



Transport-Type Models

A CRPropa model in the DEMOS format: generator → transport → post

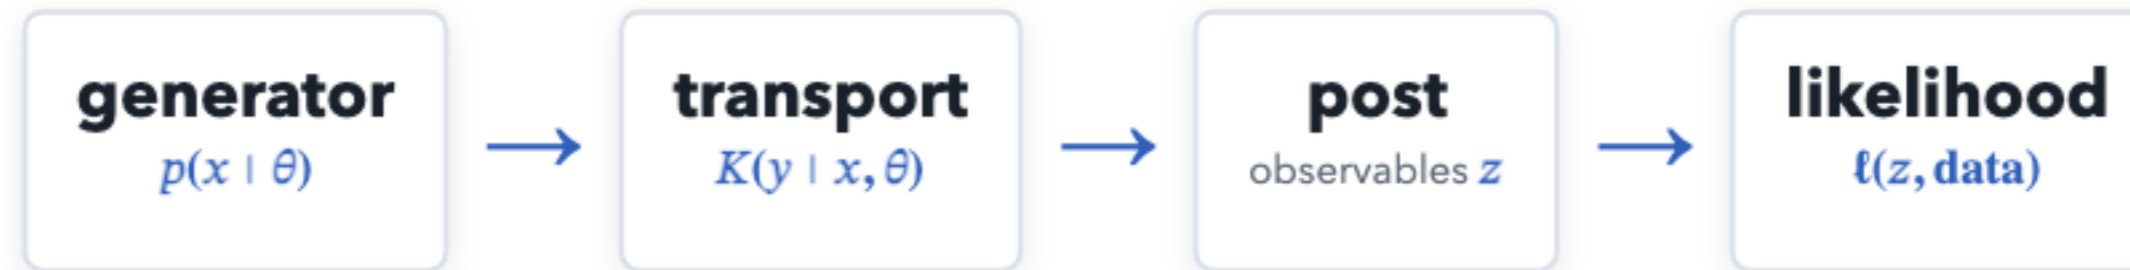
Mikhail Mikhasenko

2026-06-11

The Use Case

A researcher has performed a **parameter-constraint analysis** with CRPropa (or Geant4). The physics model is buried in scripts, containers, and cluster jobs.

DEMOS FORMAT – THE PROBABILITY MODEL



EXECUTION – STAYS WITH THE ENGINE (CRPROPA, GEANT4)

sampling algorithm

number of events

seeds

batching

containers

CPU / GPU

cluster jobs

Present the **model building** in the DEMOS format – a transparent probability model, published with the paper, reusable as a pattern. The engine keeps taking care of the execution.

Example-04 — A Minimal Transport Model

Pluto notebook (U. Hernandez Acosta, MM) – the smallest model with the full structure:

```
1  $\sigma = 0.2$ 
2 gun(; pars) = Normal(pars. $\mu$ ,  $\sigma$ ) # source
3
4 function transport(x; pars) # stochastic, heavy
5      $\Delta\alpha = (2\text{rand}() + 1) * \text{pars}.a$ 
6      $(x + \Delta\alpha)^3 * \exp(x - \text{pars}.b)$ 
7 end
8
9 post(x) = x / 2 # projection
```

```
1 function analysis(pars)
2     gun_dist = gun(; pars)
3     x = rand(gun_dist, 10_000)
4     y = transport.(x; pars)
5     z = post.(y)
6 end
7
8 glob_pars = (a = 0.1, b = 0.3,  $\mu = 1.1$ )
```

```
1 x ~ gun(; pars)
2 y ~ transport(x; pars)
3 z ~ post(y)
```

Three stages – source distribution, stochastic transport kernel, observable projection – compared with data via a likelihood.

The Format Separates Model from Execution

Example-04 shows the **typical workflow** – in a script, model and execution naturally interleave. In the format, each statement splits:

Workflow (script)	Model – what is computed	Execution – how it is computed
<code>gun_dist = gun(; pars)</code>	defines $p(x \theta)$	–
<code>x = rand(gun_dist, 10_000)</code>	$x \sim p(x \theta)$	drawing algorithm, $N = 10\,000$, RNG seed
<code>y = transport.(x; pars)</code>	kernel $y \sim K(\cdot x, \theta_T)$	element-wise loop, batching, CPU/GPU
<code>z = post.(y)</code>	projection $z = \text{post}(y)$	vectorization, histogram filling

The **model column** is the DEMOS description – the declarative computation graph. The **execution column** goes to the engine configuration – modular, exchangeable, reproducible.

The Generator Describes a Distribution

In the model format we describe **distributions** – there is no operation to sample:

$x \sim \text{generator}(\theta)$ — a declaration, not a function call

Model – the generator defines $p(x \mid \theta)$.
For CRPropa, a factorized source density:

$$p(x \mid \theta) = p(A, Z \mid \mathbf{f}) p(E_0 \mid Z, \gamma, R_{\text{cut}}) p(d_0 \mid m)$$

Execution – sampling happens here:

rand(gen, N)

drawing algorithm

number of events


seed

The same distribution can also be reweighted or evaluated as a density – the engine chooses.

Already in example-04, `gun(; pars) = Normal(pars.μ, σ)` returns a distribution object – nothing is drawn until the execution stage.

CRPropa in One Slide

Public framework for **UHECR transport**: nuclei injected at extragalactic sources, propagated through photon backgrounds (CMB, IRB) to an observer.



- Output contains **both** source and final particle IDs, energies, positions, directions, redshift, weight.
- Documented 1D example: sources at 1-1000 Mpc, $E_{\min} = 1$ EeV, mixed H / He / N / Fe injection.

CRPropa's own architecture **already matches** the generator → transport → post decomposition.


Direct Correspondence

DEMOS model concept	CRPropa interpretation
<code>pars</code>	source, transport, nuisance, and normalization parameters
<code>generator(; pars)</code>	initial cosmic-ray source distribution
<code>x ~ generator</code>	source species, energy, distance, redshift
<code>transport(x; pars)</code>	Candidate propagated through a <code>ModuleList</code>
<code>post(y)</code>	spectrum, composition, X_{\max} , ...
<code>analysis(pars)</code>	generate / propagate events, return observables
<code>likelihood</code>	compare postprocessed observables with measurements

One narrow adapter (`run_crpropa_*`) is the **only** framework-specific code – the rest of the model is engine-independent and reusable. Event records, distributions, kernel sketch → appendix.


Two Execution Plans, One Model

PLAN A – DIRECT: RERUN AT EVERY PARAMETER POINT



Transparent – but a CRPropa run for every γ , R_{cut} , composition change.

PLAN B – REFERENCE SAMPLE + IMPORTANCE REWEIGHTING



Same model, same observables – the expensive kernel leaves the inference loop.

The **model is unchanged** – the engine swaps execution plans. That is only possible because the generator is a *density*, not a sampling routine.

Design Summary

generator – defines $p(x \mid \theta)$: distributions live in the **model**

transport – stochastic kernel $K(y \mid x, \theta_T)$: a narrow adapter to the CRPropa container

post + likelihood – observable projection and comparison with data: the **analysis model**

The DEMOS format captures the probability model – the CRPropa engine takes care of the execution layer.

Appendix

From Scalars to Event Records

In example-04, x, y, z are scalars. For a real framework they become **structured event records**:

$$x = (A, Z, E_0, d_0, z_0), \quad y = (A_f, Z_f, E_f, \dots), \quad z = \text{histograms}$$

x – source event

species, injection energy, distance, redshift, weight

y – propagated event

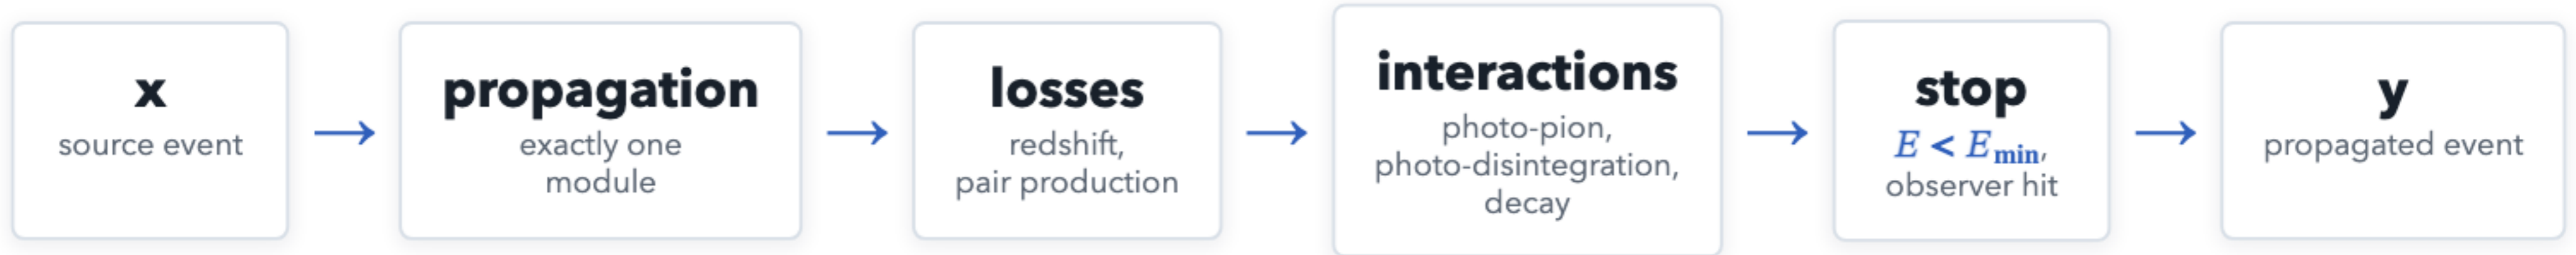
final species, final energy, detection flag, trajectory length

z – prediction

spectrum, composition vs energy, $\langle \ln A \rangle$, $\text{Var}(\ln A)$ – the pipeline shape stays **identical**, only the event type gets richer

Transport Kernel — A Sketch, Not Code

$y \sim K_{\text{CRPropa}}(\cdot \mid x, \theta_T)$ – each **Candidate** walks through the module stack until a stopping condition:



- Stochastic interactions change (A, Z, E) ; continuous losses change E ; the loop repeats per step.
- The **Observer** records the final state – detected or lost.

θ_T (photon field, step limits, E_{\min}) is **fixed per simulation** – changing interaction models or backgrounds means a new CRPropa run.

Projection and Likelihood

Projection $z = \text{post}(y)$ – weighted histograms over detected events:

spectrum

expected counts $\mu_i(\theta)$ per energy bin

composition

$\langle \ln A \rangle_i, \text{Var}(\ln A)_i$ vs energy

later

X_{\max} distributions

Likelihood – compare with measurements:

$$\begin{aligned} \ell(\theta) = & \sum_i \log \text{Pois}(n_i \mid \mu_i(\theta)) \\ & + \sum_i \log \square(\langle \ln A \rangle_i^{\text{obs}} \mid m_i(\theta), \sigma_i) \end{aligned}$$

Poisson counts for the spectrum, Gaussian constraints for composition moments.

Generator — Distributions to Support

Quantity	Distribution (model)	CRPropa source feature (engine)
Fixed species	Dirac (conceptually)	SourceParticleType
Species mixture	Categorical	SourceMultipleParticleTypes / SourceComposition
Fixed energy	Dirac (conceptually)	SourceEnergy
Power-law energy	Pareto / truncated power law	SourcePowerLawSpectrum
Custom cutoff spectrum	custom density	SourceGenericComposition
Fixed position	deterministic value	SourcePosition
Uniform 1D distance	Uniform	SourceUniform1D
Fixed redshift	deterministic value	SourceRedshift
Distance-derived redshift	deterministic transform	SourceRedshift1D
Isotropic direction	uniform sphere sampler	SourceIsotropicEmission
Directed emission	von-Mises-Fisher-like	SourceDirectedEmission

Keep the distributions in the **model**; the engine maps them to source features – it never needs to understand `Normal`, `Dirichlet`, mixtures, or hierarchies.

The Execution Layer — the Engine's Job

Sampling

drawing algorithm, RNG streams, seeds

Statistics

how many events per parameter point

Batching

serialize events, invoke the container **once**

Parallelism

candidates are independent → shared-memory / cluster

Hardware

CPU now, GPU when the engine supports it

Caching

reference samples, reweighting, lookup tables

None of this is physics – but reproducibility dies if it stays unspecified. The format should fix it **modularly, separate from the model.**

Parameter Blocks

Split by what they control – the source, the expensive transport, or the observation model:

```
1 source_pars = (  
2      $\gamma$  = 1.5,  
3     logRcut = 18.8,      # log10(Rcut / V)  
4     fractions = (  
5         H = 0.25, He = 0.25,  
6         N = 0.25, Fe = 0.25,  
7     ),  
8     source_evolution = 0.0,  
9 )  
10 observation_pars = (  
11     flux_normalization = 1.0,  
12     energy_scale_shift = 0.0,  
13 )
```

```
1 transport_pars = (  
2     dmin = 1.0,          # Mpc  
3     dmax = 1000.0,      # Mpc  
4     Emin = 1.0,         # EeV  
5     photon_field = :Kneiske04,  
6     minimum_step = 1e-3, # Mpc = 1 kpc  
7     maximum_step = 10.0, # Mpc  
8 )  
9  
10 pars = (  
11     source_pars...,  
12     transport = transport_pars,  
13     observation = observation_pars,  
14 )
```

Priors (Inference Specification)

```
1 priors = (  
2     γ = Uniform(0.0, 3.0),  
3     logRcut = Uniform(18.0, 20.5),  
4     source_evolution = Uniform(-5.0, 8.0),  
5     # Distributions.jl provides Dirichlet for simplex-valued fractions  
6     fractions = Dirichlet(fill(1.0, 4)),  
7     energy_scale_shift = Normal(0.0, 0.14),  
8 )
```

Free – source spectral index, rigidity cutoff, evolution, composition fractions; energy-scale nuisance.

Fixed – transport settings: changing interaction models, magnetic fields, or photon backgrounds requires a **new CRPropa simulation**.

Event-Record Types

```
1 struct SourceEvent
2     A::Int
3     Z::Int
4     energy::Float64    # EeV
5     distance::Float64  # Mpc
6     redshift::Float64
7     weight::Float64
8 end
```

```
1 struct PropagatedEvent
2     initial::SourceEvent
3     final_A::Int
4     final_Z::Int
5     final_energy::Float64
6     detected::Bool
7     trajectory_length::Float64
8     final_redshift::Float64
9     weight::Float64
10 end
```

In the official example `SourceUniform1D` sets the distance and `SourceRedshift1D` then derives the redshift – the order matters.

Sampling $p(x \mid \theta)$ — Execution Layer

```
1 const SPECIES = ( (:H, 1, 1), (:He, 4, 2), (:N, 14, 7), (:Fe, 56, 26)) # name, A, Z
2
3 function sample_powerlaw_cutoff(rng,  $\gamma$ , Emin, Emax, Z, Rcut)
4     proposal = truncated(Pareto(Emin, max( $\gamma - 1$ , 1e-3)), Emin, Emax)
5     while true # rejection sampler for  $p(E) \propto E^{-\gamma} \exp[-E / (Z Rcut)]$ 
6         E = rand(rng, proposal)
7         rand(rng) < exp(-E / (Z * Rcut)) && return E
8     end
9 end
10
11 function sample_source(rng = Random.default_rng(); pars)
12     name, A, Z = SPECIES[rand(rng, Categorical(collect(pars.fractions)))]
13     Rcut = 10.0^(pars.logRcut - 18) # EV  $\equiv$  EeV per unit charge
14     E = sample_powerlaw_cutoff(rng, pars. $\gamma$ , pars.transport.Emin, 1000.0, Z, Rcut)
15     d = rand(rng, Uniform(pars.transport.dmin, pars.transport.dmax))
16     SourceEvent(A, Z, E, d, distance_to_redshift(d), 1.0)
17 end
```

One possible drawing algorithm for the generator – the model itself only fixes the density.

Inside the Adapter

Essentially the module list of the official 1D example:

```
1 sim = ModuleList()
2 sim.add(SimplePropagation(1 * kpc, 10 * Mpc)) # exactly ONE propagation module
3 sim.add(Redshift())
4 sim.add(PhotoPionProduction(CMB()))
5 sim.add(PhotoPionProduction(IRB_Kneiske04()))
6 sim.add(PhotoDisintegration(CMB()))
7 sim.add(PhotoDisintegration(IRB_Kneiske04()))
8 sim.add(NuclearDecay())
9 sim.add(ElectronPairProduction(CMB()))
10 sim.add(ElectronPairProduction(IRB_Kneiske04()))
11 sim.add(MinimumEnergy(1 * EeV))
12
13 observer = Observer()
14 observer.add(ObserverPoint())
15 sim.add(observer)
```

One propagation module is required; the others implement redshift losses, stochastic interactions, continuous losses, decays, stopping conditions, and observation.

Batched Transport

Pedagogically true:

```
1 y = transport.(x; pars)
```

But implemented batched – never call the container once per event:

```
1 function transport(xs::Vector{SourceEvent}; pars)
2     run_crpropa_batch(xs, pars.transport)
3 end
4
5 xs = sample_source_batch(10_000; pars)
6 ys = transport(xs; pars)
```

1 serialize source events – Arrow, HDF5, JSONL, binary

2 invoke the CRPropa container **once**

3 propagate events in parallel (candidates are independent → shared-memory parallelism)

4 return the final event table

post() Implementation

```
1 struct Prediction
2   energy_edges::Vector{Float64}
3   expected_counts::Vector{Float64}
4   mean_logA::Vector{Float64}
5   variance_logA::Vector{Float64}
6 end
```

```
1 function analysis(pars; nevents = 10_000,
2                 energy_edges)
3   x = sample_source_batch(nevents; pars)
4   y = transport(x; pars)
5   post(y; energy_edges,
6        normalization =
7        pars.observation.flux_normalization)
8 end
```

```
1 function post(events; energy_edges,
2               normalization = 1.0)
3   det = filter(e -> e.detected, events)
4   logE = [log10(e.final_energy) + 18 for e in det]
5   logA = [log(e.final_A) for e in det]
6   nb = length(energy_edges) - 1
7   counts, s1, s2 = (zeros(nb) for _ in 1:3)
8   for (E, lA, e) in zip(logE, logA, det)
9     i = searchsortedlast(energy_edges, E)
10    1 ≤ i ≤ nb || continue
11    w = e.weight
12    counts[i] += w; s1[i] += w*lA; s2[i] += w*lA^2
13  end
14  m = s1 ./ counts
15  Prediction(energy_edges, normalization .* counts,
16            m, s2 ./ counts ./ m.^2)
17 end
```

Likelihood Implementation

```
1 data = (  
2     counts = [...],  
3     mean_logA = [...],  
4     mean_logA_errors = [...],  
5 )  
6  
7 function evaluate(pars, data;  
8     nevents, energy_edges)  
9     pred = analysis(pars;  
10     nevents, energy_edges)  
11     loglikelihood(pred, data)  
12 end
```

```
1 function loglikelihood(pred, data)  
2     ℓ_spectrum = sum(  
3         logpdf(Poisson(max( $\mu$ , eps()))), n)  
4         for ( $\mu$ , n) in zip(  
5             pred.expected_counts,  
6             data.counts))  
7     ℓ_composition = sum(  
8         logpdf(Normal( $\mu$ ,  $\sigma$ ), x)  
9         for ( $\mu$ ,  $\sigma$ , x) in zip(  
10             pred.mean_logA,  
11             data.mean_logA_errors,  
12             data.mean_logA))  
13     ℓ_spectrum + ℓ_composition  
14 end
```

Reference Sample + Reweighting (Code)

```
1 function proposal_source(rng = default_rng())
2   name, A, Z = SPECIES[rand(rng, 1:4)]
3   logE = rand(rng, # q(E) ∝ 1/E
4     Uniform(log(1.0), log(1000.0)))
5   d = rand(rng, Uniform(1.0, 1000.0))
6   SourceEvent(A, Z, exp(logE), d,
7     distance_to_redshift(d), 1.0)
8 end
9
10 # propagate ONCE
11 x_ref = [proposal_source() for _ in 1:1_000_000]
12 y_ref = transport(x_ref; ref_transport_pars)
```

```
1 function source_weight(x; pars)
2   target_density(x; pars) /
3   proposal_density(x)
4 end
5
6 function analysis_reweighted(pars, x_ref, y_ref;
7   energy_edges)
8   weighted = map(x_ref, y_ref) do x, y
9     @set y.weight = source_weight(x; pars)
10  end
11   norm = pars.observation.flux_normalization
12   post(weighted; energy_edges,
13     normalization = norm)
14 end
```

Recommended Notebook Structure

```
1 ## Parameters
2     physical source parameters, fixed transport settings, nuisance parameters
3
4 ## Source model
5     x ~ generator(; pars)
6
7 ## CRPropa transport kernel
8     y ~ transport(x; pars.transport)
9
10 ## Postprocessing
11     z = post(y; pars.observation)
12
13 ## Likelihood
14     data ~ observation_model(z)
15
16 ## Full analysis
17     prediction, log-likelihood
```