Robust and Hybrid Machine Learning

ILLS-DATAIA 2023

THOME Nicolas – Prof. at Sorbonne University
ISIR Lab, MLIA TEAM

Research activities

Research topics: machine learning (ML), deep learning

Application domains

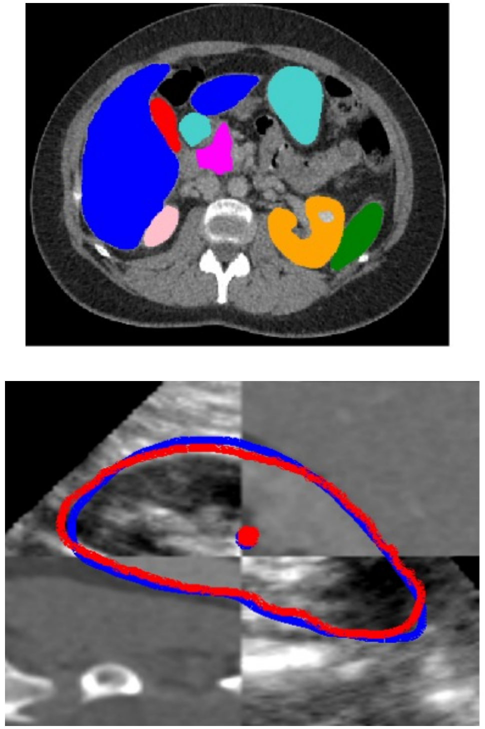

Healthcare

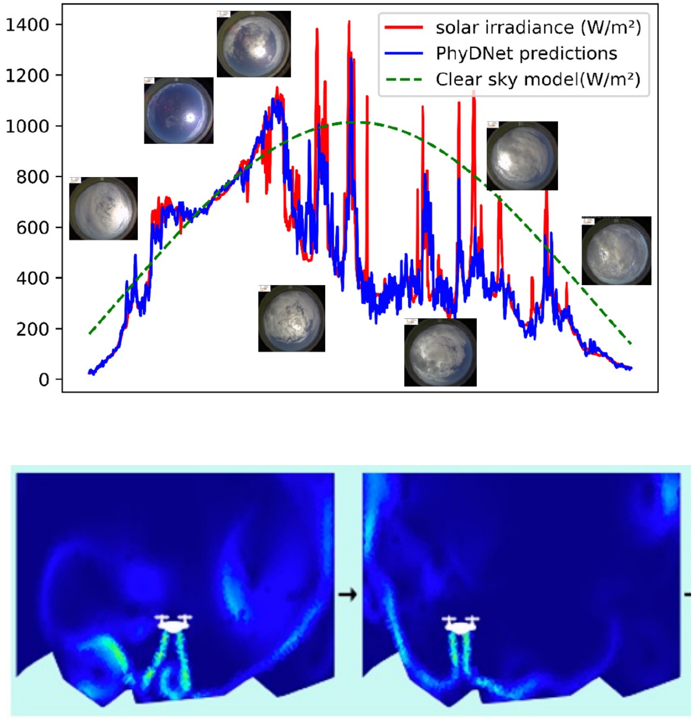

Physics

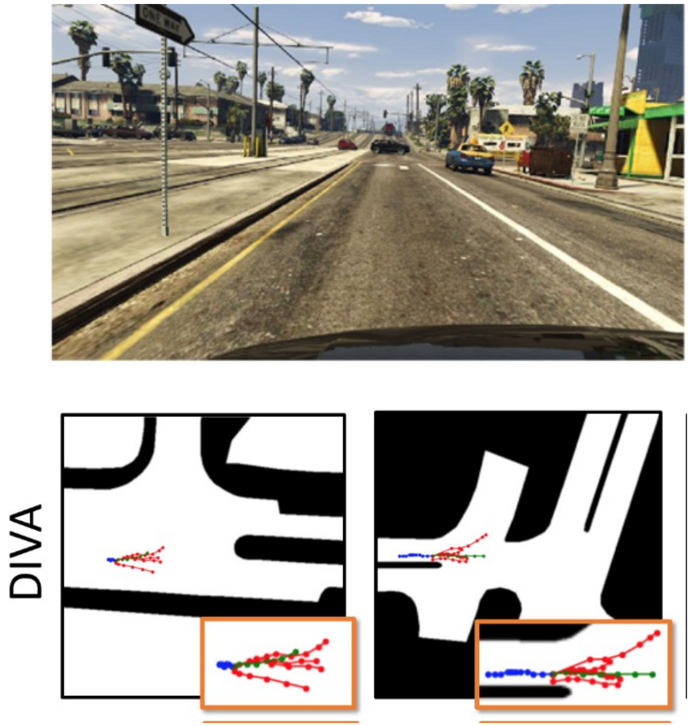

autonomous vehicles

Topics: machine learning (ML), deep learning

- **Learning formulation:** semi-supervised, weakly supervised learning
- **Theoretical M**L: robustness, optimization
- **Including various forms of knowledge in ML**

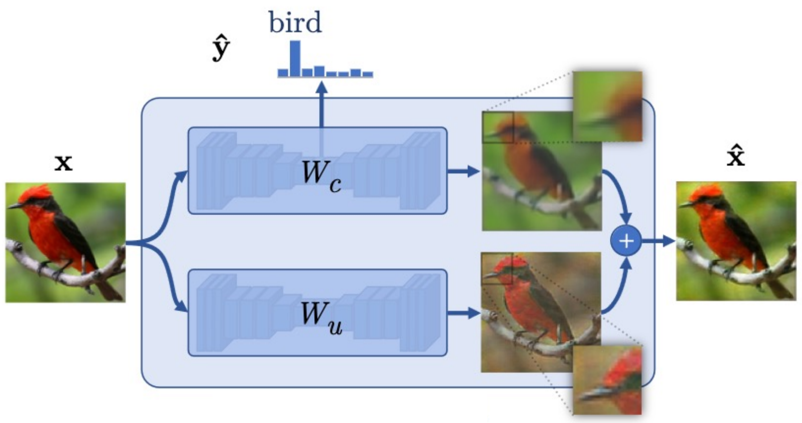

Optimization of non-decomposable losses, *e.g.* rank losses (AP)

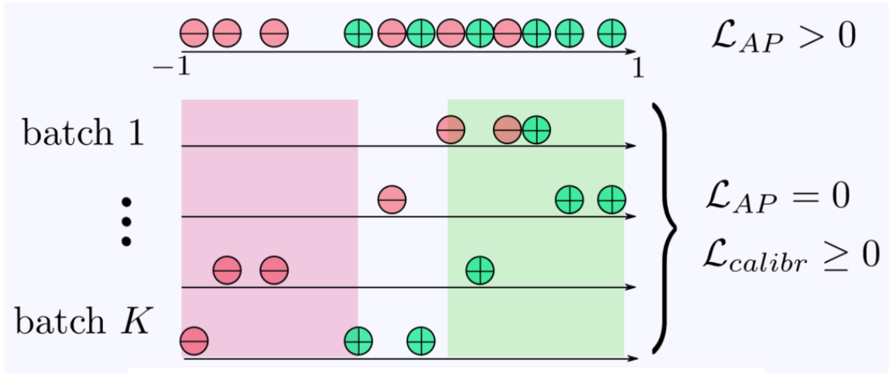

$$DG_{\mathrm{AP}}(\boldsymbol{\theta}) = \frac{1}{K}\sum_{b=1}^{K} \mathrm{AP}_i^b(\boldsymbol{\theta}) - \mathrm{AP}_i(\boldsymbol{\theta})$$

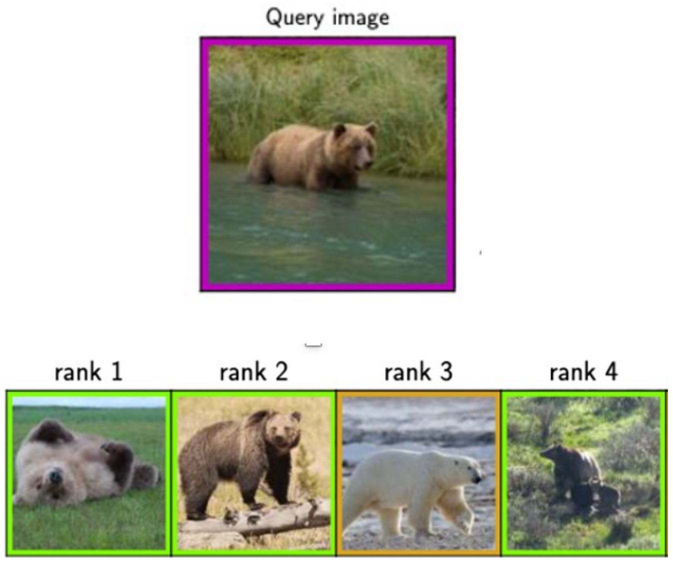

[RTC18] T. Robert, N. Thome, M. Cord. Classification and reconstruction cooperation for semi-supervised learning. ECCV 2018.
[RTR+21] E. Ramzi, N. Thome, C. Rambour, N. Audebert, X. Bitot. Robust and Decomposable Average Precision for Image Retrieval. NeurIPS 2021.
[RAT+22] E. Ramzi, N. Audebert, N. Thome, C. Rambour, X. Bitot. Hierarchical Average Precision Training for Pertinent Image Retrieval. ECCV 2022.

Outline

1. **Recent contributions**

I) Robustness in deep learning

II) Hybrid physics-informed ML

2. Open issues & perspectives

I) Robustness in deep learning

- **Uncertainty quantification: crucial in critical systems**

"Know when you do not know"

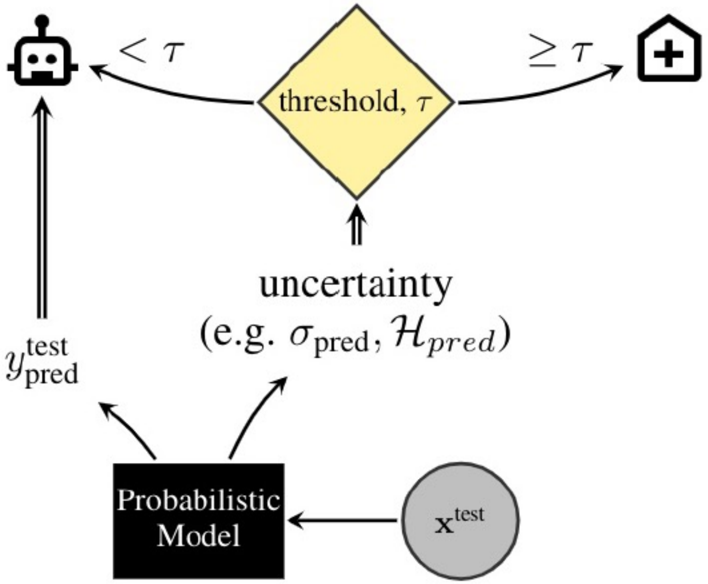

Abstain to make a prediction

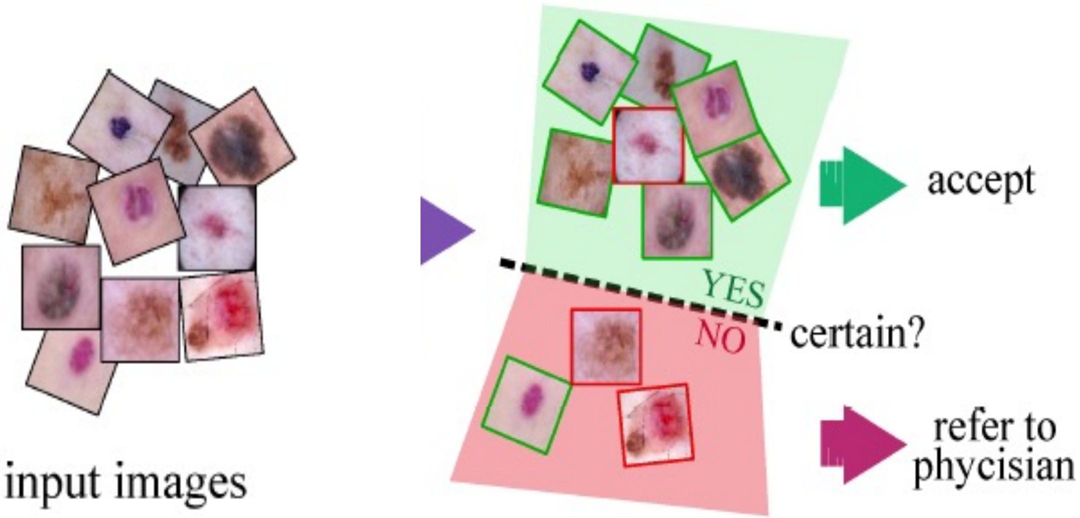

Uncertainty quantification in deep learning

- **Uncertainty for failure prediction [CBT+19]**: correct vs incorrect predictions

- **<u>Our proposal:</u> True Class probability (TCP)** vs Maximum Class Probability (MCP)

- TCP better than MCP for failure prediction

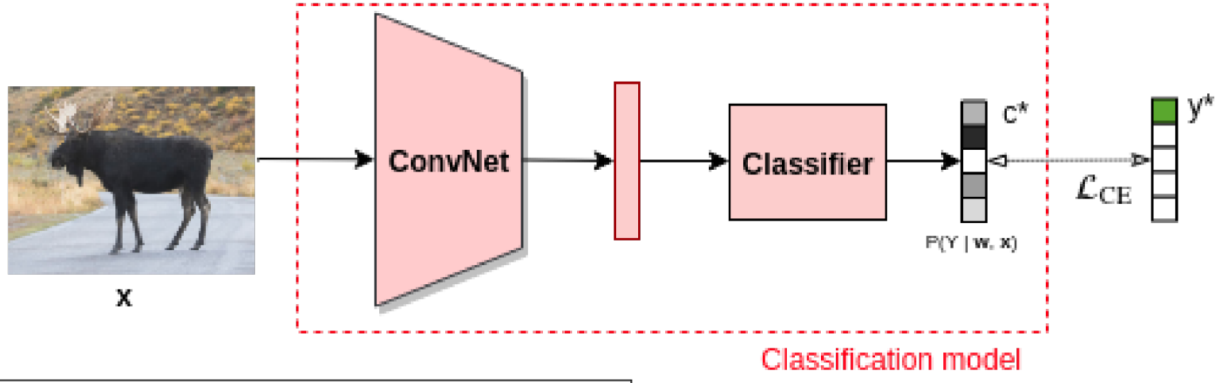

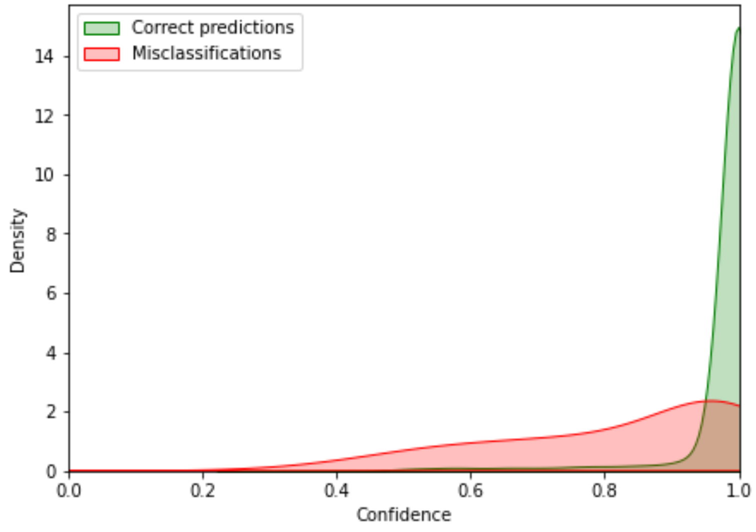

MCP

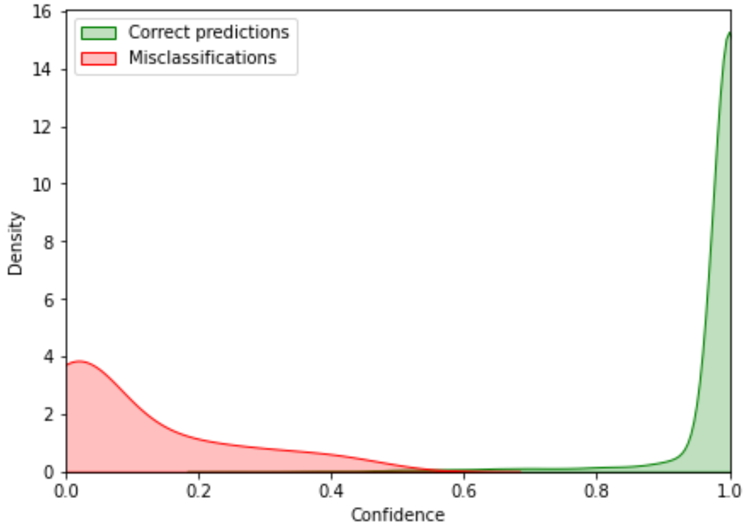

TCP

[CTB+19] C. Corbière, N. Thome, A. Bar-Hen, M. Cord, P. Pérez. Addressing Failure Detection by Learning Model Confidence. NeurIPS 2019.

Uncertainty quantification in deep learning

TCP unknown at test time: learning it!

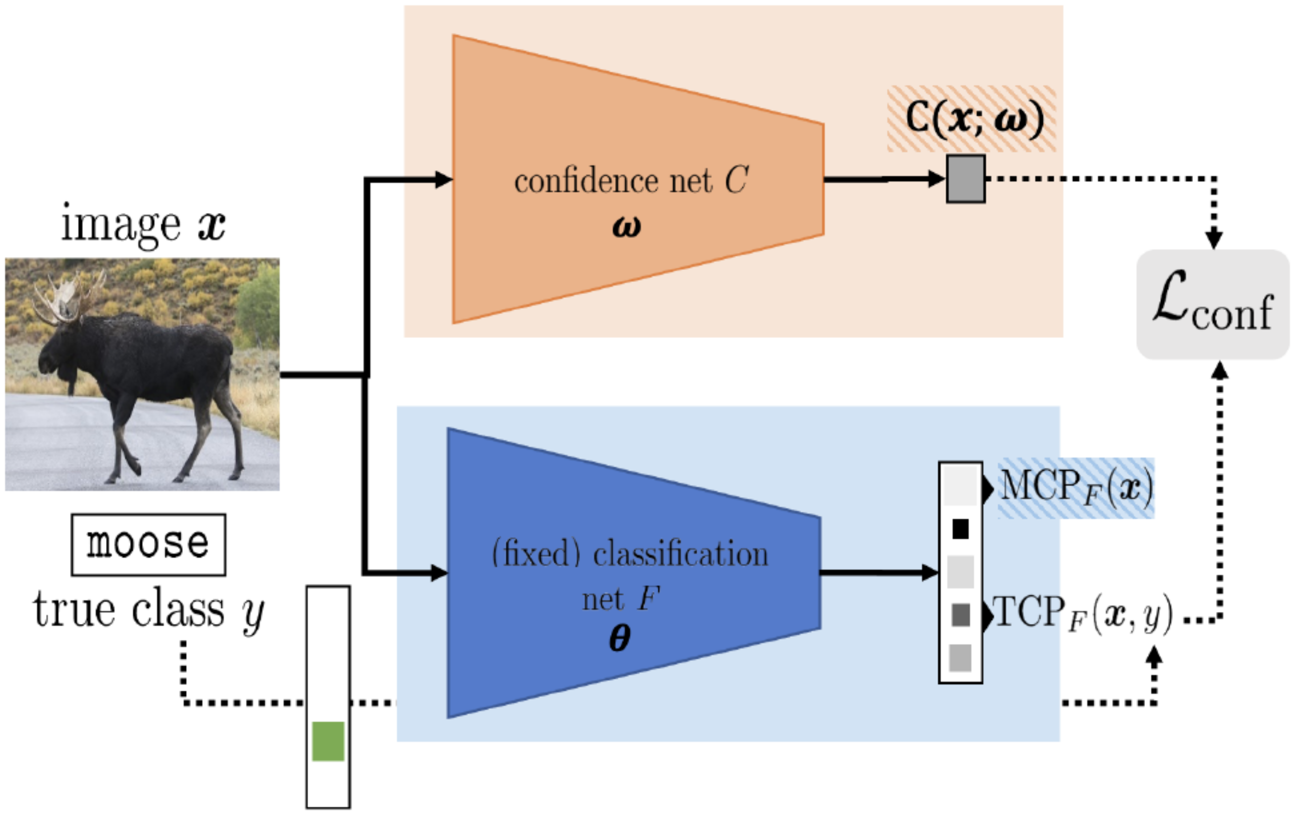

- Pre-trained prediction model (blue)
- Learning to regress TCP with an auxiliary model (orange)

$$\mathcal{L}_{conf}(\theta; \mathcal{D}) = \frac{1}{N}\sum_{i=1}^{N}\left(\hat{c}(\mathbf{x}_i, \theta) - c^*(\mathbf{x}_i, y_i^*)\right)^2$$

[CTB+19] C. Corbière, N. Thome, A. Bar-Hen, M. Cord, P. Pérez. Addressing Failure Detection by Learning Model Confidence. NeurIPS 2019.

Learning confidence for self-labelling

- Extension for domain adaptation [CTS+21]

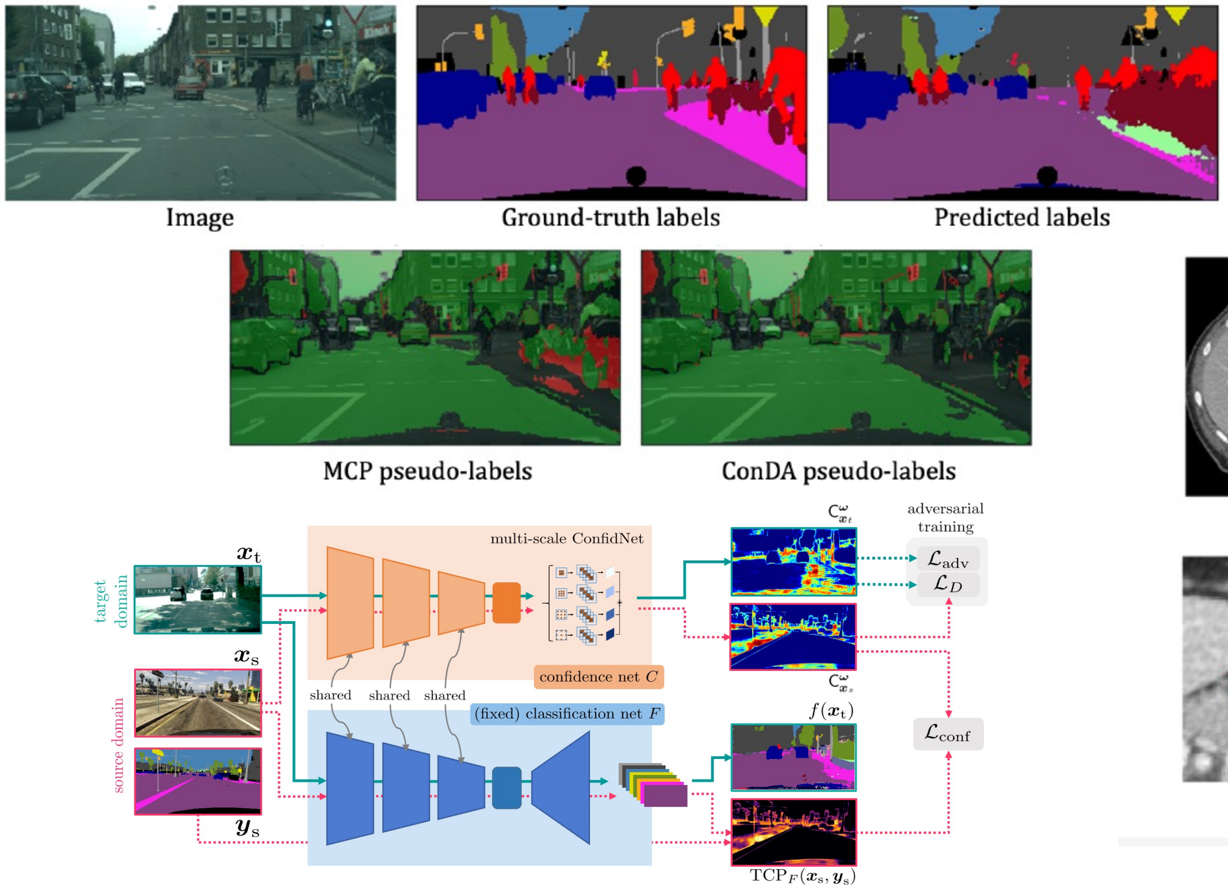

Medical image segmentation [PTS21]

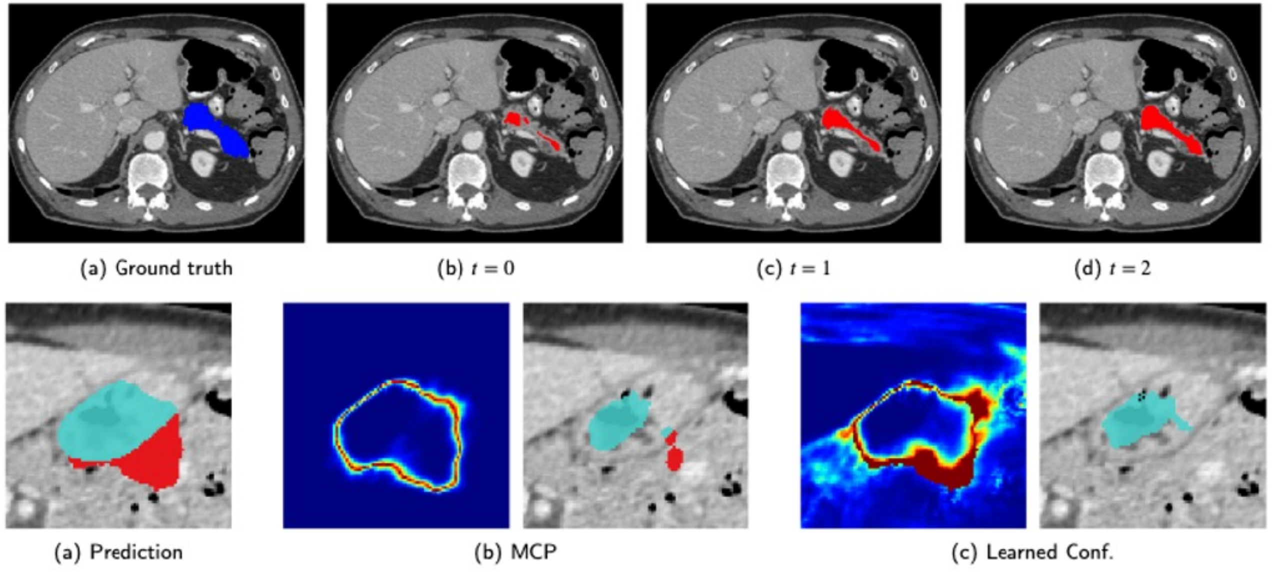

[CTS+21] C. Corbière, N. Thome, A. Saporta, T-H. Vu, M. Cord, P. Pérez. Confidence Estimation via Auxiliary Models. IEEE Transactions on Pattern Analysis and Machine Intelligence (T-PAMI), vol. 44, no. 10, pp. 6043-6055, June 2021.
[PTS21] O. Petit, N. Thome, L. Soler. 3D Spatial Priors for Semi-Supervised Organ Segmentation with Deep Convolutional Neural Networks. International Journal of Computer Assisted Radiology and Surgery, Springer Verlag, In press, 2021.

Uncertainty: Out-Of-Distribution (OOD) detection

- Accurate OOD detection ⇔ accurate density estimation
- **HEAT** [LRR+23]: **Hybrid Energy Based Model (EBM)**

- **Energy-based correction** of prior energy terms, *e.g.* Gaussians

- **Energy composition** of several terms (Gaussian, Energy Logits, std for style)

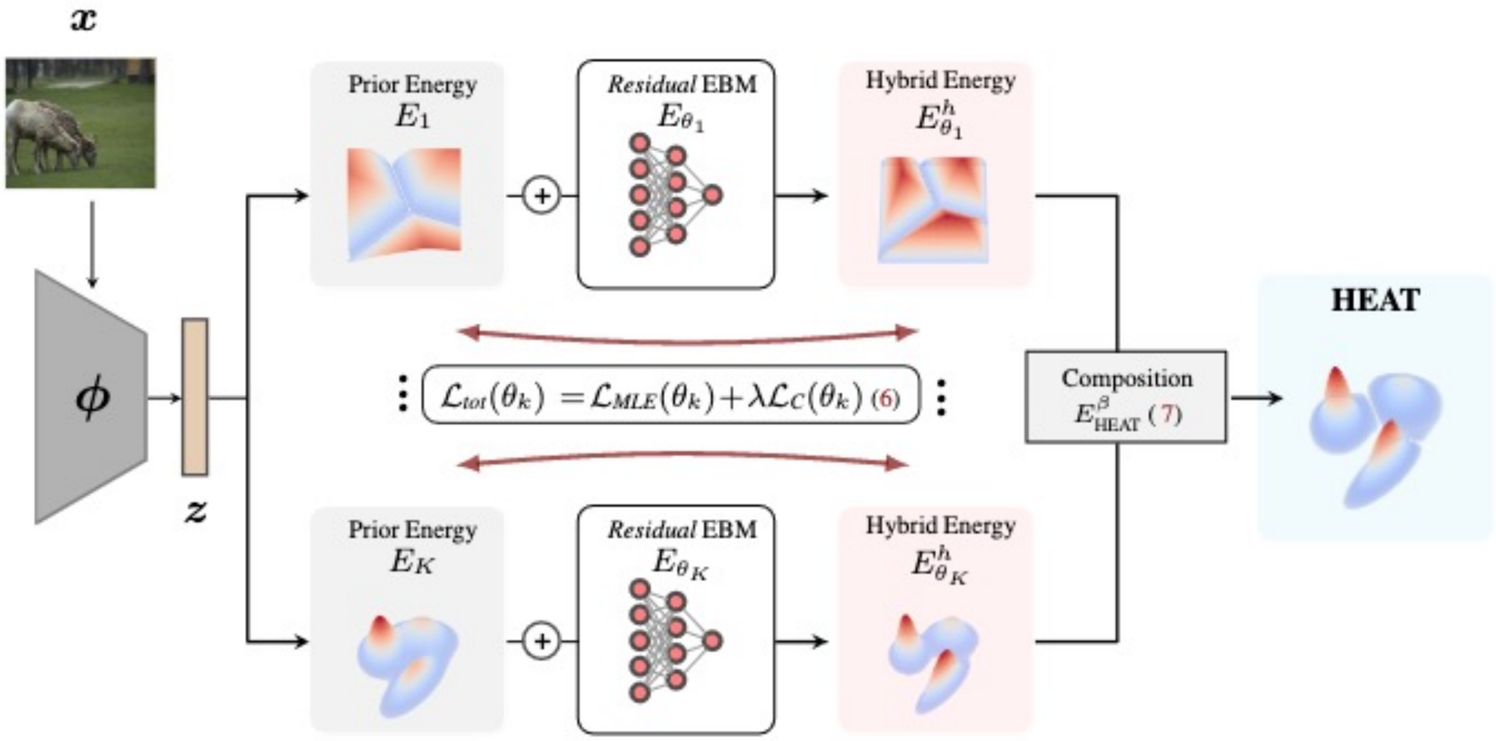

[LRR+23] M. Lafon, E. Ramzi, C. Rambour, N. Thome. Addressing Failure Detection by Learning Model Confidence. ICLR 2023.

II) Prediction for physical & dynamical systems

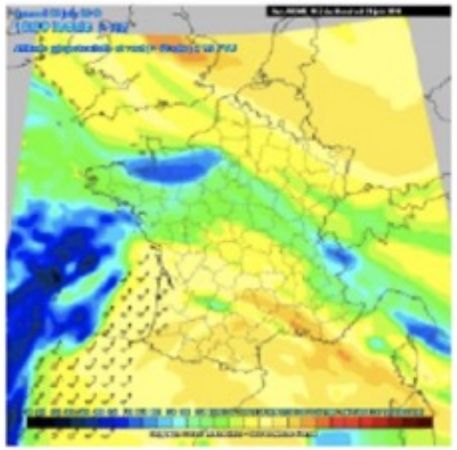

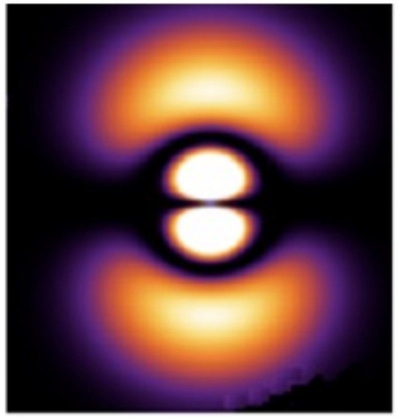

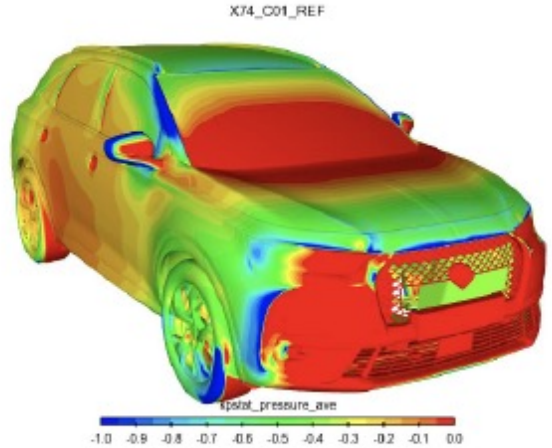

- **Model-based (MB)** approaches, *e.g.* based on ODE/PDE
  - Physical models: approximation of real world-dynamics
- **Machine Learning (ML)**: less biased BUT generalization issues

  - <u>**Contributions:**</u> **hybrid physics-informed machine learning**
    - **learning residual of approximate physical models**

Motivation: data-driven *vs* simplified models

Damped pendulum: $\frac{d^2\theta}{dt^2} + \omega_0^2 \sin\theta + \lambda \frac{d\theta}{dt} = 0$

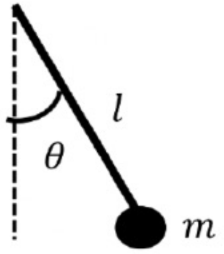

- **Data-driven models** struggle to extrapolate complex dynamics, in particular in data-scarce contexts
- **Physical models** fail to extrapolate when they are misspecified: forecasting & parameter identification failure

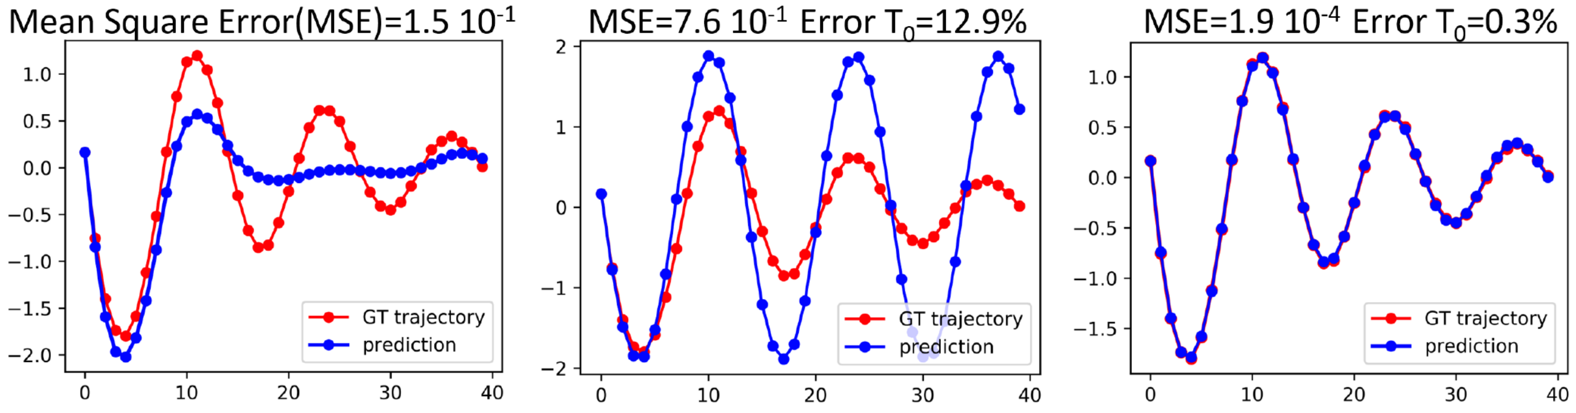

(a) Data-driven Neural ODE (b) Simple physical model (c) Our APHYNITY framework

⇒ Augmenting PHYsical models for ideNtIfying and forecasTing complex dYnamic (APHYNITY)

Augmenting physical models: APHYNITY [YLD+21]

- Representing state's derivative as $\boxed{\frac{dX_t}{dt} = F(X_t) = F_p + F_a}$
  - F_p approximate ODE/PDE, $\mathbf{F_a}$ **learned residual**
- APHYNITY objective : $\min_{F_p \in \mathcal{F}_p, F_a \in \mathcal{F}} \quad \|F_a\| \quad \text{subject to} \quad \forall X \in \mathcal{D}, \forall t, \frac{dX_t}{dt} = (F_p + F_a)(X_t)$
  - Decomposition: exists and is unique (under mild conditions)

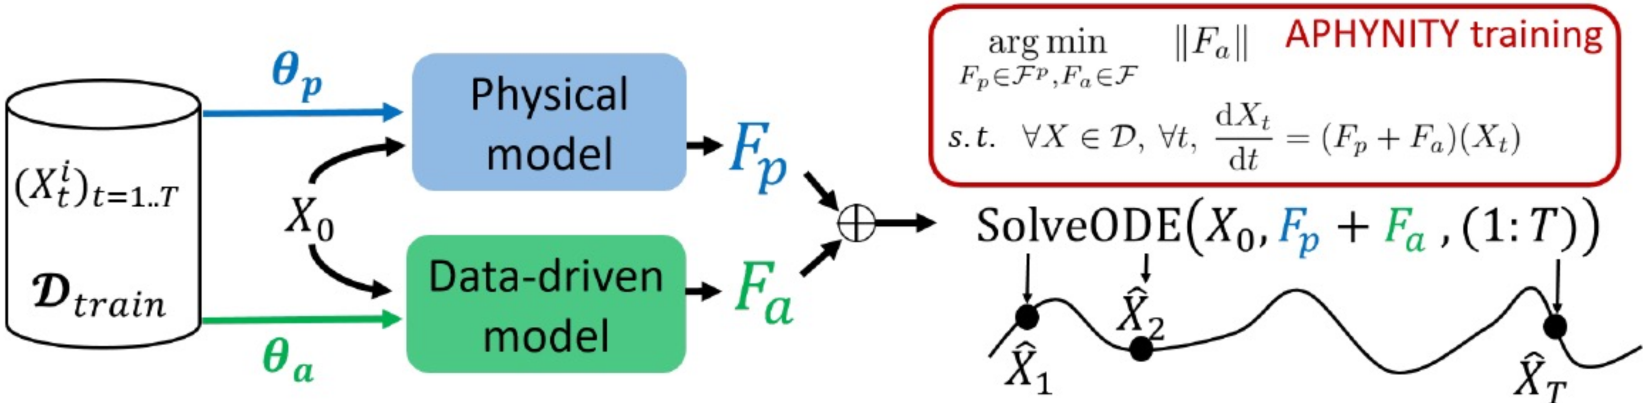

- Application to simulated physical phenomenon

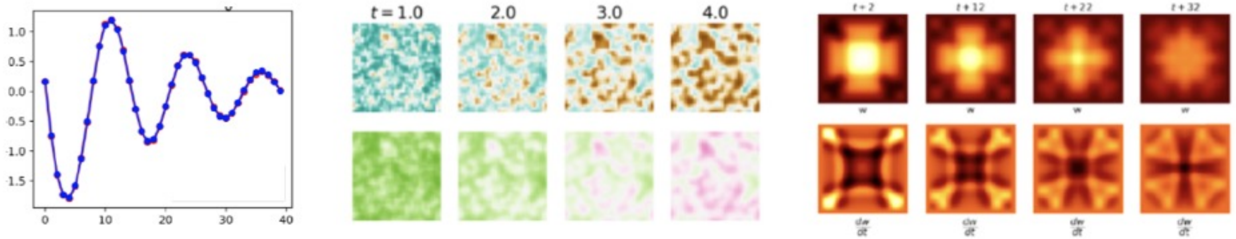

[YLD+21] Y. Yin, V. Le Guen, J. Dona, I. Ayed, E. de Bézenac, N. Thome, P. Gallinari. Augmenting Physical Models with Deep Networks for Complex Dynamics Forecasting. ICLR 2021.

Learning Residual dynamics: video prediction [LT20] and optical flow estimation [LRT22]

- Deep learning models: trained with complex curriculum, i.e. synthetic data (Chairs, Things, Sintel), real data (HD1K, Kitti)

- Traditional methods: based on brightness consistency (BC) assumption:

$$\frac{\partial I}{\partial t}(t,\mathbf{x}) + \mathbf{w}(t,\mathbf{x}) \cdot \nabla I(t,\mathbf{x}) = 0$$

  - BUT: BC violated in several usual conditions

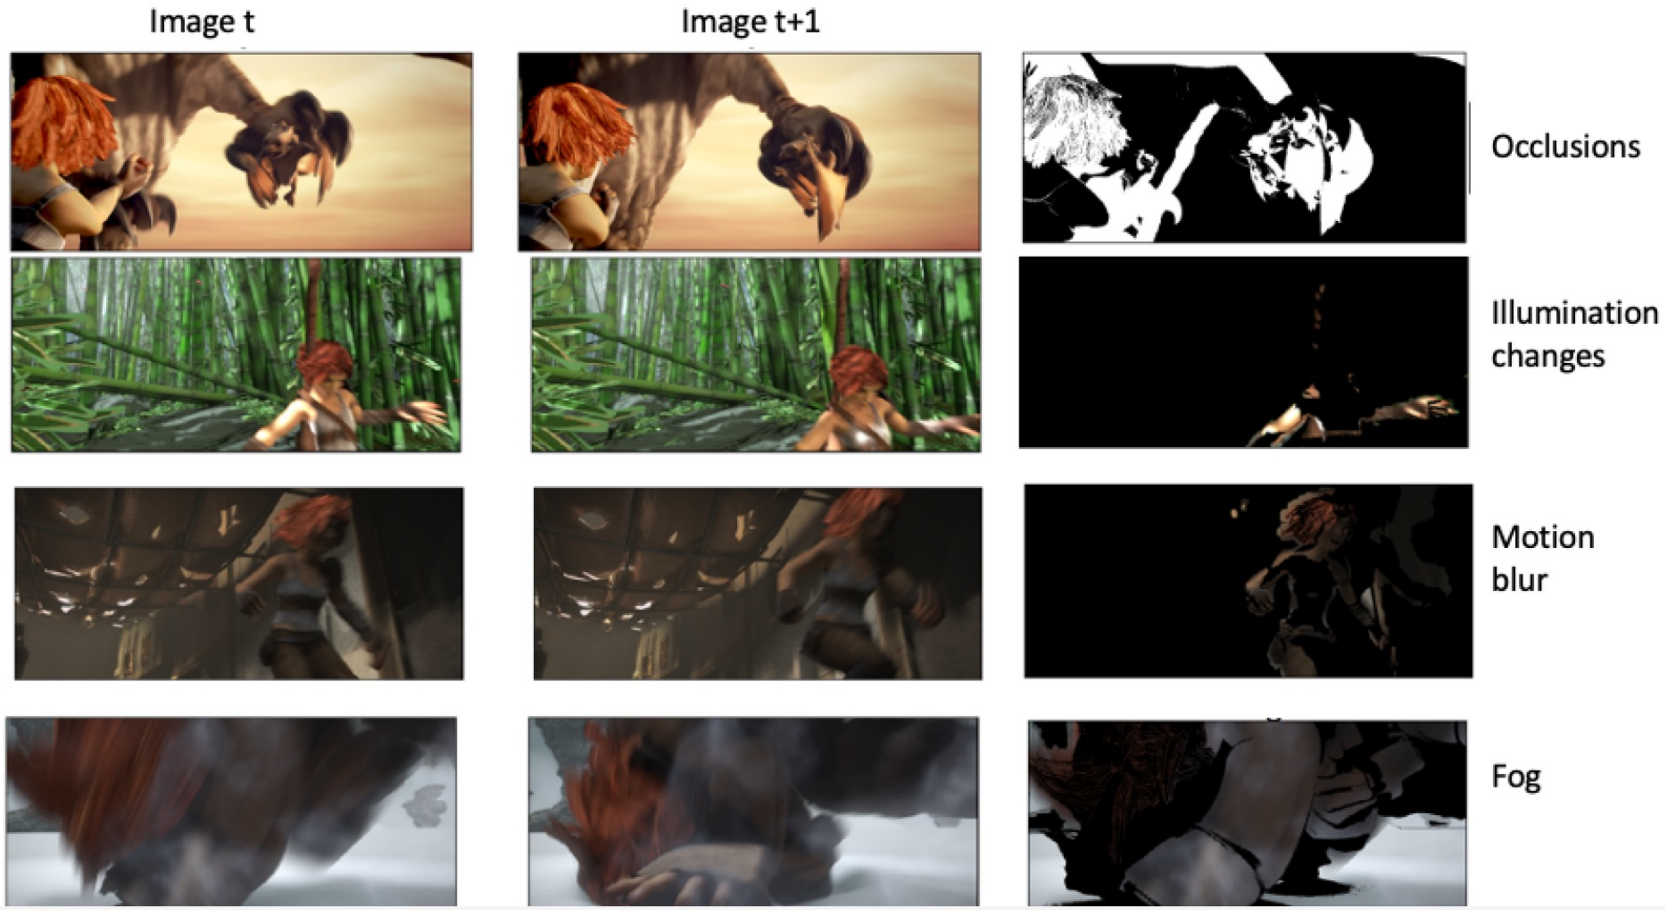

[LT20] V. Le Guen, N. Thome. Disentangling Physical Dynamics from Unknown Factors for Unsupervised Video Prediction. CVPR 2020.
[LRT22] V. Le Guen, C. Rambour N. Thome. Complementing Brightness Constancy with Deep Netwoks for Optical Flow Prediction. ECCV 2022.

COMBO model for optical flow estimation [LRT22]

- Complementing BC with deep NNs for accurate flow prediction
- **Flow decomposition:**

$$w(x) = \alpha(x)\,.\,w_p(\mathrm{x}) + (1 - \alpha(x))\,.\,w_\mathrm{a}(\mathrm{x})$$

- $\alpha(x)$ BC confidence
- $w_p(\mathrm{x})$ physical flow
- $w_\mathrm{a}(\mathrm{x})$ residual flow

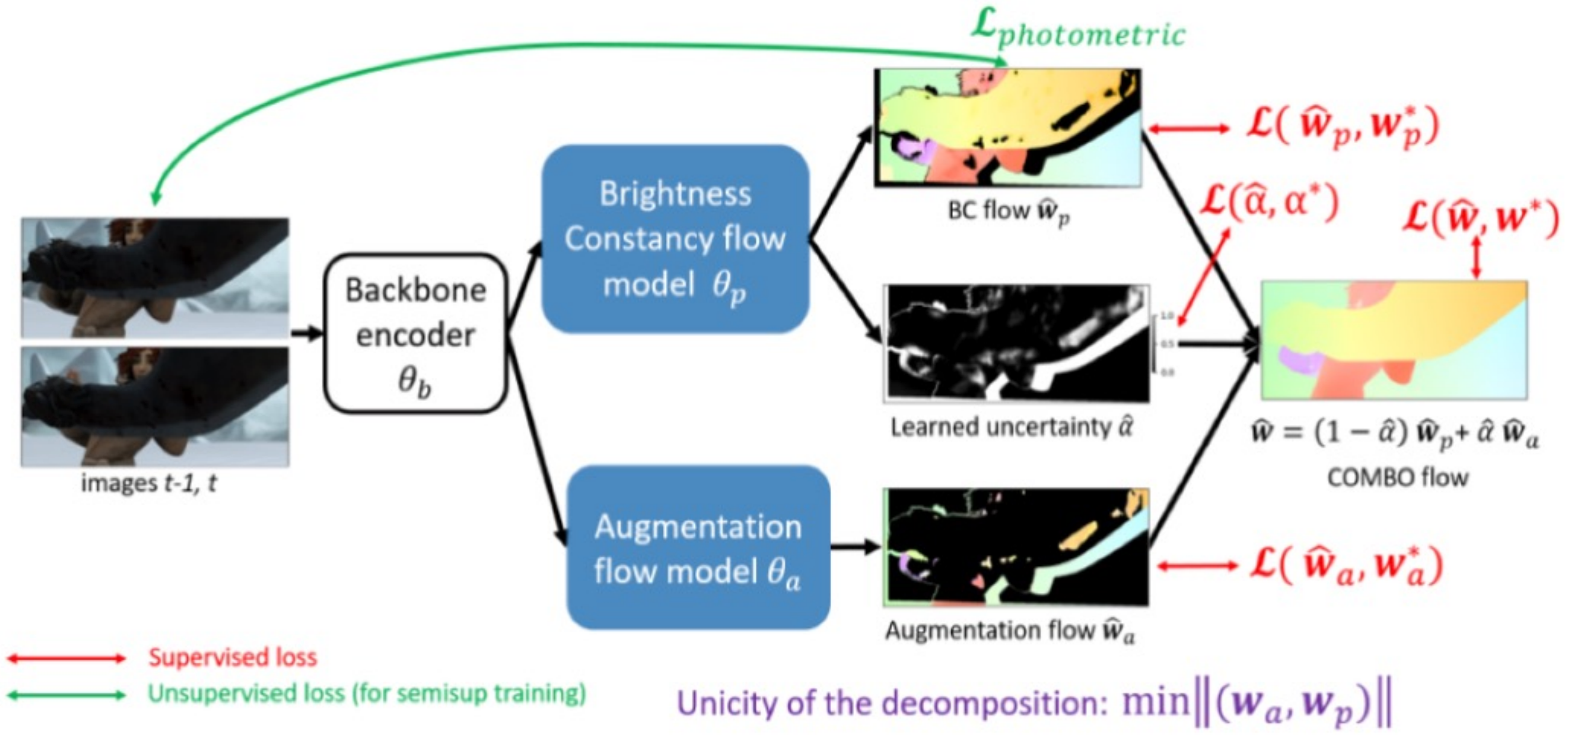

Semi-supervised: much simpler training curriculum

[LRT22] V. Le Guen, C. Rambour N. Thome. Complementing Brightness Constancy with Deep Netwoks for Optical Flow Prediction. ECCV 2022.

Outline

1. Recent contributions
2. Open issues & perspectives

Current context

- Large Language & Multi-Modal Models,
  - Huge success and buzz in the last year
  - X-modal foundation models, *e.g.* [1]
  - Flexibility of “In-Context Learning” (ICL) [2]
- MLIA Team in robotic lab (ISIR) since ’22
  - Collaborations on AI for robotics, medical

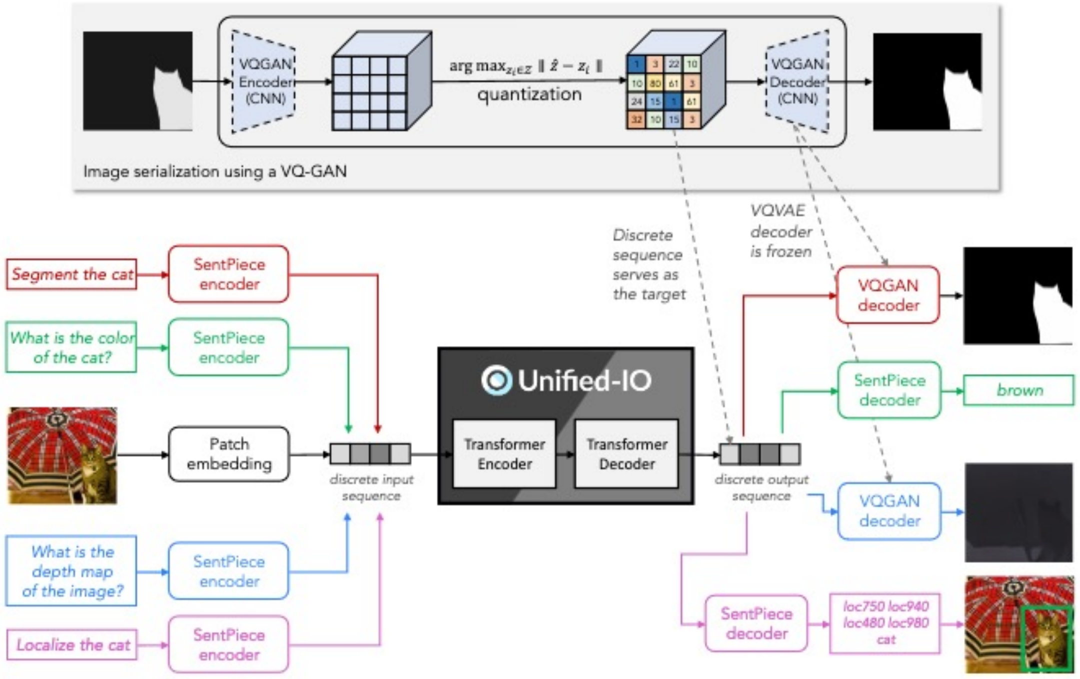

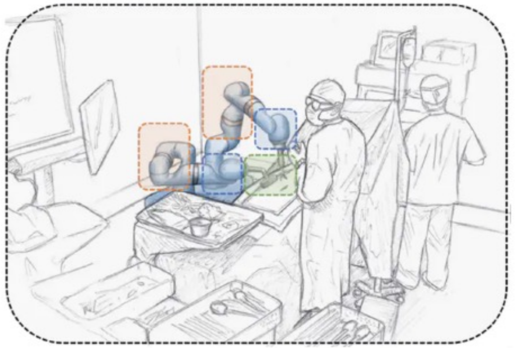

[1] Unified-IO: A Unified Model for Vision, Language, and Multi-Modal Tasks. J. Lu, C. Clark, R. Zellers, R. Mottaghi, A. Kembhavi. ICLR 2023
[2] Foundation models for generalist medical artificial intelligence. M. Moor, O. Banerjee, Z.S.H. Abad, H. M. Krumholz, J. Leskovec, E.J. Topol, P. Rajpurkar. Nature volume 616, pages 259–265, 2023.

Perspectives: learning formulation and architectures

Open questions:

- Zero/few-shot learning, pure prompt vs adapters [3]
- Instruction tuning [4]
- Multi-modal vs mono-modal pre-training
- Model compression

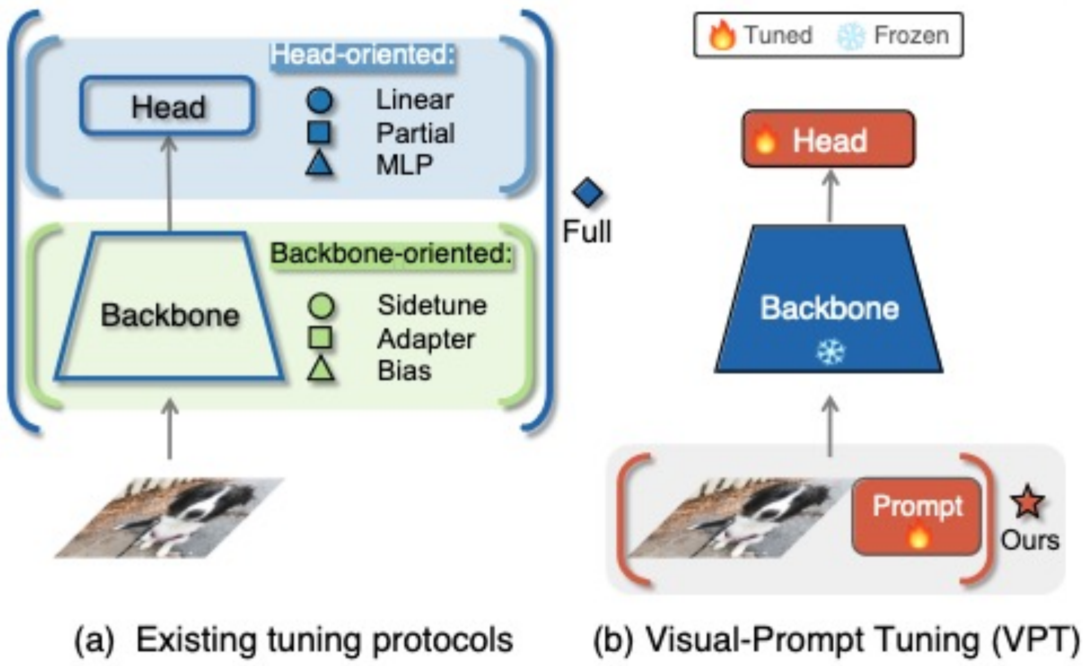

Fig. 1. Visual-Prompt Tuning (VPT) *vs*. other transfer learning methods. (a) Current transfer learning protocols are grouped based on the tuning scope: Full fine-tuning, Head-oriented, and Backbone-oriented approaches.

[3] Visual Prompt Tuning. M. Jia, L. Tang, B.C. Chen, C. Cardie, S. Belongie, B. Hariharan, S.N. Lim. ECCV 2022.
[4] MultiInstruct: Improving Multi-Modal Zero-Shot Learning via Instruction Tuning. Zhiyang Xu, Ying Shen, Lifu Huang, ArXiv, 2023.

Perspectives: Robustness

Explainability & reasoning with LLM

- High-level x-AI [5] (≠ saliency)
- Grounding explanation in images
- **Main challenge:** accurate alignment between text/image

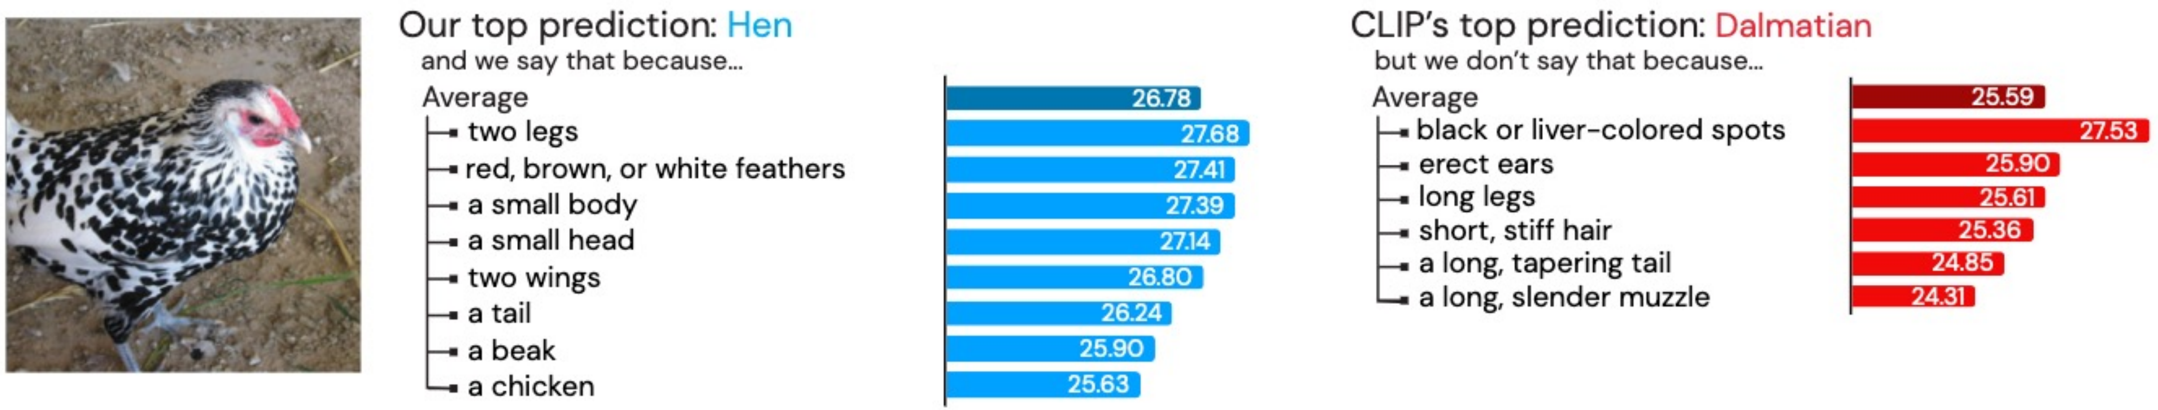

[5] Visual Classification via Description from Large Language Models. Sachit Menon, Carl Vondric, ICLR 2023

Perspectives: Hybrid prediction & control

Hybrid physical models

- Physical prior in model-based RL [6]

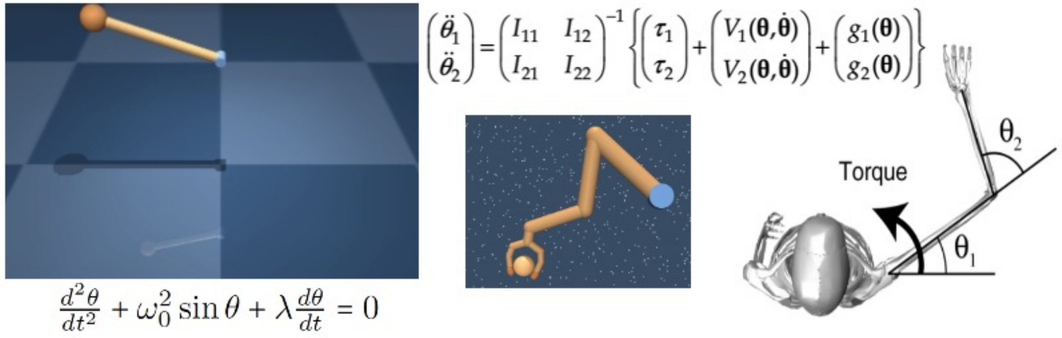

Language and control

- LLM as controllers [7]
- Hybrid methods: language, control, knowledge bases, *etc*

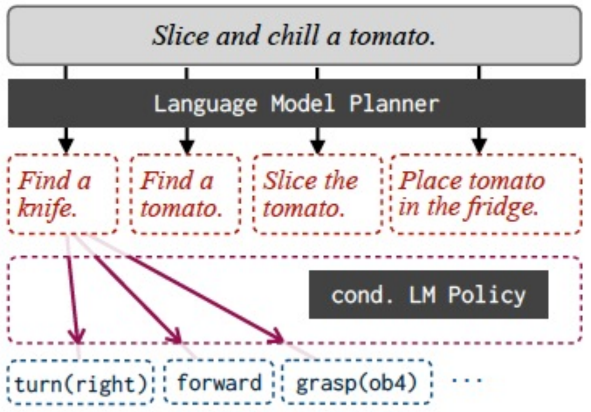

[6] Physics-Informed Model-Based Reinforcement Learning. 5th Annual Conference on Learning for Dynamics and Control, 2023
[7] Skill Induction and Planning with Latent Language. P. Sharma, A. Torralba, J. Andreas. ACL 2022.

Thank you for your attention!

Questions?

Robustness: direct metric optimization

- **Optimization of non-decomposable losses**
  - *e.g.* rank losses in image retrieval: Average Precision (AP), Recall@k, *etc*
- **Decomposability gap:** $DG_{\mathrm{AP}}(\boldsymbol{\theta}) = \frac{1}{K}\sum_{b=1}^{K} \mathrm{AP}_i^b(\boldsymbol{\theta}) - \mathrm{AP}_i(\boldsymbol{\theta})$

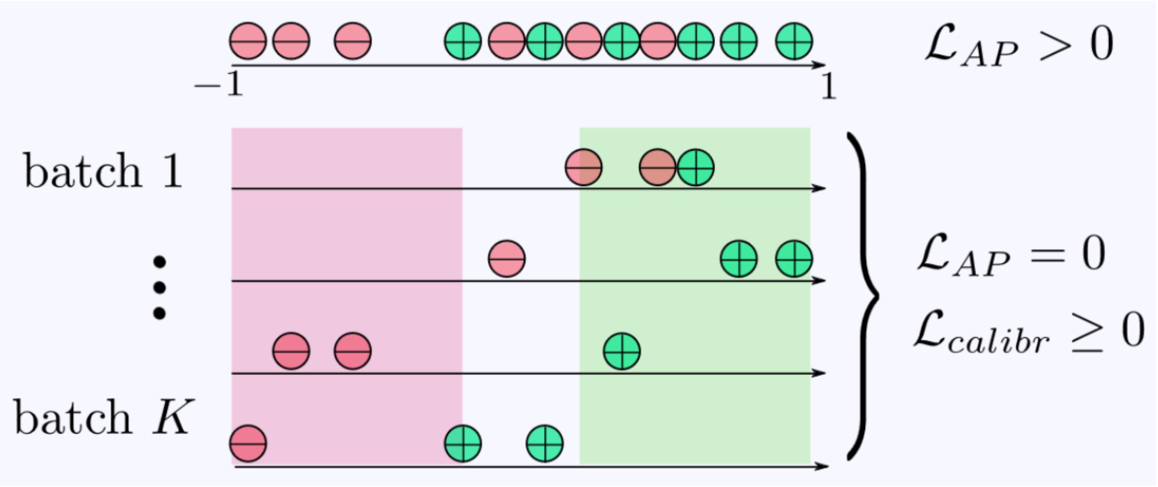

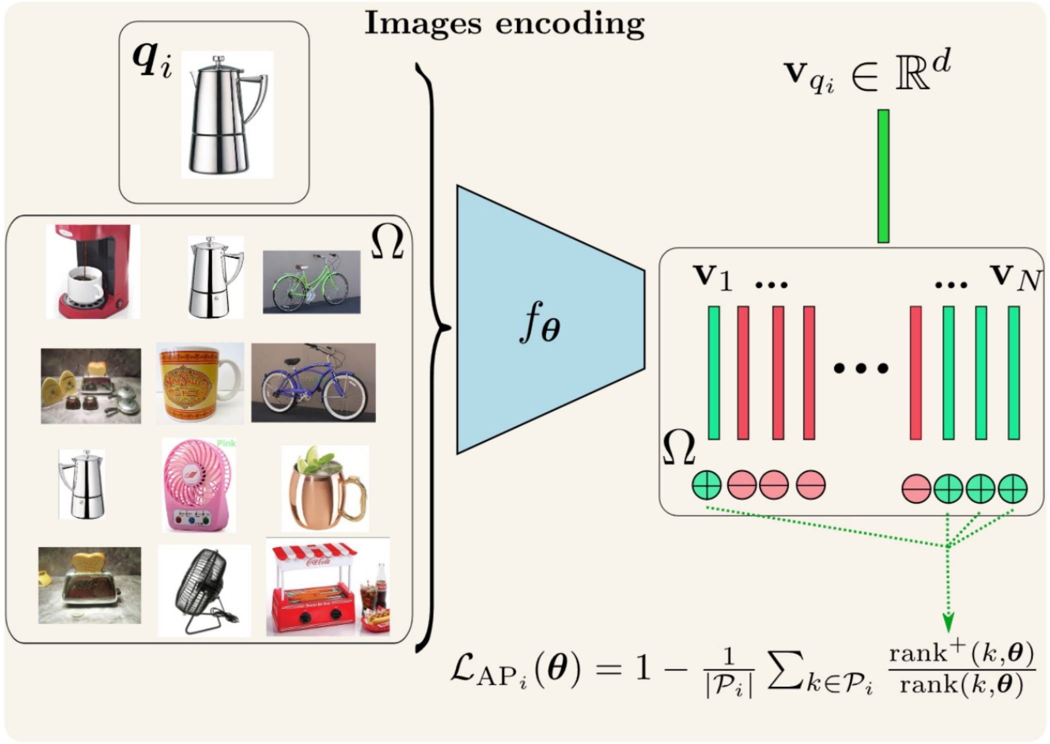

$$\mathcal{L}_{\mathrm{ROADMAP}}(\boldsymbol{\theta}) = (1-\lambda)\cdot\mathcal{L}_{\mathrm{SupAP}}(\boldsymbol{\theta}) + \lambda\cdot\mathcal{L}_{\mathrm{calibr.}}(\boldsymbol{\theta})$$

ROADMAP [RTR+21]:

- $\mathcal{L}_{\mathrm{calibr}}$: ↓ decomposability gap
- Robust wrt batch size

[RTR+21] E. Ramzi, N. Thome, C. Rambour, N. Audebert, X. Bitot. Robust and Decomposable Average Precision for Image Retrieval. NeurIPS 2021.

Optimization of hierarchical metrics

- **Extension to hierarchical metrics**
- **HAPPIER [RAT+22]:** Hierarchical Average Precision (H-AP) Training
  - **Non-binary relevance**, *e.g.* semantic tree between concepts

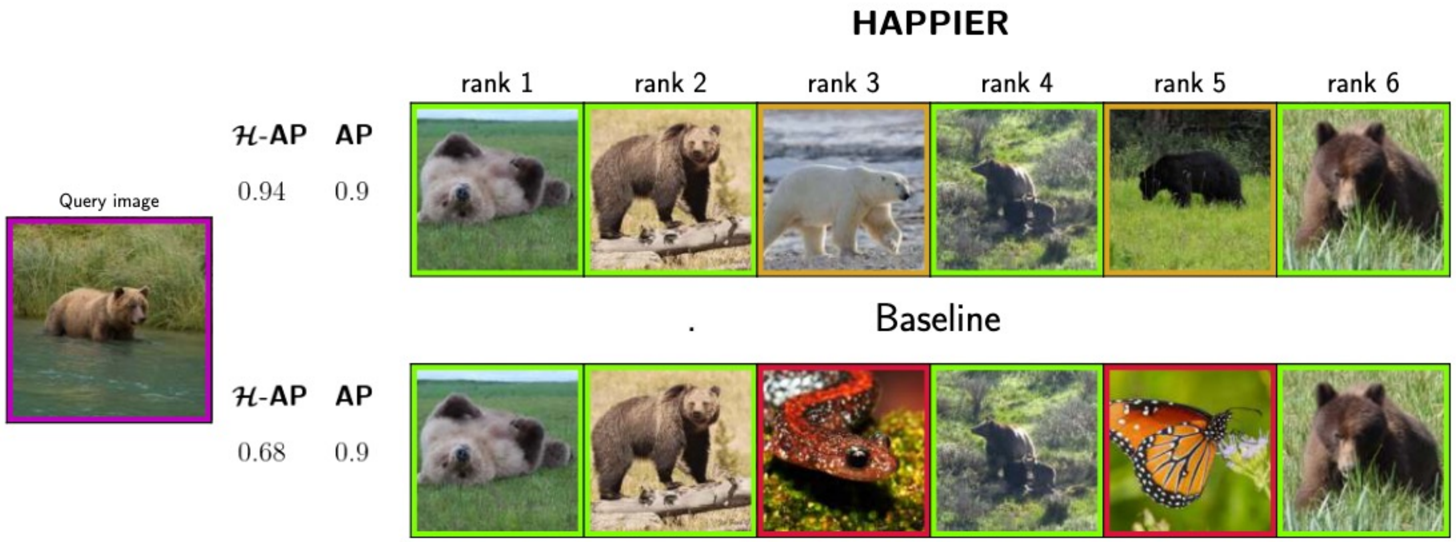

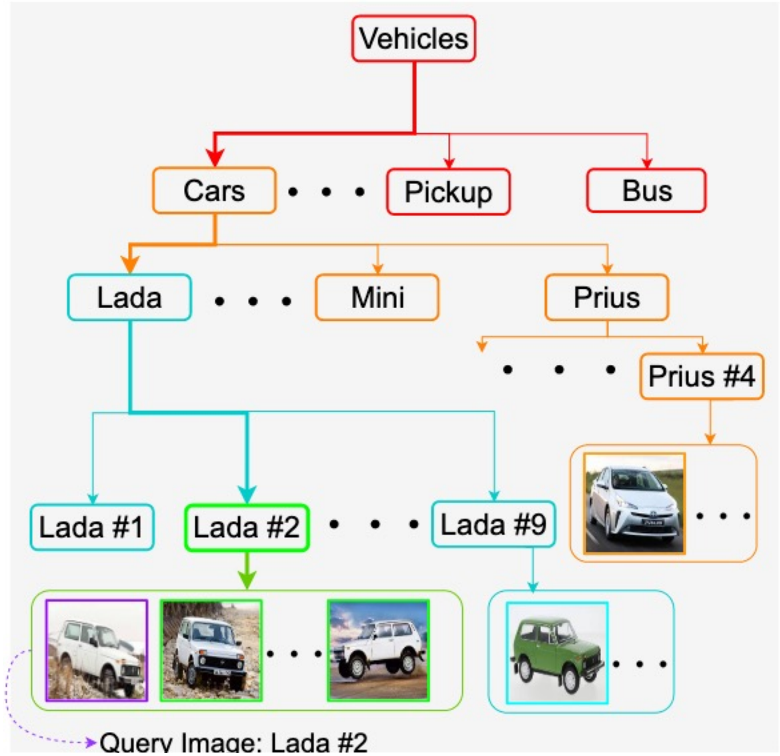

- Flat AP training: same score
- HAPPIER: favors top list

[RAT+22] E. Ramzi, N. Audebert, N. Thome, C. Rambour, X. Bitot. Hierarchical Average Precision Training for Pertinent Image Retrieval. ECCV 2022.

Data driven dynamics for CFD: EAGLE [JBN+23]

- Large-scale Learning of Turbulent Fluid Dynamics with Mesh Transformers
  - New dataset for turbulent fluid dynamics
  - New mesh-transformer architecture with graph pooling

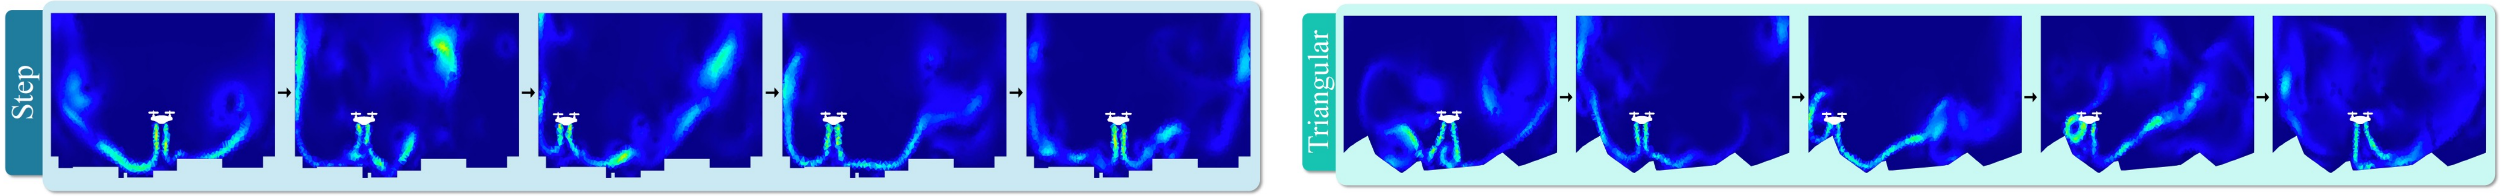

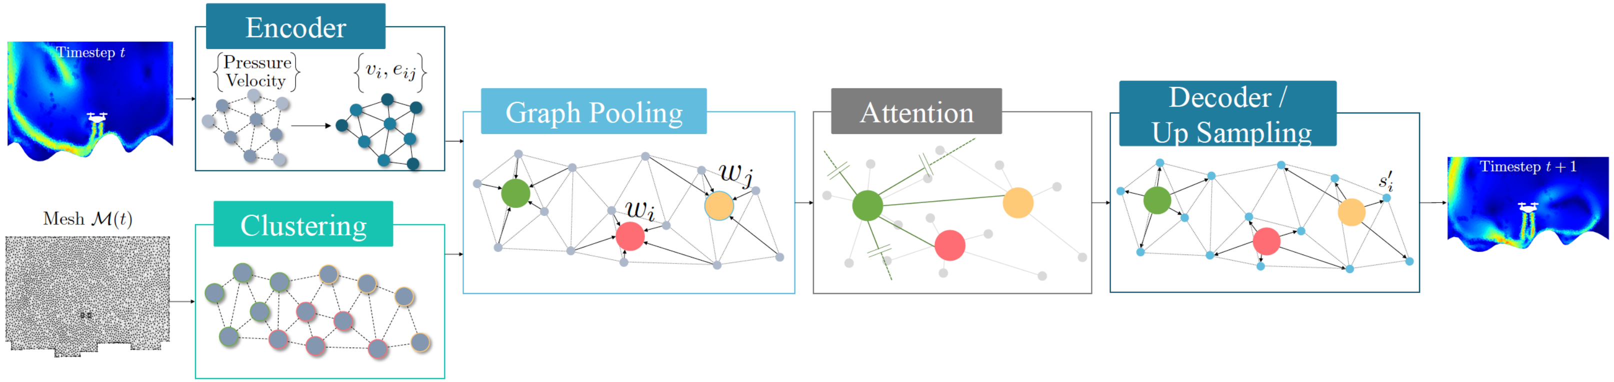

[JBN+23] S. Janny, A. Bénéteau, M. Nadri, J. Digne, N. Thome, C. Wolf. Large-scale Learning of Turbulent Fluid Dynamics with Mesh Transformers. ICLR 2023.