Full-Scale Geant4 Simulation of Cosmic-Ray Muography Through Multi-Kilometer-Scale Overburden

Detector-Centered Sampling, Empty-Terrain Flux Calibration, and Geometry-Only Response Folding

Phay Kah Seng; Chien-Chih Chen; Kai-Yu Cheng; Yan-Yu Hsieh; Chih-Hsun Lin; Shih-Hong Lo; Tzu-Hsiang Tsai; Yu-Siang Xiao; Chia-Ming Kuo

Department of Physics, National Central University

A detector-centered Geant4 workflow is used to simulate cosmic-ray muon transport through multi-kilometer-scale overburden. The method improves useful statistics with an enlarged virtual detector plane, removes source-bias using empty-terrain calibration, and maps the weighted virtual-plane flux to detector space using a geometry-only response tensor.

1–10 TeV	FTFP_BERT	10 mm	4 km	50×	$N_{\mathrm{eff}} = 6.52 \times 10^5$	2880
Muon-only energy range	Geant4 physics list	Production cut	Source distance	Virtual-plane gain	Weighted events	Concurrent CPU cores

1. Challenge: Muography at Kilometer Scale

Cosmic-ray muography infers overburden from the direction-dependent intensity of transmitted muons. The simulated terrain volume spans approximately

$$-1500\ \mathrm{m} \le x \le 1500\ \mathrm{m}, \qquad 0\ \mathrm{m} \le y \le 3000\ \mathrm{m}.$$

The terrain has roughly **700 m** of topographic relief, with an additional **200 m** subsurface material layer included below the terrain to avoid truncating muon paths near the lower boundary. The scene is much larger than the detector: the overburden is multi-kilometer scale, while the active area is only **0.4 m** × **0.4 m**. Directly transporting particles only into the physical detector would waste most generated events.

Design principle. Treat the source and virtual detector plane as importance-sampling devices, then recover the physical sea-level flux normalization by empty-terrain calibration and event-by-event reweighting. The final comparison is made in detector space, after folding the calibrated Geant4 flux through the finite detector geometry.

The objective is to construct a reusable detector-position flux map that preserves energy, incident angle, and line-of-sight thickness before the final detector-response folding step.

Detector-centered source

The source is placed **4 km** upstream and aimed at a **64 m** × **64 m** plane. The simulated particle content is μ^+ only, with no electron or gamma flux component. The generation cone is limited to $\theta_{\max} = 25^\circ$, and the analysis region is

$$-13^\circ \le \theta_x, \theta_y \le 13^\circ.$$

The source is not interpreted as a physical horizontal sky plane. It is a detector-centered proposal distribution whose bias is corrected later.

Geant4 transport settings

The full-terrain and empty-terrain simulations use the Geant4 reference physics list FTFP_BERT. The same physics configuration is used for both runs so the empty-terrain calibration removes source and acceptance effects without changing the transport model. A production cut of **10 mm** is applied to secondary-particle production thresholds throughout the geometry. This setting keeps the production configuration consistent across air, terrain, and detector-region scoring while leaving muon transport, energy loss, decay, and scattering to Geant4.

2. Log-Uniform Source and Guan Reweighting

The generated energy range is

$$E_{\min} = 1\ \mathrm{GeV}, \qquad E_{\max} = 10\ \mathrm{TeV}.$$

Energy is sampled log-uniformly,

$$q_E(E) = \frac{1}{E \ln(E_{\max}/E_{\min})}.$$

The generator-level weight is the target flux divided by the proposal density,

$$w_{\mathrm{gen}}(E,\theta) = \frac{\Phi_G(E,\theta)}{q_E(E)} = \Phi_G(E,\theta) E \ln\left(\frac{E_{\max}}{E_{\min}}\right).$$

This gives roughly equal Monte Carlo statistics per energy decade while preserving the physical sea-level spectrum after weighting. The high-energy tail remains important because only sufficiently energetic muons survive long rock paths.

Target sea-level flux

The target flux is the Guan sea-level muon parameterization,

$$\Phi_G(E,\theta) = 10^4\, 0.14\, \tilde{E}^{-2.7} \left[\frac{1}{1 + 1.1E\cos\theta^*/115} + \frac{0.054}{1 + 1.1E\cos\theta^*/850}\right],$$

$$\tilde{E} = E\left(1 + \frac{3.64}{E(\cos\theta^*)^{1.29}}\right).$$

Although the source is generated in a detector-centered coordinate system, the zenith angle θ used in Φ_G is evaluated in the world frame.

3. Empty-Terrain Flux Calibration

The flux model is defined at sea level, while the simulation source is placed **4 km** from the detector. Therefore, the generated source does not directly represent the desired detector-position flux. To calibrate this mismatch, an empty-terrain reference run is performed.

What is kept and what is removed?

The empty-terrain run keeps the air volume, source position, source direction logic, scoring plane, binning, detector-region cuts, physics list, and production cuts unchanged. Only the terrain geometry is removed. This measures the transfer from the biased source to the detector-position flux in the absence of overburden, using the same scoring and binning as the production run. Thus, artificial source-placement and acceptance effects are removed before interpreting full-terrain attenuation.

For each energy-angle bin $b = (E, \theta_x, \theta_y)$,

$$H_b^{\mathrm{empty}} = \sum_{k \in B_b} w_{\mathrm{gen},k}, \qquad \hat{\Phi}_b^{\mathrm{empty}} = \frac{H_b^{\mathrm{empty}}}{A_b \Delta\Omega_b \Delta E_b T}.$$

The empty-terrain correction factor is

$$c_b = \frac{\Phi_b^{\mathrm{target}}}{\hat{\Phi}_b^{\mathrm{empty}}}, \qquad w_k = c_b w_{\mathrm{gen},k}.$$

The same calibrated weights are then applied to full-terrain hits in the corresponding phase-space bins,

$$\hat{\Phi}_j^{\mathrm{full}} = \frac{1}{\Delta A_j \Delta\Omega_j \Delta E_j T} \sum_{k \in B_j} w_k.$$

Projected-angle solid angle

For projected angular coordinates,

$$\mathrm{n} = \frac{(\tan\theta_x, \tan\theta_y, 1)}{\sqrt{1 + \tan^2\theta_x + \tan^2\theta_y}},$$

$$J(\theta_x,\theta_y) = \frac{\sec^2\theta_x \sec^2\theta_y}{(1 + \tan^2\theta_x + \tan^2\theta_y)^{3/2}}, \qquad \Delta\Omega = \iint J\, d\theta_x d\theta_y.$$

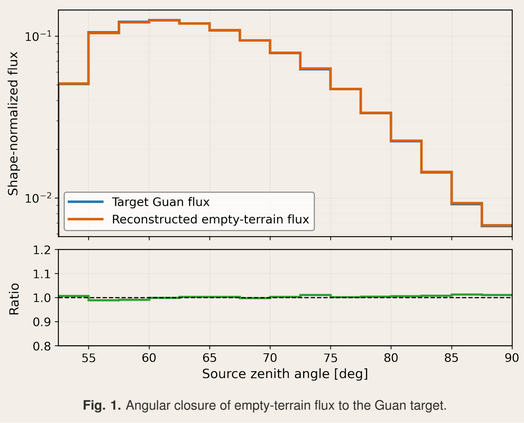

Fig. 1. Angular closure of empty-terrain flux to the Guan target.

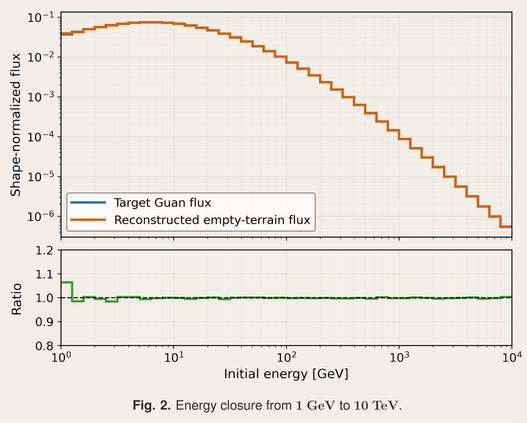

Fig. 2. Energy closure from $1\ \mathrm{GeV}$ to $10\ \mathrm{TeV}$.

Calibration result. After per-bin empty-terrain transfer correction, the reconstructed flux closes to the target Guan spectrum in both angle and energy. This validates the use of the detector-centered log-uniform source for the full-terrain production run. The full-terrain result can then be interpreted as the terrain-induced modification of an already-calibrated incident flux.

4. Virtual Plane and Geometry-Only Response

The configured Geant4 scoring plane is **4 m** × **4 m**. After analysis cuts, the effective accepted virtual area is

$$A_{\mathrm{virtual,eff}} = 2\,\mathrm{m} \times 4\,\mathrm{m} = 8\,\mathrm{m}^2.$$

The real detector active area is

$$A_{\mathrm{det}} = 0.4\,\mathrm{m} \times 0.4\,\mathrm{m} = 0.16\,\mathrm{m}^2.$$

Thus the virtual plane gives a purely geometric statistics gain of

$$\frac{A_{\mathrm{virtual,eff}}}{A_{\mathrm{det}}} = \frac{8}{0.16} = 50.$$

Because the overburden varies on scales much larger than the detector, this increases useful statistics while approximately preserving the line-of-sight thickness image.

Detector model

The detector has four square-pixel planes, plane spacing **0.5 m**, pixel pitch **50 mm**, and active area **0.4 m** × **0.4 m**. The effective output image is about 25×25 angular bins.

Geant4 already transports and scatters muons through the full terrain using the same FTFP_BERT physics configuration and **10 mm** production cut used in the calibration run. Therefore no extra scattering-ring kernel is applied after the Geant4 run. The weighted virtual-plane hits are first binned into an ideal angular image D_{uv}, then folded through a detector response tensor,

$$Y_{mn} = \sum_{u,v} R_{mn,uv} D_{uv}, \qquad R_{mn,uv} = P((m,n) \mid (u,v)).$$

Here R includes finite active area, plane spacing, pixelization, angular acceptance, and hit-pattern overlap only. The angular resolution is therefore a direct consequence of finite detector geometry, not an additional empirical smearing model. Trigger efficiency, electronics response, and empirical reconstruction bias are not included in this tensor.

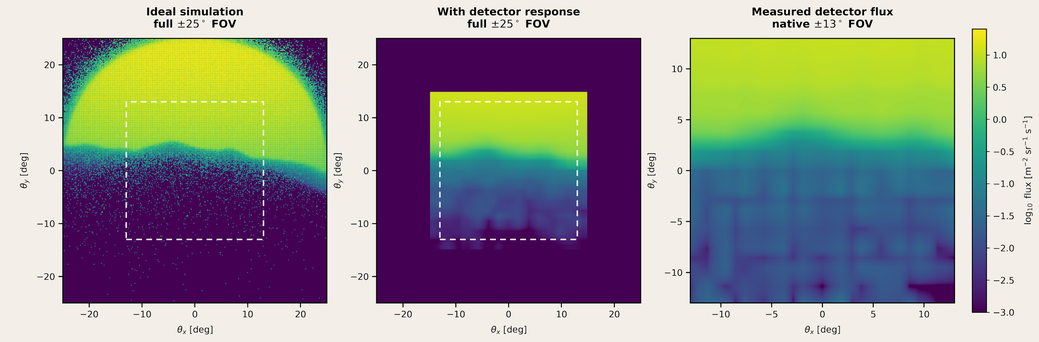

Fig. 3. Ideal angular image, geometry-response-folded image, and measured detector data in the same field of view. Finite detector geometry maps virtual-plane flux into detector space.

5. Full-Terrain Muography Result

The final full-terrain analysis compares reconstructed flux with equivalent overburden thickness along each line of sight. The data shown here are real detector measurements, and the Geant4 prediction is obtained from the calibrated full-terrain virtual-plane flux after geometry-only response folding. Increasing thickness corresponds to longer or denser material paths; the expected trend is therefore a falling transmitted muon flux.

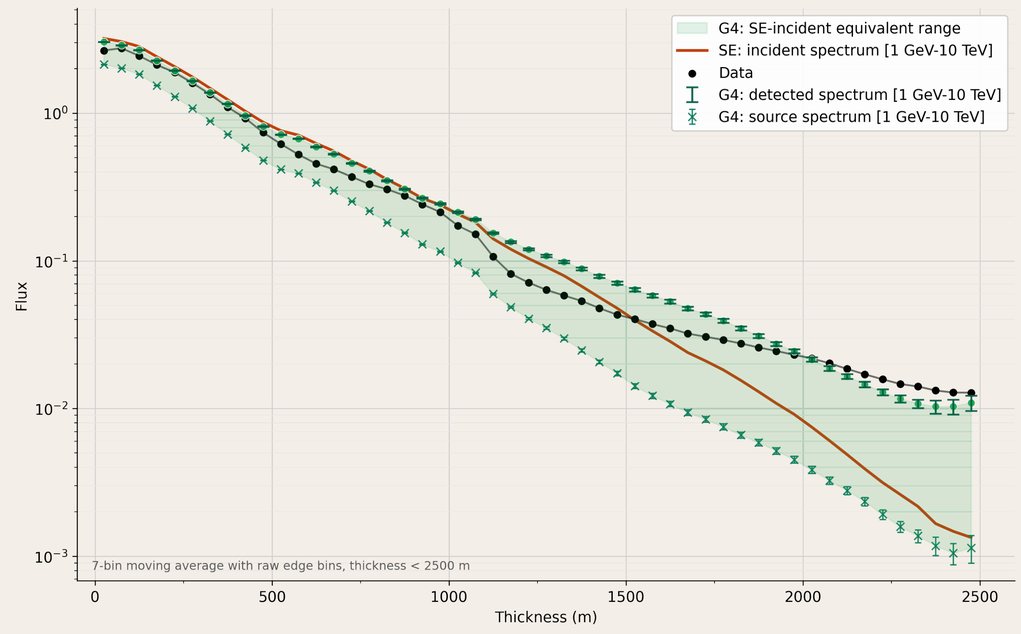

Fig. 4. Flux versus equivalent thickness. The Geant4 full-terrain simulation reproduces the measured flux-vs-thickness trend. Agreement in the trend indicates that the full-scale transport, calibration, and geometry-response stages are mutually consistent.

6. Effective Statistics and Production Scale

For the run-367 readout-style selection, the angular region is $-13^\circ \le \theta_x, \theta_y \le 13^\circ$. The selected sample contains $7{,}231{,}497$ raw events. Events with source zenith above 90° receive zero Guan weight and therefore do not affect the weighted statistics.

Weighted effective exposure

$$N_{\mathrm{eff}} = \frac{(\sum_i w_i)^2}{\sum_i w_i^2} = 651{,}504.$$

$$\frac{\sqrt{\sum_i w_i^2}}{\sum_i w_i} = 1.239 \times 10^{-3} \approx 0.124\%.$$

Using a detector rate of 1700 counts/day,

$$T_{\mathrm{eff}} = \frac{651{,}504}{1700} \approx 383\ \mathrm{days} \approx 1.05\ \mathrm{years}.$$

Computing resources

The production used 150 Geant4 jobs on 15 HTCondor nodes. Each node provided 192 CPU threads:

$$15 \times 192 = 2880$$

concurrent CPU cores for about 5 cluster days. Condor accounting reported 1.39×10^4 CPU-days, and the merged output contains 3.56×10^7 virtual-plane hit entries.

Interpretation and Limitations

The comparison should be read as a geometry-calibrated muography forward prediction. The simulation accounts for full Geant4 terrain transport with FTFP_BERT, a **10 mm** production cut, and detector geometry, but the detector response tensor is intentionally restricted to geometric effects. This separation keeps validation modular: source normalization is checked with the empty-terrain closure plots, detector mapping is checked by geometry folding, and terrain attenuation is tested through the flux-vs-thickness trend.

Primary result: the Geant4 full-terrain simulation reproduces the measured flux-vs-thickness trend. **Secondary results:** the virtual plane provides a $50\times$ geometric acceptance gain, and roughly five cluster days produce ~ 1.05 years of effective detector statistics after reweighting.

End-to-End Workflow

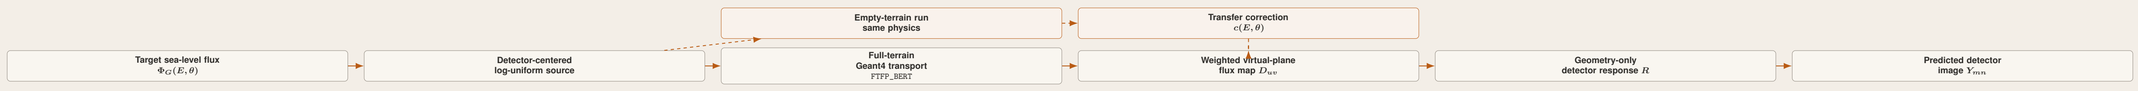

How to Read This Poster

Physical quantity

The plotted observable is not raw Geant4 counts. It is a flux-like quantity obtained after generator reweighting, empty-terrain transfer calibration, virtual-plane binning, and detector-response folding into the finite detector geometry. This is why the simulation can be compared with measured detector flux as a function of equivalent thickness.

Calibration meaning

The empty-terrain closure plots test the source model with the same physics list and production cut used in the full-terrain run. If the empty run reconstructs the Guan flux in energy and angle, then the remaining difference in the full-terrain run is attributed to terrain transport and detector geometry, not to source-placement bias.

Statistics meaning

Because the events are weighted, the raw number of selected hits is not the correct measure of statistical power. The effective count N_{eff} converts the weighted sample into an equivalent unweighted sample with the same relative uncertainty.

Conclusion

This work demonstrates a full-scale Geant4 workflow for cosmic-ray muography through kilometer-scale terrain. The detector-centered source and enlarged virtual plane improve useful statistics, while the empty-terrain calibration restores the physical sea-level flux normalization before interpreting full-terrain attenuation. Using the same FTFP_BERT physics list and **10 mm** production cut in both calibration and production runs keeps the transport configuration consistent. After geometry-only detector-response folding, the simulated flux reproduces the measured flux-versus-thickness trend, showing that the source sampling, transport, calibration, and detector mapping form a consistent forward model for terrain muography.

References. [1] S. Agostinelli et al., "Geant4 - a simulation toolkit," Nucl. Instrum. Meth. A 506, 250-303 (2003). [2] J. Allison et al., "Geant4 developments and applications," IEEE Trans. Nucl. Sci. 53, 270-278 (2006). [3] M. Guan, M.-C. Chu, J. Cao, K.-B. Luk, and C. Yang, "A parametrization of the cosmic-ray muon flux at sea-level," arXiv:1509.06176 (2015).