Title: Including astroparticle observables in global fits to new physics scenarios

Date: Oct 15, 2013 01:00 PM

URL: http://pirsa.org/13100103

Abstract: Searches for physics beyond the standard model come in
many forms, from terrestrial probes to astroparticle experiments and
cosmological observations. Efforts to
combine multiple search channels in 'global fits' to new physics scenarios
typically consider only a subset of the available channels. Astroparticle searches in particular are
usually only included in a very approximate way, if at all. In this talk I will review recent progress in
including detailed gamma-ray, neutrino and CMB searches for dark matter in
global fits. I will also preview some of
the future developments and challenges in this field, where the applicability
of global fits will move well beyond the small range of constrained
supersymmetric models they have so far mostly been applied to.

Including astroparticle observables in global fits to new physics scenarios

Pat Scott

Department of Physics, McGill University

Slides available from
http://www.physics.mcgill.ca/~patscott

Including astroparticle observables in global fits to new physics scenarios

Pat Scott

Department of Physics, McGill University

Slides available from
http://www.physics.mcgill.ca/~patscott

Outline

1. The Problem
2. Progress
   - Gamma-rays
   - Neutrinos
   - CMB constraints
3. Future Challenges
   - Respectable LHC likelihoods
   - Statistical/numerical issues
   - Parameter space $\longrightarrow$ Theory space

Searching for new physics

Many reasons to look for physics Beyond the Standard Model (BSM):

- Higgs mass (hierarchy problem + vacuum stability)
- Dark matter exists
- Baryon asymmetry
- Neutrino masses and mixings

So what do we do about it?

- Make new particles at high-E colliders
- Study rare processes at high-L colliders
- Hunt for dark matter
- Look for kooky neutrino physics

Combining searches I

Question

How do we know which models are in and which are out?

Combining searches I

Question

How do we know which models are in and which are out?

Answer

Combine the results from different searches

- Simplest method: take different exclusions, overplot them, conclude things are “allowed” or “excluded”
- Simplest BSM example: the scalar singlet model

(Cline, Kainulainen, PS & Weniger, *PRD*, 1306.4710)

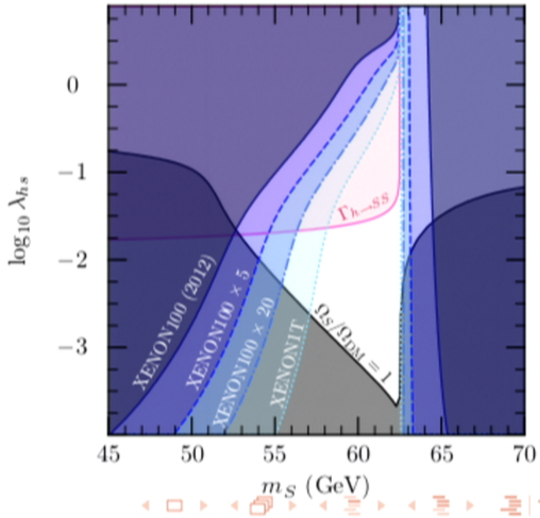

Combining searches I

Question

How do we know which models are in and which are out?

Answer

Combine the results from different searches

- Simplest method: take different exclusions, overplot them, conclude things are “allowed” or “excluded”
- Simplest BSM example: the scalar singlet model

(Cline, Kainulainen, PS & Weniger, *PRD*, 1306.4710)

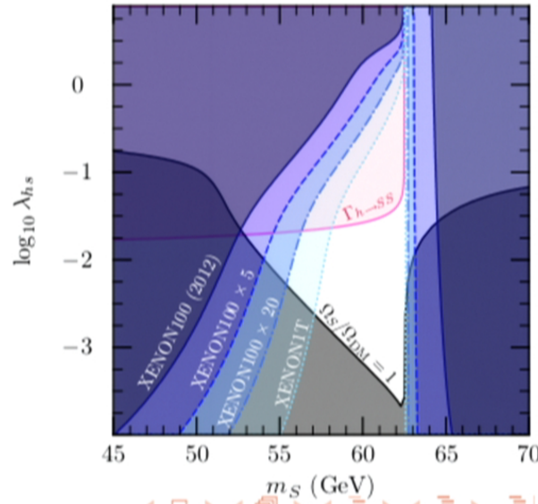

Combining searches II

That's all well and good if there are only 2 parameters and few searches. . .

Question

What if there are many different constraints?

Combining searches II

That's all well and good if there are only 2 parameters and few searches. . .

Question

What if there are many different constraints?

Answer

Combine constraints in a statistically valid way
→ composite likelihood

(Cline, Kainulainen, PS & Weniger, *PRD*, 1306.4710)

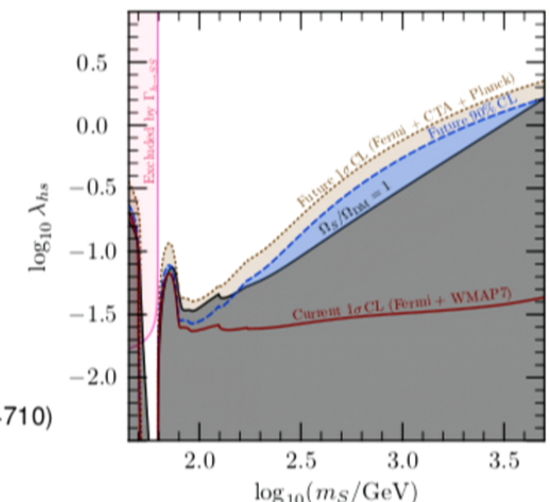

Combining searches II

That's all well and good if there are only 2 parameters and few searches. . .

Question

What if there are many different constraints?

Answer

Combine constraints in a statistically valid way
→ composite likelihood

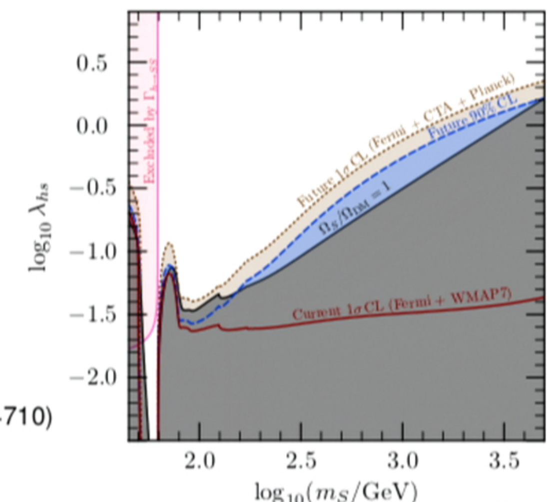

(Cline, Kainulainen, PS & Weniger, *PRD*, 1306.4710)

Combining searches III

That's all well and good if there are only 2 parameters and few searches. . .

Question

What if there are many parameters?

Combining searches III

That's all well and good if there are only 2 parameters and few searches. . .

Question

What if there are many parameters?

Answer

Need to

- scan the parameter space (smart numerics)
- interpret the combined results (Bayesian / frequentist)
- project down to parameter planes of interest (marginalise / profile)

→ global fits

Combining searches III

That's all well and good if there are only 2 parameters and few searches. . .

Question

What if there are many parameters?

Answer

Need to

- scan the parameter space (smart numerics)
- interpret the combined results (Bayesian / frequentist)
- project down to parameter planes of interest (marginalise / profile)

$\rightarrow$ global fits

Know your (SUSY) scans

Global fits:

Quantitative?
per-point: always
overall: always

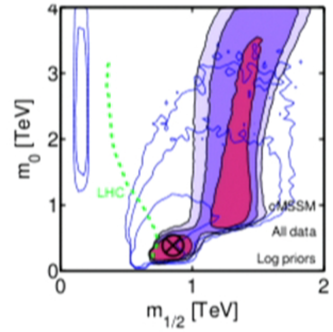

Strege et al *JCAP*, 1212.2636

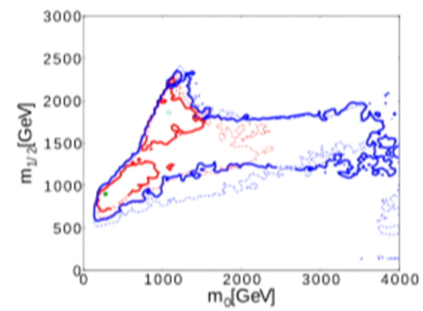

MasterCode, *EPJC*, 1207.7315

Not global fits:

Quantitative?
per-point: sometimes
overall: never

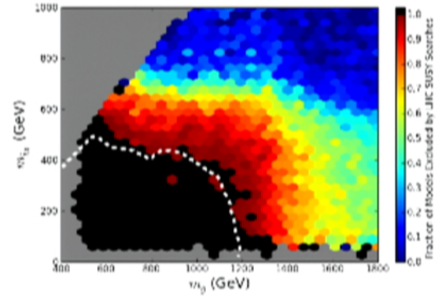

Cahill-Rowley et al, 1307.8444

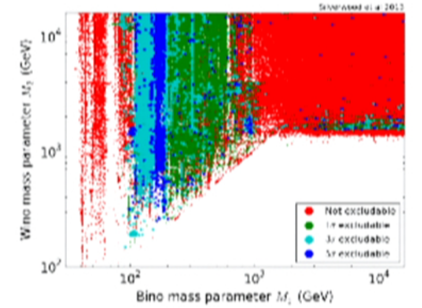

Silverwood, PS, et al, *JCAP*, 1210.0844

Know your (SUSY) scans

Global fits:

Quantitative?
- per-point: always
- overall: always

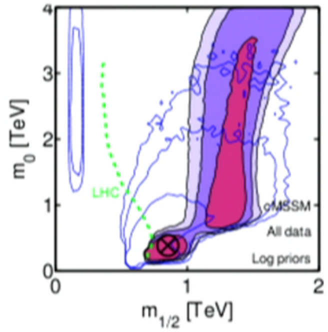

Strege et al *JCAP*, 1212.2636

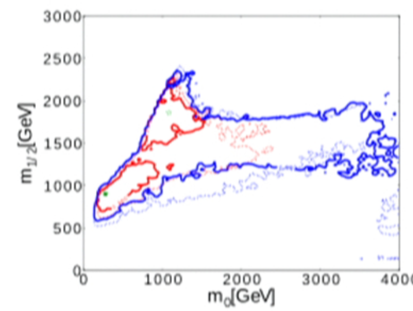

MasterCode, *EPJC*, 1207.7315

Not global fits:

Quantitative?
- per-point: sometimes
- overall: never

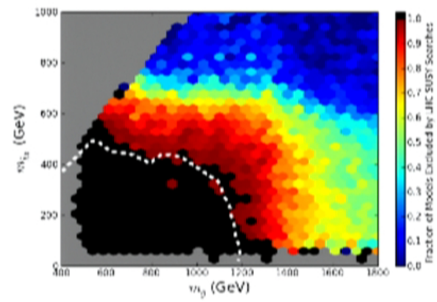

Cahill-Rowley et al, 1307.8444

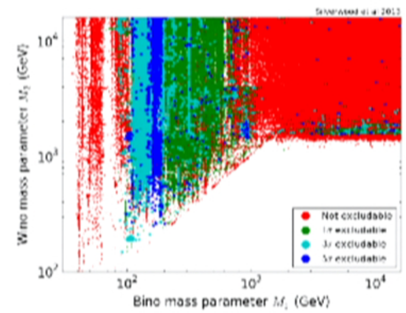

Silverwood, PS, et al, *JCAP*, 1210.0844

BSM Model Scanning – Statistics 101

Goals:

1. Given a particular theory, determine which parameter combinations fit all experiments, and how well

2. Given multiple theories, determine which fit the data better, and quantify how much better

BSM Model Scanning – Statistics 101

Goals:

1. Given a particular theory, determine which parameter combinations fit all experiments, and how well
$\Longrightarrow$ parameter estimation
2. Given multiple theories, determine which fit the data better, and quantify how much better $\Longrightarrow$ model comparison

BSM Model Scanning – Statistics 101

Goals:

1. Given a particular theory, determine which parameter combinations fit all experiments, and how well
   $\Longrightarrow$ parameter estimation
2. Given multiple theories, determine which fit the data better, and quantify how much better $\Longrightarrow$ model comparison

Why simple IN/OUT analyses are not enough…

- Only partial goodness of fit, no measure of convergence, no idea how to generalise to regions or whole space.
- Frequency/density of models in IN/OUT scans is not proportional to probability $\Longrightarrow$ means essentially nothing.

Know your (SUSY) scans

Global fits:

Quantitative?
- per-point: always
- overall: always

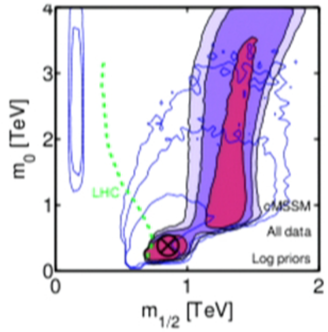

Strege et al *JCAP*, 1212.2636

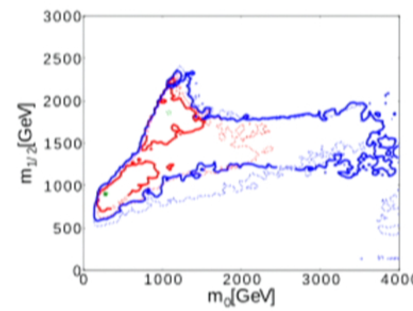

MasterCode, *EPJC*, 1207.7315

Not global fits:

Quantitative?
- per-point: sometimes
- overall: never

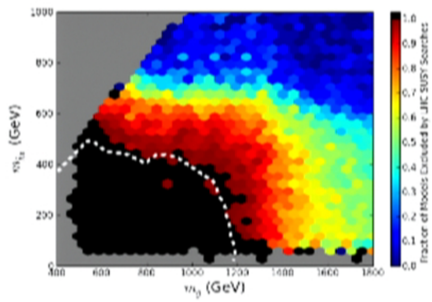

Cahill-Rowley et al, 1307.8444

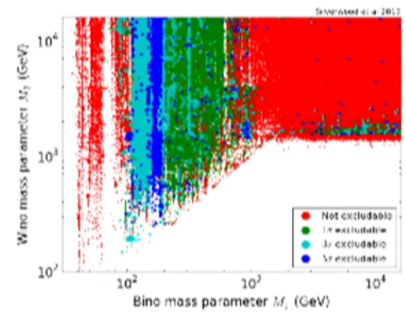

Silverwood, PS, et al, *JCAP*, 1210.0844

BSM Model Scanning – Statistics 101

Goals:

1. Given a particular theory, determine which parameter combinations fit all experiments, and how well
2. Given multiple theories, determine which fit the data better, and quantify how much better

Putting it all together: global fits

Issue 1: Combining fits to different experiments
Relatively easy – composite likelihood ($\mathcal{L}_1 \times \mathcal{L}_2 \equiv \chi_1^2 + \chi_2^2$ for simplest $\mathcal{L}$)

- dark matter relic density from WMAP
- precision electroweak tests at LEP
- LEP limits on sparticle masses
- B-factory data (rare decays, $b \to s\gamma$)
- muon anomalous magnetic moment
- rches, direct detection (only roughly implemented for now)

Putting it all together: global fits

Issue 1: Combining fits to different experiments
Relatively easy – composite likelihood ($\mathcal{L}_1 \times \mathcal{L}_2 \equiv \chi_1^2 + \chi_2^2$ for simplest $\mathcal{L}$)

- dark matter relic density from WMAP
- precision electroweak tests at LEP
- LEP limits on sparticle masses
- B-factory data (rare decays, $b \to s\gamma$)
- muon anomalous magnetic moment
- LHC searches, direct detection (only roughly implemented for now)

Putting it all together: global fits

Issue 2: Including the effects of uncertainties in input data
Easy – treat them as *nuisance parameters*

Issue 3: Finding the points with the best likelihoods
Tough – MCMCs, nested sampling, genetic algorithms, etc

Issue 4: Comparing theories
Depends – Bayesian model comparison, p values
(TS distribution? $\longrightarrow$ coverage???)

Two different approaches to including astro data in BSM scans

1. Just use the published limits on $\langle\sigma v\rangle$ (or $\sigma_{\mathrm{SI,SD}}$)
   - Fast – can cover large parameter spaces
   - Not so accurate – experimental limits are invariably based on theoretical assumptions, e.g. $b\bar{b}$ spectrum
   - Full likelihood function almost never available
2. Use the data points directly in BSM scans
   - Slow – requires full treatment of instrument profile for each point
   - Accurate – can test each point self-consistently
   - Allows marginalisation over theoretical assumptions
   - Allows construction of full multi-dimensional likelihood function

Two different approaches to including astro data in BSM scans

1. Just use the published limits on $\langle\sigma v\rangle$ (or $\sigma_{\mathrm{SI,SD}}$)
   - Fast – can cover large parameter spaces
   - Not so accurate – experimental limits are invariably based on theoretical assumptions, e.g. $b\bar{b}$ spectrum
   - Full likelihood function almost never available
2. Use the data points directly in BSM scans
   - Slow – requires full treatment of instrument profile for each point
   - Accurate – can test each point self-consistently
   - Allows marginalisation over theoretical assumptions
   - Allows construction of full multi-dimensional likelihood function

Two different approaches to including astro data in BSM scans

1. Just use the published limits on $\langle \sigma v \rangle$ (or $\sigma_{\mathrm{SI,SD}}$)
   - Fast – can cover large parameter spaces
   - Not so accurate – experimental limits are invariably based on theoretical assumptions, e.g. $b\bar{b}$ spectrum
   - Full likelihood function almost never available
2. Use the data points directly in BSM scans
   - Slow – requires full treatment of instrument profile for each point
   - Accurate – can test each point self-consistently
   - Allows marginalisation over theoretical assumptions
   - Allows construction of full multi-dimensional likelihood function
3. (indirect only: use just flux upper limits)

Gamma-rays

Gamma-ray annihilation searches in CMSSM global fits:

Fermi-LAT

Satellite pair conversion telescope

Dwarf galaxy Segue 1

(PS, Conrad et al *JCAP*, 0909.3300)

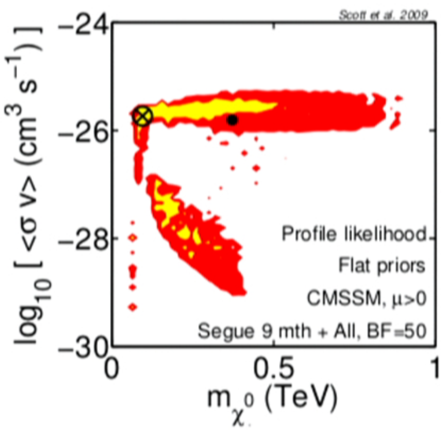

- Full binned Poissonian likelihood (no χ^2 approximation)
- Full treatment of PSF *and* energy dispersion (with fast convolution library FLATlib)
- Marginalisation over systematic error on effective area
- Diffuse BG from Fermi-LAT Galprop fits
- Isotropic BG best-fit isotropic power law
- *J*-factor from Martinez et al (*JCAP*, 0902.4715; best at the time)

Gamma-rays

Gamma-ray annihilation searches in CMSSM global fits:

Fermi-LAT

Satellite pair conversion telescope

Dwarf galaxy Segue 1

(PS, Conrad et al *JCAP*, 0909.3300)

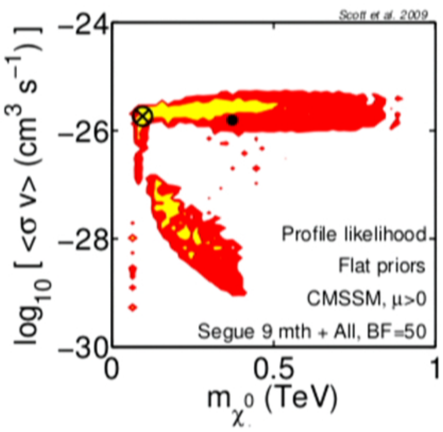

- Full binned Poissonian likelihood (no χ^2 approximation)
- Full treatment of PSF *and* energy dispersion (with fast convolution library FLATlib)
- Marginalisation over systematic error on effective area
- Diffuse BG from Fermi-LAT Galprop fits
- Isotropic BG best-fit isotropic power law
- *J*-factor from Martinez et al (*JCAP*, 0902.4715; best at the time)

Gamma-rays

Gamma-ray annihilation searches in CMSSM global fits:

HESS

Air Čerenkov telescope

Milky Way+Carina+Sculptor+Sag dwarf

(Ripken, Conrad & PS *JCAP*, 1012.3939)

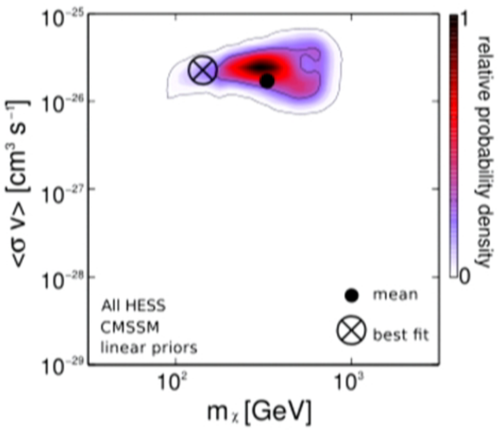

- χ^2-based analysis using public flux limits
- 'Milky Way' = halo just beyond GC (45–150 pc)
- Virtual internal bremsstrahlung from co-annihilation strip models caught at high-E by HESS
- but: J-factors for Sag dwarf rather uncertain

How to find DM with neutrino telescopes

The short version:

How to find DM with neutrino telescopes

The short version:

1. Halo WIMPs crash into the Sun

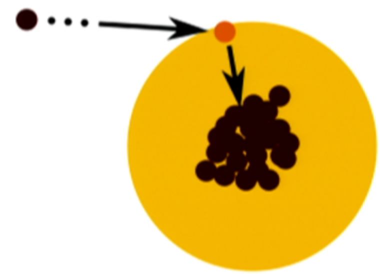

How to find DM with neutrino telescopes

The short version:

1. Halo WIMPs crash into the Sun
2. Some lose enough energy in the scatter to be gravitationally bound
3. Scatter some more, sink to the core

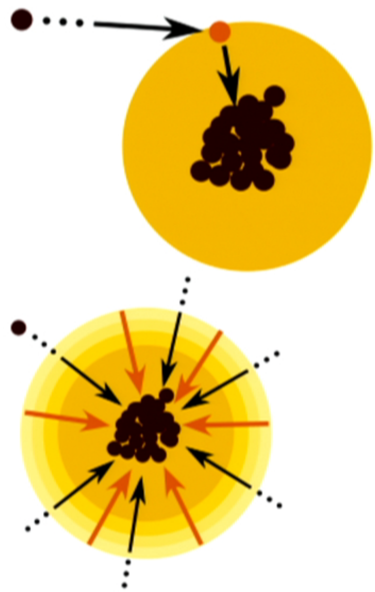

How to find DM with neutrino telescopes

The short version:

1. Halo WIMPs crash into the Sun
2. Some lose enough energy in the scatter to be gravitationally bound
3. Scatter some more, sink to the core
4. Annihilate with each other, producing neutrinos

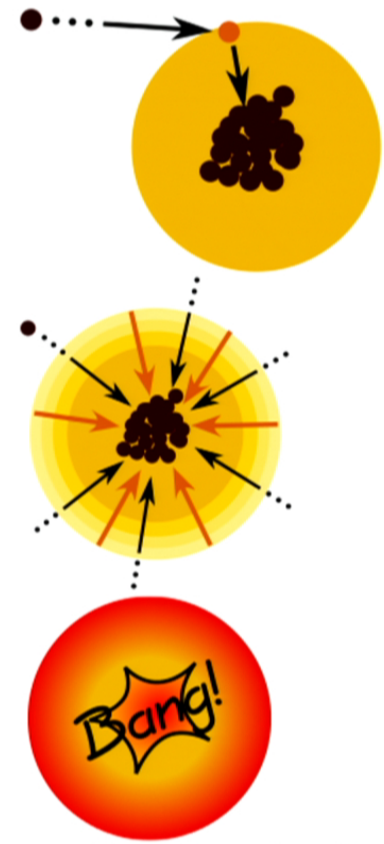

How to find DM with neutrino telescopes

The short version:

1. Halo WIMPs crash into the Sun
2. Some lose enough energy in the scatter to be gravitationally bound
3. Scatter some more, sink to the core
4. Annihilate with each other, producing neutrinos
5. Propagate+oscillate their way to the Earth, convert into muons in ice/water

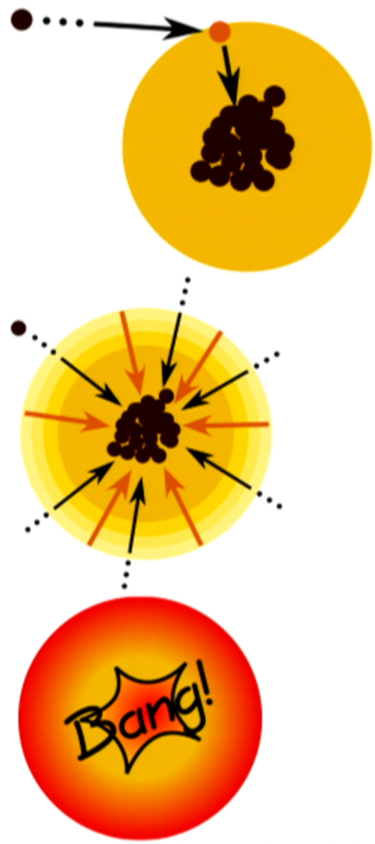

How to find DM with neutrino telescopes

The short version:

1. Halo WIMPs crash into the Sun
2. Some lose enough energy in the scatter to be gravitationally bound
3. Scatter some more, sink to the core
4. Annihilate with each other, producing neutrinos
5. Propagate+oscillate their way to the Earth, convert into muons in ice/water
6. Look for Čerenkov radiation from the muons in **IceCube**, ANTARES, etc

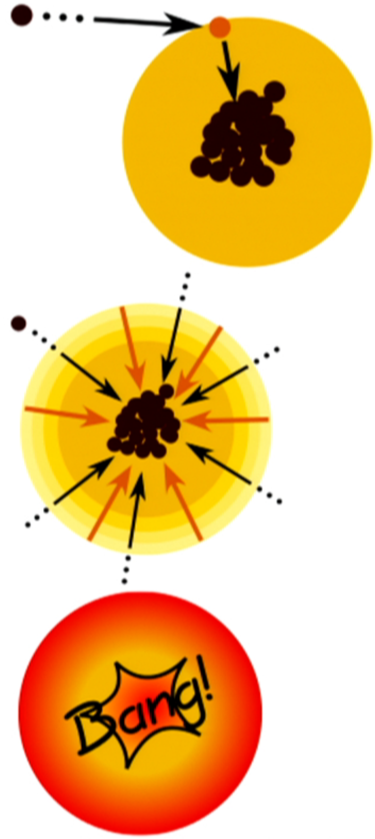

Advanced IceCube Likelihood for Model Testing

Simplest way to do anything is to first make it a counting problem. . .

Compare observed number of events n and predicted number θ for each model, taking into account error σ_ϵ on acceptance:

$$\mathcal{L}_{\rm num}(n|\theta_{\rm BG}+\theta_{\rm sig}) = \frac{1}{\sqrt{2\pi}\sigma_\epsilon}\int_0^\infty \frac{(\theta_{\rm BG}+\epsilon\theta_{\rm sig})^n e^{-(\theta_{\rm BG}+\epsilon\theta_{\rm sig})}}{n!}\frac{1}{\epsilon}\exp\left[-\frac{1}{2}\left(\frac{\ln\epsilon}{\sigma_\epsilon}\right)^2\right]\mathrm{d}\epsilon\,. \tag{1}$$

Nuisance parameter ϵ takes into account systematic errors on effective area, etc. $\sigma_\epsilon \sim 20\%$ for IceCube.

Advanced IceCube Likelihood for Model Testing

Simplest way to do anything is to first make it a counting problem...

Compare observed number of events n and predicted number θ for each model, taking into account error σ_ϵ on acceptance:

$$\mathcal{L}_{\mathrm{num}}(n|\theta_{\mathrm{BG}}+\theta_{\mathrm{sig}}) = \frac{1}{\sqrt{2\pi}\sigma_\epsilon}\int_0^\infty \frac{(\theta_{\mathrm{BG}}+\epsilon\theta_{\mathrm{sig}})^n e^{-(\theta_{\mathrm{BG}}+\epsilon\theta_{\mathrm{sig}})}}{n!}\frac{1}{\epsilon}\exp\left[-\frac{1}{2}\left(\frac{\ln\epsilon}{\sigma_\epsilon}\right)^2\right]\mathrm{d}\epsilon\,. \tag{1}$$

Nuisance parameter ϵ takes into account systematic errors on effective area, etc. $\sigma_\epsilon \sim 20\%$ for IceCube.

Advanced IceCube Likelihood for Model Testing

Simplest way to do anything is to first make it a counting problem...

Compare observed number of events n and predicted number θ for each model, taking into account error σ_ϵ on acceptance:

$$\mathcal{L}_{\rm num}(n|\theta_{\rm BG}+\theta_{\rm sig}) = \frac{1}{\sqrt{2\pi}\sigma_\epsilon}\int_0^\infty \frac{(\theta_{\rm BG}+\epsilon\theta_{\rm sig})^n e^{-(\theta_{\rm BG}+\epsilon\theta_{\rm sig})}}{n!}\frac{1}{\epsilon}\exp\left[-\frac{1}{2}\left(\frac{\ln\epsilon}{\sigma_\epsilon}\right)^2\right]\mathrm{d}\epsilon\,. \tag{1}$$

Then: upgrade to full unbinned likelihood with number ($\mathcal{L}_{\rm num}$), spectral ($\mathcal{L}_{\rm spec}$) and angular ($\mathcal{L}_{\rm ang}$) bits:

$$\mathcal{L} = \mathcal{L}_{\rm num}(n|\theta_{\rm signal+BG})\prod_{i=1}^{n}\mathcal{L}_{{\rm spec},i}\,\mathcal{L}_{{\rm ang},i} \tag{2}$$

All available in DarkSUSY v5.0.6 and later: www.darksusy.org

Advanced IceCube Likelihood for Model Testing

Simplest way to do anything is to first make it a counting problem…

Compare observed number of events n and predicted number θ for each model, taking into account error σ_ϵ on acceptance:

$$\mathcal{L}_{\rm num}(n|\theta_{\rm BG}+\theta_{\rm sig}) = \frac{1}{\sqrt{2\pi}\sigma_\epsilon}\int_0^\infty \frac{(\theta_{\rm BG}+\epsilon\theta_{\rm sig})^n e^{-(\theta_{\rm BG}+\epsilon\theta_{\rm sig})}}{n!}\frac{1}{\epsilon}\exp\left[-\frac{1}{2}\left(\frac{\ln\epsilon}{\sigma_\epsilon}\right)^2\right]\mathrm{d}\epsilon\,. \tag{1}$$

Then: upgrade to full unbinned likelihood with number ($\mathcal{L}_{\rm num}$), spectral ($\mathcal{L}_{\rm spec}$) and angular ($\mathcal{L}_{\rm ang}$) bits:

$$\mathcal{L} = \mathcal{L}_{\rm num}(n|\theta_{\rm signal+BG})\prod_{i=1}^{n}\mathcal{L}_{{\rm spec},i}\,\mathcal{L}_{{\rm ang},i} \tag{2}$$

All available in DarkSUSY v5.0.6 and later: www.darksusy.org

CMSSM model reconstruction with IceCube event data

Benchmark recovery with 22-string IceCube WIMP-search neutrino events + full likelihood:

Mock signal: 60 events, $m_\chi = 500\,\mathrm{GeV}$, 100% $\chi\chi \to W^+W^-$

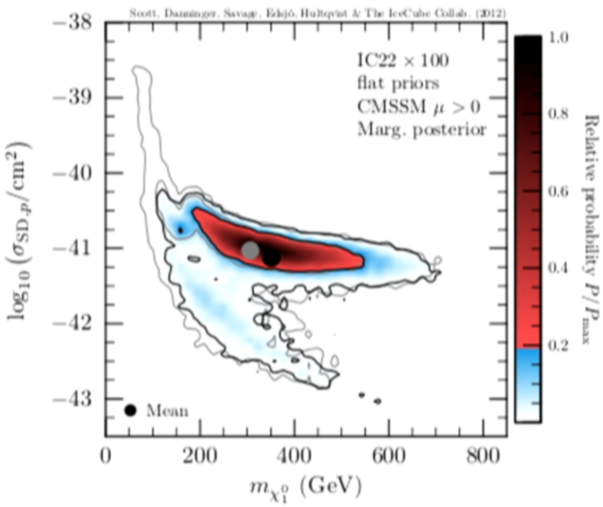

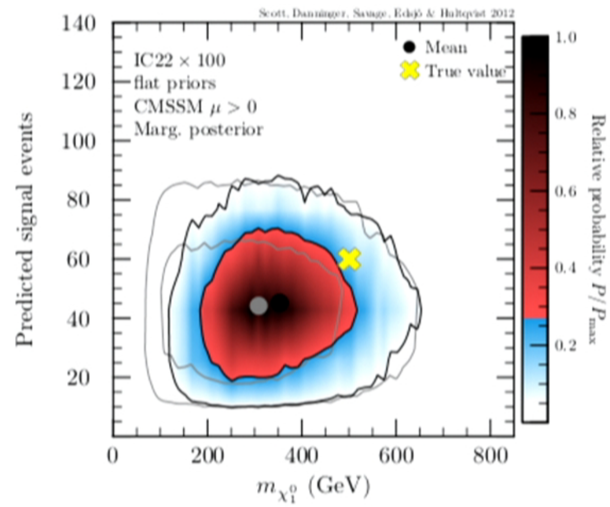

(PS, Savage, Edsjö & The IceCube Collaboration, *JCAP*, 1207.0810)

Example of Combined Direct + Indirect + LHC constraints

Base Observables

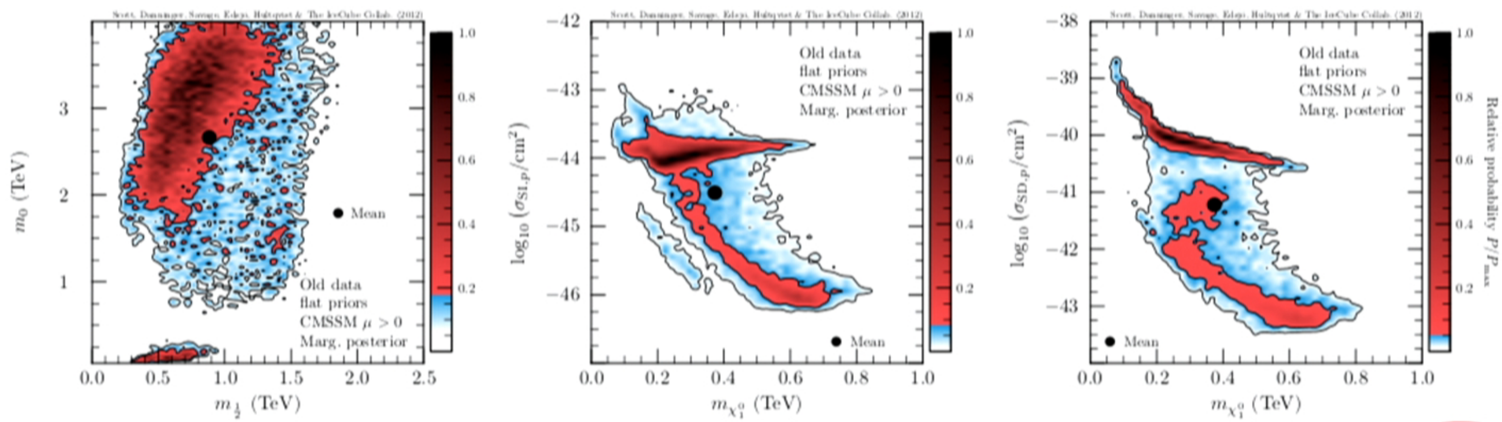

Example of Combined Direct + Indirect + LHC constraints

Base Observables + XENON-100 + CMS 5 fb^{-1}

Grey contours correspond to Base Observables *only*

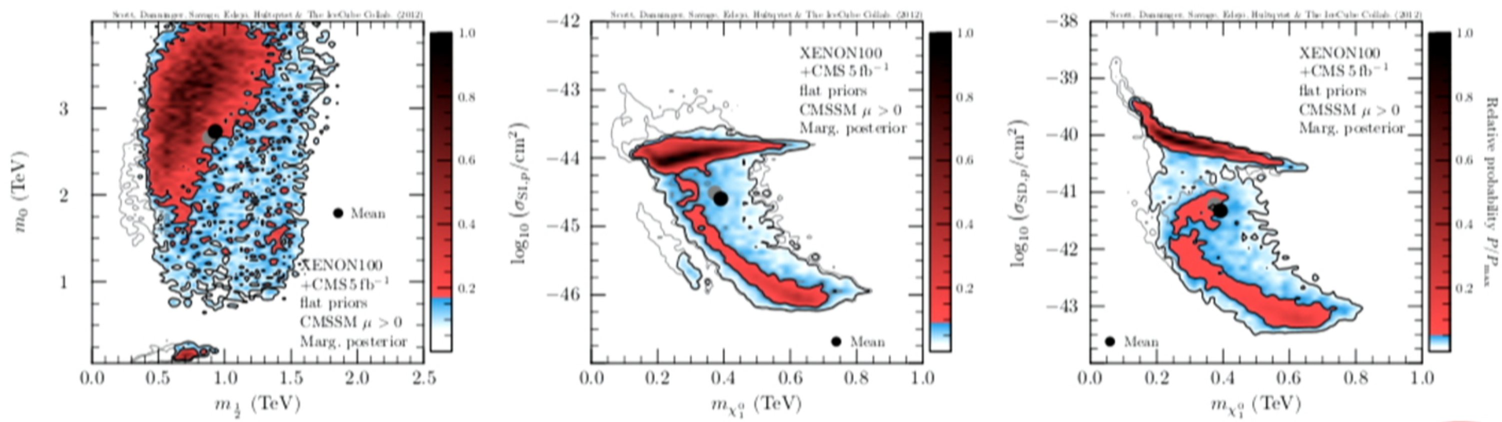

Example of Combined Direct + Indirect + LHC constraints

Base Observables + XENON-100 + CMS 5 fb^{-1} + projected IC86-DeepCore

Grey contours correspond to Base Observables *only*

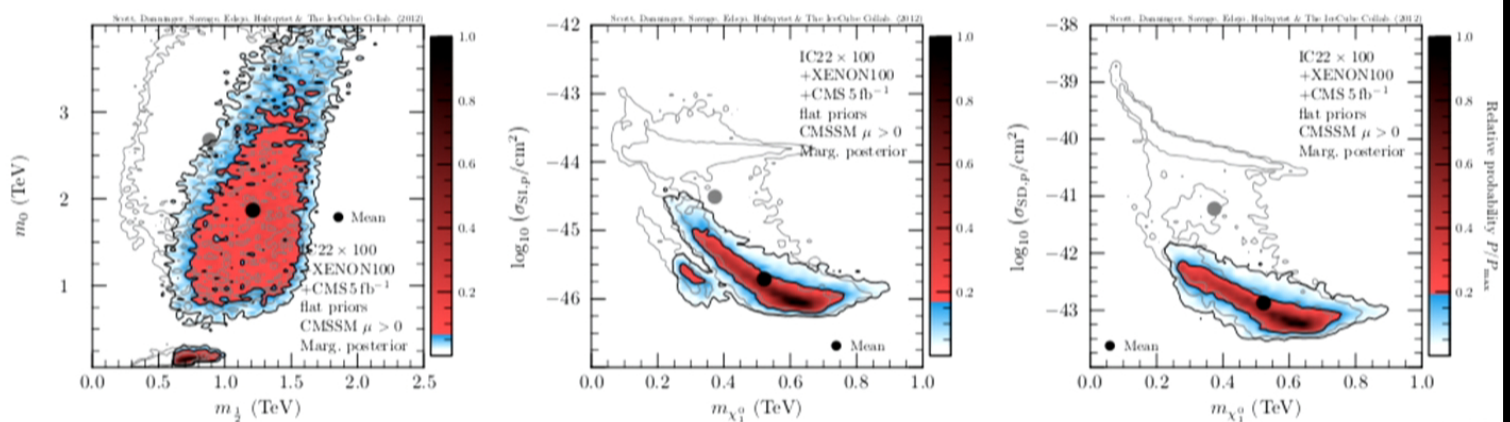

CMSSM, IceCube-22 with 100× boosted effective area
(kinda like IceCube-DeepCore)

Example of Combined Direct + Indirect + LHC constraints

Base Observables + XENON-100 + CMS 5 fb^{-1}

Grey contours correspond to Base Observables *only*

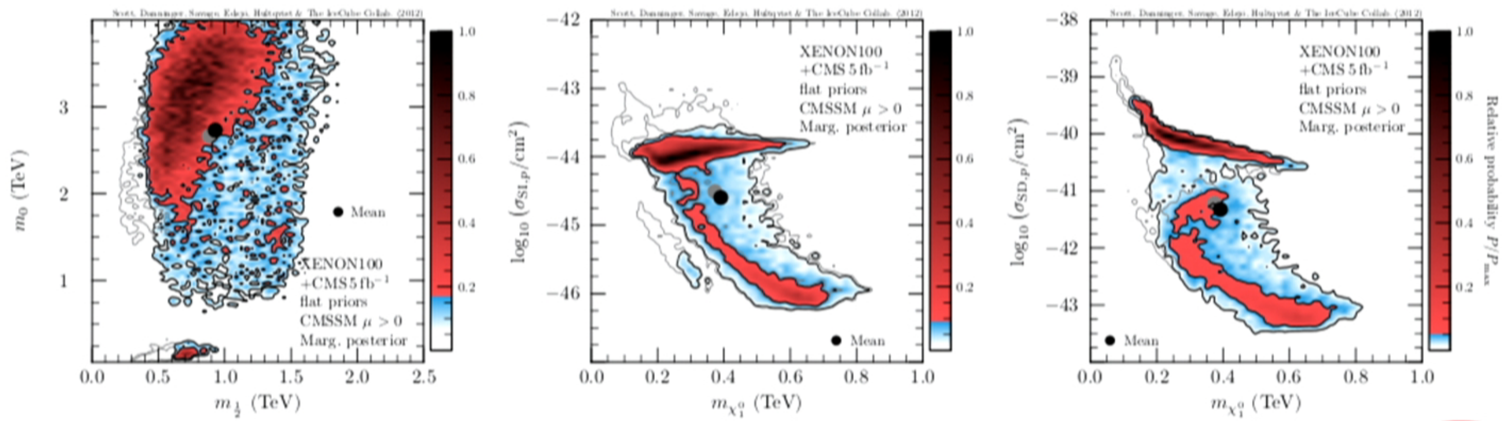

Prospects for detection in the MSSM-25

86-string IceCube vs Direct Detection (points pass $\Omega_\chi h^2$, $b \rightarrow s\gamma$, LEP)

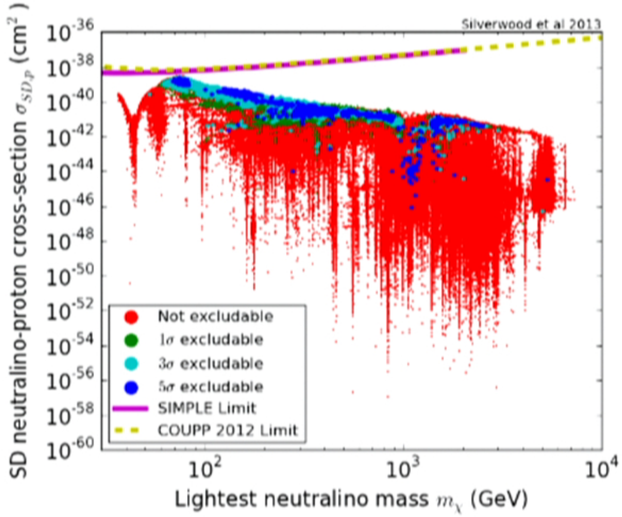

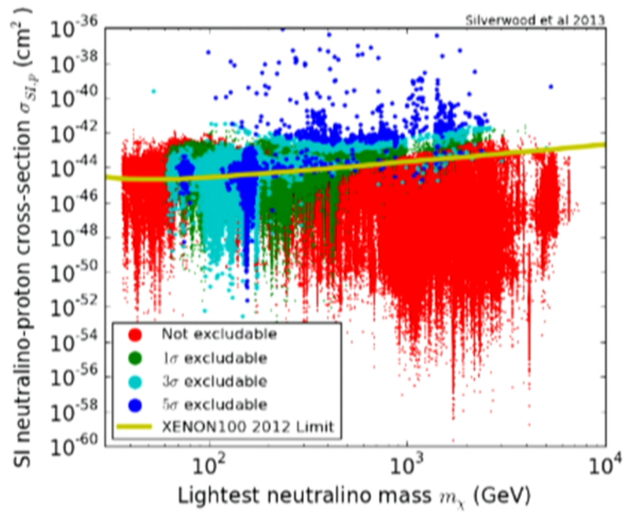

(Silverwood, PS, et al, *JCAP*, 1210.0844)

Many models that IceCube-86 can see are not accessible to direct detection. . .

Prospects for detection in the MSSM-25

86-string IceCube vs LHC (very naively)

SMS limits: 7 TeV, 4.7 fb^{-1}, jets + $E_{T,miss}$; 0 leptons (ATLAS), razor + M_{T2} (CMS)

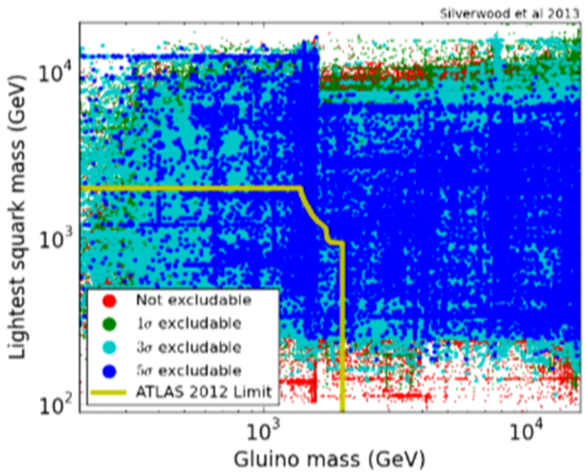

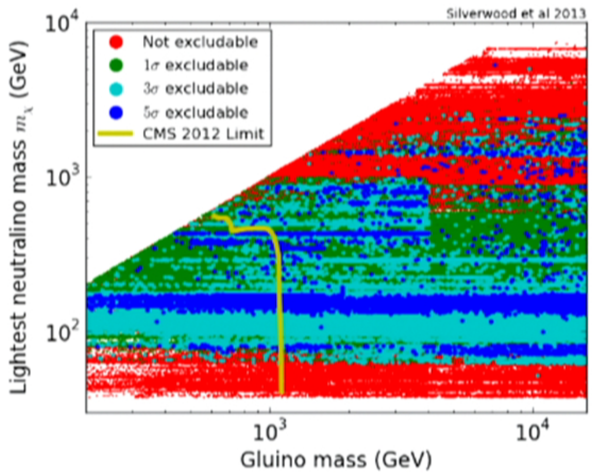

(Silverwood, PS, et al, *JCAP*, 1210.0844)

Many models that IceCube-86 can see are also not accessible at colliders.

Outline

1. The Problem
2. Progress
   - Gamma-rays
   - Neutrinos
   - CMB constraints
3. Future Challenges
   - Respectable LHC likelihoods
   - Statistical/numerical issues
   - Parameter space → Theory space

Generalised DM CMB likelihood functions

Simple CMB likelihood function, for

- Any combination of annihilation or decay channels
- Any dark matter mass
- Any decay lifetime/annihilation cross-section

$\rightarrow$ just requires interpolating one number in a table.

Cline & PS, *JCAP*, 1301.5908, using

- CMB energy deposition from
  - Slatyer (*PRD*, 1211.0283)
  - Finkbeiner et al (*PRD*, 1109.6322)
- PYTHIA annihilation/decay spectra from
  - Cirelli et al (PPPC4DMID; *JCAP*, 1012.4515)

Generalised DM CMB likelihood functions

Simple CMB likelihood function, for

- Any combination of annihilation or decay channels
- Any dark matter mass
- Any decay lifetime/annihilation cross-section

$\rightarrow$ just requires interpolating one number in a table.

f_{eff} for annihilation:

$$\ln \mathcal{L}(\langle\sigma v\rangle|m_\chi, r_i) = -\frac{1}{2} f_{\mathrm{eff}}^2(m_\chi, r_i)\lambda_1 c_1^2 \left(\frac{\langle\sigma v\rangle}{2\times 10^{-27}\mathrm{cm}^3\mathrm{s}^{-1}}\right)^2 \left(\frac{\mathrm{GeV}}{m_\chi}\right)^2 \quad (3)$$

Generalised DM CMB likelihood functions

Simple CMB likelihood function, for

- Any combination of annihilation or decay channels
- Any dark matter mass
- Any decay lifetime/annihilation cross-section

$\rightarrow$ just requires interpolating one number in a table.

$f_{\rm eff}$ for annihilation:

$$\ln \mathcal{L}(\langle\sigma v\rangle|m_\chi, r_i) = -\frac{1}{2} f_{\rm eff}^2(m_\chi, r_i)\lambda_1 c_1^2 \left(\frac{\langle\sigma v\rangle}{2\times 10^{-27}\mathrm{cm}^3\mathrm{s}^{-1}}\right)^2 \left(\frac{\mathrm{GeV}}{m_\chi}\right)^2 \quad (3)$$

η for decay:

$$\ln \mathcal{L}(\tau|m_\chi, r_i) = -\frac{1}{2}\left(\frac{\delta\Omega}{\Omega_{\rm DM}\tau}\right)^2 \eta^2(\tau, m_\chi, r_i) \quad (4)$$

Generalised DM CMB likelihood functions

Simple CMB likelihood function, for

- Any combination of annihilation or decay channels
- Any dark matter mass
- Any decay lifetime/annihilation cross-section

→ just requires interpolating one number in a table.

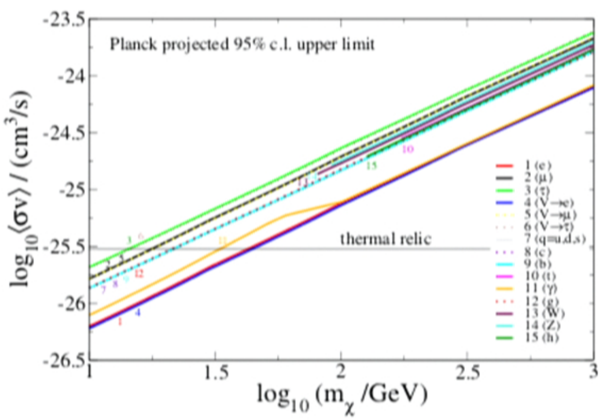

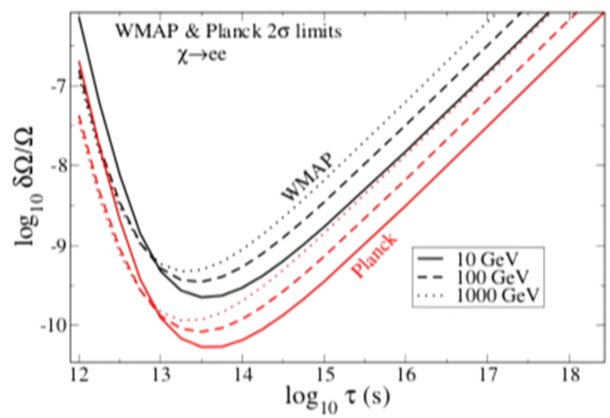

Outline

1. The Problem
2. Progress
   - Gamma-rays
   - Neutrinos
   - CMB constraints
3. Future Challenges
   - Respectable LHC likelihoods
   - Statistical/numerical issues
   - Parameter space → Theory space

The LHC likelihood monster

Time per point:

$\mathcal{O}(minute)$ in best cases

The LHC likelihood monster

Time per point:

$\mathcal{O}(minute)$ in best cases

Time per point for global fits to converge:

$\mathcal{O}(seconds)$ in worst cases

The LHC likelihood monster

Time per point:

$\mathcal{O}(minute)$ in best cases

Time per point for global fits to converge:

$\mathcal{O}(seconds)$ in worst cases

Challenge:

About 2 orders of magnitude too slow to actually include LHC data in global fits properly

Taming the LHC monster

Zeroth Order Response:

"Stuff it, just use the published limits and ignore the dependence on other parameters"

Taming the LHC monster

Zeroth Order Response:

"Stuff it, just use the published limits and ignore the dependence on other parameters"

Obviously naughty – plotted limits assume CMSSM, and fix two of the parameters

- Don't really know dependence on other parameters
- Don't have a likelihood function, just a line
- Can't use this at all for non-CMSSM global fits – e.g. MSSM-25

SuperBayeS

Taming the LHC monster

First Order Response:

"Test if things depend on the other parameters (hope not), re-simulate published exclusion curve"

Taming the LHC monster

First Order Response:

"Test if things depend on the other parameters (hope not), re-simulate published exclusion curve"

Not that great, but OK in some cases

- At least have some sort of likelihood this time
- Still a bit screwed if things do depend a lot on other parameters, but
- allows (potentially shaky) extrapolation, also to non-CMSSM models

Fittino, Mastercode

Taming the LHC monster

Second Order Response:

"That's ridiculous. I've never met a calculation I can't speed up. There must be some way to have my cake and eat it too"

Maybe – this is the challenge.

- Interpolated likelihoods (how to choose nodes?)
- Neural network functional approximation (how to train accurately?)
- Some sort of smart reduction based on event topology?
- Something else?

Balázs, Buckley, Farmer, White et al (1106.4613, 1205.1568)

Outline

1. The Problem
2. Progress
   - Gamma-rays
   - Neutrinos
   - CMB constraints
3. Future Challenges
   - Respectable LHC likelihoods
   - Statistical/numerical issues
   - Parameter space → Theory space

Taming the LHC monster

Second Order Response:

"That's ridiculous. I've never met a calculation I can't speed up. There must be some way to have my cake and eat it too"

Maybe – this is the challenge.

- Interpolated likelihoods (how to choose nodes?)
- Neural network functional approximation (how to train accurately?)
- Some sort of smart reduction based on event topology?
- Something else?

Balázs, Buckley, Farmer, White et al (1106.4613, 1205.1568)

Scanning algorithms

Convergence remains an issue, especially for profile likelihood

Messy likelihood $\Longrightarrow$ best-fit point can be (and often is) easily missed (Akrami, PS et al *JHEP*, 0910.3950, Feroz et al *JHEP*, 1101.3296)

- frequentist CLs are off, as isolikelihood levels are chosen incorrectly
- can impact coverage (overcoverage, or masking of undercoverage due to non-χ^2 *TS* distribution)
- need to use multiple priors and scanning algorithms (one optimised for profile likelihoods?)

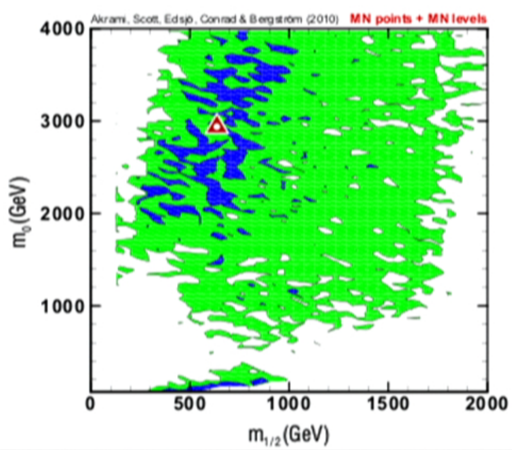

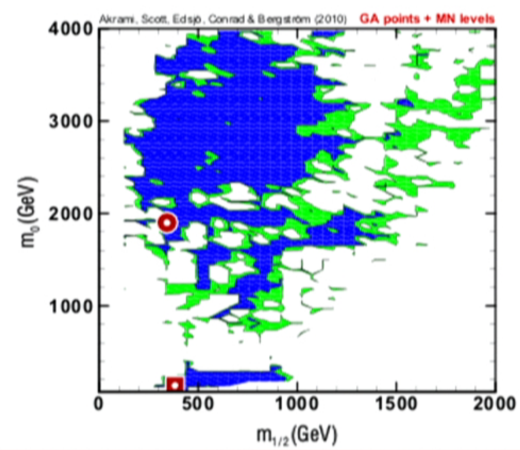

Coverage

[Statistical aside]

Test statistic: a measure on data used to construct statistical tests (e.g. χ^2, ln$\mathcal{L}$, etc.)
Coverage: the percentage of the time that a supposed 'x%' confidence region actually contains the true value

- Distribution of the test statistic and design of the test it's used in determine coverage.
- *p*-value calculation *requires* the test statistic distribution to be well known.

We don't **really** know the distribution of our test statistic in BSM global fits, as it is too expensive to Monte Carlo

- coverage is rarely spot-on unless mapping from parameters to data-space is linear

  (Akrami, Savage, PS et al *JCAP*, 1011.4297, Bridges et al *JHEP*, 1011.4306, Strege et al *PRD*, 1201.3631)
- *p*-value assessments of goodness of fit should be viewed with serious scepticism (→MasterCode)

Outline

1. The Problem
2. Progress
   - Gamma-rays
   - Neutrinos
   - CMB constraints
3. Future Challenges
   - Respectable LHC likelihoods
   - Statistical/numerical issues
   - Parameter space → Theory space

CMSSM, SMS $\neq$ BSM

(SMS = Simplified Model Spectrum)

Want to do model comparison to actually work out which theory is right...

Challenge:

How do I easily adapt a global fit to different BSM theories?

CMSSM, SMS $\neq$ BSM

(SMS = Simplified Model Spectrum)

Want to do model comparison to actually work out which theory is right...

Challenge:

How do I easily adapt a global fit to different BSM theories?

Somehow, we must recast things quickly to a new theory

- data
- likelihood functions
- scanning code 'housekeeping'
- even predictions

$\Longrightarrow$ a new, very abstract global fitting framework

Hitting the wall

Issues with current global fit codes:

- Strongly wedded to a few theories (e.g. constrained MSSM / mSUGRA)
- Strongly wedded to a few theory calculators
- All datasets and observables basically hardcoded
- Rough or non-existent treatment of most experiments (astroparticle + collider especially)
- Sub-optimal statistical methods / search algorithms
- $\Longrightarrow$ *already hitting the wall on theories, data & computational methods*

GAMBIT: a *second-generation* global fit code

GAMBIT: Global And Modular BSM Inference Tool

Overriding principles of GAMBIT: flexibility and modularity

- General enough to allow fast definition of new datasets and theoretical models
- Plug and play scanning, physics and likelihood packages
- Extensive model database – not just small modifications to constrained MSSM (NUHM, etc), and not just SUSY!
- Extensive observable/data libraries (likelihood modules)
- Many statistical options – Bayesian/frequentist, likelihood definitions, scanning algorithms
- A smart and *fast* LHC likelihood calculator
- Massively parallel
- Full open-source code release

The GAMBIT Collaboration

22 Members, 13 Institutes
8 Experiments, 3 major theory codes

Fermi-LAT	J. Conrad, J. Edsjö, G. Martinez, P. Scott†
IceCube	J. Edsjö, C. Savage, P. Scott
ATLAS	A. Buckley, P. Jackson, C. Rogan, A. Saavedra, M. White
LHCb	N. Serra
HESS	J. Conrad
AMS-02	A. Putze
CTA	C. Balázs, T. Bringmann, J. Conrad, M. White
DARWIN	J. Conrad
Theory	C. Balázs, T. Bringmann, J. Cornell, L.-A. Dal, J. Edsjö, B. Farmer, A. Krislock, A. Kvellestad, F.N. Mahmoudi, A. Raklev, C. Savage, P. Scott, C. Weniger, M. White

† PI

Closing remarks

- Robust analysis of dark matter and BSM physics requires multi-messenger global fits
- Lots of interesting astroparticle observables to include in global fits
- Quite a bit of technical (statistical/computational) detail to worry about
- GAMBIT is coming

Closing remarks

- Robust analysis of dark matter and BSM physics requires multi-messenger global fits
- Lots of interesting astroparticle observables to include in global fits
- Quite a bit of technical (statistical/computational) detail to worry about
- GAMBIT is coming