Sampling in high dimensional spaces

(with the Global and Modular BSM Inference Tool)

Martin White

Leiden 2018

Outline

- Brief introduction to global fits of beyond-SM physics theories and the Global and Modular BSM Inference Tool (GAMBIT)
- Part 1: GAMBIT comparison of methods for sampling high dimensional parameter spaces
- Part 2: Where GAMBIT might benefit from machine learning techniques (plus existing GAMBIT machine learning plans)

Where might we “see” particle dark matter?

- colliders (LHC + previous)

Where might we “see” particle dark matter?

- colliders (LHC + previous)
- measurements of the magnetic moment of the muon

Where might we "see" particle dark matter?

- colliders (LHC + previous)
- measurements of the magnetic moment of the muon
- electroweak precision tests

Where might we “see” particle dark matter?

- colliders (LHC + previous)
- measurements of the magnetic moment of the muon
- electroweak precision tests
- dark matter direct detection experiments

Where might we “see” particle dark matter?

- colliders (LHC + previous)
- measurements of the magnetic moment of the muon
- electroweak precision tests
- dark matter direct detection experiments
- searches for antimatter in cosmic rays, nuclear cosmic ray ratios
- radio astronomy data

Where might we “see” particle dark matter?

- colliders (LHC + previous)
- measurements of the magnetic moment of the muon
- electroweak precision tests
- dark matter direct detection experiments
- searches for antimatter in cosmic rays,nuclear cosmic ray ratios
- radio astronomy data
- effects of dark matter on reionisation, recombination and helioseismology

Where might we “see” particle dark matter?

- colliders (LHC + previous)
- measurements of the magnetic moment of the muon
- electroweak precision tests
- dark matter direct detection experiments
- searches for antimatter in cosmic rays, nuclear cosmic ray ratios
- radio astronomy data
- effects of dark matter on reionisation, recombination and helioseismology
- relic density (CMB + other data)

Where might we “see” particle dark matter?

- colliders (LHC + previous)
- measurements of the magnetic moment of the muon
- electroweak precision tests
- dark matter direct detection experiments
- searches for antimatter in cosmic rays, nuclear cosmic ray ratios
- radio astronomy data
- effects of dark matter on reionisation, recombination and helioseismology
- relic density (CMB + other data)
- neutrino masses and mixings

Where might we “see” particle dark matter?

- colliders (LHC + previous)
- measurements of the magnetic moment of the muon
- electroweak precision tests
- dark matter direct detection experiments
- searches for antimatter in cosmic rays, nuclear cosmic ray ratios
- radio astronomy data
- effects of dark matter on reionisation, recombination and helioseismology
- relic density (CMB + other data)
- neutrino masses and mixings
- Indirect DM searches (e.g. FERMI-LAT, HESS, CTA, IceCube, etc)

How do we tell which theories are viable?

- Combine results from all relevant experimental searches
- This is straightforward for models with few parameters:
- Simplest method:

  - overlay exclusion curves from different experiments
  - look for “excluded” and “non-excluded regions”

Beniwal, Rajec, Savage, Scott, Weniger, MJW, Williams, Phys.Rev. D93 (2016) no.11, 115016

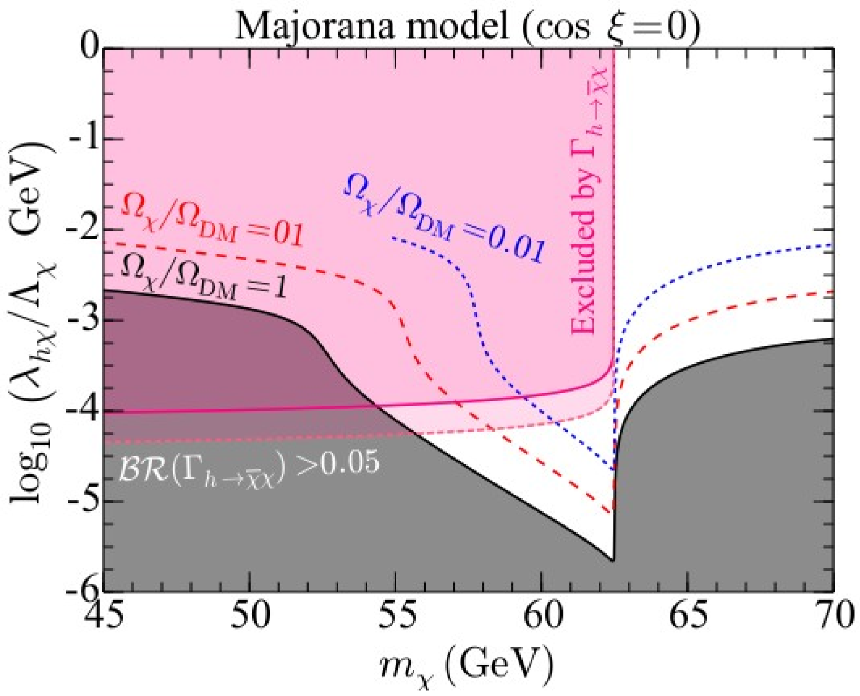

$$\mathcal{L}_\chi = \mathcal{L}_{SM} + \frac{1}{2}\overline{\chi}i\partial\!\!\!/\chi - \frac{1}{2}m_\chi\overline{\chi}\chi - \frac{1}{2}\frac{\lambda_{h\chi}}{\Lambda_\chi}\left[\cos\xi\,\overline{\chi}\chi + \sin\xi\,\overline{\chi}i\gamma_5\chi\right]\left(v_0 h + \frac{1}{2}h^2\right)$$

What if there are many constraints?

- Need to combine them properly into a *joint likelihood*

$$\mathcal{L} = \mathcal{L}_{\text{collider}}\mathcal{L}_{\text{DM}}\mathcal{L}_{\text{flavor}}\mathcal{L}_{\text{EWPO}}\cdots$$

- For two parameter models, can then continue as before

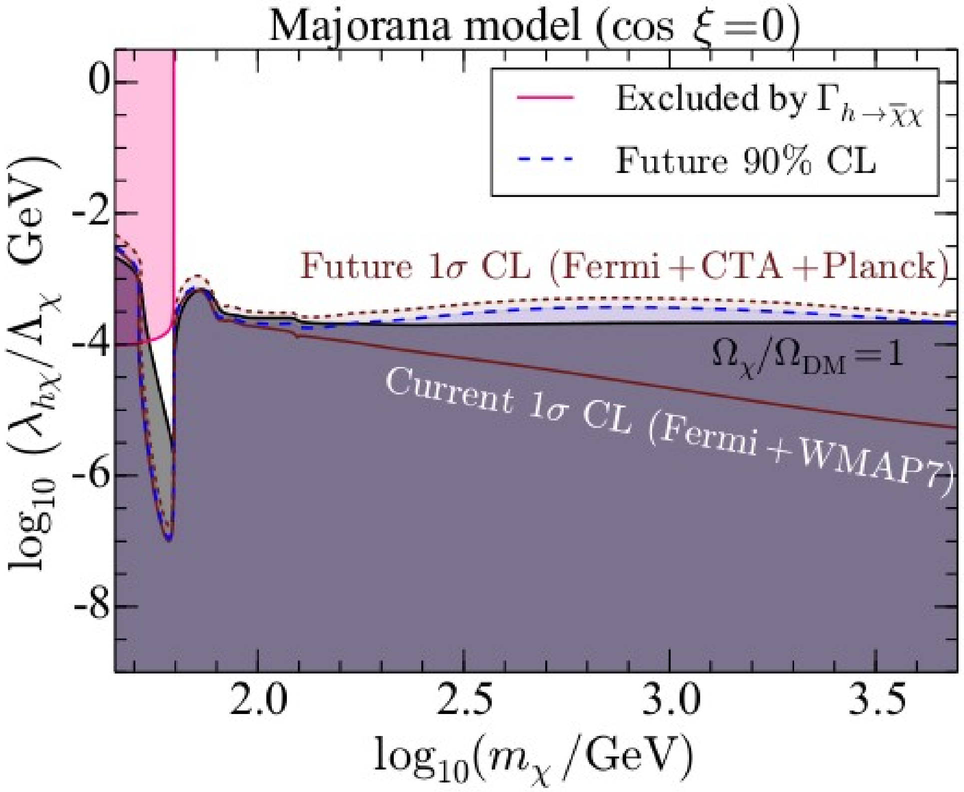

Beniwal, Rajec, Savage, Scott, Weniger, MJW, Williams, Phys.Rev. D93 (2016) no.11, 115016

$$\mathcal{L}_\chi = \mathcal{L}_{\text{SM}} + \frac{1}{2}\overline{\chi}i\partial\!\!\!/\chi - \frac{1}{2}m_\chi\overline{\chi}\chi - \frac{1}{2}\frac{\lambda_{h\chi}}{\Lambda_\chi}\left[\cos\xi\,\overline{\chi}\chi + \sin\xi\,\overline{\chi}i\gamma_5\chi\right]\left(v_0 h + \frac{1}{2}h^2\right)$$

What if there are many parameters?

- Much harder in principle
- Need to:

  - scan the space intelligently (grid scan is *precisely* the worst method)
  - interpret the results (Bayesian/frequentist)
  - project down to parameters of interest (marginalise/profile)

*i.e. need a **global statistical fit***

What if there are many models?

- Of course, there *are* many models
- Can rinse and repeat the above procedure
- Notice that we have two distinct problems:

<u>Parameter estimation</u>

Given a particular model, which set of parameters best fits the available data

(Rigorous exclusion limits and parameter measurements)

<u>Model comparison</u>

Given a set of models, which is the best description of the data, and how much better is it?

(Model *X* is now worse than model *Y*)

GAMBIT: The Global And Modular BSM Inference Tool

gambit.hepforge.org

- Fast definition of new datasets and theoretical models
- Plug and play scanning, physics and likelihood packages
- Extensive model database – not just SUSY
- Extensive observable/data libraries
- Many statistical and scanning options (Bayesian & frequentist)
- *Fast* LHC likelihood calculator
- Massively parallel
- Fully open-source

ATLAS	F. Bernlochner, A. Buckley, P. Jackson, M. White
LHCb	M. Chrząszcz, N. Serra
Belle-II	F. Bernlochner, P. Jackson
Fermi-LAT	J. Conrad, J. Edsjö, G. Martinez, P. Scott
CTA	C. Balázs, T. Bringmann, M. White
CMS	C. Rogan
IceCube	J. Edsjö, P. Scott
XENON/DARWIN	B. Farmer, R. Trotta
Theory	P. Athron, C. Balázs, S. Bloor, T. Bringmann, J. Cornell, J. Edsjö, B. Farmer, A. Fowlie, T. Gonzalo, J. Harz, S. Hoof, F. Kahlhoefer, S. Krishnamurthy, A. Kvellestad, F.N. Mahmoudi, J. McKay, A. Raklev, R. Ruiz, P. Scott, R. Trotta, A. Vincent, C. Weniger, M. White, S. Wild

31 Members in 9 Experiments, 12 major theory codes, 11 countries

GAMBIT modules

- **ColliderBit:** collider observables including Higgs + SUSY Searches from ATLAS, CMS, LEP
- **DarkBit:** dark matter observables (relic density, direct & indirect detection)
- **FlavBit:** including $g-2$, $b \to s\gamma$, B decays (new channels), angular obs., theory unc., LHCb likelihoods
- **SpecBit:** generic BSM spectrum object, providing RGE running, masses, mixings
- **DecayBit:** decay widths for all relevant SM and BSM particles
- **PrecisionBit:** precision EW tests (mostly via interface to FeynHiggs or SUSY-POPE)
- **ScannerBit:** manages stats, sampling and optimisation

What's in a module?

- Module functions (actual bits of GAMBIT C++ code)
- These can depend on other module functions
- Or can they can depend on *backends*(external codes)
- Adding new things is ***easy*** (detailed manual)
- Hooking up new backends or swapping them is ***easy***
- Module functions are **tagged** according to what they can calculate → plug and play!

How does GAMBIT work?

- You specify what to calculate and how (yaml input file)
- GAMBIT checks to see which functions can do it
- A dependency resolver stitches things together in the right order, and calculations are also ordered by speed
- GAMBIT performs the scan and writes output
- Pippi makes the plots

LHC limits: the problem

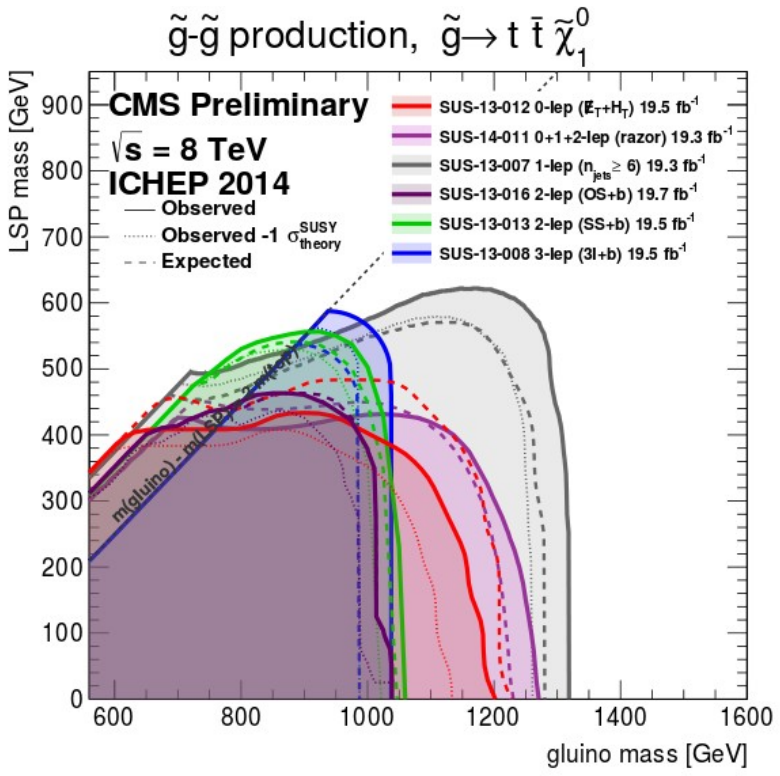

Model independent LHC limits

- Custom parallelised Pythia MC + custom detector sim
- Can generate 20,000 events on 12 cores in < 5 s
- Then apply Poisson likelihood with nuisance parameters for systematics
- Combine analyses using best expected exclusion
- The best you can do without extra public info from the experiments. CMS are getting better at this:

https://cds.cern.ch/record/2242860/files/NOTE2017_001.pdf

Astro limits: the problem

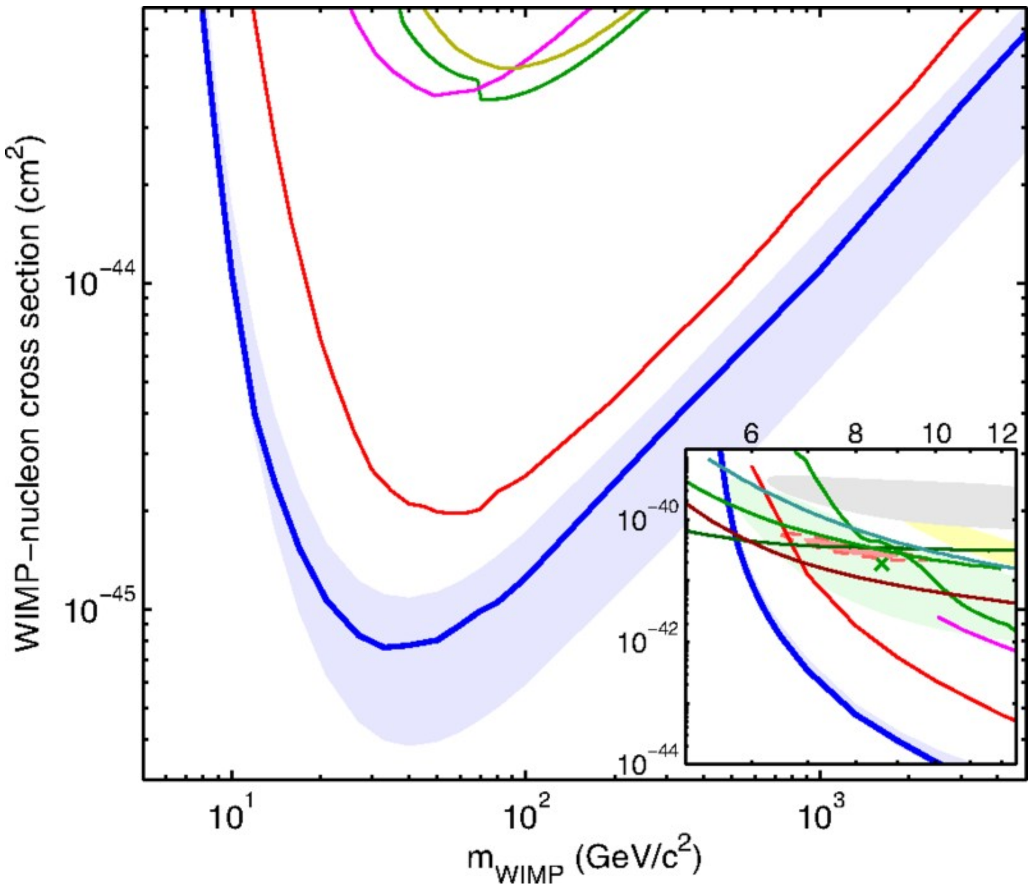

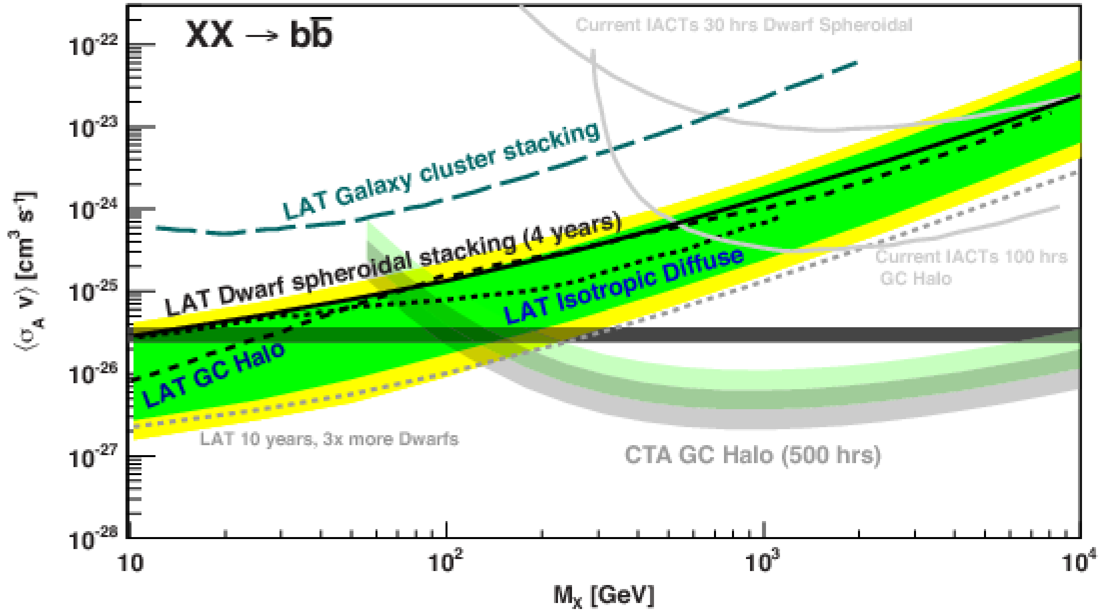

DarkBit: indirect detection

Gamma rays:

- Theoretical spectra calculated using branching fractions and tabulated gamma-ray yields
- Non-SM final state particles and Higgs are decayed on the fly with cascade Monte Carlo
- gamLike (gamlike.hepforge.org): New standalone code with likelihoods for DM searches from Fermi-LAT (dwarf spheroidals, galactic centre) and H.E.S.S. (galactic halo)

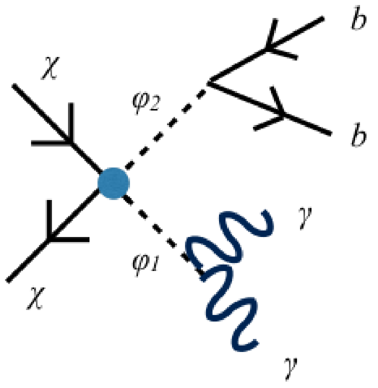

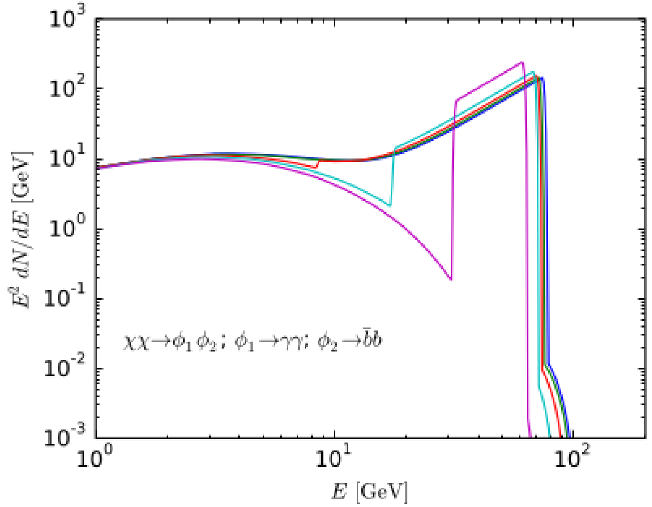

Solar neutrinos:

- Yields from DM annihilation in sun calculated by DarkSUSY. IceCube likelihoods contained in nulike (nulike.hepforge.org) standalone code.

DarkBit: direct detection

- In parallel with GAMBIT, we introduce *DDCalc* (ddcalc.hepforge.org), a tool to calculate event rates and complete likelihood functions for direct detection experiments taking into account:
  - A mix of both spin-independent and dependent contributions to the scattering rate.
  - Halo parameters (local density, DM velocity dispersion, etc.) chosen by the user.
- We currently have implemented likelihoods for Xenon(1T, 100), LUX, PandaX, SuperCDMS, PICO(60, 2L), and SIMPLE

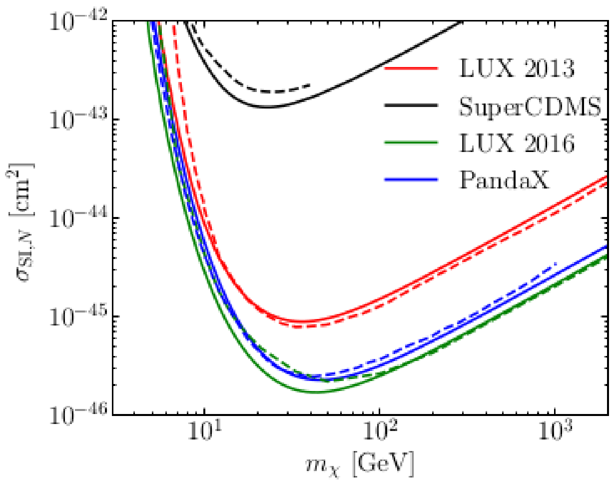

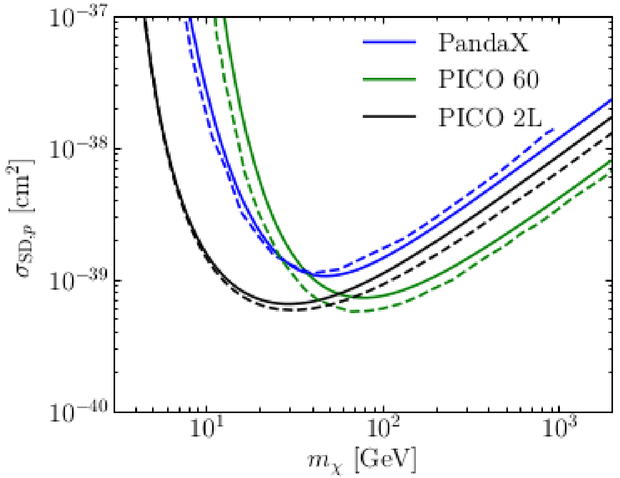

Official 90% CL limit
DDCalc

GAMBIT results

- GAMBIT has been released as an ***open source public tool***
- 9 papers published in EPJC (design, manual + physics studies of scalar singlet model, CMSSM/NUHM1/NUHM2, MSSM7)
- Feature article in *Physics World* March 2017 issue if you want a gentler introduction
- See **gambit.hepforge.org** for more info

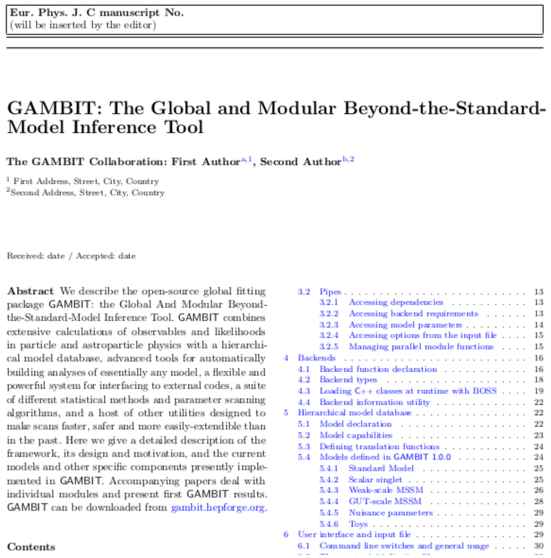

Eur. Phys. J. C manuscript No.
(will be inserted by the editor)

GAMBIT: The Global and Modular Beyond-the-Standard-Model Inference Tool

The GAMBIT Collaboration: First Author[a,1], Second Author[b,2]

[1] First Address, Street, City, Country
[2] Second Address, Street, City, Country

Received: date / Accepted: date

Abstract We describe the open-source global fitting package GAMBIT: the Global And Modular Beyond-the-Standard-Model Inference Tool. GAMBIT combines extensive calculations of observables and likelihoods in particle and astroparticle physics with a hierarchical model database, advanced tools for automatically building analyses of essentially any model, a flexible and powerful system for interfacing to external codes, a suite of different statistical methods and parameter scanning algorithms, and a host of other utilities designed to make scans faster, safer and more easily-extendible than in the past. Here we give a detailed description of the framework, its design and motivation, and the current models and other specific components presently implemented in GAMBIT. Accompanying papers deal with individual modules and present first GAMBIT results. GAMBIT can be downloaded from gambit.hepforge.org.

Contents

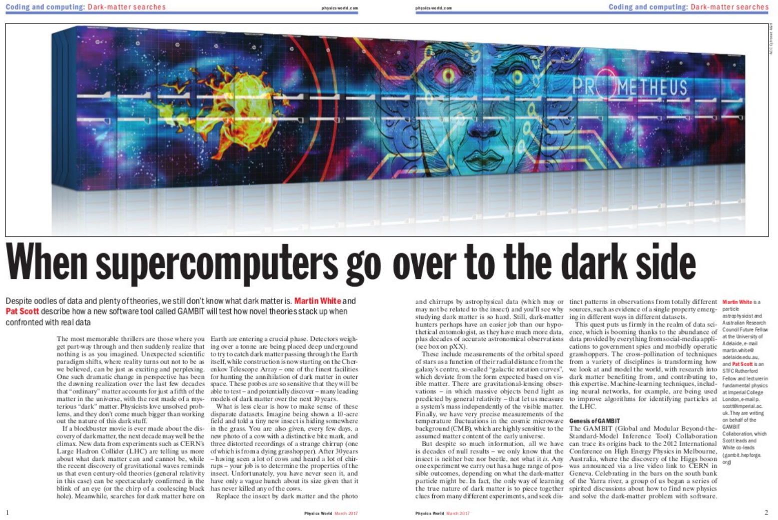

When supercomputers go over to the dark side

Despite oodles of data and plenty of theories, we still don't know what dark matter is. Martin White and Pat Scott describe how a new software tool called GAMBIT will test how novel theories stack up when confronted with real data

What's next for GAMBIT?

- More models (2HDM, axion models, more Higgs portal models, RH neutrinos, MSSM9 + 4D EW MSSM in near future, non-minimal SUSY later)
- Better interface to model building utilities such as Feynrules and SARAH
- Implementation of more complex LHC likelihoods (e.g. CMS simplified likelihood analyses for monojet and 0 lepton searches) plus Run II LHC searches
- CosmoBit

GAMBIT speed

- You may be thinking: this all sounds very nice, but how quickly can this possibly run?
- GAMBIT is made “quick enough” by:

1) Using massive parallelisation (OpenMP + MPI)

2) Using very smart sampling methods

3) Using approximations in simulations where possible

4) Ordering likelihood calculations by speed, and not doing expensive calculations if simpler likelihoods already disfavour a parameter point

Nonetheless: *many* CPU hours are required for complex models

GAMBIT sampling

Eur. Phys. J. C manuscript No.
(will be inserted by the editor)

Comparison of statistical sampling methods with ScannerBit, the GAMBIT scanning module

The GAMBIT Scanner Workgroup: Gregory D. Martinez[1,a], James McKay[2,b], Ben Farmer[3,4,c], Pat Scott[2,d], Elinore Roebber[5], Antje Putze[6], Jan Conrad[3,4]

[1]Physics and Astronomy Department, University of California, Los Angeles, CA 90095, USA
[2]Department of Physics, Imperial College London, Blackett Laboratory, Prince Consort Road, London SW7 2AZ, UK
[3]Oskar Klein Centre for Cosmoparticle Physics, AlbaNova University Centre, SE-10691 Stockholm, Sweden
[4]Department of Physics, Stockholm University, SE-10691 Stockholm, Sweden
[5]Department of Physics, McGill University, 3600 rue University, Montréal, Québec H3A 2T8, Canada
[6]LAPTh, Université de Savoie, CNRS, 9 chemin de Bellevue B.P.110, F-74941 Annecy-le-Vieux, France

Received: date / Accepted: date

Abstract We introduce ScannerBit, the statistics and sampling module of the public, open-source global fitting framework GAMBIT. ScannerBit provides a standardised interface to different sampling algorithms, enabling the use and comparison of multiple computational methods for inferring profile likelihoods, Bayesian posteriors, and other statistical quantities. The current version offers random, grid, raster, nested sampling, differential evolu- ten or more dimensions, Diver substantially outperforms the other three samplers on all metrics.

Contents

ScannerBit algorithms

- ScannerBit contains custom code or interfaces for the following methods:

  Random

  Grid

  Markov Chain Monte Carlo (MCMC)

  Ensemble Monte Carlo

  Nested Sampling

  Differential evolution

- Let's review each of these in turn...

Random and Grid scanners

- Random: just sample points randomly from the space within some box specified as a prior range on each parameter
- Grid: Scan along each axis (within some prior range)
- Not useful for serious applications: random sampling leads to biased inferences when applied to almost all problems
- Random and grid scanning both scale terribly with the number of dimensions in a problem

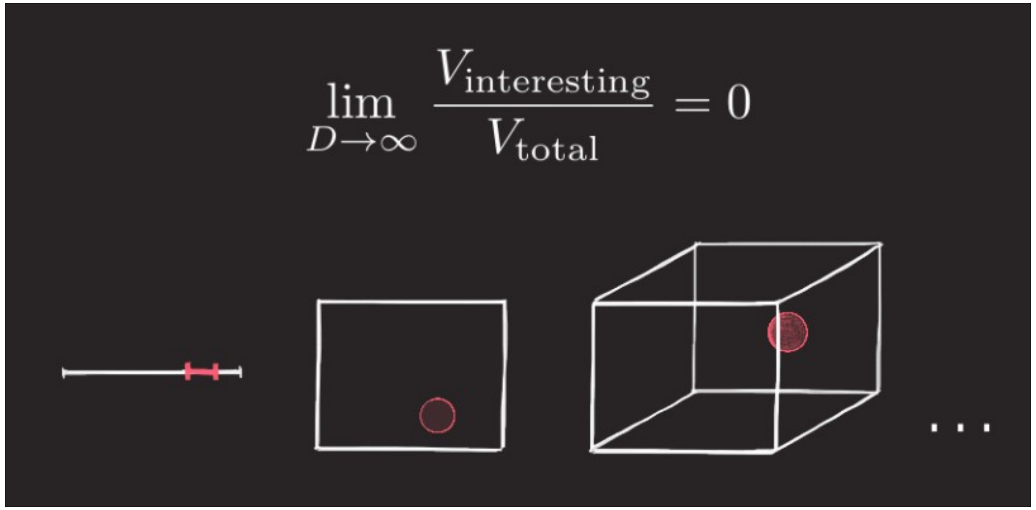

Source: Anders Kvellestad

Random and Grid scanners

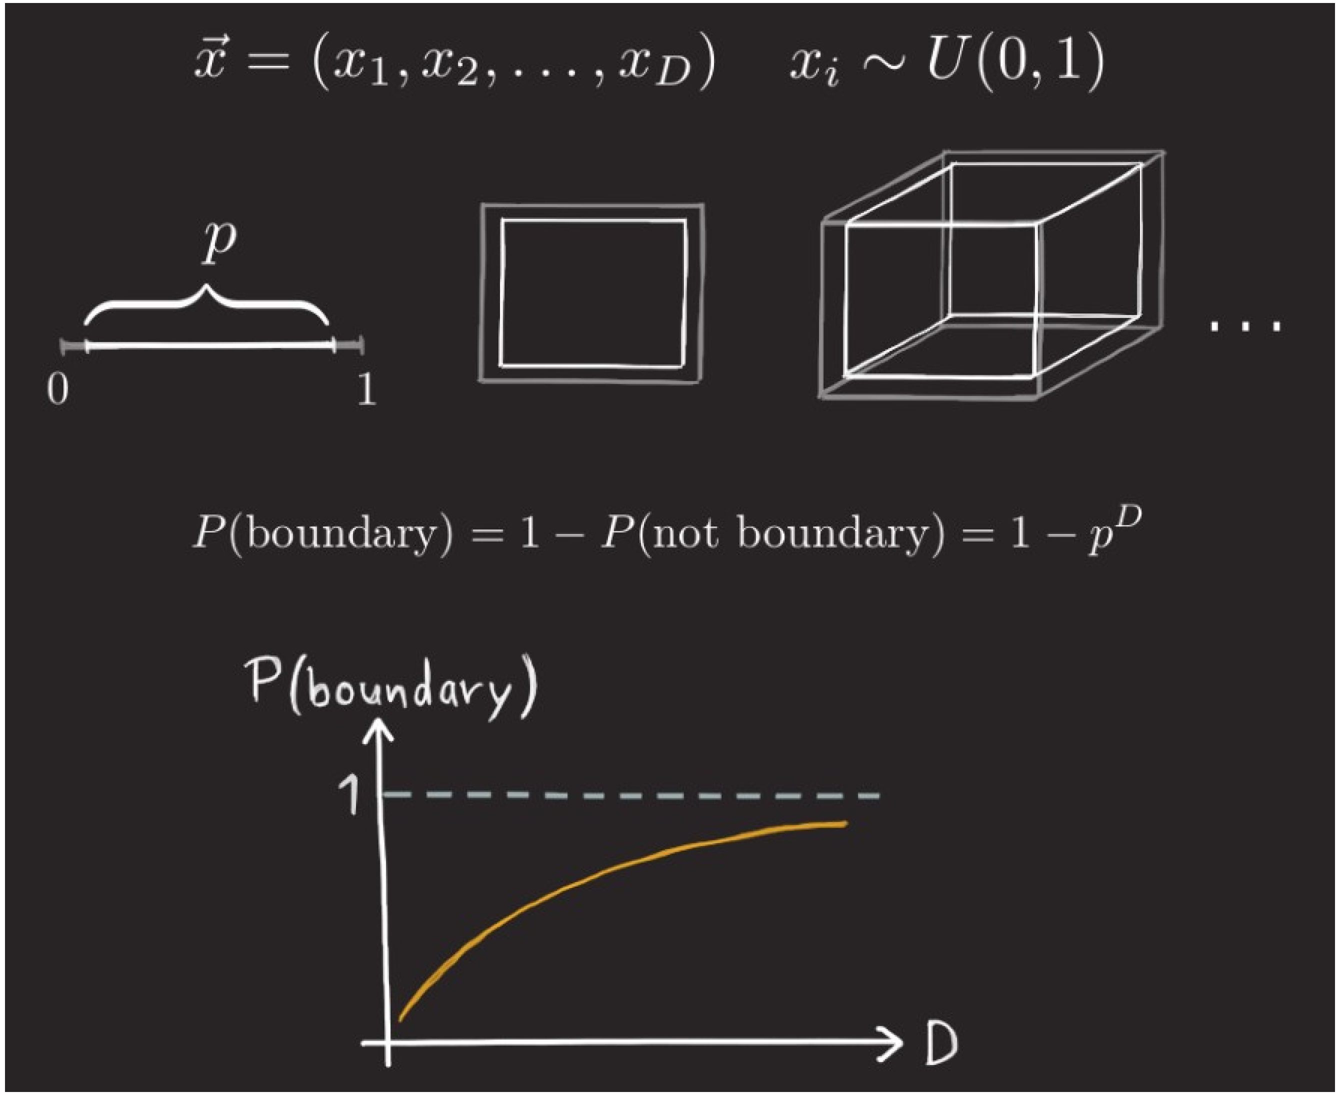

Source: Anders Kvellestad

Markov Chain Monte Carlo (MCMC) methods

- MCMC methods have been used for decades in cosmology and particle physics problems
- A popular approach is the *Metropolis-Hastings algorithm*:

1) Start at a randomly drawn initial point $\boldsymbol{\theta}_i$

2) Select another point $\boldsymbol{\theta}_{\text{trial}}$ at random using a *proposal* function $q(\boldsymbol{\theta}_{\text{trial}} | \boldsymbol{\theta}_i)$

3) The candidate point is accepted with the probability (p($\boldsymbol{\theta}$) = likelihood for flat prior on $\boldsymbol{\theta}$)

$$a(\boldsymbol{\theta}_{\text{trial}}|\boldsymbol{\theta}_i) = \min\left(1, \frac{p(\boldsymbol{\theta}_{\text{trial}})}{p(\boldsymbol{\theta}_i)} \frac{q(\boldsymbol{\theta}_i|\boldsymbol{\theta}_{\text{trial}})}{q(\boldsymbol{\theta}_{\text{trial}}|\boldsymbol{\theta}_i)}\right)$$

4) Set $\boldsymbol{\theta}_i = \boldsymbol{\theta}_{\text{trial}}$ if $\boldsymbol{\theta}_{\text{trial}}$ is accepted, else retain $\boldsymbol{\theta}_i$, then repeat procedure

Markov Chain Monte Carlo (MCMC) methods

- These points form a *Markov chain*, which spends time in the parameter space in proportion to the target posterior PDF of the parameters (given the supplied likelihood)
- For sufficiently long chains, one obtains independent samples from the target distribution p($\boldsymbol{\theta}$)
- To optimise efficiency, the proposal distribution *q* should match the (*a priori* unknown) true distribution
- GAMBIT includes an interface to the GreAT MCMC scanner that uses a multivariate Gaussian for *q*
- GreAT runs multiple chains, covariance matrix for chains is obtained from previous terminated chains (after thinning and removal of “burn-in” points)

Ensemble MCMC

- Standard MCMC is bad for high dimensional problems and/or multi-modal target functions
- Ensemble MCMC: run concurrent chains, each chain is individually advanced by contructing q from set of all points sampled by all chains
- GAMBIT includes the T-Walk ensemble MCMC method
- See ScannerBit paper for full details of how chains are advanced (depends on whether you are running the serial or parallelised version)

Nested sampling

- Nested sampling has been very popular in recent years, with many applications in particle physics, astronomy and cosmology
- It is much better at handling multimodal target functions than MCMC methods
- An efficient implementation is available in the public Multinest package, which GAMBIT makes use of

Nested sampling: a quick review of Bayesian inference

- Given a set of parameters **Θ** in a model *H*, plus some data **D**, Bayes' theorem gives:

likelihood

Posterior probability distribution → $$\Pr(\mathbf{\Theta}|\mathbf{D}, H) = \frac{\Pr(\mathbf{D}|\,\mathbf{\Theta}, H)\,\Pr(\mathbf{\Theta}|H)}{\Pr(\mathbf{D}|H)}$$ ← prior

- Denominator is a normalisation factor called the “Bayesian evidence”

$$\mathcal{Z} = \int \mathcal{L}(\mathbf{\Theta})\pi(\mathbf{\Theta})d^D\mathbf{\Theta},$$

- MCMC algorithms ignore *Z* (they give samples from the unnormalised posterior)

Nested sampling

- Nested sampling instead calculates *Z* directly by Monte Carlo integration
- Clever trick: define prior volume $dX = \pi(\mathbf{\Theta})d^D\mathbf{\Theta}$ $\quad X(\lambda) = \int_{\mathcal{L}(\mathbf{\Theta})>\lambda} \pi(\mathbf{\Theta})d^D\mathbf{\Theta}$
- Can then write evidence integral as:

$$\mathcal{Z} = \int_0^1 \mathcal{L}(X)dX,$$

- Monotonically decreasing function of X
  - Draw *N* "live points" from prior, at each iteration replace the lowest likelihood samples with higher likelihood samples, repeat until prior volume has been traversed

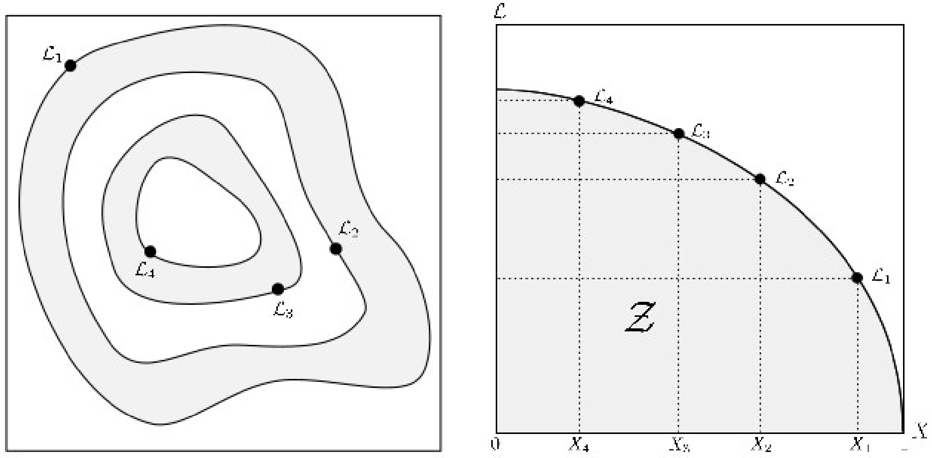

(a) (b)

Differential evolution: Diver

- Optimisation algorithm, good for multimodal posteriors in high dimensional spaces
- A simple explanation is as follows:

1) Start with a random selection of points in the parameter space (called “vectors”)

2) *Mutate* vectors by e.g. picking three random vectors and making: (**V** = “donor vector”)

$$\mathbf{V}_i = \mathbf{X}_{r1} + F(\mathbf{X}_{r2} - \mathbf{X}_{r3}).$$

3) *Crossover* the donor vectors and original vectors by making *trial vectors* **U** that have a random selection of components from the original vectors and the donor vectors

4) *Select* the vectors by computing the likelihood for the original vectors and their associated trial vectors, and choosing the highest likelihood vector for the next generation

- See ScannerBit paper for full details of GAMBIT implementation

Scanner comparisons

- GAMBIT allows the scanner to be swapped by changing one line in a yaml file
- This offers a unique test bed for comparison of scanning algorithms
- Have compared algorithms on a non-trivial physics example: scalar singlet DM

$$\mathcal{L} = \frac{1}{2}\mu_S^2 S^2 + \frac{1}{2}\lambda_{hS} S^2 |H|^2 + \frac{1}{4}\lambda_S S^4 + \frac{1}{2}\partial_\mu S \partial^\mu S$$

- Constraints from direct and indirect DM detection experiments, LHC Higgs invisible width searches, relic density upper bound plus theoretical upper bound on the Higgs-singlet coupling

Singlet DM

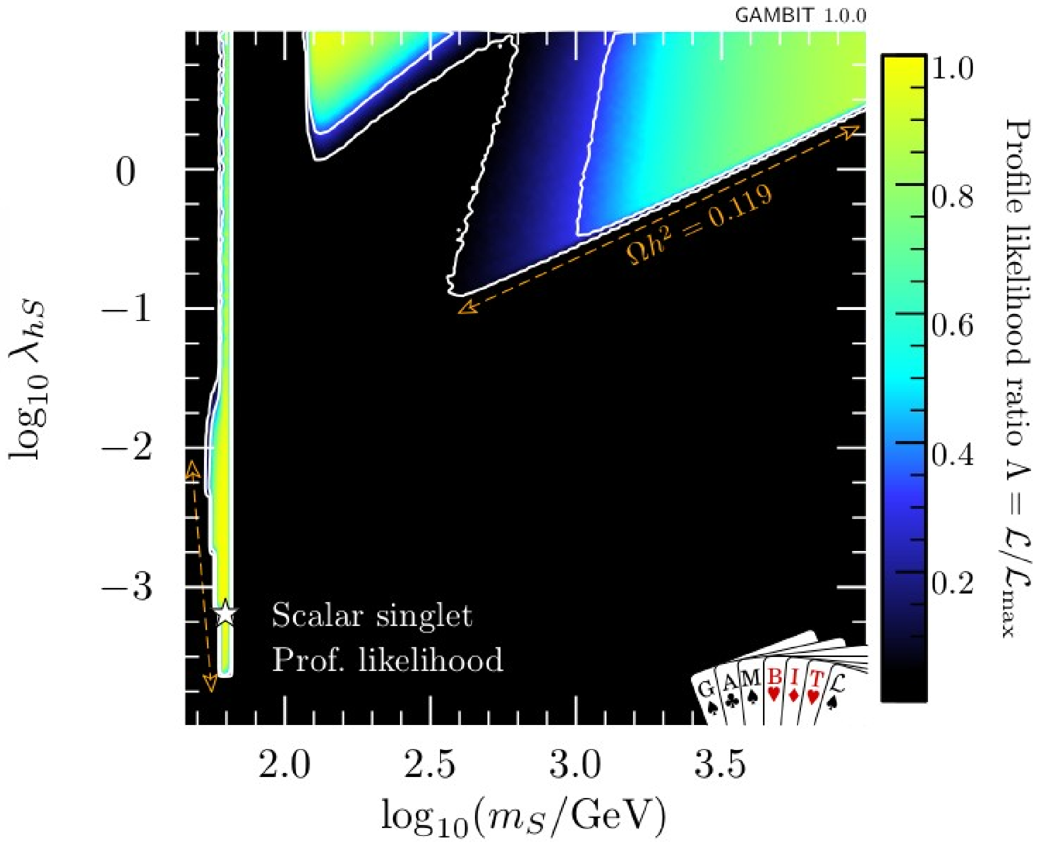

Parameter		Values
Scalar pole mass	m_S	$[45, 10^4]$ GeV
Higgs portal coupling	λ_{hS}	$[10^{-4}, 10]$
Varied in 7 and 15-dimensional scans		
Electromagnetic coupling	$1/\alpha^{\overline{MS}}(m_Z)$	127.940(42)
Strong coupling	$\alpha_s^{\overline{MS}}(m_Z)$	0.1185(18)
Top pole mass	m_t	173.34(2.28) GeV
Higgs pole mass	m_h	125.7(1.6) GeV
Local dark matter density	ρ_0	$0.4^{+0.4}_{-0.2}$ GeV cm^{-3}
Varied in 15-dimensional scans		
Nuclear matrix el. (strange)	σ_s	43(24) MeV
Nuclear matrix el. (up + down)	σ_l	58(27) MeV
Fermi coupling $\times 10^5$	$G_{F,5}$	1.1663787(18)
Down quark mass	$m_d^{\overline{MS}}(2\,\text{GeV})$	4.80(96) MeV
Up quark mass	$m_u^{\overline{MS}}(2\,\text{GeV})$	2.30(46) MeV
Strange quark mass	$m_s^{\overline{MS}}(2\,\text{GeV})$	95(15) MeV
Charm quark mass	$m_c^{\overline{MS}}(m_c)$	1.275(75) GeV
Bottom quark mass	$m_b^{\overline{MS}}(m_b)$	4.18(9) GeV

$$\mathcal{L} = \frac{1}{2}\mu_S^2 S^2 + \frac{1}{2}\lambda_{hS} S^2 |H|^2 + \frac{1}{4}\lambda_S S^4 + \frac{1}{2}\partial_\mu S \partial^\mu S$$

Scanner performance vs number of dimensions

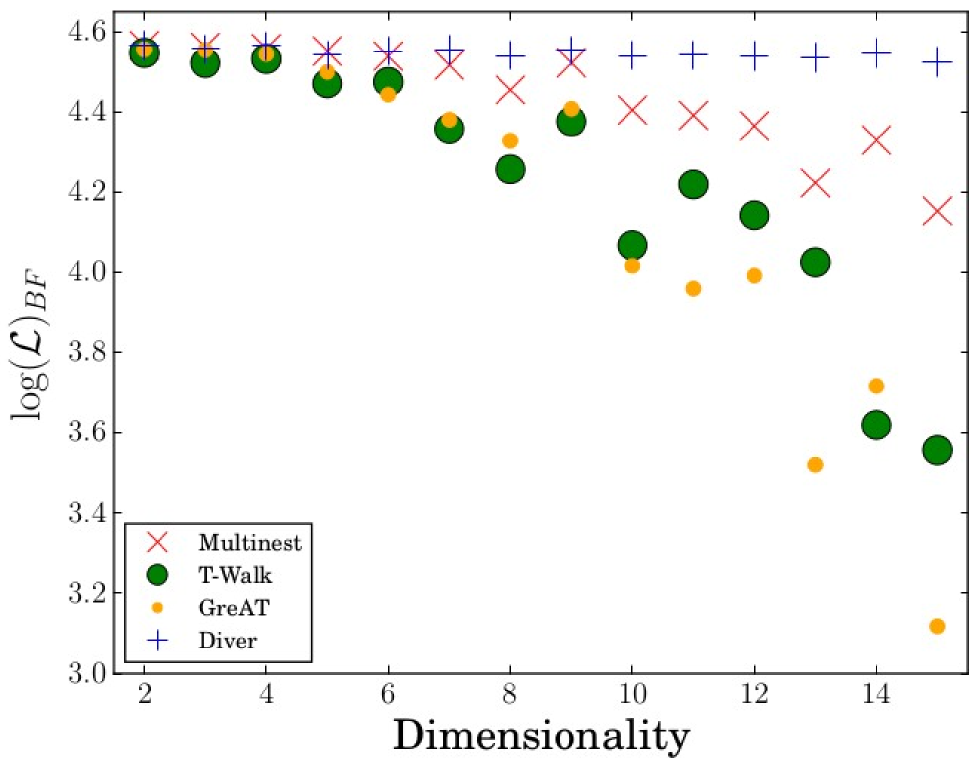

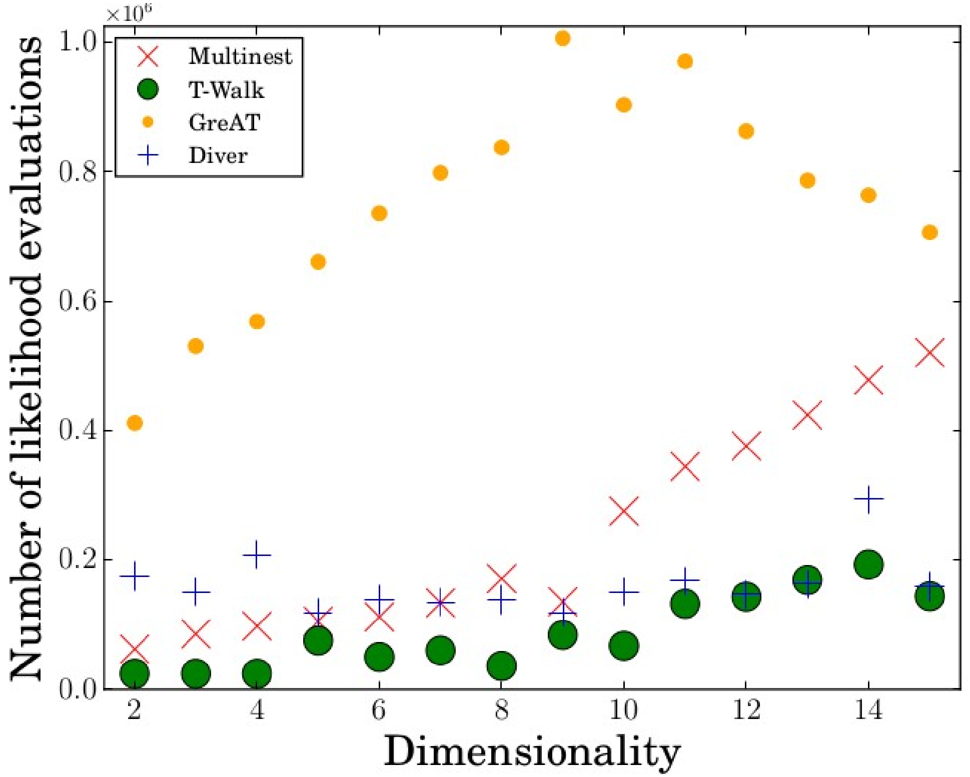

Diver: `NP` $= 20\,000$, `convthresh` $= 10^{-3}$

MultiNest: `nlive` $= 20\,000$, `tol` $= 10^{-3}$

T-Walk: `chain_number` $=$ number of MPI processes $+$ $N_{\rm dim} + 1$, $tol =$ `sqrtR` $- 1 = 0.05$

GreAT: `nTrials` $= 2000$, `nTrialsList` $= N_{\rm dim} + 1$

Real time vs number of likelihood evaluations

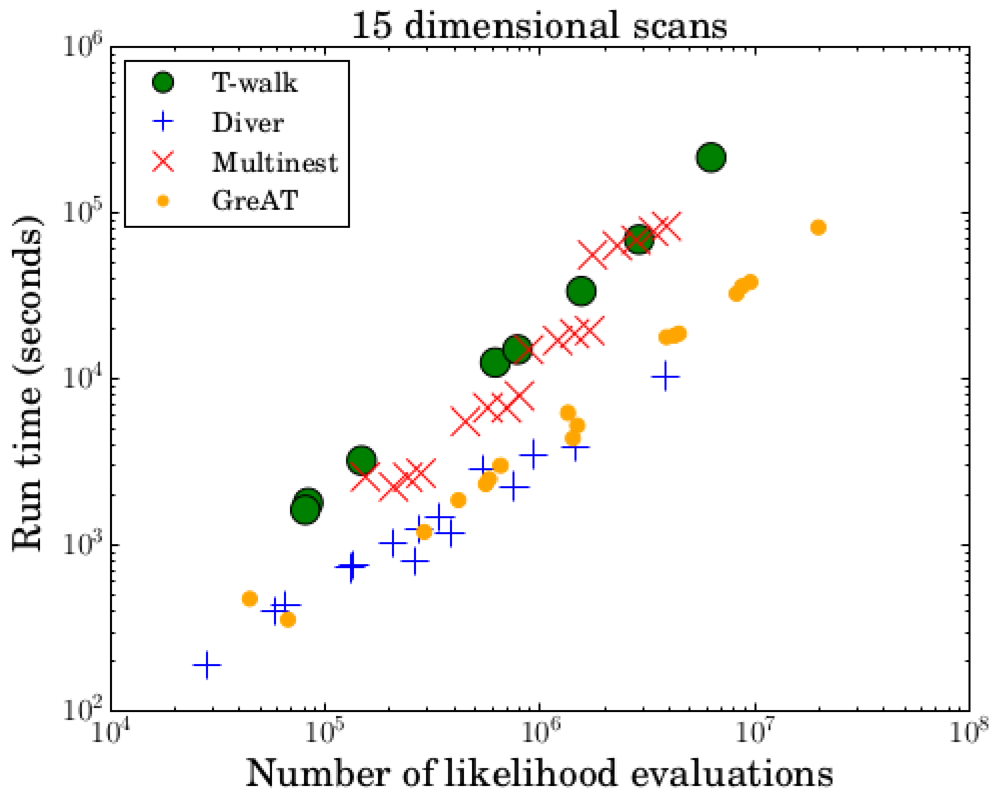

- T-Walk and Multinest are less efficient (per likelihood evaluation) than GreAT and Diver for large dimensional problems
- There are several reasons (e.g. ellipsoidal decomposition in Multinest, chain advancement calculations in T-Walk, MPI bottlenecks, etc)

Posterior mapping: 15D scan using GreAT

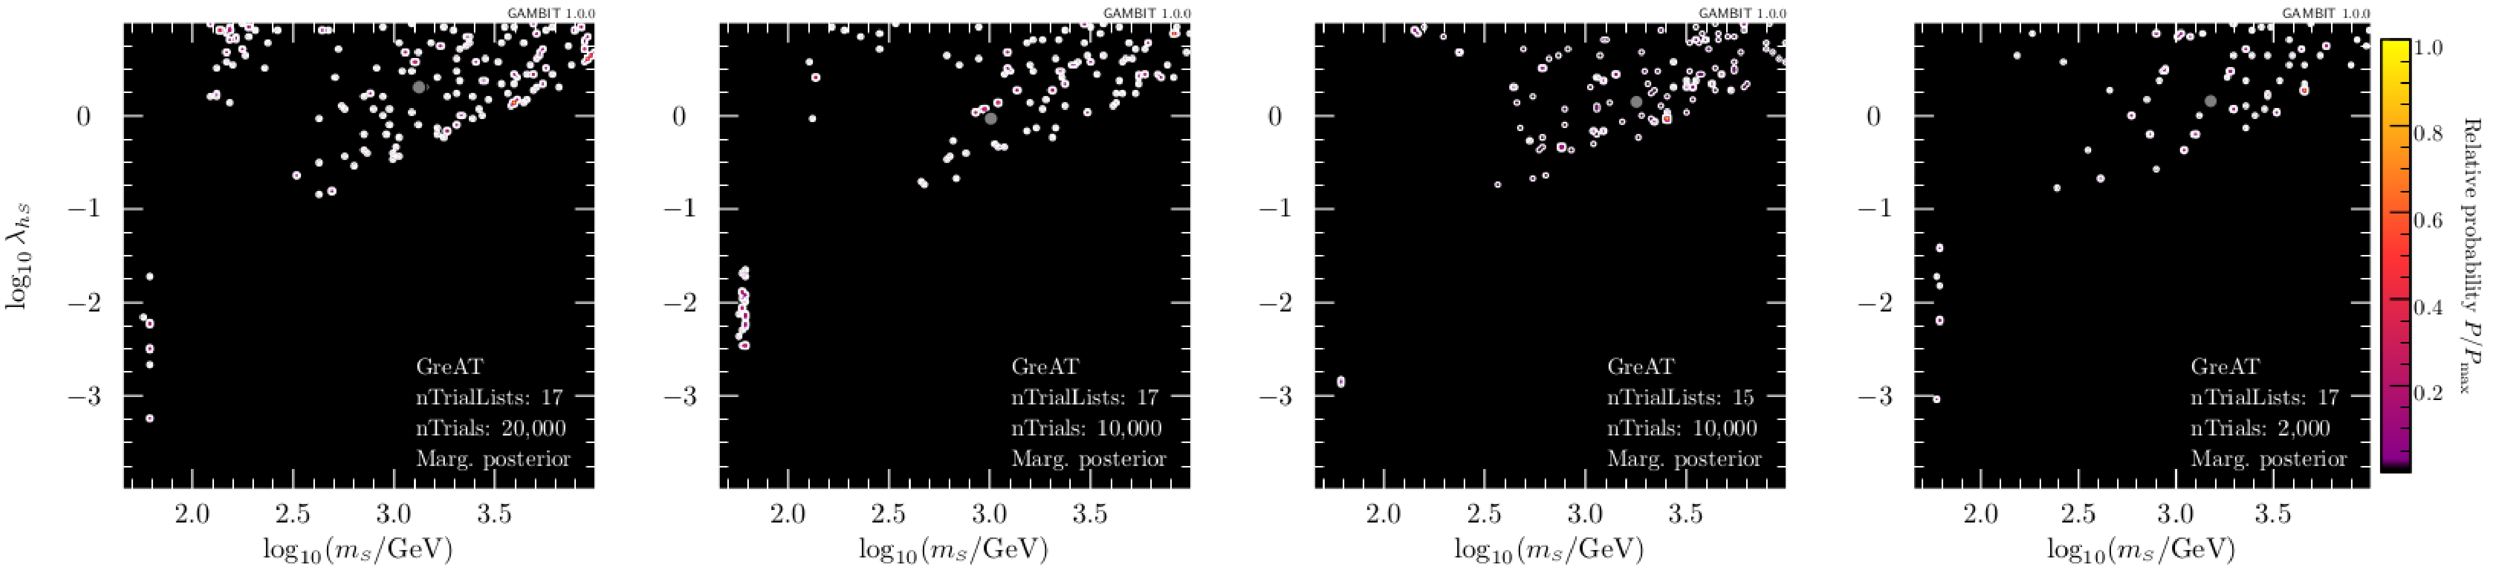

- Yikes! Validates assertion that MCMC algorithms do not cope well with multimodal posteriors?

Posterior mapping: 15D scan using T-Walk

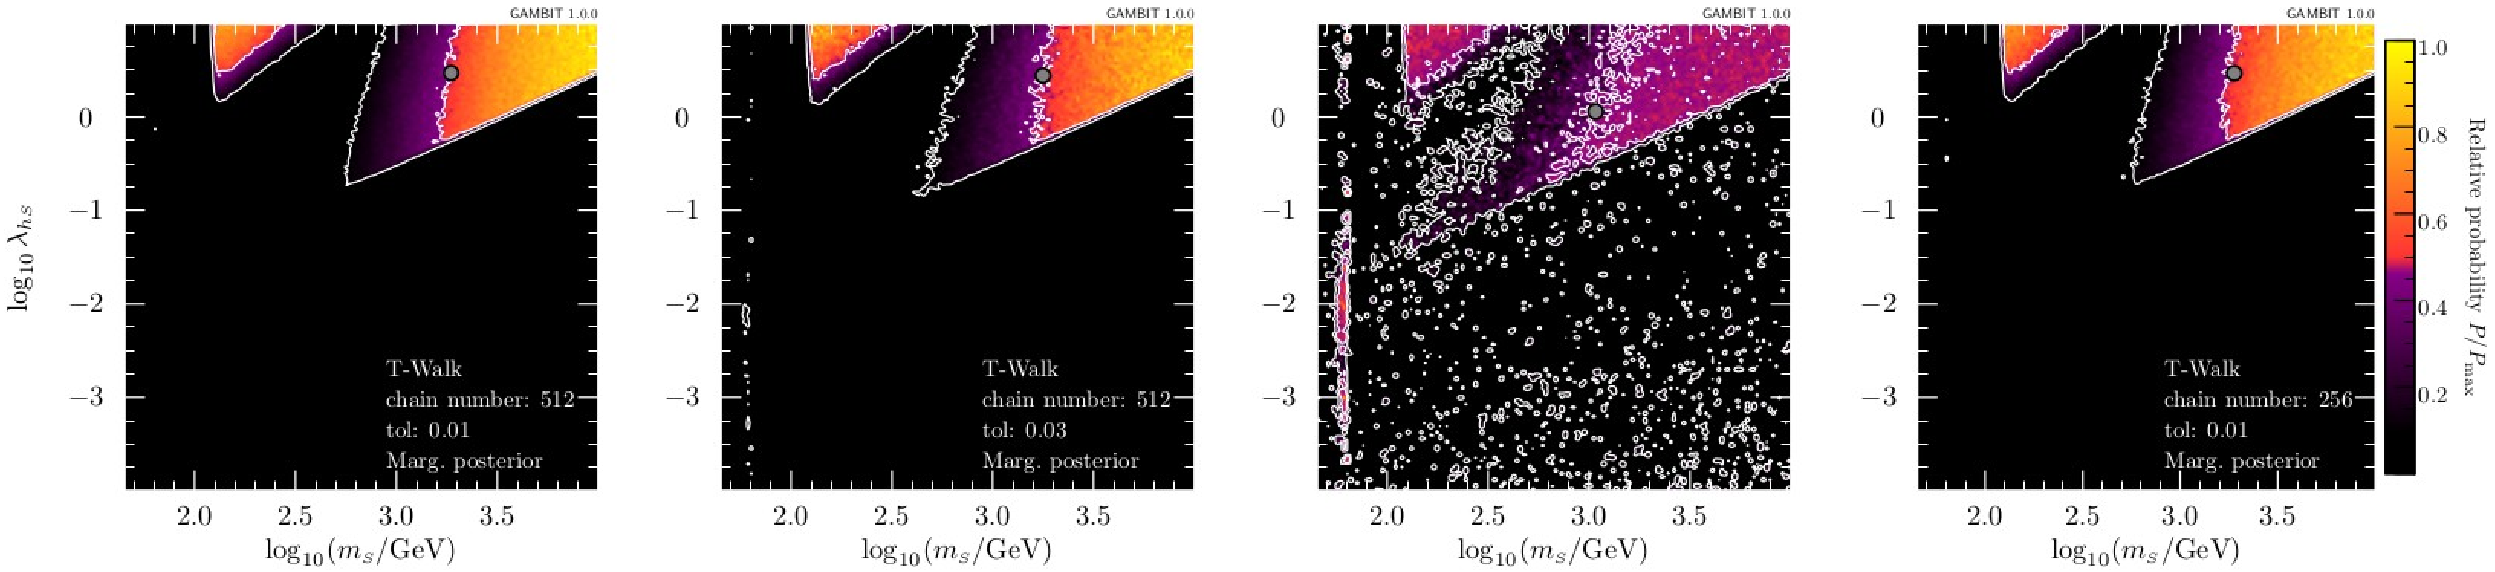

- The best scan here was the best posterior obtained, taking 9h in total
- Poorly converged scans find all modes, but don't get relative weight correct (and don't map the posterior smoothly)

Posterior mapping: 15D scan using Multinest

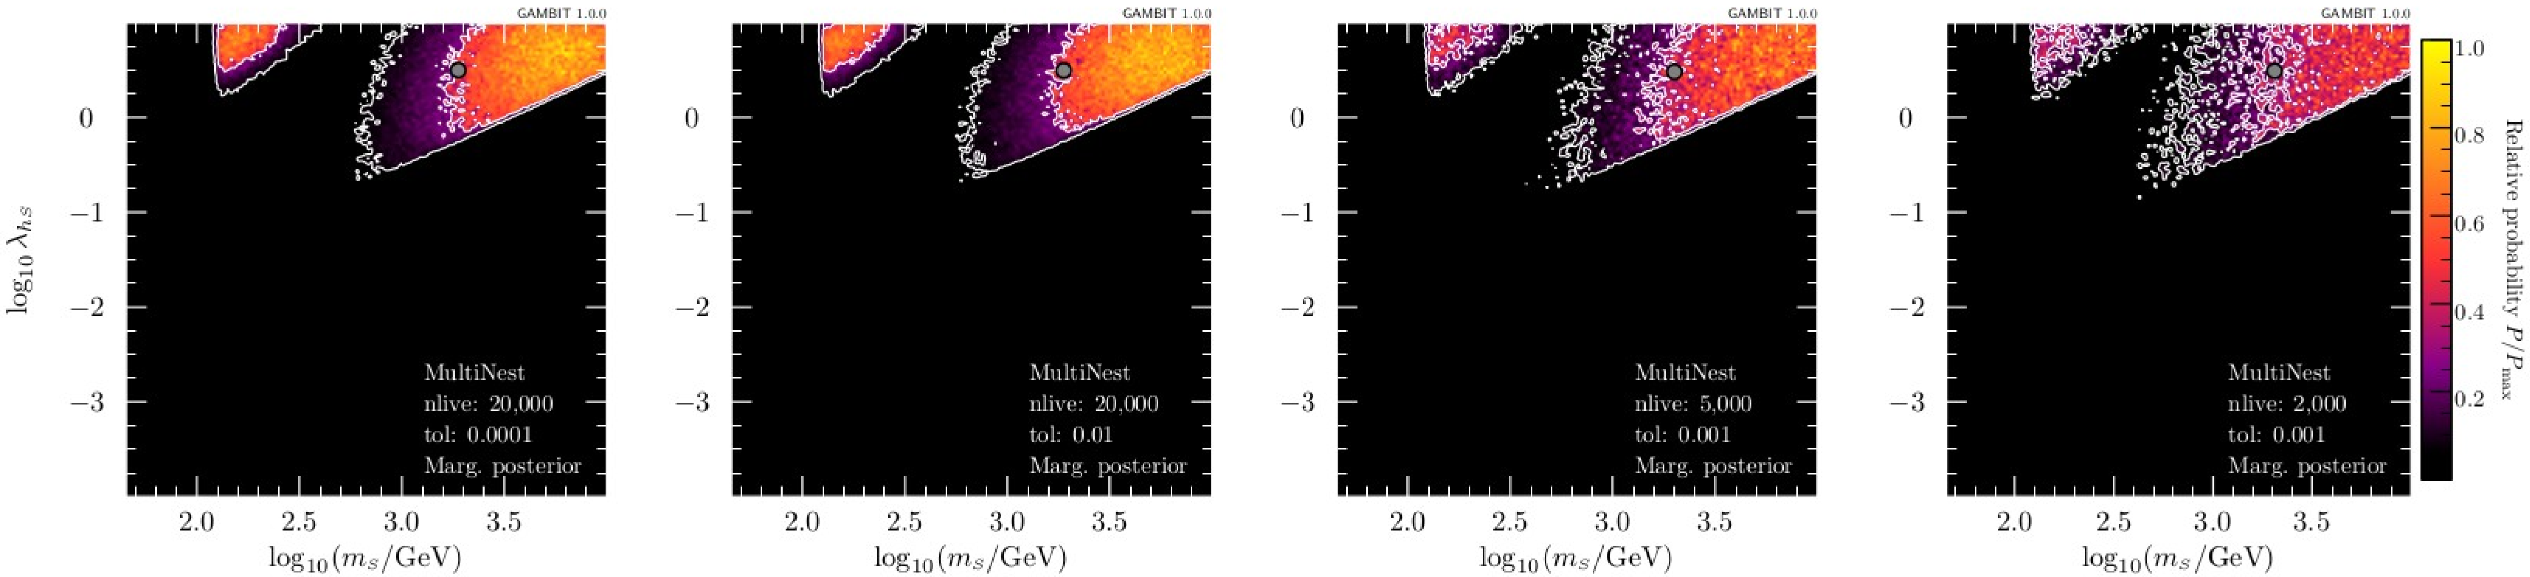

- Scans with too few live points or too high a tolerance do not find all modes
- The best scan here took > 21h, and is not as smooth as the T-Walk results
- Multinest also erroneously smooths sharp features due to its ellipsoidal sampling method

Sampling: executive summary

- Quick version: if you are a frequentist, use Diver. If you are a Bayesian, use T-Walk or Multinest.
- Longer version: Different algorithms have different quirks, can exploit this to gain insights. e.g. use T-Walk to find modes, then focus Multinest scans on those modes
- Using different algorithms with GAMBIT is both beneficial and easy!

Machine learning in GAMBIT

- GAMBIT explorations of large spaces require hundreds of millions of likelihood evaluations
- Shaving seconds from a slow likelihood evaluation can have a profound impact on the CPU hours required for a convergent fit
- Our slowest likelihood *by far* is the LHC sparticle search likelihood, the bottleneck being Pythia MC event generation
- Interpolation of a pre-simulated grid of likelihoods is a classic solution to this problem
- Downside is that this must be repeated for any physics model of interest

Machine learning for interpolation (CMSSM)

Buckley, Shilton, MJW, Comput.Phys.Commun. 183 (2012) 960-970

- Showed 5 years back that ML-based regression works in a 4D CMSSM (SVM or BNN)
- Interpolated the simulated event yield in the ATLAS 0 lepton analysis, and reproduced the exclusion limit using the ML output
- Since then: SUSY-AI, SCYNet + ??

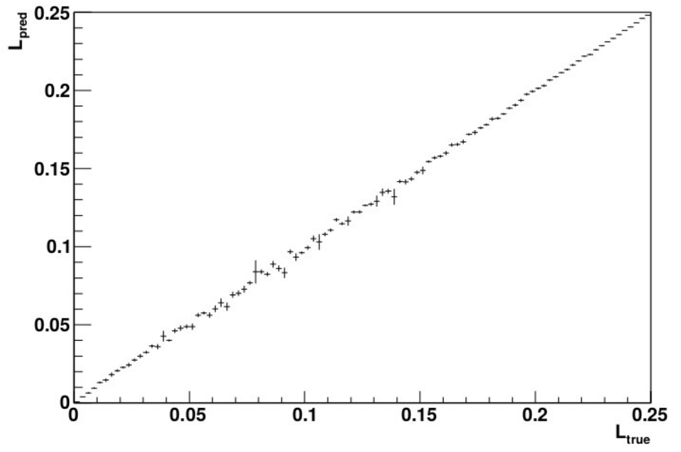

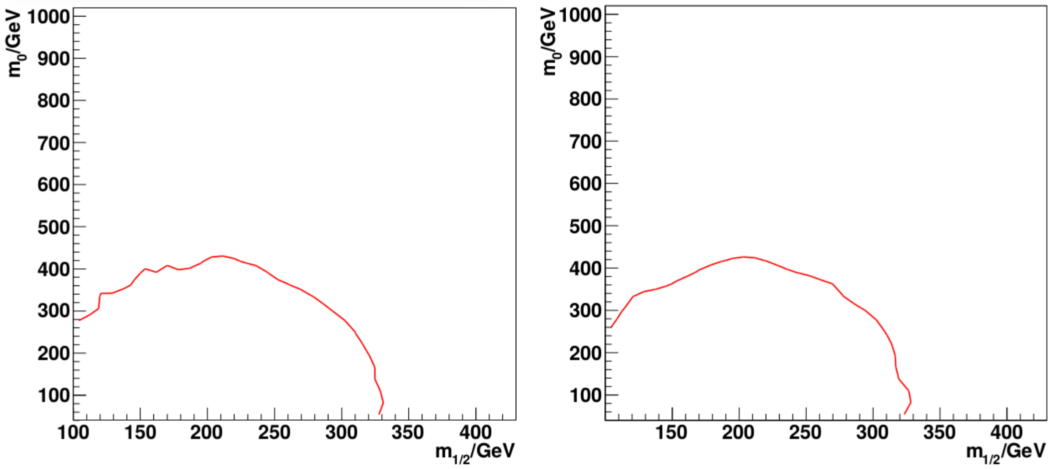

Interpolation of SUSY NLO cross-sections

- For speed, our current GAMBIT SUSY scans use the LO+LL cross-section from Pythia (already have this “for free”)
- A ML-based interpolation of NLO cross-sections was meant to be in the first release
- Progress is being made for forthcoming GAMBIT releases
- In future, can also use SUSY AI as a “fast reject” system
- Note: GAMBIT samples are public, and can be used to test and hone SUSY-AI performance

Future applications: replacing Pythia for relevant cases

- Pythia is only satisfactory if tree-level diagrams give you most of the answer
- NLO effects substantially modify the kinematics for e.g. monojet searches (see e.g. Buckley, Feld, Goncalves, Phys.Rev. D91 (2015) 015017)
- Even in SUSY, the acceptance for models with compressed spectra is highly dependent on the initial state radiation model used, and Pythia is deficient relative to e.g. Madgraph with explicit radiation of extra partons
- Can we write a code that contain interpolated yields for interesting cases? e.g. DM simplified models? Compressed SUSY EW sector? These yields could be reweighted depending on the couplings in the model.

Summary

- GAMBIT is an open source, *public* code for global statistical fits of new physics models
- It has so far proven very versatile for WIMP and non-WIMP dark matter physics, and we have an active physics programme for studies of new models
- ML can clearly have a profound impact on our total likelihood evaluation time (on a model-by-model basis)
- Reduced calculation times $\Rightarrow$ more physics quicker!