Intelligent Maintenance and Operation Systems

CH-1015 Lausanne

Prof. Dr. Olga Fink
GC A3 424 (Bâtiment GC)
Station 18

olga.fink@epfl.ch

Eidgenössisches Departement für Umwelt,
Verkehr, Energie und Kommunikation UVEK
Bundesamt für Verkehr BAV
Abteilung Politik
Sektion Direktionsgeschäfte
Mühlestrasse 6, 3063 Ittigen,

Lausanne, 24.11.2022

Progress Report «Project INITIATE»

1 Introduction

This report provides an overview of the progress made in project INITIATE. This project is a research project funded by the funding scheme for railway infrastructure "Bahninfrastrukturfonds" BIF of 2014. In this project, EPFL and SBB work together to improve the detection and prediction of critical conditions in railway power network systems by developing intelligent algorithms for control system data.

The project was initially started at ETH Zurich in 2021 and was transferred to EPFL in March 2022 due to the transfer of the research team from ETHZ to EPFL.

Figure 1: Power grids sites from which the generated data of case study I and II originates.

In June 2021, a group of domain experts from different fields collected and rated various case studies as candidates for investigation. The domain experts estimated the highest overall benefits of all case studies as well as immediate benefits if selected as the first case study. Based on these scores, two studies were selected to be the most promising and, thus, were processed first: a case study on state estimation and fault detection from power grid sensors in local areas such as power plants with connected substations (**case 1**) and a case study on hydro-power plant efficiency estimation from hydro and electrical sensors (**case 2**). Two workgroups were formed to study the different cases, led by Philipp Wenk from SBB's side and Olga Fink from EPFL's side:

- Project Lead: Philipp Wenk and (SBB) Olga Fink (EPFL)
  - Study case 1: state estimation and fault detection from power grid sensors
    - Robert Strietzel (SBB)
    - Raffael Theiler (EPFL)
  - Study case 2: hydro-power plant efficiency estimation
    - Roland Schäfer (SBB)
    - Francesco Fusaro (SBB)
    - Mengjie Zhao (EPFL)

2 Case Study 1: State Estimation and Fault Detection from Power Grid Sensors

In the 2021 report, we introduced case study 1 to detect, understand and analyze sensor and components faults of different severities that affect the power grid. In this period (2022), we found a promising way to do that by using an ensemble of forecasters with an algorithm that allows us to trace back the forecast to a set of relevant neighboring sensors that may be identified as faulty depending on an anomaly score.

Data: In this case study, we continue to focus on different datasets of *Etzelwerk* including connected substations, but we also use simulated datasets as a verification. We currently have three datasets available that cover signals of the electrical, hydraulic, and control systems as 1-minute averages from the years 2015 (all), 2018 (all), and 2021 (January to April). The dataset from 2021 has been used so far for most of the developments. The additional datasets of the years 2015 and 2018 were recently (Q4) exported as a measure against overfitting, but also to be robust to seasonality which we identified to be an influencing factor. We use the same EMS ("Energie-Management-System") data pipeline to export the two extended datasets. As already mentioned in the previous report, a key challenge remains that no verified, real fault data is available.

Table 1: Normalized test error (MSE) for current forecasts on the 2021 Etzelwerk and substations dataset.

Method	# Params	MSE
StemGNN	4.2M	**4.49e-3**
StemGNN-large	30M	4.90e-3
Spacetimeformer	9M	5.13e-3
Trivial Model	-	5.19e-3
Linear	1k	5.77e-3
A3T-GCN	523k	8.74e-3
LSTM	38k	9.24e-3
MTGAT	341k	13.29e-3

State Forecasting: In the past report, we emphasized the importance of reliable state forecasting for unsupervised anomaly detection. Such a forecasting algorithm is used as a component of anomaly-score-thresholding on the residual of the forecast with respect to the new measurements. We currently evaluate the hypothesis of such a system being able to extend the fixed threshold anomaly detection that is already implemented in EMS. To measure the forecasting algorithm performance and fine-tune the hyperparameters, knowledge of at least some confirmed faults is beneficial. Therefore, we took measures to detect sensor faults with human supervision. We found that parameter-free algorithms such as the matrix profile are a promising way to explore further.

Study of Computational Graphs for GNN Forecasting: We previously identified the power grid environment as promising to be processed with graph neural networks (GNN) due to the availability of system schematics that can be translated to a graph. In Q1/Q2, we completed the initial implementation of GNN on power grid data which was a milestone in the 2021 outlook. In this scenario, a GNN takes a structured dataset that consists of a set of nodes where each node is embodied in a graph that can be static or dynamic with respect to time. Each node has an associated feature vector that contains a window of the sensor data. The network applies graph-spectral and time-spectral filtering to the data to compute a next step ahead prediction (forecast). Initially, we translated the power grid's electrical schematic into a graph but we found that this standard GNN (we tested A3T-GCN) is outperformed by the trivial model where we use the values of the last step as the prediction for the next step.

We identified statistical properties of the power grid graph such as a low clustering-coefficient, low node degree and high diameter that explain why the power grid state cannot be forecasted with standard GNN. We explored established methods in forecasting such as trend-decomposing linear models, ANN, (bi-)LSTM because of this initial setback. Our findings are summarized in Table 1.

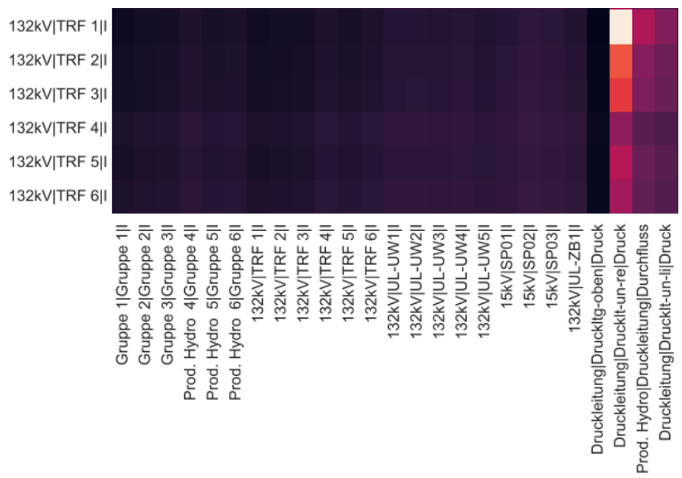

Figure 2: A segment of the learned attention of StemGNN. The colors indicate link strength. We observe that the network learns to link the hydraulic system to the electric. In StemGNN, the attention is directly interpreted as a weighted undirected graph in between signals for GCN message passing.

The best-performing model, measured by its forecasting performance, is currently StemGNN, a GNN that does not rely on an input graph, but learns internally to link signals as a graph for data patterns that correlate with a particular prediction task. Graph connectivity is established by leveraging attention, a mathematical concept derived from the neuroscientific concept of cognitive attention. Based on our finding that attention-based GNN work well, we continued exploring the similar class of attention-based transformers (Spacetimeformer). However, Spacetimeformer have a slightly worse performance compared to StemGNN and are harder to train.

Understanding and extending StemGNN: A closer look at the learned graph in StemGNN revealed a graph of signals that mirrors known physical relationships between the signals in a power plant. To give an example, the GNN learned to connect flow rate and hydraulic pressure sensors to the current forecast in the electrical system. The corresponding attention graph for this relationship is visualized in Figure 2. In Q3/Q4, we started to adapt StemGNN to power grids: With our focus on learned graph structures, we introduced different trainable attention mechanisms to generate different learned graphs. We wrote analytical code to visualize and compare the forecasting performance and stability of these graphs for different training environments. Figure 3 shows that StemGNN's attention is robust with respect to different conditions.

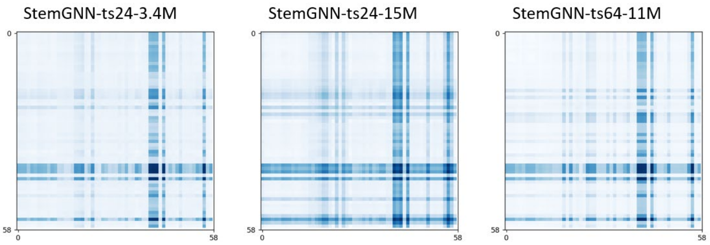

Figure 3: The learned attention of StemGNN is robust with respect to different random seeds, different model and window sizes across training runs. This leverages the path to understand how information is passed across sensors during forecasting.

To validate the robustness of attention to sensor selection and window size, we increased the parameter efficiency of StemGNN to allow training with a longer input window and a larger number of sensors. Additionally, we added backcasting as an alternative to forecasting-based anomaly detection to stabilize model training. Contrary to forecasting, in backcasting, the GNN maps the input to a temporally aligned output. The currently trained backcasting output is a reconstruction of the input. A promising untested option is to train to map to a pre-computed state estimation from a power system simulator. This could force the model to be coherent with the physics implied by the state estimation solution.

Explainability: Most forecasting algorithms do not allow us to point to individual faulty sensors and can rather be used for system-level anomaly detection. In StemGNN, the learned attention can be leveraged to address this "explainability" problem. At its core, StemGNN is a graph convolutional network (GCN) that is mathematically mappable to message passing. Therefore, we can examine the messages between sensors to understand how information is exchanged during forecasting, especially in case of a sensor fault. We empirically found that the best-performing forecasting models tend to have sparse attention patterns (Figure 3) which simplifies the explanation. Additionally, we observe that the individual sensor's forecasting performance positively correlates with the total strength of incoming attention. This can be exploited to identify neighboring sensors as a cause.

Since we have no ground truth on faults to train a classifier, we are currently testing a residual-based approach in a model ensemble that predicts the same signals from a sparse, but weakly overlapping set of sensors (Figure 4). In

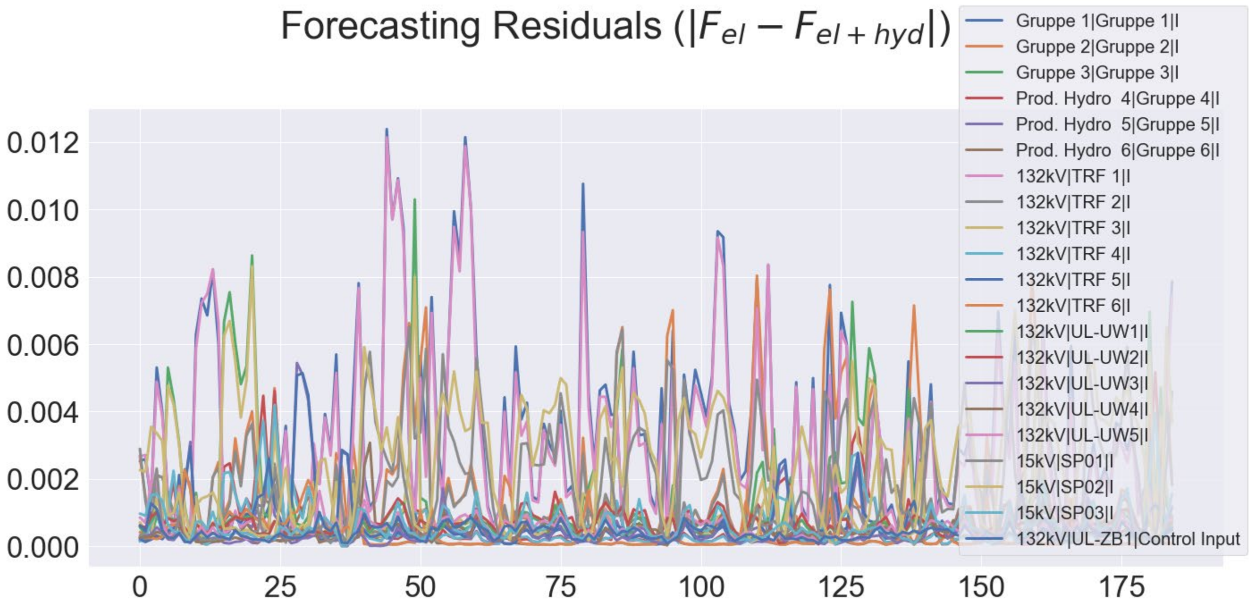

Figure 4, A StemGNN model residual approach: Here we show the (normalized) residuals of forecasted currents for the 2021 Etzelwerk test dataset. $\boldsymbol{F_{el}}$ forecasts based on electrical signals. $\boldsymbol{F_{el+hyd}}$ has access to hydraulic signals and hence strongly attends to pressure and flow rate signals (figure 3). Despite their differences, the two models should agree as they are trained on the same objective. However, we observed variability in the residuals which may be due to faulty sensors which is the subject of current investigation.

that case, a disagreement may be due to a faulty sensor on one end. In order to establish this statement, however, we must first understand and eliminate other influences on the graph and model level which is the focus of current research. The preliminary results show that the attention can be designed to be interpretable and robust to different input sizes and to model complexity. We are currently exploring the relationship between attention strength and forecasting performance which we have found to influence the forecasting residual.

3 Case Study 2: Hydropower Plant Efficiency Estimation

The goal of the case study is to understand and monitor the efficiency of hydropower plants to ensure a reliable energy supply and enable smart maintenance scheduling. The pilot case is demonstrated for the *Vernayaz* hydro power plant (240 GWh annual production with three dual-jet Pelton turbines) owned and operated by SBB. By the end of 2021, we defined the following research questions to achieve better efficiency estimation:

- Which factors affect hydropower efficiency and are feasible to be considered in this case study?
- How precise are the sensor measurements and what are their influences on efficiency estimation?
- How to ensure a reproducible efficiency estimation given the same operational conditions?
- How do generating units influence each other in terms of efficiency when multiple units are operating?

This report starts by answering the above questions and elaborates subsequent research questions. A summary of the case study progress timeline can be found in Table 2.

Signal selection: Hydro efficiency is defined as the ratio of the total electrical power produced to the theoretical hydraulic power of the water flowing through the turbine. The latter is determined by the hydraulic head (proportional to the net water fall height, determined from the turbine-level pressure) and the volumetric flowrate. Given this definition, we decided to start with the signals that appear directly in the definition, namely generated power (apparent power and reactive power), volumetric flow rates, pressure, water level, rotation speed, all extracted at 1 second sampling intervals. In the statistical analysis performed in 2021-Q4, we found that the number of active nozzles strongly affects hydro efficiency. We have further concluded that it is not only the number of active nozzle that matters, but also the position of the nozzle needle, which determines the jet area. The two aforementioned pieces of nozzle information are essential for the efficiency estimation, however they are not stored in the EMS system and must be extracted from the local storage unit. SBB is working with Axpo to integrate local operating data, including nozzle opening data, into the cloud system. The first data extraction might become available early next year. Whether the duration of the recorded signals is sufficient for the efficiency estimation needs to be verified.

Sensor measurement quality estimation: Statistical analysis performed in 2021-Q4 showed that sensor measurement quality plays a key role in the efficiency assessment. In addition, it was shown that there is a constant delay between the electrical and the hydraulic sensors. For steady state efficiency estimation, it was shown that it is essential to calibrate the signal delays. While this delay can be calibrated by a peak-to-peak alignment between volumetric flowrate and generated power, it is not possible to calibrate the pressure measurement in the same way, as it exhibits a constant oscillation behavior. This oscillation is present throughout the year with a period of 1-2min and a magnitude of 1-2 bar. In order to get to the bottom of the cause of this phenomenon, we visited the hydroelectric power station under the guidance Bernhard Roggli, SBB expert in sensor technology. On site, we observed a rapid on and off switching of the nozzles as well as a rapid change of the nozzle areas to meet rapid changes in demand. This induces a large local reactive pressure pulse with a fluctuation of around 1-2 bar, comparable to that observed in the surge tank water level fluctuation. Thus, it is not possible to define a steady state for the pressure head. This finding motivates a data-driven hydro efficiency estimation method that can self-calibrate the signal delays given a sequence of signals.

Data-driven efficiency estimation: statistical analysis requires a deep understanding of the underlying physical model. When the system becomes more complex, analytical estimation of efficiency is almost impossible due to noise and sensor inaccuracies. Data-driven methods are convenient in this case because they do not require knowledge of the exact physical system. Instead of extracting the "explicit" efficiency as expressed physically in a formula, we decided to learn the "implicit" efficiency by training the model with historical data. Here we define the "implicit" efficiency as the generated power given the sequence of volumetric flowrate and pressure signals. If the model is sufficiently accurate, the machine efficiency is implicitly stored as model weights. The difference between the predicted power and the actual power produced shows the degradation of the machine. To eliminate the influence of mixed operating conditions, only single-machine operating condition was considered. A 1D convolutional neural network (1DCNN) was applied to predict the power generated at the next time step based on a sequence of hydraulic and electrical data. In addition, the architecture was adapted to infer the nozzle area. Nevertheless, due to the small amount of single-machine operational state data (160 hours), it was difficult to

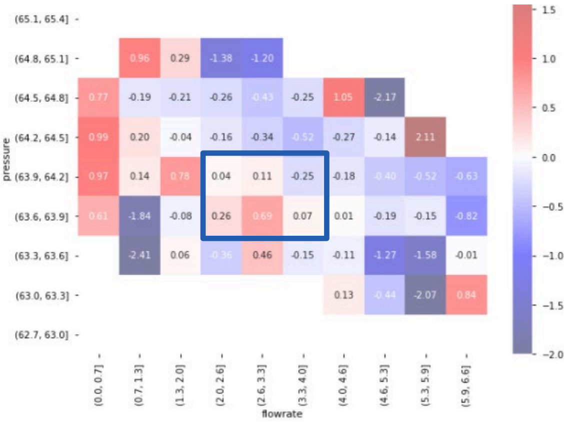

Figure 5: Difference in 1DCNN predicted power (in MW) for 2021 trained with data from 2020 vs actual generated power for MG2 in 2021. Only single-machine operational state is considered. Model prediction error for the 2020 test set (>70k data points) is 0.09MW. Blue box highlights common operating conditions with more than 5k data points for both years.

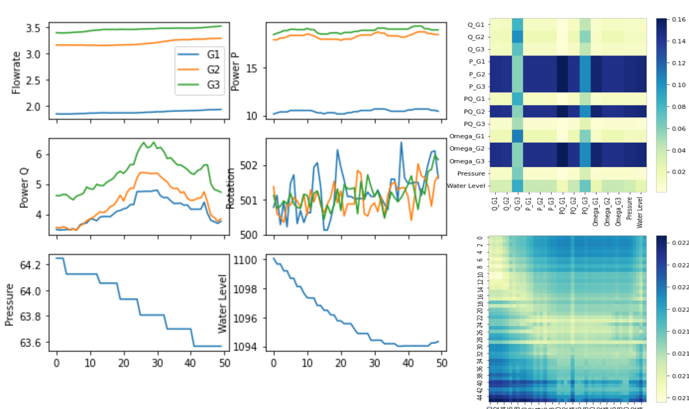

Figure 6: *Left*: input signal (flowrates, apparent and reactive power, rotational speed of three machine groups, pressure and water level). *Right*: feature attention and temporal attention obtained by the GNN. Feature attention shows how different signals interact with each other; temporal attention presents the importance of signals at different time steps

determine the nozzle area in an unsupervised manner. This could be improved as soon as the locally stored nozzle data is available. For MG2 with max. capacity of 37MW, the model was able to predict power generation 1s ahead with an absolute mean error of 0.09MW (Evaluated on the test set of December 2020 data. Due to the sparsity of singe-operational states, the training data is concentrated in May, October and November 2020). A heatmap of the differences in 1DCNN predicted power (in MW) for 2021 (data concentrated in August, September and December), trained with 2020 data, versus actual generated power for MG2 in 2021 is shown in Figure 5. Here, we face a challenge that there is a shift in operational range from high head, high flow rate in 2020 to low head, low flowrate in 2021. There is only a small region where there are enough data points for both years to evaluate efficiency. Within this region, the difference between the predicted and actual power is comparable to the model prediction error. It is, therefore, hard to justify whether this difference is due to model error or machine degradation. To verify this, further investigation could be done with data from more distant years.

Machine group interactions: Insights in 2021-Q4 study showed that the efficiency of an individual operating unit can be strongly influenced by the interaction with other operating units. Further studies in 2022-Q1 confirm that, under the same hydraulic conditions, less power is generated in multi-machine operation than in single-machine operation. This motivates studying hydro efficiency for all machine groups as a whole to better understand the interactions between different operating units. Graph Neural Networks (GNN) model the relationship between signals as edges and therefore have advantages in learning interactions between signals compared to other machine learning methods. In 2022-Q3, we developed a GNN method with Graph Attention Networks (GAT) as backbone at a feature level (for different signal sequences) and at a temporal level (for all signals at different time points). The attention layer enables an interpretation of interactions between different signals as well as the understanding of hidden signal delays. An example of the feature and the temporal attentions given 14 input signals of 50s can be found in Figure 6. The current GNN network treats feature and temporal attentions independently. This encounters limitations when there are inter temporal-feature interactions. In 2022-Q4, we have been working on developing methods that can take this factor into account on a benchmark dataset. In 2023-Q1/Q2, we plan to adapt the new methods to the Vernayaz data.

Table 2. Efficiency estimation progress timeline

	Task	Comment	Status
2021 – Q4	Data analysis, sensor delay analysis, steady state efficiency estimation	Steady state hard to determine, varying sensor quality	Done
2022 – Q1	Indirect (data-driven) efficiency estimation by training the 1DCNN model to do power prediction given operating conditions	Change of efficiency between two years is too small to evaluate, nozzle data not yet available	On pause
2022 – Q2	Understanding nozzle operating scheme and the oscillating pressure measurements by visiting the hydropower plant in Vernayaz	Pressure oscillation is driven by rapidly changing demand	Done
2022 – Q3	Learning machine group interactions with Graph Neural Networks		Ongoing
2022 – Q4	Further method development on Graph Neural Networks on a benchmark dataset	Better understanding of feature & temporal attentions	Ongoing
2023 – Q1/Q2	Adapt the new GNN method to the hydropower dataset		Not yet started

Summary: We summarize the progress timeline for Case Study II in Table 2. In 2021-Q4, we began statistical data analysis, performed steady-state efficiency estimation, and concluded that it is difficult to statistically estimate efficiency due to sensor inaccuracy and complex dynamics in operation. This was further confirmed in 2022-Q2 during the on-site visit to the Vernayez hydropower plant: Pressure fluctuates periodically due to rapid operational changes caused by fluctuating demand. 2022-Q1, we developed a data-driven method to explicitly study efficiency by predicting next time step power generation using a sequence of historical hydro and electrical signals. We tried to integrate physics into the model, but without information about the nozzle opening, this attempt was not very successful. In 2022-Q3, we developed a GNN-based method to study machine group interactions. The method was further developed in 2022-Q4 using a public benchmark dataset. We plan to finalize case study II in 2023-Q1/Q2 by adapting and applying the developed GNN method to the hydropower dataset.

4 Outlook

Case study 1

Our research suggests that model residuals could be used for unsupervised electrical sensor fault detection. Our strategy relies on finding non-overlapping sets of signals that contain sufficient information for a target forecast. In hydropower plants, the electrical signals correlate with the signals from the hydraulic domain. Hence, we started to integrate signals from the hydraulic domain into our models. We found that sets of hydraulic signals also constitute a forecast source for our approach. Since we have shown that cross-domain forecasting from hydrological to electrical data is promising, we want to further explore cross-domain modeling between electrical, mechanical, hydrological, and control signals. These domains are currently not integrated into traditional power grid state estimation and have great potential to improve current monitoring, fault detection, and forecasting solutions. With the current GNN model performance, both, online and offline data validation is possible. Finally, we have gained a better understanding of how tightly GNNs can be coupled to traditional schematic-based grid simulators and plan to use this knowledge to co-integrate synthetic data in the next period. We identify a list of milestones:

- Even with strong regularization, we observed the tendency of models to overfit. We plan to train more extensively on the new datasets containing all seasons and we plan to measure the generalization capacity from the year 2015 to 2018 to understand the dependency of the residuals on the model.
- We plan to conduct a systematic study on the properties and the stability of the learned GNN attention. A promising direction is to use rate-distortion theory in this context.
- We aim to increase the size of the model ensemble by gradually including forecasters from co-domains as well as from simulation. Here, we try to find a working trade-off in between forecasting performance (more sensors yield better MSE performance) and sensor fault detection capability (smaller models increase the number of receptive sensor fields to distinguish in case of a sensor fault).
- Labeled faults for power grid data are scarce. However, to validate model performance, we need confirmed faults. We plan to combine the anomaly scoring from our methods with the matrix profile algorithm to highlight the most anomalous segments of our datasets for expert validation.
- EMS signals are recorded every second while the EMS state estimation reference is updated every 20 seconds. This leads to false errors because signals are often very dynamic and over- or undershoot the reference in between the state estimation interval. We are currently evaluating if a model could estimate the likelihood of EMS warnings being correct or false in this context.

Case study 2

The case study 2 focuses on two aspects of efficiency estimation: a) data-driven, forecasting-based efficiency estimation for the single machine operation state and b) understanding how machine interactions can affect the efficiency of an individual machine. We plan to finalize case study 2 by 2023-Q2 for the two aspects by:

- Integrating nozzle information (the number of active nozzles as well as the nozzle needle information) to build a physics-enhanced 1DCNN model to improve single machine operation state forecasting accuracy. Hydro exploitation, the operator of the hydro powerplant, is currently working to extract this information, and the data is expected to be available early next year.
- Adapting and transferring the developed GNN from the benchmark dataset to the hydro dataset. The properties and the numbers of features of the benchmark dataset are different from those of the hydro dataset. Small modifications need to be made for the transfer of the task.