3D Line Mapping Revisited

THU-PM-080

Shaohui Liu, Yifan Yu, Rémi Pautrat, Marc Pollefeys, Viktor Larsson

Robust Scalable Pipeline for Mapping 3D Lines

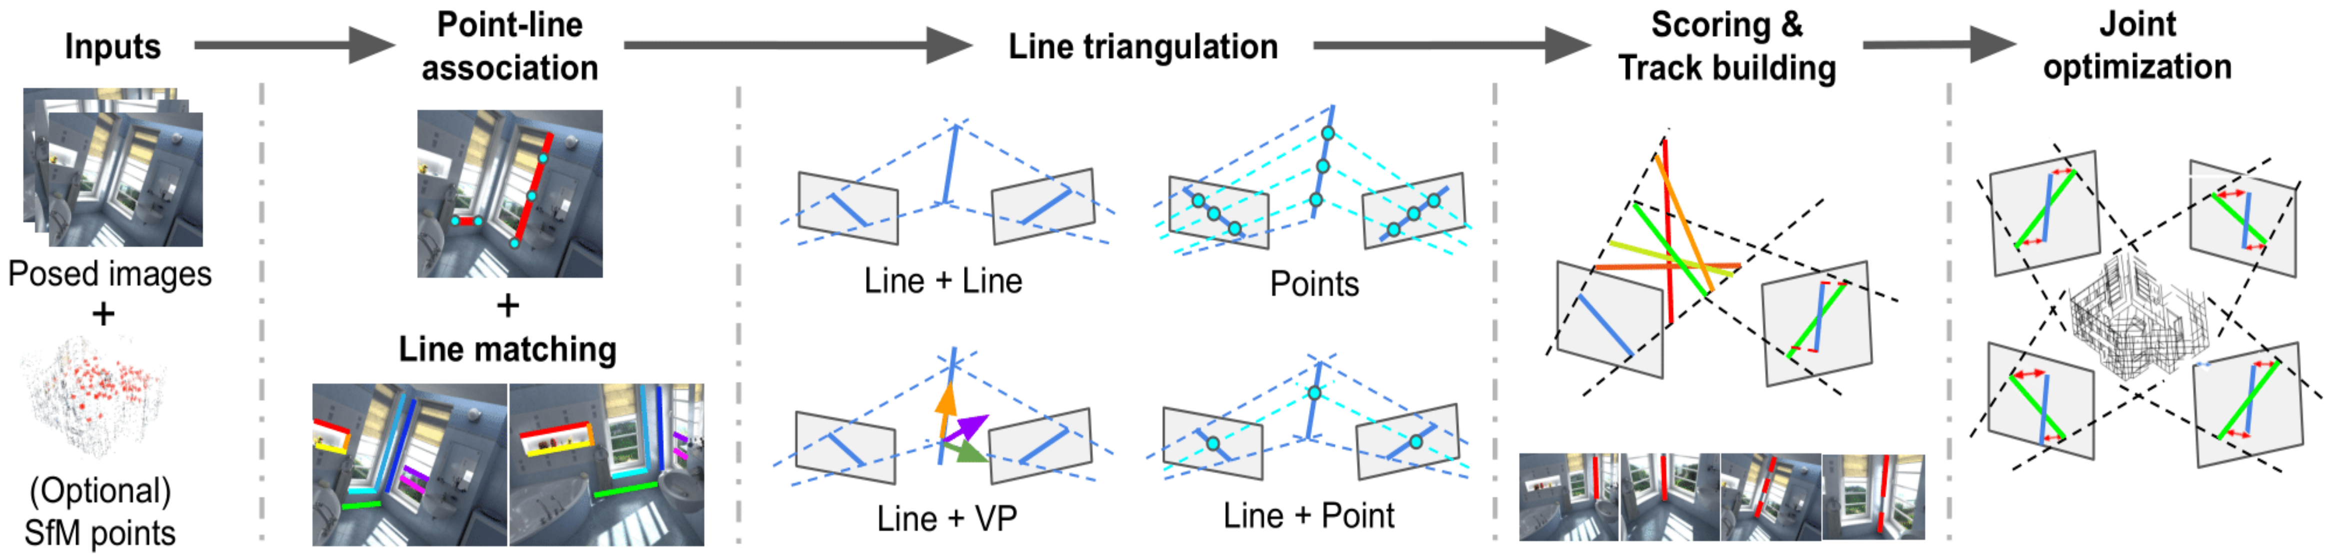

Examples of Our 3D Line Maps

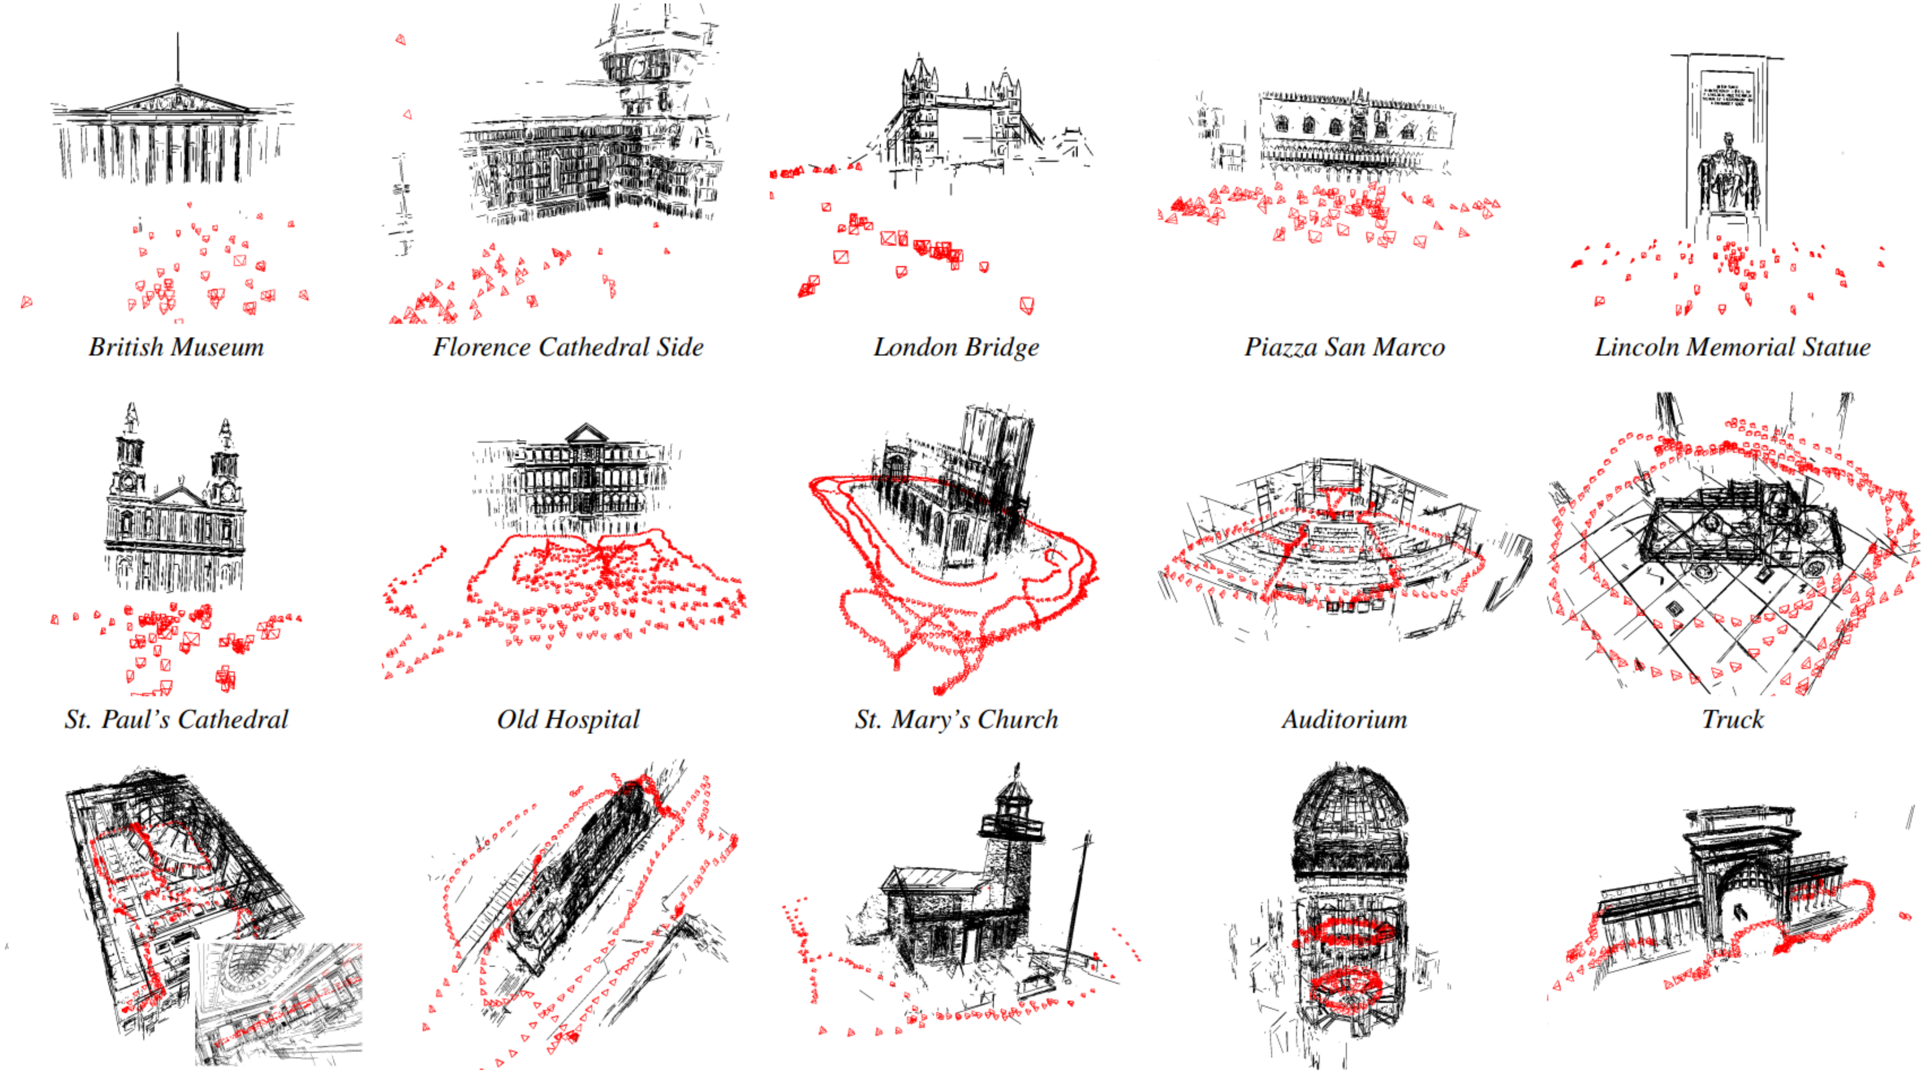
British Museum *Florence Cathedral Side* *London Bridge* *Piazza San Marco* *Lincoln Memorial Statue*
St. Paul's Cathedral *Old Hospital* *St. Mary's Church* *Auditorium* *Truck*
Courtroom (indoor and outdoor) *Train* *Lighthouse* *Museum* *Temple*

Opening up new possibilities to multiple applications

Example: Hybrid Localization

before optimization | after optimization | point-alone localization | hybrid point-line localization

Dataset	HLoc*	PtLine	Ours
Cambridge	7.0 / 0.13 / 44.0	7.4 / 0.13 / 43.5	**6.7 / 0.12 / 46.1**
7Scenes	3.3 / 1.08 / 73.0	3.3 / 1.09 / 72.7	**3.0 / 1.00 / 78.0**

Dataset	HLoc* w/ Depth	PtLine	Ours w/ Depth
7Scenes w/ depth	2.9 / 0.94 / 80.1	2.8 / 0.93 / 80.6	**2.6 / 0.87 / 83.5**

	(T / R) err. ↓	Acc. ↑
HLoc*	5.2 / 1.46	46.8
HLoc* w/ depth	4.7 / 1.25	53.4
PtLine	4.8 / 1.33	51.9
Ours w/ L3D++	4.1 / 1.14	60.8
Ours w/ LIMAP	**3.7 / 1.02**	**71.1**

Example: Hybrid Bundle Adjustment

	Med. error ↓	AUC @ (1° / 3° / 5°) ↑
COLMAP	0.188	77.3 / 89.0 / 91.6
COLMAP + LIMAP refinement	**0.146**	**82.9 / 91.2 / 93.0**

Example: Line-assisted stereo

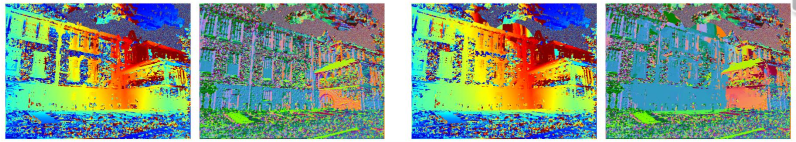

Original COLMAP MVS [45] | w. Line-based Energy

Open-sourcing – LIMAP:
a toolbox for mapping and localization with line features

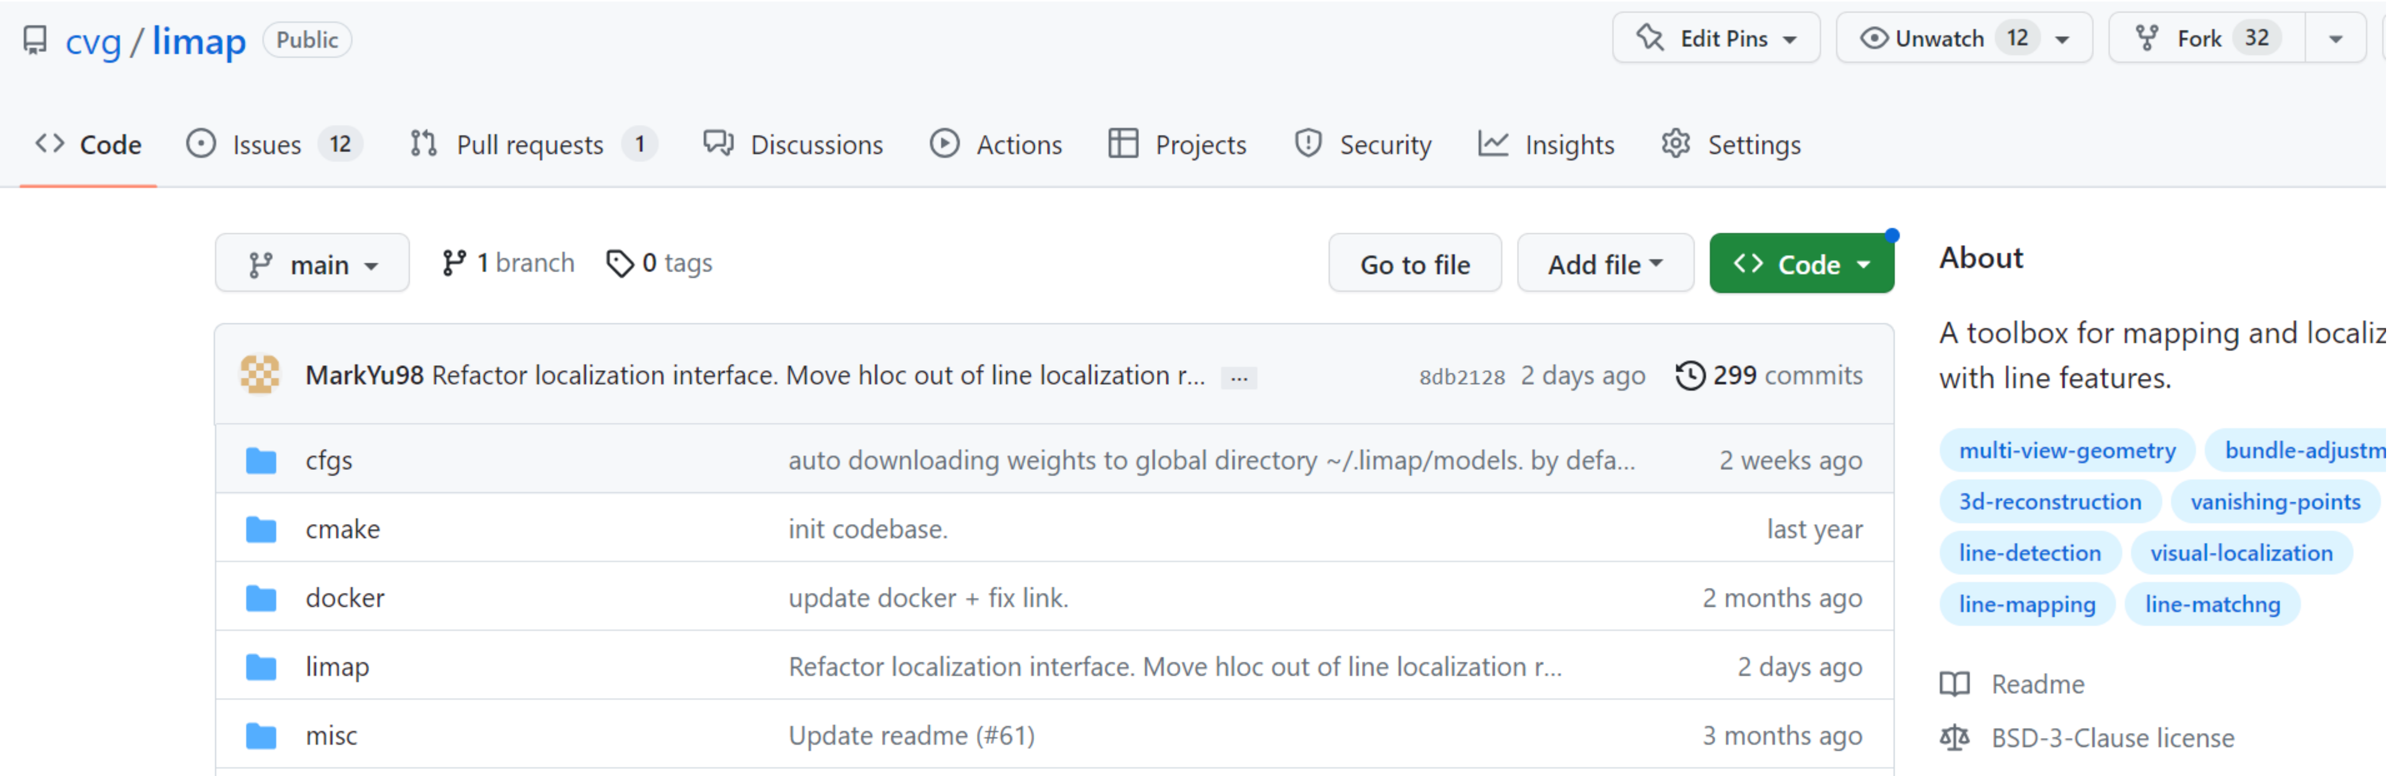

modular design + binded classes and interfaces in Python
supporting multiple line detectors, matchers and vanishing point estimators

Snavely et al. Bundler (2010)

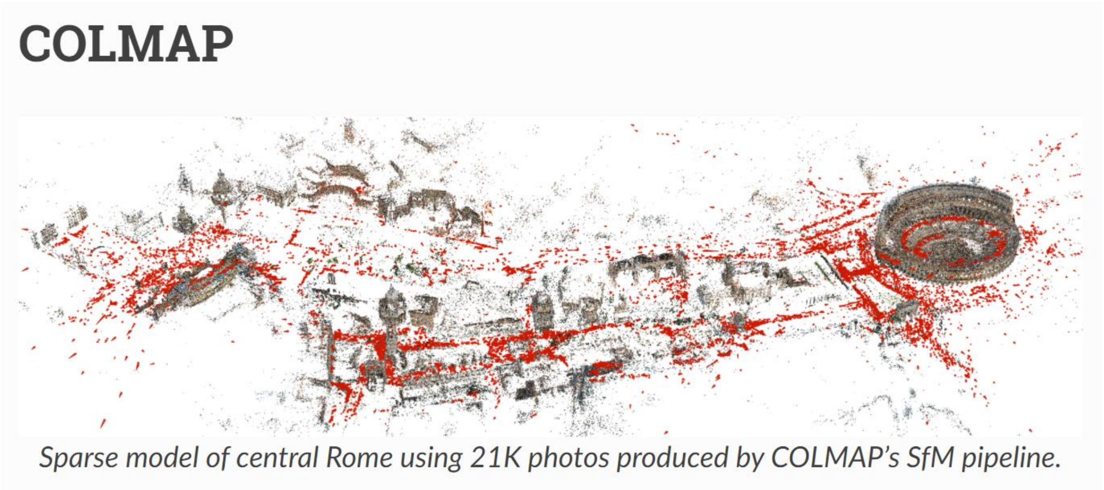

Schönberger et al. COLMAP (2016)

Modern multi-view geometry software heavily rely on feature points.

What is missing from the 3D point map?

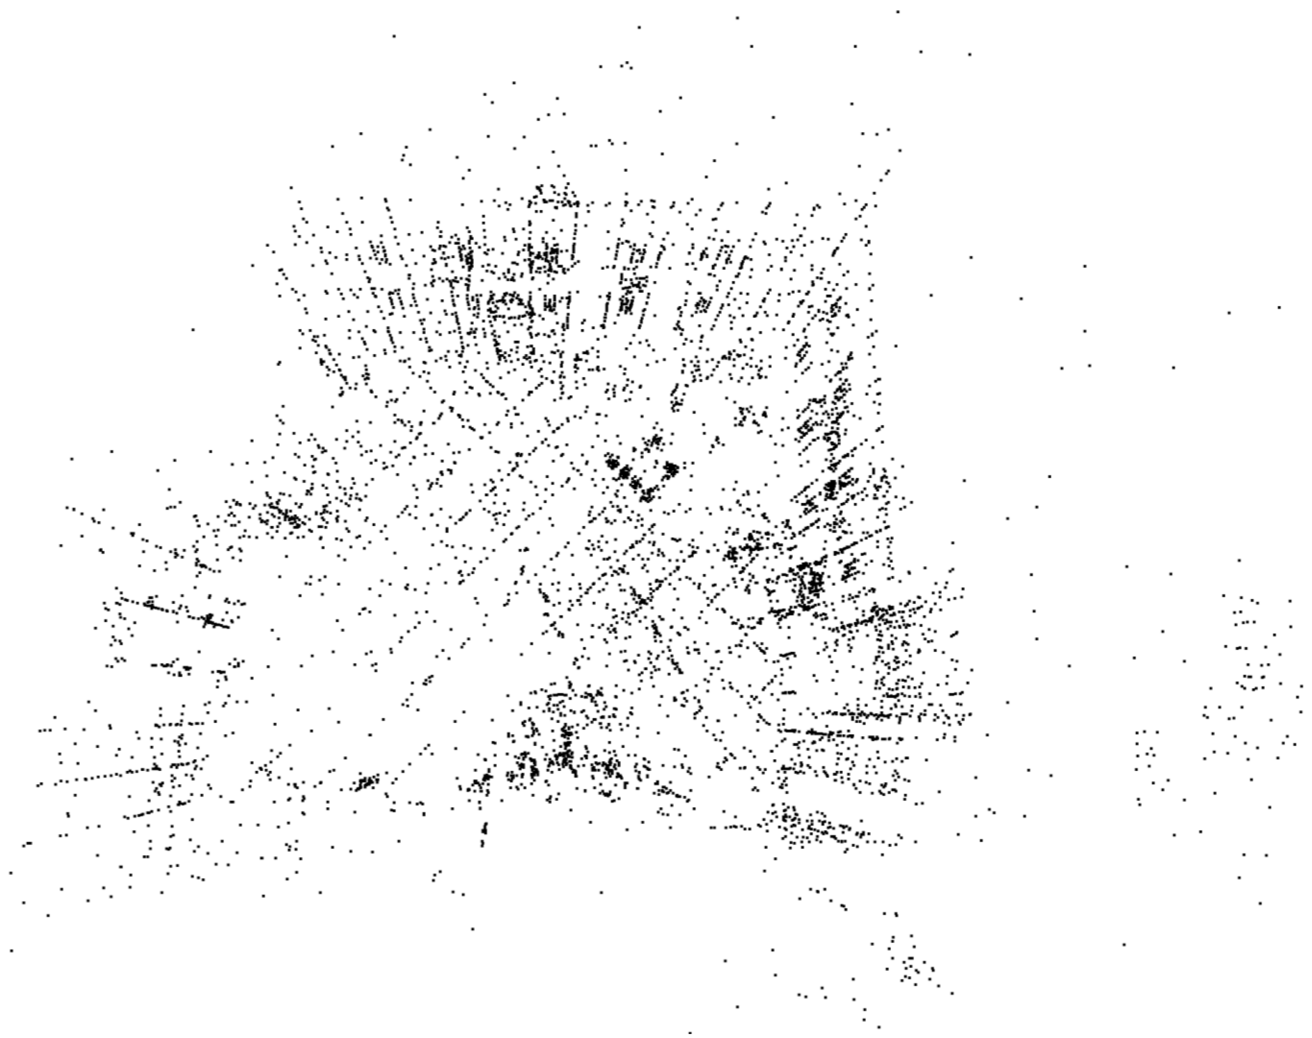

SuperPoint + COLMAP point triangulator

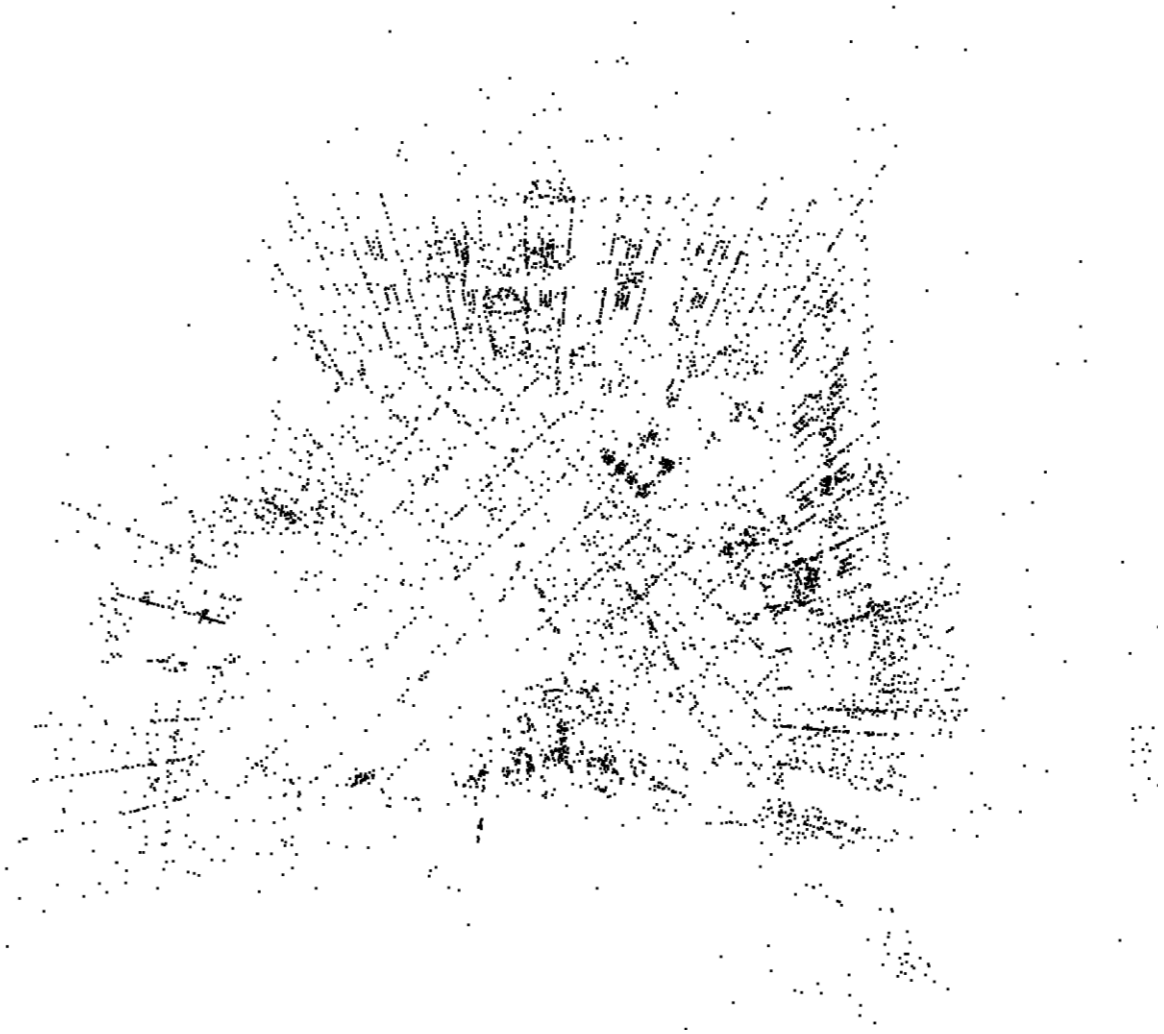

SuperPoint + COLMAP point triangulator

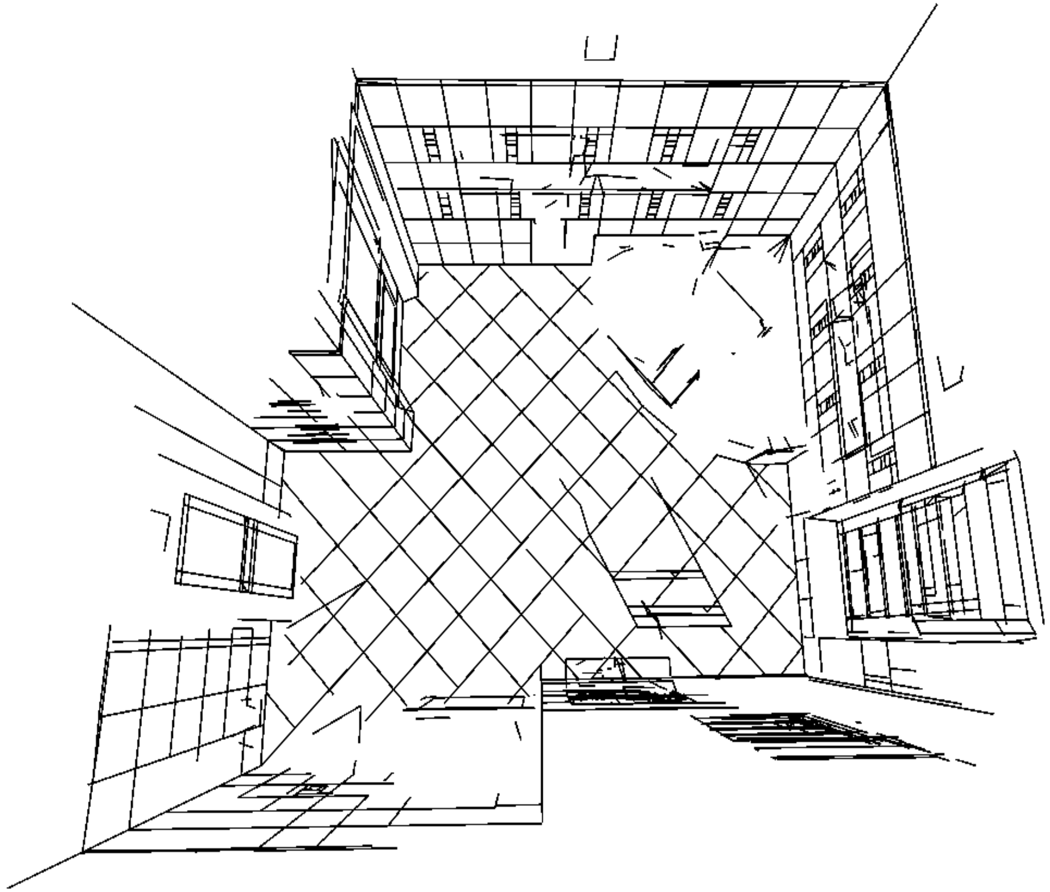

Line Mapping

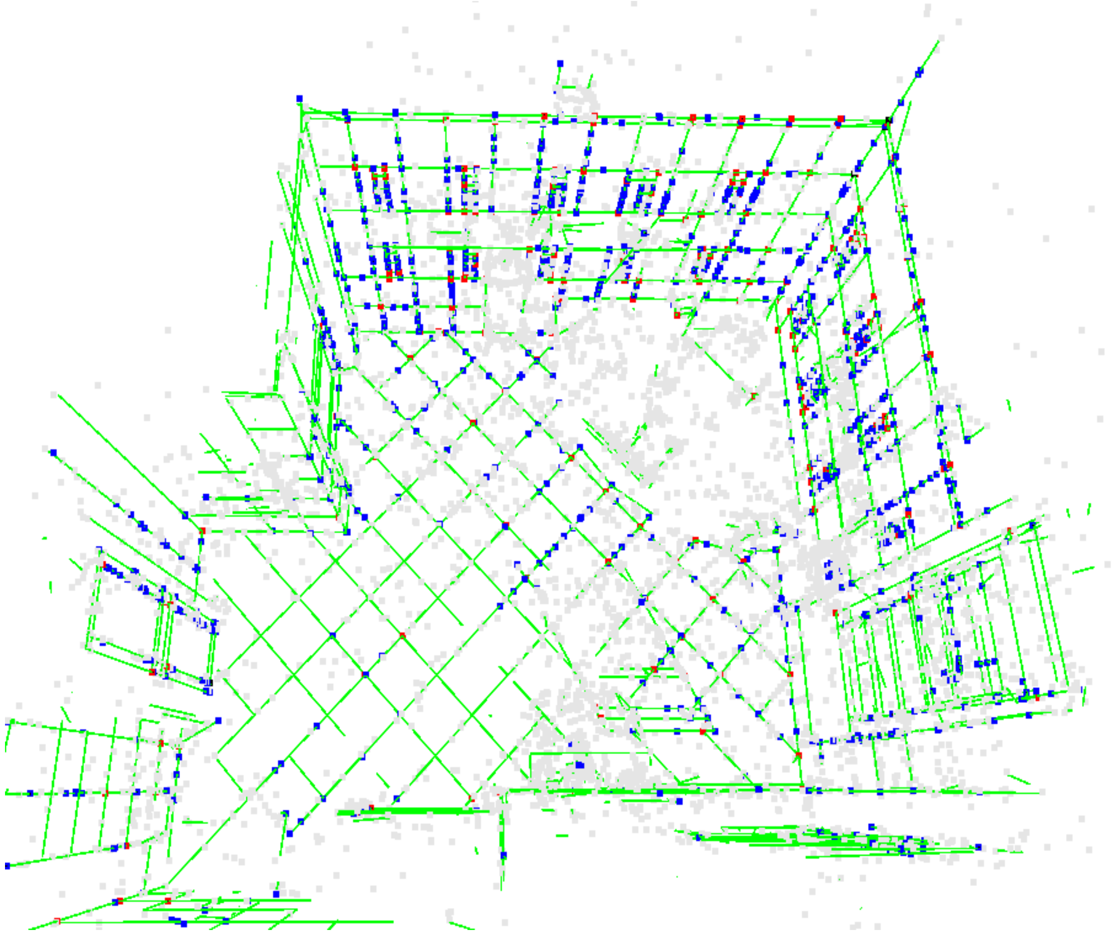

Point-line Association

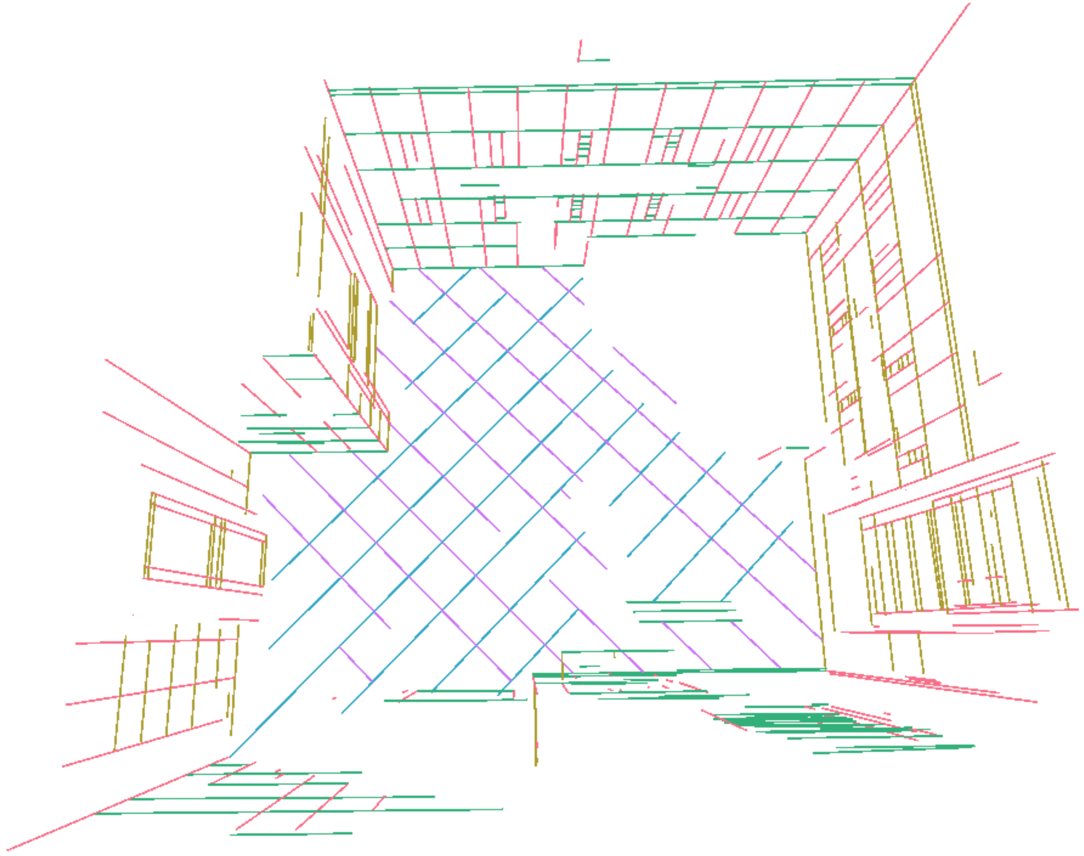

Parallelism & Orthogonality

Von Gioi et al. LSD: A fast line segment detector with a false detection control (2010)

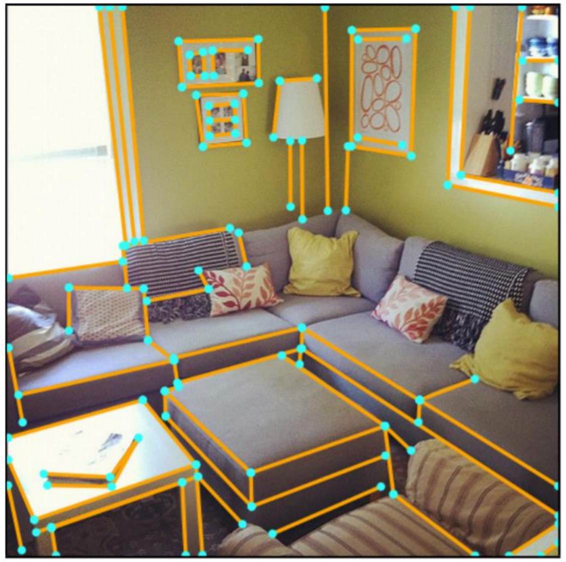

Zhou et al. LCNN (2019)

Pautrat et al. SOLD2 (2021)

Challenges on mapping lines

- Inconsistent endpoints
- Line Fragmentation
- No Two-view Geometric Verification
- Weak matchers
- Degenerate Configurations
- …………..

Challenges on mapping lines

- Inconsistent endpoints
- Line Fragmentation
- No Two-view Geometric Verification
- Weak matchers
- **Degenerate Configurations**
- …………..

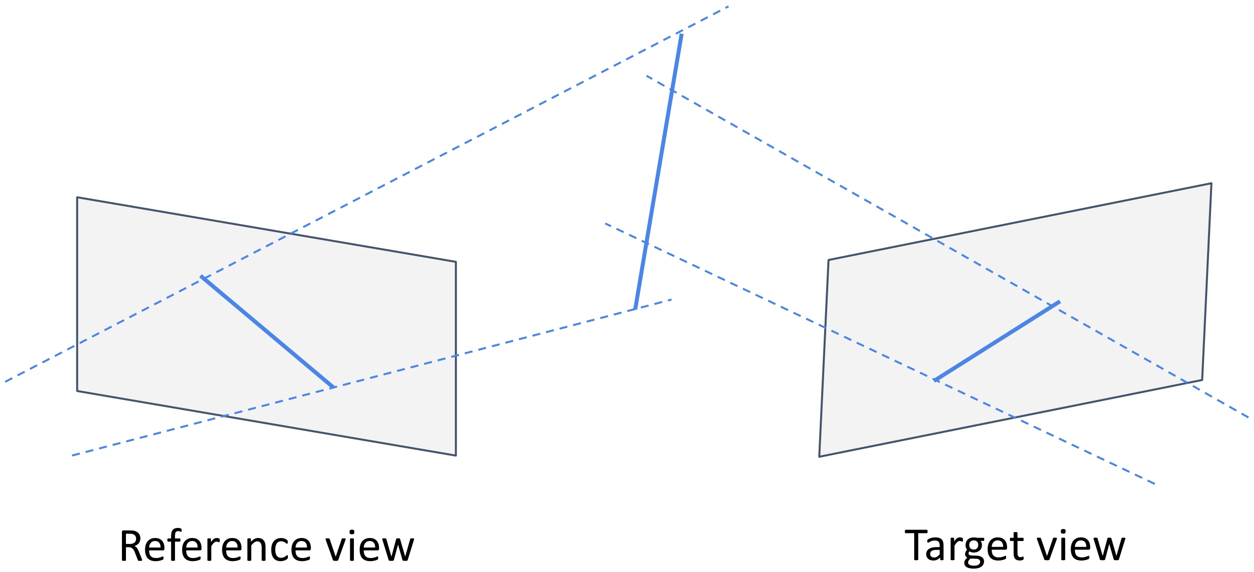

Algebraic Line triangulation

Ray-plane intersection on both endpoints respectively

Degeneracy happens when ray lies on the plane!
0 / 1 / 2 degenerate endpoints

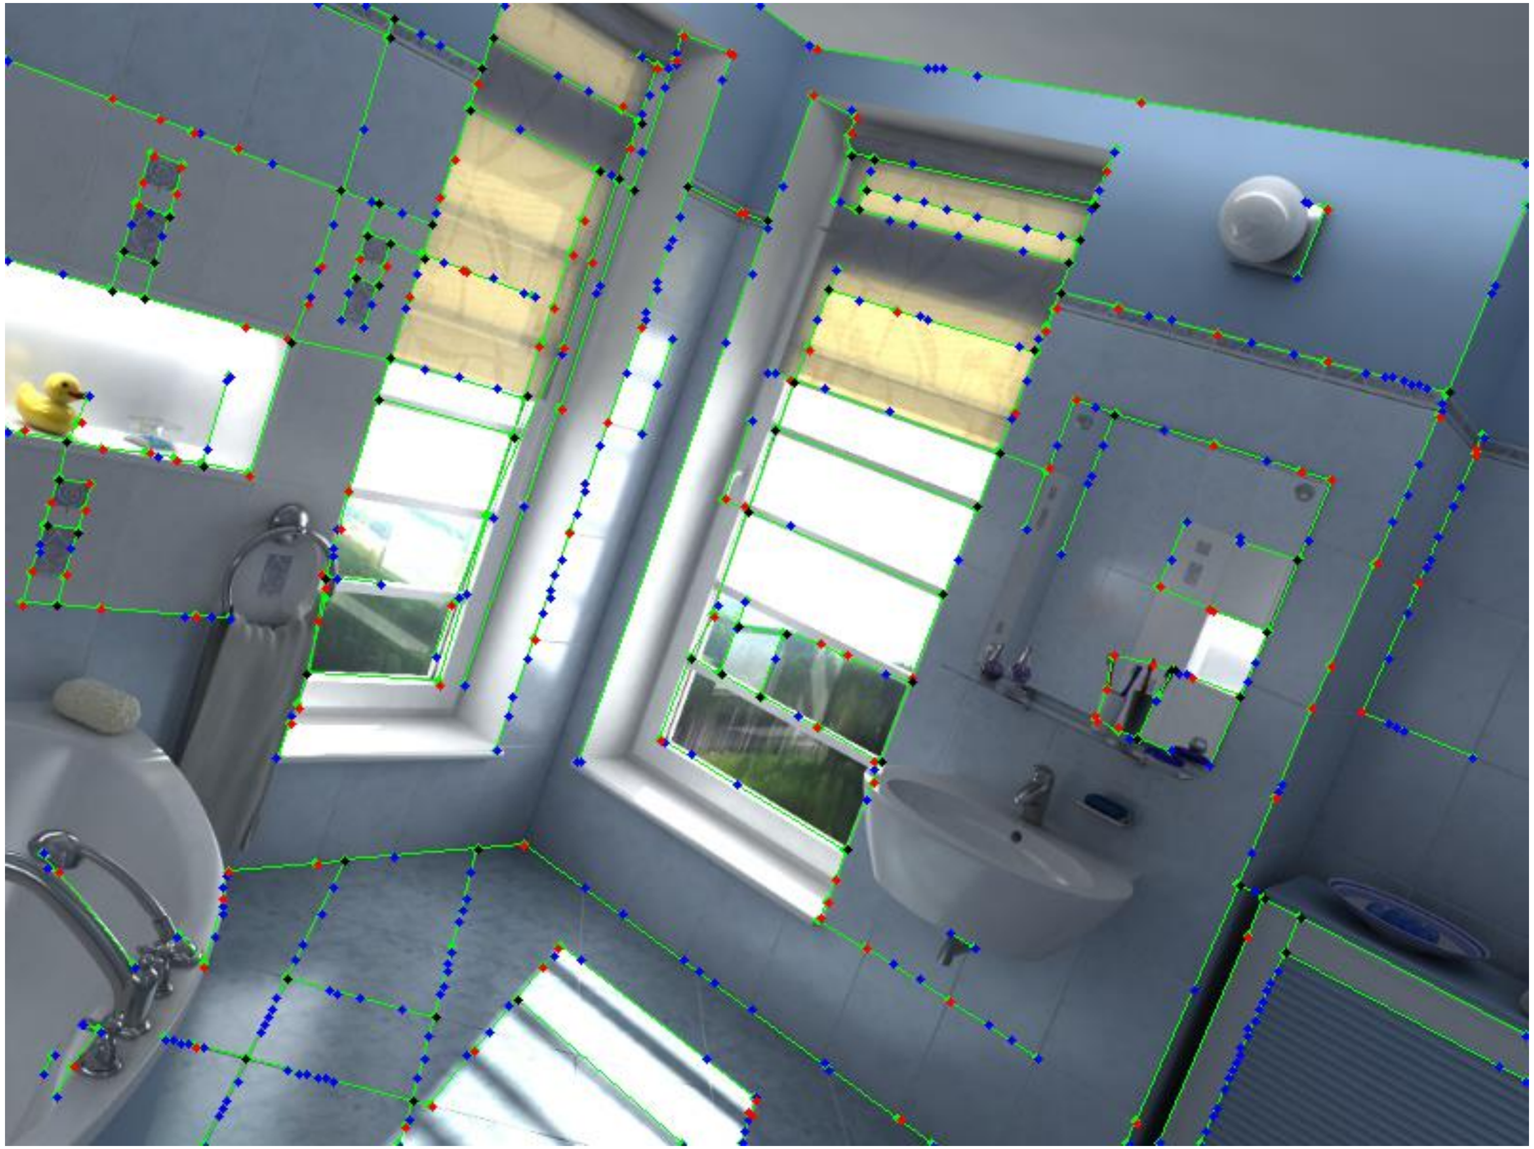

Blue – degree 1 red – degree 2 black degree >= 3

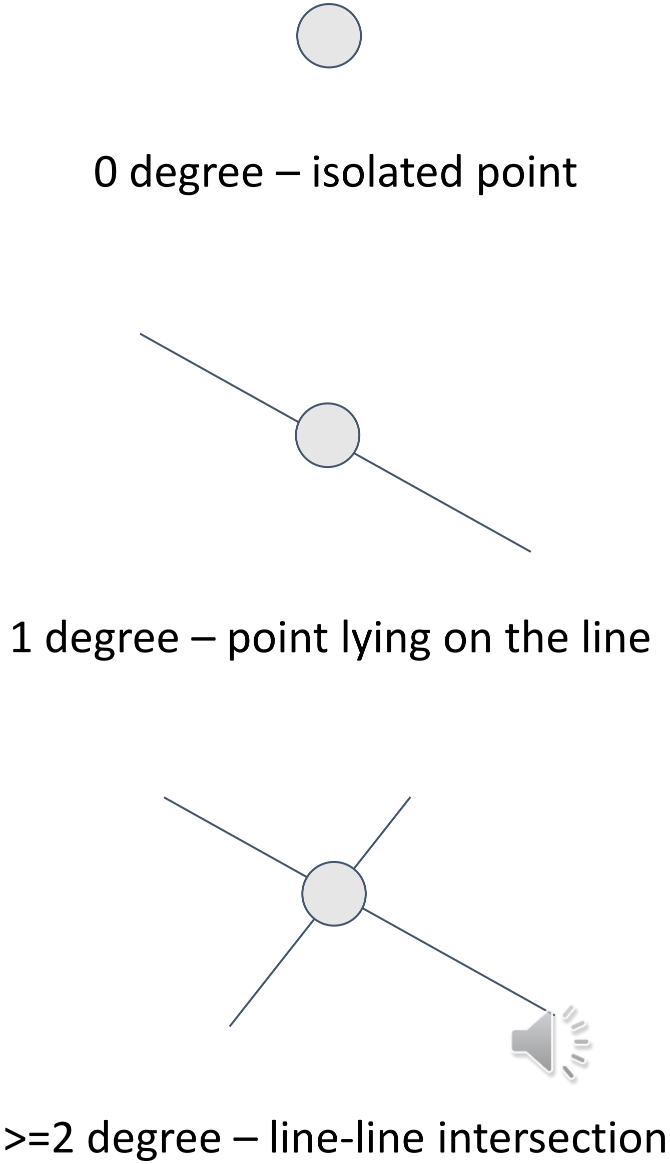

A brief overview of our mapping pipeline

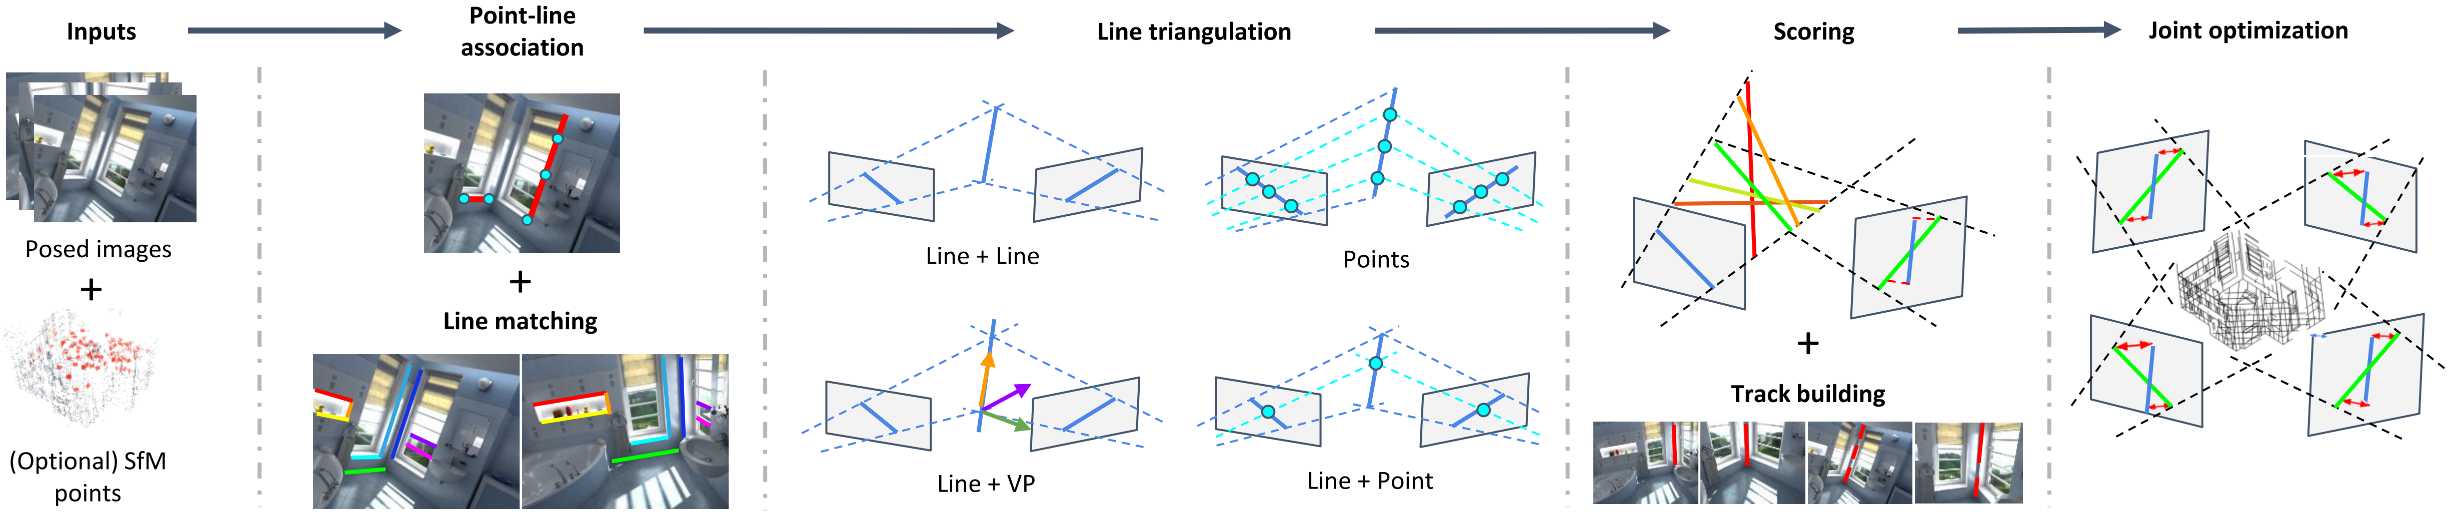

Also easily extends with available depth maps if applicable.

Triangulating proposals

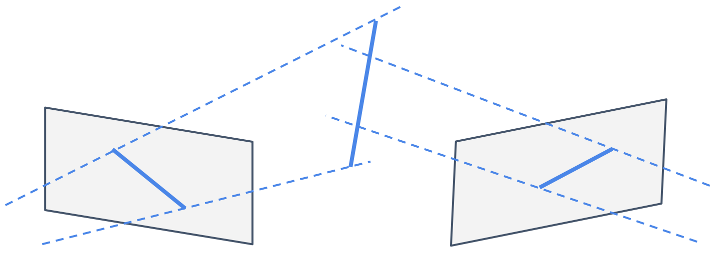

Line + Line

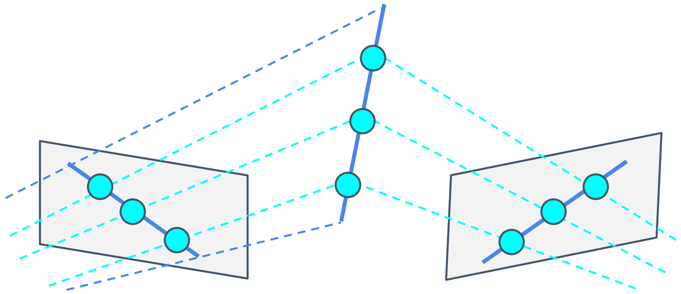

Points

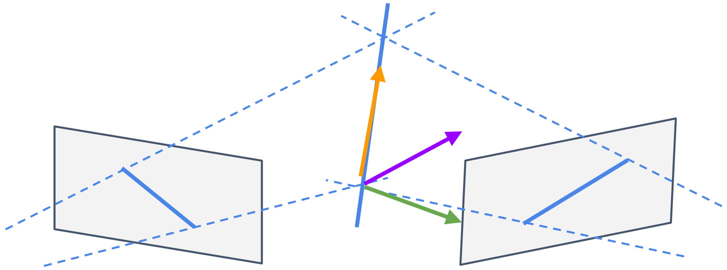

Line + VP

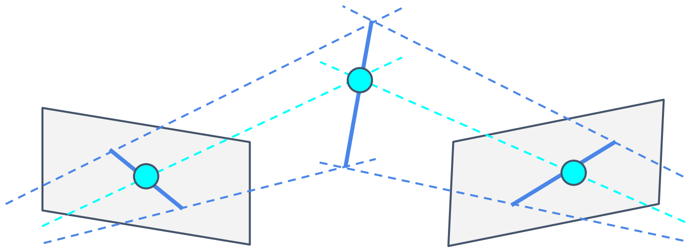

Line + Point

Using shared neighboring 3D points to help avoid degeneracy!

Scoring each proposal & Building tracks

2D + 3D metric

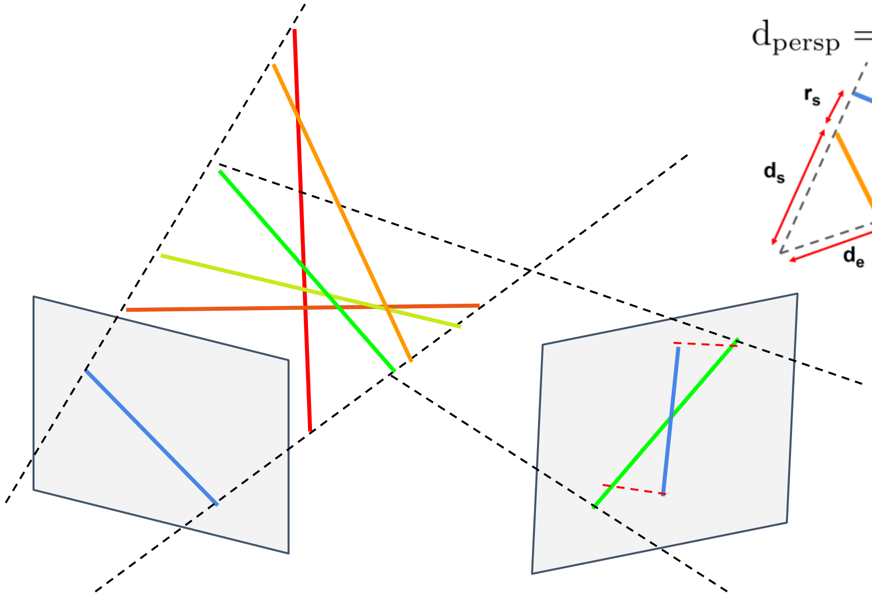

$$\mathrm{d}_{\mathrm{persp}} = \max\left(\frac{\mathbf{r_s}}{\mathbf{d_s}}, \frac{\mathbf{r_e}}{\mathbf{d_e}}\right)$$

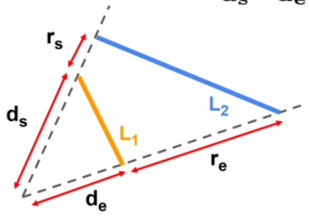

$$\mathrm{S}_{\mathrm{overlap}} = \mathbb{1}_{\frac{|\mathbf{\Pi(L_1)}|}{|\mathbf{L_2}|} \geq \tau_o}$$

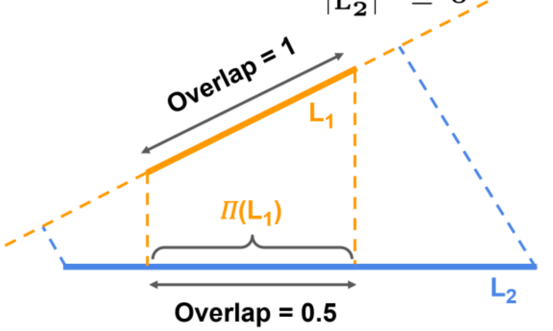

$$\mathrm{d}_{\mathrm{inner}} = \frac{\mathbf{f}\max(\mathbf{r_s},\mathbf{r_e})}{\min(\mathbf{d_1},\mathbf{d_2})}$$

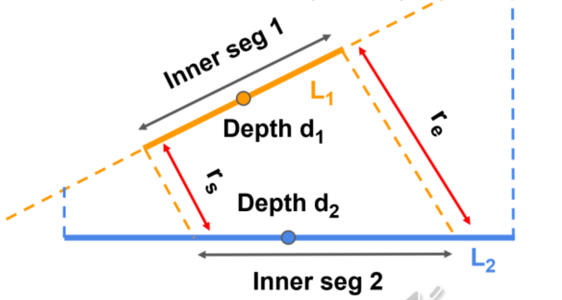

Joint Optimization over Points, Lines and VPs

- Reprojection error
- Point-line associations
- Line-line associations via VP (construct VP tracks in advance)

Weighted by analyzing connections from 2D relational graphs inside the track

Plücker coordinates!

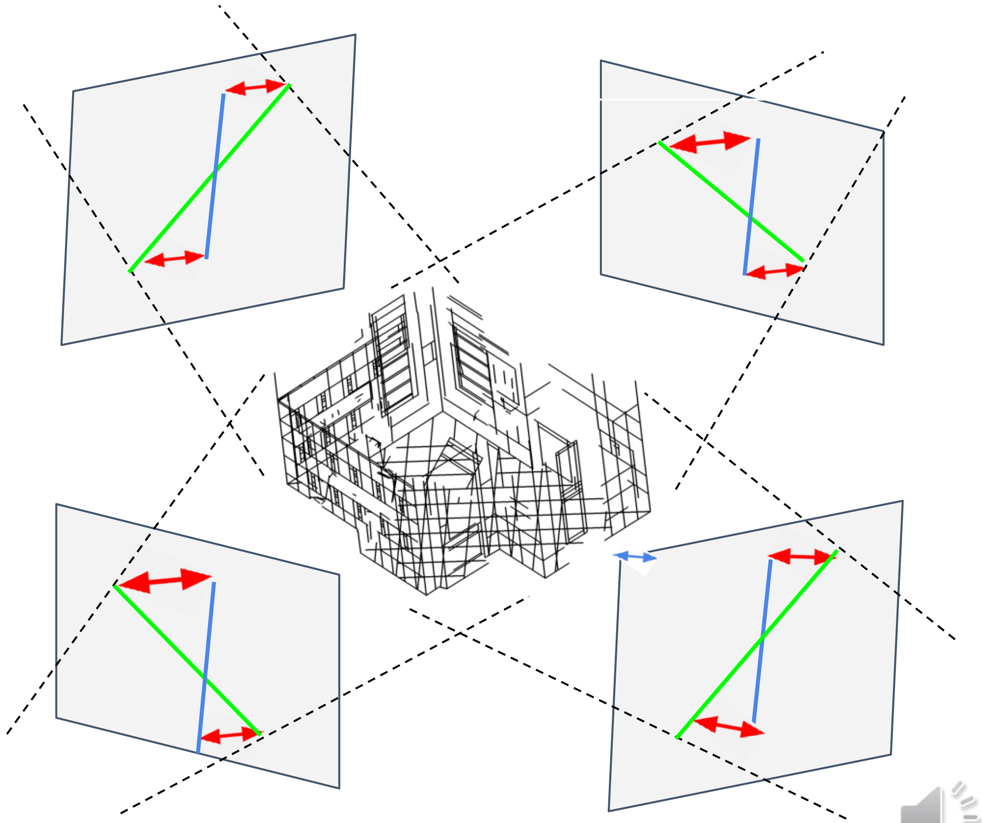

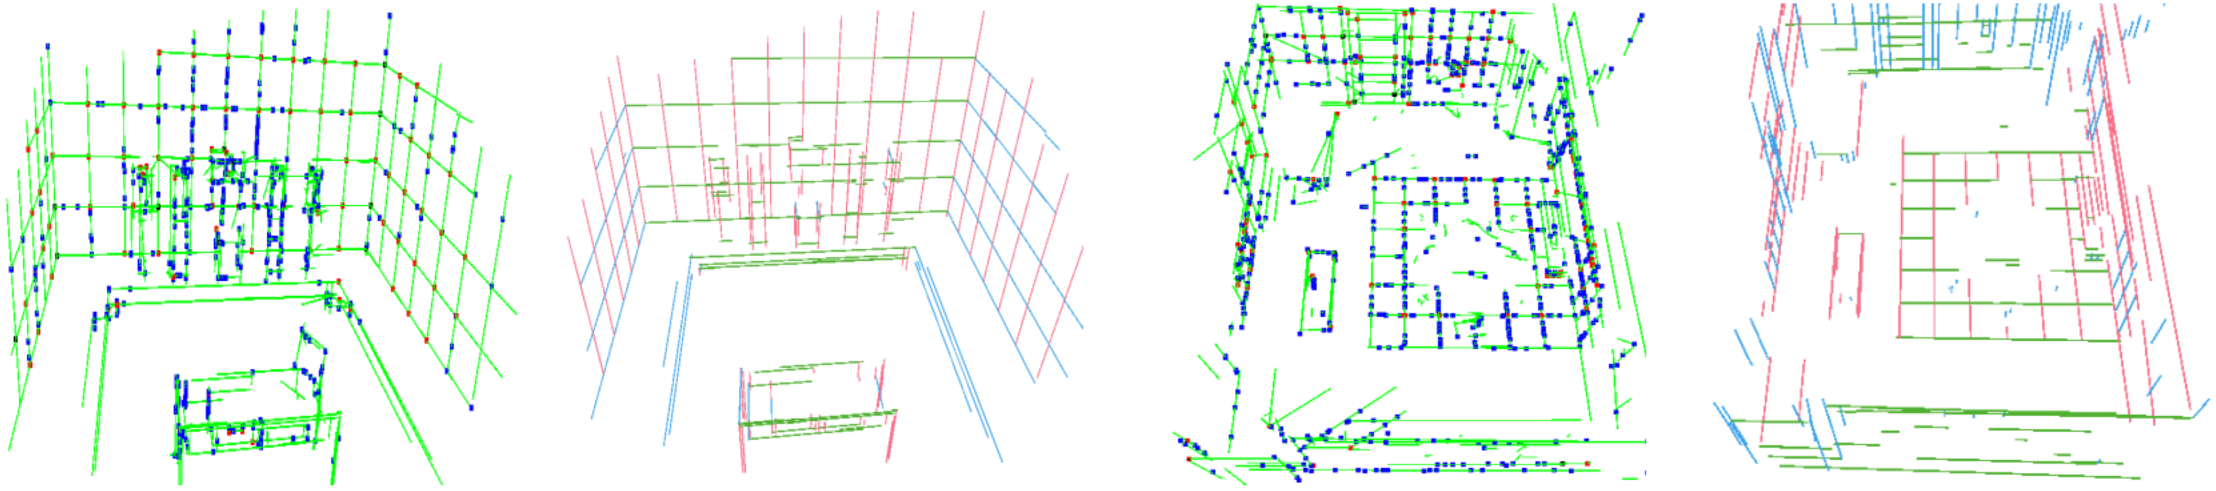

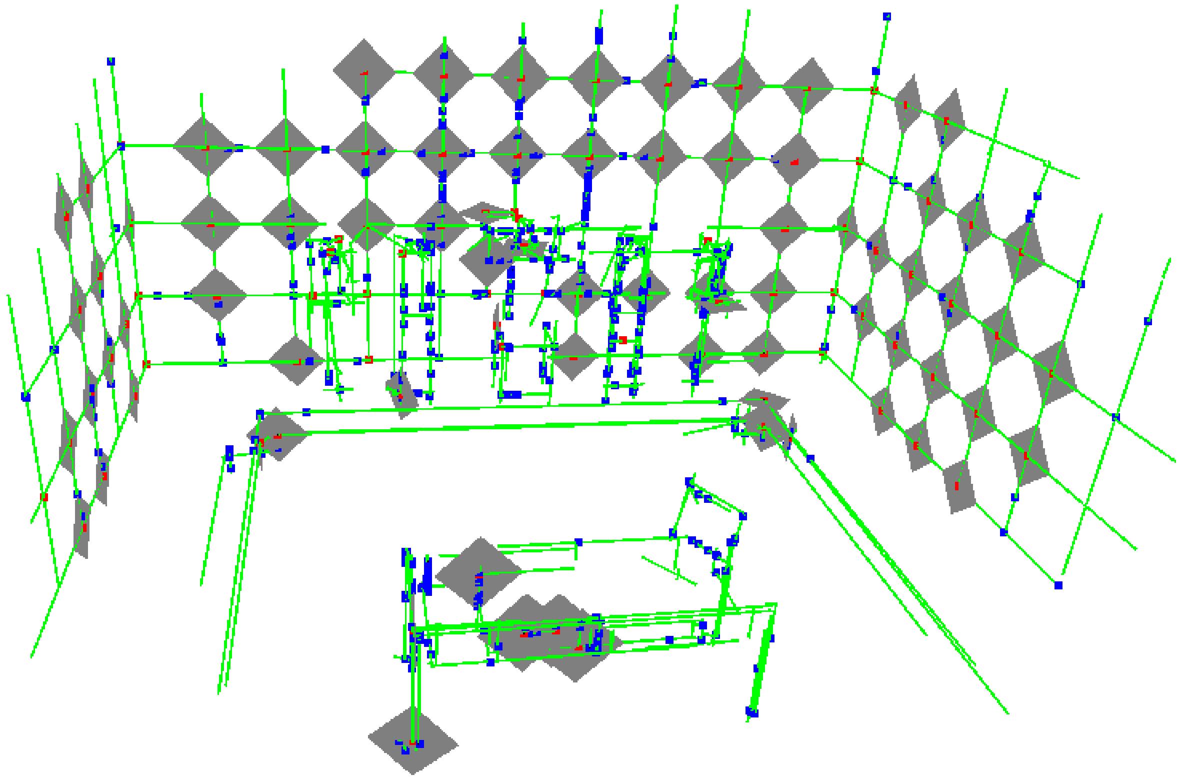

Local planes from degree-2 junctions

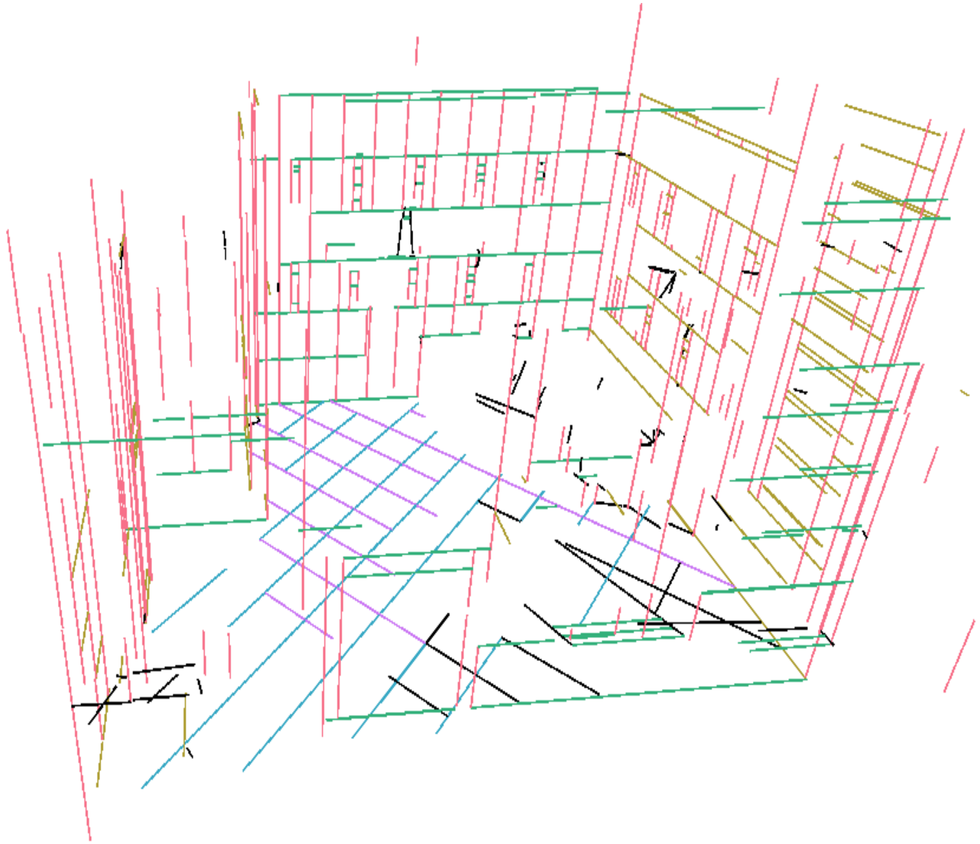

Iteratively optimize and merge VP tracks

[LOG] Orthogonal pair detected: 0 / 1, angle = 89.93
[LOG] Orthogonal pair detected: 0 / 2, angle = 89.94
[LOG] Orthogonal pair detected: 0 / 3, angle = 89.97
[LOG] Orthogonal pair detected: 0 / 4, angle = 89.90
[LOG] Orthogonal pair detected: 1 / 2, angle = 89.88
[LOG] Orthogonal pair detected: 3 / 4, angle = 90.00

-> Atlanta world with two Manhattan axis!

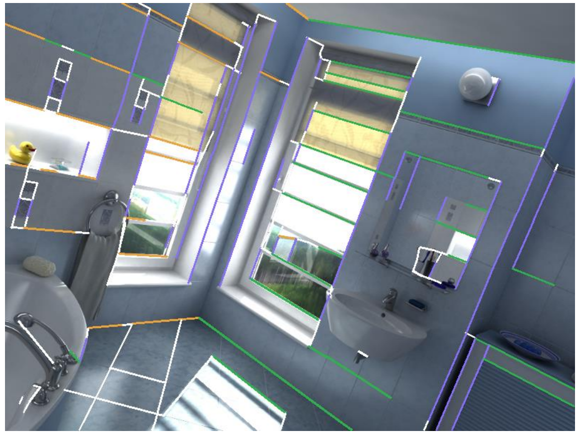

VP detection

VP track visualization (vertical)

Input

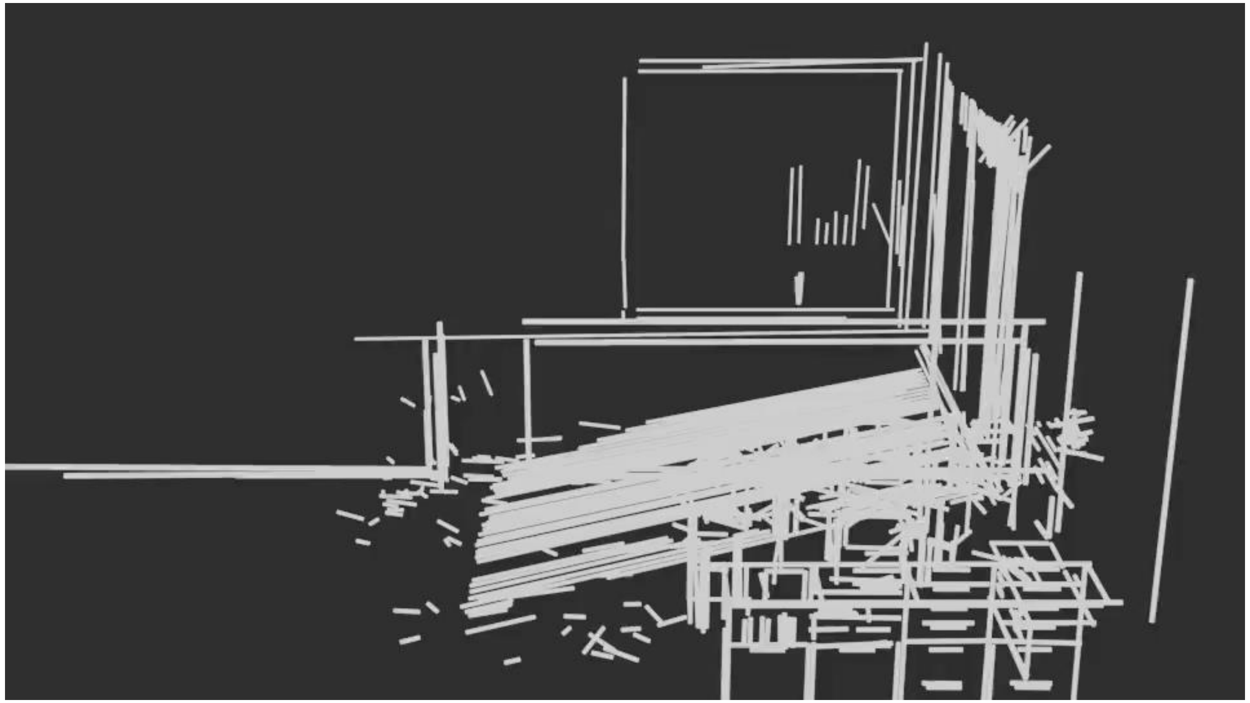

Ours

Input

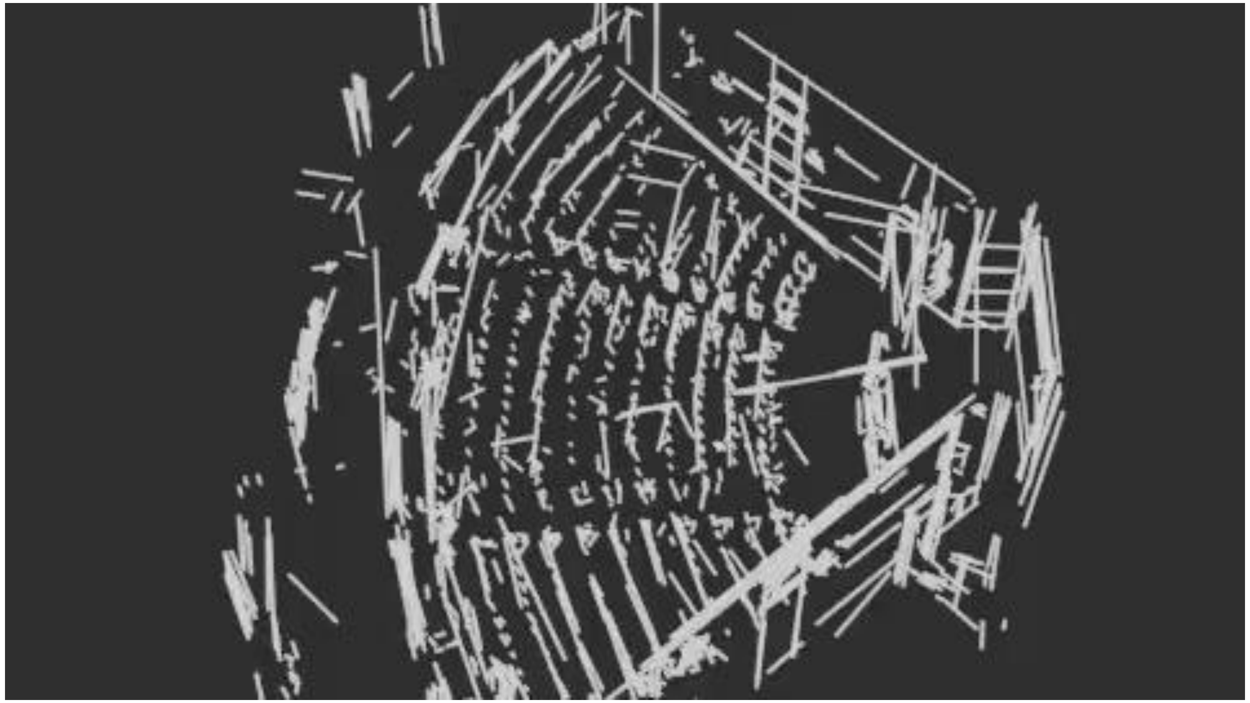

Ours

Input

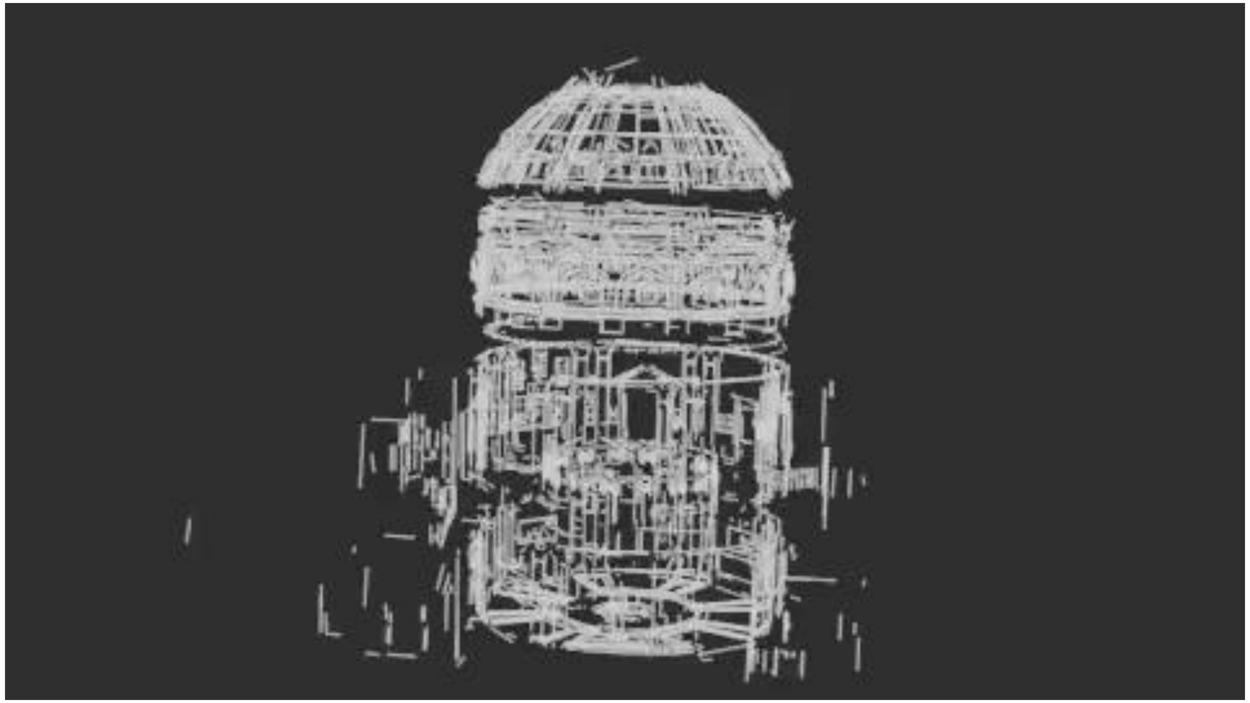

Ours

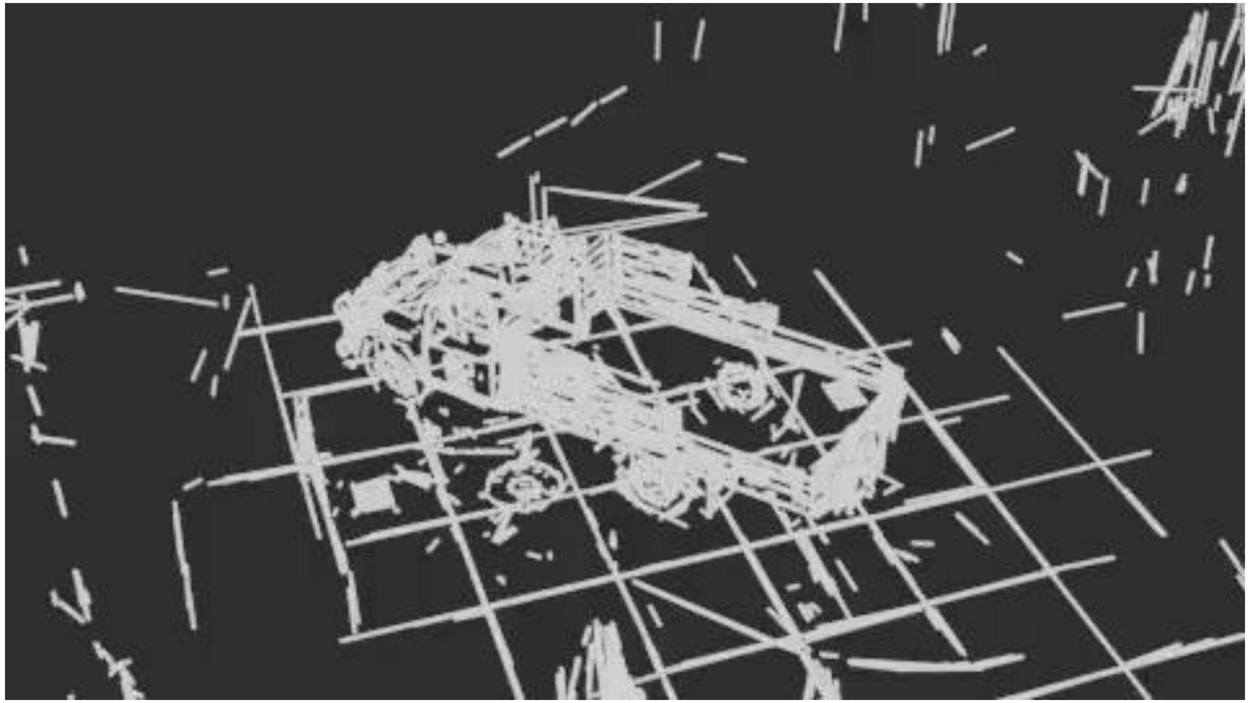

Input

Ours

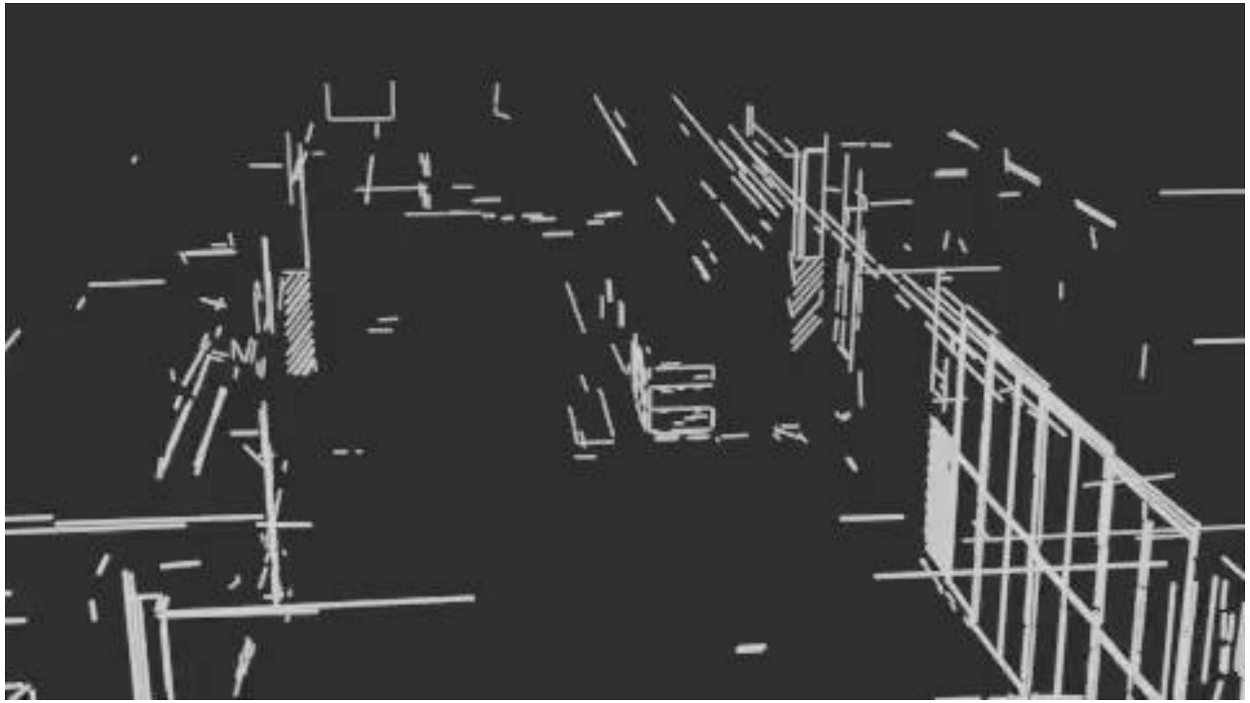

Input

Ours

Input

Line3D++

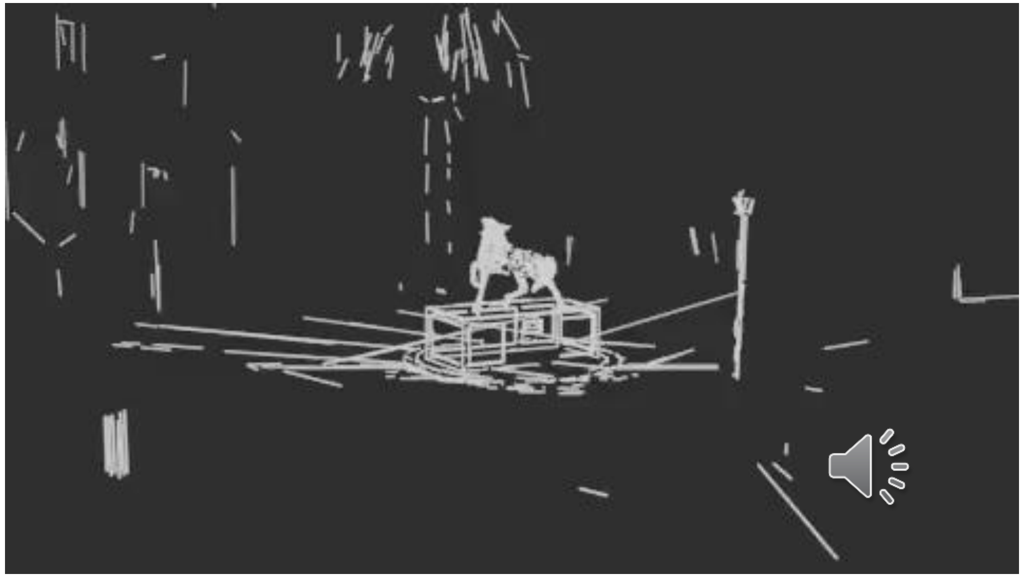

Ours

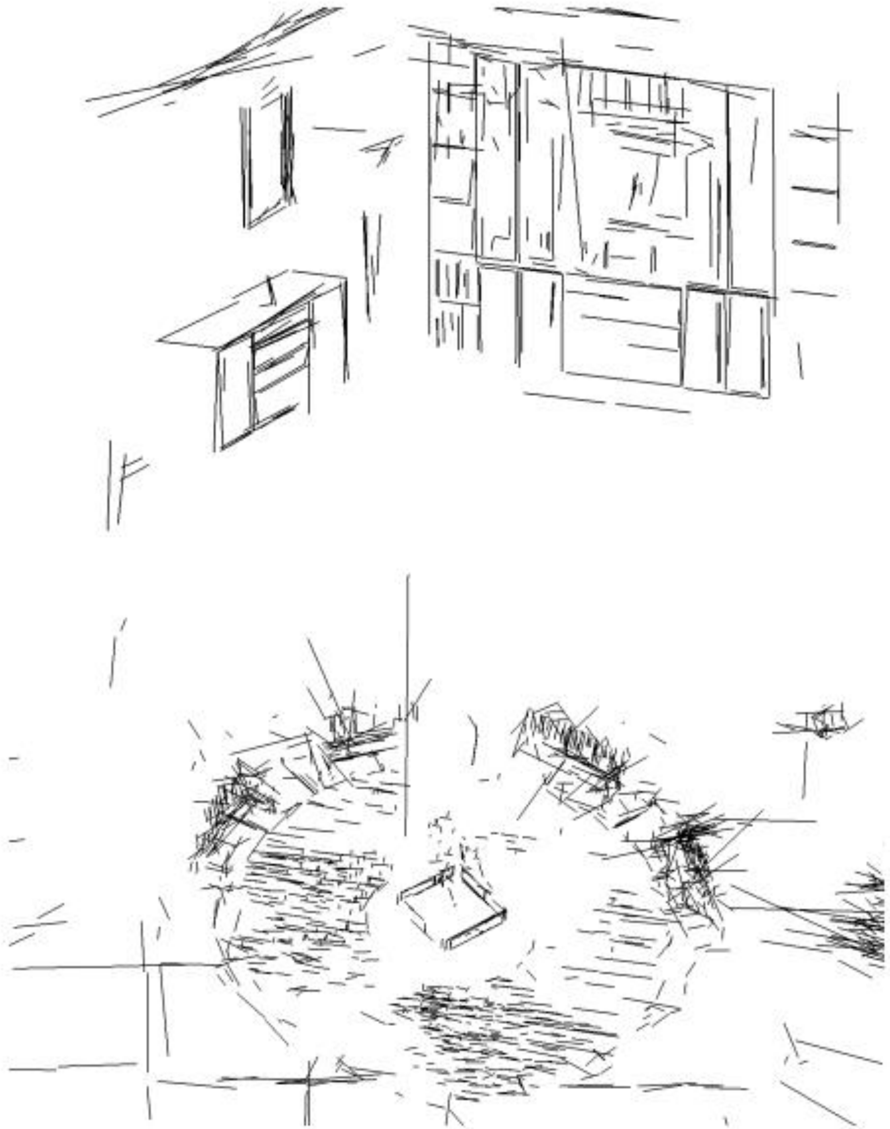

ELSR (CVPR 2022)

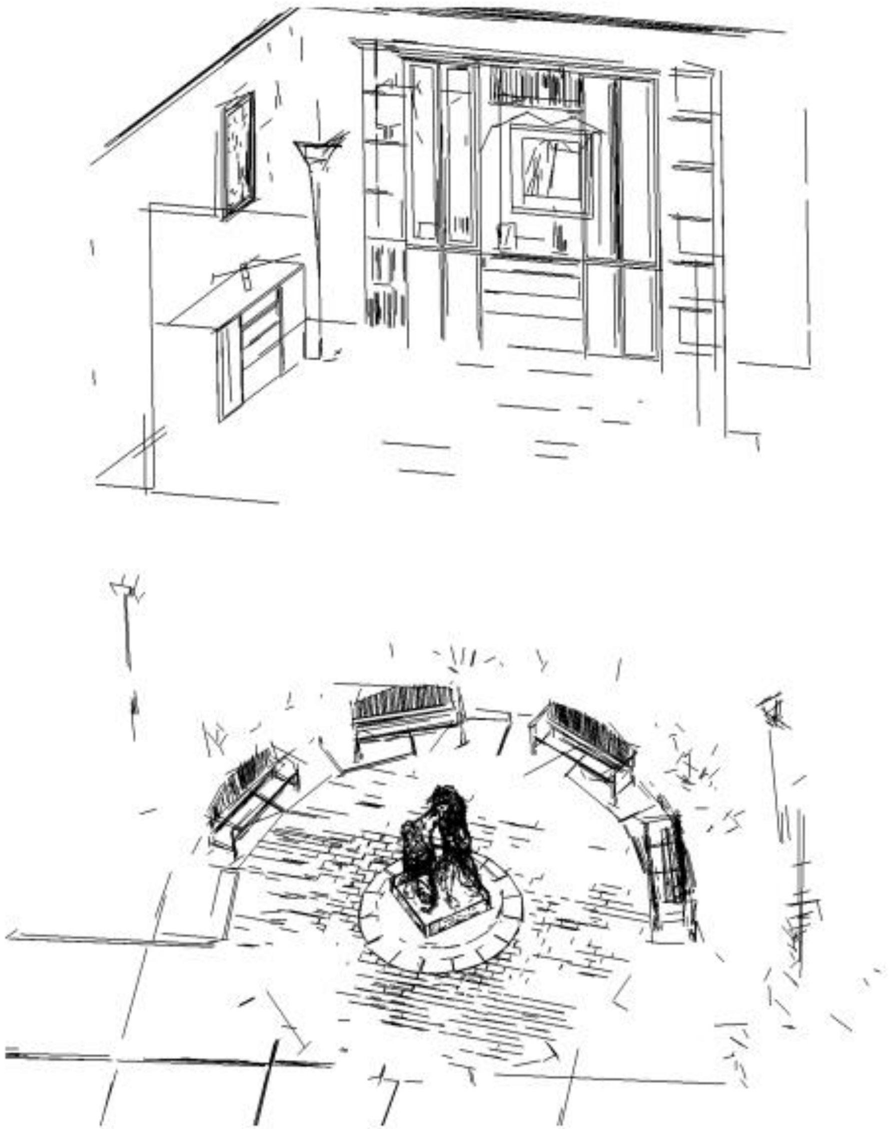

Ours

We do have quantitative evaluation

- Length Recall at certain threshold
- Precision at certain threshold
- Average number of image support / average number of line support

Line type	Method	R1	R5	R10	P1	P5	P10	# supports
LSD	L3D++	37.0	153.1	218.8	53.1	80.8	**90.6**	(14.8 / 16.8)
	ELSR	13.9	59.7	96.5	55.4	72.6	82.2	(N/A / N/A)
	Ours	**48.6**	**185.2**	**251.3**	**60.1**	**82.4**	90.0	(**16.4** / **20.5**)
SOLD2	L3D++	36.9	107.5	132.8	67.2	**86.8**	**93.2**	(13.2 / 20.4)
	Ours	**54.3**	**151.1**	**191.2**	**69.8**	84.6	90.0	(**16.5** / **38.7**)

Method	R5	R10	R50	P5	P10	P50	# supports
L3D++	373.7	831.6	2783.6	40.6	54.5	85.9	(8.8 / 9.3)
ELSR	139.2	322.5	1308.0	38.5	48.0	74.5	(N/A / N/A)
Ours (line-only)	472.1	1058.8	3720.7	**46.8**	**58.4**	**86.1**	(10.3 / 11.8)
Ours	**508.3**	**1154.5**	**4179.5**	46.0	56.9	83.7	(**10.4** / **12.0**)

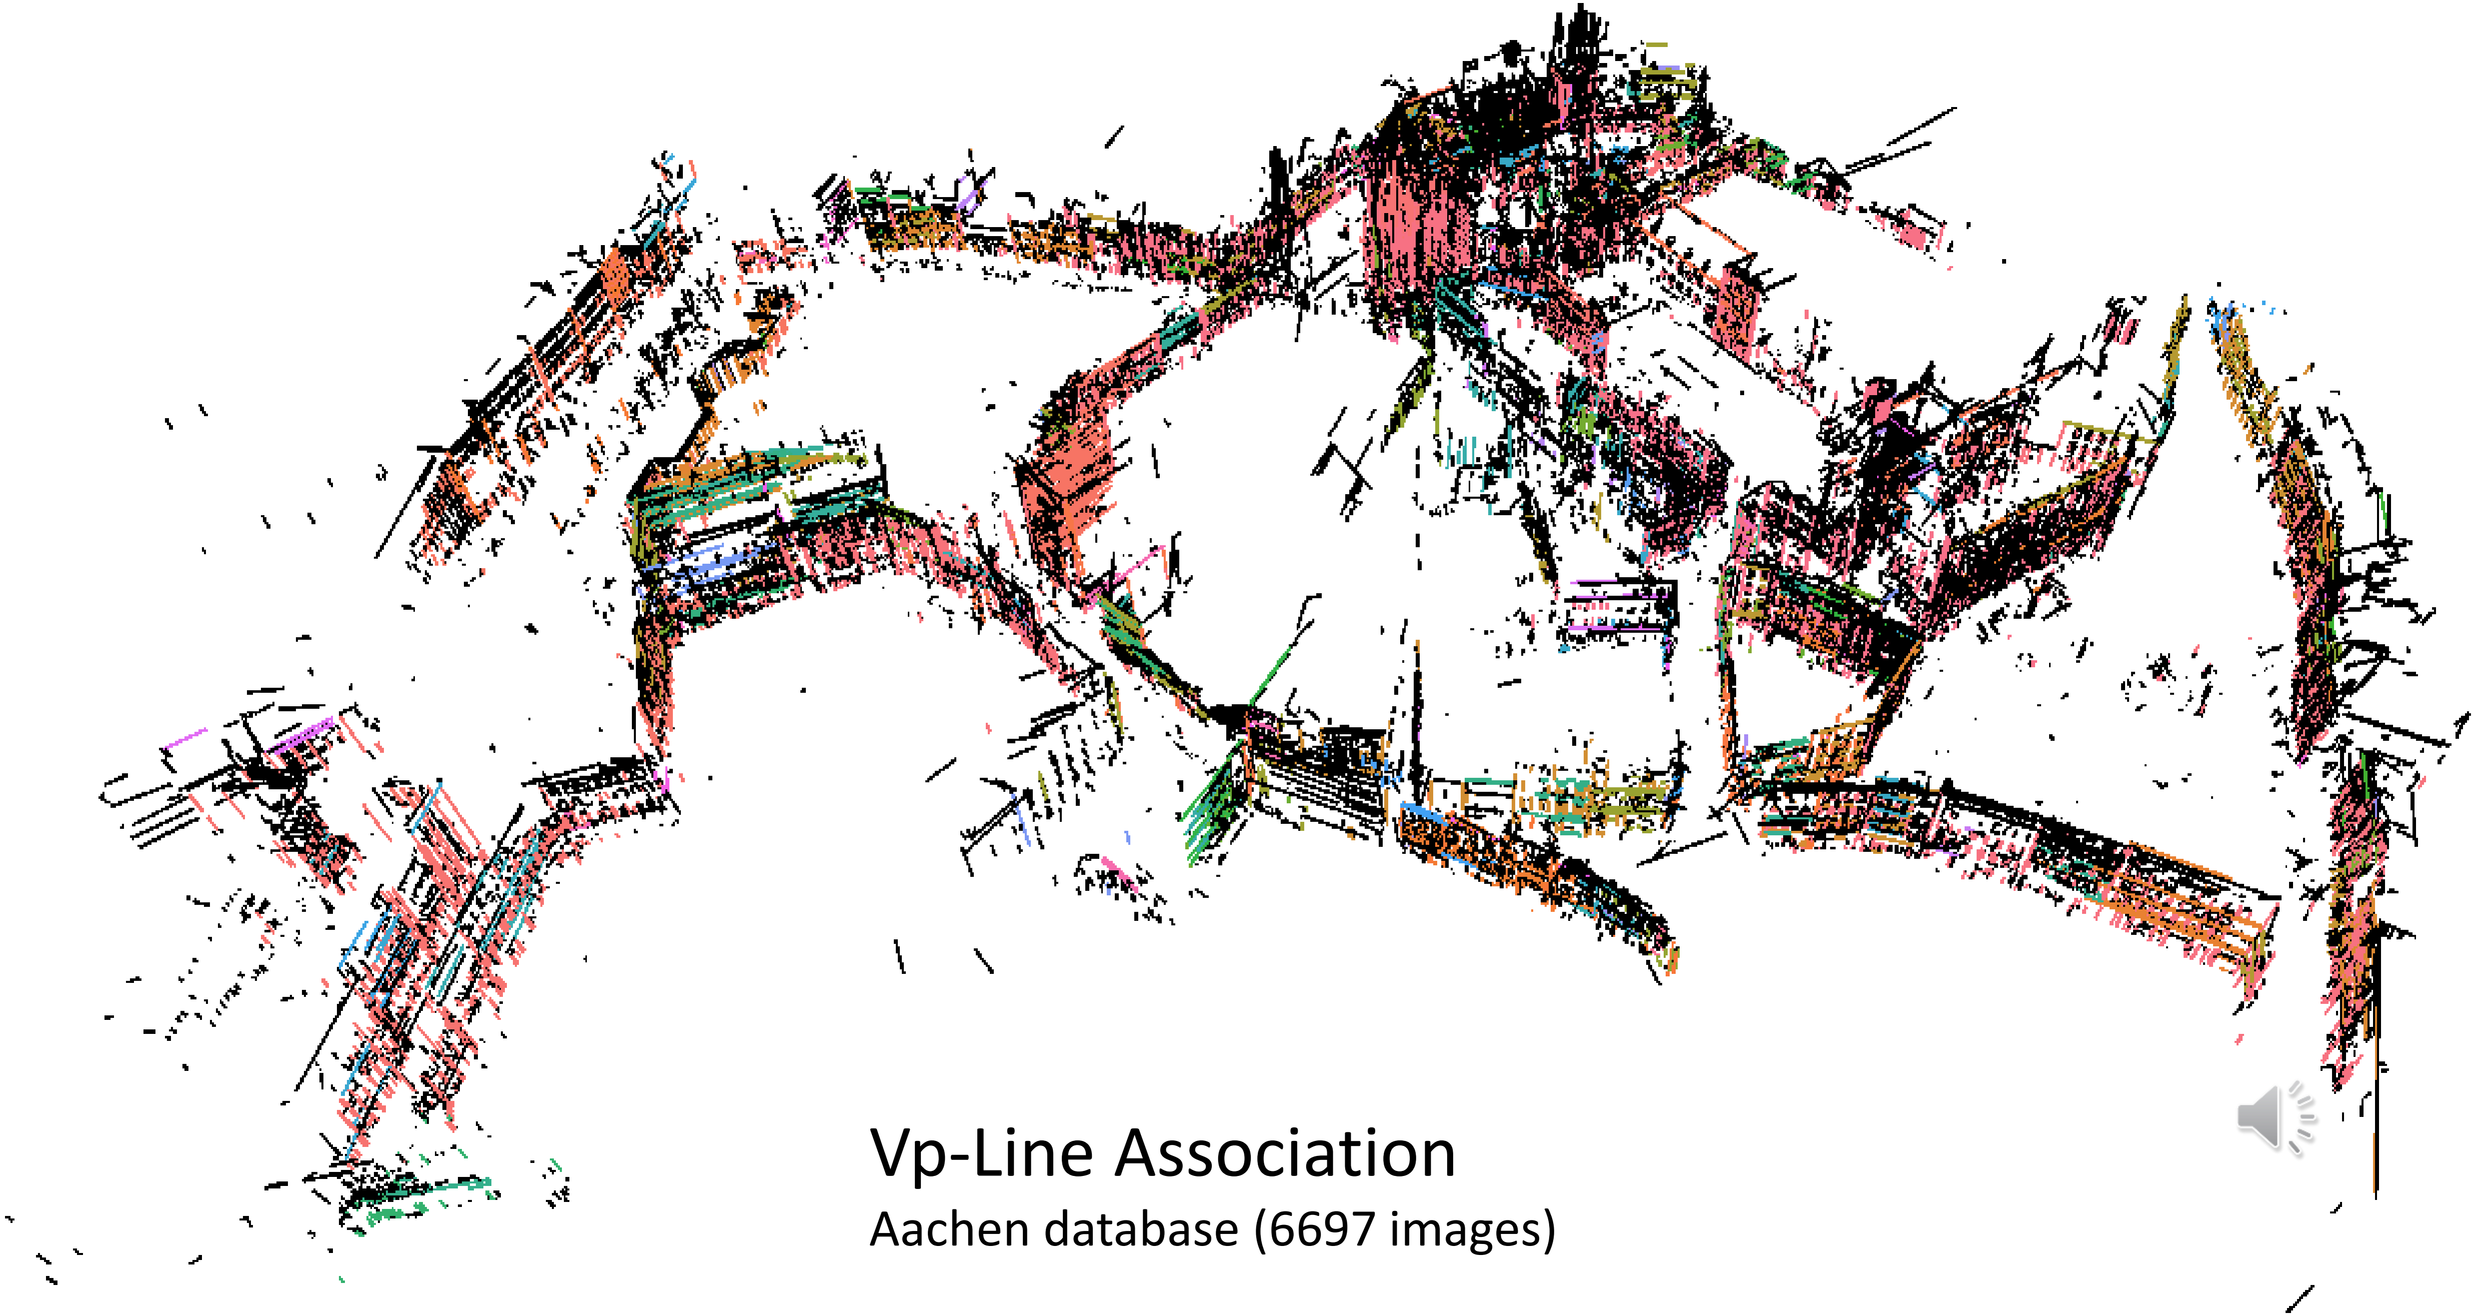

Vp-Line Association

Aachen database (6697 images)

Old saying: Every ~~road~~ line leads to Rome

Scalable to Rome 16k

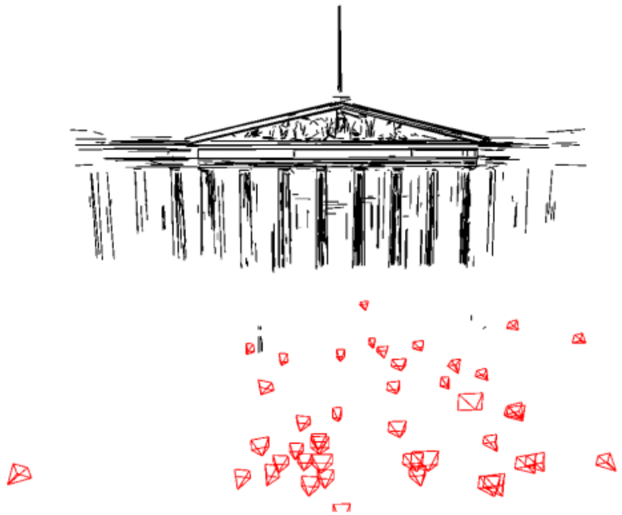

British Museum from [47]

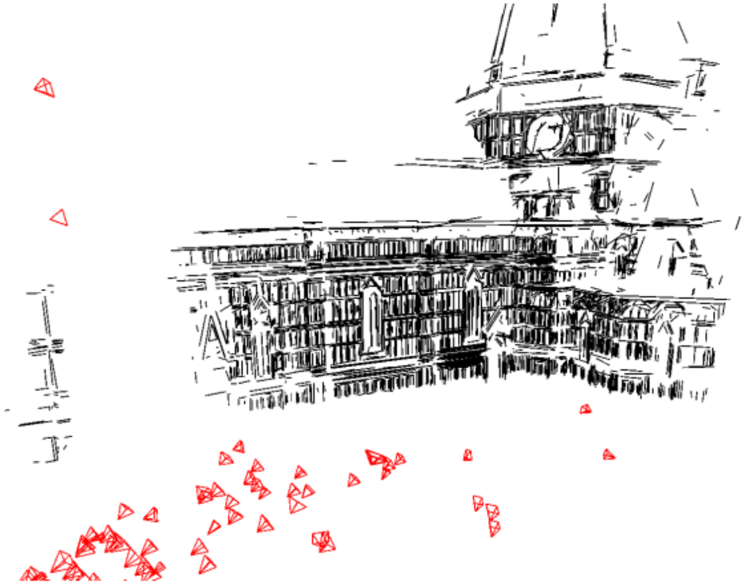

Florence Cathedral Side from [47]

London Bridge from [47]

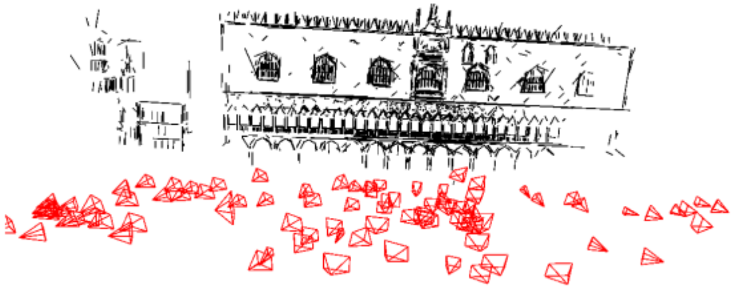

Piazza San Marco from [47]

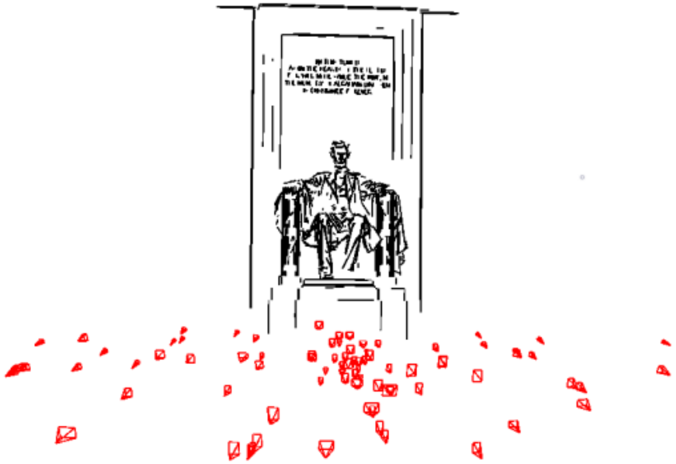

Lincoln Memorial Statue from [47]

St. Paul's Cathedral from [47]

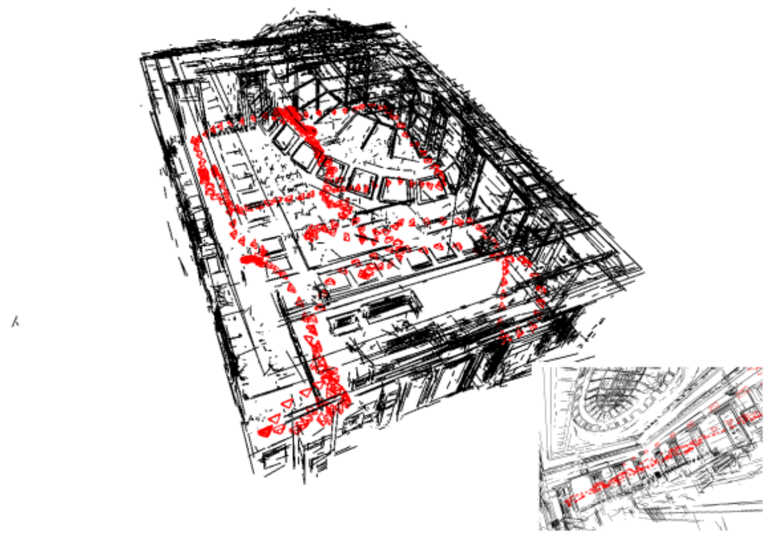

Courtroom (indoor and outdoor) from [26]

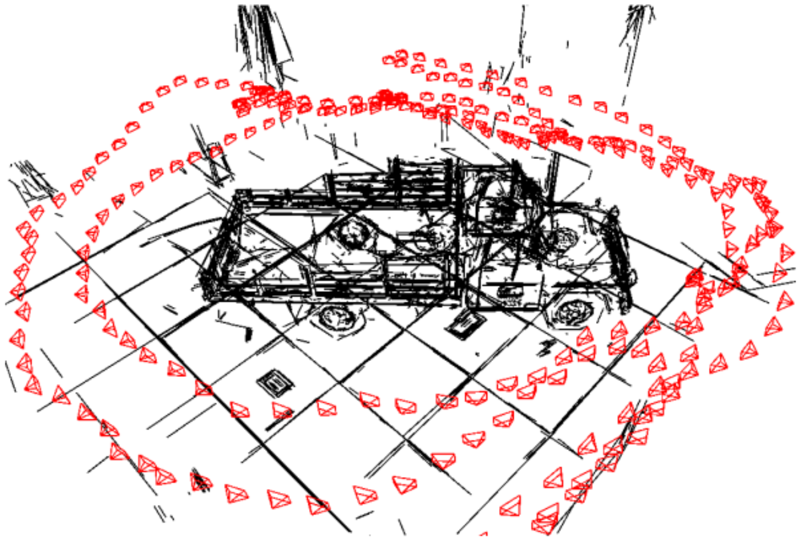

Truck from [26]

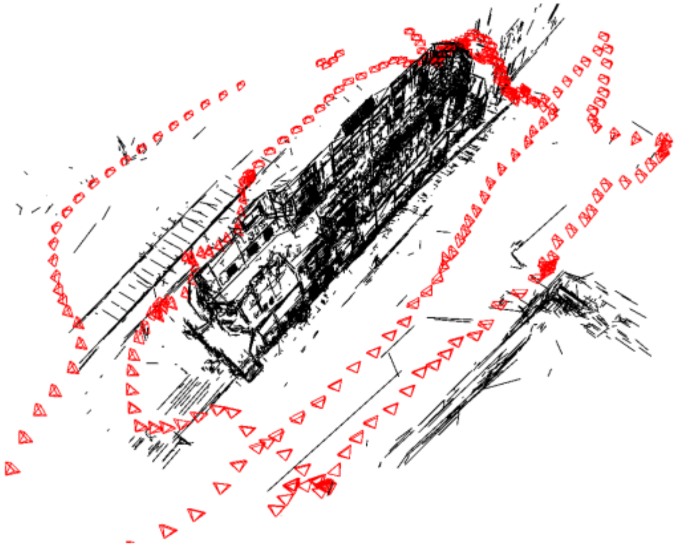

Train from [26]

Lighthouse from [26]

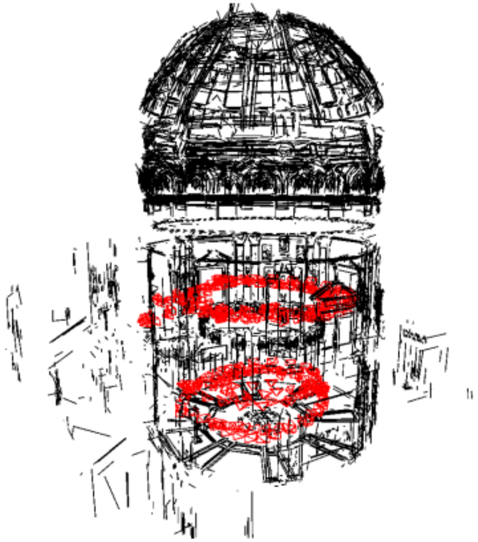

Museum from [26]

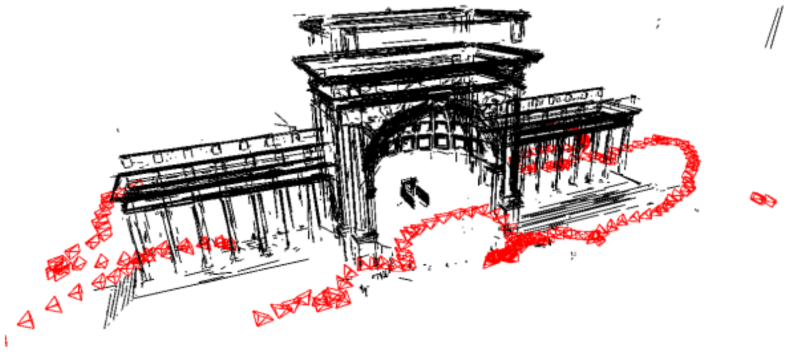

Temple from [26]

Localization

Before optimization

After optimization

Hybrid localization with points and lines

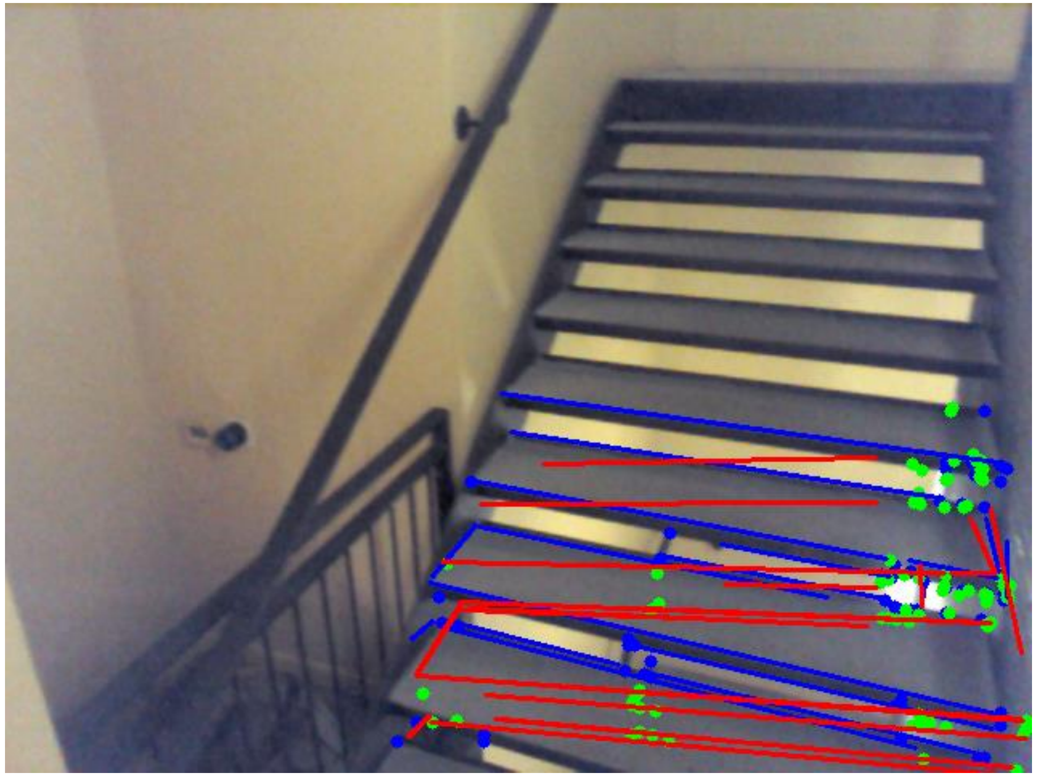

Point-alone localization with HLoc

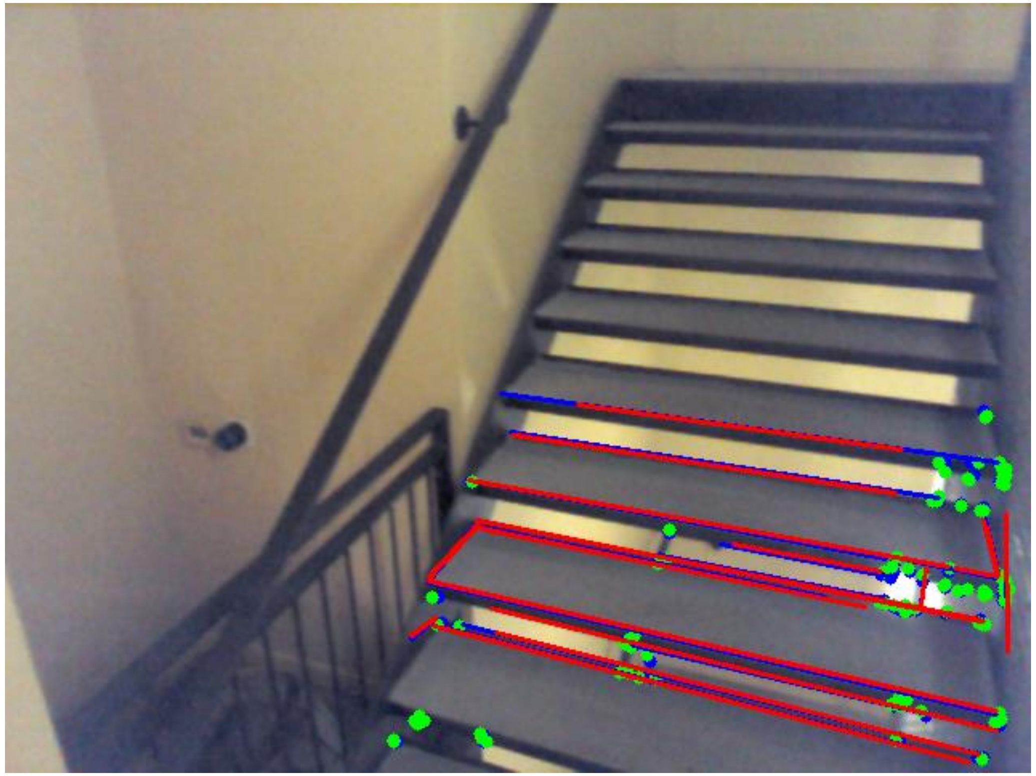

Hybrid localization with 4 solvers

Scene	HLoc	PtLine	Ours
Great Court	**9.5** / **0.05** / **20.4**	11.2 / 0.07 / 17.8	9.6 / **0.05** / 20.3
King's College	6.4 / **0.10** / 37.0	6.5 / **0.10** / 37.0	**6.2** / **0.10** / **39.4**
Old Hospital	12.5 / 0.23 / 22.5	12.7 / 0.24 / 20.9	**11.3** / **0.22** / **25.4**
Shop Facade	2.9 / 0.14 / 78.6	**2.7** / **0.12** / 79.6	**2.7** / 0.13 / **81.6**
St Mary's Church	**3.7** / 0.13 / 61.7	4.1 / 0.13 / 62.3	**3.7** / **0.12** / **63.8**
Avg.	7.0 / 0.13 / 44.0	7.4 / 0.13 / 43.5	**6.7** / **0.12** / **46.1**

Cambridge landmarks

Scene	HLoc	PtLine	Ours
Chess	**2.4** / **0.84** / **93.0**	**2.4** / 0.85 / 92.7	2.5 / 0.85 / 92.3
Fire	2.3 / 0.89 / 88.9	2.3 / 0.91 / 87.9	**2.1** / **0.84** / **95.5**
Heads	**1.1** / **0.75** / **95.9**	1.2 / 0.81 / 95.2	**1.1** / 0.76 / **95.9**
Office	3.1 / 0.91 / 77.0	3.2 / 0.96 / 74.5	**3.0** / **0.89** / **78.4**
Pumpkin	5.0 / 1.32 / 50.4	5.1 / 1.35 / 49.0	**4.7** / **1.23** / **52.9**
Redkitchen	4.2 / **1.39** / 58.9	4.3 / 1.42 / 58.0	**4.1** / **1.39** / **60.2**
Stairs	5.2 / 1.46 / 46.8	4.8 / 1.33 / 51.9	**3.7** / **1.02** / **71.1**
Avg.	3.3 / 1.08 / 73.0	3.3 / 1.09 / 72.7	**3.0** / **1.00** / **78.0**

7scenes RGB

		DUC 1	DUC 2
Points	HLoc	49.0 / 69.2 / 80.3	52.7 / **77.1** / 80.9
Points + Lines	PtLine	49.0 / 69.2 / **81.8**	56.5 / 76.3 / 80.2
	Ours	**49.5** / **72.2** / 81.3	**60.3** / 76.8 / **81.7**

InLoc

Scene	HLoc w/ Depth	PtLine	Ours w/ Depth
Chess	**2.4** / **0.81** / 94.8	**2.4** / **0.81** / **95.0**	**2.4** / 0.82 / 94.0
Fire	1.9 / 0.76 / 96.4	1.9 / 0.76 / **96.6**	**1.7** / **0.71** / **96.6**
Heads	1.1 / 0.73 / 99.0	1.1 / 0.74 / **99.4**	**1.0** / **0.72** / **99.4**
Office	2.7 / 0.83 / 83.7	2.7 / 0.83 / 83.9	**2.6** / **0.80** / **84.7**
Pumpkin	4.1 / **1.05** / **61.3**	4.0 / 1.06 / 60.8	**4.0** / **1.05** / 61.1
Redkitchen	3.3 / **1.12** / 72.1	**3.2** / **1.12** / 72.5	3.3 / **1.12** / **73.0**
Stairs	4.7 / 1.25 / 53.4	4.3 / 1.16 / 55.9	**3.2** / **0.86** / **76.0**
Avg.	2.9 / 0.94 / 80.1	2.8 / 0.93 / 80.6	**2.6** / **0.87** / **83.5**

7scenes RGBD

Consistent improvement with lines on public benchmarks

Localization on LaMAR [A]

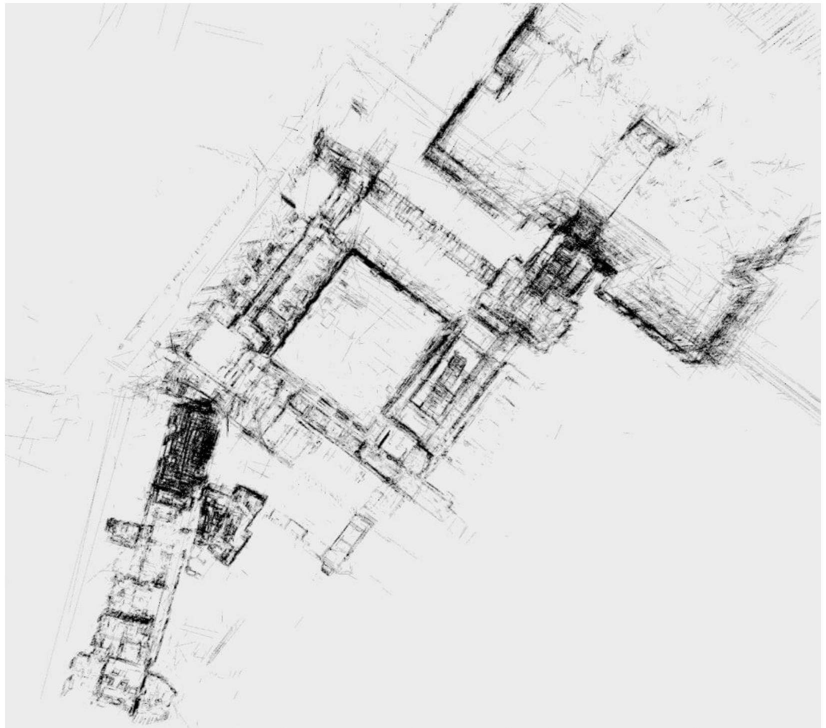

Large-scale 3D line maps of CAB building

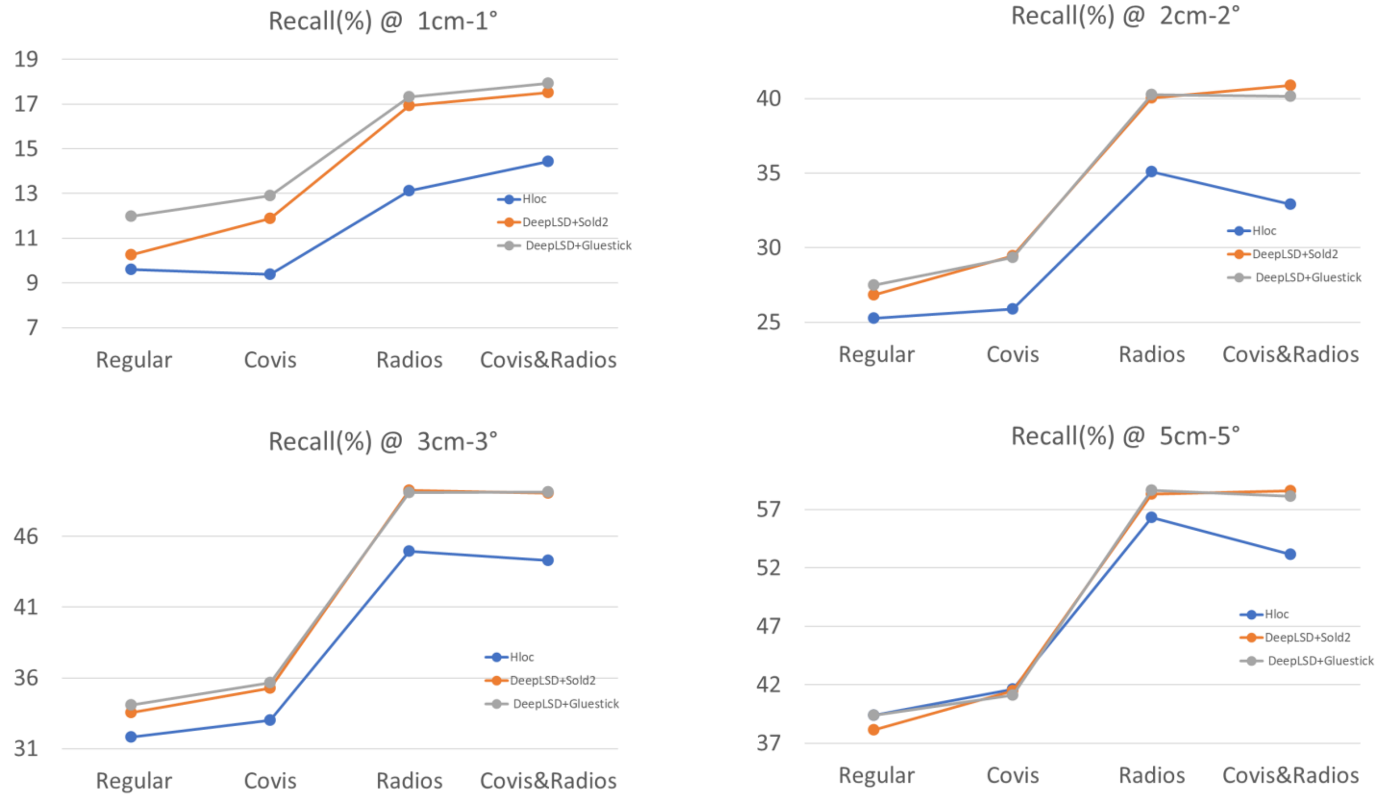

[A] Sarlin & Dusmanu et al. LaMAR: Benchmarking Localization and Mapping for AR, ECCV 2022
Slide credits: Thomas Birchler, Shinjeong Kim, Elias Salameh, and Aidyn Ubingazhibov from ETH Zurich

Hybrid bundle adjustment with points and lines

	COLMAP [42]	[42] + LIMAP (line-only)	[42] + LIMAP
ai_001_001	68.0 / 87.0 / 91.3	78.3 / 91.1 / 93.8	**80.0 / 91.7 / 94.2**
ai_001_002	75.2 / 90.2 / 94.0	87.5 / 95.6 / 97.3	**88.5 / 96.0 / 97.6**
ai_001_003	83.8 / 94.4 / 96.6	82.9 / 94.0 / 96.4	**85.7 / 95.1 / 97.1**
ai_001_004	**79.2 / 88.9 / 90.9**	67.1 / 82.1 / 86.0	77.3 / 88.3 / 90.6
ai_001_005	85.1 / 94.9 / 97.0	88.4 / 96.1 / 97.7	**90.9 / 97.0 / 98.2**
ai_001_006	83.4 / 93.1 / 95.7	80.2 / 92.9 / 95.7	**84.4 / 93.8 / 96.3**
ai_001_007	59.0 / 68.5 / 70.6	64.5 / **70.6 / 71.9**	**65.0** / 70.3 / 71.7
ai_001_008	84.9 / 94.9 / 96.9	89.5 / 96.5 / 97.9	**91.3 / 97.1 / 98.2**
Average ↑	77.3 / 89.0 / 91.6	79.8 / 89.9 / 92.1	**82.9 / 91.2 / 93.0**
Median error ↓	0.188	0.173	**0.146**

Line-assisted multi-view stereo

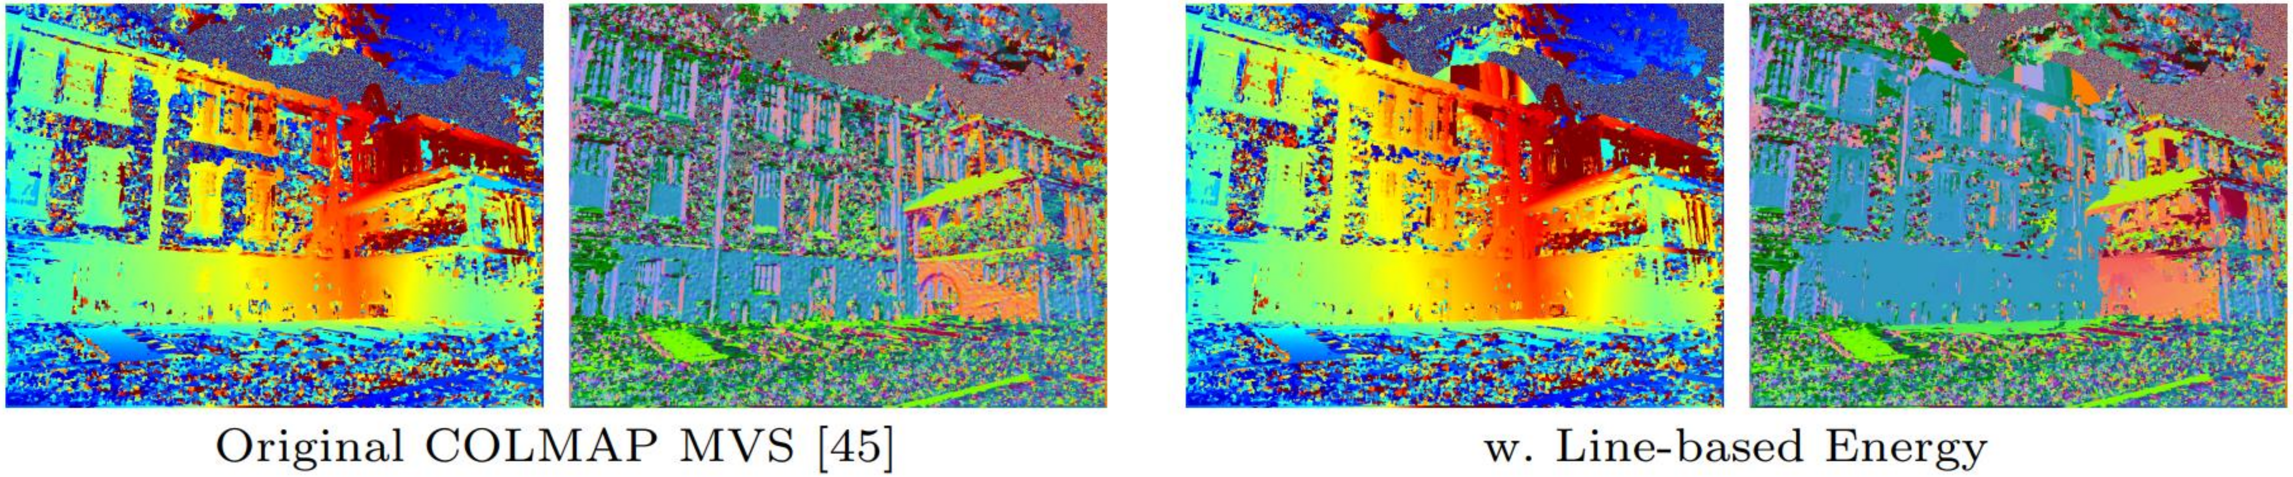

Preliminary line-assisted dense mapping.

Featuremetric Line Refinement

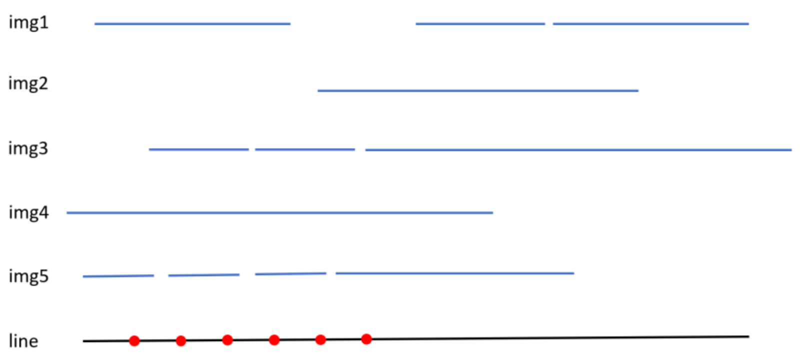

Sampling by 2-level intersection!

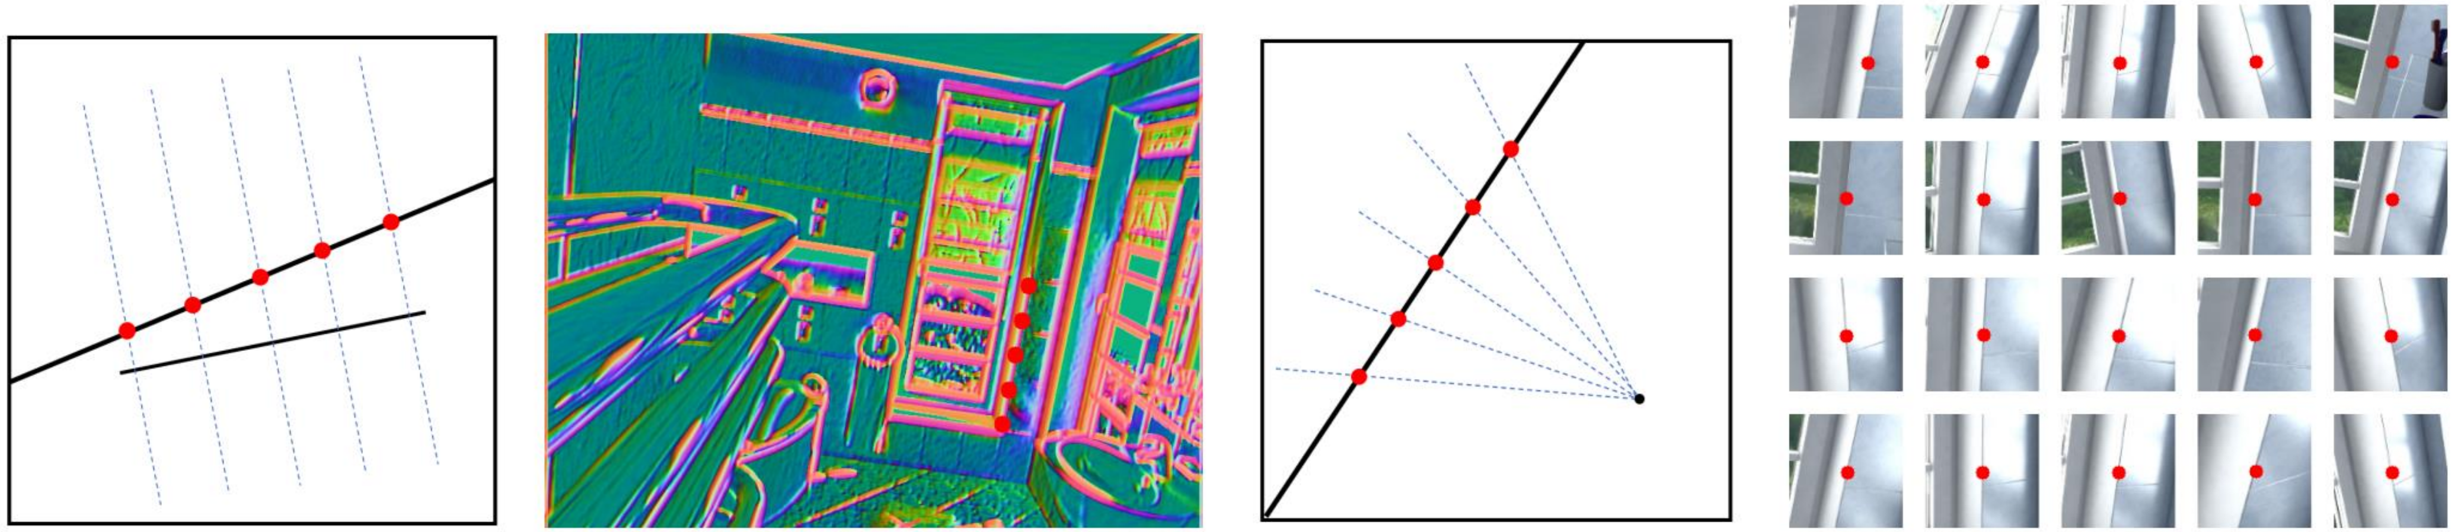

a) Reference image b) Feature map c) Target image d) Correspondences

Line patching with oriented bounding box to ensure scalability

Benchmarking

Matcher \ Detector	LSD	HAWPv3	TP-LSD	SOLD2	DeepLSD
LBD	42.2 / 58.5 / (14.0 / 14.6)	6.0 / 58.0 / (7.8 / 9.8)	21.6 / 73.2 / (9.1 / 9.3)	30.7 / 69.3 / (12.2 / 18.7)	64.6 / 70.0 / (15.8 / 18.1)
SOLD2	48.3 / 59.2 / (15.8 / 19.1)	14.7 / 62.7 / (11.2 / 20.1)	44.4 / 76.4 / (14.3 / 16.7)	50.8 / 74.4 / (15.1 / 32.2)	72.0 / 71.4 / (18.1 / 24.9)
L2D2	44.4 / 59.6 / (15.0 / 16.8)	13.5 / 63.4 / (10.7 / 18.3)	39.5 / **78.1** / (13.7 / 15.4)	43.9 / 72.8 / (13.7 / 24.9)	69.2 / 70.4 / (17.0 / 22.2)
LineTR	37.0 / 58.3 / (12.8 / 13.3)	5.4 / 60.5 / (8.4 / 10.7)	43.0 / 76.3 / (14.5 / 16.7)	29.0 / 70.1 / (12.3 / 19.9)	71.9 / 69.4 / (17.6 / 23.9)
Endpts SP + NN	48.8 / 58.6 / (15.5 / 18.2)	16.2 / 63.2 / (11.2 / 20.0)	43.7 / 75.8 / (14.3 / 16.5)	49.1 / 73.7 / (14.7 / 31.4)	72.8 / 70.3 / (17.7 / 24.0)
Endpts SP + SG	48.4 / 58.0 / (15.8 / 18.9)	16.0 / 61.9 / (11.3 / 20.9)	47.1 / 76.1 / (14.5 / 16.8)	50.0 / 72.8 / (15.5 / **34.4**)	**74.6** / 69.5 / (**18.2** / 24.8)

Pautrat et al. DeepLSD: Line Segment Detection and Refinement with Deep Image Gradients, CVPR 2023

Next Step: Hybrid Incremental SfM

JUNE 18-22, 2023
CVPR
VANCOUVER, CANADA

3D Line Mapping Revisited (THU-PM-080)

Shaohui Liu, Yifan Yu, Rémi Pautrat, Marc Pollefeys, Viktor Larsson

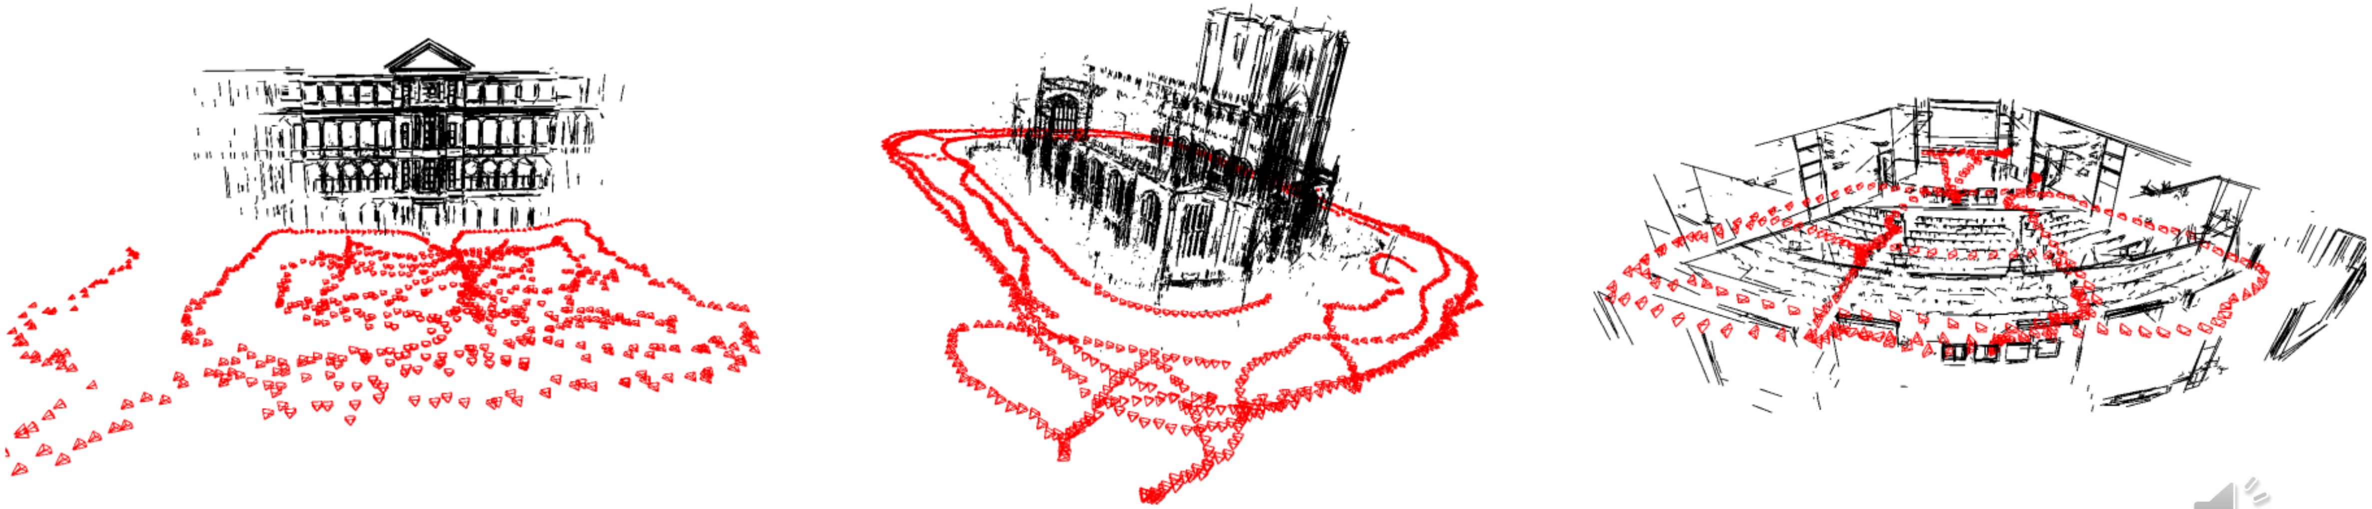

ETH zürich
Microsoft
LUND UNIVERSITY