ELIAS

European Lighthouse of AI for Sustainability

www.elias-ai.eu

Funded by the European Union

This project has received funding from the European Union's Horizon Europe Research and Innovation Programme under grant agreement No. 101120237.

Funded by
the European Union

Theme Development Workshop on Sustainability & AI

Theme Development Workshops (TDWs)

7 March, 2025
Bucharest, Romania

www.elias-ai.eu

ELiAS

Funded by the European Union

TDW on Sustainability & AI

Agenda

09:00 – 09:10 | Welcome & Introductions

- Opening remarks
- Introduction to the workshop's theme: AI and Sustainability
- Overview of the sessions and objectives

Nicu Sebe (UNITN), Filareti Tsalakanidou (CERTH)

09:10 – 09:30 | Session 1: Regulatory and Ethical Aspects

Keynote:

- Marko Milosaljevic (University of Ljubljana, Slovenia)

Q&A and Discussion: 5 mins

09:30 – 10:10 | Session 2: AI for a Sustainable Planet

Keynotes:

- Saso Dzeroski (Jozef Stefan Institute, Slovenia)
- Marius Leordeanu (National University of Science and Technology POLITEHNICA Bucharest, Romania)

Q&A and Discussion: 10 mins

10:10 – 10:50 | Session 3: AI for a Sustainable Society

Keynotes:

- Nicolò Cesa-Bianchi (University of Milan, Italy)
- Ioana Manolescu (INRIA, France)

Q&A and Discussion: 10 mins

10:50 – 11:10 | Break

11:10 – 11:30 | Session 4: Trustworthy AI for Individuals

Keynotes:

- Lorenzo Baraldi (University of Modena and Reggio Emilia, Italy)
- Nicu Sebe (University of Trento, Italy)

Q&A and Discussion: 10 mins

11:30 – 11:50 | Session 5: Fostering the Next Generation of AI Talents

Keynote:

- Charlotte Delage (Institute Polytechnique of Paris, France)

Q&A and Discussion: 5 mins

11:50 – 12:30 | Session 6: Entrepreneurship and Tech Transfer

Keynotes:

- Nina Peters (University of Tubingen, Germany)
- Isabelle Siegrist (ETH Zurich, Switzerland)

Q&A and Discussion: 10 mins

12:30 – 12:45 | Conclusion & Wrap-up

- Summary of key takeaways
- Final remarks and next steps

Nicu Sebe (UNITN), Filareti Tsalakanidou (CERTH)

ELiAS

Multi-view Self-supervised Consensus for Sustainable AI

Marius Leordeanu

University POLITEHNICA of Bucharest

1st Theme Development Workshop on Sustainability & AI

The Planets have aligned for the Romanian Artificial Intelligence

On February 28th 2025, the Government and Consortium of Universities and Companies, have aligned to sign the start of the **Romanian AI HUB**
… along with the rare alignment of Mars, Jupiter, Uranus, Saturn, Venus, Neptune and Mercury.

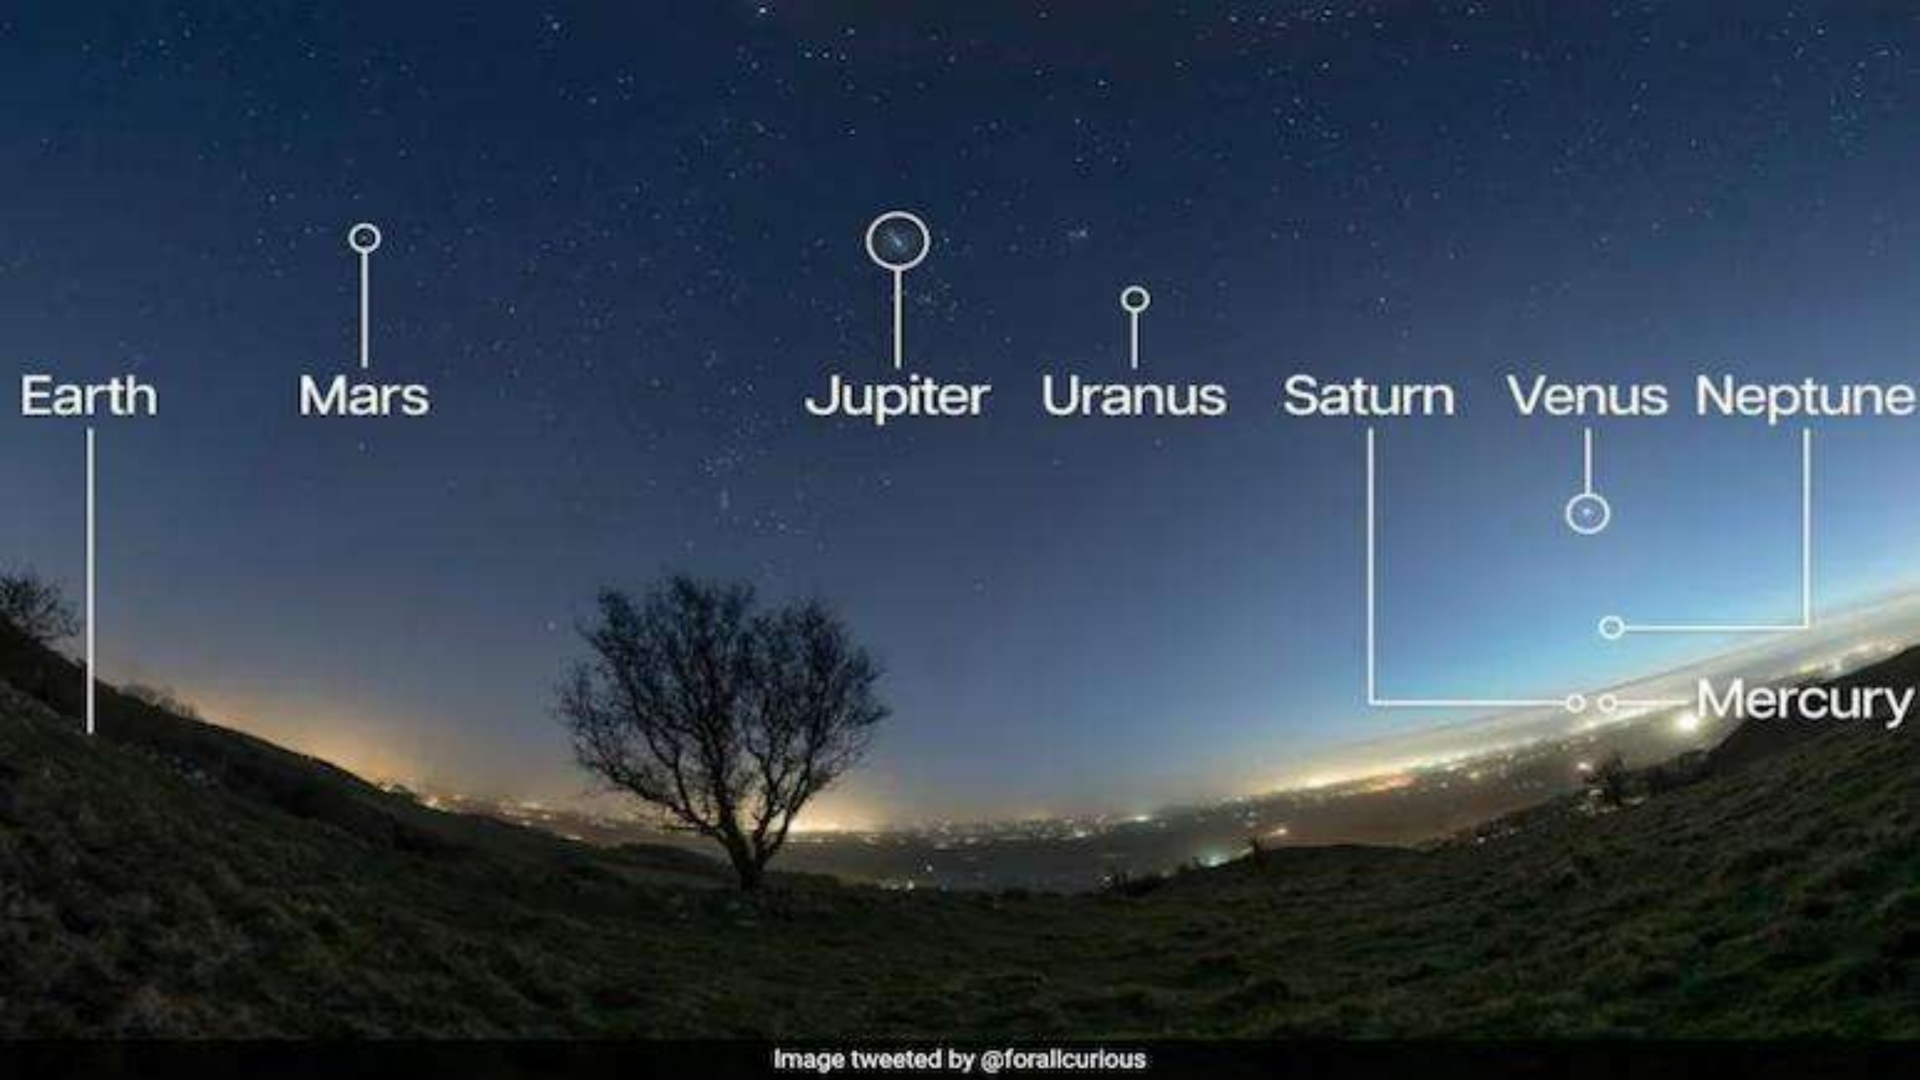
Earth
Mars
Jupiter
Uranus
Saturn
Venus
Neptune
Mercury
Image tweeted by @forallcurious

HRIA
Romanian Hub for Artificial Intelligence

Co-funded by the European Union

GUVERNUL ROMÂNIEI

HRIA Consortium – 15 partners

Leader – National University of Science and Technology POLITEHNICA București

Technical University of Cluj-Napoca

West University of Timișoara

Technical University "Gheorghe Asachi" of Iași

University Politehnica of Timișoara

University of Bucharest

University "Babeș Bolyai" of Cluj-Napoca

Beia Consult International SRL

IT Center for Science and Technology SRL

Greensoft SRL

Neural Grader SRL

Safetech Innovations SA

SoftTehnica SRL

Technology Systems and Services International SRL

Terrasigna SRL

Project in the framework "Smart Growth, Digitalization and Financial Instruments Program" (PoCIDIF)
Duration: 5 years (2025-2029)
Non-reimbursable budget: 65 Million Euro

Main research topics

- Trustworthy AI
- AI Hardware Technologies
- Intelligent and Autonomous Systems
- Intelligent Human Computer Interfaces
- AI in Medicine and Healthcare
- AI in Cybersecurity and Space
- AI Applications for Smart Cities
- AI in Environment
- Development and testing of AI prototypes and systems

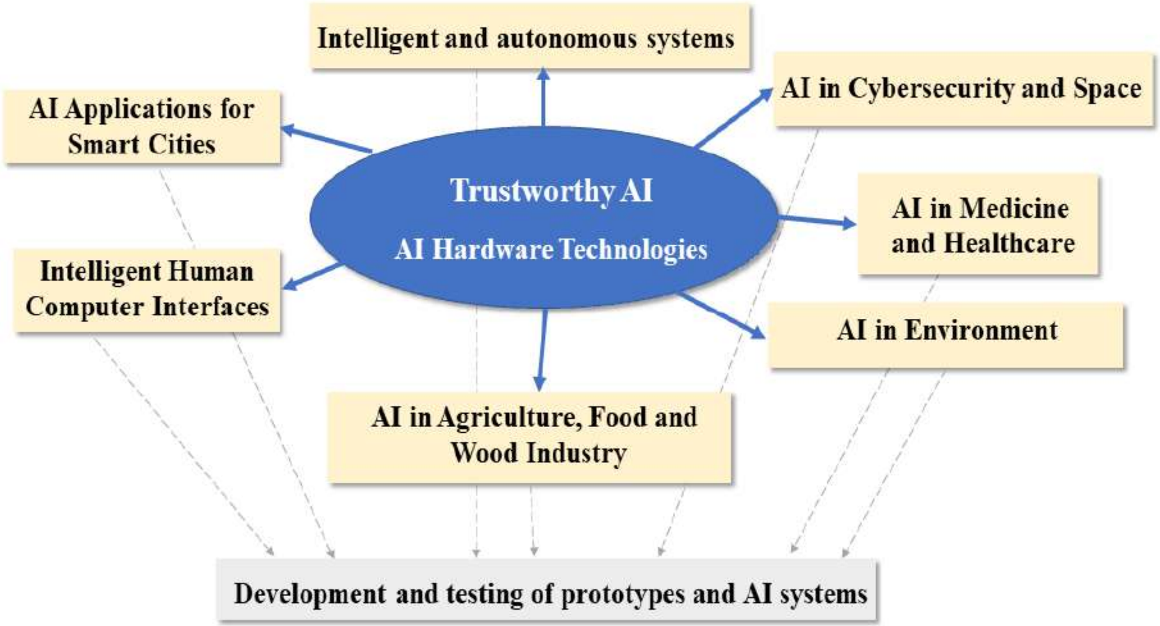

ENFIELD: European Lighthouse to Manifest Trustworthy and Green AI

HORIZON-CL4-2022-HUMAN-02-02

Vision and mission: Create a unique European Centre of Excellence that excels in fundamental research in the pillars of **Adaptive, Green, Human-Centric, and Trustworthy AI** and will further advance the research within applications of **healthcare, energy, manufacturing and space**

Project Concept

Thematic Areas

- **GREEN AI** (SINTEF)
- **ADAPTIVE AI** (IMT)
- **HUMAN-CENTRIC AI** (TUE)
- **TRUSTWORTHY AI** (NTNU)

Available Resources

- *Research Labs*
- *Data*
- *Infrastructure*
- *DIHIWARE Platform*

Strategic Initiatives

- *Policy Initiatives*
- *Networks and Alliances*
- *(Inter-) National R&I Projects*

Innovation Paradigm

- *DIHs and Incubators*
- *Industry Collaborations*
- *Open Cascading Calls: TES & TIS*

Application Domains

- **ENERGY** (INESC)
- **HEALTHCARE** (ICCS)
- **MANUFACTURING** (POLIMI)
- **SPACE** (ECOE)

Consortium

30 partners from 19 European countries

Common Research Vision and Roadmap
Safety and Security Risks Assessment Framework Regulations
Standards and Certification
Collaboration, Networking and Exchange Programmes

UPB leads WP4: ***Common Vision and Roadmaps*** and co-leads the ***Human-Centric AI*** thematic area

Consensus
in all Earth Systems, at all scales, between Different Views, Modalities and Tasks

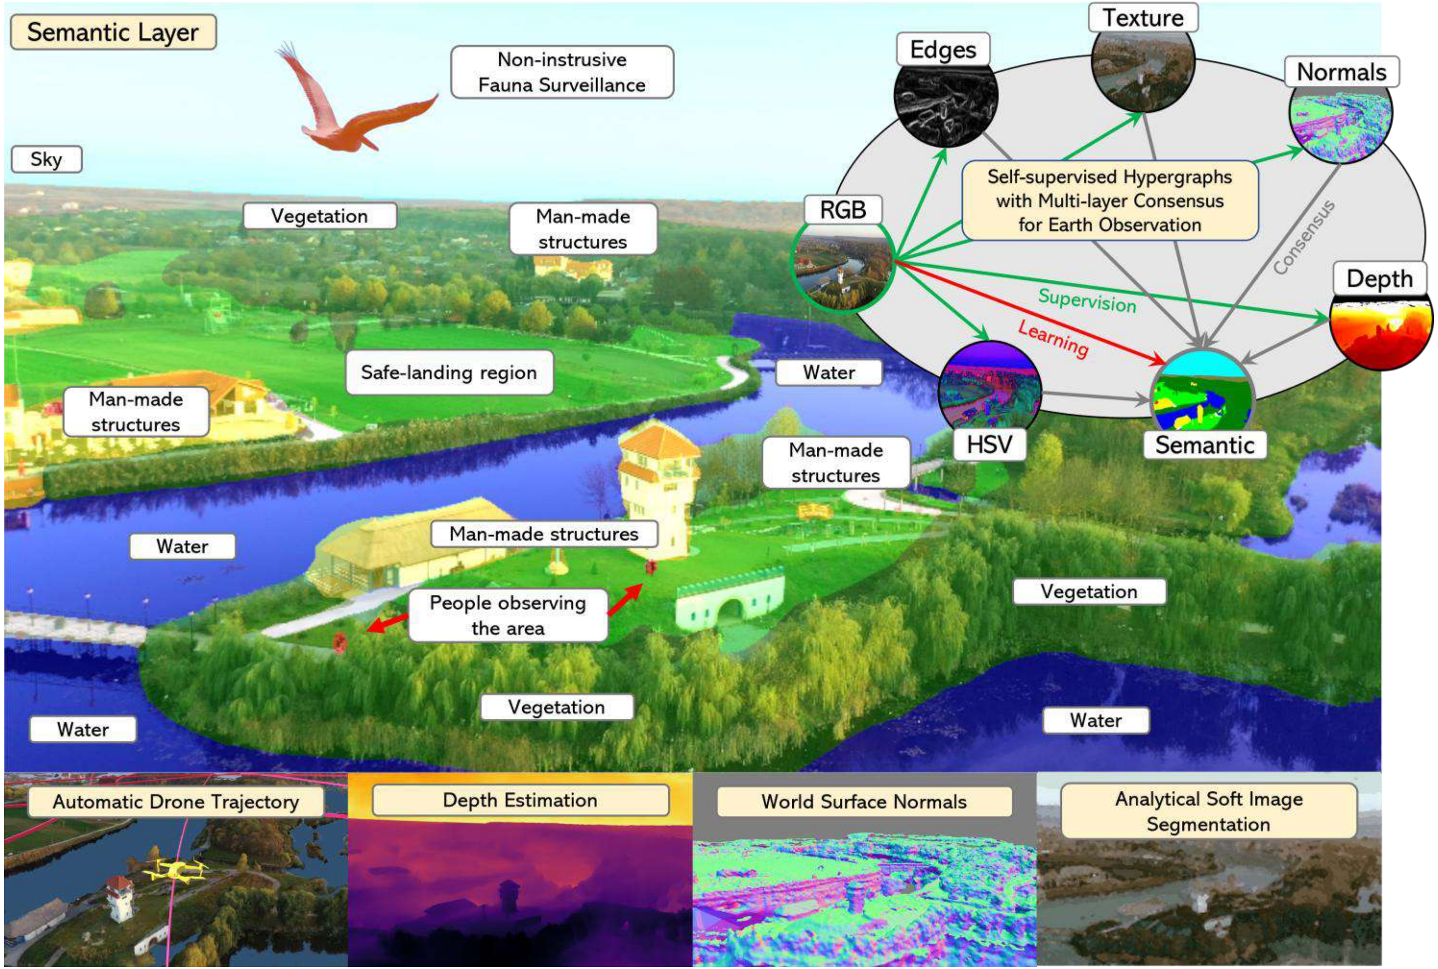

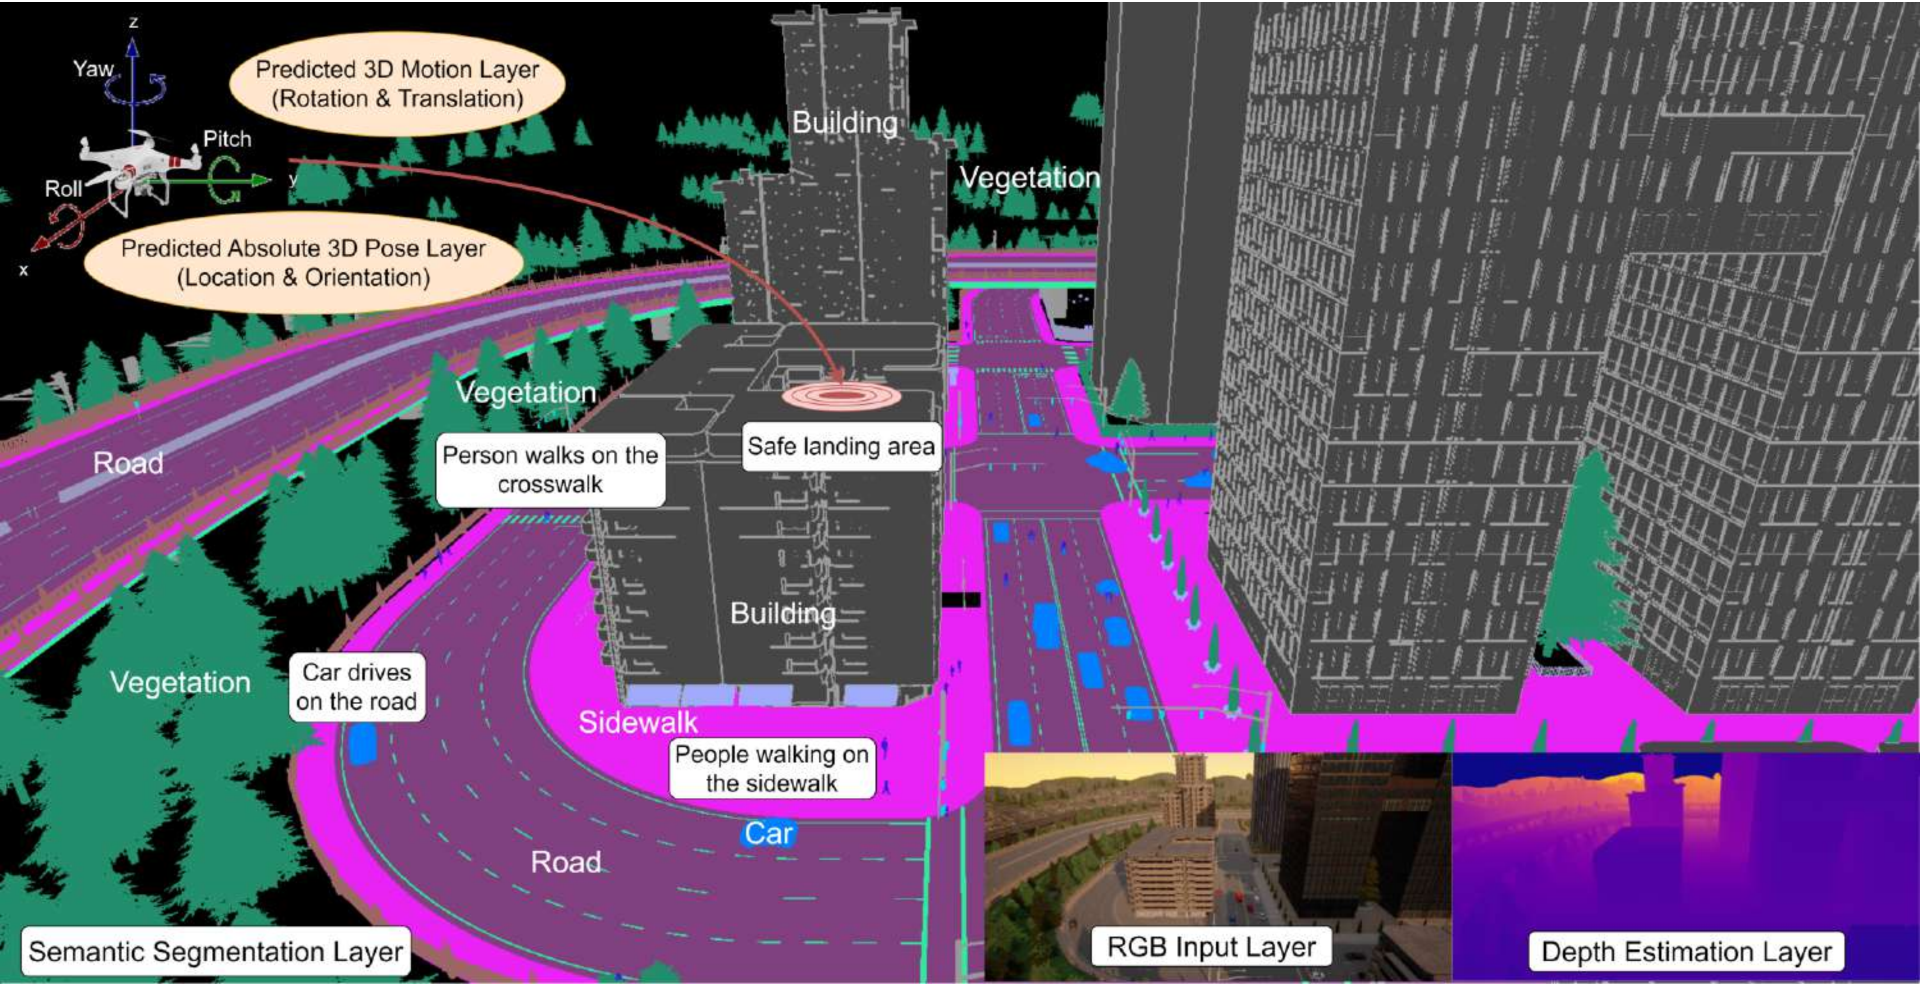

Predicted 3D Motion Layer (Rotation & Translation)
Predicted Absolute 3D Pose Layer (Location & Orientation)
Yaw
Pitch
Roll
z
y
x
Building
Vegetation
Road
Safe landing area
Person walks on the crosswalk
Car drives on the road
Sidewalk
People walking on the sidewalk
Car
Semantic Segmentation Layer
RGB Input Layer
Depth Estimation Layer

Unsupervised Learning by Neural Graph Consensus

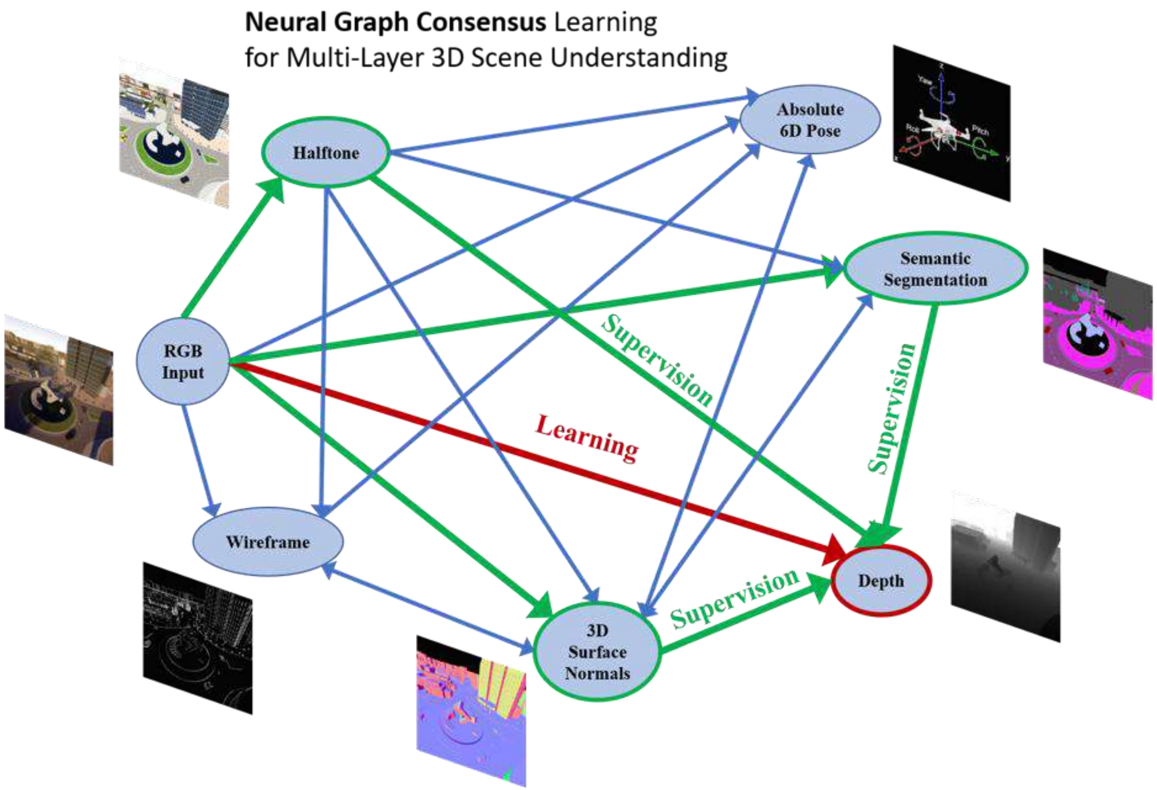


Mihai Pirvu


Dragos Costea


Alina Marcu


Emil Slusanschi


Rahul Sukthankar

Google Research

Marius Leordeanu, Mihai Pirvu, Dragos Costea, Alina Marcu, Emil Slusanschi, Rahul Sukthankar
Semi-Supervised Learning for Multi-Task Scene Understanding by Neural Graph Consensus
Accepted at AAAI Conference on Artificial Intelligence (AAAI), 2021

Theoretical and numerical analysis

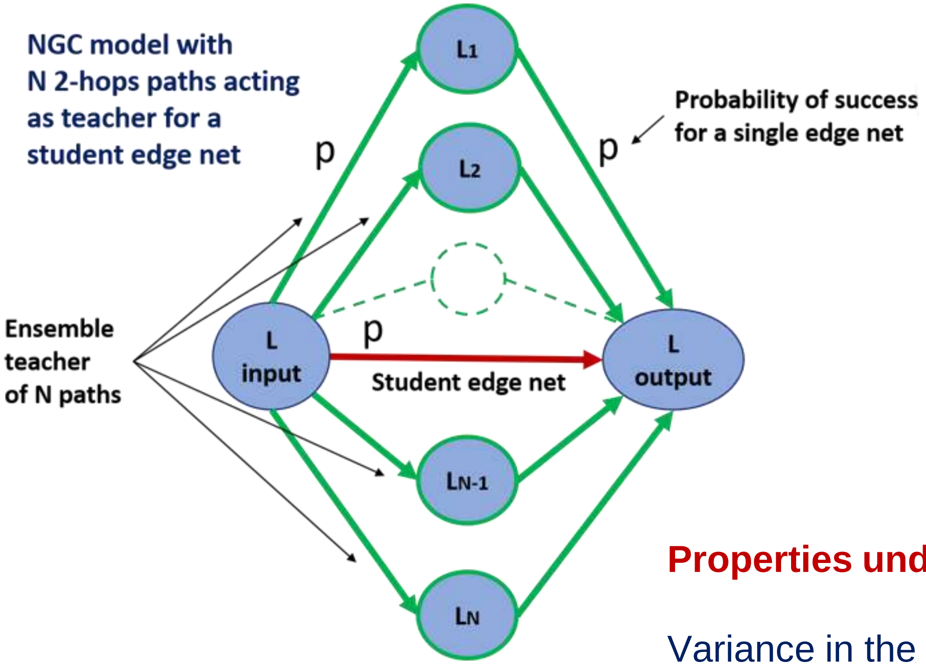

Mathematical intuition:

$$J_{\mathrm{unsup}} = \sum (x_q^c - \frac{1}{N_q} \sum_{q=1}^{N_q} x_q^{(a)})^2$$

Output of one edge

Consensual output of path ensemble

Properties under independence assumptions:

Variance in the graph will decrease

Over many iterations the model could converge to the truth

Results over several unsupervised learning iterations

Representation	Evaluation Metric	Iteration 0	Iteration 1		Iteration 2	
		EdgeNet	NGC	Distil. EdgeNet	NGC	Distil. EdgeNet
Depth	L1 (meters)	4.9844	3.4867	4.2802	**3.2994**	**3.9508**
	Pixels ↑ (%)	-	79.30	60.66	**79.69**	**61.90**
Surface Normals (C)	L1 (degrees)	8.4862	7.7914	8.2891	**7.4503**	**7.6773**
	Pixels ↑ (%)	-	74.18	53.59	**74.61**	**53.94**
Surface Normals (W)	L1 (degrees)	11.8859	8.8248	10.7500	**8.5282**	**8.6714**
	Pixels ↑ (%)	-	79.95	57.88	**81.12**	**61.14**
Semantic Segmentation	Accuracy	0.9001	0.9181	0.9019	**0.9245**	**0.9283**
	mIOU	0.4840	0.4978	0.4980	**0.5258**	**0.5159**
	Pixels ↑ (%)	-	79.46	69.62	**81.49**	**71.95**
Wireframe	Accuracy	0.9617	0.9655	0.9654	**0.9661**	**0.9655**
	Pixels ↑ (%)	-	77.71	72.57	**78.02**	**73.46**
Position	L2 (meters)	25.7597	15.5383	20.0204	**12.0764**	**15.5599**
Orientation	L1 (degrees)	3.8439	2.5001	3.3961	**2.2088**	**3.0005**

Consensus among multiple interpretations improves not only quality, but also space-time consistency, which makes the results more trustworthy

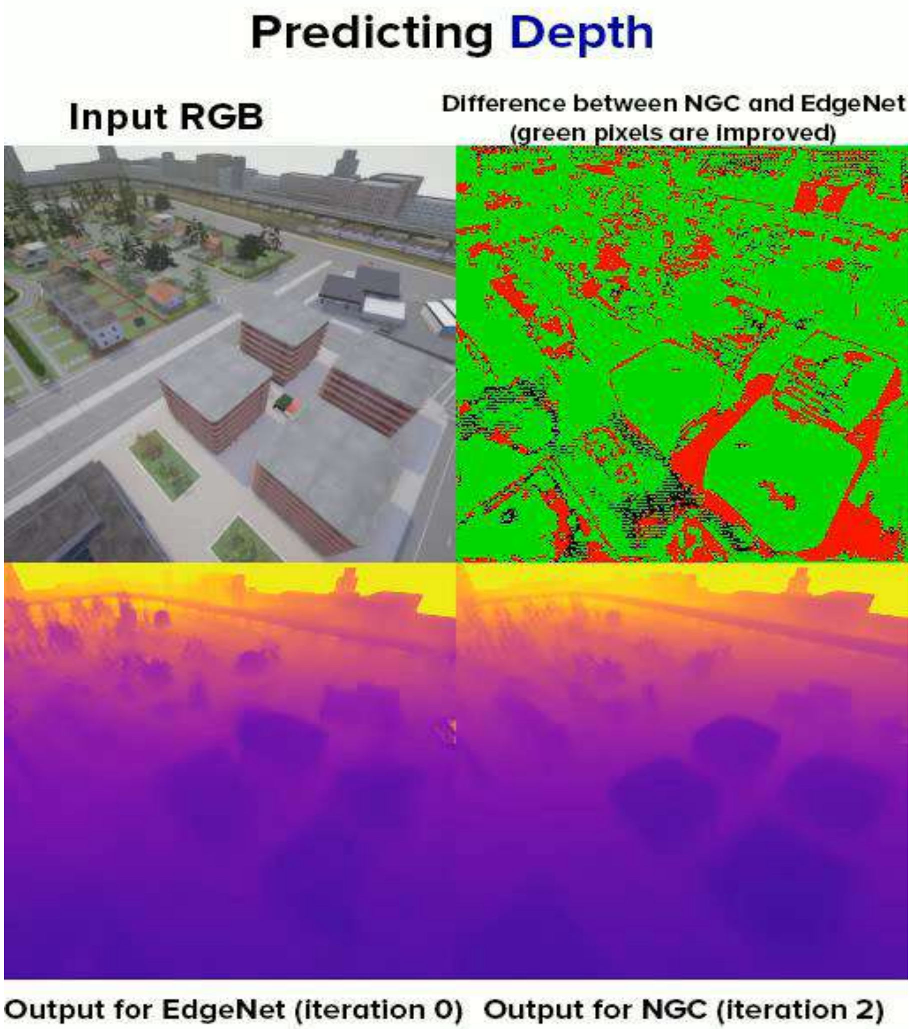

NGC learning is related to the idea of Homeostasis in living organisms

Homeostasis

The innate tendency of all living organisms towards a relatively stable equilibrium between their many interdependent elements, in order to maintain their life (Betts et al. 2016).

Betts, J. G., P. Desaix, E. Johnson, J. E. Johnson, O. Korol, D. Kruse, and K. A. Young. "Anatomy and physiology. OpenStax." (2016).

Self-supervised Multi-task Consensus in the Hydra Nervous System

One of the first multi-neuron nervous systems

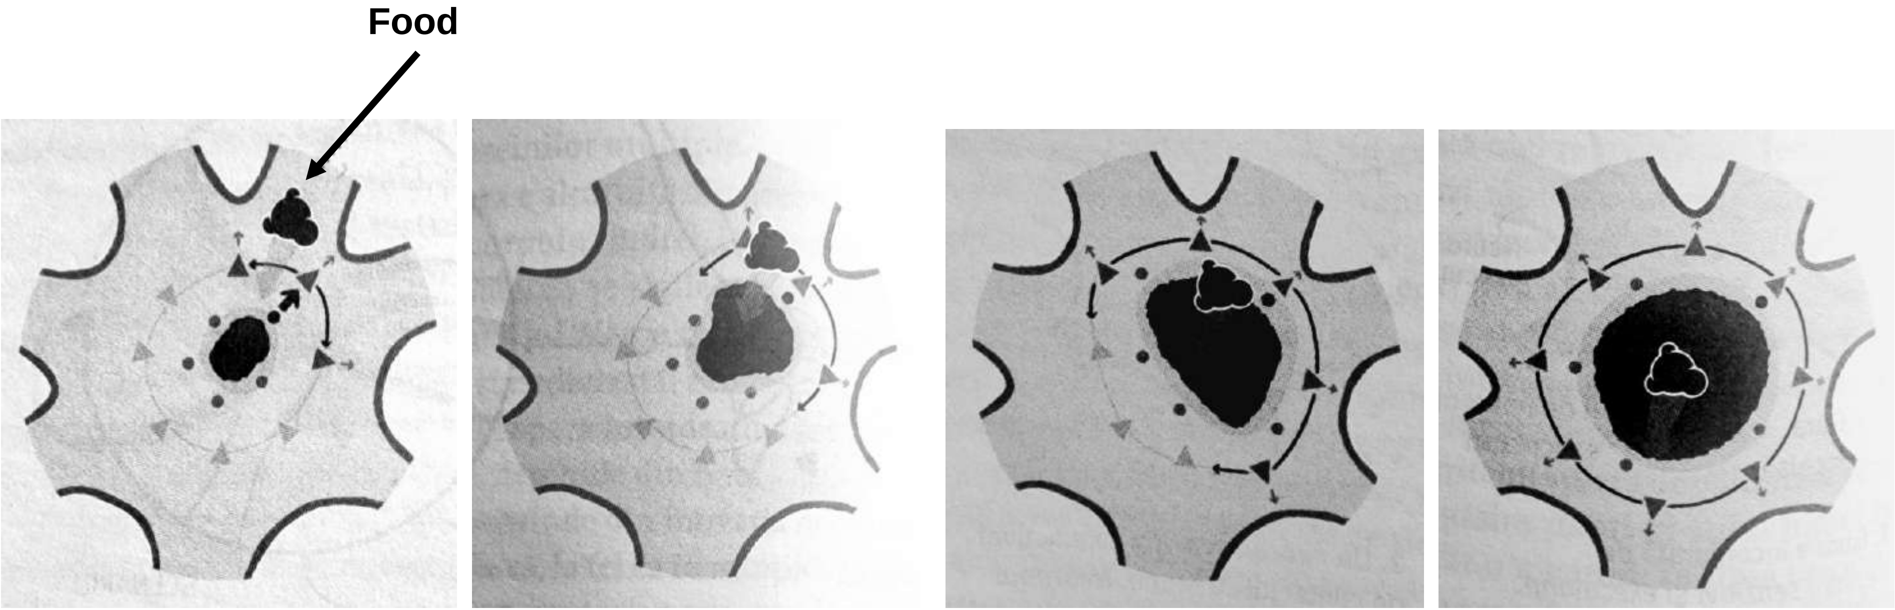

We live in a connected world!

By having a connected mind!

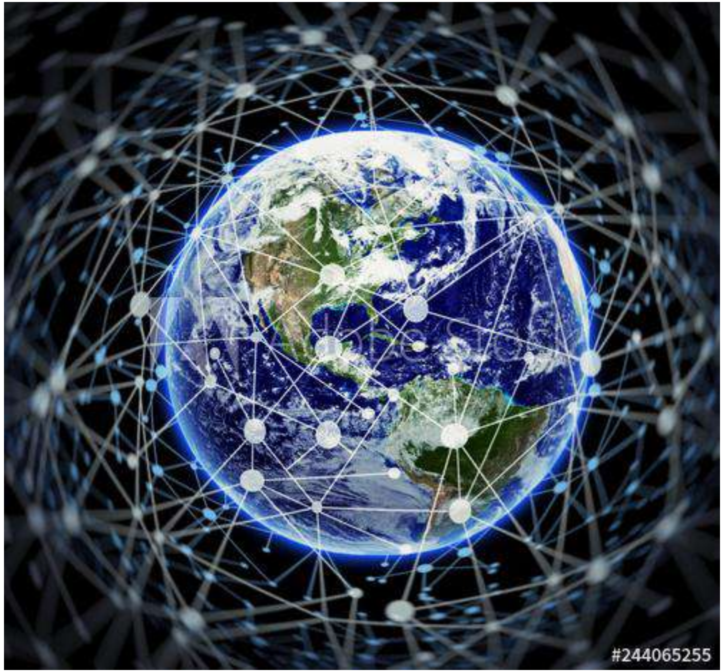

Multi-view and Multi-task Consensus could lead to a self-regulatory and self-supervised way to achieve sustainability.

Self-supervised consensus in fully connected multi-task graphs

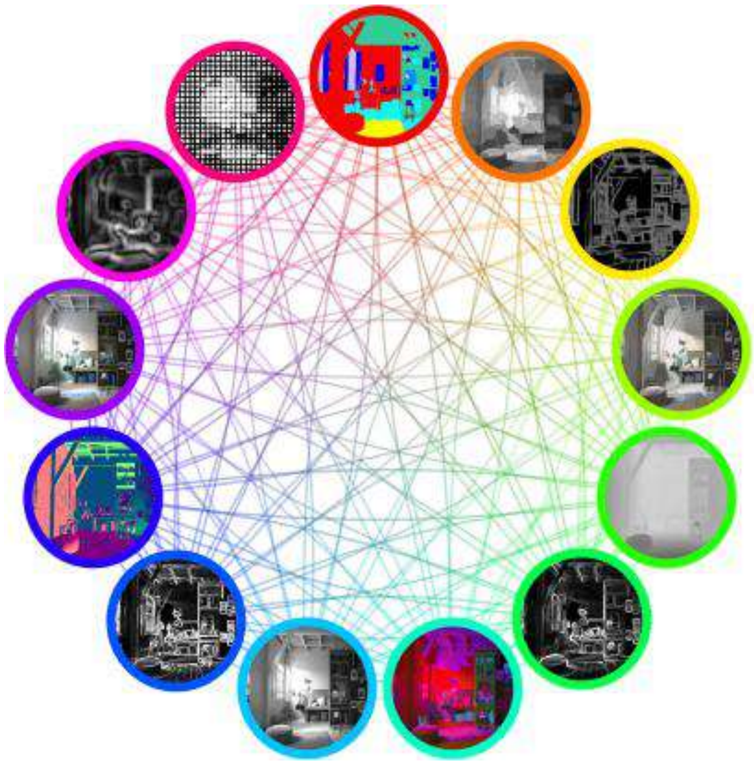

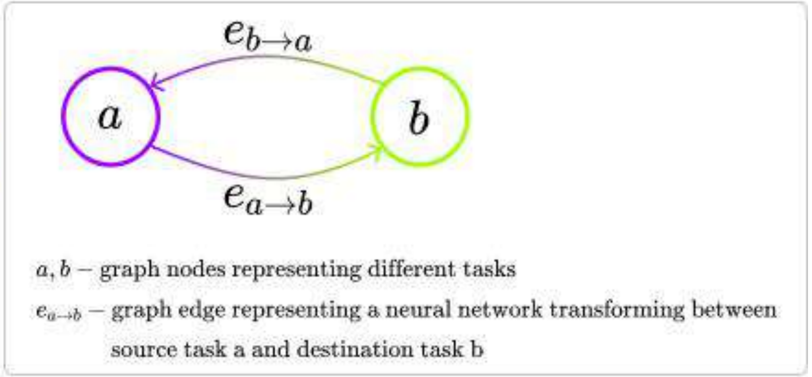

Emanuela Haller

Elena Burceanu

Emanuela Haller, Elena Burceanu and Marius Leordeanu - *Self-Supervised Learning in Multi-Task Graphs through Iterative Consensus Shift* – **British Machine Vision Conference (BMVC) 2021**

Improving over unsupervised learning iterations

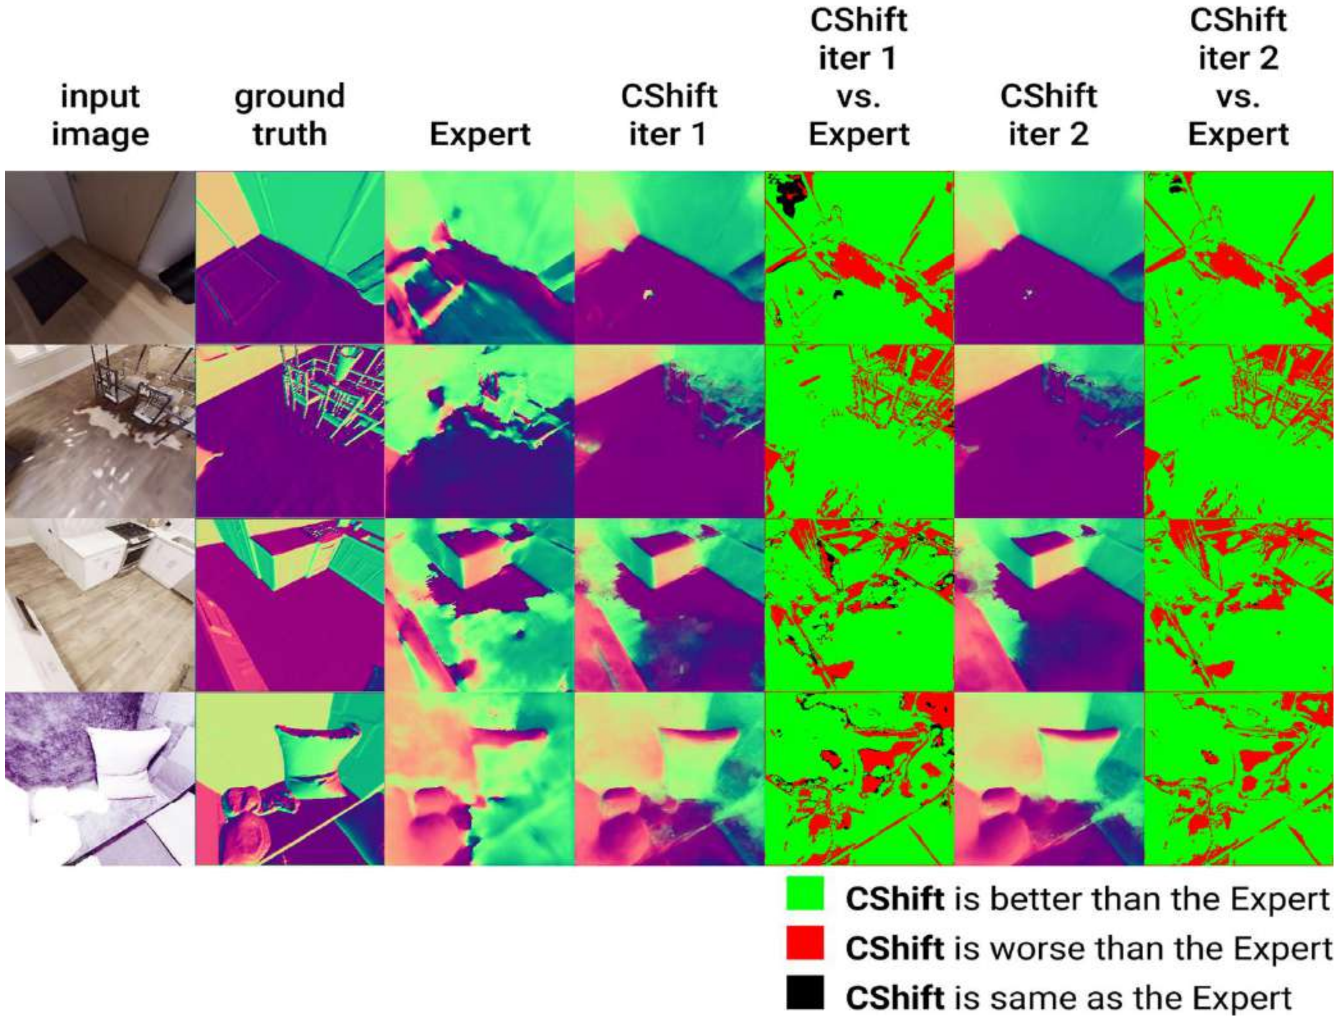

Consensus in real-world experiments with UAVs

DroneScapes Dataset

10 varied scenes, 9 in Romania and 1 in Norway

50 minutes of videos at 3840 × 2160 30 FPS

Includes GPS information, linear and angular velocities and absolute camera angles

Automatic Segmentation Propagation on Different Scenes

Only 1.6% of of frames are manually labeled

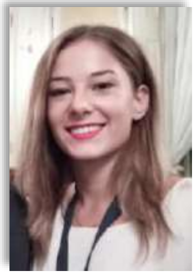

Alina Danciu

Alina Marcu

Marcu, Alina, Vlad Licaret, Dragos Costea and Marius Leordeanu. "**Semantics through time: Semi-supervised segmentation of aerial videos with iterative label propagation.**" ACCV 2020.

Automatic Segmentation Propagation on Different Scenes

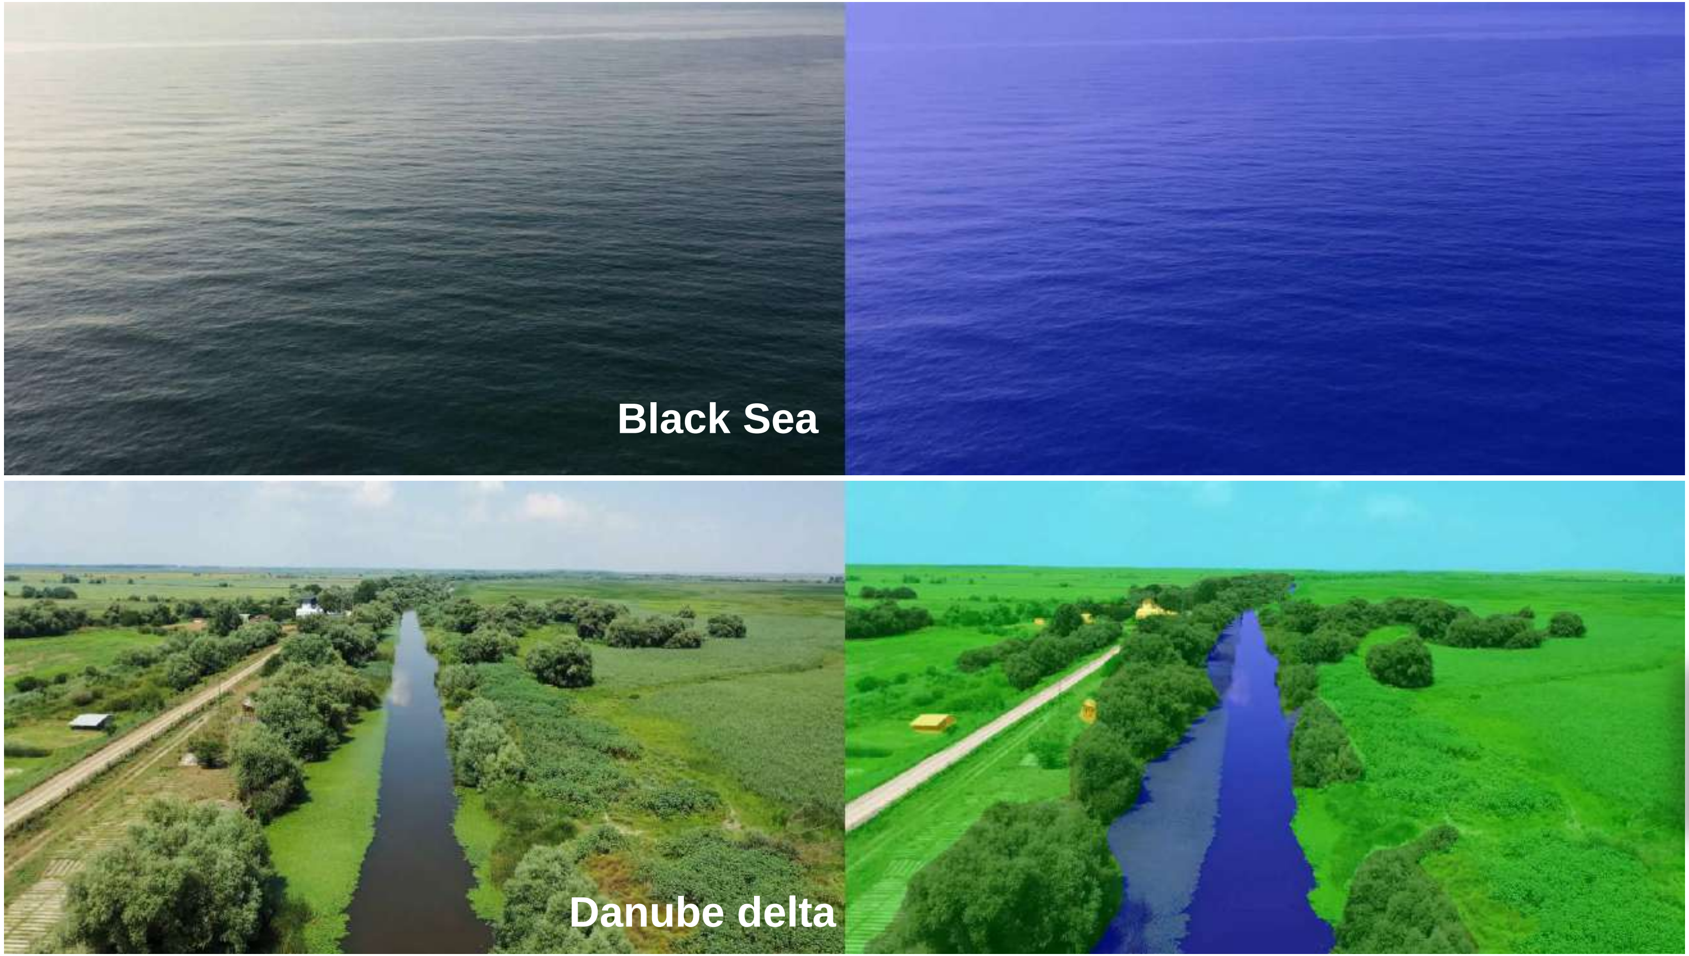

Only 1.6% of of frames are manually labeled

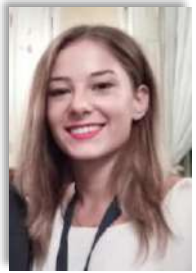

Alina Danciu

Alina Marcu

Towards total scene multi-view understanding in space and time

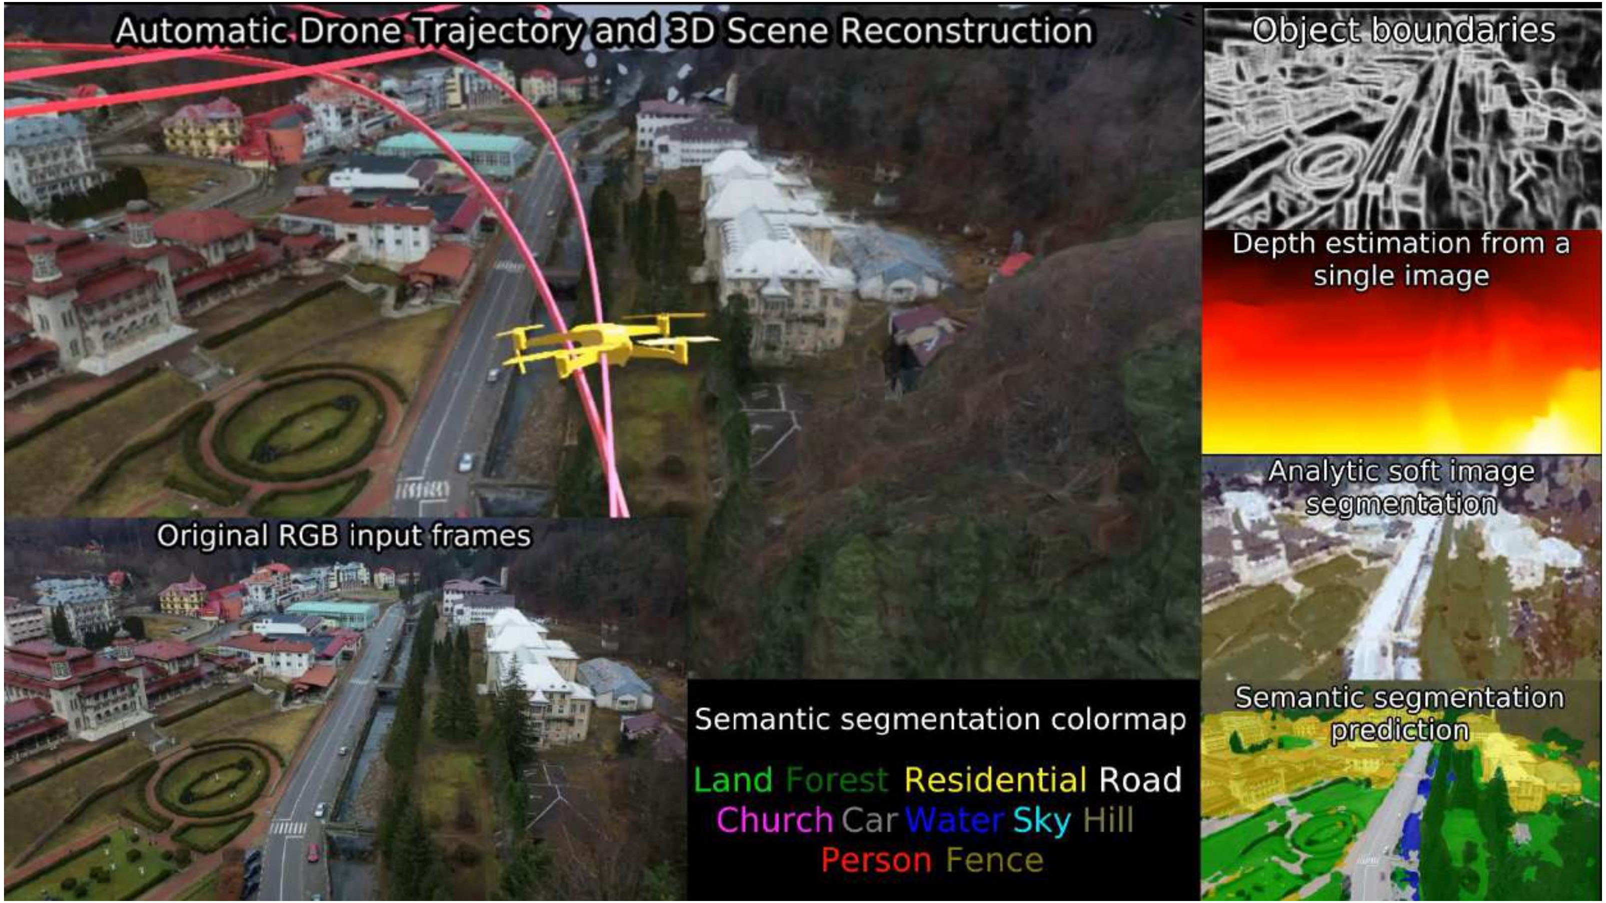

Edges and soft segmentation: [Leordeanu et al, TPAMI 2014; Leordeanu et al, ECCV 2012]
Unsupervised metric monocular depth estimation: [Licaret et al, ICRA 2022; Pirvu et al, CVPRW 2021]

Self-supervised learning with multi-view hypergraph consensus in the real world

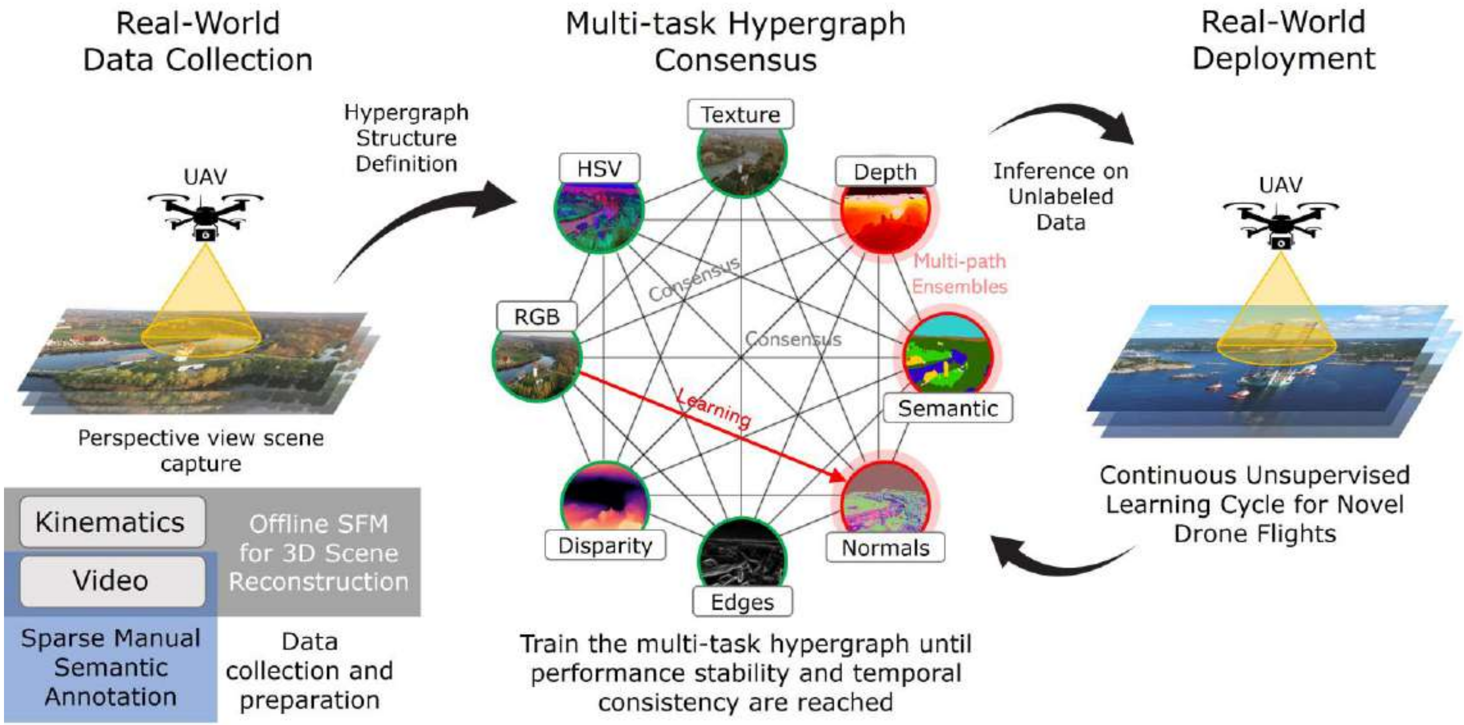

Alina Marcu

Mihai Pirvu

Dragos Costea

Emanuela Haller

Marcu, A, M. Pirvu, D. Costea, E. Haller, E. Slusanschi, AN Belbachir, R. Sukthankar and M. Leordeanu. **Self-supervised hypergraphs for learning multiple world interpretations." In Proceedings of the IEEE International Conference on Computer Vision**, 2023.

Types of Multi-modal Hyperedges

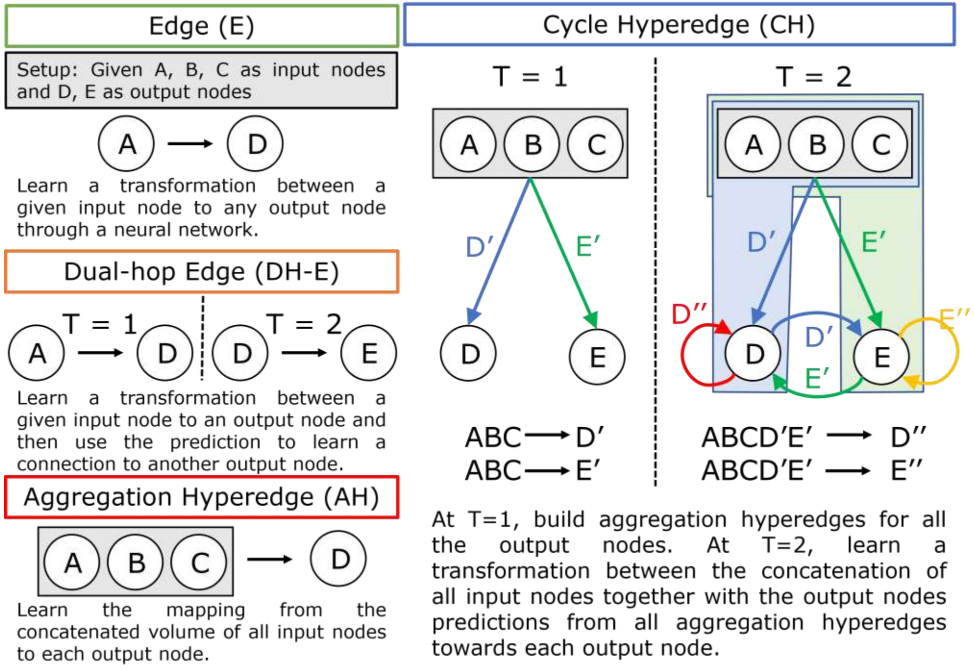

Learning Hyperedge Consensus

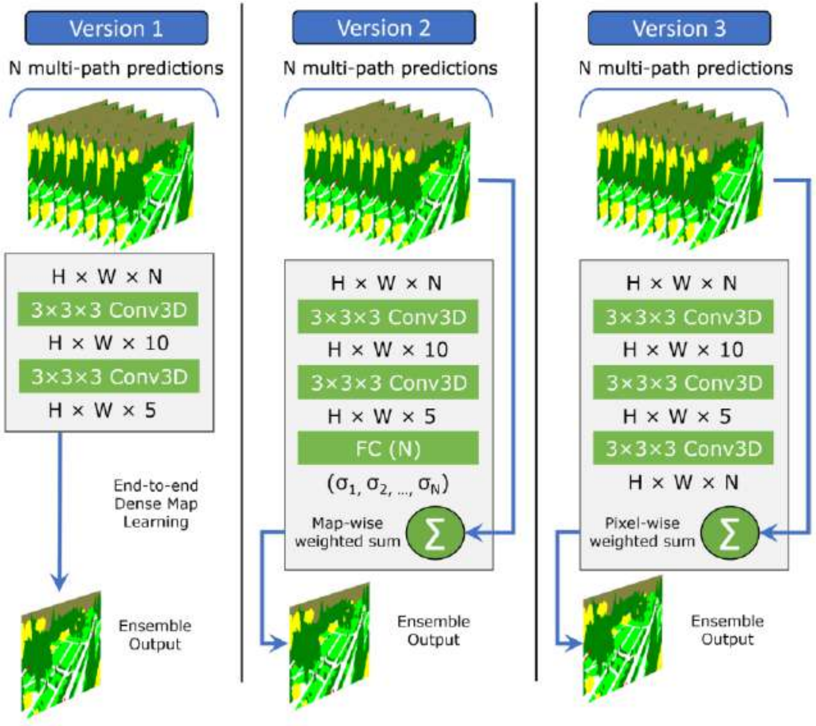

Improving accuracy and temporal consistency over iterations of self-supervised learning

Scene Name	Train Labeled	Train Unlabeled (iter 1)	Train Unlabeled (iter 2)	Train Unlabeled (iter 3)
Atanasie	76	4501 (76)	4500 (75)	-
Gradistei	18	726 (18)	484 (12)	-
Herculane	12	484 (12)	363 (9)	-
Jupiter	21	847 (21)	605 (15)	-
Olanesti	18	726 (18)	484 (12)	-
Petrova	12	484 (12)	363 (9)	-
Slanic	76	4501 (76)	4500 (75)	-
Barsana	-	-	-	1452 (36)
Comana	-	-	-	1210 (30)
Norway	-	-	-	2941 (50)
TOTAL	233	12269 (233)	11299 (207)	5603 (116)

Type	Semantic		Depth		Normals	
	IoU (↑)	Cons. (↑)	L1 (↓)	Cons. (↑)	L1 (↓)	Cons. (↑)
rgb-sup.	25.04	88.85	-	-	-	-
rgb-iter1	32.79	94.04	21.66	5.89	12.40	98.32
rgb-iter2	37.26	95.72	17.34	7.06	11.93	98.87
rgb-iter3	**40.31**	**98.13**	**16.64**	**30.26**	**11.71**	**99.30**

Input image

Mask2Former [Cheng et al, CVPR 2022]

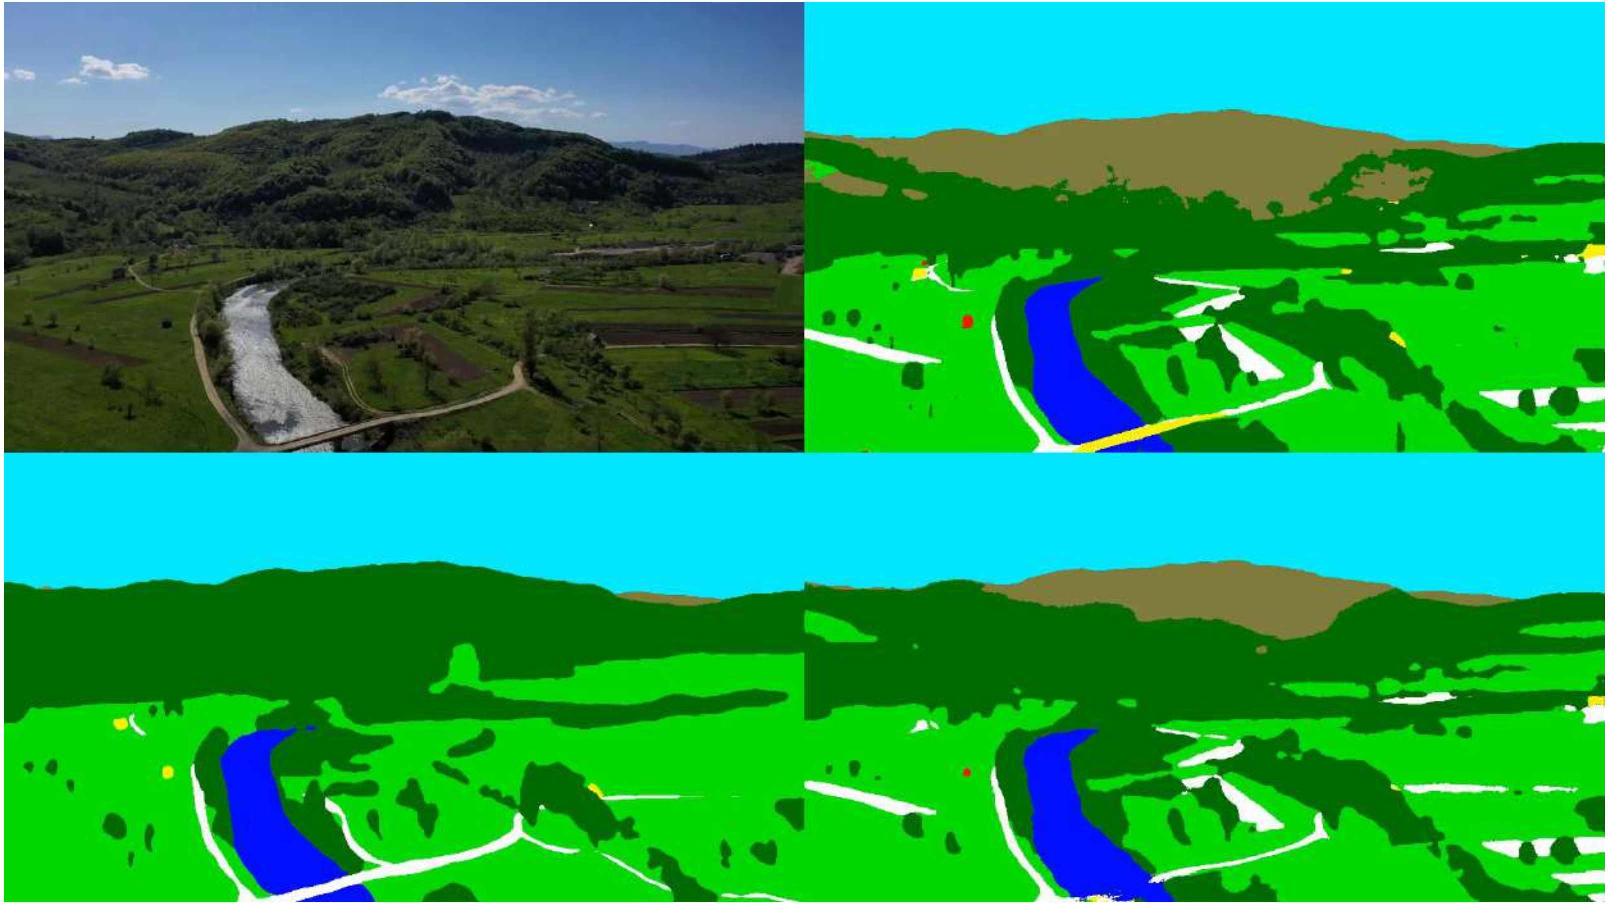

Human

Ours

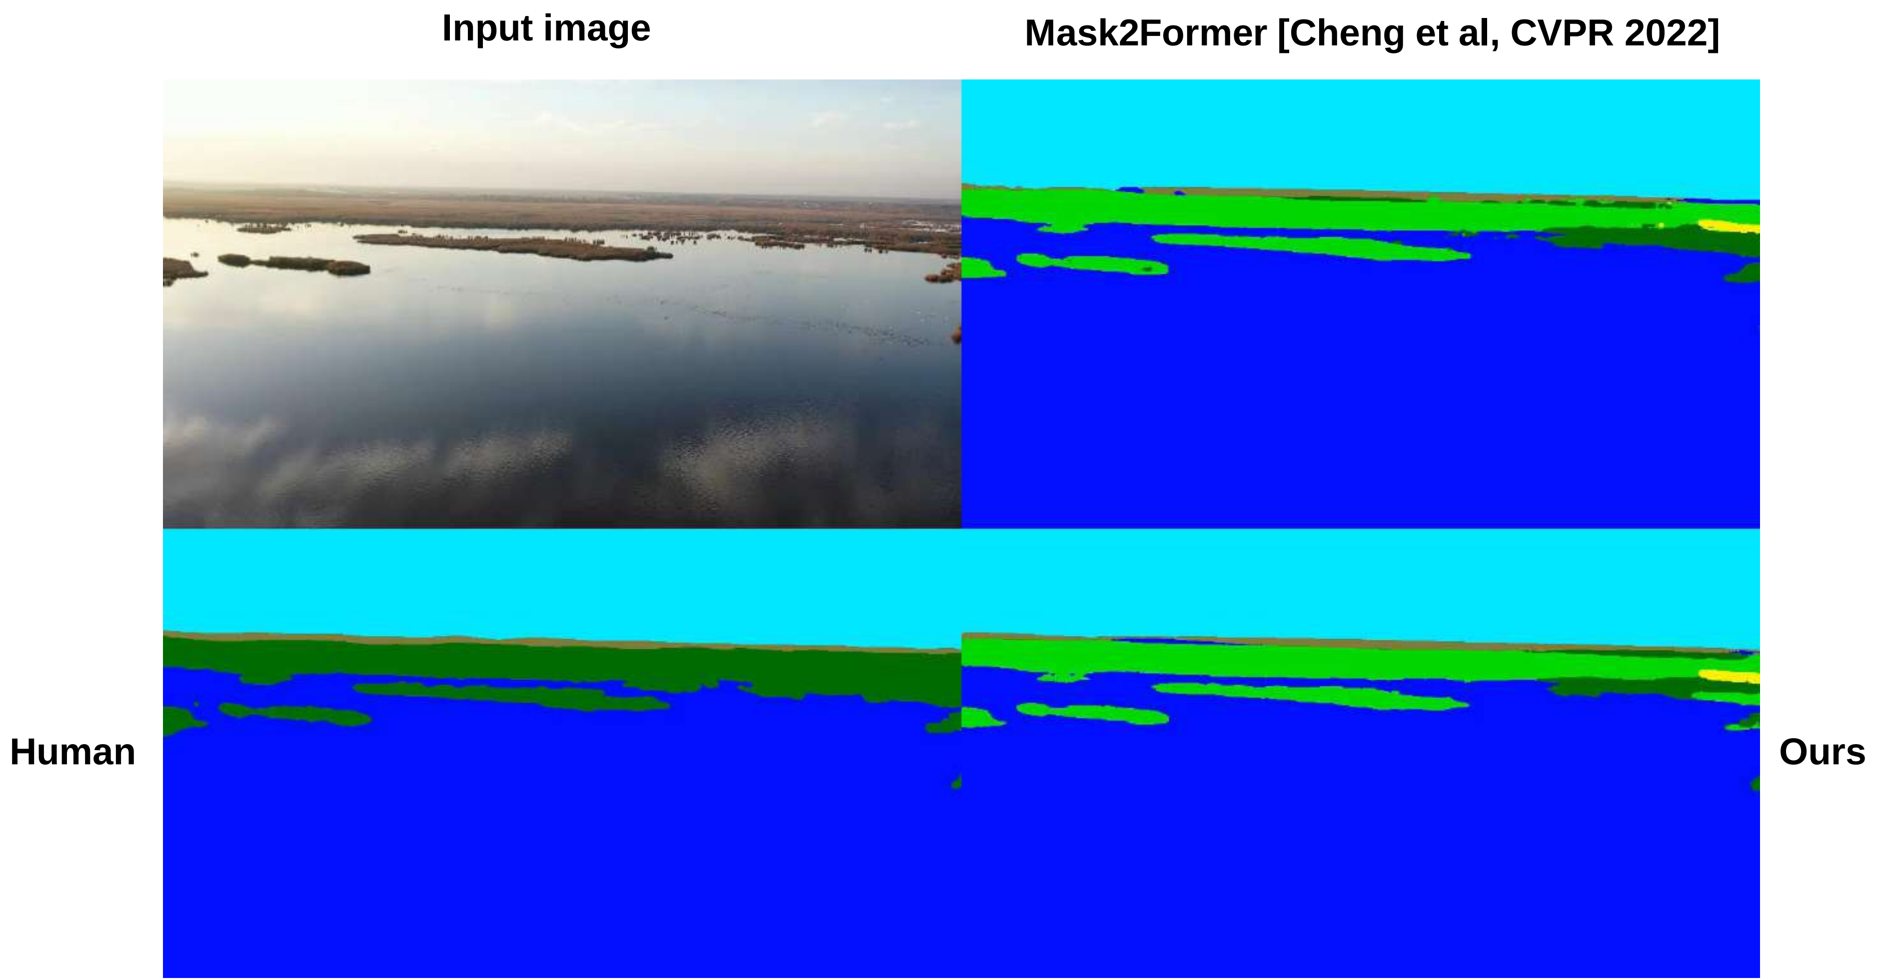
Input image
Mask2Former [Cheng et al, CVPR 2022]
Human
Ours

Input image | Mask2Former [Cheng et al, CVPR 2022]

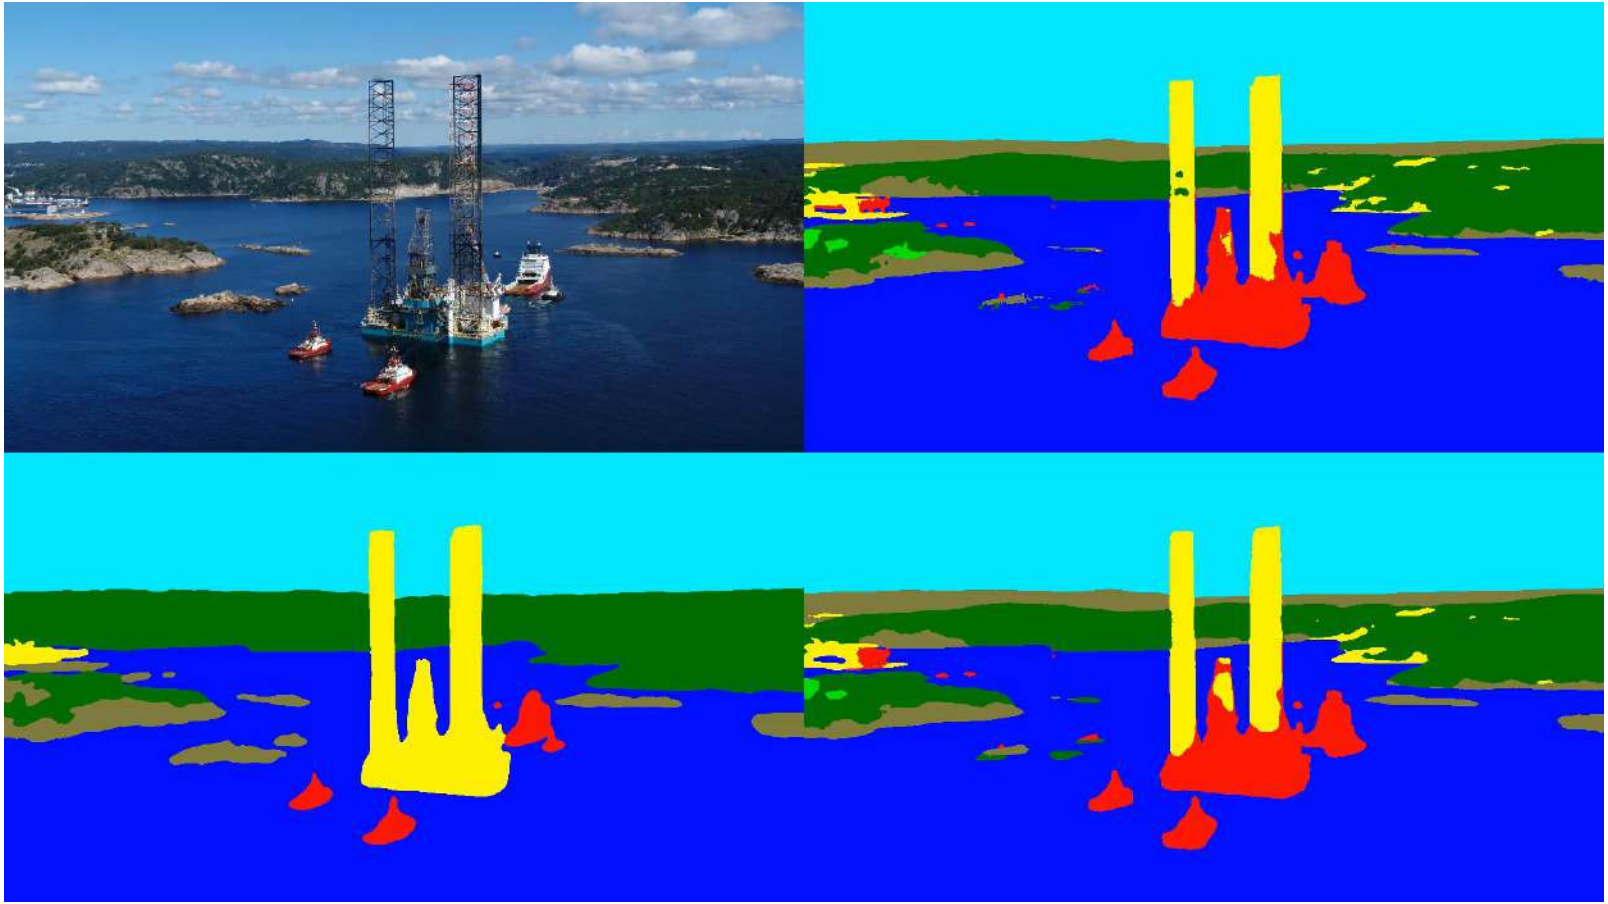

Human | Ours

Learning Multi-view Earth Consensus for Predicting Climate

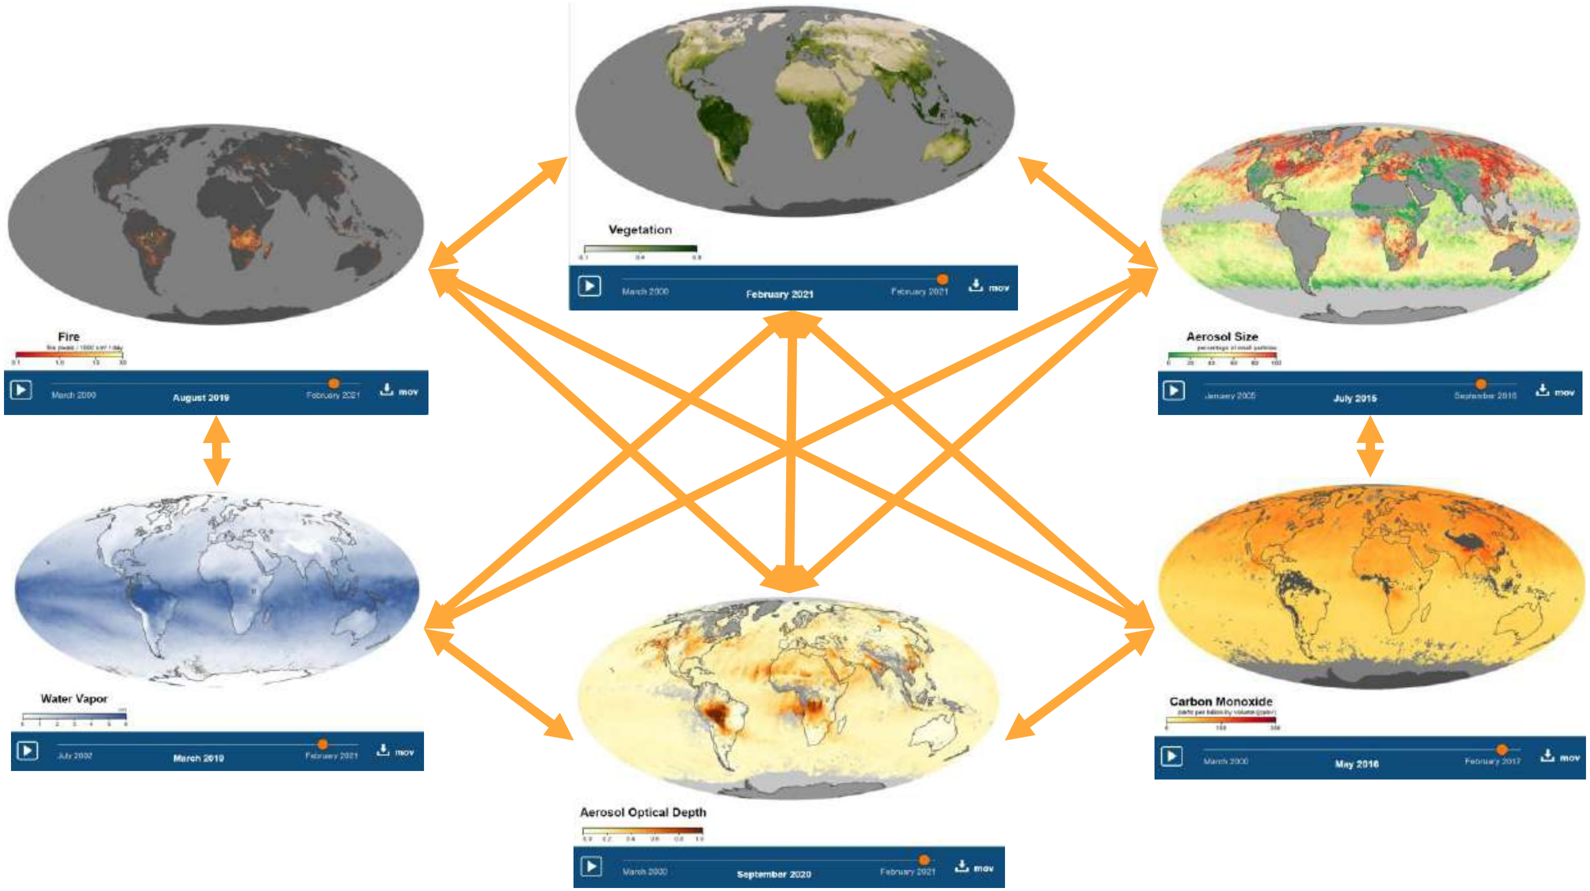

The Holocene period, representing the last 10 000 years (also marking the great development of human civilization), is a relatively stable period of the earth climate, in which different layers of the Earth System found a dynamic equilibrium (consensus) with average temperature varying less than 1 degree per year.

Consensus for Understanding Climate:
A Multi-modal Hypergraph Approach

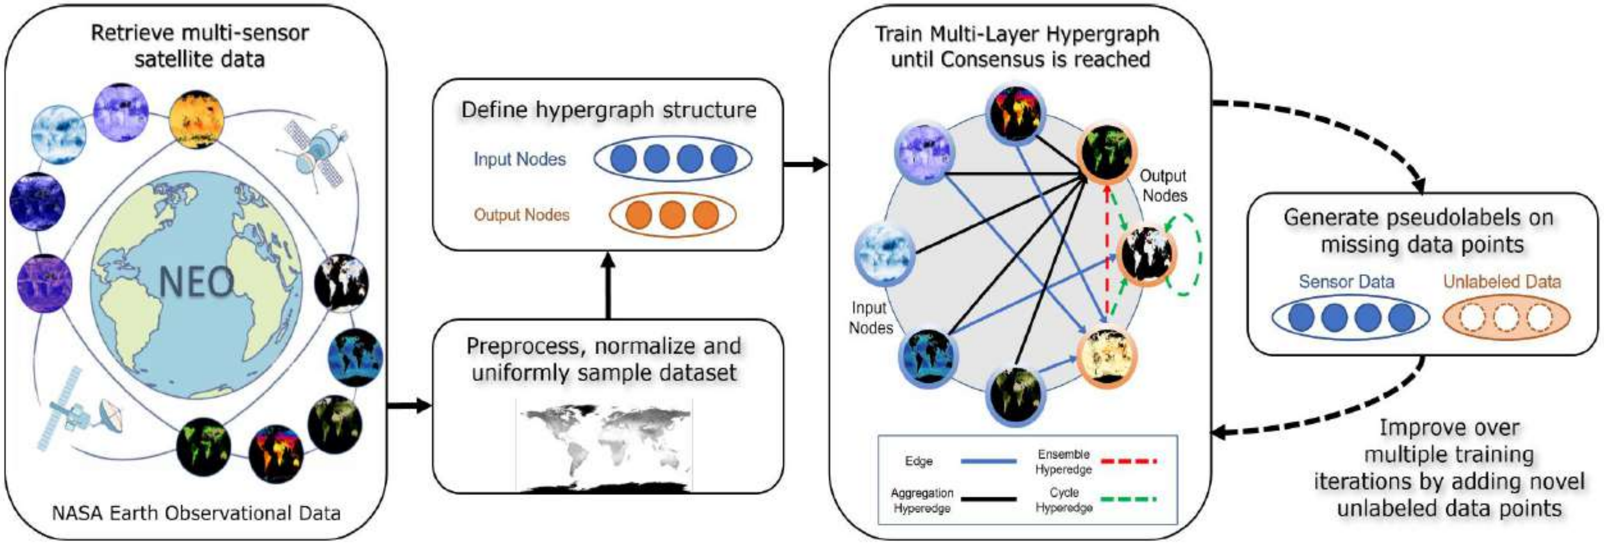

Mihai Pirvu

Alina Marcu

Alexandra Dobrescu

Pirvu, M., Alina M., A. Dobrescu, AN Belbachir and M. Leordeanu. "**Multi-Task Hypergraphs for Semi-supervised Learning using Earth Observations**." In Proceedings of the International Conference on Computer Vision, 2023.

Experiments on NASA NEO Dataset

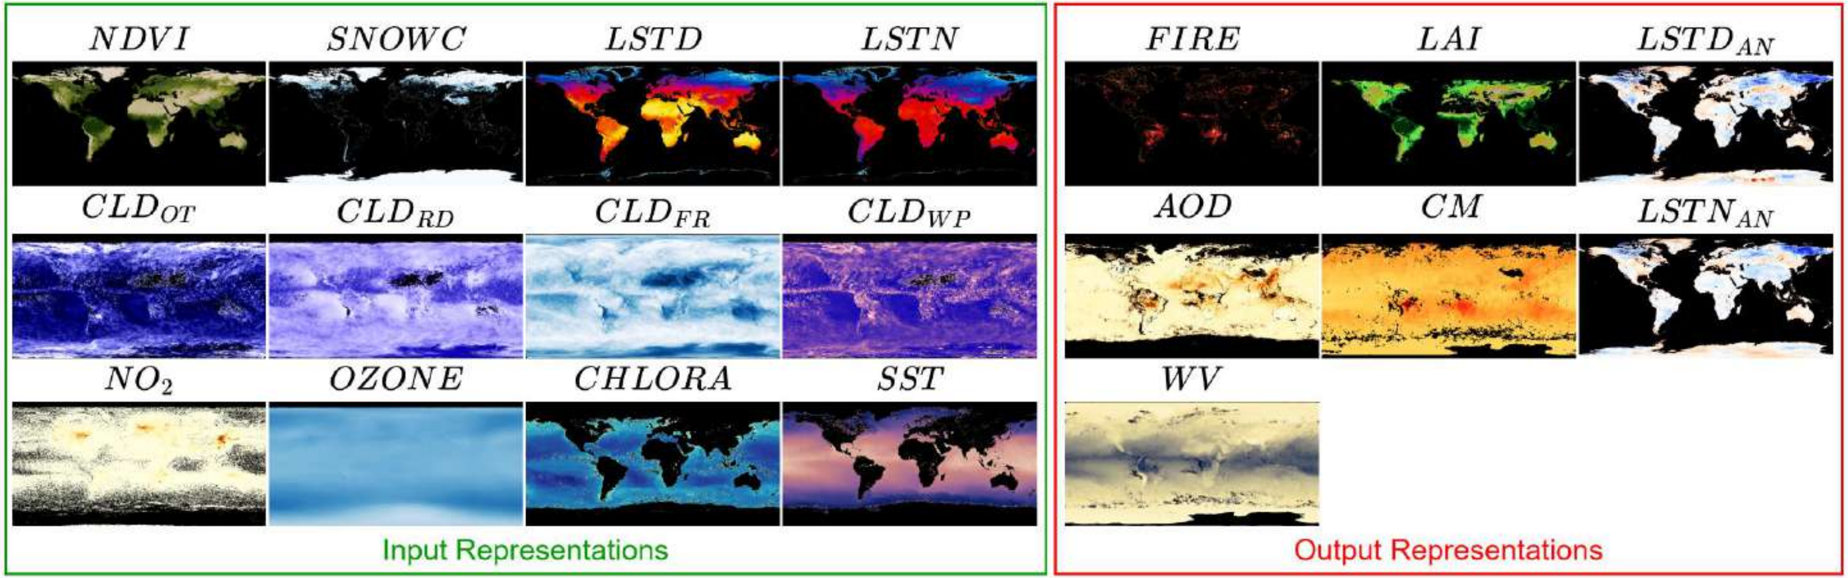

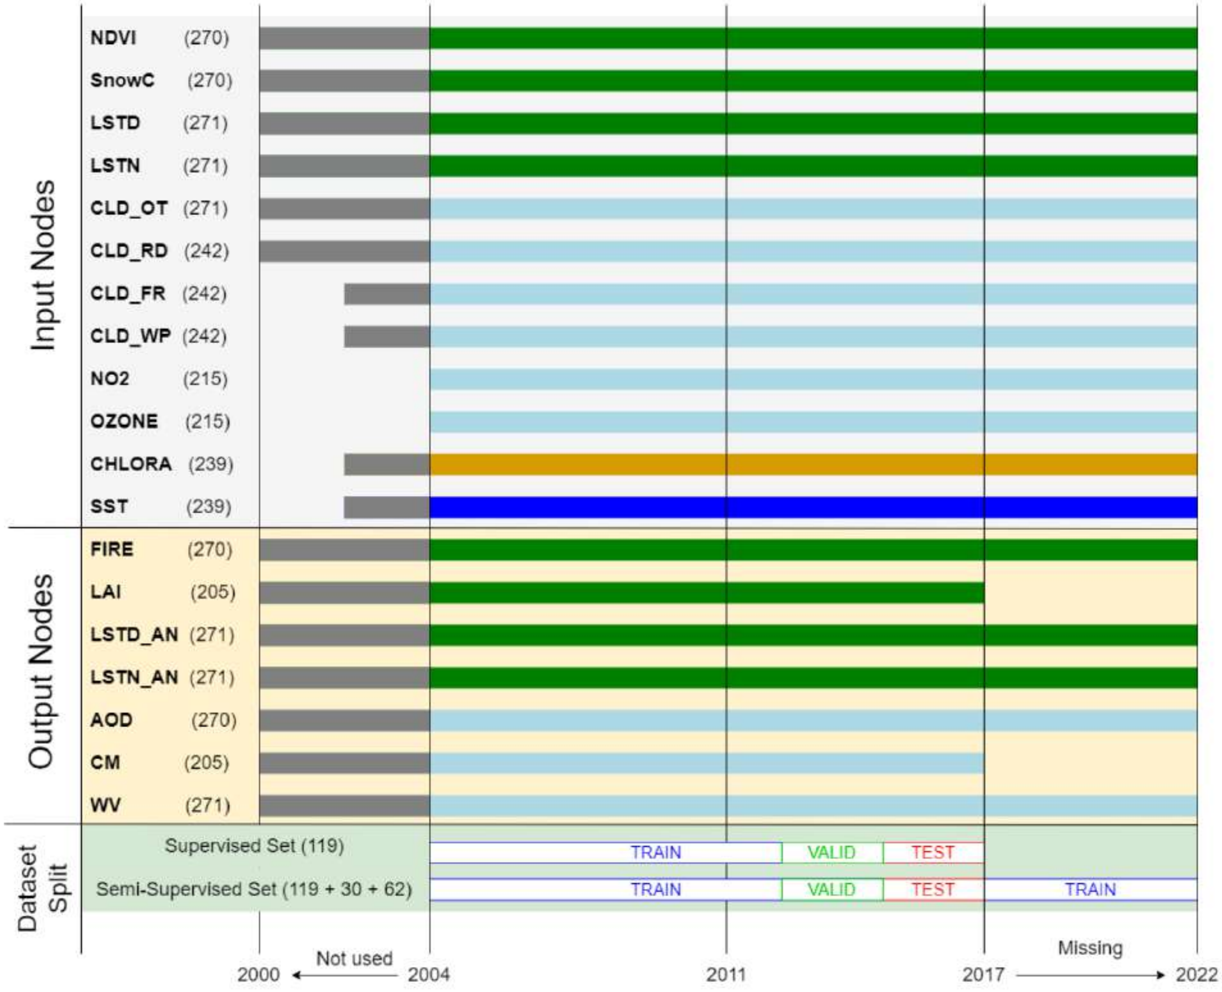
Input Nodes
NDVI (270)
SnowC (270)
LSTD (271)
LSTN (271)
CLD_OT (271)
CLD_RD (242)
CLD_FR (242)
CLD_WP (242)
NO2 (215)
OZONE (215)
CHLORA (239)
SST (239)
Output Nodes
FIRE (270)
LAI (205)
LSTD_AN (271)
LSTN_AN (271)
AOD (270)
CM (205)
WV (271)
Dataset Split
Supervised Set (119)
TRAIN
VALID
TEST
Semi-Supervised Set (119 + 30 + 62)
TRAIN
VALID
TEST
TRAIN
Not used
Missing
2000
2004
2011
2017
2022

Learning Consensus Ensembles and Self-Distillation over Iterations

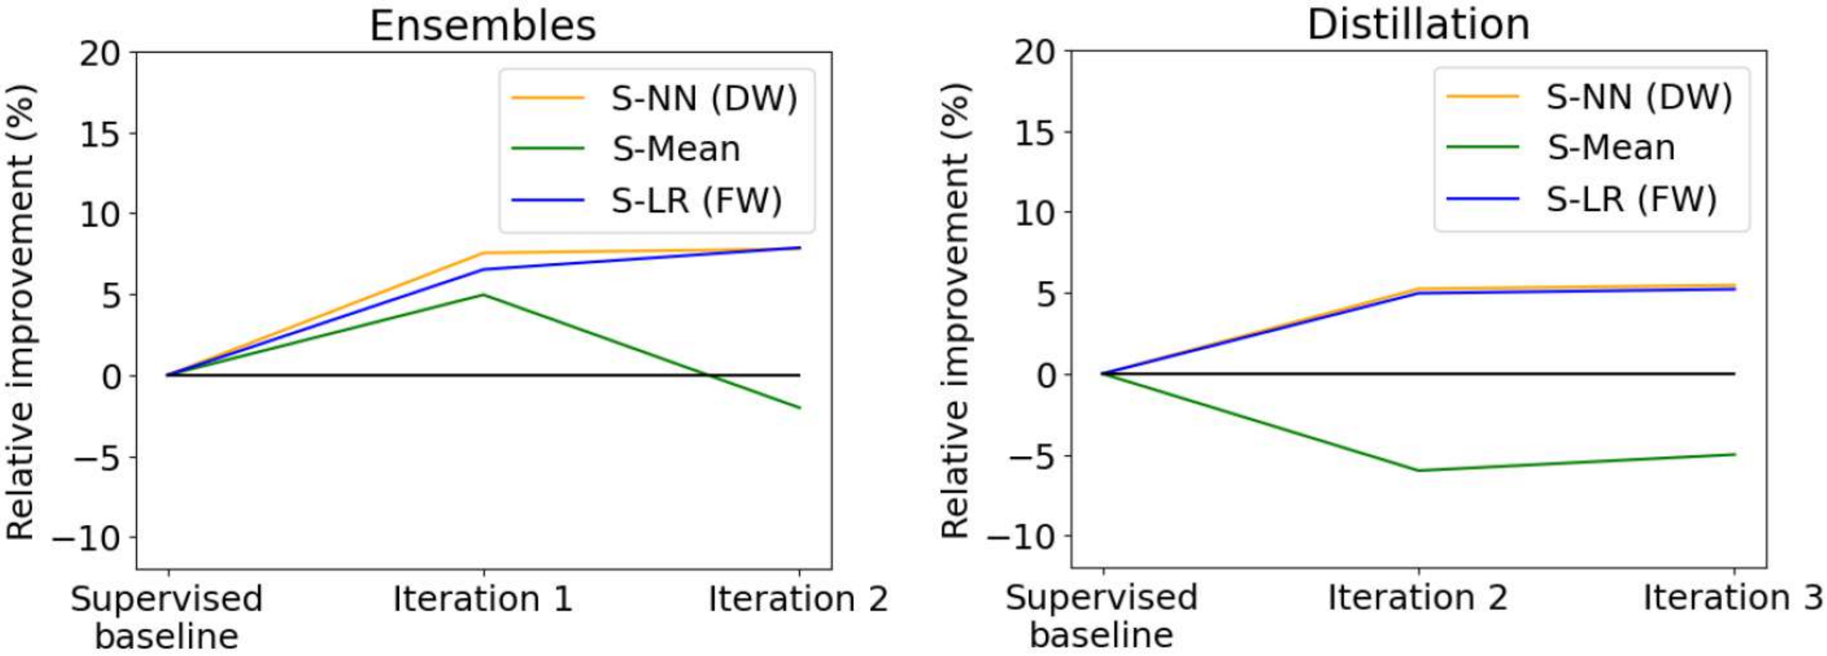

Self-supervised Learning Adapts to Climate Changes

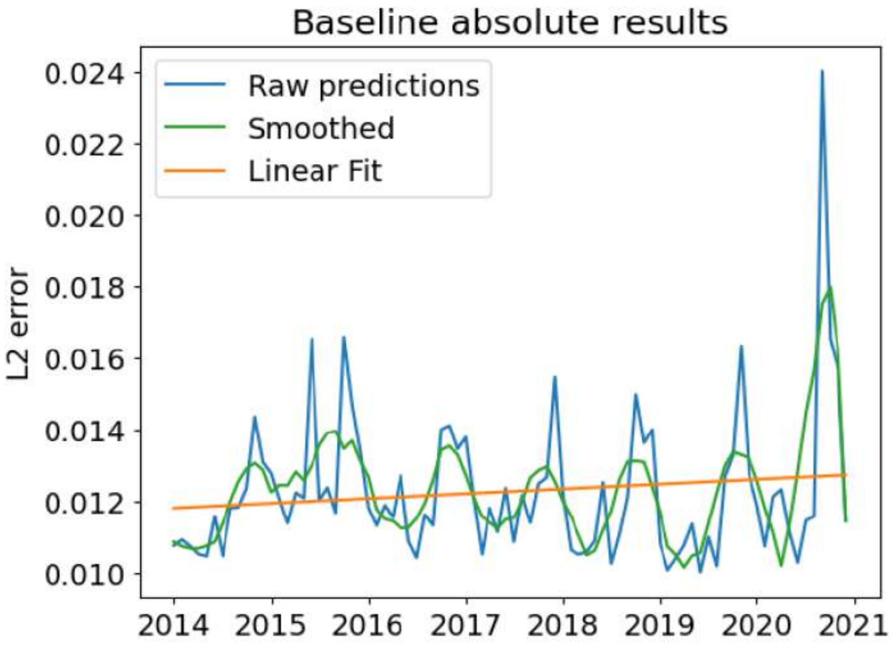

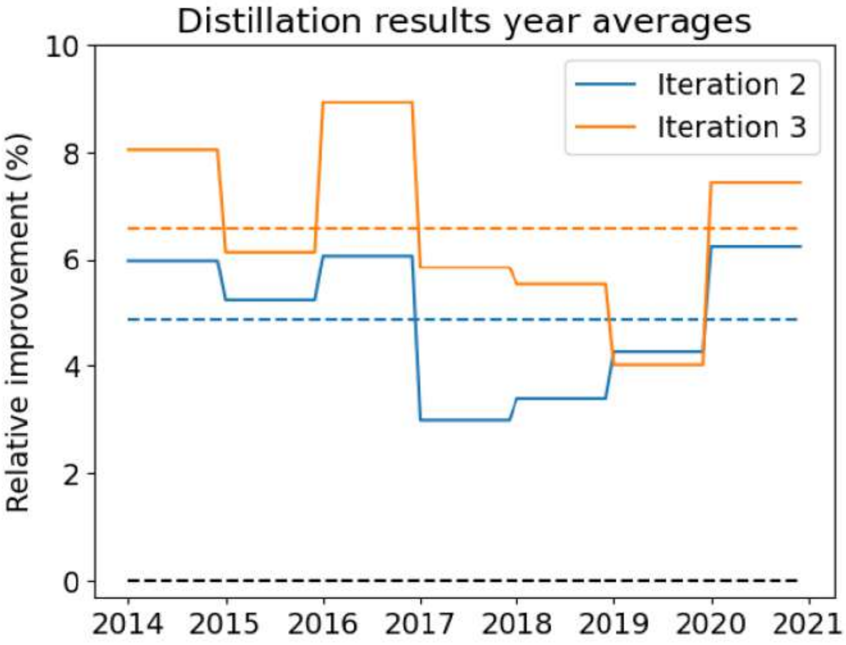

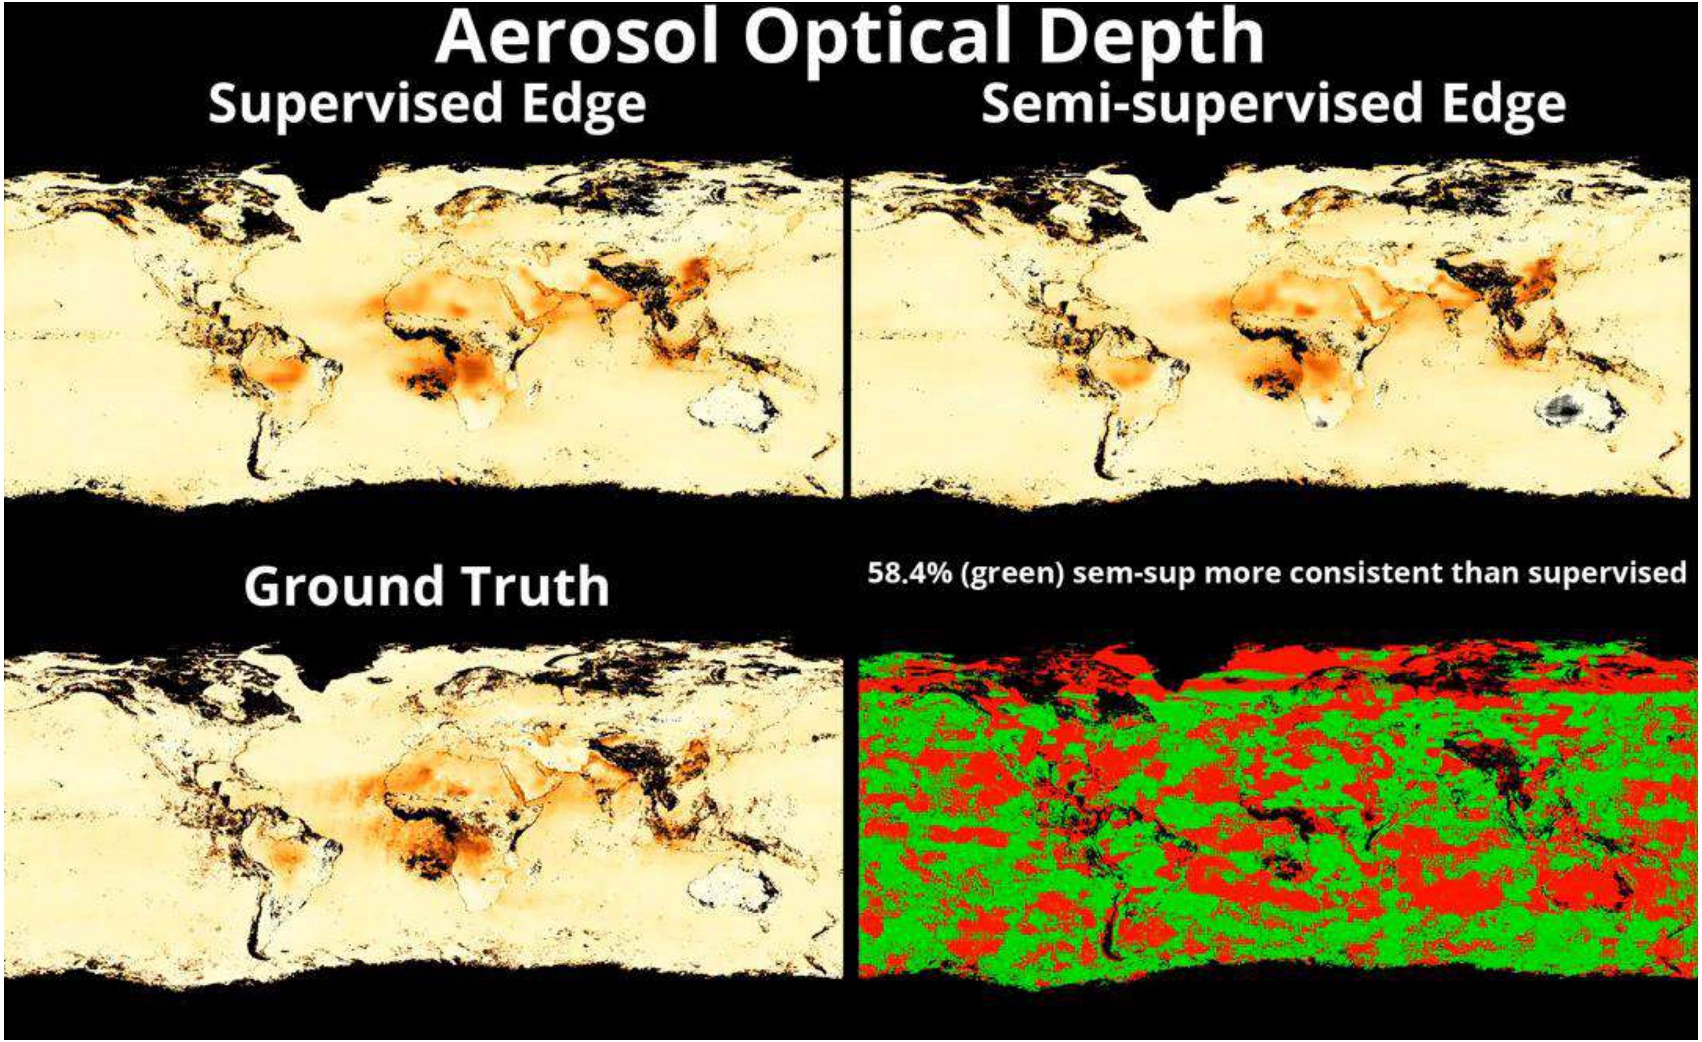
Aerosol Optical Depth
Supervised Edge
Semi-supervised Edge
Ground Truth
58.4% (green) sem-sup more consistent than supervised

Multi-view Hypergraph Consensus with
Multi-view Multiple Random Masking Auto-encoders

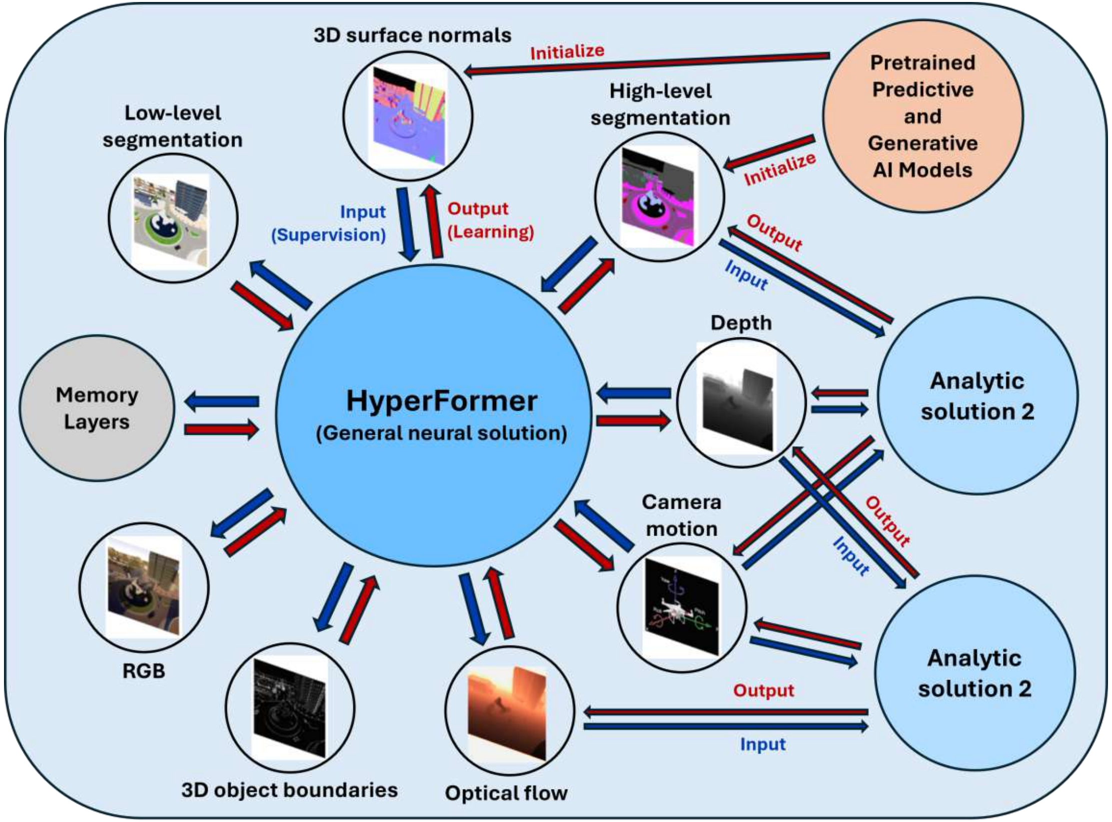

Mihai Pirvu

Introducing HyperFormer

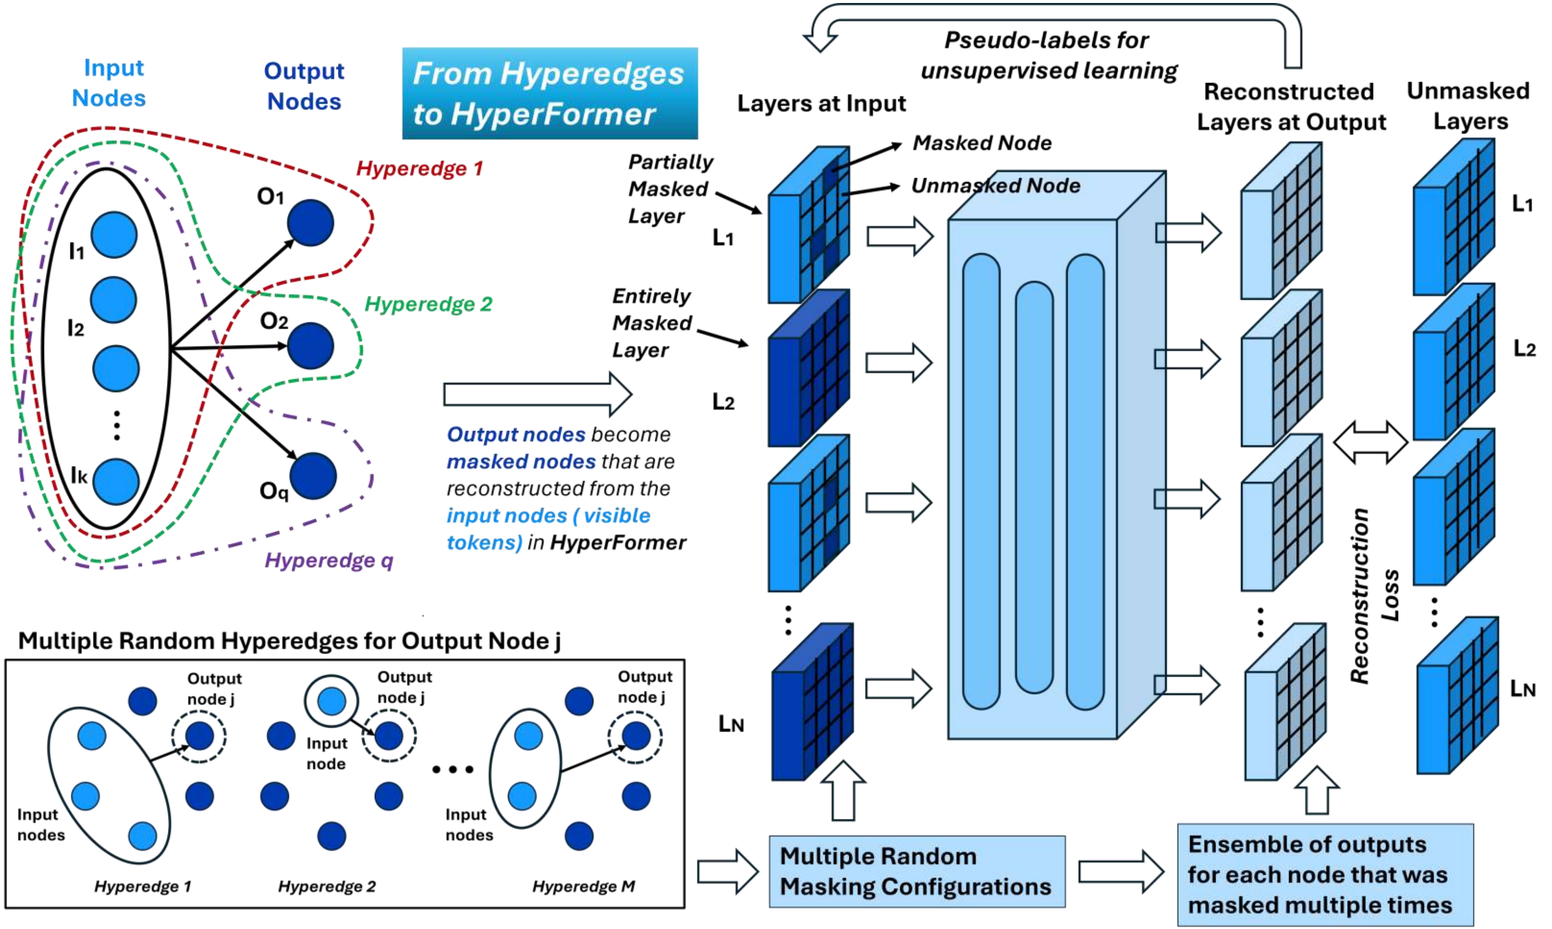

Creating Multiple Views from RGB

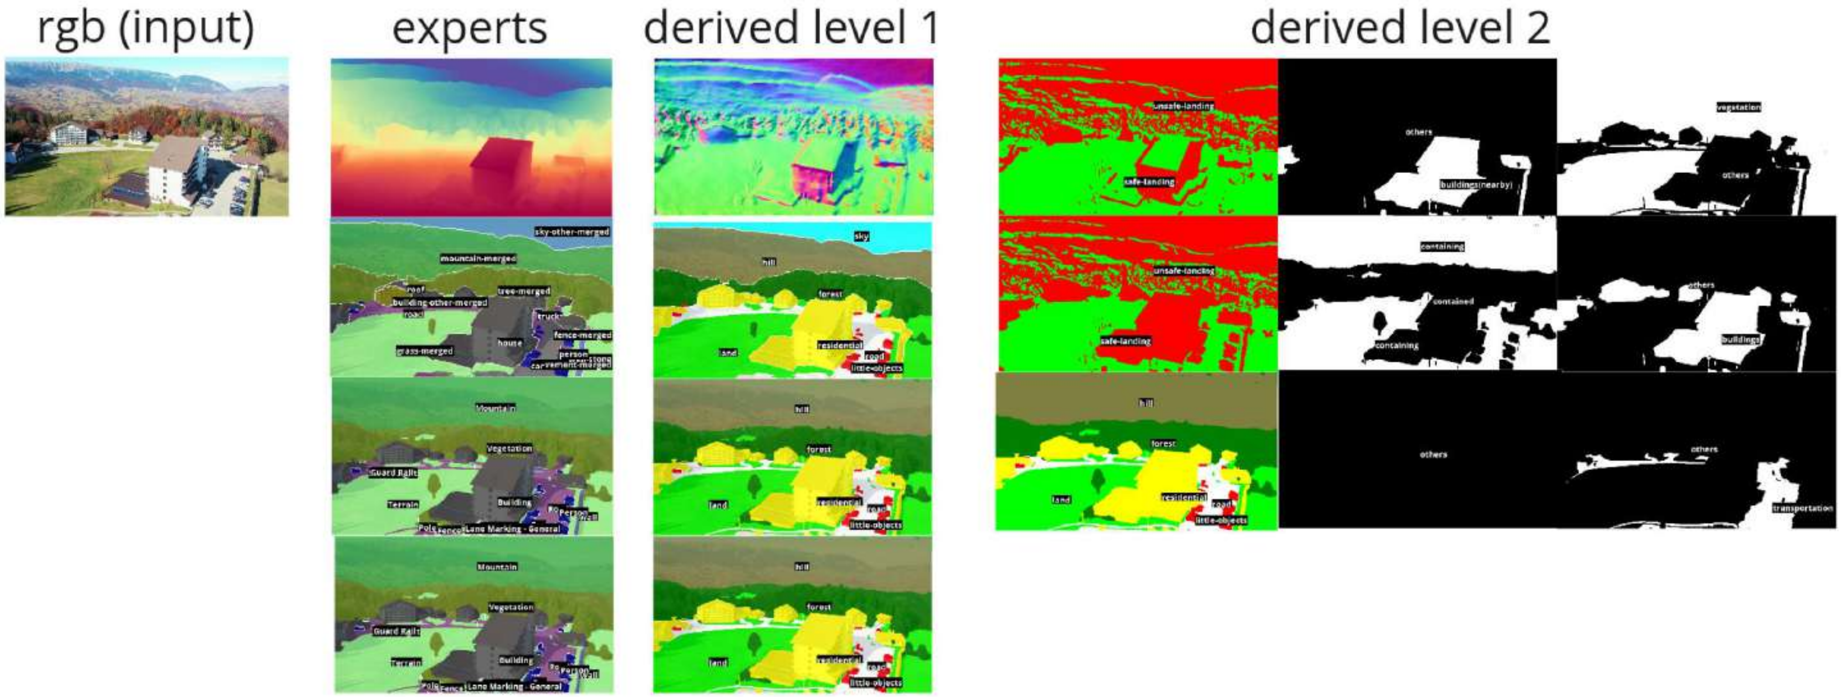

Experiments on extended DroneScapes++

Added **16 novel scenes** to initial **DronesScapes** with 10 scenes

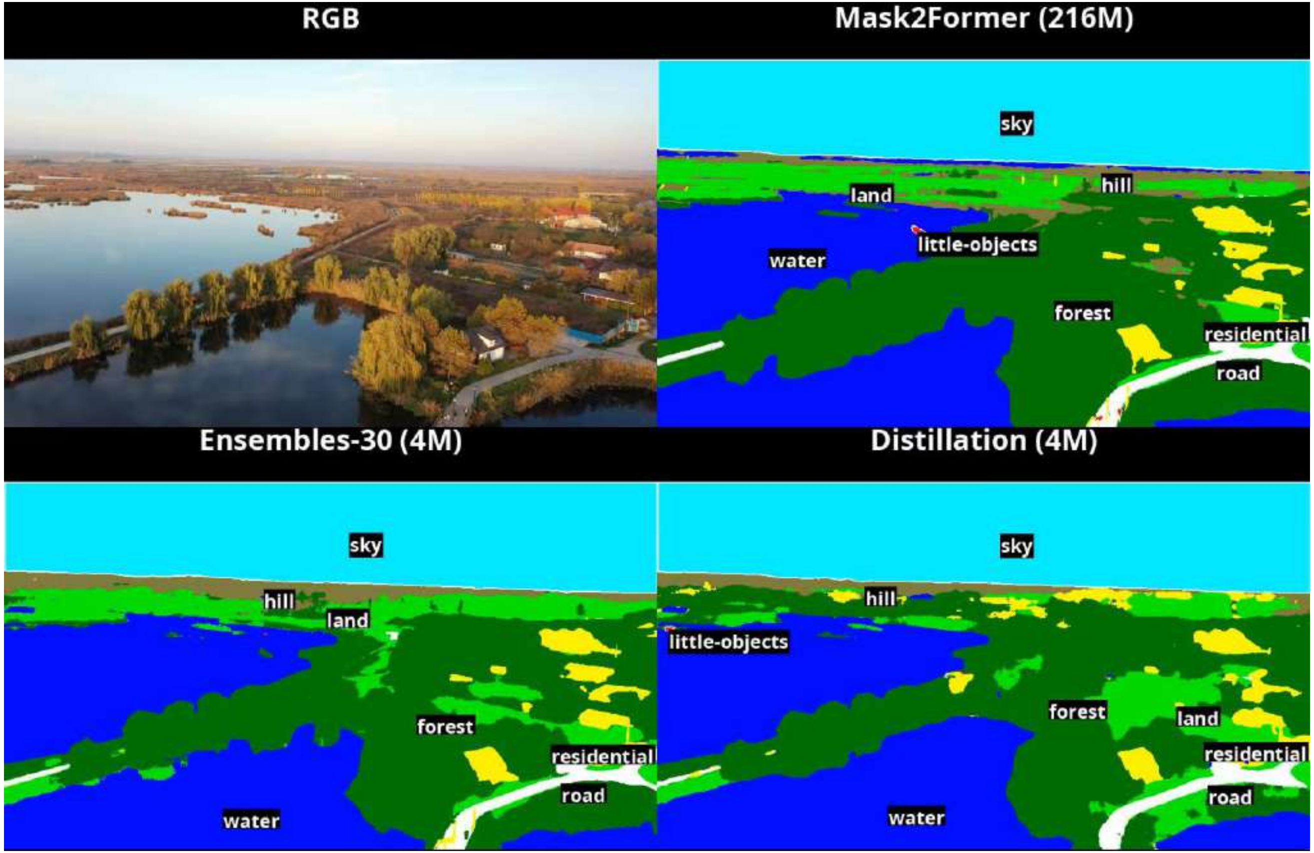

Comana scene

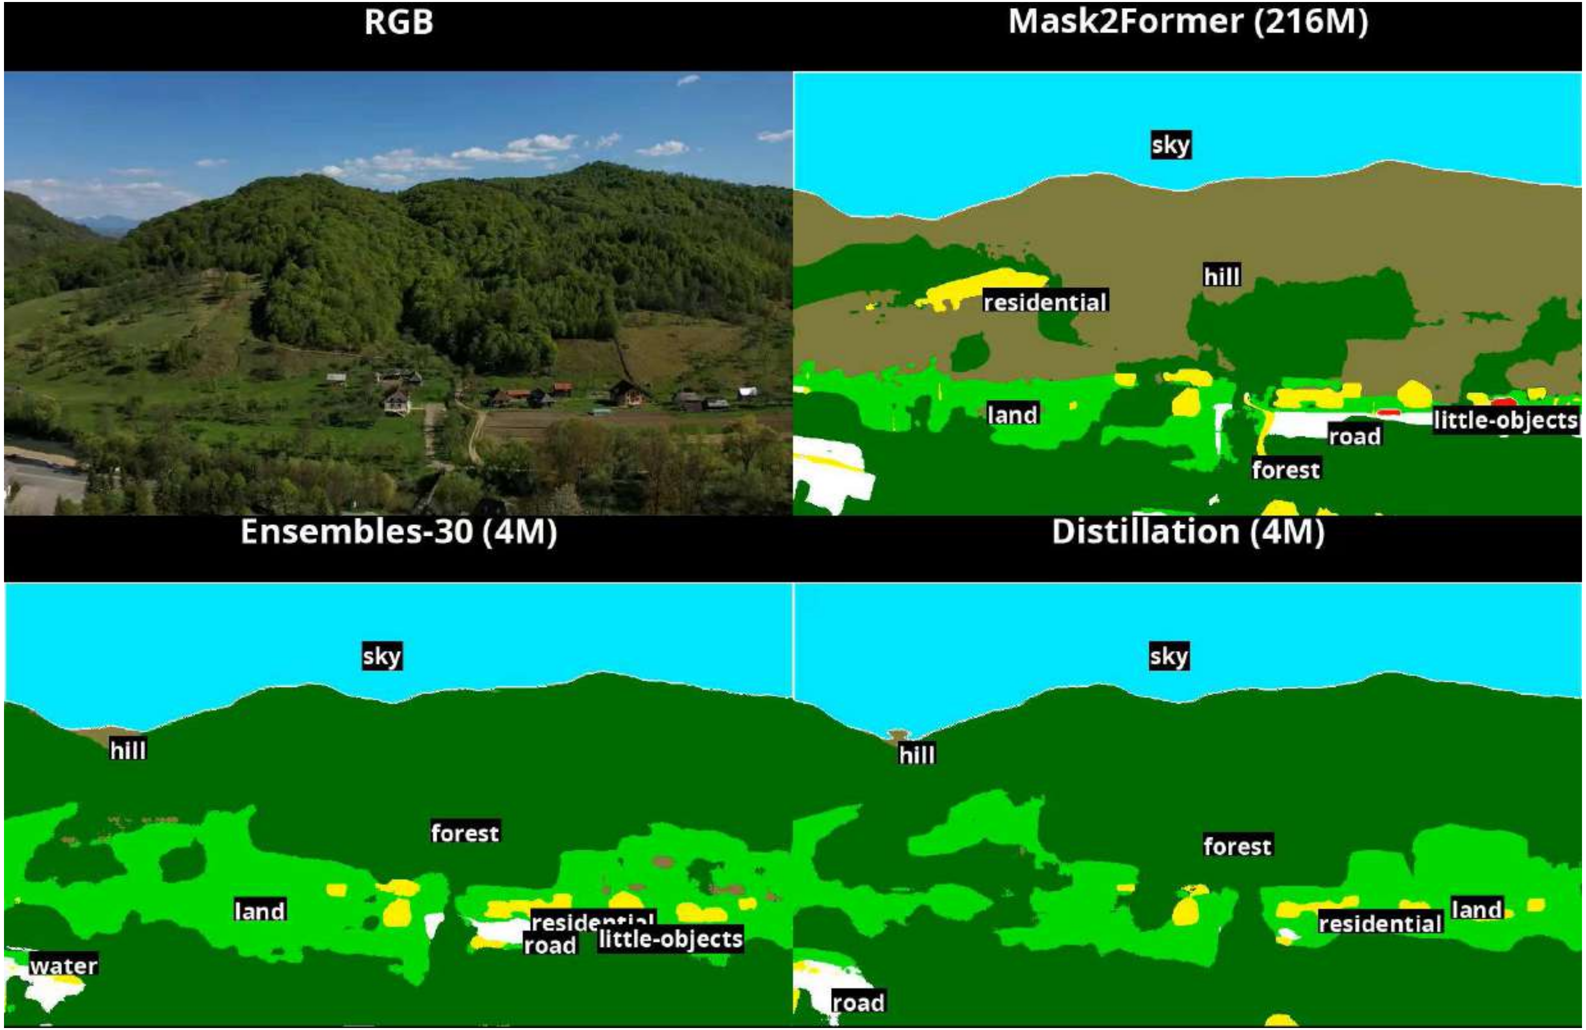

Barsana scene

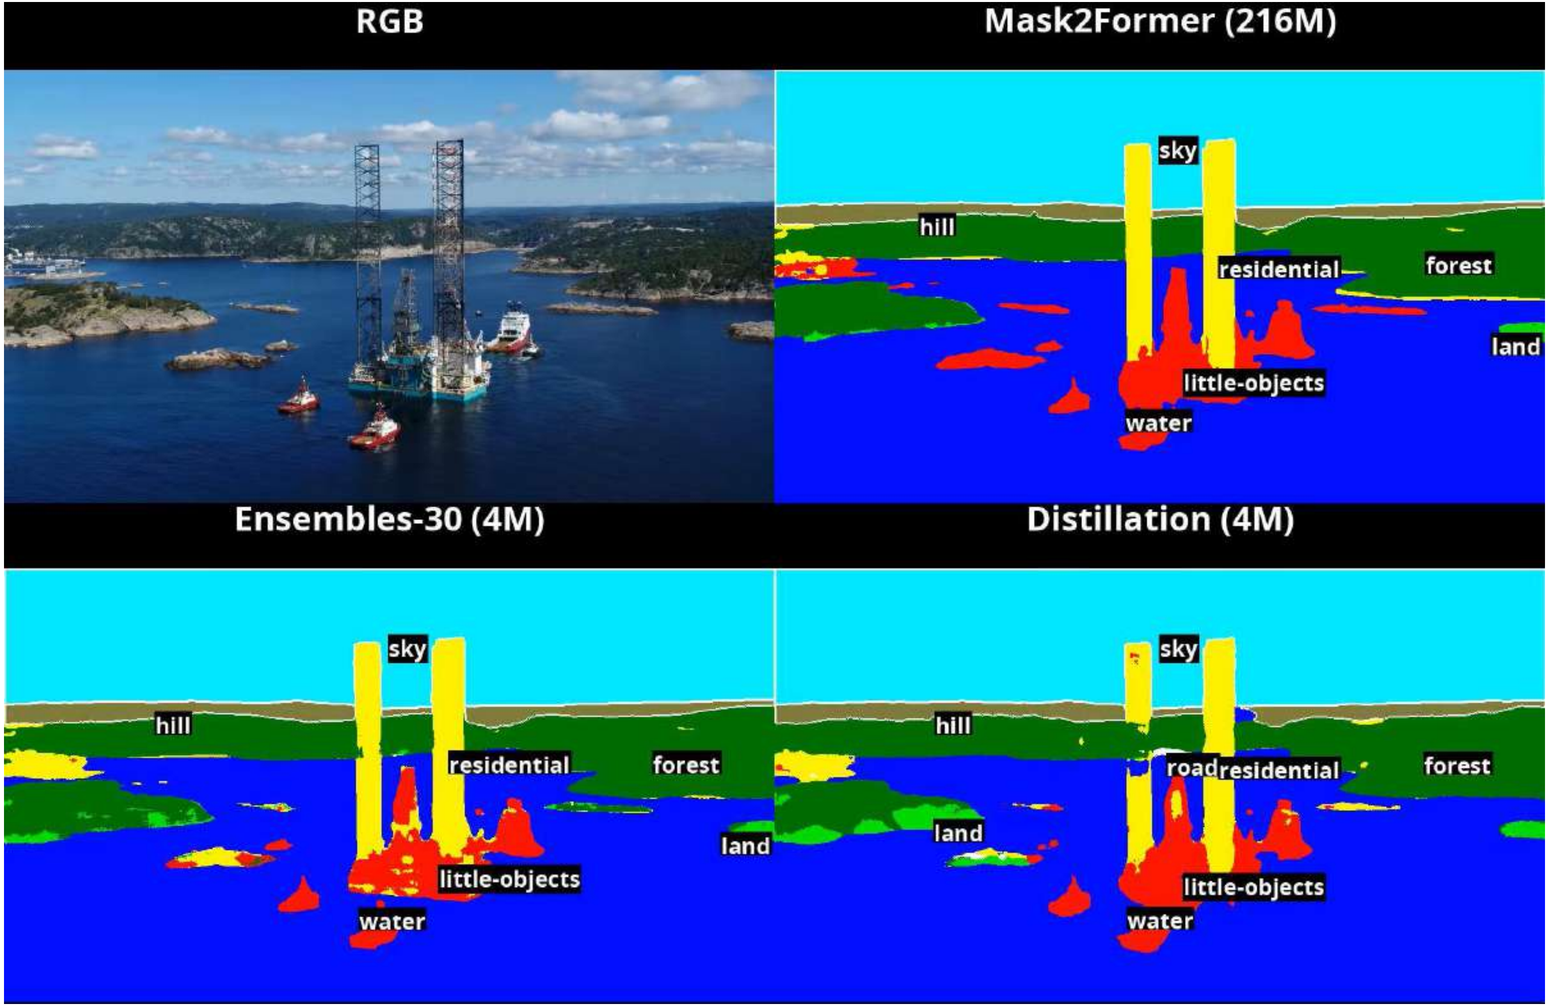

Norway scene

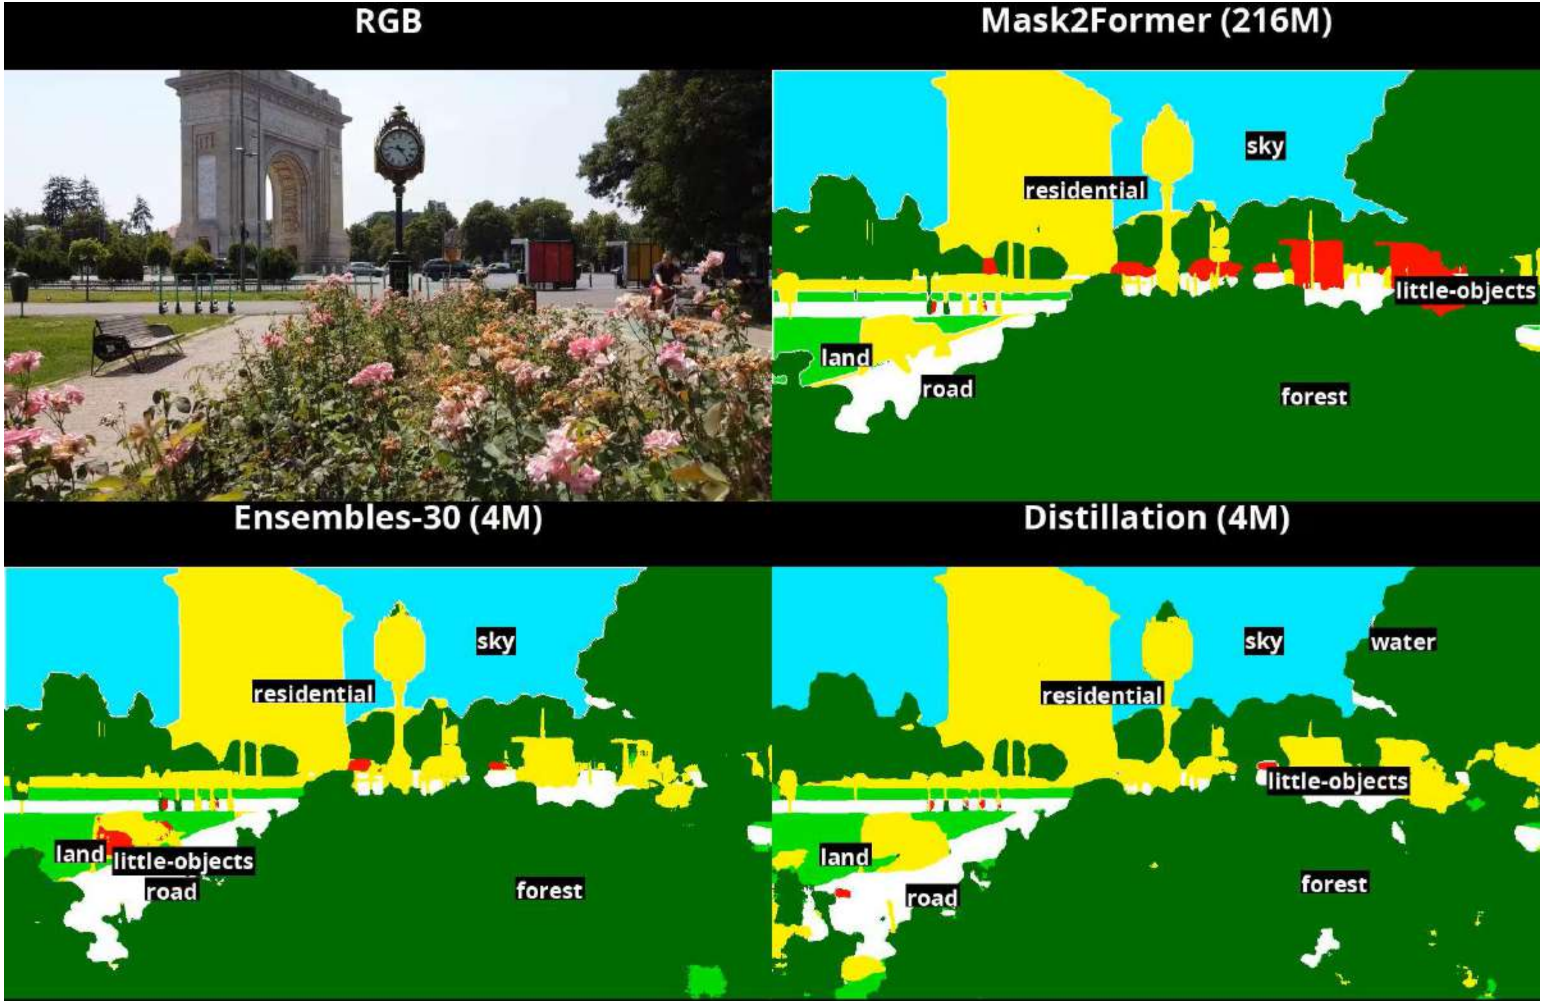

Bucharest scene

Self-supervised Multi-view Consensus improves both accuracy and consistency over time

Mean IoU	Barsana	Comana	Norway	Mean
NGC-LR ICCV-23 (32M)	46.51	45.59	30.17	40.76
Mask2Former (217M)	63.371	60.559	37.986	53.97
Ensembles-30 (4M)	**63.806**	**63.186**	**39.981**	**55.324**
Distillation (4M)	66.349	61.112	37.696	55.052

Consistency (%)	Barsana	Comana	Norway	Castelul Corvinilor	Bucharest Youtube	Mean
Mask2Former (217M)	93.97	94.68	98.42	98.37	96.79	96.44
Ensembles-30 (4M)	93.94	96.12	98.08	98.05	97.26	96.69
Distillation (4M)	**98.16**	**98.04**	**99.02**	**99.13**	**98.15**	**98.5**

Developing Self-flying Drones in Our SpaceTime Vision and Robotics Lab

Music: “Fly to the Island” (Album: “Supersonic”) , composed, performed and co-produced by Marius Leordeanu

Key ideas: Learning semantic segmentation with minimal user labeling and Unsupervised computation of metric depth

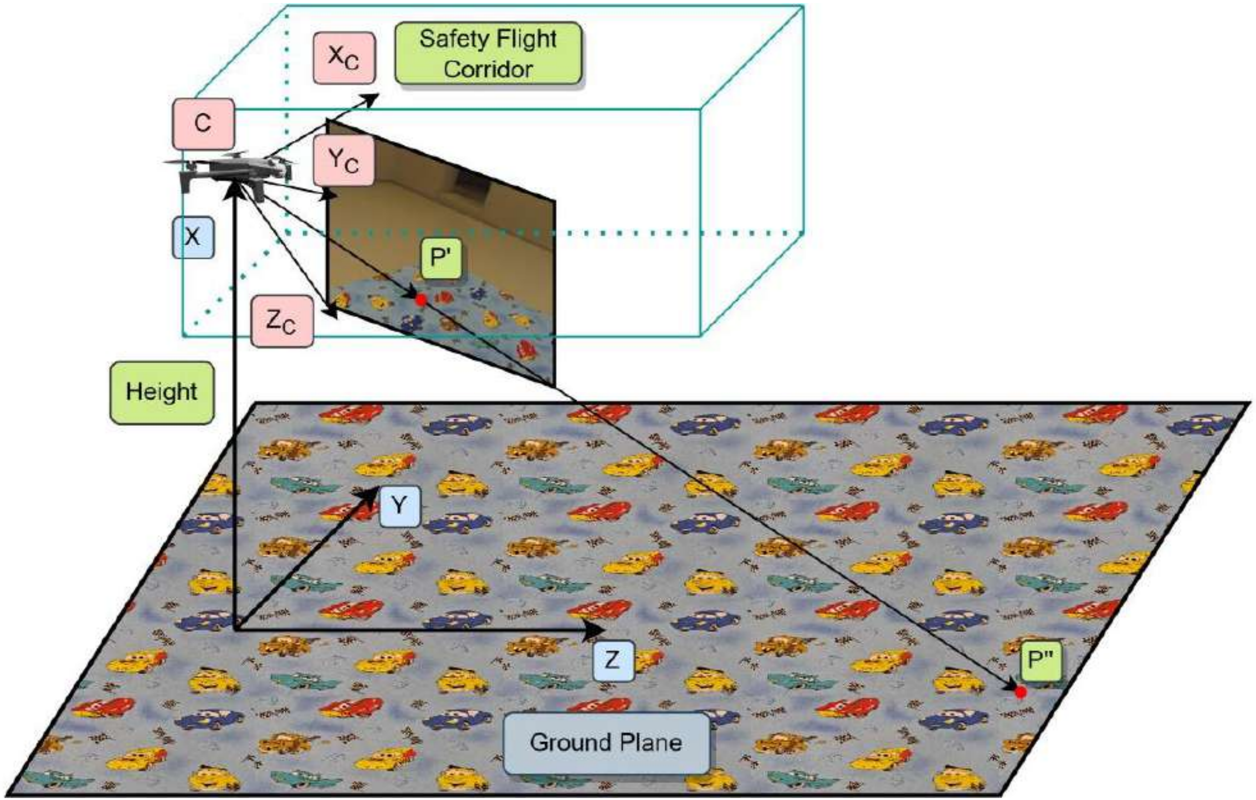

Learn and discover in the twin virtual scene – Apply in the real scene

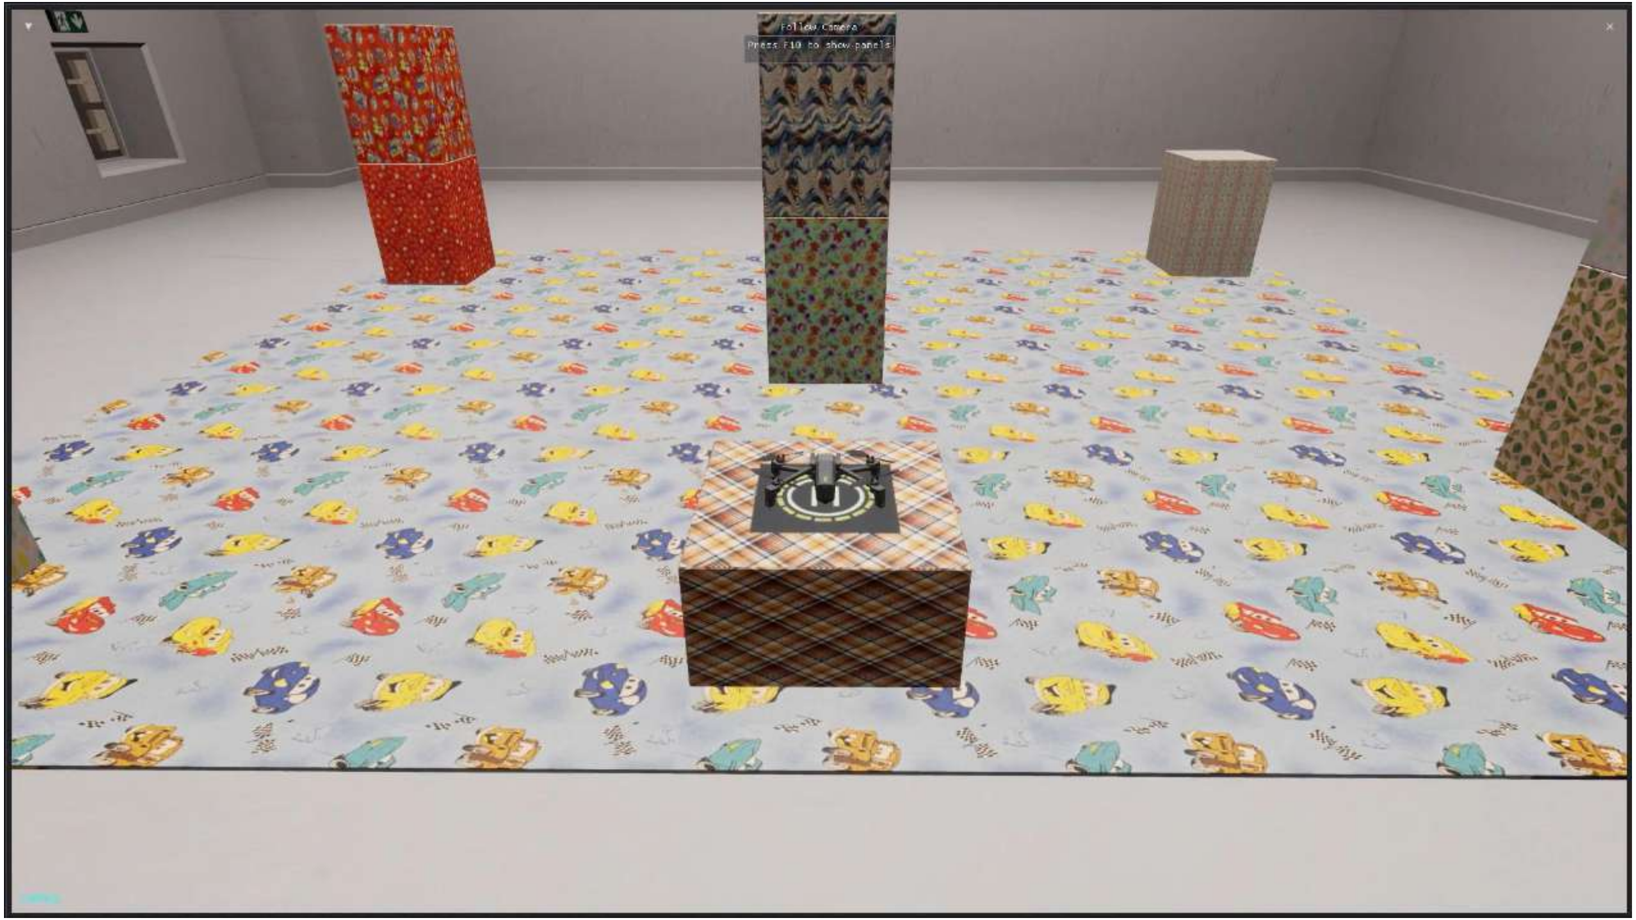

Consistency between tests in Real vs Virtual Scenes

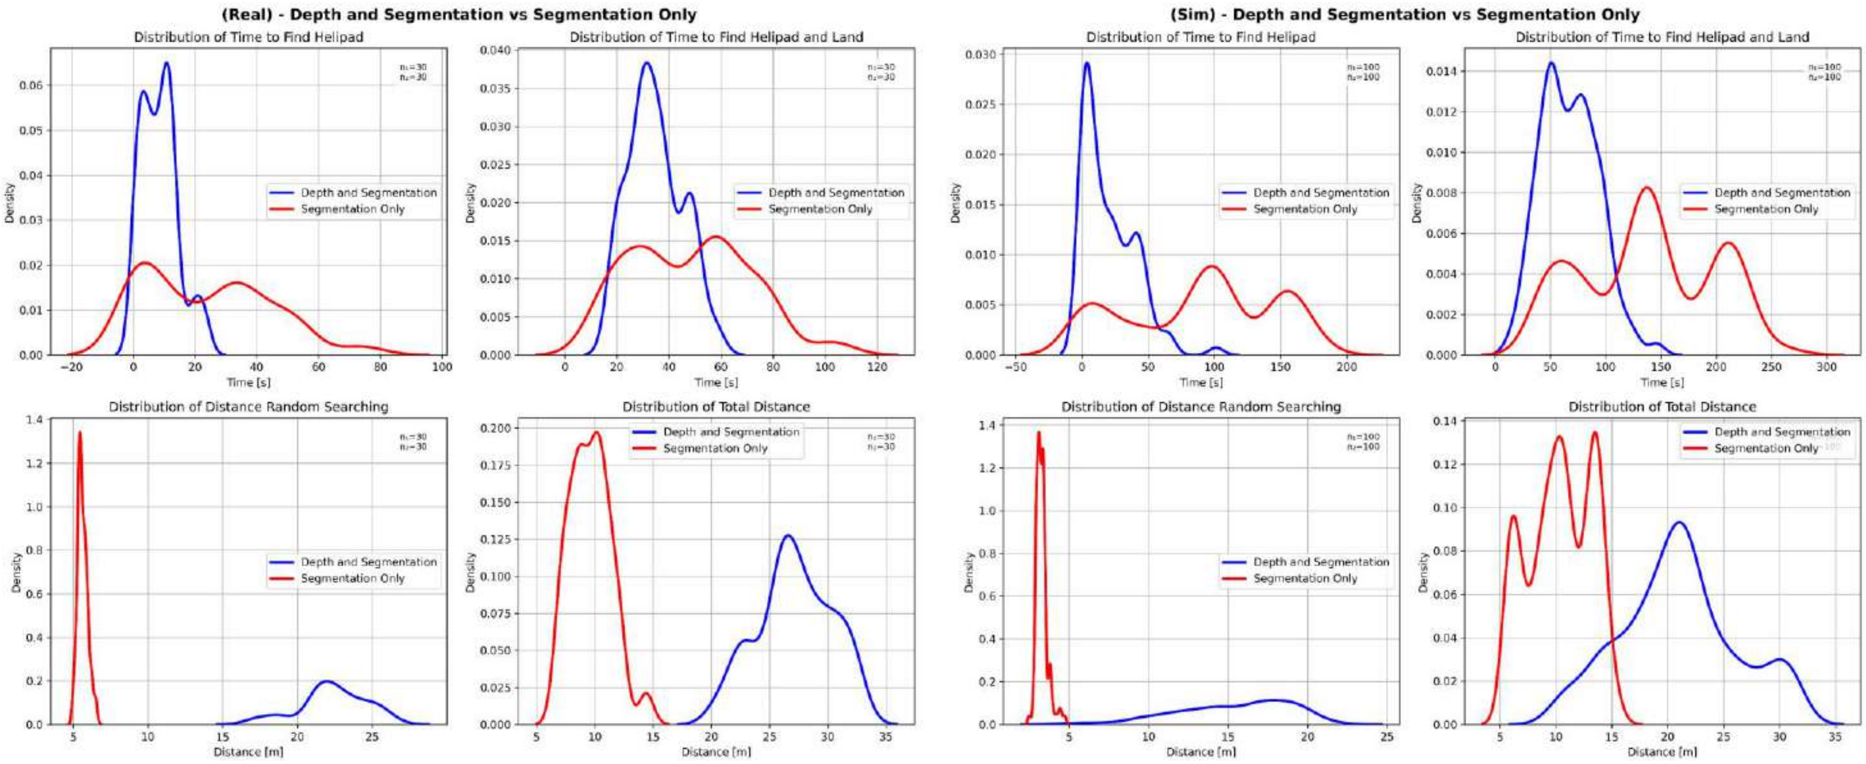

Distill the complex Depth and Segmentation Auto-pilot
into a single small neural net (1.3 M params U-Net)

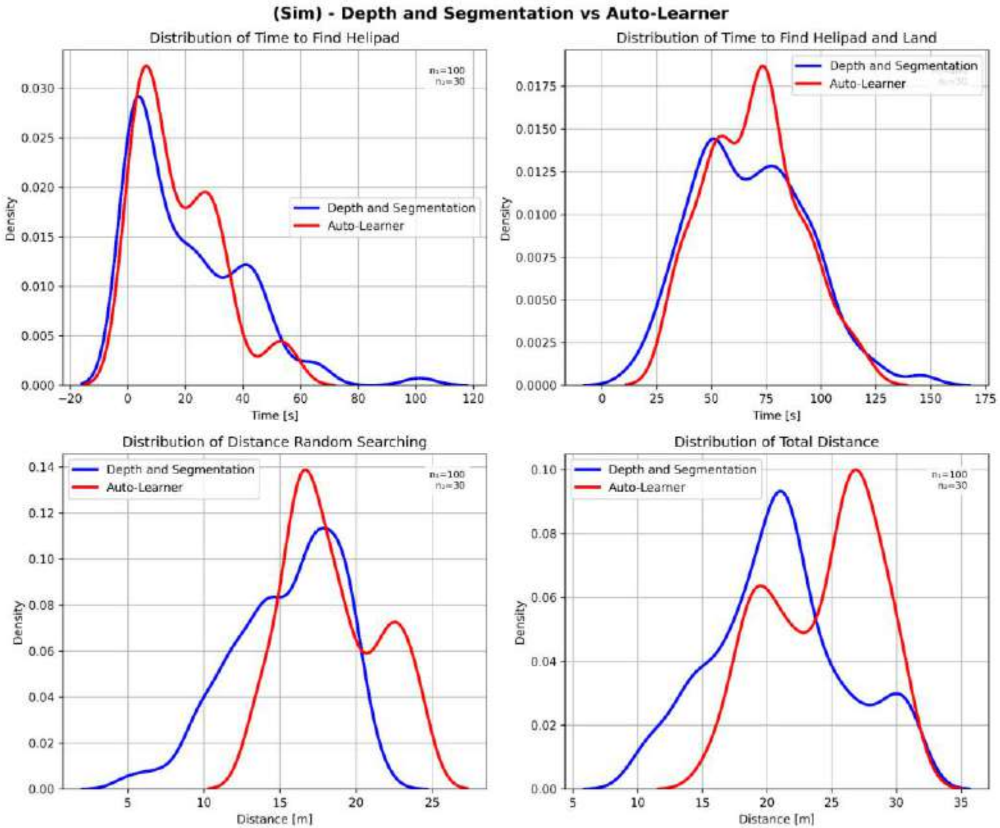

Final Small U-Net Architecture

Layer (type (var_name))	Input Shape	Output Shape	Param #	Trainable
DroneMovementNet (DroneMovementNet)	[1, 9, 256, 256]	[1, 3]	--	True
├─ConvolutionBlock (conv1)	[1, 9, 256, 256]	[1, 16, 256, 256]	--	True
│ └─Sequential (double_conv)	[1, 9, 256, 256]	[1, 16, 256, 256]	--	True
│ │ └─Conv2d (0)	[1, 9, 256, 256]	[1, 16, 256, 256]	1,312	True
│ │ └─BatchNorm2d (1)	[1, 16, 256, 256]	[1, 16, 256, 256]	32	True
│ │ └─ReLU (2)	[1, 16, 256, 256]	[1, 16, 256, 256]	--	--
│ │ └─Conv2d (3)	[1, 16, 256, 256]	[1, 16, 256, 256]	2,320	True
│ │ └─BatchNorm2d (4)	[1, 16, 256, 256]	[1, 16, 256, 256]	32	True
│ │ └─ReLU (5)	[1, 16, 256, 256]	[1, 16, 256, 256]	--	--
├─MaxPool2d (pool1)	[1, 16, 256, 256]	[1, 16, 128, 128]	--	--
├─ConvolutionBlock (conv2)	[1, 16, 128, 128]	[1, 32, 128, 128]	--	True
│ └─Sequential (double_conv)	[1, 16, 128, 128]	[1, 32, 128, 128]	--	True
│ │ └─Conv2d (0)	[1, 16, 128, 128]	[1, 32, 128, 128]	4,640	True
│ │ └─BatchNorm2d (1)	[1, 32, 128, 128]	[1, 32, 128, 128]	64	True
│ │ └─ReLU (2)	[1, 32, 128, 128]	[1, 32, 128, 128]	--	--
│ │ └─Conv2d (3)	[1, 32, 128, 128]	[1, 32, 128, 128]	9,248	True
│ │ └─BatchNorm2d (4)	[1, 32, 128, 128]	[1, 32, 128, 128]	64	True
│ │ └─ReLU (5)	[1, 32, 128, 128]	[1, 32, 128, 128]	--	--
├─MaxPool2d (pool2)	[1, 32, 128, 128]	[1, 32, 64, 64]	--	--
├─ConvolutionBlock (conv3)	[1, 32, 64, 64]	[1, 64, 64, 64]	--	True
│ └─Sequential (double_conv)	[1, 32, 64, 64]	[1, 64, 64, 64]	--	True
│ │ └─Conv2d (0)	[1, 32, 64, 64]	[1, 64, 64, 64]	18,496	True
│ │ └─BatchNorm2d (1)	[1, 64, 64, 64]	[1, 64, 64, 64]	128	True
│ │ └─ReLU (2)	[1, 64, 64, 64]	[1, 64, 64, 64]	--	--
│ │ └─Conv2d (3)	[1, 64, 64, 64]	[1, 64, 64, 64]	36,928	True
│ │ └─BatchNorm2d (4)	[1, 64, 64, 64]	[1, 64, 64, 64]	128	True
│ │ └─ReLU (5)	[1, 64, 64, 64]	[1, 64, 64, 64]	--	--
├─MaxPool2d (pool3)	[1, 64, 64, 64]	[1, 64, 32, 32]	--	--
├─ConvolutionBlock (conv4)	[1, 64, 32, 32]	[1, 128, 32, 32]	--	True
│ └─Sequential (double_conv)	[1, 64, 32, 32]	[1, 128, 32, 32]	--	True
│ │ └─Conv2d (0)	[1, 64, 32, 32]	[1, 128, 32, 32]	73,856	True
│ │ └─BatchNorm2d (1)	[1, 128, 32, 32]	[1, 128, 32, 32]	256	True
│ │ └─ReLU (2)	[1, 128, 32, 32]	[1, 128, 32, 32]	--	--
│ │ └─Conv2d (3)	[1, 128, 32, 32]	[1, 128, 32, 32]	147,584	True
│ │ └─BatchNorm2d (4)	[1, 128, 32, 32]	[1, 128, 32, 32]	256	True
│ │ └─ReLU (5)	[1, 128, 32, 32]	[1, 128, 32, 32]	--	--
├─MaxPool2d (pool4)	[1, 128, 32, 32]	[1, 128, 16, 16]	--	--
├─ConvolutionBlock (conv5)	[1, 128, 16, 16]	[1, 256, 16, 16]	--	True
│ └─Sequential (double_conv)	[1, 128, 16, 16]	[1, 256, 16, 16]	--	True
│ │ └─Conv2d (0)	[1, 128, 16, 16]	[1, 256, 16, 16]	295,168	True
│ │ └─BatchNorm2d (1)	[1, 256, 16, 16]	[1, 256, 16, 16]	512	True
│ │ └─ReLU (2)	[1, 256, 16, 16]	[1, 256, 16, 16]	--	--
│ │ └─Conv2d (3)	[1, 256, 16, 16]	[1, 256, 16, 16]	590,080	True
│ │ └─BatchNorm2d (4)	[1, 256, 16, 16]	[1, 256, 16, 16]	512	True
│ │ └─ReLU (5)	[1, 256, 16, 16]	[1, 256, 16, 16]	--	--
├─AdaptiveAvgPool2d (global_average_pooling)	[1, 256, 16, 16]	[1, 256, 1, 1]	--	--
├─Linear (fc1)	[1, 256]	[1, 512]	131,584	True
├─ReLU (relu)	[1, 512]	[1, 512]	--	--
├─Dropout (dropout)	[1, 512]	[1, 512]	--	--
├─Linear (fc2)	[1, 512]	[1, 3]	1,539	True

Total params: 1,314,739
Trainable params: 1,314,739
Non-trainable params: 0
Total mult-adds (G): 1.15

Dancing with the Drones - Towards Human-AI Art and Peformance

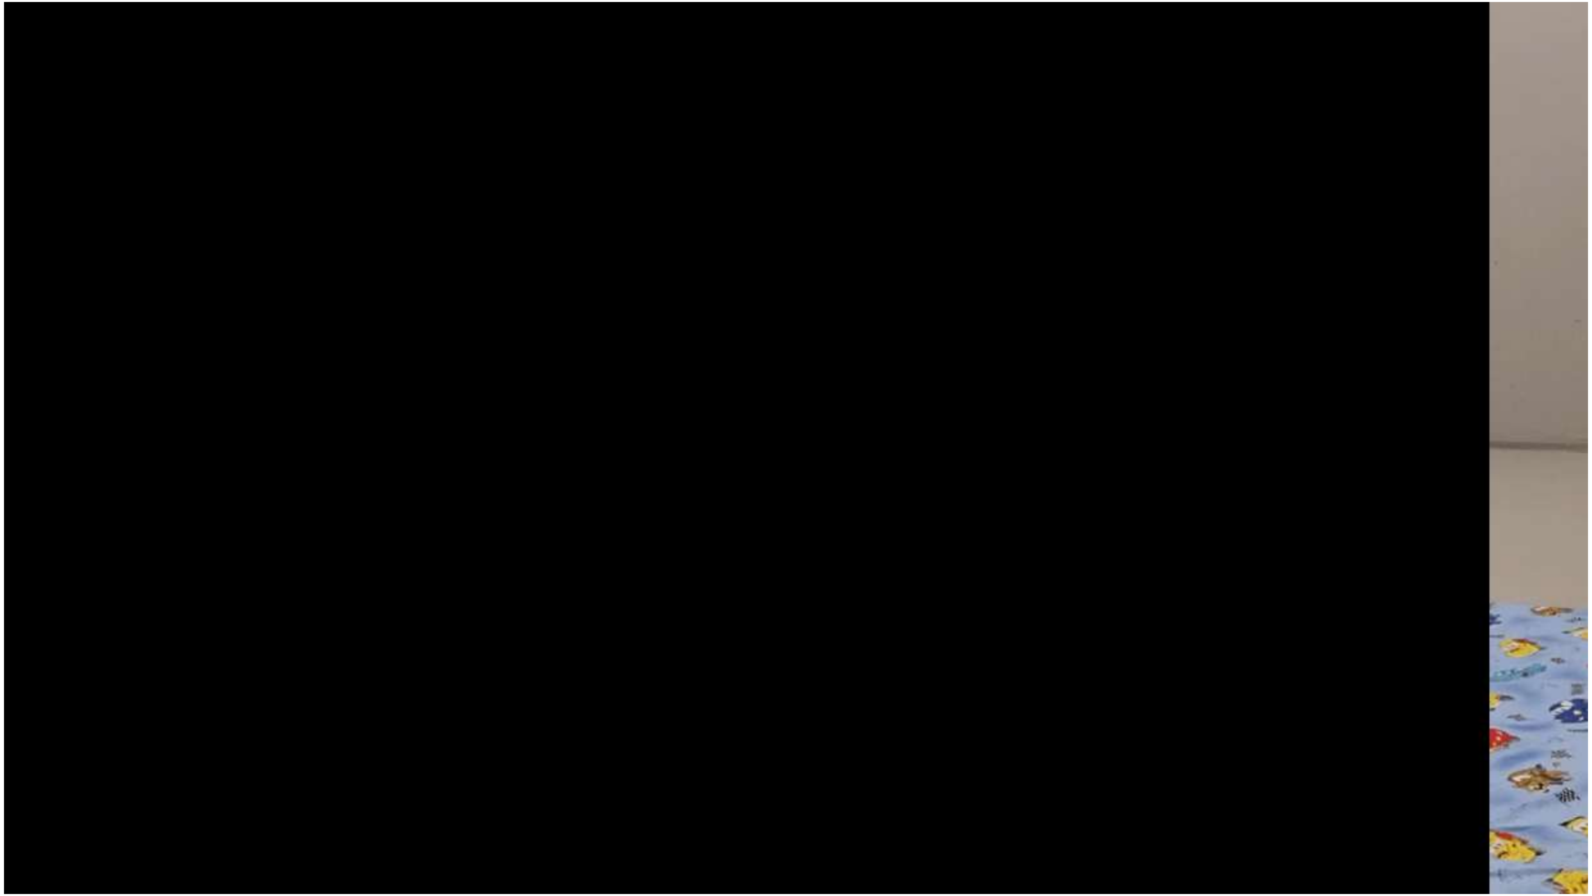

Music: “Emergence” (Album: “Supersonic”) , composed, performed and co-produced by Marius Leordeanu

Fly to the Sky – A Cosmic Attraction

Music: “Cosmic Attraction”, composed, performed and co-produced by Marius Leordeanu