2014 International Conference on Unmanned Aircraft Systems (ICUAS)
May 27-30, 2014. Orlando, FL, USA

Experimental Evaluation of Four Feature Detection Methods for Close Range and Distant Airborne Targets for Unmanned Aircraft Systems Applications*

Dan Tulpan, Nabil Belacel, Fazel Famili, Kristopher Ellis

***Abstract*— Feature detection for Unmanned Aircraft Systems (UAS) sense and avoid scenarios is a crucial preliminary step for target detection. Its importance culminates when distant (pixel size) targets representing incoming aircraft are considered. This paper presents an experimental evaluation of four popular feature detection methods using flight test data and based on evaluation criteria such as first detection distance and percentage of frames with detected target features. Our results show that for close range targets all four methods have similar performance, while for distant (pixel-size) targets, the Shi and Tomasi method outperforms the other three methods (Harris-Stephens-Plessey, SUSAN and FAST).**

I. INTRODUCTION

The fundamental barrier currently facing civil applications of unmanned aircraft beyond visual range of the pilot is known as the 'sense and avoid' problem. In order to operate in unrestricted/unsegregated airspace in visual flight conditions, aircraft are required to sense and react to avoid collision with other aircraft sharing their airspace. The sense and avoid problem has been the subject of a vast amount of global activity in terms of the development of both technology to address the problem, and regulations/standards development to ensure that the technological solutions are sound and provide an equivalent level of safety to the manned 'see and avoid system'.

Several sensor technologies are available for UAS, such as visual, infrared and radar (microwave, laser and bi-static). The use of video sensors recording information in the visual spectrum for sense-and-avoid purposes is desirable since cameras represent inexpensive, passive, small and light-weight low-energy consumption sensors that perfectly fit Unmanned Aircraft (UA) platforms [1], [2], [3], [4]. A previous study [5] suggests that visual sensors, while not perfect, have a high chance of satisfying the forthcoming regulatory requirements for UASs [6], [7]. The same study concludes that a combination of complementary sensors is highly recommended to bridge the gap between safety regulations and the use of UAS in civil airspace by increasing the sensing accuracy while reducing the number of false detections.

The extraction of accurate feature points and boundaries of potential targets such as converging airborne traffic from video streams and images is a crucial first step for target detection in online and offline image processing. The success of target detection is typically dependent on knowledge related to target features such as color, shape, size and mobility and how these features differ from background information. Nevertheless, all these criteria can only be applied when areas of interest with potential targets in an image are clearly identified. Conversely, if the area around a true target is missed during the preprocessing stage of a frame, the probability of detecting the target diminishes.

Feature detection is a key concept that underlines a large number of methods applied to identification of objects in images. For sense and avoid scenarios in the unmanned/autonomous aerial systems field, identification of pixel size mobile targets is the most challenging task since national and international regulatory agencies state that any such system must behave and have the same operational abilities as manned aircraft. Thus, detection range of distant incoming aircraft, which can be implicitly transposed to distant targets, is a key problem that must be addressed.

In this paper we perform an experimental evaluation of four feature detection methods widely used in the implementation of many target detection approaches reported in the literature [8] [9] [10] [11]: Features from Accelerated Segment Test (FAST) [12], Harris-Stephens-Plessey feature detection [13], Shi and Tomasi - Minimum Eigenvalue feature detection [14], Smallest Univalue Segment Assimilating Nucleus (SUSAN) feature detection [15]. The four methods are evaluated based on their ability to detect features corresponding to close range and distant targets in frames from three video streams recorded in October 2010 using two aircrafts and a tri-focal video recording system [16]. The "first detection" distance and the percentage of frames with features matching the target are used as evaluation criteria in this study.

II. FEATURE DETECTION METHODS

A. *Previous Work*

A large number of research papers describe feature detection and target tracking methods used in research studies included in UAS applications. Wang et al. [17] present a novel real-time stabilization method for UASs. Their method consists of three steps and relies on preliminary FAST feature detection, followed by feature

*Resarch supported by National Research Council Canada.

D. T. is with the Information and Communications Technology, National Research Council, Moncton, NB E1A 7R1 Canada (phone: +1 506-861-0958; e-mail: dan.tulpan@nrc-cnrc.gc.ca).

N. B. is with the Information and Communications Technology, National Research Council, Moncton, NB E1A 7R1 Canada (phone: +1 506-861-0963; e-mail: nabil.belacel@nrc-cnrc.gc.ca).

F. F. is with the Information and Communications Technology, National Research Council, Ottawa, ON K1A 0R6 Canada (phone: +1 613-913-0116; e-mail: fazel.famili@nrc-cnrc.gc.ca).

K. E. is with the Flight Mechanics and Avionics Group, National Research Council, Ottawa, ON K1A 0R6 Canada (phone: +1 613-998-5522; e-mail: kristopher.ellis@nrc-cnrc.gc.ca).

matching in consecutive frames, motion estimation and approximation of cumulative distortions and cubic spline parameter smoothing. Their results suggest that accuracy of feature detection influences dramatically the accuracy of the transformed images and two common causes of poor image stabilization results are image blur and noise.

Forlenza et al. [9] describe an air-to-ground target detection approach implemented in hardware on an unmanned helicopter. They compared two well-known intensity-based feature detection approaches, namely Harris-Stephens-Plessey and Shi and Tomasi feature detection. Their experimental results performed on calibration images showed good performance however no real-life data was used for further testing.

Vendra et al. [10] investigated the performance of Harris-Stephens-Plessey and SUSAN feature detection approaches for the problem of aerial refueling for unmanned aerial vehicles. The methods were implemented and tested for accuracy and robustness within a simulation environment. In this context the Harris-Stephens-Plessey algorithm provided, in general, better performance compared to the SUSAN algorithm, for the detection of the same features in every frame, and in yielding a smaller number of false positives. Furthermore, a separate study comparing the robustness and the computational speed of both algorithms showed that the Harris-Stephens-Plessey algorithm provides a higher level of robustness at the expense of increased computational effort.

Previous comparisons of feature detection algorithms are reported in the literature. Chen et al. [11] performed a comparison of Harris-Stephens-Plessey and SUSAN based on metrics such as stability, noise immunity, complexity and run-time. They concluded that Harris-Stephens-Plessey was superior to SUSAN on the whole. The main disadvantages for SUSAN feature detection were the use of a fixed rather than adaptable global threshold, a fixed shape mask, and high sensitivity to noise.

Wang and Dony [8] estimated the suitability for hardware implementation of three feature detection algorithms (SUSAN, Harris/Plessey and Wang-Brady) and compared their accuracy, computational cost and stability. Among the compared algorithms, Harris/Plessey outperformed the other two in terms of accuracy and stability, at the expense of high computational complexity and poor localization. The SUSAN algorithm required the least computational resources but suffered from localization errors.

B. *Feature Detection*

While the feature detection literature is abundant in very efficient methods such as SIFT [18], SURF [19], LK [20], BRISK [21] and BRIEF [22], here we consider four feature detection methods that were previously used in sense-and-avoid situations, namely: Features from Accelerated Segment Test (FAST), Harris-Stephen-Plessey feature detection, Shi and Tomasi feature detection and the Smallest Univalue Segment Assimilating Nucleus (SUSAN) feature detection. Each method is briefly described below.

1) *Features from Accelerated Segment Test (FAST) Feature Detection Method*

FAST is a very computationally efficient feature detection method developed by Rosten and Drummond in 2003 [12]. The FAST method relies on computations based on Bresenham circles of radius 3 containing 16 pixels centered on a point of interest are depicted in *Figure 1*.

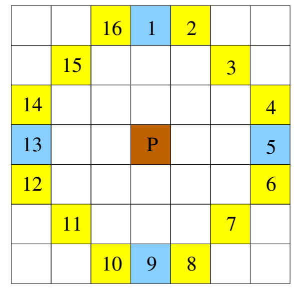

Figure 1. An interest point p surrounded by 16 pixels forming a Bresenham circle with radius 3 around pixel p. The intensities of pixels 1, 5, 9 and 13 are first compared in the FAST feature detection method against the intensity of pixel p.

The FAST algorithm functions as follows:

1. First, a pixel p with intensity I_p is selected.
2. An intensity threshold T is set and a number of pixels N ≤ 16 is chosen, such that p is considered a feature if at least N pixels will have an intensity above or below I_p.
3. Consider a Bresenham circle of radius 3 containing 16 pixels numbered clockwise centered on pixel p.
4. Compare the intensity values of 4 pixels, namely 1, 5, 9 and 13 with I_p. If at least three out of 4 pixels will have intensity above or below I_p, then the remaining pixels will be verified. Otherwise, p is not an interest point.
5. Repeat the same process for all pixels in the image.

Improved variants of the FAST algorithm based on machine learning techniques have been published since [23]. These variants solve some of the insufficiencies of the original approach such as the necessity of considering at least $N=12$ pixels to have a reduced number of false positives.

2) *Harris-Stephens-Plessey Feature Detection Method*

Rooted in the researchers need to interpret robot environments based on sequences of consecutive images, this feature detection method was developed by Harris and Stephens in 1988 [13]. The literature also refers to the same algorithm as Plessey feature detection. Throughout this paper we will refer to it as the Harris-Stephens-Plessey feature detection method. This feature detection method aims to match corresponding points in consecutive image frames and represents an improvement of the Moravec feature detection method [24] by increasing its detection and repeatability at the expense of computation time.

The Harris-Stephens-Plessey operator acts upon each pixel *(x,y)* in a gray-scale image as follows:

1. Calculate an autocorrelation matrix M for all pixels *(x,y) in image* I:

$$M = \begin{bmatrix} A & C \\ C & B \end{bmatrix}$$

Where: $A = \left(\frac{\partial I}{\partial x}\right)^2 \otimes w,\ B = \left(\frac{\partial I}{\partial y}\right)^2 \otimes w, C = \left(\frac{\partial I}{\partial x}\frac{\partial I}{\partial y}\right) \otimes w$, $\otimes$ is the convolution operator and *w* is a Gaussian window.

2. Compute the 'cornerness' measure *C(x,y)* for each pixel:

$$C(x,y) = \det(M) - k(trace(M))^2$$

Where: $\det(M) = AB - C^2 = \lambda_1\lambda_2,\ trace(M) = A + B = \lambda_1 + \lambda_2,\ \lambda_1, \lambda_2$ represent eigenvalues and $k = constant$ (typically set to 0.05).

3. Set all $C(x,y) \leq threshold$ to zero.
4. Find all local maxima in the remaining sparse matrix.
5. All non-zero points in *C* represent features.

3) Shi and Tomasi - Minimum Eigenvalue Feature Detection Method

The feature detection method developed by Shi and Tomasi [14] represents a variation of the Harris-Stephens-Plessey feature detection method and typically outperforms it. The difference between the two methods relies in the way the "cornerness measure" is defined. Shi and Tomasi suggest that only the eigenvalues must be considered, so they use the following simplified calculation:

$$C(x,y) = \min(\lambda_1, \lambda_2).$$

4) Smallest Univalue Segment Assimilating Nucleus (SUSAN) Feature Detection Method

SUSAN feature detection method was proposed in 1997 by Smith and Brady [15] and represents a significant departure from previously developed feature extraction methods. Their approach relies on non-linear filtering using circular masks. The brightness of each pixel in a mask is compared with the brightness of the mask's nucleus (central pixel) and if the two values are sufficiently similar, the area including all such pixels in the mask is labeled as a Univalue Segment Assimilating Nucleus (USAN). The presence of features of interest is marked by the area of a USAN (see Figure 2) such that the area is maximal when the nucleus of the mask lies in an image region (Figure 2-a) with uniform intensity and then successively decreases in the presence of an edge (Figure 2-b) and even more when a corner is present (Figure 2-c).

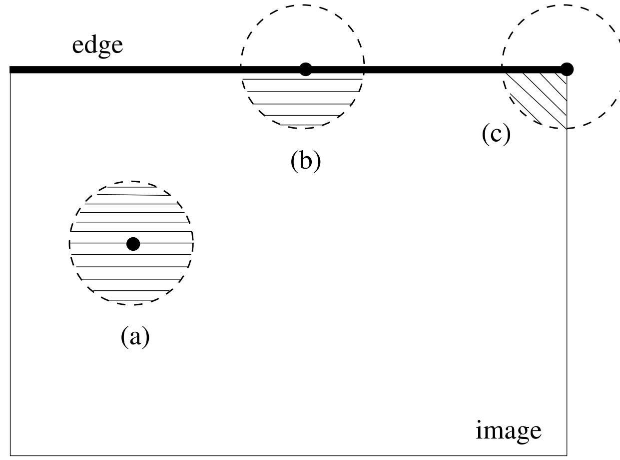

Figure 2. Three circular masks placed at different locations on a simple image.

III. Experimental Results

A. *Experimental Setup*

The experimental data was collected in October 2010. Two helicopters were used for data acquisition purposes in this project: the NRC Bell 205 Surrogate UAV, and the NRC 'intruder' aircraft (Bell 206 Jetranger).

Three Prosilica Imagers were mounted on the Bell 205. The arm affixed to the fuselage was a standard camera/gun mount designed to hold the three cameras as shown in *Figure 3*.

Figure 3. Camera mount designed to hold the 3 Prosilica Imagers on the Bell 205 aircraft.

The three Prosillica cameras were connected via Gigabit Ethernet to a pair of rack mounted PC's running the Windows XP Professional operating system. The PC's were configured to acquire the camera data stream and record uncompressed images to disk when signaled via a packet from the flight control computer. To handle the bandwidth of information, each camera's data stream was connected to a separate Gigabit Ethernet connection, and recorded to a dedicated 10,000 RPM hard disk. In addition to recording camera images, the video recording system was configured to record subtitle files which enable the individual frames to be correlated to GPS time.

The imager selected for the high resolution video recording component of this project was the Prosilica GC2450. This imager employs a 2/3 inch 5 Megapixel Kodak CCD sensor. The unit employs a global shutter, as opposed to exposing the image line-by-line, which is critical in applications where the imager is in motion as this ensures that the top of the image is temporally correlated to the bottom of the image. The maximum frame rate of the Prosilica GC2450 at full native resolution (2448 x 2050 pixels) is 15.1 frames per second.

B. Data Sets

Three video streams acquired with the experimental setup described in the previous section are used in this study (*Figure 4, Figure 5 and Figure 6*). The videos include footage of the intruder aircraft at different ranges and under different weather conditions (*Table **1***). The total number of frames in each video stream is: 2616 (VS1), 2615 (VS2) and 2616 (VS3). Frames from the video footage were selected such that each frame contains a target that can be identified by visual inspection.

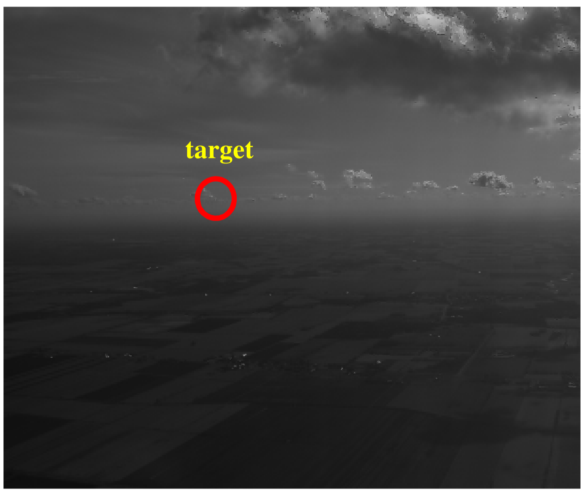

Figure 4. Frame from video stream VS1 with distant target aircraft under cloudy conditions and medium background clutter (distance to target: 0.94 km / 3084 ft).

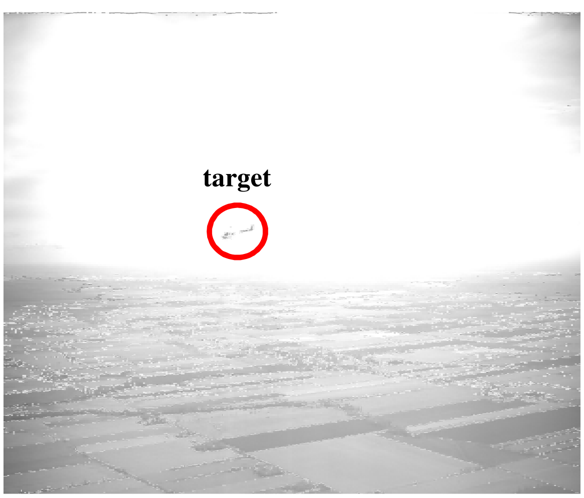

Figure 5. Frame from video stream VS2 with close range target aircraft under glare conditions and medium background clutter *(distance to target: 0.14 km / 446 ft).*

Figure 6. Frame from video stream VS3 with close range target aircraft under cloudy conditions and high background clutter *(distance to target: 0.24 km / 782 ft).*

Table 1. Data sets used in this study.

Data set	# frames with target	Target	Shooting conditions
VS1	301	Distant and medium range target	Cloudy, medium background clutter
VS2	38	Close range passing-by target	Solar glare, medium background clutter
VS3	17	Close range passing-by target	Cloudy, high background clutter

All frames containing a target (*Table **1***) were manually annotated. A ground truth data set was obtained for each video stream by recording for each frame the (x,y) image coordinates of a pixel at the center of the target. The ground truth data set is used for the evaluation of the four feature detection methods.

C. Implementation and Run Times

The four methods were implemented in Matlab R2010b and executed on a 64-bit Ubuntu 13.04 Linux (codename Raring) system with 256GB DRAM, 4 TB hard drive and two Intel Xeon 16 core 2.2 GHz CPUs. We used the Matlab - Image Processing Toolbox implementation for Harris-Stephens-Plessey and Shi and Tomasi methods, while the implementation for FAST was acquired from authors website [25] and SUSAN implementation was obtained from Matlab Central – File exchange [26]. The results reported in this paper were acquired using default initializations for Harris-Stephens-Plessey and Shi and Tomasi methods. We used a kernel size = 7 for SUSAN and a threshold of 12 pixels for the FAST method implemented in the *fast12* Matlab function.

The evaluation process was automated with Perl scripts and no parallelization was used. All methods were executed using default parameter settings. The average run-time per frame for each feature detection method is: FAST (46 seconds), Harris-Stephens-Plessey (1.6 seconds), Shi and Tomasi (1.6 seconds) and SUSAN (355 seconds). Real-time runtimes can be achieved if the methods are implemented in more computationally efficient programming languages such as C/C++ with OpenCV support or directly embedded in hardware.

D. Ratio of Frames in a Video Stream Including the Desired Target Identified by a Feature Detection Algorithm

Choosing the right feature detection method is essential for further information processing and we consider very important to use a method that detects features matching an area of interest (e.g. a target) in as many frames as possible from a given video stream during the pre-processing stage. This comes with a cost of increasing the number of false positives (detected non-target features), which can be dealt with by using additional information such as textures and saliency. Here we use the following metric to quantify this particular aspect of feature detection.

The ratio of frames in a video stream matching the desired target identified by a feature detection algorithm represents the number of frames where a specific feature detection method can help localize the target.

$$R = \frac{\#\, of\ frames\ with\ identified\ target}{total\ number\ of\ frames} * 100$$

As the ratio R increases, the utility of the feature detection method for the problem of detecting an aircraft in a sense and avoid scenario increases, too.

A radial error around each feature consisting of a fixed number of pixels (15 pixels in this study) is used for matching it with ground truth information. This is necessary since ground truth information records only the central pixel of each target.

Table 2. Ratio of frames including features localized on a desired aerial target. Radial error margin: 15 pixels.

Data set		**FAST**	**Harris-Stephens-Plessey**	**Shi and Tomasi**	**SUSAN**
VS1	First detection – frame	2225	2279	2212	2297
	First detection – distance	1868 m	1532 m	**1949 m**	1420 m
	Number of frames with detected target	195 / 301	70 / 301	276 / 301	33 / 301
	R [%]	65	23	**92**	11
VS2	First detection – frame	2510	2510	2510	2510
	First detection – distance	**105 m**	**105 m**	**105 m**	**105 m**
	Number of frames with detected target	37 / 38	38 / 38	38 / 38	37 / 38
	R [%]	97	**100**	**100**	97
VS3	First detection – frame	2549	2550	2550	2550
	First detection – distance	193 m	**199 m**	**199 m**	**199 m**
	Number of frames with detected target	17 / 17	16 / 17	16 / 17	15 / 17
	R [%]	**100**	94	94	88

Based on the information presented in Table 2, the Shi and Tomasi feature detection algorithm outperforms the other three methods both in terms of first detection distance and number/percentage of frames with detected target. For input video streams with distant targets (VS1), Shi and Tomasi is able to sense a distant target at 1.9 km, while SUSAN can only detect it the earliest at a distance of 1.4 km. The percentage of frames with detected target produced by the Shi and Tomasi method is also the highest (92%), while the second ranked method is FAST (65%) followed by Harris-Stephens-Plessey (23%) and SUSAN (11%).

For video streams with larger targets (VS2, VS3) under various filming conditions (glare for VS2 and clouds for VS3), all methods provide comparable results with statistically insignificant differences among them. The lowest detection rate (88%) was obtained with the SUSAN feature detection method applied on the VS3 video stream, which missed the target in only 2 frames (first and last) out of 17 frames. In those two frames the target was mostly occluded.

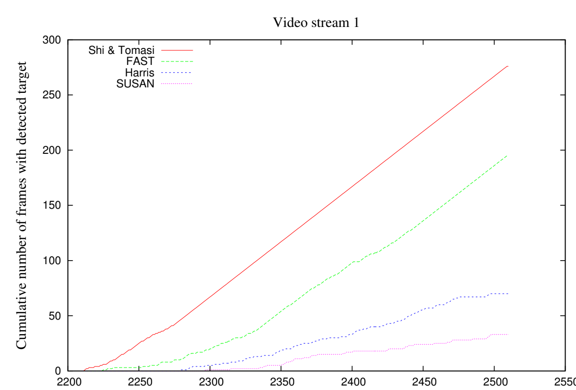

Figure 7. Cumulative distributions of target detection rates for 4 feature detection methods: FAST, Harris, SUSAN and Shi and Tomasi. Video stream 1 contains 301 frames and the target is pixel size in a majority of the frames.

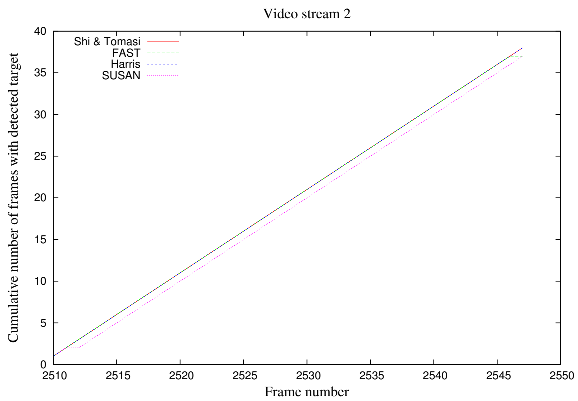

Figure 8. Cumulative distributions of target detection rates for 4 feature detection methods: FAST, Harris, SUSAN and Shi and Tomasi. Video stream 2 consists of 38 frames with a large target.

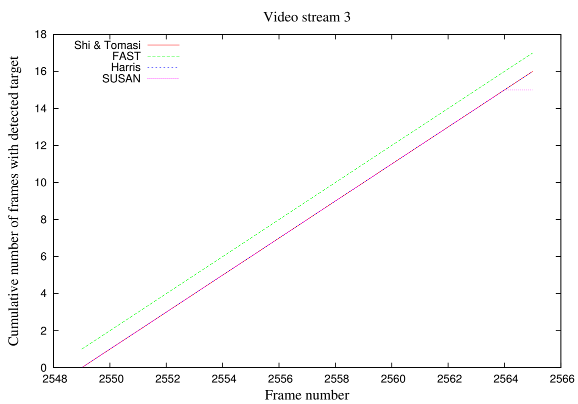

Figure 9. Cumulative distributions of target detection rates for 4 feature detection methods: FAST, Harris, SUSAN and Shi and Tomasi. Video stream 3 consists of 17 frames with a large target.

To further explore the difference in detection rates among the 4 feature detection algorithms, *Figure 7*, *Figure 8* and *Figure 9* depict the cumulative distributions of frames with detected targets. We assigned a score of 1 for each frame where a target has been marked by a detected feature while a score of 0 is assigned otherwise. Thus, if all frames would contain a detected target (all frames have score 1), a curve representing a linear function f(x) = x would be represented on the plot. Harris-Stephens-Plessey and Shi and Tomasi methods represent such examples in *Figure 8* (video stream VS2), while FAST provides the same quality for VS3 (*Figure 9*). Nevertheless, for challenging video streams including distant targets (VS1), a clear distinction can be observed between the four feature detection methods in *Figure 7*, where the Shi and Tomasi method clearly outperforms the other three methods.

IV. Conclusions

This paper presents experimental evaluation results for four widely used feature detection methods using real data of encounter scenarios captured with the aid of two aircrafts: a host and an intruder (target) aircraft. Based on this evaluation, we observed that distant (pixel size) targets pose a challenge for the investigated methods. Based on the three video streams used in this study, the Shi and Tomasi feature detection method outperforms the other three methods based on both, the detection range and the ratio of frames containing a detected target. For close range targets, we did not observe any significant difference among the four methods using the same evaluation criteria applied to distant targets. Based on the run-time analysis, Shi and Tomasi and Harris-Stephens-Plessey feature detection methods outperform FAST and SUSAN by 2 and 3 orders of magnitude, respectively. These preliminary results suggest that the Shi and Tomasi method is best suited for feature detection applied to UAS imagery and provides a good start for target detection and tracking using advanced algorithms.

Acknowledgment

The authors thank anonymous reviewers for their valuable feedback. The authors would also like to thank their colleagues at NRC for continuous support and for carefully proof reading the manuscript.

References

[1] Debadeepta Dey, Christopher Geyer, Sanjiv Singh, and Matthew Digioia, "A cascaded method to detect aircraft in video imagery," *The International Journal of Robotics Research*, vol. 30, pp. 1527-1540, 2011.

[2] John S Lai, Jason J Ford, Luis Mejias, Peter O'Shea, and Rodney A Walker, "See and avoid using on-board computer vision," in *Sense and Avoid in UAS : Research and Applications*, Plamen Angelov, Ed. Hoboken, US: John Wiley & Sons, 2012, pp. 265-294.

[3] John S Lai, Jason J Ford, Luis Mejias, and Peter J O'Shea, "Vision-based estimation of airborne target pseudobearing rate using hidden Markov model filters," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 49, no. 4, pp. 2129-2145, 2013.

[4] John Lai, Jason J Ford, Luis Mejias, and Peter O'Shea, "Characterization of sky-region morphological-temporal airborne collision detection," *Journal of Field Robotics*, vol. 30, no. 2, pp. 171-193, 2013.

[5] Brian C Karhoff, Jason I Limb, Shane W Orawsky, and Alfred D Shephard, "Eyes in the Domestic Sky: An Assessment of Sense and Avoid Technology for the Army's "Warrior" Unmanned Aerial Vehicle," in *IEEE Systems and Information Engineering Design Symposium*, 2006, pp. 36-42.

[6] (2014, April) US Federeal Aviation Administration - Unmanned Aircraft Systems (UAS). [Online]. http://www.faa.gov/about/initiatives/uas/

[7] (2010, March) Transport Canada - Unmanned Air Vehicle (UAV). [Online]. https://www.tc.gc.ca/eng/civilaviation/standards/general-recavi-brochures-uav-2270.htm

[8] Wenxin Wang and Robert D. Dony, "Evaluation of image corner detectors for hardware implementation,"

in *Canadian Conference on Electrical and Computer Engineering, 2004*, vol. 3, 2004, pp. 1285-1288.

[9] Lidia Forlenza, Patrick Carton, Domenico Accardo, Giancarmine Fasano, and Antonio Moccia, "Real Time Corner Detection for Miniaturized Electro-Optical Sensors Onboard Small Unmanned Aerial Systems," *Sensors*, vol. 12, no. 1, pp. 863-877, 2012.

[10] Soujanya Vendra et al., "Addressing corner detection issues for machine vision based UAV aerial refueling," *Machine Vision and Applications*, vol. 18, no. 5, pp. 261-273, October 2007.

[11] Jie Chen, Li-hui Zou, Juan Zhang, and Li-hua Dou, "The Comparison and Application of Corner Detection Algorithms," *Journal of multimedia*, vol. 4, no. 6, pp. 425-441, 2009.

[12] Edward Rosten and Tom Drummond, "Rapid rendering of apparent contours of implicit surfaces for realtime tracking," in *British Machine Vision Conference*, 2003, pp. 719-728.

[13] Chris Harris and Mike Stephens, "A combined corner and edge detector," in *In Proceedings of The Fourth Alvey Vision Conference*, 1988, pp. 147-151.

[14] J Shi and C Tomasi, "Good features to track," in *Proceedings Computer Vision and Pattern Recognition*, 1994, pp. 593-600.

[15] Stephen M Smith and Michael J Brady, "SUSAN—A New Approach to Low Level Image Processing," *International Journal of Computer Vision*, vol. 23, no. 1, pp. 45-78, May 1997.

[16] Jocelyn Keillor, Kris Ellis, Gregory Craig, David Rozovski, and Robert Erdos, "Studying Collision Avoidance by Nearly Colliding: A Flight Test Evaluation," in *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, vol. 55, 2011, pp. 41-45.

[17] Yue Wang, ZuJun Hou, Karianto Leman, and Richard Chang, "Real-Time Video Stabilization for Unmanned Aerial Vehicles," in *Proceedings of the MVA2011 IAPR Conference on Machine Vision Applications*, Nara, 2011, pp. 336-339.

[18] David G Lowe, "Object recognition from local scale-invariant features," *The Proceedings of the Seventh IEEE International Conference on Computer Vision*, vol. 2, pp. 1150-1157, 1999.

[19] Herbert Bay, Andreas Ess, Tinne Tuytelaars, and Luc Van Gool, "SURF: Speeded Up Robust Features," *Computer Vision and Image Understanding*, vol. 110, no. 3, pp. 346-359, 2008.

[20] Bruce D Lucas and Takeo Kanade, "An iterative image registration technique with an application to stereo vision," in *Proceedings of the DARPA Imaging Understanding Workshop*, 1981, pp. 121-130.

[21] Stefan Leutenegger, Margarita Chli, and Roland Y Siegwart, "BRISK: Binary Robust invariant scalable keypoints," in *IEEE International Conference on Computer Vision (ICCV)*, 2011, pp. 2548-2555.

[22] Michael Calonder et al., "BRIEF: Computing a Local Binary Descriptor Very Fast," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 34, no. 7, pp. 1281-1298, 2012.

[23] Edward Rosten, Reid Porter, and Tom Drummond, "Faster and better: a machine learning approach to corner detection," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 32, no. 1, pp. 105-119, 2010.

[24] Hans Moravec, "Towards Automatic Visual Obstacle Avoidance," in *Proceedings of the 5th International Joint Conference on Artificial Intelligence*, vol. 1, 1977, pp. 584-584.

[25] (2014, January) FAST Corner Detection -- Edward Rosten. [Online]. http://www.edwardrosten.com/work/fast.html

[26] (2011, March) Corner Detection using SUSAN operator. [Online]. http://www.mathworks.com/matlabcentral/fileexchange/30789-corner-detection-using-susan-operator