



# PR[AI]RIE

PaRis Artificial Intelligence Research InstitutE

Isabelle Ryl and Jean Ponce

<https://prairie-institute.fr/>



## A bit of history

- February 14, 2017: **The idea !**
- March 28, 2017: The name is picked
- October 25, 2017: The first industry partner
- March 28, 2018 : The Villani report
- March 29, 2018 : President Macron announces the French national plan for AI
- April 24, 2019: Four 3IAs are selected, ANITI, Côte d'Azur, MIAI, and PRAIRIE



- September 1, 2019: **PRAIRIE is inaugurated !**

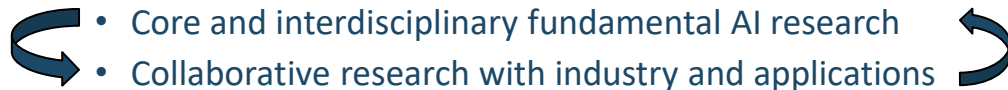


# Objectives

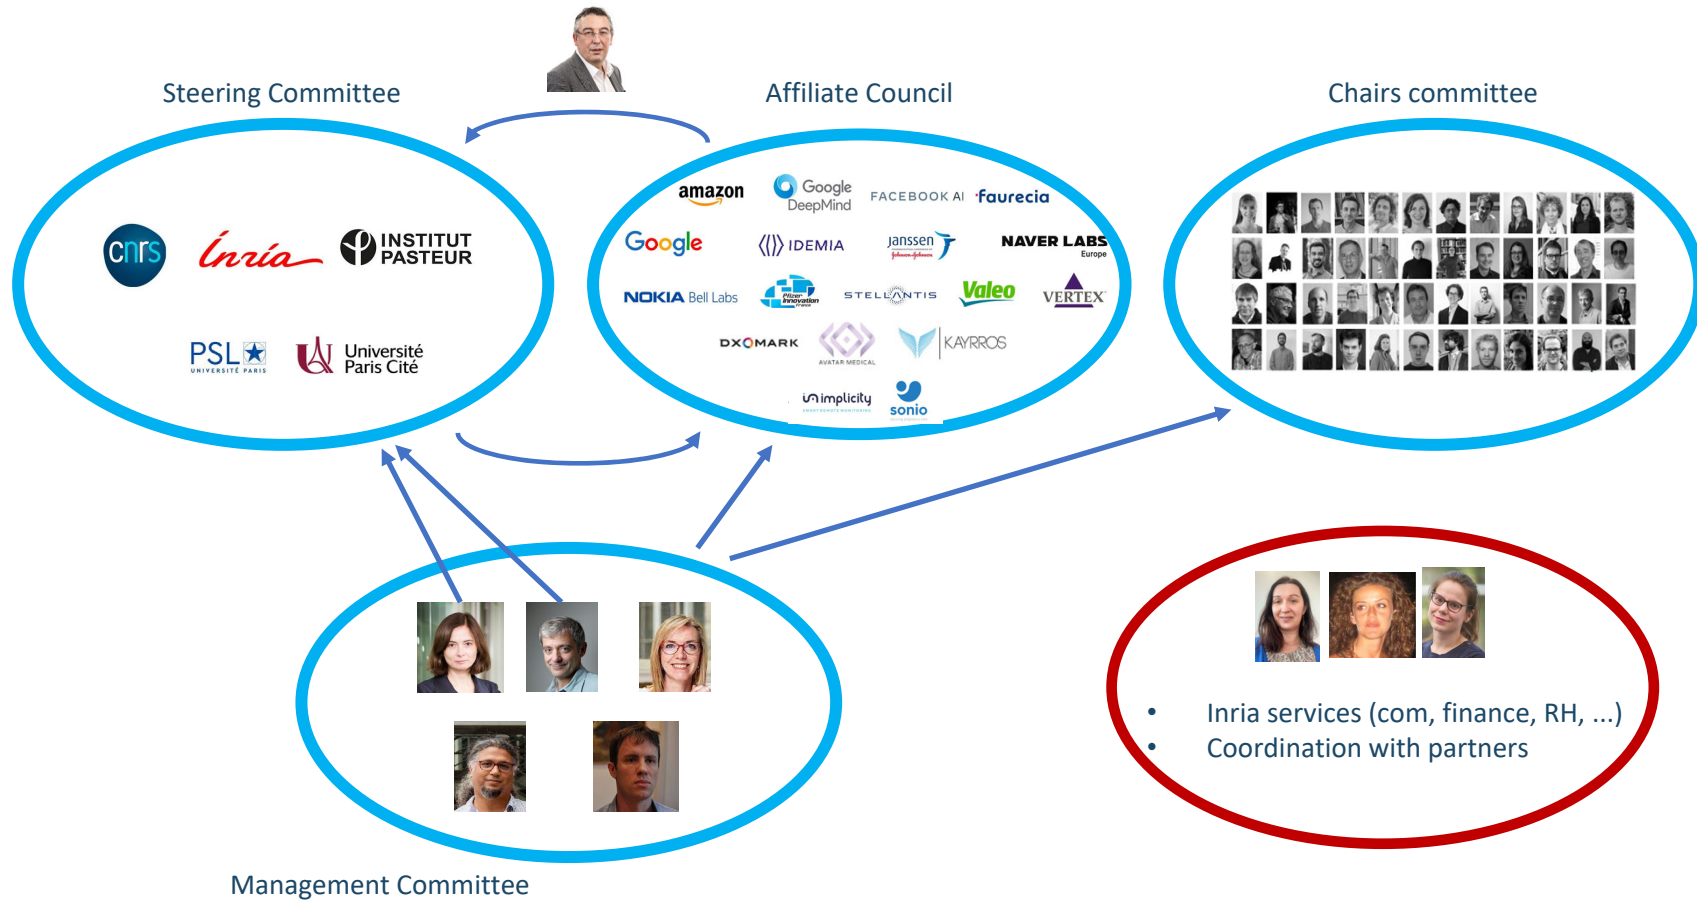
Ambition : To be a recognized international leader in AI research and education, with true socio-economic impact

## Strategy

- Excellence in research and education
  - Individuals: bring together, train, attract and retain top talents
  - Academic, industrial, and international partners
- A virtuous circle
  - Core and interdisciplinary fundamental AI research
  - Collaborative research with industry and applications
- Funding model:
  - Packages for individual chair holders from government and academia partners
  - Gifts and bilateral contracts (CIFRE PhDs, post-docs) from industry partners



# Governance





# Main focus: Science

## Core of the institute: the chairs

- Excellence is the main criterion (46 chairs, 6 members of the French Academies of Sciences and Engineering, 15(+3) ERCs, mean h-index 43, mean citations 17K)
- Research on core AI, at the interface with other disciplines, related to applications
  - **Core research:** Machine learning and optimization, autonomous agents and multi-agent systems, networked data management, cognition
  - **Integration:** Computer vision, data science, natural language processing, robotics, statistical physics, agents, HCI, NLP, psychology
  - **Applications:** Biological imaging, medical imaging, clinical decision support, genomics
  - Open to SSH, ethics, and law

## Applications fields

- Health and transportation but not just these

## Outreach

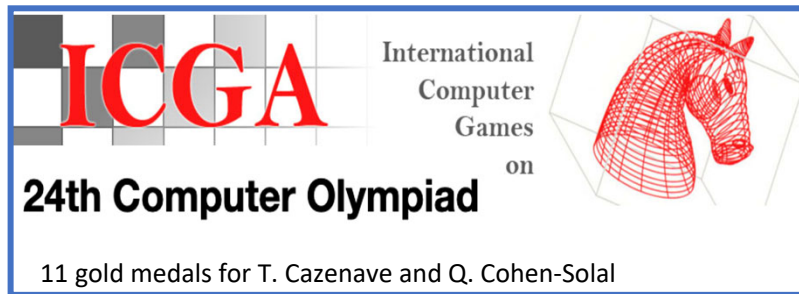
- Focus on diversity



# A few recent highlights

## Global Assessment of Oil and Gas Methane Ultra-Emitters

Lauvaux et al., *Science*, 375(6580):557-561, 2022  
with A. d'Aspremont and KAYRROS data



**ICGA** International Computer Games on  
**24th Computer Olympiad**  
11 gold medals for T. Cazenave and Q. Cohen-Solal

## Petabase-scale sequence alignment catalyses viral discovery

Edgar et al., *Nature*, 602:142-147, 2022, with R. Chikhi

OSCAR 21.09 (Sagot et al., 2021): The largest multilingual corpus available today: 150 languages, 3 TB of raw text in English, 340 GB in French, but also 14 GB in Slovak, 900 MB in Basque Basque and 49 MB in Breton!

End of 2021: 23 papers at NeurIPS, one outstanding paper award co-signed by Francis Bach & Laurent Massoulié, the 2021 test-of-time award to Francis Bach

### Online Learning for Latent Dirichlet Allocation

Matthew D. Hoffman  
Department of Computer Science  
Princeton University  
Princeton, NJ  
mdhoffma@cs.princeton.edu

David M. Blei  
Department of Computer Science  
Princeton University  
Princeton, NJ  
blei@cs.princeton.edu

Francis Bach  
INRIA—Ecole Normale Supérieure  
Paris, France  
francis.bach@ens.fr

## Catala: A Programming Language for the Law

Along with D. Merigoux, K. Barghavan has proposed Catala, a programming language for the translation of statutory laws into executable implementations, with a compiler proven to be correct using the F\* proof assistant. Catala has notably revealed a bug in the official implementation of French laws regulating family benefits.



# Overall philosophy

▣ To create a community that is more than the sum of its members

Beyond the obligation in the 3IA call to find industrial funds:

- To convince industry groups to make the bet to jointly fund, in the form of gifts, so-called “fundamental” research
- To convince the best AI researchers to talk with their industry colleagues and attack their problems

All of this within an open and multi-sector framework that offers many opportunities for cross-pollination:

- Allowing researchers to consider various application domains
- Revealing the convergence of issues between different sectors
- Fostering discussions between the industry groups themselves, along with the PRAIRIE chairs, with new application opportunities

▣ This was a gamble, but the idea caught on and the community is taking shape



# Scientific events



## Summer/winter schools

- PAISS co-organized with MIAI open to students, academics and industry (about 60% 15% 25%), 25% women
  - 2019 (Paris): 300 participants (including 22 students having received a grant, from 13 different countries, including 13 women financed by industrial partners of MIAI and PRAIRIE)
  - 2021 (Virtual) : 300 participants from 52 countries, 27% of women, youtube 726 unique viewers
- AI4Health School (2021) co-organized with HDH, MIAI and 3IA Côte d'Azur

## Seminars

- Colloquium PRAIRIE (webinar since covid, maximum of 225 simultaneous connections)
- Seminar ENS « Digital Humanities / Artificial Intelligence », <https://dhai-seminar.github.io/> (G. Peyré et T. Poibeau)

Workshop in honor of Jean-Paul Laumond, Collège de France, July 11, 2022

<https://prairie-institute.fr/workshop-in-honor-of-jean-paul-laumond-college-de-france/>



# Education

Education programs are proposed by



▀ The role of PRAIRIE? Catalyst

**The dynamics around PRAIRIE provides**

- visibility
- researchers, AI specialists, who are committed to teaching
- master grants
- building

**and creates a domino effect for experimentations and new pro**

**parisauté  
campus**

- cofund AI for the sciences (PSL): 26 PhD grants for international students
- setting up an “à la carte” program for Société de l'Industrie Automobile, Société Générale, ...
- bootcamp and SHS Data Challenge
- new “partners”: Applied Computational Social Sciences Institute Data-Intensive Governance, Institut Santé Numérique en Société

**and second phase of the national strategy on AI... EFELIA - PRAIRIE**



# European network

France: 3IA + chairs

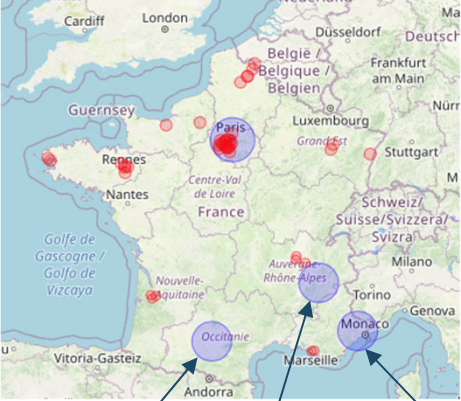
<https://instituts-3ia.fr>

Île-de-France

DIM AI4IDF

Europe

<https://ellis.eu>



# A network of international collaborations



- CIIRC Prague
- MPI Tübingen
- Turing Institute London
  
- BAIR UC Berkeley
- CDS New York University
- RI Carnegie-Mellon University
- MILA Montréal
- RIKEN AIP Tokyo
  
- Health
- AI Singapore (AISG)
- GMHC Manchester
- HMS Harvard

And this is why we are here ! To build up another bridge for collaboration !






# Beyond the Computer Vision Comfort Zone

Jean Ponce

Inria, DI/ENS (CNRS, Inria, ENS-PSL), PRAIRIE

Courant Institute and Center for Data Science, New York University

With M. Alakuijala, A. Bardes, T. Bodrito, O. Bounou, T. Chabal, N. Chahine, T. Eboli, O. Flasseur, B. Lecouat, G. Le Moing, Y. de Mont-Marin, X. Sun, E. Vincent, H.V. Vo, and M. Aubry, J. Carpentier, A.-M. Lagrange, M. Langlois, Y. LeCun, J. Mairal, P. Pérez, C. Schmid, O. Siméoni, J. Sun, Y. Wang



# A lot has been achieved in the past 10 years!




Ex: Mask R-CNN (He et al., 2018)



# Supervision: Where do the labels come from?

- A trend toward manually annotating the whole wide world with crowd sourcing
- Example: MS COCO (Lin et al., 2015) :328K images of 91 object categories



- Scaling up means a lot less supervision
- Beyond benchmarks, the real world
- And what we know about it

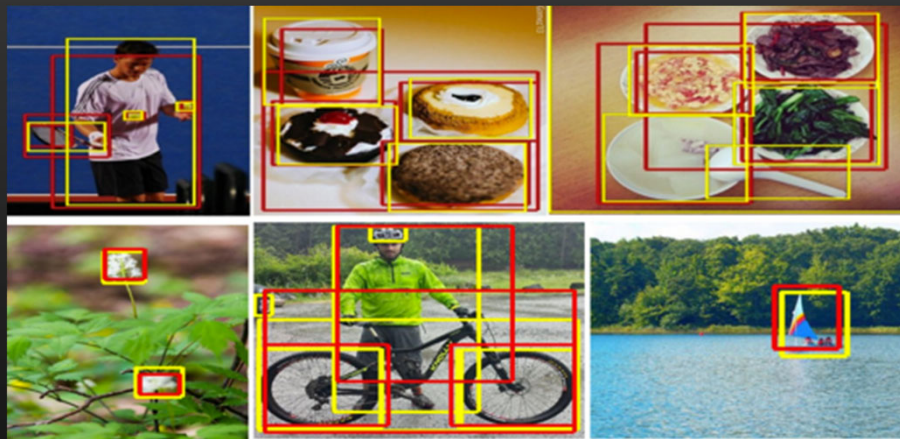
# Unsupervised object discovery

CIFRE PhD thesis of H.V. Vo with  
P. Pérez at Valeo



$$\max_{x,e} \sum_{p=1}^n \sum_{q \in \mathcal{N}(p)} e_{pq} x_p^T S_{pq} x_q \text{ such that}$$
$$\sum_{k=1}^r x_p^k \leq \nu \sum_{q \neq p} e_{pq} \leq \tau \quad \forall 1 \leq p \leq n$$

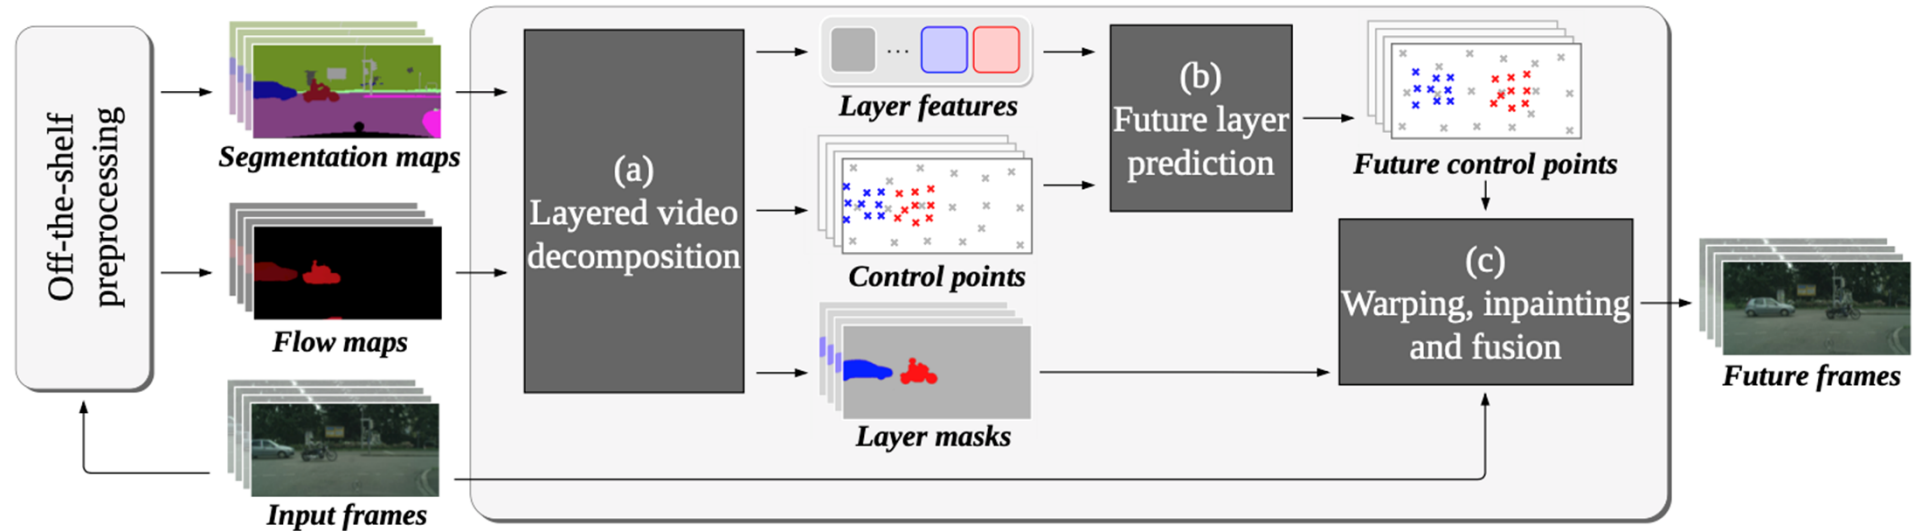
Heuristics → Combinatorial  
optimization → Ranking  
approach → 1.7M images



- (Cho et al., CVPR'15)
- (Vo et al., CVPR'19)
- (Vo et al., NeurIPS'21)
- (Siméoni et al., BMVC'21)


Perhaps more realistic: an active learning approach to weakly-supervised object detection (Vo et al., ECCV'22).

# Self-supervised video prediction (LeMoing et al., NeurIPS'21, 2023)



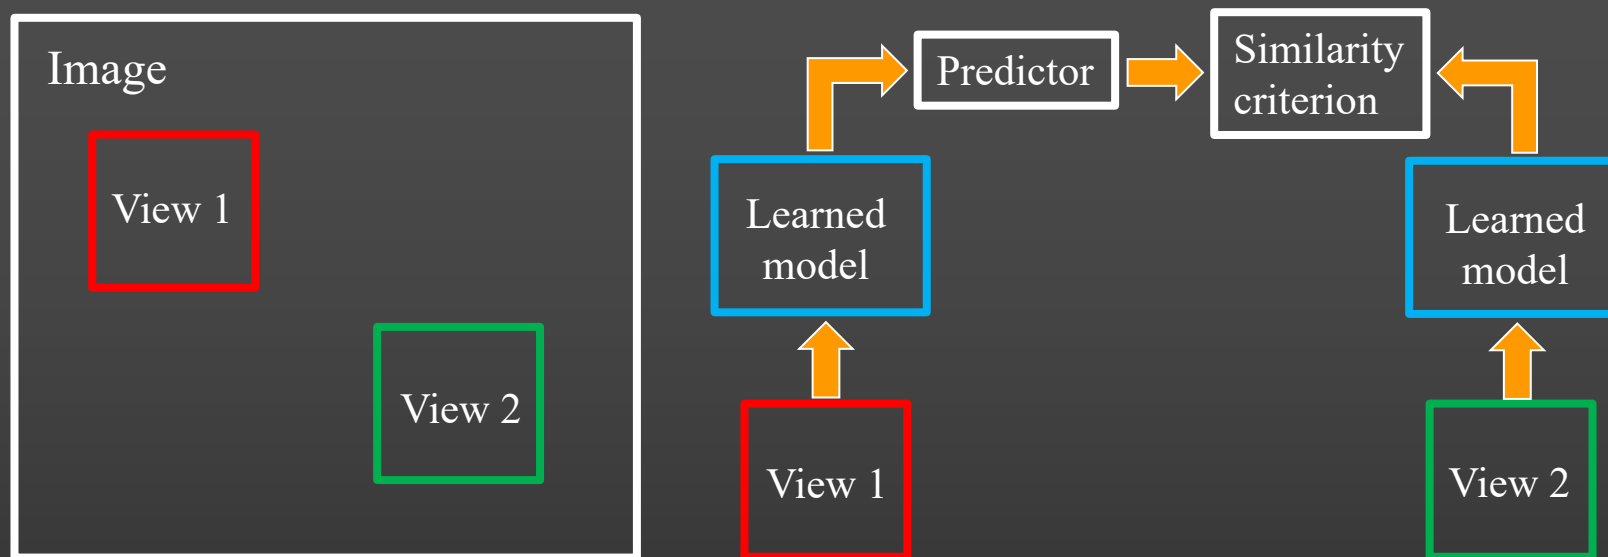
## Prediction for control

- (Koopman, PNAS, 1931)
- (Budisić et al., Chaos'12)
- (Mezić&Banaszuk, Phys. D'12)
- (Bounou et al., NeurIPS'21)
- (Bounou et al., 2022)

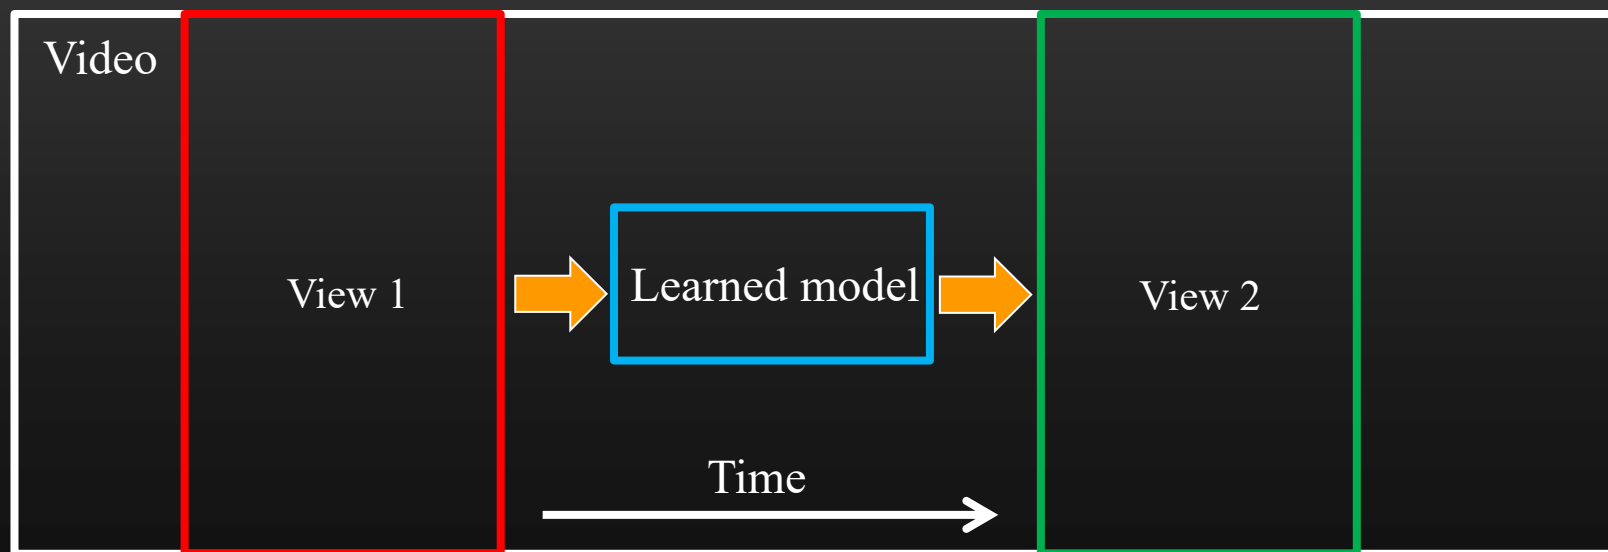


# Instances of self-supervised learning

CIFRE PhD thesis of A. Bardes  
with Y. Le Cun at FAIR



(Bardes et al., ICLR'22, NeurIPS'22)






Video prediction from 4 frames (Le Moing, Ponce, Schmid, 2023)

Next: using outside control signal and video compression




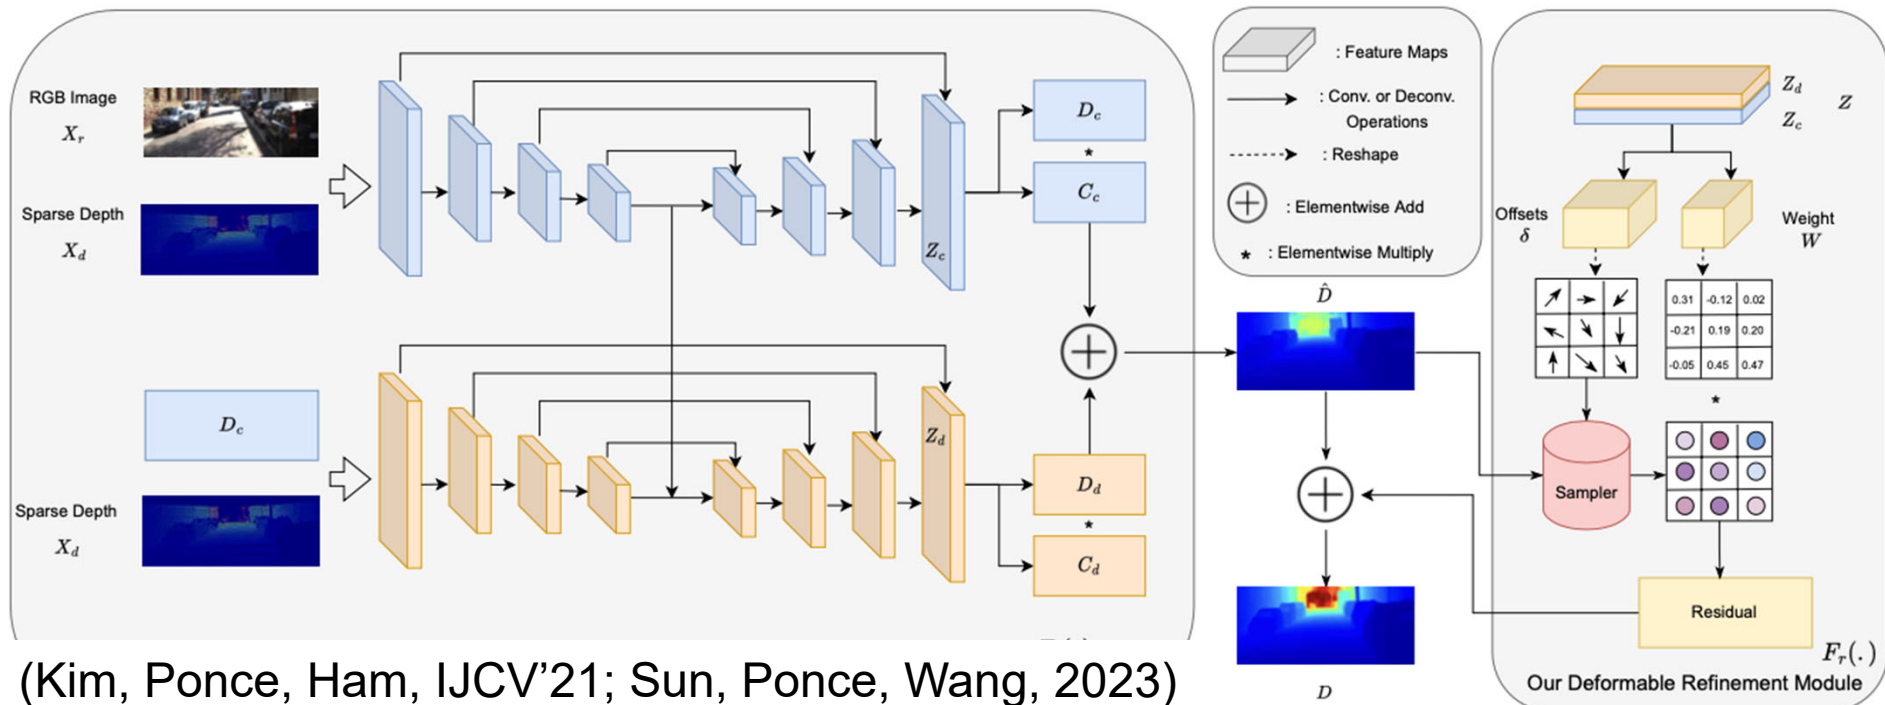
# Lidar data and the KITTI dataset (Geiger et al., 2017)



<https://velodynelidar.com/products/hdl-64e/>

Kitti depth completion task







Model	RMSE ( $\downarrow$ )	MAE ( $\downarrow$ )	iRMSE ( $\downarrow$ )	iMAE ( $\downarrow$ )
DFuseNet [51]	1206.66	429.93	3.62	1.79
CSPN [15]	1019.64	279.46	2.93	1.15
PwP [14]	777.05	235.17	2.42	1.13
FusionNet [30]	772.87	215.02	2.19	0.93
DSP [18]	766.74	220.36	2.47	1.03
MSG-CHN [34]	762.19	220.41	2.30	0.98
DeepLiDAR [13]	758.38	226.50	2.56	1.15
UberATG [52]	752.88	221.19	2.34	1.14
ACMNet [53]	744.91	206.09	2.08	0.90
CSPN++ [17]	743.69	209.28	2.07	0.90
GuideNet [31]	736.24	218.83	2.25	0.99
PENet [20]	730.08	210.55	2.17	0.94
<b>Ours</b>	<b>728.31</b>	<b>204.60</b>	<b>2.05</b>	<b>0.89</b>



# Lidar data and the KITTI dataset (Geiger et al., 2017)

From point clouds to ribbons (Sun, Heurtevent, Ponce, Wang, 2022)




And soon: "PRAIRIE's Driver in Context" public dataset:

- A Forvia, Stellantis, Valeo, PRAIRIE initiative
- Synchronous inside/outside data, Lidar & RGB
- Start acquisition end of 2022
- Release at ICCV'23 in Paris!

# The PIQ23 dataset for portrait quality assessment (Chahine et al., CVPR'23)

CIFRE PhD thesis of N. Chahine  
with S. Ferradans at DXOMark



First (as far as we know) GDPR compliant portrait data with informed consent from models

**Privacy and GDPR compliance:** We present a new portrait dataset consisting of smartphone images of people that were shot by internal Photographers. This dataset is intended for research use only. All the people appearing in the photos have explicitly granted permission for image rights. Thus, each individual has signed a transfer of image rights and has been provided with a privacy notice detailing how their image will be processed. Also, to ensure the effectiveness of people's rights, we have tagged each photo with a unique identifier assigned to each person by using a face clustering algorithm. This pseudonymisation technique prevents access to individuals' names by dataset users. That said, some scenes might include passing-by people in the background; since getting their photo rights is impossible, their faces have been blurred. Finally, we contractually require all dataset users to comply with relevant data protection laws, including the GDPR.

**Ethnic observed bias:** The annotations of the portraits have been done by comparing pairs of images of the same scene with different models, where the observer needs to choose the best image according to certain criteria (rendition of fine details, colour accuracy, exposition of the face, overall quality of the photo). These are photographic-base definitions that in the annotation guidelines were defined technically and precisely. Once every pair has been annotated, we obtain a linear scale of quality that takes into account all the comparisons made. For some scenes, we observed that deep skin models were mostly placed in the lower quality region. Cameras and AI models present in cameras, are known to be less performant on deep skin tones, and this bias might be related to this issue intrinsic to the devices we are evaluating. On the other hand, people are also known to be more precise on the ethnicities they are more common in their environment. This is why we are working on adding more ethnical diversity to the annotator pool.


Photography: Deblurring  
sharp images!  
(Eboli, Sun, Ponce, 2021)



Key idea: combine physical model of image formation, classical solutions of inverse problems, and learned image priors




# Photography: Improving spatial resolution and dynamic range (Lecouat, Ponce, Mairal, ICCV'21; Lecouat, Eboli, Ponce, Mairal, SIGGRAPH'22)





Also, N. Chahine CIFRE PhD thesis with DXOMARK and A.-S. Calarasanu  
An RGPD-compliant portrait database should be out in the next year







# Exoplanet detection from direct imaging



## Data from VLT/Sphere – Angular Differential Imaging

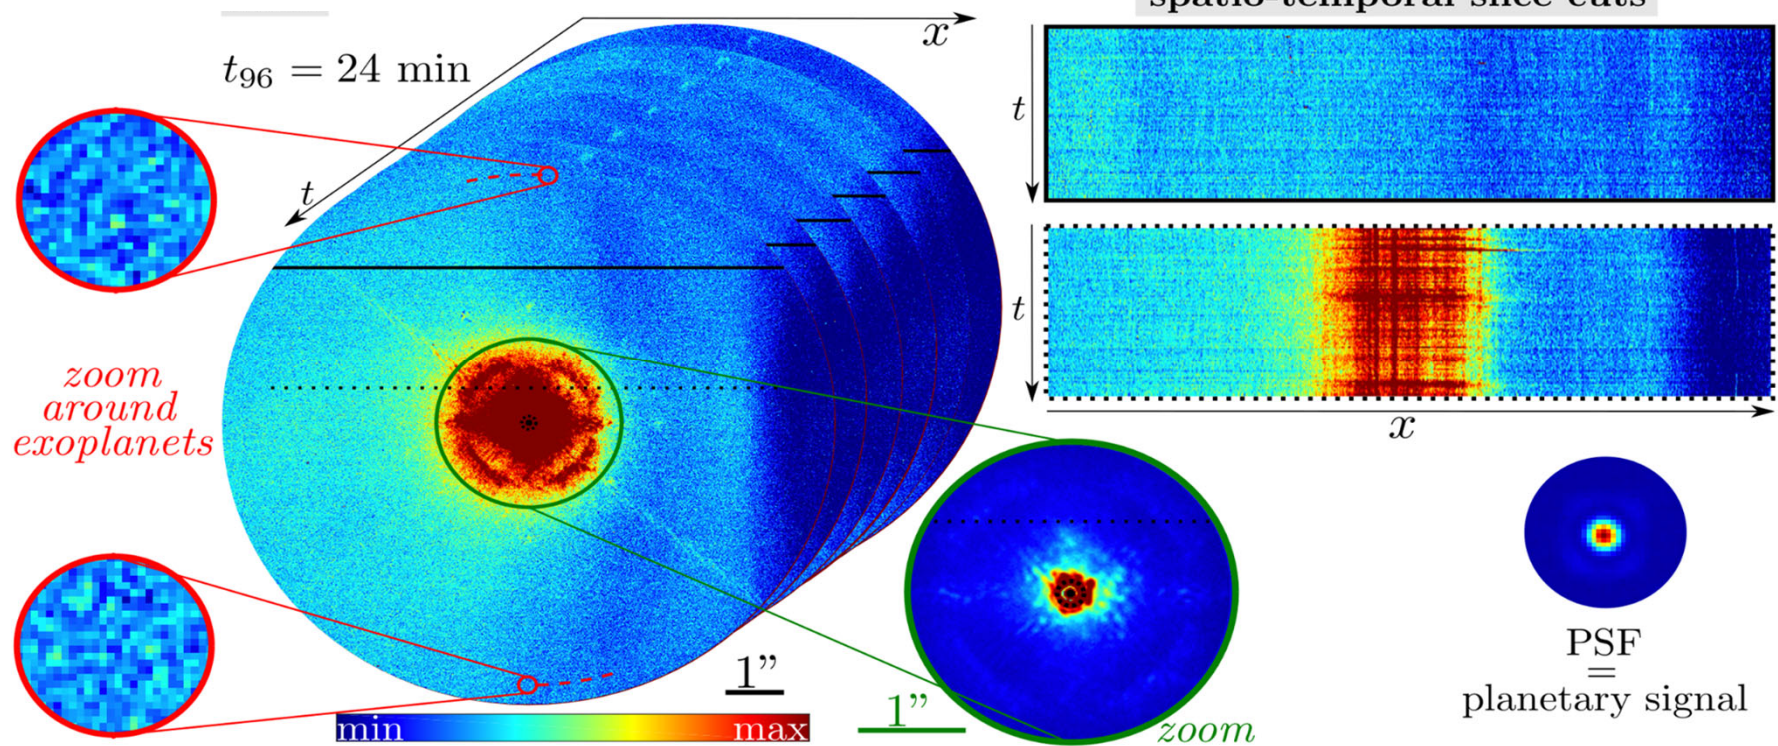


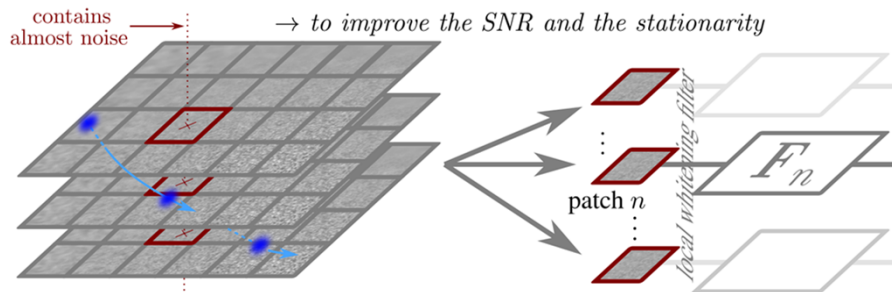
Image Credit: Olivier Flasseur

# Exoplanet detection from direct imaging - Deep PACO

T. Bodrito, O. Flasseur, J. Mairal, J. Ponce, M. Langlois, A.M. Lagrange  
PRAIRIE+MIAI+CNRS, Paris and Lyon Observatories (2022)


## Approach

1. Whitening: remove speckles structure

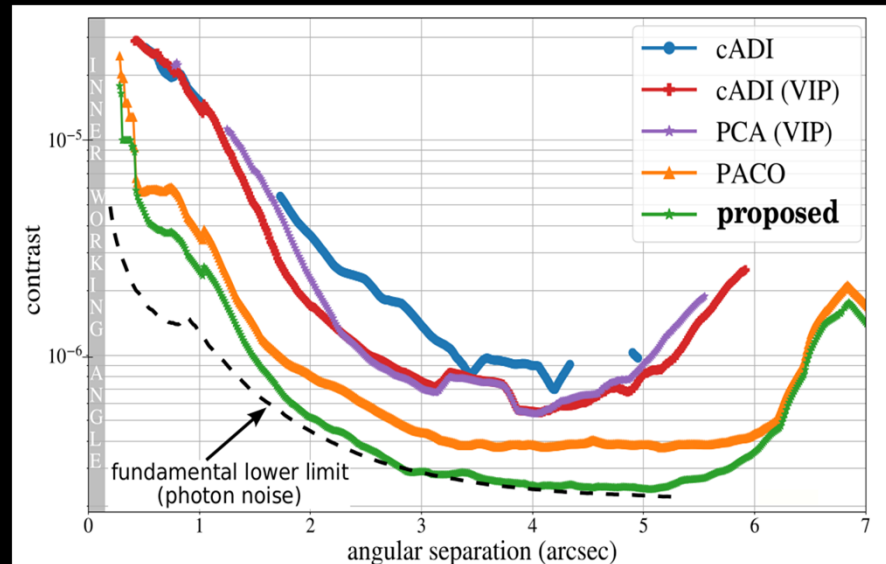


2. Derotation: sources alignment

3. Supervised Training (UNet, ResNet-18)



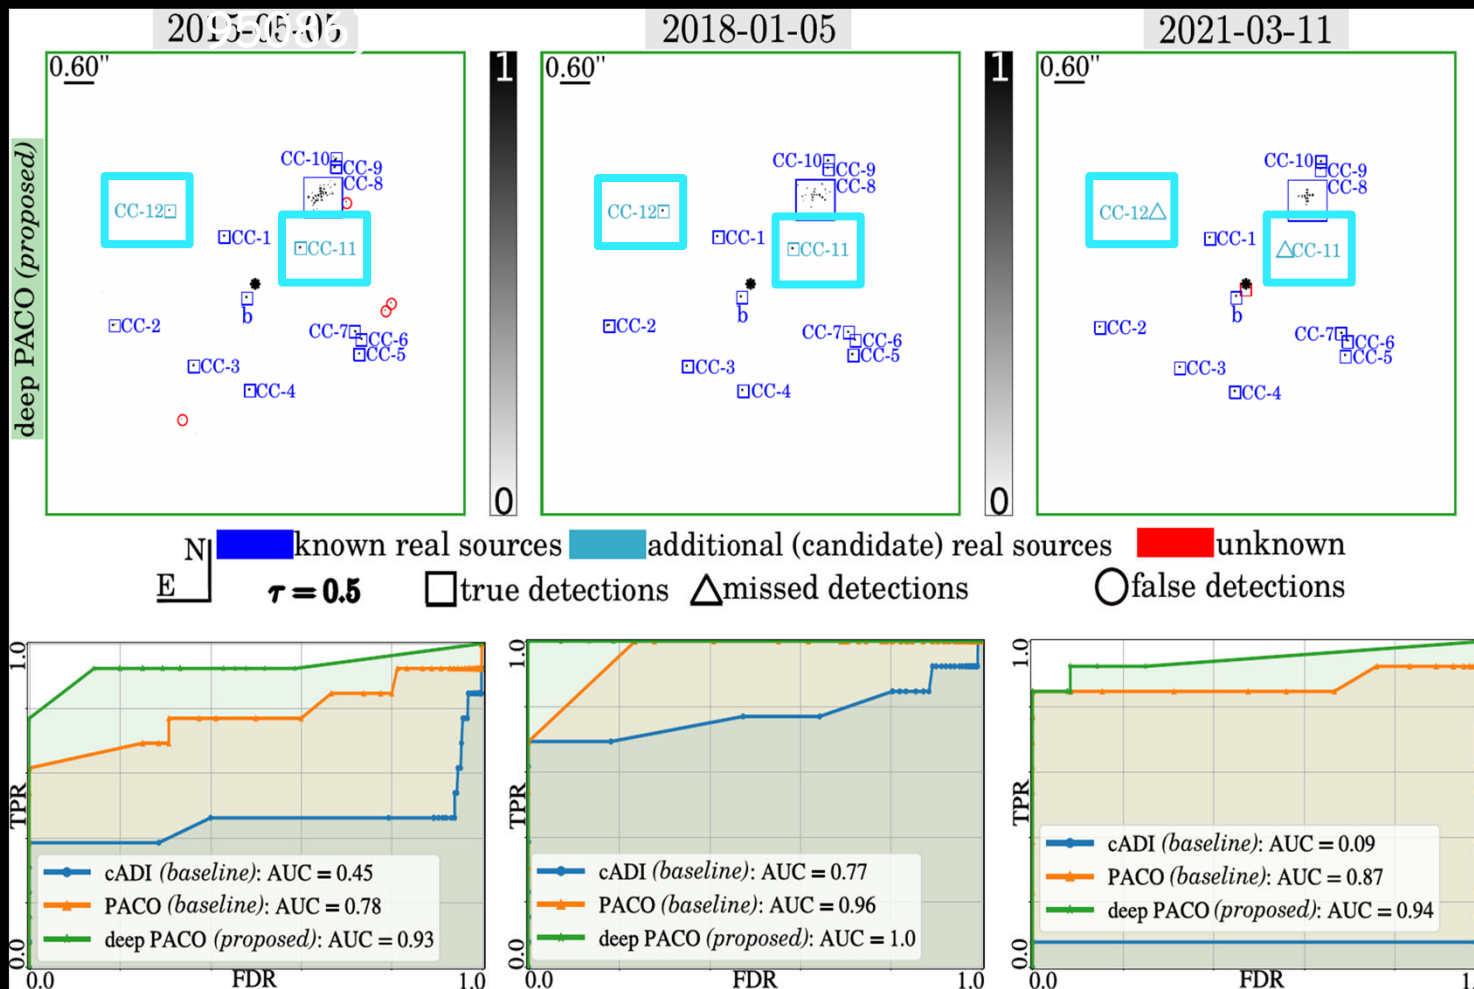
## Results on synthetic data



# Exoplanet detection from direct imaging - Deep PACO

T. Bodrito, O. Flasseur, J. Mairal, J. Ponce, M. Langlois, A.M. Lagrange  
PRAIRIE+MIAI+CNRS, Paris and Lyon Observatories (2022)


## Results on real data (HD)





Also, remote sensing with M. Aubry and E. Vincent  
and next, hopefully:

Molecular imaging (J.B. Masson, Institut Pasteur)



Palazzo Ducale, Venice (Y. Ubelmann, Iconem)



(With the kind authorization of Fondazione Musei CIVICI di Venezia)  
<https://www.grandpalais.fr/fr/evenement/venise-revelee>

# Assembly Planning from Observations under Physical Constraints

Thomas Chabal<sup>1</sup>   Robin Strudel<sup>1</sup>   Etienne Arlaud<sup>1</sup>  
Jean Ponce<sup>1,2</sup>   Cordelia Schmid<sup>1</sup>

<sup>1</sup> Inria and DIENS (ENS-PSL, CNRS, Inria), Paris, France

<sup>2</sup> Center for Data Science, New York University, New York, USA

*Inria*



PSL 



NYU

Center for  
Data Science



$$V'[\gamma] = 3.18$$

$$V'[\gamma] = 3.10$$

$$V'[\gamma] = 3.14$$

$$V'[\gamma] = 6.77$$

$$V'[\gamma] = 3.49$$

$$V'[\gamma] = 1.67$$

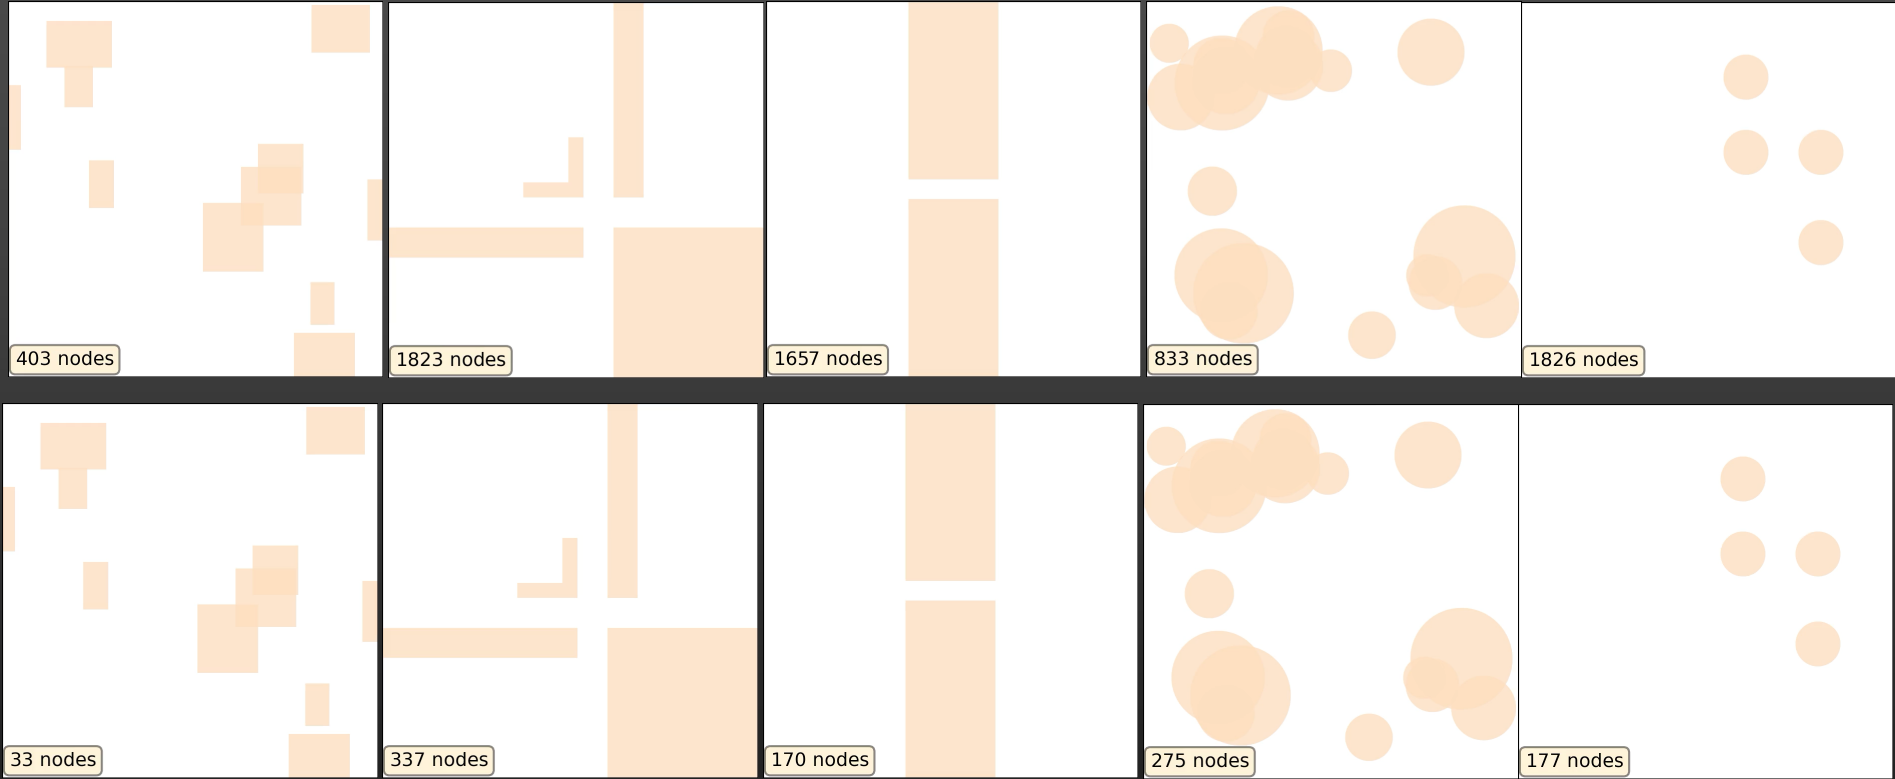
$$V'[\gamma] = 1.48$$

$$V'[\gamma] = 2.48$$

$$V'[\gamma] = 4.12$$

$$V'[\gamma] = 2.41$$





# Thank you !

Acknowledgments. This work was supported in part by the Louis Vuitton / ENS chair on artificial intelligence and the French government under management of Agence Nationale de la Recherche as part of the Investissements d'avenir program (PRAIRIE 3IA Institute).