in: KI 2009 - Advances in Artificial Intelligence (KI), Paderborn, LNCS 5803, pp. 395-402, Springer 2009

Machine Learning Techniques for Selforganizing Combustion Control

Erik Schaffernicht[1], Volker Stephan[2], Klaus Debes[1], and Horst-Michael Gross[1]

[1] Ilmenau University of Technology
Neuroinformatics and Cognitive Robotics Lab
98693 Ilmenau, Germany
[2] Powitec Intelligent Technologies GmbH
45219 Essen-Kettwig, Germany
Erik.Schaffernicht@Tu-Ilmenau.de

Abstract. This paper presents the overall system of a learning, selforganizing, and adaptive controller used to optimize the combustion process in a hard-coal fired power plant. The system itself identifies relevant channels from the available measurements, classical process data and flame image information, and selects the most suited ones to learn a control strategy based on observed data. Due to the shifting nature of the process, the ability to re-adapt the whole system automatically is essential. The operation in a real power plant demonstrates the impact of this intelligent control system with its ability to increase efficiency and to reduce emissions of greenhouse gases much better then any previous control system.

1 Introduction

The combustion of coal in huge industrial furnaces for heat and power production is a particular complex process. Suboptimal control of this kind of process will result in unneccesary emissions of nitrogen oxides (NOx) and in a reduced efficiency, which equals more carbon dioxides emitted for one unit of energy.

Hence, optimal control of this process will not only reduce the costs but decrease the environmental impact of coal combustion. But finding such a control strategy is challenging. Even today, the standard approach is a human operator, which hardly qualifies as optimal. Classical control approaches are limited to simple PID controllers and expert knowledge, sometimes augmented by computational fluid dynamics (CFD) simulations in the setup stage [1].

The main reason for this is the process itself. The hazardous environment restricts sensor placements and operations. It is not possible to measure all relevant values that are known to influence the process. High temperatures, abrasive material flows, and problems like slagging (molten ash that condenses on the furnace wall) interfere with reliability of sensor readings, thus, the measurement noise and stochasticity of the observations is very high.

In the light of all these facts, we approach the problem from another point of view. Our presumption is, that one has to try to extract more information from the furnace directly and to learn from the data with Machine Learning (ML) techniques to build the optimal controller.

This work is a complete overhaul of a previous system [2], but we are not aware of any other work about coal combustion and its control by ML techniques.

For several reasons we follow the basic idea of our approach to explicitly avoid the usage of expert knowledge. First, for many applications there is no expert knowledge available, especially if new information about the process comes into play. Secondly, industrial combustion processes usually are time varying due to changing plant properties or different and changing fuel-characteristics. Thus, initially available expert knowledge becomes invalid over time and needs to be updated too often. In consequence, we prefer a self-organizing, adaptive and learning control architecture that develops a valid control strategy based on historical data and past experiences with the process only.

We will give an overview of the complete control system (and the tasks it learns to solve) in Section 2. A more detailed explanation of the subsystems involved is given in the Sections 3 to 5. Finally, the adaptivity of the overall system is discussed in Section 6, and first results of online operation are shown in Section 7.

2 System Overview

The system that is presented in this paper is designed with the goal in mind to deploy it to existing power plants. Hence, it operates on top of the infrastructure provided by the plant and the main interface is formed between the Distributed Control System (DCS) of the power plant and the developed system. Besides additional computer hardware, special CMOS cameras are installed to observe the flames inside the furnace directly. The video data and the traditional sensor readings from the DCS are fed into the system, which is depicted in Fig. 1.

To handle the high dimensional pixel data generated with the cameras, a feature transformation is applied to preserve the information bearing parts while discarding large amounts of irrelevant data. Details and further references can be found in Section 3. Thereafter, a feature selection is performed on the extracted image data, the spectral data, and the sensory data from the DCS to reduce the number of considered inputs even more. A fairly simple filter approach is used for this task, which is presented in Section 4. For the adaptive controller itself, several approaches with different control paradigms were considered. We compared Nonlinear Model Predictive Control (MPC), Bayesian Process Control (BPC) and Neuroevolutionary Process Control (NEPC). Some basics about these approaches are presented in Section 5. Due to the nature of the combustion process (e.g. fuel changes), a certain life-long adaptivity of the system is urgently required. Hence all of the submethods must be able to adapt themselves given the current data, which we discuss in Section 6.

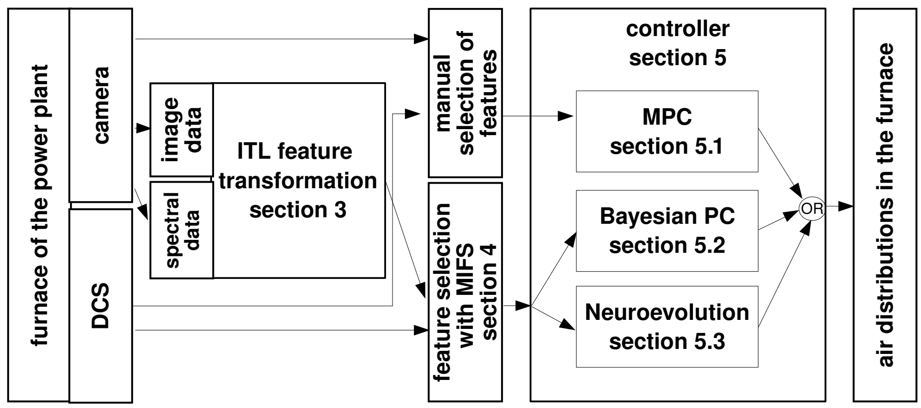

Fig. 1. This figure shows the data flow of a complete control step. The distributed control system (DCS) provides different measurements like coal feed rates or settings of the coal mill, while cameras are used to directly observe the process. Beside images from the burners, the cameras provide spectral information about the flames. Both are transformed by an Information Theoretic Learning (ITL) approach that will reduce dimensionality while maximizing the kept information about the targets. The results as well as the sensor data from the DCS are subject to a feature selection. All selected features are the inputs for two different control strategies, Bayesian Process Control (BPC) and Neuroevolutionary Process Control (NEPC). The Model Predictive Control (MPC) is fed with handselected features chosen by an human expert and is included as a reference. The selected outputs are fed back into the control system of the power plant using the DCS.

3 Image Data Processing

Using the CMOS cameras, images from the flames inside the furnace are captured. For transforming these images into information channels useful for the controller, an adaptive subspace transformation scheme based on [3] is applied. This Maximal Mutual Information approach is built upon the *Information-Theoretic Learning* (ITL) framework introduced by Principe [4]. The idea is to find a transformation that maximizes the *Quadratic Mutual Information* (QMI) between the transformed data in a certain subspace and the desired target values.

The basic adaptation loop for the optimization process is shown in Fig. 2. The original input data sample x_i is transformed by some transformation g (linear, neural network, ...) with the free parameters w into a lower dimensional space. The transformed data is denoted by y_i. The goal is to find those transformation parameters w that confer the most information into the lower dimensional space with respect to the controller's target variables.

The update rule for the parameters of the transformation is given by the following equation, where α denotes the learning rate

$$w_{t+1} = w_t + \alpha \frac{\partial I}{\partial w} = w_t + \alpha \sum_{i=1}^{N} \frac{\partial I}{\partial y_i} \frac{\partial y_i}{\partial w}. \tag{1}$$

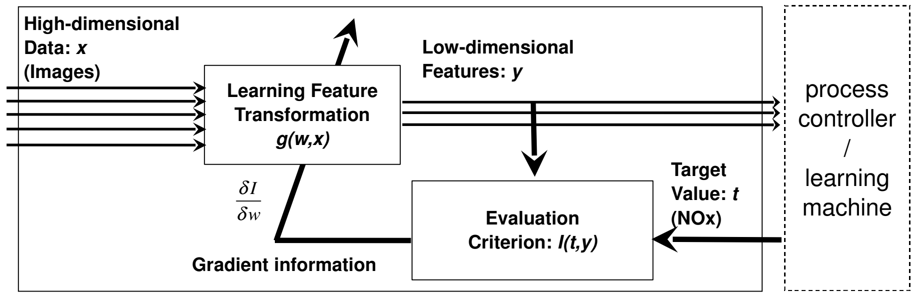

Fig. 2. The original image data is transformed into a lower-dimensional space. An evaluation criterion measures the correspondence to the desired target value (in this case NOx). From this criterion, a gradient information is derived and used to adapt the transformation parameters w.

Finding the gradient $\partial I/\partial w$ can be split into in the sample wise computation of the information forces $\partial I/\partial y_i$ and the adaption of the parameters $\partial y_i/\partial w$. The second term is dependent on the used transformation, and is quite simple for the linear case, where the transformation is computed as:

$$y_i = W^T x_i. \tag{2}$$

For the derivation of the update rule of the transformation parameters w, especially for the information forces $\partial I/\partial y_i$, we refer to details presented in [3].

Due to the intrinsic problems of learning in high dimensional input spaces like images, we compress the images to a subspace of only three dimensions, which was shown to achieve stable results. Furthermore, it was shown in [5] that this kind of transformation is superior to classic approaches like PCA for function approximation tasks in the power plant domain.

The camera system in use allows high frequency analysis of grey values, too. Image patches are observed in the frequency domain of up to 230 Hz. Using the same feature extraction technique, we were able to identify different parts of the temporal spectra that are relevant for targets like NOx.

All these new features generated from images and spectral data are sent to the feature selection module, which is described subsequently.

4 Feature and Action Selection

All features generated from images and those coming from other sensors (e.g. coal feed rates) as part of the DCS are subject to a feature selection step. The main reason for this module is the required reduction of the input space, and there are several channels that are not important or mostly redundant, which can be eliminated. A smaller input space allows an easier adaptation and generalization of the controller. For this step, we employ the MIFS algorithm [6], which is a simple approximation of the Joint Mutual Information between features f and the control targets y. This approximation uses pairwise Mutual Information only, and is hence, easily to compute.

Features f_i that maximize the following term are iterativly added to the subset of used features f_s

$$\arg\max_{f_i}(I(f_i, y) - \beta \sum_{f_s \in S} I(f_i, f_s)). \tag{3}$$

The first term describes the information between the feature and the target value and represents the relevance, while the second term represents the information between the feature and all features already in use. This represents the redundancy. The parameter β is used to balance the goals of maximizing relevance to the target value and minimize the redundancy in the subset.

With the help of this selection technique the number of inputs is trimmed down from more than 50 to approximately 20-30 features used by the controller as problem-relevant input.

The same methods were applied for action selection to reduce the number of manipulated control variables (mainly different air flows settings), but unfortunately all manipulated variables proved to be important in the 12-dimensional action space.

5 Control Strategy

In the following, we describe how to bridge the gap between the selected informative process describing features at one hand and the selected effective actions on the other hand by an adequate sensory-motor mapping realized by the controller. Practically, we tested three different control approaches.

5.1 Model Based Predictive Control

Model Based Predictive Control (MPC) is a well known approach for advanced process control of industrial processes [7]. Its basic idea is to build a mathematical model of the process to be controlled. Based on that process model, a control sequence is computed, that drives the process into the desired state within a given prediction horizon. We tested a simplified version of the MPC-family, that uses a feedforward neural network as process model. It operates with inputs defined by a human expert, rather than an automatic feature extraction.

5.2 Bayesian Process Control

Since real world processes usually are partially observable and noisy, we tested an architecture, that was designed to deal with these drawbacks. Probabilistic models [8] explicitly approximate stochastic processes. Practically, our probabilistic control approach first estimates from historical data joint probability density functions of all control, state, and target variables of the industrial process. The resulting probabilistic process model is used afterwards to build a factor graph [9]. During the so called inference process within that factor graph, a control sequence is calculated, that most probably drives the process into the desired target state, which is a minimum of NOx produced and 02 used in the combustion.

5.3 Neuroevolutionary Process Control

In addition to both control approaches described before, we also tested a neuroevolution method called Cooperative Synapse Neuroevolution (CoSyNE) [10]. CoSyNE applies evolutionary strategies to evolve a nonlinear controller based on recurrent neural networks, that drives the process as fast and accurate as possible into the desired target state. Unfortunately, evolutionary algorithms for control require extensive tests of the individual controllers with the plant in order to determine their performance, the so called fitness. Obviously, that extensive testing is not applicable for industrial processes with their real-world constraints like irreversibility and control cycle periods. Thereto, we extended the original approach proposed by [10] by a couple of alternative process models, that replace the real plant for the extensive closed-loop controller tests. Instead of a single process model, we favor a couple of different models, because that diversity ensures the development of a robust controller. On the contrary, operating on a single process model would be suboptimal, since any failure of the model may mislead the evolutionary search process for an optimal controller to a suboptimal solution or even a dead-end.

6 Adaptivity

As mentioned before, the industrial process to be controlled is time-varying for several reasons. In consequence, a suitable control system needs to be adaptive to cope with these changes automatically. Thereto, we introduced an automatic adaptivity of all components of our overall control architecture shown in Fig. 1.

Starting with the feature extraction from high dimensional image or CMOS-spectra data, we automatically retrain the underlying filters based on new data. This recalculation takes place every 6 hours and operates on image and spectral data of the last three days. This way, any change caused by either process variation, fuel change, or even camera function is taken into account automatically.

At the next stage, the feature selection is updated every 12 hours. Thereto, both process data and adaptively extracted features of the last three weeks are taken into consideration. Thus, less informative features or process measurements can be excluded from further calculation in the controller. In this way, malfunctioning sensors or even uninformative features from cameras temporarily blocked by slag will be sorted out from further observation automatically. On the other side, previously unused, but currently relevant features can be included.

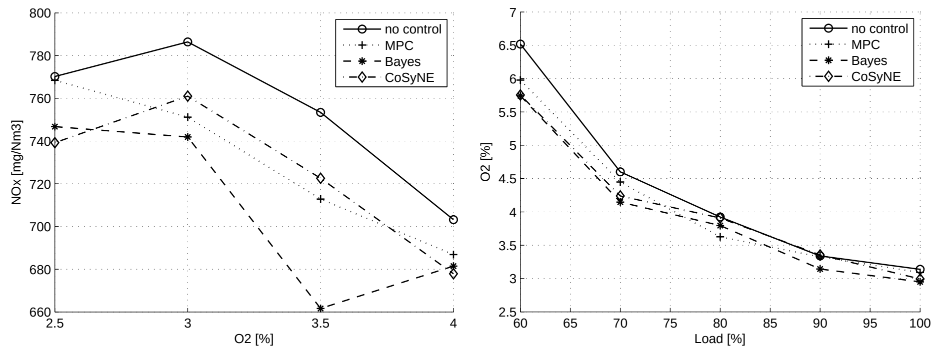

Fig. 3. Left: Comparison of NO_x emissions with different controllers plotted over O_2-concentration in the waste gas. **Right:** Comparison of efficiency factors for different controllers plotted over plant-load. The lower the O_2, the higher the efficiency factor of the plant.

Again every 12 hours or whenever the adaptive feature selection changed, also the corresponding controller is adapted to the new data by a retraining based on most recent process data of the last three weeks. In this way, the corresponding controller is adapted to new process properties of changed input-situations automatically.

7 Results of field trials in a coalfired power plant

The adaptive and self-organizing control architectures described before, are running on a real process, namely a 200MW hard-coal fired power plant in Hamburg, Germany. The task is to minimize the nitrogen emissions (NO_x) and to increase the efficiency factor of that plant. Thereto, the controllers have to adjust the total amount of used air and its distribution between the 6 burners of the boiler. In total, 12 different actuators have to be controlled. Both amount and distribution of coal were given.

For comparison, the three control approaches Model-based Predictive Control (MPC), Bayesian Process Control (Bayes), and Neuroevolutionary Control (CoSyNE) have been tested against normal plant operation without any intelligent air optimization by an AI system (no control, only basic PID control provided by the plant control system). All controllers have been tested sequentially from April, 9th until April, 20th in 2009 during normal daily operation of the plant including many load changes.

Figure 3 shows, that the Model-based Predictive Controller (MPC), operating with manually selected inputs, can outperform the operation without any air control for both targets NO_x and O_2. The Probabilistic Controller (Bayes) operating on automatically extracted and selected features performs even better. The same holds true for the evolutionary Neuro-Controller (CoSyNE), which as well as the Bayesian controller operates on automatically extracted and selected inputs.

Obviously, these self-organizing and adaptive control architectures, namely the Bayesian and the Neuroevolutionary Controller (at least for the more important overall O_2 value), are able to control complex industrial processes. Furthermore, their superior performance is also stable during permanent automatic adaptation of feature extraction, feature selection, and controller retraining using recent process data. In this way, these control architectures are able to cope with a large spectrum of typical real-world-problems ranging from sensor drift to time-varying processes.

8 Conclusion and Future Work

In this paper, an adaptive system for controlling a dynamic industrial combustion process is presented. Despite the potentially difficult handling of adaptive systems, we were able to show a superior performance compared to classical control approaches. An implementation of the methods presented here is successfully operating a real coal power plant.

Over time, the application of this system will result in considerable savings in coal consumption and emissions of greenhouse gases. Hence the application of Machine Learning techniques in this field have high economical and ecological impact.

References

1. Flynn, D. (Ed.): Thermal power plant simulation and control. *IEE London* 2003
2. Stephan, V., Debes, K., Gross, H.-M., Wintrich, F., Wintrich, H. A New Control Scheme for Combustion Processes using Reinforcement Learning based on Neural Networks. *International Journal on Computational Intelligence and Applications* **1** (2001) 121–136
3. Torkkola, K.: Feature Extraction by Non Parametric Mutual Information Maximization. *Journal of Machine Learning Research* **3** (2003) 1415–1438
4. Principe, J., Xu, D., and Fisher, J.: Information theoretic learning. In Haykin,S., editor, *Unsupervised Adaptive Filtering*. Wiley (2000) 265–319
5. Schaffernicht, E., Stephan, V., Gross, H.-M. Adaptive Feature Transformation for Image Data from Non-stationary Processes. *Int. Conf. on Artificial Neural Networks (ICANN)* 2009 To appear.
6. Battiti, R.: Using mutual information for selecting features in supervised neural net learning. *IEEE Transactions on Neural Networks* **5** (1994) 537–550.
7. Ogunnaike, B.A., and Ray, W.H. Process Dynamics, Modeling and Control. *Oxford University Press* (1994)
8. Bishop, C. Pattern Recognition and Machine Learning. *Springer* 2006.
9. Kschischang, F. R., Frey, B. J., Loeliger, H.A. (2001). Factor Graphs and the Sum-Product Algorithm. *IEEE Transactions on Information Theory* **47** (2) 498–519.
10. F. Gomez, J. Schmidhuber, and R. Miikkulainen. Efficient non-linear control through neuroevolution. In *European Conference on Machine Learning (ECML)*, 2006 654–662