FFI Norwegian Defence Research Establishment

23/00148

FFI-RAPPORT

Experimentation of Vision-aided Inertial Navigation System (VaINS) on a small fixed-wing UAV

Baheerathan Sivalingam
Ove-Kent Hagen

Experimentation of Vision-aided Inertial Navigation System (VaINS) on a small fixed-wing UAV

Baheerathan Sivalingam
Ove-Kent Hagen

Norwegian Defence Research Establishment (FFI) 16 March 2023

Keywords
Navigasjon
Bildepunkter
Gjenfinning
Sensorer
Treghetsnavigasjon
Ubemannede luftfarkoster (UAV)

FFI report
23/00148

Project number
1493

Electronic ISBN
978-82-464-3467-4

Approvers
Halvor Bjordal, *Research Manager*
Halvor Ajer, *Research Director*

The document is electronically approved and therefore has no handwritten signature

Summary

This document describes an experiment using the Vision-aided Inertial Navigation System (VaINS). The system collects and processes sensor data. It transmits the navigation solution to an operator terminal in real-time. This is the first experiment with VaINS on a fixed-wing UAV, where it was used as a data-logging box. This experiment was a collaboration activity between FFI and NORCE, who operated the UAV. The data were collected over two days with five flights overlooking buildings, vegetation, a river, a lake, and mountains. The pitch angle of the camera was varied between grazing angles 22˚, 25˚ and 90˚ downwards. The primary objective of this experiment was to experience the use of the VaINS-box on a fixed-wing UAV and to study the performance of the image-aided navigation. The navigation solution of the image-aided INS is compared with the reference solution within GPS denied time intervals. The flights with camera pitch angle at 90˚ performed slightly better than the flights with grazing angles. Comparing the image-aided solution with the free inertial solution in the GPS denied periods, the image-aided navigation performed significantly better in the linear flights.

Sammendrag

Dette dokumentet beskriver et eksperiment som bruker Vision-aided Inertial Navigation System (VaINS). Systemet samler inn og behandler sensordata. Det overfører navigasjonsløsningen til en operatørterminal i sanntid. Dette er det første eksperimentet med VaINS på en fastvinget UAV. VaINS ble i denne omgang brukt som en dataloggingboks. Eksperimentet var et samarbeid mellom FFI og NORCE, som opererte UAV-en. Dataene ble samlet inn over to dager, med i alt fem flygninger over bygninger, vegetasjon, elv, innsjø og fjell. Vertikalvinkelen («pitch angle») til kameraet ble variert mellom lavvinkler, 22˚, 25˚ og 90˚ nedover. Hovedmålet med eksperimentet var å få erfaring med bruk av VaINS-boksen på en fastvinget UAV, og å studere ytelsen til den bildestøttede navigasjonen. Navigasjonsløsningen til bildestøttet INS ble sammenlignet med en referanseløsning innenfor GPS-nektede tidsintervaller. Flygningene der kameraets vertikalvinkel var 90˚ ga litt bedre ytelse enn flygningene med lavvinkel. Sammenligningene mellom den bildestøttede løsningen og den frie treghetsløsningen («free inertial») i GPS-nektede perioder, viste at den bildestøttede navigasjonen ga betydelig bedre ytelse i de lineære delene av flygningene.

Contents

1 Introduction

For Unmanned Aerial Systems (UAS), a navigation solution is necessary to determine the vehicles' position and orientation. For the vast majority of systems, the primary means of navigation is Global Navigation Satellite System (GNSS) combined with an Inertial Navigation System (INS). Unavailability of the GNSS signal leads the navigation solution based on INS alone to an unacceptable error level after a short time. One potential solution to the problem is that of image-aided navigation in one of its various forms.

The FFI project 1493, Drone detection and counter measures, has been investigating an image-aided inertial navigation system to cope with GNSS failure. The system [1] is based on integrating inertial sensor data with position data of image-features-points over an image sequence in an error-state Extended Kalman Filter (EKF).

Earlier versions of the system was tested with real image data from a manned helicopter, an Octocopter and a Hex-copter with rotary-wings. In all the above experiments, the data was processed and analyzed offline. A real-time navigation box, named Vision-aided Inertial Navigation System (VaINS) [2], was later developed. The VaINS-box, see Figure 1.1, collects sensor data, processes the data and transmits the navigation solution to an operator-terminal in real-time. It has been tested handheld, and in a field-test on an Octocopter.

In this experiment, VaINS is used as a data-logging box on a fixed-wing UAV named Scout (Figure 1.2), from NORCE [3]. NORCE is an independent research institute that conducts research for both public and private sectors, and were responsible for mounting the VaINS-box and operating the aircraft. This is the first experiment with VaINS in a fixed-wing UAV. The experiment was carried out at Elvenes airstrip in Salangen municipality area. Flights were within a predefined area around the Elvenes airstrip overlooking buildings, vegetation, mountains, a lake and a river.

The primary objective of this experiment is to experience the use of the VaINS-box on a fixed-wing UAV, and to study the performance of the image-aided INS. This document describes the experimental setup, data collection, data analysis and summarizes the experiences and the performance.

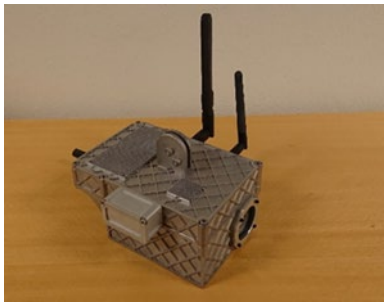

Figure 1.1 A real-time navigation box named VaINS

Figure 1.2 Fixed-wing UAV named Scout

2 Experimental setup

This experiment was a collaboration activity between FFI and NORCE. NORCE was responsible for mounting the VaINS-box in the aircraft (Figure 2.1), and planning of the flight operation. A drone pilot from NORCE was responsible for take-off, maneuvering flight and landing. Whereas the flight-coordinator from NORCE was responsible for planning, coordinating and monitoring of the flights.

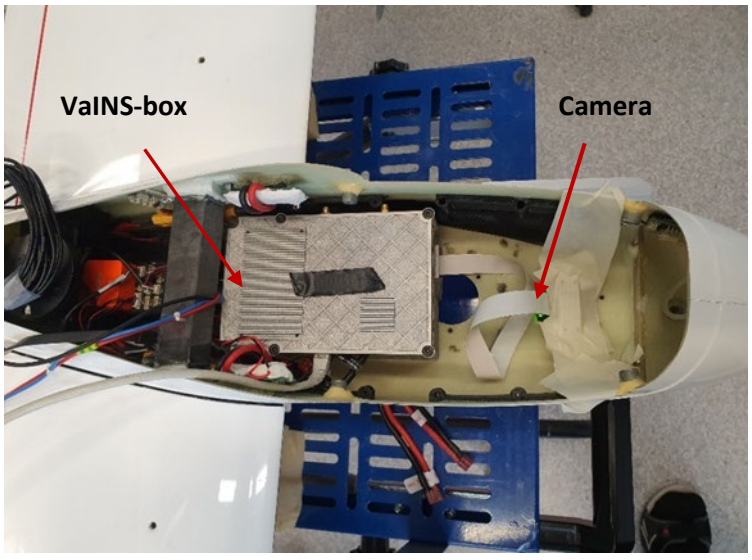

Figure 2.1 The VaINS-box was mounted inside the aircraft. The camera was mounted looking downwards with variable pitch angles

The VaINS box ended up being fixed with Velcro in the back part of the payload bay. The camera was mounted in the front, at a hole in the fuselage, on a fixture regulating the different pitch angles. This is not ideal, but rather a practical solution given the time available for preparation. Ideally, the camera and the IMU should be mounted together on a rigid fixture. Since we needed to remove the VaINS box after each flight, there is an added uncertainty of different misalignments of the IMU and camera for each flight. Because of this, we planned for making a low-altitude straight-line pass over the airfield buildings in each flight, for future data-driven calibration of the camera-IMU misalignment.

Figure 2.2 The aircraft was launched from a customized launching devise (left), and monitored from a mobile ground control station (right).

Cruise speed	15 – 20 m/s
Wingspan	2.62 m
Maximum endurance	3 hrs (Li-ION packs)
Standard empty weight	4.3 kg
Batteries - Li-Ion	2.6 kg
MTOW	10 kg
Maximum useful load	3.1 kg

Table 2.1 Aircraft specification

The aircraft was launched from a customized launching devise and monitored from a mobile ground control station (Figure 2.2). The above table, Table 2.1, tabulates the aircraft specification. FFI was responsible for guiding the flight planning, data logging and data analysis. The startup functions for logging were launched from a Linux PC via WiFi connection to VaINS. For more detail, refer section 5 in [2]. The aircraft was launched after all the functions were started.

The flight's path, cruise speed and altitude were monitored from the mobile ground control station, but we did not have real-time status report from VaINS once it was airborne. After landing, the box was connected to the Linux PC via Ethernet cable to download the data. The data was transferred to a Windows PC, where the quality and the quantity of the data was checked by a Matlab function.

3 Data collection

The data were collected in the neighborhood of Elvenes airstrip in Salangen municipality over two days, 21/6 and 22/6-2022, see Figure 3.1. Only one flight was operated on the first day due to rain, but four flights were flown on the second day.

Figure 3.1 Typical flight trajectory over Elvenes airstrip in Salangen municipality

All the flights have more or less the same trajectory within a predefine area overlooking buildings, vegetation, a river and a lake, and surrounded by mountains. However, the pitch angle of the camera was different, and the height was varied in the flights. Pitch angle 22˚ and 25˚ are selected based on the experience from the previous experiments. The following table, Table 3.1, tabulates the flight details.

Flight #	**Date & Time**	**Camera's pitch angle (degree)**	**Start heading (degree)**	**Data collection (sec)**	
				After launch	**Total**
1	21.06.2022, 11:15	22	220	670	1284
2	22.06.2022, 09:37	22	226	736	2373
3	22.06.2022, 12:22	90	238	882	1120
4	22.06.2022, 13:00	90	250	742	992
5	22.06.2022, 13:57	25	75	903	1653

Table 3.1 Flight details

Flight time, after launch, of all the flights are around 20 – 25 min. (1200 – 1500 sec.). However, the data collection stopped around 11-15 min. (660 – 900 sec.) after launch.

After some debugging on-site and reproduction of this error afterwards, we identified the problem. It turn out it is common that the requested timeout for the secure shell (SSH) terminal connections are not handled [5]. Therefore, when the terminals disconnected, the applications started in the terminals were aborted after about 17 minutes. The periods without WiFi connection in this experiment was longer than what we had experienced with VaINS before. The problem is not present if all processes are started automatically on the unit, as intended when we are not primarily logging as in this case, or by installing e.g. tmux [6] for improved terminal handling.

On the first day, initial preparation took a long time, and the first flight crashed right after the first launch. The second launch was successful. However, due to the drizzling rain, further flying was stopped for that day, because raindrops on the lens may lower the image quality.

The second day was more productive. On that day, images were taken for intrinsic camera calibration, navigation data were collected for IMU calibration and four flights were operated.

Due to a helicopter flying nearby, the fifth flight were delayed by nearly 15 min. The start heading of the flight was also different, at 75 degrees, from the other flights due to change of wind direction.

In all the five flights, the initial part of flying height is around 80 m and the main part is around 180 m. The following figure shows an example of an image for each pitch angle of the camera. The camera system in VaINS was set to provide 8-bit monochromatic images with 10 Hz frame rate at a resolution of 1600 x 1200 pixels.

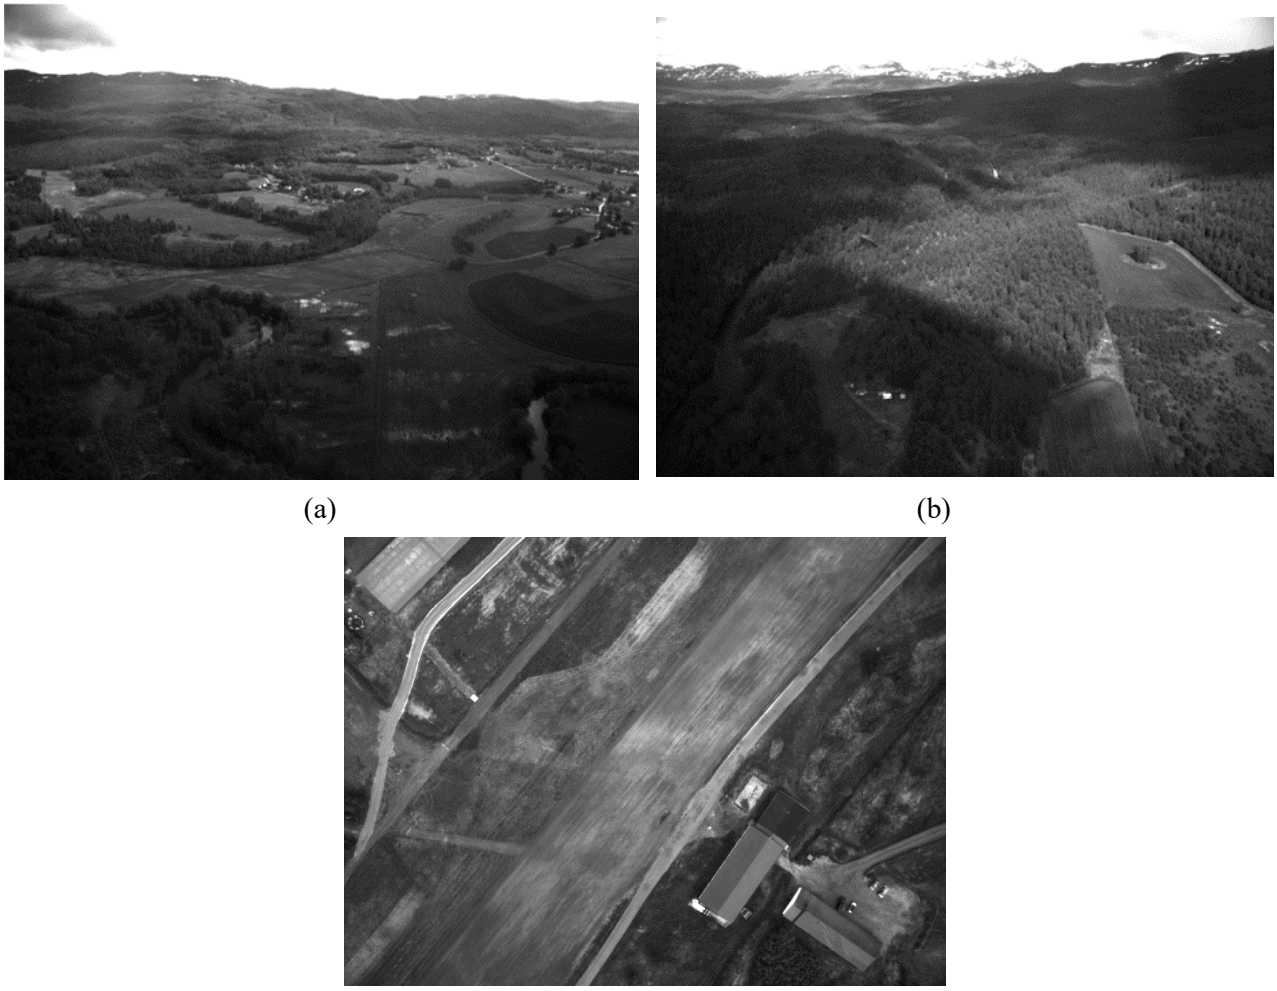

(a) (b)

(c)

Figure 3.2 Example images for each pitch angle of the camera. (a) 25˚ (b) 22˚ (c) 90˚

NORCE also provided the logs from the UAV to FFI. The logs contain the UAV navigation solution and data from its navigation sensors: three IMUs, two barometers, a GPS and a Pitot-tube for airspeed. The data from NORCE gave us an extra redundancy for the post-processing analysis. We also consider future use of some of these aiding sensors. In that respect they need to be time adjusted by subtracting a 2 hour offset to align with the VaINS data. The barometers provide a good estimate for height during GPS outage (Figure 3.3), when adjusted for initial

bias and scale factor errors, as can be estimated in the EKF. The Pitot tube airspeed may be a bit more challenging to integrate (Figure 3.4). The difference with respect to GPS speed over ground is due to both the unknown wind, and possibly a crab angle. However, it will provide a limit to the drift in velocity when integrated in the EKF, especially in the case when no image-aiding is available.

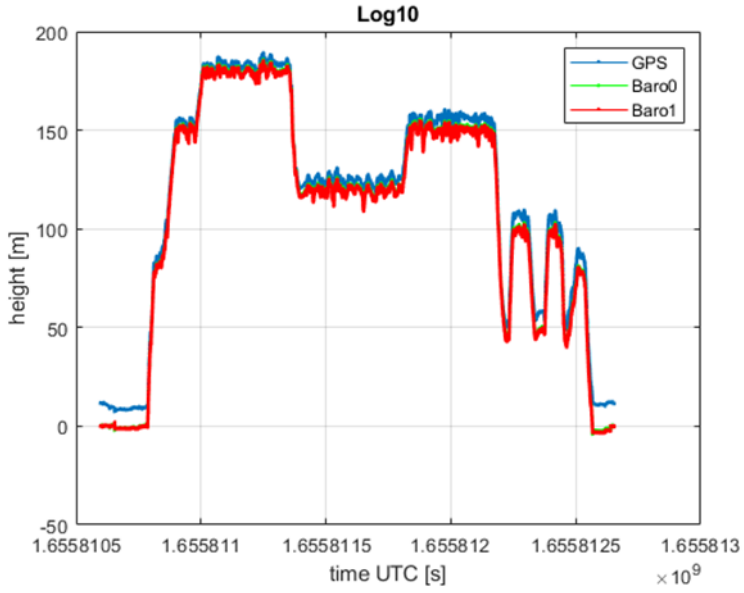

Figure 3.3 Example of GPS height versus the two barometers of the UAV in flight 1

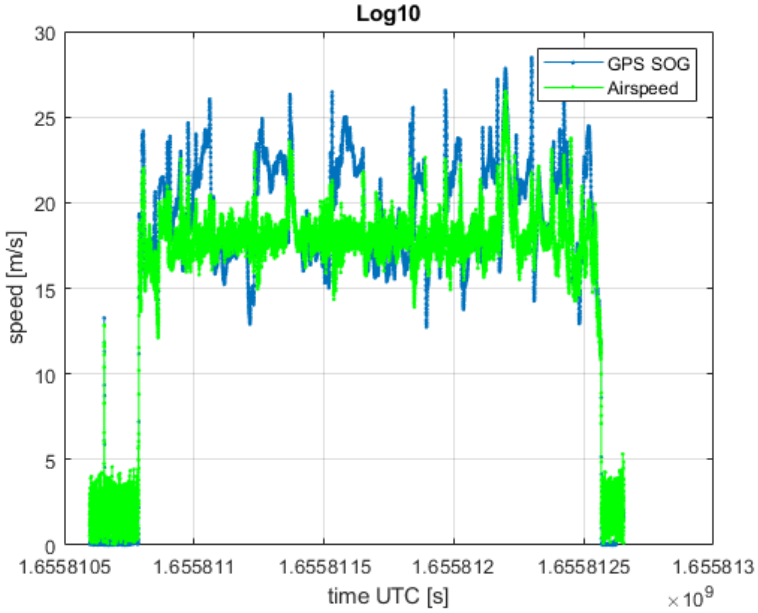

Figure 3.4 GPS speed over ground versus the Pitot tube airspeed from the UAV in flight 1

4 Analysis

NavLab4 [4] is used to process and analyze the data from VaINS offline. The IMU data and GPS data are processed with the optimal smoother in NavLab4 to provide a reference data set for position and orientation of the aircraft. To test the image-aided INS, the image-feature-points data (12 points at 10Hz) are used together with the IMU data and GPS height within the simulated GPS denied time intervals. GPS height was used for convenience only, but can be replaced by the barometer height measurements from NORCE.

The navigation solution of the image-aided INS is compared with the reference solution. It is compared separately within the linear flight and within the banking maneuver performed by the UAV during turns. The performance is assessed by the root-mean-square (RMS) error between the image-aided INS and the reference solution at the end of each GPS denied time interval. For the straight-line parts, it make sense to view the error both in per cent of travelled distance (PTD) and as a mean speed error (MSE) over this time interval, and for the banking parts we only use the mean speed error for comparison.

After some problems found in the initial analysis, we had to perform a coarse alignment estimation between the camera and the IMU. The coarse pitch alignment is estimated from identifying any systematic drift in the vertical direction, while coarse yaw alignment is estimated from any systematic body-fixed drift in the horizontal plane. Since VaINS was refitted before each flight and without assuring the same mounting alignment, we had to repeat this for each flight. The results of this misalignment estimation is therefore reported for each flight.

From this analysis we suspect the actual brackets used in flight 1 and 2 were swapped with the one used in flight 5, i.e. that the nominal pitch angle in flight 1 and 2 should be 25°, and nominal pitch in flight 5 should be 22°. This can easier explain the difference in the actual pitch alignment found in the analysis. This is confirmed by analyzing the pitch angle of the image data of the same scene from the flight 1 and 2 with flight 5

Flight	1	2	5
Pitch angle	25˚	25˚	22˚

Table 4.1 Nominal pitch angles of camera for flights 1, 2 and 5

We also found some lesser problems using the default fixed initial range guess method used in smaller scale and smaller speed scenarios. A new initial range guess method based on assuming flat topography beneath the UAV, was tested for flight 1-4, while flight 5, the one with the most grazing angle used the old method of fixed initial range.

4.1 Flight number 1

The nominal pitch angle of the camera is 25° and the flight time after launching is 670 sec. It has three linear flights and three banking maneuvers. Figure 4.1 presents the trajectory of the first two linear flights and banking maneuvers with image-aided navigation when GPS data is disabled. The rest of the maneuvers are similar to the presented ones.

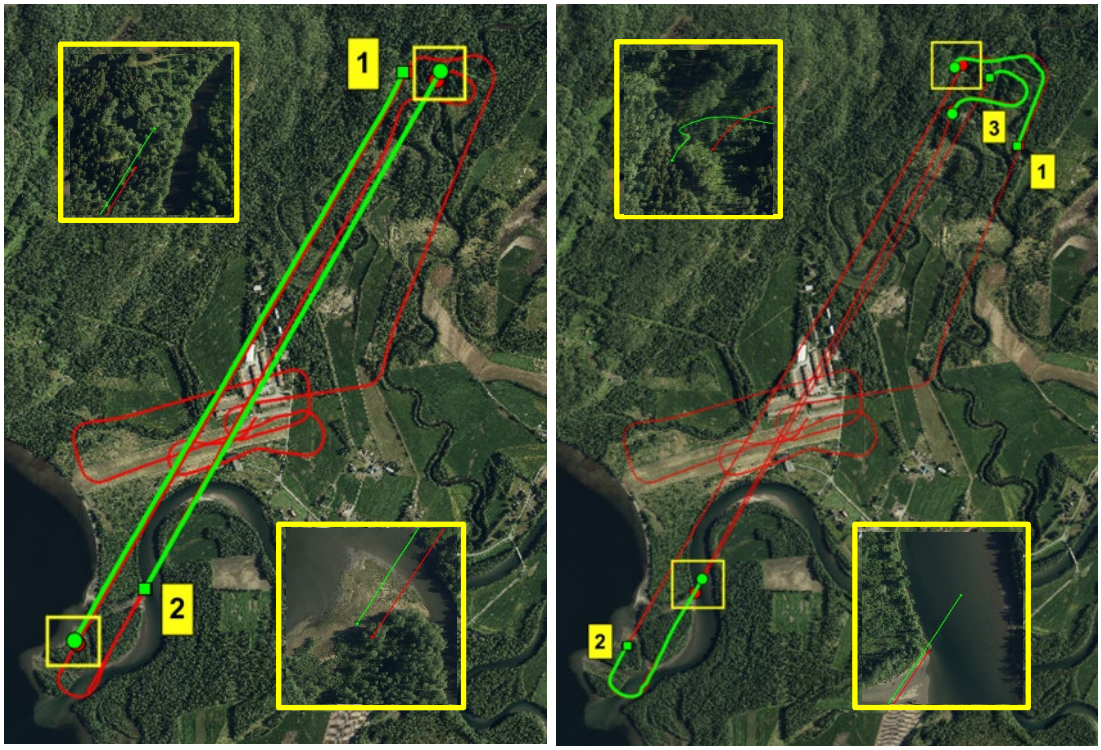

Figure 4.1 Trajectory of the linear flights (left) and banking maneuvers (right) in the first flight with image-aided navigation when GPS data is disabled. The red line represents the reference trajectory and the green line represents the image-aided navigation trajectory. Start point is marked by the square and the ending by the circle.

R_ImuCam (zyx-convention)	Nominal (°)	Misalignment (°)
x	180	0
y	-25	-2
z	0	-5

Table 4.2 Camera-IMU alignment for flight 1

The misalignment of the camera with respect to the IMU is coarsely estimated from the data, and the results for this flight is shown in Table 4.2.

The Table 4.3 tabulates the duration, distance traveled and the drift compared to reference data under the image-aided navigation in the straight line parts. The drift is presented in RMS error, in MSE and in PTD.

Duration (sec)	Distance (m)	RMS error (m)	MSE (m/s)	PTD (%)
144	2424	15.2	0.11	0.63
97	2154	32.2	0.33	1.50
134	2350	40.2	0.30	1.71

Table 4.3 Duration of image-aided navigation (flight 1) when GPS is disabled, distance traveled within the duration and the drift in RMS error, in MSE and in PTD

The duration of image-aided navigation during banking and the drift compared to reference data in RMS error and in MSE are tabulated in Table 4.3.

Duration (sec)	RMS error (m)	MSE (m/s)
30	33.0	1.10
36	53.3	1.48
23	8.0	0.35

Table 4.4 The duration of image-aided navigation (flight 1) and the drift in RMS error and in MSE under banking maneuvers.

The drift (MSE) in linear flights is lower than the drift in banking maneuvers.

4.2 Flight number 2

The camera's nominal pitch angle of this flight is 25˚. Even though this flight time is longer than the first flight, most of the time has been spent on the initial flying. The flight has only two linear flights and two banking maneuvers. The trajectories of the linear flights and banking maneuvers with image-aided navigation and the reference trajectory are shown in Figure 4.2.

The misalignment of the camera with respect to the IMU is coarsely estimated from the data, and for this flight, it is the same as flight 1, as shown in Table 4.5.

R_ImuCam (zyx-convention)	Nominal (°)	Misalignment (°)
x	180	0
y	-25	-2
z	0	-5

Table 4.5 Camera-IMU alignment for flight 2

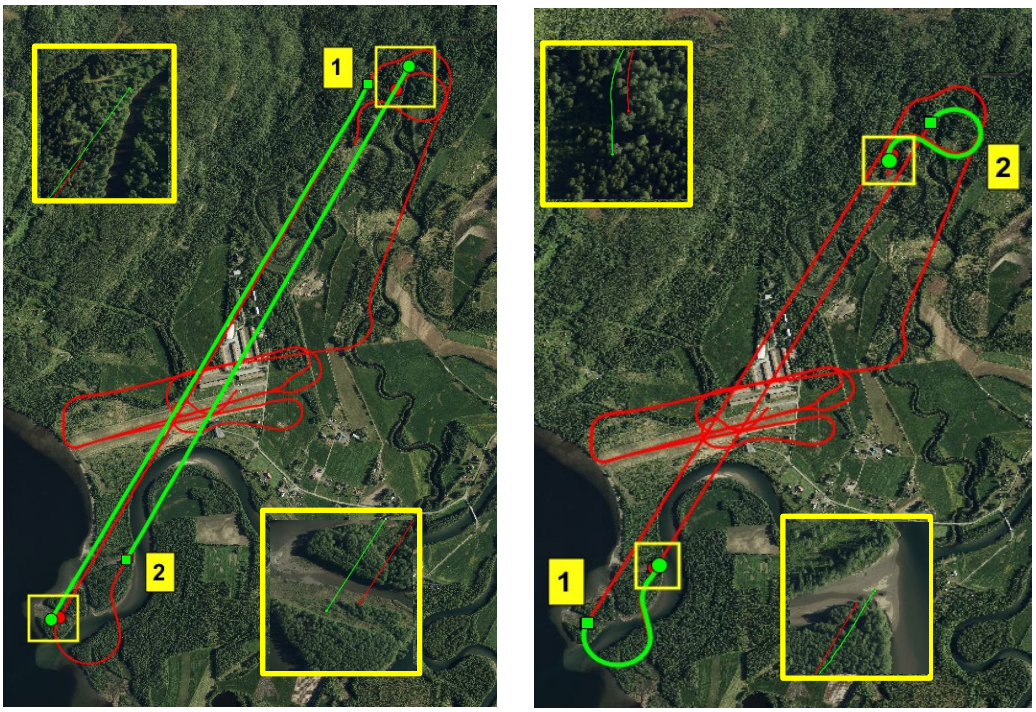

Figure 4.2 The trajectories of the linear flights (left) and banking maneuvers (right) in the second flight with image-aided navigation when GPS data is disabled (green line) and the reference trajectory (red line)

The following table tabulates the duration, distance and the drift compared to reference data under the linear flight.

Duration (sec)	Distance (m)	RMS error (m)	MSE (m/s)	PTD (%)
160	2507	37.7	0.24	1.50
113	2202	112.7	1.00	5.12

Table 4.6 Duration of image-aided navigation (flight 2) when GPS is disabled, distance traveled within the duration and the drift in RMS error, in MSE and in PTD.

Duration (sec)	RMS error (m)	MSE (m/s)
42	21.2	0.50
42	17.2	0.41

Table 4.7 The duration of image-aided navigation (flight 2) and the drift in RMS error and in MSE under banking maneuvers.

The duration and the drift in RMS error and in MES for banking maneuvers are tabulated Table 4.7. In this flight, there is generally lower drift during banking than in flight 1. The linear part varies greatly between the two laps, but is overall comparable to the banking parts.

4.3 Flight number 3

Flight number 3 has the camera's pitch angle at 90 degrees and has longer flight time after launch compared to the other flights. It has three linear flights and three banking maneuvers. The trajectories of the first two linear flights and the banking maneuvers with image-aided navigation and the reference solution are shown in Figure 4.3.

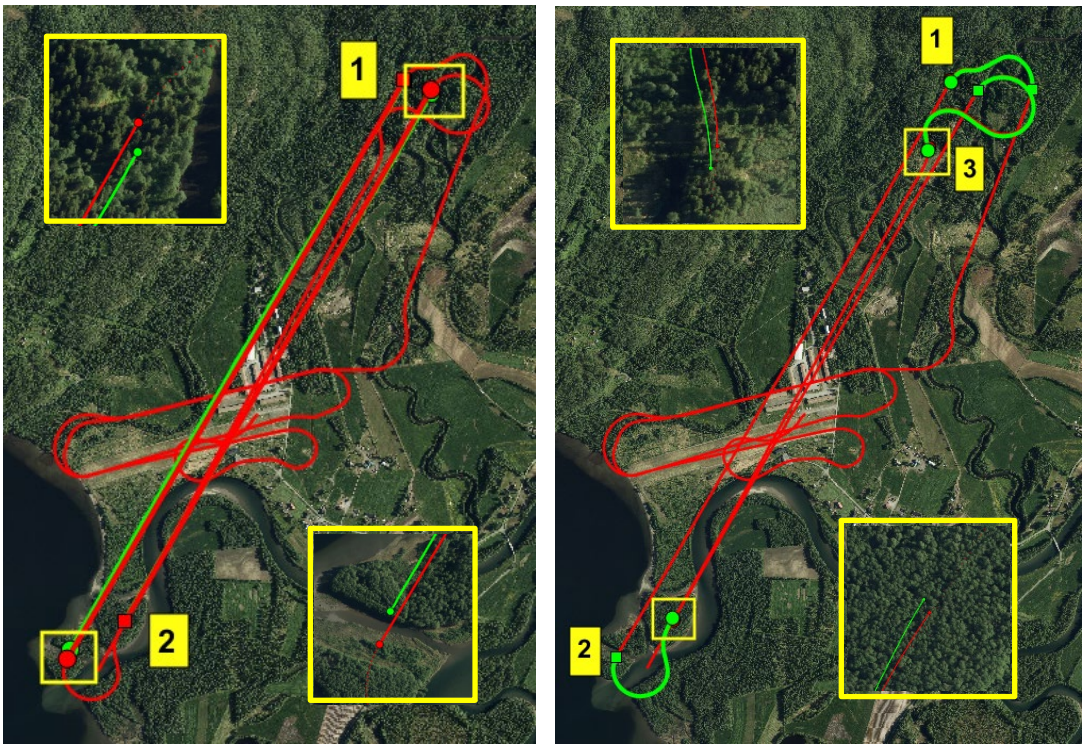

Figure 4.3 The trajectories of the linear flights (left) and banking maneuvers (right) in the third flight with image-aided navigation when GPS data is disabled (green line) and the reference trajectory (red line).

The GPS on VaINS had reception problems during the initial part of this flight, so for this flight, we use the GPS data from the UAV provided by NORCE instead.

The misalignment of the camera with respect to the IMU is coarsely estimated from the data, and for this flight, it is shown in Table 4.8.

R_ImuCam (zyx-convention)	Nominal (°)	Misalignment (°)
x	180	0
y	-90	-2.5
z	0	-3

Table 4.8 Camera-IMU alignment for flight 3

Table 4.9 tabulates the duration, distance and the drift of the linear flight.

Duration (sec)	Distance (m)	RMS error (m)	MSE (m/s)	PTD (%)
173	2555	38.6	0.22	1.51
120	2339	14.0	0.12	0.60
149	2245	62.8	0.42	2.80

Table 4.9 Duration of image-aided navigation (flight 3) when GPS is disabled, distance traveled within the duration and the drift in RMS error, in MSE and in PTD.

The first and the second lap has lower drift compared to the third lap. Table 4.10 tabulates the duration and the drift statistics of the banking maneuvers.

Duration (sec)	RMS error (m)	Drift (m/s)
28	6.1	0.22
33	8.0	0.24
47	8.5	0.18

Table 4.10 The duration of image-aided navigation (flight 3) and the drift in RMS error and in MSE under banking maneuvers.

It is interesting to note that all three laps during banking have low drift, and comparable to the linear parts.

4.4 Flight number 4

The camera's pitch angle is same as the third flight (90 °), and the flight time after launching is 742 sec. The first two trajectories of the linear flights and the banking maneuvers with image-aided navigation and the reference solution are shown in Figure 4.4.

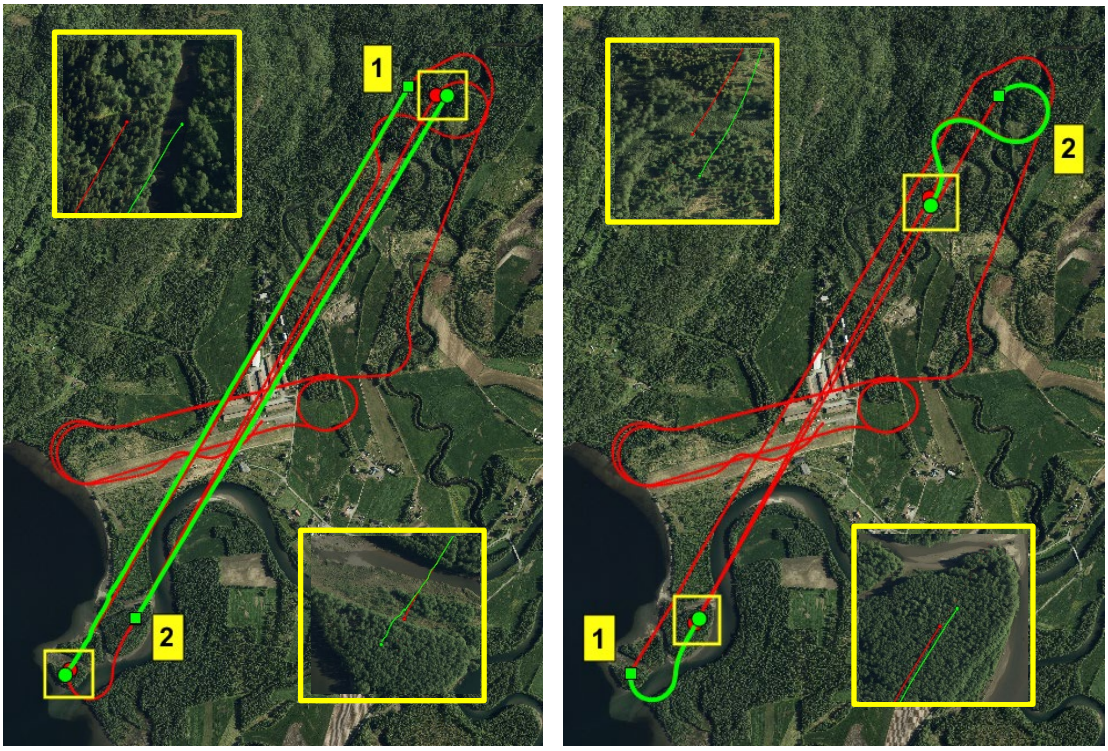

Figure 4.4 The trajectories of the linear flights (left) and banking maneuvers (right) in the fourth flight with image-aided navigation when GPS data is disabled (green line) and the reference trajectory (red line).

The misalignment of the camera with respect to the IMU is coarsely estimated from the data, and for this flight, it is found to be the same as in flight 4, as shown in Table 4.11.

R_ImuCam (zyx-convention)	Nominal (°)	Misalignment (°)
x	-180	0
y	-90	-2.5
z	0	-3

Table 4.11 Camera-IMU alignment for flight 4

The following Table 4.12 tabulates the duration, distance and the drift of the linear flights.

Duration (sec)	Distance (m)	RMS error (m)	MSE (m/s)	PTD (%)
158	2576	20.0	0.13	0.78
123	2310	36.9	0.30	1.60
67	1090	20.5	0.31	1.88

Table 4.12 Duration of image-aided navigation (flight 4) when GPS is disabled, distance traveled within the duration and the drift in RMS error, in MSE and in PTD.

Table 4.13 tabulates the duration and the drift statistics of the banking maneuvers.

Duration (sec)	RMS error (m)	MSE (m/s)
30	16.3	0.54
57	24.4	0.43

Table 4.13 The duration of image-aided navigation (flight 4) and the drift in RMS error and in MSE under banking maneuvers.

The drift is just slightly lower in the linear parts than during banking.

4.5 Flight number 5

The camera's nominal pitch angle for this flight is 22˚. Due to a helicopter flying in the neighborhood, this flight was delayed by nearly 15 min. Therefore, the total flight time was quite long. In addition, flight time after launch is longer than the other flights. However, only one flight was operated in this configuration due to the drizzle.

The misalignment of the camera with respect to the IMU is coarsely estimated from the data, and for this flight, it is shown in Table 4.14.

R_ImuCam (zyx-convention)	Nominal (°)	Misalignment (°)
x	-180	0
y	-22	-0.7
z	0	0

Table 4.14 Camera-IMU alignment for flight 5.

This flight has four linear flights and four banking maneuvers. The first two trajectories of the linear flights and the banking maneuvers with image-aided navigation and the reference solution are shown in Figure 4.5.

The drift statistics of linear and banking maneuvers are tabulated in Table 4.15 and Table 4.16 respectively.

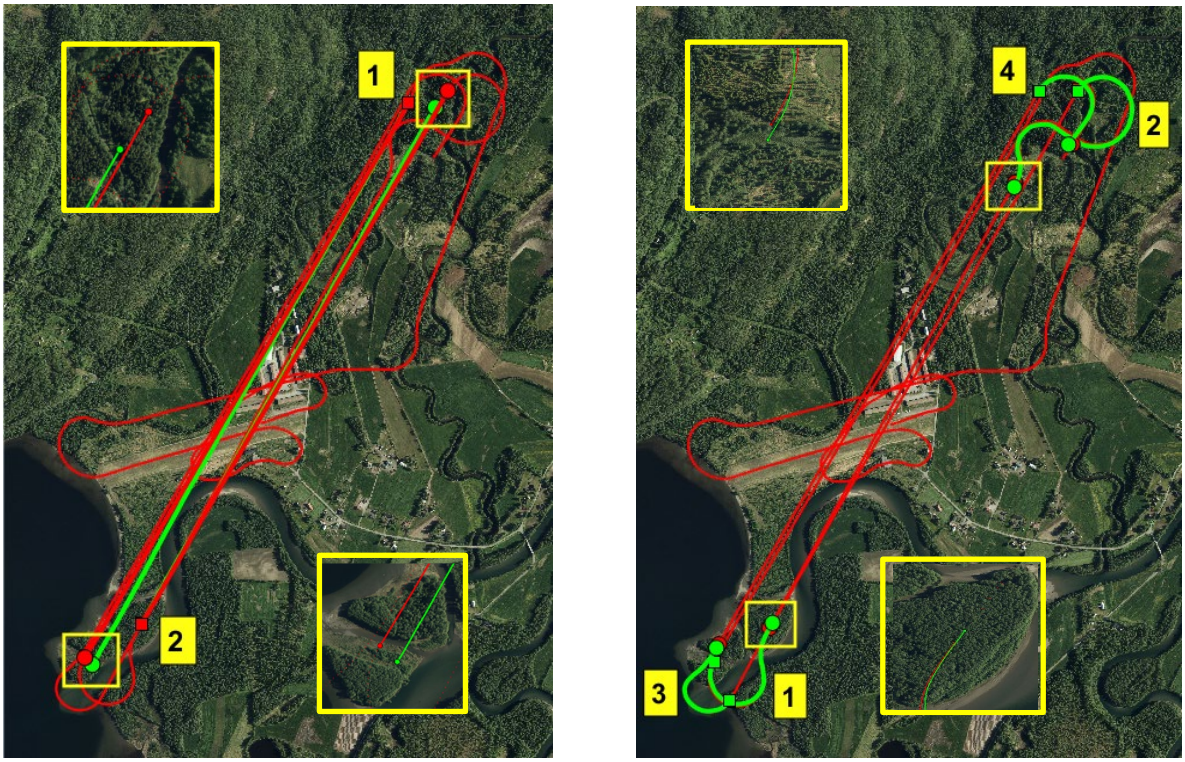

Figure 4.5 The trajectories of the linear flights (left) and banking maneuvers (right) in the fifth flight with image-aided navigation when GPS data is disabled (green line) and the reference trajectory (red line).

Duration (sec)	Distance (m)	RMS error (m)	MSE (m/s)	PTD (%)
136	2439	37.6	0.28	1.54
130	2338	82.4	0.63	3.52
122	2231	50.3	0.41	2.25
133	2411	56.7	0.43	2.35

Table 4.15 Duration of image-aided navigation (flight 5) when GPS is disabled, distance traveled within the duration and the drift in RMS error, in MSE and in PTD.

Duration (sec)	RMS error (m)	MSE (m/s)
33	18.4	0.56
53	14.1	0.27
26	18.9	0.73
24	15.0	0.63

Table 4.16 The duration of image-aided navigation (flight 5) and the drift in RMS error and in MSE under banking maneuvers.

The drift in banking maneuvers are slightly larger than the linear flights.

4.6 Image Feature Points on water surface

The image navigation algorithm integrates inertial sensor data with position data of image-features-points over an image sequence. The image-features-points need to be static and need to be present in the image sequence for a reasonable time. In all the five flights, the image-feature-points satisfied the conditions for most of the time. Especially when flying over the buildings and fields. However, it is harder to find the image-feature-points that satisfy the conditions when flying over the river and lake area. The following images shows the selected image-feature-points over the water.

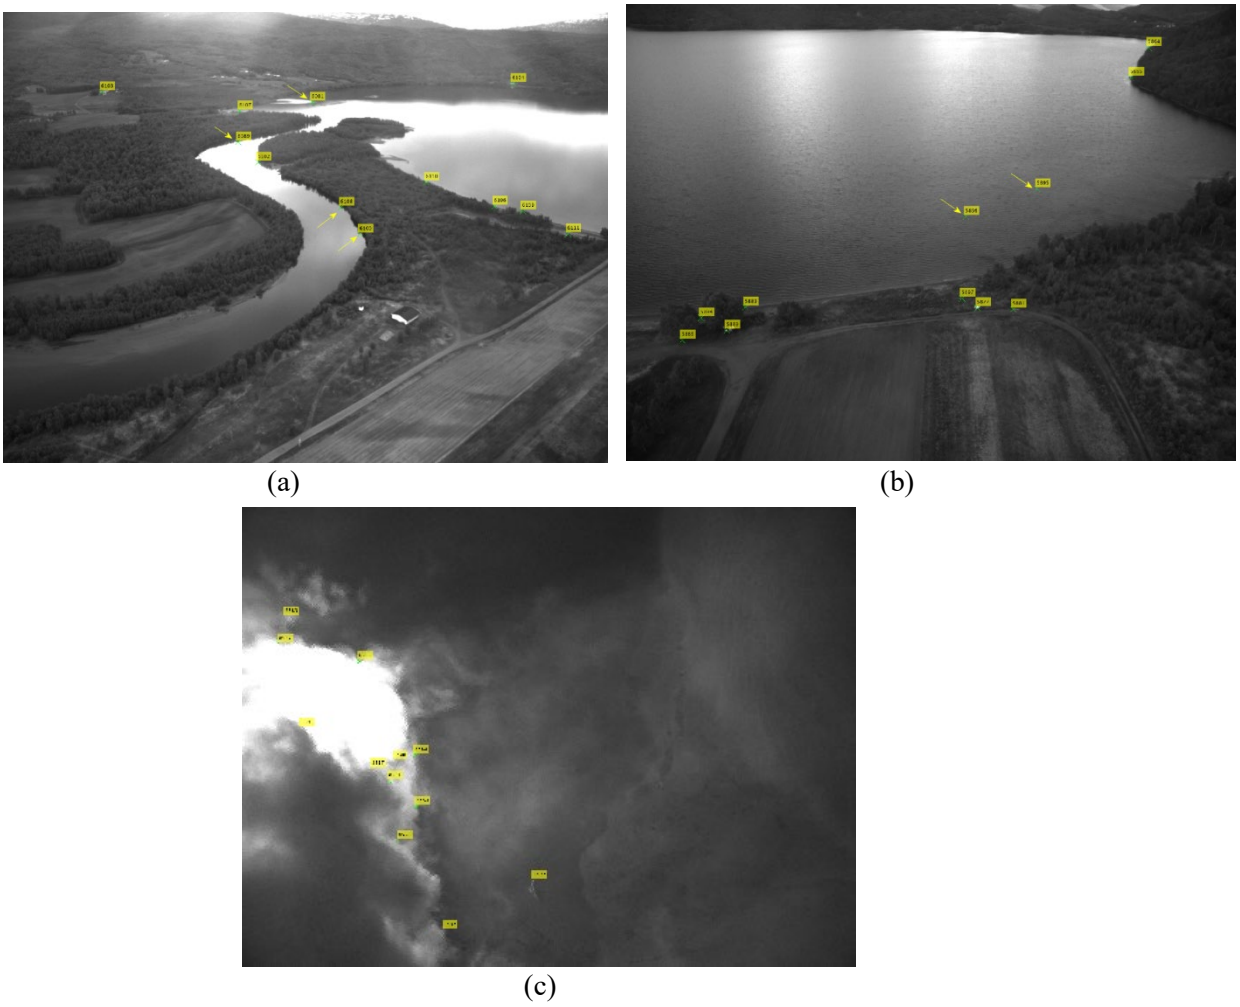

(a) (b)

(c)

Figure 4.6 Example images of image-feature-points on the river and the lake.

The first image shows the image-feature-points that are selected on the reflection of the trees on the water, which is not static. The second image shows that the image-feature-points are selected on the waves of the water. These points could be assumed static for a reasonable short time. The third image shows the points that are selected on the reflection of the clouds on the water, which are not static when the clouds are on move.

4.7 Discussion

The results of image-aided navigation in the GPS-denied time intervals, and the performance from the free inertial in the same period are given in Appendix B. It is clear that image-aided navigation gives a great overall improvement. In the linear parts the free inertial RMS error ranges from 26 to 2094 meters, while image-aided navigation RMS error ranges from 14 to 112 meters. In the banking parts, however, the results are more similar, 9.5 to 159.7 meters for free inertial, and from 6.1 to 53.3 meters for image-aided navigation. In periods of maneuvers, such as the banking turns, the free inertial drift is naturally lower, due to cancellation of errors. Over greater time intervals the free inertial error drift becomes more non-linear, and increases faster as seen in the long straight-line parts.

When comparing the navigation solution of the image-aided INS with the reference solution, the drift in banking maneuvers are overall slightly larger than the linear flights. Under banking maneuvers, due to the combination of low pitch angle and high roll angle, the image-feature-points are selected far away in the scene, see Figure 4.7. In addition, the feature-points disappeared quickly out of the field-of-view of the camera. This led to larger drift in image-aided navigation.

Figure 4.7 Image-feature-points selected far away in the scene under banking maneuver.

When a given number of image-feature-points disappear from the field of view, new feature-points are selected. This take more processing time than simply tracking the feature-points, which leads to skipping a number of images. Again, this leads to larger drift in image-aided navigation. The tables in Appendix A tabulate duration of maneuvering, how many images have been skipped, time without image-feature data in seconds and in percentage compared to duration for straight-line maneuvering and banking.

The time without image-feature data (in percentage) under straight-line maneuver is significantly lower than the banking maneuver in all the flights.

When the pitch angle is low (22˚, 25˚), the image covers larger area of the scene compare to the pitch angle of 90˚. For a given pixel resolution, a pixel of a low pitch image represents a larger area in the scene compare to a pixel of high pitch (90˚). This makes the high pitch image sharper than the low pitch image, which leads to stable image-feature-points tracking.

There seems to be slightly better comparable performance with the pitch 90° (nadir) configuration, than the grazing angle configurations. In the nadir configuration the whole image scene is in use, even during banking, the range to all the tracked landmarks are similar, and we have good initial guess from altitude. The accuracy of the range estimate is directly correlated with the accuracy of estimated speed in the INS. On the downside, the image-feature-points are not tracked over a long time.

For the grazing angle configurations, there is a certain amount of the image scene, above the horizon, that are never in use. There is also a great deal of variation in the range to the different landmarks, making it difficult to have a good initial guess. It seems that this makes it difficult to achieve good accuracy in the INS speed estimate. In this scenario, it is probably a good idea to dynamically update and estimate the initial guess from the first few frames of each image-feature-point track, before actually initializing it in the EKF, see e.g. [7] for a similar approach to this problem. On the upside, the landmarks are tracked for potentially a longer time, but during banking almost all is lost.

The processing is highly sensitive to the alignment between the IMU and the camera. The misalignment should therefore be estimated from the data from the pass over the buildings before the racetracks. This is not a problem in the VaINS box, where the camera and the IMU are usually rigidly mounted next to each other. An improvement to the current integration into the fixed-wing, would then be to also mount the IMU with steering pins on the same bracket as the camera, both outside the VaINS box in front in the payload bay. In this way, the alignment can be estimated once.

5 Summary

We have described an experiment using the Vision-aided Inertial Navigation System (VaINS). It collects sensor data, processes the data and transmits the navigation solution to an operator-terminal in real-time. The navigation solution is based on integrating inertial sensor data with position data of image-features-points over an image sequence in an error-state Extended Kalman Filter.

This is the first experiment with VaINS on a fixed-wing UAV and it was used as a data-logging box, not a real-time navigation box. This experiment was a collaboration activity between FFI and NORCE. The primary objective of this experiment was to experience the use of the VaINS-box on a fixed-wing UAV and to study the performance of the image-aided navigation.

The data were collected over two days with five flights overlooking buildings, vegetation, river, lake, and mountains. The pitch angle of the camera was varied between grazing angles 22˚, 25˚downwards and 90˚ nadir. We analyzed the performance in post-processing, where IMU and GPS data were processed with an optimal smoother to provide a reference navigation solution. Image-feature-points data together with IMU data and GPS height were used within simulated GPS denied time intervals to provide the image-aided navigation solution.

The drift between the reference solution and the image-aided solution in these time intervals was compared for all the flights, separately within the linear flights and within the banking maneuvers during turns. The flights with camera pitch angle at nadir performed overall better (less drift) than the other flights with grazing angles. The drift in the banking maneuvers are slightly larger than in the straight-line maneuvers.

During banking at grazing angles, the feature-points disappeared quickly out of the field-of-view due to the combination of low pitch angle of the camera and high roll angle of the UAV. Selecting new feature-points takes more processing time than simply tracking the feature-points, which leads to skipping of a number of consecutive images. This leads the image-aided navigation into free-inertial state and causes a larger drift.

However, comparing the results of image-aided navigation in the GPS-denied time intervals with the results from the free inertial in the same period, the image-aided navigation performed significantly better in the straight-line maneuvers. In periods of banking turns, the free inertial drift is lower, due to cancellation of errors and shorter time intervals, and therefore more similar in performance compared to image-aided navigation.

We have two main lessons learned from this experiment. Since the IMU and the camera were not rigidly mounted together, this caused an added uncertainty in their misalignment that had not be calibrated for each flight. The solution for this should be improved, maybe by mounting the IMU on the same fixture as the camera outside VaINS. In addition, we should have had a minimum of real-time feedback on the status of VaINS. This might be done by interfacing the communication that NORCE has with its ground control station.

Appendix

A Skipped image analysis

As mentioned in discussion, new feature-points selection take more processing time than simply tracking the existing feature-points. This leads to skipping a number of images. The tables in this section tabulate how many images have been skipped, time without image-feature data in seconds and in percentage compared to flying time for all the flights.

Table A.1 presents the time statistics of the straight-line maneuver and the Table A.2 presents the time statistics of the banking maneuver of the first light.

Interval # (Duration in sec.)	Number of skipped images	Time without image-feature data	
		Seconds	Percentage
1 (144)	23	2.3	1.6
2 (97)	34	3.4	3.5
3 (134)	52	5.2	3.9

Table A.1 The number of images skipped, time without image-feature data in seconds and in percentage under straight-line maneuver in the first flight.

Interval # (Duration in sec.)	Number of skipped images	Time without image-feature data	
		Seconds	Percentage
1 (30)	60	6	20
2 (36)	57	5.7	16
3 (23)	68	6.8	29

Table A.2 The number of images skipped, time without image-feature data in seconds and in percentage under banking maneuver in the first flight

Table A.3 presents the time statistics of the straight-line maneuver and the Table A.4 presents the time statistics of the banking maneuver of the second flight.

Interval # (Duration in sec.)	Number of skipped images	Time without image-feature data	
		Seconds	Percentage
1 (160)	27	2.7	1.7
2 (113)	9	0.9	0.8

Table A.3 The number of images skipped, time without image-feature data in seconds and in percentage under straight-line maneuver in the second flight

Interval # (Duration in sec.)	Number of skipped images	Time without image-feature data	
		Seconds	Percentage
1 (42)	79	7.9	18.8
2 (42)	51	5.1	12,1

Table A.4 The skipped number of images skipped, time without image-feature data in seconds and in percentage under banking maneuver in the second flight.

Table A.5 presents the time statistics of the straight-line maneuver and the Table A.6 presents the time statistics of the banking maneuver of the third flight.

Interval # (Duration in sec.)	Number of skipped images	Time without image-feature data	
		Seconds	Percentage
1 (173)	71	7.1	4.1
2 (120)	25	2.4	2.0
3 (149)	55	5.5	3.7

Table A.5 The number of images skipped, time without image-feature data in seconds and in percentage under straight-line maneuver in the third flight.

Interval # (Duration in sec.)	Number of skipped images	Time without image-feature data	
		Seconds	Percentage
1 (28)	51	5.1	18.2
2 (33)	62	6.2	18.8
3 (47)	76	7.6	16.2

Table A.6 The number of images skipped, time without image-feature data in seconds and in percentage under banking maneuver in the third flight.

Table A.7 presents the time statistics of the straight-line maneuver and the Table A.8 presents the time statistics of the banking maneuver of the fourth flight

Interval # (Duration in sec.)	Number of skipped images	Time without image-feature data	
		Seconds	Percentage
1 (173)	71	7.1	4.1
2 (120)	25	2.4	2.0
3 (149)	55	5.5	3.7

Table A.7 The number of images skipped, time without image-feature data in seconds and in percentage under straight-line maneuver in the fourth flight.

Interval # (Duration in sec.)	Number of skipped images	Time without image-feature data	
		Seconds	Percentage
1 (28)	13	1,3	4.7
2 (33)	19	1.9	5.8
3 (47)	24	2,4	5.1

Table A.8 The number of images skipped, time without image-feature data in seconds and in percentage under banking maneuver in the fourth flight.

Table A.9 presents the time statistics of the straight-line maneuver and the Table A.10 presents the time statistics of the banking maneuver of the fifth flight.

Interval # (Duration in sec.)	Number of skipped images	Time without image-feature data	
		Seconds	Percentage
1 (136)	28	2.8	2.1
2 (130)	18	1.8	1.4
3 (122)	19	1.9	1.6
4 (133)	34	3.4	2.6

Table A.9 The number of images skipped, time without image-feature data in seconds and in percentage under straight-line maneuver in the fifth flight.

Interval # (Duration in sec.)	Number of skipped images	Time without image-feature data	
		Seconds	Percentage
1 (33)	24	2.4	7.3
2 (53)	98	9.8	18.5
3 (26)	27	2.7	10.3
4 (24)	38	3.8	15.8

Table A.10 The number of images skipped, time without image-feature data in seconds and in percentage under banking maneuver in the fifth flight

B Free-Inertial performance

For comparison, we analyzed also the performance of free inertial in the same GPS-denied time intervals as used for image-aided navigation. The results of the root-mean-square (RMS) error at the end of the interval for every flight, and straight-line laps is given in Table- B.1. The corresponding results for the banking laps are given in Table- B.2.

Flight	Lap	Duration (s)	Distance (m)	Free-Inertial RMS error (m)
1	1	144	2424	1299
	2	97	2154	427
	3	134	2350	1483
2	1	160	2507	660
	2	113	2202	476
3	1	173	2555	2094
	2	120	2339	416
	3	149	2245	1346
4	1	158	2576	960
	2	123	2310	26
	3	67	1090	59
5	1	136	2439	340
	2	130	2338	855
	3	122	2231	627
	4	133	2411	1105

Table B.1 The RMS error at end of GPS-denied intervals for straight-line laps in each flight.

Flight	Lap	Duration (s)	Free-Inertial RMS error (m)
1	1	30	13.2
	2	36	28.5
	3	23	12.9
2	1	42	12.3
	2	42	17.5
3	1	28	39.9
	2	33	28.8
	3	47	29.4
4	1	30	159.7
	2	57	39.5
5	1	33	68.9
	2	53	39.2
	3	26	9.5
	4	24	10.3

Table B.2 The RMS error at the end of GPS-denied intervals for banking laps in each flight.

References

[1] Sivalingam. B, and Hagen O.K, “Image-aided inertial navigation for an Octocopter”, SPIE Defence + Commercial Sensing Symposium, “Unmanned Systems Technology”, 2018.

[2] Sivalingam. B, Hagen O.K and Skaugen A, “VaINS Technical documentation”, FFI Report 22/01877, 2022.

[3] NORCE Norwegian Research Centre, https://www.norceresearch.no

[4] Gade, K., “NAVLAB, a generic simulation and post-processing tool for navigation”, European Journal of Navigation, vol. 2(4), pp. 51-59, 2004

[5] Nvidia Developer Forum, https://forums.developer.nvidia.com/t/nano-ssh-server-disconnect-the-ssh-tunnel-after-18mn/167333, accessed 20.12.2022.

[6] N. Marriot, tmux, https://github.com/tmux, accessed 20.12.2022.

[7] Munguia, Rodrigo, and Antoni Grau. "Delayed inverse depth monocular SLAM." IFAC Proceedings Volumes 41.2 (2008): 2365-2370.

About FFI

The Norwegian Defence Research Establishment (FFI) was founded 11th of April 1946. It is organised as an administrative agency subordinate to the Ministry of Defence.

FFI's mission

FFI is the prime institution responsible for defence related research in Norway. Its principal mission is to carry out research and development to meet the requirements of the Armed Forces. FFI has the role of chief adviser to the political and military leadership. In particular, the institute shall focus on aspects of the development in science and technology that can influence our security policy or defence planning.

FFI's vision

FFI turns knowledge and ideas into an efficient defence.

FFI's characteristics

Creative, daring, broad-minded and responsible.

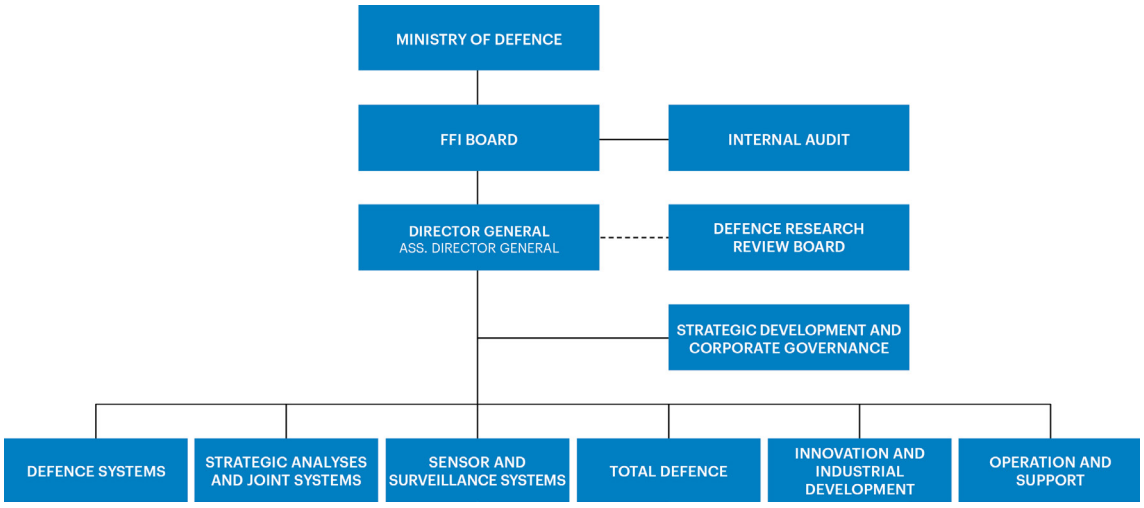

Forsvarets forskningsinstitutt (FFI)
Postboks 25
2027 Kjeller

Besøksadresse:
Kjeller: Instituttveien 20, Kjeller
Horten: Nedre vei 16, Karljohansvern, Horten

Telefon: 91 50 30 03
E-post: post@ffi.no
ffi.no

Norwegian Defence Research Establishment (FFI)
PO box 25
NO-2027 Kjeller
NORWAY

Visitor address:
Kjeller: Instituttveien 20, Kjeller
Horten: Nedre vei 16, Karljohansvern, Horten

Telephone: +47 91 50 30 03
E-mail: post@ffi.no
ffi.no/en