



Interpretation of high-dimensional linear regression: Effects of nullspace and regularization demonstrated on battery data

Joachim Schaeffer^{a,b}, Eric Lenz^a, William C. Chueh^c, Martin Z. Bazant^b, Rolf Findeisen^a, Richard D. Braatz^{b,*}

^a Technical University of Darmstadt, Control and Cyber-Physical Systems Laboratory, Karolinenpl. 5, Darmstadt, 64289, Germany

^b Massachusetts Institute of Technology, 77 Massachusetts Avenue, Cambridge, 02139, MA, USA

^c Stanford University, 450 Jane Stanford Way, Stanford, 94305, CA, USA

ARTICLE INFO

Dataset link: <https://github.com/JoachimSchaeffer/HDRRegAnalytics>, <https://data.matr.io/1/>

Keywords:

Interpretable machine learning
Linear regression
High dimensions
Nullspace
Functional data
Lithium-ion batteries

ABSTRACT

High-dimensional linear regression is important in many scientific fields. This article considers discrete measured data of underlying smooth latent processes, as is often obtained from chemical or biological systems. Interpretation in high dimensions is challenging because the nullspace and its interplay with regularization shapes regression coefficients. The data's nullspace contains all coefficients that satisfy $\mathbf{X}\mathbf{w} = \mathbf{0}$, thus allowing very different coefficients to yield identical predictions. We developed an optimization formulation to compare regression coefficients and coefficients obtained by physical engineering knowledge to understand which part of the coefficient differences are close to the nullspace. This nullspace method is tested on a synthetic example and lithium-ion battery data. The case studies show that regularization and z-scoring are design choices that, if chosen corresponding to prior physical knowledge, lead to interpretable regression results. Otherwise, the combination of the nullspace and regularization hinders interpretability and can make it impossible to obtain regression coefficients close to the true coefficients when there is a true underlying linear model. Furthermore, we demonstrate that regression methods that do not produce coefficients orthogonal to the nullspace, such as fused lasso, can improve interpretability. In conclusion, the insights gained from the nullspace perspective help to make informed design choices for building regression models on high-dimensional data and reasoning about potential underlying linear models, which are important for system optimization and improving scientific understanding.

1. Introduction

Many important regression problems have the dimensionality of the data p much larger than the sample size n (Bühlmann and Van De Geer, 2011; Johnstone and Titterton, 2009; Hastie et al., 2009; Kobak et al., 2020). Consequently, $p \gg n$ and the matrix $\mathbf{X} \in \mathbb{R}^{n \times p}$ of predictors is “wide”. This case arises, for example, in most spectroscopies, lithium-ion batteries (Ralbovsky and Lednev, 2020; Schaeffer and Braatz, 2022), brain imaging, and computational biology (Boulesteix and Strimmer, 2007; Hastie et al., 2009; Bühlmann and Van De Geer, 2011). Classical literature on linear regression (Groß, 2003; Montgomery et al., 2012; Seber and Lee, 2003) focuses mainly on the case where $p < n$ and mostly assumes full column rank; however, many linear regression methods work well with wide predictor matrices. While Ordinary Least Squares (OLS) is not defined for wide data matrices because $\mathbf{X}^T\mathbf{X}$ is singular, the related minimum norm solution (e.g., Monticelli (1999)) can be used instead. Ridge Regression (RR)

and other shrinkage-based regression methods (e.g., Least Absolute Shrinkage and Selection Operator (lasso), Elastic Net (EN)) do not suffer from this problem due to the penalty term that is added to the main diagonal of $\mathbf{X}^T\mathbf{X}$. The fused lasso, a generalization of the lasso, adds an L_1 -norm penalty of adjacent regression coefficient differences to the objective function (Tibshirani et al., 2005). This additional penalty encourages piecewise constant regression coefficients, i.e., sparsity in regression coefficient differences. Thus it is required that the predictors can be ordered in some meaningful way. Latent variable methods such as Partial Least Squares (PLS) and Principal Component Regression (PCR) are popular choices for high-dimensional regression in the chemometrics community.

A key question is how to interpret high-dimensional linear regression results and the corresponding regression coefficients. In particular, how to reason about an underlying (linear) model for scientific insights and system optimization? Technically, regression coefficients

* Corresponding author.

E-mail address: braatz@mit.edu (R.D. Braatz).

for a linear model can be analyzed and compared to engineering or scientific expectations in terms of shape (e.g., peaks, plateaus, slopes), which is often done implicitly by engineers when looking at regression coefficients. However, as shown in this article, such an interpretation can lead to misleading conclusions.

This article develops a method, based on the nullspace of the predictor matrix $\mathcal{N}(\mathbf{X})$, for comparing coefficients obtained by different methods with each other for the case of high-dimensional data that was generated by a smooth latent process, also called *functional data* (Ramsey and Silverman, 2005).¹ We use the fact that $\mathcal{N}(\mathbf{X})$ consists of all solutions to $\mathbf{X}\mathbf{w} = 0$ and thus the predictions do not change when adding a vector of the nullspace to the regression coefficients $\mathbf{X}(\boldsymbol{\beta} + \mathbf{w}) = \mathbf{X}\boldsymbol{\beta}$. The nullspace and its interplay with regularization significantly influence the shape of the regression coefficients. Therefore, an understanding of the effect of the nullspace is needed for interpretation and scientific understanding. Our objective is to support such an understanding with this article.

The next section briefly introduces the key linear regression methods used in this article. Then the nullspace approach is derived. Subsequently, case studies are presented on fully synthetic data, lithium-ion battery data with two different synthetic linear responses, and the measured nonlinear cycle life response (Severson et al., 2019). The conclusion section summarizes the key learnings. All code and data used in this article are open-source and open-access, allowing the reproduction of results.

2. Motivation and linear regression

Linear, static models, assuming mean-centered data, have the general form

$$\mathbf{y} = \mathbf{X}\boldsymbol{\beta}^* + \boldsymbol{\epsilon} \quad (1)$$

where the input data matrix $\mathbf{X} \in \mathbb{R}^{n \times p}$, n is the number of observations, p is the number of predictors, and we assume that $p \gg n$. Our work is motivated by measurements of chemical or biochemical systems, i.e., discrete, noisy measurements of an assumed smooth underlying process. Consequently, we assume a latent model structure, i.e., \mathbf{X} (independently of \mathbf{y}) can be approximated in a lower dimensional space, and \mathbf{X} is not sparse. Most of the analysis in this article is technically not limited to this assumption. However, the nullspace perspective is motivated by a latent model structure and the high multicollinearity of columns that arises from functional data. The coefficients $\boldsymbol{\beta} \in \mathbb{R}^p$ contain the relation between \mathbf{X} and \mathbf{y} . The errors $\boldsymbol{\epsilon} \in \mathbb{R}^n$ are assumed to be homoscedastic, to have zero means, and to be uncorrelated.

Linear regression denotes statistical methods to determine $\hat{\boldsymbol{\beta}}$ from data \mathbf{X} and \mathbf{y} minimizing the error $\hat{\boldsymbol{\epsilon}}$ concerning a defined measure of the error,

$$\mathbf{y} = \hat{\mathbf{y}} + \hat{\boldsymbol{\epsilon}} = \mathbf{X}\hat{\boldsymbol{\beta}} + \hat{\boldsymbol{\epsilon}}. \quad (2)$$

The objective of regression methods is to find a $\hat{\boldsymbol{\beta}}$ that yields predictions that are reasonably close to the predictions of $\boldsymbol{\beta}^*$ when applied to independent data, i.e., were not available during training. When a true underlying linear model exists, interpretation and scientific insights would be supported by achieving a different goal, which is to reconstruct the true coefficients, i.e., $\hat{\boldsymbol{\beta}} = \boldsymbol{\beta}^*$, where $\boldsymbol{\beta}^*$ are the true coefficients of the model.

¹ Measured data from chemical and other systems often exhibit a certain degree of smoothness and can be considered to originate from discretized functions (similar to the assumptions made by Dette and Tang (2021)). The term *smoothness*, as used in this article, refers to data in which neighboring values are linked to each other to some extent, are not too different from one another, and there exists an underlying function that is differentiable once or multiple times.

As shown by Schaeffer and Braatz (2022), often columns in high-dimensional functional data are correlated, and regularized regression will find a solution that is optimal for its objective function; however, the resulting regression coefficients can be visually very different from $\boldsymbol{\beta}^*$ due to the interplay of the regularization and the nullspace, $\mathcal{N}(\mathbf{X})$. Furthermore, in practice, $\boldsymbol{\beta}^*$ is not known and the true underlying system might be nonlinear, requiring a thorough understanding of the interplay of regularization and the nullspace to draw reasonable conclusions about the underlying model.

Generally, the regression coefficients associated with $\hat{\boldsymbol{\beta}}^{\text{Model}}$ are random variables because \mathbf{X} and \mathbf{y} are realizations from a system that contains randomness (e.g., measurement errors, random system processes, etc.). One approach to model the regression coefficients probabilistically is Bayesian linear regression which places a prior on the regression coefficients and yields their posterior distribution, conditioned on data, which can then be analyzed (e.g., see Makalic and Schmidt (2016) for more information on Bayesian linear regression for high-dimensional data). While probabilistic modeling of regression coefficients is important, we focus on analyzing linear regression methods that do not model regression coefficients probabilistically because chemical engineers commonly use non-probabilistic models. We use $\boldsymbol{\beta}^{\text{Model}}$ to denote that it is a realization of the random variable by the ‘‘Model’’ and specific training data. From here, we drop the ‘‘hat’’ notation because it is clear from the model name in the superscript that the coefficients were obtained by regression from data.

Ordinary Least Squares (OLS) regression estimates with the closed-form solution $\boldsymbol{\beta}^{\text{OLS}} = (\mathbf{X}^T\mathbf{X})^{-1}\mathbf{X}^T\mathbf{y}$ for the case $p < n$ have low bias and are optimal under the assumption of the Gauss–Markov theorem. However, the regression coefficients $\boldsymbol{\beta}$ have a very large variance if the condition number of $\mathbf{X}^T\mathbf{X}$ is large, as is the case for many real-world data analytics problems, resulting in low prediction accuracy on unseen data. Ridge Regression (RR) addresses this problem by adding the squared L_2 -norm of the weights as a penalty to the least-squares objective (Hoerl and Kennard, 1970):

$$\min_{\boldsymbol{\beta}} \|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}\|_2^2 + \lambda\|\boldsymbol{\beta}\|_2^2, \quad (3)$$

yielding the closed-form solution $\boldsymbol{\beta}^{\text{RR}} = (\mathbf{X}^T\mathbf{X} + \lambda\mathbf{I})^{-1}\mathbf{X}^T\mathbf{y}$. The regularization penalty adds to the main diagonal of $\mathbf{X}^T\mathbf{X}$ and ensures that the resulting matrix is also invertible in the case $p > n$. RR improves the model’s generalization by introducing a bias that reduces variance in the estimated parameters (Zou and Hastie, 2005). For $p < n$ and $\lambda \rightarrow 0$, RR converges to OLS. In the more general case, without making assumptions about the dimensionality and rank of the real matrix \mathbf{X} , Singular Value Decomposition (SVD), $\mathbf{X} = \mathbf{U}\boldsymbol{\Sigma}\mathbf{V}^T$, can be used to show that

$$\boldsymbol{\beta}_0 = \lim_{\lambda \rightarrow 0} \boldsymbol{\beta}_\lambda = \mathbf{X}^\dagger \mathbf{y}. \quad (4)$$

The full derivation and further information can be found in the Supplementary Information (SI), Sec.S1 and in Kobak et al. (2020). For the case $p > n$, the Moore–Penrose-Inverse \mathbf{X}^\dagger can be written as

$$\boldsymbol{\beta}_0 = \mathbf{X}^T(\mathbf{X}\mathbf{X}^T)^{-1}\mathbf{y}. \quad (5)$$

This expression is known as the minimum norm solution (e.g., Monticelli (1999)),

$$\boldsymbol{\beta}_0 = \arg \min_{\boldsymbol{\beta}} \left\{ \|\boldsymbol{\beta}\|_2^2 \mid \|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}\|_2^2 = 0 \right\}. \quad (6)$$

For any $\tilde{\boldsymbol{\beta}}$ that fulfills $\mathbf{X}\tilde{\boldsymbol{\beta}} = \mathbf{y}$ (i.e., regression coefficients that fit the data \mathbf{X} and \mathbf{y} perfectly including the noise), (5) can be used to show that $\mathbf{X}(\tilde{\boldsymbol{\beta}} - \boldsymbol{\beta}_0) = \mathbf{0}$, and that

$$\begin{aligned} (\tilde{\boldsymbol{\beta}} - \boldsymbol{\beta}_0)^T \boldsymbol{\beta}_0 &= (\tilde{\boldsymbol{\beta}} - \boldsymbol{\beta}_0)^T \mathbf{X}^T (\mathbf{X}\mathbf{X}^T)^{-1} \mathbf{y} \\ &= (\mathbf{X}(\tilde{\boldsymbol{\beta}} - \boldsymbol{\beta}_0))^T (\mathbf{X}\mathbf{X}^T)^{-1} \mathbf{y} \\ &= \mathbf{0} \end{aligned} \quad (7)$$

and consequently $(\hat{\beta} - \beta_0) \perp \beta_0$ which is equivalent to $\mathcal{N}(\mathbf{X}) \perp \beta_0$ (Boyd and Lal, 2022). Thus there exists a set of regression coefficients \mathcal{S} that all fulfill $\mathbf{X}\hat{\beta} = \mathbf{y}$ with $\hat{\beta} \in \mathcal{S}$.

Most regularized regression methods solve an optimization of the form

$$\min_{\beta} \|\mathbf{y} - \mathbf{X}\beta\|_2^2 + F(\beta). \quad (8)$$

Regularized methods trade the perfect fit to the training data against the objective of keeping regression coefficients small. This trade-off is seen in the objective used to define the regularization methods. The orthogonality of the regression coefficients to $\mathcal{N}(\mathbf{X})$ does not hold for arbitrary regularization terms $F(\beta)$. Orthogonality holds for RR, because the regularization term in (3) is always smaller for coefficients orthogonal to $\mathcal{N}(\mathbf{X})$. The PCR coefficients are orthogonal to the $\mathcal{N}(\mathbf{X})$ because all eigenvectors of $\mathbf{X}^T\mathbf{X}$ that correspond to nonzero eigenvalues are orthogonal to each other and to the nullspace. Similarly, the PLS coefficients are also orthogonal to the nullspace by construction. The pathological case of \mathbf{y} being the nullvector must be excluded and is not relevant. Proofs for orthogonality between regression coefficients and nullspace for RR, PCR, and PLS are included in the SI, Sec.S2. However, regression coefficients obtained by the lasso and EN are not orthogonal to $\mathcal{N}(\mathbf{X})$ because of the L_1 -norm. For more information on regularized high-dimensional regression, see Schaeffer and Braatz (2022) and Tibshirani (1996). Depending on the function $F: \mathbb{R}^p \rightarrow \mathbb{R}$, regression coefficients obtain different shapes. Usually, methods such as RR, PCR, and PLS yield solutions that are not sparse, which can make interpretation difficult. An alternative method is the lasso, however, sparsity is often not a reasonable assumption for functional data. A generalization of the lasso is

$$\min_{\beta} \frac{1}{2} \|\mathbf{y} - \mathbf{X}\beta\|_2^2 + \lambda \|\mathbf{D}\beta\|_1 \quad (9)$$

where choosing \mathbf{D} as the identity matrix recovers the lasso. For

$$\mathbf{D}_1 = \begin{bmatrix} 1 & -1 & 0 & \dots & 0 \\ 0 & 1 & -1 & \ddots & \vdots \\ \vdots & \ddots & \ddots & \ddots & 0 \\ 0 & \dots & 0 & 1 & -1 \end{bmatrix}, \quad (10)$$

the resulting model is called the one-dimensional fused lasso which penalizes the L_1 -norm of the regression coefficients as well as their differences, but thus requires predictors that can be ordered (Tibshirani et al., 2005), a characteristic of functional data (Ramsay and Silverman, 2005). The choice of \mathbf{D} can incorporate expectations about the underlying model structure (Tibshirani and Taylor, 2011), and can thus yield models that should be interesting for many chemical engineering problems for its flexibility to incorporate assumptions, and its potential to yield easier-to-interpret regression coefficients. In the next section, we derive the nullspace method to compare the regression coefficients of different regularized models thoroughly.

3. Nullspace method

The addition of a vector $\mathbf{v} \in \mathcal{N}(\mathbf{X})$, i.e., a vector in the nullspace, to any β yields coefficients with unchanged predictions. The vectors in the nullspace affect only the regularization term in the objective function. We are interested in a method for understanding the effects of the nullspace when comparing different coefficients and how such a comparison can be used to reason about underlying relationships. Consider a regularized regression model called A,

$$\hat{\mathbf{y}}^A = \mathbf{X}\beta^A, \quad (11)$$

where β^A is associated with method A. For any vector \mathbf{v} in the nullspace, this equation implies that

$$\hat{\mathbf{y}}^A = \mathbf{X}(\beta^A + \mathbf{v}). \quad (12)$$

We want to compare the regression coefficients β^A with other coefficients β^B . The coefficients β^B can either be another estimator obtained by another regression method or instead be chosen for engineering or scientific reasons (e.g., constant regression coefficients). Thus, we propose finding coefficients $\mathbf{v}^* \in \mathcal{N}(\mathbf{X})$ that are closest to the difference between the coefficients under comparison $\beta_\Delta = \beta^A - \beta^B$. This approach can be formalized by

$$\min_{\mathbf{v}} \|\beta_\Delta + \mathbf{v}\|_2^2 \quad (13a)$$

$$\text{subject to } \mathbf{X}\mathbf{v} = \mathbf{0}, \quad (13b)$$

This optimization is a convex quadratic program with linear constraints. The solution is the projection of β_Δ onto the nullspace,

$$\mathbf{v}^* = (\mathbf{X}^T(\mathbf{X}\mathbf{X}^T)^{-1}\mathbf{X} - \mathbf{I})\beta_\Delta, \quad (14)$$

where $\mathbf{X}\mathbf{X}^T$ is assumed to be invertible. The derivation is included in the SI, Sec. S3. The expression can be simplified by inserting the singular value decomposition $\mathbf{X} = \mathbf{U}\Sigma\mathbf{V}^T$,

$$\mathbf{v}^* = (\mathbf{V}\Sigma^T(\Sigma\Sigma^T)^{-1}\Sigma\mathbf{V}^T - \mathbf{I})\beta_\Delta, \quad (15)$$

which can be used to improve the numerical efficiency. Simplifying (15) leads to

$$\mathbf{v}^* = \left(\mathbf{v} \begin{bmatrix} \mathbf{I}_n & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{bmatrix} \mathbf{V}^T - \mathbf{I} \right) \beta_\Delta. \quad (16)$$

The property that \mathbf{V} is an orthogonal matrix leads to

$$\mathbf{v}^* = -\mathbf{V} \begin{bmatrix} \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{I}_{p-n} \end{bmatrix} \mathbf{V}^T \beta_\Delta. \quad (17)$$

The projection onto the nullspace can be a hard requirement that might yield a vector \mathbf{v}^* that is dominated by noise and difficult to interpret, in particular, if $\mathbf{X}\mathbf{X}^T$ is ill-conditioned as is often the case for many real-world chemical engineering problems. Furthermore, regularization shapes regression coefficients by trading their variance against a bias towards zero to improve generalization. However, regularized regression coefficients usually differ from the true coefficients (if they exist), and their difference is not expected to lie exactly within the nullspace but might be close to it, motivating the relaxed optimization.

We propose to reformulate the optimization in (13) to allow deviations from the nullspace,

$$\min_{\mathbf{v}} \|\beta_\Delta + \mathbf{v}\|_2^2 + \gamma \|\mathbf{X}\mathbf{v}\|_2^2, \quad (18)$$

where γ is a nonnegative scalar. Setting the derivative of (18) with respect to \mathbf{v} to zero gives

$$\mathbf{v}_\gamma = -(\gamma\mathbf{X}^T\mathbf{X} + \mathbf{I})^{-1}\beta_\Delta. \quad (19)$$

For $\gamma = 0$, the nullspace is not considered and $\beta_\Delta = \mathbf{v}_0$. For $\gamma \rightarrow \infty$, the optimization converges to (17), as seen by

$$\begin{aligned} \lim_{\gamma \rightarrow \infty} -(\gamma\mathbf{X}^T\mathbf{X} + \mathbf{I})^{-1}\beta_\Delta &= \lim_{\gamma \rightarrow \infty} -\mathbf{V}(\gamma\Sigma^T\Sigma + \mathbf{I})^{-1}\mathbf{V}^T\beta_\Delta \\ &= -\mathbf{V} \begin{bmatrix} \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{I}_{p-n} \end{bmatrix} \mathbf{V}^T \beta_\Delta. \end{aligned} \quad (20)$$

Analyzing the nullspace, i.e., comparing the coefficients β^A and $\beta^A + \mathbf{v}_\gamma$ with β^B , allows us to identify which differences can be removed with a vector that is close to the nullspace and which differences would require significant deviations from the nullspace and are thus mainly responsible for the differences of the associated predictions. We propose to select γ , i.e., the penalization strength for deviations from the nullspace, based on a change in prediction accuracy to make it easier to interpret the result. That is, we define γ based on the Normalized-Root-Mean-Square Error (NRMSE) defined by

$$s = \max_i \{y_i\} - \min_i \{y_i\} \quad (21)$$

$$L(\hat{\mathbf{y}}, \mathbf{y}) = \frac{1}{s\sqrt{n}} \|\hat{\mathbf{y}} - \mathbf{y}\|_2, \quad (22)$$

leading to the heuristic:

$$\begin{aligned} & \max_{\gamma} \quad (23) \\ & \text{subject to } \left| L(\mathbf{X}(\beta^A + \mathbf{v}_{\gamma}), \mathbf{y}) - L(\mathbf{X}(\beta^A), \mathbf{y}) \right| \leq c \\ & \quad \mathbf{v}_{\gamma} = -(\gamma \mathbf{X}^T \mathbf{X} + \mathbf{I})^{-1} \beta_{\Delta}, \end{aligned}$$

where c defines the maximum loss function change introduced by the nullspace approach that is considered acceptable. The optimization (23) is not convex for most practical examples but is easy to solve because it only has one degree of freedom, γ .²

4. Case studies

This section demonstrates the nullspace method on several example cases to derive insights for interpretation of regression coefficients. The data \mathbf{X} and \mathbf{y} are generated synthetically for the first example. The second and third examples are on data from lithium-ion batteries (Severson et al., 2019), where we use constructed response variables by assuming different linear relationships to showcase the differences between regression coefficients and true coefficients. The last example uses the measured cycle life response where the true relationship between \mathbf{X} and \mathbf{y} is unknown.

4.1. Synthetic parabolic data

The parabolic example is inspired by measurements of some quantity over a continuous domain (e.g., time, concentration, voltage). To keep the data and relationships simple, the data are drawn from

$$\mathbf{x}_i = a_i \mathbf{d} \odot \mathbf{d}, \quad i \in \{1, 2, 3, \dots, 50\}, \quad (24)$$

$$\mathbf{d} = [1.0, 1.01, 1.02, \dots, 3.0], \quad (25)$$

$$\mathbf{X}^* = [\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n]^T,$$

where \mathbf{d} is the vector of discretizations on the underlying domain, with a constant spacing of 0.01 and a length of $p = 201$, and \odot is the element-wise product. The parameters $a_i \sim \mathcal{N}(\mu, \sigma^2)$ with $\mu = 0.3$ and $\sigma = 0.3$. Consequently, $\mathbf{X} \in \mathbb{R}^{50 \times 201}$. We define the response as

$$\begin{aligned} \mathbf{y}^* &= \mathbf{X}^* \beta^*, \\ \text{with } \beta^* &= \frac{1}{p} \mathbf{I}. \end{aligned} \quad (26)$$

The true coefficients are thus equal. Subsequently, we add white Gaussian noise to the data and response

$$\mathbf{x}_i = \mathbf{x}_i^* + \epsilon_{x_i} \quad (27)$$

$$\mathbf{y} = \mathbf{y}^* + \epsilon_y \quad (28)$$

yielding the matrices \mathbf{X} and \mathbf{y} for use in regression. The added noise $\epsilon_{x_i} \in \mathbb{R}^{201}$ is chosen such that the average SNR of each sample (\mathbf{x}_i) is 50 and the SNR of \mathbf{y} is 50 as well. Fig. 1a shows the mean-centered data, where each line corresponds to a matrix row. The 201 individual data points of each row are connected with a line, which is a reasonable visualization because of the underlying functional structure. We picked a PLS model with one component to learn the relationship between \mathbf{X} and \mathbf{y} . The PLS method is popular among chemical engineers, and its regularization parameter, the number of components, is discrete and simple to choose. Fig. 1b shows that the true coefficients and the PLS coefficients have very different shapes. However, their predictions and prediction accuracies are almost identical (cf. SI, Sec.S4). The noise leads to a prediction error even when the true coefficients are used (i.e., 0.105% NRMSE). We hand-selected $\gamma = 10$ for this case study. The resulting vector \mathbf{v}_{10} is very close to the nullspace, i.e., does

² For example, the optimization can be solved by plotting the left-hand side of the inequality with respect to γ .

not significantly change the prediction accuracy. The wrinkles indicate independent identically distributed noise, in line with the data generation. The adjusted coefficients $\beta_1^{\text{PLS}} + \mathbf{v}_{10}$ are very similar to the true coefficients β^* . While β_1^{PLS} is orthogonal to the nullspace, $\beta_1^{\text{PLS}} + \mathbf{v}_{10}$ is not orthogonal to the nullspace. Due to the simple underlying structure of the data and the model, the PLS coefficients yield a similar prediction accuracy. However, the PLS coefficients have a smaller L₂-norm, i.e., $\|\beta_1^{\text{PLS}}\|_2^2 < \|\beta^*\|_2^2$, due to the implicit regularization of PLS.

Assume that the coefficients are expected to be piecewise constant for physical reasons. We can then reformulate the regression as a generalized lasso problem with the matrix \mathbf{D} in (9) set to \mathbf{D}_1 . Fig. 2 shows the regression coefficients associated with RR and the fused lasso. The regularization parameter was chosen by CV and the one-standard-error rule (Hastie et al., 2009). Here we hand-selected $\gamma = 0.1$, as a different example. The resulting vector $\mathbf{v}_{0.1}$ is smoother than \mathbf{v}_{10} associated with Fig. 1, while still being very close to the nullspace and only changing the associated NRMSE prediction error by 0.001%. Fig. 2 looks remarkably similar to Fig. 1. The fused lasso coefficients are nearly identical to the true coefficients, and the ridge coefficients are similar to the PLS coefficients with one component but slightly noisier.

From the data alone, it is not possible to state whether \mathbf{y} was constructed from constant or parabolic coefficients. Furthermore, regression coefficients obtained from methods that are orthogonal to the nullspace can yield coefficients that appear to disagree with prior knowledge at first sight. As this example shows, methods that are not orthogonal to the nullspace, such as the fused lasso, can be advantageous for interpretation and conclusions if selected based on prior knowledge.

4.2. Lithium-ion battery data

As a real-world measurement data example, we consider a LFP battery data set, which contains cycling data for 124 batteries (Severson et al., 2019). Each battery has a fixed charging and discharging protocol. The charging protocols vary widely between the cells, whereas the discharge is constant and identical for all cells. The objective of the original paper was the prediction of the cycle life, i.e., the number of cycles until the battery's capacity drops below 80% of its nominal capacity. Features based on the difference between the discharge capacity of voltage curves for two cycles, subsequently called ΔQ_{a-b} , were shown to linearly correlate well with the logarithm of the cycle life. For this case study, we use the cycle pair $a = 100$ and $b = 10$, as suggested by Severson et al. (2019). Furthermore, we denote by a tilde ($\tilde{\Delta Q}_{100-10}$) that the columns are mean centered. The dimensionality $\tilde{\Delta Q}_{100-10} \in \mathbb{R}^{41 \times 1000}$ due to the high resolution of the discharge capacity over the voltage domain. More information about the data set and reasoning about the modeling objective can be found in Severson et al. (2019). Fig. 3a shows the LFP data set, partitioned into training, primary test, and secondary test data as suggested by Severson et al. (2019). Fig. 3b shows the mean subtracted training data. The data of the shortest-lived battery is clearly separated from the remainder of the data set. However, we keep this battery in the data set, as its influence on the training is benign. Fig. 3c shows the z-scored training data (i.e., standardized data, yielding mean-centered unit variance columns). The unit (Ah) is lost by z-scoring the data. Usually, z-scoring is recommended for data with features that have different units and thus might vary by orders of magnitude. However, for functional high-dimensional data, the unit of all columns is the same. Nevertheless, the measured values can vary by order of magnitude. Fig. 3c shows that the noise in the voltage region 3.2–3.5 V is amplified by rescaling because of a lower SNR in this voltage region. A more detailed analysis of the SNR can be found in the SI, Sec.S5. However, whether z-scoring is useful does not only depend on the data matrix \mathbf{X} but also on its underlying relationship with \mathbf{y} , which we explore next on synthetic responses \mathbf{y} before moving to the cycle life response.

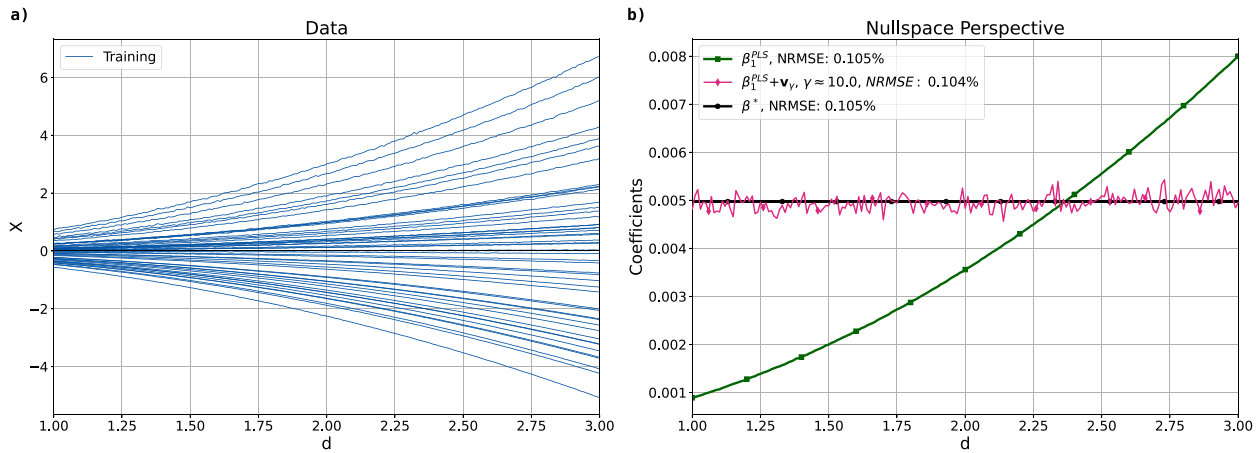


Fig. 1. (a) Mean-centered parabolic data, with white Gaussian noise corresponding to an SNR of 50 added to X and y prior to mean centering, (b) True coefficients in black and regression coefficients in green and nullspace-modified PLS coefficients in magenta.




Fig. 2. Ridge coefficients in green and fused lasso coefficients in black based on CV and the one-standard-error rule. Nullspace-modified ridge coefficients in magenta.

4.2.1. Synthetic response

Constant coefficients. The response for this example is the sample mean defined in (26), with $p = 1000$ to match the dimensionality of the LFP data set with added white Gaussian noise corresponding to an SNR of 50. Fig. 4a shows the nullspace perspective for the constant coefficient response with the data on the original scale, and Fig. 4b based on z-scored data. The number of PLS components is determined by CV and the one-standard-error rule. The PLS model associated with the z-scored data needs more components. The nullspace penalization parameter γ was chosen in both cases to allow for $c = 0.01\%$ NRMSE change, i.e., approximately half of the NRMSE difference between the two models. Fig. 4a shows the differences between the true coefficients and the PLS model's regression coefficients in the section from 2.0–3.1 V are relatively close to one another; however, some differences remain. The differences in the voltage region from 3.2 to 3.5 V only have a minor effect on the prediction results. Most of the difference between the regression coefficients in this area is associated with the nullspace, indicated by the large difference between the nullspace-modified PLS coefficients in magenta and the original PLS coefficients in green. Thus, the differences in the region 3.2 to 3.5 V do not change the prediction results on the training data significantly and arise due to the interplay of the regularization objective with the nullspace.

When the data are z-scored, most of the visible differences between the PLS coefficients in green and the true coefficients in black are contained in the enlarged nullspace (Fig. 4b). The modified coefficients

match the true coefficients very well. The prediction error difference between the PLS coefficients in green and the nullspace-modified coefficients in magenta is 0.01% NRMSE. The remaining differences between the true coefficients in black and the modified coefficient in magenta are barely visible but are responsible for another 0.01% NRMSE prediction error change, highlighting the effect of the nullspace. Comparing Fig. 4a and b shows that, in case the true coefficients are constant (i.e., all columns are equally important), z-scoring can help regression to yield coefficients that are more similar to the true coefficients.

Column mean coefficients. The true coefficient vector β^* for the next synthetic example is the column mean of the data X prior to column centering

$$\beta_j^* = \frac{1}{n} \sum_{i=1}^n x_{i,j}, \quad (29)$$

$$\beta^* = [\beta_1^*, \beta_2^*, \dots, \beta_p^*]^T.$$

The PLS model with 6 components associated with the z-scored data picks up a high amount of noise in the voltage regions from 3.3 to 3.5 V (Fig. 5b). In contrast, the PLS model with 3 components associated with the data on the original scale converges well to the true coefficients over the entire voltage region (Fig. 5a). The small differences are very closely associated with the nullspace. The nullspace penalization parameter γ was chosen in both cases to allow for $c = 0.05\%$ NRMSE change, i.e., approximately half of the NRMSE difference between the two models. Here, z-scoring amplifies and feeds noise into the model, manifesting as the spiky regression coefficients, with the most extreme spikes in the voltage regions from 3.2 to 3.5 V (Fig. 5b). Still, the PLS model associated with the z-scored data has approximately the same prediction accuracy as the PLS model associated with the original data.

Suppose there was some prior evidence or physical intuition that the true coefficients are constant or at least of similar magnitude. Then, z-scoring feeds the assumption that all the columns' importance is in the same order of magnitude to the model. However, if the coefficients are expected to vary by an order of magnitude (e.g., as is the case for the true coefficients being the column mean of the data), then not z-scoring the data accounts for the assumption that the scale of the columns is correlated with the assumed underlying true coefficients. The two examples show that the potential effects of z-scoring on the regression coefficients should be considered carefully for functional data. When data are z-scored, the model can become better at learning the underlying relationship, but noise may be amplified, depending on the noise structure.




Fig. 3. (a) LFP Discharge capacity difference between cycle 100 and 10, data split into training, primary and secondary test set, (b) Mean-centered training data, (c) z-scored training data.




Fig. 4. (a) True coefficients in black, PLS coefficients based on CV and the one-standard-error deviation rule in green, nullspace-modified coefficients in magenta, (b) Nullspace perspective similar to a, with PLS coefficients estimated with the one-standard-error rule corresponding to z-scored data.

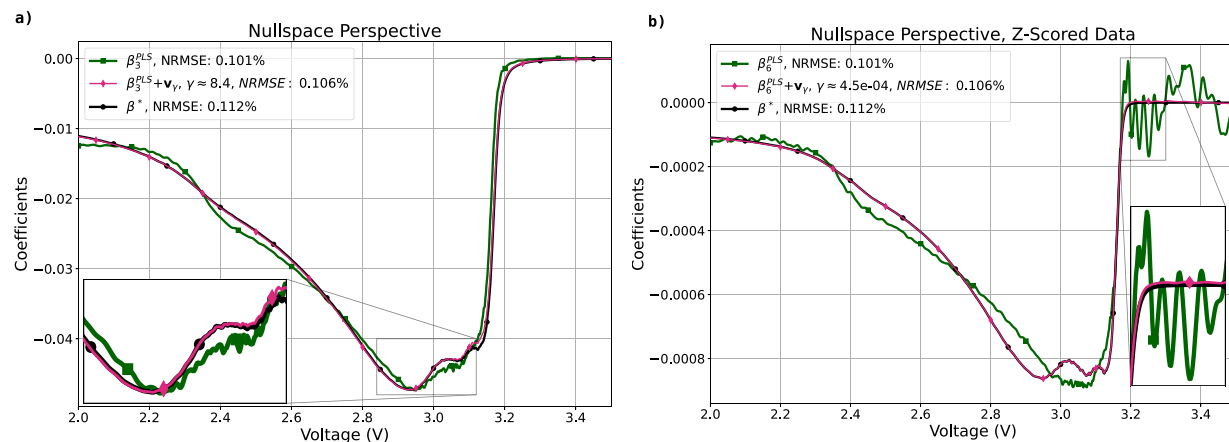


Fig. 5. (a) True coefficients in black, PLS coefficients based on CV and the one-standard-error rule in green, nullspace-modified coefficients in magenta, (b) Nullspace perspective similar to a, with PLS coefficients estimated with the one-standard-error rule corresponding to z-scored data.

4.2.2. Measured cycle life response

The measured response associated with the LFP battery data is the cycle life. We are interested in building models to predict the cycle life based on only discharge data from two cycles, which is a strategy to reduce the risk of data leakage (Geslin et al., 2023). We train the models by using the logarithm of the cycle life and use the same training, primary test, and secondary test set as suggested by Severson et al. (2019). We determine the regularization parameter based on the

minimum CV error and do not employ the one-standard-error rule.³ The PLS coefficients with five components have a similar shape as the fused lasso coefficients (cf. Fig. 6ab). However, the PLS coefficients have high-frequency perturbations, in particular, in the voltage range

³ The standard deviation of the CV error is large due to the long-living cells that heavily influence the prediction performance, which would lead to overly conservative regularization estimates.




Fig. 6. Cross-validated regression coefficients, original data (cf. Fig. 3b): (a) PLS coefficients, (b) Fused lasso coefficients.




Fig. 7. Cross-validated regression coefficients, z-scored data (cf. Fig. 3c): (a) PLS coefficients, (b) Fused lasso coefficients.

from 2.9–3.2 V, which is likely due to noise, making the PLS coefficients harder to interpret. The fused lasso coefficients (Fig. 6b) clearly indicate three regions of importance, enabling a physical interpretation. The range around 2.0–2.1 V is associated with the capacity change of the cell between cycles 10 and 100. Around 2.4 V, a different pattern can be seen in the data (Fig. 3), corresponding to the negative peak in the regression coefficients, which may correspond to LFP cathode degradation associated with iron anti-site defects, as the free energy of reaction (overpotential times charge) exceeds their formation energy ~ 0.55 eV (Malik et al., 2010). This interpretation is also consistent with experiments showing that chemical reduction of LFP by citric acid is able to heal iron anti-site defects with a similar free energy of reaction of 0.58 eV (Xu et al., 2020). The voltage range around 2.9–3.3 V contains most of the regression coefficient peaks. The two dominant plateaus of the Open-Circuit Voltage (OCV), which result from the single broad plateau of LFP superimposed with two more narrow plateaus of graphite, are located here, and most of the cell's capacity is discharged in this voltage range. These voltage plateaus correspond to phase transformations of the porous electrodes (Ferguson and Bazant, 2012), specifically between the low and high-density stable phases of LFP, as well as between stages 1, 2, and 3 of lithiated graphite (Ferguson and Bazant, 2014). The fused-lasso coefficients showcase three distinct negative and positive peaks, corresponding to changes in the rate-dependent tilt of the voltage plateaus, which may result from changes in particle-size-dependent nucleation barriers and population dynamics of reaction-controlled phase transformations (Ferguson and Bazant, 2014; Li et al., 2014). The peak width and height can be interpreted as a weighted sum of the average slopes of the data between the respective peaks. On low-rate data, the position and magnitude of peaks in the incremental capacity analysis correspond to different degradation modes (Dubarry et al., 2012). The peaks and peak shifts of the incremental capacity analysis blur out at higher C-rates, as expected from the suppression of phase separation by driven auto-inhibitory reactions (Bazant, 2017). In particular, the decreasing reaction rate with increasing lithium concentration in the LFP cathode, which has been predicted theoretically (Fraggedakis et al., 2021) and confirmed experimentally (Zhao et al., 2023), erases the voltage

plateaus at high rates and causes more homogeneous reactions that are likely favorable for battery lifetime (Ferguson and Bazant, 2012, 2014; Li et al., 2014). However, the obtained regression coefficients indicate that there is degradation information in this region even in ΔQ_{100-10} (i.e., the discharge capacity difference of cycle 100 and 10, both at 4C) that is important for capturing past degradation and forecasting future degradation. On the other hand, if the 4C current is well into the regime of suppressed phase separation, then we would expect a negative correlation between lifetime and internal resistance of the intercalation reaction, which in turn is correlated with larger ΔQ_{100-10} .

The coefficients regressed on the z-scored data have similar peaks and characteristics as the coefficients regressed on the original data (cf. Figs. 6ab, and 7ab). The z-scoring of columns introduces a linear transformation that significantly changes the regression coefficients in the range from 3.2 and 3.4 V.

The fused lasso based on the z-scored columns yields high prediction accuracy and interpretable coefficients (cf. Table 1 and Fig. 7b). In the higher voltage region, an additional peak appears around 3.35 V, which could not be learned from the original data because of the very small variance of the data prior to rescaling in combination with regularization. Moreover, the coefficients estimated on the z-scored data have the highest prediction performance on the training, primary, and secondary test sets (Table 1), showcasing that there is valuable information in the higher voltage region above 3.2 V. Furthermore, both models on the z-scored data outperform the variance model suggested by Severson et al. (2019). While being interpretable, fused lasso provides comparable prediction accuracy than PLS. The PLS model with nine components, suggested by Attia et al. (2021), slightly outperforms the fused lasso model when all cells are considered. However, the fused lasso yields the lowest Root-Mean-Square Error (RMSE) error for both test sets when only evaluated on the shorter-lived cells. The higher performance of the PLS model with nine components on the secondary test set is thus mainly associated with the longest-living cells that are more difficult to predict (cf. Severson et al. (2019) and Attia et al. (2021)). But, the coefficients associated with the PLS model are challenging to interpret because their sign changes frequently. What is more, the secondary test set was impacted by a longer calendar aging due to an extended storing

Table 1

Root-Mean-Square Error (RMSE) prediction accuracies associated with the coefficients in Figs. 6–7. Low Cycle Life (CL): $y_i \leq 1200$ cycles; high Cycle Life (CL): $y_i > 1200$ cycles. All models were trained on the entire training data. Scatter plots of prediction results are shown in the Supplementary Information (SI) Sec. S6.

Set	Original scale		Z-Scored		Feature Variance model (Severson et al., 2019)
	FL D1	PLS 5 Comp.	FL D1	PLS 9 Comp. (Attia et al., 2021)	
Training (41)	68	83	62	57	104
Test 1 (42)	115	116	105	102	138
Test 2 (40)	198	217	192	174	196
Training low CL (39)	62	82	53	50	103
Test 1 low CL (39)	96	101	76	80	96
Test 2 low CL (34)	135	202	115	132	119
Training high CL (2)	138	106	150	139	115
Test 1 high CL (3)	258	231	280	252	385
Test 2 high CL (6)	395	285	412	322	419

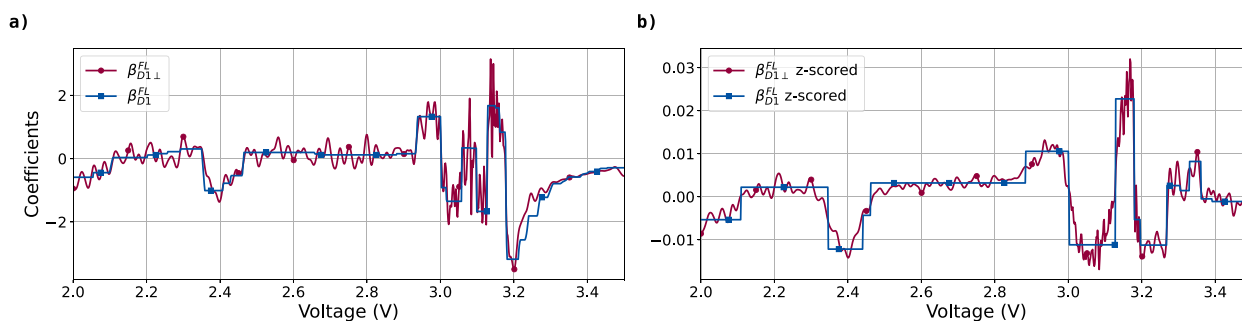


Fig. 8. Comparison of fused lasso coefficients (blue) and their component orthogonal to the nullspace (red): (a) Original data, (b) z-scored data.

period before the cycling started (cf. SI of Attia et al. (2020)), making it tough to understand the higher prediction accuracy of the PLS model on the secondary test set. In particular, the PLS coefficients show further peaks in the voltage region above 3.4 V influenced at least partially by noise because the SNR in this region is very low (cf. SI Sec. S5).

Similarly to the parabolic data set, we observe that the fused lasso coefficients are more interpretable than the PLS coefficients. Not requiring the coefficients orthogonal to the nullspace improves interpretability (Fig. 8ab). The component of the regression coefficients orthogonal to the nullspace (cf. red coefficients in Fig. 8ab) are less interpretable while making identical predictions on the training data. The improved interpretability of the fused lasso has also been observed in chemical engineering and other applications (e.g., Nowak et al. (2011), Martin-Barragan et al. (2014), Park and Kim (2016), Ross Kunz and She (2013) and Yu et al. (2015)).

5. Conclusion

The article proposes a nullspace perspective for gaining insights to help make informed design choices for building regression models on high-dimensional data and for reasoning about potential underlying linear models. We demonstrate the nullspace method on a fully synthetic data set and lithium-ion battery data with a designed linear response. Applying the nullspace insights for predicting the cycle life led to further insights into how degradation manifests itself for LFP batteries during discharge at 4C.

The nullspace allows different-looking regression coefficients to yield similar predictions (Fig. 1). While z-scoring for high-dimensional functional data can be beneficial, it should be an active design choice because it can increase noise by scaling up columns with low SNR (Fig. 5). Appropriate regularization can mitigate increased noise after z-scoring. Furthermore, regularization and z-scoring must be carefully considered and correspond to prior physical knowledge to obtain interpretable regression results (Figs. 6, 7). Otherwise, the combination of the nullspace and regularization can hinder interpretability and

potentially make it impossible to obtain regression coefficients close to the true coefficients.

Regression methods which yield coefficients orthogonal to the nullspace, such as RR, PCR, and PLS, can be challenging to interpret. Methods that yield regression coefficients not orthogonal to the nullspace, such as the fused lasso, can be advantageous for interpretability (Fig. 8).

The learnings from the nullspace analysis help to build and interpret linear regression models for high-dimensional functional data. The case studies show how to reason about underlying linear relationships between X and y , which is important for system optimization and to improve scientific understanding.

CRedit authorship contribution statement

Joachim Schaeffer: Conceptualization, Methodology, Software, Validation, Formal analysis, Investigation, Data curation, Writing – original draft, Writing – review & editing, Visualization, Funding acquisition. **Eric Lenz:** Formal analysis, Writing – review & editing. **William C. Chueh:** Writing – review & editing. **Martin Z. Bazant:** Writing – review & editing. **Rolf Findeisen:** Resources, Writing – review & editing, Funding acquisition. **Richard D. Braatz:** Conceptualization, Methodology, Resources, Writing – original draft, Writing – review & editing, Supervision, Project administration, Funding acquisition.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

The code for this article is available in the corresponding GitHub repository, *HDRRegAnalytics*, <https://github.com/JoachimSchaeffer/HDRRegAnalytics>

RegAnalytics. The repository contains the source code and notebooks to visualize the results. The repository contains a small subset of the LFP data set that was published with [Severson et al. \(2019\)](#) and is available at <https://data.matr.io/1/>.

Acknowledgments and funding

Initial ideas for this work were conceptualized during Joachim Schaeffer's time at ETH Zürich, for which we acknowledge financial support from the German Academic Exchange Service (DAAD) within the scholarship program for Master studies abroad. The main work was carried out by Joachim Schaeffer at the Technical University of Darmstadt. The work was refined and extended during Joachim Schaeffer's time at the Massachusetts Institute of Technology, for which we acknowledge financial support by a fellowship within the IFI program of the German Academic Exchange Service (DAAD), funded by the Federal Ministry of Education and Research (BMBF), Germany. Furthermore, financial support is acknowledged from the Toyota Research Institute, Germany through the D3BATT Center on Data-Driven-Design of Rechargeable Batteries.

Code and data availability

The code for this article is available in the corresponding GitHub repository, *HDRegAnalytics*, <https://github.com/JoachimSchaeffer/HDRegAnalytics>. The repository contains the source code and notebooks to visualize the results. The repository contains a small subset of the LFP data set that was published with [Severson et al. \(2019\)](#) and is available at <https://data.matr.io/1/>.

Appendix A. Supplementary data

Supplementary material related to this article can be found online at <https://doi.org/10.1016/j.compchemeng.2023.108471>.

References

- Attia, P.M., Grover, A., Jin, N., Severson, K.A., Markov, T.M., Liao, Y.-H., Chen, M.H., Cheong, B., Perkins, N., Yang, Z., et al., 2020. Closed-loop optimization of fast-charging protocols for batteries with machine learning. *Nature* 578 (7795), 397–402. <https://dx.doi.org/10.1038/s41586-020-1994-5>.
- Attia, P.M., Severson, K.A., Witmer, J.D., 2021. Statistical learning for accurate and interpretable battery lifetime prediction. *J. Electrochem. Soc.* 168 (9), 090547. <https://dx.doi.org/10.1149/1945-7111/ac2704>.
- Bazant, M.Z., 2017. Thermodynamic stability of driven open systems and control of phase separation by electro-autocatalysis. *Faraday Discuss.* 199, 423–463. <https://dx.doi.org/10.1039/C7FD00037E>.
- Boulesteix, A.-L., Strimmer, K., 2007. Partial least squares: A versatile tool for the analysis of high-dimensional genomic data. *Brief. Bioinform.* 8 (1), 32–44. <https://dx.doi.org/10.1093/bib/bbl016>.
- Boyd, S., Lal, S., 2022. EE263: Introduction to Linear Dynamical Systems. In: Lecture 16: Least-norm Solutions of Underdetermined Equations, Department of Electrical Engineering, Stanford University, Stanford, California, URL: <http://ee263.stanford.edu/lectures.html>, Accessed: 09/31/2023.
- Bühlmann, P., Van De Geer, S., 2011. *Statistics for High-Dimensional Data: Methods, Theory and Applications*. Springer, Berlin-Heidelberg.
- Detle, H., Tang, J., 2021. Statistical inference for function-on-function linear regression. <https://dx.doi.org/10.48550/ARXIV.2109.13603>, arXiv preprint, <https://arxiv.org/abs/2109.13603>.
- Dubarry, M., Truchot, C., Liaw, B.Y., 2012. Synthesize battery degradation modes via a diagnostic and prognostic model. *J. Power Sources* 219, 204–216. <https://dx.doi.org/10.1016/j.jpowsour.2012.07.016>.
- Ferguson, T.R., Bazant, M.Z., 2012. Nonequilibrium thermodynamics of porous electrodes. *J. Electrochem. Soc.* 159 (12), A1967–A1985. <https://dx.doi.org/10.1149/2.048212jes>.
- Ferguson, T.R., Bazant, M.Z., 2014. Phase transformation dynamics in porous battery electrodes. *Electrochim. Acta* 146, 89–97. <https://dx.doi.org/10.1016/j.electacta.2014.08.083>.
- Fraggedakis, D., McEldrew, M., Smith, R.B., Krishnan, Y., Zhang, Y., Bai, P., Chueh, W.C., Shao-Horn, Y., Bazant, M.Z., 2021. Theory of coupled ion-electron transfer kinetics. *Electrochim. Acta* 367, 137432. <https://dx.doi.org/10.1016/j.electacta.2020.137432>.

- Geslin, A., van Vlijmen, B., Cui, X., Bhargava, A., Asinger, P.A., Braatz, R.D., Chueh, W.C., 2023. Selecting the appropriate features in battery lifetime predictions. *Joule* 7, P1956–P1965. <http://dx.doi.org/10.1016/j.joule.2023.07.021>.
- Groß, J., 2003. *Linear Regression*. Springer Verlag, Berlin Heidelberg.
- Hastie, T., Tibshirani, R., Friedman, J.H., 2009. *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*, second ed. Springer, New York.
- Hoerl, A.E., Kennard, R.W., 1970. Ridge regression: Biased estimation for nonorthogonal problems. *Technometrics* 12 (1), 55–67. <https://dx.doi.org/10.2307/1271436>.
- Johnstone, I.M., Titterton, D.M., 2009. Statistical challenges of high-dimensional data. *Phil. Trans. R. Soc. A* 367 (1906), 4237–4253. <https://dx.doi.org/10.1098/rsta.2009.0159>.
- Kobak, D., Lomond, J., Sanchez, B., 2020. The optimal ridge penalty for real-world high-dimensional data can be zero or negative due to the implicit ridge regularization. *J. Mach. Learn. Res.* 21 (169), 1–16.
- Li, Y., El Gabaly, F., Ferguson, T.R., Smith, R.B., Bartelt, N.C., Sugar, J.D., Fenton, K.R., Cogswell, D.A., Kilcoyne, A.D., Tylliszczak, T., Bazant, M.Z., Chueh, W.C., 2014. Current-induced transition from particle-by-particle to concurrent intercalation in phase-separating battery electrodes. *Nature Mater.* 13 (12), 1149–1156. <https://dx.doi.org/10.1038/nmat4084>.
- Makalic, E., Schmidt, D.F., 2016. High-dimensional Bayesian regularised regression with the BayesReg package. arXiv:1611.06649, arXiv preprint, <https://arxiv.org/abs/1611.06649>.
- Malik, R., Burch, D., Bazant, M., Ceder, G., 2010. Particle size dependence of the ionic diffusivity. *Nano Lett.* 10 (10), 4123–4127. <https://dx.doi.org/10.1021/nl1023595>, PMID: 20795627.
- Martin-Barragan, B., Lillo, R., Romo, J., 2014. Interpretable support vector machines for functional data. *European J. Oper. Res.* 232 (1), 146–155. <https://dx.doi.org/10.1016/j.ejor.2012.08.017>.
- Montgomery, D.C., Peck, E.A., Vining, G.G., 2012. *Introduction to linear regression analysis*, fifth ed. John Wiley & Sons, Hoboken, N.J..
- Monticelli, A., 1999. Least-squares and minimum norm problems. In: *State Estimation in Electric Power Systems: A Generalized Approach*. Springer US, Boston, pp. 15–37, chapter 2.
- Nowak, G., Hastie, T., Pollack, J.R., Tibshirani, R., 2011. A fused lasso latent feature model for analyzing multi-sample aCGH data. *Biostatistics* 12 (4), 776–791. <https://dx.doi.org/10.1093/biostatistics/kxr012>.
- Park, C., Kim, S.B., 2016. Virtual metrology modeling of time-dependent spectroscopic signals by a fused lasso algorithm. *J. Process Control* 42, 51–58. <https://dx.doi.org/10.1016/j.jprocont.2016.04.002>.
- Ralbovsky, N.M., Lednev, I.K., 2020. Towards development of a novel universal medical diagnostic method: Raman spectroscopy and machine learning. *Chem. Soc. Rev.* 49 (20), 7428–7453. <https://dx.doi.org/10.1039/D0CS01019G>.
- Ramsay, J.O., Silverman, B.W., 2005. *Functional Data Analysis*, second ed. Springer, New York, p. 38.
- Ross Kunz, M., She, Y., 2013. Multivariate calibration maintenance and transfer through robust fused LASSO. *J. Chemometr.* 27 (9), 233–242. <https://dx.doi.org/10.1002/cem.2529>.
- Schaeffer, J., Braatz, R.D., 2022. Latent variable method demonstrator – Software for understanding multivariate data analytics algorithms. *Comput. Chem. Eng.* 167, 108014. <https://dx.doi.org/10.1016/j.compchemeng.2022.108014>.
- Seber, G.A.F., Lee, A.J., 2003. *Linear Regression Analysis*. John Wiley & Sons, Hoboken, NJ.
- Severson, K.A., Attia, P.M., Jin, N., Perkins, N., Jiang, B., Yang, Z., Chen, M.H., Aykol, M., Herring, P.K., Fraggedakis, D., Bazant, M.Z., Harris, S.J., Chueh, W.C., Braatz, R.D., 2019. Data-driven prediction of battery cycle life before capacity degradation. *Nat. Energy* 4 (5), 383–391.
- Tibshirani, R., 1996. Regression shrinkage and selection via the lasso. *J. R. Stat. Soc. Ser. B Stat. Methodol.* 58 (1), 267–288. <https://dx.doi.org/10.1111/j.2517-6161.1996.tb02080.x>.
- Tibshirani, R., Saunders, M., Rosset, S., Zhu, J., Knight, K., 2005. Sparsity and smoothness via the fused lasso. *J. R. Stat. Soc. Ser. B Stat. Methodol.* 67 (1), 91–108. <https://dx.doi.org/10.1111/j.1467-9868.2005.00490.x>.
- Tibshirani, R.J., Taylor, J., 2011. The solution path of the generalized lasso. *Ann. Statist.* 39 (3), 1335–1371. <https://dx.doi.org/10.1214/11-AOS878>.
- Xu, P., Dai, Q., Gao, H., Liu, H., Zhang, M., Li, M., Chen, Y., An, K., Meng, Y.S., Liu, P., Li, Y., Spangenberg, J.S., Gaines, L., Lu, J., Chen, Z., 2020. Efficient direct recycling of lithium-ion battery cathodes by targeted healing. *Joule* 4 (12), 2609–2626. <https://dx.doi.org/10.1016/j.joule.2020.10.008>.
- Yu, D., Lee, S.J., Lee, W.J., Kim, S.C., Lim, J., Kwon, S.W., 2015. Classification of spectral data using fused lasso logistic regression. *Chemometr. Intell. Lab. Syst.* 142, 70–77. <https://dx.doi.org/10.1016/j.chemolab.2015.01.006>.
- Zhao, H., Deng, H.D., Cohen, A.E., Lim, J., Li, Y., Fraggedakis, D., Jiang, B., Storey, B.D., Chueh, W.C., Braatz, R.D., Bazant, M.Z., 2023. Learning heterogeneous reaction kinetics from X-ray videos pixel by pixel. *Nature* 621 (7978), 289–294. <https://dx.doi.org/10.1038/s41586-023-06393-x>.
- Zou, H., Hastie, T., 2005. Regularization and variable selection via the elastic net. *J. R. Stat. Soc. Ser. B Stat. Methodol.* 67 (2), 301–320. <https://dx.doi.org/10.1111/j.1467-9868.2005.00503.x>.