Computationally Efficient Optimization of Long Term Energy Storage Using Machine Learning ⋆

Simon Karlsson, Stavros Vouros, Kristian Sandström, Konstantinos Kyprianidis

Mälardalen University, Västerås, CO 72123 Sweden (e-mail: simon.karlsson@mdu.se

Abstract:
Energy storage can be charged when energy is cheap and discharged when it is expensive to make an energy system more profitable or used to make the plant operation more efficient to reduce CO_2 emissions. To optimize long term energy storage with conventional methods a long time horizon must be used. When the long term energy storage is combined with a complex energy system the computational cost becomes large when using conventional methods. To reduce the time horizon, an algorithm will be used to decide the state of charge of the long term energy storage at the end of the day. This algorithm is trained using machine learning with data of the optimal state of charge obtained by running computationally heavy long time mixed integer linear programming ahead of time. Then a one-day or week mixed integer linear programming optimization will be done for the production planning. The seasonal patterns of the long term energy storage can then be captured while giving the plant operator a simple one-day or week production plan. A case study will be done with a combined heat and power plant system with 4 boilers, a long-term thermal storage, and a hydrogen storage system. Using this method the complexities of a multi energy system with long term energy storage can be captured while doing day ahead production planning.

Keywords: Energy, Optimization, Energy Storage, Machine Learning, Unit Commitment, Production Planning

1. INTRODUCTION

Energy storage is an important technology in the transition to more sustainable energy system since the energy generated from variable renewable energy sources will not match up with demand. This leads to energy having to be stored to meet demand without oversizing the energy generation and curtailing energy. Some types of renewable energy generation such as solar or wind also have seasonal patterns which can require long term energy storage (LTES) for efficient operation of the energy system International Energy Agency (2024). Because of this, the optimization of LTES is important to help the efficient transition towards a more sustainable energy system. For example, Brey et al. investigate how hydrogen could be used as seasonal energy storage in Spain and conclude that it could be used to smooth out seasonal imbalances Brey (2021).

There are different kinds of electricity markets, in some of these markets like Nordpool in northern Europe. Trading is done with both electricity users and suppliers placing bids and then a price is decided depending on where these bids meet Nordpool (2024). In this system, the bidding period is 1 hour and because of this, there are requirements on the computational speed of the optimization process for electricity suppliers. To optimize LTES with conventional methods like mixed integer linear programming (MILP) a long time horizon must be used which can make the optimization computationally expensive. This time can be too long to make bids on the electricity market especially if the optimization has to be run several times to run different uncertainty scenarios.

Saletti et al. use linear programming (LP) for the long time horizon (LTH) while MILP is used for the short time horizon (STH) Saletti et al. (2022). This method has a fast solution time, however, the solving time of the (LTH still depends on system complexity. The objective of the optimization is to meet the heat and electricity demand of a hospital and not maximize profit by selling to the electricity market. Marzi et al. use MILP to do day ahead scheduling of a multi energy system with LTES considering uncertainty Marzi et al. (2023). The computation time for their method is, however, too long to do bidding in less than one hour.

In a study by Bischi et al. a rolling horizon is used together with typical weeks to optimize plant operation with MILP considering the entire year. The goal of this optimization is however not to consider how the state of charge (SOC)

⋆ This work is supported by project ProPlan under the RESILIENT competence center, financed by the Swedish Energy Agency (2021-90273) and co-financed by Mälardalen University and industrial partners. The support received by Mälarenergi AB and Eskilstuna Strängnäs Energi och Miljö AB is appreciated.

of the storage will change over the year but to make sure that yearly emission constraints are met optimally.

Bruninx et al. optimize an energy system with energy storage using unit commitment by considering reserve capacity in a computationally efficient way Bruninx and Delarue (2017). The time horizon in this study is however 24 hours so the focus is not on LTES. Optimization of a compressed air energy storage is done by Ghaljehei et al. by using using stochastic programming and mixed integer nonlinear programming. Here the time horizon is also 24 hours so it is not fit for LTES Ghaljehei and Golkar (2017).

System states are used to optimize medium and long term energy storage in a study by Worgin et al. Wogrin et al. (2016). Here some states of the system are defined and clustered and based on what cluster the system is in the storage is operated accordingly.

A life cycle analysis with energy storage optimization is done by Dong et al. where the energy storage is optimized using a representative day for each season Dong et al. (2023). This representative day is used to calculate how much the storages will be charged or discharged during each season to store energy over seasons. This method will however not work when doing operational optimization since it will just have the same operation every day and not ex discharge the storages more for a day with high electricity prices

Mi et al. use multi timescale optimization to do generation and expansion planning where the longest timescale is one month Mi et al. (2021). Here the longer timescale is however used to optimize capacity credits and not to optimize LTES. Zhang et al. also use multi timescale optimization to optimize the operation of an energy system with hydrogen energy storage Zhang et al. (2023). Here a rolling horizon optimization is used where different kinds of energy have different time resolutions. Here two days ahead is used to optimize the energy storage using MILP. Su et al. use multiple timescales and add a flexibility requirement to make the energy system more prepared for uncertain future disruptions Su et al. (2023). Here a short, medium, and long time horizon is used where the long time horizon is one week.

In a study by Bahlawan et al. the design and operational operation of an energy system with long term thermal energy storage is optimized Bahlawan et al. (2022). Here switch on priority is used to do the operational optimization where one energy conversion technology is used first and only if this technology can not supply the demand the next conversion technology is used. This method will however not work well for a system where the operational cost of different technologies changes and there is no electricity demand but instead electricity is sold to the grid.

Reinforcement learning is used by Alabi et al. to control an energy system with energy storage and carbon capture Alabi et al. (2023). Here reinforcement learning is used to control the power output of the energy units and not to optimize any kind of LTES. Sleptchenko et al. use LP to optimize multiple different energy storage technologies as a part of an energy system Sleptchenko and Sgouridis (2019). In this study, the focus is not on computational speed but on the seasonal patterns of the storage operation.

Water value is an optimization method to optimize how hydropower resovoirs are used. Here a value of the water in the reservoir is calculated and used to determine if the reservoir should be discharged Helseth et al. (2017); Jahns et al. (2020). This method is quite computationally costly if used for daily production planning with the optimization by Helseth et al. taking between 28 and 40 hours.

1.1 Current Work

In this paper day ahead planning will be done for an energy system with LTES where the goal is to make as much profit as possible by selling electricity to the grid while supplying the required district heating (DH) demand. Instead of using LP or MILP to optimize the long term behavior of the system a machine learning (ML) model will be used to predict the end of day SOC of storages and then MILP will be used to optimize the daily operation with this SOC as a constraint.

The contribution of this work will be (i) to develop a new faster method to optimize LTES which allows for scenario analysis in production planning or be used in studies where the optimization has to be used many times. (ii) Analyze the effect of system complexity and optimization horizon. (iii) Test which input features give the best prediction.

2. METHODOLOGY

Because MILP is slow over long time horizons, a ML algorithm is used to predict the SOC of the storage's at the end of the day or week so the MILP can run for one day or week instead of a longer time. The ML algorithm is trained using optimal SOC data obtained by running the MILP on historical electricity price and DH demand data. Because only a few years of DH demand data was available some synthetic electricity price and DH demand data were also generated for training. This was done by using the probability density function (PDF) which can be seen in Eq. (1) to decide how much the scenarios should deviate from the real data like in Marzi et al. (2023) but with some changes. These changes are, instead of using the PDF to decide the deviation from the real data the PDF is used to decide the change in deviation at each timestep. Some of the spikes in the electricity price were also randomly removed and new ones were added so the spikes in the electricity price would not occur at the same time of year in all the generated scenarios. This synthetic data was run through the MILP to get optimal SOC data for training the ML algorithm. A flowchart of how the training and optimization are done can be seen in **Fig. 1**.

$$PDF(x) = \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{(x-\mu)^2}{2\sigma^2}} \tag{1}$$

A case study based on the system seen in **Fig.** 2 where the full system has four combined heat and power (CHP) plants, one TES (thermal energy storage) which uses water to store heat. There is also a hydrogen energy storage (HES) with an electrolyzer to convert electricity to hydrogen, a hydrogen storage tank, and then a fuel cell (FC) to convert the hydrogen back to electricity For both the electrolyzer and FC there are some losses in the form of heat which is

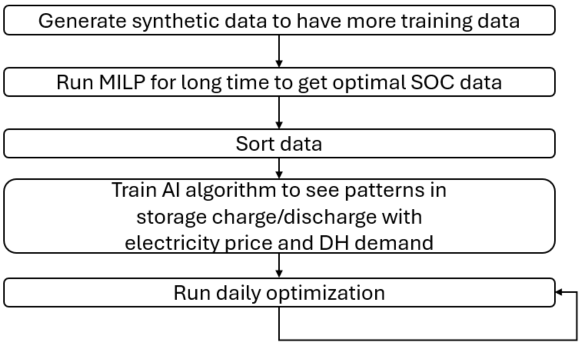

Fig. 1. Flowchart of method

used to both charge the TES and provide heat to the DH network. The system is used to provide the DH demand to the district heating network and sell electricity to the electricity grid. To evaluate how system complexity affects the current methods performance some other cases were also evaluated, these are one case with one boiler, the HES, and the TES, one case with one boiler and the TES, and one case with one boiler and the HES.

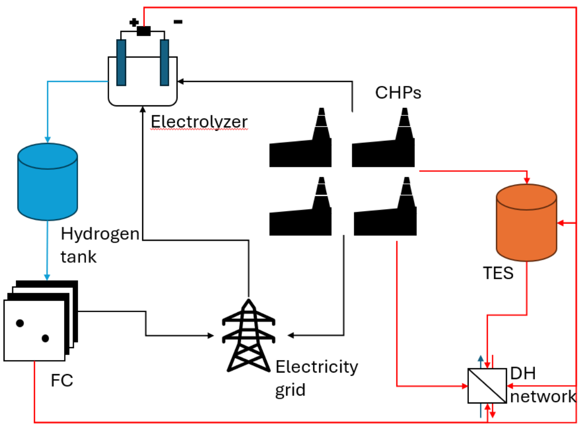

Fig. 2. Full system layout

2.1 Mixed Integer Linear Programming (MILP)

The general MILP formulation can be seen in Eq. (2), where x is a vector of the decision variables, c is a vector with the relationship between the decision variables, and A and b make the constraints where A is a matrix and b a vector. When running the MILP the binary constraints were relaxed to increase computational speed since generating training data without relaxing binaries was too computationally costly with the used hardware. However, if the model is simple enough or there is enough computing power the full model could be run with binary constraints. Other methods could also be used to increase the computational speed of the MILP.

$$\begin{aligned} &\min(c^T x) \\ st. \quad &Ax <= b \end{aligned} \tag{2}$$

The objective function can be seen in Eq. (3) where C_{eco} is the economic cost and C_{change} is a penalty to punish uneven operation of the storages and boilers. In the results when profit is referred to it refers to C_{eco}. The constraints that are considered in the MILP model can be seen in **Table 1** with what constraints apply to each unit.

$$C = C_{eco} + C_{change} \tag{3}$$

Table 1. List of constraints for MILP model

Constraint	CHP	HES	TES
Max/min power	✓	✓	✓
Ramp up/down	✓	✓	✓
SOC	×	✓	✓
Min up/down time	✓	×	×
on/off status	✓	✓	✓
electricity to heat ratio	✓	×	×
SOC start and end of time horizon	×	✓	✓
Heat loss to enviroment	×	×	✓
Startup status	✓	×	×
DH demand met	-	-	-
Transmission capacity out of plant	-	-	-

The MILP was tested in three different ways, the first is to just run the MILP for 1 year to get the optimal behaviour of the system. The second way is to give the MILP a constraint at the end of day SOC and then run the MILP for 36 hours but only taking the operation from the first 24. The third option is to use a rolling horizon optimization Bischi et al. (2019); Marquant et al. (2015) where the optimization is done daily with a one week time horizon. For this method the constraint on the SOC on the storages was also set to happen after one week. A optimality gap of 1% was used for the MILP otpimization

2.2 Machine Learning (ML) Algorithm

Some different ML algorithms were tested these are deep neural network (DNN), random forest (RF), historic gradient boosting (HGB), and Gaussian regression (GR). For all of these hyper parameter optimization was done and for the DNN different architectures of the network were also tested. The variables being predicted are the optimal daily or weekly charge and discharge from the HES and TES where training data is retrieved by running the MILP with a long time horizon. Some different input features were tested to get the lowest prediction error possible. The training and testing data were split by having training data be the data generated based on the first year and the testing data be the real data from the second year.

Prepossessing of the data was done before passing it to the ML algorithm. This prepossessing consisted of calculating the mean, max, and minimum daily electricity price and DH demand and monthly and weekly mean electricity price and DH demand. The data was also scaled with the electricity price, DH demand, SOC of the storages, day of the year, and weekday being scaled between 0 and 1 and the charge/discharge of the storages being scaled between -1 and 1 where -1 is fully discharging and 1 is fully charging. The loss metric used during the training of the ML models is mean square error.

A lot of the charge and discharge data of the storages is distributed around 0 to avoid any bias in the model training weights were used in the loss function to make all charge and discharge amounts be equally represented. This was done using DenseWeight which applies weights to different values of the training data based on kernel density estimation Steininger et al. (2021).

3. RESULTS

3.1 Long term Mixed Integer Linear Programming MILP

Figure 3 shows how the SOC of the HES and TES change when using MILP to optimize the system. The data used for optimization is from 2017, the year the ML algorithm makes its prediction. The storages does not start and end at the same SOC since the MILP optimization was done over 3 (2016-2018) with the SOC being constrained to be the same at the beginning of 2016 and the end of 2018.

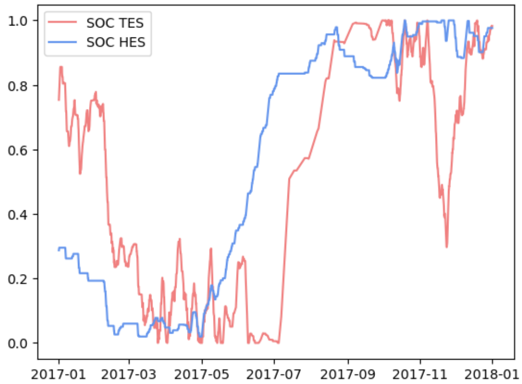

Fig. 3. SOC of storages based on MILP over predicted year

3.2 Full System One Day Prediction Horizon

The MAE as a percentage of the maximum occured charge/discharge of the storages for different ML algorithms can be seen in **Table 2**. As can be seen, the MAE of HGB is the lowest, however, this MAE is achieved by having the charge/discharge around zero all the time which does not lead to a good operation of the storages. The DNN on the other hand makes predictions that are based on the features and most of the time the decision to charge or discharge the storage is correct. The amount charged or discharged is however often wrong. This leads to the DNN operating the storages in a better way than HGB even though the MAE is higher. RF operates the storages in a similar way as HGB in that it tries to keep the charge/discharge around zero. GR operates the storages in a way that is somewhere between the strategy of the DNN and HGB. Because of this, the DNN is used as the ML algorithm for the rest of the results.

Table 2. Prediction performance of ML algorithms

ML method	MAE HES	MAE TES
DNN	25%	20%
RF	22%	19%
HGB	14%	16%
GR	23%	19%

The ML algorithm predicts the optimal SOC of the storages at the end of the day or week based on the features that can be seen in **Table 3**. Different combinations of features were tested but these were chosen since they gave the lowest mean absolute error (MAE). The data from the long term MILP and ML prediction using a DNN (deep neaural network) can be seen in **Figs. 4** and **5**, here a one-day prediction horizon was used. As can be seen, the prediction error is evenly spread except for predicting too low values when the HES is charged at maximum power. The HES has a MAE of 4500 kWh and the TES has a MAE of 332 000 kWh. This MAE is quite high, around 25% and 20% of the maximum daily charge/discharge power that occurred for the HES and TES. This error is however not important as long as the ML algorithm can give predictions that have a good operation of the storages in the daily MILP optimization.

In future research, this error could be reduced either by using a more complicated method such as first classifying if the storage will be charged, discharged, or not used, and then after that having 2 different specialized models for charging and discharging for each of the storages. Different ML models could also be used for different periods of the year. Another way to improve the results could be to use reinforcement learning and a one-day or week MILP model to more directly optimize based on the objective function.

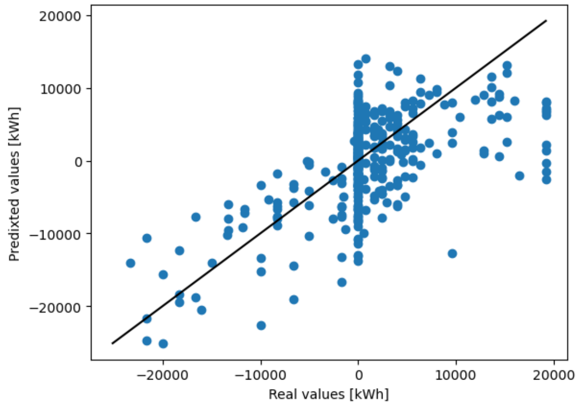

Fig. 4. Correlation between predicted and real charge/dis-charge for HES using DNN with a one-day prediction horizon

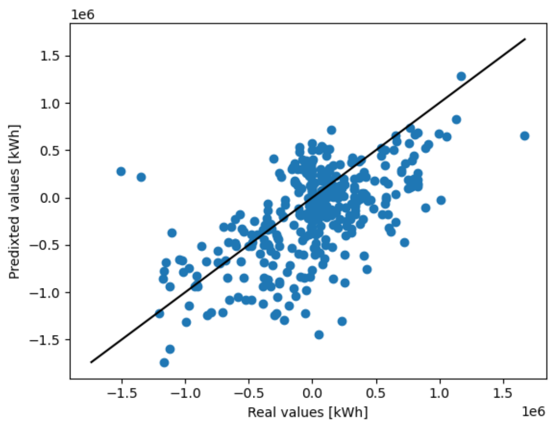

Fig. 5. Correlation between predicted and real charge/dis-charge for TES using DNN with a one-day prediction horizon

Table 3. Features used for ML

Feature	HES	TES
Current day mean electricity price and DH heat demand	✓	✓
Current day max electricity price and DH heat demand	✓	✓
Current day min electricity price and DH heat demand	✓	✓
Two weeks of mean electricity price	✓	✓
Two weeks of mean DH demand	×	✓
Two months of mean electricity price and DH demand	✓	✓
Time of year	✓	✓
Day of week	✓	✓
SOC of storages	✓	✓

The DNN performs better in some parts of the year and worse in others as can be seen in **Figs. 6** and **7**. The accuracy might be able to be improved if multiple ML models were trained for different parts of the year. The algorithm does however still mostly charge and discharge the storages at the correct time but the amount charged or discharged is often wrong. For both figures the DNN was trained using data generated based on data from 2016 and then tested using real real data from 2017. The optimality gap used for the MILP for both training and testing is 1%.

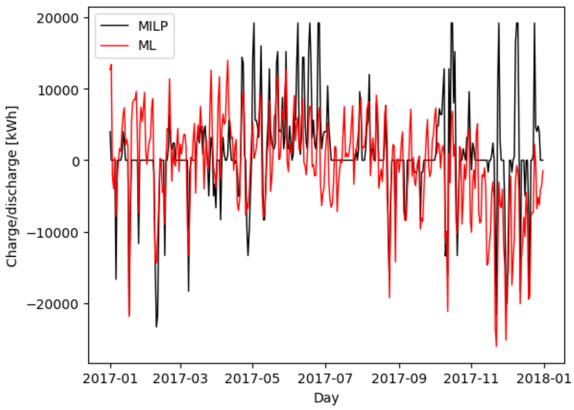

Fig. 6. Comparison charge/discharge HES MILP and DNN

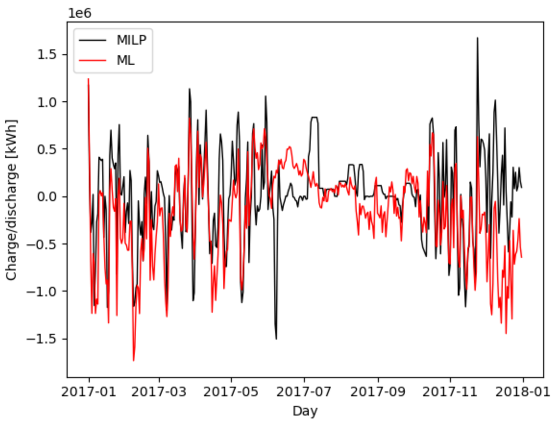

Fig. 7. Comparison charge/discharge TES MILP and DNN

The SOC of the HES for MILP and with constraints from the ML model can be seen in **Fig. 8** and the same for the TES in **Fig. 9**. Both the HES and TES SOC are quite different between using MILP and using DNN model constraints. The important thing here however is not that the SOC of the storages are the same but how profitable the operation of the entire energy system is in both scenarios. This will be discussed in the next section. The SOC pattern for the TES is however similar between the MILP and DNN model constraint with it discharging during the winter and charging during the summer. There is a difference in when the storages is being charged/discharged between the MILP and DNN. The reason for this could be that the DNN gets a low electricity price as an input and therefore charges the storage while the MILP does not charge the storage since it has all the data and knows that there will be an even cheaper electricity price in the future. In reality, a forecast for the electricity price would have to be used to operate the MILP in this way which could make the results of the MILP and DNN more similar.

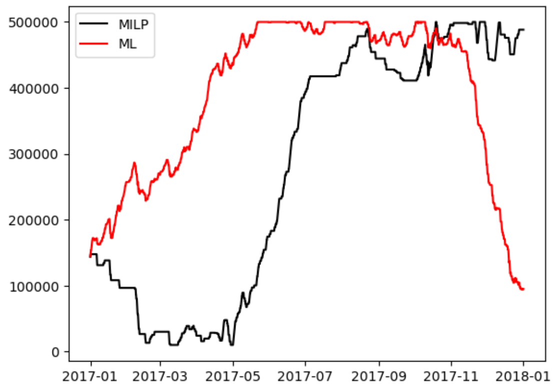

Fig. 8. Comparison SOC HES MILP and DNN

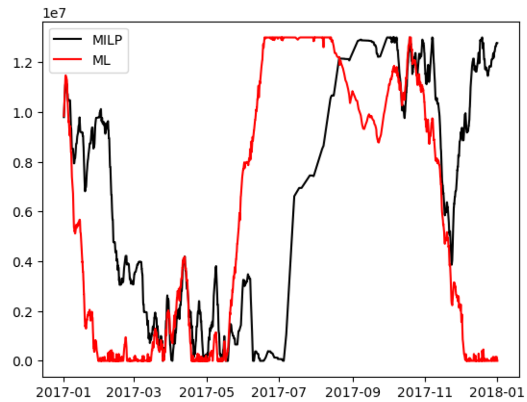

Fig. 9. Comparison SOC TES MILP and DNN

3.3 Comparison Cases

Table 4 compares the profitability of the daily MILP model with the ML constraints and when constraints are taken from the optimal operation of the previous year. The comparison is made as a ratio profitability compared to the long time MILP optimization results. When a one day time horizon is used the daily MILP model with constraints from DNN outperforms the other models with constraints from the optimal operation of the last year and other ML methods. All of them are also close to the long time horizon MILP optimization being 1, 4, and 6 percent away. Of note here is that the energy storages is only one part of the system so it is not the only factor effecting profitability.

When a one week time horizon is used, both the models with constraints from ML algorithms and last year's optimal operation have the same profitability as the long time MILP optimization. This is because a rolling time horizon is used with the constraint placed at the end of the week but only the operation from the first day is used, then the next day the optimization is run again. This means that even if there is some error in the prediction the operation of the first day can be good since there is no constraint for the SOC at the end of the first day.

Even though the MAE of the prediction is high the profitability is not greatly impacted. This is because the prediction is based on the electricity price DH demand and time of year. This gives a good operation even if the prediction is different from the value from the long time MILP.

Another benefit to using ML to give constraints to the MILP compared to using values from last year is that if a change in the electricity price or DH demand were to occur the operational plan can change. This makes the constraints given by the ML model more robust. The diversity and amount of training data generated and the number of years used to generate training data will also affect how robust the optimization is.

One thing to note when looking at **Table 4** is that the SOC the storages is not constrained to be the same at the end of the year which affects the profitability. The final SOC with the DNN constraints can be seen in **Figs. 8** and **9**. The final SOC when using the optimal results from the last year can be seen in the same figures but looking at the beginning of the year.

Table 4. Profitability comparison with constraints from ML models and taking SOC values from last year MILP optimization

Method \ Time horizon	1 day	1 week
DNN	0.97	1
Last year MILP	0.96	1
GR	0.94	1
HGB	0.94	1
RF	0.94	1

In **Table 5** the MAE and computational speed of the method using daily and weekly MILP with DNN constraints and the long term MILP model can be seen. Here the MAE is a percentage of the maximum occured charge/discharge of the storages. As can be seen, the optimization is fast both for the one day and one week time horizon when using constraints from the DNN. When running the MILP for a year with the full model the optimization time is over 30 hours which does not allow for day ahead planning. Another problem when running the model in this way is that a forecast for electricity price and DH demand is needed for the entire year. When using the ML algorithm only a forecast for the average electricity price for the next 2 weeks and next 2 months is needed. Even this can be removed with some increase in the prediction error.

The MAE of the HES prediction increases when doing a one week prediction while the MAE of the TES prediction decreases. This is likley because the seasonal patterns of the TES are stronger which makes a one week prediction easier since any irregular spikes in temperature will have a lower effect. The HES is more driven by the electricity price which has a less seasonal pattern so in this case the increase in features for the DNN only increase the MAE.

Table 5. MAE for different time resolutions and horizons

Time horizon	MAE HES	MAE TES	Computational time
1 day	25%	20%	1.78 s
1 week	30%	18%	12.33 s
Long term MILP	-	-	180.2 s*

*With binary constraints relaxed, the full model takes over 30 hours to run

Table 6 shows the error when doing predictions based on data from a simpler system. The prediction is slightly better for the system with only one boiler. The prediction is better on simpler systems since the behavior of the system becomes less complex and therefore easier to predict.

Table 6. Comparison different systems

System	MAE HES	MAE TES
One boiler only TES	-	20%
One boiler only HES	22%	-
One boiler HES and TES	22%	18%
Four boiler HES and TES	25%	20%

4. DISCUSSION

This method is fast enough to implement in real-time, when doing so retraining of the ML model should be done to catch any new patterns in electricity price or DH demand. The period between retraining will have to be decided based on testing different periods. When retraining the algorithm data could be generated again to increase the training data since the generated data is created based on real data and will therefore have some similar patterns.

Some things are required for it to be possible to use this method, the first is some historical data that can be used for training and creation of synthetic data or a way of creating realistic synthetic data without any real data. Some long-term energy storage is also needed for this method to be effective, if no long-term energy storage exists conventional methods are more suitable for optimization.

The use case for this kind of optimization method is in cases where the optimization has to be done in a short time or where the optimization has to be done a lot of times, for both of these cases LTES should also be a part of the energy system. For the case where optimization has to be done fast it could be at a powerplant where the MILP model is too complex to optimize over a long time horizon, then this method can be used to speed up the optimization. For a case where optimization has to be done many times, there could be a case where the MILP optimization is part of an inner loop where it has to run many times per iteration of some other optimization layer.

Since the synthetic data is only used for training the ML algorithms and is generated without using any of the testing data the use of synthetic data should not have any negative impact on the results. The use of synthetic data might also not be needed if enough historical data is available, other methods of generating synthetic data could also be used.

5. CONCLUSIONS

Using ML to reduce the time horizon of a MILP model by constraining the SOC of LTES gives a similar economic operation to letting the MILP run over a long time horizon. The MAE of the prediction is large but the economic operation is still good with this method. This method outperforms using the past years storage operation when running the MILP daily and has an equal performance when running the MILP weekly.

The ML method that gives the best operation of the storages is a DNN.

ACKNOWLEDGEMENTS

This work is supported by project ProPlan under the RESILIENT competence center, financed by the Swedish Energy Agency (2021-90273) and co-financed by Mälardalen University and industrial partners. The support received by Mälarenergi AB and Eskilstuna Strängnäs Energi och Miljö AB is appreciated.

REFERENCES

Alabi, T.M., Lawrence, N.P., Lu, L., Yang, Z., and Gopaluni, R.B. (2023). Automated deep reinforcement learning for real-time scheduling strategy of multi-energy system integrated with post-carbon and direct-air carbon captured system. *Applied Energy*, 333. doi: 10.1016/j.apenergy.2022.120633.

Bahlawan, H., Losi, E., Manservigi, L., Morini, M., Pinelli, M., Spina, P.R., and Venturini, M. (2022). Optimization of a renewable energy plant with seasonal energy storage for the transition towards 100 *Renewable Energy*, 198, 1296–1306. doi:10.1016/j.renene.2022.08.126.

Bischi, A., Taccari, L., Martelli, E., Amaldi, E., Manzolini, G., Silva, P., Campanari, S., and Macchi, E. (2019). A rolling-horizon optimization algorithm for the long term operational scheduling of cogeneration systems. *Energy*, 184, 73–90. doi:10.1016/j.energy.2017.12.022.

Brey, J.J. (2021). Use of hydrogen as a seasonal energy storage system to manage renewable power deployment in spain by 2030. *International Journal of Hydrogen Energy*, 46, 17447–17457. doi:10.1016/j.ijhydene.2020.04.089.

Bruninx, K. and Delarue, E. (2017). Improved energy storage system unit commitment scheduling. doi: 10.1109/PTC.2017.7981000.

Dong, H., Fu, Y., Jia, Q., Zhang, T., and Meng, D. (2023). Low carbon optimization of integrated energy microgrid based on life cycle analysis method and multi time scale energy storage. *Renewable Energy*, 206, 60–71. doi:10.1016/j.renene.2023.02.034.

Ghaljehei, M. and Golkar, M.A. (2017). Effect of optimal generation scheduling of compressed air energy storage and wind power generation on economic and technical issues. doi:10.1109/SGC.2017.8308854.

Helseth, A., Fodstad, M., Askeland, M., Mo, B., Nilsen, O.B., Pérez-Díaz, J.I., Chazarra, M., and Guisández, I. (2017). Assessing hydropower operational profitability considering energy and reserve markets. *IET Renewable Power Generation*, 11, 1640–1647. doi:10.1049/iet-rpg.2017.0407.

International Energy Agency (2024). Managing the seasonal variability of electricity demand and supply. URL `https://www.iea.org/reports/managing-the-seasonal-variability-of-electricity-demand-and-supply`.

Jahns, C., Podewski, C., and Weber, C. (2020). Supply curves for hydro reservoirs – estimation and usage in large-scale electricity market models. *Energy Economics*, 87. doi:10.1016/j.eneco.2020.104696.

Marquant, J.F., Evins, R., and Carmeliet, J. (2015). Reducing computation time with a rolling horizon approach applied to a milp formulation of multiple urban energy hub system. volume 51, 2137–2146. Elsevier B.V. doi:10.1016/j.procs.2015.05.486.

Marzi, E., Morini, M., Saletti, C., Vouros, S., Zaccaria, V., Kyprianidis, K., and Gambarotta, A. (2023). Power-to-gas for energy system flexibility under uncertainty in demand, production and price. *Energy*, 284. doi: 10.1016/j.energy.2023.129212.

Mi, Y., Liu, C., Yang, J., Zhang, H., and Wu, Q. (2021). Low-carbon generation expansion planning considering uncertainty of renewable energy at multi-time scales. *Global Energy Interconnection*, 4, 261–272. doi: 10.1016/j.gloei.2021.07.005.

Nordpool (2024). Nordpool. URL `https://www.nordpoolgroup.com/en/`.

Saletti, C., Morini, M., and Gambarotta, A. (2022). Smart management of integrated energy systems through co-optimization with long and short horizons. *Energy*, 250. doi:10.1016/j.energy.2022.123748.

Sleptchenko, A. and Sgouridis, S. (2019). Joint optimization of energy production storage. doi: 10.1109/IEA.2019.8714874.

Steininger, M., Kobs, K., Davidson, P., Krause, A., and Hotho, A. (2021). Density-based weighting for imbalanced regression. *Machine Learning*, 110, 2187–2211. doi:10.1007/s10994-021-06023-5.

Su, S., Tan, D., Li, X., Li, X., Chen, Z., Lian, R., and Soronga, B.K.L. (2023). Multi-time scale coordinated optimization of new energy high permeability power system considering flexibility requirements. *Journal of Electrical Engineering and Technology*, 18, 815–828. doi:10.1007/s42835-022-01244-7.

Wogrin, S., Galbally, D., and Reneses, J. (2016). Optimizing storage operations in medium- and long-term power system models. *IEEE Transactions on Power Systems*, 31, 3129–3138. doi:10.1109/TPWRS.2015.2471099.

Zhang, L., Dai, W., Zhao, B., Zhang, X., Liu, M., Wu, Q., and Chen, J. (2023). Multi-time-scale economic scheduling method for electro-hydrogen integrated energy system based on day-ahead long-time-scale and intra-day mpc hierarchical rolling optimization. *Frontiers in Energy Research*, 11. doi:10.3389/fenrg.2023.1132005.