2017 25th Mediterranean Conference on Control and Automation (MED)
July 3-6, 2017. Valletta, Malta

# Data-driven Modelling, Learning and Stochastic Predictive Control for the Steel Industry*

Domagoj Herceg[a], George Georgoulas[b], Pantelis Sopasakis[c,*], Miguel Castaño[b], Panagiotis Patrinos[c], Alberto Bemporad[a], Jan Niemi[d] and George Nikolakopoulos[b]

***Abstract*— The steel industry involves energy-intensive processes such as combustion processes whose accurate modelling via first principles is both challenging and unlikely to lead to accurate models let alone cast time-varying dynamics and describe the inevitable wear and tear. In this paper we address the main objective which is the reduction of energy consumption and emissions along with the enhancement of the autonomy of the controlled process by online modelling and uncertainty-aware predictive control. We propose a risk-sensitive model selection procedure which makes use of the modern theory of risk measures and obtain dynamical models using process data from our experimental setting: a walking beam furnace at Swerea MEFOS. We use a scenario-based model predictive controller to track given temperature references at the three heating zones of the furnace and we train a classifier which predicts possible drops in the excess of Oxygen in each heating zone below acceptable levels. This information is then used to recalibrate the controller in order to maintain a high quality of combustion, therefore, higher thermal efficiency and lower emissions.**

***Index Terms*— Advanced Process Control; Machine Learning; Stochastic Model Predictive Control; Risk-sensitive Model Selection; Cyber-Physical Systems.**

## I. Introduction

The process industry is undergoing a remarkable transformation led by the proliferation of devices with sensing, communication and computation capabilities which has marked the advent of cyber-physical systems which are ensembles of physical processes and computational agents [1].

In the process industry, one often encounters energy intensive processes which operate in highly uncertain contexts. The main objective is the design of control systems which lead to higher performance and lower consumption of resources, lower emissions, resilience to the inevitable underlying uncertainty and, as a result, high autonomy. It has nowadays become evident that the torrents of data which are produced by industrial processes can be exploited to reach these objectives.

[a] IMT School for Advanced Studies Lucca, Piazza S. Francesco 19, 55100 Lucca, Italy.

[b] KU Leuven, Department of Electrical Engineering (ESAT), STADIUS Center for Dynamical Systems, Signal Processing and Data Analytics, Kasteelpark Arenberg 10, 3001 Leuven, Belgium.

[c] Luleå University of Technology, Control Engineering Group SE–971 87 Luleå, Sweden.

[d] Swerea MEFOS, Box 812, SE–971 25 Luleå, Sweden.

* This work was supported by the EU-funded H2020 research project DISIRE "Integrated process control based on in-situ sensors into raw material and energy feedstock," grant agreement No. 636834. The work of P. Patrinos was supported by KU Leuven Research Council BOF/STG-15-043.

* Corresponding author – pantelis.sopasakis@kuleuven.be.

Model predictive control (MPC) has become the golden standard in the process industry for its ability to account for multiple-input multiple-output dynamics, problem objectives and constraints [2]. MPC consists in solving an optimal control problem along a prediction horizon minimising a performance index subject to the system dynamics and constraints; this yields an optimal sequence of control actions the first one of which is applied to the system [3]. Often, uncertainty is accounted for in MPC in a worst-case fashion and the performance index is minimised for the worst-case realisation of the uncertainty leading, naturally, to conservative behaviour and poor performance. Advanced variants of MPC, such as its stochastic variants, have been proven powerful tools in dealing with uncertain systems [4]–[6].

In this paper we propose an integrated process control scheme using a walking beam furnace of MEFOS Swerea — a common process in the steel industry — as our case study. To the best of our knowledge, this is the only large-scale experimental furnace in Europe. A rough sketch of our scheme is depicted in Figure 2. We explore several control-oriented data-driven modelling approaches which can be applied adaptively to update the obtained dynamical models and detect changes in the dynamic behaviour of the system. We propose a novel risk-based model selection methodology which is suitable for MPC.

## II. Heat Exchange Dynamics & Data-Driven Modelling

### A. Walking Beam Furnace

A *walking beam furnace* is a furnace where steel slabs are walked through several heating zones until they reach a desired temperature. Heat is provided in each of the zones via combustion of oil. The temperature increases as the slabs move from zone 1 towards zone 3, so heat flows inside the furnace upstream. Apart from heat losses through the wall of the furnace, heat flows to the ambient by the charge and discharge gates in zones 1 and 3 respectively which open and close at regular intervals to load and unload the slabs (or, sometimes, for manual inspection). Heat flow is further influenced by the combustion air flow; a stream of air convecting heat upstream. A schematic overview of the walking beam furnace of MEFOS is shown in Figure 1.

The walking beam furnace is designed to heat the steel products slabs (large steel beams) to a specific temperature prior to the rolling process. Depending on the product the discharge temperature varies, typically between $1100°C$ to $1300°C$, with an acceptable discharge temperature window

978-1-5090-4533-4/17/$31.00 ©2017 IEEE

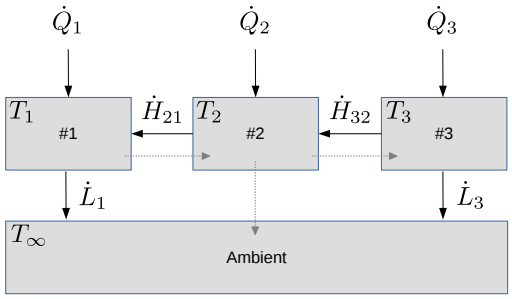


**Fig. 1:** Schematic of the walking beam furnace of MEFOS comprising three heating zones. Heat is provided to each zone with rate $\dot{Q}_i$ and, because of temperature difference and convection, there is a heat flow $\dot{H}_{21}$ from zone 2 to zone 1 and $\dot{H}_{32}$ from zone 3 to zone 2. Heat escapes zones 1 and 3 at rates $\dot{L}_1$ and $\dot{L}_3$ mainly because of the input and output doors.

of a few degrees. In the furnace the slabs travel through by using a combination of revolving and stationary beams. During this passage, the slabs are directly exposed to the heat produced by burners located inside the furnace. Typical residence time inside the furnace is up to a few hours.

Highly turbulent flow profiles and heat transfer by conduction, convection and radiation call for accurate modelling techniques involving partial differential equations [7] which are not suitable for controller design purposes. Traditionally, simple models are used and such plants are controlled by simple PID-type controllers [8]. It has, however, nowadays become evident that MPC can afford industrial furnaces high performance and reliability [9], [10].

The walking beam furnace is equipped with temperature sensors, two at each combustion zone ( one at the ceiling and one on the wall) and Oxygen sensors at each zone; these define the output variables of the system. The manipulated variables are the supply rates of fuel (atomised oil) and combustion air.

### *B. Risk-sensitive model selection*

Models, regardless of their form and structure, need to be appraised for their predictive ability. For this purpose we introduced and used a new performance indicator which we present in this section. At time $k$, using past information (past observations of the system state and input values) and given the current control action $u_k$, the model predicts a value $\hat{y}_{k+1|k}$. Recursively, and given a sequence of future (predicted) control actions $u_{k+1}, u_{k+2}, \ldots, u_{k+N-1}$, the model is used to predict $\hat{y}_{k+i|k}$ for $i = 2, \ldots, N$. Once the actual output realisations have occurred, that is, at the end of the horizon, we may compute the errors

$$\epsilon_{k+j|k} = \hat{y}_{k+j|k} - y_{k+j} \tag{1}$$

Thus, at every $k$ we may create a vector $\epsilon_k := (\epsilon_{k+1|k}, \ldots, \epsilon_{k+N|k}) \in \mathbb{R}^N$ which is, essentially, a random variable. We quantify the predictive ability of the model *at time $k$* using the *predicted mean square error* defined as

$$\mathrm{PMSE}_k = \frac{1}{N} \sum_{j=1}^{N} \|\epsilon_{k+j|k}\|^2. \tag{2}$$

We then define the *predicted root mean square error at time $k$* which we denote by $\mathrm{PRMSE}_k$ as the square root of $\mathrm{PMSE}_k$,

$$\mathrm{PRMSE}_k = \sqrt{\mathrm{PMSE}_k}. \tag{3}$$

This is a random variable. It should be clear that performing a single-shot prediction and computing $\mathrm{PRMSE}_k$ at a particular $k$ is not indicative of the model's predictive ability at other time instants $k$. Computing the expected value of $\mathrm{PRMSE}_k$ is again not really indicative because we will be disregarding all other higher-order moments of $\mathrm{PRMSE}_k$.

In order to extract a meaningful characteristic value out of the random variable $\mathrm{PRMSE}_k$ we need to employ an appropriate *risk measure*. In particular, since we need to minimise the prediction error – rather than to maximise it – we are looking for a *convex* risk measure.

The theory of risk measures has been well established the last decade and certain well-posedness axioms have been postulated. A risk measure is called *coherent* if it satisfies the four *coherency axioms* which can be found in [11]. The most popular risk measure which enjoys a series of favourable properties is the *average value-at-risk* of an (integrable) random variable $Z$ which is defined as

$$\mathrm{AV@R}_\alpha[Z] = \inf_{t \in \mathbb{R}} \{t + \alpha^{-1} \mathbb{E}[Z - t]_+\}, \tag{4}$$

where $[X]_+ = \max\{0, X\}$ and $\alpha$ is the significance level at which $\mathrm{AV@R}_\alpha$ is estimated.

In order to understand the meaning and practical significance of this operator, we need first to define the *value-at-risk* of a real-valued random variable $Z$ at significance level $\alpha$ which is $\mathrm{V@R}_\alpha[Z] = \inf\{z : \mathrm{P}[Z \leq z] \geq \alpha\}$. It can be easily shown that $\mathrm{AV@R}_\alpha[Z]$ is the expectation of $Z$ conditioned by $Z \geq \mathrm{V@R}_\alpha[Z]$, i.e.,

$$\mathrm{AV@R}_\alpha[Z] = \mathbb{E}[Z \mid Z \geq \mathrm{V@R}_\alpha[Z]]. \tag{5}$$

We need to highlight that the minimisation of $\mathrm{AV@R}_\alpha[Z]$ leads to different choices than the minimisation of $\mathbb{E}[Z]$. It should be neither assumed that $\mathrm{AV@R}_\alpha[Z_1] \leq \mathrm{AV@R}_\alpha[Z_2]$ implies $\mathbb{E}[Z_1] \leq \mathbb{E}[Z_2]$, nor the converse. Following (5), $\mathrm{AVAR}_0[Z] = \sup[Z]$ and $\mathrm{AVAR}_1[Z] = \mathbb{E}[Z]$ — therefore, $\mathrm{AV@R}_\alpha$ interpolates between sup (the worst-case operator) and $\mathbb{E}$ (the expectation).

In our analysis, in order to select a reliable dynamical model, we assess its predictive ability on a set of data which has not been used for training; this we call the *test dataset* and thereon we compute $\mathrm{AV@R}_\alpha[\mathrm{PRMSE}_k]$. Note that the models we produce may not be suitable for long-term open-loop predictions, but they are suitable for control applications where a predictive ability of a few time instants is required, notwithstanding, but the availability of feedback obviates the need for very accurate models.

### *C. Heat transfer dynamics*

After an experimental campaign at MEFOS, a dataset was obtained with measurements of temperature, oxygen, combustion air flow, fuel supply rate and gate status in intervals of 10s; a dataset of 35000 measurements was thus compiled. This was split into a training set of 25000

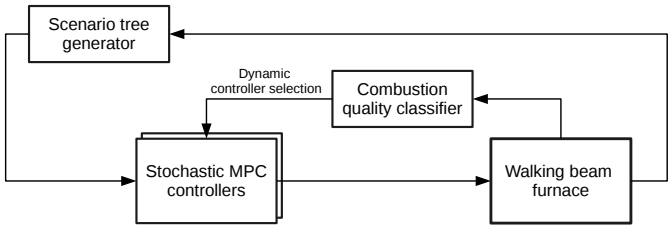


**Fig. 2:** Data-driven control concept for the industrial process of the walking beam furnace.

measurements and a test set of 10000 measurements. Models were tested on the test set for their predictive ability using the indexes presented in the previous section.

Hereafter we present some modelling techniques we used along with validation results.

*1) ARX model description:* Auto-regressive model with exogenous inputs (ARX) provide a simple and efficient modelling framework. We assume that the temperatures $y(k) = (T_1(k), T_2(k), T_3(k))$ at the three zones of the furnace are predicted by the model

$$\hat{y}(k) = \sum_{i=1}^{N_y} \sum_{j=1}^{N_y^i} a_j^i y^i(k-j) + \sum_{i=1}^{N_u} \sum_{j=1}^{N_u^i} b_j^i u^i(k-j), \quad (6)$$

that is, they are given as a linear combination of past temperatures, and past values of the fuel supply rate, combustion air supply and door statuses (which are packed in the vector $u$). History lengths are selected as shown in Tables I and II, and the results in Table III.

| | $T_1$ | $T_2$ | $T_3$ |
|---|---|---|---|
| $T_1$ | $q$ | $l$ | |
| $T_2$ | $l$ | $q$ | $l$ |
| $T_3$ | | $l$ | $q$ |

**TABLE I:** Output-output correspondence of variables and the lags. We use the parametrisation $q = \lfloor 1.67l \rfloor$.

| | $u_1$ | $u_2$ | $u_3$ | $u_4$ | $u_5$ | $u_6$ | $u_7$ | $u_8$ |
|---|---|---|---|---|---|---|---|---|
| $T_1$ | $l$ | $l$ | | | | | $l$ | |
| $T_2$ | | | $l$ | $l$ | | | | |
| $T_3$ | | | | | $l$ | $l$ | | $l$ |

**TABLE II:** Output-input correspondence of variables and the lags used for each variables. *Input varialbes.* $u_1, u_3, u_5$: Oil consumption rate at zones 1, 2 and 3 respectively. $u_2, u_4, u_6$: Combustion air flow at zones 1, 2 and 3 respectively. $u_7$: load hatch-door, and $u_8$: unload hatch-door.

Equation (6) can be concisely written as

$$\hat{y}(k) = \phi(k)^\top p, \quad (7)$$

where $\phi(k)$ is called a vector of regressors at time $k$ and is a collection of $y^i(k-j)$ and $u^i(k-j)$, while $p$ is the vector of corresponding model coefficients. For convenience we will define the following

$$A = \left[ \phi(k_0) \ \phi(k_0+1) \ \cdots \ \phi(k_0+D) \right]^\top, \quad (8)$$

here $D$ is number of data points in our training data set. The standard approach to system identification consists in determining $p$ by solving the least-squares problem [12],

$$\text{minimize}_p \ \tfrac{1}{2}\|Ap - y\|_2^2. \quad (9)$$

| $q$ | $l$ | $\text{AV@R}_{0.05}$ | $\text{AV@R}_{0.1}$ |
|---|---|---|---|
| 8 | 5 | 8.7980 | 6.8428 |
| 20 | 12 | 8.8208 | 6.8744 |
| 33 | 20 | 8.8798 | 6.9127 |
| 50 | 30 | 8.9429 | 6.9885 |

**TABLE III:** Summary of system identification results for $K = 10$ steps-ahead prediction.

*2) Sparse models — LASSO and Elastic net:* In this section will identify sparse models of the underlying physical process. We aim at finding *simple* models that still explain the data with sufficient accuracy. By doing so we hope to obtain a model that will have superior generalisation properties, based on the reasoning that simple models do not capture noise. A well-known approach is to solve a $\ell_1$-regularised least squares problem, where penalisation is imposed on models parameters. This technique is otherwise known as Least Absolute Shrinkage and Selection Operator (LASSO) [13] and enables us to select model structure based on the data only. Furthermore, results of LASSO can help us to better interpret the model.

The regularised problem is

$$\text{minimise} \ \tfrac{1}{2}\|Ap - y\|_2^2 + \lambda\|p\|_1. \quad (10)$$

Parameter $\lambda$ enforces sparsity in the model — higher $\lambda$ will tend to force more values of $p$ towards zeros and is usually chosen by cross-validation. Once we have solved the LASSO problem to identify the model structure (the nonzeros of $p$), we solve a standard least-squares problem for the reduced-order system known as *debiasing*.

Similar to LASSO, elastic net solves the optimisation problem which involves both $\ell_1$ and $\ell_2$ penalty terms

$$\text{minimise} \ \tfrac{1}{2}\|Ap - y\|_2^2 + \lambda\|p\|_1 + \frac{\lambda^2}{2}\|p\|_2^2, \quad (11)$$

and it overcomes certain limitations of LASSO such as its sensitivity to highly correlated variables.

Results are reported in Table IV. It can be seen that Elastic net method produces models that give lower $\text{AV@R}_\alpha$ score on the test data. It is worth pointing out that LASSO and elastic net models can be trained online as new data arrive using streaming methodologies such as [14].

## III. Uncertainty propagation

We previously presented various predictive dynamical models and proposed a method to evaluate their predictive ability. When, however, it comes to using these models to make predictions it is expedient to know, not only how good

| $\lambda$ | Method | LASSO | Elastic net |
|---|---|---|---|
| | 30 | 8.9751 | - |
| | 60 | 9.3528 | 7.9545 |
| | 90 | 8.9414 | 8.3693 |
| | 120 | 8.2951 | 8.2951 |
| | 150 | 9.1920 | 8.4822 |

**TABLE IV:** Summary of model selection algorithms. Performance is assessed by $\text{AV@R}_{0.1}[\text{PRMSE}]$. In all of the models history of $l = 80$ was used with different $\lambda$.

we expect the prediction to be, but what are some probable realisations of the future evolution of the predicted trajectory. Such a collection of future state trajectories is known as a *scenario fan* [15].

Informative scenario fans can be directly constructed from historical data — possibly the entire history of available data. Such a voluble representation of the uncertainty is, however, unnecessarily complicated and will lead to the formulation of huge scale optimisation problem in the context of stochastic MPC. Scenario fans may be *compressed* into scenario trees — discrete filtered probability distribution such as the one shown in Figure 3 — which can be constructed following [16].

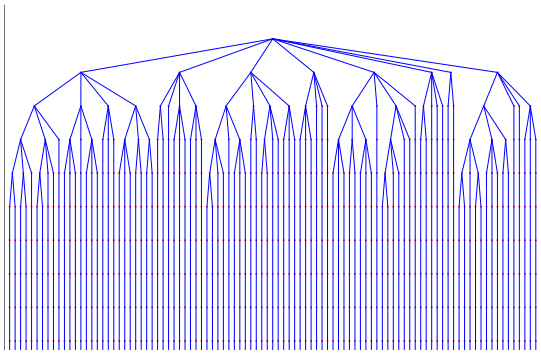

**Fig. 3:** Scenario tree structure

## IV. Combustion quality models

Combustion quality can be estimated by the amount of oxygen in each zone of the furnace. However, the construction of dynamical models for the concentration of Oxygen in each zone is hindered by the highly fluctuating and noisy available measurements. Instead, here we opt for a predictive model for the expected quality of combustion — a nominal (class) variable.

An estimate of that measure, can be given by averaging the oxygen concentration over a predefined Horizon $H$. Based on expert knowledge, instead of providing a continues measurement, only predefined intervals matter. Therefore the oxygen state prediction can be cast as a classification problem, involving three classes:

- **Class 1:** (Poor combustion) if the $H$-step-ahead predicted average $[O_2]$ is in [0; 1.5]
- **Class 2:** (Adequate combustion) if the $H$-step-ahead predicted average $[O_2]$ is in (1.5; 5.0]
- **Class 3:** (Excess of Oxygen) Otherwise.

The developed predictive model is based on current and previous values of the measured variables described in Section II-C. As in any other machine learning and data driven approach the problem of the "curse of dimensionality" is present, which can make the learning/training procedure problematic. Therefore it is common to have a feature/attribute selection stage, to try to pick the most informative/relevant features.

On the other hand feature selection is not an easy task. It is a search problem to find a subset of $l$ features from an original set of $d$ features, such that $l < d$, which achieves the maximum targeted outcome. Therefore, the "goodness" of a particular feature subset is evaluated by using an objective function, $J(Y_m)$, where $Y_m$ is a feature subset of size $m$. This kind of objective function can be defined primarily using either filters or wrappers [17] . One one hand, filters rate the "goodness" of features without employing any learning algorithm, while one the other hand, wrappers using the performance of a learning algorithm as a measure of the "goodness" of a feature (sub)set. A quite powerful hybrid scheme combines filter ranking with wrapper subset evaluation [18], adding features sequentially based on their ranking, forming nested subsets:

$$S_1 = \{f_{i_1}\},\ S_2 = \{f_{i_1}, f_{i_2}\}, \ldots,\ S_n = \{f_{i_1}, f_{i_2}, \ldots, f_{i_n}\}$$

and selecting the one that achieves the best results once it is fed to the predictive model.

In this work, due to the imbalance nature of the data set, the ranking was based using the Area Under the Receiver Operating Characteristic (ROC) curve (AUC) [19], which is immune to the problem of class imbalance. AUC is given by:

$$\text{AUC} = \frac{1}{mn} \sum_{i=1}^{m} \sum_{j=1}^{n} I(r_i^-, r_j^+) \tag{12}$$

where

$$I(r_i^-, r_j^+) = \begin{cases} 1, & \text{if } r_i^- \geq r_j^+ \\ 0, & \text{if } r_i^- = r_j^+ \\ 3, & \text{if } r_i^- \leq r_j^+ \end{cases} \tag{13}$$

where $m$ is the number of the negative cases $r^-$, $r_i^-$ is the value (of the feature) of the $i$-th negative case and $n$ is the number of positive cases $r^+$ and $r_j^+$ the value (of the feature) of the $j$-th positive case. For multi-class problems AUC is given as the average of all class pairs.

The predictive model involved both in the wrapper stage as well as the final oxygen prediction is the Nearest Neighbour (NN) classifier, which as its name implies, given a labeled training set, assigns unseen examples to the class of its nearest neighbour. A set of initial experiments revealed that the imbalanced nature of the dataset made very difficult the training of the predictive model (the under-sampled class 1 was practically never predicted correctly). In order to combat the imbalanced nature of the data set, the Synthetic Minority Oversampling TEchnique (SMOTE) was employed [20],

which creates stochastically synthetic instances for the minority class along the lines connecting an instance of the minority class with its k nearest neighbours.

The results of the proposed approach for the three different zones are depicted in the following confusion matrices, given in Table V, which are translated into overall accuracies per zones 1, 2 and 3, equal to 96.33%, 94.65% and 88.57% respectively.

| CM Zone 1 | | Predicted | | |
|---|---|---|---|---|
| | | Class 1 | Class 2 | Class 3 |
| | Class 1 | **350** | 16 | 0 |
| True | Class 2 | 13 | **1908** | 199 |
| | Class 3 | 0 | 138 | **7347** |
| CM Zone 2 | | Predicted | | |
| | | Class 1 | Class 2 | Class 3 |
| | Class 1 | **564** | 8 | 0 |
| True | Class 2 | 10 | **2100** | 77 |
| | Class 3 | 0 | 64 | **148** |
| CM Zone 3 | | Predicted | | |
| | | Class 1 | Class 2 | Class 3 |
| | Class 1 | **42** | 4 | 0 |
| True | Class 2 | 49 | **5410** | 699 |
| | Class 3 | 0 | 388 | **3379** |

**TABLE V:** Confusion matrices (CM) for the combustion quality classifier for the three zones.

## V. Data-driven operating management

### A. Stochastic model predictive control

For the model predictive control problem formulation we identify two main objectives: (i) to retain the temperatures at each zone as close as possible to the desired temperature set-points with higher importance on the tracking error at zones 2 and 3, (ii) to penalise the consumption of fuel at each zone. For this purpose, we use the stage cost function

$$\ell(x, u; x^{sp}) = \tfrac{1}{2}\|x - x^{sp}\|_Q^2 + \tfrac{1}{2}\|u\|_R^2, \qquad (14)$$

where $Q = \mathrm{diag}(Q_1, Q_2, Q_3)$ is a diagonal matrix where $Q_i > 0$ is the tracking error weight for zone $i$ with $Q_1 < Q_2 < Q_3$, $x^{sp}$ is the temperature set-point which is specified by the operator and $R$ is a diagonal matrix with zeros on the diagonal at the positions which don't correspond to the fuel consumption. The notation $\|\cdot\|_Q^2$ is meant as $\|x\|_Q^2 = x^\top Q x$. We will discuss a data-driven procedure to select $Q$ and $R$ in Section V-B.

The stochastic MPC controller needs to take into account the input constraints. Let $u_k = (u_k^1, u_k^2)$ where $u_k^1 \in \mathbb{R}^3$ are the fuel supply rates and $u_k^2 \in \mathbb{R}^3$ are the combustion air flows at each zone. Then

$$\begin{bmatrix}0\\0\\0\end{bmatrix} \le u_k^1 \le \begin{bmatrix}50\\50\\50\end{bmatrix}, \begin{bmatrix}0\\0\\0\end{bmatrix} \le u_k^2 \le \begin{bmatrix}1000\\900\\1000\end{bmatrix}. \qquad (15)$$

The exact times at which load and unload doors are opened is scheduled in advanced and those are fed to the MPC controller as *prescient* inputs. This is because the doors open and close based on a predefined production schedule.

At this point note that the stage cost $\ell$ is a random variable as the future evolution of the system state is uncertain. As customary in MPC we define the following value function to be optimised

$$V = \sum_{k=0}^{N-1} \ell(x_{k+j|k}, u_{k+j|k}; x_{k+j|k}^{sp}), \qquad (16)$$

where $x_{k+j|k}$ is the future state of the system, at time $k+j$ as predicted at time $k$ and $u_{k+j|k}$ are decisions which are to be taken at time $k+j$ using the information that will be available to the controller at that time (that is $x_{k|k}$, $x_{k+1|k}, \ldots, x_{k+j|k}$.) The total cost $V$ is again a random variable. In stochastic MPC we minimise $V$ over the space of all possible *causal control function* $u(x)$ instead of minimising over *sequences of control actions*. Here, for every possible future realisation of $x_{k+j|k}$ we make a decision $u_{k+j|k}$. This control scheme is known as *scenario-based model predictive control* and the interested reader can find more information in [15], [21]. The stochastic MPC problem now becomes

$$\text{minimise}_{\mathbf{u}}\, \mathbb{E}V \qquad (17)$$

subject to the system dynamics, constraints (15) and $x_{k|k} = x_k$.

### B. Data-driven tuning

The tuning parameters $Q$ and $R$ can be adapted on-line from process data so as to strike a good balance between quality of set-point tracking and quality of combustion (availability of oxygen). Besides the oil and the combustion air, another crucial factor for the quality of combustion is the oxygen. If the level of oxygen in the furnace is satisfying, the combustion process can go on, otherwise we will have lower quality combustion which leads to poorer temperature tracking.

To this end we employ a switching scheme for the control of the furnace based on the predicted oxygen levels. If the level of oxygen is predicted to be acceptable, we use a nominal MPC controller. Likewise, if the classifier predicts low oxygen levels, we switch to a more conservative MPC controller that penalises oil expenditure more than the nominal one. This is regulated by increasing the values of entries of $R$ matrix in (14).

We proceed in the following fashion. At each time step oxygen state predictor returns a good (Class 2 or 3) or bad (Class 1) flag for the oxygen levels in each furnace zone. Based on the prediction we adjust the corresponding entries of the matrix $R$. In effect, we will have eight different MPC controllers where, at each time step, a corresponding controller is selected based on the classifications given by the algorithm. Diagonal entries of matrix $Q$ are always set to $Q_{11} = 3$, $Q_{22} = 4$ and $Q_{33} = 5$, whereas, depending on the high or low predicted oxygen content, we have $R_{ii}^H = 10^{-2}$ or $R_{ii}^L = 10^{-3}$ for the diagonal entries of the matrix $R$.

Next, model predictive controller is selected based on the predicted zones oxygen levels with different stage-wise costs due to selection of matrices $Q$ and $R$.

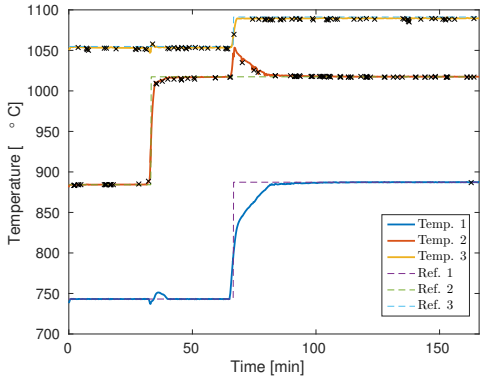


**Fig. 4:** Simulation of the closed-loop operation of the walking beam furnace. Black crosses mark times at which a different controller (with higher $R$ values) is used to save up oxygen in the zone.

### C. Simulations

Here we provide simulations to support the concept presented in the previous section. At each time step we select a controller from the set of eight possible ones based on the classifier prediction regarding the zones oxygen levels at each zone. We then apply the first computed input and repeat the procedure at the next time instant.

In order to demonstrate our concept we have taken a simple ARX model with $l = 5$ and converted it to state space representation. The resulting model has $56$ states, but is also significantly sparse.

Closed-loop simulations are shown in Figure 4 where we employ the proposed control scheme. Most of the time, the combustion quality is high. In particular, in the low-temperature zone 1, there is only one instance of using a high-$R$ controller. The temperature spike around time $70min$ is due to the increase of temperature in zone 3 because heat is transferred from zone 3 to zone 2. Moreover, it can be observer that the temperature response in zone 1 is less aggressive that the response in zones 2 and 3 and temperature tracking in zone 3 is constant with a negligible fluctuation of a few degrees.

Since dynamical models are not perfect and an estimation in the simulations shown here (cf. Figure 4) the underlying system is taken to be uncertain. This allows us to perform realistic simulations even with imperfect dynamical models by replaying the modelling errors.

## VI. Conclusions

In this paper we have put together three novel tools to address issues related to the control of industrial processes. First, we proposed a robust data-driven model selection procedure based on the average value-at-risk which can be used to automatically and reliably train and select process models. Second, we used a stochastic variant of model predictive control to deal with uncertainty as a result of modelling and measurement errors. Third, we proposed an on-line automatic tuning methodology where the parameters $Q$ and $R$ are allowed to change online based on the predictions of a machine learning model for the expected combustion quality. The proposed data-driven control scheme augurs the advent of industrial processes of higher performance, resilience and autonomy.

Our ongoing work focuses on the utilisation of information from immersible sensors which travel inside the furnace with the slabs and provide temperature measurements from their surface and core. This will lead to the acquisition of further experimental data that will allow us to build higher resolution temperature and oxygen models.

## References

[1] K.-D. Kim and P. R. Kumar, "Cyber–physical systems: A perspective at the centennial," *Proc. IEEE*, vol. 100, pp. 1287–1308, 2012.

[2] M. G. Forbes, R. S. Patwardhan, H. Hamadah, and R. B. Gopaluni, "Model predictive control in industry: Challenges and opportunities," *IFAC-PapersOnLine*, vol. 48, no. 8, pp. 531 – 538, 2015.

[3] J. Rawlings and D. Mayne, *Model predictive control : theory and design*. Madison, Wis. Nob Hill Pub., 2009.

[4] A. Mesbah, "Stochastic model predictive control: An overview and perspectives for future research," *IEEE Control Systems Magazine*, vol. 36, no. 6, pp. 30–44, 2016.

[5] C. Hans, P. Sopasakis, A. Bemporad, J. Raisch, and C. Reincke-Collon, "Scenario-based model predictive operation control of islanded microgrids," in *54 IEEE Conf. Decision and Control*, (Osaka, Japan), Dec 2015.

[6] P. Patrinos, P. Sopasakis, H. Sarmiveis, and A. Bemporad, "Stochastic model predictive control for constrained discrete-time Markovian switching systems," *Automatica*, vol. 50, pp. 2504–2514, October 2014.

[7] E. M. Madsen, "Steeltemp – a program for temperature analysis in steel plants," *Journal of Materials Processing Technology*, vol. 42, pp. 187–195, 1994.

[8] B. Leden, "A control system for fuel optimization of reheating furnaces," *Scandinavian Journal of Metallurgy*, vol. 15, pp. 16–24, 1986.

[9] R. Speets, "Model predictive control on slab reheating furnace - master thesis system control engineering," Master's thesis, Nijmegen, The Netherlands, 2015.

[10] G. van Ditzhuijzen, D. Staalman, and A. Koorn, "Identification and model predictive control of a slab reheating furnace," in *Proceedings of the International Conference on Control Applications*, vol. 1, pp. 361–366 vol.1, 2002.

[11] A. Shapiro, D. Dentcheva, and Ruszczyński, *Lectures on stochastic programming: modeling and theory*. Society for Industrial Mathematics, 2009.

[12] S. Boyd and L. Vandenberghe, *Convex Optimization*. New York, NY, USA: Cambridge University Press, 2004.

[13] T. Hastie, R. Tibshirani, and J. Friedman, *The Elements of Statistical Learning*. Springer Series in Statistics, New York, NY, USA: Springer New York Inc., 2001.

[14] P. Sopasakis, N. Freris, and P. Patrinos, "Accelerated reconstruction of a compressively sampled data stream," in *24th European Signal Processing conference (EUSIPCO)*, (Budapest, Hugary), 2016.

[15] D. Bernardini and A. Bemporad, "Scenario-based model predictive control of stochastic constrained linear systems," in *Proceedings of the 48h IEEE Conference on Decision and Control (CDC) held jointly with 2009 28th Chinese Control Conference*, pp. 6333–6338, Dec 2009.

[16] H. Heitsch and W. Römisch, "Scenario tree modeling for multistage stochastic programs," *Mathematical Programming*, vol. 118, no. 2, pp. 371–406, 2009.

[17] I. Guyon and A. Elisseeff, "An introduction to variable and feature selection," *Journal of machine learning research*, vol. 3, no. Mar, pp. 1157–1182, 2003.

[18] I. Guyon, "Practical feature selection: from correlation to causality," *NATO science for peace and security*, vol. 19, pp. 27–43, 2008.

[19] M. Wasikowski and X.-W. Chen, "Combating the small sample class imbalance problem using feature selection," *IEEE Transactions on knowledge and data engineering*, vol. 22, no. 10, pp. 1388–1400, 2010.

[20] N. V. Chawla, K. W. Bowyer, L. O. Hall, and W. P. Kegelmeyer, "SMOTE: synthetic minority over-sampling technique," *Journal of artificial intelligence research*, vol. 16, pp. 321–357, 2002.

[21] M. Maiworm, T. Bäthge, and R. Findeisen, "Scenario-based model predictive control: Recursive feasibility and stability," *IFAC-PapersOnLine*, vol. 48, no. 8, pp. 50 – 56, 2015.