LUND UNIVERSITY

A Virtual Sensor for Predicting Diesel Engine Emissions from Cylinder Pressure Data

Henningsson, Maria; Tunestål, Per; Johansson, Rolf

Published in:
3rd IFAC Workshop on Engine and Powertrain Control, Simulation and Modeling

DOI:
10.3182/20121023-3-FR-4025.00063

2012

Link to publication

Citation for published version (APA):
Henningsson, M., Tunestål, P., & Johansson, R. (2012). A Virtual Sensor for Predicting Diesel Engine Emissions from Cylinder Pressure Data. In *3rd IFAC Workshop on Engine and Powertrain Control, Simulation and Modeling* (20 ed., Vol. 45, pp. 424-431). (IFAC Proceedings Volumes; Vol. 45, No. 30). https://doi.org/10.3182/20121023-3-FR-4025.00063

Total number of authors:
3

General rights
Unless other specific re-use rights are stated the following general rights apply:
Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.
• Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
• You may not further distribute the material or use it for any profit-making activity or commercial gain
• You may freely distribute the URL identifying the publication in the public portal

Read more about Creative commons licenses: https://creativecommons.org/licenses/

Take down policy
If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

LUND UNIVERSITY

PO Box 117
221 00 Lund
+46 46-222 00 00

Download date: 18. Dec. 2025

A Virtual Sensor for Predicting Diesel Engine Emissions from Cylinder Pressure Data

Maria Henningsson * **Per Tunestål** ** **Rolf Johansson** *

* *Department of Automatic Control, Lund University*
** *Department of Energy Sciences, Lund University*

Abstract: Cylinder pressure sensors provide detailed information on the diesel engine combustion process. This paper presents a method to use cylinder-pressure data for prediction of engine emissions by exploiting data-mining techniques. The proposed method uses principal component analysis to reduce the dimension of the cylinder-pressure data, and a neural network to model the nonlinear relationship between the cylinder pressure and emissions. An algorithm is presented for training the neural network to predict cylinder-individual emissions even though the training data only provides cylinder-averaged target data. The algorithm was applied to an experimental data set from a six-cylinder heavy-duty engine, and it is verified that trends in emissions during transient engine operation are captured successfully by the model.

1. INTRODUCTION

Engine manufacturers are constantly forced by legislators and the market to make engines cleaner and more fuel efficient. To comply with these demands, engines become more complex with more actuators and thus more degrees of freedom to be calibrated. Since the time required for calibration increases exponentially with the number of variables, engine calibration and control design is becoming a bottleneck in the development process. To remain competitive and comply with legislation, engines must make optimal use of their complex and expensive set of actuators. Another issue is that as the emission targets are reduced, a larger part of the total drive-cycle emissions are due to spikes during transients. While it was previously sufficient to calibrate the engine at steady-state, more attention needs to be paid to the transient behavior in the future.

One way to approach the calibration challenge is to add more sensors to the engine. With more on-line sensor information, focus may be moved from off-line calibration to on-line closed-loop control. Instead of detailed pre-calibrated maps, fine-tuning of actuator settings to optimize engine behavior could be performed on-line. To that purpose, sensors that provide information on the high-level engine performance objectives such as emissions and fuel consumption are needed. Today, on-line feedback control of the engine is based mainly on indirect measurements such as air flow or inlet manifold pressure. Sensors for NO_x have been added to engines for SCR control. However, these sensors are normally placed downstream of the turbine where the pressure and temperature is lower, and therefore they only provide low-pass-filtered, cylinder-averaged information. Particulate matter sensors have also recently been announced, targeted at on-board-diagnostics of diesel particulate filter performance but speed and accuracy of these sensors are so far insufficient for on-line feedback control.

A large number of methods to model emissions based on other engine variables have been presented in the literature, see for example (Hirsch et al. [2008], Westlund and Ångström [2009], Aithal [2010], Mrosek et al. [2010], Tschanz et al. [2010], Sequenz and Isermann [2011], Shi et al. [2011], Grahn et al. [2012]) and references therein. The complexity of these models range from extremely detailed computational fluid dynamics models, to very simple empirical input-output maps. In order to use models for on-line feedback control, the computational complexity must be sufficiently low, and the input to the model must be available from sensors with sufficient speed and accuracy.

Cylinder pressure sensors provide information with high temporal resolution on the combustion process in individual cylinders. These sensors have long been common-place in engine research labs, and some announcements indicate that they may eventually be a realistic option also for production engines (Birch [2008], Pudenz [2007], General Motors [2007], Shahroudi [2008]). To evaluate the market potential of cylinder pressure sensors, it is important to understand the potential benefits and limitations in the kind of combustion information they provide. In this paper, we assume that pressure sensor signals are available for each individual cylinder, and investigate the quality of information they can provide on engine emissions.

A method is proposed to use the detailed information on the combustion process provided by the cylinder pressure to predict engine emissions. The idea of using cylinder pressure to predict other engine outputs such as emissions has been exploited in many contexts before, both using empirical data-based models, and physical models. For example, Seykens et al. [2009] proposed a physical model for predicting NO_x and soot where cylinder pressure was used as one out of several inputs. Wilhelmsson et al. [2009]

⋆ This work was supported by the Swedish Energy Agency and Volvo Powertrain Corp.

proposed a physical two-zone model for NO_x prediction based on cylinder pressure data. Traver et al. [1999] defined a set of physical quantities extracted from the cylinder pressure trace (such as ignition delay, combustion duration, peak pressure) and trained a neural network to predict HC, CO, CO_2 and NO_x emissions from these physical quantities. Çebi et al. [2011] also used a set of physical quantities extracted from the cylinder pressure trace, and a parametric empirical model to predict PM emissions.

The approach of this paper is based on data-mining techniques. Cylinder pressure is normally sampled at a rate of several kHz, and efficient information-processing techniques are needed to extract the useful information from this data in a systematic way. Most research publications that use cylinder pressure data focus on a few parameters with a physical interpretation that are derived from the cylinder pressure, such as combustion phasing or IMEP. However, a large number of different such parameters may be defined, many of which may be redundant. Selecting a subset of such parameters to characterize the combustion process becomes somewhat heuristic. Principal component analysis (PCA) is a systematic method to reduce dimensionality of data, and has been used in the combustion engine context for ion-current signals, see (Rizzoni [1997]), and recently for cylinder pressure signals (Stadlbauer et al. [2012], Henningsson et al. [2012]). In (Henningsson et al. [2012]), principal component analysis was suggested as a systematic approach to characterizing the cylinder pressure with a reduced set of variables, and a neural network was introduced to model emissions from this reduced set of variables. In this paper, we extend this method with a modified cost function for training the neural network that is shown to result in realistic cylinder-individual emissions predictions. Also, a larger data set is used to test the method.

The paper is organized as follows. Section 2 introduces the proposed model for dimensionality reduction using principal component analysis, and emissions prediction using a neural network. The algorithm to train the neural network is also presented. Section 3 describes the experimental data used to evaluate the algorithm. Section 4 shows results for using the virtual sensor scheme for cylinder-individual prediction of λ, NO_x, and opacity. Finally, some concluding remarks are given.

2. MODEL

The scheme for predicting emissions from cylinder pressure data consists of two steps. In the first step, the dimension of the cylinder-pressure data is reduced, and in the second step, a nonlinear model is introduced from the reduced-dimension variable to the target variable.

In this section, the method is presented in general terms. In Section 4, where the algorithm is evaluated on a specific data set, the exact input- and output variables are specified. Depending on the application, the method could be used with other sets of inputs and outputs.

We will distinguish between five different sets of variables:

- *Cylinder pressure* p: The measured pressure, high-dimensional data.

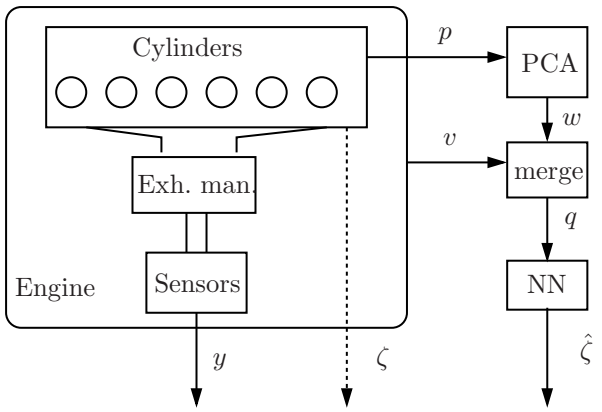

Fig. 1. Overview of of information flow.

- *Reduced-dimension cylinder pressure representation* w: Reduced-dimension representation of p.
- *Input to nonlinear model* q: w combined with other measured engine variables v such as engine speed, injection timing, etc.
- *Target variables* ζ: The target variables for prediction, i.e., the cylinder-individual instantaneous cylinder-out emissions (not available for measurement).
- *Measured variables* y: The measured emissions used for inference of model parameters. These differ from ζ in that they are averaged over cylinders and corrupted by the dynamics in the exhaust manifold and the sensors.

The relationship between the variables is illustrated in Fig. 1.

2.1 Dimensionality reduction from p to w

Principal component analysis (PCA) is used to find a reduced-order representation of p, as was described in (Henningsson et al. [2012]). From a large set of cylinder pressure data collected at a variety of operating conditions, principal component analysis may be used to find an optimized orthonormal set of basis functions $\{\phi\}_{l=1}^{L}$ to approximate the data. For details on PCA, see for example (Bishop [2006]).

Once the optimal basis functions have been determined, each observation p_n of the cylinder pressure $p \in \mathbb{R}^P$ can be approximated by a weighted sum of the basis functions

$$p_n = \phi_0 + \sum_{i=1}^{D} w_{nl}\phi_i + e_n \tag{1}$$

where the approximation error e_n decreases as the number of basis functions L increases.

Because the basis functions $\{\phi\}_{i=1}^{D}$ are orthonormal, the weights $w_n \in \mathbb{R}^D$ can be computed from p_n through a straight-forward matrix multiplication

$$w_n = \Phi(p_n - \phi_0) \tag{2}$$

The weights w_n is a D-dimensional representation of the P-dimensional cylinder pressure data p_n. With a sampling interval of 0.2 CAD, $P = 3600$. In (Henningsson et al. [2012]) it was shown that the approximation error decreases rapidly with increasing D, and that $D < 10$ may be sufficient to approximate the cylinder pressure

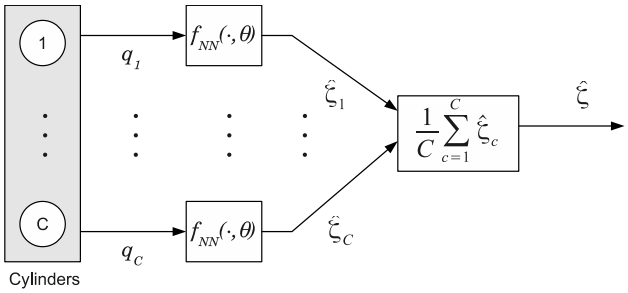

Fig. 2. Structure of neural network for cylinder-individual prediction of emissions. Note that the functions $f_{NN}(\cdot)$ and parameters θ are the same for all cylinders.

N	number of data
K	dimension of ζ_n (number of network outputs)
L	dimension of q_n (number of network inputs)
M	dimension of z_n (number of hidden nodes)
C	number of cylinders

Table 1. Summary of notation for the neural network model.

data. There is consequently a significant dimensionality reduction in going from p to w.

2.2 Neural network model from q to ζ

A neural network is proposed to predict the target variable ζ from the set of reduced variables q. We assume that q is available for each cylinder individually, whereas the training data only provide y measured in the exhaust flow common to all cylinder. We wish to train a neural network model $f_{NN}(q_c, \theta)$ that predicts ζ_c individually for each cylinder by assuming that the network parameters θ are common to all cylinders, and that mixing of the exhaust flow averages emissions over the cylinders,

$$\zeta = \frac{1}{C}\sum_{c=1}^{C}\zeta_c = \frac{1}{C}\sum_{c=1}^{C} f_{NN}(q_c, \theta) \tag{3}$$

The model is illustrated in Fig. 2.

Define:

- ζ_{nk}: output k at time n, measured or estimated from sensor data
- $\hat{\zeta}_{nkc}$: predicted output k at time n for cylinder c
- $\bar{\hat{\zeta}}_{nk}$: predicted output k at time n, averaged over cylinders.

Network model The neural network structure is built from a standard feedforward network with a single hidden layer and sigmoidal basis functions (Hastie et al. [2009]), and the model for averaging over cylinders (3). The notation for indices used in this section is summarized in Table 1.

The neural network model is given by

$$\xi_{nmc} = \alpha_{m0} + \sum_{l=1}^{L}\alpha_{ml}q_{nlc} \tag{4}$$

$$z_{nmc} = \sigma(\xi_{nmc}) \tag{5}$$

$$\bar{z}_{nm} = \frac{1}{C}\sum_{c=1}^{C} z_{nmc} \tag{6}$$

$$\hat{\zeta}_{nkc} = \beta_{k0} + \sum_{m=1}^{M}\beta_{km}z_{nmc} \tag{7}$$

$$\bar{\hat{\zeta}}_{nk} = \frac{1}{C}\sum_{c=1}^{C}\hat{\zeta}_{nkc} = \beta_{k0} + \sum_{m=1}^{M}\beta_{km}\bar{z}_{nm} \tag{8}$$

Here, z represent the hidden layer variables, and α_0, α, β_0, β are the network parameters.

Cost function For training of the network, a cost function R built from a prediction error term R_{pe} and a regularization term R_{reg} is defined

$$R = R_{pe} + \gamma R_{reg}, \tag{9}$$

where γ is constant weight. The prediction error cost is given by

$$R_{pe} = \frac{1}{2}\sum_{n=1}^{N}\sum_{k=1}^{K}(\zeta_{nk} - \bar{\hat{\zeta}}_{nk})^2 \tag{10}$$

and the regularization cost by

$$R_{reg} = \frac{1}{2C}\sum_{n=1}^{N}\sum_{k=1}^{K}\sum_{c=1}^{C}(\hat{\zeta}_{nkc} - \bar{\hat{\zeta}}_{nk})^2. \tag{11}$$

The regularization term is included to avoid overfitting by penalizing the difference between the predicted output of the individual cylinders $\hat{\zeta}_{nkc}$ and the predicted average $\bar{\hat{\zeta}}_{nk}$.

Network training Optimization of R with respect to $\theta = \{\beta_0, \beta, \alpha_0, \alpha\}$ can be performed through a gradient descent method

$$\theta^{(\tau+1)} = \theta^{(\tau)} - \eta^{(\tau)}\frac{dR}{d\theta}(\theta^{(\tau)}). \tag{12}$$

Algebraic expressions for the derivative of R with respect to θ are found in Appendix A. The step size η^{τ} is chosen iteratively such that R is decreased in each iteration τ. The training scheme is summarized in Algorithm 1.

3. EXPERIMENTAL DATA

Experimental data was collected on a six-cylinder heavy-duty Volvo D13 engine. The engine was equipped with a low-pressure EGR loop, a VGT, and cylinder-individual water-cooled pressure sensors. A data set of approximately 30,000 engine cycles was collected in transient operation at three different engine speeds and a range of loads at each speed. The load-speed map is shown in Fig. 3. At each operating point, the fuel injection duration, the injection timing, and the EGR and VGT actuators were manipulated in open loop.

A Siemens VDO/NGK Smart NOx sensor (Siemens VDO [2005]) was used to measure NO_x and λ. An opacimeter was used to measure PM emissions. Piezo-electrical, water-cooled pressure transducers of type Kistler 7061B were used to measure the cylinder pressure.

Algorithm 1 Scheme for training neural network from measured variables q to target variables ζ.

1: Rescale each input q_k to zero mean and unit sample variance.
2: Draw 20 initial estimates $\theta^{(0)}_{1:20} \in \mathcal{N}(0,1)$.
3: Compute $R(\theta)$ for each $\theta^{(0)}$ and select the one that gives the smallest cost.
4: Let $\eta^{(0)} = 1$.
5: **for** $\tau = 0$ to τ_{max} **do**
6: Compute the gradient of R with respect to θ for $\theta = \theta^{(\tau)}$,

$$\nabla R^{(\tau)} = \frac{dR}{d\theta}(\theta^{(\tau)})$$

according to Appendix A.
7: Compute $\theta^* = \theta^{(\tau)} - \eta^{(\tau)} \nabla R^{(\tau)}$
8: **while** $R(\theta^*) > R(\theta^\tau)$ **do**
9: $\eta^{(\tau)} := \eta^{(\tau)}/2$
10: $\theta^* = \theta^{(\tau)} - \eta^{(\tau)} \nabla R^{(\tau)} dR/d\theta(\theta^\tau)$
11: **end while**
12: $\theta^{(\tau+1)} = \theta^*$
13: $\eta^{(\tau+1)} = 2\eta^{(\tau)}$
14: **end for**

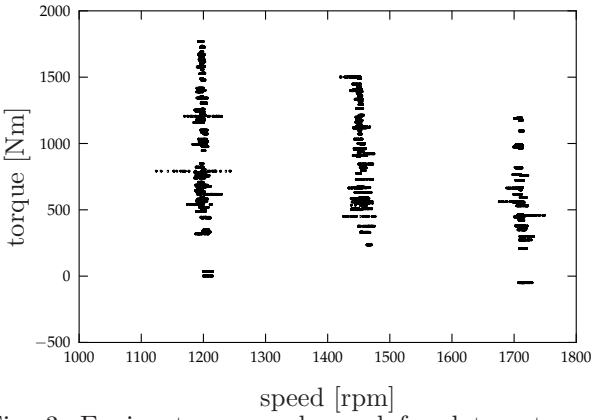

Fig. 3. Engine torque and speed for data set used for network training.

3.1 Pre-processing of the cylinder pressure p

The cylinder pressure sensor offset was computed by simultaneously estimating the polytropic exponent and pressure offset for a range of crank angle degrees in the compression stroke prior to combustion, as described in Tunestål [2009].

3.2 Pre-processing of measured emissions data y

The exhaust flow- and sensor dynamics were modeled as first-order low-pass filters with a time delay, and an optimal non-causal Wiener filter was used to estimate ζ from y to obtain neural network training data. For details of this procedure, see (Henningsson et al. [2012]). The measured output y and Wiener filter estimates $\hat{\zeta}_{wiener}$ for a step increase in load is illustrated in Fig. 4.

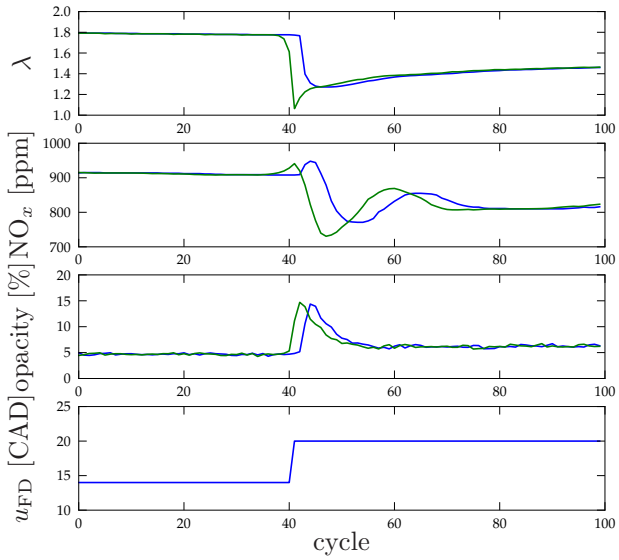

Fig. 4. Measurements y of NO_x, λ, and opacity (blue) when load is increased, and Wiener filter output $\hat{\zeta}_{\text{wiener}}$ (green). Note that the time delay of the measurements has been removed without increasing the visual noise levels.

4. RESULTS

4.1 Network structure

The scheme was tested on prediction of λ, NO_x, and opacity

$$\zeta = (\lambda \ X_{NOx} \ X_{op})^T \tag{13}$$

Opacity was chosen rather than particulate matter mass or particle number, as this signal was available for transient measurements in the laboratory setup. The purpose was to illustrate to what extent the cylinder pressure can provide information on soot emissions.

Input to the network was the first eight principal component weights $w_{1:8}$, as well as injection duration u_{FD}, injection timing u_{SOI}, and engine speed N_{engine}

$$q = (w_{1:8} \ u_{FD} \ u_{SOI} \ N_{engine})^T \tag{14}$$

The choice of the number of principal component weights to include was made considering computational complexity for training the network, and it was also concluded that the higher-order components mostly represent combustion noise. The additional inputs u_{FD}, u_{SOI}, and N_{engine} were chosen because they are available on a cycle-to-cycle, cylinder-individual basis, and they have a direct effect on the in-cylinder combustion process. Other control variables such as EGR valve and VGT actuator positions would not be appropriate since they only have an indirect, dynamic effect on combustion through the gas flows.

The number of hidden nodes was chosen as $M = 10$, and the regularization weight to $\gamma = 0.5$.

From the total data set of 30,000 cycles, a subset was chosen for network training. The subset was selected by

considering the distribution of the target variables ζ. The desired size N_{train} of the training data was determined. For each of the outputs ζ_k, a subset $\mathcal{Z}_k$ of size $2N_{train}$ was selected such that the estimated probability distribution $p(\zeta_i)$ would be approximately uniform in $\mathcal{Z}_i$. The training subset $\mathcal{D}$ was finally selected as N_{train} random samples over the union of $\mathcal{Z}_k$ over the outputs. The purpose was to include more samples from transients with large peaks in emissions.

The network training algorithm was implemented in Python.

4.2 Validation

For network training, 10% of the data was used, and the gradient descent network training algorithm was run for 5,000 iterations (corresponding to approximately 5 minutes on a standard PC). Figure 5 shows neural network prediction $\hat{\zeta}_{NN}$ of the target variables ζ over the entire data set. Note that the targets ζ were obtained through the non-causal Wiener filter from the measured variables y, and should therefore not be interpreted as the 'truth'. The R^2 and root mean square error statistics for prediction over the entire data set were

$$\begin{aligned} R^2_\lambda &= 0.92 \qquad \mathrm{RMSE}_\lambda = 0.074 \\ R^2_{\mathrm{NOx}} &= 0.90 \quad \mathrm{RMSE}_{\mathrm{NOx}} = 74 \text{ ppm} \\ R^2_{\mathrm{op}} &= 0.74 \qquad \mathrm{RMSE}_{\mathrm{op}} = 1.9\% \end{aligned} \tag{15}$$

4.3 Example: EBP transient

To exemplify the performance of the prediction scheme, we look closer at a segment of the data. This data set consists of 1000 consecutive cycles from an experiment at load $\mathrm{IMEP}_n = 7$ bar and speed $N_{engine} = 1700$ rpm. All actuators were kept constant except for the exhaust back pressure (EBP) valve that controls the low-pressure EGR flow. This valve was varied as shown in Fig. 6, which caused λ to vary between approximately 1.4 to 2.1. Also shown in Fig. 6 is the cylinder pressure from cylinder 2 during this transient. In Fig. 8, the cylinder pressure for all six cylinders is shown at cycle 300 and 500 to illustrate the cylinder-to-cylinder variability. During the transient, the principal component weights w computed from the cylinder pressure are the only variables in the neural network input q that vary. We can therefore see how well emissions are predicted using cylinder pressure data only. Figure 7 shows how the principal component weights w_1 to w_8 for the six cylinders vary during the transient.

Figure 9 shows the neural network predictions. The predictions are fairly good; notably the trends are captured rather well. Cylinder-to-cylinder variations are significant. However, the differences appear to be systematic; as an example we can see that cylinder 3 (red) appears to have slightly higher λ compared to the other cylinders and at the same time higher NO_x and lower opacity. This is in line with what would be expected from physical insight.

In some parts of the data, the cycle-to-cycle noise is significant. There are a number of possible explanations for this,

- The number of nodes in the neural network was too small to provide good models at all operating points.

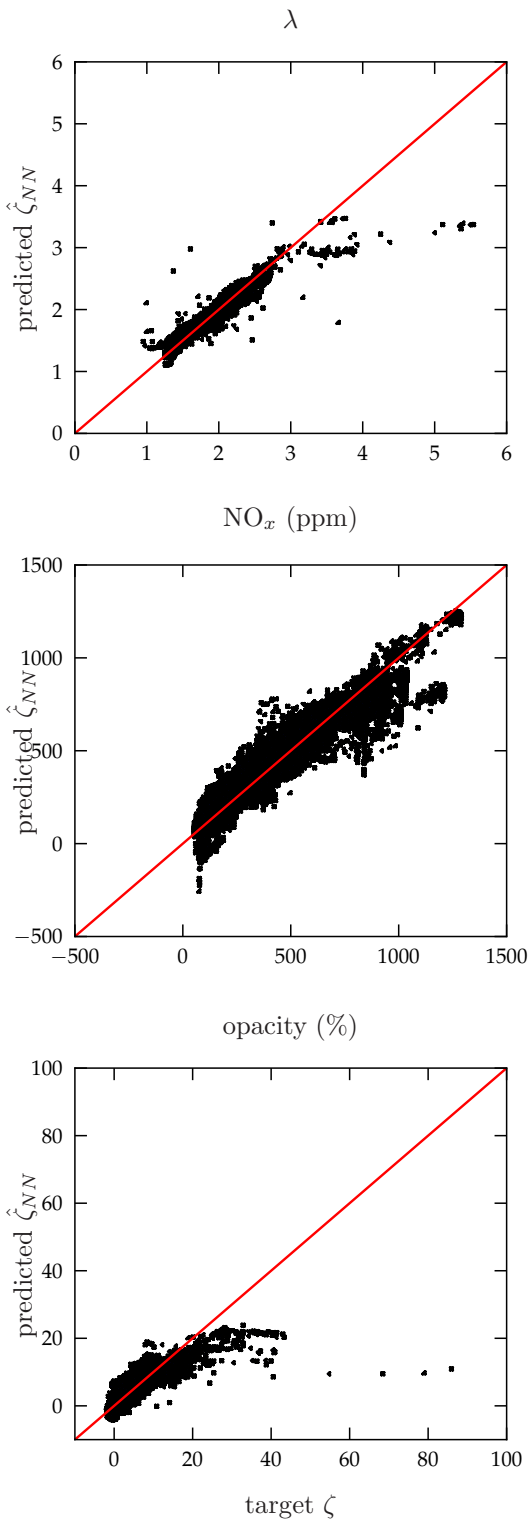

Fig. 5. Predicted outputs vs. target outputs.

- The gradient descent algorithm converged to a local minimum.
- The training data did not provide sufficient information to train a successful model.
- The inputs to the neural network do not provide enough information to predict NO_x, regardless of model.
- The true cycle-to-cycle variations in NO_x are large.

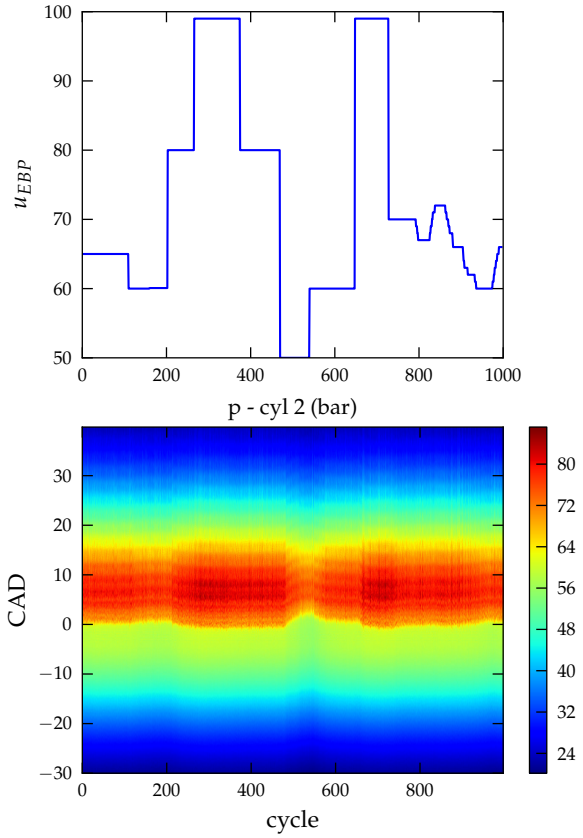

Fig. 6. Data from transient experiment used to illustrate virtual sensor performance. We can see the effect on the cylinder pressure when the exhaust back pressure valve position u_{EBP} was changed.

5. CONCLUSIONS

A virtual sensor that predicts emissions based on cylinder pressure data was presented. The PCA approach was used to find a reduced-dimension approximative representation of the cylinder pressure data that minimized the approximation error. It was shown that a neural network could be trained to predict emissions from the reduced set of variables. Cylinder-to-cylinder differences in predicted emissions were rather substantial. However, the differences were systematic and agreed with physical intuition.

In contrast to previously published work on predicting emissions from cylinder pressure data, the method presented here has been validated on a cycle-to-cycle cylinder-individual basis during transient operation for prediction of both λ, NO_x, and PM. None of the references mentioned in the introduction (Traver et al. [1999], Seykens et al. [2009], Wilhelmsson et al. [2009], Çebi et al. [2011]) present cycle-to-cycle transient predictions. In (Wilhelmsson et al. [2009]), a very small validation data set of 150 cycles collected at steady-state in a single-cylinder engine was used. Cycle-to-cycle variations in NO prediction were shown to be considerable. In (Traver et al. [1999]), only two out of eight cylinders were equipped with pressure sensors and

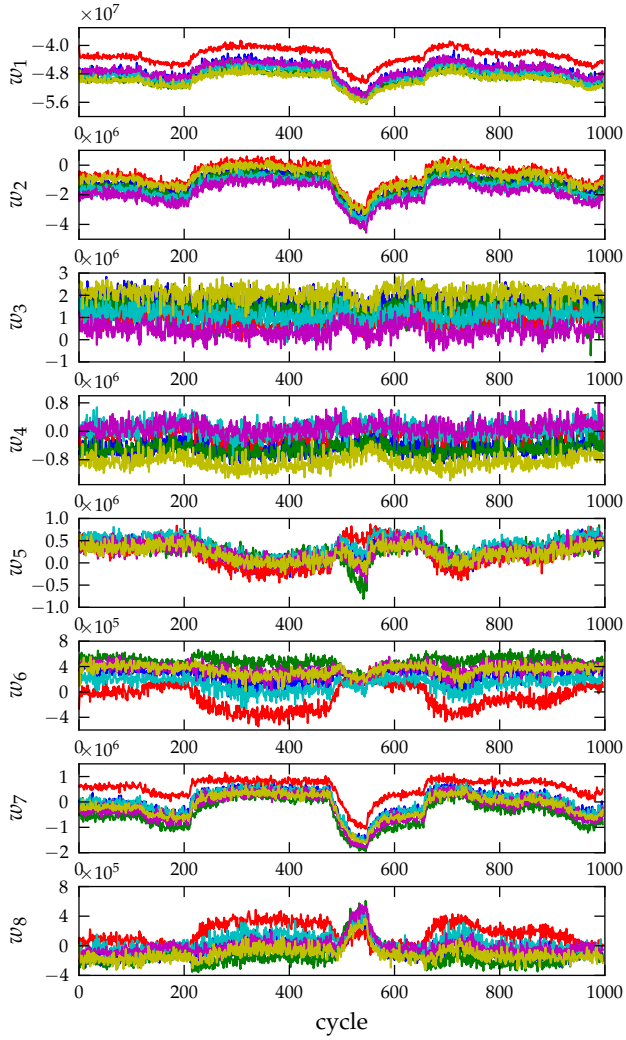

Fig. 7. Principal component weights $w_{1:8}$ for the six cylinders for the transient in Fig. 6.

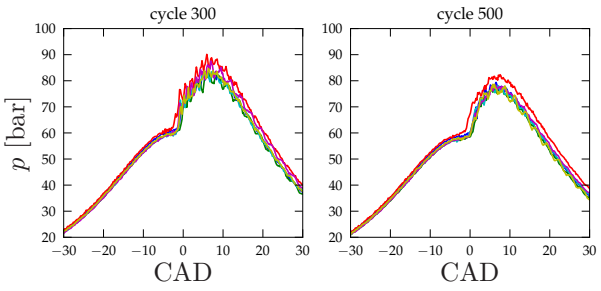

Fig. 8. Cylinder pressure for all six cylinders at engine cycles 300 and 500 for the transient in Fig. 6.

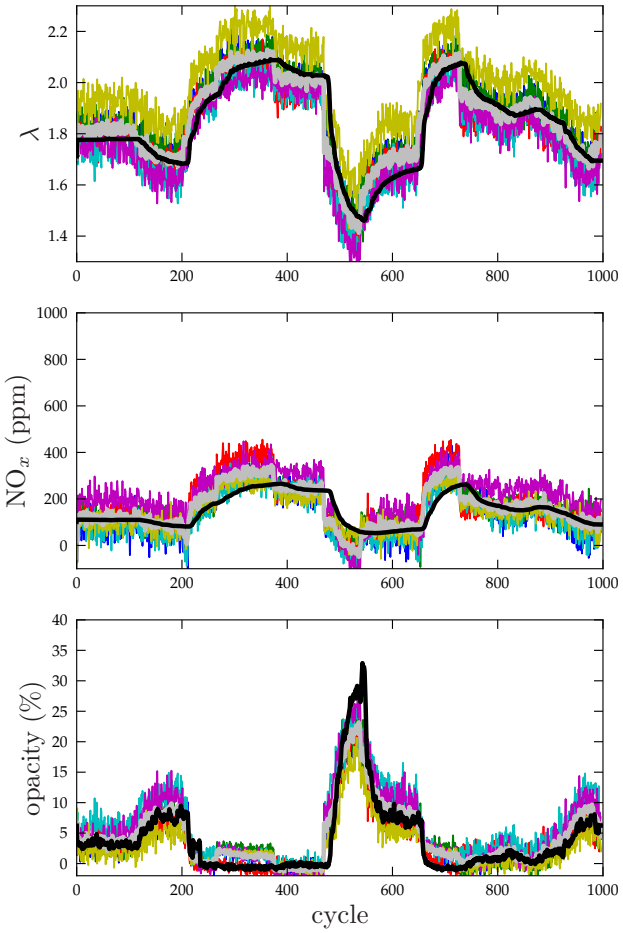

Fig. 9. Target output ζ (black), predicted output averaged over cylinders $\hat{\zeta}$ (grey), and cylinder-individual predicted output $\hat{\zeta}_c$ (colors) for the transient in Fig. 6.

no results on cylinder-to-cylinder differences were shown. The cycle-to-cycle data were low-pass filtered using a five-point moving-average filter. The model was trained at steady-state but still proved to give qualitatively good predictions of NO_x, HC, and CO_2 emissions during an FTP transient. In (Seykens et al. [2009]), only steady-state data was used. Cylinder-to-cylinder variations were shown to be fairly large, but it is not clear from the paper how large the cycle-to-cycle differences were or how the data was averaged. In (Çebi et al. [2011]), an empirical model for PM calibrated at steady-state proved to give very good transient performance. However, it is unclear how many actuators that were manipulated independently in the training and validation data, i.e., which variables that were varied in open loop for excitation as opposed to being determined from other variable through the standard ECU calibration. It is therefore difficult to estimate how well the model would perform in general.

In light of the previous work on cylinder-pressure based estimation of esmissions, the proposed model is rather promising in terms of prediction accuracy. Prediction of λ and NO_x were quantitatively reasonably accurate, and prediction of opacity qualitatively accurate. A key point of using cylinder-pressure based estimates of emissions is to obtain cycle-to-cycle cylinder-individual estimates that could be used for feedback control, for example on-line tuning of start of injection to the optimal NO_x–PM trade-off for each cylinder individually. It is then essential that the emissions estimates are sufficiently accurate for each individual cycle and cylinder, so that no extensive low-pass filtering that decreases the bandwidth of the sensor is required. It should be noted that the validation results presented in the theses are based on an entirely static model. To predict emissions at any given cycle, no information from measurements in previous cycles is used. This leads to fast but noisy measurements. The results could likely be significantly improved by cycle-to-cycle filtering at the expense of a slower response. To avoid steady-state error, a sensor fusion scheme could be introduced where the fast cylinder-individual virtual sensor output is combined with slower sensors that are placed downstreams of the turbine.

All empirical models are sensitive to the data set used for training, and extrapolation performance is generally poor. For future work, it would be of interest to study the robustness of the model to disturbances in the operating mode due to e.g. different intake air temperature or humidity, or aging of engine components. Also, using the virtual sensor for feedback emissions control will be considered.

REFERENCES

S. M. Aithal. Modeling of NO_x formation in diesel engines using finite-rate chemical kinetics. *Applied Energy*, 87: 2256–2265, 2010.

S. Birch. Audi diesel targets Bin 5, Euro 6. *SAE Automotive Engineering International Online*, 2008.

C. M. Bishop. *Pattern Recognition and Machine Learning*. Springer, 2006.

E. Cihan Çebi, G. Rottenkolber, and E. Uyar. In-cylinder pressure based real-time estimation of engine-out particulate matter emissions of a diesel engine. In *SAE Technical Paper 2011-01-1440*, 2011.

General Motors. GM takes new combustion technology out of the lab and on to the road, Press release 2007-08-24, archives.media.gm.com, 2007.

M. Grahn, K. Johansson, C. Vartia, and T. McKelvey. A structure and calibration method for data-driven modeling of NO_x and soot emissions from a diesel engine. In *SAE Technical Paper 2012-01-0355*, 2012.

T. Hastie, R. Tibshirani, and J. Friedman. *The Elements of Statistical Learning*. Springer, New York, 2009.

M. Henningsson, B. Bernhardsson, P. Tunestål, and R. Johansson. A machine-learning approach to information extraction from cylinder pressure sensors. In *SAE Technical Paper 2012-01-0440*, 2012.

M. Hirsch, D. Alberer, and L. del Re. Grey-box control oriented emissions models. In *Proceedings of the 17th IFAC World Congress, Jul 6-11, Seoul, Korea*, pages 8514–8519, 2008.

M. Mrosek, H. Sequenz, and R. Isermann. Control oriented NO_x and soot models for diesel engines. In *Proc. for*

the Sixth IFAC Symposium on Advances in Automotive Control, pages 234–239, Munich, Germany, 2010.

K. Pudenz. Audi sportback concept with combustion chamber sensors for lower emissions. *ATZ online*, 2007.

G. Rizzoni. Methods and apparatus for performing combustion analysis in an internal combustion engine utilizing ignition voltage analysis. United States Patent 5,687,082, 1997.

H. Sequenz and R. Isermann. Emission model structures for an implementation on engine control units. In *Proceedings of the 18th IFAC World Congress, Aug 28–Sep 2, Milan, Italy*, pages 11851–11856, 2011.

X. L. J. Seykens, R. S. G. Baert, L. M. T. Somers, and F. P. T. Willems. Experimental validation of extended NO and soot model for advanced HD diesel engine combustion. In *SAE Technical Paper 2009-01-0683*, 2009.

K. E. Shahroudi. Robust design evolution and impact of in-cylinder pressure sensors to combustion control and optimization: A systems and strategy perspective. Master's thesis, Massachussetts Institute of Technology, Cambridge, MA, USA, 2008.

Y. Shi, R. D. Reitz, and H.-W. Ge. *Computational Optimization of Internal Combustion Engines*. Springer, London, 2011.

Siemens VDO. Smart NOx sensor. www.siemensvdo.com, 2005.

S. Stadlbauer, D. Alberer, M. Hirsch, S. Formentin, C. Benatzky, and L. del Re. Evaluation of virtual NOx sensor models for off road heavy duty diesel engines. In *SAE Technical Paper 2012-01-0358*, 2012.

M. L. Traver, R. J. Atkinson, and C. M. Atkinson. Neural-network-based diesel engine emissions prediction using in-cylinder combustion pressure. In *SAE Technical Paper 1999-01-1532*, 1999.

F. Tschanz, A. Amstutz, C. H. Onder, and L. Guzzella. A real-time soot model for emission control of a diesel engine. In *Proc. for the Sixth IFAC Symposium on Advances in Automotive Control*, pages 222–227, Munich, Germany, 2010.

P. Tunestål. Self-tuning gross heat release computation for internal combustion engines. *Control Engineering Practice*, 17:518–524, 2009.

A. Westlund and H.-E. Ångström. Fast physical prediction of NO and soot in diesel engines. In *SAE Technical Paper 2009-01-1121*, 2009.

C. Wilhelmsson, P. Tunestål, B. Johansson, A. Widd, and R. Johansson. A physical two-zone NO_x model intended for embedded implementation. In *SAE Technical Paper 2009-01-1509*, 2009.

Appendix A. COST FUNCTION GRADIENTS

Algebraic expressions for the gradients of the two cost functions R_{pe} and R_{reg} used for neural network training are presented here. The input q is a three-dimensional tensor, with the dimensions representing sample index, input index, and cylinder index, respectively.

The derivatives with respect to $\alpha_0 \in \mathbb{R}^M$, $\alpha \in \mathbb{R}^{M\times L}$, $\beta_0 \in \mathbb{R}^K$, $\beta \in \mathbb{R}^{K\times M}$ are derived element-wise to avoid introducing notation for tensor operations. Note that computations can be implemented efficiently using multi-dimensional arrays and tensor multiplication.

$$\frac{\partial R_{pe}}{\partial \alpha_{m0}} = -\frac{1}{C}\sum_{n=1}^{N}\sum_{k=1}^{K}\sum_{c=1}^{C}(\zeta_{nk}-\bar{\hat{\zeta}}_{nk})\beta_{km}\sigma'(\xi_{nmc}) \tag{A.1}$$

$$\frac{\partial R_{pe}}{\partial \alpha_{ml}} = -\frac{1}{C}\sum_{n=1}^{N}\sum_{k=1}^{K}\sum_{c=1}^{C}(\zeta_{nk}-\bar{\hat{\zeta}}_{nk})\beta_{km}\sigma'(\xi_{nmc})q_{nlc} \tag{A.2}$$

$$\frac{\partial R_{pe}}{\partial \beta_{k0}} = -\sum_{n=1}^{N}(\zeta_{nk}-\bar{\hat{\zeta}}_{nk}) \tag{A.3}$$

$$\frac{\partial R_{pe}}{\partial \beta_{km}} = -\sum_{n=1}^{N}(\zeta_{nk}-\bar{\hat{\zeta}}_{nk})\bar{z}_{nm} \tag{A.4}$$

$$\frac{\partial R_{reg}}{\partial \alpha_{m0}} = \frac{1}{C}\sum_{n=1}^{N}\sum_{k=1}^{K}\sum_{c=1}^{C}\hat{\zeta}_{nkc}\beta_{km}(\sigma'(\xi_{nmc}) \tag{A.5}$$

$$-\frac{1}{C}\sum_{d=1}^{C}\sigma'(\xi_{nmc})) \tag{A.6}$$

$$\frac{\partial R_{reg}}{\partial \alpha_{ml}} = \frac{1}{C}\sum_{n=1}^{N}\sum_{k=1}^{K}\sum_{c=1}^{C}\hat{\zeta}_{nkc}\beta_{km}(\sigma'(\xi_{nmc})q_{nlc} \tag{A.7}$$

$$-\frac{1}{C}\sum_{d=1}^{C}\sigma'(\xi_{nmd})q_{nld}) \tag{A.8}$$

$$\frac{\partial R_{reg}}{\partial \beta_{k0}} = 0 \tag{A.9}$$

$$\frac{\partial R_{reg}}{\partial \beta_{km}} = \frac{1}{C}\sum_{n=1}^{N}\sum_{c=1}^{C}\hat{\zeta}_{nkc}(z_{nmc}-\bar{z}_{nm}) \tag{A.10}$$