International Journal of Data Science and Analytics (2026) 22:13
https://doi.org/10.1007/s41060-025-00906-9

REGULAR PAPER

Overfitting, regularisation and condition estimation in regression

Joab R. Winkler[1]

Received: 22 April 2025 / Accepted: 15 October 2025

Abstract
Overfitting is a problem in regression and deep neural networks, and it is often stated that Tikhonov regularisation minimises its adverse effects, but the relationship between regularisation and overfitting has not been established. The theory of regularisation is well developed, but overfitting has a qualitative description and it is not defined mathematically. This paper addresses the relationship between overfitting, regularisation and condition estimation by considering underdetermined and overdetermined least squares (LS) problems that arise in regression. This study is important because regularisation is not benign since its use when a condition on the decay of the singular values of the coefficient matrix in the LS minimisation is not satisfied leads to a large error in the solution of the regularised LS problem. Examples in which the regression curve overfits the data are shown, but regularisation must not be applied because the LS problem is well conditioned. Also, an ill conditioned LS problem whose solution does not display overfitting is shown, but its ill conditioned nature implies regularisation should be applied in order to obtain a numerically stable solution. It is concluded that regularisation does not solve the problem of overfitting in regression.

Keywords Regression · Overfitting · Regularisation · Condition estimation

1 Introduction

Table 1 lists the symbols used in this paper.

The development of a model of data $\mathcal{D}$ in machine learning requires that it be partitioned into a training set $\mathcal{D}_{\text{train}}$ and a test set $\mathcal{D}_{\text{test}}$, where $\mathcal{D} = \mathcal{D}_{\text{train}} \cup \mathcal{D}_{\text{test}}$ and $\mathcal{D}_{\text{train}} \cap \mathcal{D}_{\text{test}} = \emptyset$. The model is developed using the data set $\mathcal{D}_{\text{train}}$ and tested using the data set $\mathcal{D}_{\text{test}}$, and it should be sufficiently rich to represent the true properties of the data, but not too rich such that it overfits the data. Overfitting occurs, for example, when a polynomial $p(x)$ of degree greater than two models data whose exact form varies quadratically. It manifests itself in the model by many local minima and maxima that arise from the high degree of $p(x)$, but they are not present in the data. Classical theory suggests that models that overfit training data should perform badly on test data, but this is not observed. The causes of this good generalisation property of rich models on test data and its inconsistency with classical theory are the focus of much research because of their implications for machine learning [1–3].

✉ Joab R. Winkler
j.r.winkler@sheffield.ac.uk

[1] School of Computer Science, The University of Sheffield, Regent Court, 211 Portobello, Sheffield S1 4DP, United Kingdom

It is often claimed that regularisation solves the problem of overfitting, and it arises from the bias-variance trade-off, which defines the relationship between the complexity and accuracy of a model. The trade-off states that the mean squared error (MSE) of a model satisfies

$$\text{MSE} = (\text{bias})^2 + \text{variance},$$

where the bias in the model arises from incorrect assumptions in the learning algorithm and the variance is the error due to the sensitivity of the model to perturbations in the data. A low value of the variance is associated with little or no overfitting, and a high value of the variance is associated with extensive overfitting. It follows that, for a given value of the MSE, a reduction in the extent that a regression curve overfits the data is associated with an increase in its bias.

Overfitting is a major problem in neural networks, and several methods, including data augmentation, early stopping and dropout, may be used to implement regularisation. The performance of a neural network is dependent on many parameters, including the number of layers, the number of neurons in each layer, the batch size and the number of epochs, and it is therefore difficult to identify a relationship, if it exists, between regularisation and overfitting. These issues do not arise in regression, and thus the analysis in this paper

Published online: 08 January 2026

Table 1 Notation and symbols

Symbol	Representation
A	Coefficient matrix of order $m \times n$
$A^\dagger$	Pseudo-inverse of A
p	$\min(m, n)$
b	Vector of function values
x_{ls}	$\arg\min_z \|Az - b\|^2$
$U\Sigma V^T$	Singular value decomposition of A
σ_i	ith singular value of A
u_i, v_i	ith column of U, V
c	$U^T b$
$\kappa(A)$	Condition number of A
$\eta(A, b)$	Effective condition number of the least squares problem
e_{ls}	Error in the least squares solution
λ	Regularisation parameter
λ^*	Optimal value of the regularisation parameter
$x_0(\lambda)$	Regularised solution, with regularisation parameter λ
$e_{\text{reg}}(\lambda)$	Error in the regularised solution, with regularisation parameter λ

is restricted to underfitting and overfitting in regression [1, §1].

Tikhonov regularisation imposes stability on the solution x_{ls} of the least squares (LS) problem,

$$x_{\text{ls}} = \arg\min_z \|Az - b\|^2 , \tag{1}$$

where $\|\cdot\| = \|\cdot\|_2$, when it is sensitive to a perturbation in b, that is, the LS problem is ill conditioned.[1] The claim that regularisation solves the problem of overfitting requires, therefore, that the connection between ill conditioned LS problems and overfitting be considered. Evidence for this connection has not been established and thus the objective of this paper is an investigation into the association between regularisation and overfitting in order to determine a relationship between them, assuming it exists. It is stated in [4, §2.3] with reference to physics informed neural networks (neural networks that are used to solve partial differential equations), and in [5, §1.2] with reference to deep neural networks, that regularisation has little or no effect in minimising overfitting, but details are not provided. This paper extends the work in these papers by the provision of theoretical analysis, examples of underdetermined and overdetermined LS problems that arise in regression and consideration of overfitting that may, or may not, occur. This leads to the aim of this paper:

- An investigation into the association between overfitting and regularisation, using theoretical analysis and computational examples of regression, to determine this association, if it exists.

This investigation is important because (i) overfitting is a problem in machine learning since it causes a reduction in the quality of the computed model of the data, and (ii) regularisation is not benign. In particular, its application to the LS problem (1) requires that A and b satisfy the discrete Picard condition [6–9], which is a condition on the rate of decay of the singular values of A. If this condition is not satisfied, the error between the exact and regularised solutions of the LS problem is large.

The result of this paper is summarised:

- Regularisation does not solve the problem of overfitting in regression. In particular, well conditioned LS problems that lead to regression curves that overfit the data are shown, but the discrete Picard condition is not satisfied and thus regularisation must not be applied. Also, an ill conditioned LS problem that leads to a regression curve that does not overfit the data is shown, but regularisation must be applied in order to impose stability on its solution.

Section 2 includes an example of regression that motivates the work in this paper. Condition estimation of the LS problem is considered in Section 3 and it is shown that the condition number $\kappa(A)$ of A may yield a large overestimate of the numerical condition of the LS problem (1). A more refined measure, the effective condition number $\eta(A, b)$, of the condition of the LS problem is developed and its advantages with respect to $\kappa(A)$ are discussed. Regularisation is considered

[1] Tikhonov regularisation is known as ridge regression in machine learning, and it will, for brevity, be termed regularisation.

in Section 4, and the L-curve [7, §4.6], which is a graphical method for calculating the optimal value of the regularisation parameter, is considered. It is shown that the shape of the graph is dependent on the satisfaction, or lack of satisfaction, of the discrete Picard condition. Examples of regularisation, overfitting, and the condition numbers $\kappa(A)$ and $\eta(A, b)$ are in Section 5 and the paper is summarised in Section 6.

2 Background

The problem addressed in this paper is demonstrated in Example 1 by considering regression using a given matrix A and two forms $\tilde{b}_1$ and $\tilde{b}_2$ of b:

- The LS problem (1) is ill conditioned for $b = \tilde{b}_1$ and regularisation must therefore be applied. The regression curve has few local minima and maxima, and it does not overfit the data.
- The LS problem (1) is well conditioned for $b = \tilde{b}_2$ and thus regularisation must not be applied. The regression curve has many local minima and maxima, and it overfits the data.

Regression of the data $(x, y) = \{x_j, y_j\}_{j=1}^m$ by the basis functions $\phi(x) = \{\phi_i(x)\}_{i=1}^n$,

$$y_j = \sum_{i=1}^{n} a_i \phi_i(x_j), \qquad j = 1, \ldots, m,$$

leads to the LS problem, whose solution $x_{\text{ls}} \in \mathbb{R}^n$ contains the coefficients $\{a_i\}_{i=1}^n$,

$$\begin{aligned} x_{\text{ls}} &= \arg\min_z \|Az - b\|^2 \\ &= A^\dagger b \\ &= \begin{cases} (A^T A)^{-1} A^T b, & m > n, \\ A^{-1} b, & m = n, \\ A^T (AA^T)^{-1} b, & m < n, \end{cases} \end{aligned} \tag{2}$$

where $A_{i,j} = \phi_j(x_i)$ and $A = \{A_{i,j}\}_{i,j=1}^{m,n}$, $\operatorname{rank} A = \min(m, n)$ and $b = \{y_j\}_{j=1}^m$. The solution of (2) is

$$x_{\text{ls}} = \sum_{i=1}^{p} \left(\frac{c_i}{\sigma_i} \right) v_i, \tag{3}$$

where $p = \operatorname{rank} A = \min(m, n)$,

$$c = \{c_i\}_{i=1}^m = U^T b, \tag{4}$$

the singular value decomposition (SVD) of A is $U \Sigma V^T$, v_i is the ith column of V and $\Sigma \in \mathbb{R}^{m \times n}$ is the diagonal matrix of the singular values $\sigma_i, i = 1, \ldots, p$, of A. The solution x_{ls} of (2) is a function of A and b, but the condition number $\kappa(A)$ of A is a function of A only, and it is independent of b. This is problematic because it implies that $\kappa(A)$ may not be an accurate measure of the condition of the LS problem. This issue is addressed by the effective condition number $\eta(A, b)$, which is a function of A and b, and it therefore provides more accurate information on the condition of the LS problem. An expression for this condition number is developed in Section 3 and its importance for the analysis of the LS problem follows from the following points:

- The LS problem is well conditioned if $\eta(A, b) = O(1)$, even if $\kappa(A) \gg 1$, and regularisation must not be applied.
- The LS problem is ill conditioned if $\eta(A, b) \gg 1$ and regularisation must be applied if the discrete Picard condition is satisfied.
- The effective condition number $\eta(A, b)$ may be infinite, which must be compared with $\kappa(A)$, which is finite if A has full rank.

Example 1 shows that a study of the relationship between regularisation and overfitting requires that the more refined information on the condition of the LS problem provided by $\eta(A, b)$ be used, and that $\kappa(A)$ does not provide this refined information.

Example 1 Consider the regression of 100 points (x_j, y_j), $j = 1, \ldots, 100$, in the interval $I = [0, \ldots, 20]$ using 33 exponential basis functions, where the 100 points are randomly distributed in I. It follows that $A \in \mathbb{R}^{100 \times 33}$, and two sets of data, $b = \tilde{b}_1$ and $b = \tilde{b}_2$, which are shown in Figure 1, are considered.

(i) The data (x_j, y_j), $j = 1, \ldots, 100$, are shown in Fig. 1(i) and the LS problem is ill conditioned because

$$\eta(A, \tilde{b}_1) = 4.63 \times 10^8,$$

where $\tilde{b}_1 = \{y_j\}_{j=1}^{100}$, and thus

$$\eta(A, \tilde{b}_1) \approx \kappa(A) = \frac{\sigma_1}{\sigma_{33}} = 5.14 \times 10^8.$$

Regularisation is required and it is shown in Section 3 that the dominant components of $\tilde{b}_1$ lie in the space spanned by the columns of U associated with the large singular values of A. It therefore follows from (3) and (4) that x_{ls} lies in the space spanned by the columns of V that are also associated with the large singular values of A. Regularisation removes the effects of the small singular values from x_{ls} and the consequent error in the regularised solution is small because these singular values make a minor contribution to x_{ls}. The solution x_{ls} of the LS problem is exact because $\tilde{b}_1$ lies in the column space of

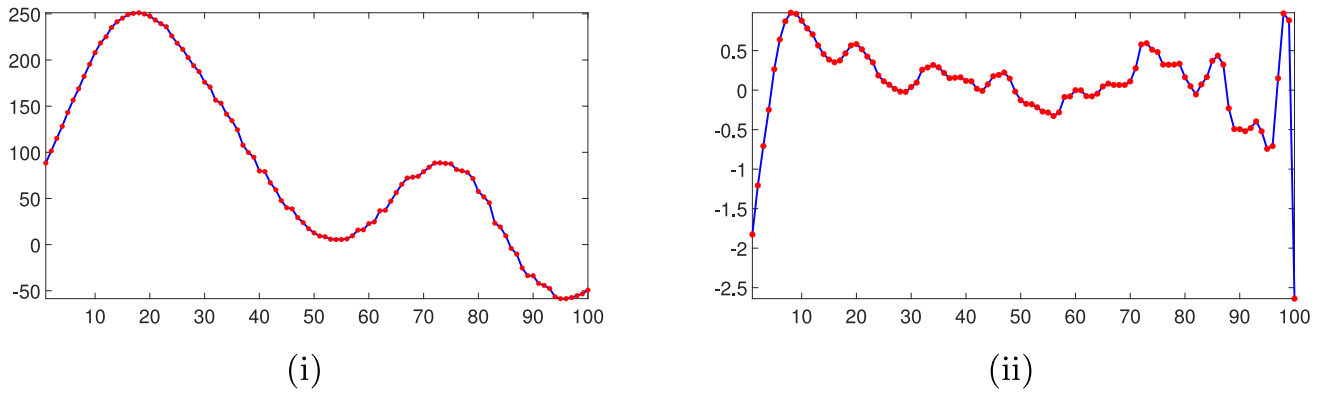

Fig. 1 (i) The data $b = \tilde{b}_1$ and the regression curve for an ill conditioned LS problem, and (ii) the data $b = \tilde{b}_2$ and the regression curve for a well conditioned LS problem, for Example 1

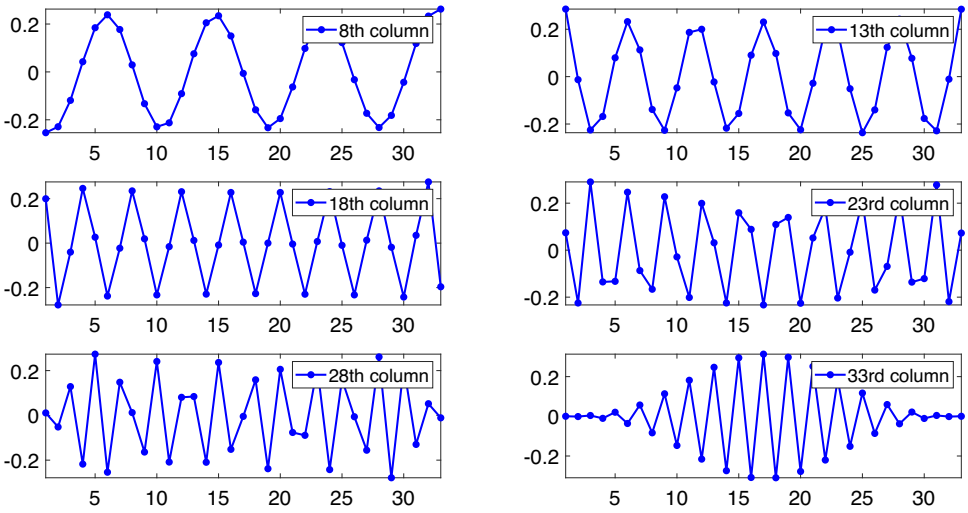

Fig. 2 The 8th, 13th, 18th, 23rd, 28th and 33rd columns of V for Example 1

Table 2 Results of Example 1. The vectors u_i, $i = 1, \ldots, 100$, and v_i, $i = 1, \ldots, 33$, are the columns of U and V, respectively

	$b = \tilde{b}_1$	$b = \tilde{b}_2$
$\eta(A, b)$	4.63×10^8	7.96
$b = Uc$	$\sum_{i=1}^{t} c_i u_i, \ t \ll 33$	$\sum_{i=t}^{33} c_i u_i, \ t \gg 1$
x_{ls}	$\sum_{i=1}^{t} \left(\frac{c_i}{\sigma_i}\right) v_i, \ t \ll 33$	$\sum_{i=t}^{33} \left(\frac{c_i}{\sigma_i}\right) v_i, \ t \gg 1$
Sign changes in the columns of V	Few sign changes in the first t columns	Many sign changes in the last $34 - t$ columns
Regularise x_{ls}	Yes	No
Overfitting	Not present	Present

A and thus the regression curve interpolates the data points, as shown in Fig. 1(i). The curve does not overfit the data, but as noted above, regularisation must be applied because x_{ls} is unstable with respect to a perturbation in $\tilde{b}_1$.

(ii) The data (x_j, y_j), $j = 1, \ldots, 100$, are shown in Fig. 1(ii) and the LS problem is well conditioned because

$$\eta(A, \tilde{b}_2) = 7.96 \approx \min_X \kappa(X) = 1,$$

where $\tilde{b}_2 = \{y_j\}_{j=1}^{100}$ and X is an arbitrary matrix of full rank. This situation arises because the dominant components of $\tilde{b}_2$ lie in the space spanned by the columns of U associated with the small singular values of A, and it follows from (3) and (4) that x_{ls} lies in the space spanned by the columns of V that are also associated with the small singular values of A. Regularisation must not be applied because it removes these singular values from x_{ls}, which would therefore yield a large error. The solution of the LS problem is exact because

$\tilde{b}_2$ lies in the column space of A and thus the regression curve interpolates the data points, as shown in Fig. 1(ii). The curve overfits the data, but this problem cannot be solved by regularisation because, as noted above, it must not be applied.

Figure 2 shows six columns of V and it is seen that its ith column v_i has $i-1$ sign changes, and thus:

- $x_{\rm ls} = A^\dagger \tilde{b}_1 = \left(A^T A\right)^{-1} A^T \tilde{b}_1$ is a linear combination of vectors v_i whose entries have few sign changes because it is a linear combination of the first few columns of V.
- $x_{\rm ls} = A^\dagger \tilde{b}_2 = \left(A^T A\right)^{-1} A^T \tilde{b}_2$ is a linear combination of vectors v_i whose entries have many sign changes because it is a linear combination of the last few columns of V.

The results for $b = \tilde{b}_1$ and $b = \tilde{b}_2$ are summarised in Table 2. □

Theorem 1, [10, p. 87] and [11, §5], establishes that the oscillatory property of the columns of V shown in Fig. 2 is a property of the eigenvectors of an oscillation matrix. It is included in the examples in Sect. 5 and it is shown that a distinction must be made between underdetermined and overdetermined LS problems.

Theorem 1 *If (λ_k, w_k) is the kth eigenpair of an oscillation matrix of order n, where the eigenvalues satisfy $\lambda_i > \lambda_j > 0$ for $i < j$, then for every set of coefficients $e_p, e_{p+1}, \ldots, e_q$, $1 \le p \le q \le n$, such that $\sum_{i=p}^{q} e_i^2 > 0$, the number of sign changes $\mathcal{S}$ in the entries of r,*

$$r = e_p w_p + e_{p+1} w_{p+1} + \cdots + e_q w_q,$$

satisfies $p - 1 \le \mathcal{S} \le q - 1$. It follows that if $i := p = q$, the entries of $r = e_i w_i$ have exactly $i - 1$ sign changes.

3 Condition estimation

The stability of $x_{\rm ls} = A^\dagger b$ is usually quantified by the condition number $\kappa(A) = \sigma_1/\sigma_p$, $p = \min(m, n)$, of A. It is a function of A only, but $x_{\rm ls}$ is a function of A and b, from which it follows that $\kappa(A)$ may not be a true measure of the stability of $x_{\rm ls}$. Condition numbers of the LS problem when A and b are perturbed are considered in [12, §5.3.6], but the analysis in this section is simpler because it is assumed there are errors in b only.

Example 2 Consider the two sets of data in Example 1. The solution $x_1 = A^\dagger \tilde{b}_1$ is sensitive to a change $\delta\tilde{b}_1$ in the data $\tilde{b}_1$, but the solution $x_2 = A^\dagger \tilde{b}_2$ is not sensitive to a change $\delta\tilde{b}_2$ in the data $\tilde{b}_2$. The condition number of A is $\kappa(A) = 5.14 \times 10^8$, which suggests that x_1 and x_2 are unstable. This is incorrect because x_2 is stable and thus the large value of $\kappa(A)$ cannot explain the difference between the results for the two data sets. □

Example 2 shows that the formula for the stability of $x_{\rm ls}$ with respect to a perturbation in b must be extended, such that it is a function of A and b because this will allow the difference between the stability of the solutions x_1 and x_2 to be quantified. This extension of $\kappa(A)$, which is called the effective condition number $\eta(A, b)$ [13, 14], is defined in Definition 1 and an expression for it is developed in Theorem 2. The forms of the expression when the LS problem is ill conditioned, and when it is well conditioned, are derived and the geometric conditions that lead to these situations are established. The inequality between $\eta(A, b)$ and $\kappa(A)$ is derived and it is shown that $\kappa(A)$ is finite if A has full rank, and that $\eta(A, b)$ may be infinite.

Definition 1 (The effective condition number) The effective condition number of the LS problem (2) is

$$\eta(A, b) = \max_{\delta b \in \mathbb{R}^m} \frac{\Delta x_{\rm ls}}{\Delta b},$$

where

$$\Delta x_{\rm ls} = \frac{\|\delta x_{\rm ls}\|}{\|x_{\rm ls}\|} \quad \text{and} \quad \Delta b = \frac{\|\delta b\|}{\|b\|}.$$

Theorem 2 is established in [15, §2] and [16, §5.1].

Theorem 2 *The effective condition number $\eta(A, b)$ is*

$$\begin{aligned} \eta(A, b) &= \frac{\left\|A^\dagger\right\| \|b\|}{\|x_{\rm ls}\|} \\ &= \frac{\|c\|}{\sigma_p \left\|\Sigma^\dagger c\right\|} \\ &= \frac{1}{\sigma_p} \left(\frac{\sum_{i=1}^{m} c_i^2}{\sum_{i=1}^{p} \left(\frac{c_i}{\sigma_i}\right)^2} \right)^{\frac{1}{2}}, \end{aligned} \tag{5}$$

where $c = \{c_i\}_{i=1}^m$ is defined in (4).

The superiority of $\eta(A, b)$ with respect to $\kappa(A)$ was considered above, and $\eta(A, b)$ allows the dominant columns of U and V spanned by b and $x_{\rm ls}$, respectively, when the LS problem is ill conditioned, and when it is well conditioned, to be calculated.

(i) *The* LS *problem is ill conditioned* It was stated in Sect. 1 that the LS problem is ill conditioned if the discrete Picard condition (6) is satisfied. This condition requires that

the constants $|c_i|$ decay to zero faster than the singular values decay to zero,

$$\frac{|c_i|}{\sigma_i} \to 0 \quad \text{as} \quad i \to p = \min(m, n). \tag{6}$$

If this condition is satisfied such that $\left\|\Sigma^{-1}c\right\| \approx |c_1|/\sigma_1$ because the terms $|c_i|/\sigma_i$ decay sufficiently rapidly, then

$$\begin{aligned}\eta(A, b) &\approx \frac{\sigma_1}{\sigma_p}\left(\frac{c_1^2 + \sum_{i=p+1}^{m} c_i^2}{c_1^2}\right)^{\frac{1}{2}} \\ &= \kappa(A)\left(1 + \frac{\sum_{i=p+1}^{m} c_i^2}{c_1^2}\right)^{\frac{1}{2}} \\ &\geq \kappa(A),\end{aligned}$$

where $\eta(A, b) \approx \kappa(A)$ if $p = m = n$, and thus the LS problem is ill conditioned if the discrete Picard condition is satisfied. It follows from (3) and (4) that if the LS problem is ill conditioned, then:

- The dominant components of b lie in the space spanned by the columns of U associated with the large singular values of A.
- The dominant components of x_{ls} lie in the space spanned by the columns of V associated with the large singular values of A.

(ii) *The* LS *problem is well conditioned* This condition is defined by the situation in which the constants $|c_i|$ are either approximately constant or increase, in which case

$$\frac{|c_{i+1}|}{\sigma_{i+1}} \gg \frac{|c_i|}{\sigma_i}, \qquad i = 1, \ldots, p-1.$$

It follows that

$$\begin{aligned}\eta(A, b) &\approx \left(\frac{\sum_{i=1}^{p} c_i^2 + \sum_{i=p+1}^{m} c_i^2}{c_p^2}\right)^{\frac{1}{2}} \\ &\approx \left(1 + \frac{\sum_{i=p+1}^{m} c_i^2}{c_p^2}\right)^{\frac{1}{2}},\end{aligned}$$

if

$$|c_{i+1}| \gg |c_i|, \quad i = 1, \ldots, p-1,$$

and thus the LS problem is well conditioned if this inequality is satisfied and

$$\left(\frac{\sum_{i=p+1}^{m} c_i^2}{c_p^2}\right)^{\frac{1}{2}} = O(1).$$

It follows from (3) and (4) that the spaces spanned by b and x_{ls} can be specified:

- The dominant components of b lie in the space spanned by the columns of U associated with the small singular values of A.
- The dominant components of x_{ls} lie in the space spanned by the columns of V associated with the small singular values of A.

These results on the spaces spanned by b and x_{ls} for ill conditioned and well conditioned LS problems confirm the entries in Table 2 for Example 1.

An upper bound for $\eta(A, b)$ is established in Theorem 3 [17, §4]. This theorem shows that $\eta(A, b)$ may be infinite, which marks a difference between it and $\kappa(A)$, which is finite if A has full rank.

Theorem 3 *Let θ be the angle between b and its component that lies in the column space of A. The condition numbers $\kappa(A)$ and $\eta(A, b)$ are related by*

$$\eta(A, b) \leq \frac{\kappa(A)}{\cos\theta}, \tag{7}$$

where, from (4) and the SVD of A,

$$\cos\theta = \frac{\|Ax_{\text{ls}}\|}{\|b\|} = \left(\frac{\sum_{i=1}^{p} c_i^2}{\sum_{i=1}^{m} c_i^2}\right)^{\frac{1}{2}}.$$

The ratio $\cos\theta$ is related to the error e_{ls} in x_{ls},

$$e_{\text{ls}}^2 = \frac{\|b - Ax_{\text{ls}}\|^2}{\|b\|^2} = 1 - \cos^2\theta,$$

and thus e_{ls} is approximately equal to its maximum value of one if θ is large. Also, it follows from (7) that the upper bound of $\eta(A, b)$ increases rapidly as θ approaches 90 degrees. In particular, $\eta(A, b)$ is infinite when $\cos\theta = 0$, that is, b is orthogonal to the space spanned by the columns of A, even if A has full rank, but $\kappa(A)$ is finite if A has full rank. The difference between the condition numbers $\kappa(A)$ and $\eta(A, b)$ for the Hilbert matrix is considered in [15, Ex. 2.2].

The effective condition number $\eta(A, b)$ would appear to overcome the disadvantage of the condition number $\kappa(A)$ because it is a function of A and b, but it follows from (5) that the denominator of the expression for $\eta(A, b)$ contains the term $\|x_{\text{ls}}\|$. It therefore follows that if x_{ls} is unstable with respect to a perturbation in b, then $\eta(A, b)$ is also unstable [18, §4]. This instability is quantified in Theorem 4 by considering the relative error in $\eta(A, b)$ due to a perturbation in b.

Theorem 4 *The relative error in $\eta(A,b)$ due to a first order perturbation δb in b is*

$$\max_{\delta b} \frac{\Delta\eta(A,b)}{\Delta b} = 1 + \eta(A,b), \tag{8}$$

where

$$\Delta\eta(A,b) = \frac{|\delta\eta(A,b)|}{\eta(A,b)},$$

and

$$\Delta b = \frac{\|\delta b\|}{\|b\|}.$$

PROOF It follows from (5) that

$$\eta^2(A,b+\delta b) = \frac{\left\|A^\dagger\right\|^2 \|b+\delta b\|^2}{\left\|A^\dagger(b+\delta b)\right\|^2},$$

where

$$\begin{aligned} \|b+\delta b\|^2 &= b^T b + 2b^T\delta b, \\ \left\|A^\dagger(b+\delta b)\right\|^2 &= b^T C b\left(1+\frac{2\delta b^T C b}{b^T C b}\right), \end{aligned}$$

to first order in δb, and

$$C = \left(A^\dagger\right)^T A^\dagger \quad \text{and} \quad b^T C b = \left\|A^\dagger b\right\|^2.$$

It follows that to first order in δb,

$$\begin{aligned} \eta^2(A,b+\delta b) &= \frac{\left\|A^\dagger\right\|^2 \|b+\delta b\|^2}{\left\|A^\dagger(b+\delta b)\right\|^2} \\ &= \left(\frac{\left\|A^\dagger\right\|^2 \left(b^T b + 2\delta b^T b\right)}{b^T C b}\right)\left(1 - \frac{2\delta b^T C b}{b^T C b}\right) \\ &= \left(\frac{\left\|A^\dagger\right\|^2}{b^T C b}\right)\left(b^T b - \frac{2\delta b^T C b\left(b^T b\right)}{b^T C b} + 2\delta b^T b\right) \\ &= \eta^2(A,b) + \delta b^T\left(\frac{2\left\|A^\dagger\right\|^2}{b^T C b}\right)\left(b - \frac{C b\left(b^T b\right)}{b^T C b}\right) \\ &= \eta^2(A,b)\left(1 + 2\left(\frac{\delta b^T}{\|b\|^2}\right)\left(I - \frac{C b b^T}{b^T C b}\right) b\right), \end{aligned}$$

and thus to first order in δb,

$$\eta(A,b+\delta b) = \eta(A,b)\left(1 + \left(\frac{\delta b^T}{\|b\|^2}\right)\left(I - \frac{C b b^T}{b^T C b}\right) b\right).$$

The relative error in $\eta(A,b)$ is therefore

$$\frac{\eta(A,b+\delta b) - \eta(A,b)}{\eta(A,b)} = \left(\frac{\delta b^T}{\|b\|^2}\right)\left(I - \frac{C b b^T}{b^T C b}\right) b.$$

The magnitude of the left-hand side is $\Delta\eta(A,b)$ and thus

$$\begin{aligned} \frac{\Delta\eta(A,b)}{\Delta b} &\le \left\| I - \frac{C b b^T}{b^T C b}\right\| \le 1 + \frac{\left\|\left(A^\dagger\right)^T A^\dagger b b^T\right\|}{\left\|A^\dagger b\right\|^2} \\ &\le 1 + \eta(A,b), \end{aligned}$$

which establishes the result (8). □

Theorem 4 shows that $\eta(A,b)$ is unstable with respect to a perturbation in b if $\eta(A,b) \gg 1$, which is problematic because it implies that the denominator $\left\|\Sigma^{-1}(c+\delta c)\right\|$ of $\eta(A,b+\delta b)$ cannot be used to determine if the discrete Picard condition is satisfied in the presence of noise $\delta c = U^T\delta b$. In particular, it is shown in [17, §4] that if the exact solution satisfies the discrete Picard condition (6), then this condition is not satisfied by data that are perturbed by noise,

$$\frac{|c_i + \delta c_i|}{\sigma_i} \not\to 0 \quad \text{as} \quad i \to p,$$

and thus it cannot be determined if regularisation is required when b is perturbed by noise. This problem is addressed by assuming there exists prior information that the exact solution $x_{\rm ls}$ satisfies the discrete Picard condition, in which case it is shown in Sect. 4 that regularisation yields a stable solution whose error is small. This prior information is satisfied by many exact images and it explains the success of regularisation in image deblurring [19, §5.6].

4 Regularisation

It was shown in Sect. 3 that $x_{\rm ls}$ is unstable if the discrete Picard condition is satisfied. Regularisation must therefore be applied and it leads to the regularised LS problem whose solution is $x_0(\lambda)$,

$$\begin{aligned} x_0(\lambda) &= \arg\min_z \left\{\|Az-b\|^2 + \lambda\|z\|^2\right\} \\ &= (A^T A + \lambda I)^{-1} A^T b, \end{aligned} \tag{9}$$

if $m \ge n$, where $x_0(0) = x_{\rm ls}$ and $\lambda \ge 0$ is the regularisation parameter. Equation (9) defines a family of functions $x_0(\lambda)$ parameterised by λ, and methods for calculating the optimal value of λ are considered in Sect. 4.1. It cannot be used if $m < n$ because it requires the inverse of a singular matrix when $\lambda = 0$. The equivalent expression for $m < n$ follows from the equality $(A^T A)A^T \equiv A^T(AA^T)$ and the addition

of λA^T to both sides,

$$\begin{aligned}(A^TA)A^T + \lambda A^T &\equiv A^T(AA^T) + \lambda A^T \\ &\equiv A^T(AA^T + \lambda I),\end{aligned}$$

and thus

$$(A^TA + \lambda I)A^T \equiv A^T(AA^T + \lambda I).$$

It follows that

$$(A^TA + \lambda I)^{-1}A^T \equiv A^T(AA^T + \lambda I)^{-1},$$

and thus $x_0(\lambda)$ can be written as

$$x_0(\lambda) = \begin{cases} (A^TA + \lambda I)^{-1}A^Tb, & m \geq n, \\ A^T(AA^T + \lambda I)^{-1}b, & m < n. \end{cases} \tag{10}$$

The minimisation (9) yields the regularised solution $x_0(\lambda^*)$ where λ^* is the optimal value of λ, and thus $x_0(\lambda^*) \neq x_{\text{ls}}$. There is therefore an error in the regularised solution $x_0(\lambda^*)$ with respect to the exact solution x_{ls}. This error is known as the regularisation error, and it is acceptable if (i) it is small, and (ii) $x_0(\lambda^*)$ is stable, and much more stable than x_{ls}. This leads to the premise on which regularisation is based:

There is a trade-off between the error and stability of the regularised solution $x_0(\lambda^*)$: The solution $x_0(\lambda^*)$ is accepted because (i) its error with respect to x_{ls} is small, and (ii) it is stable, and much more stable than x_{ls}.

This is the trade-off that was discussed in Sect. 1, where the bias and variance are, respectively, the error and stability of the regularised solution. The trade-off is shown in Fig. 3, which shows the problem $\mathcal{P}$ that it is desired to solve. This problem, whose solution is x_{ls}, is ill conditioned and the neighbouring problem $\mathcal{P}^*$, whose solution is $x_0(\lambda^*)$, is well conditioned, and $x_0(\lambda^*)$ is a very good approximation to x_{ls}. It follows that if $e_{\text{reg}}(\lambda)$ is the regularisation error, then the trade-off requires that $x_0(\lambda^*)$ satisfy the following conditions in order that regularisation be applied to x_{ls},

$$e_{\text{reg}}(\lambda^*) = \frac{\|x_0(\lambda^*) - x_0(0)\|}{\|x_0(0)\|} \ll 1,$$

and

$$\eta(A, b, \lambda^*) \ll \eta(A, b, 0) = \eta(A, b),$$

where $\eta(A, b, \lambda)$ is the effective condition number of the regularised LS problem (9). It is shown in Theorem 5 [17, §5.1] and Theorem 6 [17, §5.2] that these conditions are satisfied for $\lambda = \lambda^*$ if the discrete Picard condition is satisfied, which is therefore a necessary condition for the application of regularisation.

Theorem 5 *The regularisation error $e_{\text{reg}}(\lambda)$ is*

$$e_{\text{reg}}(\lambda) = \lambda \left(\frac{\sum_{i=1}^{p} \left(\frac{c_i}{\sigma_i}\right)^2 \frac{1}{(\sigma_i^2+\lambda)^2}}{\sum_{i=1}^{p} \left(\frac{c_i}{\sigma_i}\right)^2} \right)^{\frac{1}{2}},$$

and if the optimal value λ^* of λ satisfies

$$\begin{aligned} &\lambda^* \ll \sigma_i^2, \quad i = 1, \ldots, r-1, \\ &\lambda^* \approx \sigma_i^2, \quad i = r, \\ &\lambda^* \gg \sigma_i^2, \quad i = r+1, \ldots, p, \end{aligned} \tag{11}$$

then

(i)

$$e_{\text{reg}}(\lambda^*) \approx \frac{\lambda^*}{\sigma_1^2} \ll 1,$$

if the discrete Picard condition (6) is satisfied.

(ii)

$$e_{\text{reg}}(\lambda^*) \approx \left(\frac{p-r}{p}\right)^{\frac{1}{2}} < 1,$$

if

$$\frac{|c_i|}{\sigma_i} \approx 1, \qquad i = 1, \ldots, p.$$

(iii)

$$e_{\text{reg}}(\lambda^*) \approx \frac{\lambda^*}{\sigma_p^2 + \lambda^*} \approx 1,$$

if

$$\frac{|c_{i+1}|}{\sigma_{i+1}} \gg \frac{|c_i|}{\sigma_i}, \qquad i = 1, \ldots, p-1.$$

It follows that $e_{\text{reg}}(\lambda^*) \ll 1$ if the discrete Picard condition is satisfied, and that the two other conditions on the ratios $|c_i|/\sigma_i$ yield large errors $e_{\text{reg}}(\lambda^*)$. It is also necessary that the effective condition number $\eta(A, b) = \eta(A, b, \lambda = 0)$ of the LS problem be extended to the regularised LS problem (9). This extension follows from Theorem 2 and it is stated in Theorem 6.

Theorem 6 *The effective condition number of the regularised LS problem (9) is*

$$\eta(A, b, \lambda) = \max_{\delta b \in \mathbb{R}^m} \frac{\Delta x_0(\lambda)}{\Delta b} = \begin{cases} \frac{\|(\Sigma^T\Sigma+\lambda I)^{-1}\Sigma^T\|\|c\|}{\|(\Sigma^T\Sigma+\lambda I)^{-1}\Sigma^Tc\|}, & m \geq n, \\ \frac{\|\Sigma^T(\Sigma\Sigma^T+\lambda I)^{-1}\|\|c\|}{\|\Sigma^T(\Sigma\Sigma^T+\lambda I)^{-1}c\|}, & m < n, \end{cases}$$

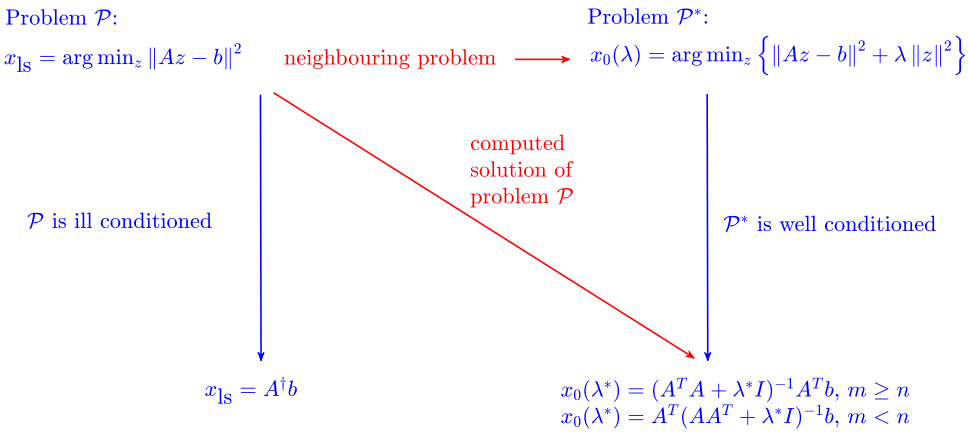

Fig. 3 The problem $\mathcal{P}$ that it is desired to solve and the problem $\mathcal{P}^*$ that is solved by regularisation

and if the discrete Picard condition is satisfied, then

$$\eta(A, b, \lambda^*) \approx \gamma \left(\frac{\sigma_1}{\sigma_r}\right) \ll \frac{\sigma_1}{\sigma_p} = \kappa(A), \tag{12}$$

where $1/2 \leq \gamma \leq 1$, *and* λ^* *and* r *are defined in (11).*

The result (12) follows from (11) because regularisation removes the singular values σ_i, $i = r+1, \ldots, p$, from $x_0(0)$.

The regularisation parameter λ is continuous, but it is an integer in truncated SVD (TSVD), which can also be used to impose regularisation on x_{ls} [7, §3.2]. In this method, the sum (3) is truncated at an integer $t < p$,

$$x_1(t) = \sum_{i=1}^{t} \left(\frac{c_i}{\sigma_i}\right) v_i, \tag{13}$$

where $x_1(p) = x_{\text{ls}}$. If the discrete Picard condition is satisfied, this truncation yields a stable solution whose regularisation error is small.

Tikhonov regularisation and TSVD achieve the same objective—the imposition of stability on the solution of an ill conditioned LS problem—but the methods achieve it slightly differently. This is most easily seen by writing (10) as

$$x_0(\lambda) = \sum_{i=1}^{p} \left(f_i(\lambda)\left(\frac{c_i}{\sigma_i}\right)\right) v_i,$$

where v_i is the ith column of V and $f_i(\lambda)$ is a filter,

$$f_i(\lambda) = \frac{\sigma_i^2}{\sigma_i^2 + \lambda}, \qquad i = 1, \ldots, p.$$

It follows that Tikhonov regularisation removes the small singular values from x_{ls} by the filter $f_i(\lambda)$ whose cut-off is defined by the continuous parameter λ. This must be compared with TSVD because it follows from (13) that this removal is achieved by a discrete filter, that is, a filter with an infinitely sharp cut-off. If the singular values of A are well separated, then the two methods yield very similar results. The results from the methods differ, however, if the singular values are closely spaced because a singular value that is removed by TSVD will not contribute to the regularised solution, but its contribution to the solution from Tikhonov regularisation will reduce smoothly to zero as λ increases from λ^*.

4.1 The optimal value of the regularisation parameter

The effectiveness of regularisation depends on the value of the regularisation parameter λ. If it is too small, then $x_0(\lambda)$ has too much noise, but a large portion of the exact solution is filtered from $x_0(\lambda)$ if it is too large. Two methods for computing the optimal value λ^* of λ are the L-curve [7, §4.6] and generalised cross validation (GCV) [7, §7.4]. It is shown in [7, §7.5.1] that if the discrete Picard condition is satisfied and the noise is white, a parametric plot of $\log_{10} \|x_0(\lambda)\|$ against $\log_{10} \|Ax_0(\lambda) - b\|$ as a function of λ is a curve in the form of an L, as shown in Fig. 4(i) and which gives rise to the name of the method. As λ increases from $\lambda = 0$, $\|x_0(\lambda)\|$ decreases and $\|Ax_0(\lambda) - b\|$ is approximately constant, until $\lambda = \lambda^*$, which is the value of λ in the corner of the L. As λ increases further, $\|x_0(\lambda)\|$ is approximately constant and $\|Ax_0(\lambda) - b\|$ increases, and thus $\|x_0(\lambda)\|$ and $\|Ax_0(\lambda) - b\|$ are approximately minimised simultaneously for $\lambda = \lambda^*$, which is therefore the optimal value of λ.

If the discrete Picard condition is not satisfied, a parametric plot of $\log_{10} \|x_0(\lambda)\|$ against $\log_{10} \|Ax_0(\lambda) - b\|$ as a function of λ has the form shown in Fig. 4(ii). As λ increases from $\lambda = 0$, $\|x_0(\lambda)\|$ is approximately constant and $\|Ax_0(\lambda) - b\|$ increases until $\lambda = \lambda_0$, and $\|x_0(\lambda)\|$ decreases and $\|Ax_0(\lambda) - b\|$ is approximately constant as λ increases further. It follows there does not exist a value of λ that approximately minimises simultaneously $\|x_0(\lambda)\|$ and

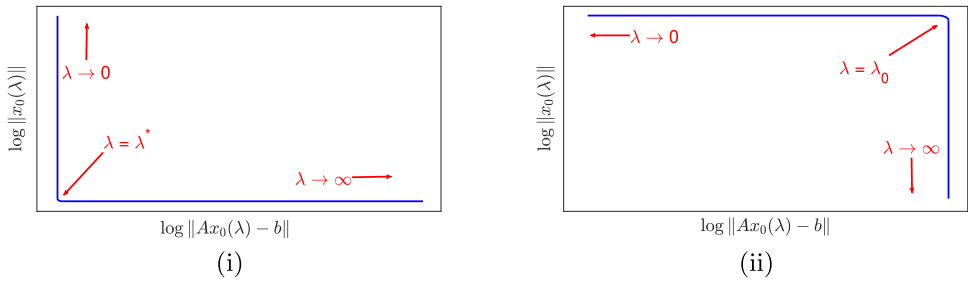

Fig. 4 A parametric plot of $\log \|x_0(\lambda)\|$ against $\log \|Ax_0(\lambda) - b\|$ when (i) the discrete Picard condition is satisfied, and (ii) the discrete Picard condition is not satisfied

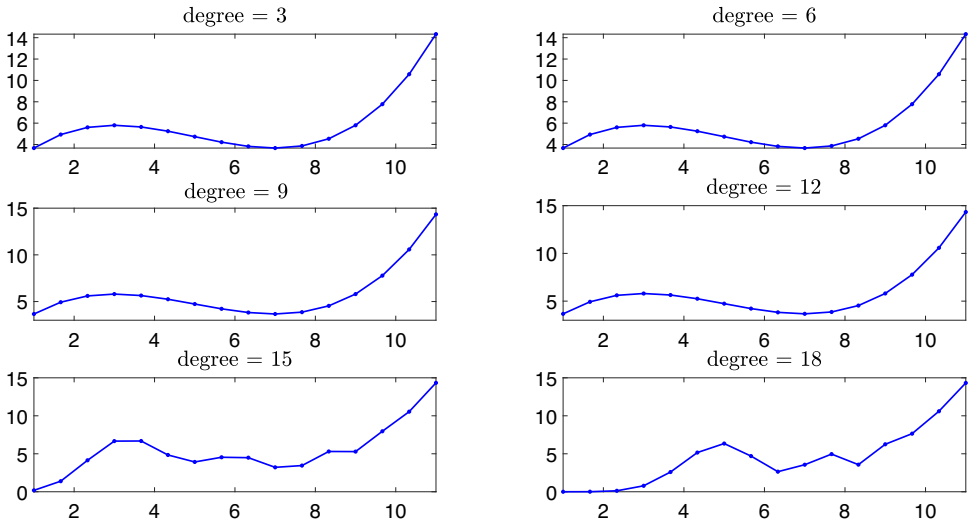

Fig. 5 The computed polynomials $\tilde{f}_d(x)$ for $d = 3, 6, 9, 12, 15, 18$, for Example 3

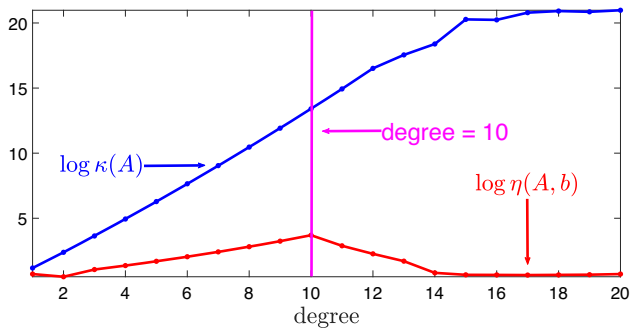

Fig. 6 The variation of the condition number $\log_{10} \kappa(A)$ and effective condition number $\log_{10} \eta(A, b)$ with the degree d of the approximating polynomial, for Example 3

$\|Ax_0(\lambda) - b\|$, and thus $\lambda^* = 0$, that is, regularisation must not be applied.

The GCV is different because it is based on the premise that if an arbitrary component b_k of b is omitted, then the regularised solution with $\lambda = \lambda^*$ obtained with this omission should yield a value of b_k whose error is small. The GCV also requires that the discrete Picard condition be satisfied. Comparisons of the L-curve, GCV and other methods for the determination of the value of λ^* are in [20, 21] and the references therein, and it is shown it is dependent on several parameters, including the noise level.

5 Examples

This section contains examples of regression that consider the effective condition number, regularisation and overfitting. Example 3 considers the situation in which data that can be modelled by a polynomial of degree three are represented by polynomials of degree d, where $1 \leq d \leq 20$. Overfitting occurs for large values of d but the LS problem is well conditioned for these values of d and thus regularisation must not be applied. Examples 4 and 5 consider solutions that underfit and overfit the data, and the approximating polynomials of high degree overfit the data but regularisation must not be applied because the LS problems are well conditioned. Example 6 is different because regularisation must be applied since the LS problem is ill conditioned and a small level of overfitting occurs, but a reduction in the number of basis functions yields a well conditioned LS problem whose solution does not exhibit overfitting.

All computations were performed using MATLAB 2022a.

Example 3 The LS problem was solved in order to compute the coefficients of polynomials $\tilde{f}_d(x)$ of degrees $d = 1, 2, \ldots, 19, 20$, that model the data defined by $f(x)$,

$$f(x) = \frac{1}{15}x^3 - x^2 + \frac{21}{5}x + \frac{2}{5}.$$

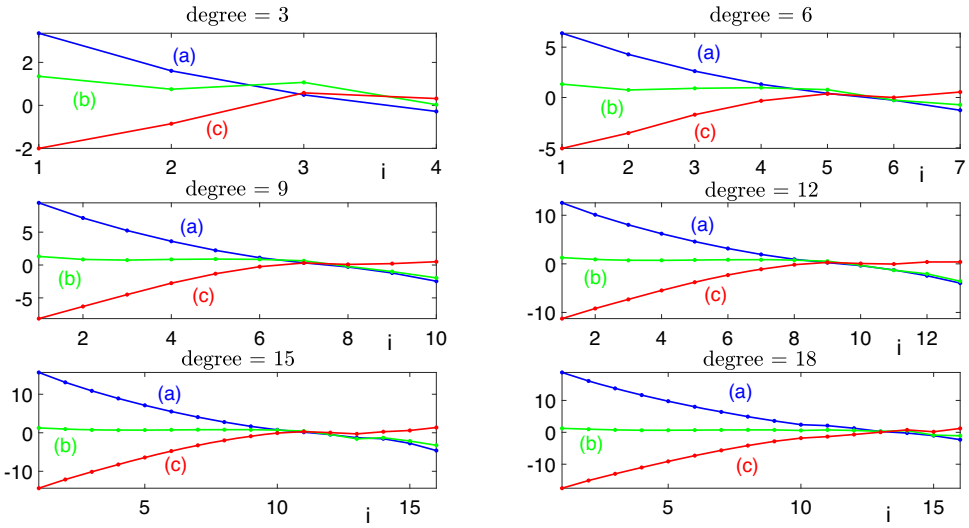

Fig. 7 The variation of **a** $\log_{10}\sigma_i$, **b** $\log_{10}|c_i|$ and **c** $\log_{10}|c_i|/\sigma_i$, for Example 3

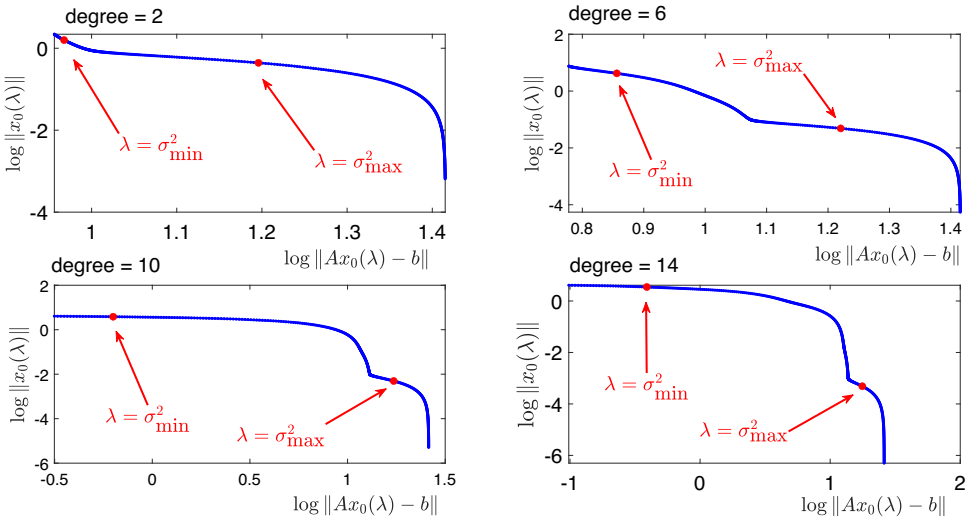

Fig. 8 The L-curves for polynomials of degrees $d = 2, 6, 10, 14$, for Example 3. The values of λ defined by the square of the minimum and maximum singular values of A are marked

The function $f(x)$ was represented by 16 points x_i, i=1, ..., 16, uniformly distributed in the interval I=[1, ..., 11]. Figure 5 shows the polynomials $\tilde{f}_d(x)$ of degrees d = 3, 6, 9, 12, 15, 18, and

Figure 6 shows the variation of the condition number $\log_{10}\kappa(A)$ and effective condition number $\log_{10}\eta(A,b)$ with d. Overfitting does not occur for $3 \le d \le 12$, even though $\eta(A,b) \approx 10$ for $d = 3$ and $\eta(A,b) \approx 10^4$ for $d = 10$. The maximum value of $\eta(A,b)$ occurs when $d = 10$, after which it decreases and the LS problem is well conditioned ($\eta(A,b) \approx 1$) for $d \ge 14$. Furthermore, Fig. 5 shows that overfitting occurs for large values of d and Fig. 6 shows that $\eta(A,b) \approx 1$ for these values of d.

Figure 7 shows the variation of $\log_{10}\sigma_i$, $\log_{10}|c_i|$ and $\log_{10}|c_i|/\sigma_i$ with i for polynomials of degrees $d = 3, 6, 9, 12, 15, 18$. The dominant components of $|c_i|/\sigma_i$ are defined by large values of i and thus the discrete Picard condition is not satisfied. Figure 8 shows the L-curves for polynomials of degrees 2, 6, 10, 14, and they are similar to Fig. 4(ii). The curves for degrees 6, 10, 14 exhibit an L but they do not display a rapid decrease in $\|x_0(\lambda)\|$ as λ increases from $\lambda = 0$, as shown in Fig. 4(i). It follows that regularisation cannot be used to solve the problem of overfitting discussed above.

Figure 9 shows the number of sign changes in each column of V for $d = 1, \ldots, 20$. It is seen that A^TA is an oscillation matrix for $d = 1, \ldots, 13$, because the ith column of V has $i - 1$ sign changes for these values of d. Furthermore, it follows from Theorem 1 that the entries of $x_{\rm ls}$ have many sign changes for $1 \ll d \le 12$, but Fig. 5 shows that the regression curves are smooth.

Noise was not added to the coefficients of $f(x)$, and Figs. 6 and 7 show that the LS problem is well conditioned. It follows that the effect of noise on the computed coefficients would be minor. □

Example 4 Figure 10 shows 25 points, corresponding to a sampling interval of 0.075, of the function $f(x)$, which is

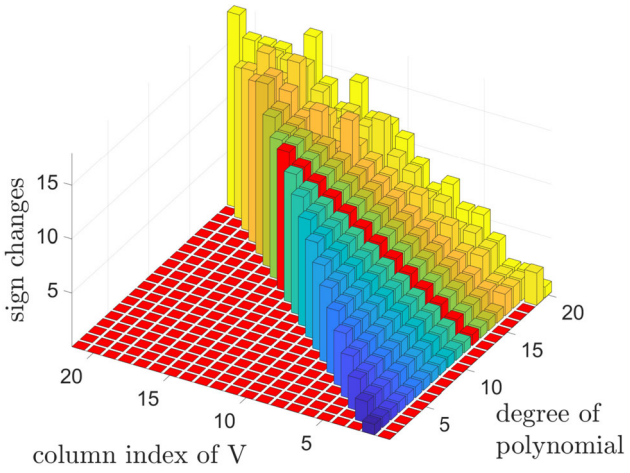

Fig. 9 The number of sign changes in each column of V for polynomials of degrees $d = 1, \ldots, 20$, for Example 3. The red bars show the maximum degree d of the approximating polynomial for which $A^T A$ is an oscillation matrix

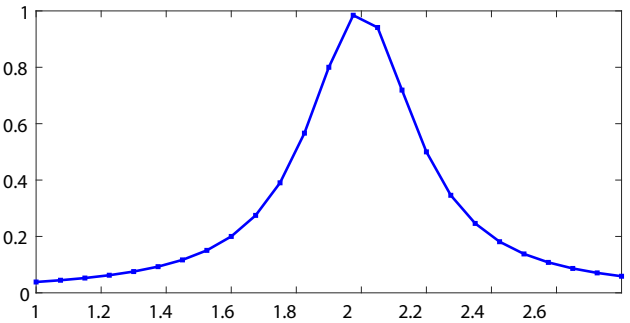

Fig. 10 The function $f(x)$ defined in (14) for Example 4

similar to the function in [22, §2.3],

$$f(x) = \frac{1}{1 + 25(x-2)^2}, \tag{14}$$

for $1 \le x \le 2.8$. The sampling interval was increased to 0.150, which corresponds to 13 points, and approximating polynomials of degrees $d = 1, \ldots, 24$, were computed. It follows that $A \in \mathbb{R}^{13\times(d+1)}$, and Fig. 11 shows the variation of $\log_{10} \kappa(A)$ and $\log_{10} \eta(A, b)$ with d. It is seen that $\log_{10} \kappa(A)$ increases almost linearly for $d \le 12$, it attains its maximum value when $d = 12$, that is, A is square, and it is then approximately constant at a slightly lower value. The effective condition number displays much less variation, its maximum value is about 100 and this figure is therefore similar to Fig. 6.

Figure 12 shows the variation of $\log_{10} \sigma_i$, $\log_{10} |c_i|$ and $\log_{10} |c_i|/\sigma_i$ with i for polynomials of degrees $d = 4, 8, 12, 16, 20, 24$. The graphs in the figure show that the constants $|c_i|$ are approximately independent of i and d, and that the ratios $|c_i|/\sigma_i$ increase with i, which shows that the discrete Picard condition is not satisfied. It follows that $x_{\rm ls}$

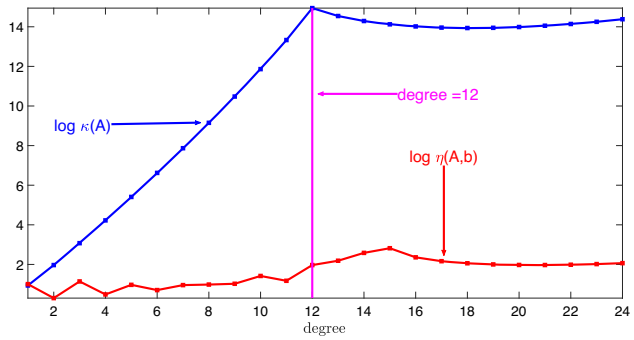

Fig. 11 The variation of the condition number $\log_{10} \kappa(A)$ and effective condition number $\log_{10} \eta(A, b)$ with the degree d of the approximating polynomial, for Example 4

is dominated by the small singular values of A and regularisation must not be applied because it would remove the contribution of these singular values from $x_{\rm ls}$. Figure 13 shows the L-curves for polynomials of degrees 5, 10, 15, 20 and they are similar to Fig. 4(ii). The curve for a polynomial of degree 20 has an L, but as for Example 3, $\|x_0(\lambda)\|$ does not decrease rapidly as λ increases from $\lambda = 0$.

The LS problem was solved for polynomials of degrees $d = 1, \ldots, 24$, and the number of sign changes in each column of V, for each LS problem, was computed. The results are shown in Fig. 14 and it is seen that for each value of $d = 1, \ldots, 12$, there are $i - 1$ sign changes in the ith column of V for the overdetermined problem ($1 \le d \le 11$) and the interpolation problem ($d = 12$).

Figure 15 shows the function (14), a noisy form $\tilde{f}(x)$, with signal-to-noise ratio (SNR) $\|b\|/\|\delta b\| = 7$, of $f(x)$, approximations with polynomials of degree $d = 22$ that overfit $f(x)$ and $\tilde{f}(x)$, and an approximation with a polynomial of degree $d = 2$ that underfits $\tilde{f}(x)$. Overfitting of $f(x)$ and $\tilde{f}(x)$ occurs at the tails of the curve, but regularisation cannot be used to reduce it because Fig. 12 shows that the ratios $|c_i|/\sigma_i$ increase approximately monotonically and thus the exact solution $x_{\rm ls}$ is dominated by the small singular values of A. Regularisation removes these singular values from $x_{\rm ls}$ and thus its application would lead to a large error. It follows that regularisation cannot be used to solve the problem of overfitting of a high degree polynomial approximation to the function $f(x)$ defined in (14).

The oscillatory curves in Fig. 15 demonstrate Runge's phenomenon, also known as the Gibbs phenomenon, which arises when a polynomial of high degree approximates a function. It is characterised by large oscillations and errors near the edges of the interval, and it shows that interpolation by polynomials of high degree leads to overfitting.

Noise was not added to the 25 points, that is, to b, in Fig. 10. Figure 12 shows that the ratios $|c_i|/\sigma_i$ increase with i and thus the solution of the LS problem is dominated by the small

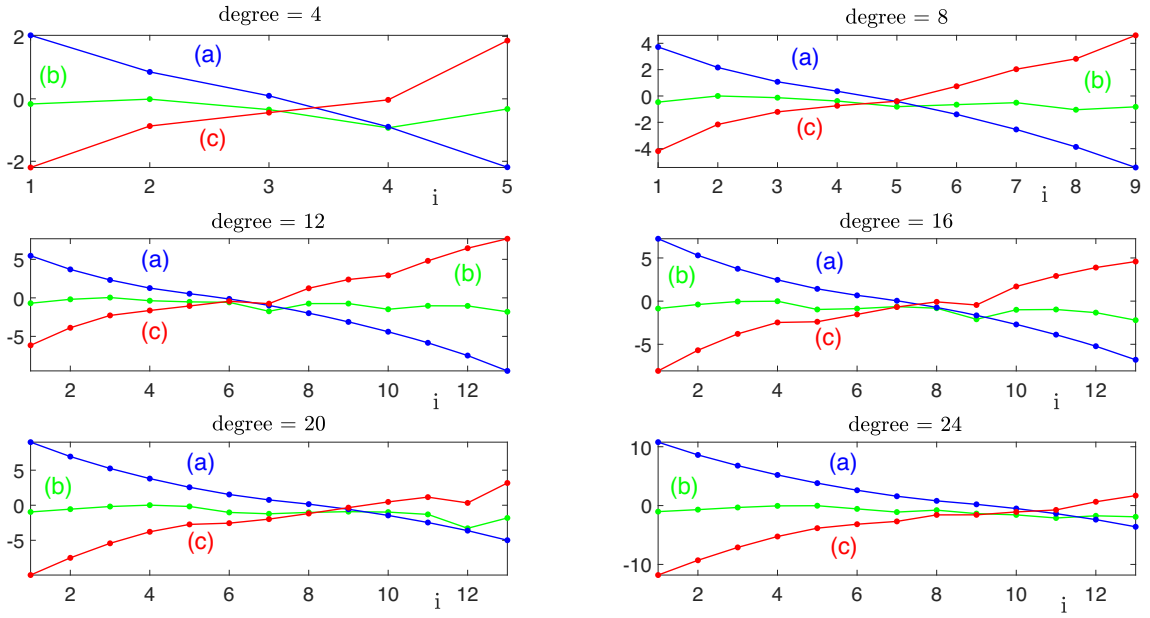

Fig. 12 The variation of **a** $\log_{10}\sigma_i$, **b** $\log_{10}|c_i|$, and **c** $\log_{10}|c_i|/\sigma_i$ with i, for polynomials of degrees $d = 4, 8, 12, 16, 20, 24$, for Example 4

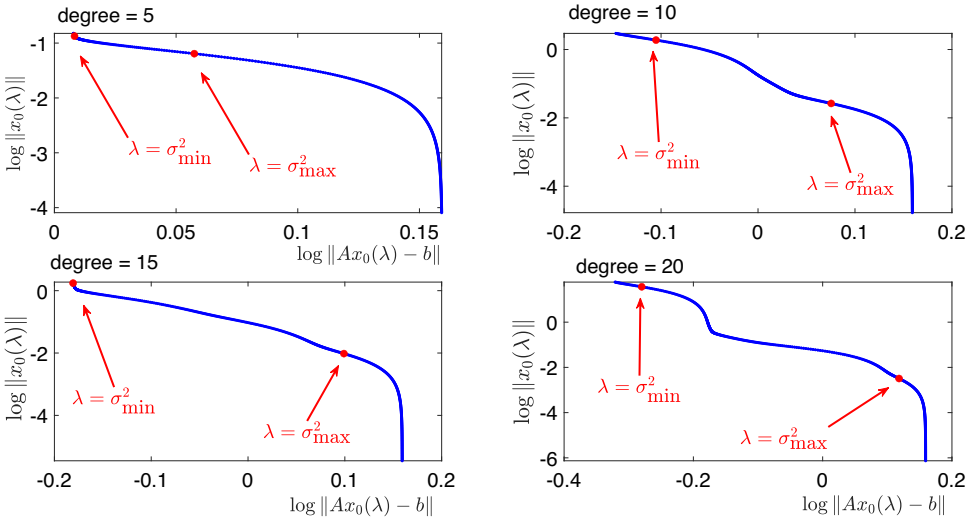

Fig. 13 The L-curves for polynomials of degrees $d = 5, 10, 15, 20$, for Example 4. The values of λ defined by the square of the minimum and maximum singular values of A are marked

singular values of A. The addition of noise would therefore cause a small error in this solution. □

Example 5 Example 4 was repeated for the function [22, §2.3],

$$f(x) = \frac{1}{1 + 25x^2}, \tag{15}$$

for $-1 \le x \le 0.8$, which differs from the function (14) by a shift of two units of the independent variable x but it is otherwise identical. Figure 16 shows the variation of $\log_{10}\kappa(A)$ and $\log_{10}\eta(A, b)$ with the degree d of the approximating polynomial. The maximum value of $\eta(A, b)$ is $10^{1.5} = 32$ and thus the LS problem is well conditioned. The figure is similar to Figs. 6 and 11 but $\kappa(A)$ is much smaller, by several orders of magnitude, and the variation of $\eta(A, b)$ with d is very similar for the functions (14) and (15). Figure 17 shows the variation of $\log_{10}\sigma_i$, $\log_{10}|c_i|$ and $\log_{10}|c_i|/\sigma_i$ with i for polynomials of degrees $d = 4, 8, 12, 16, 20, 24$, for the function (15). Figure 16 is confirmed by the graphs in Fig. 17, which show that the ratios $\log_{10}|c_i|/\sigma_i$ increase with i, albeit not monotonically. The graphs in the figure are very similar to the graphs in Fig. 12 but the singular values span a smaller range, which is expected from comparison of $\kappa(A)$ in Figs. 11 and 16. Figure 18 shows the L-curves for polynomials of degrees 5, 10, 15, 20, and they are similar to Fig. 4(ii). They confirm that the discrete Picard condition is not satisfied, and thus regularisation must not be applied.

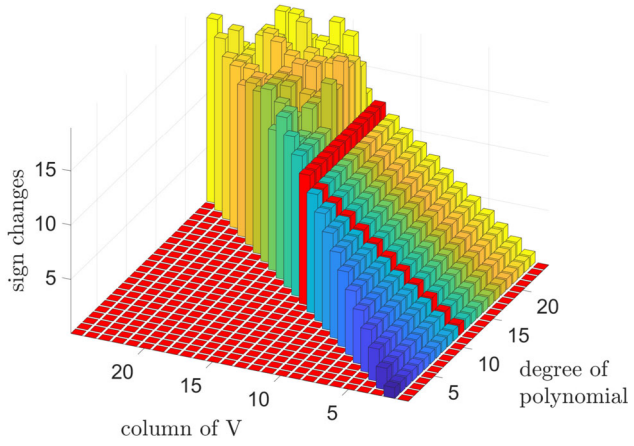

Fig. 14 The number of sign changes in each column of V for polynomials of degrees $d = 1, \ldots, 24$, for Example 4. The coefficient matrix A is square when $d = 12$, which is shown by the red bars

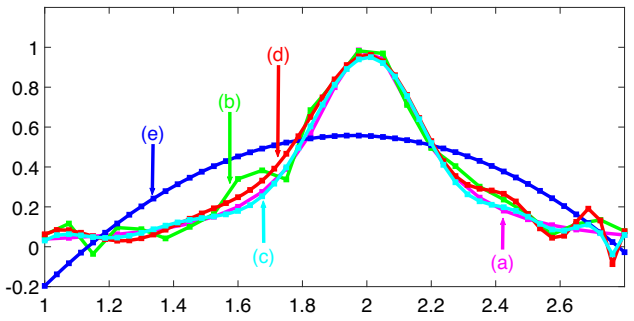

Fig. 15 **a** The exact function $f(x)$, **b** a noisy form $\tilde{f}(x)$ of $f(x)$, **(c)** an approximation that overfits $f(x)$, **d** an approximation that overfits $\tilde{f}(x)$ and **(e)** an approximation that underfits $\tilde{f}(x)$, for Example 4

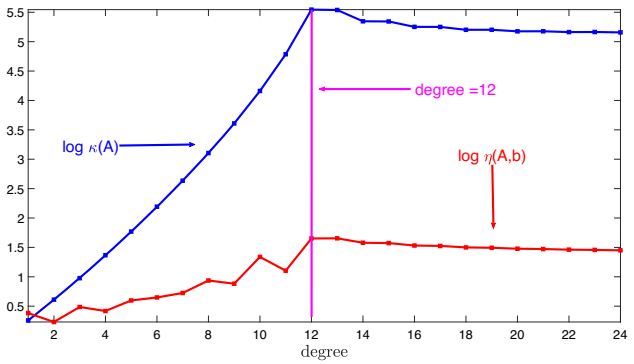

Fig. 16 The variation of the condition number $\log_{10} \kappa(A)$ and effective condition number $\log_{10} \eta(A, b)$ with the degree d of the approximating polynomial, for Example 5

Figure 19 shows the exact function $f(x)$, (15), a noisy form $\tilde{f}(x)$ of $f(x)$ when noise with SNR $= 25$ is added to b, approximations that overfit $f(x)$ and $\tilde{f}(x)$, and an approximation that underfits $\tilde{f}(x)$. There is considerable overfitting, and Figs. 16 and 17 show that $x_{\rm ls}$ is stable and thus regularisation must not be imposed. It follows that regularisation cannot be used to solve the problem of overfitting that arises when a polynomial of high degree approximates the function (15). Runge's phenomenon is evident in Fig. 19 to a much greater degree than in Fig. 15 and thus the extent of overfitting is dependent on the interval of the independent variable.

Noise was not added to the function values $f(x_i), i = 1, \ldots, 13$, of the function (15), and its effect would be minor because $\eta(A, b) \approx 32$. Figure 17 confirms that the solution of the LS problem is stable because it is dominated by the small singular values of A. □

Example 6 Figure 20 shows $m = 26$ points $\{x_i, f(x_i)\}_{i=1}^m$ uniformly distributed in the interval $I = [1, \ldots, 21]$ and a curve formed from $n = 17$ exponential basis functions,

$$f(x_i) = \sum_{j=1}^{n} a_j \exp\left(\frac{-(x_i - d_j)^2}{2\sigma^2}\right), \tag{16}$$

for $i = 1, \ldots, m$, that interpolates the points, where $\sigma = 2.35$ and the centres $d = \{d_j\}_{j=1}^n$ of the exponential functions are uniformly distributed in I.

Figure 21(i) shows the variation of the constants $\log_{10} |c_i|$, the singular values $\log_{10} \sigma_i$ and the ratios $\log_{10} |c_i|/\sigma_i$ with i, and it is seen that the constants $|c_i|$ decay to zero faster than the singular values σ_i decay to zero, from which it follows that $x_{\rm ls}$ is unstable because the discrete Picard condition (6) is satisfied. Figure 22 shows the L-curve and it is very similar to the L-curve in Fig. 4(i), which confirms that the discrete Picard is satisfied. The instability of $x_{\rm ls}$ is apparent in Fig. 21(ii), which shows the ratios $\log_{10} |c_i|/\sigma_i$ and $\log_{10} |c_i + \delta c_i|/\sigma_i$, after the addition of noise with SNR $= 25$ to b, because $\log_{10} |c_i + \delta c_i|/\sigma_i$ is dominated by noise for $i \geq 7$.

It follows that the noise dominates the solution $x_{\rm ls} + \delta x_{\rm ls}$ and thus only the singular values σ_{16} and σ_{17} need be considered in this solution,

$$\frac{|c_i + \delta c_i|}{\sigma_i} \approx \frac{|\delta c_i|}{\sigma_i} \gg \frac{|c_i|}{\sigma_i}, \qquad i = 16, 17,$$

and thus

$$\begin{aligned} x_{\rm ls} + \delta x_{\rm ls} &\approx \sum_{i=n-1}^{n} \left(\frac{c_i + \delta c_i}{\sigma_i}\right) v_i \\ &\approx \left(\frac{\delta c_{16}}{\sigma_{16}}\right) v_{16} + \left(\frac{\delta c_{17}}{\sigma_{17}}\right) v_{17}. \end{aligned}$$

Figure 21(ii) shows that $|c_i + \delta c_i|/\sigma_i$ and $|c_i|/\sigma_i$ decrease at the same rate for $i = 1, \ldots, 6$, and that the noise δc_i manifests itself for $i = 7, \ldots, 17$. The regularised approximation of $x_{\rm ls}$ by TSVD (13) is therefore obtained by removing the singular

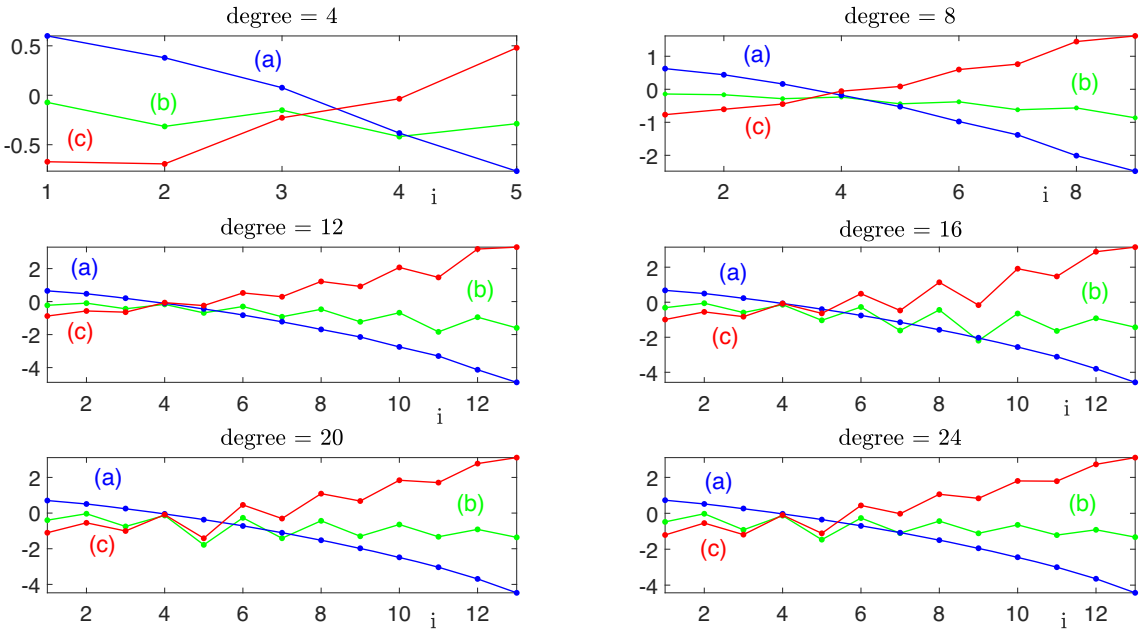

Fig. 17 The variation of **a** $\log_{10}\sigma_i$, **b** $\log_{10}|c_i|$, and **c** $\log_{10}|c_i|/\sigma_i$ with i, for polynomials of degrees $d = 4, 8, 12, 16, 20, 24$, for Example 5

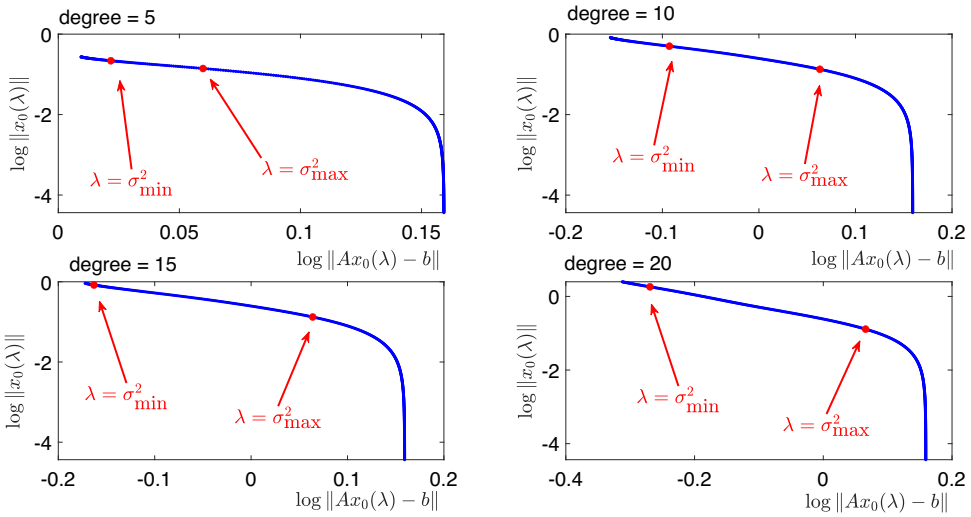

Fig. 18 The L-curves for polynomials of degrees $d = 5, 10, 15, 20$ for Example 5. The values of λ defined by the square of the minimum and maximum singular values of A are marked

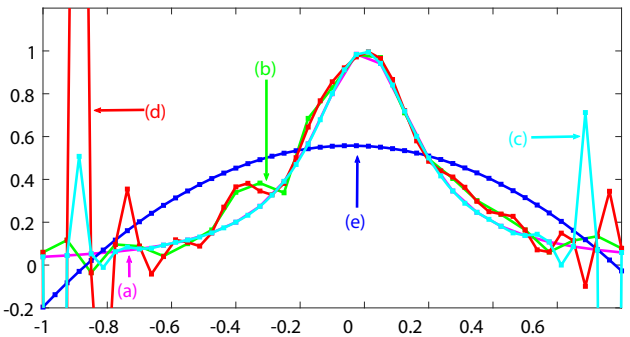

Fig. 19 **a** The exact function $f(x)$, **b** a noisy form $\tilde{f}(x)$ of $f(x)$, **c** an approximation that overfits $f(x)$, **d** an approximation that overfits $\tilde{f}(x)$, and **e** an approximation that underfits $\tilde{f}(x)$, for Example 5

values σ_i, $i = 7, \ldots, 17$, from the noisy solution,

$$x_1(t=6) = \sum_{i=1}^{6}\left(\frac{c_i + \delta c_i}{\sigma_i}\right)v_i \approx \sum_{i=1}^{6}\left(\frac{c_i}{\sigma_i}\right)v_i. \tag{17}$$

Figure 23(i) shows the coefficients $\log_{10}|a_i|$ of the exact solution x_{ls} and the coefficients $\log_{10}|a_i + \delta a_i|$ of the perturbed solution $x_{\text{ls}} + \delta x_{\text{ls}}$. Figure 23(ii) shows the coefficients a_i of the exact solution x_{ls}, the coefficients of the regularised solutions in which λ^* is computed from the L-curve and the GCV, and the coefficients of the regularised solution from

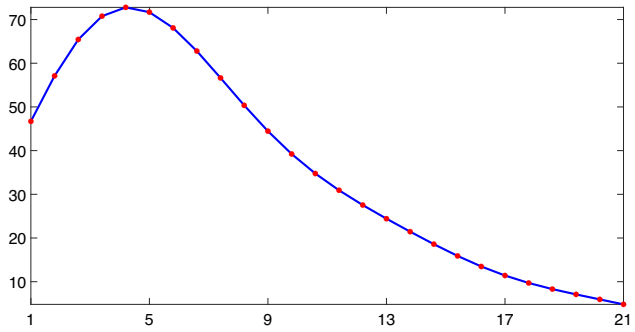

Fig. 20 A curve that interpolates $m = 26$ points in the interval $I = [1, \ldots, 21]$, for Example 6

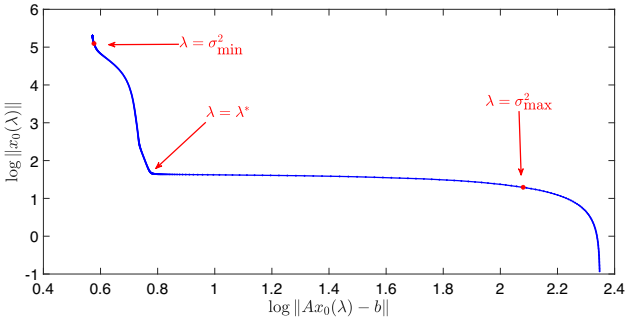

Fig. 22 The L-curve for Example 6. The values of λ defined by the square of the minimum and maximum singular values of A, and the optimal value $\lambda^* = 0.0187$ of λ, are marked

TSVD (17). The figure shows that TSVD yields the best regularised solution and that the regularised solution in which λ^* is computed from the L-curve has more oscillations than the regularised solution in which λ^* is computed from the GCV. The failure of the L-curve to yield a satisfactory regularised solution has been observed in some problems whose solutions are smooth [7, p. 190]. The decay of the value of λ^* computed from the L-curve as the noise level decreases to zero is considered in [23] and it is shown it is too rapid, which yields a regularised solution that does not converge to the exact solution. Experiments show that under smoothing may occur when the GCV is used, and this is observed in the regularised solution obtained from the GCV in Fig. 23(ii) [7, p. 197].

Theorem 4 shows that if x_{ls} is unstable with respect to a perturbation in b, then $\eta(A, b)$ is also unstable. This is confirmed in Fig. 24, which shows the variation of $\log_{10} \kappa(A)$, $\log_{10} \eta(A, b)$, and $\log_{10} \eta(A, b + \delta b)$ when noise δb with SNR $= 25$ is added to b, for $m = 26$ and $n = 1, \ldots, 32$. The noise causes a reduction in the value of $\eta(A, b)$ and the maximum value of this reduction is about seven orders of magnitude. Figure 24 shows that $\kappa(A)$ attains its maximum value when A is square ($m = n = 26$), after which it is approximately constant at a slightly lower value, and it therefore has the same properties as the graphs in Figs. 11 and 16 for the monomial basis.

Figure 25 shows the number of sign changes in the entries of the columns v_i of V for $m = 26$ as the number of basis functions increases from $n = 1$ to $n = 32$. There are $i - 1$ sign changes in v_i for $n \leq m$ and this sign change property is not present for $n > m$. The condition $n = m$ therefore marks the division between the satisfaction, and lack of satisfaction, of the sign change property. The matrix $A^T A$ is singular for $n > m$ and thus it is not an oscillation matrix for these values of n.

Figure 26 shows the curve that interpolates the $m = 26$ points, and the points $\tilde{b} = Ax_{\text{ls}} = AA^{\dagger}b$ for $n = 7$ basis functions and $n = 30$ basis functions. Figure 24 shows that $\eta(A, b) = 12.40$ and $\kappa(A) = 20.23$ when $n = 7$, and Fig. 26 shows that the error between this curve and the given curve is small, and that overfitting does not occur. The figure also shows there is a small level of overfitting when $n = 30$, but significantly less overfitting than in Figs. 15 and 19 for Examples 4 and 5, respectively. Regularisation must be applied when $n = 30$ in order to obtain a computationally reliable solution because Fig. 24 shows that $\eta(A, b) = 4.57 \times 10^9$ and $\kappa(A) = 9.77 \times 10^{13}$, and thus x_{ls} is unstable. □

6 Summary

This paper has considered the application of regularisation to regression and it has been shown it cannot solve the problem of overfitting. This result, which is consistent with previous results on physics informed neural networks and deep neural networks, was derived using a refined condition number of the LS problem, from which the discrete Picard condition was established. The satisfaction of this condition implies

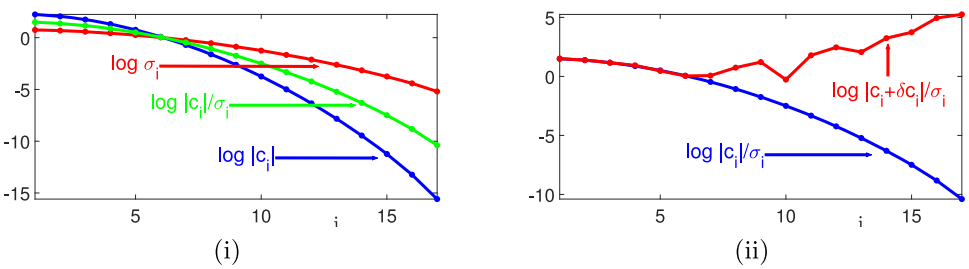

Fig. 21 (i) The constants $\log_{10} |c_i|$, the singular values $\log_{10} \sigma_i$ and the ratios $\log_{10} |c_i|/\sigma_i$, and (ii) the ratios $\log_{10} |c_i|/\sigma_i$ and the ratios $\log_{10} |c_i + \delta c_i|/\sigma_i$, for Example 6

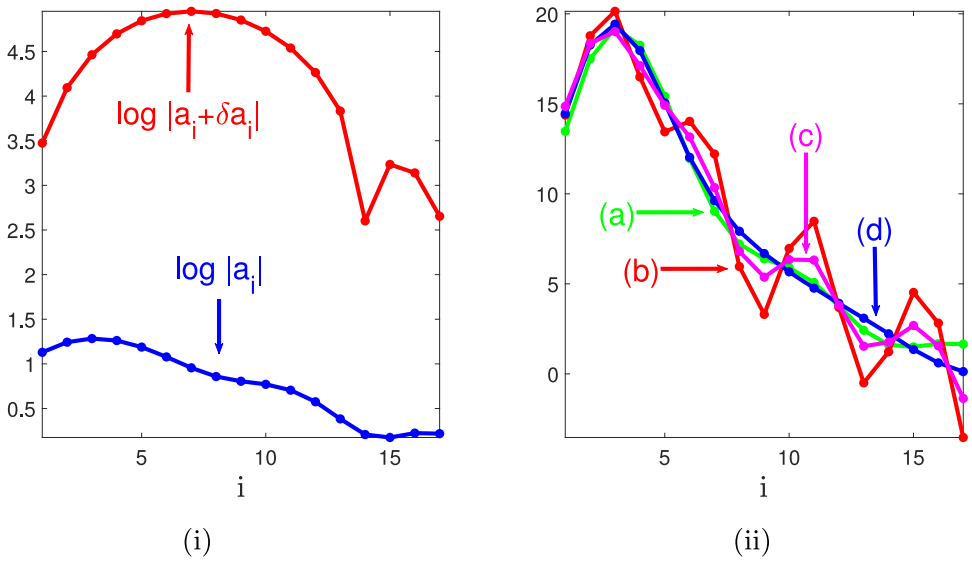

Fig. 23 (i) The coefficients $\log_{10}|a_i|$ and $\log_{10}|a_i + \delta a_i|$ of the exact and perturbed solutions $x_{\rm ls}$ and $x_{\rm ls} + \delta x_{\rm ls}$, respectively, and (ii) the coefficients of **(a)** the exact solution, **(b)** the regularised solution with λ^* computed from the L-curve, **(c)** the regularised solution with λ^* computed from the GCV, and **(d)** the regularised solution (17) computed by TSVD, for Example 6

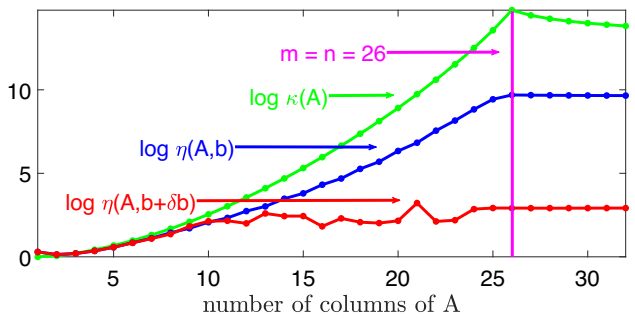

Fig. 24 The effective condition number $\log_{10}\eta(A, b)$, the effective condition number $\log_{10}\eta(A, b + \delta b)$ and the condition number $\log_{10}\kappa(A)$, for Example 6

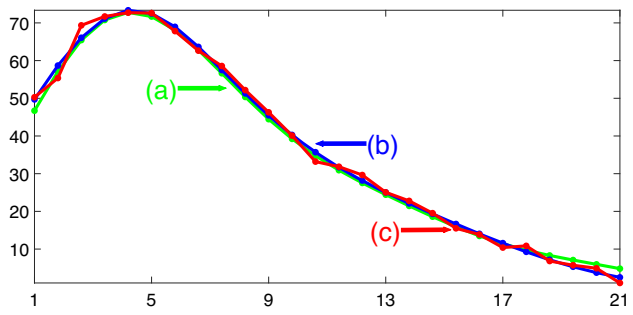

Fig. 26 **a** The function (16), **b** the computed data using $n = 7$ basis function, and **c** the computed data using $n = 30$ basis functions, for Example 6

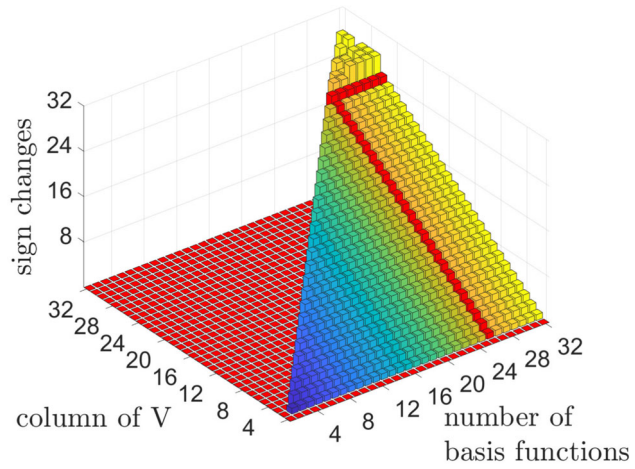

Fig. 25 The number of sign changes in each column of V for $n = 1, \ldots, 32$, basis functions, for Example 6. The coefficient matrix A is square when $n = m = 26$, which is shown by the red bars

that regularisation can be applied and it leads to the premise on which regularisation is based. Specifically, it is based on a trade-off between the error and condition of the LS problem because a significant reduction in its condition is associated with a small and acceptable error in the regularised solution. If, however, the discrete Picard condition is not satisfied, then the error in the regularised solution is large, which is therefore unacceptable. This trade-off between the condition of the LS problem and the error in the regularised solution is known as the bias-variance trade-off in machine learning.

Well conditioned LS problems for which the regression curve displays overfitting were presented, but it follows from their well conditioned property that regularisation must not be applied. The condition numbers of the coefficient matrices in these examples are large, which suggests the LS problems are ill conditioned. The effective condition numbers are, however, much smaller and they show that regularisation must not be applied because it would lead to a large error. Also, an ill conditioned LS problem that requires regularisation and for

which the regression curve exhibits a small level of overfitting was shown.

It has been shown that regularisation does not solve the problem of overfitting in regression, but the problem of overfitting persists. Research is therefore required to identify the cause of overfitting and to develop a method that minimises its adverse effects.

Author Contributions Dr Joab Winkler carried out all the work: (a) Theoretical analysis (b) The software that demonstrates the examples (c) The writing of the manuscript and proof reading

Data Availability No datasets were generated or analysed during the current study.

Declarations

Conflict of interest The authors declare no conflict of interest.

References

1. Bartlett, P., Long, P., Lugosi, G., Tsigler, A.: Benign overfitting in linear regression. Proc. Nat. Acad. Sci. **117**(48), 30063–30070 (2020)
2. Belkin, M., Hsu, D., Ma, S., Mandal, S.: Reconciling modern machine-learning practice and the classical bias-variance trade-off. Proc. Nat. Acad. Sci. **116**(32), 15849–15854 (2019)
3. Huang, T., Hogg, D.W., Villar, S.: Dimensionality reduction, regularization, and generalization in overparameterized regressions. SIAM J. Math. Data Sci. **4**(1), 126–152 (2022)
4. Bajaj, C., McLennan, L., Andeen, T., Roy, A.: Recipes for when physics fails: recovering robust learning of physics informed neural networks. Mach. Learn.: Sci. Technol. **4**, 015013 (2023)
5. Zhang, C., Bengio, S., Hardt, M., Recht, B., Vinyals, O.: Understanding deep learning (still) requires rethinking generalization. Comm. ACM **64**(3), 107–115 (2021)
6. Hansen, P.C.: The discrete Picard condition for discrete ill-posed problems. BIT **30**, 658–672 (1990)
7. Hansen, P.C.: Rank-Deficient and Discrete Ill-Posed Problems: Numerical Aspects of Linear Inversion. SIAM, Philadelphia, USA (1998)
8. Richards, D., Mourtada, J., Rosasco, L.: Asymptotics of ridge (less) regression under general source condition. Proc. Mach. Learn. Res. **130**, 3889–3897 (2021)
9. Wu, D., Xu, J.: On the optimal weighted ℓ_2 regularization in overparameterized linear regression. Adv. Neural. Inf. Process. Syst. **33**, 10112–10123 (2020)
10. Gantmacher, F.R., Krein, M.G.: Oscillation Matrices and Kernels and Small Vibrations of Mechanical Systems. AMS Chelsea Publishing, Rhode Island, USA (2002)
11. Price, H.: Monotone and oscillation matrices applied to finite difference approximations. Math. Comput. **22**, 489–516 (1968)
12. Golub, G.H., Van Loan, C.F.: Matrix Computations. John Hopkins University Press, Baltimore, USA (2013)
13. Li, Z.-C., Huang, H.-T.: Effective condition number for numerical partial differential equations. Numer. Linear Algebra Appl. **15**, 575–594 (2008)
14. Li, Z.-C., Huang, H.-T., Chen, J.-T., Wei, Y.: Effective condition number and its applications. Computing **89**, 87–112 (2010)
15. Lakshmi, M.V., Winkler, J.R.: Numerical properties of LASSO regression. Appl. Numer. Math. **208**, 297–309 (2025)
16. Winkler, J.R.: Error analysis and condition estimation of the pyramidal form of the Lucas-Kanade method in optical flow. Electronics **13**(5) (2024)
17. Winkler, J.R., Mitrouli, M.: Condition estimation for regression and feature estimation. J. Comp. Appl. Math. **373**, 112212 (2020)
18. Winkler, J.R., Mitrouli, M., Koukouvinos, C.: The application of regularisation to variable selection in statistical modelling. J. Comp. Appl. Math. **404**, 113884 (2022)
19. Hansen, P.C., Nagy, J.G., O'Leary, D.P.: Deblurring Images: Matrices, Spectra, and Filtering. SIAM, Philadelphia, USA (2006)
20. Choi, H.G., Thite, A.N., Thompson, D.J.: Comparison of methods for parameter selection in Tikhonov regularization with application to inverse force determination. J. Sound Vib. **304**, 894–917 (2007)
21. Zare, H., Hajarian, M.: Determination of regularization parameter via solving a multi-objective optimization problem. App. Num. Math. **156**, 542–554 (2020)
22. Bollt, E.: Regularized kernel machine learning for data driven forecasting of chaos. Ann. Rev. Chaos Theory, Bifurcat. Dyn. Syst. **9**, 1–26 (2020)
23. Hanke, M.: Limitations of the L-curve method in ill-posed problems. BIT **36**, 287–301 (1996)