January 2020

Airship-based security and monitoring with machine learning

Yan Mayster

Brian Shucker

Airship-based security and monitoring with machine learning

ABSTRACT

Large open spaces, e.g., beaches, cattle-grazing grounds, etc., often need to be monitored for security or other reasons. In many cases, monitoring from the air provides the widest and most effective coverage. However, due to the cost of air-based monitoring, such aerial monitoring is often not possible, putting individuals or assets at risk or causing loss of revenue. This disclosure describes techniques to deploy small aircraft, e.g., kites, blimps, low-powered drones, etc. to monitor a given area. A video feed from the aircraft is analyzed using a trained machine-learning model to automatically detect situations that meet safety or compliance criteria, e.g., a swimmer at a beach getting far from the coast, cattle straying into forbidden territory, etc., and to automatically provide alerts.

KEYWORDS

- Aerial monitoring
- Aerial imagery
- Semantic segmentation
- Instance segmentation
- Drone
- Blimp
- Video understanding
- Image understanding
- Machine learning

BACKGROUND

Large open spaces, e.g., private or public beaches, cattle-grazing farm grounds, fenced or unfenced areas restricted to authorized individuals, etc. need to be monitored for security and other reasons. In many cases, monitoring from the air provides the widest and most effective coverage. However, due to the cost of air-based monitoring or lack of cheap technological solutions, such spaces may often have no aerial monitoring. In many cases, this can put individuals or assets at risk, or cause lost revenue, e.g., due to security fears of potential customers. Recent incidents, e.g., of a shark that was fortunately spotted as it approached a popular swimming area, underscore the importance of airborne monitoring. However, airborne monitoring is most effective when continuous, or at least periodic, and automated, e.g., without the need for a human monitor.

DESCRIPTION

This disclosure describes techniques to deploy airborne craft that are more or less stationary and require minimal or no power, e.g., kites, blimps, a rotating set of drones or copters, etc. at a height sufficient for on-board cameras to cover the area to be monitored. A multi-type, machine-learned, instance-segmentation model is trained on aerial imagery data relevant to the deployment and can detect situations that trigger alerts, e.g., for human action. Training data can be gathered by deploying the aircraft for a period of time prior to field deployment to accumulate sufficient training-data footage.

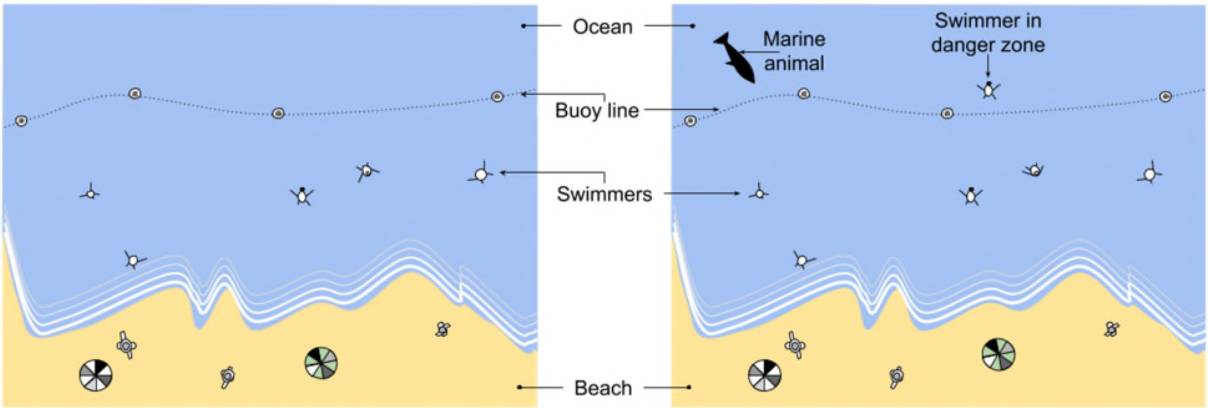

Fig. 1: (a) Negative; and (b) Positive training images for a beach swimming area use case

For specific use cases, a list of common feature types is identified. For example, in the case of beach swimming areas illustrated in Fig. 1, negative training images, e.g., images that do not trigger alarm, comprise images such as those shown in Fig. 1(a), where all humans are on the shore or close to the coast (or within the buoy line) and there are no dangerous marine animals or watercraft near the swimmers. Positive training images, e.g., images that do trigger an alarm, comprise images such as those shown in Fig. 1(b), where one or more swimmers are too far from the coast or in a known danger zone (potentially signaling that the swimmers may be caught in a rip current), a dangerous marine animal such as a shark is detected, etc.

To achieve a high detection rate at a low false-alarm rate, the set of training images can include training images that are difficult to classify. For example, in the beach use case, the machine-learned model can be trained to distinguish between dangerous marine animals like sharks from other marine animals like dolphins. The precision/recall of the model can be selected such that the false-alarm rate is as low as possible while accurately identifying situations where an alert is to be generated.

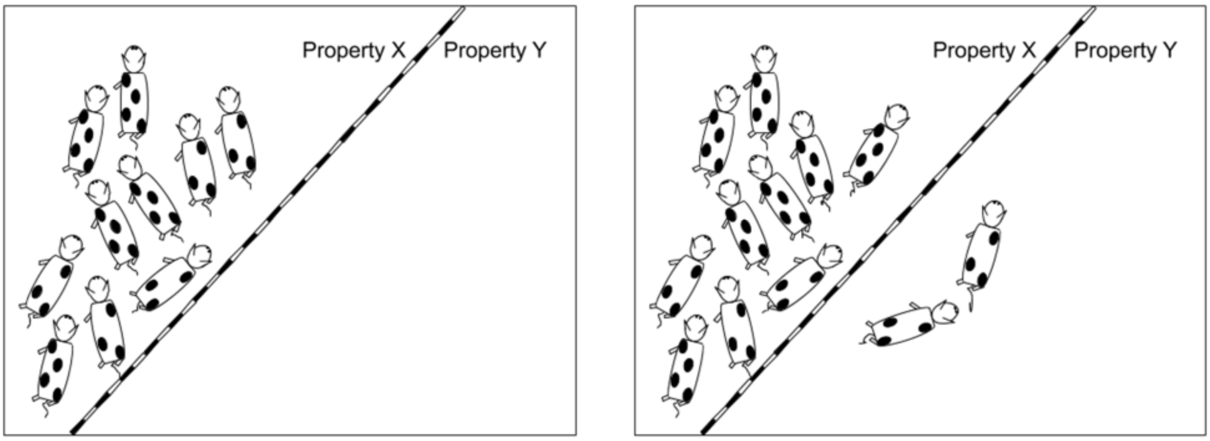

Fig. 2: (a) Negative; and (b) Positive training images for tracking cattle

Another example use case, illustrated in Fig. 2, is the tracking of cattle. In this example, all the cattle are under normal circumstances to remain on the left side of the fence, e.g., in Property X. Fig. 2(a) represents one of a set of negative training images, e.g., images that do not trigger an alarm. Fig. 2(b), in which some of the cattle have strayed to the forbidden Property Y, represents one of a set of positive training images, e.g., images that trigger an alarm.

The described techniques can be used for other situations that benefit from airborne monitoring, such as detection and location to high precision of trespassers, search-and-rescue missions, monitoring of geographically widespread assets, vehicles, etc. Airborne craft and cameras can be deployed in compliance with local regulations and with individuals in the area being notified, e.g., via prominent signboards, that the area is being monitored.

Processing of the aerially acquired imagery can be done either in a computing device onboard the aircraft or by sending the imagery to a server. If a situation that meets the criteria to raise an alert is detected, information is relayed to a human agent. A variety of alert options can be utilized. For example, in public areas such as a beach, a special siren or voice recording can be played via loudspeaker, or information can be posted on a public LED display. The posted information can be detailed, or can be as simple as a color-coded alert based on the severity of the situation. In cases requiring rescue or recovery (swimmer swept out to sea, stray cow roaming onto someone else's land, etc.), the location of the detected instance is determined based on the imagery and relayed to the human agent or other help dispatched. The detected instance can be kept in continuous focus.

Information relating to the overall number of incidents in the monitored area can be posted publicly to help people better gauge their safety (e.g., on the beach), expenses (e.g., in the cattle ranching scenario), to determine the need for additional security measures or investments, etc. Airborne monitoring can be combined with other strategies such as human patrolling, safety signs, fencing, land or sea-based cameras, etc. such that coverage of the area is sufficient for the type of area and activities.

CONCLUSION

This disclosure describes techniques to deploy small aircraft, e.g., kites, blimps, low-powered drones, etc. to monitor a large area. A video feed from the aircraft is analyzed using a trained machine-learning model to automatically detect situations that meet safety or compliance criteria, e.g., a swimmer at a beach getting far from the coast, cattle straying into forbidden territory, etc., and to automatically provide alerts.

REFERENCES

[1] Ashley Boucher, "Shark lurks near children swimming at Florida beach in chilling drone images captured by their dad" https://people.com/human-interest/shark-kids-florida-beach-drone-images/ Accessed Jan. 3, 2020.