

# The DRPM Strikes Back: More Flexibility for a Bayesian Spatio-Temporal Clustering Model

Candidate: Federico Angelo Mor

Advisor: Prof. Alessandra Guglielmi  
Coadvisor: Prof. Alessandro Carminati



**POLITECNICO**  
MILANO 1863

December 11, 2024

## What is the thesis about?

The Dependent Random Partition Model from (Page et al., 2022) is a Bayesian spatio-temporal clustering model which directly models the temporal dependencies in the sequence of clusters over time.

Currently, the model's implementation

- produces up to spatially-informed clusters  
→ I additionally introduced covariates information
- only accepts complete datasets  
→ I made it work with missing values in the target variable
- has quite slow execution times (especially on large datasets)  
→ I developed a brand-new and more efficient implementation in Julia rather than C

① Description of the problem

What is the DRPM

How did we improve it

② Implementation and optimizations


③ Analysis of the models

④ Conclusion

# Clustering

The Dependent Random Partition Model from (Page et al., 2022) is a Bayesian spatio-temporal **clustering** model which directly models the temporal dependencies in the sequence of clusters over time.


Clustering is a fundamental technique of unsupervised learning where a set of data points has to be divided into homogeneous groups of units which exhibit a similar behaviour.



# Why going Bayesian?

The Dependent Random Partition Model from (Page et al., 2022) is a **Bayesian** spatio-temporal clustering **model** which directly models the temporal dependencies in the sequence of clusters over time.


Bayesian models incorporate prior information on the model parameters and allow to assess uncertainty when performing inference on the results.



## A bit of (spatio-temporal) context

The Dependent Random Partition Model from (Page et al., 2022) is a Bayesian **spatio-temporal** clustering model which directly models the temporal dependencies in the sequence of clusters over time.


In spatio-temporal datasets, observations are collected over time and across various spatial locations. So we will have  $n$  units that have to be clustered at all time instants  $t = 1, \dots, T$ .



# Why should we care about temporal dependencies?

The **Dependent** Random Partition Model from (Page et al., 2022) is a Bayesian spatio-temporal clustering model which **directly models the temporal dependencies in the sequence of clusters over time**.

This allows to derive a more gentle and interpretable evolution of clusters.



**Figure:**  $\text{ARI}(\hat{p}_t, \hat{p}_{t+k})$  computed for DRPM (left) and for a competitor model (right).

# Why should we care about temporal dependencies?

The **Dependent** Random Partition Model from (Page et al., 2022) is a Bayesian spatio-temporal clustering model which **directly models the temporal dependencies in the sequence of clusters over time**.

This allows to derive a more gentle and interpretable evolution of clusters.




Figure:  $\text{ARI}(\hat{p}_t, \hat{p}_{t+k})$  computed for DRPM (left) and for a competitor model (right).



# Modelling the temporal dependence (Page et al., 2022)

To introduce temporal dependence in a collection of partitions, Page et al. (2022) assumed a first-order Markovian structure, leading to

$$P(\rho_1, \dots, \rho_T) = P(\rho_T | \rho_{T-1}) \cdots P(\rho_2 | \rho_1) P(\rho_1)$$

where  $P(\rho_1)$  is an exchangeable partition probability function (EPPF). Page et al. (2022) chose the EPPF induced by the Dirichlet Process

$$P(\rho_1) \propto \prod_{j=1}^{k_1} M \cdot (|S_{j1}| - 1)!$$

# Modelling the temporal dependence (Page et al., 2022)

To characterize the other terms  $P(\rho_t|\rho_{t-1})$ , the following auxiliary variables need to be introduced. For all units  $i = 1, \dots, n$  we define

$$\gamma_{it} = \begin{cases} 1 & \text{if unit } i \text{ is } \textit{not} \text{ reallocated when moving from time } t-1 \text{ to } t \\ 0 & \text{otherwise (namely, the unit } i \text{ is reallocated)} \end{cases}$$

These parameters model the similarity between  $\rho_{t-1}$  and  $\rho_t$ :

- if  $\rho_{t-1}$  and  $\rho_t$  are highly dependent
  - $\implies$  their cluster configurations will change minimally
  - $\implies$  the majority of  $\gamma_{it}$  will be 1
- if  $\rho_{t-1}$  and  $\rho_t$  exhibit low dependence
  - $\implies$  their cluster configurations will change significantly
  - $\implies$  the majority of  $\gamma_{it}$  will be 0

# Modelling the temporal dependence (Page et al., 2022)

Page et al. (2022) assumed  $\gamma_{it} \stackrel{\text{ind}}{\sim} \text{Ber}(\alpha_t)$ , where  $\alpha_t \in [0, 1]$  serves as a temporal dependence parameter, spanning from perfect temporal correlation ( $\alpha_t = 0$ ) to full independence ( $\alpha_t = 1$ ).

In this way the formulation of the joint model becomes

$$P(\gamma_1, \rho_1, \dots, \gamma_T, \rho_T) = P(\rho_T | \gamma_T, \rho_{T-1}) P(\gamma_T) \times \dots \\ \times P(\rho_2 | \gamma_2, \rho_1) P(\gamma_2) P(\rho_1)$$

Once the model for the partition is specified, there is considerable flexibility in how to define the remainder of the Bayesian model.

# The DRPM (Page et al., 2022)

DRPM formulation according to Page et al. (2022) (henceforth, CDRPM, with C as the C language used for the model's implementation).

$$Y_{it} | Y_{it-1}, \boldsymbol{\mu}_t^*, \boldsymbol{\sigma}_t^{2*}, \boldsymbol{\eta}, \mathbf{c}_t \stackrel{\text{ind}}{\sim} \mathcal{N}(\mu_{c_{it}}^* + \eta_{1i} Y_{it-1}, \sigma_{c_{it}}^{2*} (1 - \eta_{1i}^2))$$

$$Y_{i1} \stackrel{\text{ind}}{\sim} \mathcal{N}(\mu_{c_{i1}}^*, \sigma_{c_{i1}}^{2*})$$

$$\xi_i = \text{Logit}(\frac{1}{2}(\eta_{1i} + 1)) \stackrel{\text{ind}}{\sim} \text{Laplace}(a, b)$$

$$(\mu_{jt}^*, \sigma_{jt}^*) \stackrel{\text{ind}}{\sim} \mathcal{N}(\vartheta_t, \tau_t^2) \times \mathcal{U}(0, A_\sigma)$$

$$\vartheta_t | \vartheta_{t-1} \stackrel{\text{ind}}{\sim} \mathcal{N}((1 - \varphi_1)\varphi_0 + \varphi_1\vartheta_{t-1}, \lambda^2(1 - \varphi_1^2))$$

$$(\vartheta_1, \tau_t) \stackrel{\text{iid}}{\sim} \mathcal{N}(\varphi_0, \lambda^2) \times \mathcal{U}(0, A_\tau)$$

$$(\varphi_0, \varphi_1, \lambda) \sim \mathcal{N}(m_0, s_0^2) \times \mathcal{U}(-1, 1) \times \mathcal{U}(0, A_\lambda)$$

$$\{\mathbf{c}_t, \dots, \mathbf{c}_T\} \sim \text{tRPM}(\boldsymbol{\alpha}, M) \text{ with } \alpha_t \stackrel{\text{iid}}{\sim} \text{Beta}(a_\alpha, b_\alpha)$$

## Our generalized model

DRPM formulation according to our generalization (henceforth, JDRPM, with J as the Julia language used for the model's implementation).

$$Y_{it} | Y_{it-1}, \boldsymbol{\mu}_t^*, \boldsymbol{\sigma}_t^{2*}, \boldsymbol{\eta}, \mathbf{c}_t \stackrel{\text{ind}}{\sim} \mathcal{N}(\mu_{c_{it}}^* + \eta_{1i} Y_{it-1} + \mathbf{x}_{it}^T \boldsymbol{\beta}_t, \sigma_{c_{it}}^{2*} (1 - \eta_{1i}^2))$$

$$Y_{i1} \stackrel{\text{ind}}{\sim} \mathcal{N}(\mu_{c_{i1}}^* + \mathbf{x}_{i1}^T \boldsymbol{\beta}_1, \sigma_{c_{i1}}^{2*})$$

$$\boldsymbol{\beta}_t \stackrel{\text{ind}}{\sim} \mathcal{N}_p(\mathbf{b}, s^2 I)$$

$$\xi_i = \text{Logit}(\frac{1}{2}(\eta_{1i} + 1)) \stackrel{\text{ind}}{\sim} \text{Laplace}(a, b)$$

$$(\mu_{jt}^*, \sigma_{jt}^{2*}) \stackrel{\text{ind}}{\sim} \mathcal{N}(\vartheta_t, \tau_t^2) \times \text{invGamma}(a_\sigma, b_\sigma)$$

$$\vartheta_t | \vartheta_{t-1} \stackrel{\text{ind}}{\sim} \mathcal{N}((1 - \varphi_1)\varphi_0 + \varphi_1\vartheta_{t-1}, \lambda^2(1 - \varphi_1^2))$$

$$(\vartheta_1, \tau_1^2) \stackrel{\text{iid}}{\sim} \mathcal{N}(\varphi_0, \lambda^2) \times \text{invGamma}(a_\tau, b_\tau)$$

$$(\varphi_0, \varphi_1, \lambda^2) \sim \mathcal{N}(m_0, s_0^2) \times \mathcal{U}(-1, 1) \times \text{invGamma}(a_\lambda, b_\lambda)$$

$$\{\mathbf{c}_t, \dots, \mathbf{c}_T\} \sim \text{tRPM}(\boldsymbol{\alpha}, M) \text{ with } \alpha_t \stackrel{\text{iid}}{\sim} \text{Beta}(a_\alpha, b_\alpha)$$

## Our updated formulation

1) We inserted a regression term in the likelihood and changed the prior distributions of the variance parameters

$$Y_{it}|Y_{it-1}, \boldsymbol{\mu}_t^*, \boldsymbol{\sigma}_t^{2*}, \boldsymbol{\eta}, \mathbf{c}_t \stackrel{\text{ind}}{\sim} \mathcal{N}(\mu_{c_{it}}^* + \eta_{1i} Y_{it-1} + \mathbf{x}_{it}^T \boldsymbol{\beta}_t, \sigma_{c_{it}}^{2*} (1 - \eta_{1i}^2))$$

$$Y_{i1} \stackrel{\text{ind}}{\sim} \mathcal{N}(\mu_{c_{i1}}^* + \mathbf{x}_{i1}^T \boldsymbol{\beta}_1, \sigma_{c_{i1}}^{2*})$$

$$\boldsymbol{\beta}_t \stackrel{\text{ind}}{\sim} \mathcal{N}_p(\mathbf{b}, s^2 I)$$

$$\xi_i = \text{Logit}(\frac{1}{2}(\eta_{1i} + 1)) \stackrel{\text{ind}}{\sim} \text{Laplace}(a, b)$$

$$(\mu_{jt}^*, \sigma_{jt}^{2*}) \stackrel{\text{ind}}{\sim} \mathcal{N}(\vartheta_t, \tau_t^2) \times \text{invGamma}(a_\sigma, b_\sigma)$$

$$\vartheta_t | \vartheta_{t-1} \stackrel{\text{ind}}{\sim} \mathcal{N}((1 - \varphi_1)\varphi_0 + \varphi_1 \vartheta_{t-1}, \lambda^2 (1 - \varphi_1^2))$$

$$(\vartheta_1, \tau_t^2) \stackrel{\text{iid}}{\sim} \mathcal{N}(\varphi_0, \lambda^2) \times \text{invGamma}(a_\tau, b_\tau)$$

$$(\varphi_0, \varphi_1, \lambda^2) \sim \mathcal{N}(m_0, s_0^2) \times \mathcal{U}(-1, 1) \times \text{invGamma}(a_\lambda, b_\lambda)$$

$$\{\mathbf{c}_t, \dots, \mathbf{c}_T\} \sim \text{tRPM}(\boldsymbol{\alpha}, M) \text{ with } \alpha_t \stackrel{\text{iid}}{\sim} \text{Beta}(a_\alpha, b_\alpha)$$

## Our updated formulation

The regressor term  $\beta_t$  provides more flexibility to the model formulation through the insertion of covariates into the likelihood.

Prior:  $\beta_t \sim \mathcal{N}_p(\mathbf{b}, s^2 I)$

Update rule:

for  $t = 1$ :  $f(\beta_t | -) \propto$  kernel of a  $\mathcal{N}\text{Canon}(\mathbf{h}_{(\text{post})}, J_{(\text{post})})$  with

$$\mathbf{h}_{(\text{post})} = \left( \frac{\mathbf{b}}{s^2} + \sum_{i=1}^n \frac{(Y_{it} - \mu_{c_{it}}^*) \mathbf{x}_{it}}{\sigma_{c_{it}}^{2*}} \right) \quad J_{(\text{post})} = \left( \frac{1}{s^2} I + \sum_{i=1}^n \frac{\mathbf{x}_{it} \mathbf{x}_{it}^T}{\sigma_{c_{it}}^{2*}} \right)$$

for  $t > 1$ :  $f(\beta_t | -) \propto$  kernel of a  $\mathcal{N}\text{Canon}(\mathbf{h}_{(\text{post})}, J_{(\text{post})})$  with

$$\mathbf{h}_{(\text{post})} = \left( \frac{\mathbf{b}}{s^2} + \sum_{i=1}^n \frac{(Y_{it} - \mu_{c_{it}}^* - \eta_{1i} Y_{it-1}) \mathbf{x}_{it}}{\sigma_{c_{it}}^{2*}} \right) \quad J_{(\text{post})} = \left( \frac{1}{s^2} I + \sum_{i=1}^n \frac{\mathbf{x}_{it} \mathbf{x}_{it}^T}{\sigma_{c_{it}}^{2*}} \right)$$

where  $\mathcal{N}\text{Canon}(\mathbf{h}, J)$  is the canonical formulation of the  $\mathcal{N}(\boldsymbol{\mu}, \Sigma)$ , with  $\mathbf{h} = \Sigma^{-1} \boldsymbol{\mu}$  and  $J = \Sigma^{-1}$ .

## Our updated formulation

The choice of the inverse gamma distribution for the variance parameters recovers conjugacy within the model.

Prior:  $\sigma_{jt}^{2*} \sim \text{invGamma}(a_\sigma, b_\sigma)$

Update rule:

for  $t = 1$ :  $f(\sigma_{jt}^{2*} | -) \propto$  kernel of a  $\text{invGamma}(a_{\sigma(\text{post})}, b_{\sigma(\text{post})})$  with

$$a_{\tau(\text{post})} = a_\sigma + \frac{|S_{jt}|}{2} \quad b_{\tau(\text{post})} = b_\sigma + \frac{1}{2} \sum_{i \in S_{jt}} (Y_{it} - \mu_{jt}^* - \mathbf{x}_{it}^T \beta_t)^2$$

for  $t > 1$ :  $f(\sigma_{jt}^{2*} | -) \propto$  kernel of a  $\text{invGamma}(a_{\sigma(\text{post})}, b_{\sigma(\text{post})})$  with

$$a_{\tau(\text{post})} = a_\sigma + \frac{|S_{jt}|}{2} \quad b_{\tau(\text{post})} = b_\sigma + \frac{1}{2} \sum_{i \in S_{jt}} (Y_{it} - \mu_{jt}^* - \eta_{1i} Y_{it-1} - \mathbf{x}_{it}^T \beta_t)^2$$

Similar derivations apply to  $\tau_t^2$  and  $\lambda^2$ .



## Additional information level

2) We introduced covariates information inside the prior for the partition.

Page et al. (2022) introduced spatial information by moving to a PPM structure with a *cohesion function*  $C(S_{jt}, \mathbf{s}_{jt}^* | M, \mathcal{S})$ , which measures the compactness of the spatial coordinates  $\mathbf{s}_{jt}^*$ .

$$P(\rho_t | M, \mathcal{S}) \propto \prod_{j=1}^{k_t} C(S_{jt}, \mathbf{s}_{jt}^* | M, \mathcal{S})$$

We implemented the spatial cohesion functions proposed in Page et al. (2016).

## Additional information level

2) We introduced covariates information inside the prior for the partition.

In a similar way, we inserted covariates information by introducing a term composed by a *similarity function*  $g(S_{jt}, \mathbf{x}_{jtr}^* | \mathcal{C})$ , which measures the similarity of the  $r$ -th covariate values  $\mathbf{x}_{jtr}^*$ .

$$P(\rho_t | M, \mathcal{S}, \mathcal{C}) \propto \prod_{j=1}^{k_t} C(S_{jt}, \mathbf{s}_{jt}^* | M, \mathcal{S}) \left( \prod_{r=1}^p g(S_{jt}, \mathbf{x}_{jtr}^* | \mathcal{C}) \right)$$

We implemented the covariates similarity functions proposed in Page et al. (2018).

## Missing data

3) We let the model accept missing data in the target variable through the derivation of an update rule for the missing  $Y_{it}$ 's.

for  $t = 1$ :  $f(Y_{it}|-) \propto \mathcal{N}(\mu_{Y_{it}(\text{post})}, \sigma_{Y_{it}(\text{post})}^2)$  with

$$\sigma_{Y_{it}(\text{post})}^2 = 1 / \left( \frac{1}{\sigma_{c_{it}}^{2*}} + \frac{\eta_{1i}^2}{2\sigma_{c_{it+1}t+1}^{2*}(1 - \eta_{1i}^2)} \right)$$

$$\mu_{Y_{it}(\text{post})} = \sigma_{Y_{it}(\text{post})}^2 \left( \frac{\mu_{c_{it}}^* + \mathbf{x}_{it}^T \beta_t}{\sigma_{c_{it}}^{2*}} + \frac{\eta_{1i}(Y_{it+1} - \mu_{c_{it+1}t+1}^* - \mathbf{x}_{it+1}^T \beta_{t+1})}{\sigma_{c_{it+1}t+1}^{2*}(1 - \eta_{1i}^2)} \right)$$

for  $1 < t < T$ :  $f(Y_{it}|-) \propto \mathcal{N}(\mu_{Y_{it}(\text{post})}, \sigma_{Y_{it}(\text{post})}^2)$  with

$$\sigma_{Y_{it}(\text{post})}^2 = (1 - \eta_{1i}^2) / \left( \frac{1}{\sigma_{c_{it}}^{2*}} + \frac{\eta_{1i}^2}{\sigma_{c_{it+1}t+1}^{2*}} \right)$$


$$\mu_{Y_{it}(\text{post})} = \sigma_{Y_{it}(\text{post})}^2 \left( \frac{\mu_{c_{it}}^* + \eta_{1i} Y_{it-1} + \mathbf{x}_{it}^T \beta_t}{\sigma_{c_{it}}^{2*}(1 - \eta_{1i}^2)} + \frac{\eta_{1i}(Y_{it+1} - \mu_{c_{it+1}t+1}^* - \mathbf{x}_{it+1}^T \beta_{t+1})}{\sigma_{c_{it+1}t+1}^{2*}(1 - \eta_{1i}^2)} \right)$$

for  $t = T$ :  $f(Y_{it}|-)$  is just the likelihood of  $Y_{it}$

# New implementation

4) We developed a brand-new and more efficient implementation for the updated MCMC algorithm which we now describe in the following section.

original MCMC algorithm  
of Page et al. (2022)



# Why Julia?

- combines the ease and expressiveness of high-level languages (e.g. R, python, matlab) with the efficiency and performance of low-level languages (e.g. C, C++, Fortran)
- code can be tested interactively, as with interpreted languages. . . . . but performance is guaranteed through just-in-time compilation
- linear algebra computations are optimized through BLAS and LAPACK libraries
- vast collection of optimized and complete scientific packages, e.g. `Statistics`, `Distributions` (Besançon et al., 2021), and `BenchmarkTools` (Chen et al., 2016)

## Spoiler alert

Using Julia, we obtained **improved computational performance** compared to the original C implementation of Page et al. (2022), with peaks up to a 2x speedup.

This performance gain was achieved through several optimizations steps.

# Optimizing covariates similarities

Problem: optimizing covariates similarity  $g_4$ .

```
function similarity4(X_jt::AbstractVector{<:Real}, mu_c::Real,
↳ lambda_c::Real, a_c::Real, b_c::Real, lg::Bool)
    n = length(X_jt); nm = n/2
    xbar = mean(X_jt)
    aux2 = 0.
    for i in eachindex(X_jt)
        aux2 += X_jt[i]^2
    end
    aux1 = b_c + 0.5 * (aux2 - (n*xbar + lambda_c*mu_c)^2/(n+lambda_c) +
↳ lambda_c*mu_c^2 )
    out = -nm*log2pi + 0.5*log(lambda_c/(lambda_c+n)) + lgamma(a_c+nm) -
↳ lgamma(a_c) + a_c*log(b_c) + (-a_c-nm)*log(aux1)
    return lg ? out : exp(out)
end
```

## Optimizing covariates similarities


Problem: optimizing covariates similarity  $g_4$ .

Solution: applying some optimizing macros on the inner loop.

```
function similarity4(X_jt::AbstractVector{<:Real}, mu_c::Real,
↳ lambda_c::Real, a_c::Real, b_c::Real, lg::Bool)
  n = length(X_jt); nm = n/2
  xbar = mean(X_jt)
  aux2 = 0.
  @inbounds @fastmath @simd for i in eachindex(X_jt)
    aux2 += X_jt[i]^2
  end
  aux1 = b_c + 0.5 * (aux2 - (n*xbar + lambda_c*mu_c)^2/(n+lambda_c) +
↳ lambda_c*mu_c^2)
  out = -nm*log2pi + 0.5*log(lambda_c/(lambda_c+n)) + lgamma(a_c+nm) -
↳ lgamma(a_c) + a_c*log(b_c) + (-a_c-nm)*log(aux1)
  return lg ? out : exp(out)
end
```



# Optimizing covariates similarities



# Analysis of the models

To evaluate the numerical performance of both algorithms, we will analyse posterior samples and clusters estimates in two scenarios

- ① on a synthetic dataset, with kind-of-randomly generated data, that includes only the response variable
- ② on a real-world spatio-temporal dataset about air pollution in Lombardy, derived from the AgrImOnIA project (Fassò et al., 2023)

We will start with an assessment question about our implementation, followed by a specific analysis of the upgrades brought by our work.

## Correctness check

Is our generalized model working as expected?

We fitted both the original and our updated models under identical conditions: same datasets, same MCMC setup (burnin, thinning, number of iterations), and same hyperparameters.

# Correctness check

Is our generalized model working as expected? *yes!*

Everything works nicely, JDRPM produces similar results to CDRPM.

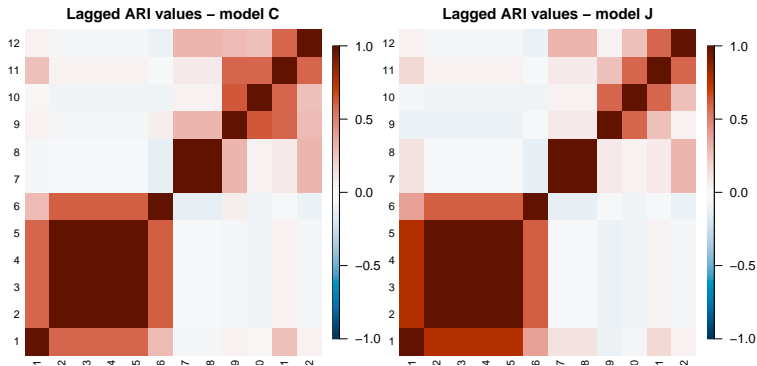
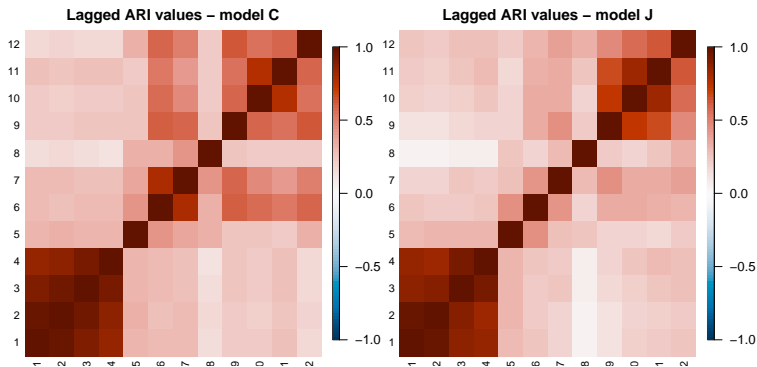


Figure: Lagged ARI values of the two models, fitted on simulated data.

# Correctness check

Is our generalized model working as expected? *yes!*

Everything works nicely, JDRPM produces similar results to CDRPM.



**Figure:** Lagged ARI values of the two models, fitted on real-world data.

## Correctness check

Is our generalized model working as expected? *yes!*

Under identical testing conditions, the models are morally equivalent.

... *but our implementation is faster* 😊

simulated data	MSE mean	MSE median	execution time
CDRPM	1.6221	1.5823	19s (3.8 ms/it)
JDRPM	<b>1.2634</b>	<b>1.2034</b>	<b>13s (2.6 ms/it)</b>

real-world data	MSE mean	MSE median	execution time
CDRPM	0.0142	0.0149	1h38m (54 ms/it)
JDRPM	<b>0.0131</b>	<b>0.0138</b>	<b>48m (26 ms/it)</b>

## Performance with missing values

How will our generalized model perform in presence of missing values?

We randomly removed 10% of the target values  $Y_{it}$  from the dataset, so that they would become our “missing values”, and then repeated the fits in both scenarios.

# Performance with missing values

How will our generalized model perform in presence of missing values?  
*very well!* All the true values lie within the credible intervals.

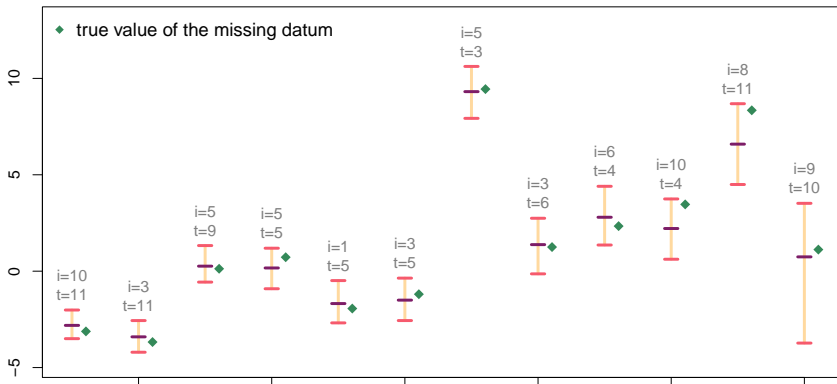


Figure: 95% CIs of the fitted estimates for the missing  $Y_{it}$ 's, on simulated data.



# Performance with missing values

How will our generalized model perform in presence of missing values?  
*quite well!* 74% of the true values lie within the credible intervals.

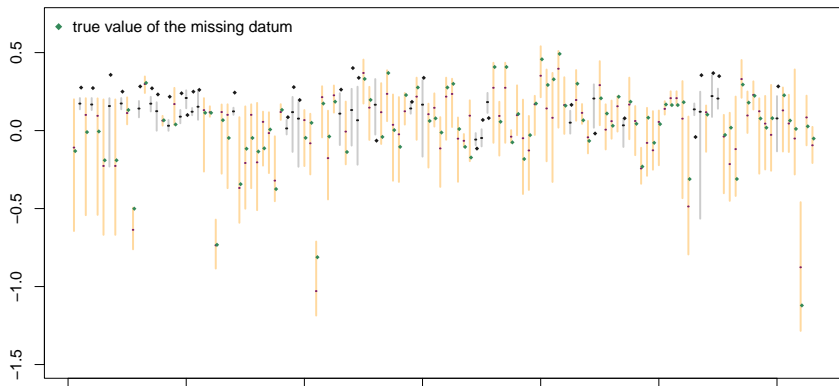


Figure: 95% CIs of the fitted estimates for the missing  $Y_{it}$ 's, on real-world data.

# Performance with missing values

How will our generalized model perform in presence of missing values?  
 Also the temporal trend and the clusters estimates remain very similar.

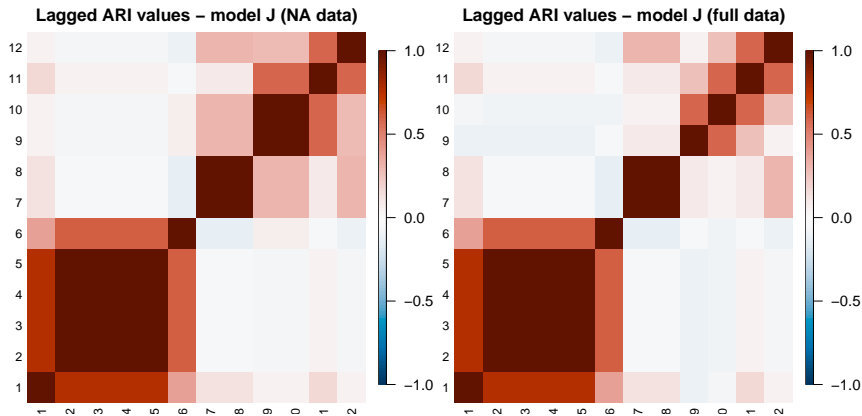


Figure: Lagged ARI values of the two JDRPM fits, on simulated data.

# Performance with missing values

How will our generalized model perform in presence of missing values?  
Also the temporal trend and the clusters estimates remain very similar.




Figure: Lagged ARI values of the two JDRPM fits, on real-world data.

# Covariates in the likelihood

What is the effect of including covariates in the likelihood?

We repeated the missing values analysis, this time including covariates in the likelihood.

## Covariates in the likelihood

What is the effect of including covariates in the likelihood?

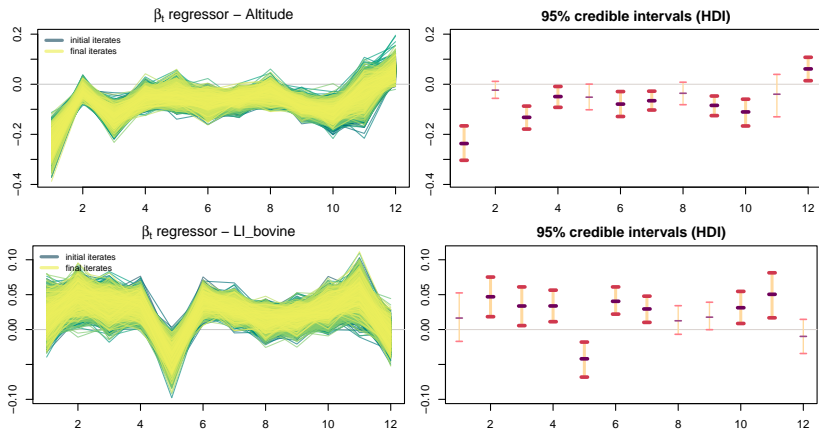
They improve the fitted estimates of  $Y_{it}$  and of the model parameters.

(JDRPM)	MSE mean	MSE median	LPML	WAIC	exec. time
full data	0.0131	0.0138	624.91	-1898.05	48m
NA data	0.0160	0.0170	502.86	-1793.64	<b>43m</b>
NA data + Xlk	<b>0.0127</b>	<b>0.0130</b>	<b>625.81</b>	<b>-1902.74</b>	58m

# Covariates in the likelihood

What is the effect of including covariates in the likelihood?

They provide insights about the included covariates.



# Covariates in the prior

What is the effect of including covariates in the prior?

We repeated the real-world analysis, this time including covariates in the prior.

## Covariates in the prior

What is the effect of including covariates in the prior?

They improve the quality and interpretability of the results.


	MSE mean	MSE median	LPML	WAIC	exec. time
CDRPM	0.0142	0.0149	694.81	-1768.42	1h38m
JDRPM	0.0131	0.0138	624.91	-1898.05	<b>48m</b>
JDRPM + Xcl	<b>0.0126</b>	<b>0.0135</b>	<b>677.71</b>	<b>-1969.76</b>	1h20m



# Covariates in the prior

What is the effect of including covariates in the prior?

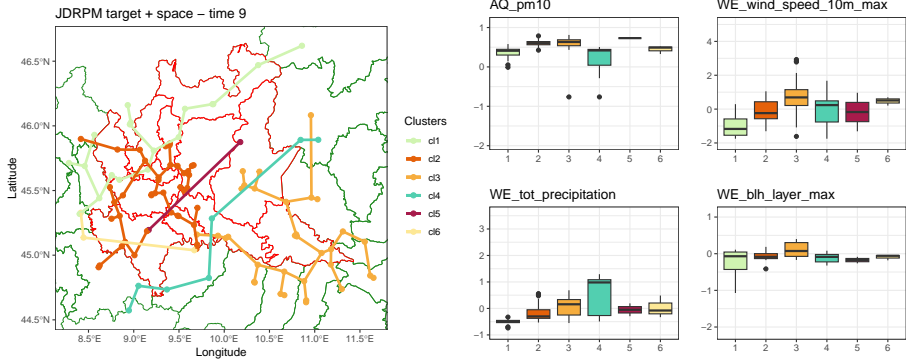
They improve the quality and interpretability of the results.



# Covariates in the prior

What is the effect of including covariates in the prior?

They improve the quality and interpretability of the results.



# Covariates in the prior

What is the effect of including covariates in the prior?

They improve the quality and interpretability of the results.

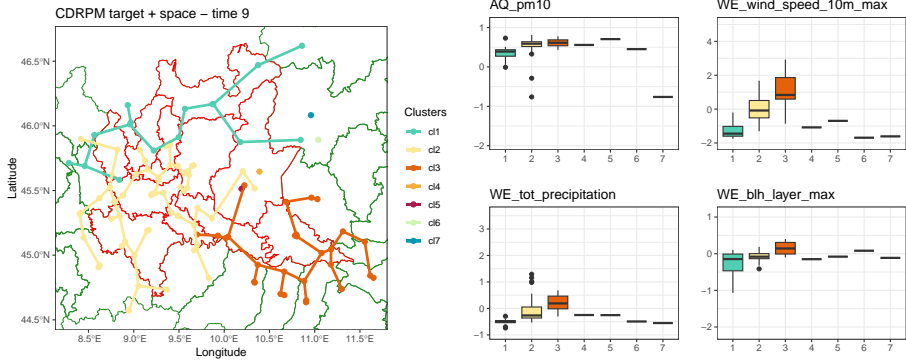


Figure: CDRPM spatially-informed fit.


# Scaling performances

By how much is our implementation faster?

We fitted both models across a “mesh” of dataset sizes, with  $n$  and  $T$  ranging through  $\{10, 50, 100, 250\}$  and with information layers inserted incrementally on top of each other.

# Scaling performances

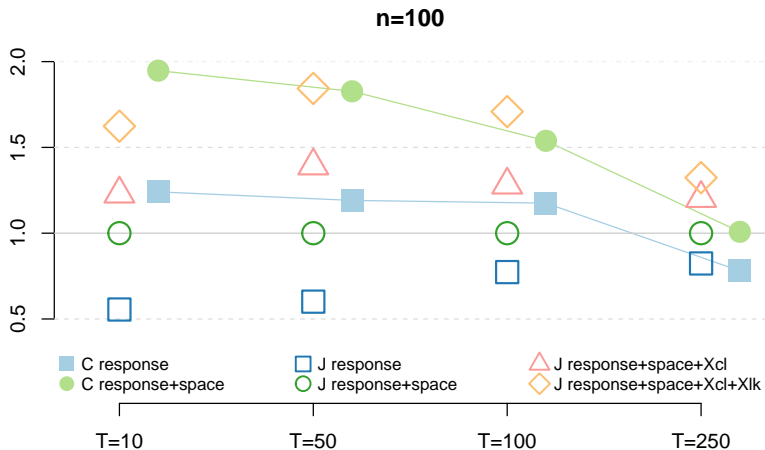
By how much is our implementation faster? *By a lot* 😎



**Figure:** Slowdown factors computed for all fits, relative to JDPRM spatially-informed fit.

# Scaling performances

By how much is our implementation faster? *Hardware is the limit*



**Figure:** Slowdown factors computed for all fits, relative to JDPRM spatially-informed fit.

## Strengths of our work

### Our generalization of the DRPM model

- will allow more flexibility in the real-world researches, thanks to the introduction of covariates and the handling of missing values
- despite the increased complexity, provides more efficiency in the implementation, thus significantly reducing execution times
- is implemented in Julia, facilitating easier code variations and future developments
- is already ready-to-use on R  
(all the Julia and R codes of the thesis are available at <https://github.com/federicomor/Tesi/tree/main/src/JDRPM>)

## Drawbacks of our work






- the estimates of the model parameters will be sensible to the choice of hyperparameters, covariates similarities, spatial cohesions, etc. However this is a characteristic of all complex Bayesian models.
- reaching an appropriate balance between spatial and covariates information may require empirical testing.  
However, the Julia implementation already provides an optional argument, `cv_weight` (defaulted to 1), which allows to scale the contribute of covariates similarities.





*That's all Folks!*

# Bibliography I

-  Besançon, Mathieu et al. (2021). “Distributions.jl: Definition and Modeling of Probability Distributions in the JuliaStats Ecosystem”. In: *Journal of Statistical Software* 98.16, pp. 1–30. ISSN: 1548-7660. DOI: 10.18637/jss.v098.i16.
-  Chen, Jiahao and Jarrett Revels (Aug. 2016). “Robust benchmarking in noisy environments”. In: *arXiv e-prints*, arXiv:1608.04295.
-  Fassò, A. et al. (2023). *AgriImOnIA: Open Access dataset correlating livestock and air quality in the Lombardy region, Italy (3.0.0)*. DOI: <https://doi.org/10.5281/zenodo.7956006>.
-  Page, Garritt L. and Fernando A. Quintana (Sept. 2018). “Calibrating covariate informed product partition models”. In: *Statistics and Computing* 28, pp. 1–23. DOI: 10.1007/s11222-017-9777-z.
-  — (2016). “Spatial Product Partition Models”. In: *Bayesian Analysis* 11.1, pp. 265–298. DOI: 10.1214/15-BA971.

## Bibliography II



Page, Garritt L., Fernando A. Quintana, and David B. Dahl (2022).  
“Dependent Modeling of Temporal Sequences of Random Partitions”.  
In: *Journal of Computational and Graphical Statistics* 31.2,  
pp. 614–627. DOI: [10.1080/10618600.2021.1987255](https://doi.org/10.1080/10618600.2021.1987255).