NVIDIA

MACHINE LEARNING FOR WEATHER

David M. Hall   Senior Solution Architect   ECMWF-UEF   May 2020

AGENDA

OVERVIEW

What is machine learning? And why is it useful?

TOOLS

What do we need, to make it work?

APPLICATIONS

What precisely can we do with it?

CHALLENGES

What challenges remain, and how might they be addressed?

MACHINE LEARNING: A NEW SET OF TOOLS FOR SCIENCE

Machine learning provides a new approach for building software, by reverse-engineering functions from a set of examples. This approach complements traditional algorithm development, providing a means of devising algorithms too complex, subtle, or unintuitive to code by hand.

REVERSE-ENGINEERING FUNCTIONS FROM EXAMPLES

Find f, given x and y

MACHINE LEARNING

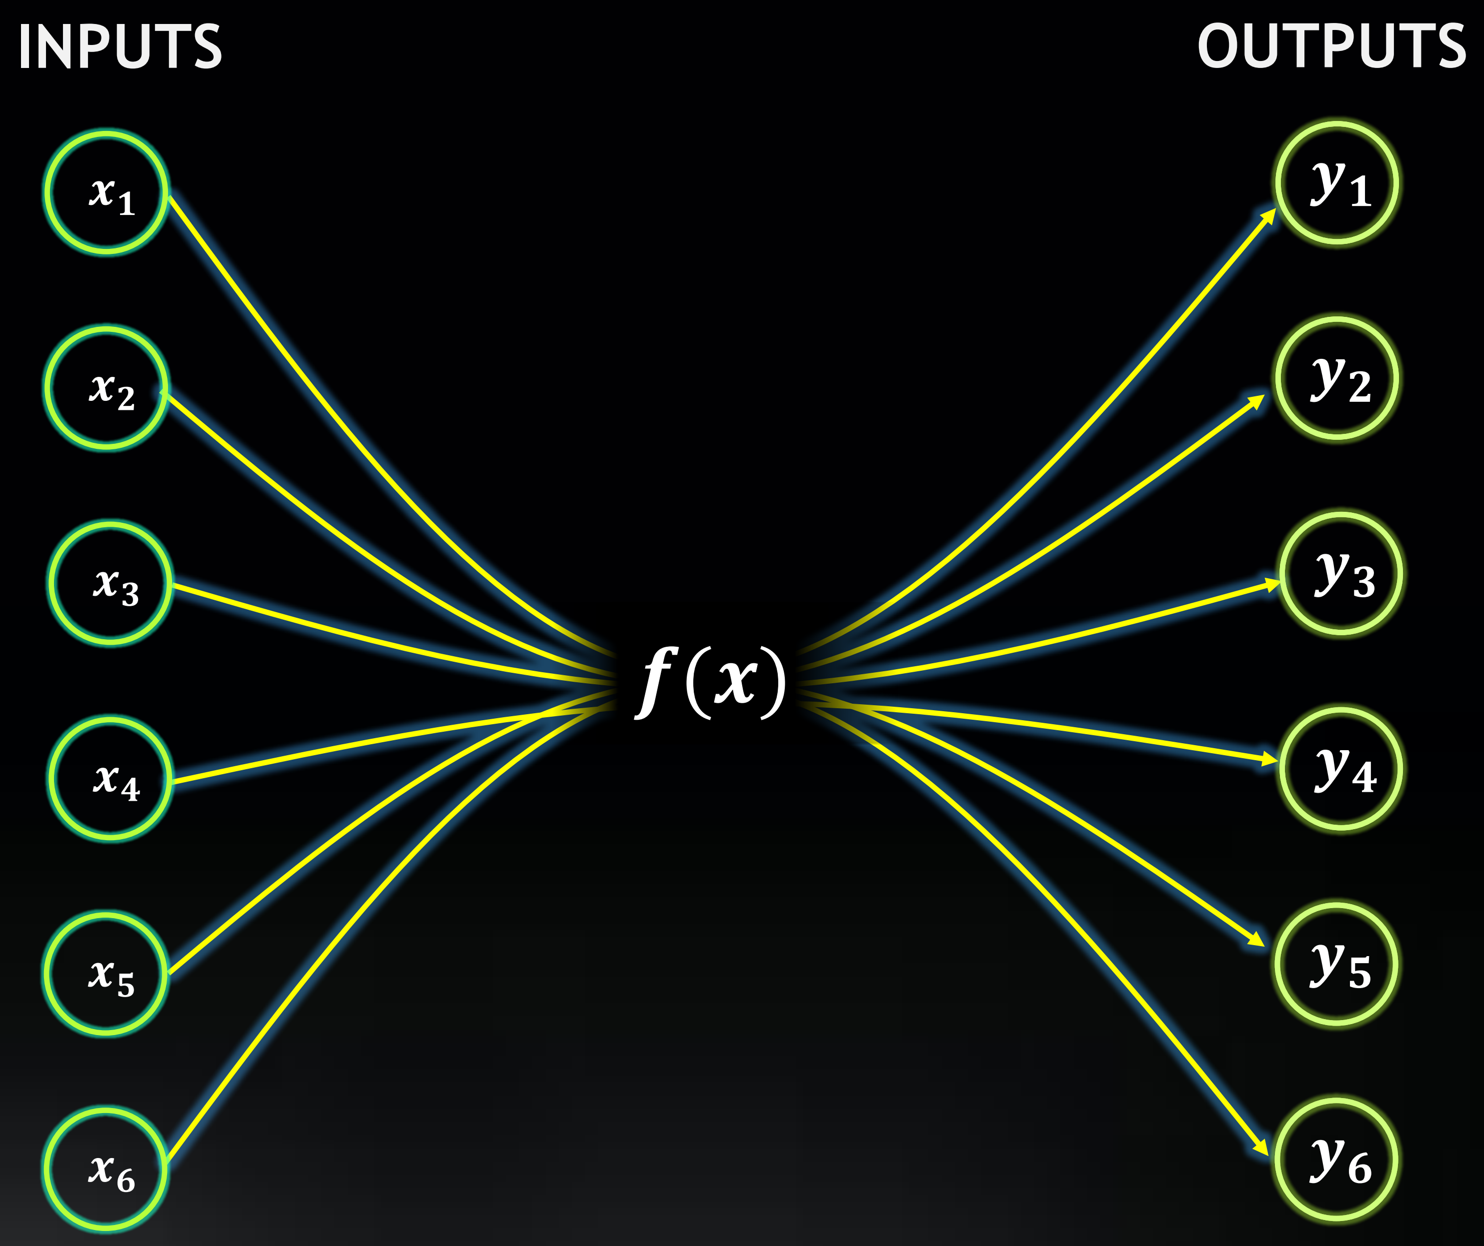

A GENERALIZATION OF CURVE FITTING

Find f, given x and y

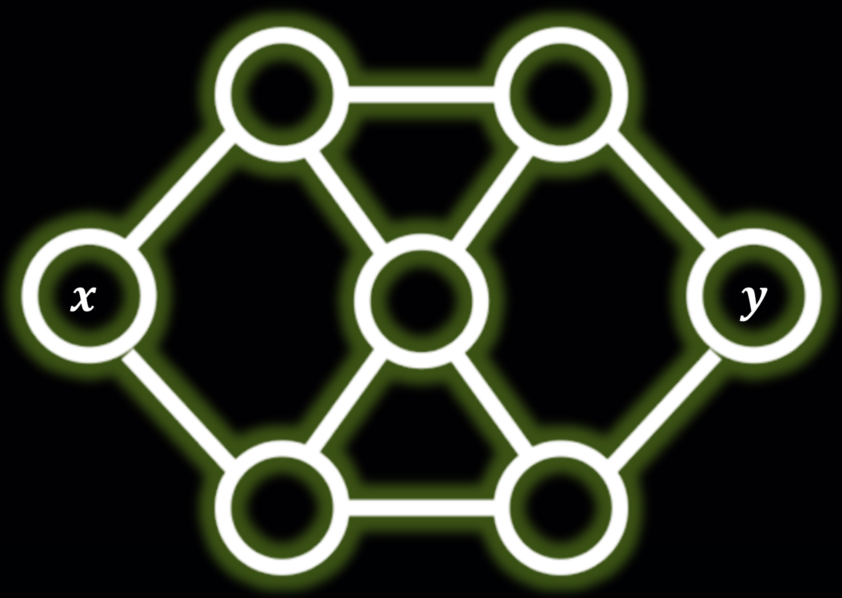

Machine Learning

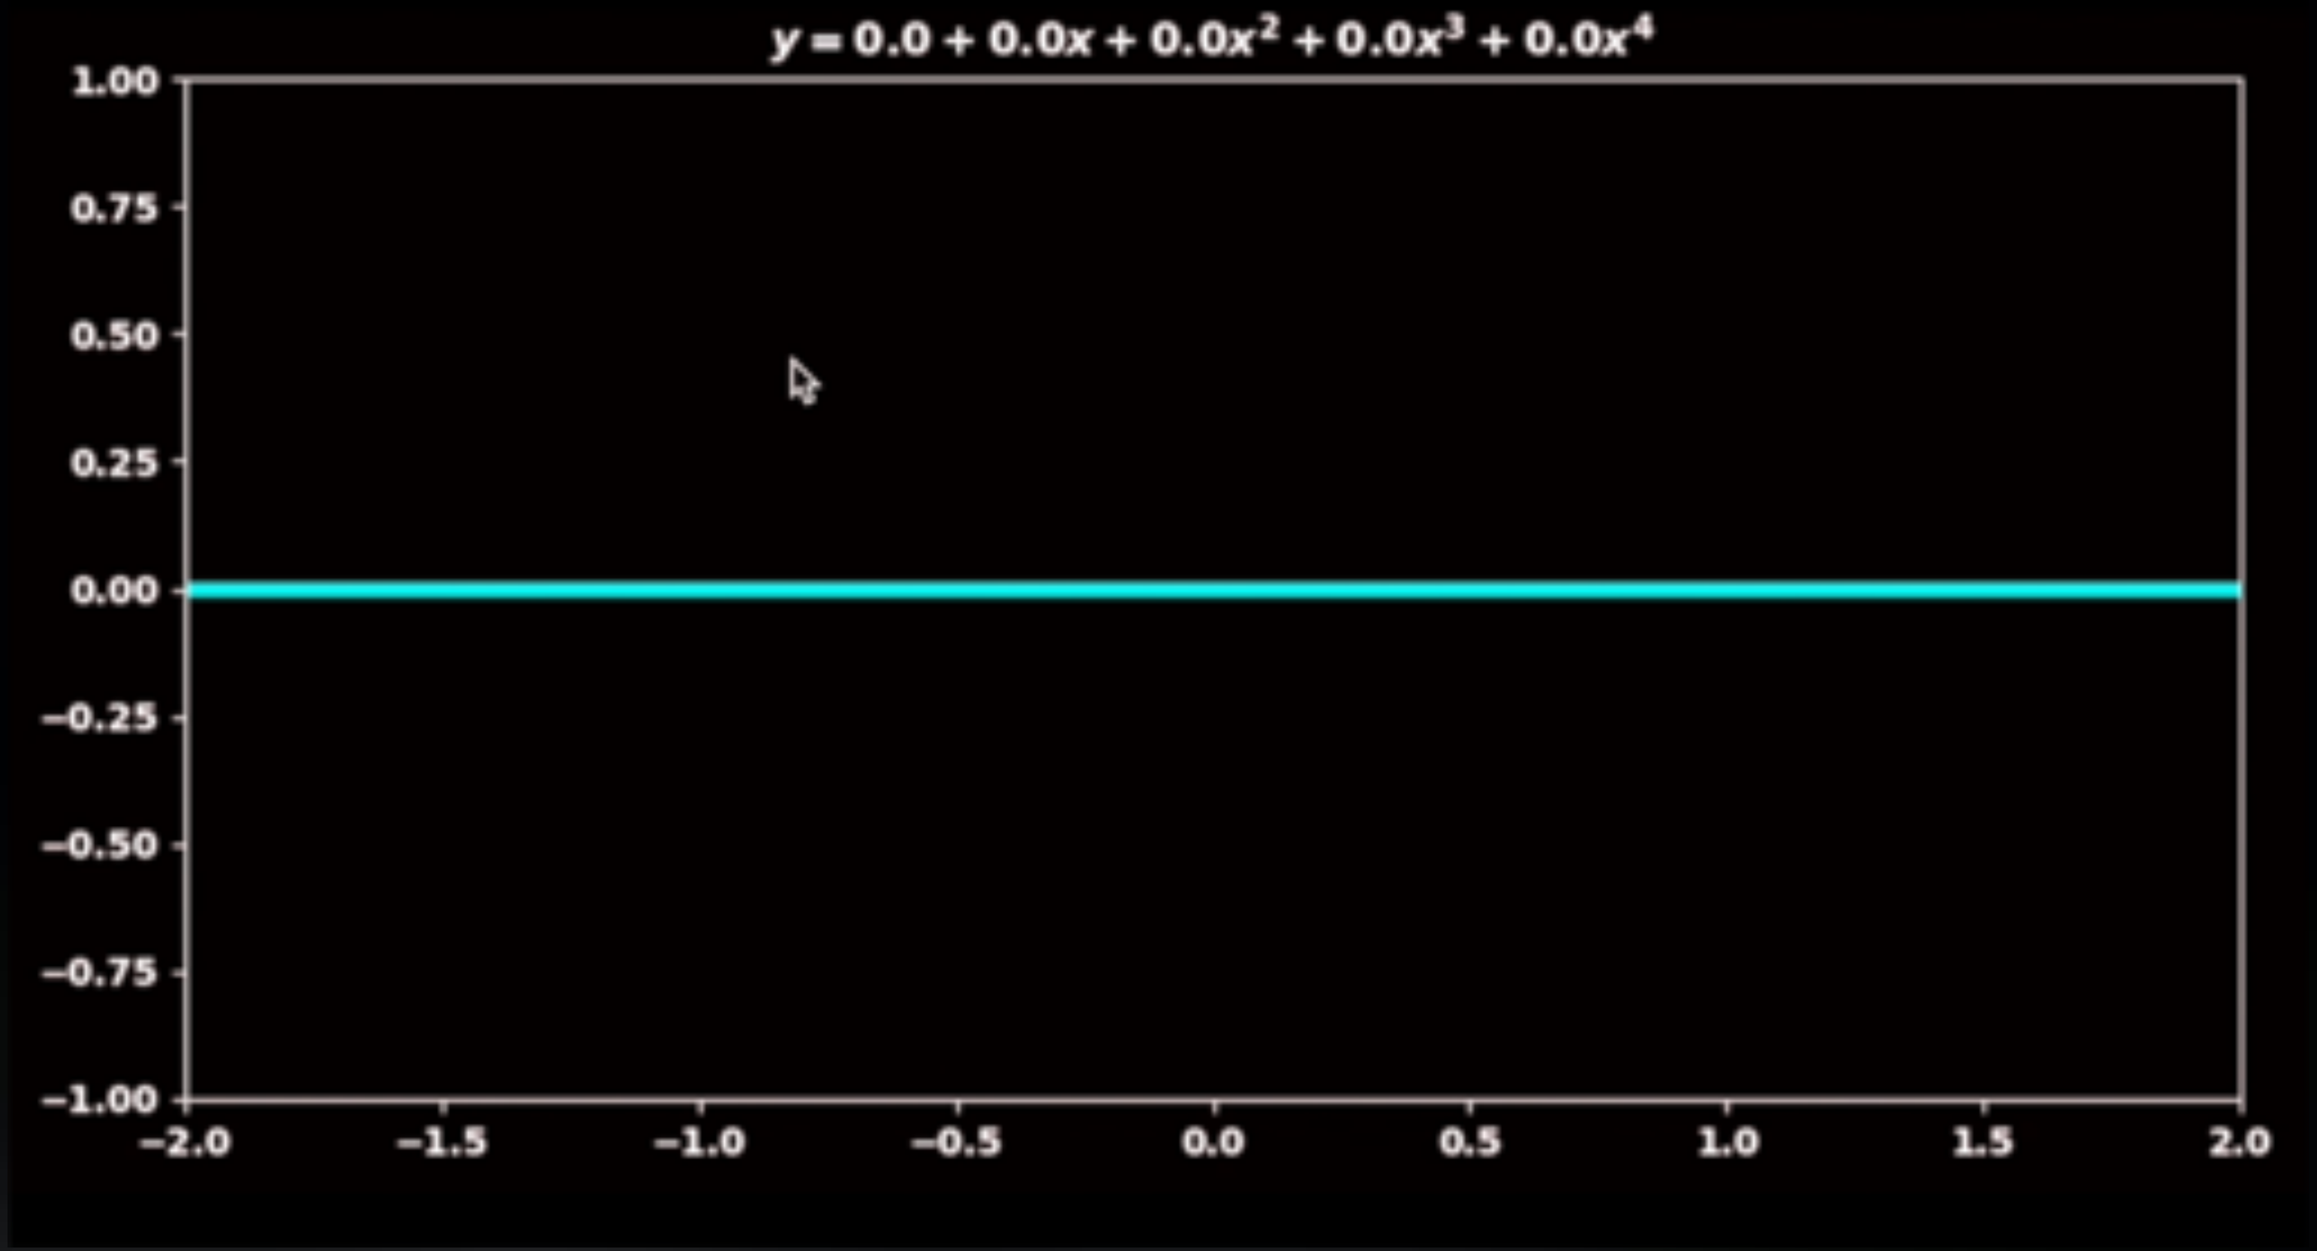

HIGH DIMENSIONS, LARGE MODELS

Find f, given x and y

Machine Learning

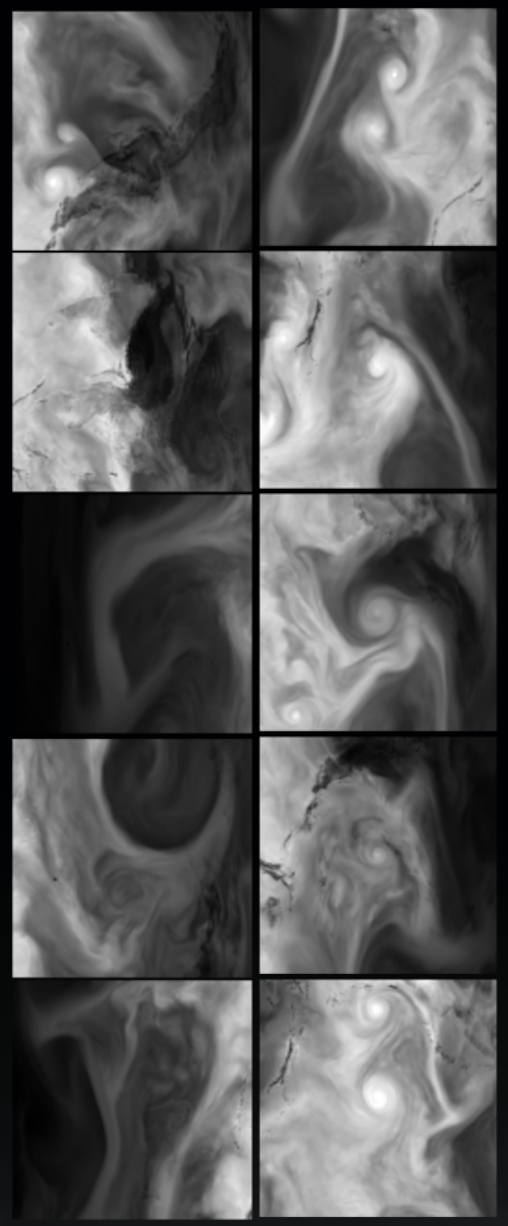

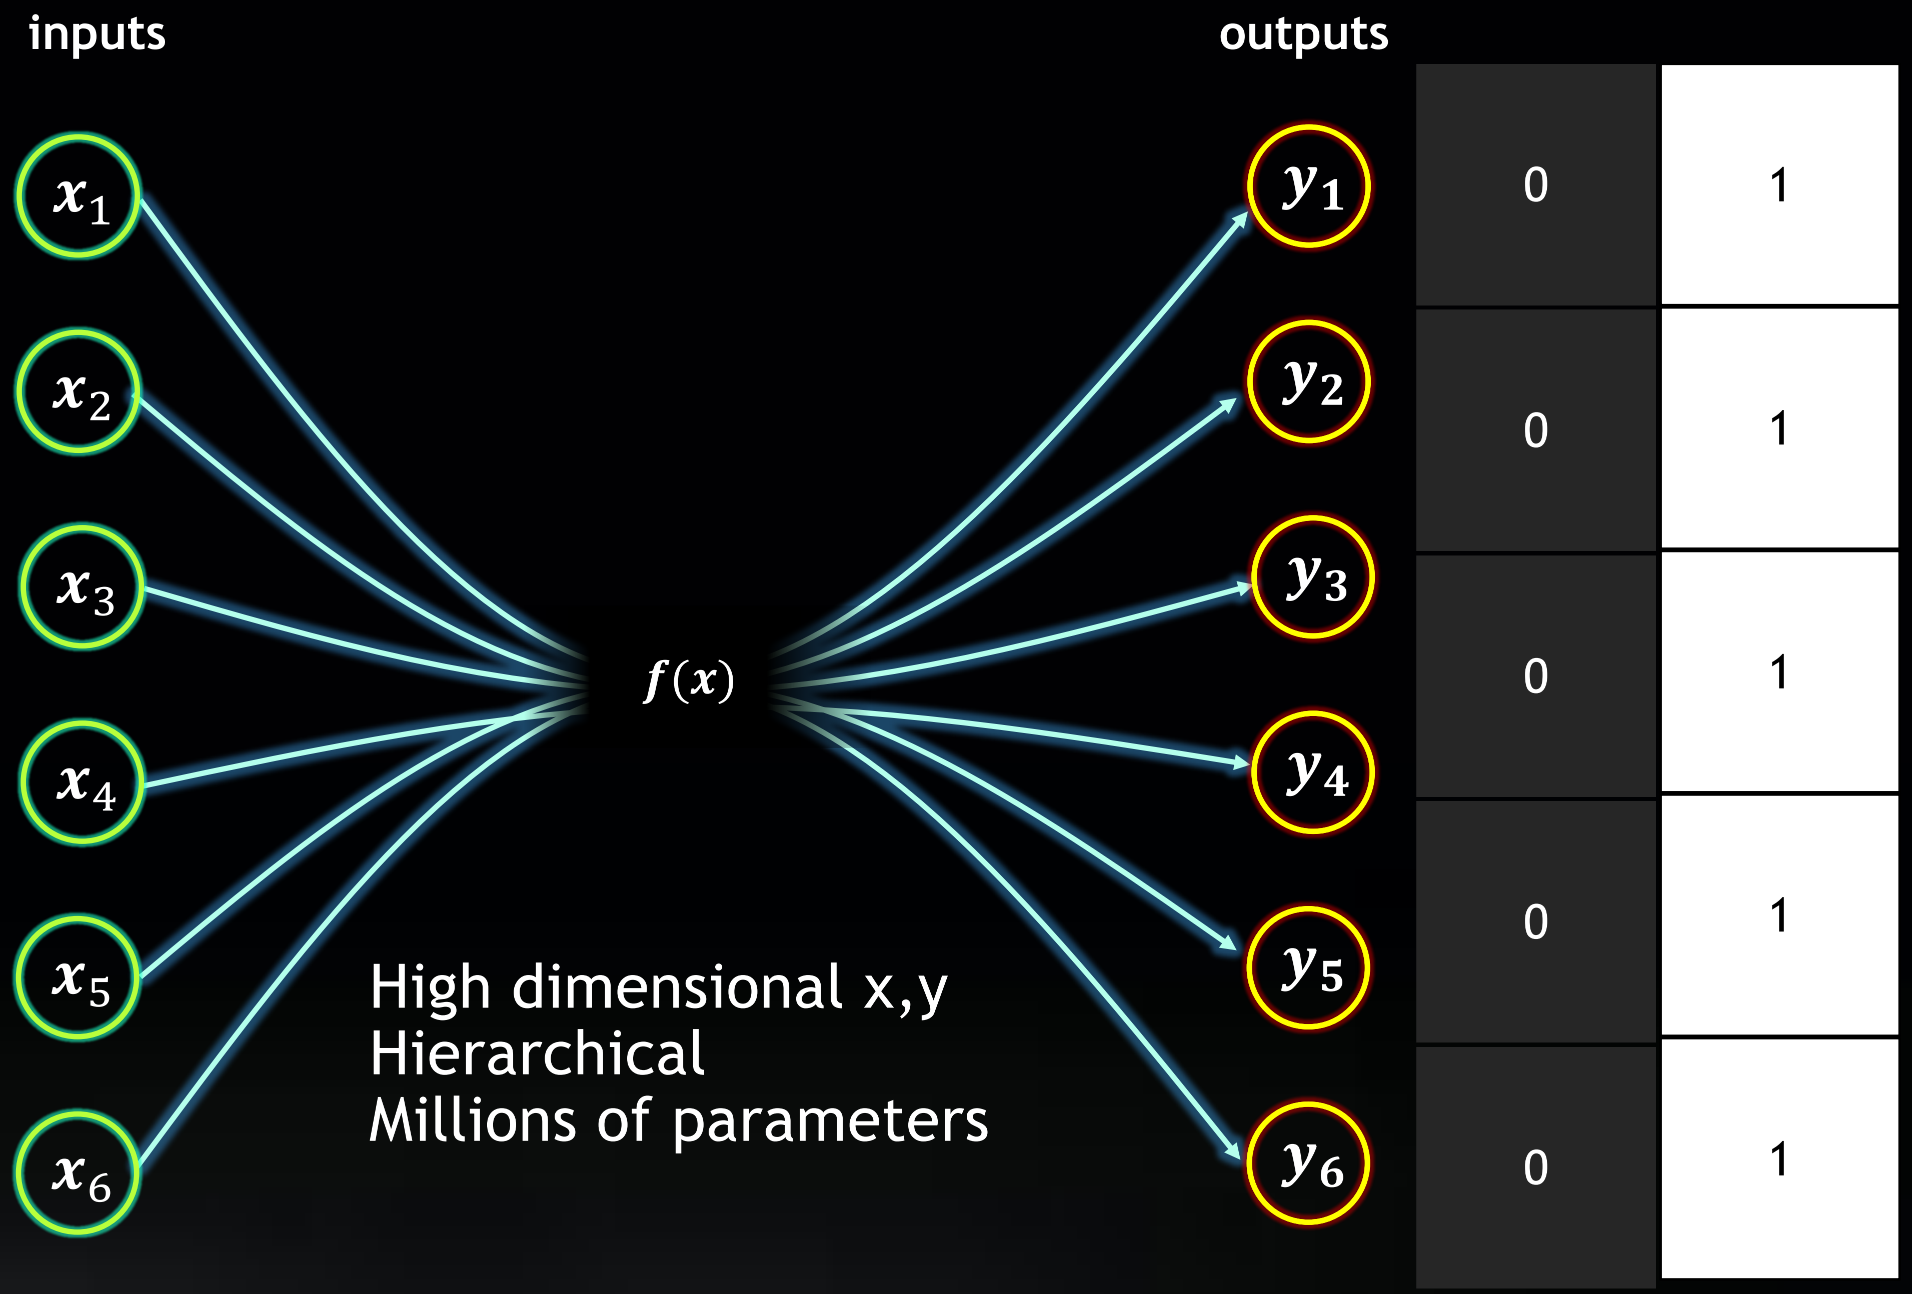

MACHINE LEARNING IS THE NEXT STEP IN SOFTWARE ENGINEERING

TEMP, PRESSURE, MOISTURE

PROBABILITY OF RAIN

HAND-WRITTEN FUNCTION

```
Function1(T,P,Q)
update_mass()
update_momentum()
update_energy()
do_macrophysics()
do_microphysics()
y = get_precipitation()
return y
```

Convert expert knowledge into a function

LEARNED FUNCTION

```
Function1(T,P,Q)
A = relu( w1 * [T,P,Q] + b1)
B = relu( w2 * A       + b2)
C = relu( w3 * B       + b3)
D = relu( w4 * C       + b4)
E = relu( w5 * D       + b5)
y = sigmoid(w6 * E     + b6)
return y
```

Reverse-engineer a function from inputs / outputs

ML CAN DESCRIBE COMPLEX, REAL-WORLD PHENOMENA

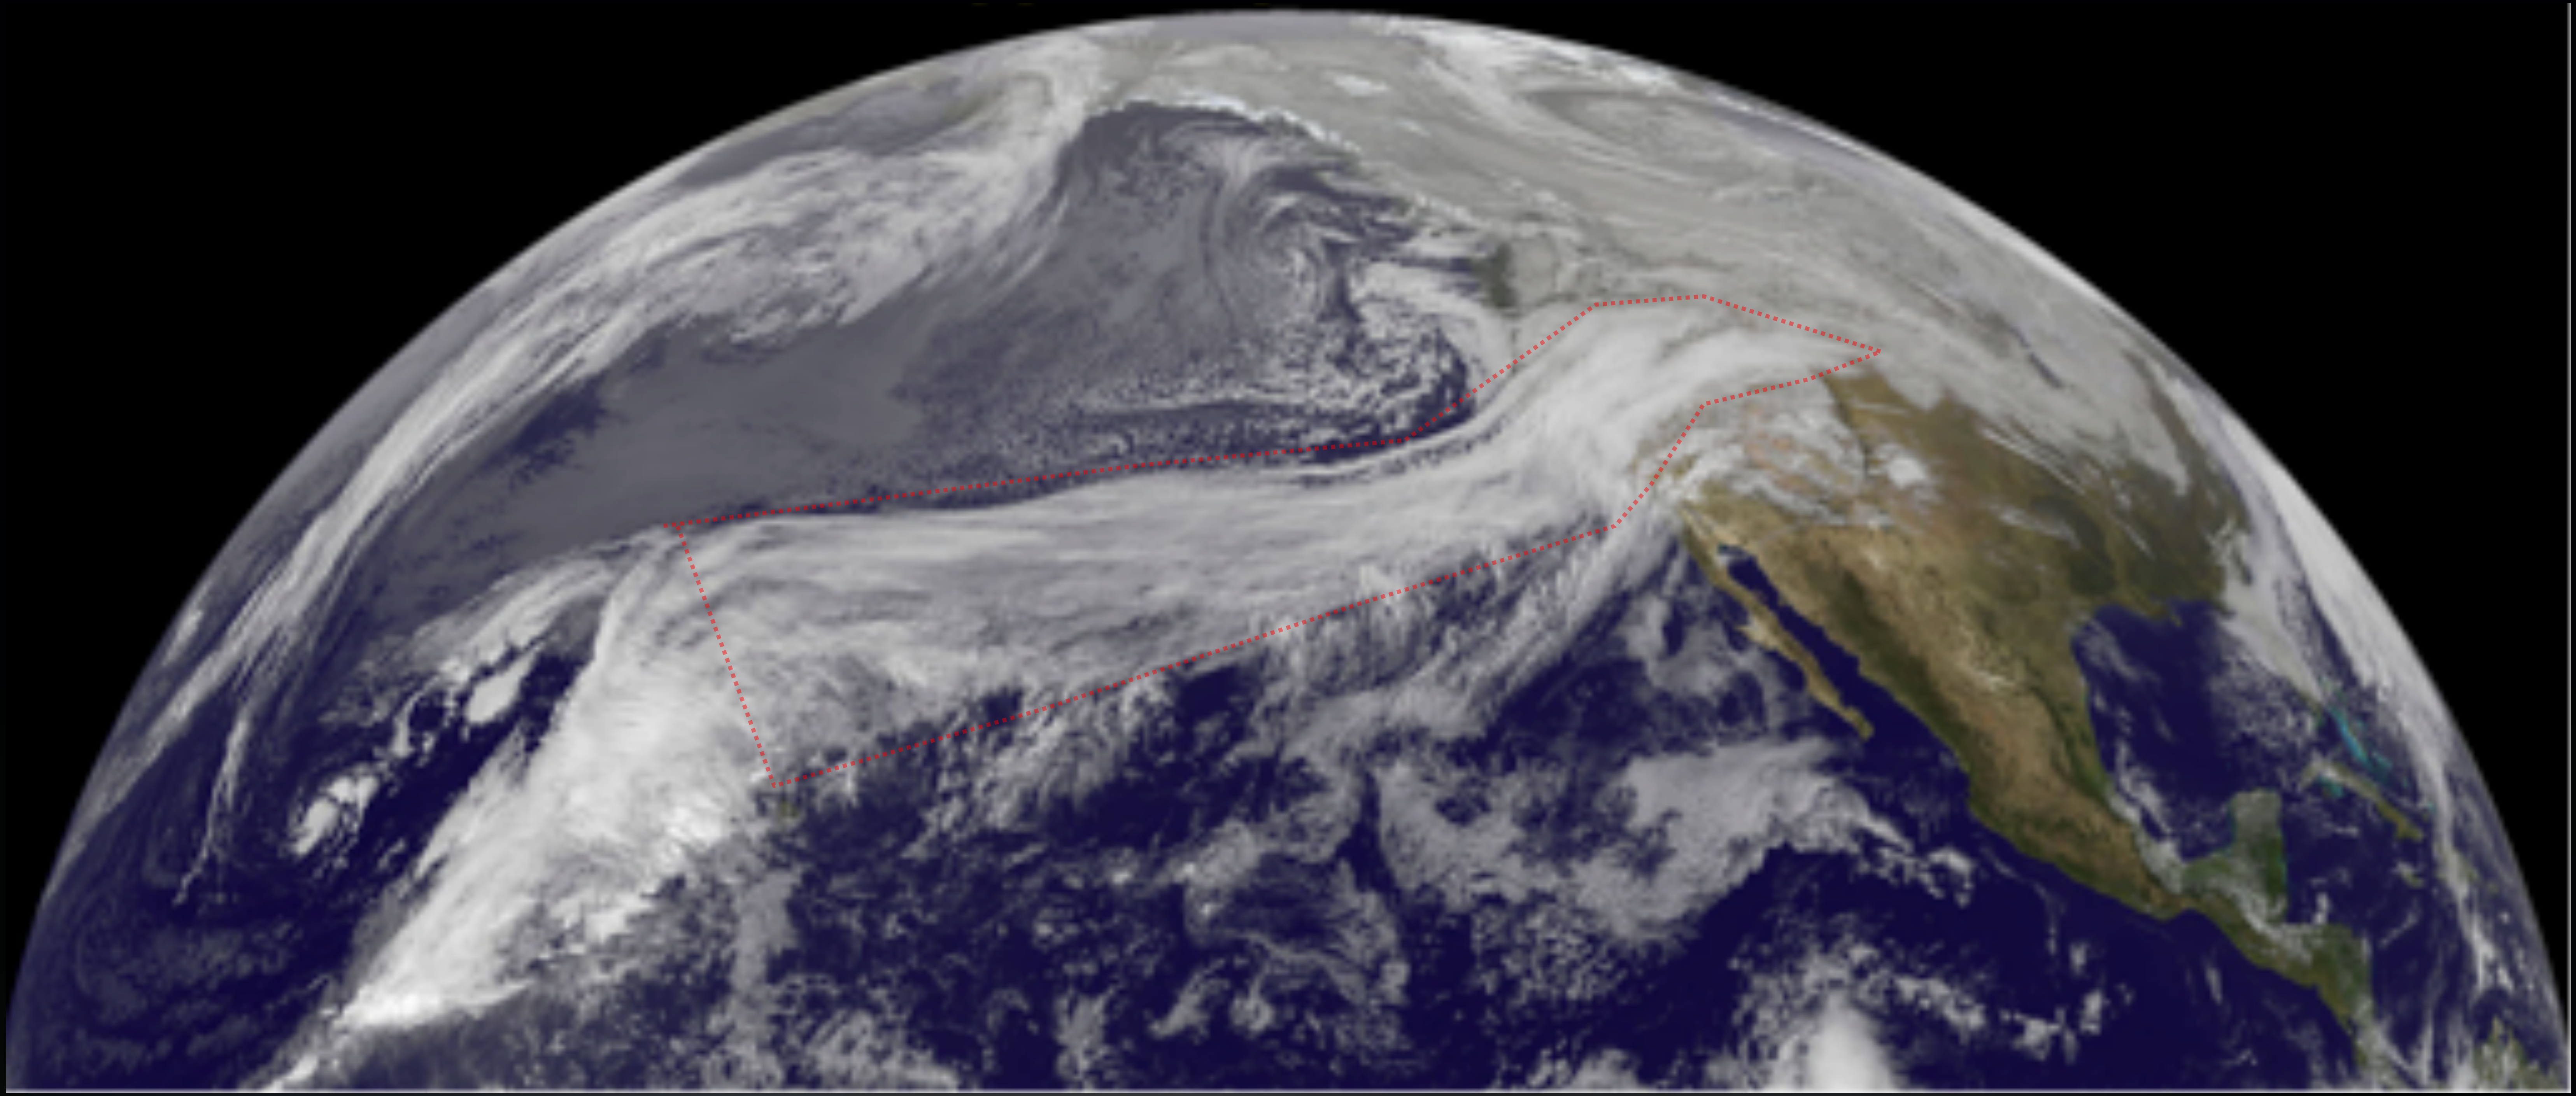

EXAMPLE: ATMOSPHERIC RIVER

ML CAN IMPROVE EXISTING APPLICATIONS

Improve all stages of numerical weather prediction

OBSERVATION THINNING ASSIMILATION DYNAMICS PARAMETRIZATION DISSEMINATION

TOOLS

WHAT YOU NEED TO MAKE IT WORK

LARGE QUANTITIES OF DATA

ML FRAMEWORK

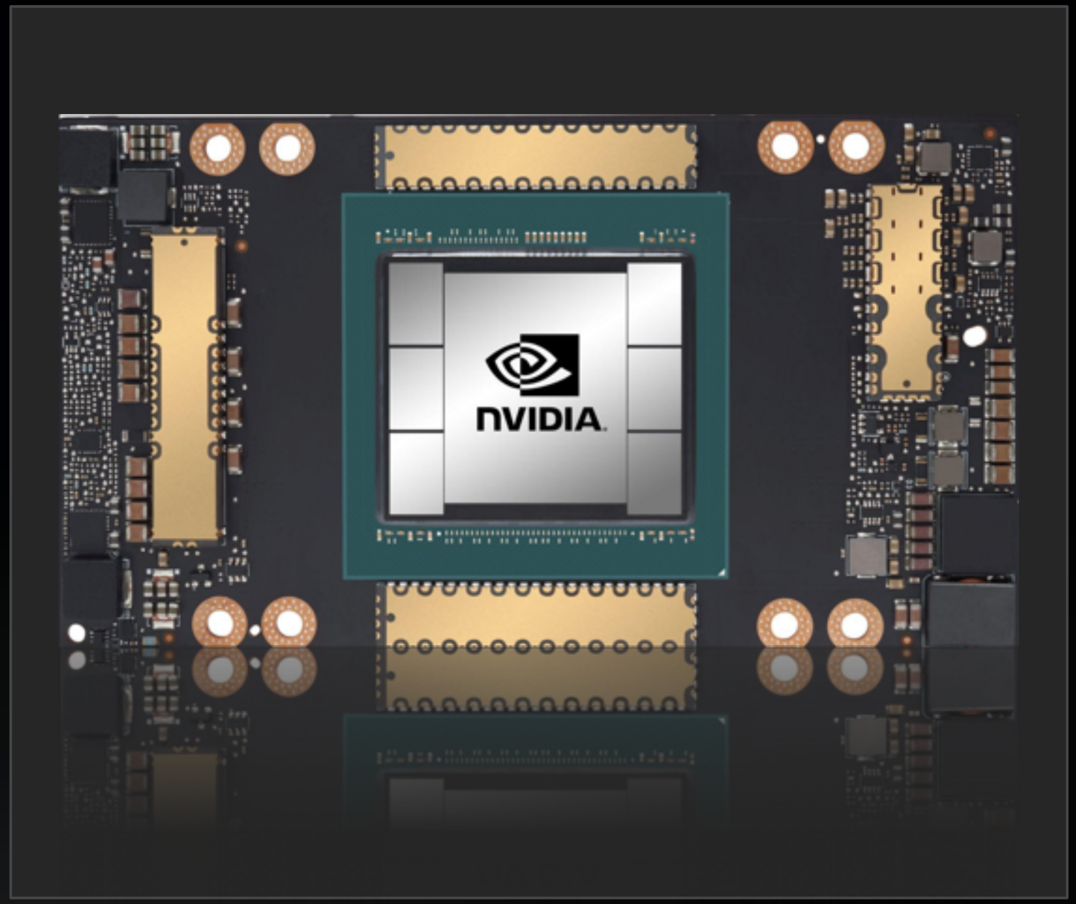

GPU ACCELERATOR

DEEP LEARNING FRAMEWORK

PyTorch

K + TensorFlow

Python

mxnet

GLUON

C++

flux

Julia

Chainer mxnet Caffe2 Microsoft CNTK TensorFlow Keras GLUON PYTORCH theano

GPUs and Machine Learning

The Imagenet competition: Automatically classify images from 1000 different categories

GPUS MAKE MACHINE LEARNING PRACTICAL

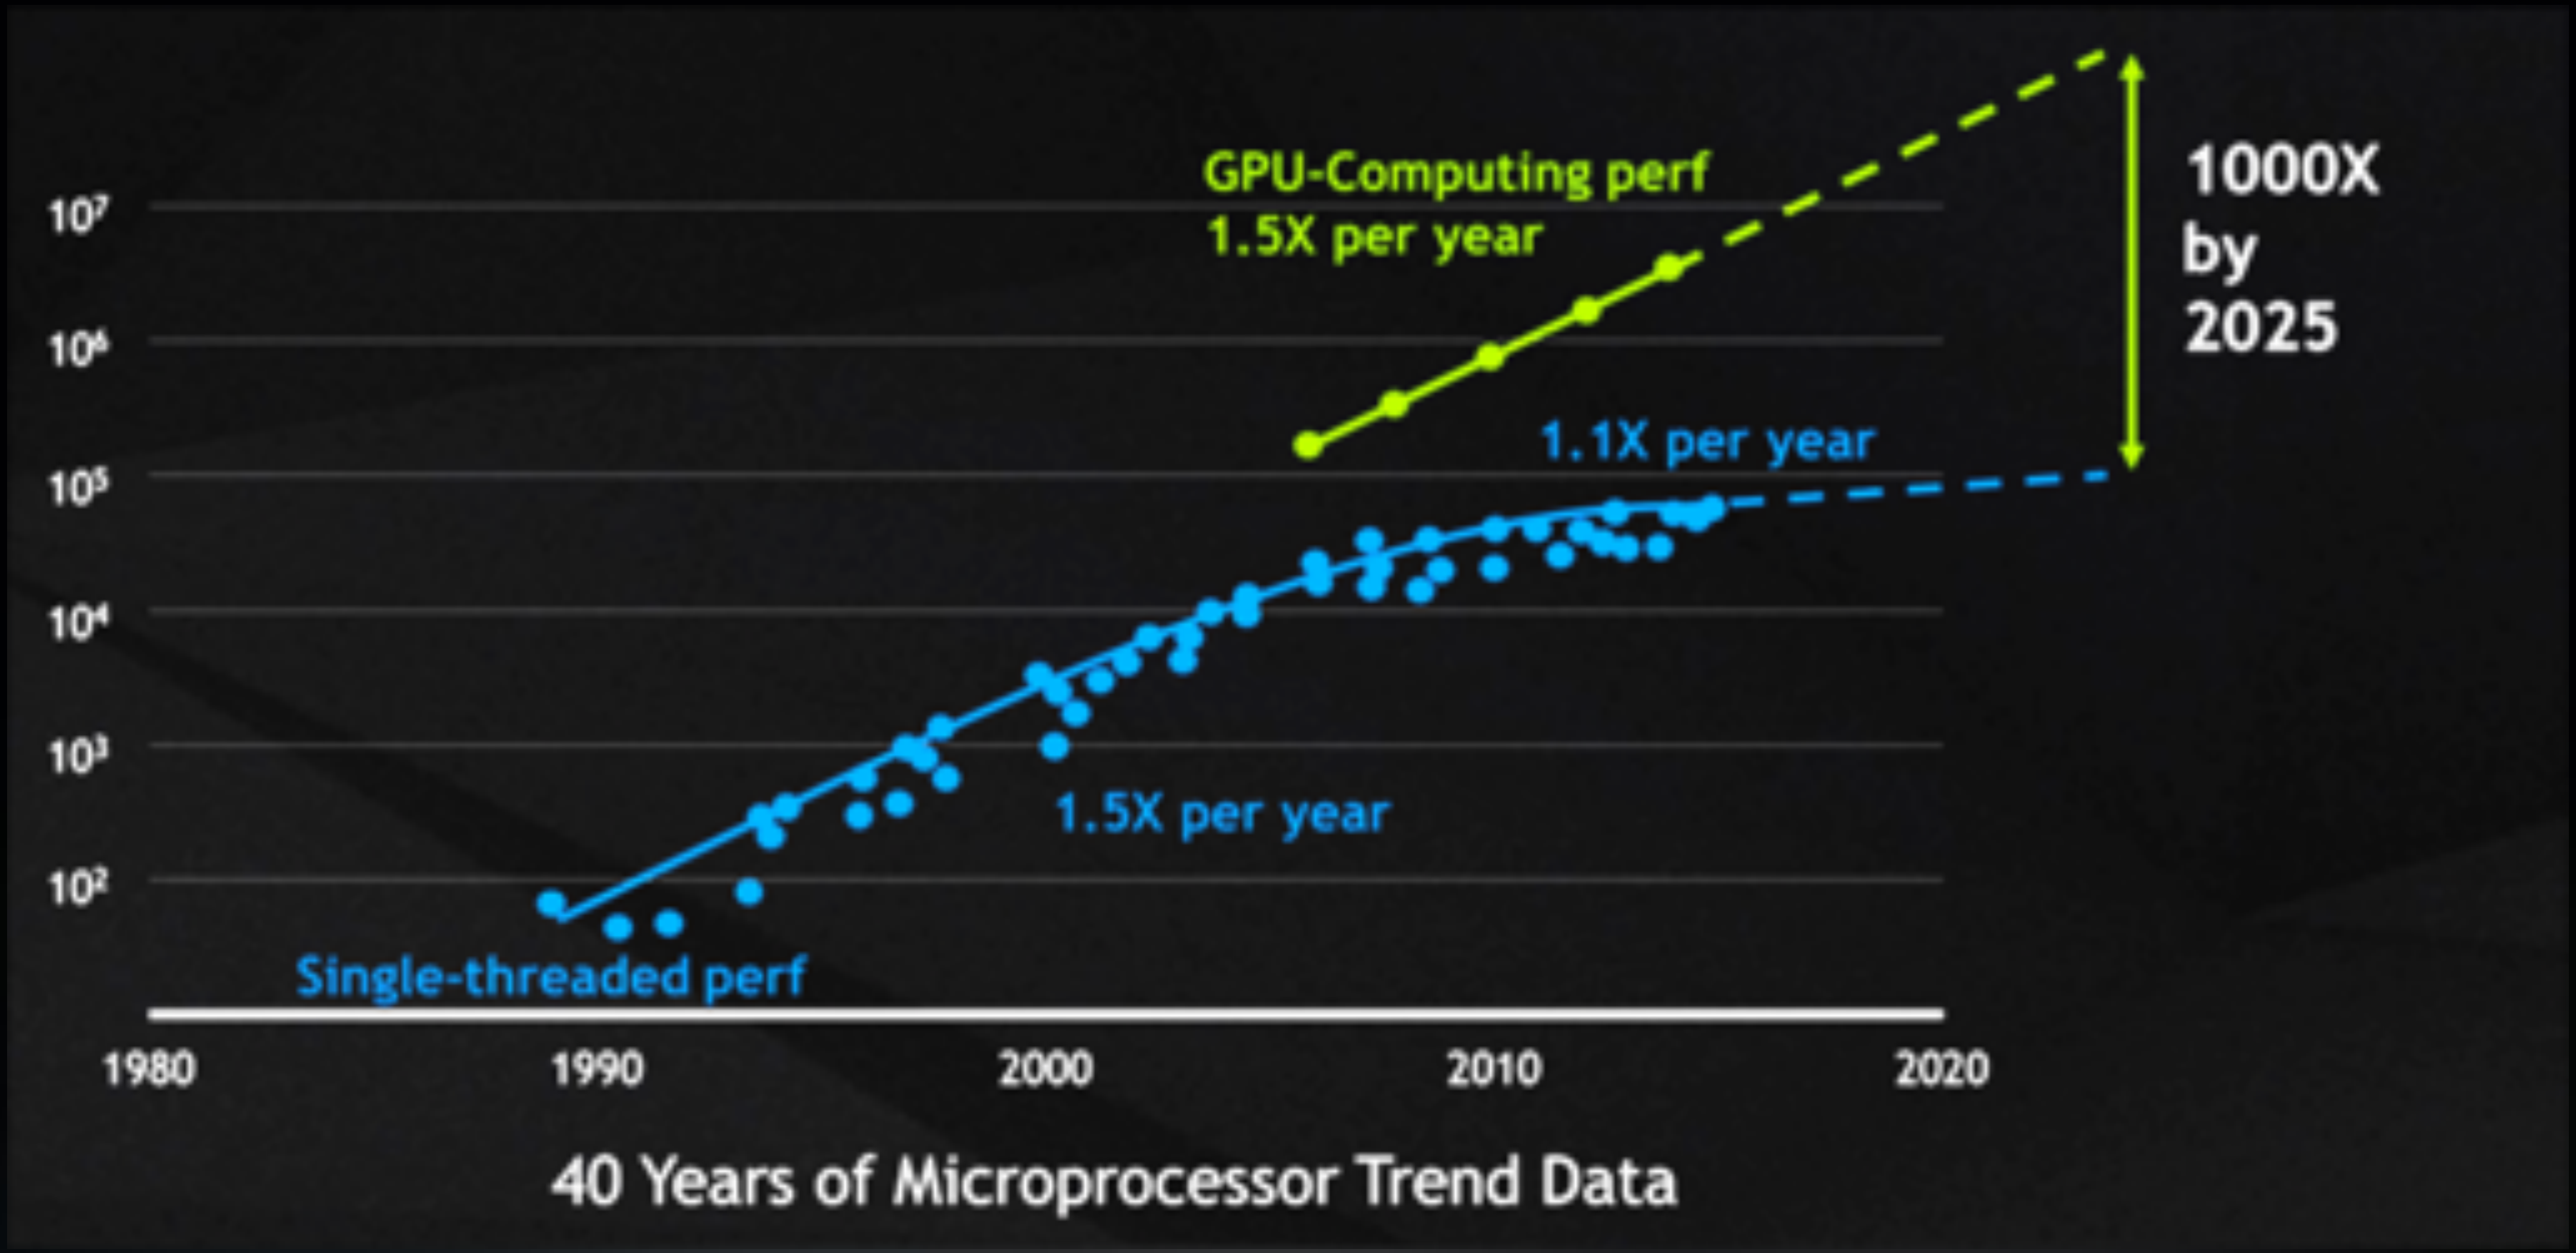

LEARNED FUNCTIONS ARE GPU ACCELERATED

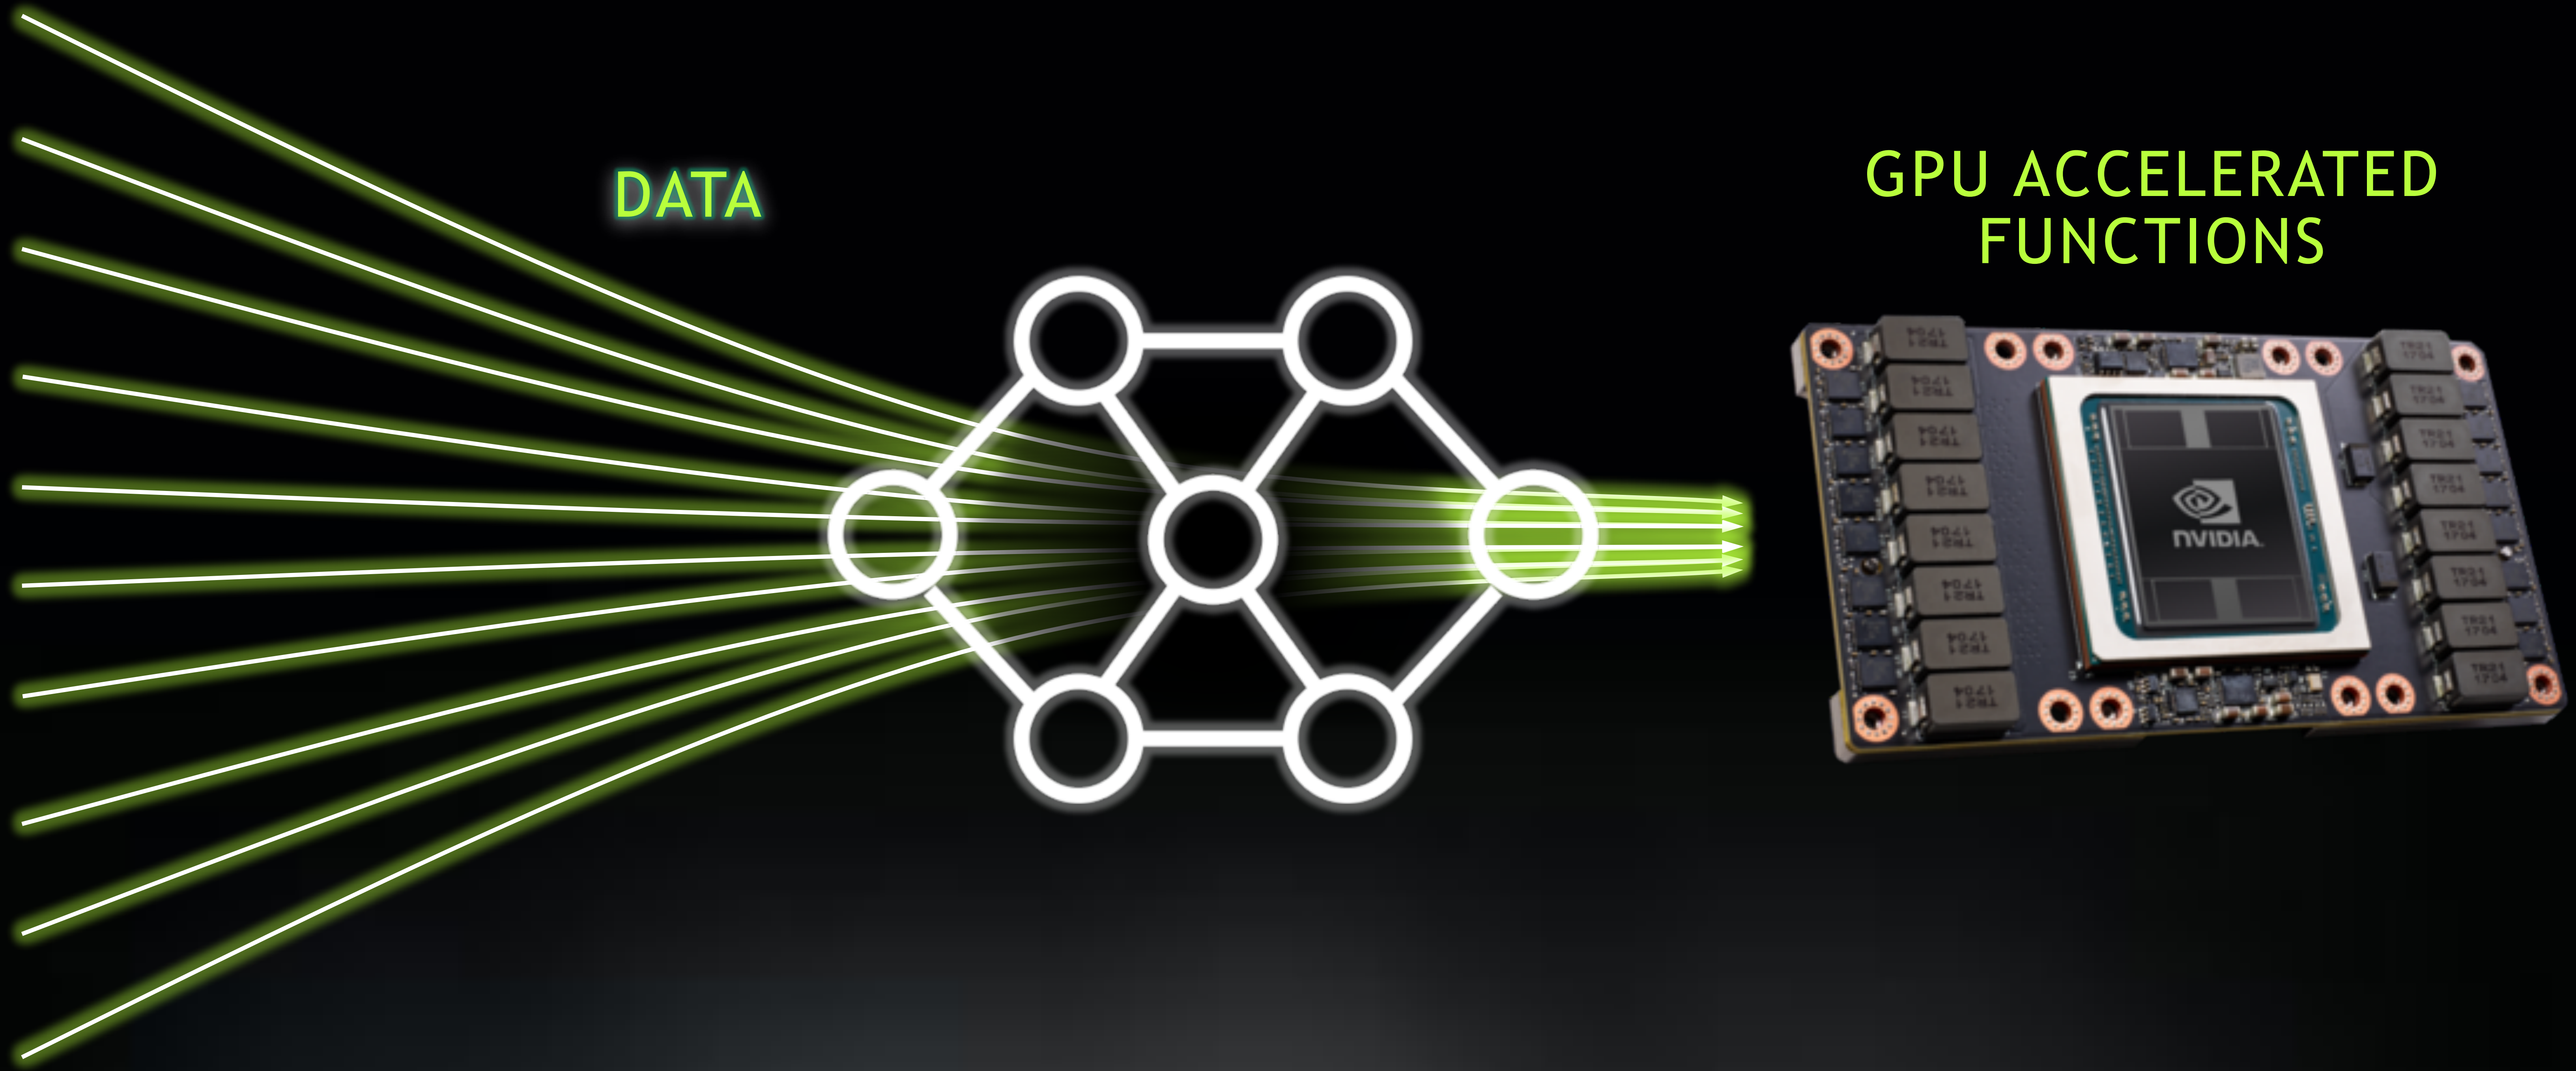

APPLICATIONS

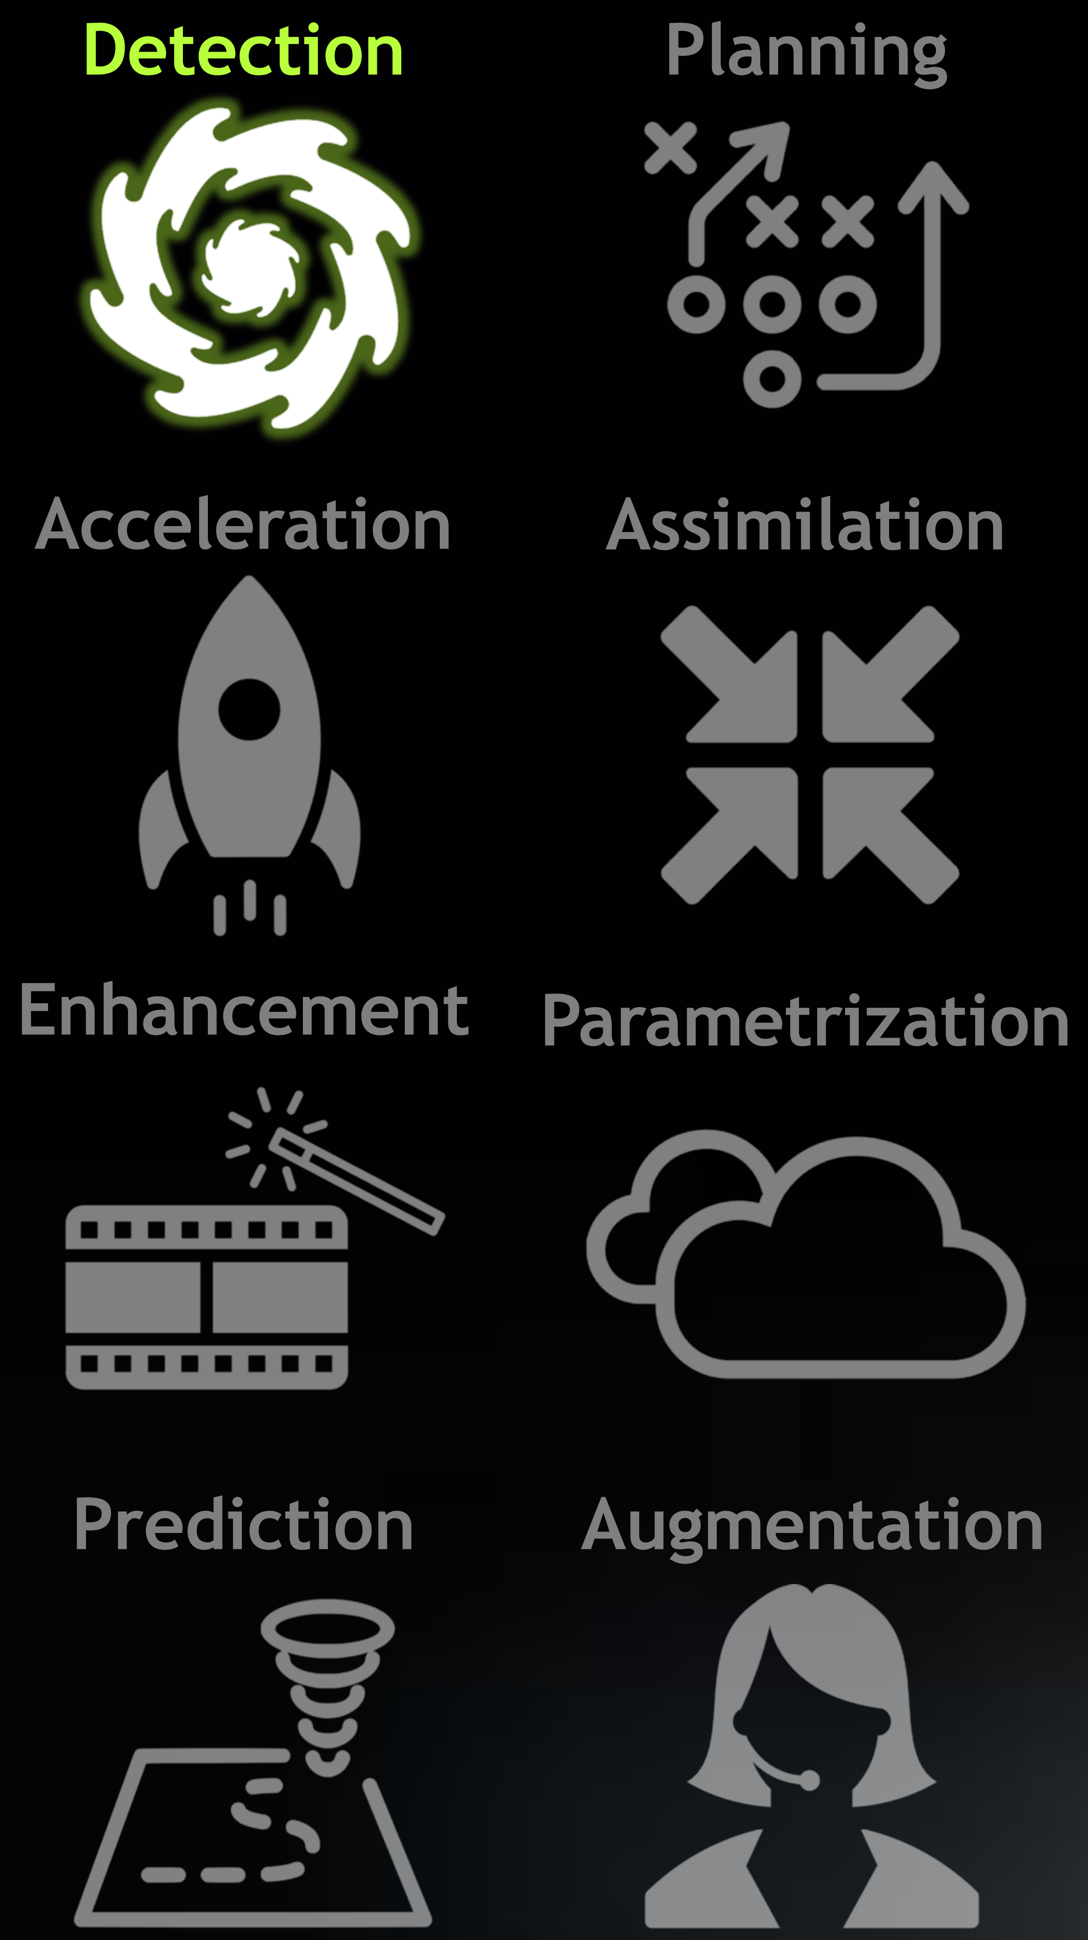

Feature Detection

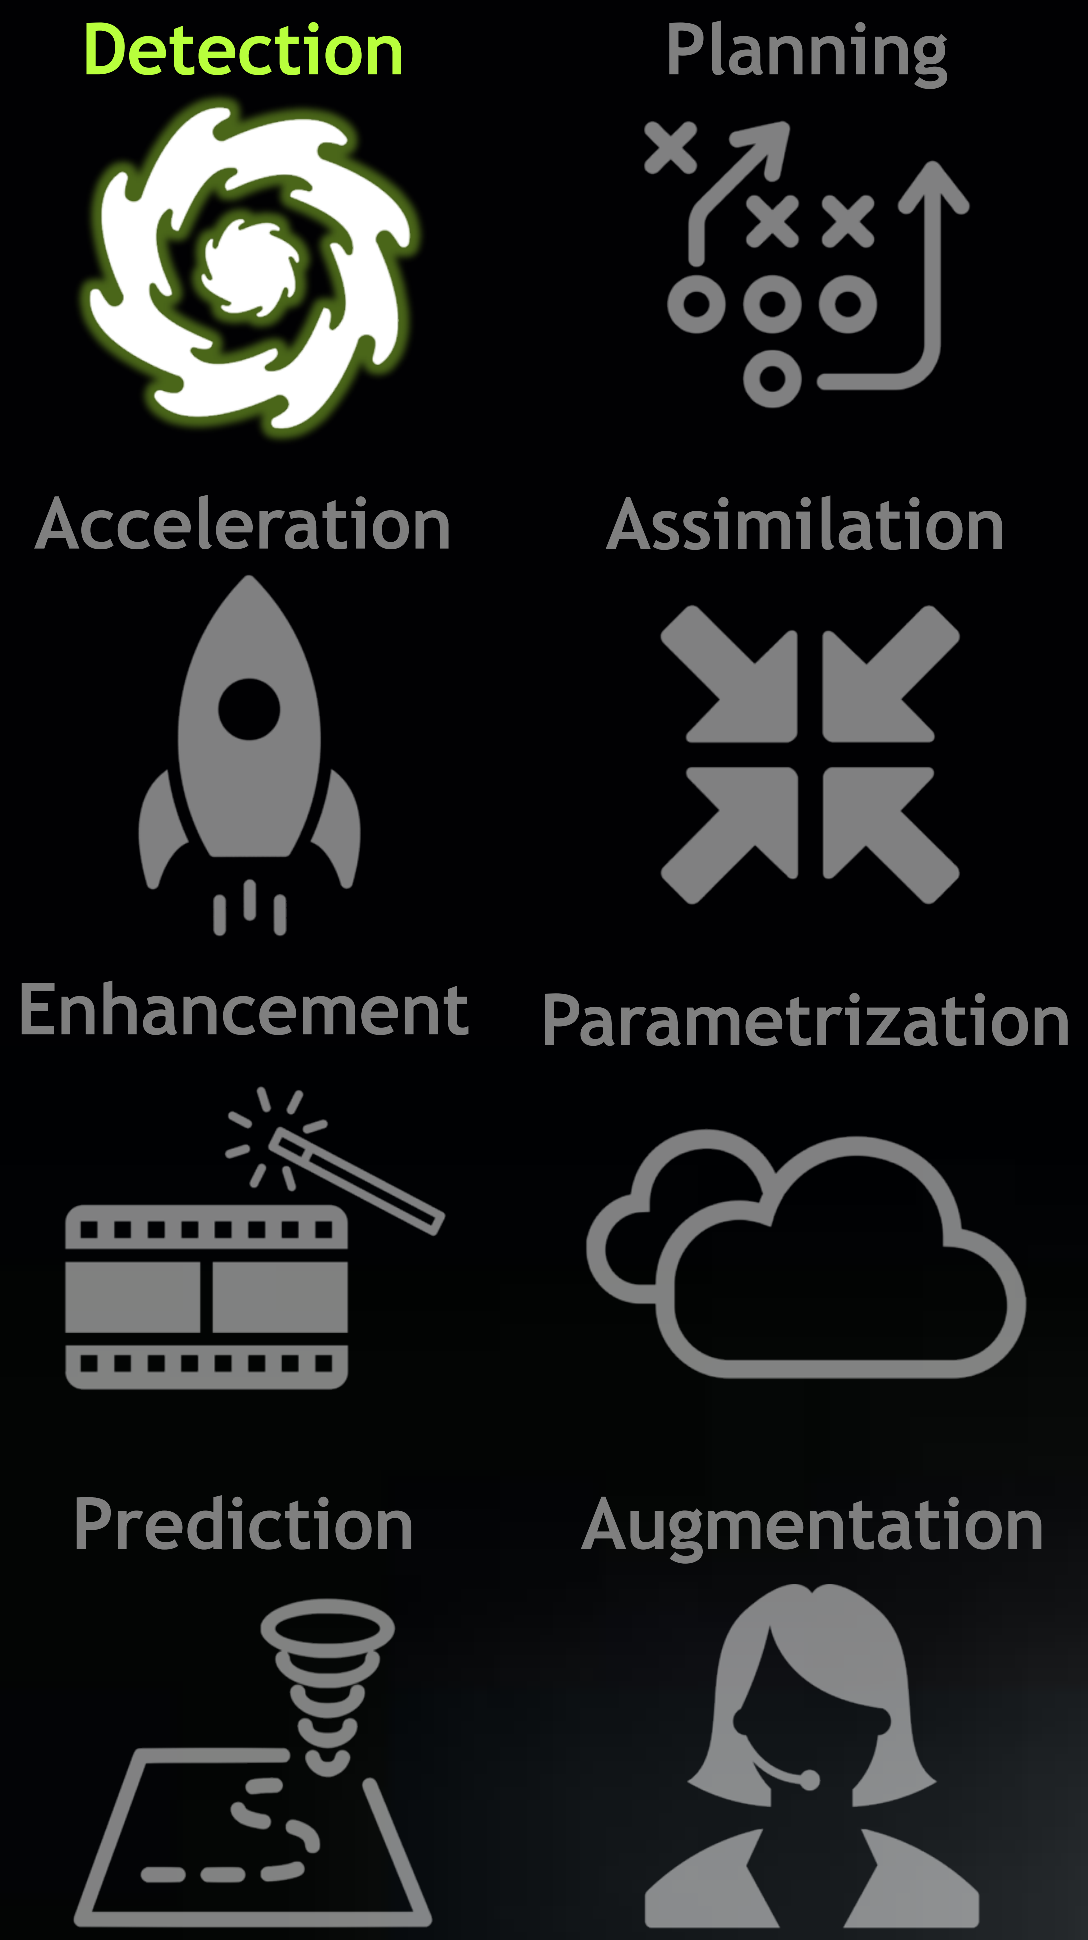

Track Extreme Weather in Real Time

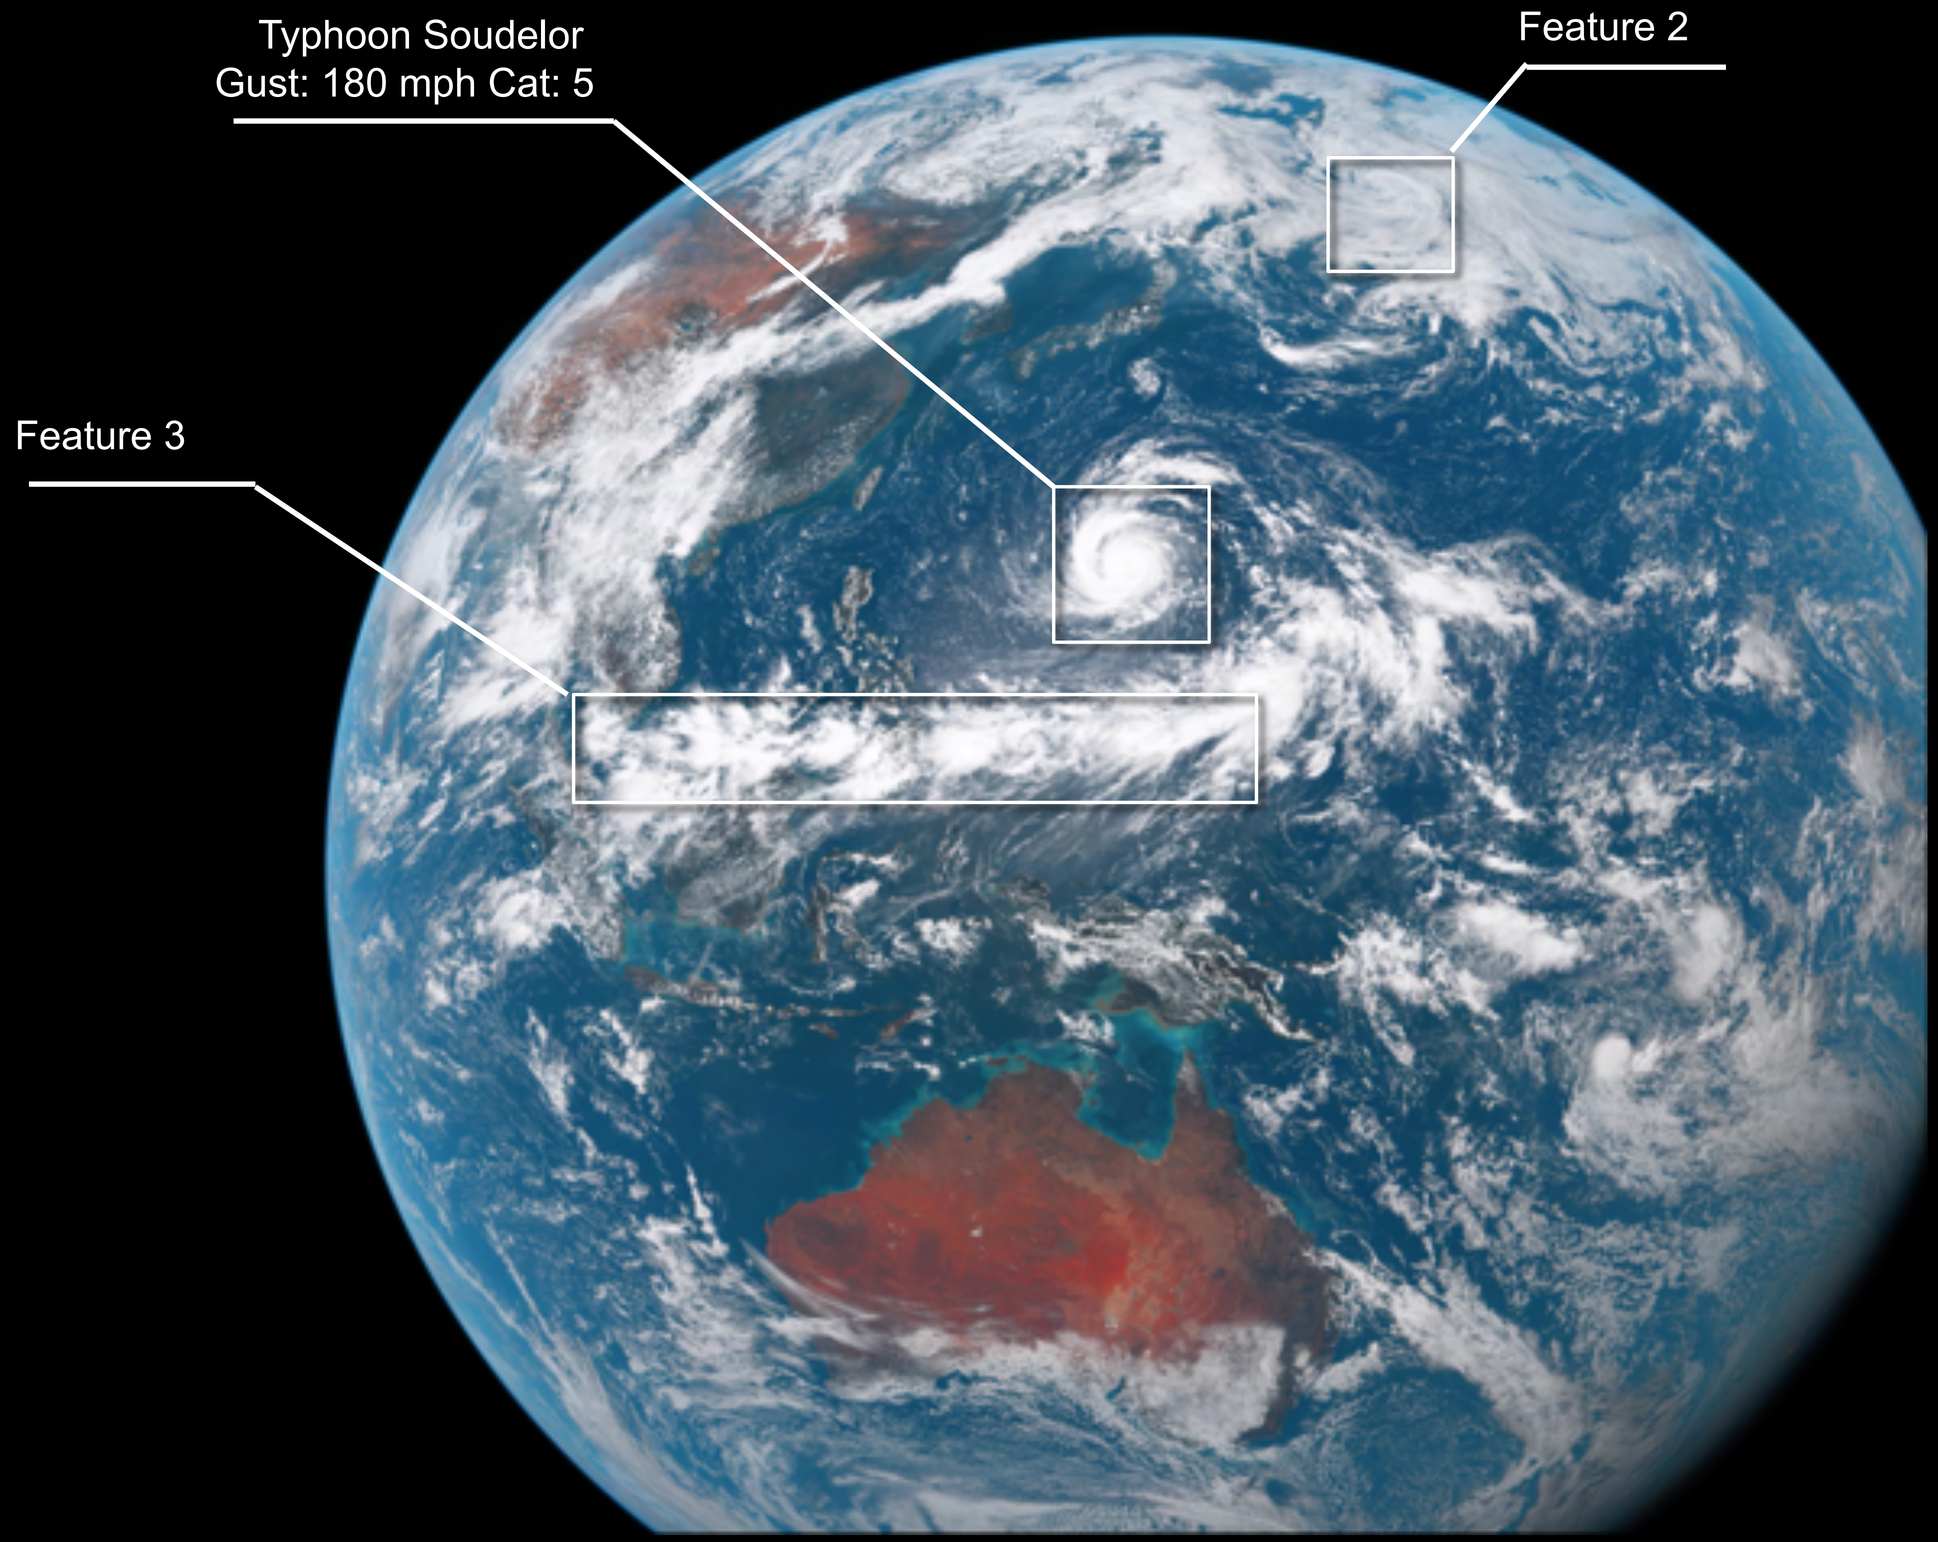

Exascale deep learning for climate analytics. In Proceedings of the International Conference for High Performance Computing, Networking, Storage, and Analysis(SC '18). IEEE Press, Article 51, 1-12. DOI:https://doi.org/10.1109/SC.2018.00054

Detection

Planning

Acceleration

Assimilation

Enhancement

Parametrization

Prediction

Augmentation

Monitor Environmental Change

Helber, Patrick, et al. "Eurosat: A novel dataset and deep learning benchmark for land use and land cover classification." *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing* 12.7 (2019): 2217-2226.

Detection

Planning

Acceleration

Assimilation

Enhancement

Parametrization

Prediction

Augmentation

Strategy and Planning

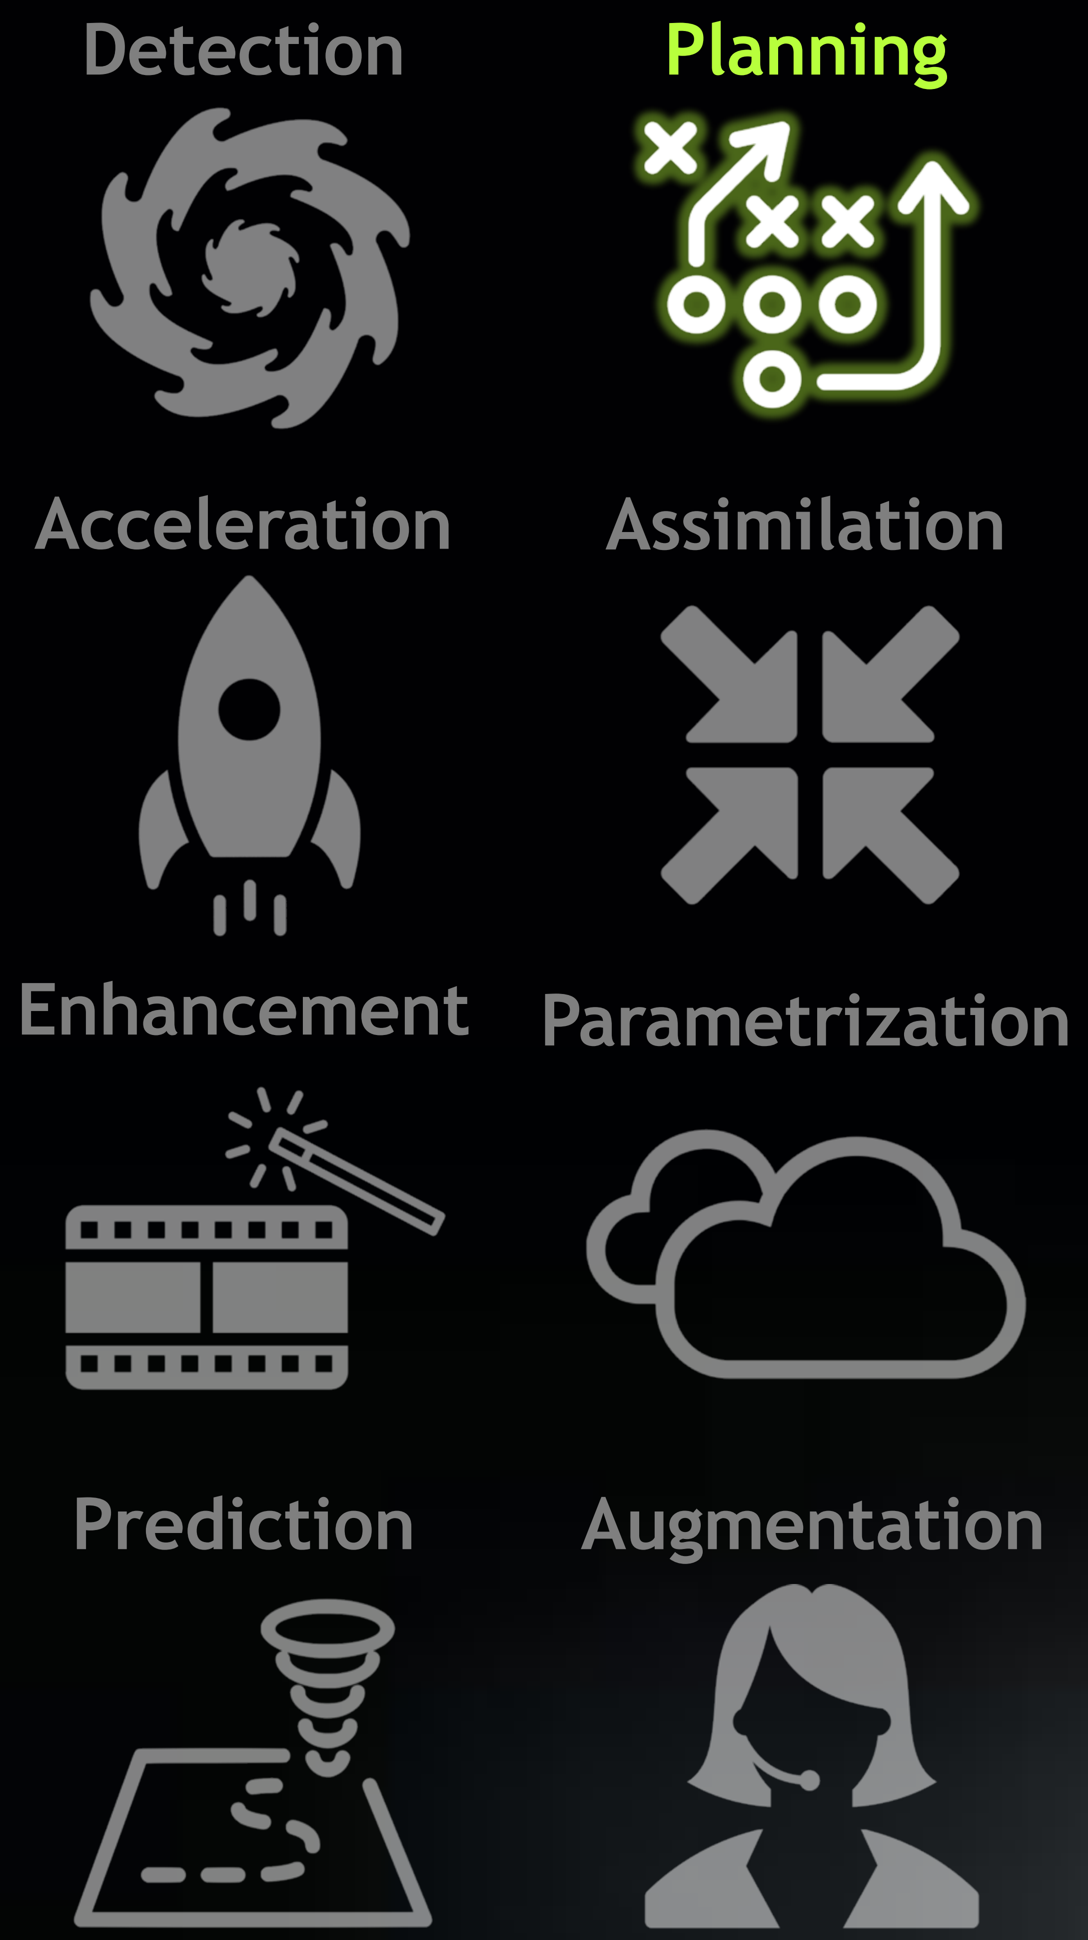

Optimize Disaster Planning

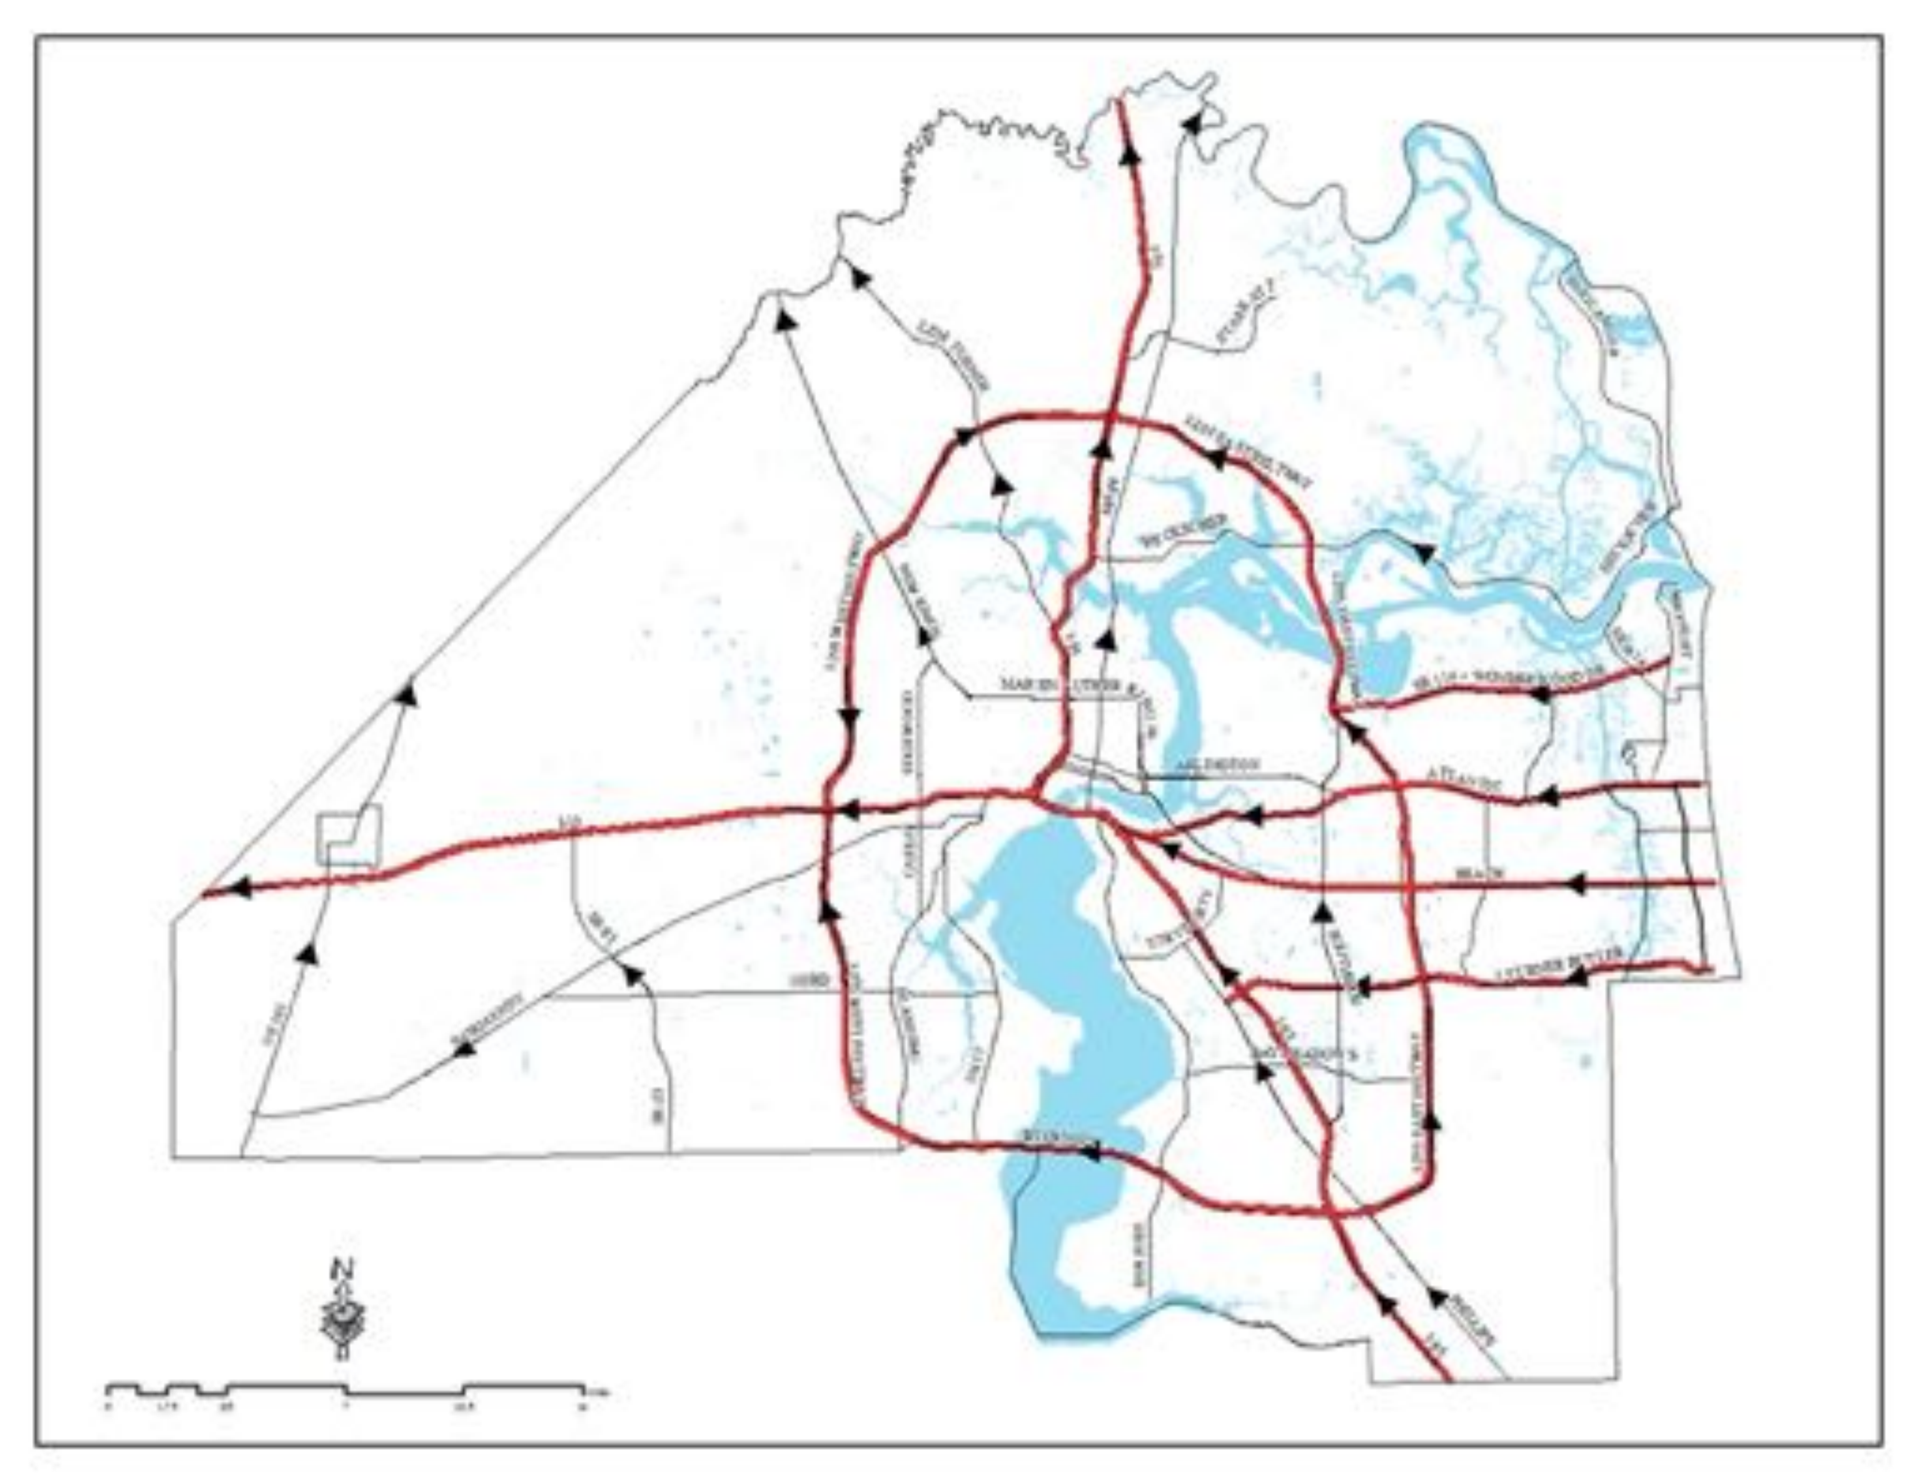

J. Sharma, P. Andersen, O. Granmo and M. Goodwin, "Deep Q-Learning With Q-Matrix Transfer Learning for Novel Fire Evacuation Environment," in *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, doi: 10.1109/TSMC.2020.2967936.

Detection

Planning

Acceleration

Assimilation

Enhancement

Parametrization

Prediction

Augmentation

Autonomous Vehicles and Drones

E. T. Steimle, R. R. Murphy, M. Lindemuth and M. L. Hall, "Unmanned marine vehicle use at Hurricanes Wilma and Ike," *OCEANS 2009*, Biloxi, MS, 2009, pp. 1-6, doi: 10.23919/OCEANS.2009.5422201.

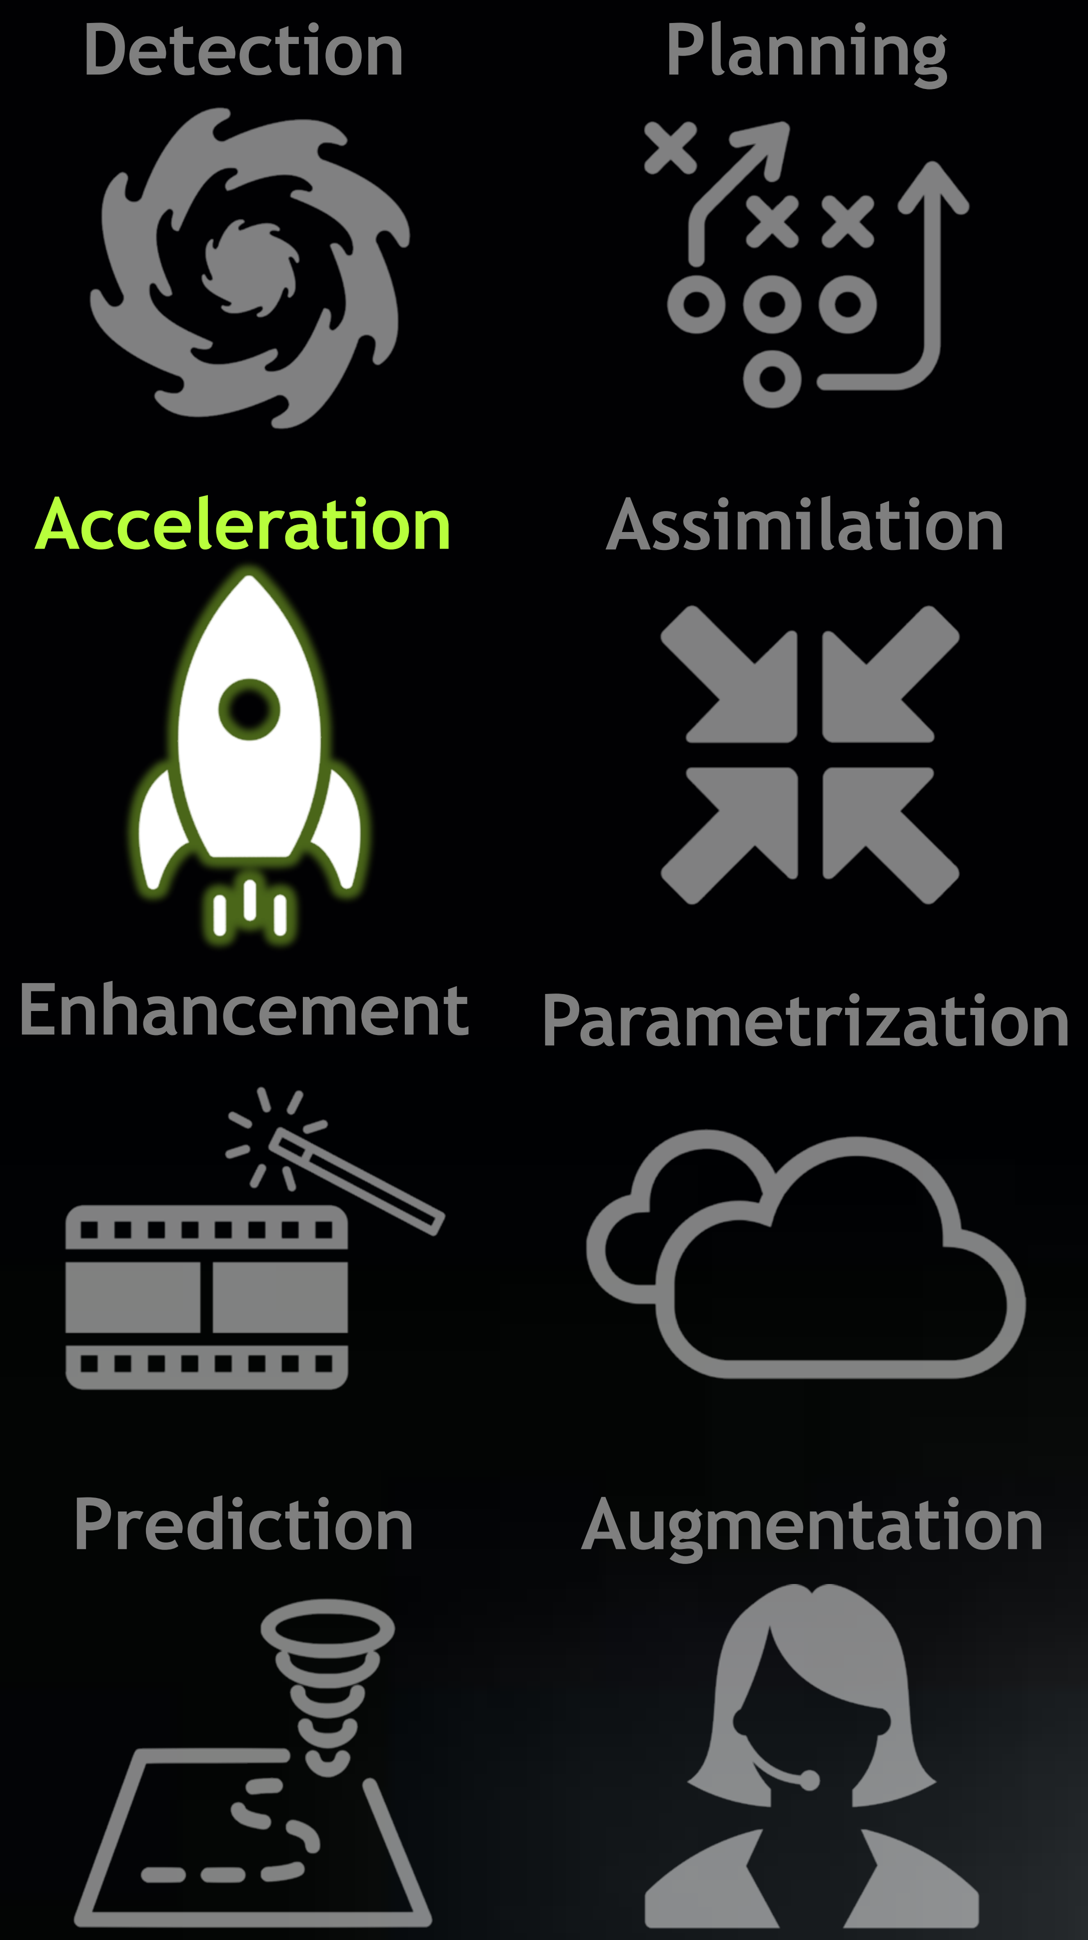

Accelerate Expensive Models

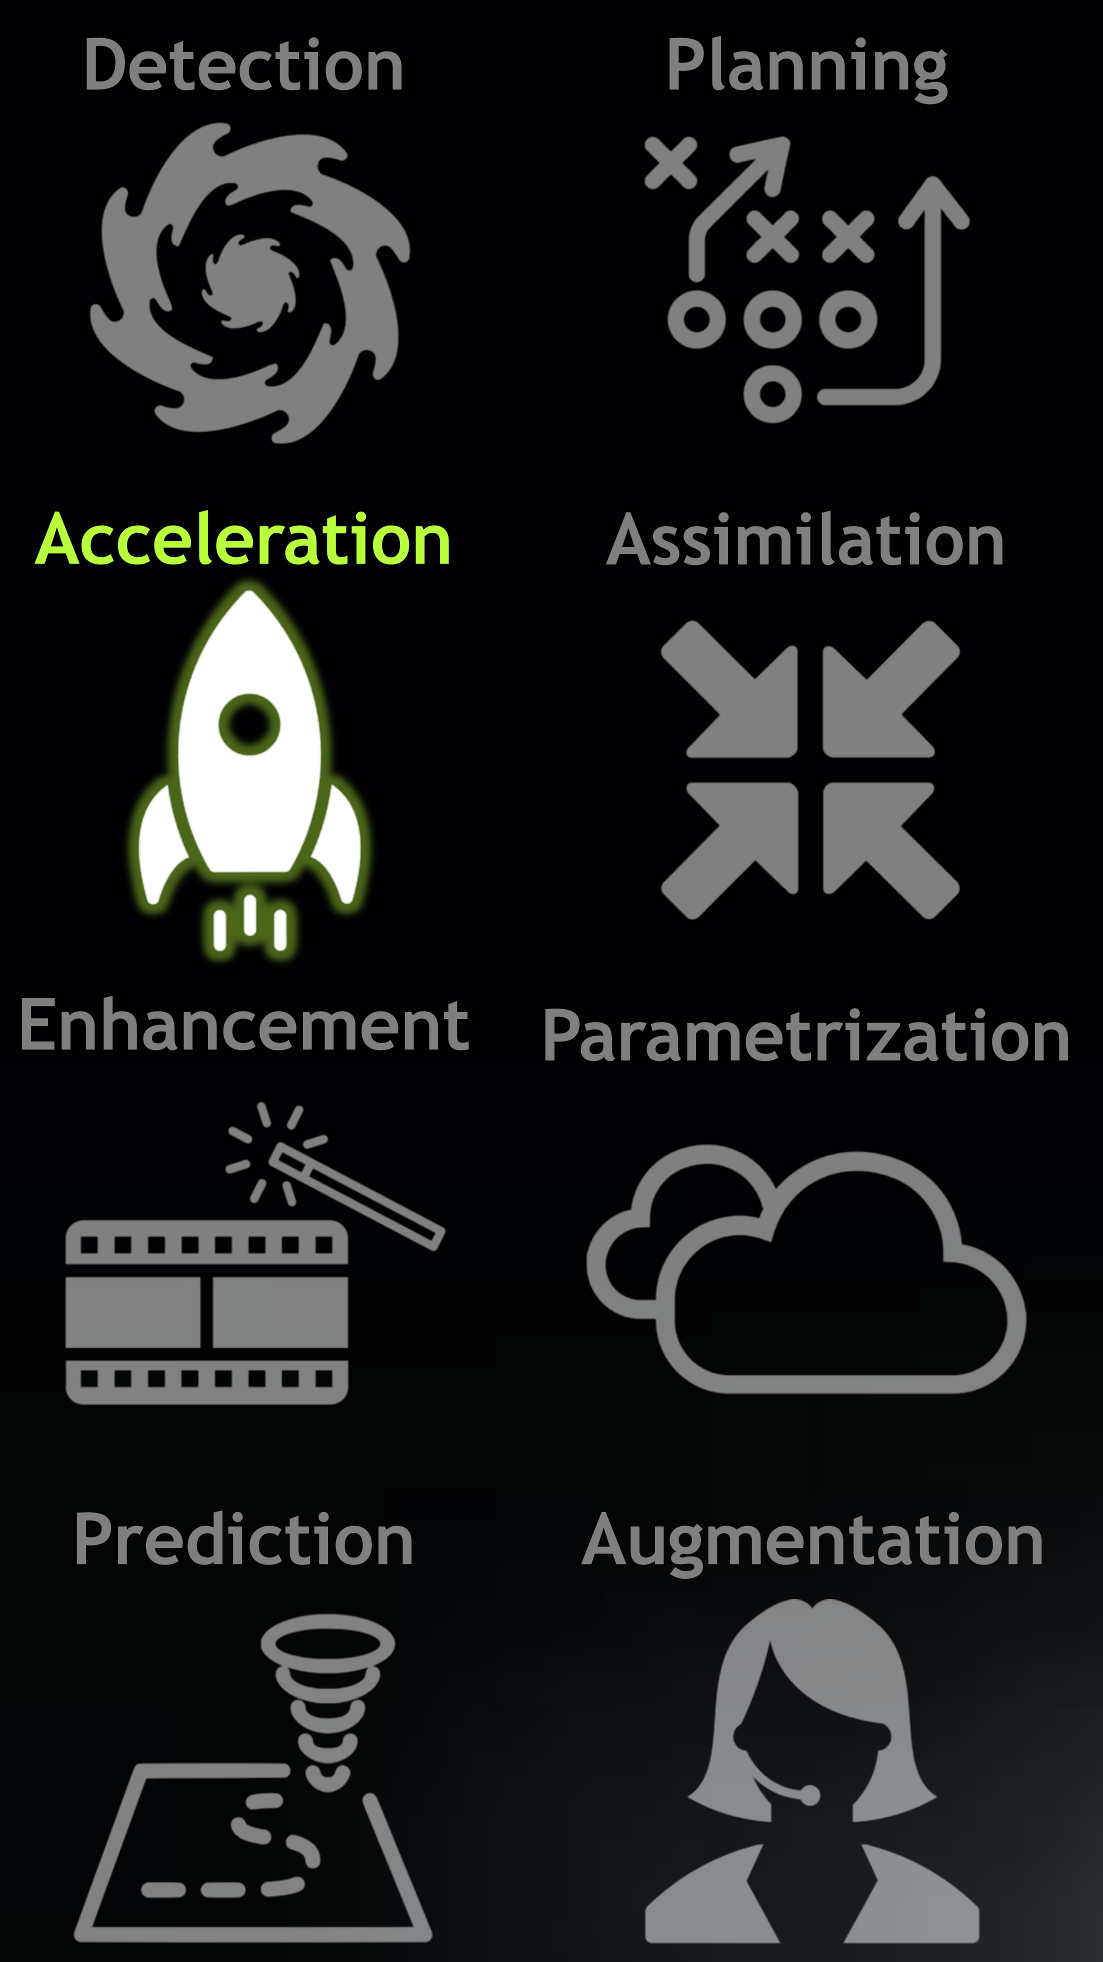

Emulation: AI powered approximation

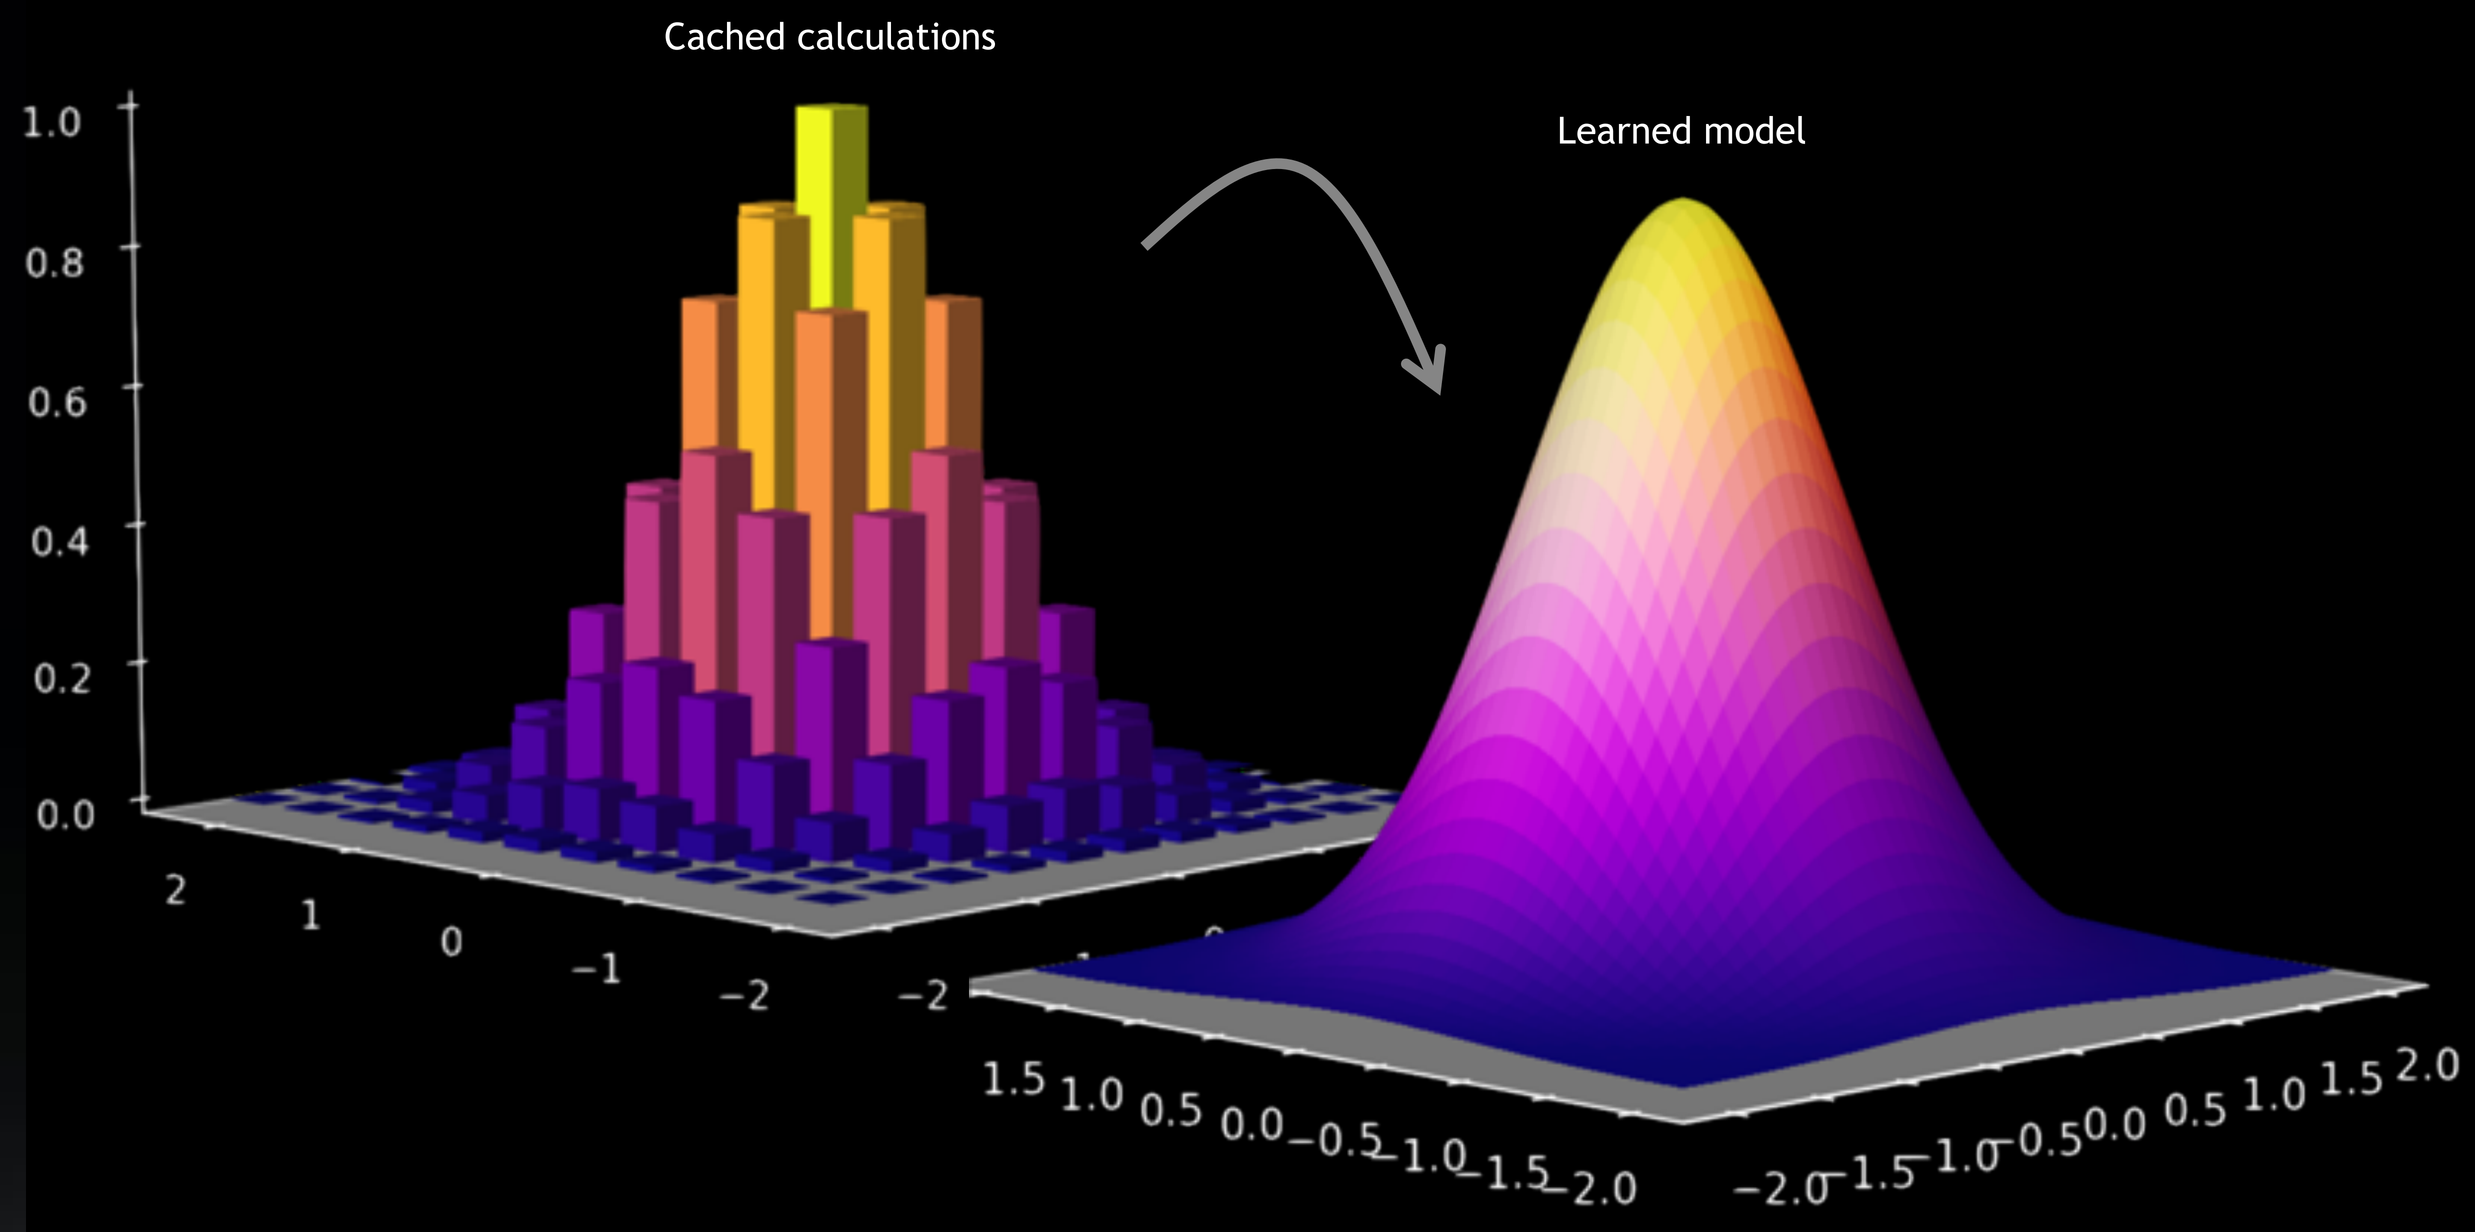

https://arxiv.org/pdf/1811.06533.pdf

Detection

Planning

Acceleration

Assimilation

Enhancement

Parametrization

Prediction

Augmentation

Accelerate Data Assimilation via Emulation

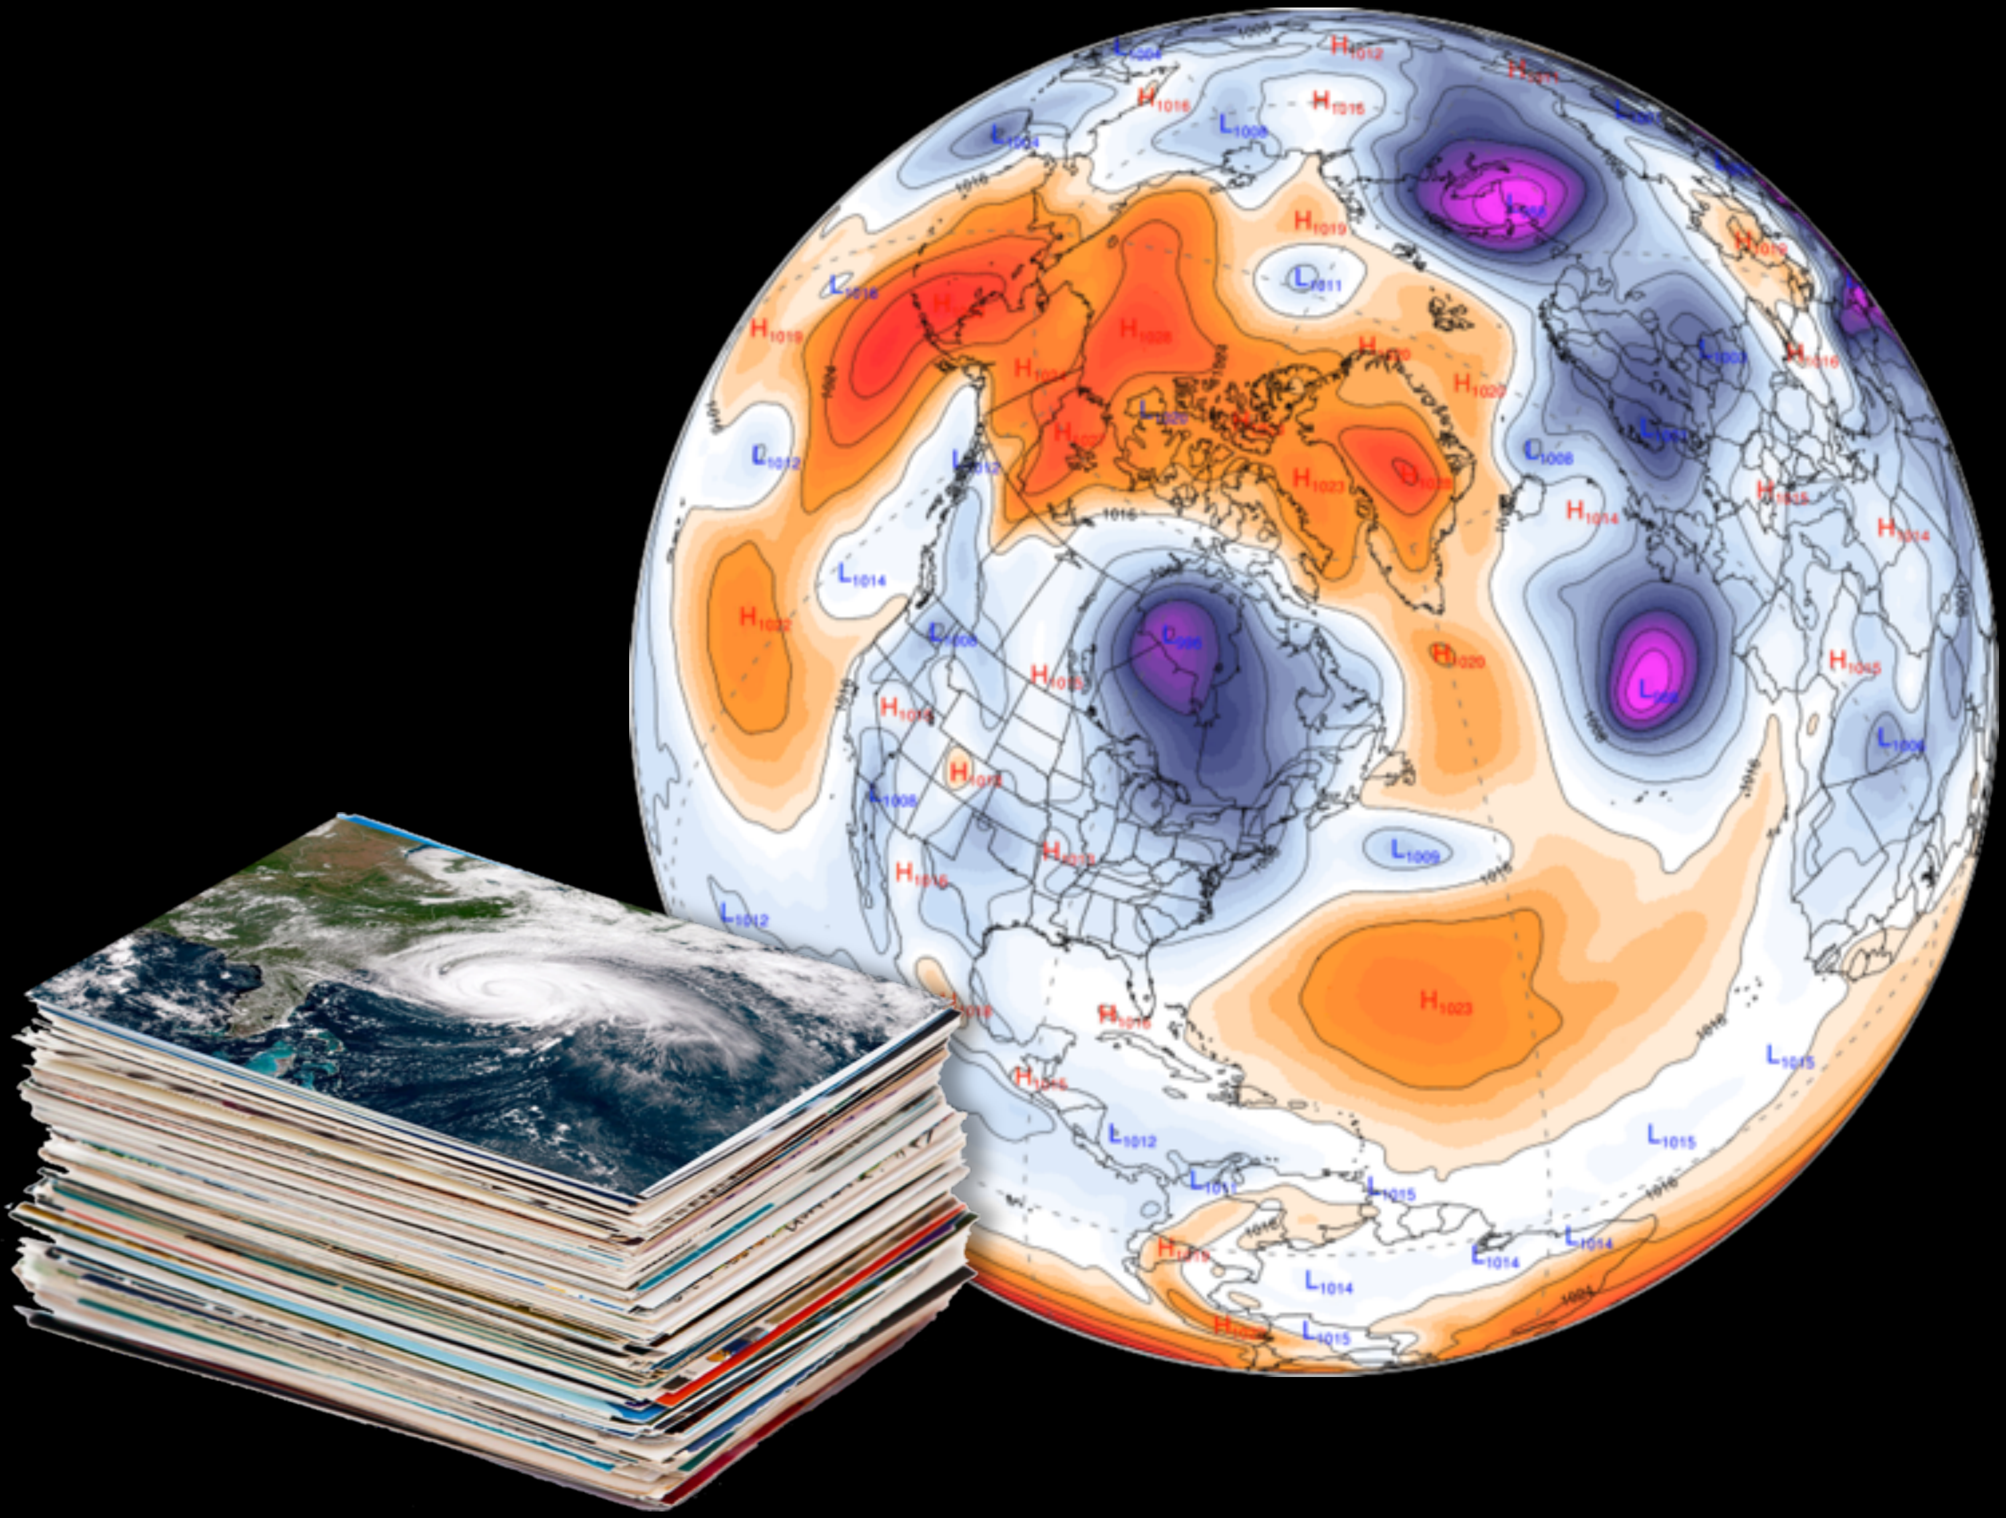

Dueben, Hogan, Bauer @ECMWF and Progsch, Angerer @NVIDIA

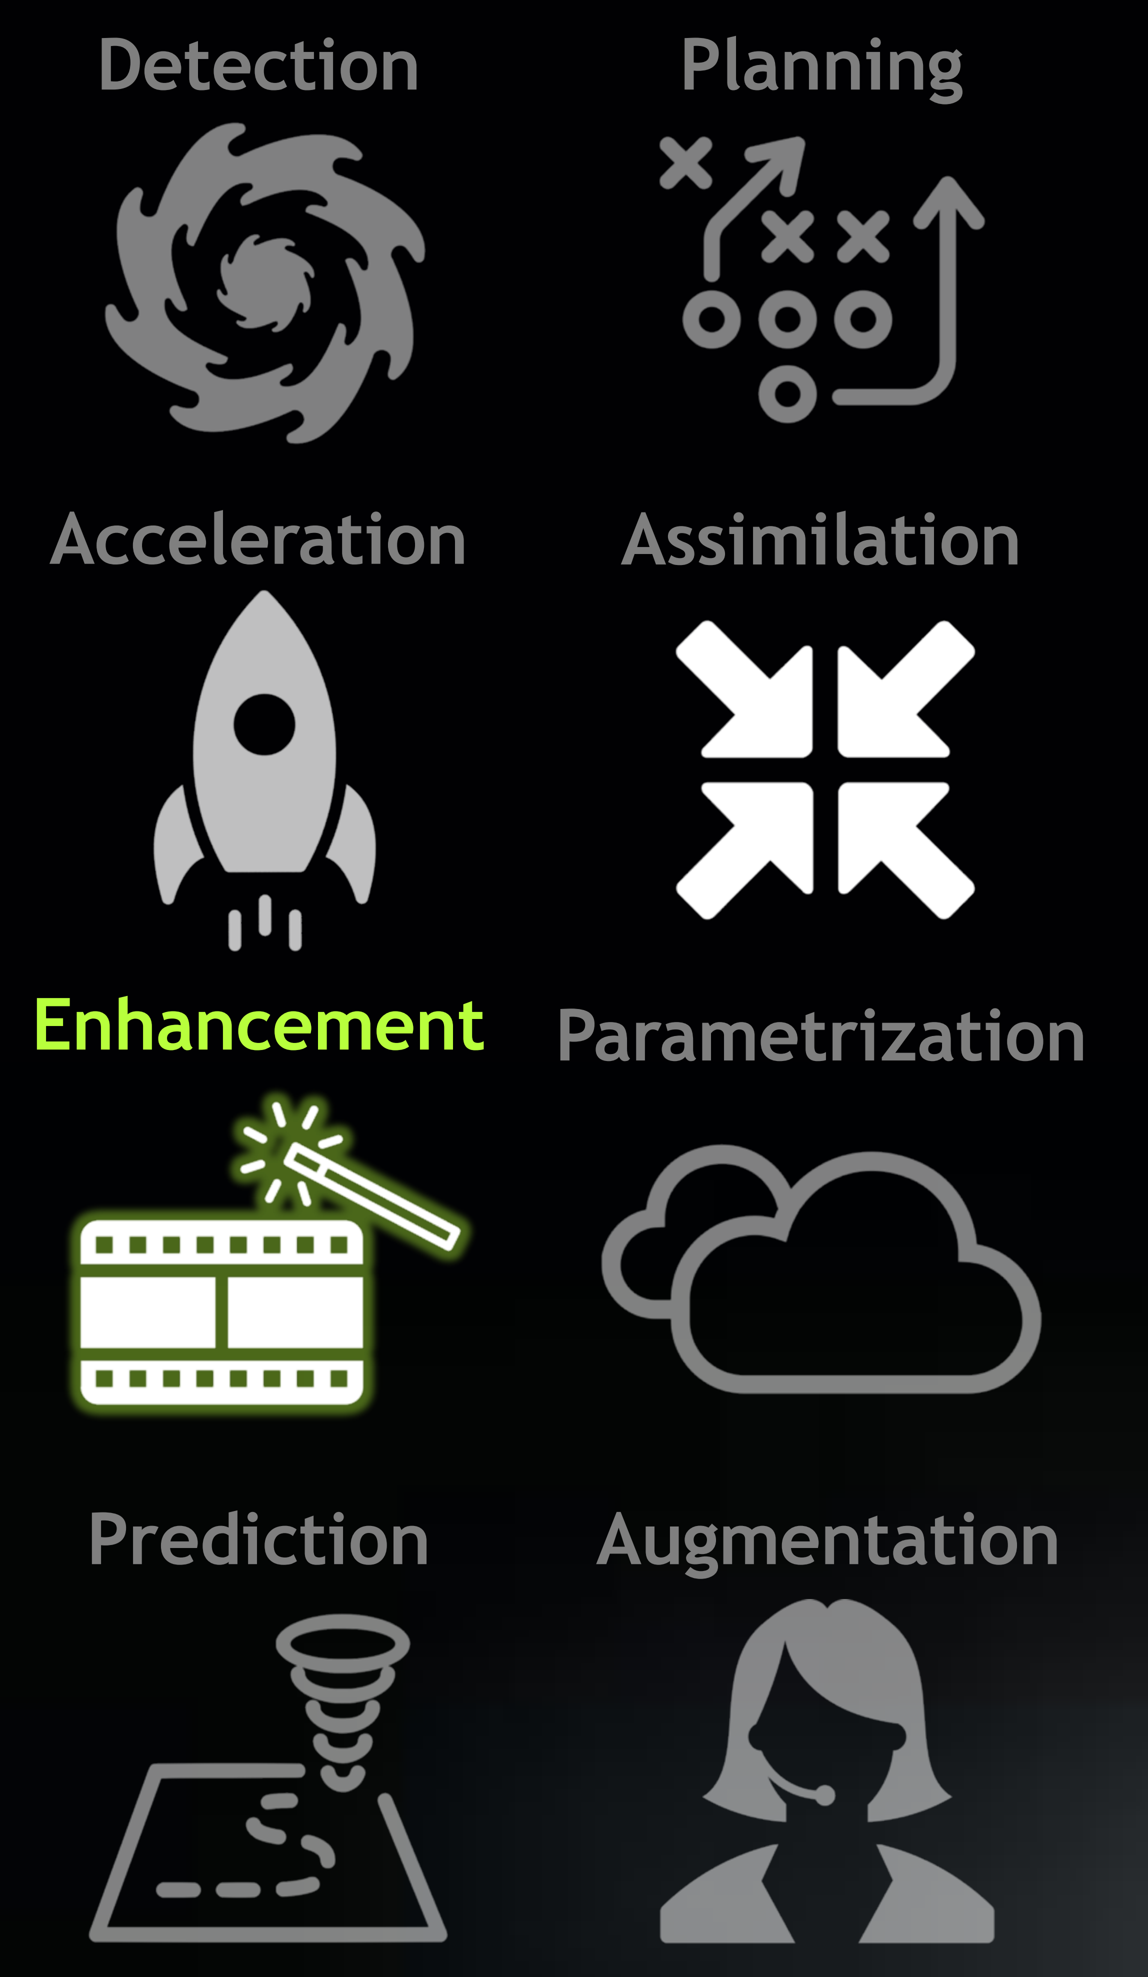

Enhance and Repair Your Data

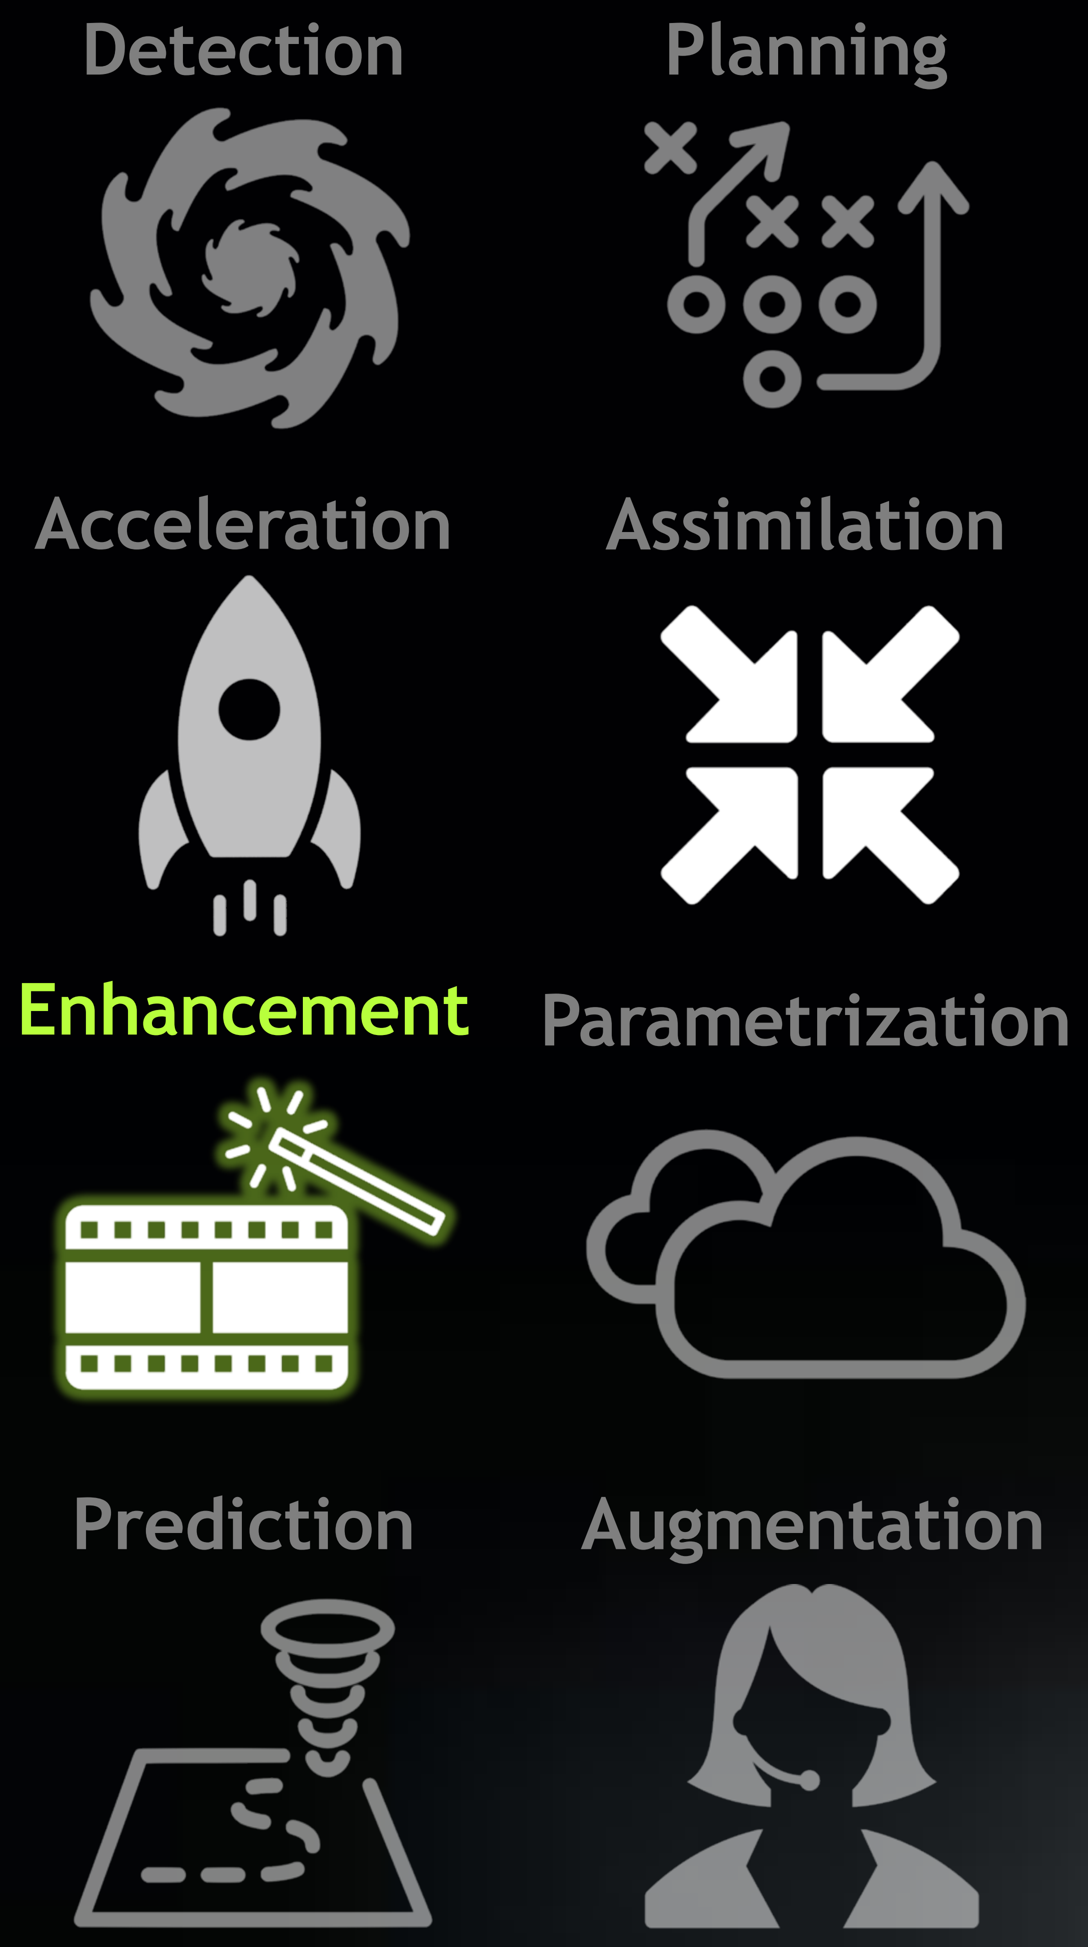

Use In-Painting to fill in Missing Data

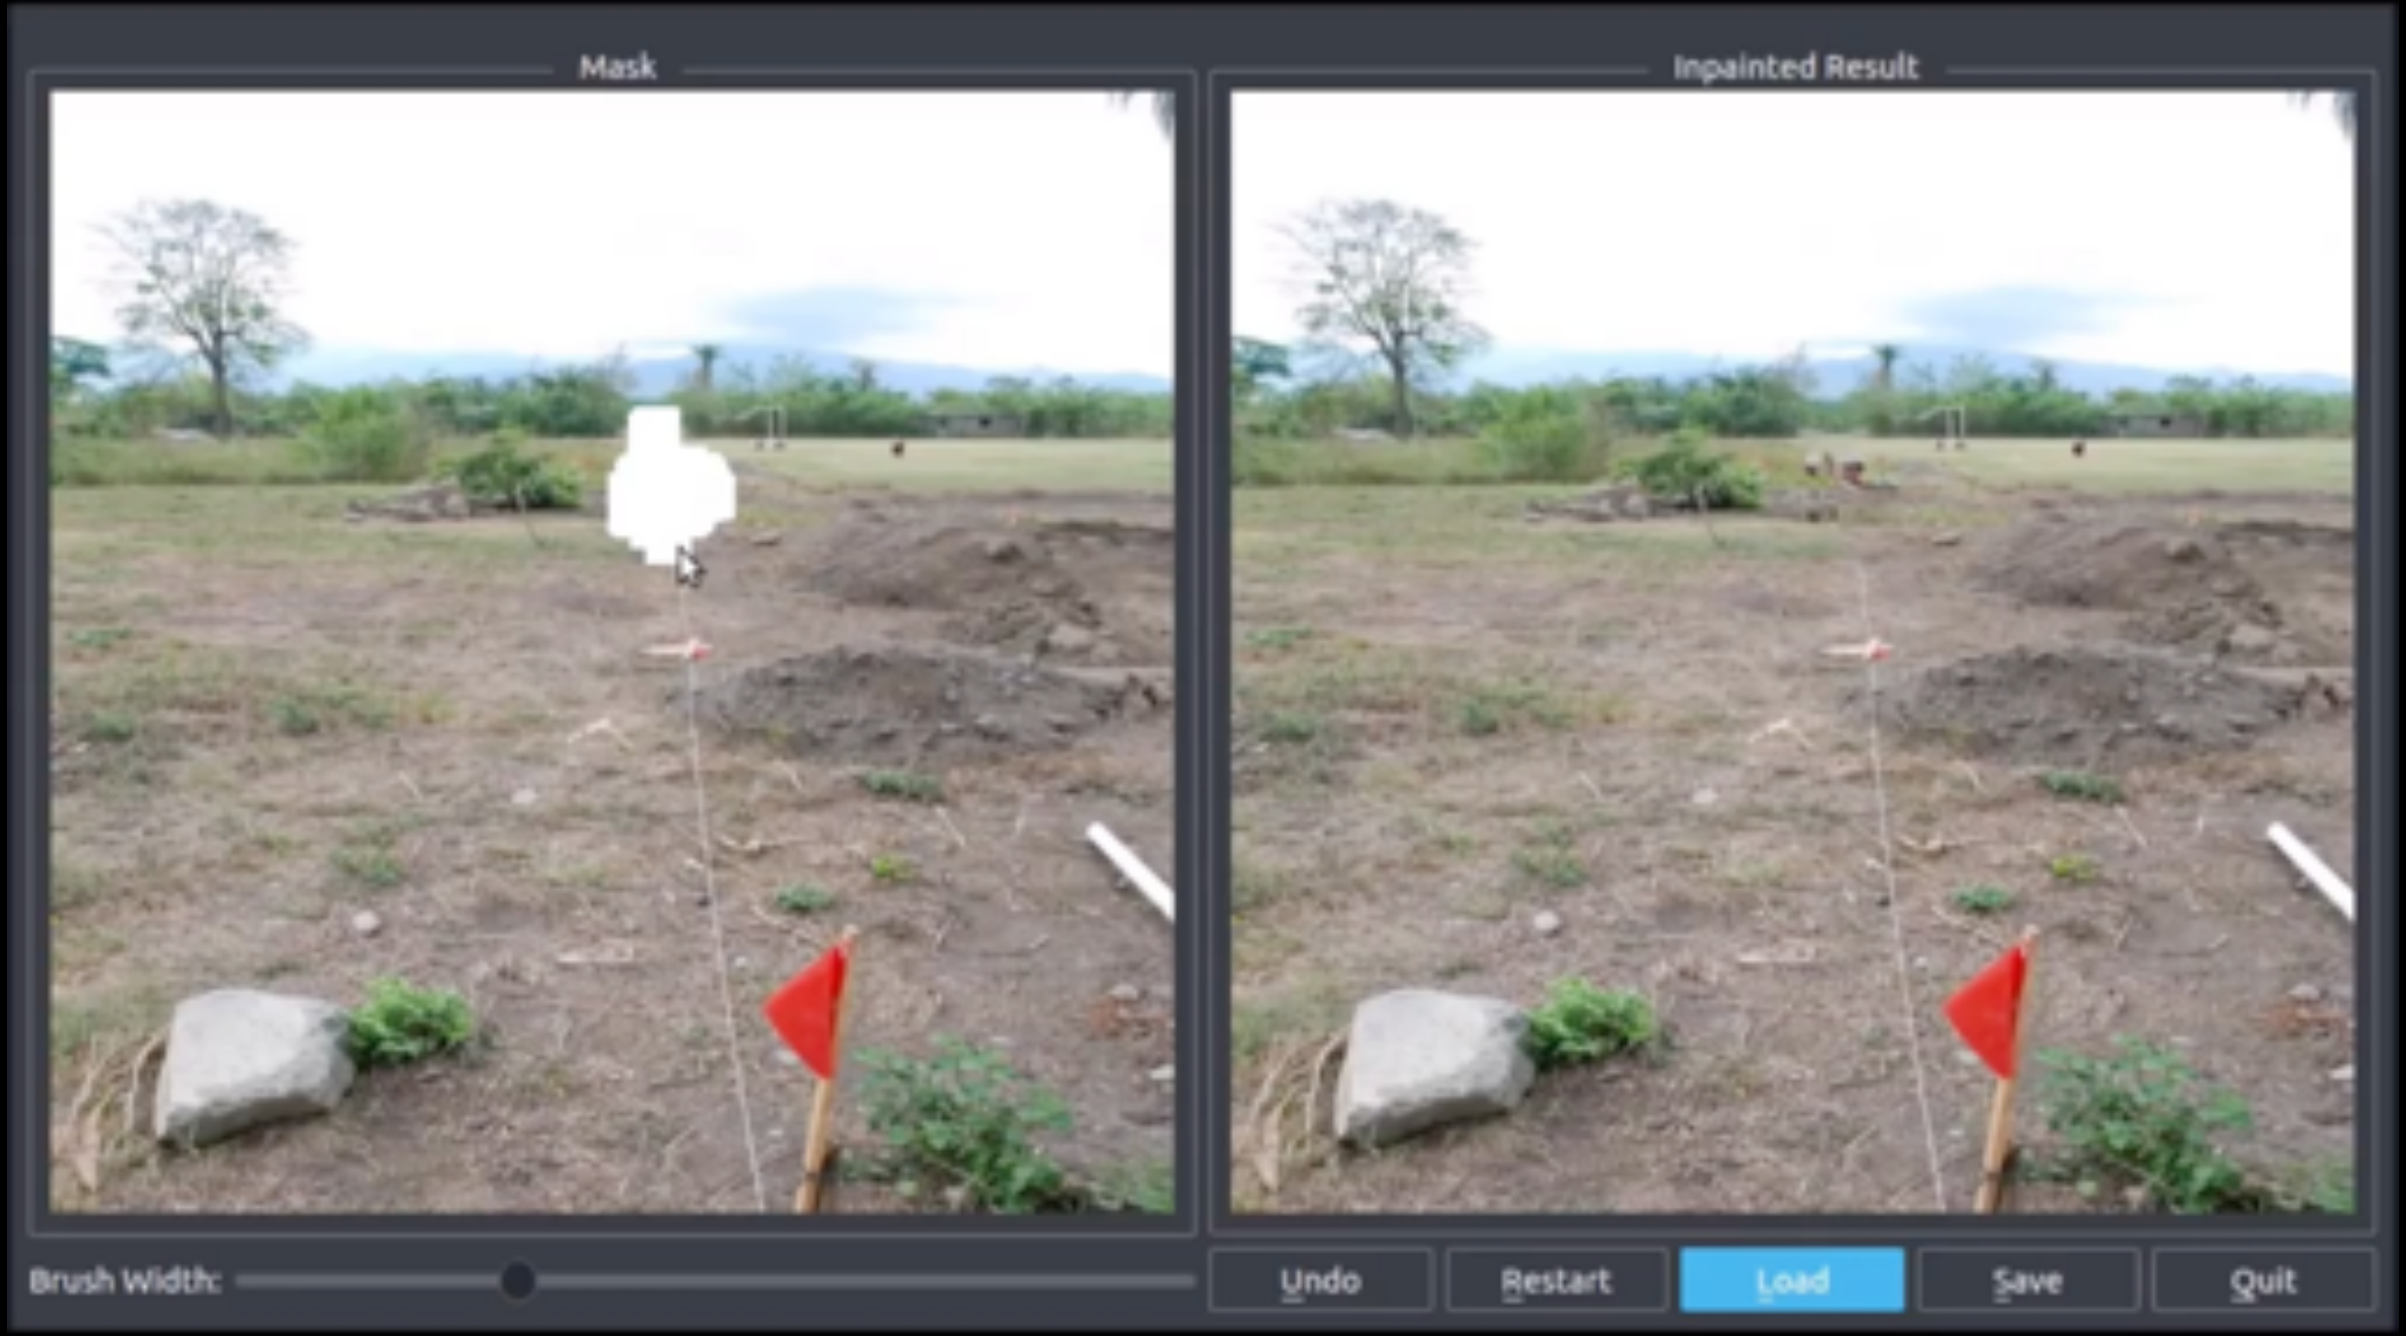

https://www.nvidia.com/research/inpainting/
https://arxiv.org/abs/1804.07723

Detection

Planning

Acceleration

Assimilation

Enhancement

Parametrization

Prediction

Augmentation

Use Inpainting to Repair Damaged GOES-17 Observations

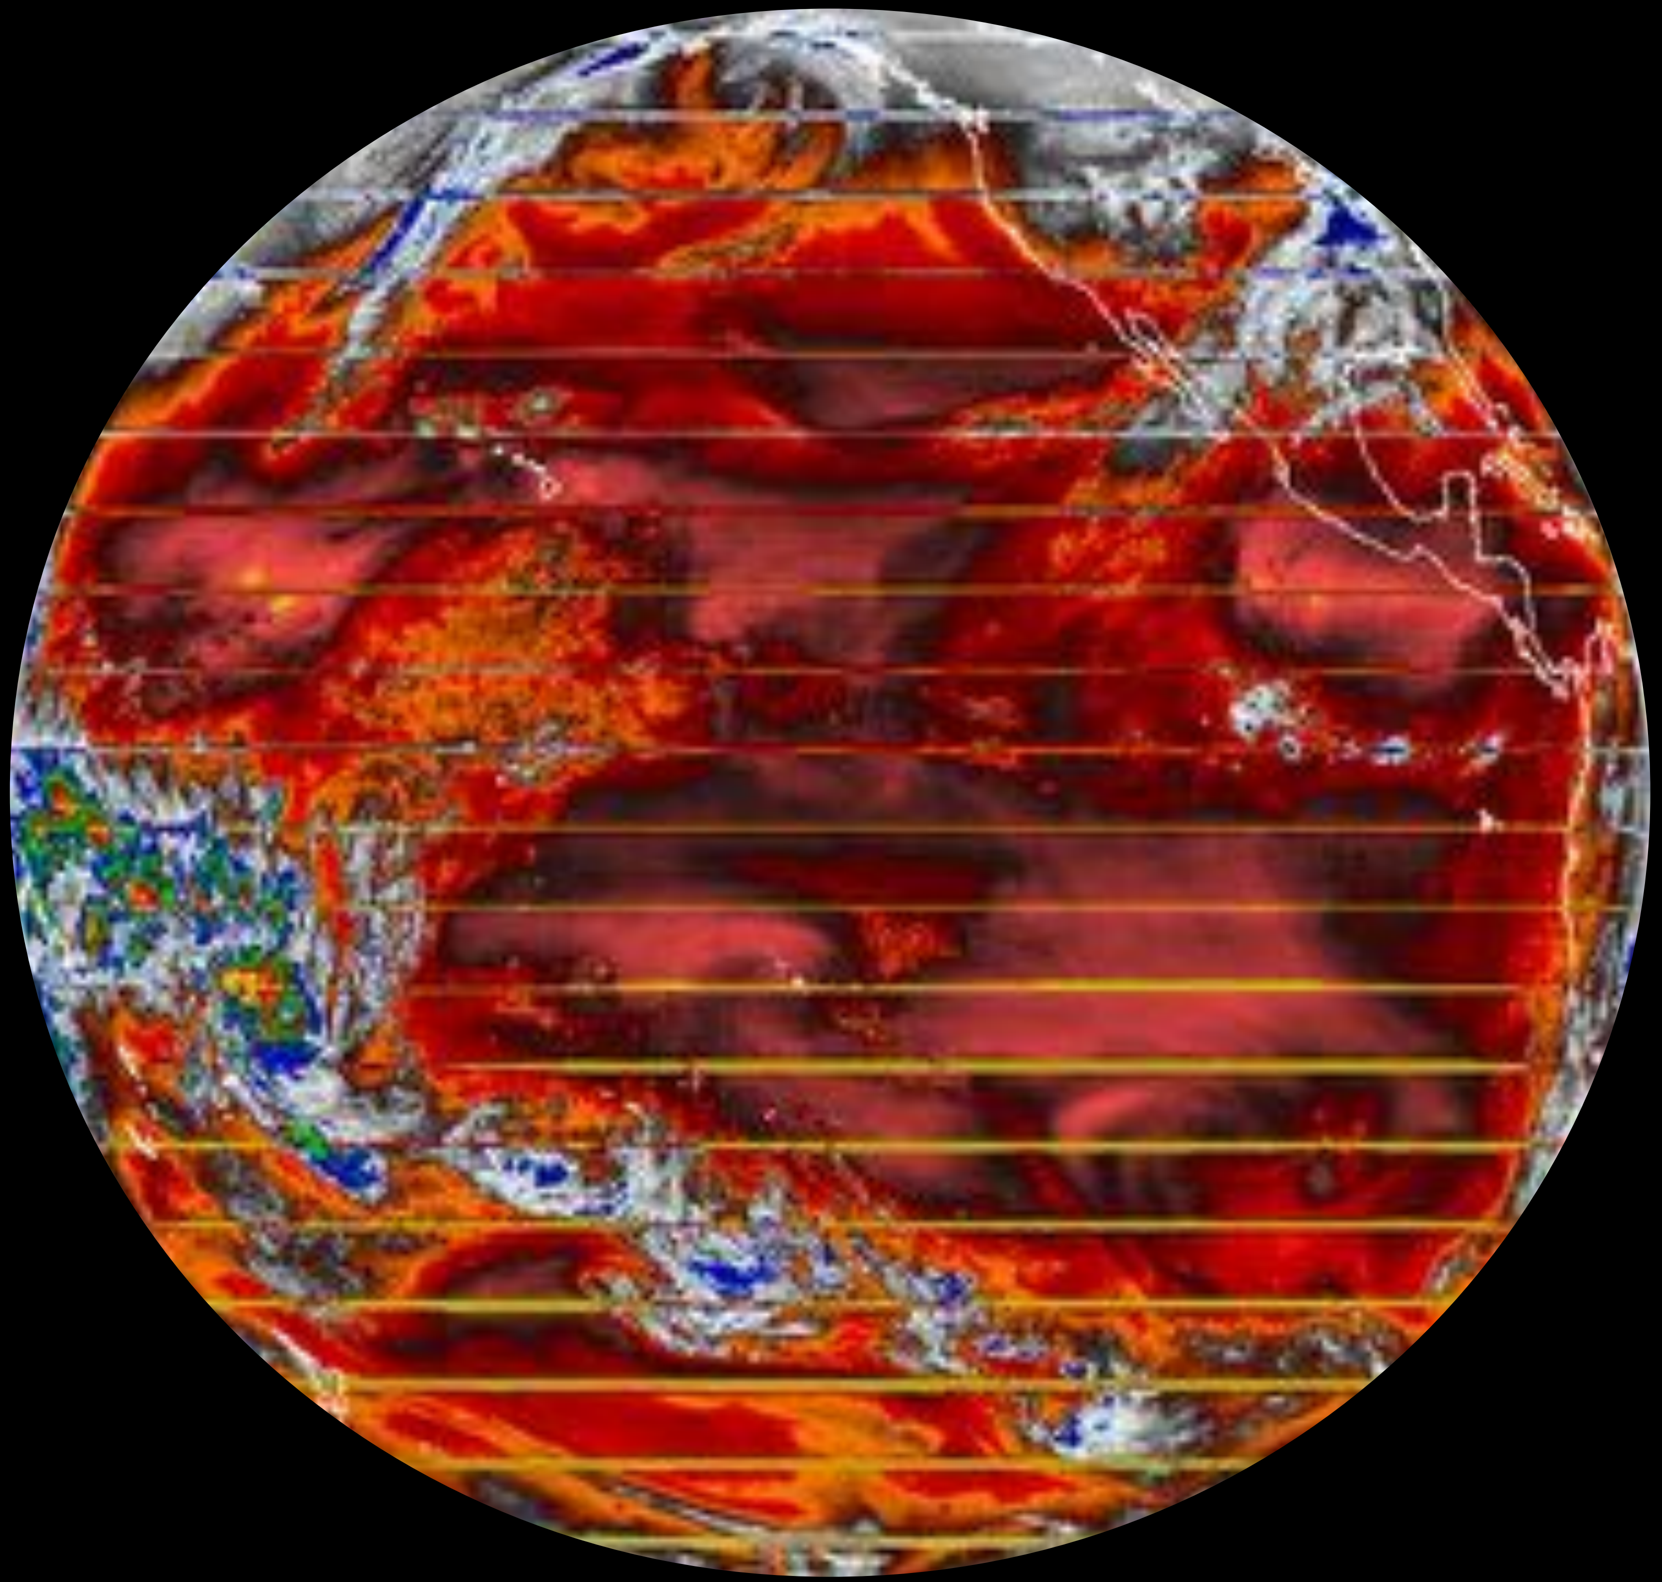

Matthew Pennybacker, Alexander Ignatov, Olafur Jonasson, Irina Gladkova, Boris Petrenko, Yury Kihai, "Mitigation of the GOES-17 ABI performance issues in the NOAA ACSPO SST products," Proc. SPIE 11014, Ocean Sensing and Monitoring XI, 110140Q (30 May 2019);https://doi.org/10.1117/12.2521051

Detection

Planning

Acceleration

Assimilation

Enhancement

Parametrization

Prediction

Augmentation

Use Inpainting to Reconstruct missing Climate Data

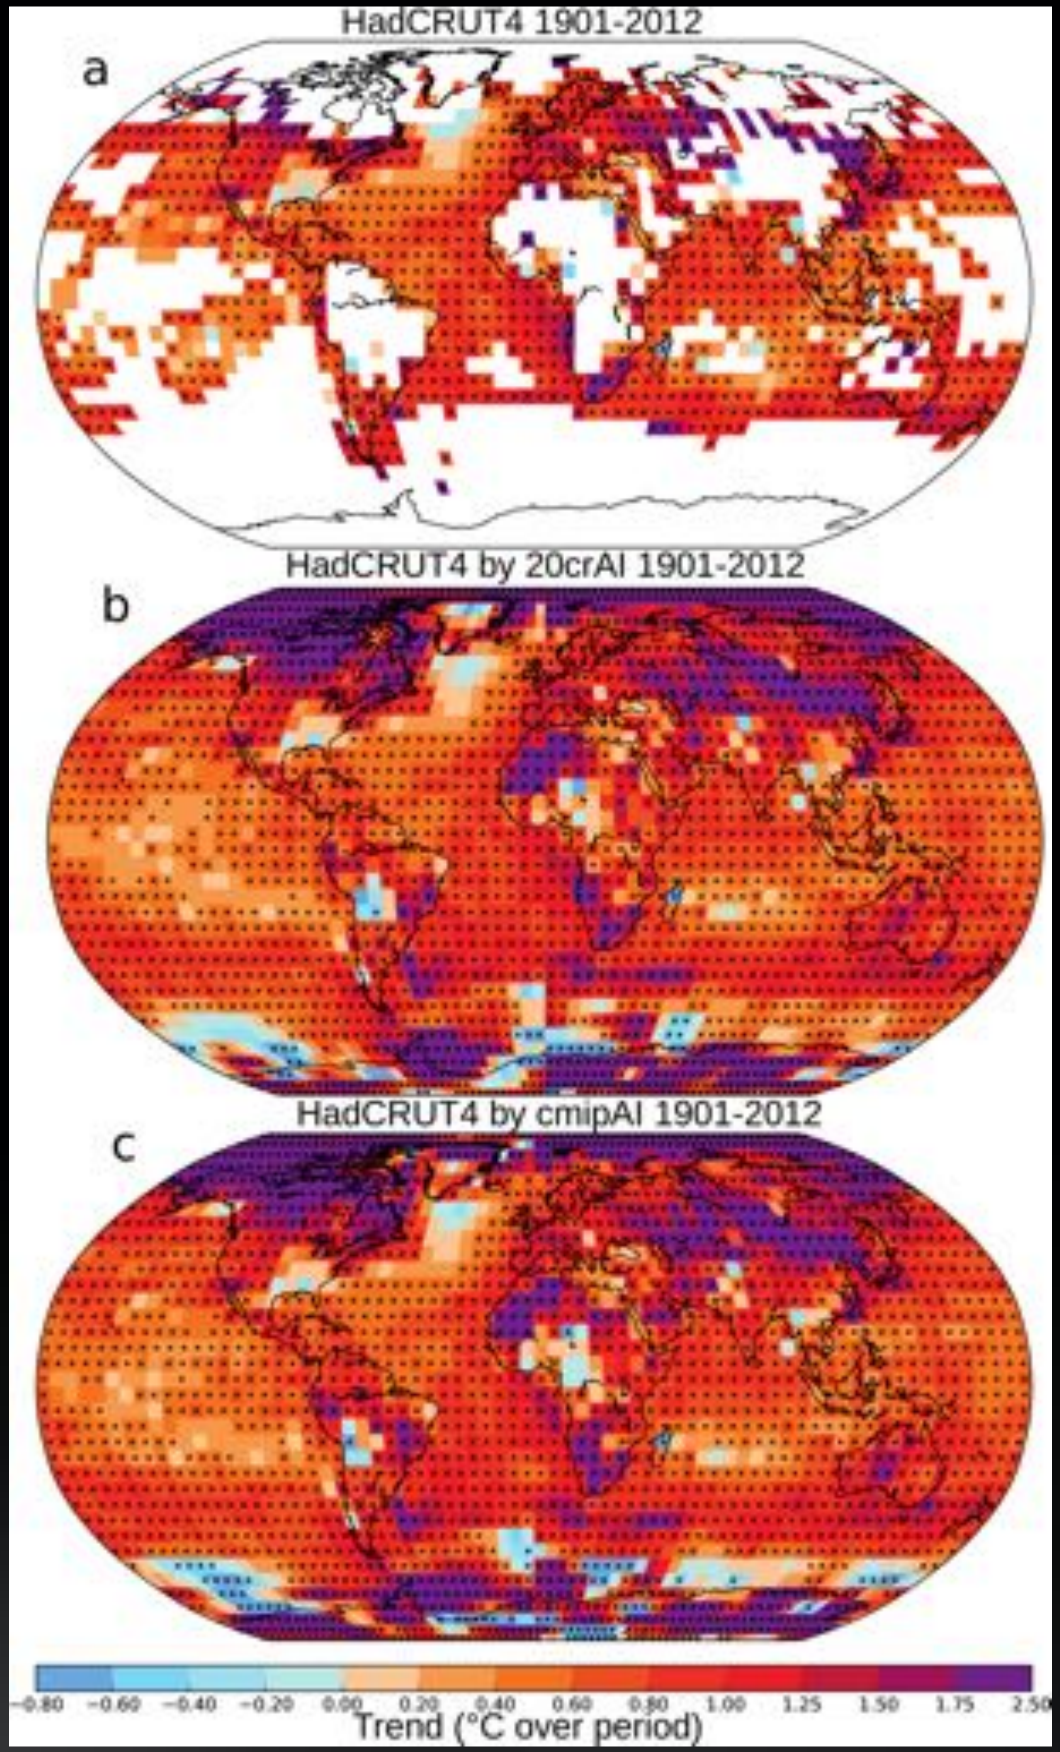

Artificial intelligence reconstructs missing climate information,
Christopher Kadow, David Matthew Hall and Uwe Ulbrich
Nature Geoscience

Detection

Planning

Acceleration

Assimilation

Enhancement

Parametrization

Prediction

Augmentation

Deep Learning Super Resolution to Fill in Details

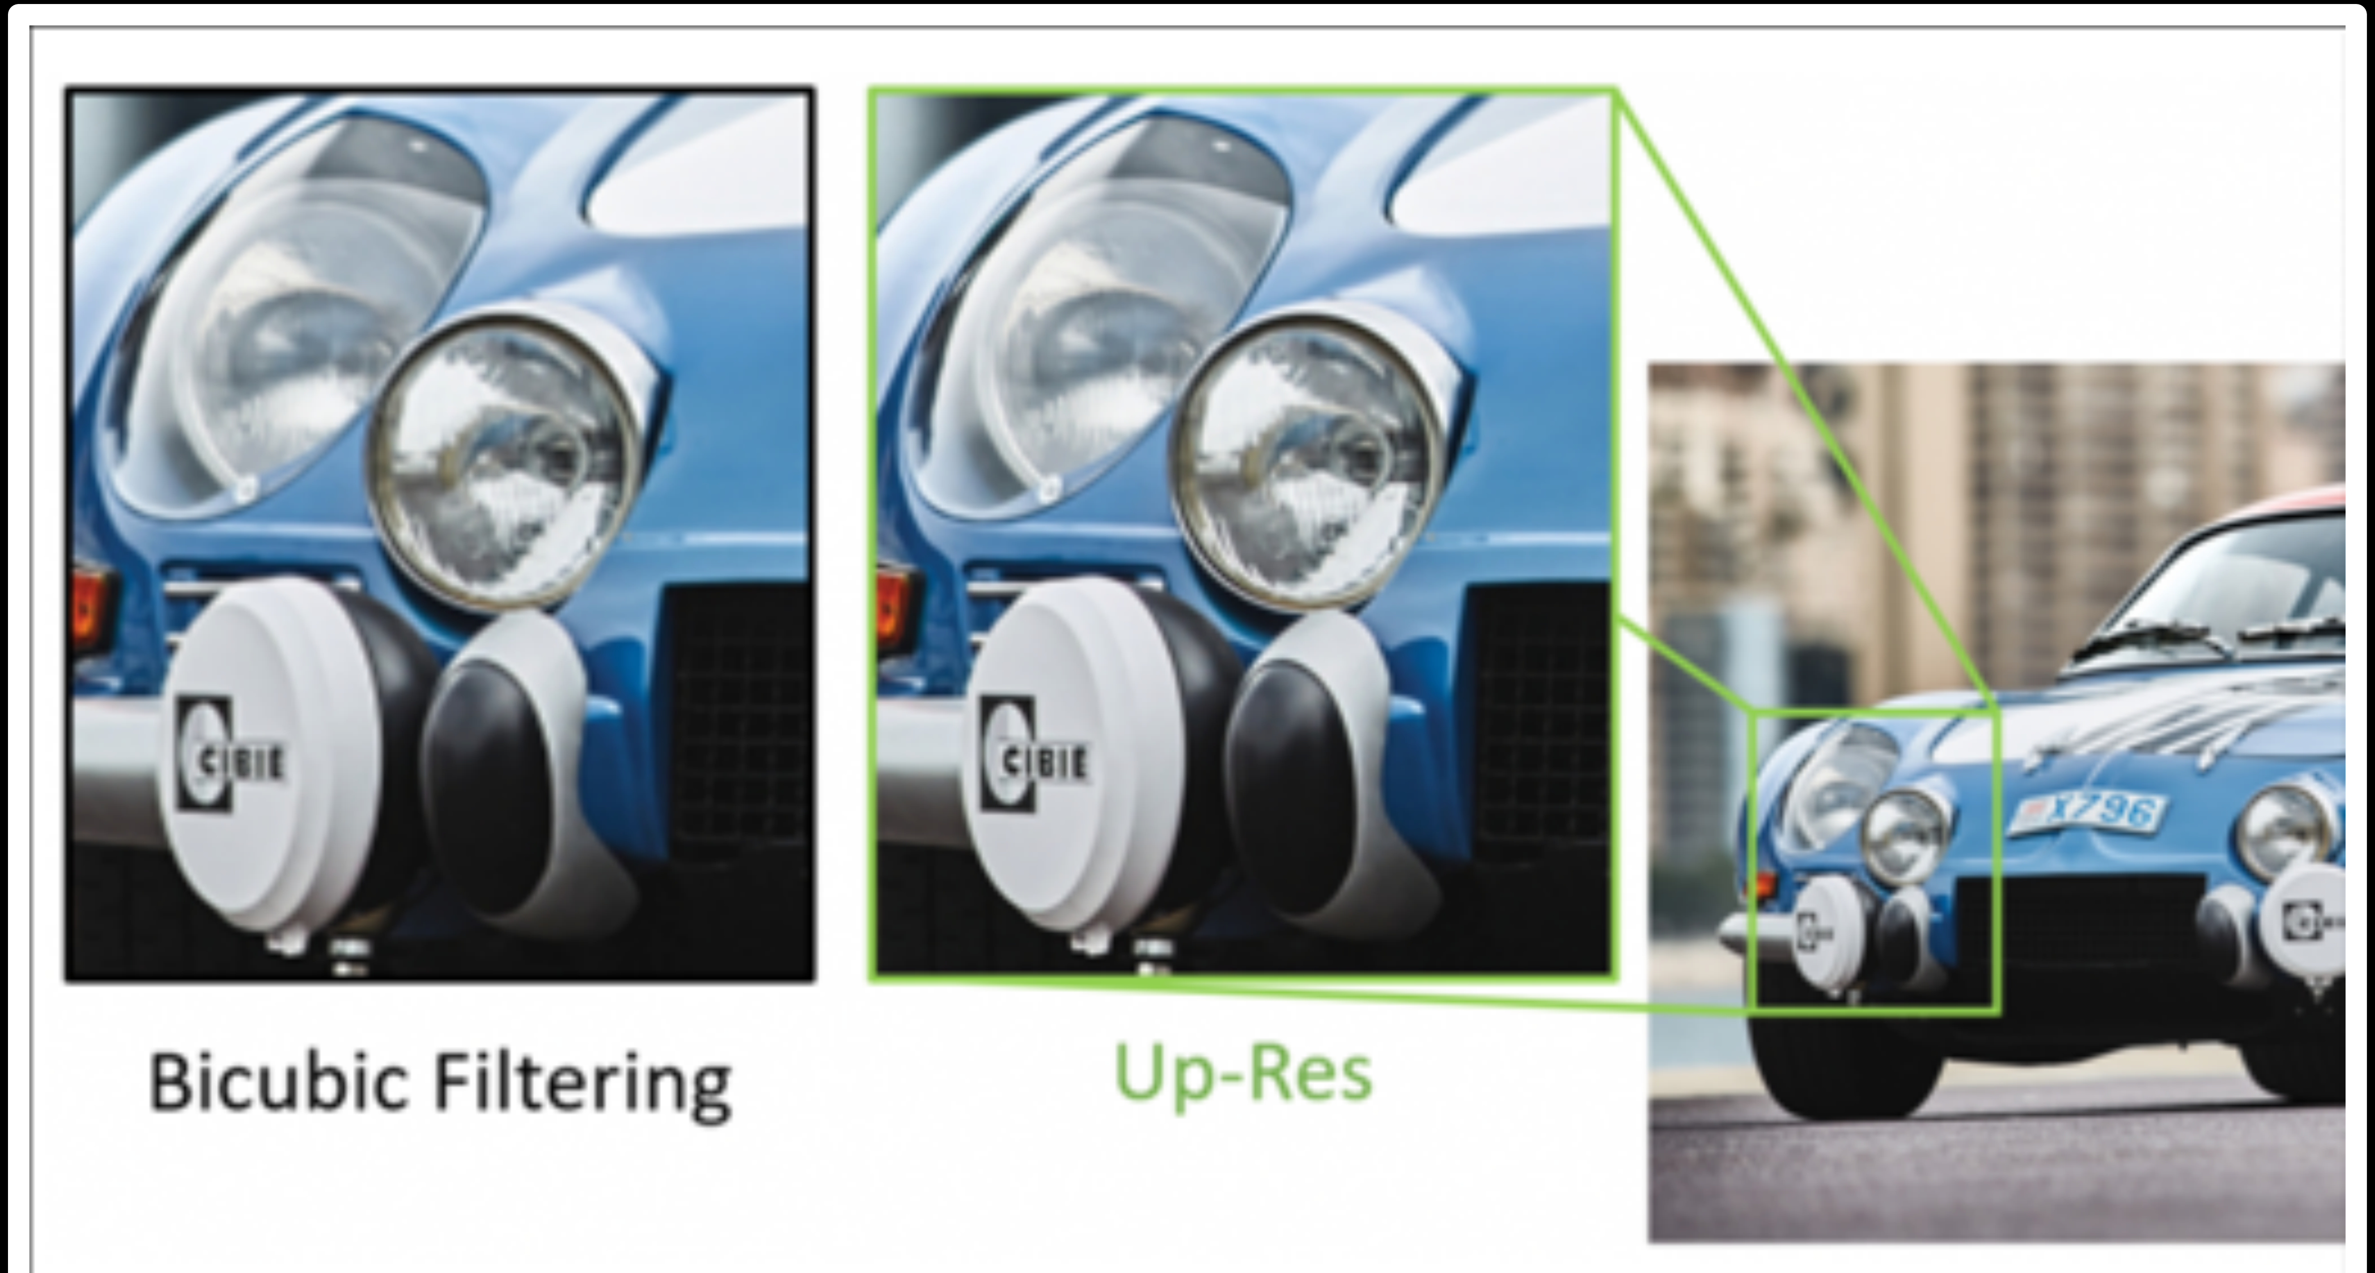

Image above: Up-Res provides improved image clarity over Bicubic filtering.

Detection

Planning

Acceleration

Assimilation

Enhancement

Parametrization

Prediction

Augmentation

Super Resolution techniques for More Accurate Downscaling

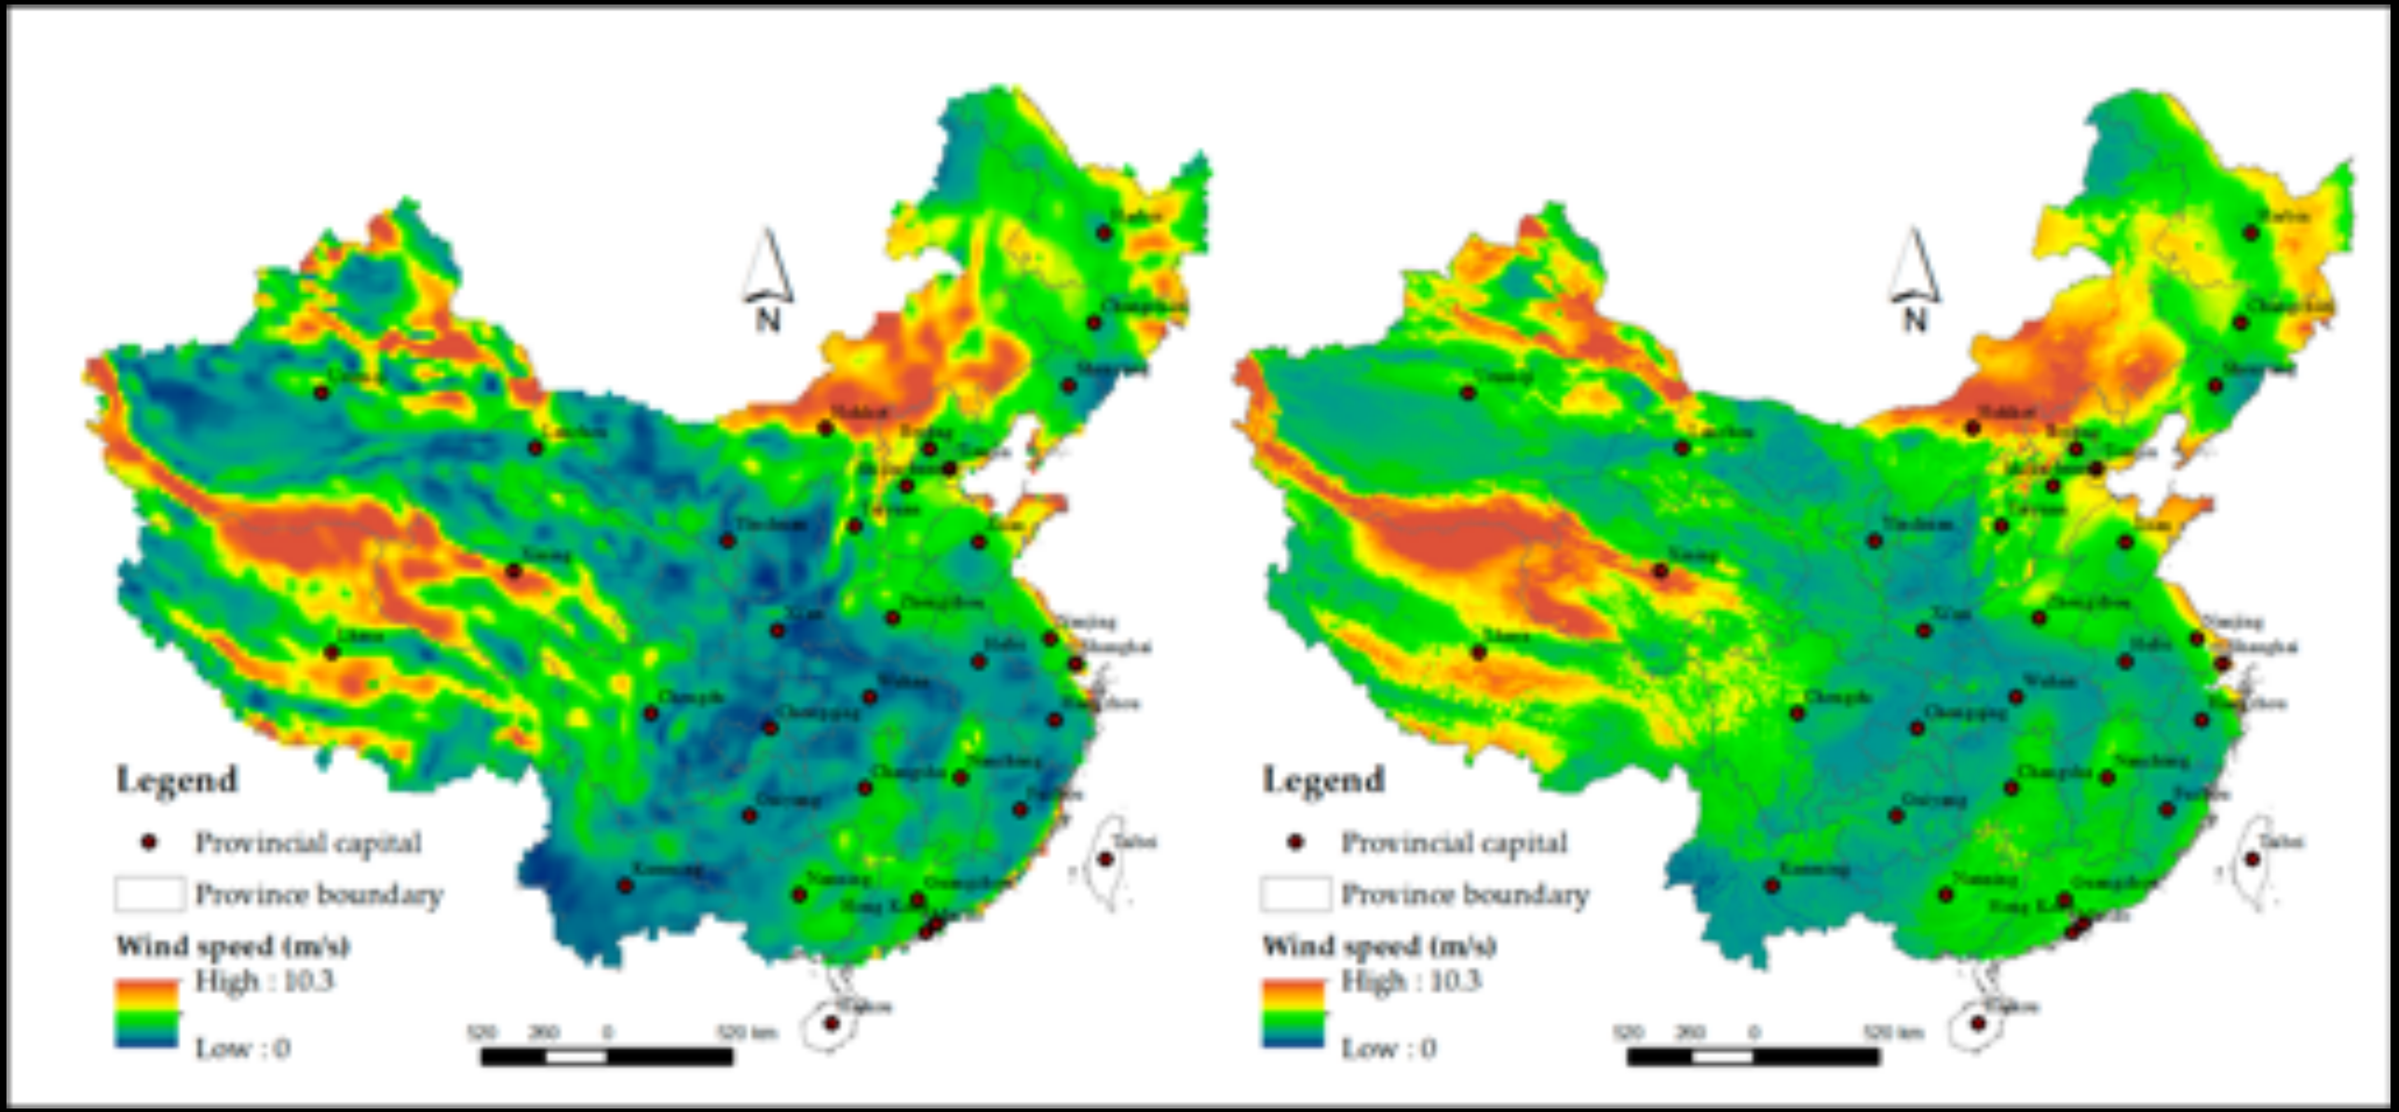

Remote Sens. **2019**, *11*, 1378; doi:10.3390/rs11111378

Detection

Planning

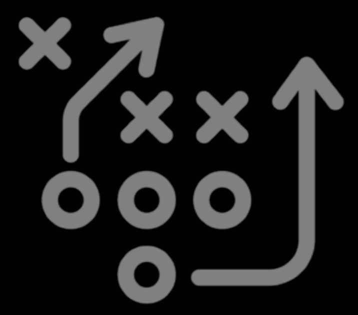

Acceleration

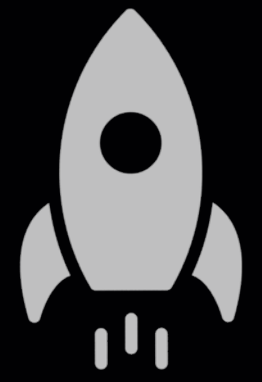

Assimilation

Enhancement

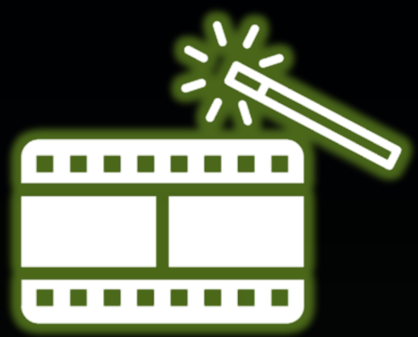

Parametrization

Prediction

Augmentation

Slow-motion interpolation of video

Detection

Planning

Acceleration

Assimilation

Enhancement

Parametrization

Prediction

Augmentation

High accuracy temporal interpolation

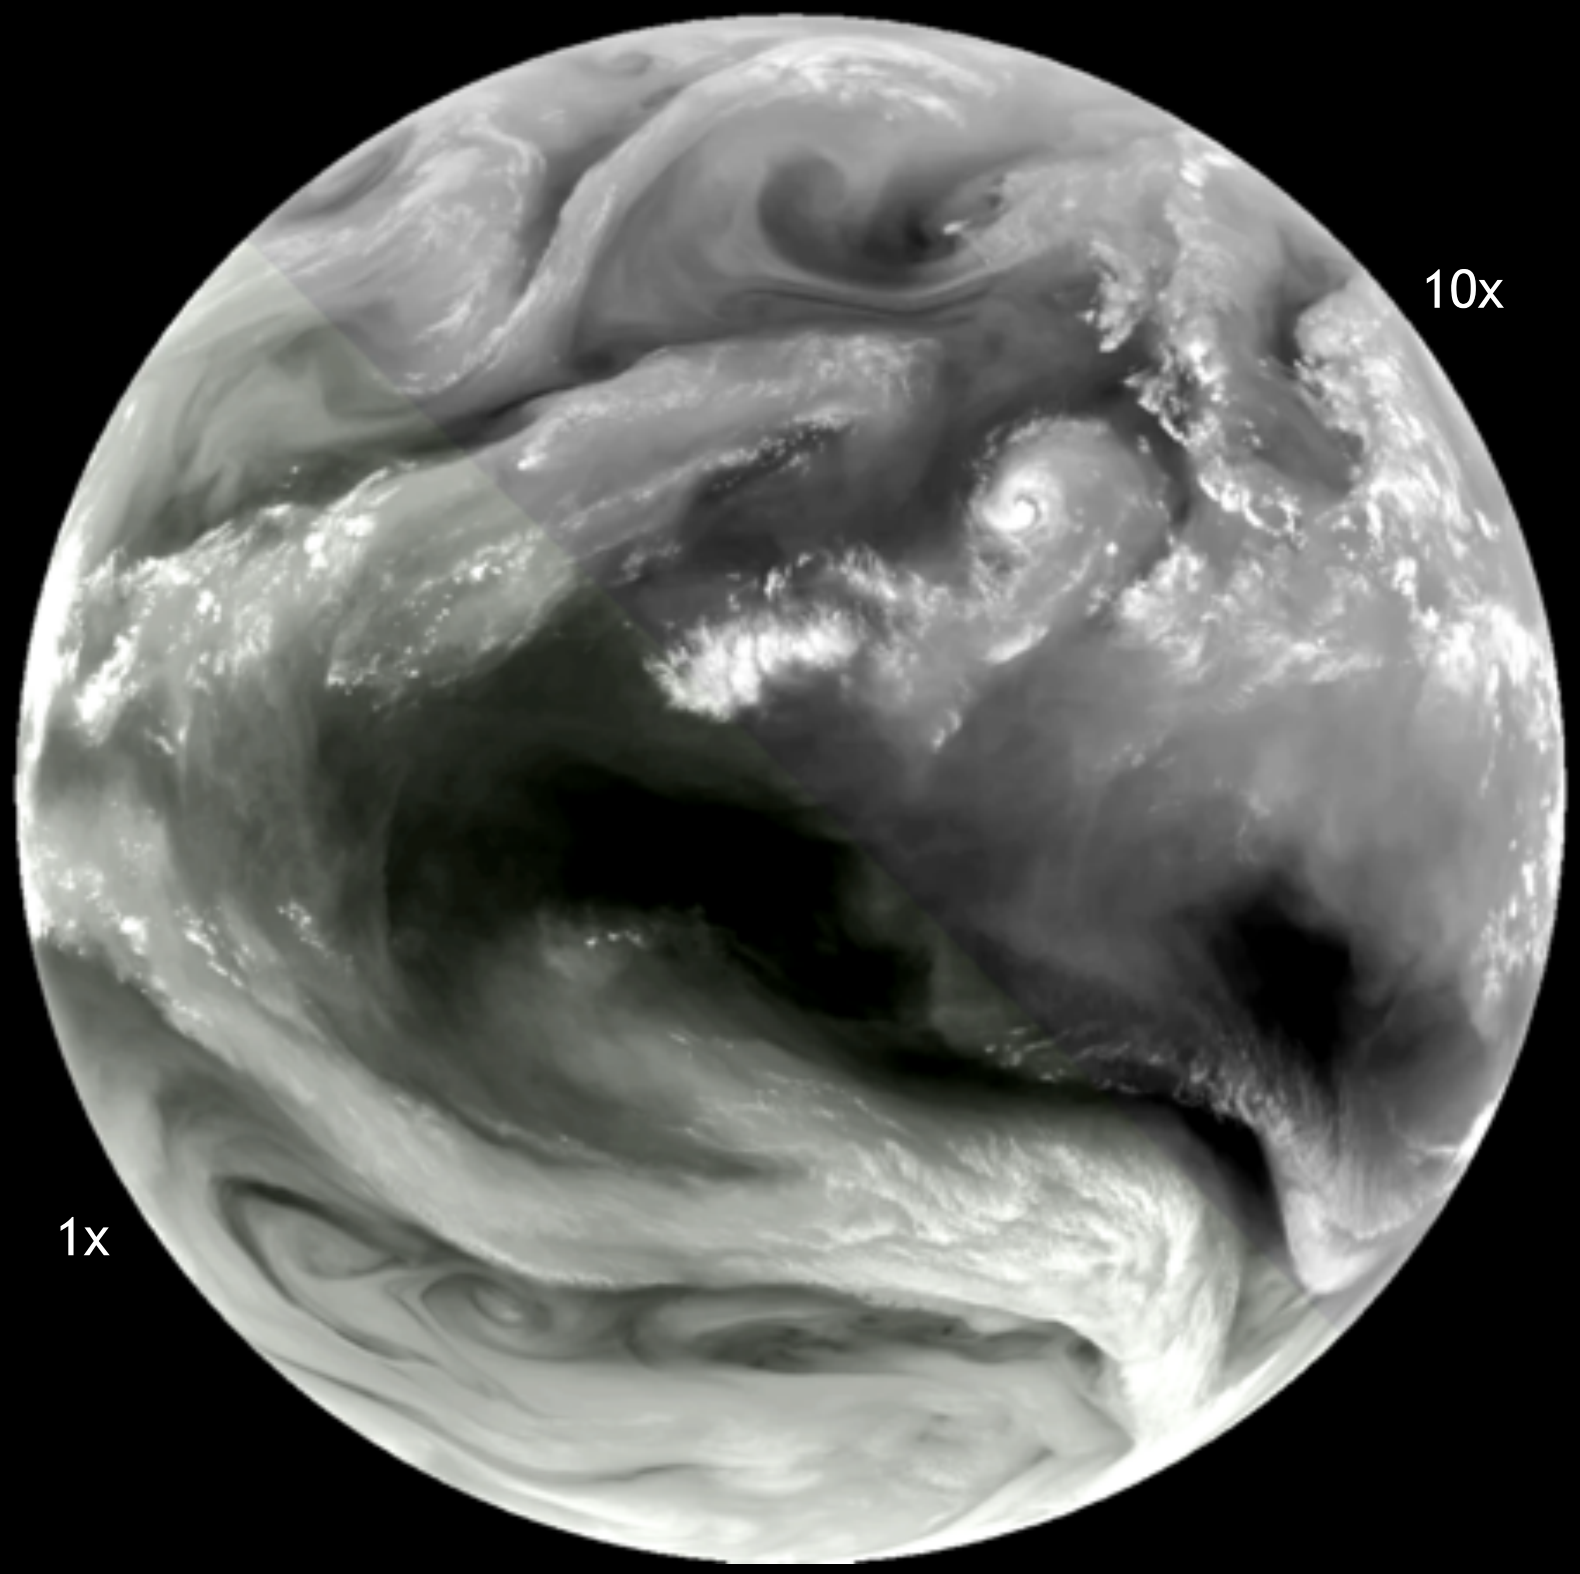

Temporal Interpolation of Geostationary Satellite Imagery with Task Specific Optical Flow, Thomas Vandal, Ramakrishna Nemani, https://arxiv.org/abs/1907.12013

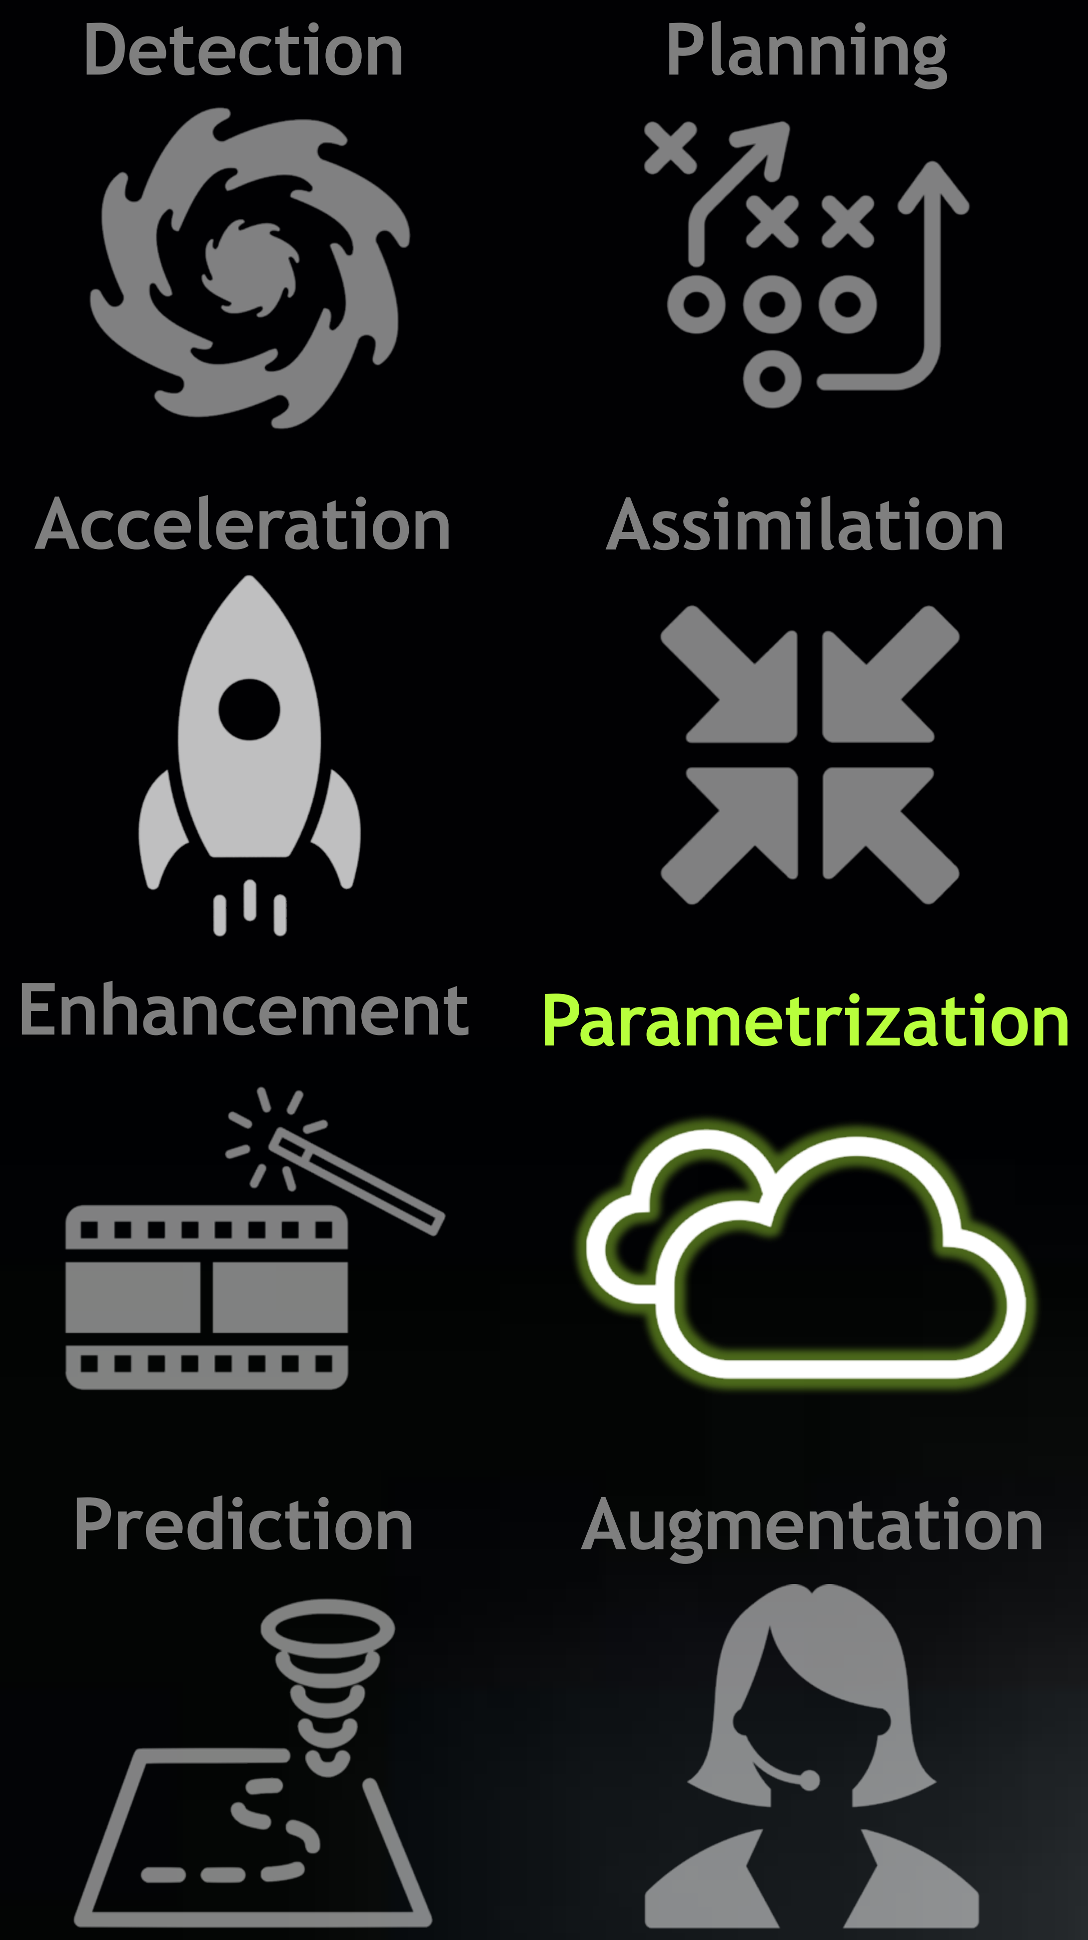

Accelerate and Improve Physical Parametrizations

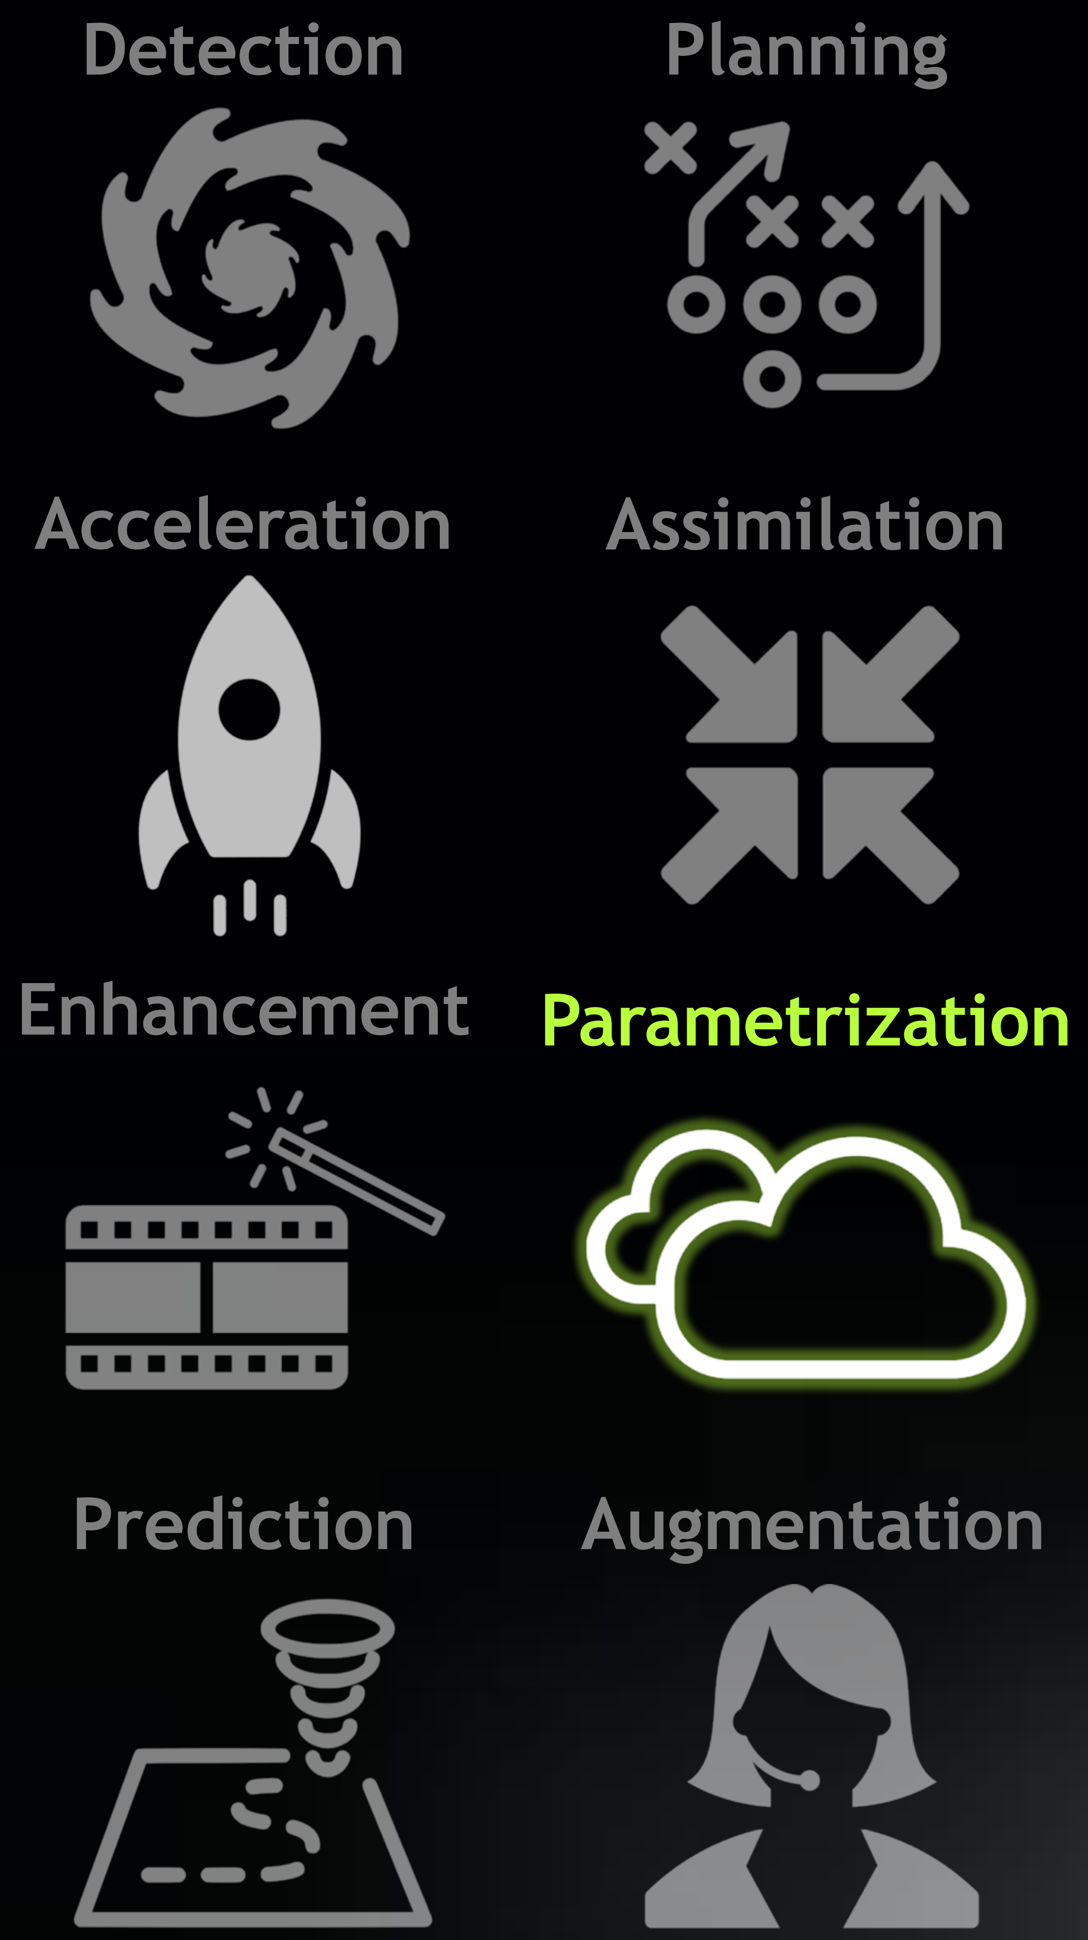

Accelerate Existing Parametrizations

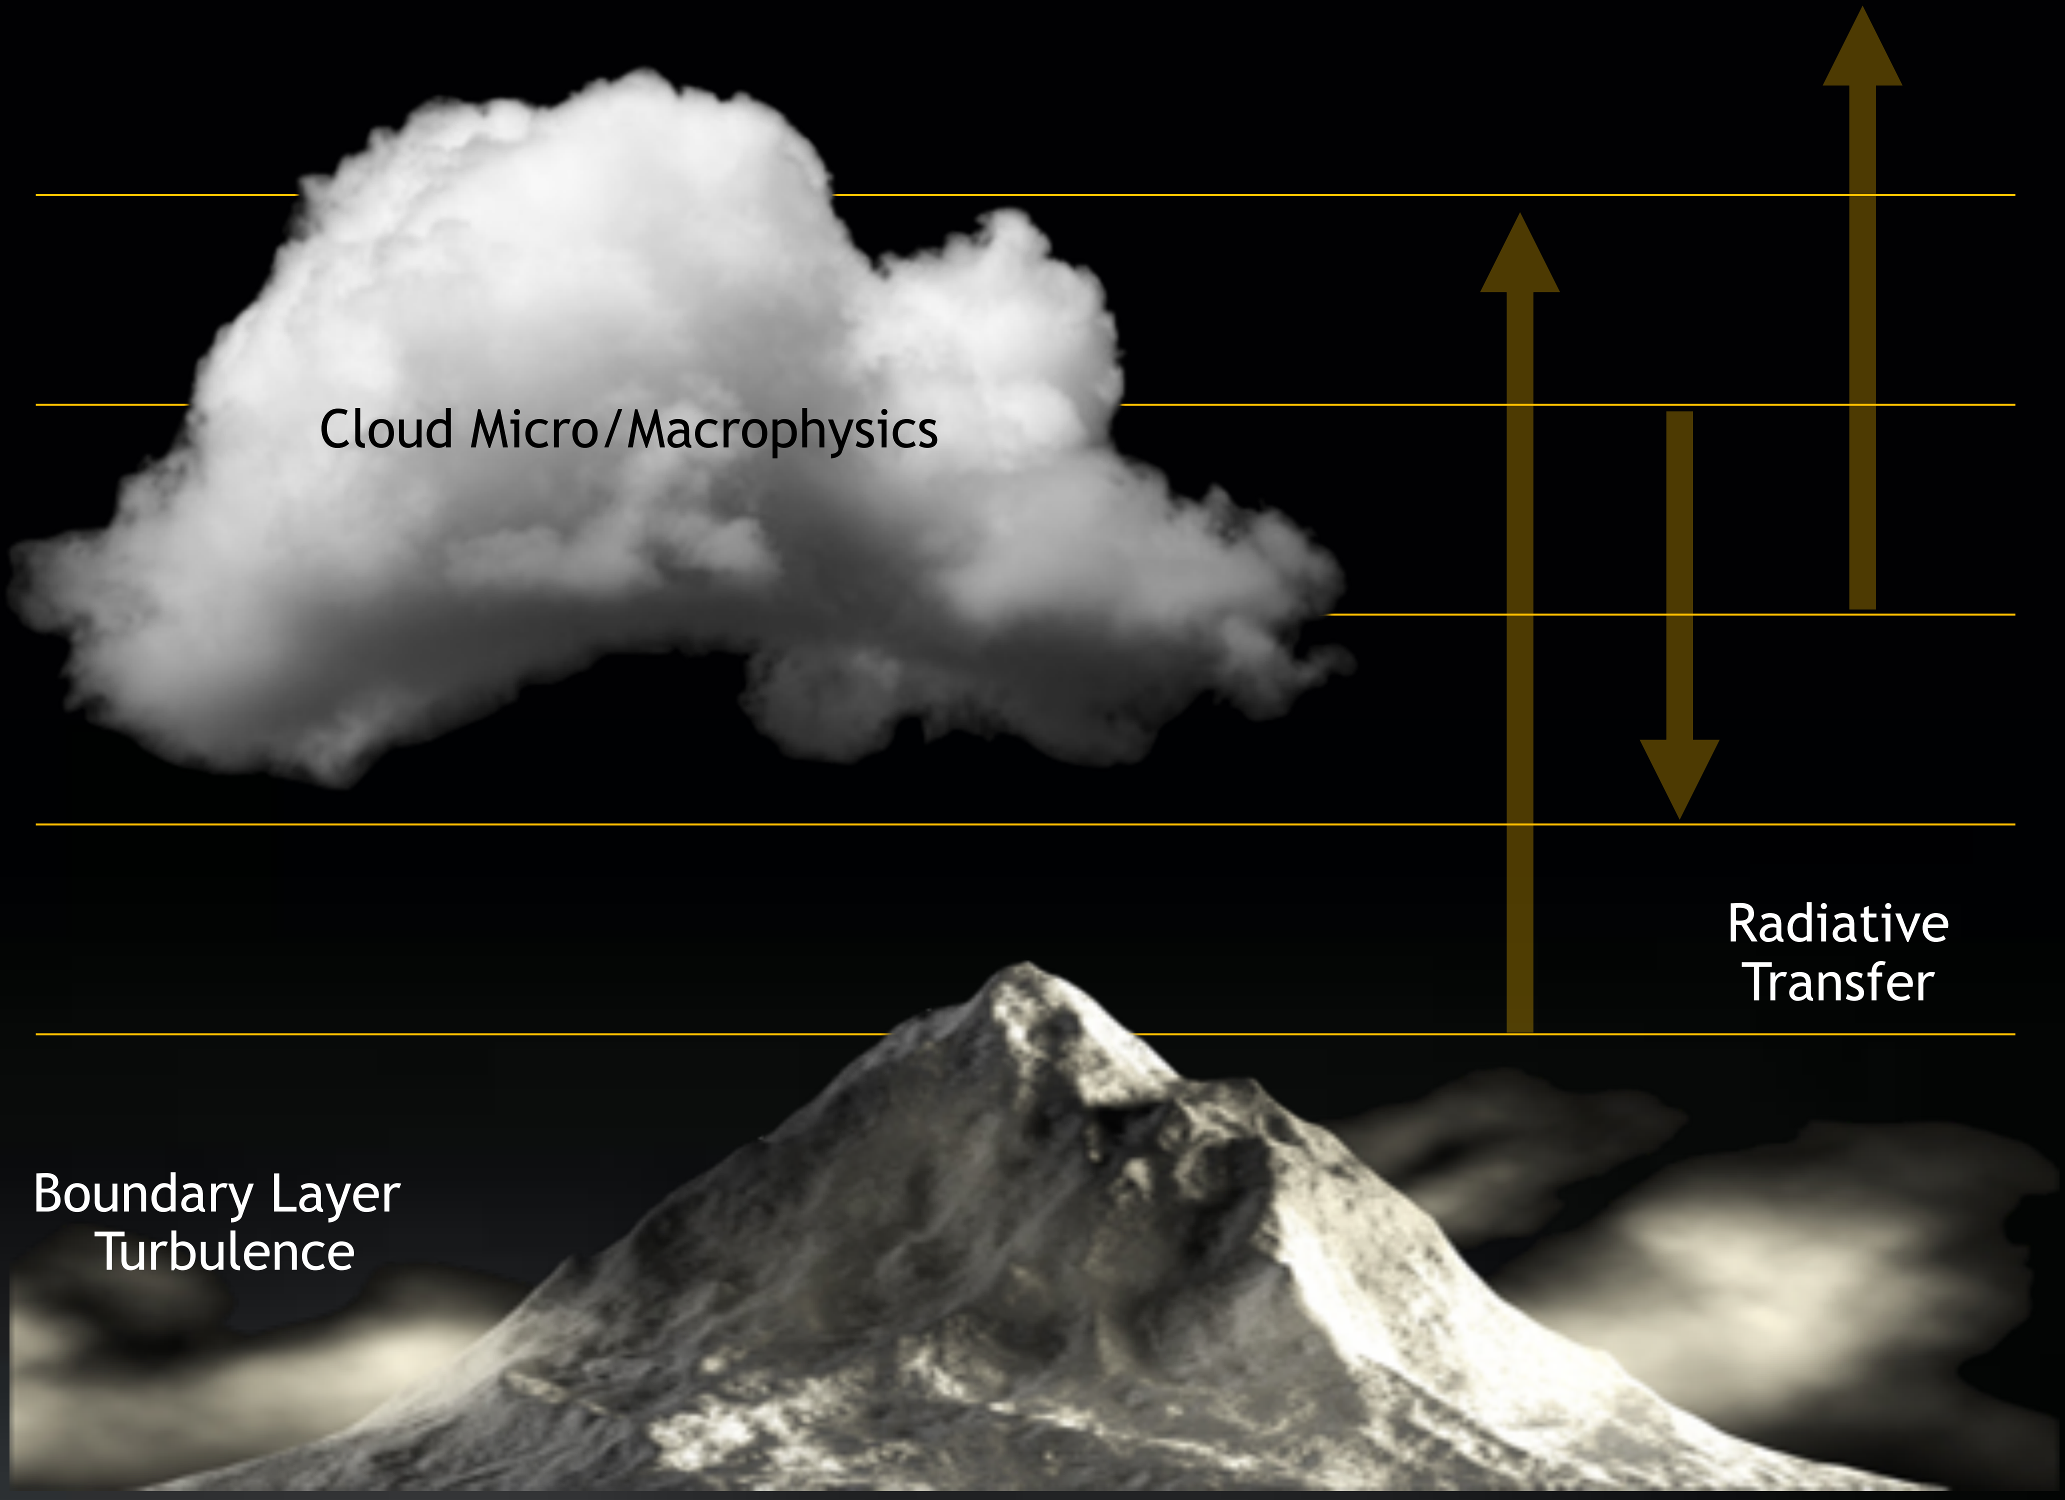

Dueben, Hogan, Bauer @ECMWF and Progsch, Angerer @NVIDIA

Detection

Planning

Acceleration

Assimilation

Enhancement

Parametrization

Prediction

Augmentation

Parametrizations from high-res simulations

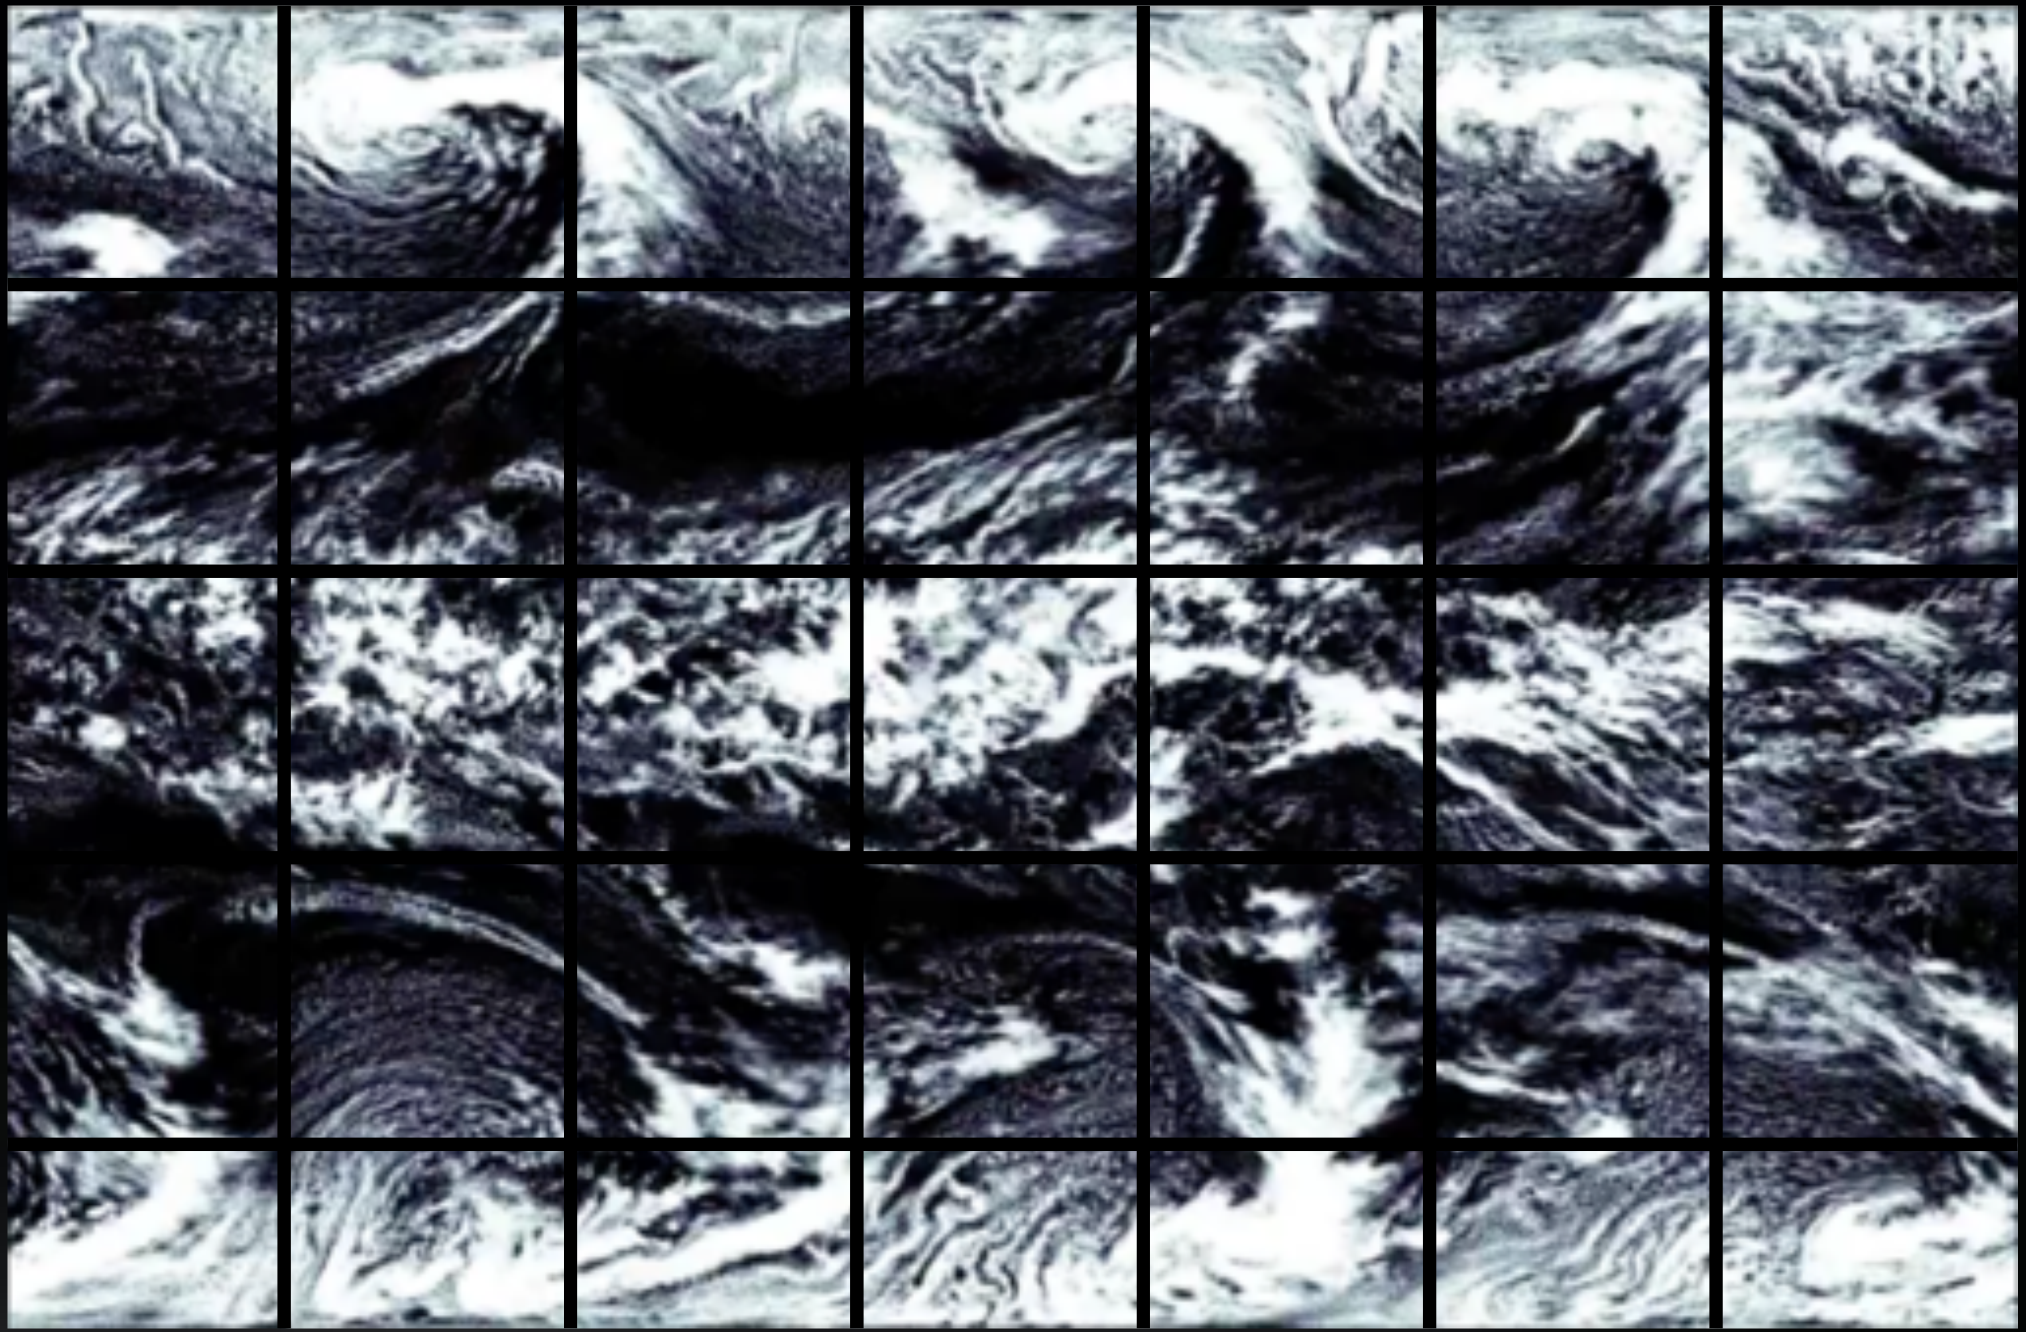

Noah Brenowitz and Cristopher Bretherton, University of Washington

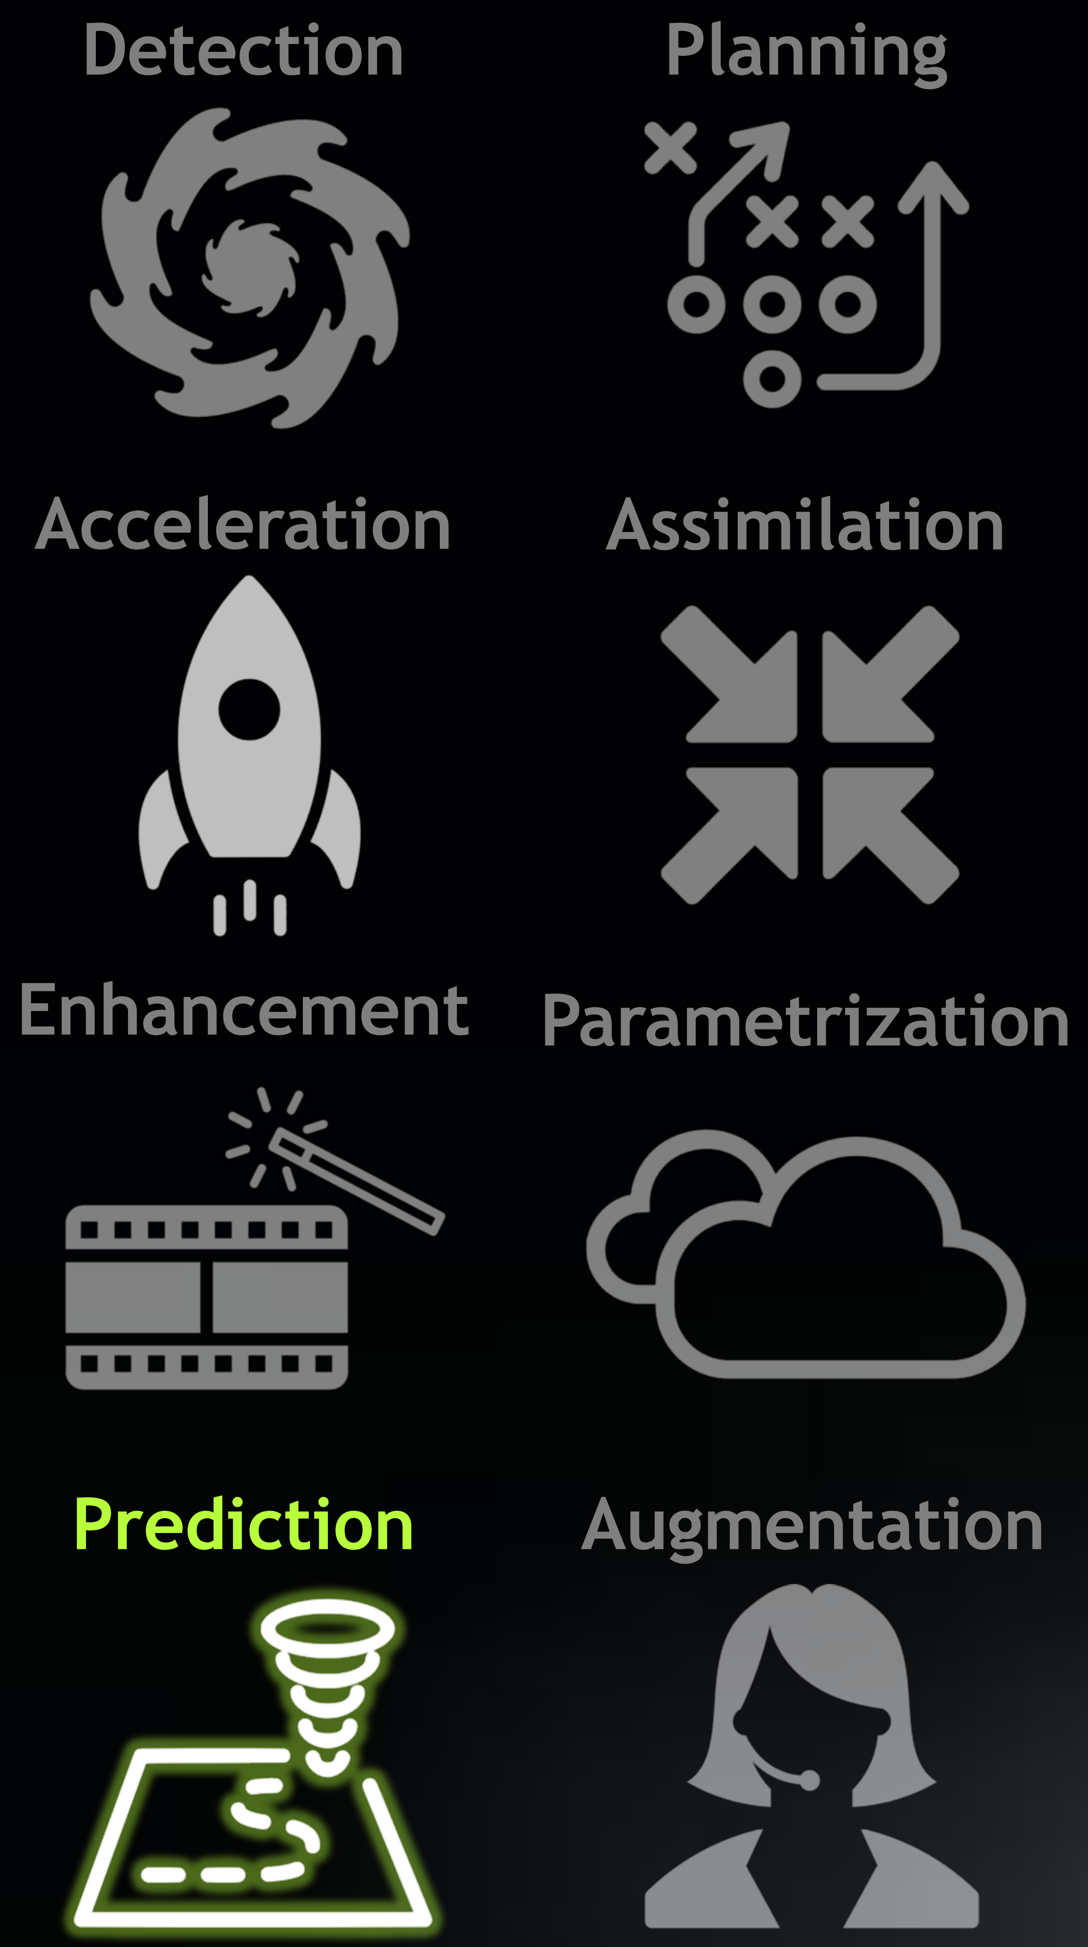

Create More Accurate Time-Series Predictions

Detection

Planning

Acceleration

Assimilation

Enhancement

Parametrization

Prediction

Augmentation

Improve storm track / intensity forecasts

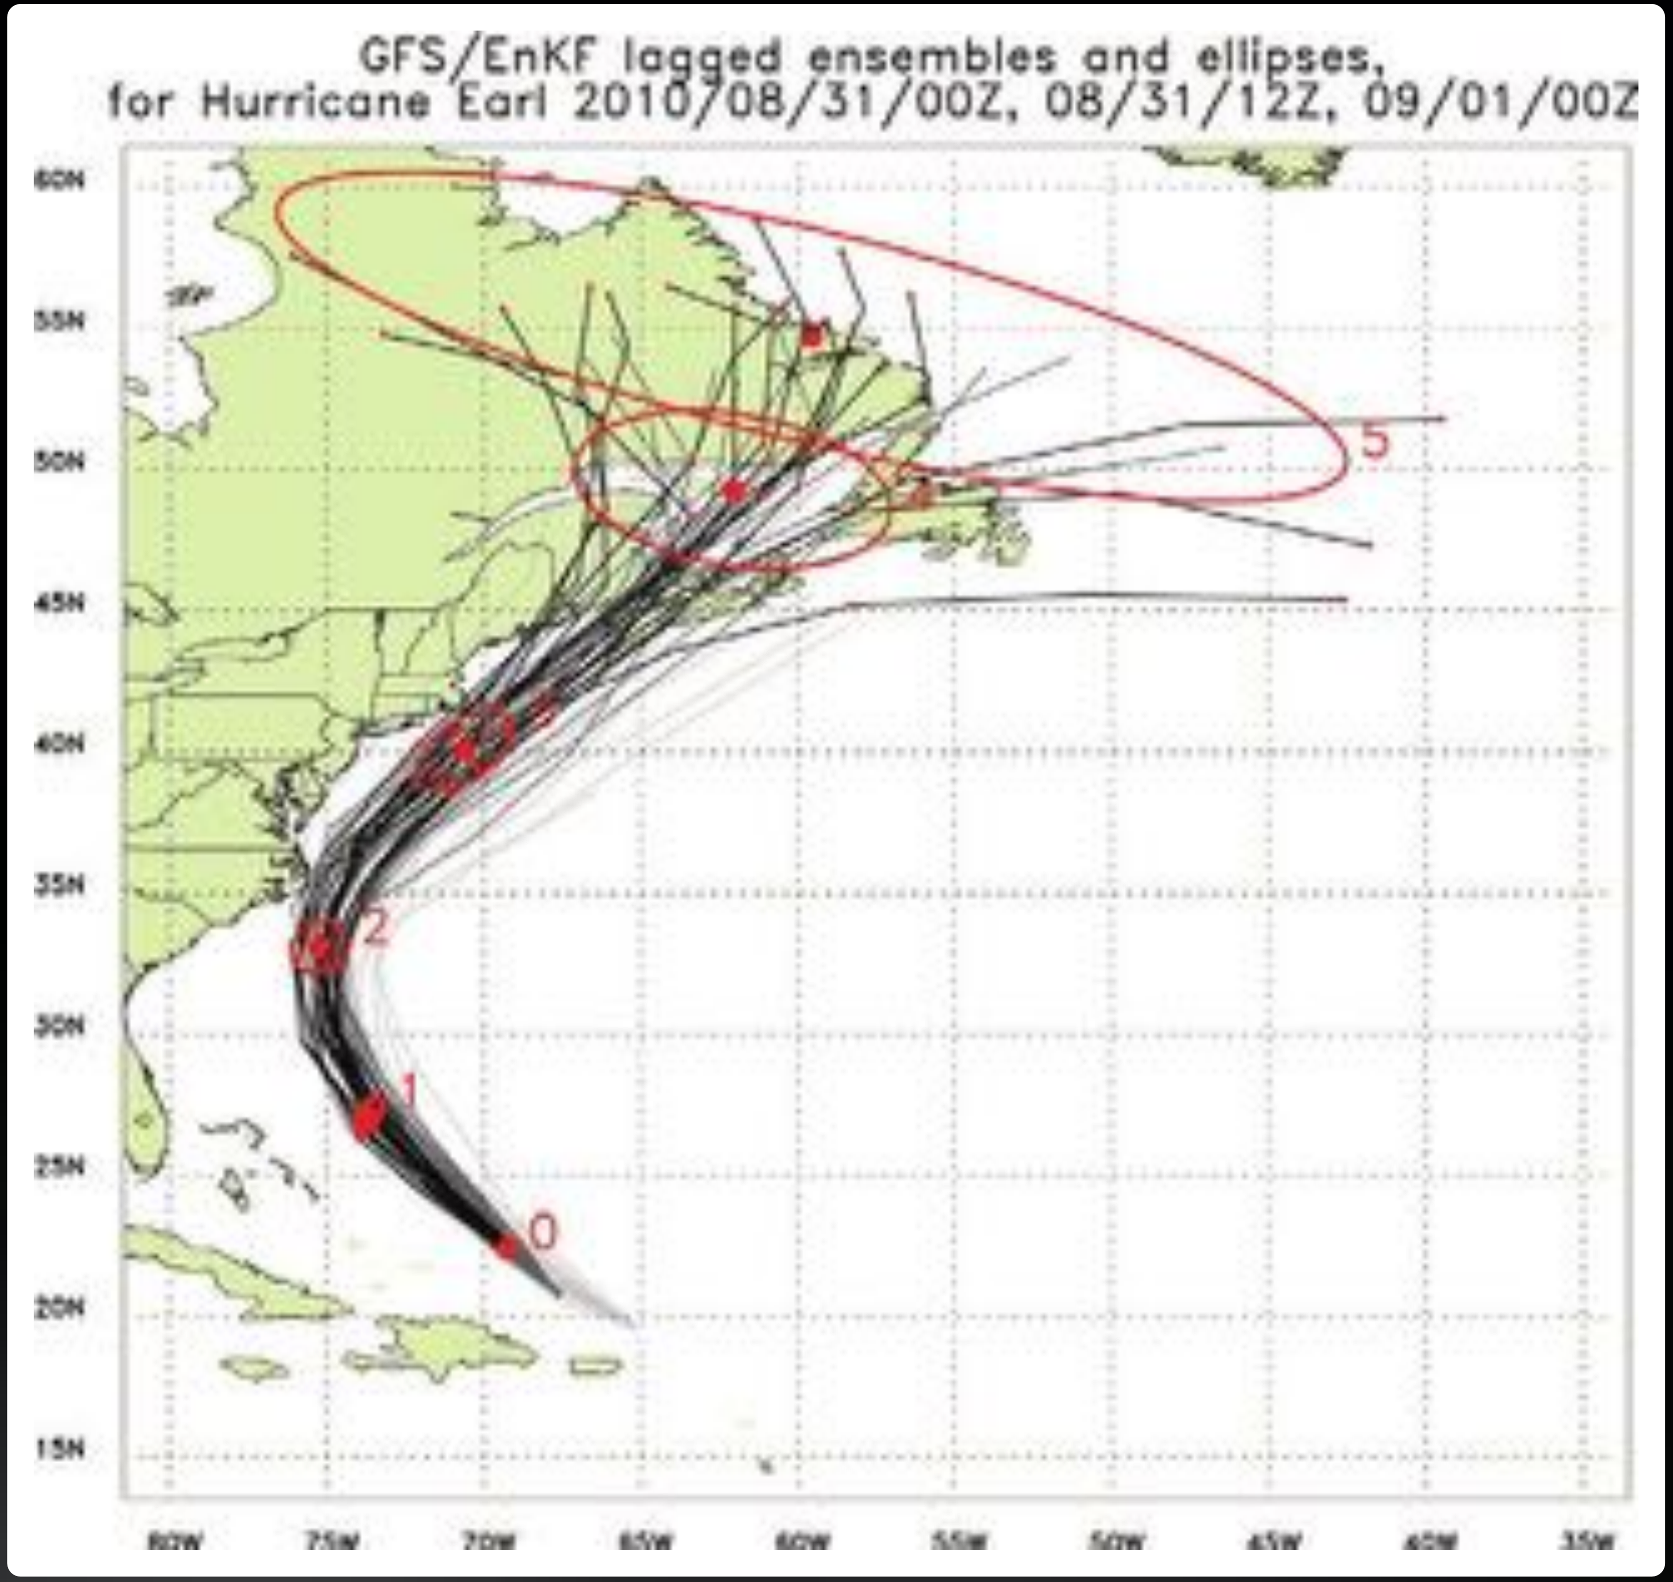

https://www.colorado.edu/faculty/claire-monteleoni/sites/default/files/attached-files/workshop_nips_2018_preprint_0.pdf

Detection

Planning

Acceleration

Assimilation

Enhancement

Parametrization

Prediction

Augmentation

Forecast Bias Correction

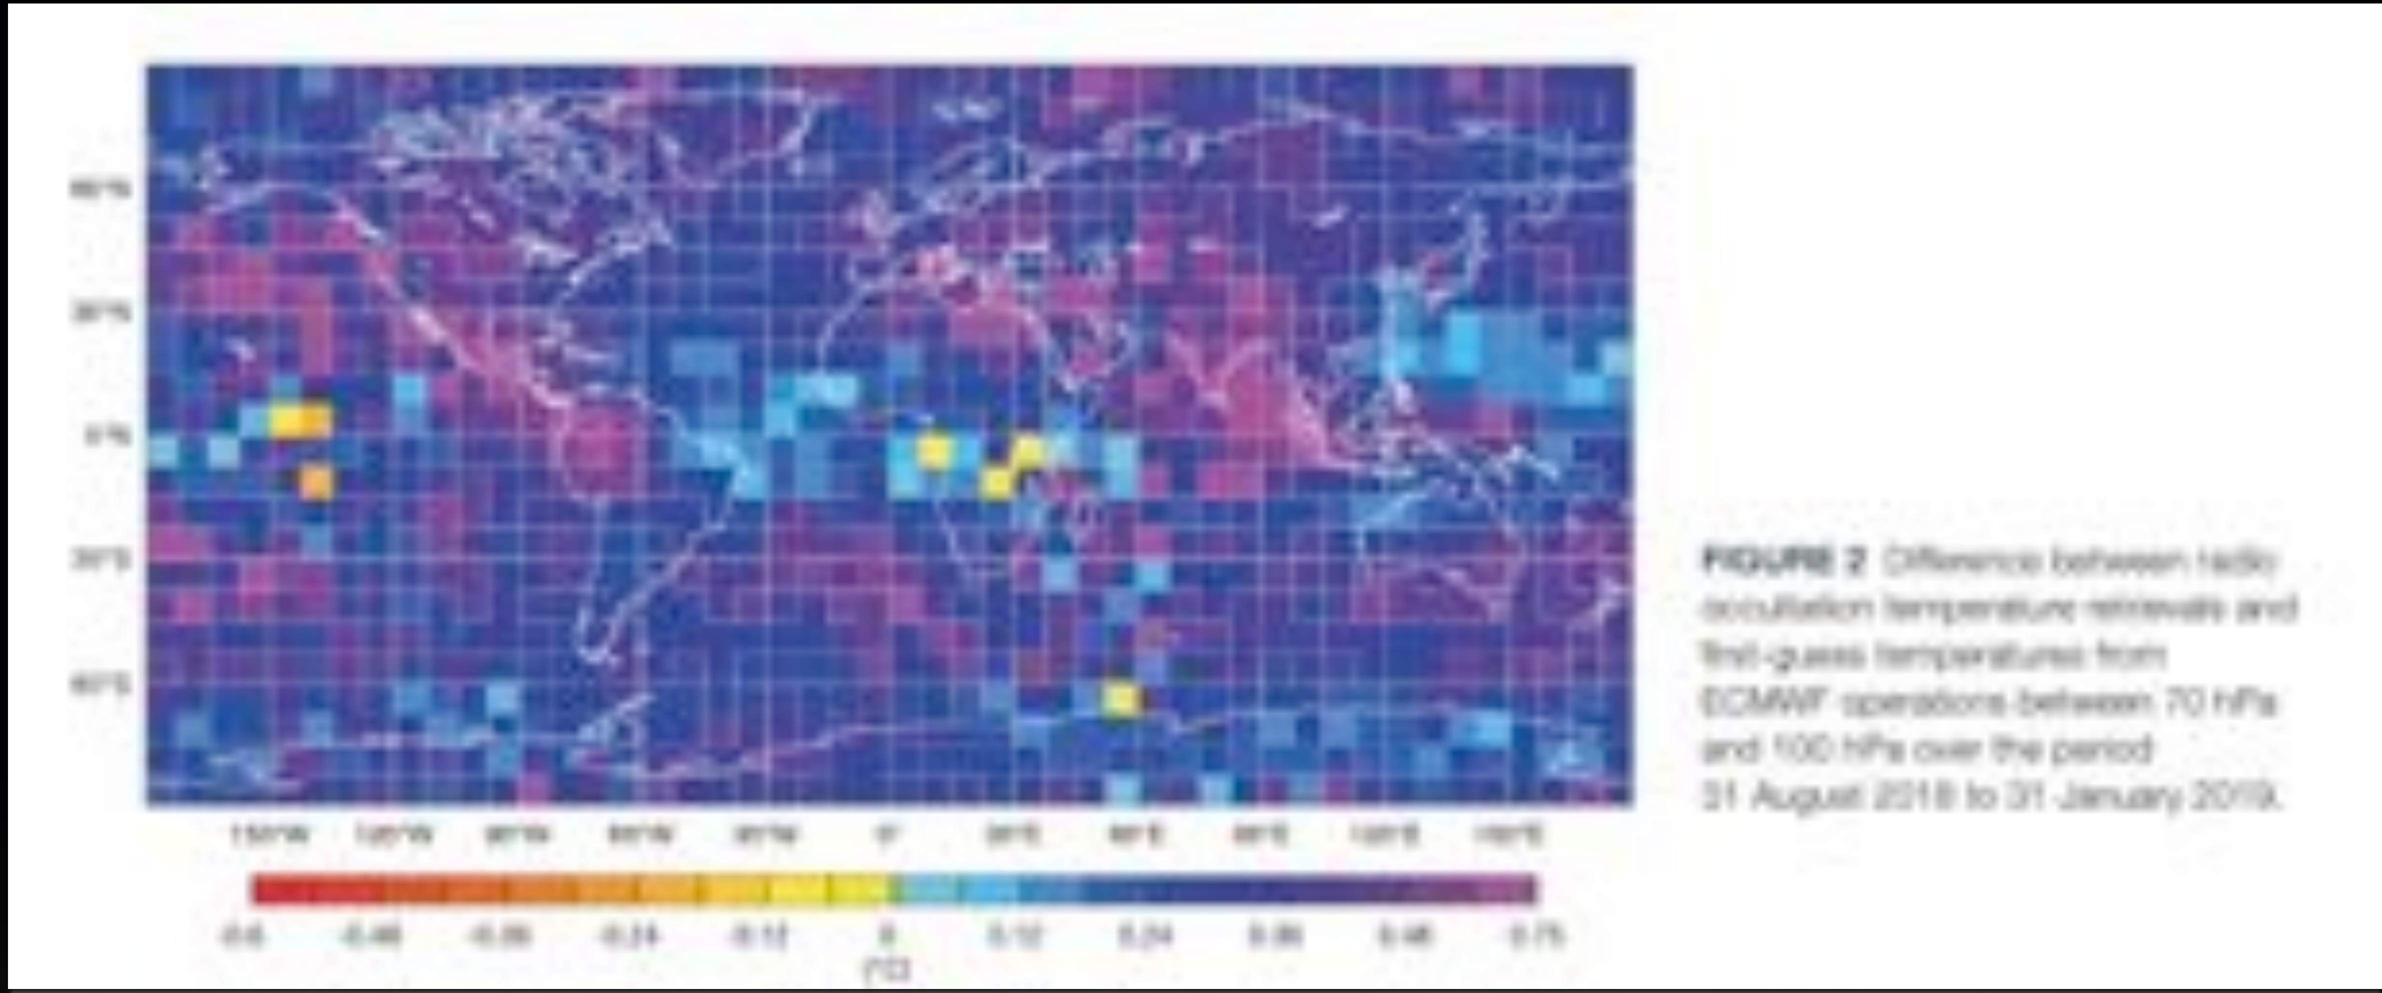

Improving the handling of model bias in data assimilation
Patrick Laloyaux, Massimo Bonavita Peter Dueben, Thorston Kurth, David Hall
https://www.ecmwf.int/sites/default/files/elibrary/2020/19508-newsletter-no-163-spring-2020.pdf

Detection

Planning

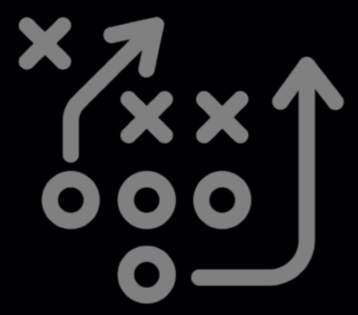

Acceleration

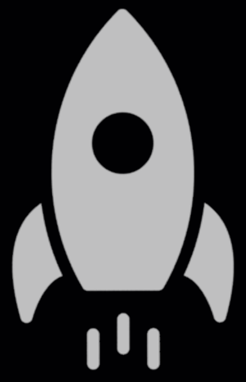

Assimilation

Enhancement

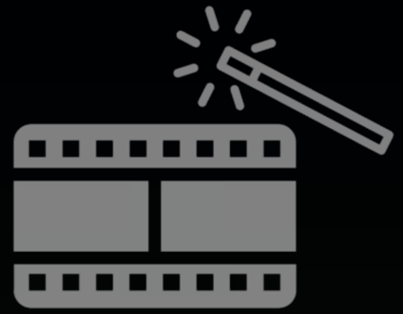

Parametrization

Prediction

Augmentation

AI powered Nowcasting

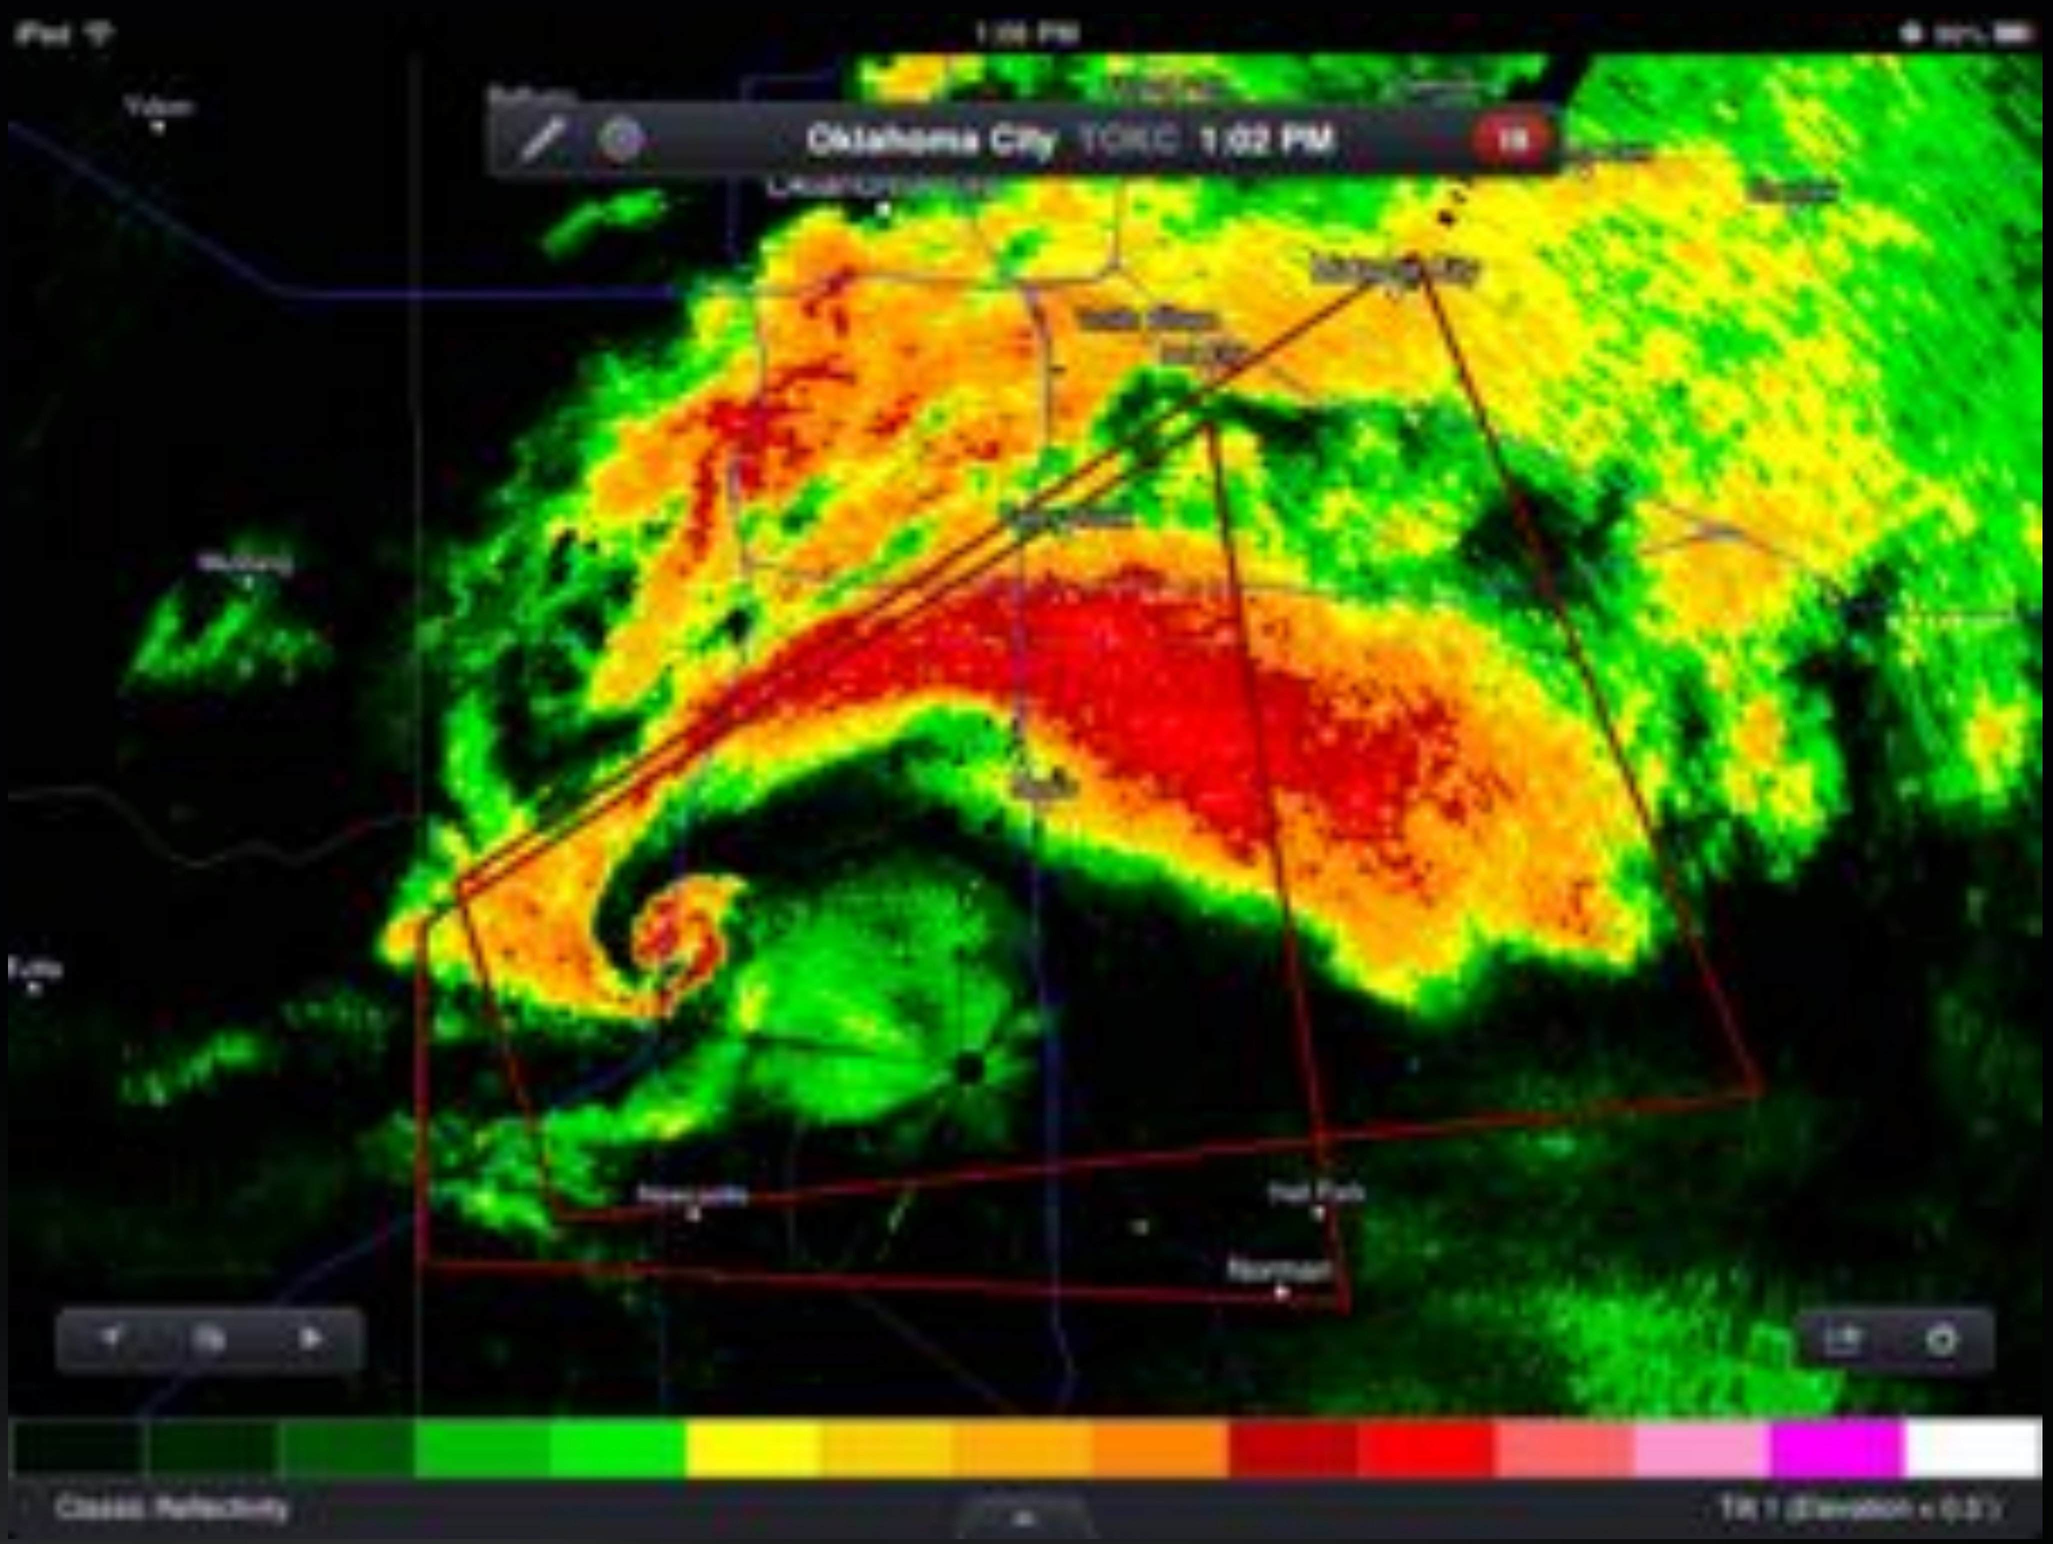

MetNet: A Neural Weather Model for Precipitation Forecasting
https://arxiv.org/abs/2003.12140

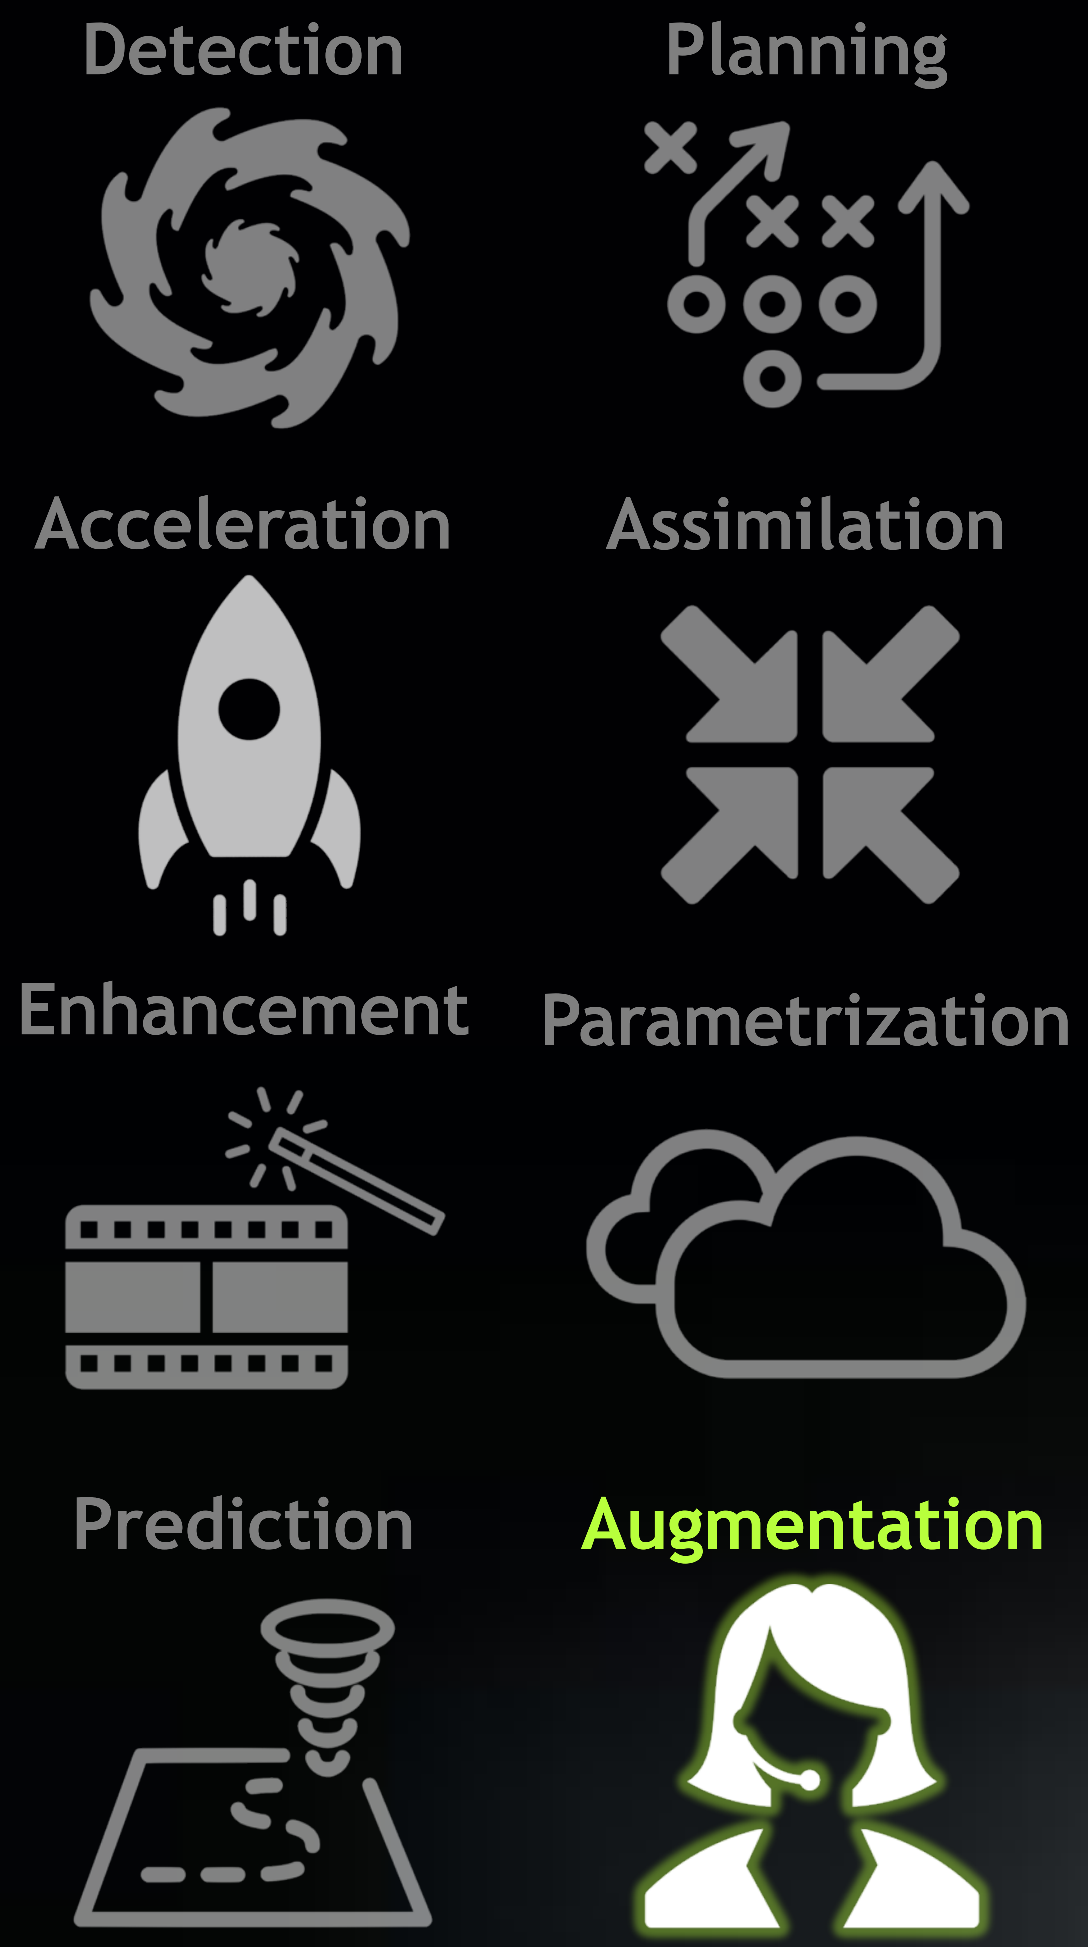

Augment your tools and get intelligent assistance

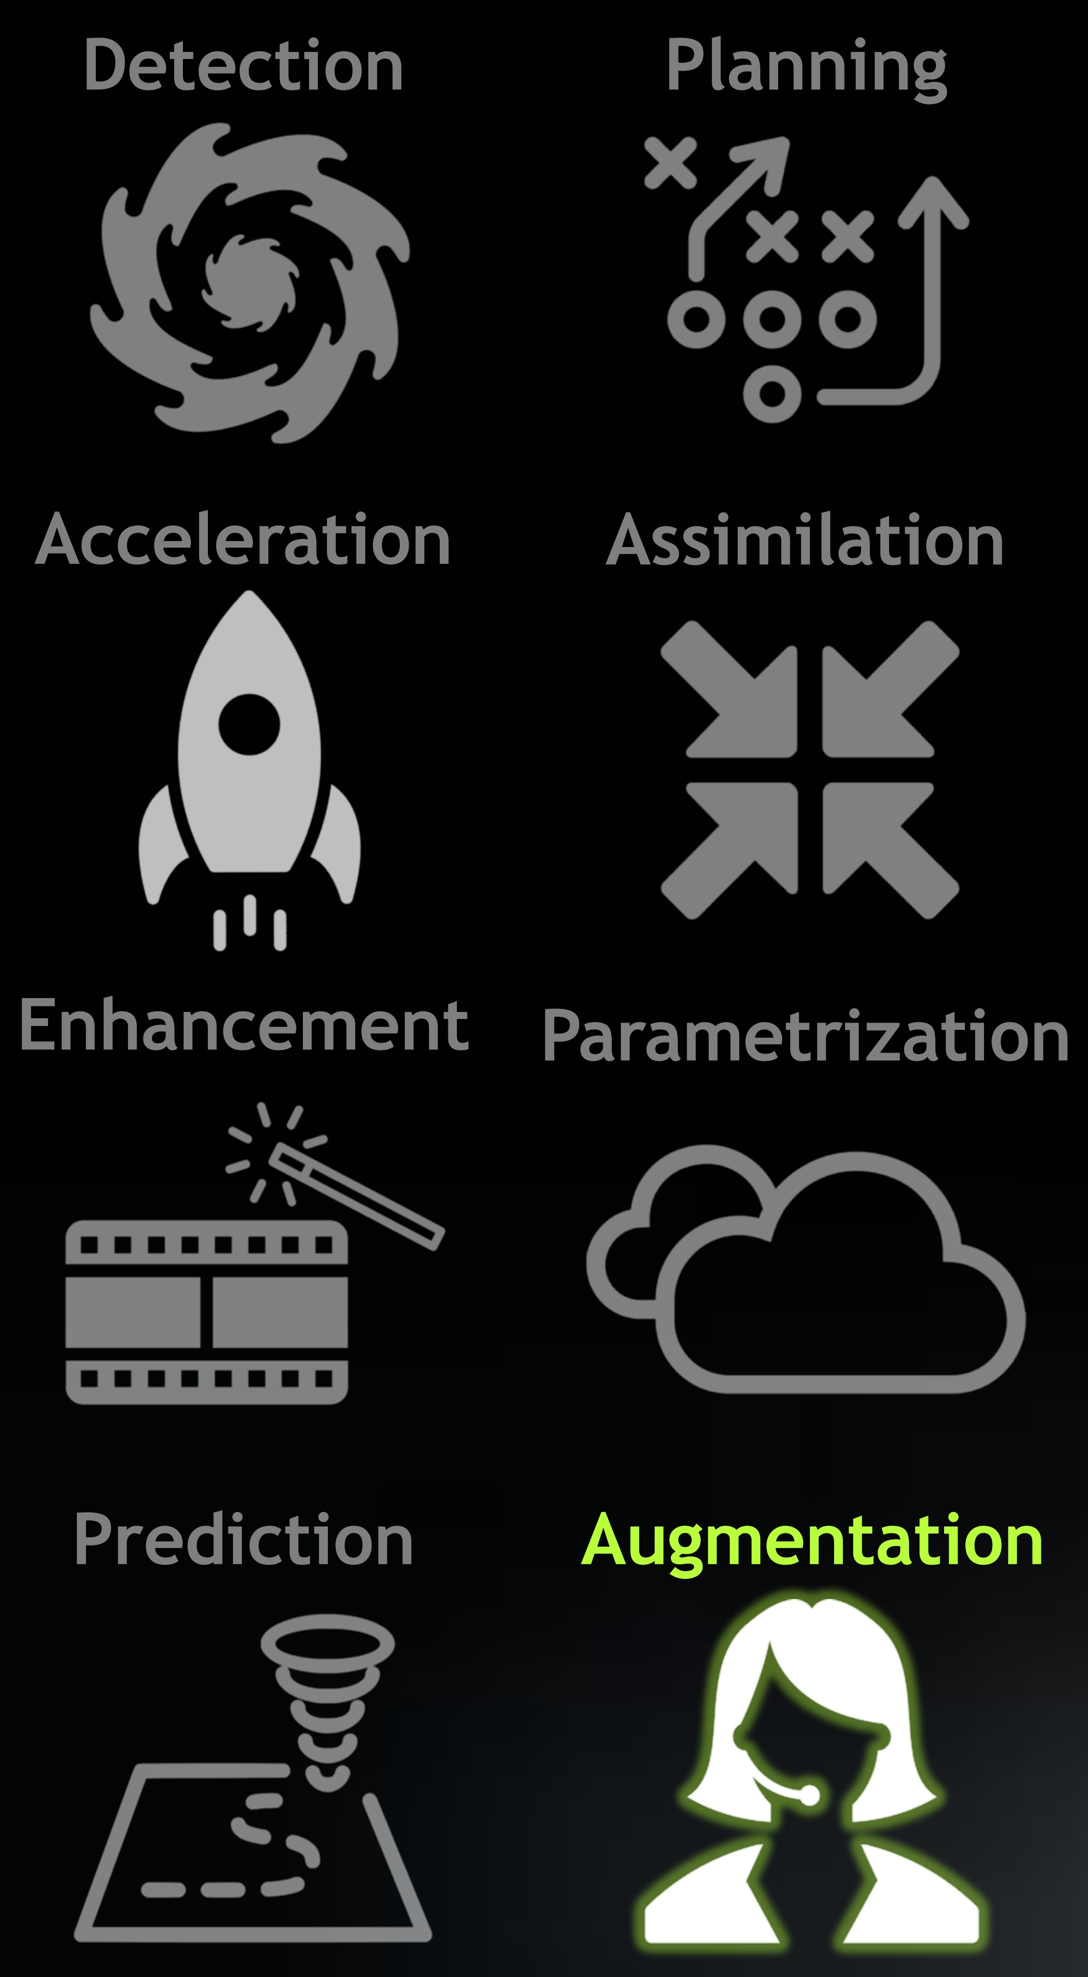

Build AI powered tools

```python
import os
import sys

# Count lines of code in the given directory, separated by file extension
def main(directory):
    line_count = {}
```

https://tabnine.com/blog/deep

Detection

Planning

Acceleration

Assimilation

Enhancement

Parametrization

Prediction

Augmentation

JARVIS

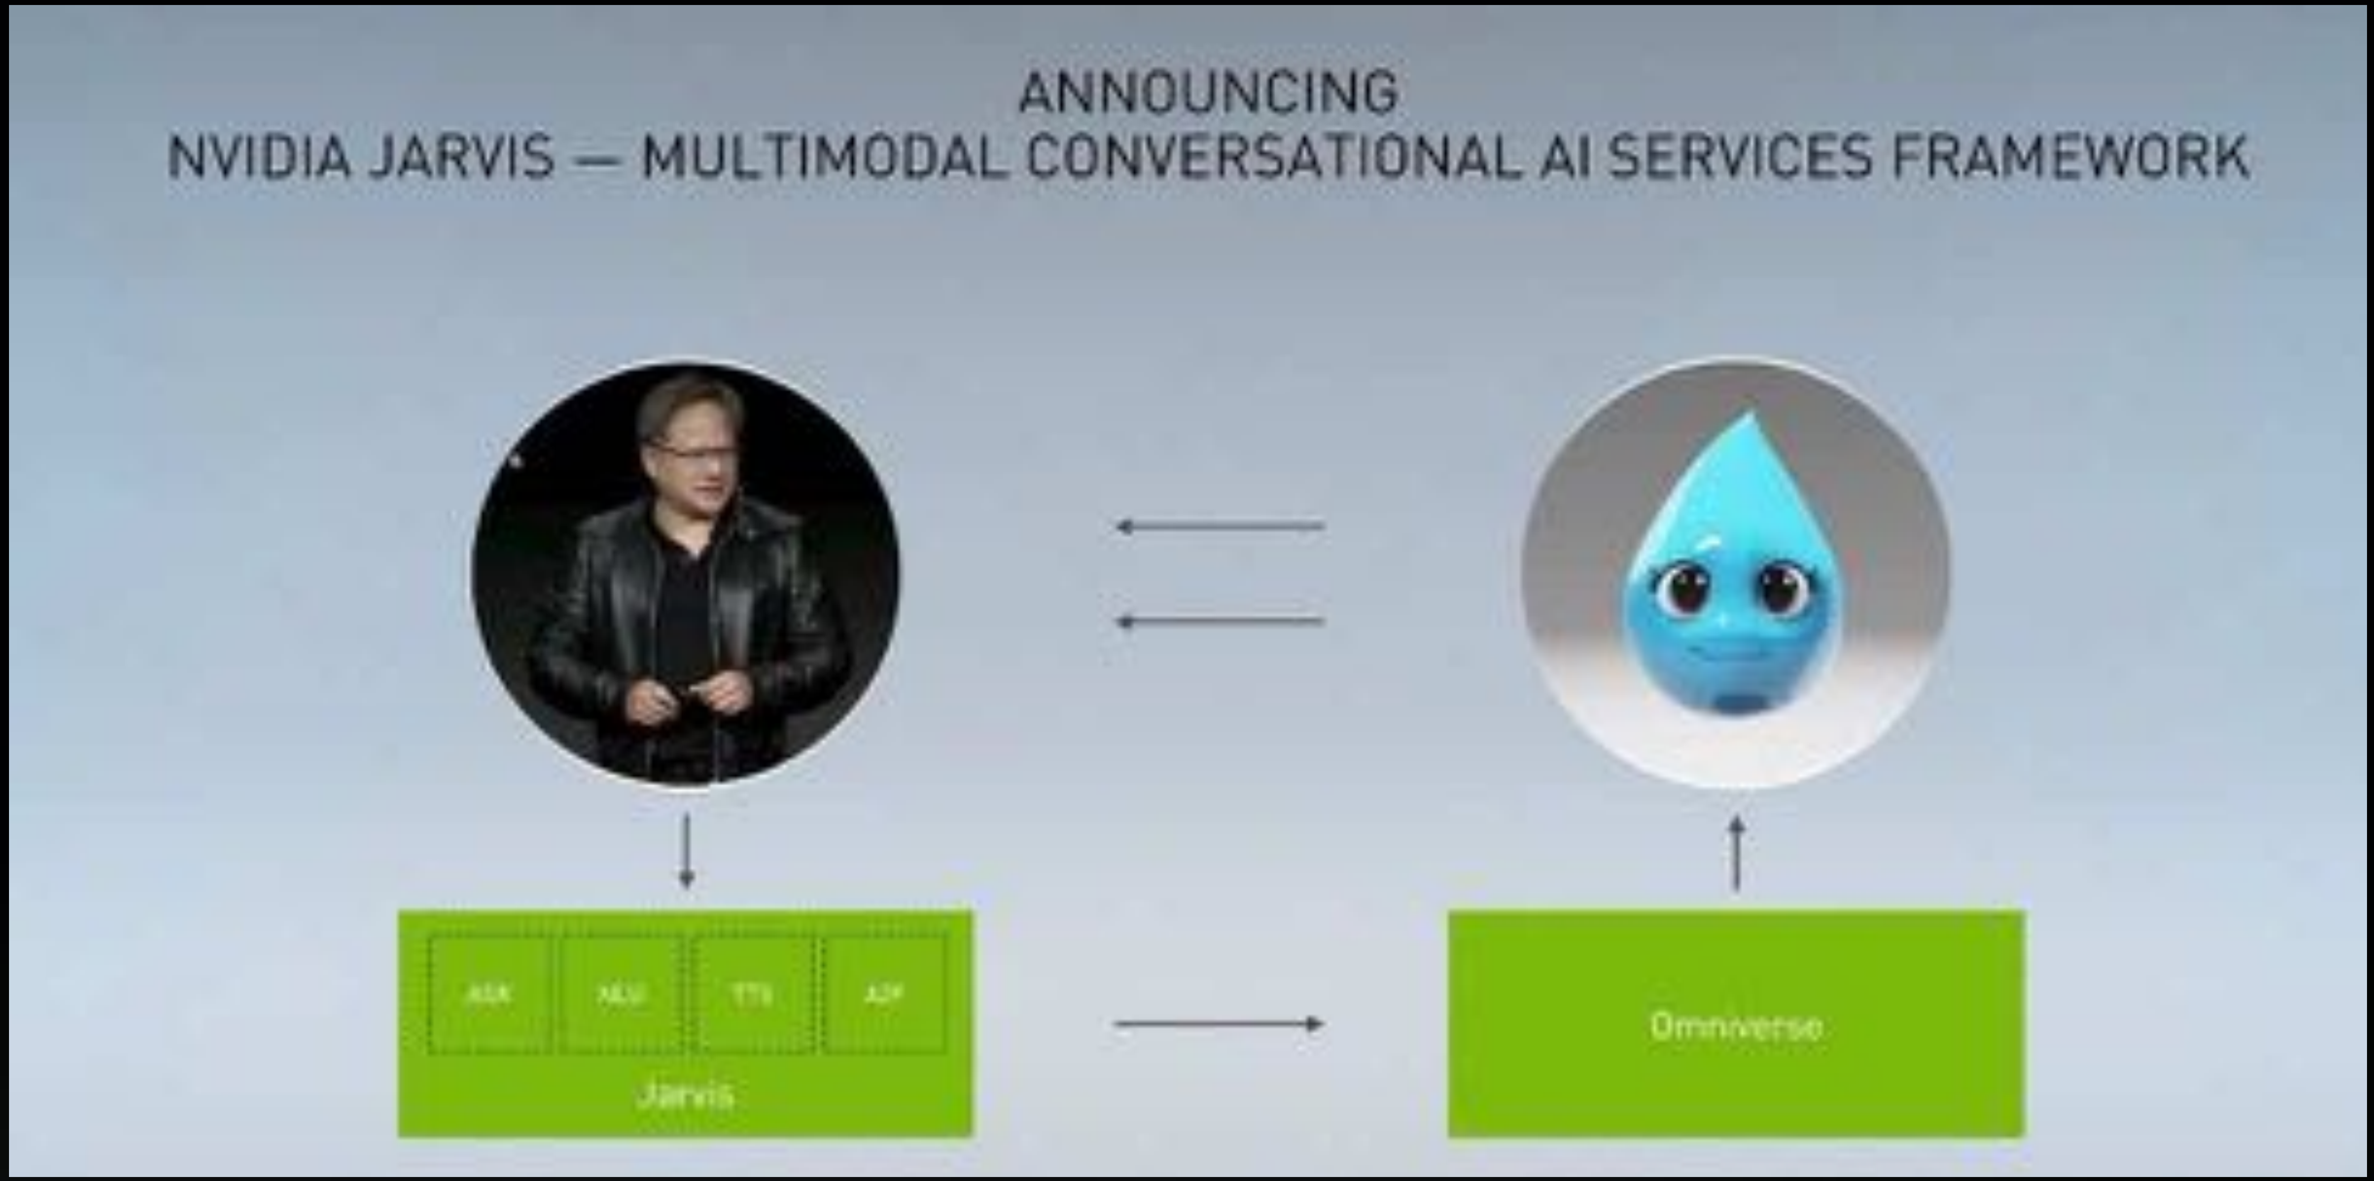

https://developer.nvidia.com/nvidia-jarvis

CHALLENGES AND POTENTIAL SOLUTIONS

LABELLING LARGE QUANTITIES OF DATA

How can we overcome the need for manual labelling?

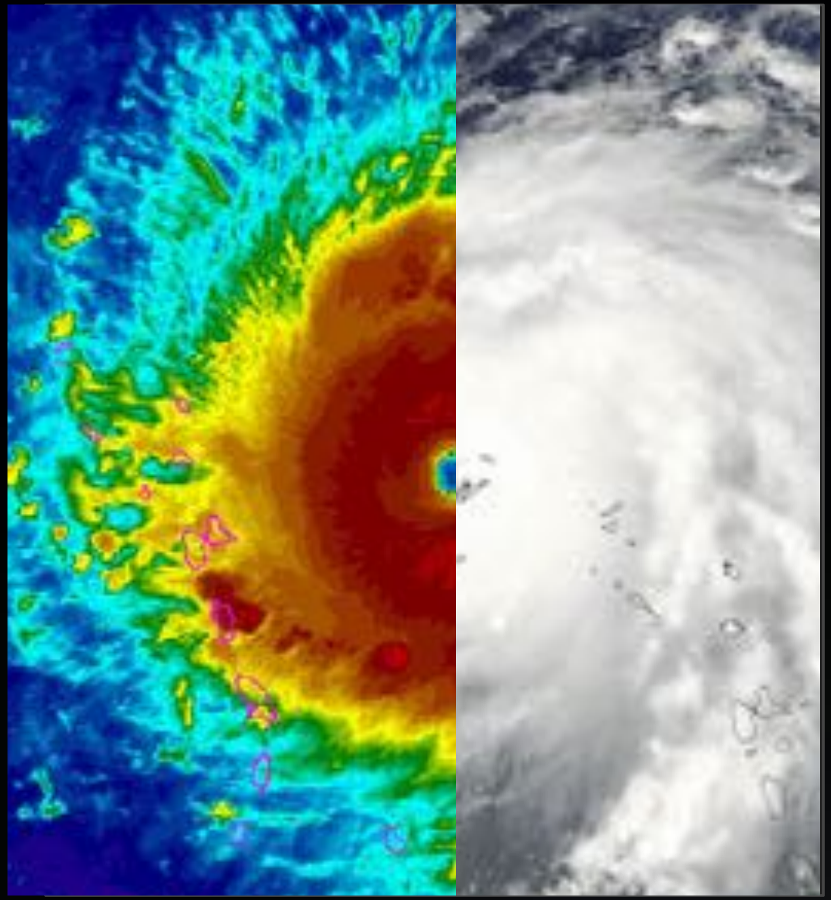

Data Fusion
Using one data source as the label for another

Self-Supervised Learning
Predicting input B from input A

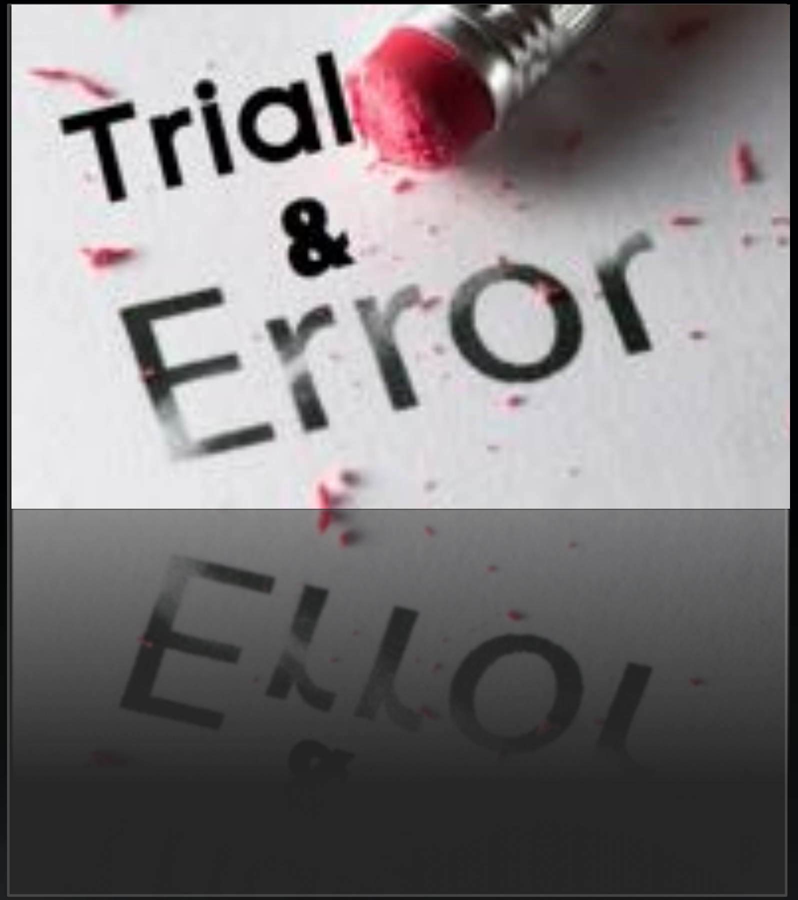

Reinforcement Learning
Obtaining labels directly from the environment or simulation

Human-in-the-loop
Using human machine iteration to make labelling easier

TRANSFER LEARNING: DON'T START FROM SCRATCH

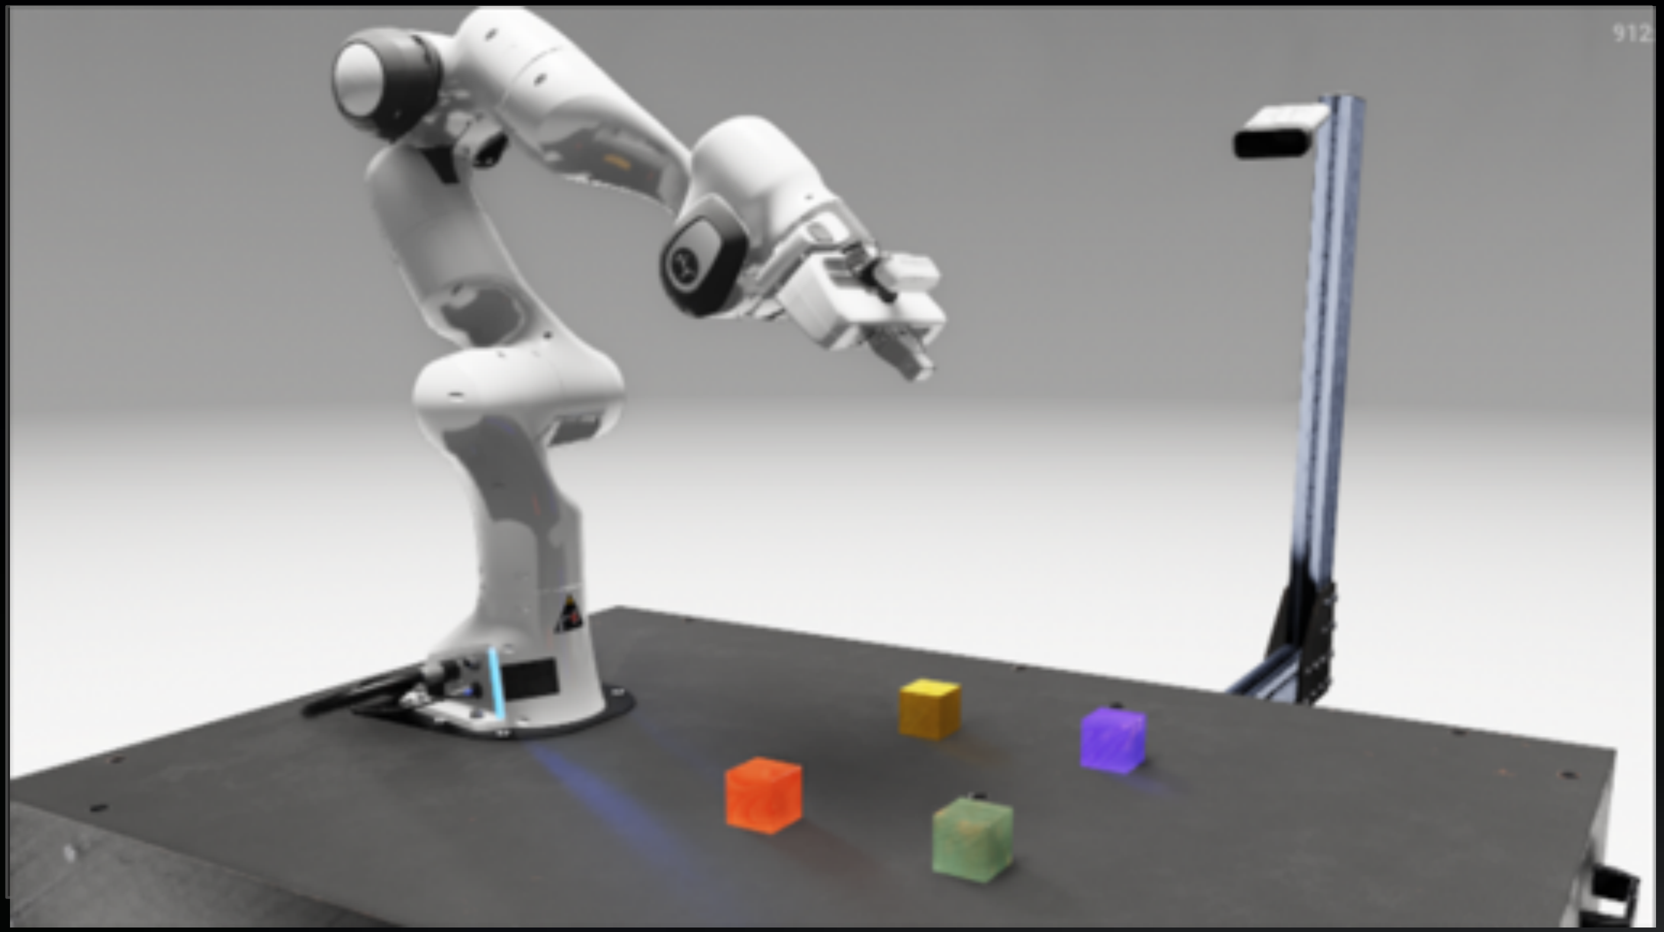

Train on simulated or related data

Fine-tune on the real data

BENCHMARKS: THE NEED FOR A COMMON GOAL

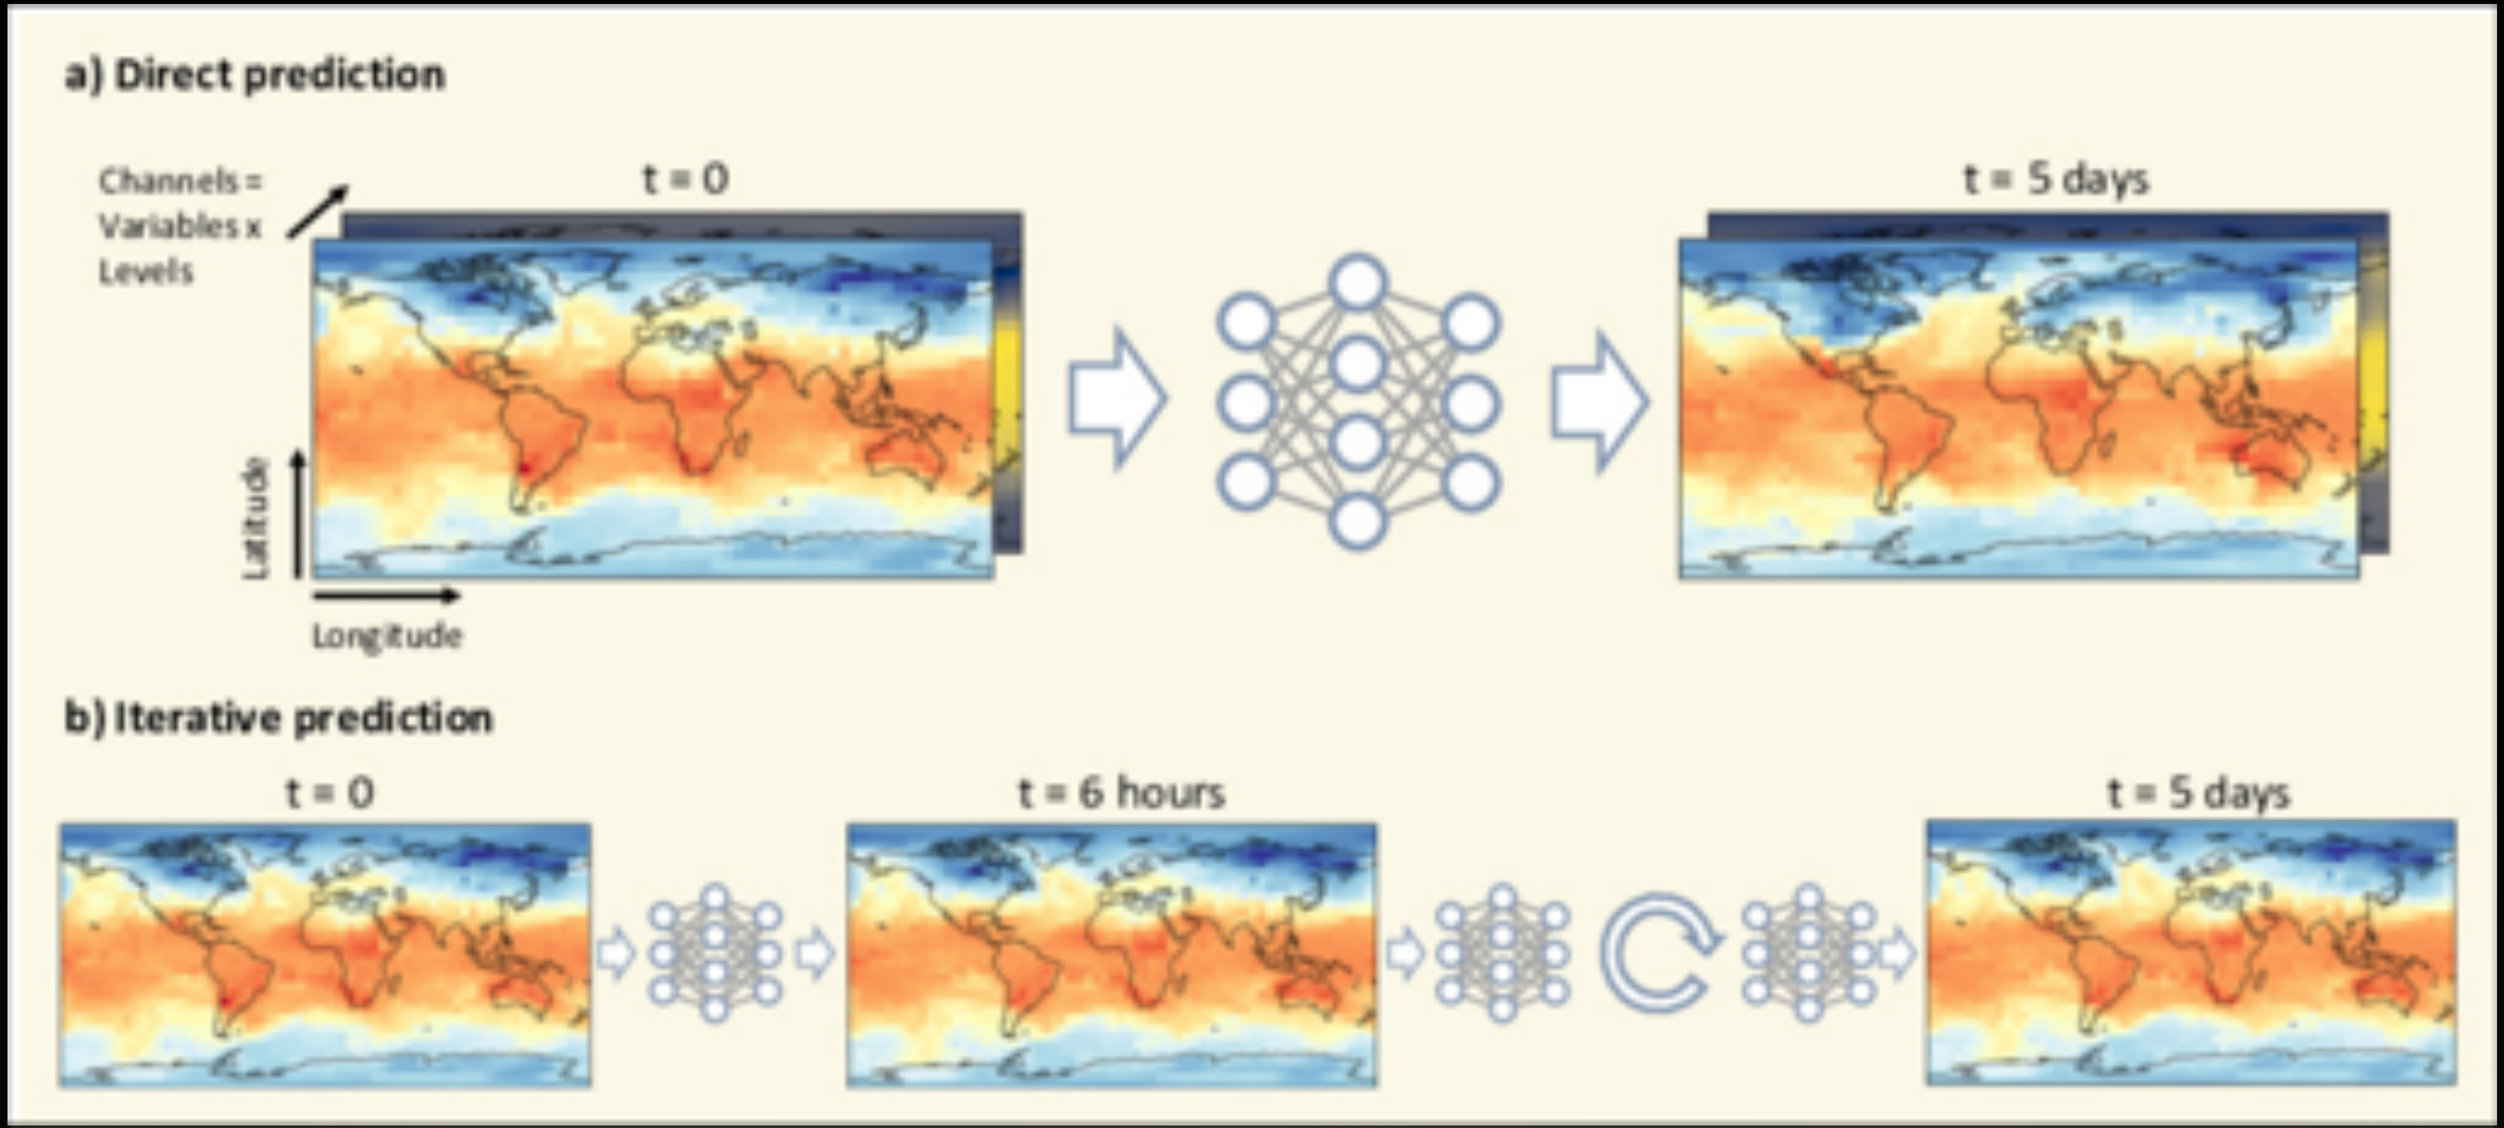

WeatherBench: A benchmark dataset for data-driven weather forecasting
Stephan Rasp, Peter D. Dueben, Sebastian Scher, Jonathan A. Weyn, Soukayna Mouatadid, Nils Thuerey
https://arxiv.org/abs/2002.00469

ENFORCING PHYSICAL CONSTRAINTS

Conservation of Mass, Momentum, Energy, Incompressibility, Turbulent Energy Spectra, Translational Invariance

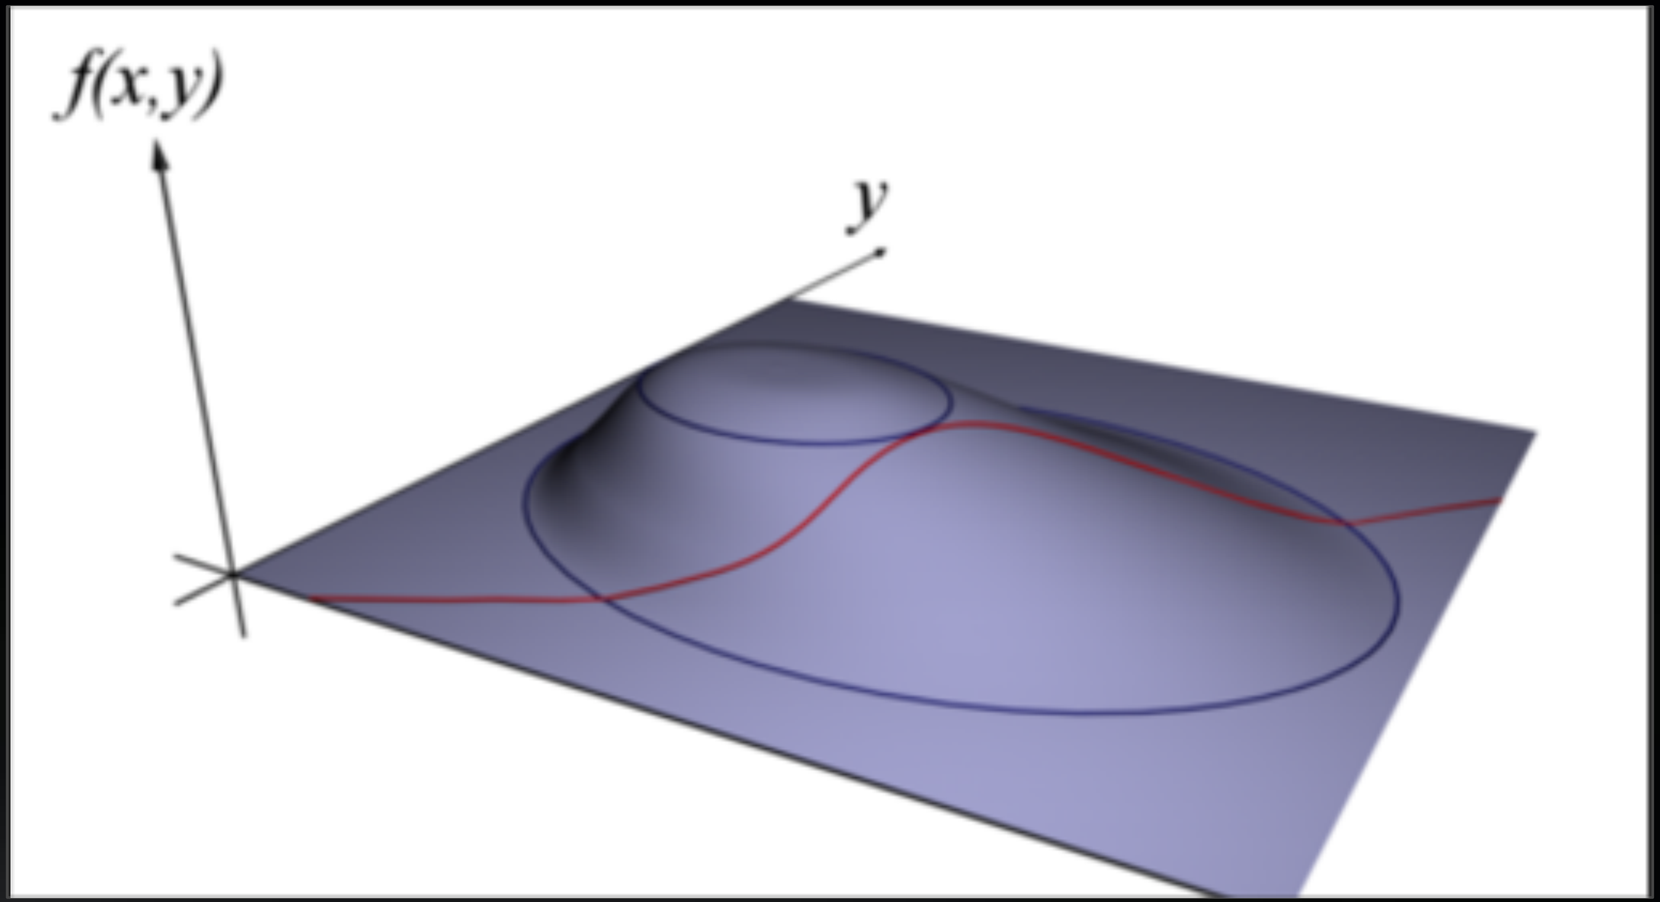

Lagrange multipliers (penalization), Hard Constraints, Projective Methods, Differentiable Programming

INTERPRETABILITY: EXPLAINABLE AI

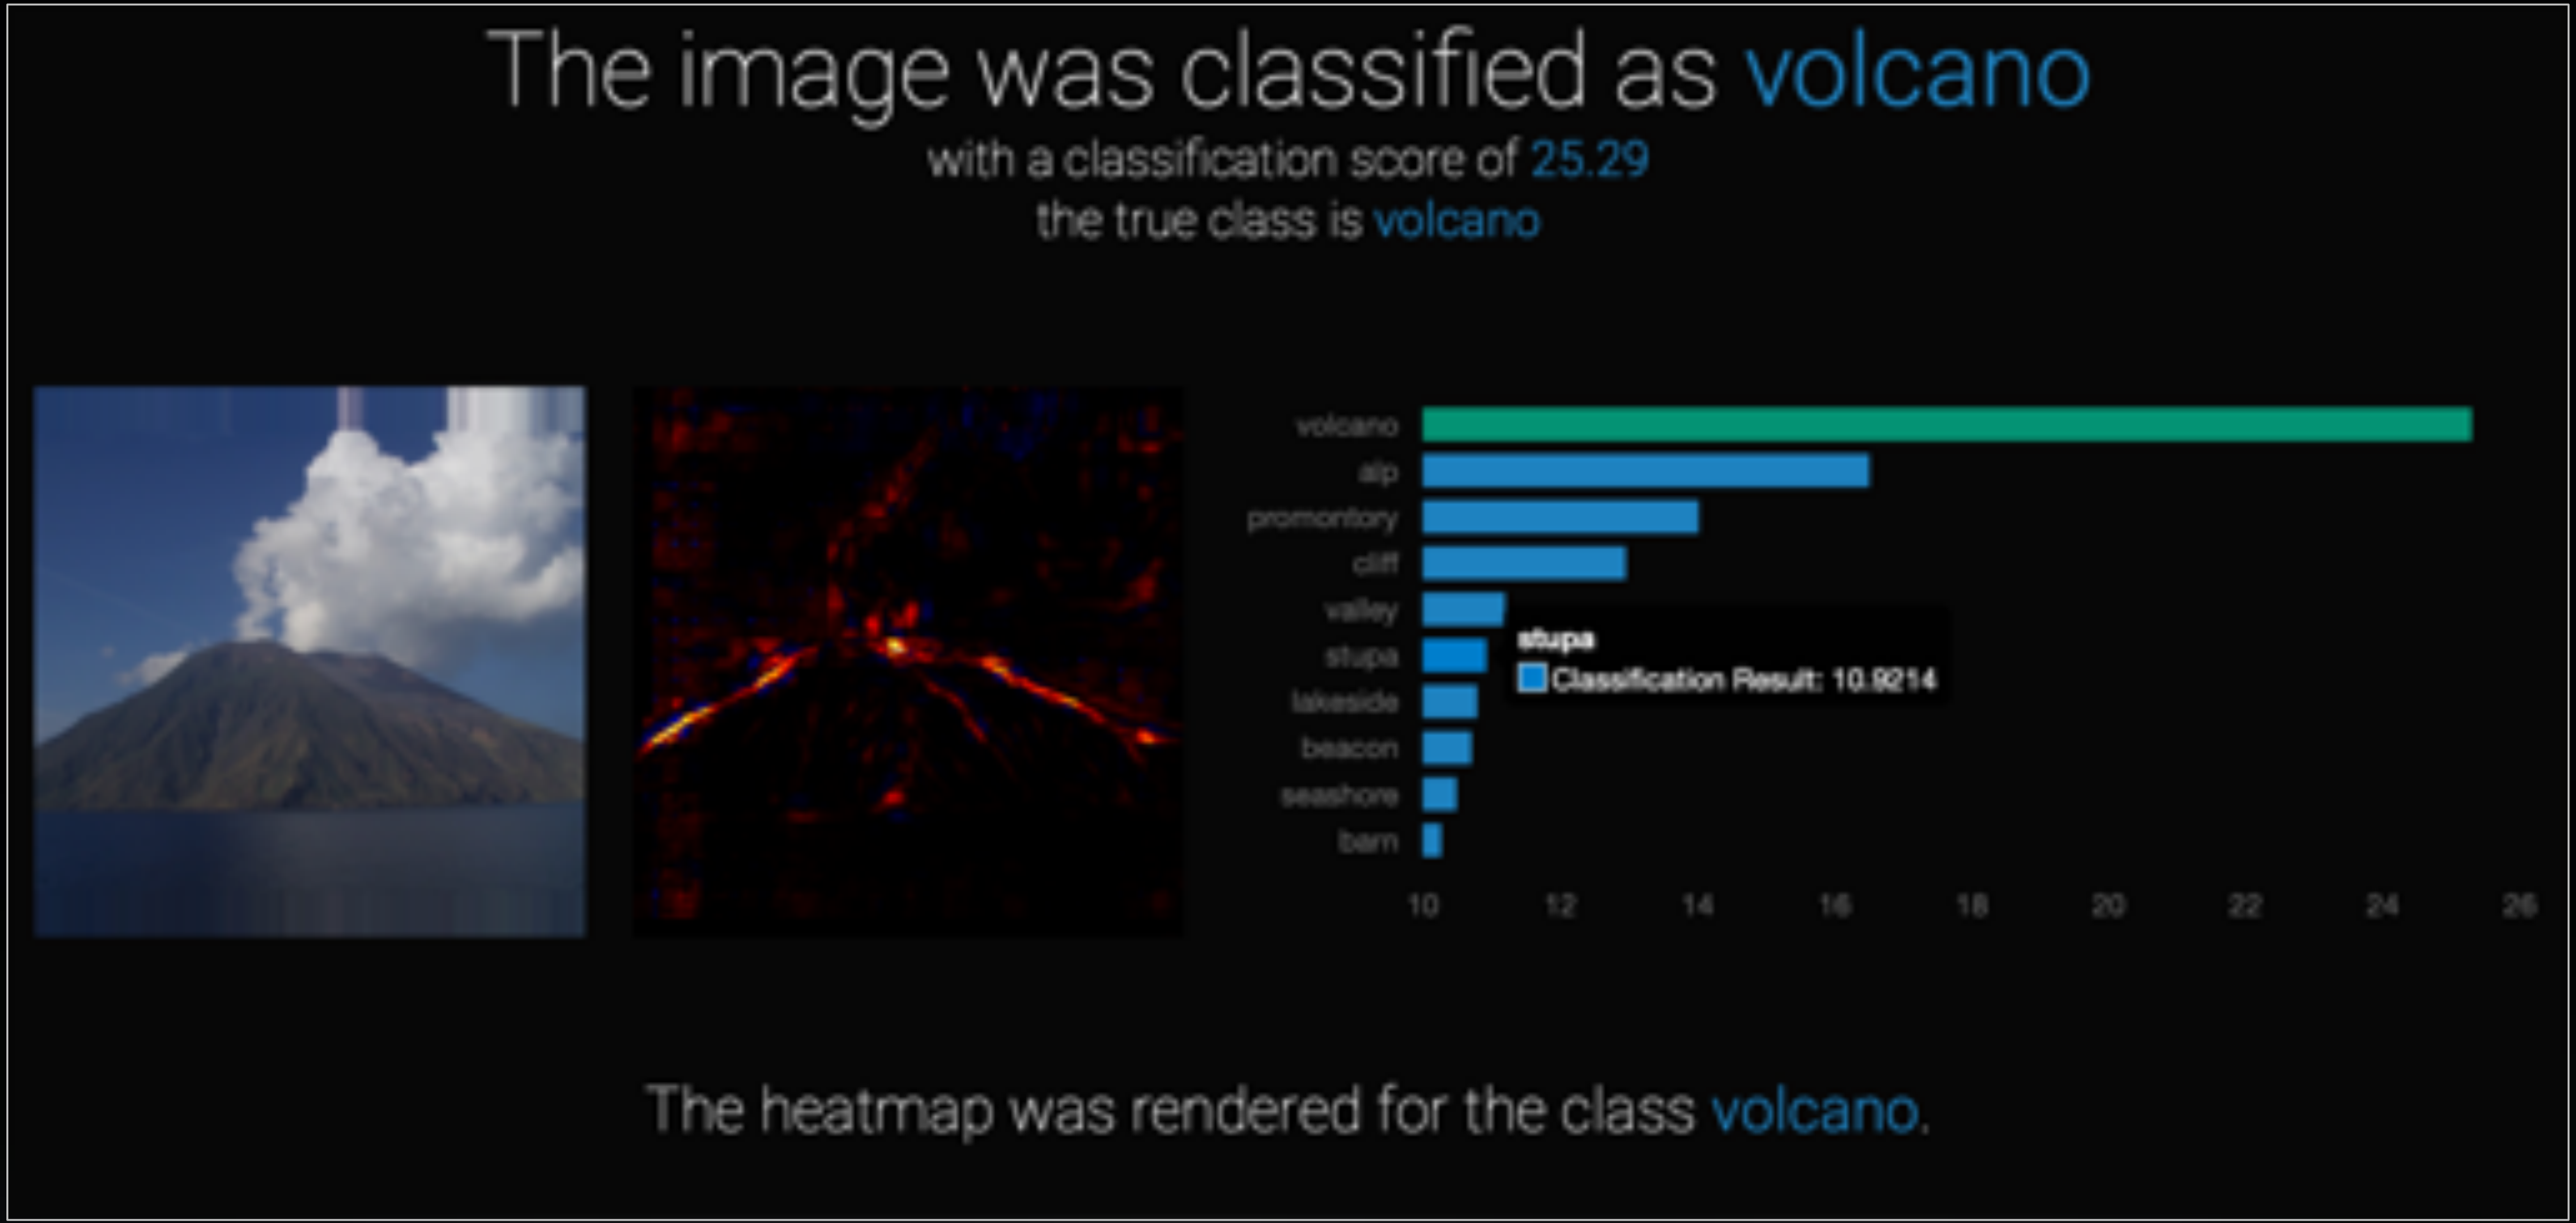

Layer-wise Relevance Propagation

https://lrpserver.hhi.fraunhofer.de/image-classification

UNCERTAINTY ESTIMATION

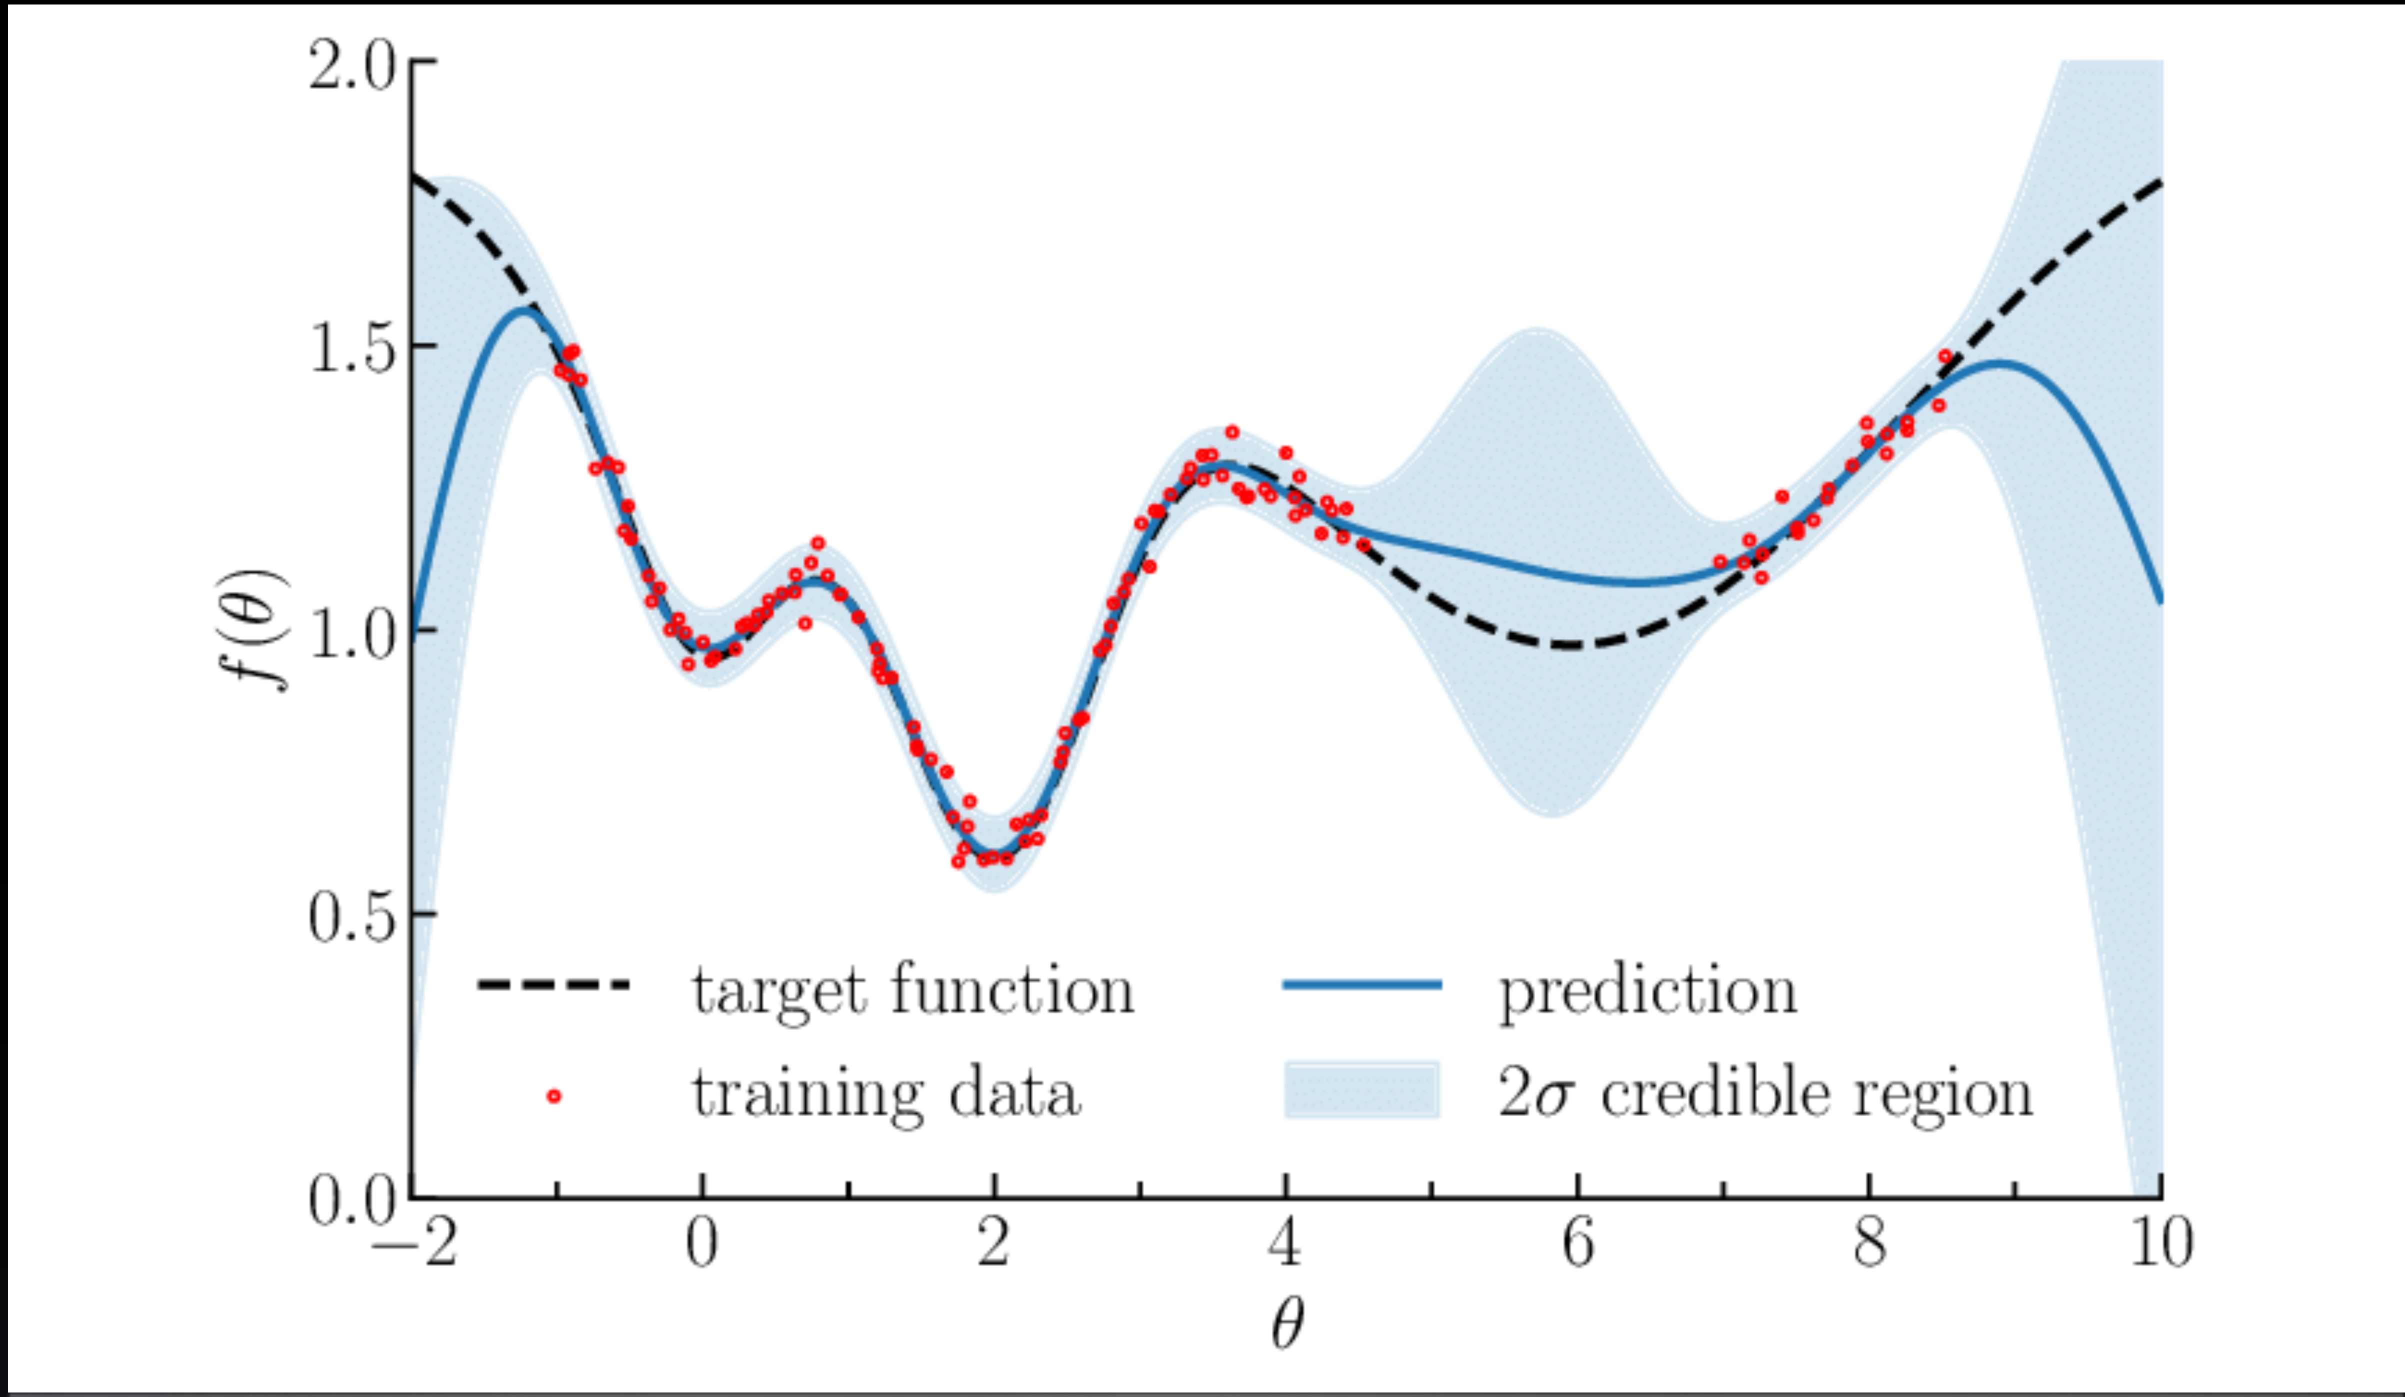

https://www.researchgate.net/figure/Illustration-of-Gaussian-process-regression-in-one-dimension-for-the-target-test_fig1_327613136

- ML tools provide a powerful new way to build software
- I expect many breakthroughs will from this direction in the near future.
- GPUs makes ML practical, while ML makes GPUs more accessible.
- ML of tomorrow might be radically different than today. Tools and hardware are evolving rapidly.
- Challenges exist. These tools are new. But we have barely scratched the surface of their potential.

dhall@nvidia.com

NVIDIA