ARTIFICIAL NEURAL NETWORKS FOR REAL-TIME MODEL PREDICTIVE CONTROL OF ORGANIC RANKINE CYCLES FOR WASTE HEAT RECOVERY

Yannic Vaupel[1], Adrian Caspari[1], Nils C. Hamacher[1], Wolfgang R. Huster[1], Adel Mhamdi[1], Ioannis G. Kevrekidis[2], Alexander Mitsos[1*]

[1] RWTH Aachen University - Process Systems Engineering (AVT.SVT)
Forckenbeckstr. 51, 52074 Aachen, Germany
amitsos@alum.mit.edu

[2] Johns Hopkins University, Department of Chemical and Biomolecular Engineering
Baltimore, MD, USA
yannis@princeton.edu

* Corresponding Author

ABSTRACT

Recovering waste heat from the exhaust gas of heavy-duty diesel trucks using a bottoming organic Rankine cycle (ORC) is a promising option to reduce fuel consumption. In contrast to most other applications of ORCs, e.g., geothermal or solar-thermal, the heat source in automotive applications is subject to strong fluctuations with limited predictability. Consequently, controlling the ORC system to maintain safe and efficient operation is a challenging task. Nonlinear model predictive control (NMPC) has been proposed for ORC systems and showed promising *in silico* results. It suffers, however, from a high computational expense and real-time capable implementation on vehicle hardware is questionable. Several methods are available that aim at shifting the majority of the computational load to the design phase of the controller, reducing on-line resource demand. In this work, we apply artificial neural networks (ANN) *in silico* to learn the control law of the NMPC controller off-line. We obtain training data from various NMPC scenarios with different initial conditions and heat source conditions using our in-house dynamic optimization tool DyOS. Subsequently, we apply the ANN-based controller *in silico* to different scenarios of transient heat source conditions. We compare the results to the NMPC solution obtained with DyOS and findings indicate that performance loss associated with the ANN-based controller is marginal while the control policy can be obtained at negligible computational cost.

1. INTRODUCTION

Control of organic Rankine cycle (ORC) systems for waste heat recovery in heavy-duty diesel trucks is a challenging task as the vehicles are operated in street traffic, and hence, the heat source fluctuates strongly. Moreover, predictability of these fluctuations is limited. To maintain efficient and safe operation under those conditions, several authors have proposed model predictive control (MPC) schemes ranging from linear tracking MPC, e.g., (Feru et al., 2014), to (economic) nonlinear MPC, e.g., (Petr et al., 2015, Peralez et al., 2015). Nonlinear MPC in particular promises good disturbance rejection, yet it is computationally expensive. Albeit CPU times of less than 100 ms for solving an NMPC problem for WHR in a truck have recently been reported on a desktop computer (Guerrero Merino, 2018), implementation on in-vehicle hardware remains questionable. Most authors try to reduce model complexity to maintain computational tractability (Petr et al., 2015, Peralez et al., 2015).

Several ideas to reduce the computational costs of NMPC are available. Bemporad et al. (2002) proposed to exploit that the solution of MPC is parametric in the systems state variables and to solve the

parametric problem *a priori*. While the approach exhibits strong theoretical foundations in terms of its stability properties, solving the parametric problem scales badly for increasing state vector size. Another method, which has been revisited recently, is to learn the NMPC control policy using machine learning techniques e.g., (Åkesson and Toivonen, 2006), (Lucia and Karg, 2018), (Lucia et al., 2018), (Karg and Lucia, 2018), (Lovelett et al., 2018). The approach is straightforward to implement, however, it suffers from the typical weaknesses of machine learning approaches, i.e., the course of dimensionality and little or no extrapolation capability.

Hence, it is of great significance to obtain a dense sampling of the state space and to ensure that the controller does not operate outside the sampled region. Lucia and Karg (2018) showed that in practice an ANN-based controller might exhibit some extrapolation capabilities. Lovelett et al. (2018) proposed to combine an ANN-based controller with manifold learning techniques to achieve a simpler correlation to learn, especially for large-scale systems. In this manuscript we apply an ANN-based controller to an ORC waste heat recovery system for a heavy-duty diesel truck. Our focus is on maintaining a desired superheat set-point.

The remainder of this manuscript is structured as follows. We present the system considered and its model in Sec. 2. The method for gathering the training data and the training procedure are described in Sec. 3 followed by an evaluation of the controller performance in Sec. 4. We give our conclusions and provide an outlook on future research in Sec. 5.

2. MODEL DESCRIPTION

We consider an ORC for waste heat recovery in a heavy-duty diesel truck (cf. Fig. 1.) The working

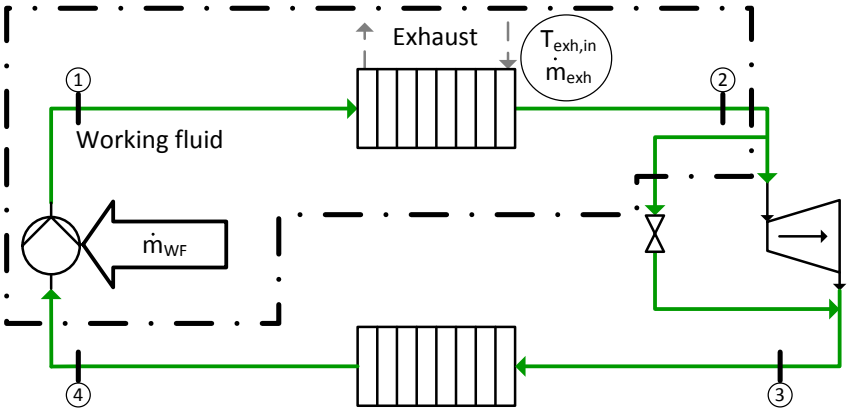

Figure 1: Topology of the examined system. The WF is indicated by the solid green line and the exhaust gas by the dashed gray line. The manipulated variable is indicated by the arrow and the disturbances by the circle. The system boundary for this work is indicated by the dot-dashed black line.

fluid (WF) ethanol is compressed in a pump (4→1) and then preheated, evaporated and superheated in a single heat exchanger (1→2). The superheated WF is expanded in a turbine (2→3) and then condensed in a condenser (3→4).

We focus on maintaining a desired superheat set-point at the heat exchanger WF outlet. For simplicity, we assume that the condenser operates at ambient pressure and is adequately controlled to cool the WF to fixed subcooling which renders a sophisticated condenser model unnecessary. Under this assumption, the turbine speed solely serves for optimizing power output. Hence, we did not implement a turbine or condenser model and the WF inlet massflow $\dot{m}_{WF,in}$ is the sole degree of freedom. The evaporator is represented by a dynamic model using the moving boundary approach and the pump by a pseudo-

stationary model as suggested in Huster et al. (2018). However, for simplicity we assume constant heat transfer coefficients and a simple geometry typical for waste heat recovery in a truck. Key parameters for the model are specified in Tab.1. We assume that the WF leaving the evaporator is always in super-

$\alpha_{i,0}$	$\alpha_{i,1}$	$\alpha_{i,2}$	α_{exh}	A	A_w	A_{Laval}	b_{WF}	b_{exh}	l_{tot}
	$\mathrm{W/(m^2K)}$				$\mathrm{m^2}$			m	
100	900	50	45	0.004	0.03	0.00002	12	80	0.3

Table 1: Key heat exchange and geometry model parameters.

heated state so that no discrete events occur which would complicate the use of dynamic optimization for NMPC (Vaupel et al., 2019).

3. METHOD: DATA ACQUISITION AND TRAINING

To learn the NMPC control law, it is of great significance to provide a set of data for training which adequately samples the state space. Recent efforts on ANN-based MPC, e.g., (Lucia and Karg, 2018), (Lovelett et al., 2018), present chemical reactions in stirred tank reactors as a case study. Sampling the state space by providing a set of practically relevant initial conditions is straightforward for such systems and can be achieved by simply altering initial concentrations and reactor volume.

In contrast, it is complicated to provide a variety of physically meaningful initial conditions for the ORC system, in particular for wall temperatures and zone lengths. For example, it would make little sense to initialize the wall temperature of the superheated zone at a smaller value than the wall temperature of the subcooled zone. Such an initialization might even result in an infeasible DAE initialization or integration failure. We use an optimization-based procedure to address this issue. We solve 200 dynamic optimization problems using our in-house tool DyOS (Caspari et al., 2019), each starting from the same feasible initial point but with a different time-invariant exhaust gas massflow $\dot{m}_{exh}$. The optimizer minimizes the deviation of the superheat T_{sup} at final time t_f from a desired superheat T_{sup}^{des}, which is different in each run, thus achieving a set of well distributed initial points. As a degree of freedom, the optimizer can choose $\dot{m}_{WF,in}$ which is constant for each run. Thereby, and by choosing t_f = 2000 s we achieve that the system is effectively at steady-state at t_f. We provide 200 combinations of $\dot{m}_{exh} \in [0.1\ \mathrm{kg/s}, 0.6\ \mathrm{kg/s}]$ and a desired superheat $T_{sup}^{des} \in [10\ \mathrm{K}, 50\ \mathrm{K}]$ through latin hypercube sampling (LHS) and assume a constant exhaust gas inlet temperature of $T_{exh,in}$ = 600 K. The dynamic optimization problem reads

$$\min_{u,x(t)} \quad \left(T_{sup}\left(t_f\right) - T_{sup}^{des,k}\right)^2 \tag{1}$$

$$\text{s.t.} \quad M\dot{x}\left(t\right) = f\left(x\left(t\right), y\left(t\right), u, d\right) \tag{2}$$

$$y\left(t\right) = g\left(x\left(t\right), y\left(t\right), u, d\right) \tag{3}$$

$$x\left(t=0\right) = x_0 \tag{4}$$

$$u_{min} \le u \le u_{max}\,. \tag{5}$$

where (1) is a Mayer-type objective function that minimizes deviation of the superheat T_{sup} from T_{sup}^{des} at final time t_f. The model is described by the differential and and algebraic equations (2)-(3) where M is the constant mass matrix of the DAE-system, x the differential state variables, y the algebraic state variables, u the time-invariant input and d the time-invariant disturbance, i.e., $\dot{m}_{exh}$ and $T_{exh,in}$. The initial state of the model is specified in (4) and the problem is input-constrained with (5).

We use the resulting steady-state optimal state vectors $x^{opt,k}\left(d^k\right)$ as initial values for the following NMPC runs and combine them with a different LHS of $\dot{m}_{exh}$. We then carry out NMPC with these sets, minimizing the devations of the superheat trajectory from desired superheat of 30 K. The NMPC problem

is solved repeatedly with a sampling time of $\Delta t = 5$ s, i.e., at every instant k and minimizes the integral over the prediction horizon N_P of the deviation from desired superheat.

$$\min_{u(t),x(t)} \quad \int_{t_k}^{t_k+N_P\Delta t} \left(T_{sup}(t) - 30\ \mathrm{K}\right)^2 \mathrm{d}t \tag{6}$$

$$\text{s.t.} \quad M\dot{x}(t) = f(x(t), y(t), u(t), d) \tag{7}$$

$$y(t) = g(x(t), y(t), u(t), d) \tag{8}$$

$$x(t = t_k) = x_0 \tag{9}$$

$$u_{min} \leq u(t) \leq u_{max} \tag{10}$$

Here, we chose a piecewise-constant discretization of the inputs $u(t)$. We chose a control horizon $N_C = 8$ and a prediction horizon $N_P = 10$. The dynamic optimization problem was solved with DyOS at every sampling instant and the total duration of each run was 200 s, resulting in 40 samples per simulation.

A projection of the trajectories in the state space on pressure p and WF outlet enthalpy h_{out} of all 200 NMPC simulations is depicted in Fig. 2. We can see that we achieve a dense sampling in the state-

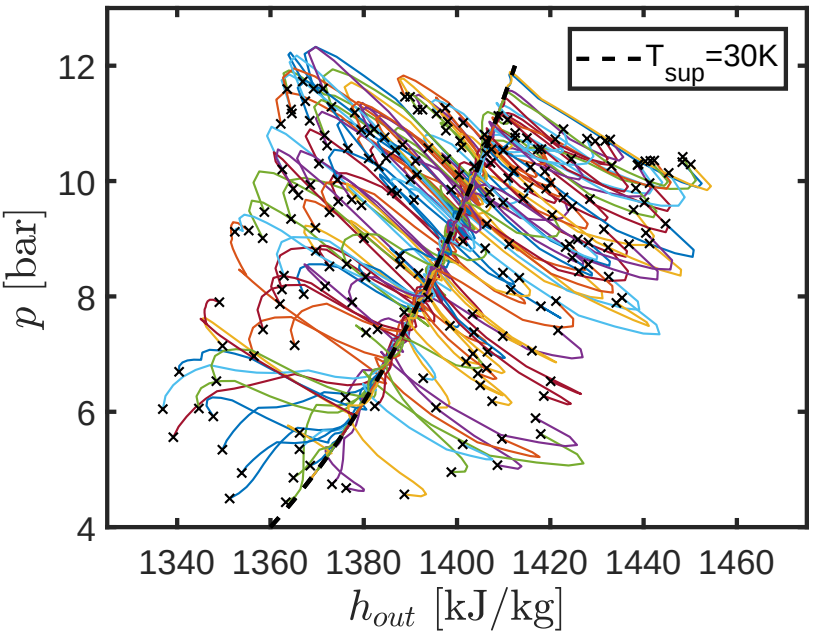

Figure 2: State space projection on p and h_{out} of the 200 NMPC simulations. Crosses indicate initial values. The controllers regulates the system to achieve the desired superheat.

space and that the trajectories converge to the desired superheat. The data is used for training the ANN and the amount of data is considered to be sufficiently dense to allow for effective interpolation. In total, we obtained the datasets $\mathbf{u}^* \in \mathbb{R}^{8000\times1}$, $\mathbf{x}^* \in \mathbb{R}^{8000\times7}$ and $\mathbf{d}^* \in \mathbb{R}^{8000\times1}$. The solution to the NMPC problem for our scenario is parametric in x and d. Consequently, we learn the mapping from those quantities to the optimal control policy u. The NMPC controller provides the optimal sequence $\mathbf{u_k} := (u(t_k), u(t_k + \Delta t), \ldots, u(t_k + (N_c - 1)\Delta t))$ over the control horizon consisting of N_C elements at every instance. However, as only the first calculated input of this sequence is applied to the plant before the NMPC problem is resolved, we only learn the mapping for this quantity. The training process is illustrated in Fig. 3. We did not include T_{sup} as an input to the ANN as it is fixed to 30 K. We trained the ANNs using the Levenberg-Marquardt algorithm in Matlab's Neural Network Toolbox (Beale et al., 2010) and assumed all states to be measurable, thus eliminating the need to implement state estimation.

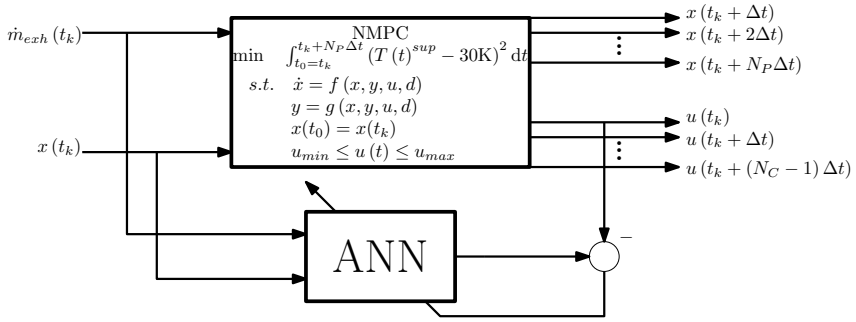

Figure 3: Illustration of the training process of the ANN. Only the first control input $u(t_k)$ is considered.

Due to the parametric nature, the mapping is only dependent on the measured state and disturbance at the current time. Hence, we used a feed-forward neural network with the hyperbolic tangent activation function. We tested several different ANN architectures by varying the number of hidden layers and the number of neurons per layer. In total we tested architectures consisting of one to four hidden layers consisting of 10, 20 and 30 neurons each. To obtain reliable results, we executed five training runs for each architecture.
We found that all trained ANNs achieved good results, differences in performance were small and even a shallow network with only one hidden layer consisting of ten neurons exhibited satisfactory performance. While this is an interesting observation, we do not claim that it is generally true for learning NMPC policies as we solved a rather simple problem with only one manipulated variable here.

4. CASE STUDY: ANN-BASED CONTROLLER VS. NMPC

We test our proposed ANN-based controller in further scenarios *in silico*. To assess its performance, we compare the resulting control policy and deviations from desired superheat to the solution obtained by solving (6)-(10). We used an ANN consisting of two hidden layers of 20 neurons each and we examined a total of ten structurally similar scenarios. All scenarios include two steps in the exhaust gas mass flow which, however, do occur at different times. To generate the different exhaust gas profiles, we sampled a LHS with three values for the different levels of $\dot{m}_{exh}$ and another LHS for the time-points where the steps occur. We obtained initial values with the procedure described in Sec. 3. However, to test the robustness of the ANN-based controller with respect to conditions not encountered during training, we sampled $T_{sup}^{des} \in [10\text{ K}, 60\text{ K}]$ and $\dot{m}_{exh} \in [0.1\text{ kg/s}, 0.65\text{ kg/s}]$.

We then solved these scenarios with the ANN-based controller and with NMPC and assumed that the solution can be computed without time delay. To compare both approaches we use the average deviation from the desired superheat ε_{avg}.

$$\varepsilon_{avg} = \frac{\int_0^{t_f} \sqrt{\left(T_{sup}(t) - 30\text{ K}\right)^2} \mathrm{d}t}{t_f} \tag{11}$$

Tab. 2 presents the values of ε_{avg} for all executed simulations. In general, both controllers perform well

$\boldsymbol{\varepsilon_{avg}}$ **[K]**										
Run	1	2	3	4	5	6	7	8	9	10
ANN	.109	.722	.447	.241	.996	.410	.847	.101	.280	.276
NMPC	.096	.636	.431	.241	.848	.319	.809	.100	.271	.244

Table 2: Average deviation from desired superheat for all ten runs.

in all scenarios with the average deviation always smaller than 1 K. However, NMPC exhibited slightly better performance for each run. Higher average deviation in certain runs does not necessarily indicate bad controller performance. In fact, it is predominantly due to deviations of the superheat from 30 K at the initial point of the respective run. The average CPU time for obtaining the ANN-based control policy was 20 ms on a desktop computer with an Intel(R) Core(TM) i7-4790 CPU and 16 GB RAM.

For further assessment, we examine the scenario with the highest relative deviation between the ANN-based controller and NMPC controller, i.e., Run 6. The exhaust gas mass flow is provided in Fig. 4a. Interestingly, it is a scenario where one of the exhaust gas mass flow levels is outside the training set. This is in agreement with intuition as it can be expected that the loss in controller performance is larger for disturbances outside the training set. As in all scenarios in this manuscript, $T_{exh,in}$ remains constant at 600 K. Fig. 4c depicts the control action taken by the ANN-based controller and the NMPC controller. The resulting superheat is shown in Fig. 4d. The control policy computed by the ANN is initially iden-

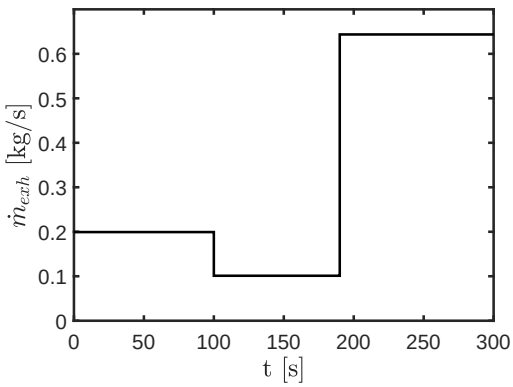

(a) Exhaust gas mass flow disturbance profile

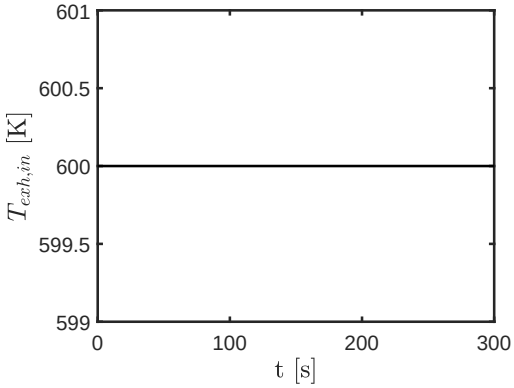

(b) Exhaust gas temperature disturbance profile

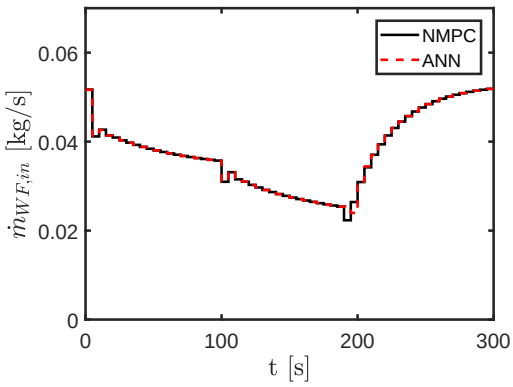

(c) Control input $\dot{m}_{WF,in}$

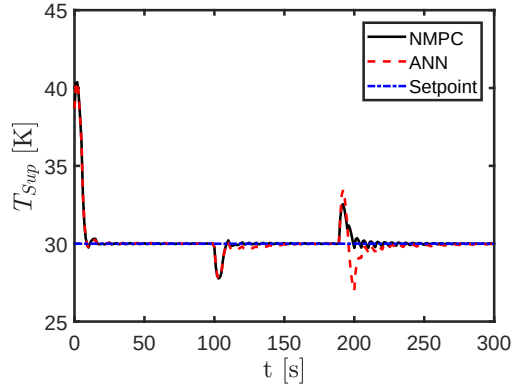

(d) Superheat

Figure 4: Disturbances, controls and resulting superheat for Run 6.

tical to the sequence provided by NMPC. Both controllers react well to the change in $\dot{m}_{exh}$ at t = 100 s though the ANN-based controller takes marginally longer to reach the desired superheat. The second step in $\dot{m}_{exh}$ occurs at t = 190 s and reaches a level outside the training set. Though the ANN-based controller is capable to reject this disturbance, the NMPC exhibits superior performance. The deviations for the ANN-based Controller remain, however, sufficiently small.

5. CONCLUSION AND OUTLOOK

Nonlinear model predictive control of ORCs for waste heat recovery in vehicles is a challenging problem due to the transient and unpredictable heat source conditions, the limited computational resources and

necessary high sampling rates. NMPC has been proposed for this system by several authors but typically requires significant model simplification. Large computational cost is a well known issue in many other NMPC applications and several approaches to overcome it have been developed.

In this work, we presented a machine learning approach, where an ANN was used to learn the NMPC control policy as function of the system state and the measured disturbance. To obtain data for training, we solved to 200 NMPC simulations with varying initial conditions and exhaust gas mass flow. In those cases, we aimed to track a superheat set-point of 30 K. We trained the ANN-based controller using shallow and deep ANNs consisting of up to four hidden layers and up to 30 neurons per layer. Differences in performance were small and even the result for the smallest ANN with one hidden layer of ten neurons showed satisfactory performance. We then performed a case study where we executed ten NMPC simulations and compared the results to the ANN-based controller. Two steps of random height at random time-points occurred in the exhaust gas mass flow during these simulations. Further we included initial values and exhaust gas mass flows which were slightly outside of the training set. Our results show that the control policy learned by the ANN exhibits only small deviations from the solution of NMPC while requiring negligible computational expense. However, small deviations do occur, in particular when the disturbance changes while the system is at desired superheat but not yet at steady-state. This might improve when the training scenarios are modified.

The issue at large is that there might be a lack of confidence in the method as long as it solely relies on machine learning techniques. Hence, future work should aim at combining machine learning with NMPC to obtain a good compromise between computational tractability and rigorous results. The method used herein has the benefit of being able to provide control signals virtually without delay at the cost of solving many problems *a priori*. However, these calculations have to be done again if for example the sampling rate or control horizon is to be changed. Therefore, possibilities to make the method more flexible with respect to the controller parameters and reduce the need for retraining should be explored. We used many simplifying assumptions to present the method in this manuscript. Future work should test the performance under more realistic assumptions, e.g., introducing measurement noise, plant-model mismatch, state estimation and further degrees of freedom and disturbances. Finally, it is questionable how the method works for large-scale problems. Lovelett et al. (2018) proposed a method to exploit the fact that there might be a variable transformation which provides a more adequate space to sample the data. This could for example be tested with an ORC model using a finite volume evaporator model.

NOMENCLATURE

Latin		
d	disturbance	(-)
h	specific enthalpy	(kJ/kg)
$\dot{m}$	mass flow	(kg/s)
M	mass matrix	(-)
N_C	control horizon	(-)
N_P	prediction horizon	(-)
T	temperature	(K)
t	time	(s)
u	input variable	(-)
x	differential state variable	(-)
y	algebraic state variable	(-)

Greek		
ε	temperature deviation	(K)

Subscript	
0	initial
avg	average
exh	exhaust gas
f	final
in	inlet
k	counter variable
max	maximal
min	minimal
out	outlet
WF	working fluid

Superscript	
des	desired
opt	optimal

REFERENCES

Åkesson, B. M. and Toivonen, H. T. (2006). A neural network model predictive controller. *Journal of Process Control*, 16(9):937–946.

Beale, M. H., Hagan, M. T., and Demuth, H. B. (2010). Neural network toolbox. *User's Guide, MathWorks*, 2:77–81.

Bemporad, A., Morari, M., Dua, V., and Pistikopoulos, E. N. (2002). The explicit linear quadratic regulator for constrained systems. *Automatica*, 38(1):3–20.

Caspari, A., Hannemann-Tamás, R., Bremen, A., Faust, J. M. M., Jung, F., Kappatou, C. D., Sass, S., Vaupel, Y., Mhamdi, A., and Mitsos, A. (2019). DyOS - A Framework for Optimization of Large-Scale Differential Algebraic Equation Systems. In *29th European Symposium on Computer-Aided Process Engineering (ESCAPE-29)*.

Feru, E., Willems, F., de Jager, B., and Steinbuch, M. (2014). Modeling and control of a parallel waste heat recovery system for euro-VI heavy-duty diesel engines. *Energies*, 7(10):6571–6592.

Guerrero Merino, E. E. (2018). Real-time optimization for estimation and control: Application to waste heat recovery for heavy duty trucks.

Huster, W. R., Vaupel, Y., Mhamdi, A., and Mitsos, A. (2018). Validated dynamic model of an organic rankine cycle (ORC) for waste heat recovery in a diesel truck. *Energy*, 151:647–661.

Karg, B. and Lucia, S. (2018). Efficient representation and approximation of model predictive control laws via deep learning.

Lovelett, R. J., Dietrich, F., Lee, S., and Kevrekidis, I. G. (2018). Some manifold learning considerations towards explicit model predictive control.

Lucia, S. and Karg, B. (2018). A deep learning-based approach to robust nonlinear model predictive control. *IFAC-PapersOnLine*, 51(20):511–516.

Lucia, S., Navarro, D., Karg, B., Sarnago, H., and Lucia, O. (2018). Deep learning-based model predictive control for resonant power converters.

Peralez, J., Tona, P., Nadri, M., Dufour, P., and Sciarretta, A. (2015). Optimal control for an organic rankine cycle on board a diesel-electric railcar. *Journal of Process Control*, 33:1–13.

Petr, P., Schröder, C., Köhler, J., and Gräber, M. (2015). Optimal Control of Waste Heat Recovery Systems Applying Nonlinear Model Predictive Control. In Lemort, V., Quoilin, S., de Paepe, M., and van den Broek, M., editors, *Proceedings of the 3rd International Seminar on ORC Power Systems*, pages 1183–1192.

Vaupel, Y., Huster, W. R., Holtorf, F., Mhamdi, A., and Mitsos, A. (2019). Analysis and improvement of dynamic heat exchanger models for nominal and start-up operation. *Energy*, 169:1191–1201.

ACKNOWLEDGEMENT

The authors gratefully acknowledge the financial support of the Kopernikus project SynErgie by the Bundesministerium für Bildung und Forschung (BMBF) and the project supervision by the project management organization Projektträger Jülich (PtJ). IGK acknowledges the partial support of DARPA.