

Openvins feature initialization

Friday, April 3, 2020 4:48 PM

In code `ov-core/src/feat/FeatureInitializer.cpp`:

```
bool FeatureInitializer::single_triangulation(Feat *feat,
std::unordered_map<size_t, std::unordered_map<double, ClonePose>> &clones(CAM) {
    int total_meas = 0;
    size_t anchor_most_meas = 0;
    size_t most_meas = 0;
    for (auto const & pair : feat->timesteps) {
        total_meas += (int) pair.second.size();
        if (pair.second.size() > most_meas) {
            anchor_most_meas = pair.first;
            most_meas = pair.second.size();
        }
    }
}
```

In `feature.h`, there's a variable

`std::unordered_map<size_t, std::vector<double>> timestamps;`
`size_t` is the camera id, and `std::vector<double>` holds a vector of timestamps correspond to observations.

In `FeatureDatabase.h`:

```
void update_feature(size_t id, double timestamp, size_t cam_id,
float u, float v, float u_n, float v_n) {
```

```
    std::unique_lock<std::mutex> lck(mtx);
    if (features_id_lookup.find(id) != features_id_lookup.end()) {
        Feature *feat = features_id_lookup[id];
        feat->uvs[cam_id].emplace_back(Eigen::Vector2f(u, v));
        feat->uvs_norm[cam_id].emplace_back(Eigen::Vector2f(u_n, v_n));
        feat->timestamps[cam_id].emplace_back(timestamp);
        return;
    }
```

`emplace_back` adds the value to the end of vector.

`update_feature` is used in `track/TrackKLT.cpp`

```
void TrackKLT::feed_monocular(double timestamp, CV::Mat bImg, size_t cam_id) {
```

```
    for (size_t i = 0; i < good_left.size(); ++i) {
```

```
        cv::Point2f npt_l = undistort_point(good_left.at(i).pt, cam_id);
        data_base->update_feature(good_ids_left.at(i), timestamp, cam_id,
                                   good_left.at(i).pt.x, good_left.at(i).pt.y,
                                   npt_l.x, npt_l.y);
```

In `ov-core/src/test-tracking.cpp`

```
handle_Stereo(double time0, double time1, CV::Mat img0, CV::Mat img1) {
```

`extractor` → `feed_Stereo` (`time0, img0, img1, 0, 1`);

`extractor` → `feed_Monocular` (`time0, img0, 0`);

`extractor` → `feed_Monocular` (`time0, img1, 1`);

∴ `cam_id` means which camera, e.g. for monocular `cam_id` is 0,

for Stereo `cam_id` is 0,1, and select the camera that has

most observations across time to be anchor.

`Eigen::MatrixXd A = Eigen::MatrixXd::Zero(2 * total_meas, 3);`

`Eigen::MatrixXd b = Eigen::MatrixXd::Zero(2 * total_meas, 1);`

`size_t c = 0;`

`Clone Pose Anchor clone = clones.CAM.at(feat->anchor.cam_id).at(feat->anchor_clone_timestamp);`

`Eigen::Matrix<double, 3, 3> & R_GtoA = anchorclone.Rot();`

`Eigen::Matrix<double, 3, 1> & p_AinG = anchorclone.pos();`

`for (auto const & pair : feat->timesteps) {`

`for (size_t m = 0; m < feat->timesteps.at(pair.first).size(); ++m) {`

`Eigen::Matrix<double, 3, 3> & R_GtoCi = clones.CAM.at(pair.first).at(feat->timesteps.at(pair.first).at(m)).Rot();`

`Eigen::Matrix<double, 3, 1> & p_CinG = clones.CAM.at(pair.first).at(feat->timesteps.at(pair.first).at(m)).pos();`

`Eigen::Matrix<double, 3, 1> R_AtoCi;`

`R_AtoCi.noalias() = R_GtoCi * R_GtoA.transpose();`

`Eigen::Matrix<double, 3, 1> p_CinA;`

`p_CinA.noalias() = R_GtoA * (p_CinG - p_AinG);`

feature is triangulated in Anchor frame A, we can have the transformation from Ci:

$$\begin{aligned} C_i P_f &= A^T (P_f - A P_{C_i}) \\ A P_f &= A^T C_i P_f + A P_{C_i} \end{aligned} \quad (1)$$

`Eigen::Matrix<double, 3, 1> b_i;`

`b_i <-> feat->uvs_norm.at(pair.first).at(m)(0), feat->uvs_norm.at(pair.first).at(m)(1),`

`b_i = b_i / b_i.norm();`

`Eigen::Matrix<double, 2, 3> B_prep = Eigen::Matrix<double, 2, 3>::Zero();`

`B_prep <-> b_i(2, 0), 0, b_i(0, 0), b_i(2, 0), -b_i(1, 0);`

`A.block(2 * c, 0, 2, 3) = B_prep;`

`b.block(2 * c, 0, 2, 1).noalias() = B_prep * p_CinA;`

`c++;`

The user can indicate with the `noalias()` function that there is no aliasing in Eigen, as follows:

`matB.noalias() = matA * matA.`

This allows Eigen to evaluate the matrix product `matA * matA` directly into `matB`.

`Eigen::block`: block of size (p, q) starting at (i, j).

`matrix::block(i, j, p, q)`

The measurement in current frame is unknown bearing θ_b and depth z_f , we can map the feature in current frame via

$$C_i P_f = Z_f b_f = \begin{bmatrix} u_n \\ v_n \end{bmatrix}$$

where u_n, v_n are undistorted normalized coordinates.

The bearing can be warped into the anchor frame by

substituting eq. (1):

$$A^T C_i P_f = A^T (A^T C_i P_f + A P_{C_i})$$

$$= A^T A^T C_i P_f + A^T A P_{C_i}$$

$$= Z_f A^T b_f + A P_{C_i}$$

To remove the need to estimate the extra degree of freedom of depth Z_f , we define

$$A_{n_1} = \begin{bmatrix} -A_{bf}(3) & D & A_{bf}(1) \end{bmatrix}^T$$

$$A_{n_2} = \begin{bmatrix} 0 & A_{bf}(3) & -A_{bf}(2) \end{bmatrix}^T$$

these are perpendicular with vector A_{bf} , ∴

$$A_{n_1}^T A_{bf} = 0$$

$$A_{n_2}^T A_{bf} = 0$$

We can then multiply the eqs to form two eqs only depend on 3 dof unknown $A P_f$:

$$\begin{bmatrix} A_{n_1}^T \\ A_{n_2}^T \end{bmatrix} A P_f = \underbrace{\begin{bmatrix} A_{n_1}^T \\ A_{n_2}^T \end{bmatrix} Z_f b_f}_{\text{zero}} + \begin{bmatrix} A_{n_1}^T \\ A_{n_2}^T \end{bmatrix} A P_{C_i}$$

$$\begin{bmatrix} A_{n_1}^T \\ A_{n_2}^T \end{bmatrix} A P_f = \begin{bmatrix} A_{n_1}^T \\ A_{n_2}^T \end{bmatrix} A P_{C_i}$$

By stacking all measurements: B_{perp} in code

$$\begin{bmatrix} \vdots \\ A_{n_1}^T \\ A_{n_2}^T \\ \vdots \end{bmatrix} A P_f = \begin{bmatrix} \vdots \\ A_{n_1}^T \\ A_{n_2}^T \\ \vdots \end{bmatrix} A P_{C_i}$$

$$A \nearrow \begin{bmatrix} \vdots \\ A_{n_1}^T \\ A_{n_2}^T \\ \vdots \end{bmatrix} A P_f = \begin{bmatrix} \vdots \\ A_{n_1}^T \\ A_{n_2}^T \\ \vdots \end{bmatrix} A P_{C_i}$$

$$\uparrow \text{unknown} \nearrow \begin{bmatrix} \vdots \\ A_{n_1}^T \\ A_{n_2}^T \\ \vdots \end{bmatrix} b$$

`Eigen::MatrixXd p_f = A.colPivHouseholderQr().solve(b);`

since each pixel measurement provides two constraints, as long as $m \geq 3$, we have enough constraints to triangulate.

In practice, more views of feature the better, so we choose at least five views.

We also check that the triangulated feature is valid and in front of the camera and not too far.

The condition number of the above linear system is used to reject features that are sensitive to noise.

`Eigen::JacobiSVD<Eigen::MatrixXd> svd(A, Eigen::ComputeThinU | Eigen::ComputeThinV);`

`Eigen::MatrixXd singularValues;`

`singularValues.resize(svd.singularValues().rows(), 1);`

`singularValues = svd.singularValues();`

`double condA = singularValues(0, 0) / singularValues(singularValues.rows() - 1, 0);`

`if (std::abs(condA) > options.max_cond_number ||`

`p_f(2, 0) < -options.min_dist ||`

`p_f(2, 0) > -options.max_dist ||`

`std::isnan(p_f.norm())) {`

`return false;`

`}`

`feat->p_FinA = p_f;`

`feat->p_FinG = R_GtoA.transpose() * feat->p_FinA + p_AinG;`

`return true;`