Efficient propagation of systematic uncertainties from calibration to analysis with the SnowStorm method in IceCube — Source link ⬈

M. G. Aartsen, Markus Ackermann, Jenni Adams, J. A. Aguilar ...+354 more authors

Related papers:

- Energy Reconstruction Methods in the IceCube Neutrino Telescope

- Measurement of South Pole ice transparency with the IceCube LED calibration system

- The design and performance of IceCube DeepCore

- The IceCube Neutrino Observatory: Instrumentation and Online Systems

- Accounting for Calibration Uncertainties in X-ray Analysis: Effective Areas in Spectral Fitting

Efficient propagation of systematic uncertainties from calibration to analysis with the SnowStorm method in IceCube

M. G. Aartsen,[15] M. Ackermann,[53] J. Adams,[15] J. A. Aguilar,[11] M. Ahlers,[19] C. Alispach,[25] B. Al Atoum,[3] K. Andeen,[36] T. Anderson,[50] I. Ansseau,[11] G. Anton,[23] C. Argüelles,[13] J. Auffenberg,[0] S. Axani,[13] P. Backes,[0] H. Bagherpour,[15] X. Bai,[42] A. Balagopal V.,[28] A. Barbano,[25] S. W. Barwick,[27] B. Bastian,[53] V. Baum,[35] S. Baur,[11] R. Bay,[7] J. J. Beatty,[17,18] K.-H. Becker,[52] J. Becker Tjus,[10] S. BenZvi,[44] D. Berley,[16] E. Bernardini,[53,a] D. Z. Besson,[29,b] G. Binder,[7,8] D. Bindig,[52] E. Blaufuss,[16] S. Blot,[53] C. Bohm,[45] M. Börner,[20] S. Böser,[35] O. Botner,[51] J. Böttcher,[0] E. Bourbeau,[19] J. Bourbeau,[34] F. Bradascio,[53] J. Braun,[34] S. Bron,[25] J. Brostean-Kaiser,[53] A. Burgman,[51] J. Buscher,[0] R. S. Busse,[37] T. Carver,[25] C. Chen,[5] E. Cheung,[16] D. Chirkin,[34] S. Choi,[47] K. Clark,[30] L. Classen,[37] A. Coleman,[38] G. H. Collin,[13] J. M. Conrad,[13] P. Coppin,[12] P. Correa,[12] D. F. Cowen,[49,50] R. Cross,[44] P. Dave,[5] C. De Clercq,[12] J. J. DeLaunay,[50] H. Dembinski,[38] K. Deoskar,[45] S. De Ridder,[26] P. Desiati,[34] K. D. de Vries,[12] G. de Wasseige,[12] M. de With,[9] T. DeYoung,[21] A. Diaz,[13] J. C. Díaz-Vélez,[34] H. Dujmovic,[47] M. Dunkman,[50] E. Dvorak,[42] B. Eberhardt,[34] T. Ehrhardt,[35] P. Eller,[50] R. Engel,[28] P. A. Evenson,[38] S. Fahey,[34] A. R. Fazely,[6] J. Felde,[16] K. Filimonov,[7] C. Finley,[45] A. Franckowiak,[53] E. Friedman,[16] A. Fritz,[35] T. K. Gaisser,[38] J. Gallagher,[33] E. Ganster,[0] S. Garrappa,[53] L. Gerhardt,[8] K. Ghorbani,[34] T. Glauch,[24] T. Glüsenkamp,[23] A. Goldschmidt,[8] J. G. Gonzalez,[38] D. Grant,[21] Z. Griffith,[34] S. Griswold,[44] M. Günder,[0] M. Gündüz,[10] C. Haack,[0] A. Hallgren,[51] L. Halve,[0] F. Halzen,[34] K. Hanson,[34] A. Haungs,[28] D. Hebecker,[9] D. Heereman,[11] P. Heix,[0] K. Helbing,[52] R. Hellauer,[16] F. Henningsen,[24] S. Hickford,[52] J. Hignight,[22] G. C. Hill,[1] K. D. Hoffman,[16] R. Hoffmann,[52] T. Hoinka,[20] B. Hokanson-Fasig,[34] K.

Hoshina,[34,c] F. Huang,[50] M. Huber,[24] T. Huber,[28,53] K. Hultqvist,[45] M. Hünnefeld,[20] R. Hussain,[34] S. In,[47] N. Iovine,[11] A. Ishihara,[14] G. S. Japaridze,[4] M. Jeong,[47] K. Jero,[34] B. J. P. Jones,[3] F. Jonske,[0] R. Joppe,[0] D. Kang,[28] W. Kang,[47] A. Kappes,[37] D. Kappesser,[35] T. Karg,[53] M. Karl,[24] A. Karle,[34] U. Katz,[23] M. Kauer,[34] J. L. Kelley,[34] A. Kheirandish,[34] J. Kim,[47] T. Kintscher,[53] J. Kiryluk,[46] T. Kittler,[23] S. R. Klein,[7,8] R. Koirala,[38] H. Kolanoski,[9] L. Köpke,[35] C. Kopper,[21] S. Kopper,[48] D. J. Koskinen,[19] M. Kowalski,[9,53] K. Krings,[24] G. Krückl,[35] N. Kulacz,[22] N. Kurahashi,[41] A. Kyriacou,[1] M. Labare,[26] J. L. Lanfranchi,[50] M. J. Larson,[16] F. Lauber,[52] J. P. Lazar,[34] K. Leonard,[34] A. Leszczyńska,[28] M. Leuermann,[0] Q. R. Liu,[34] E. Lohfink,[35] C. J. Lozano Mariscal,[37] L. Lu,[14] F. Lucarelli,[25] J. Lünemann,[12] W. Luszczak,[34] Y. Lyu,[7,8] W. Y. Ma,[53] J. Madsen,[43] G. Maggi,[12] K. B. M. Mahn,[21] Y. Makino,[14] P. Mallik,[0] K. Mallot,[34] S. Mancina,[34] I. C. Mariş,[11] R. Maruyama,[39] K. Mase,[14] R. Maunu,[16] F. McNally,[32] K. Meagher,[34] M. Medici,[19] A. Medina,[18] M. Meier,[20] S. Meighen-Berger,[24] T. Menne,[20] G. Merino,[34] T. Meures,[11] J. Micallef,[21] D. Mockler,[11] G. Momenté,[35] T. Montaruli,[25] R. W. Moore,[22] R. Morse,[34] M. Moulai,[13] P. Muth,[0] R. Nagai,[14] U. Naumann,[52] G. Neer,[21] H. Niederhausen,[24] S. C. Nowicki,[21] D. R. Nygren,[8] A. Obertacke Pollmann,[52] M. Oehler,[28] A. Olivas,[16] A. O'Murchadha,[11] E. O'Sullivan,[45] T. Palczewski,[7,8] H. Pandya,[38] D. V. Pankova,[50] N. Park,[34] P. Peiffer,[35] C. Pérez de los Heros,[51] S. Philippen,[0] D. Pieloth,[20] E. Pinat,[11] A. Pizzuto,[34] M. Plum,[36] A. Porcelli,[26] P. B. Price,[7] G. T. Przybylski,[8] C. Raab,[11] A. Raissi,[15] M. Rameez,[19] L. Rauch,[53] K. Rawlins,[2] I. C. Rea,[24] R. Reimann,[0] B. Relethford,[41] M. Renschler,[28] G. Renzi,[11] E. Resconi,[24] W. Rhode,[20] M. Richman,[41] S. Robertson,[8] M. Rongen,[0] C. Rott,[47] T. Ruhe,[20] D. Ryckbosch,[26] D. Rysewyk,[21] I. Safa,[34] S. E. Sanchez Herrera,[21] A. Sandrock,[20] J. Sandroos,[35] M. Santander,[48] S. Sarkar,[40] S. Sarkar,[22] K. Satalecka,[53] M. Schaufel,[0] H. Schieler,[28] P. Schlunder,[20] T. Schmidt,[16] A. Schneider,[34] J. Schneider,[23] F. G. Schröder,[28,38] L. Schumacher,[0] S. Sclafani,[41] D. Seckel,[38] S. Seunarine,[43] S. Shefali,[0] M. Silva,[34] R. Snihur,[34] J. Soedingrekso,[20] D. Soldin,[38] M. Song,[16] G. M. Spiczak,[43] C. Spiering,[53] J. Stachurska,[53] M. Stamatikos,[18] T. Stanev,[38] R. Stein,[53] P. Steinmüller,[28] J. Stettner,[0] A. Steuer,[35] T. Stezelberger,[8] R. G. Stokstad,[8] A. Stößl,[14] N. L. Strotjohann,[53] T. Stürwald,[0] T. Stuttard,[19] G. W. Sullivan,[16] I. Taboada,[5] F.

Tenholt,[10] **S. Ter-Antonyan,**[6] **A. Terliuk,**[53] **S. Tilav,**[38] **K. Tollefson,**[21] **L. Tomankova,**[10] **C. Tönnis,**[47] **S. Toscano,**[11] **D. Tosi,**[34] **A. Trettin,**[53] **M. Tselengidou,**[23] **C. F. Tung,**[5] **A. Turcati,**[24] **R. Turcotte,**[28] **C. F. Turley,**[50] **B. Ty,**[34] **E. Unger,**[51] **M. A. Unland Elorrieta,**[37] **M. Usner,**[53] **J. Vandenbroucke,**[34] **W. Van Driessche,**[26] **D. van Eijk,**[34] **N. van Eijndhoven,**[12] **S. Vanheule,**[26] **J. van Santen,**[53] **M. Vraeghe,**[26] **C. Walck,**[45] **A. Wallace,**[1] **M. Wallraff,**[0] **N. Wandkowsky,**[34] **T. B. Watson,**[3] **C. Weaver,**[22] **A. Weindl,**[28] **M. J. Weiss,**[50] **J. Weldert,**[35] **C. Wendt,**[34] **J. Werthebach,**[34] **B. J. Whelan,**[1] **N. Whitehorn,**[31] **K. Wiebe,**[35] **C. H. Wiebusch,**[0] **L. Wille,**[34] **D. R. Williams,**[48] **L. Wills,**[41] **M. Wolf,**[24] **J. Wood,**[34] **T. R. Wood,**[22] **K. Woschnagg,**[7] **G. Wrede,**[23] **D. L. Xu,**[34] **X. W. Xu,**[6] **Y. Xu,**[46] **J. P. Yanez,**[22] **G. Yodh,**[27] **S. Yoshida,**[14] **T. Yuan**[34] **and M. Zöcklein**[0]

[0]III. Physikalisches Institut, RWTH Aachen University, D-52056 Aachen, Germany

[1]Department of Physics, University of Adelaide, Adelaide, 5005, Australia

[2]Dept. of Physics and Astronomy, University of Alaska Anchorage, 3211 Providence Dr., Anchorage, AK 99508, USA

[3]Dept. of Physics, University of Texas at Arlington, 502 Yates St., Science Hall Rm 108, Box 19059, Arlington, TX 76019, USA

[4]CTSPS, Clark-Atlanta University, Atlanta, GA 30314, USA

[5]School of Physics and Center for Relativistic Astrophysics, Georgia Institute of Technology, Atlanta, GA 30332, USA

[6]Dept. of Physics, Southern University, Baton Rouge, LA 70813, USA

[7]Dept. of Physics, University of California, Berkeley, CA 94720, USA

[8]Lawrence Berkeley National Laboratory, Berkeley, CA 94720, USA

[9]Institut für Physik, Humboldt-Universität zu Berlin, D-12489 Berlin, Germany

[10]Fakultät für Physik & Astronomie, Ruhr-Universität Bochum, D-44780 Bochum, Germany

[11]Université Libre de Bruxelles, Science Faculty CP230, B-1050 Brussels, Belgium

[12]Vrije Universiteit Brussel (VUB), Dienst ELEM, B-1050 Brussels, Belgium

[13]Dept. of Physics, Massachusetts Institute of Technology, Cambridge, MA 02139, USA

[14]Dept. of Physics and Institute for Global Prominent Research, Chiba University, Chiba 263-8522, Japan

[15]Dept. of Physics and Astronomy, University of Canterbury, Private Bag 4800, Christchurch, New Zealand

[16]Dept. of Physics, University of Maryland, College Park, MD 20742, USA

[17]Dept. of Astronomy, Ohio State University, Columbus, OH 43210, USA

[18]Dept. of Physics and Center for Cosmology and Astro-Particle Physics, Ohio State University, Columbus, OH 43210, USA

[19]Niels Bohr Institute, University of Copenhagen, DK-2100 Copenhagen, Denmark

[20]Dept. of Physics, TU Dortmund University, D-44221 Dortmund, Germany

[21]Dept. of Physics and Astronomy, Michigan State University, East Lansing, MI 48824, USA

[22] Dept. of Physics, University of Alberta, Edmonton, Alberta, Canada T6G 2E1

[23] Erlangen Centre for Astroparticle Physics, Friedrich-Alexander-Universität Erlangen-Nürnberg, D-91058 Erlangen, Germany

[24] Physik-department, Technische Universität München, D-85748 Garching, Germany

[25] Département de physique nucléaire et corpusculaire, Université de Genève, CH-1211 Genève, Switzerland

[26] Dept. of Physics and Astronomy, University of Gent, B-9000 Gent, Belgium

[27] Dept. of Physics and Astronomy, University of California, Irvine, CA 92697, USA

[28] Karlsruhe Institute of Technology, Institut für Kernphysik, D-76021 Karlsruhe, Germany

[29] Dept. of Physics and Astronomy, University of Kansas, Lawrence, KS 66045, USA

[30] SNOLAB, 1039 Regional Road 24, Creighton Mine 9, Lively, ON, Canada P3Y 1N2

[31] Department of Physics and Astronomy, UCLA, Los Angeles, CA 90095, USA

[32] Department of Physics, Mercer University, Macon, GA 31207-0001

[33] Dept. of Astronomy, University of Wisconsin, Madison, WI 53706, USA

[34] Dept. of Physics and Wisconsin IceCube Particle Astrophysics Center, University of Wisconsin, Madison, WI 53706, USA

[35] Institute of Physics, University of Mainz, Staudinger Weg 7, D-55099 Mainz, Germany

[36] Department of Physics, Marquette University, Milwaukee, WI, 53201, USA

[37] Institut für Kernphysik, Westfälische Wilhelms-Universität Münster, D-48149 Münster, Germany

[38] Bartol Research Institute and Dept. of Physics and Astronomy, University of Delaware, Newark, DE 19716, USA

[39] Dept. of Physics, Yale University, New Haven, CT 06520, USA

[40] Dept. of Physics, University of Oxford, Parks Road, Oxford OX1 3PU, UK

[41] Dept. of Physics, Drexel University, 3141 Chestnut Street, Philadelphia, PA 19104, USA

[42] Physics Department, South Dakota School of Mines and Technology, Rapid City, SD 57701, USA

[43] Dept. of Physics, University of Wisconsin, River Falls, WI 54022, USA

[44] Dept. of Physics and Astronomy, University of Rochester, Rochester, NY 14627, USA

[45] Oskar Klein Centre and Dept. of Physics, Stockholm University, SE-10691 Stockholm, Sweden

[46] Dept. of Physics and Astronomy, Stony Brook University, Stony Brook, NY 11794-3800, USA

[47] Dept. of Physics, Sungkyunkwan University, Suwon 16419, Korea

[48] Dept. of Physics and Astronomy, University of Alabama, Tuscaloosa, AL 35487, USA

[49] Dept. of Astronomy and Astrophysics, Pennsylvania State University, University Park, PA 16802, USA

[50] Dept. of Physics, Pennsylvania State University, University Park, PA 16802, USA

[51] Dept. of Physics and Astronomy, Uppsala University, Box 516, S-75120 Uppsala, Sweden

[52] Dept. of Physics, University of Wuppertal, D-42119 Wuppertal, Germany

[53] DESY, D-15738 Zeuthen, Germany

[a] also at Università di Padova, I-35131 Padova, Italy

[b] also at National Research Nuclear University, Moscow Engineering Physics Institute (MEPhI), Moscow 115409, Russia

^{c}Earthquake Research Institute, University of Tokyo, Bunkyo, Tokyo 113-0032, Japan

E-mail: analysis@icecube.wisc.edu

Abstract. Efficient treatment of systematic uncertainties that depend on a large number of nuisance parameters is a persistent difficulty in particle physics experiments. Where low-level effects are not amenable to simple parameterization or re-weighting, analyses often rely on discrete simulation sets to quantify the effects of nuisance parameters on key analysis observables. Such methods may become computationally untenable for analyses requiring high statistics Monte Carlo with a large number of nuisance degrees of freedom, especially in cases where these degrees of freedom parameterize the shape of a continuous distribution. In this paper we present a method for treating systematic uncertainties in a computationally efficient and comprehensive manner using a single simulation set with multiple and continuously varied nuisance parameters. This method is demonstrated for the case of the depth-dependent effective dust distribution within the IceCube Neutrino Telescope.

Contents

1 Introduction

A pervasive challenge in experimental particle physics is the development of robust treatments for systematic uncertainties. There is no single correct approach to treating systematic uncertainties in any analysis, but it is nevertheless recognized that a credible parameterization of ignorance about the experiment being undertaken is a critical ingredient in any measurement. Furthermore, for analyses yielding unified confidence intervals or Bayesian intervals in multi-dimensional spaces, neither an overly tight nor an overly loose parameterization of ignorance can generally be considered to be conservative - either may enhance the likelihood of a spurious signal or overly strong exclusion under certain circumstances, and so it is vital to parameterize ignorance as accurately as possible. Given that no parameterization of uncertainty is exact, and that un-modelled effects and simplifying assumptions will always be present, treatment of systematic uncertainties is be a frequent source of difficulty in particle physics and beyond.

The goal of this work is to present a method used at the IceCube South Pole Neutrino Observatory [1, 2] to propagate a particularly complex form of systematic uncertainty from calibration constraints to analyses of atmospheric neutrinos, including searches for sterile neutrinos [3]. The uncertainty in question is that of the effective dust concentration as a function of depth within the ice of the IceCube neutrino telescope. For the purpose of this work, the term "dust" will be used to refer to all forms of optical impurity within the ice, although mineral dust itself is only one major component.

The dust distribution within IceCube and the methods used to characterize it have been described at length elsewhere [4] and ongoing campaigns continue to provide deeper understanding of its optical properties [5–10]. In past generations of IceCube analyses, the detailed properties of the bulk ice in the array could be shown to be a sub-dominant concern (for example, in Ref [11]), with their effect on analysis distributions being comfortably less than statistical uncertainty in most samples. The latest generation of atmospheric neutrino

analyses at IceCube, on the other hand, consider samples containing more than 250,000 muon neutrino candidates [12]. For these analyses, percent-scale precision on event distributions is required. This mandates a detailed understanding of the effect of uncertainty on ice properties such as dust distribution within analyses.

The reasons why the dust distribution in IceCube is a particularly challenging uncertainty to parameterize include the following:

- The effect of the dust concentration vs depth on variables such as reconstructed muon energy and zenith angle is highly non-trivial to estimate. The only known way to establish this effect is through Monte Carlo simulation of an event within a given ice model, and it is computationally unfeasible to generate Monte Carlo samples given every reasonable ice model configuration.

- The dust concentration varies continuously as a function of depth. The uncertainty on this continuous function must be constrained and propagated without introducing an overwhelming number of nuisance parameters.

These two challenges are not unique to this source of uncertainty or to IceCube, and the method presented here may be used to treat uncertainties on continuous functions in highly multidimensional spaces in other contexts. The organization of this paper is such that we treat these two challenges in order. The first challenge is addressed in Sec. 2, via construction of the SnowStorm method, a Monte Carlo technique whereby constraints on many nuisance variables can be mapped from calibration data to analysis space in a computationally efficient manner. The second challenge is then addressed in Sec. 3, focusing on the specific example of the effective dust distribution in IceCube. A Fourier parameterization is developed that allows for propagation of the most important uncertainty contributions from calibration to analysis space, and the SnowStorm approach is applied to generate a covariance matrix for use in physics analysis.

2   The SnowStorm MC Method

2.1   Measurements and systematic uncertainties

A wide class of measurements in particle physics are made by comparing some observed distribution of events $\phi(E_\alpha)$ - for example, a measured neutrino spectrum as a function of energy - against a predicted distribution $\Psi_{\vec{\rho}}(E_\alpha)$, generally derived from calculations and Monte Carlo simulations, informed by calibration data. The distribution to be measured is a function of some variable E_α, where α may be either a simple index (for example, labelling energy bin number), a compound index (for example, energy and zenith bin numbers), or a continuous variable (in a non-binned analysis). The prediction depends on some "physics parameters" to be measured $\vec{\rho}$, and carries some uncertainty on associated with systematic and statistical effects. In the regime where Gaussian propagation of errors is valid, the uncertainty on $\Psi_{\vec{\rho}}(E_\alpha)$ can be completely described as a covariance matrix in the analysis space, $\Sigma_{\vec{\rho}}(E_\alpha, E_\beta)$. A measurement of $\vec{\rho}$ is made by comparing $\phi(E_\alpha)$ with $\Psi_{\vec{\rho}}(E_\alpha)$ given $\Sigma_{\vec{\rho}}(E_\alpha, E_\beta)$ at each $\vec{\rho}$. Exactly what the outcome of this measurement represents and how the final result is constructed depends on whether a frequentist or Bayesian probability construction is preferred, and we will not review this topic here. What is important is that in either case, the proper estimation of $\Sigma_{\vec{\rho}}(E_\alpha, E_\beta)$ is central to calculation of the final result.

Any effect that influences the analysis distribution and that depends on some parameters that are imperfectly known will contribute to $\Sigma_{\vec{\rho}}(E_\alpha, E_\beta)$. The unknown parameters on which this effect depends are typically called "nuisance parameters" and we represent them as a vector $\vec{\eta}$. Given some choice of $\vec{\eta}$, the predicted event distribution can be calculated, and we denote this as $\psi_{\vec{\rho},\vec{\eta}}(E_\alpha)$. Choosing coordinates such that the "central" prediction is one corresponding to $\vec{\eta} = \vec{0}$, we impose the definition $\Psi_{\vec{\rho}}(E_\alpha) \equiv \psi_{\vec{\rho},\vec{0}}(E_\alpha)$.

Broadly speaking, there are two typical things an analyzer may want to know about the effects of the nuisance parameters on the prediction:

1. What is the effect of varying $\vec{\eta}$ on the predicted distribution of events? In other words: how does $\psi_{\vec{\rho},\vec{\eta}}(E_\alpha)$ depend on $\vec{\eta}$?

2. Given some set of constraints on $\vec{\eta}$ from calibration data or other experiments, what is the uncertainty on $\Psi_{\vec{\rho}}(E_\alpha)$? In other words, given available knowledge of $\vec{\eta}$, what is $\Sigma_{\vec{\rho}}(E_\alpha, E_\beta)$?

For some nuisance parameters the dependence of $\psi_{\vec{\rho},\vec{\eta}}(E_\alpha)$ upon η_i is straightforward to model, and techniques such as re-weighting can be used to answer the first question. However, for other classes of systematic uncertainty, for example, effects that change detector response in non-trivial ways, the effect on final analysis-level distributions is highly non-trivial to estimate. A common approach for these classes of uncertainty involves production of many distinct Monte Carlo sets with different choices of nuisance parameters. This approach, sometimes described as a "multiple Universes" scheme, becomes quickly computationally untenable as the dimensionality of $\vec{\eta}$ becomes large, especially for high-statistics samples where the corresponding Monte Carlo samples must have high event counts. It can also lead to neglecting of important correlations between different η_i variables.

For this reason and others, enumeration of a suitable suite of $\vec{\eta}$ parameters and answering of the two questions posed above is often the most time-consuming aspect of a carefully constructed particle physics analysis. The SnowStorm Monte Carlo method described in this paper is a recipe for answering both questions (1) and (2), given an arbitrarily large number of $\vec{\eta}$ parameters using a single Monte Carlo ensemble.

2.2   The SnowStorm Monte Carlo ensemble

A SnowStorm Monte Carlo ensemble[1] consists of a single simulation set where every event (or every group of a few events) is generated with a unique set of nuisance parameters $\vec{\eta}$. Under some loose assumptions, given a suitably prepared SnowStorm Monte Carlo sample, it is possible to assess the effects of these variations on an analysis by appropriately manipulating this ensemble. The structure of the SnowStorm ensemble in $\vec{\eta}$ space is compared schematically to a "multiple Universes" based approach in Fig. 1, left.

An assumption underlying the SnowStorm method is that the effects of systematic uncertainties on analysis variables are sufficiently small that they can be treated perturbatively. Such an assumption also underlies, explicitly or implicitly, the majority of methods for systematic uncertainty estimation and error analysis in common use. We also assume that the statistical uncertainty on Monte Carlo event counts is very small compared to that on the data. Under such conditions, summation over a Monte Carlo ensemble with different values

---

[1]The name SnowStorm reflects the idea that each event in the sample is distinct, like each snowflake in a snowstorm
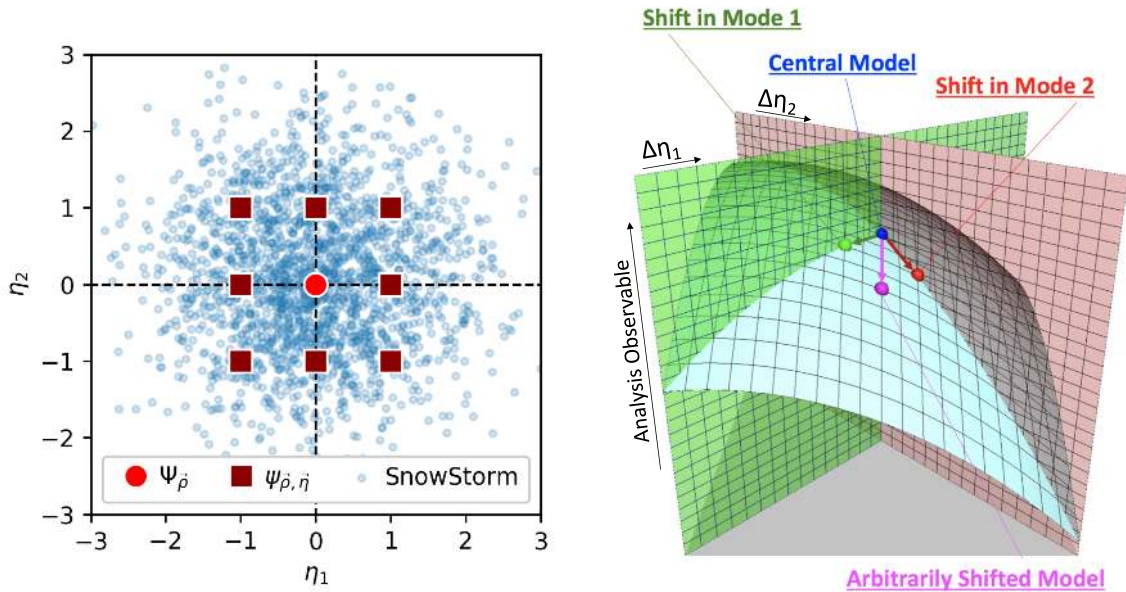
Figure 1. Left: Cartoon illustrating the difference between many discrete shifts in nuisance parameters (red squares) – each requiring an entire Monte Carlo set – compared with a single SnowStorm Monte Carlo (small blue dots). Right: the method of extracting independent gradients and building up a covariance matrix is accurate in the limit that the analysis space is locally flat as the nuisance parameters are varied. This corresponds to linear perturbativity.

of $\vec{\eta}$ distributed according to a distribution function $P(\vec{\eta})$ can be replaced by an integration:

$$\sum_{all\ \vec{\eta}} \rightarrow \int d^N\eta P(\vec{\eta}). \tag{2.1}$$

The distribution function $P(\vec{\eta})$ is specified when generating the Monte Carlo ensemble. We enforce on $P(\vec{\eta})$ two properties: it is normalized and it is symmetric in every variable η_i such that:

$$P(\eta_i, \vec{\eta}_{j\neq i}) = P(-\eta_i, \vec{\eta}_{j\neq i}) \quad \forall i, \tag{2.2}$$

$$\int d^N\eta\, P(\vec{\eta}) = 1. \tag{2.3}$$

Much of what follows conceptually relies only on these assumptions. For simplicity, however, we make an explicit choice of $P(\vec{\eta})$ choosing a product of Gaussian functions of equal width σ:

$$P(\vec{\eta}) = \prod_i \frac{1}{\sigma\sqrt{2\pi}} e^{-\eta_i^2/2\sigma^2}. \tag{2.4}$$

The equal-width requirement corresponds to a choice of coordinates for $\vec{\eta}$, rather than a constraint on the size of any physical effects to be modelled.

If the effects of the nuisance parameters $\vec{\eta}$ are perturbative, this implies that the distribution function at any $\vec{\eta}$ can be written as a Taylor expansion around the central distribution:

$$\psi_{\vec{\rho},\vec{\eta}} = \Psi_{\vec{\rho}} + \vec{\eta}.\vec{\nabla}_\eta \left[\psi_{\vec{\rho},\vec{\eta}}\right]_{\vec{\eta}=\vec{0}} + \mathcal{O}(\eta^2), \tag{2.5}$$

with the largest effects present in the first term. Neglecting the $\mathcal{O}(\eta^2)$ terms implies making a locally Euclidian approximation to the nuisance space as illustrated in Fig 1, right. This is appropriate for a wide class of effects, and always true for sufficiently small perturbations in $\vec{\eta}$. This assumption should not be taken for granted, however. Its validity can be tested given availability of both a central Monte Carlo set $\Psi_{\vec{\rho}}$ and a SnowStorm ensemble, by considering the integrated prediction of the SnowStorm sample:

$$\psi_{\vec{\rho}}^{SnowStorm} = \int d\vec{\eta} P(\vec{\eta}) \psi_{\vec{\rho},\vec{\eta}} = \int d\vec{\eta} P(\vec{\eta}) \left[\Psi_{\vec{\rho}} + \vec{\eta}.\vec{\nabla}_\eta \left[\psi_{\vec{\rho},\vec{\eta}} \right]_{\vec{\eta}=\vec{0}} + \mathcal{O}(\eta^2) \right]. \qquad (2.6)$$

The integral over the distribution function in the first term normalizes to 1 since $\int d\vec{\eta} P(\vec{\eta}) = 1$. The second term is an odd function of $\vec{\eta}$ integrated between even limits and as such gives no contribution; The final term encodes all the residual higher order effects. Thus:

$$\psi_{\vec{\rho}}^{SnowStorm} = \Psi_{\vec{\rho}} + \int d\vec{\eta} P(\eta) \mathcal{O}(\eta^2). \qquad (2.7)$$

The prediction of the SnowStorm ensemble and the central model are identical up to $\mathcal{O}(\eta^2)$ effects. If comparison shows they are equivalent within the available statistical uncertainty, it is appropriate to neglect the non-linear terms in $\vec{\eta}$ to this same precision. Assuming this test passes and the linear and perturbative approximation is a good one, Eq. 2.5 reduces to:

$$\psi_{\vec{\rho},\vec{\eta}} = \Psi_{\vec{\rho}} + \vec{\eta}.\vec{G}_{\vec{\rho}}, \qquad \vec{G}_{\vec{\rho}} \equiv \vec{\nabla}_\eta \left[\psi_{\vec{\rho},\vec{\eta}} \right]_{\vec{\eta}=\vec{0}}. \qquad (2.8)$$

The prediction of any model in η space can be obtained given $\Psi_{\vec{\rho}}$ and a list of N "nuissance gradient" coefficients $\vec{G}_{\vec{\rho}}$.

For analyses where the effects of the physics parameters are relatively small in the region of interest, it is possible to make a further approximation that $\vec{G}_{\vec{\rho}} = \vec{G}_{\vec{0}}$, since the relevant dependencies on $\vec{\rho}$ would be second order perturbative effects in Eq. 2.8. This reduces computational load to calculating only one universal set of N gradients, rather than one set for each value of the physics parameters $\vec{\rho}$. This assumption is not a necessary condition for use of the SnowStorm method, however, and its suitability should be assessed on a per-case basis.

2.3 Nuisance gradient extraction from the SnowStorm ensemble

Under the assumption of linearity and perturbativity (testable by comparing the SnowStorm ensemble to a central model), a method of extracting the N nuisance gradients G_i would provide an answer to Question (1) of Sec 2.1: that is, the effect of making small changes in each of the nuisance parameters on the distribution of the analysis observable. There are two approaches to extracting nuisance gradients, which will now be described. In the limit of infinite statistical precision and perfect linearity, these two methods should provide identical results. In subsequent analysis will use only the first approach, which offers better statistical precision on $\vec{G}$ for our case of interest.

First, consider making a cut of the SnowStorm ensemble along one direction $\eta_i = 0$. The sample is then divided into two sub-samples, $\psi_{\vec{\rho}}^{i+}$ and $\psi_{\vec{\rho}}^{i-}$, which each encode predictions:

$$\psi_{\vec{\rho}}^{i+} = \int_0^\infty d\eta_i \int_{-\infty}^\infty d^{N-1}\eta P(\vec{\eta}) \left[\Psi_{\vec{\rho}} + \vec{\eta}.\vec{\nabla}_\eta \left[\psi_{\vec{\rho},\vec{\eta}} \right]_{\vec{\eta}=\vec{0}} \right], \qquad (2.9)$$

$$\psi_{\vec{\rho}}^{i-} = \int_{-\infty}^0 d\eta_i \int_{-\infty}^\infty d^{N-1}\eta P(\vec{\eta}) \left[\Psi_{\vec{\rho}} + \vec{\eta}.\vec{\nabla}_\eta \left[\psi_{\vec{\rho},\vec{\eta}} \right]_{\vec{\eta}=\vec{0}} \right]. \qquad (2.10)$$
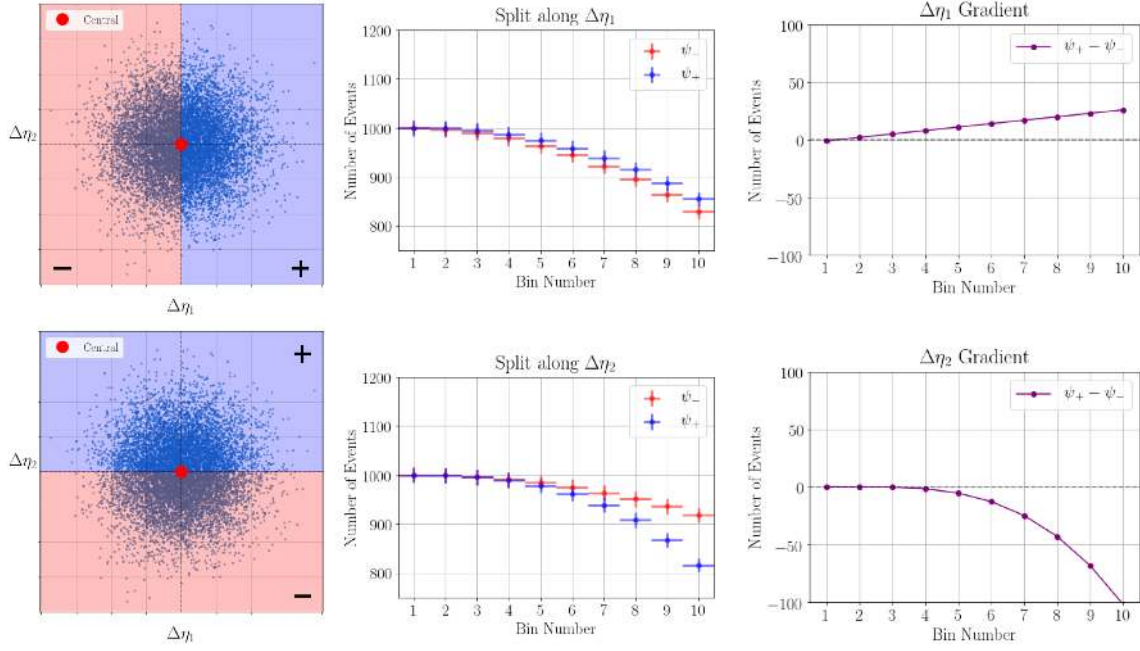
Figure 2. Cartoon illustrating procedure for extracting gradients. The distribution is split along a given nuisance parameter direction into positively perturbed and negatively perturbed sets. The gradients are then proportional to the difference of the two sets.

The $N - 1$ integrals along directions other than i give no contribution from the linear term, by symmetry. However, symmetry of the integral along the i direction is broken by the cut, and a finite contribution can be calculated by evaluating the integral over $\eta_i P(\eta)$ in this half space, yielding:

$$\psi_{\vec{\rho}}^{i\pm} = \frac{1}{2}\Psi_{\vec{\rho}} \pm \frac{\sigma}{\sqrt{2\pi}}G_{\vec{\rho},i}. \tag{2.11}$$

The gradient function $G_{\vec{\rho},i}$ is then straightforwardly extracted by subtraction:

$$G_{\vec{\rho},i} = \frac{1}{\sigma}\sqrt{\frac{\pi}{2}}\left(\psi_{\vec{\rho}}^{i+} - \psi_{\vec{\rho}}^{i-}\right). \tag{2.12}$$

Cutting the SnowStorm ensemble in each direction one by one, the N dimensional vector $\vec{G}_{\vec{\rho}} \equiv \vec{\nabla}_\eta\left[\psi_{\vec{\rho},\vec{\eta}}\right]_{\vec{\eta}=\vec{0}}$ can be extracted completely from one Monte Carlo ensemble.

There is a second approach to extract $\vec{G}$, involving weighting rather than cutting. In this case, consider constructing a prediction where each event in the SnowStorm ensemble is weighted by a factor of η_i, for one chosen direction in $\vec{\eta}$ space. This prediction takes the form:

$$\tilde{\psi}_{\vec{\rho}}^i = \int_{-\infty}^{\infty} d^N\eta\,\eta_i P(\vec{\eta})\left[\Psi_{\vec{\rho}} + \vec{\eta}.\vec{\nabla}_\eta\left[\psi_{\vec{\rho},\vec{\eta}}\right]_{\vec{\eta}=\vec{0}}\right]. \tag{2.13}$$

Now, the left term integrates to zero by symmetry, as do all the terms in the dot product except the one for the i component, which leaves a residual contribution due to its even integrand:

$$\tilde{\psi}_{\vec{\rho}}^i = \frac{\sigma^2}{2}G_{\vec{\rho},i}. \tag{2.14}$$

Which is trivially inverted to obtain a second, different expression for the elements of $\vec{G}$. We thus see that either by weighting or cutting the sample, the full nuisance gradient vector $\vec{G}$ can be extracted, and from this, the prediction of any nearby model in $\vec{\eta}$ space can be obtained via Eq. 2.8.

Our experience has shown that somewhat better statistical precision on $\vec{G}$ is obtained using the cutting method, since the weighting method exaggerates the contribution of events in the tails of the P distribution, leading to larger fluctuations. A schematic illustration of the cutting method applied on a two-dimensional $\vec{\eta}$ space is shown in Fig. 2.

2.4   Propagation of calibration constraints and nuisance priors to analysis space

Implicitly, all distribution functions ϕ, Ψ, ψ and the gradients $\vec{G}$ are functions of the analysis variable E_α, though we had suppressed this dependence for notational convenience. Here we restore it, as:

$$G_i(E_\alpha) \equiv G_{i;\alpha}, \quad \psi_{\vec{\rho},\vec{\eta}}(E_\alpha) \equiv \psi_{\vec{\rho},\vec{\eta};\alpha}, \quad \text{etc...} \tag{2.15}$$

In general, the nuisance parameters $\vec{\eta}$ will not be totally unknown, but subject to some constraints either from calibration data or from previous measurements. These constrains can be represented in the form of a covariance matrix Ξ, which would be diagonal if all η_i were independent, but generally will not be if the $\vec{\eta}$ values are constrained by common calibration data. A basic theorem (proof reviewed in Appendix) can be invoked to map Ξ, the covariance matrix on the nuisance parameters $\vec{\eta}$, to Σ, the covariance matrix in the binned analysis space:

$$\Sigma_{\alpha,\beta} = G_{i;\alpha} \Xi_{ij} G_{j;\beta}, \tag{2.16}$$

where $\vec{G}$ is the gradient vector (representing the effects of small shifts in the nuisance parameters on analysis-space distributions) derived in the previous section. Thus, given constraints on $\vec{\eta}$ which may or may not involve correlations, and the gradient vector $\vec{G}$, the analysis covariance $\Sigma_{\alpha,\beta}$ can be extracted, providing an answer to Question 2 of Sec 2.1. Exactly how the matrix Ξ can be derived will vary by experiment and by effect. It is common to use analysis side-bands or calibration data for this purpose, to constrain the nuisance parameters in the sample. An illustrative example where an independent calibration data set is used to construct Ξ will be presented in the following section.

3   The SnowStorm method applied to the IceCube dust distribution

The optical properties of the ice at the South Pole influence the propagation and detection of photons within the IceCube detector. It has long been established that the ice has properties that are depth dependent and correlate with the climate history on Earth over the period of its accumulation, of approximately 100,000 years [13, 14]. The primary optical effects that influence photon propagation are absorption and scattering, and these are observed to exhibit a strong depth dependent correlation with one-another, suggesting a common origin. This origin is generally understood to be small concentrations of dust throughout the array. Since the dust concentration is very low, the dust distribution is best probed via its effects on light propagation, using both dedicated dust-logger devices [14] and in-situ calibration using LED flashers within the IceCube array [4]. Based on this calibration data, a layered ice optical model has been constructed, with absorption and scattering coefficients assigned for each 10m layer, their values continuing to be refined as understanding of the ice properties improves through evolving calibration techniques and the incorporation of increasingly comprehensive
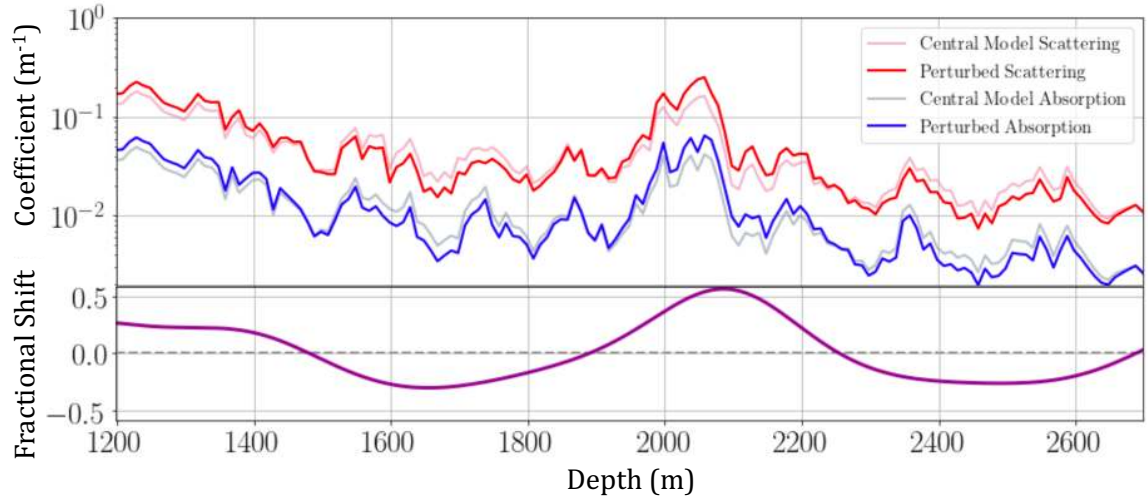
Figure 3. An illustrative (but relatively large) perturbation of the ice parameters demonstrating depth dependant features. This example uses a relative shift of +50% in amplitudes 2 and 5 and -50% in amplitude 3.

catalogues of optical effects. The structure and tuning of the ice model is described at length elsewhere, and a detailed exposition of its construction is outside the scope of this work. The absorption and scattering coefficients from a recent generation of IceCube ice model are shown in Fig 3, as well as an example of a perturbed model illustrating the continuous freedom in these parameters which must be explored to properly account for its uncertainty.

   IceCube analyses typically rely on comparison of data to Monte Carlo, with this Monte Carlo ideally generated using the most up-to-date generation of ice model available. As an illustration of the SnowStorm method we present its application to a specific form of ice uncertainty- that induced by imprecise knowledge depth-dependent distribution of dust inducing scattering and absorption within the IceCube array. We do not treat here the uncertainty on the form or composition of dust (i.e. the possibility that some areas of dust could create more absorption and less scattering, for example), uncertainty on the source and form of optical anisotropy [9], uncertainty on the properties of refrozen ice in the vicinity of the Digital Optical Modules, or other possible sources of un-modeled uncertainty. The resulting uncertainty estimate should thus not be considered as a comprehensive uncertainty budget for all IceCube optical effects, or as a strong statement about glaciology. Rather, we seek to construct a method for propagating a known, significant form of optical uncertainty for high-statistics atmospheric neutrino analyses in IceCube, that has been hitherto very difficult to assess.

3.1 Model parameter reduction via Fourier analysis

Due to the large number of parameters in the ice model combined with the substantial computational load required to accurately simulate light propagation within the ice, traditionally the most efficient means of assessing the effects of ice uncertainties on high-energy and high-statistics analyses has been to create discrete simulation sets where the absolute scales of bulk ice properties (e.g. scattering and absorption) are shifted by the same fraction across the entire detector. This method is fundamentally restricted in its ability to accurately assess depth-dependent contributions to systematic uncertainty.
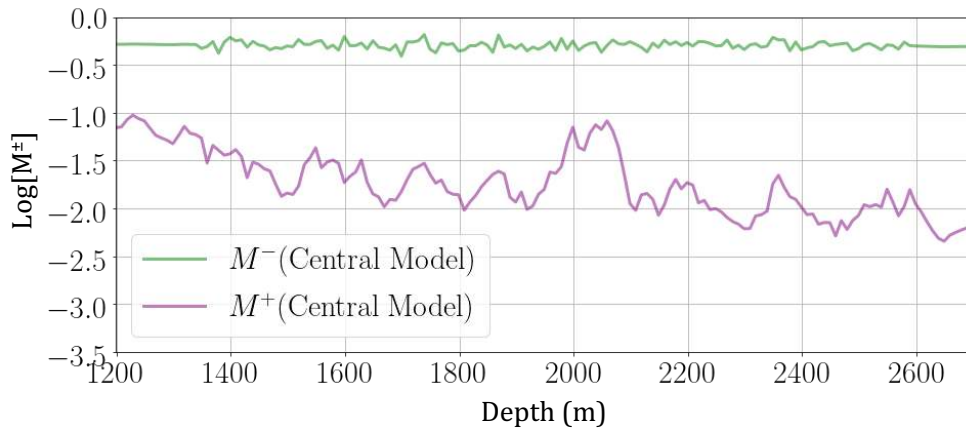
Figure 4. The spectrum to be Fourier decomposed from central model for the re-parametrization. Note the flatness of M^- indicating a strong correlation between scattering an absorption over all depths.

One possibility would be to vary individual layers one by one, and consider the scale of these variations as entries in the $\vec{\eta}$ vector. However, it is clear that varying only a single layer can have at most a very modest effect on the distribution of analysis variables such as reconstructed energy when integrated over the full sample of detected events. The ice model variants of most interest, which may strongly impact analysis outcomes, are those where correlated shifts in the ice model are made over certain larger regions of the detector. These could be accounted for with a layer-by-layer variation protocol, given a suitably derived correlated covariance matrix for $\vec{\eta}$. This would require many thousands of independent parameter constrains to be derived from flasher data, two for each layer and their covariances, and very strong local anti-correlations would be expected.

To avoid generating constraints within the aforementioned large layer-by-layer nuisance covariance matrix, a significantly more efficient method has been developed. Rather than describing the ice model layer-by-layer, it can equivalently be described in Fourier space, with entries in $\vec{\eta}$ being the amplitudes and phases of Fourier modes. A fully covariant constraint on $\vec{\eta}$ derived in Fourier space would yield an equivalent uncertainty envelope to one derived in depth space. However, the benefit of Fourier space is that it offers a natural ordering principle to select the most important uncertainty contributions. Low frequency Fourier modes have large amplitudes, and their variation reflects macroscopic changes over large regions of the IceCube detector. These are expected to contribute the dominant parts of analysis-space uncertainties. The high frequency modes, on the other hand, have smaller amplitudes and their variation is expected to have a minimal effect in analysis space, since the effect of the fast-varying ice properties effectively cancels when integrating over the detector and accounting for resolution effects. It will be demonstrated that this intuitive picture is valid in Sec. 3.2.

A notable feature of absorption and scattering coefficients shown in Fig. 3 is that they are highly correlated with each other as a function of depth. These two distributions can thus be represented in terms of a completely correlated part, which contains most of the information about their shapes (physically, the effective dust concentration vs. depth), and a completely anti-correlated part, which encodes the far sub-leading effects that determine their differences (physically, changes in dust composition or shape throughout the array).
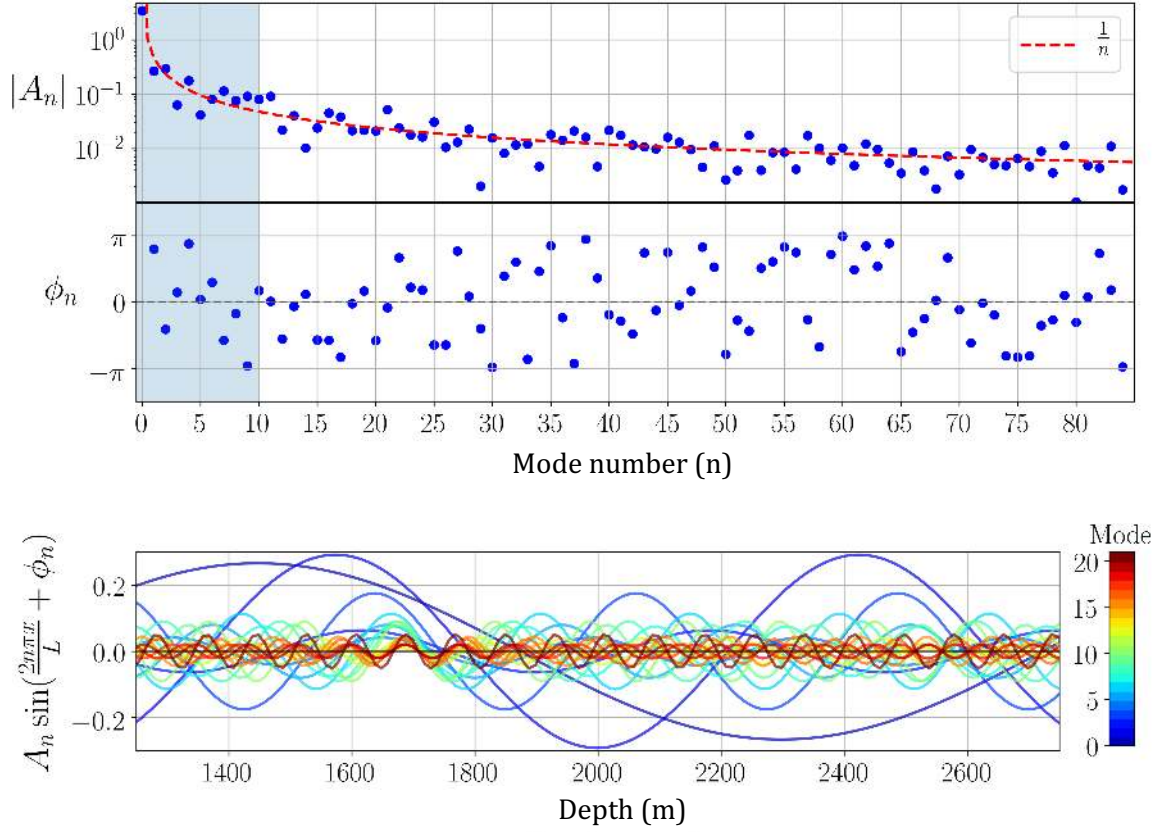
Figure 5. Top: The parameters from the Fourier decomposition of the central model. The blue region highlights the leading modes considered in this work. Note the rate of convergence of the amplitudes goes as $1/n$, consistent with a well-behaved series. Bottom: the depth dependant structure of each term in the Fourier expansion (with mode zero omitted).

To this end, we take the discrete Fourier decomposition of the sum (M^+) and difference (M^-) of the logarithms of the scattering and absorption coefficients at their central-model values. These can be considered to represent the totally correlated and totally anti-correlated parts of the absorption and scattering profiles, respectively. The depth profiles of M^+ and M^- are shown in Fig. 4.

$$M^+(x) \equiv \frac{1}{2}\log_{10}\left(Abs \times Sca\right) = \frac{A_0}{2} + \sum_n A_n \sin\left(\frac{2n\pi x}{L} + \phi_n\right), \tag{3.1}$$

$$M^-(x) \equiv \frac{1}{2}\log_{10}\left(Abs \div Sca\right) = \frac{B_0}{2} + \sum_n B_n \sin\left(\frac{2n\pi x}{L} + \gamma_n\right). \tag{3.2}$$

From these variables, we may recover the physical parameters via the transformation:

$$Abs = 10^{M^+ + M^-} \qquad Sca = 10^{M^+ - M^-}, \tag{3.3}$$

with all of Abs, Sca, M^+ and M^- being functions of depth, x. The zeroth Fourier mode of M^- encodes the absolute relative scale of absorption vs. scattering, which is substantial, whereas the higher modes encode their depth-dependent shape differences, and are very small. This can be seen in the flatness of M^- in Fig. 4, which reflects very little variation in

the ratio of absorption to scattering across the detector within the best-fit ice model. The variations in $M+$ that increase or reduce absorption or scattering in unison as a function of depth are focus of this work.

3.2 Determination of nuisance covariance from IceCube calibration data

As is common in experimental particle physics, IceCube uses calibration data to constrain systematic uncertainties that may impact physics analyses. Given a set of nuisance parameters $\vec{\eta}$, the constraints from calibration data can, under a suitable and testable assumption of Gaussianity, be expressed as a covariance matrix Ξ relative to the best fit point. The nuisance parameters of interest here the amplitudes and phases of the lowest-frequency modes of $M+$. The covariance matrix in this space encodes the allowed variations in the nuisance parameters, as well as their correlations.

While ice model calibrations can in principle be performed in many ways, a particularly powerful suite of calibration samples for IceCube are data taken using the LED flashers installed on each Digital Optical Module. The arrival time distribution and intensity of detected light detected by pairs of flashing and receiving Digital Optical Modules gives a geometry-dependent constraint on absorption and scattering. These data can be used to derive Ξ for the $M+$ modes of relevance to physics analysis samples.

The covariance matrix construction to describe constraints on the nuisance parameters (for example, properties of the ice) rests on the assumption of a multivariate Gaussian likelihood profile L in the vicinity of the best fit point. This is to be expected based on the central limit theorem, in most cases. The covariance matrix then encodes the shape of the likelihood surface around its minimum:

$$L_{\eta_0,\Xi}(\vec{\eta}) = \frac{1}{\sqrt{(2\pi)^d det(\Xi)}} exp\left\{ -\frac{1}{2}(\vec{\eta} - \vec{\eta}_0)^T \Xi^{-1}(\vec{\eta} - \vec{\eta}_0) \right\}. \tag{3.4}$$

We measure Ξ by probing L. The approach we take here is to pick specific directions in which to profile L around the best fit point, in particular in the directions along single nuisance parameter variations, and in fully correlated variations of pairs of nuisance parameters, and find points where the likelihood L has shifted by an amount corresponding to a 1σ variation in each direction. From these points we extract $\mathcal{M} \equiv \Xi^{-1}$, which can then be inverted to yield the nuisance covariance matrix Ξ. The details of the likelihood construction used to compare ice models to calibration data in IceCube is discussed elsewhere [4]. In the vicinity of the minimum it has been validated that this profile is Gaussian to a very good approximation, in all scanned single-parameter and covariant directions. The one-sigma points are defined as vectors $\vec{\sigma}$ that may point in any direction, as:

$$\vec{\sigma}^T \mathcal{M} \vec{\sigma} = \mathcal{M}_{ab}\sigma_a\sigma_b = 1. \tag{3.5}$$

We define two classes of one-sigma vectors $\vec{\sigma}$. First, for the unit vector directions, $\vec{\sigma}^i_{unit}$ for each i. Each vector $\vec{\sigma}^i$ with superscript i points in direction $\hat{i}$, that is:

$$\vec{\sigma}^i_{unit} = \sigma^i_{unit}\hat{i} \qquad [no\,sum\,on\,i], \tag{3.6}$$

or if we include the vector index b explicitly, $\sigma^i_{unit,b} = \sigma^i_{unit}\delta_{ib}$ We also define the one-sigma points in the correlated directions. These are labelled by two indices $\vec{\sigma}^{ij}_{corr}$, and the vector $\vec{\sigma}^{ij}$ with subscripts i and j points in a direction $(\hat{i} + \hat{j})/\sqrt{2}$, that is:

$$\vec{\sigma}^{ij}_{corr} = \sigma^{ij}_{corr}\left(\frac{\hat{i} + \hat{j}}{\sqrt{2}}\right) \qquad [no\,sum\,on\,ij]. \tag{3.7}$$
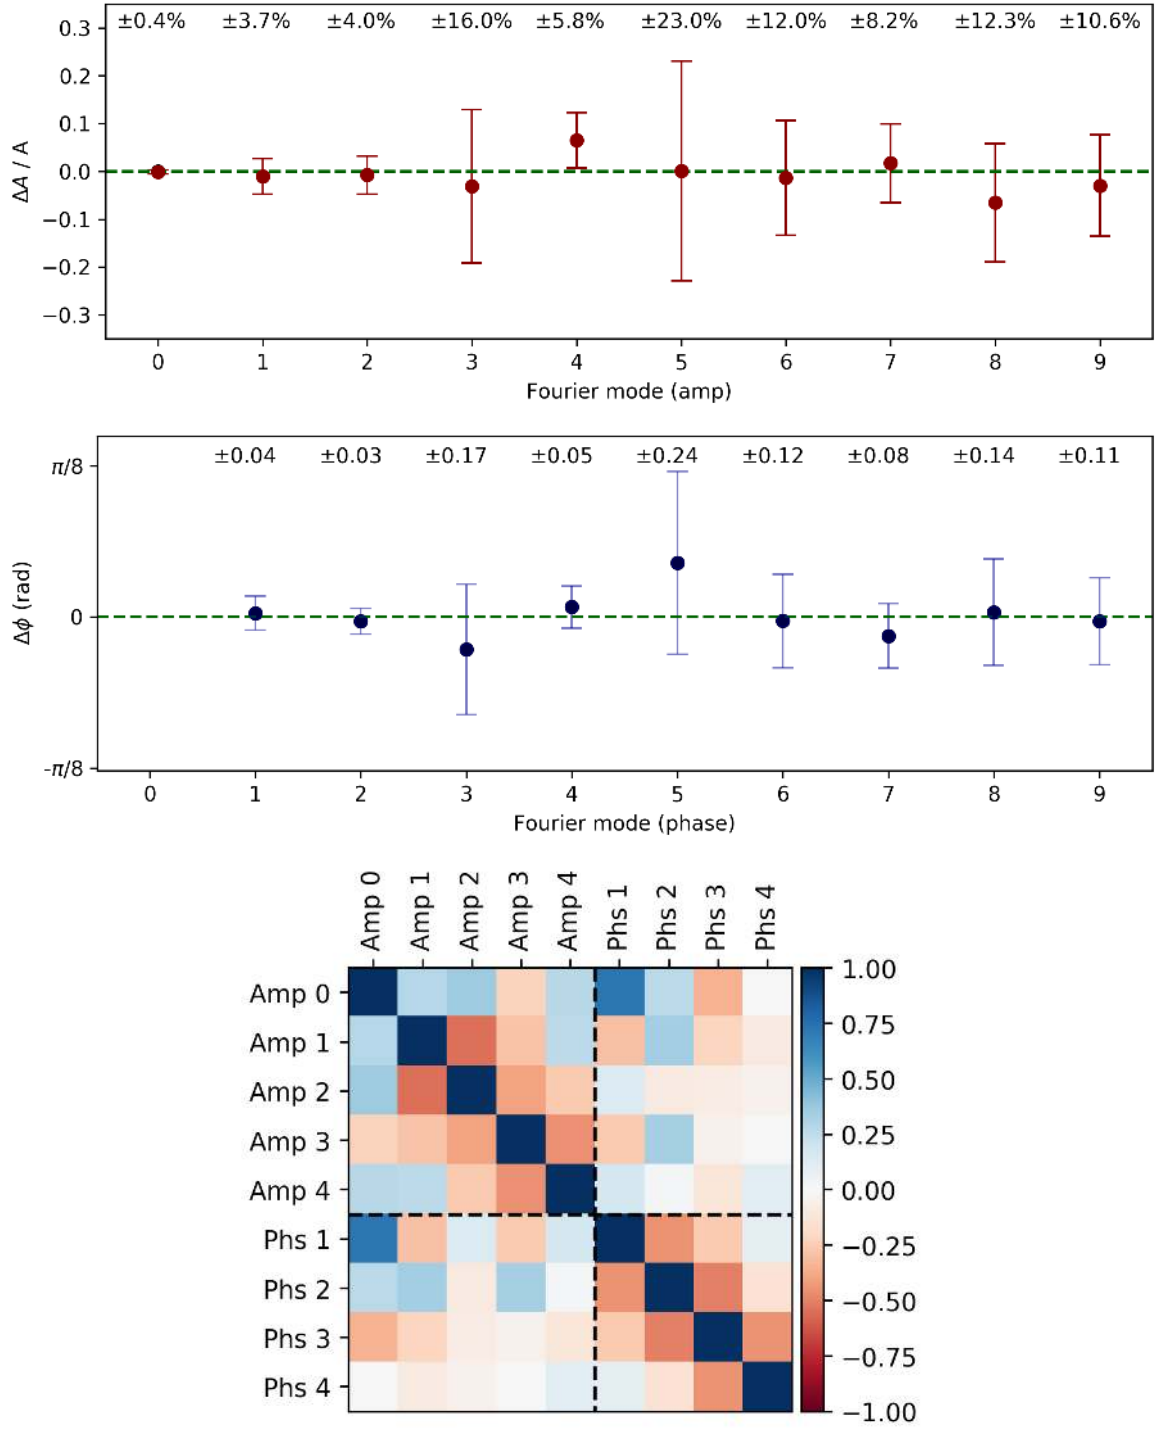
Figure 6. Top: Diagonal Fourier parameter constraints derived using flasher calibration data in amplitudes and phases. Bottom: Hessian matrix in nuisance space including correlated parameter widths for a mode 0-4 from flasher calibration data. The dashed lines separate amplitude and phase parameters.

Despite being labelled by two superscripts, each $\vec{\sigma}_{corr}^{ij}$ is still a vector with a single vector index, which we can again write as b, as $\sigma_{corr,b}^{ij} = \sigma_{corr}^{ij} \left(\delta_{ib} + \delta_{jb} \right) / \sqrt{2}$. The unit direction one-sigma points allow extraction of the diagonal elements of $\mathcal{M}$, via Eq. 3.6:

$$\mathcal{M}_{ab}\sigma_{unit,a}^{i}\sigma_{unit,b}^{i} = \mathcal{M}_{ab}\delta_{ia}\delta_{ib}\left(\sigma_{unit}^{i}\right)^2 = 1 \quad \rightarrow \quad \mathcal{M}_{ii} = \frac{1}{\sigma_{i,unit}^2} \quad [no\,sum\,on\,i]. \qquad (3.8)$$

The correlated directions similarly give the following constraints on $\mathcal{M}$:

$$1 = \sum_{ab} \mathcal{M}_{ab}\sigma_{corr,a}^{ij}\sigma_{corr,b}^{ij} \qquad [no\,sum\,on\,ij], \qquad (3.9)$$

$$= \frac{1}{2}\left(\sigma_{corr}^{ij}\right)^2\left(\mathcal{M}_{ii} + \mathcal{M}_{jj} + \mathcal{M}_{ij} + \mathcal{M}_{ji}\right). \qquad (3.10)$$

Since $\mathcal{M}$ can be defined without loss of generality as a symmetric matrix, and we have already obtained equations for the diagonal terms, we can invert this expression to find the off-diagonal elements $\mathcal{M}_{ij}$:

$$\mathcal{M}_{ij} = \frac{1}{\left(\sigma_{corr}^{ij}\right)^2} - \frac{1}{2}\left(\frac{1}{\sigma_{i,unit}^2} + \frac{1}{\sigma_{j,unit}^2}\right). \qquad (3.11)$$

Thus with the unit and correlated one-sigma points we can specify the entire matrix $\mathcal{M}_{ij}$ that defines the multivariate Gaussian, via $\Xi = \mathcal{M}^{-1}$. The diagonal 1σ widths for the leading Fourier amplitudes and phases from IceCube flasher data and the full Hessian matrix $\mathcal{M}$ normalized to these values are shown in Fig. 6.

The relevance of the correlation terms is well illustrated by considering the envelope of allowed perturbed models in depth space. Fig. 7 shows the envelope of allowed correlated variations within the measured covariace matrix, both including off-diagonal correlation terms (top), and treating each mode as independent (bottom). Red lines mark the edges of frequentist envelopes containing 68%, 95% and 99% of sampled models at each depth. The emergence of regions of high uncertainty outside of the detector volume, and in the strongly absorbing "dust layer" that is known to reside at around -100m in IceCube coordinates, are only manifest when nuisance parameter correlations are properly accounted for. This demonstrates a case where the often overlooked correlations between nuisance parameters constrained by common calibration data are crucial for properly encoding physical features of the underlying detector model.

The envelope described above was generated with a zero-to-four mode model (9 nuisance parameters), which limits the sharpness of features in the depth envelope. The choice of a fourth-mode cut-off is motivated in Sec. 3.3 in the context of the IceCube atmospheric neutrino sample, where modes above the fourth are shown to have a rapidly diminishing effect on the analysis distributions within flasher-constrained uncertainty. Because the size of the envelope in Fig. 7 depends on the high frequency cutoff, caution should be exercised against over-interpreting as an absolute statement of dust uncertainty vs depth. In principle, very high frequency modes are neither constrained by flasher calibration nor relevant for the atmospheric neutrino sample, and so any envelope so constructed relies on some notion of smoothness or regularization. Our goal here is not to construct a measurement of dust concentration vs. depth, but rather to map the physically important constraints on low frequency modes from calibration data to analysis space, and it is the effect of these modes only that is encoded in Fig. 7.
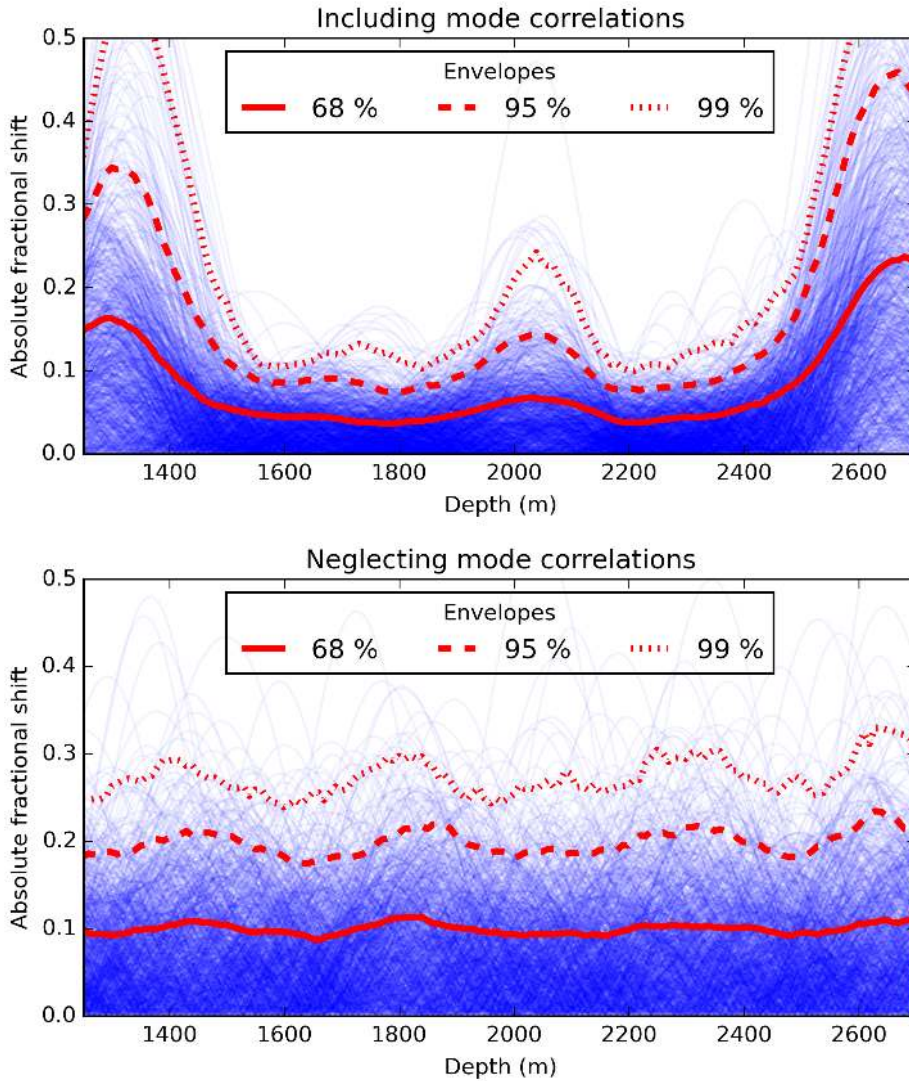
Figure 7. Fractional excursions arising from multiple sampling of 5,000 models. The largest excursion are allowed in regions where the properties of the ice are poorly constrained, such as outside the detector and in the dust layer. Note that this structure is embodied in the correlations between nuisance parameters, rather than their individual spreads - this can be seen by comparing the correlated sampling approach (top) using full covariance matrix with an uncorrelated sampling (bottom).

3.3   Extraction of Fourier nuisance gradients

A SnowStorm ensemble was generated with variations in the first 12 Fourier modes, and processed through photon propagation, detector simulation, Level 1 and Level 2 processing and event selection for atmospheric neutrinos. This processing and selection chain derives from the one described in [15] with several improvements, and is used for the forthcoming IceCube high energy sterile neutrino analysis. That analysis, which represents the multi-year extension of previous work [3], selects events between 500 GeV and 10 TeV and compares the energy and zenith distribution of up-going muons against predictions made under different exotic oscillation scenarios. Hypothetical eV-scale sterile neutrinos, which have been invoked
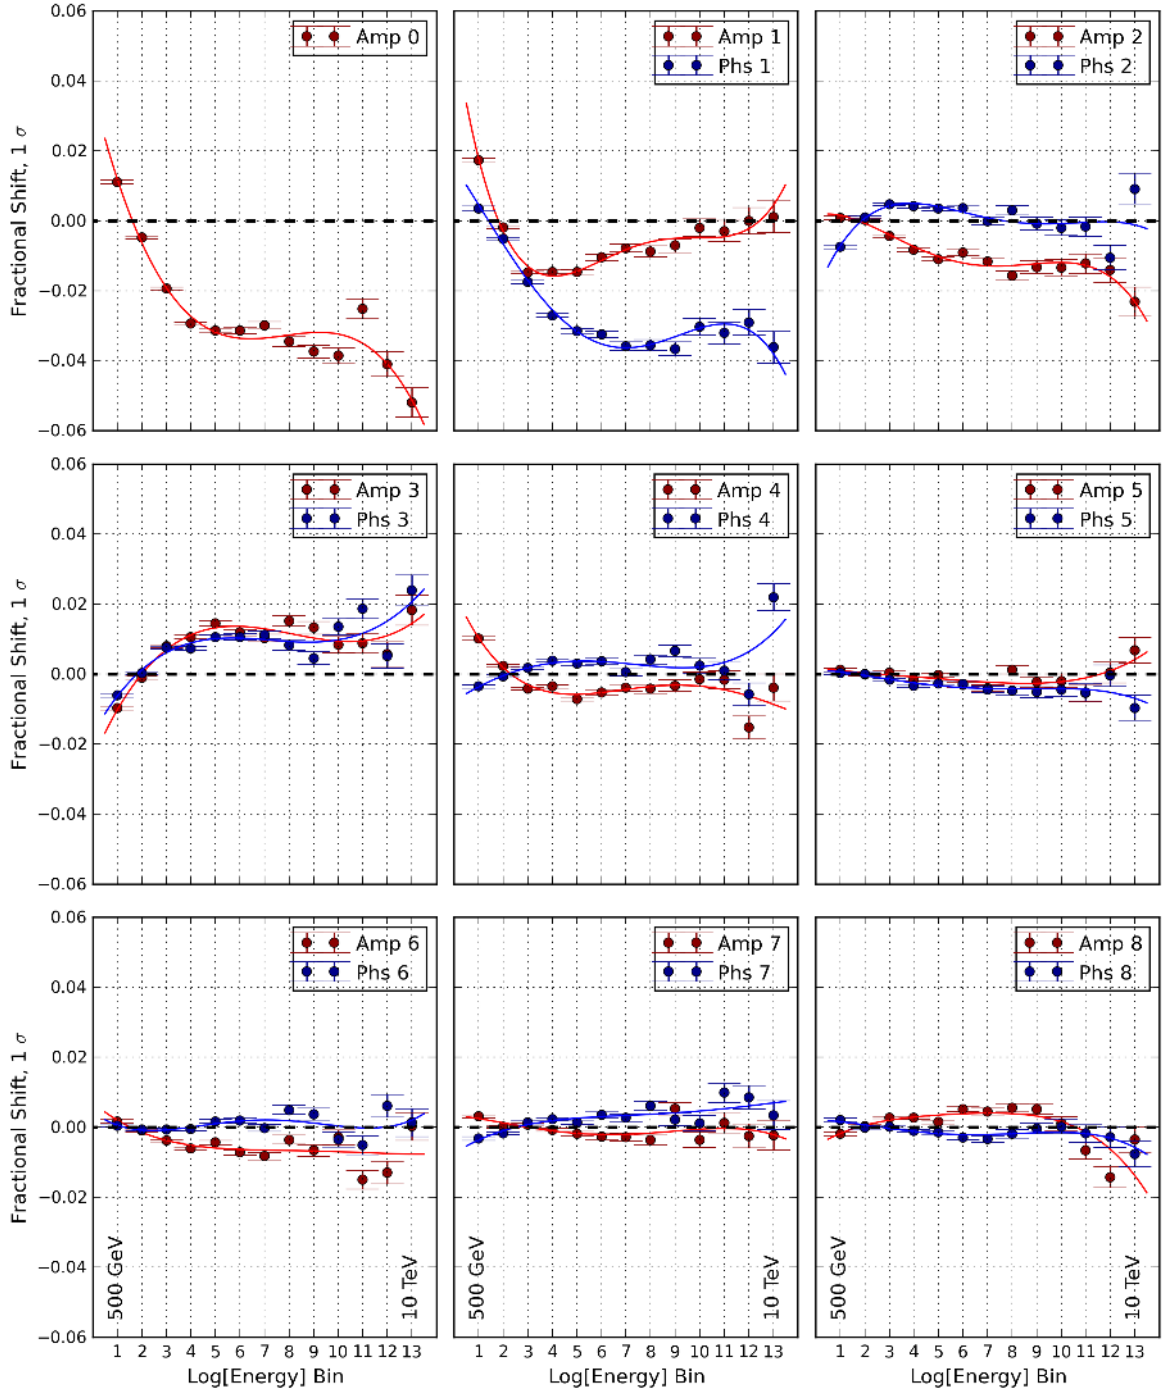
Figure 8. Fractional energy gradients split along amplitudes 0 to 8 (red), and phases 1 to 8 (blue). Note the decreasing impact of increasing mode number, verifying the assumption of negligible impact of higher modes.

to explain previous short-baseline neutrino oscillation anomalies [16, 17], would induce a characteristic disappearance of muon neutrinos due to matter-resonant oscillations. Such a disappearance effect would be observable as a distortion in the reconstructed spectrum,

given sufficiently large sterile neutrino mixing. Accuracy of this analysis relies on a proper description of the effects of systematic uncertainties at the few-percent level over the full spectrum.

Using the cutting based procedure described in Sec. 2, the vector of gradients $\vec{G}_{\vec{\rho}} \equiv \vec{\nabla}_\eta \left[\psi_{\vec{\rho},\vec{\eta}} \right]_{\vec{\eta}=\vec{0}}$ were extracted. Since existing constraints limit the sterile neutrino parameter space to that where only relatively small perturbations to the atmospheric neutrino spectrum are viable, a single $\vec{G}$ evaluated at $\rho = \vec{0}$ is appropriate.

The first nine gradients for Fourier amplitude and phase parameters as a function of reconstructed energy are shown in Fig. 8. The scale of effect of perturbing each mode within its 1σ diagonal uncertainty quickly diminishes with increasing mode number. Modes beyond approximately mode 4 become irrelevant in terms of their contributions to analysis uncertainty. This motivates restriction of the nuisance covariance and derived analysis covariance to four-mode model. Physically, this represents the principle that the effects of rapid changes in absorption and scattering strength averaging to zero are generally integrated out in the final energy distribution.

The gradients were also examined in zenith space, and found to be negligible there. The dominant effect of the dust distribution uncertainty is thus demonstrated to be primarily exhibited as distortions to the energy spectrum, and not the zenith distribution of events. The final covariance matrix is thus constructed in energy space only.

Although the trends are clear, the gradients retain some scatter from statistical precision of the SnowStorm ensemble set. In order to produce smooth gradients and avoid artifacts due to Monte Carlo discretization effects, we fit smoothed curves to the extracted fractional gradients. Empirically, a fourth order polynomial provides an adequate fit in all cases. The best fit curves are shown in Fig. 8 as solid lines. These smoothed gradients are propagated into the analysis covariance matrix.

3.4 Determination of analysis covariance

The gradients $\vec{G}$ and the nuisance covariance matrix from calibration data Ξ can finally be combined into an analysis space covariance matrix Σ, via Eq. 2.16. The smallness of the effect in zenith space allows to concentrate the statistical precision of the sample on achieving the most accurate possible estimate of the energy shape covariance. The final analysis covariance matrix is shown in Fig. 9. This matrix represents the depth-dependent dust uncertainty within the atmospheric neutrino sample between 500 GeV and 10 TeV. It incorporates the effects of the physically important modes of a depth-dependent continuous function representing effective dust distribution in the array, properly including effects of correlations between parameters, to yield a fully covariant estimate of uncertainty on the final sample. Most notably, despite depending on a large number of free parameters, this was achieved using a single Monte Carlo ensemble, via the SnowStorm method.

3.5 Tests of validity

Several assumptions are made in the construction of the SnowStorm-derived covariance matrix. Several of these assumptions can be explicitly tested. Here we present two examples.

First, the principle that the gradients $\vec{G}_{\vec{\rho}} \equiv \vec{\nabla}_\eta \left[\psi_{\vec{\rho},\vec{\eta}} \right]_{\vec{\eta}=\vec{0}}$ encapsulate the effects of small variations in the parameters η relies on approximate linearity, or equivalently, that the effects of the higher order terms in Eq. 2.5 are negligible. Were this not the case, a distribution of the variable of interest made using the full SnowStorm ensemble would not be equivalent to a similar distribution constructed using an unperturbed Monte Carlo set (see Eq. 2.7).
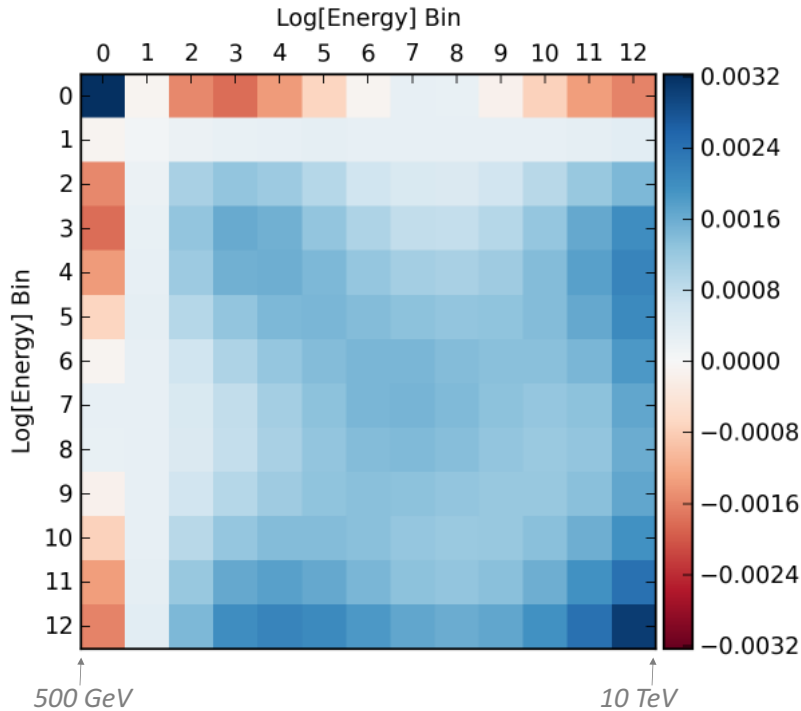
Figure 9. Analysis covariance matrix for application in atmospheric neutrino binned analysis space.

An explicit test of linearity is thus provided by comparing these two distributions, with any non-equivalence indicating the scale of neglected higher-order effects. A direct comparison of these two distributions for this analysis is shown in Fig. 10, left, demonstrating agreement at the sub-percent level, at all energies.

Second, the principle that the constraints on nuisance parameters can be described by a covariance matrix Ξ requires that the likelihood profile from calibration data is adequately described by a multivariate Gaussian near the best fit. For this study we found the likelihood from calibration data to be well described by Gaussian profiles in all single- and two-parameter directions, within statistical precision. Fig. 10, right shows the delta-log-likelihood distribution between several model perturbations and the best-fit point for variations of the amplitude of mode 4, as an example case. This is overlaid on a Gaussian profile, demonstrating a high degree of consistency. All other inspected profiles showed a similar level of consistency with the multivariate Gaussian assumption.

4 Conclusions

We have presented the SnowStorm method, an approach for efficiently propagating multivariate systematic uncertainties from calibration constraints to analysis space. Using a perturbative treatment of nuisance parameters, construction of a single Monte Carlo ensemble with randomly distributed nuisance variables allows for construction of predictions given generic nuisance vectors. The nuisance gradients extracted can then be combined with constraints in the nuisance space derived from calibration data. The power of this approach lies in the fact that only a single Monte Carlo ensemble is required for treatment of uncertainties describable by an in principle arbitrarily large number of nuisance parameters.
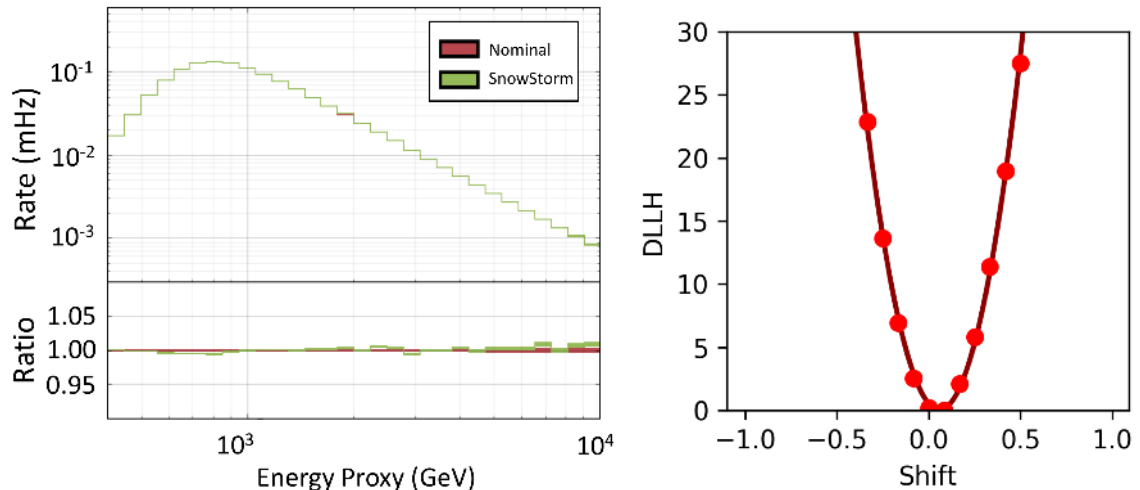
Figure 10. Two tests of validity of the SnowStorm method. Left: comparison of integrated Snow-Storm ensemble to default Monte Carlo. Lack of substantial difference illustrates that higher order terms can be neglected. Right: Example delta-log-likelihood (DLLH) profile from calibration data in one of the amplitude directions (mode 4). All such likelihood profiles are found to be Gaussian within statistical precision.

As an example case, an application of the SnowStorm method to IceCube atmospheric neutrino analyses was described. There, a class of bulk ice optical uncertainties related to depth-dependent effective dust concentration within the IceCube array was propagated from LED calibration data to the analysis space of the energy distributions of atmospheric neutrinos. This application is especially illustrative of the power of this method, since it involves a large number of nuisance parameters that have physically important correlations in their calibration constraints, mapped to analysis covariance via a single Monte Carlo ensemble. This method is being used by forthcoming high statistics IceCube analyses, and may find wider utility for treating complicated systematic uncertainties, both within IceCube and beyond.

Acknowledgements

Appendix: General error propagation and extraction of analysis covariance

We begin with the most general form of error propagation in a function vector $\vec{f}(\vec{x})$ and proceed via Taylor expansion so that for a small shift in the input parameters $\vec{x} \rightarrow \vec{x} + \delta\vec{x}$ we have:

$$\delta f_i = f_i(x) + \frac{\partial f_i}{\partial x_j}\delta x_j = J_{ij}\delta x_j, \tag{4.1}$$

where J is the Jacobian matrix:

$$J_{ij} = \frac{df_i}{dx_j}. \tag{4.2}$$

For the propagation of uncertainties, we seek the covariance Σ_f of f with respect to its components. We define this as:

$$\Sigma_f = \left\langle \left[\vec{f} - E[\vec{f}]\right] \left[\vec{f} - E[\vec{f}]\right] \right\rangle. \tag{4.3}$$

Expanding, we find:

$$\Sigma_f = \left\langle f_i f_j - \bar{f}_i f_j - f_i \bar{f}_j + \bar{f}_i \bar{f}_j \right\rangle = \left\langle f_i f_j \right\rangle - \left\langle \bar{f}_i f_j \right\rangle - \left\langle f_i \bar{f}_j \right\rangle + \left\langle \bar{f}_i \bar{f}_j \right\rangle, \tag{4.4}$$

with all three of the last terms reducing to the same expression:

$$\left\langle \bar{f}_i f_j \right\rangle = \left\langle f_i \bar{f}_j \right\rangle = \left\langle \bar{f}_i \bar{f}_j \right\rangle = \bar{f}_i \bar{f}_j, \tag{4.5}$$

and so:

$$\Sigma_f = \left\langle f_i f_j \right\rangle - \bar{f}_i \bar{f}_j. \tag{4.6}$$

If the variability in f around its mean is due to small perturbations in x we can write $f = \bar{f} + J.\delta x$ for randomly fluctuating δx and so:

$$\Sigma_f = \left\langle \left(\bar{f}_i + J_{ik}\delta x_k\right)\left(\bar{f}_j + J_{jl}\delta x_l\right) \right\rangle - \bar{f}_i \bar{f}_j, \tag{4.7}$$

$$\Sigma_f = \bar{f}_i \bar{f}_j + \bar{f}_i \mathcal{J}_{ik} \left\langle \delta x_k \right\rangle + \bar{f}_j J_{jl} \left\langle \delta x_l \right\rangle - \bar{f}_i \bar{f}_j + J_{ik}J_{jl} \left\langle \delta x_k \delta x_l \right\rangle. \tag{4.8}$$

Since the δx_i all fluctuate around 0, we know:

$$\left\langle \delta x_i \right\rangle = 0 \quad \forall \quad i. \tag{4.9}$$

We can also identify the last ensemble average as the covariance in x:

$$\langle \delta x_k \delta x_l \rangle \equiv [\Sigma_x]_{kl}, \qquad (4.10)$$

leading us to the central identity, as required:

$$\Sigma_f = J \Sigma_x J^T. \qquad (4.11)$$

In our case of interest we associate $f \to \psi_{\vec{\rho}, \vec{\eta}}$ and $x \to \vec{\eta}$. Then Σ_f is the covariance matrix in analysis space, $\Sigma_f \to \Sigma$, and Σ_x is the covariance matrix in nuisance space, $\Sigma_x \to \Xi$. Then $J_{\alpha,i} = \partial \psi_{\vec{\rho}, \vec{\eta}}(E_\alpha) / \partial \eta_i = G_{\vec{\rho}, i; \alpha}$, proving the theorem

$$\Sigma_{\alpha, \beta} = G_{i;\alpha} \Xi_{ij} G_{j;\beta}. \qquad (4.12)$$

References

[1] F. Halzen and S. R. Klein. IceCube: An Instrument for Neutrino Astronomy. *Rev. Sci. Instrum.*, 81:081101, 2010.

[2] M. G. Aartsen et al. The IceCube Neutrino Observatory: Instrumentation and Online Systems. *JINST*, 12(03):P03012, 2017.

[3] M. G. Aartsen et al. Searches for Sterile Neutrinos with the IceCube Detector. *Phys. Rev. Lett.*, 117(7):071801, 2016.

[4] M. G. Aartsen et al. Measurement of south pole ice transparency with the icecube led calibration system. *Nucl.Instrum.Meth.*, A(711):73–89, 2013.

[5] P. Askebjer, S. W. Barwick, L. Bergstrom, A. Bouchta, S. Carius, et al. Optical properties of deep ice at the South Pole: Absorption. *Appl.Opt.*, 36:4168–4180, 1997.

[6] P. Askebjer et al. Optical properties of the South Pole ice at depths between 0.8-km and 1-km. *Science*, 267:1147–1150, 1995.

[7] P. Buford Price and L. Bergström. Optical properties of deep ice at the south pole: scattering. *Appl. Opt.*, 36(18):4181–4194, Jun 1997.

[8] N. E. Bramall, R. C. Bay, K. Woschnagg, R. A. Rohde, and P. B. Price. A deep high-resolution optical log of dust, ash, and stratigraphy in south pole glacial ice. *Geophysical Research Letters*, 32(21), 2005.

[9] D. Chirkin. Evidence of optical anisotropy of the South Pole ice. *Proceedings, 33rd International Cosmic Ray Conference (ICRC2013): Rio de Janeiro, Brazil, July 2-9, 2013*, page 0580, 2013.

[10] M. Ackermann et al. Optical properties of deep glacial ice at the south pole. *Journal of Geophysical Research: Atmospheres*, 111(D13), 2006.

[11] B. J. P. Jones. *Sterile Neutrinos in Cold Climates*. PhD thesis, MIT, 2015.

[12] B. J. P. Jones for the IceCube Collaboration. IceCube Sterile Neutrino Searches. *EPJ Web Conf.*, 207:04005, 2019.

[13] M. Ackermann et al. South Pole glacial climate reconstruction from multi-borehole laser particulate stratigraphy. *J. Glaciol.*, 59(218):1117–1128, 2013.

[14] P. B. Price, K. Woschnagg, and D. Chirkin. Age vs depth of glacial ice at south pole. *Geophysical Research Letters*, 27(14):2129–2132, 2000.

[15] M. G. Aartsen et al. Evidence for Astrophysical Muon Neutrinos from the Northern Sky with IceCube. *Phys. Rev. Lett.*, 115(8):081102, 2015.

[16] C. Athanassopoulos et al. Evidence for anti-muon-neutrino to anti-electron-neutrino oscillations from the LSND experiment at LAMPF. *Phys. Rev. Lett.*, 77:3082–3085, 1996.

[17] A. A. Aguilar-Arevalo et al. Improved Search for $\bar{\nu}_\mu \to \bar{\nu}_e$ Oscillations in the MiniBooNE Experiment. *Phys. Rev. Lett.*, 110:161801, 2013.