

Functional ANOVA for Carbon Flux Estimates from Remote Sensing Data

Jonathan Hobbs¹, Matthias Katzfuss², Hai Nguyen¹, Vineet Yadav¹, and Junjie Liu¹

¹Jet Propulsion Laboratory, California Institute of Technology

²Texas A&M University

Correspondence: Jonathan Hobbs (Jonathan.M.Hobbs@jpl.nasa.gov)

1 Supplement

This supplement includes additional exposition of the statistical modeling assumptions in the functional analysis of variance (ANOVA) methodology. This is accompanied by an overview of the Markov chain Monte Carlo (MCMC) algorithm for Bayesian estimation of the parameters and functional ANOVA components.

5 1.1 Spatial Statistical Model

The general functional ANOVA and corresponding notation follow Kaufman and Sain (2010),

$$y_{ijk}(\mathbf{s}) - y_{ijk}^{(0)}(\mathbf{s}) = \mu(\mathbf{s}) + \alpha_i(\mathbf{s}) + \beta_j(\mathbf{s}) + (\alpha\beta)_{ij}(\mathbf{s}) + \epsilon_{ijk}(\mathbf{s}), \quad (1)$$

where $y_{ijk}(\mathbf{s})$ is a carbon flux estimate and $y_{ijk}^{(0)}(\mathbf{s})$ is the corresponding prior flux at spatial location \mathbf{s} . In this study, there are realizations of the spatial fields at N locations for each of the $n_\alpha \times n_\beta$ combinations of the experimental factors and for all n_ϵ replicates. For this collection of realizations, the model can be represented in a form with the N locations assembled into vectors,

$$\mathbf{y}_{ijk} - \mathbf{y}_{ijk}^{(0)} = \boldsymbol{\mu} + \boldsymbol{\alpha}_i + \boldsymbol{\beta}_j + (\boldsymbol{\alpha}\boldsymbol{\beta})_{ij} + \boldsymbol{\epsilon}_{ijk} \quad (2)$$

Further, for convenience, let Σ_μ represent the covariance for the full spatial field $\boldsymbol{\mu}$, with analogous notation for the other ANOVA processes. Functional ANOVA models for spatial data often specify Gaussian process (GP) models with Matérn covariance structure as prior distributions for the ANOVA processes (Kaufman and Sain, 2010; Kang and Cressie, 2013). This approach provides a flexible representation of spatial correlation as a function of distance, but these GP models yield dense covariance matrices (Σ_μ), making the computations for MCMC expensive. Therefore, we use a Vecchia approximation (Katzfuss and Guinness, 2021) for each GP. This approximation facilitates computational demand for MCMC.

Each GP is parameterized by a standard deviation σ , a spatial range parameter λ , and a smoothness parameter ν . These parameters are likely to vary across the GPs for the different ANOVA components, but they can be estimated from the data and can be included in the MCMC procedure. The fully Bayesian inference includes prior distributions for the GP parameters.

In our implementation, we assume the same functional form for a GP parameter prior across the ANOVA components. These forms are

$$\begin{aligned}
& \sigma_\mu \sim \text{Half-normal}(r_\mu) \\
25 \quad \pi(\sigma_\mu) &= \frac{\sqrt{2}}{r_\mu \sqrt{\pi}} \exp\left(-\frac{\sigma_\mu^2}{2r_\mu^2}\right), \sigma_\mu \geq 0 \\
& \lambda_\mu \sim \text{Lognormal}(v_\mu, w_\mu) \\
\pi(\lambda_\mu) &= \frac{1}{\lambda_\mu w_\mu \sqrt{2\pi}} \exp\left(-\frac{(\log \lambda_\mu - v_\mu)^2}{2w_\mu^2}\right), \lambda_\mu > 0 \\
& \nu_\mu \sim \text{Scaled Beta}(b_{1,\mu}, b_{2,\mu}, c_{1,\mu}, c_{2,\mu}) \\
\pi(\nu_\mu) &= \left(\frac{\nu_\mu - c_{1,\mu}}{c_{2,\mu} - c_{1,\mu}}\right)^{b_{1,\mu}-1} \left(\frac{c_{2,\mu} - \nu_\mu}{c_{2,\mu} - c_{1,\mu}}\right)^{b_{2,\mu}-1} \frac{\Gamma(b_{1,\mu} + b_{2,\mu})}{\Gamma(b_{1,\mu}) \Gamma(b_{2,\mu})}, c_{1,\mu} \leq \nu_\mu \leq c_{2,\mu}
\end{aligned}$$

30 where the hyperparameters in parentheses are chosen to vary across ANOVA components. The half-normal form for the GP standard deviations is motivated by the literature on variance components in hierarchical model (Gelman, 2006). The scaled Beta distribution for ν constrains the support for the smoothness parameter and improves convergence of the MCMC sampler in our experience. The resulting prior distributions are shown in Figure S.1.

1.2 Parameterization via Contrasts

35 In classical ANOVA, the factorial model (1) is over-parameterized. There are only $n_\alpha \times n_\beta$ unique group means for the various combinations of factors. This can be reconciled in the model by imposing constraints on the main effects and interactions. Kaufman and Sain (2010) extend this to functional ANOVA by imposing sum-to-zero constraints for each spatial location. For the first main effect, define $n_\alpha - 1$ contrast effects $\alpha_{i'}^*$. Then the original main effects can be related to the contrast effects through a $n_\alpha \times (n_\alpha - 1)$ contrast matrix \mathbf{H}_α . The collection is a set of Helmert contrasts that are scaled such that $\mathbf{H}_\alpha^T \mathbf{H}_\alpha = \mathbf{I}$.

40 Tables S.1 and S.2 provide the contrast matrices for the OCO-2 MIP examples where $n_\alpha = 4$ and $n_\beta = 2$.

Table S.1. Matrix of contrast coefficients, \mathbf{H}_α , for the OCO-2 MIP functional ANOVA with four levels, $n_\alpha = 4$.

Level	Contrast		
	$i' = 1$	$i' = 2$	$i' = 3$
$i = 1$	-0.707107	-0.408248	-0.288675
$i = 2$	0.707107	-0.408248	-0.288675
$i = 3$	0	0.816496	-0.288675
$i = 4$	0	0	0.866025

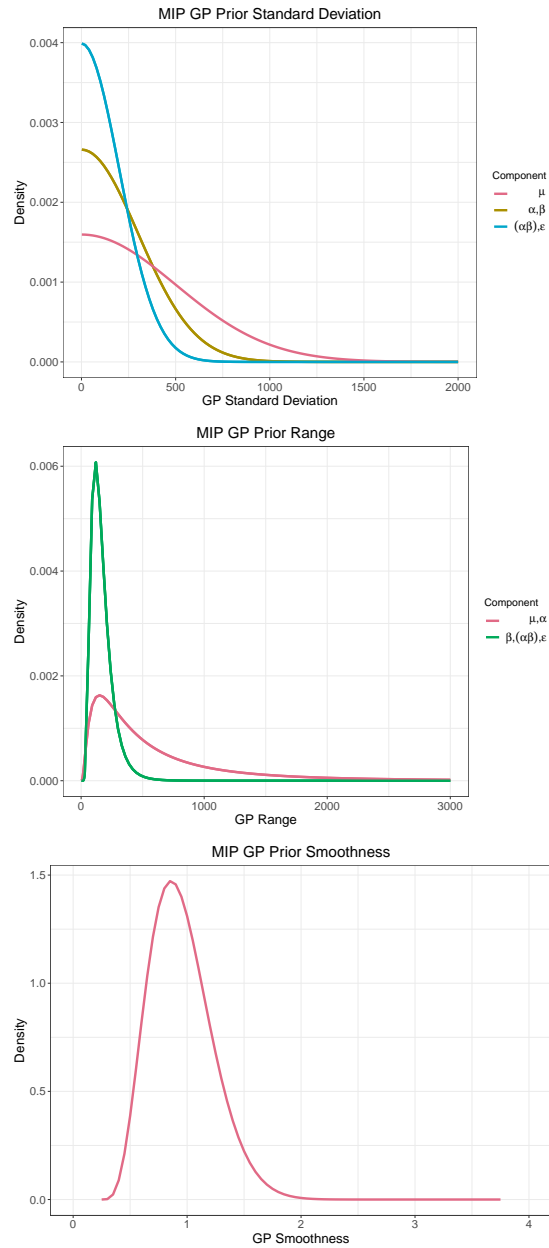


Figure S.1. Prior densities for Gaussian process parameters. Top panel depicts half-normal priors for GP standard deviation σ (units are $\text{gC m}^{-2} \text{yr}^{-1}$). Middle panel depicts lognormal priors for GP range λ (units are km). Bottom panel depicts scaled Beta priors for GP smoothness ν . ANOVA components have different prior parameters as indicated by colors in each panel.

Table S.2. Matrix of contrast coefficients, \mathbf{H}_β , for the OCO-2 MIP functional ANOVA with four levels, $n_\beta = 2$.

Level	Contrast
	$j' = 1$
$j = 1$	-0.707107
$j = 2$	0.707107

1.3 Gibbs Sampler

For many Bayesian hierarchical models, the posterior distribution is not readily available in closed form, but it can be sampled via MCMC. We implement a Gibbs sampler, which iterates conditional sampling among the stochastic nodes given all other elements and the data (Gelman et al., 2013). For the functional ANOVA model, this sampling alternates between updates for the various ANOVA component GPs and for the GP spatial dependence parameters and variances.

1.3.1 Algorithm

1. Update for functional ANOVA spatial processes

The ANOVA structure for a complete balanced design, with all levels of all factors observed, allows for convenient forms for conditional distributions that rely on marginal sums, sometimes termed “row” and “column” sums in classical ANOVA. Because of the zero-sum constraints, the MCMC updates are performed for the reduced rank GPs α^* , β^* , and $(\alpha\beta^*)$. In each case, the conditional posterior distribution, denoted e.g. $p(\boldsymbol{\mu}|\mathbf{Y}, \boldsymbol{\theta})$, given the data and GP model parameters $\boldsymbol{\theta}$, is a multivariate Gaussian distribution, with means and covariances defined below. At each iteration of the Gibbs sampler, these calculations are performed given the current values of the GP dependence parameters, and a draw from each of these multivariate Gaussian distributions is performed to update the ANOVA components.

– For the mean field,

$$p(\boldsymbol{\mu}|\mathbf{y} - \mathbf{y}^{(0)}, \boldsymbol{\theta}) \sim \mathcal{N}(\mathbf{a}_\mu, \mathbf{S}_\mu)$$

$$\mathbf{S}_\mu = (\boldsymbol{\Sigma}_\mu^{-1} + n_\alpha n_\beta n_\epsilon \boldsymbol{\Sigma}_\epsilon^{-1})^{-1}$$

$$\mathbf{a}_\mu = \mathbf{S}_\mu \left(\boldsymbol{\Sigma}_\epsilon^{-1} \sum_{i=1}^{n_\alpha} \sum_{j=1}^{n_\beta} \sum_{k=1}^{n_\epsilon} (\mathbf{y}_{ijk} - \mathbf{y}_{ijk}^{(0)}) \right)$$

– For the Factor A main effect(s), α^* , the following level-specific sums are computed for each $i = 1, \dots, n_\alpha$,

$$\mathbf{y}_{i\cdot\cdot} = \sum_{j=1}^{n_\beta} \sum_{k=1}^{n_\epsilon} (\mathbf{y}_{ijk} - \mathbf{y}_{ijk}^{(0)})$$

Similarly, the contrast coefficients are applied to yield contrast sums,

$$\mathbf{y}_{i'\cdot\cdot} = \sum_{i=1}^{n_\alpha} h_{\alpha(i,i')} \mathbf{y}_{i\cdot\cdot},$$

where $h_{\alpha(i,i')}$ is the row i , column i' element of \mathbf{H}_α

Then the conditional posterior distribution for the contrast effect $\alpha_{i'}$ *

65

$$p(\alpha_{i'}^* | \mathbf{y} - \mathbf{y}^{(0)}, \boldsymbol{\theta}) \sim \mathcal{N}(\mathbf{a}_{\alpha,i'}, \mathbf{S}_\alpha)$$

$$\mathbf{S}_\alpha = (\boldsymbol{\Sigma}_\alpha^{-1} + n_\alpha n_\beta n_\epsilon \boldsymbol{\Sigma}_\epsilon^{-1})^{-1}$$

$$\mathbf{a}_{\alpha,i'} = \mathbf{S}_\alpha (\boldsymbol{\Sigma}_\epsilon^{-1} \mathbf{y}_{i'..})$$

– For the Factor B main effect(s), β^* , the following level-specific sums are computed for each $j = 1, \dots, n_\beta$,

$$\mathbf{y}_{.j\cdot} = \sum_{i=1}^{n_\alpha} \sum_{k=1}^{n_\epsilon} (\mathbf{y}_{ijk} - \mathbf{y}_{ijk}^{(0)})$$

70

The contrast coefficients are applied to yield contrast sums,

$$\mathbf{y}_{.j'\cdot} = \sum_{j=1}^{n_\beta} h_{\beta(j,j')} \mathbf{y}_{.j\cdot},$$

where $h_{\alpha(j,j')}$ is the row j , column j' element of \mathbf{H}_β

Then the conditional posterior distribution for the contrast effect $\beta_{j'}$ *

75

$$p(\beta_{j'}^* | \mathbf{y} - \mathbf{y}^{(0)}, \boldsymbol{\theta}) \sim \mathcal{N}(\mathbf{a}_{\beta,j'}, \mathbf{S}_\beta)$$

$$\mathbf{S}_\beta = (\boldsymbol{\Sigma}_\beta^{-1} + n_\alpha n_\beta n_\epsilon \boldsymbol{\Sigma}_\epsilon^{-1})^{-1}$$

$$\mathbf{a}_{\beta,j'} = \mathbf{S}_\beta (\boldsymbol{\Sigma}_\epsilon^{-1} \mathbf{y}_{.j'\cdot})$$

– For the interaction effect(s), sums are computed for across the treatment combinations,

$$\mathbf{y}_{ij\cdot} = \sum_{k=1}^{n_\epsilon} (\mathbf{y}_{ijk} - \mathbf{y}_{ijk}^{(0)})$$

In a similar fashion, the contrast coefficients are assembled,

80

$$\mathbf{y}_{i'j'\cdot} = \sum_{i=1}^{n_\alpha} \sum_{j=1}^{n_\beta} h_{\alpha(i,i')} h_{\beta(j,j')} \mathbf{y}_{ij\cdot}$$

Then the conditional posterior distribution for the contrast effect $(\alpha\beta)_{i'j'}$ *

$$p((\alpha\beta)_{i'j'}^* | \mathbf{y} - \mathbf{y}^{(0)}, \boldsymbol{\theta}) \sim \mathcal{N}(\mathbf{a}_{(\alpha\beta),i'j'}, \mathbf{S}_{(\alpha\beta)})$$

$$\mathbf{S}_{(\alpha\beta)} = (\boldsymbol{\Sigma}_{(\alpha\beta)}^{-1} + n_\alpha n_\beta n_\epsilon \boldsymbol{\Sigma}_\epsilon^{-1})^{-1}$$

$$\mathbf{a}_{(\alpha\beta),i'j'} = \mathbf{S}_{(\alpha\beta)} (\boldsymbol{\Sigma}_\epsilon^{-1} \mathbf{y}_{i'j'\cdot})$$

85 For each of these updates, if either $n_\alpha > 2$ or $n_\beta > 2$, the posterior means for multiple contrast effects can be solved simultaneously. The Vecchia approximation (Katzfuss and Guinness, 2021) enables efficient construction of the sparse Cholesky factors needed for simulating realizations from these multivariate distributions. These computations are implemented with the `GPvecchia` package in the R statistical computing environment (Katzfuss et al., 2020).

2. Update for spatial model parameters

90 The spatial covariance parameters, $\theta \equiv (\sigma, \lambda, \nu)$, for each ANOVA component are sampled with Metropolis-Hastings (MH) steps. The MH proposal distributions are multivariate Gaussian for the transformation $(\log(\sigma), \log(\lambda), \text{logit}(\nu))$. During a burn-in period for the Markov chain, the proposal covariance is adapted periodically (Gelman et al., 2013). These MCMC steps involve evaluation of the prior densities $\pi(\theta)$ and multivariate Gaussian likelihoods, e.g. $f(\boldsymbol{\mu}|\boldsymbol{\Sigma}_\mu)$, for the relevant GPs.

95 – For the mean field, the conditional posterior includes the likelihood for the corresponding ANOVA GP,

$$p(\sigma_\mu, \lambda_\mu, \nu_\mu | \boldsymbol{\mu}) \propto \pi(\sigma_\mu) \pi(\lambda_\mu) \pi(\nu_\mu) f(\boldsymbol{\mu} | \boldsymbol{\Sigma}_\mu)$$

– For the main and interaction effects, the conditional posterior includes the likelihood for each of the contrast effects. For example, this posterior for the Factor A main effect is

$$p(\sigma_\alpha, \lambda_\alpha, \nu_\alpha | \boldsymbol{\alpha}^*) \propto \pi(\sigma_\alpha) \pi(\lambda_\alpha) \pi(\nu_\alpha) \prod_{i'=1}^{n_\alpha-1} f(\boldsymbol{\alpha}_{i'}^* | \boldsymbol{\Sigma}_\alpha)$$

100 – For the ANOVA error term, the conditional posterior combines the likelihood across all replicates,

$$p(\sigma_\epsilon, \lambda_\epsilon, \nu_\epsilon | \mathbf{y} - \mathbf{y}^{(0)}) \propto \pi(\sigma_\epsilon) \pi(\lambda_\epsilon) \pi(\nu_\epsilon) \prod_{i=1}^{n_\alpha} \prod_{j=1}^{n_\beta} \prod_{k=1}^{n_\epsilon} f(\mathbf{y}_{ijk} - \mathbf{y}_{ijk}^{(0)} | \boldsymbol{\mu}, \boldsymbol{\alpha}^*, \boldsymbol{\beta}^*, (\boldsymbol{\alpha}\boldsymbol{\beta})^*, \boldsymbol{\Sigma}_\epsilon)$$

1.3.2 Initialization

The full MCMC algorithm can successfully converge if the parameters and functional ANOVA states are initialized to reasonable values. In addition, experience indicates that fixing some GP parameters in the preliminary MCMC burn-in also aids convergence. The procedure outlined below has proved successful for the examples in this work as well as tests with simulated data.

1. Compute single-location ANOVA estimates. Location-specific ANOVA models (1) are fit by ordinary least squares.
2. Perform grid-search for preliminary estimates of GP parameters. Given the initial ANOVA component estimates, the GP likelihood is evaluated over a coarse grid of values for the GP parameters θ for all components (mean, main and interaction effects, noise).
3. Run burn-in for ANOVA elements only. The functional ANOVA states in the MCMC are initialized to modest random perturbations of the least-squares estimates, and a short MCMC is run with GP parameters fixed at the values obtained in the grid search. This MCMC is a Gibbs sampler only with no MH updates.

- 115
4. Run full adaptive MCMC. Starting from the final MCMC states from the previous step, the ANOVA components and GP parameters are sampled with the full Metropolis-within-Gibbs algorithm. The MH proposal covariances are updated adaptively in this stage. This burn-in was run for 20,000 iterations.
 5. Run post burn-in sampling with fixed MH proposals. Suitable MH proposal covariances are fixed for the GP parameter updates, and the full MCMC is run for an additional 50,000 iterations. The posterior summaries are constructed from only this stage of the MCMC, where the chain has reached the stationary distribution.

120 **References**

- Gelman, A.: Prior distributions for variance parameters in hierarchical models, *Bayesian Anal.*, 1, 515–534, <https://doi.org/10.1214/06-BA117A>, 2006.
- Gelman, A., Carlin, J. B., Stern, H. S., Dunson, D. B., Vehtari, A., and Rubin, D. B.: *Bayesian Data Analysis*, Chapman & Hall/CRC, Boca Raton, FL, third edn., 2013.
- 125 Kang, E. L. and Cressie, N.: Bayesian Hierarchical ANOVA of Regional Climate-Change Projections from NARCCAP Phase II, *Int. J. Appl. Earth Obs. Geoinf.*, 22, 3–15, <https://doi.org/10.1016/j.jag.2011.12.007>, 2013.
- Katzfuss, M. and Guinness, J.: A general framework for Vecchia approximations of Gaussian processes, *Stat. Sci.*, 36, 124–141, <https://doi.org/10.1214/19-STS755>, 2021.
- Katzfuss, M., Jurek, M., Zilber, D., and Gong, W.: *GPvecchia: Scalable Gaussian-Process Approximations*, R package version 0.1.3, 2020.
- 130 Kaufman, C. G. and Sain, S. R.: Bayesian Functional ANOVA Modeling Using Gaussian Process Prior Distributions, *Bayesian Anal.*, 5, 123–150, <https://doi.org/10.1214/10-BA505>, 2010.