

1 Supplementary Material

1.1 Derivation of equation (2)

The multi-information is defined as

$$I[\mathbf{Y} : \mathbf{U} : \mathbf{V}] = \left\langle \log \frac{p(\mathbf{y}, \mathbf{u}, \mathbf{v})}{p(\mathbf{y})p(\mathbf{u})p(\mathbf{v})} \right\rangle_{\mathbf{Y}, \mathbf{U}, \mathbf{V}}.$$

It satisfies the chain rule

$$I[\mathbf{Y} : \mathbf{U} : \mathbf{V}] = I[(\mathbf{Y}, \mathbf{U}) : \mathbf{V}] + I[\mathbf{Y} : \mathbf{U}].$$

Therefore,

$$\begin{aligned} I[\mathbf{Y} : \mathbf{U} : \mathbf{V}] &= I[(\mathbf{Y}, \mathbf{U}) : \mathbf{V}] + I[\mathbf{Y} : \mathbf{U}] \\ &= I[\mathbf{Y} : (\mathbf{U}, \mathbf{V})] + I[\mathbf{U} : \mathbf{V}] \\ \Leftrightarrow I[(\mathbf{Y}, \mathbf{U}) : \mathbf{V}] &= I[\mathbf{Y} : (\mathbf{U}, \mathbf{V})] + I[\mathbf{U} : \mathbf{V}] - I[\mathbf{Y} : \mathbf{U}] \\ &= I[\mathbf{Y} : \mathbf{X}] + I[\mathbf{U} : \mathbf{V}] - I[\mathbf{Y} : \mathbf{U}]. \end{aligned}$$

1.2 Kernels and their derivatives

RBF kernel The RBF kernel is given by

$$k(\mathbf{x}_i, \mathbf{x}_j) = \exp\left(-\frac{\|\mathbf{x}_i - \mathbf{x}_j\|^2}{\sigma^2}\right).$$

Its derivative w.r.t. \mathbf{x}_i is

$$\frac{\partial}{\partial \mathbf{x}_i} k(\mathbf{x}_i, \mathbf{x}_j) = k(\mathbf{x}_i, \mathbf{x}_j) \cdot -\frac{2}{\sigma^2} (\mathbf{x}_i - \mathbf{x}_j).$$

RBF tensor kernel The RBF tensor kernel is given by

$$\begin{aligned} k((\mathbf{x}_1, \mathbf{y}_1), (\mathbf{x}_2, \mathbf{y}_2)) &= \exp\left(-\frac{\|\mathbf{x}_1 \otimes \mathbf{y}_1 - \mathbf{x}_2 \otimes \mathbf{y}_2\|_2^2}{\sigma^2}\right) \\ \|\mathbf{x}_1 \otimes \mathbf{y}_1 - \mathbf{x}_2 \otimes \mathbf{y}_2\|_2^2 &= \langle \mathbf{x}_1 \otimes \mathbf{y}_1, \mathbf{x}_1 \otimes \mathbf{y}_1 \rangle - 2 \langle \mathbf{x}_1 \otimes \mathbf{y}_1, \mathbf{x}_2 \otimes \mathbf{y}_2 \rangle + \langle \mathbf{x}_2 \otimes \mathbf{y}_2, \mathbf{x}_2 \otimes \mathbf{y}_2 \rangle \\ &= \langle \mathbf{x}_1, \mathbf{x}_1 \rangle \langle \mathbf{y}_1, \mathbf{y}_1 \rangle - 2 \langle \mathbf{x}_1, \mathbf{x}_2 \rangle \langle \mathbf{y}_1, \mathbf{y}_2 \rangle + \langle \mathbf{x}_2, \mathbf{x}_2 \rangle \langle \mathbf{y}_2, \mathbf{y}_2 \rangle. \end{aligned}$$

The derivative of k w.r.t. \mathbf{x}_1 and \mathbf{y}_2 are given by

$$\begin{aligned} \frac{\partial}{\partial \mathbf{x}_1} k((\mathbf{x}_1, \mathbf{y}_1), (\mathbf{x}_2, \mathbf{y}_2)) &= k((\mathbf{x}_1, \mathbf{y}_1), (\mathbf{x}_2, \mathbf{y}_2)) \cdot -\frac{2}{\sigma^2} (\langle \mathbf{y}_1, \mathbf{y}_1 \rangle \mathbf{x}_1 - \langle \mathbf{y}_1, \mathbf{y}_2 \rangle \mathbf{x}_2) \\ \frac{\partial}{\partial \mathbf{y}_1} k((\mathbf{x}_1, \mathbf{y}_1), (\mathbf{x}_2, \mathbf{y}_2)) &= k((\mathbf{x}_1, \mathbf{y}_1), (\mathbf{x}_2, \mathbf{y}_2)) \cdot -\frac{2}{\sigma^2} (\langle \mathbf{x}_1, \mathbf{x}_1 \rangle \mathbf{y}_1 - \langle \mathbf{x}_1, \mathbf{x}_2 \rangle \mathbf{y}_2). \end{aligned}$$

1.3 Computation of J

For the regular case For HSIC, the matrix J can be computed in terms of the partial derivatives

$$\left(D_\eta^{(u)}\right)_{ij} = \left(\frac{\partial}{\partial u_{i\eta}} k((\mathbf{u}_i, \mathbf{v}_i, \mathbf{y}_i), (\mathbf{u}_j, \mathbf{v}_j, \mathbf{y}_j))\right)_{ij}$$

of the kernel with respect to the η^{th} dimension of \mathbf{u} (and analogously for \mathbf{v}) in the *first* argument, even if $i = j$.

In general, consider any function f that depends on a kernel matrix K which in turn depends on set of data points \mathbf{u}_i collected in the rows of a matrix Υ . Since K_{ij} only depends on the i^{th} and j^{th} example, the derivative $\frac{\partial f}{\partial u_{\nu\eta}}$ can be written as

$$\frac{\partial f}{\partial u_{\nu\eta}} = \sum_{i,j=1}^m \left(\frac{df}{dK}\right)_{ij} \frac{dk_{ij}}{du_{\nu\eta}} (\delta_{i\nu} + \delta_{j\nu}) \quad \text{or} \quad \frac{\partial f}{\partial \Upsilon_\eta} = \text{diag}\left(\left(\frac{\partial f}{\partial K} + \frac{\partial f}{\partial K}^\top\right) D_\eta^{(u)\top}\right), \quad (1.1)$$

where $\Upsilon_{\cdot,\eta}$ denotes the η th column of Υ . With $f = \text{tr}$ and $\frac{\partial}{\partial K_1} \text{tr}(K_1 H K_2 H) = H K_2 H$, the derivatives in $J = (J^{(u)}, J^{(v)})$ can be generically computed as a function of the derivatives of kernels $D_\eta^{(u)}$ and $D_\eta^{(v)}$:

$$\begin{aligned} J_\eta^{(u)} &= \frac{2}{(m-1)^2} \text{diag} \left(H K_2 H D_\eta^{(u)\top} \right) \\ J_\eta^{(v)} &= \frac{2}{(m-1)^2} \text{diag} \left(H K_1 H D_\eta^{(v)\top} \right), \end{aligned}$$

since $\text{tr}(K_1 H K_2 H) = \text{tr}(K_2 H K_1 H)$.

For the incomplete Cholesky decomposition When computing the derivative of $\hat{\gamma}_{hs}^2$ with the incomplete Cholesky decomposition, we need to take into account that (i) each entry in the kernel matrix might now be a function of more than a pair of data points, and we (ii) want to avoid having to compute the whole kernel matrix. In order to compute the derivative note that the approximation $\tilde{K} = LL^\top$ to K is given by

$$K \approx \tilde{K} = LL^\top = \begin{pmatrix} K_{\mathbf{ii}} & K_{\mathbf{ij}} \\ K_{\mathbf{ii}}^\top & K_{\mathbf{ij}}^\top K_{\mathbf{ii}}^{-1} K_{\mathbf{ij}} \end{pmatrix},$$

where \mathbf{i} is an index set containing the indices of the pivot elements used to compute the incomplete Cholesky decomposition and $\mathbf{j} = \{1, \dots, m\} \setminus \mathbf{i}$ is its complement [1]. Therefore,

$$\text{tr} \left(\tilde{K} \underbrace{H \tilde{K}_2 H}_{=: A^{(2)}} \right) = \text{tr} \left(K_{\mathbf{ii}} A_{\mathbf{ii}}^{(2)} \right) + \text{tr} \left(K_{\mathbf{ij}} A_{\mathbf{ji}}^{(2)} \right) + \text{tr} \left(K_{\mathbf{ji}} A_{\mathbf{ij}}^{(2)} \right) + \text{tr} \left(K_{\mathbf{ij}}^\top K_{\mathbf{ii}}^{-1} K_{\mathbf{ij}} A_{\mathbf{jj}}^{(2)} \right),$$

where indexing with the index sets \mathbf{i} and \mathbf{j} denotes the extraction of a sub-matrix of the respective matrix.

We can now take the derivatives of $\hat{\gamma}_{hs}^2$ with respect to the pivot and non-pivot elements (corresponding to the index sets \mathbf{i} and \mathbf{j} and—equivalently—to rows of J). Note that equation (1.1) becomes $\frac{\partial f}{\partial \Upsilon_{\cdot,\eta}} = \text{diag} \left(\frac{\partial f}{\partial K} D_\eta^\top \right)$ in the case of the cross-kernel matrix $K_{\mathbf{ij}}$. Using the product rule for matrix derivatives [2], this reduces the derivative of the approximate case to the one above since

$$\begin{aligned} \frac{\partial}{\partial K_{\mathbf{ji}}} \text{tr} \left(K_{\mathbf{ij}}^\top K_{\mathbf{ii}}^{-1} K_{\mathbf{ij}} A_{\mathbf{jj}}^{(2)} \right) &= K_{\mathbf{ii}}^{-1} K_{\mathbf{ij}} A_{\mathbf{jj}}^{(2)} + \left(A_{\mathbf{jj}}^{(2)} K_{\mathbf{ji}} K_{\mathbf{ii}}^{-1} \right)^\top \\ \frac{\partial}{\partial K_{\mathbf{ii}}} \text{tr} \left(K_{\mathbf{ji}} K_{\mathbf{ii}}^{-1} K_{\mathbf{ij}} A_{\mathbf{jj}}^{(2)} \right) &= -K_{\mathbf{ii}}^{-1} K_{\mathbf{ij}} A_{\mathbf{jj}}^{(2)} K_{\mathbf{ji}} K_{\mathbf{ii}}^{-1}. \end{aligned}$$

Let $K := K_1$, $A^{(2)} := H \tilde{K}_2 K$, and \mathbf{i} and \mathbf{j} the pivot and non-pivot indices of K_1 . Then the first k columns (corresponding to the features \mathbf{u}_i) of J are given by

$$\begin{aligned} J_{i\eta}^{(u)} &= \frac{2}{(m-1)^2} \left(\text{diag} \left(A_{\mathbf{ii}}^{(2)} D_{i\eta}^{(u)\top} \right) + \text{diag} \left(A_{\mathbf{ij}}^{(2)} D_{i\eta}^{(u)\top} \right) + \text{diag} \left(K_{\mathbf{ii}}^{-1} K_{\mathbf{ij}} A_{\mathbf{jj}}^{(2)} D_{i\eta}^{(u)\top} \right) - \text{diag} \left(K_{\mathbf{ii}}^{-1} K_{\mathbf{ij}} A_{\mathbf{jj}}^{(2)} K_{\mathbf{ij}}^\top K_{\mathbf{ii}}^{-1} D_{i\eta}^{(u)\top} \right) \right) \\ J_{j\eta}^{(u)} &= \frac{2}{(m-1)^2} \left(\text{diag} \left(A_{\mathbf{ji}}^{(2)} D_{j\eta}^{(u)\top} \right) + \text{diag} \left(A_{\mathbf{jj}}^{(2)} K_{\mathbf{ij}}^\top K_{\mathbf{ii}}^{-1} D_{j\eta}^{(u)\top} \right) \right). \end{aligned}$$

Let $K := K_2$, $A^{(1)} := H \tilde{K}_1 H$, and \mathbf{i} and \mathbf{j} the pivot and non-pivot indices of K_2 . Then the last $n - k$ columns (corresponding to the features \mathbf{v}_i) of J are given by

$$\begin{aligned} J_{i\eta}^{(v)} &= \frac{2}{(m-1)^2} \left(\text{diag} \left(A_{\mathbf{ii}}^{(1)} D_{i\eta}^{(v)\top} \right) + \text{diag} \left(A_{\mathbf{ij}}^{(1)} D_{i\eta}^{(v)\top} \right) + \text{diag} \left(K_{\mathbf{ii}}^{-1} K_{\mathbf{ij}} A_{\mathbf{jj}}^{(1)} D_{i\eta}^{(v)\top} \right) - \text{diag} \left(K_{\mathbf{ii}}^{-1} K_{\mathbf{ij}} A_{\mathbf{jj}}^{(1)} K_{\mathbf{ij}}^\top K_{\mathbf{ii}}^{-1} D_{i\eta}^{(v)\top} \right) \right) \\ J_{j\eta}^{(v)} &= \frac{2}{(m-1)^2} \left(\text{diag} \left(A_{\mathbf{ji}}^{(1)} D_{j\eta}^{(v)\top} \right) + \text{diag} \left(A_{\mathbf{jj}}^{(1)} K_{\mathbf{ij}}^\top K_{\mathbf{ii}}^{-1} D_{j\eta}^{(v)\top} \right) \right). \end{aligned}$$

References

- [1] F. R. Bach and M. I. Jordan. Predictive low-rank decomposition for kernel methods. In *Proceedings of the 22nd international conference on Machine learning - ICML '05*, pages 33–40, New York, New York, USA, 2005. ACM Press.
- [2] T. P. Minka. Old and New Matrix Algebra Useful for Statistics. *MIT Media Lab Note*, pages 1–19, 2000.