Generic Merging of Structure from Motion Maps with a Low Memory Footprint

Gabrielle Flood, David Gillsjö, Patrik Persson, Anders Heyden, Kalle Åström

Centre for Mathematical Sciences, Lund University, Sweden

gabrielle.flood@math.lth.se

Motivation and Background

- We present a memory efficient method for merging of maps.
- The method also allows for non-rigid transformations within a map.
- The method can perform loop closure and it can be used in a divide and conquer manner for increased robustness during map merging.

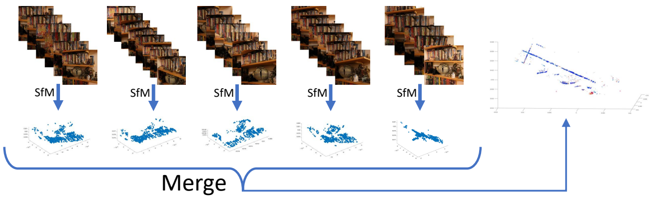

- A map is a set of 3D points, each with a position and a feature vector.
- Merged 3D points, camera matrices and coordinate system are estimated simultaneously in an efficient bundle over an approximative residual.
- The optimised parameters are denoted $\mathbf{z}$ and the residuals $\mathbf{r}$. The ML estimate of $\mathbf{z}$ is found by minimising the sum of squared residuals, $\mathbf{z}^* = \text{argmin}_{\mathbf{z}} \mathbf{r}^T\mathbf{r}$.

Pre-Processing

Separate maps of the same scene can be merged to a single, more accurate map. The separate maps are created using e.g. SLAM or SfM. The optimal residuals from these bundles can be linearised to decrease the memory footprint, according to in our previous paper Efficient Merging of Maps and Detection of Changes.

The parameters in $\mathbf{z}$ are ordered s.t. $\Delta\mathbf{z} = \begin{bmatrix}\Delta\mathbf{q} & \Delta\mathbf{s}\end{bmatrix}^T$. The Jacobian J is divided correspondingly, with J_a and J_b corresponding to $\mathbf{q}$ and $\mathbf{s}$, respectively. The auxiliary parameters in $\mathbf{s}$ will depend on the points in $\mathbf{q}$ as

$$\frac{\partial\mathbf{s}}{\partial\mathbf{q}} = -(J_b^T J_b)^{-1}(J_a^T J_b)^T.$$

That derivative can be used to express how a change in $\mathbf{q}$ affect the residuals,

$$\Delta\mathbf{r} = \underbrace{\left(J_a + J_b\cdot\frac{\partial\mathbf{s}}{\partial\mathbf{q}}\right)}_{J_q}\Delta\mathbf{q}.$$

Viewing the residual as a function of an update $\Delta\mathbf{q}$ and linearising it around an optimal point o gives the following approximation of the squared residual

$$\mathbf{r}^T\mathbf{r} \approx a^2 + \Delta\mathbf{q}^T R^T R \Delta\mathbf{q},$$

where $a^2 = \mathbf{r}|_o^T\mathbf{r}_o$ and R is a triangular matrix originating from QR-decomposition of $J_q|_o$. Note that R is much smaller than $J_q|_o$.

Merging Separate Maps

For the merge, unknown parameters are collected in a structure $\mathbf{w} = (\mathbf{q}, T_1, T_2, \ldots, T_N)$, where T_i is the tranformation connected to map i. We do local optimisation using an LM approach to minimise $\mathbf{r}^T\mathbf{r} \approx \hat{\mathbf{r}}^T\hat{\mathbf{r}}$ with

$$\hat{\mathbf{r}} = \begin{bmatrix} a^{(1)} \\ R^{(1)}(T_1 p_1(\mathbf{q}) - \mathbf{q}^{(1)}) \\ \vdots \\ a^{(N)} \\ R^{(N)}(T_N p_N(\mathbf{q}) - \mathbf{q}^{(N)}) \end{bmatrix}.$$

The last seven rows of the matrices $R^{(i)}$ are zero, due to gauge freedom. This is an issue, since the solution has to be close to the point we have linearised around and also because $\Delta\mathbf{q}^{(i)}$ can get large without affecting the error. To compensate for this we add a penalty, setting these rows othogonal to the other rows of $R^{(i)}$.

Using Merging for Increased Robustness

Denote $\tilde{a} = \bar{a}^2 - \sum_k \left(a^{(k)}\right)^2$ and let κ_i be the number of points that are common in i individual maps. Then, $\tilde{a}$ should be Γ distributed with mean

$$\mathbf{E}[\tilde{a}] = \sigma^2\left(\sum_k d_{dof}^{(k)} - d_{dof}\right) = \sigma^2\left(\left(\sum_{i=1}^{N} 3\kappa_i(i-1)\right) - 7\cdot(N-1)\right).$$

This can be used as a hypothesis test. If $\tilde{a}$ does not seem to come from this distribution after a merge, something is wrong. The maps can then be divided into other, smaller parts, for which the test is repeated.

Results

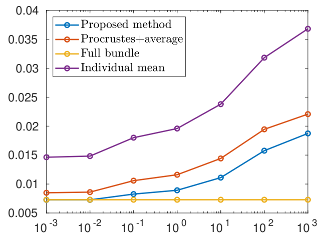

Simulated experiement. The plot shows how the RMSE for the final map achieved using different merging methods change when the individual map bundles are terminated at different levels. The x-axis shows at which level the individual bundles were stoppend in Euclidean norm of the gradient and the y-axis the RMSE. Our method performs as well as a full bundle for accurate individual maps and better than Procrustes in all cases.

Real experiment. The table below shows how the distance between some points in the map changed after different types of merging. The rightmost column shows the true distance measured using a tape. The experiment data is shown below.

Pt 1 ind	Pt 2 ind	Dist (mm) one map	Dist (mm) merge Pro.	Dist (mm) merge our	Dist (mm) gt
52	766	365	365	220	213
52	839	589	589	512	516
52	840	1358	1296	1264	1260
60	839	825	825	834	840
60	840	879	1023	860	857

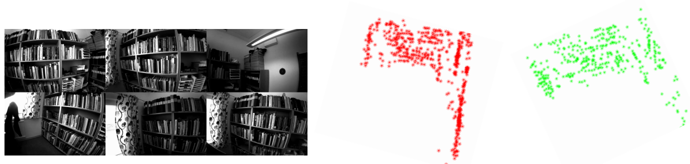

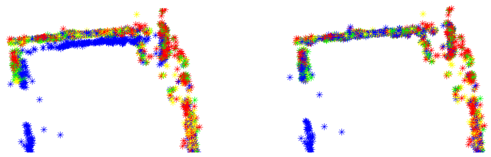

The top row shows a few of the used images and the 3D reconstructions. The bottom row shows parts of the merged map using Procrustes to the left and our proposed method to the right. Note that the top and left walls are doubled after the Procrustes registration, while our method solves that problem.

Conclusions

- With our proposed method, maps can be merged efficiently and robustly.
- No pre-alignment is necessary and the gauge freedom in the error function is accounted for by adding an additional penalty in those directions.
- Due to the non-rigid transformations problems like loop closure can be solved.
- The proposed method is faster than doing one large bundle and more accurate than Procrustes followed by averaging.