



3DF Samantha: user manual v 1.60

3Dflow srl

February 23, 2014

1 Changelog

1.1 Changes from v 1.50 to v 1.60

- Improved speed and memory footprint with large datasets
- Improved matching accuracy and speed
- Improved autocalibration

1.2 Changes from v 1.41 to v 1.50

- Improved autocalibration.
- Minor reconstruction improvements.
- Improved matching speed.
- Minimum requirement for Cuda Capabilities raised to 2.0.
- Automatically switch to CPU if Cuda not found.

1.3 Changes from v 1.40 to v 1.41

- Fixed a (rare) bug in autocalibration.
- Changed server connection settings for the online updater.
- Updated to the last version of the external library Eigen.

1.4 Changes from v 1.31 to v 1.4

- Added 32 bit support.
- Improved keypoint extraction and description (robustness and accuracy improved).

- The KeypointDensity and KeypointMaxN parameters have been removed and have been substituted with KeypointMeanN, i.e. the mean number of keypoint to extract for each image. The extraction process has become a bit slower, but more stable.
- Faster matching phase.
- Fixed a bug in the two step reconstruction (accuracy improved).
- Fixed a bug in two views triangulation (robustness improved).
- Image undistortion is now skipped if no distortion is present (improved speed).
- Fixed a crash during ANN computation.
- Added exporter for the bundler v3 format.

1.5 Changes from v 1.30 to v 1.31

- Fixed a bug that made the application crash when exporting for CMVS.
- Added online check for new versions (this check can be disabled from settings).

1.6 Changes from v 1.20 to v 1.30

- Local bundle adjustment - global speed improvements on the geometric reconstruction phase.
- Point Recovery Strategy - improvements on global robustness.
- Final Point Cloud cleanup heuristics.
- Minor fixes and improvements on projective pipeline.
- Exterior fitting procedure changed (using Procrustean algorithm).
- Autocalibration is now the default choice in settings.

1.7 Changes from v 1.11 to v 1.20

- Major improvements on autocalibration and variable intrinsics bundle adjustment.
- Optional tangential radial distortion coefficient Adjustment.
- Novel two-Step autocalibration technique.
- Minor fix on keypoint descriptor computation near the images border.

- The maximum number of keypoints (KeypointMaxN parameter) forces the keypoints to span as more image area as possible.
- Points triangulation speed improvements.
- Paths management and default settings changed.

1.8 Changes from v 1.1 to v 1.11

- Fixed a bug in KeypointDensity property.
- Fixed a minor bug in descriptor computation.
- Added the possibility to switch multithreading off during ANN computation.

1.9 Changes from v 1.0 to v. 1.1

- Novel Scale-Space robust keypoint extractor and descriptor.
- Multicore ANN stage.
- Minor fixes in Multimatching.
- Improved clusterization.
- Minor fix on MSE computation.
- Minor Bundle Adjustment improvements.

2 Introduction

3DF Samantha is the commercial implementation provided by 3Dflow¹ of a hierarchical structure from motion pipeline pipeline. The software takes in input a set of images and produces a 3D sparse reconstruction of the scene.

3DF Samantha uses a hierarchical approach similar to [?, ?]. With respect to the above papers there are several differences and improvements. In particular, the matching phase has been completely revisited and uses a proprietary keypoint extractor and descriptor and the geometrical clustering phase is different in many parts.

The binary software is distributed for research purposes only and cannot be redistributed or used for commercial purposes. See the file `License.pdf` for a detailed description.

¹<http://www.3dflow.net>

3 Usage

3DF Samantha is distributed in the form of compiled binaries that can run on a Windows 64 bit or 32 bit platform. To speed up the computation, it is highly recommended to have a system equipped with a mid/high-end Cuda enabled gpu.

The usage is:

```
3DFSamantha [setting_file]
```

All the parameters, including the folder where the input images are to be sought are specified in the `[setting_file]` (a plain text file).

By default a file called `settings.ini` is loaded, unless the user specifies another name at the command line. If the file is not present in the current folder, a default `settings.ini` is created and used.

Before running the program the user needs to:

- put all the images in a folder
- specify the desired settings in `settings.ini`
- (optional) provide camera calibration data (if known) in a file called `calibrations.txt`. The content of this file overrides the default calibration data given in `settings.ini`.

The program runs the auto-calibrated motion and structure algorithm by default, so it does not need the camera parameters.

In order to activate the calibrated version with autocalibration, one should set `AdjustIntrinsic=false` and `ProjectiveReconstruction=false` in `settings.ini` (see next Section).

By default a `ply` file is produced which can be opened with your favourite visualizer (Meshlab [?], for example). A simple SaM visualizer, `Ogrettha` is also included in the distribution. Please press F1 inside `Ogrettha` for a quick help. For the other available output formats please see the next section.

4 Settings

The default `settings.ini` is the following:

```
[Workspace]
Path=./Test/

[Export]
CmvsPmvs2=false
Ogrettha=false
PlyTxt=true
TextMatlab=false

[CameraCalibration]
FocalX=1000
FocalY=1000
PrincipalPtX=1000
PrincipalPtY=1000
Skew=0
RadialDistK1=0
RadialDistK2=0
RadialDistK3=0
RadialDistP1=0
RadialDistP2=0

[Cuda]
Device=0

[Samantha]
MinimumMatches=10
MinimumTrackLength=3
KeypointMeanN=7500
ANNNumberofNeighbors=8
ANNNumberofFeatures=500
ANNApproximationThreshold=0.0
CheckUpdates=true
NumberOfMatchViews=0
UseSymmetricMatches=true
StructSequence=0
ReprojectionThreshold=3.0
MaxBAIterations=200
LinkageMinSearchRange=1
ImportKeypointsFromTxt=false
ForceSequentialPipeline=false
AdjustIntrinsic=true
AdjustTangentialDistortion=false
AdjustRadialDistortion=true
ProjectiveReconstruction=true
MultithreadANN=true
FinalMaxReprojectionError=1.0
FinalMinimumTrackLength=3
PointRecovery=true
LocalBa=true
```

Path Folder where the input is supposed to be.

CmvsPmvs2 If set to **true** a file compatible with Furukawa's CMVS/PMVS is produced. Results are exported in [Path]/Export.

Ogrettha If set to **true** a file (output.ply) in a format compatible with our viewer (Ogrettha) is produced. Results are exported in [Path]/Export.

PlyTxt If set to **true** a file ply format (SamPoints.ply points only) and txt files containing camera ppms are produced. Every .ppm files contains the camera projection matrix (first 3 rows) and the radial distortion coefficient in the last row (k1 k2 c1 c2 k3). Results are exported in [Path]/Export.

TextMatlab If set to **true** a (large!) MATLAB file is produced, containing the following variables:

ppm camera projection matrices

pts 3d points

ptsDescr mapping from 3d points to 2d descriptors

descr descriptors properties (position, scale, angle, camera, histogram)

distanceMatrix view distance matrix

clusterInfo dendrogram of views in the classic matlab 'linkage' format (doc linkage for info)

[**CameraCalibration**] The following 10 items refers to the default internal parameters of the cameras. They are overridden by those specified in **calibrations.txt** and **ignored if autocalibration is on** (**ProjectiveReconstruction=false**). The radial distortion parameters follows the classical Bown's model and are consistent with those produced by the well known camera calibration toolbox².

FocalX Default camera focal length in X

FocalY Default camera focal length in Y

PrincipalPtX Default camera principal point X coordinate

PrincipalPtY Default camera principal point y coordinate

Skew Default camera skew

RadialDistK1

RadialDistK2

RadialDistK3

RadialDistP1

²<http://www.vision.caltech.edu/bouguetj/calibdoc/>

RadialDistP2

Device Id of the preferred CUDA graphics card to use. Set to -1 to use the cpu. Set to 0 to use the primary video adapter.

KeypointMeanN The mean number of keypoints that are to be extracted for each image (0=unlimited).

ImportKeypointsFromTxt Set to `true` to import keypoints from files. This is useful to bypass the built-in keypoint extractor and use any other detector / descriptor. The keypoints files have to be put in the images directory. Each keypoints file must have the same image basename with a `.float` (in case of floating point descriptors) or `.bool` extension (in case of boolean descriptors). The files have to be organized in the same format of the SIFT32 keypoints executable (freely downloadable³). Please note however that image dimension is limited to 1800 with that program. The included SIFTWin32Driver.exe utility automatically generates the sift keypoints `.float` text files for each image. The program requires the SIFT32.exe to be downloaded and put in the same directory.

MultithreadANN Turn threading on or off during Neighbors view search stage. Turn off if you are having problems.

ANNApproximationThreshold This is the ϵ parameters of ANN (see).

ANNNumberOfFeatures In order to establish which image is to be matched with which, a fixed number of keypoints is considered in each image;

ANNNumberOfNeighbours How many neighbours are considered for each keypoint. This is the ℓ parameters in [?].

NumberOfMatchViews This the number of views to be matched, m parameter in [?]. If set to 0 is automatically computed as $4 + (\#\text{views})/100$. Beware that in this case the cost of matching is not linear in the number of views any more.

StructSequence If images are known to be in a sequence, the discovery of which image matches with which can be by-passed: 1=linear, 2=circular, 0=no structure (default).

UseSymmetricMatches Specifies which algorithm is used to match keypoints. Symmetric, as opposed to one-way, means that a match from image I to image J must be also verified from image J to image I.

NBucketsPerWidth Keypoint matching follows a bucketing strategy. This parameter specifies the number of buckets in the horizontal dimension.

³<http://www.cs.ubc.ca/~lowe/keypoints/>

LinkageMinSearchRange This parameter specify the linkage strategy (the ℓ parameter in [?]). A number higher than 1 favors a more balanced tree and consequently the computing time decreases.

ReprojectionThreshold Used to discard points when their reprojection error in an image is too large during reconstruction. The parameter is expressed in pixels.

MaxBAIterations Max number of BA iterations. The default value (=1000) is only a safeguard, usually BA converges in few iterations.

MinimumMatches Two clusters that have less than **MinimumMatches** in common are not merged.

MinimumTrackLength Only tracks whose length is greater or equal to this value are considered during reconstruction. Default is 3 and normally it should not be changed.

CheckUpdates If set to **true**, Samantha will check online if the current version is the latest available.

AdjustIntrinsic If set to **true** intrinsic parameters are modified during BA, otherwise they are kept fixed. If they come from precise calibration it is recommended to keep them fixed; If they are guessed from EXIF data or computed via autocalibration (see below) they should be adjusted.

AdjustRadialDistortion If set to **true** the radial distortion is let to vary during BA, otherwise it is kept fixed.

AdjustTangentialDistortion If set to **true** the tangential radial distortion coefficient are let to vary during BA, otherwise they are kept fixed. Ignored if **AdjustRadialDistortion** is set to **false**.

ProjectiveReconstruction If set to **true** activates autocalibration (needed when camera parameters are unknown).

LocalBa If set to **true** activates the local bundle adjustment strategy. The geometric reconstruction will go faster with no noticeable differences.

PointRecovery If set to **true** activates a bad tracks/point recovery strategy during reconstruction. The geometric reconstruction phase will be slower but generally more robust.

FinalMinimumTrackLength Only tracks whose final length is greater or equal to this value gets triangulated and exported in the end.

FinalMaxReprojectionError Used to discard points when their reprojection error in an image is too large before exporting the final point cloud. The parameter is expressed in pixels.

5 Camera parameters

The `calibrations.txt` file contains one line for each different group of cameras sharing the same intrinsic parameters (the group can be made of one camera). Each group must be identified by a unique prefix in the file name. These parameters overrides those specified in `settings.ini` -> `[CameraCalibration]`.

In the case of autocalibration, `calibrations.txt` is used to provide information on which cameras share the same parameters, but the parameters themselves are ignored.

If `calibrations.txt` do not exist, all the cameras are considered to have equal intrinsics (either specified in `settings.ini` or autocalibrated).

The line has the following format:

```
[Prefix] [FocalX] [FocalY] [PrincipalPtX] [PrincipalPtY] [RadialDistK1]
[RadialDistK2] [RadialDistP1] [RadialDistP2]
```

For example:

```
IMGA 2881.65 2874.8 1394.35 1127.74 0 0 0 0
IMGB 3731.1 3725.93 1685.77 1284.09 0 0 0 0
```

Please note that `[Skew]` and `[RadialDistK3]` are considered to be zero.

6 Suggestions

The most common failure mode is obtaining a partial reconstruction, meaning that a relevant fraction of cameras have been not added to the (one) final reconstruction or that there is not a single reconstruction at all but an archipelago of smaller, independent reconstructions. The cause may be inherent in the data if images do not have sufficient overlap or may depend on parameters settings. In the latter case, a “sloppier” setting may help. This can be obtained by changing the following parameters:

`KeypointMeanN` Increase

`NumberOfMatchViews` Increase

`MatchingType` Use one-way matching

`LinkageMinSearchRange` Set to 1

`PInlier` Decrease

`SigmaNoise` Increase

`ReprojectionThreshold` Increase

`MinimumMatches` Decrease (not less than 5)

`MinimumTrackLength` Set to 2 (but only if everything else fails)

A change in the opposite direction makes the algorithm more conservative (few but good): the final reconstruction will have less points and less views but the result is more accurate and reliable.

Please note that the computing time is likely to be longer with the sloppier setting.

7 Contacts

Please report any bug/suggestion to: samantha@3dflow.net

For information regarding a commercial license of 3DF Samantha or related software please contact 3Dflow srl at info@3dflow.net.

References

- [1] M. Farenzena, A. Fusiello, R. Gherardi, Structure-and-motion pipeline on a hierarchical cluster tree, in: *IEEE International Workshop on 3-D Digital Imaging and Modeling*, Kyoto, Japan, 2009.
- [2] R. Gherardi, M. Farenzena, A. Fusiello, Improving the efficiency of hierarchical structure-and-motion, in: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR 2010)*, San Francisco, CA, 2010, pp. 1594 – 1600.
- [3] <http://meshlab.sourceforge.net/>