You may wish to pick out 3 or 4 questions from those given below. In all of the following, assume that the design matrix $\mathbf{X} \in \mathbb{R}^{n \times p}$ has had its columns mean-centred and then scaled to have $\ell_{2}$-norm $\sqrt{n}$.

## Ridge regression

1. Consider performing ridge regression when $\mathbf{Y}=\mathbf{X} \boldsymbol{\beta}^{0}+\varepsilon-\mathbf{1} \bar{\varepsilon}$ (so the response has been centred), where $\mathbf{X} \in \mathbb{R}^{n \times p}$ has full column rank, and $\operatorname{Var}(\varepsilon)=\sigma^{2} \mathbf{I}$. Let the (thin) SVD of $\mathbf{X}$ be $\mathbf{U D V}^{T}$ and write $\mathbf{U}^{T} \mathbf{X} \boldsymbol{\beta}^{0}=\boldsymbol{\gamma}$. Show that

$$
\frac{1}{n} \mathbb{E}\left\|\mathbf{X} \boldsymbol{\beta}^{0}-\mathbf{X} \hat{\boldsymbol{\beta}}_{\lambda}^{R}\right\|_{2}^{2}=\frac{1}{n} \sum_{j=1}^{p}\left(\frac{\lambda}{\lambda+D_{j j}^{2}}\right)^{2} \gamma_{j}^{2}+\frac{\sigma^{2}}{n} \sum_{j=1}^{p} \frac{D_{j j}^{4}}{\left(\lambda+D_{j j}^{2}\right)^{2}} .
$$

Now suppose the size of the signal is $n$, so $\left\|\mathbf{X} \boldsymbol{\beta}^{0}\right\|_{2}^{2}=n$. For what $\boldsymbol{\gamma}$ is the mean squared prediction error above minimised? For what $\gamma$ is it maximised?
2. In the following, assume that forming $\mathbf{A B}$ where $\mathbf{A} \in \mathbb{R}^{a \times b}, \mathbf{B} \in \mathbb{R}^{b \times c}$ requires $O(a b c)$ computational operations, and that if $\mathbf{M} \in \mathbb{R}^{d \times d}$ is invertible, then forming $\mathbf{M}^{-1}$ requires $O\left(d^{3}\right)$ operations.
(a) Suppose we wish to apply ridge regression to data $(\mathbf{Y}, \mathbf{X}) \in \mathbb{R}^{n} \times \mathbb{R}^{n \times p}$ with $n \gg p$. A complication is that the data is split into $m$ separate datasets of size $n / m \in \mathbb{N}$,

$$
\mathbf{Y}=\left(\begin{array}{c}
\mathbf{Y}^{(1)} \\
\vdots \\
\mathbf{Y}^{(m)}
\end{array}\right) \quad \mathbf{X}=\left(\begin{array}{c}
\mathbf{X}^{(1)} \\
\vdots \\
\mathbf{X}^{(m)}
\end{array}\right)
$$

with each dataset located on a different server. Moving large amounts of data between servers is expensive. Explain how one can produce ridge estimates $\hat{\boldsymbol{\beta}}_{\lambda}^{\mathrm{R}}$ by communicating only $O\left(p^{2}\right)$ numbers from each server to some central server. What is the total order of the computation time required at each server, and at the central server for your approach?
(b) Now suppose instead that $p \gg n$ and it is instead the variables that are split across $m$ servers, so each server has only a subset of $p / m \in \mathbb{N}$ variables for each observation, and some central server stores $\mathbf{Y}$. Explain how one can obtain the fitted values $\mathbf{X} \hat{\boldsymbol{\beta}}_{\lambda}^{\mathrm{R}}$ communicating only $O\left(n^{2}\right)$ numbers from each server to the central server. What is the total order of the computation time required at each server, and at the central server for your approach?
3. Suppose we have a matrix of predictors $\mathbf{X} \in \mathbb{R}^{n \times p}$ where $p \gg n$. Explain how to obtain the fitted values of the following ridge regression using the kernel trick:

Minimise over $\boldsymbol{\beta} \in \mathbb{R}^{p}, \boldsymbol{\theta} \in \mathbb{R}^{p(p-1) / 2}, \boldsymbol{\gamma} \in \mathbb{R}^{p}$,

$$
\sum_{i=1}^{n}\left(Y_{i}-\sum_{k=1}^{p} X_{i k} \beta_{k}-\sum_{k=1}^{p} \sum_{j=1}^{k-1} X_{i k} X_{i j} \theta_{j k}-\sum_{k=1}^{p} X_{i k}^{2} \gamma_{k}\right)^{2}+\lambda_{1}\|\boldsymbol{\beta}\|_{2}^{2}+\lambda_{2}\|\boldsymbol{\theta}\|_{2}^{2}+\lambda_{3}\|\boldsymbol{\gamma}\|_{2}^{2} .
$$

Note we have indexed $\boldsymbol{\theta}$ with two numbers for convenience.

## The Lasso

4. (a) When proving results on the Lasso, we started with the so-called basic inequality that

$$
\frac{1}{2 n}\left\|\mathbf{X}\left(\boldsymbol{\beta}^{0}-\hat{\boldsymbol{\beta}}\right)\right\|_{2}^{2} \leq \frac{1}{n} \boldsymbol{\varepsilon}^{T} \mathbf{X}\left(\hat{\boldsymbol{\beta}}-\boldsymbol{\beta}^{0}\right)+\lambda\left\|\boldsymbol{\beta}^{0}\right\|_{1}-\lambda\|\hat{\boldsymbol{\beta}}\|_{1}
$$

Show that in fact we can improve this to

$$
\begin{equation*}
\frac{1}{n}\left\|\mathbf{X}\left(\boldsymbol{\beta}^{0}-\hat{\boldsymbol{\beta}}\right)\right\|_{2}^{2} \leq \frac{1}{n} \boldsymbol{\varepsilon}^{T} \mathbf{X}\left(\hat{\boldsymbol{\beta}}-\boldsymbol{\beta}^{0}\right)+\lambda\left\|\boldsymbol{\beta}^{0}\right\|_{1}-\lambda\|\hat{\boldsymbol{\beta}}\|_{1} \tag{1}
\end{equation*}
$$

Hint: Start from the KKT conditions for the Lasso.
(b) Under the assumptions of Theorem 7 on the prediction and estimation properties of the Lasso under a compatibility condition, show that, with probability at least $1-2 p^{-\left(A^{2} / 8-1\right)}$, we have

$$
\frac{1}{n}\left\|\mathbf{X}\left(\boldsymbol{\beta}^{0}-\hat{\boldsymbol{\beta}}\right)\right\|_{2}^{2} \leq \frac{9 A^{2} \log (p)}{4 \phi^{2}} \frac{\sigma^{2} s}{n}
$$

5. Consider once more the setup of Theorem 7 and its proof. Let $\hat{S}=\left\{j: \hat{\beta}_{j} \neq 0\right\}$ and set $\hat{s}=|\hat{S}|$. Show that on the event $\Omega=\left\{2\left\|\mathbf{X}^{T} \varepsilon\right\|_{\infty} / n \leq \lambda\right\}$, for any non-empty subset $B$ of $\hat{S}$, we have

$$
\begin{equation*}
\frac{1}{n} \operatorname{sgn}\left(\hat{\boldsymbol{\beta}}_{B}\right)^{T} \mathbf{X}_{B}^{T} \mathbf{X}\left(\boldsymbol{\beta}^{0}-\hat{\boldsymbol{\beta}}\right) \geq \frac{\lambda|B|}{2} \tag{2}
\end{equation*}
$$

Let $\kappa_{m}^{2}$ be the maximum eigenvalue of $\mathbf{X}_{M}^{T} \mathbf{X}_{M} / n$ over all $M \subset\{1, \ldots, p\}$ with $|M|=m$. Let

$$
m^{*}=\min \left\{m \geq 1: m>9 \kappa_{m}^{2} s / \phi^{2}\right\}
$$

with $m^{*}=\infty$ if there does not exist any $m$ satisfying the condition defining the set above. Prove that on $\Omega$, we have $\hat{s}<m^{*}$. Hint: First try to obtain an upper bound on the LHS of (2) involving $\kappa_{|B|}$ by making use of the (1) and the Cauchy-Schwarz inequality.
By considering the minimality of $m^{*}$, show furthermore that on $\Omega$, we have $\hat{s} \leq 9 \kappa_{m^{*}}^{2} s / \phi^{2}$. [In words, the number of non-zero coefficients of the Lasso is of the same order as the number of true non-zeroes.]

## Graphical modelling

7. Let $\mathbf{Z}=\left(Z_{1}, \ldots, Z_{p}\right)^{T} \in\{0,1\}^{p}$ be a binary random vector with probability mass function given by

$$
\mathbb{P}\left(Z_{1}=z_{1}, \ldots, Z_{p}=z_{p}\right)=\exp \left(\Theta_{00}+\sum_{k=1}^{p} \Theta_{0 k} z_{k}+\sum_{k=1}^{p} \sum_{j=1}^{k-1} \Theta_{j k} z_{j} z_{k}-\Phi(\mathbf{\Theta})\right)
$$

where $\exp (-\Phi(\boldsymbol{\Theta}))$ is a normalising constant. Show that

$$
\operatorname{logit}\left(\mathbb{P}\left(Z_{k}=1 \mid \mathbf{Z}_{-k}=\mathbf{z}_{-k}\right)\right)=\Theta_{0 k}+\sum_{j: j<k} \Theta_{j k} z_{j}+\sum_{j: j>k} \Theta_{k j} z_{j},
$$

where $\operatorname{logit}(q)=\log \{q /(1-q)\}$ for $q \in(0,1)$. Conclude that, for $j<k$,

$$
Z_{j} \Perp Z_{k} \mid \mathbf{Z}_{-j k} \Longleftrightarrow \Theta_{j k}=0
$$

Note that for discrete random variables we can replace the densities in our definition of conditional independence with probability mass functions (which are in any case densities with respect to counting measure). How might we go about estimating the $\Theta_{j k}$ ?

## High-dimensional inference

In the following questions, suppose there are $m$ null hypotheses being tested, $H_{1}, \ldots, H_{m}$. Let $p_{1}, \ldots, p_{m}$ be the associated $p$-values, and let $p_{(1)} \leq \cdots \leq p_{(m)}$ be the ordered $p$ values (so ( $i$ ) is the index of the $i$ th smallest $p$-value). Further let $I_{0}$ be the set of true null hypotheses.
8. Consider the closed testing procedure applied to $m$ hypotheses and let $\mathcal{R}$ be the collection of all $I \subseteq\{1, \ldots, m\}$ for which for all $J \supseteq I$, the local test $\phi_{J}=1$. Now suppose that (perhaps after having looked at the results of the $\phi_{I}$ ), we decide we want to reject a set of hypotheses indexed by $B \subseteq\{1, \ldots, m\}$. Let

$$
t_{\alpha}(B)=\max \{|I|: I \subseteq B, I \notin \mathcal{R}\} .
$$

Show that $\left\{0,1, \ldots, t_{\alpha}(B)\right\}$ gives a $1-\alpha$ confidence set for the number of false rejections in $B$. That is, show that

$$
\mathbb{P}\left(\left|B \cap I_{0}\right|>t_{\alpha}(B)\right) \leq \alpha,
$$

and that this is true no matter how $B$ is chosen. Hint: Argue by working on the event $\left\{\phi_{I_{0}}=0\right\}$. This question is based on the cherry picking procedure of Goeman and Solari [2011].
9. Suppose we have a family of intersection hypotheses $H_{I}: I \in \mathcal{I}$ that is hierarchical in the sense that for any $I, J \in \mathcal{I}$, we either have $I \cap J=\emptyset$ or $I \subseteq J$ or $J \subseteq I$. Suppose that for each $H_{I}, I \in \mathcal{I}$ we have a $p$-value $p_{I}$. Define the adjusted $p$-value of $H_{I}$ to be

$$
p_{I}^{\mathrm{adj}}=\max _{J: J \in \mathcal{I}, J \supseteq I} \frac{m}{|J|} p_{J .} .
$$

Consider the procedure Meinshausen, 2008 that rejects all hypotheses $H_{I}$ for which $p_{I}^{\text {adj }} \leq \alpha$. Let $\mathcal{I}_{0}$ be the subset of $\mathcal{I}$ consisting of true intersection hypotheses, so

$$
\mathcal{I}_{0}=\left\{I \in \mathcal{I}: I \subseteq I_{0}\right\} .
$$

Show that with probability at least $1-\alpha$, this procedure makes no false rejections, so no intersection hypothesis indexed by $\mathcal{I}_{0}$ is rejected. Hint: Consider the set $\mathcal{T}_{0} \subseteq \mathcal{I}$ of $I \in \mathcal{I}_{0}$ that are maximal, that is

$$
\mathcal{T}_{0}=\left\{I \in \mathcal{I}_{0}: \text { if } J \in \mathcal{I}_{0} \text { then } J \subseteq I \text { or } J \cap I=\emptyset\right\} .
$$

10. The Benjamini-Hochberg procedure allows us to control the FDR when the $p$-values of true null hypotheses are independent of each other, and independent of the false null hypotheses. The following variant of the method, known as the Benjamini-Yekutieli procedure Benjamini and Yekutieli, 2001 allows us to control the FDR under arbitrary dependence of the $p$-values, and works as follows. Define

$$
\gamma_{m}=1+\frac{1}{2}+\cdots+\frac{1}{m} .
$$

Let $\hat{k}=\max \left\{i: p_{(i)} \leq \alpha i /\left(m \gamma_{m}\right)\right\}$ and reject $H_{(1)}, \ldots, H_{(\hat{k})}$. First show that the FDR of this procedure satisfies

$$
\mathrm{FDR}=\sum_{i \in I_{0}} \mathbb{E}\left(\frac{1}{R} \mathbb{1}_{\left\{p_{i} \leq \alpha R /\left(m \gamma_{m}\right)\right\}} \mathbb{1}_{\{R>0\}}\right) .
$$

Now go on to prove that $\mathrm{FDR} \leq \alpha m_{0} / m \leq \alpha$. Hint: Verify that that for any $r \in \mathbb{N}$ we have

$$
\frac{1}{r}=\sum_{j=1}^{\infty} \frac{\mathbb{1}_{\{j \geq r\}}}{j(j+1)},
$$

and use this to replace $1 / R$.

## References

Y. Benjamini and D. Yekutieli. The control of the false discovery rate in multiple testing under dependency. Annals of Statistics, 29:1165-1188, 2001.
J. J. Goeman and A. Solari. Multiple testing for exploratory research. Statistical Science, pages 584-597, 2011.
N. Meinshausen. Hierarchical testing of variable importance. Biometrika, 95:265-278, 2008.

