LAMARR

INSTITUTE FOR MACHINE LEARNING AND ARTIFICIAL INTELLIGENCE

Efficient Communication between AI and Domain Experts

How to?

Jens Buß & Ramsés Sanchéz & Sebastian Buschjäger

Partner institutions:

technische universität dortmund

Fraunhofer IAIS

UNIVERSITÄT BONN

Fraunhofer IML

Institutionally funded by:

Bundesministerium für Bildung und Forschung

Ministerium für Kultur und Wissenschaft des Landes Nordrhein-Westfalen

The Players

Domain Expert:

Astrophysicist

aiming to use ML to achieve analysis objective!

AI Expert:

Computer Scientist

looking for an interesting Use-case for ML

Typical Approach in Science

Specific Problem

- A trader wants to transport her goods from A to B

Mathematical Modelling

- We map streets and cities to a Graph

Mathematical Algorithm

- We formulate an optimization problem on top of the Graph
- We search for an optimization algorithm

Implementation

- We implement the algorithm in a concrete system (e.g. Google Maps)

AI: Same general approach but with extra steps!

Typical Approach in AI

Fact: Some problems are difficult to describe mathematically

But: For some problems it is easy to give examples (= data) of what we want

Specific Problem
- Distinguish between a cat and not a cat

Gather Data
- Search a lot of pictures with cats and other animals

Formulate a Learning Task
- Given a picture, predict if there is a cat or not

Model + Optimization
- Choose a model class (linear, tree, deep learning, …)
- Formulate an optimization algorithm that finds the optimal model for the data and tasks

Implementation
- Implement the optimization algorithm in a computer and let it search for the best model given the data

Downstream task

AI Training

Model Deployment

Typical Approach in AI – A closer look

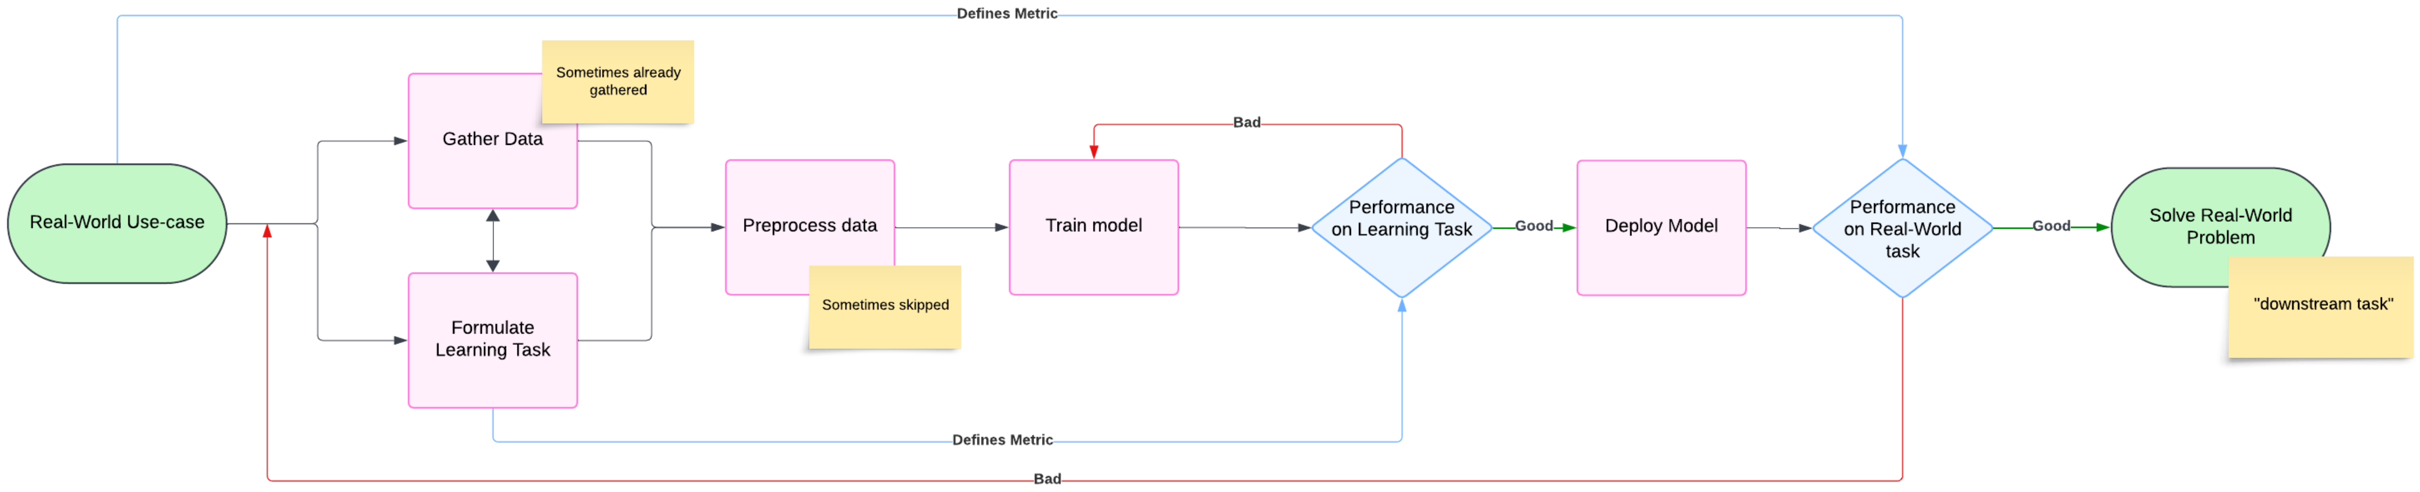

Efficient Communication between AI and domain experts

Obviously
- A domain expert who knows the entire AI pipeline does not need an AI expert
- An AI expert who knows the entire real-world use-case does not need a domain expert

But clearly: Nobody knows everything

So, what should an AI expert tell the domain expert and vice-versa?

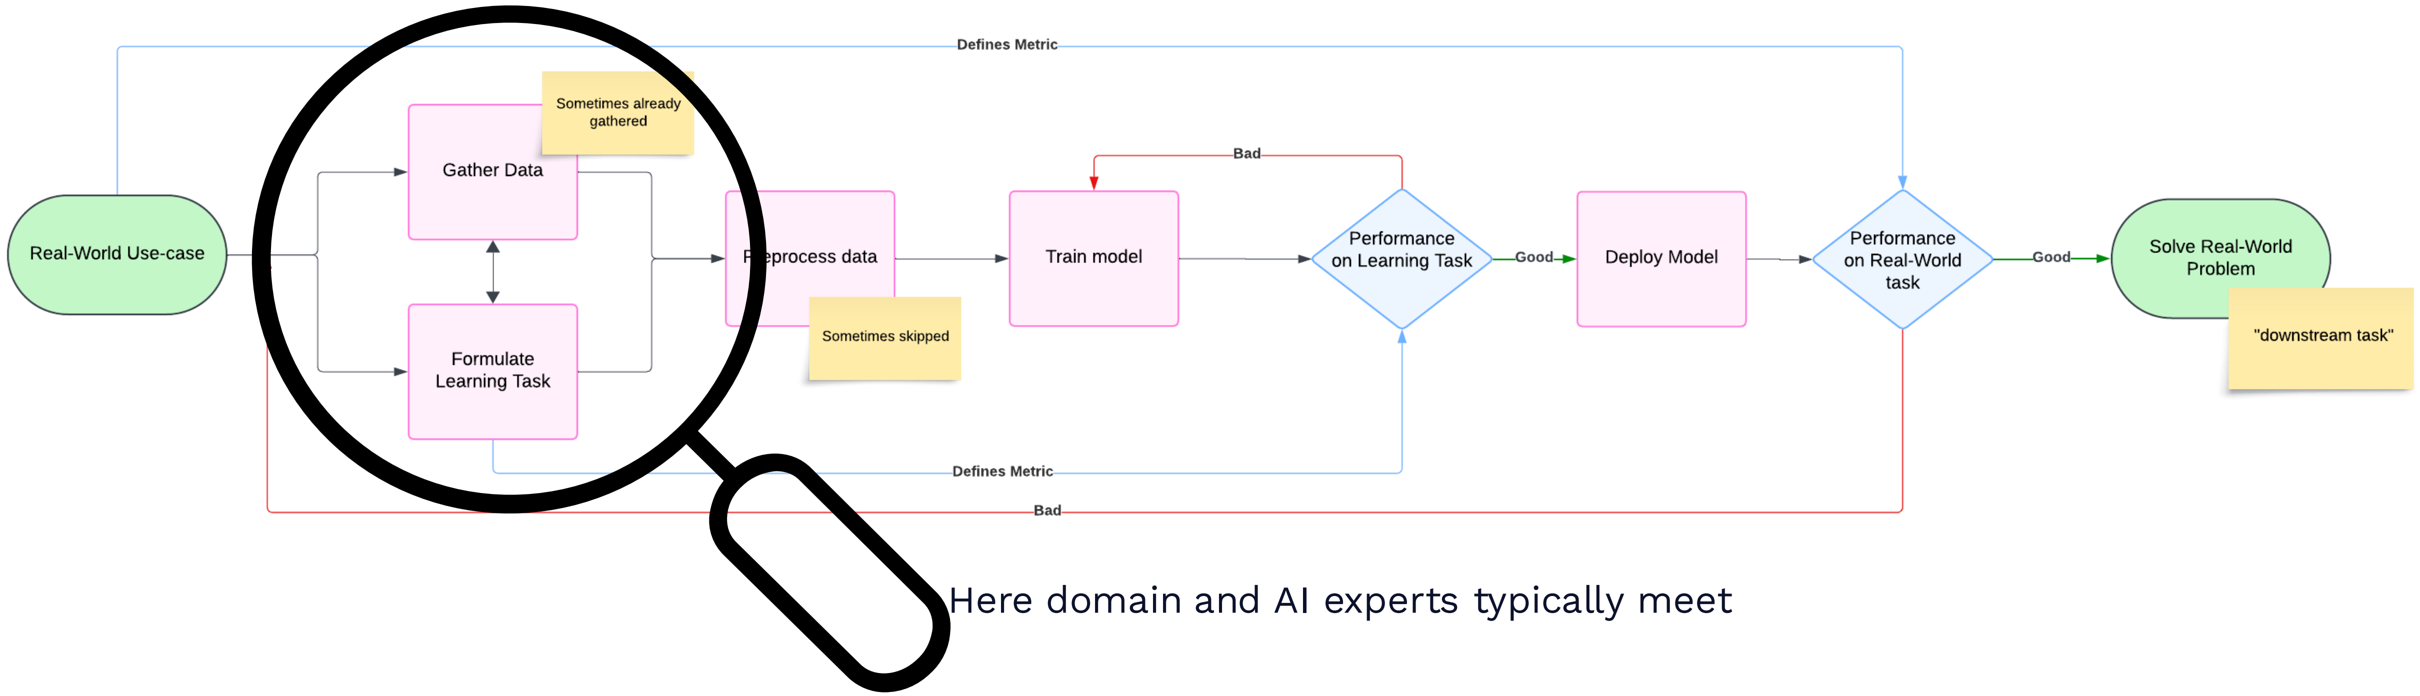

Here domain and AI experts typically meet


Domain Expert:
Astrophysicist
AI Expert:
Computer Scientist
Black Hole

Which information are needed for the players to collaborate?

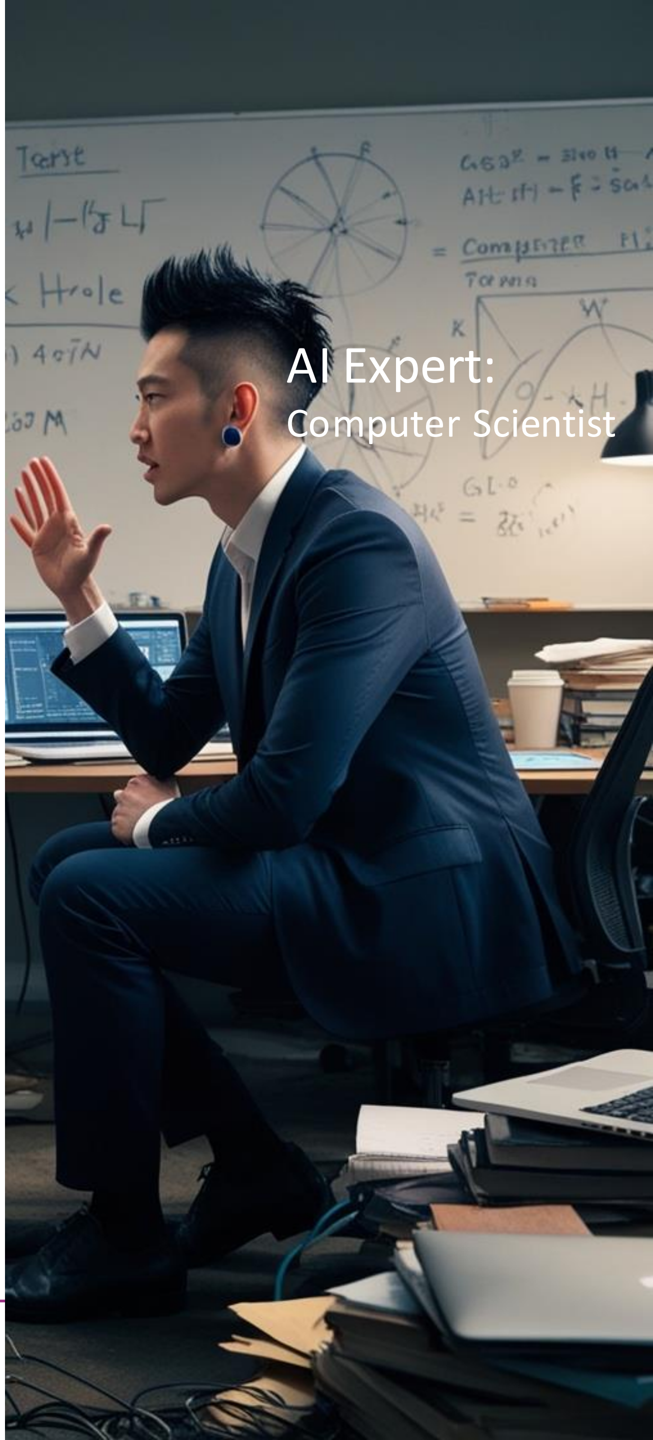

Which information are needed for the players to collaborate?

- Real-World use-case (Motivation)
- Performance metric for real-world task
- Downstream Task
- Performance metric for real-world task
- Data type: e.g. time series, images, tabular
- Learning Task
- Ground Truth / Labels
- Applicable learning algorithms and their properties

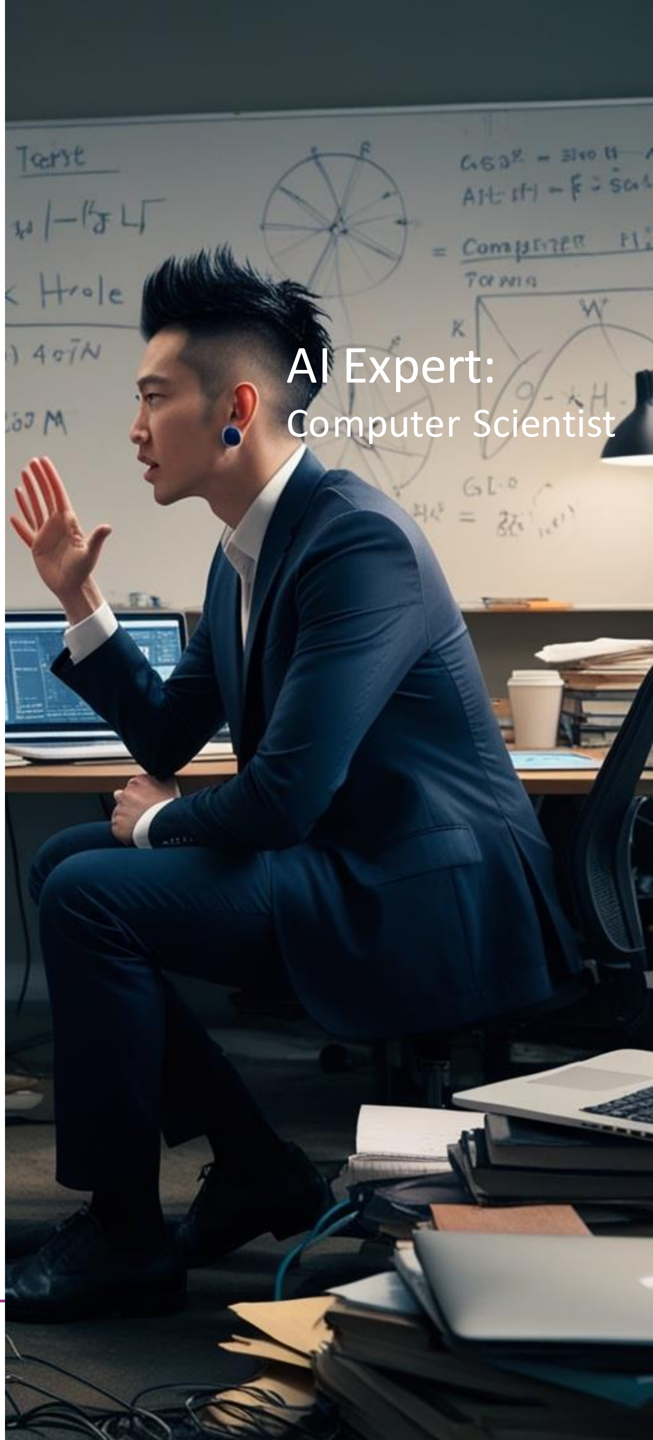

Efficient Communication between AI and domain experts

As domain expert: We have to communicate *some* information about the real-world task its core metrics and the data

As AI expert: We have to communicate *some* information about the AI pipeline, data processing and models used

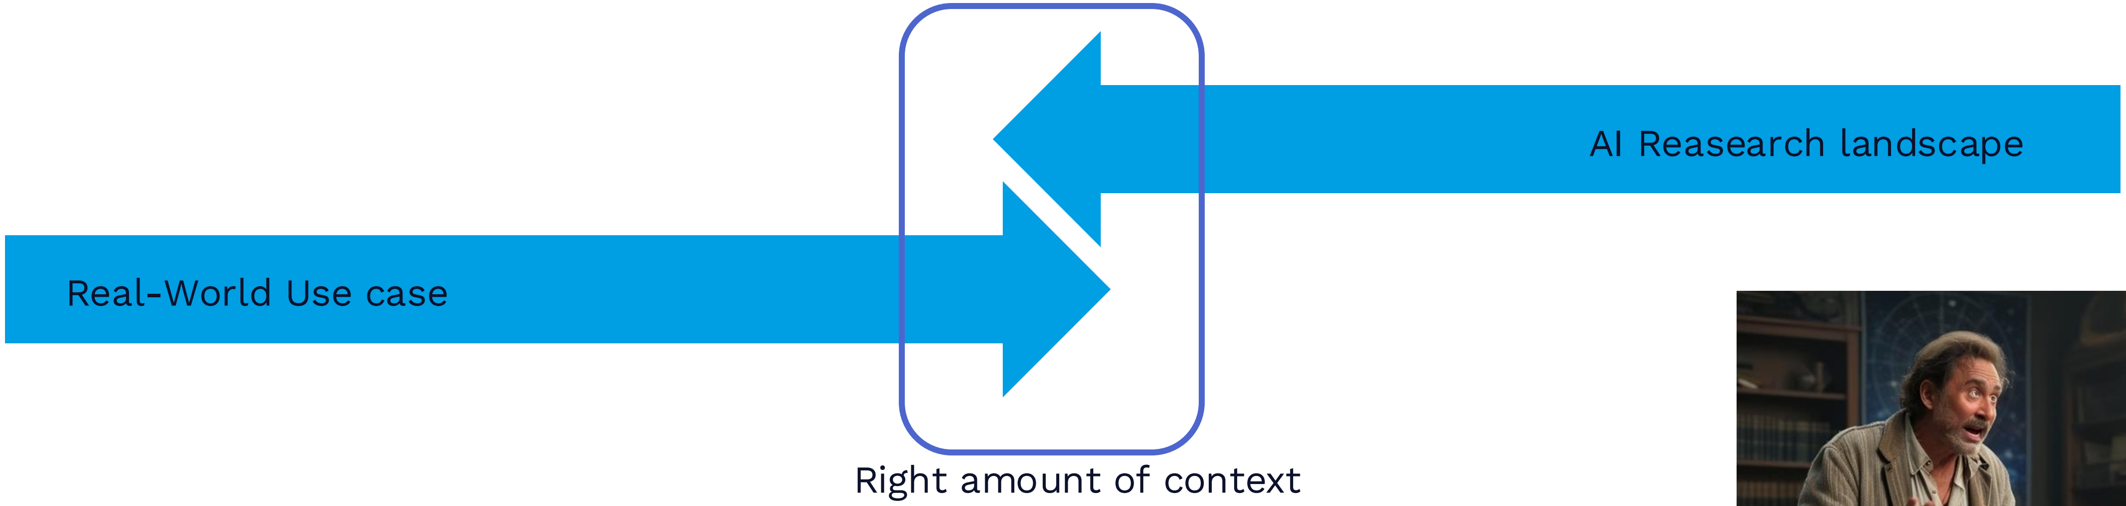

An (incomplete) Overview of Learning Tasks

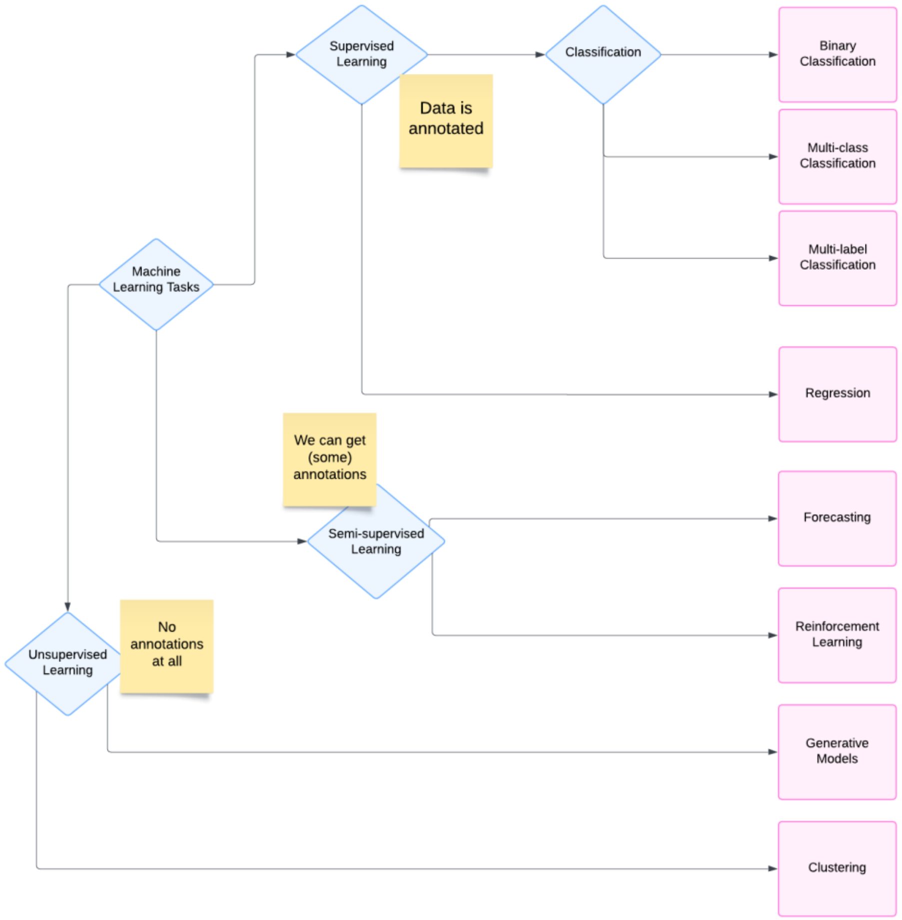

An (incomplete) Overview of Learning Tasks

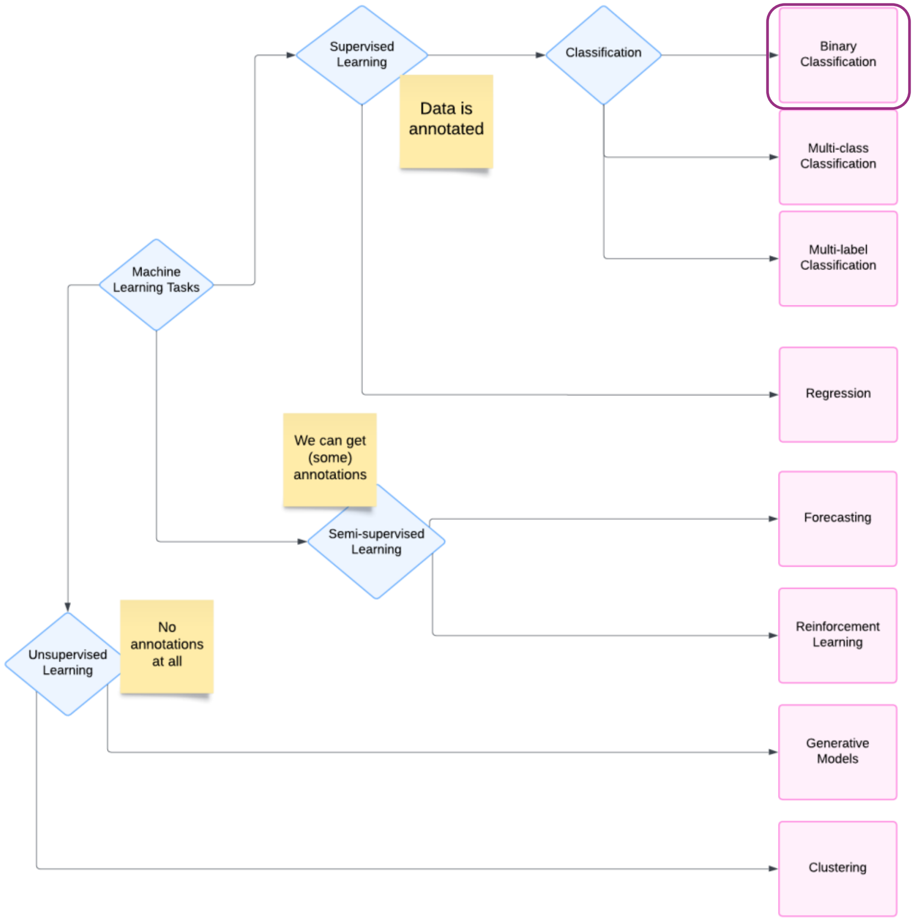

Binary Classification

A task where the model predicts one of two possible classes (labels) based on input data.

Examples

A spam filter that classifies emails as either "spam" or "not spam".

A model that classifies whether a detected particle in an airshower telescope is either a gamma or a hardronic event.

An (incomplete) Overview of Learning Tasks

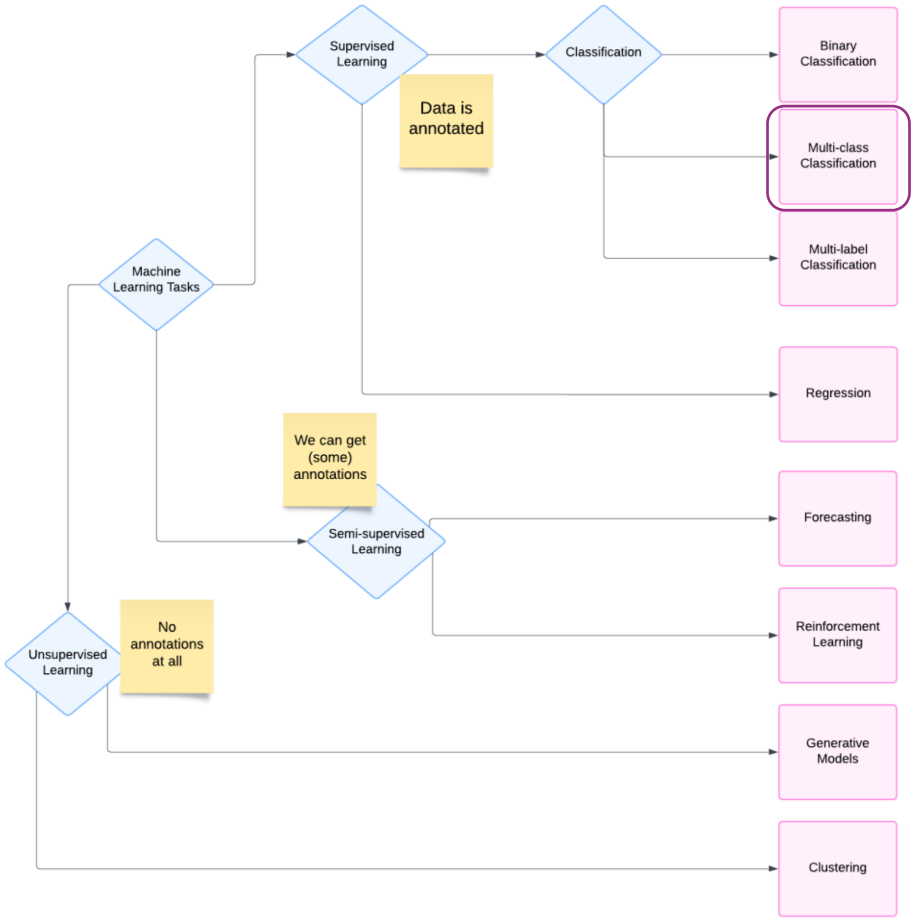

Multi-class Classification

A task where the model predicts one class out of several possible classes.

Examples

An image recognition system that classifies a picture as either "cat," "dog," or "bird".

A model used in a gamma-ray observatory to classify incoming particles as either gamma-rays, hardrons (e.g. protons and heavy nucliei), electrons, or muons, based on their energy and interaction signatures.

An (incomplete) Overview of Learning Tasks

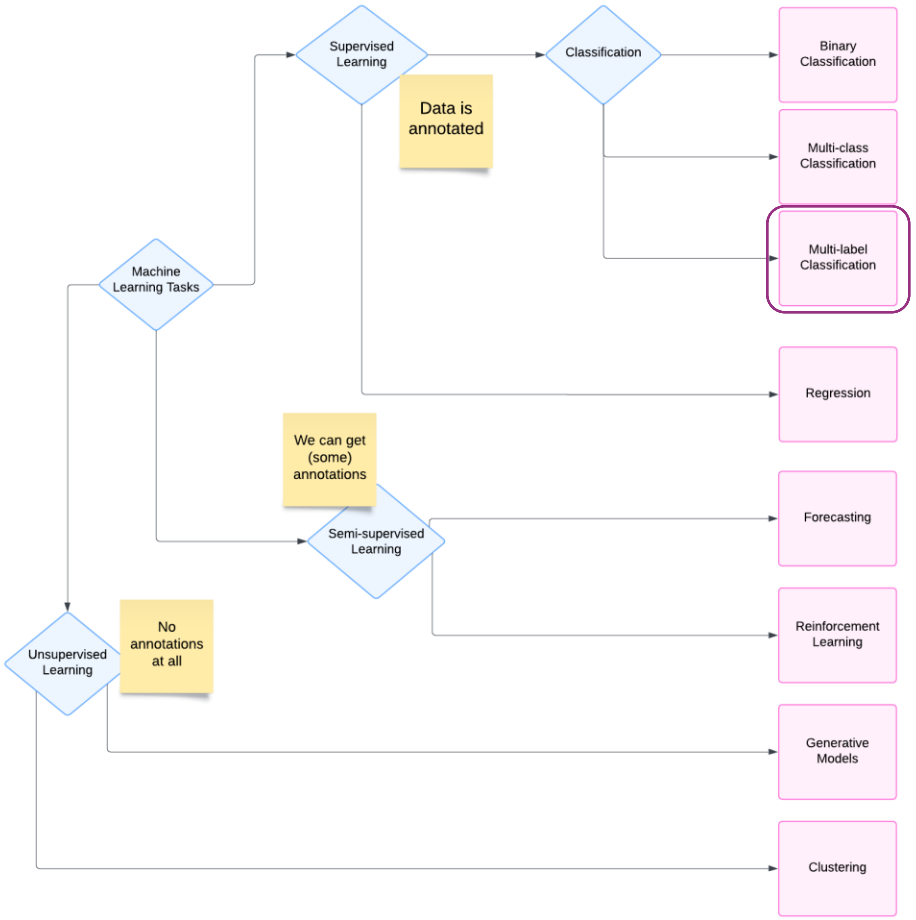

Multi-label Classification

A task where the model predicts multiple labels (classes) for each input, where each label is independent of others.

Examples

A movie recommendation system where each movie can have multiple genres (e.g., "Action," "Comedy," and "Adventure").

A model in an air shower experiment (detecting cosmic rays) that labels each particle as both a proton and a secondary electron or muon, depending on the multiple components involved in the event.

An (incomplete) Overview of Learning Tasks

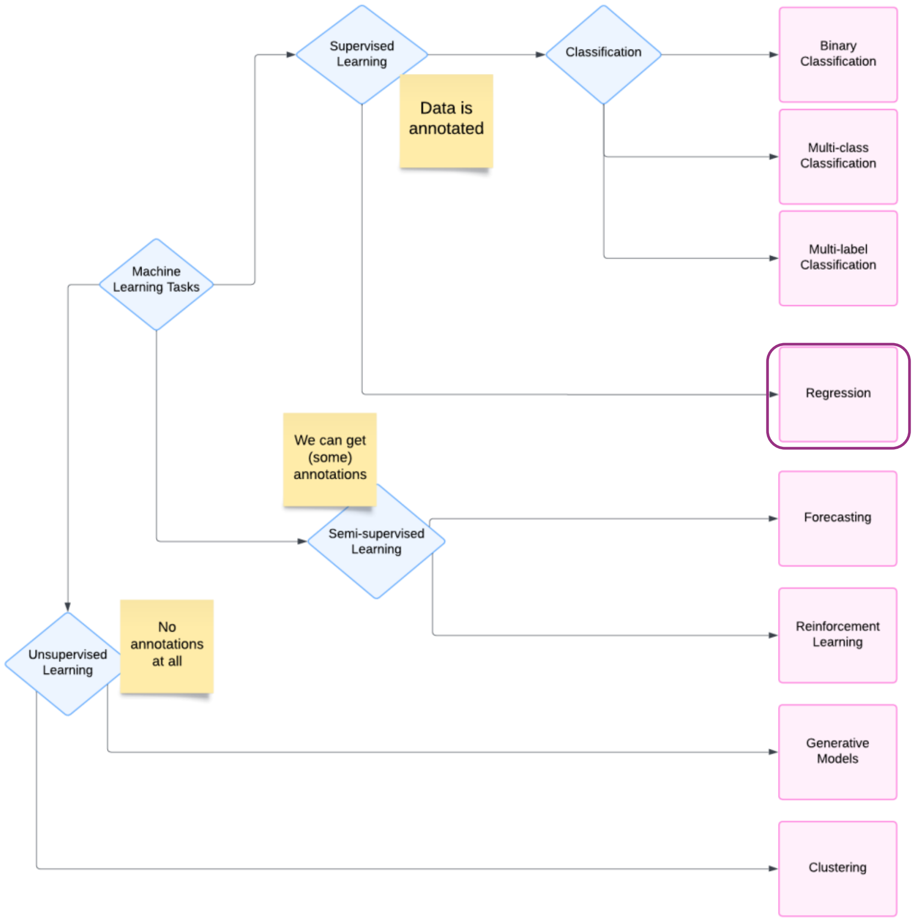

Regression

A task where the model predicts a continuous value (usually real numbers) based on input data.

Examples

A house price prediction model that takes in features like square footage, number of bedrooms, and location, and predicts the price.

A model that estimates the energy of a gamma ray based on its shower‘s light-yield and spacial properties.

An (incomplete) Overview of Learning Tasks

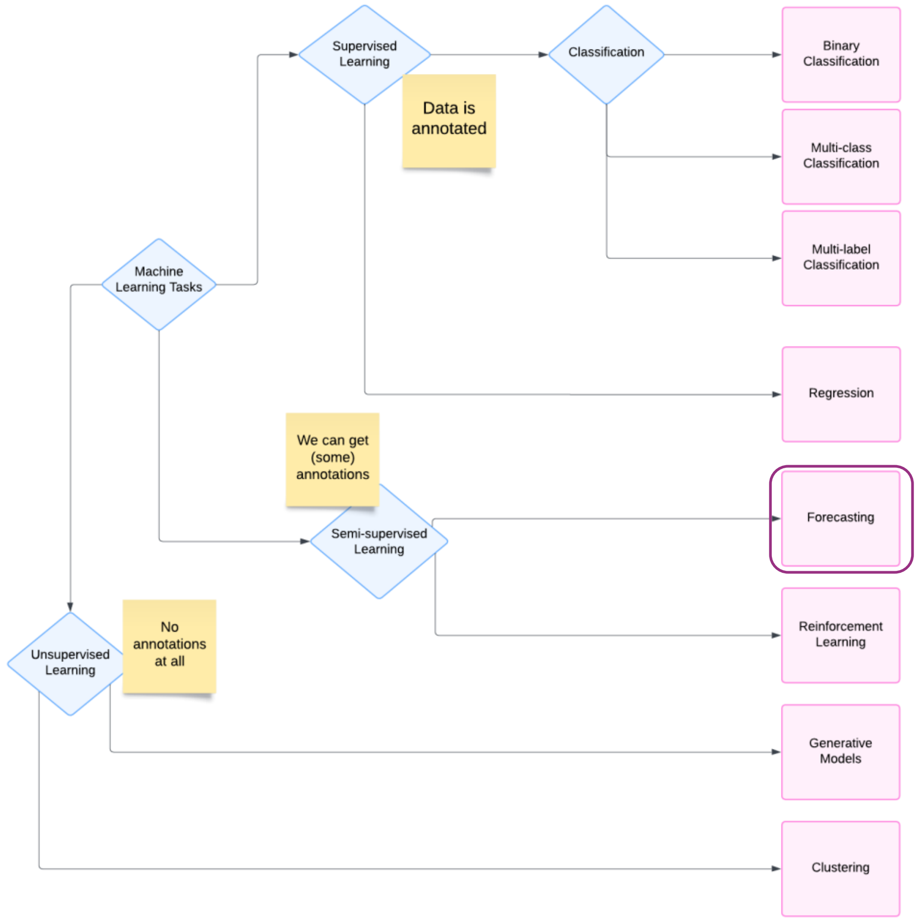

Forecasting

A specialized form of regression where the model predicts future values based on historical data.

Examples

A stock price prediction model that uses past stock prices to forecast future stock prices.

A model that forecasts the solar activity cycle (e.g., sunspots or solar flares) based on historical data, helping predict solar storms.

An (incomplete) Overview of Learning Tasks

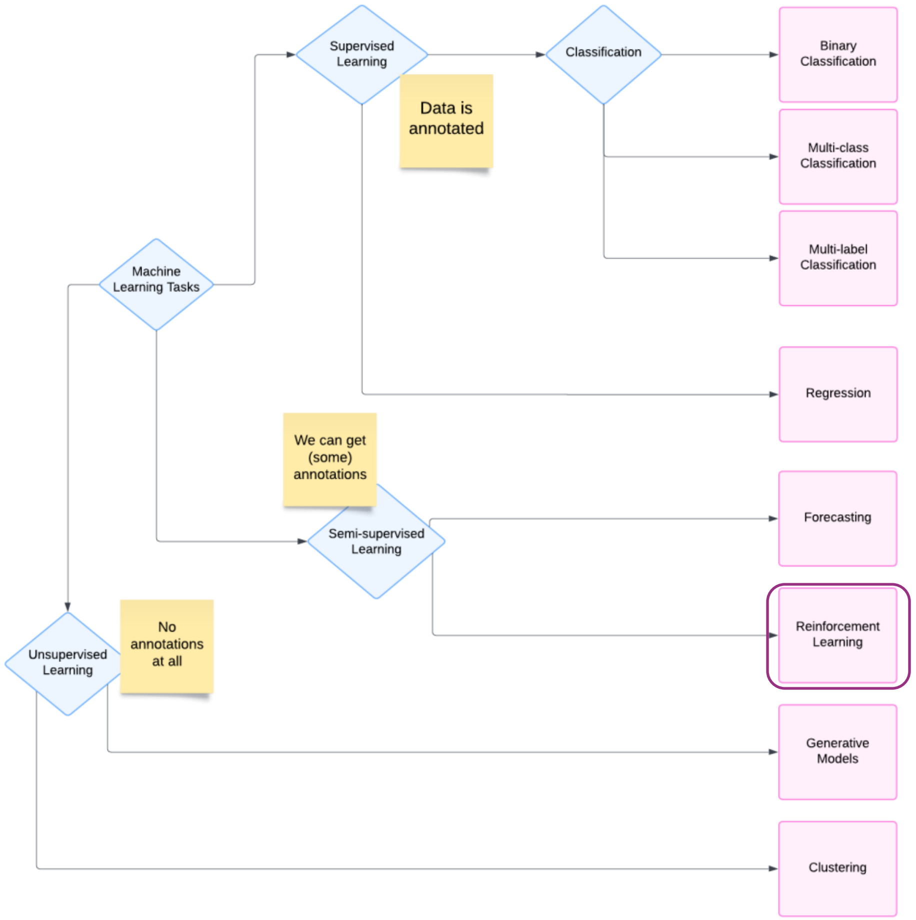

Reinforcement Learning

A task where an agent learns to make decisions by interacting with an environment and receiving rewards or penalties for its actions.

Examples

A self-driving car learning to navigate by making decisions and receiving rewards for avoiding obstacles and reaching its destination safely.

A robotic rover on Mars that uses reinforcement learning to navigate the planet's surface and explore geological features, optimizing its path for maximum exploration.

An (incomplete) Overview of Learning Tasks

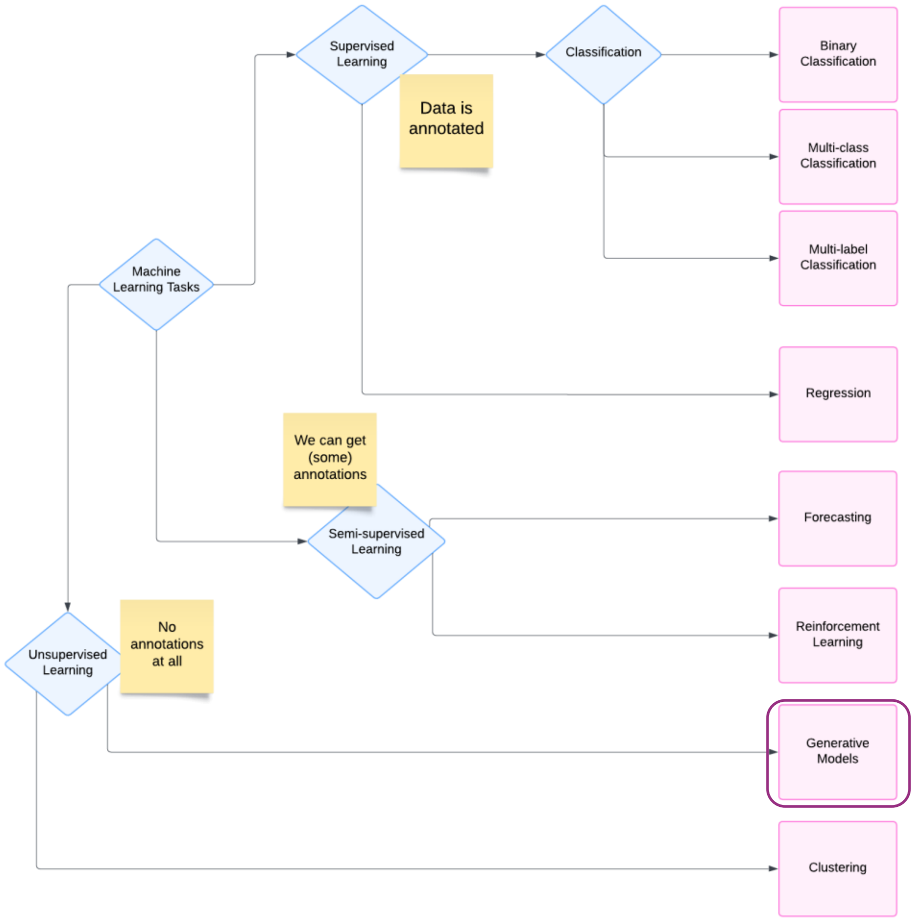

Generative Models

Models that learn to generate new data similar to the input data, such as images, text, or music.

Examples

A model that generates realistic images of people who don't exist, similar to the images found in GAN (Generative Adversarial Networks).

A Generative Adversarial Network (GAN) generating synthetic data of neutrino interactions in a detector, based on patterns in real event data, to simulate rare particle interactions that are hard to observe directly.

An (incomplete) Overview of Learning Tasks

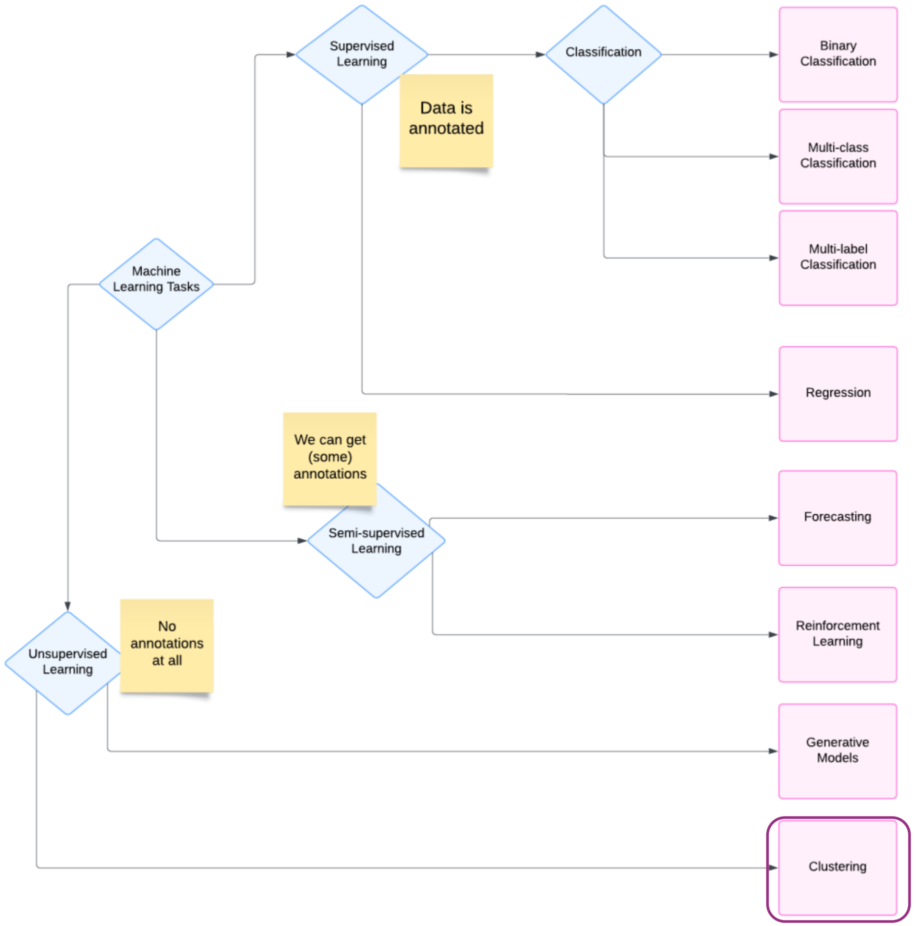

Clustering

A task where the model groups data points into clusters based on similarities, without any predefined labels.

Examples

A customer segmentation model that groups customers into clusters based on purchasing behavior (e.g., "high spenders," "frequent buyers," etc.).

A model that clusters galaxies into different groups based on their morphology (e.g., elliptical, spiral, irregular)

LAMARR
INSTITUTE FOR MACHINE LEARNING AND ARTIFICIAL INTELLIGENCE

Efficient Communication between AI and Domain Experts

How to do it!

Jens Buß & Ramsés Sanchéz & Sebastian Buschjäger

Partner institutions:
technische universität dortmund
Fraunhofer IAIS
Universität Bonn
Fraunhofer IML

Institutionally funded by:
Bundesministerium für Bildung und Forschung
Ministerium für Kultur und Wissenschaft des Landes Nordrhein-Westfalen

LAMARR
INSTITUTE FOR MACHINE LEARNING AND ARTIFICIAL INTELLIGENCE

Efficient Communication between AI and domain experts

OR RATHER how NOT to do it!

Jens Buß & Ramsés Sanchéz & Sebastian Buschjäger

Partner institutions:
technische universität dortmund
Fraunhofer IAIS
UNIVERSITÄT BONN
Fraunhofer IML

Institutionally funded by:
Bundesministerium für Bildung und Forschung
Ministerium für Kultur und Wissenschaft des Landes Nordrhein-Westfalen

Vision: AI-driven Gamma-ray Astronomy

Help us to improve our data analysis with AI. You‘re the experts!

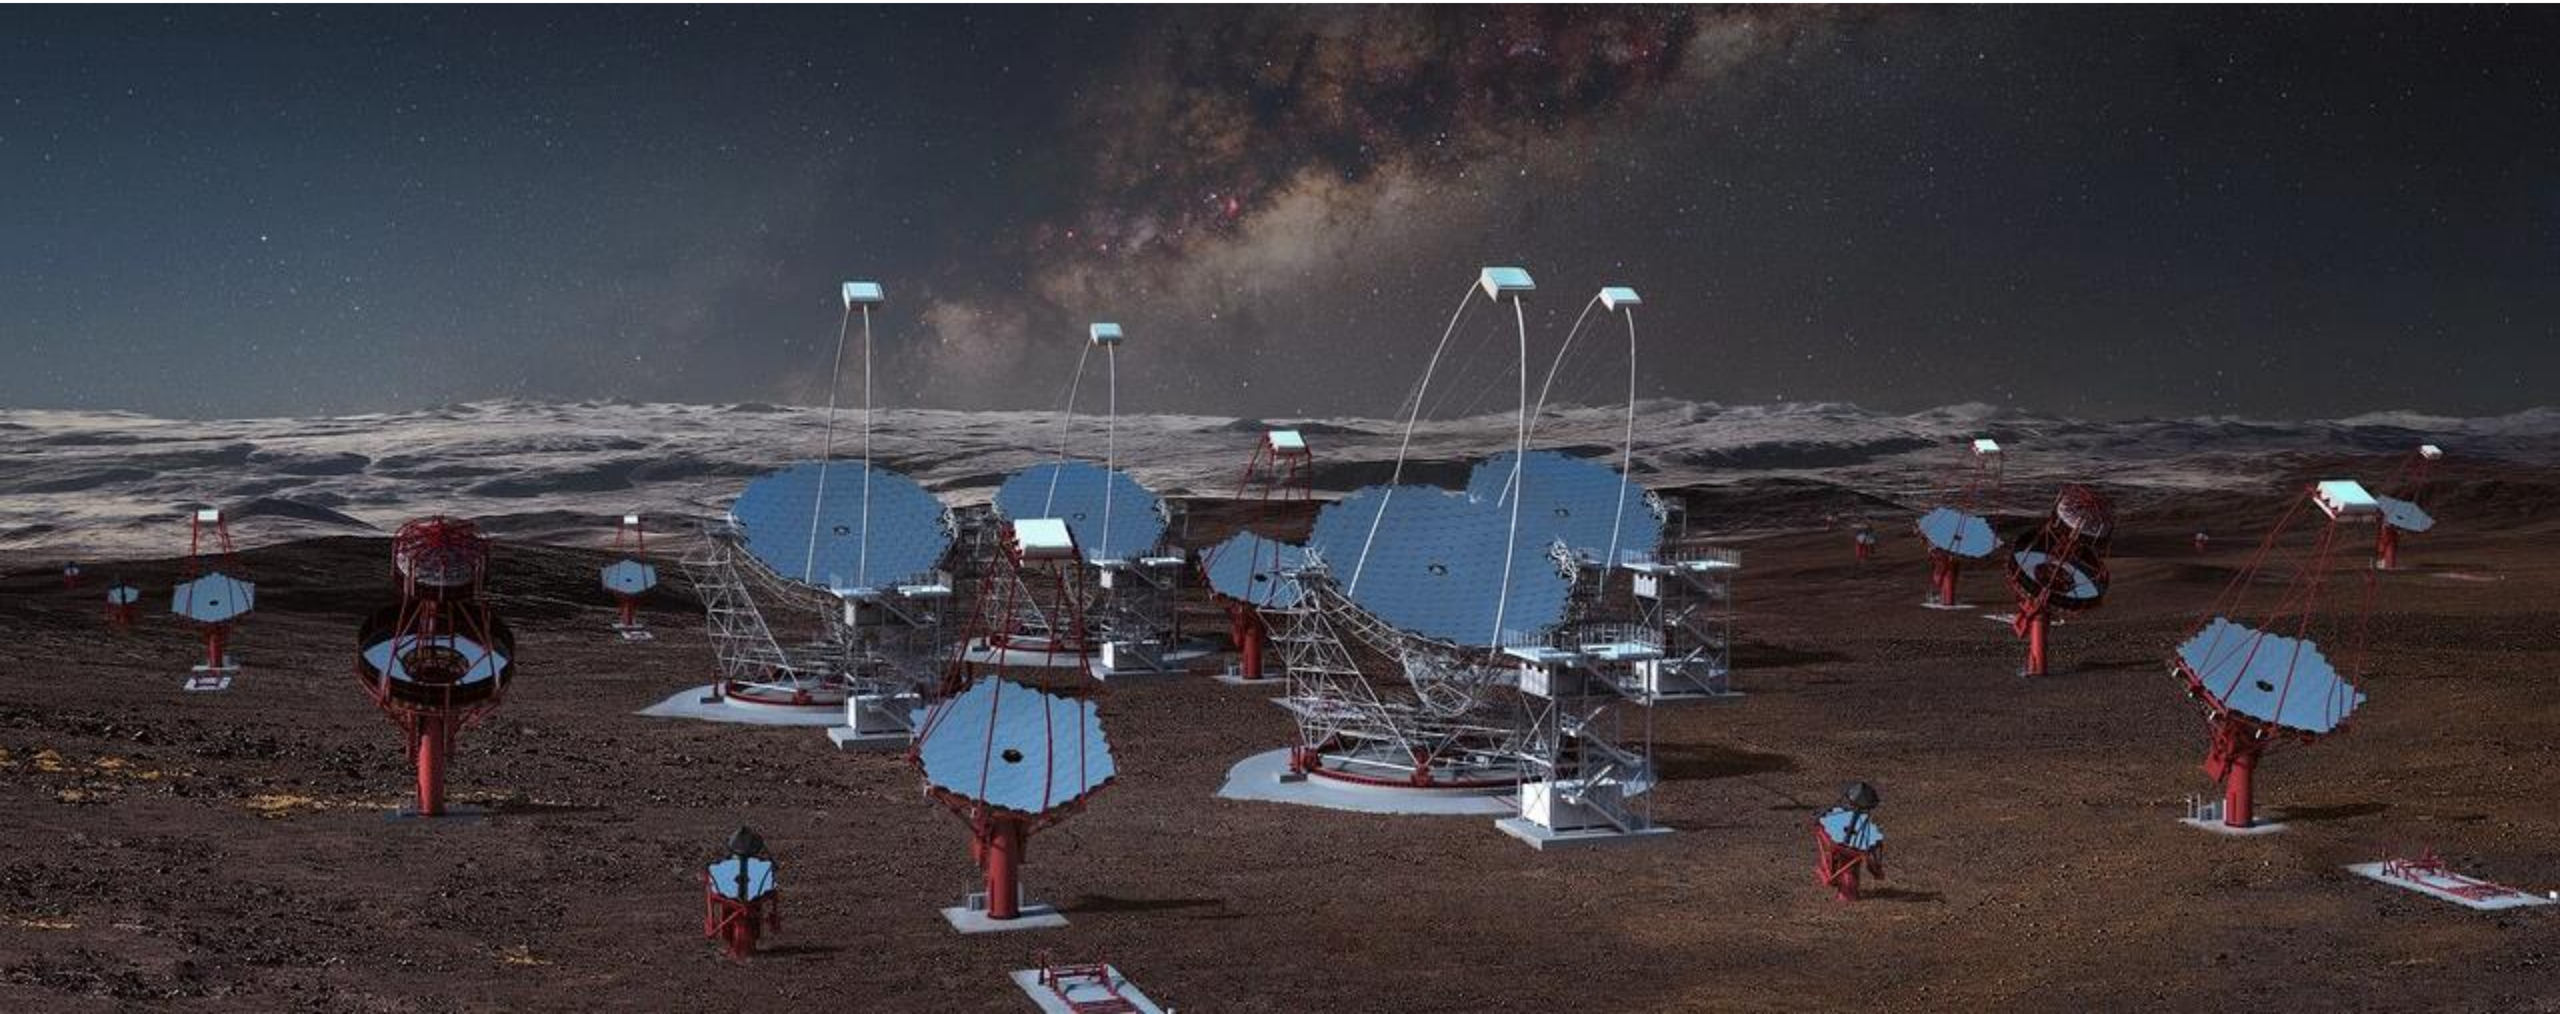

Vision: AI-driven Gamma-ray Astronomy

Help us to improve our data analysis with AI. You‘re the experts!

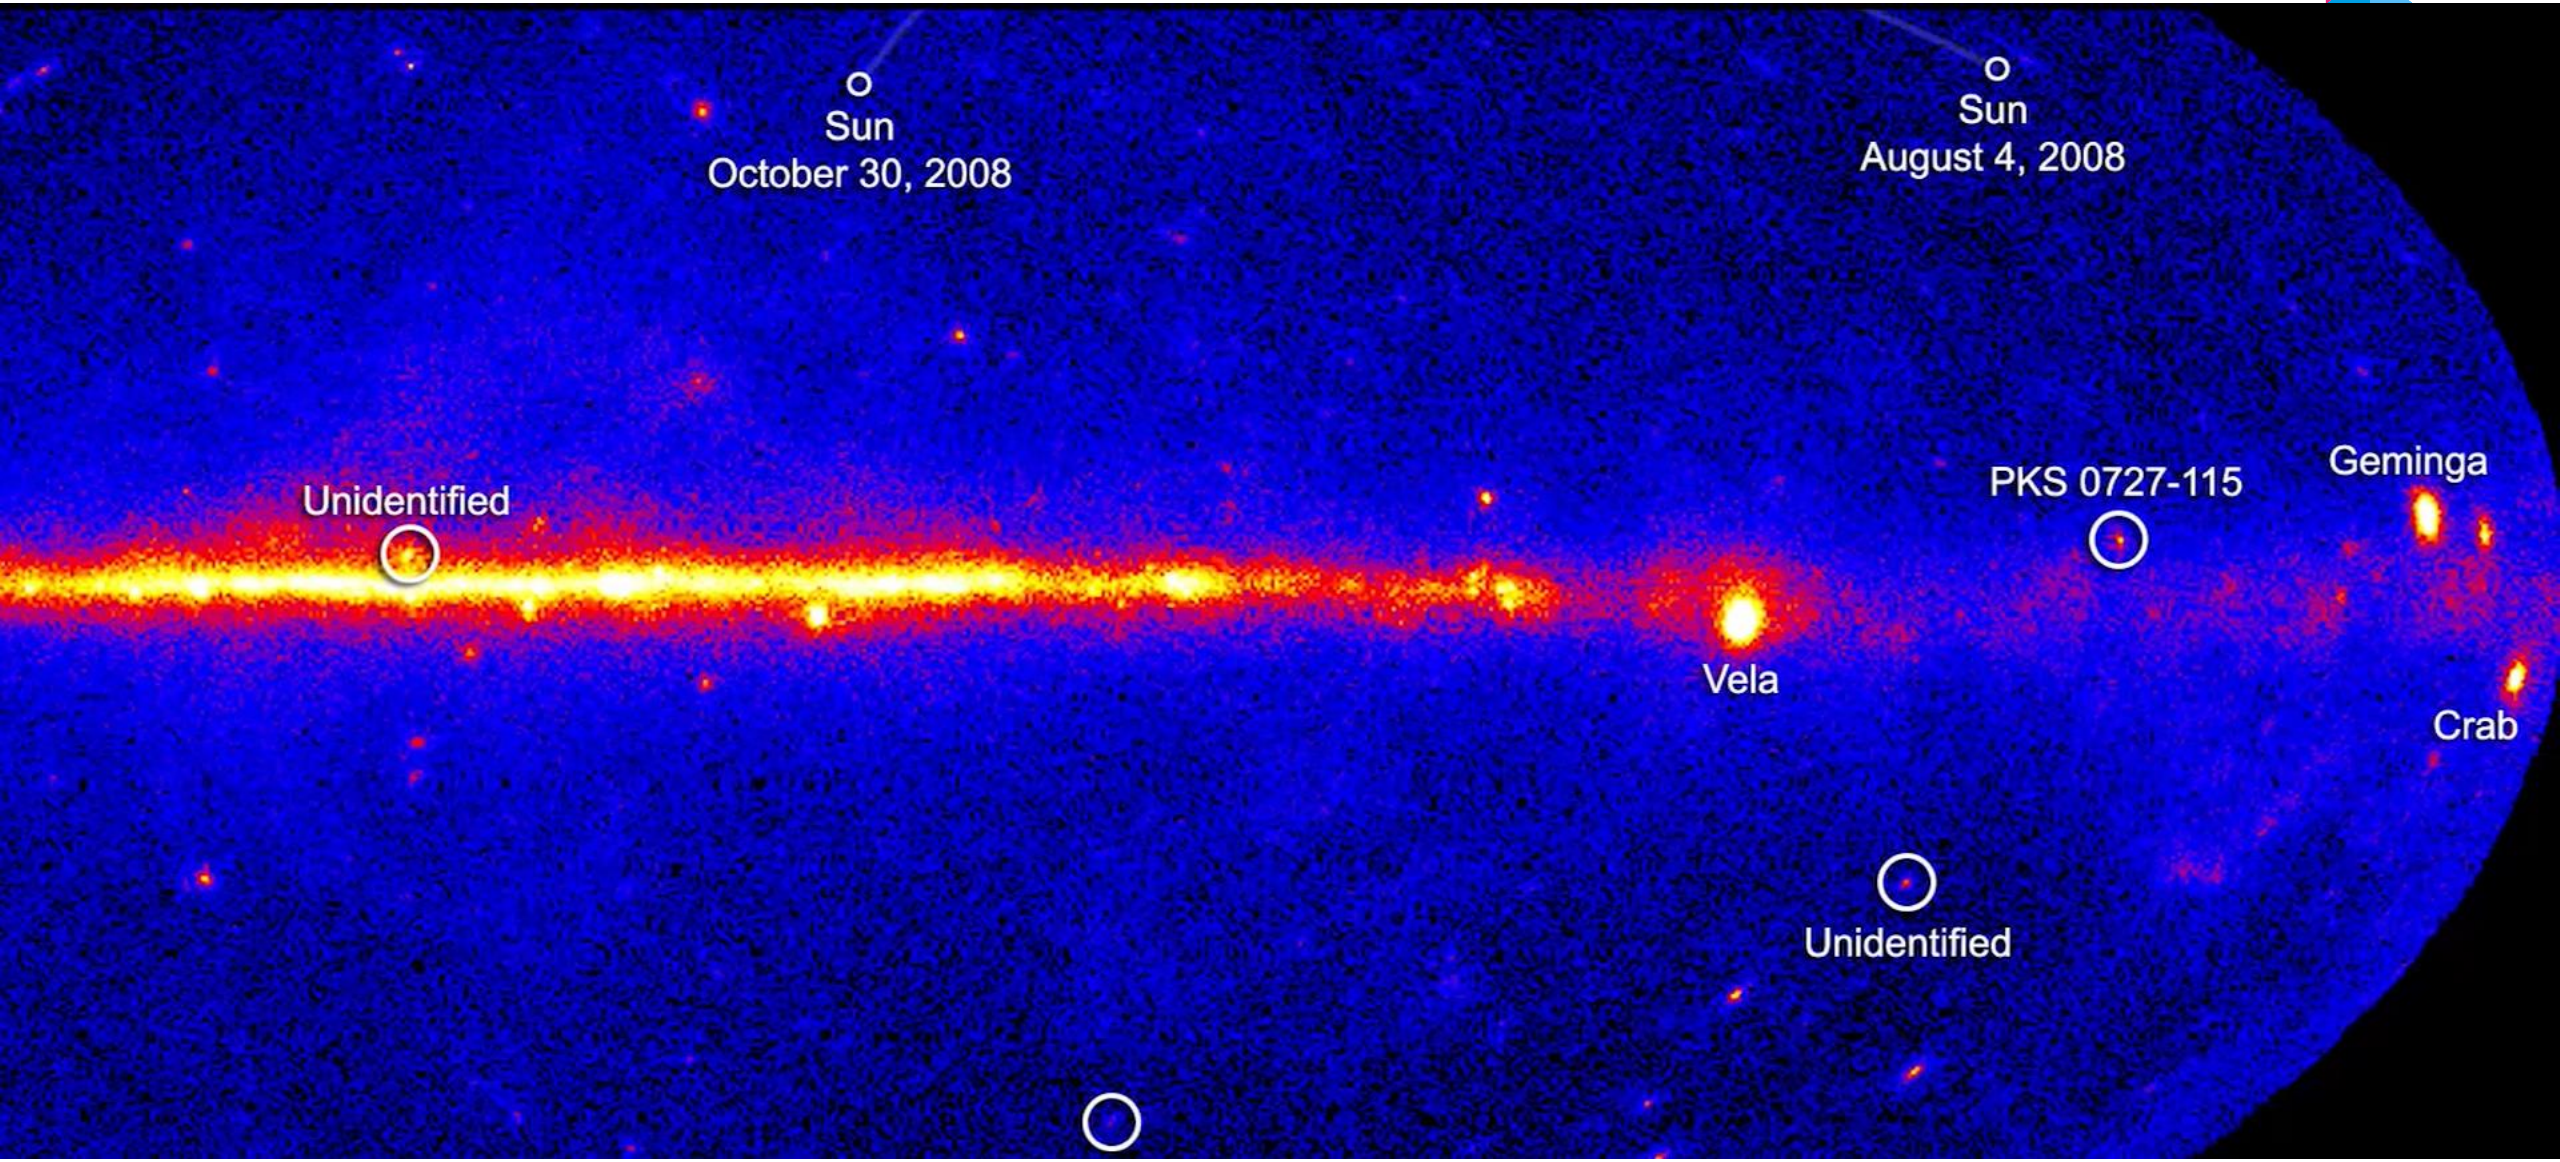
Sun
October 30, 2008
Sun
August 4, 2008
Unidentified
PKS 0727-115
Geminga
Vela
Crab
Unidentified

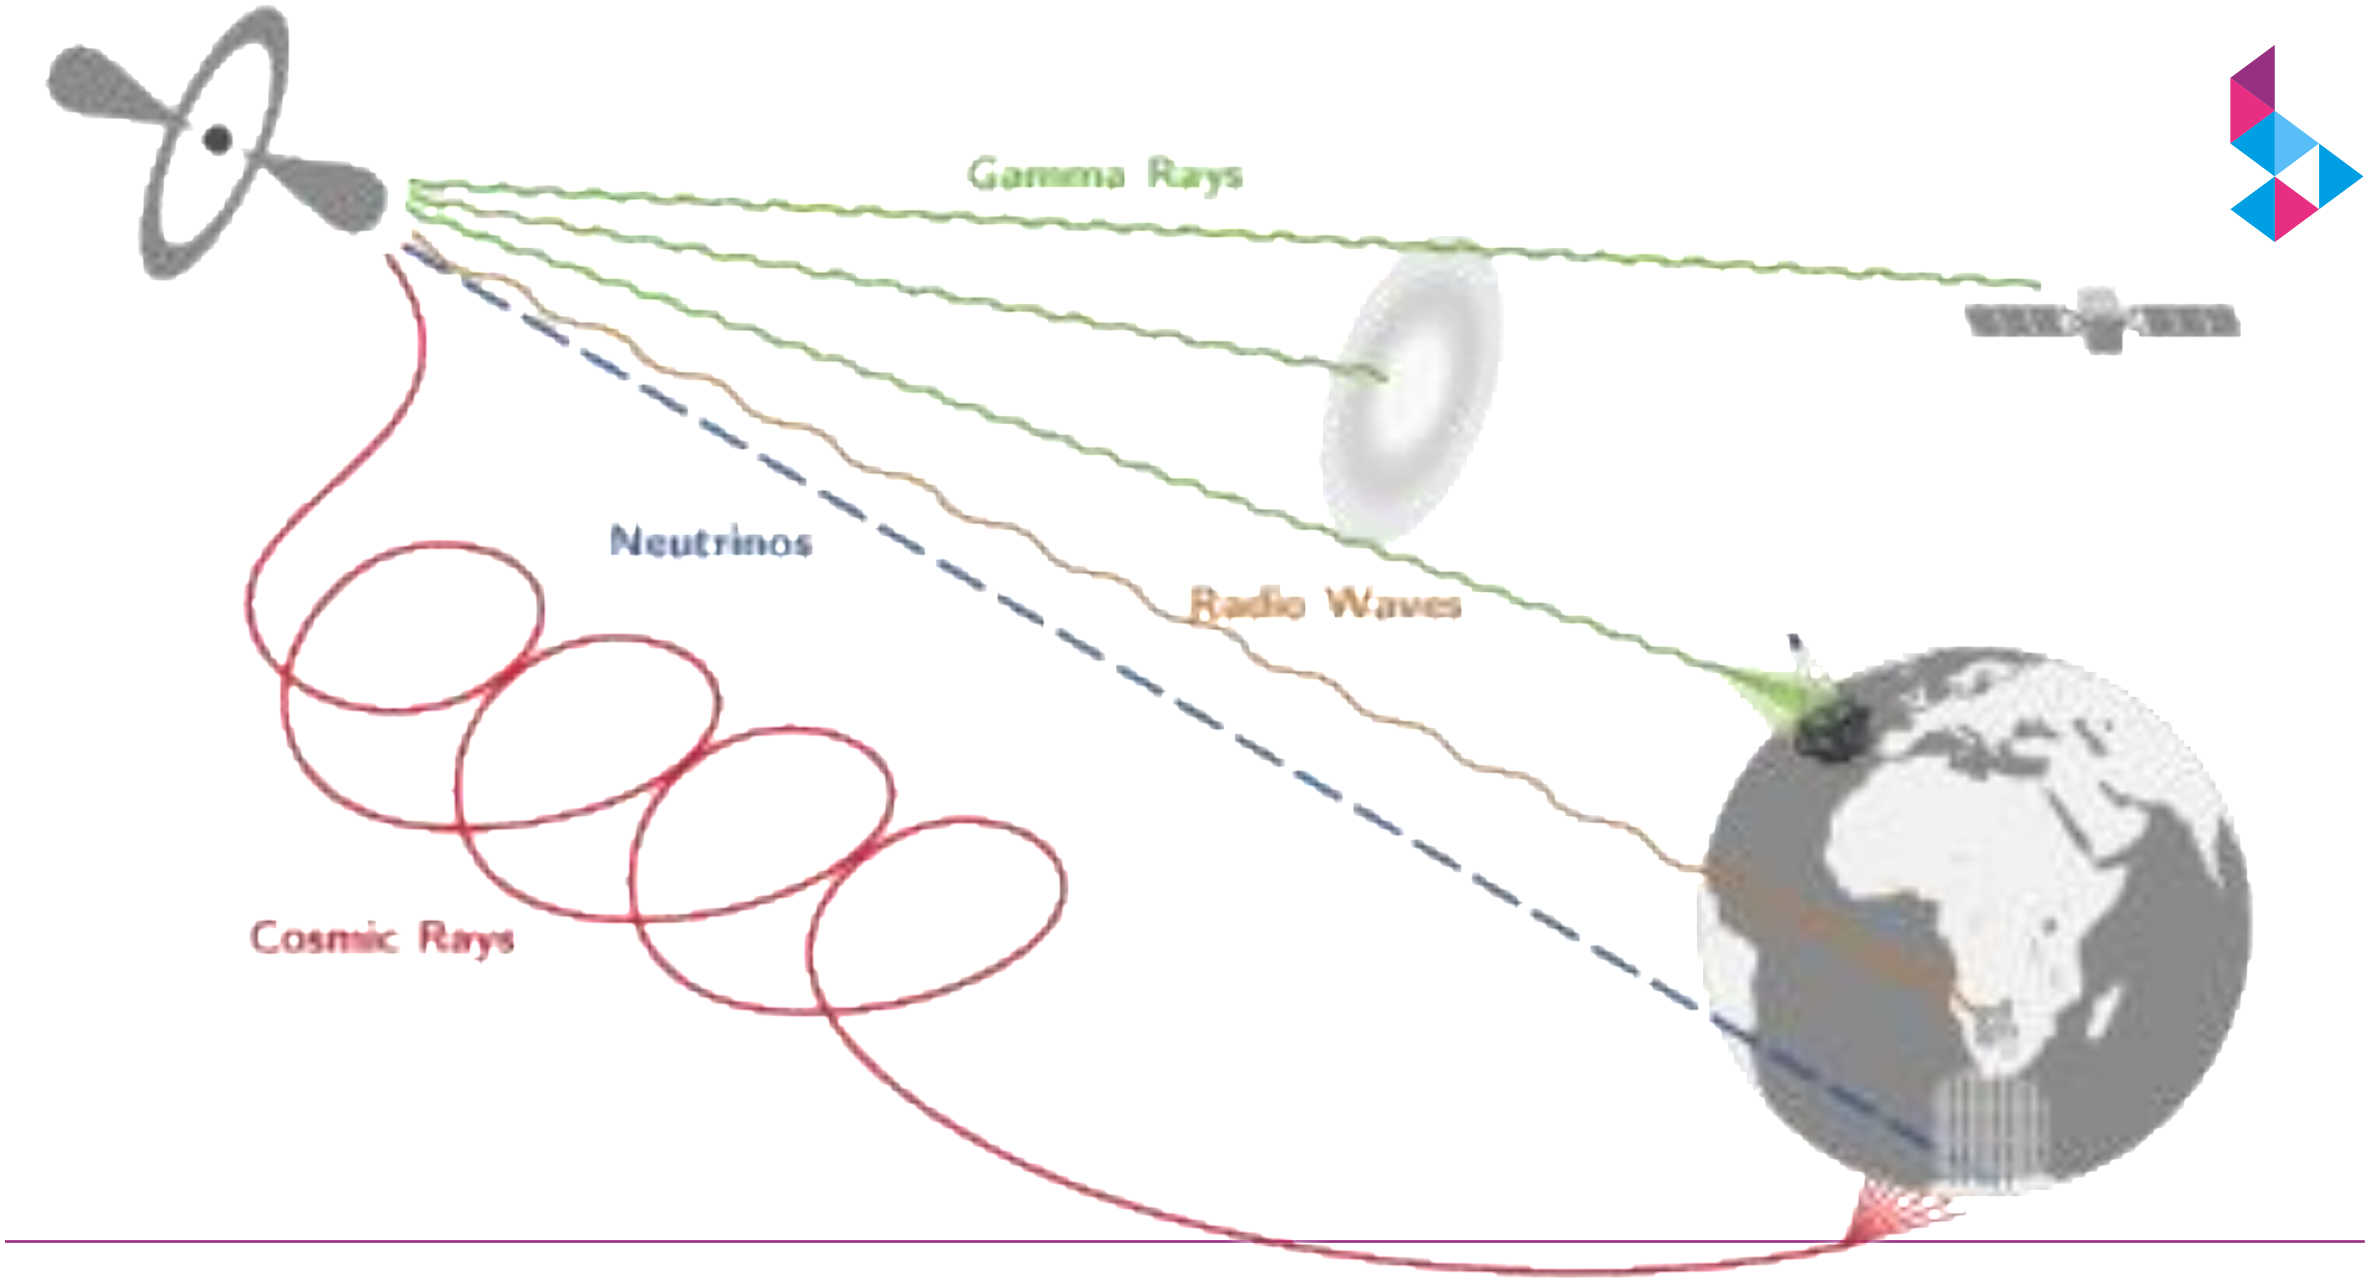
Gamma Rays
Neutrinos
Radio Waves
Cosmic Rays

Airshowers

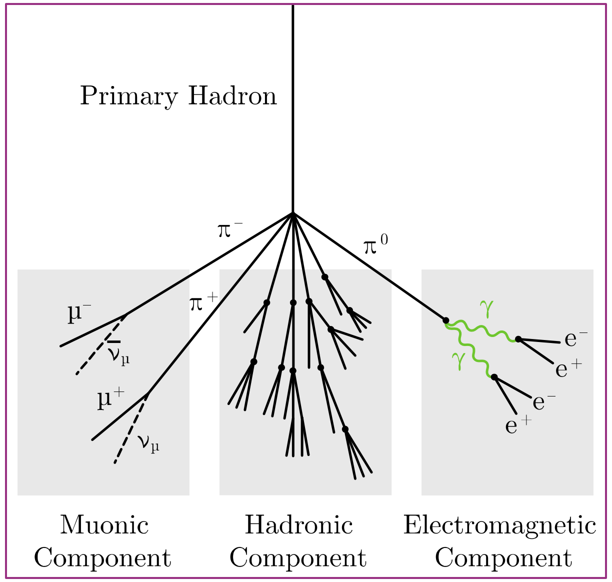

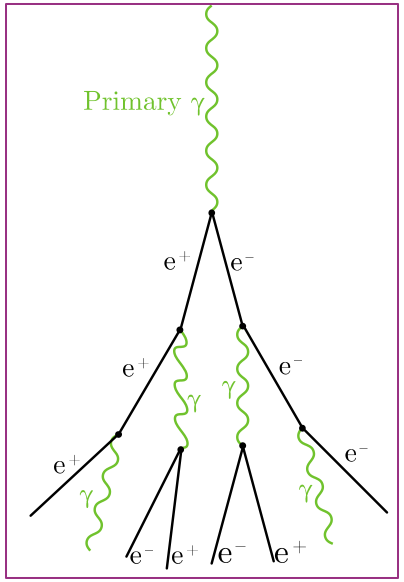

Detection principle

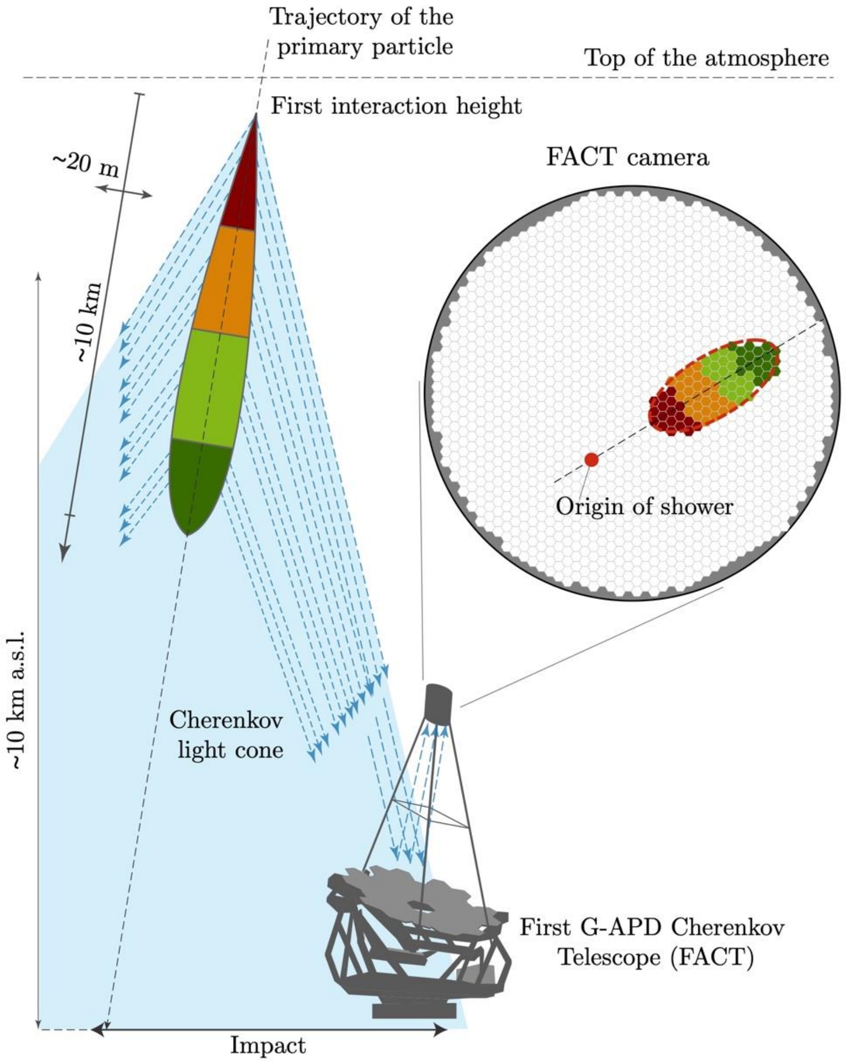

This is what we are looking for...

Spectral energy distributions (SEDs)

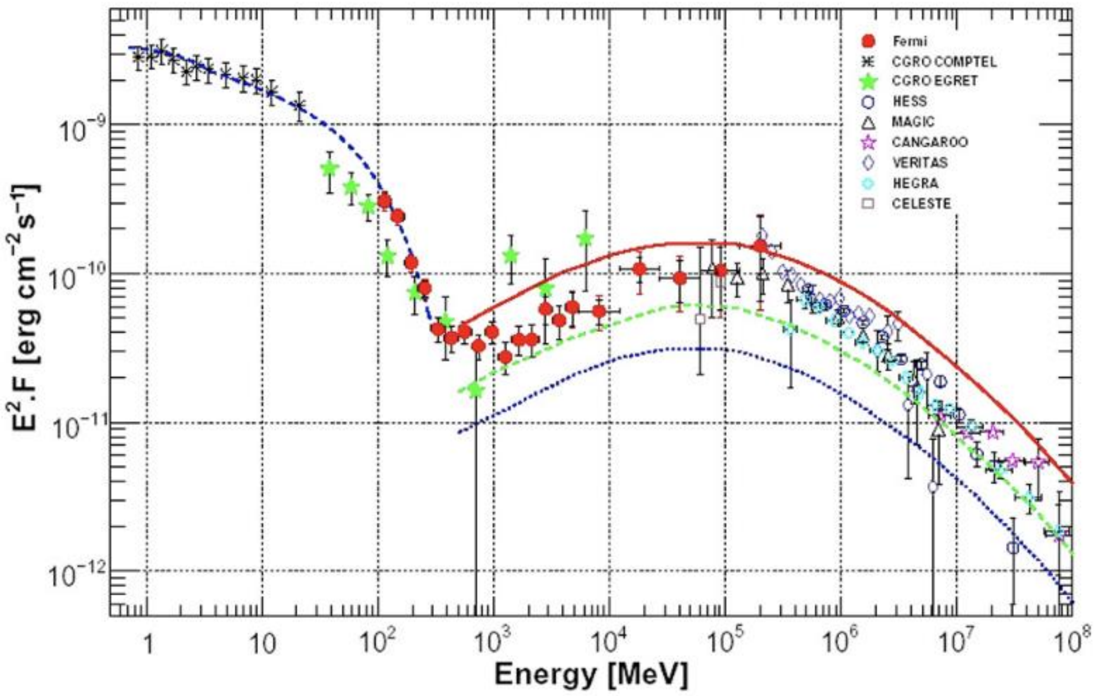

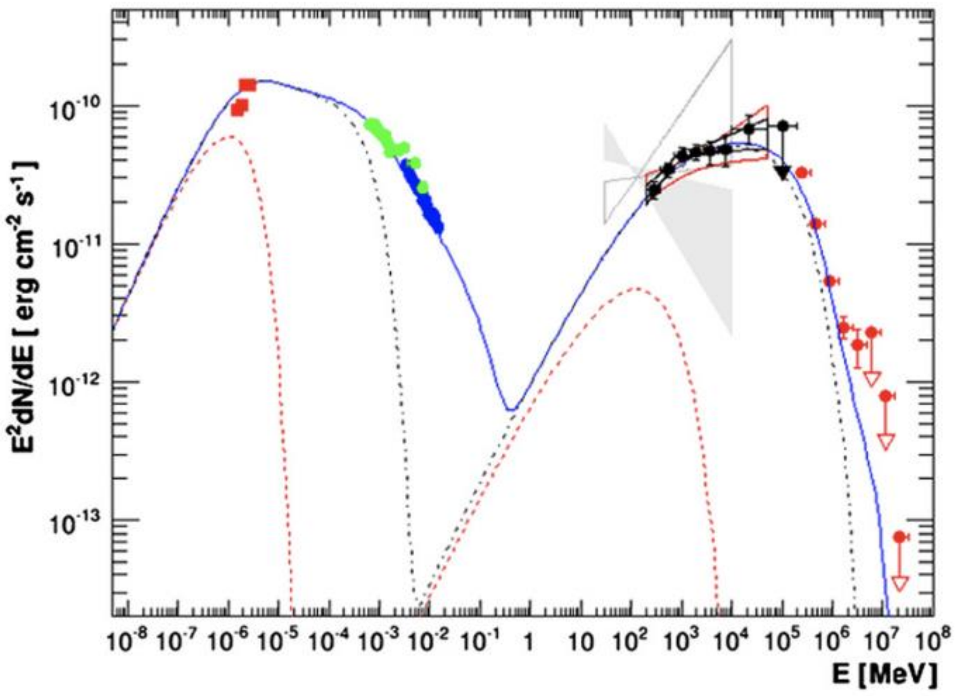

This is what we are looking for… Lightcurves

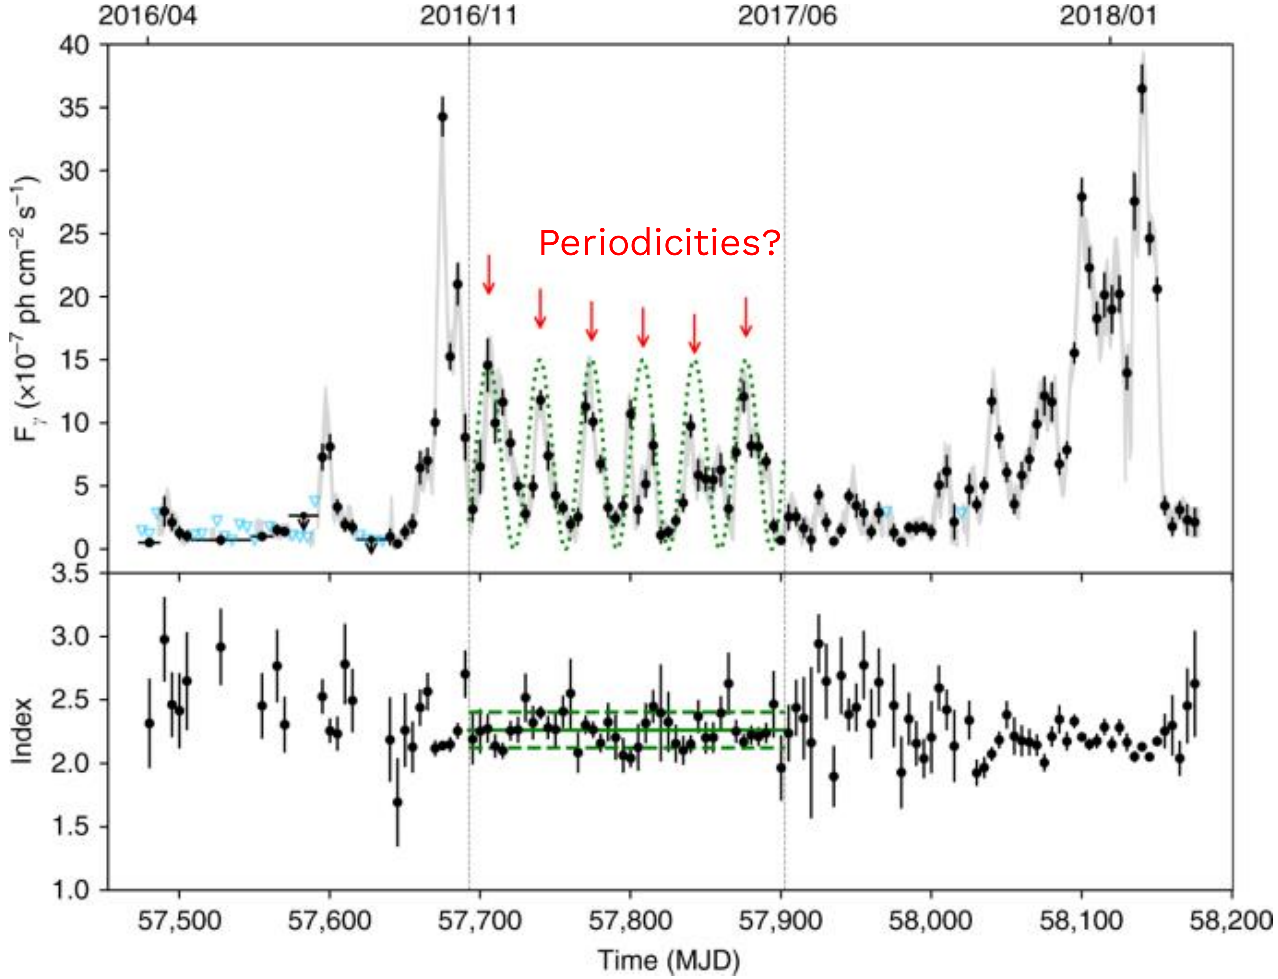

…But

High-Energy Gamma-Ray Lightcurves

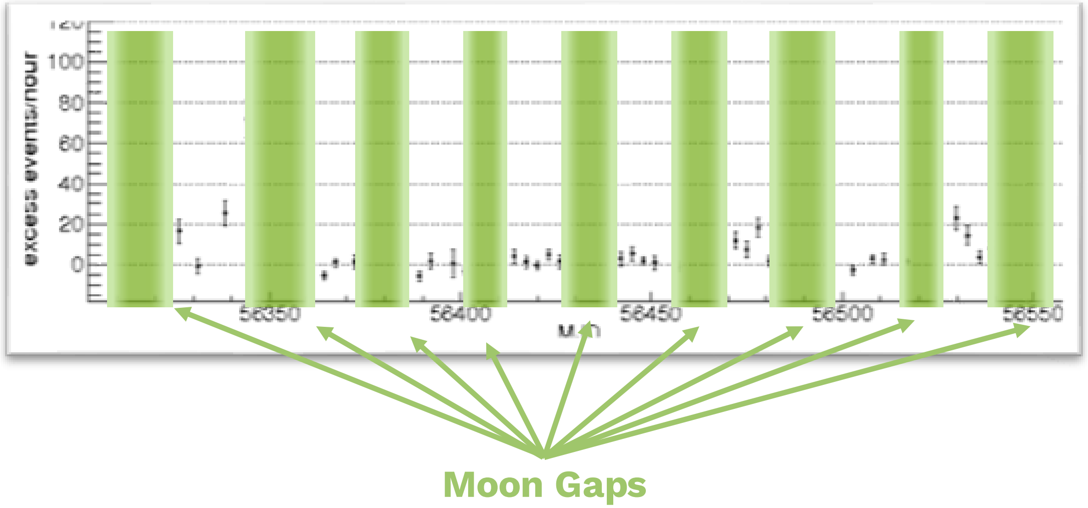

Bad Moon Rising?

PMT

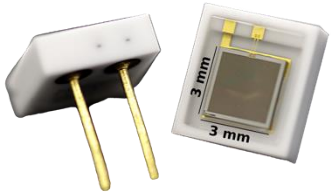

SiPM

- No HV (~70V)
- PDE comparable to PMTs
- Robust

The Experiment

30 Hexagon
Mirrors
9,5 m² Mirror Area

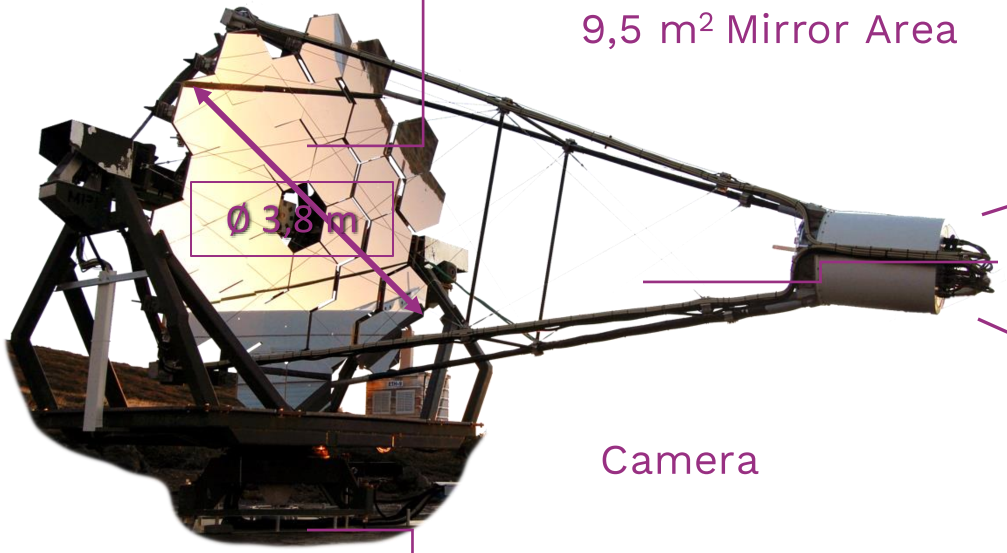

Camera

La Palma
~ 2200 m alt.

4,5°
FOV

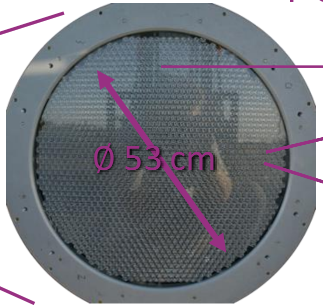

Hexagon
Pixel
1440 SiPMs

Analysis chain

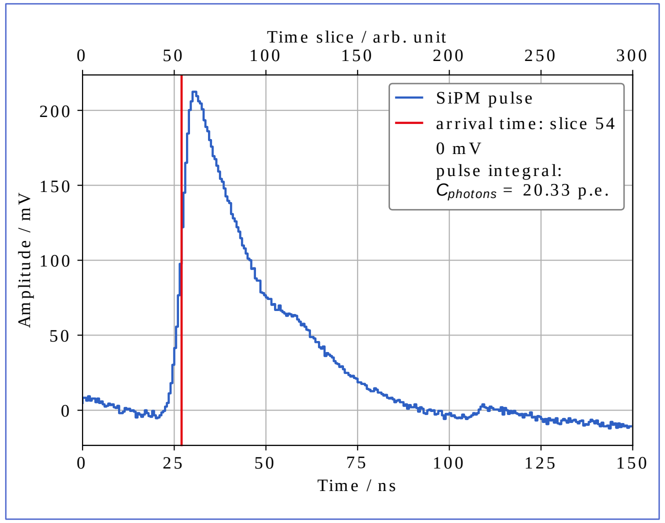

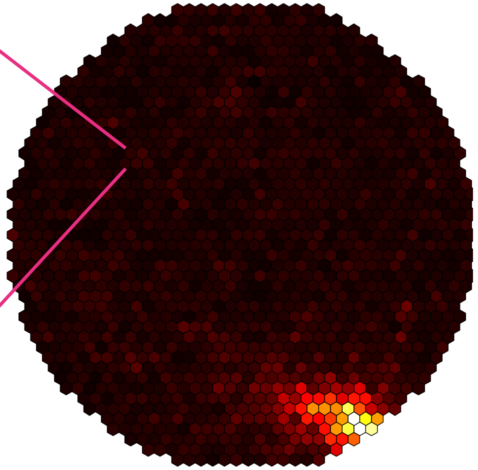

Analysis chain

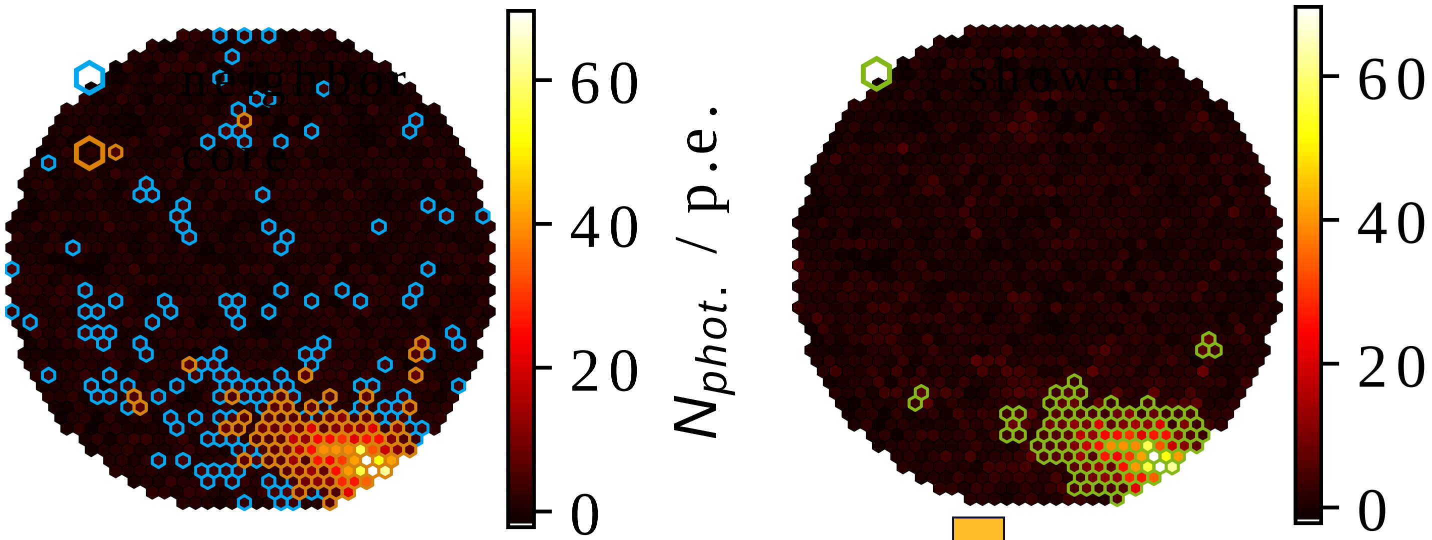

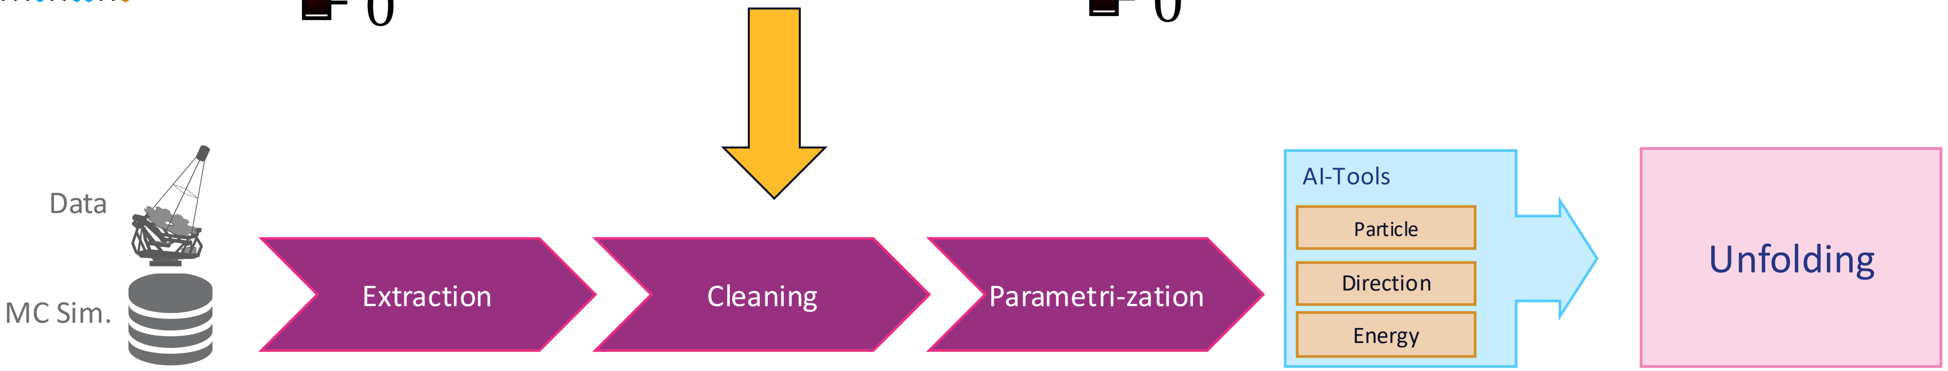

Analysis chain

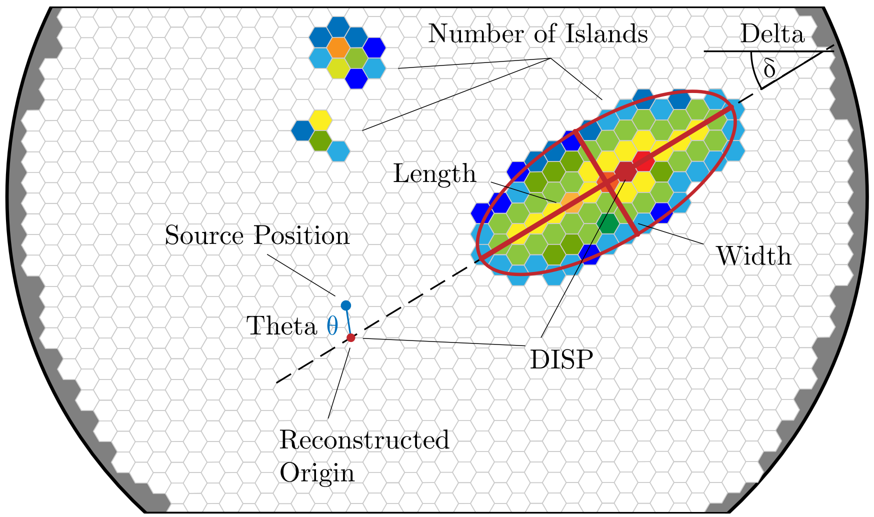

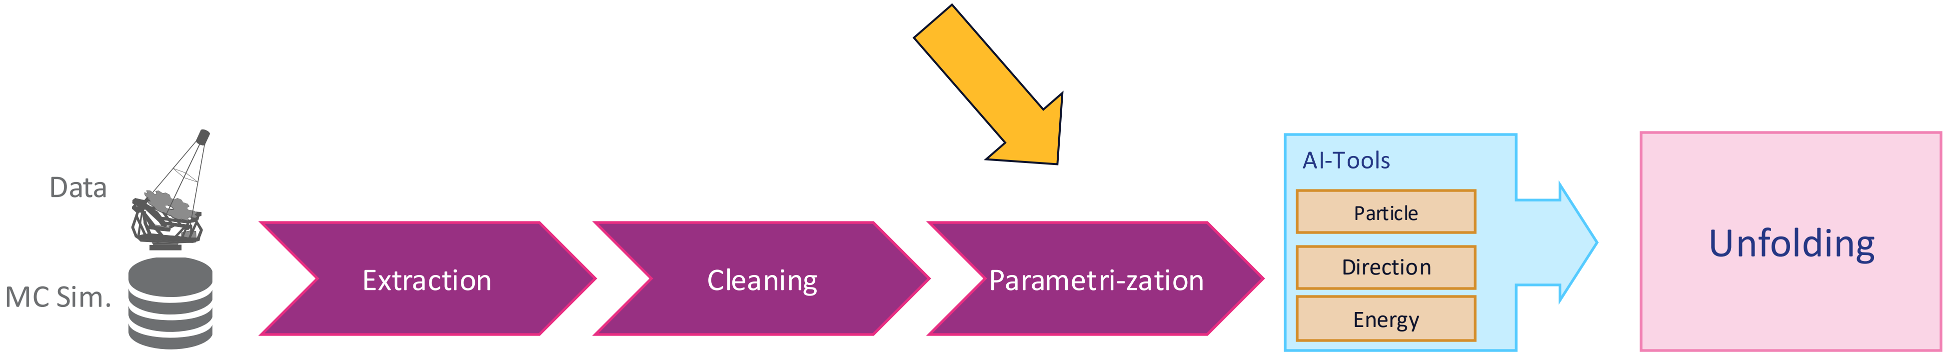

Analysis chain

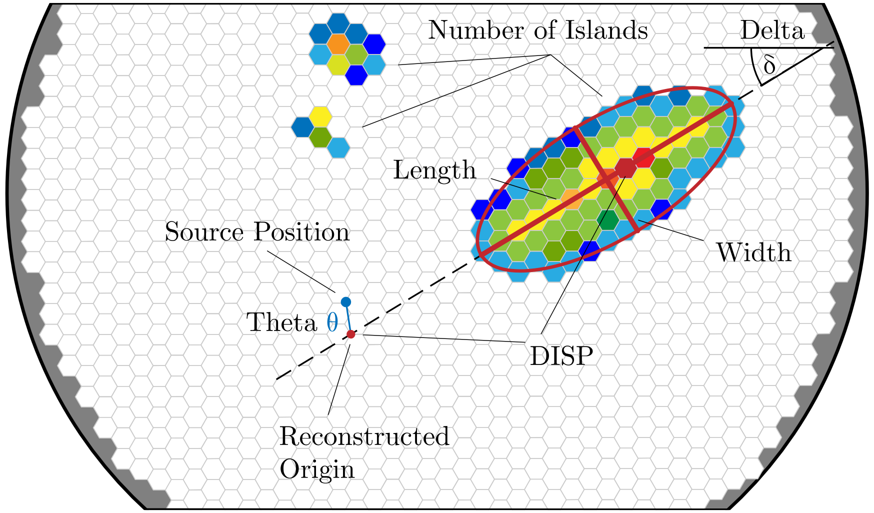

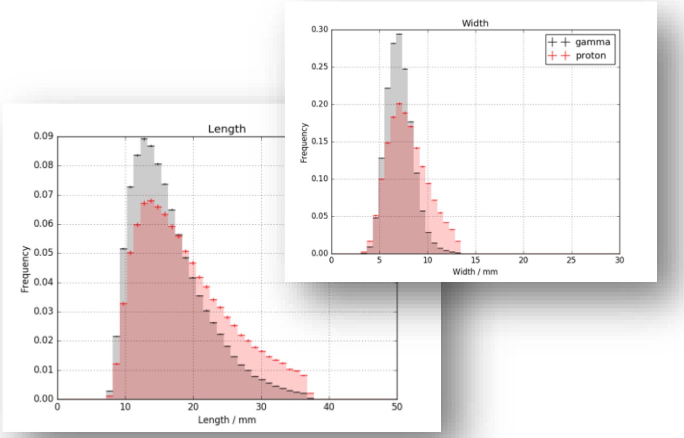

Data

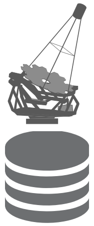

MC Sim.

Extraction → Cleaning → Parametri-zation

AI-Tools
- Particle
- Direction
- Energy

Unfolding

Analysis chain

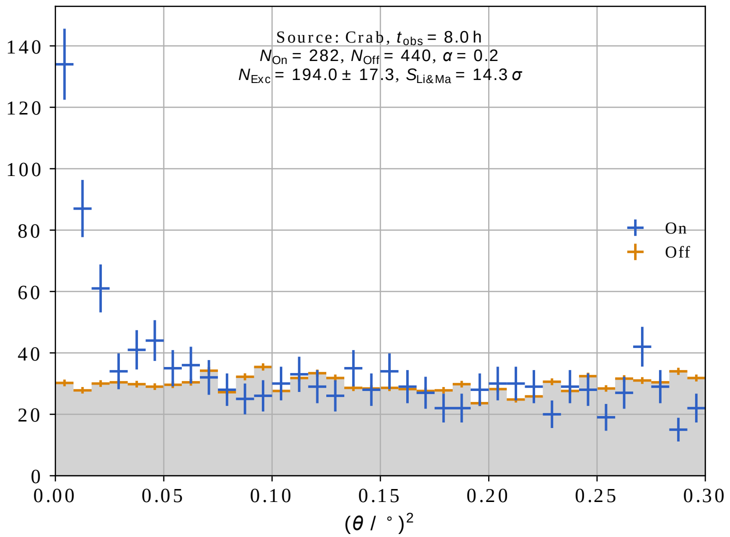

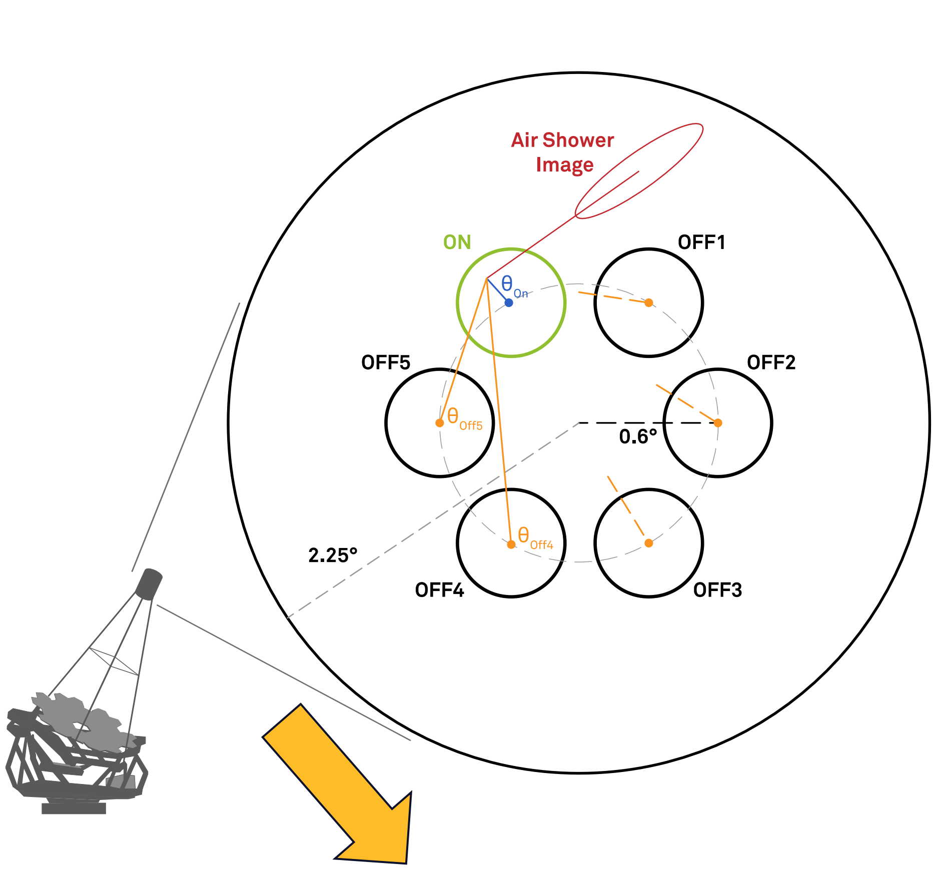

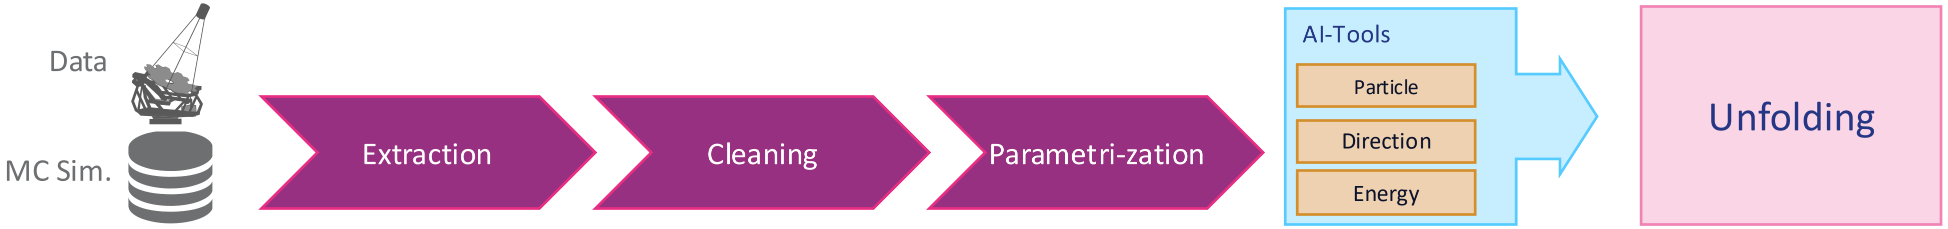

Analysis chain

Antwortfunktion

Untergrund

$$g\left(y\right) = \int A\left(x, y\right)\ f\left(x\right) dx + b\left(y\right)$$

Observablen

wahres Spektrum

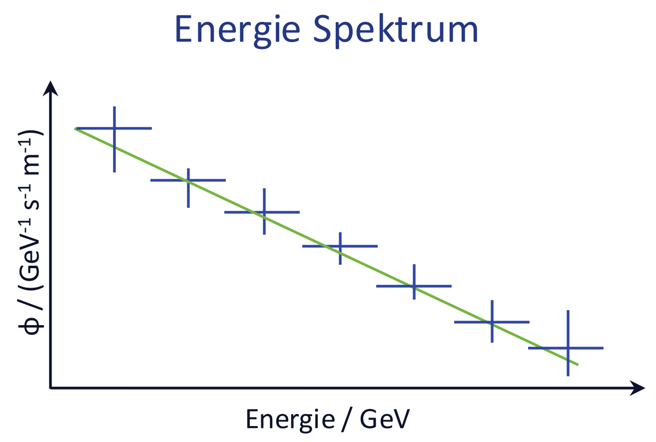

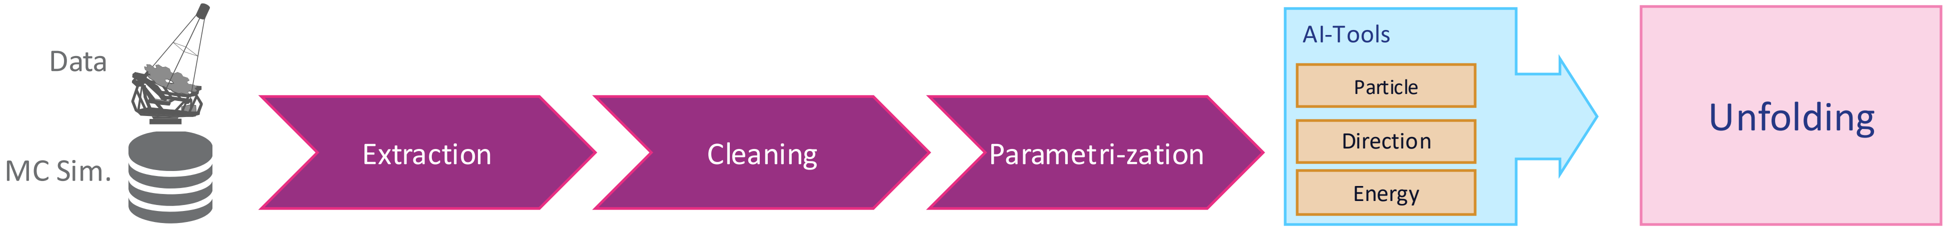

Simulations

We don't have a testbeam so we need simulations to know the truth

CORSIKA Samples

Simulation settings	Proton	Gamma
Energy range	100 GeV to 200 TeV	200 GeV to 50 TeV
Slope of the simulated spectrum	-2.7	-2.7
Viewcone	5°	0°
Event reuse	20	
Maximum impact parameter	400 m	270 m
Zenith range	$(0-30)^\circ$	
Azimuth range	$(0-0)^\circ$	
Atmosphere	Atmospheric Model 11 (MAGIC winter)	
Number of resulting events	780 046 520	18 000 000

Questions?

Which Information did you gather?

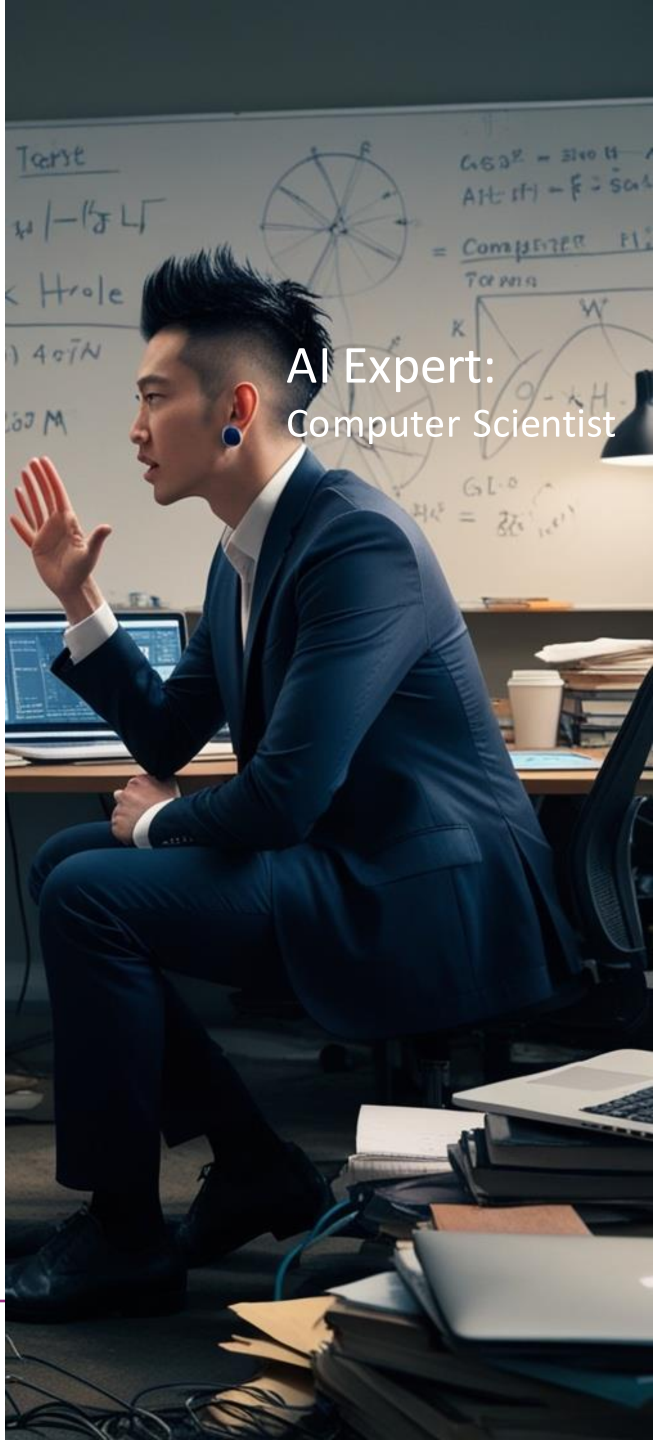

- Real-World use-case (Motivation)
- Performance metric for real-world task
- Downstream Task
- Performance metric for real-world task
- Data type: e.g. time series, images, tabular
- Learning Task
- Ground Truth / Labels

Assessment of the provided information