

Simulation-based Regularized Logistic Regression

Robert B. Gramacy* and Nicholas G. Polson†

Abstract. In this paper, we develop a simulation-based framework for regularized logistic regression, exploiting two novel results for scale mixtures of normals. By carefully choosing a hierarchical model for the likelihood by one type of mixture, and implementing regularization with another, we obtain new MCMC schemes with varying efficiency depending on the data type (binary v. binomial, say) and the desired estimator (maximum likelihood, maximum *a posteriori*, posterior mean). Advantages of our omnibus approach include flexibility, computational efficiency, applicability in $p \gg n$ settings, uncertainty estimates, variable selection, and assessing the optimal degree of regularization. We compare our methodology to modern alternatives on both synthetic and real data. An R package called `reglogit` is available on CRAN.

Keywords: logistic regression, regularization, z -distributions, data augmentation, classification, Gibbs sampling, lasso, variance-mean mixtures, Bayesian shrinkage.

1 Introduction

Large scale logistic regression has numerous modern day applications from text classification to genetics. We develop a flexible framework for maximum likelihood, maximum *a posteriori*, and full Bayesian posterior inference for regularized models. Our motivations stem from a desire to find common ground between point estimation in “large- p ” settings (Krishnapuram et al. 2005; Genkin et al. 2007), where p is the number of predictors, and full Bayesian inference for “small- p ” (Holmes and Held 2006; Frühwirth-Schnatter and Frühwirth 2007; Frühwirth-Schnatter et al. 2009; Fahrmeir et al. 2010; Frühwirth-Schnatter and Frühwirth 2010). Collecting such distinct methods into a unifying framework facilitates a number of novel enhancements including posterior inference for the amount of regularization, and an efficient handling of binomial data.

We start by framing a typical regularized optimization criterion as a powered-up posterior, or *power-posterior* (Friel and Pettitt 2008), with a shrinkage prior such as the lasso (Tibshirani 1996). We then show how inference may proceed by employing two (heretofore unrelated) data augmentation schemes: one for the powered-up logistic likelihood; and the other for the prior. The combined effect is a fully Gibbs Markov chain Monte Carlo (MCMC) sampler which, among other advantages, allows estimators previously requiring custom algorithms to be calculated via a single simulated annealing (Kirkpatrick et al. 1983) scheme.

Specifically, consider a set of binary responses, y_i , encoded as ± 1 , regressed on p -

*Booth School of Business, Chicago, IL, rgramacy@chicagobooth.edu

†Booth School of Business, Chicago, IL, nicholas.polson@chicagobooth.edu

dimensional predictors x_i via the logistic model $\Pr(y_i = \pm 1 | x_i, \beta) = (1 + e^{-y_i x_i^\top \beta})^{-1}$, for $i = 1, \dots, n$. When p is large it is paramount to infer β under regularization or penalization. A common formulation (e.g., Genkin et al. 2007; Park and Hastie 2008) involves finding regularized point-estimates $\hat{\beta}$ under an L_α -norm penalty, where parameters $\sigma^2 = (\sigma_1^2, \dots, \sigma_p^2)$ control the relative penalization applied to each predictor, and \top is a transpose operation

$$\hat{\beta} = \operatorname{argmin}_\beta \sum_{i=1}^n \ln \left(1 + e^{-y_i x_i^\top \beta} \right) + \nu^{-\alpha} \sum_{j=1}^p \left| \frac{\beta_j}{\sigma_j} \right|^\alpha, \quad \alpha > 0. \quad (1)$$

The parameter ν dictates the amount of regularization, the relative pull (ν^{-1}) or shrinkage of the β_j 's towards zero. Depending on the choice of α , a number of algorithms have been proposed to solve for $\hat{\beta}$. For example, Madigan and Ridgeway (2004) discuss how the Least Angle Regression algorithm (LARS) can be useful as a subroutine for the popular case of $\alpha = 1$. It is typical to work with x_i pre-scaled to have unit L_2 -norm with $\sigma_1 = \dots = \sigma_p = 1$ so that inference for β is equivariant under a re-scaling of the covariates. We follow this convention in application but develop much of the discussion in the general case for completeness. The special setting $\sigma_j^2 = \infty$ indicates no shrinkage for β_j . At least $\max\{0, p - n\}$ of the σ_j^2 's must be finite to obtain stable estimators. If there is an intercept in the model, denoted by β_0 , then it is common practice to set $\sigma_0^2 = \infty$. Throughout we begin the j -indexing at $j = 1$, ignoring the 0th term for simplicity.

Our approach offers a fully probabilistic alternative by viewing the objective function (1) as a log posterior distribution whose maximum *a posteriori* (MAP) estimator coincides with $\hat{\beta}$. A multiplicity parameter κ can then be introduced to help find the MAP via simulation. Our key insight, which makes the simulation efficient, is that the logistic likelihood component of the posterior can be written hierarchically using z -distributions (Barndorff-Nielsen et al. 1982), leading to a data augmentation scheme that generalizes that of Holmes and Held (2006) [HH hereafter]. Combining this with a standard data augmentation for the prior yields a highly blocked Gibbs MCMC algorithm for logistic regression. Z -distributions also suggest a new representation of the likelihood that is equivalent (to HH) but requires n fewer latent variables. Finally, we recognize that κ has a secondary use for binomial data (multiple y observed for each x) which otherwise would require more latent variables.

A distinctive feature of our framework is how it deals with the amount of regularization, ν , which is traditionally chosen by cross validation (CV). As an alternative, we may extend the hierarchical model to include a prior for ν so that the marginal likelihood can be computed and used to set $\nu = \hat{\nu}$, or to integrate ν out. Posterior expectations, thus obtained, can give superior point-estimators for β in large- p linear regression contexts (Hans 2009), and we show how this extends to logistic regression.

The rest of the paper is outlined as follows. Section 2 provides our data augmentation strategies for sparse high dimensional logistic regression, and Section 3 develops an MCMC scheme for estimation. Section 4 illustrates our approach with empirical comparisons to modern competitors. Finally, Section 5 concludes with simple extensions

and directions for future research. A supporting R package, `reglogit`, is available on CRAN.

2 Regularized logistic regression via power-posteriors

The central problem is to find the MLE, MAP, or posterior mean estimator in logistic regression. To do this, consider the following power-posterior distribution inspired by Eq. (1):

$$\pi_{\kappa,\alpha}(\beta|y, \nu, \sigma^2) = C_{\kappa,\alpha}(\nu) \exp \left\{ -\kappa \left(\sum_{i=1}^n \ln \left(1 + e^{-y_i x_i^\top \beta} \right) + \nu^{-\alpha} \sum_{j=1}^p \left| \frac{\beta_j}{\sigma_j} \right|^\alpha \right) \right\}. \quad (2)$$

The placement of κ and α as subscripts in $\pi_{\kappa,\alpha}$ and $C_{\kappa,\alpha}(\nu)$, a normalization factor, signals that these are user specified, not parameters to be estimated. The α setting indicates the type of L_α regularization, e.g., L_1 for absolute, and L_2 for quadratic. The *multiplicity* (or thermodynamic) parameter κ , is a tool borrowed from the power-posterior and simulated annealing literature (see, e.g., Pincus 1968; Kirkpatrick et al. 1983; Doucet et al. 2002; Jacquier et al. 2007; Friel and Pettitt 2008), that facilitates several types of simulation based inference, as we shall describe.

Power-posterior analysis can be helpful for calculating modes and posterior means from complex optimization criteria, and marginal likelihoods for Bayesian estimators. Larger values of κ cause the density to concentrate near the modes, whereas small κ distributes it away from the modes, in the troughs. This motivates two types of estimators. First, $E_{\kappa,\alpha}\{\beta|y, \sigma^2, \nu\}$ can be estimated for choices of ν by allowing κ to vary as in simulated annealing. When $\nu = 0$ the estimator converges to the MLE as $\kappa \rightarrow \infty$. When $\nu > 0$, it converges to a posterior mode, or equivalently the regularized estimator, $\hat{\beta}$ solving Eq. (1). Furthermore, setting $\kappa = 1$ yields the posterior mean estimator. Second, we recognize that κ can be used to obtain an efficient computational framework for binomial regression, where multiple binary responses are recorded for each predictor. In what immediately follows, we regard κ as fixed—a further discussion is deferred to Section 3.

Observe that the likelihood–prior combination below yields Eq. (2) via Bayes’ rule.

$$L_\kappa(y|\beta) = e^{-\kappa \sum_{i=1}^n \ln \left(1 + e^{-y_i x_i^\top \beta} \right)} = \prod_{i=1}^n \left(1 + e^{-y_i x_i^\top \beta} \right)^{-\kappa} \quad (3)$$

$$p_{\kappa,\alpha}(\beta|\nu, \sigma^2) \propto \exp \left(-\kappa \nu^{-\alpha} \sum_{j=1}^p \left| \beta_j / \sigma_j \right|^\alpha \right) = \prod_{j=1}^p \exp \left\{ -\kappa \left| \frac{\beta_j}{\nu \sigma_j} \right|^\alpha \right\}.$$

The following subsections provide data augmentation schemes for this likelihood and prior. They primarily concentrate on the $\alpha = 1$ case, i.e., the double–exponential or lasso prior, although results are developed in generality when possible. Section 5 briefly touches on the simpler $\alpha = 2$ case.

2.1 Hierarchical representation of the logistic

Extending a well-known technique for generating logistic regression (e.g., [Andrews and Mallows 1974](#); [Holmes and Held 2006](#)), we represent the powered-up likelihood (3) for β as a marginal quantity obtained after integrating over latent variables (z, λ) , where $z = (z_1, \dots, z_n)$ and $\lambda = (\lambda_1, \dots, \lambda_n)$. That is, each element of the product of independent logistic terms can be written as a two-dimensional integral:

$$L_\kappa(y|\beta) = \prod_{i=1}^n \int_0^\infty \int_0^\infty p_\kappa(z_i|\beta, \lambda_i, y_i) p_\kappa(\lambda_i) d\lambda_i dz_i. \quad (4)$$

This suggests a hierarchical representation in terms of latent variables, z_i for each y_i , mixed over λ_i . It remains to determine the appropriate form of $p_\kappa(z_i|\beta, \lambda_i, y_i)$ and $p_\kappa(\lambda_i)$ so that $(1 + e^{-y_i x_i^\top \beta})^{-\kappa} = \iint p_\kappa(z_i|\beta, \lambda_i, y_i) p_\kappa(\lambda_i) d\lambda_i dz_i$.¹

Our key result, generalizing HH, relies on a scale mixture representation of z -distributions ([Barndorff-Nielsen et al. 1982](#)). These are characterized by their pdf as:

$$\begin{aligned} Z(z; a, b, \sigma, \mu) &\equiv f_Z(z|a, b, \sigma, \mu) = \frac{1}{\sigma B(a, b)} \frac{e^{a(z-\mu)/\sigma}}{(1 + e^{(z-\mu)/\sigma})^{a+b}} \\ &= \int_0^\infty \frac{1}{\sqrt{2\pi\lambda\sigma^2}} \exp\left\{-\frac{1}{2\lambda\sigma^2} \left(z - \mu - \frac{1}{2}(a-b)\lambda\sigma\right)^2\right\} q_{a,b}(\lambda) d\lambda \end{aligned} \quad (5)$$

where $q_{a,b}(\lambda)$ is a Polya distribution, i.e., an infinite sum of exponentials:

$$q_{a,b}(\lambda) = \sum_{k=0}^{\infty} w_k e^{-\frac{1}{2}\psi_k \lambda} \quad \text{where} \quad \psi_k = (a+k)(b+k), \quad (6)$$

and the weights w_k are determined via $\delta = (a+b)/2$ and $\theta = (a-b)/2$ as

$$w_k = \binom{-2\delta}{k} \frac{(\delta+k)}{B(\delta+\theta, \delta-\theta)} = \frac{(-1)^k (2\delta) \dots (2\delta+k-1)}{k!} \frac{(\delta+k)}{B(\delta+\theta, \delta-\theta)}. \quad (7)$$

This prior has a simple generative form:

$$\lambda \stackrel{D}{=} \sum_{k=0}^{\infty} 2\psi_k^{-1} \epsilon_k, \quad \text{where} \quad \epsilon_k \sim \text{Exp}(1). \quad (8)$$

Then, each component $(1 + e^{y x^\top \beta})^{-\kappa}$ of the likelihood (dropping i subscripts) can be written as the cumulative distribution function (cdf) evaluation (at zero) of a particular z -distribution.

¹The notation reserves $\pi(\cdot)$ for the marginal posterior β as a visual queue for the quantity of primary interest. All other probability densities use $p(\cdot)$, including the joint for latent (z, λ) and all priors.

Theorem 2.1. *The (powered up) logistic function may be represented as follows.*

$$\begin{aligned} (1 + e^{-yx^\top \beta})^{-\kappa} &= \\ &\int_0^\infty \int_0^\infty \frac{1}{\sqrt{2\pi\lambda}} \exp\left\{-\frac{1}{2\lambda} \left(z - yx^\top \beta - \frac{1}{2}(1 - \kappa)\lambda\right)^2\right\} q_{1,\kappa}(\lambda) d\lambda dz. \end{aligned} \tag{9}$$

Proof. If $z \sim Z(1, \kappa, 1, yx^\top \beta)$, then $F_Z(z) = 1 - (1 + e^{z - yx^\top \beta})^{-\kappa}$, giving $1 - F_Z(0) = (1 + e^{-yx^\top \beta})^{-\kappa}$. In other words,

$$(1 + e^{-yx^\top \beta})^{-\kappa} = \int_0^\infty Z(z; 1, \kappa, 1, yx^\top \beta) dz, \tag{10}$$

establishing the outer integration, over z , in Eq. (9). Applying the representation in Eq. (5) yields the desired result. QED.

The statistical implication of this is a hierarchical model which we summarize in the following corollary.

Corollary 2.2. *The conditional distribution $p_\kappa(z_i|\beta, \lambda_i, y_i)$ and the mixing distribution $q_{1,\kappa}(\lambda_i)$ imply that the latent z_i follow*

$$p_\kappa(z_i|\beta, \lambda_i, y_i) \equiv \mathcal{N}^+\left(y_i x_i^\top \beta + \frac{1}{2}(1 - \kappa)\lambda_i, \lambda_i\right), \tag{11}$$

where \mathcal{N}^+ is the normal distribution truncated to the positive real line.

In more compact notation, $z|\beta, \lambda, y \sim \mathcal{N}_n^+((y.X)\beta + \frac{1}{2}(1 - \kappa)\lambda, \Lambda)$, where $y = (y_1, \dots, y_n)^\top$, $y.X = \text{diag}(y)X$, $\Lambda = \text{diag}(\lambda_1, \dots, \lambda_n)$, and the truncation is to the all-positive orthant. Observe that, when $\kappa = 1$, the above formulation is identical to the generative model described by HH. Given predictors x_i and regression coefficients β , generate $y_i \in \{-1, +1\}$ as $y_i = \text{sign}(z_i)$, where

$$z_i \sim \mathcal{N}(x_i^\top \beta, \lambda_i) \quad \text{and} \quad \lambda_i = \sum_{k=1}^\infty \frac{2}{(1+k)^2} \epsilon_k, \quad \epsilon_k \stackrel{\text{iid}}{\sim} \text{Exp}(1). \tag{12}$$

When $\kappa > 1$, the asymmetry of the z -distribution makes it harder to extract y_i from $y_i x_i^\top \beta + \frac{1}{2}(1 - \kappa)\lambda_i$, the mean of the truncated normal in Eq. (11). In Section 3.3, we indirectly suggest that one can interpret κy_i as a binomial response when κ is an integer.

An alternative z -representation:

Theorem 2.1 shows how components of the powered-up logistic likelihood can be represented hierarchically by the cdf of z -distributions. We therefore call that multiplicity extension to HH the *cdf representation*. However, further inspection reveals that it is

possible to eliminate an integral in Eq. (4) and thus n latent variables, and use the representation

$$\begin{aligned} (1 + e^{z_i - \mu_i})^{-\kappa} &\equiv Z(z_i; a = 0, b = \kappa, 1, \mu_i) \\ &= \int_0^\infty \frac{1}{\sqrt{2\pi\lambda_i}} \exp \left\{ -\frac{1}{2\lambda_i} \left(z_i - \mu_i + \frac{1}{2}\kappa\lambda_i \right)^2 \right\} q_{0,\kappa}(\lambda_i) d\lambda_i \end{aligned}$$

which avoids integrating over z_i . Instead, set them to zero (and $\mu_i = y_i x_i^\top \beta$) and directly obtain $(1 + e^{y_i x_i^\top \beta})^{-\kappa}$. By analogy, we call this a *pdf representation* as it involves evaluating a particular z -density function. This simple representation is problematic, however, since the Polya mixing density $q_{0,\kappa}$ is improper. In particular, note that $\psi_0 = 0$, resulting in a infinite weight in the generative formulation (8).

Fortunately, a similar representation may be generated

$$\begin{aligned} (1 + e^{-\mu})^{-\kappa} &\equiv Z(z; a, b, 1, \mu) \Big|_{z=0} \\ &= e^{a\mu} \int_0^\infty \frac{1}{\sqrt{2\pi\lambda}} \exp \left\{ -\frac{1}{2\lambda} \left(-\mu - \frac{1}{2}(a-b)\lambda \right)^2 \right\} q_{a,b}(\lambda) d\lambda \end{aligned} \quad (13)$$

which involves a proper Polya mixing density as long as $(a, b) > 0$ and $a + b = \kappa$. In Sections 3.1 and 4, we show how the extra $e^{a\mu} \equiv e^{a y_i x_i^\top \beta}$ poses no problem for efficient inference, and that $(a = \frac{1}{2}, b = \kappa - \frac{1}{2})$ works well in practice. But first, we complete the power-posterior specification with a family of regularization priors on β .

2.2 Prior regularization

Regularization is achieved via a family of priors, $p_{\kappa,\alpha}(\beta|\nu, \sigma^2)$, implementing the L_α -norm via the decomposition $p_{\kappa,\alpha}(\beta_j|\nu, \sigma^2) = \int p_{\kappa,\alpha}(\beta_j|\omega_j, \nu, \sigma^2) p_\alpha(\omega_j) d\omega_j$, following Carlin and Polson (1991) and Park and Casella (2008) in regularized (Bayesian) linear regression context. The idea is that, given $\beta_j = \frac{\nu}{\kappa^{1/\alpha}} \sigma_j \sqrt{\omega_j} \epsilon_j$ and $\epsilon_j \stackrel{\text{iid}}{\sim} \mathcal{N}(0, 1)$, small ν (i.e., heavy regularization) and large κ (i.e., heavy concentration of power-posterior density around the mode at the origin) both shrink β_j towards zero. We provide $p_\alpha(\omega)$ yielding the desired regularization penalty which, after unpacking factors from $C_{\kappa,\alpha}(\nu)$ in Eq. (3), is

$$p_{\kappa,\alpha}(\beta|\nu, \sigma^2) = \prod_{j=1}^p p_{\kappa,\alpha}(\beta_j|\nu, \sigma_j^2) \propto \nu^{-p\kappa} \exp \left(-\kappa \sum_{j=1}^p \left| \frac{\beta_j}{\nu\sigma_j} \right|^\alpha \right). \quad (14)$$

Box and Tiao (1973) provide a general discussion of (14) in the linear regression context. Some notable special cases in the recent literature on sparse logistic regression include the following: when $\nu = 1$, $\alpha = 1$, and $\sigma_j = \lambda_j$ it is the Laplace prior used in Genkin et al. (2007); when $\alpha = 2$, $\sigma_j = 1$ and $\nu = \sigma^2$ it is the Gaussian prior, and when $\alpha = 2$, $\sigma_j = 1$ and $\nu^{-1} = \lambda$ it is the Laplace prior from Krishnapuram et al. (2005).²

²The λ_j and λ variables correspond to the shrinkage parameters so named in our references. They should not be confused with the latent λ_i used in our hierarchical likelihood representation.

Inference for ν in these cases typically proceeds by CV, or by inspecting the paths of $\hat{\beta}_\nu$ solutions for varying ν . Assessing the uncertainty in estimators $\hat{\beta}_\nu$ on the final choice of $\hat{\nu}$ can pose difficulties.

Power posterior analysis offers an intriguing third option by providing the potential for tractable marginalization over prior uncertainty $\nu \sim p_{\kappa,\alpha}(\nu)$. Two particular choices in the $\alpha = 1$ case lead to efficient inference by Gibbs sampling [Section 3.1]. One option is an inverse gamma (IG) prior for ν^2 with shape $r_\kappa = \kappa(r + 1) - 1$ and scale $d_\kappa = \kappa d$, where $\kappa = 1$ yields a base case $\text{IG}(\nu^2; r, d)$ prior. The second option is IG for ν , with identical powering-up identities. It has lighter tails in ν^{-1} , thus providing more aggressive shrinkage.

The prior in Eq. (14)—for the purposes of efficient inference [Section 3]—is an adaptation of a scale mixture of normals result from West (1987) to account for κ . Specifically,

$$p_{\kappa,\alpha}(\beta_j | \nu, \sigma_j^2) = \int_{\mathbb{R}_+} \mathcal{N}\left(\beta_j; 0, \omega_j \cdot \frac{\nu^2 \sigma_j^2}{\kappa^2/\alpha}\right) p_\alpha(\omega_j) d\omega_j, \tag{15}$$

where $p_\alpha(\omega_j) \propto \omega_j^{-\frac{3}{2}} \text{St}_{\frac{\alpha}{2}}^+(\omega_j^{-1})$ and $\text{St}_{\alpha/2}^+$ is the density function of a positive stable random variable of index $\alpha/2$. In compact notation, $\beta | \sigma^2, \omega, \nu, \kappa \sim \mathcal{N}_p(0, \nu^2/\kappa^{2/\alpha} \Sigma \Omega)$ where $\Sigma = \text{diag}(\sigma_1^2, \dots, \sigma_p^2)$ and $\Omega = \text{diag}(\omega_1, \dots, \omega_p)$. An important corollary, obtained by adapting an Andrews and Mallows (1974) result, is that if $\alpha = 1$, $\omega_j \stackrel{\text{iid}}{\sim} \text{Exp}(2)$, and $\sigma_j = 1$ for $j = 1, \dots, p$ then $p_\kappa(\beta | \nu)$ is double exponential (Laplace) with a mean zero and scale ν^2/κ^2 .

3 Simulation-based logistic regression

We develop a Gibbs sampling algorithm [Section 3.2] for sampling the augmented power-posterior $p_\kappa(\beta, z, \omega, \lambda, \nu | y, \sigma^2)$, for any κ . We first derive the relevant posterior conditionals [Section 3.1], treating cdf and pdf representations in turn. When $\kappa = 1$ the marginal samples of β summarize the posterior distribution of the main parameters of interest. Obtaining the MAP or MLE requires an inhomogeneous Markov chain [Section 3.2]. Finally, we describe how a vectorized κ can facilitate efficient Bayesian binomial regression [Section 3.3].

3.1 Posterior conditionals

To begin, consider the latent z and λ variables in the cdf and pdf representations, in turn, followed by the coefficients β and corresponding regularization prior parameters (ω, ν) .

Latent likelihood parameters (z, λ)

By construction [Eq. (11) of Corollary 2.2], the posterior full conditional for the latents, $p_\kappa(z_i|\beta, \lambda_i, y_i)$, is a truncated (non-negative) normal distribution. Obtaining samples, independently for $i = 1, \dots, n$, is straightforward following the methods of Robert (1995).

Sampling from the full conditional $p_\kappa(\lambda_i|\beta, z_i, y_i)$ is complicated by the infinite sum in the expression for the prior (6), which precludes a naïve approach via truncation since certain combinations of λ_i and $b \equiv \kappa$ can give highly inaccurate, even negative, evaluations. HH derive an expression for this conditional when $\kappa = 1$ and provide a rejection sampling algorithm by squeezing (Devroye 1986). Although adaptable for general κ , we prefer a Rao–Blackwellized approach. Interchanging the order of integration in Eq. (9) suggests a corollary to Theorem 2.1 that is helpful in constructing a Metropolis–Hastings (MH) scheme for obtaining λ_i draws.

Corollary 3.1. *The following is an alternate integral representation of the logistic function*

$$\exp\left\{-\kappa \ln\left(1 + e^{-y_i x_i^\top \beta}\right)\right\} = \int_0^\infty \Phi\left(\frac{-y_i x_i^\top \beta - \frac{1}{2}(1 - \kappa)\lambda_i}{\sqrt{\lambda_i}}\right) q_{1,\kappa}(\lambda_i) d\lambda_i,$$

where Φ is the cdf of the standard normal distribution.

Proposals $\lambda'_i \sim q_{1,\kappa}(\lambda)$ can then be accepted via MH with probability $\min\{1, A_i\}$ where

$$A_i = \frac{\Phi\left\{\frac{-y_i x_i^\top \beta - \frac{1}{2}(1 - \kappa)\lambda'_i}{\sqrt{\lambda'_i}}\right\}}{\Phi\left\{\frac{-y_i x_i^\top \beta - \frac{1}{2}(1 - \kappa)\lambda_i}{\sqrt{\lambda_i}}\right\}}. \quad (16)$$

Good proposals may be obtained by truncating the sum in Eq. (8) at $K = 100$ for $\kappa = b = 1$, with improvements for larger κ . Direct sampling is also possible (e.g., Weron 1996).

Empirically, the MH acceptance rate is high (> 90%) for $\kappa = 1$ because the posterior is similar to the prior ($q_{1,1}$). Therefore the MH scheme may be preferable to the rejection/squeezing method of HH who report acceptance rates as low as 25%. Both rates decline as κ is increased, but the MH rate is still above 1% for $\kappa = 20$. A good rule of thumb is to thin $\lceil \kappa \rceil$ draws for each draw saved, which is reasonable from a computational standpoint as sampling from $q_{a,b}$ is fast. Even when thinning more than 10-fold, the MH sampler is competitive to HH/Devroye in terms of sheer speed. The MH requires two Φ evaluations, a few arithmetic operations, and two square roots. HH/Devroye, by contrast, can perform dozens (or more) expensive operations such as `pow` before the squeeze is made. Finally, drawing λ_i unconditional on z_i yields lower autocorrelation in the overall joint MCMC sampling scheme.

The pdf representation is simpler since z_i is set to zero. Proposed $\lambda'_i \sim q_{a,b}$ may be accepted or rejected via MH by exchanging a cdf for a pdf in Eq. (16) and replacing $\frac{1}{2}(1 - \kappa)$ with $\frac{1}{2}(a - b)$. Another feature that works well for the pdf representation

is an adaptation of the slice sampler of [Godsill \(2000\)](#). Given λ_i , the next sample λ'_i may be obtained via an auxiliary uniform random variable as follows. Let $\phi_i \equiv \phi\{-y_i x_i^\top \beta + \frac{1}{2}(a-b)\lambda_i\}/\sqrt{\lambda_i}$, where ϕ is the pdf of a standard normal distribution. Then sample

$$u|\lambda_i, x_i, y_i, \beta \sim U[0, \phi_i], \quad \text{followed by} \quad \lambda'_i|u, x_i, y_i, \beta \sim q_{a,b}(\lambda'_i)I_{\{\phi'_i > u\}}, \quad (17)$$

where the second step is facilitated by accept/rejects following random draws from the Polya mixing density. Although more automatic in that it does not require thinning, we show in [Section 4.2](#) that the MH scheme is faster overall. The two methods behave similarly when κ gets large, causing the rate of rejections/required thinning to increase.

Regularized regression coefficient parameters (β, ω, ν)

In the cdf representation, the multivariate normal priors for z [[Section 2.1](#)] and β [[Section 2.2](#)] combine to give $\beta|z, \omega, \lambda, \nu, \kappa \sim \mathcal{N}_p(\tilde{\beta}, V)$ with hyperparameters

$$\begin{aligned} \tilde{\beta} &= V(y.X)^\top \Lambda^{-1} \left(z - \frac{1}{2}(1-\kappa)\lambda \right), \quad \text{and} \\ V^{-1} &= (\nu/\kappa^{1/\alpha})^{-2} \Sigma^{-1} \Omega^{-1} + (y.X)^\top \Lambda^{-1} (y.X). \end{aligned}$$

Obtaining V from V^{-1} is generally $O(p^3)$, which represents a significant computational burden in the $p \gg n$ context. By employing the Sherman–Morrison–Woodbury formula (e.g., [Bernstein 2005](#), pp. 67), it is possible to use an $O(n^3)$ operation instead, which could represent significant savings. In the pdf representation a similar combination of regularization penalties and likelihoods gives an identical V^{-1} expression, but a new $\tilde{\beta} = (a - \frac{1}{2}[a-b])VX^\top y$ [see [Appendix 5](#)]. Choosing $(a = \frac{1}{2}, b = \kappa - \frac{1}{2})$ gives $\tilde{\beta} = \frac{\kappa}{2}VX^\top y$, a particularly simple expression that may be used for $\kappa > \frac{1}{2}$. It is interesting to observe that the parameters (λ, ω, ν) only enter into the conditional for β through V in the pdf representation.

The full conditional distribution of each latent ω_j is proportional to the integrand of [Eq. \(15\)](#). When $\alpha = 1$ we have the following adaptation of a standard result.

Corollary 3.2. *For $\alpha = 1$, the full conditional distribution of the reciprocal of ω_j^{-1} follows an inverse Gaussian distribution: $\omega_j^{-1}|\beta_j, \nu, \kappa \sim \text{IN}(\frac{\nu}{\kappa}|\frac{\beta_j}{\sigma_j}|^{-1}, 1)$.*

Proof. From the integrand in [Eq. \(15\)](#) with $\alpha = 1$ we have

$$p_\kappa(\omega_j|\beta_j, \nu) \propto \frac{1}{\sqrt{2\pi\omega_j}} \exp \left\{ -\frac{1}{2} \left(\frac{\kappa^2 \beta_j^2}{\nu^2 \sigma_j^2 \omega_j} + \omega_j \right) \right\} \equiv \text{GIG} \left(\omega_j; \frac{1}{2}, 1, \frac{\kappa^2 \beta_j^2}{\nu^2 \sigma_j^2} \right),$$

implying that $\omega_j^{-1} \sim \text{IN} \left(\frac{\nu}{\kappa} \left| \frac{\beta_j}{\sigma_j} \right|, 1 \right)$. [See [Appendix 5](#) for IN/GIG definitions]. *QED.*

Our IG priors for ν are both conditionally conjugate. An IG prior for ν^2 and the

representation in Eq. (15) gives

$$\nu^2 | \beta, \omega, \kappa \sim \text{IG} \left(r_\kappa + \frac{\kappa p}{2}, d_\kappa + \frac{\kappa^2}{2} \sum_{j=1}^p \frac{\beta_j^2}{\sigma_j^2 \omega_j} \right).$$

An IG prior for ν leads to efficiency gains (in addition to better tail properties) since there is no conditioning on ω . Using Eq. (14) directly then gives

$$\nu | \beta, \kappa \sim \text{IG} \left(r_\kappa + \kappa p, d_\kappa + \kappa \sum_{j=1}^p \left| \frac{\beta_j}{\sigma_j} \right| \right),$$

extending the analysis of Park and Casella (2008).

3.2 Gibbs sampling and annealing for point estimators

A full Gibbs sampling algorithm for both cdf and pdf representations is outlined in Figure 1. For compactness, variations with slice sampling for λ [in the pdf case] or a prior on ν^2 are not shown. The former requires replacing each iteration of step 2 by the method surrounding Eq. (17). The latter requires drawing

$$\nu^{2(s)} \sim \text{IG} \left(r_\kappa + \frac{\kappa p}{2}, d_\kappa + \frac{\kappa^2}{2} \sum_{j=1}^p \frac{\beta_j^{2(s)}}{\sigma_j^2 \omega_j^{(s)}} \right)$$

in step 5 and specification of $\nu^{2(0)}$ on input. Initial latent ω_j values are not required.

The samples obtained may be used to approximate expectations under the power-posterior distribution with multiplicity κ . If $\kappa = 1$ then these are samples from a well-defined posterior distribution which may be used, e.g., to approximate the posterior mean of β or provide samples from the posterior predictive distribution. Both take the full uncertainties of all parameters (including ν) into account—a feature unique to full Bayesian analysis.

Settings of $\kappa > 0$ are useful for finding other popular estimators via *simulated annealing* (SA). In our context, SA establishes an inhomogeneous Markov chain over a sequence of power-posteriors, starting with $\kappa = 1$ and then increasing according to a pre-determined schedule. Except when Gibbs sampling is possible for all κ (as for our power-posterior), it is usually difficult to ensure that the Markov chain mixes well, particularly when κ increases. A pragmatic approach starts at $\kappa \approx 1$, and systematically makes modest increases in κ until Monte Carlo variation in the power-posterior expectations of the quantities of interest is below a pre-determined threshold. Each annealing iteration is initialized with the last value $\beta^{(S)}$, $\nu^{(S)}$, $\lambda^{(S)}$ and $z^{(S)}$, from the previous iteration, thereby stitching the inhomogeneous Markov chains together. The chain for each κ must have enough iterations to establish convergence to its particular power-posterior.

Annealed procedures such as ours present an MCMC alternative to EM-style algorithms. Importantly, SA is known to converge to the global optima in certain conditions

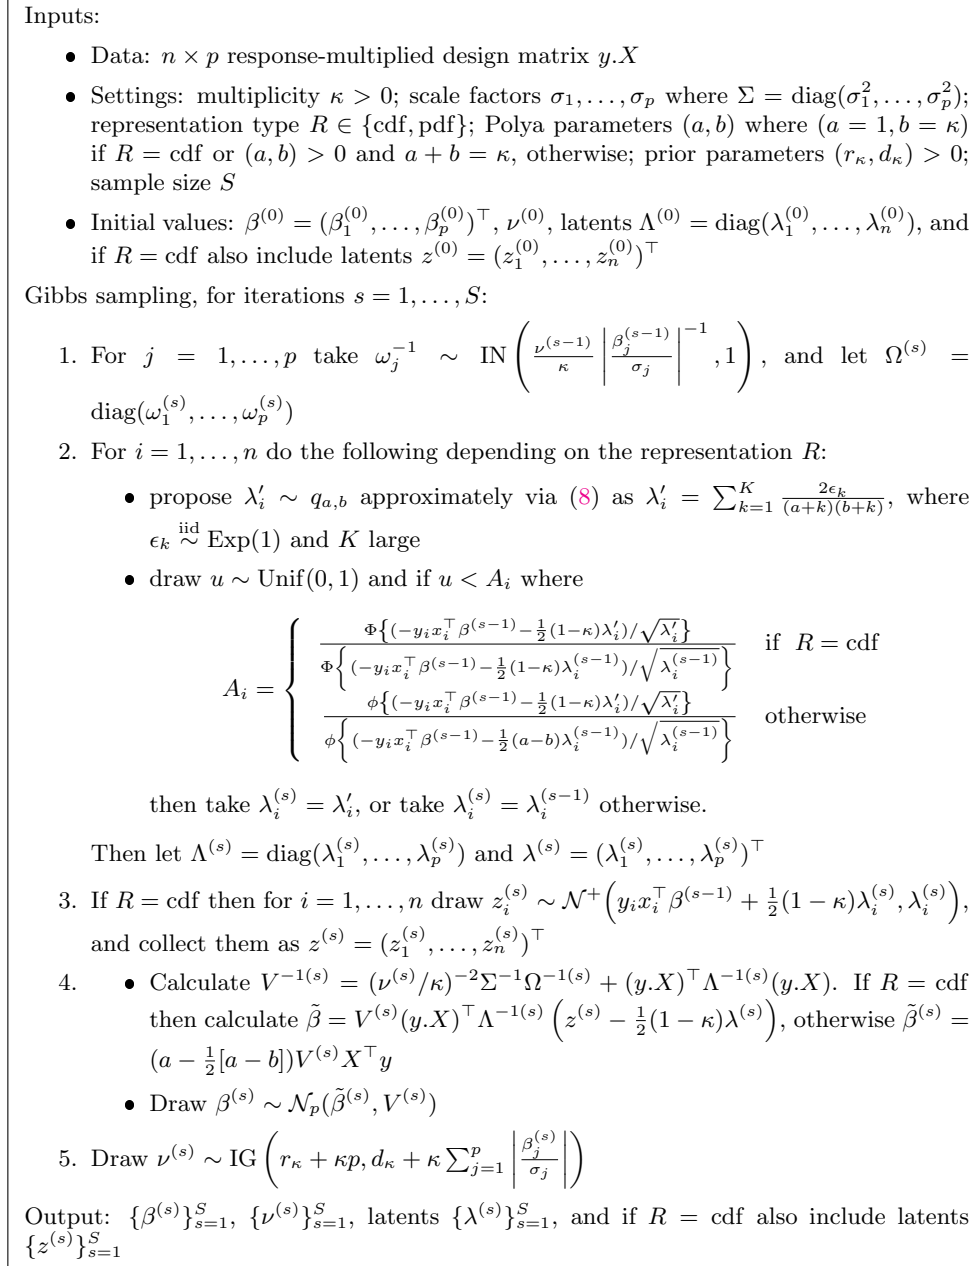


Figure 1: Pseudocode for simulation based regularized logistic regression.

(when $\kappa \rightarrow \infty$), whereas EM is only guaranteed to find a local optima. Although convergence for EM is usually quick, there are no guarantees that it will be so and indeed there are examples, particularly in high dimensional settings, where convergence can be arbitrarily slow. SA however, comes with the burden of choosing the schedule for increasing κ . We have found that for our regularized logistic regression scheme, convergence is fast and mixing so good that short schedules such as $\kappa = 1, 5, 10, 20$ are a safe default [see Section 4]. Even jumping immediately to modest κ (≈ 20), skipping $\kappa = 1$, can very often yield cheap and accurate approximations.

But perhaps the most noteworthy difference between our simulation approach and previous methods (like EM) is the myriad of options (beyond CV) for inferring ν . One option, in the classical context, is to use annealing to find the joint mode of (β, ν) . Another option is to first use samples from the posterior marginal $p(\nu|X, y)$ to estimate the posterior mean $\hat{\nu} = E\{\nu|X, y\}$, and then proceed to estimate $\hat{\beta} = E_{\kappa}\{\beta|\hat{\nu}, X, y\}$ as before. In Figure 1 this would be facilitated by inputting $\nu^{(0)} = \hat{\nu}$ and replacing step 5 with $\nu^{(s)} = \nu^{(s-1)}$. Our experience is that the former works well for small p problems, and the latter for large p . When p is large, the joint prior for (β, ν) dominates near the posterior mode of ν , which tends to zero and yields $\hat{\beta} = 0$, which is not helpful. The marginal posterior mean is far less sensitive to the regularization prior, and represents a more convenient choice for large p applications. In Section 4.4 we provide an example where the joint mode is easy to find with a few dozen predictors, whereas an interaction expanded version using thousands requires more care.

3.3 Efficient handling of binomial data

Another advantage of our approach is the extension to binomial data, where binary responses are collected repeatedly and independently, n_i times for subjects with the same covariates x_i . Contingency tables are one important example. A typical (unregularized) logistic regression model is $y_i|x_i \sim \text{Bin}(n_i, \mu_i)$, where $\mu_i = e^{\eta_i}/(1+e^{\eta_i})$ and η_i is linear in x_i . One way to situate such data within this article's regularized logistic regression framework is to *flatten* it, so that n_i components appear in the likelihood for each subject i : $\prod_{j=1}^{n_i} (1 + e^{-y_{ij}x_i^T\beta})^{\kappa}$, using the binary encoding $y_{ij} \in \{-1, 1\}$ giving $\sum_{j=1}^{n_i} |y_{ij}| = n_i$. This allows inference to proceed as described in Section 3, but it can lead to an inefficient MCMC scheme if the n_i are large due to the n_i latents required for each i . It turns out that it is possible to use only two latents for each i , echoing a feature of methods described by Frühwirth-Schnatter et al. (2009).

Observe that the component of the likelihood for subject i may be equivalently written with just two terms as $(1 + e^{-x_i^T\beta})^{\kappa y_i} (1 + e^{x_i^T\beta})^{\kappa(n_i - y_i)}$, which is proportional to the i^{th} component of a typical binomial likelihood with logit link. Hence the full likelihood, with m unique subjects, can be written as $\prod_{i=1}^m (1 + e^{-x_i^T\beta})^{\kappa_{i+}} (1 + e^{x_i^T\beta})^{\kappa_{i-}}$, where $\kappa_{i+} = \kappa y_i$ and $\kappa_{i-} = \kappa(n_i - y_i)$. This is identical to a z -distribution representation of the logistic likelihood with $2m$ terms, which may be much less than the $\sum_{i=1}^m n_i$ produced by flattening. The first m terms use response "data" $y'_i = +1$ with multiplicity parameter κ_{i+} , and the second m terms use $y'_i = -1$ with κ_{i-} . A *multiplicity imple-*

mentation is therefore facilitated by forming vectors y' and κ' , each of length $n = 2m$, and using $\prod_{i=1}^n (1 + e^{-y'_i x_i^\top \beta}) \kappa'_i$.

The MCMC scheme proceeds as in Section 3 by vectorization. For example, steps 2 and 3 for z_i and λ_i would use κ'_i instead of κ . For β in step 4 with $(a = 0.5, b = \kappa - 0.5)$ say, replace $\kappa 1_n$ with the κ' vector in the expression for $\tilde{\beta}$. Terms can be eliminated from the likelihood, thus eliminating the corresponding latents, where $\kappa'_i = 0$, as is the case when $y_i \in \{0, n_i\}$. The original, scalar, κ is used for the conditionals corresponding to the parameters of the prior. For example, the posterior conditional covariance V of β is unchanged.

4 Applications

4.1 Pima Indian data

The Pima Indian diabetes data [UCI Machine Learning Repository ([Asuncion and Newman 2007](#))] includes outcomes for diabetes tests performed on $n = 768$ women of Pima heritage with 8 real-valued predictors. Some of the predictors have many zeros, which may reasonably be interpreted as “missing” values. To remain consistent with the treatment of this data by HH, and other authors, we do not treat these values in any special way. The following analysis highlights properties of regularized estimators of $\beta = (\beta_0 \equiv \mu, \beta_1, \dots, \beta_8)$ obtained with $\alpha = 1, \sigma_j = 1$ for $j = 1, \dots, 8$, and $T = 1000$ samples from the resulting posterior (the first 100 as burn-in).




Figure 2: Power-posteriors for the Pima Indian data: $\nu = 6, \kappa \in \{1, 5, 20\}$.

Figure 2 summarizes the marginal power posterior(s) for β with boxplots. Three settings of $\kappa \in \{1, 5, 20\}$ (each panel) were used, and heavy regularization (fixing $\nu = 6$) was applied. Only the first panel ($\kappa = 1$) summarizes samples from the true posterior. The $\kappa > 1$ settings are useful for obtaining other estimators. The MLE, obtained from the `glm` command in R ([R Development Core Team 2009](#)), and the MAP as estimated from the sample(s), are also shown. Shrinkage is apparent in the divergence between the MAP and MLE values in all panels. Observe how the quartiles and outliers converge on

the MAP as κ is increased, reflecting higher precision Monte Carlo estimates of those values. Convergence is particularly rapid for the intercept term, and the two coefficients with considerable mass near zero (β_4 and β_5). These columns of X have the highest concentration of “missing” values (30% and 49% respectively), so it is not surprising that the MAP estimator excludes them.




Figure 3: Illustrating the concentration of posterior mass of β_2 and β_4 on the Pima Indian data for $\kappa \in \{1, 5, 10, 20\}$.

Figure 3 illustrates how mass concentrates on the MAP in two disparate cases for varying values of κ . For β_2 (*left* panel), which is decidedly non-zero in the power posterior(s), the convergence to the MAP (apparently around $\beta_2 = 15$) is modest. In the case of β_4 (*right* panel) the convergence to the MAP (to zero) is more rapid as κ is increased, allowing for confident variable de-selection in a way similar to the lasso for linear regression.

Finally, we consider the case where ν is also inferred by MCMC, jointly with the other parameters in the model. We use the IG prior on ν with ($r = 2, d = 0.1$), a typical default choice for linear regression (e.g., [Gramacy and Pantaleo 2010](#)). Figure 4 shows




Figure 4: Concentration of posterior mass of ν on the Pima Indian data for $\kappa \in \{1, 5, 10, 20\}$. The histogram extends to $\nu = 100$ when $\kappa = 1$, but the figure is trimmed.

the marginal posterior for ν under our settings of κ . The rate of convergence is modest, with the spread of samples in the $\kappa = 20$ case being only half that of the $\kappa = 1$ case.

4.2 Comparing cdf and pdf representations on binomial data

To illustrate the efficient handling of binomial data and, simultaneously, to compare the cdf and pdf representations, consider the following simple binomial logistic regression problem. The true linear predictor is $\eta_i = 1 + x_i^\top \beta$ where $\beta = (2, -3, 2, -4, 0, 0, 0, 0)^\top$, and the $p = 9$ dimensional x_i are uniform in $[0, 1]^p$. The responses, $y_i \in \{0, \dots, n_i\}$, are sampled with $y_i|x_i \sim \text{Bin}(n_i, \mu_i)$ where $n_i = 20$ and $\mu_i = e^{\eta_i} / (1 + e^{\eta_i})$.

	RMSE (sd)				time (sd)			
	flat		multi		flat		multi	
cdf	0.2117	(0.0602)	0.2120	(0.0606)	570.4	(37.8)	64.6	(0.82)
pdf	0.2119	(0.0613)	0.2121	(0.0602)	570.2	(28.7)	64.4	(0.99)

Table 1: Comparing RMSEs (left) and timings in seconds (right) of cdf and pdf representations and flattened/multiplicity treatments of binomial regression modeling.

Table 1 compares four different implementations of regularized binomial logistic regression ($\alpha = 1$) based on the output of 100 repeated experiments with $\sum n_i = 2000$ (i.e., $m = 100$ distinct x_i predictors). The metrics for comparison are root mean squared error (RMSE) between the true and posterior mean β s, and overall computing time of the respective MCMC samplers. In all cases, we use $T = 1000$ MCMC rounds with MH sampling of λ_i at thinning level(s) set by κ' (i.e., via κ'_i for each λ_i) as described in Section 3.1. The first 100 rounds were discarded as burn-in. The left table shows that there is no significant difference between the cdf and pdf representations, or between the flattened or multiplicity handling of binomial data, in terms of RMSE. The right table portrays a more interesting story in terms of CPU times. The many fewer latent variables needed by the multiplicity implementation leads to a much (9x) faster execution compared to flattening, with no cost in accuracy (via RMSE). In contrast, there is no speed gain to using n fewer latent z_i variables in the pdf representation.




Figure 5: Comparing MH (left) and slice (right) samplers for a λ_i in the pdf representation.

Figure 5 illuminates the differences in behavior between the MH and slice sampler

for the λ_i draws (in the pdf representation). A particularly “sticky” case, as chosen from output of the experiment, had $\kappa'_i = 14$. The *left* panel shows that many proposals from $q_{a,b}$ can be rejected under the MH ratio, even when the chain is automatically thinned. The *right* panel shows the chain obtained for the same λ_i under the slice sampler, which never saves any rejected draws. However, this comes at the expense of many rejections in the inner-loop of the slice, resulting in a slow overall sampler. The median was four, but the mean was 81 owing to a heavy right-hand tail in the distribution of rejections whose central 95% quantile spanned to 114 and maximum reached 140,600. The overall MCMC scheme based on the slice sampler took four times longer than the one based on MH. Despite the absence of rejections, the mixing in slice sampler chain (assessed visually) was no better than MH. Indeed, their effective sample size due to autocorrelation (Kass et al. 1998) was nearly identical: 223 for slice sampling, and 221 for MH. Therefore, MH is recommended for speed considerations.

4.3 A simulated $p \gg n$ experiment

We turn now to a predictive comparison of the methods of this paper, both fully Bayesian and full/joint MAP (including ν), benchmarked against other modern approaches to regularized logistic regression. Consider a synthetic data experiment like the one in Section 4.2 except: $n_i = 5$ for each of 20 unique predictors x_i , so that $\sum n_i = 100$. Three variations on the data-generating β vectors were used. In the first case $p = 9$ and $\beta = (2, -3, 0.74, -0.9, 0, 0, 0, 0)^\top$; in the second case $p = 100$, augmenting β from the first case with 91 more zeros; and in the third $p = 1000$ with 900 more zeros still. Each experiment involves a new random training design in the unit p -cube. Random testing sets are created similarly, except that $n'_i = 100$ so $\sum n'_i = 10000$. The metrics of comparison are (approximated) expected log likelihood (ELL)³ and misclassification rates.

Fully Bayesian posterior mean estimators (i.e., $\kappa = 1$) are derived via priors/MCMC exactly as described in the preceding sections with (100, 1000), (500, 1500), (1000, 2000) burn-in and total MCMC rounds in each of the cases $p = 9, 100, 1000$, respectively. MAP estimators are found by running a $\kappa = 10$ chain initialized at (β, λ, ν) -values from the $\kappa = 1$ chain used for the mean estimators, except in the $p = 1000$ case where ν was fixed to its posterior mean for reasons laid out in Section 3.2. Comparators include: the MLE obtained via the `glm` command in R; a binomial fit from the `glmnet` package (Friedman et al. 2010); and the estimator of Krishnapuram et al. (2005)⁴ [“krish” for short]. The MLE was unstable in the $p = 100$ and $p = 1000$ cases, so these results were omitted. CV was used to choose the penalty parameter in the $p = 9$ and $p = 100$ cases for `glmnet`, via `cv.glmnet`. The same procedure gave fatal errors in the $p = 1000$ case so we plugged in the estimate obtained from the corresponding $p = 100$ run for this final case. Reliably setting the penalty parameter for “krish”, via CV or otherwise, was too computationally intensive for the $p = 100, 1000$ cases so we picked a setting by hand

³Specifically, the average of $(1 - p_i) \log(1 - \hat{p}_i) + p_i \log \hat{p}_i$ over all testing locations i , where p_i and \hat{p}_i are the true and estimated predictive probabilities of the first label, respectively.

⁴This is equivalent to the Genkin et al. (2007) estimator but computationally less efficient.

using out-of-sample simulations from the $p = 9$ case.




Figure 6: Expected log likelihood (ELL) and misclassification rates in boxplot (left) and tabular (right) form. In both cases there are three sections, depending on the number of irrelevant predictors in the design matrix, wherein the same estimators are applied. The vertical dashed-red lines in the boxplots indicate the same demarkation as the horizontal lines in the tables.

The results of the Monte Carlo experiment are summarized in Figure 6 by boxplots, and numerically. The best estimators have high ELL, low miss rates, and lower variability across the 100 repetitions. The fully Bayesian and “krish” methods are the best when $p = 9$ (left-hand region of the boxplots and the top region of the tables). The former wins by ELL, having fewer low values, and the latter wins on miss rate, having more small ones. The “krish” method wins by both metrics on average, since it employs a fortuitously hand-chosen setting of the penalty parameter. The MLE is good on average, but has some extreme ELL and miss rate values. The `glmnet` and MAP estimators are positioned in between.

Distinctions in performance between the methods increase with p . See the right-hand regions of the boxplots and the bottom regions of the tables. The “krish” method suffers from high variability due to the fixed choice of the penalty parameter. The

`glmnet` variability is much lower, but there are many extreme outliers. Behavior in both $p = 100$ and 1000 cases is qualitatively similar for this estimator even though the former used CV to set the penalty parameter and the latter used the same fixed value. The MAP and fully Bayesian estimators have similar average behavior compared to other estimators, but with lower variability. Apparently, choosing the penalty parameter via the posterior offers the most stability in high dimensional settings. The fully Bayesian approach appears preferable to the MAP in all cases, but this distinction is harder to make out as p increases.

4.4 Spam data with interactions

For a similar real-data experiment, consider the Spambase data set from UCI. It contains the binary classifications of 4601 emails based on 57 attributes which are treated as predictors. An interaction-expanded version of the predictor set contains approximately 1700 predictors. We performed a Monte Carlo experiment comprised of 20 random 5-fold CV training and testing sets using both the original and expanded predictors. Estimators were fit on the 100 training sets, and validated by misclassification rate on the testing ones. The Bayes estimators used (500,1500) MCMC (burn-in, total) rounds with the original 57 attributes, and (1000,2000) with the expanded set. The MAP and `glmnet` calculations were exactly as described for the $p = 100$ case in Section 4.3 for the original predictor set, and like the $p = 1000$ case for the expanded one. And “krish” was like $p = 9$ and $p = 100$, respectively.




Figure 7: Misclassification rates in boxplot (*left*) and tabular (*right*) form. In both cases there are two sections, depending absence or presence of interaction terms in the design matrix, wherein the same estimators are applied. The vertical dashed-red line in the boxplot indicates the same demarcation as the horizontal line in the table.

The results of the experiment are summarized in Figure 7. The first thing we notice is that, in contrast with the results in Section 4.3, the performance improves as the predictor set expands since some of the interaction terms make good predictors. The

MLE is unstable, and so the regularized estimators offer an improvement even when the number of predictors is small relative to the number of instances. The Bayesian methods outperform `glmnet` across the board, and using the posterior to set the value of the regularization parameter is important in high dimensional settings. The estimator “krish” with fortuitous regularization is the best on the original predictor set, but worst on the expanded one where a revised setting of regularization could not be automated efficiently.

5 Discussion and extension

We provide a simulation-based approach to regularized logistic regression that facilitates a variety of inferential goals under a single framework. Most of the development of the methodology, and all of the applications, involved the $\alpha = 1$ case. Everything extends to the ridge prior ($\alpha = 2$), i.e., an independent normal prior for each coefficient β_j with variance $\sigma_j^2 \nu^2 / \kappa$. Then, $p_\kappa(\omega_j | \beta, \nu)$ is a point mass at $\omega_j = 1$. Thus similar conjugacy results hold for the gamma prior on ν and ν^2 .

From a computational perspective, our methods are competitive with the state-of-the-art in un-regularized (and $\kappa = 1$) contexts too. For example, we compared the efficiency of our methods to the “dRUM” MH sampler described by [Frühwirth-Schnatter and Frühwirth \(2010\)](#). This method is attractive because it is fast and easy to implement. For example, on the Pima data it takes about 32s to generate 10,000 samples from the posterior which is about 7x faster than our pdf representation, which took 230s. However, the MH acceptance rate of the dRUM method was 46% which lead to a marginal effective sample size (ESS) of 957 averaged over the nine β_j coefficients. Our pdf representation had an average ESS that was about 5x better, at 4518. So the methods work out to have similar overall efficiencies in that example. But in higher dimensions like the 57-d spam data, our Gibbs sampling approach is much more attractive. The acceptance rate for dRUM was extremely low at 0.4%, which leads to ESSs that are essentially nil. Although our pdf representation is (again) 7x slower, faster convergence due to better movement in the chain leads to reasonable ESSs around 500.

There are several extensions of our methodology that readily present themselves. For example, handling polychotomous data (i.e., > 2 classes) is straightforward. Following the setup in HH we may introduce C collections of coefficients $\beta^{(1)}, \dots, \beta^{(C)}$ for C classes with the convention that $\beta^{(C)} = 0$ so that logistic regression is recovered in the $C = 2$ case. Then, we simply work with the conditional likelihoods $L(\beta^{(j)} | y, \beta^{(-j)})$ which turn out to have exactly the form of a logistic regression likelihood for the class indicator that each $y_i = j$, independently for $i = 1, \dots, n$. If there are $n_i > 1$ trials for predictors x_i , then our algorithm for binomial logistic regression is applicable via a vectorized multiplicity parameter as described in Section 3.3. Extending the methods to ordinal responses is even easier. [Johnson and Albert \(1999, Chapter 4\)](#) describe a Bayesian probit model which may be adapted for the logit case following either HH or our cdf representation. The pdf representation may not be readily applicable because the latent z_i are useful for efficient sampling of the ordinal break points.

A further direction is to other classes of regularization priors. Implementing the Normal–Gamma extension (Griffin and Brown 2010) requires adding an extra (conjugate) parameter. A promising new approach is the *horseshoe* prior (Carvalho et al. 2010), which can be implemented with the addition of a slice sampler. Often variable selection is a primary goal of regularization, for which our methods would require further extension. For example, HH describe an approach to variable selection for logistic regression via Reversible Jump MCMC (Green 1995) which is adaptable to our framework. A similar regularized approach in a linear regression is provided by Gramacy and Pantaleo (2010). For variable selection for logistic regression using spike-and-slab priors, see Tüchler (2008).

Acknowledgments

This research was partially funded by EPSRC grant EP/D065704/1 to RBG, and was done in part while RBG was at the Statistical Laboratory, University of Cambridge. The authors would like thank Matt Taddy for interesting discussions on the efficient handling of Binomial data, extensions to Multinomial regression, and EM code for the MAP estimator(s). We are grateful to two referees and an associate editor for valuable comments.

References

- Andrews, D. and Mallows, C. (1974). “Scale Mixtures of Normal Distributions.” *Journal of the Royal Statistical Society, Series B*, 36, 99–102. 570, 573
- Asuncion, A. and Newman, D. (2007). “UCI Machine Learning Repository.” [Http://www.ics.uci.edu/~mllearn/MLRepository.html](http://www.ics.uci.edu/~mllearn/MLRepository.html). 579
- Barndorff-Nielsen, O., Kent, J., and Sorensen, M. (1982). “Normal Variance-Mean Mixtures and z -distributions.” *International Statistical Review*, 50, 145–159. 568, 570
- Bernstein, D. (2005). *Matrix Mathematics*. Princeton, NJ: Princeton University Press. 575
- Box, G. and Tiao, G. (1973). *Bayesian Inference in Statistical Analysis*. Mass: Addison Wesley. 572
- Carlin, B. P. and Polson, N. G. (1991). “Inference for Nonconjugate Bayesian Models using the Gibbs sampler.” *The Canadian Journal of Statistics*, 19, 4, 399–405. 572
- Carvalho, C., Polson, N., and Scott, J. (2010). “The horseshoe estimator for sparse signals.” *Biometrika*, 9, 2, 465–480. 586
- Devroye, L. (1986). *Non-Uniform Random Variate Generation*. Springer-Verlag. 574, 589

- Doucet, A., Godsill, S., and Robert, C. (2002). “Marginal maximum a posteriori estimation using Markov chain Monte Carlo.” *Statistics and Computing*, 21, 77–84. [569](#)
- Fahrmeir, L., Kneib, T., and Konrath, S. (2010). “Bayesian regularisation in structured additive regression: A unifying perspective on shrinkage, smoothing and predictor selection.” *Statistics and Computing*, 203–219. [567](#)
- Friedman, J. H., Hastie, T., and Tibshirani, R. (2010). “Regularization Paths for Generalized Linear Models via Coordinate Descent.” *Journal of Statistical Software*, 33, 1, 1–22. [582](#)
- Friel, N. and Pettitt, A. (2008). “Marginal likelihood estimation via power posteriors.” *Journal of the Royal Statistical Society, Series B.*, 70, 3, 589–607. [567](#), [569](#)
- Frühwirth-Schnatter, S. and Frühwirth, R. (2007). “Auxilliary Mixture Sampling with Applications to Logistic Models.” *Computational Statistics and Data Analysis*, 51, 7, 3509–3528. [567](#)
- (2010). “Data augmentation and MCMC for binary and multinomial logit models.” In *Statistical Modelling and Regression Structures Festschrift in Honour of Ludwig Fahrmeir*, eds. T. Kneib and G. Tutz, 111–132. Physica-Verlag. [567](#), [585](#)
- Frühwirth-Schnatter, S., R., Frühwirth, Held, L., and Rue, H. (2009). “Improved auxiliary mixture sampling for hierarchical models of non-Gaussian data.” *Statistics and Computing*, 19, 479–492. [567](#), [578](#)
- Genkin, A., Lewis, D., and Madigan, D. (2007). “Large-Scale Bayesian Logistic Regression for Text Categorization.” *Technometrics*, 49, 3, 291–304. [567](#), [568](#), [572](#), [582](#)
- Godsill, S. (2000). “Inference in symmetric alpha-stable noise using MCMC and the slice sampler.” In *IEEE International Conference on Acoustics, Speech and Signal Processing*, vol. VI, 3806–3809. [575](#)
- Gramacy, R. and Pantaleo, E. (2010). “Shrinkage regression for multivariate inference with missing data, and an application to portfolio balancing.” *Bayesian Analysis*, 5, 2, 237–262. [580](#), [586](#)
- Green, P. (1995). “Reversible Jump Markov Chain Monte Carlo Computation and Bayesian Model Determination.” *Biometrika*, 82, 711–732. [586](#)
- Griffin, J. E. and Brown, P. J. (2010). “Inference with Normal–Gamma prior distributions in regression problems.” *Bayesian Analysis*, 5, 1, 171–188. [586](#)
- Hans, C. (2009). “Bayesian Lasso Regression.” *Biometrika*, 96, 836–845. [568](#)
- Holmes, C. and Held, K. (2006). “Bayesian Auxilliary Variable Models for Binary and Multinomial Regression.” *Bayesian Analysis*, 1, 1, 145–168. [567](#), [568](#), [570](#)

- Jacquier, E., Johannes, M., and Polson, N. (2007). “MCMC Maximum Likelihood for Latent State Models.” *Journal of Econometrics*, 137, 615–640. 569
- Johnson, V. and Albert, J. (1999). *Ordinal Data Modeling*. Springer-Verlag. 585
- Kass, R. E., Carlin, B. P., Gelman, A., and Neal, R. M. (1998). “Markov Chain Monte Carlo in Practice: A Roundtable Discussion.” *The American Statistician*, 52, 2, 93–100. 582
- Kirkpatrick, S., Gelatt, C., and Vecchi, M. (1983). “Optimization by simulated annealing.” *Science*, 220, 671–680. 567, 569
- Krishnapuram, B., Carin, L., Figueiredo, M., and Hartemink, A. (2005). “Sparse Multinomial Logistic Regression: Fast Algorithms and Generalization Bounds.” *IEEE Pattern Analysis and Machine Intelligence*, 27, 6, 957–969. 567, 572, 582
- Madigan, D. and Ridgeway, G. (2004). “Discussion of ‘Least Angle Regression’ by B. Efron, T. Hastie, I. Johnstone, and R. Tibshirani.” *Annals of Statistics*, 32, 2, 465–469. 568
- Park, M. and Hastie, T. (2008). “Penalized Logistic Regression for Detecting Gene Interactions.” *Biostatistics*, 9, 1, 30–50. 568
- Park, T. and Casella, G. (2008). “The Bayesian Lasso.” *Journal of the American Statistical Association*, 103, 482, 681–686. 572, 576
- Pincus, M. (1968). “A Closed Form Solution of Certain Programming Problems.” *Operations Research*, 18, 1225–1228. 569
- Robert, C. (1995). “Simulation of Truncated Normal Variables.” *Statistics and Computing*, 5, 2, 121–125. 574
- R Development Core Team (2009). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria. ISBN 3-900051-07-0. 579
- Tibshirani, R. (1996). “Regression shrinkage and Selection via the Lasso.” *Journal of the Royal Statistical Society, Series B.*, 58, 1, 267–288. 567
- Tüchler, R. (2008). “Bayesian variable selection for logistic models using auxiliary mixture sampling.” *Journal of Computational and Graphical Statistics*, 17, 76–94. 586
- Weron, R. (1996). “On the Chambers-Mallows-Stuck Method for Simulating Skewed Stable Random Variables.” *Statistics and Probability Letters*, 28, 2, 165–171. 574
- West, M. (1987). “On Scale Mixtures of Normal Distributions.” *Biometrika*, 74, 3, 646–648. 573

Appendix 1 Posterior conditional for β in the pdf representation

For a particular λ , i.e., ignoring the integral in Eq. (13), we have the following expression for the likelihood in vector/matrix form.

$$\prod_{i=1}^n \left(1 + e^{-y_i x_i^\top \beta}\right) = e^{a y^\top X \beta} \exp \left\{ -\frac{1}{2} \left((y \cdot X) \beta + \frac{1}{2} (a - b) \lambda \right)^\top \Lambda^{-1} \left((y \cdot X) \beta + \frac{1}{2} (a - b) \lambda \right) \right\}$$

An expression for the posterior conditional density for β can then be obtained by multiplying by the kernel of the MVN prior given ω , provided below Eq. (15), namely: $\exp\{-\frac{1}{2}\beta^\top (\frac{\kappa^2}{\nu^2} \Sigma^{-1} \Omega^{-1} \beta)\}$. Combining the terms in the three exponents gives the following quadratic form:

$$-\frac{1}{2} \left[-2a y^\top X \beta + \left((y \cdot X) \beta + \frac{1}{2} (a - b) \lambda \right)^\top \Lambda^{-1} \left((y \cdot X) \beta + \frac{1}{2} (a - b) \lambda \right) + \beta^\top \left(\frac{\kappa^2}{\mu^2} \Sigma^{-1} \Omega^{-1} \right) \beta \right].$$

Collecting terms for β yields

$$\beta^\top \left((y \cdot X)^\top \Lambda^{-1} (y \cdot X) + \frac{\kappa^2}{\nu^2} \Sigma^{-1} \Omega^{-1} \right) \beta - (2a y^\top X - (a - b) (y \cdot X) \Lambda^{-1} \lambda) \beta.$$

Therefore we deduce that the conditional is $\mathcal{N}_p(\tilde{\beta}, V)$ where $V^{-1} = (y \cdot X)^\top \Lambda^{-1} (y \cdot X) + \frac{\kappa^2}{\nu^2} \Sigma^{-1} \Omega^{-1}$. Recognizing that $(y \cdot X) \Lambda^{-1} \lambda = X^\top y$ gives that $\tilde{\beta} = V(a - \frac{1}{2}[a - b]) X^\top y$.

Appendix 2 Generalized Inverse Gaussian distribution

The pdf of a Generalized Inverse Gaussian, $GIG(\lambda, \chi, \psi)$ is

$$g(x; \lambda, \chi, \psi) = \frac{(\psi/\chi)^{\lambda/2}}{2K_\lambda(\sqrt{\psi\chi})} x^{\lambda-1} \exp \left\{ -\frac{1}{2}(\psi x + \chi/x) \right\},$$

where K_λ is a modified Bessel function of the second kind. If $X \sim GIG(\frac{1}{2}, \chi, \psi)$ then $X^{-1} \sim IN(\mu = \sqrt{\psi/\chi}, \lambda = \psi)$ where where IN is the inverse Gaussian distribution with pdf

$$f(x; \mu, \lambda) = \sqrt{\frac{\lambda}{2\pi x^3}} \exp \left\{ -\frac{\lambda(x - \mu)^2}{2\mu^2 x} \right\}.$$

The mean and variance are $E\{x\} = \mu$ and $\text{Var}[x] = \mu^3/\lambda$. A generalized inverse Gaussian $GIG(\frac{1}{2}, \chi, \psi)$ is an inverse of an Inverse Gaussian. For simulation from GIG and IN distributions see Devroye (1986).

