@ashwinvis
ashwinvishnu
https://fluid.quest
@ashwinvis@mastodon.acc.sunet.se

ASHWIN VISHNU MOHANAN / 2023-11-21

Towards use of AI as an accessory and accelerator at SMHI

About me (or how I got here)

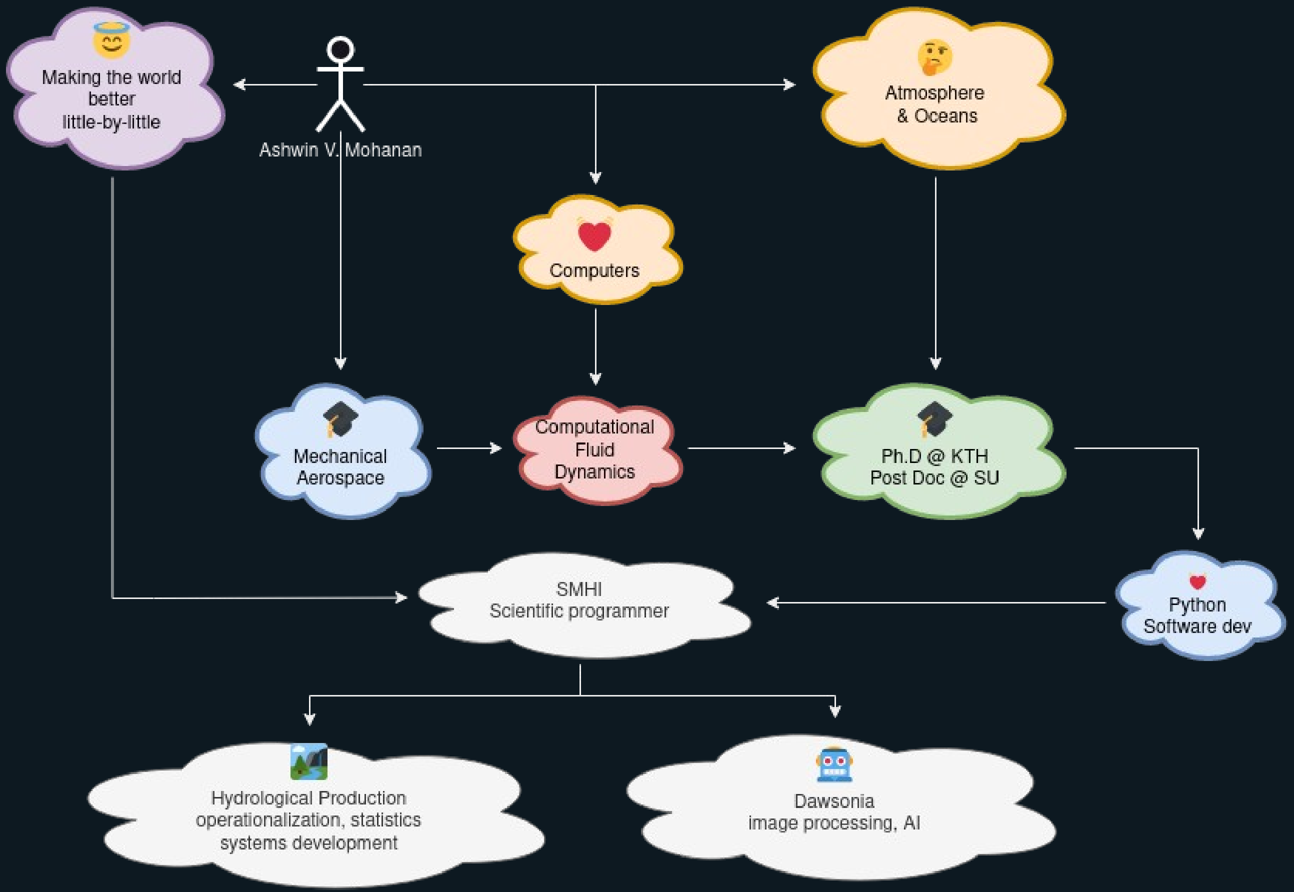

Acts

I. Dawn of AI in weather
II. Dawsonia: (Digitize hAndWritten obServatiONs In weather journAls)
III. The HPC perspective

Dawn of AI in weather services

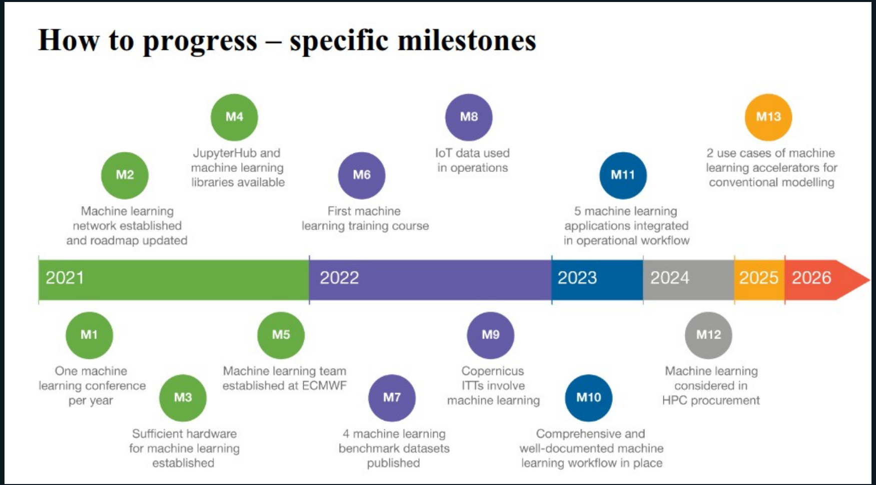

- ECMWF starts a roadmap for ML on 2021
- Planned to focus on infrastructure, data and training in beginning
- First lightweight ML applications in 2023
- ML in modelling (i.e. forecast) *was planned* for 2025

Source: https://doi.org/10.21957/ge7ckgm

Dawn of AI in weather services

Ryan Keisler @RyanKeisler 16 Feb 2022

📢 Time to share a project I've been working on:

Forecasting Global Weather with Graph Neural Networks

🧵 (1/N)

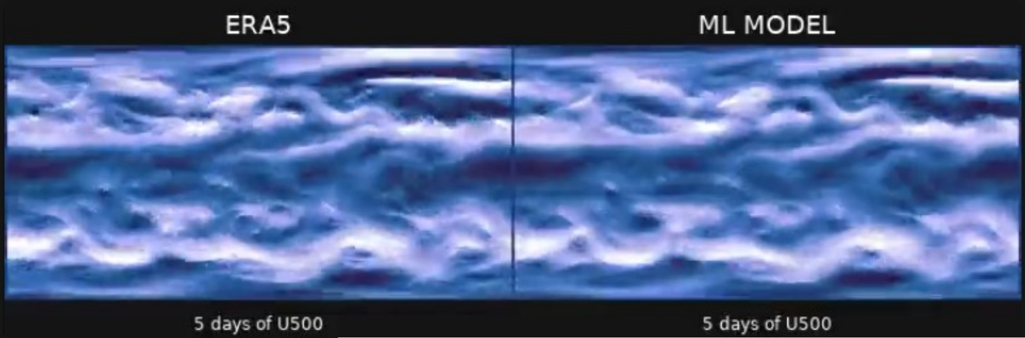

💬 43 🔁 315 ❞ 63 ♥ 2,

Fair use:
@RyanKeisler on Twitter
https://arxiv.org/pdf/2202.07575.pdf

3.3 Training

We trained our final model using the Adam optimizer and a 3-round training schedule with progressively smaller learning rates: 3.5 days of training at lr=3e-4, 1 day at lr=3e-5, and 1 day at lr=3e-6. The total training time was 5.5 days on a single NVIDIA A100 GPU, which cost approximately $370 using the Google Cloud Platform. The training procedure used multi-resolution training data, a multi-step loss, and a specific loss normalization, as described below.

Dawn of AI in weather services

SMHI

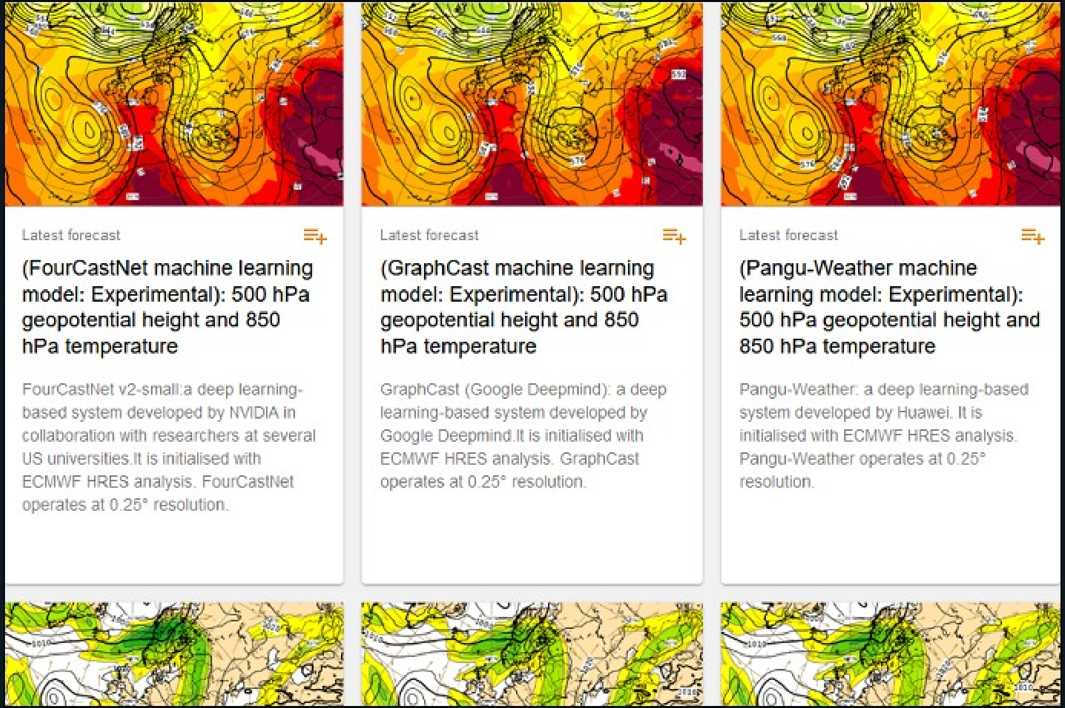

[Submitted on 22 Feb 2022]

FourCastNet: A Global Data-driven High-resolution Weather Model using Adaptive Fourier Neural Operators

Jaideep Pathak, Shashank Subramanian, Peter Harrington, Sanjeev Raja, Ashesh Chattopadhyay, Morteza Mardani, Thorsten Kurth, David Hall, Zongyi Li, Kamyar Azizzadenesheli, Pedram Hassanzadeh, Karthik Kashinath, Animashree Anandkumar

[Submitted on 3 Nov 2022]

Pangu-Weather: A 3D High-Resolution Model for Fast and Accurate Global Weather Forecast

Kaifeng Bi, Lingxi Xie, Hengheng Zhang, Xin Chen, Xiaotao Gu, Qi Tian

[Submitted on 24 Dec 2022 (v1), last revised 4 Aug 2023 (this version, v2)]

GraphCast: Learning skillful medium-range global weather forecasting

Remi Lam, Alvaro Sanchez-Gonzalez, Matthew Willson, Peter Wirnsberger, Meire Fortunato, Ferran Alet, Suman Ravuri, Timo Ewalds, Zach Eaton-Rosen, Weihua Hu, Alexander Merose, Stephan Hoyer, George Holland, Oriol Vinyals, Jacklynn Stott, Alexander Pritzel, Shakir Mohamed, Peter Battaglia

Fair use: https://www.ecmwf.int/en/about/media-centre/news/2023/how-ai-models-are-transforming-weather-forecasting-showcase-data

See also: https://charts.ecmwf.int

Dawn of AI in weather and climate

SMHI

An active area, even today

Science

News Home | All News | ScienceInsider | News Features

HOME > NEWS > ALL NEWS > AI CHURNS OUT LIGHTNING-FAST FORECASTS AS GOOD AS THE WEATHER AGENCIES'

NEWS | CLIMATE

AI churns out lightning-fast forecasts as good as the weather agencies'

Running in mere minutes, AI forecasts are surpassing supercomputers in speed and accuracy

14 NOV 2023 • 10:00 AM ET • BY PAUL VOOSEN

npj | climate and atmospheric science

Explore content | About the journal | Publish with us

nature > npj climate and atmospheric science > articles > article

Article | Open access | Published: 16 November 2023

FuXi: a cascade machine learning forecasting system for 15-day global weather forecast

Lei Chen, Xiaohui Zhong, Feng Zhang, Yuan Cheng, Yinghui Xu, Yuan Qi & Hao Li

npj Climate and Atmospheric Science **6**, Article number: 190 (2023) | Cite this article

261 Accesses | **4** Altmetric | Metrics

Examples of AI research at SMHI

SMHI

- Digitization of old observation data (what I do and part of this talk)
- Quality control for data from personal weather stations using AI (EUMETNET)
- Text generation
- Limited Area Models for forecasting in the Nordics and Baltics: MEPS (MetCoOp Ensemble Prediction System)

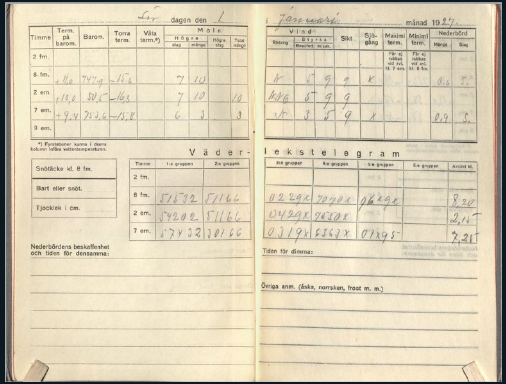

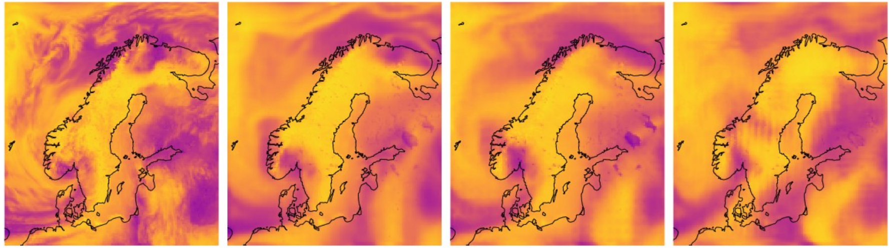

(a) Ground Truth (b) Hi-LAM (c) GC-LAM (d) 1L-LAM

Figure 3: Ground truth and example forecasts of `nlwrs` at lead time 57 h.

[Submitted on 29 Sep 2023 (v1), last revised 14 Nov 2023 (this version, v2)]

Graph-based Neural Weather Prediction for Limited Area Modeling

Joel Oskarsson, Tomas Landelius, Fredrik Lindsten

Sources:
Internal presentation by A. Yong & M. Hansson
https://arxiv.org/abs/2309.17370

The hope

SMHI

- **Faster Weather Predictions -** faster and more immediate weather predictions, crucial during rapidly evolving weather situation
- **More Accurate Predictions -** enabling individuals and organizations to make better decisions
- **Where do humans fit in with AI weather prediction? -** humans remain essential for interpreting and communicating the impact
- **Will AI replace Meteorologist?** No, the combination of humans and AI will be a unique and effective combination that will result in a very robust system

Sources:
Internal presentation by A. Yong & M. Hansson

Python package: `dawsonia`

Aim to develop a semi-automated process to digitize handwritten weather journals from PDF to machine-readable data

https://dawsonia.readthedocs.io/en/latest/

https://git.smhi.se/ai-for-obs/dawsonia

Data rescue: an unsolved problem

SMHI

Ed Hawkins
@ed_hawkins@fediscience.org

Ever wondered about how much old weather data is not available to climate scientists to use?

Just in the UK we have around 100 million daily rainfall measurements for the period 1850-1960 sitting in the Archives on hundreds of thousands of sheets of paper.

Where is the AI tool to read these?!

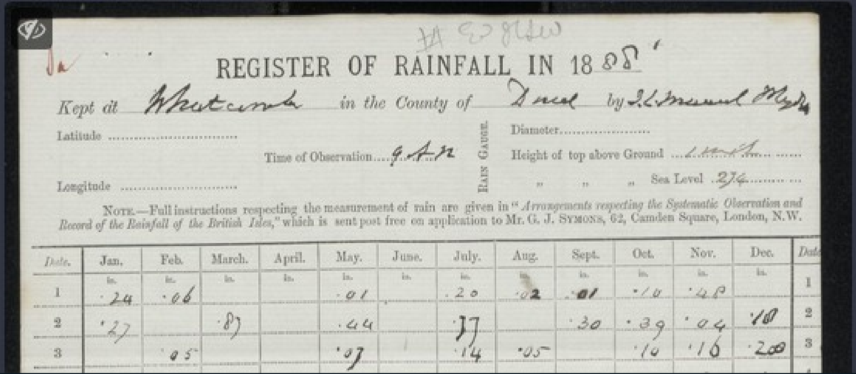

"Because the records are handwritten, there is no possibility to automate the digitization process. - Skrynyk et al. (2020)

Fair use:
@ed_hawkins@fediscience.org
Sources:
https://doi.org/10.15407/uhmi.report.01
See also:
https://datarescue.climate.copernicus.eu/best_practice_guidelines

SMHI

Pipeline (step-by-step transformation)

PDF file →

Preprocess:
- Remove noise
- Straighten
- Improve contrast

→

Table detection:
- Subtract lines
- Detect boxes
- Save position

→

OCR with ML:
- Detect numbers

→

Output:
- Values
- Confidence

Pipeline (step-by-step transformation)

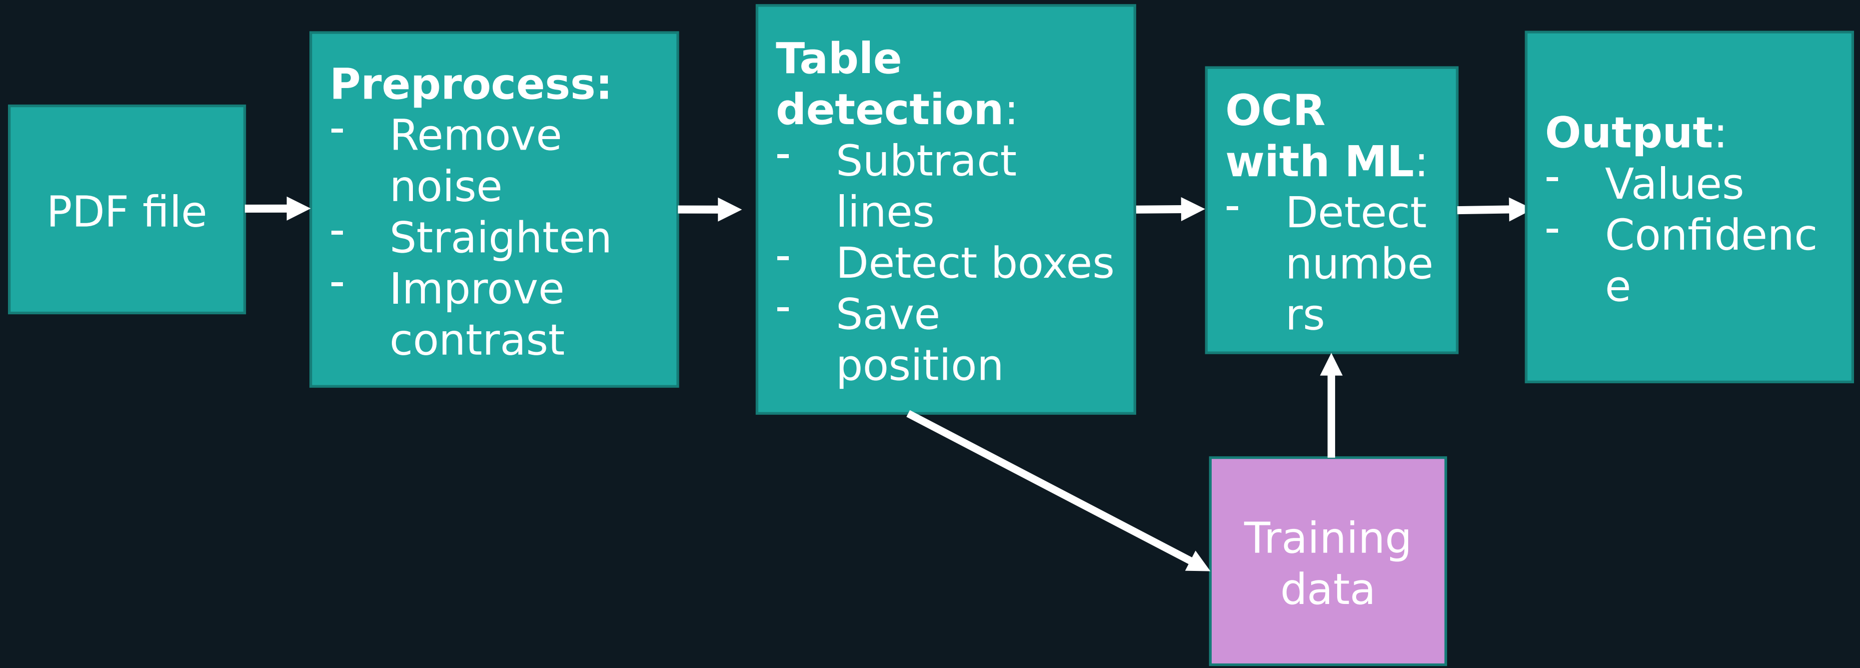

dawsonia label

Preprocessing and table detection

Original scan

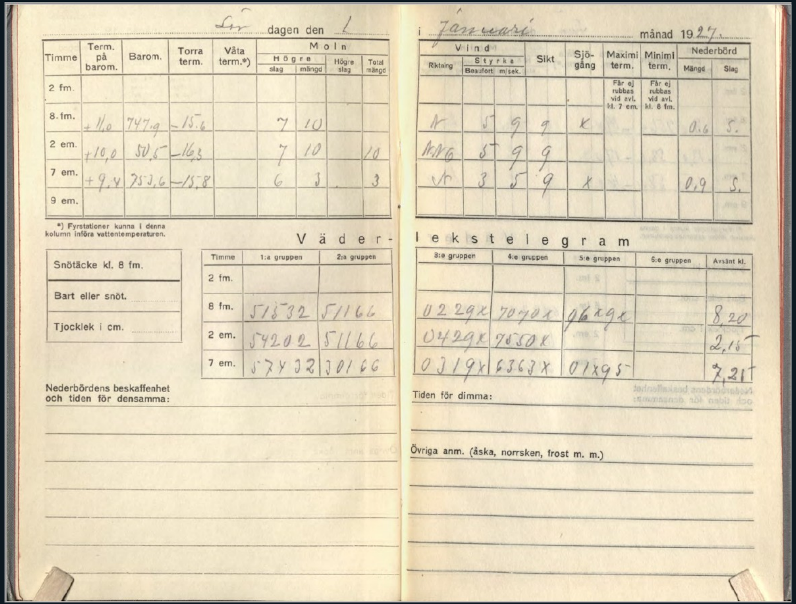

Lör dagen den 1 i Januari månad 1927.

Timme	Term. på barom.	Barom.	Torra term.	Våta term.*)	Moln: Högre slag	Moln: Högre mängd	Högre slag	Total mängd
2 fm.								
8 fm.	+11,0	747,9	−15,6		7	10		
2 em.	+10,0	50,5	−16,3		7	10		10
7 em.	+9,4	753,6	−15,8		6	3	.	3
9 em.								

Vind: Riktning	Styrka: Beaufort	Styrka: m/sek.	Sikt	Sjögång	Maximi term. (Får ej rubbas vid avl. kl. 7 em.)	Minimi term. (Får ej rubbas vid avl. kl. 8 fm.)	Nederbörd: Mängd	Nederbörd: Slag
N	5	9	9	K			0,6	S.
NNO	5	9	9					
N	3	5	9	K			0,9	S.

*) Fyrstationer kunna i denna kolumn införa vattentemperaturen.

Snötäcke kl. 8 fm.

Bart eller snöt.

Tjocklek i cm.

Väder-lekstelegram

Timme	1:a gruppen	2:a gruppen	3:e gruppen	4:e gruppen	5:e gruppen	6:e gruppen	Avsänt kl.
2 fm.							
8 fm.	51532	51166	0229x	7070x	96x9x		8,20
2 em.	54202	51166	0429x	7650x			2,15
7 em.	57432	30166	0319x	6363x	01x95		7,21

Nederbördens beskaffenhet och tiden för densamma:

Tiden för dimma:

Övriga anm. (åska, norrsken, frost m. m.)

Detect boxes to identify tabular structure

dagen den 9 i Mars månad 19 36

Tid	Term. på barom.	Barometern	Torra term.	Våta term.*)	Moln: Låga slag	Moln: Låga mängd	Moln: Medelhöga slag	Moln: Höga slag	Total molnighet, solsken, dimma och nederbörd	Vind: Riktning	Vind: Styrka Beaufort	Vind: Styrka m/sek.	Sikt	Sjögång	Maximiterm.	Minimiterm.	Nederbörd: Mängd	Nederbörd: Slag
kl. 2															Får ej rubbas vid avl. kl. 19!	Får ej rubbas vid avl. kl. 8!		
kl. 8	+16	758,2	-2		5	10	x	x	10	S	6	10	4	x				
kl. 14	+14,5	57,1	-2,2		5	10	x	x	10	SW	6	10	4	x				
kl. 19	+13	56,3	-5		x	x	x	x	10	SW	3	5	9	x				
kl. 21	+15	55,9	-4,5		x	x	x	x	8	SW	2	3	9	x			**)	

*) Fyrstationer, som sakna psykrometer, kunna i denna kolumn införa vattentemperaturen.

**) Här införes avläsning på minimitermometerns spritpelare.

Väderlekstelegram

Snötäcke kl. 8

B, BS, SB eller S: ....

Tjocklek i cm.: ....

Tid	1:a gruppen	2:a gruppen	3:e gruppen	4:e gruppen	5:e gruppen	6:e gruppen	Avsänt kl.
kl. 2							
kl. 8	0625x	08458	16628	13352	xx813	00x49	8
kl. 14	0625x	08468	20628	13252	xx600		14
kl. 19	062xx	03949	20328	12255	xx803	00x99	19

Nederbördens beskaffenhet och tiden för densamma:

Tiden för dimma:

Övriga anm. (åska, regn, norrsken, frost m. m.):

corr_rotate: end

im.min()=0.0 im.max()=0.8662192 median(im)=0.7622919 im.dtype=dtype('float32')

dagen den 9 i Mars månad 19 36

Tid	Term. på barom.	Barometern	Torra term.	Våta term.*)	Moln: Låga slag	Moln: Låga mängd	Moln: Medel-höga slag	Moln: Höga slag
kl. 2								
kl. 8	+16	758,2	-2		5	10	x	L
kl. 14	+14,5	57,1	-2,2		5	10	x	x
kl. 19	+13	56,3	-5		x	x	x	x
kl. 21	+15	55,9	-4,5		x	x	x	x

Total molnighet, solsken, dimma och nederbörd	Vind: Riktning	Vind: Styrka Beaufort	Vind: Styrka m/sek.	Sikt	Sjögång	Maxi-mi-term.	Mini-mi-term.	Nederbörd: Mängd	Nederbörd: Slag
						Får ej rubbas vid avl. kl. 19!	Får ej rubbas vid avl. kl. 8!		
10	S	6	10	4	x				
10	SW	6	10	4	x				
10	SW	3	5	9	x				
8	SW	2	3	9	x			**)	

*) Fyrstationer, som sakna psykrometer, kunna i denna kolumn införa vattentemperaturen.

**) Här införes avläsning på minimitermometerns spritpelare.

Väder-lekstelegram

Tid	1:a gruppen	2:a gruppen	3:e gruppen	4:e gruppen	5:e gruppen	6:e gruppen	Avsänt kl.
kl. 2							
kl. 8	0625x	08458	16628	13352	xx813	00x49	8
kl. 14	0625x	08468	20628	13252	xx600		14
kl. 19	062xx	03949	20328	12255	xx803	00x99	19

Snötäcke kl. 8

B, BS, SB eller S:

Tjocklek i cm.:

Nederbördens beskaffenhet och tiden för densamma:

Tiden för dimma:

Övriga anm. (åska, regn, norrsken, frost m. m.):

ligne_filter

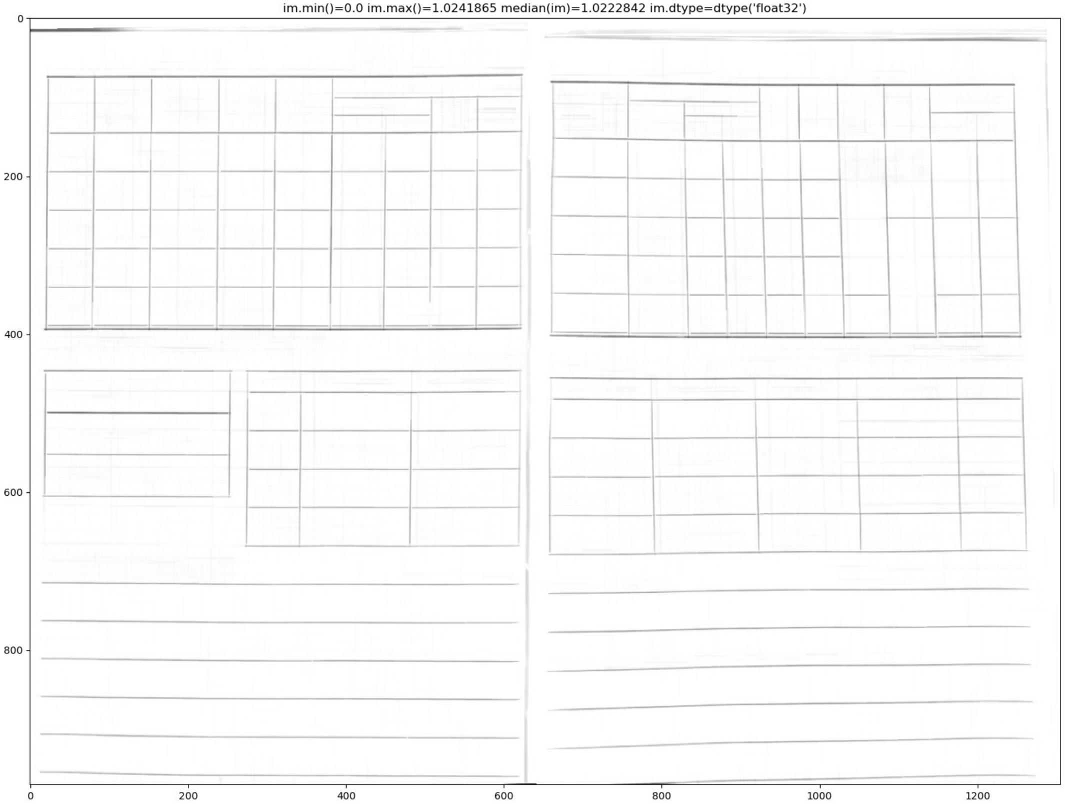

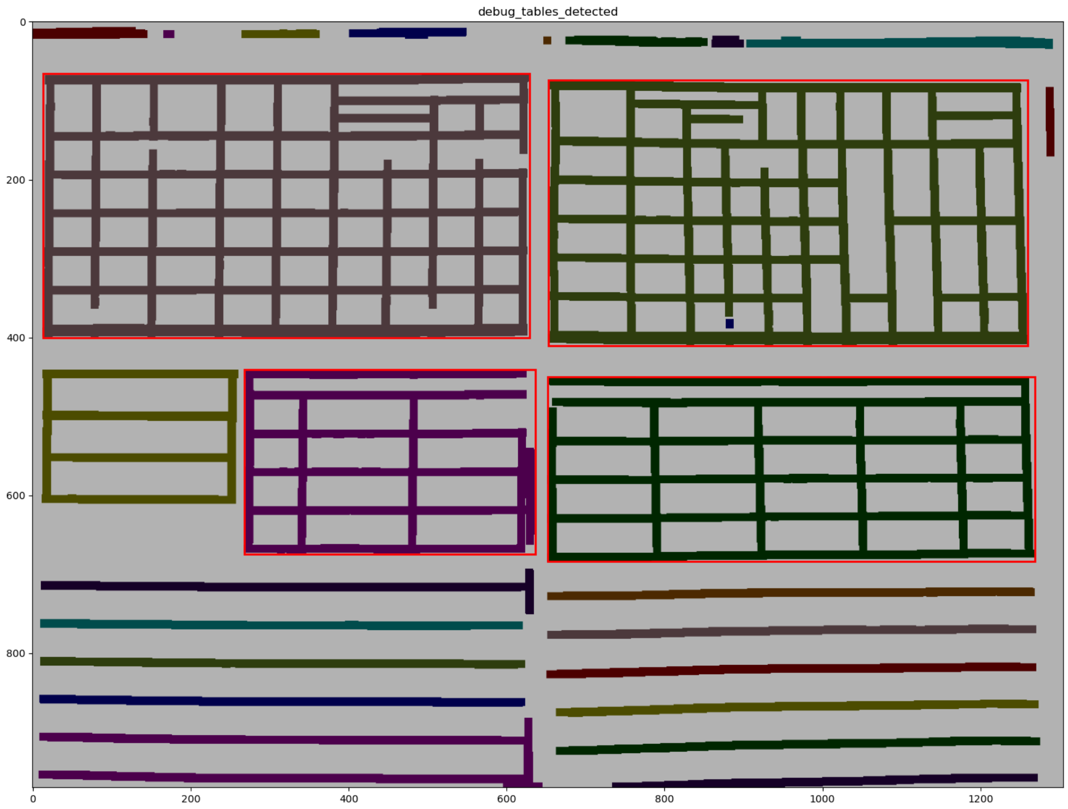
debug_tables_detected
0
200
400
600
800
0
200
400
600
800
1000
1200

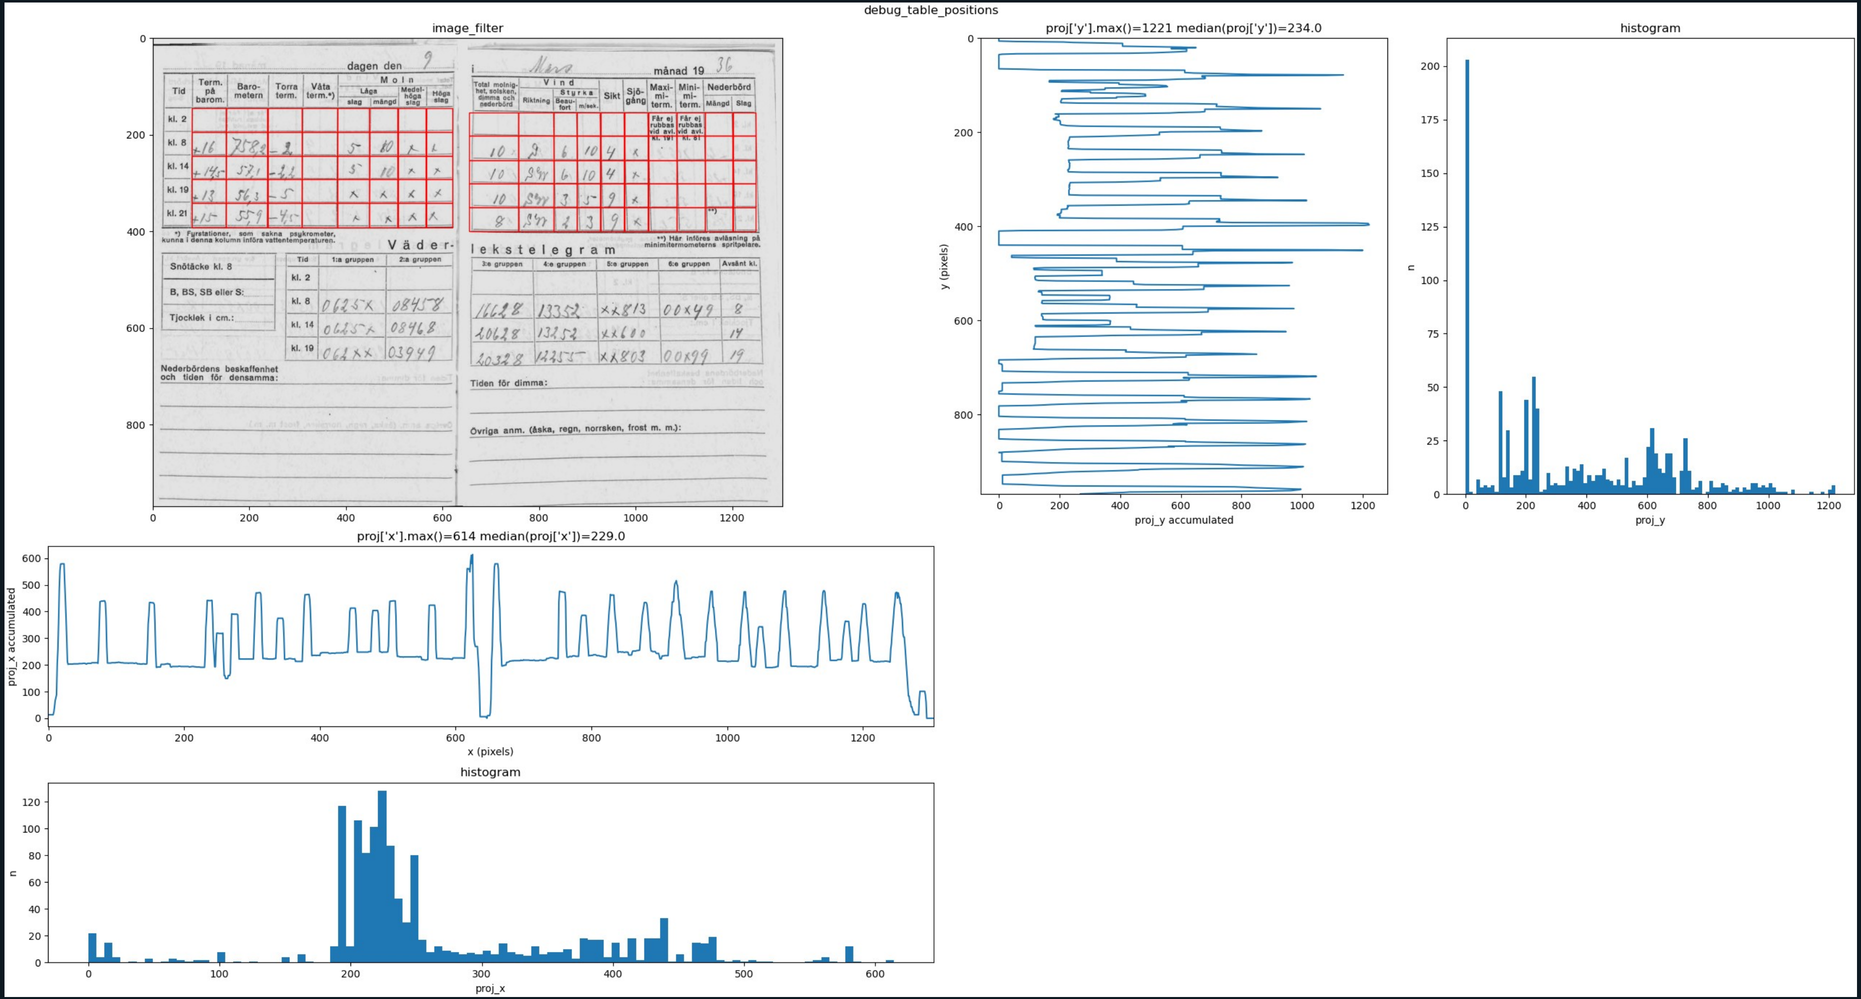
debug_table_positions
image_filter
proj['y'].max()=1221 median(proj['y'])=234.0
y (pixels)
proj_y accumulated
histogram
n
proj_y
proj['x'].max()=614 median(proj['x'])=229.0
proj_x accumulated
x (pixels)
histogram
n
proj_x

Flexible configuration

… ❯ dev ❯ ai-for-ob… ❯ dawsonia | bjuröklubb.toml | haparanda.toml | härnösand.toml |

```toml
[1926]
version = 1

[1927]
version = 1

[1928]
version = 1

[1929]
version = 1

[1938-0]
# Jan-Jun
version = 0

[1938-1]
# Jul-Dec
version = 2

[default]
version = 0

```

```toml
[version.0]
columns = [
  [
    "term_på_baro",
    "barom",
    "torra_term",
    "våta_term",
    "moln_slag_lägre",
    "moln_mängd_lägre",
    "moln_slag_medel",
    "moln_slag_högre"
  ],
  [
    "moln_het_sol_dimma_nederbörd_total",
    "vind_riktning",
    "vind_beaufort",
    "vind_m_sek",
    "sikt",
    "sjögang",
    "maximi_term",
    "minimi_term",
    "nederbörd_mängd",
    "nederbörd_slag"
  ]
]
name_idx = "tid"
rows = [2, 8, 14, 19, 21]
tables = [
  [5, 8],
  [5, 10],
  [3, 1],
  [4, 2],
  [4, 5]
]

```

dawsonia ml (HTR)

- Training set = 7000 images
- Validation set = 1000 images
- Test set = 1000 images
- Duration = 10 hours 41 mins with on an 8 core Intel Xeon machine. ~1 hour on a GPU

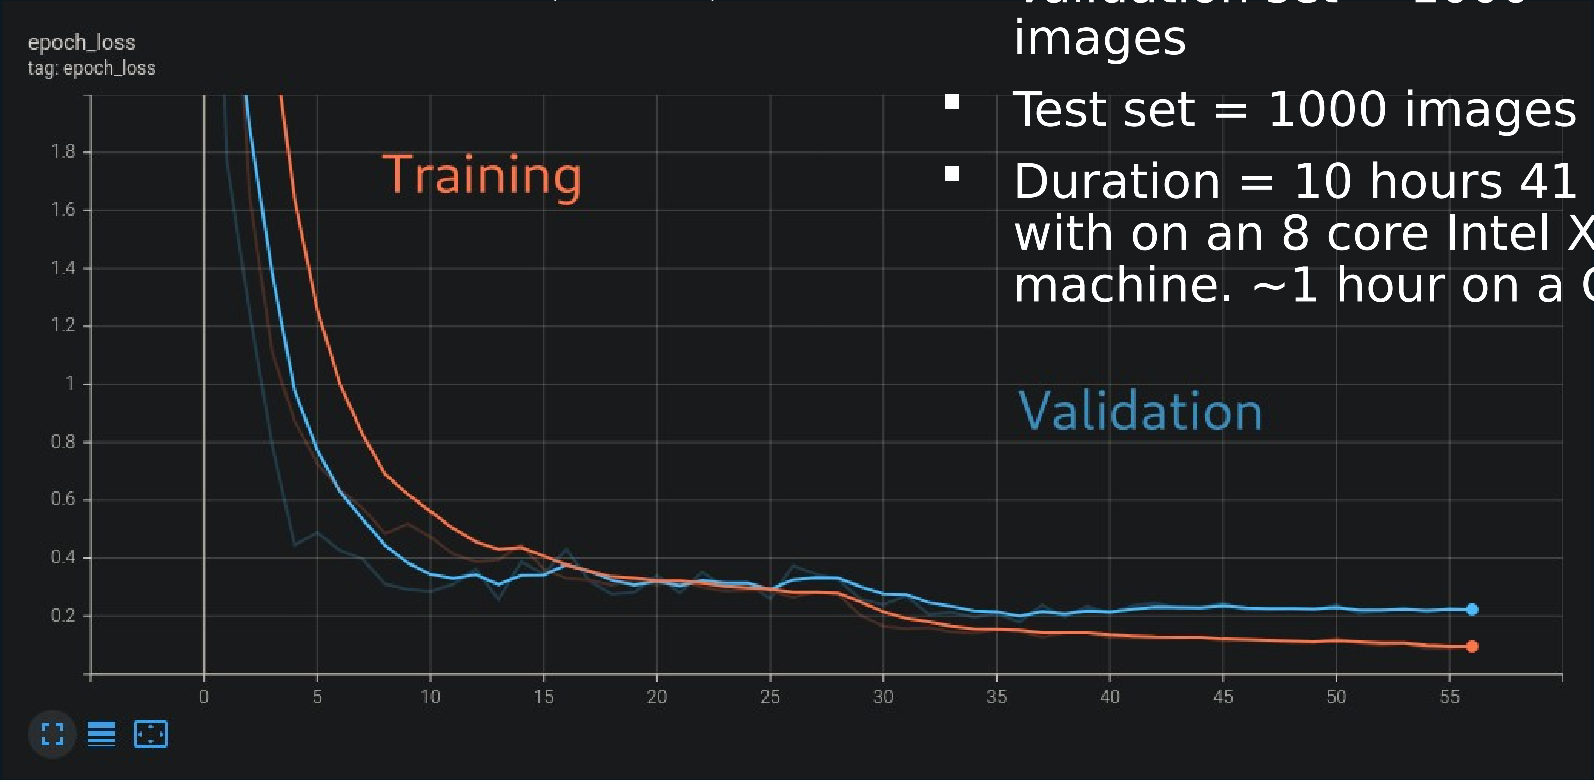

dawsonia digitize

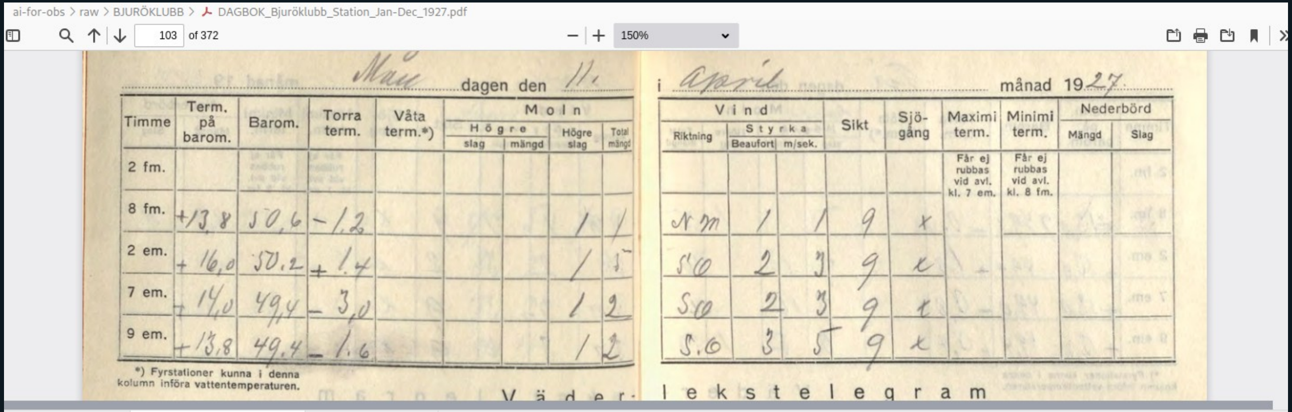

1927-04-11.md | Preview 1927-04-11.md | 1927-04-12.md

103: 1927-04-11

tid	term_på_baro	barom	torra_term	våta_term	moln_slag_lägre	moln_mängd_lägre	moln_slag_högre	moln_mängd_total	vind_riktning	vind_beaufort	vind_m_sek	sikt	sjögang
02:00:00													
08:00:00	+13.8	50.6	-1.2				1	1	+7	1	1	9	x
14:00:00	+16.0	50.2	+1.4				1	5	510	2	3	9	x
19:00:00	+14.0	49.4	-3.0				1	9	510	2	3	9	x
21:00:00	+13.8	49.4	-1.6				1	2	5.0	3	5	9	x

Technologies

- Python
- TensorFlow
- OpenCV
- Scikit-image
- Typer

Hosted in SMHI’s GitLab server as a fully open-source project

The HPC perspective

How I think HPC resources can be of help

Road ahead for us (where HPC is needed)

- Fix steps in the pipeline which **lead to problematic results** and **low model confidence**.
  - Improving **preprocessing** and **table detection** stage (if it cannot be solved with image-processing)
  - Expanding HTR **training dataset** using DIDA
  - Transfer-learning to **other** open-source **HTR** detection neural networks
  - **Hyper-tuning** of **current HTR** model

Developer experience

- **Running a simple JupyterLab server requires deep knowledge**
  Esp. multiple GPUs requires mastery of SLURM, tools to build Singularity, network communication layers

- **Now, much better at least for a single GPU job,** with the on-demand interface

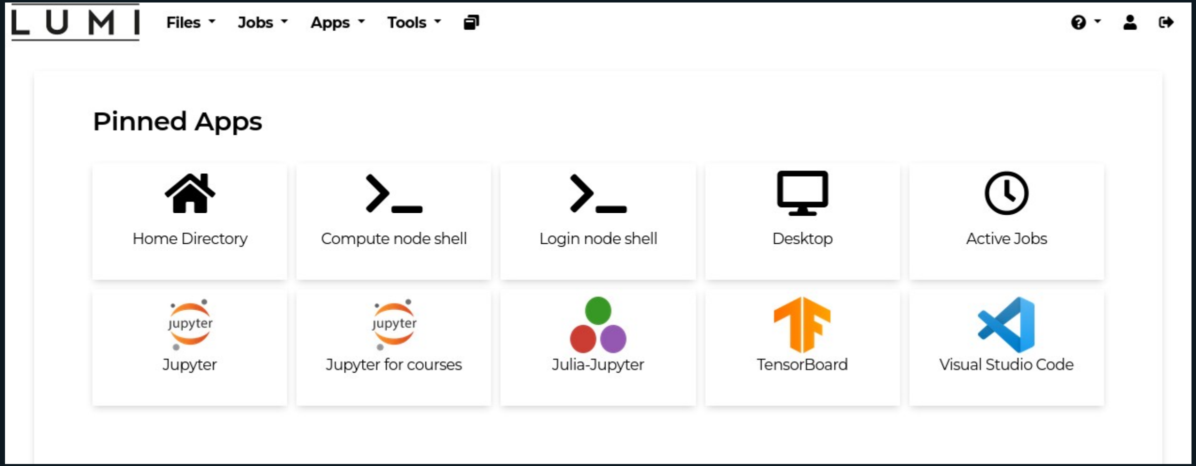

What about MLops?

- This approach works fine for one-off development with HPC.
  With small ML models, once trained, the inference pipeline can be executed on CPUs
- How would we operationalize (data, training, inference pipelines) for bigger models?
  Think temporally varying input data.

SMHI

@ashwinvis

ashwinvishnu

https://fluid.quest

@ashwinvis@mastodon.acc.sunet.se

Thank you for your attention!