

Data-Driven Modelling of Wind Turbines

Gijs van der Veen, Jan-Willem van Wingerden and Michel Verhaegen

Abstract—In this paper we present a novel approach that allows global modelling of the power control dynamics of a wind turbine based on measured data. The approach is based on the assumption that all the nonlinearities over the operating range arise from a static aerodynamic mapping, which is interconnected with linear, time-invariant dynamics. This so-called Hammerstein structure is exploited to simplify the model identification procedure. The global model is suited to control design methods such as model predictive control or can be used to extract local linear models. The approach is demonstrated on a benchmark example, the 5MW NREL/Upwind reference turbine and is shown to work well. Tools from convex optimisation and the recently introduced nuclear norm techniques prove to be instrumental to the successful implementation of the algorithms.

I. INTRODUCTION

Current controllers for power regulation in wind turbine applications are all based on PID feedback or related methods. It is common practice to design such controllers based on linear models [1], [2]. To deal with the time-varying dynamics of the system, gain scheduling is often employed on the controller gains by which the controller is tuned in several discrete operating points and the gains are interpolated during operation. In this paper we introduce a procedure to identify a globally valid model of the dynamics of a wind turbine involved in power production [3]. Such a model can be used to extract local-linear models in several operating points, but also to design a controller that deals intrinsically with the time-varying dynamics. Furthermore, such a model has value in fault detection applications and it results in an estimate of the true torque curve of a wind turbine.

In the wind energy community, linear models are usually obtained from aeroelastic models or, more recently, derived from measurement data [4], [5], with the advantage of not having to rely on the theoretical approximations made in first-principles models. In recent years, advances have been made in identification of linear parameter-varying models of wind turbines [6]. Such methods, however, are still severely hampered by computational complexity issues. For this reason, we introduce a method to identify a global model of a wind turbine that can easily cope with large amounts of sampled data. Whereas pure LTI identification of wind turbines from operational data is made difficult by a constant need to maintain a steady operating point, this is no longer necessary for the proposed method, which is another

advantage.

The proposed identification method basically involves a two-step procedure. In a first step, a theoretical prediction of the aerodynamic characteristics (such as the torque coefficient) is used to choose a suitable set of radial basis functions that can be used to describe it. In a second step, standard LTI system identification is performed to retrieve the linear, time-invariant portion of the dynamics and the coefficients of the basis functions.

First, the problem definition is introduced in section II. Then, in sections III and IV the method is described in more detail. In section V the method is applied to a simulation example. The paper ends with a number of conclusions and suggestions for future work.

II. PROBLEM STATEMENT

The power production control loop of a wind turbine is commonly designed around the dynamics governed by the pitch angle β , rotor speed Ω and generator torque T_g . This is shown schematically in Fig. 1. The rotor speed can be controlled by adjusting the demanded generator torque and by adjusting the pitch angles of the blades.




Fig. 1. Schematic showing a typical model of a wind turbine between pitch angle β , generator torque T_g and rotor speed Ω .




Fig. 2. Typical shape of the torque coefficient $C_Q(\lambda, \beta)$.

This work was supported by Vestas Wind Systems A/S
The authors are with the Delft Center for Systems and Control, Delft University of Technology, 2628 CD, The Netherlands {G.J.vanderVeen, J.W.vanWingerden, M.Verhaegen}@tudelft.nl

Referring to Fig. 1, the aerodynamic torque T_a on a wind turbine under steady operating conditions can be described in terms of the torque coefficient $C_Q(\lambda, \beta)$ which appears in Eq. (1).

$$T_a = f(\lambda, \beta, V^2) = \frac{1}{2} \rho \pi R^3 C_Q(\lambda, \beta) V^2, \quad (1)$$

where T_a is the aerodynamic torque, λ is the tip speed ratio, β is the pitch angle, ρ is the air density, V is the wind speed, R is the rotor radius and $C_Q(\lambda, \beta)$ is a smooth surface (Fig. 2) depending on the pitch angle β and the tip-speed ratio $\lambda = (\Omega R)/V$. The aerodynamic torque further depends on the wind speed V , the rotor radius R and the air density ρ . It is clear that the dynamics of the wind turbine subject to the aerodynamic torque in Eq. (1) are non-linear.

We wish to identify a non-linear dynamic model of this system on the basis of input and output measurements obtained from the system operating in closed-loop, based on a set of measurements $\{u_k, y_k\}_{k=1}^N$. The system to be identified is assumed to admit an innovation state-space representation given by:

$$x_{k+1} = Ax_k + B_1 f(w_k) + B_2 u_k + K e_k, \quad (2a)$$

$$y_k = Cx_k + e_k. \quad (2b)$$

The vectors $x_k \in \mathbb{R}^n$, $u_k \in \mathbb{R}^m$, $y_k \in \mathbb{R}^l$ and $e_k \in \mathbb{R}^l$ are the state sequence, input, output and innovation, respectively. The matrices A , B_1 , B_2 and C are of compatible dimensions. K is a Kalman gain. The innovation sequence e_k is an ergodic zero-mean white noise sequence with covariance matrix $\mathbb{E}\{e_j e_k^T\} = W \delta_{jk}$, with $W \succ 0$. The vector $w_k \in \mathbb{R}^3$ represents the inputs affecting the aerodynamic torque T_a through a memoryless non-linear mapping $f(\cdot) : \mathbb{R}^3 \rightarrow \mathbb{R}$ (in this case λ , β and V). We note that the nonlinearity is only determined up to a nonzero scalar multiplication without affecting the input/output behaviour of the system, that is, we may estimate $\hat{f}(\cdot) \approx T^{-1} f(\cdot)$ instead of $\hat{f}(\cdot) \approx f(\cdot)$. Further note that the linear and nonlinear inputs to the LTI system have been separated. A causal model is enforced by disallowing a direct feedthrough component.

The non-linear mapping $f(\lambda, \beta, V^2)$ defined in Eq. (1) will be parametrised in terms of basis functions $\varphi_i(\lambda, \beta) : \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$, defined by

$$f(\lambda, \beta, V^2) = \sum_{i=1}^{n_b} \alpha_i \varphi_i(\lambda, \beta) V^2,$$

where n_b denotes the number of basis functions and $\alpha_i \in \mathbb{R}$ are scalar parameters. These parameters will be estimated, while the basis functions are chosen a priori. In fact, the parametrisation of the non-linearity in terms of basis functions can be viewed as an over-parametrisation of the input, that is, a higher dimensional input $\bar{w} \in \mathbb{R}^{n_b}$ can be defined according to

$$\bar{w} = \begin{bmatrix} \varphi_1(\lambda, \beta) \\ \varphi_2(\lambda, \beta) \\ \vdots \\ \varphi_{n_b}(\lambda, \beta) \end{bmatrix} V^2. \quad (3)$$

Then, a new input matrix can be defined as

$$B_1 \equiv B_1 [\alpha_1 \quad \alpha_2 \quad \cdots \quad \alpha_{n_b}] \equiv B_1 \alpha^T. \quad (4)$$

The over-parametrised input signals arising from the non-linearity will serve as the linear inputs for the identification method.

III. MODELLING THE STATIC TORQUE COEFFICIENT

In this section, we start by considering the static model of the torque coefficient in Eq. (1). Once we succeed in approximating this nonlinear static mapping in terms of a set of 2-dimensional basis functions, we will proceed to the dynamic behaviour of the turbine.

A. Selection of basis functions

In general, the choice of suitable basis functions to model a nonlinear function is not at all trivial. A commonly used function is the Gaussian radial basis function:

$$\varphi_i(x) = e^{-\|\Sigma^{\frac{1}{2}}(x-c_i)\|_2^2}, \quad (5)$$

where x (here, a 2-vector) is the coordinate where the function is evaluated and c_i is the i -th radial basis centre. The diagonal matrix Σ is used to control the shape of the basis functions in the different coordinate directions. The parameters to be designed a priori are the number of basis functions n_b , their centres c_i and the shape parameter Σ .

Usually, an aeroelastic model of a wind turbine is available. Using such a model or a blade element momentum approach, one can usually obtain the theoretical surface $C_Q(\lambda, \beta)$. A typical C_Q -surface is shown in Fig. 2. This theoretically obtained surface will now be used to determine a candidate set of basis functions to allow the modelling of the true surface.

To determine the ranges of the inputs β and λ , data was acquired through simulations for typical operating conditions under varying wind speeds. In Fig. 3, a typical example of such a trajectory of β and λ is shown. The area spanned by the ranges of β and λ is then gridded using a uniform grid as also shown in Fig. 3. It is crucial that we end up with a limited number of basis functions to limit computational complexity and numerical ill-conditioning. A large number of $N = 1000$ random points within operating area was generated. Using the uniformly spaced set of basis functions a least-squares fit of the surface was obtained by solving the following least-squares problem:

$$\alpha = \arg \min_{\alpha} \sum_{k=1}^N \|C_Q(\lambda_k, \beta_k) - \alpha \varphi(\lambda_k, \beta_k)\|_2^2, \quad (6)$$

where N is the number of points available for fitting and $C_Q(\lambda_k, \beta_k)$ is the theoretical torque coefficient evaluated at (λ_k, β_k) . Each vector $\varphi(\lambda_k, \beta_k)$ contains the basis functions evaluated at the centres 1 through n_b . At this point, the shape parameter Σ and the spacing of the basis functions are tuned to achieve a good accuracy of fit. That is, while the number of RBFs is kept fixed, the elements of Σ are adjusted until

the accuracy is highest. Note that this could also be solved as a nonlinear least-squares problem, as done in e.g. [7]. To achieve a compact basis, only basis functions close to the operational data were generated. An alternative way to reduce the size of the basis is to add the ℓ_1 norm of the coefficients, $\kappa\|\alpha\|_1$, to the criterion in Eq. (6), see e.g. [8, Ch. 3]. Depending on the magnitude of κ , this will force some of the coefficients to zero, allowing the corresponding basis functions to be discarded. A reduced basis results in better conditioning and less sensitivity to variations in the data. Applying the procedure outlined here to torque curve data of the 5 MW NREL reference turbine resulted in the following parameters:

$$n_b = 21, \quad \Sigma = \text{diag}(0.128, 0.016).$$

On average, a variance-accounted-for[9] (VAF) of 97% was obtained. Here, the VAF¹ is defined as 100% minus the relative mean-square error, which corresponds to the coefficient of determination (R^2) in the literature on statistics. A more accurate fit could be obtained with a more dense basis, but this compact basis was favoured in view of computational complexity issues. The result is a compact basis, as shown in Fig. 3. Figure 4 shows the distribution of unit radial basis functions. Validation on an independent set of randomly generated points on the curve confirms that the surface can be modelled accurately using this set of basis functions.




Fig. 3. Radial basis function centers. The dashed line indicates the convex hull of the operating points.




Fig. 4. The distribution of radial basis functions.

¹VAF = $\max \left\{ 0, \left(1 - \frac{\text{var}(y_k - \hat{y}_k)}{\text{var}(y_k)} \right) \times 100\% \right\}$, where \hat{y} is the prediction.

IV. NONLINEAR SYSTEM IDENTIFICATION

It is assumed that a set of input and output data $\{u_k, w_k, y_k\}_{k=1}^N$ has been obtained through measurements. In this case the rotor speed Ω is considered as output:

$$y_k = \Omega_k.$$

The inputs to the true system are λ , β , V and T_g . However, based on the expansion of the nonlinearity in the previous section, the inputs for identification are defined by the elements of $\bar{w}_k = \varphi(\lambda_k, \beta_k)V_k^2$ (Eq. 3) and $T_{g,k}$, so that

$$\tilde{u}_k = \begin{bmatrix} \varphi(\lambda_k, \beta_k)V_k^2 \\ T_{g,k} \end{bmatrix},$$

can be defined as the input for the system identification procedure, where the coefficients α have been absorbed into the system to be identified. This results in an input vector $\tilde{u}_k \in \mathbb{R}^{n_b+1}$. In fact, the procedure in the previous section allows one to now consider a MIMO identification problem of a linear, time-invariant system. Here, we use the closed-loop MOESP subspace identification method [10], [11] to estimate the state-space system matrices. In fact, any LTI identification method suitable for MIMO closed-loop identification could be applied here. We favour the subspace identification method here for its numerical robustness and efficiency. Note that considering a closed-loop procedure is necessary, since the rotor-speed signal that determines the tip-speed ratio is also an output of the system. The system identification procedure will not be described here in more detail but we mention here that it results in a set of state-space matrices (A, B_1, B, C, K) defining the system as described above in Eqs. (2,4). The model class that can be identified using the CL-MOESP method encompasses all innovation state-space models. ‘‘Model selection’’ is effected by selecting a system order based on the order information that is typically conveyed by the class of subspace methods (in the singular value decomposition step) and the choice is verified on the basis of obtained VAF values.

A. Recovering the nonlinearity

After identification, an estimate of the matrix B_1 (Eq. (4)) is available. Let \bar{W} denote a matrix in which all available inputs \bar{w}_k (Eq. (3)) as function of time index k are stacked next to each other. Then the following product has rank 1 according to the definition in (Eqs. (3-4)):

$$\text{rank}(B_1 \bar{W}) = 1. \quad (7)$$

An SVD can be performed on this product:

$$B_1 \bar{W} = [U \quad U_\perp] \begin{bmatrix} \Sigma & 0 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} \mathcal{V}^T \\ \mathcal{V}_\perp^T \end{bmatrix}. \quad (8)$$

The actual B_1 matrix relating to the nonlinear input can then be retrieved as $B_1 \approx U\Sigma$, whereas the coefficients of the nonlinearity are found to be $\alpha = \mathcal{V}^T \bar{W}^+$, where \bar{W}^+ denotes the pseudoinverse of \bar{W} .

1) *Retrieving the physical signal T_a* : It was already mentioned that the nonlinear function can only be estimated upto an unknown scaling, since such a scaling cancels in the input-output behaviour of the system. However, one can make use of some knowledge about the system to retrieve this scaling approximately. Based on a power balance during steady operation (constant speed), the following relation must hold:

$$\frac{T_a}{T_g} = \frac{G}{\eta},$$

where G is the gearbox ratio and η is the gearbox efficiency. Since in this steady condition the rotor speed does not change, this ratio determines the relative static gains of the inputs T_a and T_g to the linear part of the system. Thus, knowing the ratio, the system and nonlinearity can both be scaled such the magnitude of input T_g corresponds to the actual physical rotor torque. The accuracy is limited by how well the static gain of the true system has been estimated and the accuracy of the gearbox efficiency. A key result is that we can effectively estimate the torque curve of a turbine from measurement data.

2) *Improving the rank condition*: In general there is no guarantee that the low-rank condition Eq. (7) is satisfied in practice with a finite amount of data. Noise and model mismatch also contribute to this effect. In subspace identification methods, the system matrices and the initial state x_0 are obtained from a least-squares problem. Here, the example of MOESP subspace identification is used, but in fact the procedure applies to most subspace methods. In order to enforce a low rank of the identified matrix \mathcal{B} , the following regularised least-squares regression is performed:

$$\arg \min_{\mathcal{B}, \mathcal{B}_2, K, x_0} \left\| \underbrace{Y - S \begin{bmatrix} \text{vec}(\mathcal{B}_1) \\ \text{vec}(\mathcal{B}_2) \\ \text{vec}(K) \\ x_0 \end{bmatrix}}_{\| \cdot \|_2} \right\|_2^2 + \lambda \|\mathcal{B}_1\|_* \quad (9)$$

The underbraced term represents the standard least-squares problem that is solved in MOESP subspace identification (for details consult [10]), where Y and S depend on the measured data. The additional term represents the *nuclear norm* [12] of \mathcal{B}_1 . The nuclear norm is defined as:

$$\|\mathcal{B}_1\|_* \equiv \sum_{i=1}^{n_\sigma} \sigma_i(\mathcal{B}_1),$$

and can therefore be used as a heuristic for the rank [12] of \mathcal{B}_1 . Eq. (9) is a convex problem and can be solved using a tailored solver within the CVX framework [13]. Recently, we have been able to improve the result further by explicitly parameterizing the low-rank matrix product and performing a separable least-squares regression on \mathcal{B}_1 , using the nuclear norm solution as initial condition for the nonlinear optimisation. The main advantage is that, since the obtained solution \mathcal{B}_1 is already of low rank, the truncated SVD in Eq. (8) does not result in further loss of information.

V. IDENTIFYING THE POWER CONTROL LOOP OF THE 5 MW NREL REFERENCE TURBINE

A. Identification design

To be able to identify an accurate model from input-output data under turbulent wind conditions, the problem of identification design is important [14], [9]. The excitation signal should be such that it excites all the relevant modes of the system while observing the system's limitations, yet it should provide sufficient excitation to result in a satisfactory signal-to-noise ratio. Common identification signals are white noise signals, broadband multisine signals and pseudo-random binary sequences. The latter type has the advantage of being strictly limited in amplitude, while delivering maximal signal energy to the system within the amplitude constraints. Furthermore, its low-frequency content can be emphasised without violating amplitude constraints by simply modifying its sample rate.

In the current example, excitations have been superimposed on the torque and pitch references. For the torque and pitch references, a binary signal of amplitude 1200 Nm with a sampling time of 0.1 s and a binary signal of amplitude 0.5° with a sampling time of 0.5 s were used, respectively. At this stage, a uniform time-varying wind field was used with a hub wind speed that varied between 5 and 23 m/s, see Fig. V-A. The wind speed consists of a step change every 60 seconds and a stochastic component. The stochastic component is represented by a zero-mean white noise with $\sigma = 2.5$ low-pass filtered at 0.05 Hz. Next steps will include a full 3D turbulent wind field.




Fig. 5. Time history of the hub wind speed.

B. Practical aspects

The procedure outlined in previous sections suggests that a data set encapsulating the entire operating range of the turbine should be available. In practice, this is not achievable in a single experiment due to unpredictable temporal variations in wind conditions. Therefore it is emphasized that the subspace identification framework used here can deal with multiple batches of measurement data in a straightforward way. Data sets can easily be concatenated, and the only technical implication is that for each data set a distinct corresponding initial condition must be identified.

C. Results

To test the performance of the devised method, a model of the 5 MW NREL/UPWIND reference turbine implemented in the aeroelastic environment of Bladed[®] [15] was used.

Based on the wind speed trajectory and a sampling frequency of 20 Hz, 12400 samples (≈ 10 min) of data

were obtained. As a first step in the subspace identification approach, order detection was performed. Based on this order information and trying several model orders, the identification was carried out for an order of $n = 8$.

The resulting prediction of the output is shown in Fig. 6. As a quality measure, the variance-accounted-for [9] (VAF) was again used, which gives a measure of how well the variability of the output signal is predicted by the model. A VAF of around 99% was obtained. Ideally, one would simulate the identified system with a fresh dataset to verify that the model has not been fitted to the noise realisation of the training data. However, since the system is marginally stable (an integrator is present which corresponds to the rigid-body mode of the drivetrain), any small deviations will accumulate and the simulations slowly diverge. Nevertheless, the signals did show the same dynamic behaviour when this was attempted. A solution is to use the obtained model as an observer for the true system.




Fig. 6. The predicted output (blue) based on the identification data.

Figure 7 shows the estimated nonlinear function. Comparing this figure with the theoretical curve in Fig. 2 a reasonable correspondence is observed. It is emphasised that the actual local behaviour of the system is characterised by the gradient of the torque curve, so the gradient in this case is more crucial to the model accuracy than the absolute value, which merely results in a static offset of the local equilibrium (operating point).




Fig. 7. The estimated torque coefficient $\hat{C}_Q(\lambda, \beta)$.

D. Validation

As a means of validating the obtained nonlinear Hammerstein model of the wind turbine, we have already shown

the simulated outputs on the training dataset in Fig. 6. A second way of validating the model is by linearising the nonlinear model in several operating points and comparing the local linear models to those obtained from the linearisation feature in Bladed. Based on the nonlinear model description (Eq. 2), we can linearise around some steady-state $(\bar{x}, \bar{w}, \bar{u}) = (\bar{x}, \bar{\lambda}, \bar{\beta}, \bar{V}, \bar{T}_g)$ as follows:

$$x_{k+1} \approx [A\bar{x} + Bf(\bar{u})] + A(x_k - \bar{x}) + B_1 \left. \frac{\partial f}{\partial w} \right|_{w=\bar{w}} (w_k - \bar{w}) + B_2(u_k - \bar{u}) + Ke_k.$$

If we redefine $x_k \leftarrow x_k - \bar{x}$, $w_k \leftarrow w_k - \bar{w}$ and $u_k \leftarrow u_k - \bar{u}$, we obtain the local linear model as:

$$x_{k+1} = Ax_k + B_1 \left. \frac{\partial f}{\partial w} \right|_{w=\bar{w}} w_k + B_2 u_k + Ke_k$$

$$y_k = Cx_k + e_k.$$

To be able to evaluate the partial derivatives of the nonlinear function, we need to calculate the partial derivatives of the radial basis functions (Eq. 5). To this end, we have the following result:

$$\frac{\partial f(w)}{\partial w_j} = \sum_{i=1}^{n_b} \alpha_i \frac{\partial \varphi_i(w)}{\partial w_j}, \quad \text{where}$$

$$\frac{\partial \varphi_i(w)}{\partial w_j} = -2e^{-\|\Sigma^{\frac{1}{2}}(w-c_i)\|_2^2} |\Sigma_{j,j}(w_j - c_{i,j})| \times$$

$$\times \text{sgn}(\Sigma_{j,j}(w_j - c_{i,j})).$$

Note that based on the model structure in Fig. 1 there is in fact a feedback loop from the rotor speed to the nonlinearity via λ . This implies that for the the actual local linear model to be precise we should take this into account by calculating the closed-loop. In practice, however, no effect is observed when this step is skipped, due to the low-pass nature of this feedback path. It is also clear from the linearisation that the local dynamics are governed by the gradient of C_Q .

In Fig. 8 the LTI model of the demanded torque to rotor speed is shown, which is compared to the analytical linearised models from Bladed[®]. The analytical models demonstrate that the drive-train dynamics themselves are almost LTI. Further, a good correspondence is observed between the identified and analytical models. The small peak that is not fitted by the identification procedure corresponds to an interaction between the drivetrain and the tower modes. In these simulations no attempt was made to excite this mode and it turned out that it was almost invisible in the identification data, thus obscuring it in the result. A next step would possibly be to modify the excitation signals so as to excite this mode.

In Figs. 9 and 10 the LTI models of the demanded pitch to rotor speed are compared. The models clearly show an increasing gain with windspeed. The identified and analytical models show very similar trends. Again, the small peak at 1.5 Hz was not observable in the data and hence not modelled. The zero around 0.3 Hz is believed to be an artifact of the linearisation. Overall, the above-rated transfer




Fig. 8. The estimated transfer function from generator torque demand to rotor speed (black) compared to the analytical models in different operating points ($V = \{5, 7, \dots, 25 \text{ m/s}\}$).

functions match very well, whereas the below-rated ones underestimate the damping of the main resonance mode and show some low-frequency mismatch.




Fig. 9. The transfer functions from pitch angle demand to rotor speed (analytically obtained with Bladed[®], $V = \{5, 7, \dots, 25 \text{ m/s}\}$).




Fig. 10. The transfer functions from pitch angle demand to rotor speed (estimated, $V = \{5, 7, \dots, 25 \text{ m/s}\}$).

VI. CONCLUSIONS AND FUTURE WORK

This work has demonstrated that if the torque curve $C_Q(\lambda, \beta)$ of a wind turbine is properly parametrised, a Hammerstein model can be fitted to measured data. Thus, a global model describing the wind turbine can be obtained from data with a reasonable accuracy. The model can be used to extract local linear models as well as torque curve data. The model described in this paper can be extended with further inputs and outputs. In view of controlling power production using this model, a first logical choice is to include the rotor thrust curve $C_T(\lambda, \beta)$ together with a nacelle or tower displacement output to achieve load control. First attempts have already demonstrated that this is indeed feasible.

Next steps will include the design of a model predictive

control strategy to control the power production of the wind turbine across the operational envelope using the nonlinear model. Alternatively, the methodologies sketched here could be developed into a monitoring tool to track changes in aerodynamic performance (e.g. the “power curve”) of the wind turbine. Goals for the near future will be to add a full 3D turbulent wind field and a nacelle displacement measurement. It is further hoped that we have the opportunity to apply the method to operational data from a real wind turbine to assess its practical feasibility.

VII. ACKNOWLEDGMENTS

The authors gratefully acknowledge the support of Vestas Wind Systems A/S and the help of Ivo Houtzager in setting up the simulation environment.

REFERENCES

- [1] E. Bossanyi, “The design of closed loop controllers for wind turbines,” *Wind Energy*, vol. 3, no. 3, pp. 149–163, 2000.
- [2] J. Laks, L. Pao, and A. Wright, “Control of wind turbines: Past, present, and future,” in *Proc. of American Control Conference, 2009. ACC '09.*, jun. 2009, pp. 2096–2103.
- [3] E. van der Hooft, P. Schaak, and T. van Engelen, “Wind turbine control algorithms,” ECN, Tech. Rep. DOWEC-F1W1-EH-030094/0, 2003.
- [4] J. W. van Wingerden, A. W. Hulskamp, T. Barlas, I. Houtzager, H. E. N. Bersee, G. A. M. van Kuik, and M. Verhaegen, “Two-degree-of-freedom active vibration control of a prototyped “smart” rotor,” in *IEEE Transactions on Control System Technology (submitted for publication)*, 2010.
- [5] M. Iribas and I.-D. Landau, “Identification of wind turbines in closed-loop operation in the presence of three-dimensional turbulence wind speed: torque demand to measured generator speed loop,” *Wind Energy*, vol. 12, no. 7, pp. 660–675, 2009.
- [6] J. W. van Wingerden, I. Houtzager, F. Felici, and M. Verhaegen, “Closed-loop identification of the time-varying dynamics of variable-speed wind turbines,” *International Journal of Robust and Nonlinear Control*, special issue on Wind turbines: New challenges and advanced control solutions, vol. 19, no. 1, pp. 4–21, 2009.
- [7] J. Sjöberg and M. Viberg, “Separable non-linear least-squares minimization—possible improvements for neural net fitting,” in *Neural Networks for Signal Processing [1997] VII. Proceedings of the 1997 IEEE Workshop*, sep. 1997, pp. 345–354.
- [8] C. M. Bishop, *Pattern Recognition and Machine Learning (Information Science and Statistics)*, 1st ed. Springer, October 2007.
- [9] M. Verhaegen and V. Verdult, *Filtering and System Identification; A Least Squares Approach*, 1st ed. Cambridge, UK: Cambridge University Press, 2007.
- [10] G. J. van der Veen, J.-W. van Wingerden, and M. Verhaegen, “Closed-loop MOESP subspace model identification with parametrisable disturbances,” in *Proc. of the 49th IEEE Conference on Decision and Control, Atlanta, 2010. CDC 2010.*, dec. 2010.
- [11] —, “Closed-loop system identification of wind turbines in the presence of periodic effects,” in *Proc. of the 3rd conference, The Science of Making Torque from Wind*, 2010.
- [12] Z. Liu and L. Vandenberghe, “Interior-point method for nuclear norm approximation with application to system identification,” *SIAM Journal on Matrix Analysis and Applications*, vol. 31, no. 3, pp. 1235–1256, 2009.
- [13] M. Grant and S. Boyd, “CVX: Matlab software for disciplined convex programming, version 1.21,” May 2010. [Online]. Available: <http://cvxr.com/cvx>
- [14] L. Ljung, *System Identification: Theory for the User (2nd Edition)*. Prentice Hall PTR, December 1998.
- [15] *Bladed*, Garrad Hassan and Partners Ltd., accessed: May 19, 2010. [Online]. Available: <http://www.garradhassan.com>