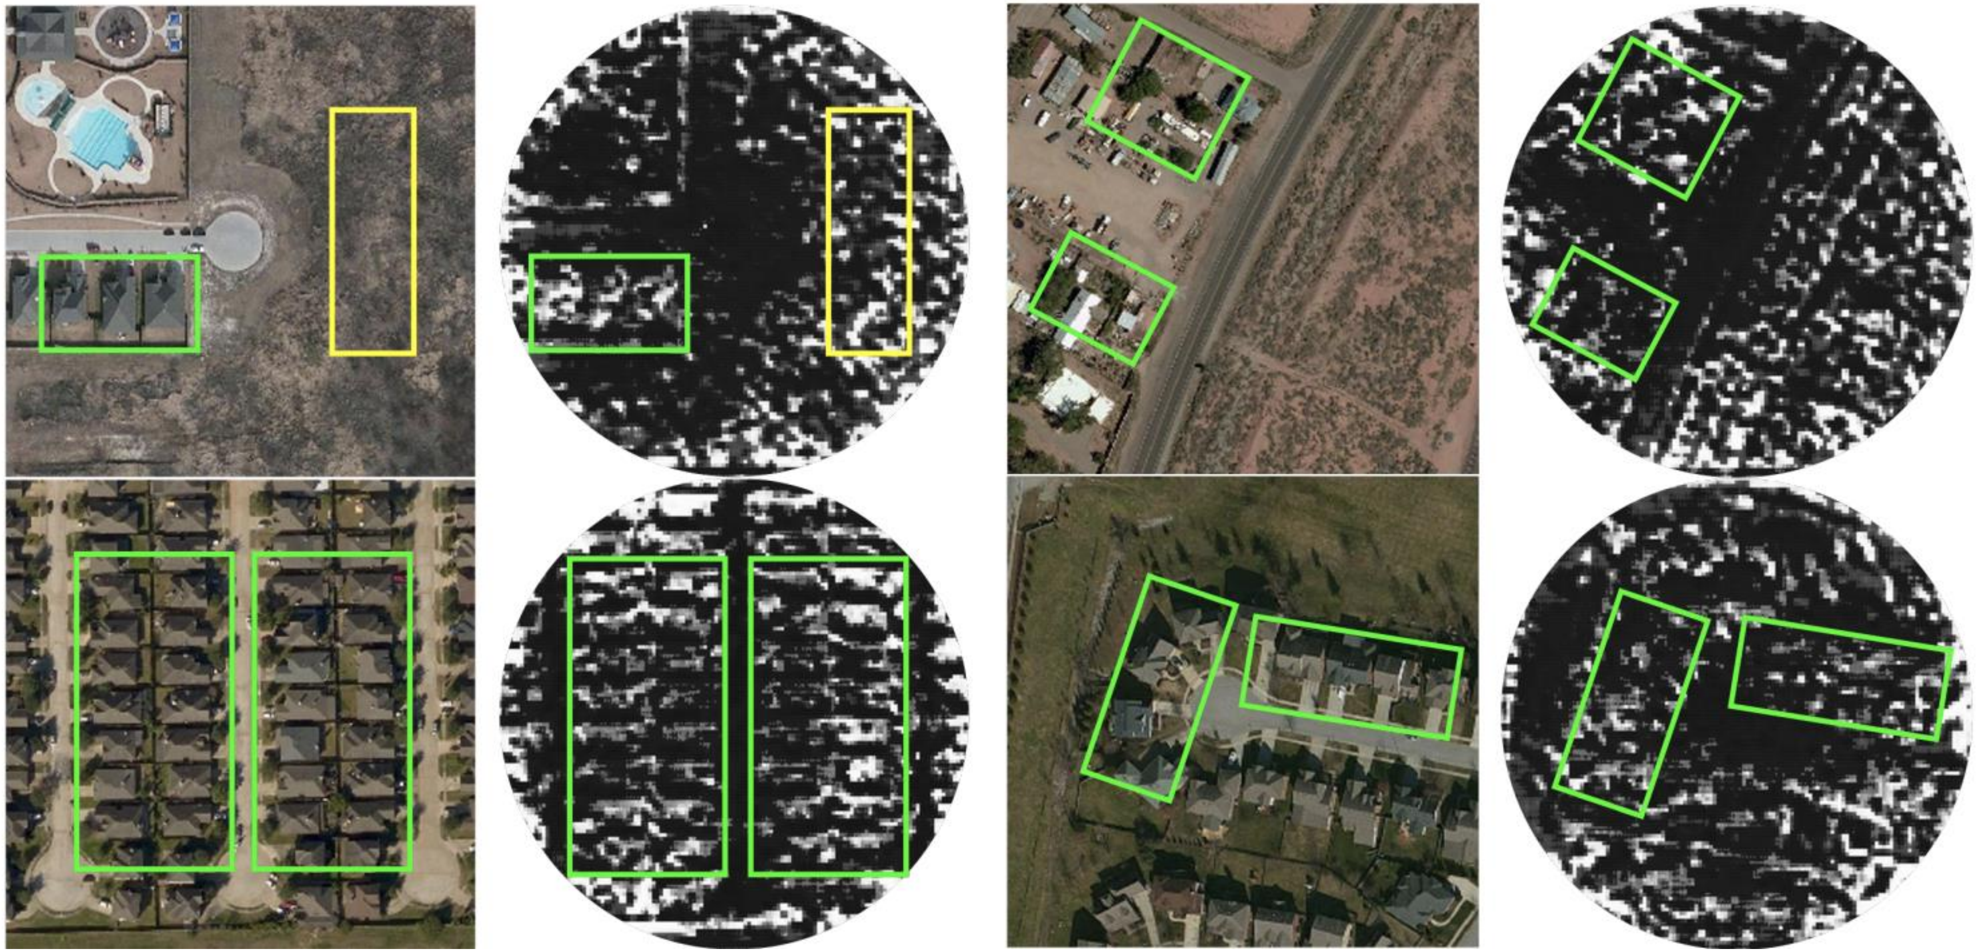

W or W/O Height Estimation

(a) Satellite Image (b) Ours w/o Height (Projection) (c) Ours (d) Ground Truth

Comparison with Other Works

Satellite Image   Pix2Pix [23]   XFork [10]   Ours   Ground Truth

Comparison with Other Works

Satellite Image | Pix2Pix [23] | XFork [10] | Ours | Ground Truth

Other Related Works

1. Vo, Nam N., and James Hays. "Localizing and orienting street views using overhead imagery." ECCV 2016.
2. Tian, Yicong, Chen Chen, and Mubarak Shah. "Cross-view image matching for geo-localization in urban environments." CVPR 2017.
3. Hu, Sixing, et al. "Cvm-net: Cross-view matching network for image-based ground-to-aerial geo-localization." CVPR 2018.
4. Liu, Liu, and Hongdong Li. "Lending orientation to neural networks for cross-view geo-localization." CVPR 2019.
5. Regmi, Krishna, and Mubarak Shah. "Bridging the domain gap for ground-to-aerial image matching." ICCV 2019.
6. Cai, Sudong, et al. "Ground-to-aerial image geo-localization with a hard exemplar reweighting triplet loss." ICCV 2019.
7. Toker, Aysim, et al. "Coming down to earth: Satellite-to-street view synthesis for geo-localization." CVPR 2021.
8. Zhu, Sijie, Taojiannan Yang, and Chen Chen. "Vigor: Cross-view image geo-localization beyond one-to-one retrieval." CVPR 2021.
9. Yang, Hongji, Xiufan Lu, and Yingying Zhu. "Cross-view geo-localization with layer-to-layer transformer." NeurIPS 2021.
10. Zhu, Sijie, Taojiannan Yang, and Chen Chen. "Revisiting street-to-aerial view image geo-localization and orientation estimation." WACV 2021.
11. Zhu, Sijie, Mubarak Shah, and Chen Chen. "Transgeo: Transformer is all you need for cross-view image geo-localization." CVPR 2022.
12. Fervers, Florian, et al. "Uncertainty-aware Vision-based Metric Cross-view Geolocalization." CVPR 2023.
13. …

JUNE 18-22, 2023 CVPR VANCOUVER, CANADA

Australian National University

Thank you!