Pose-graph via Adaptive Image Re-ordering

Daniel Barath, Jana Noskova, Ivan Eichhardt, Jiri Matas

BMVC 2022

Motivation

- Pose graph generation is **slow** in SfM. Especially, in global SfM.
- This is due to the „accept-or-reject" strategy of two-view estimations.

Many RANSAC iterations are spent on non-matchable image pairs.

Global Structure-from-Motion Pipeline

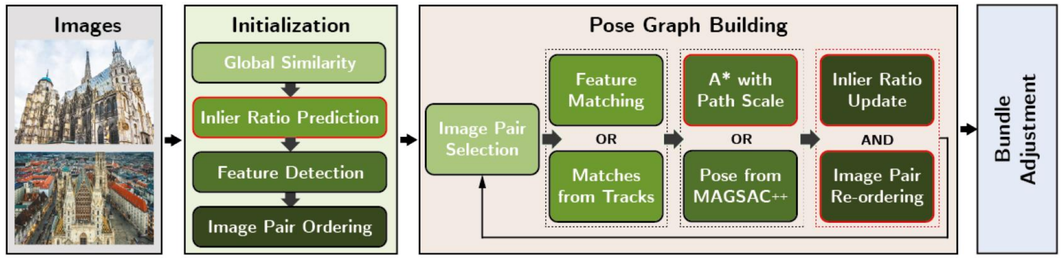

New components:

- Image pair revisiting (re-ordering) strategy.
- A*-based pose estimation on *scaled* paths.
- Expected inlier ratio update.

Image Pair Re-ordering Strategy

(a) Traditional "accept-or-reject" strategy.

(b) Proposed re-ordering strategy

- Always the *most-likely-to-match* image pair is selected for RANSAC.
- Only a few iterations in RANSAC, determined by the *expected inlier ratio*.
- Inlier ratio is prior comes from a deep network.

Expected Inlier Ratio (Bayesian) Update

- Given a prior inlier ratio $\mu_0 \in [0,1]$ for each image. Function $k(\mu_0)$ is the RANSAC iter. number.
- The random number of all-inlier-samples N_{all} in $k(\mu_t)$ follows a binomial distribution.
- Usual conjugate prior is a beta distribution with params. $a(t)$ and $b(t)$, where

Expectation in t-th iter.: $\mu_t = \left(\frac{a(t)}{a(t)+b(t)}\right)^{1/m}$ *Variance*: $v_t = \frac{a(t)b(t)}{\left(a(t)+b(t)\right)^2\left(a(t)+b(t)+1\right)}$

Real-world Experiments

Dataset: 13 scenes from 1DSfM (614366 pairs).

Run-times projected to a single CPU core.

Results:

	# edges	# inliers	total time (hours)	MAGSAC++ pose estimation t_{total}	MAGSAC++ pose estimation t_{avg}	MAGSAC++ pose estimation # runs	A* pose estimation t_{total}	A* pose estimation t_{avg}	A* pose estimation # runs
Baseline	417572	56148287	55.11	50.93	4.63	614366		–	
A* w/o scale [5]	524221	65176541	16.69	13.99	1.27	217109	**0.042**	**0.004**	525831
A* + scale + re-ord.	**554182**	68442654	**2.06**	**1.06**	**0.10**	301105	0.306	0.028	593712
Ablation study									
A* + no scale + re-ord.	532947	61900737	5.94	1.23	0.11	348344	0.057	0.005	634217
A* + scale + no re-ord.	538119	**70991857**	10.96	10.12	0.92	183904	0.281	0.026	529280
Baseline + re-ord.	392779	48070653	6.82	2.14	0.19	1174609		–	

- 8 times speed-up only with re-ordering.
- 27 times speed-up both with re-ordering and A*.
- Similar number of final inliers and pose-graph edges.

- Global SfM with the Theia library.
- Applied to pose-graph generated as proposed.

Results:

- Similar number of recon. cameras.
- Similar rotation accuracy, slightly better positions.

	# views	# points	AVG ε_R (°)	MED ε_R (°)	AVG ε_p (m)	MED ε_p (m)
Baseline	820	108161	9.83	7.41	3.14	2.19
A* w/o scale [5]	815	106336	9.80	7.41	3.18	2.25
A* + scale + re-ordering	**821**	**106810**	9.61	7.27	**3.05**	**2.04**
A* + no scale + re-ordering	816	106408	**9.45**	**7.02**	3.17	2.28
A* + scale + no re-ordering	**821**	107827	9.95	7.62	3.20	2.27
Baseline + re-ordering	819	107750	9.53	7.11	3.13	2.14

Conclusion:

- With the same accuracy,
- we achieve an order-of-magnitude speed-up, by
- a clever image pair re-selection / re-ordering strategy