






Georgios Voulgaris

11 Turner Drive, Oxford, OX2 9GX · (+44) 07449486247

✉ georgios.voulgaris@biology.ox.ac.uk |  <https://github.com/gvsam7> |  <https://georgiosvoulgaris.com/>
 <https://orcid.org/0000-0003-4597-7352> |  [linkedin.com/in/georgios-voulgaris-4b9625249](https://www.linkedin.com/in/georgios-voulgaris-4b9625249) |  @GeorgiosVoulga1

I am a Postdoctoral Researcher in Deep Learning, and a member of the Sheldon Group at the University of Oxford. My research focuses on robustness in Computer Vision and Deep Learning for Remote Sensing datasets.

SKILLS

- Proficient in Deep Learning for Computer Vision.
- Strong coding skills; efficient code, as well as debugging.
- Experienced use of modern source control (Git).
- Experienced in working with GPU's (CUDA).
- Experienced in using Remote Sensing data.
- Python, NumPy, Scipy, Scikit-learn, Pandas.
- PyTorch
- Embedded C
- MATLAB

RESEARCH PROJECTS

Embedding Differential Signal Processing Priors to Deep Learning Models for Semantic Segmentation

Department of Biology, University of Oxford (October 2023 – Present)

I propose PerceptiveNet, a semantic segmentation model built around a novel Logarithmic Gabor-parameterised convolutional layer, whose weights are fully learnable via backpropagation. This layer enforces key Log-Gabor properties such as zero DC component, uniform Fourier domain coverage, and strong spatial localisation. These characteristics were shown to improve segmentation performance, particularly in capturing small, irregular structures within densely packed and visually complex scenes. The results of this research are published at this year's CVPR EarthVision (Single author).



<https://github.com/gvsam7/GabUNet>

Multimodal, Multi-Temporal Deep Learning Data Fusion

Department of Biology, University of Oxford (November 2024 – Present)

I developed a physics-informed mid-level data fusion architecture (FusionNet) for multi-modal integration of thermal and short-wave infrared satellite data. While the application was environmental, the methodological focus on structured fusion and modality-aware feature design is directly relevant to challenges in biomedical imaging, particularly where multiple imaging modalities or spatial contrasts are involved. The results of this research were accepted at this year's IGARSS (first author, oral presentation).



<https://github.com/gvsam7/GeoMap>

Tree Crown Instance Segmentation from UAV Images

Department of Biology, University of Oxford (July 2024 – Present)

I propose a Deep Learning model that extracts more salient feature representations and investigate how the instance segmentation performance is affected when detecting individual tree crowns in dense aerial scenes.



<https://github.com/gvsam7/TreeCrownInstSeg>

Deep Learning Robustness in Semantic Segmentation Tasks for Aerial/Satellite Images

School of Engineering and Informatics, University of Sussex (May 2022 – August 2023)

I propose a Deep Learning architecture that focuses on more salient image features and investigate how it affects the performance of semantic segmentation in detecting water bodies and scenes under canopy/shadows. The results of this research were published at IGARSS23 (first author, oral presentation).



<https://github.com/qvsam7/GabUNet>

British Academy funded research project: Satellite/Aerial image Scene Segmentation

Science Policy Research Unit & the Predictive Analytics Lab, University of Sussex (January 2020 – August 2023)

This was an interdisciplinary project involving the Science Policy Research Unit (SPRU) at the University of Sussex. The aim was to apply deep learning techniques, to map peri-urban agriculture in Ghaziabad India, and research ways of integrating multiple types of data through a web-based mapping and visualisation tool. Thus, support research and stakeholder engagement to understand the trade-offs between Sustainable Development goals (SDGs) in urbanising contexts. For this project I designed a classifier that classifies scenes from aerial images from Ghaziabad India. Classifier predictions are imported to the web application for visualisation.



https://github.com/qvsam7/Deep_Augmentations



<https://wearepal.ai/projects/ssrp>

Deep Learning Coupled with Earth Observation Satellite Data Exploration of Polluting Plants

Satellite Applications Catapult (Internship), Oxford (January 2022 – April 2022)

Applied Deep Learning on remote sensing data, to classify cement plants in China, by exploiting physical properties such as, plant surrounding temperature and soil moisture.

Datasets: Landsat-8 thermal infrared (bands 10 & 11), short wave infrared (bands 6 & 7), and a geological ratio of the short-wave infrared image chips of two classes (cement plants and the surrounding land cover).

Deep Learning architectures: VGG, ResNet, EfficientNet.

The findings of this work were published at the International Geoscience and Remote Sensing Symposium (IGARSS2022) (Oral).



<https://github.com/qvsam7/GeoMap>



<https://youtu.be/OxkWbdjIjWk>

Deep Learning Robustness to Domain Shifts During Seasonal Variations

School of Engineering and Informatics, University of Sussex (August 2021 – March 2023)

In South Asia, the landscape changes dramatically between dry and wet seasons. The main factor responsible for this variation is the flora that transforms the landscape between seasons. These transformations can affect the performance of deep learning models trained to analyse satellite images, especially if there are domain shifts between training and testing data distributions. Here we show that a deep network with a Gabor convolutional layer as the first layer focuses on more salient parts of the image than one which uses a standard convolutional layer meaning that removing colour information is less damaging than for the standard network. Further we show that the proposed architecture is robust in the presence of domain shifts due to seasonal data variations. The findings of this work were published at the IGARSS2022 (first author) and EarthVision CVPR2023 (first author).



https://github.com/qvsam7/Deep_Augmentations



<https://youtu.be/Zci4eASXmkQ>

Water Detection from Satellite/Aerial Images

School of Engineering and Informatics, University of Sussex (October 2018 – December 2019)

Applied Computer Vision and Machine Learning techniques to segment water bodies from aerial images.

PUBLICATIONS

Voulgaris, 2025, June. Bridging Classical and Modern Computer Vision: PerceptiveNet for Tree Crown Semantic Segmentation. In CVPR 2025 Computer Vision and Pattern Recognition Conference.

Voulgaris, Bayaraa, & Rossi, 2025, August. Detecting Cement Plants with Landsat-8: A Physics-Informed, Multi-Temporal, and Multi-Spectral Deep Learning Fusion Approach. In IGARSS 2025 IEEE International Geoscience and Remote Sensing Symposium (Accepted, Oral).

Voulgaris, Philippides, Dolley, Reffin, Marshall, & Quadrianto, 2023, June. Seasonal Domain Shift Dataset in the Global South: A Deep Features Scene Analysis. In CVPR 2023 Computer Vision and Pattern Recognition Conference.

Voulgaris, Philippides, & Quadrianto, 2023, July. Water Physics Aware Semantic Segmentation Through Texture-Biased U-NET Architectures. In IGARSS 2023 IEEE International Geoscience and Remote Sensing Symposium (Oral).

Voulgaris, Philippides, & Quadrianto, 2022, July. Deep Learning Robustness to Domain Shifts During Seasonal Variations. In IGARSS 2022-2022 IEEE International Geoscience and Remote Sensing Symposium (pp. 417-420). IEEE.

Rossi, Tkachenko, Bayaraa, Foster, Reece, Scott, Voulgaris, Christiaen and McCarten, 2022, July. Detection and Characterisation of Pollutant Assets with AI and EO to Prioritise Green Investments: The Geoasset Framework. In IGARSS 2022-2022 IEEE International Geoscience and Remote Sensing Symposium (pp. 7717-7720). IEEE. (Oral)

PROFESSIONAL EXPERIENCE

Technical Authority (Reviewer)

CVPR EarthVision (2024, 2025)

Workshop Co-Organiser

ICIP 2025 Workshop on Computer Vision for Ecological and Biodiversity Monitoring (2025)

Deep Learning Exploration for Polluting Plants

Satellite Applications Catapult (Internship), Oxford (January 2022 – April 2022)

Teaching Assistant

School of Engineering and Informatics, University of Sussex (October 2018 – June 2020)

Teaching Assistant in: -

- Computer Vision.
- Industrial Automation Mechatronics.
- Advance Electronic Systems.

Designed and Developed a Fully Automated Green House

Brighton (May 2013 – October 2018)

Design Engineer Control Systems Oil and Gas

Rockwell Automation, Horsham (September 2006 – May 2013)

EDUCATION

EPSRC funded PhD Informatics

University of Sussex (October 2018 – August 2023)

Deep Learning Aerial Scene Analysis: Extracting Salient Features for Domain Adaptation and Semantic Segmentation Tasks.

DISCNET Scholarship

Consortium – University of Sussex, University of Southampton, University of Portsmouth, QMUL, and the Open University (January 2021 – December 2021)

Doctoral training in Machine Learning, High Performance Computing, Big Data, and industrial placement.

Machine Learning Summer School

MLSS-Indonesia (August 2020)

Awards received: - Best Research Proposal, Most Active Participants

MSC Modern Digital Communication Systems

University of Sussex (September 2005 – September 2006)

Obtained a merit.

BEng Hons Electronics with Communication Engineering Modern Digital Communication Systems

University of Brighton (September 2001 – July 2005)

Obtained a 2.2

REFERENCES

Will be given upon request.