



Analysis of Interpolating Regression Models and the Double Descent Phenomenon

Downloaded from: <https://research.chalmers.se>, 2025-12-05 04:39 UTC

Citation for the original published paper (version of record):

McKelvey, T. (2023). Analysis of Interpolating Regression Models and the Double Descent Phenomenon. IFAC-PapersOnLine, 56(2): 5869-5874. <http://dx.doi.org/10.1016/j.ifacol.2023.10.084>

N.B. When citing this work, cite the original published paper.

Analysis of Interpolating Regression Models and the Double Descent Phenomenon

Tomas McKelvey*

* Chalmers University of Technology, Gothenburg, Sweden
(e-mail: tomas.mckelvey@chalmers.se).

Abstract: A regression model with more parameters than data points in the training data is overparametrized and has the capability to interpolate the training data. Based on the classical bias-variance tradeoff expressions, It is commonly assumed that models which interpolate noisy training data are poor to generalize. In some cases, this is not true. The best models obtained are overparametrized and the testing error exhibits the double descent behavior as the model order increases. In this contribution, we provide some analysis to explain the double descent phenomenon, first reported in the machine learning literature. We focus on interpolating models derived from the minimum norm solution to the classical least-squares problem and also briefly discuss model fitting using ridge regression. We derive a result based on the behavior of the smallest singular value of the regression matrix that explains the peak location and the double descent shape of the testing error as a function of model order.

Copyright © 2023 The Authors. This is an open access article under the CC BY-NC-ND license (<https://creativecommons.org/licenses/by-nc-nd/4.0/>)

Keywords: machine learning, non-linear regression, matrix analysis, least-squares

1. INTRODUCTION

Linearly parametrized regression models that interpolate the training data have recently attracted significant attention Hastie et al. (2022), mainly due to the close connections to many state-of-the-art machine learning models Zhang et al. (2016); Belkin et al. (2019). Interpolation of the training data is obtained when the number of estimated/trained parameters in the model is equal to or larger than the number of training data used to estimate the model. Such models are hence overparametrized and there exists an infinite number of solutions that interpolate the data. It has been noticed in recent publications Hastie et al. (2022); Ribeiro et al. (2021) that the quality of an estimated model with a linear parametrization often follows a so-called double descent curve. This means that the test data error first decreases with an increasing model order and then increases again to a maximum when the number of samples of the data is equal to the number of parameters and then gradually decreases again with an increasing model order. The double descent phenomenon has previously been noticed in the deep-learning area Belkin et al. (2019); Loog et al. (2020).

In this paper, we will show that this behavior can be explained with a minimum of theory and complement the available literature on the subject. Furthermore, we point out that the extent of the behavior is closely tied to how the model class is constructed, i.e., the ordering of the basis functions employed in the regression model.

The paper is structured as follows. In Section 2, the regression problem is formulated and necessary notation and assumptions are explained. In Section 3, we derive explicit expressions for the bias and variance contributions for predictions obtained using the estimated model and

show that the smallest singular value of the regression matrix plays a key role. We further provide a theorem that predicts the behavior of the smallest singular value as the model order increases. A few numerical examples are given in Section 4 that illustrate the connection between the model quality and the smallest singular value. In Section 5, the paper is concluded with a summary of the findings.

2. PROBLEM FORMULATION

We consider a general regression problem where we seek a function that maps a value in the domain \mathcal{X} to the co-domain \mathcal{Y} , i.e. $f : \mathcal{X} \rightarrow \mathcal{Y}$ based on a given set of samples of training data

$$\mathcal{D} = \{(y(t), x(t))\}_{t=1}^M \quad (1)$$

where $x(t) \in \mathcal{X}$ and $y(t) \in \mathcal{Y}$ and we consider a linearly parametrized function class

$$f(x; \theta) = \sum_{i=1}^n \theta_i \phi_i(x) = \phi(x) \theta \quad (2)$$

where $\phi_i : \mathcal{X} \rightarrow \mathcal{Y}$ are given distinct basis functions and

$$\begin{aligned} \phi(x) &= [\phi_1(x) \ \phi_2(x) \ \cdots \ \phi_n(x)] \\ \theta &= [\theta_1 \ \theta_2 \ \cdots \ \theta_n]^T \in \mathcal{T}. \end{aligned} \quad (3)$$

The model order is equal to the size of the parameter vector θ and is denoted by the integer n . In the analysis that follows we assume the sets \mathcal{X} and \mathcal{Y} can be real or complex spaces with finite dimensions m and p respectively and the parameter set \mathcal{T} can be real or complex valued with finite dimension n .

Based on the training data set the model parameters are determined by minimizing the sum of squared errors

$$\hat{\theta} = \arg \min_{\theta} \sum_{t=1}^M \|y(t) - \sum_{i=1}^n \theta_i \phi_i(x(t))\|^2 \quad (4)$$

The minimization problem above can be written as the least-squares (LS) problem

$$\hat{\theta} = \arg \min_{\theta} \|y - \Phi \theta\|^2 \quad (5)$$

with the *regression matrix*

$$\Phi = \begin{bmatrix} \phi(x(1)) \\ \phi(x(2)) \\ \vdots \\ \phi(x(M)) \end{bmatrix} = \begin{bmatrix} \phi_1(x(1)) & \phi_2(x(1)) & \cdots & \phi_n(x(1)) \\ \phi_1(x(2)) & \phi_2(x(2)) & \cdots & \phi_n(x(2)) \\ \vdots & \vdots & \ddots & \vdots \\ \phi_1(x(M)) & \phi_2(x(M)) & \cdots & \phi_n(x(M)) \end{bmatrix} \quad (6)$$

which is a matrix with $N \triangleq pM$ rows and n columns and the vector

$$y = \begin{bmatrix} y(1) \\ y(2) \\ \vdots \\ y(M) \end{bmatrix}. \quad (7)$$

has N elements. We note that if $n \leq N$ we have an under-parametrized problem and the solution to (5) is unique if Φ has full rank. If $n > N$ the problem is overparametrized and there exist infinite many solutions to (5). The singular value decomposition (SVD) Golub and Van Loan (1989) of the regression matrix Φ in (3) plays a key role in the analysis in this paper and we denote it as

$$\Phi = \sum_{k=1}^r \sigma_k u_k v_k^H \quad (8)$$

where $r = \min(n, N)$, the ordered singular values are $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_r \geq 0$, and u_k and v_k are the left and right singular vectors respectively and $(\cdot)^H$ denote the Hermitian transpose. We assume Φ has full rank and hence $\sigma_r > 0$.

For analysis purposes we assume there exists a function $f_0 : \mathcal{X} \rightarrow \mathcal{Y}$ such that the data generated can be described as

$$y(t) = f_0(x(t)) + z(t) \quad (9)$$

where $z(t)$ is an i.i.d. zero mean white noise process with variance

$$\mathbf{E} z(t) z(t)^H \triangleq \mathbf{R}_z.$$

3. ANALYSIS

In this section, we discuss and provide some analytical results on the bias and variance of the estimated model that is given by the minimum norm solution and the ridge regression solution. We show that the largest variance contribution is proportional to $1/\sigma_r^2$, the inverse of the square of the smallest singular value of the regression matrix Φ . Finally, we show that when $n < N$ and the model order is increased to $n + 1$ the smallest singular value will decrease or stay unchanged. This implies that the variance increases or stays constant with the increase of the model order. Furthermore, when $n \geq N$ and the model order is increased to $n + 1$ we show that the smallest singular value is increased or stays the same. This implies that the variance decreases or stays constant with an increase in the model order.

3.1 The minimum norm solution

The unique *minimum norm solution* to (5) is obtained with the Moore-Penrose pseudo-inverse and can be expressed using the SVD as (see, e.g. Golub and Van Loan (1989); Laub (2005))

$$\hat{\theta} = \Phi^+ y = \sum_{k=1}^r \frac{1}{\sigma_k} v_k u_k^H y. \quad (10)$$

By introducing the notation

$$x = \begin{bmatrix} x(1) \\ x(2) \\ \vdots \\ x(M) \end{bmatrix}, \quad f_0(x) = \begin{pmatrix} f_0(x(1)) \\ f_0(x(2)) \\ \vdots \\ f_0(x(M)) \end{pmatrix}, \quad z = \begin{pmatrix} z(1) \\ z(2) \\ \vdots \\ z(M) \end{pmatrix} \quad (11)$$

the estimate in (10) can be decomposed into

$$\hat{\theta} = \sum_{k=1}^r \frac{1}{\sigma_k} v_k u_k^H (f_0(x) + z) = \theta_* + \tilde{\theta} \quad (12)$$

where $\theta_* = \Phi^+ f_0(x) = \sum_{k=1}^r \frac{1}{\sigma_k} v_k u_k^H f_0(x)$ is the noise free minimum norm solution and $\tilde{\theta} = \sum_{k=1}^r \frac{1}{\sigma_k} v_k u_k^H z$ is the contribution due to the noise. The zero mean assumption on $z(t)$ yields

$$\mathbf{E} \tilde{\theta} = \theta_* \quad (13)$$

The error of the output $f(x'; \hat{\theta})$ of the estimated model given a new test data sample x' is

$$\begin{aligned} e(x') &= \phi(x') \hat{\theta} - f_0(x') \\ &= \phi(x') \theta_* - f_0(x') + \phi(x') \sum_{k=1}^r \frac{1}{\sigma_k} v_k u_k^H z \end{aligned} \quad (14)$$

From above it is clear that the error $e(x')$ is composed of a bias part $\phi(x') \theta_* - f_0(x')$ and a zero mean stochastic part $\phi(x') \sum_{k=1}^r \frac{1}{\sigma_k} v_k u_k^H z$ that contributes with variance

$$\begin{aligned} R_e(x') &\triangleq \text{cov}(e(x')) = \\ &\phi(x') \left[\sum_{k=1}^r \frac{1}{\sigma_k} v_k u_k^H \right] (I_N \otimes R_z) \left[\sum_{k=1}^r \frac{1}{\sigma_k} v_k u_k^H \right]^H \phi^H(x') \end{aligned} \quad (15)$$

If the measurement noise is uncorrelated between the channels and have equal variance r_z , then $R_z = r_z \mathbf{I}$. In this case, the covariance expression (15) is simplified to

$$\begin{aligned} R_e(x') &= r_z \sum_{k=1}^r \frac{1}{\sigma_k^2} (\phi(x') v_k) (\phi(x') v_k)^H \\ &= r_z \phi(x') \left(\sum_{k=1}^r \frac{1}{\sigma_k^2} v_k v_k^H \right) \phi^H(x') \end{aligned} \quad (16)$$

since $u_k^H u_l = 0$ for $k \neq l$. From (14) (and (16)) it is clear that if the smallest singular value of Φ is close to zero the error (and variance) can become arbitrarily large unless $\phi(x')$ is perpendicular to the right singular vector corresponding to the smallest singular value. Further we note that for any singular value distribution we have the inequality $\sum_{k=1}^r \frac{1}{\sigma_k^2} \geq \frac{r^2}{\sum_{k=1}^r \sigma_k^2}$ with equality if all singular values are equal, see e.g. Xia et al. (1999). A selection of basis functions that results in equal singular values can hence be regarded as *variance optimal*.

3.2 Bias

The size of the bias contribution $\phi(x')\theta_* - f_0(x')$ in (14) depends on several factors. If we start by assuming that the model is correctly specified, i.e. there exists a vector θ_0 such that for all x we have $f_0(x) = \phi(x)\theta_0$ then the bias part of the error is

$$\begin{aligned} \mathbf{E}e(x') &= \phi(x')\theta_* - f_0(x') = \phi(x')\Phi^+\Phi\theta_0 - f_0(x') \\ &= \phi(x')(\Phi^+\Phi - \mathbf{I})\theta_0 \end{aligned} \quad (17)$$

- In the underparametrized case ($n \leq N$), we see that $\Phi^+\Phi - \mathbf{I} = 0$, since Φ has full rank, thus the bias is zero.
- If $n > N$ then $\Phi^+\Phi - \mathbf{I}$ is a rank $n - N$ projection matrix which projects onto the nullspace of Φ . The bias of $e(x')$ hence depends on the co-linearity of the projected true parameter vector and the row vector(s) $\phi(x')$. It also follows that the bias is zero if $\theta_0 = \Phi^T \mathbf{p}$ for some vector \mathbf{p} . In effect, this tells us that of all possible correctly specified models $f_0(x) = \phi(x)\theta_0$ only the N dimensional subset of models given by a parameter that can be expressed as $\Phi^T \mathbf{p}$ will have zero bias since $(\Phi^+\Phi - \mathbf{I})\Phi^T \mathbf{p} = 0$. The set of zero bias models hence depends explicitly on the training data through the properties of the matrix Φ . As n increases the size of the set of true functions with a non-zero bias also increases as the size of the nullspace of Φ increases with n .

For the misspecified case, when true function $f_0(x)$ is not part of the parametrized model class $f(x; \hat{\theta}) = \phi(x)\hat{\theta}$ the bias is

$$\mathbf{E}e(x') = \phi(x')\theta_* - f_0(x') = \phi(x')\Phi^+f_0(x) - f_0(x') \quad (18)$$

For the overparametrized case when $n \geq N$ the estimated model interpolates the training data. Hence, if x' is very close to one of the training data inputs $x(t)$ in \mathcal{D} , (see (1)), we can expect the bias $\mathbf{E}e(x')$ to be very small if the basis functions $\phi_i(x)$ are continuous. In general when x' is further away from the training samples the size of the bias error is difficult to characterize beyond the expression given in (18).

3.3 Ridge-Regression

If we add a squared penalty of the parameters to the LS norm in (5) we obtain the ridge regression solution. For a positive scalar λ , the parameter estimate is given by

$$\begin{aligned} \hat{\theta} &= \arg \min_{\theta} \|\mathbf{y} - \Phi\theta\|^2 + \lambda\|\theta\|^2 \\ &= \arg \min_{\theta} \left\| \begin{bmatrix} \mathbf{y} \\ 0 \end{bmatrix} - \begin{bmatrix} \Phi \\ \sqrt{\lambda}\mathbf{I} \end{bmatrix} \theta \right\|^2 \end{aligned} \quad (19)$$

When $\lambda > 0$ the extended regression matrix always has full rank and hence the LS problem has a unique solution. Using the SVD of Φ defined in (8) we can explicitly write the solution as

$$\begin{aligned} \hat{\theta} &= \left(\begin{bmatrix} \Phi \\ \sqrt{\lambda}\mathbf{I} \end{bmatrix}^H \begin{bmatrix} \Phi \\ \sqrt{\lambda}\mathbf{I} \end{bmatrix} \right)^{-1} \Phi^H \mathbf{y} = (\Phi^H \Phi + \lambda \mathbf{I})^{-1} \Phi^H \mathbf{y} \\ &= \left(\sum_{k=1}^r (\sigma_k^2 + \lambda) \mathbf{v}_k \mathbf{v}_k^H + \sum_{k=r+1}^n \lambda \mathbf{v}_k \mathbf{v}_k^H \right)^{-1} \times \\ &\quad \left(\sum_{k=1}^r \sigma_k \mathbf{v}_k \mathbf{u}_k^H \right) \mathbf{y} = \sum_{k=1}^r \frac{\sigma_k}{\sigma_k^2 + \lambda} \mathbf{v}_k \mathbf{u}_k^H \mathbf{y} \end{aligned} \quad (20)$$

where the sum $\sum_{k=r+1}^n \lambda \mathbf{v}_k \mathbf{v}_k^H$ vanishes if $n \leq N$. It is clear from (20) that the ridge regression solution converges to the minimum norm solution (10) as $\lambda \rightarrow 0$.

Following the same analysis as above we have $\hat{\theta} = \theta_* + \tilde{\theta}$ where the noise free solution is given by

$$\theta_* = \sum_{k=1}^r \frac{\sigma_k}{\sigma_k^2 + \lambda} \mathbf{v}_k \mathbf{u}_k^H f_0(x) \quad (21)$$

and the noise-induced error is given by

$$\tilde{\theta} = \sum_{k=1}^r \frac{\sigma_k}{\sigma_k^2 + \lambda} \mathbf{v}_k \mathbf{u}_k^H \mathbf{z} \quad (22)$$

The expression on the variance of the model for a new value x' corresponding to (16) is given by

$$R_e(x') = r_z \sum_{k=1}^r \frac{\sigma_k^2}{(\sigma_k^2 + \lambda)^2} (\phi(x') \mathbf{v}_k) (\phi(x') \mathbf{v}_k)^H \quad (23)$$

It is clear that the variance can be reduced by increasing λ . However, if $\lambda > 0$ and $n \geq N$ then

$$\begin{aligned} \Phi \hat{\theta} &= \left[\sum_{k=1}^r \sigma_k \mathbf{u}_k \mathbf{v}_k^H \right] \sum_{k=1}^r \frac{\sigma_k}{\sigma_k^2 + \lambda} \mathbf{v}_k \mathbf{u}_k^H \mathbf{y} \\ &= \sum_{k=1}^r \frac{\sigma_k^2}{\sigma_k^2 + \lambda} \mathbf{u}_k \mathbf{u}_k^H \mathbf{y} \neq \mathbf{y} \end{aligned} \quad (24)$$

which means that the estimated model does not interpolate the training data. This effect is commonly known as shrinkage since the estimated model parameters are smaller in magnitude than the LS minimum norm solution. The shrinkage effect will hence add to the total bias of the estimated model.

3.4 Analysis of the smallest singular value

In this section we derive results on the behaviour of the smallest singular value as the model order increases. The results gives a direct explanation to the double descent phenomenon. We will show

- that when $n > N$ then the minimum singular value of Φ for increasing model orders is non-decreasing.
- that for $n < N$ then the minimum singular value is non-increasing for increasing model orders.

The result is based on the following general matrix result. *Theorem 1.* Let Φ denote a matrix with n columns and N rows and define $\tilde{\Phi} = [\Phi \ \phi_{n+1}]$ where ϕ_{n+1} is an arbitrary vector.

- (1) Assume $n < N$ and let $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_n$ denote the singular value of Φ and let $\bar{\sigma}_1 \geq \bar{\sigma}_2 \geq \dots \geq \bar{\sigma}_{n+1}$ denote the singular values of $\tilde{\Phi}$. Then

$$\bar{\sigma}_1 \geq \sigma_1 \geq \bar{\sigma}_2 \geq \sigma_2 \geq \dots \geq \sigma_n \geq \bar{\sigma}_{n+1} \quad (25)$$

- (2) Assume $n \geq N$ and let $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_N$ denote the singular values of Φ and let $\bar{\sigma}_1 \geq \bar{\sigma}_2 \geq \dots \geq \bar{\sigma}_N$ denote the singular values of $\bar{\Phi}$. Then
- $$\bar{\sigma}_1 \geq \sigma_1 \geq \bar{\sigma}_2 \geq \sigma_2 \geq \dots \geq \bar{\sigma}_N \geq \sigma_N \quad (26)$$

Proof. The result follows from (Horn and Johnson, 1985, Theorem 4.3.15) or (Horn and Johnson, 1991, Corollary 3.1.3)

Corollary 2.

- (1) If the model order satisfies $n < N$ then a model order increase will result in a non-increase (unchanged or decreased size) of the smallest singular value.
- (2) If the model order satisfies $n \geq N$ then a model order increase will result in a non-decrease (unchanged or increased size) of the smallest singular value.

The presented result shows that if the inverse of the smallest singular value has a maximum, then it will appear when $n = N$. As the level of the variance is highly dependent on the smallest singular value as shown in Section 3 the maximum variance will in general appear for model order equal to N and the double descent curve will peak at $n = N$ if the error is dominated by the variance contribution.

3.5 Does overparametrization give any advantages?

The key finding in the section above is that for $n > N$ the smallest singular value σ_r will not decrease with n . It will stay the same or increase. For the miss-specified case where the noise-free solution is given by

$$\theta_* = \phi^+ f_0(x) = \sum_{k=1}^r \frac{1}{\sigma_k} v_k u_k^H f_0(x) \quad (27)$$

we can conclude that the magnitude of the elements in θ_* is highly influenced by the value of $1/\sigma_r$. As the model order increases $1/\sigma_r$ will in general decrease and the magnitude of the elements in θ_* will decrease. From the predictive point of view of the estimated models, i.e. values outside the training set, it seems more natural that models with a smaller norm of the parameter vector are better than models that have very large norm of the parameter vector. A second, more clear, benefit is that the noise sensitivity is decreased as the model order increases, at least for moderate values above N . This effect is of course most pronounced when the regression matrix is close to singular when $n = N$ and hence, the smallest singular value is closest to zero.

4. NUMERICAL ILLUSTRATIONS

In this section we examine some simulated numerical examples and interpret the results with help of the results derived in the analysis section. In all examples we will estimate regression models with varying orders using the scalar valued (i.e. $p = 1$) complex exponential $\phi(x) = e^{j2\pi f x}$ as basis functions. A model structure is defined by the set of frequencies $\mathcal{F} = \{f_i\}_{i=1}^n$ and is given by

$$f(x; \theta) = \sum_{i=1}^n \theta_i e^{j2\pi f_i x}. \quad (28)$$

We generate the training data according to (9) where we select $x(t) = t$ and let $t = 0, 1, \dots, N-1$ where $N = 10$

Case	ε	f_0	n_{\max}	N
A	0.5	$f_{0,\text{lin}}(x)$	30	10
B	0	$f_{0,\text{lin}}(x)$	30	10
C	0.5	$f_{0,\text{opt}}(x)$	30	10
D	0	$f_{0,\text{opt}}(x)$	30	10

Table 1. Definition of the different cases in the numerical examples

and the noise $z(t)$ are i.i.d. samples drawn from a zero mean circular symmetric complex Gaussian distribution with variance 0.1. To evaluate the quality of an estimated regression model $f(x; \hat{\theta})$ we derive the normalized mean square error of the predictions at the test data samples $x'(t)$ for $t = 0, 1, \dots, N-1$

$$\text{NMSE} = \frac{\sum_t |f_0(x'(t)) - f(x'(t); \hat{\theta})|^2}{\sum_t |f_0(x'(t))|^2}. \quad (29)$$

We let $x'(t) = t + \varepsilon$, $t = 0, \dots, N-1$ and vary ε in the different experiments. If $\varepsilon = 0$ the estimated model is evaluated in the same data points as used for the learning. If $\varepsilon = 0.5$ we evaluate the quality of the model's ability to predict values in between training samples, i.e. ability to generalize.

In the experiments we use two different model structures. We set $n_{\max} = 3N$ as the maximum model order. For the model structure denoted as *linear ordering* we define the frequency set as

$$\mathcal{F}_{\text{lin},n} = \{k/n_{\max}\}_{k=0}^{n-1}. \quad (30)$$

For the model structure denoted as *optimal ordering* we define the frequency set as

$$\mathcal{F}_{\text{opt},n} = \begin{cases} \{k/N\}_{k=0}^{n-1}, & n \leq N \\ \{k/N\}_{k=0}^{N-1} \cup \mathcal{S}_{n-N}, & n > N \end{cases} \quad (31)$$

where the set \mathcal{S}_k is the first k elements in the ordered set $\{k/n_{\max}\}_{k=0}^{n_{\max}-1} - \{k/N\}_{k=0}^{N-1}$. If $n \leq N$ then the columns in the regression matrix Φ defined in (6) are orthogonal to each other and have equal norms. This in turn shows that all singular values are non-zero and equal and hence this model structure is *variance optimal* as discussed in Section 3. We note that that $\mathcal{F}_{\text{lin},n_{\max}} = \mathcal{F}_{\text{opt},n_{\max}}$, i.e. the two model structures are identical for $n = n_{\max}$ but with a different ordering of the basis functions.

We will use the same type of basis functions to define two data generating functions where the first one is given by

$$f_{0,\text{lin}}(x) = \sum_{k=1}^{10} \alpha_k e^{j2\pi \frac{k-1}{n_{\max}} x} \quad (32)$$

and the second one is

$$f_{0,\text{opt}}(x) = \sum_{k=1}^{10} \alpha_k e^{j2\pi \frac{k-1}{N} x} \quad (33)$$

In the Monte-Carlo Simulations below we will generate the coefficients α_k by sampling from a zero mean circular symmetric complex Gaussian distribution with unit variance. The construction of the data generating systems implies that for $f_{0,\text{lin}}(x)$ then all model structures defined by $\mathcal{F}_{\text{lin},n}$ for $n \geq 10$ will include the data generating system in the model class. Along the same lines as above we notice that $f_{0,\text{opt}}(x)$ is included in all model structures defined by the set $\mathcal{F}_{\text{opt},n}$ when $n \geq N$. In Table 1 we define four experimental cases. For each case we generate a data generating system using the function according to the

f_0 column and create 500 training datasets with added noise and 500 data sets without noise. For each dataset a model is estimated and the NMSE is evaluated at the x values defined by $t + \varepsilon$ for $t = 0, \dots, N - 1$.

The average NMSE over the Monte-Carlo simulations for the two model structures as a function of the model orders are reported in the graphs in Figure 1 to Figure 4. In Figure 5 the inverse of the smallest singular value of the regression matrix Φ is illustrated as a function of model order for the two model structures.

4.1 Discussion

The NMSE testing error is in the figures shown for models trained on noise free data as well as trained on the noisy data. The testing results on models trained on noise free data give direct information on the NMSE caused by the bias contribution given by (17) and (18). The testing results on models trained on noisy data give information about the total MSE caused by the bias contribution and the variance contribution given by (16).

In Case A the true system is in the linear ordering model class for model orders $n \geq 10$. Hence for noise free data and $n = 10$ we recover the true model as seen in Figure 1. For the noise free case the test data error has an increase again for model orders larger than 20. This is the effect when the true model parameters are not in the row space of Φ as discussed below (17). For models estimated from noisy data the double descent phenomenon is clearly visible. A comparison with the top graph in Figure 5 show the qualitative agreement between the NMSE and the inverse of the smallest singular value. The peaks are located for $n = N = 10$ in both graphs and the behaviour for the singular values are in agreement with Corollary 2. The result for Case A for the optimal ordering model class is shown in the bottom graph in Figure 1. For this case the true model is in the model set for $n \geq 16$. Hence, even for noise free data this model structure has a non-zero error that for the highest model orders increases again for the same reasons as discussed before. However, for the noisy case the performance is significantly improved for this model structure and is in par with the performance of the model estimated from the noise free data. The reason for this is found in the bottom graph in Figure 5. The inverse of the smallest singular value is much smaller than the linear ordering model structure for all model orders. This implies that the variance as given by (16) is much smaller as compared with the other model structure. In Case B the same setup is used except that the test data points are the same as the training data points. For the noise free case we obtain zero error for both model structures for $n \geq 10$. For the noisy data case an error which is equal to the noise level is obtained since all models interpolate the training data when $n \geq 10$. For case C and D the true system is now given by the optimal order data structure. Hence, for $n = 10$ the optimal ordering model structure recovers the true model for noise free data and give the best NMSE for the noisy data. For higher model orders the performance is slightly reduced which again is attributed to an increase in the bias as discussed before. For the linear ordering model structure it is only at $n = 28$ the true system is in the model class and it is at this model order the best NMSE on test data are achieved. The ill-conditioning of this model




Fig. 1. Case A: Graphs show normalized mean squared error as a function of model order for the linear ordering model structure (top) and the optimal ordering (bottom).




Fig. 2. Case B: Graphs show normalized mean squared error as a function of model order for the linear ordering model structure (top) and the optimal ordering (bottom).

structure is for the lower model orders clearly visible and the NMSE again has a peak at $n = N = 10$. For this model structure we can conclude that overparametrization produces a model with reasonable performance as compared to the solutions for model orders around 10.

5. CONCLUSIONS

The existence of a double descent behaviour is closely related to the inverse of the smallest singular value of the associated regression matrix. A model structure with a near singular regression matrix when $n = N$ results in a double descent behavior for the NMSE on test data at other locations than the training data.




Fig. 3. Case C: Graphs show normalized mean squared error as a function of model order for the linear ordering model structure (top) and the optimal ordering (bottom).




Fig. 4. Case D: Graphs show normalized mean squared error as a function of model order for the linear ordering model structure (top) and the optimal ordering (bottom).

To estimate overparametrized models, i.e. more parameters than training data using the pseudo inverse solution can be reasonable ($NMSR < 1$) if the true parameter is close to the row space of the regression matrix. If this is not the case the solutions will have poor performance.

To obtain robust overparametrized solutions it is important to select a model class such that the minimum singular value of the associated regression matrix is as large as possible.

ACKNOWLEDGEMENTS

The author would like to thank Daniel McKelvey for giving valuable comments on the manuscript and the reviewers for their constructive comments.




Fig. 5. The graphs show the inverse of the smallest singular value for the regressor matrix Φ as a function of model order for the linear ordering model structure (top) and the optimal ordering (bottom).

REFERENCES

- Belkin, M., Hsu, D., Ma, S., and Mandal, S. (2019). Reconciling modern machine-learning practice and the classical bias-variance trade-off. *Proceedings of the National Academy of Sciences*, 116(32), 15849–15854.
- Golub, G.H. and Van Loan, C.F. (1989). *Matrix Computations*. The Johns Hopkins University Press, Baltimore, Maryland.
- Hastie, T., Montanari, A., Rosset, S., and Tibshirani, R.J. (2022). Surprises in high-dimensional ridgeless least squares interpolation. *The Annals of Statistics*, 50(2), 949–986. doi:10.1214/21-AOS2133.
- Horn, R.A. and Johnson, C.R. (1985). *Matrix Analysis*. Cambridge University Press, Cambridge, NY.
- Horn, R.A. and Johnson, C.R. (1991). *Topics In Matrix Analysis*. Cambridge University Press, Cambridge, NY.
- Laub, A.J. (2005). *Matrix Analysis for Scientists and Engineers*, volume 91. Siam.
- Loog, M., Viering, T., Mey, A., Krijthe, J.H., and Tax, D.M. (2020). A brief prehistory of double descent. *Proceedings of the National Academy of Sciences*, 117(20), 10625–10626.
- Ribeiro, A.H., Hendriks, J.N., Wills, A.G., and Schön, T.B. (2021). Beyond Occam's razor in system identification: Double-descent when modeling dynamics. *IFAC-PapersOnLine*, 54(7), 97–102.
- Xia, D.F., Xu, S.L., and Qi, F. (1999). A proof of the arithmetic mean-geometric mean-harmonic mean inequalities. *RGMA research report collection*, 2(1).
- Zhang, C., Bengio, S., Hardt, M., Recht, B., and Vinyals, O. (2016). Understanding deep learning requires rethinking generalization. doi: 10.48550/ARXIV.1611.03530.