

Introduction to Anemoi

Jesper Dramsch

Scientist for Machine Learning


Jesper.Dramsch@ecmwf.int





Goals

- Develop an open-source framework to facilitate:
 - the development of data-driven weather forecasts...
 - ... and their running in operations
- In collaboration with the Centre's Member States and others
 - Support both global and limited area models





Goals (cont.)

- Rely on existing tools
 - PyTorch, Lightning, Hydra, Zarr, Xarray, earthkit, and many more
- Focus on best use of resources (File systems, GPUs, ...)
 - Do not starve the GPUs during training, due to slow I/Os
- Makes R2O as simple as possible
 - Meteorological evaluations
 - Delivery of code and weights to production

User Base

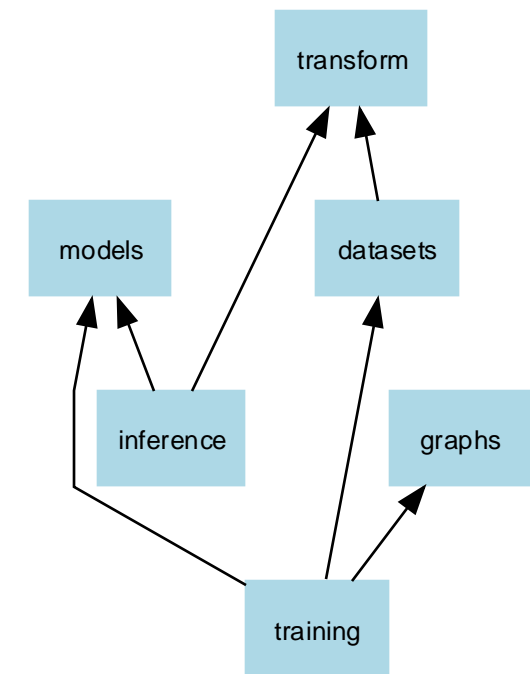
Modifies Configs for
Experimentation and
Improvement of Anemoi Model

Modifies Codebase to
implement new Features and
Augment Anemoi Libraries

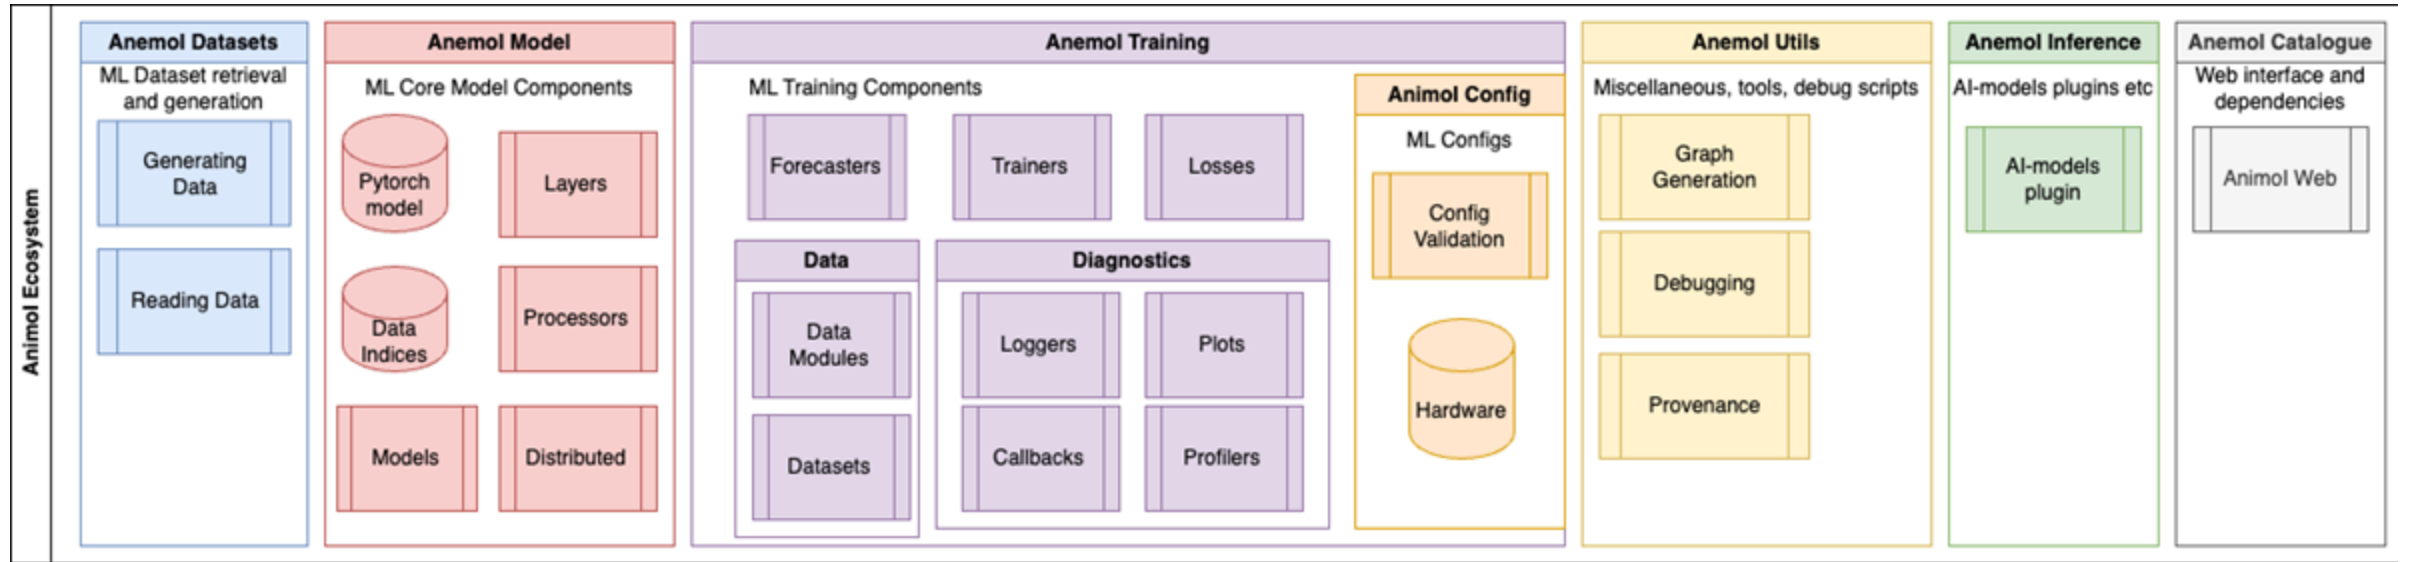
Runs the Anemoi Model in a
common interface on reliable
infrastructure

Main components - Design decisions

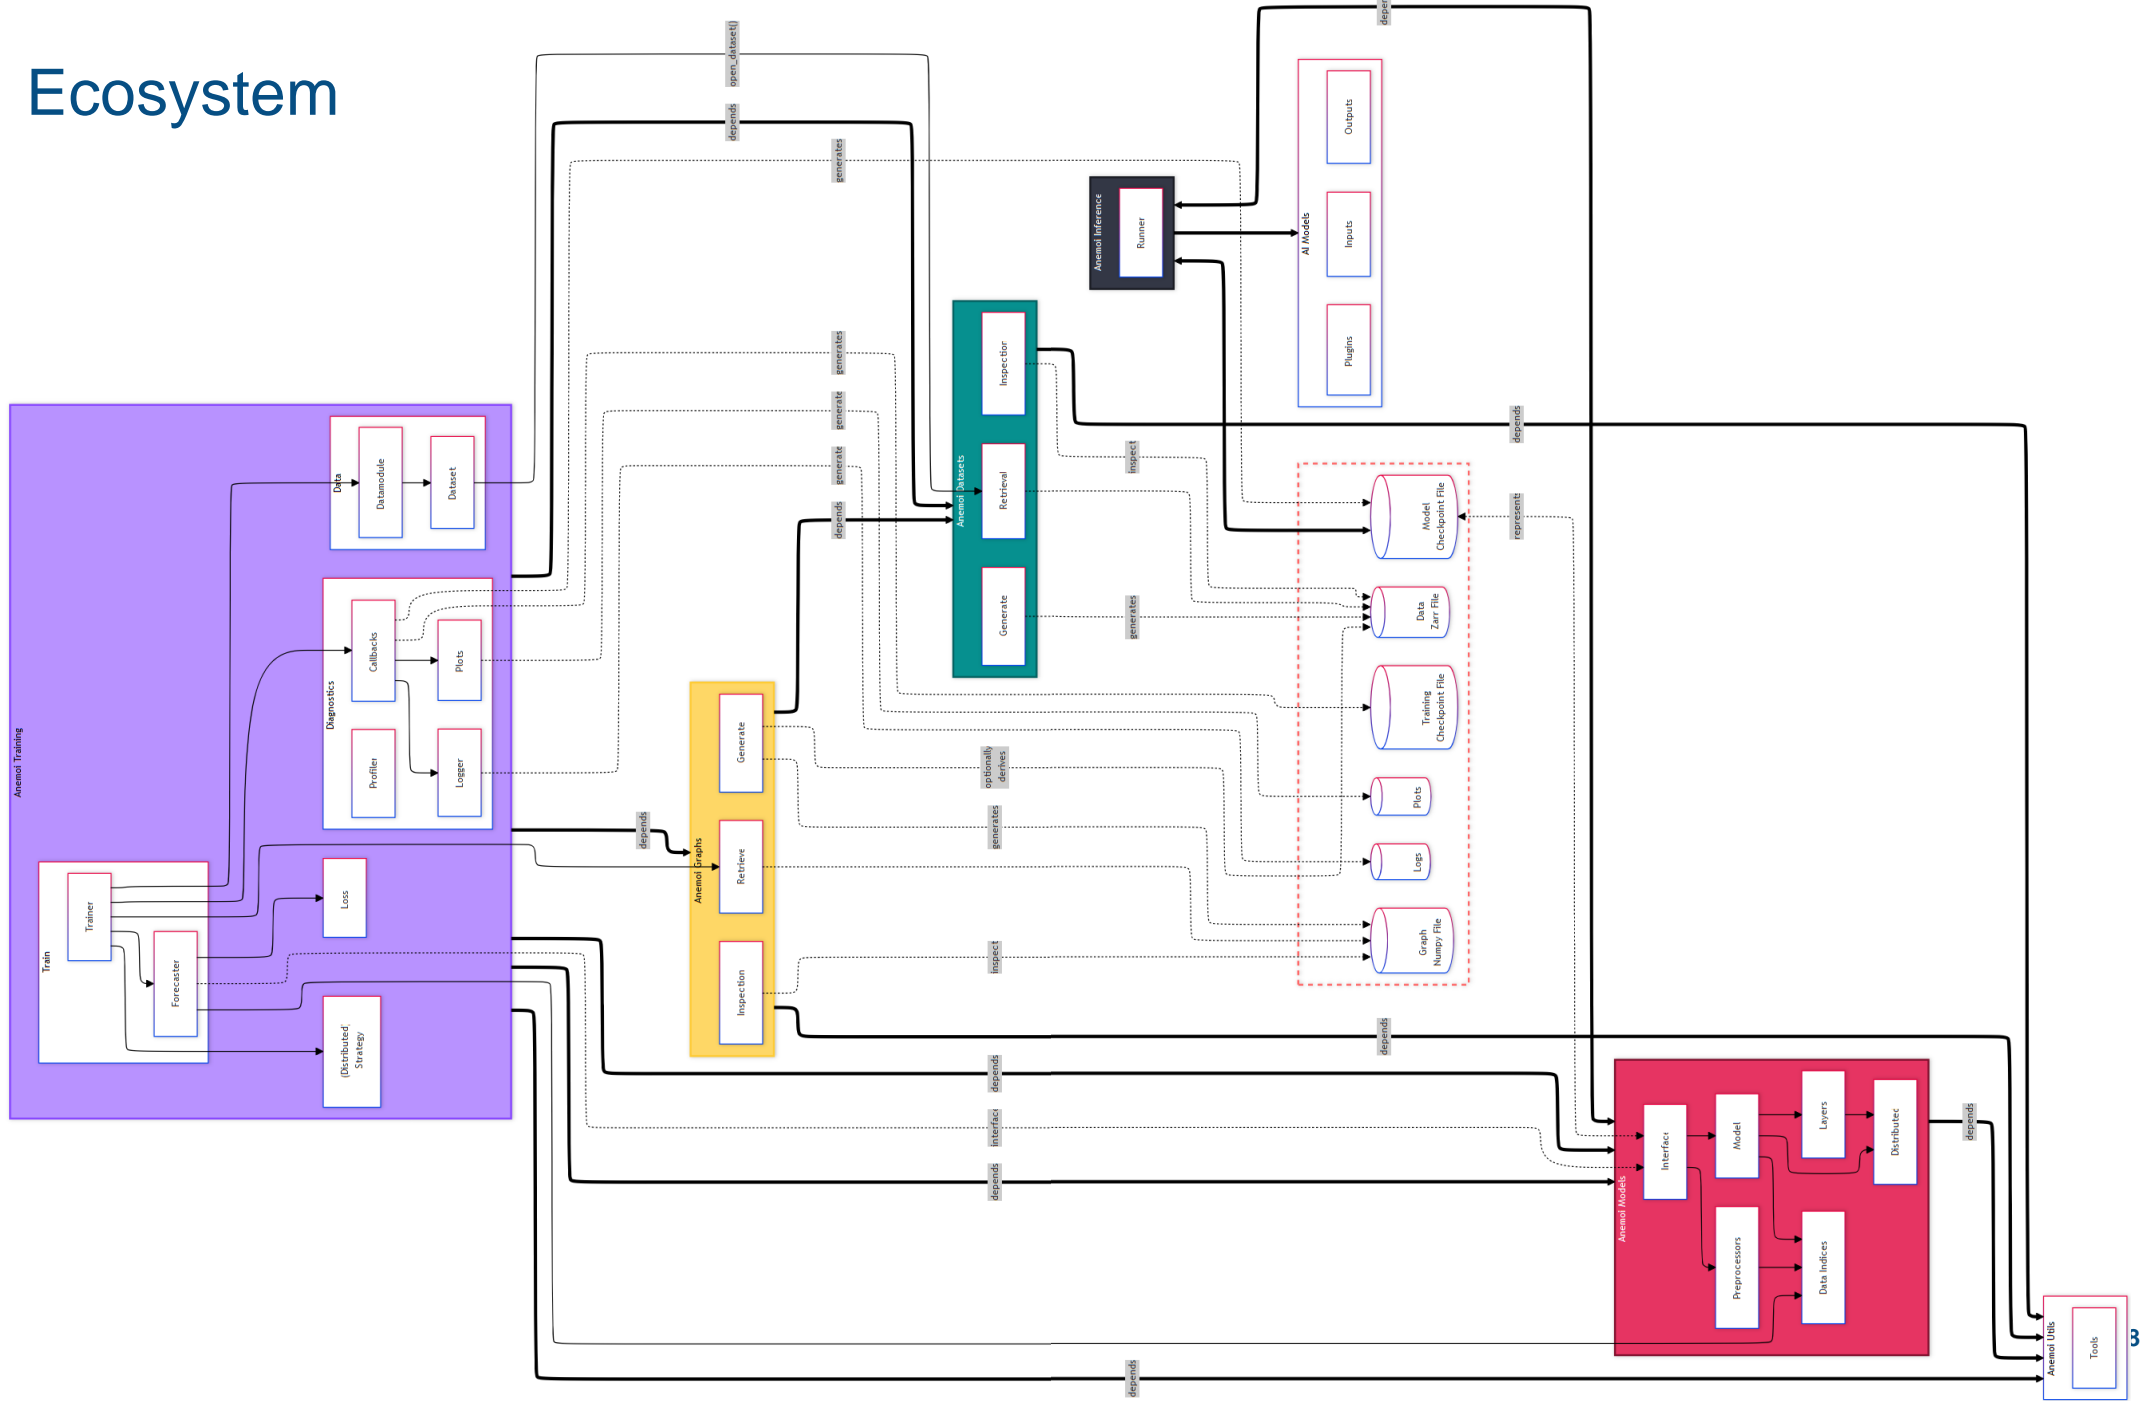
- OO design, heavy use of factories to instantiate object from config files
- Each component provides one or more command line tools
 - `anemoi-datasets create data-config.yaml out.zarr`
 - `anemoi-training train train-config.yaml`
 - `anemoi-inference run run-config.yaml`
- Minimise software dependencies to facilitate R2O
 - Inference and training are independent
- Each component collects metadata that can be used by the others



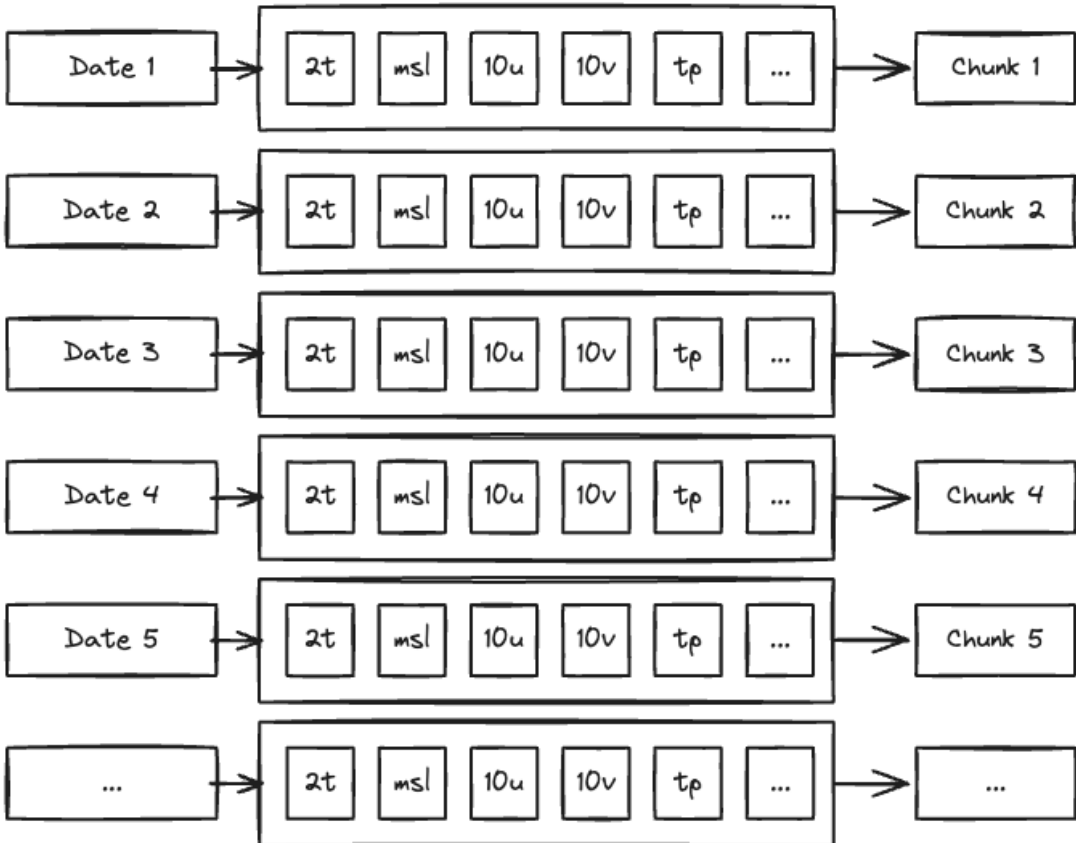

A tour through the components of Anemoi



Ecosystem



anemoi-datasets



anemoi-datasets

Goals

- Ensuring data loading is hidden beneath training costs.
- Porting **metadata** to provide **traceability** and inference use of trained models.
- Opening Zarrs within training code.
 - Handling merges of datasets on the fly.

Features

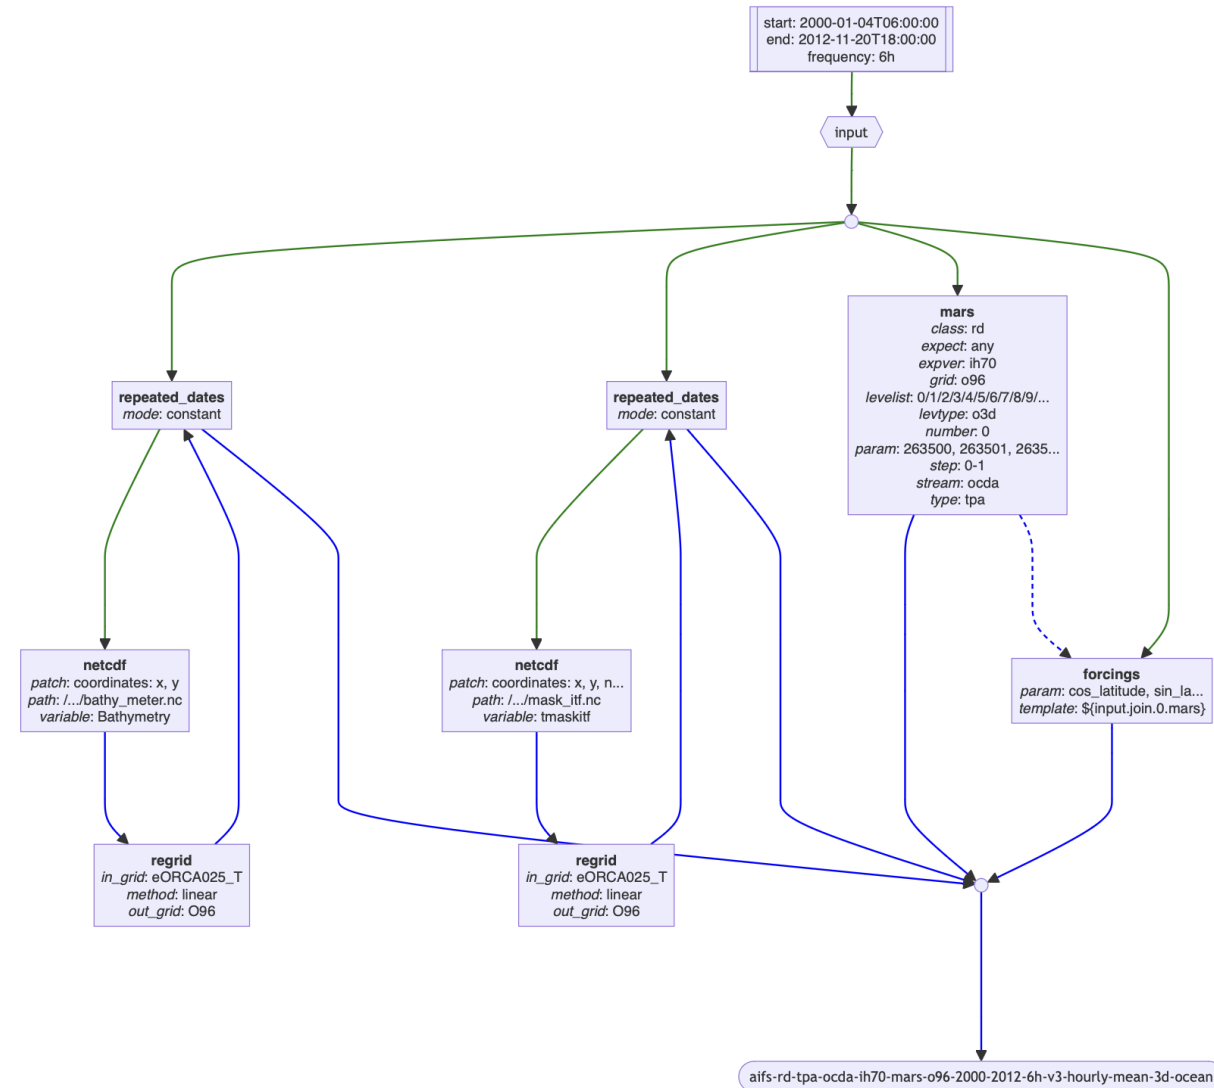
- Creation of Zarr(s) from
 - MARS/Grib/netCDF
 - Chunk for speed, not for interpretability
- Calculation of statistics for normalisation.
- Handle multiple data sources, can be combined in a single array.

How it works:
Built on ecflow and python tools
for handling data sources.

Open source, a read the docs in progress <https://anemoi-datasets.readthedocs.io/>

anemoui-datasets

- Tools to build Zarr files from a list of sources and filters
 - Supports GRIB, NetCDF, Zarr input format
 - Filters allows data transformations
 - e.g. wind speed/direction to u/v
- Easily extendable with new sources and new filters
 - See anemoui-transform
- Large datasets:
 - Building can be done in parallel
 - Building can be done incrementally




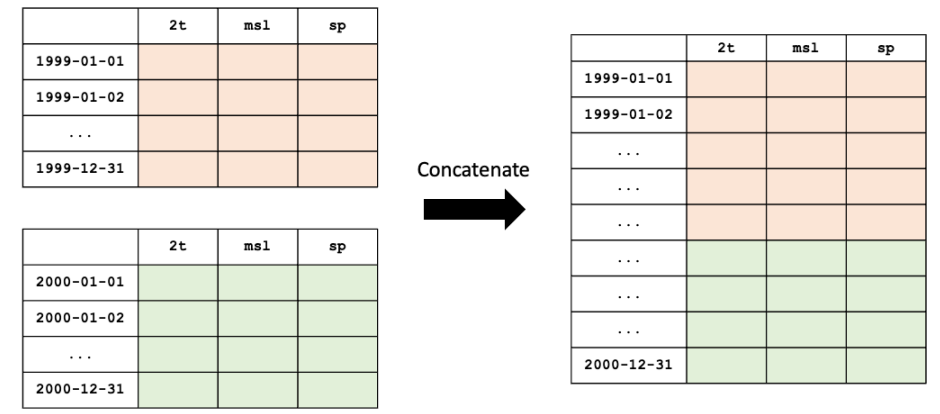
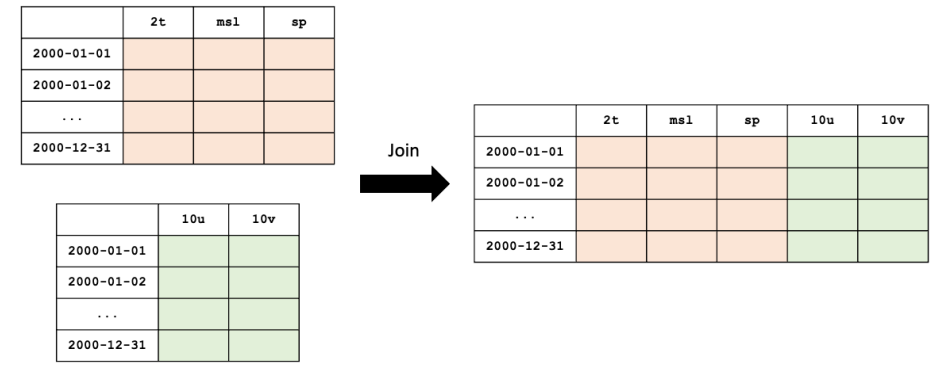
anemoi-datasets (usage)

- When used, datasets can be lazily combined into “virtual” datasets
 - Allowing researchers to find the best combinations of variables for their training needs

```

from anemoi.datasets import open_dataset

ds = open_dataset(
    [
        {
            "dataset": "aifs-od-an-oper-0001-mars-o96-2016-2023-6h-v6",
            "drop": ["sp", "2d", "skt", "tcw", "cp"],
            "end": 2023,
            "frequency": "6h",
        },
        {
            "dataset": "aifs-od-an-oper-0001-mars-o96-2016-2023-6h-v1-precipitations",
            "end": 2023,
            "frequency": "6h",
            "rename": {"tp_0h_12h": "tp"},
            "select": ["tp_0h_12h"],
        },
    ],
    end=202311,
)
    
```





anemoui-training

- Code to train models, using torch-lightning and Hydra
 - Multi-node/multi-GPU training support
 - Deterministic training with probabilistic training coming soon...
 - Callbacks for profiling evaluating, plotting and logging intermediate results
 - Implement various losses, more can be easily added
 - Interfaces with trackers such as mlflow
- Highly configurable
- Interfaces with:
 - anemoui-dataset via data loaders
 - anemoui-models via Hydra configuration
 - anemoui-inference via metadata-rich checkpoints

anemoi-training (cont.)

```
Model / GNN.yml
activation: GELU
num_channels: 512

model:
  _target_:
  anemoi.models.models.encoder_processor_decoder
  .AnemoiModelEncProcDec

processor:
  _target_:
  anemoi.models.layers.processor.GNNProcessor
  _convert_: all
  activation: ${model.activation}
  trainable_size:
  ${model.trainable_parameters.hidden2hidden}
  sub_graph_edge_attributes:
  ${model.attributes.edges}
  num_layers: 16
  num_chunks: 2
  mlp_extra_layers: 0

encoder:
  _target_:
  anemoi.models.layers.mapper.GNNForwardMapper
```

anemoi-training train model=gnn

```
Model / GraphTransformer.yml
activation: GELU
num_channels: 1024

model:
  _target_:
  anemoi.models.models.encoder_processor_decoder.Ane
  moiModelEncProcDec

processor:
  _target_:
  anemoi.models.layers.processor.GraphTransformerPro
  cessor
  _convert_: all
  activation: ${model.activation}
  trainable_size:
  ${model.trainable_parameters.hidden2hidden}
  sub_graph_edge_attributes:
  ${model.attributes.edges}
  num_layers: 16
  num_chunks: 2
  mlp_hidden_ratio: 4 # GraphTransformer
  num_heads: 16 # GraphTransformer

encoder:
```

anemoi-training train model=graphtransformer

```
Model / Transformer.yml
activation: GELU
num_channels: 1024

model:
  _target_:
  anemoi.models.models.encoder_processor_decoder.AnemoiModelEn
  cProcDec

processor:
  _target_:
  anemoi.models.layers.processor.TransformerProcessor
  _convert_: all
  activation: ${model.activation}
  num_layers: 16
  num_chunks: 2
  mlp_hidden_ratio: 4 # Transformer only
  num_heads: 16 # Transformer only
  window_size: 512
  dropout_p: 0.0

encoder:
  _target_: anemoi.models.layers.mapper.GraphTransformerFor
```


anemoi-training train model=transformer

- Make it easy to switch components
- Allow for reproduceable training
- Easy to extend with new models and components





anemoi-training

Example: current AIFS configuration


Graph Encoder



Graph Decoder



Anemoi.training – MLFlow servers for Experiment Tracking







anemoi-models


- Building blocks to implement data-driven forecasting systems
 - Models
 - Blocks (e.g. encoders/decoders/processors)
 - Layers
- Design choice: graph-based (or similar) outer layers.
 - Provides high levels of flexibility on input & output data.
 - Inner layers/blocks can be either attention- or graph-based.
- Multi-node, multi-GPUs to scale to large model
 - Multiple model-sharded implementations supported, e.g. 1-hop graph, head sharding.
 - Ensures model design is not limited by GPU-memory.
- Supports GraphCast-style, Oskarson-style and AIFS models.
- Further models can be easily created

anemoui-graphs

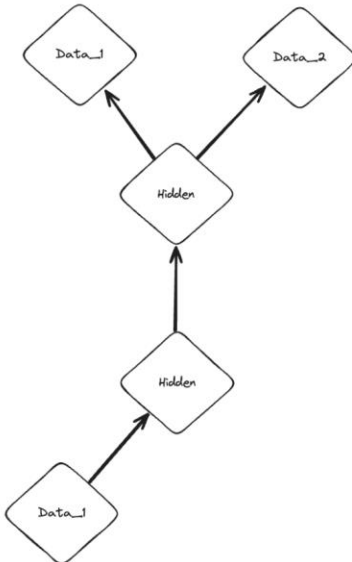
- Highly configurable
- Integrated with Anemoui ecosystem
- Includes graph inspection tools
- Optimized for GPU
- Supports multiple applications:
 - Global weather forecasting
 - Limited area modelling
 - Stretched-grid graphs
 - Hierarchical graphs
 - Dynamic graphs in progress




What comes next?




Hierarchical graph



Multiple decoder

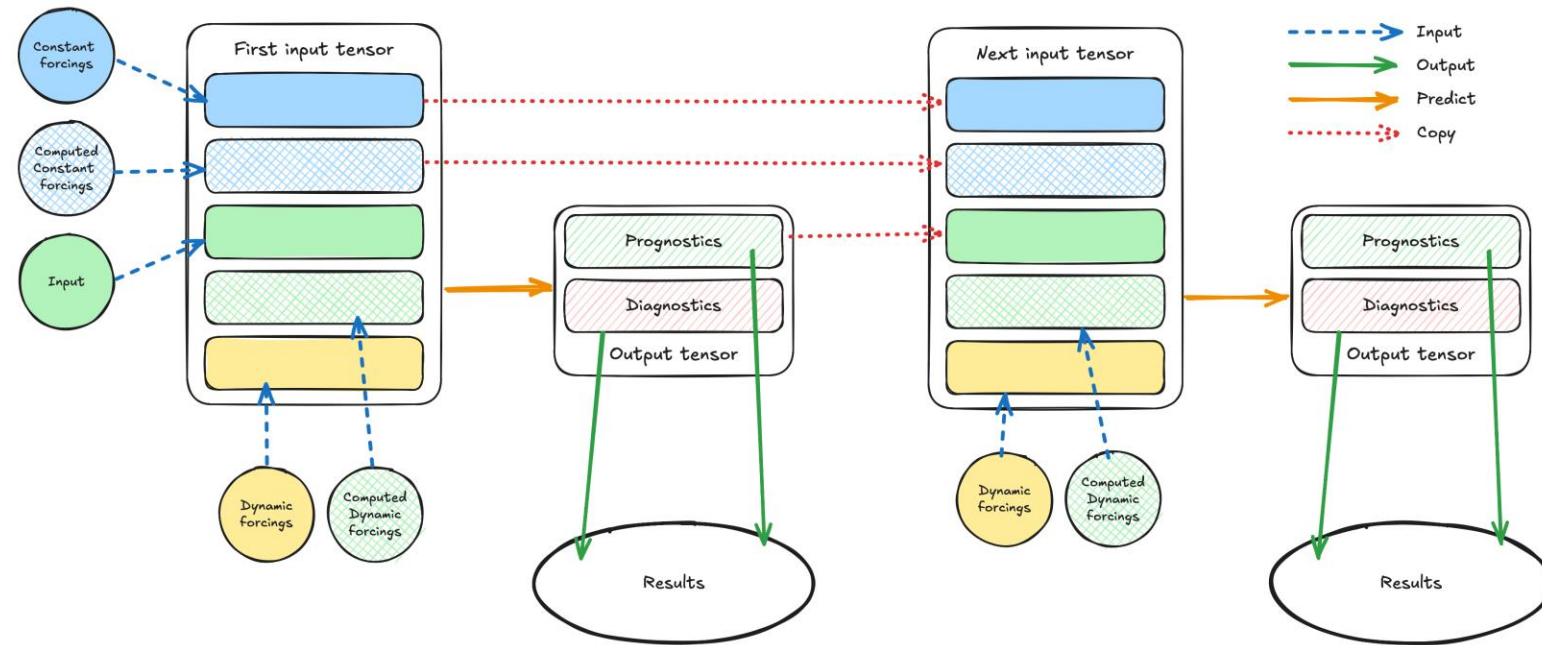


Multiple encoder



Multiple encoder & decoder

anemoi-inference



- Running trained models in inference mode
- Auto-discovery of the model parameters from the metadata embedded in the checkpoint
 - List of prognostic, diagnostic and forcing variables
- Provide full control of input and output to the user
 - E.g. using the output of physics-based NWP analyses or ensembles, as initial conditions
- In progress: allowing coupling of two or more data-driven forecasts
 - E.g. ocean/atmosphere

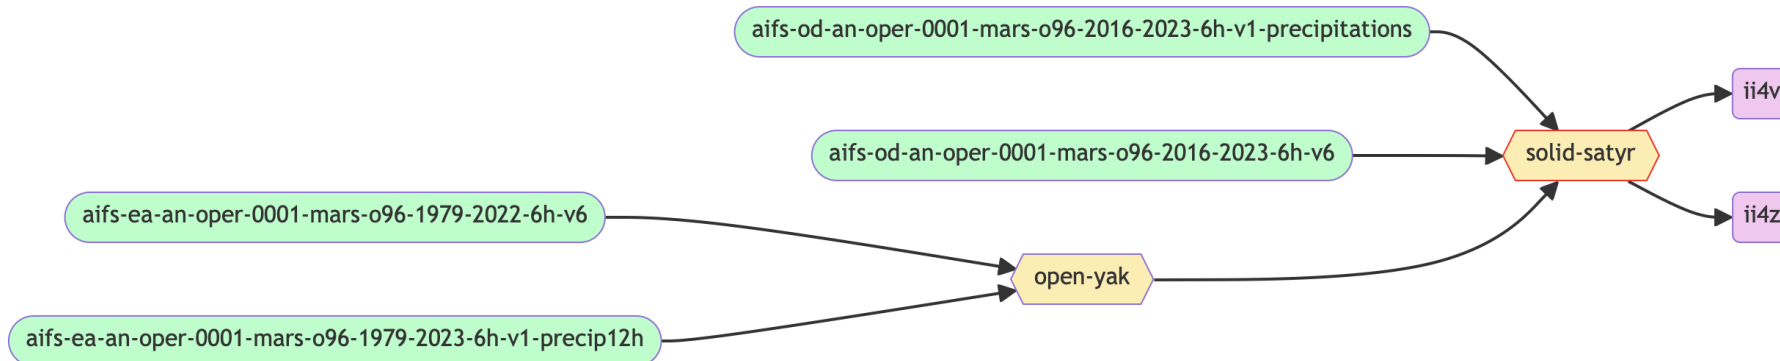
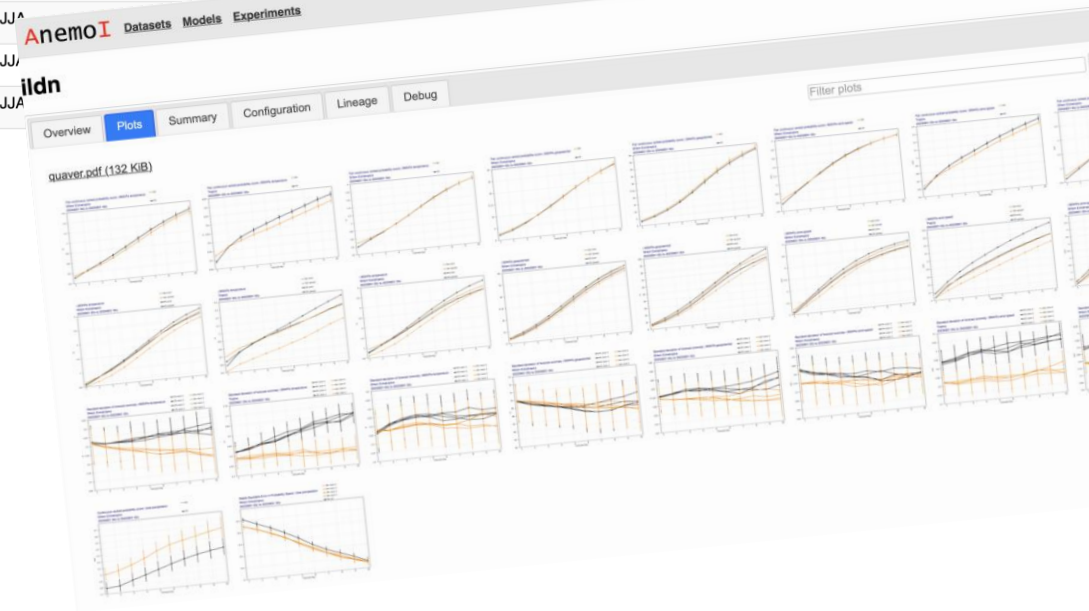
anemoi-registry

- code to manage a catalogue
 - training datasets,
 - model configurations,
 - weights, model evaluations, etc,
- With a strong emphasis on recording lineage dependencies.

AnemoI Datasets Models Experiments

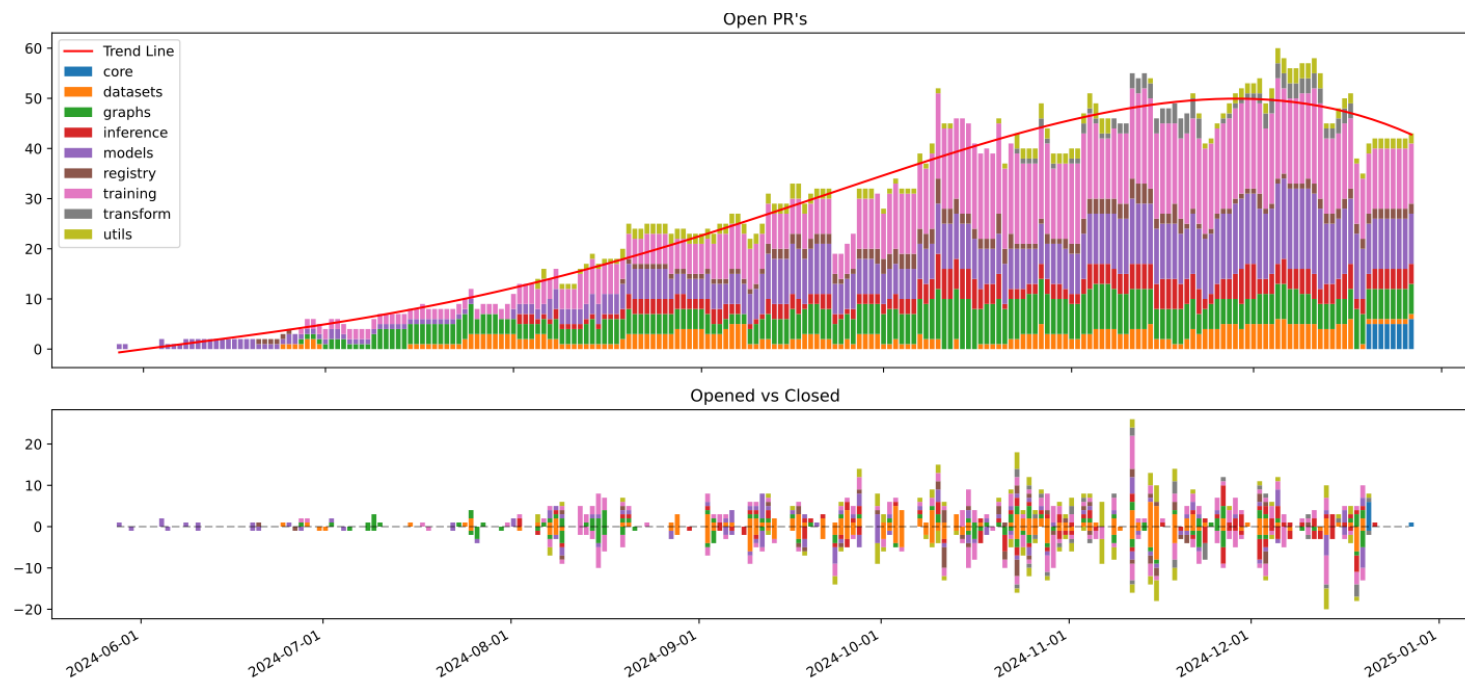
Show 20 entries Search:

Name	Owner	Type	Dates	Runs	Created	Description
ilee	Sara	single	2023 JJA	183	2025-01-03	202306-08, aifs with waves oper finetune, from era5, n320
ileh	Sara	single	2023 JJA	183	2025-01-03	202306-08, aifs with waves oper finetune, from oper, n320
ijbm	Gareth	ensemble	2020 JJ	31	2025-01-02	Test ensemble
ile2	Jakob	ensemble	2023 JJA	92	2025-01-02	Ensemble
ile1	Jakob	ensemble	2023 JJ^			
ile0	Jakob	ensemble	2023 JJ			
ildn	Mariana	ensemble	2023 JJA			




Collaboration

- Code reviews via pull requests
- Pre-commit checks
- GitHub actions
- CI/CD
- Anyone can contribute!





What Anemoi enables?



Training deterministic & ensemble models

Lang et al. 2024a & b

Building coupled earth-system models




Easy incorporation of more atmospheric fields

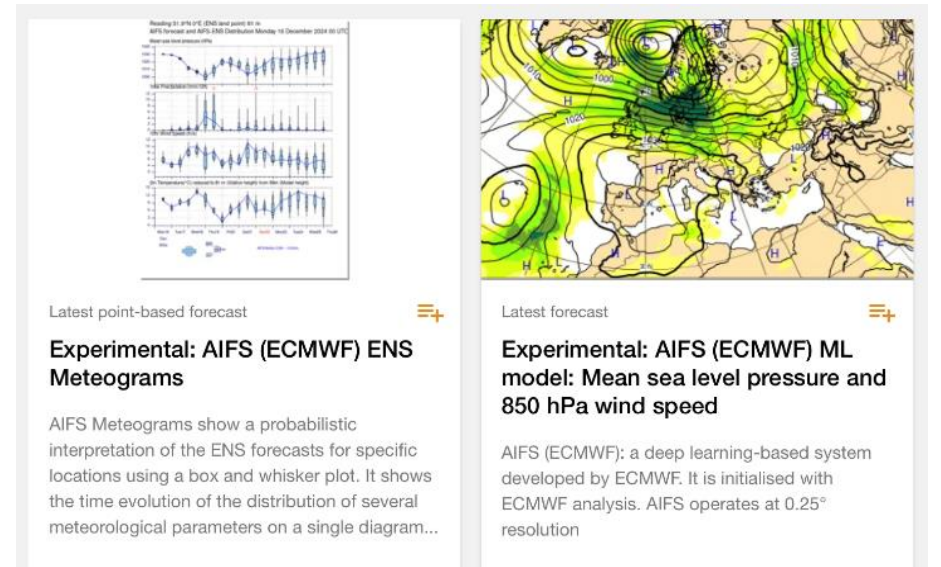
Regional high-resolution modelling:

See Bris poster (970) on Thursday afternoon


Nipen et al. 2024



Real-time running at operation centres
Anemoi will power ECMWF's operationalisation of AIFS in 2025.




Collaboration, without aggregation around a single model




Roadmap – what's coming next for Anemoi

- Ensembles, supporting score optimisation & diffusion training



Roadmap – what's coming next for Anemoi

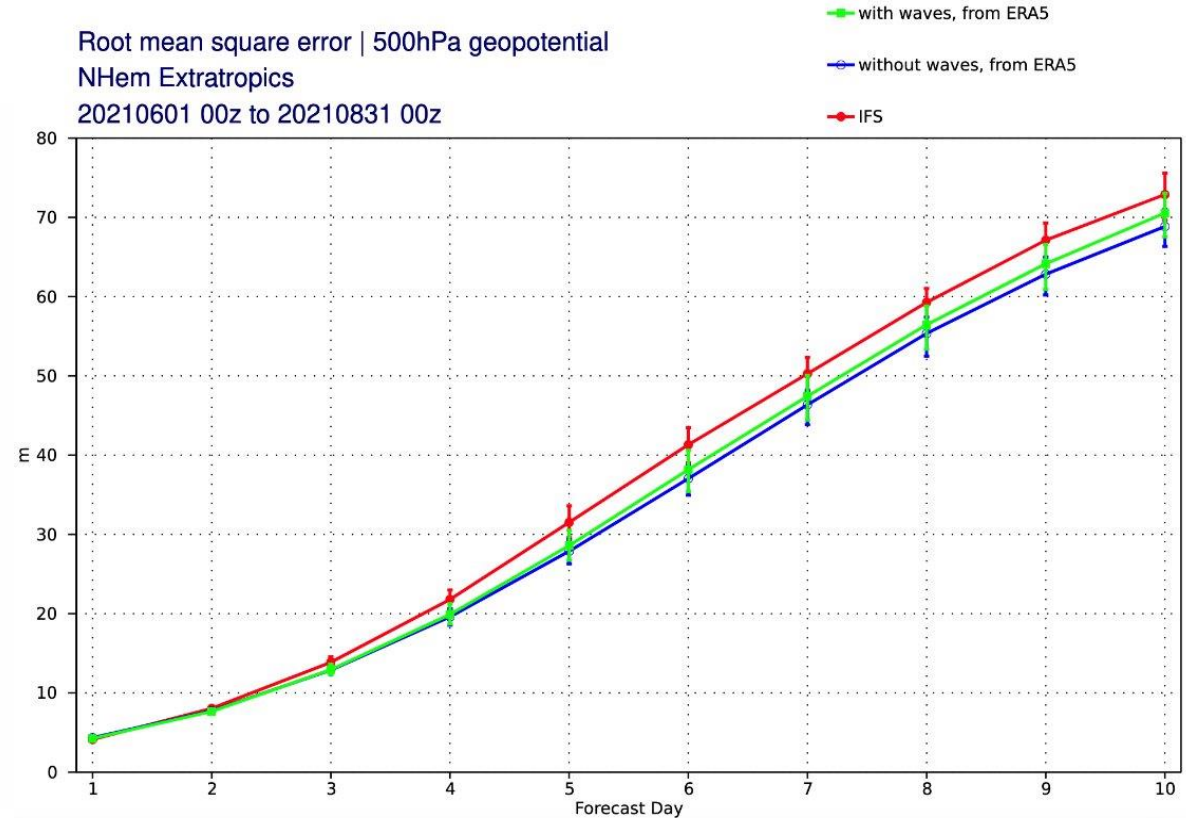

- Dynamic graphs for training towards observations



Roadmap – Earth system components

- Preliminary Result: 14-day forecast of wave fields without deteriorating the performance in the atmosphere (o96)

Mean wave direction (14 day forecast)



the lower the better



What's next?

- Ensembles: both CRPS-optimization and diffusion (similar to GenCast)
- Better support for observations
- GPU-parallel inference
- Support for environmental data (ocean, land, waves)
- Global/LAM coupling
- Automated training and validation

Thank you