

# The Role of GMRFs in Statistical Modeling and Scalable Bayesian Inference (part II)



King Abdullah University of  
Science and Technology



جامعة الملك عبد الله  
للعلوم والتقنية

Håvard Rue  
King Abdullah University of Science and Technology  
Saudi Arabia

May 2026



- **Continuous time**
- Continuous space
- Where do GMRFs come now?
- Generalisation: how simple ideas can solve hard problems
- ...and the solution is still a GMRF
- (Approximate) Bayesian inference
- Variational Bayes
- Numerics for sparse (precision-)matrices
- New sparse solver: *sTiles*



- Continuous time
- Continuous space
- Where do GMRFs come now?
- Generalisation: how simple ideas can solve hard problems
- ...and the solution is still a GMRF
- (Approximate) Bayesian inference
- Variational Bayes
- Numerics for sparse (precision-)matrices
- New sparse solver: *sTiles*



- Continuous time
- Continuous space
- Where do GMRFs come now?
  - Generalisation: how simple ideas can solve hard problems
  - ...and the solution is still a GMRF
  - (Approximate) Bayesian inference
  - Variational Bayes
  - Numerics for sparse (precision-)matrices
  - New sparse solver: *sTiles*



- Continuous time
- Continuous space
- Where do GMRFs come now?
- Generalisation: how simple ideas can solve hard problems
- ...and the solution is still a GMRF
- (Approximate) Bayesian inference
- Variational Bayes
- Numerics for sparse (precision-)matrices
- New sparse solver: *sTiles*



- Continuous time
- Continuous space
- Where do GMRFs come now?
- Generalisation: how simple ideas can solve hard problems
- ...and the solution is still a GMRF
- (Approximate) Bayesian inference
- Variational Bayes
- Numerics for sparse (precision-)matrices
- New sparse solver: *sTiles*



- Continuous time
- Continuous space
- Where do GMRFs come now?
- Generalisation: how simple ideas can solve hard problems
- ...and the solution is still a GMRF
- (Approximate) Bayesian inference
  - Variational Bayes
  - Numerics for sparse (precision-)matrices
  - New sparse solver: *sTiles*



- Continuous time
- Continuous space
- Where do GMRFs come now?
- Generalisation: how simple ideas can solve hard problems
- ...and the solution is still a GMRF
- (Approximate) Bayesian inference
- Variational Bayes
- Numerics for sparse (precision-)matrices
- New sparse solver: *sTiles*



- Continuous time
- Continuous space
- Where do GMRFs come now?
- Generalisation: how simple ideas can solve hard problems
- ...and the solution is still a GMRF
- (Approximate) Bayesian inference
- Variational Bayes
- Numerics for sparse (precision-)matrices
- New sparse solver: *sTiles*



- Continuous time
- Continuous space
- Where do GMRFs come now?
- Generalisation: how simple ideas can solve hard problems
- ...and the solution is still a GMRF
- (Approximate) Bayesian inference
- Variational Bayes
- Numerics for sparse (precision-)matrices
- New sparse solver: *sTiles*



# Example: Markov processes in time

- Auto-regressive process of order 1


$$x_t = \phi x_{t-1} + \epsilon_t, \quad 0 \leq \phi < 1, \quad t = 1, \dots, n$$

with Gaussian noise  $\epsilon_t$

- $Q$  is tridiagonal
- $C$  is  $\phi^{-|i-j|}$  (exp-correlation function)
- Cont.time: Ornstein-Uhlenbeck process

$$dx_t = -\theta x_t dt + \sigma dW_t,$$

Wiener process  $W_t$







# Example: Markov processes in time

- Auto-regressive process of order 1

$$x_t = \phi x_{t-1} + \epsilon_t, \quad 0 \leq \phi < 1, \quad t = 1, \dots, n$$

with Gaussian noise  $\epsilon_t$

- **Q** is tridiagonal
- **C** is  $\phi^{-|i-j|}$  (exp-correlation function)
- Cont.time: Ornstein-Uhlenbeck process

$$\begin{pmatrix} 1 & \phi & \phi^2 & \phi^3 & \phi^4 & \phi^5 & \phi^6 \\ \phi & 1 & \phi & \phi^2 & \phi^3 & \phi^4 & \phi^5 \\ \phi^2 & \phi & 1 & \phi & \phi^2 & \phi^3 & \phi^4 \\ \phi^3 & \phi^2 & \phi & 1 & \phi & \phi^2 & \phi^3 \\ \phi^4 & \phi^3 & \phi^2 & \phi & 1 & \phi & \phi^2 \\ \phi^5 & \phi^4 & \phi^3 & \phi^2 & \phi & 1 & \phi \\ \phi^6 & \phi^5 & \phi^4 & \phi^3 & \phi^2 & \phi & 1 \end{pmatrix}$$

$$dx_t = -\theta x_t dt + \sigma dW_t,$$

Wiener process  $W_t$



- Auto-regressive process of order 1


$$x_t = \phi x_{t-1} + \epsilon_t, \quad 0 \leq \phi < 1, \quad t = 1, \dots, n$$

with Gaussian noise  $\epsilon_t$

- $\mathbf{Q}$  is tridiagonal
- $\mathbf{C}$  is  $\phi^{-|i-j|}$  (exp-correlation function)
- Cont.time: Ornstein-Uhlenbeck process

$$dx_t = -\theta x_t dt + \sigma dW_t,$$

Wiener process  $W_t$





$$E(x_t | \{x_s, s \leq 0\}) = x_0 \exp(-\theta t)$$

$$\text{Var}(x_t | \{x_s, s \leq 0\}) = \frac{\sigma^2}{2\theta} (1 - \exp(-2\theta t))$$

With timesteps equal to 1, then the OU-process equal to AR1 with

$$\phi = \exp(-\theta)$$

and

$$\text{Var}(\epsilon_t) = \frac{\sigma^2}{2\theta} (1 - \exp(-2\theta))$$



Informally, we have

$$x'_t = -\theta x_t + \epsilon_t$$

It should be no surprise, that we can do higher order with

$$\dots + a_2 x''_t + a_1 x'_t + a_0 x_t = \epsilon_t$$

or

$$\sum_{k=0}^p a_k D^{(k)} x_t = DW_t$$



$$\begin{pmatrix} x \\ x' \\ x'' \\ \vdots \end{pmatrix}_t \text{ depends on } \begin{pmatrix} x \\ x' \\ x'' \\ \vdots \end{pmatrix}_0 + \text{noise}$$



where the state vector  $\mathbf{X}(t) = (X_0(t), \dots, X_{p-1}(t))^T$  satisfies the Itô equation,

$$(2.3) \quad d\mathbf{X}(t) = A\mathbf{X}(t)dt + \mathbf{e}dW(t),$$

with

$$A = \begin{bmatrix} 0 & 1 & 0 & \cdots & 0 \\ 0 & 0 & 1 & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \cdots & 1 \\ -a_p & -a_{p-1} & -a_{p-2} & \cdots & -a_1 \end{bmatrix} \quad \text{and} \quad \mathbf{e} = \begin{bmatrix} 0 \\ 0 \\ \vdots \\ 0 \\ 1 \end{bmatrix}.$$

So we're back at the state-space formulation



$$\mathbf{x}_{t+1} = \mathbf{A}\mathbf{x}_t + \epsilon_t$$

- Kalman recursions
- $\text{Var}(\epsilon_t)$  is singular, so implementation is annoying and wasteful
- Markov, as  $\mathbf{x}_{t+1}$  only depends on  $\mathbf{x}_t$
- In continuous time, this include also its derivatives
- Again, Markov in a extended space



# Matérn fields and the reference SPDE

The solution of this SPDE (Whittle, 1954/1963)

$$(\kappa^2 - \Delta)^{\alpha/2} \mathbf{x}(\mathbf{s}) = \epsilon(\mathbf{s}), \quad \alpha = \nu + \text{dim}/2$$

is a Gaussian field with Matérn Covariance function.

The SPDE is harder to work with, but knowing how, then earlier hard problems becomes easy!

- manifolds
- non-stationary
- computational properties
- +++





# Matérn fields and the reference SPDE

The solution of this SPDE (Whittle, 1954/1963)

$$(\kappa^2 - \Delta)^{\alpha/2} \mathbf{x}(\mathbf{s}) = \epsilon(\mathbf{s}), \quad \alpha = \nu + \text{dim}/2$$

is a Gaussian field with Matérn Covariance function.

The SPDE is harder to work with, but knowing how, then earlier hard problems becomes easy!

- manifolds
- non-stationary
- computational properties
- +++





# Matérn fields and the reference SPDE

The solution of this SPDE (Whittle, 1954/1963)

$$(\kappa^2 - \Delta)^{\alpha/2} \mathbf{x}(\mathbf{s}) = \epsilon(\mathbf{s}), \quad \alpha = \nu + \text{dim}/2$$

is a Gaussian field with Matérn Covariance function.

The SPDE is harder to work with, but knowing how, then earlier hard problems becomes easy!

- manifolds
- non-stationary
- computational properties
- +++





# Matérn fields and the reference SPDE

The solution of this SPDE (Whittle, 1954/1963)

$$(\kappa^2 - \Delta)^{\alpha/2} \mathbf{x}(\mathbf{s}) = \epsilon(\mathbf{s}), \quad \alpha = \nu + \text{dim}/2$$

is a Gaussian field with Matérn Covariance function.

The SPDE is harder to work with, but knowing how, then earlier hard problems becomes easy!

- manifolds
- non-stationary
- computational properties
- +++





# Matérn fields and the reference SPDE

The solution of this SPDE (Whittle, 1954/1963)

$$(\kappa^2 - \Delta)^{\alpha/2} \mathbf{x}(\mathbf{s}) = \epsilon(\mathbf{s}), \quad \alpha = \nu + \text{dim}/2$$

is a Gaussian field with Matérn Covariance function.


The SPDE is harder to work with, but knowing how, then earlier hard problems becomes easy!

- manifolds
- non-stationary
- computational properties
- +++






# Piece-wise linear representations






# Piece-wise linear representations



$$x(\mathbf{u}) = \sum_k \psi_k(\mathbf{u}) x_k$$





# How to “work” with the SPDE?

Hilbert space representation/finite element method


$$\mathbf{x}(\mathbf{u}) = \sum_{k=1}^N \psi_k(\mathbf{u}) w_k$$

for basis-functions  $\{\psi_k\}$  and (Gaussian) weights  $\{w_k\}$






# Piecewise linear representations






# Piecewise linear representations



$$x(\mathbf{u}) = \sum_k \psi_k(\mathbf{u}) x_k$$





# (Stochastic) Weak solution

$$(\kappa^2 - \Delta)^{\alpha/2} \mathbf{x}(\mathbf{s}) = \boldsymbol{\epsilon}(\mathbf{s})$$

(Stochastic) Weak solution

$$\left\{ \langle \phi_k, (\kappa^2 - \Delta)^{\alpha/2} \mathbf{x} \rangle \right\}_k \stackrel{D}{=} \left\{ \langle \phi_k, \boldsymbol{\epsilon} \rangle \right\}_k$$

for all test functions  $\{\phi_k\}_k$ .



# Results for $\alpha = 1, 2, \dots$ (I)

$$\alpha = 1 : \quad \phi_k = (\kappa^2 - \Delta)^{1/2} \psi_k$$

$$\alpha = 2 : \quad \phi_k = \psi_k$$

Define matrices **C**, **G** and **K**

$$C_{ij} = \langle \phi_i, \phi_j \rangle \quad i \neq j$$

$$G_{ij} = \langle \nabla \phi_i, \nabla \phi_j \rangle$$

$$\mathbf{K} = \kappa^2 \mathbf{C} + \mathbf{G}$$



## Results for $\alpha = 1, 2, \dots$ (II)

The weights are Gaussian with precision matrix  $\mathbf{Q}_{\alpha, \kappa}$

$$\mathbf{Q}_{1, \kappa} = \kappa^2 \mathbf{C} + \mathbf{G}$$

$$\mathbf{Q}_{2, \kappa} = \mathbf{K} \mathbf{C}^{-1} \mathbf{K}$$

$$\vdots$$

$$\mathbf{Q}_{\alpha, \kappa} = \mathbf{K} \mathbf{C}^{-1} \mathbf{Q}_{\alpha-2, \kappa} \mathbf{C}^{-1} \mathbf{K}$$


Replace  $\mathbf{C}$  with a appropriate diagonal matrix (“mass lumping”): This is OK



# Non-integer values of $\alpha$

## Non-integer values of $\alpha$

- Not Markov
- Similar to ‘long memory’ from Part I
- Can be approach in the same way: as a weighted sum of integer- $\alpha$  fields




[github.com/davidbolin/rSPDE](https://github.com/davidbolin/rSPDE)



# Non-integer values of $\alpha$

## Non-integer values of $\alpha$

- Not Markov
- Similar to ‘long memory’ from Part I
- Can be approach in the same way: as a weighted sum of integer- $\alpha$  fields




[github.com/davidbolin/rSPDE](https://github.com/davidbolin/rSPDE)



# Non-integer values of $\alpha$

## Non-integer values of $\alpha$

- Not Markov
- Similar to ‘long memory’ from Part I
- Can be approach in the same way: as a weighted sum of integer- $\alpha$  fields



[github.com/davidbolin/rSPDE](https://github.com/davidbolin/rSPDE)



# Markov properties in $\text{dim} \geq 2$

- Similar to time-case
- Iff the Spectrum is of the form

$$S(\mathbf{w}) \propto \frac{1}{\text{even.poly}(\mathbf{w})}$$


- It is intuitive

$$\sum \text{part.deriv}(\mathbf{x}(\mathbf{s})) = \text{noise}$$

then numerical solution is  $\mathbf{Ax} = \mathbf{w}$  for some **local A**, and

$$\text{Prec}(\mathbf{x}) \propto \mathbf{AA}^T$$

which is local as **A** is local.





## Markov properties in $\text{dim} \geq 2$

- Similar to time-case
- Iff the Spectrum is of the form

$$S(\mathbf{w}) \propto \frac{1}{\text{even.poly}(\mathbf{w})}$$


- It is intuitive

$$\sum \text{part.deriv}(\mathbf{x}(s)) = \text{noise}$$

then numerical solution is  $\mathbf{Ax} = \mathbf{w}$  for some **local A**, and

$$\text{Prec}(\mathbf{x}) \propto \mathbf{AA}^T$$

which is local as **A** is local.





## Markov properties in $\dim \geq 2$

- Similar to time-case
- Iff the Spectrum is of the form

$$S(\mathbf{w}) \propto \frac{1}{\text{even.poly}(\mathbf{w})}$$


- It is intuitive

$$\sum \text{part.deriv}(\mathbf{x}(\mathbf{s})) = \text{noise}$$

then numerical solution is  $\mathbf{Ax} = \mathbf{w}$  for some **local**  $\mathbf{A}$ , and

$$\text{Prec}(\mathbf{x}) \propto \mathbf{AA}^T$$

which is local as  $\mathbf{A}$  is local.





# Extensions

- space-time
- non-stationary
- sphere
- and mix thereof

Spatial mesh with 1251 nodes, and the 13581 stations.





# Extensions

- space-time
- non-stationary
- sphere
- and mix thereof

Spatial mesh with 1251 nodes, and the 13581 stations.





# Extensions

- space-time
- non-stationary
- sphere
- and mix thereof

Spatial mesh with 1251 nodes, and the 13581 stations.





# Extensions

- space-time
- non-stationary
- sphere
- and mix thereof

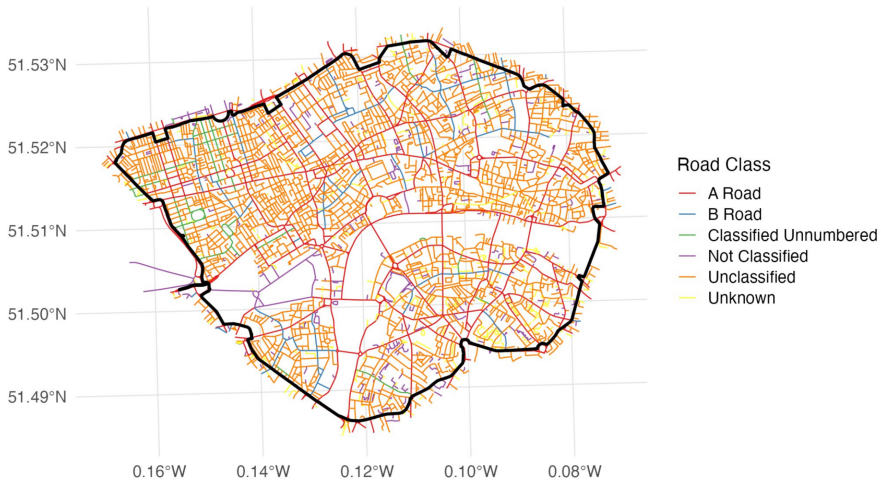
Spatial mesh with 1251 nodes, and the 13581 stations.





# Extensions

## Road Network by Classification



Source: OS Open Roads



# Implicite, not explicite!





# Extensions

MetricGraph 1.5.1.9000 [Get started](#) [Functions](#) [Vignettes](#) [Discussions](#) [Changelog](#)

## MetricGraph



CRAN 1.6.0 - 2026-05-06 [R-CMD-check](#) **falling** [R-CMD-check-windows](#) **falling**

**MetricGraph** is an R package used for working with data and random fields on metric graphs, such as street or river networks. The main functionality is contained in the `metric_graph` class, which is used for specifying metric graphs, adding data to them, visualization, and other basic functions that are needed for working with data and random fields on metric graphs. The package also implements various Gaussian fields on metric graphs, and in particular the Whittle–Matérn fields introduced in the references below.

Research group of Prof David Bolin, KAUST



# Transparent barrier





# Transparent barrier





Figure 1: Study area in the Red Sea for Dugong conservation




# Transparent barrier






# Transparent barrier









# Transparent barrier









# Transparent barrier









# Transparent barrier









# Transparent barrier









# Transparent barrier





# Transparent barrier






# Bayesian inference: Latent Gaussian models (LGMs)

**The** most important class of models...

- Conditionally independent data  $y_i | \eta_i, \theta$
- Latent

$$\eta_i = \sum_j \beta_j c_{ij} + \sum_k f_{ki}$$

- Hyperparameters  $\theta$  (latent and/or likelihood)
- Each component is a GMRF (or a “small” dense)






# Bayesian inference: Latent Gaussian models (LGMs)

**The** most important class of models...

- Conditionally independent data  $y_i | \eta_i, \theta$
- Latent

$$\eta_i = \sum_j \beta_j c_{ij} + \sum_k f_{ki}$$

- Hyperparameters  $\theta$  (latent and/or likelihood)
- Each component is a GMRF (or a “small” dense)






# Bayesian inference: Latent Gaussian models (LGMs)

**The** most important class of models...

- Conditionally independent data  $y_i | \eta_i, \theta$
- Latent

$$\eta_i = \sum_j \beta_j c_{ij} + \sum_k f_{ki}$$

- Hyperparameters  $\theta$  (latent and/or likelihood)
- Each component is a GMRF (or a “small” dense)






# Bayesian inference: Latent Gaussian models (LGMs)

**The** most important class of models...

- Conditionally independent data  $y_i | \eta_i, \theta$
- Latent

$$\eta_i = \sum_j \beta_j c_{ij} + \sum_k f_{ki}$$


- Hyperparameters  $\theta$  (latent and/or likelihood)
- Each component is a GMRF (or a “small” dense)





# Examples


inla\_graph 8rtKSK  
n = 7,860 nnz = 58,670 avg deg = 7.4






# Examples

`inla_graph_ayaLRw`  
n = 69,224    nnz = 267,944    avg deg = 3.9






# Examples






# Examples






# Examples






# Examples





# Examples


inla\_graph\_net814381  
n = 81,438    nnz = 297,626    avg deg = 3.7






# Examples

`inla_graph_spacetime`  
n = 24,931    nnz = 1,263,615    avg deg = 50.7





# Examples





# Inference: simple ideas

$$\pi(\boldsymbol{\theta}|\mathbf{y}) \propto \int \pi(\mathbf{y}|\mathbf{x}, \boldsymbol{\theta}) \pi(\mathbf{x}|\boldsymbol{\theta}) d\mathbf{x}$$

$$\pi(x_i|\mathbf{y}) \propto \iint \pi(\mathbf{x}|\boldsymbol{\theta})\pi(\boldsymbol{\theta}|\mathbf{y}) d\mathbf{x}_{-i} d\boldsymbol{\theta}$$

NORGES TEKNISK-NATURVITENSKAPELIGE  
UNIVERSITET

**Implementing Approximate Bayesian Inference using Integrated Nested  
Laplace Approximation: a manual for the InLa program**

by

Sara Martino & Håvard Rue

PREPRINT  
STATISTICS NO. 2/2008



# Poisson example

- Count-data  $y_j \sim \text{Poisson}(\exp(x_j))$
- Latent GMRF  $\mathbf{x} \sim \mathcal{N}(\mathbf{0}, \mathbf{Q})$

Need to approximate

$$\pi(\mathbf{x}|\mathbf{y}) \propto \exp\left(-\frac{1}{2}\mathbf{x}^T\mathbf{Q}\mathbf{x} + \sum_j y_j x_j - \exp(x_j)\right)$$

with something “manageable”, and also approximate marginals

$$\pi(x_i|\mathbf{y}) \propto \int \exp\left(-\frac{1}{2}\mathbf{x}^T\mathbf{Q}\mathbf{x} + \sum_j y_j x_j - \exp(x_j)\right) d\mathbf{x}_{-i}$$



# Poisson example

- Count-data  $y_j \sim \text{Poisson}(\exp(x_j))$
- Latent GMRF  $\mathbf{x} \sim \mathcal{N}(\mathbf{0}, \mathbf{Q})$

Need to approximate

$$\pi(\mathbf{x}|\mathbf{y}) \propto \exp\left(-\frac{1}{2}\mathbf{x}^T\mathbf{Q}\mathbf{x} + \sum_j y_j x_j - \exp(x_j)\right)$$

with something “manageable”, and also approximate marginals

$$\pi(x_i|\mathbf{y}) \propto \int \exp\left(-\frac{1}{2}\mathbf{x}^T\mathbf{Q}\mathbf{x} + \sum_j y_j x_j - \exp(x_j)\right) d\mathbf{x}_{-i}$$



# Poisson example

- Count-data  $y_j \sim \text{Poisson}(\exp(x_j))$
- Latent GMRF  $\mathbf{x} \sim \mathcal{N}(\mathbf{0}, \mathbf{Q})$

Need to approximate

$$\pi(\mathbf{x}|\mathbf{y}) \propto \exp\left(-\frac{1}{2}\mathbf{x}^T \mathbf{Q} \mathbf{x} + \sum_j y_j x_j - \exp(x_j)\right)$$

with something “manageable”, and also approximate marginals

$$\pi(x_i|\mathbf{y}) \propto \int \exp\left(-\frac{1}{2}\mathbf{x}^T \mathbf{Q} \mathbf{x} + \sum_j y_j x_j - \exp(x_j)\right) d\mathbf{x}_{-i}$$



# Poisson example

- Count-data  $y_j \sim \text{Poisson}(\exp(x_j))$
- Latent GMRF  $\mathbf{x} \sim \mathcal{N}(\mathbf{0}, \mathbf{Q})$

Need to approximate

$$\pi(\mathbf{x}|\mathbf{y}) \propto \exp\left(-\frac{1}{2}\mathbf{x}^T\mathbf{Q}\mathbf{x} + \sum_j y_j x_j - \exp(x_j)\right)$$

with something “manageable”, and also approximate marginals

$$\pi(x_i|\mathbf{y}) \propto \int \exp\left(-\frac{1}{2}\mathbf{x}^T\mathbf{Q}\mathbf{x} + \sum_j y_j x_j - \exp(x_j)\right) d\mathbf{x}_{-i}$$



## 2nd order approximations

The starting point is that

$$\exp\left(-\frac{1}{2}\mathbf{x}^T\mathbf{Q}\mathbf{x} + \sum_j y_j x_j - \exp(x_j)\right)$$

$$\approx \text{constant} \times \exp\left(-\frac{1}{2}(\mathbf{x} - \boldsymbol{\mu})^T(\mathbf{Q}_{\text{prior}} + \mathbf{Q}_{\text{loglike}})(\mathbf{x} - \boldsymbol{\mu})\right)$$

- We need many of those
- Constant graph; that is good
- We can do better than Taylor-expansion



## 2nd order approximations

The starting point is that

$$\exp\left(-\frac{1}{2}\mathbf{x}^T\mathbf{Q}\mathbf{x} + \sum_j y_j x_j - \exp(x_j)\right)$$

$$\approx \text{constant} \times \exp\left(-\frac{1}{2}(\mathbf{x} - \boldsymbol{\mu})^T(\mathbf{Q}_{\text{prior}} + \mathbf{Q}_{\text{loglike}})(\mathbf{x} - \boldsymbol{\mu})\right)$$

- We need many of those
- Constant graph; that is good
- We can do better than Taylor-expansion



## 2nd order approximations

The starting point is that

$$\exp\left(-\frac{1}{2}\mathbf{x}^T\mathbf{Q}\mathbf{x} + \sum_j y_j x_j - \exp(x_j)\right)$$

$$\approx \text{constant} \times \exp\left(-\frac{1}{2}(\mathbf{x} - \boldsymbol{\mu})^T(\mathbf{Q}_{\text{prior}} + \mathbf{Q}_{\text{loglike}})(\mathbf{x} - \boldsymbol{\mu})\right)$$

- We need many of those
- Constant graph; that is good
- We can do better than Taylor-expansion



## 2nd order approximations

The starting point is that

$$\begin{aligned} \exp\left(-\frac{1}{2}\mathbf{x}^T\mathbf{Q}\mathbf{x} + \sum_j y_j x_j - \exp(x_j)\right) \\ \approx \text{constant} \times \exp\left(-\frac{1}{2}(\mathbf{x} - \boldsymbol{\mu})^T(\mathbf{Q}_{\text{prior}} + \mathbf{Q}_{\text{loglike}})(\mathbf{x} - \boldsymbol{\mu})\right) \end{aligned}$$

- We need many of those
- Constant graph; that is good
- We can do better than Taylor-expansion



# Bayes thm from optimal information processing

From the *American Statistician*, 1988 (with discussion)

## Optimal Information Processing and Bayes's Theorem

ARNOLD ZELLNER\*

In this article statistical inference is viewed as information processing involving input information and output information. After introducing information measures for the input and output information, an information criterion functional is formulated and optimized to obtain an optimal information processing rule (IPR). For the particular information measures and criterion functional adopted, it is shown that Bayes's theorem is the optimal IPR. This optimal IPR is shown to be 100% efficient in the sense that its use leads to the output information being exactly equal to the given input information. Also, the analysis links Bayes's theorem to maximum-entropy considerations.

KEY WORDS: Information theory; Maximum entropy; Statistical inference.

$\pi_a(\theta|I)$   $\equiv$  given prior or antedata probability density function (pdf) for  $\theta \in \Theta$ , based on prior information  $I$ ,

$\pi_p(\theta|D)$   $\equiv$  postdata pdf for  $\theta \in \Theta$ , where  $D = (y, I)$ , the given sample,  $y$ , and prior information,  $I$ ,

$p(y|I)$   $\equiv$  pdf for  $y$ , given by  
 $p(y|I) = \int_{\Theta} \pi_a(\theta|I) f(y|\theta) d\theta$ , (2.1)

where  $f(y|\theta)$  is the pdf for  $y$  given  $\theta$ . [The term *postdata pdf* is employed instead of *posterior pdf* to emphasize that the optimal form of  $\pi_p(\theta|D)$  is to be derived.]

Note that (2.1) is a definition that does not involve the assumption that the product rule of probability theory necessarily holds. See Jeffreys (1967, pp. 25, 52) for a discussion of assumptions needed for the product rule to be valid.

The inputs and outputs of any information processing rule (IPR) are depicted graphically in Figure 1, where  $I(\theta|y)$ ,



# Bayes thm from optimal information processing




Figure 1. Inputs and Outputs of Any Information Processing Rule.



# Bayes thm from optimal information processing (VB)

$$q^*(\theta) = \arg \min_{q \in \mathcal{P}(\Theta)} \left\{ -\mathbb{E}_{q(\theta)} \log \pi(\mathbf{y}|\theta) + \text{KLD}(q(\theta) \parallel \pi(\theta)) \right\}$$

where

- $\mathcal{P}(\Theta)$  is the set of all probability distributions on  $\Theta$
- $\text{KLD}(q(\theta) \parallel \pi(\theta))$  is the Kullback-Leibler divergence  $\mathbb{E}_{q(\theta)} \log \frac{q(\theta)}{\pi(\theta)}$

Bayes thm gives

$$\pi(\theta|\mathbf{y}) \propto \pi(\mathbf{y}|\theta)\pi(\theta)$$

and of course

$$q^*(\theta) = \pi(\theta|\mathbf{y})$$



# Bayes thm from optimal information processing (VB)

$$q^*(\theta) = \arg \min_{q \in \mathcal{P}(\Theta)} \left\{ -E_{q(\theta)} \log \pi(\mathbf{y}|\theta) + \text{KLD}(q(\theta) \parallel \pi(\theta)) \right\}$$

where

- $\mathcal{P}(\Theta)$  is the set of all probability distributions on  $\Theta$
- $\text{KLD}(q(\theta) \parallel \pi(\theta))$  is the Kullback-Leibler divergence  $E_{q(\theta)} \log \frac{q(\theta)}{\pi(\theta)}$

Bayes thm gives

$$\pi(\theta|\mathbf{y}) \propto \pi(\mathbf{y}|\theta)\pi(\theta)$$

and of course

$$q^*(\theta) = \pi(\theta|\mathbf{y})$$



# Bayes thm from optimal information processing (VB)

$$q^*(\theta) = \arg \min_{q \in \mathcal{P}(\Theta)} \left\{ -\mathbb{E}_{q(\theta)} \log \pi(\mathbf{y}|\theta) + \text{KLD}(q(\theta) \parallel \pi(\theta)) \right\}$$

where

- $\mathcal{P}(\Theta)$  is the set of all probability distributions on  $\Theta$
- $\text{KLD}(q(\theta) \parallel \pi(\theta))$  is the Kullback-Leibler divergence  $\mathbb{E}_{q(\theta)} \log \frac{q(\theta)}{\pi(\theta)}$

Bayes thm gives

$$\pi(\theta|\mathbf{y}) \propto \pi(\mathbf{y}|\theta)\pi(\theta)$$

and of course

$$q^*(\theta) = \pi(\theta|\mathbf{y})$$



# Variational inference

$$\hat{q}^*(\theta) = \arg \min_{q \in \hat{\mathcal{P}}(\Theta)} \left\{ -\mathbb{E}_{q(\theta)} \log \pi(\mathbf{y}|\theta) + \text{KLD}(q(\theta) \parallel \pi(\theta)) \right\}$$

where

- $\hat{\mathcal{P}}(\Theta)$  is the set of all approximate probability distributions on  $\Theta$
- same objective as optimising ELBO/VB (in ML)
- Zellner(1988) provides the justification



# Variational inference

$$\hat{q}^*(\boldsymbol{\theta}) = \arg \min_{q \in \hat{\mathcal{P}}(\boldsymbol{\Theta})} \left\{ -\mathbb{E}_{q(\boldsymbol{\theta})} \log \pi(\mathbf{y}|\boldsymbol{\theta}) + \text{KLD}(q(\boldsymbol{\theta}) \parallel \pi(\boldsymbol{\theta})) \right\}$$

where

- $\hat{\mathcal{P}}(\boldsymbol{\Theta})$  is the set of all approximate probability distributions on  $\boldsymbol{\Theta}$
- same objective as optimising ELBO/VB (in ML)
- Zellner(1988) provides the justification



# Variational inference

$$\hat{q}^*(\boldsymbol{\theta}) = \arg \min_{q \in \hat{\mathcal{P}}(\Theta)} \left\{ -\mathbb{E}_{q(\boldsymbol{\theta})} \log \pi(\mathbf{y}|\boldsymbol{\theta}) + \text{KLD}(q(\boldsymbol{\theta}) \parallel \pi(\boldsymbol{\theta})) \right\}$$

where

- $\hat{\mathcal{P}}(\Theta)$  is the set of all approximate probability distributions on  $\Theta$
- same objective as optimising ELBO/VB (in ML)
- Zellner(1988) provides the justification



# VB is valid Bayesian inference



Forrest Gump 1994



# Example: VB correcting Poisson regression

$$y_i \sim \text{Poisson}(\exp(\beta x_i)), \quad \beta \sim \mathcal{N}(0, 1)$$

log-posterior with constant prior

$$U(\beta) = -\frac{1}{2}\beta^2 + \beta \sum_i x_i y_i - \sum_i \exp(\beta x_i) + \text{constant}$$

Gaussian approximation/“Laplace method” found using MLE/glm

$$\beta | \mathbf{y} \approx \mathcal{N}(\hat{\beta}, \hat{\sigma}^2)$$

$$\hat{\beta} = \arg \max U(\beta), \quad \frac{1}{\hat{\sigma}^2} = -\text{Hessian}(U(\beta)) |_{\beta=\hat{\beta}}$$



# Example: VB correcting Poisson regression

$$y_i \sim \text{Poisson}(\exp(\beta x_i)), \quad \beta \sim \mathcal{N}(0, 1)$$

log-posterior with constant prior

$$U(\beta) = -\frac{1}{2}\beta^2 + \beta \sum_i x_i y_i - \sum_i \exp(\beta x_i) + \text{constant}$$

Gaussian approximation/“Laplace method” found using MLE/glm

$$\beta | \mathbf{y} \approx \mathcal{N}(\hat{\beta}, \hat{\sigma}^2)$$

$$\hat{\beta} = \arg \max U(\beta), \quad \frac{1}{\hat{\sigma}^2} = -\text{Hessian}(U(\beta)) \Big|_{\beta=\hat{\beta}}$$



## Example: VB correcting Poisson regression

$$y_i \sim \text{Poisson}(\exp(\beta x_i)), \quad \beta \sim \mathcal{N}(0, 1)$$

log-posterior with constant prior

$$U(\beta) = -\frac{1}{2}\beta^2 + \beta \sum_i x_i y_i - \sum_i \exp(\beta x_i) + \text{constant}$$

Gaussian approximation/“Laplace method” found using MLE/glm

$$\beta | \mathbf{y} \approx \mathcal{N}(\hat{\beta}, \hat{\sigma}^2)$$

$$\hat{\beta} = \arg \max U(\beta), \quad \frac{1}{\hat{\sigma}^2} = -\text{Hessian}(U(\beta)) \Big|_{\beta=\hat{\beta}}$$



## Example: VB correcting Poisson regression (II)

Mean correction with a  $\mathcal{N}(0, 1)$  prior

$$\beta \approx \mathcal{N}(\hat{\beta} + \delta, \hat{\sigma}^2)$$

Since the MLE-approximation is already useful, then

- $|\delta|$  should be small
- optimization should be fast and numerical stable



## Example: VB correcting Poisson regression (II)

Mean correction with a  $\mathcal{N}(0, 1)$  prior

$$\beta \approx \mathcal{N}(\hat{\beta} + \delta, \hat{\sigma}^2)$$

Since the MLE-approximation is already useful, then

- $|\delta|$  should be small
- optimization should be fast and numerical stable



# Example: VB correcting Poisson regression (III)

The VB correction will minimise

$$-E_{q(\delta)} \log\text{-likelihood} + \text{KLD}(q(\delta) \parallel \text{prior})$$



## Example: VB correcting Poisson regression (IV)

The KLD between two multivariate Gaussians is

$$D_{\text{KL}}(\mathcal{N}_0 \parallel \mathcal{N}_1) = \frac{1}{2} \left( \text{tr}(\Sigma_1^{-1} \Sigma_0) - k + (\mu_1 - \mu_0)^\top \Sigma_1^{-1} (\mu_1 - \mu_0) + \ln \left( \frac{\det \Sigma_1}{\det \Sigma_0} \right) \right)$$

Here,  $(\cdot)_1$  is the “prior” and  $(\cdot)_0$  is the “posterior” (the notation is backwards in the Wikipedia page...)

Since we keep the variance constant, the KLD simplifies into

$$\text{KLD}(\delta) = \frac{1}{2} (\hat{\beta} + \delta)^2 + \text{constant}$$



## Example: VB correcting Poisson regression (IV)

The KLD between two multivariate Gaussians is

$$D_{\text{KL}}(\mathcal{N}_0 \parallel \mathcal{N}_1) = \frac{1}{2} \left( \text{tr}(\Sigma_1^{-1} \Sigma_0) - k + (\mu_1 - \mu_0)^\top \Sigma_1^{-1} (\mu_1 - \mu_0) + \ln \left( \frac{\det \Sigma_1}{\det \Sigma_0} \right) \right)$$

Here,  $(\cdot)_1$  is the “prior” and  $(\cdot)_0$  is the “posterior” (the notation is backwards in the Wikipedia page...)

Since we keep the variance constant, the KLD simplifies into

$$\text{KLD}(\delta) = \frac{1}{2} (\hat{\beta} + \delta)^2 + \text{constant}$$



## Example: VB correcting Poisson regression (V)

The negative expected likelihood is

$$-E_{q(\delta)} \left[ \beta \sum_i x_i y_i - \sum_i \exp(\beta x_i) \right] + \text{constant}$$

Since  $q(\delta)$  is Normal, we can do this directly

$$-(\hat{\beta} + \delta) \sum_i x_i y_i + \sum_i \exp \left( (\hat{\beta} + \delta) x_i + \frac{1}{2} x_i^2 \hat{\sigma}^2 \right)$$



## Example: VB correcting Poisson regression (V)

The negative expected likelihood is

$$-E_{q(\delta)} \left[ \beta \sum_i x_i y_i - \sum_i \exp(\beta x_i) \right] + \text{constant}$$

Since  $q(\delta)$  is Normal, we can do this directly

$$-(\hat{\beta} + \delta) \sum_i x_i y_i + \sum_i \exp \left( (\hat{\beta} + \delta) x_i + \frac{1}{2} x_i^2 \hat{\sigma}^2 \right)$$



## Example: VB correcting Poisson regression (V)

The negative expected likelihood is

$$-E_{q(\delta)} \left[ \beta \sum_i x_i y_i - \sum_i \exp(\beta x_i) \right] + \text{constant}$$

Since  $q(\delta)$  is Normal, we can do this directly

$$-(\hat{\beta} + \delta) \sum_i x_i y_i + \sum_i \exp \left( (\hat{\beta} + \delta) x_i + \frac{1}{2} x_i^2 \hat{\sigma}^2 \right)$$



## Example: VB correcting Poisson regression (VI)

Adding the two terms, we get

$$\hat{\delta} = \arg \min \left( A\delta + \frac{1}{2}B\delta^2 + \sum_i C_i \exp(D_i\delta) \right) = \arg \min V(\delta)$$

for some  $A, B, \{C_i\}, \{D_i\}$ .

Since we are already pretty close to the solution, then

$$\hat{\delta} \approx -\frac{V'(0)}{V''(0)} = -\frac{A + \sum_i C_i D_i}{B + \sum_i C_i D_i^2}$$

(can iterate if needed)



## Example: VB correcting Poisson regression (VI)

Adding the two terms, we get

$$\hat{\delta} = \arg \min \left( A\delta + \frac{1}{2}B\delta^2 + \sum_i C_i \exp(D_i\delta) \right) = \arg \min V(\delta)$$

for some  $A, B, \{C_i\}, \{D_i\}$ .

Since we are already pretty close to the solution, then

$$\hat{\delta} \approx -\frac{V'(0)}{V''(0)} = -\frac{A + \sum_i C_i D_i}{B + \sum_i C_i D_i^2}$$

(can iterate if needed)



# Computational issues

## Numerical methods for sparse matrices

- Factorise  $\mathbf{Q} = \mathbf{L}\mathbf{L}^T$
- Partial inverse  $\mathbf{Q}^{-1}$  for all  $i \sim j$  or  $i = j$ .
- Determinant  $\log |\mathbf{Q}|$
- Solve  $\mathbf{Q}\mathbf{x} = \mathbf{b}$  and  $\mathbf{Q}\mathbf{X} = \mathbf{B}$

```

for (i = 0; i < dimensionSize; i++) {
  for (j = 0; j <= i; j++) {
    float sum = 0;
    for (k = 0; k < j; k++)
      sum += L[i][k] * L[j][k];

    if (i == j)
      L[i][j] = sqrt(A[i][i] - sum);
    else
      L[i][j] = (1.0 / L[j][j]) * (A[i][j] - sum);
  }
}

```

**Most important ingredient to scalable inference, is to handle sparse matrices well!**



# Computational issues

## Numerical methods for sparse matrices

- Factorise  $\mathbf{Q} = \mathbf{L}\mathbf{L}^T$
- Partial inverse  $\mathbf{Q}^{-1}$  for all  $i \sim j$  or  $i = j$ .
- Determinant  $\log |\mathbf{Q}|$
- Solve  $\mathbf{Q}\mathbf{x} = \mathbf{b}$  and  $\mathbf{Q}\mathbf{X} = \mathbf{B}$

```

for (i = 0; i < dimensionSize; i++) {
  for (j = 0; j <= i; j++) {
    float sum = 0;
    for (k = 0; k < j; k++)
      sum += L[i][k] * L[j][k];

    if (i == j)
      L[i][j] = sqrt(A[i][i] - sum);
    else
      L[i][j] = (1.0 / L[j][j]) * (A[i][j] - sum);
  }
}

```

Most important ingredient to scalable inference, is to handle sparse matrices well!



# Computational issues

## Numerical methods for sparse matrices

- Factorise  $\mathbf{Q} = \mathbf{L}\mathbf{L}^T$
- Partial inverse  $\mathbf{Q}^{-1}$  for all  $i \sim j$  or  $i = j$ .
- Determinant  $\log |\mathbf{Q}|$
- Solve  $\mathbf{Q}\mathbf{x} = \mathbf{b}$  and  $\mathbf{Q}\mathbf{X} = \mathbf{B}$

```

for (i = 0; i < dimensionSize; i++) {
  for (j = 0; j <= i; j++) {
    float sum = 0;
    for (k = 0; k < j; k++)
      sum += L[i][k] * L[j][k];

    if (i == j)
      L[i][j] = sqrt(A[i][i] - sum);
    else
      L[i][j] = (1.0 / L[j][j]) * (A[i][j] - sum);
  }
}

```

Most important ingredient to scalable inference, is to handle sparse matrices well!



# Computational issues

## Numerical methods for sparse matrices

- Factorise  $\mathbf{Q} = \mathbf{L}\mathbf{L}^T$
- Partial inverse  $\mathbf{Q}^{-1}$  for all  $i \sim j$  or  $i = j$ .
- Determinant  $\log |\mathbf{Q}|$
- Solve  $\mathbf{Q}\mathbf{x} = \mathbf{b}$  and  $\mathbf{Q}\mathbf{X} = \mathbf{B}$

```

for (i = 0; i < dimensionSize; i++) {
  for (j = 0; j <= i; j++) {
    float sum = 0;
    for (k = 0; k < j; k++)
      sum += L[i][k] * L[j][k];

    if (i == j)
      L[i][j] = sqrt(A[i][i] - sum);
    else
      L[i][j] = (1.0 / L[j][j]) * (A[i][j] - sum);
  }
}

```

Most important ingredient to scalable inference, is to handle sparse matrices well!



# Computational issues

## Numerical methods for sparse matrices

- Factorise  $\mathbf{Q} = \mathbf{L}\mathbf{L}^T$
- Partial inverse  $\mathbf{Q}^{-1}$  for all  $i \sim j$  or  $i = j$ .
- Determinant  $\log |\mathbf{Q}|$
- Solve  $\mathbf{Q}\mathbf{x} = \mathbf{b}$  and  $\mathbf{Q}\mathbf{X} = \mathbf{B}$

```

for (i = 0; i < dimensionSize; i++) {
  for (j = 0; j <= i; j++) {
    float sum = 0;
    for (k = 0; k < j; k++)
      sum += L[i][k] * L[j][k];

    if (i == j)
      L[i][j] = sqrt(A[i][i] - sum);
    else
      L[i][j] = (1.0 / L[j][j]) * (A[i][j] - sum);
  }
}

```

Most important ingredient to scalable inference, is to handle sparse matrices well!



# Computational issues

## Numerical methods for sparse matrices

- Factorise  $\mathbf{Q} = \mathbf{L}\mathbf{L}^T$
- Partial inverse  $\mathbf{Q}^{-1}$  for all  $i \sim j$  or  $i = j$ .
- Determinant  $\log |\mathbf{Q}|$
- Solve  $\mathbf{Q}\mathbf{x} = \mathbf{b}$  and  $\mathbf{Q}\mathbf{X} = \mathbf{B}$

```

for (i = 0; i < dimensionSize; i++) {
  for (j = 0; j <= i; j++) {
    float sum = 0;
    for (k = 0; k < j; k++)
      sum += L[i][k] * L[j][k];

    if (i == j)
      L[i][j] = sqrt(A[i][i] - sum);
    else
      L[i][j] = (1.0 / L[j][j]) * (A[i][j] - sum);
  }
}

```

**Most important ingredient to scalable inference, is to handle sparse matrices well!**



# Not that much around...

- HPC community like “dense” more than “sparse”
- TAUCS (outdated > 20 years ago)
- CHOLMOD (what is used in R, MATLAB, does not scale well in parallel do not handle inverse)
- PARDISO (good, not open-source and but turned more commercial, so...)
- etc...
- With the right person around, we can “just write our own”, right?



Esmail Abdul Fattah



# Not that much around...

- HPC community like “dense” more than “sparse”
- TAUCS (outdated  $> 20$  years ago)
- CHOLMOD (what is used in R, MATLAB, does not scale well in parallel do not handle inverse)
- PARDISO (good, not open-source and but turned more commercial, so...)
- etc...
- With the right person around, we can “just write our own”, right?



Esmail Abdul Fattah



# Not that much around...

- HPC community like “dense” more than “sparse”
- TAUCS (outdated  $> 20$  years ago)
- CHOLMOD (what is used in R, MATLAB, does not scale well in parallel do not handle inverse)
- PARDISO (good, not open-source and but turned more commercial, so...)
- etc...
- With the right person around, we can “just write our own”, right?



Esmail Abdul Fattah



# Not that much around...

- HPC community like “dense” more than “sparse”
- TAUCS (outdated  $> 20$  years ago)
- CHOLMOD (what is used in R, MATLAB, does not scale well in parallel do not handle inverse)
- PARDISO (good, not open-source and but turned more commercial, so...)
- etc...
- With the right person around, we can “just write our own”, right?



Esmail Abdul Fattah



# Not that much around...

- HPC community like “dense” more than “sparse”
- TAUCS (outdated  $> 20$  years ago)
- CHOLMOD (what is used in R, MATLAB, does not scale well in parallel do not handle inverse)
- PARDISO (good, not open-source and but turned more commercial, so...)
- etc...
- With the right person around, we can “just write our own”, right?



Esmail Abdul Fattah



# Not that much around...

- HPC community like “dense” more than “sparse”
- TAUCS (outdated  $> 20$  years ago)
- CHOLMOD (what is used in R, MATLAB, does not scale well in parallel do not handle inverse)
- PARDISO (good, not open-source and but turned more commercial, so...)
- etc...
- With the right person around, we can “just write our own”, right?



Esmail Abdul Fattah




# Cholesky factorisation

```
for (i = 0; i < dimensionSize; i++) {
  for (j = 0; j <= i; j++) {
    float sum = 0;
    for (k = 0; k < j; k++)
      sum += L[i][k] * L[j][k];

    if (i == j)
      L[i][j] = sqrt(A[i][i] - sum);
    else
      L[i][j] = (1.0 / L[j][j] * (A[i][j] - sum));
  }
}
```



We can, because that person participated in the Valencia summer school!






# 3 years later...

Sparse • Smart • Structured • Scalable

## A High-Performance Framework for Sparse Matrices

sTiles targets the full **Cholesky** → **solve** → **selected inverse** pipeline on a single shared-memory node, using tile-based parallelism and GPU acceleration. Today's release handles symmetric positive-definite matrices across the entire spectrum, from very sparse to fully dense, with one unified solver. Distributed-memory support is on the roadmap.

[View Examples](#) [API Documentation](#)



Sparse   Semi-Sparse   Semi-Dense   Dense





# 3 years later...

- We're getting there
- Announced in Glasgow'25 in May
- ...the we found some new problematic cases
- that took another year to fix
- Written for modern CPU and will have GPU acceleration (for hard cases)
- Results are *very good*, with excellent parallel scaling
- Handle hard cases particularly well
- Open source when done
- Stay tuned!





## 3 years later...

- We're getting there
- Announced in Glasgow'25 in May
- ...the we found some new problematic cases
- that took another year to fix
- Written for modern CPU and will have GPU acceleration (for hard cases)
- Results are *very* good, with excellent parallel scaling
- Handle hard cases particularly well
- Open source when done
- Stay tuned!





## 3 years later...

- We're getting there
- Announced in Glasgow'25 in May
- ...the we found some new problematic cases
- that took another year to fix
- Written for modern CPU and will have GPU acceleration (for hard cases)
- Results are *very good*, with excellent parallel scaling
- Handle hard cases particularly well
- Open source when done
- Stay tuned!





## 3 years later...

- We're getting there
- Announced in Glasgow'25 in May
- ...the we found some new problematic cases
- that took another year to fix
- Written for modern CPU and will have GPU acceleration (for hard cases)
- Results are *very* good, with excellent parallel scaling
- Handle hard cases particularly well
- Open source when done
- Stay tuned!





## 3 years later...

- We're getting there
- Announced in Glasgow'25 in May
- ...the we found some new problematic cases
- that took another year to fix
- Written for modern CPU and will have GPU acceleration (for hard cases)
- Results are *very good*, with excellent parallel scaling
- Handle hard cases particularly well
- Open source when done
- Stay tuned!





## 3 years later...

- We're getting there
- Announced in Glasgow'25 in May
- ...the we found some new problematic cases
- that took another year to fix
- Written for modern CPU and will have GPU acceleration (for hard cases)
- Results are *very good*, with excellent parallel scaling
- Handle hard cases particularly well
- Open source when done
- Stay tuned!





## 3 years later...

- We're getting there
- Announced in Glasgow'25 in May
- ...the we found some new problematic cases
- that took another year to fix
- Written for modern CPU and will have GPU acceleration (for hard cases)
- Results are *very* good, with excellent parallel scaling
- Handle hard cases particularly well
- Open source when done
- Stay tuned!





## 3 years later...

- We're getting there
- Announced in Glasgow'25 in May
- ...the we found some new problematic cases
- that took another year to fix
- Written for modern CPU and will have GPU acceleration (for hard cases)
- Results are *very* good, with excellent parallel scaling
- Handle hard cases particularly well
- Open source when done
- Stay tuned!





## 3 years later...

- We're getting there
- Announced in Glasgow'25 in May
- ...the we found some new problematic cases
- that took another year to fix
- Written for modern CPU and will have GPU acceleration (for hard cases)
- Results are *very* good, with excellent parallel scaling
- Handle hard cases particularly well
- Open source when done
- Stay tuned!



Thank you • شكرا



جامعة الملك عبد الله  
للعلوم والتقنية  
King Abdullah University of  
Science and Technology