

# Automating Insect Detection in Videos with Metacentrum


Jan Filip

Insect Communities Research Group  
Charles University, Prague

Principal Investigator: Robert Tropek

Design and Development: Petr Chlup

Deployment and Testing: Jan Filip, Dominik Anýž



cesnet  
**metacentrum**

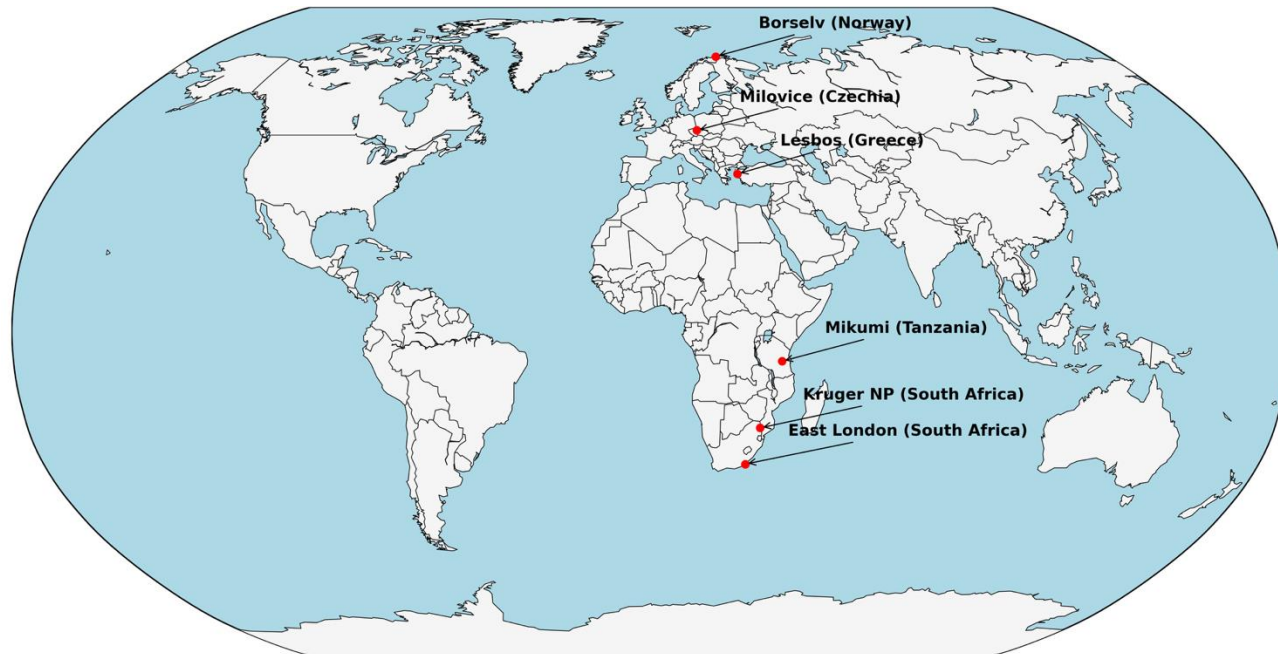


UNIVERZITA  
KARLOVA



# Scientific Motivation


- We study **pollination networks** (who pollinates whom, how often)
  - help quantify ecosystem interactions and specialization
- Current project: **latitudinal gradient** (Junior Star GAČR, 2021-2025)



Advertisement  
a) sensory biases  
b) reward motivation



Pollen Transfer  
a) anthers → pollinator  
b) pollinator → stigma



# The Data Challenge

- Method: **CCTV cameras on flowers**
  - recording full 24h cycles across many flowering plant species



- Advantages:
- Captures **diurnal** and **nocturnal** activity (infrared at night)
  - Ensures **statistical power** (huge recording times)
  - Deployable at **ecosystem scale** (30-50 flowering species per site)
- Drawbacks:
- 10 years** of accumulated recording time
  - ~3,500 full-day recordings = 350,000 video files = **~8 billion frames**
  - Video **quality varies** (day vs night, movement of flowers, small insects)

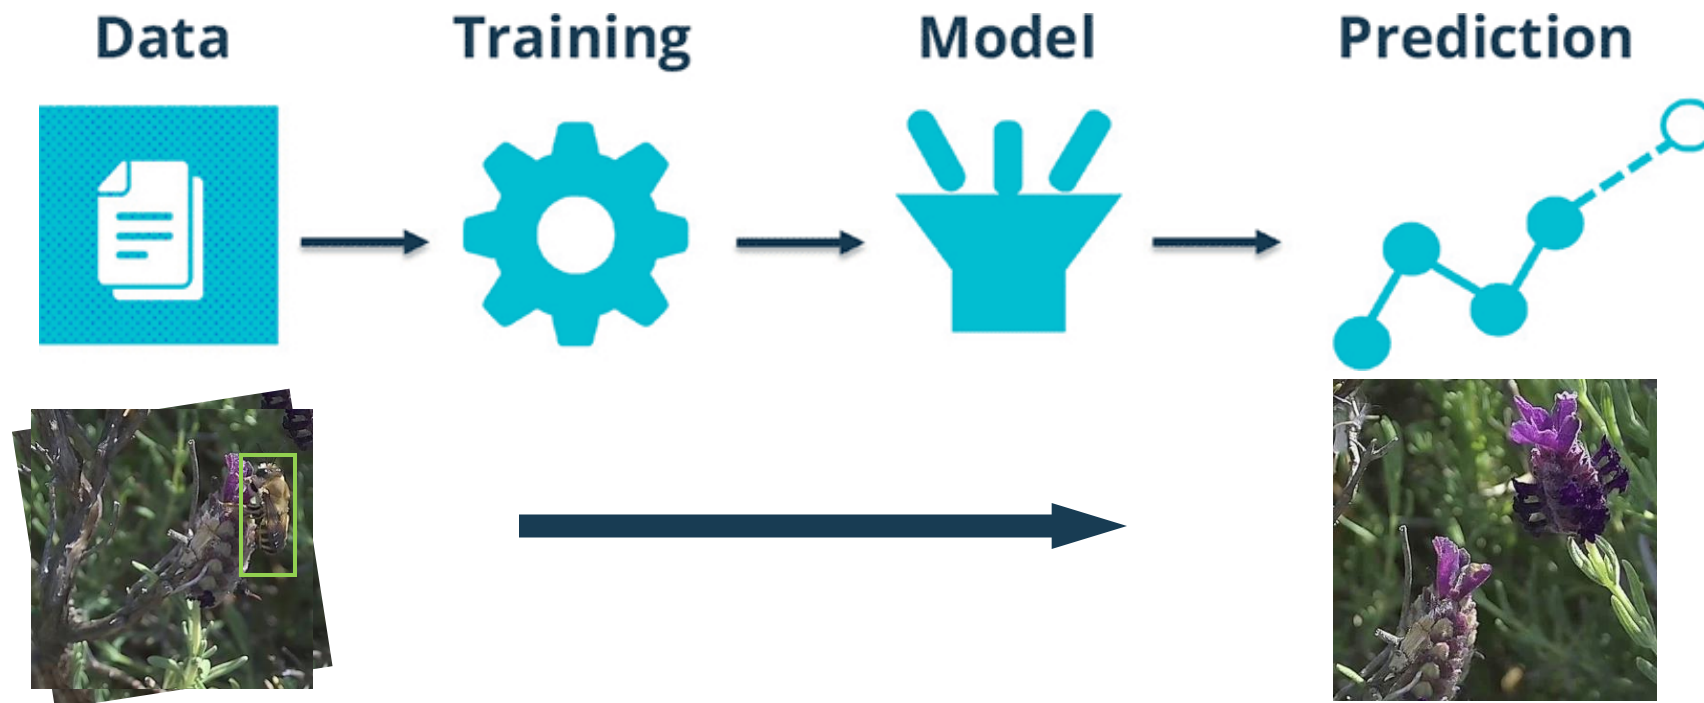



# Data Processing Goals

- **Raw videos** -> detect insect visits

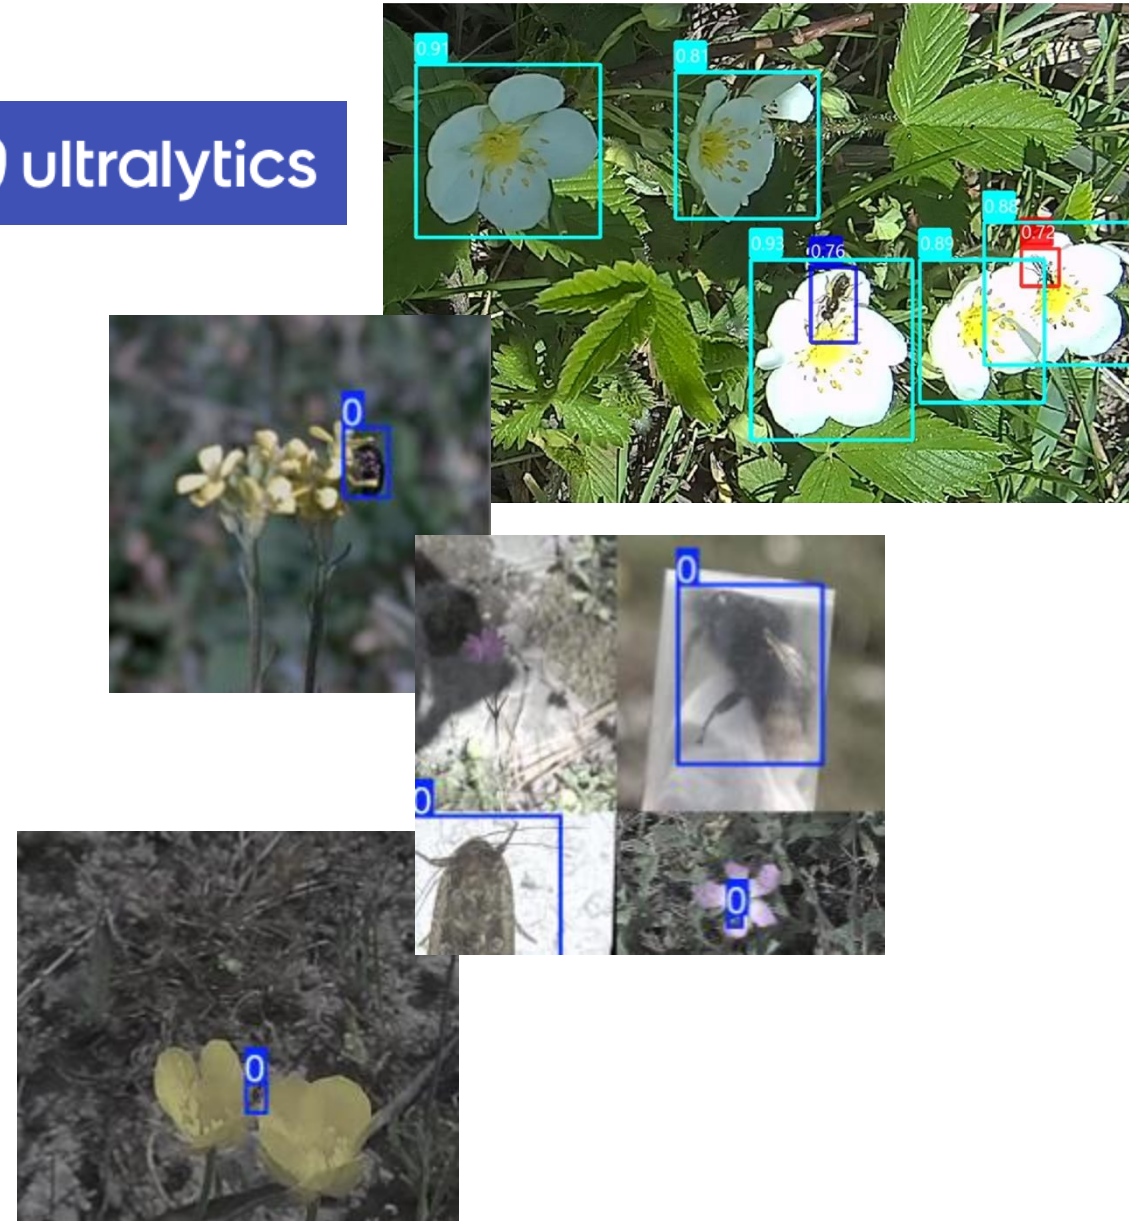
locate insects relative to flowers

extract **visit-level records** (arrival/departure)



# Model Training

- Chosen algorithm: **YOLO** (Ultralytics)
- **Training data:**
  - Manually annotated frames
  - Semi-automated labeling
  - Enriched with internet insect images
- Strategy: pretrain „**insect expert**“  
-> fine-tune on flower videos
- **Flower x insect** model






# Why MetaCentrum?

- Needed **HPC resources** for:
  - Large training runs (GPUs)
  - Huge inference workload (CPUs)
  - Long-term S3 storage of TBs of video
  - Jupyter Notebooks with processing resources
- MetaCentrum **enabled** scaling to billions of frames

Jupyter -> OnDemand -> PBS jobs -> GPUs, CPUs, Singularity

# Training on GPUs

- Used GPUs nodes with high **VRAM (~64 GB)**  
(**adan**, **galdor** – **gpu\_long**)
- Avoided AI-specific cluster (queues too long)
- Built and ran models in Singularity **containers**
- Solved CUDA/package version issues by pinning dependencies



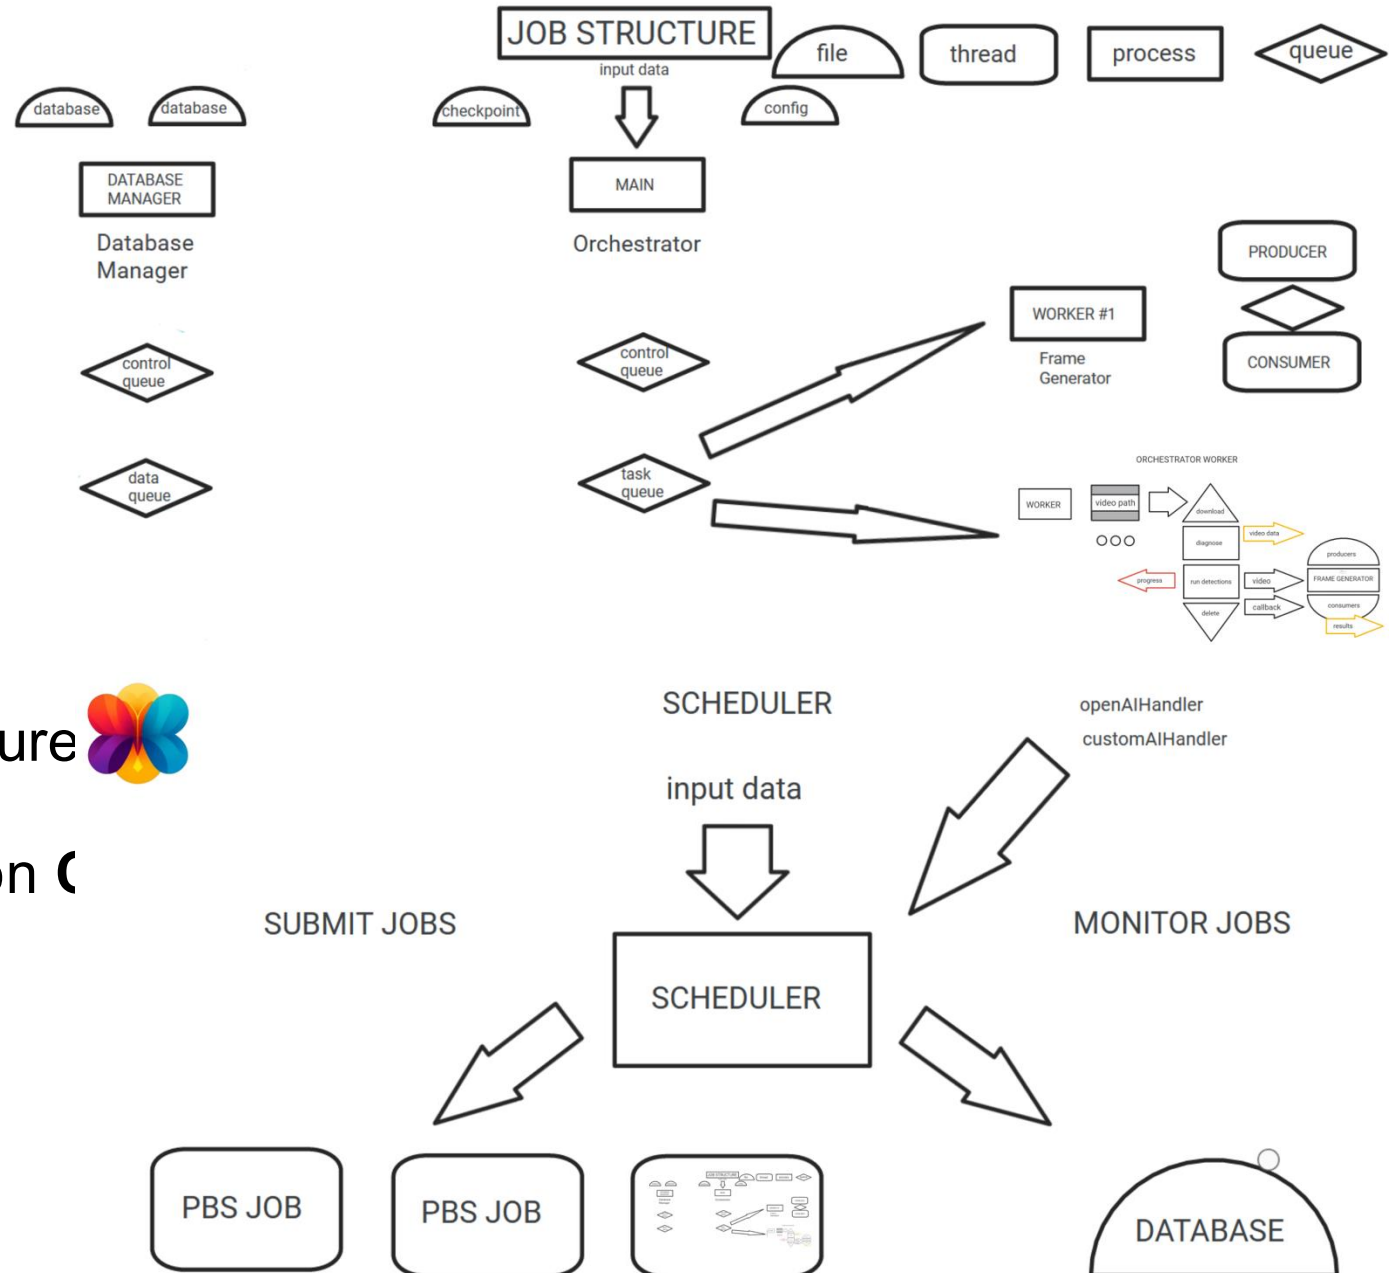
# Inference on CPUs

- **PBS** job design:

- Each job = batch of tens of vic
- Dozens of jobs in parallel -> h

- Jobs restartable to handle failure
- All large-scale inference ran on C only:

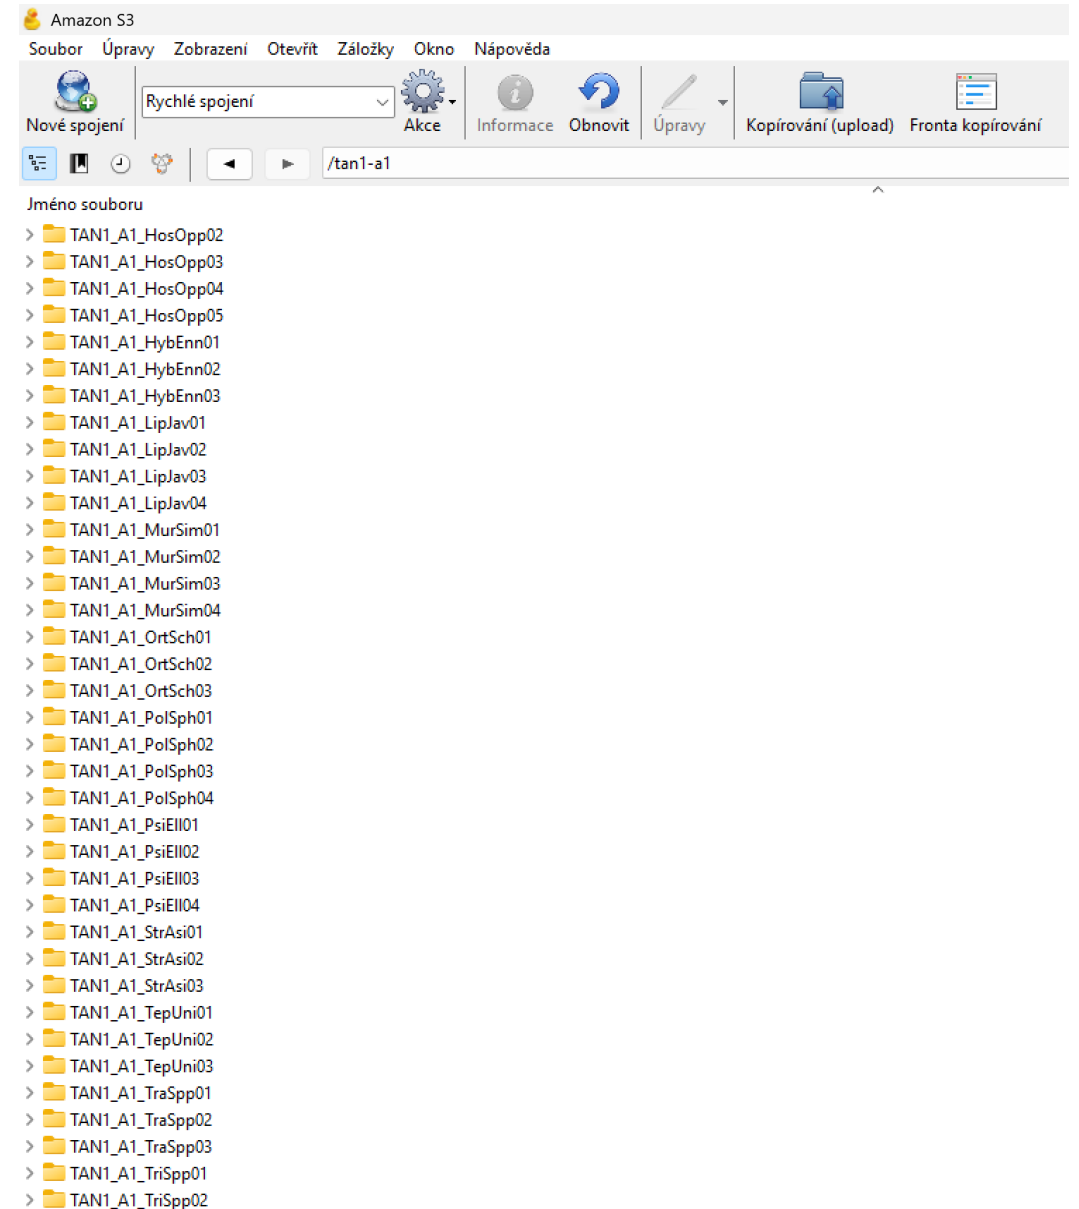
- GPU queues too slow
- CPU jobs easy to parallelize






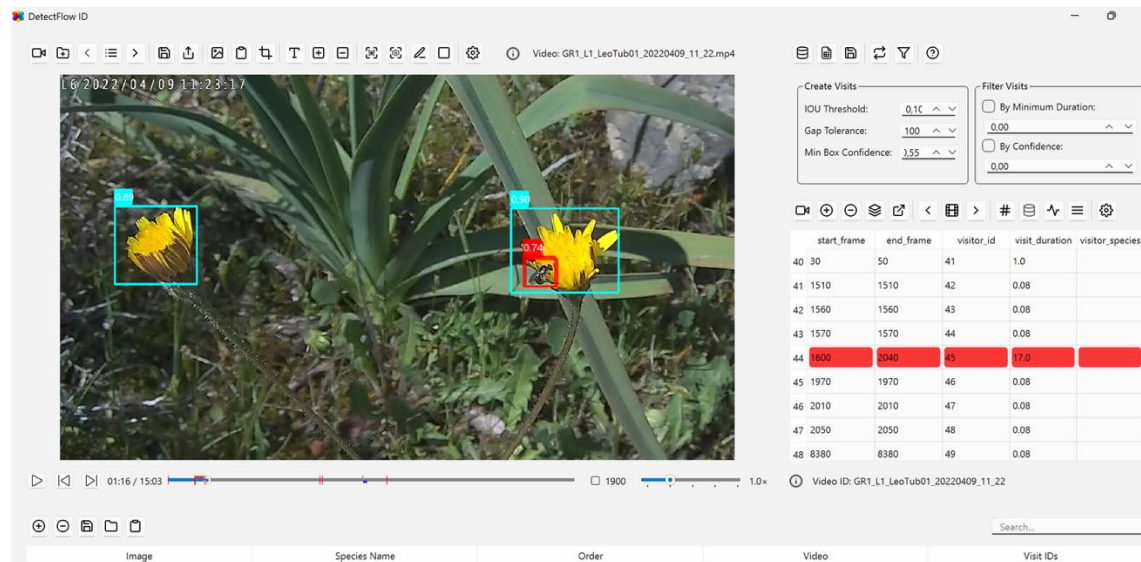
# Storage & Data Handling

- **S3 storage**: long-term archive, easy upload
- Local cluster storage only for temporary files
- Essential for handling terabytes efficiently
- Perfect protection through **credentials**
- **Shareable** in virtual organizations, work groups




# Outcomes

- HPC saved years of manual work -> project otherwise impossible
- Processing of **~8 billion frames** feasible
- Produced **visit-level dataset** from raw video
- Custom GUI app allows refining detections into final data



# Lessons Learned

- **Containerization** critical for reproducibility
- **S3 storage** indispensable for large video workflows
- **PBS job parallelization** allowed scaling
- **Standardised** sampling crucial for easier AI learning



# Acknowledgements

- Many thanks  <sup>cesnet</sup> for infrastructure and help!
- Thanks to the whole team and worldwide collaborators!

- Contact: [jan.filip55@gmail.com](mailto:jan.filip55@gmail.com)

Open to work 😊

