Real-Time, Adaptive Machine Learning for Non-Stationary, Near Chaotic Gasoline Engine Combustion Time Series[☆]

Adam Vaughan^{a,*}, Stanislav V. Bohac^a

^a*Dept. of Mechanical Engineering, University of Michigan, Ann Arbor, MI 48109, USA*

Abstract

Fuel efficient Homogeneous Charge Compression Ignition (HCCI) engine combustion timing predictions must contend with non-linear chemistry, non-linear physics, period doubling bifurcation(s), turbulent mixing, model parameters that can drift day-to-day, and air-fuel mixture state information that cannot typically be resolved on a cycle-to-cycle basis, especially during transients. In previous work, an abstract cycle-to-cycle mapping function coupled with ϵ -Support Vector Regression was shown to predict experimentally observed cycle-to-cycle combustion timing over a wide range of engine conditions, despite some of the aforementioned difficulties. The main limitation of the previous approach was that a partially acausal randomly sampled training dataset was used to train proof of concept offline predictions. The objective of this paper is to address this limitation by proposing a new online adaptive Extreme Learning Machine (ELM) extension named Weighted Ring-ELM. This extension enables fully causal combustion timing predictions at randomly chosen engine set points, and is shown to achieve results that are as good as or better than the previous offline method. The broader objective of this approach is to enable a new class of real-time model predictive control strategies for high variability HCCI and, ultimately, to bring HCCI's low engine-out NO_x and reduced CO_2 emissions to production engines.

Keywords: non-linear, non-stationary, time series, chaos theory, dynamical system, adaptive extreme learning machine

1. Introduction

Since the 1800s, gasoline engines have largely been operated by (1) controlling power output with a throttle that restricts airflow, (2) using a simple spark to control burn timing, and (3) operating close to fuel-air stoichiometry for reliable spark ignition and so catalysts can reduce NO_x , HC, and CO emissions. The throttle hurts fuel efficiency with pumping losses (especially at low-load), and the stoichiometric mixtures used are thermodynamically less fuel efficient than mixtures diluted with air or exhaust gases.

With the broad availability of enabling technologies (e.g. variable valve timing), a relatively new type of combustion called Homogeneous Charge Compression Ignition (HCCI) has received increased research interest over the past decade. HCCI uses autoignition to burn lean (excess air) mixtures and can produce ultra-low NO_x quantities that do not require expensive catalyst aftertreatment. Instead of a spark, combustion timing is controlled by the thermodynamic trajectory of the mixture and complex chemical kinetics. With both ultra-low NO_x production and freedom from the stoichiometric shackles of spark ig-

nition, HCCI achieves greater fuel efficiency through thermodynamically ideal lean mixtures and unthrottled operation. This improved fuel economy, has real-world relevance to near-term sustainability, national oil independence, and greenhouse gas initiatives that seek to curb petroleum usage.

The primary challenge of HCCI autoignition is to ensure that the burn timing is synchronized against the motion of the piston. This is important for efficient extraction of mechanical work from the fuel-air mixture and to avoid unsafe, noisy combustion or unstable (near chaotic) combustion oscillations. This synchronization is so important that combustion researchers do not use normal units of time. Instead, they use the angle of the crank, which (1) describes the position of the piston, and (2) represents time because each crank angle takes a certain amount of time at a fixed crank rotation speed. The thermodynamic trajectory of the air-fuel mixture is driven by the piston varying the cylinder volume as function of crank angle (Fig. 1). These angles are measured relative to when the piston is at the top of the cylinder, or Top Dead Center (TDC). In a four-stroke engine, TDC occurs twice per cycle. In different regions, the piston may be compressing or expanding the mixture, or, if a valve is open, moving the mixture into or out of the intake or exhaust manifolds.

Highlighted on the cylinder volume curve are two regions, one for when the exhaust valve is open and the other for when the intake valve is open. Note that the two

[☆]There are no fundamental changes in this updated version of the original October, 2013 pre-print paper [1]. This version includes algebraic simplifications, minor corrections, and improved body text.

*Corresponding author

Email addresses: vaughana@umich.edu (Adam Vaughan), sbohac@umich.edu (Stanislav V. Bohac)

Figure 1: A schematic of key engine cycle variables.

valve events are separated by a number of crank angle degrees, termed Negative Valve Overlap or NVO. Unlike conventional engines, NVO prevents some of the hot exhaust gases from leaving the cylinder (typically 20-60% [2, 3]). This stores “residual gases” for the next cycle, offering a practical way to raise the mixture temperature to ensure HCCI autoignition [4]. By changing the amount of NVO, one can affect the mixture temperature and dilution and ultimately control the chemical kinetics behind combustion timing. Temperature and dilution work in opposite directions, but typically temperature dominates [5]. NVO is not instantly adjustable with common variable valve timing systems, and the reader is cautioned that many researchers publish results with fully variable (lift and timing) electric or hydraulic valve actuation systems that are expensive to implement in production engines.

The use of NVO residual gases introduces strong cycle-to-cycle coupling on top of the already non-linear chemistry and physics that occur throughout a complete engine cycle [6]. Further compounding the issues with residual gases is that neither the airflow to the cylinder(s) nor the quantity of residual gases in the cylinder can be accurately resolved before a burn happens on a *cycle-to-cycle* (not mean value) basis with commonly available sensors, especially during transients. Beyond residual gas influences, there are also complex secondary influences on combustion behavior such as turbulent mixing, manifold resonance effects, combustion deposits, different varieties of fuel and even ambient temperature variations [7, 8].

While HCCI is already a significant challenge given the above complexity, the combustion mode also exhibits a period doubling bifurcation cascade to chaos [6, 9, 10], similar to what is seen in high residual spark ignition engines [11]. When nearly chaotic, HCCI is still deterministic, but becomes oscillatory and very sensitive to parameter variations (e.g. residual gas fraction fluctuations [9, 10]). This oscillatory “stability limit” behavior is commonly referred

to as high Cyclic Variability (CV) and it severely constrains the available load limits of HCCI.

1.1. Motivation and goals

A primary constraint for HCCI is the need to keep combustion timing between the ringing and combustion stability limits [12]. At the ringing limit, excessive pressure rise rates are encountered, and at the stability limit, high CV in combustion timing is observed [12]. Since these limits play a key role in constraining HCCI’s usable operating range, it is desirable to explore new methods to predict the behavior at and beyond these constraints. In particular, the ability to predict and correct for high CV might enable the use of late phased combustion to mitigate the excessive pressure rise rates that currently constrain HCCI’s high-load operation [13], while also potentially addressing the high CV experienced at low-load. Towards the end goal of expanding the HCCI load envelope, this paper builds on previous work [6] by describing a new online adaptive machine learning method that enables fully causal cycle-to-cycle combustion timing predictions across randomly chosen engine set point transients that include both the stable and near chaotic bifurcation behavior.

1.2. Experimental Observations

In the authors’ previous publication [6], an abstract mapping function for engine combustion was created within the framework of a discrete dynamical system:

$$\begin{aligned} \text{next combustion} = \\ \text{function}(\text{previous combustion}, \text{parameters}) \end{aligned} \quad (1)$$

This simple abstraction is intended to convey a conceptual understanding of the experimental cycle-to-cycle behavior seen in Fig. 2’s return maps. These return maps show the experimentally observed combustion timing for a given cycle n along the abscissa and the next cycle $n + 1$ along the ordinate under random engine actuator set points [6]. The value CA90 is the time in Crank Angle Degrees (°CA) where 90% of the fuel’s net heat release is achieved, and thus measures the timing of the end of the burn in relation to piston’s position as the crank rotates.¹

The reader should note that there is structure to the cycle-to-cycle behavior despite the random actuator set points used to generate Fig. 2. This structure shows a deterministic transition to oscillatory high CV behavior as combustion moves towards later combustion timing that can be viewed as at least a single period doubling bifurcation with sensitive dependence on the engine set point [6, 9, 10]. While mathematically interesting, this oscillatory high CV structure undesirably constrains practical HCCI engine operation. A more thorough description of these data is provided in [6].

¹While it is not shown here, there is similar structure in the CA10 and CA50 percent burn metrics, although less pronounced (especially in CA10, see [6]).

Figure 2: Return map probability histograms of CA90 generated from 129,964 cycles and 2,221 random engine set points. Outliers are omitted and total only $\sim 3\%$ of the data. The colormap is $\log_{10}$ to show order of magnitude differences. [6]

1.3. Modeling Approaches

In [6], a skeletal functional form for the abstract mapping function was built out using measurable quantities, thermodynamics, and known correlations. Then, unlike the physics-based approaches that are usually discussed in the engine literature, the machine learning technique of ϵ -Support Vector Regression (ϵ -SVR) was combined with the abstract functional form to provide quantitative predictions. The primary motivation for this machine learning approach was not that existing chemical kinetics with Computational Fluid Dynamics (CFD) cannot capture engine behavior (see [14] and gasoline mechanism validation [15]) but that the methods are:

- Too computationally intensive for real-time predictions.
- Subject to experimental uncertainties in the *cycle-to-cycle* (not mean value) mixture state and composition.

As a point of reference, the simulation time of a 2,500 RPM, 48 millisecond engine cycle is measured in $\sim \text{day(s)}$ for a single core of a modern computer.

At the other computational complexity extreme, low-order approximation models of HCCI for control have been developed since at least the early 2000s [16], based on spark ignition engine knock models developed in the 1950s [17]. Recently, efforts have been made to extend this type of model to the high CV regions of HCCI by injecting random residual gas fraction noise to capture uncertainties in the mixture state and composition [10]. This model was tuned for a limited set of steady-state conditions and only “predictive” in the sense that when injected with random (i.e. unpredictable) noise could it generate a qualitative

return map shape similar to what is seen experimentally. Time series predictions were not shown, only a cloud of possible combustion timings ranging from stable to oscillatory. This fact was highlighted in an extension of the work [18] that again used random residual noise because “the actual time series of disturbances to the experiments are unknown.” That said, these models are useful for showing that a period doubling cascade to chaos driven by residual gas fraction can explain the observed high CV behavior.

In the context of the above, machine learning provides a computationally efficient way to capture complex combustion patterns while simultaneously avoiding explicit knowledge of the underlying mixture state and composition (provided an appropriate abstract mapping function is chosen). While there are clearly benefits to this machine learning approach, a key issue is that machine learning is data driven, and relatively large quantities of data are needed to adequately cover large dimensional spaces. As shown conceptually in Fig. 3, these high dimensional data might be viewed as a “porcupine” [19]. Each engine operating condition might be viewed as a “quill” of Eq. 2’s six-dimensional “porcupine,” and machine learning algorithms know nothing about the ideal gas law or chemical kinetics, so their ability to extrapolate between “quills” is limited, especially when provided sparse data. Previous work [6] used a random sampling of cycle time series for training to ensure the data driven model had data to fit the “quills,” and then assessed the model’s ability to predict on the remaining (randomly chosen) cycles. Thus, the training dataset was partially acausal and that the model itself wasn’t shown to adapt to new conditions.

Figure 3: High dimensional data might be viewed conceptually as a “porcupine” [19]. The primary goal of this paper is to design an online adaptive algorithm to fit new data between the “quills.”

1.4. Contribution

The primary contribution of this work is the development of a new online learning method to provide real-time adaptive, fully causal predictions of near chaotic HCCI combustion timing. This method, called Weighted Ring - Extreme Learning Machine (WR-ELM), enables robust online updates to an Extreme Learning Machine (ELM) model that is trained to fit the “quills” of offline data.

2. Methods

2.1. Mapping Function Modifications

In previous work [6], engine combustion was abstracted to the following mapping function:

$$CA50_{n+1} = \text{function}(CA90_n, TI, SOI, P_{IVC}, P_{EVO}, P_{NVO}) \quad (2)$$

where n is the cycle iteration index, $CA50$ is the time in $^\circ CA$ where 50% of net heat release has occurred, $CA90$ is the time in $^\circ CA$ when 90% of net heat release has occurred, TI is the injection pulse width in milliseconds, SOI is Start of Injection in $^\circ CA$ Before Top Dead Center ($^\circ BTDC$), and the pressure variables measurements are mean pressures during specific regions of the combustion cycle (see Fig. 1).² Details of the simplified net heat release algorithm are available in [6]. Fuel rail pressure is constant; however, the reader should note that the pressure drop to cylinder pressure during NVO injections varies with each transient step and during high CV regions. The cylinder pressure variables P_{IVC} , P_{EVO} , and P_{NVO} were chosen to capture cycle-to-cycle residual coupling and air flow without the difficulties of explicitly modeling those quantities. To meet real-time engine controller timing requirements, P_{IVC} and P_{NVO} have been modified from [6]. P_{IVC} has been moved to the previous cycle, and the range of P_{NVO} 's mean has been shortened. P_{IVC} has also been moved closer to TDC to take advantage of the inherent signal amplification provided by the compression process. The subscripts IVC, EVO, and NVO refer to the general timing regions Intake Valve Close, Exhaust Valve Open, and Negative Valve Overlap, respectively.

2.2. WR-ELM Overview

The primary benefits of an ELM approach over the ϵ -SVR method used in [6] are:

- An ELM is easily adapted to online adaptation [20].
- An ELM provides good model generalization when the data are noisy [21].
- An ELM is extremely computationally efficient [20, 21].

WR-ELM is developed in this work as a weighted least squares extension to Online Sequential - ELM [20]. While developed independently, a similar derivation is available in [22]. The difference between this work and the classification application in [22] is the use of a ring buffer data

structure “chunk” for online updates to an offline trained regression model.

The data in the WR-ELM ring buffer can be weighted more heavily than the data originally used to fit the offline model. This allows emphasis to be placed on recent measurements that might be between the “quills” of Fig. 3's offline trained model or the result of day-to-day engine parameter variation. Thus, this approach allows one to prescribe a partitioned balance between the offline model fit against the need to adapt to the most recent conditions. It also explicitly avoids over adaptation to the local conditions (that could compromise global generality) by “forgetting” old ring buffer data that eventually exit the buffer. Fig. 4 gives a schematic representation of this approach.

Figure 4: A schematic of overview WR-ELM.

Other differences from [20, 22] are that the derivation below lacks a bias vector b , uses the Gaussian distribution for $\mathbf{a}$, drops the unnecessary logistic function exponential negative, and uses a Padé approximant for $\exp(x)$. It was found empirically that the computation of the bias b addition step could be removed with no loss of fitting performance if $\mathbf{a}$'s elements were drawn from the Gaussian distribution $\mathcal{N}(0, 1)$. ELM theory only requires the distribution to be continuous [21], although the ability to remove the bias is likely problem specific.

2.3. WR-ELM Core Algorithm

The basic goal of an Extreme Learning Machine (ELM) is to solve for the output layer weight vector β that scales the transformed input $\mathbf{H}$ to output $\mathbf{T}$:

$$\mathbf{H}\beta = \mathbf{T} \quad (3)$$

where $\mathbf{H}$ is the hidden layer output matrix of a given input matrix and $\mathbf{T}$ is the target vector.

For a set of n input-output data pairs and $\tilde{N}$ neurons at the n th cycle timestep, these variables are given by

$$\mathbf{H}(\mathbf{a}, \mathbf{x}) = \begin{bmatrix} G(\mathbf{a}_1, \mathbf{x}_1) & \dots & G(\mathbf{a}_{\tilde{N}}, \mathbf{x}_1) \\ \vdots & \ddots & \vdots \\ G(\mathbf{a}_1, \mathbf{x}_n) & \dots & G(\mathbf{a}_{\tilde{N}}, \mathbf{x}_n) \end{bmatrix}_{n \times \tilde{N}} \quad (4)$$

where $G(\mathbf{a}_i, \mathbf{x})$ is the neuron activation function, chosen to be a commonly used logistic function, but without the unnecessary negative:

$$G(\mathbf{a}_i, \mathbf{x}) = \frac{1}{1 + \exp(\mathbf{x} \cdot \mathbf{a}_i)} \quad (5)$$

²Since $CA90$ is a stronger indicator of the oscillatory behavior seen in Fig. 2, it is used as the model input. $CA50$ is more commonly encountered in the engine literature and is thus used as the model output. The two are related quantities, and in terms of model fit statistics there was no significant benefit of using one over the other. That said, a few isolated instances were observed where $CA90$ did a better job predicting large oscillations.

Using a random input weight vector $\mathbf{a}_i$ that is composed of random variable (r.v.) samples from a Gaussian distribution for each of the z input variables gives

$$\mathbf{a}_i = \begin{bmatrix} r.v. \sim \mathcal{N}(0, 1) \\ \vdots \\ r.v. \sim \mathcal{N}(0, 1) \end{bmatrix}_{z \times 1} \quad (6)$$

The use of a random $\mathbf{a}_i$ that is re-sampled for each of the $\tilde{N}$ individual neurons during initialization is the main difference of an Extreme Learning Machine versus conventional neural networks that iteratively train each $\mathbf{a}_i$ [21]. These $\mathbf{a}_i$ vectors can then be collected into a single input weight matrix $\mathbf{a}$, which is held fixed across all n input row vectors $\mathbf{x}$

$$\mathbf{x} = [CA90_n \quad TI \quad SOI \quad P_{IVC} \quad P_{EVO} \quad P_{NVO}]_{n \times z} \quad (7)$$

and n output values

$$\mathbf{T} = [CA50_{n+1}]_{n \times 1} \quad (8)$$

While the above logistic works well, one modification improves the computational efficiency on processors without a dedicated $\exp(x)$ instruction, such as the Raspberry Pi[®]. The modification is to replace the exponential with the following Padé approximant:

$$\exp(y) \approx p(y) = \frac{120 + 60y + 12y^2 + y^3}{120 - 60y + 12y^2 - y^3}, \quad (9)$$

which has the following simple logistic relations:

$$\frac{1}{1 + \exp(y)} \approx \frac{1}{1 + p(y)} = \frac{(120 + 12y^2) - 60y - y^3}{2 \cdot (120 + 12y^2)} \quad (10)$$

The small number of floating point operations used, reused intermediate terms, known boundedness of the normalized inputs, and known $\mathbf{a}$ weights make this approximant work well in this application. No significant degradation in model performance was found, and as such it is used in all implementations described hereafter.

The normal equations can then be used to solve for the least squares solution β of Eq. 3 with

$$\beta = (\mathbf{H}^\top \mathbf{H})^{-1} \mathbf{H}^\top \mathbf{T} \quad (11)$$

To extend this to a weighted least squares solution, one can incorporate a diagonal weight matrix $\mathbf{W}$ to the normal equations [23]:

$$\beta = (\mathbf{H}^\top \mathbf{W} \mathbf{H})^{-1} \mathbf{H}^\top \mathbf{W} \mathbf{T} \quad (12)$$

The solution can then be split between an offline and online “chunk” of recent input-output data pairs to avoid both the computational and storage burden of using offline data directly. To do this, the matrices are partitioned with

subscript 0 and 1 denoting the offline and online updated components, respectively:

$$\mathbf{H} = \begin{bmatrix} \mathbf{H}_0 \\ \mathbf{H}_1 \end{bmatrix}_{n \times \tilde{N}} \quad \mathbf{W} = \begin{bmatrix} \mathbf{W}_0 & \mathbf{0} \\ \mathbf{0} & \mathbf{W}_1 \end{bmatrix}_{n \times N} \quad (13)$$

$$\mathbf{T} = \begin{bmatrix} \mathbf{T}_0 \\ \mathbf{T}_1 \end{bmatrix}_{n \times 1}$$

Then, following a similar derivation in [20] for recursive least squares but adding the weight matrix, the inversion portion $\mathbf{H}^\top \mathbf{W} \mathbf{H}$ of the weighted normal equations Eq. 12 can be re-written in terms of $\mathbf{K}_0$ and $\mathbf{K}_1$:

$$\begin{aligned} \mathbf{K}_1 &= \\ &= \mathbf{H}^\top \mathbf{W} \mathbf{H} \\ &= \begin{bmatrix} \mathbf{H}_0 \\ \mathbf{H}_1 \end{bmatrix}^\top \begin{bmatrix} \mathbf{W}_0 & \mathbf{0} \\ \mathbf{0} & \mathbf{W}_1 \end{bmatrix} \begin{bmatrix} \mathbf{H}_0 \\ \mathbf{H}_1 \end{bmatrix} \\ &= [\mathbf{H}_0^\top \quad \mathbf{H}_1^\top] \begin{bmatrix} \mathbf{W}_0 & \mathbf{0} \\ \mathbf{0} & \mathbf{W}_1 \end{bmatrix} \begin{bmatrix} \mathbf{H}_0 \\ \mathbf{H}_1 \end{bmatrix} \\ &= [\mathbf{H}_0^\top \mathbf{W}_0 \quad \mathbf{H}_1^\top \mathbf{W}_1] \begin{bmatrix} \mathbf{H}_0 \\ \mathbf{H}_1 \end{bmatrix} \\ &= \mathbf{H}_0^\top \mathbf{W}_0 \mathbf{H}_0 + \mathbf{H}_1^\top \mathbf{W}_1 \mathbf{H}_1 \\ &= \mathbf{K}_0 + \mathbf{H}_1^\top \mathbf{W}_1 \mathbf{H}_1 \end{aligned} \quad (14)$$

The non-inverted portion of the normal equations can similarly be re-written using existing relations:

$$\begin{aligned} \mathbf{H}^\top \mathbf{W} \mathbf{T} &= \\ &= \begin{bmatrix} \mathbf{H}_0 \\ \mathbf{H}_1 \end{bmatrix}^\top \begin{bmatrix} \mathbf{W}_0 & \mathbf{0} \\ \mathbf{0} & \mathbf{W}_1 \end{bmatrix} \begin{bmatrix} \mathbf{T}_0 \\ \mathbf{T}_1 \end{bmatrix} \\ &= [\mathbf{H}_0^\top \quad \mathbf{H}_1^\top] \begin{bmatrix} \mathbf{W}_0 & \mathbf{0} \\ \mathbf{0} & \mathbf{W}_1 \end{bmatrix} \begin{bmatrix} \mathbf{T}_0 \\ \mathbf{T}_1 \end{bmatrix} \\ &= [\mathbf{H}_0^\top \mathbf{W}_0 \quad \mathbf{H}_1^\top \mathbf{W}_1] \begin{bmatrix} \mathbf{T}_0 \\ \mathbf{T}_1 \end{bmatrix} \\ &= \mathbf{H}_0^\top \mathbf{W}_0 \mathbf{T}_0 + \mathbf{H}_1^\top \mathbf{W}_1 \mathbf{T}_1 \\ &= \mathbf{K}_0 \mathbf{K}_0^{-1} \mathbf{H}_0^\top \mathbf{W}_0 \mathbf{T}_0 + \mathbf{H}_1^\top \mathbf{W}_1 \mathbf{T}_1 \\ &= \mathbf{K}_0 \beta_0 + \mathbf{H}_1^\top \mathbf{W}_1 \mathbf{T}_1 \\ &= (\mathbf{K}_1 - \mathbf{H}_1^\top \mathbf{W}_1 \mathbf{H}_1) \beta_0 + \mathbf{H}_1^\top \mathbf{W}_1 \mathbf{T}_1 \\ &= \mathbf{K}_1 \beta_0 - \mathbf{H}_1^\top \mathbf{W}_1 \mathbf{H}_1 \beta_0 + \mathbf{H}_1^\top \mathbf{W}_1 \mathbf{T}_1 \end{aligned} \quad (15)$$

Substituting Eq. 15 into the full online solution

$$\begin{aligned} \beta_1 &= \\ &= (\mathbf{K}_1^{-1}) (\mathbf{H}^\top \mathbf{W} \mathbf{T}) \\ &= \beta_0 - \mathbf{K}_1^{-1} \mathbf{H}_1^\top \mathbf{W}_1 \mathbf{H}_1 \beta_0 + \mathbf{K}_1^{-1} \mathbf{H}_1^\top \mathbf{W}_1 \mathbf{T}_1 \\ &= \beta_0 + \mathbf{K}_1^{-1} \mathbf{H}_1^\top \mathbf{W}_1 (\mathbf{T}_1 - \mathbf{H}_1 \beta_0) \end{aligned} \quad (16)$$

yields the online solution without the need for the offline dataset. To trade the computational burden of the $\tilde{N} \times \tilde{N}$ sized $\mathbf{K}$ inverse for an inverse that scales with a smaller

sized ring buffer, one can let $\mathbf{P} = \mathbf{K}^{-1}$

$$\mathbf{P}_0 = \mathbf{K}_0^{-1} = \left(\mathbf{H}_0^\top \mathbf{W}_0 \mathbf{H}_0 \right)^{-1} \quad (17)$$

$$\mathbf{P}_1 = \mathbf{K}_1^{-1} = \left(\mathbf{P}_0^{-1} + \mathbf{H}_1^\top \mathbf{W}_1 \mathbf{H}_1 \right)^{-1} \quad (18)$$

and use the matrix inversion lemma on Eq. 18 to yield:

$$\mathbf{P}_1 = \mathbf{P}_0 - \mathbf{P}_0 \mathbf{H}_1^\top \left(\mathbf{W}_1^{-1} + \mathbf{H}_1 \mathbf{P}_0 \mathbf{H}_1^\top \right)^{-1} \mathbf{H}_1 \mathbf{P}_0. \quad (19)$$

Unlike OS-ELM and WOS-ELM [20, 22], additional simplification is possible because the WR-ELM algorithm does not propagate $\mathbf{P}_1$ in time. To begin, append the $\mathbf{H}_1^\top \mathbf{W}_1$ portion of Eq. 16 to Eq. 19 and distribute $\mathbf{H}_1^\top$ to give:

$$\begin{aligned} \mathbf{P}_1 \mathbf{H}_1^\top \mathbf{W}_1 = & \left[\mathbf{P}_0 \mathbf{H}_1^\top - \mathbf{P}_0 \mathbf{H}_1^\top \left(\mathbf{W}_1^{-1} + \mathbf{H}_1 \mathbf{P}_0 \mathbf{H}_1^\top \right)^{-1} \right. \\ & \left. \cdot \mathbf{H}_1 \mathbf{P}_0 \mathbf{H}_1^\top \right] \mathbf{W}_1. \end{aligned} \quad (20)$$

Eq. 20 can then be simplified with the substitutions $\mathbf{A} = \mathbf{P}_0 \mathbf{H}_1^\top$, $\mathbf{B} = \mathbf{H}_1 \mathbf{A}$ and then distributing $\mathbf{W}_1$ to provide:

$$\mathbf{A} \left[\mathbf{W}_1 - \left(\mathbf{W}_1^{-1} + \mathbf{B} \right)^{-1} \mathbf{B} \mathbf{W}_1 \right]. \quad (21)$$

Transforming Eq. 21 with the identity $(\mathbf{X} + \mathbf{Y})^{-1} \mathbf{Y} = \mathbf{X}^{-1} (\mathbf{X}^{-1} + \mathbf{Y}^{-1})^{-1}$ [24] gives:

$$\mathbf{A} \left[\mathbf{W}_1 - \mathbf{W}_1 \left(\mathbf{W}_1 + \mathbf{B}^{-1} \right)^{-1} \mathbf{W}_1 \right]. \quad (22)$$

Eq. 22 is then in a form where the identity $\mathbf{X} - \mathbf{X}(\mathbf{X} + \mathbf{Y})^{-1} \mathbf{X} = (\mathbf{X}^{-1} + \mathbf{Y}^{-1})^{-1}$ [24] can be applied to yield a substantially simpler form with a ring buffer sized inverse:

$$\mathbf{P}_1 \mathbf{H}_1^\top \mathbf{W}_1 = \mathbf{A} \left(\mathbf{W}_1^{-1} + \mathbf{B} \right)^{-1}. \quad (23)$$

Finally, noting that $\mathbf{P}_1 \mathbf{H}_1^\top \mathbf{W}_1 = \mathbf{K}_1^{-1} \mathbf{H}_1^\top \mathbf{W}_1$, one can then substitute Eq. 23 into Eq. 16 and arrive at the following algorithm summary:

OFFLINE TRAINING

$$\begin{aligned} \mathbf{P}_0 &= \left[\left(\mathbf{H}_0^\top \mathbf{W}_0 \mathbf{H}_0 \right)^{-1} \right]_{\tilde{N} \times \tilde{N}} \\ \beta_0 &= \left[\mathbf{P}_0 \mathbf{H}_0^\top \mathbf{W}_0 \mathbf{T}_0 \right]_{\tilde{N} \times 1} \end{aligned} \quad (24)$$

ONLINE ADAPTATION

$$\begin{aligned} \mathbf{A} &= \mathbf{P}_0 \mathbf{H}_1^\top, \quad \mathbf{B} = \mathbf{H}_1 \mathbf{A} \\ \beta_1 &= \beta_0 + \mathbf{A} \left(\mathbf{W}_1^{-1} + \mathbf{B} \right)^{-1} (\mathbf{T}_1 - \mathbf{H}_1 \beta_0) \end{aligned} \quad (25)$$

ONLINE PREDICTIONS

$$\text{CA50}_{n+2} = \mathbf{T}_{n+1} = \mathbf{H}(\mathbf{a}, \mathbf{x}_{n+1}) \beta_1 \quad (26)$$

The reader should note that only $\mathbf{P}_0$ and β_0 are needed for online adaptation, and the size of these matrices scales only with an increasing number of neurons $\tilde{N}$. **None of the original offline data are needed.** Additionally, note that Eq. 26 is simply the reverse of Eq. 3 with the most recent $\mathbf{x}_{n+1}$ cycle vector and β_1 updated from the weighted ring buffer. Finally, it should be mentioned that the resulting update law Eq. 25 is structurally similar that of the steady-state Kalman filter [23], which also uses recursive least squares. Future work should look at applying Kalman filtering algorithm improvements (e.g. square root filtering) to WR-ELM.

2.4. Usage procedure

1. Scale $\mathbf{x}$ and $\mathbf{T}$ columns between zero and unity for each variable. For the combustion implementation, column variable values below the 0.1% and above the 99.9% percentile were saturated at the respective percentile value, and then normalized between zero and unity between these percentile based saturation limits. This was done to both adequately represent the distribution tails and avoid scaling issues.
2. The random non-linear transformation that enables the low computational complexity of the WR-ELM algorithm may result in ill-conditioned matrices. All numerical implementations should use double precision. Additionally, one should consider using Singular Value Decomposition for ill-conditioned matrix inversions.
3. Using the $\mathcal{N}(0, 1)$ Gaussian distribution, initialize the $z \times \tilde{N}$ ELM input weights $\mathbf{a}$ and hold them fixed for all training / predictions. For the combustion implementation, this was done with MATLAB®'s built-in *randn()* function and the Mersenne Twister pseudo random number generator with seed 7,898,198. An $\tilde{N}$ of 64 was used based on initial trials, and each cylinder's individually computed WR-ELM model used an identical input weight matrix $\mathbf{a}$.
4. Build $\mathbf{H}_0(\mathbf{a}, \mathbf{x}_0)$ from previously acquired samples that cover a wide range of conditions with Eq. 4 using an input matrix $\mathbf{x}_0$ and output target vector $\mathbf{T}_0$ (the formats of these are given in Eqs. 7 and 8, respectively). For the combustion implementation, the initial training data were ~40 minutes of random engine set points covering 53,884 cycles and 1,179 random engine set points at a single engine speed; however, it appears that only ~20 minutes of data may be sufficient. Pruning the training data to only include ~6 cycles before and ~9 cycles after a transient set point step provided a small model fitting performance improvement.

5. Specify a weight matrix $\mathbf{W}_0$ for offline measurements. For the combustion implementation, a simple scalar value $\mathbf{W}_0 = 3.5 \times 10^{-3}$ was chosen using a design of experiments. While this weight works well as a proof of concept, future work should more rigorously determine the weight(s), perhaps with optimization techniques. Note that $\mathbf{W}_0$ allows weighting to be applied offline and that a small offline weighting is equivalent to a large online weighting.
6. Solve for the offline solution $\mathbf{P}_0$ and β_0 using Eqs. 24 and hold these values constant for all future predictions.
7. Populate a ring buffer of size r with recently completed input-output pairs using:

$$\begin{aligned} \text{input ring buffer} = \mathbf{x}_1 &= \begin{bmatrix} \mathbf{x}_{n-r+1} \\ \vdots \\ \mathbf{x}_n \end{bmatrix}_{r \times z} \\ \text{output ring buffer} = \mathbf{T}_1 &= \begin{bmatrix} CA50_{n-r+2} \\ \vdots \\ CA50_{n+1} \end{bmatrix}_{r \times 1} \end{aligned} \quad (27)$$

Then execute the WR-ELM update algorithm between combustion cycle $n + 1$ and $n + 2$ as shown in Fig. 4. For the combustion implementation, r was taken to be 8 cycles after tuning with existing datasets. If desired, r can vary cycle-to-cycle.

8. As with the offline data, build $\mathbf{H}_1(\mathbf{a}, \mathbf{x}_1)$ with Eq. 4 using an input matrix $\mathbf{x}_1$ and output target vector $\mathbf{T}_1$. Specify a weight matrix $\mathbf{W}_1$. For the combustion implementation the identity matrix ($\mathbf{W}_1 = \mathbf{I}$) was chosen since weighting was already applied to the offline data in step 5. Gradually increased weighting on the most recent time steps in the ring buffer was explored; however, it did not net a significant improvement to model fitting performance over a simple scalar value on offline data. Although not explored in the current implementation, $\mathbf{W}_1$ can vary cycle-to-cycle.
9. Solve for the updated β_1 solution using Eqs. 25.
10. After cycle $n + 1$'s input vector $\mathbf{x}_{n+1}$ is fully populated, transform vector into $\mathbf{H}_{n+1}$ using Eq. 4 and solve for a predicted target value $\mathbf{T}_{n+1}$ or $CA50_{n+2}$ using Eq. 26.
11. Repeat steps 7-10 for each new time step, caching results (e.g. hidden layer outputs) from previous time steps to reduce computational requirements.

2.5. Real-Time Implementation

A collection of unoptimized MATLAB[®] software routines was developed using the techniques described in the previous sections. The offline solution provided by Eqs. 24 was solved at an average rate of $1.1 \mu s$ per combustion cycle per cylinder on an Intel[®] i7 860 2.8 GHz desktop

computer running Gentoo GNU/Linux[®]. The online predictions from Eqs. 25, 27, and 26 were recast into a *parfor* loop that automatically parallelized the code across four worker threads to provide predictions at an average rate of $66 \mu s$ per combustion cycle per cylinder. This level of performance is more than adequate for real-time.

Although algorithm development is the main focus of this paper, a real-time implementation of the WR-ELM algorithm has been built using custom, 18-bit Raspberry Pi[®] data acquisition hardware. A two-minute video demonstrating both predictions and control is available at [25]. The software for this system is comprised of:

- A PREEMPT_RT patched Linux[®] kernel with minor patches to completely disable Fast Interrupt request (FIQ) usage by the USB driver.
- ARM assembly code for high-speed pressure data acquisition using the FIQ (up to ~ 240 kilosamples per second, total all cylinder channels).
- A C code Linux[®] kernel module that contains the FIQ assembly code with page-based memory allocation and standard mmap, ioctl, and fasync hooks to Linux[®] user space.
- A multi-threaded Linux[®] user space application that runs heat release calculations and WR-ELM. This software leverages the Eigen C++ matrix library and a custom assembly code matrix multiply³ for the Raspberry Pi 1's VFPv2 double precision floating point unit. Asynchronous fasync notification is used for specific crank angle events published from the FIQ code to synchronize the user space software's execution against the crank's rotation.
- A low-priority second user space thread that uses standard WebSockets from the libwebsockets C code library to stream processed data to a web-based user interface coded in JavaScript with the d3.js library.
- A minimal Raspbian GNU/Linux[®] distribution [26].

After the adaptation routine is run, it is possible to perform 11 model predictive control predictions within a worst case task context switch and calculation latency window of $\sim 300 \mu s$. This level of real-time performance is needed to ensure control authority with the *SOI* actuator immediately after P_{NVO} is measured.

2.6. Experimental Setup

Table 1 provides a summary of the experimental setup and conditions visited. In-cylinder pressure was acquired on a $1.0^\circ CA$ basis and pegged thermodynamically for each cycle after IVC using a polytropic exponent of 1.35. This

³This code was benchmarked to be 29% faster than OpenBLAS's VFPv2 GEMM and 126% faster than Eigen C++.

exponent was chosen to most closely match the pegging results achieved using the single intake runner high speed pressure sensor on cylinder 1. For the purpose of computing cycle-to-cycle net Indicated Mean Effective Pressure (IMEP), a cycle was defined as starting at 360 °BTDC firing and ending at 359 °ATDC firing. The reference cam lift for timing and duration in Table 1 is 0.5 mm. The air-fuel ratio range indicated in Table 1 was measured post-turbine, and represents a mixture from all four cylinders. Fuel mass per cycle was *estimated* using the fuel’s lower heating value, assuming that the gross heat release was 20% greater than the net heat release, and that the combustion efficiency was 100%.

2.7. Dataset Description

The full collection of 129,964 cycles is comprised of five ~20 minute random test subsequences. Each random subsequence covers the same nominal ranges listed in Table 1; however, one subsequence holds *SOI* fixed. The sequence with *SOI* fixed is only used as part of Fig. 2, and not during the model training and testing presented here. Total

cycle counts are reported after outliers are removed. The outlier criteria (detailed in [6]) are intended to remove misfires and partial burns. These criteria are fairly permissive and remove only ~3% of the data.

The offline solution was trained using ~40 minutes of test cell time covering 53,884 cycles and 1,179 random engine set points at 2,500 rpm (two random subsequences). These subsequences are comprised of random, transient set point steps occurring approximately every 0.5 - 10 seconds with occasional misfires covering the nominal variable ranges given in Table 1. The training data were pruned to only include 6 cycles before and 9 cycles after a transient set point step for a small model fitting performance improvement. The online solution was run with a separate random subsequence and fed unseen cycles one-by-one, similar to what would be experienced in a real-time implementation. This online dataset is comprised of 25,323 consecutive cycles with 521 random engine set points. Longer online sequences were also tested, and achieved similar results.

3. Results and Discussion

The fitting performance on a 25,323 cycle dataset (excluding outliers as defined in [6]) is shown in Table 2 and in Figs. 5, 6, and 7. The minimum coefficient of determination (R^2) given in Table 2 shows that at least 80% of the cycle-to-cycle variance can be explained by the model as it is currently defined for a dataset with random transient steps occurring approximately every 0.5 - 10 seconds and occasional misfires. This is better than the 76% achieved with ϵ -Support Vector Regression (ϵ -SVR) on the same dataset in [6]. However, this is not a 1:1 comparison because WR-ELM is fully predicting the entire 25,323 cycle dataset, whereas ϵ -SVR’s training strategy ensured the data-driven model had partially seen the operating points it was trying to predict. Steady-state Root Mean Squared Error (RMSE) in Table 2 was assessed at a single set point with a mean CA50 of 3.9 °ATDC and a net Indicated Mean Effective Pressure (IMEP) of 2.8 bar before the transient sequence started.

Fig. 5 shows the distribution of model errors. It is clear that there is a slight positive bias to the predictions. Fig. 6 provides insight into the tails of Fig. 5 and shows

Table 1: Experimental setup and test conditions

Engine	
Make / model	GM / LNF Ecotec
Cylinder layout	in-line 4
Overall displacement	2.0 L
Bore / stroke	86 / 86 mm
Geometric compression ratio ^a	11.2 : 1
Cam lift ^a	3.5 mm
Cam duration ^a	93 °CA
Cam phaser type	hydraulic
Fuel injector type	direct, side mounted, wall guided
Fuel	
Designation	Haltermann HF0437, EPA Tier II EEE
Description	U.S. Federal Emission Cert. Gasoline
Research Octane Number	97.0
Motor Octane Number	88.1
ASTM D240 heating value	42.8 MJ / kg
Aromatic / olefin / saturate fractions	28 / 1 / 71 % volume
Test conditions	
Throttle position	wide open
Turbocharger	wastegate open
Supercharger	bypassed
Residual retention strategy	negative valve overlap
I/O set point range ^b	78.6 / 128 °ATDC
EVC set point range ^b	-118 / -83.0 °ATDC
<i>SOI</i> set point range ^b	272 / 378 °BTDC
<i>TI</i> set point range ^b	0.582 / 1.01 ms
Net IMEP values visited ^b	1.85 / 3.62 bar
Air-fuel ratios visited ^b	0.90 / 1.6
Estimated fuel per cycle ^b	6 / 11 mg
Intake runner temps., all cyls.	$\mu = 52.6$ °C, $\sigma = 1.6$ °C
Fuel injection pressure	$\mu = 70.0$ bar, $\sigma = 0.85$ bar
Coolant temperature	$\mu = 89.5$ °C, $\sigma = 3.4$ °C
Engine speed	$\mu = 2,500$ RPM, $\sigma = 6$ RPM

^a Modified from stock engine

^b First to 99th percentile

Table 2: WR-ELM model of Eq. 2 error statistics

Cylinder #	Overall ^c R^2	Overall RMSE [°CA]	Steady-State RMSE [°CA]
1	0.81	1.85	0.84
2	0.81	2.06	0.97
3	0.80	2.17	0.97
4	0.83	1.64	0.86

^c 25,323 consecutive cycles with random transient steps occurring approx. every 0.5 - 10 sec. and occasional misfires.

Figure 5: Error histograms for WR-ELM model of Eq. 2 across 25,323 consecutive cycles with random transient steps occurring approx. every 0.5 - 10 seconds and occasional misfires.

Figure 6: Predicted versus measured WR-ELM model of Eq. 2 across 25,323 consecutive cycles with random transient steps occurring approximately every 0.5 - 10 seconds and occasional misfires. Late combustion timing is under predicted, but almost all prediction outliers capture the correct directionality.

that model errors still generally capture the correct directionality. Fig. 6 also shows that late combustion timing is under predicted and that the positive bias is largely from the midrange values of CA50. Fig. 7a-d shows the cycle-to-cycle time series predictions, which can be computed as early as 358 °BTDC firing. Missing segments are the outliers described earlier. Fig. 7e-h provide qualitative insight into the model's β_1 weights under online adaptation. The

neurons are sorted by the 2-norm of their respective input weight vector $\mathbf{a}_i$. The same non-linear transformation specified by $\mathbf{a}$ is used for each cylinder, and any cylinder-to-cylinder differences in the cycle-to-cycle β_1 are due to different characteristics of each cylinder. Fig. 7i shows the IMEP and Fig. 7j-l shows the random engine actuator inputs that are driving each engine set point transient.

The model predictions of Fig. 7a-d generally show good agreement; however, there are occasional tracking errors. It is unclear what the source of these tracking errors is (e.g. is it a fundamental model limitation, the need for more inputs,⁴ the influence of other misfiring cylinders going through a harsh reignite, the need for more weight tuning or offline training data, or perhaps something else?). Future work will try to answer these questions. Overall, however, the authors believe the level of fit shown in Table 2 and in Figs. 5, 6, and 7 is very good considering that the dataset includes both transients and operating points with high CV, right up to complete engine misfire.

4. Summary and Conclusion

This work presents a new online adaptation algorithm named Weighted Ring - Extreme Learning Machine. The approach uses a weighted ring buffer data structure of recent measurements to recursively update an offline trained Extreme Learning Machine solution. It is shown that WR-ELM can be used to approximate the combustion mapping function developed in [6] and provide reasonably accurate, causal predictions of near chaotic combustion behavior. In the combustion application only in-cylinder pressure and crank encoder sensors are needed for predictions, and these predictions can be computed as early as 358 °BTDC firing. The algorithm is fast, and has been implemented in real-time on the low-cost Raspberry Pi® platform (a two-minute video demonstrating this is available at [25]).

Future work will explore optimal selection of weight(s) and try to better understand the situations that lead to the occasional model tracking errors. Finally, the broader objective of this new modeling approach is to enable a new class of cycle-to-cycle model predictive control strategies that could potentially bring HCCI's low engine-out NO_x and reduced CO_2 emissions (higher fuel efficiency) to production gasoline engines.

Acknowledgments

This material is based upon work supported by the Department of Energy [National Energy Technology Laboratory] under Award Number(s) DE-EE0003533. This work is performed as a part of the ACCESS project consortium (Robert Bosch LLC, AVL Inc., Emitec Inc., Stanford University, University of Michigan) under the direction of PI

⁴Adding additional inputs is not necessarily practical computationally, experimentally, or even advisable given Occam's razor.

Figure 7: The WR-ELM model of Eq. 2 can track CA50 through transients every 0.5 - 10 seconds, operating points with near chaotic high CV, and at steady-state during a particularly harsh region of the 25,323 cycle dataset that includes misfires. The colormaps (linearly scaled) provide qualitative insight into the level of cycle-to-cycle adaptation and into cylinder-to-cylinder model differences.

Hakan Yilmaz and Co-PI Oliver Miersch-Wiemers, Robert Bosch LLC.

The authors thank Vijay Janakiraman for providing the raw data analyzed in this paper. The authors also thank Jeff Sterniak for both his test cell and project support. A. Vaughan thanks Nisar Ahmed for many helpful discussions, and his advisors S. V. Bohac and Claus Borgnakke for the freedom to explore the interesting topic of near chaotic combustion.

References

- [1] A. Vaughan, S. V. Bohac, An Extreme Learning Machine Approach to Predicting Near Chaotic HCCI Combustion Phasing in Real-Time, arXiv pre-print (1310.3567), 2013.
- [2] A. Babajimopoulos, V. Challa, G. Lavoie, D. Assanis, Model-Based Assessment of Two Variable CAM Timing Strategies for HCCI Engines: Recompression vs. Rebreathing, in: Proceedings of the ASME Internal Combustion Engine Division 2009 Spring Technical Conference, ICES2009-76103, 2009.
- [3] E. A. O. Soto, J. Vavra, A. Babajimopoulos, Assessment of Residual Mass Estimation Methods for Cylinder Pressure Heat Release Analysis of HCCI Engines With Negative Valve Overlap, Journal of Engineering for Gas Turbines and Power 134 (8) (2012) 082802.
- [4] L. M. Olesky, J. Vavra, D. Assanis, A. Babajimopoulos, Effects of Charge Preheating Methods on the Combustion Phasing Limitations of an HCCI Engine With Negative Valve Overlap, Journal of Engineering for Gas Turbines and Power 134 (11) 112801.
- [5] S. Saxena, I. D. Bedoya, Fundamental phenomena affecting low temperature combustion and HCCI engines, high load limits and strategies for extending these limits, Progress in Energy and Combustion Science 39 (5) (2013) 457 – 488.
- [6] A. Vaughan, S. V. Bohac, A Cycle-to-Cycle Method to Predict HCCI Combustion Phasing, in: Proceedings of the ASME Internal Combustion Engine Division 2013 Fall Technical Conference, ICEF2013-19203, 2013.
- [7] J. S. Lacey, The effects of advanced fuels and additives on homogeneous charge compression ignition combustion and deposit formation, Ph.D. thesis, University of Michigan, 2012.
- [8] A. Vaughan, G. Delagrammatikas, A High Performance, Continuously Variable Engine Intake Manifold SAE Paper 2011-01-0420, 2011.
- [9] Understanding the transition between conventional spark-ignited combustion and HCCI in a gasoline engine, Proceedings of the Combustion Institute 31 (2) (2007) 2887 – 2894.
- [10] E. Hellström, A. G. Stefanopoulou, L. Jiang, Cyclic Variability and Dynamical Instabilities in Autoignition Engines with High Residuals, IEEE Transactions on Control Systems Technology 21 (5) (2013) 1527–1536.
- [11] J. C. Kantor, A Dynamical Instability of Spark Ignited Engines, Science 224 (4654) (1984) 1233–1235.
- [12] L. Manofsky, J. Vavra, D. Assanis, A. Babajimopoulos, Bridging the Gap between HCCI and SI: Spark-Assisted Compression Ignition SAE Paper 2011-01-1179, 2011.
- [13] M. Sjöberg, J. E. Dec, A. Babajimopoulos, D. Assanis, Comparing Enhanced Natural Thermal Stratification Against Retarded Combustion Phasing for Smoothing of HCCI Heat-Release Rates SAE Paper 2004-01-2994, 2004.
- [14] M. Mehl, W. J. Pitz, M. Sjöberg, J. E. Dec, Detailed Kinetic Modeling of Low-Temperature Heat Release for PRF Fuels in an HCCI Engine SAE Paper 2009-01-1806, 2009.
- [15] Experimental and surrogate modeling study of gasoline ignition in a rapid compression machine, Combustion and Flame 159 (10) (2012) 3066 – 3078.
- [16] G. M. Shaver, Physics-Based Modeling and Control of Residual-Affected HCCI Engines using Variable Valve Actuation, Ph.D. thesis, Stanford University, 2005.
- [17] J. Livengood, P. Wu, Correlation of autoignition phenomena in internal combustion engines and rapid compression machines, Symposium International on Combustion 5 (1) (1955) 347–356.
- [18] S. Jade, J. Larimore, E. Hellström, A. Stefanopoulou, L. Jiang, Controlled Load and Speed Transitions in a Multicylinder Recompression HCCI Engine, Control Systems Technology, IEEE Transactions on PP (99) (2014) 1–1.
- [19] V. Cherkassky, F. Mulier, Learning from Data: Concepts, Theory, and Methods, Wiley, 2007.
- [20] N.-Y. Liang, G.-B. Huang, P. Saratchandran, N. Sundararajan, A fast and accurate online sequential learning algorithm for feedforward networks, Neural Networks, IEEE Transactions on 17 (6) (2006) 1411–1423.
- [21] G.-B. Huang, Q.-Y. Zhu, C.-K. Siew, Extreme learning machine: Theory and applications, Neurocomputing 70 (1-3) (2006) 489 – 501.
- [22] B. Mirza, Z. Lin, K.-A. Toh, Weighted Online Sequential Extreme Learning Machine for Class Imbalance Learning, Neural Processing Letters (2013) 1–22.
- [23] D. Simon, Optimal State Estimation: Kalman, H Infinity, and Nonlinear Approaches, Wiley-Interscience, 2006.
- [24] John Halleck's Matrix Identities (also Ring Identities), URL: <http://www.cc.utah.edu/~nahaj/math/matrix-identities.html>, [Online; accessed 12/31/2014], 2014.
- [25] Version 0.1 Raspberry Pi Engine Control with Real-Time Adaptive Extreme Learning Machine, URL: <https://www.youtube.com/watch?v=qQG70cnE3EA>, [Online; accessed 1/8/2015], 2015.
- [26] Minimal Raspbian unattended netinstaller, URL: <https://github.com/debian-pi/raspbian-ua-netinst>, [Online; accessed 9/5/2014], 2014.

DISCLAIMER: This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.