

Analysis of overfitting in the regularized Cox model

Mansoor Sheikh^{1,2}  and Anthony C C Coolen^{1,2,3}

¹ Institute for Mathematical and Molecular Biomedicine, King's College London, Hodgkin Building, London SE1 1UL, United Kingdom

² Department of Mathematics, King's College London, The Strand, London WC2R 2LS, United Kingdom

³ Saddle Point Science Ltd, London, United Kingdom

E-mail: mansoor.sheikh@kcl.ac.uk and ton.coolen@kcl.ac.uk

Received 26 April 2019, revised 19 July 2019

Accepted for publication 31 July 2019

Published 26 August 2019



CrossMark

Abstract

The Cox proportional hazards model is ubiquitous in the analysis of time-to-event data. However, when the data dimension p is comparable to the sample size N , maximum likelihood estimates for its regression parameters are known to be biased or break down entirely due to overfitting. This prompted the introduction of the so-called regularized Cox model. In this paper we use the replica method from statistical physics to investigate the relationship between the true and inferred regression parameters in regularized multivariate Cox regression with L_2 regularization, in the regime where both p and N are large but with $\zeta = p/N \sim \mathcal{O}(1)$. We thereby generalize a recent study from maximum likelihood to maximum *a posteriori* inference. We also establish a relationship between the optimal regularization parameter and ζ , allowing for straightforward overfitting corrections in time-to-event analysis.

Keywords: Cox proportional hazards model, survival analysis, overfitting, replica method

(Some figures may appear in colour only in the online journal)

1. Introduction

Inference of parameters for generalized linear models using the maximum likelihood (ML) protocol becomes increasingly biased due to overfitting as the ratio $\zeta = p/N$ increases, where p is the number of covariates and N the number of training data. Overfitting occurs when model parameters seek to explain not only the 'signal' but also the 'noise' in training data, and is characterized by a difference in outcome prediction accuracy between training and validation samples. See e.g. [1, 3–5] for examples from logistic regression, gamma distributions

and Cox models [1]. Hence, standard statistical significance tests for regression coefficients, being usually based on asymptotic results derived for fixed p , become increasingly inaccurate [6]. Unfortunately, in post-genome medicine, having large ratios ζ is the rule rather than the exception. This prompted epidemiologists to formulate heuristic rules for avoiding overfitting, such as limits on the number of events per variable [7–11]. The Cox proportional hazards model [2], commonly used in epidemiological studies and clinical trials, predicts the continuous time-to-event random variable by combining an unspecified baseline hazard rate with a function of patient covariates. The canonical form for the covariate-dependent hazard rate of this model is $\lambda(t) = \lambda_0(t)e^{\beta \cdot \mathbf{z}}$. It was originally developed for use with life-tables where N is large (population-wide data) and the number of covariates p is small. Maximizing a likelihood function is a valid inference method in this regime. Approximate recipes for correcting ML estimates were developed in e.g. [12, 13]. Alternative methods of addressing the overfitting problem include feature selection and regularization. In feature selection one seeks to identify a subset of covariates that are informative of outcomes [14–16]. Its advantages include reduction in the required computational resources, and increased interpretability. In regularization one adds a penalty term to the objective function of ML inference (which can alternatively be derived from a prior in Bayesian inference) to suppress the number or magnitude of the model parameters [17, 18]. Application of regularization to survival analysis with high-dimensional covariates is studied widely, see e.g. [19, 20] and references therein.

A recent study [1] provided a new approach to overfitting in survival analysis. It showed how the replica method from statistical physics can be used to model ML inference analytically as the zero noise limit of a suitably defined stochastic minimization, starting from an information-theoretic measure of overfitting. The theory predicted the quantitative relation between ML-inferred and true parameters in the Cox model [1], and a phase transition at $\zeta = 1$.

Let us denote the set of model parameters as ϑ , and the data as \mathcal{D} . The observation that ML inference is equivalent to minimization of the Kullback–Leibler divergence between the empirical data distribution $\hat{P}_{\mathcal{D}}$ and the parametrized distribution P_{ϑ} assumed as a model of the data, suggests [1] using $E(\vartheta, \mathcal{D}) \equiv D(\hat{P}_{\mathcal{D}} \| P_{\vartheta}) - D(\hat{P}_{\mathcal{D}} \| P_{\vartheta^*})$ as a measure of overfitting⁴, in which ϑ^* are the true (but *a priori* unknown) parameter values. Perfect regression implies $E = 0$, underfitting implies $E > 0$, and overfitting implies $E < 0$. To gain more intuition for this measure, we generate synthetic data from a simple logistic regression model, find the ML estimators of its parameters, and calculate E . Here the parameters are $\{\beta_{\mu}\}_{\mu=0}^p$, the data are $\mathcal{D} = \{(t_1, \mathbf{z}_1), \dots, (t_N, \mathbf{z}_N)\}$, with $\mathbf{z}_i \in \mathbb{R}^{p+1}$ and $t_i \in \{0, 1\}$, and we use the shorthand $\beta \cdot \mathbf{z} = \sum_{\mu=0}^p \beta_{\mu} z_{\mu}$ (with the convention $z_0 = 1$). The outcome likelihood $P_{\beta}(t|\mathbf{z})$ and the measure E are given by

$$P_{\beta}(t|\mathbf{z}) = \left(\frac{1}{1 + e^{-\beta \cdot \mathbf{z}}} \right)^t \left(\frac{1}{1 + e^{\beta \cdot \mathbf{z}}} \right)^{1-t} \quad (1.1)$$

$$E(\beta^*, \mathcal{D}) = \frac{1}{N} \sum_{i=1}^N \left\{ t_i \log \left(\frac{1 + e^{-\beta^* \cdot \mathbf{z}_i}}{1 + e^{-\beta \cdot \mathbf{z}_i}} \right) + (1 - t_i) \log \left(\frac{1 + e^{\beta^* \cdot \mathbf{z}_i}}{1 + e^{\beta \cdot \mathbf{z}_i}} \right) \right\}. \quad (1.2)$$

Results are shown in figure 1. When $\zeta = 0.025$, E converges towards zero during the minimization, indicating perfect parameter recovery. As the number of samples in the data set is reduced, giving $\zeta = 0.25$ and $\zeta = 2.5$, E converges to increasingly negative values. Since

⁴ A similar idea for comparing estimators of probability distributions was used in [22], using the Lévy distance rather than the KL divergence. Other measures of overfitting can be found in e.g. [20, 23].

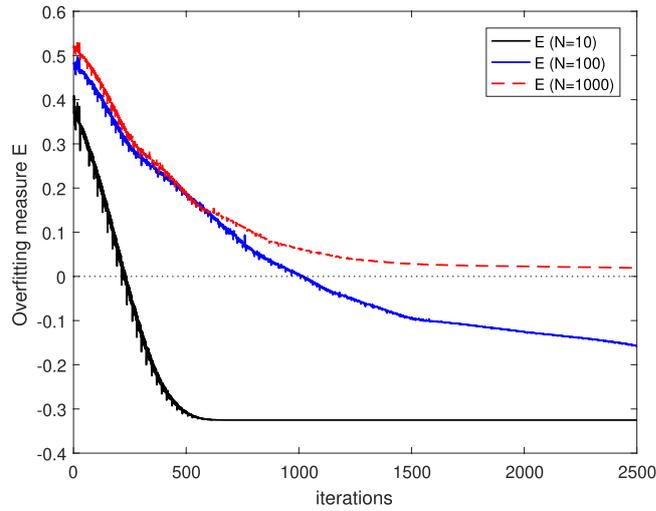


Figure 1. Synthetic data with dimension $p=25$ and $N = \{10, 100, 1000\}$ are generated using the logistic regression model. The ML estimate of model parameters is found numerically using the Nelder–Mead algorithm. The overfitting measure E is plotted after each iteration. The starting value model parameters $\boldsymbol{\beta}$ in the minimization search is the zero vector, giving a positive value of E (implying an underfitted model).

there is no model mismatch (the data were generated from a logistic model), the negative values of E indicate overfitting.

Switching from maximum likelihood to maximum *a posteriori* estimators implies adding a penalty term to the likelihood: $D(\hat{P}_{\mathcal{D}}\|P_{\boldsymbol{\vartheta}}) \rightarrow D(\hat{P}_{\mathcal{D}}\|P_{\boldsymbol{\vartheta}}) - \log p(\boldsymbol{\vartheta})$ where $p(\boldsymbol{\vartheta})$ represents a parameter prior, giving

$$\begin{aligned} E(\boldsymbol{\vartheta}^*, \mathcal{D}) &\equiv \min_{\boldsymbol{\vartheta}} \left\{ D(\hat{P}_{\mathcal{D}}\|P_{\boldsymbol{\vartheta}}) - \log p(\boldsymbol{\vartheta}) \right\} - \left\{ D(\hat{P}_{\mathcal{D}}\|P_{\boldsymbol{\vartheta}^*}) - \log p(\boldsymbol{\vartheta}^*) \right\} \\ &= \min_{\boldsymbol{\vartheta}} \left\{ \frac{1}{N} \sum_{i=1}^N \log \frac{p(t_i|\mathbf{z}_i, \boldsymbol{\vartheta}^*) p(\boldsymbol{\vartheta}^*)}{p(t_i|\mathbf{z}_i, \boldsymbol{\vartheta}) p(\boldsymbol{\vartheta})} \right\}. \end{aligned} \quad (1.3)$$

MAP regression is equivalent to minimizing the quantity (1.3). This minimization should in principle be over all $\boldsymbol{\vartheta}$, but may in practice be constrained to simplify the calculation (see e.g. [24–26]). For generalized linear models, commonly used priors are $p(\boldsymbol{\beta}) \propto \exp[-\eta \sum_{\mu=1}^p |\beta_{\mu}|]$ (giving L_1 regularization⁵, or ‘LASSO’ regression [18]) and $p(\boldsymbol{\beta}) \propto \exp[-\eta \sum_{\mu=1}^p \beta_{\mu}^2]$ (giving L_2 regularization, or ‘ridge’ regression).

In the present paper we generalize the replica analysis of [1] from ML to MAP inference, upon adding an L_2 regularization term to the log-likelihood function. This term suppresses overfitting effects, and removes the ML phase transition of the Cox model [1] at $\zeta = 1$; see e.g. figure 2. In the presence of an L_2 regularizer, correlations between covariates can no longer be transformed away, as was done in [1], leading to the appearance in the theory of the population covariance matrix \mathbf{A} of the covariates. Under mild restrictions on the eigenvalue

⁵This choice promotes sparsity in the regression coefficient vector $\boldsymbol{\beta}$, which would result in a horizontal line segment passing through the origin in figure 2. Since our theory aims to predict the slope of the data clouds in figure 2, we will not pursue L_1 regularizers in this paper.

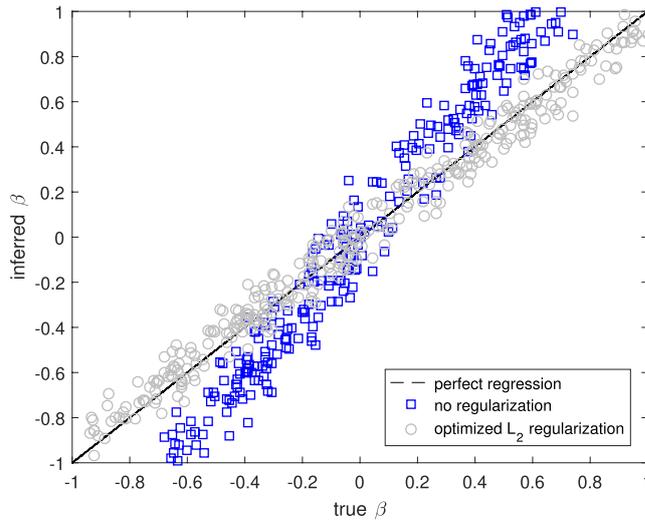


Figure 2. Comparison of true and inferred regression coefficients for the Cox proportional hazards model. A systematic bias is found for non-zero values of ζ , which can be corrected with regularization. Synthetic survival data were generated [21] using Gaussian covariates ($p = 500$, $N = 833$, $\zeta = 0.6$). The regression coefficients are inferred [17] using ML (no regularization) or maximum *a posteriori* probability regression (MAP, with regularization). Data points on the diagonal imply perfect inference.

spectrum of this matrix, we show how an accurate theory of overfitting for the regularized Cox proportional hazards model can be developed, in spite of such additional mathematical complications, including for the previously inaccessible regime $\zeta > 1$. We find, as in [1], that the replica symmetric version of the theory is sufficient to explain accurately the behaviour of interest. The resulting equations can also be used to predict the amount of regularization needed for unbiased regression, expressed in term of spectrum of \mathbf{A} and the ratio ζ .

2. Replica analysis of regularized Cox regression

2.1. Generalized replica formalism to include priors

Following [1], we interpret minimization of (1.3) as computing the ground state energy of a statistical mechanical system with degrees of freedom ϑ and Hamiltonian $H(\vartheta|\vartheta^*, \mathcal{D})$, at inverse temperature γ , where $\mathcal{D} = \{(t_1, \mathbf{z}_1), \dots, (t_N, \mathbf{z}_N)\}$ and

$$H(\vartheta|\vartheta^*, \mathcal{D}) = \log \prod_{i=1}^N \left[\frac{p(t_i|\mathbf{z}_i, \vartheta^*) p(\vartheta^*)}{p(t_i|\mathbf{z}_i, \vartheta) p(\vartheta)} \right]. \quad (2.1)$$

We define the associated free energy, which we average over the disorder (the microscopic realization of \mathcal{D}), and can compute the disorder-averaged ground state energy as the $\gamma \rightarrow \infty$ limit of the disorder-averaged energy density $E_\gamma(\vartheta^*)$, where

$$E_\gamma(\vartheta^*) = -\frac{1}{N} \frac{\partial}{\partial \gamma} \left\langle \log \int d\vartheta \prod_{i=1}^N \left[\frac{p(t_i|\mathbf{z}_i, \vartheta) p(\vartheta)}{p(t_i|\mathbf{z}_i, \vartheta^*) p(\vartheta^*)} \right]^\gamma \right\rangle_{\mathcal{D}}. \quad (2.2)$$

The replica identity $\langle \log Z \rangle = \lim_{n \rightarrow 0} n^{-1} \log \langle Z^n \rangle$ is subsequently used to simplify the average of the logarithm (see [1] for details and references), giving in the present case

$$\begin{aligned}
E_\gamma(\boldsymbol{\vartheta}^*) &= -\frac{\partial}{\partial \gamma} \lim_{n \rightarrow 0} \frac{1}{Nn} \log \left\langle \left\{ \int d\boldsymbol{\vartheta} \prod_{i=1}^N \left[\frac{p(t_i|\mathbf{z}_i, \boldsymbol{\vartheta}) p(\boldsymbol{\vartheta})}{p(t_i|\mathbf{z}_i, \boldsymbol{\vartheta}^*) p(\boldsymbol{\vartheta}^*)} \right]^\gamma \right\}^n \right\rangle_{\mathcal{D}} \\
&= -\frac{\partial}{\partial \gamma} \lim_{n \rightarrow 0} \frac{1}{Nn} \log \int \left\{ \prod_{\alpha=1}^n d\boldsymbol{\vartheta}^\alpha \left[\frac{p(\boldsymbol{\vartheta}^\alpha)}{p(\boldsymbol{\vartheta}^*)} \right]^\gamma \right\} \left\langle \prod_{i=1}^N \prod_{\alpha=1}^n \left[\frac{p(t_i|\mathbf{z}_i, \boldsymbol{\vartheta}^\alpha)}{p(t_i|\mathbf{z}_i, \boldsymbol{\vartheta}^*)} \right]^\gamma \right\rangle_{\mathcal{D}} \\
&= -\frac{\partial}{\partial \gamma} \lim_{n \rightarrow 0} \frac{1}{Nn} \log \int \left\{ \prod_{\alpha=1}^n d\boldsymbol{\vartheta}^\alpha \left[\frac{p(\boldsymbol{\vartheta}^\alpha)}{p(\boldsymbol{\vartheta}^*)} \right]^\gamma \right\} \\
&\quad \times \left\{ \int d\mathbf{z} dt p(\mathbf{z}) p(t|\mathbf{z}, \boldsymbol{\vartheta}^*) \prod_{\alpha=1}^n \left[\frac{p(t|\mathbf{z}, \boldsymbol{\vartheta}^\alpha)}{p(t|\mathbf{z}, \boldsymbol{\vartheta}^*)} \right]^\gamma \right\}^N. \tag{2.3}
\end{aligned}$$

Equation (2.3) is applicable to any parametric model $p(t|\mathbf{z}, \boldsymbol{\vartheta})$ and any prior $p(\boldsymbol{\vartheta})$. See also [27] for alternative results on the use of the replica method in statistical inference. We will now make a specific choice for $p(t|\mathbf{z}, \boldsymbol{\vartheta})$, and use (2.3) to develop a theory for regression and overfitting in regularized Cox models with Gaussian priors.

2.2. Application to the regularized Cox proportional hazards model

Cox's proportional hazards model originally described in [2] assumes a parametrization of the form⁶

$$p(t|\mathbf{z}, \boldsymbol{\vartheta}) = \lambda(t) e^{\boldsymbol{\beta} \cdot \mathbf{z} - \exp(\boldsymbol{\beta} \cdot \mathbf{z}) \int_0^t dt' \lambda(t')} \tag{2.4}$$

where the random variable $t \in \mathbb{R}^+$ represents the time-to-event/failure for the sample. Its parameters are the coefficients $\boldsymbol{\beta} \in \mathbb{R}^p$, and a base hazard rate $\lambda(t)$ (a nonnegative function defined for $0 \leq t < \infty$). For practical use, the focus is often on the so-called hazard ratios which compare the values of the factors $\exp(\boldsymbol{\beta}_\mu \cdot \mathbf{z}_\mu)$ for different covariate values. In this case, no assumptions are required for the unknown $\lambda(t)$ beyond $\lambda(t) \geq 0$. For our replica analysis, we make a variational approximation described in section 3.5. Substituting $\boldsymbol{\vartheta} = \{\boldsymbol{\beta}, \lambda\}$ translates (2.3) into

$$\begin{aligned}
E_\gamma(\boldsymbol{\beta}^*, \lambda^*) &= -\frac{\partial}{\partial \gamma} \lim_{n \rightarrow 0} \frac{1}{Nn} \log \int \{d\lambda^1 \dots d\lambda^n\} \int d\boldsymbol{\beta}^1 \dots d\boldsymbol{\beta}^n \left\{ \prod_{\alpha=1}^n \left[\frac{p(\boldsymbol{\beta}^\alpha)}{p(\boldsymbol{\beta}^*)} \right]^\gamma \right\} \\
&\quad \times \left\{ \int d\mathbf{z} dt p(\mathbf{z}) p(t|\mathbf{z}, \boldsymbol{\beta}^*, \lambda^*) \prod_{\alpha=1}^n \left[\frac{p(t|\mathbf{z}, \boldsymbol{\beta}^\alpha, \lambda^\alpha)}{p(t|\mathbf{z}, \boldsymbol{\beta}^*, \lambda^*)} \right]^\gamma \right\}^N. \tag{2.5}
\end{aligned}$$

Functional integrals are written as $\int \{d\lambda\}$, the true parameters responsible for the data are written as $\{\boldsymbol{\beta}^*, \lambda^*\}$, and we follow the standard convention for regularized Cox models of only including a prior for the association parameters (equivalently, assuming an improper, or 'flat', prior for the base hazard rate). Our L_2 prior is $p(\boldsymbol{\beta}) \propto \exp(-p\eta\boldsymbol{\beta}^2)$, and we will find in our analysis that this form indeed gives the appropriate scaling with p . To proceed with the analytical treatment, we assume that the covariate vectors \mathbf{z}_i are drawn independently from a population distribution with zero mean and covariance matrix \mathbf{A} . The introduction of

⁶The Cox proportional hazards model can be written in terms of the probability density function for t , the survival function and the cumulative hazard rates. Derivations of the relationships between these functions can be found in e.g. [32].

regularization means that the regression equations for correlated covariates can no longer be transformed to those corresponding to uncorrelated ones. This leads to a more complex theory than [1], and ultimately to conditions on the eigenvalue spectrum of \mathbf{A} .

Our analysis is carried out in the regime where both $N, p \rightarrow \infty$ but with fixed ratio $\zeta = p/N \sim \mathcal{O}(1)$. To retain non-zero event times, even for $p \rightarrow \infty$, we must rescale the regression coefficients according to $\boldsymbol{\beta} \rightarrow \boldsymbol{\beta}/\sqrt{p}$, resulting in $\boldsymbol{\beta} \cdot \mathbf{z} \sim \mathcal{O}(1)$. Without this rescaling we would have event time distributions with all weight concentrated on $t \rightarrow 0$ and $t \rightarrow \infty$. We also replace $\boldsymbol{\beta}^*$ by $\boldsymbol{\beta}^0$, to allow for more compact notation. Following [1] we next introduce

$$p(\mathbf{y}|\boldsymbol{\beta}^0, \dots, \boldsymbol{\beta}^n) = \int d\mathbf{z} p(\mathbf{z}) \prod_{\alpha=0}^n \delta\left[y^\alpha - \frac{\boldsymbol{\beta}^\alpha \cdot \mathbf{z}}{\sqrt{p}}\right] \quad (2.6)$$

where $\mathbf{y} = \{y^0, y^1, \dots, y^n\} \in \mathbb{R}^{n+1}$. The magnitude of $e^{\boldsymbol{\beta} \cdot \mathbf{z}}$ represents the relative risk of failure, compared to that of an ‘average’ individual (with $\mathbf{z} = \mathbf{0}$). Therefore \mathbf{y} can be considered a vector of risk scores. Our energy density then becomes

$$\begin{aligned} E_\gamma(\boldsymbol{\beta}^*, \lambda^*) &= -\frac{\partial}{\partial \gamma} \lim_{n \rightarrow 0} \frac{1}{Nn} \log \int \{d\lambda^1 \dots d\lambda^n\} \int d\boldsymbol{\beta}^1 \dots d\boldsymbol{\beta}^n \prod_{\alpha=1}^n \left[\frac{p(\boldsymbol{\beta}^\alpha)}{p(\boldsymbol{\beta}^0)} \right]^\gamma \\ &\quad \times \left\{ \int d\mathbf{y} p(\mathbf{y}|\boldsymbol{\beta}^0, \dots, \boldsymbol{\beta}^n) \int dt p(t|y^0, \lambda^0) \prod_{\alpha=1}^n \left[\frac{p(t|y^\alpha, \lambda^\alpha)}{p(t|y^0, \lambda^0)} \right]^\gamma \right\}^N \end{aligned} \quad (2.7)$$

in which now

$$p(t|y, \lambda) = \lambda(t) e^{y - \exp(y) \int_0^t dt' \lambda(t')}. \quad (2.8)$$

To proceed we assume that $p(\mathbf{y}|\boldsymbol{\beta}^0, \dots, \boldsymbol{\beta}^n)$ is Gaussian. This holds for any N and p as soon as $p(\mathbf{z})$ is Gaussian, and for non-Gaussian covariate statistics it will generally hold due to the Central Limit Theorem if the correlations among the covariates are weak, and N and p are large. Since we assumed $\int d\mathbf{z} p(\mathbf{z}) \mathbf{z} = \mathbf{0}$, the risk score distribution is now given by

$$p(\mathbf{y}|\boldsymbol{\beta}^0, \dots, \boldsymbol{\beta}^n) = \frac{e^{-\frac{1}{2} \mathbf{y} \cdot \mathbf{C}^{-1}[\{\boldsymbol{\beta}\}] \mathbf{y}}}{\sqrt{(2\pi)^{n+1} \det \mathbf{C}[\{\boldsymbol{\beta}\}]}}. \quad (2.9)$$

It is determined in full by the $(n+1) \times (n+1)$ covariance matrix $\mathbf{C}[\{\boldsymbol{\beta}\}]$, with entries

$$C_{\alpha\rho}[\{\boldsymbol{\beta}\}] = \int d\mathbf{z} p(\mathbf{z}) \left(\frac{\boldsymbol{\beta}^\alpha \cdot \mathbf{z}}{\sqrt{p}} \right) \left(\frac{\boldsymbol{\beta}^\rho \cdot \mathbf{z}}{\sqrt{p}} \right) = \frac{1}{p} \boldsymbol{\beta}^\alpha \cdot \mathbf{A} \boldsymbol{\beta}^\rho. \quad (2.10)$$

The entries of \mathbf{A} are given by $A_{\mu\nu} = \int d\mathbf{z} p(\mathbf{z}) z_\mu z_\nu$. The $\{C_{\alpha\rho}[\{\boldsymbol{\beta}\}]\}$ measure the similarity between the p -dimensional vectors formed by the regression parameters in different replicas. For each replica pair (α, ρ) we use the integral representation of the Dirac delta function, and rescale the conjugate integration parameter by p ,

$$1 = \int dC_{\alpha\rho} \delta\left[C_{\alpha\rho} - \frac{1}{p} \boldsymbol{\beta}^\alpha \cdot \mathbf{A} \boldsymbol{\beta}^\rho\right] = \int \frac{dC_{\alpha\rho} d\hat{C}_{\alpha\rho}}{2\pi/p} e^{ip\hat{C}_{\alpha\rho}(C_{\alpha\rho} - \frac{1}{p} \boldsymbol{\beta}^\alpha \cdot \mathbf{A} \boldsymbol{\beta}^\rho)} \quad (2.11)$$

in order to simplify expression (2.7) to

$$\begin{aligned}
E_\gamma(\boldsymbol{\beta}^*, \lambda^*) &= -\frac{\partial}{\partial \gamma} \lim_{n \rightarrow 0} \frac{1}{Nn} \log \int \{d\lambda^1 \dots d\lambda^n\} \int d\mathbf{C} d\hat{\mathbf{C}} \frac{e^{ip \sum_{\alpha,\rho=0}^n \hat{C}_{\alpha\rho} C_{\alpha\rho}}}{(2\pi/p)^{(n+1)^2}} \\
&\times \left[\int \frac{d\mathbf{y} e^{-\frac{1}{2}\mathbf{y}^T \mathbf{C}^{-1} \mathbf{y}}}{\sqrt{(2\pi)^{n+1} \det \mathbf{C}}} \int dt p(t|y^0, \lambda^0) \prod_{\alpha=1}^n \left[\frac{p(t|y^\alpha, \lambda^\alpha)}{p(t|y^0, \lambda^0)} \right]^\gamma \right]^N \\
&\times \int d\boldsymbol{\beta}^1 \dots d\boldsymbol{\beta}^n e^{-\eta\gamma \sum_{\alpha=1}^n [(\boldsymbol{\beta}^\alpha)^2 - (\boldsymbol{\beta}^0)^2] - i \sum_{\alpha,\rho=0}^n \hat{C}_{\alpha\rho} \boldsymbol{\beta}^\alpha \cdot \mathbf{A} \boldsymbol{\beta}^\rho}.
\end{aligned} \tag{2.12}$$

The quadratic nature of the exponent in the $\boldsymbol{\beta}$ integral, a consequence of having chosen L_2 regularization, allows for a closed form solution. Changing the penalty term to L_1 or L_q with $q > 2$ would significantly complicate the integrals.

2.3. Conversion into a saddle point problem

With a modest amount of foresight we transform $\hat{\mathbf{C}} = -\frac{1}{2}i\mathbf{D}$, and introduce the short-hand $\tilde{\boldsymbol{\beta}} \equiv \mathbf{A}^{\frac{1}{2}}\boldsymbol{\beta}$. To evaluate the Gaussian $\boldsymbol{\beta}$ integral in (2.12) we define the $np \times np$ matrix Ξ and the np -dimensional vector $\boldsymbol{\xi}$, with entries

$$\Xi_{\alpha\mu;\beta\nu} = 2\eta\gamma\delta_{\alpha\beta}(\mathbf{A}^{-1})_{\mu\nu} + \delta_{\mu\nu}D_{\alpha\beta}, \quad \xi_\mu^\alpha = -D_{0\alpha}\tilde{\beta}_\mu^0. \tag{2.13}$$

With these definitions we may write the Gaussian integral in (2.12) as

$$\begin{aligned}
&\int \left(\prod_{\alpha=1}^n d\tilde{\boldsymbol{\beta}}^\alpha e^{-\eta\gamma\tilde{\boldsymbol{\beta}}^\alpha \cdot \mathbf{A}^{-1}\tilde{\boldsymbol{\beta}}^\alpha} \right) e^{-\frac{1}{2} \sum_{\alpha,\rho=1}^n D_{\alpha\rho}\tilde{\boldsymbol{\beta}}^\alpha \cdot \tilde{\boldsymbol{\beta}}^\rho - \sum_{\rho=1}^n D_{0\rho}\tilde{\boldsymbol{\beta}}^0 \cdot \tilde{\boldsymbol{\beta}}^\rho} \\
&= e^{\frac{1}{2}\boldsymbol{\xi} \cdot \Xi^{-1} \boldsymbol{\xi}} \int d\tilde{\boldsymbol{\beta}} e^{-\frac{1}{2}(\tilde{\boldsymbol{\beta}} - \Xi^{-1}\boldsymbol{\xi}) \cdot \Xi (\tilde{\boldsymbol{\beta}} - \Xi^{-1}\boldsymbol{\xi})} = \frac{(2\pi)^{\frac{np}{2}}}{\sqrt{\det \Xi}} e^{\frac{1}{2}\boldsymbol{\xi} \cdot \Xi^{-1} \boldsymbol{\xi}}.
\end{aligned} \tag{2.14}$$

Let $\{a_\mu\}$ and $\{b_\alpha\}$ denote the eigenvalues of \mathbf{A} and \mathbf{D} , respectively. The two terms \mathbf{P} and \mathbf{Q} of the matrix Ξ , with components $P_{\alpha\mu;\beta\nu} = 2\eta\gamma\delta_{\alpha\beta}(\mathbf{A}^{-1})_{\mu\nu}$ and $Q_{\alpha\mu;\beta\nu} = \delta_{\mu\nu}D_{\alpha\beta}$, clearly commute. The complete set of eigenvectors of Ξ can therefore be written as $\{\hat{\mathbf{u}}^{\mu\alpha}\}$, with components $\hat{u}_{\nu\rho}^{\mu\alpha} = u_\rho^\alpha v_\nu^\mu$, and where $\sum_{\rho \leq n} D_{\lambda\rho} u_\rho^\alpha = b_\alpha u_\rho^\alpha$ and $\sum_{\nu \leq p} A_{\lambda\nu} v_\nu^\mu = a_\mu v_\lambda^\mu$, and where both are normalised according to $\sum_{\rho \leq n} (u_\rho^\alpha)^2 = \sum_{\nu \leq p} (v_\nu^\mu)^2 = 1$. The eigenvalues of Ξ are then $\xi_{\mu\alpha} = 2\eta\gamma/a_\mu + b_\alpha$, and

$$\det \Xi = \prod_{\mu=1}^p \prod_{\alpha=1}^n \left(\frac{2\eta\gamma}{a_\mu} + b_\alpha \right), \quad (\Xi^{-1})_{\alpha\mu;\alpha'\mu'} = \sum_{\beta=1}^n \sum_{\nu=1}^p \frac{u_\alpha^\beta v_\mu^\nu u_{\alpha'}^\beta v_{\mu'}^\nu}{2\eta\gamma/a_\nu + b_\beta}. \tag{2.15}$$

Hence the integral (2.14) can be written as

$$\frac{(2\pi)^{\frac{np}{2}}}{\sqrt{\det \Xi}} e^{\frac{1}{2}\boldsymbol{\xi} \cdot \Xi^{-1} \boldsymbol{\xi}} = e^{\frac{1}{2}np \log(2\pi) - \frac{1}{2}np \langle \log(2\eta\gamma/a+b) \rangle + \frac{1}{2}np \langle (\boldsymbol{\xi} \cdot \hat{\mathbf{u}})^2 (2\eta\gamma/a+b)^{-1} \rangle} \tag{2.16}$$

where the averages in the exponents are over the eigenvalues and orthonormal eigenvectors of Ξ , i.e. $\langle f(a, b, \hat{\mathbf{u}}) \rangle = (np)^{-1} \sum_{\mu=1}^p \sum_{\alpha=1}^n f(a_\mu, b_\alpha, \hat{\mathbf{u}}^{\mu\alpha})$. Since $p = \zeta N$ with $\zeta > 0$, the integrals over \mathbf{C} , $\hat{\mathbf{C}}$ and the base hazard rates in (2.12) can for $N \rightarrow \infty$ be evaluated by steepest descent, provided the limits $n \rightarrow 0$ and $N \rightarrow \infty$ commute. Expression (2.16) then enables us to write the result as

$$\lim_{N \rightarrow \infty} E_\gamma(\beta^*, \lambda^*) = \frac{\partial}{\partial \gamma} \lim_{n \rightarrow 0} \frac{1}{n} \text{extr} \Psi(\mathbf{C}, \mathbf{D}, \lambda^1 \dots \lambda^n) \quad (2.17)$$

in which

$$\begin{aligned} \Psi(\mathbf{C}, \mathbf{D}, \lambda^1 \dots \lambda^n) &= -\frac{1}{2} \zeta \left[\sum_{\alpha, \rho=0}^n D_{\alpha\rho} C_{\alpha\rho} - \frac{1}{p} D_{00} (\tilde{\beta}^0)^2 \right] + \frac{1}{2} (n+1 - n\zeta) \log(2\pi) \\ &+ \frac{1}{2} \log \det \mathbf{C} - n\eta\zeta\gamma S^2 + \frac{1}{2} n\zeta \left\langle \log \left(\frac{2\eta\gamma}{a} + b \right) \right\rangle - \frac{1}{2} n\zeta \left\langle \frac{(\boldsymbol{\xi} \cdot \hat{\mathbf{u}})^2}{2\eta\gamma/a + b} \right\rangle \\ &- \log \int d\mathbf{y} e^{-\frac{1}{2} \mathbf{y} \cdot \mathbf{C}^{-1} \mathbf{y}} \int dt p(t|y^0, \lambda^0) \prod_{\alpha=1}^n \left[\frac{p(t|y^\alpha, \lambda^\alpha)}{p(t|y^0, \lambda^0)} \right]^\gamma \end{aligned} \quad (2.18)$$

where we have defined $S^2 = \lim_{p \rightarrow \infty} p^{-1} (\beta^0)^2$. Differentiating Ψ with respect to D_{00} removes D_{00} from the problem, and gives $C_{00} = p^{-1} \beta^0 \cdot \mathbf{A} \beta^0 \equiv \tilde{S}^2$.

3. Replica symmetric theory

3.1. Replica symmetric saddle points

To proceed, we make the replica symmetric ansatz, which implies assuming ergodicity of the stochastic regression process, and translates into invariance of all order parameters under all permutations of the replicas $\{1, \dots, n\}$. Now, for all $1 \leq \alpha, \rho \leq n$:

$$\lambda^\alpha(t) = \lambda(t), \quad \begin{array}{l} C_{0\alpha} = c_0 \\ D_{0\alpha} = d_0 \end{array}, \quad \begin{array}{l} C_{\alpha\rho} = C\delta_{\alpha\rho} + c(1 - \delta_{\alpha\rho}) \\ D_{\alpha\rho} = D\delta_{\alpha\rho} + d(1 - \delta_{\alpha\rho}) \end{array}. \quad (3.1)$$

Both \mathbf{C} and \mathbf{D} are positive definite, so $C > c$ and $D > d$. We may now write

$$\mathbf{C} = \begin{pmatrix} C_{00} & c_0 & \dots & \dots & c_0 \\ c_0 & C & c & \dots & c \\ \vdots & c & C & \dots & c \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ c_0 & c & c & \dots & C \end{pmatrix}, \quad \mathbf{C}^{-1} = \begin{pmatrix} B_{00} & b_0 & \dots & \dots & b_0 \\ b_0 & B & b & \dots & b \\ \vdots & b & B & \dots & b \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ b_0 & b & b & \dots & B \end{pmatrix}.$$

The eigenvalues and eigenvectors of \mathbf{C} , \mathbf{C}^{-1} and \mathbf{D} are found in [1]. \mathbf{C} has two nondegenerate eigenvalues λ_\pm with $\lambda_+ \lambda_- = [C + (n-1)c]C_{00} - nc_0^2$, and a further $n-1$ fold degenerate eigenvalue $\lambda_0 = C - c$. Hence

$$\begin{aligned} \log \det \mathbf{C} &= \log \left([C + (n-1)c]C_{00} - nc_0^2 \right) + (n-1) \log(C-c) \\ &= \log C_{00} + n \log(C-c) + \frac{n(c-c_0^2/C_{00})}{C-c} + \mathcal{O}(n^2). \end{aligned} \quad (3.2)$$

The entries of \mathbf{C}^{-1} are found to be

$$B_{00} = \frac{C + (n-1)c}{C_{00}[C + (n-1)c] - nc_0^2}, \quad b_0 = -\frac{c_0}{C_{00}[C + (n-1)c] - nc_0^2} \quad (3.3)$$

$$B = b + \frac{1}{C - c}, \quad b = \frac{c_0^2 - cC_{00}}{(C_{00}[C + (n - 1)c] - nc_0^2)(C - c)}. \quad (3.4)$$

Hence

$$\mathbf{y} \cdot \mathbf{C}^{-1} \mathbf{y} = B_{00}(y^0)^2 + (B - b) \sum_{\alpha=1}^n (y^\alpha)^2 + b \left(\sum_{\alpha=1}^n y^\alpha \right)^2 + 2b_0 y^0 \sum_{\alpha=1}^n y^\alpha. \quad (3.5)$$

Next we turn to terms in (2.18) that involve the spectrum of \mathbf{D} . This matrix has one eigenvalue $D + (n - 1)d$ with eigenvector $\mathbf{v} = (1, \dots, 1)$, and the $n - 1$ fold degenerate eigenvalue $D - d$ with eigenspace $\sum_{i=1}^n \mathbf{v}_i = 0$. Hence

$$\begin{aligned} \left\langle \log \left(\frac{2\eta\gamma}{a} + b \right) \right\rangle &= \frac{1}{np} \sum_{\mu=1}^p \left[(n-1) \log \left(\frac{2\eta\gamma}{a_\mu} + D - d \right) + \log \left(\frac{2\eta\gamma}{a_\mu} + D - d + nd \right) \right] \\ &= \left\langle \log \left(\frac{2\eta\gamma}{a} + D - d \right) \right\rangle + \left\langle \frac{da}{2\eta\gamma + (D - d)a} \right\rangle + \mathcal{O}(n). \end{aligned} \quad (3.6)$$

Similarly, using the RS form of $\xi_\mu^\alpha = -d_0(\mathbf{A}^{\frac{1}{2}}\boldsymbol{\beta}^0)_\mu$, we may write

$$\begin{aligned} \left\langle \frac{(\boldsymbol{\xi} \cdot \hat{\mathbf{u}})^2}{2\eta\gamma/a + b} \right\rangle &= \frac{1}{np} \sum_{\mu=1}^p \left(\sum_{\nu=1}^p \sum_{\rho=1}^n (\mathbf{A}^{\frac{1}{2}}\boldsymbol{\beta}^0)_\nu v_\nu^\mu \frac{1}{\sqrt{n}} \right)^2 \frac{d_0^2}{2\eta\gamma/a_\mu + D + (n-1)d} \\ &= \frac{1}{p} \sum_{\mu=1}^p (\boldsymbol{\beta}^0 \cdot \mathbf{v}^\mu)^2 \frac{d_0^2 a_\mu}{2\eta\gamma/a_\mu + D - d} + \mathcal{O}(n) \\ &= d_0^2 \left\langle \frac{a^2 (\boldsymbol{\beta}^0 \cdot \mathbf{v})^2}{2\eta\gamma + (D - d)a} \right\rangle + \mathcal{O}(n). \end{aligned} \quad (3.7)$$

The averages in (3.6) and (3.7) are now over the joint distribution of eigenvalues and eigenvectors of \mathbf{A} only. Inserting the above RS expressions into (2.18), and using $C_{00} = \bar{S}^2$, then gives us, with the short-hand $D\mathbf{z} = (2\pi)^{-\frac{1}{2}} \mathbf{e}^{-\frac{1}{2}\mathbf{z}^2} d\mathbf{z}$,

$$\begin{aligned} \frac{1}{n} \Psi_{\text{RS}}(\dots) &= -\frac{1}{2} \zeta (2d_0 c_0 + DC - dc) + \frac{1}{2} (1 - \zeta) \log(2\pi) - \eta \zeta \gamma S^2 + \mathcal{O}(n) \\ &+ \frac{1}{2} \left[\log(C - c) + \frac{c - c_0^2/C_{00}}{C - c} \right] - \frac{1}{2} \zeta d_0^2 \left\langle \frac{a^2 (\boldsymbol{\beta}^0 \cdot \mathbf{v})^2}{2\eta\gamma + (D - d)a} \right\rangle \\ &+ \frac{1}{2} \zeta \left\langle \log \left(\frac{2\eta\gamma}{a} + D - d \right) \right\rangle + \frac{1}{2} \zeta \left\langle \frac{da}{2\eta\gamma + (D - d)a} \right\rangle + \frac{1}{n} \log \bar{S} \\ &- \frac{1}{n} \log \int D\mathbf{z} \int \frac{dy_0}{\sqrt{2\pi}} \mathbf{e}^{-\frac{1}{2} B_{00} y_0^2} \int dt p(t|y_0, \lambda^0) \\ &\times \left[\int d\mathbf{y} \mathbf{e}^{-\frac{1}{2}(B-b)y^2 + y(i\mathbf{z}\sqrt{b} - b_0 y_0)} \frac{p^\gamma(t|\mathbf{y}, \lambda)}{p^\gamma(t|y_0, \lambda^0)} \right]^n \end{aligned}$$

$$\begin{aligned}
 &= -\frac{1}{2}\zeta(2d_0c_0 + DC - dc) + \frac{1}{2}(1-\zeta)\log(2\pi) - \eta\zeta\gamma S^2 + \mathcal{O}(n) \\
 &+ \frac{1}{2}\left[\log(C-c) + \frac{c-c_0^2/\tilde{S}^2}{C-c}\right] - \frac{1}{2}\zeta d_0^2 \left\langle \frac{a^2(\boldsymbol{\beta}^0 \cdot \mathbf{v})^2}{2\eta\gamma + (D-d)a} \right\rangle \\
 &+ \frac{1}{2}\zeta \left\langle \log\left(\frac{2\eta\gamma}{a} + D-d\right) \right\rangle + \frac{1}{2}\zeta \left\langle \frac{da}{2\eta\gamma + (D-d)a} \right\rangle + \frac{1}{2n} \log(\tilde{S}^2 B_{00}) \\
 &- \frac{1}{n} \log \int Dz Dy_0 \int dt p(t|y_0/\sqrt{B_{00}}, \lambda^0) \\
 &\times \left[\int dy e^{-\frac{1}{2}(B-b)y^2 + y(iz\sqrt{b}-b_0y_0/\sqrt{B_{00}})} \frac{p^\gamma(t|y, \lambda)}{p^\gamma(t|y_0/\sqrt{B_{00}}, \lambda^0)} \right]^n. \tag{3.8}
 \end{aligned}$$

We note that

$$B_{00}^{-1} = \tilde{S}^2 - nc_0^2/(C-c) + \mathcal{O}(n^2), \quad B - b = 1/(C-c) \tag{3.9}$$

$$b_0 = -c_0/\tilde{S}^2(C-c) + \mathcal{O}(n), \quad b = \frac{c_0^2 - c\tilde{S}^2}{\tilde{S}^2(C-c)^2} + \mathcal{O}(n) \tag{3.10}$$

and these identities enable us, after some simple rearrangements, to write the limit $\Psi_{RS}(\dots) = \lim_{n \rightarrow 0} n^{-1} \Psi_{RS}(\dots)$ in the much simpler form

$$\begin{aligned}
 \Psi_{RS}(\dots) &= -\frac{1}{2}\zeta \left\{ 2d_0c_0 + DC - dc + \log(2\pi) + 2\eta\gamma S^2 \right. \\
 &+ d_0^2 \left\langle \frac{a^2(\boldsymbol{\beta}^0 \cdot \mathbf{v})^2}{2\eta\gamma + (D-d)a} \right\rangle - \left\langle \log\left(\frac{2\eta\gamma}{a} + D-d\right) \right\rangle - \left\langle \frac{da}{2\eta\gamma + (D-d)a} \right\rangle \left. \right\} \\
 &- \int Dz Dy_0 \int dt p(t|\tilde{S}y_0, \lambda^0) \log \int Dy \frac{p^\gamma(t|y\sqrt{C-c} + z(c-c_0^2/\tilde{S}^2)^{\frac{1}{2}} + y_0c_0/\tilde{S}, \lambda)}{p^\gamma(t|\tilde{S}y_0, \lambda^0)}. \tag{3.11}
 \end{aligned}$$

3.2. Simplification of the theory and interpretation of order parameters

Expression (3.11) can readily be extremized over d_0 , which removes a further order parameter from our theory, and we carry out a suitable transformation of the remaining order parameters,

$$u = \sqrt{C-c}, \quad v = \sqrt{c-(c_0/\tilde{S})^2}, \quad w = c_0/\tilde{S}, \quad f = d, \quad g = D-d \tag{3.12}$$

with $u, v, w \in [0, \infty)$ and with the inverse transformations

$$c_0 = \tilde{S}w, \quad c = v^2 + w^2, \quad C = u^2 + v^2 + w^2. \tag{3.13}$$

These steps result after some simple rearrangements, and upon removing the term in Ψ_{RS} that will vanish upon differentiation with respect to γ , in

$$\lim_{N \rightarrow \infty} E_\gamma(\boldsymbol{\beta}^*, \lambda^*) = \frac{\partial}{\partial \gamma} \text{extr}_{u,v,w,f,g,\lambda} \Psi_{RS}(u, v, w, f, g, \lambda) \tag{3.14}$$

in which

$$\begin{aligned}
\Psi_{\text{RS}}(\dots) = & -\frac{1}{2}\zeta(g+f)u^2 - \frac{1}{2}\zeta g(v^2+w^2) - \zeta\eta\gamma S^2 \\
& + \frac{1}{2}\zeta \left\{ \tilde{S}^2 w^2 \left\langle \frac{a^2(\boldsymbol{\beta}^0 \cdot \mathbf{v})^2}{2\eta\gamma + ga} \right\rangle^{-1} + \left\langle \log \left(\frac{2\eta\gamma + ga}{a} \right) \right\rangle + f \left\langle \frac{a}{2\eta\gamma + ga} \right\rangle \right\} \\
& - \int \text{DzDy}_0 \int \text{dt} p(t|\tilde{S}y_0, \lambda^0) \log \int \text{Dy} \frac{p^\gamma(t|uy+wy_0+ vz, \lambda)}{p^\gamma(t|\tilde{S}y_0, \lambda^0)}. \quad (3.15)
\end{aligned}$$

In principle we could also extremise over f , leading to a simple expression with which to remove not just f but also either u or g . The true association parameters $\boldsymbol{\beta}^0$ are seen to enter the asymptotic theory only in two places: in $\tilde{S}^2 = \lim_{p \rightarrow \infty} p^{-1} \boldsymbol{\beta}^0 \cdot \mathbf{A} \boldsymbol{\beta}^0$ and in $\langle a^2(\boldsymbol{\beta}^0 \cdot \mathbf{v})^2 / (2\eta\gamma + ga) \rangle$. Both are quadratic functions of $\boldsymbol{\beta}^0$. In appendix we show that, if the true associations $\{\beta_\mu^0\}$ are drawn randomly and independently from a zero-average distribution, and under mild conditions on the spectrum $\varrho(a)$ of the covariate correlation matrix \mathbf{A} , both terms will be self-averaging with respect to the realization of $\boldsymbol{\beta}^0$. Consequently, with $S^2 = \lim_{p \rightarrow \infty} p^{-1}(\boldsymbol{\beta}^0)^2$ we may then write

$$\tilde{S}^2 = S^2 \langle a \rangle, \quad \left\langle \frac{a^2(\boldsymbol{\beta}^0 \cdot \mathbf{v})^2}{2\eta\gamma + ga} \right\rangle = \left\langle \frac{S^2 a^2}{2\eta\gamma + ga} \right\rangle \quad (3.16)$$

(where we used the fact that the eigenvectors \mathbf{v} of \mathbf{A} were defined to be normalized). Our replica symmetric theory thereby becomes

$$\begin{aligned}
\lim_{N \rightarrow \infty} E_\gamma(\boldsymbol{\beta}^0, \lambda^0) = & \int \text{Dy}_0 \int \text{dt} p(t|S\langle a \rangle^{\frac{1}{2}} y_0, \lambda^0) \log p(t|S\langle a \rangle^{\frac{1}{2}} y_0, \lambda^0) - \zeta\eta S^2 \\
& + \eta\zeta \left\{ w^2 \langle a \rangle \left\langle \frac{a^2}{2\eta\gamma + ga} \right\rangle^{-2} \left\langle \frac{a^2}{(2\eta\gamma + ga)^2} \right\rangle + \left\langle \frac{1}{2\eta\gamma + ga} \right\rangle - f \left\langle \frac{a}{(2\eta\gamma + ga)^2} \right\rangle \right\} \\
& - \int \text{DzDy}_0 \int \text{dt} p(t|S\langle a \rangle^{\frac{1}{2}} y_0, \lambda^0) \frac{\int \text{Dy} p^\gamma(t|uy+wy_0+ vz, \lambda) \log p(t|uy+wy_0+ vz, \lambda)}{\int \text{Dy} p^\gamma(t|uy+wy_0+ vz, \lambda)}. \quad (3.17)
\end{aligned}$$

The scalar order parameters (u, v, w, f, g) and the function $\lambda(t)$ are computed by extremization of the following function, from which we removed any constant terms:

$$\begin{aligned}
\Psi_{\text{RS}}(\dots) = & -\frac{1}{2}\zeta(g+f)u^2 - \frac{1}{2}\zeta g(v^2+w^2) \\
& + \frac{1}{2}\zeta \left\{ w^2 \langle a \rangle \left\langle \frac{a^2}{2\eta\gamma + ga} \right\rangle^{-1} + \left\langle \log(2\eta\gamma + ga) \right\rangle + f \left\langle \frac{a}{2\eta\gamma + ga} \right\rangle \right\} \\
& - \int \text{DzDy}_0 \int \text{dt} p(t|S\langle a \rangle^{\frac{1}{2}} y_0, \lambda^0) \log \int \text{Dy} p^\gamma(t|uy+wy_0+ vz, \lambda). \quad (3.18)
\end{aligned}$$

The physical meaning of the RS order parameters can be inferred by adapting the route followed in [1]. Upon defining averages over the stochastic MAP minimisation process as $\langle \dots \rangle$ and those over the realisations of the data set as $\langle \dots \rangle_{\mathcal{D}}$, this results in

$$C = \lim_{p \rightarrow \infty} \frac{1}{p} \langle \langle \boldsymbol{\beta} \cdot \mathbf{A} \boldsymbol{\beta} \rangle \rangle_{\mathcal{D}}, \quad c = \lim_{p \rightarrow \infty} \frac{1}{p} \langle \langle \boldsymbol{\beta} \rangle \cdot \mathbf{A} \langle \boldsymbol{\beta} \rangle \rangle_{\mathcal{D}}, \quad c_0 = \lim_{p \rightarrow \infty} \frac{1}{p} \boldsymbol{\beta}^0 \cdot \mathbf{A} \langle \langle \boldsymbol{\beta} \rangle \rangle_{\mathcal{D}}. \quad (3.19)$$

These order parameters can be used to predict the slope and width of the association parameter cloud in figure 2. A plausible model for this cloud is $\langle \boldsymbol{\beta} \rangle = \kappa \boldsymbol{\beta}^0 + \boldsymbol{\omega}$, in which $\boldsymbol{\omega}$ denotes

a zero-average inference noise contribution that depends on the data set \mathcal{D} . The entries $\Omega_{\mu\nu} = \langle \omega_\mu \omega_\nu \rangle_{\mathcal{D}}$ of the $p \times p$ noise covariance matrix Ω have the same dimension as those of \mathbf{A}^{-1} , which prompt us to postulate that $\Omega = \sigma^2 \mathbf{A}^{-1}$. Inserting the above expression for $\langle \beta \rangle$ into (3.19) leads to the following identities:

$$c = \kappa^2 \tilde{S}^2 + \sigma^2, \quad c_0 = \kappa \tilde{S}^2 \quad (3.20)$$

$$u^2 = \lim_{p \rightarrow \infty} \frac{1}{p} \sum_{\mu\nu=1}^p A_{\mu\nu} \langle \langle \beta_\mu \beta_\nu \rangle - \langle \beta_\mu \rangle \langle \beta_\nu \rangle \rangle_{\mathcal{D}}. \quad (3.21)$$

Using the transformations (3.12) we then obtain the following simple expressions for the two dominant characteristics κ and σ of the simulation data clouds:

$$\kappa = w/\tilde{S}, \quad \sigma = v. \quad (3.22)$$

3.3. Scaling of order parameters with γ

We will only be interested in the limit $\gamma \rightarrow \infty$, where the stochastic process becomes deterministic MAP inference. Following [1], and with a modest amount of foresight regarding the behaviour of the new order parameters that did not feature in [1], we make the following ansatz for the scaling with γ of the scalar order parameters:

$$u = \tilde{u}/\sqrt{\gamma}, \quad v, w = \mathcal{O}(1), \quad g = \tilde{g}\gamma, \quad f = \tilde{f}\gamma^2. \quad (3.23)$$

Insertion into (3.18), followed by taking the limit $\gamma \rightarrow \infty$, gives

$$\begin{aligned} \lim_{\gamma \rightarrow \infty} \frac{1}{\gamma} \Psi_{\text{RS}}(\dots) &= \frac{1}{2} \zeta \left\{ w^2 \langle a \rangle \left\langle \frac{a^2}{2\eta + \tilde{g}a} \right\rangle^{-1} + \tilde{f} \left[\left\langle \frac{a}{2\eta + \tilde{g}a} \right\rangle - \tilde{u}^2 \right] - \tilde{g}(v^2 + w^2) \right\} \\ &\quad - \int \text{DzDy}_0 \int \text{dt} p(t|S\langle a \rangle^{\frac{1}{2}} y_0, \lambda^0) \lim_{\gamma \rightarrow \infty} \frac{1}{\gamma} \log \int \text{dy} e^{\gamma \left[\log p(t|\tilde{u}y + wy_0 + vz, \lambda) - \frac{1}{2}y^2 \right]} \\ &= \frac{1}{2} \zeta \left\{ w^2 \langle a \rangle \left\langle \frac{a^2}{2\eta + \tilde{g}a} \right\rangle^{-1} + \tilde{f} \left[\left\langle \frac{a}{2\eta + \tilde{g}a} \right\rangle - \tilde{u}^2 \right] - \tilde{g}(v^2 + w^2) \right\} \\ &\quad - \int \text{DzDy}_0 \int \text{dt} p(t|S\langle a \rangle^{\frac{1}{2}} y_0, \lambda^0) \max_y \left[\log p(t|\tilde{u}y + wy_0 + vz, \lambda) - \frac{1}{2}y^2 \right]. \end{aligned} \quad (3.24)$$

The maximization over y proceeds as in [1], giving

$$\text{argmax}_y \left[\log p(t|\tilde{u}y + wy_0 + vz, \lambda) - \frac{1}{2}y^2 \right] = \tilde{u} - \frac{1}{\tilde{u}} W \left(\tilde{u}^2 e^{\tilde{u}^2 + wy_0 + vz} \Lambda(t) \right) \quad (3.25)$$

$$\begin{aligned} \max_y \left[\log p(t|\tilde{u}y + wy_0 + vz, \lambda) - \frac{1}{2}y^2 \right] &= \frac{1}{2} (\tilde{u}^2 + \tilde{u}^{-2}) + wy_0 + vz \\ &\quad + \log \lambda(t) - \frac{1}{2\tilde{u}^2} \left[W \left(\tilde{u}^2 e^{\tilde{u}^2 + wy_0 + vz} \Lambda(t) \right) + 1 \right]^2 \end{aligned} \quad (3.26)$$

in which $W(x)$ denotes Lambert's W -function, i.e. the inverse of $f(x) = xe^x$. This then results in

$$\begin{aligned} \lim_{\gamma \rightarrow \infty} \frac{1}{\gamma} \Psi_{\text{RS}}(\dots) &= \frac{1}{2} \zeta \left[w^2 \langle a \rangle \left\langle \frac{a^2}{2\eta + \tilde{g}a} \right\rangle^{-1} + \tilde{f} \left[\left\langle \frac{a}{2\eta + \tilde{g}a} \right\rangle - \tilde{u}^2 \right] - \tilde{g}(v^2 + w^2) \right] \\ &+ \frac{1}{2\tilde{u}^2} \int \text{DzDy}_0 \int \text{dt} p(t|S\langle a \rangle^{\frac{1}{2}} y_0, \lambda^0) \left[W\left(\tilde{u}^2 e^{\tilde{u}^2 + w y_0 + v z} \Lambda(t)\right) + 1 \right]^2 \\ &- \frac{1}{2} (\tilde{u}^2 + \tilde{u}^{-2}) - \int \text{Dy}_0 \int \text{dt} p(t|S\langle a \rangle^{\frac{1}{2}} y_0, \lambda^0) \log \lambda(t). \end{aligned} \quad (3.27)$$

Similarly, working out (3.17) in the limit $\gamma \rightarrow \infty$ gives

$$\begin{aligned} \lim_{N \rightarrow \infty} E_{\infty}(\beta^0, \lambda^0) &= \int \text{Dy}_0 \int \text{dt} p(t|S\langle a \rangle^{\frac{1}{2}} y_0, \lambda^0) \left[\log p(t|S\langle a \rangle^{\frac{1}{2}} y_0, \lambda^0) - \log \lambda(t) \right] \\ &- \tilde{u}^2 - \zeta \eta S^2 + \eta \zeta \left[w^2 \langle a \rangle \left\langle \frac{a^2}{2\eta + \tilde{g}a} \right\rangle^{-2} \left\langle \frac{a^2}{(2\eta + \tilde{g}a)^2} \right\rangle - \tilde{f} \left\langle \frac{a}{(2\eta + \tilde{g}a)^2} \right\rangle \right] \\ &+ (1 + \tilde{u}^{-2}) \int \text{DzDy}_0 \int \text{dt} p(t|S\langle a \rangle^{\frac{1}{2}} y_0, \lambda^0) W\left(\tilde{u}^2 e^{\tilde{u}^2 + w y_0 + v z} \Lambda(t)\right). \end{aligned} \quad (3.28)$$

What remains in our RS analysis is to determine the order parameters $\{\tilde{u}, v, w, \tilde{f}, \tilde{g}, \lambda\}$ by extremization of (3.27), and to substitute the result into (3.28).

3.4. Scalar saddle point equations

Partial differentiation of (3.27) with respect to the five scalar order parameters $\{\tilde{u}, v, w, \tilde{f}, \tilde{g}\}$ is now straightforward and gives, upon using identities such as $W'(z) = W(z)/z[1 + W(z)]$ and after manipulations similar to those used in [1]:

$$\zeta \tilde{f} \tilde{u}^4 = - \int \text{DzDy}_0 \int \text{dt} p(t|S\langle a \rangle^{\frac{1}{2}} y_0, \lambda^0) \left[W\left(\tilde{u}^2 e^{\tilde{u}^2 + w y_0 + v z} \Lambda(t)\right) - \tilde{u}^2 \right]^2 \quad (3.29)$$

$$\zeta \tilde{g} \tilde{u}^2 = \int \text{DzDy}_0 \int \text{dt} p(t|S\langle a \rangle^{\frac{1}{2}} y_0, \lambda^0) \frac{W\left(\tilde{u}^2 e^{\tilde{u}^2 + w y_0 + v z} \Lambda(t)\right)}{1 + W\left(\tilde{u}^2 e^{\tilde{u}^2 + w y_0 + v z} \Lambda(t)\right)} \quad (3.30)$$

$$\begin{aligned} 0 &= \zeta w \left[\langle a \rangle \left\langle \frac{a^2}{2\eta + \tilde{g}a} \right\rangle^{-1} - \tilde{g} \right] \\ &+ \frac{1}{\tilde{u}^2} \int \text{DzDy}_0 \int \text{dt} p(t|S\langle a \rangle^{\frac{1}{2}} y_0, \lambda^0) W\left(\tilde{u}^2 e^{\tilde{u}^2 + w y_0 + v z} \Lambda(t)\right) \end{aligned} \quad (3.31)$$

$$\tilde{u}^2 = \left\langle \frac{a}{2\eta + \tilde{g}a} \right\rangle \quad (3.32)$$

$$v^2 = w^2 \left[\langle a \rangle \left\langle \frac{a^2}{2\eta + \tilde{g}a} \right\rangle^{-2} \left\langle \frac{a^3}{(2\eta + \tilde{g}a)^2} \right\rangle - 1 \right] - \tilde{f} \left\langle \frac{a^2}{(2\eta + \tilde{g}a)^2} \right\rangle. \quad (3.33)$$

Compared to the simpler scenario of [1], the present RS theory involves two additional order parameters, \tilde{f} and \tilde{g} . As a simple test we can set $\eta = 0$, i.e. remove the priors for association parameters. This reduces the last two of the above saddle point equations to $\tilde{g} = 1/\tilde{u}^2$ and $\tilde{f} = -v^2/\tilde{u}^4$, removes all dependencies of the theory on the spectrum $\varrho(a)$ of the covariate correlation matrix (other than via $\langle a \rangle$), and simplifies the remaining three scalar order parameter equations correctly to those derived in [1].

3.5. Functional saddle point equation

The equation from which to solve the functional order parameter $\lambda(t)$ is derived by functional differentiation of (3.27). Upon using the short-hand $p(t) = \int \text{D}y_0 \int \text{d}t p(t|S\langle a \rangle^{\frac{1}{2}}y_0, \lambda^0)$ for the typical distribution of the event times in the data, this equation takes the form

$$\frac{p(t)}{\lambda(t)} = \int \text{D}z \text{D}y_0 \int_t^\infty \frac{\text{d}t'}{\tilde{u}^2 \Lambda(t')} p(t'|S\langle a \rangle^{\frac{1}{2}}y_0, \lambda^0) W\left(\tilde{u}^2 e^{\tilde{u}^2 + w y_0 + v z} \Lambda(t')\right). \quad (3.34)$$

It differs only minimally from the one in [1], and is hence equally difficult to solve analytically. Following [1] we will therefore follow a variational approach, motivated by the asymptotic form of the solution for large times (see [1] for details), and choose the functional ansatz $\Lambda(t) = k[\Lambda^0(t)]^\rho$, leaving two variational parameters (k, ρ) to be solved, instead of a function. Inserting this ansatz into (3.27), followed by partial differentiation with respect to k and ρ then leads to the following two equations:

$$\tilde{u}^2 = \int \text{D}z \text{D}y_0 \int \text{d}t p(t|S\langle a \rangle^{\frac{1}{2}}y_0, \lambda^0) W\left(k\tilde{u}^2 e^{\tilde{u}^2 + w y_0 + v z} [\Lambda^0(t)]^\rho\right) \quad (3.35)$$

$$0 = \frac{1}{\tilde{u}^2} \int \text{D}z \text{D}y_0 \int \text{d}t p(t|S\langle a \rangle^{\frac{1}{2}}y_0, \lambda^0) W\left(k\tilde{u}^2 e^{\tilde{u}^2 + w y_0 + v z} [\Lambda^0(t)]^\rho\right) \log \Lambda^0(t) \\ - \frac{1}{\rho} - \int \text{D}y_0 \int \text{d}t p(t|S\langle a \rangle^{\frac{1}{2}}y_0, \lambda^0) \log \Lambda^0(t). \quad (3.36)$$

These have to be solved numerically alongside (3.29)–(3.33). We will compactify our equations by using instead of k the variable $q = k\tilde{u}^2 \exp(\tilde{u}^2)$. As a further benefit of our variational ansatz, the time integrations in the saddle point equations can be simplified significantly upon switching to the new integration variable $s = \exp[-\exp(S\langle a \rangle^{\frac{1}{2}}y_0)\Lambda^0(t)] \in [0, 1]$, which gives $\text{d}s = -p(t|S\langle a \rangle^{\frac{1}{2}}y_0, \lambda^0)\text{d}t$. After this transformation we can also combine the Gaussian variables into a single one, giving

$$\zeta \tilde{f} \tilde{u}^4 = - \int \text{D}x \int_0^1 \text{d}s \left[W(qe^{\sigma x} \log^\rho(1/s)) - \tilde{u}^2 \right]^2 \quad (3.37)$$

$$\zeta \tilde{g} \tilde{u}^2 = \int \text{D}x \int_0^1 \text{d}s \frac{W(qe^{\sigma x} \log^\rho(1/s))}{1 + W(qe^{\sigma x} \log^\rho(1/s))} \quad (3.38)$$

$$w = \frac{\tilde{g}\rho S}{\langle a \rangle^{\frac{1}{2}}} \left\langle \frac{a^2}{2\eta + \tilde{g}a} \right\rangle \quad (3.39)$$

$$\tilde{u}^2 = \left\langle \frac{a}{2\eta + \tilde{g}a} \right\rangle \quad (3.40)$$

$$v^2 = w^2 \left[\langle a \rangle \left\langle \frac{a^2}{2\eta + \tilde{g}a} \right\rangle^{-2} \left\langle \frac{a^3}{(2\eta + \tilde{g}a)^2} \right\rangle - 1 \right] - \tilde{f} \left\langle \frac{a^2}{(2\eta + \tilde{g}a)^2} \right\rangle \quad (3.41)$$

$$\tilde{u}^2 = \int \text{D}x \int_0^1 \text{d}s W(qe^{\sigma x} \log^\rho(1/s)) \quad (3.42)$$

$$\frac{\tilde{u}^2}{\rho} = \int \mathbf{D}x \int_0^1 ds W\left(qe^{\sigma x} \log^\rho\left(\frac{1}{s}\right)\right) \log \log\left(\frac{1}{s}\right) - \zeta \tilde{g} \tilde{u}^2 S\langle a \rangle^{\frac{1}{2}} (w - \rho S\langle a \rangle^{\frac{1}{2}}) + \tilde{u}^2 C_E \quad (3.43)$$

in which $\sigma^2 = (w - \rho S\langle a \rangle^{\frac{1}{2}})^2 + v^2$, C_E denotes Euler's constant, and where we used the integral $\int_0^1 ds \log \log(1/s) = \int_0^\infty dx e^{-x} \log x = -C_E$.

3.6. The limits $\eta \rightarrow 0$, $\zeta \rightarrow 0$ and $\zeta \rightarrow \infty$

Here we investigate the order parameter behaviour in the small and large ζ limits, to confirm the shape of the order parameter plots, both analytically and by numerical analysis in the next section. In particular, since the order parameters v and w increase with ζ for small ζ and tend to zero for large ζ , we conclude that there must be a stationary point between these two extremes. This analytical argument is validated by numerical solutions of (3.37)–(3.43) (see figure 3) and by synthetic data studies (figure 4). An explanation of this phenomenon in terms of model complexity and the emergence of statistical constraints can be constructed [33]. In the limit $\eta \rightarrow 0$, describing a fully flat prior for association parameters, the regression changes from MAP to ML, and our RS equations should therefore reduce to those of [1]. Upon setting $\eta \rightarrow 0$ in equations (3.37)–(3.43), we immediately find that

$$w = \rho S\langle a \rangle^{\frac{1}{2}}, \quad \tilde{g} = 1/\tilde{u}^2, \quad \tilde{f} = -v^2/\tilde{u}^4. \quad (3.44)$$

From the first of these it follows that $\sigma = v$, and that the remaining RS scalar order parameter equations from which to solve $\{v, \rho, \tilde{u}, q\}$ hence simplify to

$$\zeta v^2 = \int \mathbf{D}x \int_0^1 ds \left[W(qe^{vx} \log^\rho(1/s)) - \tilde{u}^2 \right]^2 \quad (3.45)$$

$$\zeta = \int \mathbf{D}x \int_0^1 ds \frac{W(qe^{vx} \log^\rho(1/s))}{1 + W(qe^{vx} \log^\rho(1/s))} \quad (3.46)$$

$$\tilde{u}^2 = \int \mathbf{D}x \int_0^1 ds W(qe^{vx} \log^\rho(1/s)) \quad (3.47)$$

$$\frac{\tilde{u}^2}{\rho} = \int \mathbf{D}x \int_0^1 ds W\left(qe^{vx} \log^\rho\left(\frac{1}{s}\right)\right) \log \log\left(\frac{1}{s}\right) + \tilde{u}^2 C_E. \quad (3.48)$$

We observe that, as a consequence of having modified our present derivation compared to the one in [1] (we changed the order of integral transformations and partial differentiations), the above expressions provide the proof for the simplifying identity $w = \rho S$ (which holds for the case where $\langle a \rangle = 1$), that was suggested by numerical analysis but not yet proven in [1]. For $\eta \rightarrow 0$ we can thus retrieve from our present results in an even more satisfactory manner the variational RS theory of [1].

For $\zeta \rightarrow 0$ (no overfitting) we expect to find $v \rightarrow 0$ and $w, \rho, k \rightarrow 1$. In analogy with [1] we now make the ansätze that $\tilde{u}, v = \mathcal{O}(\sqrt{\zeta})$ and $\rho = 1 + \mathcal{O}(\zeta)$ for $\zeta \rightarrow 0$, and expand our equations (3.37)–(3.43) in leading order for small ζ , using $W(z) = z + \mathcal{O}(z^2)$ for $z \rightarrow 0$. After expanding the various integrals, whose leading orders in ζ can all be done analytically, this results in

$$\tilde{u}^2/\zeta = 1 + \mathcal{O}(\zeta), \quad \zeta \tilde{g} = 1 + \mathcal{O}(\zeta), \quad w = S\langle a \rangle^{\frac{1}{2}} + \mathcal{O}(\zeta), \quad (3.49)$$

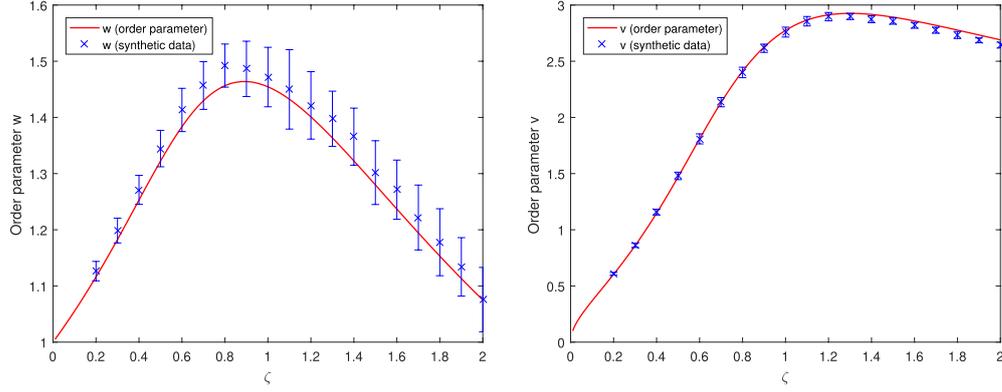


Figure 3. Predicted and measured values of the order parameters w and v (solid lines and markers, respectively), for $\mathbf{A} = \mathbb{I}_p$, $S = 1$ and $p = 2000$, shown versus $\zeta = p/N \in (0, 2]$. Measurements are determined via MAP regression, with regularization parameter $\eta = 0.025$. Simulations are repeated 50 times with independent data sets (generated according to [28], with constant hazard rates), and results shown as averages with error bars indicating one standard deviation. Note that for these settings, slope and the width of the association parameter cloud equal w and v , respectively.

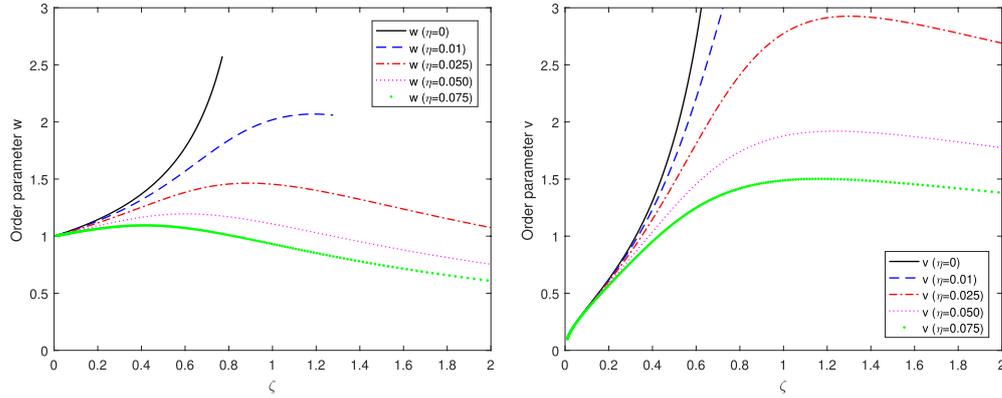


Figure 4. Predicted values of the order parameters w (left) and v (right), shown versus $\zeta = p/N$. They are obtained by solving numerically the RS equations (3.37)–(3.43) for $\mathbf{A} = \mathbb{I}_p$ and $S = 1$, with the variational approximation for $\lambda(t)$, and different choices of the regularization parameter η .

$$\tilde{f}\zeta = -1 + \mathcal{O}(\zeta), \quad k = 1 + \mathcal{O}(\zeta), \quad v^2/\zeta = 1 + \mathcal{O}(\zeta), \quad (3.50)$$

which confirms that in the absence of overfitting we indeed recover the correct values of the order parameters from our RS equations.

Finally we inspect the behaviour of the RS equations (3.37)–(3.43) in the limit $\zeta \rightarrow \infty$ of a diverging imbalance between the number of covariates and the number of samples. Note that in [1] (i.e. for $\eta = 0$) this limit was inaccessible, due to a phase transition at $\zeta = 1$, where $v, w \rightarrow \infty$. In the present theory, describing the regularized version of the Cox model, this phase transition is suppressed by the Bayesian prior, provided we choose $\eta > 0$. We now make the ansatz that $\tilde{g} \rightarrow 0$ for $\zeta \rightarrow \infty$, giving $\tilde{u}^2 \rightarrow \langle a \rangle / 2\eta$, $v \rightarrow 0$, $w \rightarrow 0$, $\tilde{f} \rightarrow 0$,

$\sigma^2 \rightarrow \rho^2 S^2 \langle a \rangle$, and upon introducing $Q = \lim_{\zeta \rightarrow \infty} \zeta \tilde{g} \tilde{u}^2$, the remaining trio $\{Q, q, \rho\}$ is for $\zeta \rightarrow \infty$ to be solved from the remaining three coupled equations

$$Q = \int \mathbf{D}x \int_0^1 ds \frac{W(qe^{\rho S \langle a \rangle \frac{1}{2} x} \log^\rho(1/s))}{1 + W(qe^{\sigma x} \log^\rho(1/s))} \tag{3.51}$$

$$\frac{\langle a \rangle}{2\eta} = \int \mathbf{D}x \int_0^1 ds W(qe^{\rho S \langle a \rangle \frac{1}{2} x} \log^\rho(1/s)) \tag{3.52}$$

$$\frac{\langle a \rangle}{2\eta\rho} = \int \mathbf{D}x \int_0^1 ds W\left(qe^{\rho S \langle a \rangle \frac{1}{2} x} \log^\rho\left(\frac{1}{s}\right)\right) \log \log\left(\frac{1}{s}\right) + QS^2 \langle a \rangle \rho + \frac{\langle a \rangle}{2\eta} C_E. \tag{3.53}$$

For $\zeta \rightarrow \infty$ we thus expect to find, as a consequence of $\lim_{\zeta \rightarrow \infty} v = \lim_{\zeta \rightarrow \infty} w = 0$, vanishing inferred association parameters in the present regularized Cox model, with the assumed scaling of the width of the prior.

3.7. Expression for the overfitting measure

Finally, using the variational approximation for the cumulative hazard rate, the simple manipulations applied to the RS saddle point equations, and the actual order parameter equations themselves, the overfitting measure (3.28) can be simplified to the transparent form

$$\begin{aligned} \lim_{N \rightarrow \infty} E_\infty(\beta^0, \lambda^0) &= \eta\zeta \left[w^2 \langle a \rangle \left\langle \frac{a^2}{2\eta + \tilde{g}a} \right\rangle^{-2} \left\langle \frac{a^2}{(2\eta + \tilde{g}a)^2} \right\rangle - \tilde{f} \left\langle \frac{a}{(2\eta + \tilde{g}a)^2} \right\rangle \right] \\ &\quad + \int dt p(t) \log \left(\frac{\lambda^0(t)}{\lambda(t)} \right) - \zeta \eta S^2 \end{aligned} \tag{3.54}$$

with the short-hand $p(t) = \int \mathbf{D}y_0 p(t|S \langle a \rangle^{\frac{1}{2}} y_0, \lambda^0)$. Our variational ansatz $\Lambda(t) = k[\Lambda^0(t)]^\rho$ implies that $\lambda(t) = k\rho\lambda^0(t)[\Lambda^0(t)]^{\rho-1}$, hence

$$\begin{aligned} \int dt p(t) \log \left(\frac{\lambda^0(t)}{\lambda(t)} \right) &= -\log k - \log \rho - (\rho-1) \int dt p(t) \log \Lambda^0(t) \\ &= -\log k - \log \rho - (\rho-1) \int \mathbf{D}y_0 \int_0^1 ds \left[e^{-S \langle a \rangle \frac{1}{2} y_0} \log\left(\frac{1}{s}\right) \right] \\ &= -\log k - \log \rho - (\rho-1) \int_0^\infty dx e^{-x} \log x \\ &= -\log k - \log \rho + (\rho-1) C_E. \end{aligned} \tag{3.55}$$

Our final result for the asymptotic overfitting measure $E(S) = \lim_{N \rightarrow \infty} E_\infty(\beta^0, \lambda^0)$ is therefore

$$\begin{aligned} E(S) &= \eta\zeta \left[w^2 \langle a \rangle \left\langle \frac{a^2}{2\eta + \tilde{g}a} \right\rangle^{-2} \left\langle \frac{a^2}{(2\eta + \tilde{g}a)^2} \right\rangle - \tilde{f} \left\langle \frac{a}{(2\eta + \tilde{g}a)^2} \right\rangle \right] \\ &\quad - \log k - \log \rho + (\rho-1) C_E - \zeta \eta S^2. \end{aligned} \tag{3.56}$$

We observe that, as was the case in [1] (without regularization), both the RS order parameter equations and the overfitting measure have within the variational approximation become completely independent of the true base hazard rate $\lambda^0(t)$.

4. Numerical experiments

4.1. Numerical solution of order parameter equations

Numerical solution of the RS saddle point equations (3.37)–(3.43), with the variational approximation for the base hazard rate, results in data as shown in figure 3. This figure corresponds to $\mathbf{A} = \mathbb{I}_p$, i.e. uncorrelated and normalized covariates, and $S = 1$. The phase transition at $\zeta = 1$ of [1] (corresponding to ML regression) is for $\eta > 0$ no longer present, due to the regularization incorporated into MAP regression. As η increases, we find the slope κ (which for the present parameter settings is identical to w) and the variance v of the data cloud decreasing.

To test the above predictions, we generated synthetic time-to-event data using zero mean covariate vectors \mathbf{z} with covariance matrix \mathbf{A} , and Gaussian random and zero-average association vectors β^0 , for different values of N and p . Base hazard rates were chosen to be constant. Event times were generated from the Cox proportional hazards model following [28], and from the simulated data we then extracted estimates of the association parameters via penalized Cox regression (using the R package, *glmnet* [17]). Upon solving our RS order parameter equations (3.37)–(3.43) for the chosen values of $\zeta = p/N$ and $S^2 = p^{-1}(\beta^0)^2$, we compared the solution with the regression outcomes via (3.19), under various conditions. By construction, there is no model mismatch, since the data are generated from the model assumed in parameter inference. Our theoretical predictions for the slope and variance agree remarkably well with the simulations; see figure 3.

The effect of covariate collinearity on the inferred regression coefficients [29] was investigated with two non-diagonal covariance matrices \mathbf{A} , both with $\lim_{p \rightarrow \infty} \langle a \rangle = 1$ and $\lim_{p \rightarrow \infty} \langle a^2 \rangle = 1 + \epsilon^2$ (hence with spectra of finite width), and $\epsilon \sim \mathcal{O}(1)$. This ensures that the requirements for self-averaging of the RS theory on the eigenvalue spectrum $\varrho(a)$ of \mathbf{A} are fulfilled. Our first choice was $A_{\mu\nu} = \delta_{\mu\nu} + (1 - \delta_{\mu\nu})\epsilon/\sqrt{p}$, with eigenvalues $1 - \epsilon/\sqrt{p}$ (multiplicity $p-1$) and $1 + (p-1)\epsilon/\sqrt{p}$ (multiplicity 1). Upon working out the spectrum-dependent quantities in the RS equations, we find that for this matrix choice they are independent of ϵ . Hence the order parameters are predicted to be identical to those for data with uncorrelated covariates. Simulations (not shown here) confirm that this is indeed the case, modulo finite size fluctuations. Our second choice for \mathbf{A} had again $A_{\mu\mu} = 1$ for all μ , but now covariates are correlated in ordered pairs: $A_{\mu,\mu+1} = A_{\mu+1,\mu} = \epsilon$ for all μ odd, with $A_{\mu\nu} = 0$ for all other $\mu \neq \nu$ (with $0 \leq \epsilon \leq 1$). This is a block diagonal matrix with $\varrho(a) = \frac{1}{2}\delta(a-1-\epsilon) + \frac{1}{2}\delta(a-1+\epsilon)$, and the RS order parameters *will* depend on the strength ϵ of the covariate correlations. In figure 5, we show the values of the order parameters v and w , as solved from the RS equations, for $S = 1$, $\eta = 0.025$ and different values of the correlation parameter ϵ , as functions of ζ . Here we again have $\kappa = w$. In the same figure we show the results of numerical simulations carried out for $\epsilon = \{0.0, 0.5, 1.0\}$ and $Np = 400\,000$. The error bars of approximately $\pm 10\%$ were not displayed for clarity. The covariates were generated according to: $z_\mu^i = y_{i\mu}$ for μ odd, and $z_\mu^i = \epsilon y_{i\mu-1} + \sqrt{1-\epsilon^2} y_{i\mu}$, in which all $\{y_{i\mu}\}$ are independent Gaussian random variables, with $\langle y_{i\mu} \rangle = 0$ and $\langle y_{i\mu}^2 \rangle = 1$. This choice generates the above covariate correlations $A_{\mu\nu}$. The markers each represent averages over 32 regressions with distinct covariate and association realizations. The agreement between theory and simulations is seen to be quite satisfactory. We observe that the effect of covariate correlations on the overfitting noise is always a reduction (v decreases with ϵ).

Covariates of real survival data can obviously be distributed in many different ways. The assumption in (2.9) of Gaussian distributed risk scores is a direct consequence of working in the limit $p \rightarrow \infty$, in combination with the Central Limit Theorem. More specifically, there is no

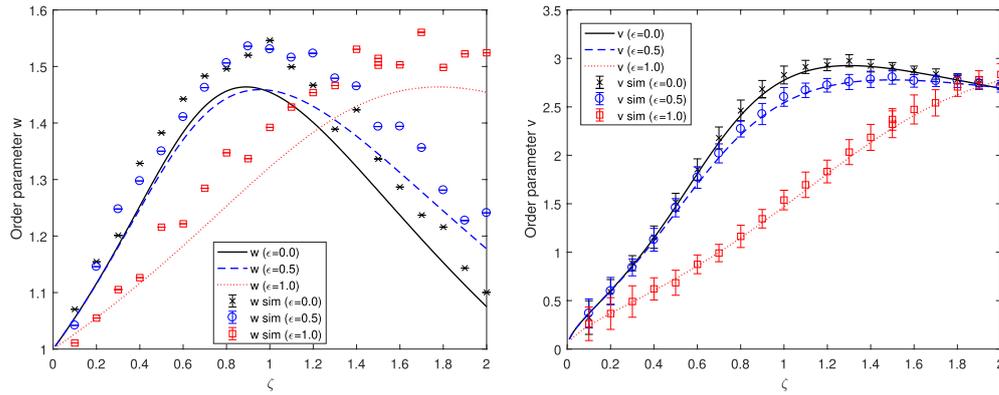


Figure 5. Predicted values of the order parameters w (left) and v (right), shown versus $\zeta = p/N$. They are obtained by solving numerically the RS equations (3.37)–(3.43) for $\eta = 0.025$ and $S = 1$, with the variational approximation for $\lambda(t)$. Here the covariates are pairwise correlated according to $A_{\mu,\mu+1} = A_{\mu+1,\mu} = \epsilon$ for all μ odd, with $A_{\mu\nu} = 0$ for all other $\mu \neq \nu$, with $\epsilon \in [0, 1]$. Note that for these settings, w and v are the slope and the width of the association data cloud. For the left w plot, only mean simulation values are shown since including the errors bars of approximately $\pm 10\%$ led to cluttered plots. Error bars can be displayed clearly for all values of ϵ on the right v plot. The markers each represent averages over 32 regressions with distinct covariate and association realizations and we fix the value of $Np = 400\,000$.

need to assume Gaussian covariate statistics. To verify the validity of Gaussian risk score statistics, we carried out simulations with four common covariate distributions, all with identical first two moments: normal, $p(z_i) = \mathcal{N}(0, 1)$, Rademacher, $p(z_i) = \frac{1}{2}\delta(z_i - 1) + \frac{1}{2}\delta(z_i + 1)$, uniform, $z_i \sim \mathcal{U}(-\sqrt{3}, \sqrt{3})$ and the student t-distribution, $z_i \sim t\text{-dist}(\nu)/\sqrt{\nu/(\nu-2)}$ (with degrees of freedom $\nu = 5$). The deviations between predictions using Gaussian covariates and the above distributions were indeed small ($<1\%$ for w and $<0.5\%$ for v) validating our asymptotic assumption that our theory admits a range of covariates distributions.

In MAP analyses the regularization parameter η is usually determined by k -fold cross-validation, or via the Generalized Cross Validation (GCV) estimator [30]. A fraction of the data is set aside for this purpose, leaving fewer samples available for inference of model parameters. This has a detrimental effect on inference accuracy. Our present theory, in contrast, suggests a more data efficient method of estimating the amount of regularization needed, without the need to sacrifice any samples. By fixing the slope parameter to unbiased recovery of the regression coefficients, i.e. $w/\tilde{S} = 1$, and solving the order parameter equations (3.37)–(3.43) with η as a parameter to be determined (instead of w), the *optimal* values of η can be estimated without any cross-validation; see figure 6. The optimal values in figure 6 are seen to match those (ζ, η) pairs in figure 3 where $w/\tilde{S} = 1$, as they should. For example, when $\zeta = 1$, the required amount of regularization to compensate for high covariate dimensionality can be read off from figure 6 to be $\eta \approx 0.05$. In interpreting this figure, however, we should note our rescaling of our association parameters, prompted by the observation that $\beta^2 \sim O(1)$ is required to avoid non-finite event times for large p . This implied that our L_2 prior in MAP inference is of the form $p(\beta) \propto \exp(-\eta p \beta^2)$.

The upward sloping region of figure 6, for small ζ , matches our intuition of requiring an increasing amount of regularization for an increasing ζ (up to $\zeta \approx 0.01$). However, as ζ is increased further, we see that optimal regularization now requires a decreasing value of η . To

test this less intuitive prediction, we chose four larger values of ζ , read off the required values of η for unbiased inference from figure 6, and calculated the slope of the association parameter cloud from 100 simulations. These predictions were made with $p = 250$ suggesting our theory is valid for relatively low values of p and N . The results show that the slope of the association parameter cloud is indeed unity, i.e. for the η values proposed by the RS theory, the overfitting-induced inference bias is indeed suppressed as predicted; see the table below:

ζ	Required η	Corresponding <i>glmnet</i> λ	Mean slope \pm 1 s.d
0.110	0.165	0.036	1.007 ± 0.028
0.552	0.100	0.110	1.009 ± 0.081
1.055	0.062	0.131	1.013 ± 0.094
2.001	0.031	0.124	0.956 ± 0.139

Note that figure 6, together with our confirmation in regression simulations that the predicted optimal values of η indeed induce unbiased MAP estimators for regression coefficients (i.e. slopes $\kappa = 1$ in the association parameter clouds), confirm *a posteriori* the correctness of the chosen scaling with p of our L_2 prior $p(\beta) \propto \exp(-\eta p \beta^2)$. In those situations where the conditions for our theory to apply are not met, other properties may of course affect the optimal value of η . For instance, our simulated data are generated from the Cox model where the ground truth association vector β^* is not sparse. Equally, the data could be generated from a model with fewer nonzero associations, including choices for which the Central Limit Theorem no longer guarantees that the risk scores $\beta^* \cdot \mathbf{z}$ have Gaussian statistics.

5. Discussion

Failure to correct multivariate ML or MAP regression results for overfitting can lead to serious inference errors. The inferred regression coefficients of the multivariate Cox model are known to be increasingly biased as the ratio of data dimension p to the sample size N increases. For medical time-to-event analysis, where it is possible to obtain (and common to have) large numbers of measurements per patient, such as genomic, epigenetic and imaging covariates, this bias is quite problematic. It induces false positive associations, which will inevitably turn out to be non-reproducible. This leads to a preventable waste of time and health funds, and frustrates the translation of the significant progress made in recent decades in medical data acquisition into effective data-driven personalized medicine. In this paper, which builds on the recent study [1], we have built successfully a theory to predict this bias for the multivariate Cox model in the presence of ridge regularization, when the data dimension scales as $p \sim N$. This paves the way further for effective overfitting corrections in multivariate MAP inference. Alternatively, our analysis allows for a straightforward analytical determination of the optimal regularization needed to correct the overfitting bias, without having to sacrifice valuable training data to cross-validation. In addition to overfitting-induced inference bias, there is a further effect of overfitting on inferred error bars. To determine the statistical significance of inferred regression coefficients, p-values are typically used. These rely on asymptotic results which do not hold in the regime where both p and N are large with $\zeta = p/N \sim \mathcal{O}(1)$ [4, 6], leading to incorrect rejections of the null hypothesis. Our theory shows that the variance of the inferred regression coefficients around the true value is a function of p/N , necessitating an adjustment to traditional test statistics used in p-value calculations.

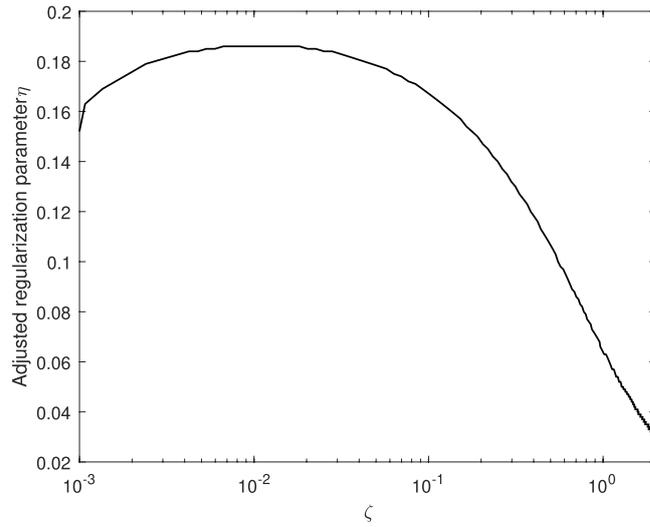


Figure 6. The present theory allows for the analytical identification of the optimally adjusted MAP regularization parameter for Cox regression, by solving the RS order parameter equations (3.37)–(3.43) upon demanding unbiased recovery of regression coefficients, $\kappa = 1$, with η as parameter to be solved instead of w . Here we show the result versus $\zeta = p/N$. It is not straightforward to solve the order parameter equations close to $\zeta = 0$, but we know that the curve should tend to the origin for $\zeta = 0$ (where ML inference is asymptotically exact).

The aim of our theory is to provide epidemiologists and clinical trials practitioners with a means of analysing data where $p \sim \mathcal{O}(N)$. This paper considers a student-teacher learning problem where we assume the data-generating model is known. A practical overfitting correction protocol for multivariate MAP regression on high-dimensional time-to-event data, based on our present theory, requires knowledge of the values of S (the magnitude $|\beta^0|$ of the true association parameter vector) and of the eigenvalue spectrum of the covariate correlation matrix \mathbf{A} . For synthetic data, these are available by assumption. For real data, S can be computed from the inferred regression parameters $\hat{\beta}$, alongside the RS order parameters, using (3.19), from which one infers the relation $v^2 + w^2 = \hat{\beta} \cdot \mathbf{A} \hat{\beta}$ (in non-rescaled notation). The value of $\hat{\beta}$ is available in practice as it is the outcome of the regression. We typically only have access to the empirical covariance matrix from which to infer the covariate correlation matrix \mathbf{A} . A possible solution for this problem is to use the link between the empirical and population level eigenvalue distributions in the Marchenko–Pastur equation [31]. The population spectrum can be estimated from its empirical counterpart in [22], by applying convex optimisation to the inverted Marchenko–Pastur equation. This method is an improvement on naively using the sample eigenvalue spectrum as an estimator of its population counterpart when $p > N$.

There are many directions for extension of the present line of research. For instance, time-to-event data, whether from observational studies or clinical trials, are typically censored. Censoring may reflect the impact of competing risks, patients withdrawing from studies, or finite study durations. The incorporation of censoring into our theory is an obvious next research target, together with investigation of regimes where the risk score are no longer Gaussian distributed. Finally, the overfitting measure in (2.3) is quite general, and can be applied to many other survival analysis models [32]. Equally, the theory developed in this paper is directly applicable to time-to-event studies outside medical data such as credit risk analysis.

Acknowledgments

MS is supported by the Biotechnology and Biological Sciences Research Council (award 1668568) and GSK Ltd. He would like to thank the Isaac Newton Institute for Mathematical Sciences for support and hospitality during the programme ‘Statistical Scalability’ when some of the work in this paper was undertaken. Both authors are grateful to Alexander Mozeika and Fabian Aguirre Lopez for constructive discussions.

Appendix. Self-averaging with respect to true associations

Here we investigate properties of random variables of the form $\mathcal{R} = p^{-1}\beta^0 \cdot \mathbf{P}\beta^0$ in the limit $p \rightarrow \infty$, where the true association vectors $\beta^0 = \{\beta_\mu^0\}$ are drawn randomly from some distribution $p(\beta^0)$ and \mathbf{P} is a fixed symmetric positive definite $p \times p$ matrix, which is independent of β^0 . In particular, we wish to determine under which conditions \mathcal{R} will be self-averaging, i.e. $\lim_{p \rightarrow \infty} \langle \mathcal{R} \rangle > 0$ exists, and $\lim_{p \rightarrow \infty} [\langle \mathcal{R}^2 \rangle - \langle \mathcal{R} \rangle^2] = 0$. Brackets will in this appendix denote averaging over $p(\beta^0)$, and we will write the eigenvalue distribution of \mathbf{P} as $\varrho(\lambda)$. We make the following assumptions⁷:

- (i) The $\{\beta_\mu^0\}$ are independent and identically distributed, i.e. $p(\beta^0) = \prod_{\mu=1}^p p(\beta_\mu^0)$.
- (ii) $p(\beta_\mu^0)$ is symmetric in β_μ^0 , with finite second and fourth order moments.

In view of our earlier definition $S^2 = \lim_{p \rightarrow \infty} p^{-1} \langle (\beta_\mu^0)^2 \rangle$, we must identify $\langle (\beta_\mu^0)^2 \rangle = S^2$. We will write $\Sigma = \langle (\beta_\mu^0)^4 \rangle$. It then follows that

$$\lim_{p \rightarrow \infty} \langle \mathcal{R} \rangle = \lim_{p \rightarrow \infty} \frac{1}{p} \sum_{\mu\nu=1}^p \langle \beta_\mu^0 \beta_\nu^0 \rangle P_{\mu\nu} = S^2 \lim_{p \rightarrow \infty} \int d\lambda \varrho(\lambda) \lambda \quad (\text{A.1})$$

$$\begin{aligned} \lim_{p \rightarrow \infty} \langle \mathcal{R}^2 \rangle &= \lim_{p \rightarrow \infty} \frac{1}{p^2} \sum_{\mu\nu\kappa\tau=1}^p \langle \beta_\mu^0 \beta_\nu^0 \beta_\kappa^0 \beta_\tau^0 \rangle P_{\mu\nu} P_{\kappa\tau} \\ &= \lim_{p \rightarrow \infty} \frac{1}{p^2} \left\{ S^4 \text{Tr}^2(\mathbf{P}) + 2S^4 \text{Tr}(\mathbf{P}^2) + (\Sigma - 3S^4) \sum_{\mu=1}^p (P_{\mu\mu})^2 \right\} \\ &\leq \left(\lim_{p \rightarrow \infty} \langle \mathcal{R} \rangle \right)^2 + \lim_{p \rightarrow \infty} \frac{1}{p} (\Sigma - S^4) \int d\lambda \varrho(\lambda) \lambda^2. \end{aligned} \quad (\text{A.2})$$

We conclude that \mathcal{R} will be self-averaging in the limit $p \rightarrow \infty$ if $\lim_{p \rightarrow \infty} \int d\lambda \varrho(\lambda) \lambda$ exists and $\lim_{p \rightarrow \infty} p^{-1} \int d\lambda \varrho(\lambda) \lambda^2 = 0$. Equivalently,

$$\lim_{p \rightarrow \infty} \frac{1}{p} \sum_{\mu=1}^p P_{\mu\mu} \in \mathfrak{R} \quad \text{and} \quad \lim_{p \rightarrow \infty} \frac{1}{p^2} \sum_{\mu\nu=1}^p P_{\mu\nu}^2 = 0. \quad (\text{A.3})$$

The two relevant quadratic expression for which we seek to demonstrate self-averaging are the following:

- Application to $\mathbf{P} = \mathbf{A}$ tells us that if $\lim_{p \rightarrow \infty} \langle a \rangle \in \mathfrak{R}$ and $\lim_{p \rightarrow \infty} p^{-1} \langle a^2 \rangle = 0$ (i.e. the covariate correlations are not excessive), then

⁷ Assuming distinct variances for each β_μ^0 complicates various equations but ultimately leads to similar final conditions on the eigenvalue spectrum of \mathbf{A} .

$$\tilde{S}^2 = \lim_{p \rightarrow \infty} \frac{1}{p} \beta^0 \cdot \mathbf{A} \beta^0 = S^2 \langle a \rangle. \tag{A.4}$$

- Our second application is to the following matrix, in which the vectors $\{\mathbf{v}^\mu\}$ are the orthogonal and normalised eigenvectors of \mathbf{A} , with eigenvalues a_μ :

$$P_{\mu\nu} = \sum_{\rho=1}^p \frac{a_\rho^2 v_\mu^\rho v_\nu^\rho}{2\eta\gamma + ga_\rho}. \tag{A.5}$$

Here we find, anticipating that $g > 0$ and using $\eta\gamma > 0$,

$$\frac{1}{p} \sum_{\mu=1}^p P_{\mu\mu} = \frac{1}{p} \sum_{\rho=1}^p \frac{a_\rho^2}{2\eta\gamma + ga_\rho} \leq \frac{\langle a \rangle}{g} \tag{A.6}$$

$$\begin{aligned} \frac{1}{p^2} \sum_{\mu\nu=1}^p P_{\mu\nu}^2 &= \frac{1}{p^2} \sum_{\rho\rho'\mu\nu=1}^p \frac{a_\rho^2 a_{\rho'}^2 v_\mu^\rho v_\nu^\rho v_\mu^{\rho'} v_\nu^{\rho'}}{(2\eta\gamma + ga_\rho)(2\eta\gamma + ga_{\rho'})} \\ &= \frac{1}{p^2} \sum_{\rho=1}^p \frac{a_\rho^4}{(2\eta\gamma + ga_\rho)^2} \leq \frac{\langle a^2 \rangle}{pg}. \end{aligned} \tag{A.7}$$

We conclude, provided $g > 0$, that the same two conditions on \mathbf{A} that guarantee self-averaging of \tilde{S}^2 for $p \rightarrow \infty$ will also imply self-averaging here:

$$\lim_{p \rightarrow \infty} \frac{1}{p} \sum_{\rho=1}^p \frac{a_\rho^2 (\beta^0 \cdot \mathbf{v}^\rho)^2}{2\eta\gamma + ga_\rho} = \left\langle \frac{S^2 a^2}{2\eta\gamma + ga} \right\rangle. \tag{A.8}$$

Thus, for our RS theory to be self-averaging with respect to the realisation of the true association vector β^0 (given our mild assumptions on the distribution from which β^0 is drawn), it is sufficient that average and width of the eigenvalue distribution $\varrho(a)$ of the covariate correlation matrix \mathbf{A} remain finite in the limit $p \rightarrow \infty$.

ORCID iDs

Mansoor Sheikh  <https://orcid.org/0000-0001-6566-5562>

References

[1] Coolen A C C, Barrett J E, Paga P and Perez-Vicente C J 2017 Replica analysis of overfitting in regression models for time-to-event data *J. Phys. A: Math. Gen.* **50** 375001
 [2] Cox D R 1972 Regression models and life-tables *J. R. Stat. Soc. B* **34** 187–202
 [3] Anderson J A and Richardson S C 1979 Logistic discrimination and bias correction in maximum likelihood estimation *Technometrics* **21** 71–8
 [4] Sur P and Candès E J 2019 A modern maximum-likelihood theory for high-dimensional logistic regression *Proc. Natl Acad. Sci. USA* **116** 14516–25
 [5] Giles D E *et al* 2009 Bias of the maximum likelihood estimators of the two-parameter gamma distribution revisited *Technical Report* Department of Economics, University of Victoria
 [6] Fan Y, Demirkaya E and Lv J 2019 Nonuniformity of p-values can occur early in diverging dimensions *J. Mach. Learn. Res.* **20** 1–33

- [7] Concato J, Peduzzi P, Holford T R and Feinstein A R 1995 Importance of events per independent variable in proportional hazards analysis I. Background, goals, and general strategy *J. Clin. Epidemiol.* **48** 1495–501
- [8] Peduzzi P, Concato J, Feinstein A R and Holford T R 1995 Importance of events per independent variable in proportional hazards regression analysis II. Accuracy and precision of regression estimates *J. Clin. Epidemiol.* **48** 1503–10
- [9] Vittinghoff E and McCulloch C E 2007 Relaxing the rule of ten events per variable in logistic and Cox regression *Am. J. Epidemiol.* **165** 710–8
- [10] Courvoisier D S, Combescure C, Agoritsas T, Gayet-Ageron A and Perneger T V 2011 Performance of logistic regression modeling: beyond the number of events per variable, the role of data structure *J. Clin. Epidemiol.* **64** 993–1000
- [11] Greenland S, Mansournia M A and Altman D G 2016 Sparse data bias: a problem hiding in plain sight *BMJ Br. Med. J.* **352** i1981
- [12] Firth D 1993 Bias reduction of maximum likelihood estimates *Biometrika* **80** 27–38
- [13] Cordeiro G M and McCullagh P 1991 Bias correction in generalized linear models *J. R. Stat. Soc. B* **53** 629–43
- [14] Fan J and Lv J 2010 A selective overview of variable selection in high dimensional feature space *Stat. Sinica* **20** 101
- [15] Battey H S and Cox D R 2018 Large numbers of explanatory variables: a probabilistic assessment *Proc. Royal Soc. A* **474** 20170631
- [16] George E I and McCulloch R E 1997 Approaches for Bayesian variable selection *Stat. Sinica* **7** 339–73
- [17] Friedman J, Hastie T and Tibshirani R 2009 Glmnet: Lasso and elastic-net regularized generalized linear models *R Package Version* **1**(4)
- [18] Tibshirani R 1996 Regression shrinkage and selection via the lasso *J. R. Stat. Soc. B* **58** 267–88
- [19] Witten D M and Tibshirani R 2010 Survival analysis with high-dimensional covariates *Stat. Methods Med. Res.* **19** 29–51
- [20] Huang J and Harrington D 2002 Penalized partial likelihood regression for right-censored data with bootstrap selection of the penalty parameter *Biometrics* **58** 781–91
- [21] Therneau T M and Lumley T 2017 Package ‘survival’ Published on CRAN *R Package Version* 2–41
- [22] El Karoui N *et al* 2008 Spectrum estimation for large dimensional covariance matrices using random matrix theory *Ann. Stat.* **36** 2757–90
- [23] Fan Y and Tang C Y 2013 Tuning parameter selection in high dimensional penalized likelihood *J. R. Stat. Soc. B* **75** 531–52
- [24] Fyodorov Y V 2019 A spin glass model for reconstructing nonlinearly encrypted signals corrupted by noise *J. Stat. Phys.* **175** 789–818
- [25] Shinzato T 2018 Replica approach for minimal investment risk with cost *J. Phys. Soc. Japan* **87** 064801
- [26] Ciliberti S and Mézard M 2007 Risk minimization through portfolio replication *Eur. Phys. J. B* **57** 175–80
- [27] Barbier J and Macris N 2017 The stochastic interpolation method: a simple scheme to prove replica formulas in Bayesian inference (arXiv:1705.02780)
- [28] Bender R, Augustin T and Blettner M 2005 Generating survival times to simulate Cox proportional hazards models *Stat. Med.* **24** 1713–23
- [29] Stewart G W 1987 Collinearity and least squares regression *Stat. Sci.* **2** 68–84
- [30] Craven P and Wahba G 1978 Smoothing noisy data with spline functions *Numer. Math.* **31** 377–403
- [31] Marchenko V A and Pastur L A 1967 Distribution of eigenvalues for some sets of random matrices *Math. USSR SB +* **114** 507–36
- [32] Klein J P and Moeschberger M L 2006 *Survival Analysis: Techniques for Censored and Truncated Data* (Berlin: Springer)
- [33] Bulso N, Marsili M and Roudi Y 2019 On the complexity of logistic regression models (arXiv:1903.00386)