A Quantitative Measure for the Evaluation of Drone-Based Video Quality on a Target

Daniela Doroftei
Department of Mechanics
Royal Military Academy
Brussels, Belgium
daniela.doroftei@rma.ac.be

Geert De Cubber
Department of Mechanics
Royal Military Academy
Brussels, Belgium
geert.de.cubber@rma.ac.be

Hans De Smet
Department of Economy, Management and Leadership
Royal Military Academy
Brussels, Belgium
hans.de.smet@rma.ac.be

***Abstract*—For the evaluation of drone operators, it is important to assess their capability to produce high-quality video of a certain object. However, traditional video quality assessment methodologies are in general more geared towards video compression and thus focus on the correct image representation, and not on the real content of the produced data. In this paper, we therefore propose a methodology to define a video quality metric, specifically geared towards drone operations. Using this quantitative measure as a baseline, we also propose a methodology which proposes optimal drone trajectories for obtaining a maximum amount of qualitative video data about a given object or target in a minimum amount of time. The proposed methodologies are validated within a virtual pilot training environment.**

***Index Terms*—video production; quality assessment; drones; unmanned aircraft systems; situational awareness**

I. Introduction

Drones are used nowadays for a plenitude of tasks, including often missions where the drone operator is required to produce a high-quality video of a certain object or target. These kind of operations can range from wedding photography over search and rescue [1] or bridge inspection [2] to military operations [3]. Obviously, in order to fulfil such tasks in a proper way, the drone operator requires a specific form of training and skills development. Moreover, for mission-critical applications, it is essential to assess on beforehand that the drone operator has a sufficient level of skills with respect to this task. However, to date the quantitative evaluation of this drone pilot training expertise remains problematic, as it is not really possible to quantify the quality of the piloting skills, notably related to the capability of producing high-quality video data of a given target.

Indeed, as we will develop in Section 2 on the state of the art, there does not really exist a tool which allows to tell whether a video produced by a certain drone operator contains enough information about a certain target or not. Therefore, we propose in this paper a methodology to quantitatively assess the content of drone-based video data.

It should be noted that the methodology, which is explained in detail in Section 3 of this paper, is *not* performed by an analysis of the video signal, as this would render the approach very difficult to port from one type of application or mission scenario to another. Instead, the methodology is based on the analysis of the position data, which drones typically receive via their positioning sensors. As such, the methodology is task-agnostic and can be applied to a wide range of applications.

We do focus in this research study mostly on military operations, where aim is to gather a maximum amount of data about a target in a minimum amount of time. A drawback of this choice as an application is that the proposed approach ignores cinematographic constraints (rule of thirds, etc.) as they are commonly used for professional video photography, which makes it less useful for these kinds of applications.

In the fourth section of this paper, we show how the proposed evaluation metric can be used inside an optimization scheme in order to automatically generate drone trajectories that maximize the amount of high-quality video data obtained from a certain target.

In the fifth section of this paper, we validate the proposed methodologies in two use cases, highlighting the novel contributions of this paper:

- A methodology for content-based video-quality analysis, used for the assessment of the performance of drone operators
- A methodology for the automatic generation of optimal drone trajectories for maximizing the information gathered of a certain subject

In relation to these validation experiments, it should be noted that the quantitative drone-based video quality assessment methodology presented in this paper is not an isolated development. It is developed within the framework of the ALPHONSE project by Belgian Defence [4], which has as a goal to develop a virtual training environment for the training of drone pilots of security services (e.g., police, firefighters, civil protection, military) and to study the human factors that influence the performance of these operatives. Within the ALPHONSE training environment, drone operators perform regular virtual training missions and the goal is to track their performance related to high-quality video production with the tools presented in this paper.

 ISBN: 978-1-61208-966-9

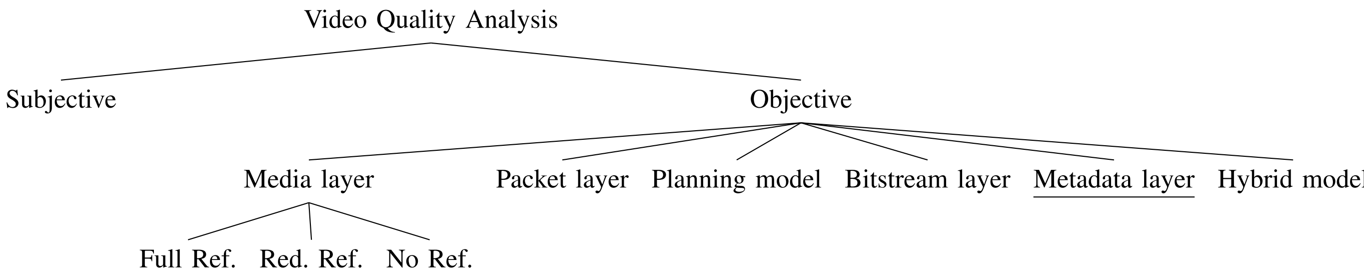

Fig. 1: Taxonomy of video quality analysis methodologies.

II. Related Work

Video quality analysis methodologies can be generically classified into two different categories. On the one hand there are the subjective video quality analysis methodologies [5], which directly express the video quality as experienced by humans. In order to do so, subjective video analysis methodologies require the video sequences to be shown to groups of viewers. The subjective opinion of the audience is recorded and averaged into a so-called Mean Opinion Score (MOS) to evaluate the quality of the video sequence. While subjective video analysis methods render excellent results, the problem is that they are extremely labour-intensive and therefore difficult to deploy in a practical context. As a result, more objective - and therefore also more automated - have been developed.

Objective video quality analysis methods have been classified by the International Telecommunications Union (ITU), according to the input data employed by the algorithms [6].

Media layer models directly use the video signal to define a quality measure. These methods do not require a priori information about the system under test and are therefore often used for the comparison of different compression methodologies. Depending on the type of source data the processed video is compared with, 3 types of submethods can be further identified:

- Full-reference methods extract data from high-quality undegraded source signals. This is the type of methodology (often a PSNR-derivative [7]), which is used for video codec evaluation.
- Reduced-reference methods extract data from a side-channel with signal parameter data.
- No-reference methods do not employ any source information.

Parametric packet-layer models estimate the video quality using only information from the signal packet header.

Parametric planning models use quality planning parameters for networks as an input in order to estimate a video quality measure.

Bitstream layer models estimate the video quality by combining encoded bitstream information with packet-layer information.

Hybrid models use a combination of the previously discussed methodologies in order to estimate a measure for the video quality.

All these traditional methodologies for video quality analysis estimate the video quality based upon the assumption that there exists a perfect input signal, wich is then degraded, due to encoding, network transmission, decoding, display constraints, etc. In our application, this is not really the case: the question we are interested in is more whether a certain subject has been perceived sufficiently within the video material. As such, it is more a content-based analysis which is required.

Video content-based planning for drones has been shown before by Hulens and Goedemé in [8]. Within this paper, an autonomous drone is presented that automatically adjusts its position in order to keep a subject (in this application an interviewee) within view under certain cinematographic constraints. In comparison with our research question, where it is the aim to maximize the information gain about generic subjects, this approach is very much geared towards one application (the subjects are always human faces).

Building upon the ITU taxonomy for objective video quality analysis methods (see Fig. 1), we therefore introduce a new category of models which are based upon a processing of the *metadata layer*. Indeed, drones typically have accurate GPS sensors on-board, enabling to geo-localize all image and video data produced by these systems. As we will develop in the following section, we propose a methodology which uses this meta-data to estimate a content-based video quality metric.

Obviously, the new metadata-layer ignores important aspects of the video quality analysis paradigm, as it does not consider any errors that may be induced in the *encoding - transmission - decoding - display* - pipeline. We acknowledge this and advocate that it is - in a realistic deployment - probably the best idea to incorporate the proposed metadata layer model together with another model in a hybrid architecture. However, in this paper, we will focus fully on the elaboration and validation of the metadata layer model by itself.

Drone trajectory optimization is a research field that has received a lot of attention in recent years, as researchers have started employing drones for a wide range of applications [9], [10]. In its essence, this problems boils down to a constrained optimisation problem, as there is an objective function (e.g., a number of targets that need to be reached) that is to be minimized, while taking into consideration the constraints imposed by the flight dynamics of the drone. This is the same in this paper, where we employ in Section 4 the proposed video quality metric as a basis for the optimization function.

III. Methodology towards the Quantitative Evaluation of Drone-based Videos

The video quality analysis methodology presented here is developed to be as task-agnostic as possible. However, this makes it necessary to define some key basic assumptions made by the algorithm.

- We assume that the drone camera is always looking straight at the target. This assumption is made in order not to over-complicate the algorithm with (dynamic) viewpoint changes in function of the drone movement. This is a realistic assumption, as in real operations, it is very often the case that there is a separate camera gimbal operator, whose task it is to point the camera at the target. Obviously, this is a task that can be automated as well using visual servoing methodologies [11]. In this paper, we assume that such a method has been implemented.
- In order to ensure that the target object is perceived equally from any different viewing angle, we assume that the target has a perfect spherical shape. This will be obviously an approximation, and for realistic objects with a very different shape, it may lead to a different evaluation. However, this is the most generic assumption which can be made, and it can - if required - be further refined if certain specific target shapes are better suited for specific applications.
- As the zoom factor is a piece of information which is not dynamically available to the algorithm, we assume here a static zoom factor.
- The sole input parameters used by the video quality assessment algorithm are the position of the drone at a certain time instance $\mathbf{x}_i = (x_i, y_i, z_i)$ and the position of the target $\mathbf{x}_t = (x_t, y_t, z_t)$, which is assumed to be static during the entire video sequence.

The methodology proposed here towards quantitative video quality analysis considers three sub-criteria, which determine together the overall measure of video quality. These three metrics are:

1) *The number of pixels on target ϕ_p.* It is well-known that for machine vision image interpretation algorithms (e.g., human detection [12], vessel detection [13]), the number of pixels on target is a key parameter to predict the success of the image interpretation algorithm. Also, for human image interpretation, the so-called Johnson's criteria [14] state clearly that the ability of human observers to perform visual tasks (detection, recognition, identification) is a function of the image resolution on a target.
The number of pixels on target is obviously correlated to the zoom factor of the camera. However, as the zoom factor is assumed to be constant, the number of pixels on target is inversely proportional to the distance between the target and the drone, such that:
$$\phi_p = \frac{\lambda}{|\overline{\mathbf{x}_i \mathbf{x}_t}|}, \tag{1}$$
where λ is a constant parameter ensuring that $0 \leqslant \phi_p \leqslant 1$. The parameter λ is dependent on the minimum distance between the drone and the target, the resolution of the camera and the focal length.
2) *The data innovation ϕ_d.* As expressed in the introduction, we want to assess the capability of drone operators to obtain a maximum amount of information about a given target in a minimum amount of time. What is therefore very important is that the operators are able to produce high-quality *new* video data of a target. The data innovation metric is there to evaluate this innovation quality.
This is performed by building up a viewpoint history memory $\boldsymbol{\theta}_j$ with $j = 1...i-1$, which contains a memory of all normalized incident angles of previous viewpoints. The current incident angle $\boldsymbol{\theta}_i$ is then compared to this memory. In practice this is done by taking the norm of the difference between the current incident angle and each of the previous incident angles. The data innovation is then equal to the smallest of these norms, as this represents the distance to the closest viewpoint on a unit sphere:
$$\phi_d = \min_{j=1}^{i-1} \left(|\boldsymbol{\theta}_i - \boldsymbol{\theta}_j| \right) \tag{2}$$
As the idea is to generate as much as possible new data, the new viewpoint $\boldsymbol{\theta}_i$ should be as far away as possible from existing viewpoints, which is expressed by (2).
3) *The trajectory smoothness ϕ_t.* In order to achieve a high-quality video, it is important that the trajectory of the drone is smooth over time. Indeed, if the drone follows an irregular motion pattern, then the resulting video signal would be hard to interpret by human operators or by machine vision algorithms. The metric ϕ_t therefore evaluates the trajectory smoothness, by building up a velocity profile $\dot{\boldsymbol{x}}_j$ with $j = 1...i-1$, which contains a memory of all velocities at previous time instances. The current drone velocity $\dot{\boldsymbol{x}}_i$ is then compared to the n most recent iterations in this velocity memory. This is done by taking the norm of the difference between the current velocity and each of the previous velocities. In order to make more recent data count more in the evaluation, this norm is weighted according to the recency of the information. The weighted and normed sum of the n most recent velocity differences is a measure for the changes in the motion profile and is thus inversely proportional to the trajectory smoothness, as expressed by (3).
$$\phi_t = \frac{1}{\sum\limits_{j=i-n}^{i-1} \frac{1}{i-n} |\dot{\boldsymbol{x}}_j - \dot{\boldsymbol{x}}_i|} \tag{3}$$

All 3 video quality subcriteria have been constructed such that they produce values between 0 and 1. According an equal importance to each of these subcriteria, the overall proposed measure for drone-based video quality can be written as:

$$\phi = |(\phi_p, \phi_d, \phi_t)| \tag{4}$$

Obviously, it is possible to attach weights to this global metric in order to prioritize one or two of the subcriteria, in function of the requirements of a given application. In this paper, we study the generic case and do not apply any of such weights.

IV. Methodology for Generating Optimal Drone Trajectories for Target Observation

As advocated in Section 2, the drone trajectory optimization problem is essentially a constrained optimization problem, where the constraints are given by the flight dynamics of the drone and where the optimization function expresses some application-specific goal. Therefore, we need to define in the first place the drone model and the application scenario.

As described in the introduction, one of the goals of this project is to develop a methodology to automatically generate drone trajectories such that a maximum amount of information can be gathered about a subject in a minimum amount of time. The application scenario is thus clearly a target observation mission.

In this research work, we assume to be dealing with a rotorcraft drone. This is a reasonable assumption, as rotorcraft are in most cases also the types of unmanned aircraft that would be used for short inspection or target observation tasks. While the execution of complex dynamic flight behaviours with rotorcraft drones requires a complex motion model and control architecture [15], this motion model can be quite well simplified for low-speed and quite static observation applications as is the case in the context of this paper. Therefore, we adopt a very simple motion model [16] for the drone to generate possible locations to move to.

Another assumption that we make is that we do not account for weather effects such as wind. Obviously, such external influencing factors can be incorporated into the system later, but here we wanted to validate in the first place the effectiveness of the proposed trajectory generation approach.

A pseudo-code representation of the general framework of the algorithm for generating drone trajectories is given in Algorithm 1. We will now explain this methodology line by line:

- *Line 2*: As stated above, the algorithm starts from a simple drone motion model, which proposes a number of possible discrete locations where the drone can move to, taking into account the flight dynamics constraints.
  In a first step, we perform a search over all possible new locations in order to assess which one is the best to move to. This means that a brute brute-force search is followed for searching for the optimal position. This is a quite simplistic approach, but we have opted for this option as the number of possible locations is not so enormous and it is therefore not required to incorporate some advanced optimization scheme.
- *Line 3*: In a second step, the safety of the proposed new drone location is assessed. This analysis considers in fact two different aspects:
  - The physical safety of the drone, which is in jeopardy if the drone comes too close to the ground. Therefore, a minimal distance from the ground will be imposed and proposed locations too close to the ground are disregarded.
  - The safety of the (stealth) observation operation, which is in jeopardy if the drone comes too close to the target, which means that the target (in a military context often an enemy) could hear / perceive the drone and the stealthiness of the operation would thus be violated. Therefore, a minimal distance between the drone and the target will be imposed and proposed locations too close to the target will be disregarded.
- *Lines 4-6*: The different sub-criteria are assessed, following equations (1), (2) and (3).
- *Line 7*: The global objective video quality measure ϕ at the newly proposed location is calculated, following the equation (4).
- *Line 8*: The point with the highest video quality score ϕ is recorded.
- *Line 9-10*: At this point, an optimal point for the drone to move to has been selected ($\boldsymbol{x}_b$). The viewpoint history memory $\boldsymbol{\theta}_j$ and the velocity history memory $\dot{\boldsymbol{x}}_j$ are updated to include this new point.
- *Line 11*: The drone is moved to the new point $\boldsymbol{x}_b$, in order to prepare for the next iteration.
- *Line 12*: The point $\boldsymbol{x}_b$ is appended to the drone trajectory profile.

Algorithm 1: Trajectory generation algorithm.

Input: drone position $\mathbf{x}_i = (x_i, y_i, z_i)$
target position $\mathbf{x}_t = (x_t, y_t, z_t)$
Output: Drone trajectory $\boldsymbol{Y}$
1 **while** *not at end of the iteration* **do**
2 $\quad \mathbf{x}_n \leftarrow CalculatePossibleNewPositions(\mathbf{x}_i)$
$\quad$ **forall** *proposed positions* $\mathbf{x}_n$ **do**
3 $\quad\quad$ **if** $EvaluateSafety(\mathbf{x_n})$ **then**
4 $\quad\quad\quad \phi_t \leftarrow CalculatePixelsOnTarget(\mathbf{x}_n, \mathbf{x}_t)$
5 $\quad\quad\quad \phi_d \leftarrow CalculateInnovation(\mathbf{x}_n, \mathbf{x}_t, \boldsymbol{\theta}_j)$
6 $\quad\quad\quad \phi_t \leftarrow CalculateSmoothness(\mathbf{x}_n, \mathbf{x}_t, \dot{\boldsymbol{x}}_j)$
7 $\quad\quad\quad \phi \leftarrow |(\phi_p, \phi_d, \phi_t)|$
8 $\quad\quad\quad \mathbf{x}_b \leftarrow RecordBestPoint(\mathbf{x}_n, \phi)$
9 $\quad \boldsymbol{\theta}_j \leftarrow UpdateDataInnovation(\boldsymbol{\theta}_j, \boldsymbol{x}_b)$
10 $\quad \dot{\boldsymbol{x}}_j \leftarrow UpdateTrajectorySmoothness(\dot{\boldsymbol{x}}_j, \boldsymbol{x}_b)$
11 $\quad \boldsymbol{x}_i \leftarrow AdvanceDrone(\boldsymbol{x}_b)$
12 $\quad \boldsymbol{Y} \leftarrow RecordDronePosition(\boldsymbol{x}_b)$

V. Results & Discussion

For the validation of the proposed methodologies, we started with the assessment of the performance of drone operators in a simulation environment [4]. Multiple simulated operators were asked to produce a high-quality video of a target within the simulation environment and the resulting total ϕ scores they obtained were recorded, as shown in Fig. 2. The figure shows that the algorithm is capable of discriminating between proficient users (e.g., number 3) and less proficient users. However, it is certainly the case that more research is still required in order to validate the relationship between the subjective quality assessment and this objective metric.

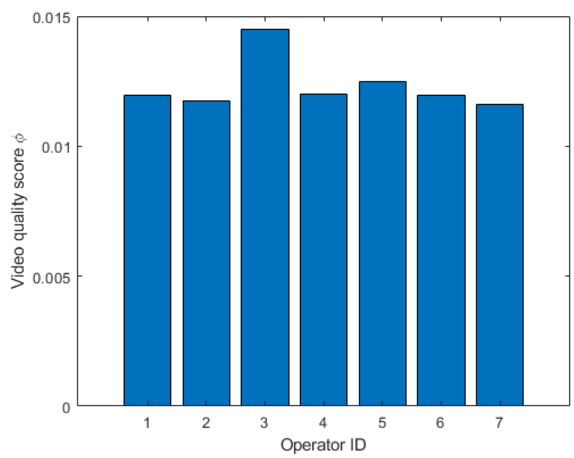

Fig. 2: Video quality scores obtained by seven operators.

In a second phase of evaluation, we validated the proposed automated drone trajectory generation, as presented in Section 3. For this validation process, we let the drone start from a random location and evaluate the optimal trajectories estimated by the algorithm.

An example of this analysis is shown in Fig. 3. In the case of the experiment depicted in Fig. 3, the drone starts from a location which is high above the target. The solution proposed by the proposed automatic trajectory generation methodology is shown in Fig. 3e, where the target position is depicted by the large sphere on the bottom. As can be noted, the proposed solution in this case consists of a spiraling downwards movement, which ensures that the target is well-perceived from all angles. Once the safety distance (from the ground and from the target) is reached, the movement pattern switches more towards an outwards extending rectangular pattern. This movement pattern is both economic for the drone and ensures that the target is perceived from ever more oblique angles.

Figs. 3a-3c show the evolution of the subcriteria ϕ_p, ϕ_d and ϕ_t during different steps of the drone trajectory. As can be noted, the algorithm achieves attaining a relatively high amount of pixels on target in the first half of the trajectory, while the drone is spiraling downwards. In the second half, the number of pixels on target decreases as the drone goes further away to obtain more oblique views.

The data innovation ϕ_d, shown on Fig. 3b shows a mostly decreasing evolution, which is due to the fact that it becomes ever more difficult to find new information.

As can be seen on Fig. 3c, the trajectory smoothness is quite constant during the majority of the trajectory, which means that the algorithm achieves in choosing smooth trajectories. Only near the end, there are peaks and valleys, which are related to the rectangular pattern where 90° turns are interspersed with straight lines.

Summing up the data innovation ϕ_d over time allows to define a so-called scan completeness measure, shown in Fig. 3d. This metric gives an idea of the amount of new data which is gathered per step in the trajectory. In all experiments we have conducted, this scan completeness metric shows an asymptotic evolution, as shown in Fig. 3d. This is also to be expected, as it becomes after some time harder and harder to obtain new data. This metric can therefore be very useful for drone operatives to evaluate in real-time whether it has sense to continue the observation task or whether it is more sensible to stop the mission.

VI. Conclusions and Future Work

In this paper, we have developed a novel metric for assessing the video quality for drone-based observation or inspection tasks. This measure is based upon an analysis of the GNSS positioning metadata embedded in the signal. The metric is essentially based upon three criteria: the number of pixels on target, the data innovation and the trajectory smoothness. The metric was embedded in an automated trajectory generation approach, which finds optimal trajectories for maximizing the amount of information perceived from a given target. The metric and the trajectory optimization methodology were validated in the framework of a drone training and simulation environment.

The validation showed that the proposed video quality metric is capable of discriminating between different levels of users. However, more research is certainly required in this domain to assess the viability of the proposed metric. As the metric is now incorporated in the drone training environment [4], it is now the idea to start larger-scale user-testing to address this issue.

An obvious shortcoming of the proposed metric, is that it only takes into account (GNSS) positioning metadata. We will therefore in a future iteration integrate this approach in a hybrid video quality analysis model, in order to come to a more comprehensive metric. Furthermore, we aim to address some of the assumptions made in this work, making the approach work also for non-spherical objects and while also accounting for weather effects.

As discussed in Section 5, the proposed trajectory generation approach succeeds in finding optimal trajectories for target observation and inspection. Moreover, having a real-time view on the scan completeness allows to know when it is the best time to end a mission. Both of these contributions of this paper

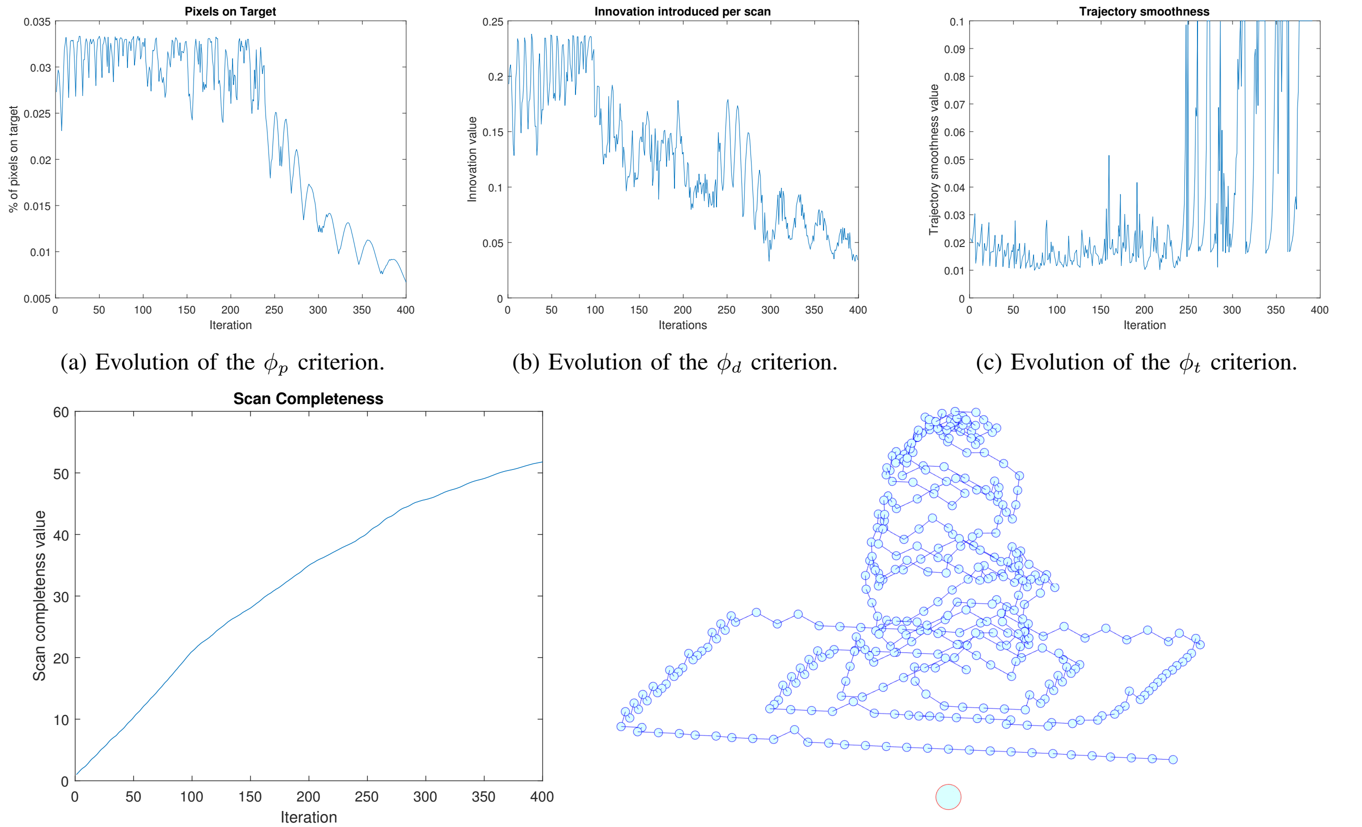

(a) Evolution of the ϕ_p criterion.

(b) Evolution of the ϕ_d criterion.

(c) Evolution of the ϕ_t criterion.

(d) Scan completeness evolution.

(e) Drone trajectory.

Fig. 3: Drone trajectory generation validation.

can be important time-savers for drone operators, or they can form the basis for automated target observation missions.

ACKNOWLEDGMENTS

This research was funded by the Belgian Royal Higher Institute for Defense and the VIAS Institute in the framework of the research study HFM19/05 (ALPHONSE).

REFERENCES

[1] G. De Cubber, D. Doroftei, K. Rudin, K. Berns, D. Serrano, J. Sanchez, S. Govindaraj, J. Bedkowski, and R. Roda, *Search and rescue robotics-from theory to practice*. InTechOpen, 2017.

[2] J. Seo, L. Duque, and J. Wacker, "Drone-enabled bridge inspection methodology and application," *Automation in Construction*, vol. 94, pp. 112 – 126, 2018.

[3] I. Lahouli, E. Karakasis, R. Haelterman, Z. Chtourou, G. De Cubber, A. Gasteratos, and R. Attia, "Hot spot method for pedestrian detection using saliency maps, discrete chebyshev moments and support vector machine," *IET Image Processing*, vol. 12, no. 7, pp. 1284–1291, 2018.

[4] D. Doroftei, G. De Cubber, and H. De Smet, "Reducing drone incidents by incorporating human factors in the drone and drone pilot accreditation process," in *11th International Conference on Applied Human Factors and Ergonomics (AHFE 2020)*, vol. 2020-July, Springer, 2020.

[5] T. Hoßfeld, C. Keimel, M. Hirth, B. Gardlo, J. Habigt, K. Diepold, and P. Tran-Gia, "Best practices for qoe crowdtesting: Qoe assessment with crowdsourcing," *IEEE Transactions on Multimedia*, vol. 16, pp. 541–558, Feb 2014.

[6] A. Takahashi, D. Hands, and V. Barriac, "Standardization activities in the itu for a qoe assessment of iptv," *IEEE Communications Magazine*, vol. 46, pp. 78–84, February 2008.

[7] S. Winkler and P. Mohandas, "The evolution of video quality measurement: From psnr to hybrid metrics," *IEEE Transactions on Broadcasting*, vol. 54, pp. 660–668, Sep. 2008.

[8] D. Hulens and T. Goedemé, "Autonomous flying cameraman with embedded person detection and tracking while applying cinematographic rules," in *2017 14th Conference on computer and robot vision (CRV 2017)*, vol. 2018-January, pp. 56–63, IEEE, 2017.

[9] P. Perazzo, F. Sorbelli, M. Conti, G. Dini, and C. M. Pinotti, "Drone path planning for secure positioning and secure position verification," *IEEE Transactions Mob. Computing*, vol. 16, no. 9, pp. 2478–2493, 2017.

[10] K. S. Yakovlev, D. A. Makarov, and E. S. Baskin, "Automatic path planning for an unmanned drone with constrained flight dynamics," *Scientific and Technical Information Processing*, vol. 42, no. 5, pp. 347–358, 2015.

[11] G. De Cubber, S. A. Berrabah, and H. Sahli, "Color-based visual servoing under varying illumination conditions," *Robotics and Autonomous Systems*, vol. 47, no. 4, pp. 225–249, 2004.

[12] G. De Cubber and G. Marton, "Human victim detection," in *Third International Workshop on Robotics for risky interventions and Environmental Surveillance-Maintenance, RISE*, 2009.

[13] J. S. Marques, A. Bernardino, G. Cruz, and M. Bento, "An algorithm for the detection of vessels in aerial images," in *2014 11th IEEE International Conference on Advanced Video and Signal Based Surveillance (AVSS)*, pp. 295–300, Aug 2014.

[14] J. Johnson, "Analysis of image forming systems," in *Image Intensifier Symposium*, p. 244–273, 1958.

[15] M. Kamel, K. Alexis, M. Achtelik, and R. Siegwart, "Fast nonlinear model predictive control for multicopter attitude tracking on so(3)," in *2015 IEEE Conference on Control Applications (CCA)*, pp. 1160–1166, Sep. 2015.

[16] C. Powers, D. Mellinger, and V. Kumar, "Quadrotor kinematics and dynamics," in *Handbook of Unmanned Aerial Vehicles*, pp. 307–328, Springer, 2014.