Anomaly aware machine learning for dark matter direct detection at DARWIN

Andre Scaffidi and Roberto Trotta for the DARWIN collaboration.

Overview

- ⇒ New deep learning pipeline to improve upon traditional likelihood approaches.
- Can improve sensitivity over standard approach.
- Methods directly applicable to any detector!

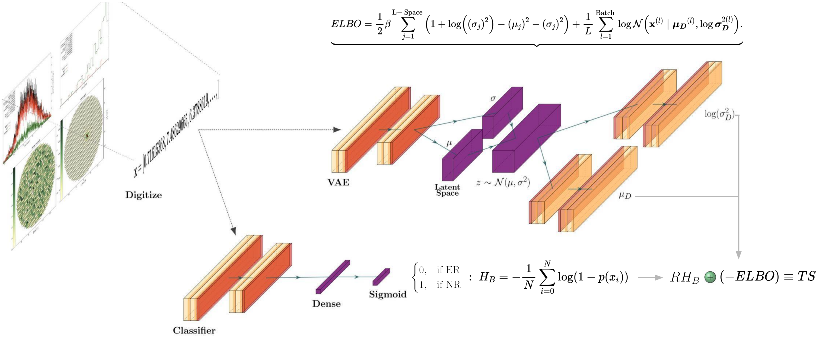

DARWIN collaboration: Proposal

∼ 200 members

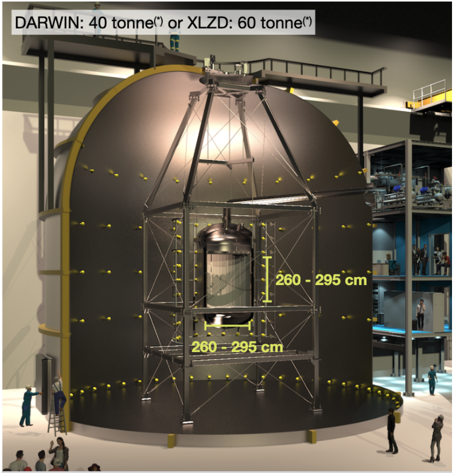

DARWIN collaboration

Direct detection: Traditional likelihood-based analysis

$$\prod_{i=1}^{n} \frac{d\left(N_s + N_b\right)}{dE}\left(E_i \mid \theta\right) \rightarrow \text{ 2D pdf derived from 'templates'}$$

Relies heavily on high-level summary statistics cS1,cS2:
$\Rightarrow \mathbf{E} = \mathbf{g}(\mathbf{cS1}, \mathbf{cS2})$

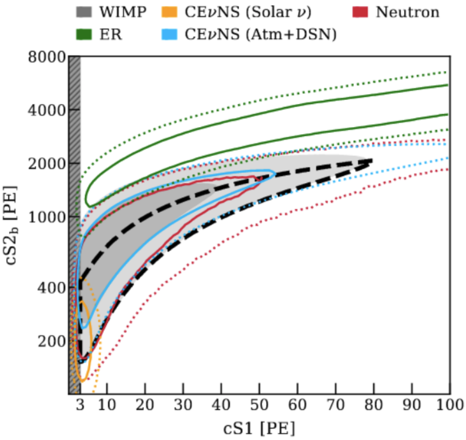

Direct detection: Traditional likelihood-based analysis

$$\prod_{i=1}^{n} \frac{d\left(N_s + N_b\right)}{dE}\left(E_i \mid \theta\right) \rightarrow$$ 2D pdf derived from 'templates'

Relies heavily on high-level summary statistics cS1,cS2:
$\Rightarrow \mathbf{E} = \mathbf{g}(\mathbf{cS1}, \mathbf{cS2})$

↳ Fitted analytically

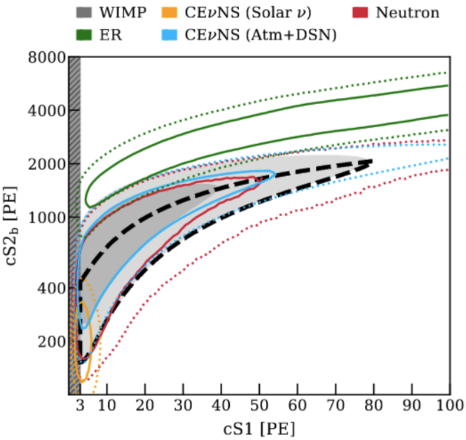

Simulation based inference (SBI)

Benefits of SBI

Bypass need for high level stats. Do inference directly from data.

- Can handle complex models with intractable likelihoods.
- Use deep neural nets to learn underlying features of simulated data/summary stats.
- Once a simulator has been established, possible to include arbitrarily complicated simulations into analysis: prompt readouts $\rightarrow$ high level summary stats.
- Makes no assumptions regarding the analytical form of the likelihood.
- Need no special treatment of nuisance parameters.
  - Can in principle simulate/calibrate any detector effects and learn them directly.

Benefits of SBI

Bypass need for high level stats. Do inference directly from data.

- Can handle complex models with intractable likelihoods.
- Use deep neural nets to learn underlying features of simulated data/summary stats.
- Once a simulator has been established, possible to include arbitrarily complicated simulations into analysis: prompt readouts $\rightarrow$ high level summary stats.
- Makes no assumptions regarding the analytical form of the likelihood.
- Need no special treatment of nuisance parameters.
  - Can in principle simulate/calibrate any detector effects and learn them directly.

Benefits of SBI

Bypass need for high level stats. Do inference directly from data.

- Can handle complex models with intractable likelihoods.
- Use deep neural nets to learn underlying features of simulated data/summary stats.
- Once a simulator has been established, possible to include arbitrarily complicated simulations into analysis: prompt readouts $\rightarrow$ high level summary stats.
- Makes no assumptions regarding the analytical form of the likelihood.
- Need no special treatment of nuisance parameters.
  - Can in principle simulate/calibrate any detector effects and learn them directly.

Benefits of SBI

Bypass need for high level stats. Do inference directly from data.

- Can handle complex models with intractable likelihoods.
- Use deep neural nets to learn underlying features of simulated data/summary stats.
- Once a simulator has been established, possible to include arbitrarily complicated simulations into analysis: prompt readouts $\rightarrow$ high level summary stats.
- Makes no assumptions regarding the analytical form of the likelihood.
- Need no special treatment of nuisance parameters.
  - Can in principle simulate/calibrate any detector effects and learn them directly.

Benefits of SBI

Bypass need for high level stats. Do inference directly from data.

- Can handle complex models with intractable likelihoods.
- Use deep neural nets to learn underlying features of simulated data/summary stats.
- Once a simulator has been established, possible to include arbitrarily complicated simulations into analysis: prompt readouts $\rightarrow$ high level summary stats.
- Makes no assumptions regarding the analytical form of the likelihood.
- Need no special treatment of nuisance parameters.
    - Can in principle simulate/calibrate any detector effects and learn them directly.

Benefits of SBI

Bypass need for high level stats. Do inference directly from data.

- Can handle complex models with intractable likelihoods.
- Use deep neural nets to learn underlying features of simulated data/summary stats.
- Once a simulator has been established, possible to include arbitrarily complicated simulations into analysis: prompt readouts $\rightarrow$ high level summary stats.
- Makes no assumptions regarding the analytical form of the likelihood.
- Need no special treatment of nuisance parameters.
  - Can in principle simulate/calibrate any detector effects and learn them directly.

Simulation-Based Inference with Neural Nets

We have a variety of data/summary stats available to us.

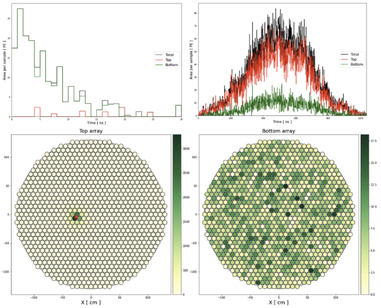

Neural net ‘learns’ underlying likelihood function directly from data.

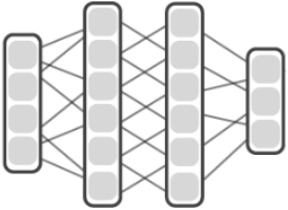

Analysis pipeline 1: Classification of recoil events

Underground TPCs: Two types of events

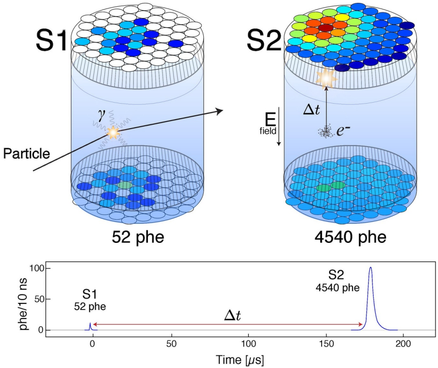

- First primary objective in an analysis is to veto the dominant ER $\Rightarrow$ Binary classification!
- Previous ML studies Sanz et. al, Herrero-Garcia et. al `arXiv:1911.09210, 2110.12248` for XENONnT.

Training data: Simulations

RAW event output S1, S2 PMT deposits (4-fold coincidence, 200 ns):

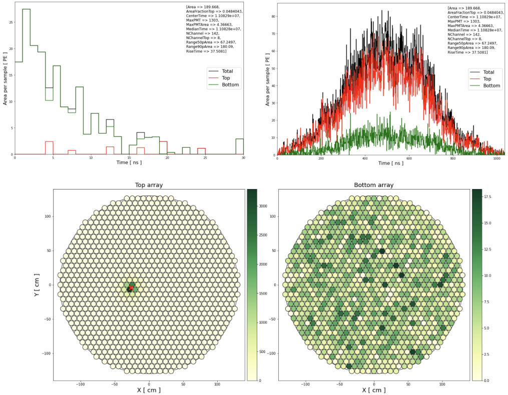

$\Longrightarrow$ $\mathbf{x} = \left[\texttt{S1WaveformTotal}, \texttt{S2WaveformTotal}, \texttt{S2Pattern}\right]$

Two distinct quanta: **Electron Recoil** (ER) and **Nuclear Recoil** (NR)

Training data: Simulations

RAW event output S1, S2 PMT deposits (4-fold coincidence, 200 ns):

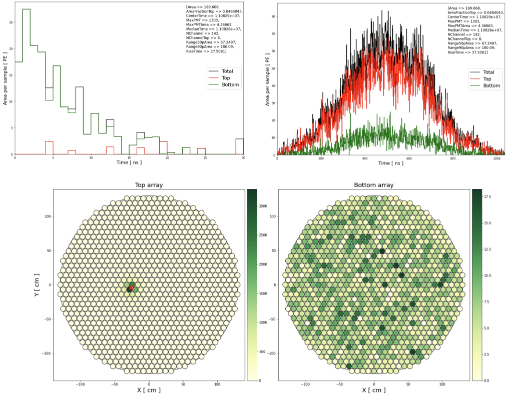

$\Longrightarrow \mathbf{x} = [\texttt{S1WaveformTotal}, \texttt{S2WaveformTotal}, \texttt{S2Pattern}]$

Two distinct quanta: **Electron Recoil** (ER) and **Nuclear Recoil** (NR)

Classification: ER vs. NR Results

- Train on ~ 40000 ER/NR events with $E \in [0, 100]$ keV.
- Check performance $\rightarrow$ confusion matrix:

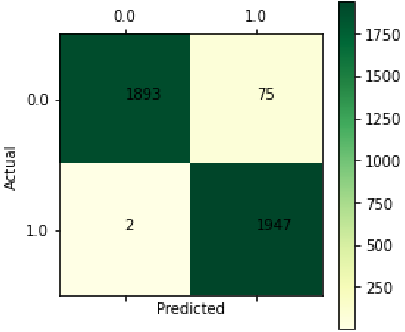

- Takeaway $\Rightarrow$ **98.03% accuracy**. (Recall = 98.07%, Precision = 96.39%)
- **This works regardless WIMP properties: NR/ER are what matter.** (Originally thought not i.e Sanz 1911.0921)

Classification: ER vs. NR Results

- Train on $\sim$ 40000 ER/NR events with $E \in [0, 100]$ keV.
- Check performance $\rightarrow$ confusion matrix:

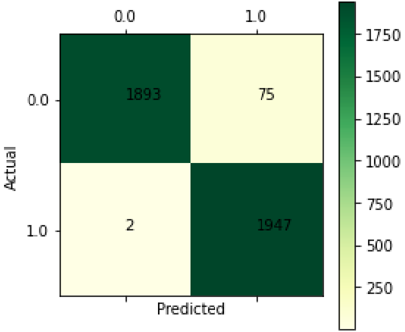

- Takeaway $\Rightarrow$ **98.03% accuracy**. (Recall = 98.07%, Precision = 96.39%)
- **This works regardless WIMP properties: NR/ER are what matter.** (Originally thought not i.e Sanz 1911.0921)

Classification: ER vs. NR Results

- Train on ~ 40000 ER/NR events with $E \in [0, 100]$ keV.
- Check performance $\rightarrow$ confusion matrix:

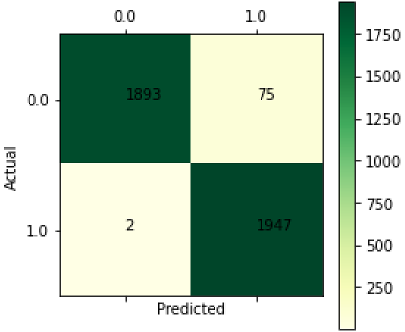

- Takeaway $\Rightarrow$ **98.03% accuracy**. (Recall = 98.07%, Precision = 96.39%)
- **This works regardless WIMP properties: NR/ER are what matter.** (Originally thought not i.e Sanz 1911.0921)

Classification: ER vs. NR Results

- Train on ~ 40000 ER/NR events with $E \in [0, 100]$ keV.
- Check performance $\rightarrow$ confusion matrix:

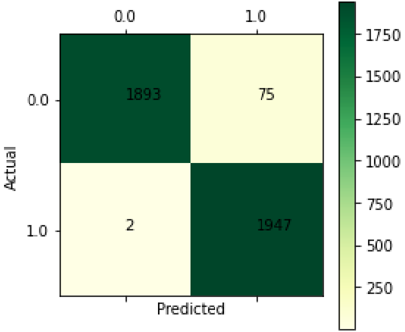

- Takeaway $\Rightarrow$ **98.03% accuracy**. (Recall = 98.07%, Precision = 96.39%)
- **This works regardless WIMP properties: NR/ER are what matter.** (Originally thought not i.e Sanz 1911.0921)

Great! But...

- Trained an effective event-by-event ER veto: In standard DARWIN pipeline, need to sacrifice NR acceptance!
    - Current DARWIN estimate 99.98% ER rejection. 30% NR acceptance.
    - **No need with SBI NN classifier.**
- However, no information regarding the energy of events: WIMPs manifest through number of events + spectral distribution!
- Can we learn the spectral information?

Great! But...

- Trained an effective event-by-event ER veto: In standard DARWIN pipeline, need to sacrifice NR acceptance!
  - Current DARWIN estimate 99.98% ER rejection. 30% NR acceptance.
  - **No need with SBI NN classifier.**
- However, no information regarding the energy of events: WIMPs manifest through number of events + spectral distribution!
- Can we learn the spectral information?

Great! But...

- Trained an effective event-by-event ER veto: In standard DARWIN pipeline, need to sacrifice NR acceptance!
    - Current DARWIN estimate 99.98% ER rejection. 30% NR acceptance.
    - **No need with SBI NN classifier.**
- However, no information regarding the energy of events: WIMPs manifest through number of events + spectral distribution!
- Can we learn the spectral information?

Great! But...

- Trained an effective event-by-event ER veto: In standard DARWIN pipeline, need to sacrifice NR acceptance!
    - Current DARWIN estimate 99.98% ER rejection. 30% NR acceptance.
    - **No need with SBI NN classifier.**
- However, no information regarding the energy of events: WIMPs manifest through number of events + spectral distribution!
- Can we learn the spectral information?

Great! But...

- Trained an effective event-by-event ER veto: In standard DARWIN pipeline, need to sacrifice NR acceptance!
    - Current DARWIN estimate 99.98% ER rejection. 30% NR acceptance.
    - **No need with SBI NN classifier.**
- However, no information regarding the energy of events: WIMPs manifest through number of events + spectral distribution!
- Can we learn the spectral information?

Analysis pipeline 2: Unsupervised approach

Generative Deep Learning: The Variational Auto-Encoder

- Variety of studies in HEP use these for anomaly detection tasks.
- Goal: Learn low dimensional representation (encoding) of data via dimensional reduction.
- Latent space (bottleneck) layer is a bunch of normal distributions parameterized by some $\boldsymbol{\mu}$ and $\boldsymbol{\sigma}$.
- Network should return accurate representations of the input.

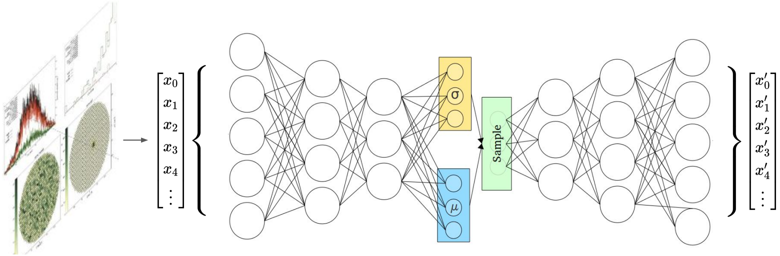

Generative Deep Learning: The Variational Auto-Encoder

- Variety of studies in HEP use these for anomaly detection tasks.
- Goal: Learn low dimensional representation (encoding) of data via dimensional reduction.
- Latent space (bottleneck) layer is a bunch of normal distributions parameterized by some $\boldsymbol{\mu}$ and $\boldsymbol{\sigma}$.
- Network should return accurate representations of the input.

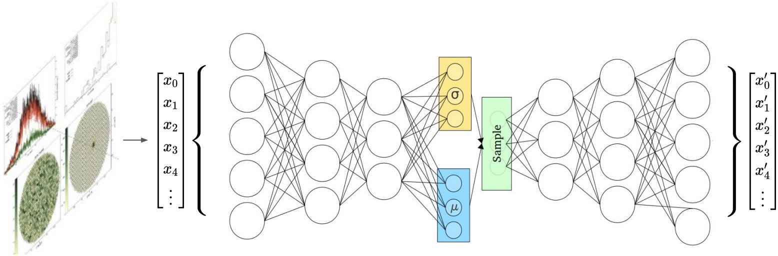

Generative Deep Learning: The Variational Auto-Encoder

- Variety of studies in HEP use these for anomaly detection tasks.
- Goal: Learn low dimensional representation (encoding) of data via dimensional reduction.
- Latent space (bottleneck) layer is a bunch of normal distributions parameterized by some $\boldsymbol{\mu}$ and $\boldsymbol{\sigma}$.
- Network should return accurate representations of the input.

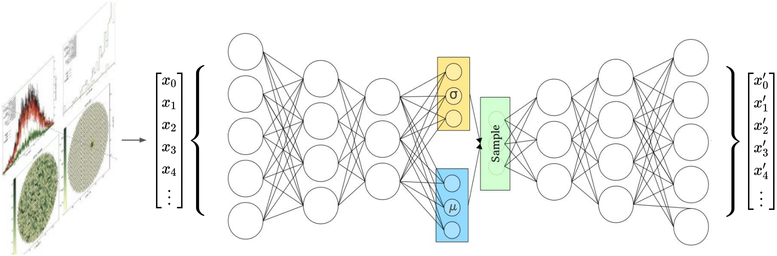

Generative Deep Learning: The Variational Auto-Encoder

- Variety of studies in HEP use these for anomaly detection tasks.
- Goal: Learn low dimensional representation (encoding) of data via dimensional reduction.
- Latent space (bottleneck) layer is a bunch of normal distributions parameterized by some $\boldsymbol{\mu}$ and $\boldsymbol{\sigma}$.
- Network should return accurate representations of the input.

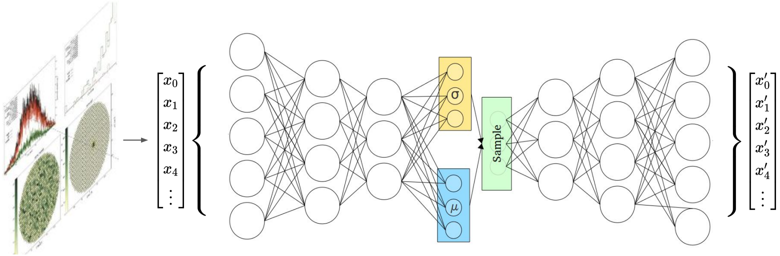

Variational-Auto-Encoder: Training

- Train by maximising evidence lower bound (ELBO):

$$\log p(x) \geq \text{ELBO} = \mathbb{E}_{q(z|x)}\left[\log \frac{p(x,z)}{q(z \mid x)}\right]$$
$$= E[\log p(x|z)] - \beta D_{KL}(q(z|x)||p(z))$$
$$x = \text{Input}$$
$$z = \text{Latent vector}$$
$$\beta = \text{Regularization parameter}$$

- **Loss** $\equiv -\text{ELBO}$

Variational-Auto-Encoder: Training

- Train by maximising evidence lower bound (ELBO):

$$\log p(x) \geq \text{ELBO} = \mathbb{E}_{q(z|x)}\left[\log \frac{p(x,z)}{q(z \mid x)}\right]$$
$$= E[\log p(x|z)] - \beta D_{KL}(q(z|x)||p(z))$$
$$x = \text{Input}$$
$$z = \text{Latent vector}$$
$$\beta = \text{Regularization parameter}$$

- **Loss** $\equiv -\text{ELBO}$

VAE: Training

- Train the network for 200 epochs.

$$Loss \equiv -\text{ELBO}$$

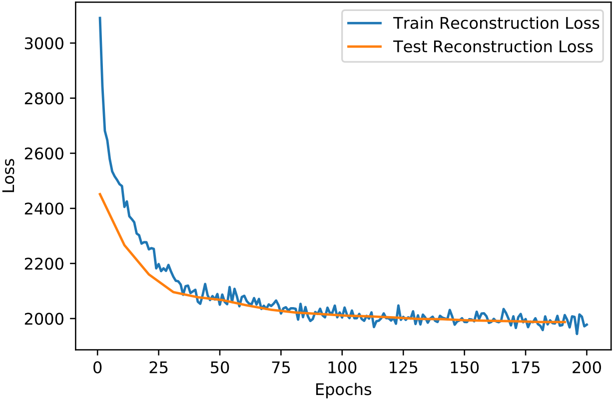

VAE trained on ER realisations: Spectral information

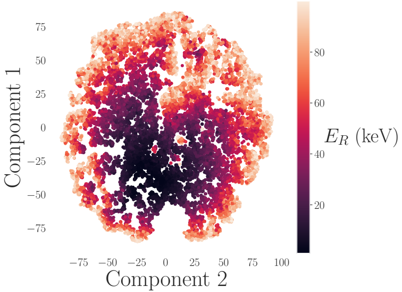

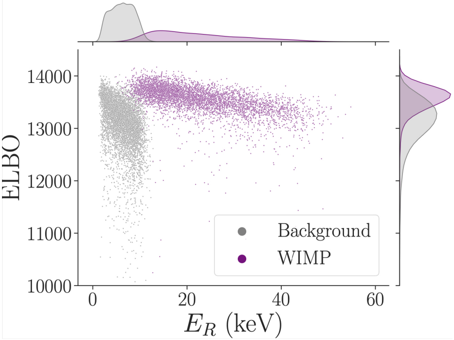

- Auto-encoder can learn underlying spectral information of events $\Rightarrow$ Sensitivity to WIMP mass.
- Can we also just fully reconstruct the energy of an event straight from the data? Yes! See later.

Anomaly detection (Looking for non background-like events)

Anomaly detection

- Accept/Reject

$$\mathcal{H}_0 : TS \sim \mathcal{P}\left(x \mid \text{No signal}\right)$$

  $\mathcal{P}$: Conditional process generating some statistic TS, under the assumption that no WIMP signal is present.

- $\mathcal{P}$ intractable: Use neural networks to derive optimal TS.

Anomaly detection

- Accept/Reject

$$\mathcal{H}_0 : TS \sim \mathcal{P}\left(x \mid \text{No signal}\right)$$

  $\mathcal{P}$: Conditional process generating some statistic TS, under the assumption that no WIMP signal is present.

- $\mathcal{P}$ intractable: Use neural networks to derive optimal TS.

Anomaly detection

What is a suitable TS?

- If VAE has learned the underlying properties of ER events, any **non-ER** events will in general have **higher loss** (smaller ELBO).
- Furthermore, non trivial spectral info learned $\Rightarrow$ ELBO **distribution** shape can further inform about NR background.
- Loss distribution of anomalous data (new physics) will show as an **excess** over background only loss distribution.
- Try distribution of $= ELBO$

Anomaly detection

What is a suitable TS?

- If VAE has learned the underlying properties of ER events, any **non-ER** events will in general have **higher loss** (smaller ELBO).
- Furthermore, non trivial spectral info learned $\Rightarrow$ ELBO **distribution** shape can further inform about NR background.
- Loss distribution of anomalous data (new physics) will show as an **excess** over background only loss distribution.
- Try distribution of $= ELBO$

Anomaly detection

What is a suitable TS?

- If VAE has learned the underlying properties of ER events, any **non-ER** events will in general have **higher loss** (smaller ELBO).
- Furthermore, non trivial spectral info learned $\Rightarrow$ ELBO **distribution** shape can further inform about NR background.
- Loss distribution of anomalous data (new physics) will show as an **excess** over background only loss distribution.
- Try distribution of $= ELBO$

Anomaly detection

What is a suitable TS?

- If VAE has learned the underlying properties of ER events, any **non-ER** events will in general have **higher loss** (smaller ELBO).
- Furthermore, non trivial spectral info learned $\Rightarrow$ ELBO **distribution** shape can further inform about NR background.
- Loss distribution of anomalous data (new physics) will show as an **excess** over background only loss distribution.
- Try distribution of $= ELBO$

Anomaly detection

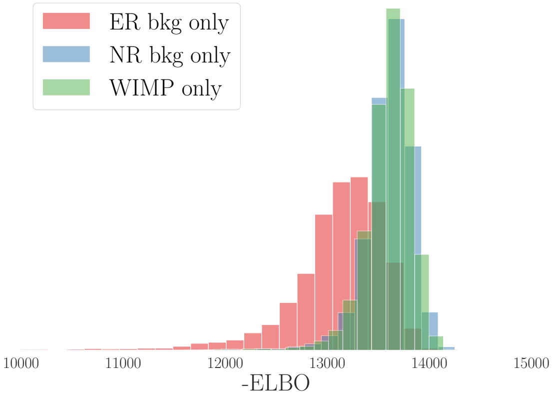

- Background loss distributions + WIMP loss distribution.
- Any* anomalous signal will show up as statistical deviation in (pseudo)data loss vs. (known) background loss.

Semi-unsupervised anomaly detection: New distance metric

Cool. But...

- A bit rubbish: Can we get greater separation (anomaly awareness) between these distributions?
- New 'anomaly score' that utilizes pre-trained supervised NN classifier:

$$TS = -ELBO + R\,H_B\ ,$$

where

- $H_B = -\frac{1}{N}\sum_{i=0}^{N}\log\left(1 - p\left(x_i\right)\right)$ (Binary cross-sentropy.)
- R scales the contribution of the cross-entropy term $\rightarrow$ makes it more/less supervised.

Semi-unsupervised anomaly detection: New distance metric

Cool. But...

- A bit rubbish: Can we get greater separation (anomaly awareness) between these distributions?
- New 'anomaly score' that utilizes pre-trained supervised NN classifier:

$$TS = -ELBO + R\,H_B\ ,$$

where

- $H_B = -\frac{1}{N}\sum_{i=0}^{N}\log\left(1-p\left(x_i\right)\right)$ (Binary cross-sentropy.)
- R scales the contribution of the cross-entropy term $\rightarrow$ makes it more/less supervised.

Semi-unsupervised anomaly detection: New distance metric

Cool. But...

- A bit rubbish: Can we get greater separation (anomaly awareness) between these distributions?
- New 'anomaly score' that utilizes pre-trained supervised NN classifier:

$$TS = -ELBO + R\,H_B\ ,$$

where

- $H_B = -\frac{1}{N}\sum_{i=0}^{N}\log\left(1 - p\left(x_i\right)\right)$ (Binary cross-sentropy.)
- R scales the contribution of the cross-entropy term $\rightarrow$ makes it more/less supervised.

Semi-unsupervised anomaly detection: New distance metric

Cool. But...

- A bit rubbish: Can we get greater separation (anomaly awareness) between these distributions?
- New 'anomaly score' that utilizes pre-trained supervised NN classifier:

$$TS = -ELBO + R\, H_B \ ,$$

where

- $H_B = -\frac{1}{N}\sum_{i=0}^{N} \log\left(1 - p\left(x_i\right)\right)$ (Binary cross-sentropy.)
- R scales the contribution of the cross-entropy term $\rightarrow$ makes it more/less supervised.

Semi-unsupervised anomaly detection: Full pipeline

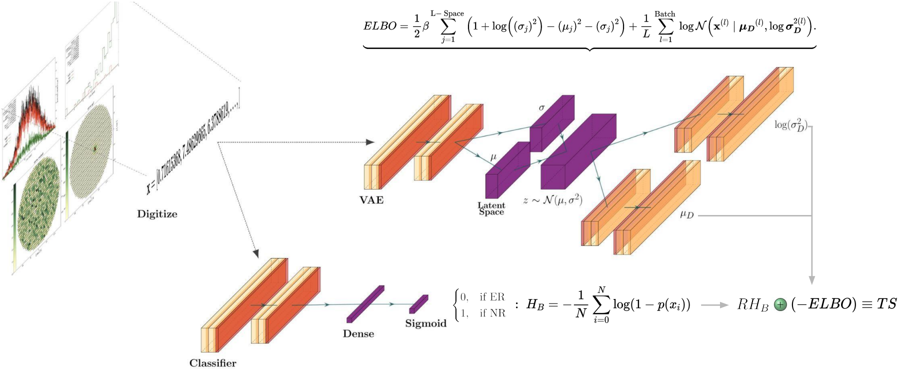

Semi-unsupervised anomaly detection: New distance metric

$$TS = (-ELBO) + R\,H_B\ ,$$

⇒ Semi-unsupervised. Much greater anomaly awareness!

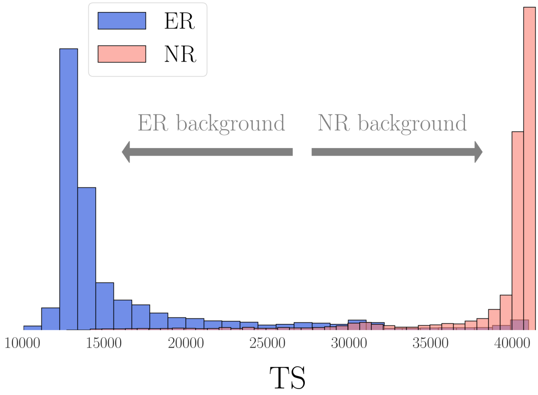

Pseudo-data sets

Re-weight anomaly score distributions TS according to expected ER+NR backgrounds and inject some WIMP signal: ER [2-10] keVee, NR [5-35] keVnr

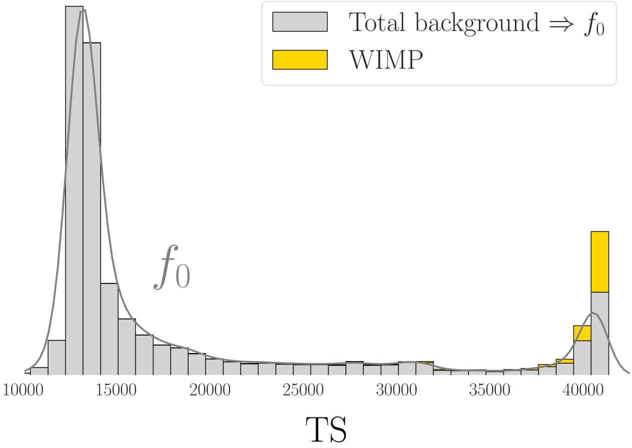

Pseudo-data sets

Re-weight anomaly score distributions TS according to expected ER+NR backgrounds and inject some WIMP signal: ER [2-10] keVee, NR [5-35] keVnr

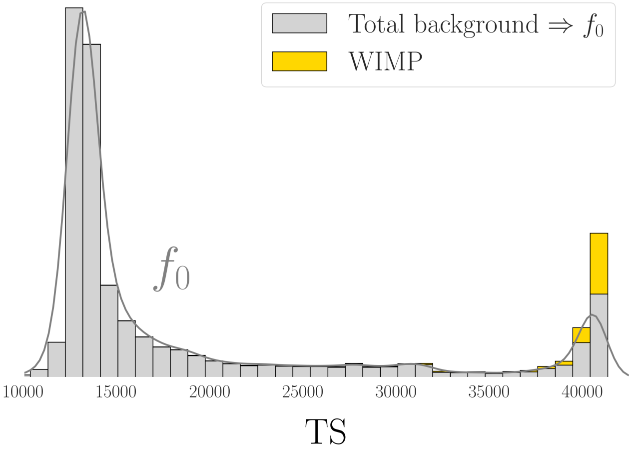

Dimensionally reduced two-sample test

- $\Rightarrow$ 1D analysis in TS space: Accept/reject $\mathcal{H}_0$.

$$\mathcal{P}\left(x \mid \text{No signal}\right) \equiv \mathcal{L}(\mathbf{TS}|\mathcal{H}_0) \propto e^{-B} \prod_{i=1}^{N} \left(B f_0 \left(TS_i\right)\right)$$

- Unbinned.
- Parametrically independent on WIMP model.
- No auxiliary terms/nuisance parameters required *assuming* simulations have suitably descriptive coverage.
- Capability to conduct ER only searches with same machinery.
- In principal can propagate uncertainties on the bkg from simulation (or even better, calibration).

Dimensionally reduced two-sample test

- $\Rightarrow$ 1D analysis in TS space: Accept/reject $\mathcal{H}_0$.

$$\mathcal{P}\left(x \mid \text{No signal}\right) \equiv \mathcal{L}(\mathbf{TS}|\mathcal{H}_0) \propto e^{-B} \prod_{i=1}^{N} \left(B f_0\left(TS_i\right)\right)$$

- Unbinned.
- Parametrically independent on WIMP model.
- No auxiliary terms/nuisance parameters required *assuming* simulations have suitably descriptive coverage.
- Capability to conduct ER only searches with same machinery.
- In principal can propagate uncertainties on the bkg from simulation (or even better, calibration).

Dimensionally reduced two-sample test

- $\Rightarrow$ 1D analysis in TS space: Accept/reject $\mathcal{H}_0$.

$$\mathcal{P}\left(x \mid \text{No signal}\right) \equiv \mathcal{L}(\mathbf{TS}|\mathcal{H}_0) \propto e^{-B} \prod_{i=1}^{N} \left(B f_0\left(TS_i\right)\right)$$

- Unbinned.
- Parametrically independent on WIMP model.
- No auxiliary terms/nuisance parameters required *assuming* simulations have suitably descriptive coverage.
- Capability to conduct ER only searches with same machinery.
- In principal can propagate uncertainties on the bkg from simulation (or even better, calibration).

Dimensionally reduced two-sample test

- $\Rightarrow$ 1D analysis in TS space: Accept/reject $\mathcal{H}_0$.

$$\mathcal{P}\left(x \mid \text{No signal}\right) \equiv \mathcal{L}(\mathbf{TS}|\mathcal{H}_0) \propto e^{-B} \prod_{i=1}^{N} \left(B f_0 \left(TS_i\right)\right)$$

- Unbinned.
- Parametrically independent on WIMP model.
- No auxiliary terms/nuisance parameters required *assuming* simulations have suitably descriptive coverage.
- Capability to conduct ER only searches with same machinery.
- In principal can propagate uncertainties on the bkg from simulation (or even better, calibration).

Dimensionally reduced two-sample test

- $\Rightarrow$ 1D analysis in TS space: Accept/reject $\mathcal{H}_0$.

$$\mathcal{P}\left(x \mid \text{No signal}\right) \equiv \mathcal{L}(\mathbf{TS}|\mathcal{H}_0) \propto e^{-B} \prod_{i=1}^{N} \left(B f_0\left(TS_i\right)\right)$$

- Unbinned.
- Parametrically independent on WIMP model.
- No auxiliary terms/nuisance parameters required *assuming* simulations have suitably descriptive coverage.
- Capability to conduct ER only searches with same machinery.
- In principal can propagate uncertainties on the bkg from simulation (or even better, calibration).

Dimensionally reduced two-sample test

- $\Rightarrow$ 1D analysis in TS space: Accept/reject $\mathcal{H}_0$.

$$\mathcal{P}\left(x \mid \text{No signal}\right) \equiv \mathcal{L}(\mathbf{TS}|\mathcal{H}_0) \propto e^{-B} \prod_{i=1}^{N} \left(B f_0 \left(TS_i\right)\right)$$

- Unbinned.
- Parametrically independent on WIMP model.
- No auxiliary terms/nuisance parameters required *assuming* simulations have suitably descriptive coverage.
- Capability to conduct ER only searches with same machinery.
- In principal can propagate uncertainties on the bkg from simulation (or even better, calibration).

Median Sensitivity (50 GeV, $\sigma_{SI} = 4.9 \times 10^{-49}$ cm^2)

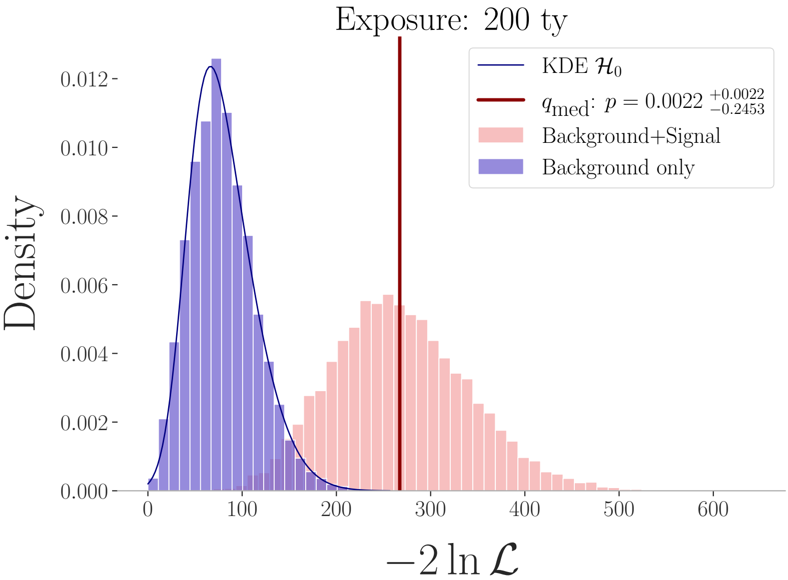

As a function of exposure

- Neural net
- Binned likelihood based: Median sensitivity [30% NR acceptance, 99% ER rejection]

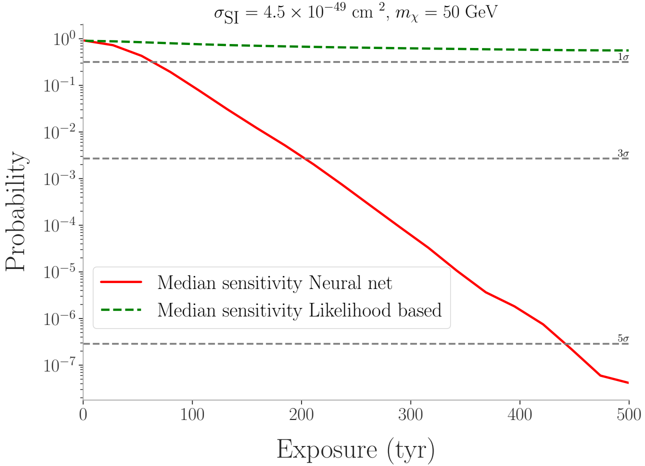

Full sensitivity (Preliminary)

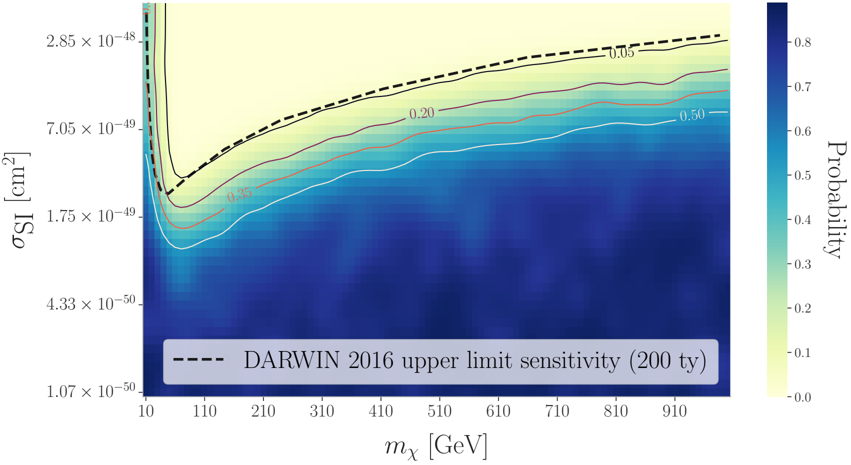

Caution: 90% C.L upper limit is model dependent → 'weaker' test.

Energy reconstruction

Energy reconstruction SBI with masked autoregressive flows

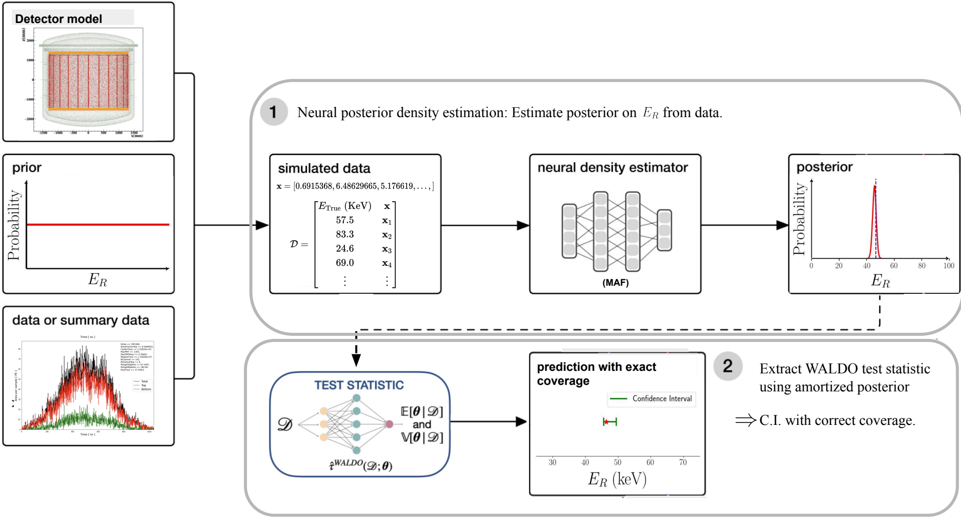

$$\tau^{\text{WALDO}}(\mathcal{D};\boldsymbol{\theta}_0) = \left(\mathbb{E}[\boldsymbol{\theta} \mid \mathcal{D}] - \boldsymbol{\theta}_0\right)^T \mathbb{V}[\boldsymbol{\theta} \mid \mathcal{D}]^{-1} \left(\mathbb{E}[\boldsymbol{\theta} \mid \mathcal{D}] - \boldsymbol{\theta}_0\right)$$

Follow up work: E reconstruction

Neural posterior density estimation (Masked auto-regressive flows)

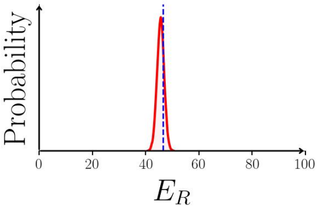

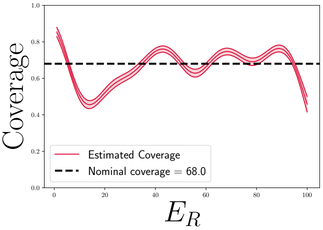

Neural posterior density estimation + WALDO

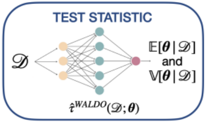

arXiv:2205.15680

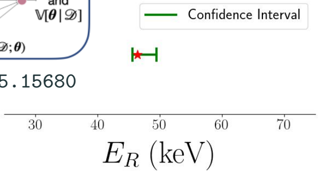

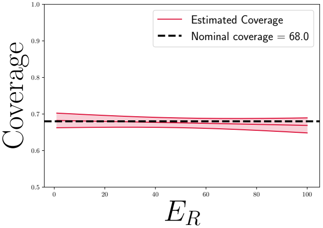

On directional detection

On directional detection

- SBI methods directly applicable.
- Depending on data:
    - 3D Structured: MLP, Unstructured Point clouds → PointNet, GNN's

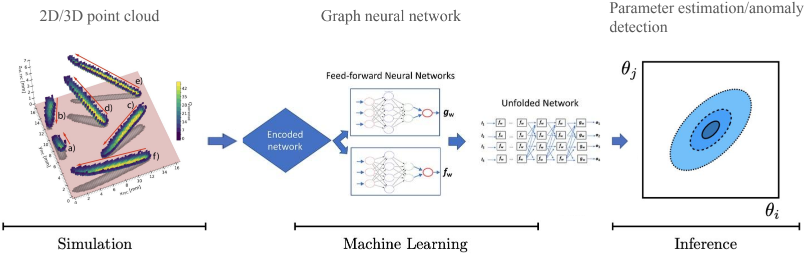

Thank you!

Backup Slides

Simulation-Based Inference (in a nutshell)

Simulation-based inference is a statistical technique that allows us to make inferences about a population or process based on simulated/calibrated data. It involves the following steps:

1. Generate simulated data.
2. Use deep neural nets to learn underlying features of simulated data.
3. Use trained models to inference.

Simulation-Based Inference (in a nutshell)

Simulation-based inference is a statistical technique that allows us to make inferences about a population or process based on simulated/calibrated data. It involves the following steps:

1. Generate simulated data.
2. Use deep neural nets to learn underlying features of simulated data.
3. Use trained models to inference.

Simulation-Based Inference (in a nutshell)

Simulation-based inference is a statistical technique that allows us to make inferences about a population or process based on simulated/calibrated data. It involves the following steps:

1. Generate simulated data.
2. Use deep neural nets to learn underlying features of simulated data.
3. Use trained models to inference.

Simulation-Based Inference (in a nutshell)

Simulation-based inference is a statistical technique that allows us to make inferences about a population or process based on simulated/calibrated data. It involves the following steps:

1. Generate simulated data.
2. Use deep neural nets to learn underlying features of simulated data.
3. Use trained models to inference.

DARWIN: Simulation pipeline

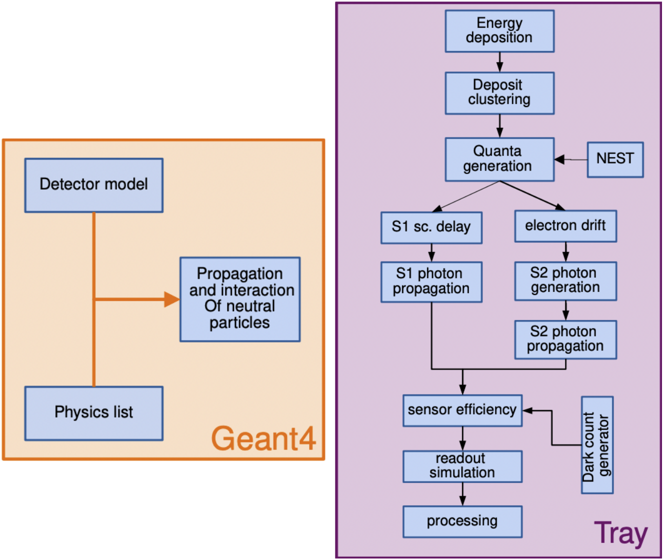

Backgrounds

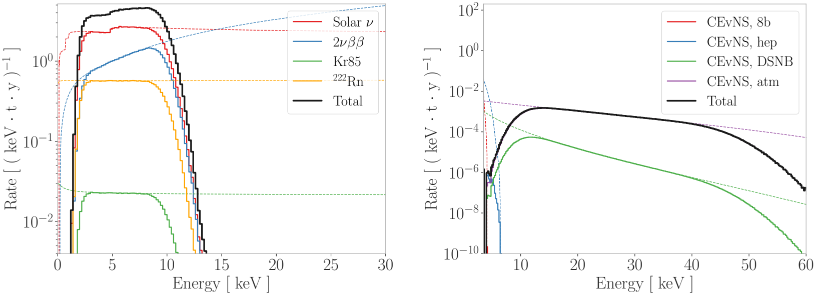

- Intrinsic and extrinsic.
- Coherent neutrino scattering provides dominant background for WIMP searches.

Binned likelihood based approach

$$\mathcal{L}(\mathbf{x}) = \prod_{i=1}^{bins} \frac{\lambda_i^{n_i}}{n_i!} e^{-\lambda_i} \quad : \text{ ER veto (99.98\%), fidiucilization etc.}$$

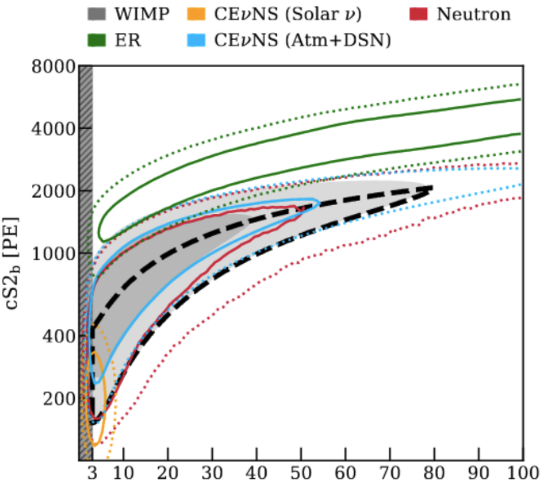

Effect of R

- Explore effect of the R parameter.
- Three mock data sets corresponding to 10, 500 and 1000 GeV at fixed $\sigma = 10^{-45}\text{cm}^2$, 5 $t{\cdot}$yr exposure.
- Best result for $R \sim 170$, but generally free to choose!

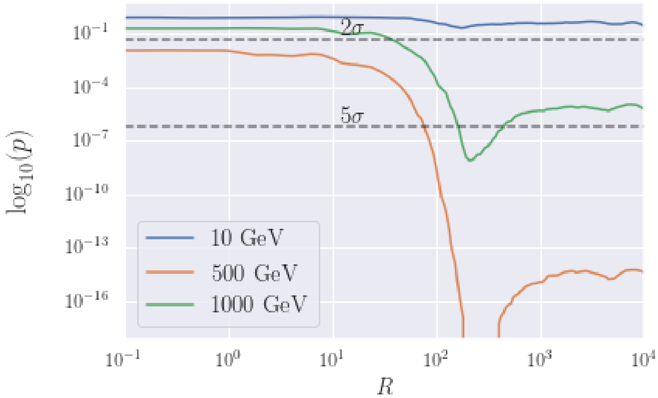

Effect of R

- Explore effect of the R parameter.
- Three mock data sets corresponding to 10, 500 and 1000 GeV at fixed $\sigma = 10^{-45}\text{cm}^2$, 5 $t{\cdot}$yr exposure.
- Best result for $R \sim 170$, but generally free to choose!

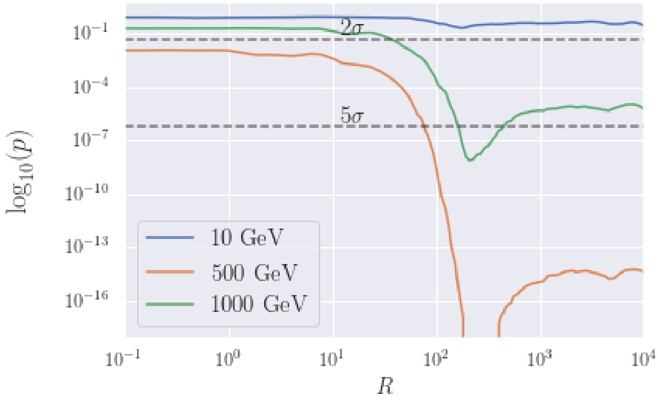

Effect of R

- Explore effect of the R parameter.
- Three mock data sets corresponding to 10, 500 and 1000 GeV at fixed $\sigma = 10^{-45}\text{cm}^2$, 5 $t{\cdot}$yr exposure.
- Best result for $R \sim 170$, but generally free to choose!

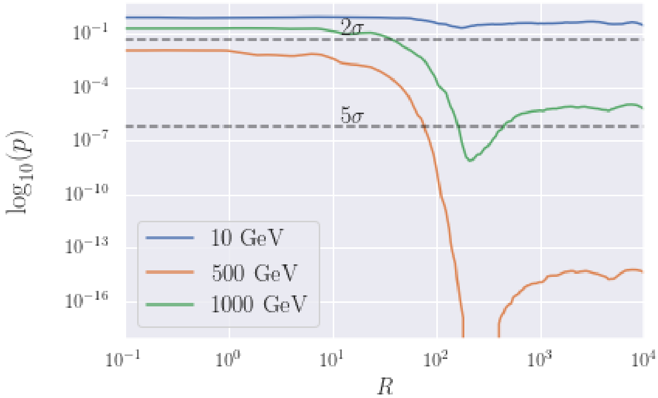

Forecasting sensitivity

Median sensitivity

- Probability to accept/reject $\mathcal{H}_0$ after some exposure.
- Model independent.
- Simulate $\sim 10^4$ realisations of $-2\ln\mathcal{L}(\mathbf{TS}\,|\,\mathcal{H}_0)$ to ascertain the asymptotic form of $\mathcal{H}_0$.

$$p = \int_{q_{\text{med}}}^{\infty} dq\, \mathcal{H}_0(q)\ .$$

Median sensitivity

- Probability to accept/reject $\mathcal{H}_0$ after some exposure.
- Model independent.
- Simulate $\sim 10^4$ realisations of $-2\ln\mathcal{L}(\mathbf{TS}\,|\,\mathcal{H}_0)$ to ascertain the asymptotic form of $\mathcal{H}_0$.

$$p = \int_{q_{\mathrm{med}}}^{\infty} dq\; \mathcal{H}_0(q)\ .$$

Median sensitivity

- Probability to accept/reject $\mathcal{H}_0$ after some exposure.
- Model independent.
- Simulate $\sim 10^4$ realisations of $-2\ln\mathcal{L}(\mathbf{TS}\,|\,\mathcal{H}_0)$ to ascertain the asymptotic form of $\mathcal{H}_0$.

$$p = \int_{q_{\text{med}}}^{\infty} dq\; \mathcal{H}_0(q)\;.$$