

# Bridging Physics and Learning: application to ocean dynamics

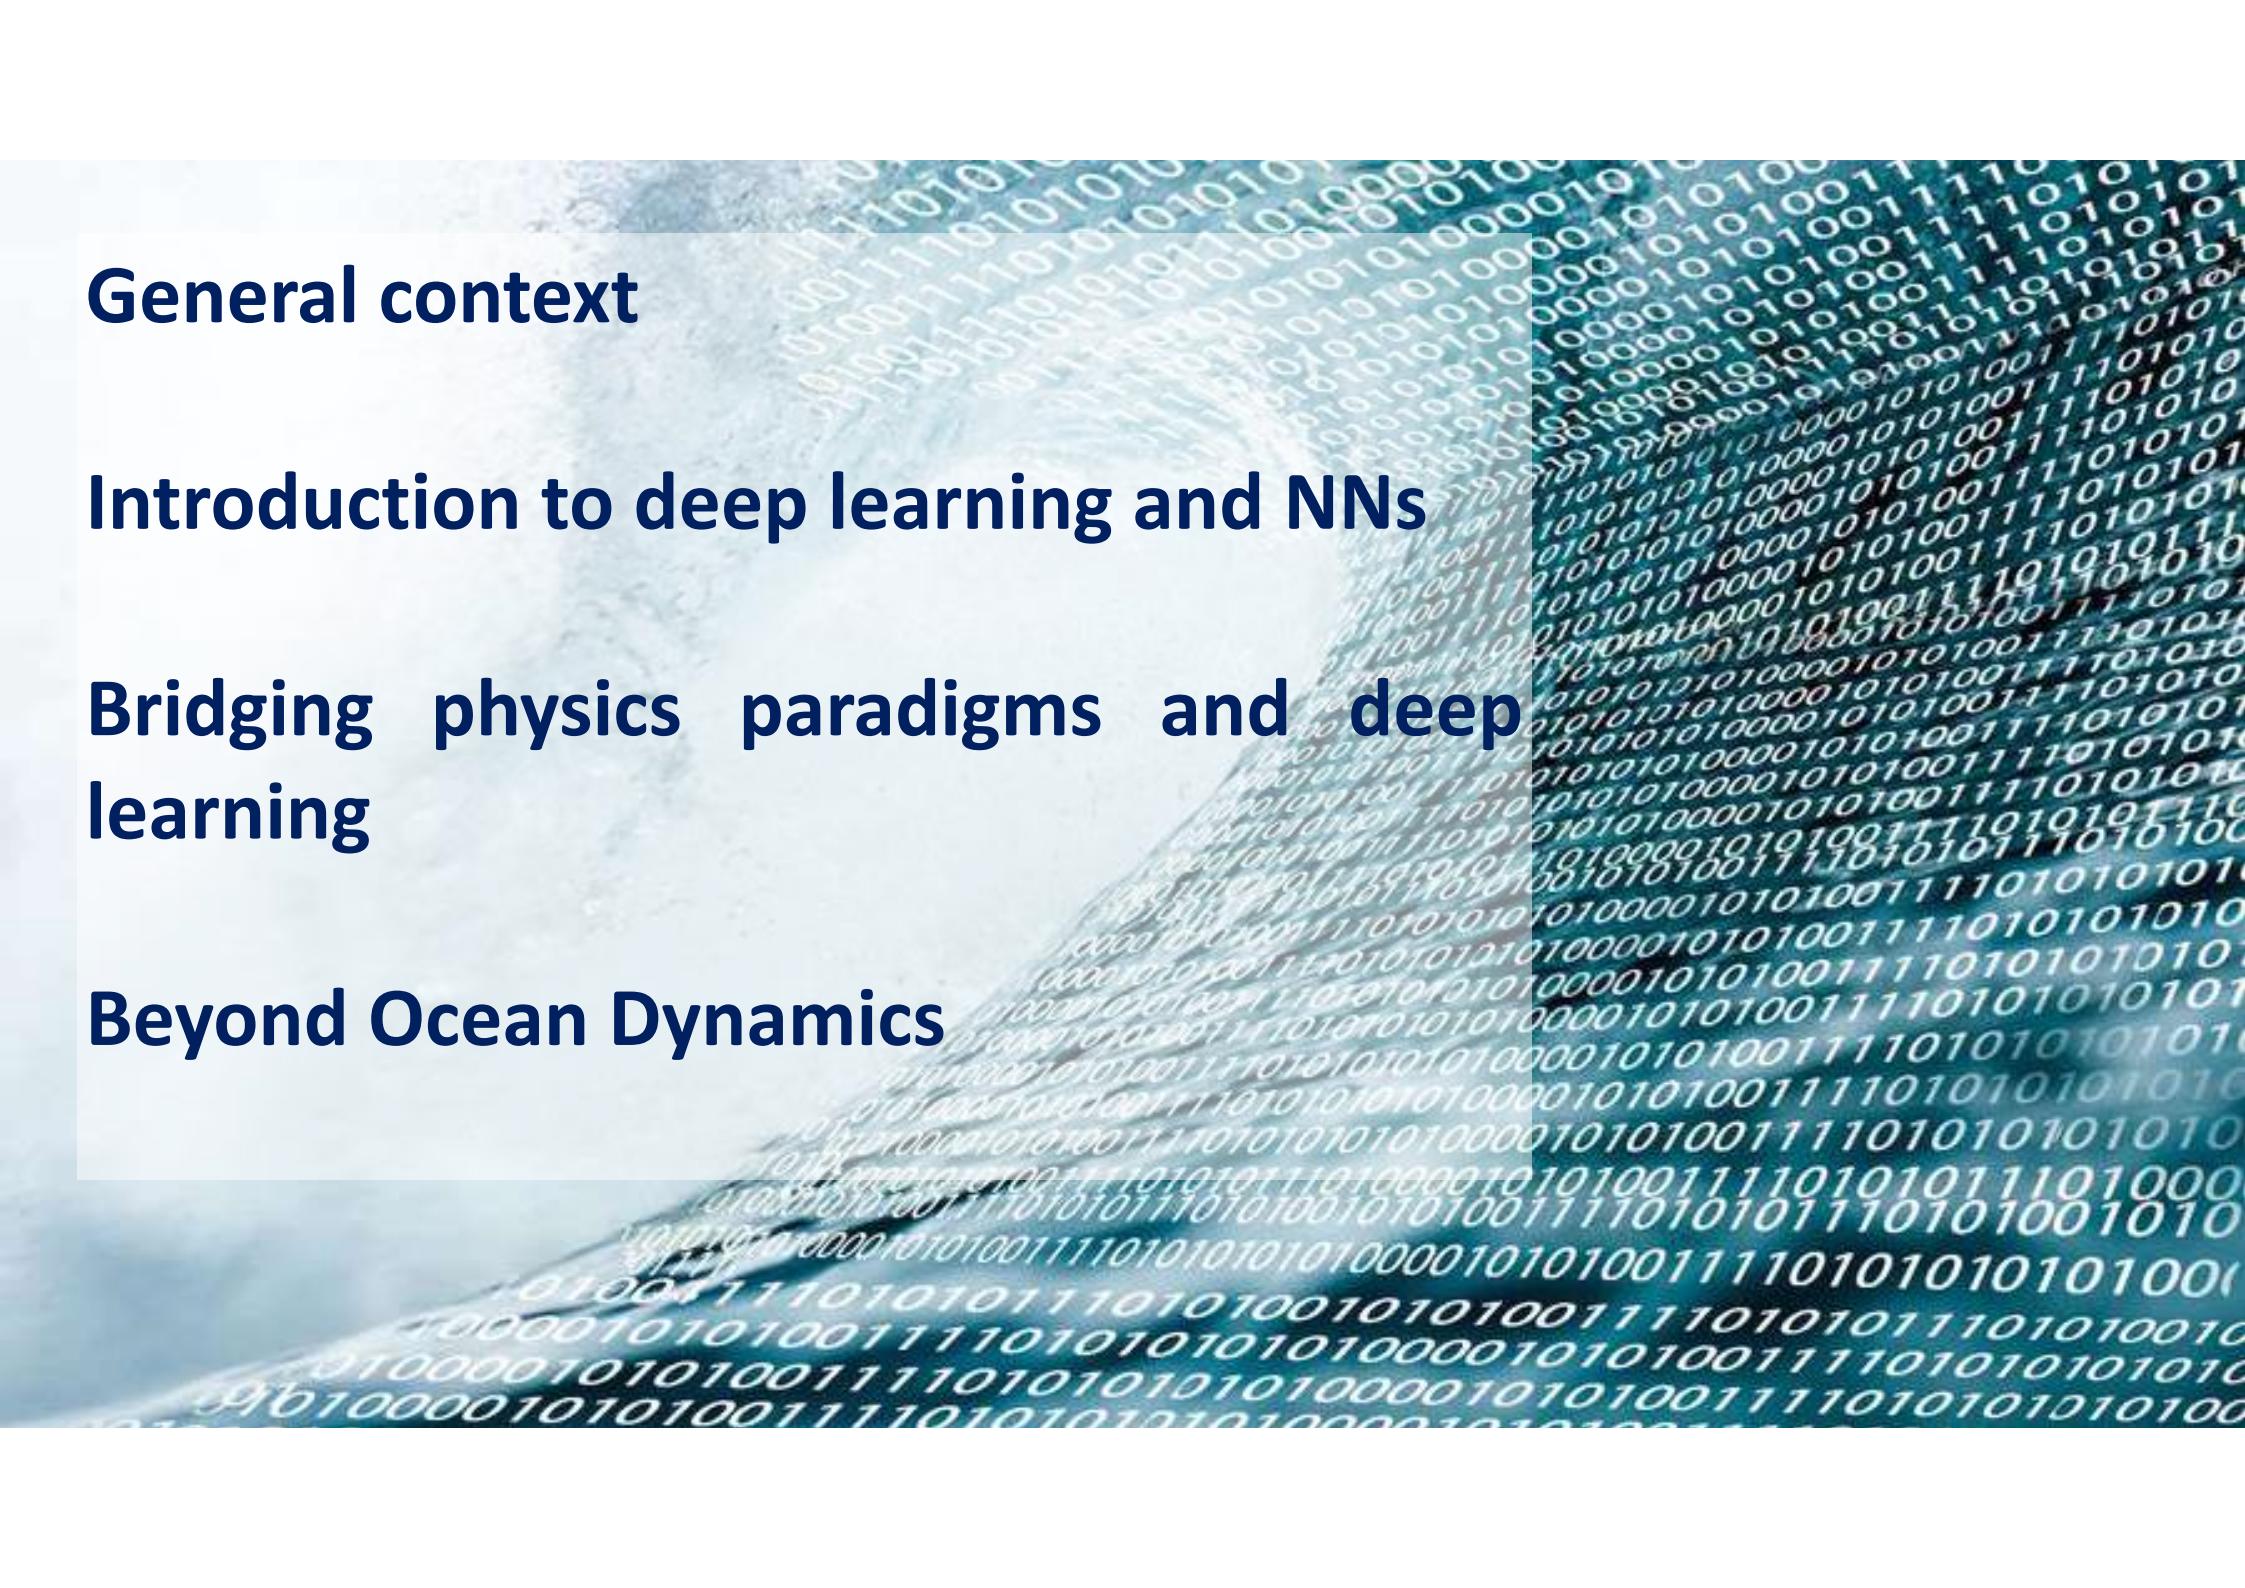
R. Fablet et al.

[ronan.fablet@imt-atlantique.fr](mailto:ronan.fablet@imt-atlantique.fr)

web: [rfablet.github.io](http://rfablet.github.io)

Webinar IMT Data & AI, July 2020





## General context

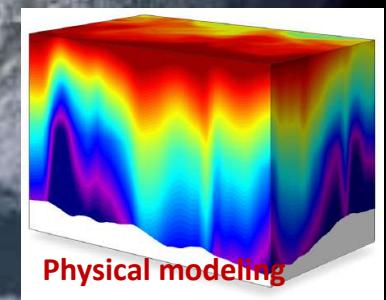
## Introduction to deep learning and NNs

## Bridging physics paradigms and deep learning

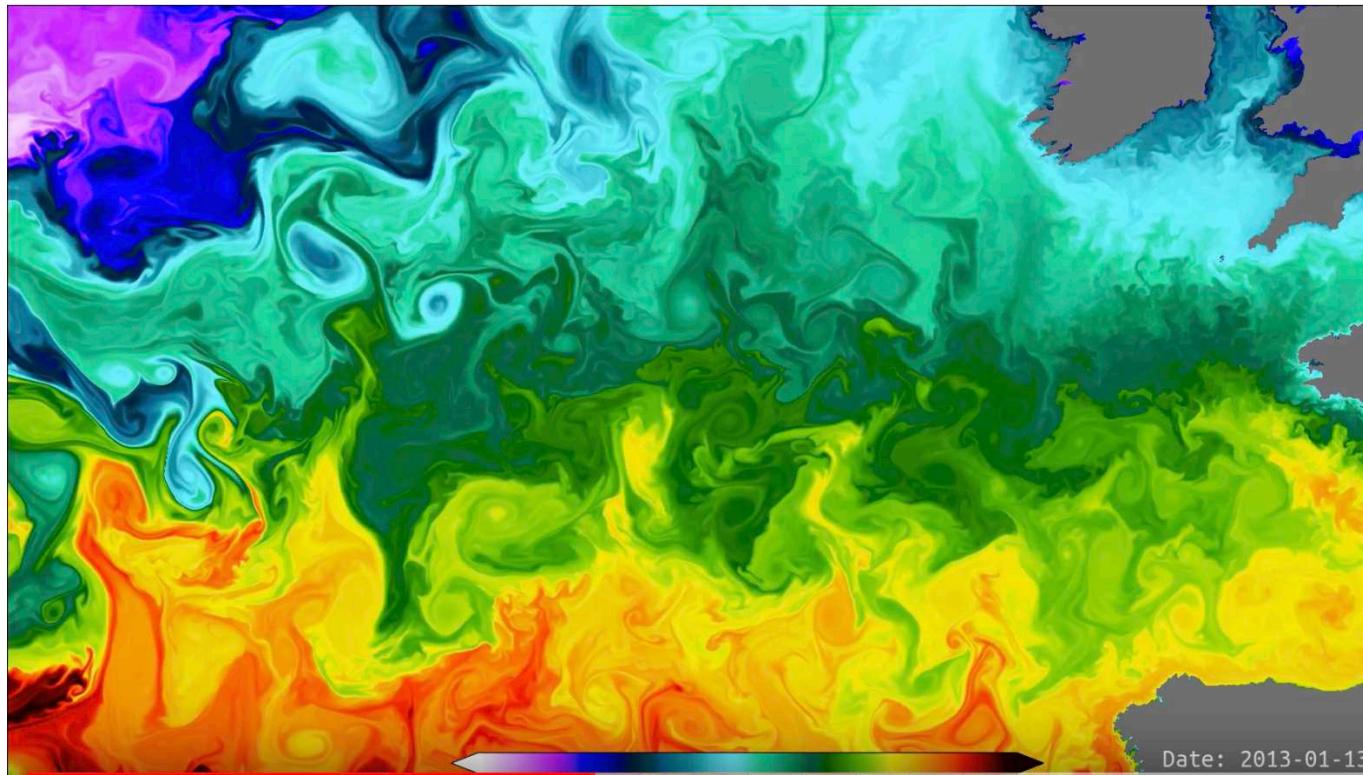
## Beyond Ocean Dynamics

# General question

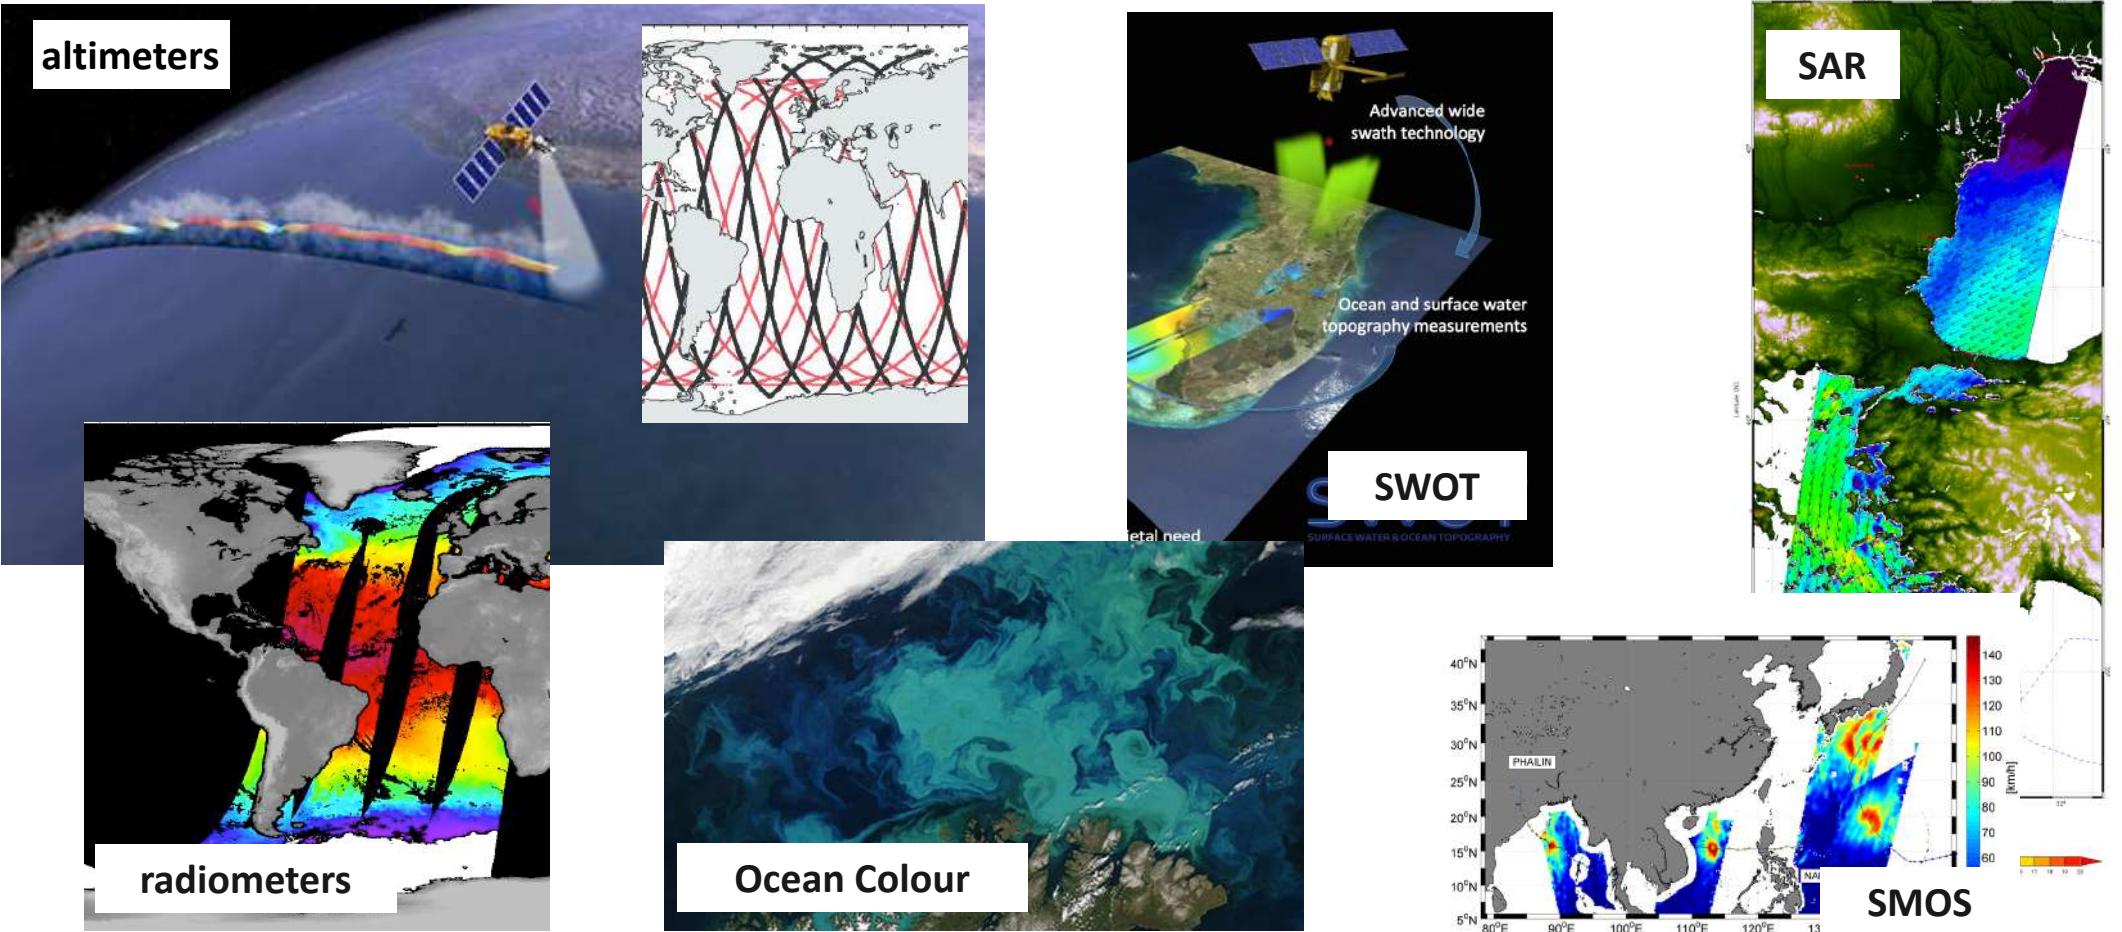
How to solve sampling gaps and extract high-level information for ocean monitoring and surveillance ?



**Context: No observation / simulation system to resolve all scales and processes simultaneously**



# Illustration of satellite-derived sea surface observations

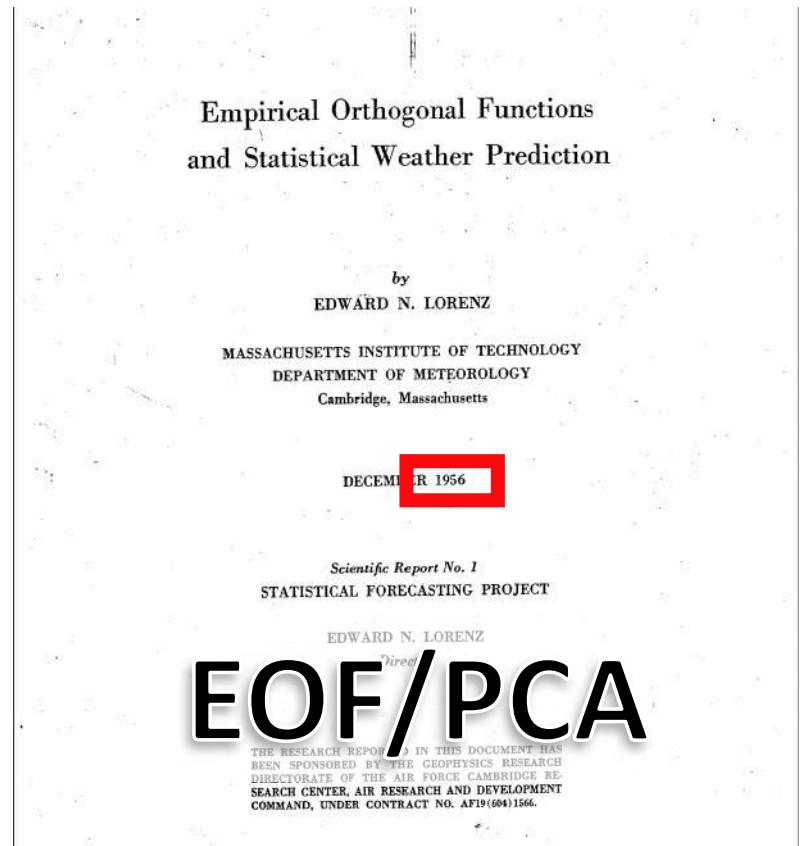


28-03-2008 08:19:12.3 (UTC)

CLS



# Learning & Geoscience: an old story ?



## Deterministic Nonperiodic Flow<sup>1</sup>

EDWARD N. LORENZ

Massachusetts Institute of Technology

(Manuscript received 18 November 1962, in revised form January 1963)

### ABSTRACT

Finite systems of deterministic ordinary nonlinear differential equations may be designed to represent forced dissipative hydrodynamic flow. Solutions of these equations can be identified with trajectories in phase space. For those systems with bounded solutions, it is found that nonperiodic solutions are ordinarily unstable with respect to small modifications, so that slightly differing initial states can evolve into considerably different states. Systems with bounded solutions are shown to possess bounded numerical solutions. A simple system representing atmospheric convection is solved numerically. All of the solutions are found to be unstable, and almost all of them are nonperiodic.

The feasibility of very-long-range weather prediction is examined in the light of these results.

### 1. Introduction

Certain hydrodynamical systems exhibit steady-state flow patterns, while others oscillate in a regular periodic fashion. Still others vary in an irregular, seemingly haphazard manner, and, even when observed for long periods of time, do not appear to repeat their previous history.

These modes of behavior can be illustrated by familiar rotating fluid systems. As pointed out by Lorenz et al. (1959), for example, if a cylinder is heated from below, a cylindrical vortex forms in the rotating fluid about its axis, and is heated near its rim and cooled near its center in a steady symmetrical fashion. Under certain conditions the resulting flow is as symmetric and steady as the heating which gives rise to it. Under different conditions a system of regularly spaced waves develops, and progresses at a uniform speed without changing its shape. Under still different conditions an irregular flow, termed "turbulence," develops, and moves in a haphazard manner.

Lack of periodicity is one of the distinguishing characteristics of turbulent flow.

Because instantaneous turbulent flow patterns are so irregular, attention is often confined to the statistics of turbulence, which, in contrast to the details of turbulence, often behave in a regular well-organized manner. The short-range weather forecaster, however, is forced unwillingly to predict the details of the large-scale turbulent flow, and must do so with confidence which continually fluctuates in accordance with the

Thus there are occasions when more than the statistics of irregular flow are of very real concern.

In this study we shall work with systems of deterministic equations which are idealizations of hydrodynamical systems. We shall be interested principally in nonperiodic solutions, i.e., solutions which never repeat their past history exactly, and where approximate repetitions are of finite duration. These we shall be interested in because they are associated with turbulent flow conditions. We shall also consider systems of finite mass which consist of a collection of molecules, usually a very large finite collection—in which case the governing laws are expressible as a finite set of ordinary differential equations. These equations are generally highly intractable, and the set of molecules is usually approximated by a continuous distribution function. The governing laws are then expressed as a set of partial differential equations, containing such quantities as velocity and pressure as dependent variables.

It is sometimes possible to obtain particular solutions of these equations analytically, especially when the solutions are periodic or invariant with time, and, indeed, much work has been devoted to obtaining such solutions by one scheme or another. Ordinarily, however, nonperiodic solutions cannot readily be determined, either by analytical procedures or by numerical calculations. It is the purpose of this paper to introduce a new method of attack on this problem which may be the basis for further developments.

<sup>1</sup> The research reported in this paper has been sponsored by the Air Force Geophysics Research Directorate of the Air Force Cambridge Research Center, under Contract No. AF 19(60)1566.

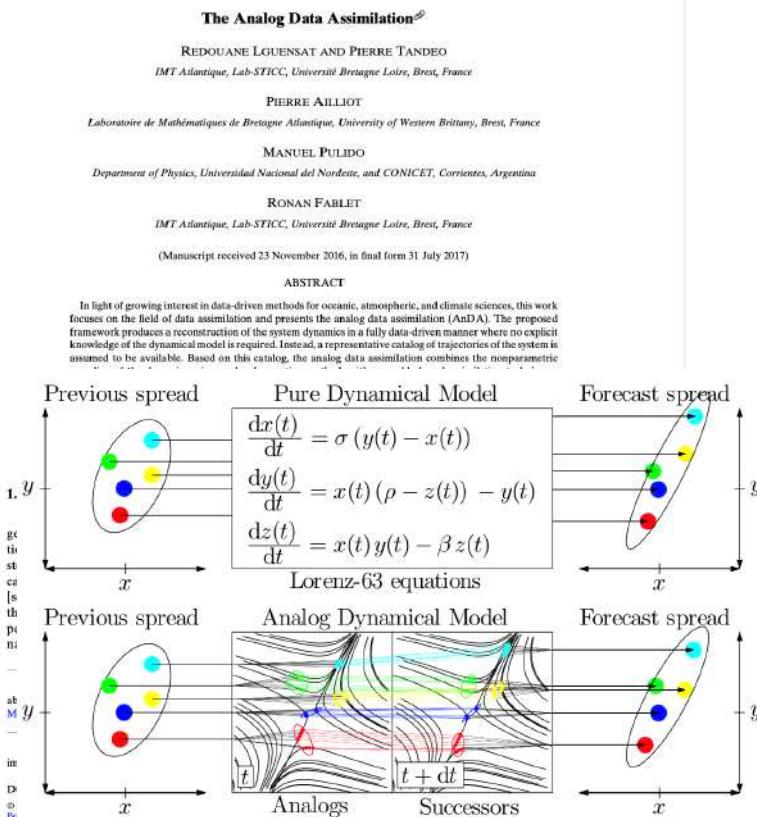
# Analogs/ Nearest- neighbors

# Learning & Geoscience: Data-driven approaches for data assimilation

OCTOBER 2017

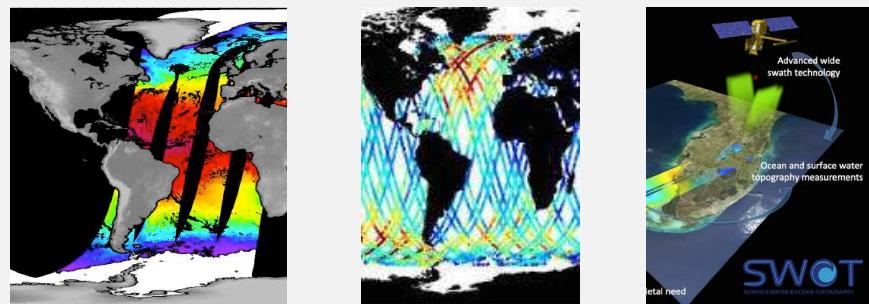
LGUENSAT ET AL.

4093



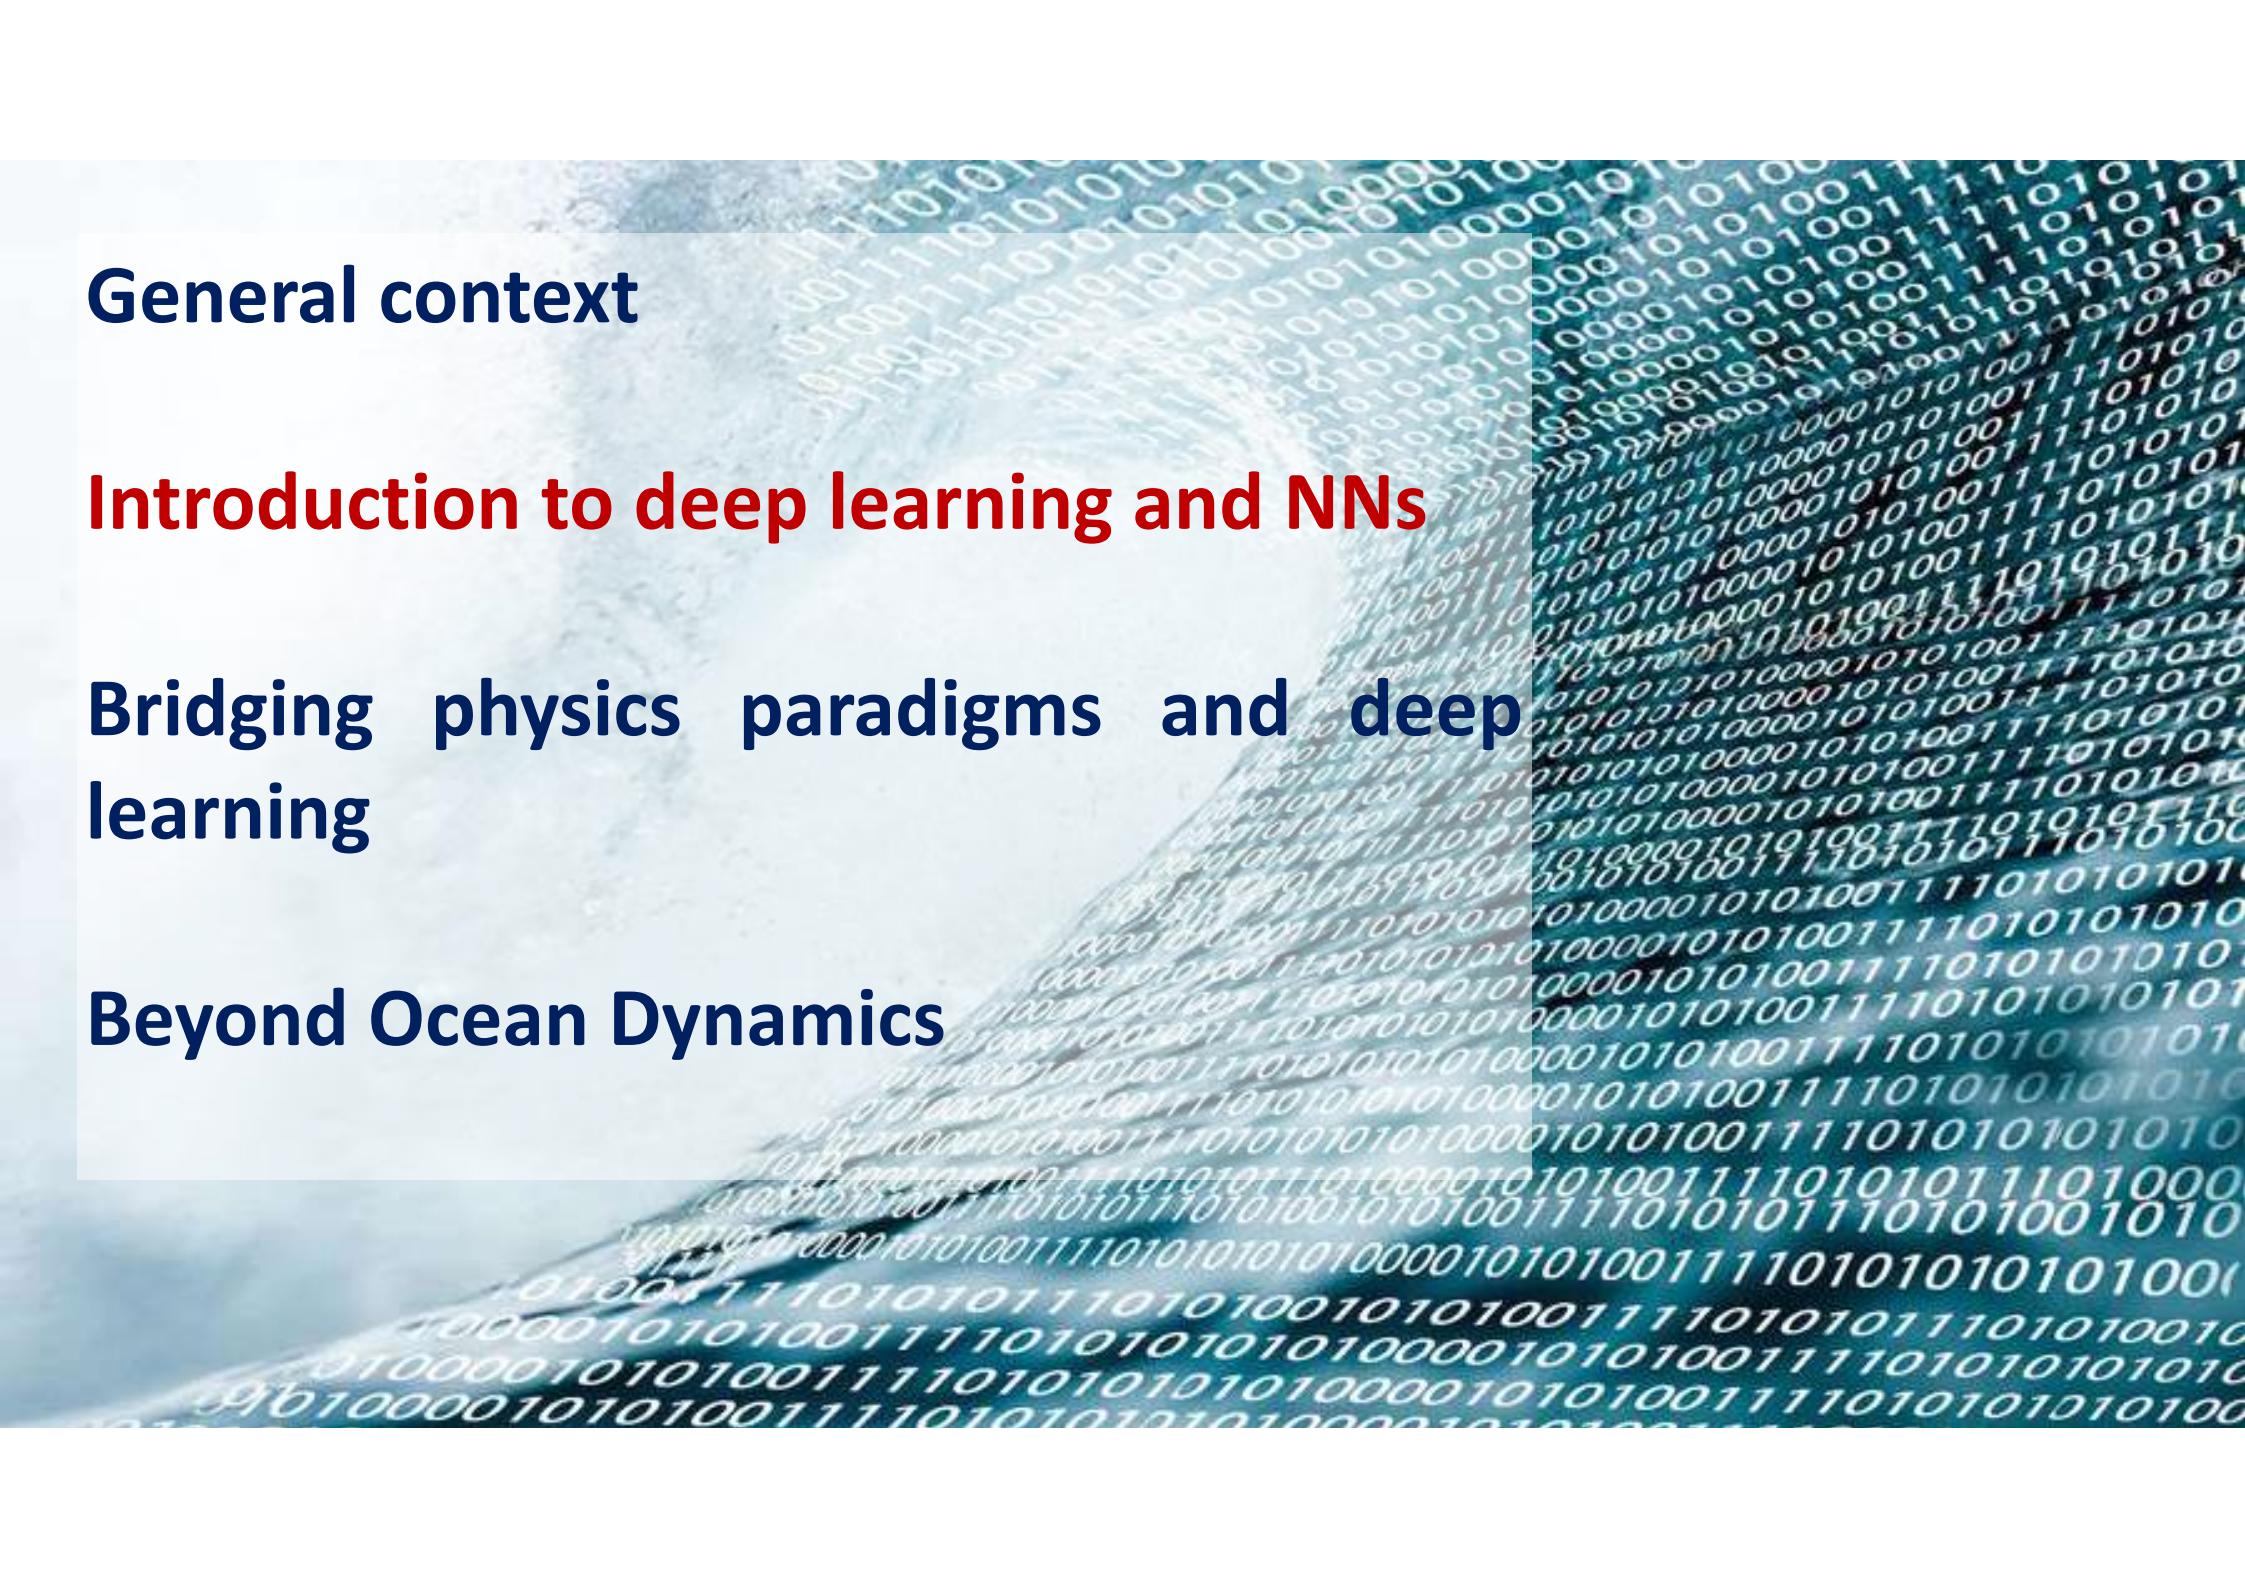
## The analog data assimilation [Lguensat et al., 2017]

- Combination of analog forecasting strategies and EnKF assimilation schemes
- Extension to 2D+ $t$  geophysical dynamics



## Open questions

- Bridging model-driven and data-driven paradigms
- Learning data-driven representations from real observation data



## General context

## Introduction to deep learning and NNs

## Bridging physics paradigms and deep learning

## Beyond Ocean Dynamics

# Introduction to Deep Learning

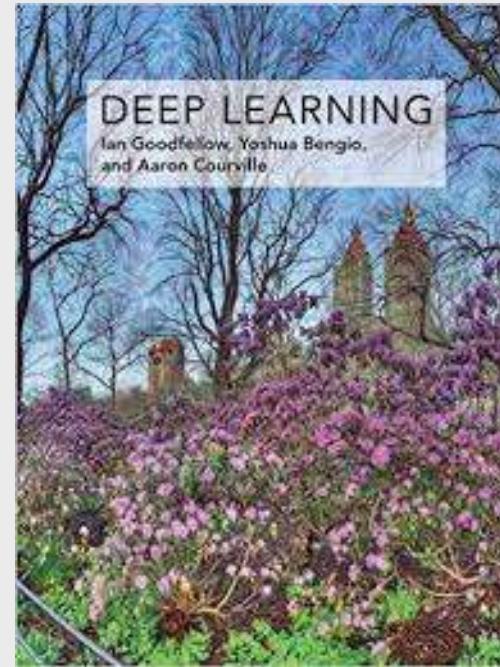


## Resources

- Book: Deep Learning

Goodfellow, Bengio, Courville, MIT Press

Online version <http://www.deeplearningbook.org/>



- Online course by Andrew Ng (Stanford/Baidu)

Youtube: [link](#)

Online course on Coursera: [link](#)

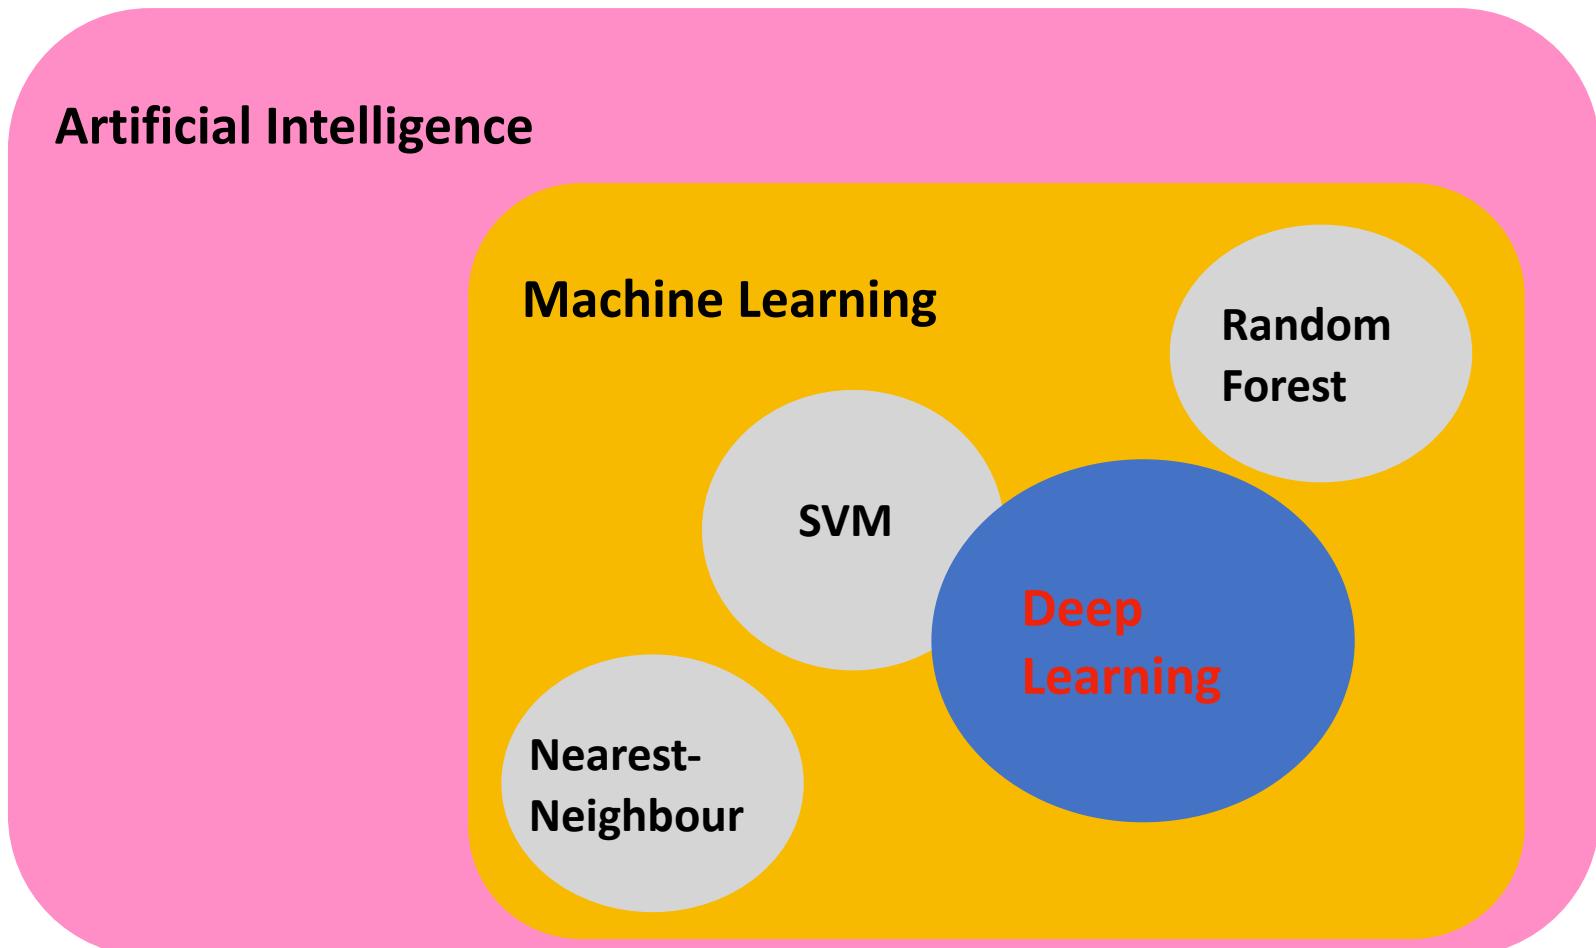
- Review paper: Deep learning in neural networks by

J. Schmidhuber

pdf: [link](#)

The image shows the abstract page of a research paper. At the top, it says 'Neural Networks 61 (2015) 89–127' and 'Contents lists available at ScienceDirect'. Below that is the journal logo 'ELSEVIER' and the journal name 'Neural Networks'. The abstract is titled 'Deep learning in neural networks: An overview' and is authored by 'Jürgen Schmidhuber'. It includes a short summary of the paper's content, which discusses deep artificial neural networks, their applications in pattern recognition and machine learning, and their historical context. The abstract also mentions various types of neural networks like convolutional, recurrent, and generative models, and their applications in fields like image and speech recognition, robotics, and game playing. The page footer indicates it was published in 'Volume 61, Part A, November 2014' and 'Available online 21 October 2014'. It also lists the author's affiliation at the University of Lugano, Switzerland, and the copyright information '© 2014 Published by Elsevier Ltd'.

# What's Deep Learning?



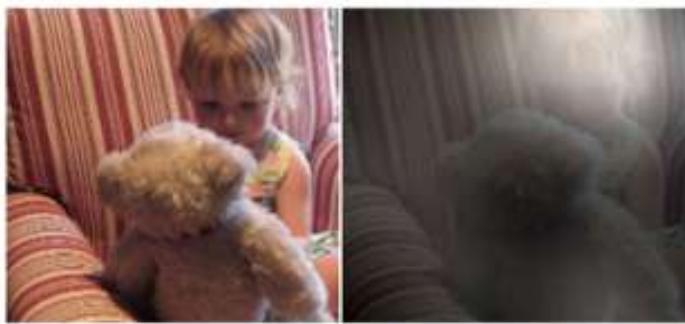
# From Image to Text (image captioning)



A woman is throwing a **frisbee** in a park.



A **dog** is standing on a hardwood floor.



A little **girl** sitting on a bed with a teddy bear.



A group of **people** sitting on a boat in the water.

# Automatic Translation

The image displays two side-by-side automatic translation interfaces. On the left is the Google Translate interface, showing the input "Ce cours sur l'apprentissage profond est génial." and the output "This deep learning course is great." On the right is the DeepL interface, showing the input "Ce cours sur l'apprentissage profond est génial." and the output "This course on deep learning is great." A red stamp with the text "Early 2017" is overlaid on the Google Translate interface.

Google

Traduction

Anglais Français Arabe Détecter la langue

French English Arabic Detect language

Ce cours sur l'apprentissage profond est génial. This deep learning course is great.

48/5000

Désactiver la traduction instantanée

Early 2017

Google

Traduction

Anglais Français Arabe Détecter la langue

French English Arabic Detect language

Ce cours sur l'apprentissage profond est génial. This deep learning course is great.

48/5000

Désactiver la traduction instantanée

Suggérer une modification

DeepL

Translate from FRENCH (detected) ▾

Translate into ENGLISH ▾

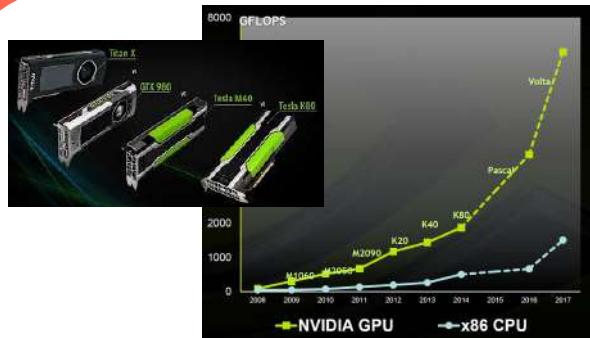
Ce cours sur l'apprentissage profond est génial.

This course on deep learning is great.

To look up words in the dictionary, just click on them.

and many more.....

# Key reasons for the emergence of DL



High-performance computing (GPU)



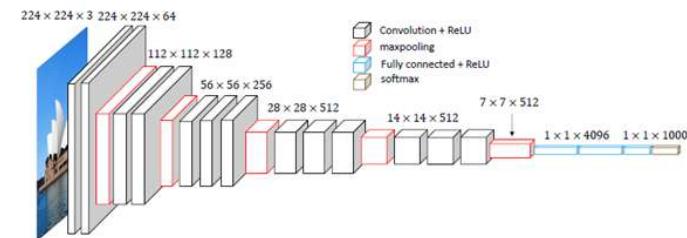
Large annotated dataset (> 1M)



TensorFlow

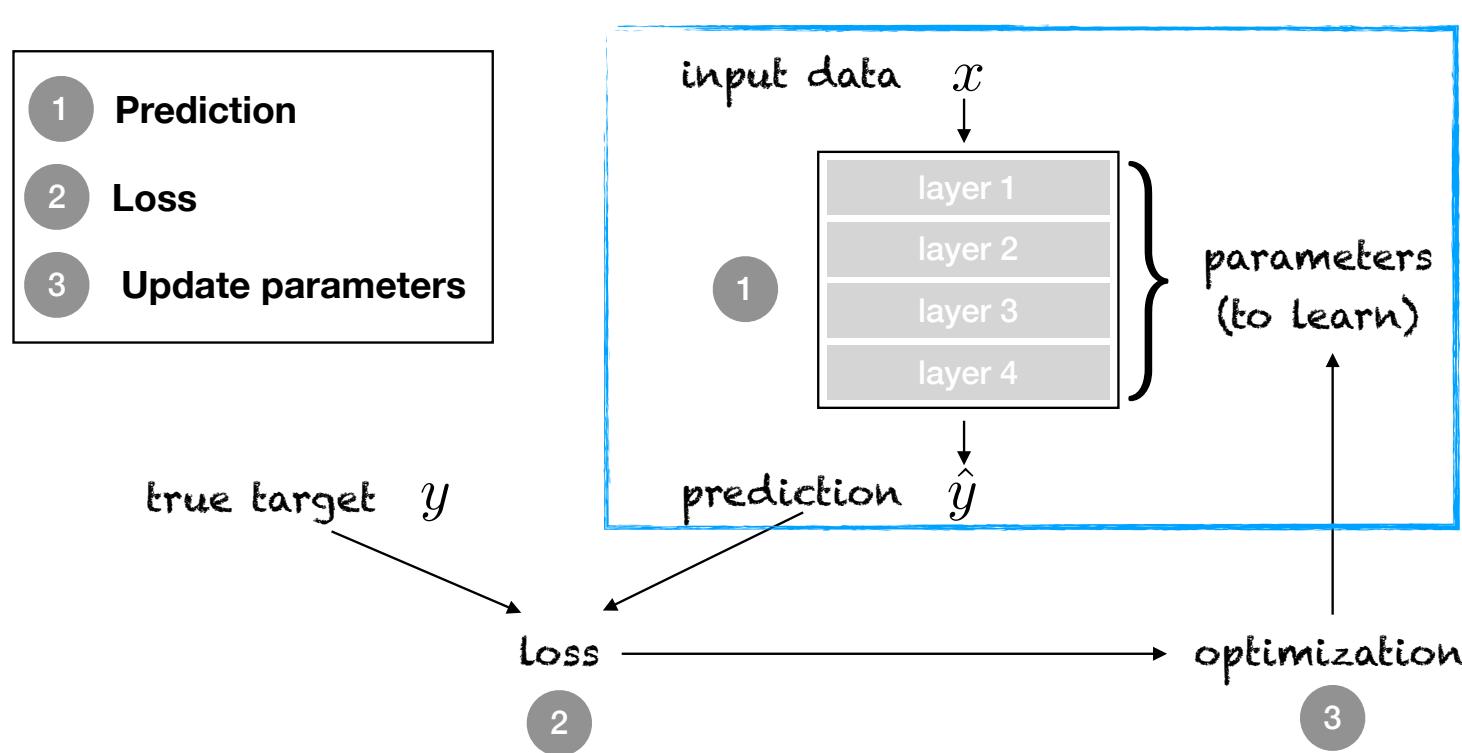


Efficient & easy-to-use frameworks



End-to-end learning

# What's learning (for a computer) ?



# What's learning (for a computer) ?

**Mathematical formulation:** Learning comes to minimising some loss function given w.r.t. model parameters and training data

$$\widehat{\theta} = \arg \min_{\theta} \mathcal{L} (\{x_i, y_i\}_{i \in \{1, \dots, N\}}; f_{\theta})$$

**Key questions:**

- Which parameterisation for model  $f$  ?
- Which loss function ?

# Deep learning models

# Neural networks: composition idea

- Approximation through the composition of (simple) elementary functions:

$$f_{\theta}(x) = f_{\theta_N} \circ \dots \circ f_{\theta_2} \circ f_{\theta_1}(x)$$

- Key features:
  - Any continuous function can be approximated as the composition of elementary functions
  - Analytical/exact computation of the derivative of  $f$  with respect to parameters and input variables
  - Direct exploitation of gradient-based optimisation schemes for learning

# Neural networks: automatic differentiation

- Given the general composition idea:

$$f_{\theta}(x) = f_{\theta_N} \circ \dots \circ f_{\theta_2} \circ f_{\theta_1}(x)$$

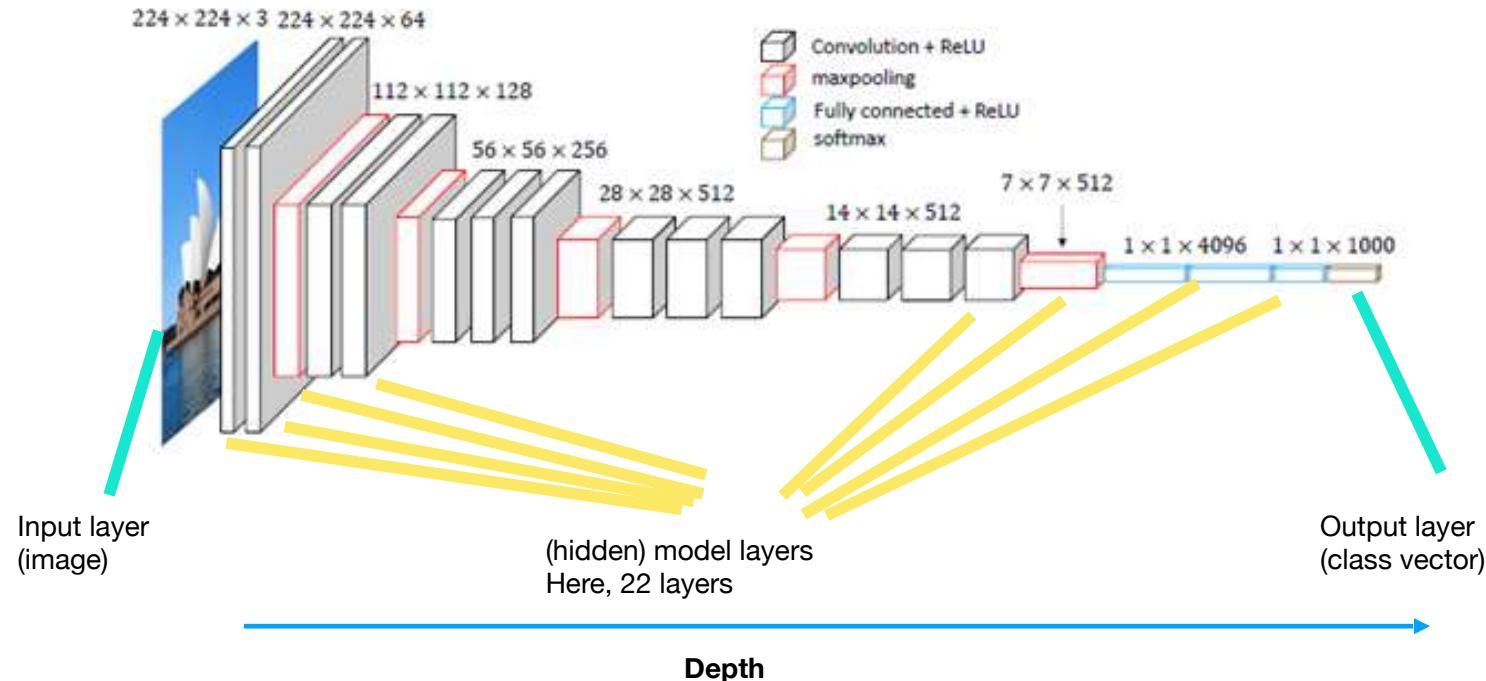
- NNs can implement automatic differentiation knowing the symbolic differentiation for elementary functions:

- AD example: [https://pytorch.org/tutorials/beginner/blitz/autograd\\_tutorial.html](https://pytorch.org/tutorials/beginner/blitz/autograd_tutorial.html)
- Basis of the backpropagation algorithm (similar to adjoint method)
- Resulting gradient descent for learning model parameters:

$$\widehat{\theta} = \arg \min_{\theta} \mathcal{L}(\{x_i, y_i\}_i; f_{\theta}) \quad \Longrightarrow \quad \theta^{(k+1)} = \theta^{(k+1)} + \lambda_k \nabla_{\theta} \mathcal{L}(\{x_i, y_i\}_i; f_{\theta})$$

# Deep learning models

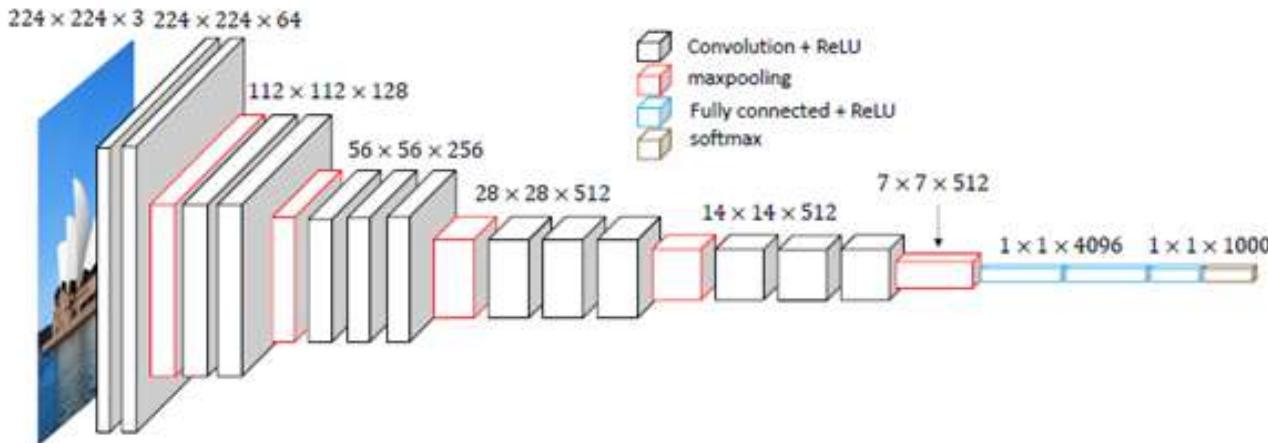
DL models are (in general) feedforward models. VGG16 as an illustration



The more layers, the deeper..... Some models may have up to several hundreds to thousands of layers.

# Basics of DL models

DL models are (in general) feedforward models. VGG16 as an illustration



Elementary components

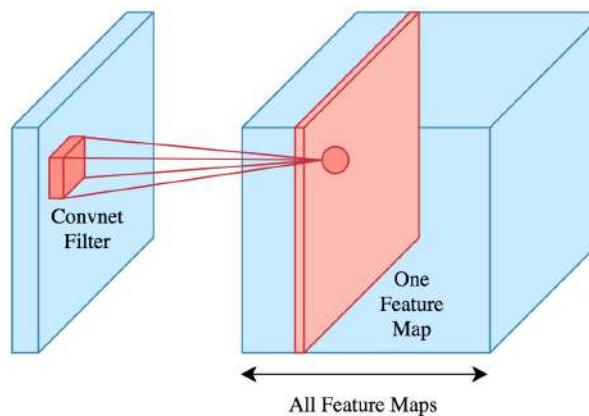
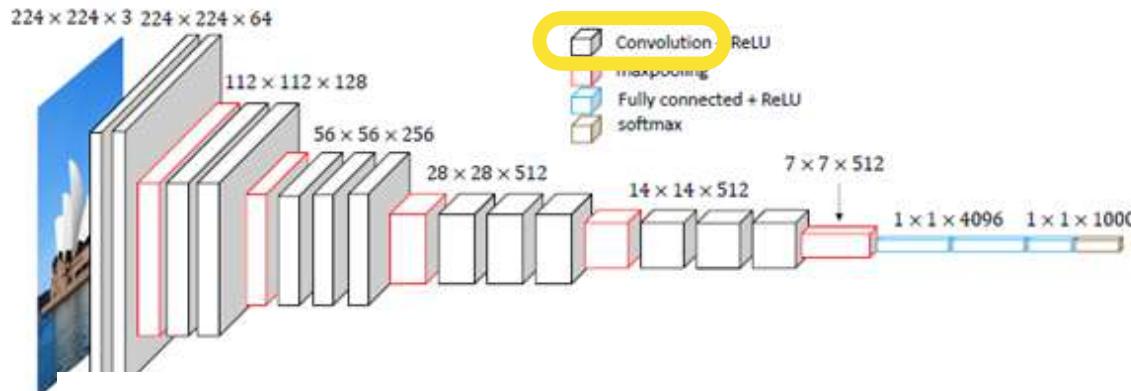
Convolution layers

Activation layers

Pooling layers

Dense layers

# Basics of DL models

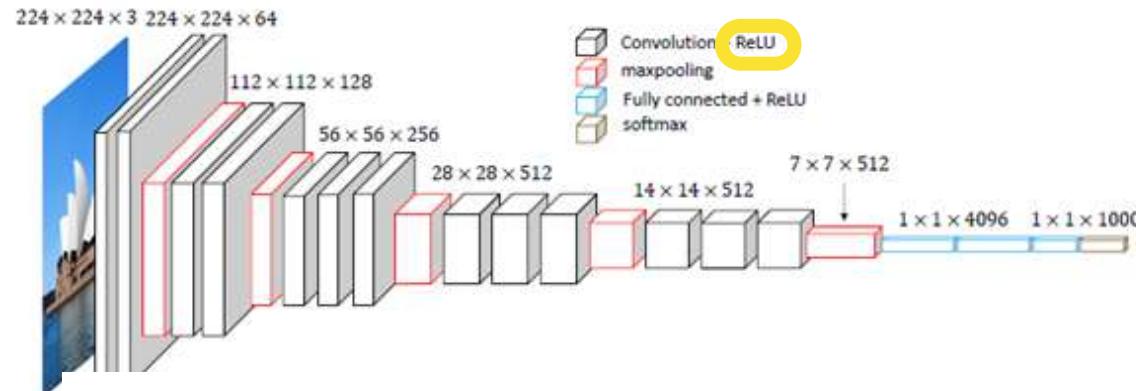


Number of parameters:  
Filter size x Number of filters

e.g.  $3 \times 3 \times K \times N_{\text{filt}}$

**Independent on the sizes of the input  
and output layer**

# Basics of DL models



Elementary components

Convolution layers

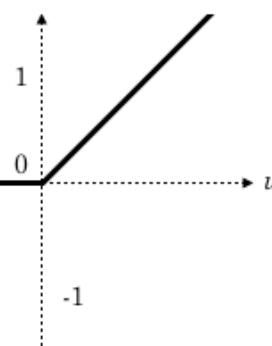
Activation layers

Pooling layers

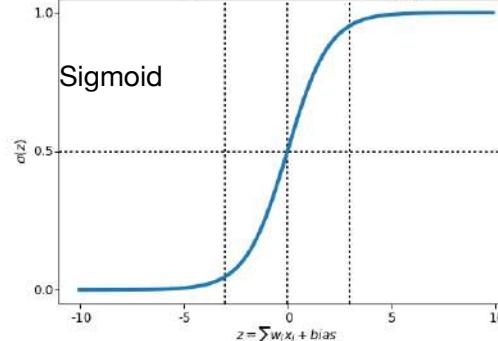
Dense layers

ReLU  
(Rectified Linear Unit)

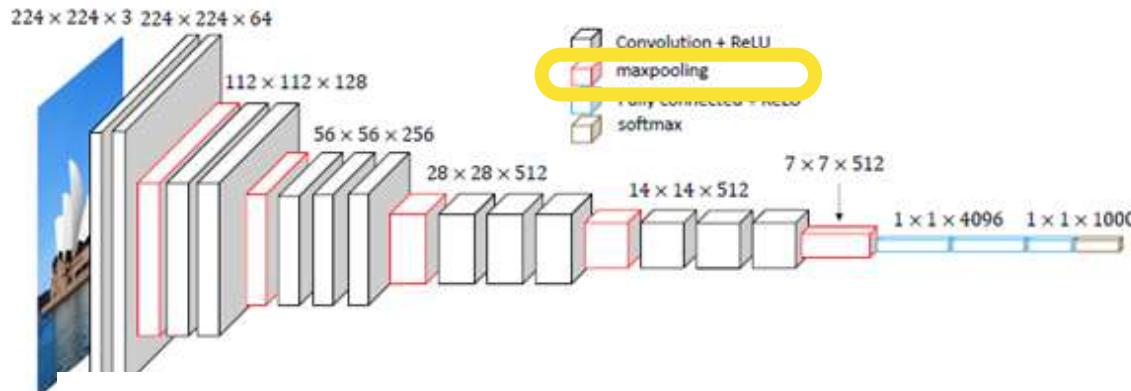
$$f(u) = \max(0, u)$$



Sigmoid Function  $\sigma(z) = \frac{1}{1+e^{-z}}$



# Basics of DL models



Elementary components

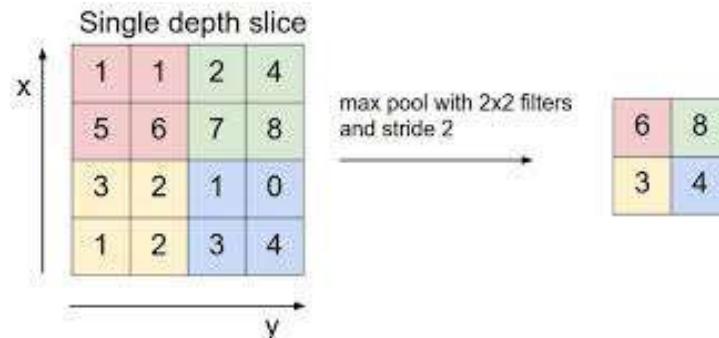
Convolution layers

Activation layers

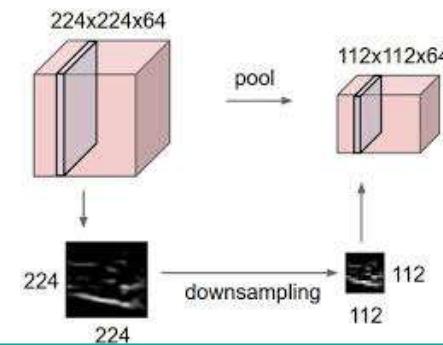
Pooling layers

Dense layers

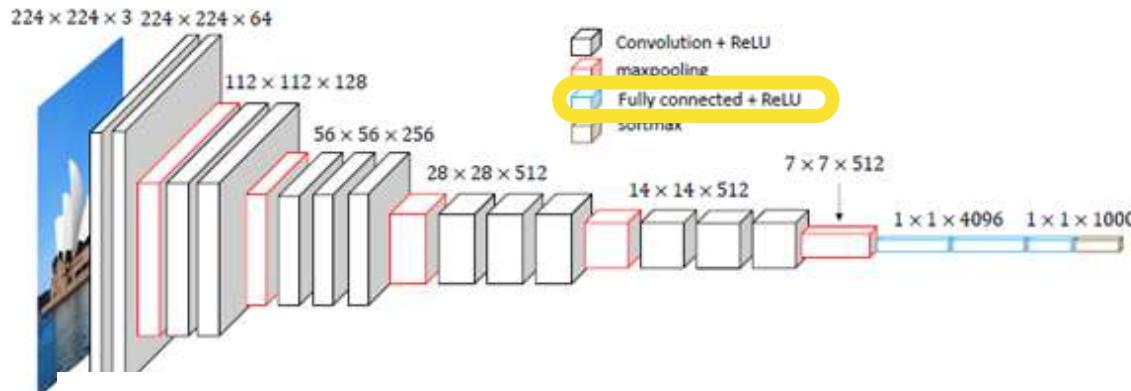
An example of max-pooling operator



Pooling downsamples the input layer



# Basics of DL models



Elementary components

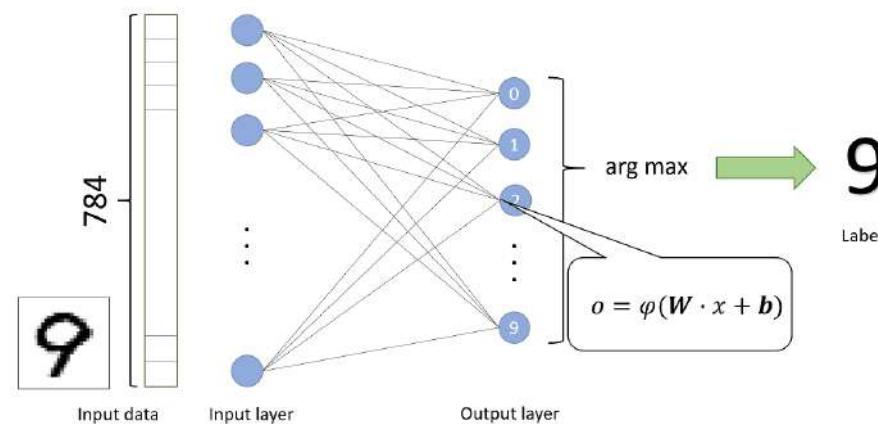
Convolution layers

Activation layers

Pooling layers

Dense layers

Dense layers  
or  
Fully-connected (FC) layer  
as in a classic MLP



**Let's play with MLPs using  
tensorflow playground**

<http://playground.tensorflow.org/>

## General context

## Introduction to deep learning and NNs

## Bridging physics paradigms and deep learning

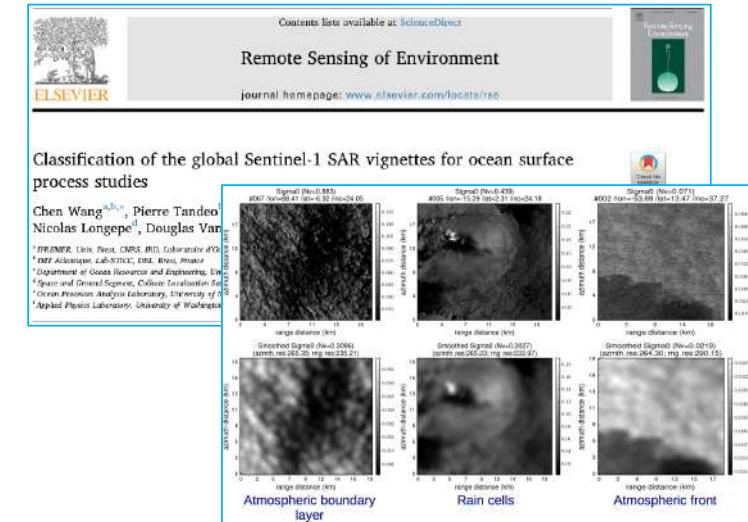
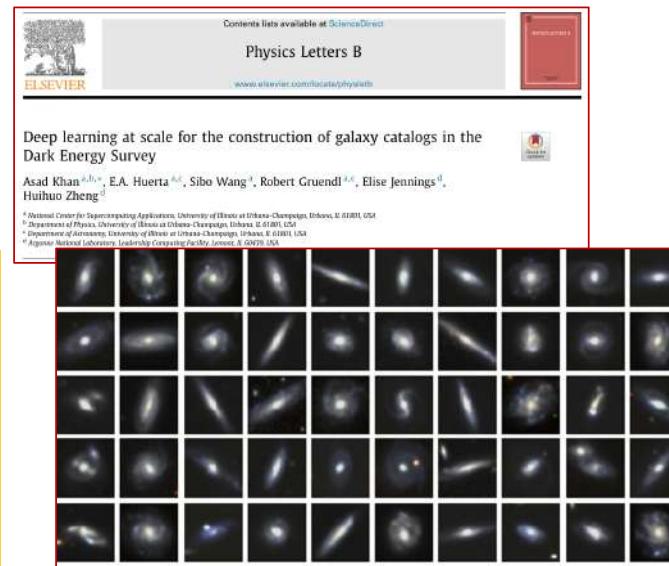
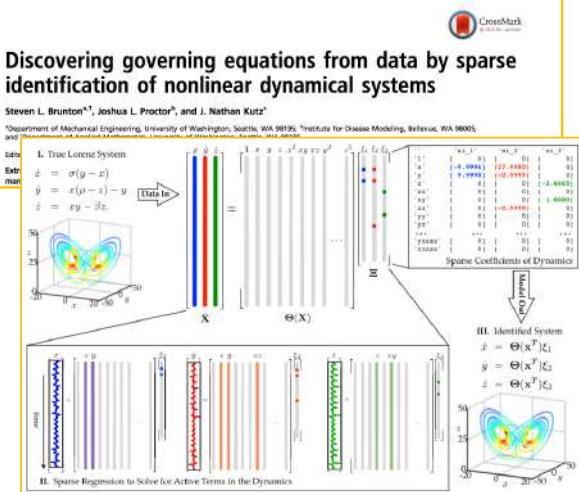
## Beyond Ocean Dynamics

# Deep Learning applied to physics

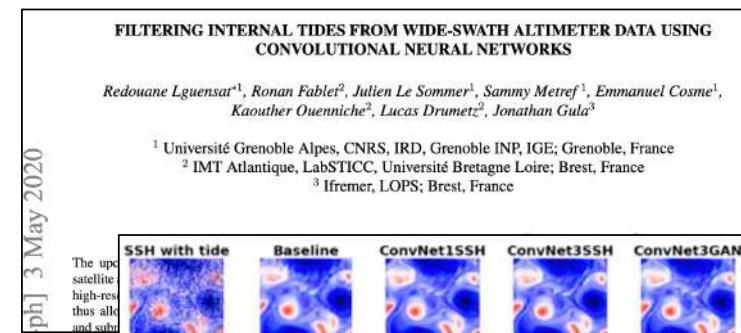
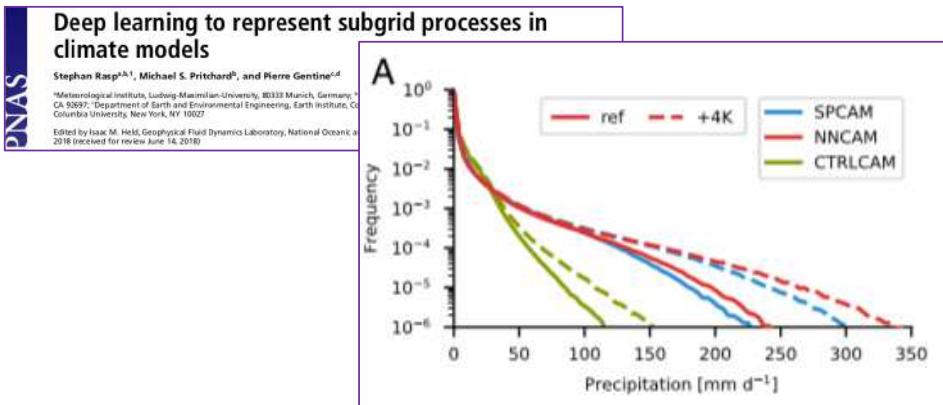


# Direct applications of DL schemes to physics-related issues

PNAS



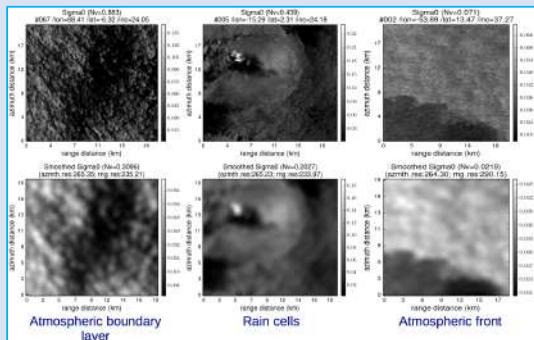
PNAS



# How to (directly) apply DL to physics-related issues ?

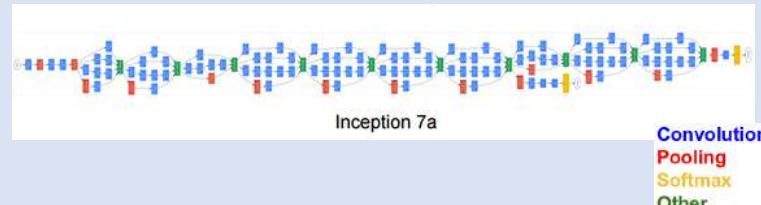
1

Build a groundtruthed dataset



2

Design or choose an architecture



Many architectures available online  
(VGG, ResNet, U-Net,...)

3

Choose a DL framework



K Keras

PYTORCH

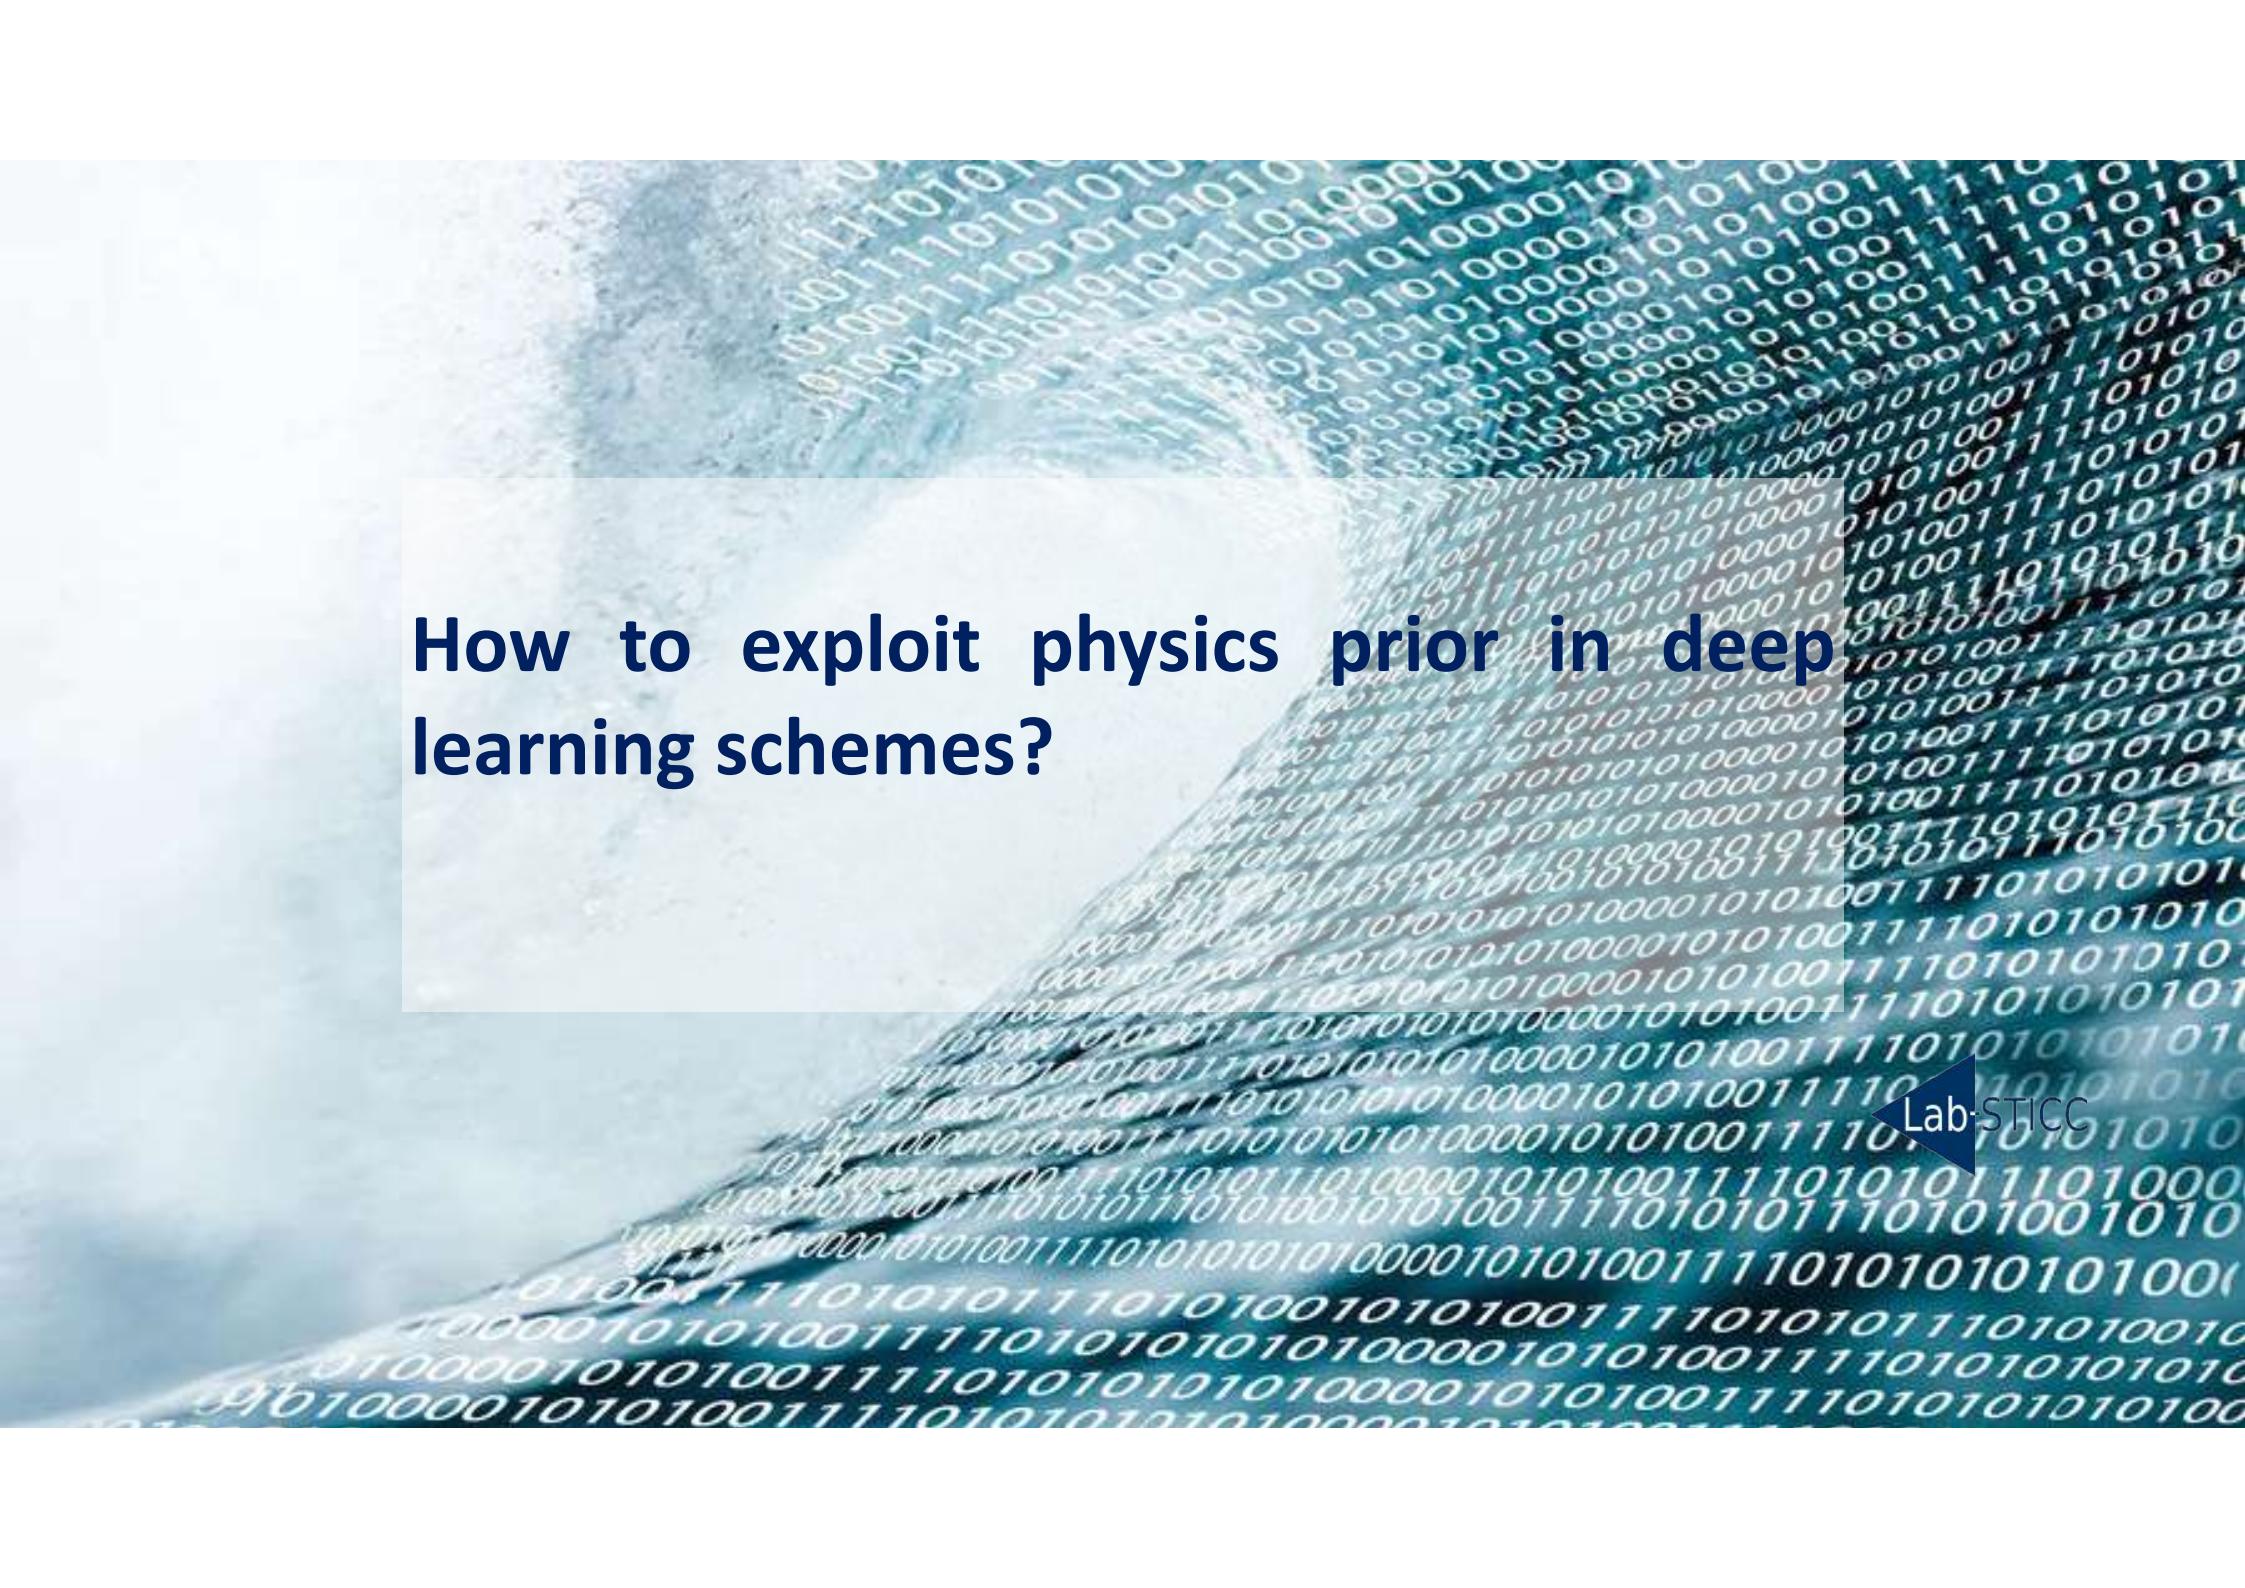
4

Learn the model

Select a training loss

Optimize model parameters using optimizers embedded in DL frameworks

A toy example using Tensorflow-Keras: <https://www.tensorflow.org/tutorials/keras/classification>



# How to exploit physics prior in deep learning schemes?

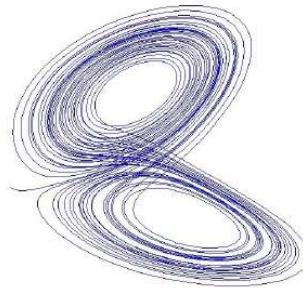


# How to embed physics-driven priors in DL models ?

An illustration through neural ODE for Lorenz-63 dynamics (Fablet et al., 2018)

$$\begin{aligned}\frac{dx(t)}{dt} &= \sigma (y(t) - x(t)) \\ \frac{dy(t)}{dt} &= x(t) (\rho - z(t)) - y(t) \\ \frac{dz(t)}{dt} &= x(t) y(t) - \beta z(t)\end{aligned}$$

Lorenz-63 equations



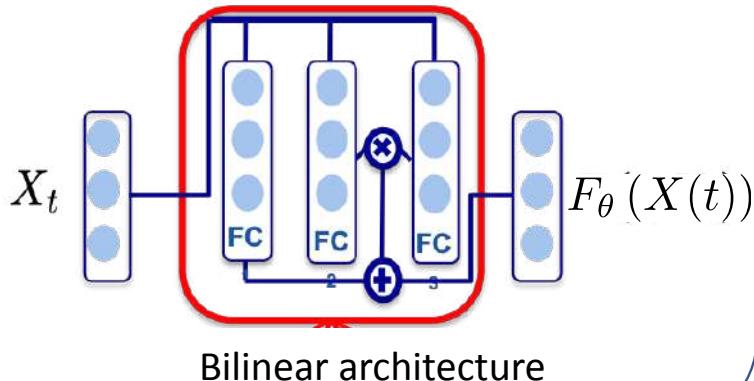
*Associated Euler integration scheme*

$$d_t X(t) = F_\theta (X(t))$$

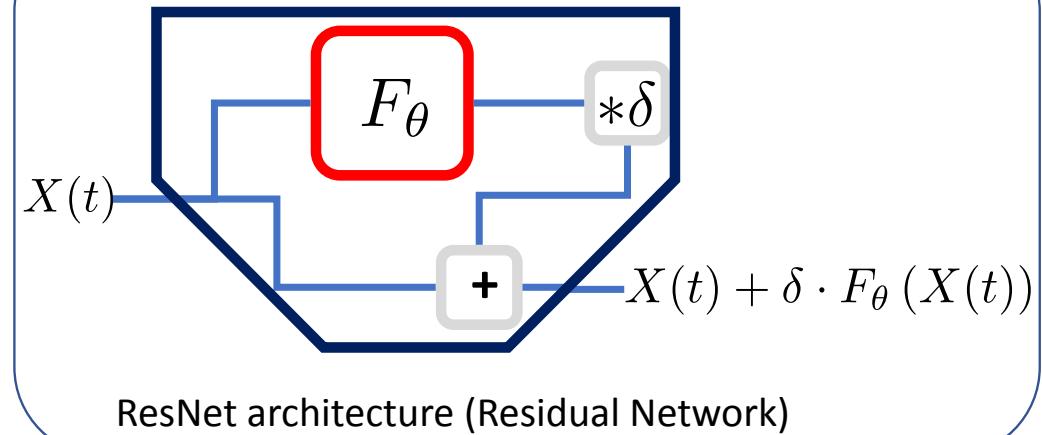


$$X(t + \delta) = X(t) + \delta \cdot F_\theta (X(t))$$

*NN architecture for differential operator*



*NN architecture for integration scheme*



# How to embed physics-driven priors in DL models ?

An illustration through neural ODE for Lorenz-63 dynamics (Fablet et al., 2018)

$$\frac{dx(t)}{dt} = \sigma (y(t) - x(t))$$

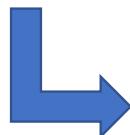
$$\frac{dy(t)}{dt} = x(t) (\rho - z(t)) - y(t)$$

$$\frac{dz(t)}{dt} = x(t) y(t) - \beta z(t)$$

Lorenz-63 equations

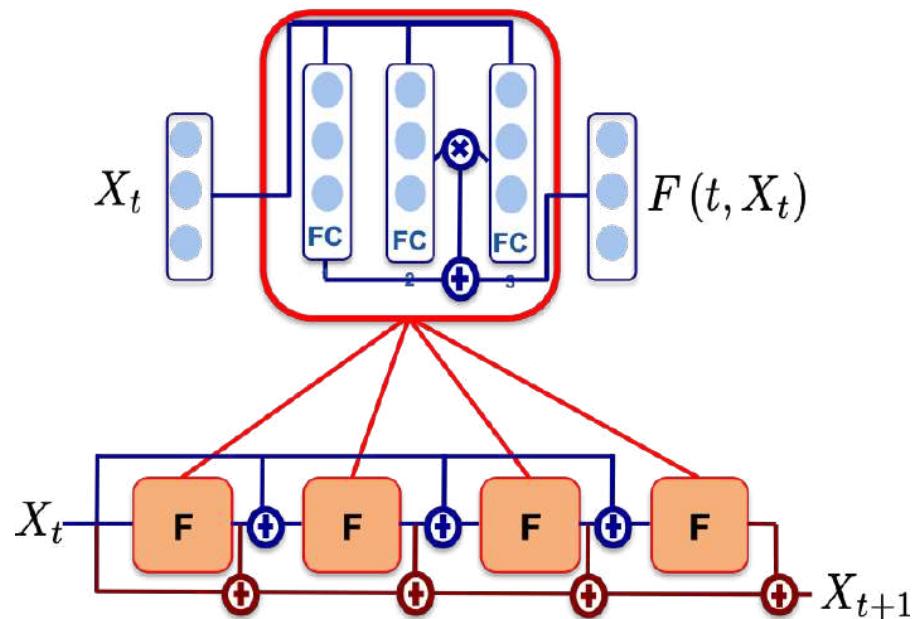
**Generalization to higher-order integration schemes (eg, RK4)**

$$d_t X(t) = F_\theta (X(t))$$



$$X(t + \delta) = X(t) + \sum_i \beta_i k_i$$

$$\text{with } k_i = F_\theta (X(t) + \delta \alpha_i k_{i-1})$$



**NB: Same number of trainable model parameters as the Euler-based architecture**

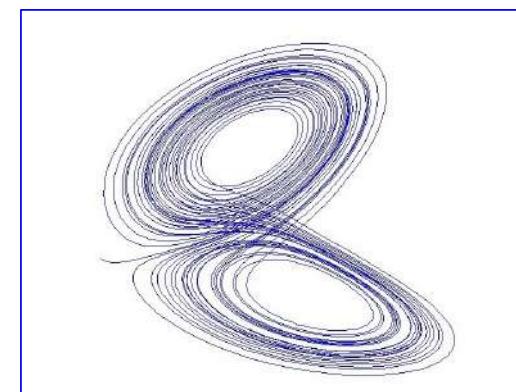
# How to embed physics-driven priors in DL models ?

An illustration through L63 dynamics: numerical experiments (Fablet et al., 2018)

Forecasting experiments			
Noise-free training data			
Forecasting time step	$t_0+h$	$t_0+4h$	$t_0+8h$
Analog forecasting	$<10^{-6}$	0.002	0.005
Sparse regression	$<10^{-6}$	0.002	0.006
MLP	$<10^{-6}$	0.018	0.044
<b><i>Bi-NN(4)</i></b>	<b><math>&lt;10^{-6}</math></b>	<b><math>&lt;10^{-6}</math></b>	<b><math>&lt;10^{-6}</math></b>

Noisy training data ( $\sigma=0.5$ )					
Forecasting time step	$t_0+h$	$t_0+4h$	$t_0+8h$	in	
Analog forecasting	$<10^{-6}$	2.01	2.2	0	0.25
<b><i>Bi-NN(4)</i></b>	<b><math>&lt;10^{-6}</math></b>	<b>0.054</b>	<b>0.14</b>	<b>1</b>	



Assimilation experiment (1 obs. every 8 time steps)					
Noise	standard deviation	in	0	0.25	1
training data					
<i>True model</i>	<u>0.50</u>	-	-		
Analog forecasting		0.65	1.17	1.81	
<b><i>Bi-NN(4)</i></b>	<b>0.60</b>	<b>0.75</b>	<b>0.86</b>		

# Bridging physics & AI: a broader picture

Physical model

$$\frac{\partial u}{\partial t} + \langle \nabla u, v \rangle = \kappa \Delta u$$

Representation learning

Trainable representation



## Making the most of AI and Physics Theory

- Model-Driven/Theory-Guided & Data-Constrained schemes
- Data-Driven & Physics-Aware schemes (eg, Ouala et al., 2019)

# Bridging physics & AI: a broader picture

Model-driven

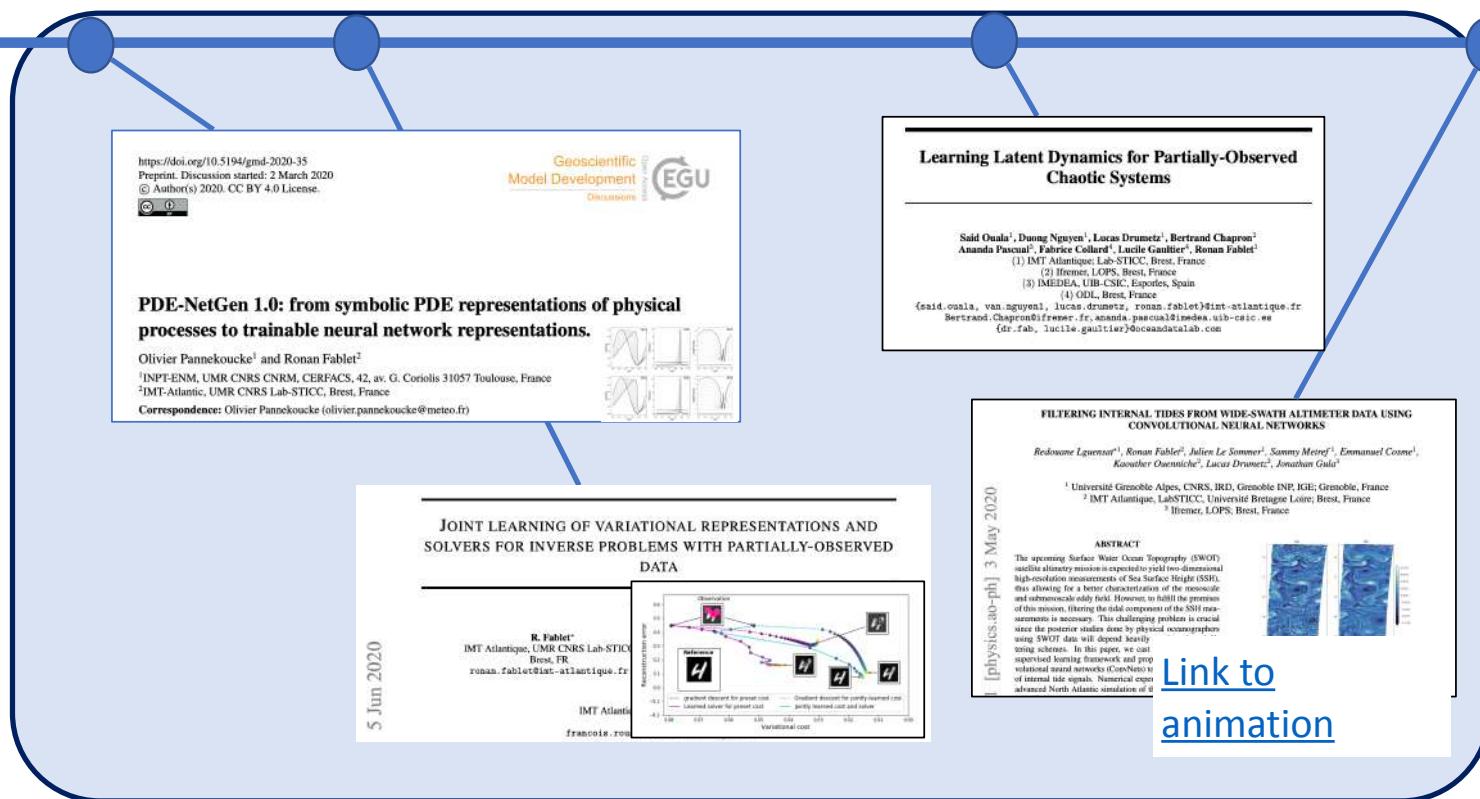
ODEs/PDEs  
Variational models  
State-space models  
Data Assimilation

...

Data-driven

Analog methods  
Kernel methods  
Neural Networks

....



Physics-informed &  
Data-constrained

Data-driven &  
Physics-aware

# Bridging physics & AI: Expected breakthroughs

Physical model

$$\frac{\partial u}{\partial t} + \langle \nabla u, v \rangle = \kappa \Delta u$$

Representation learning

Data-driven representation



## Making the most of AI and Physics Theory

- **Model-Driven/Theory-Guided & Data-Constrained schemes**
- Data-Driven & Physically-Sound schemes (eg, Ouala et al., 2019)

# NN Generator from Symbolic PDEs (Pannekoucke et al., 2020)

$$\partial_t u + u \partial_x u = \kappa \partial_x^2 u$$



**Symbolic calculus  
(Simplicy)**



**PDE-GenNet  
(keras)**

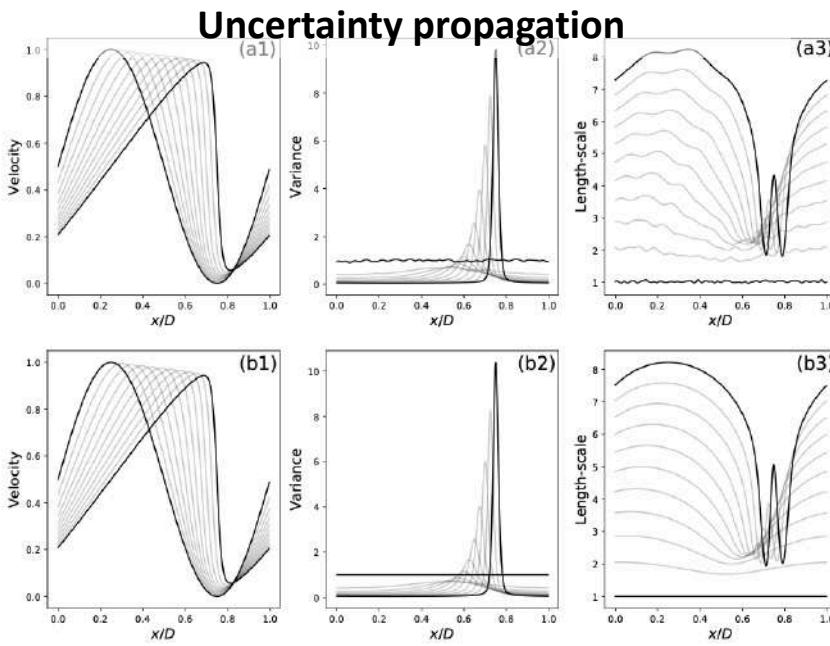


$$\left( \begin{array}{c} \mu_u(t) \\ \Sigma_u(t) \end{array} \right) \rightarrow \text{ResNet} \rightarrow \left( \begin{array}{c} \mu_u(t+1) \\ \Sigma_u(t+1) \end{array} \right)$$

```
# Example of computation of a derivative
kernel_Du_x_01 = np.asarray([[0.0, 0.0, 0.0, 0.0, 0.0, 0.0], [0.0, 1/(2*self.dx[self.coordinates.index('x')]), 0.0]]).reshape((3, 3)+(1,1))
Du_x_01 = DerivativeFactory((3, 3), kernel=kernel_Du_x_01, name='Du_x_01')(u)

# Computation of trend_u
mul_1 = keras.layers.multiply([Dkappa_11_x_01, Du_x_01], name='MulLayer_1')
mul_2 = keras.layers.multiply([Dkappa_12_x_01, Du_y_01], name='MulLayer_2')
mul_3 = keras.layers.multiply([Dkappa_12_y_01, Du_x_01], name='MulLayer_3')
mul_4 = keras.layers.multiply([Dkappa_22_y_01, Du_y_01], name='MulLayer_4')
mul_5 = keras.layers.multiply([Du_x_02, kappa_11], name='MulLayer_5')
mul_6 = keras.layers.multiply([Du_y_02, kappa_22], name='MulLayer_6')
mul_7 = keras.layers.multiply([Du_x_01_y_01, kappa_12], name='MulLayer_7')
sc_mul_1 = keras.layers.Lambda(lambda x: 2.0*x, name='ScalarMullLayer_1')(mul_7)
trend_u = K
```

**Generated code**



**Ensemble-based prediction**

**NN prediction**

# Bridging physics & AI: Expected breakthroughs

Physical model

$$\frac{\partial u}{\partial t} + \langle \nabla u, v \rangle = \kappa \Delta u$$

Representation learning

Data-driven representation



## Making the most of AI and Physics Theory

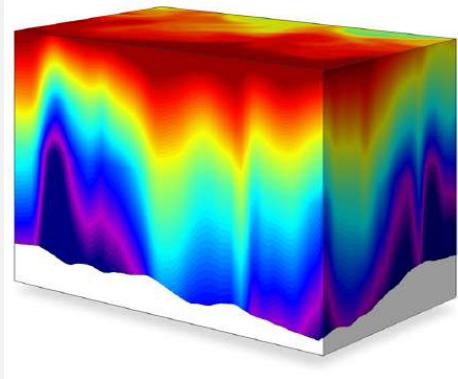
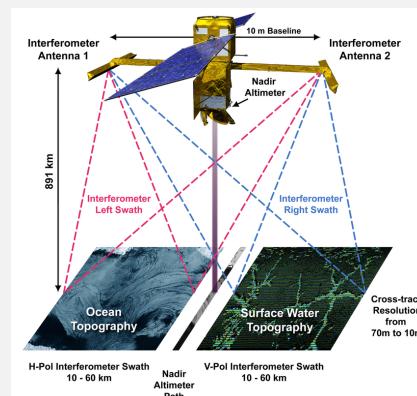
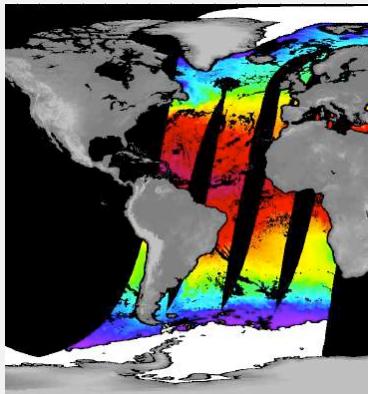
- Model-Driven/Theory-Guided & Data-Constrained schemes
- Data-Driven & Physically-Sound schemes (eg, Ouala et al., 2019)

# Deep learning for irregularly-sampled and partially-observed systems



# learning for partially-observed systems / from irregularly-sampled data [Ouala et al., 2020; Fablet et al., 2020]

Can we learn directly from observation data ?



*Generic issue:  
Joint identification and inversion*

*Dynamical model*

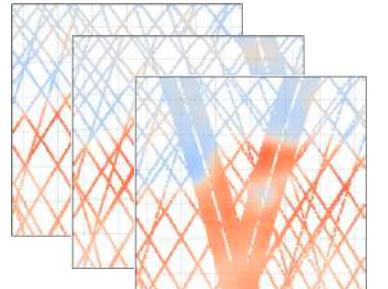
$$X_t \rightarrow \partial_t X = F(X, \xi, t, \theta) \rightarrow X_{t+1}$$



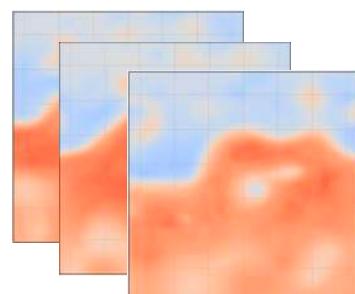
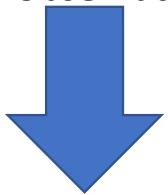
*Observation model*

$$Y_t = H(X, \zeta, t, \phi)$$

# (Weak constraint) 4DVar Data Assimilation (DA) formulation



Partial observations  $y$



True states  $x$

**State-space formulation:**

$$\begin{cases} \frac{\partial x(t)}{\partial t} = \mathcal{M}(x(t)) \\ y(t) = x(t) + \epsilon(t), \forall t \in \{t_0, t_0 + \Delta t, \dots, t_0 + N\Delta t\} \end{cases}$$

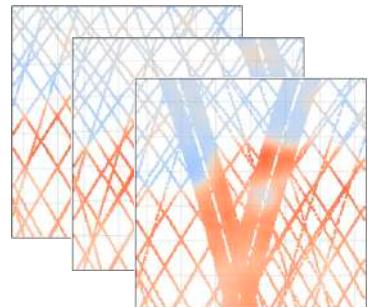
**Associated variational formulation:**

$$\arg \min_x \lambda_1 \sum_i \|x(t_i) - y(t_i)\|_{\Omega_{t_i}}^2 + \lambda_2 \sum_n \|x(t_i) - \Phi(x)(t_i)\|^2$$

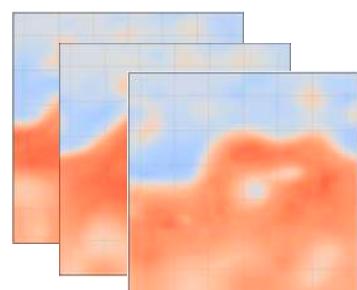
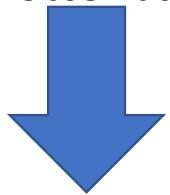
with  $\Phi(x)(t) = x(t - \Delta) + \int_{t-\Delta}^t \mathcal{M}(x(u)) du$

$$\boxed{\arg \min_x \lambda_1 \|x - y\|_{\Omega}^2 + \lambda_2 \|x - \Phi(x)\|^2}$$

# 4DVar Data Assimilation (DA) formulation



Partial observations  $y$



True states  $x$

**Model-driven schemes:** 
$$\arg \min_x \underbrace{\lambda_1 \|x - y\|_{\Omega}^2 + \lambda_2 \|x - \Phi(x)\|^2}_{U_{\Phi}(x, y, \Omega)}$$

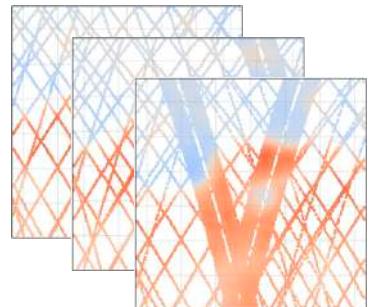
**Gradient-based solver (adjoint/Euler-Lagrange method):** 
$$U_{\Phi}(x, y, \Omega)$$

$$x^{(k+1)} = x^{(k)} - \alpha \nabla_x U_{\Phi} \left( x^{(k)}, y, \Omega \right)$$

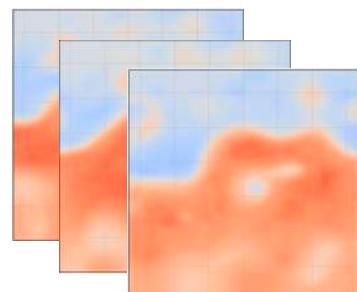
## Implementation of 4DVar Data Assimilation using Deep Learning frameworks

- Implement operator  $\Phi$  as a neural network (e.g., neural ODE architecture w.r.t. a known dynamical operator)
- Use embedded automatic differentiation tool to implement the above gradient descent without deriving analytically the adjoint operator
- Link to an example for L63 dynamics: [https://colab.research.google.com/drive/1SyU6PzOAzgBz4\\_U\\_5WbPMriwT4bLXKvh?usp=sharing](https://colab.research.google.com/drive/1SyU6PzOAzgBz4_U_5WbPMriwT4bLXKvh?usp=sharing)

# End-to-end learning for inverse problems (Fablet et al., 2020)



Partial observations  $y$



True states  $x$

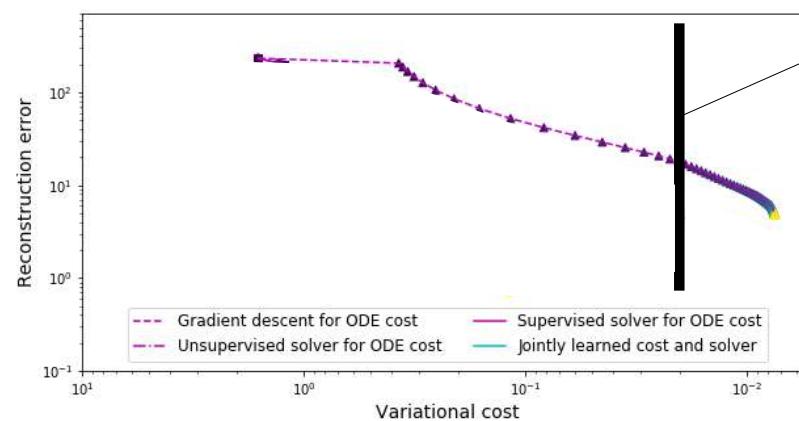
**Model-driven schemes:**  $\hat{x} = \arg \min_x \underbrace{\lambda_1 \|x - y\|_{\Omega}^2 + \lambda_2 \|x - \Phi(x)\|^2}_{U_{\Phi}(x^{(k)}, y, \Omega)}$

**Gradient-based solver (adjoint/Euler-Lagrange method):**  $x^{(k+1)} = x^{(k)} - \alpha \nabla_x U_{\Phi}(x^{(k)}, y, \Omega)$

$$x^{(k+1)} = x^{(k)} - \alpha \nabla_x U_{\Phi}(x^{(k)}, y, \Omega)$$

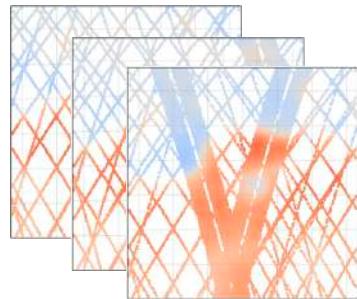
No control on the reconstruction error

$$x^{true} \neq \arg \min_x U_{\Phi}(x^{(k)}, y, \Omega)$$

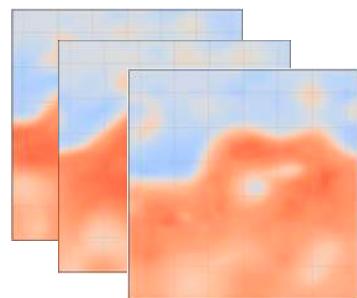


Variational cost  
for the true state

# End-to-end learning for inverse problems (Fablet et al., 2020)



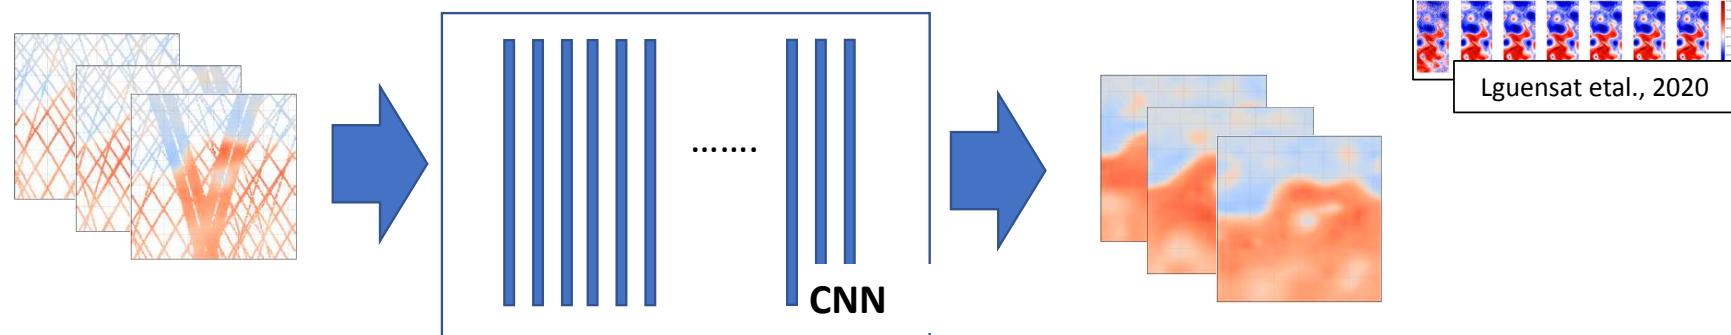
Partial observations  $y$



True states  $x$

**Model-driven schemes:**  $\hat{x} = \arg \min_x \lambda_1 \|x - y\|_{\Omega}^2 + \lambda_2 \Phi(x)$

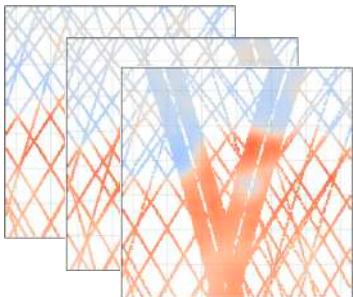
**Direct learning for inverse problems:**  $\hat{x} = \Psi(y)$   $y \rightarrow \text{CNN} \rightarrow x$



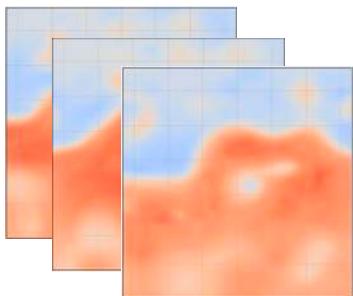
**Examples of CNN architectures:** Reaction-Diffusion architectures, ADMM-inspired architectures,...

**Good performance but possibly weak interpretability/generalization capacities of the solution beyond the training cases**

# End-to-end learning for inverse problems (Fablet et al., 2020)



Partial observations  $y$



True states  $x$

**Model-driven schemes:**  $\hat{x} = \arg \min_x \lambda_1 \|x - y\|_{\Omega}^2 + \lambda_2 \Phi(x)$

---

**Direct learning for inverse problems:**  $\hat{x} = \Psi(y)$   $y \rightarrow \text{CNN} \rightarrow x$

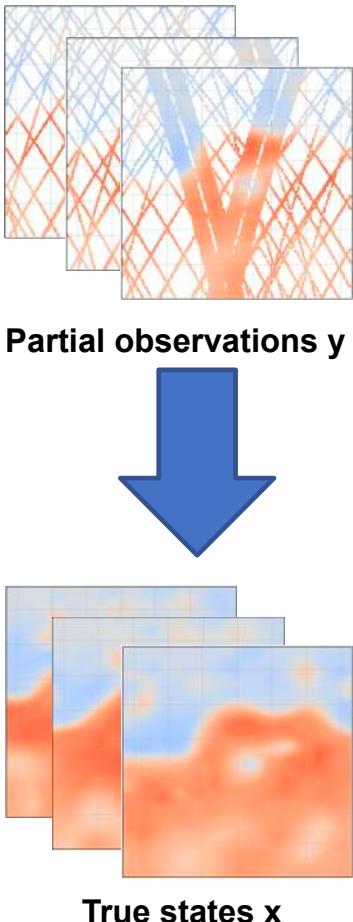
---

**Proposed scheme: joint learning of the variational model and solver**

- **Theoretical bi-level optimization**

$$\arg \min_{\Phi} \sum_n \|x_n - \tilde{x}_n\|^2 \text{ s.t. } \tilde{x}_n = \arg \min_{x_n} U_{\Phi}(x_n, y_n, \Omega_n)$$

# End-to-end learning for inverse problems (Fablet et al., 2020)



**Model-driven schemes:**  $\hat{x} = \arg \min_x \lambda_1 \|x - y\|_{\Omega}^2 + \lambda_2 \Phi(x)$

**Direct learning for inverse problems:**  $\hat{x} = \Psi(y)$      $y \rightarrow \text{CNN} \rightarrow x$

**Proposed scheme: joint learning of the variational model and solver**

- **Theoretical bi-level optimization**

$$\arg \min_{\Phi} \sum_n \|x_n - \tilde{x}_n\|^2 \text{ s.t. } \tilde{x}_n = \arg \min_{x_n} U_{\Phi}(x_n, y_n, \Omega_n)$$

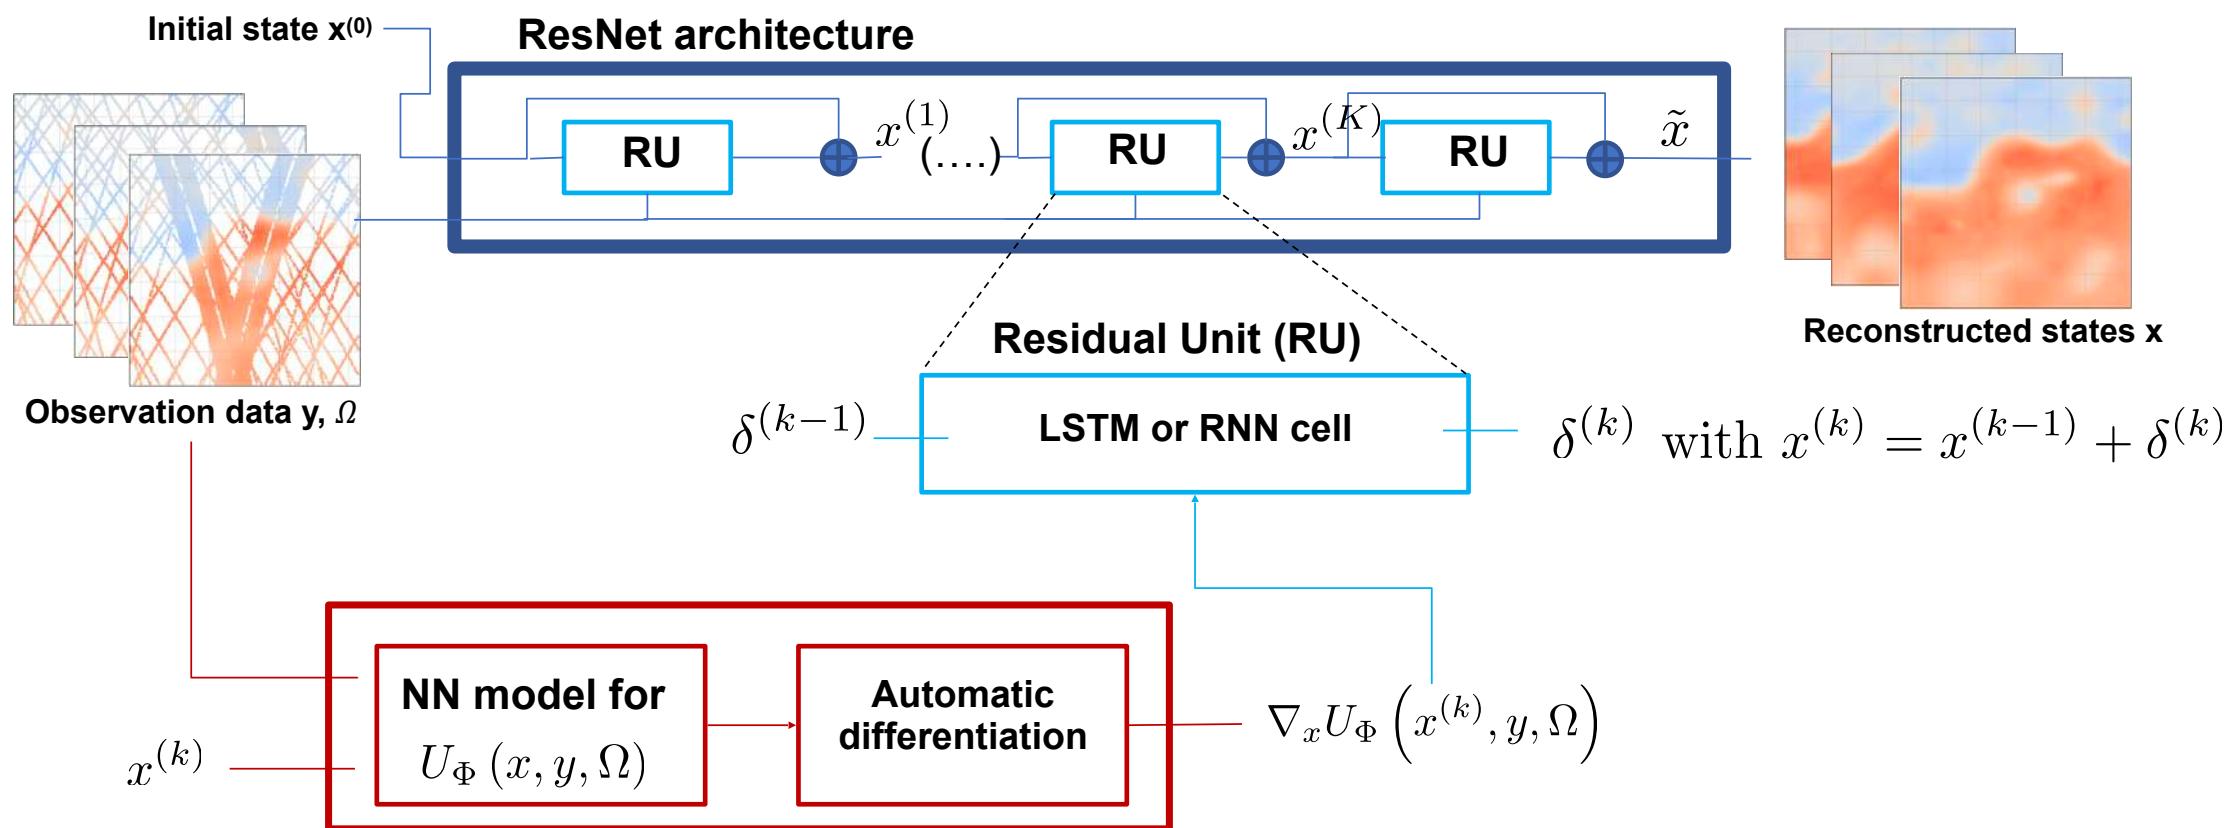
- **Restated with a gradient-based NN solver for inner minimization**

$$\arg \min_{\Phi, \Gamma} \sum_n \|x_n - \tilde{x}_n\|^2 \text{ s.t. } \tilde{x}_n = \Psi_{\Phi, \Gamma}(x_n^{(0)}, y_n, \Omega_n)$$

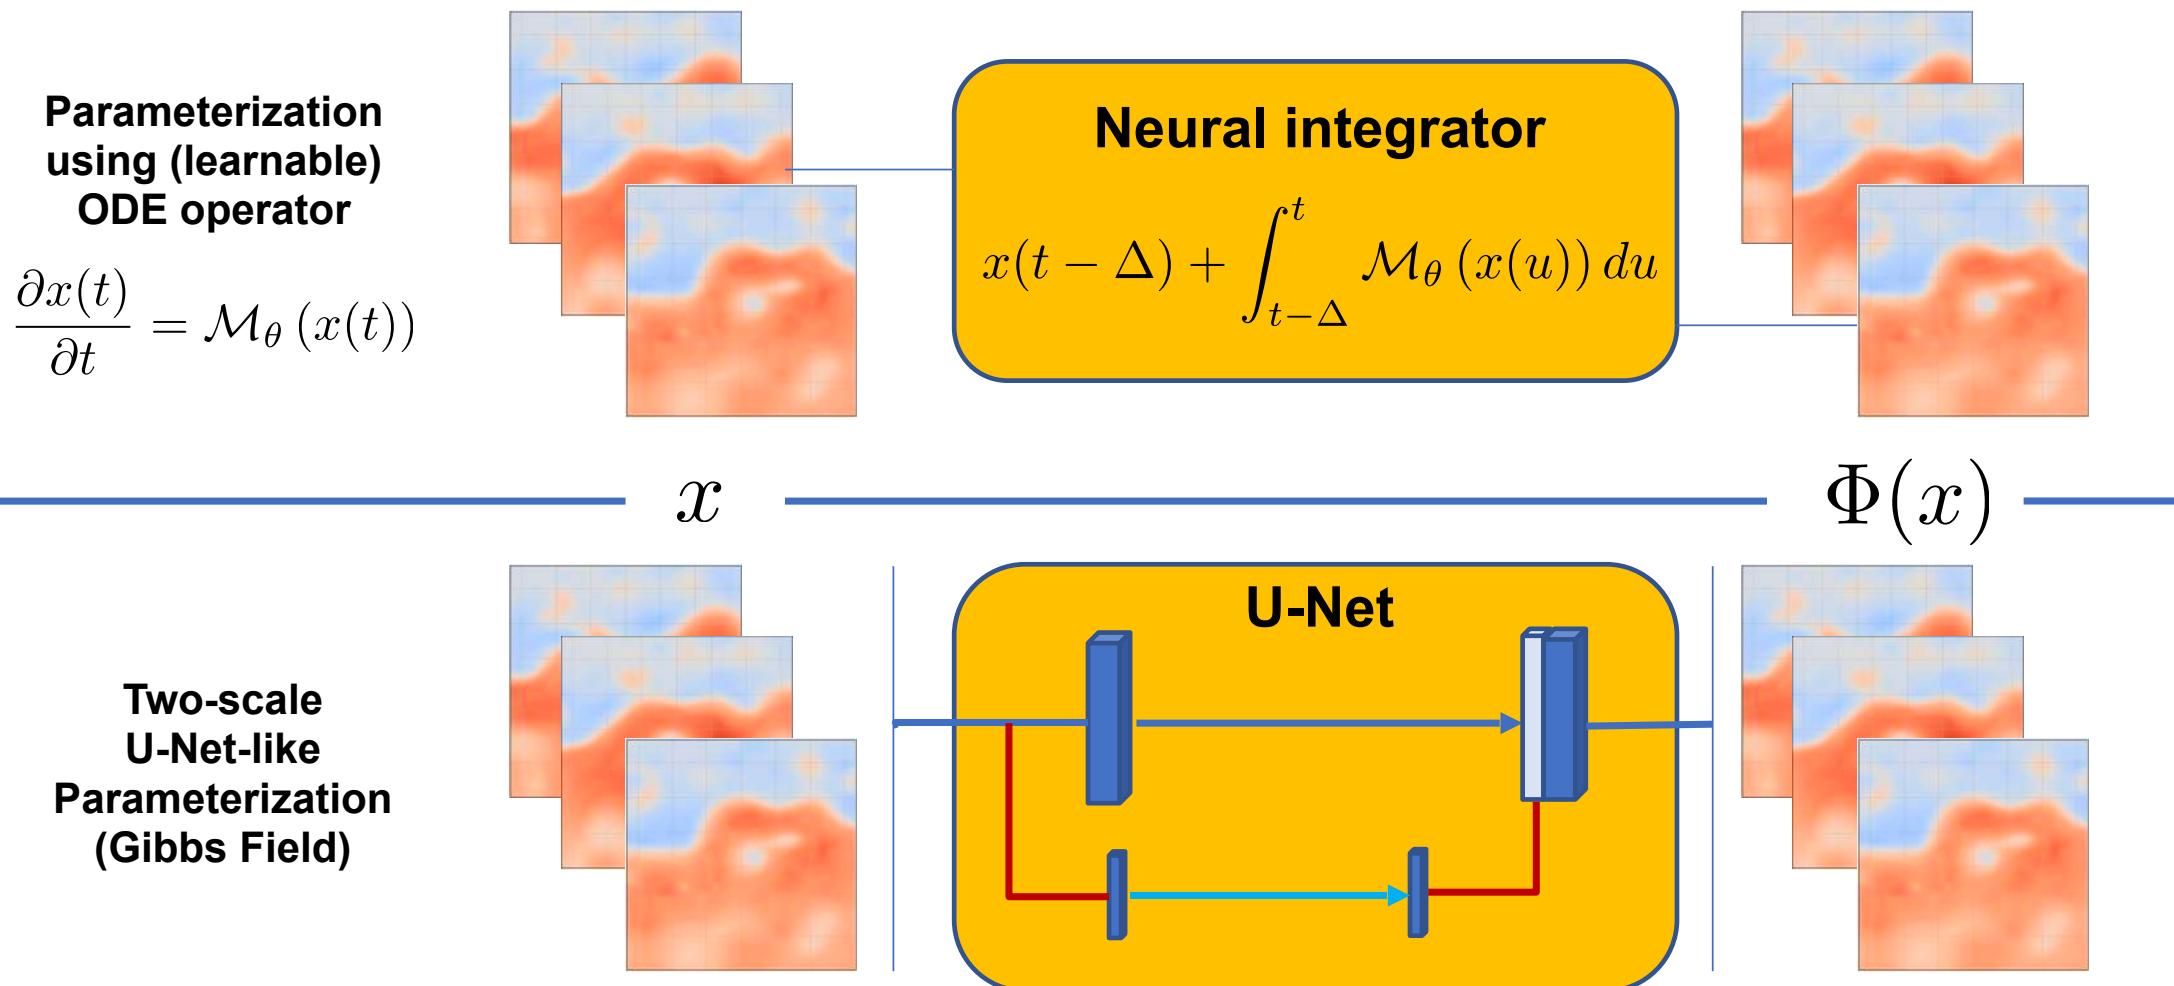
Iterative NN solver using automatic differentiation to compute gradient  $\nabla_x U_{\Phi}(x^{(k)}, y, \Omega)$

# End-to-end learning for inverse problems (Fablet et al., 2020)

## Proposed scheme: associated NN architecture

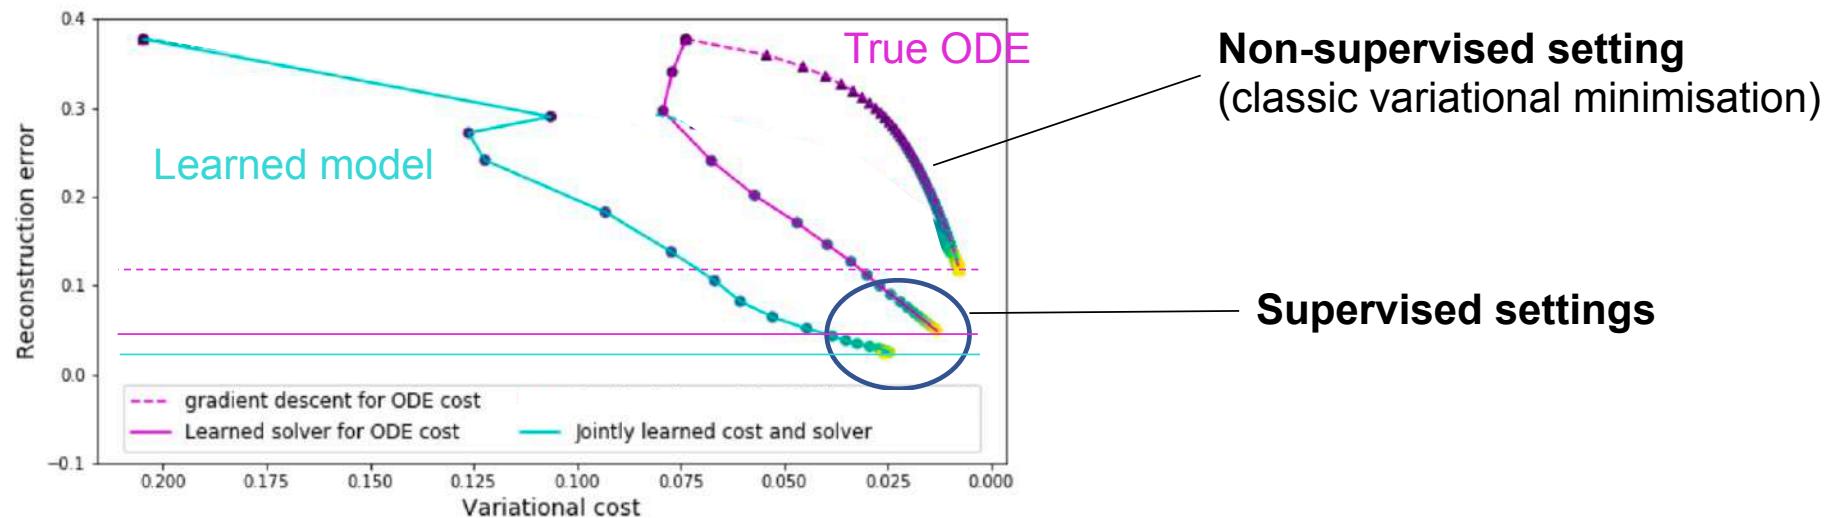


# End-to-end learning for 4DVar DA: projection operator $\Phi$

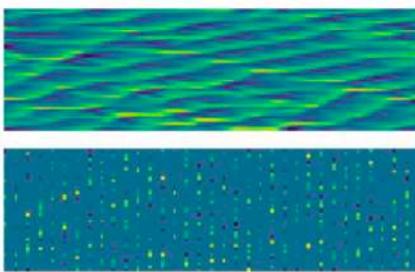


# End-to-end learning for inverse problems (Fablet et al., 2020)

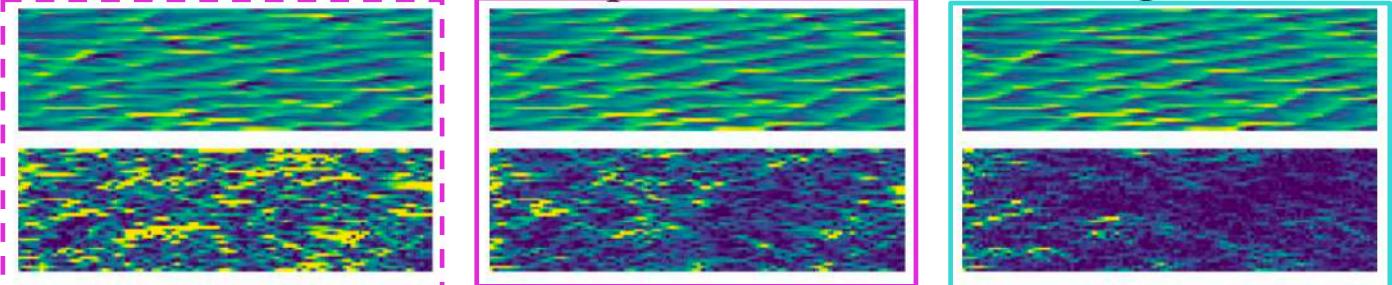
Illustration on Lorenz-96 dynamics (Bilinear ODE)



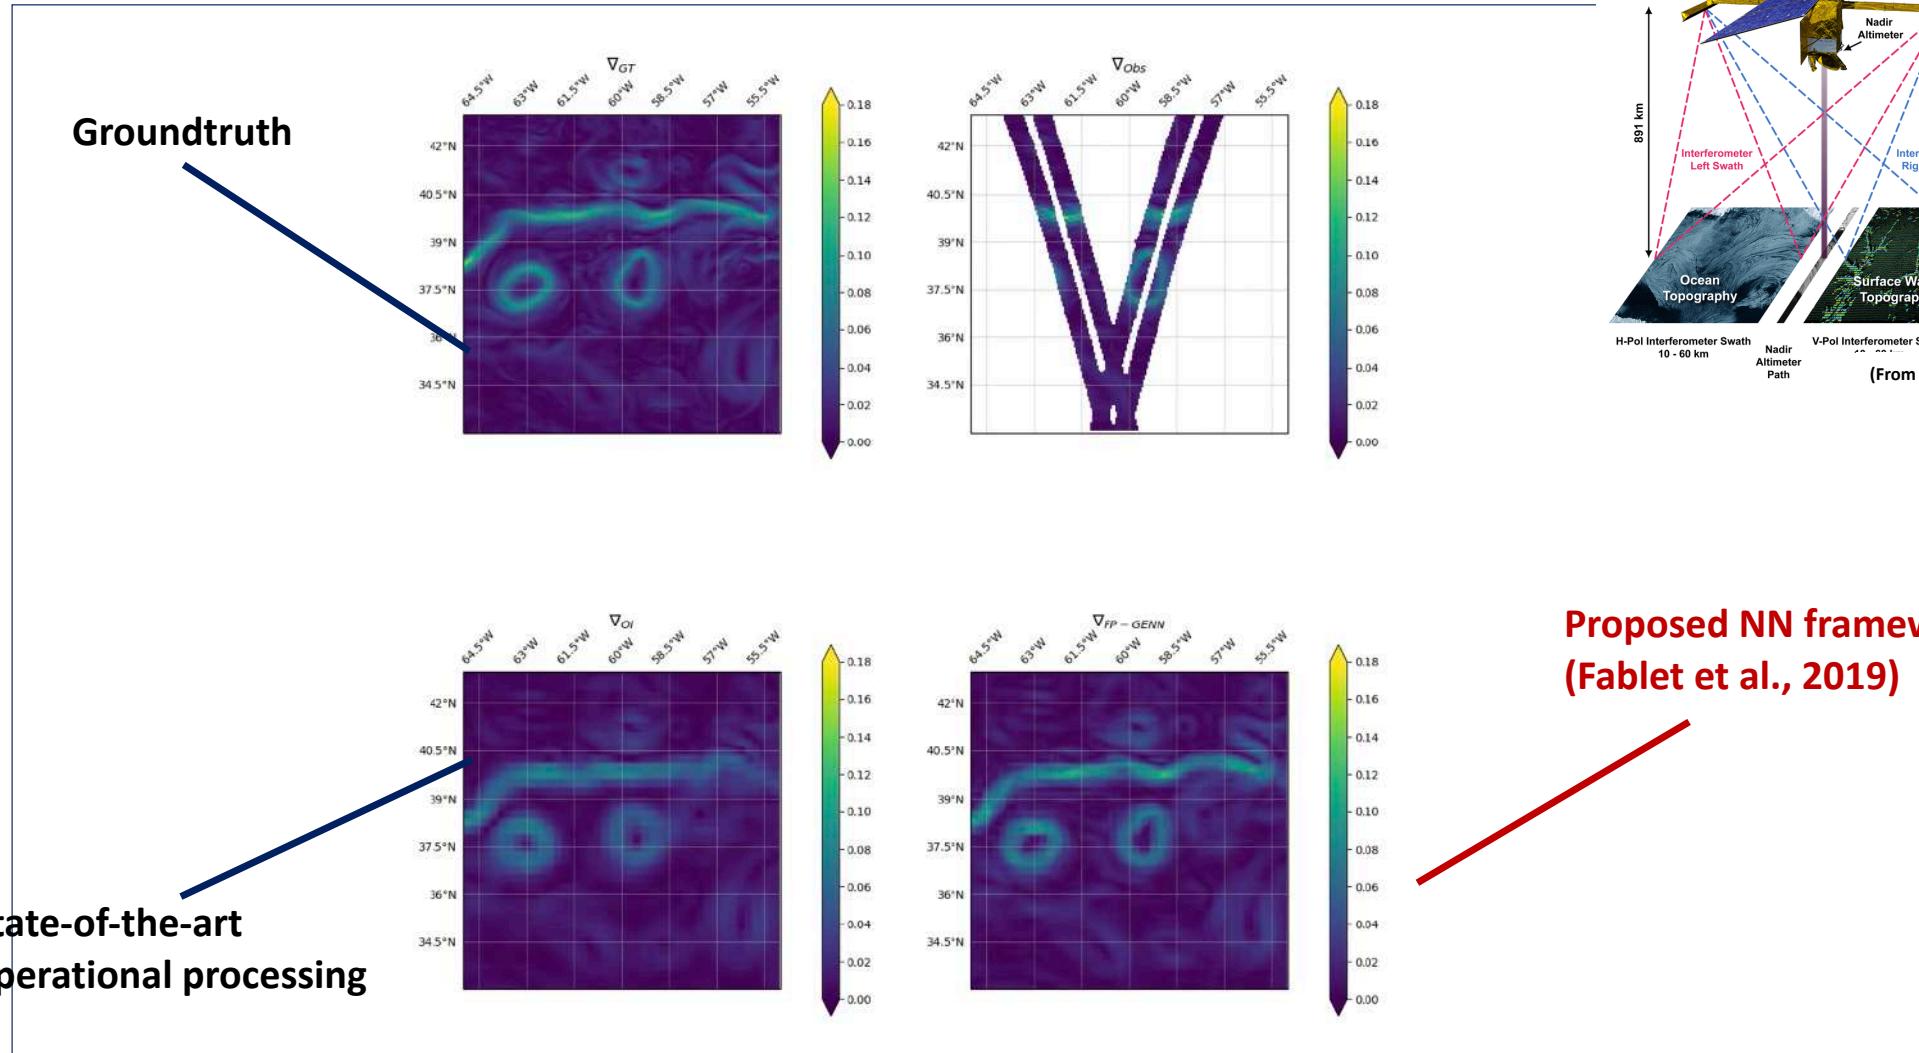
True and observed states



Reconstruction examples and associated error maps



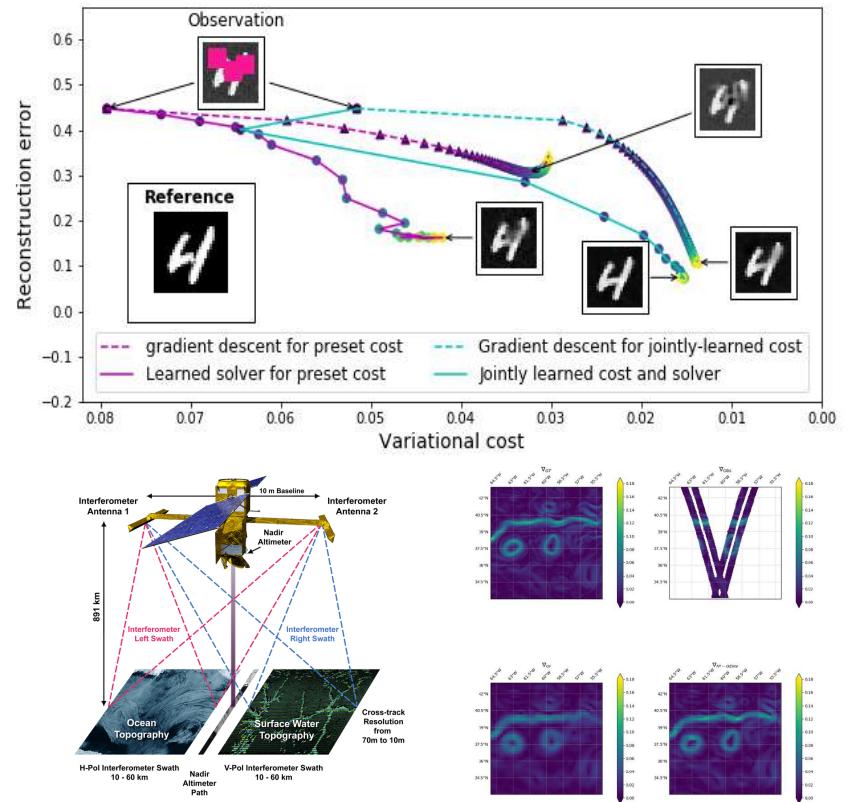
# An application for upcoming SWOT mission



# End-to-end learning for inverse problems (Fablet et al., 2020)

## Key messages

- We can bridge DNN and variational models to solve inverse problems
- Learning both variational priors and solvers using groundtruthed (simulation) or observation-only data
- The best model may not be the TRUE one for inverse problems
- Generic formulation/architecture beyond space-time dynamics

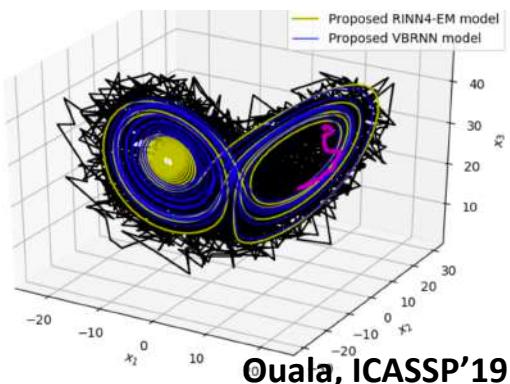


Preprint: <https://arxiv.org/abs/2006.03653>

Code: <https://github.com/CIA-Oceanix>

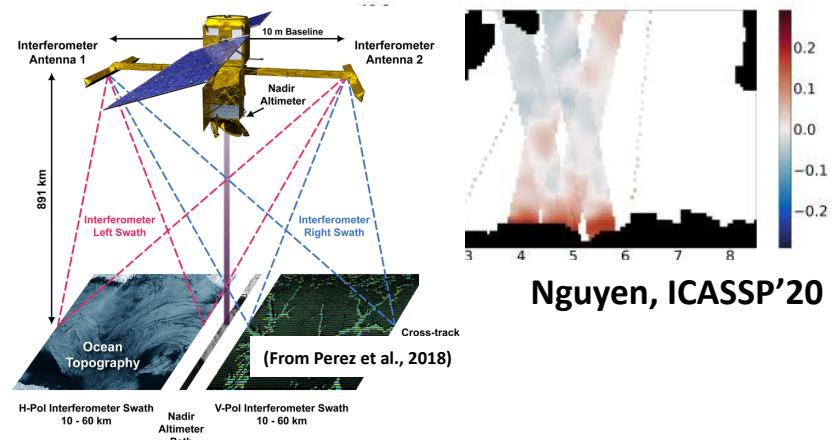
# End-to-end learning from real observation data ?

## Scarce time sampling



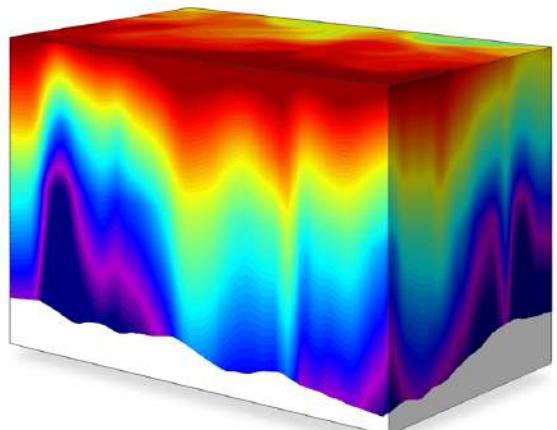
Ouala, ICASSP'19

## Noisy and irregular sampling



Nguyen, ICASSP'20

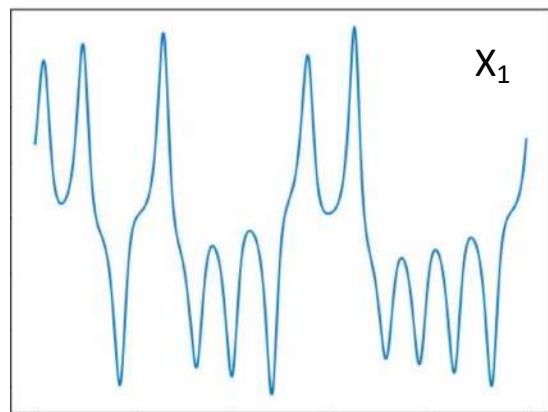
## Partially-observed system



Ouala, preprint 2019

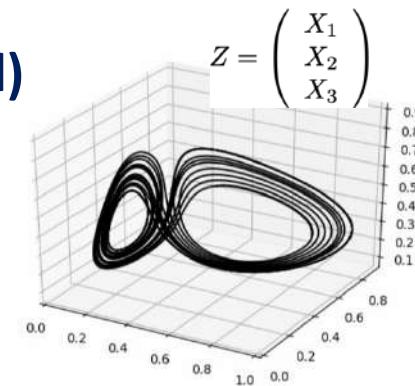
# Neural ODE for partially-observed systems [Ouala et al., 2020]

Illustration for L63 assuming only the first components is observed

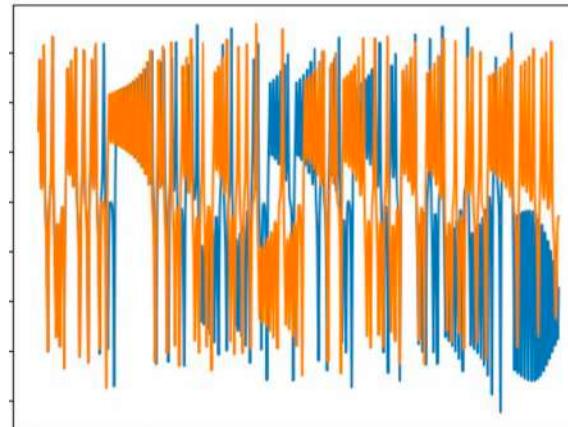


Learning Latent (unobserved) dynamics

$$d_t Z_t = \Phi_\theta(Z_t)$$



**Objectives:** accurate short-term forecast and realistic « long-term » patterns for  $X_1$



**Approach:** trainable variational formulation with latent dynamics

# Neural ODE for partially-observed systems [Ouala et al., 2019]

**Problem statement:** end-to-end learning of the latent (augmented) space and of the associated dynamics

$$X_t = \begin{pmatrix} x_t \\ z_t \end{pmatrix}$$

Observed variables  
Unknown variables

Dynamical model in the latent space

$$\partial_t X_t = f_\theta (X_t)$$

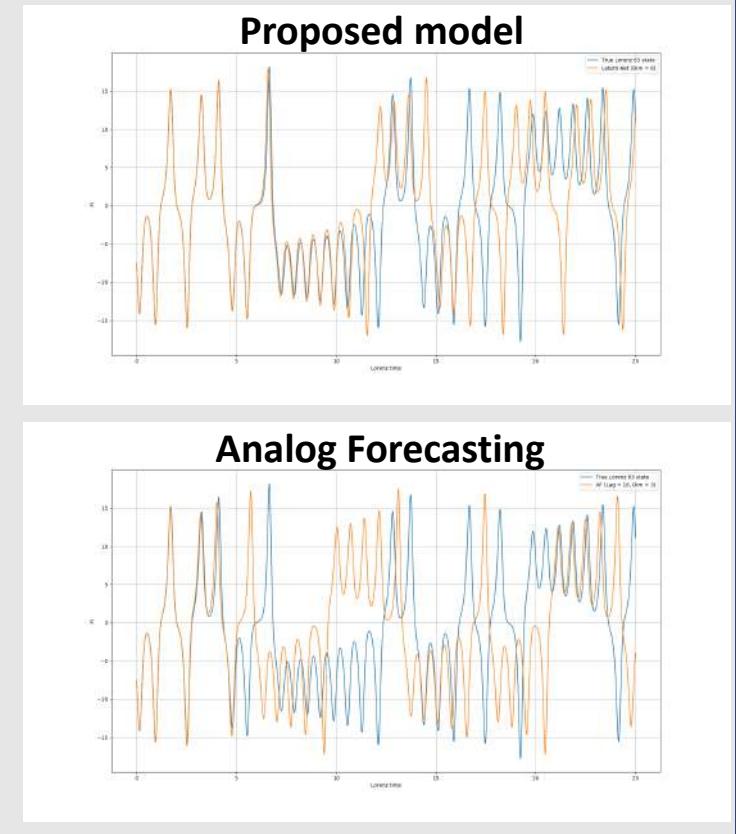
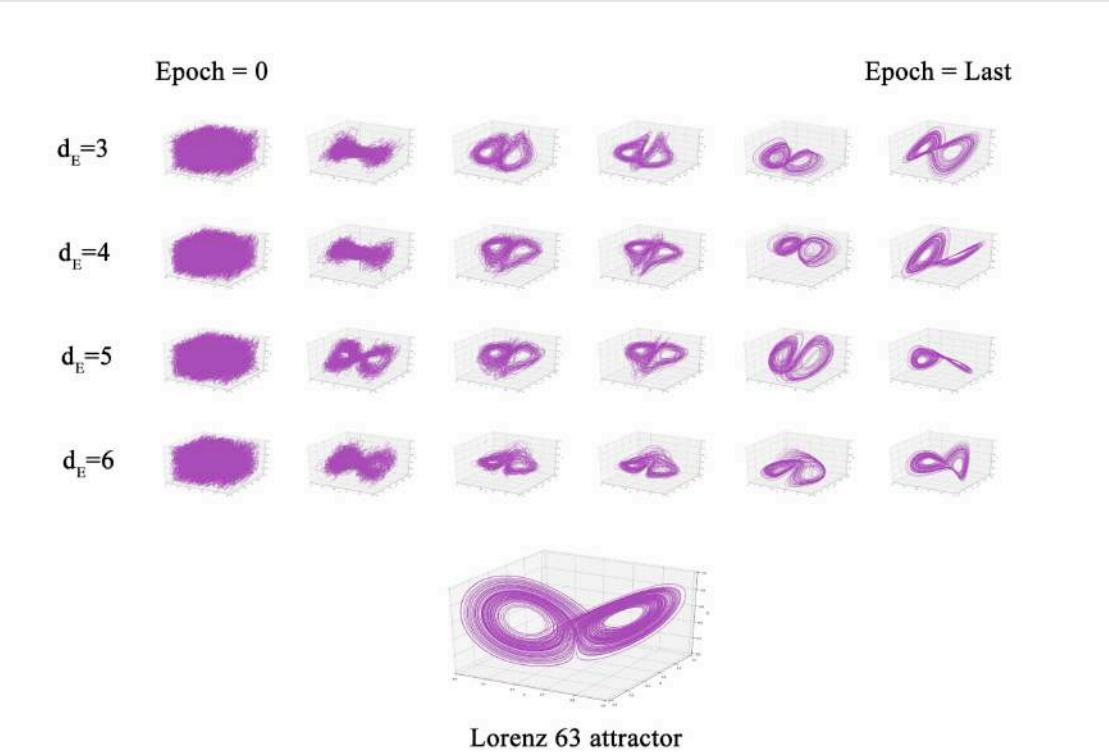
**Goals:**

1. Learn model parameters  $\theta$  from observed time series
2. Forecast future observed states given previous ones

**Proposed approach:** WC 4DVar formulation with an unknown dynamical model

# Neural ODE for partially-observed systems [Ouala et al., 2020]

## Illustration on Lorenz-63 dynamics



# Summary

- *NNs as numerical schemes for ODE/PDE/energy-based representations of geophysical flows*
- *Embedding geophysical priors in NN representations* (e.g., Lguensat et al., 2019; Ouala et al., 2019)
- *End-to-end architecture for jointly learning a representation (eg, ODE) and a solver* (e.g., Fablet et al., 2020)
- *Towards stochastic representations embedded in NN architectures* (e.g., Pannekoucke et al., 2020, Nguyen et al., 2020)

# Beyond Ocean Dynamics

Learning stochastic hidden dynamics

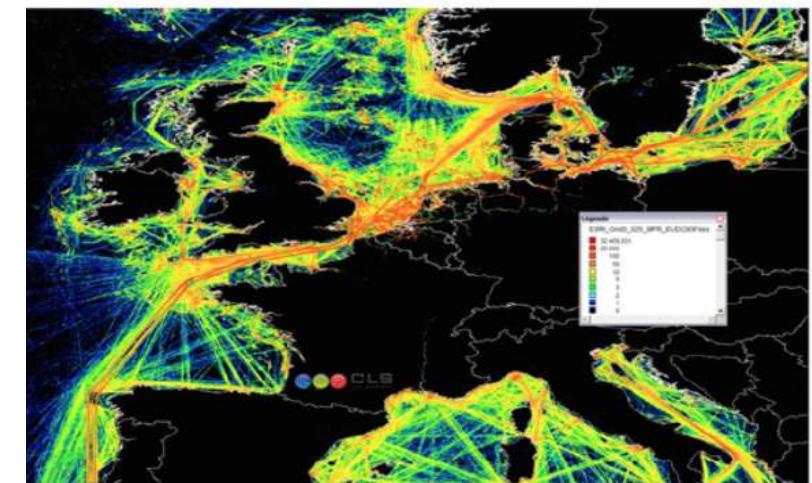


# Learning stochastic hidden dynamics [Nguyen et al., 2018]

The example of AIS Vessel trajectory data

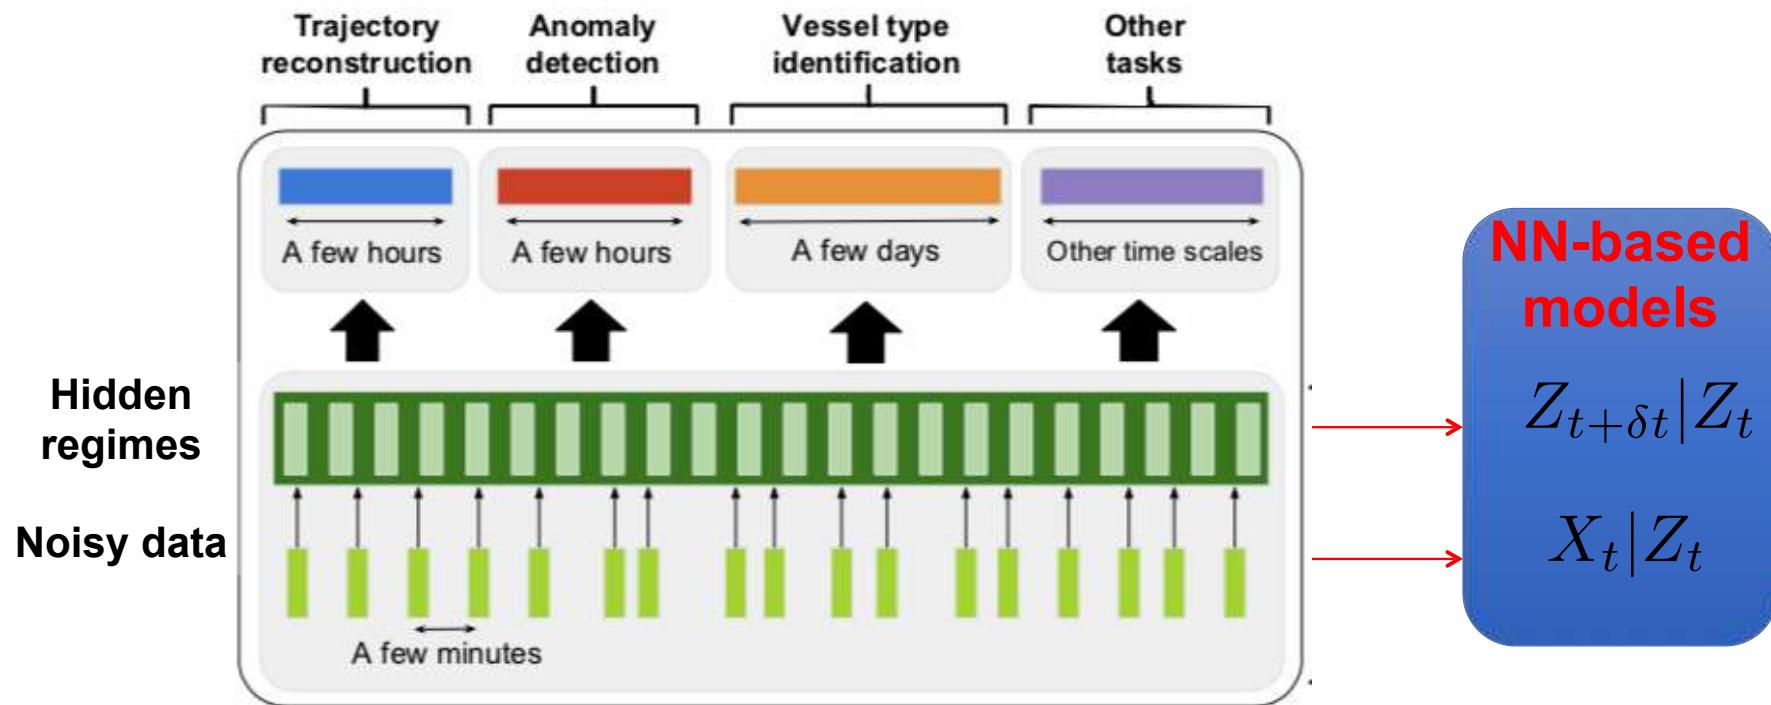


- Millions of AIS positions daily
- Noisy data: irregular sampling, corrupted data



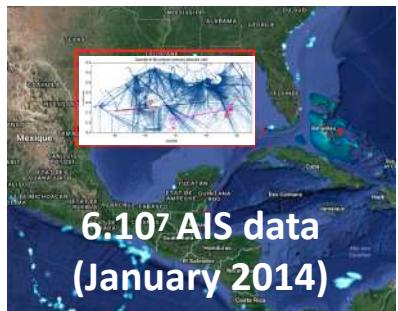
How can we learn from AIS data streams ?

# Learning stochastic hidden dynamics [Nguyen et al., 2018]



Model training from noisy AIS streams using variational Bayesian approximation

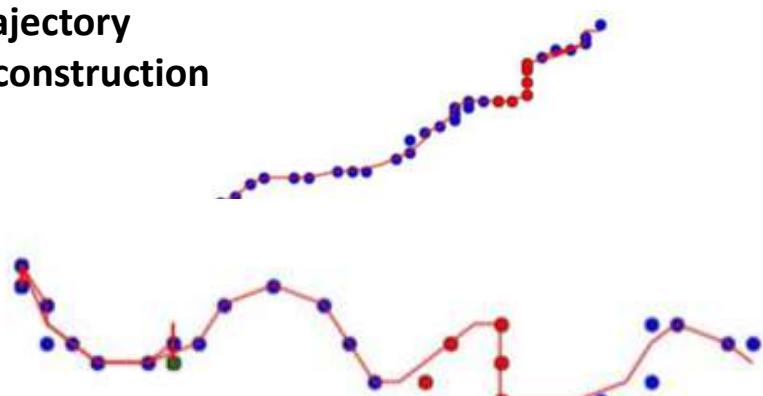
# Learning stochastic hidden dynamics [Nguyen et al., 2018]



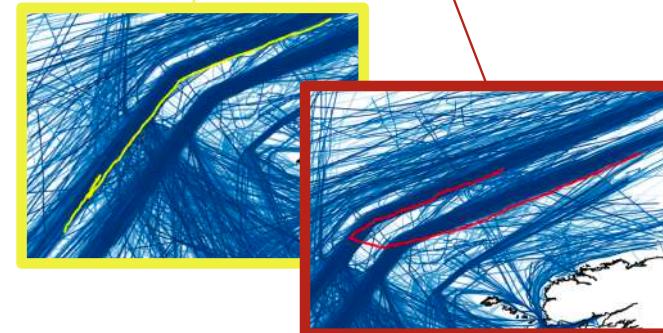
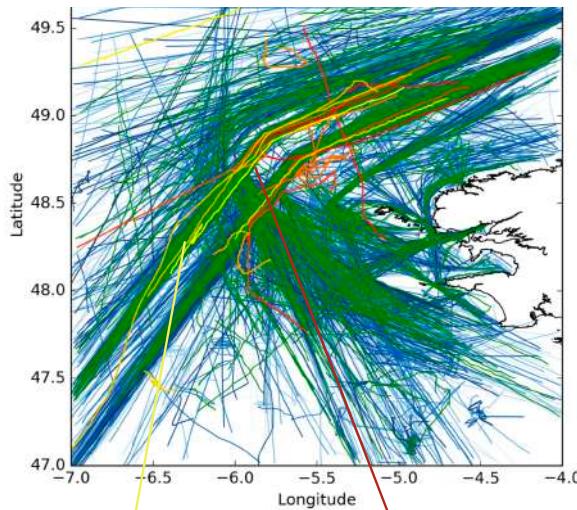
Vessel type recognition

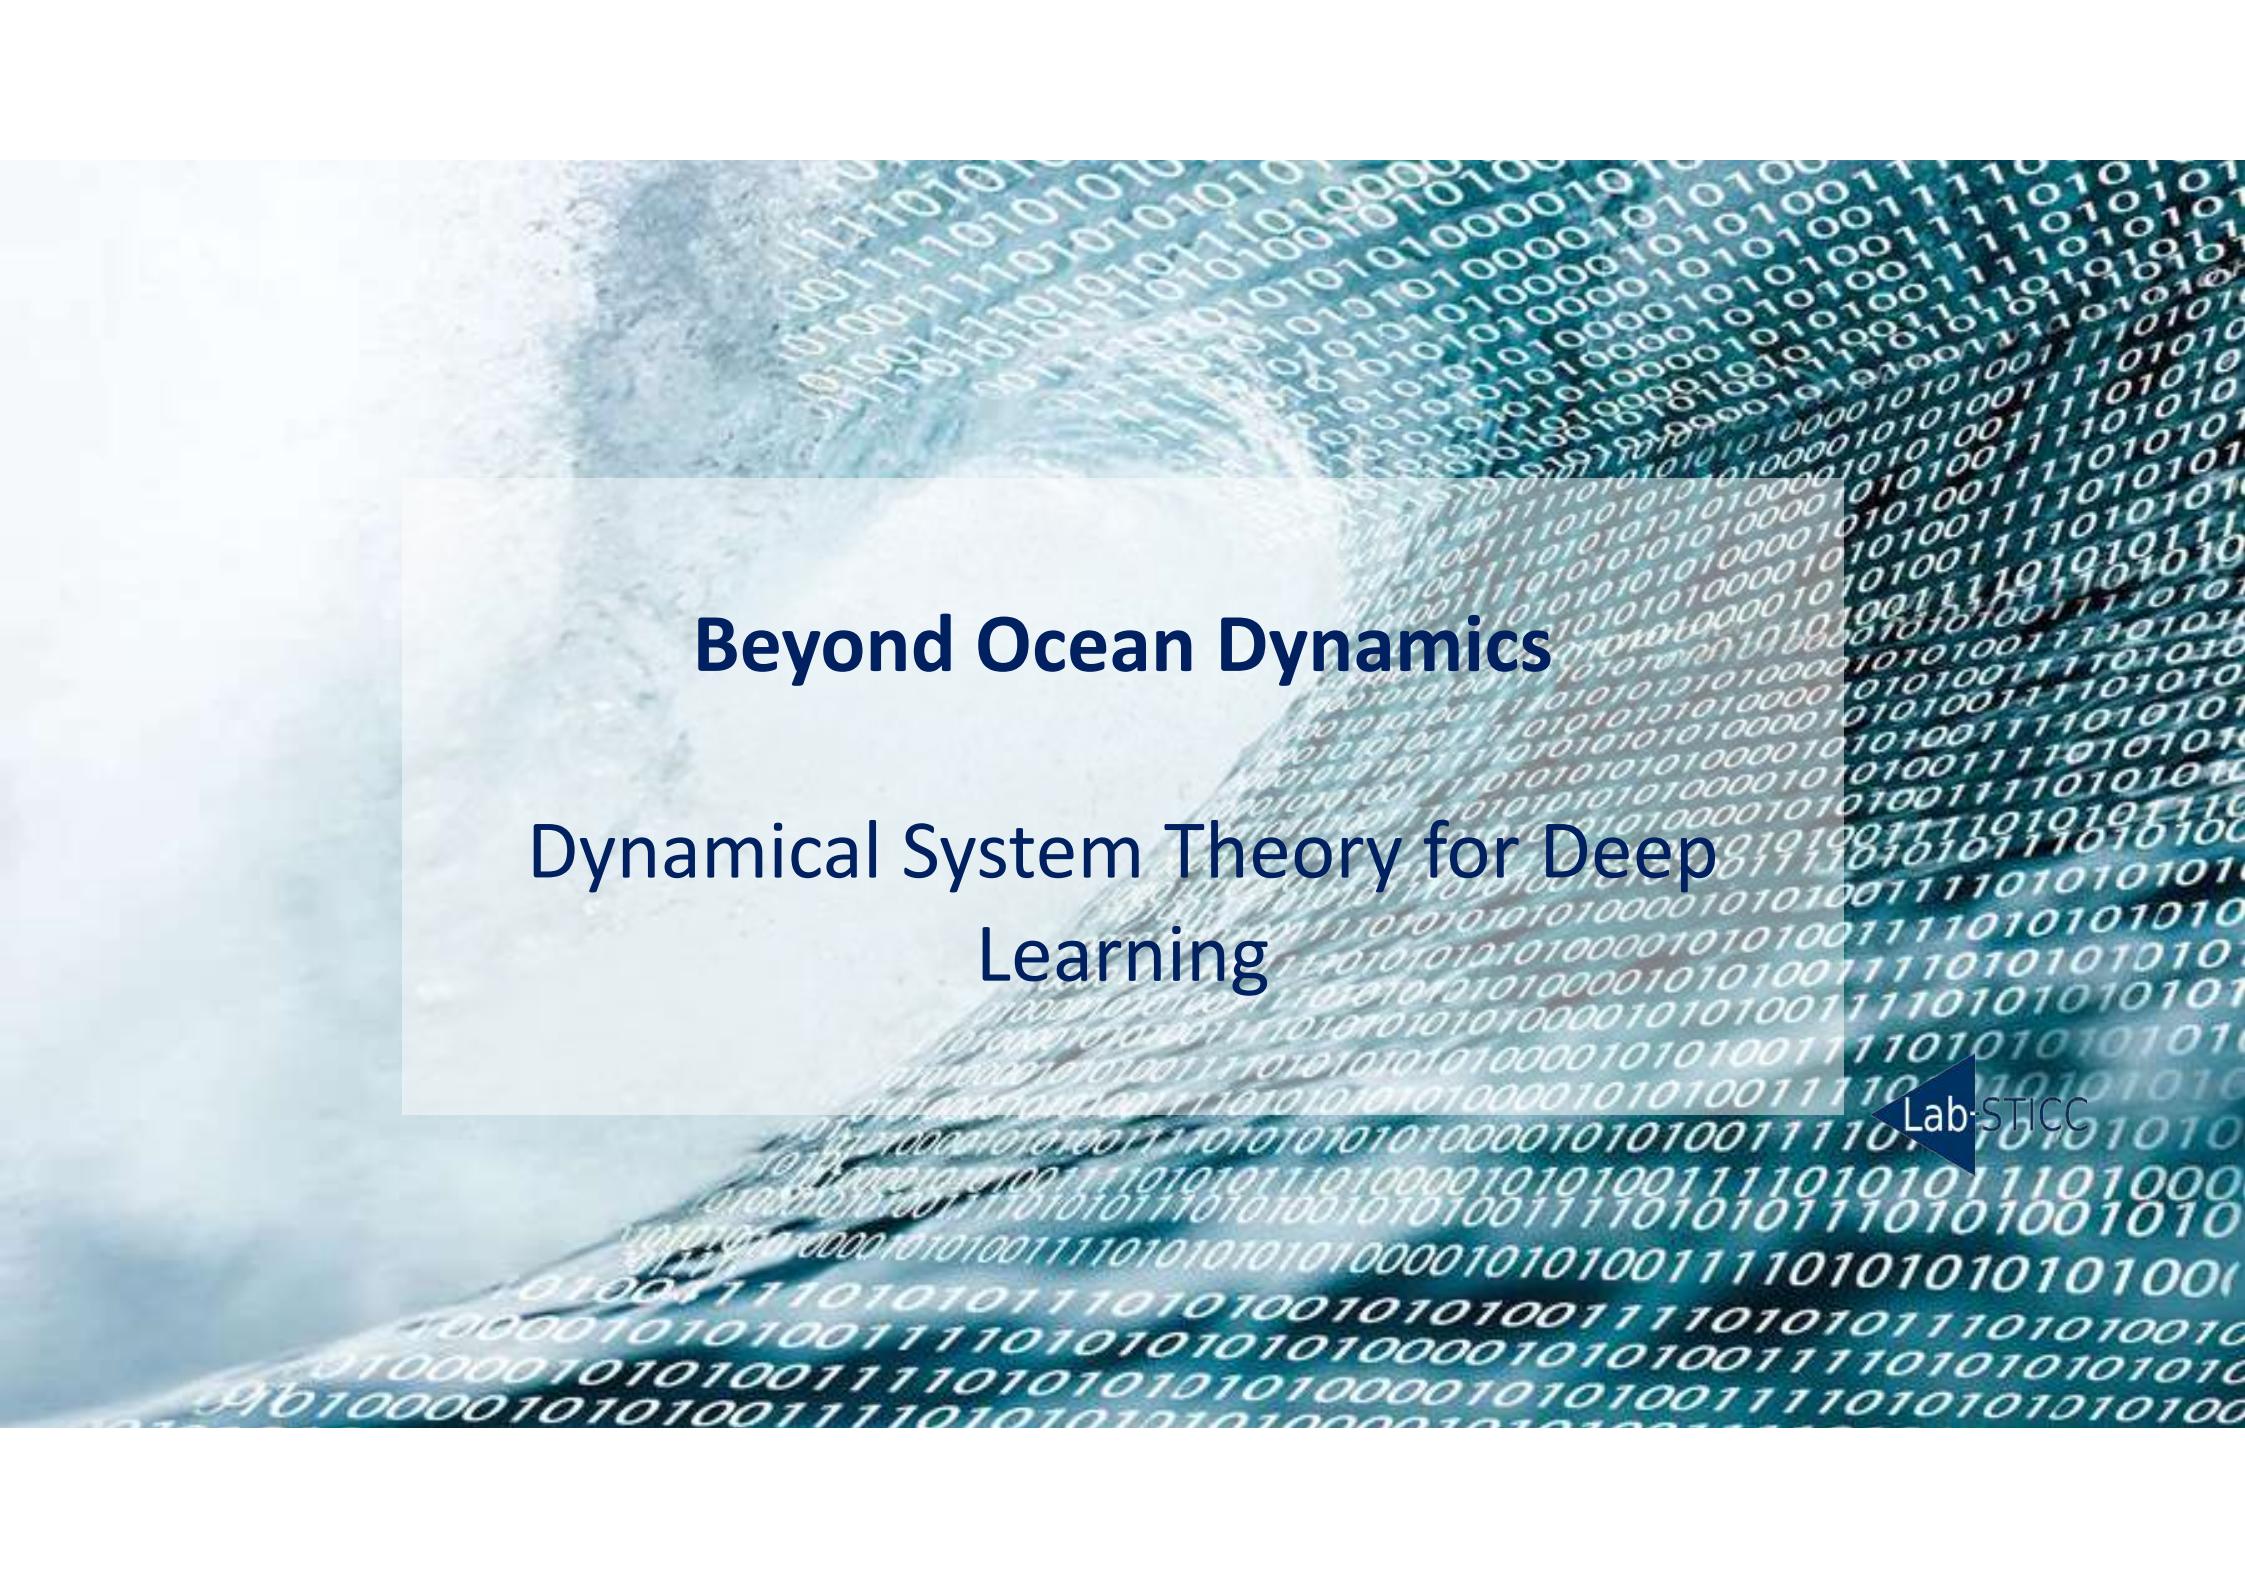
~88% of correct  
recognition

Trajectory  
reconstruction



Abnormal behaviour detection





# Beyond Ocean Dynamics

## Dynamical System Theory for Deep Learning



# Understanding DL models ?



88% **tabby cat**

adversarial  
perturbation

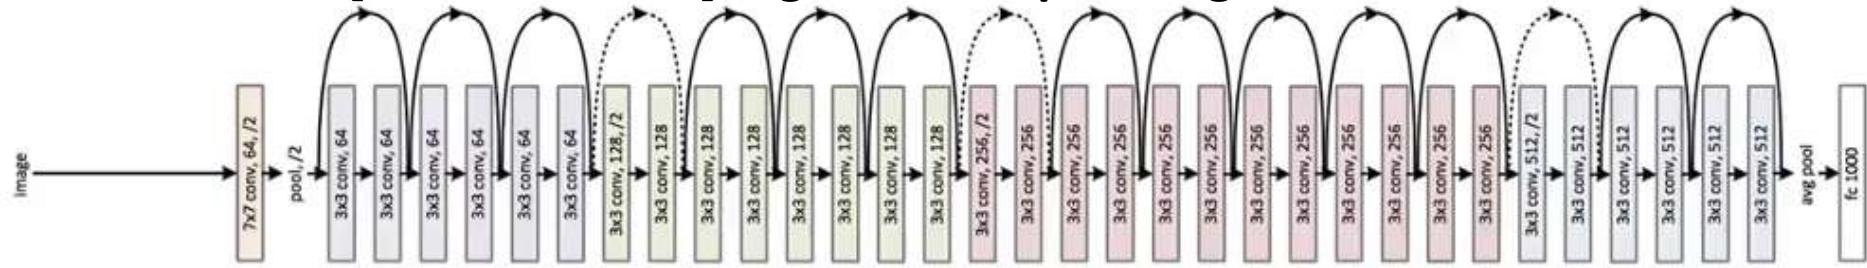


99% **guacamole**

Szegedy et al., 2015

# Understanding ResNets [Rousseau et al., 2019]

ResNet [He et al., 2015] regarded as space registration machines



- Image registration examples

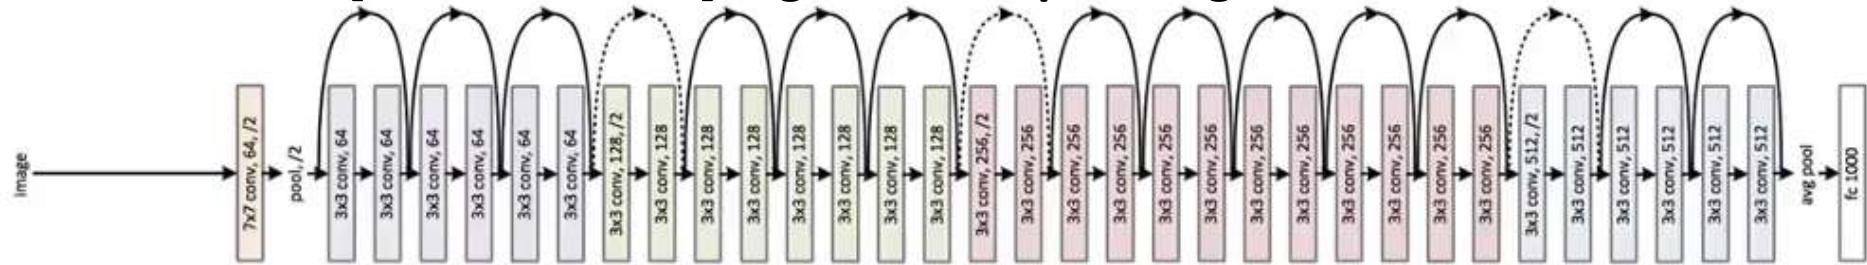


[Matlab tutorial]

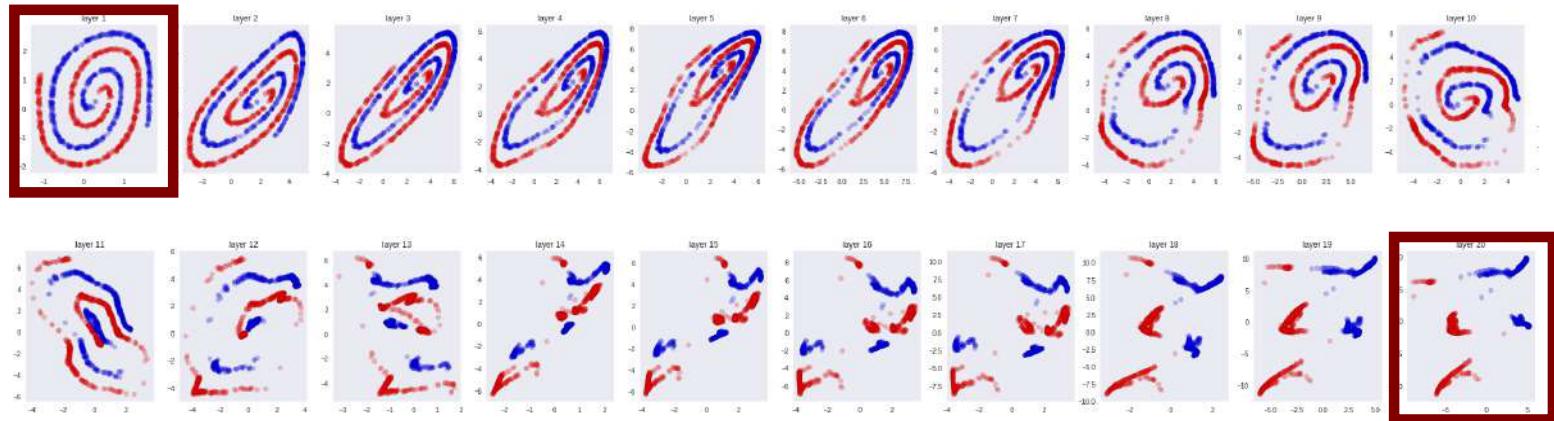
[Dramms tutorial]

# Understanding ResNets [Rousseau et al., 2019]

ResNet [He et al., 2015] regarded as space registration machines



Original  
feature  
space



Registered space to make feasible a linear  
separation between classes



**Thank you.**

# AI Chair OceaniX 2020-2024

Physics-informed AI for Observation-  
Driven Ocean AnalytiX

PI: **R. Fablet**, Prof. IMT Atlantique, Brest

Web: <https://cia-oceanix.github.io/>



**Internship, PhD  
and postdoc  
opportunities**

