Color test

I just want to be

sure

that all my favourite colors

are

being displayed correctly on this

new

device. If not I'll modify them.

Inferring cosmic particles

LSD seminar

Rémi Bardenet

LAL, LRI, University of Paris-Sud XI

April 13th, 2012

Contents

Google Earth tour

launch

The generative process of an air shower

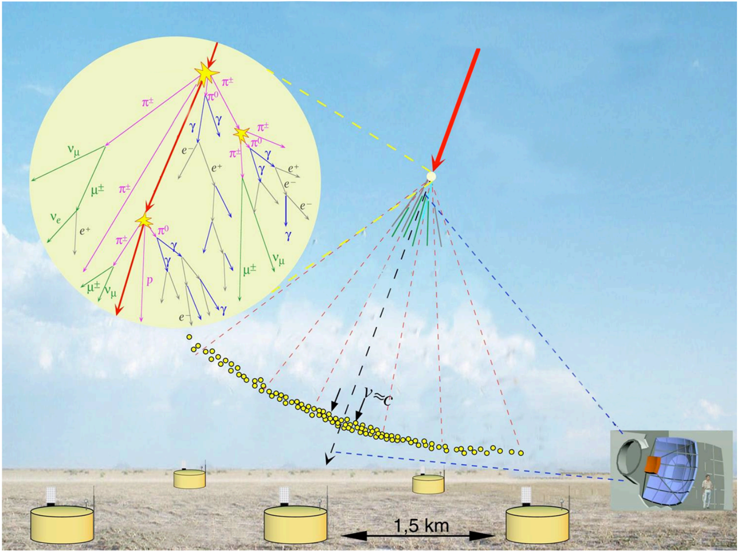

First task: write down the model

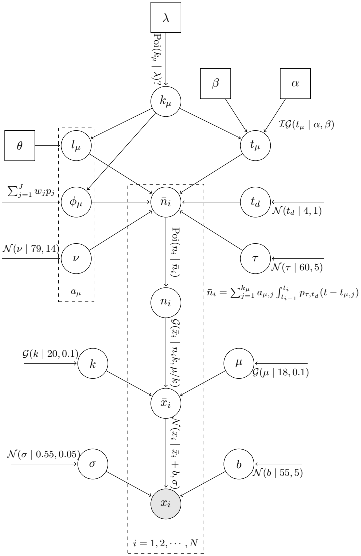

A glance at tank signals

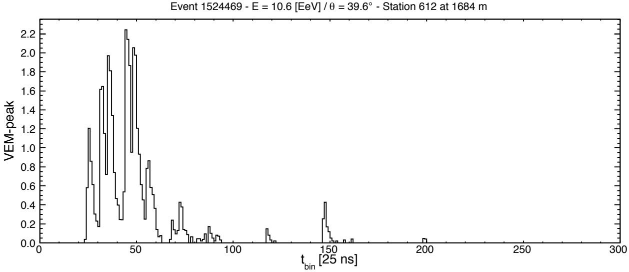

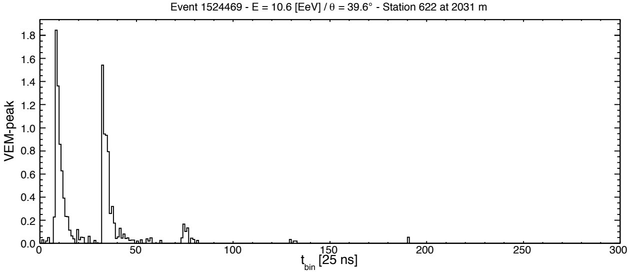

A glance at tank signals

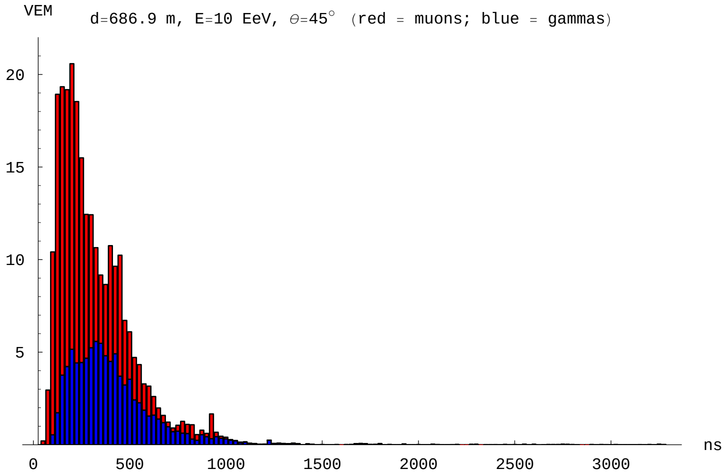

A glance at tank signals

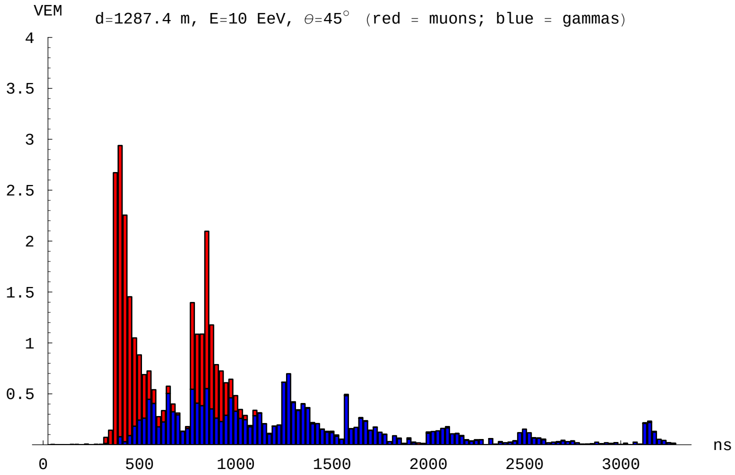

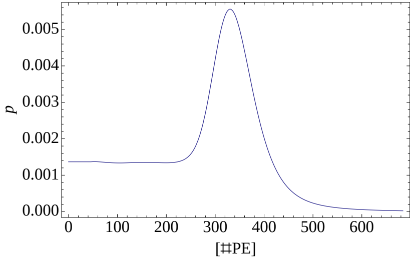

(a) Muonic signal amplitude distribution

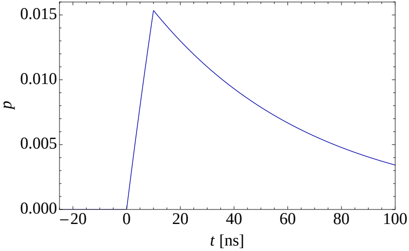

(b) Muonic time response model p_{τ,t_d}

Mean number of Photo-electrons per bin & per muon

$$\bar{n}_i(A_{\boldsymbol{\mu}}, t_{\boldsymbol{\mu}}) = A_\mu \int_{t_{i-1}}^{t_i} p_{\tau,t_d}(t - t_\mu)dt.$$

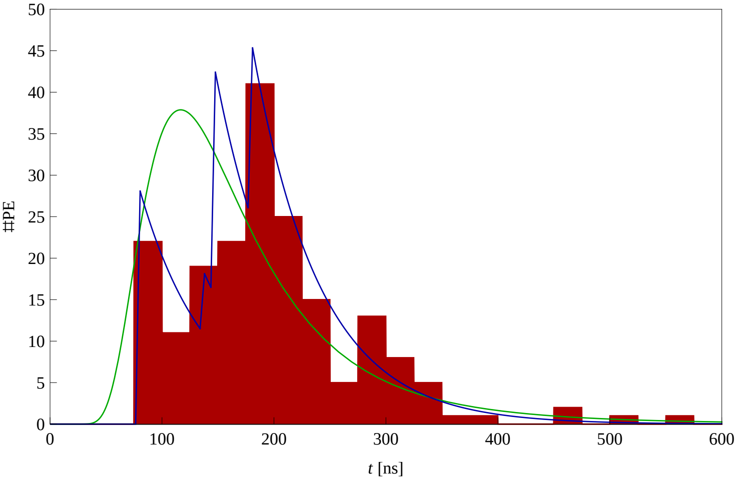

n_i Poisson with mean $\bar{n}_i(\mathbf{A}_{\boldsymbol{\mu}}, \mathbf{t}_{\boldsymbol{\mu}}) = \sum_{j=1}^{N_\mu} \bar{n}_i(A_{\boldsymbol{\mu}_j}, t_{\boldsymbol{\mu}_j}),$

MCMC 101: Metropolis

- Bayesian inference: obtain

$$\pi(\mathbf{A}_{\boldsymbol{\mu}}, \mathbf{t}_{\boldsymbol{\mu}}) = p(\mathbf{A}_{\boldsymbol{\mu}}, \mathbf{t}_{\boldsymbol{\mu}} | \text{signal}) \propto p(\text{signal} | \mathbf{A}_{\boldsymbol{\mu}}, \mathbf{t}_{\boldsymbol{\mu}}) p(\mathbf{A}_{\boldsymbol{\mu}}, \mathbf{t}_{\boldsymbol{\mu}}).$$

- MCMC methods sample from π.

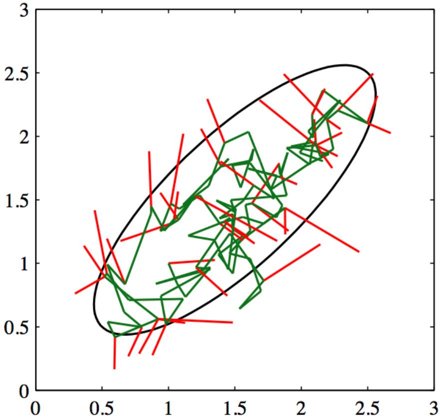

The Metropolis algorithm (blue steps)

AMOR$(\pi(x), X_0, T, \mu_0, \Sigma_0, c)$

$\mathcal{S} \leftarrow \emptyset$
for $t \leftarrow 1$ **to** T
  $\Sigma \leftarrow c\Sigma_{t-1}$ ▷ *scaled adaptive covariance*
  $\widetilde{X} \sim \mathcal{N}(\,\cdot\,|X_{t-1}, \Sigma)$ ▷ *proposal*
  $\widetilde{P} \sim \arg\min_{P \in \mathfrak{P}} L_{(\mu_{t-1}, \Sigma_{t-1})}(P\widetilde{X})$ ▷ *pick an optimal permutation*
  $\widetilde{X} \leftarrow \widetilde{P}\widetilde{X}$ ▷ *permute*
  if $\dfrac{\pi(\widetilde{X})\sum_{P \in \mathfrak{P}} \mathcal{N}(PX_{t-1}|X, \Sigma)}{\pi(X_{t-1})\sum_{P \in \mathfrak{P}} \mathcal{N}(PX|X_{t-1}, \Sigma)} > \mathcal{U}[0,1]$ **then**
    $X_t \leftarrow X$ ▷ *accept*
  else
    $X_t \leftarrow X_{t-1}$ ▷ *reject*
  $\mathcal{S} \leftarrow \mathcal{S} \cup \{X_t\}$ ▷ *update posterior sample*
  $\mu_t \leftarrow \mu_{t-1} + \frac{1}{t}(X_t - \mu_{t-1})$ ▷ *update running mean and covariance*
  $\Sigma_t \leftarrow \Sigma_{t-1} + \frac{1}{t}\left((X_t - \mu_{t-1})(X_t - \mu_{t-1})^{\intercal} - \Sigma_{t-1}\right)$
return $\mathcal{S}$

- Possibly high dimensions but also highly correlated model.
  - Use adaptive proposals.
- The number of muons N_μ is unknown.
  - Use a nonparametric prior or
  - use a Reversible Jump sampler.
- Likelihood $p(\mathbf{n}|\mathbf{A}_\mu, \mathbf{t}_\mu)$ is permutation invariant.
  - Marginals are useless, a problem known as label-switching.

- Possibly high dimensions but also highly correlated model.
  - Use adaptive proposals.
- The number of muons N_μ is unknown.
  - Use a nonparametric prior or
  - use a Reversible Jump sampler.
- Likelihood $\mathcal{P}(\mathbf{n}|\mathbf{A}_\mu, \mathbf{t}_\mu)$ is permutation invariant.
  - Marginals are useless, a problem known as label-switching.

- Possibly high dimensions but also highly correlated model.
  - Use adaptive proposals.
- The number of muons N_μ is unknown.
  - Use a nonparametric prior or
  - use a Reversible Jump sampler.
- Likelihood $\mathcal{P}(\mathbf{n}|\mathbf{A}_{\boldsymbol{\mu}}, \mathbf{t}_{\boldsymbol{\mu}})$ is permutation invariant.
  - Marginals are useless, a problem known as label-switching.

Adaptive Metropolis (blue and green steps)

AMOR$\left(\pi(x), X_0, T, \mu_0, \Sigma_0, c\right)$

$\mathcal{S} \leftarrow \emptyset$

for $t \leftarrow 1$ **to** T

  $\Sigma \leftarrow c\Sigma_{t-1}$ ▷ *scaled adaptive covariance*

  $\widetilde{X} \sim \mathcal{N}\left(\,\cdot\,|X_{t-1}, \Sigma\right)$ ▷ *proposal*

  $\widetilde{P} \sim \arg\min_{P \in \mathfrak{P}} L_{(\mu_{t-1}, \Sigma_{t-1})}\left(P\widetilde{X}\right)$ ▷ *pick an optimal permutation*

  $\widetilde{X} \leftarrow \widetilde{P}\widetilde{X}$ ▷ *permute*

  if $\dfrac{\pi(\widetilde{X}) \sum_{P \in \mathfrak{P}} \mathcal{N}\left(PX_{t-1}|X, \Sigma\right)}{\pi(X_{t-1}) \sum_{P \in \mathfrak{P}} \mathcal{N}\left(PX|X_{t-1}, \Sigma\right)} > \mathcal{U}[0,1]$ **then**

    $X_t \leftarrow X$ ▷ *accept*

  else

    $X_t \leftarrow X_{t-1}$ ▷ *reject*

  $\mathcal{S} \leftarrow \mathcal{S} \cup \{X_t\}$ ▷ *update posterior sample*

  $\mu_t \leftarrow \mu_{t-1} + \frac{1}{t}\left(X_t - \mu_{t-1}\right)$ ▷ *update running mean and covariance*

  $\Sigma_t \leftarrow \Sigma_{t-1} + \frac{1}{t}\left(\left(X_t - \mu_{t-1}\right)\left(X_t - \mu_{t-1}\right)^{\intercal} - \Sigma_{t-1}\right)$

return $\mathcal{S}$

Reversible Jump MCMC

- Nonparametric prior: needs simple model with conjugacy properties for easy sampling, Neal 00.
- RJMCMC Green 05: MCMC kernel on

$$\bigcup_{N_\mu=1}^{N_\mu^{\max}} \{N_\mu\} \times \{ \text{ Parameter space for } N_\mu \text{ muons } \}.$$

- RJMCMC needs careful design of transdimensional moves.

Label-Switching

- If the prior is exchangeable, then $\pi(\mathbf{A}_{\boldsymbol{\mu}}, \mathbf{t}_{\boldsymbol{\mu}})$ is permutation-invariant.

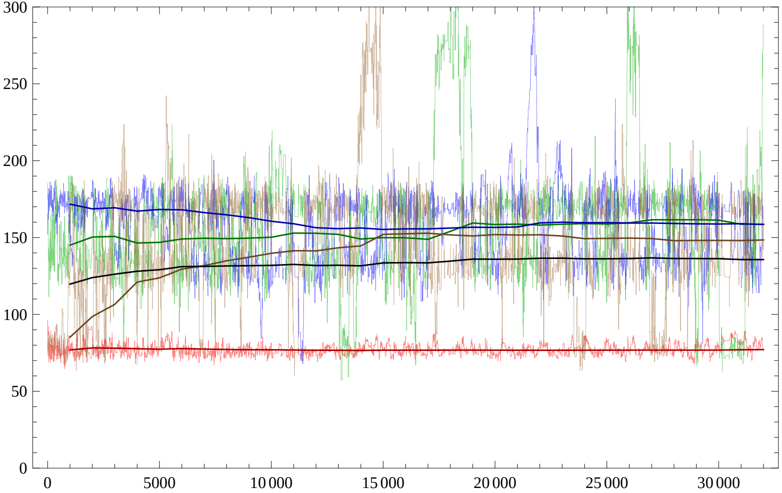

Label-Switching

- If the prior is exchangeable, then $\pi(\mathbf{A}_{\boldsymbol{\mu}}, \mathbf{t}_{\boldsymbol{\mu}})$ is permutation-invariant.

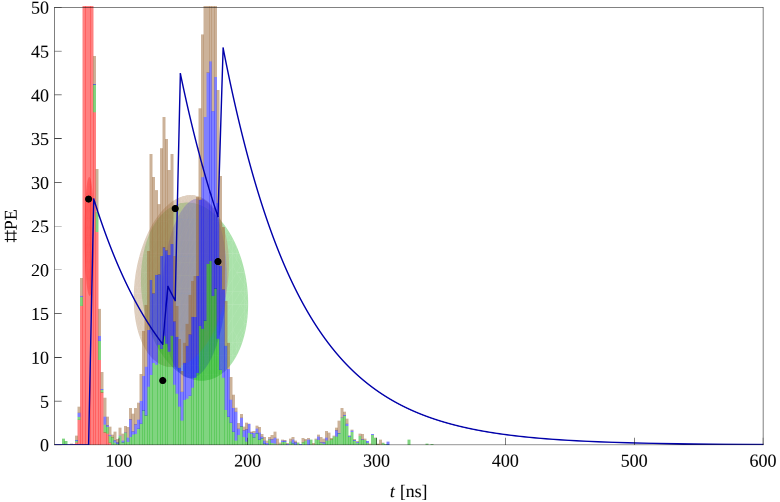

Label-Switching

- If the prior is exchangeable, then $\pi(\mathbf{A}_{\boldsymbol{\mu}}, \mathbf{t}_{\boldsymbol{\mu}})$ is permutation-invariant.

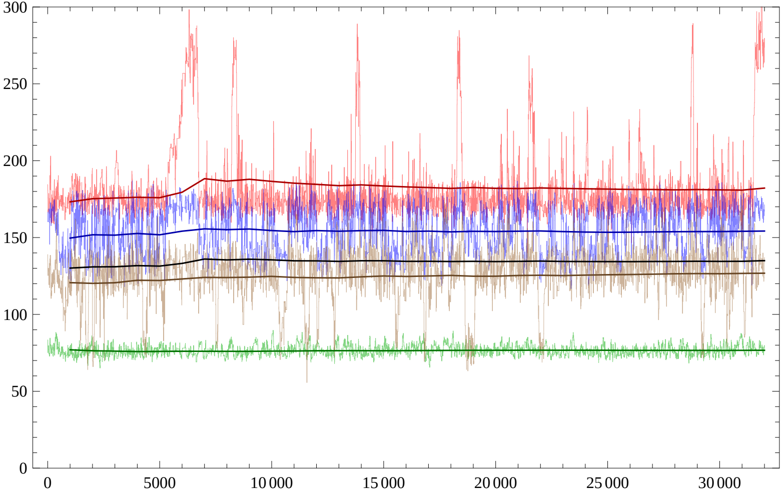

Label-Switching

- If the prior is exchangeable, then $\pi(\mathbf{A}_{\boldsymbol{\mu}}, \mathbf{t}_{\boldsymbol{\mu}})$ is permutation-invariant.

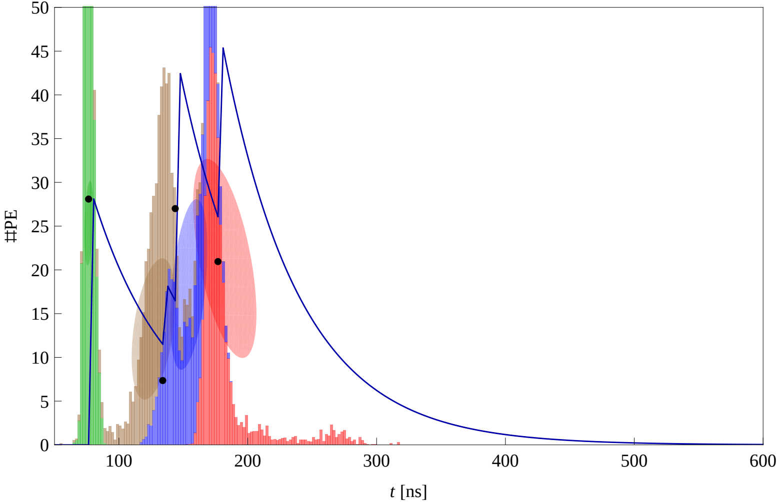

The AMOR sampler

AMOR$(\pi(x), X_0, T, \mu_0, \Sigma_0, c)$

```
S ← ∅
for t ← 1 to T
    Σ ← cΣ_{t-1}                                      ▷ scaled adaptive covariance
    X̃ ~ N( · | X_{t-1}, Σ)                            ▷ proposal
    P̃ ~ argmin_{P∈𝔓} L_{(μ_{t-1}, Σ_{t-1})}(PX̃)      ▷ pick an optimal permutation
    X̃ ← P̃X̃                                           ▷ permute
    if [π(X̃) Σ_{P∈𝔓} N(PX_{t-1} | X, Σ)] / [π(X_{t-1}) Σ_{P∈𝔓} N(PX | X_{t-1}, Σ)] > U[0, 1] then
        X_t ← X                                       ▷ accept
    else
        X_t ← X_{t-1}                                 ▷ reject
    S ← S ∪ {X_t}                                     ▷ update posterior sample
    μ_t ← μ_{t-1} + (1/t)(X_t − μ_{t-1})              ▷ update running mean and covariance
    Σ_t ← Σ_{t-1} + (1/t)((X_t − μ_{t-1})(X_t − μ_{t-1})^T − Σ_{t-1})
return S
```

Where:

$$\widetilde{P} \sim \arg\min_{P \in \mathfrak{P}} L_{(\mu_{t-1}, \Sigma_{t-1})}(P\widetilde{X})$$

$$\frac{\pi(\widetilde{X}) \sum_{P \in \mathfrak{P}} \mathcal{N}(PX_{t-1} | X, \Sigma)}{\pi(X_{t-1}) \sum_{P \in \mathfrak{P}} \mathcal{N}(PX | X_{t-1}, \Sigma)} > \mathcal{U}[0,1]$$

$$\mu_t \leftarrow \mu_{t-1} + \frac{1}{t}(X_t - \mu_{t-1})$$

$$\Sigma_t \leftarrow \Sigma_{t-1} + \frac{1}{t}\left((X_t - \mu_{t-1})(X_t - \mu_{t-1})^{\mathsf{T}} - \Sigma_{t-1}\right)$$

Conclusion and ads

- Particle Physics is cool,
- MCMC is neat,
- Mixing the two is great.
- To learn more on cosmic rays and Auger: Karim Louedec's thesis (link).
- On the model & AMOR: our last AISTATS paper (link) and a submitted Physics paper I can probably send you.
- On adaptive MCMC: see Fort et al 2011 (link)

If you want to know more about MCMC (and Stats in general)