Privacy Preserving Partial Localization

Marcel Geppert[1] Viktor Larsson[1,2] Johannes L. Schönberger[3] Marc Pollefeys[1,3]
[1]ETH Zurich [2]Lund University [3]Microsoft

Introduction

Cloud-based localization systems inevitably carry **privacy risks** for the user since potentially private information needs to be shared with a server. In this work we aim to **hide the user's exact pose from the server** while still being able to obtain this exact pose on the user side. To this end we

- Use geometric lifting to add degrees of freedom (DoF) in the map representation.
- Intentionally create an **underconstrained pose estimation problem**.
- **Query multiple, independent servers** for partial poses.
- Locally combine the partial poses into a full 6 DoF pose.

In contrast to previous work [2], our map representation **can not be inverted using density analysis attacks** [1].

Mathematical constraints

Partial constraints are derived directly from Point-to-Point alignment

$$R\,\boldsymbol{X}_q + \boldsymbol{t} = \boldsymbol{X}_m$$

Point constraint

$$\begin{bmatrix} \boldsymbol{r}_1 \\ \boldsymbol{r}_2 \\ \boldsymbol{r}_3 \end{bmatrix} \begin{bmatrix} x_q \\ y_q \\ z_q \end{bmatrix} + \begin{bmatrix} t_1 \\ t_2 \\ t_3 \end{bmatrix} = \begin{bmatrix} x_m \\ y_m \\ z_m \end{bmatrix}$$

Plane constraints

$$\begin{bmatrix} \boldsymbol{r}_1 \\ \boldsymbol{r}_2 \\ \boldsymbol{r}_3 \end{bmatrix} \begin{bmatrix} x_q \\ y_q \\ z_q \end{bmatrix} + \begin{bmatrix} t_1 \\ t_2 \\ t_3 \end{bmatrix} = \begin{bmatrix} x_m \\ y_m \\ z_m \end{bmatrix}$$

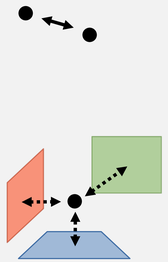

Minimal system

$$\begin{bmatrix} \boldsymbol{X}_{q1}^T & 1 \\ \boldsymbol{X}_{q2}^T & 1 \\ \boldsymbol{X}_{q3}^T & 1 \end{bmatrix} \begin{bmatrix} \boldsymbol{r}^T \\ t \end{bmatrix} = \begin{bmatrix} o_{m1} \\ o_{m2} \\ o_{m3} \end{bmatrix} \quad \text{s.t.} \|\boldsymbol{r}\| = 1$$

Method Overview

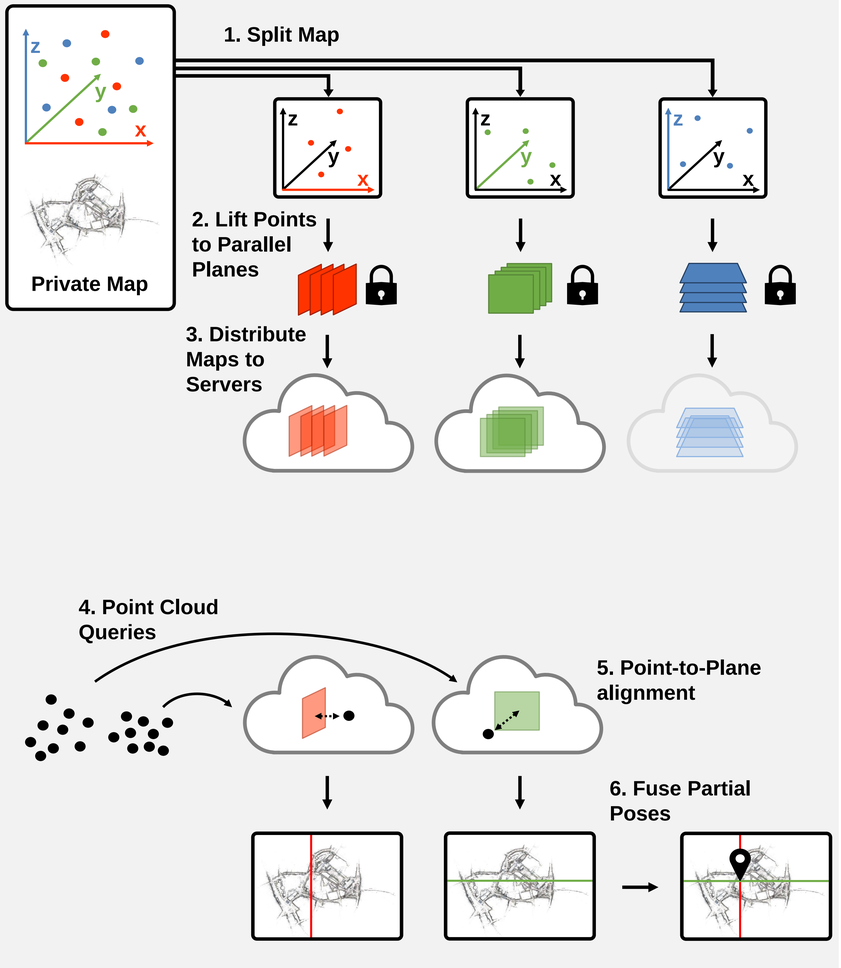

SfM Query Results

Scene	Point-to-Point [5] τ_1	Point-to-Point [5] τ_2	Point-to-Point [5] τ_3	Point-to-Plane [4] τ_1	Point-to-Plane [4] τ_2	Point-to-Plane [4] τ_3	Ours τ_1	Ours τ_2	Ours τ_3
Alamo	21.3	60.8	86.1	12.6	54.5	86.1	22.1	62.6	79.5
Gendarmenmarkt	7.7	40.8	72.4	4.6	33.4	61.6	5.6	31.1	56.7
Madrid Metropolis	4.7	32.1	68.2	1.8	23.0	59.9	6.6	36.1	62.8
Roman Forum	11.3	53.0	79.2	7.6	47.5	76.8	12.2	43.8	66.8
Tower of London	5.9	43.7	73.9	3.3	37.1	72.3	8.7	37.1	58.9

$\tau_1 = (0.05m\ /\ 2^\circ)$ $\tau_2 = (0.2m\ /\ 5^\circ)$ $\tau_3 = (0.5m\ /\ 10^\circ)$

Active Depth Query Results

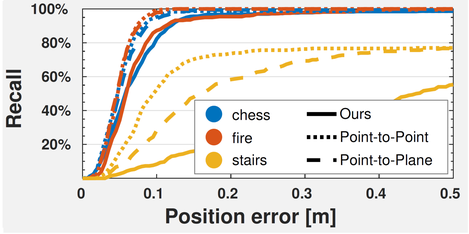

Limitations

- Query point clouds carry the risk of revealing scene appearance [2].
- Collaborating servers can still learn the user's location.
- Servers could accumulate query point clouds over time to circumvent lifted representation.
- Fewer constraints obtained.

Conclusion

- We present **a cloud-based localization system that divides the map representation and query** between three independent servers.
- The user's **exact pose remains hidden from the servers**.
- Prior knowledge allows pose estimation even from only two servers.
- The system exhibits only **small reductions in accuracy** compared to established baselines.

While the approach still experiences some practical disadvantages, we hope that this work further highlights the importance of privacy in visual localization and motivates future research in this area.

References

[1] Chelani et al., How Privacy-Preserving are Line Clouds? Recovering Scene Details from 3D Lines, CVPR 2021
[2] Pittaluga et al., Revealing scenes by inverting structure from motion reconstruction, CVPR 2019
[3] Speciale et al., Privacy Preserving Image-Based Localization, CVPR 2019
[4] Speciale et al., Privacy preserving image queries for camera localization, ICCV 2019
[5] Umeyama, Least-squares estimation of transformation parameters between two point patterns, PAMI 1991