Imaging for drone safety

Contributors: M. Tzelepi, E. Kakaletsis, A. Tefas, N. Nikolaidis
(Aristotle University of Thessaloniki, Greece)

Presenter: Prof. Ioannis Pitas
Aristotle University of Thessaloniki
pitas@aiia.csd.auth.gr
www.multidrone.eu
Presentation version 1.0

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 731667 (MULTIDRONE)

Imaging for drone safety

- **Human crowd detection for safe autonomous drones**
- Emergency landing site detection.

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 731667 (MULTIDRONE)

Human crowd detection for safe autonomous drones

MultiDrone

- Detect where crowd exists.
  - Comply with legislation.
  - Detect emergency landing points.
  - Provide heatmaps of the estimated probability of crowd presence in each location.

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 731667 (MULTIDRONE)

Human crowd detection for safe autonomous drones

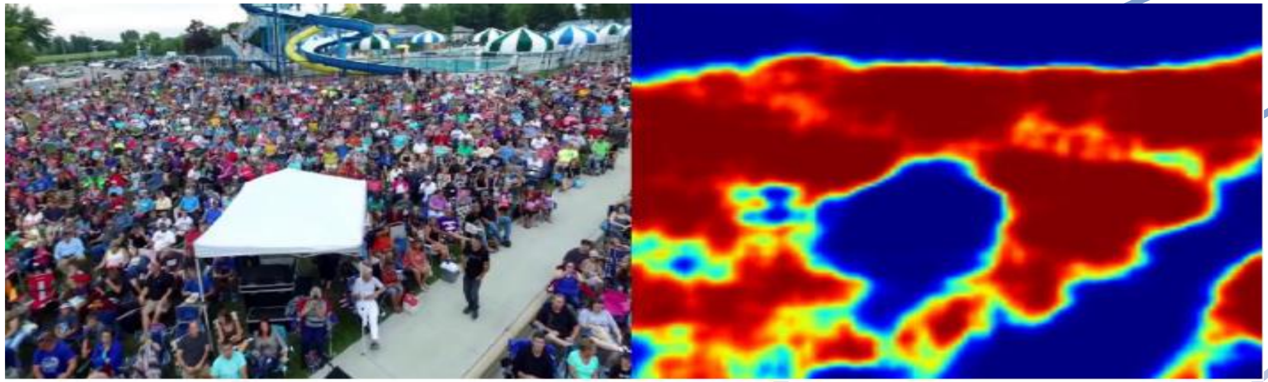

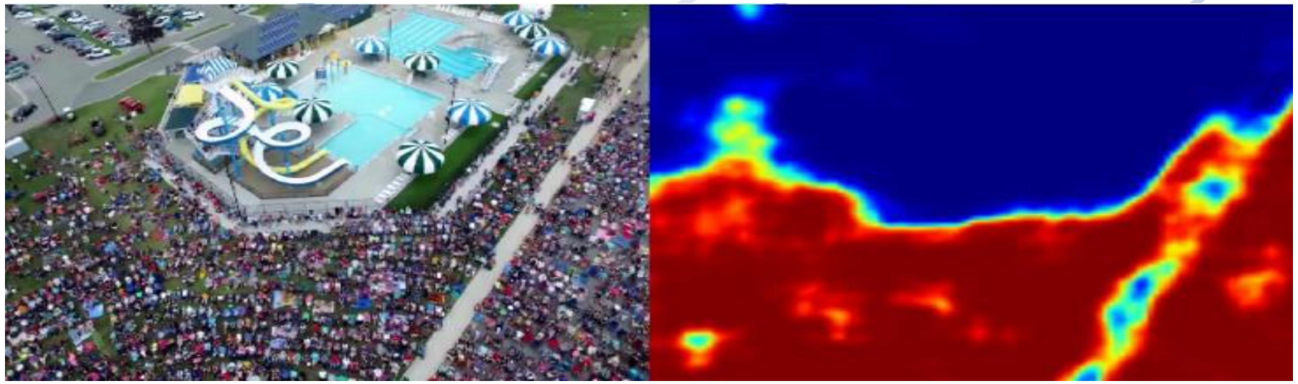

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 731667 (MULTIDRONE)

Human crowd detection for safe autonomous drones

MultiDrone

- Limited previous efforts on crowd detection, using computer vision techniques.
- Crowded scenes are considered in related research works involving crowds, e.g.,:
  - crowd understanding,
  - crowd counting,
  - human detection and tracking in crowds.

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 731667 (MULTIDRONE)

Human crowd detection for safe autonomous drones

- State-of-the art approaches on crowd analysis utilize deep learning techniques:
  - In [1] an effective Multi-column Convolutional Neural Network architecture is proposed to map the image to its crowd density map.
  - In [2] a switching convolutional neural network for crowd counting is proposed, aiming to leverage the variation of crowd density within an image.

[1] Zhang, Yingying, et al. "*Single-image crowd counting via multi-column convolutional neural network.*“, Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2016.

[2] Sam, Deepak Babu, Shiv Surya, and R. Venkatesh Babu. "*Switching convolutional neural network for crowd counting.*“, Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, Vol.1, No.3, 2017.

This project has received funding from the European Union’s Horizon 2020 research and innovation programme under grant agreement No 731667 (MULTIDRONE)

Human crowd detection for safe autonomous drones

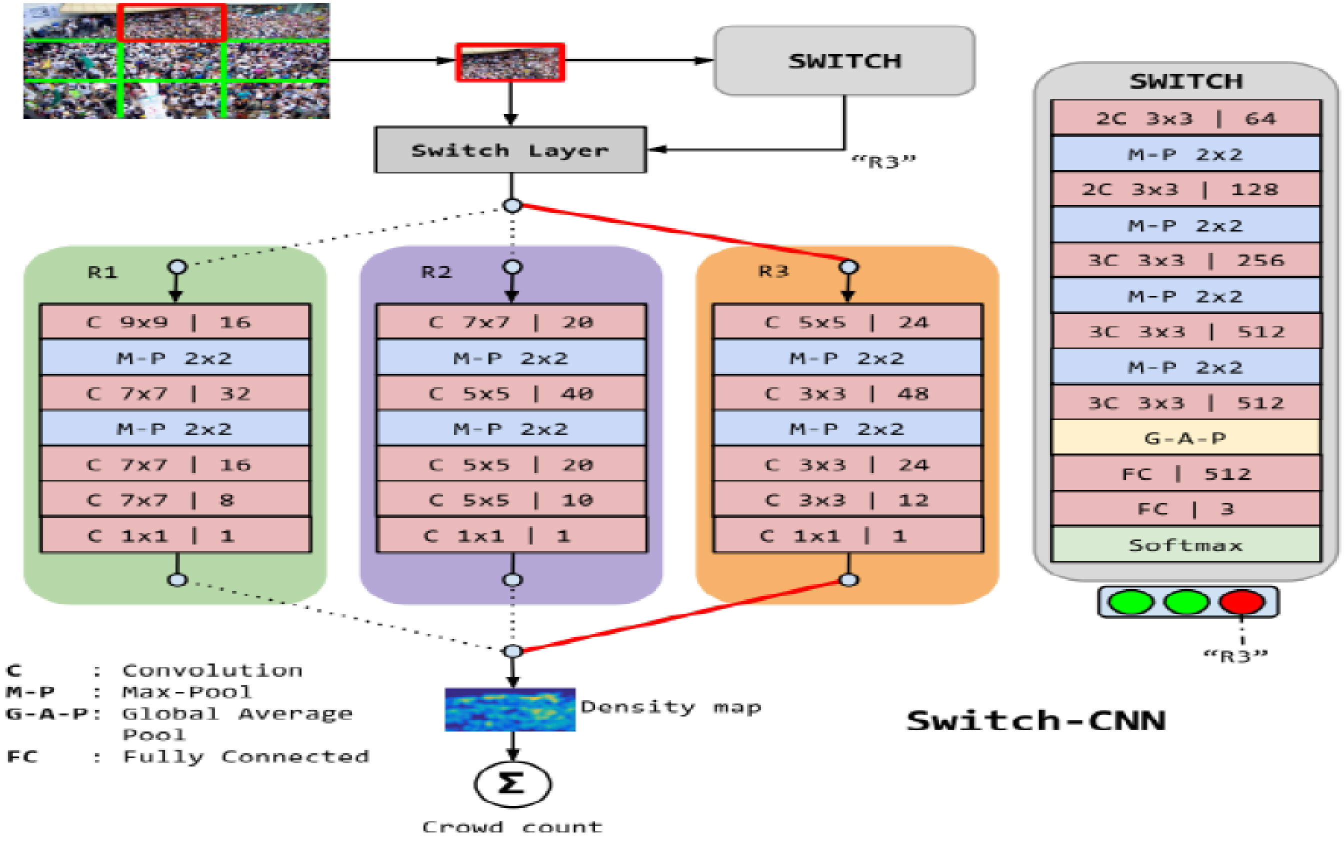

Switch CNN: Sam, Deepak Babu, Shiv Surya, and R. Venkatesh Babu. "*Switching convolutional neural network for crowd counting.*", Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, Vol.1, No.3, 2017 Crowd

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 731667 (MULTIDRONE)

Human crowd detection for drone flight safety using CNNs

MultiDrone

- In [1], a method utilizing Convolutional Neural Networks (CNNs) for crowd detection is proposed.
  - Two approaches:
    - transforming a pre-trained CNN to a fast, fully-convolutional network,
    - devising a two-loss-training model, enhancing the separability of the crowd and non-crowd classes.

[1] Tzelepi, Maria, and Anastasios Tefas, "*Human Crowd Detection for Drone Flight Safety Using Convolutional Neural Networks.*" in European Signal Processing Conference (EUSIPCO), Kos, Greece, 2017.

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 731667 (MULTIDRONE)

Human crowd detection for drone flight safety using CNNs

MultiDrone

- Provide lightweight models, as imposed by the computational restrictions of the application.
- Effectively distinguish between crowded and non-crowded scenes.
- Provide crowd heatmaps to semantically enhance flight maps by defining no-fly zones.

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 731667 (MULTIDRONE)

Human crowd heatmaps

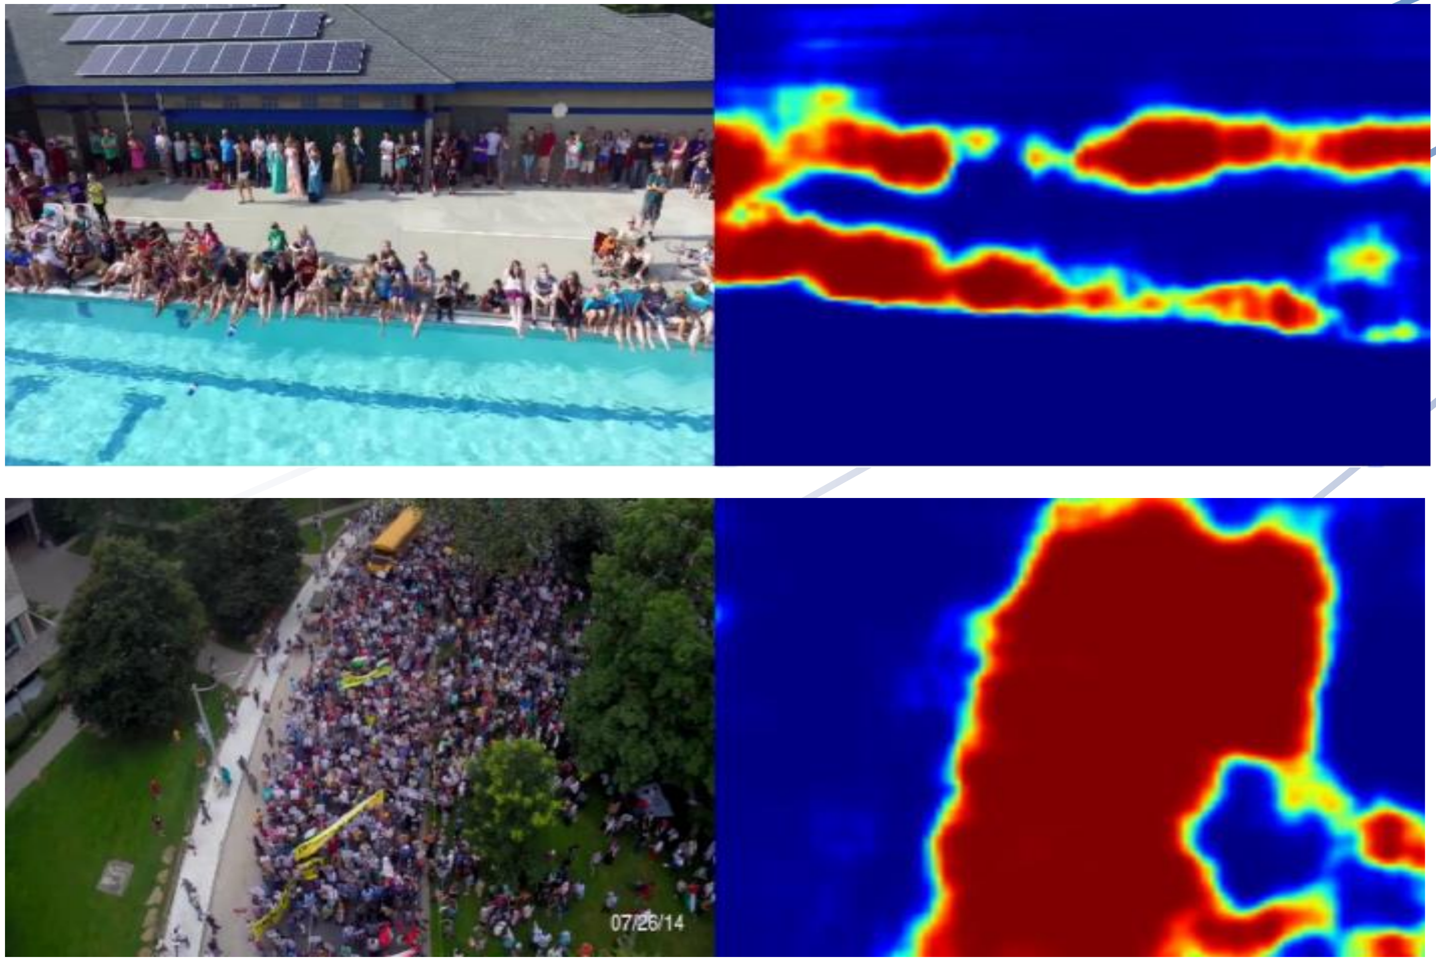

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 731667 (MULTIDRONE)

Human crowd heatmaps

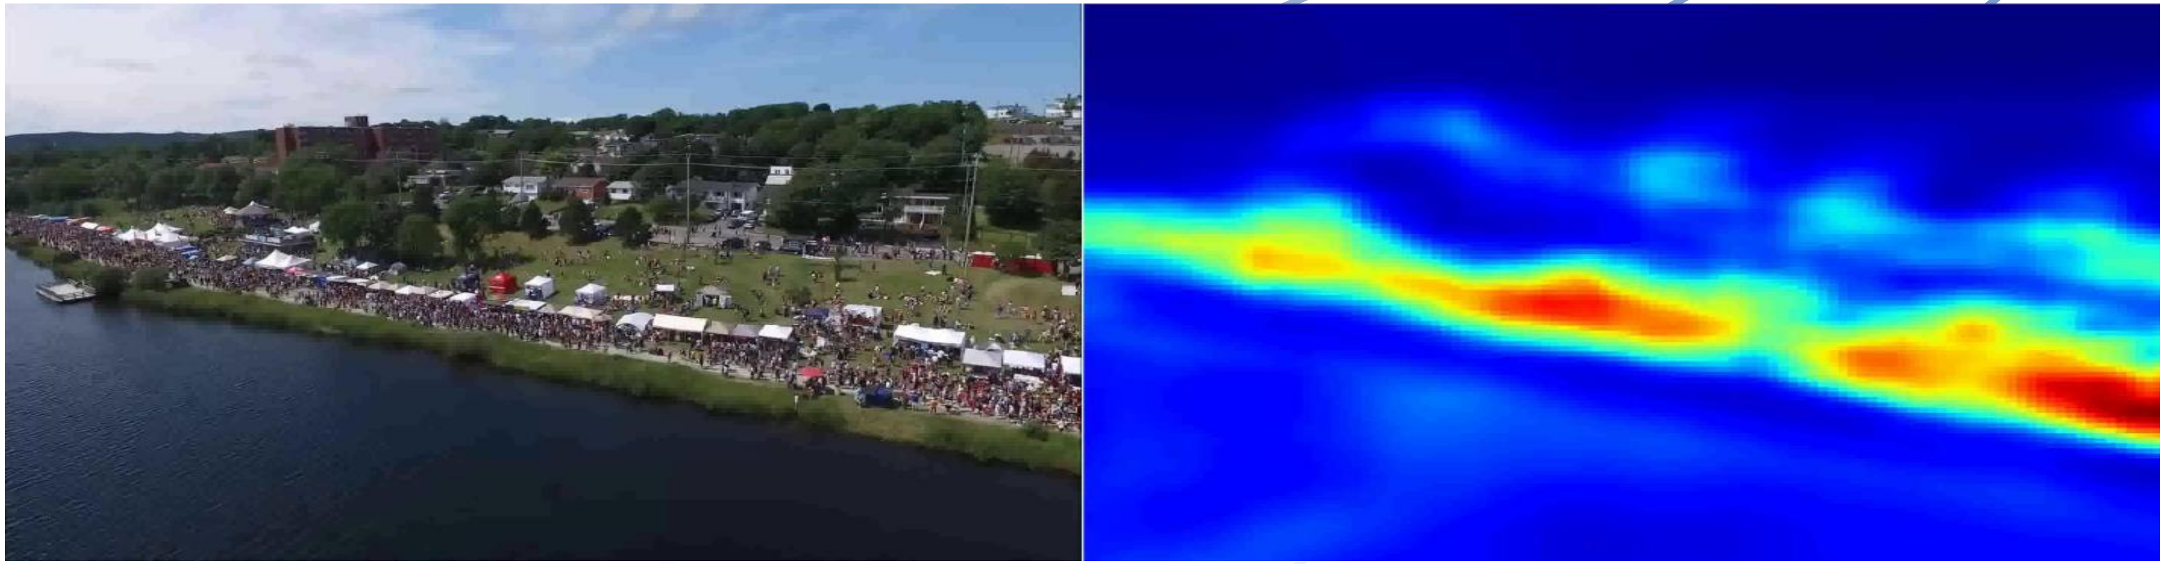

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 731667 (MULTIDRONE)

One-loss convolutional model

- Pre-trained model: BVLC Reference CaffeNet.
  - trained on the ImageNet Large Scale Visual Recognition Challenge (ILSVRC) 2012 to classify 1M+ images to 1000 ImageNet classes,
- New model *One-Loss Convolutional*:
  - discarding the fully-connected portion of the network,
  - attaching an extra convolutional layer, *CONV6,* with receptive field equal to the whole input,
  - using softmax loss during training.

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 731667 (MULTIDRONE)

One-loss convolutional model

- Proposed modification:
  - Drastically reduces the amount of the model parameters.
  - Restricts the computational cost.
  - Allows arbitrary-sized input images, since the fixed-length input requirement concerns the fully-connected layers.

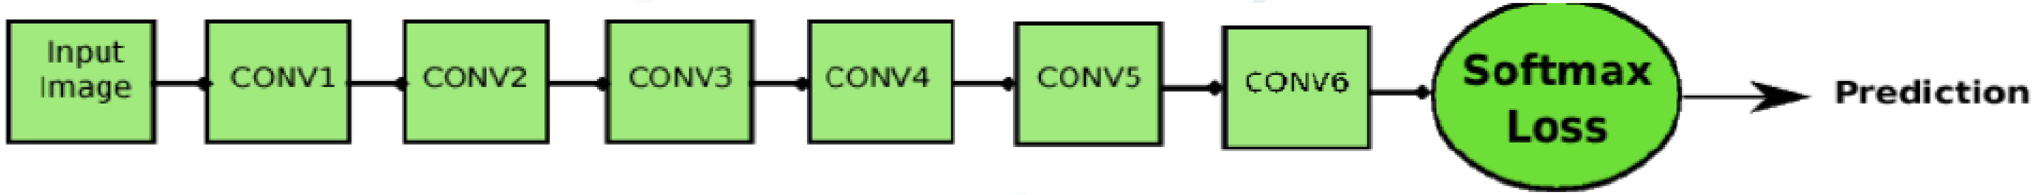

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 731667 (MULTIDRONE)

Human crowd heatmaps

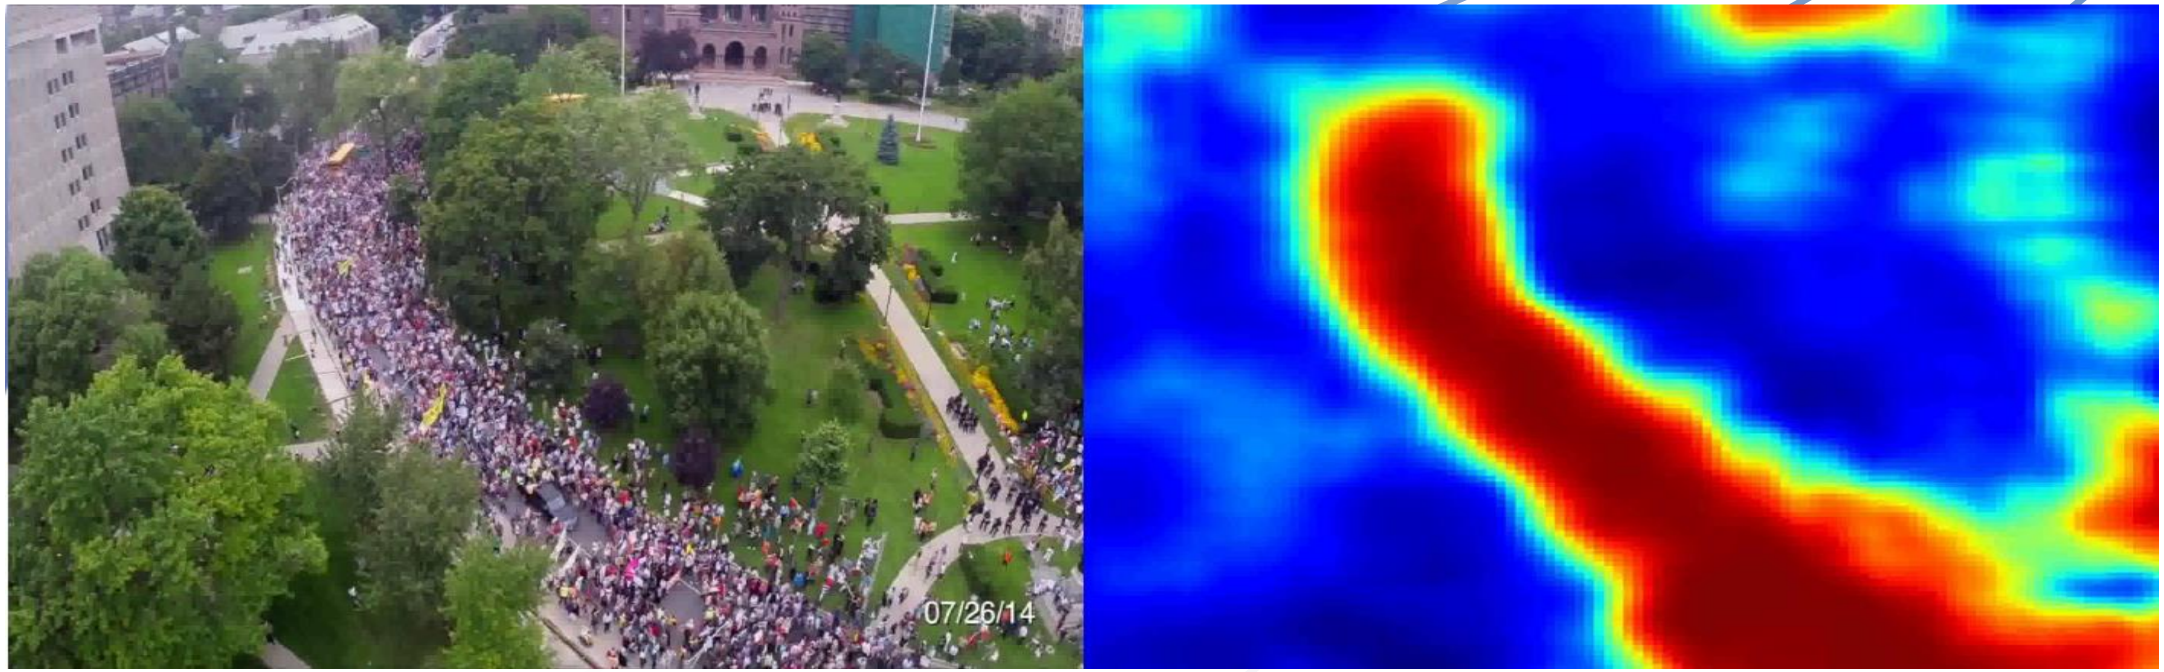

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 731667 (MULTIDRONE)

Human crowd heatmaps

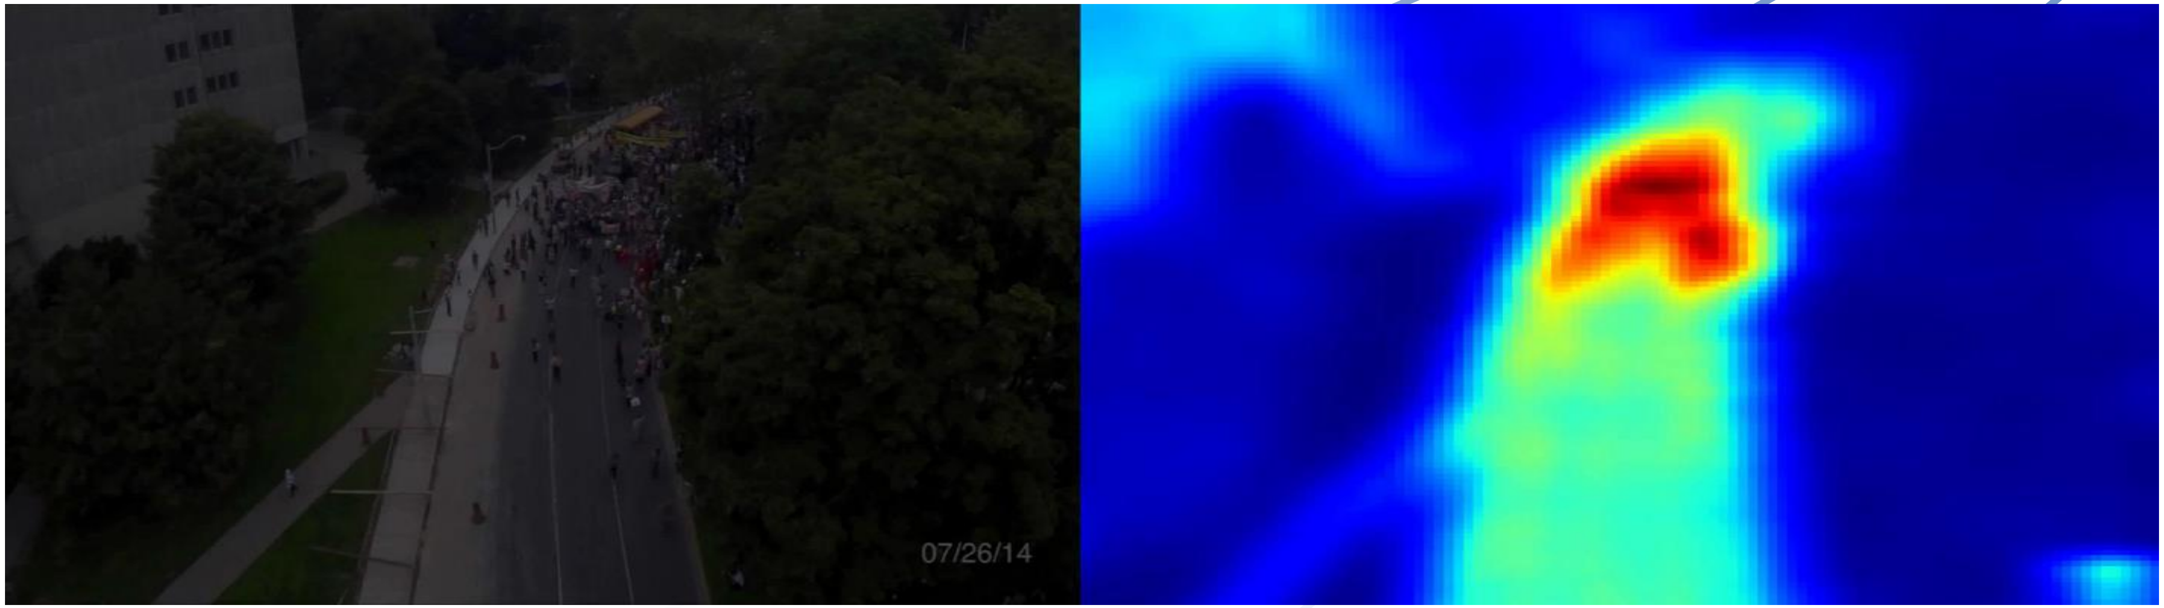

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 731667 (MULTIDRONE)

Two-loss convolutional model

- Inspired by the Linear Discriminant Analysis (LDA) method:
  - aims at best separating samples of different classes, projecting them into a new low-dimensional space,
  - maximizes between-class separability,
  - minimizing within-class variability.

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 731667 (MULTIDRONE)

Two-loss convolutional model

MultiDrone

- Proposed model employs:
  - two loss layers:
    - softmax: preserving the between class separability,
    - Euclidean: aiming at bringing the samples of the same class closer to each other.
  - a pooling layer:
    - MAC_5 : implementing the Maximum Activations of Convolutions (MAC), over the width and height of the output volume.

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 731667 (MULTIDRONE)

Dataset Construction

- Construction of dataset querying Youtube search engine with keywords describing crowded events (e.g., parade, festival, marathon, protests, political rally, etc.):
  - 60 drone crowded videos selected,
  - non-crowded videos also gathered.

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 731667 (MULTIDRONE)

Dataset Construction

MultiDrone

- Entire video sequences were left out of the training set, with their corresponding extracted frames formulating the test set.

	Train	**Test**
Crowd	2184	727
Non Crowd	1914	429
Total	4098	1156

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 731667 (MULTIDRONE)

Evaluation

- Adopted performance metric: Classification Accuracy.
- Heatmaps for the class Crowd of the proposed classifier were provided.
- Test images of size 1024×1024 were fed to the network and the output for label “Crowd”, which is the desired heatmap, was computed at the layer CONV6.

This project has received funding from the European Union’s Horizon 2020 research and innovation programme under grant agreement No 731667 (MULTIDRONE)

Experimental Results

Training Approach	Parameters	Layers	Accuracy
CaffeNet	61M	8	0.9299
One-Loss Convolutional	2.3M	6	0.91
Two-Loss Convolutional	2.3M	6	**0.9532**

- One-Loss-Convolutional model: slightly worse than the refined CaffeNet model, with a drastic reduction of model parameters.
- Two-Loss-Convolutional training procedure: considerably improved performance against the baselines, with a significantly lighter architecture.

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 731667 (MULTIDRONE)

Imaging for drone safety

MultiDrone

- Human crowd detection for safe autonomous drones
- **Emergency landing site detection.**

This project has received funding from the European Union’s Horizon 2020 research and innovation programme under grant agreement No 731667 (MULTIDRONE)

Mapping for UAVs

- Maps play a crucial role in UAV navigation:
  - they can include terrain information to be used for the navigation and control of the UAV in normal and emergency situations in the form of Digital Elevation Models (DEM):
    - Digital Terrain Models (DTM) include information regarding the height variations of an area’s bare ground without any man-made structures or vegetation.
    - Digital Surface Models (DSM) include information for both the ground and the manmade structures or vegetation lying on it.

This project has received funding from the European Union’s Horizon 2020 research and innovation programme under grant agreement No 731667 (MULTIDRONE)

Mapping for UAVs

DTM

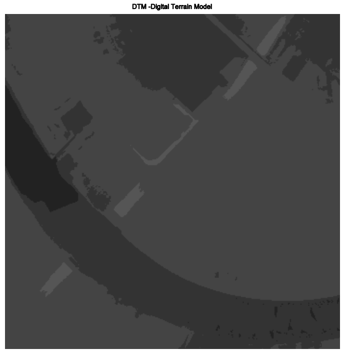

DSM

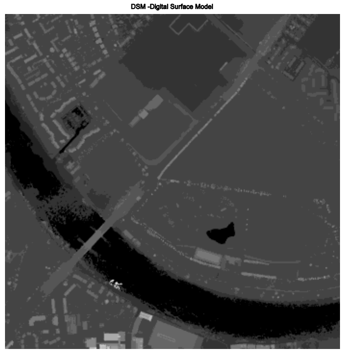

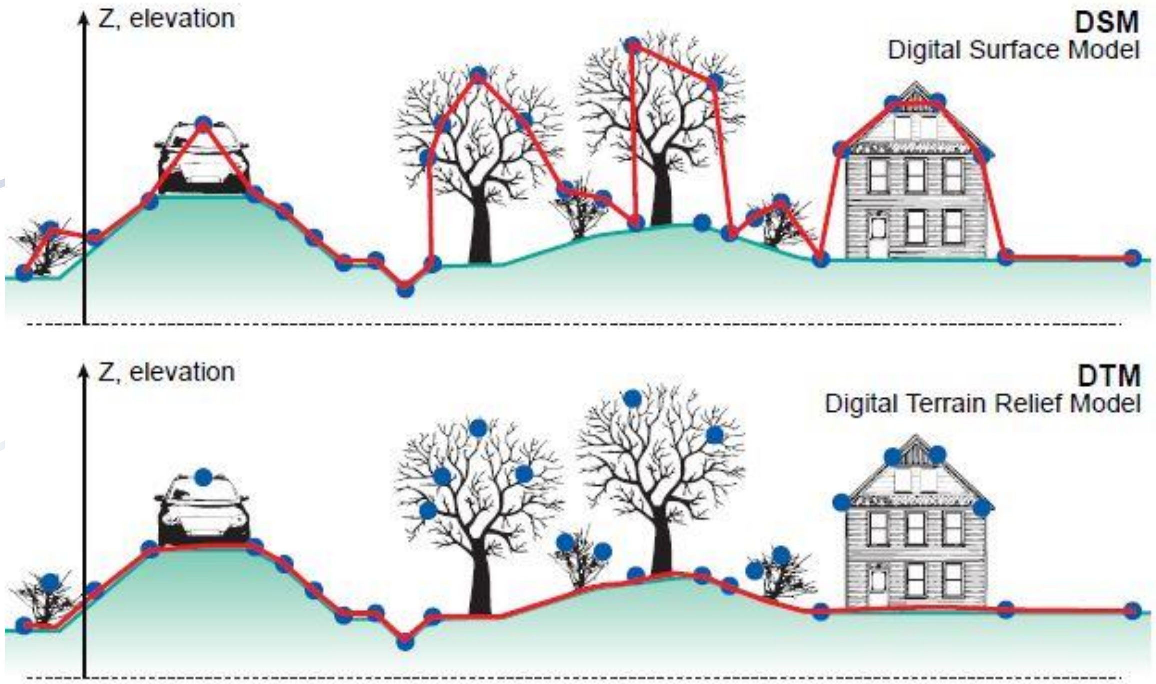

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 731667 (MULTIDRONE)

Safe UAV landing

- Identified potential landing areas both for normal and emergency landing, should be:
  - flat enough,
  - sufficiently large,
  - not occupied by vegetation or buildings.

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 731667 (MULTIDRONE)

Safe UAV landing

- In order to be actually used for UAV landing at a certain time instance, potential landing sites have to be free from:
  - Water,
  - people/crowds,
  - Cars.

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 731667 (MULTIDRONE)

Existing Solutions for landing site detection on Digital Elevation Models (DEM):

- Garg, Mayank, Abhishek Kumar, and P. B. Sujit. "***Terrain-based landing site selection and path planning for fixed-wing UAVs***." International Conference on Unmanned Aircraft Systems (ICUAS). IEEE, 2015.
  - Use of the average height and height variance inside quadtree based DEM partitions.
  - Partitions whose height variance is below a limit are selected as landing sites and merged with neighboring partitions if they have similar average heights.

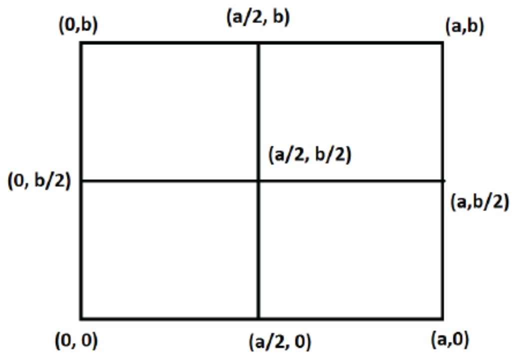

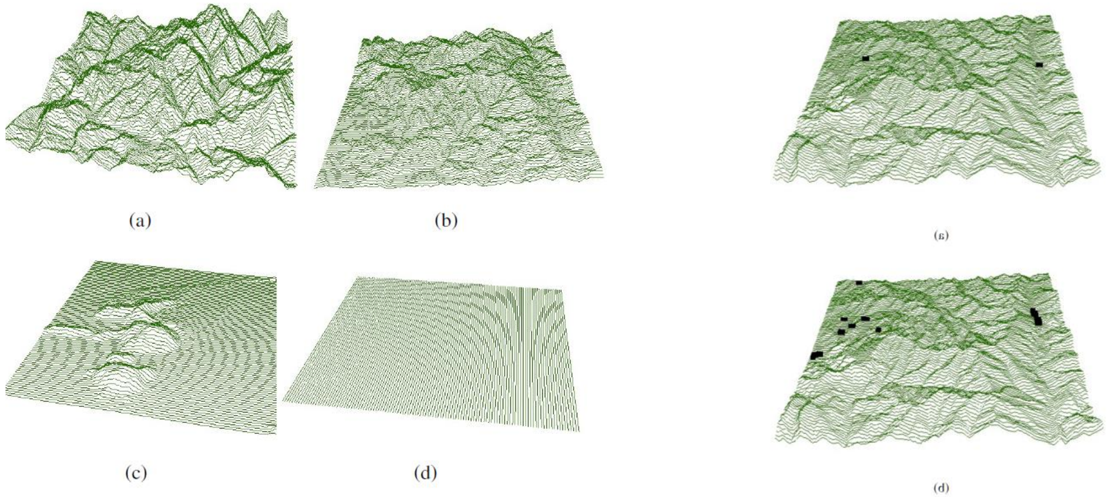

Existing solutions for landing site detection on Digital Elevation Models (DEM):

- Surface fitting on coarse elevation models using Least Squares Error.
- Slope calculation to specify landing areas.

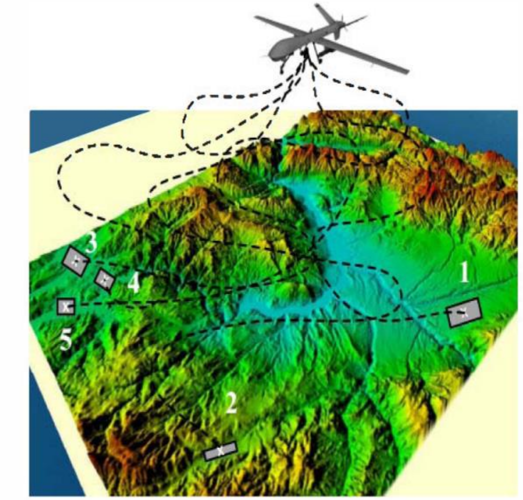

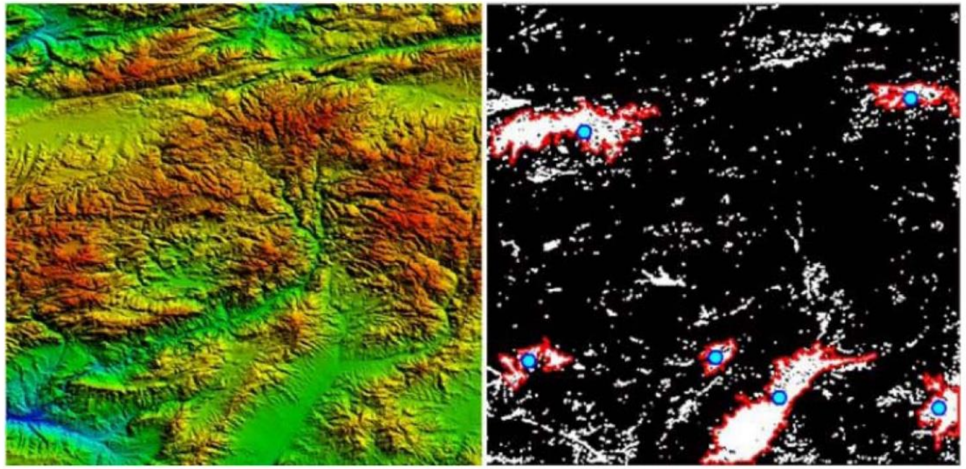

Aydin, Musa, and Emin Kugu, "*Finding smoothness area on the topographic maps for the unmanned aerial vehicle's landing site estimation.*", Sixth International Conference on Digital Information and Communication Technology and its Applications (DICTAP), IEEE, 2016.

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 731667 (MULTIDRONE)

Potential landing site detection algorithm

DSM & DTM

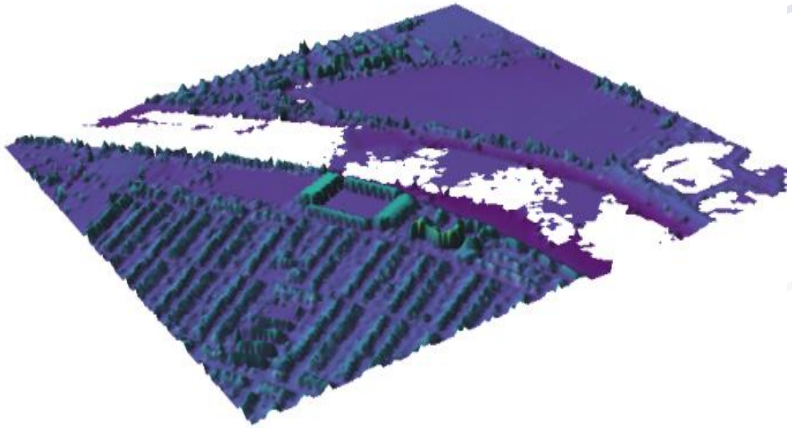

INPUT: 2D DSM & DTM projection (height)

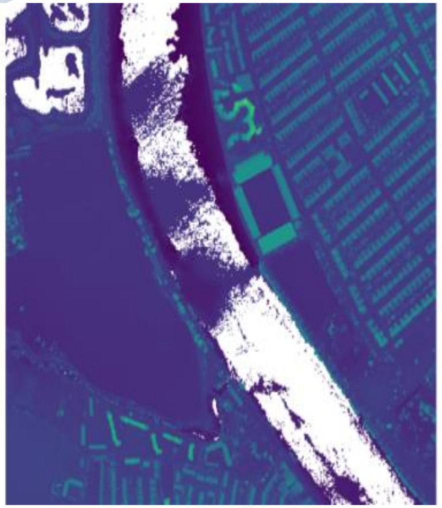

OUTPUT: Safe landing areas in blue color

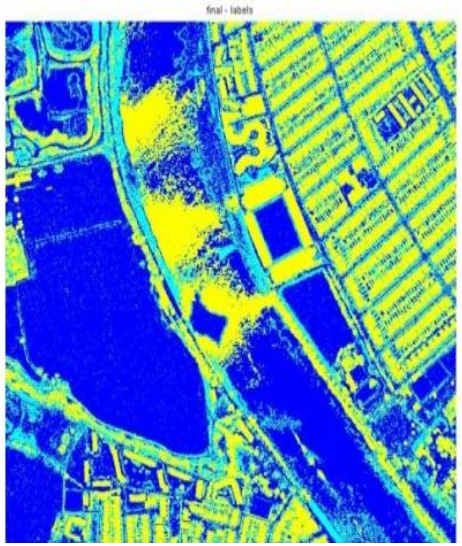

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 731667 (MULTIDRONE)

Potential landing site detection algorithm

- The proposed algorithm consists of 5 discrete steps and its input consists of two digital elevation models in raster format, i.e., as a regular grid of elevation values of a depicted terrain:
  - the digital surface model (DSM) and
  - the digital terrain model (DTM) of a region.

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 731667 (MULTIDRONE)

Potential landing site detection algorithm

- Step 1. Detection of man-made structures and vegetation:
  - Subtracting DTM from DSM and applying a threshold to the outcome, a binary image is derived, marking pixels depicting man-made structures and vegetation of height above a selected (small) threshold.

Binary Mask of Buildings, Vegetation etc

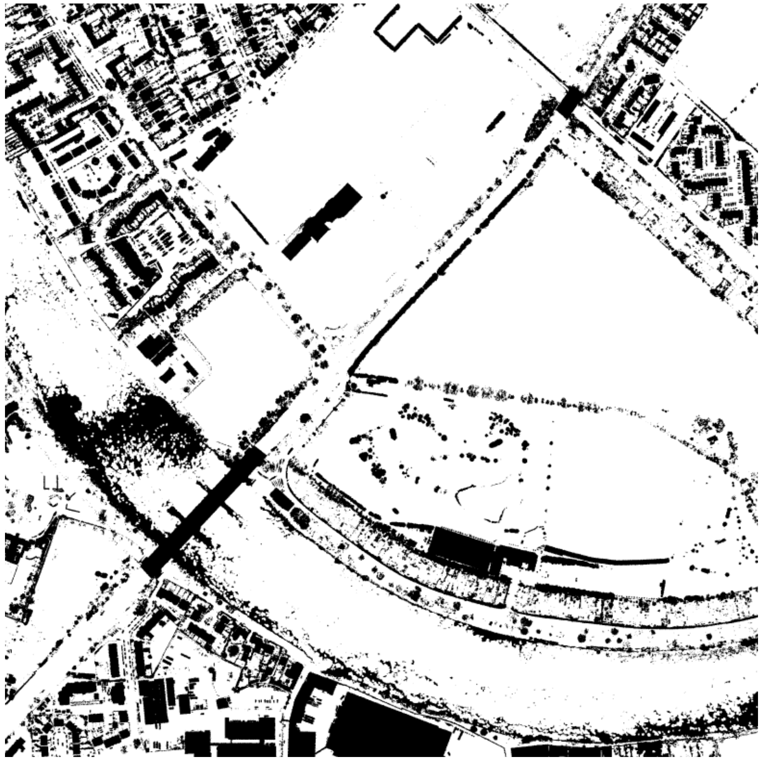

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 731667 (MULTIDRONE)

Potential landing site detection algorithm

- Step 2. Terrain slope determination:
  - Local slope of the depicted areas in the DSM.
  - The maximum rate of change in value (elevation) from a pixel (cell) to its neighbors.
  - The horizontal and vertical derivative approximated by a Sobel operator, scaled by a factor of $8 * cellsize$.

Sobel gradient

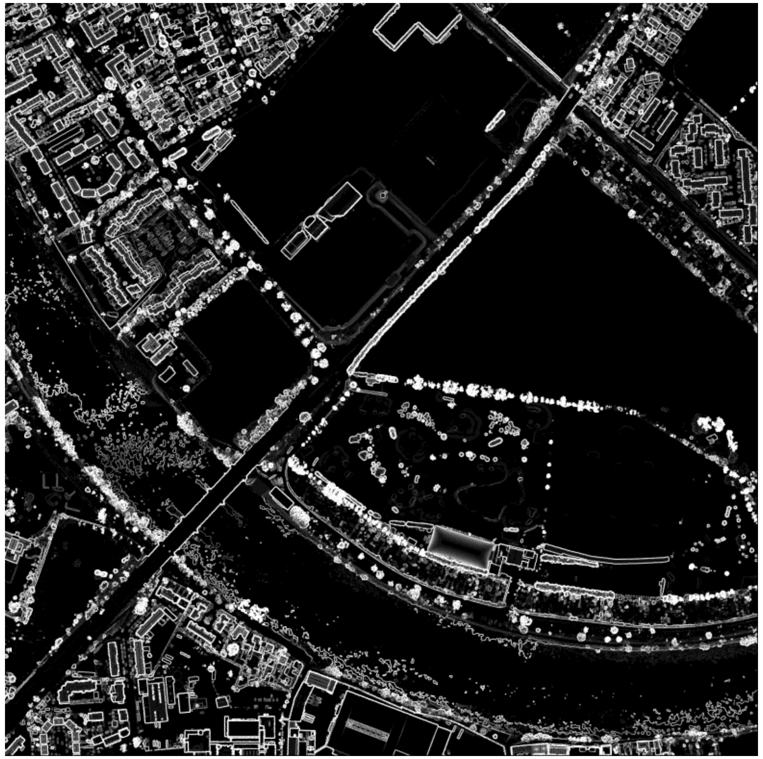

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 731667 (MULTIDRONE)

Potential landing site detection algorithm

$$slope_{degrees} = \frac{180}{\pi} \arctan \sqrt{\left(\left[\frac{dz}{dx}\right]^2 + \left[\frac{dz}{dy}\right]^2\right)}$$

$$\frac{dz}{dx} = \frac{(c + 2f + i) - (a + 2d + g)}{8 * x_{cellsize}}$$

$$\frac{dz}{dy} = \frac{(g + 2h + i) - (a + 2b + c)}{8 * y_{cellsize}}$$

a	b	c
d	e	f
g	h	i

8-neighborhood of a DSM

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 731667 (MULTIDRONE)

Potential landing site detection algorithm

- Horizontal and vertical derivative images, calculated using the Sobel operator, for a digital image $f(n,m)$:

$$G_x(n,m) = \begin{bmatrix} +1 & 0 & -1 \\ +2 & 0 & -2 \\ +1 & 0 & -1 \end{bmatrix} * f \qquad G_y(n,m) = \begin{bmatrix} +1 & +2 & +1 \\ 0 & 0 & 0 \\ -1 & -2 & -1 \end{bmatrix} * f$$

- DSM image (elevation) gradient magnitude approximation according to Sobel:

$$G_(n,m) = \sqrt{G_x^2(n,m) + G_y^2(n,m)}$$

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 731667 (MULTIDRONE)

Potential landing site detection algorithm

- Step 3. Sobel operator gradient image thresholding:
  - The elevation gradient magnitude image is thresholded so that DSM pixels can be classified in flat or non-flat areas based on the local slope.
  - Near flat areas are retained as potential landing areas.

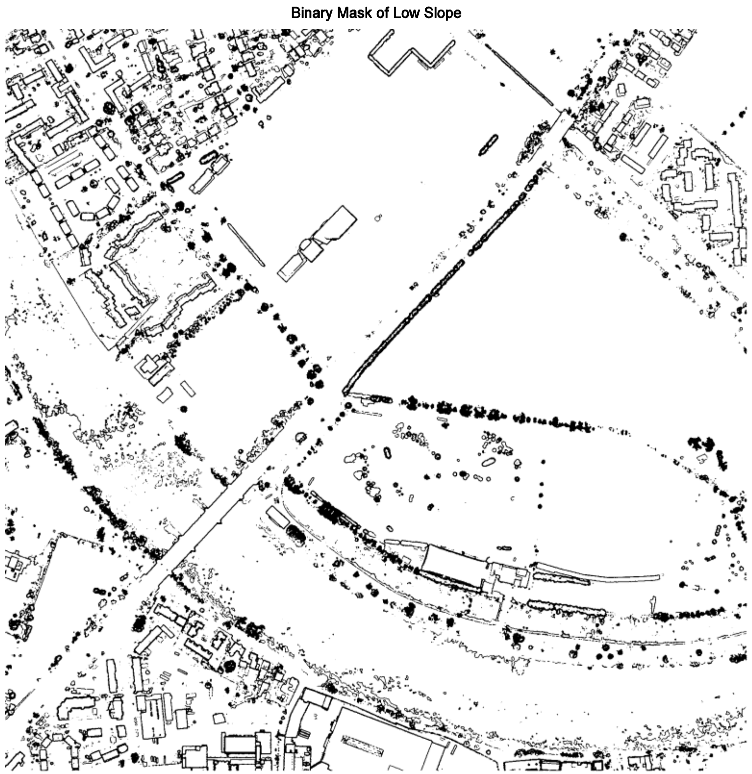

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 731667 (MULTIDRONE)

Potential landing site detection algorithm

- Step 4: Binary image connected components evaluation:
  - Connected components analysis is applied on the binary image resulting from the previous step.
  - Connected components of sufficiently large number of pixels, i.e. of sufficient area for landing are formed.

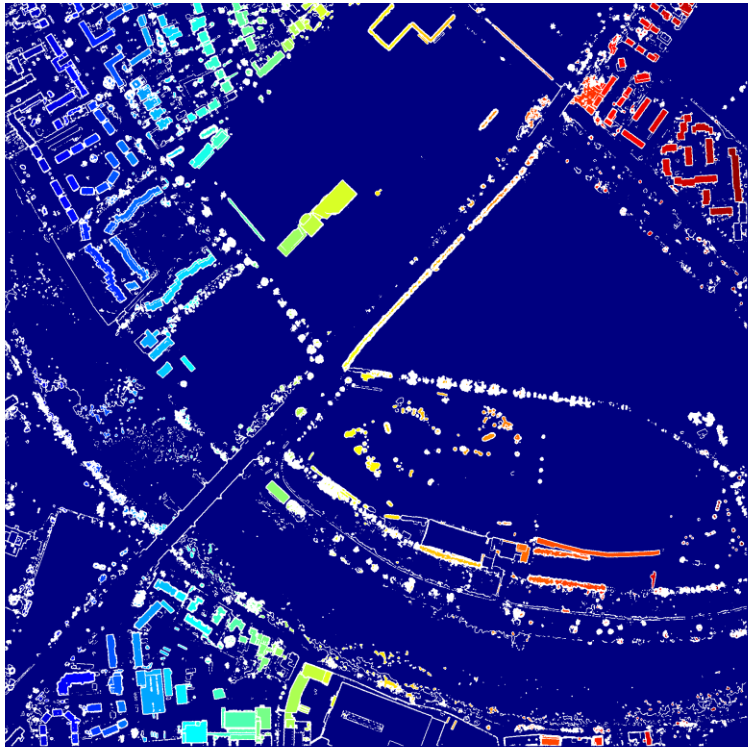

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 731667 (MULTIDRONE)

Potential landing site detection algorithm

- Step 5. Creation of the final landing map:
  - Areas overlapping with buildings and vegetation, found in step 1, are removed from the large low slope areas found in the previous step.
- **Blue pixels:** landing zones, i.e., small slope and enough pixels
- **Light blue pixels:** no landing zones, i.e., large slope or very few pixels
- **Yellow pixels:** no landing zones due to buildings and vegetation

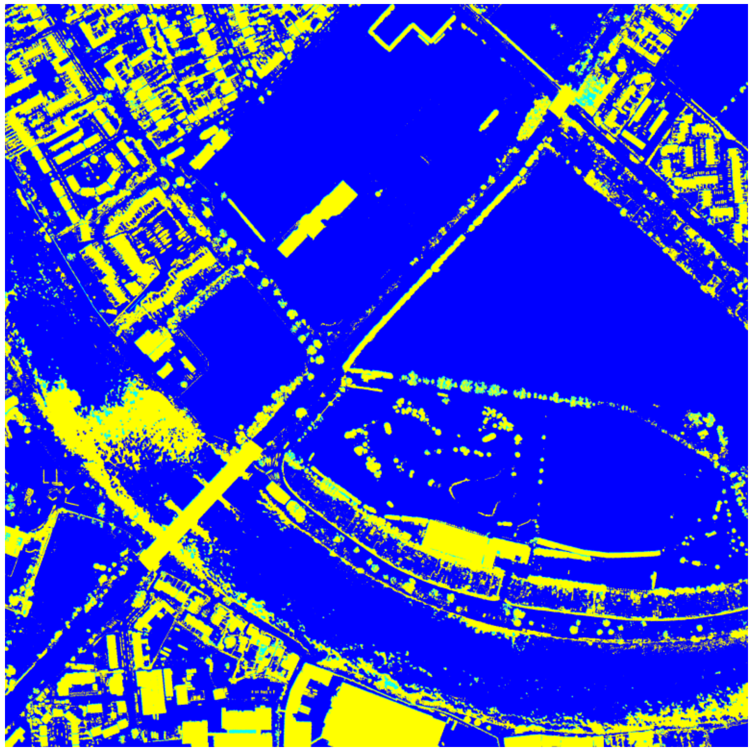

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 731667 (MULTIDRONE)

Proposed potential landing site detection algorithm

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 731667 (MULTIDRONE)

Map Datasets

MultiDrone

- Three of the areas from the publicly available dataset provided by UKs Environment Agency, depicted in the DEM data:
  - Covering urban, suburban, rural and bush areas
  - Spatial resolutions (pixel size per dimension) ranging from $0.25m$ to $2m$.

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 731667 (MULTIDRONE)

Map Datasets

MultiDrone

- Areas no 1 and 2: depicting an urban environment with many obstacles such as buildings and trees (resolution $0.25m/pixel$).
- Area no 3: a rural environment with steep downhill descent parts (resolution $2m/pixel$).
- Ground truth (potential landing sites, areas not suitable for landing): manually constructed through visual inspection of the DEMs and satellite images (the latter obtained by Google Maps).

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 731667 (MULTIDRONE)

Map Dataset

DSM

DSM -Digital Surface Model

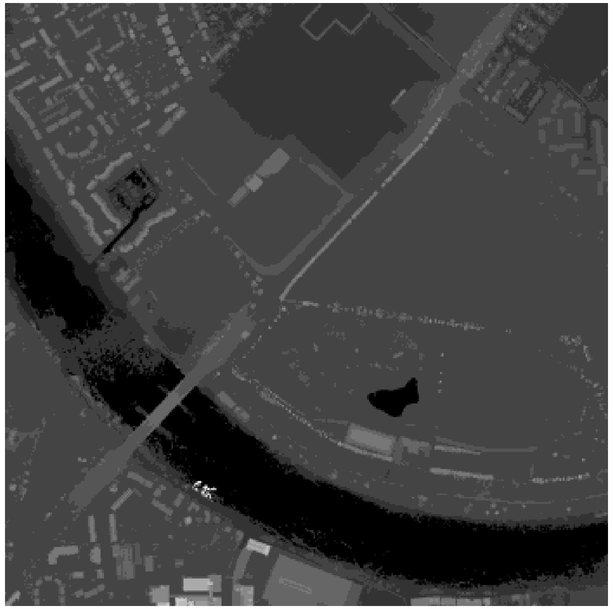

Ground
Truth

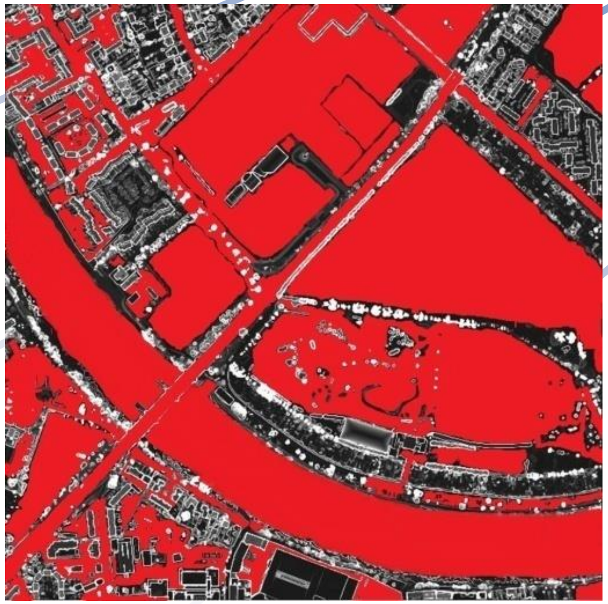

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 731667 (MULTIDRONE)

Experimental evaluation

	No1	No2	No3
precision	0.8008	0.7925	0.8037
recall	0.7668	0.7189	0.7778
f-measure	0.7834	0.7539	0.7905

- The algorithm can identify potential landing sites with satisfactory precision.
  - Precision is much more important than recall, in this case, since landing a UAV in an unsuitable area should be avoided.
- Recall values show room for improvement.

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 731667 (MULTIDRONE)

Q & A

Thank you very much for your attention!

Contact: Prof. I. Pitas
pitas@aiia.csd.auth.gr
www.multidrone.eu

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 731667 (MULTIDRONE)