Intro to Unsupervised ML and Anomaly Detection

COFI Winter School 2023
Old San Juan, Puerto Rico

David Shih
December 11, 2023

RUTGERS
THE STATE UNIVERSITY OF NEW JERSEY

Outline

1. Motivation: new physics and the LHC
2. Model-agnostic new physics searches with modern ML
    - Outlier detection
    - Overdensities
3. Common tools, cross-cutting domains: Anomaly Detection from LHC to Astro

By now, there are countless searches for new physics at the LHC.
Almost all cut-based, model-specific.
No sign of new physics yet…

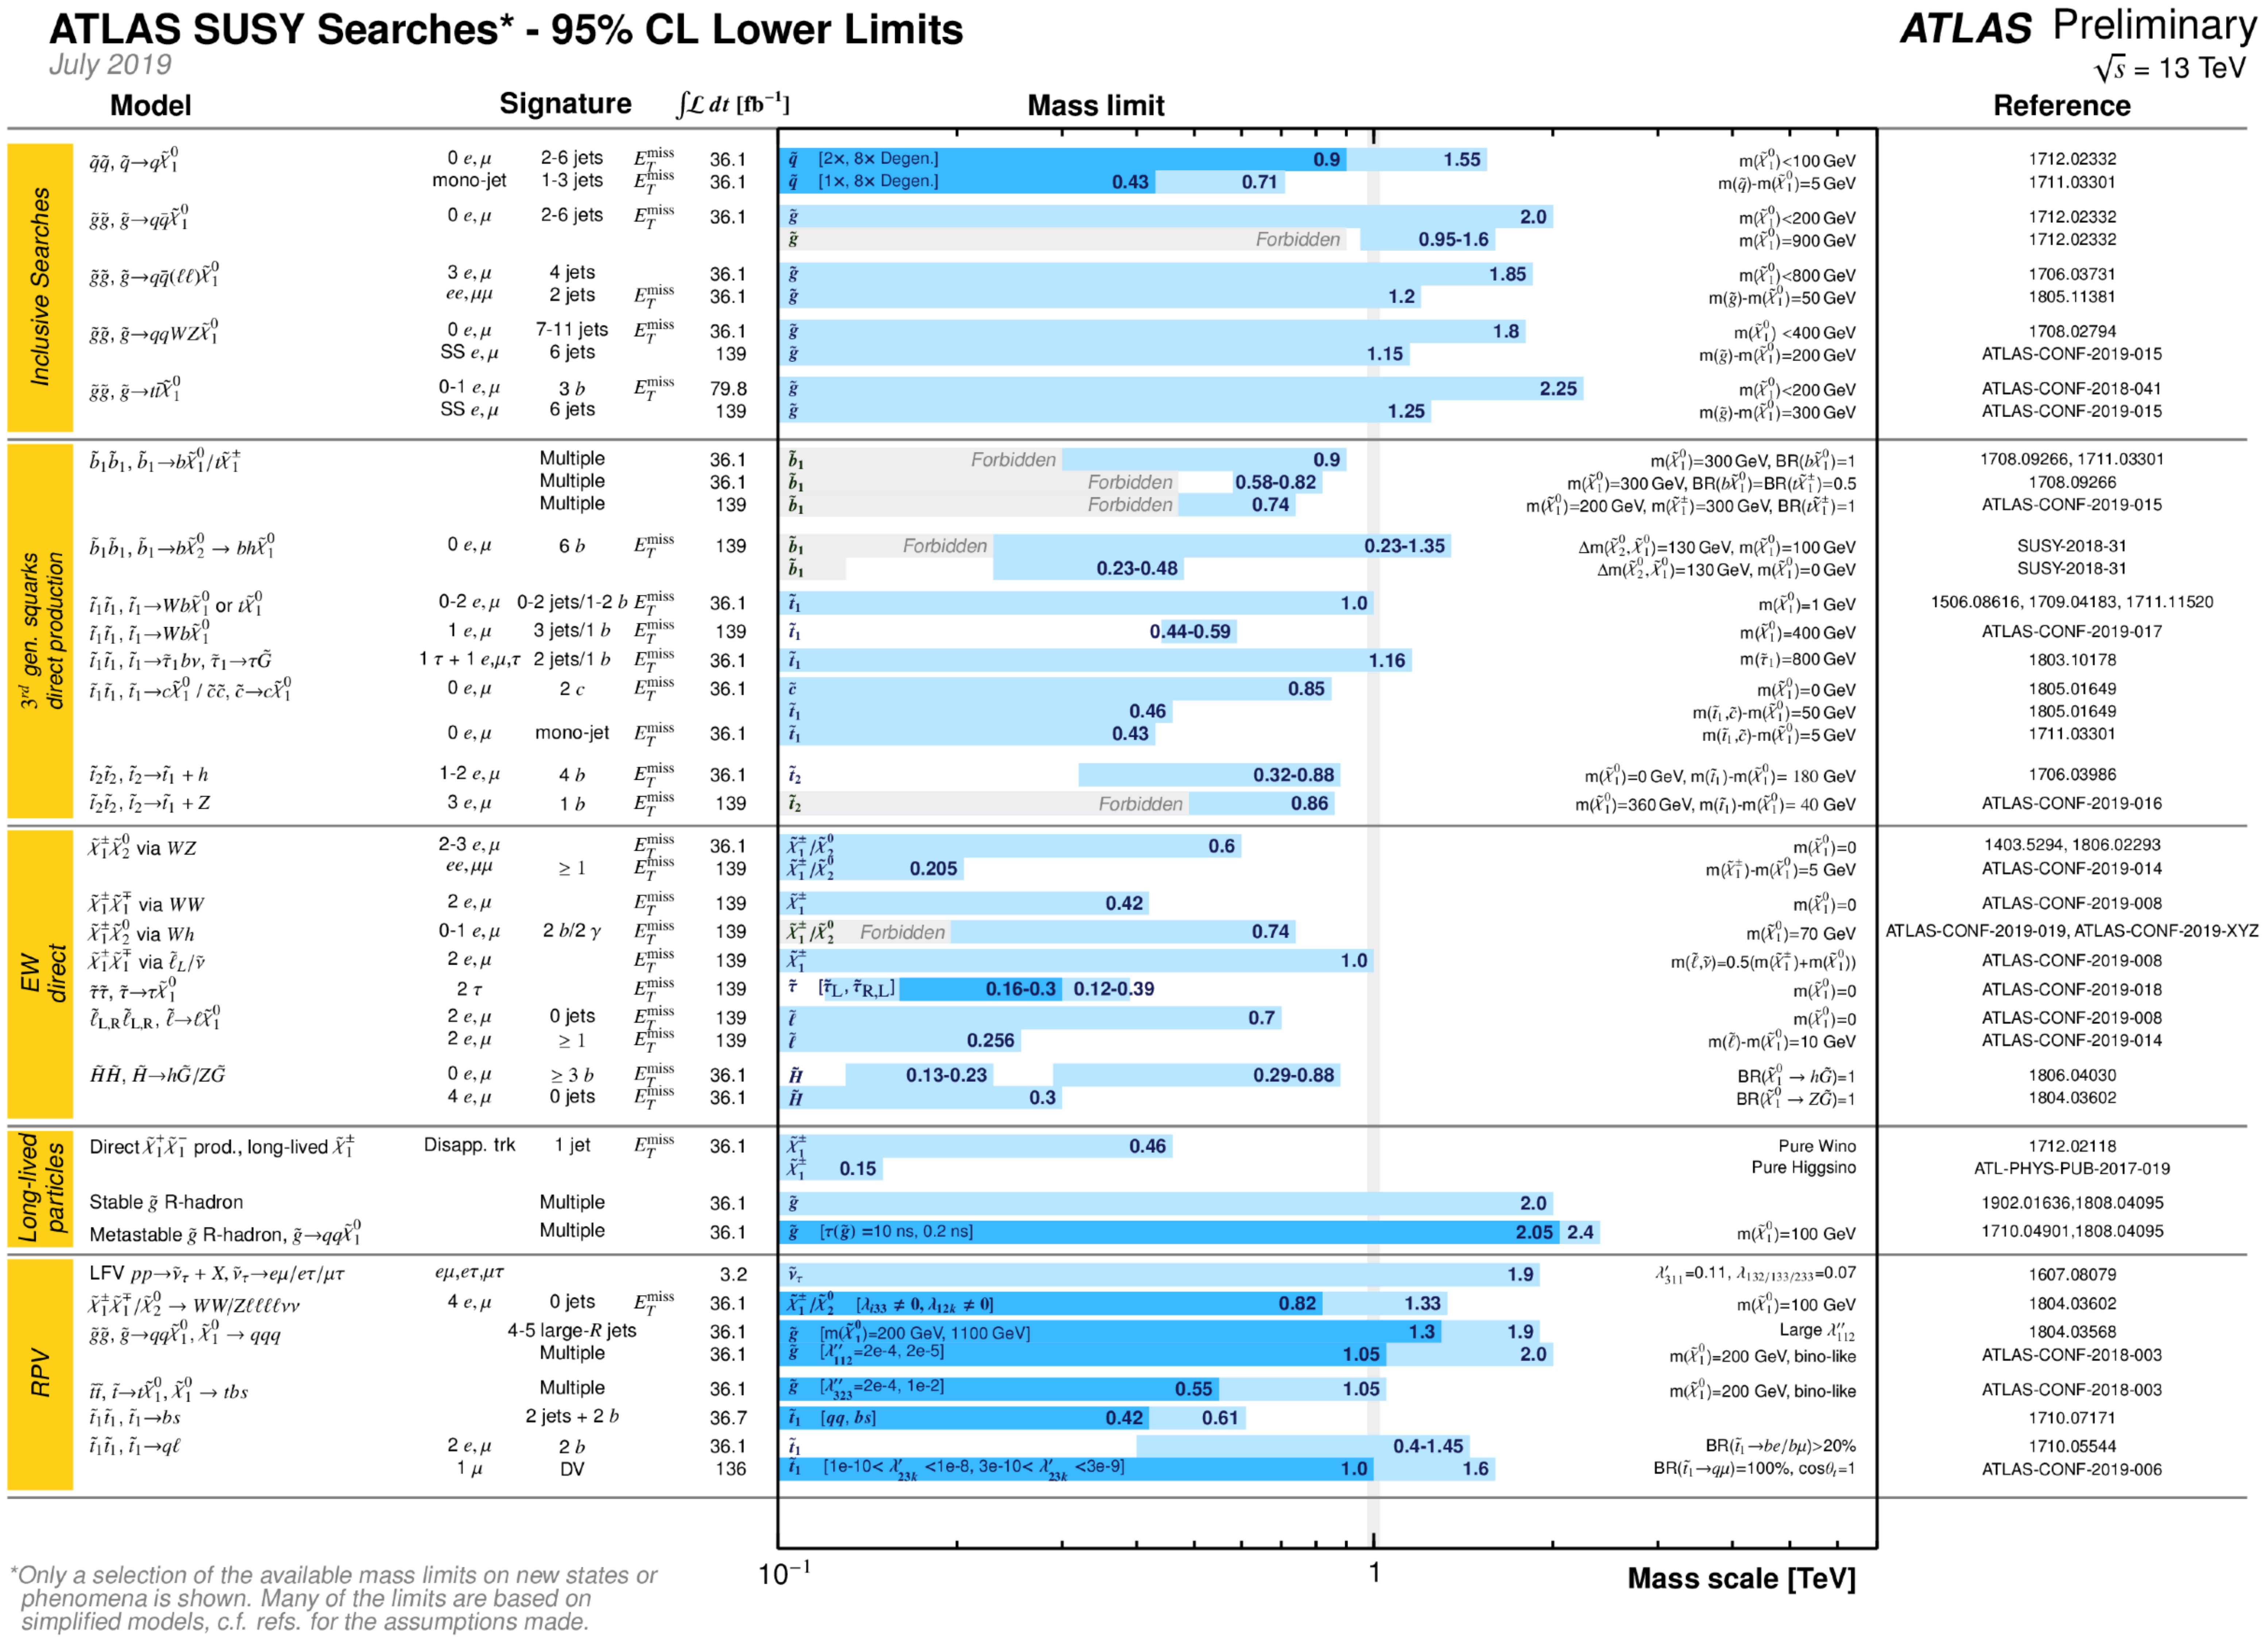

By now, there are countless searches for new physics at the LHC.
Almost all cut-based, model-specific.
No sign of new physics yet…

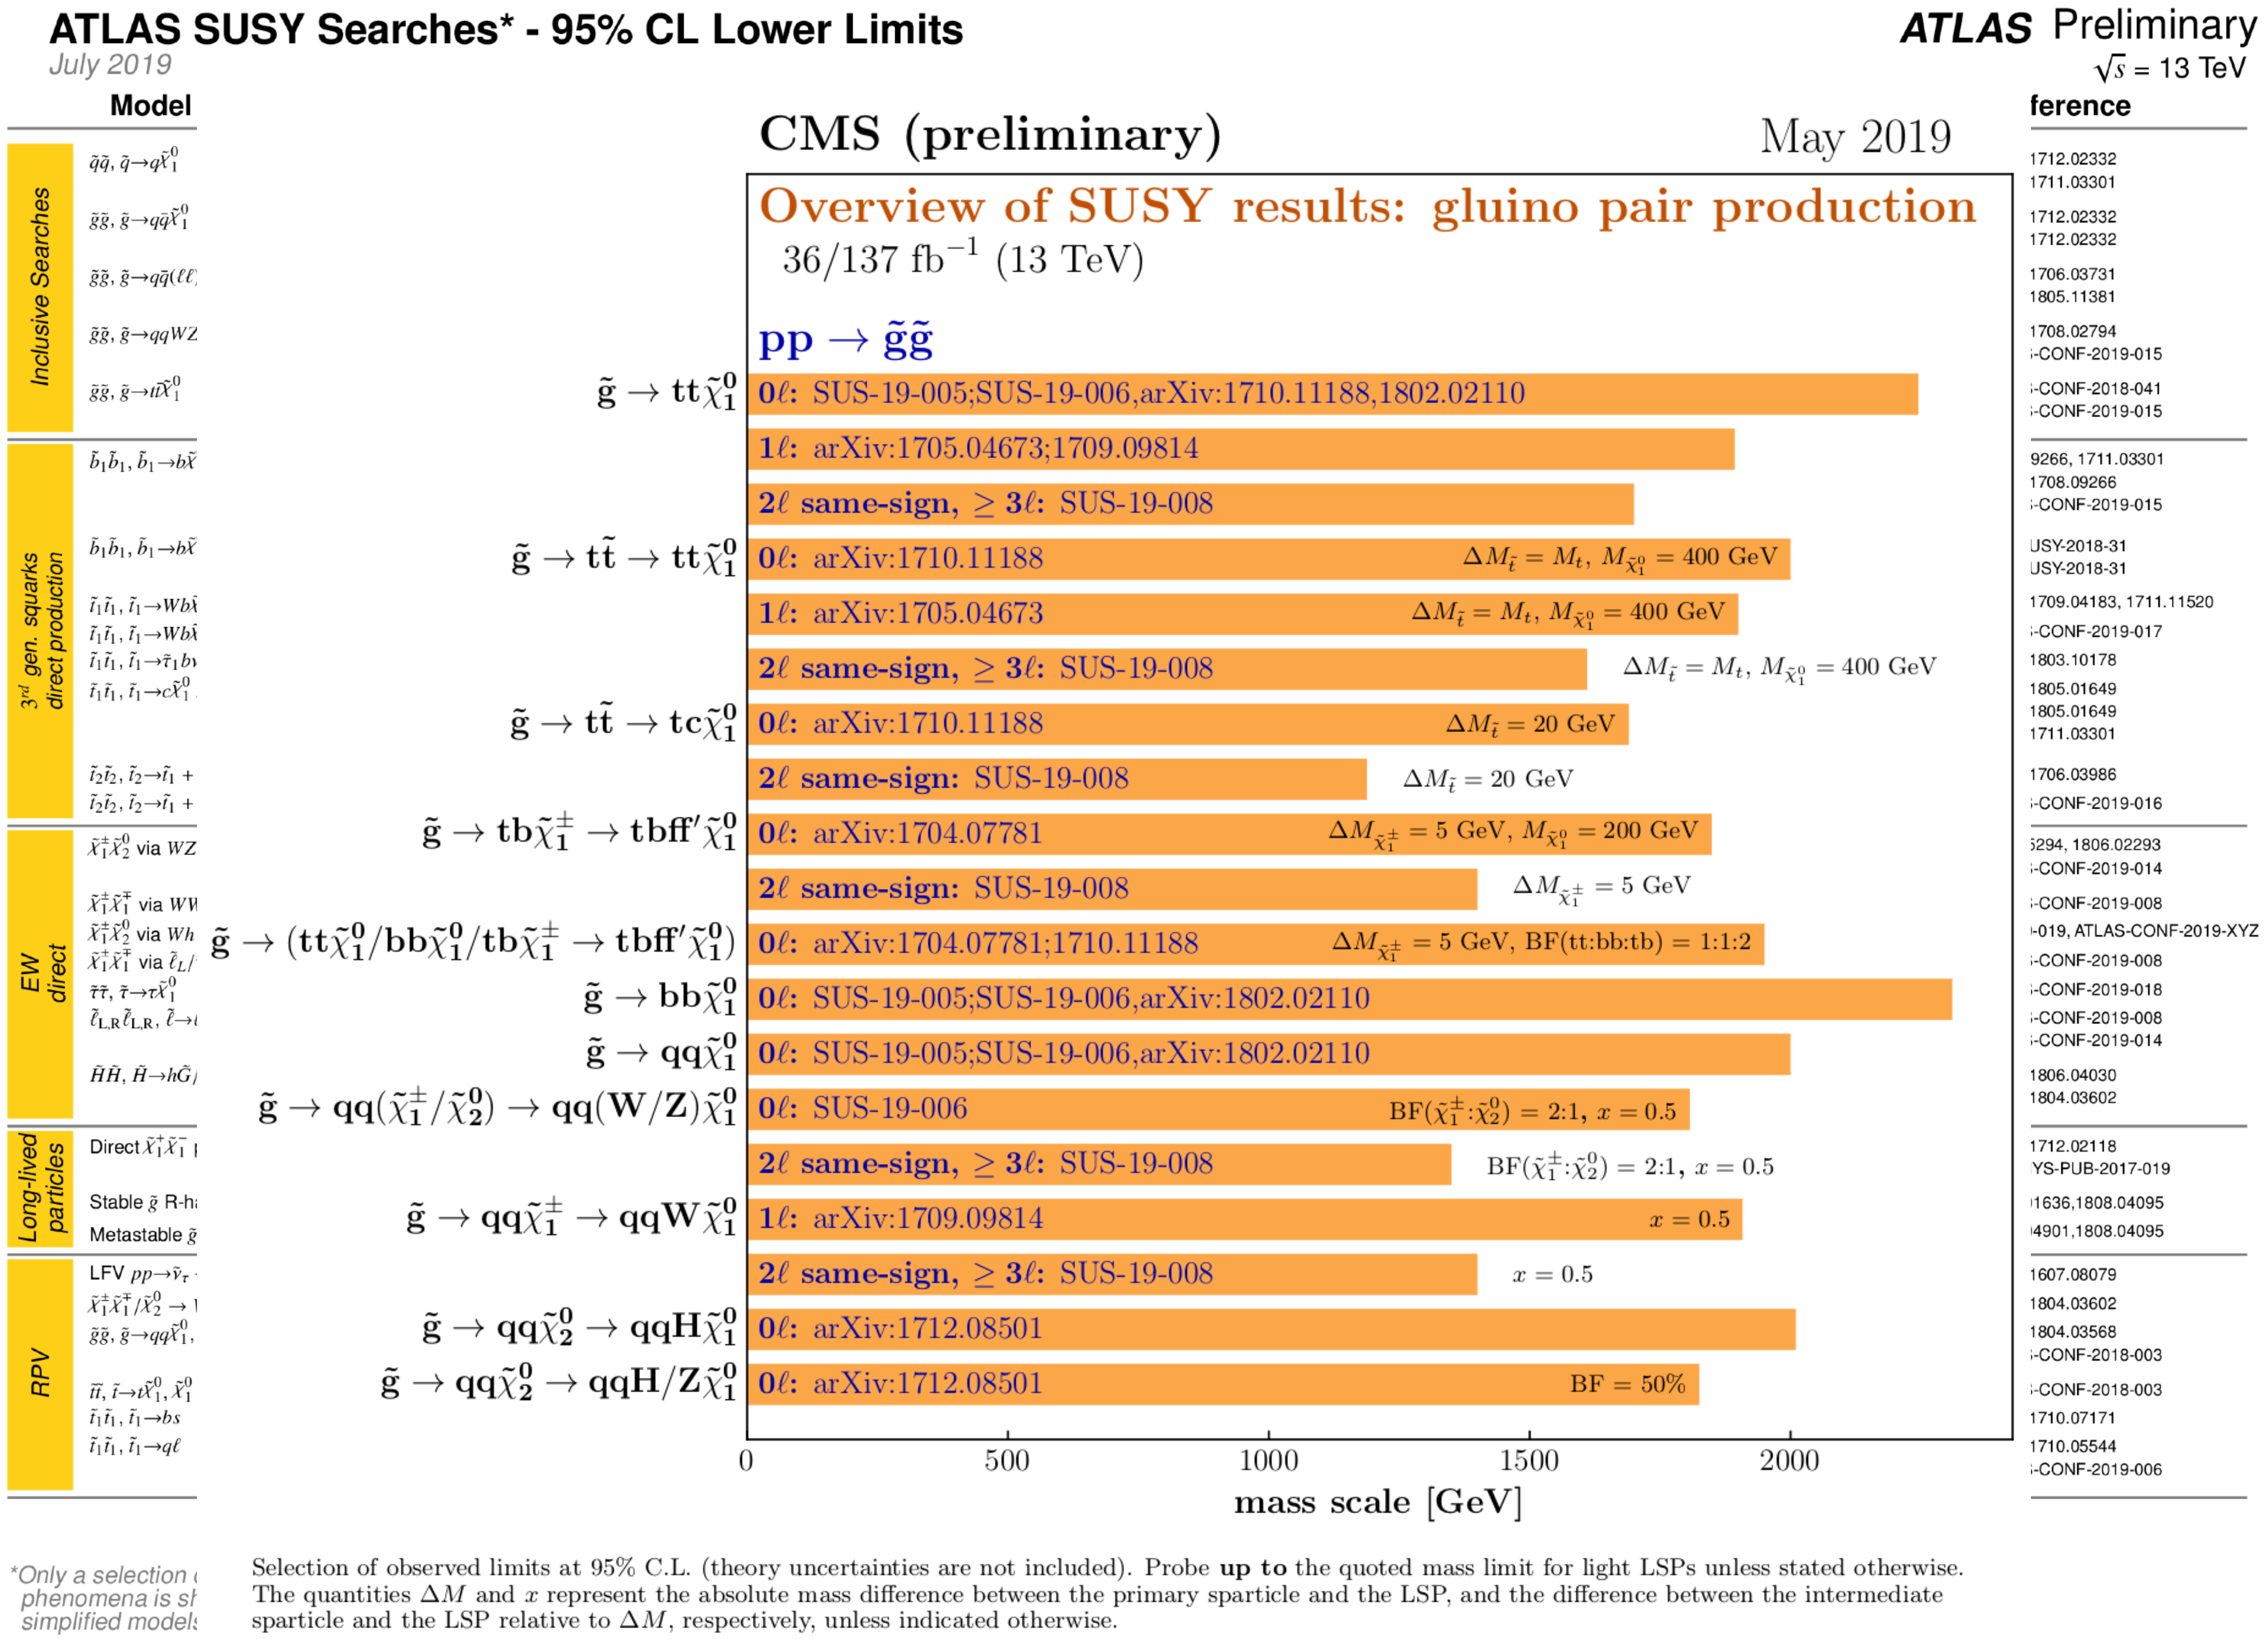

By now, there are countless searches for new physics at the LHC.
Almost all cut-based, model-specific.
No sign of new physics yet…

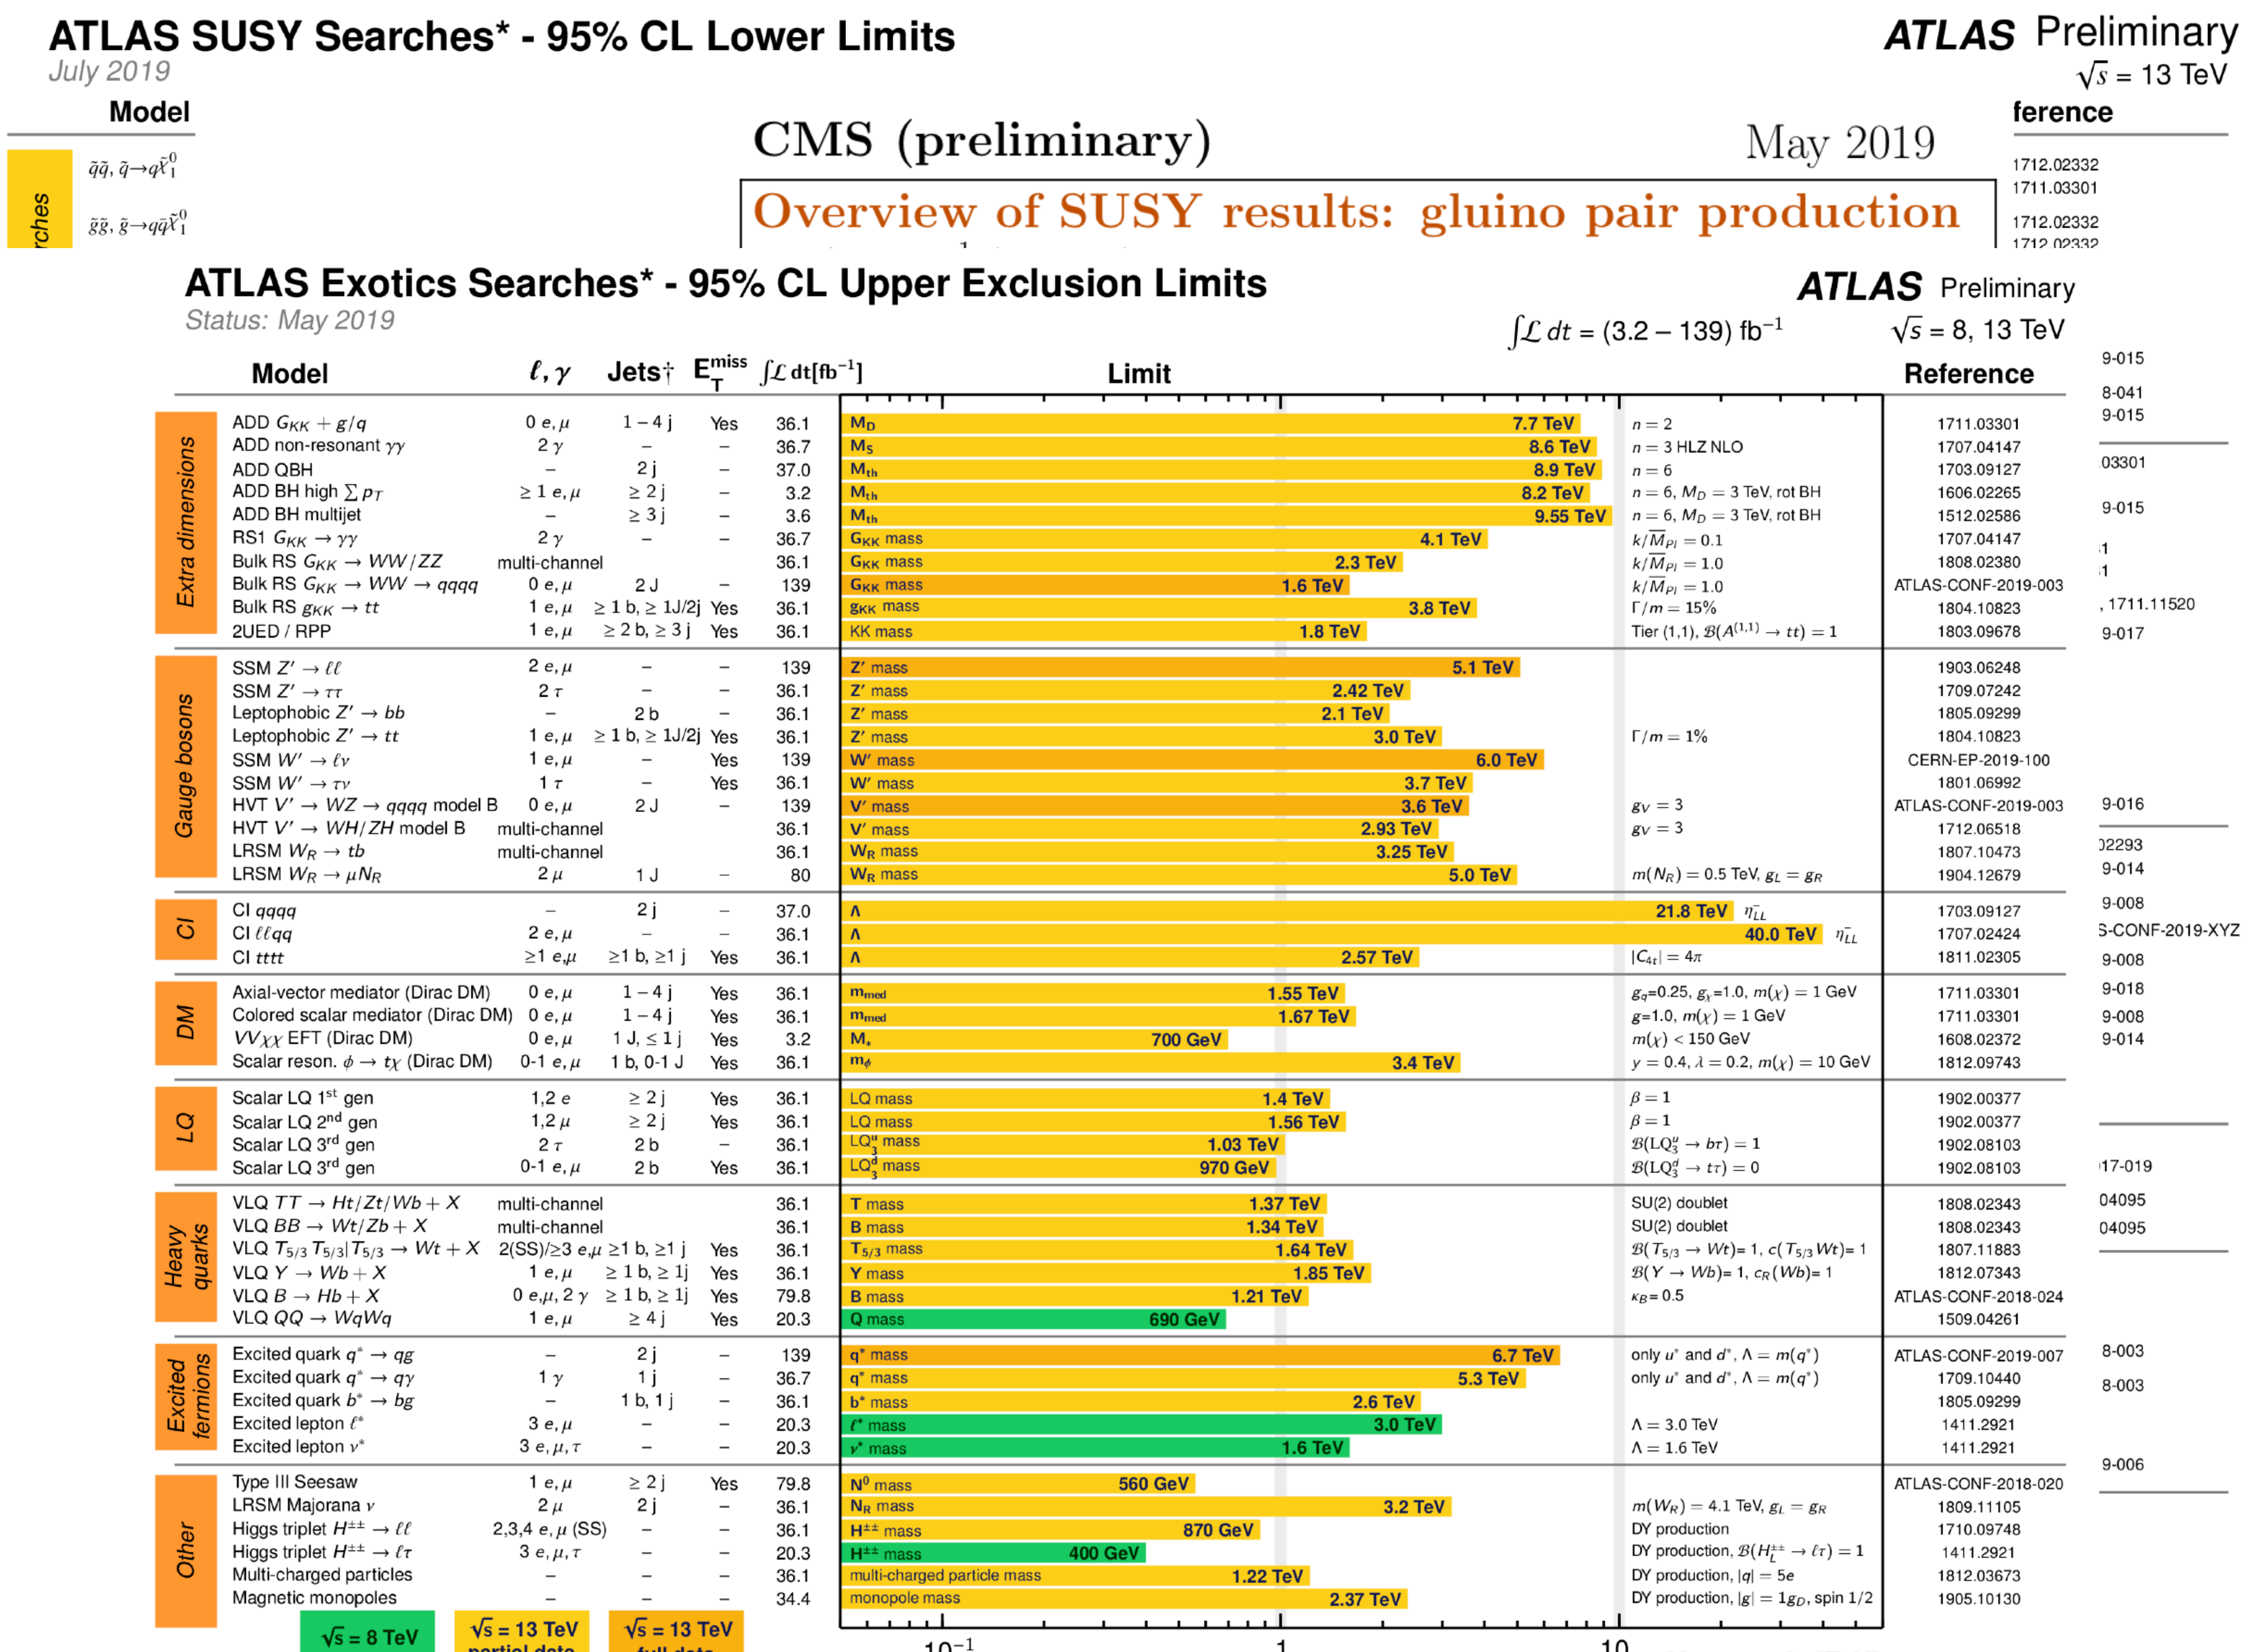

By now, there are countless searches for new physics at the LHC.
Almost all cut-based, model-specific.
No sign of new physics yet…

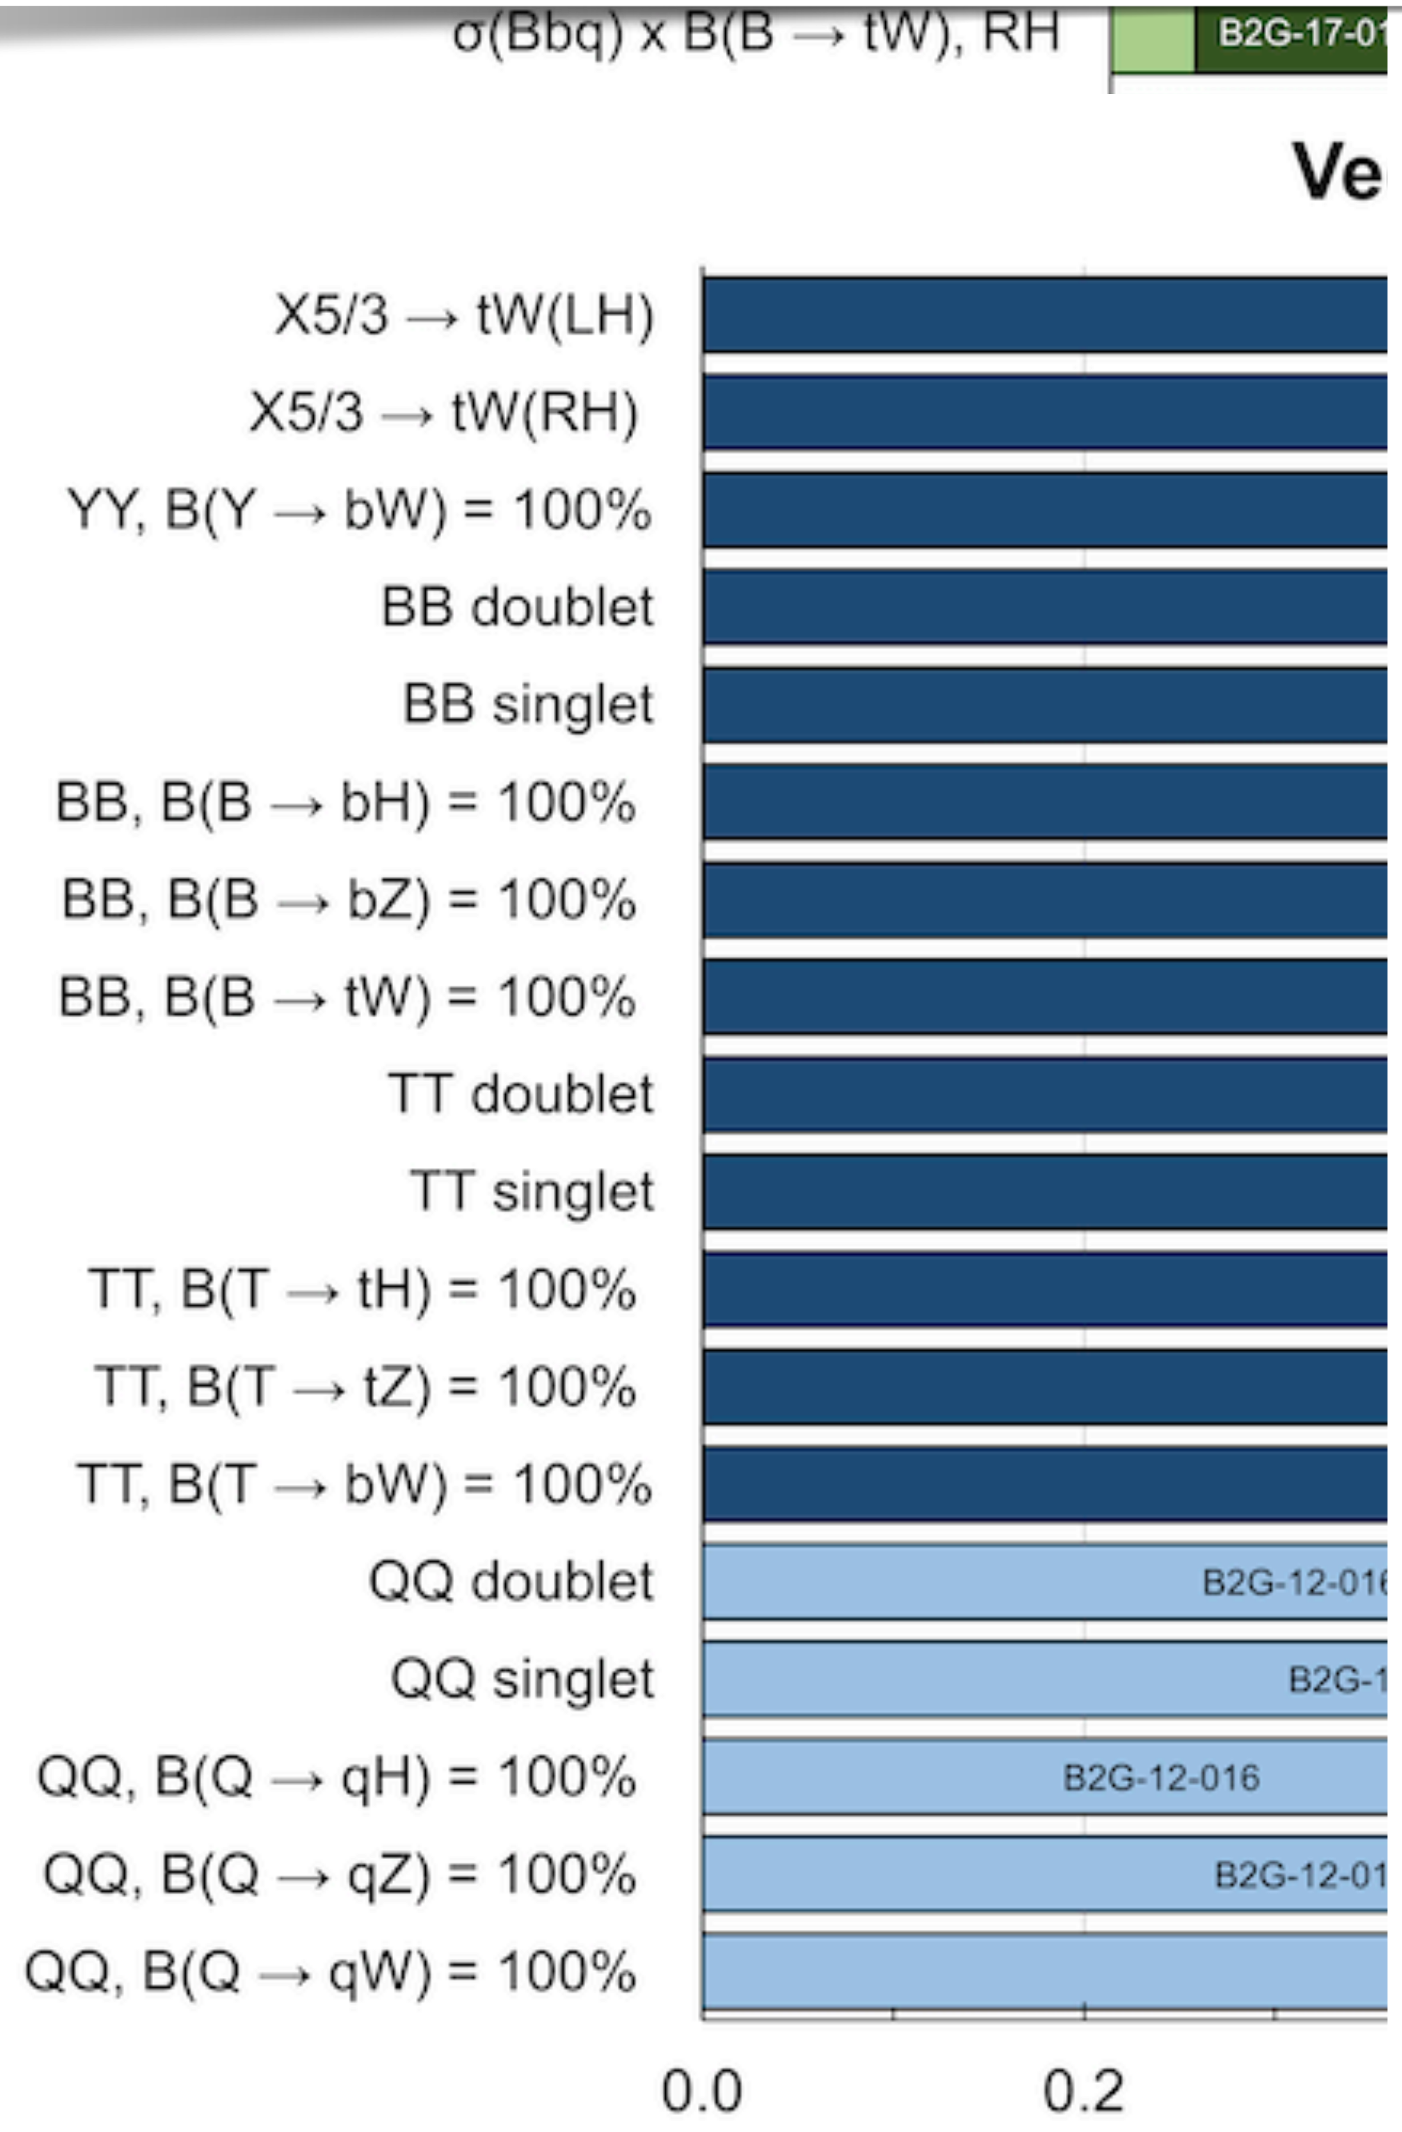

CMS, EPS-HEP 2019

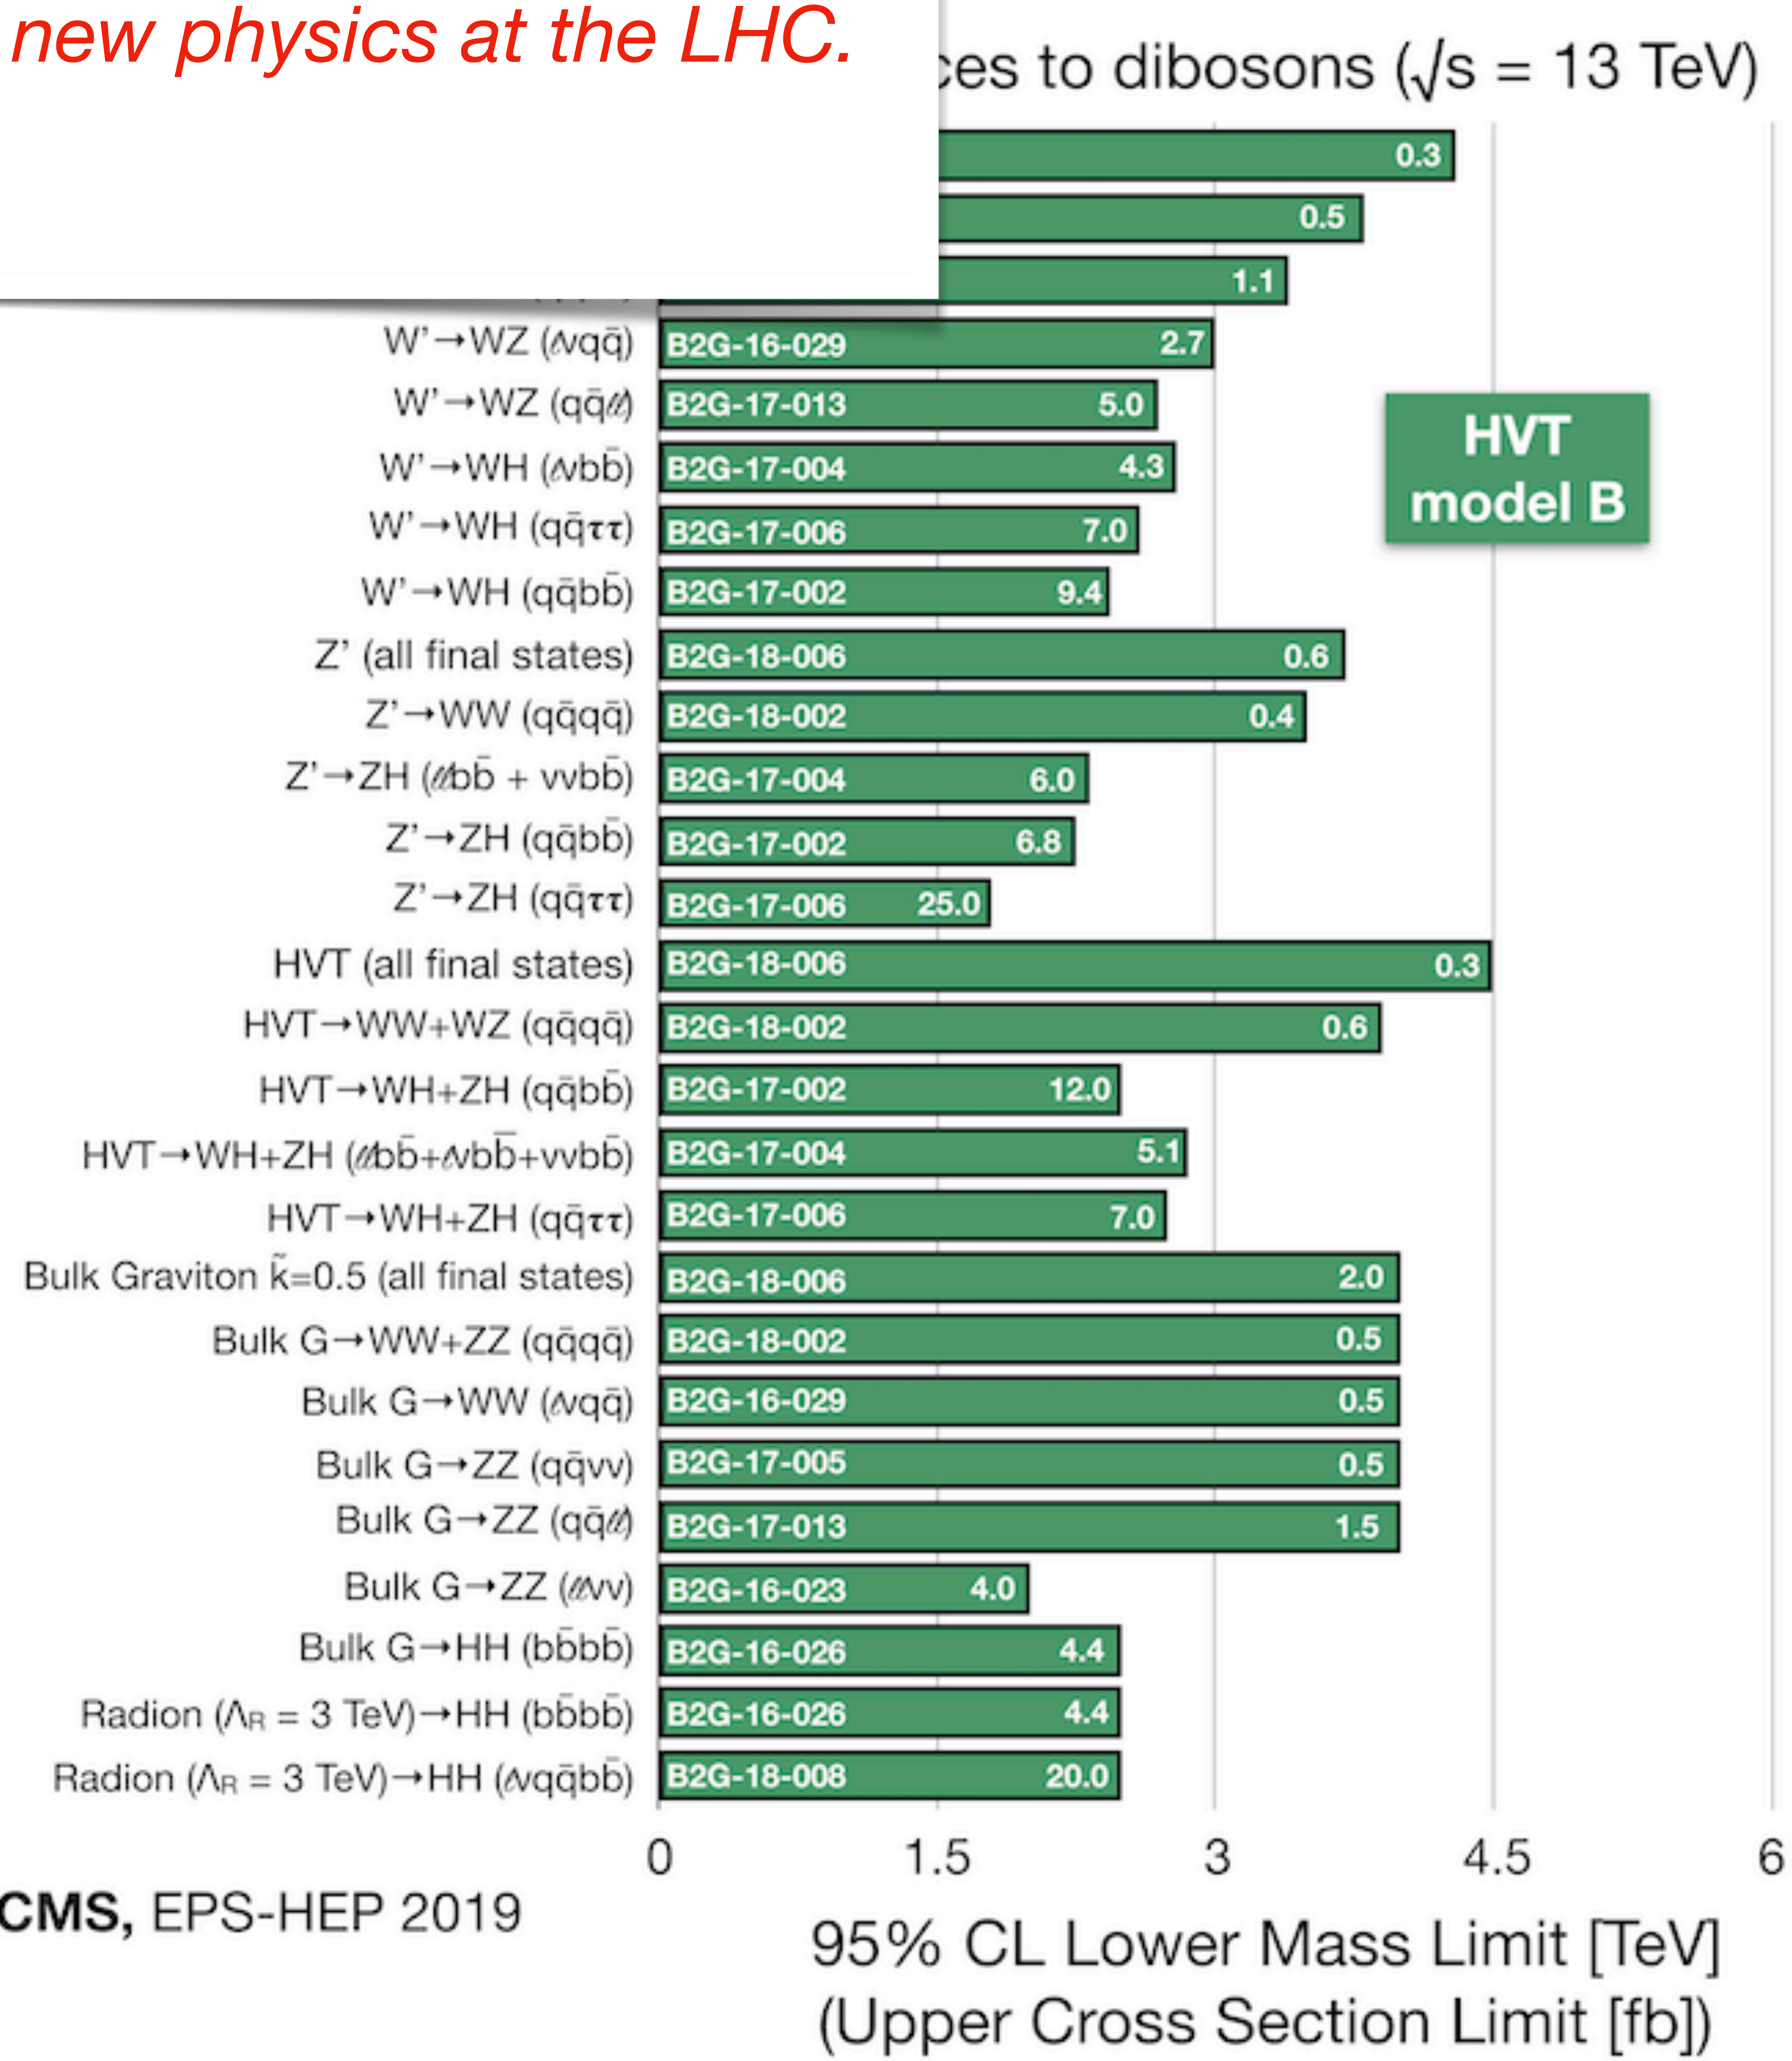

CMS, EPS-HEP 2019

Model-agnostic NP searches?

TABLE I. Existing two-body exclusive final state resonance searches at $\sqrt{s} = 8$ TeV. The $\varnothing$ symbol indicates no existing search at the LHC.

	e	μ	τ	γ	j	b	t	W	Z	h
e	$\pm\mp$[4],$\pm\pm$[5]	$\pm\pm$[5, 6] $\pm\mp$[6, 7]	[7]	$\varnothing$	$\varnothing$	$\varnothing$	$\varnothing$	$\varnothing$	$\varnothing$	$\varnothing$
μ		$\pm\mp$[4],$\pm\pm$[5]	[7]	$\varnothing$	$\varnothing$	$\varnothing$	$\varnothing$	$\varnothing$	$\varnothing$	$\varnothing$
τ			[8]	$\varnothing$	$\varnothing$	$\varnothing$	[9]	$\varnothing$	$\varnothing$	$\varnothing$
γ				[10]	[11–13]	$\varnothing$	$\varnothing$	[14]	[14]	$\varnothing$
j					[15]	[16]	[17]	[18]	[18]	$\varnothing$
b						[16]	[19]	$\varnothing$	$\varnothing$	$\varnothing$
t							[20]	[21]	$\varnothing$	$\varnothing$
W								[22–25]	[23, 24, 26, 27]	[28–30]
Z									[23, 25, 31]	[28, 30, 32, 33]
h										[34–37]

From Craig, Draper, Kong, Ng & Whiteson 1610.09392

What if NP is not like any of the models we have searched for?

There are a lot of NP scenarios that are not covered by existing searches!

	e	μ	τ	q/g	b	t	γ	Z/W	H	BSM $\to$ SM$_1$ $\times$ SM$_1$: q/g	γ/π^0's	b	$\cdots$	BSM $\to$ SM$_1$ $\times$ SM$_2$: tZ/H	bH	$\cdots$	BSM $\to$ complex: $\tau qq'$	eqq'	$\mu qq'$	$\cdots$
e	[37, 38]	[39, 40]	[39]	$\varnothing$	$\varnothing$	$\varnothing$	[41]	[42]	$\varnothing$	$\varnothing$	$\varnothing$	$\varnothing$		$\varnothing$	$\varnothing$	$\varnothing$	$\varnothing$	[43, 44]	$\varnothing$	
μ		[37, 38]	[39]	$\varnothing$	$\varnothing$	$\varnothing$	[41]	[42]	$\varnothing$	$\varnothing$	$\varnothing$	$\varnothing$		$\varnothing$	$\varnothing$	$\varnothing$	$\varnothing$	$\varnothing$	[43, 44]	
τ			[45, 46]	$\varnothing$	[47]	$\varnothing$	$\varnothing$	$\varnothing$	$\varnothing$	$\varnothing$	$\varnothing$	$\varnothing$		$\varnothing$	$\varnothing$	$\varnothing$	[48, 49]	$\varnothing$	$\varnothing$	
q/g				[29, 30, 50, 51]	[52]	$\varnothing$	[53, 54]	[55]	$\varnothing$	$\varnothing$	$\varnothing$	$\varnothing$		$\varnothing$	$\varnothing$	$\varnothing$	$\varnothing$	$\varnothing$	$\varnothing$	
b					[29, 52, 56]	[57]	[54]	[58]	[59]	$\varnothing$	$\varnothing$	$\varnothing$		[60]	$\varnothing$	$\varnothing$	$\varnothing$	$\varnothing$	$\varnothing$	
t						[61]	$\varnothing$	[62]	[63]	$\varnothing$	$\varnothing$	$\varnothing$		[64]	[60]	$\varnothing$	$\varnothing$	$\varnothing$	$\varnothing$	
γ							[65, 66]	[67–69]	[68, 70]	$\varnothing$	$\varnothing$	$\varnothing$		$\varnothing$	$\varnothing$	$\varnothing$	$\varnothing$	$\varnothing$	$\varnothing$	
Z/W								[71]	[71]	$\varnothing$	$\varnothing$	$\varnothing$		$\varnothing$	$\varnothing$	$\varnothing$	$\varnothing$	$\varnothing$	$\varnothing$	
H									[72, 73]	[74]	$\varnothing$	$\varnothing$		$\varnothing$	$\varnothing$	$\varnothing$	$\varnothing$	$\varnothing$	$\varnothing$	
BSM $\to$ SM$_1$ $\times$ SM$_1$: q/g										$\varnothing$	$\varnothing$	$\varnothing$		$\varnothing$	$\varnothing$	$\varnothing$	$\varnothing$	$\varnothing$	$\varnothing$	
γ/π^0's											[75]	$\varnothing$		$\varnothing$	$\varnothing$	$\varnothing$	$\varnothing$	$\varnothing$	$\varnothing$	
b												[76, 77]		$\varnothing$	$\varnothing$	$\varnothing$	$\varnothing$	$\varnothing$	$\varnothing$	
$\vdots$																				
$\vdots$																				

From Kim, Kong, Nachman & Whiteson 1907.06659

Model-agnostic NP searches?

Why aren’t there more model-agnostic new physics searches?

Modern ML can help!

- ***model-specific ~ supervised ML***
- ***model-agnostic ~ less-than-supervised ML***
- ***unbinned analysis of high dimensional feature spaces (100s or 1000s of features)***
- ***(nearly) optimal classifiers***
- ***density estimation***
- ***generative modeling***
- ***…***

2. Model-agnostic NP searches with modern ML

Model-agnostic NP Searches @ LHC

The LHC Olympics 2020

A Community Challenge for Anomaly Detection in High Energy Physics

Gregor Kasieczka (ed),[1] Benjamin Nachman (ed),[2,3] David Shih (ed),[4] Oz Amram,[5] Anders Andreassen,[6] Kees Benkendorfer,[2,7] Blaz Bortolato,[8] Gustaaf Brooijmans,[9] Florencia Canelli,[10] Jack H. Collins,[11] Biwei Dai,[12] Felipe F. De Freitas,[13] Barry M. Dillon,[8,14] Ioan-Mihail Dinu,[5] Zhongtian Dong,[15] Julien Donini,[16] Javier Duarte,[17] D. A. Faroughy[10] Julia Gonski,[9] Philip Harris,[18] Alan Kahn,[9] Jernej F. Kamenik,[8,19] Charanjit K. Khosa,[20,30] Patrick Komiske,[21] Luc Le Pottier,[2,22] Pablo Martín-Ramiro,[2,23] Andrej Matevc,[8,19] Eric Metodiev,[21] Vinicius Mikuni,[10] Inês Ochoa,[24] Sang Eon Park,[18] Maurizio Pierini,[25] Dylan Rankin,[18] Veronica Sanz,[20,26] Nilai Sarda,[27] Uroš Seljak,[2,3,12] Aleks Smolkovic,[8] George Stein,[2,12] Cristina Mantilla Suarez,[5] Manuel Szewc,[28] Jesse Thaler,[21] Steven Tsan,[17] Silviu-Marian Udrescu,[18] Louis Vaslin,[16] Jean-Roch Vlimant,[29] Daniel Williams,[9] Mikaeel Yunus[18]

https://arxiv.org/abs/2101.08320

https://arxiv.org/abs/2105.14027

The Dark Machines Anomaly Score Challenge: Benchmark Data and Model Independent Event Classification for the Large Hadron Collider

T. Aarrestad[a] M. van Beekveld[b] M. Bona[c] A. Boveia[e] S. Caron[d] J. Davies[c] A. De Simone[f,g] C. Doglioni[h] J. M. Duarte[i] A. Farbin[j] H. Gupta[k] L. Hendriks[d] L. Heinrich[a] J. Howarth[l] P. Jawahar[m,a] A. Jueid[n] J. Lastow[h] A. Leinweber[o] J. Mamuzic[p] E. Merényi[q] A. Morandini[r] P. Moskvitina[d] C. Nellist[d] J. Ngadiuba[s,t] B. Ostdiek[u,v] M. Pierini[a] B. Ravina[l] R. Ruiz de Austri[p] S. Sekmen[w] M. Touranakou[x,a] M. Vaškevičiūte[l] R. Vilalta[y] J.-R. Vlimant[t] R. Verheyen[z] M. White[o] E. Wulff[h] E. Wallin[h] K.A. Wozniak[α,a] Z. Zhang[d]

A lot of new ideas for model-agnostic searches!

Model-agnostic NP Searches @ LHC

Proofs-of-concept are becoming actual LHC searches!

CWoLa Hunting
ATLAS, PRL **125** 131801 (2020)

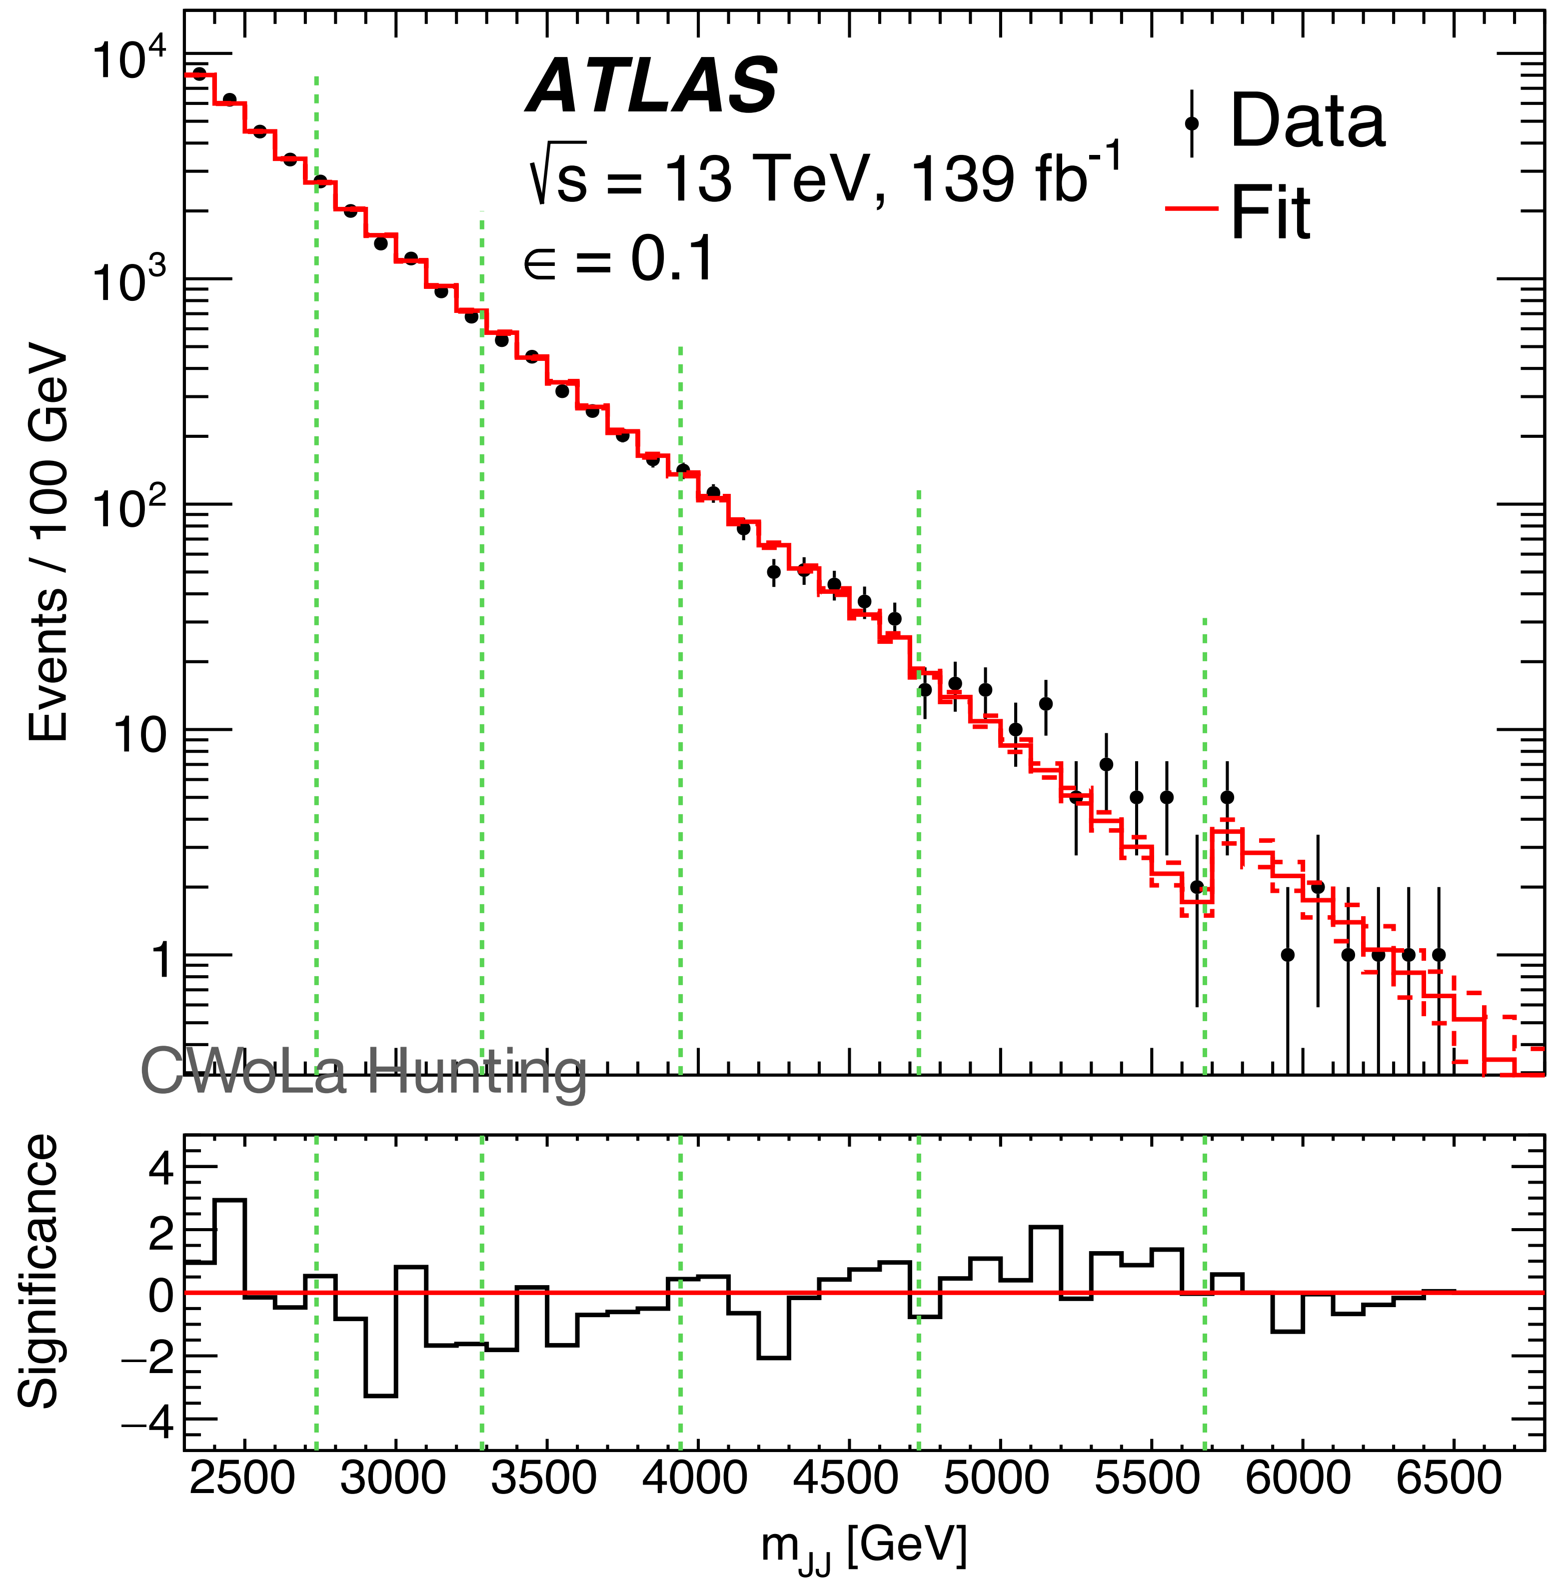

RNN VAE
ATLAS-CONF-2022-045

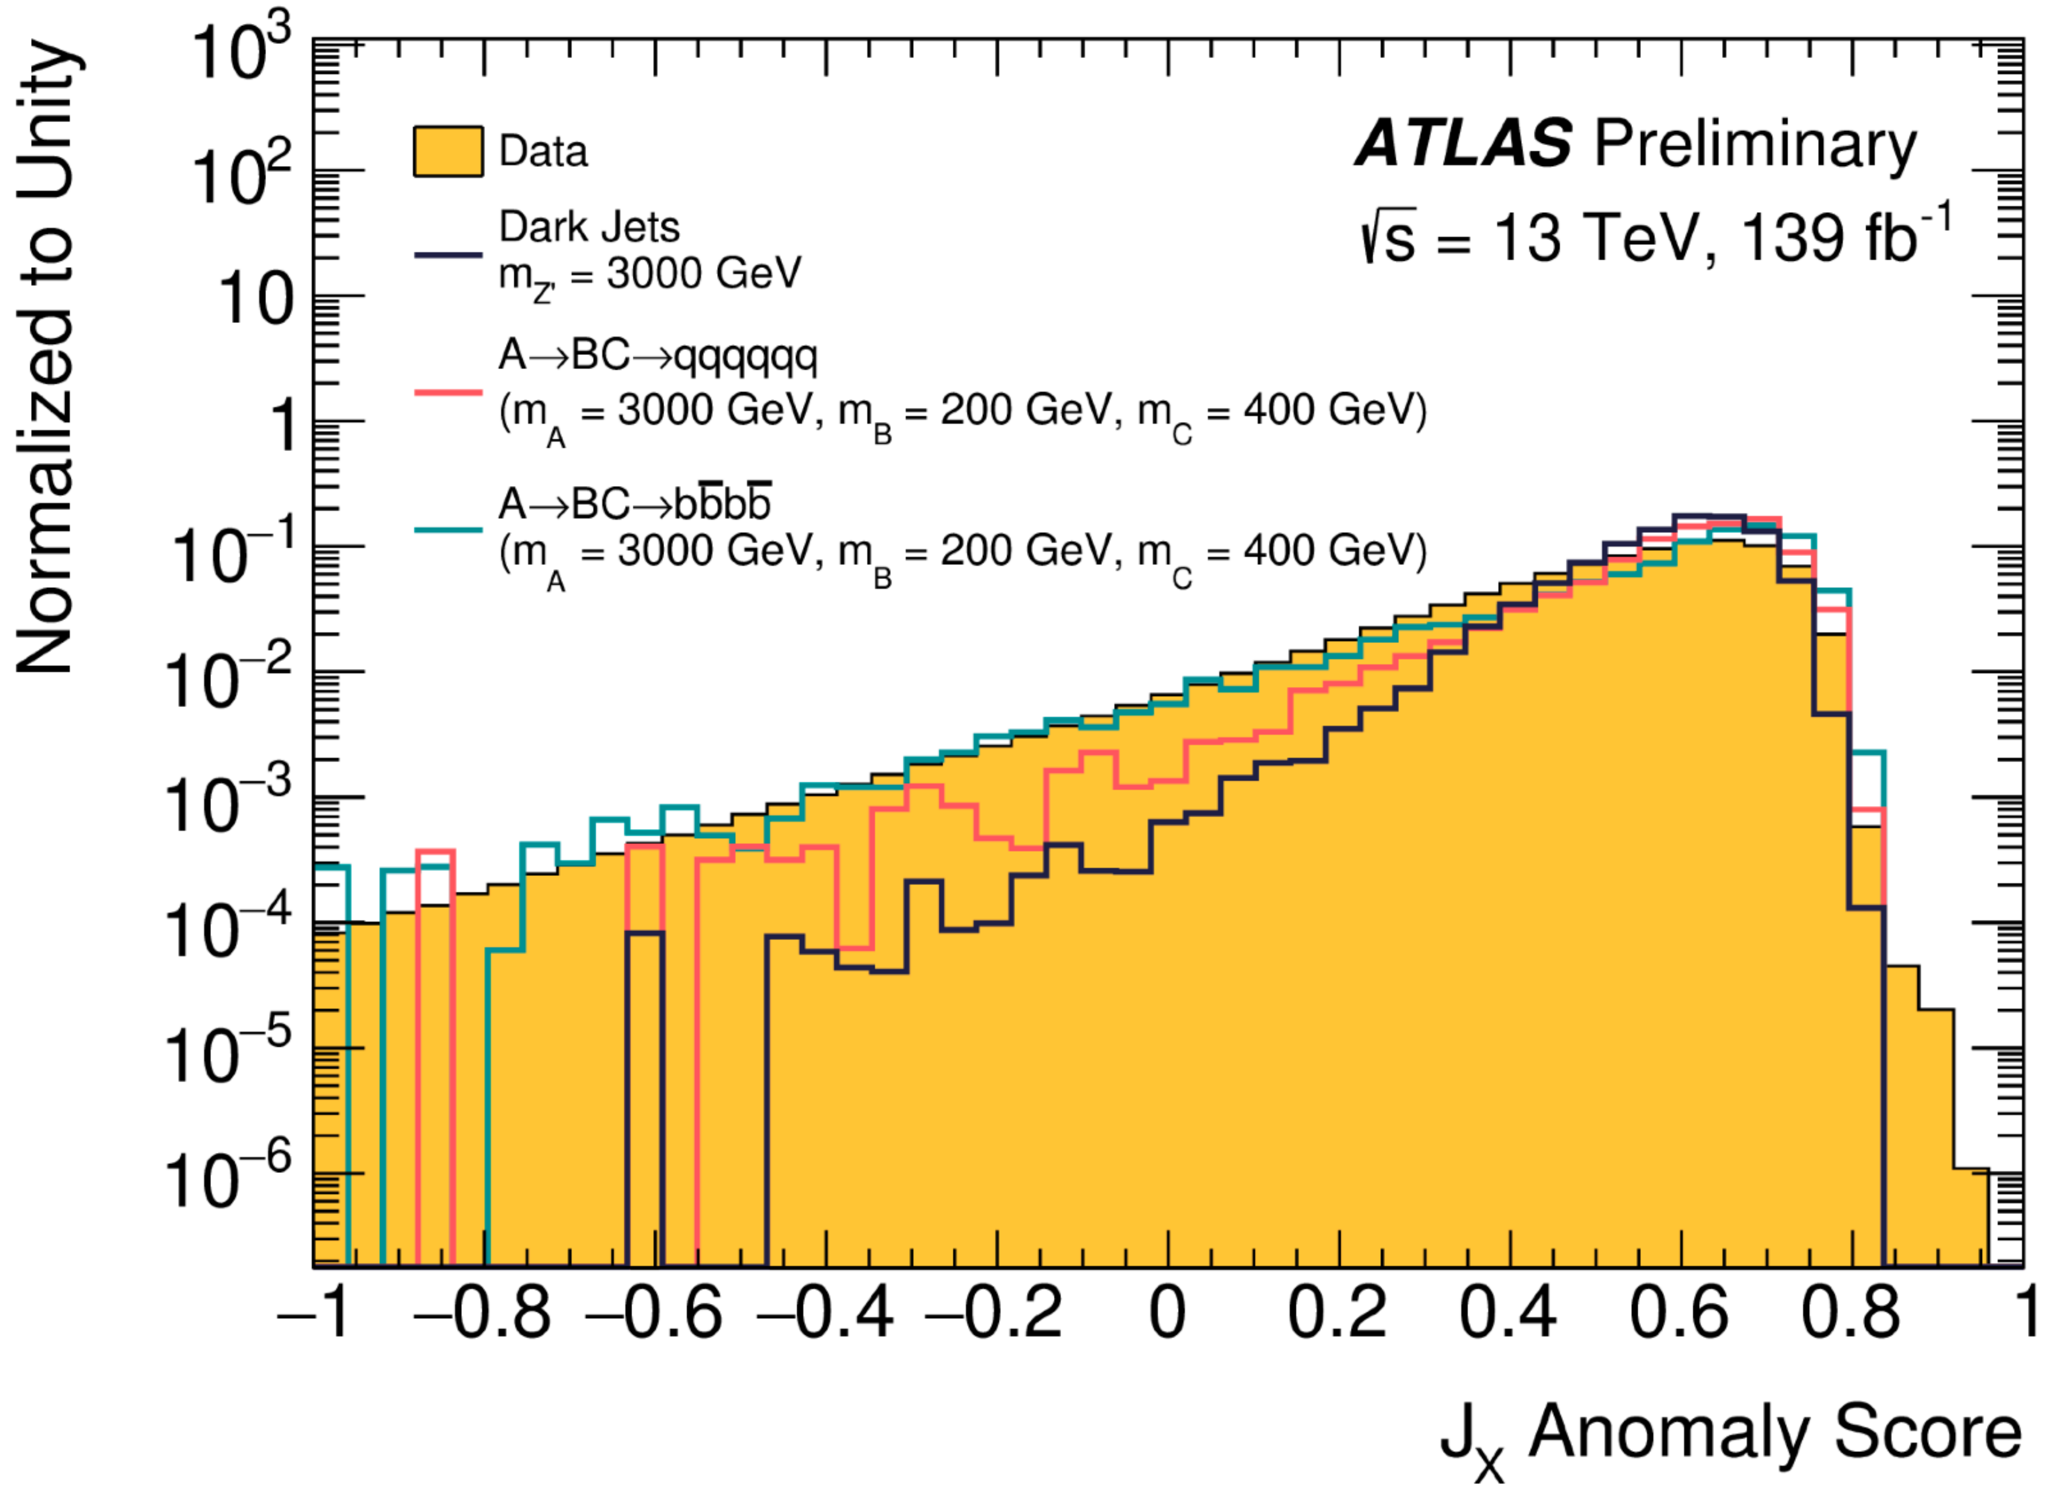

Model-agnostic NP Searches @ LHC

Proofs-of-concept are becoming actual LHC searches!

CWoLa Hunting
ATLAS, PRL **125** 131801 (2020)

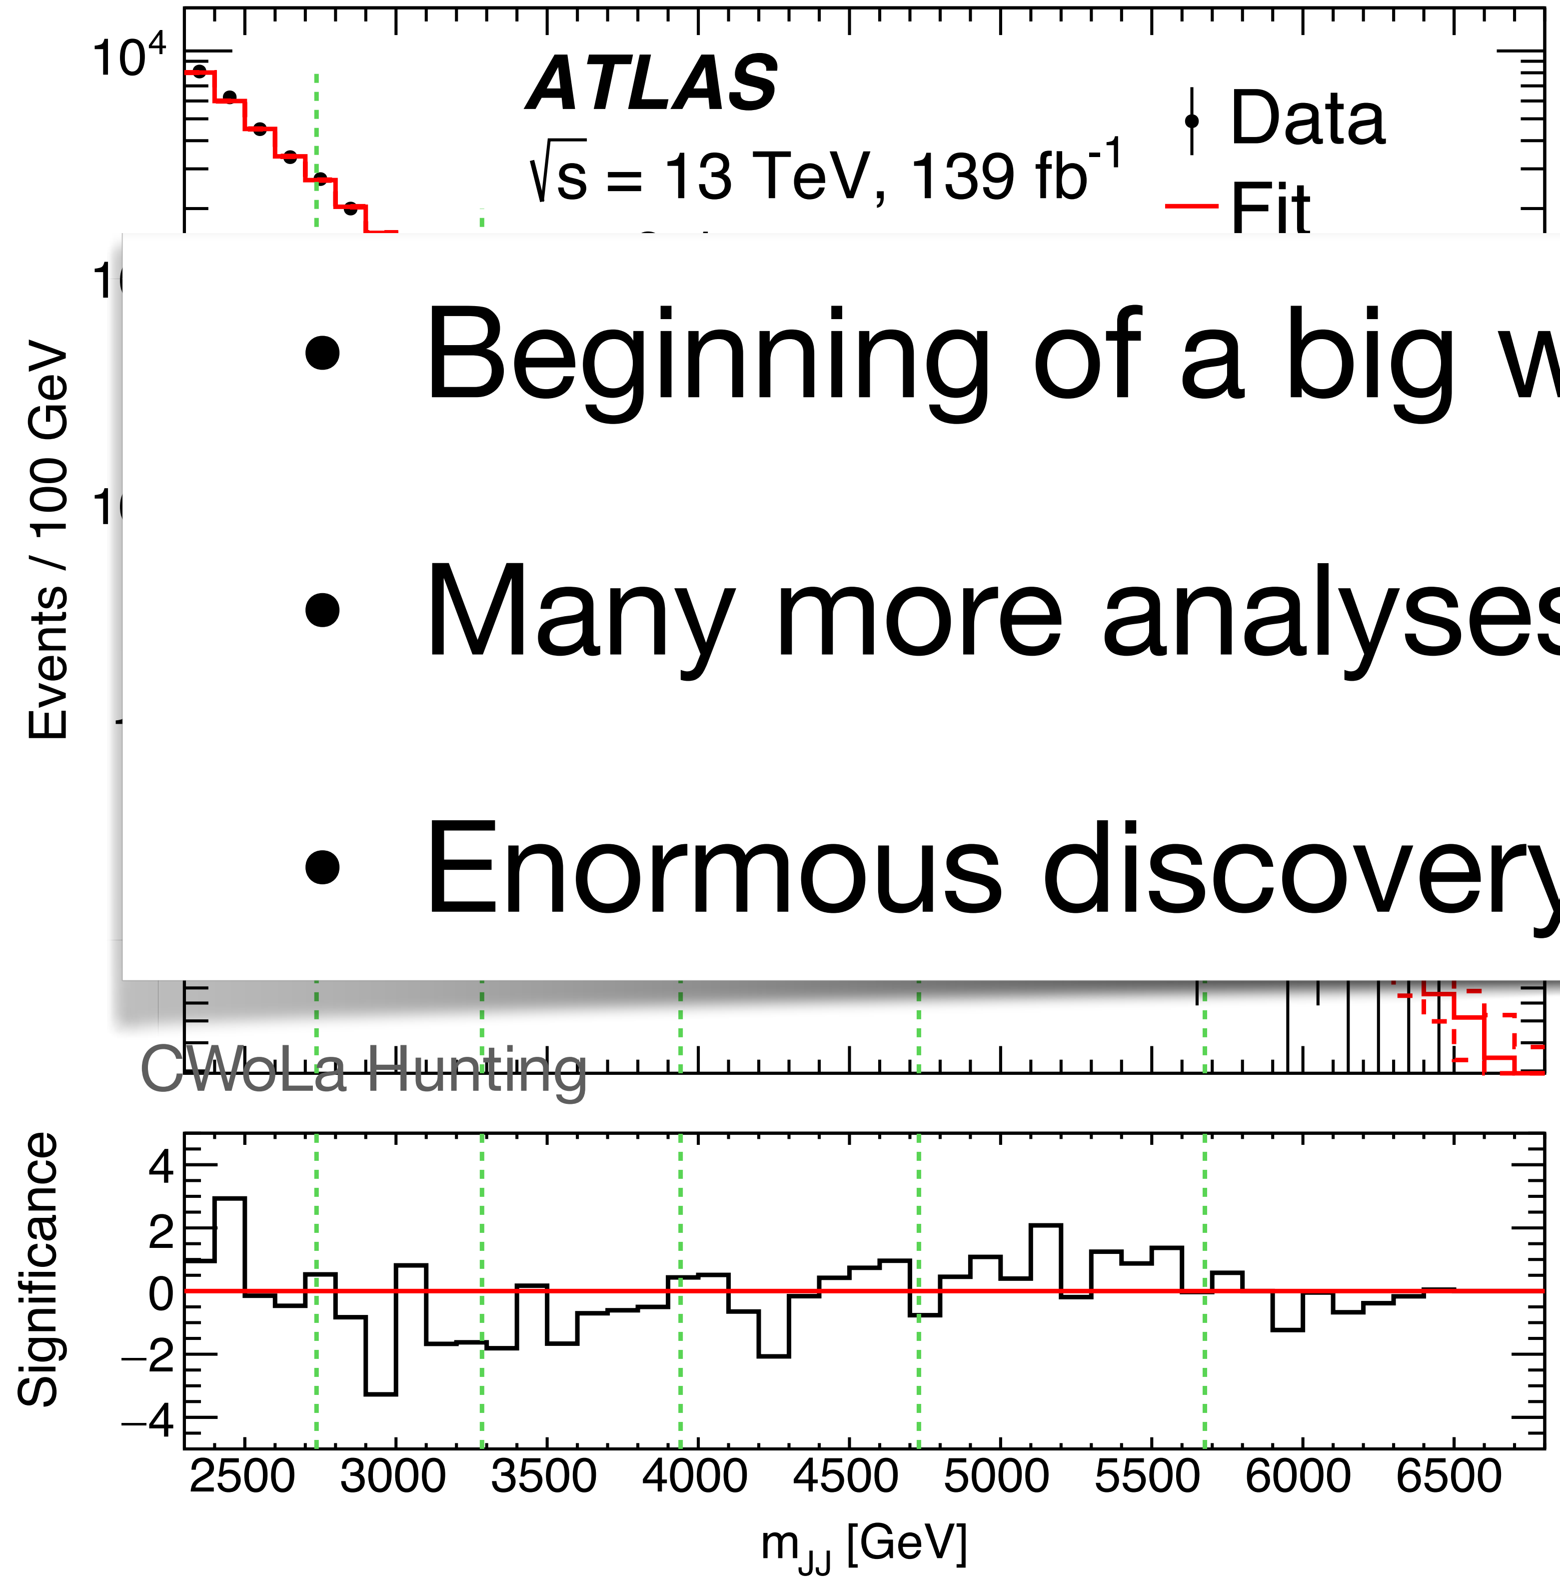

RNN VAE
ATLAS-CONF-2022-045

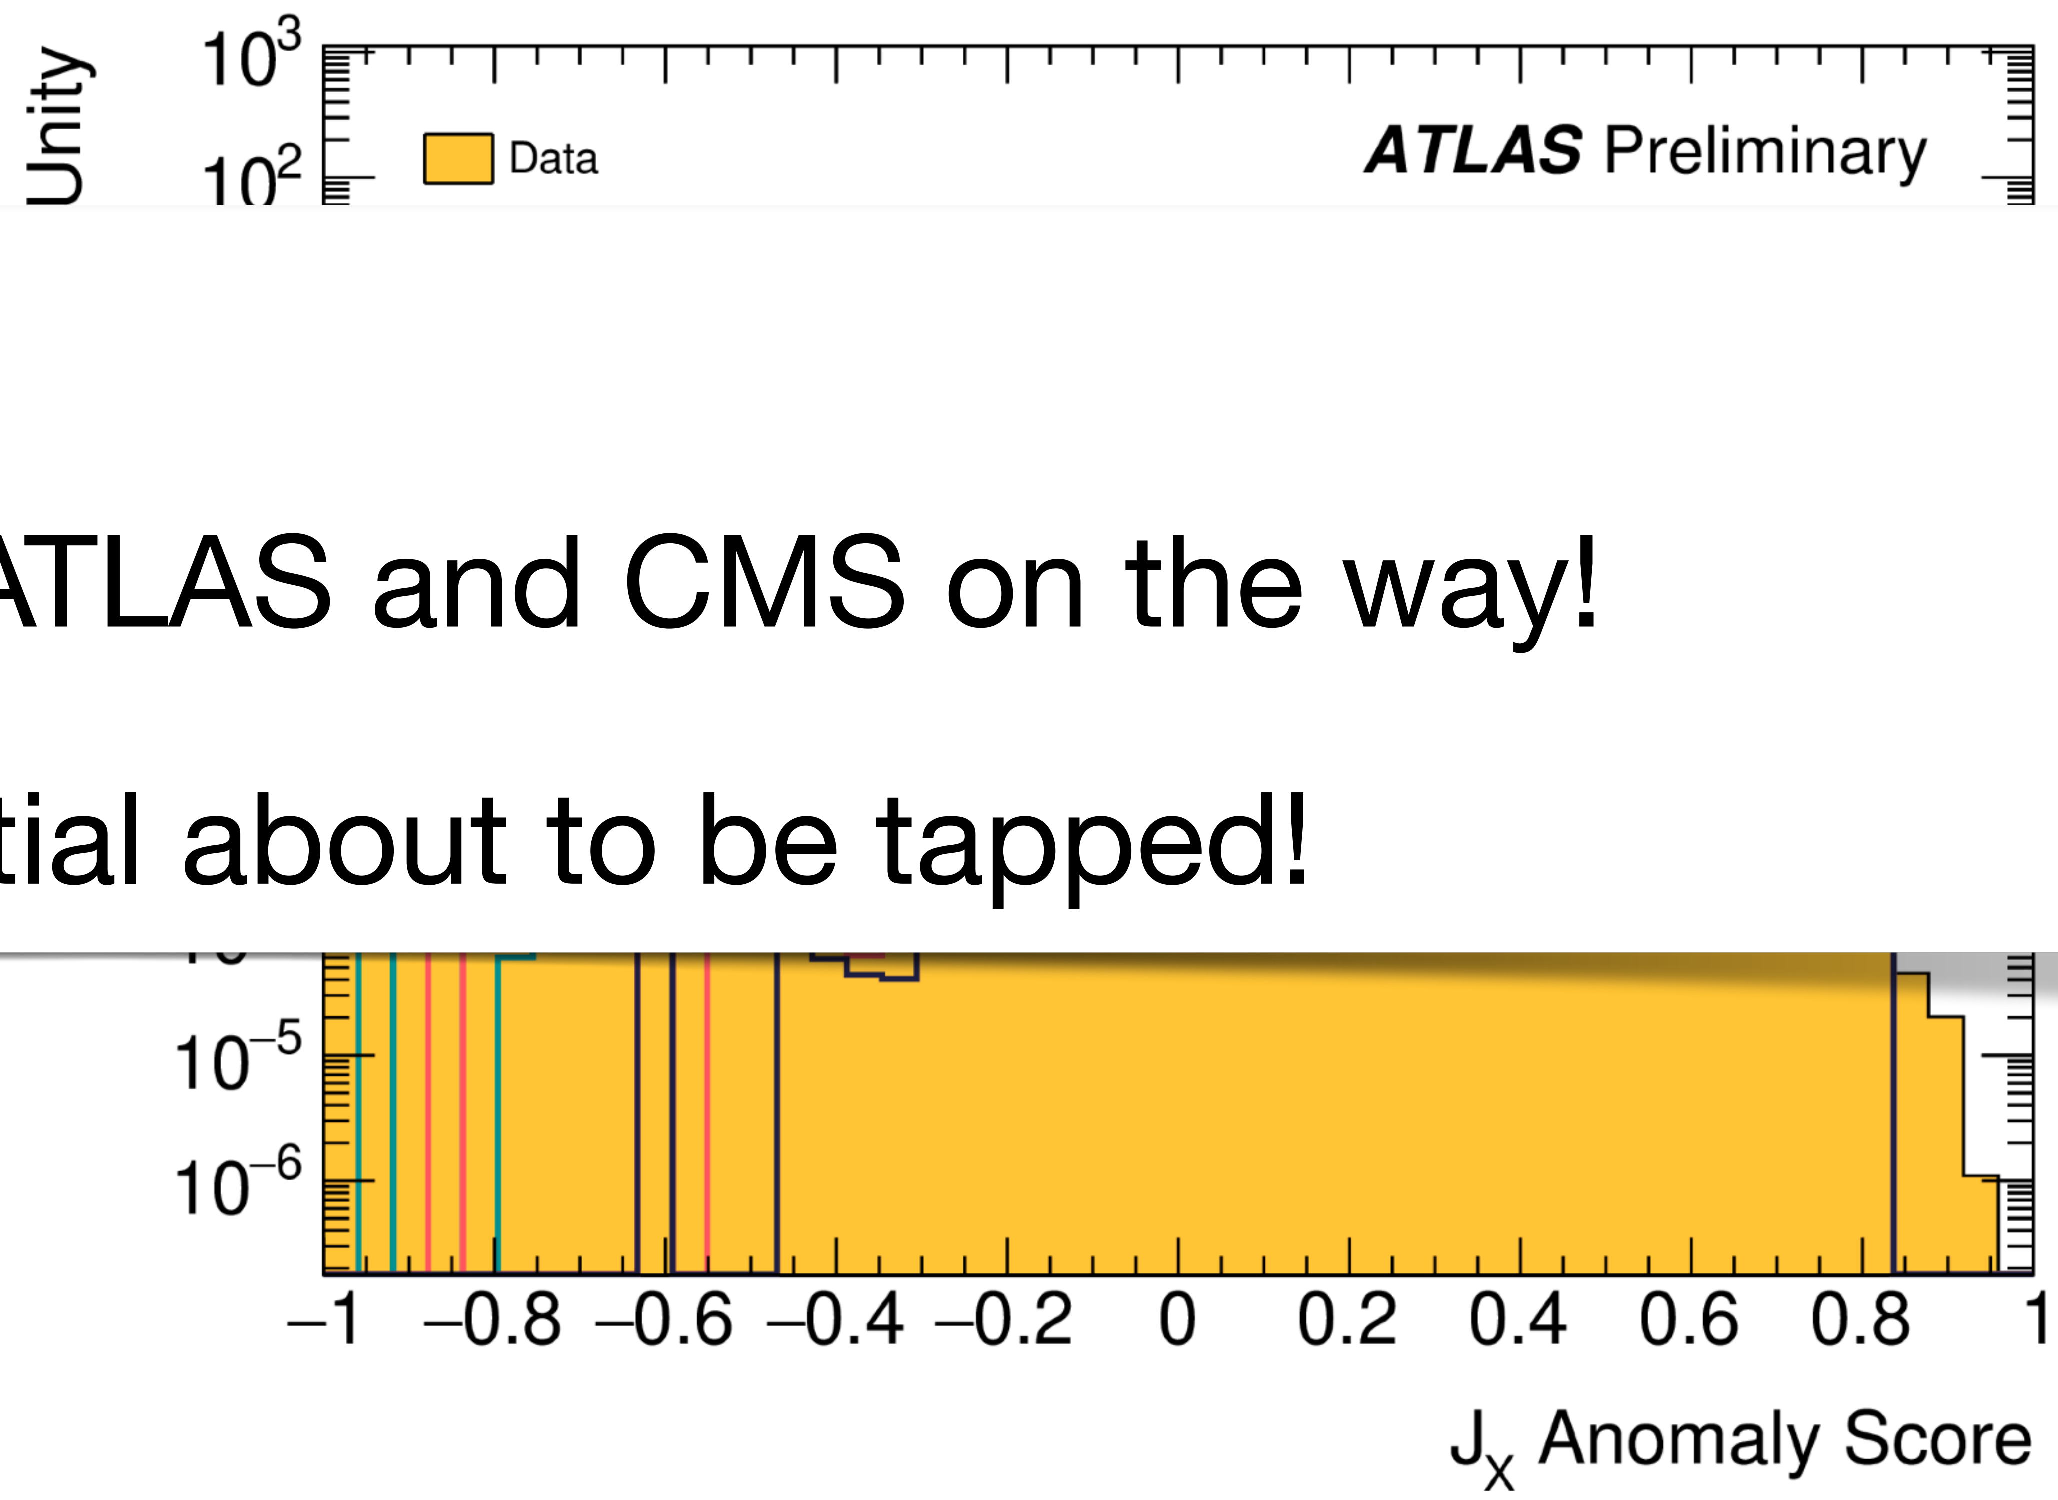

- Beginning of a big wave?
- Many more analyses from ATLAS and CMS on the way!
- Enormous discovery potential about to be tapped!

Two modes of anomaly detection

a. Outlier detection ("point anomalies")

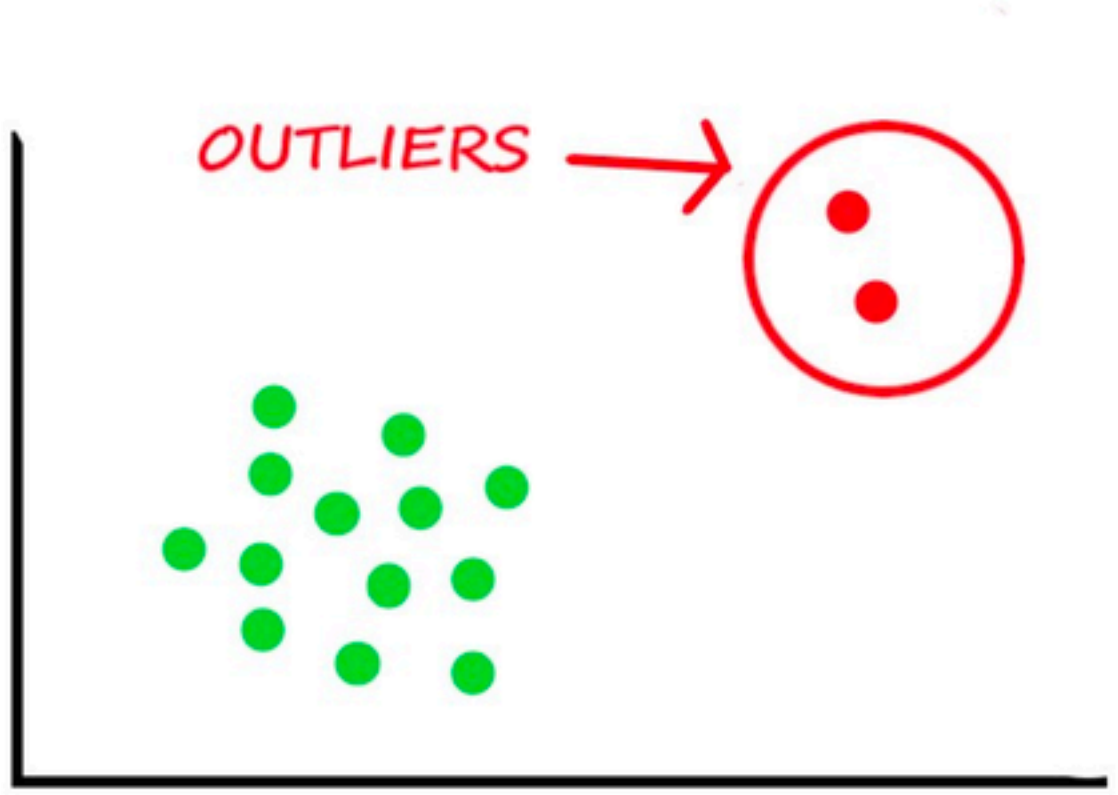

low $p(x)$

Autoencoders

Farina, Nakai & **DS** 1808.08992
Heimel et al 1808.08979
Cerri et al 1811.10276
…

Density estimation

Caron, Hendriks, Verheyen 2106.10164
…

b. Overdensity detection ("group anomalies")

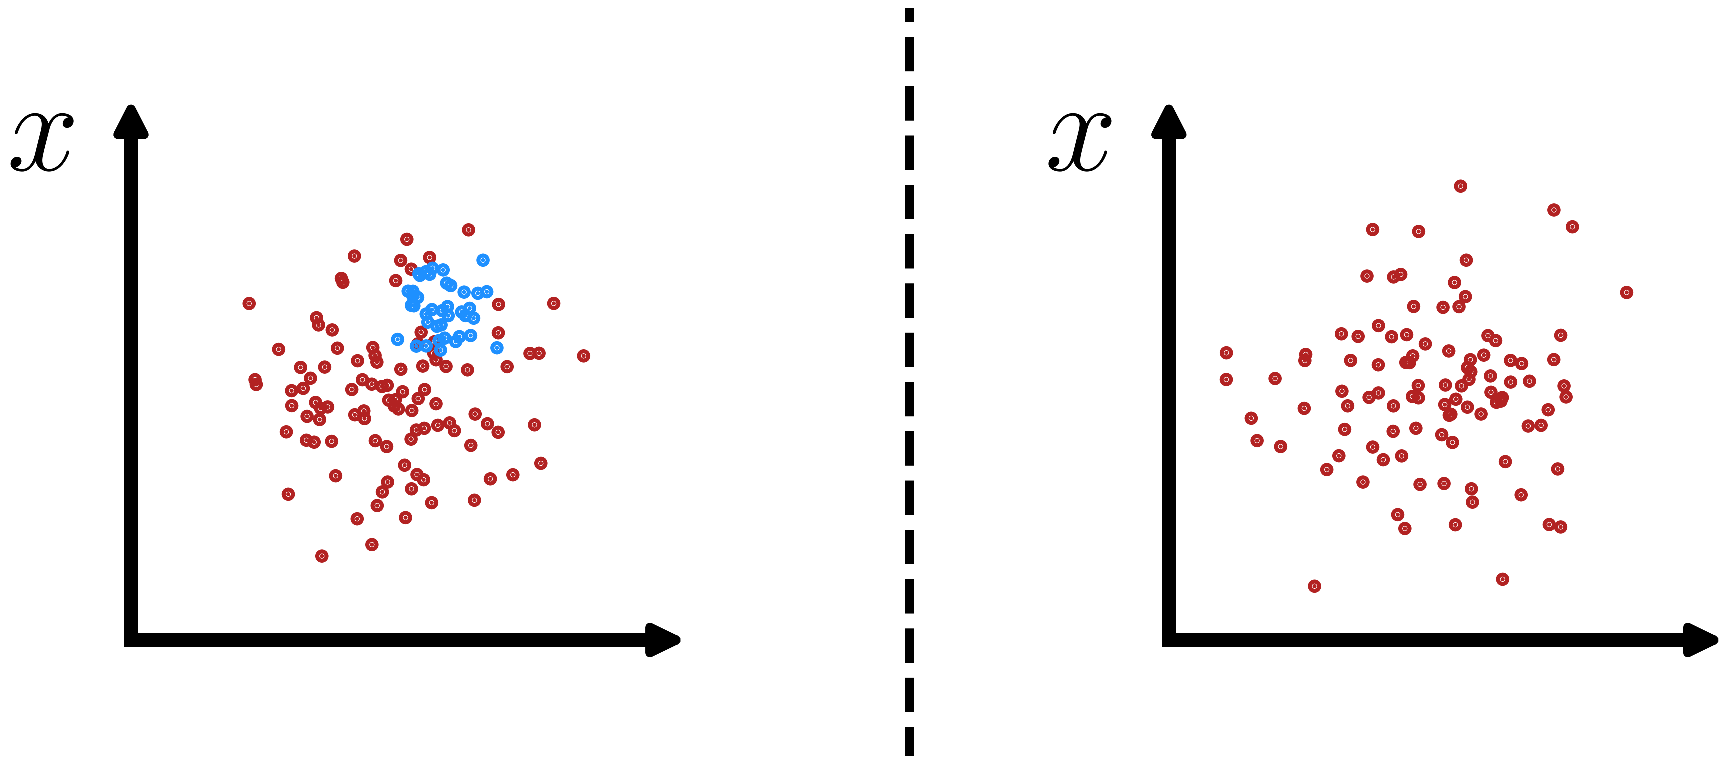

high $\frac{p_{data}(x)}{p_{bg}(x)}$

Data vs bg test statistic

D'Agnolo et al 1806.02350,1912.12155, 2111.13633

Enhanced bump hunts

CWoLa Hunting [Collins, Howe & Nachman 1805.02664, 1902.02634]
ANODE [Nachman & **DS** 2001.04990]
CATHODE [**DS+** Hallin et al 2109.00546, 2210.14924]
CURTAINS [Raine et al 2203.09470]
…

2a. Outlier detection

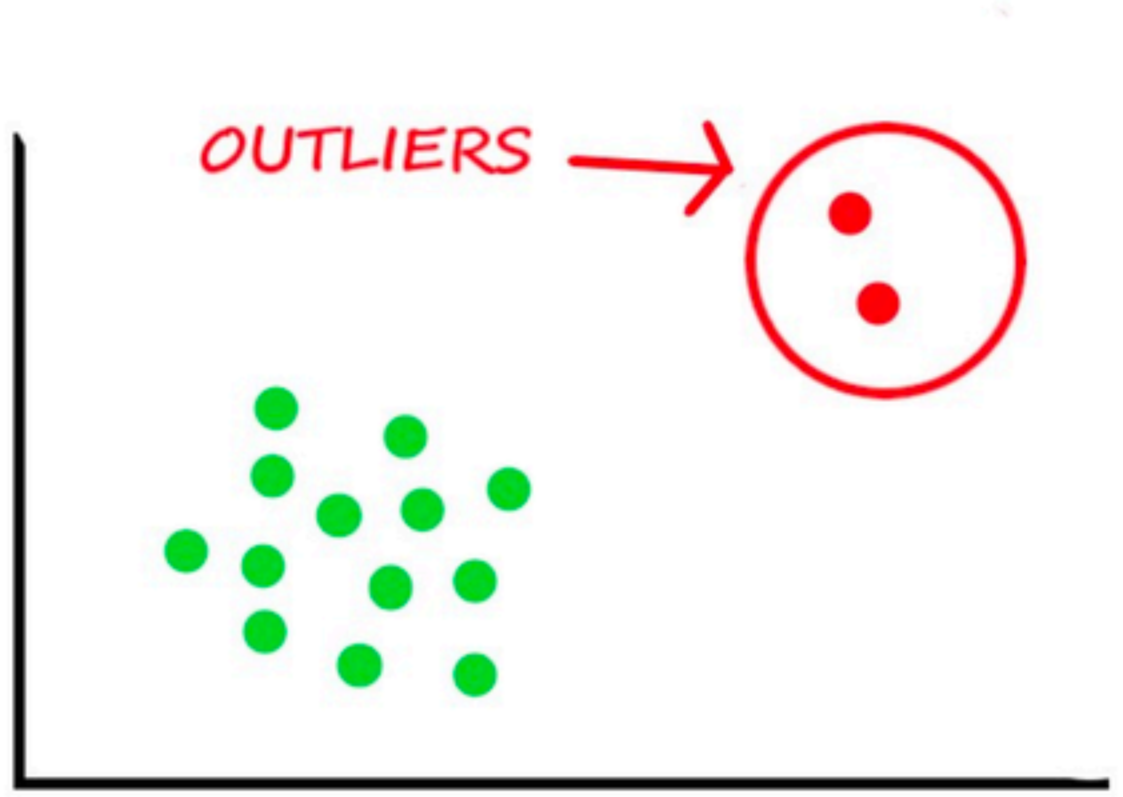

low $p(x)$

- Pros:
  - can be fully unsupervised
  - can potentially find very rare anomalies
- Cons:
  - "low p(x)" is coordinate dependent! An event can be anomalous or not depending on parametrization of features
    [Le Lan & Dinh 2012.03808, **DS**+ Kasieczka et al 2209.06225]

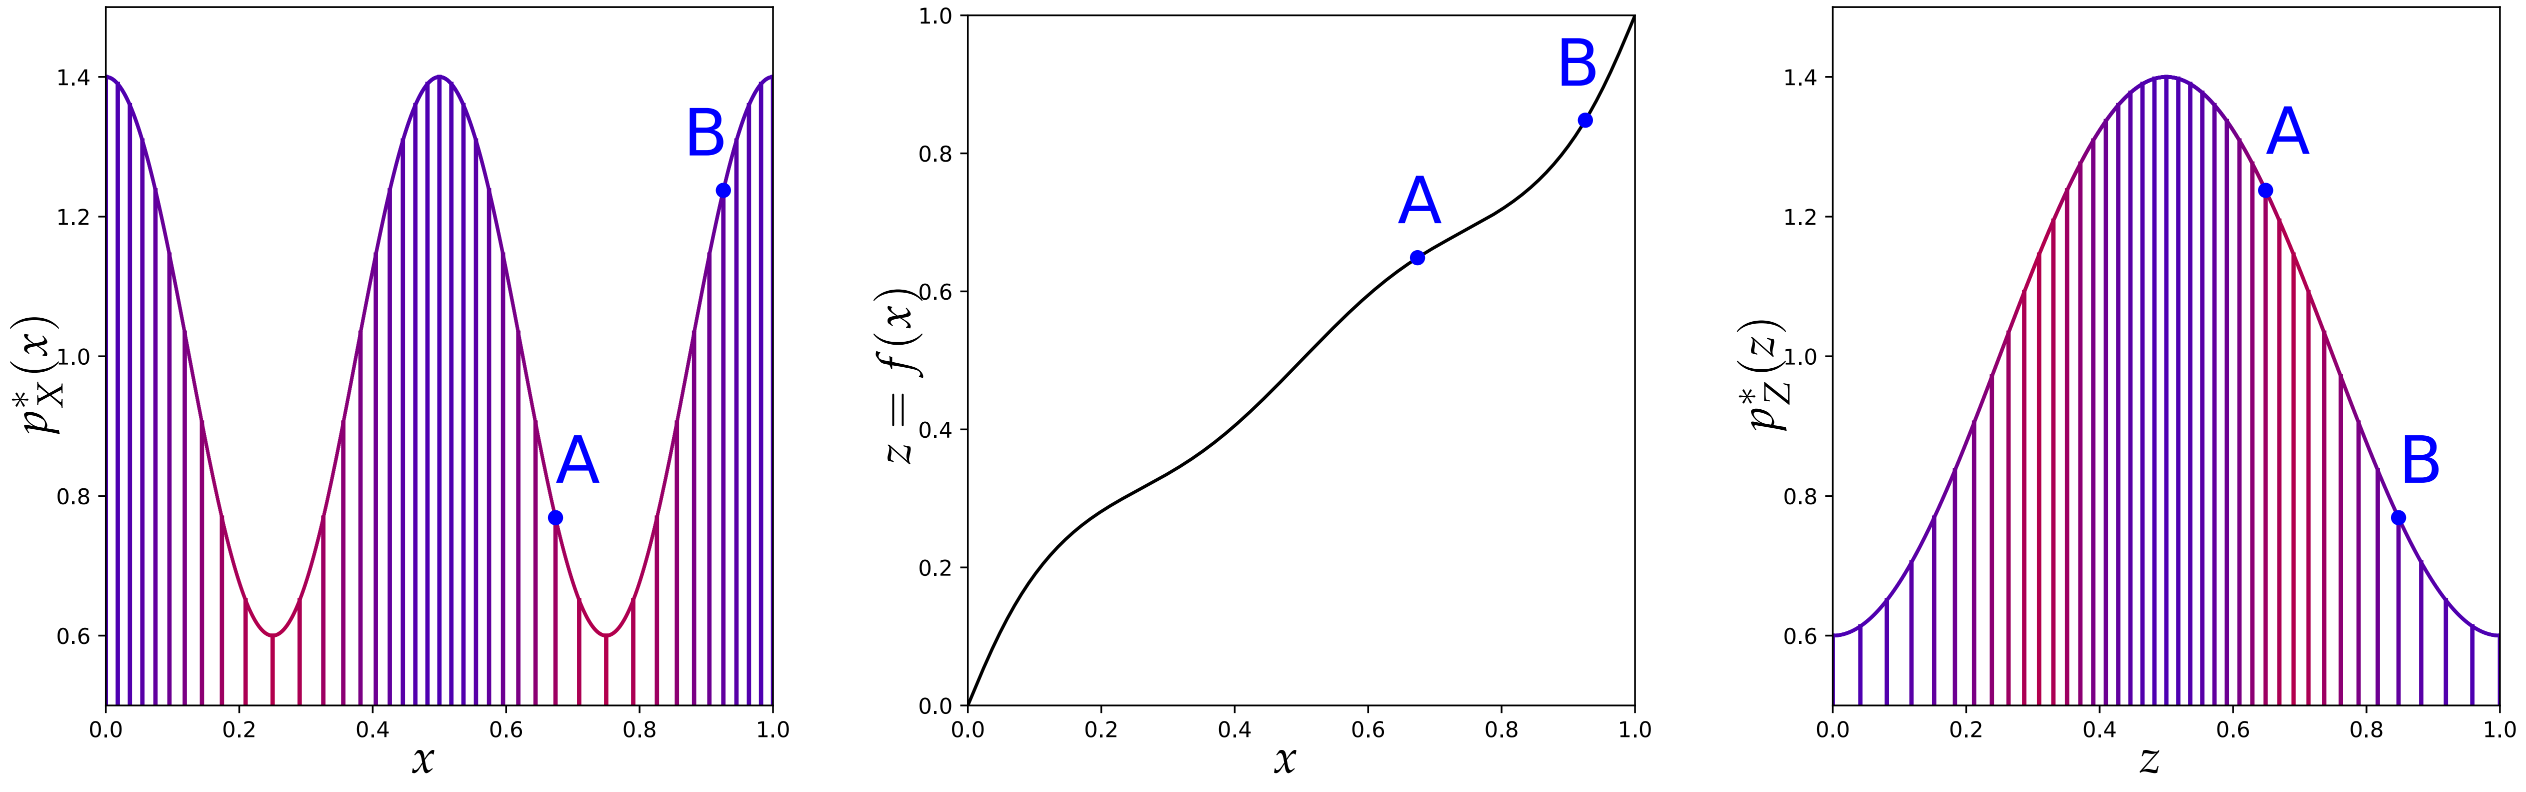

(a) (b) (c)

Example: searching for NP with autoencoders

Farina, Nakai & **DS** 1808.08992; Heimel, Kasieczka, Plehn & Thompson 1808.08979; Cerri et al 1811.10276; and many more…

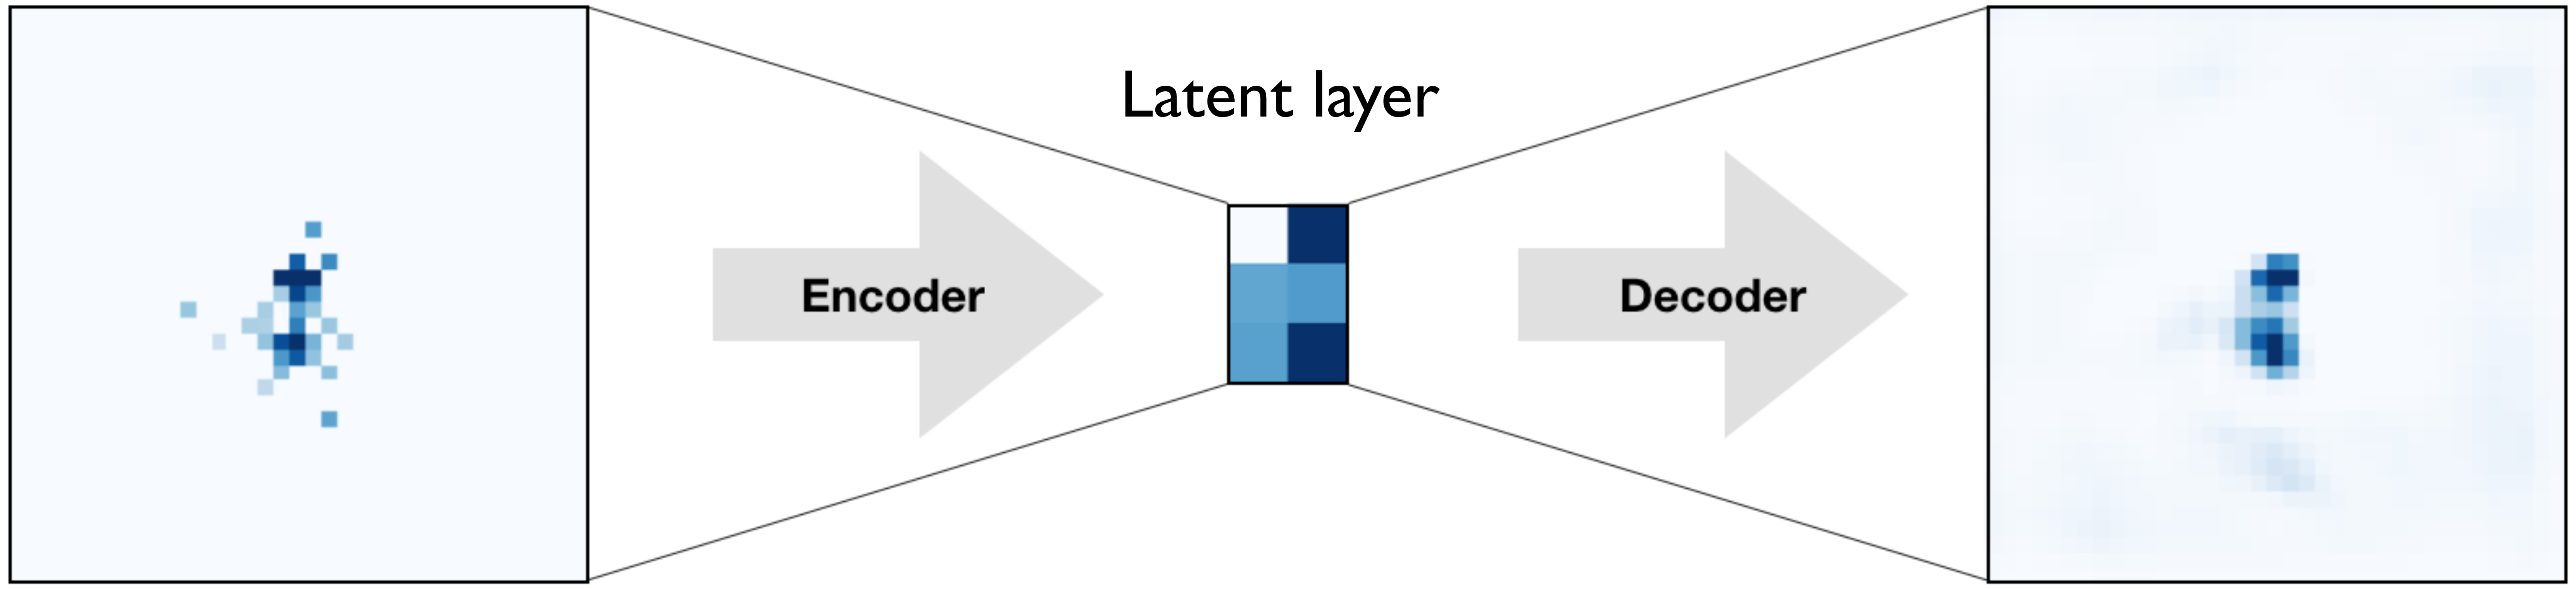

An autoencoder maps an input into a "latent representation" and then attempts to reconstruct the original input from it.

The encoding is lossy, so the reconstruction is not perfect.

Many real world applications of autoencoders, including anomaly detection, fraud detection, denoising, compression, generation, density estimation

Example: searching for NP with autoencoders

Farina, Nakai & **DS** 1808.08992; Heimel, Kasieczka, Plehn & Thompson 1808.08979; Cerri et al 1811.10276; and many more…

Loss function for autoencoder: $$L = \frac{1}{N}\sum_{i=1}^{N}(x_i^{in} - x_i^{out})^2$$ "reconstruction error"

Example: searching for NP with autoencoders

Farina, Nakai & **DS** 1808.08992; Heimel, Kasieczka, Plehn & Thompson 1808.08979; Cerri et al 1811.10276; and many more…

Loss function for autoencoder: $$L = \frac{1}{N}\sum_{i=1}^{N}(x_i^{in} - x_i^{out})^2$$ "reconstruction error"

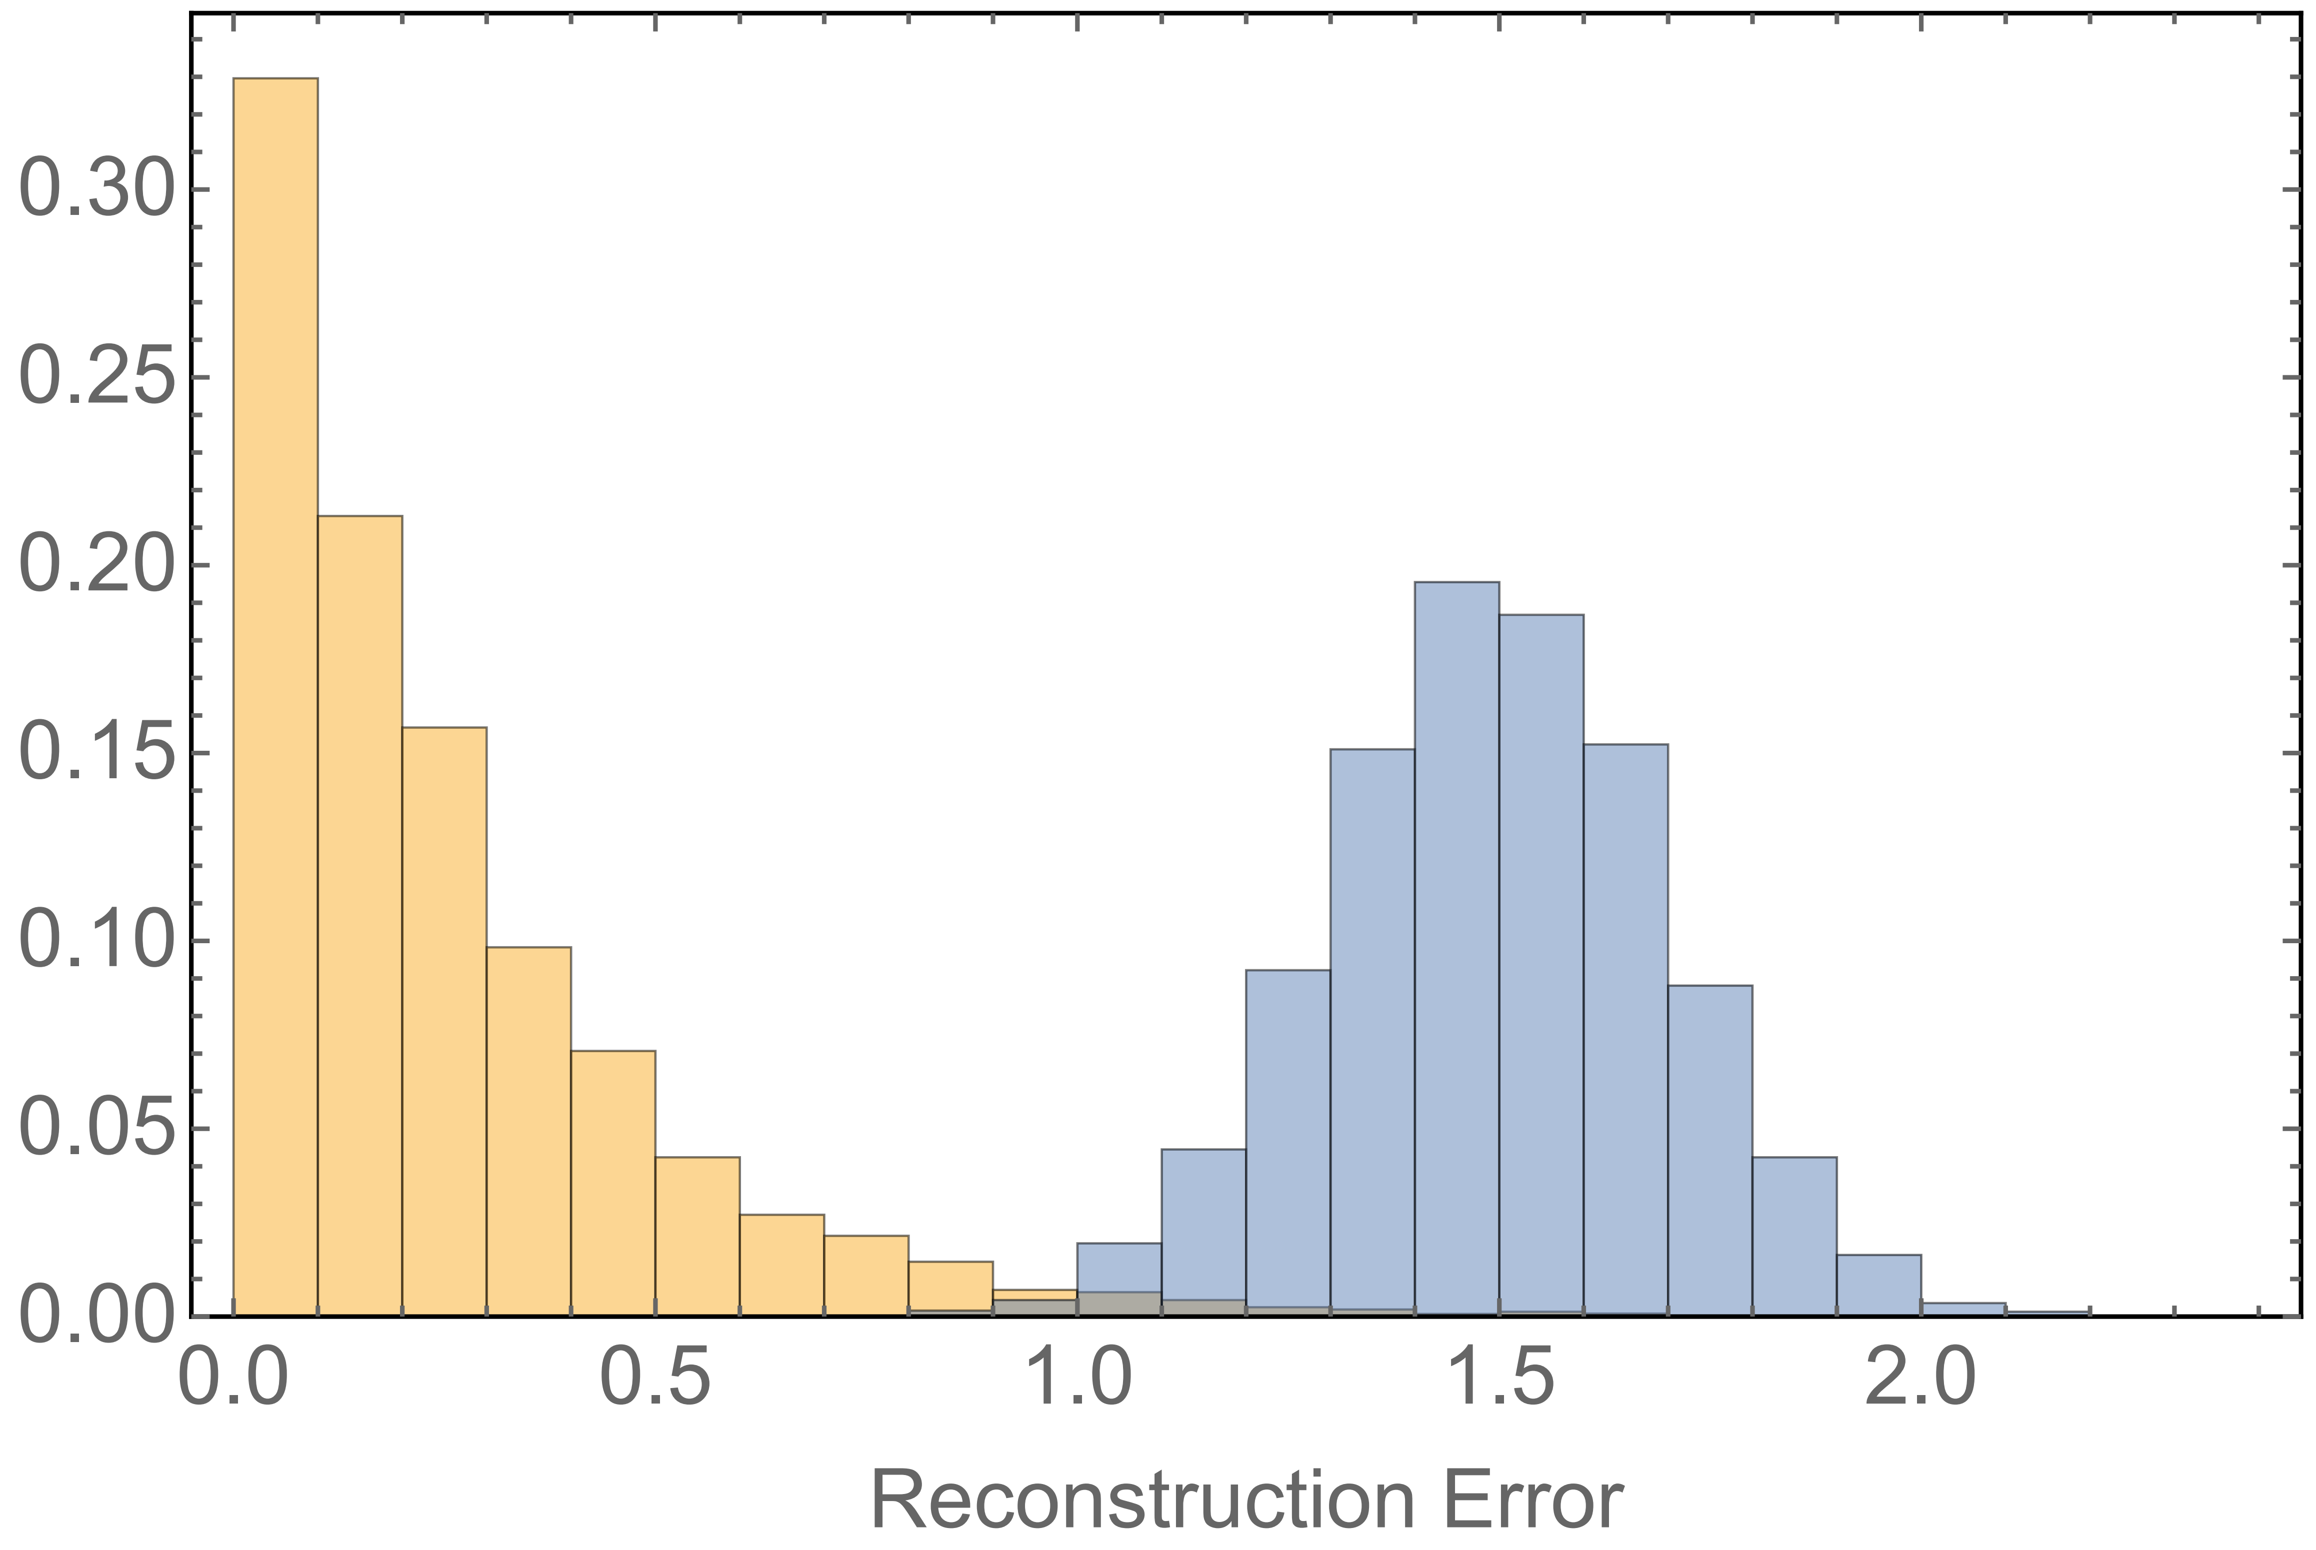

Example: searching for NP with autoencoders

Farina, Nakai & **DS** 1808.08992; Heimel, Kasieczka, Plehn & Thompson 1808.08979; Cerri et al 1811.10276; and many more…

Loss function for autoencoder: $$L = \frac{1}{N}\sum_{i=1}^{N}(x_i^{in} - x_i^{out})^2$$ "reconstruction error"

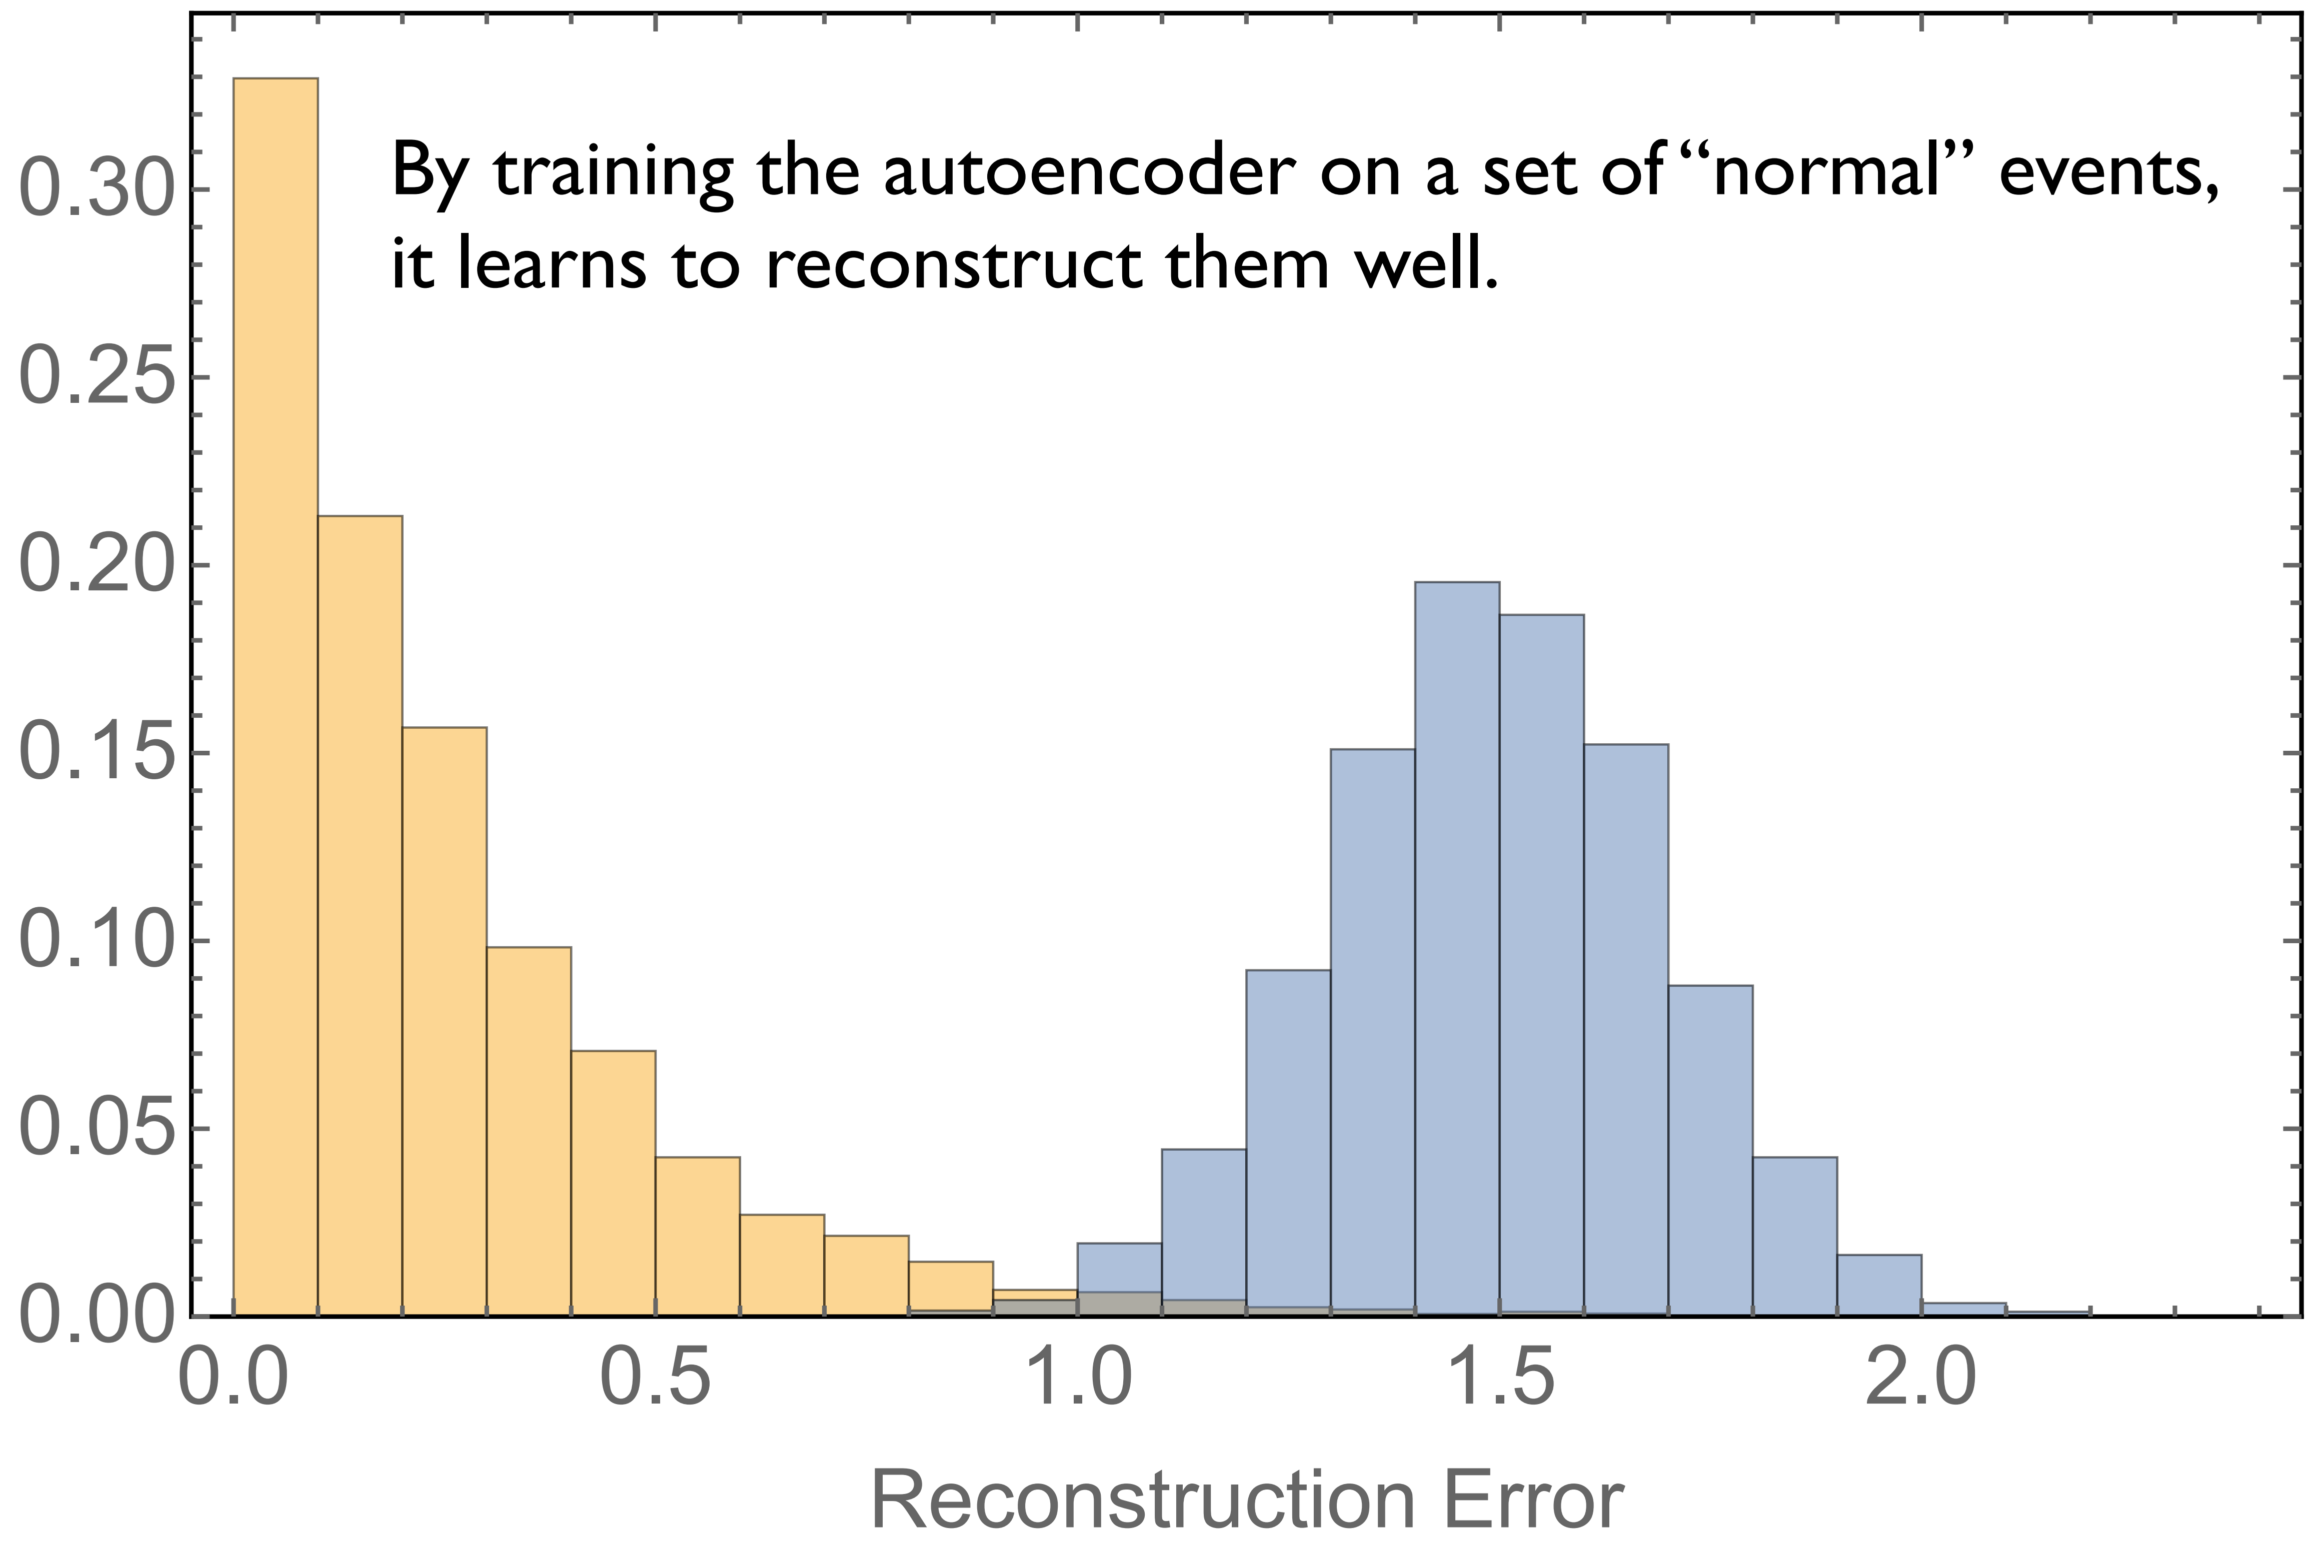

Example: searching for NP with autoencoders

Farina, Nakai & **DS** 1808.08992; Heimel, Kasieczka, Plehn & Thompson 1808.08979; Cerri et al 1811.10276; and many more…

Loss function for autoencoder: $$L = \frac{1}{N}\sum_{i=1}^{N}(x_i^{in} - x_i^{out})^2$$ "reconstruction error"

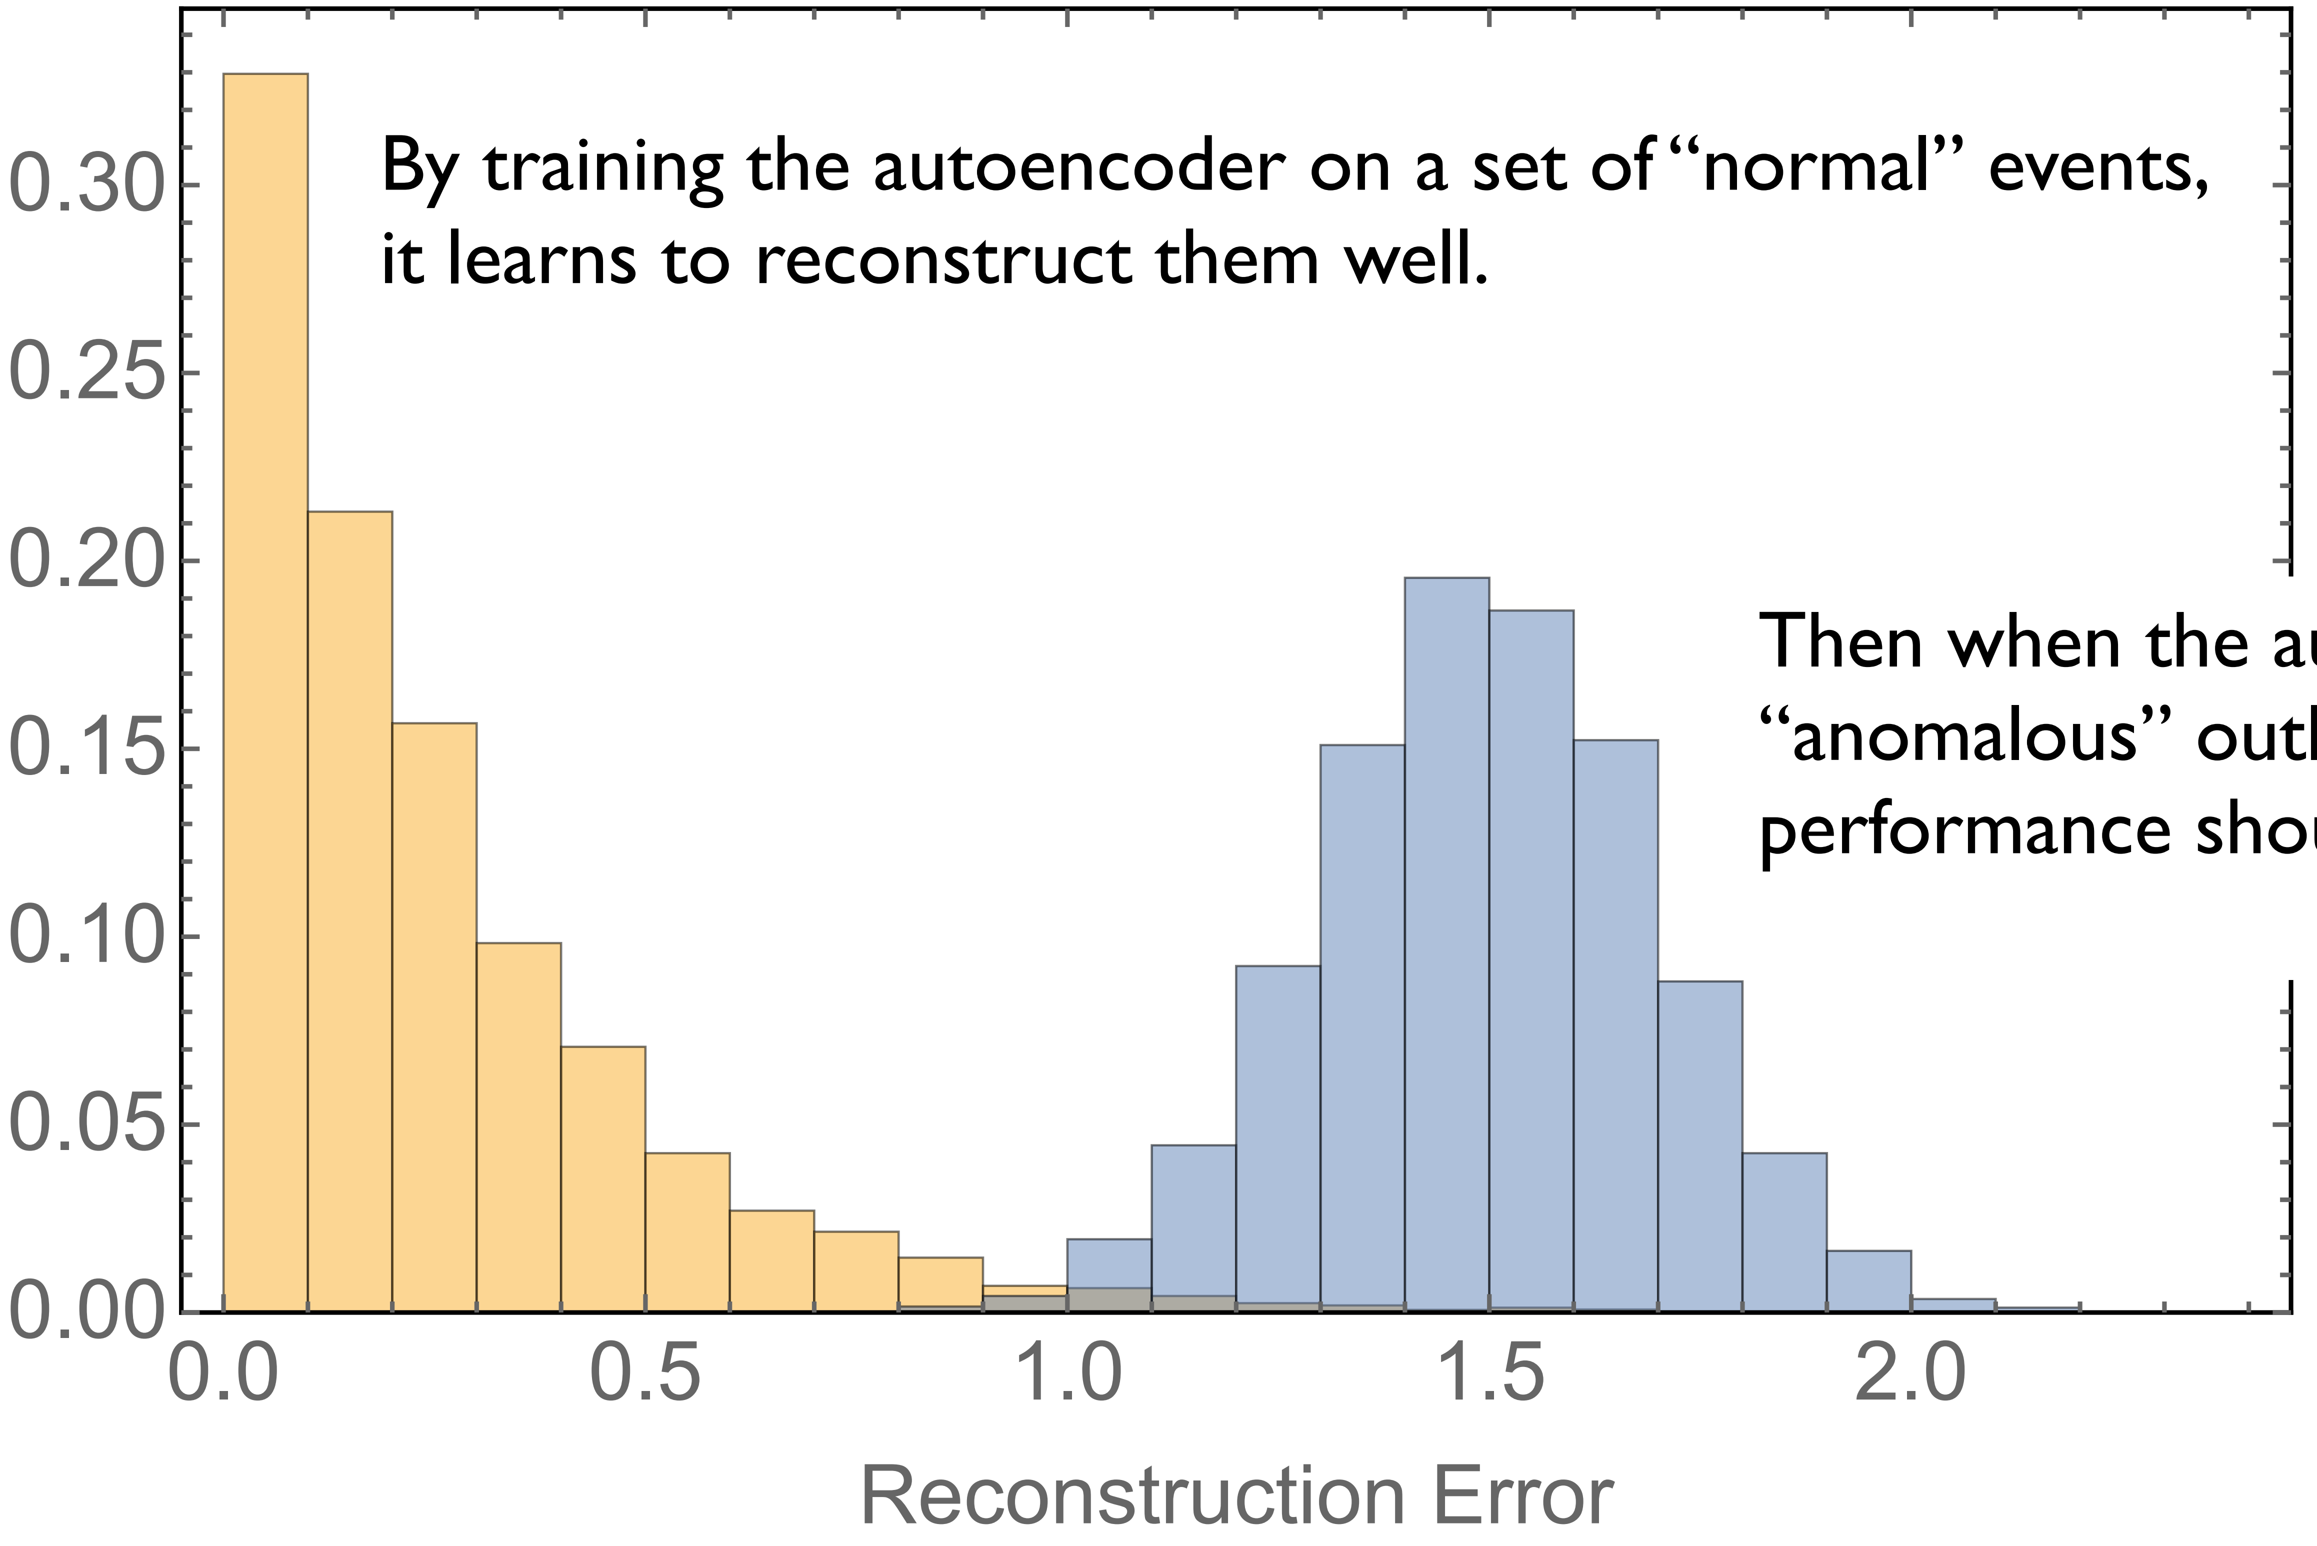

By training the autoencoder on a set of "normal" events, it learns to reconstruct them well.

Then when the autoencoder encounters "anomalous" outlier events, its performance should be worse.

Example: searching for NP with autoencoders

Farina, Nakai & **DS** 1808.08992; Heimel, Kasieczka, Plehn & Thompson 1808.08979; Cerri et al 1811.10276; and many more…

Loss function for autoencoder: $$L = \frac{1}{N}\sum_{i=1}^{N}(x_i^{in} - x_i^{out})^2$$ "reconstruction error"

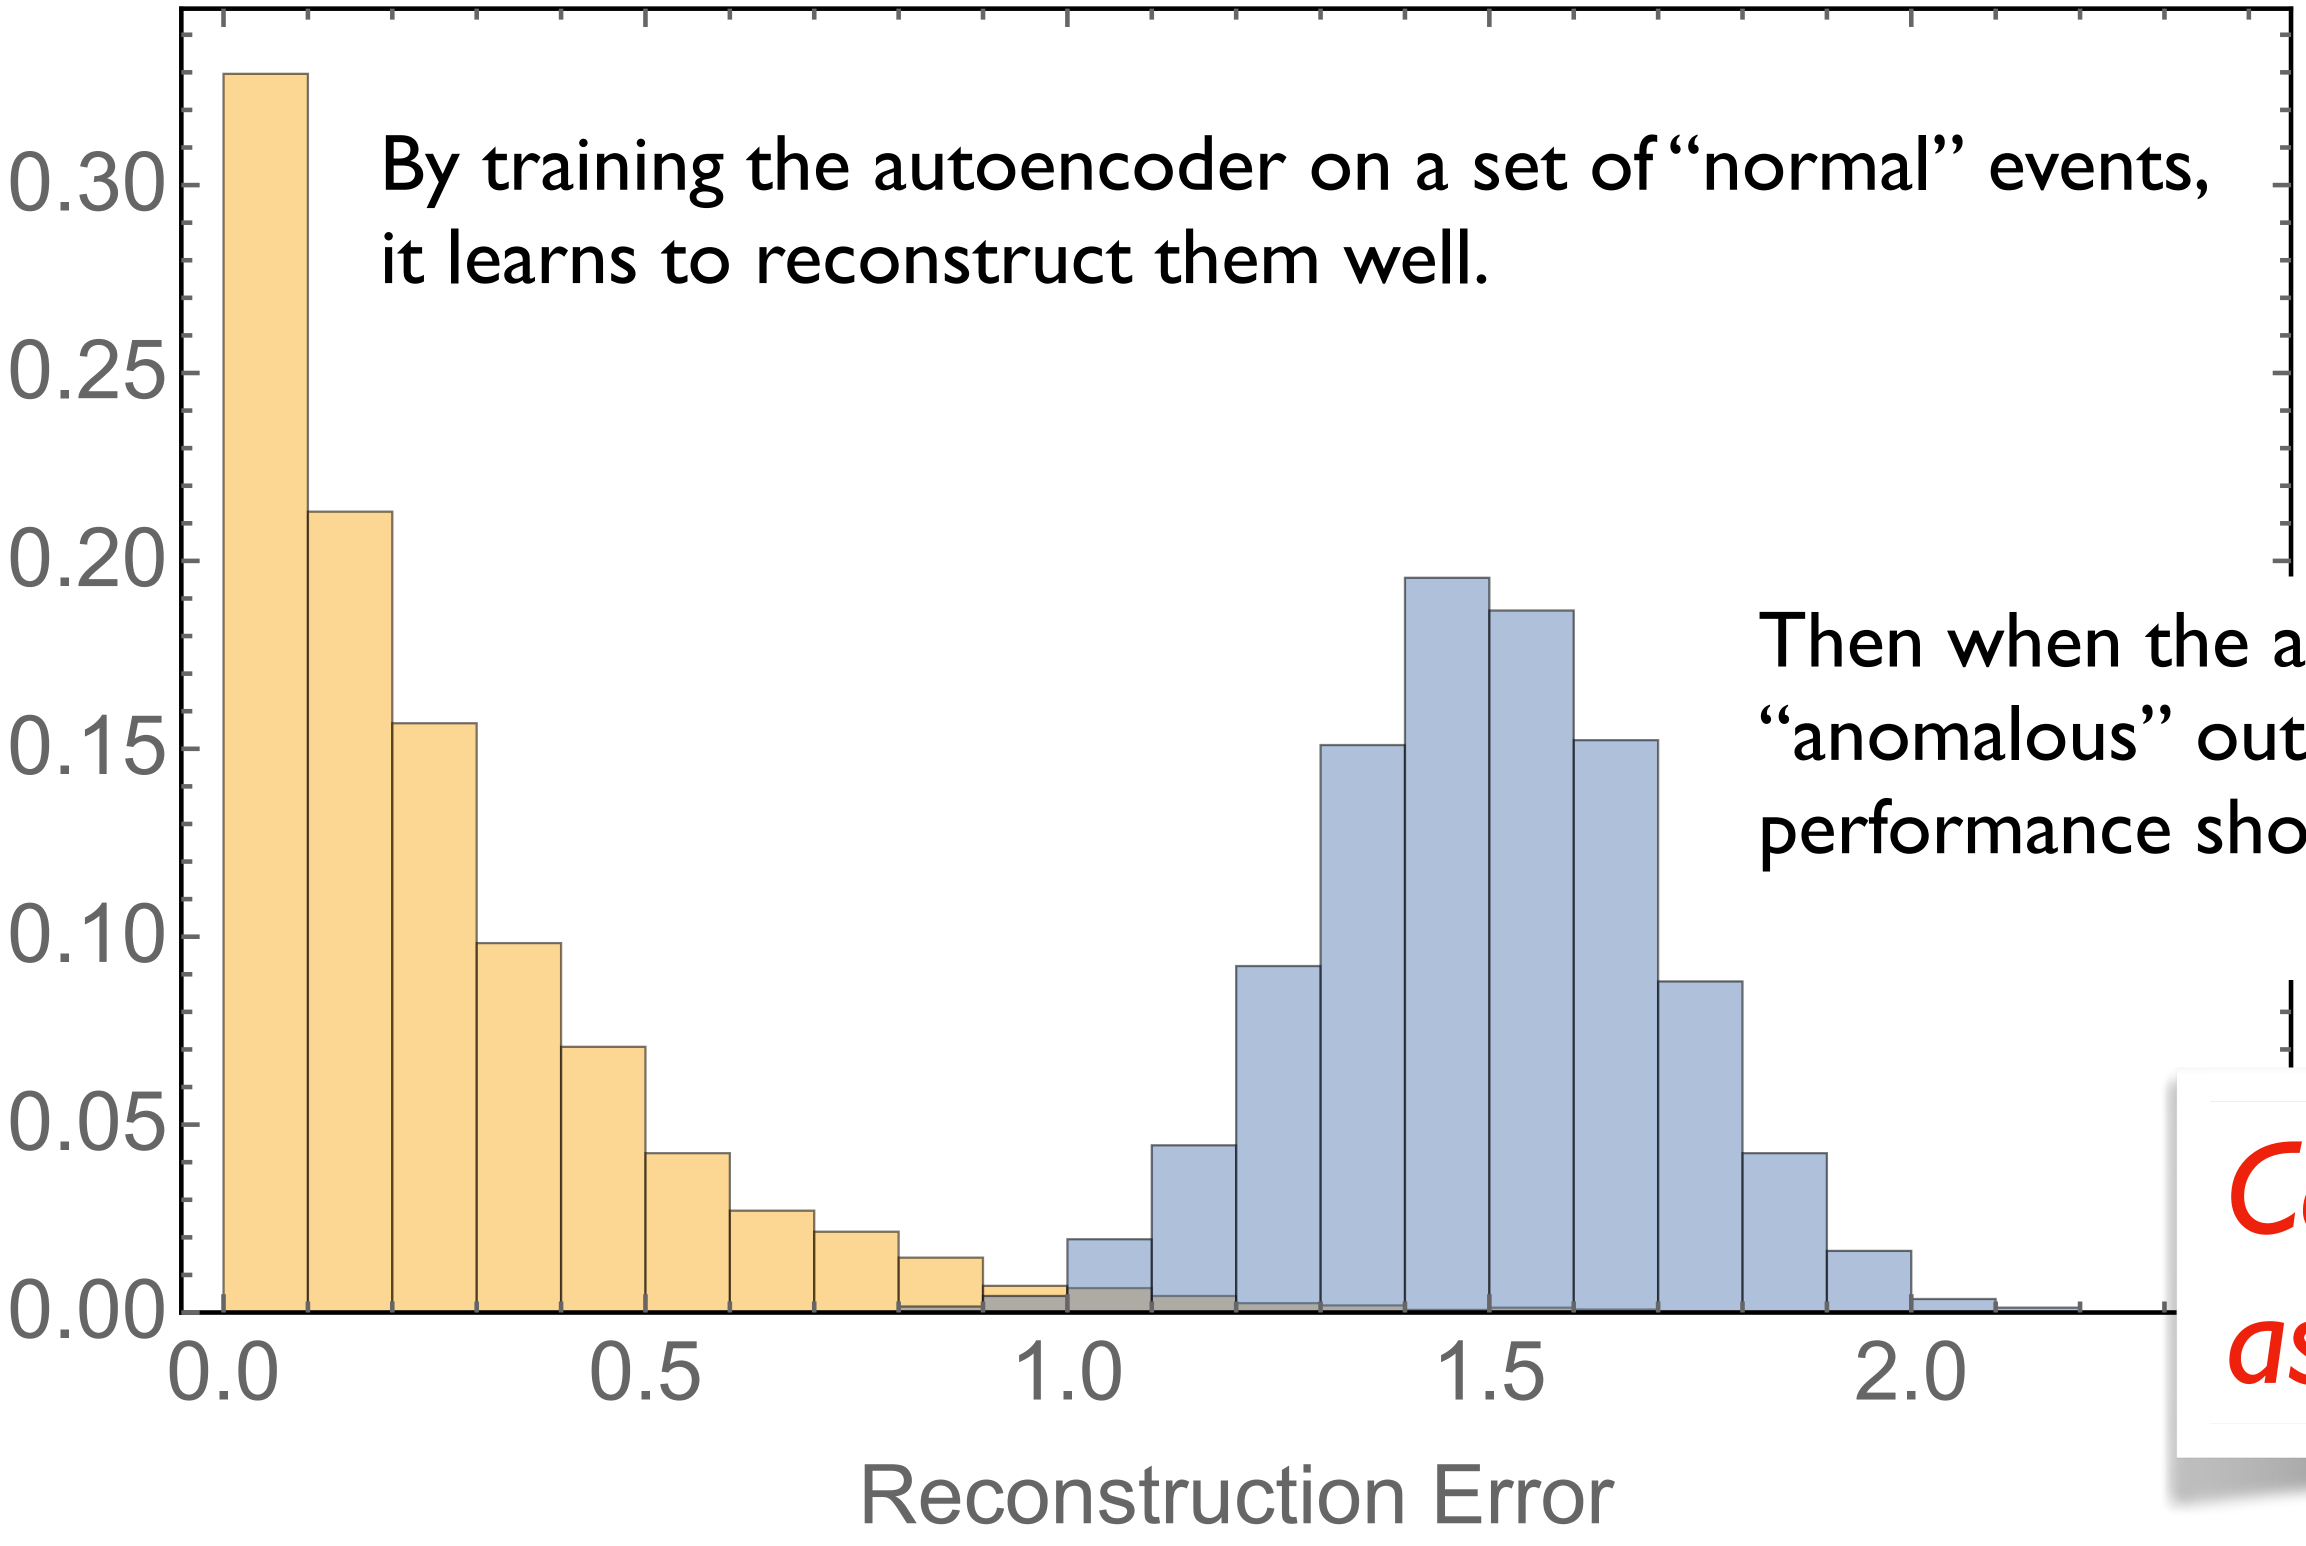

By training the autoencoder on a set of "normal" events, it learns to reconstruct them well.

Then when the autoencoder encounters "anomalous" outlier events, its performance should be worse.

Can use reconstruction error as an anomaly score!

Example: searching for NP with autoencoders

Farina, Nakai & **DS** 1808.08992; Heimel, Kasieczka, Plehn & Thompson 1808.08979; Cerri et al 1811.10276; and many more…

Loss function for autoencoder: $$L = \frac{1}{N}\sum_{i=1}^{N}(x_i^{in} - x_i^{out})^2$$ “reconstruction error”

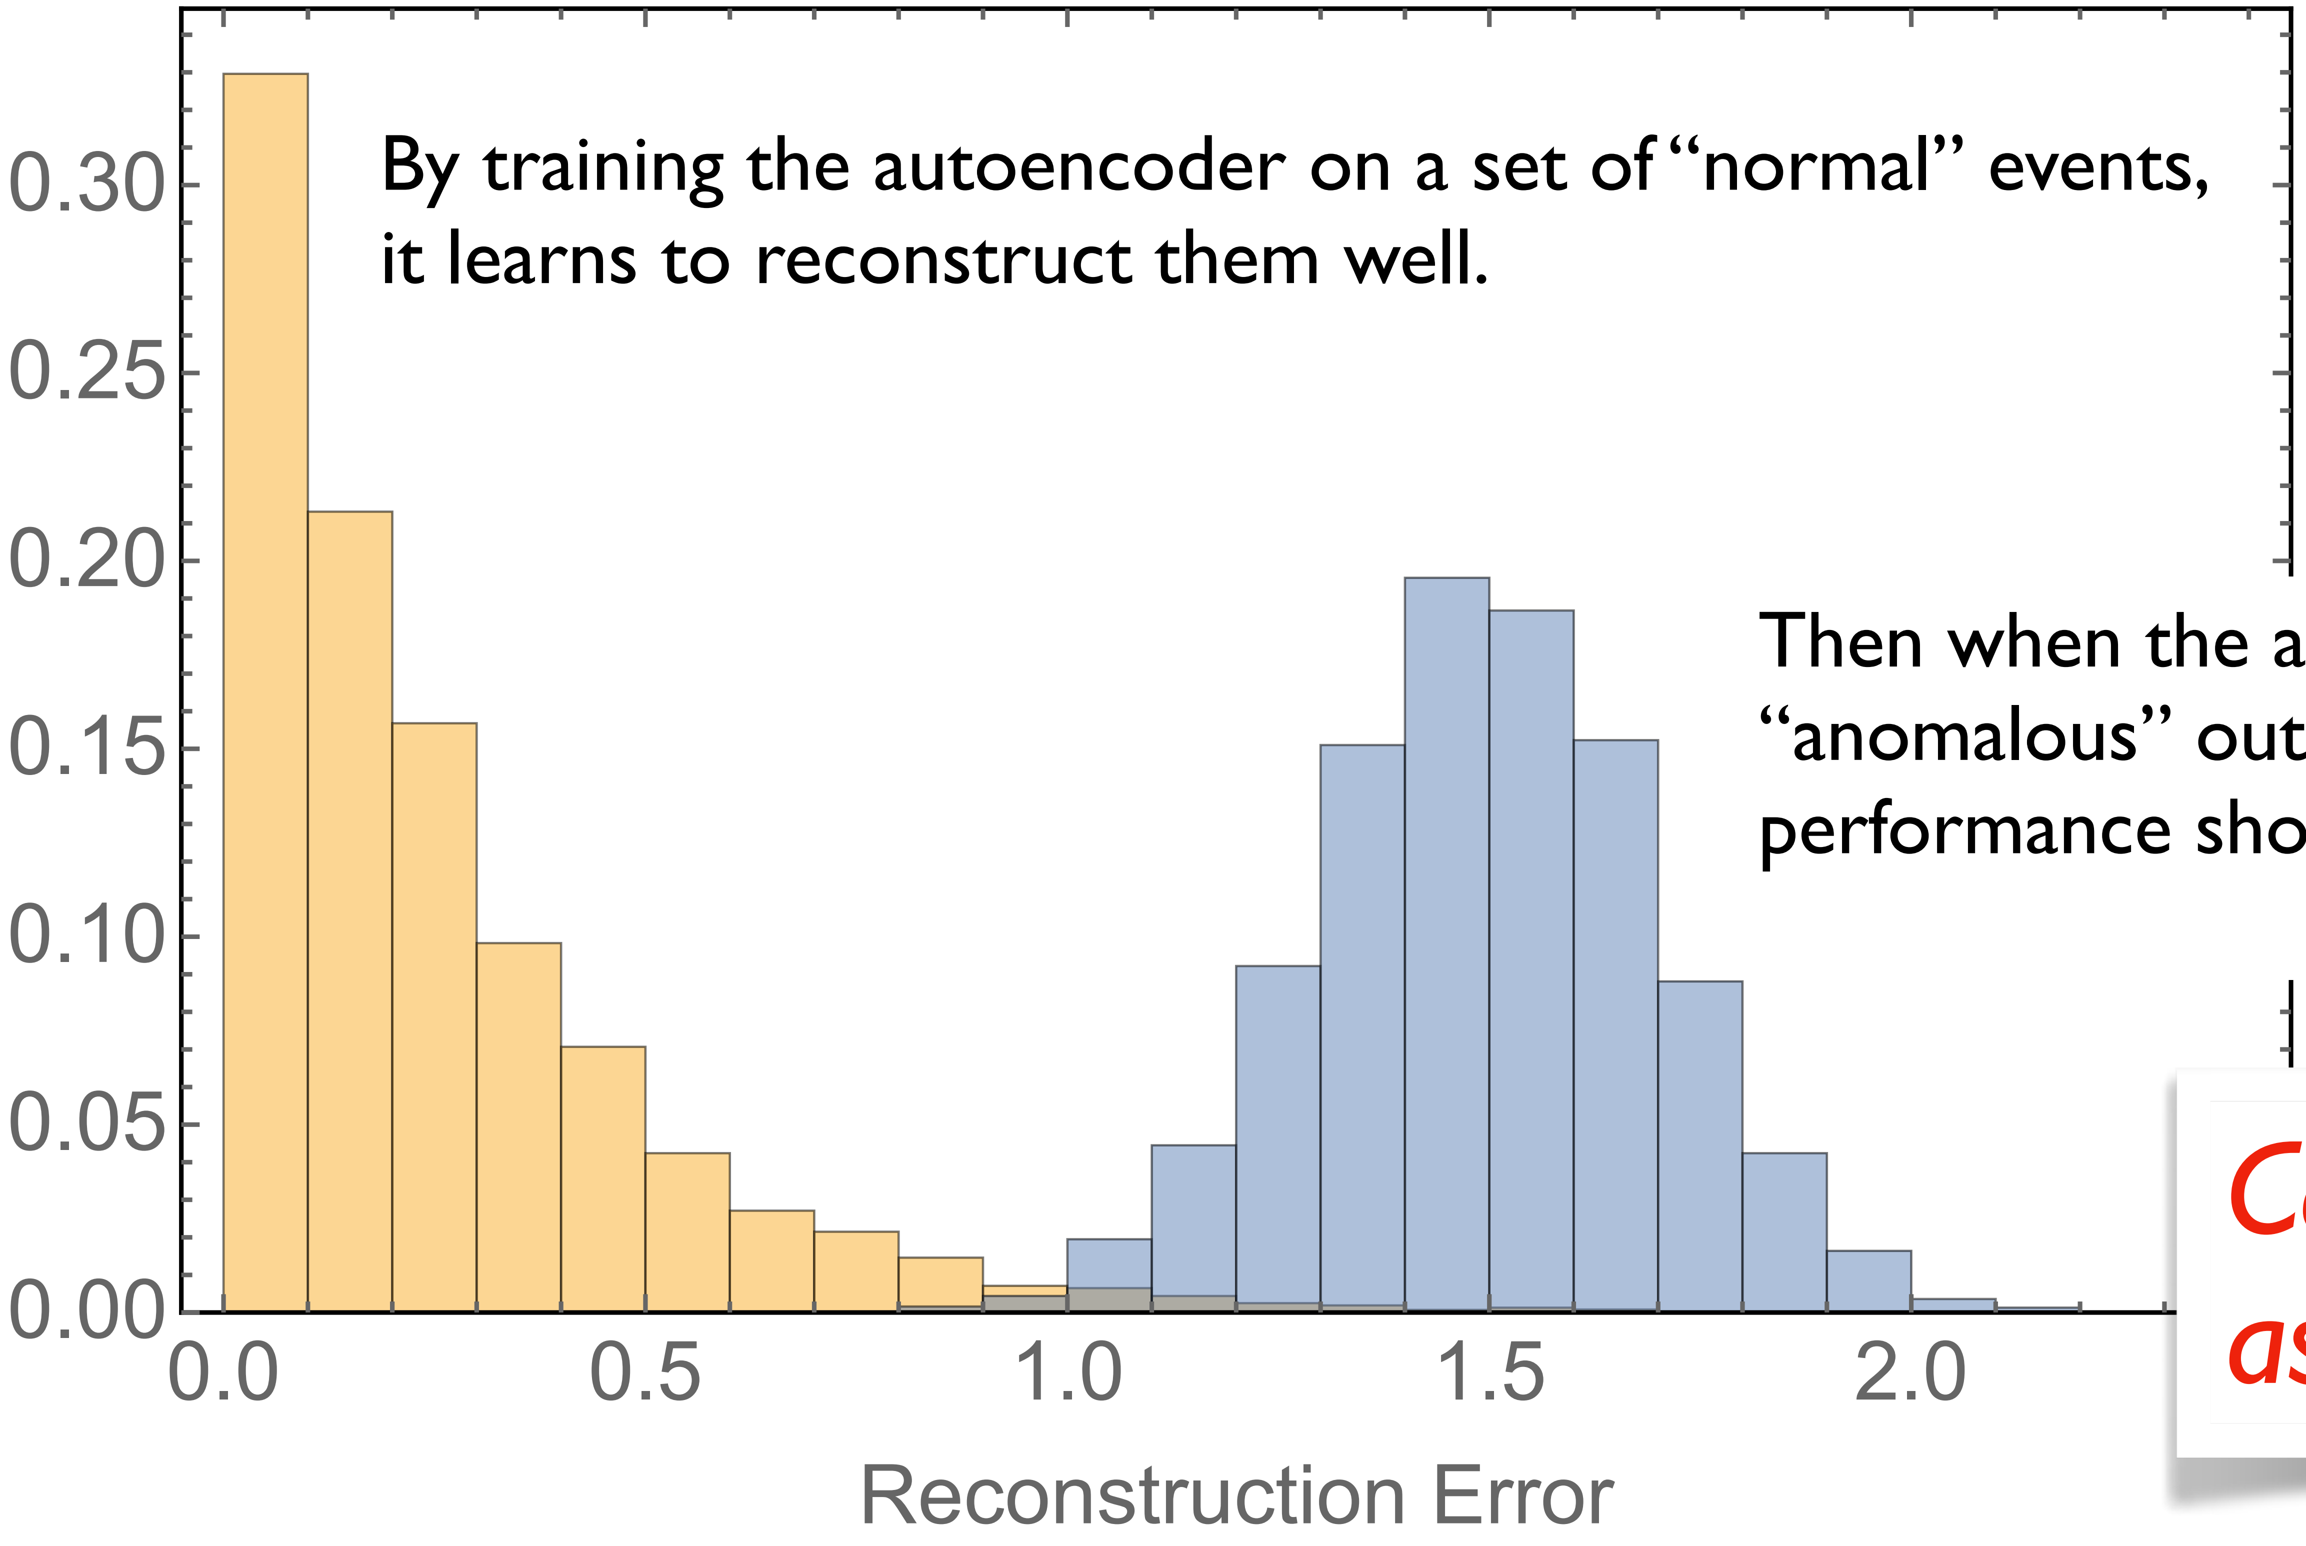

By training the autoencoder on a set of “normal” events, it learns to reconstruct them well.

Then when the autoencoder encounters “anomalous” outlier events, its performance should be worse.

*[See **Maria’s talk** next for other ways to use (variational) autoencoders for anomaly detection]*

Can use reconstruction error as an anomaly score!

Example: searching for NP with autoencoders

Farina, Nakai & **DS** 1808.08992; Heimel, Kasieczka, Plehn & Thompson 1808.08979; Cerri et al 1811.10276; and many more…

from 1808.08992

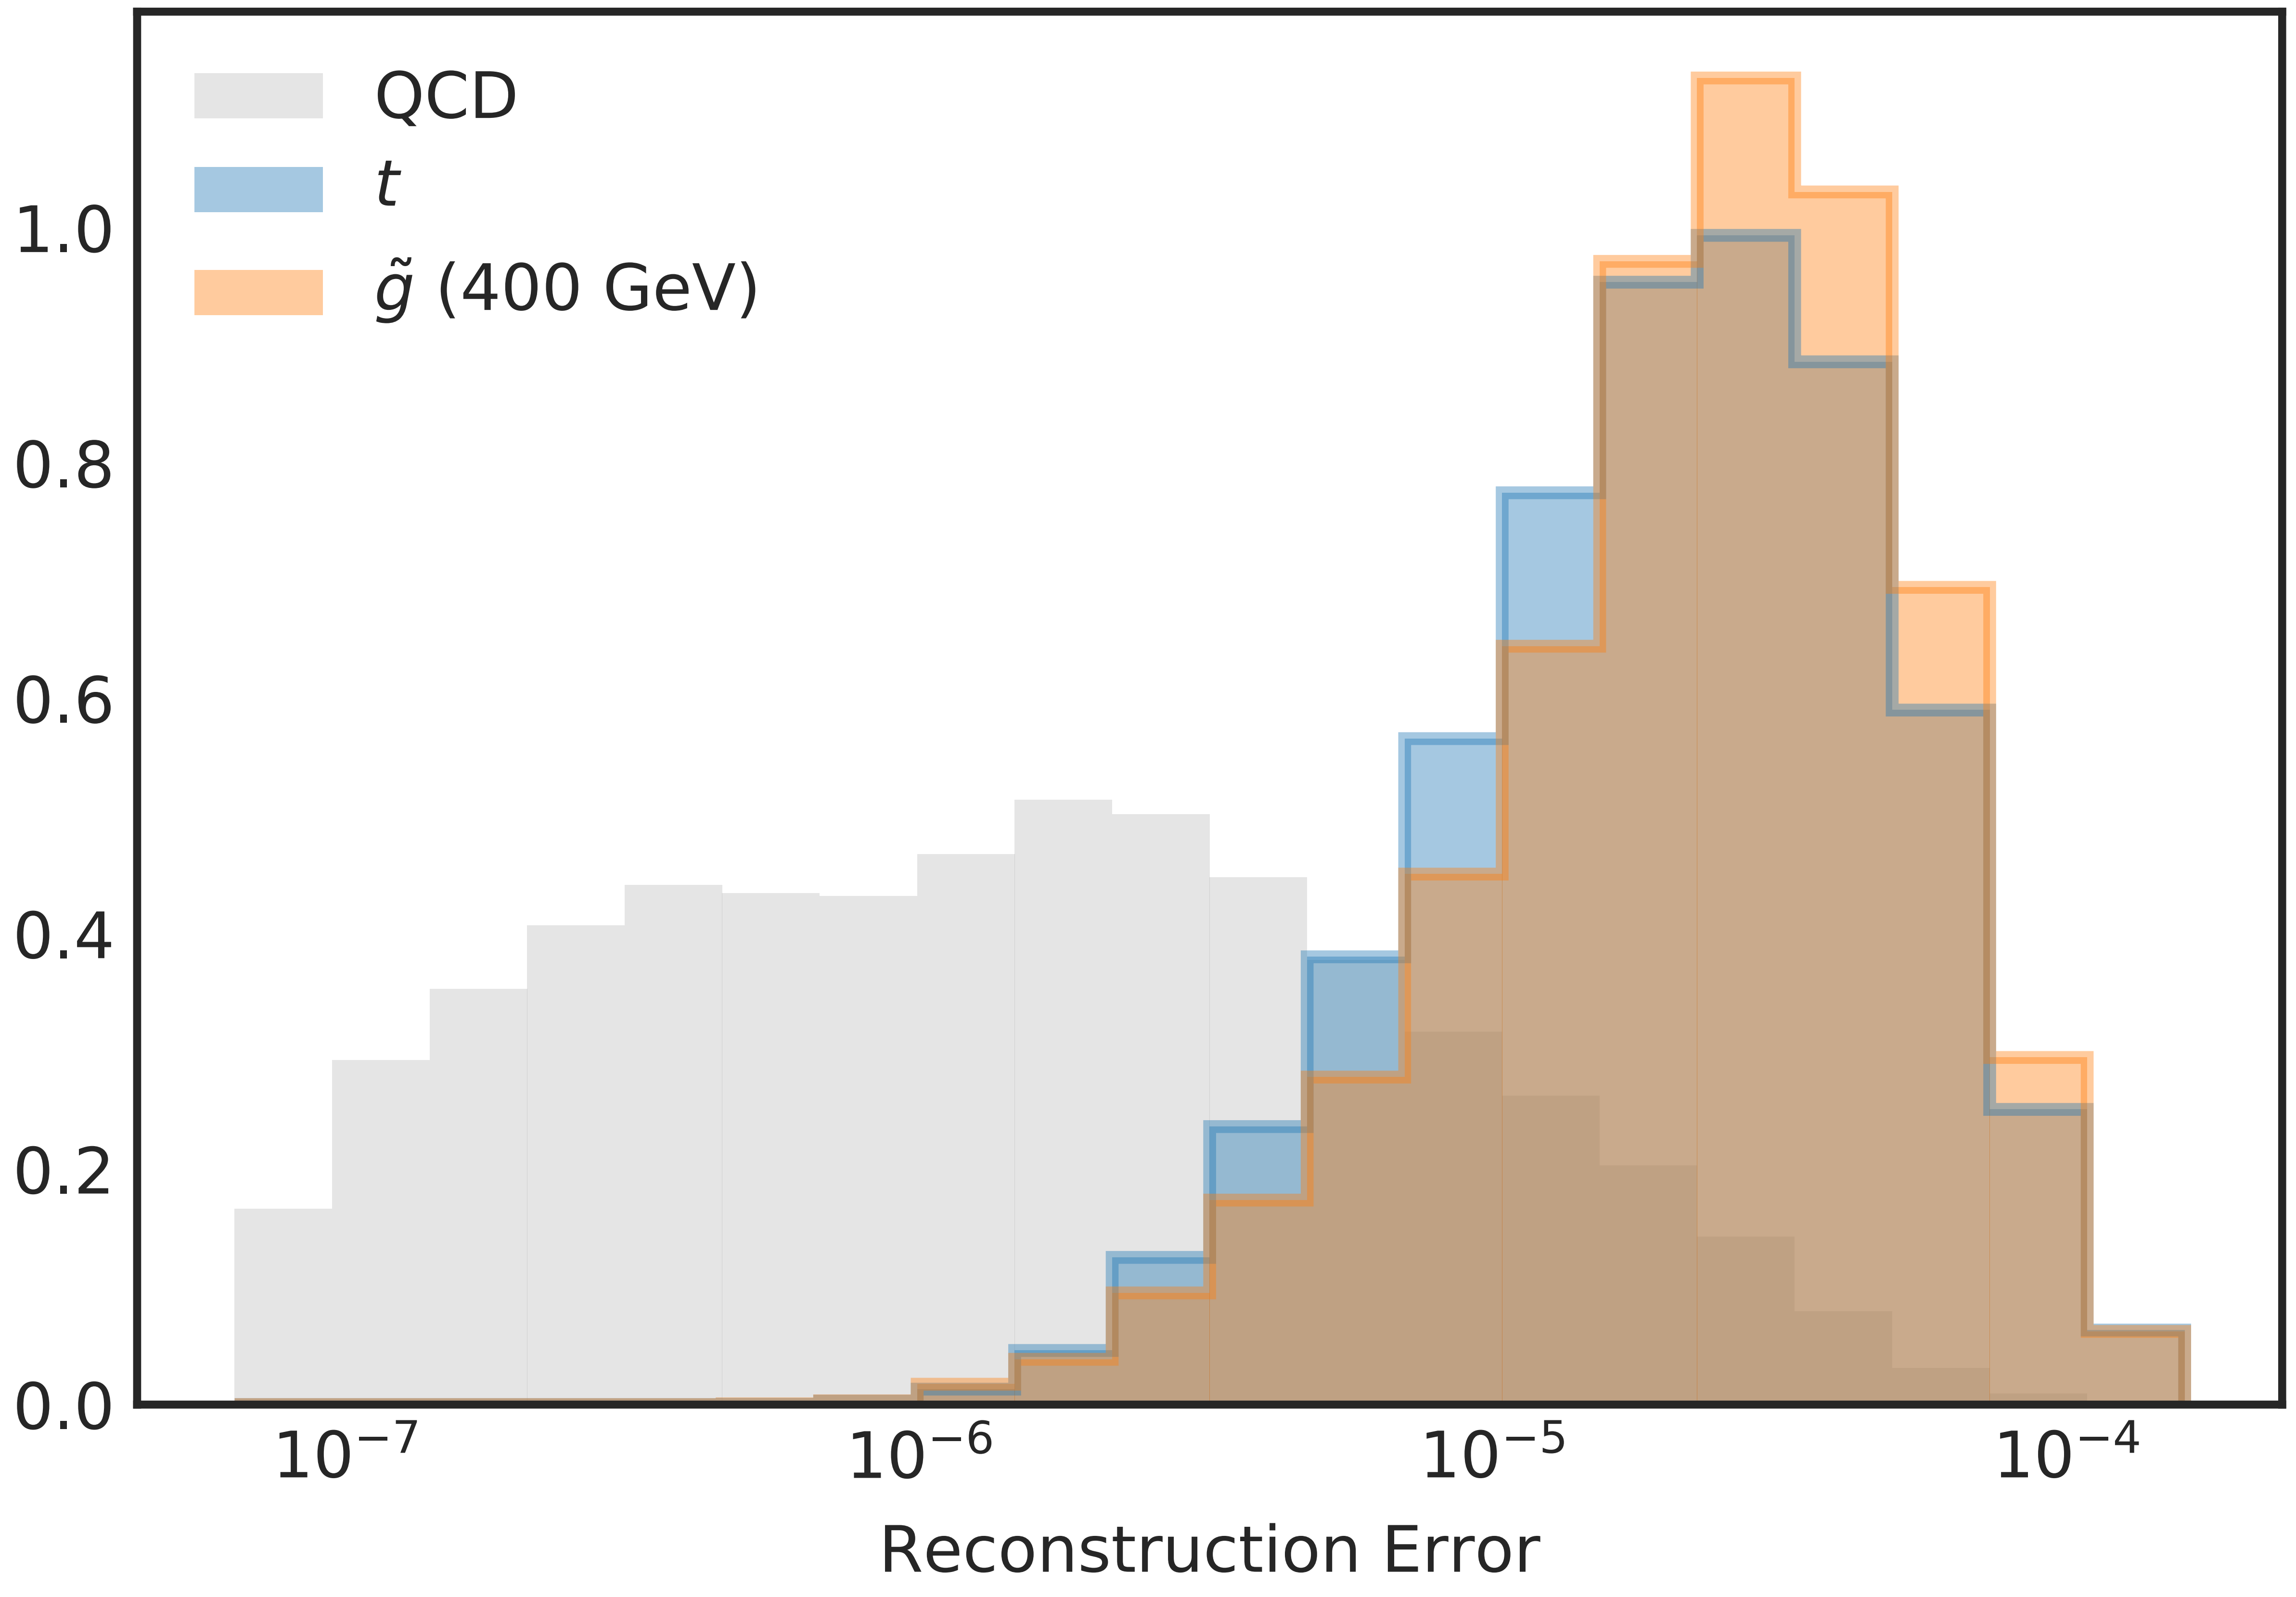

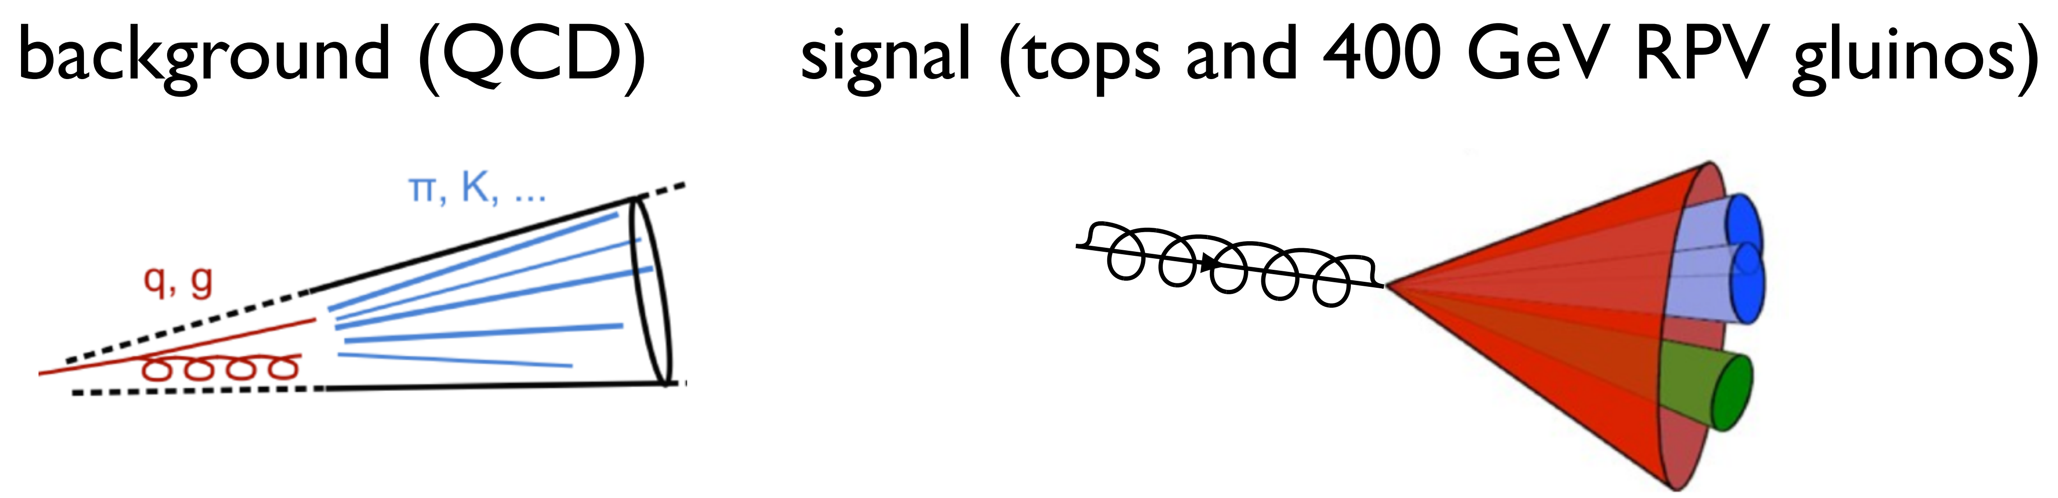

We showed that the AE could detect interesting physics anomalies.

Train the AE on QCD jets only.
Can detect top and gluino jets as anomalous!

Autoencoders: challenges

Autoencoders: challenges

Challenges:

Autoencoders: challenges

Challenges:

- Uncontrolled, not very sensitive, optimality not guaranteed — the AE will find what it finds…

Autoencoders: challenges

Challenges:

- Uncontrolled, not very sensitive, optimality not guaranteed — the AE will find what it finds…
- The AE can fail to detect outliers if they are “simpler” than the background
  T. Weber MSc Thesis [G. Kasieczka]; Dillon et al 2104.08291; Finke et al 2104.09051

from 2104.09051

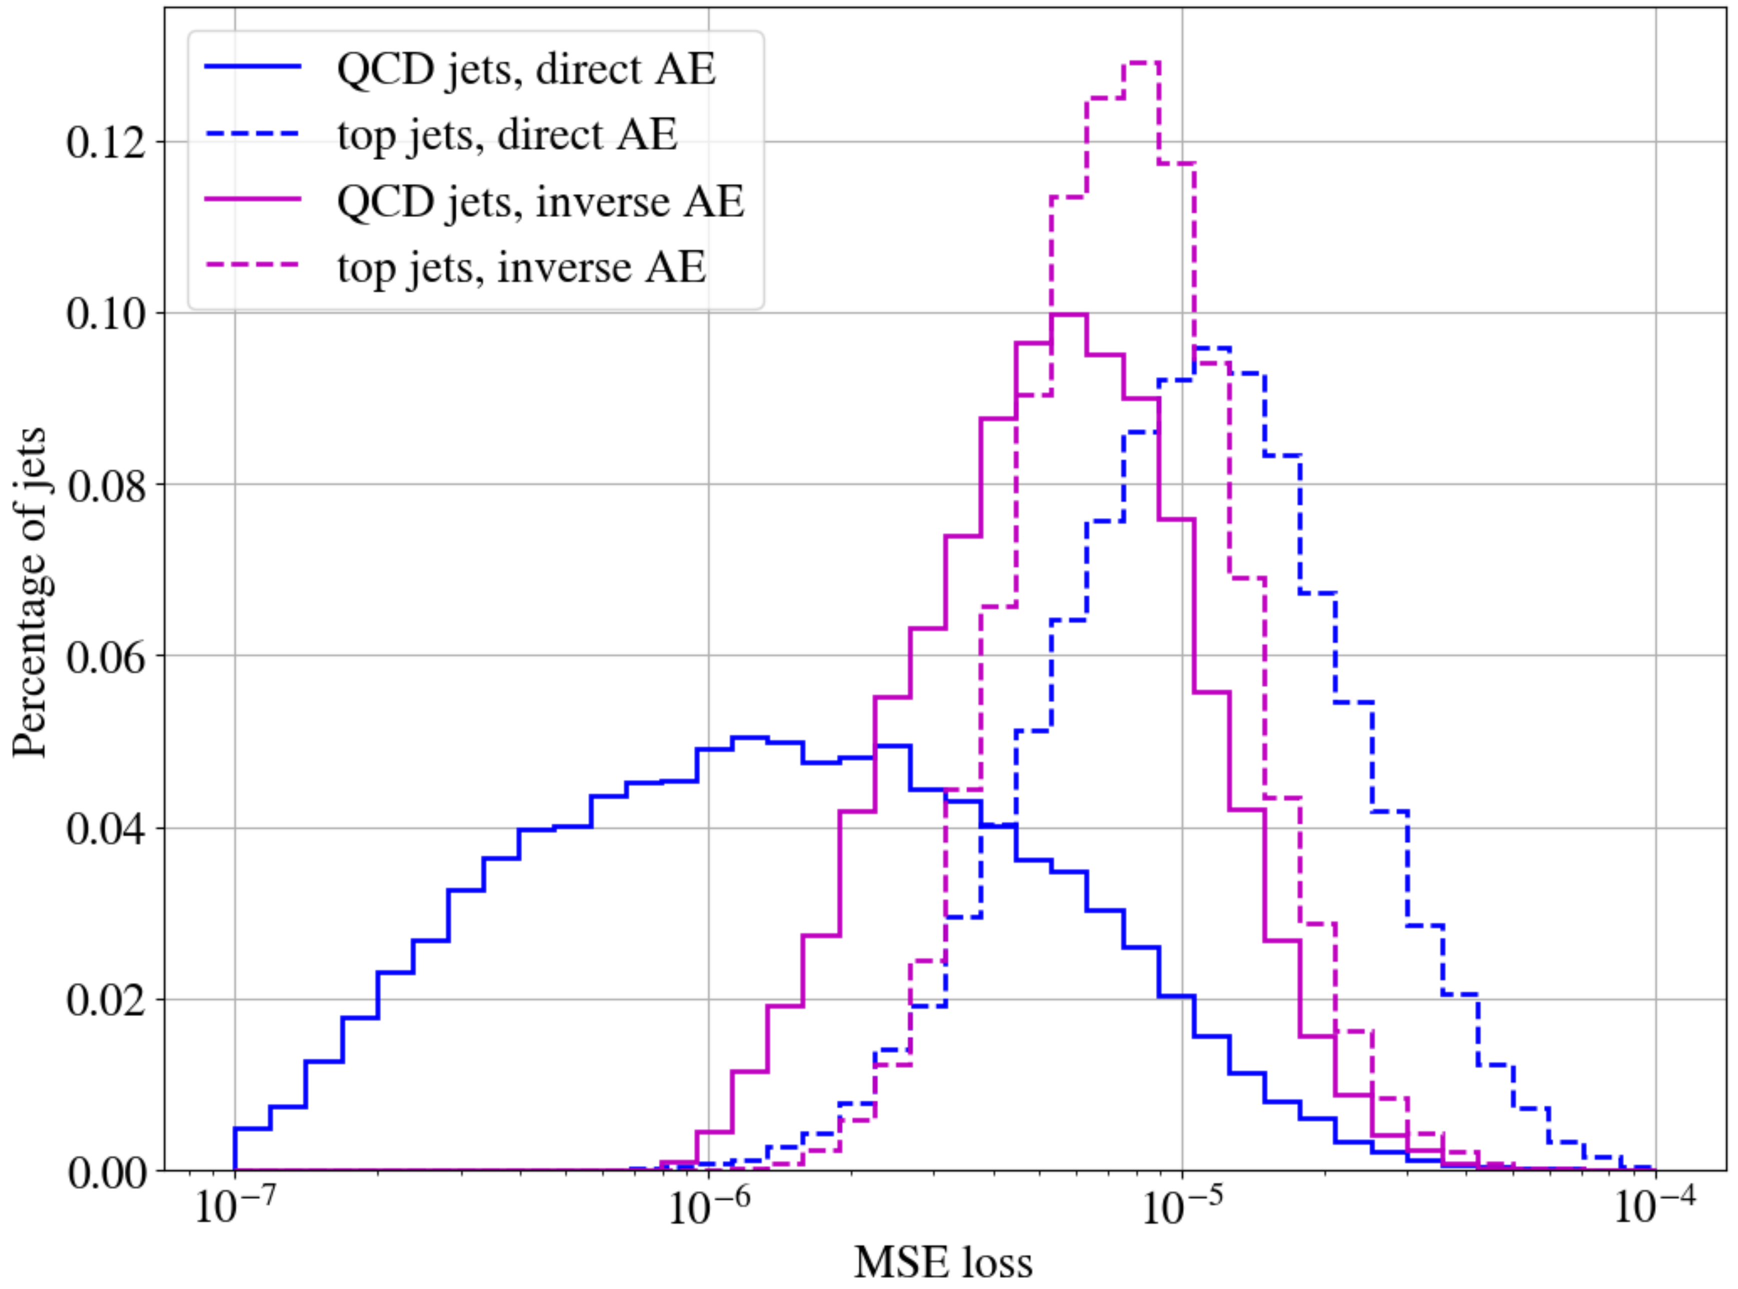

Top jets (more complex) are identified as anomalous when AE trained on QCD jets (simpler)

But not vice versa

Normalized autoencoders

Yoon et al 2105.05735, Dillon et al 2206.14225

Add additional normalization term to usual AE loss to further penalize outliers during training

$$\mathcal{L}(x) = -\log p_\theta(x) = E_\theta(x) + \log Z_\theta \quad \Rightarrow \quad \mathcal{L} = \left\langle E_\theta(x) + \log Z_\theta \right\rangle_{x \sim p_{\text{data}}}$$

from 2206.14225

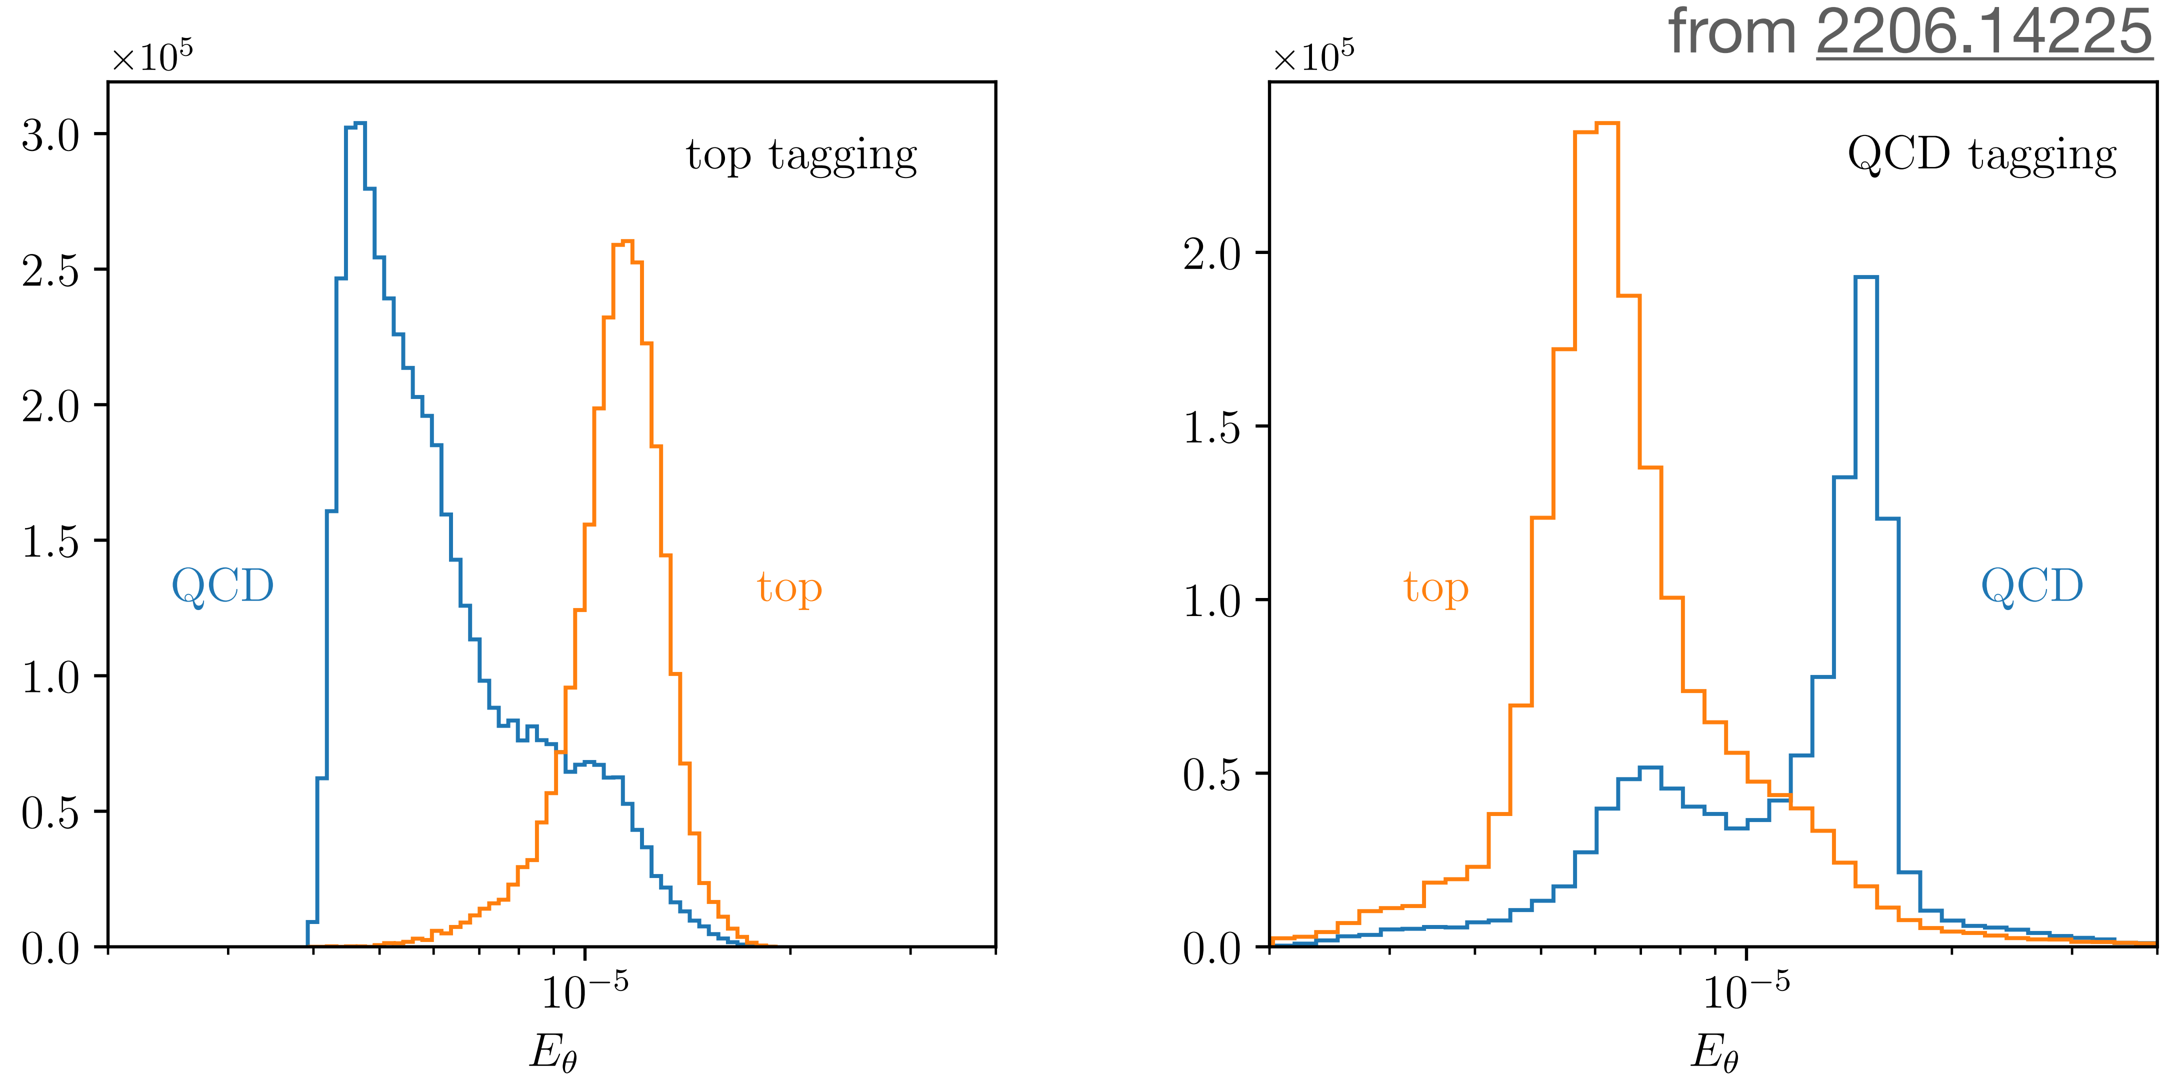

Figure 3: Distribution of the energy or MSE after training on QCD jets (left) and on top jets (right). We show the energy for QCD jets (blue) and top jets (orange) in both cases.

Now performance of AE is "symmetrical"!

- Tops are identified as anomalous when AE trained on QCD
- QCD are identified as anomalous when AE trained on Tops

Autoencoders: challenges

Autoencoders: challenges

<u>**Challenges:**</u>

- Background estimation with outlier anomaly detection

Autoencoders: challenges

Challenges:

- Background estimation with outlier anomaly detection
  - Can combine with bump hunt at cost of model-independence
    Farina, Nakai & **DS** 1808.08992; Heimel, Kasieczka, Plehn & Thompson 1808.08979

Autoencoders: challenges

Challenges:

- Background estimation with outlier anomaly detection
  - Can combine with bump hunt at cost of model-independence
    Farina, Nakai & **DS** 1808.08992; Heimel, Kasieczka, Plehn & Thompson 1808.08979
  - **New idea: Double Decorrelated AE**
    Mikuni, Nachman & **DS** 2111.06417

Autoencoders: challenges

Challenges:

- Background estimation with outlier anomaly detection
  - Can combine with bump hunt at cost of model-independence
    Farina, Nakai & **DS** 1808.08992; Heimel, Kasieczka, Plehn & Thompson 1808.08979
  - **New idea: Double Decorrelated AE**
    Mikuni, Nachman & **DS** 2111.06417
    - Train *two* autoencoders and force them to be statistically independent of one another

$$L[f_1, f_2, g_1, g_2] = \sum_i R_1(x_i)^2 + \sum_i R_2(x_i)^2 + \lambda\,\mathrm{DisCo}^2[R_1(X), R_2(X)]$$

"DisCo Decorrelation"
*Kasieczka & **DS** 2001.05310*

Szekely et al 0803.4101 et seq

Autoencoders: challenges

Challenges:

- Background estimation with outlier anomaly detection
  - Can combine with bump hunt at cost of model-independence
    Farina, Nakai & **DS** 1808.08992; Heimel, Kasieczka, Plehn & Thompson 1808.08979
  - **New idea: Double Decorrelated AE**
    Mikuni, Nachman & **DS** 2111.06417
    - Train *two* autoencoders and force them to be statistically independent of one another
    - Use ABCD method for fully data-driven background estimation

$$L[f_1, f_2, g_1, g_2] = \sum_i R_1(x_i)^2 + \sum_i R_2(x_i)^2 + \lambda\, \mathrm{DisCo}^2[R_1(X), R_2(X)]$$

"DisCo Decorrelation"
*Kasieczka & **DS** 2001.05310*

Szekely et al 0803.4101 et seq

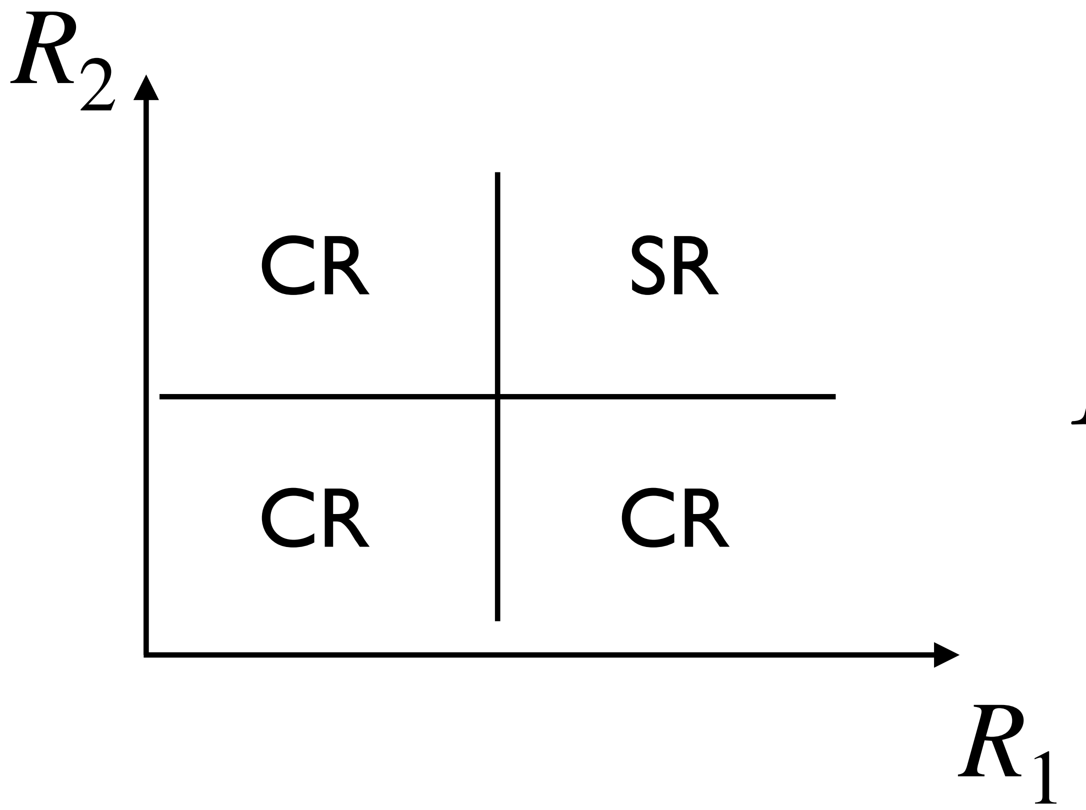

$$N^{\text{predicted}}_{>,>}(\vec{c}) = \frac{N_{>,<}(\vec{c})N_{<,>}(\vec{c})}{N_{<,<}(\vec{c})}$$

Double Decorrelated AE

Mikuni, Nachman & **DS** 2111.06417

The method works!

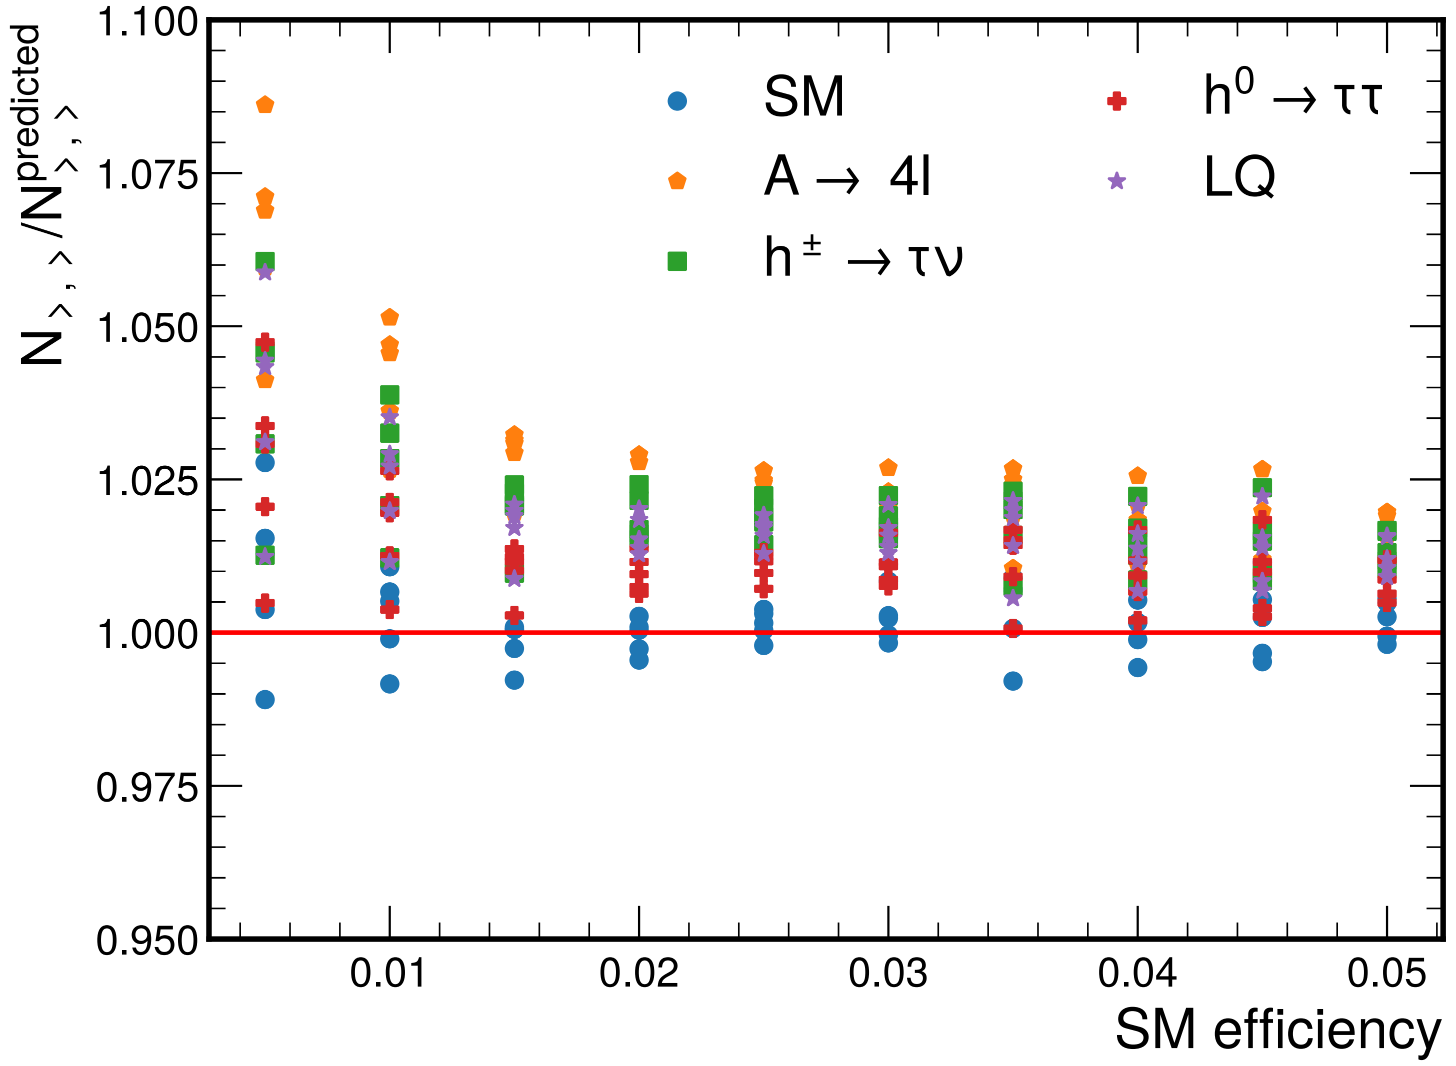

ADC2021 1-lepton dataset
Govorkova et al 2107.02157

(Initial NP significance: 0.8σ)

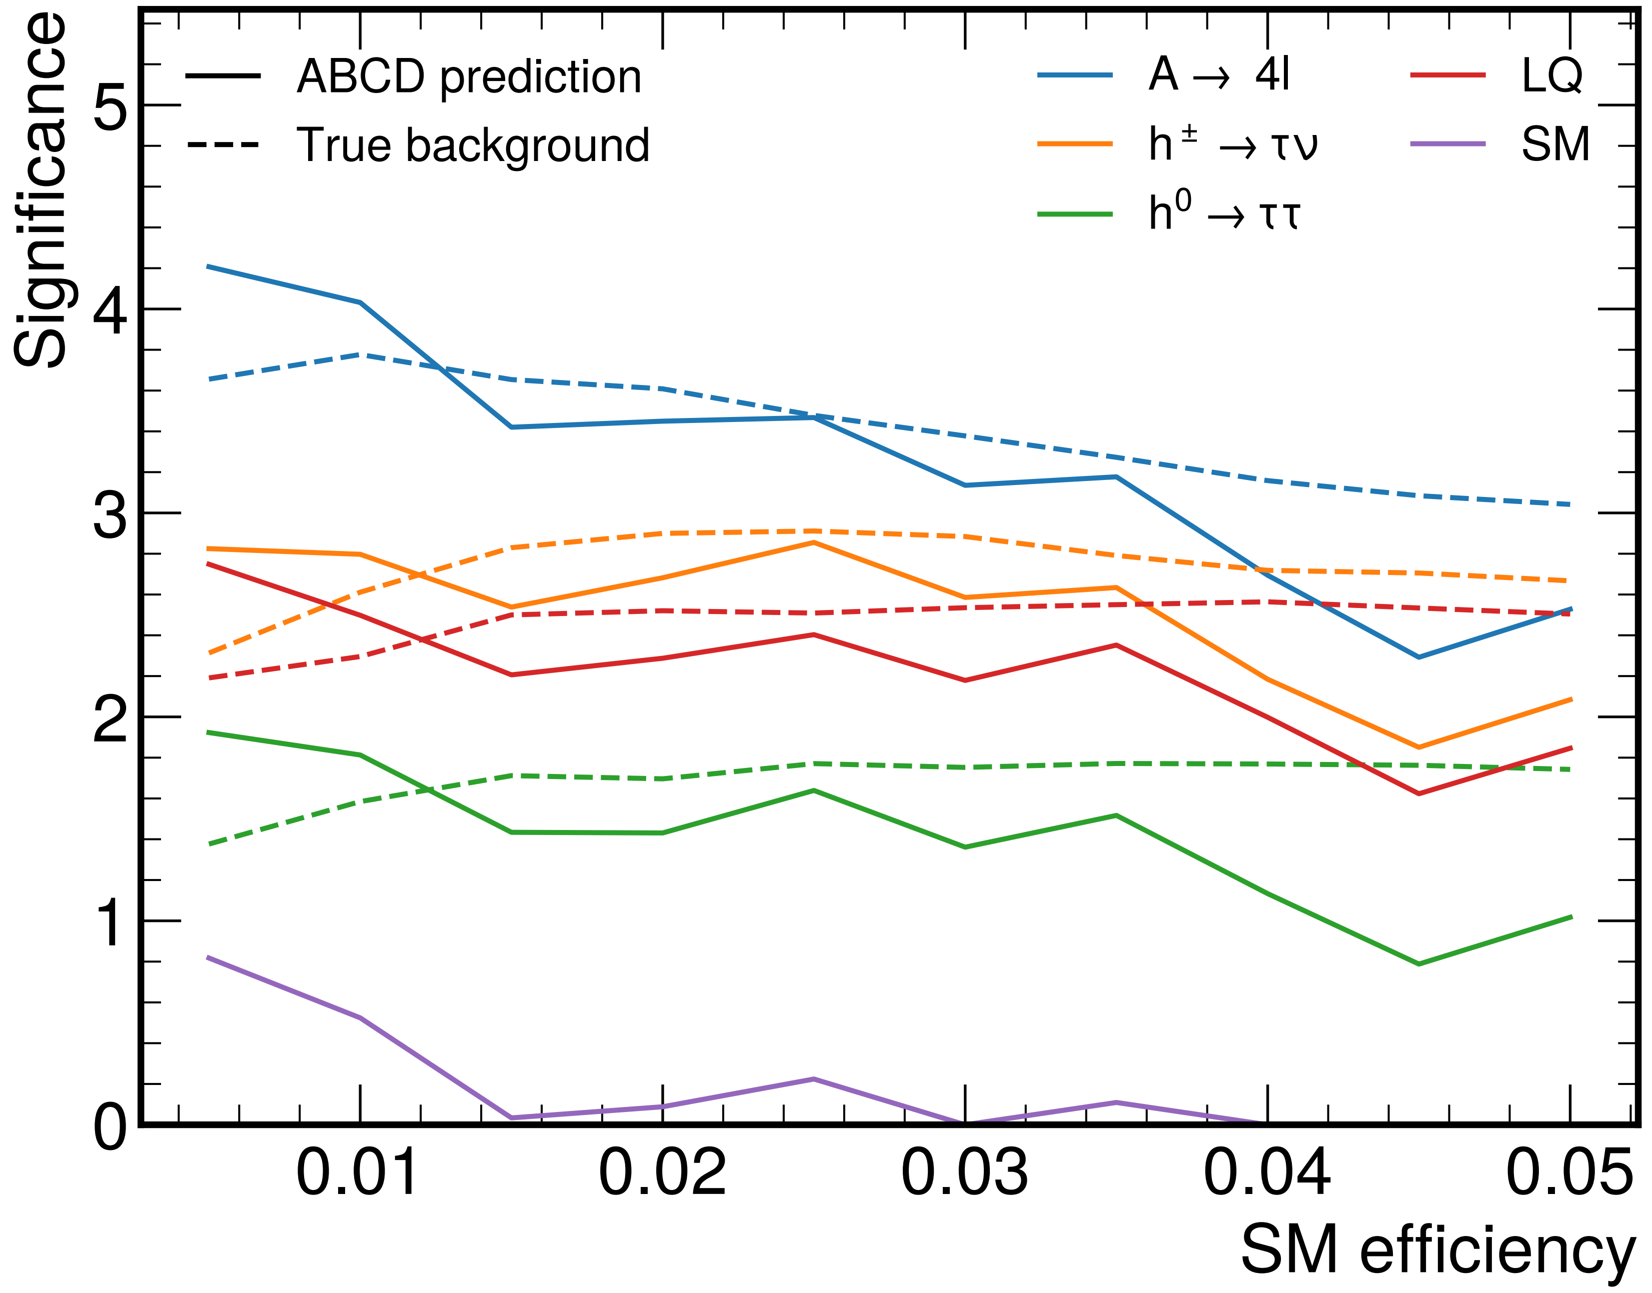

First complete strategy for unsupervised, non-resonant anomaly detection

Can also be used online as an anomaly trigger

2b. Overdensity detection

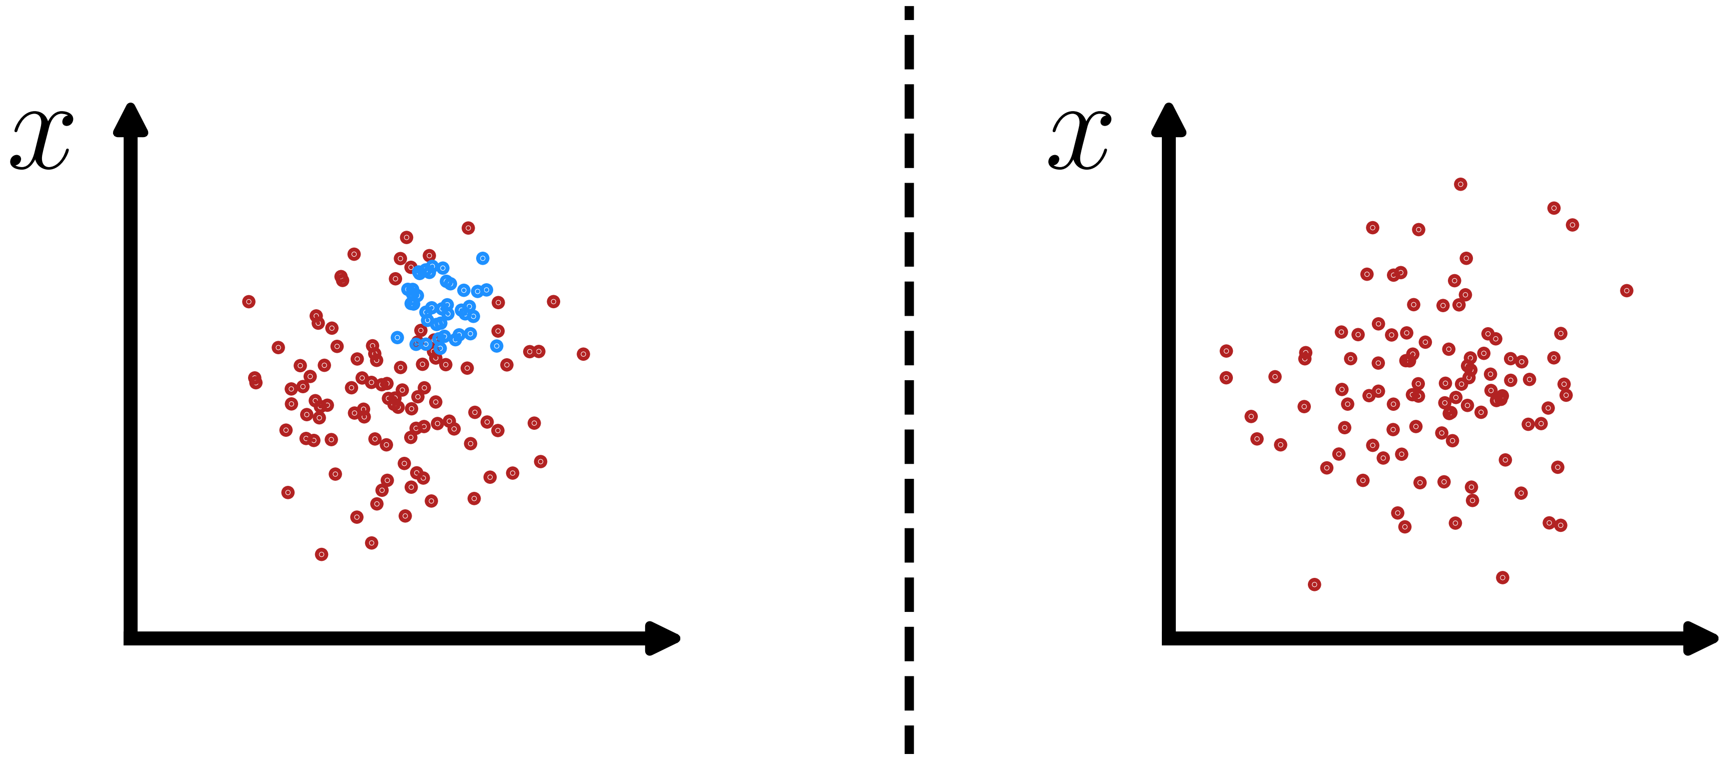

high $\frac{p_{data}(x)}{p_{bg}(x)}$

- Pros:
  - reparametrization invariant
  - asymptotically optimal
- Cons:
  - Requires more precise knowledge of background (reference) distribution
  - Performance suffers when signal is too rare

Classic Overdensity Search

1D Bump Hunt

Idea: assume signal is localized in some feature (usually invariant mass) while background is smooth.

Interpolate from **sidebands** into **signal region** (eg window in invariant mass), search for an excess.

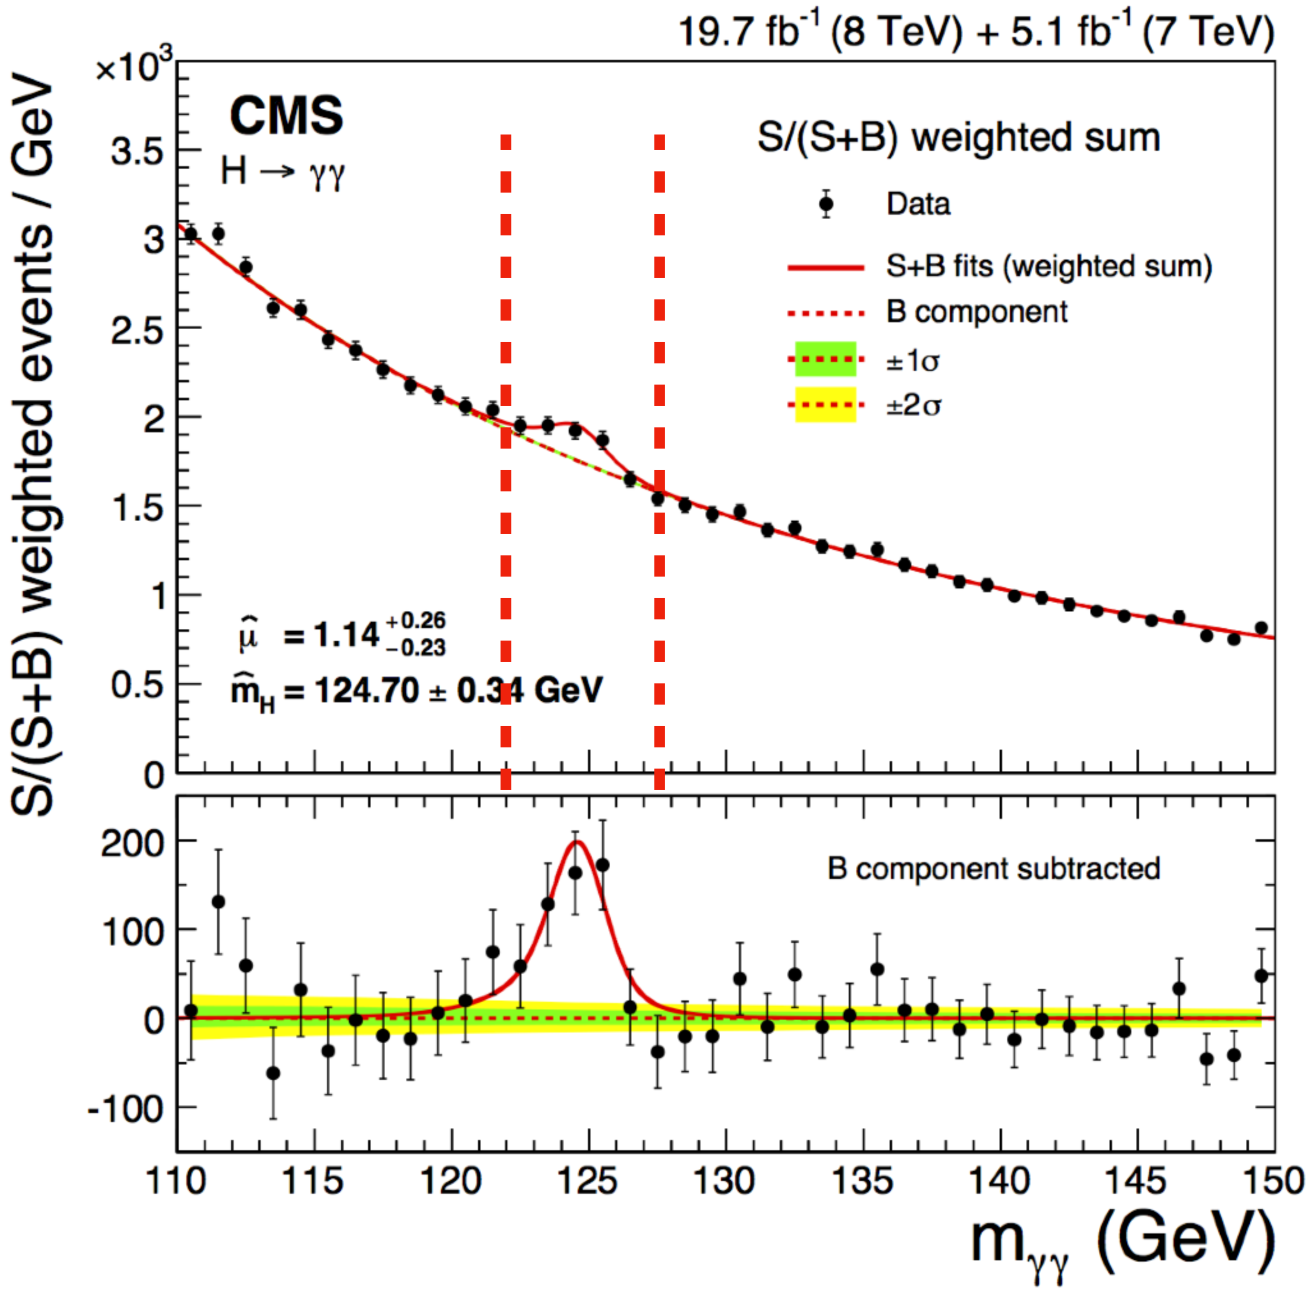

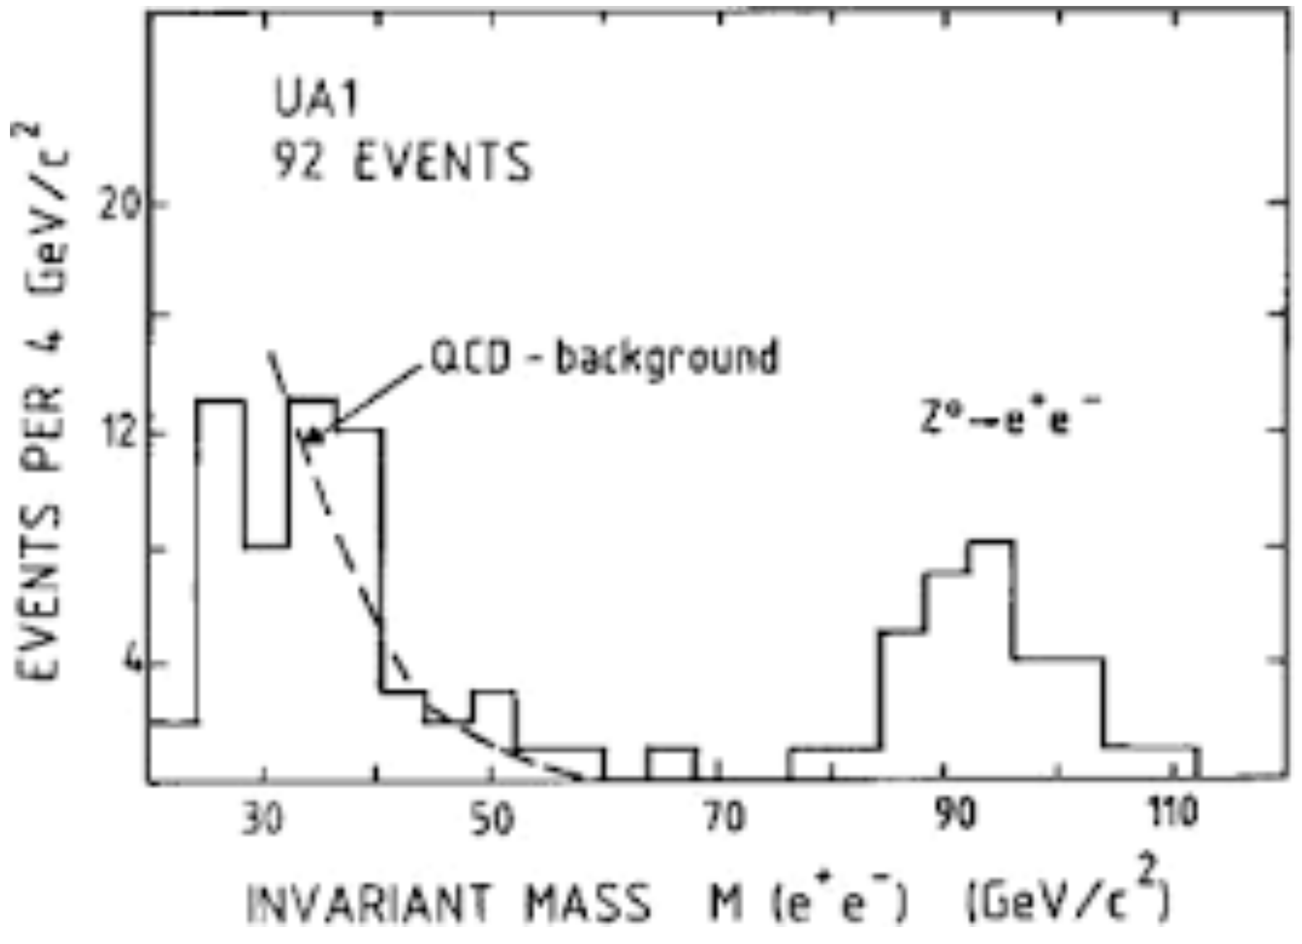

Used in many discoveries!

ML-enhanced bump hunts

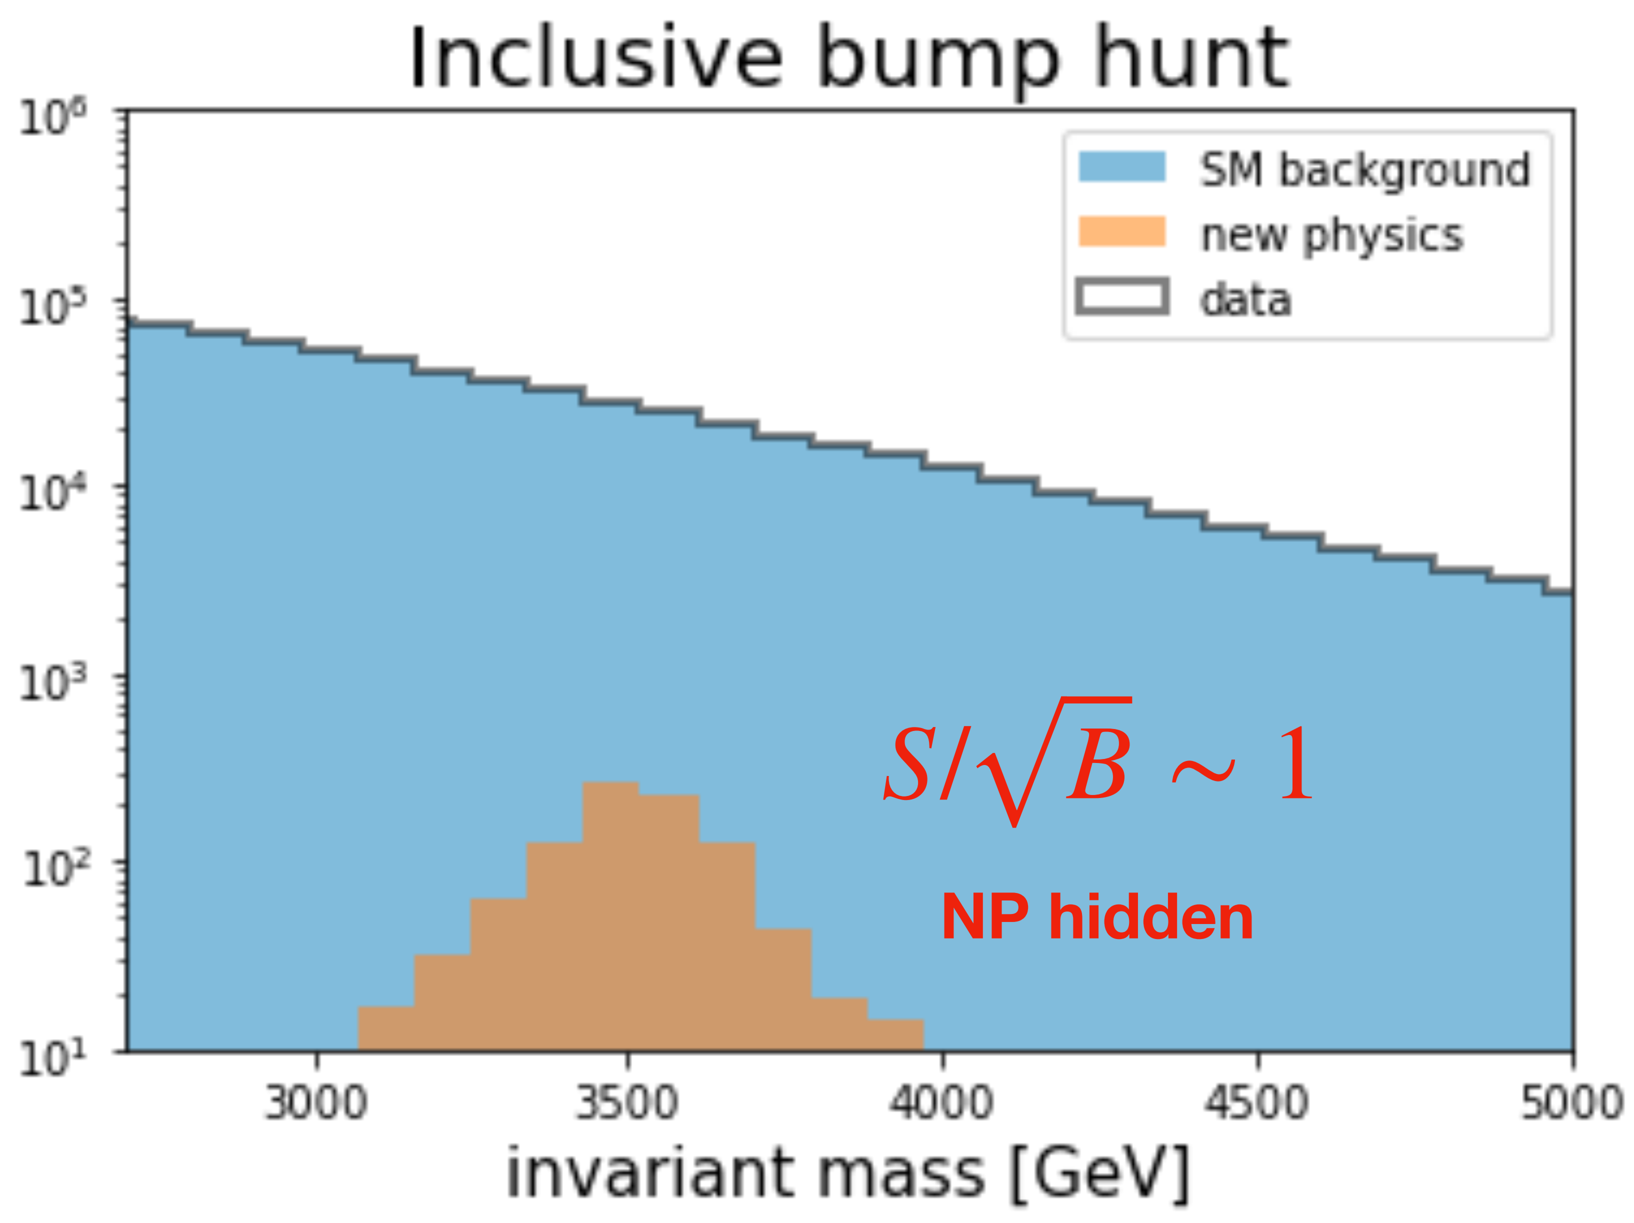

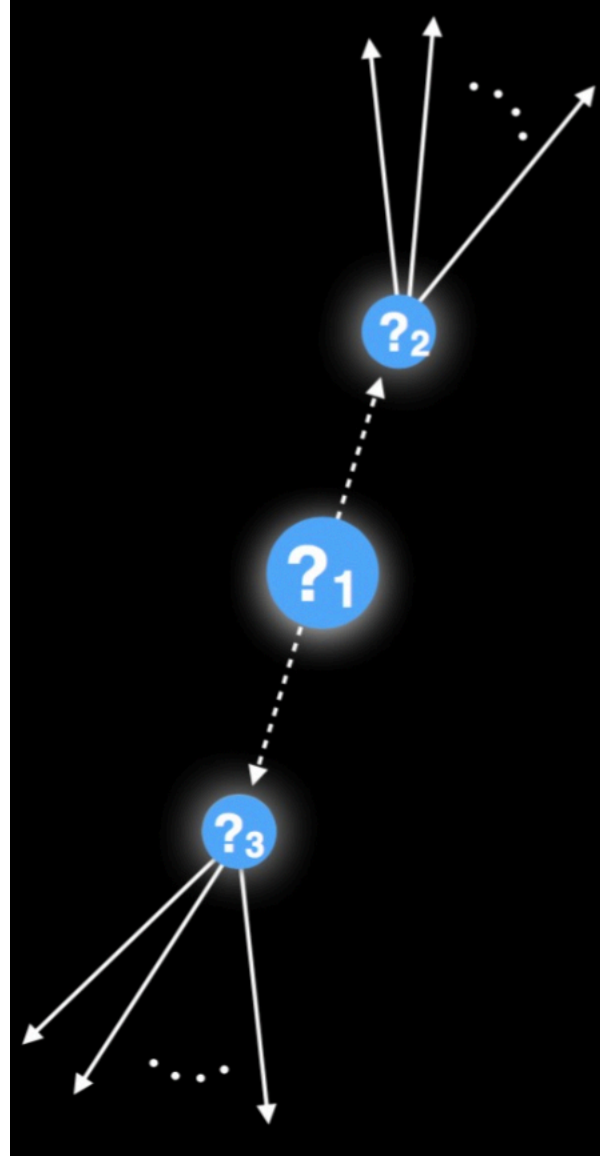

ML-enhanced bump hunts

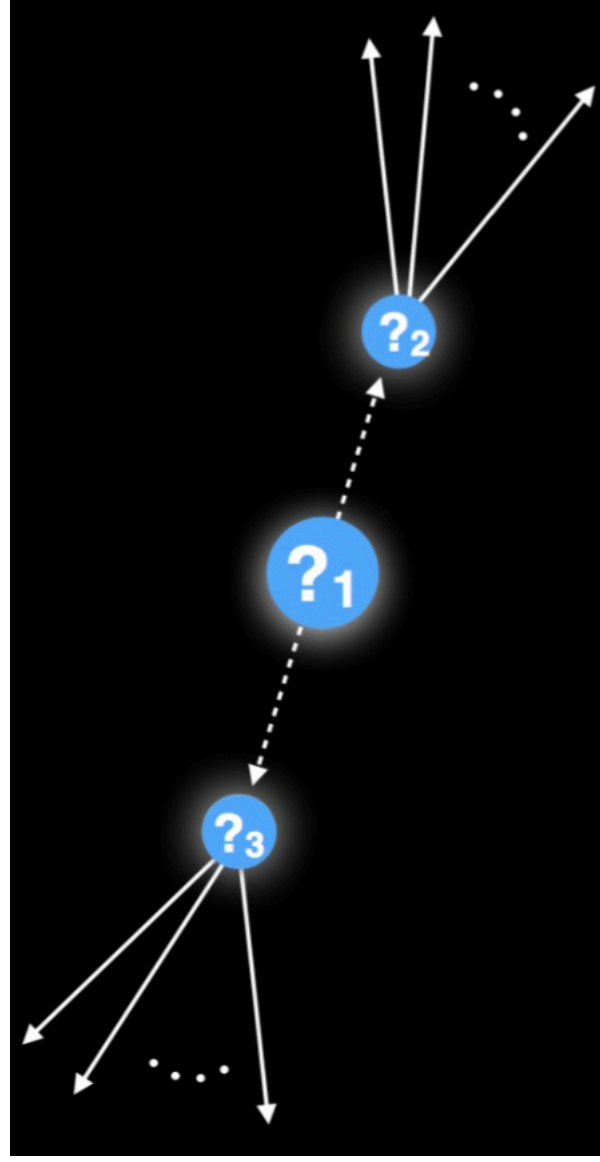

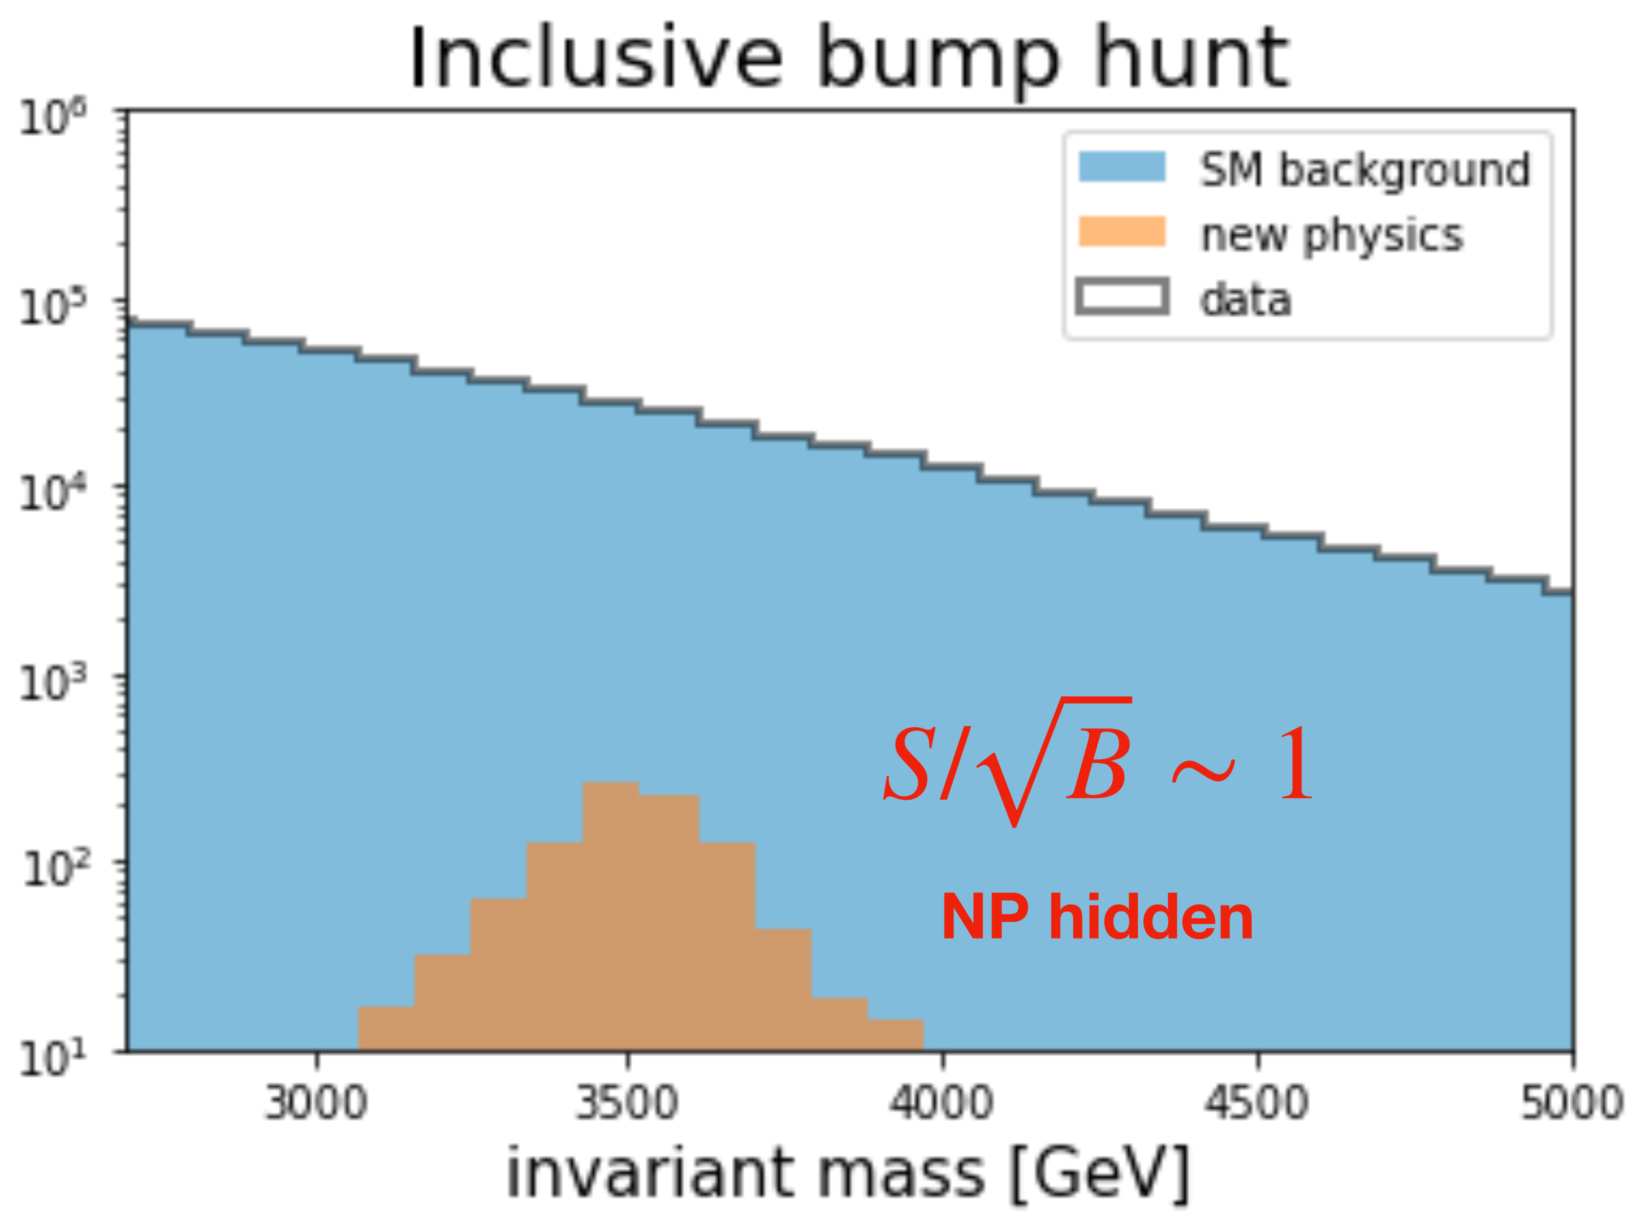

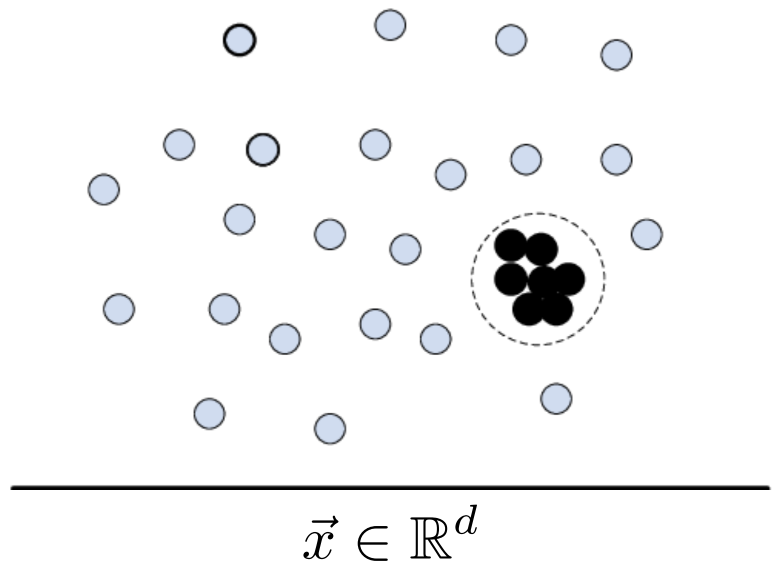

x: ***additional*** features where NP could be localized

ML-enhanced bump hunts

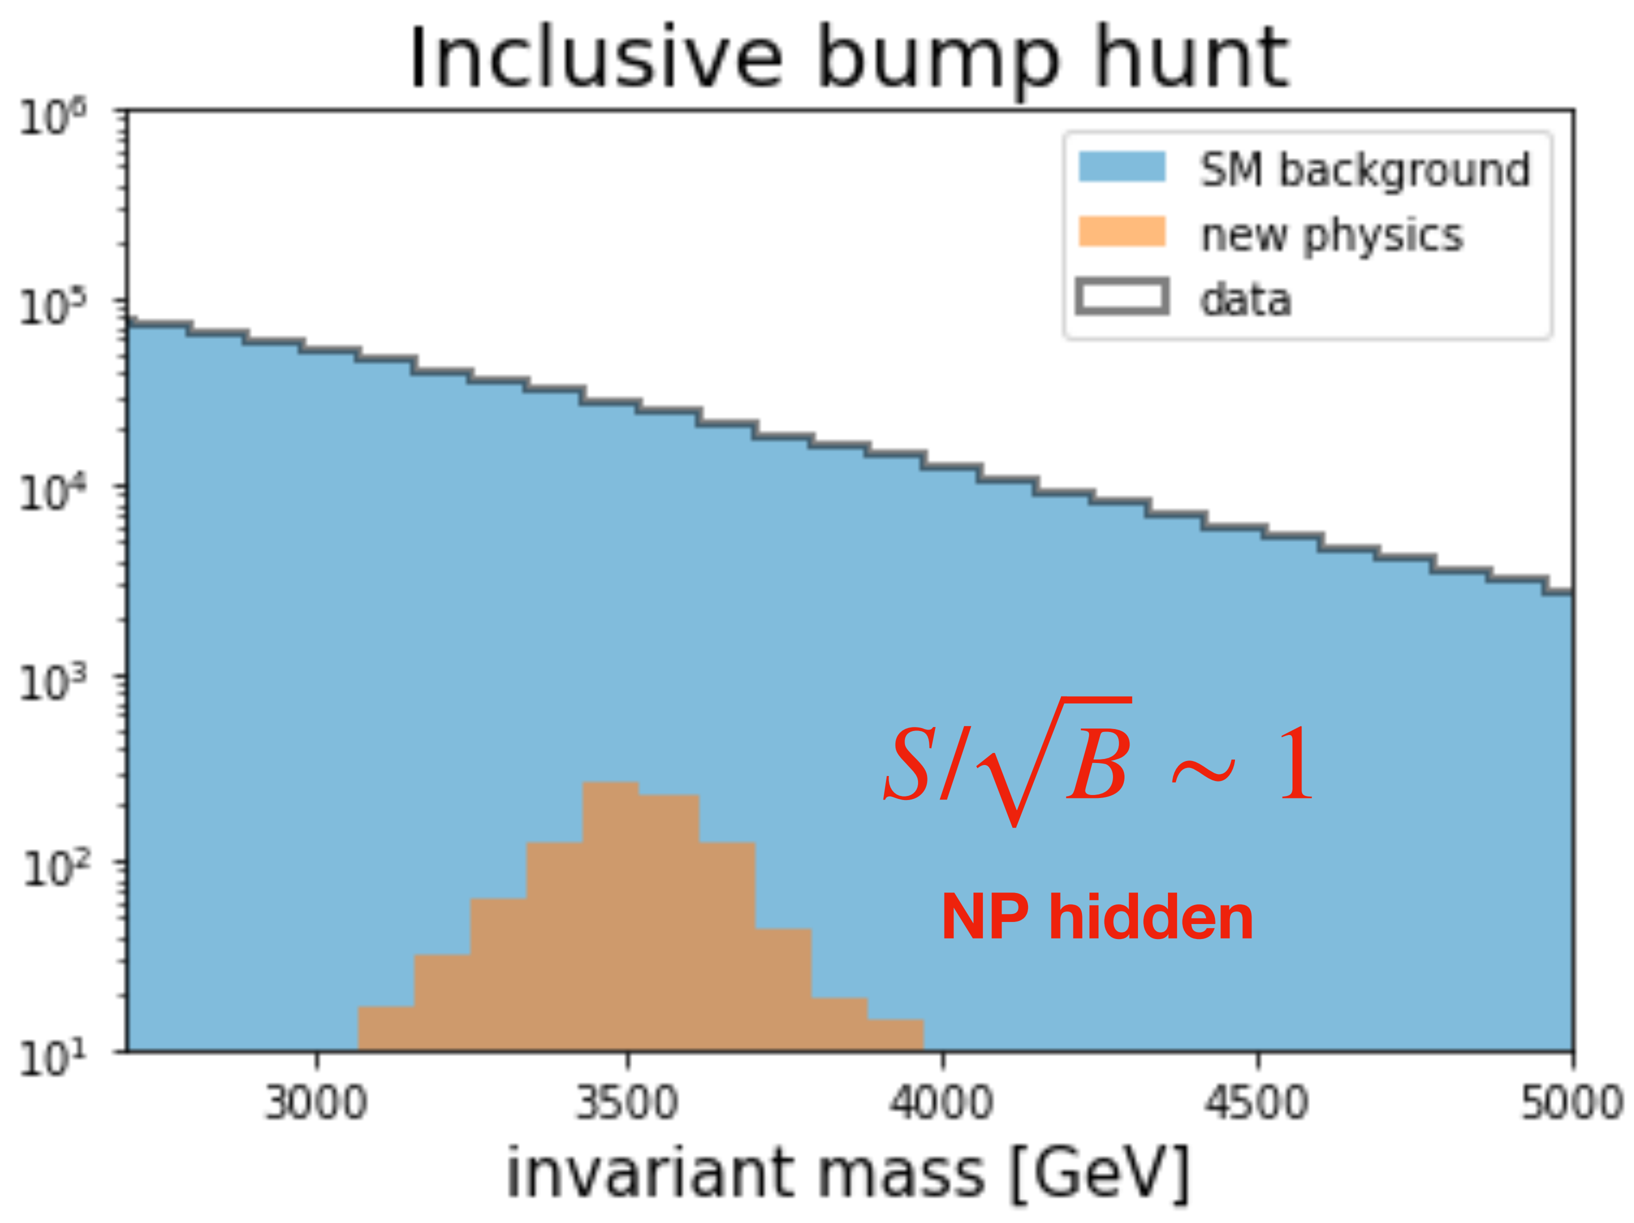

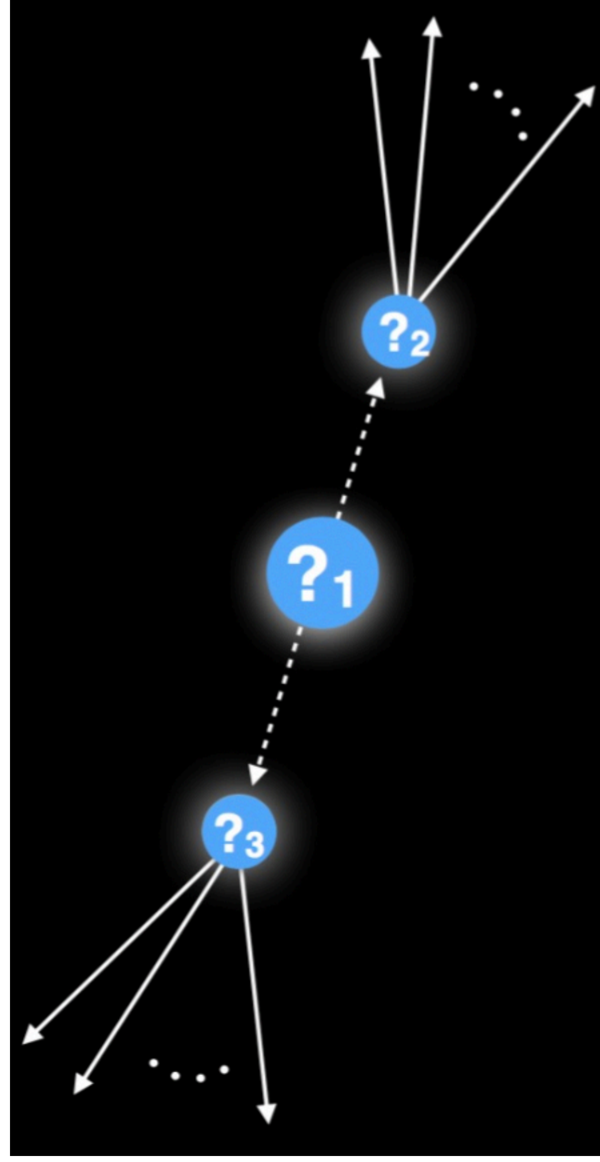

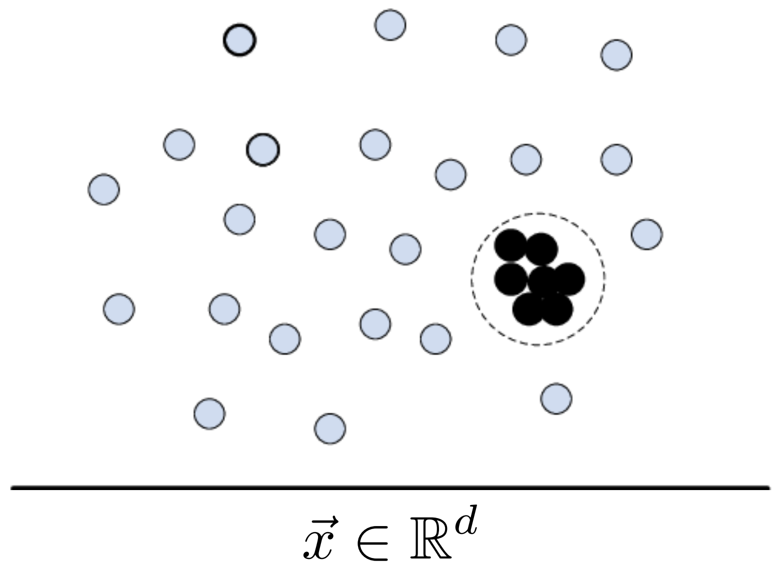

x: ***additional*** features where NP could be localized

Learn model-agnostic **anomaly score $R(x)$** from data

ML-enhanced bump hunts

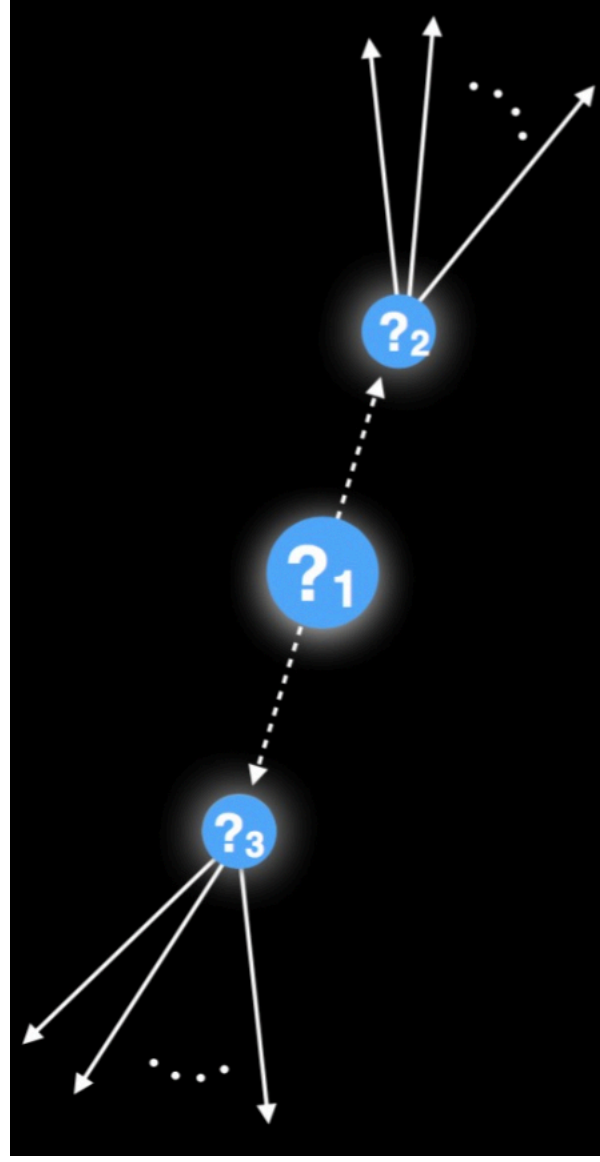

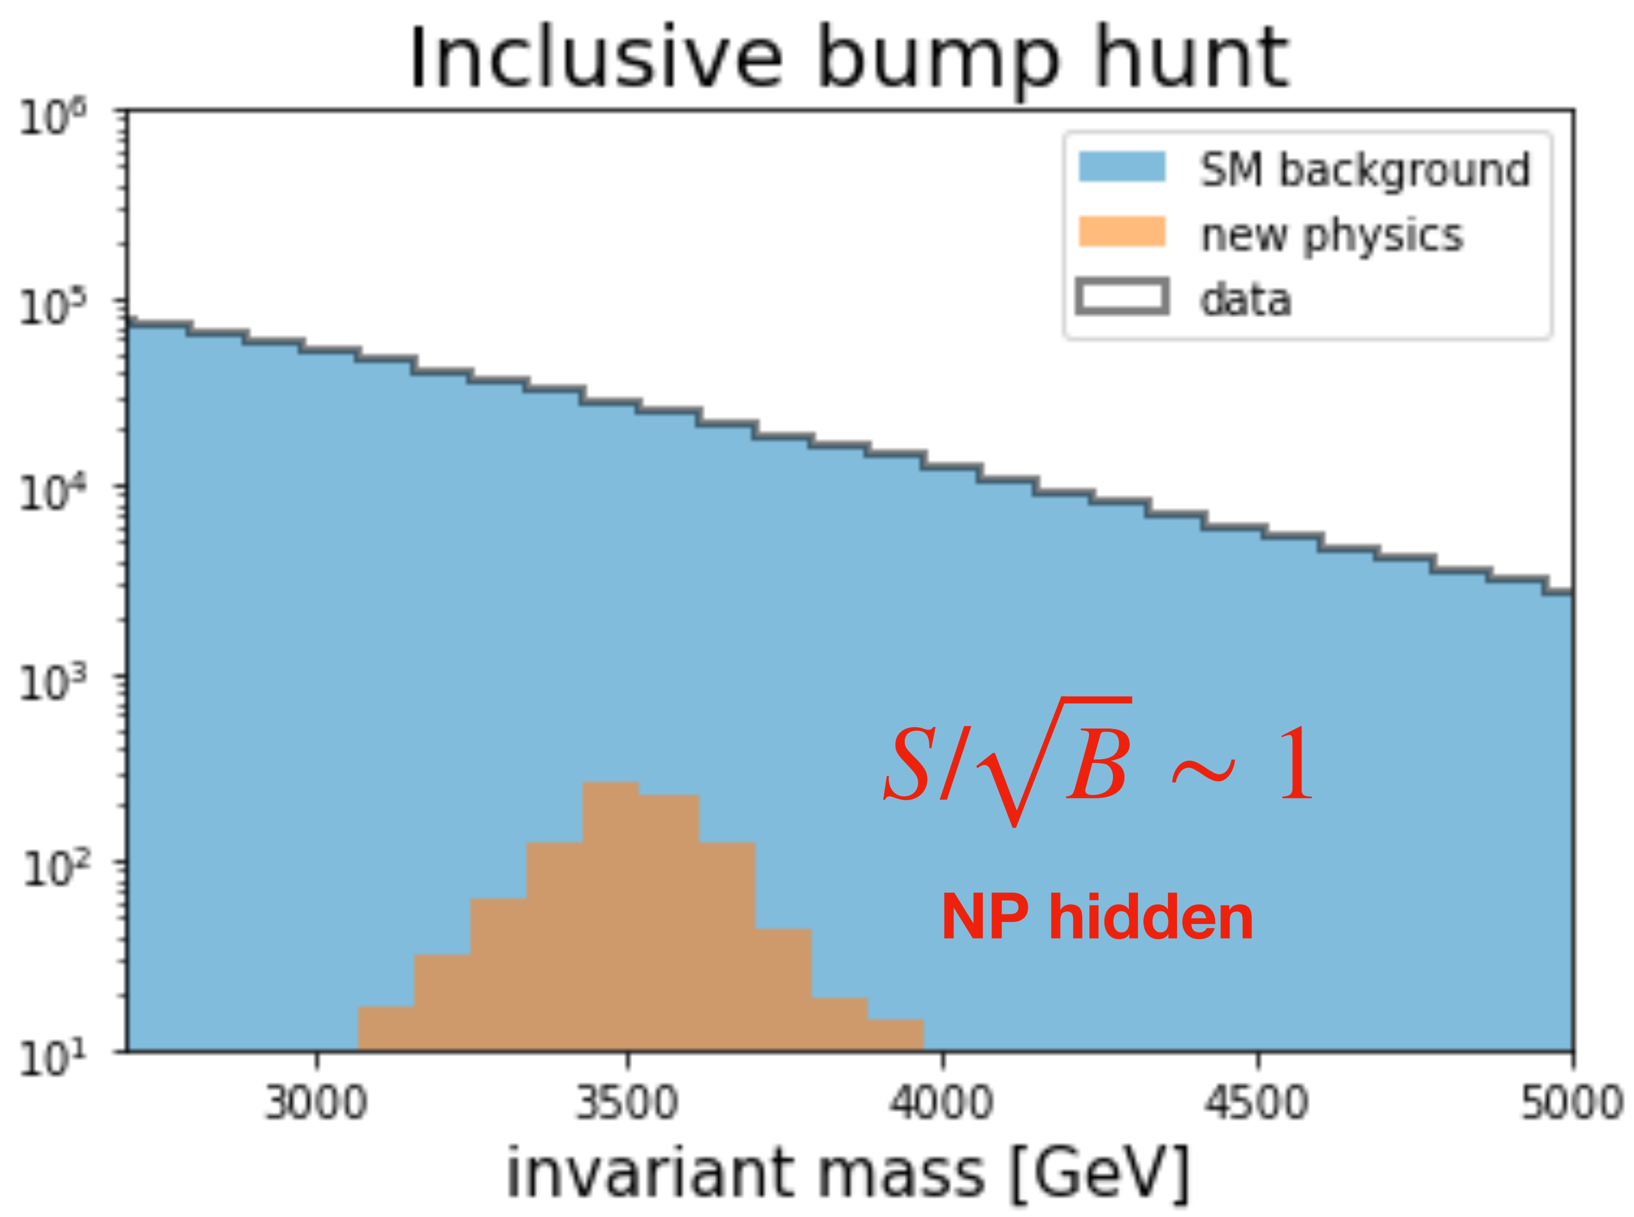

Cut on anomaly score $R(x)$

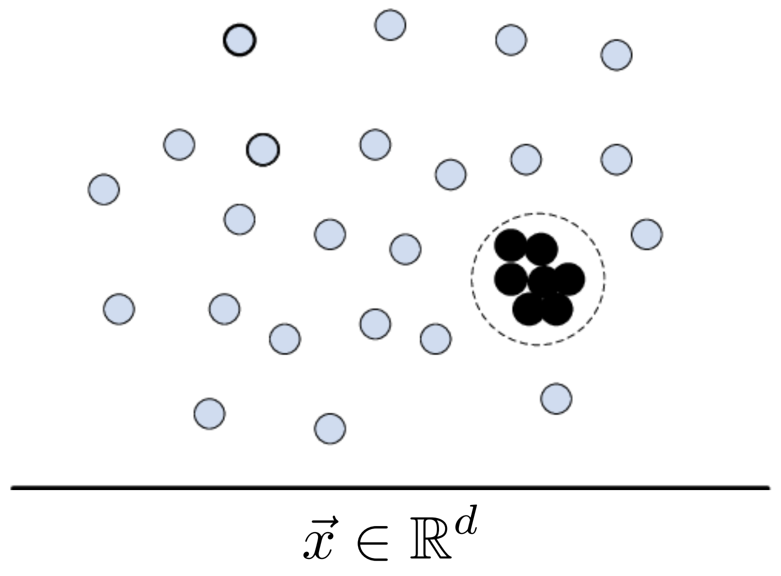

x: ***additional*** features where NP could be localized

Learn model-agnostic **anomaly score** $R(x)$ from data

ML-enhanced bump hunts

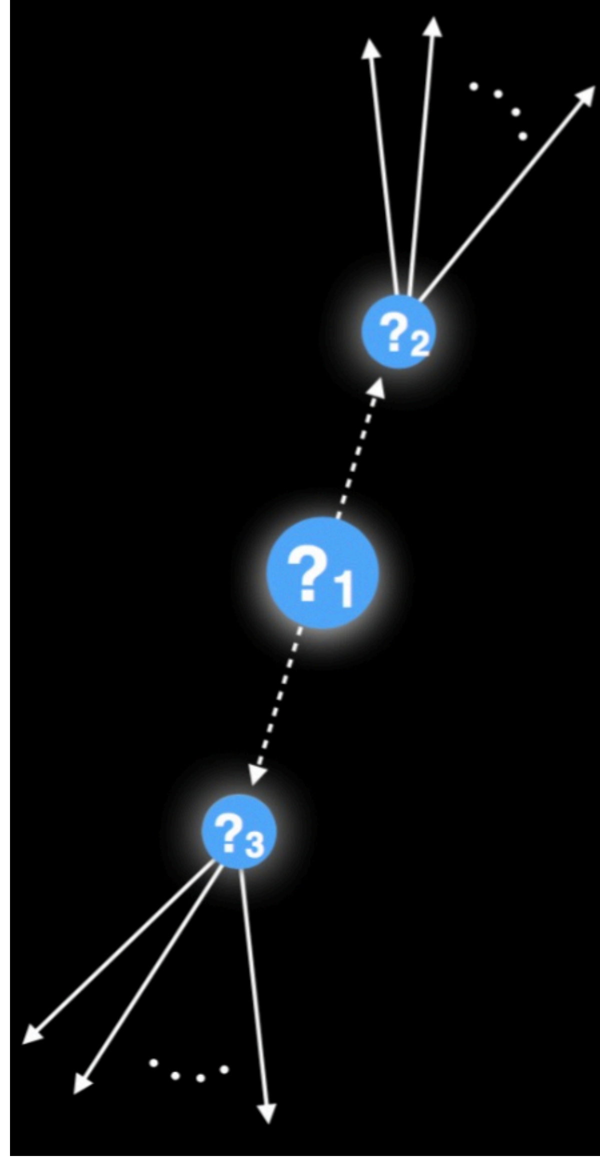

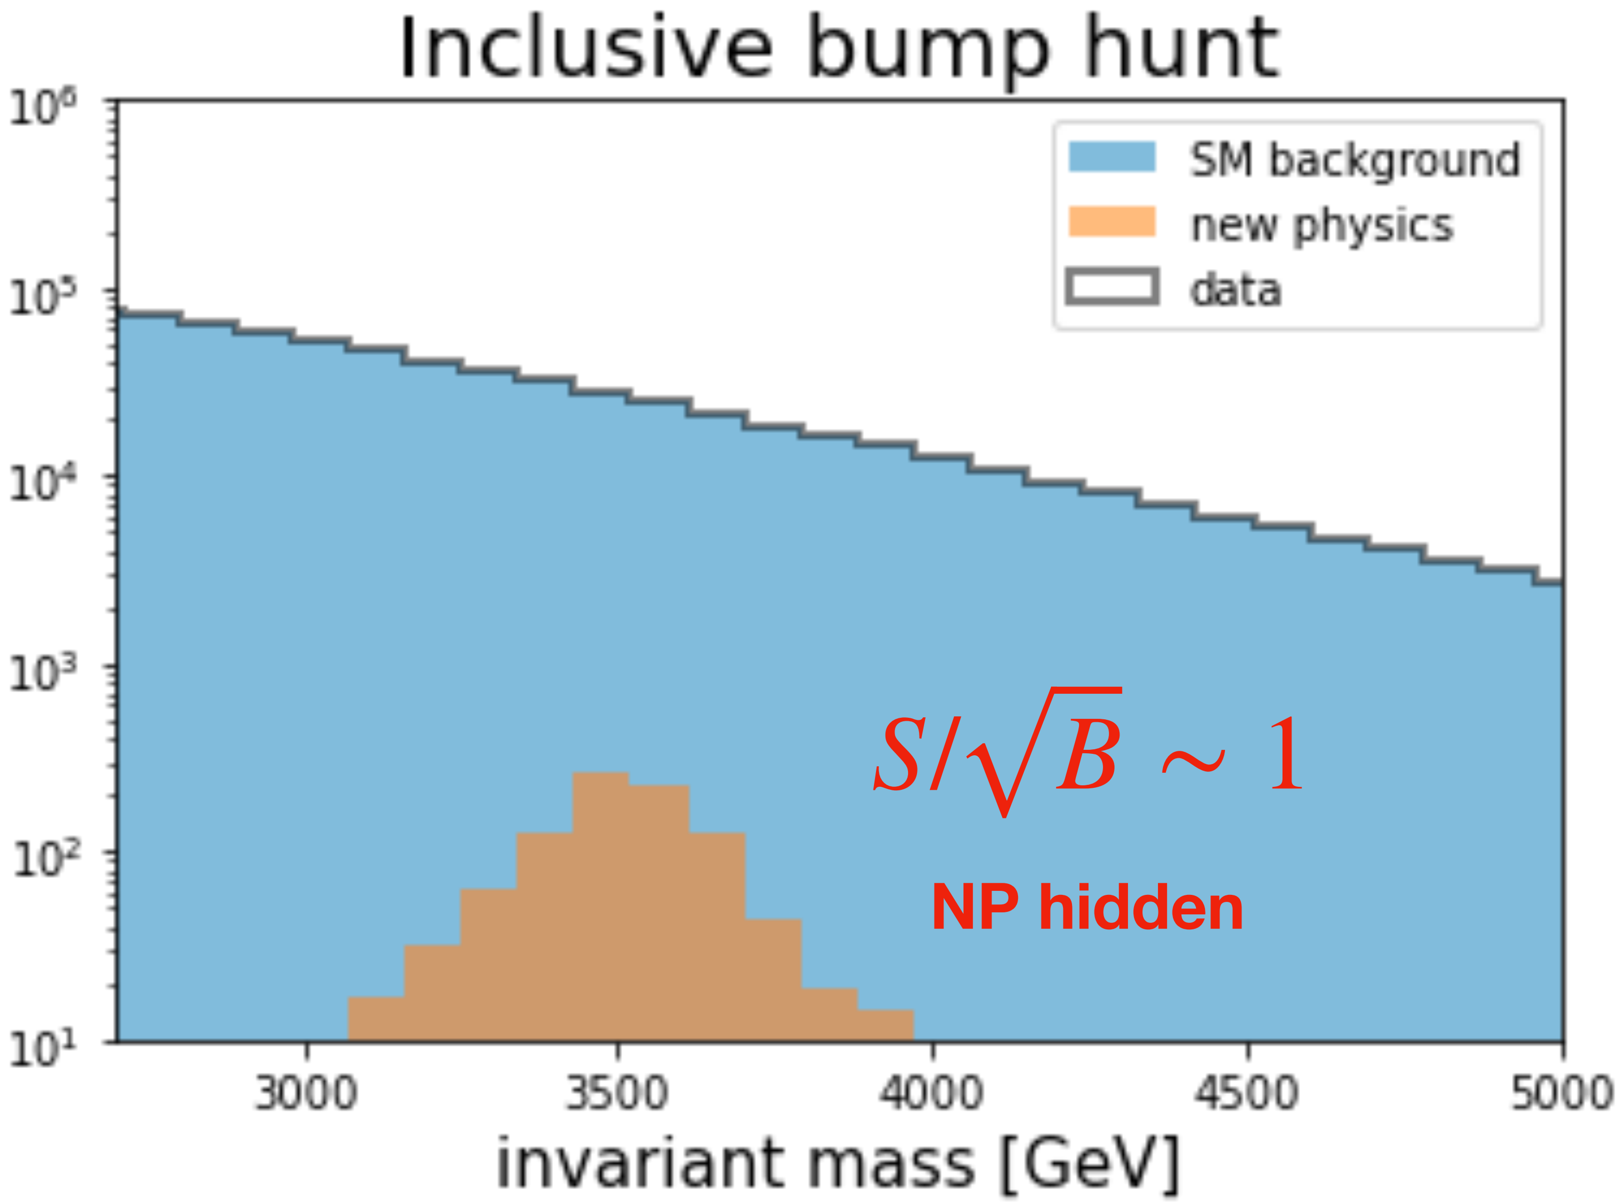

Cut on anomaly score $R(x)$

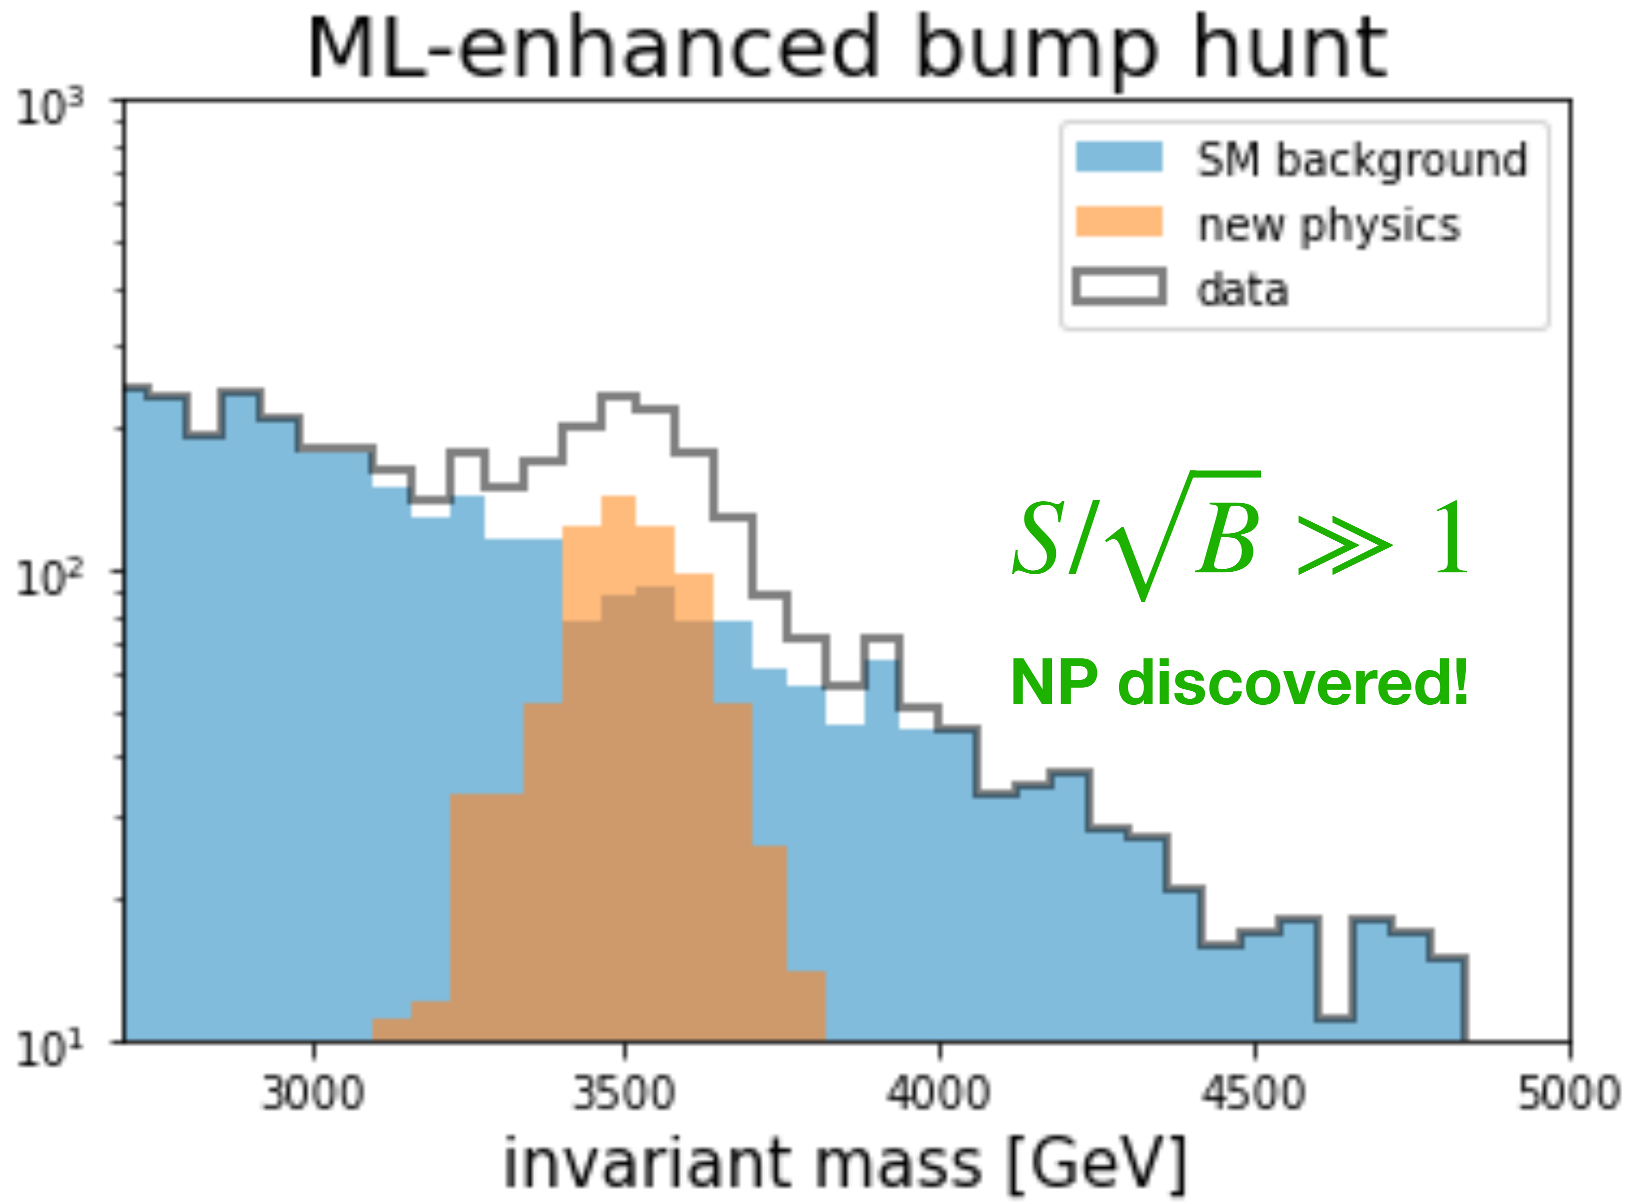

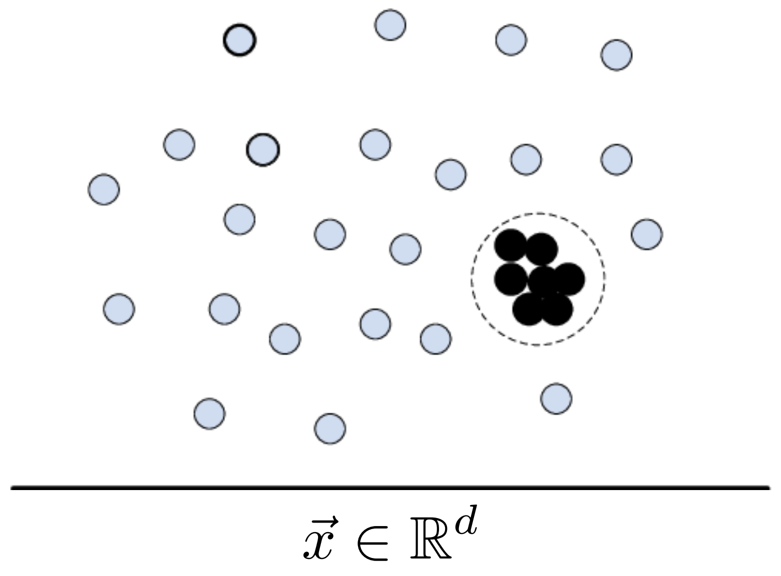

x: ***additional*** features where NP could be localized

Learn model-agnostic **anomaly score $R(x)$** from data

Idealized Anomaly Detector

Claim: the optimal model-agnostic discriminant would be (Neyman & Pearson)

$$\boxed{R(x) = \frac{p_{data}(x)}{p_{bg}(x)}}$$

"Idealized Anomaly Detector"

Idealized Anomaly Detector

Claim: the optimal model-agnostic discriminant would be (Neyman & Pearson)

$$\boxed{R(x) = \frac{p_{data}(x)}{p_{bg}(x)}}$$

"Idealized Anomaly Detector"

Proof:

$$p_{data}(x) = \epsilon_{sig} p_{sig}(x) + (1 - \epsilon_{sig}) p_{bg}(x)$$

$$R(x) = (1 - \epsilon_{sig}) + \epsilon_{sig} \frac{p_{sig}(x)}{p_{bg}(x)}$$

Idealized Anomaly Detector

Claim: the optimal model-agnostic discriminant would be (Neyman & Pearson)

$$\boxed{R(x) = \frac{p_{data}(x)}{p_{bg}(x)}}$$

"Idealized Anomaly Detector"

Proof:

$$p_{data}(x) = \epsilon_{sig} p_{sig}(x) + (1 - \epsilon_{sig}) p_{bg}(x)$$

$$R(x) = (1 - \epsilon_{sig}) + \epsilon_{sig} \frac{p_{sig}(x)}{p_{bg}(x)}$$

$R(x)$ is monotonic with signal-to-background likelihood ratio ***regardless of unknown, arbitrary signal strength and probability density***

Idea: data vs simulation classifier

Train a neural network to classify data vs MC simulation of the SM.

If the NN classifier is optimal, its output should be (monotonic with) the data vs MC likelihood ratio (Neyman-Pearson).

$$R_{classifier}(x) = \frac{p_{data}(x)}{p_{MC}(x)}$$

Idea: data vs simulation classifier

Train a neural network to classify data vs MC simulation of the SM.

If the NN classifier is optimal, its output should be (monotonic with) the data vs MC likelihood ratio (Neyman-Pearson).

$$R_{classifier}(x) = \frac{p_{data}(x)}{p_{MC}(x)}$$

"The likelihood-ratio trick"

Idea: data vs simulation classifier

Train a neural network to classify data vs MC simulation of the SM.

If the NN classifier is optimal, its output should be (monotonic with) the data vs MC likelihood ratio (Neyman-Pearson).

$$R_{classifier}(x) = \frac{p_{data}(x)}{p_{MC}(x)}$$

“The likelihood-ratio trick”

This can work if simulations are reliable and their systematic uncertainties are well-understood.

D’Agnolo et al
2111.13633

Idea: data vs simulation classifier

Train a neural network to classify data vs MC simulation of the SM.

If the NN classifier is optimal, its output should be (monotonic with) the data vs MC likelihood ratio (Neyman-Pearson).

$$R_{classifier}(x) = \frac{p_{data}(x)}{p_{MC}(x)}$$

"The likelihood-ratio trick"

This can work if simulations are reliable and their systematic uncertainties are well-understood.

D'Agnolo et al 2111.13633

Alternatively, can we get $p_{data}(x)$ and $p_{bg}(x)$ in a data-driven way?

Idea: data vs sideband classifier

Collins, Howe & Nachman 1805.02664,1902.02634

from 2109.00546

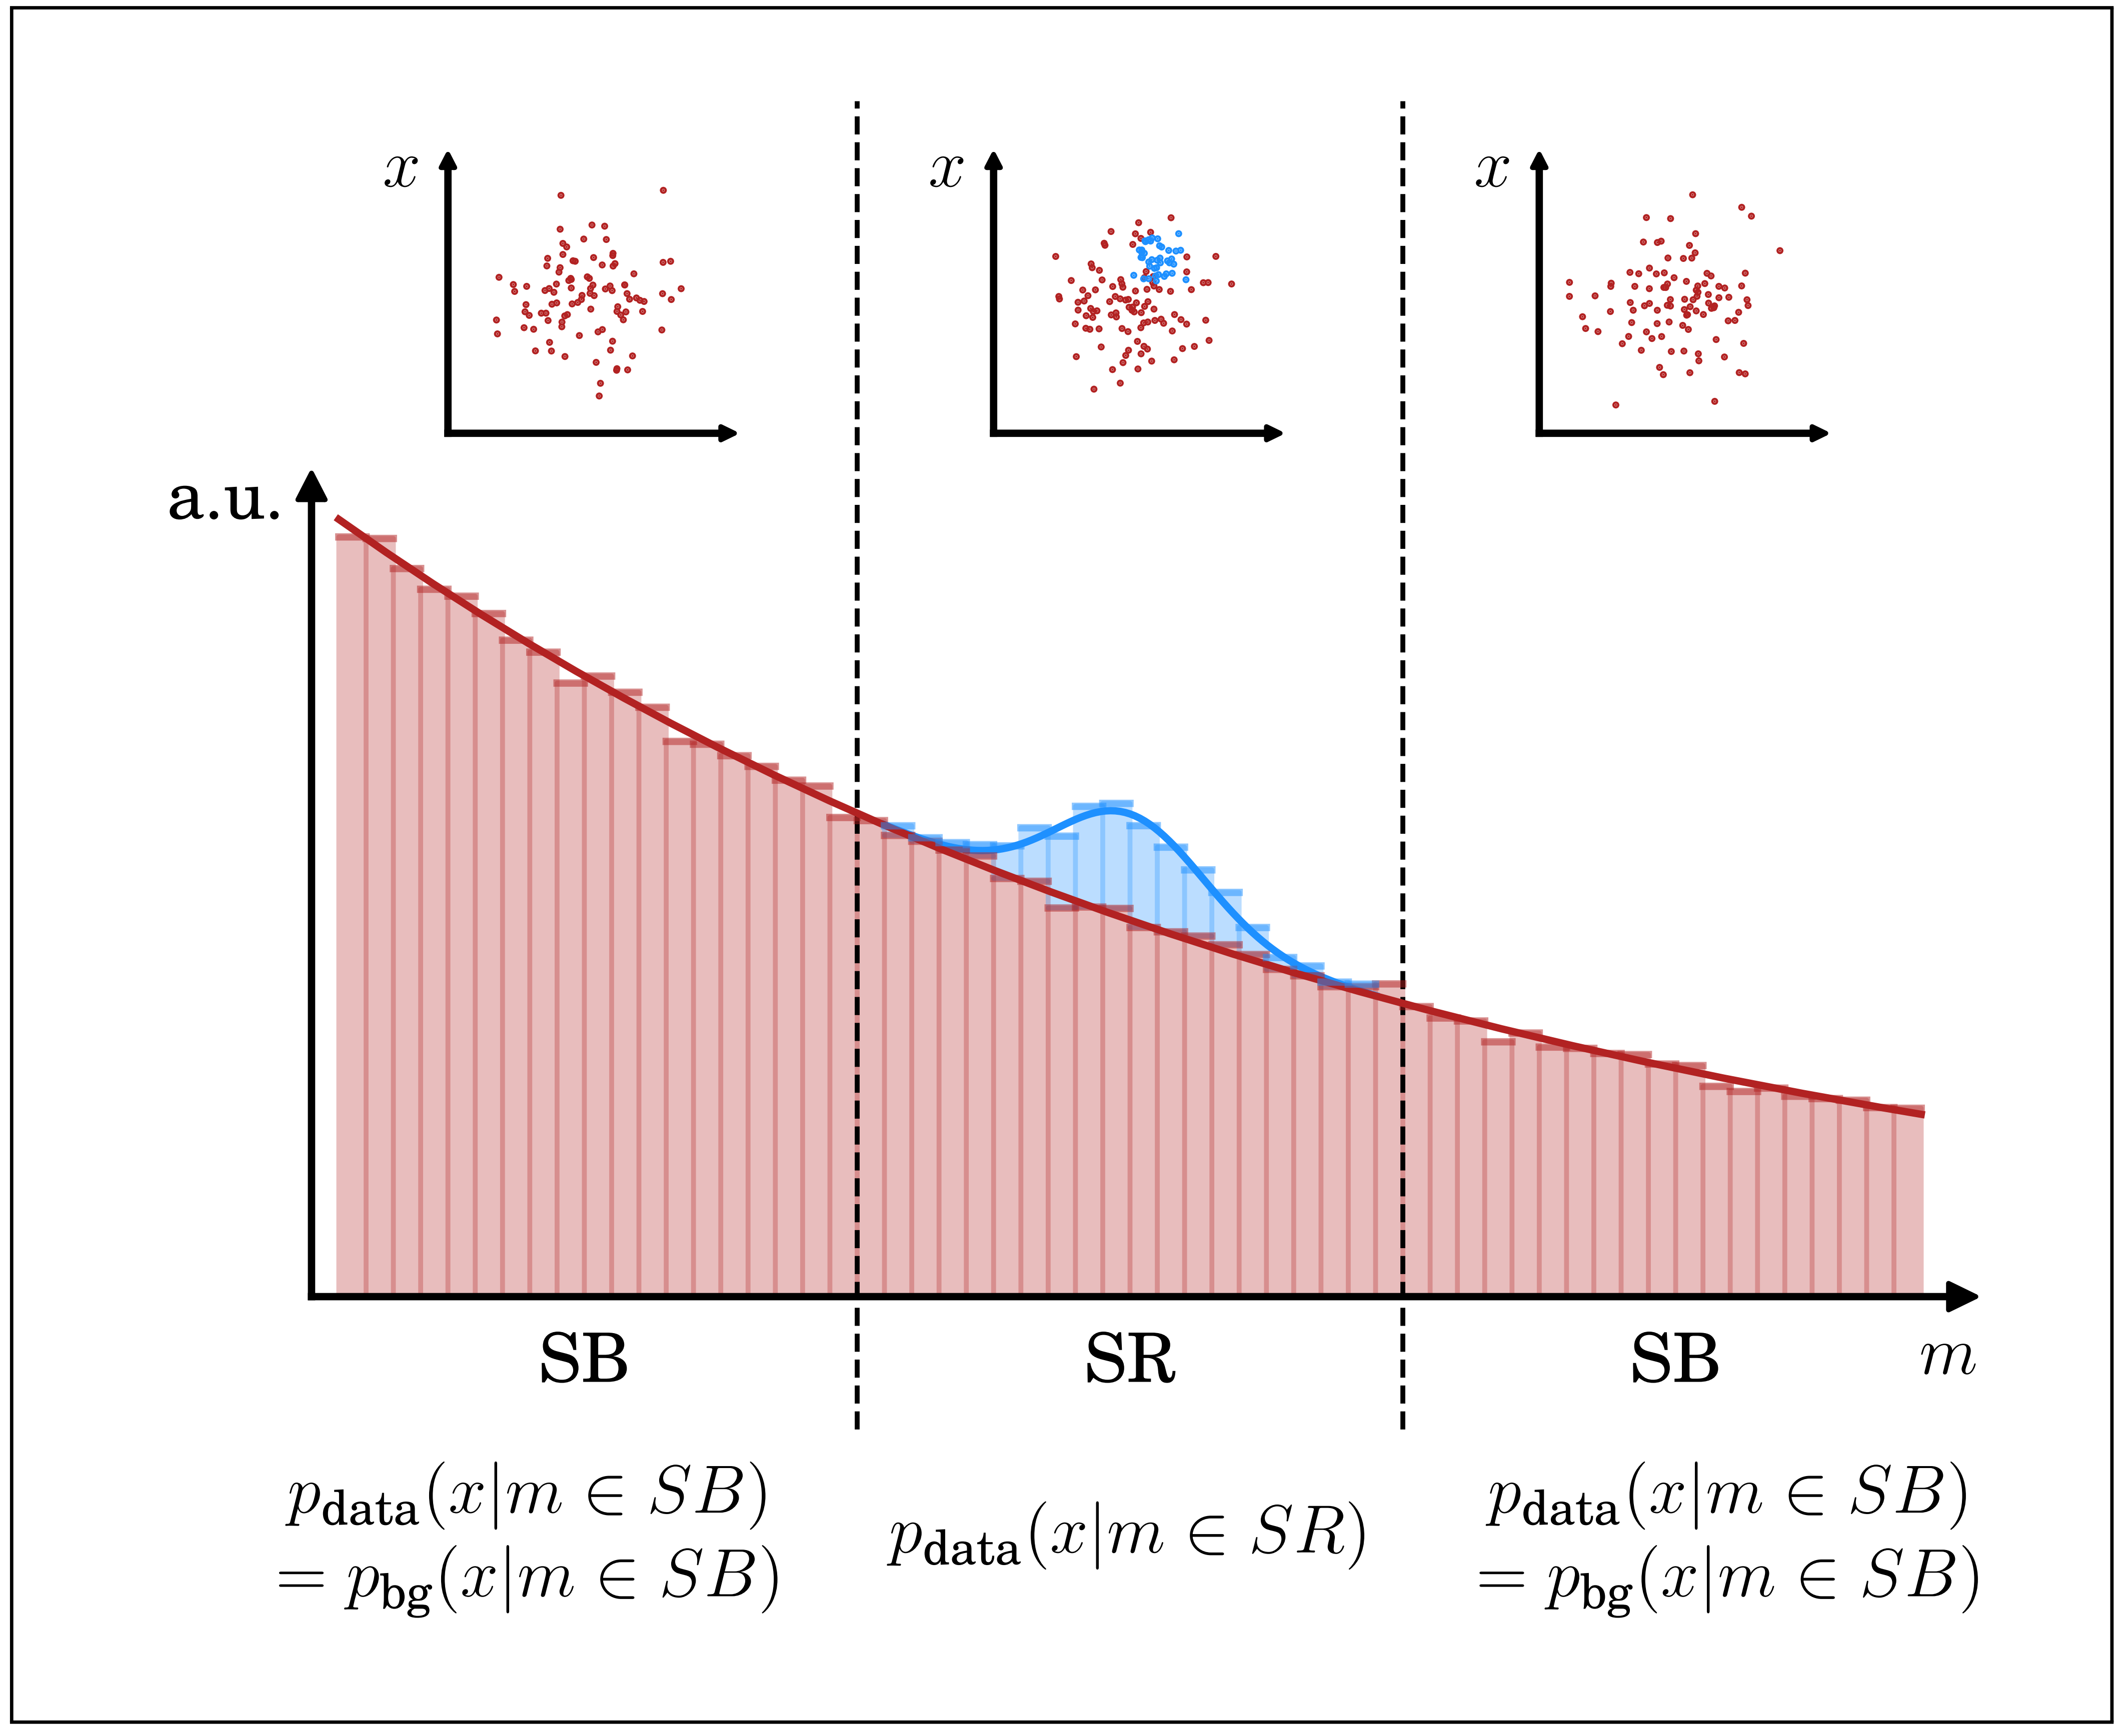

Idea: data vs sideband classifier

Collins, Howe & Nachman 1805.02664,1902.02634

from 2109.00546

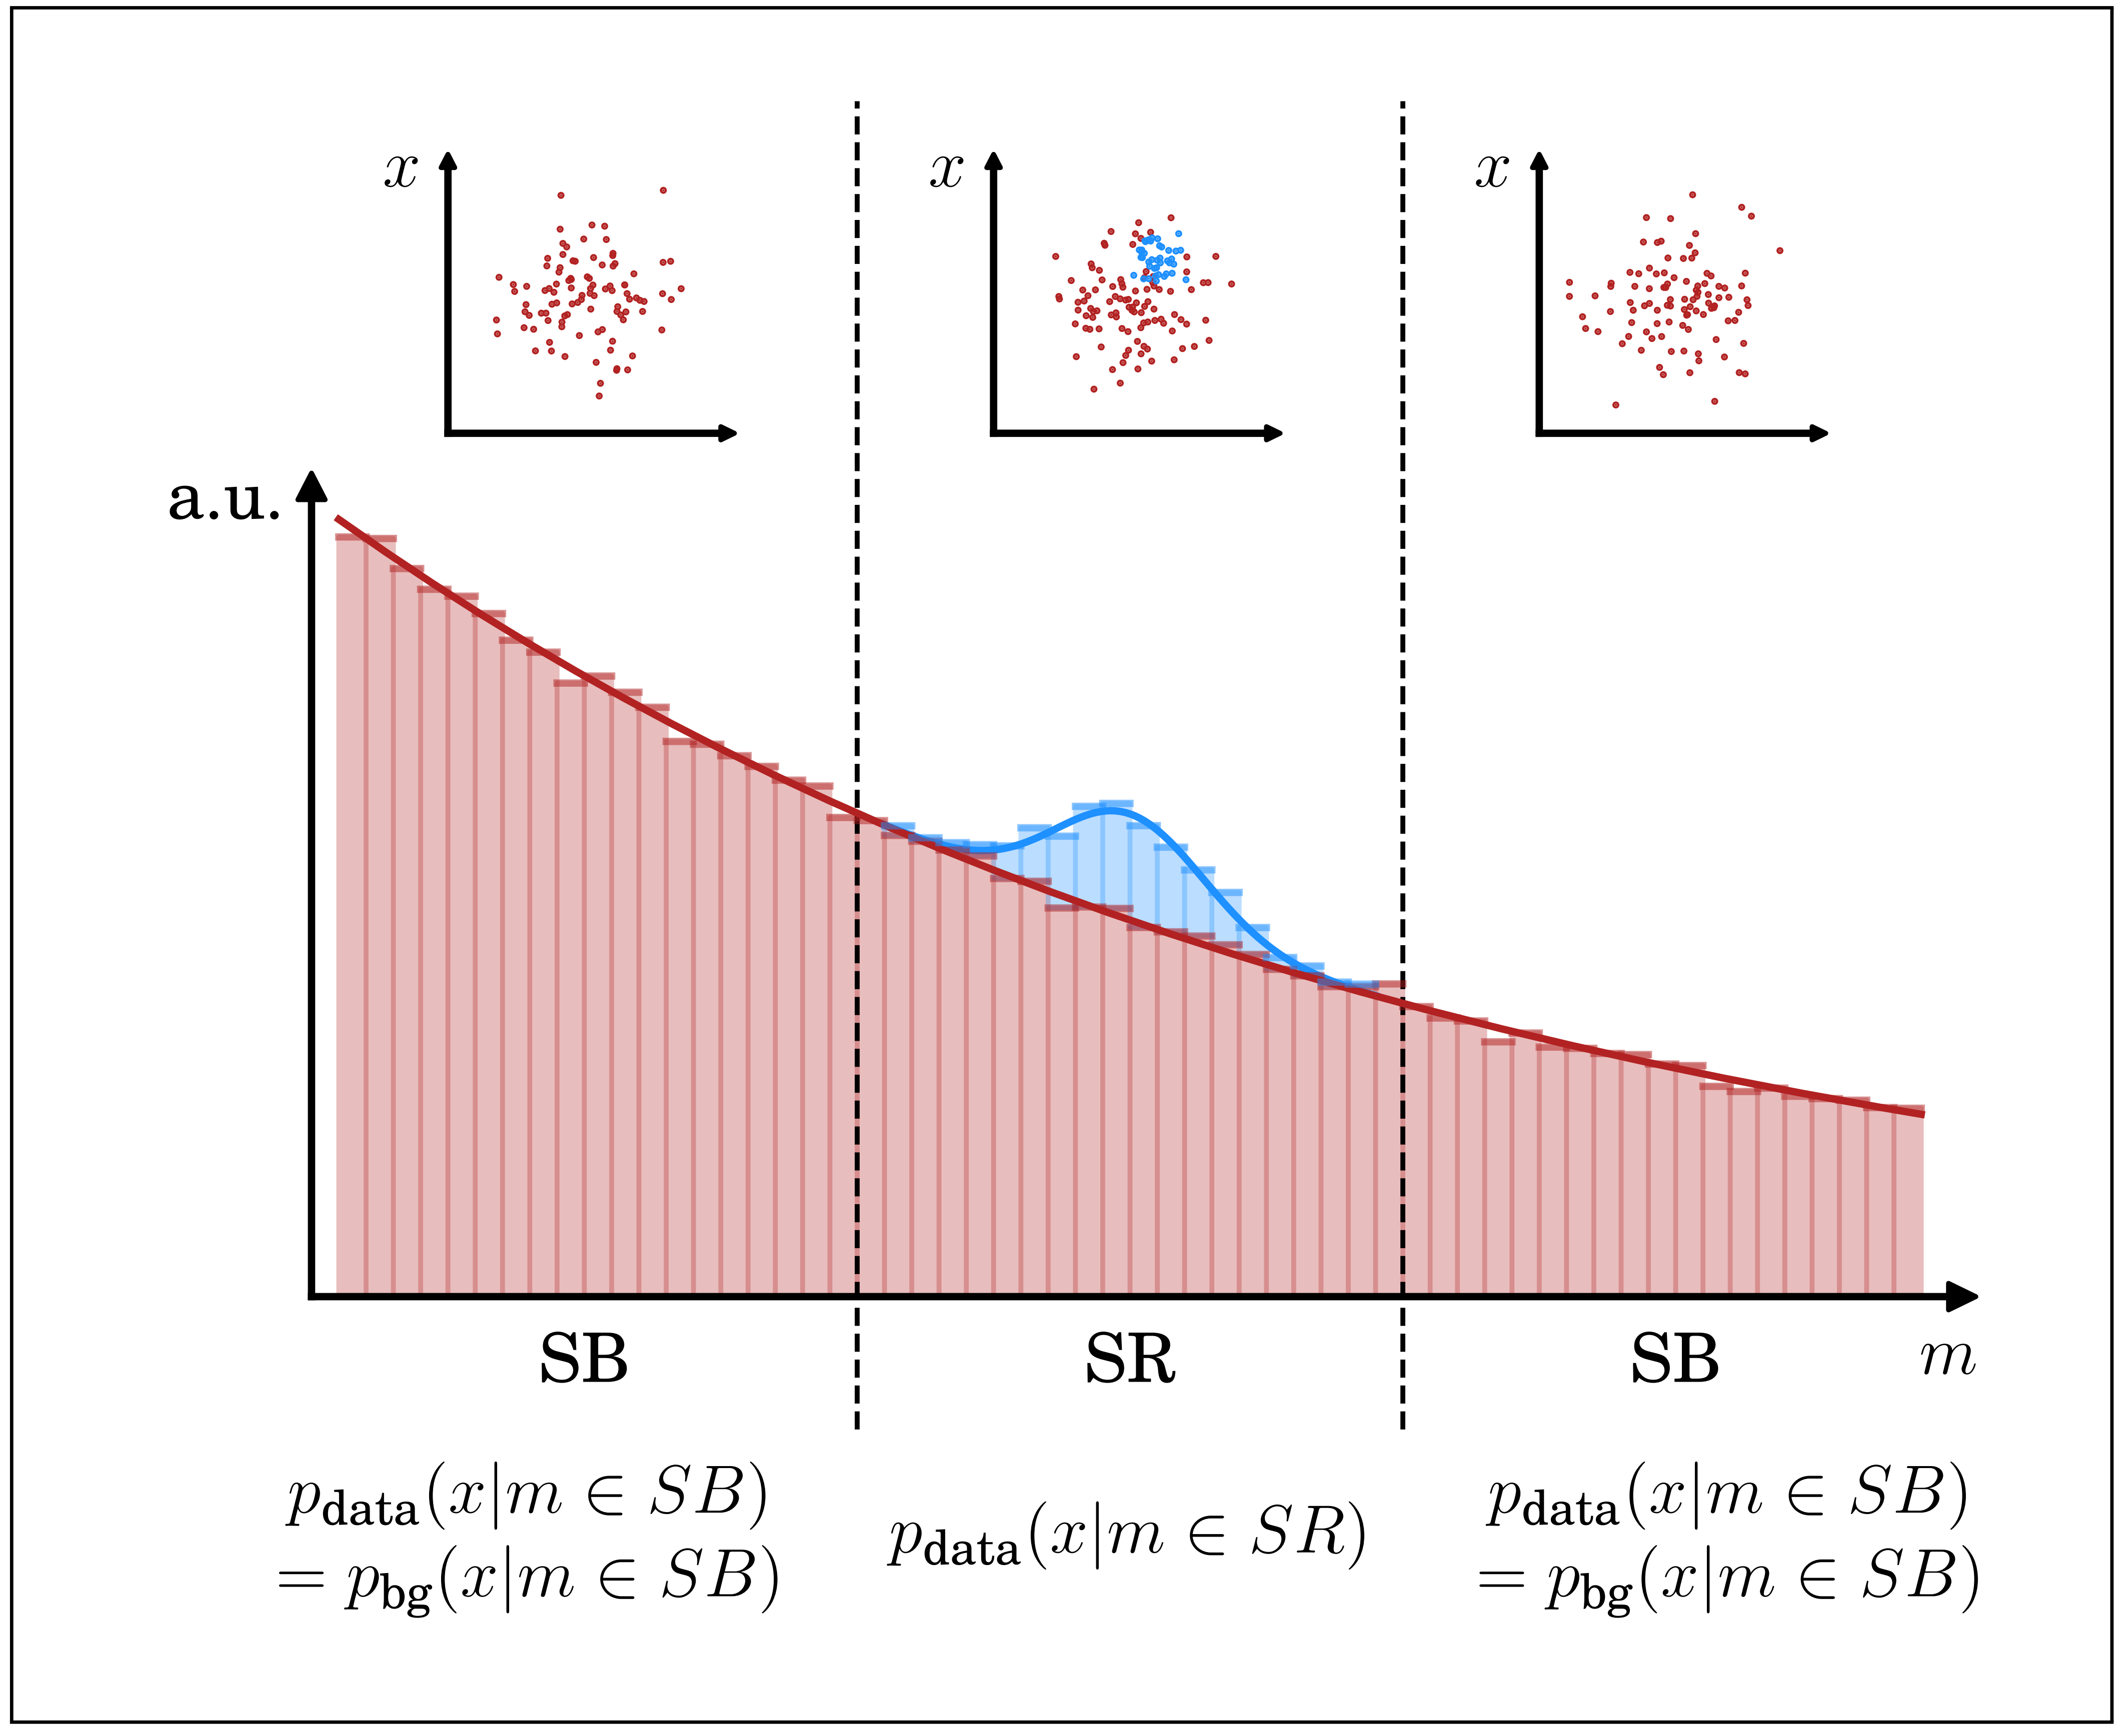

Train a NN classifier on SR vs SB data, learn

$$R_{classifier}(x) \approx \frac{p_{data,SR}(x)}{p_{data,SB}(x)} = \frac{p_{data,SR}(x)}{p_{bg,SB}(x)}$$

Idea: data vs sideband classifier

Collins, Howe & Nachman 1805.02664,1902.02634

from 2109.00546

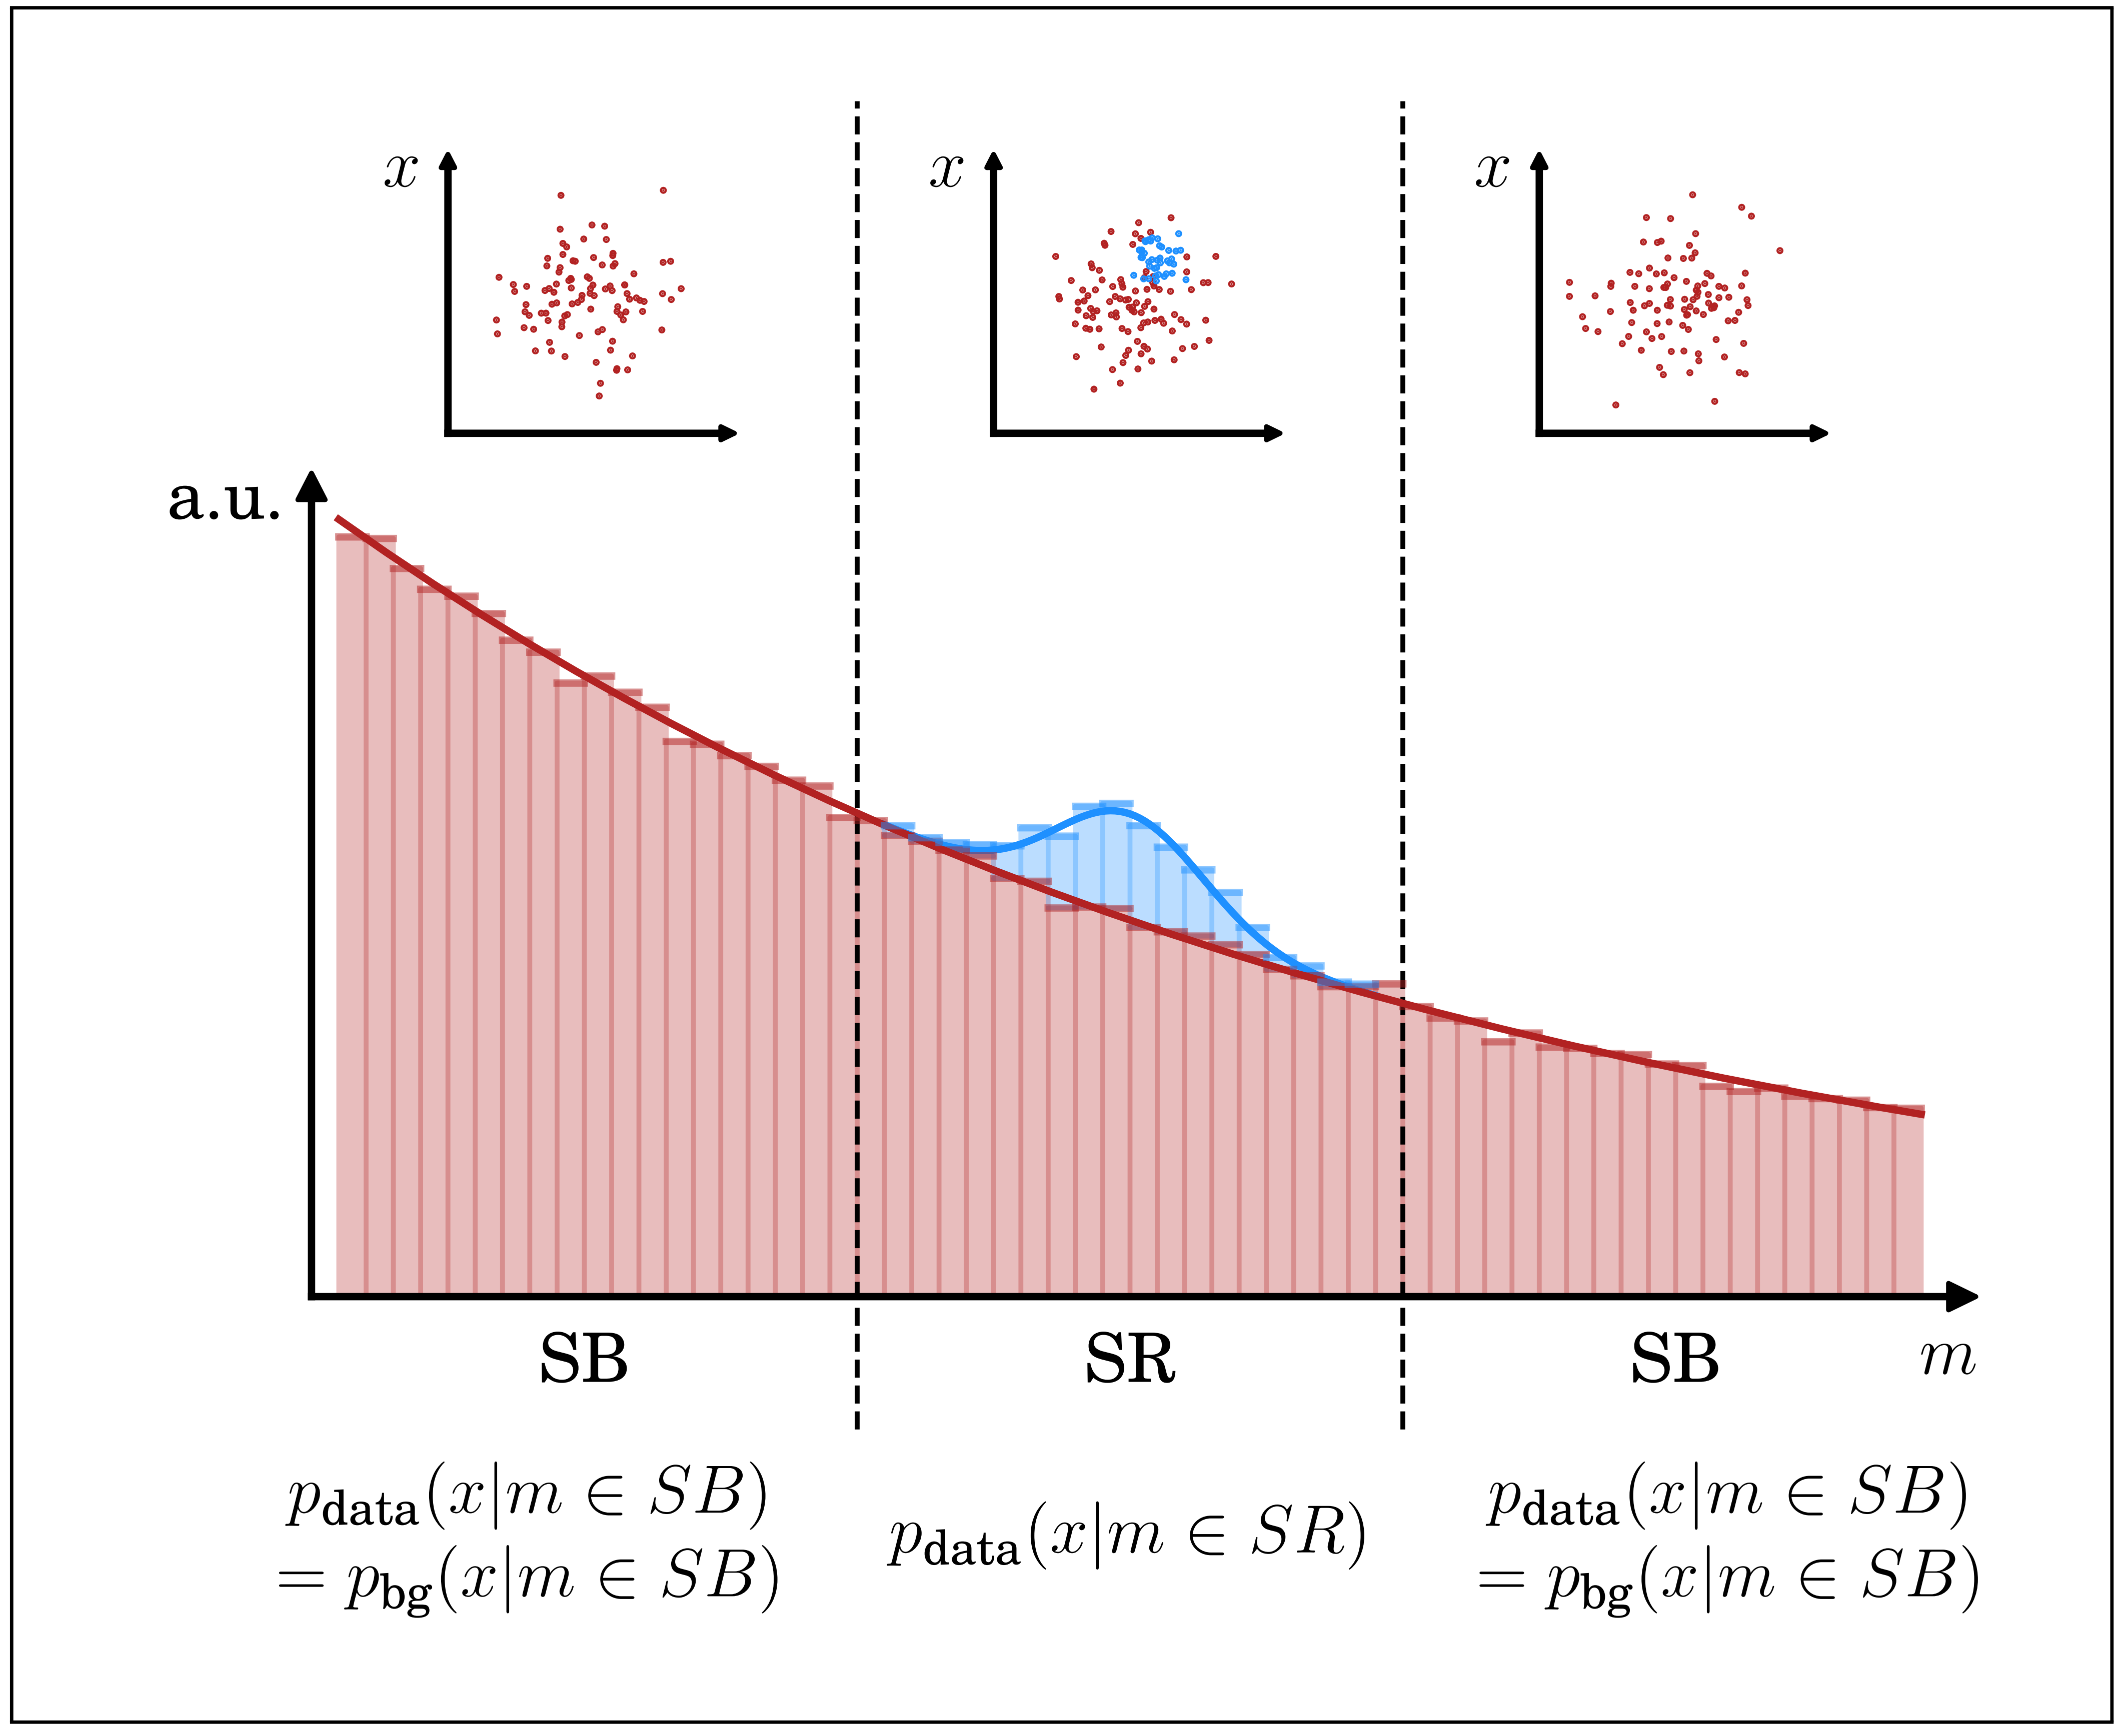

Train a NN classifier on SR vs SB data, learn

$$R_{classifier}(x) \approx \frac{p_{data,SR}(x)}{p_{data,SB}(x)} = \frac{p_{data,SR}(x)}{p_{bg,SB}(x)}$$

If $p_{bg,SB}(x) = p_{bg,SR}(x)$ [i.e. features x are independent of m in the background] then the classifier gives the desired likelihood ratio.

$$R_{classifier}(x) \to \frac{p_{data,SR}(x)}{p_{bg,SR}(x)}$$

Idea: data vs sideband classifier

Collins, Howe & Nachman 1805.02664,1902.02634

from 2109.00546

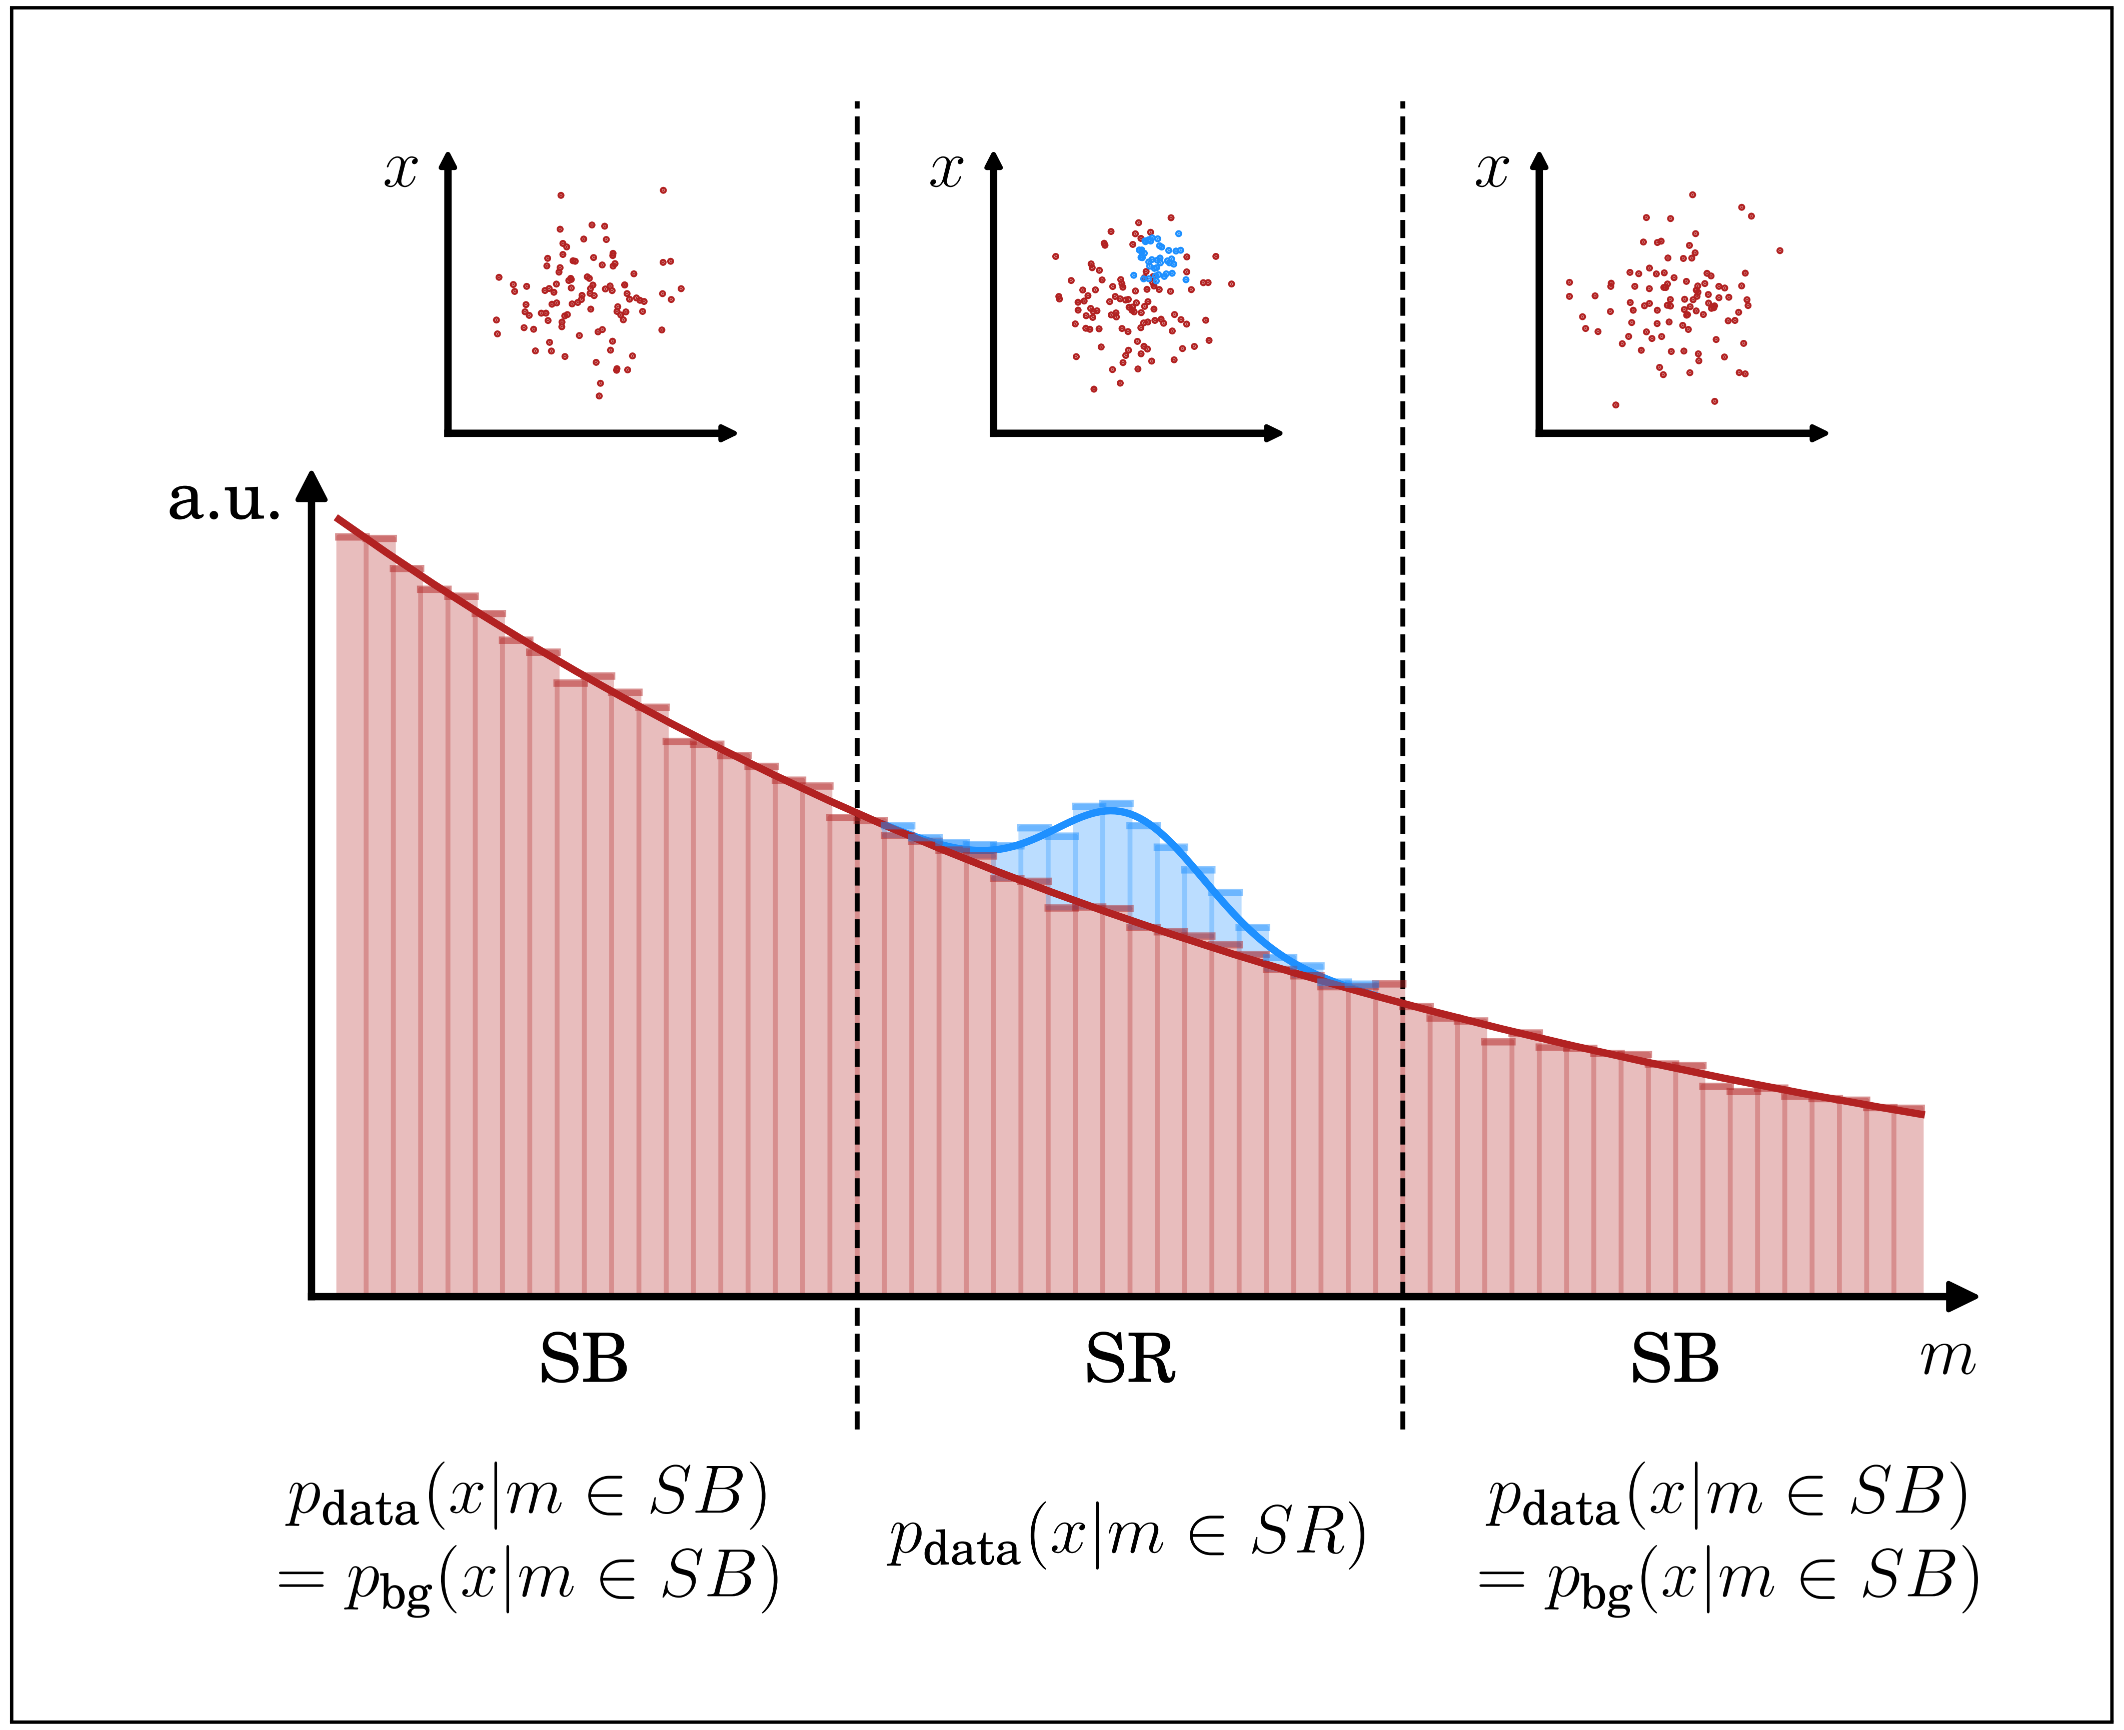

Train a NN classifier on SR vs SB data, learn

$$R_{classifier}(x) \approx \frac{p_{data,SR}(x)}{p_{data,SB}(x)} = \frac{p_{data,SR}(x)}{p_{bg,SB}(x)}$$

If $p_{bg,SB}(x) = p_{bg,SR}(x)$ [i.e. features x are independent of m in the background] then the classifier gives the desired likelihood ratio.

$$R_{classifier}(x) \to \frac{p_{data,SR}(x)}{p_{bg,SR}(x)}$$

"CWoLa Hunting"

Idea: conditional density estimation

Nachman & **DS** 2001.04990

Tight connection between ***generative models*** and ***anomaly detection***!

from 2109.00546

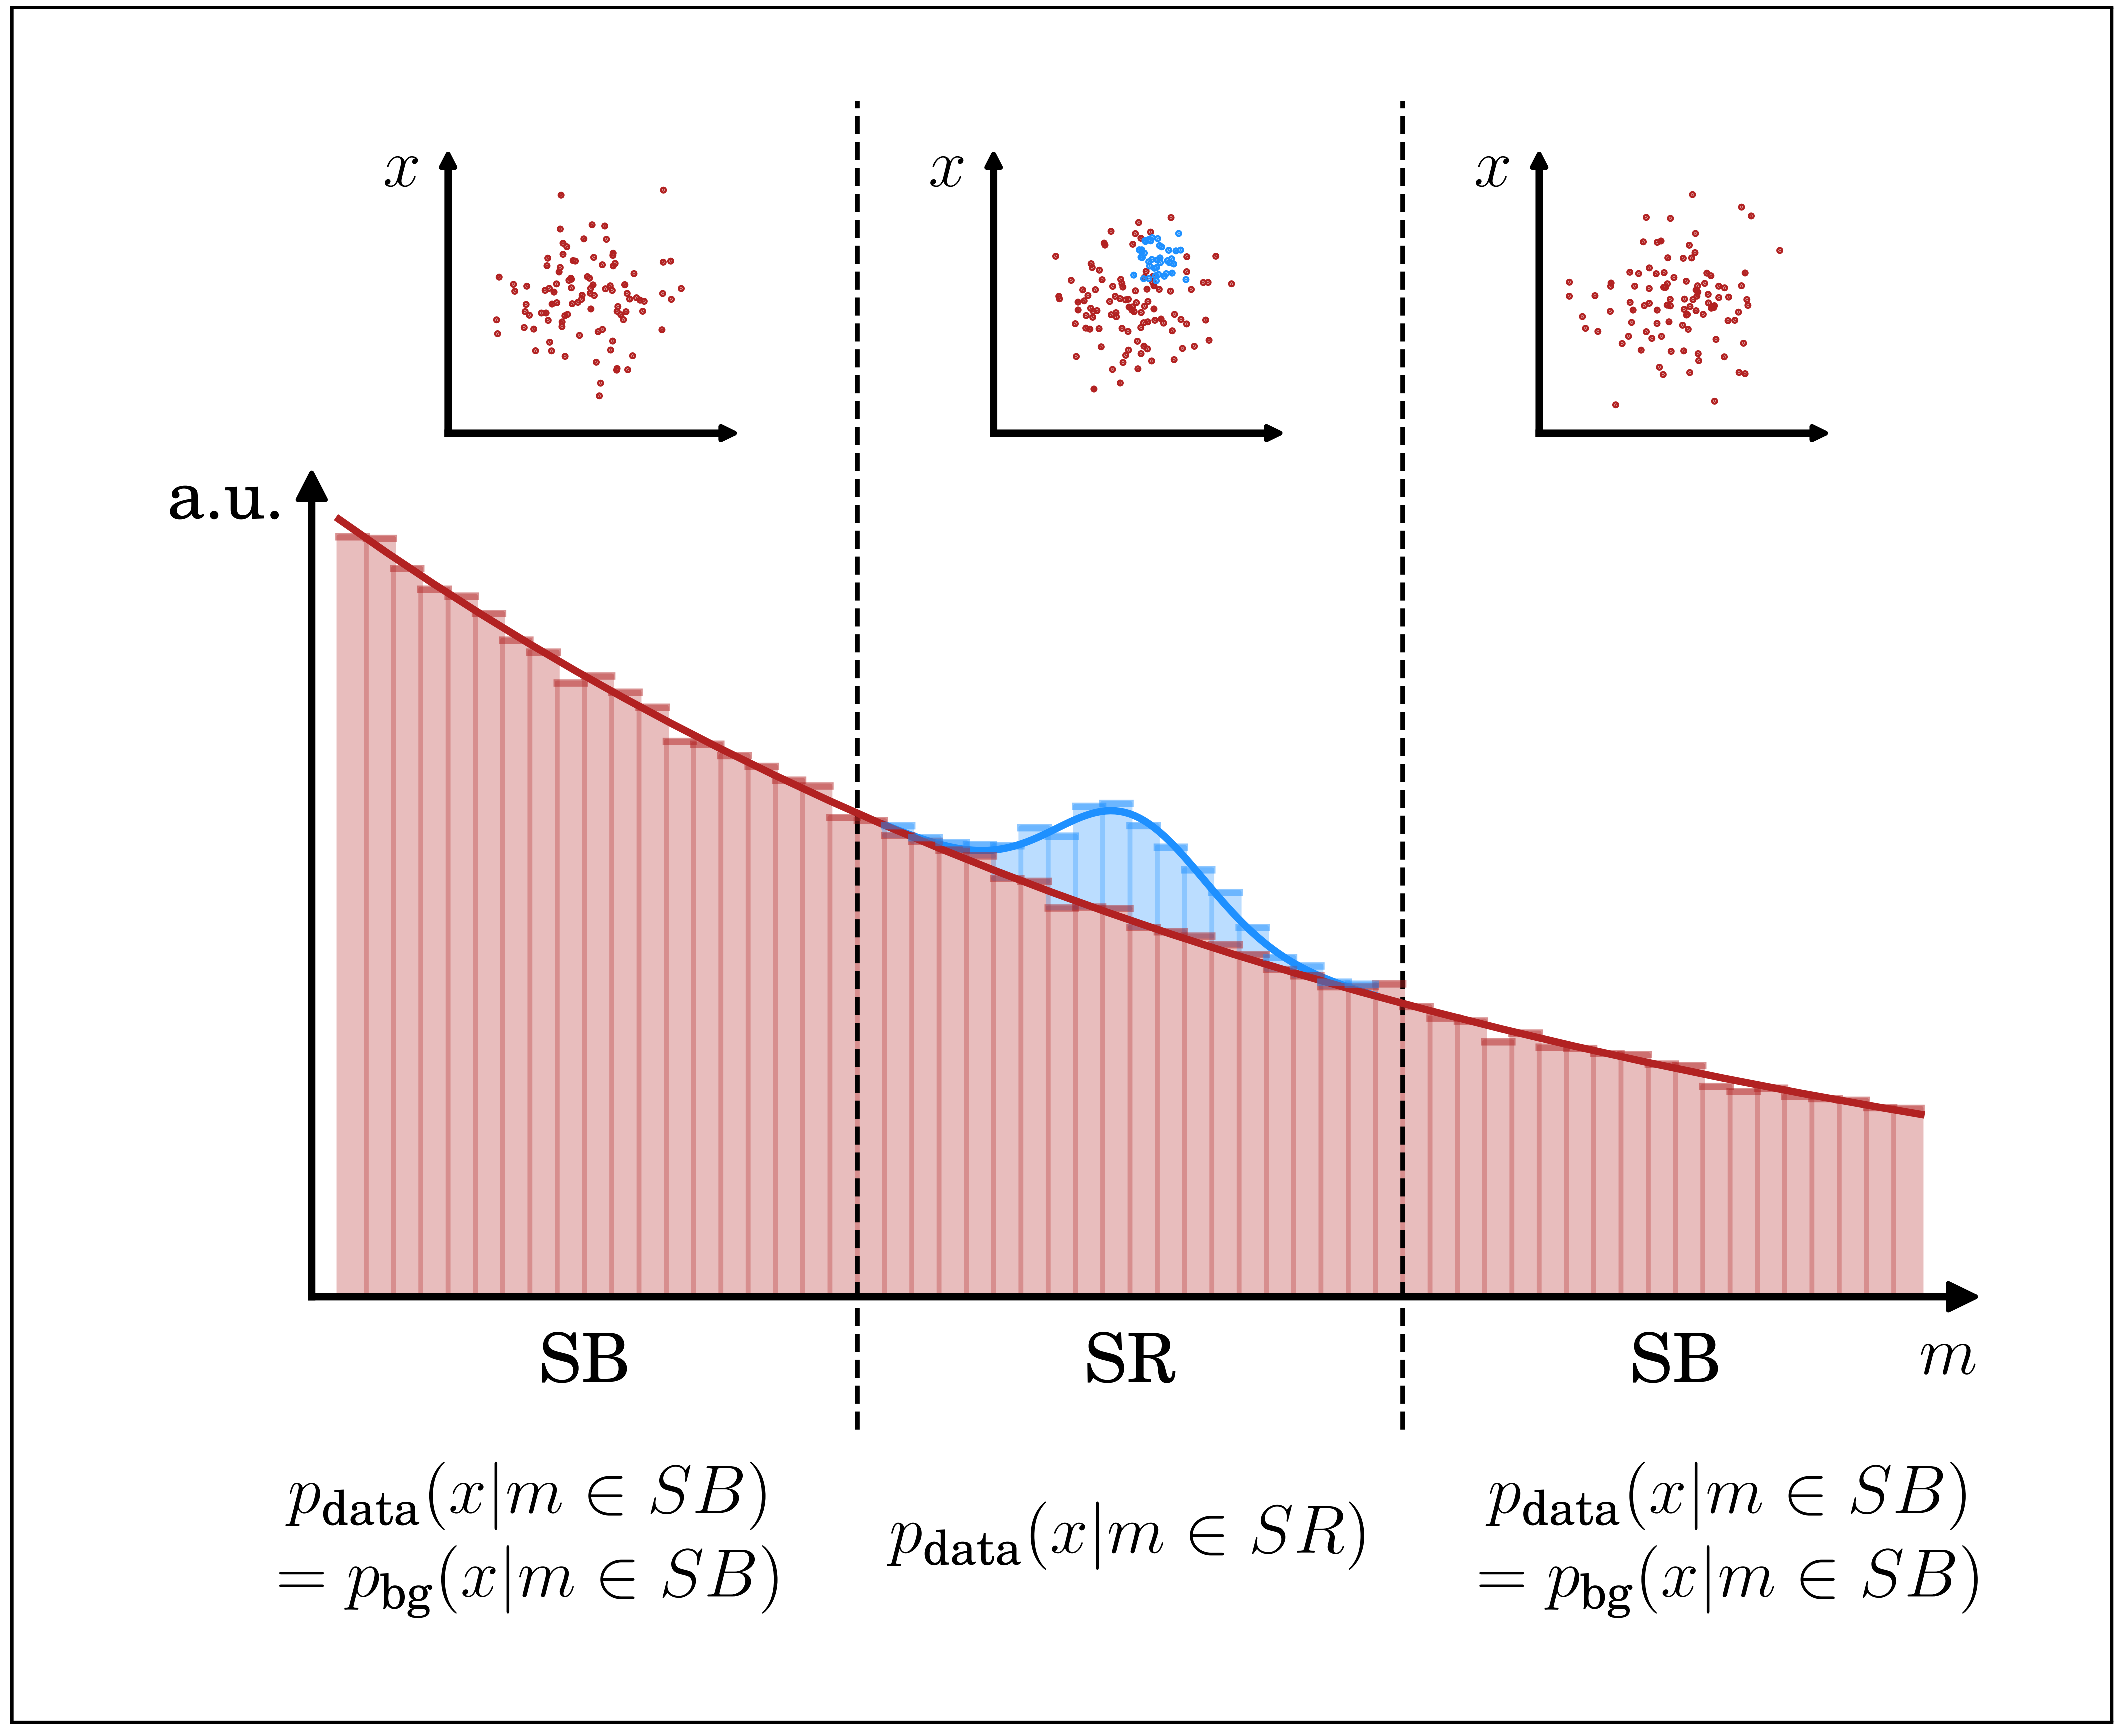

Idea: conditional density estimation

Tight connection between ***generative models*** and ***anomaly detection***!

Nachman & **DS** 2001.04990

from 2109.00546

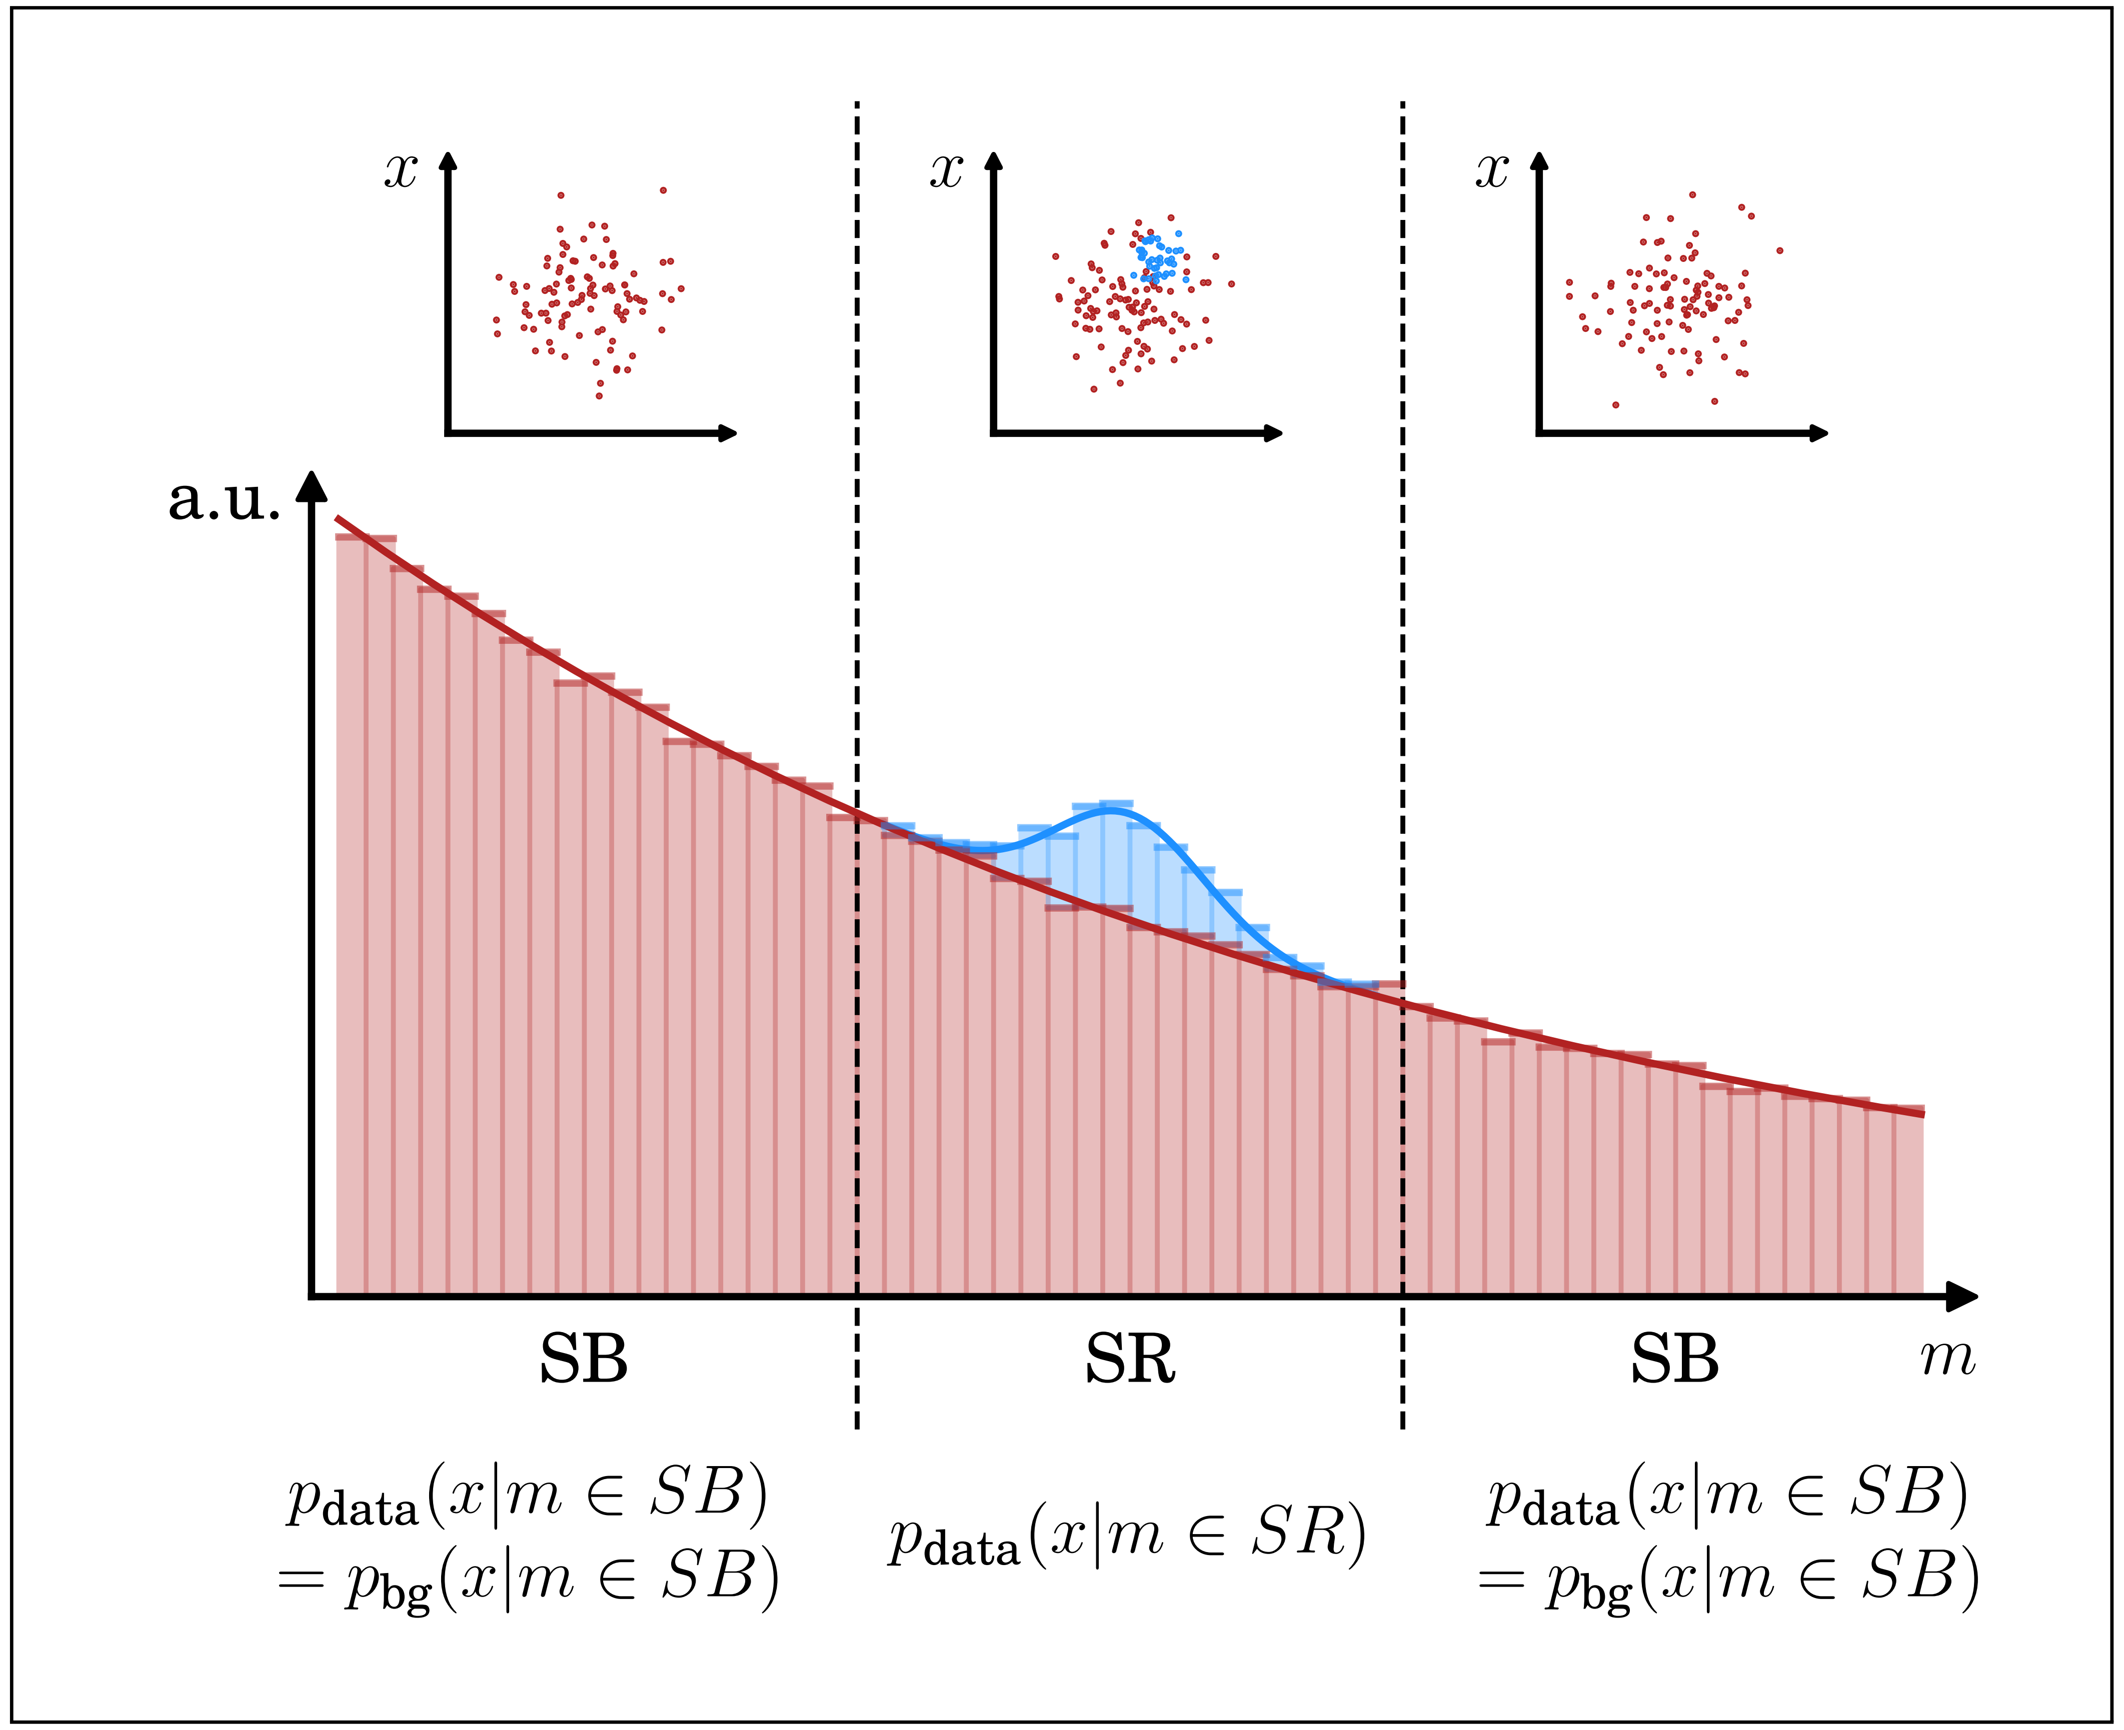

Train two separate ***normalizing flows*** on SR and SB events to learn $p_{data}(x\,|\,m \in SR)$ and $p_{data}(x\,|\,m \in SB) = p_{bg}(x\,|\,m \in SB)$.

Idea: conditional density estimation

Tight connection between ***generative models*** and ***anomaly detection***!

Nachman & **DS** 2001.04990

from 2109.00546

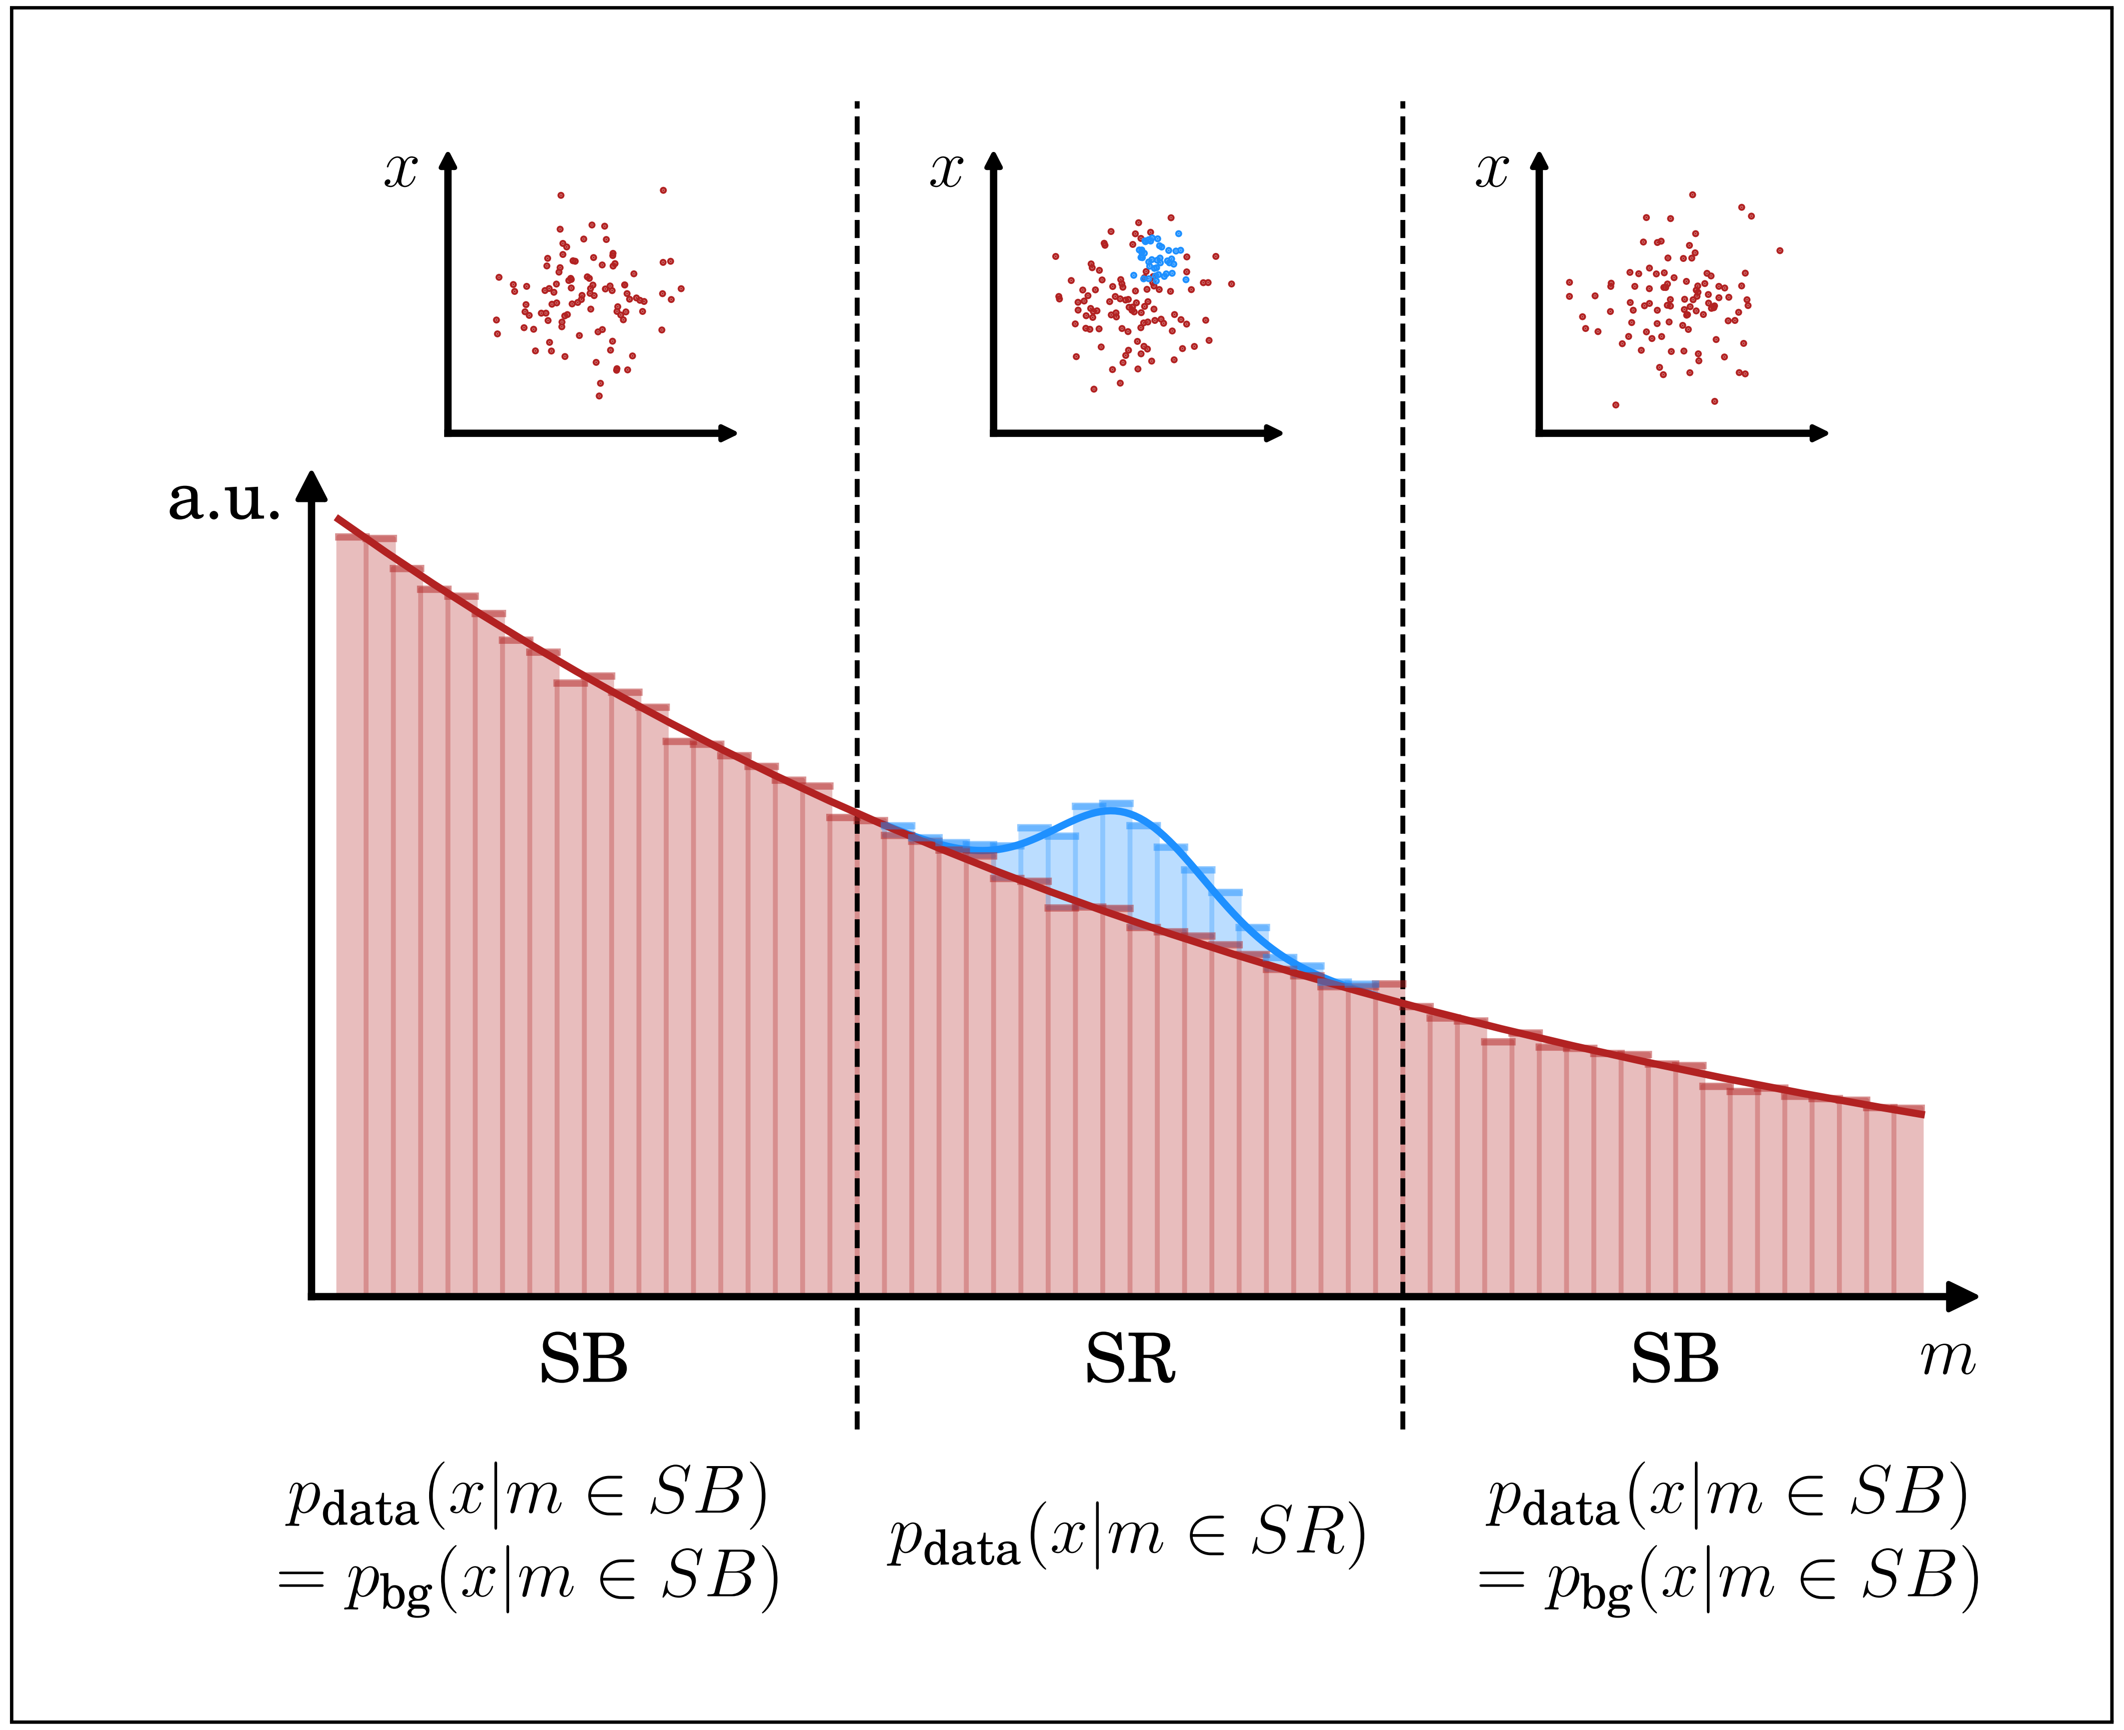

Train two separate ***normalizing flows*** on SR and SB events to learn $p_{data}(x\,|\,m \in SR)$ and $p_{data}(x\,|\,m \in SB) = p_{bg}(x\,|\,m \in SB)$.

The SB NF automatically interpolates into the SR, giving an estimate of $p_{bg}(x\,|\,m \in SR)$.

Idea: conditional density estimation

Nachman & **DS** 2001.04990

Tight connection between ***generative models*** and ***anomaly detection***!

from 2109.00546

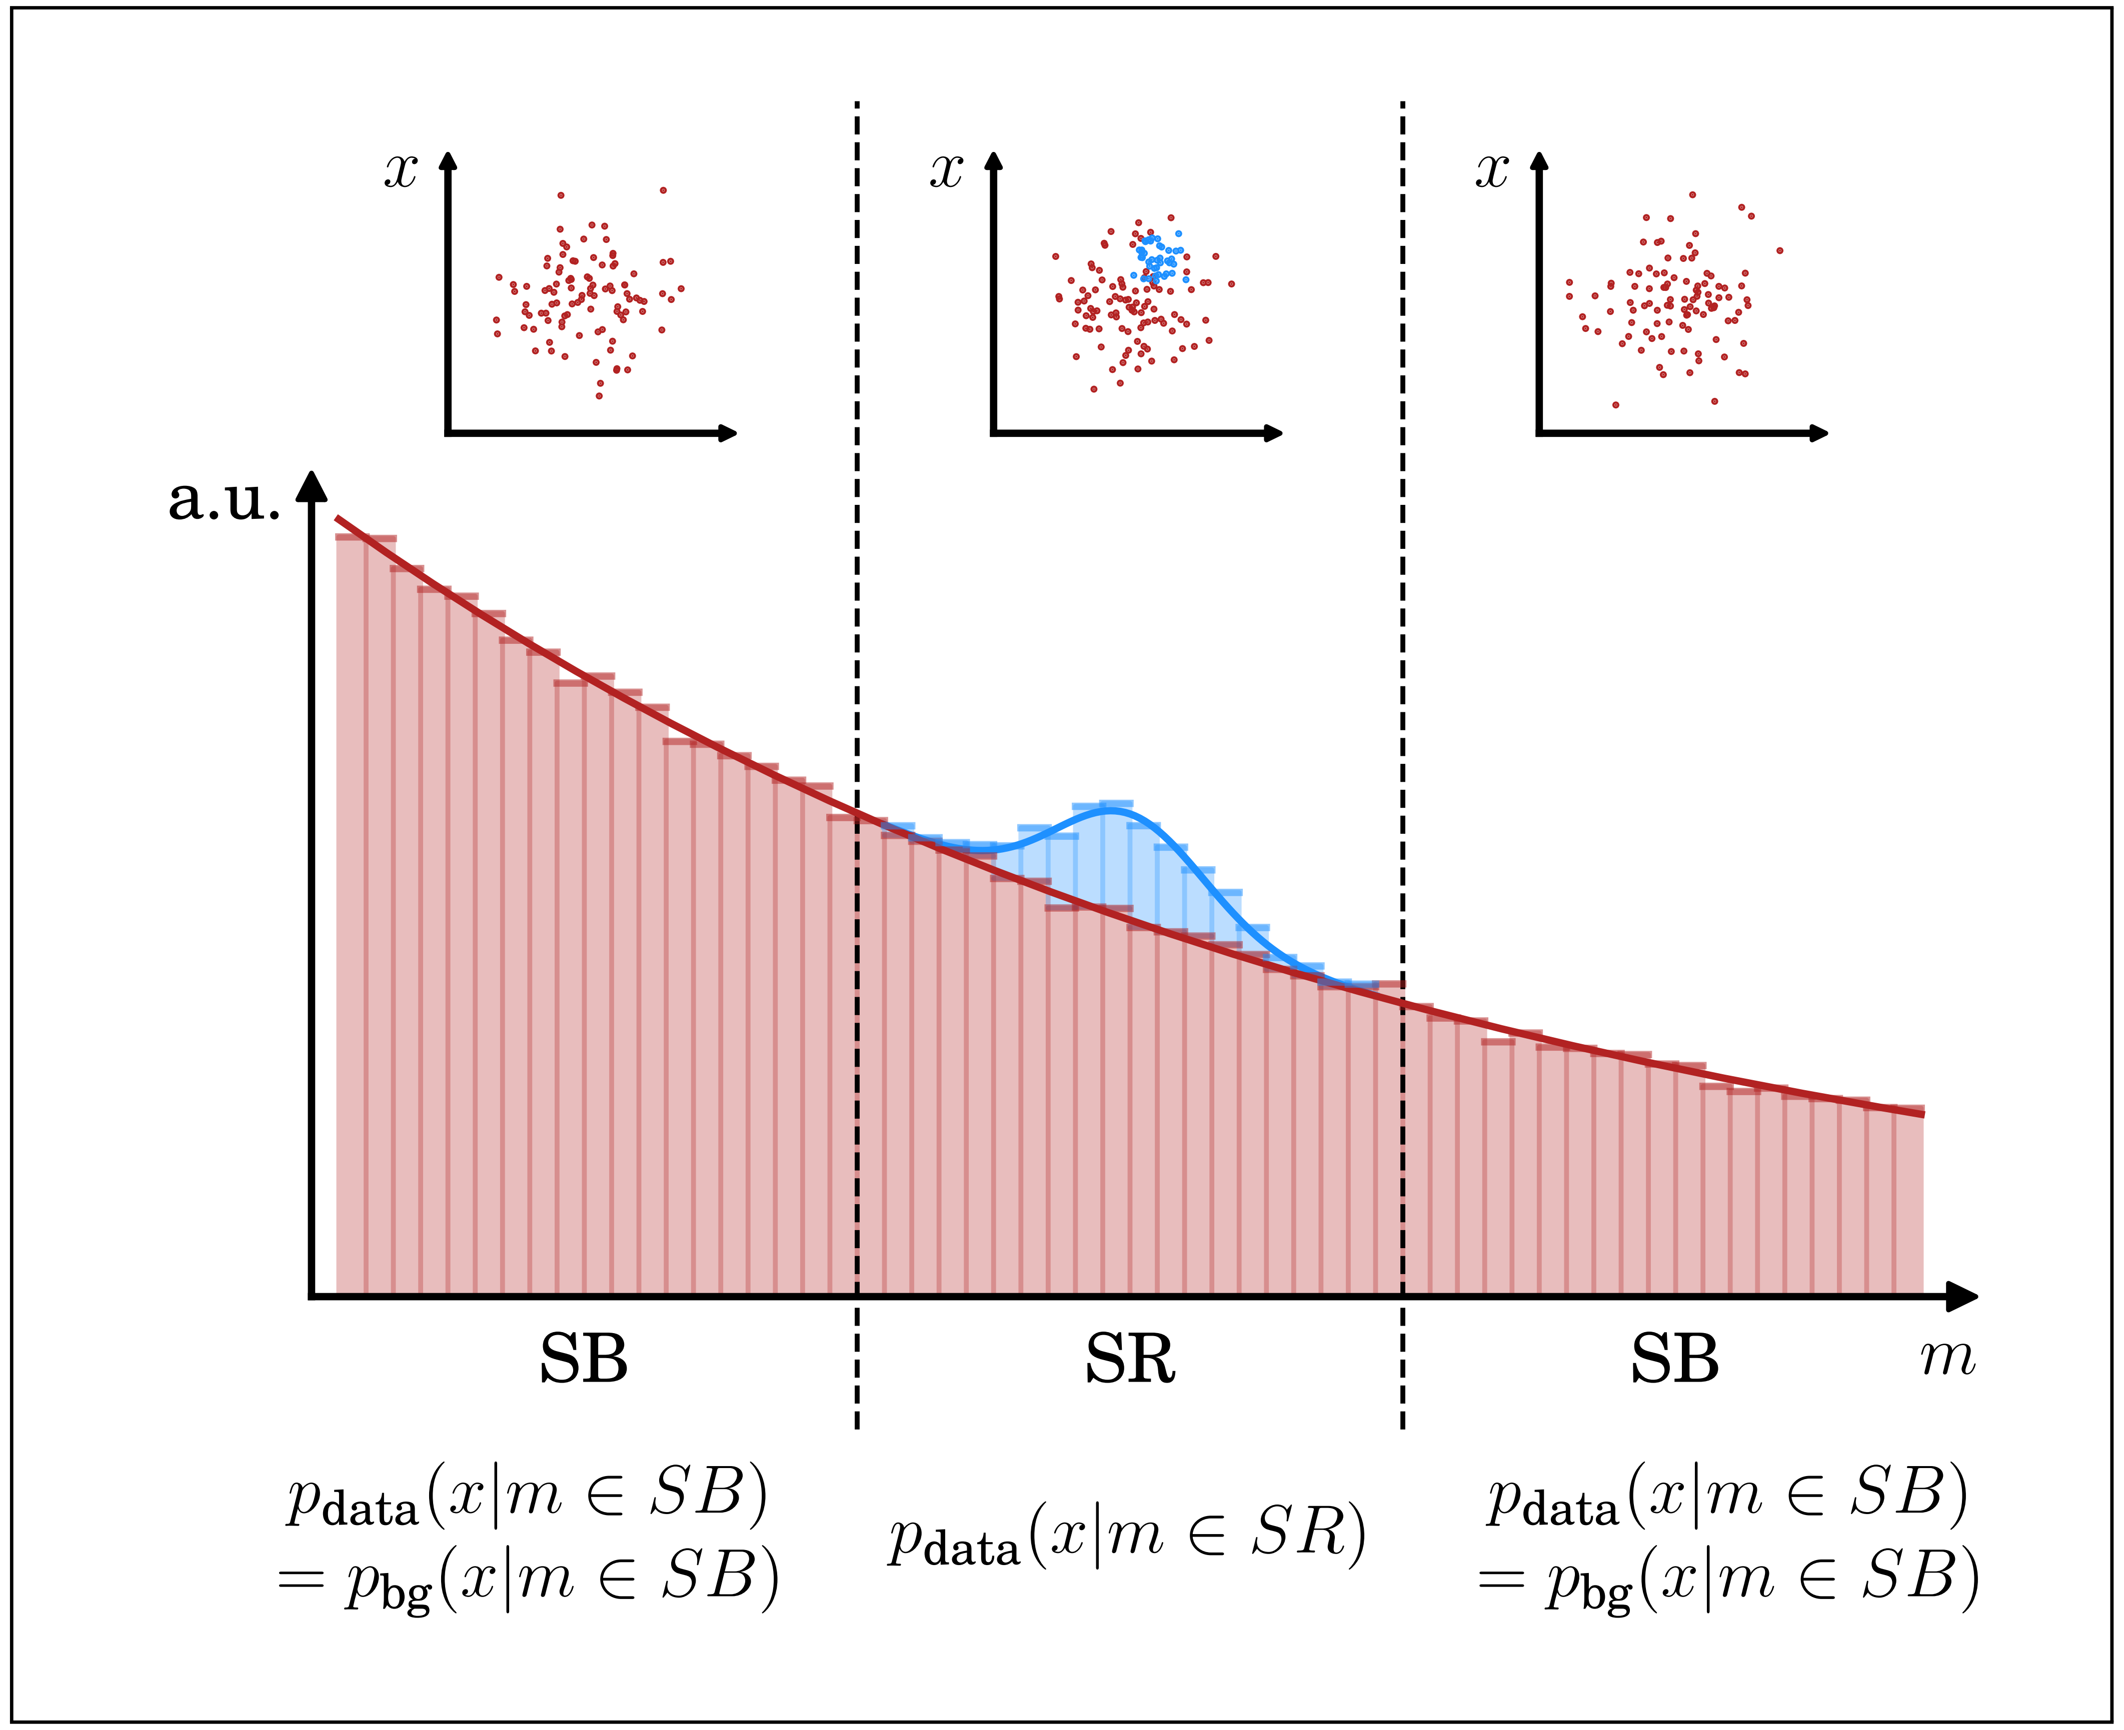

Train two separate ***normalizing flows*** on SR and SB events to learn $p_{data}(x\,|\,m \in SR)$ and $p_{data}(x\,|\,m \in SB) = p_{bg}(x\,|\,m \in SB)$.

The SB NF automatically interpolates into the SR, giving an estimate of $p_{bg}(x\,|\,m \in SR)$.

Construct likelihood ratio explicitly.

Idea: conditional density estimation

Tight connection between ***generative models*** and ***anomaly detection***!

Nachman & **DS** 2001.04990

from 2109.00546

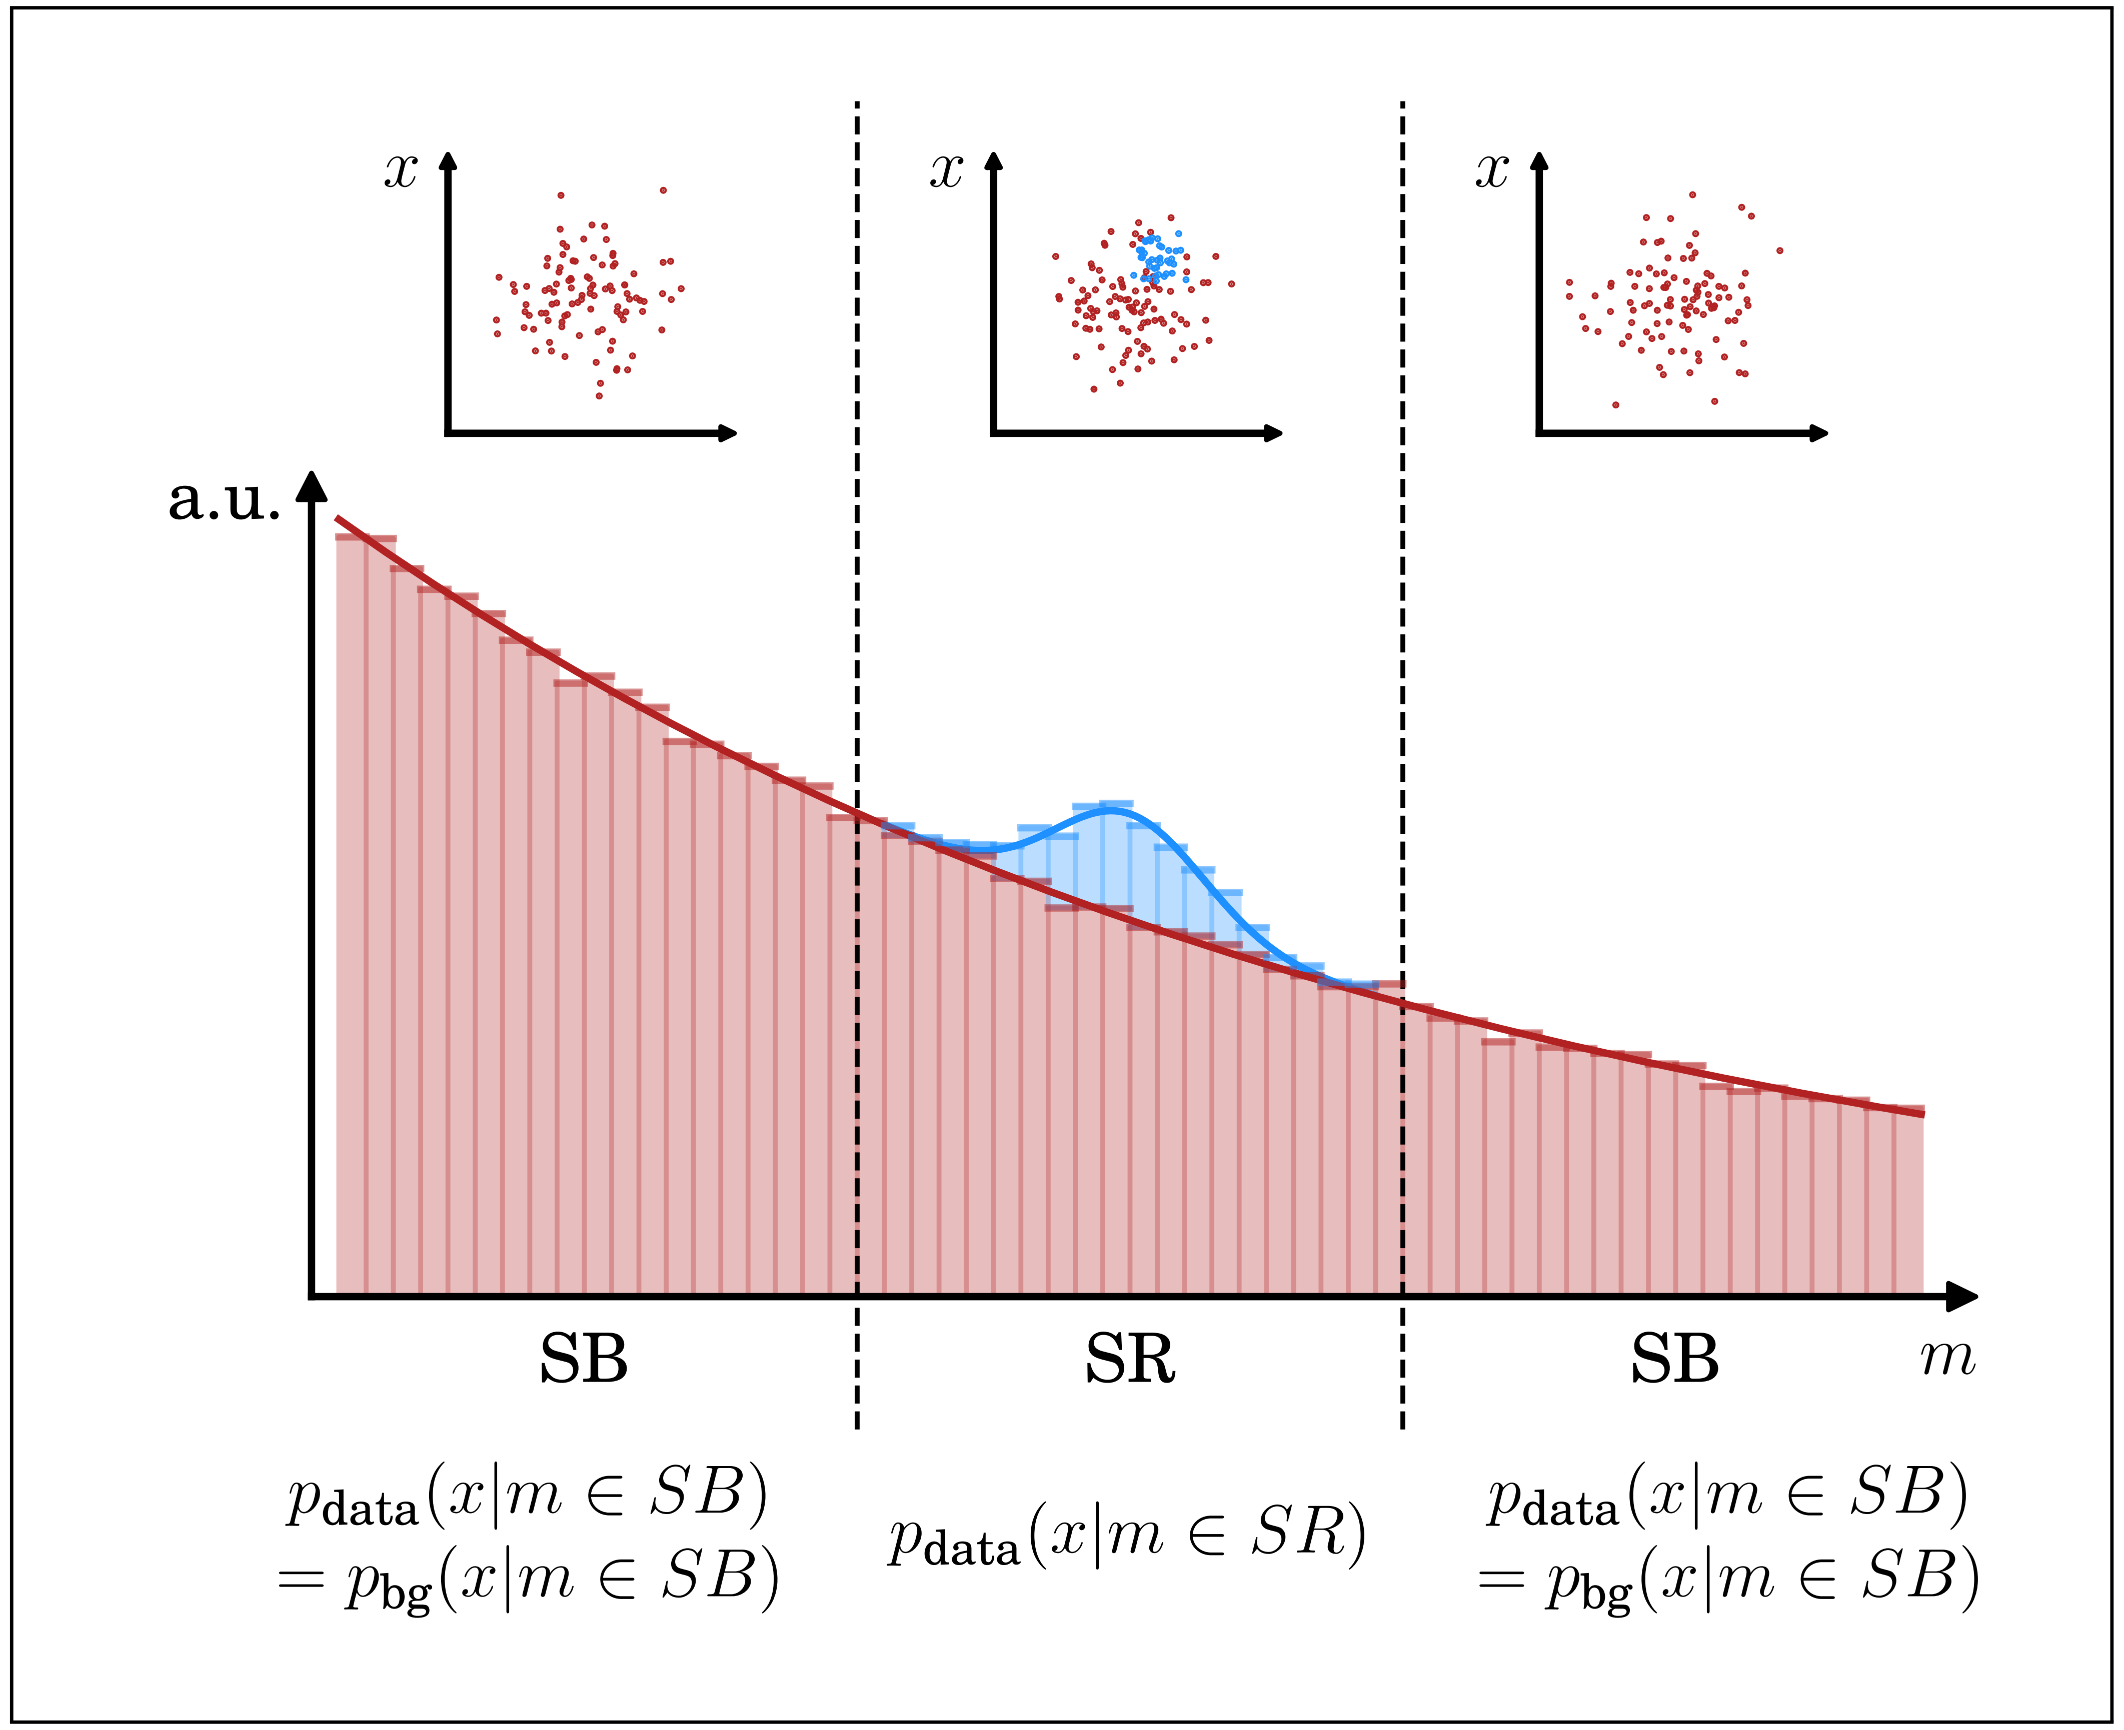

Train two separate ***normalizing flows*** on SR and SB events to learn $p_{data}(x\,|\,m \in SR)$ and $p_{data}(x\,|\,m \in SB) = p_{bg}(x\,|\,m \in SB)$.

The SB NF automatically interpolates into the SR, giving an estimate of $p_{bg}(x\,|\,m \in SR)$.

Construct likelihood ratio explicitly.

- Pros: robust against correlations!

Idea: conditional density estimation

Tight connection between ***generative models*** and ***anomaly detection***!

Nachman & **DS** 2001.04990

from 2109.00546

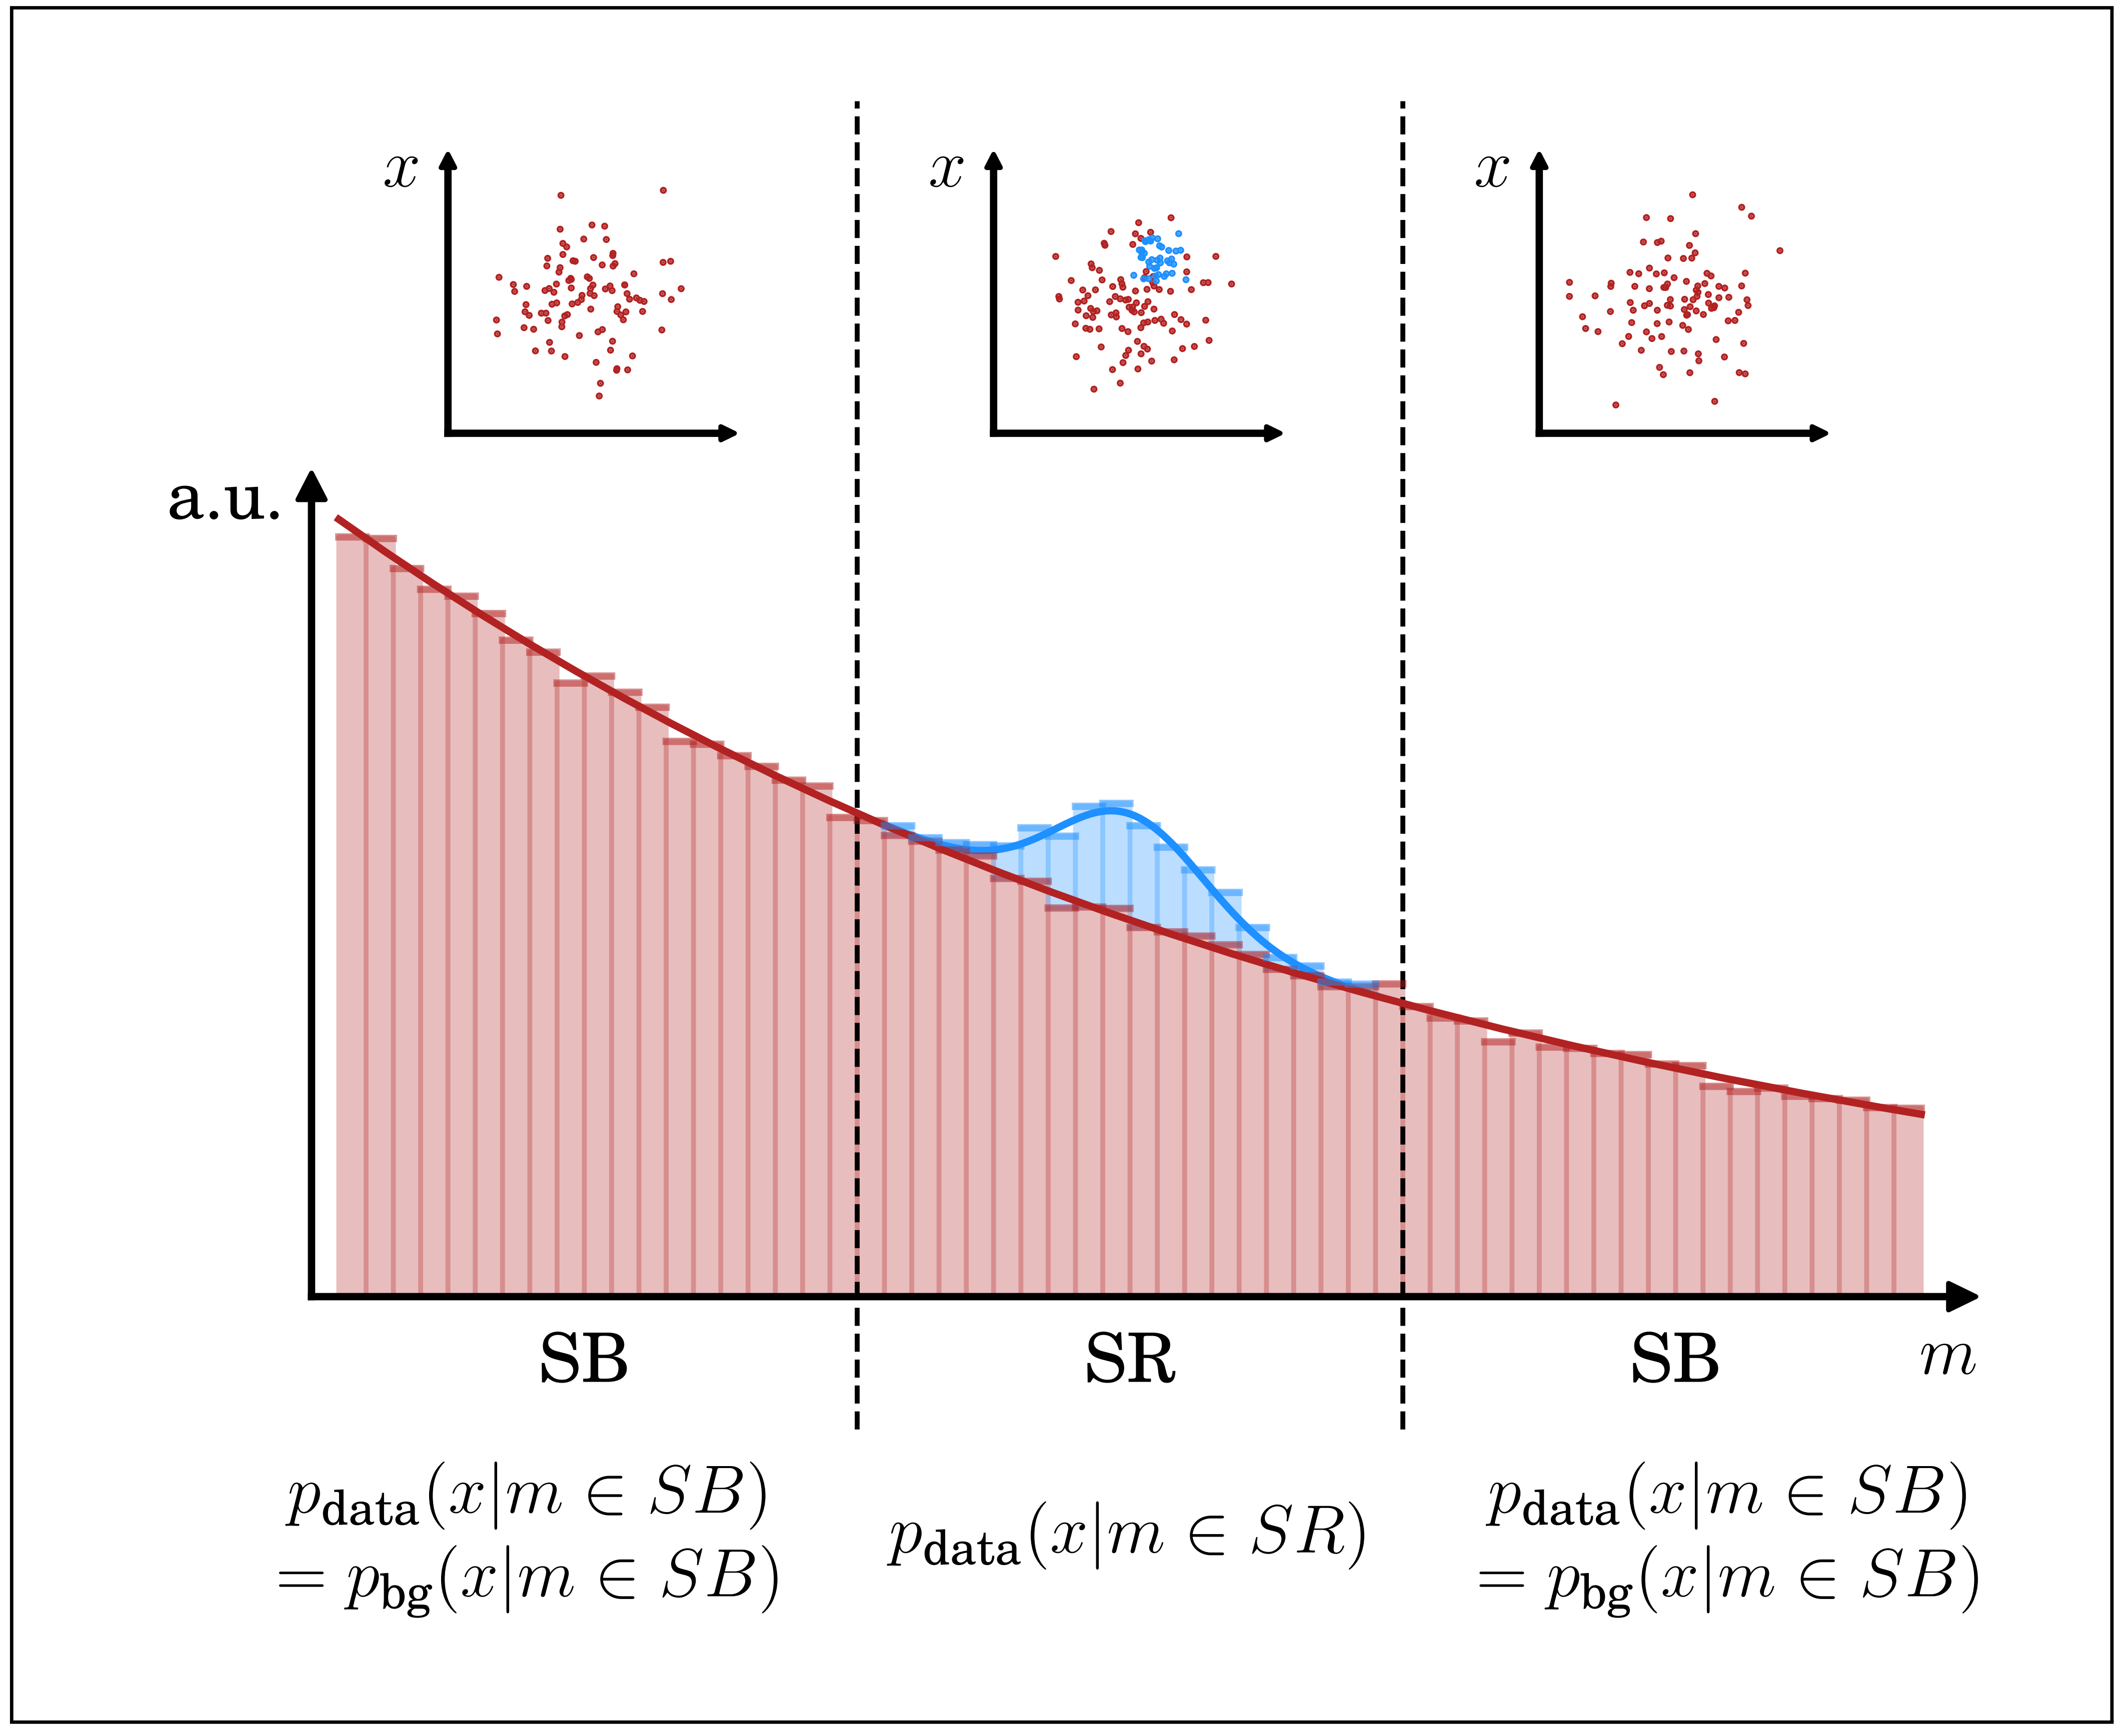

Train two separate ***normalizing flows*** on SR and SB events to learn $p_{data}(x\,|\,m \in SR)$ and $p_{data}(x\,|\,m \in SB) = p_{bg}(x\,|\,m \in SB)$.

The SB NF automatically interpolates into the SR, giving an estimate of $p_{bg}(x\,|\,m \in SR)$.

Construct likelihood ratio explicitly.

- Pros: robust against correlations!
- Cons: density estimation much harder than classification

Idea: conditional density estimation

Tight connection between ***generative models*** and ***anomaly detection***!

Nachman & **DS** 2001.04990

from 2109.00546

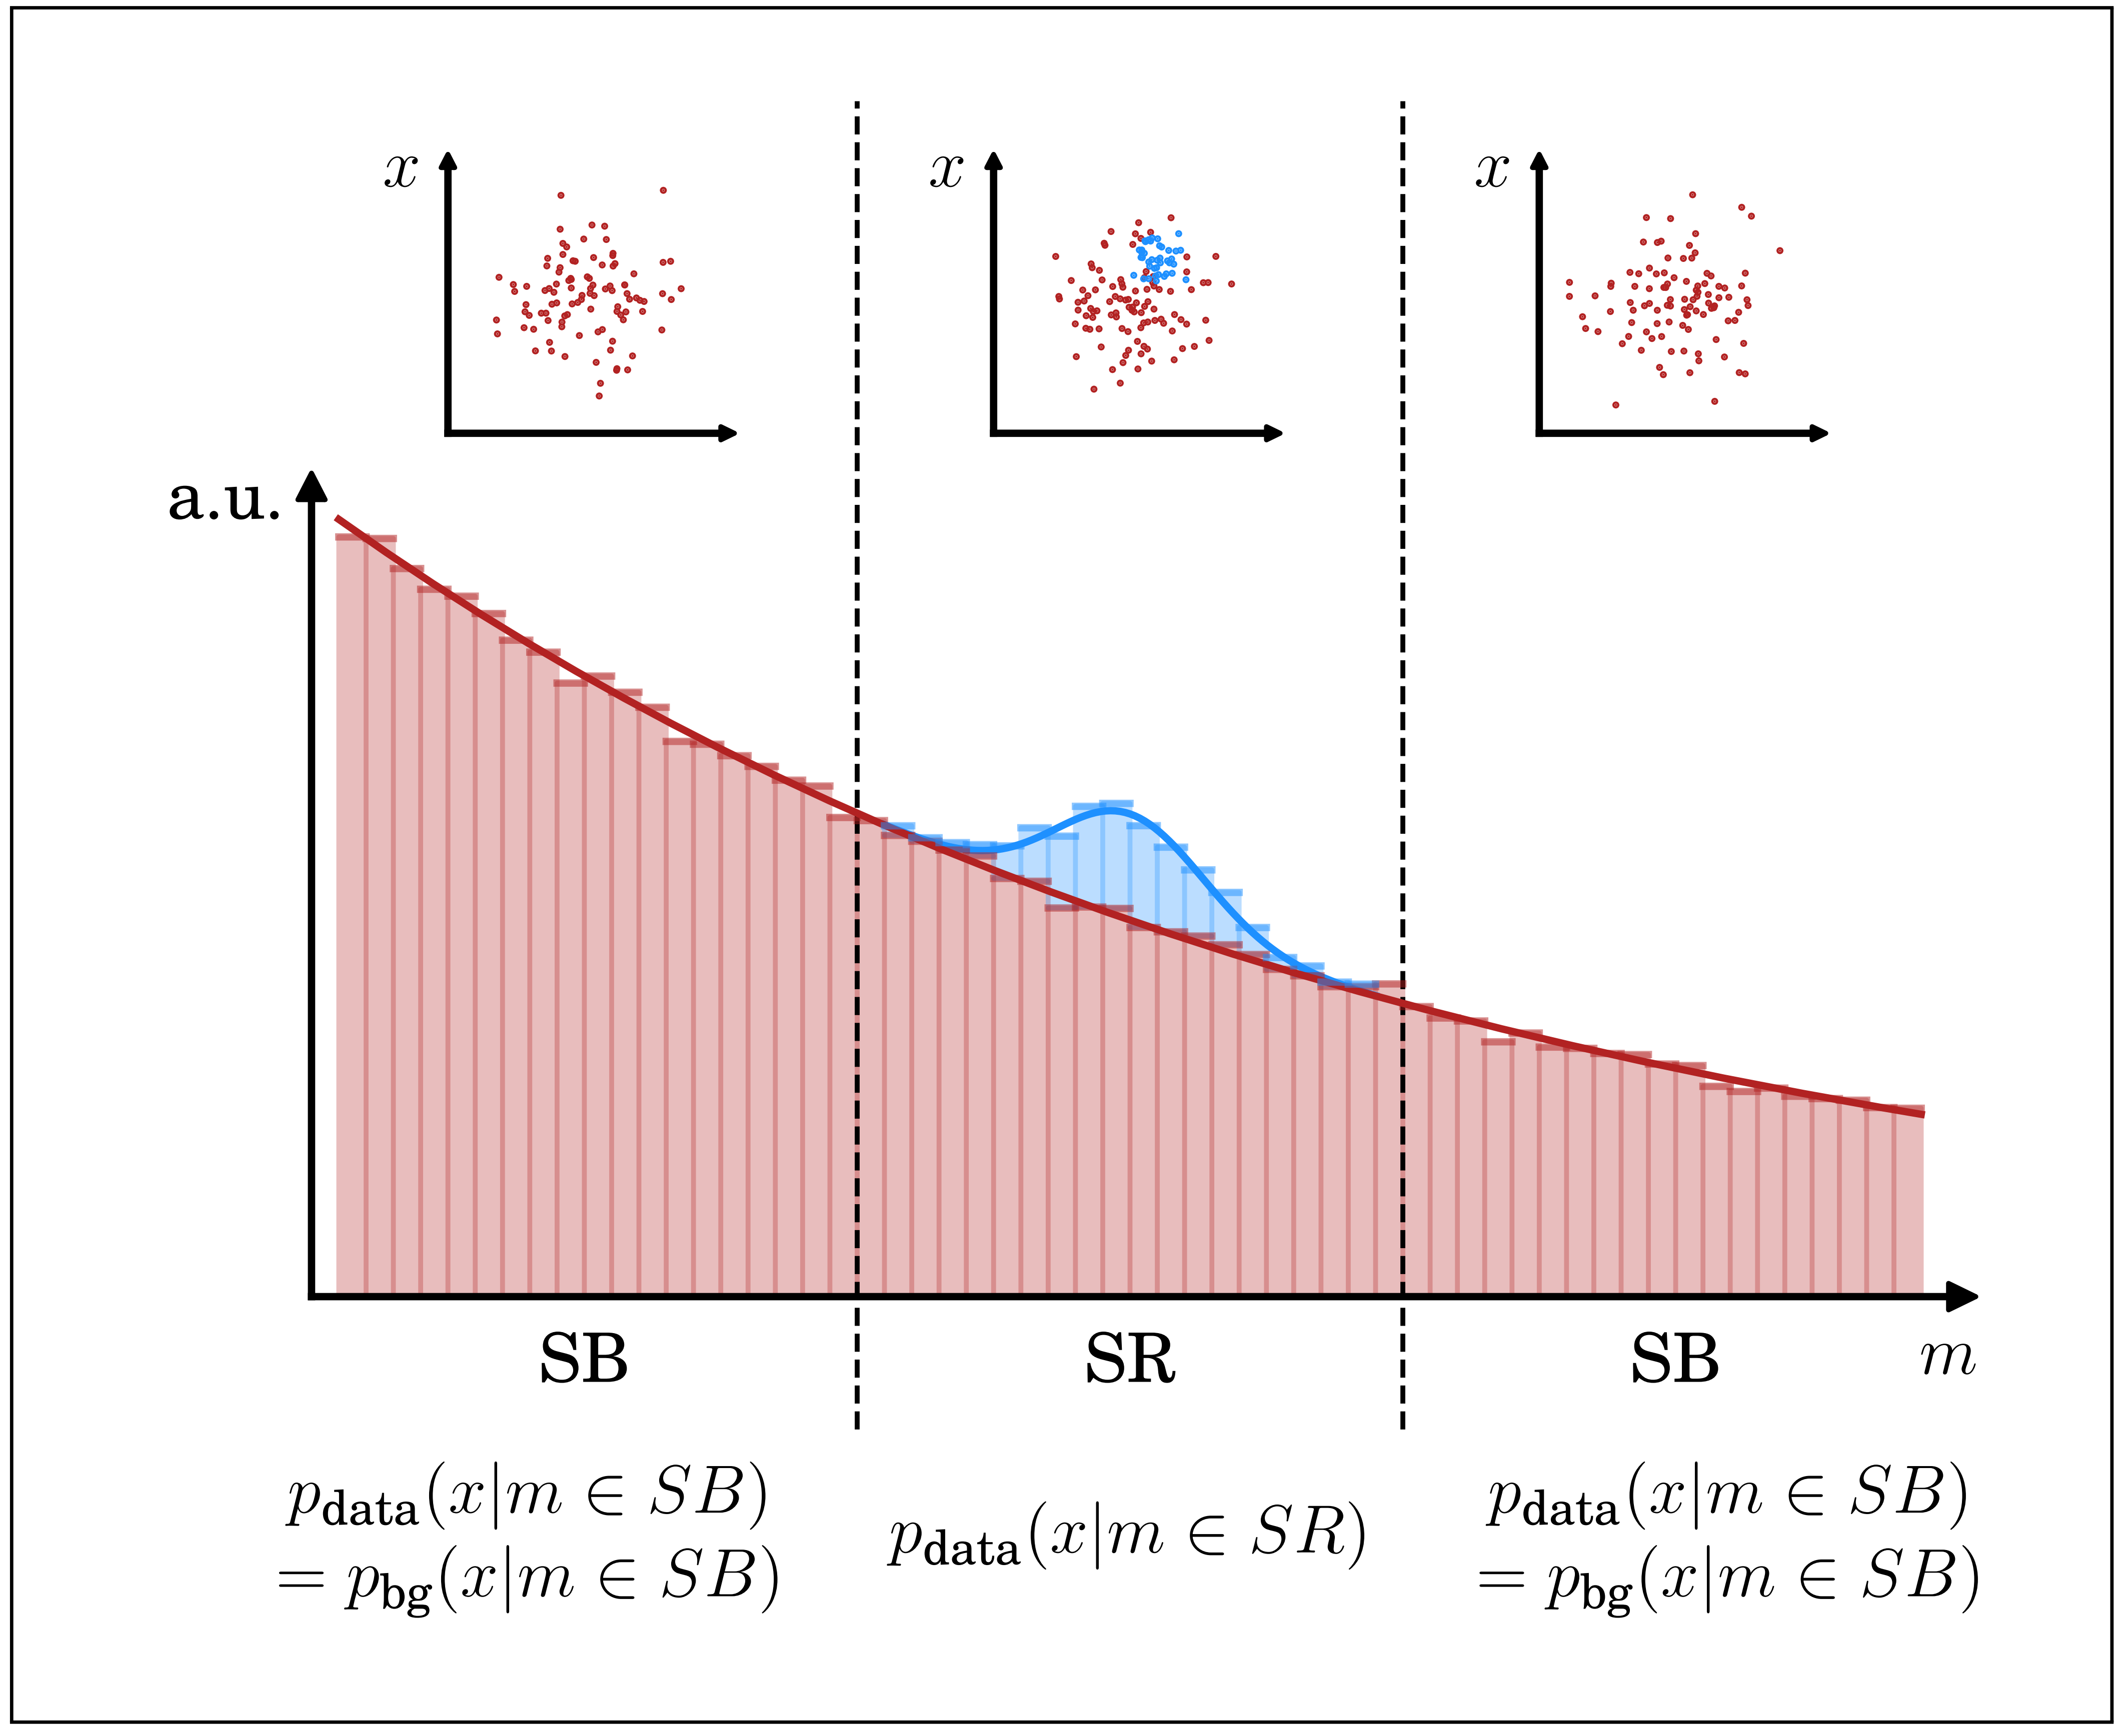

Train two separate ***normalizing flows*** on SR and SB events to learn $p_{data}(x\,|\,m \in SR)$ and $p_{data}(x\,|\,m \in SB) = p_{bg}(x\,|\,m \in SB)$.

The SB NF automatically interpolates into the SR, giving an estimate of $p_{bg}(x\,|\,m \in SR)$.

Construct likelihood ratio explicitly.

- Pros: robust against correlations!
- Cons: density estimation much harder than classification

"ANOmaly detection with **D**ensity **E**stimation **(ANODE)"**

Idea: density estimation+classification

DS+ Hallin et al 2109.00546, 2210.14924

from 2109.00546

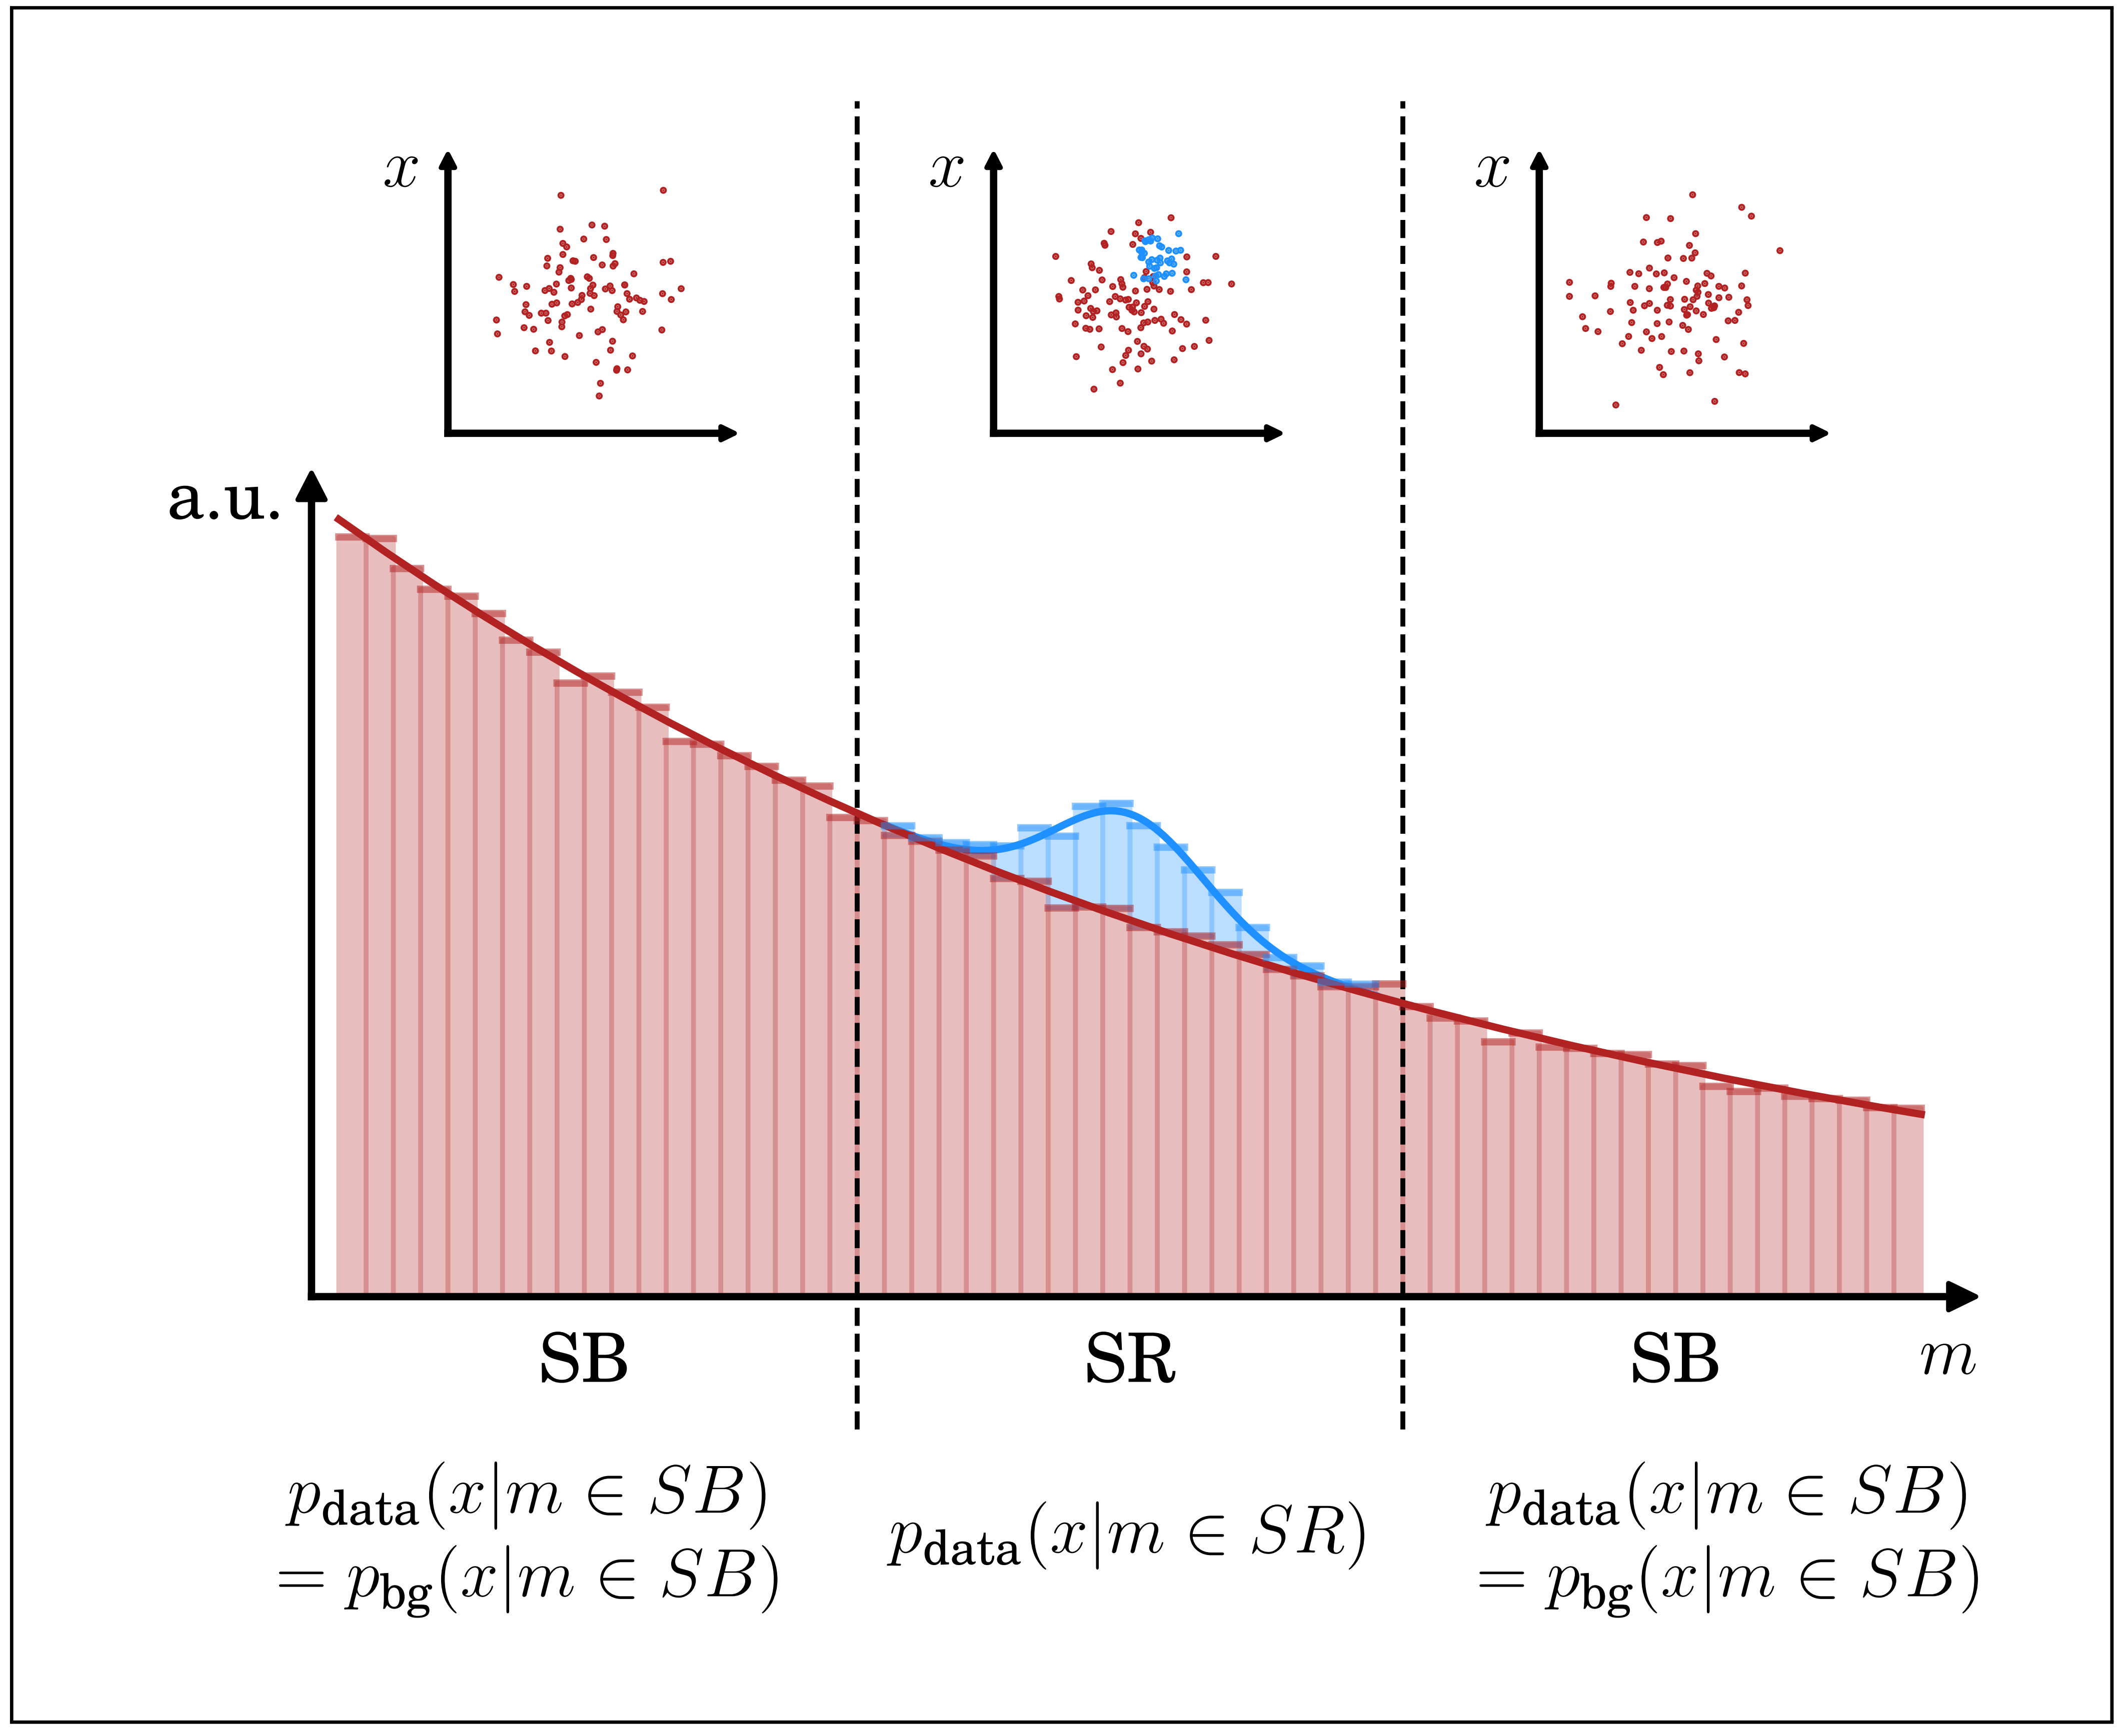

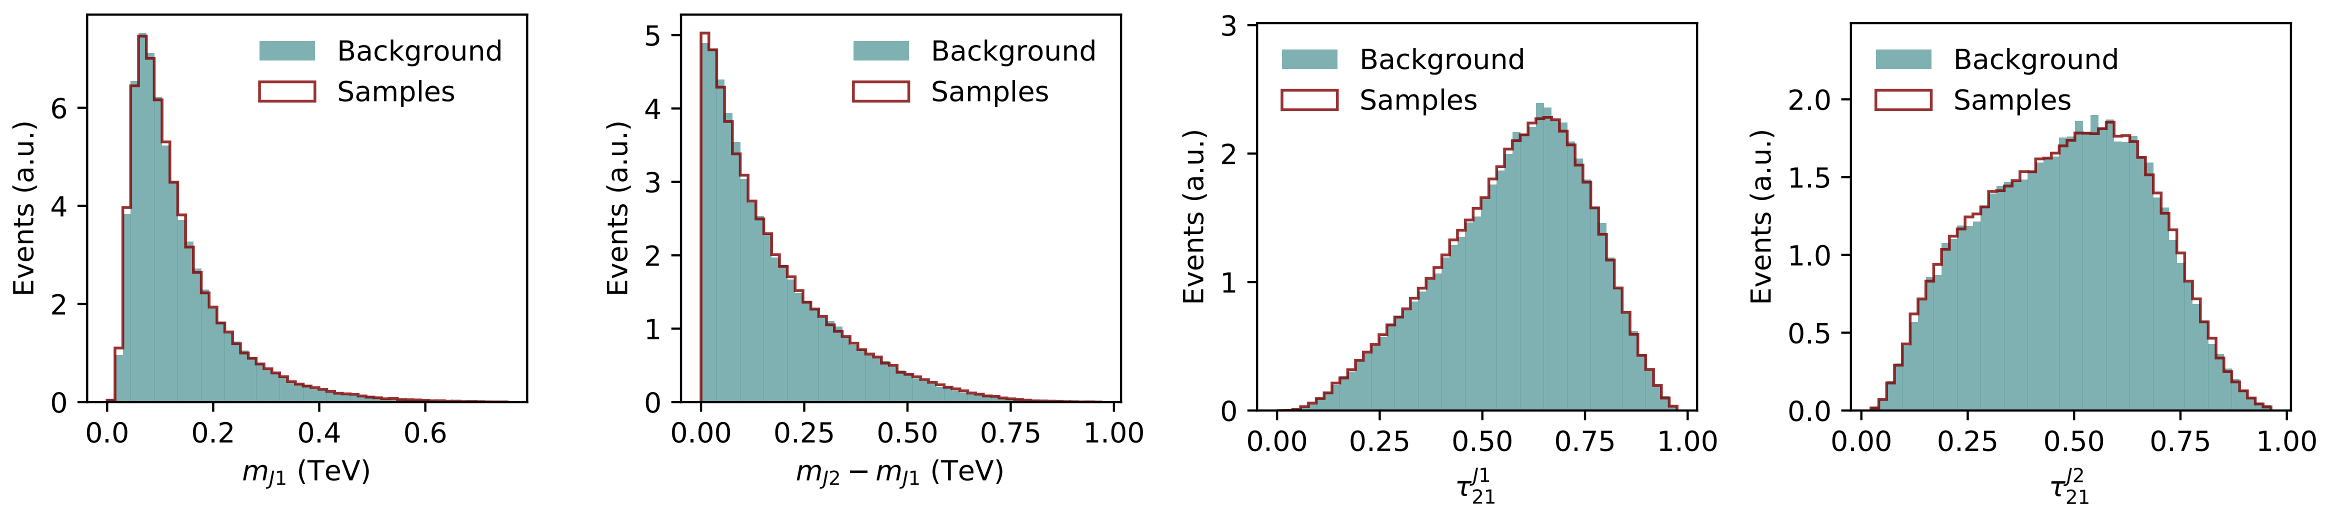

Idea: density estimation+classification

DS+ Hallin et al 2109.00546, 2210.14924

from 2109.00546

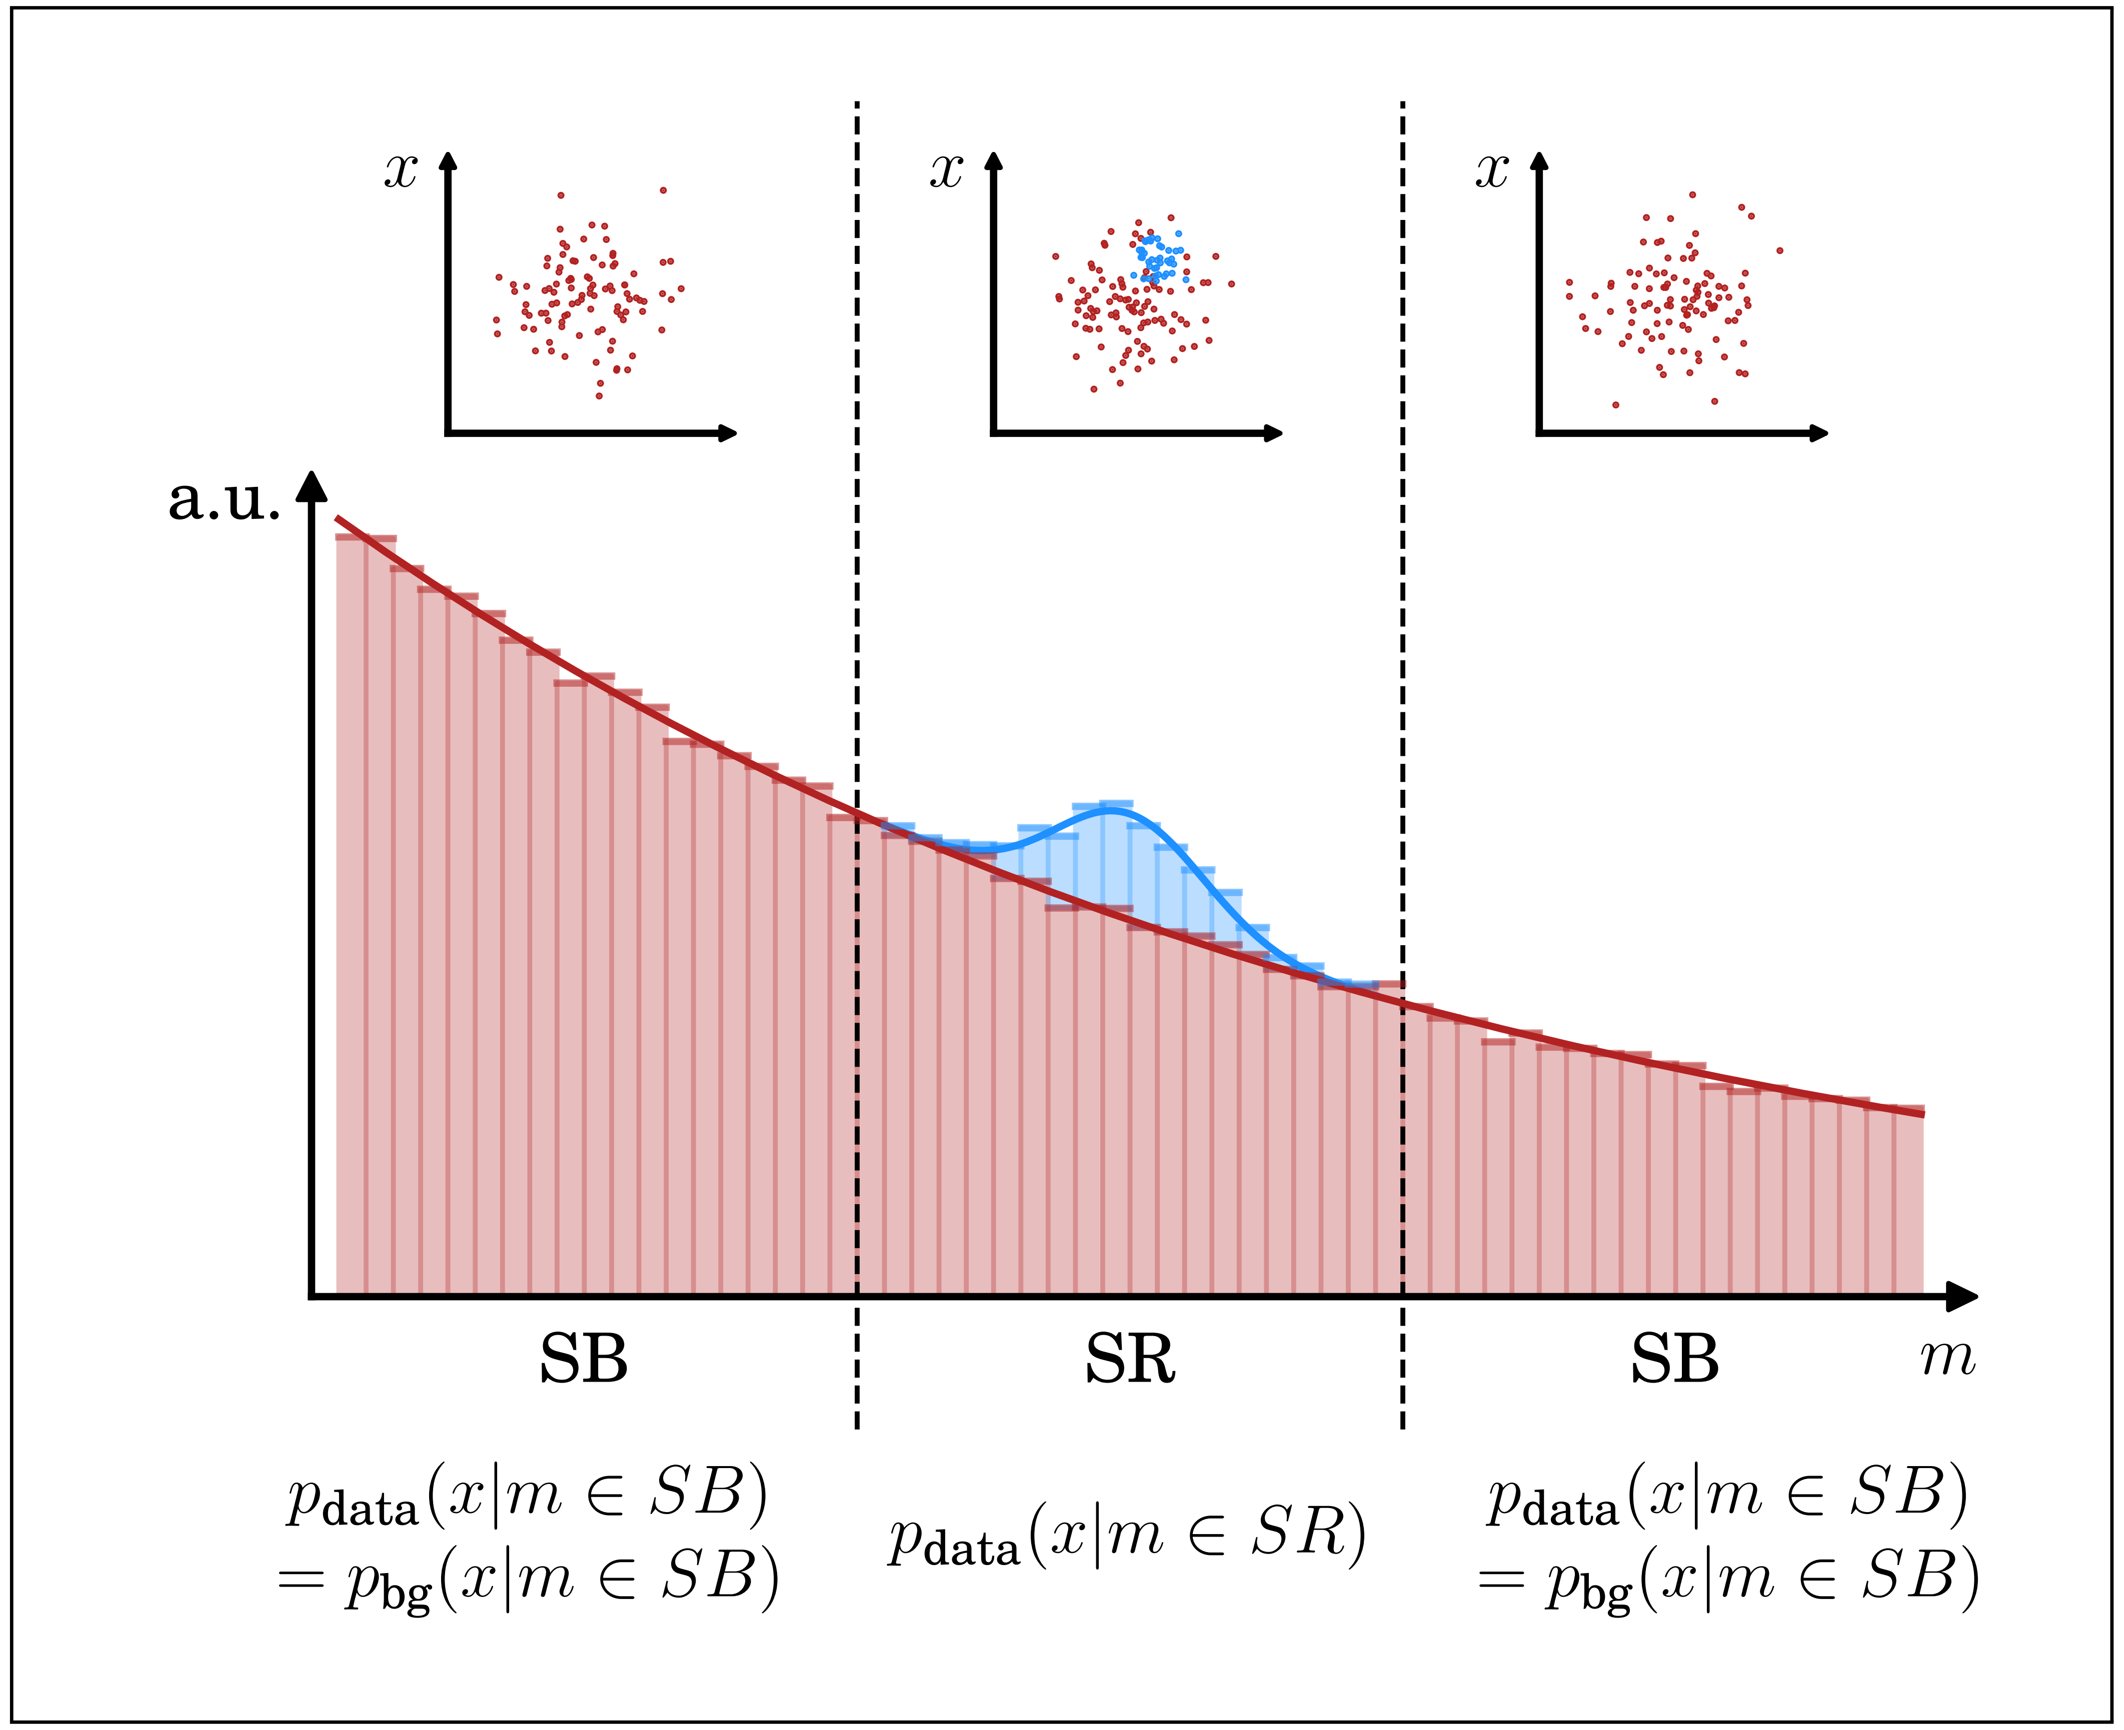

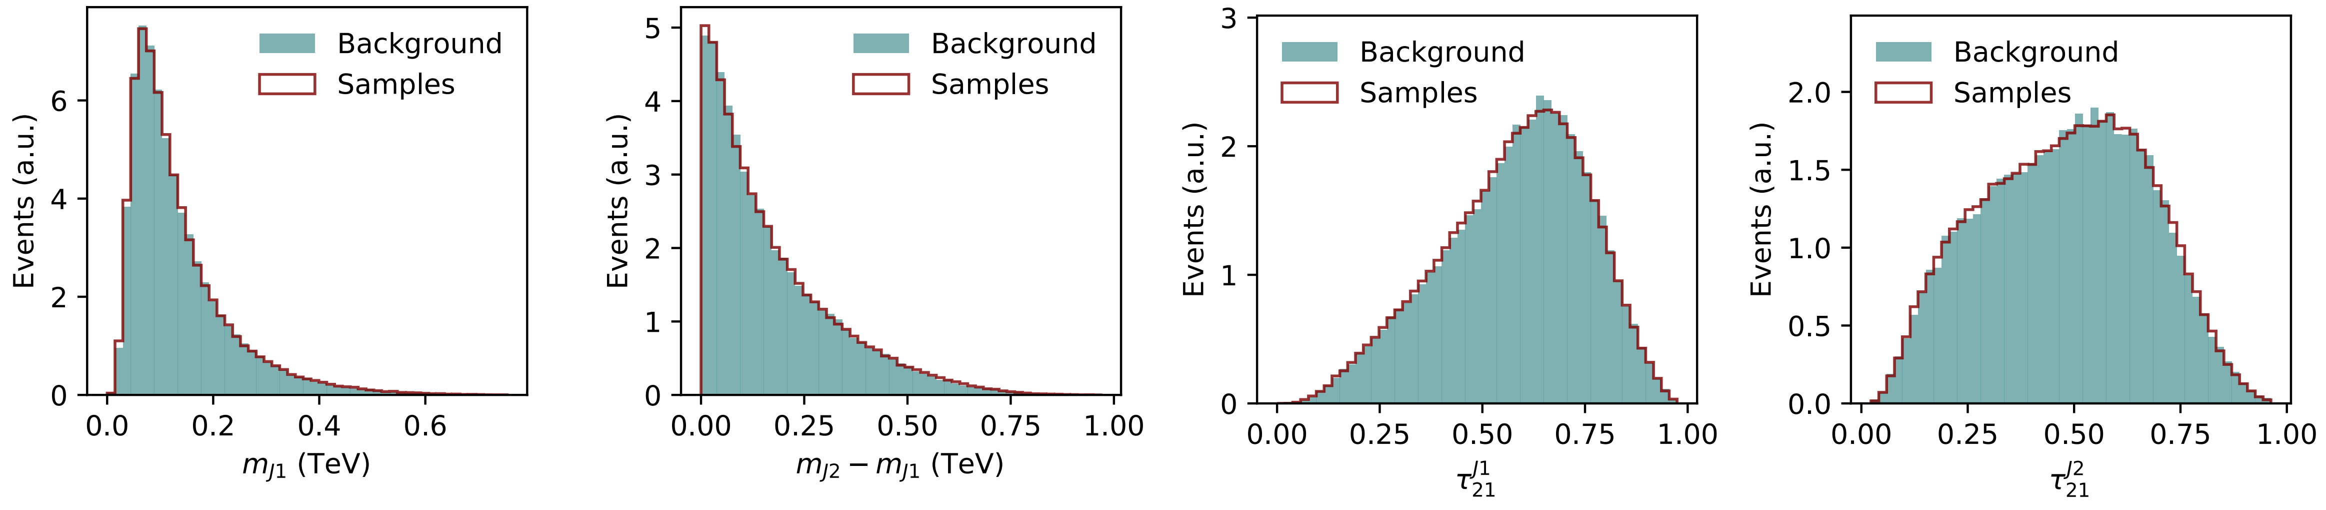

Can also sample ***synthetic bg events*** from the interpolated SB model!

Idea: density estimation+classification

DS+ Hallin et al 2109.00546, 2210.14924

from 2109.00546

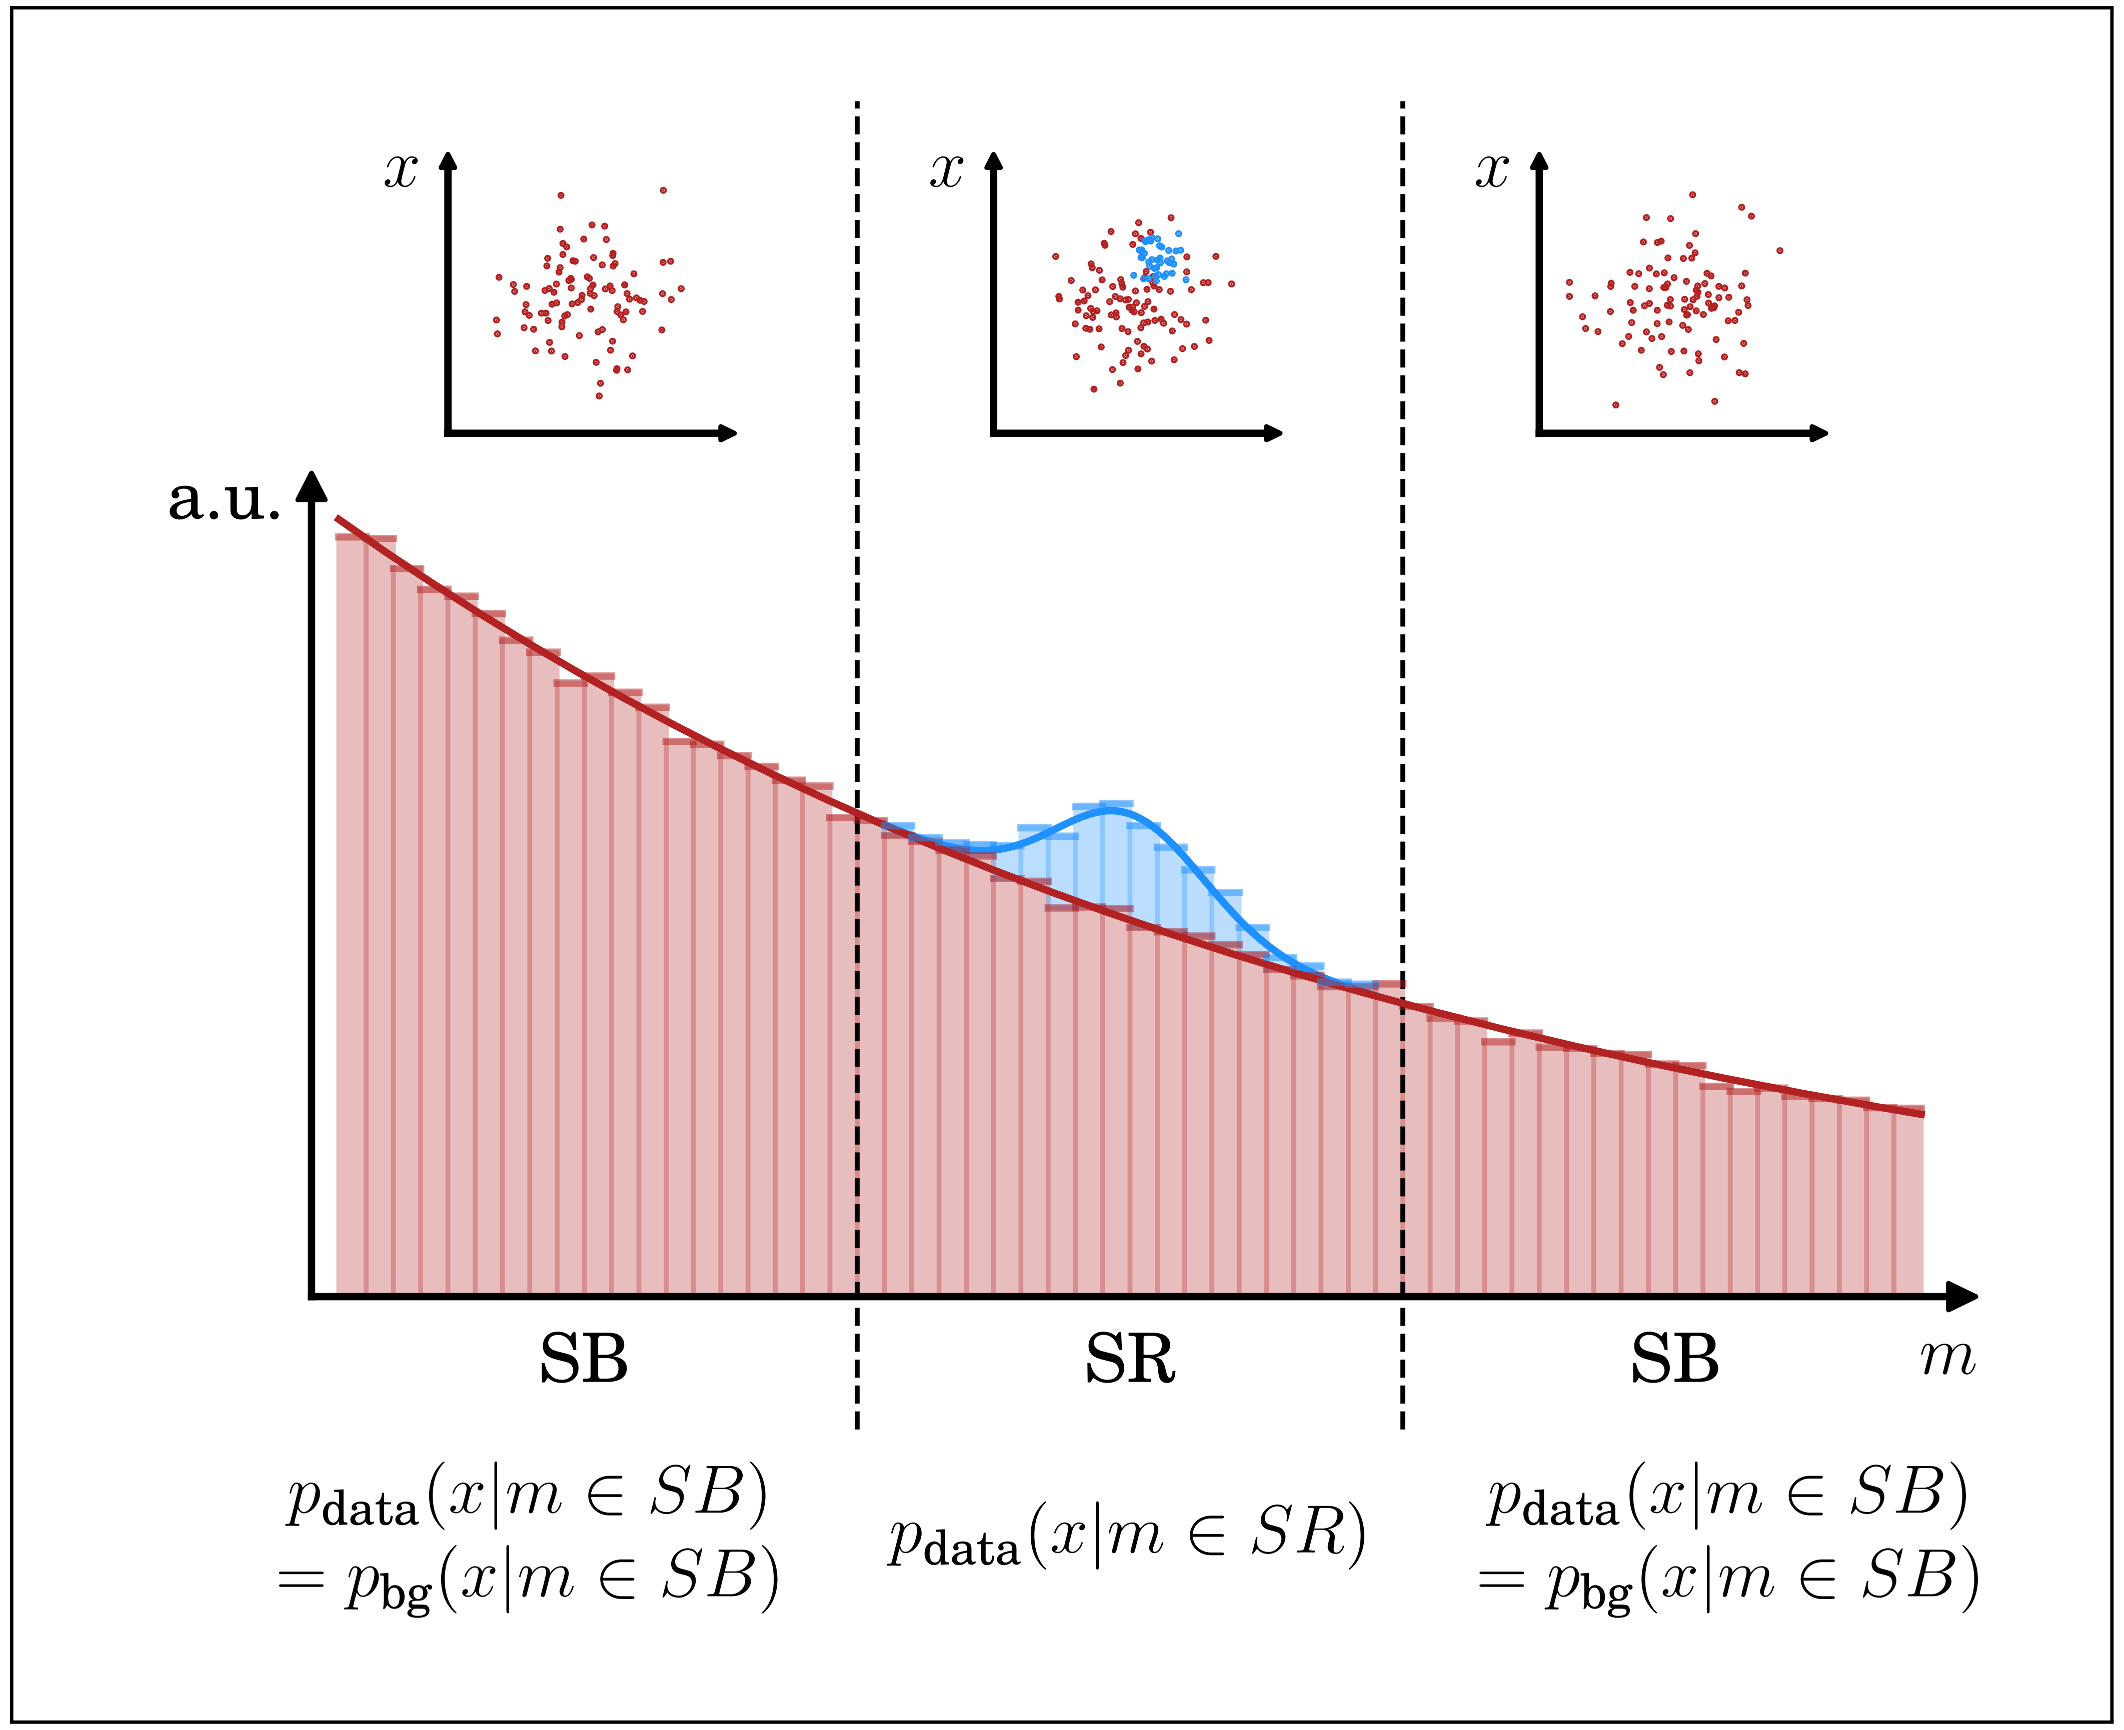

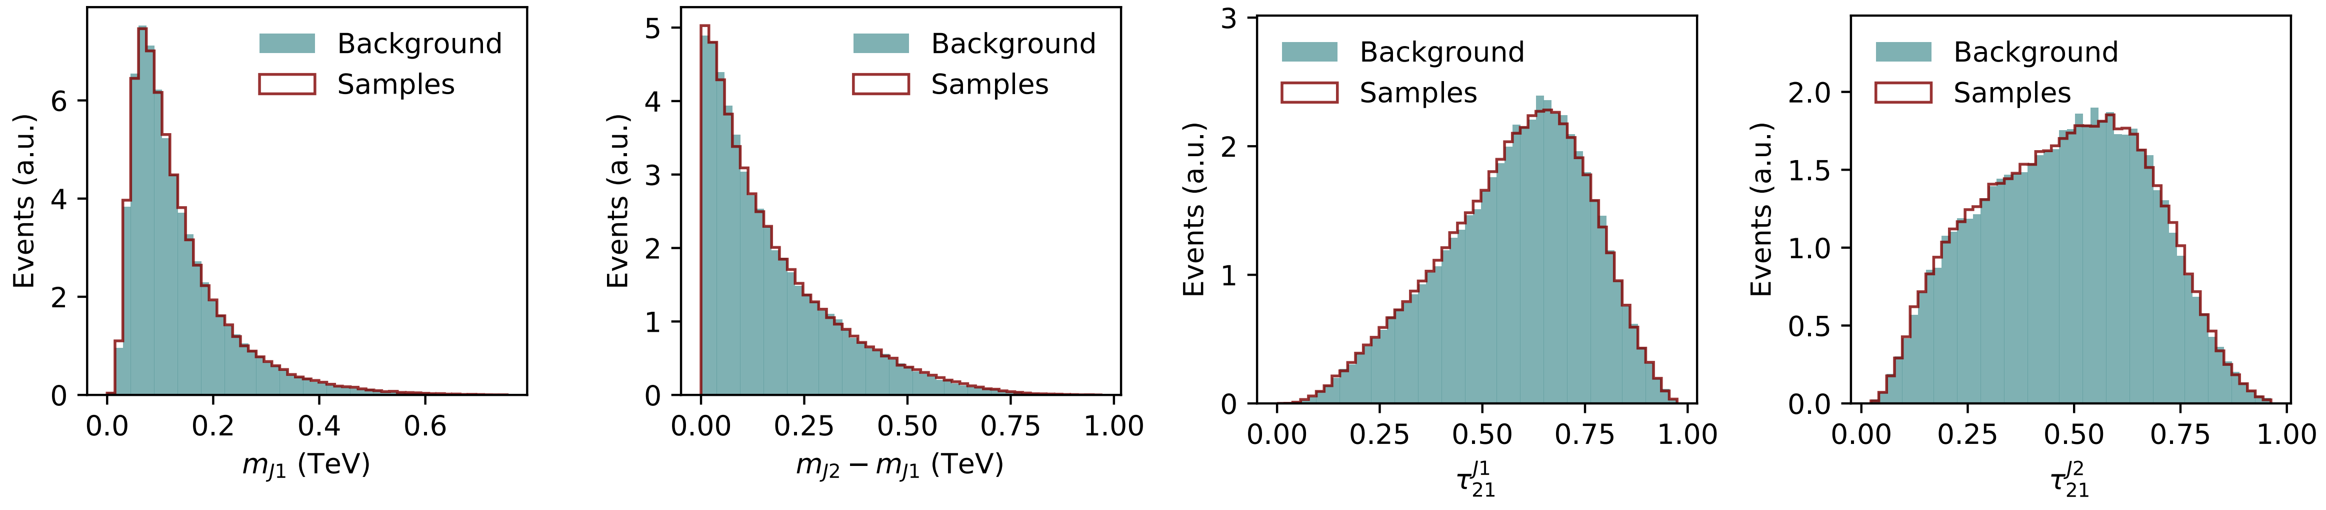

Can also sample ***synthetic bg events*** from the interpolated SB model!

Then learn data-vs-bg classifier as in CWoLa Hunting.

Idea: density estimation+classification

DS+ Hallin et al 2109.00546, 2210.14924

from 2109.00546

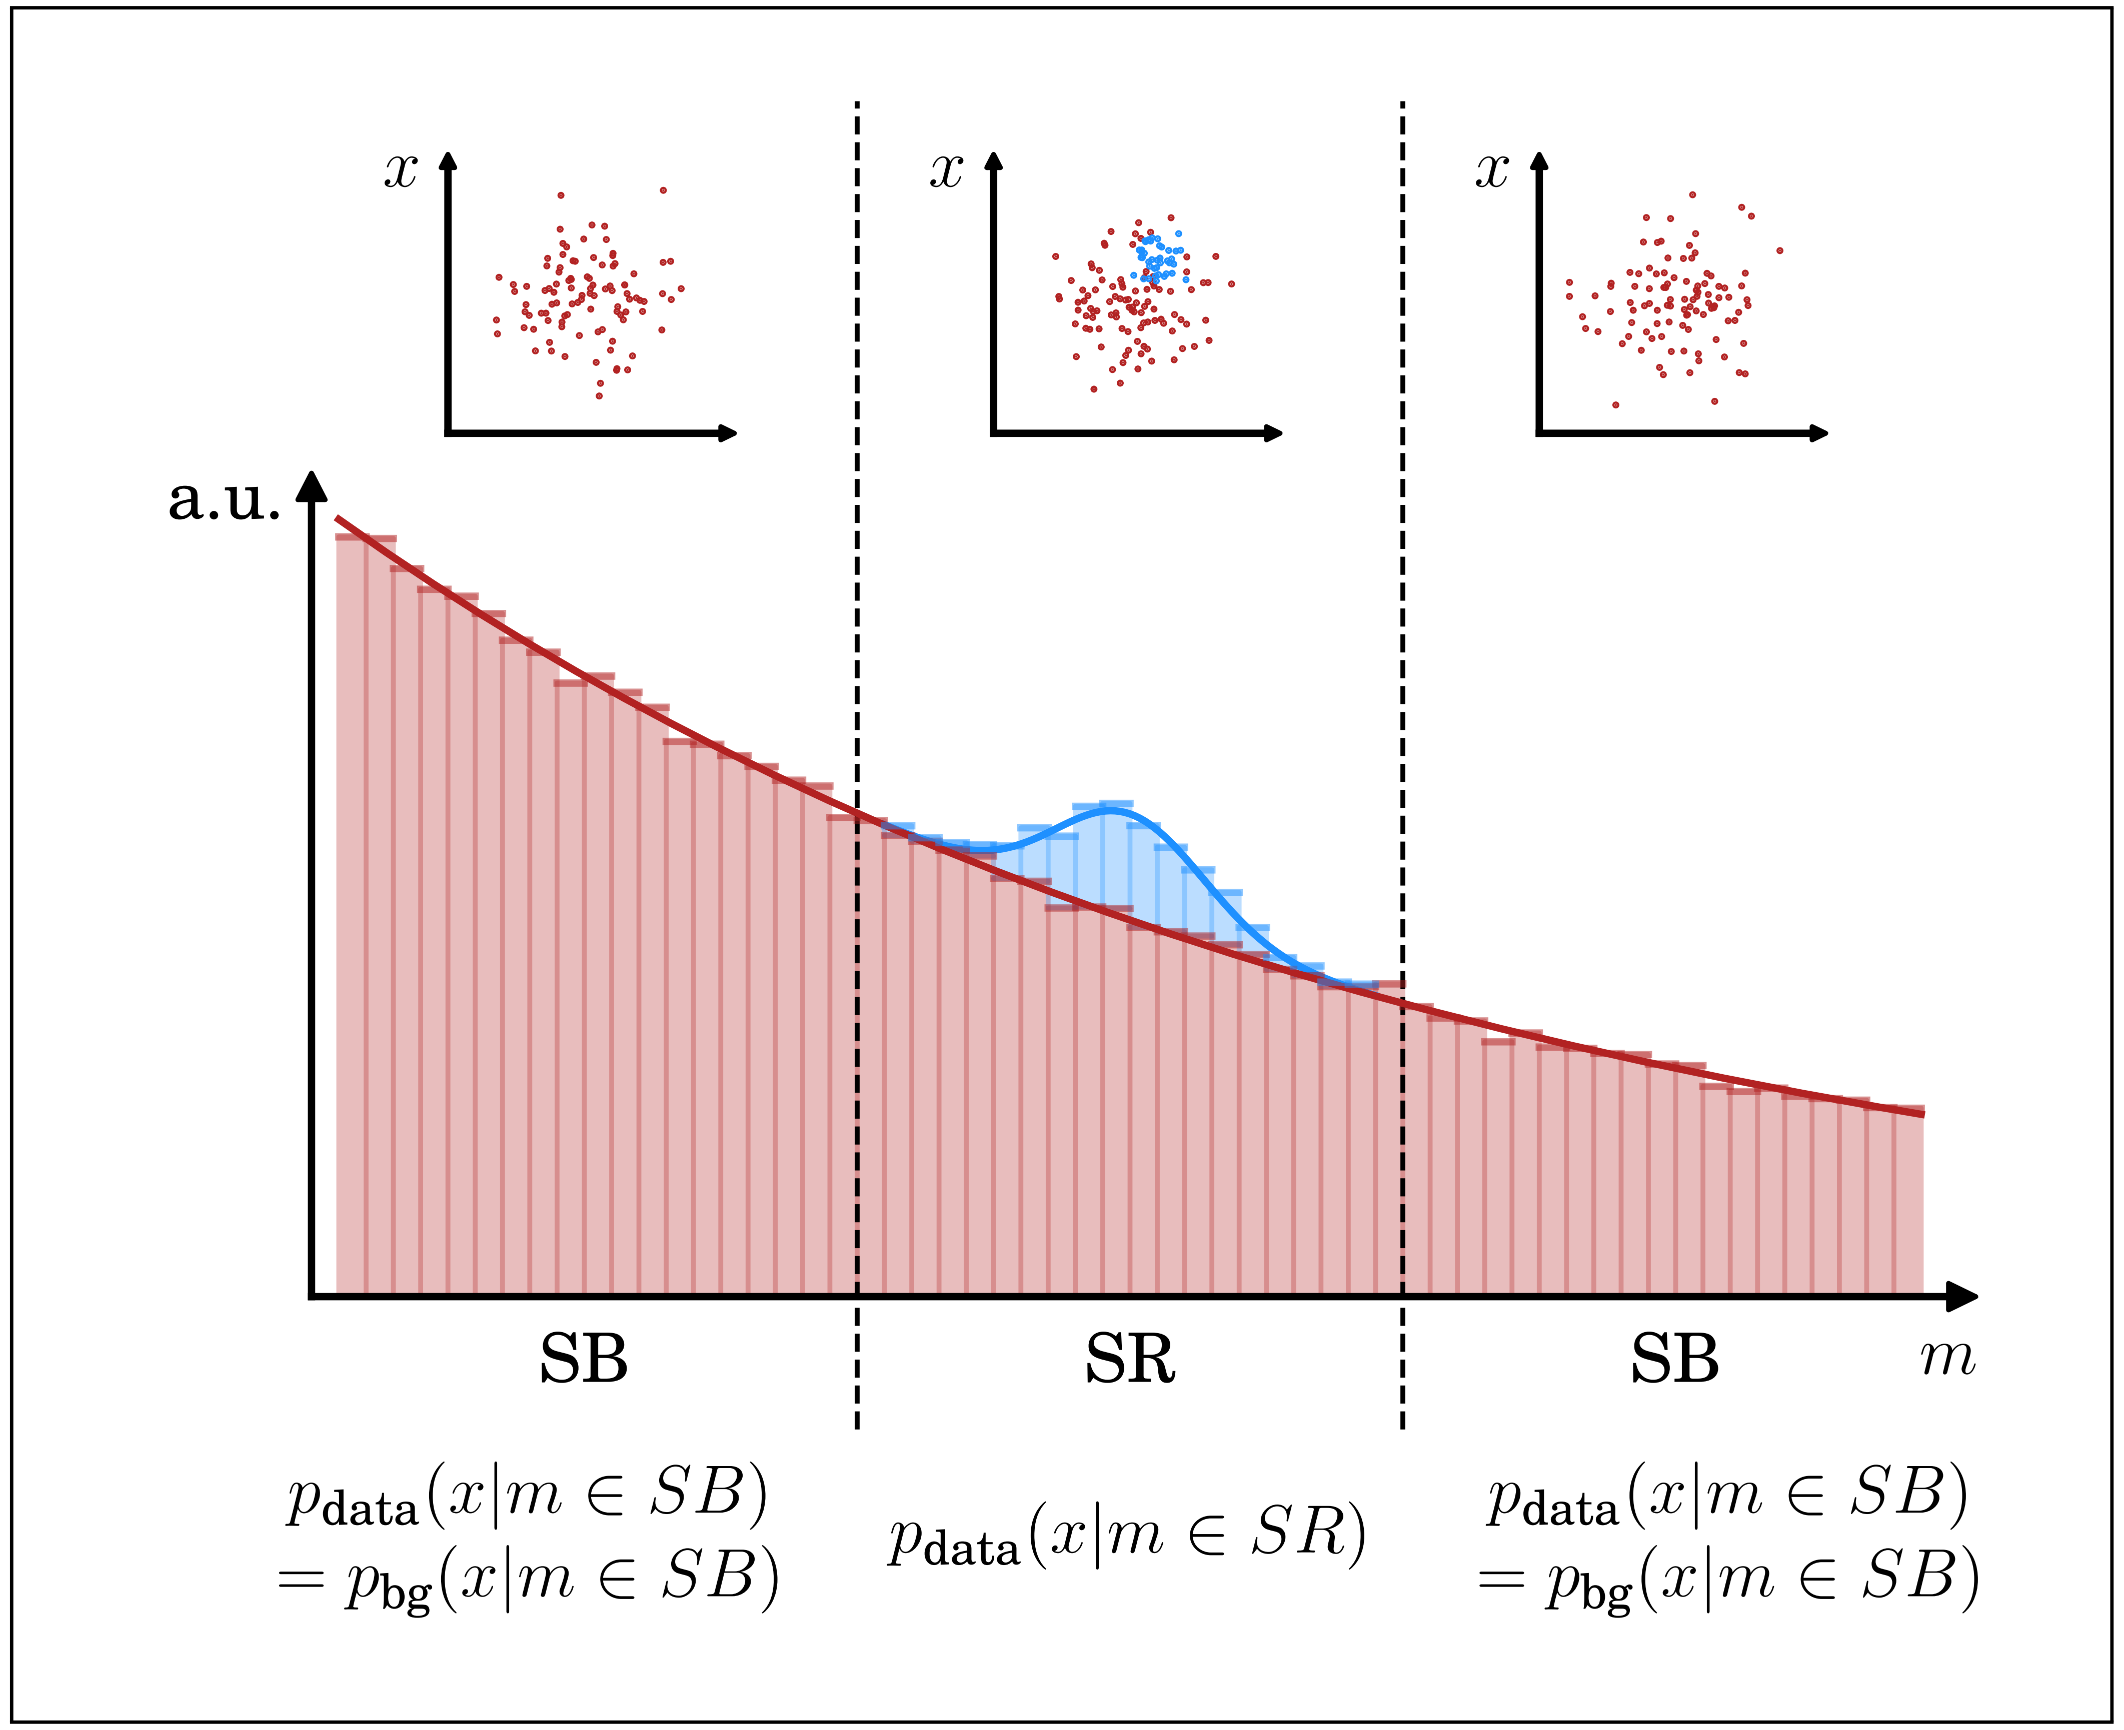

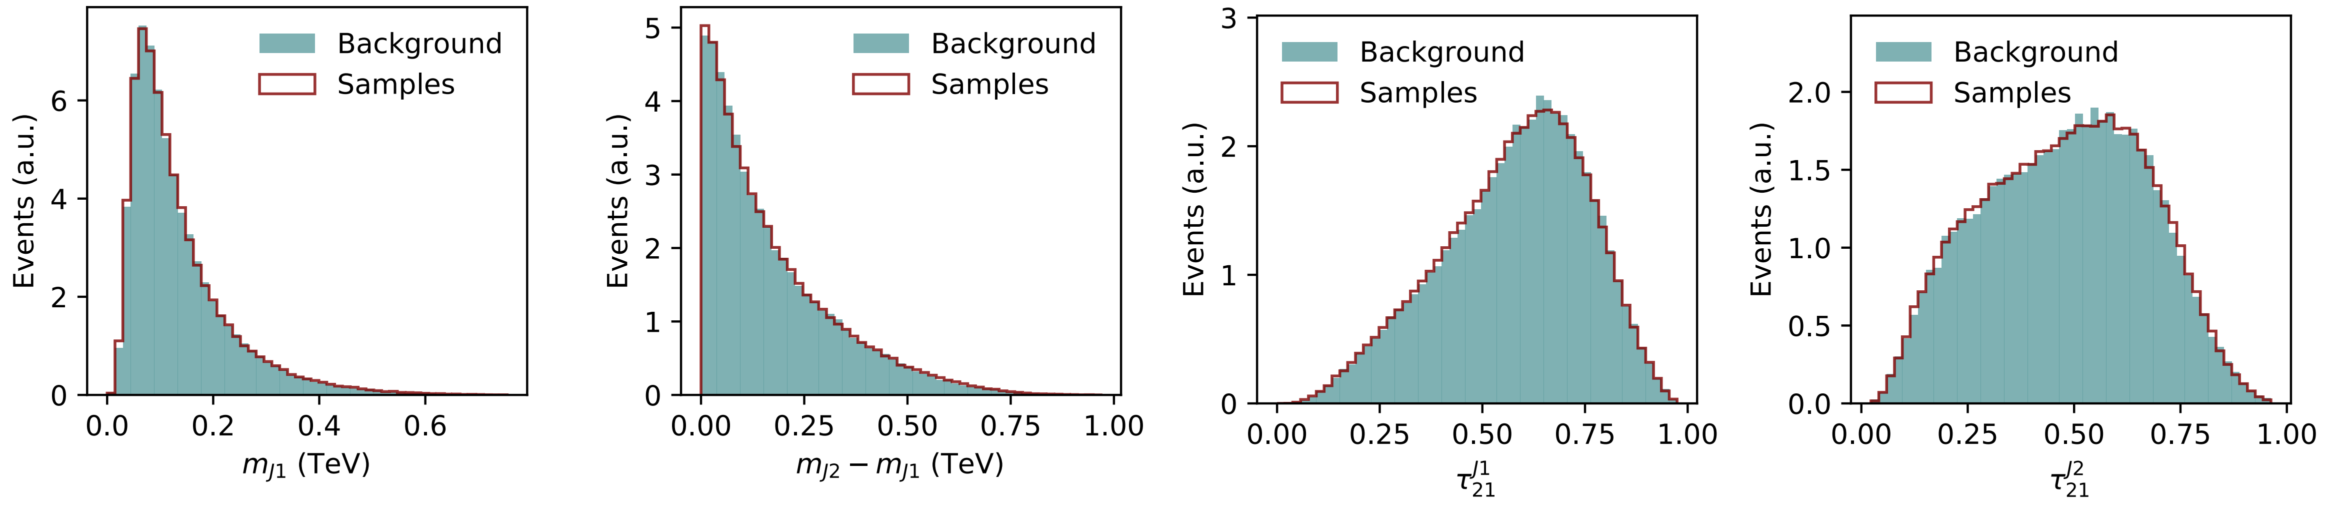

Can also sample ***synthetic bg events*** from the interpolated SB model!

Then learn data-vs-bg classifier as in CWoLa Hunting.

- Robust against correlations ***and*** don't have to learn separate density estimator for $p_{data}(x \,|\, m \in SR)$

Idea: density estimation+classification

DS+ Hallin et al 2109.00546, 2210.14924

from 2109.00546

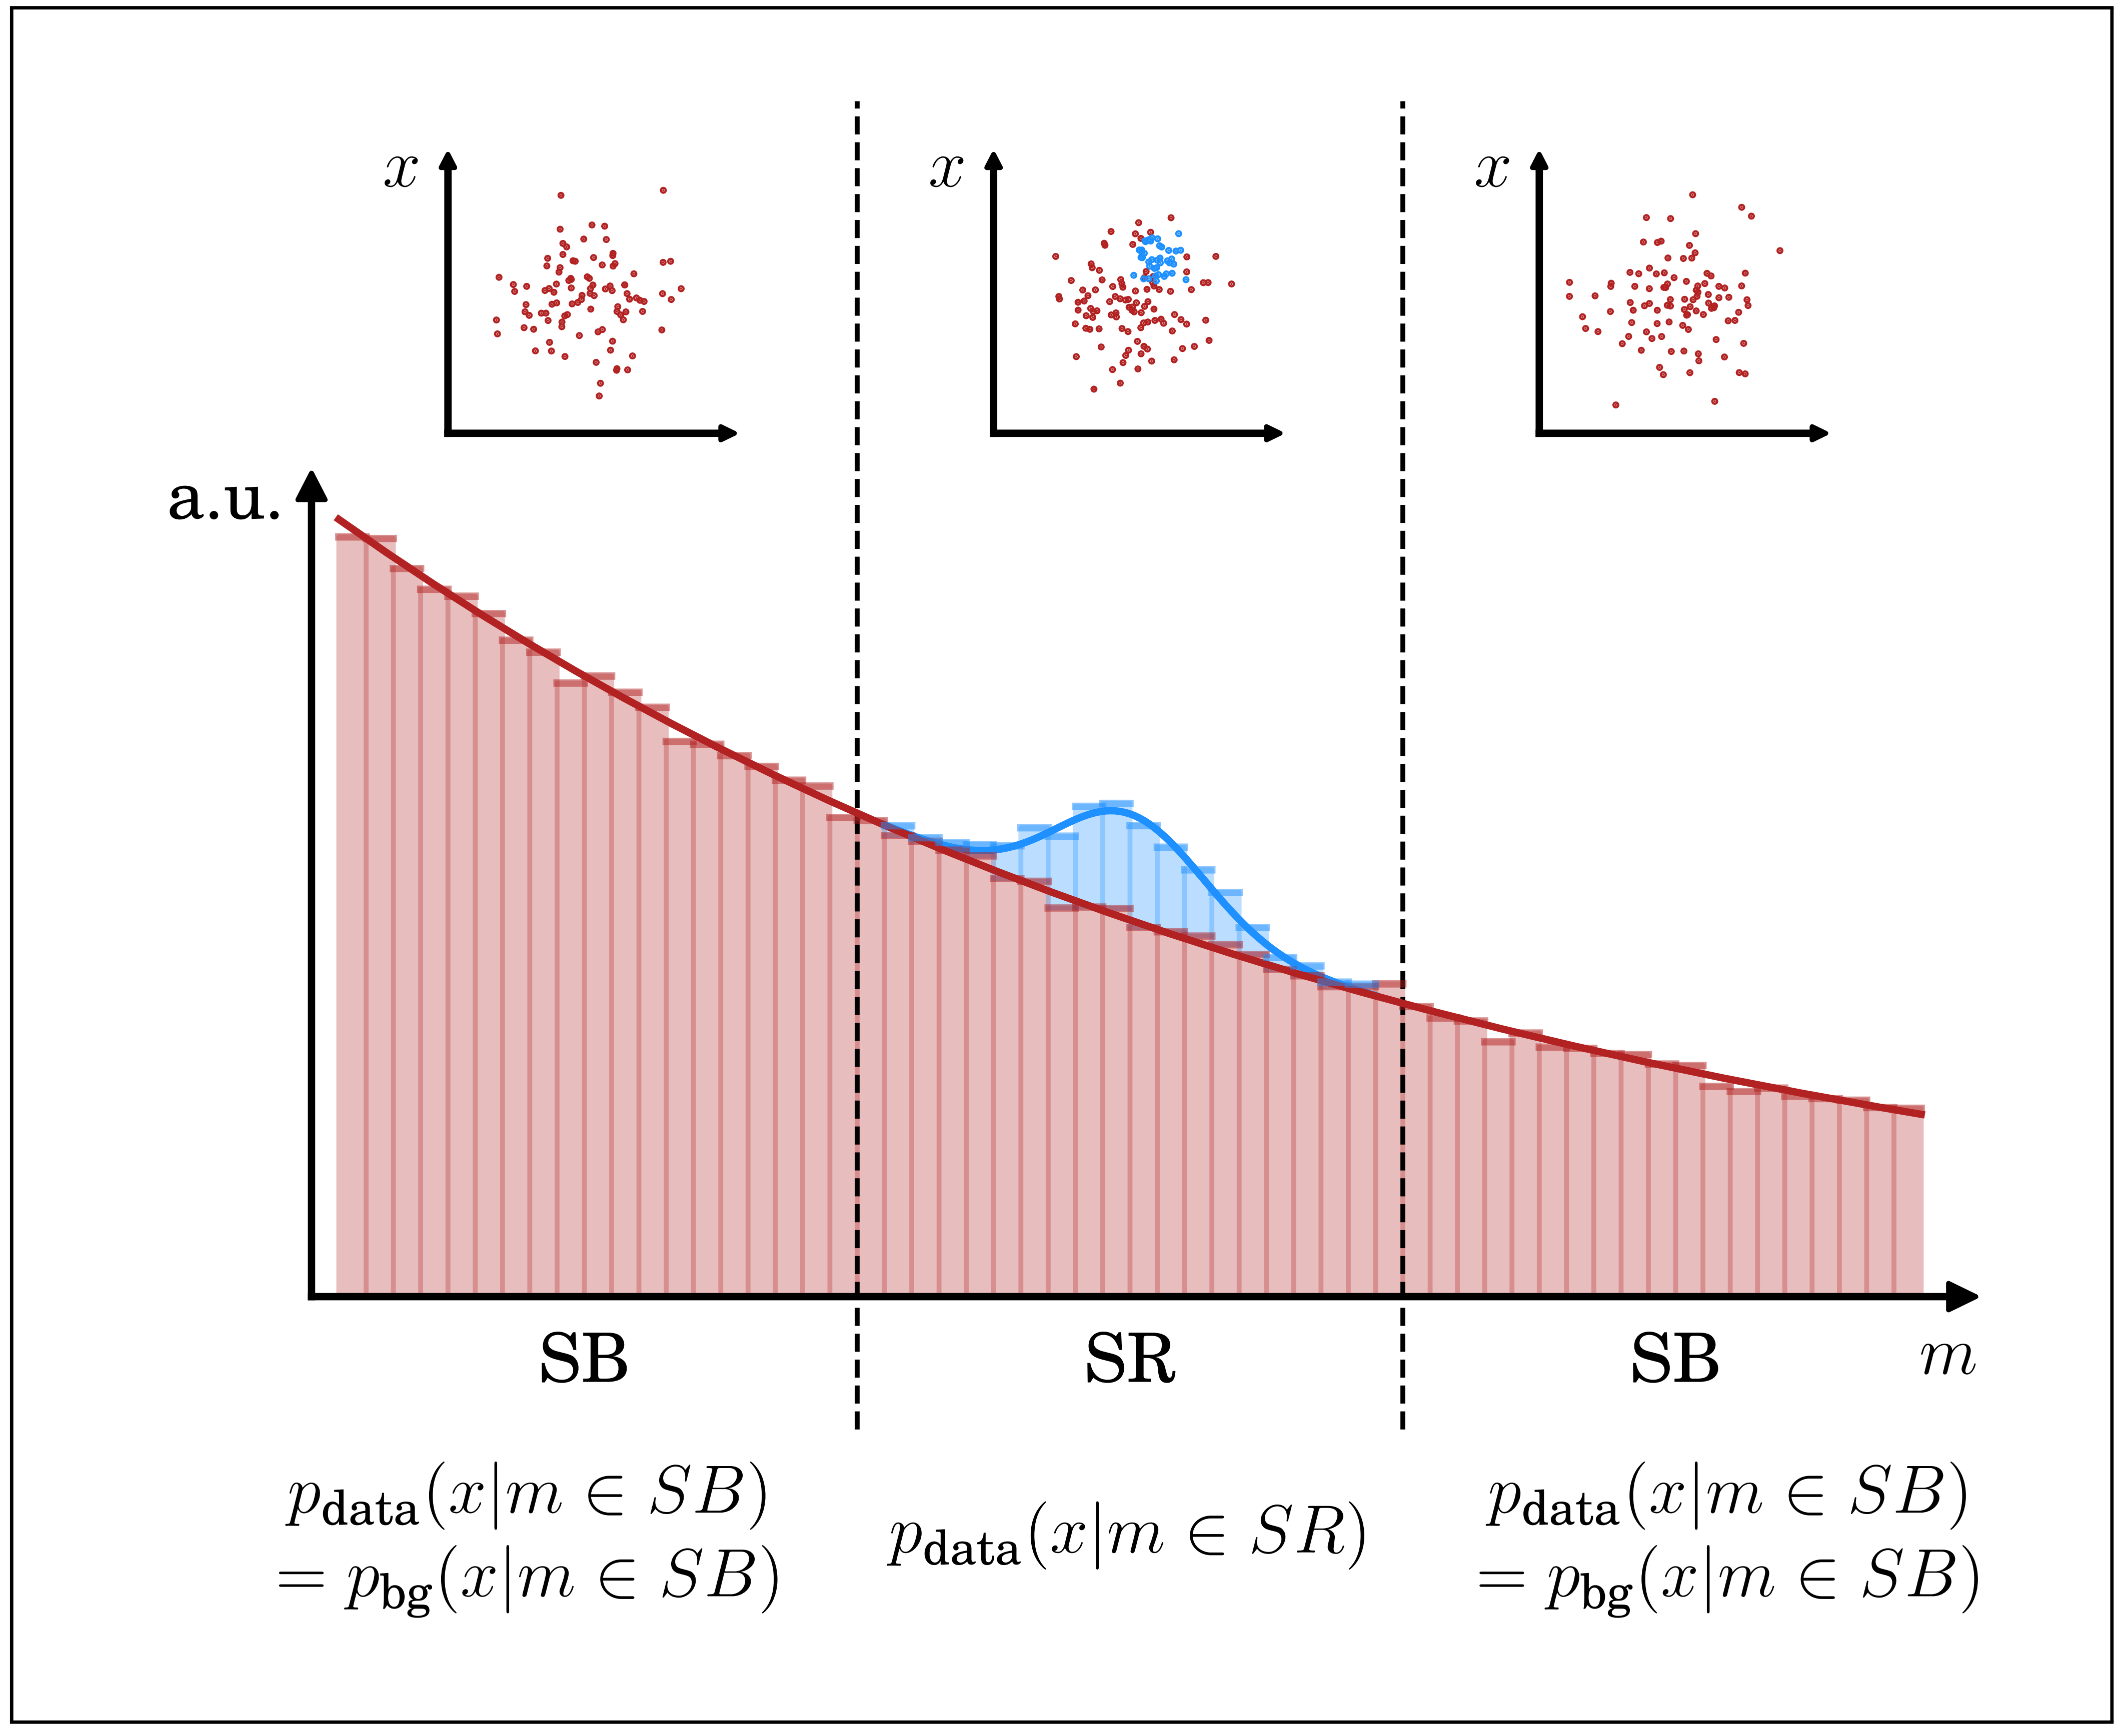

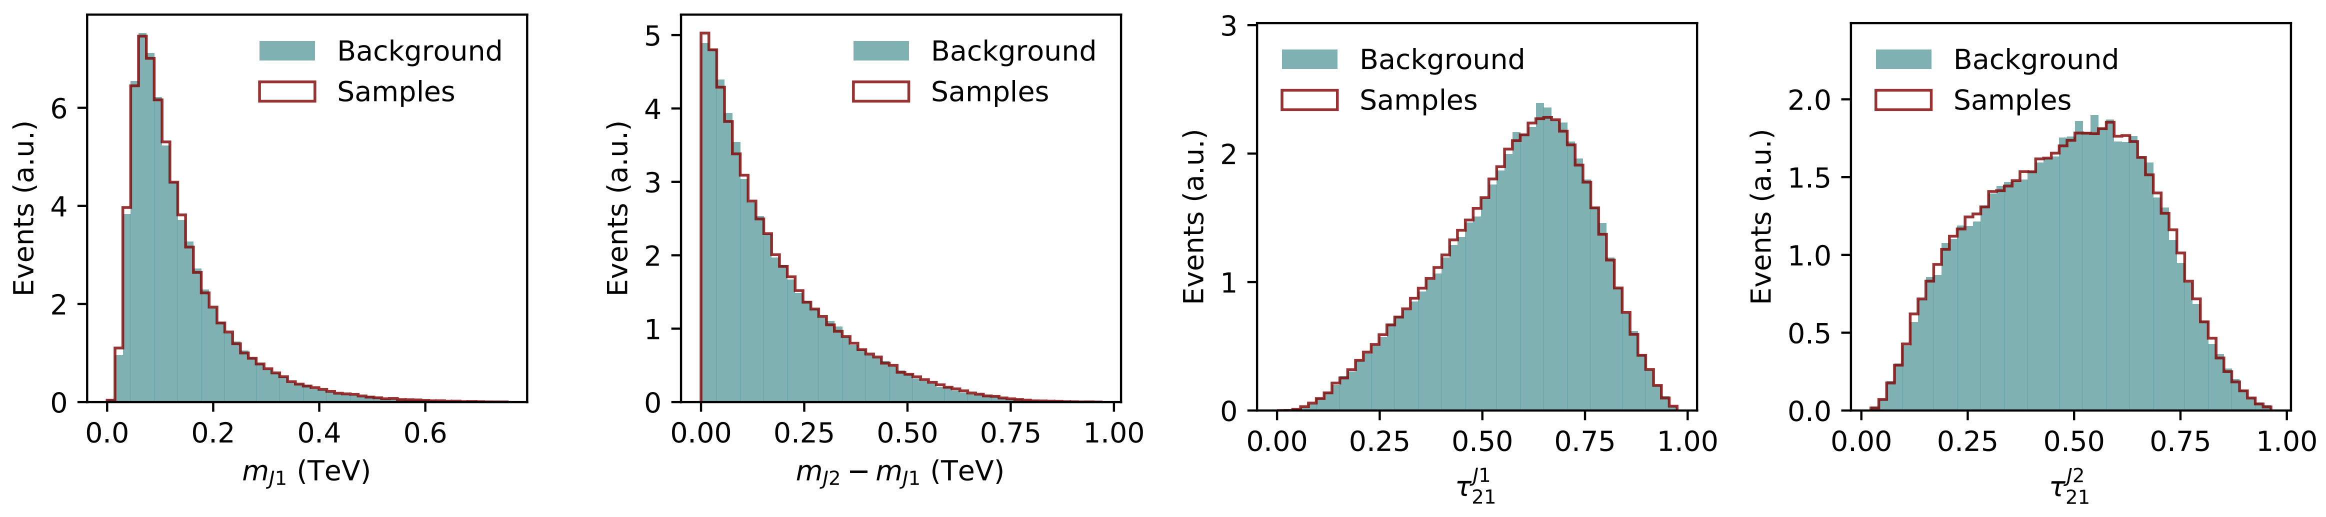

Can also sample ***synthetic bg events*** from the interpolated SB model!

Then learn data-vs-bg classifier as in CWoLa Hunting.

- Robust against correlations ***and*** don't have to learn separate density estimator for $p_{data}(x \,|\, m \in SR)$

"**C**lassifying **A**nomalies **TH**rough **O**uter **D**ensity **E**stimation (CATHODE)"

Summary of methods

from 2109.00546

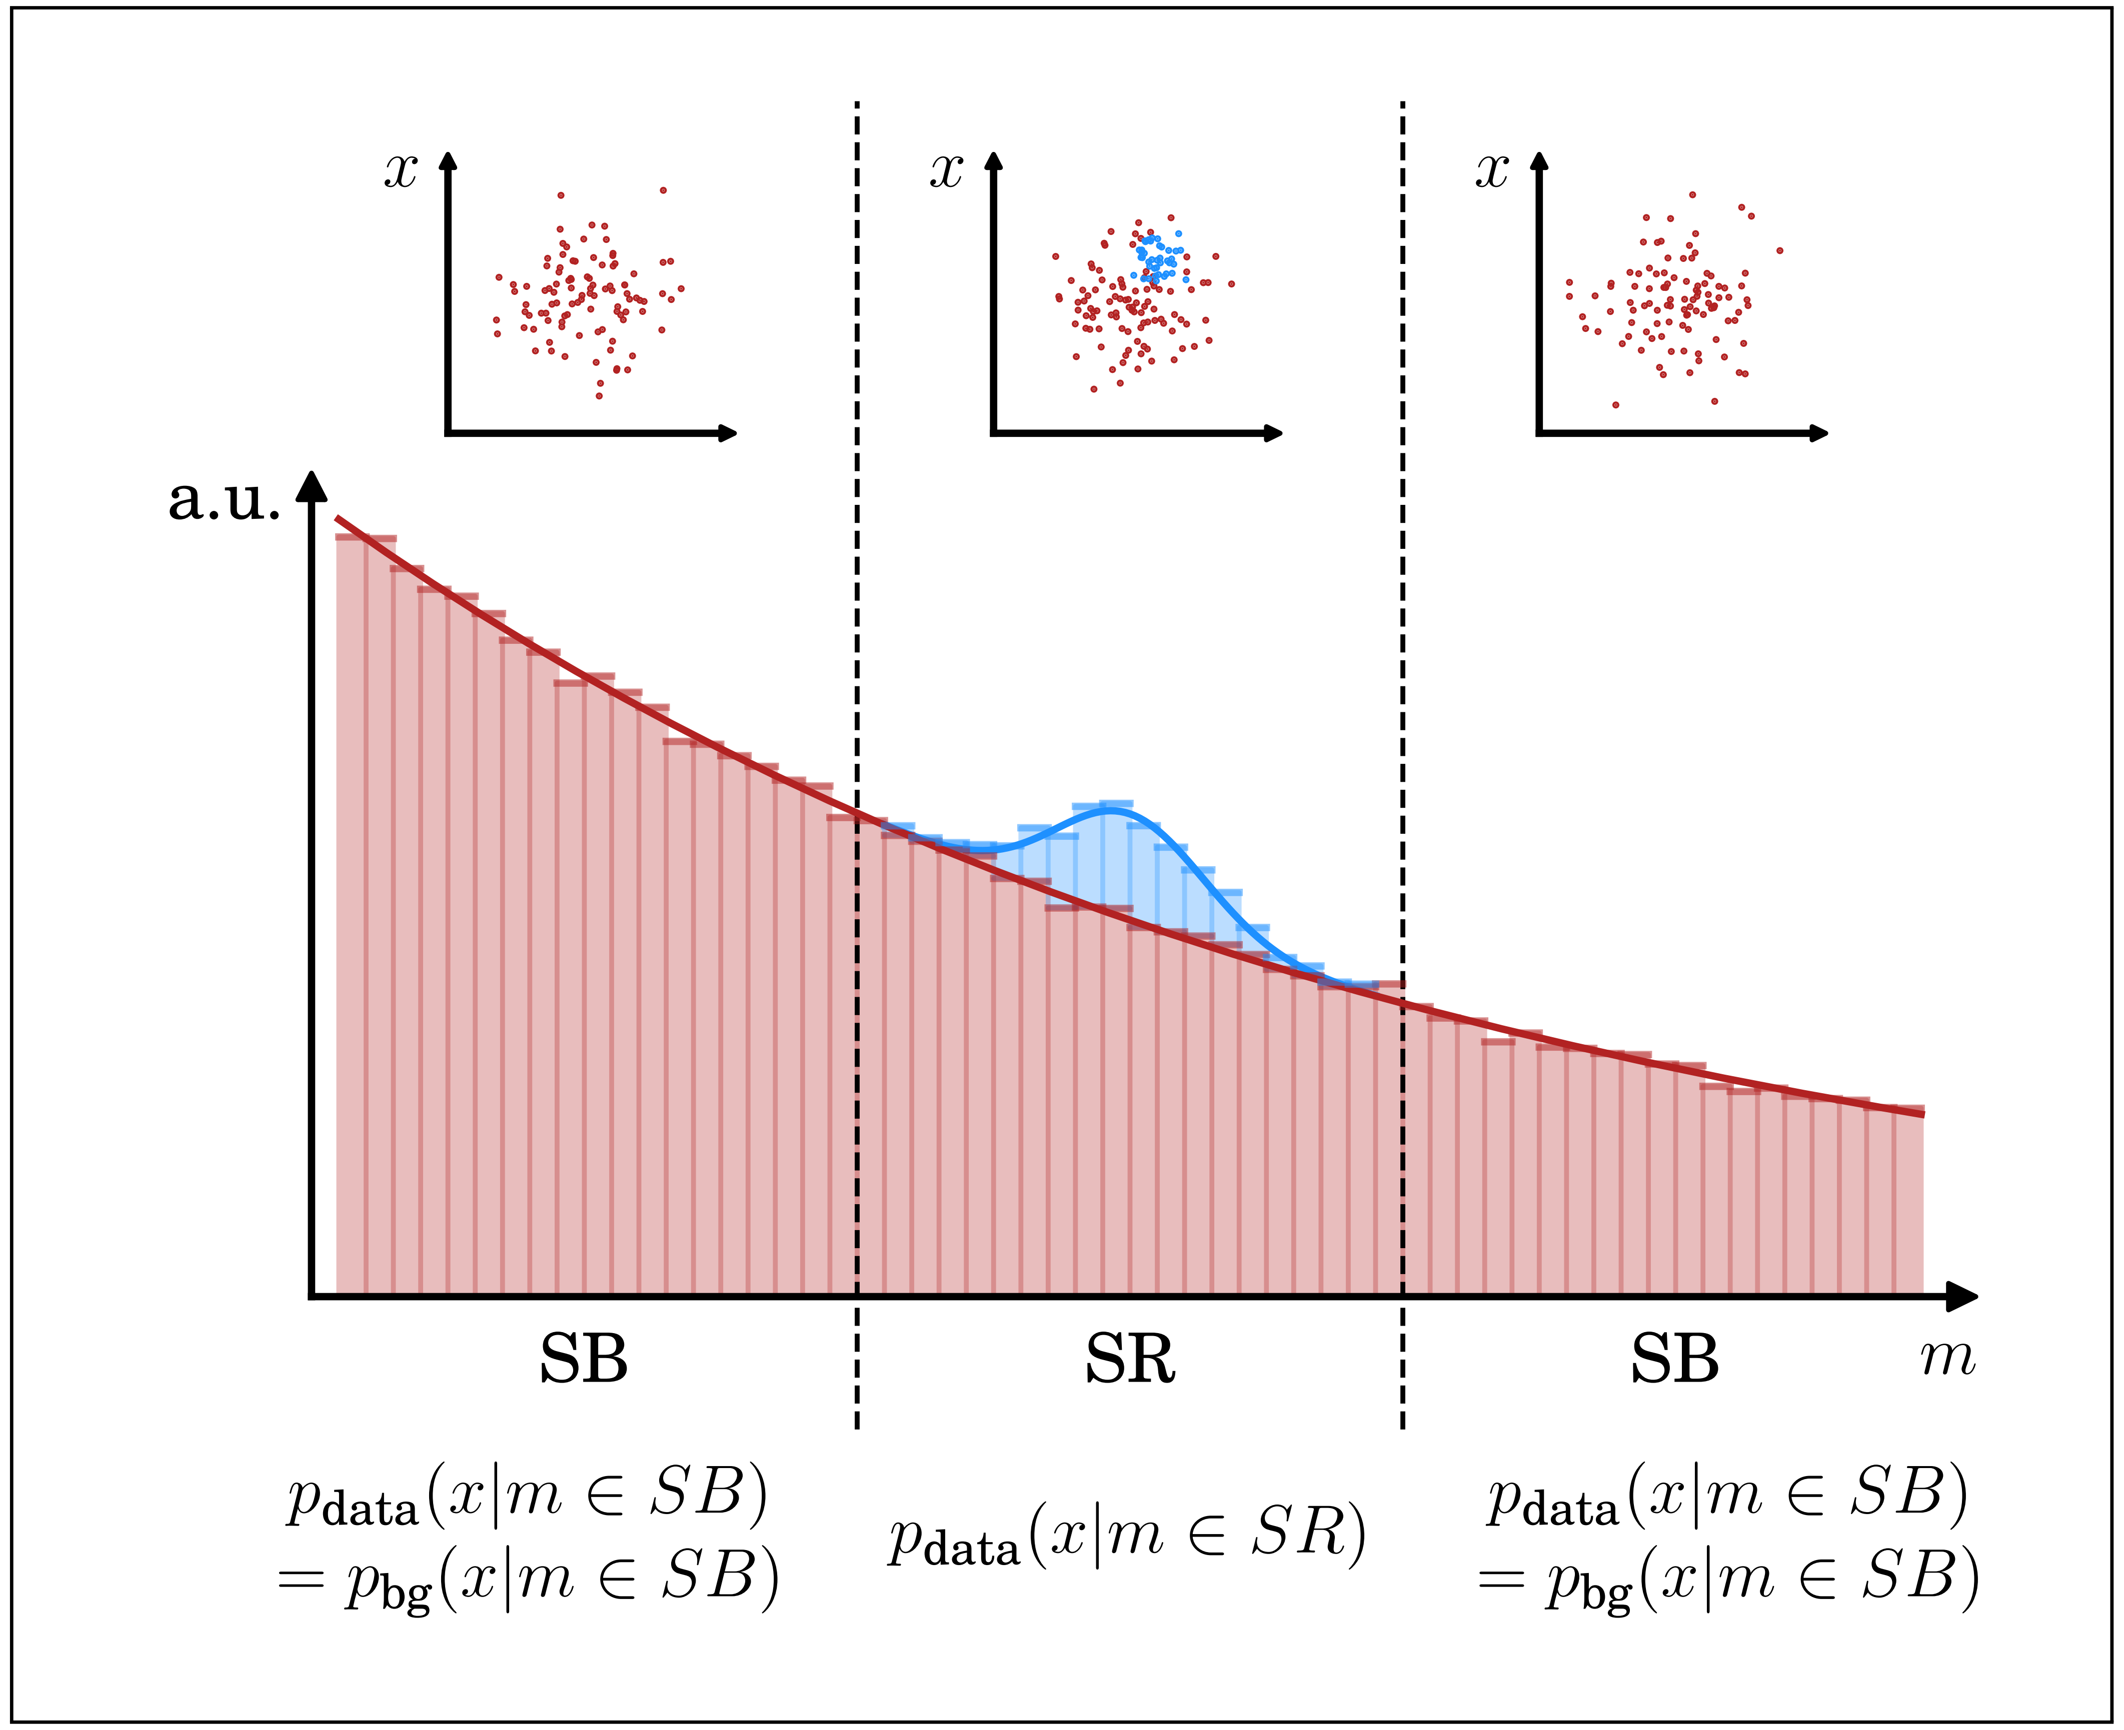

- CWoLa Hunting: **classifier** between SB and SR data
- ANODE: two **conditional density estimators** on SB and SR data; interpolate SB density estimator into SR
- CATHODE: single **conditional density estimator** on SB data; sample interpolated SB density estimator in SR; **classifier** between sampled events and data in SR
- Many other approaches also proposed!
  - CURTAINS: invertible NN for SB->SB interpolation [Raine et al 2203.09470]
  - Simulation assisted resonant anomaly detection: SALAD [Andreassen, Nachman & **DS** 2001.05001], SA-CWoLa [Benkendorfer et al 2009.02205], FETA [Golling et al 2212.11285]

Comparison of methods

- Simulated with Pythia8 + Delphes
- pT(J1)>1.2 TeV trigger
- 4-vectors of every reconstructed particle in the event

We compared the methods on a common toy dataset:

LHC Olympics 2020 R&D dataset [https://doi.org/10.5281/zenodo.2629072]

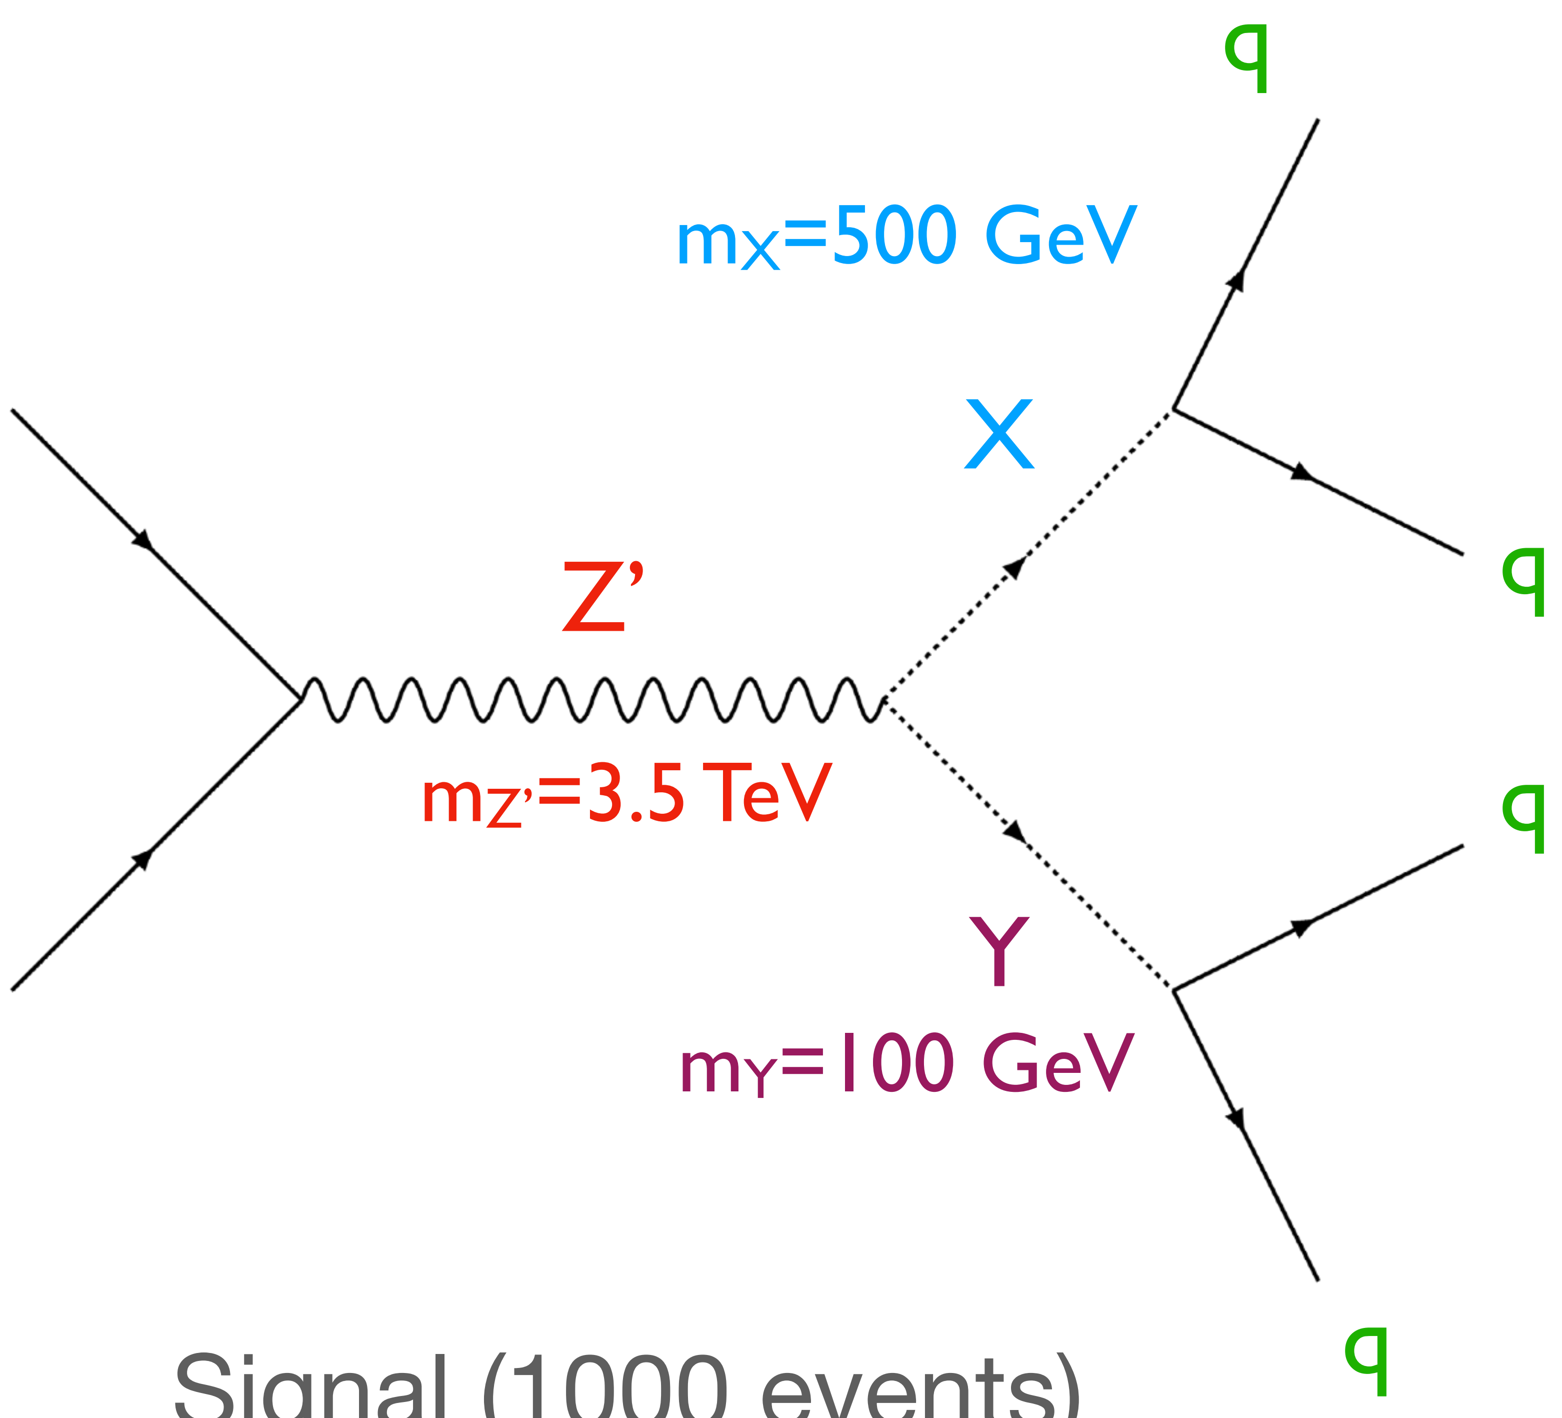

Signal (1000 events)

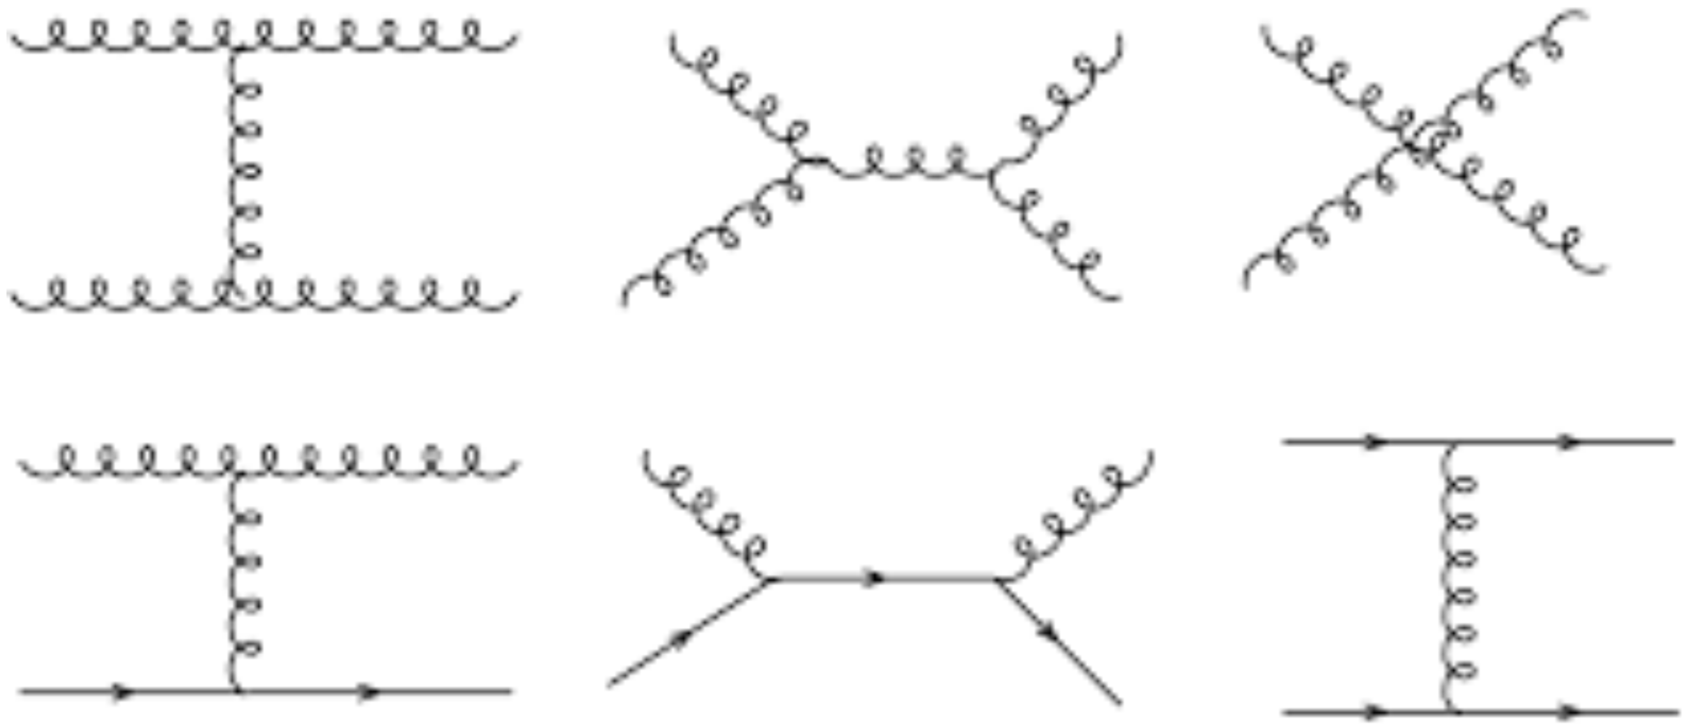

Background: QCD dijets (1M events)

No explicit search at the LHC for this scenario!

LHCO2020 R&D Dataset

Resonant feature

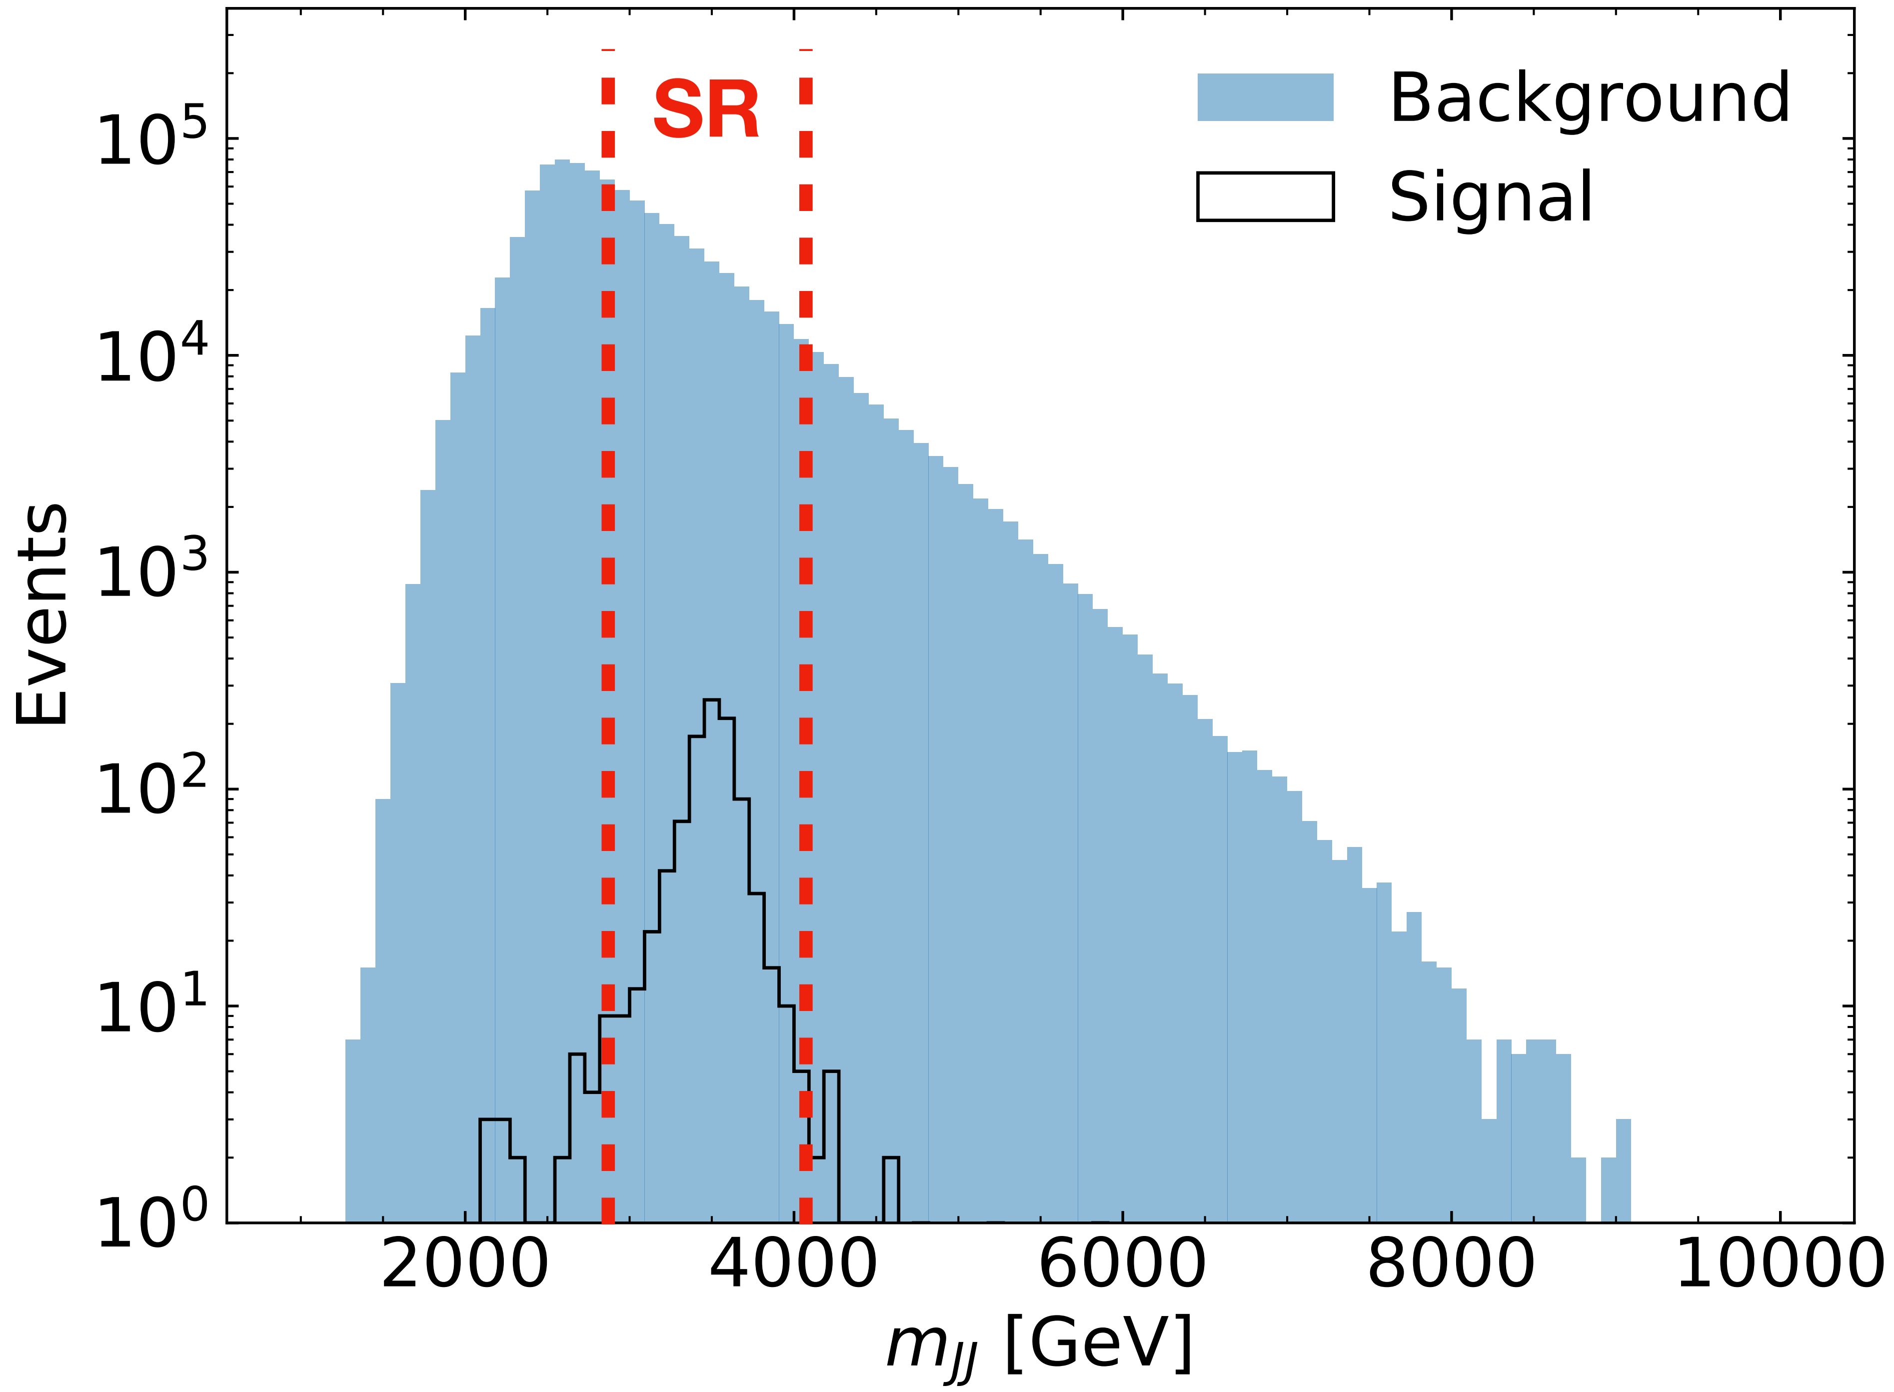

$$m = m_{JJ}$$

Additional features

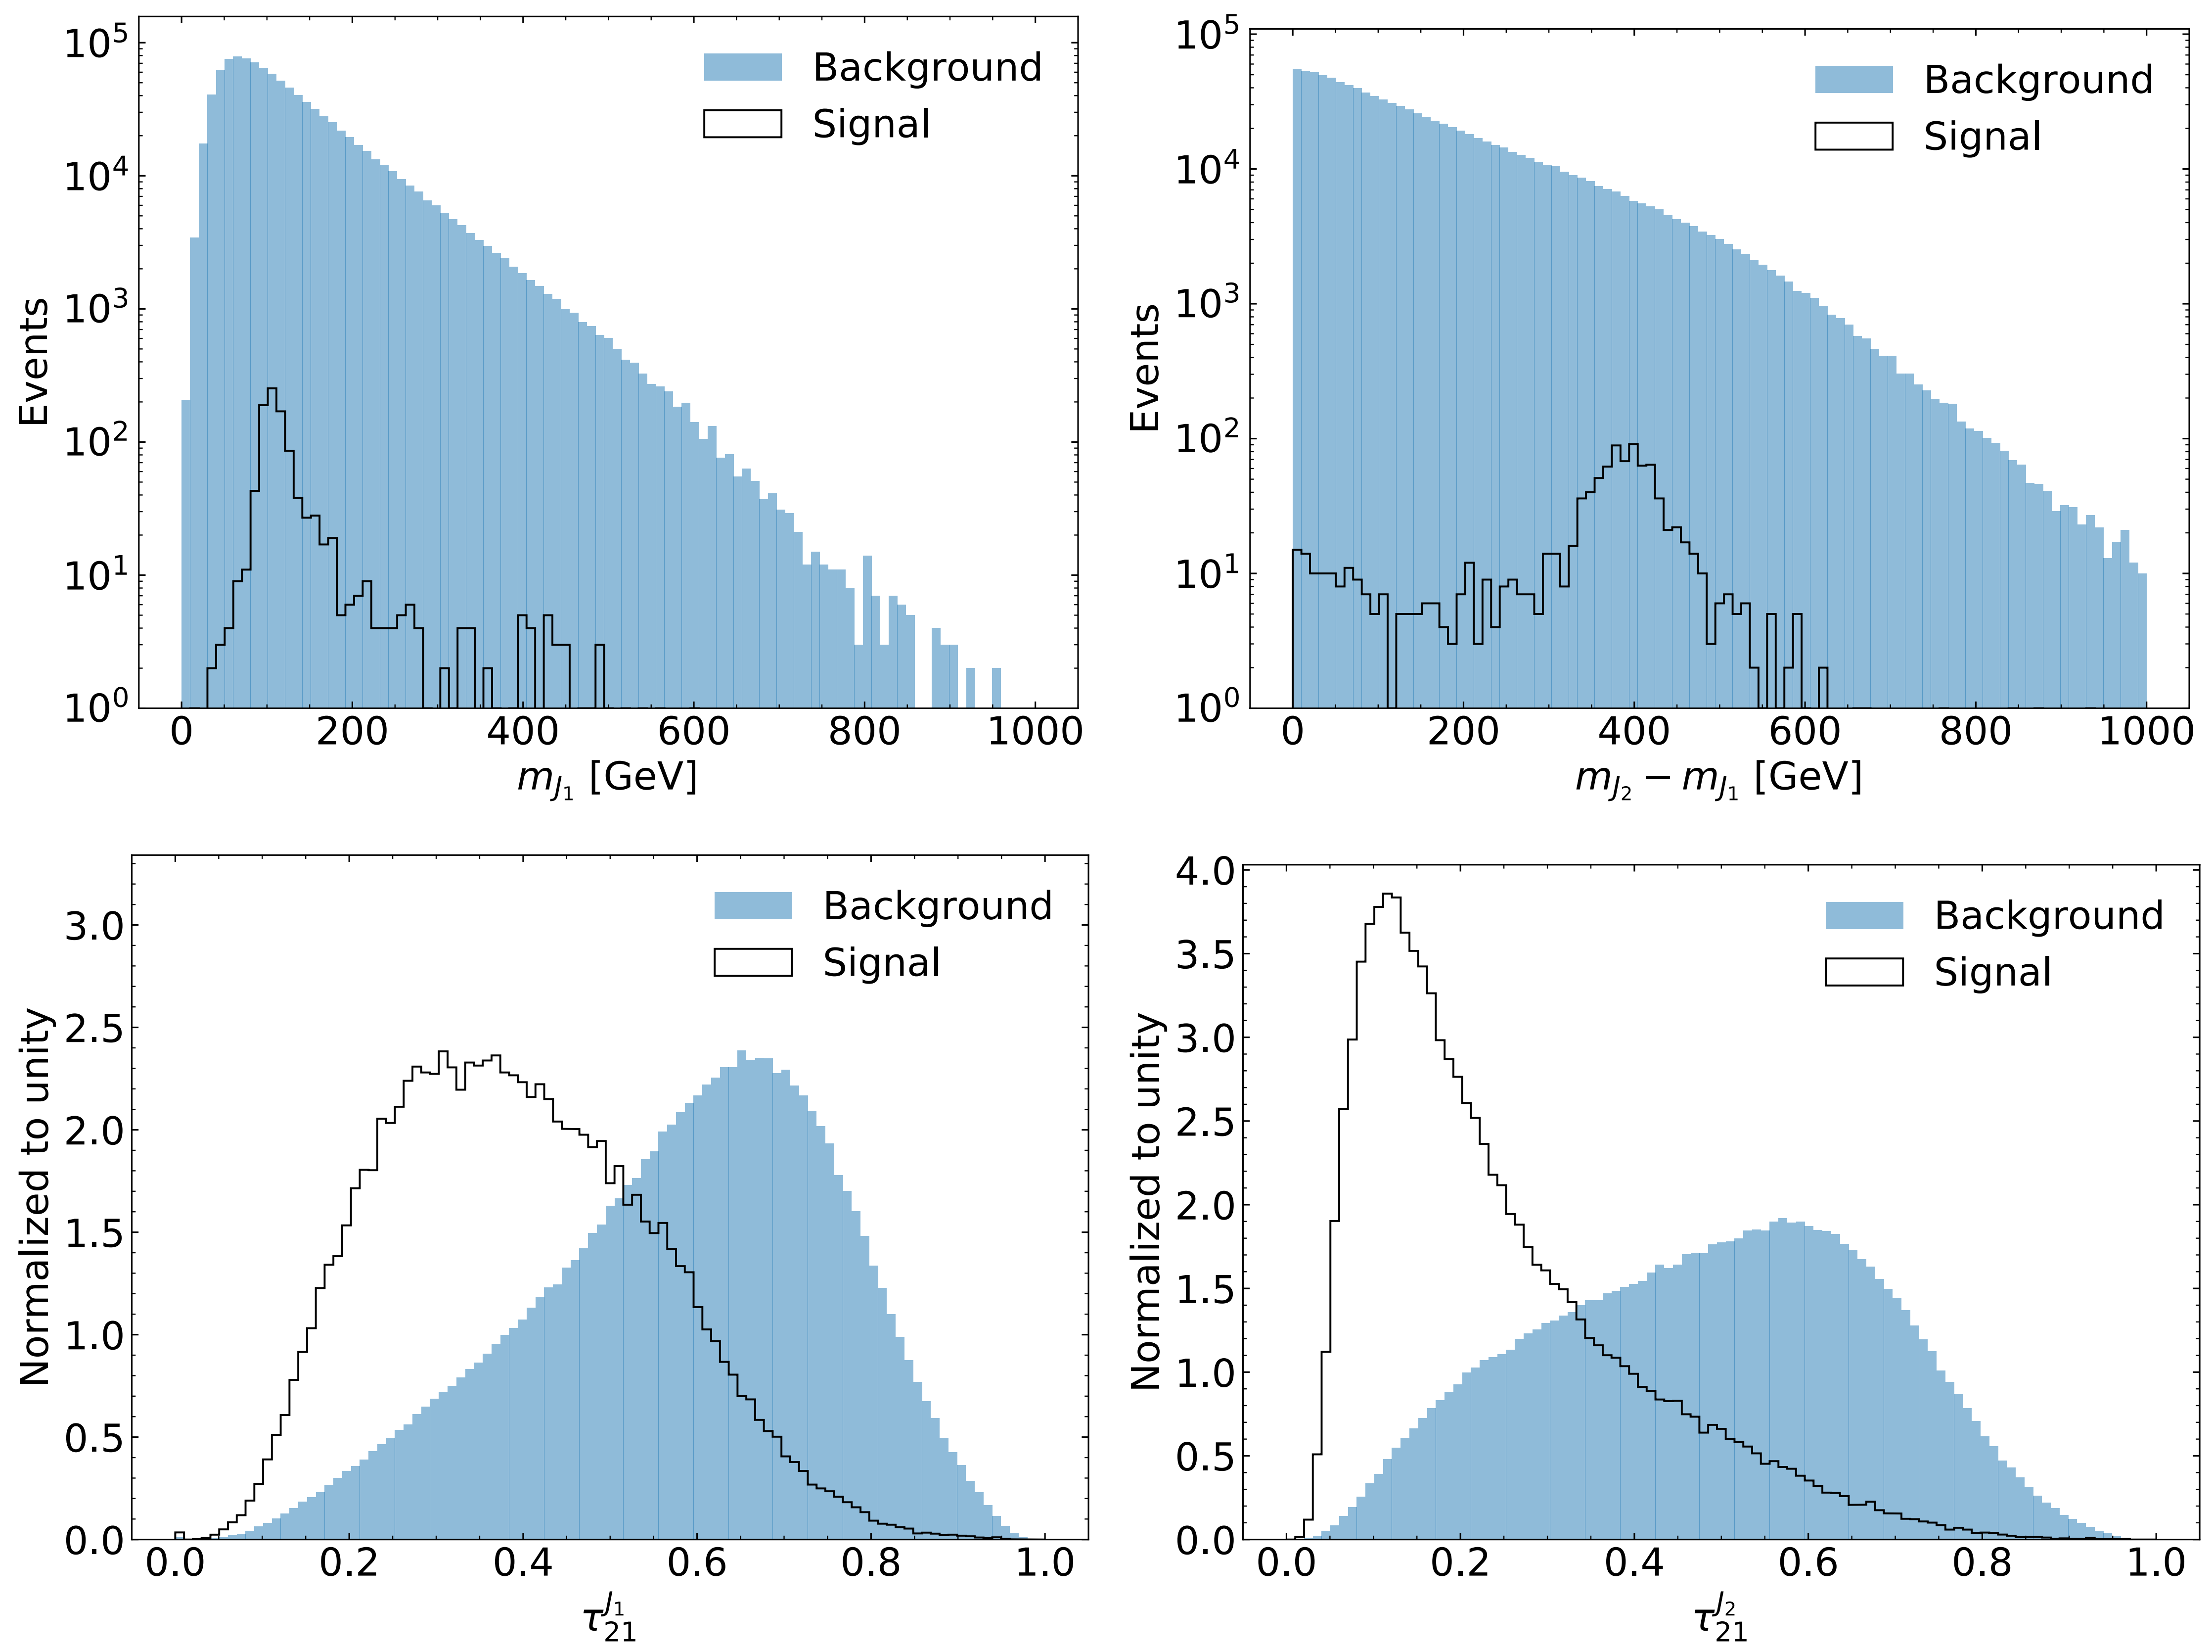

$$x = (m_{J_1}, m_{J_2}, \tau_{21}^{J_1}, \tau_{21}^{J_2})$$

Benchmark signal strength: S_{SR}=772, B_{SR}=120,000

S_{SR}/B_{SR}~0.6%, $S_{SR}/\sqrt{B_{SR}}$~2.2

For further comparisons of these methods and more, see 2307.11157!

Comparison of methods

Significance improvement characteristic (SIC): $\epsilon_S/\sqrt{\epsilon_B}$

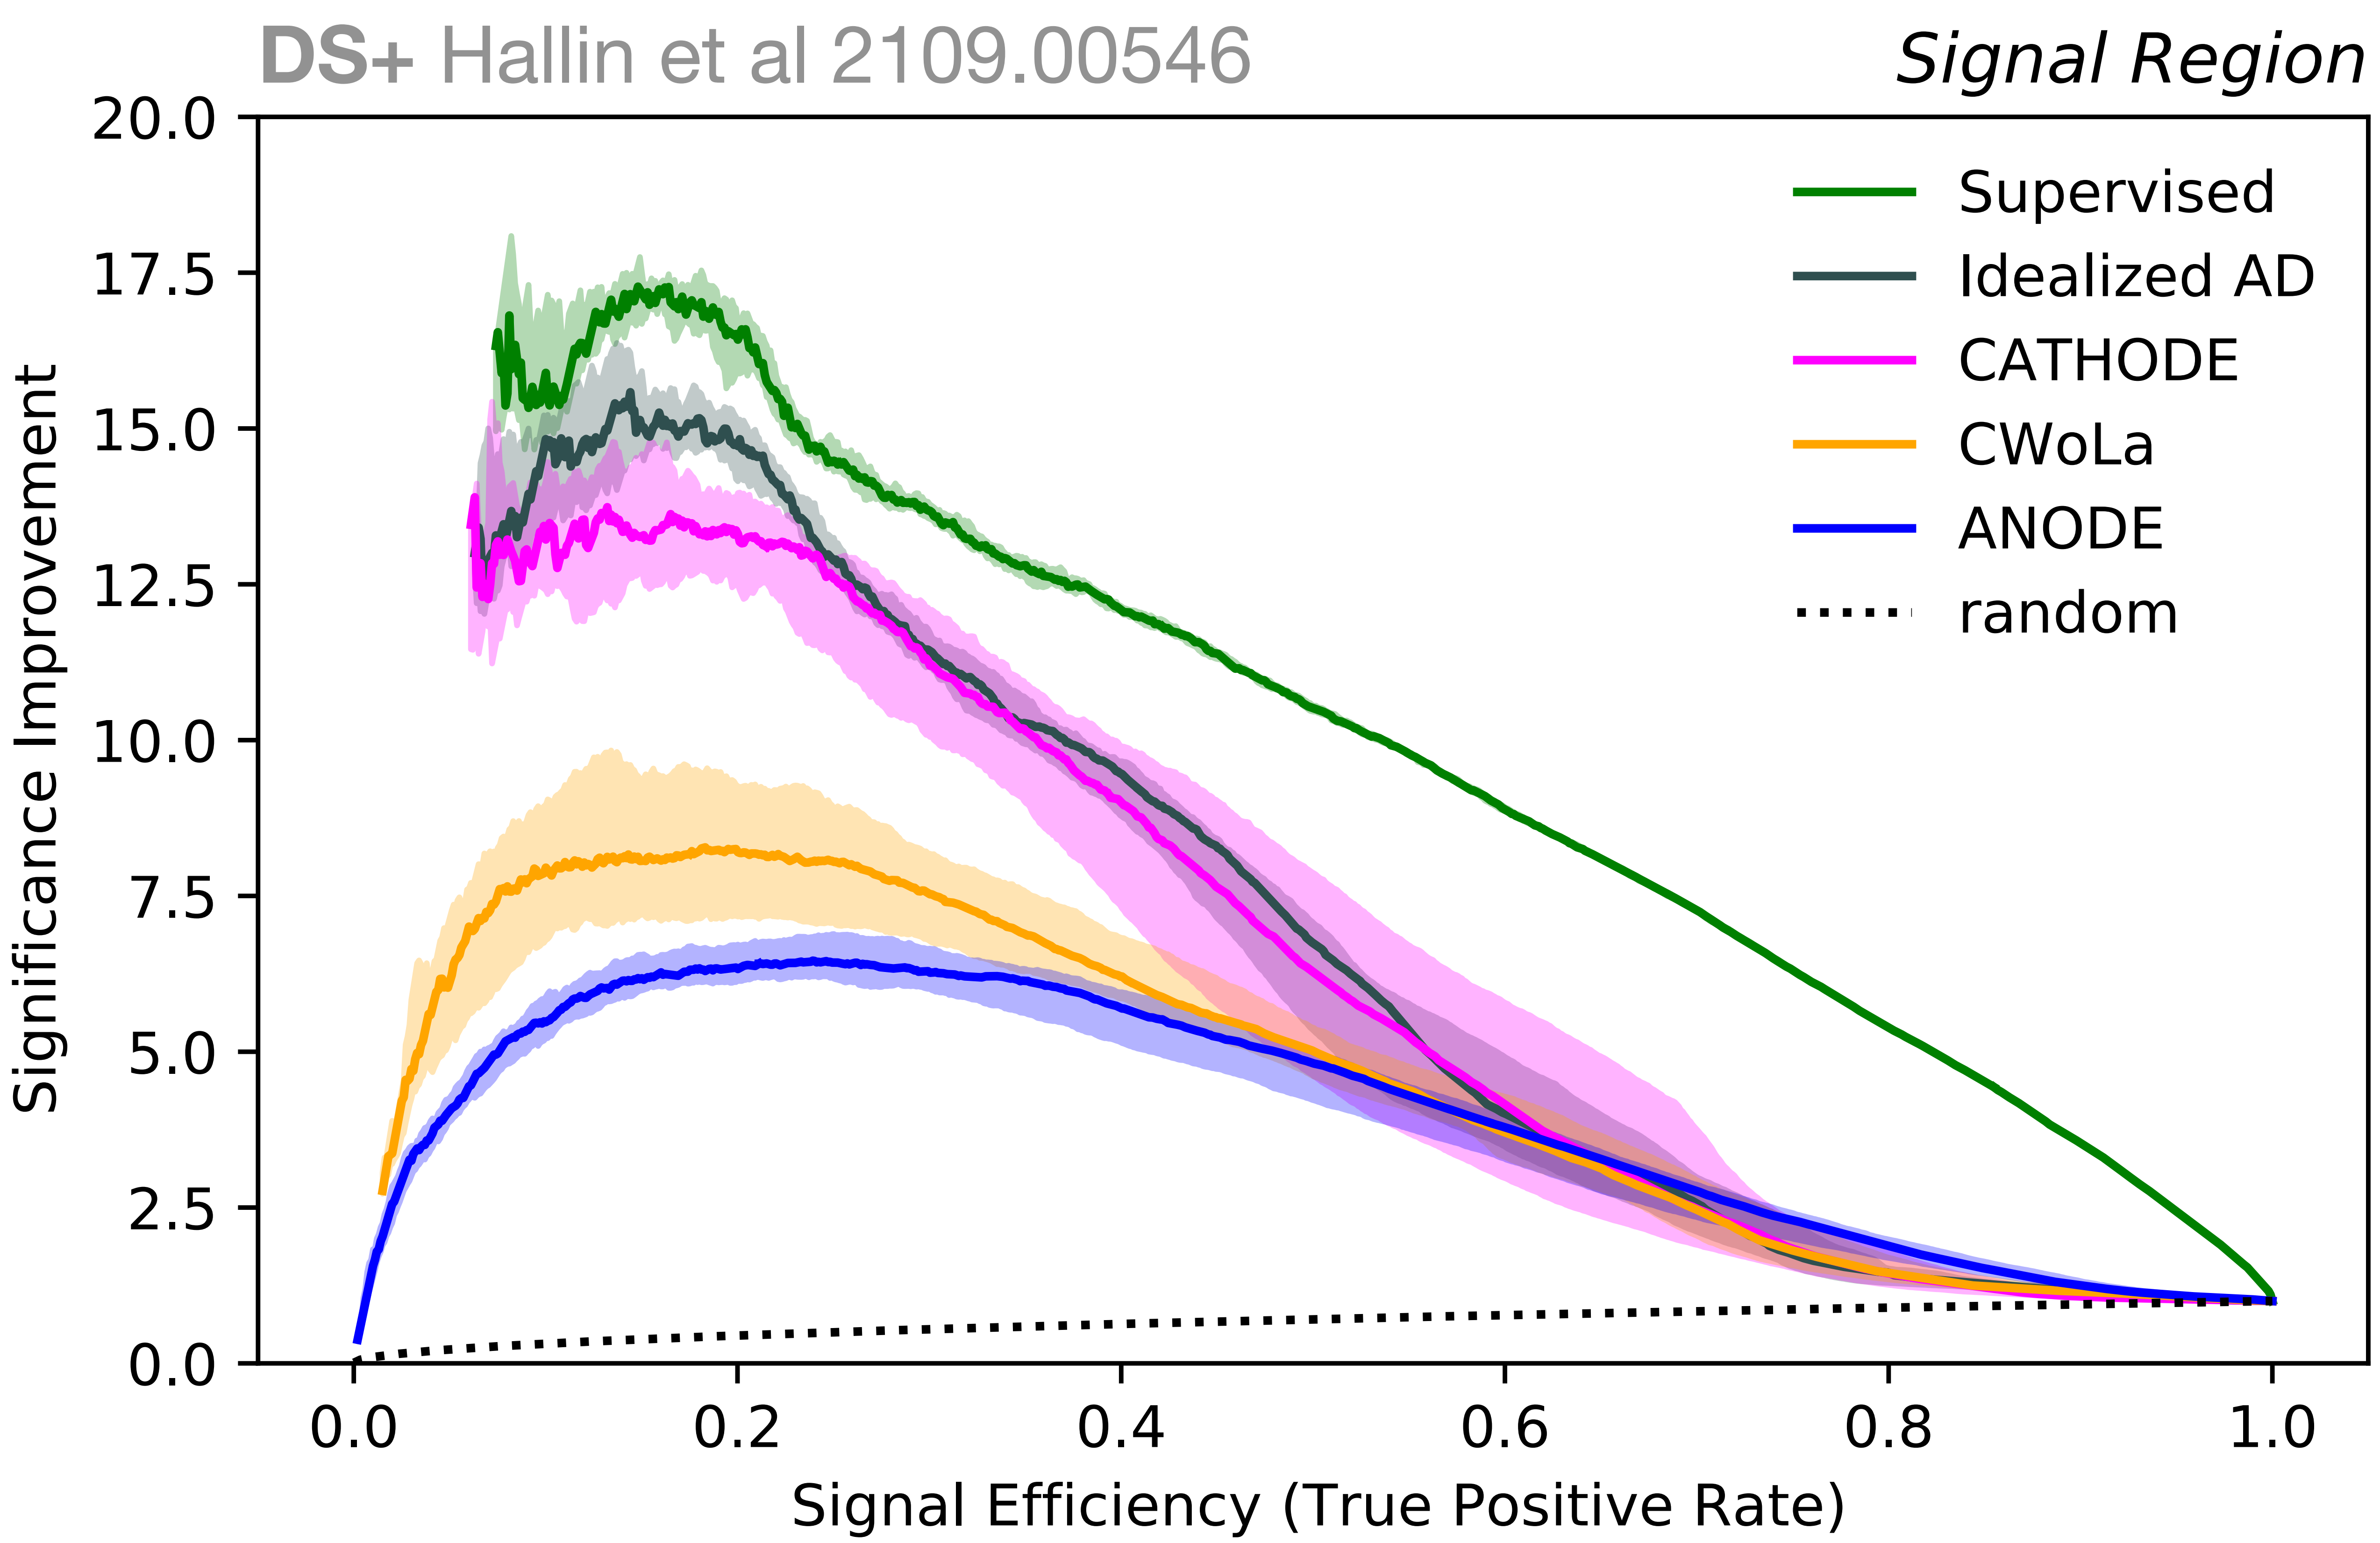

For further comparisons of these methods and more, see 2307.11157!

Comparison of methods

Significance improvement characteristic (SIC): $\epsilon_S/\sqrt{\epsilon_B}$

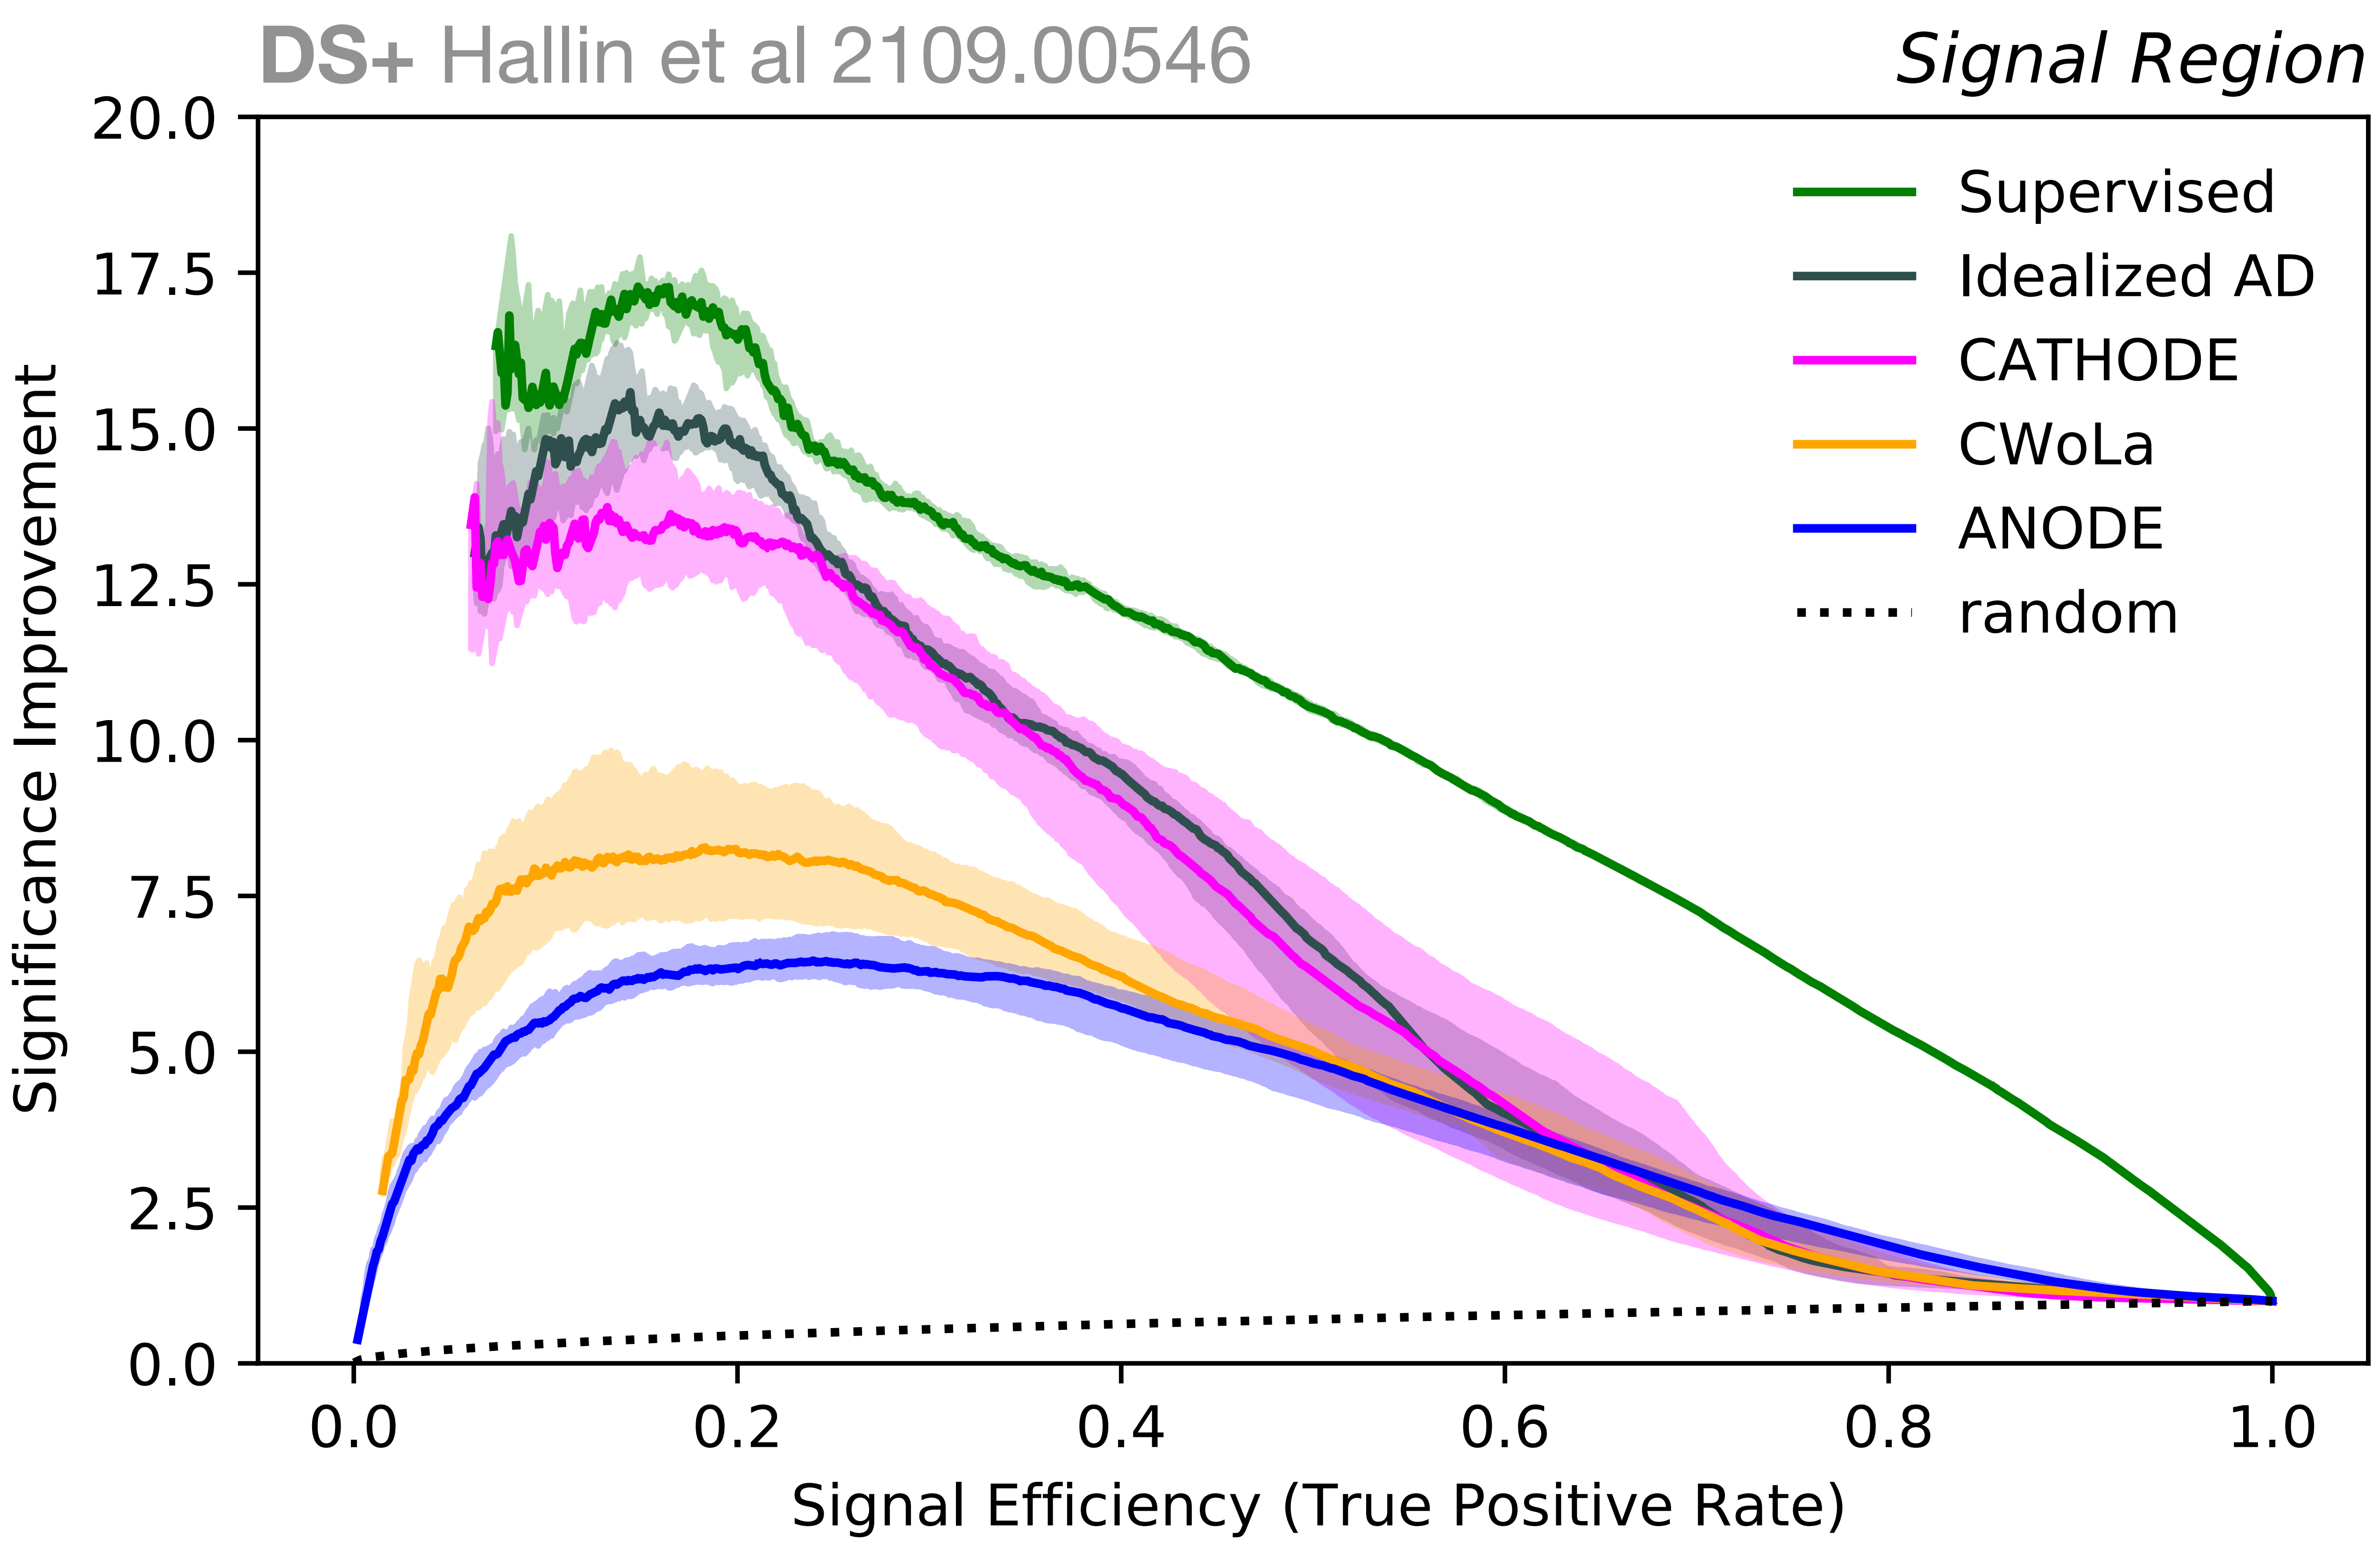

Comparison of methods

For further comparisons of these methods and more, see 2307.11157!

Significance improvement characteristic (SIC): $\epsilon_S / \sqrt{\epsilon_B}$

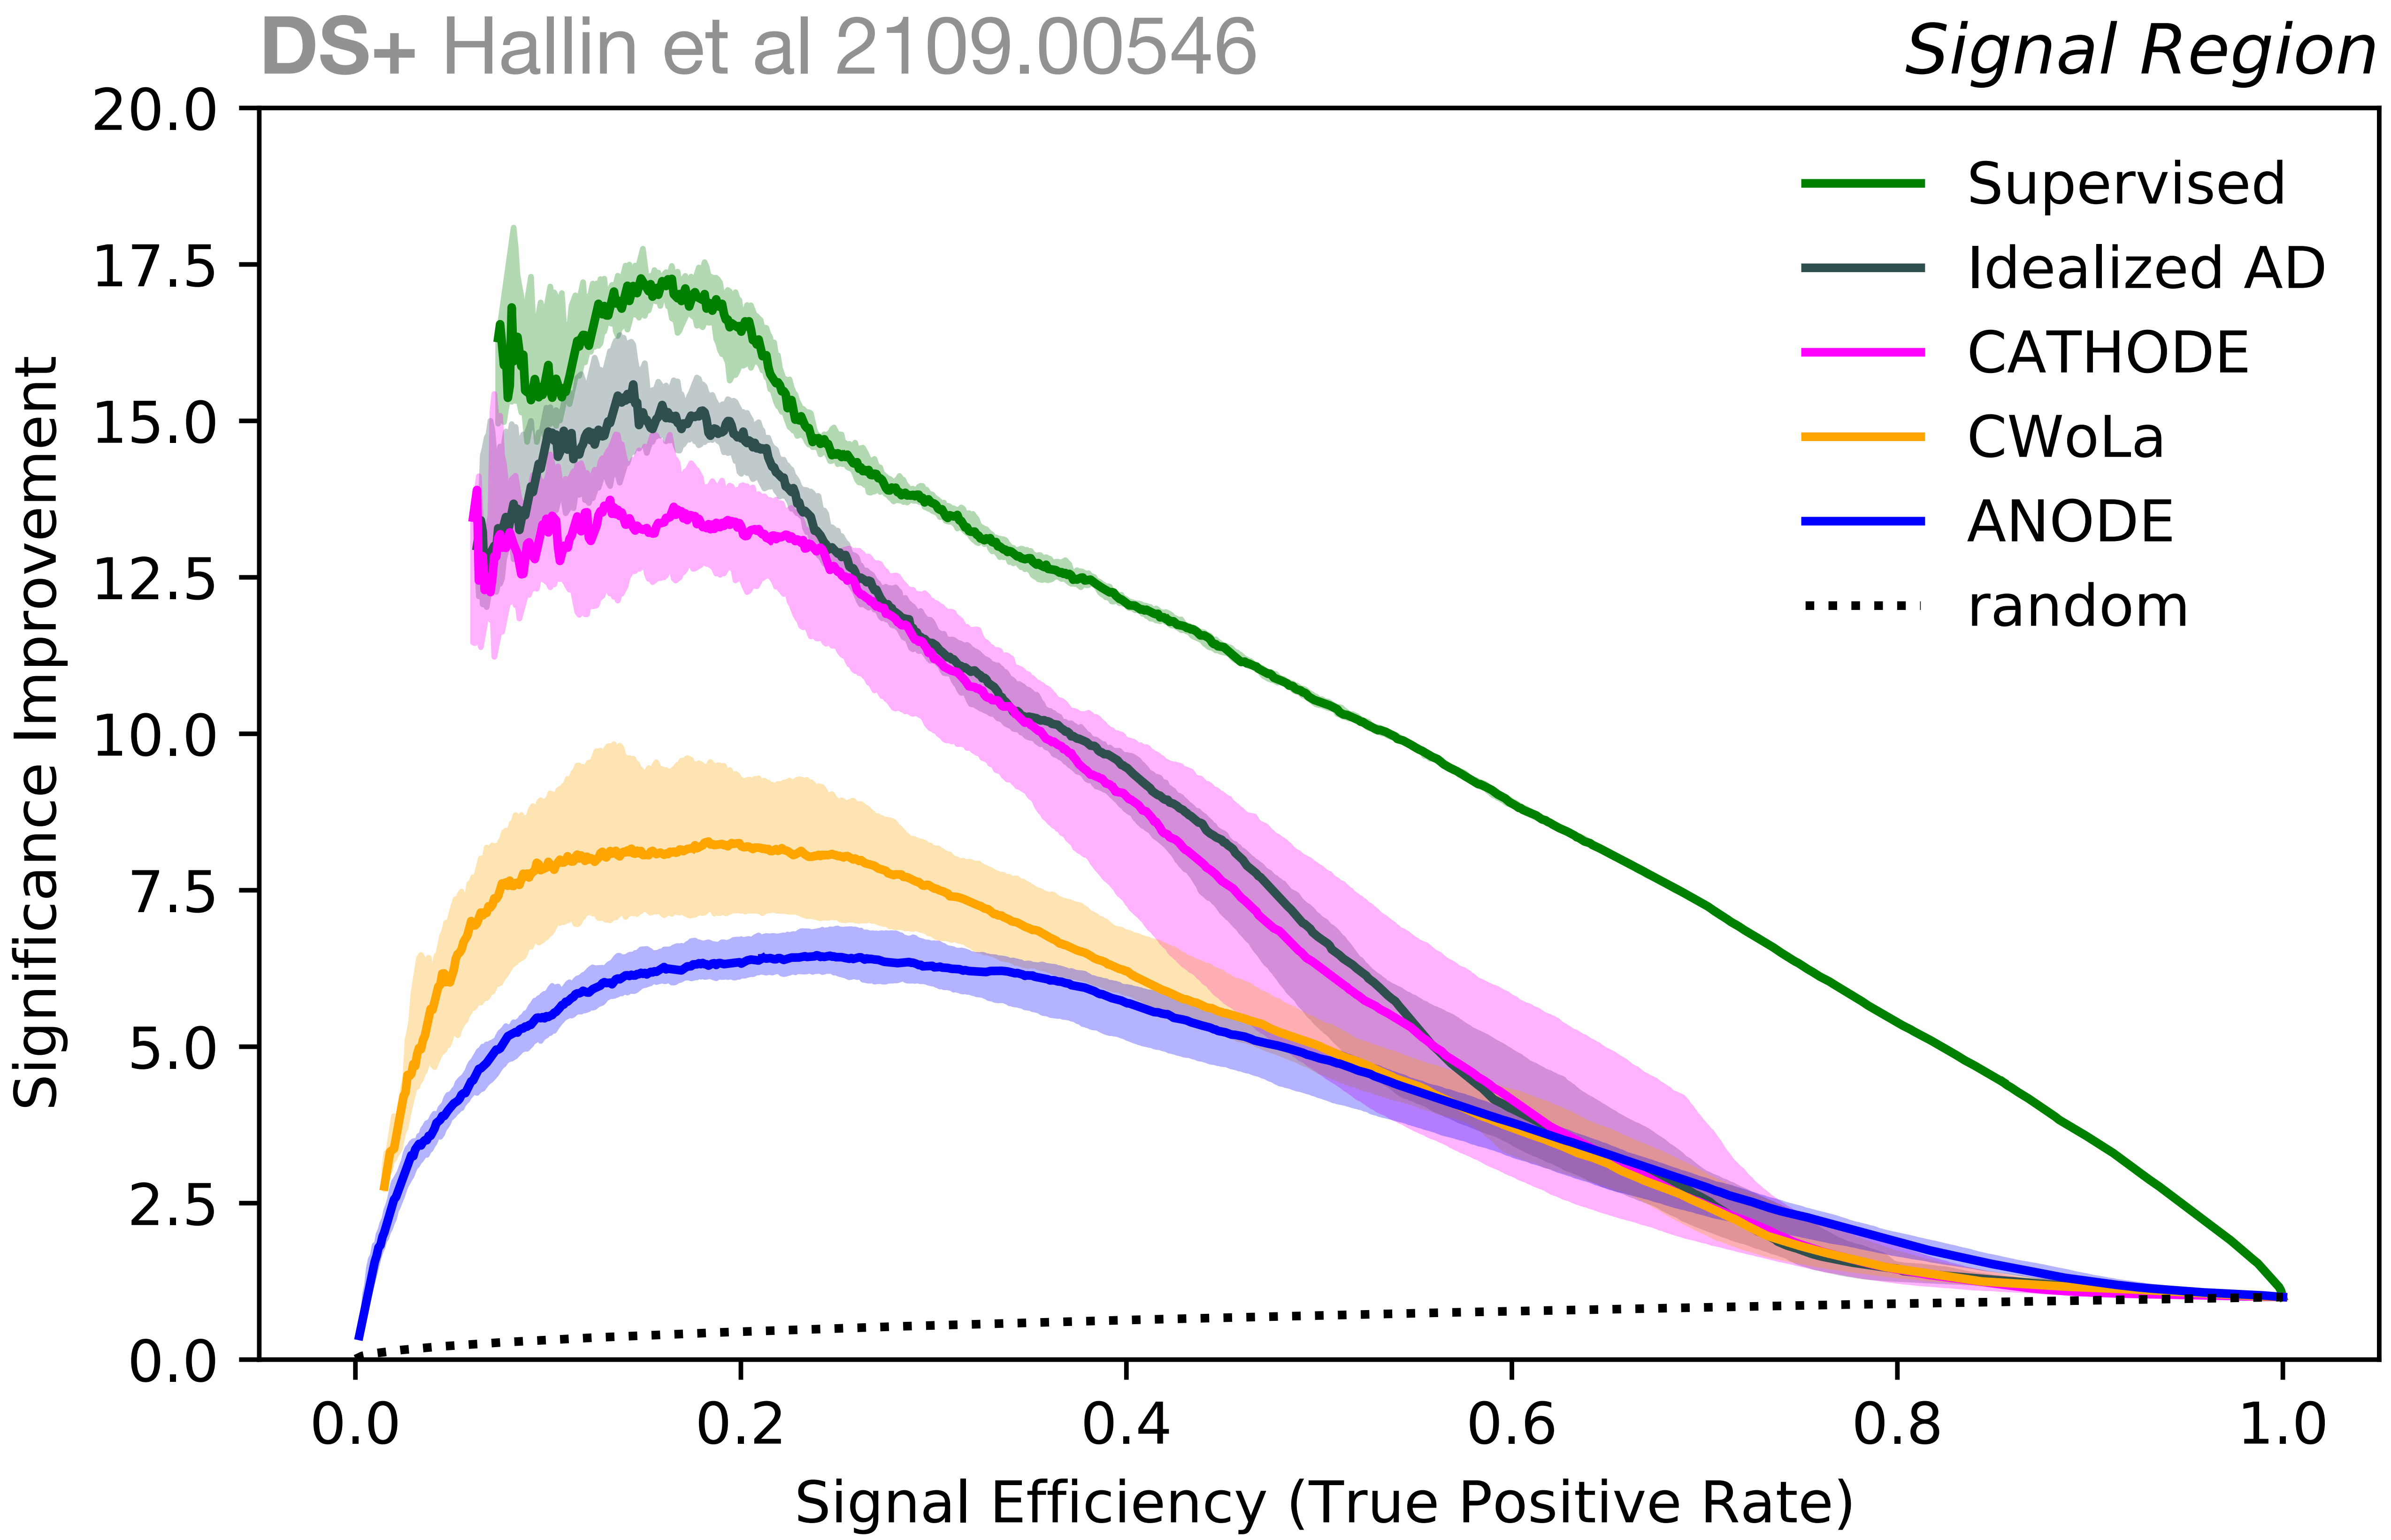

CATHODE outperforms CWoLa and ANODE and nearly saturates the idealized anomaly detector!

Comparison of methods

For further comparisons of these methods and more, see 2307.11157!

Significance improvement characteristic (SIC): $\epsilon_S/\sqrt{\epsilon_B}$

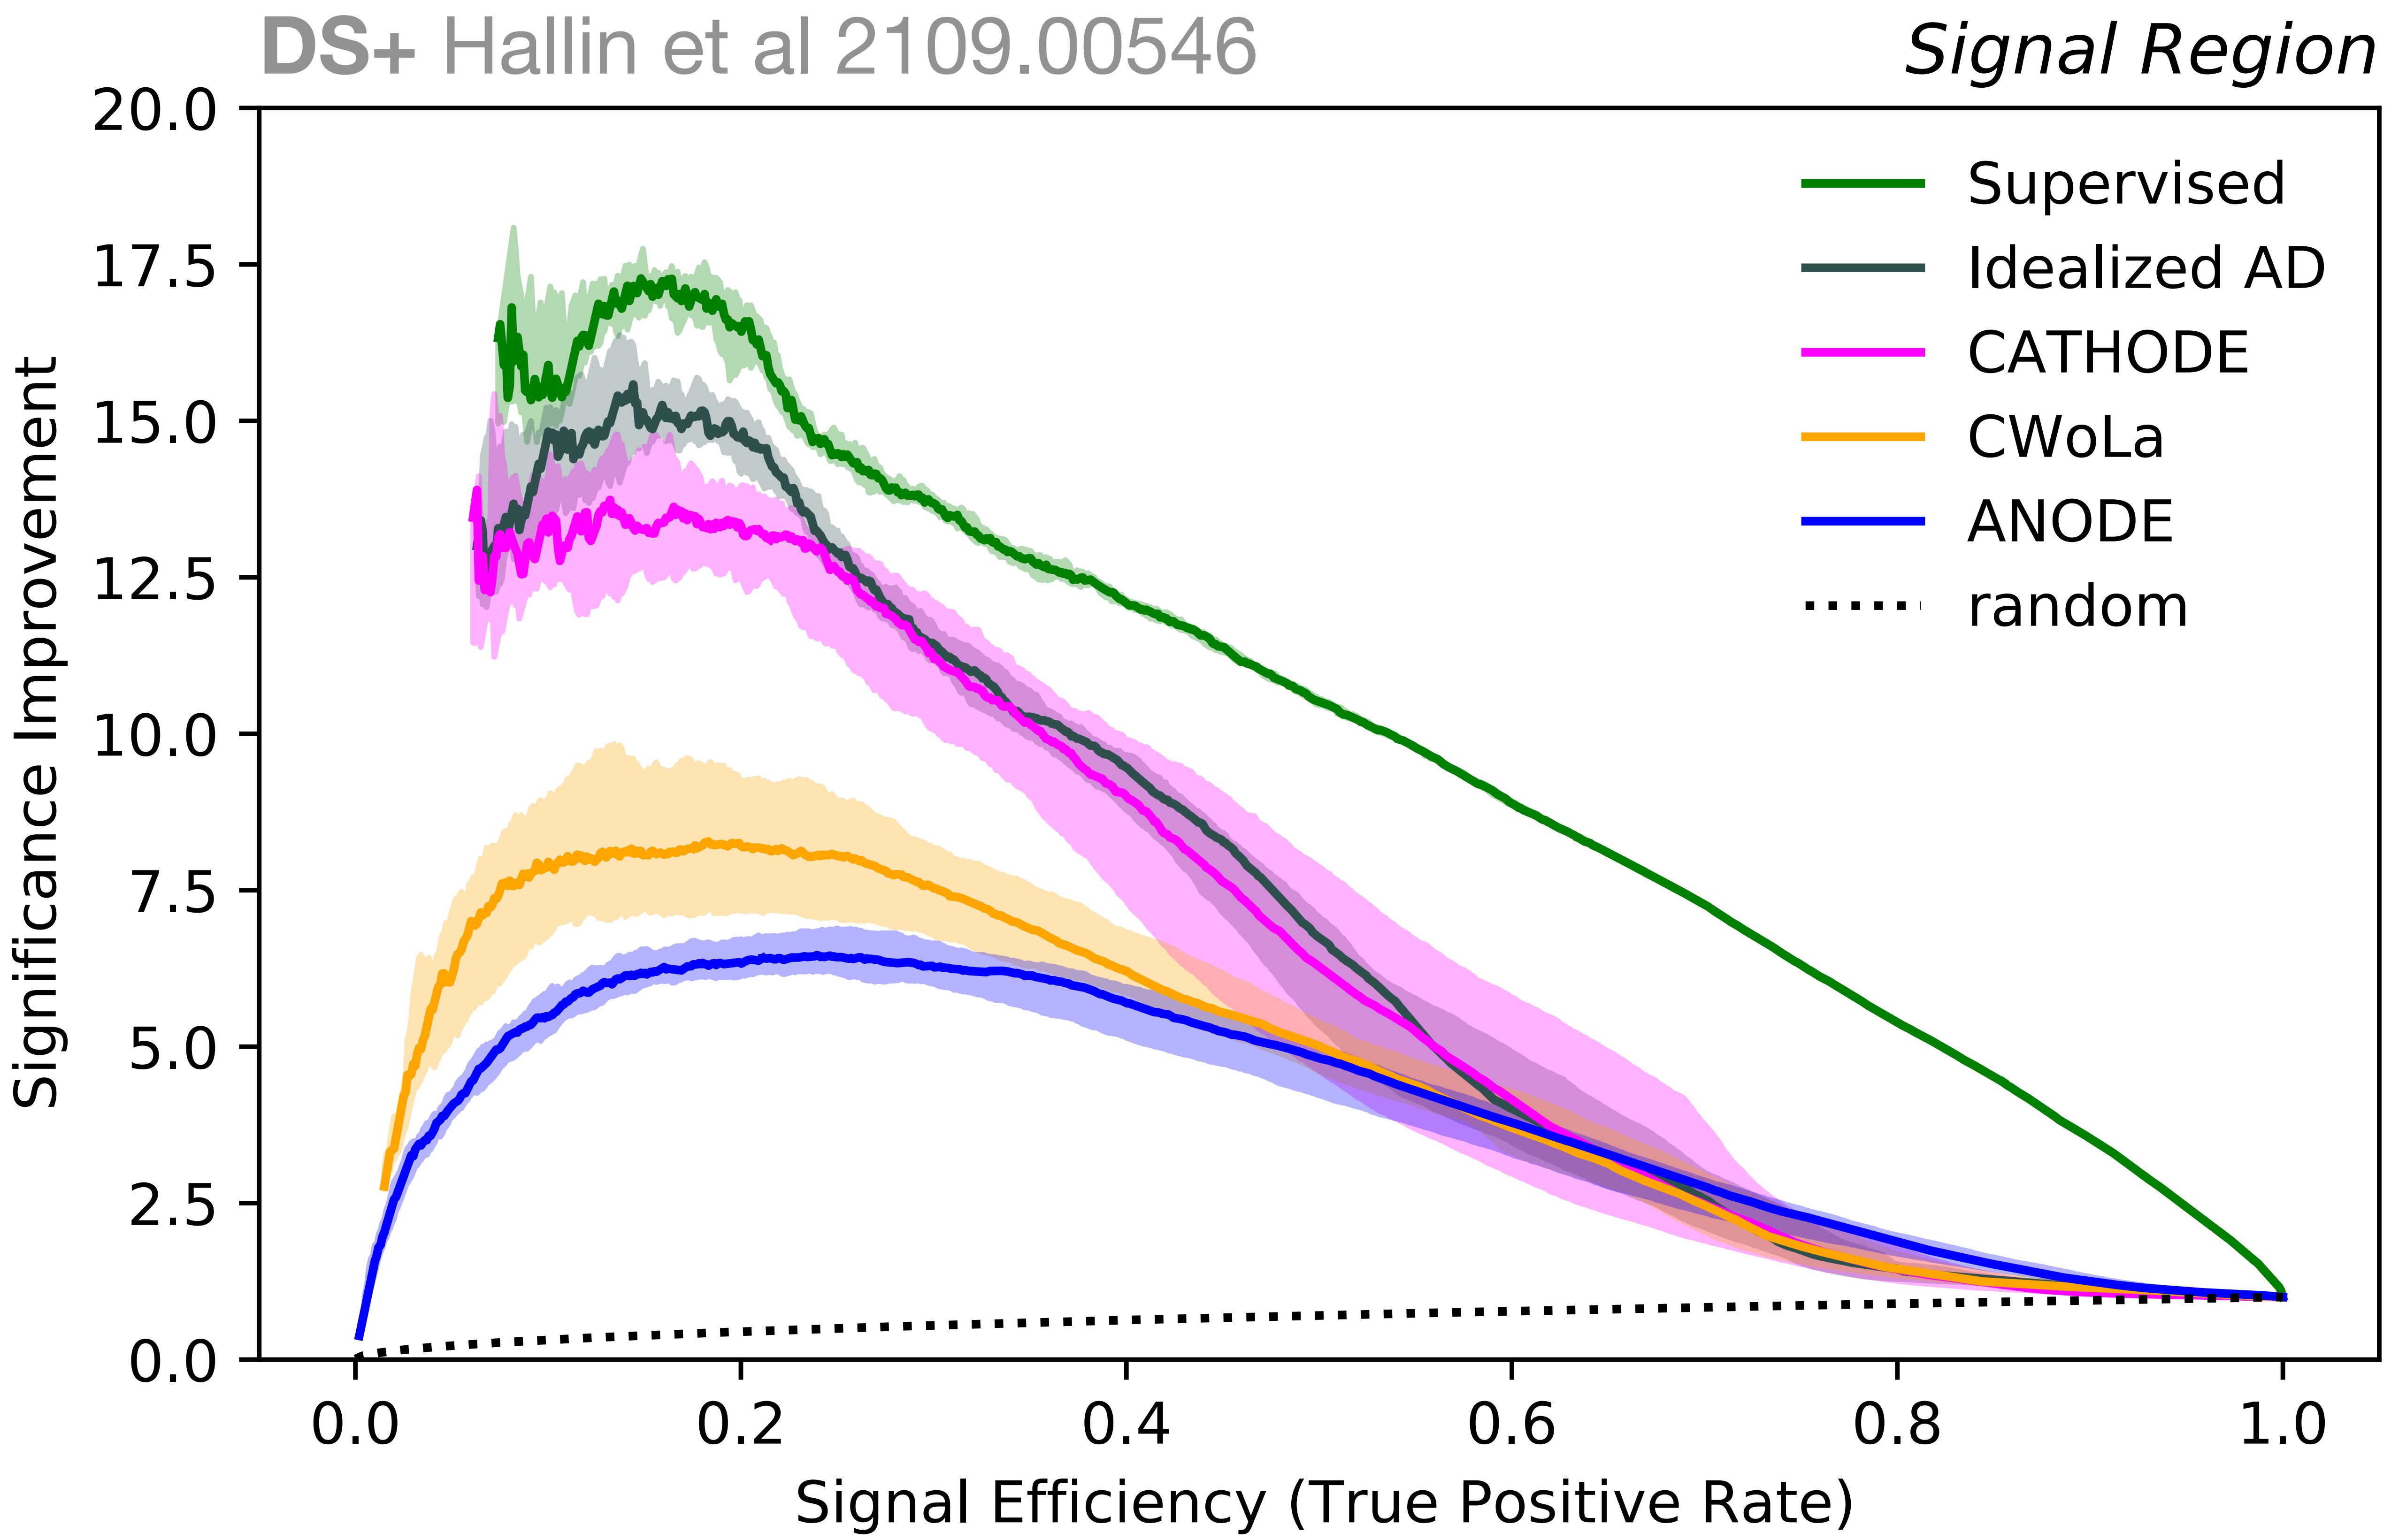

CATHODE outperforms CWoLa and ANODE and nearly saturates the idealized anomaly detector!

Initial significance was ~2.2σ

=> **a ~30σ anomaly could be hiding in the data right now!**

Current frontiers of resonant anomaly detection

Robustness

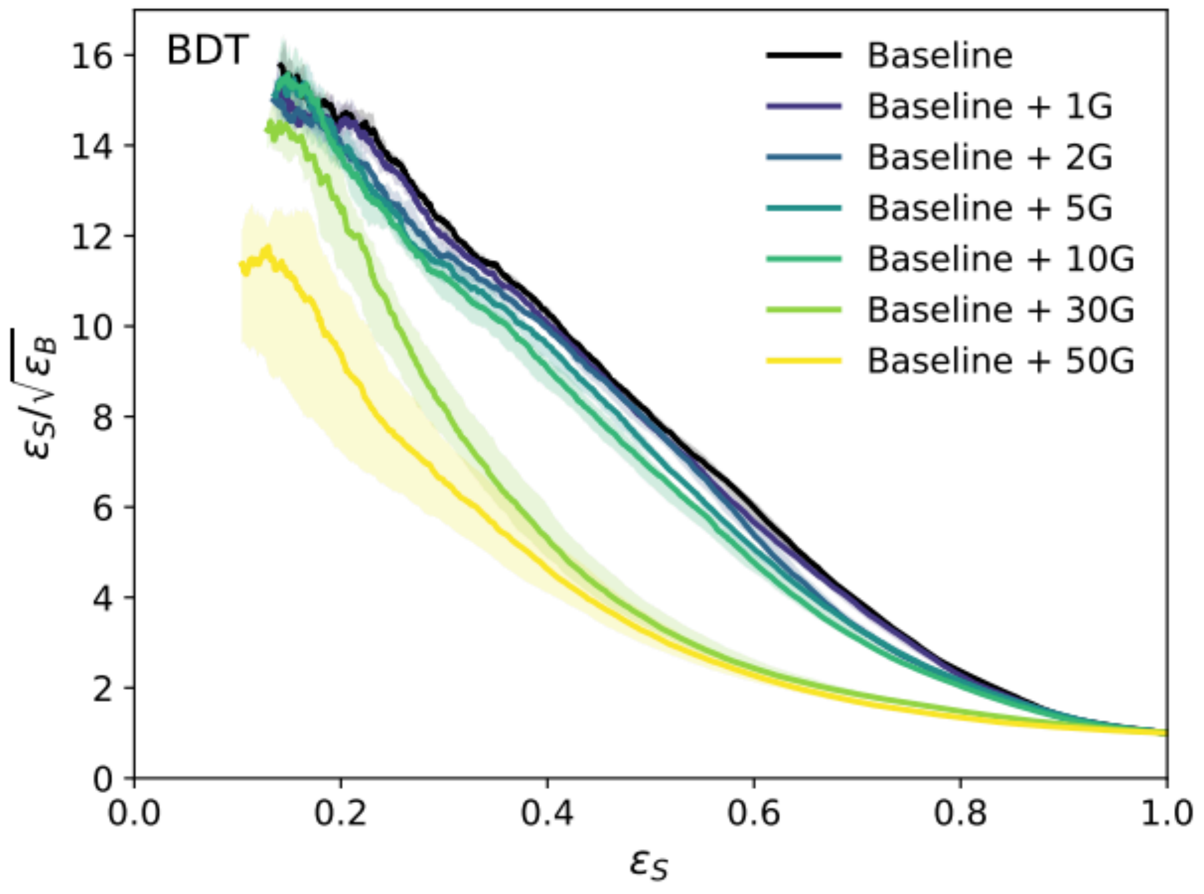

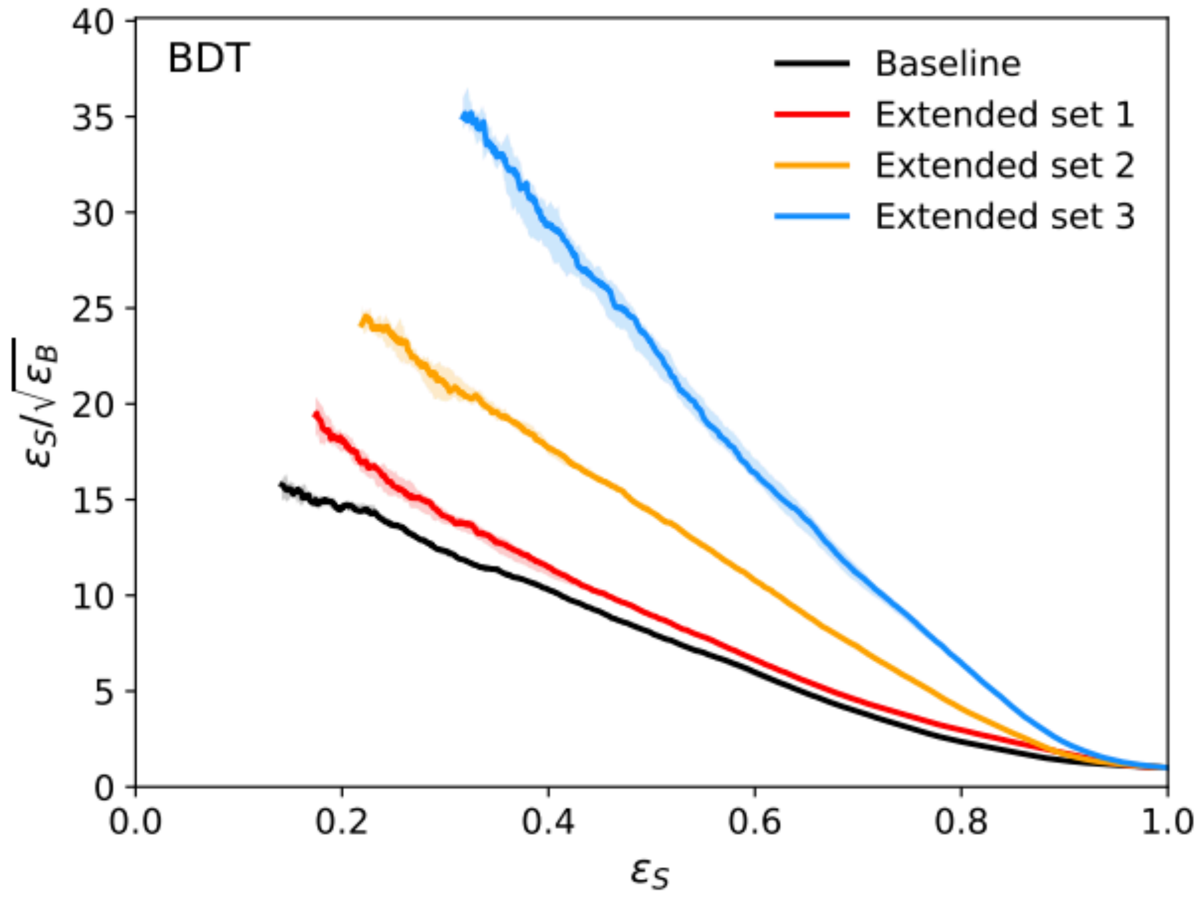

DS+ Finke et al 2309.13111

Full phase space

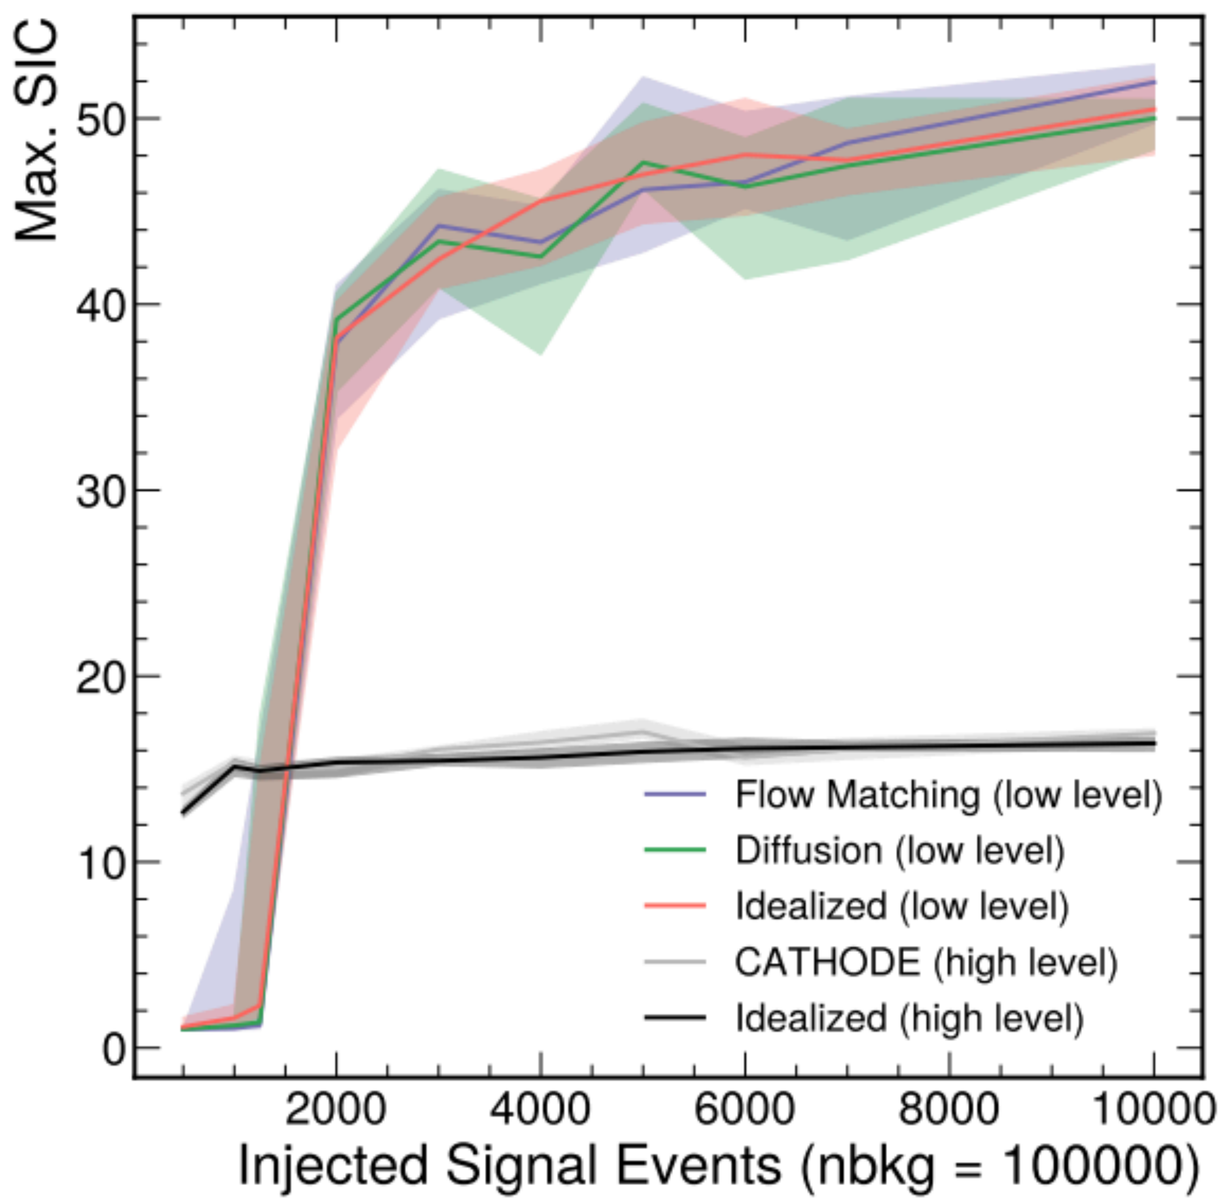

DS+ Buhmann et al 2310.06897

Interpretability

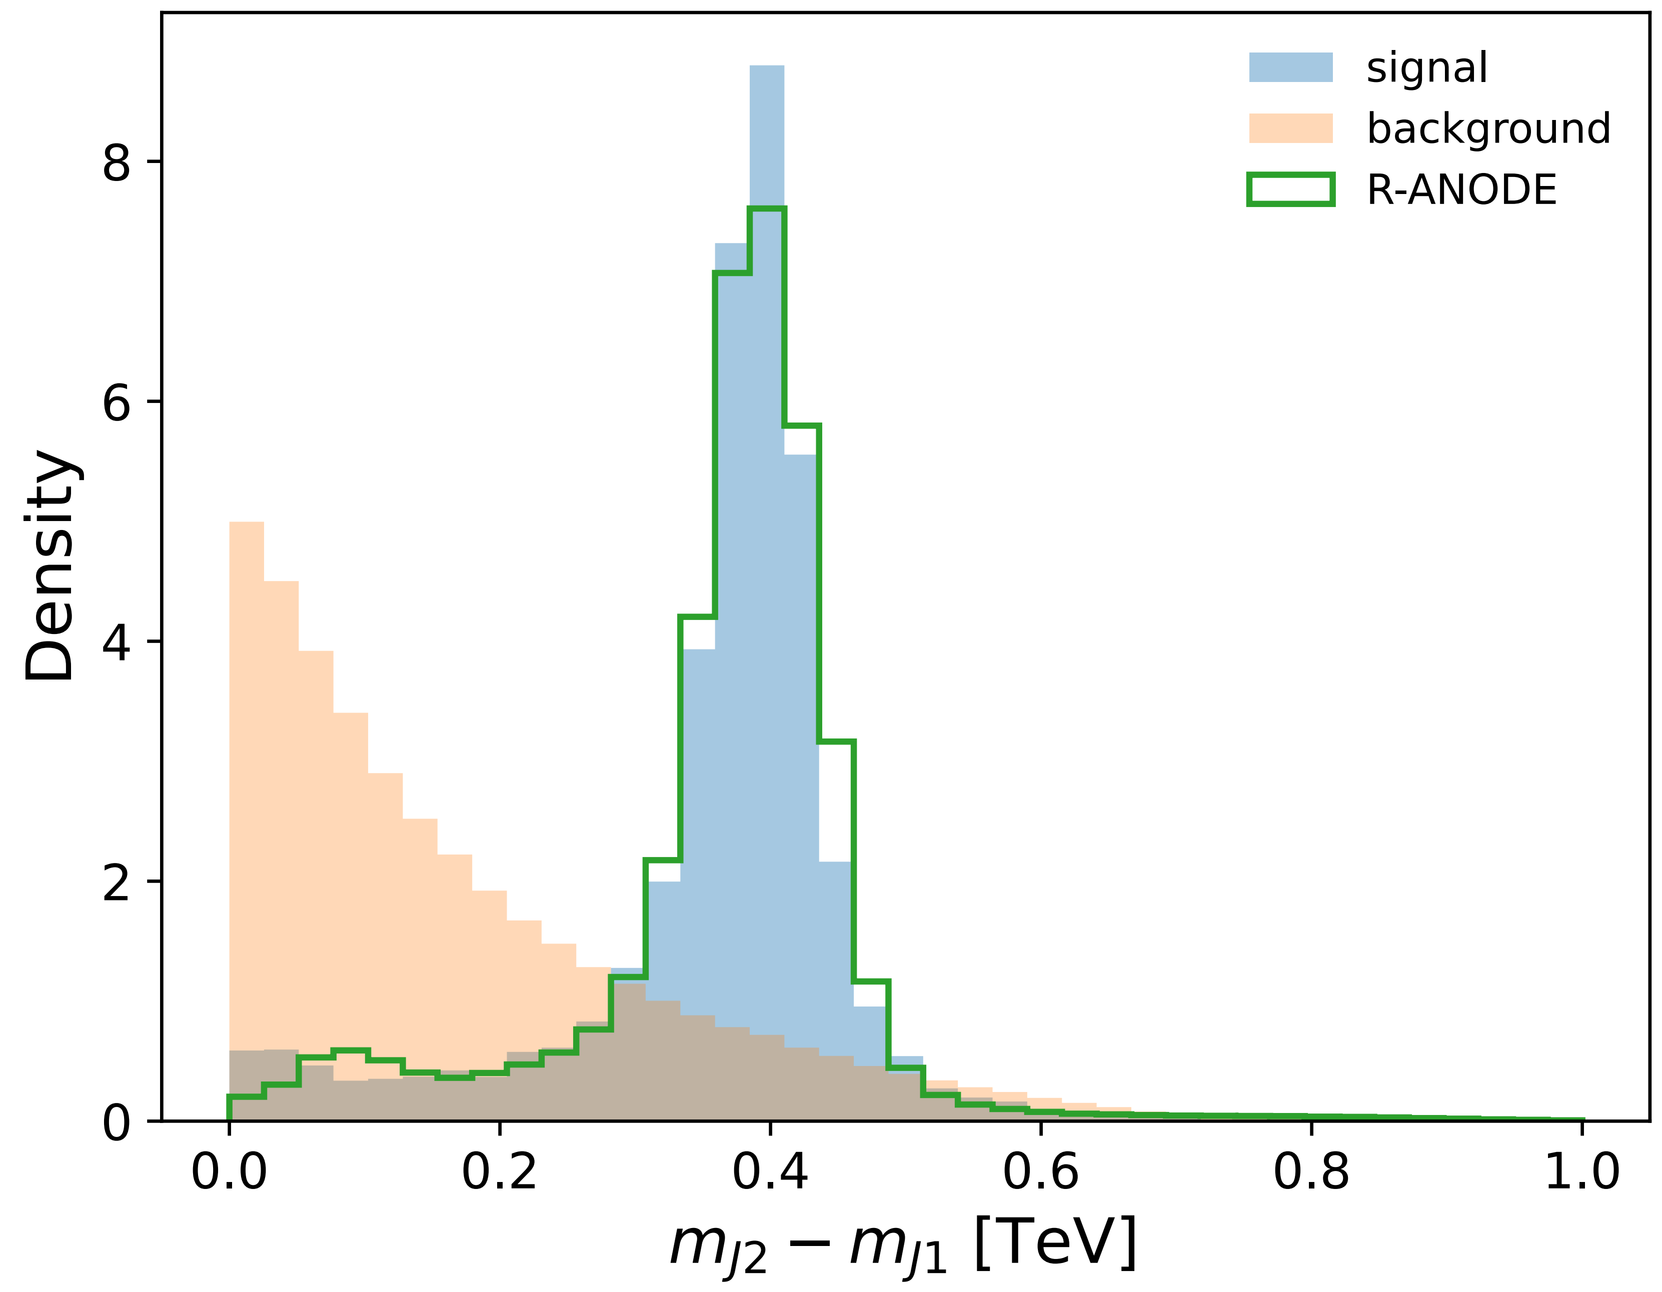

Das, Kasieczka & **DS**, 2311.nnnnn

3. Anomaly Detection from LHC to Astro

Searching for Stellar Streams in Gaia

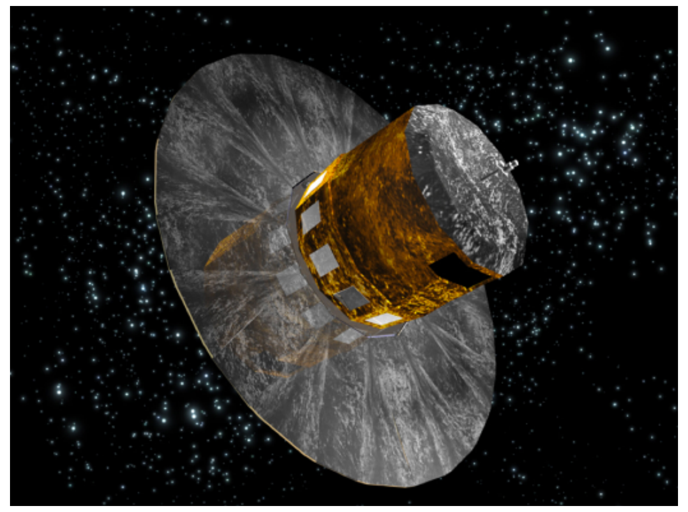

- We realized the same ML-enhanced bump hunt methods developed for LHC could be applied to **Gaia data** to search for **stellar streams**
    - *An example of power of ML to cut across domains!*

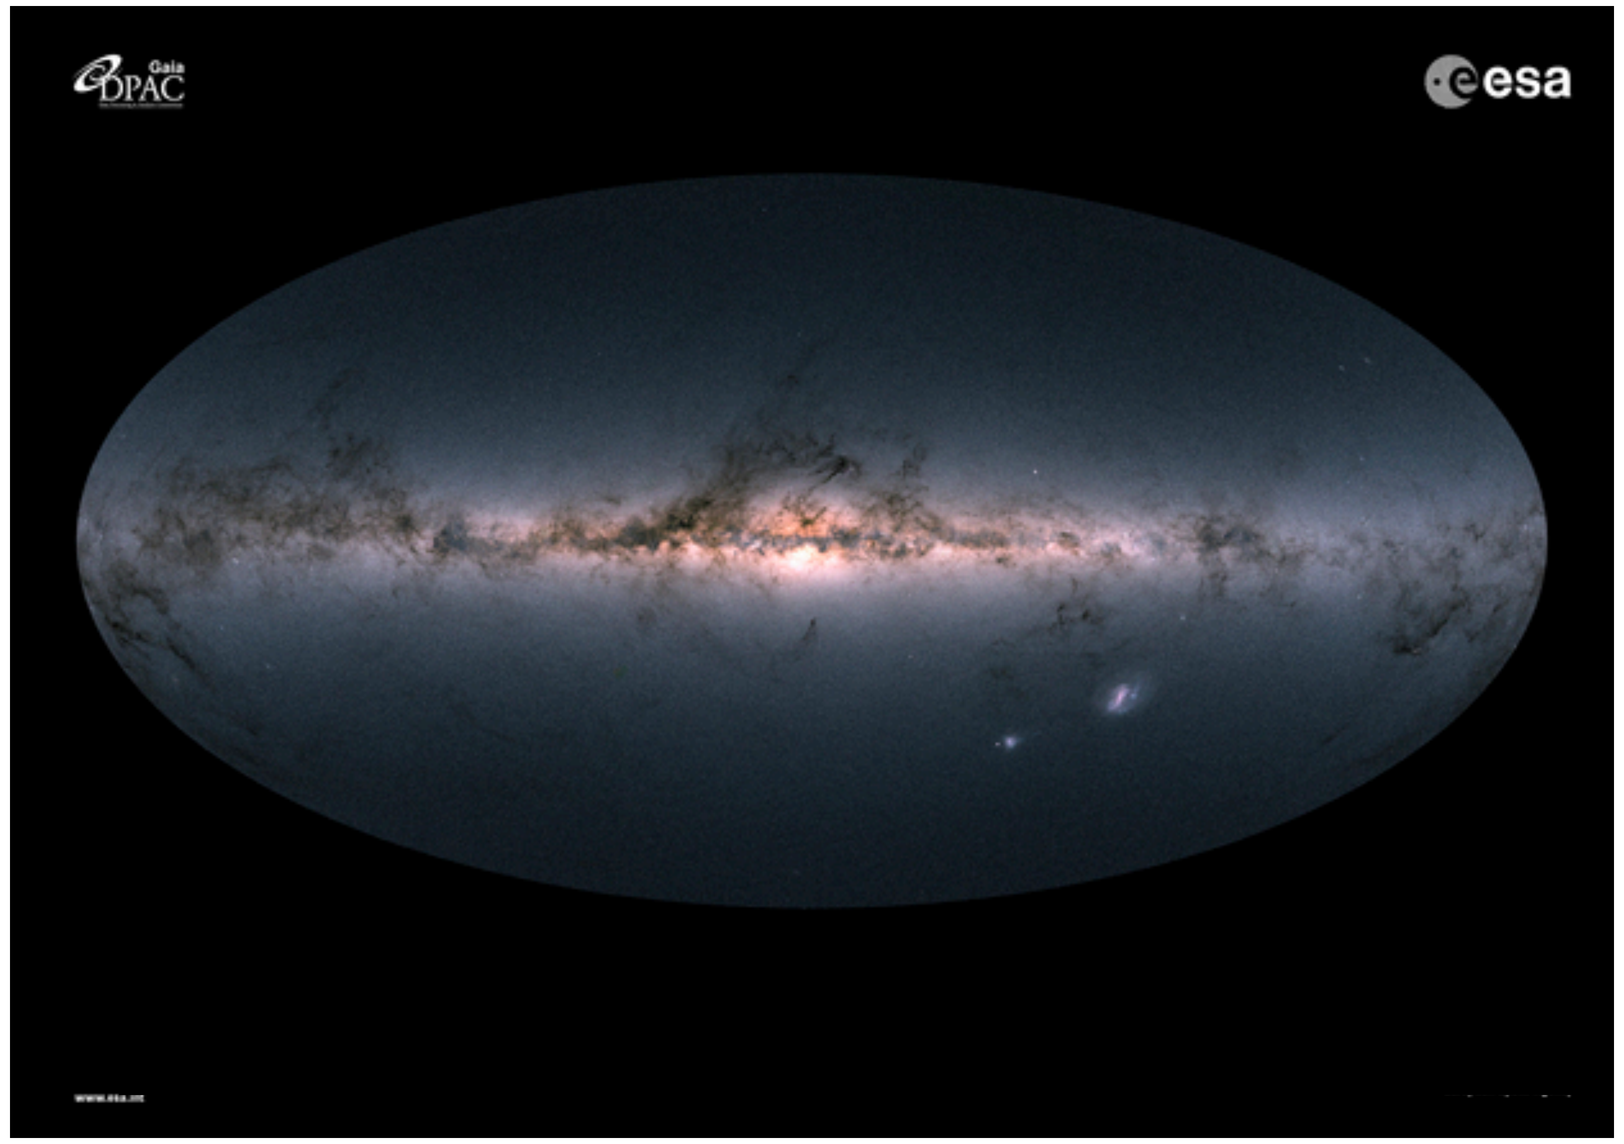

Gaia satellite:

- Launched in 2013; ongoing
- Angular positions, proper motions, color and magnitude of over **1 billion stars** in our Galaxy
- Distances and radial velocities for a smaller subset of nearby stars

Gaia data

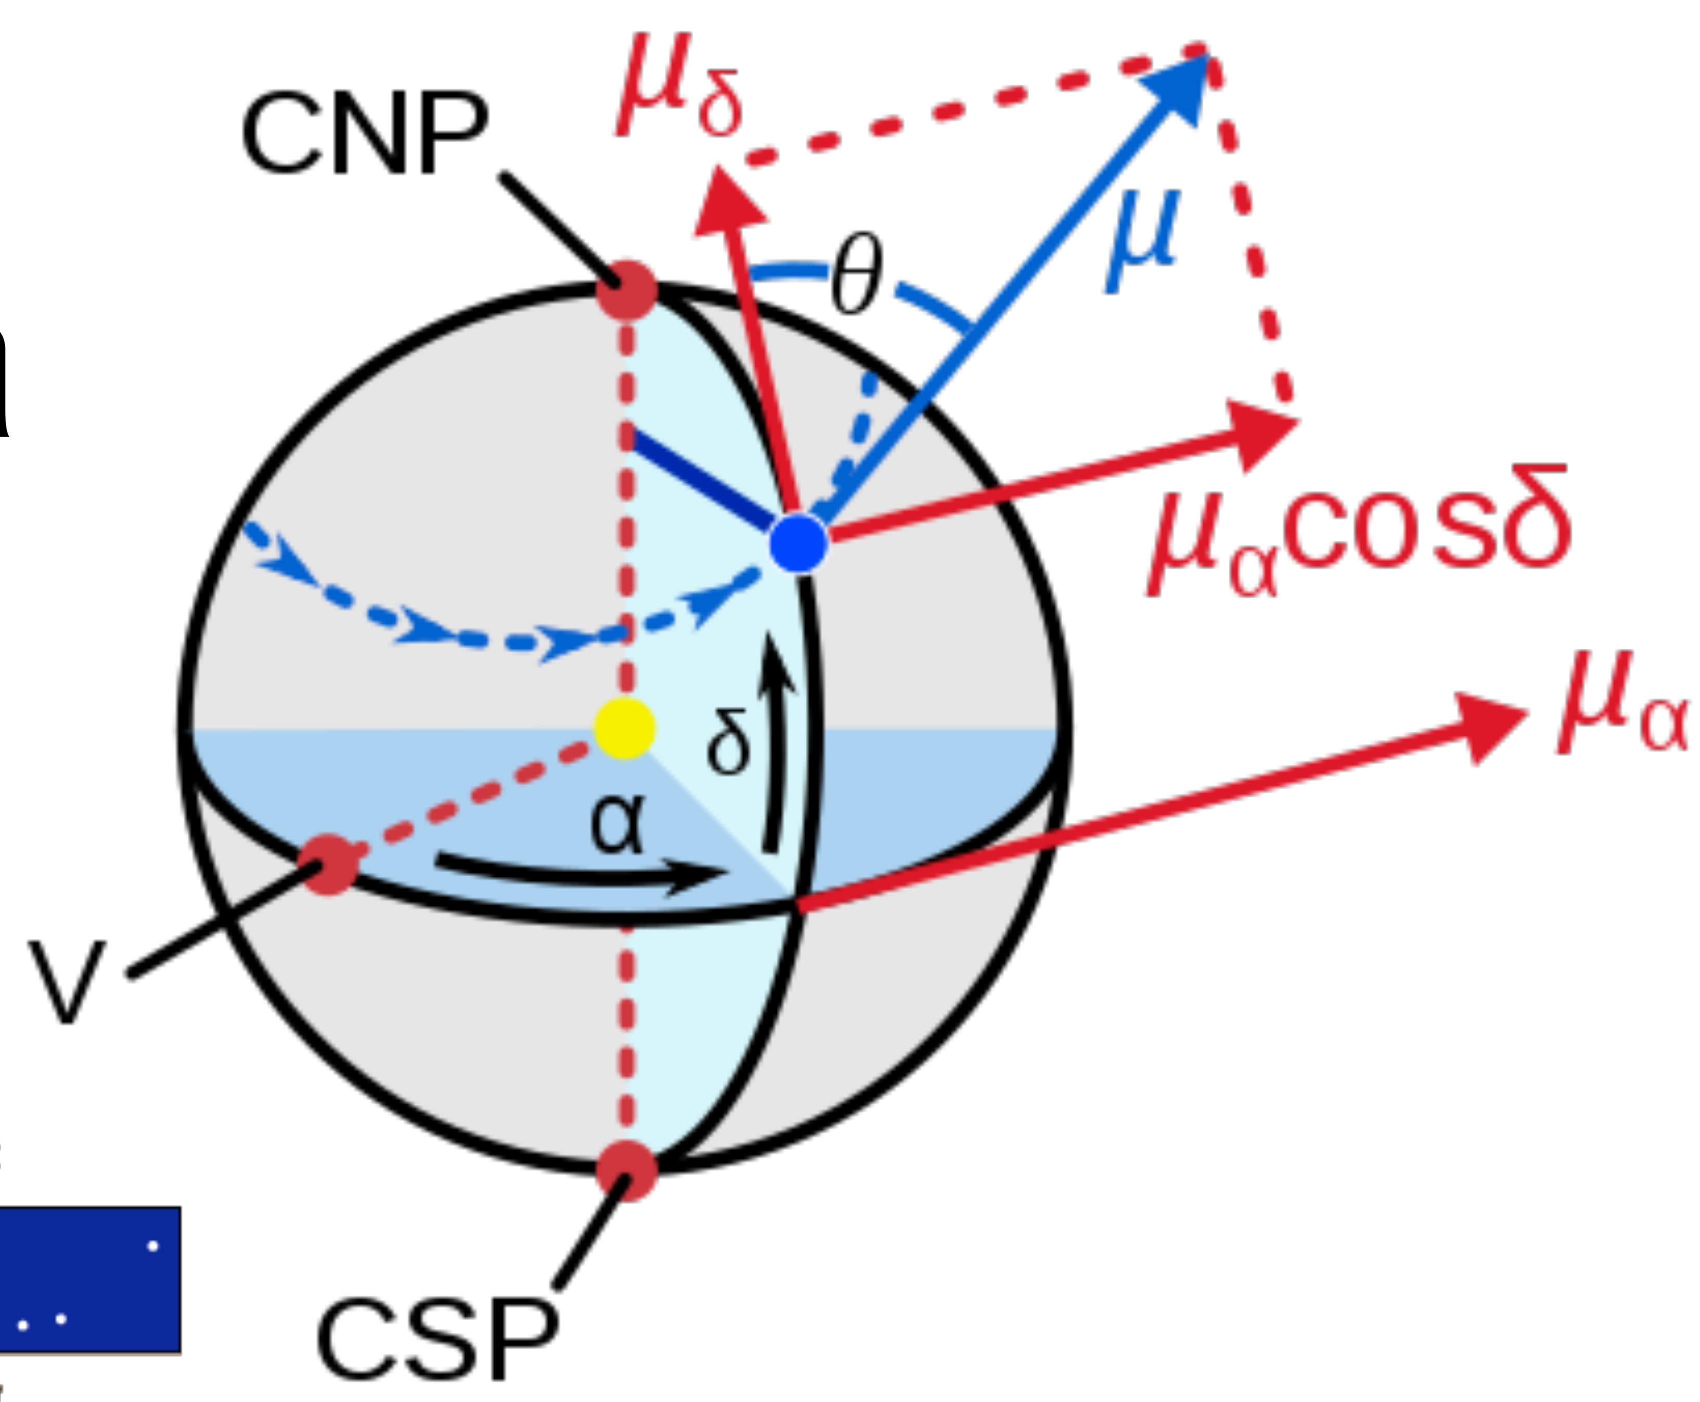

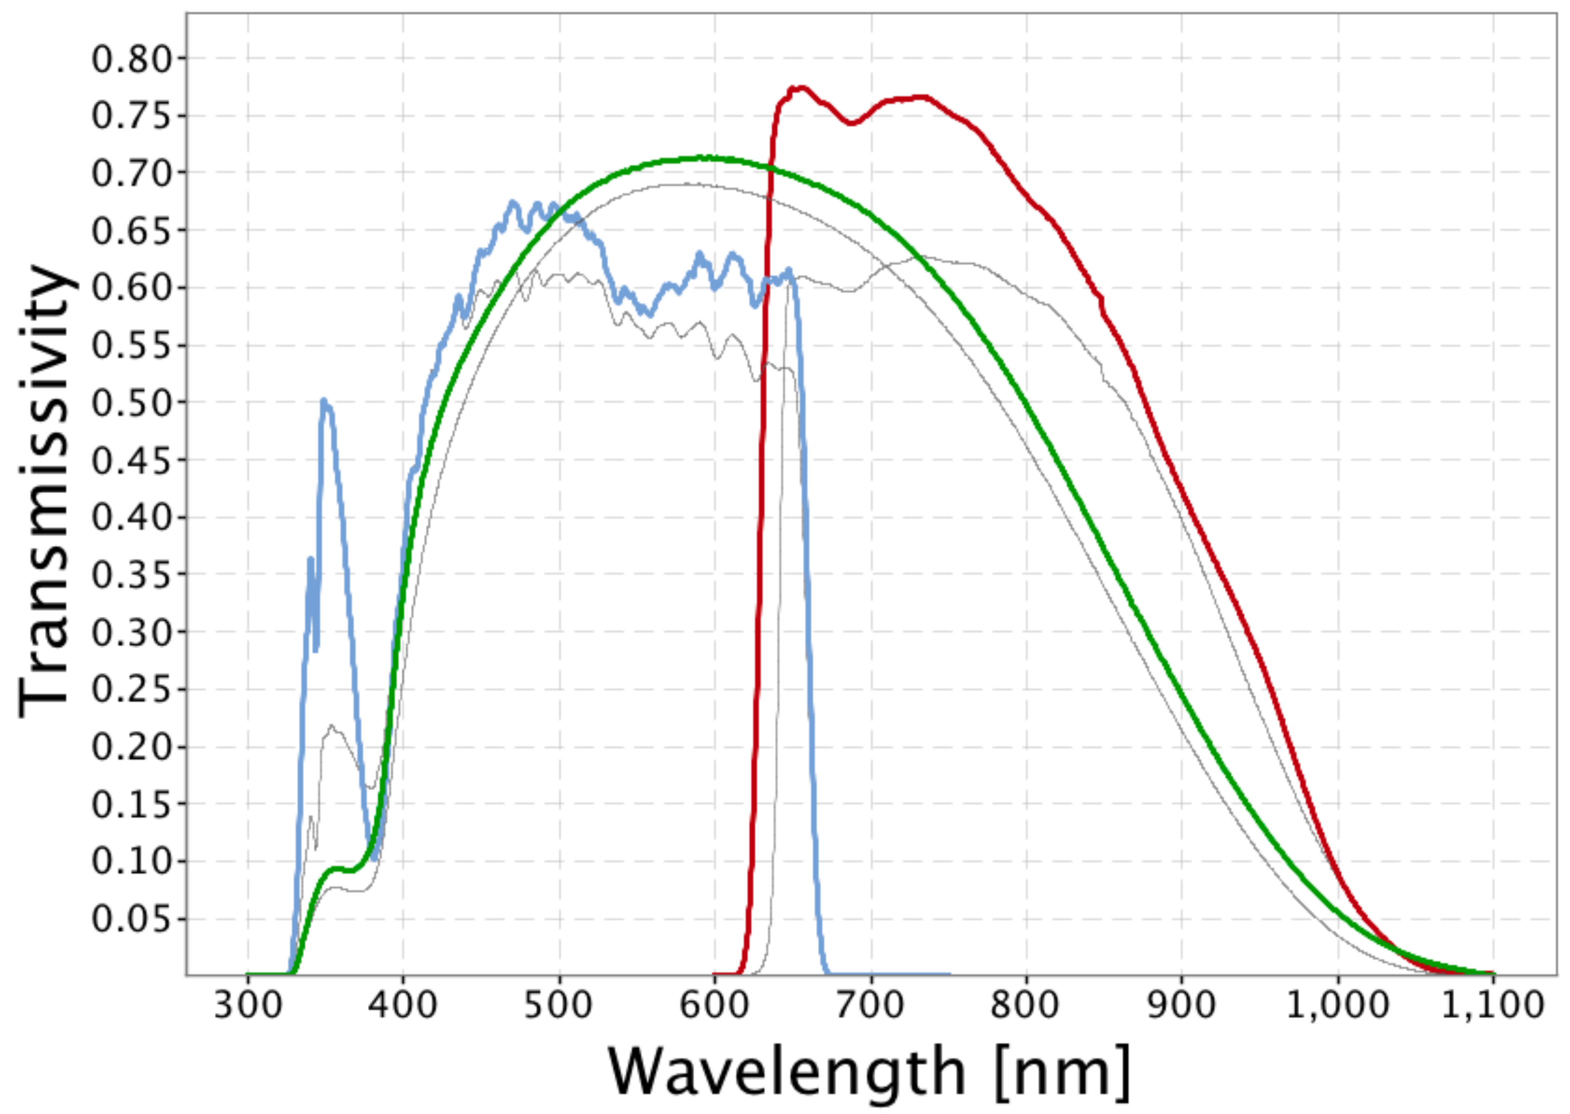

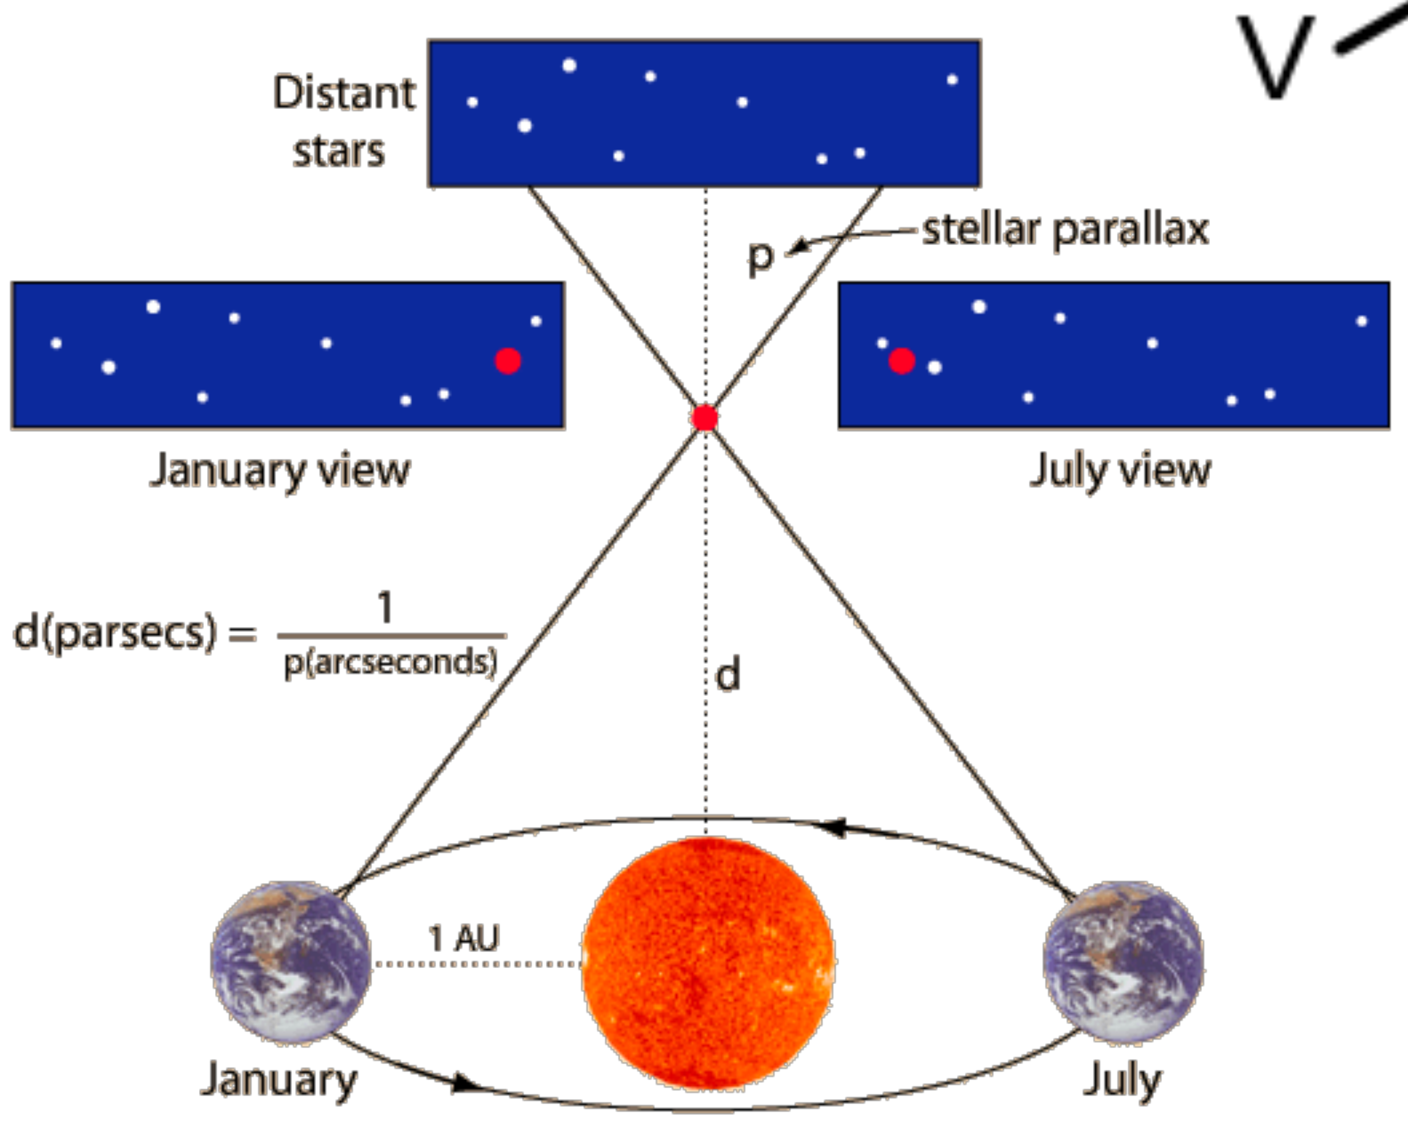

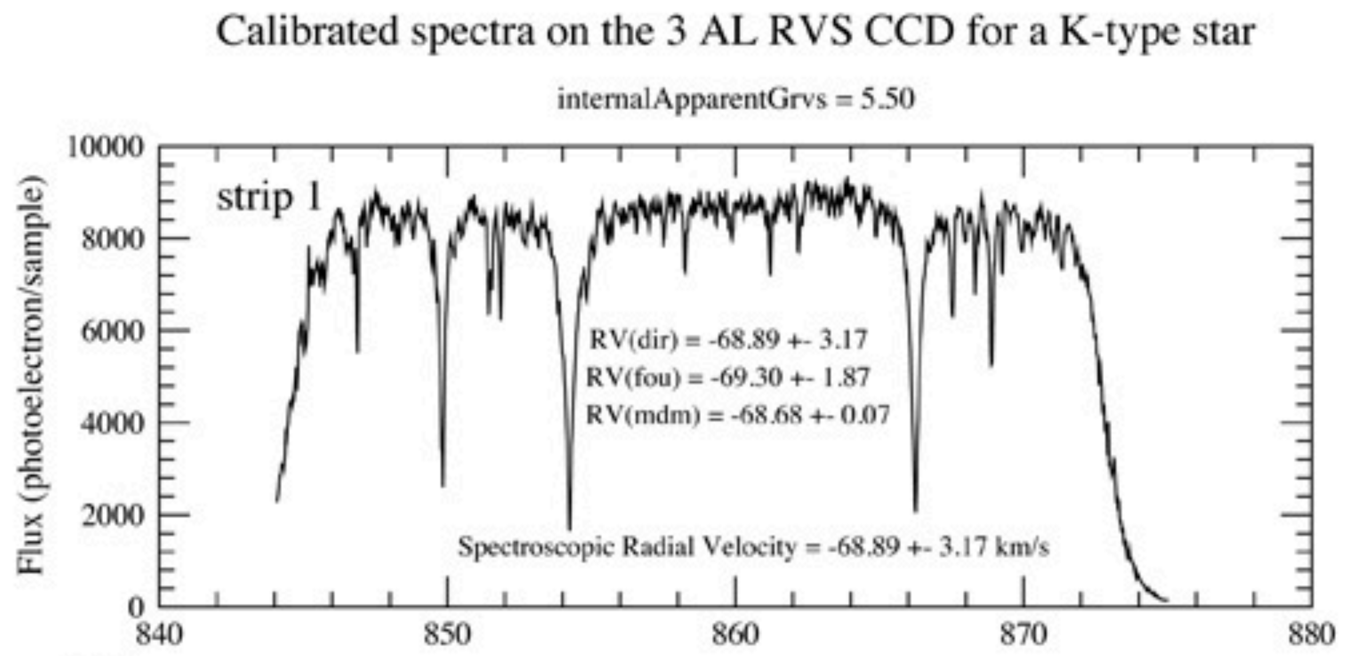

- Angular positions (α, δ) of stars on the sky (very precise)
- Through repeated observations:
  - *proper motions* (μ_α, μ_δ) (generally well-measured)
  - *parallax* (well-measured only for nearby stars)
- Magnitude in 3 passbands (colors)
- Radial velocities (using spectrometer)

Gaia data

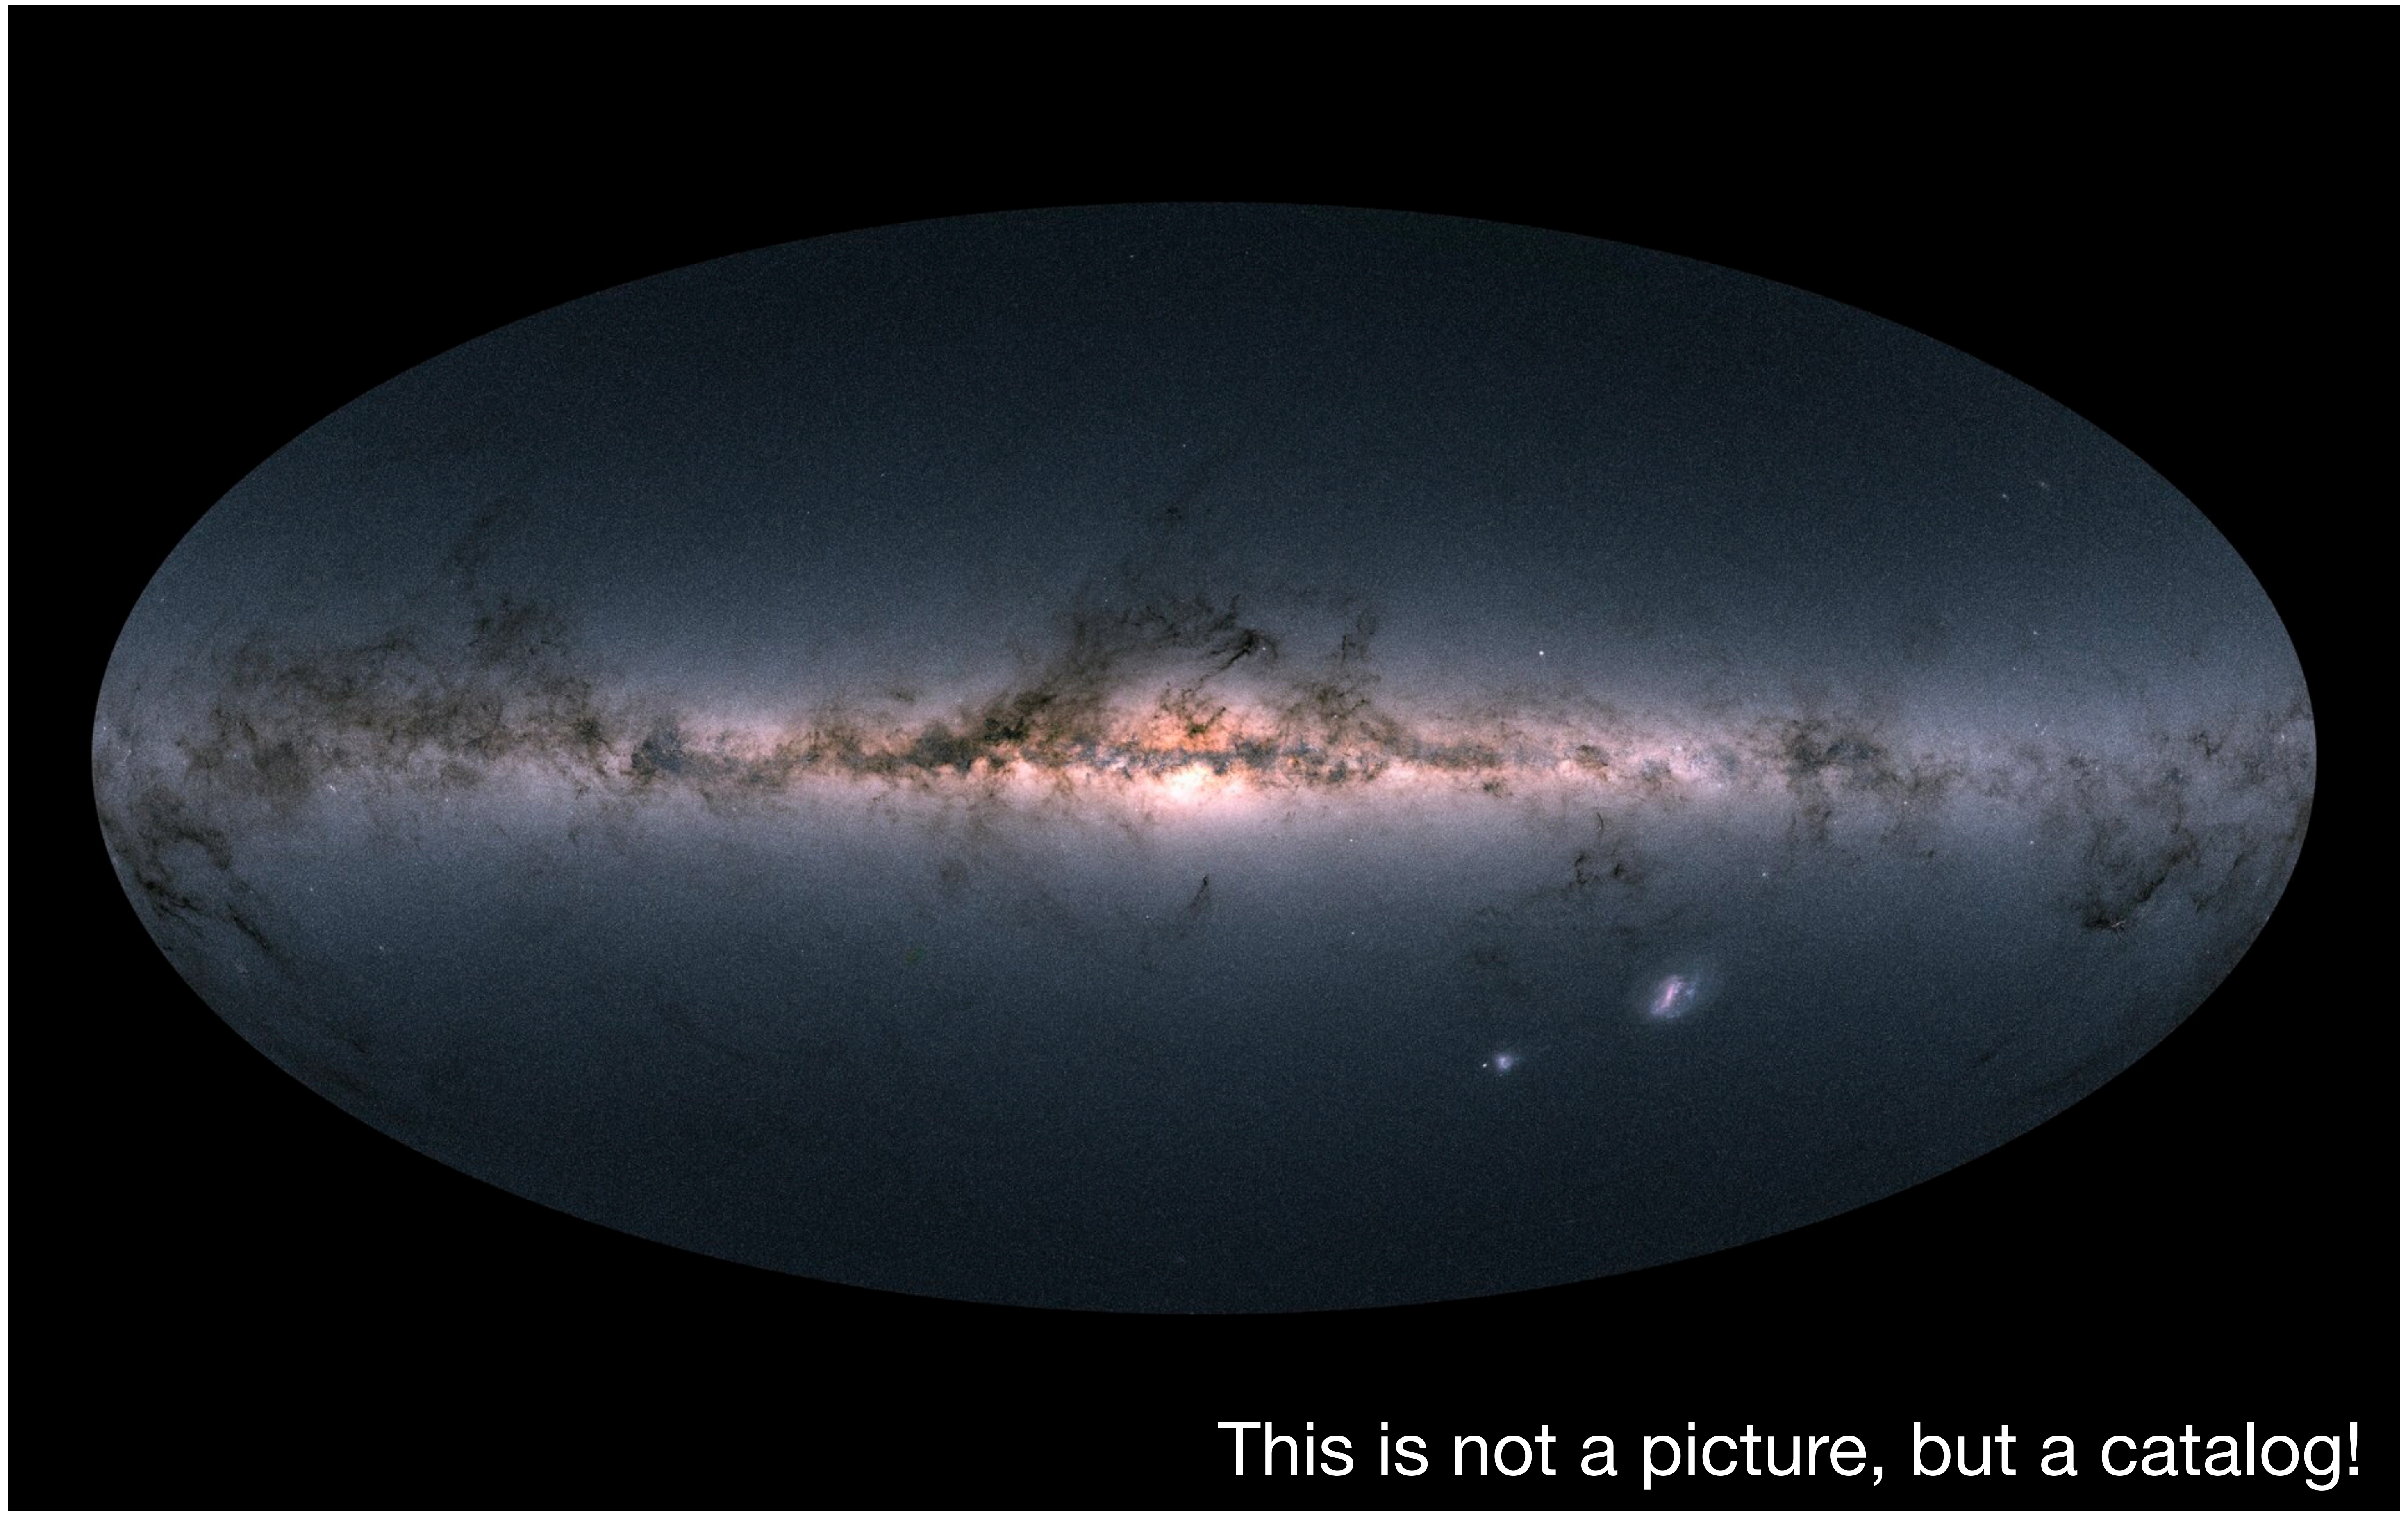
This is not a picture, but a catalog!

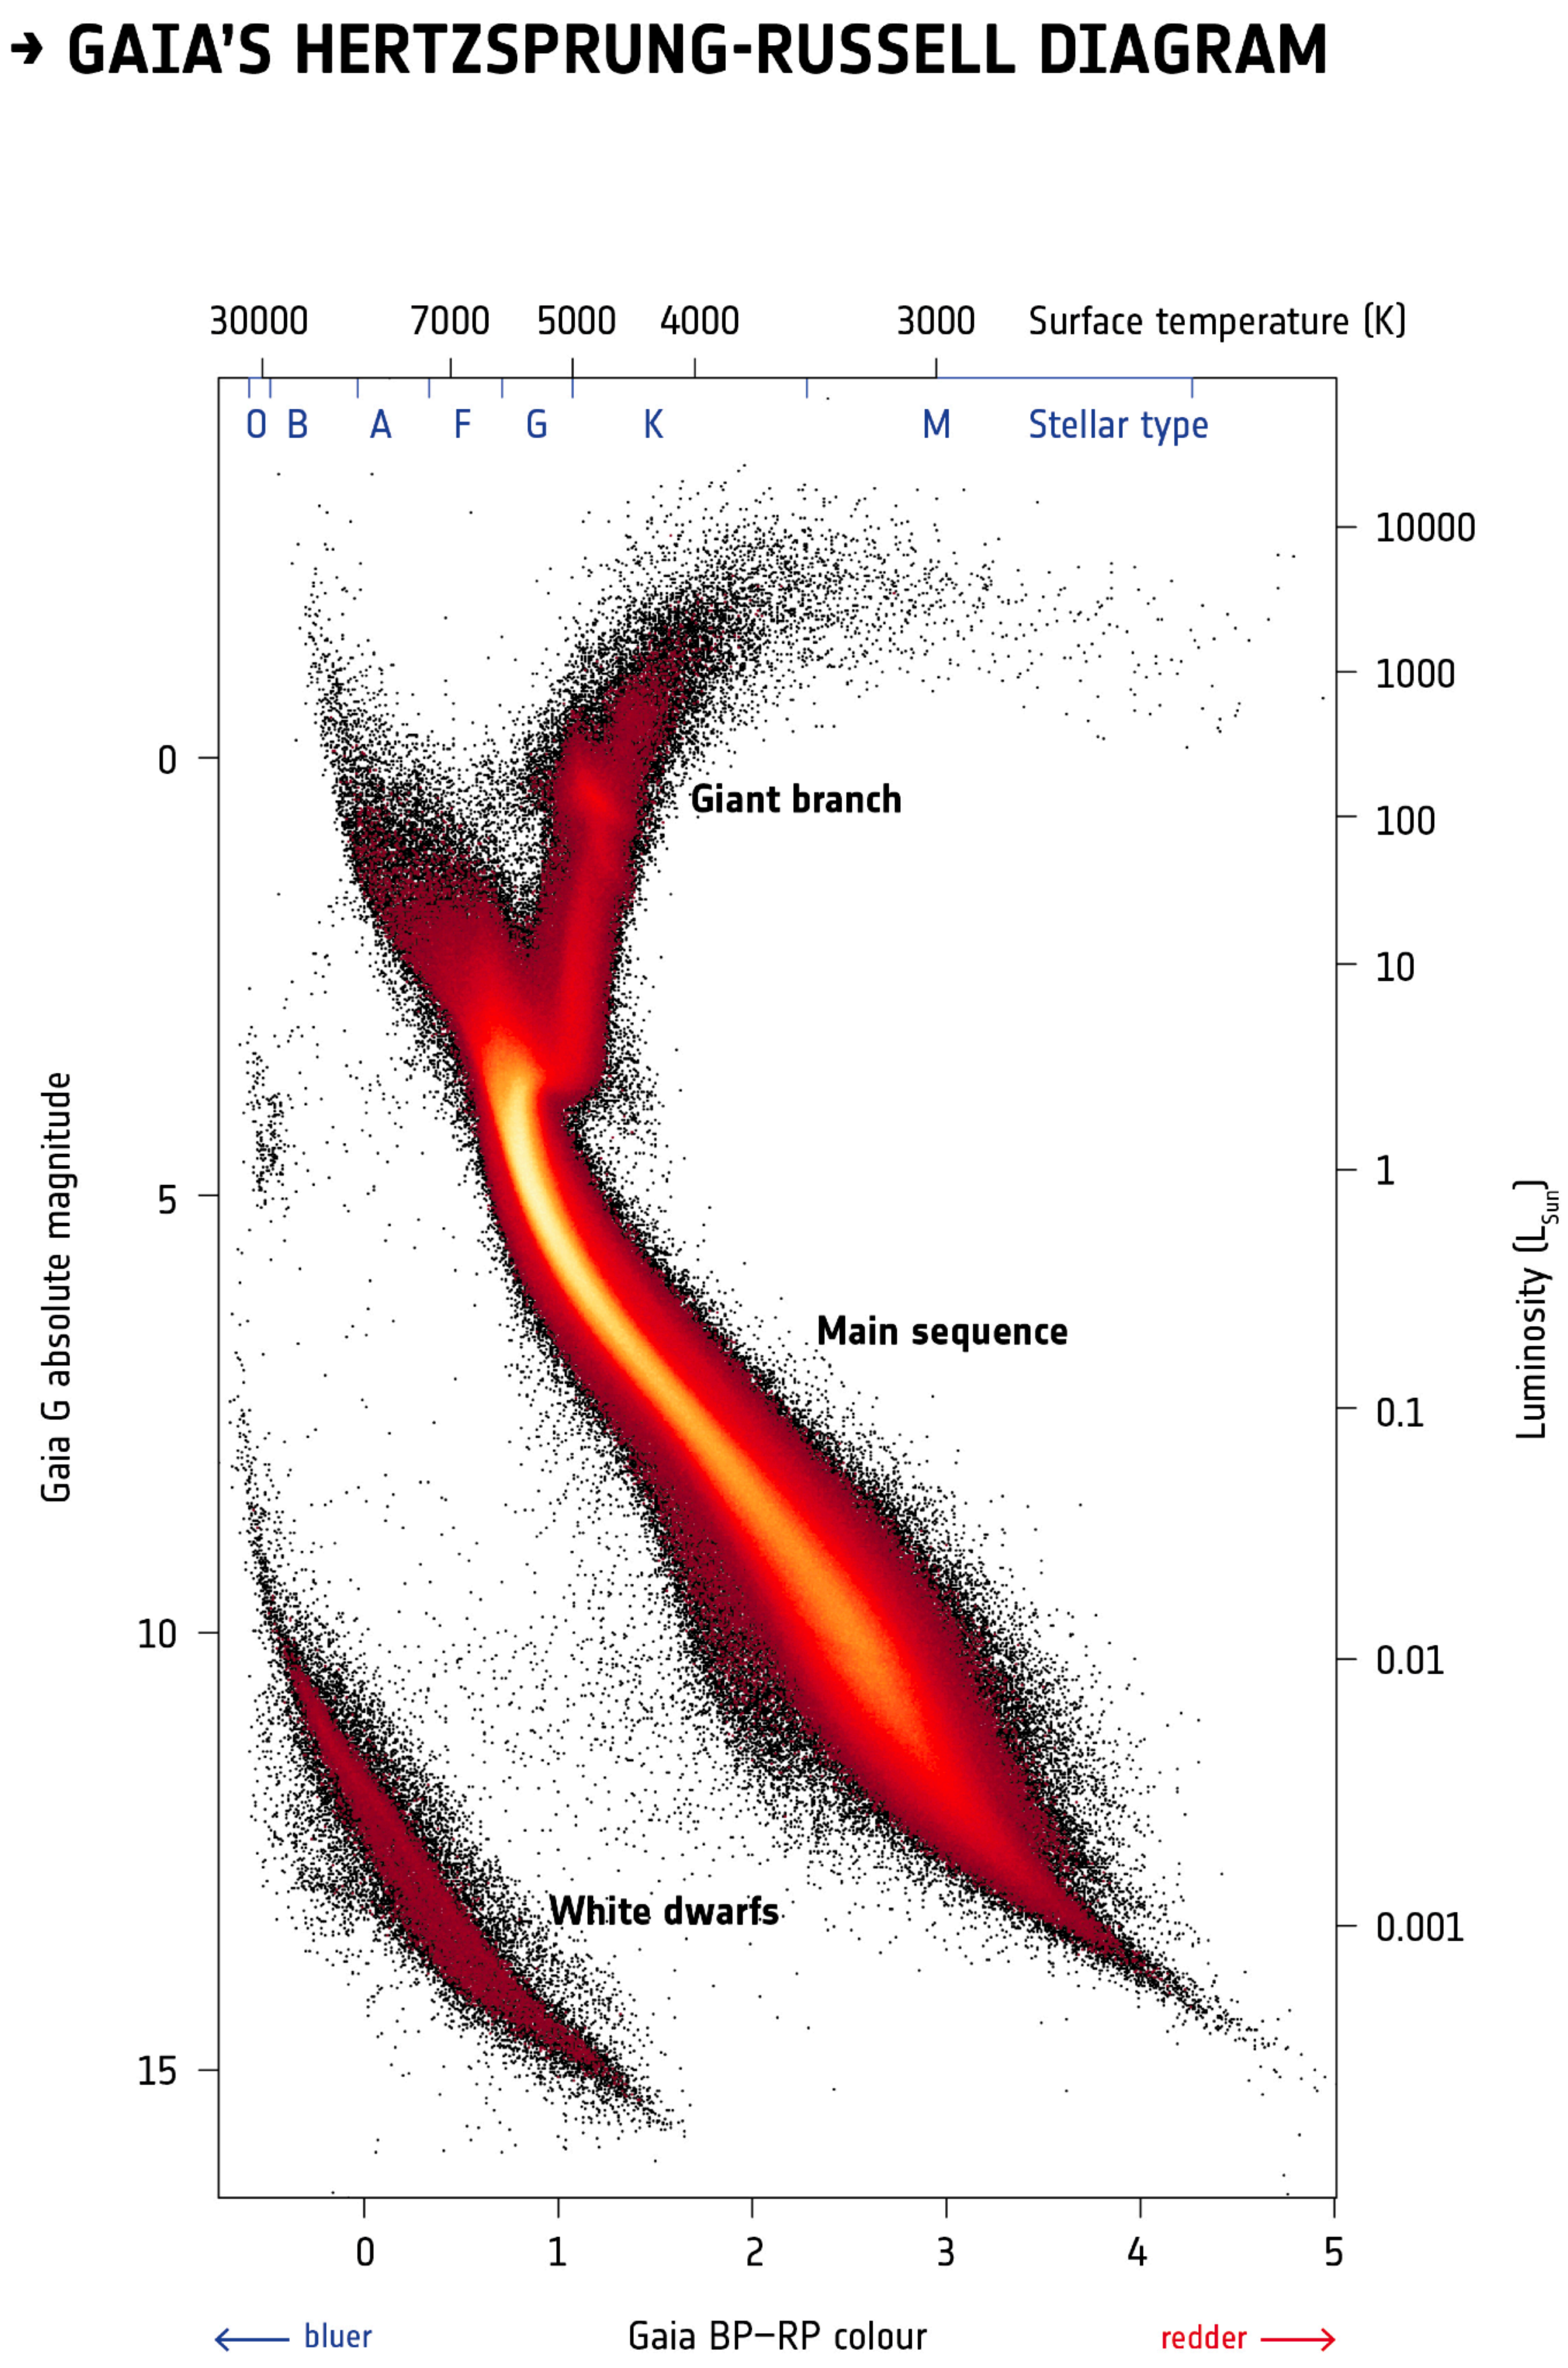
→ GAIA'S HERTZSPRUNG-RUSSELL DIAGRAM
30000
7000
5000
4000
3000
Surface temperature (K)
O B A F G K M Stellar type
Giant branch
Main sequence
White dwarfs
Gaia G absolute magnitude
0
5
10
15
Luminosity (L_{Sun})
10000
1000
100
10
1
0.1
0.01
0.001
0 1 2 3 4 5
← bluer
Gaia BP−RP colour
redder →

Gaia data

	# sources in Gaia DR3	# sources in Gaia DR2
Total number of sources	**1,811,709,771**	**1,692,919,135**
	Gaia Early Data Release 3	
Number of sources with full astrometry	1,467,744,818	1,331,909,727
Number of 5-parameter sources	585,416,709	
Number of 6-parameter sources	882,328,109	
Number of 2-parameter sources	343,964,953	361,009,408
Gaia-CRF sources	1,614,173	556,869
Sources with mean G magnitude	1,806,254,432	1,692,919,135
Sources with mean G_{BP}-band photometry	1,542,033,472	1,381,964,755
Sources with mean G_{RP}-band photometry	1,554,997,939	1,383,551,713
	New in Gaia Data Release 3	Gaia DR2
Sources with radial velocities	33,812,183	7,224,631
Sources with mean G_{RVS}-band magnitudes	32,232,187	-
Sources with rotational velocities	3,524,677	-
Mean BP/RP spectra	219,197,643	-
Mean RVS spectra	999,645	-
Variable-source analysis	10,509,536	550,737

Gaia data

	# sources in Gaia DR3	# sources in Gaia DR2
Total number of sources	**1,811,709,771**	**1,692,919,135**
	Gaia Early Data Release 3	
Number of sources with full astrometry	1,467,744,818	1,331,909,727
Number of 5-para	585,416,709	
Number of 6-para	882,328,109	
Number of 2-parameter sources	343,964,953	361,009,408
Gaia-CRF sources	1,614,173	556,869
Sources with mean G magnitude	1,806,254,432	1,692,919,135
Sources with mean G_{BP}-band photometry	1,542,033,472	1,381,964,755
Sources with mean G_{RP}-band photometry	1,554,997,939	1,383,551,713
	New in Gaia Data Release 3	Gaia DR2
Sources with radial velocities	33,812,183	7,224,631
Sources with mean G_{RVS}-band magnitudes	32,232,187	-
Sources with rotational velocities	3,524,677	-
Mean BP/RP spectra	219,197,643	-
Mean RVS spectra	999,645	-
Variable-source analysis	10,509,536	550,737

1.5B stars with 5d information (angular positions, proper motions, parallax)

Gaia data

	# sources in Gaia DR3	# sources in Gaia DR2
Total number of sources	**1,811,709,771**	**1,692,919,135**
	Gaia Early Data Release 3	
Number of sources with full astrometry	1,467,744,818	1,331,909,727
Number of 5-para	585,416,709	
Number of 6-para	882,328,109	
Number of 2-parameter sources	343,964,953	361,009,408
Gaia-CRF sources	1,614,173	556,869
Sources with mean G	1,806,254,432	1,692,919,135
Sources with mean G_{BP}-band photometry	1,542,033,472	1,381,964,755
Sources with mean G_{RP}-band photometry	1,554,997,939	1,383,551,713
	New in Gaia Data Release 3	Gaia DR2
Sources with radial velocities	33,812,183	7,224,631
Sources with mean G_{RVS}-band magnitudes	32,232,187	-
Sources with rotational velocities	3,524,677	-
Mean BP/RP spectra	219,197,643	-
Mean RVS spectra	999,645	-
Variable-source analysis	10,509,536	550,737

1.5B stars with 5d information (angular positions, proper motions, parallax)

(much smaller subset of nearby stars with “well-measured” parallaxes)

Gaia data

	# sources in Gaia DR3	# sources in Gaia DR2
Total number of sources	**1,811,709,771**	**1,692,919,135**
	Gaia Early Data Release 3	
Number of sources with full astrometry	1,467,744,818	1,331,909,727
Number of 5-para	585,416,709	
Number of 6-para	882,328,109	
Number of 2-parameter sources	343,964,953	361,009,408
Gaia-CRF sources	1,614,173	556,869
Sources with mean G	1,806,254,432	1,692,919,135
Sources with mean G_{BP}-band photometry	1,542,033,472	1,381,964,755
Sources with mean G_{RP}-band photometry	1,554,997,939	1,383,551,713
	New in Gaia Data Release 3	Gaia DR2
Sources with radial velocities	33,812,183	7,224,631
Sources with mean G_{RVS}-band magnitudes	32,232,187	-
Sources with rotationa	3,524,677	-
Mean BP/RP spectra	219,197,643	-
Mean RVS spectra	999,645	-
Variable-source analysis	10,509,536	550,737

1.5B stars with 5d information (angular positions, proper motions, parallax)

(much smaller subset of nearby stars with “well-measured” parallaxes)

“only” 33M stars with full 6d information (5d+radial velocities)

Stellar Streams

Stellar streams are the very old remnants of tidally disrupted globular clusters and dwarf galaxies.

Collection of stars moving together along a common orbit — concentrated spatially and in velocity.

Unique probes into the formation history and gravitational potential of the Galaxy, and into dark matter substructure.

Stellar Streams

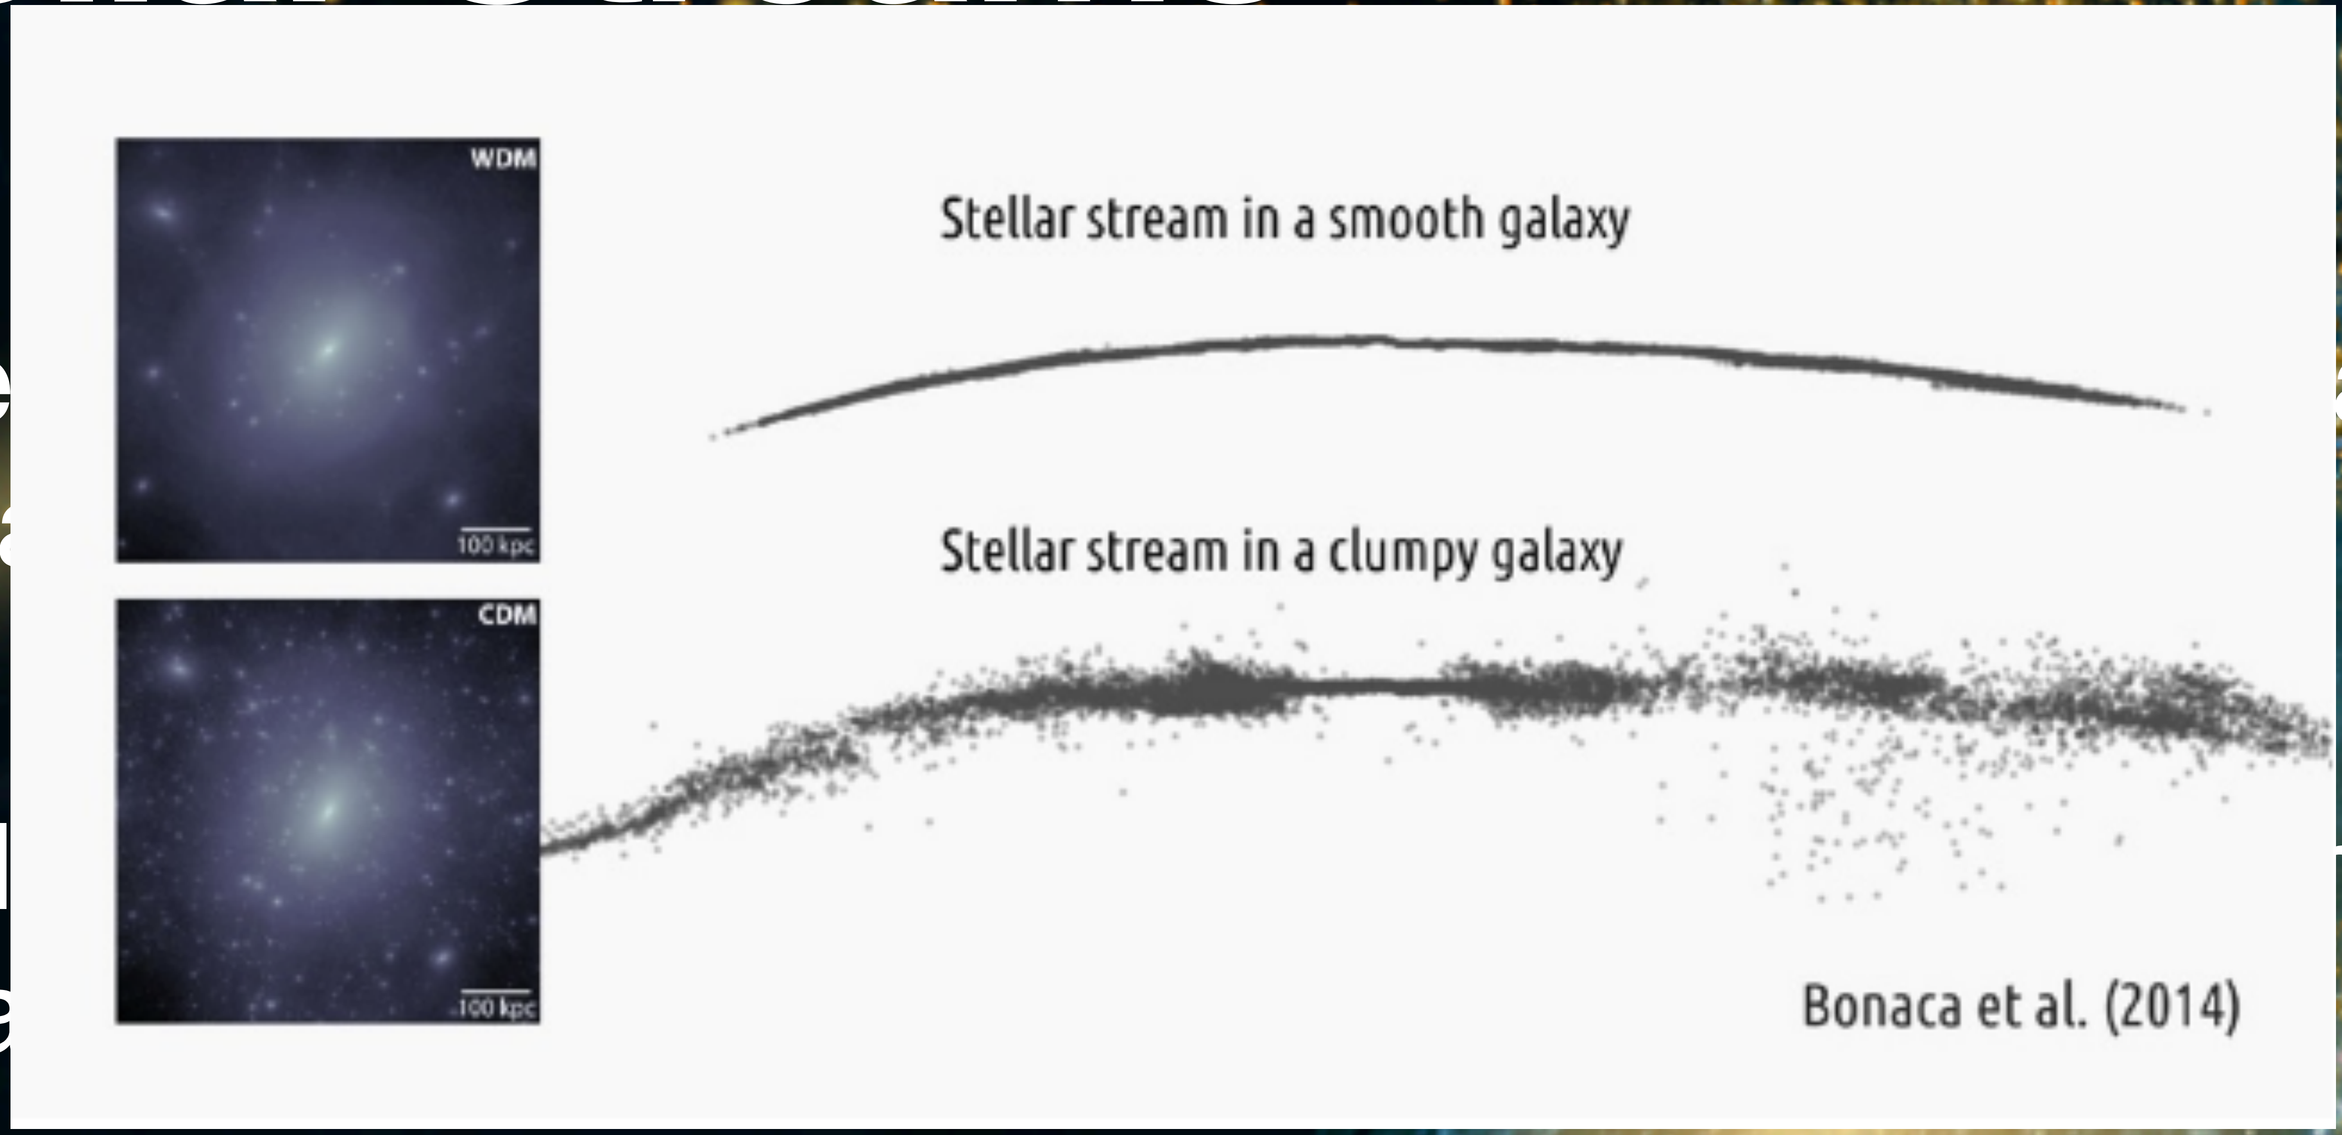

Ste ally disrupted globular clusters and dw

Col mmon orbit — concentrated spa

Unique probes into the formation history and gravitational potential of the Galaxy, and into dark matter substructure.

Stellar Streams

Ste… ally disrupted globular clusters and dwa…

Col… mmon orbit — concentrated spa…

Unique probes into the formati… the Galaxy, and into dark matte…

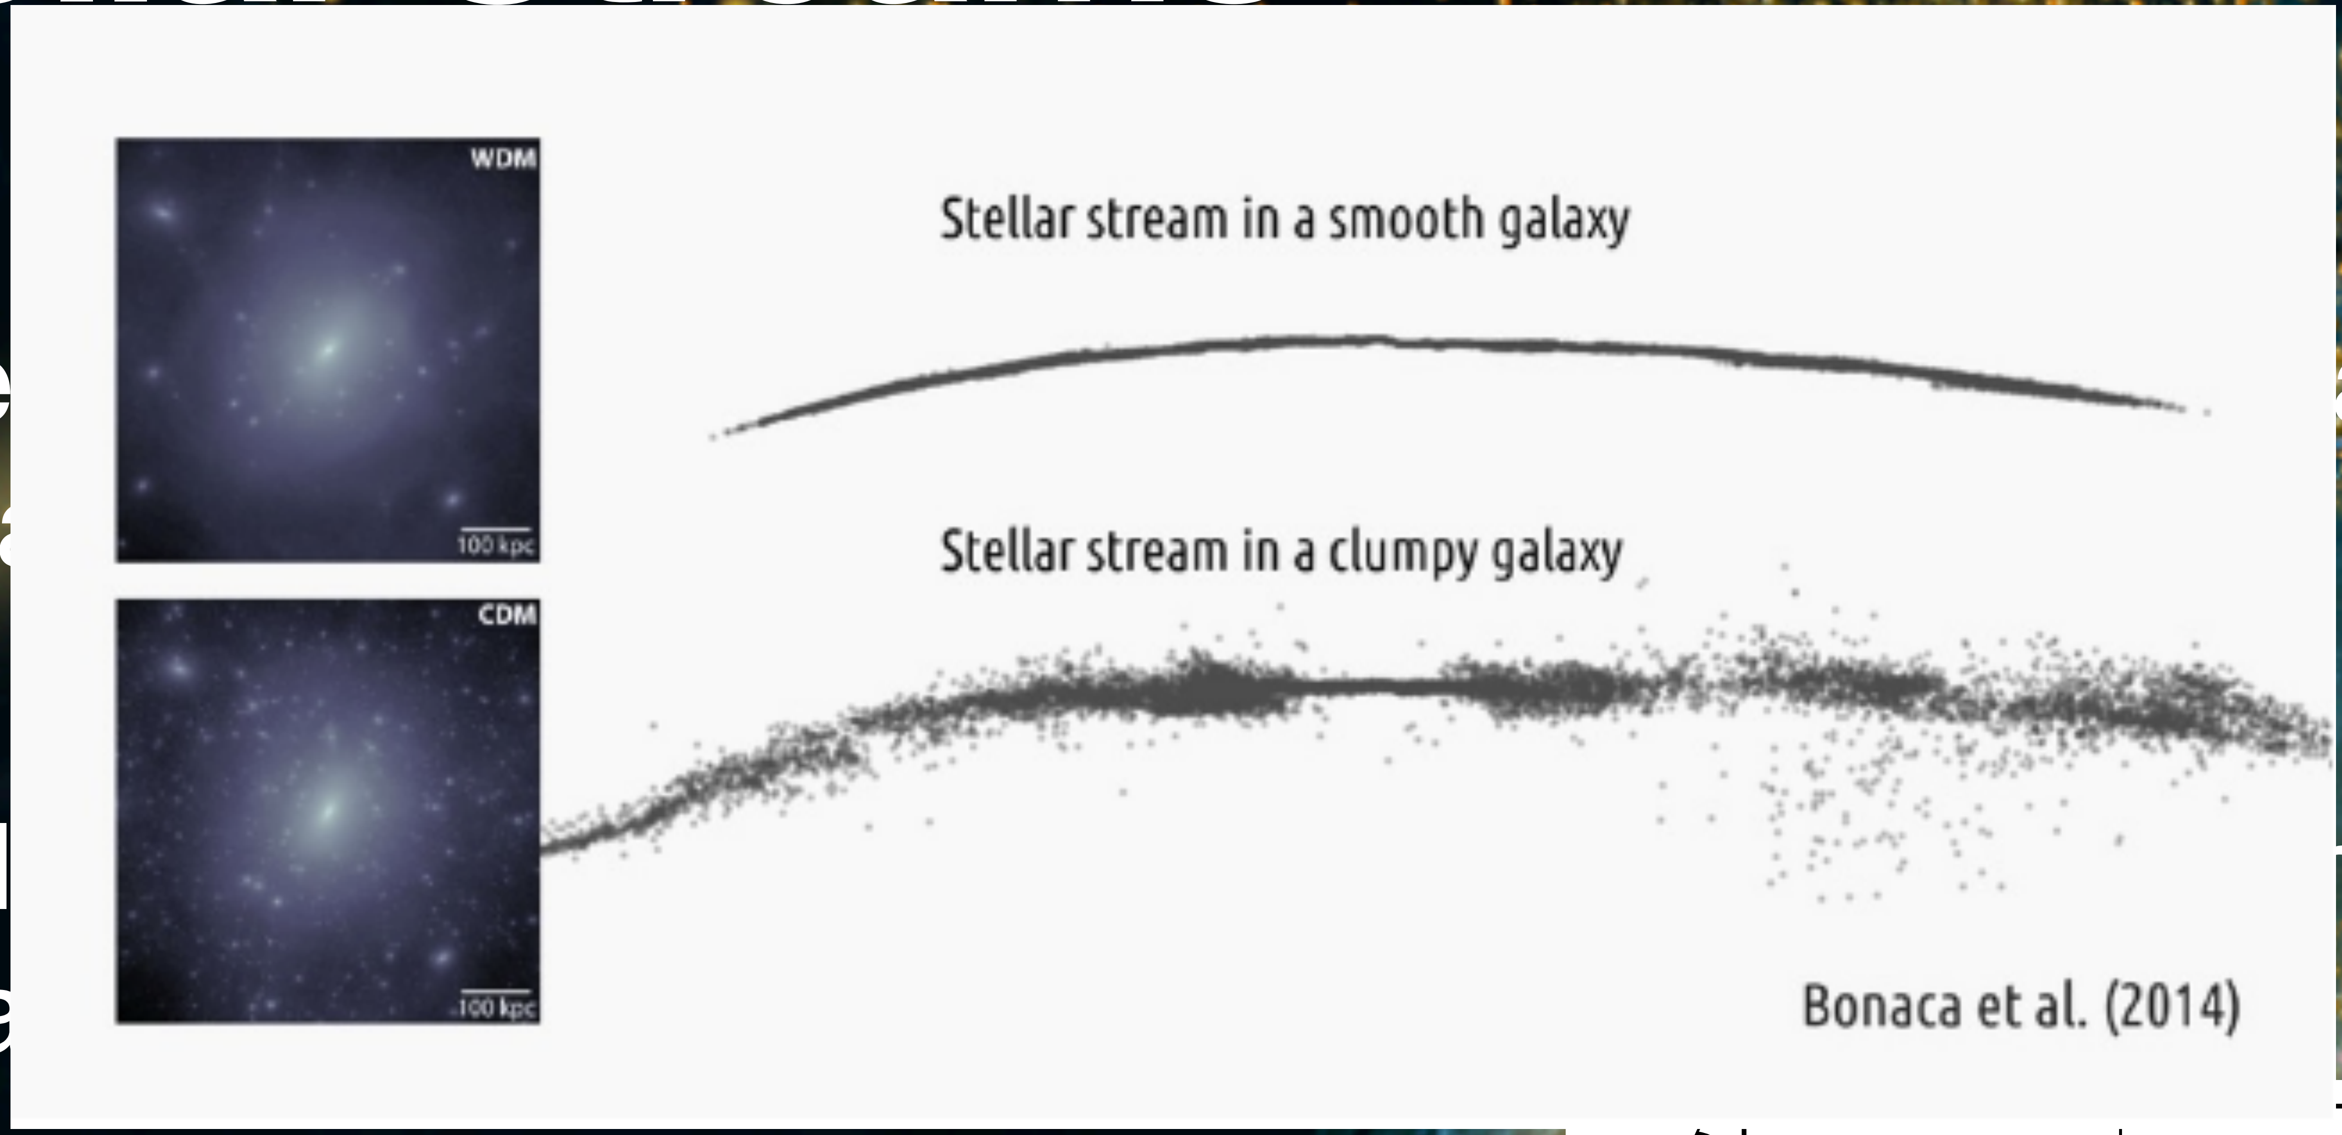

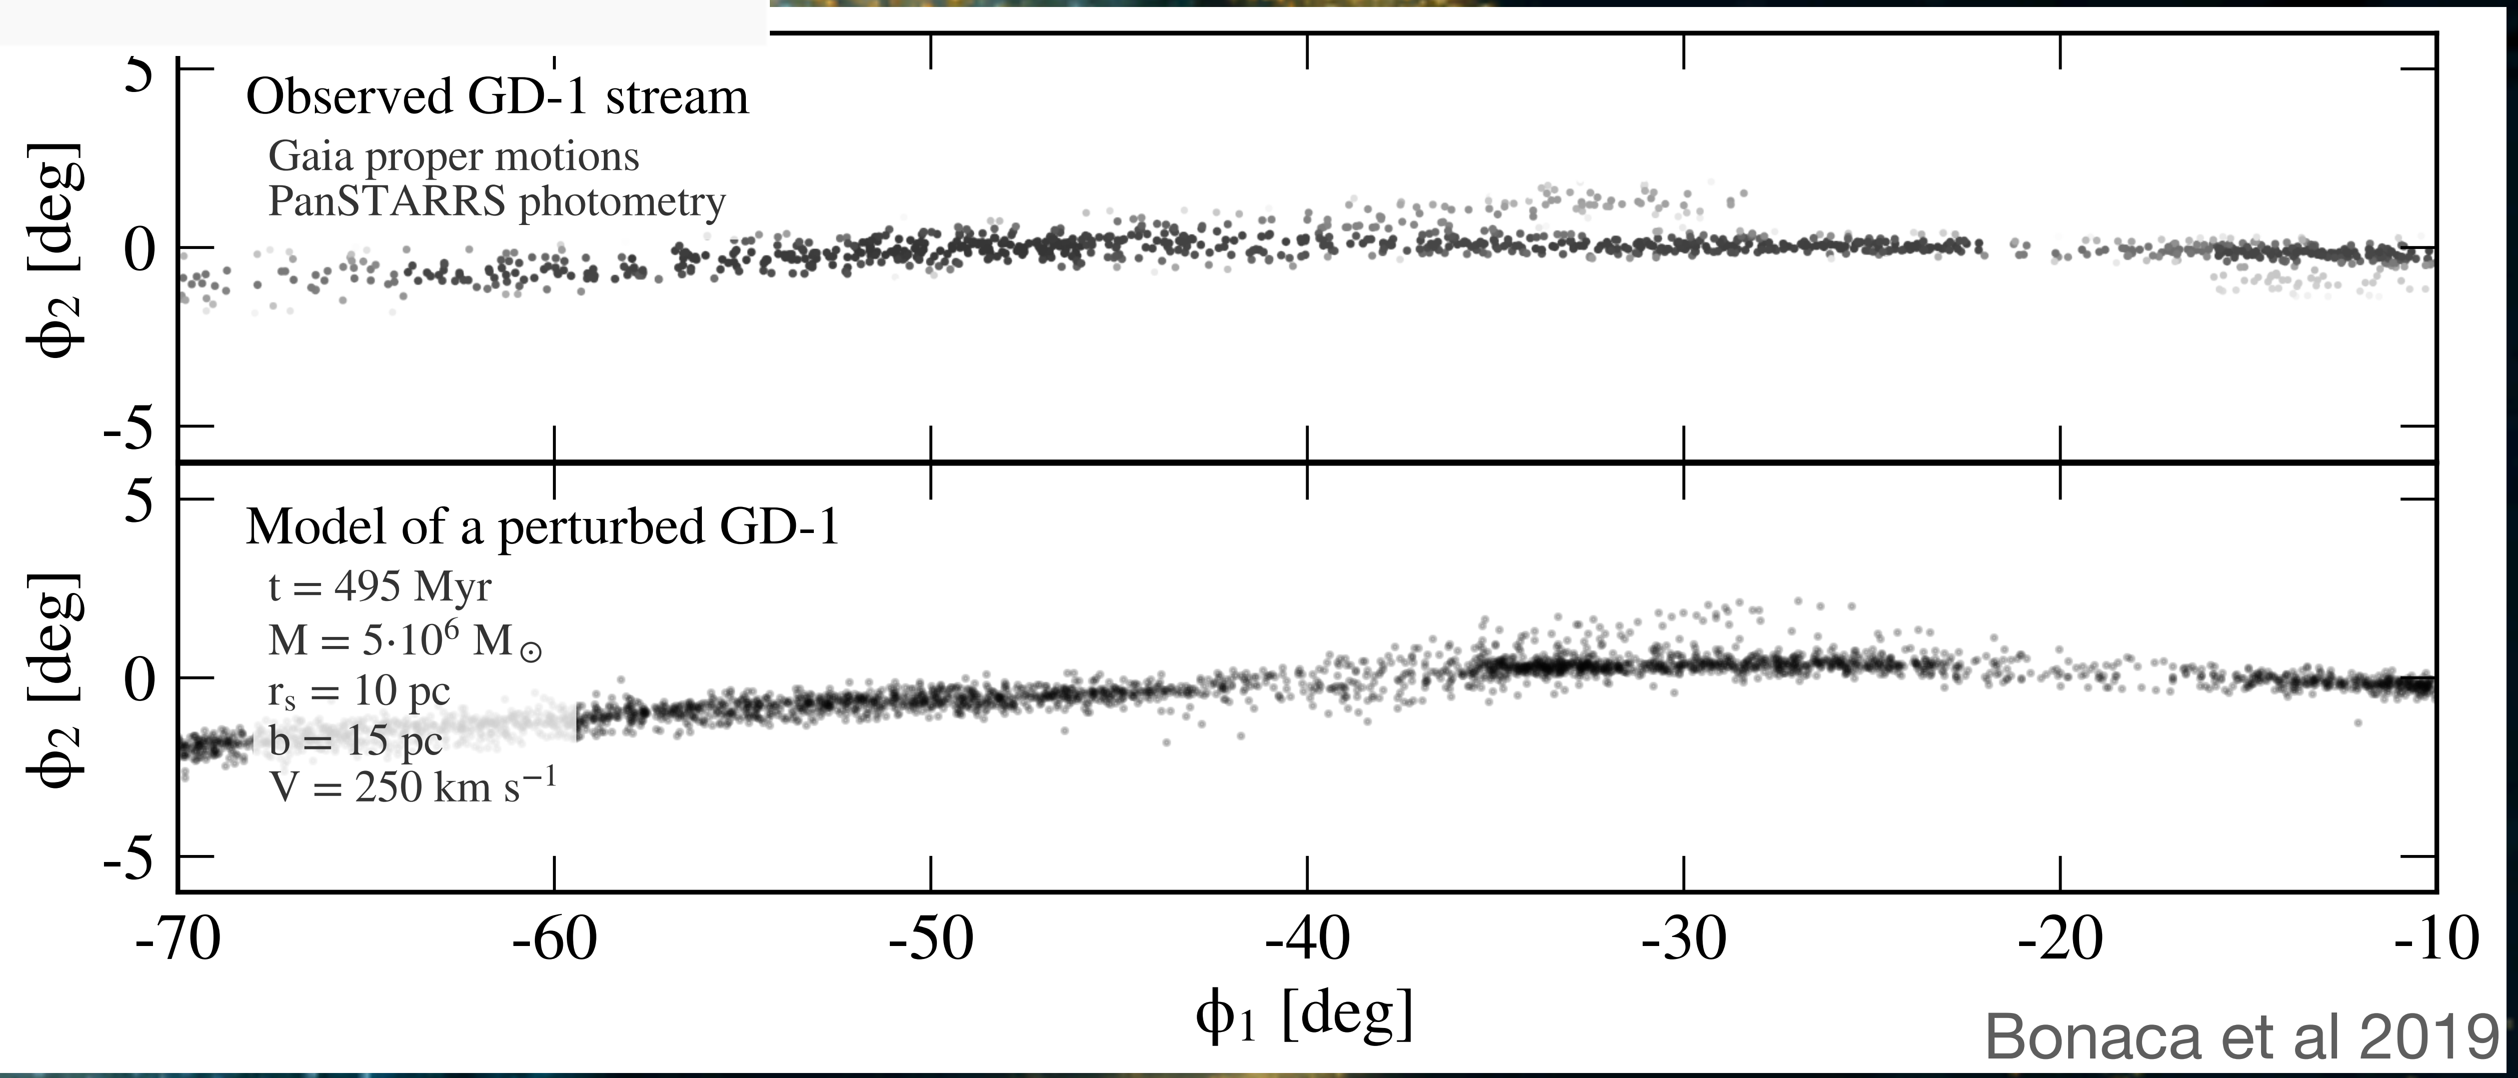

Known Stellar Streams of the Milky Way

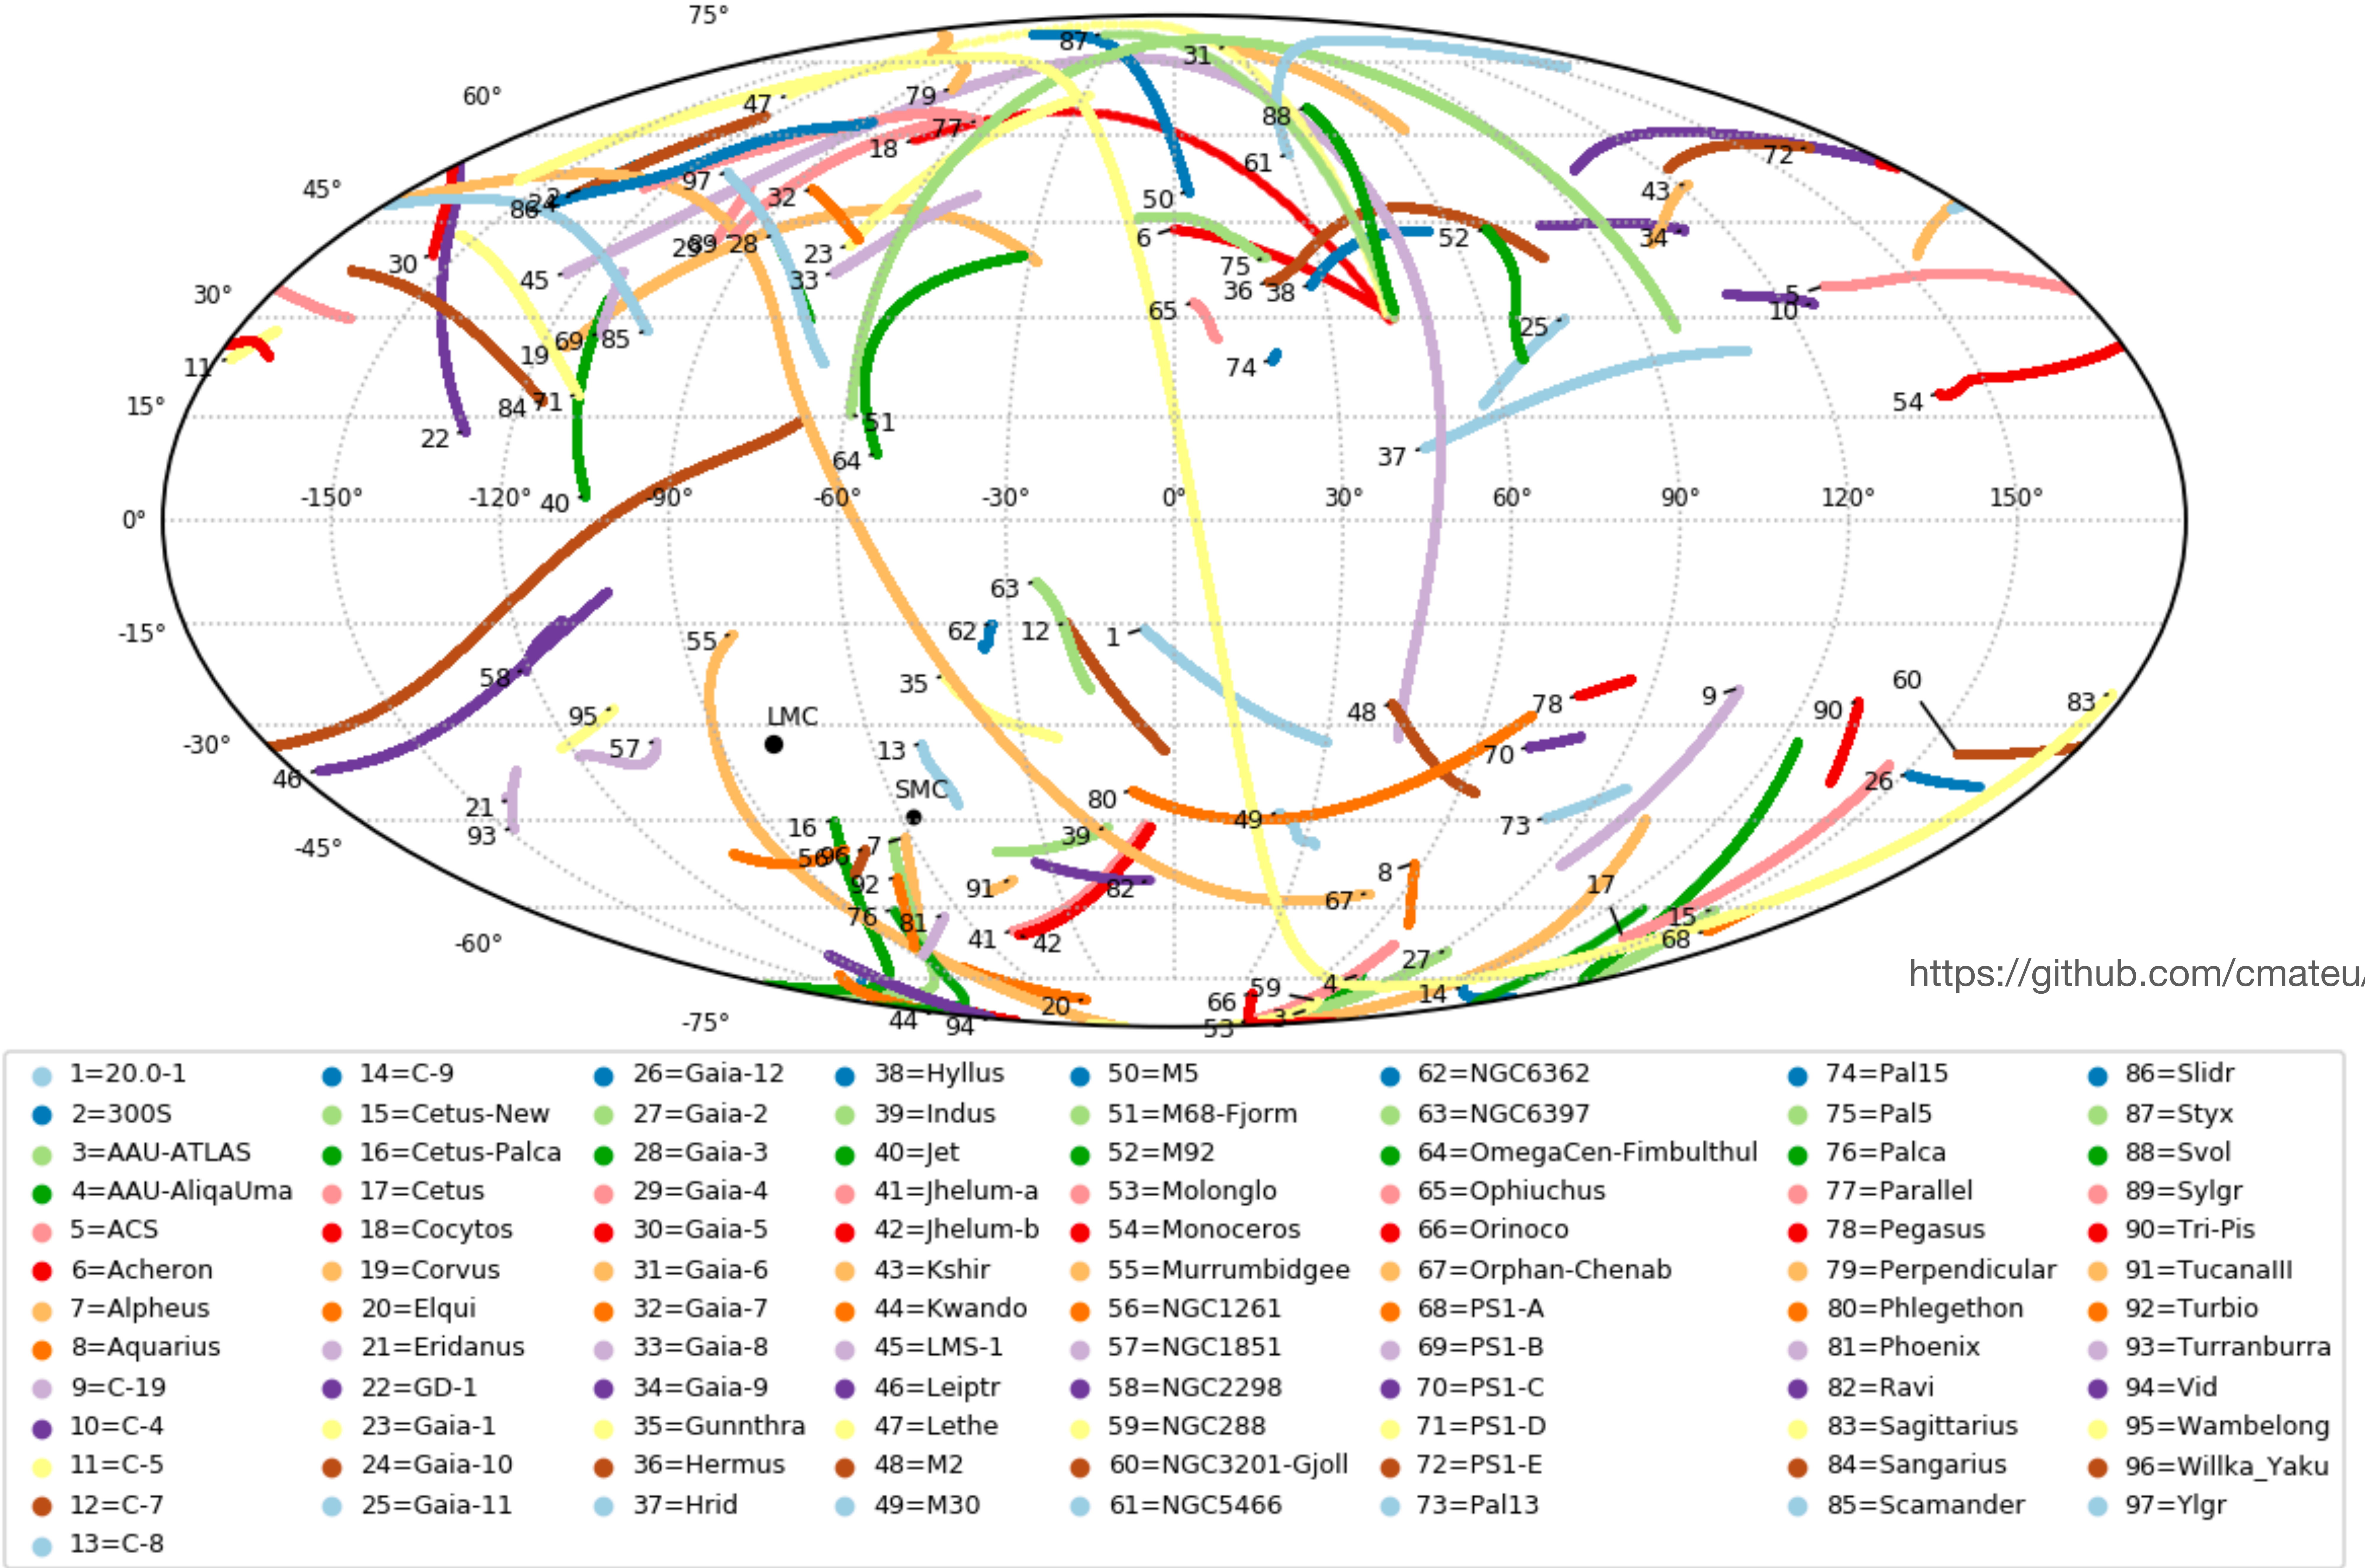

https://github.com/cmateu/galstreams

Via Machinae

[DS, Buckley, Necib ‘23] [DS, Buckley, Necib, Tamanas ‘21]

Angular position on the sky

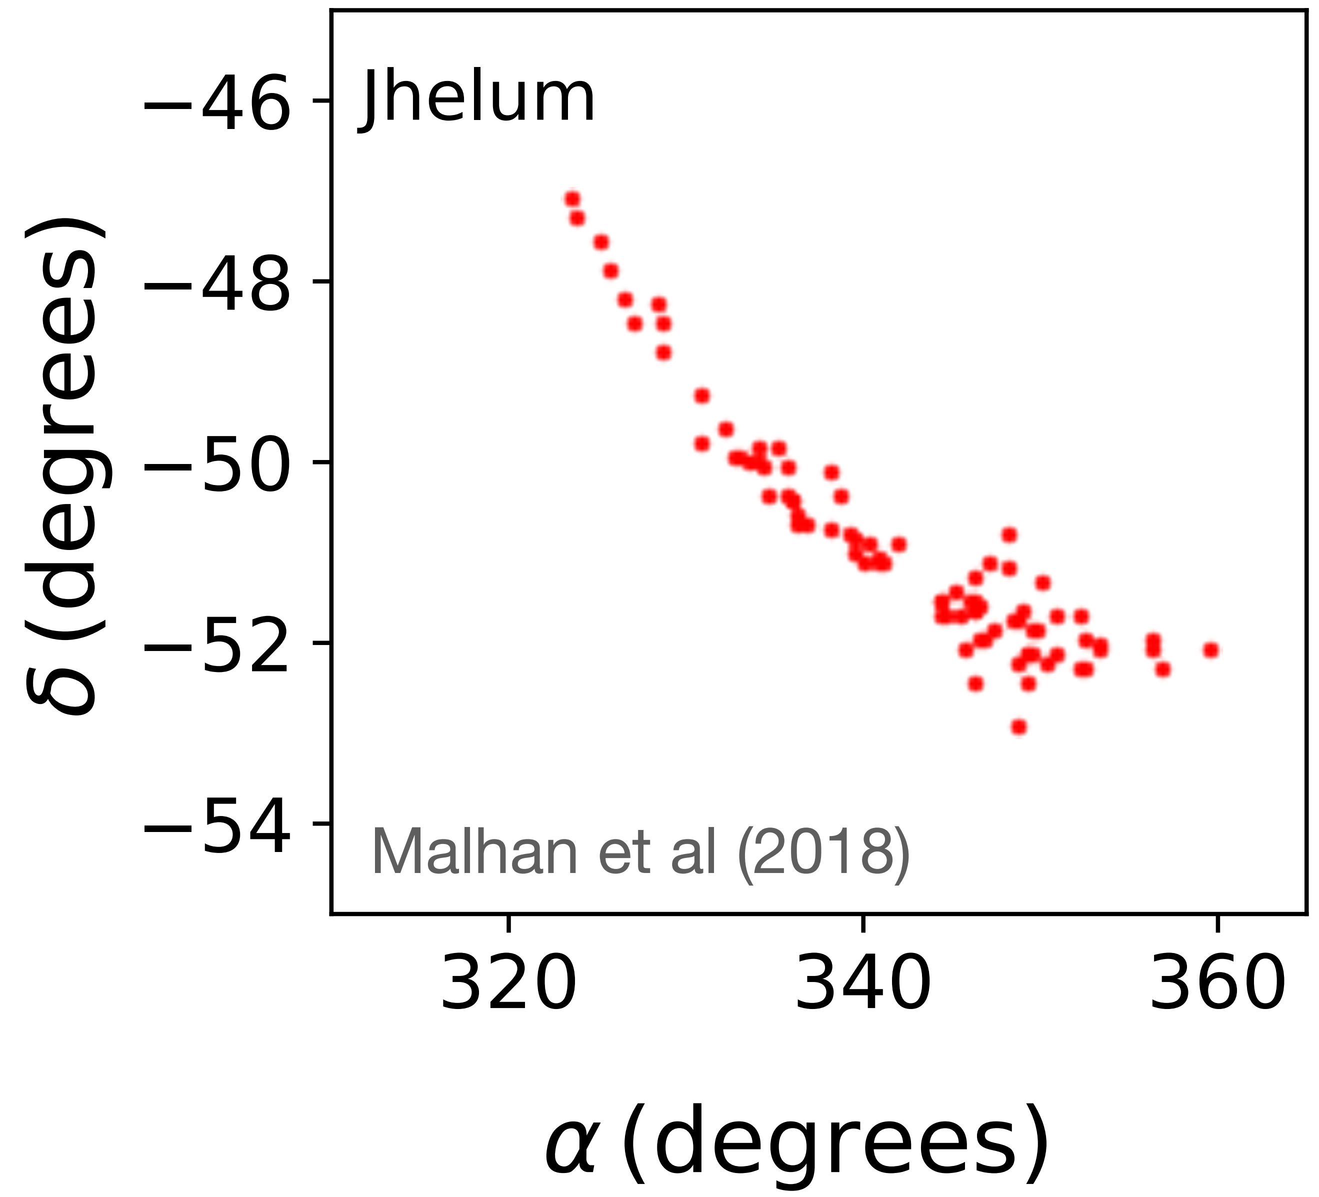

Velocity on the sky (proper motion)

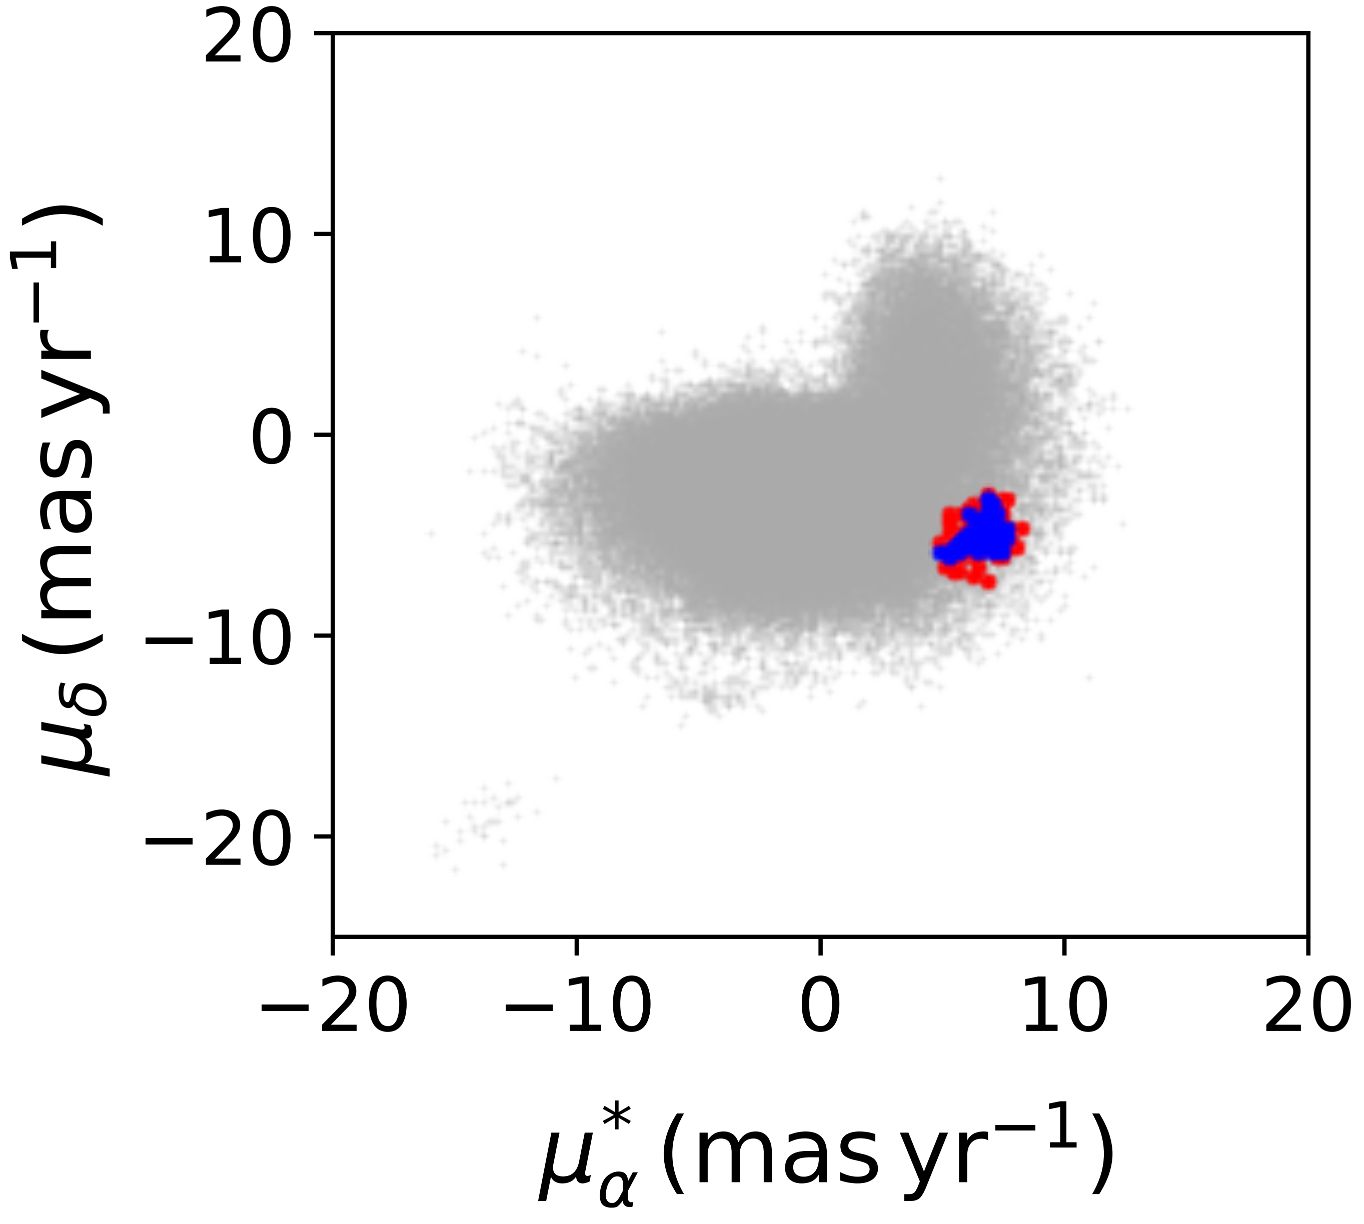

Photometry

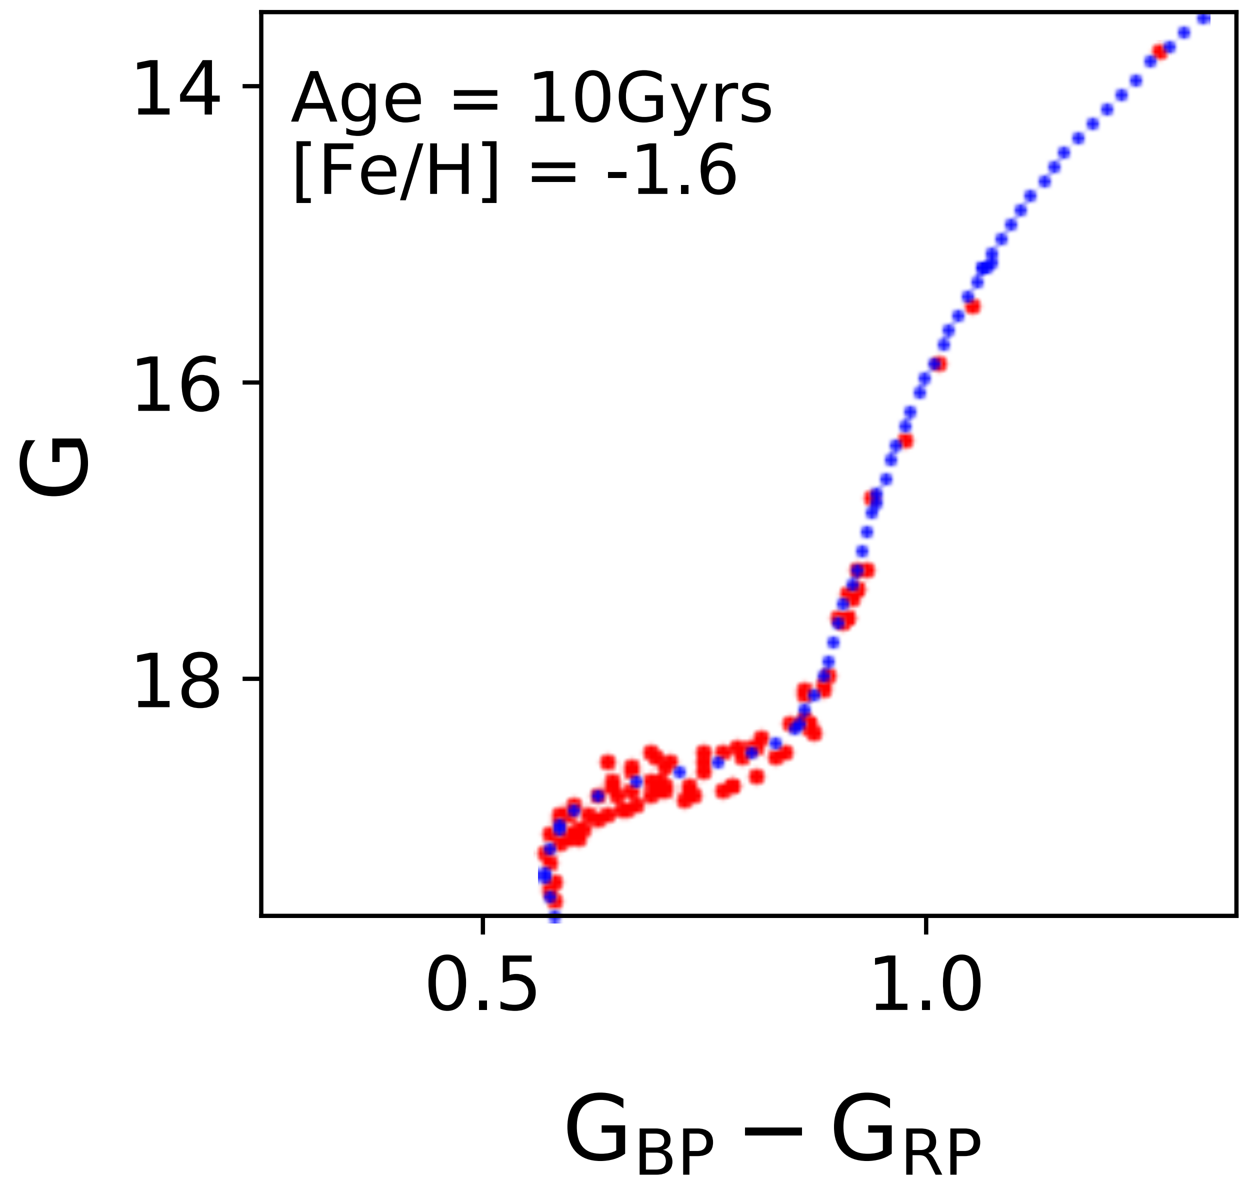

Via Machinae

[DS, Buckley, Necib ‘23] [DS, Buckley, Necib, Tamanas ‘21]

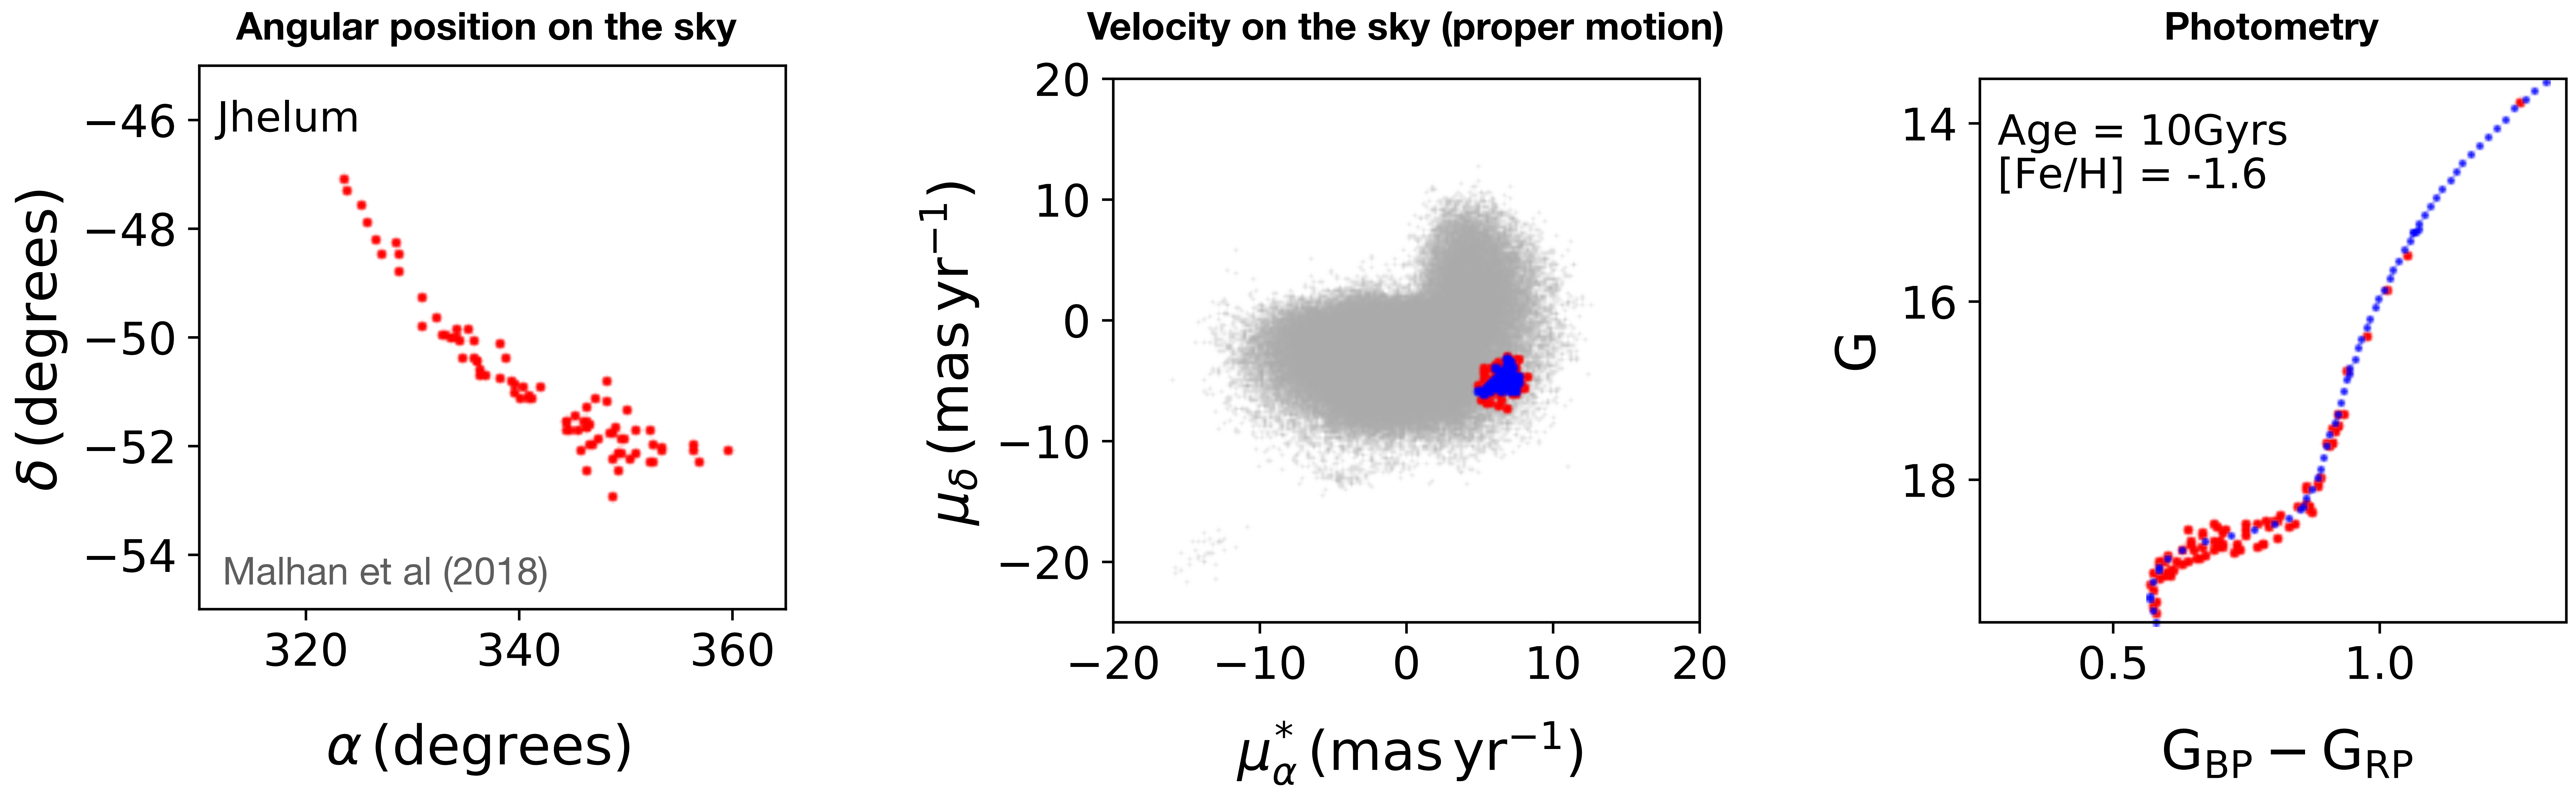

- Streams are local overdensities in multiple features — ideal for **enhanced bump hunt methods**!

Via Machinae

[DS, Buckley, Necib ‘23] [DS, Buckley, Necib, Tamanas ‘21]

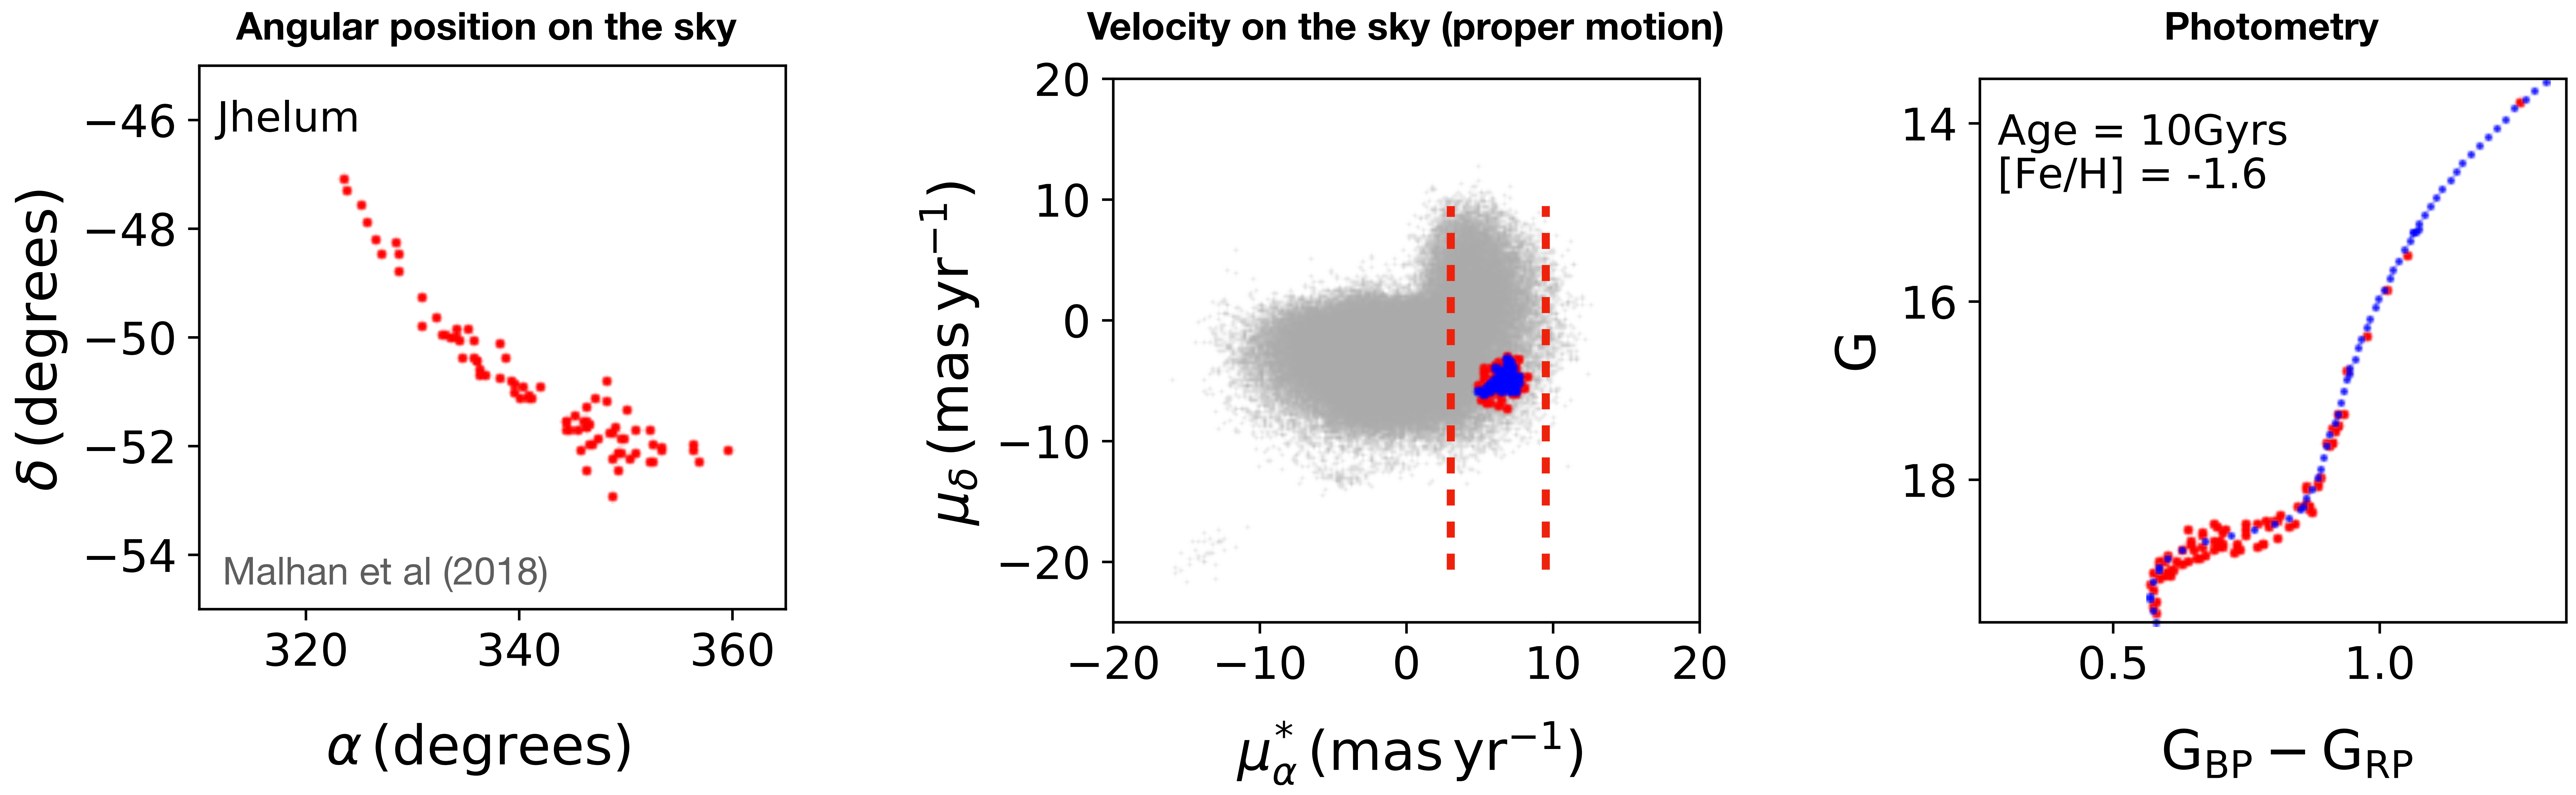

- Streams are local overdensities in multiple features — ideal for **enhanced bump hunt methods**!
- Choose either proper motion coordinate as resonant feature

Via Machinae

[DS, Buckley, Necib '23] [DS, Buckley, Necib, Tamanas '21]

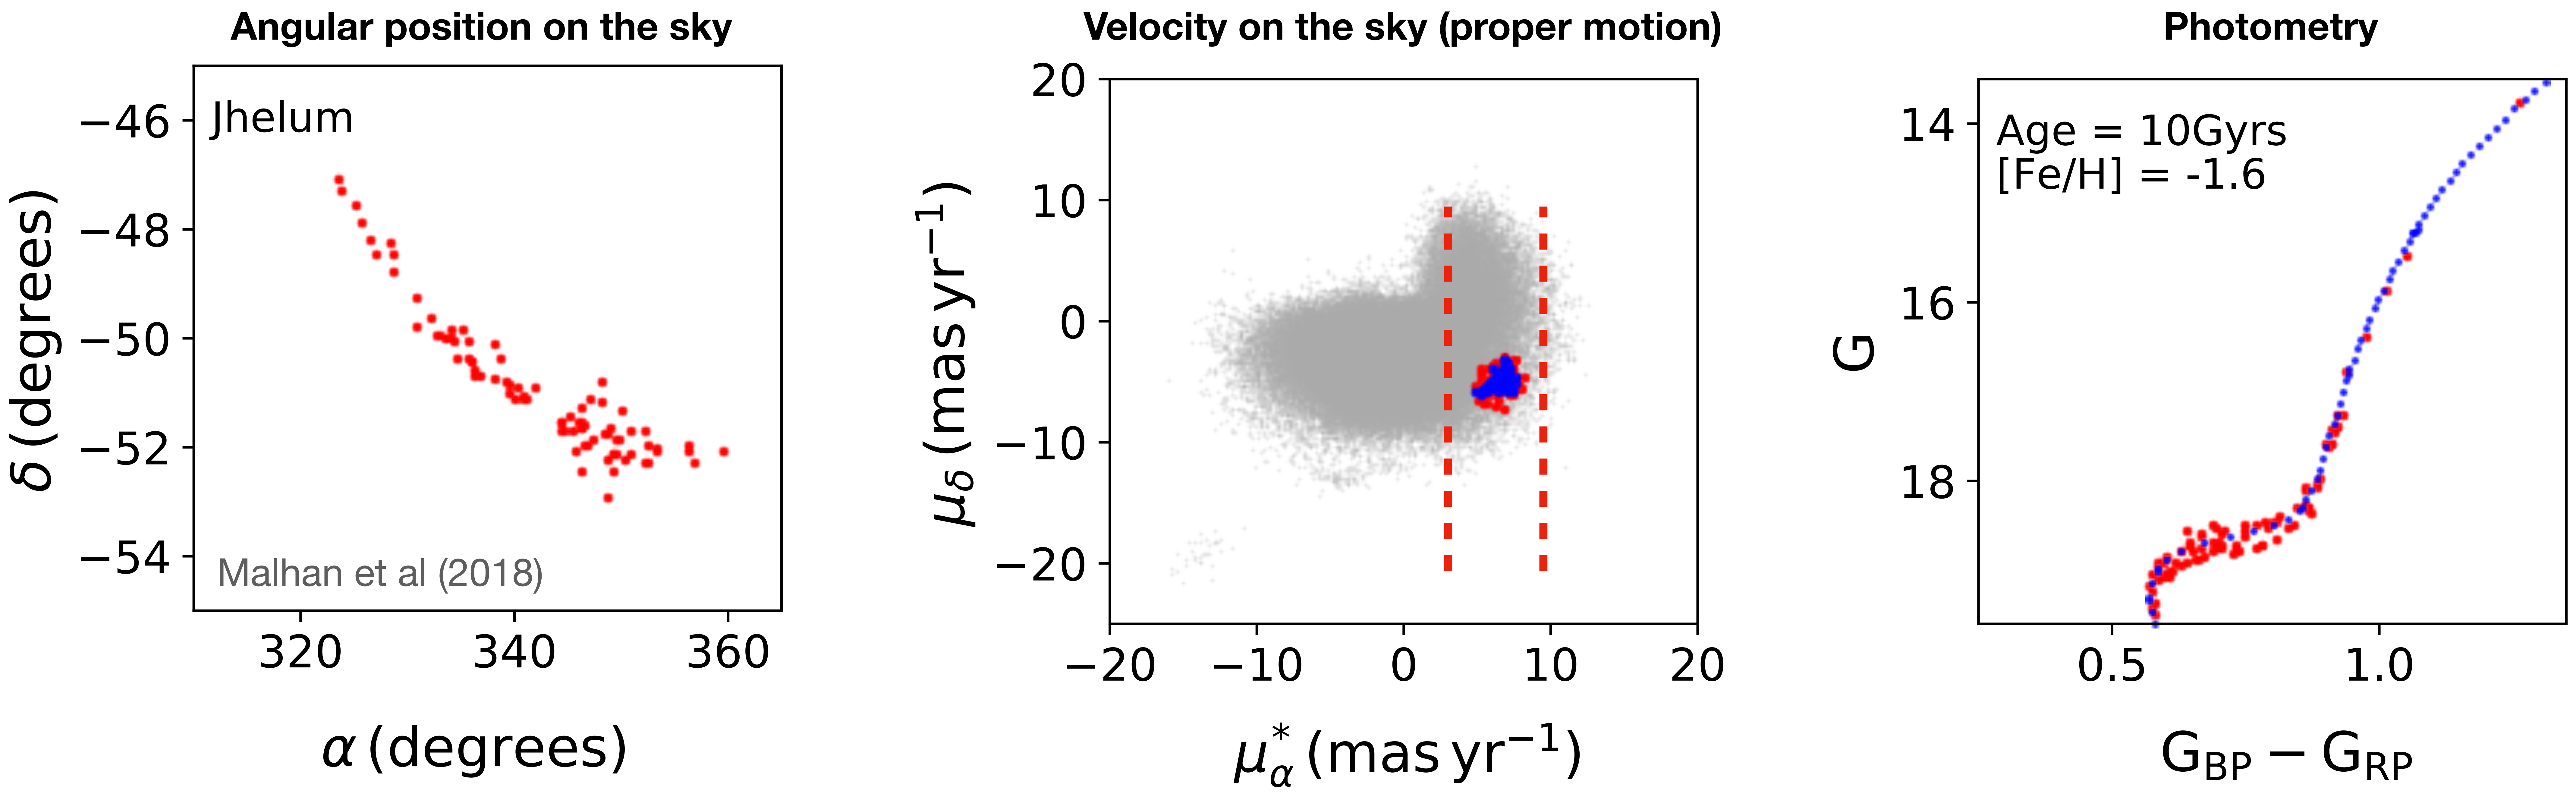

- Streams are local overdensities in multiple features — ideal for **enhanced bump hunt methods**!
- Choose either proper motion coordinate as resonant feature
- Use ANODE method to learn anomaly score with remaining five features

Via Machinae

[DS, Buckley, Necib ‘23] [DS, Buckley, Necib, Tamanas ‘21]

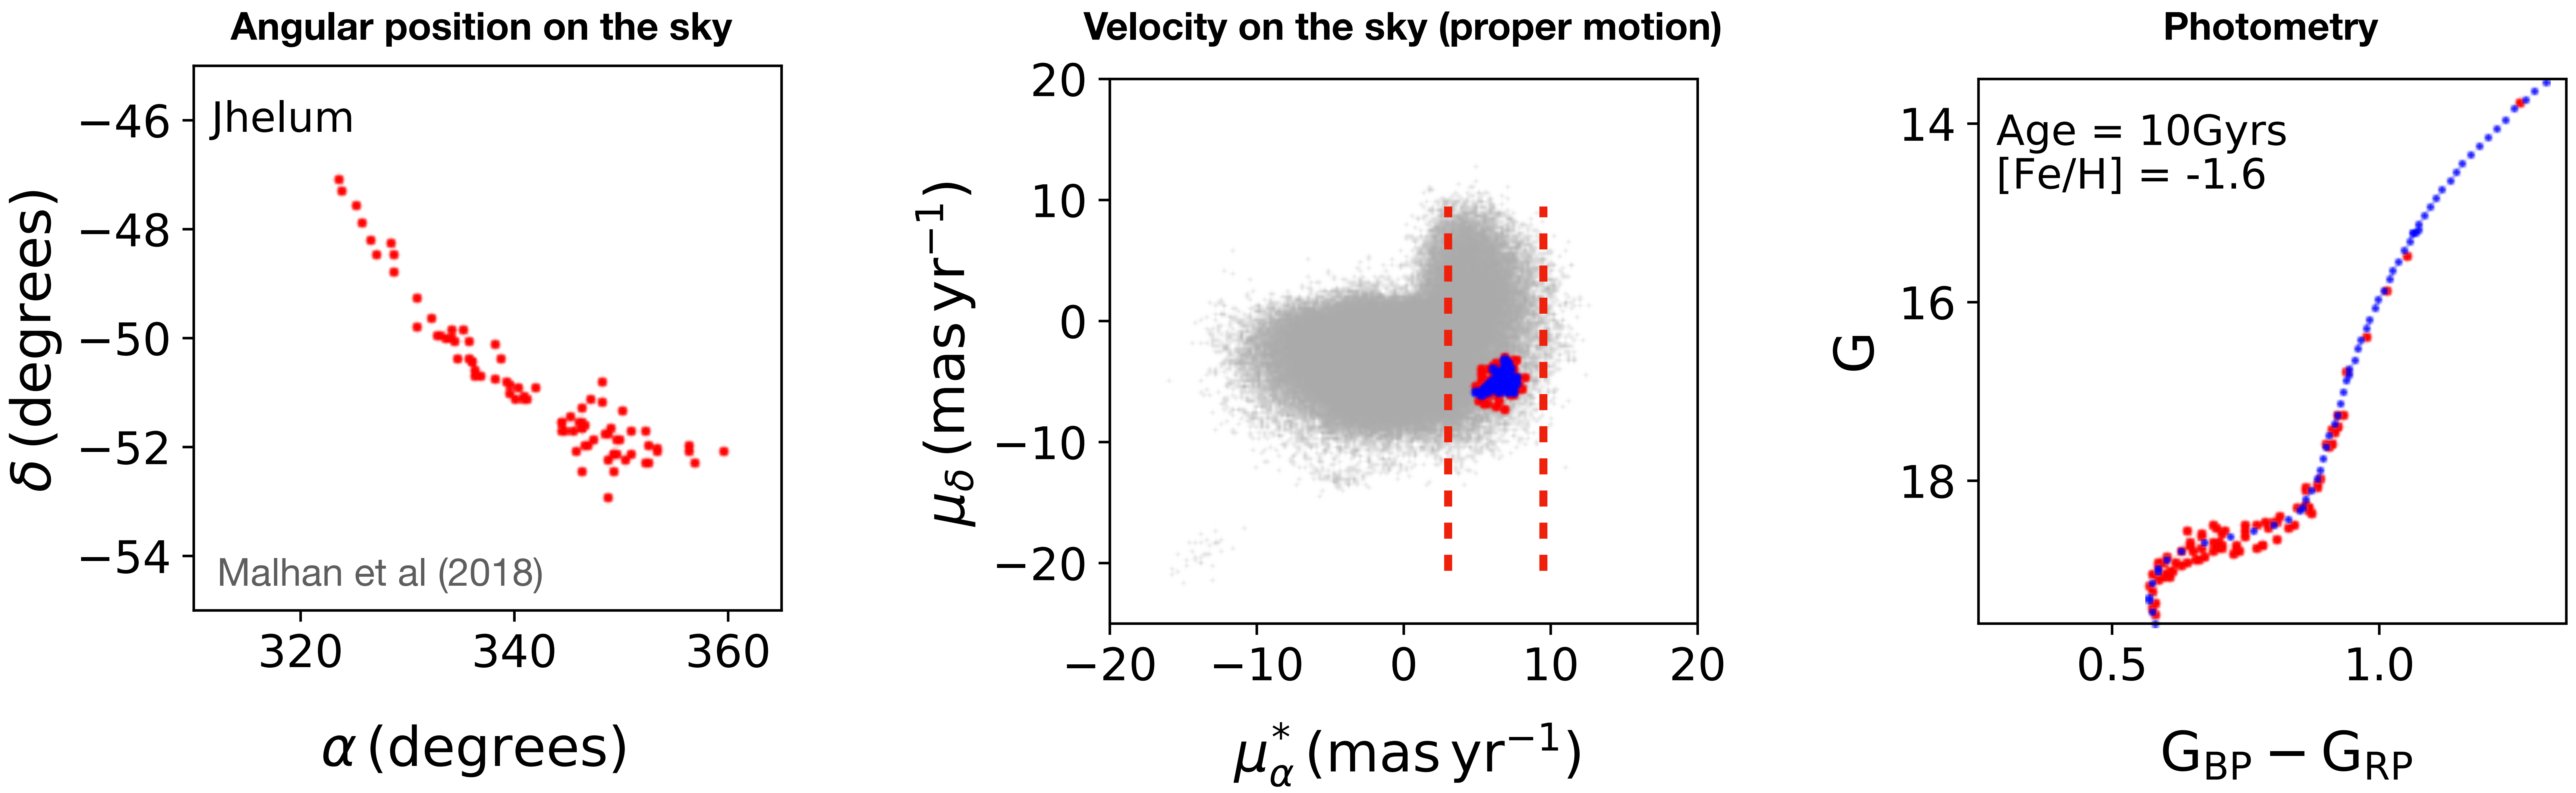

- Streams are local overdensities in multiple features — ideal for **enhanced bump hunt methods**!
- Choose either proper motion coordinate as resonant feature
- Use ANODE method to learn anomaly score with remaining five features

see also: DS+ Pettee et al (2305.03761) [CWoLa], DS+ Hallin et al (23xx.xxxxx) [CATHODE]

Core method — illustrated with GD-1 Stream

[DS, Buckley, Necib, Tamanas '21]

Fully data driven, simulation independent!

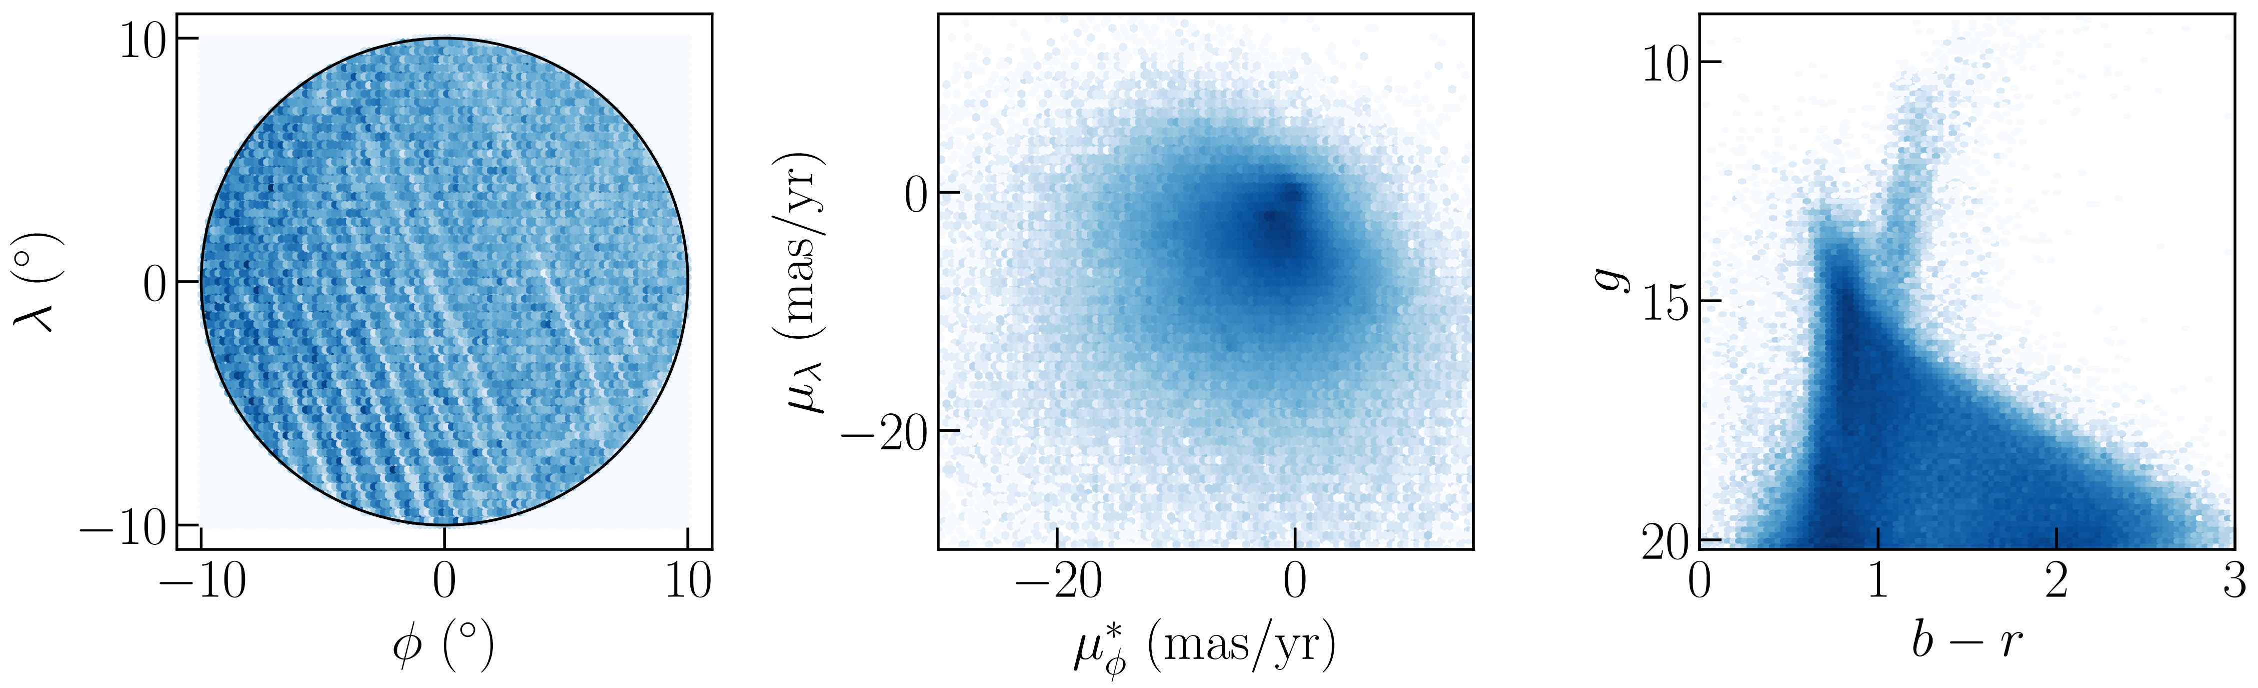

All stars in a patch of the sky containing (part of) GD-1 (ra,dec)=(148.6,24.2)

Core method — illustrated with GD-1 Stream

[DS, Buckley, Necib, Tamanas '21]

Fully data driven, simulation independent!

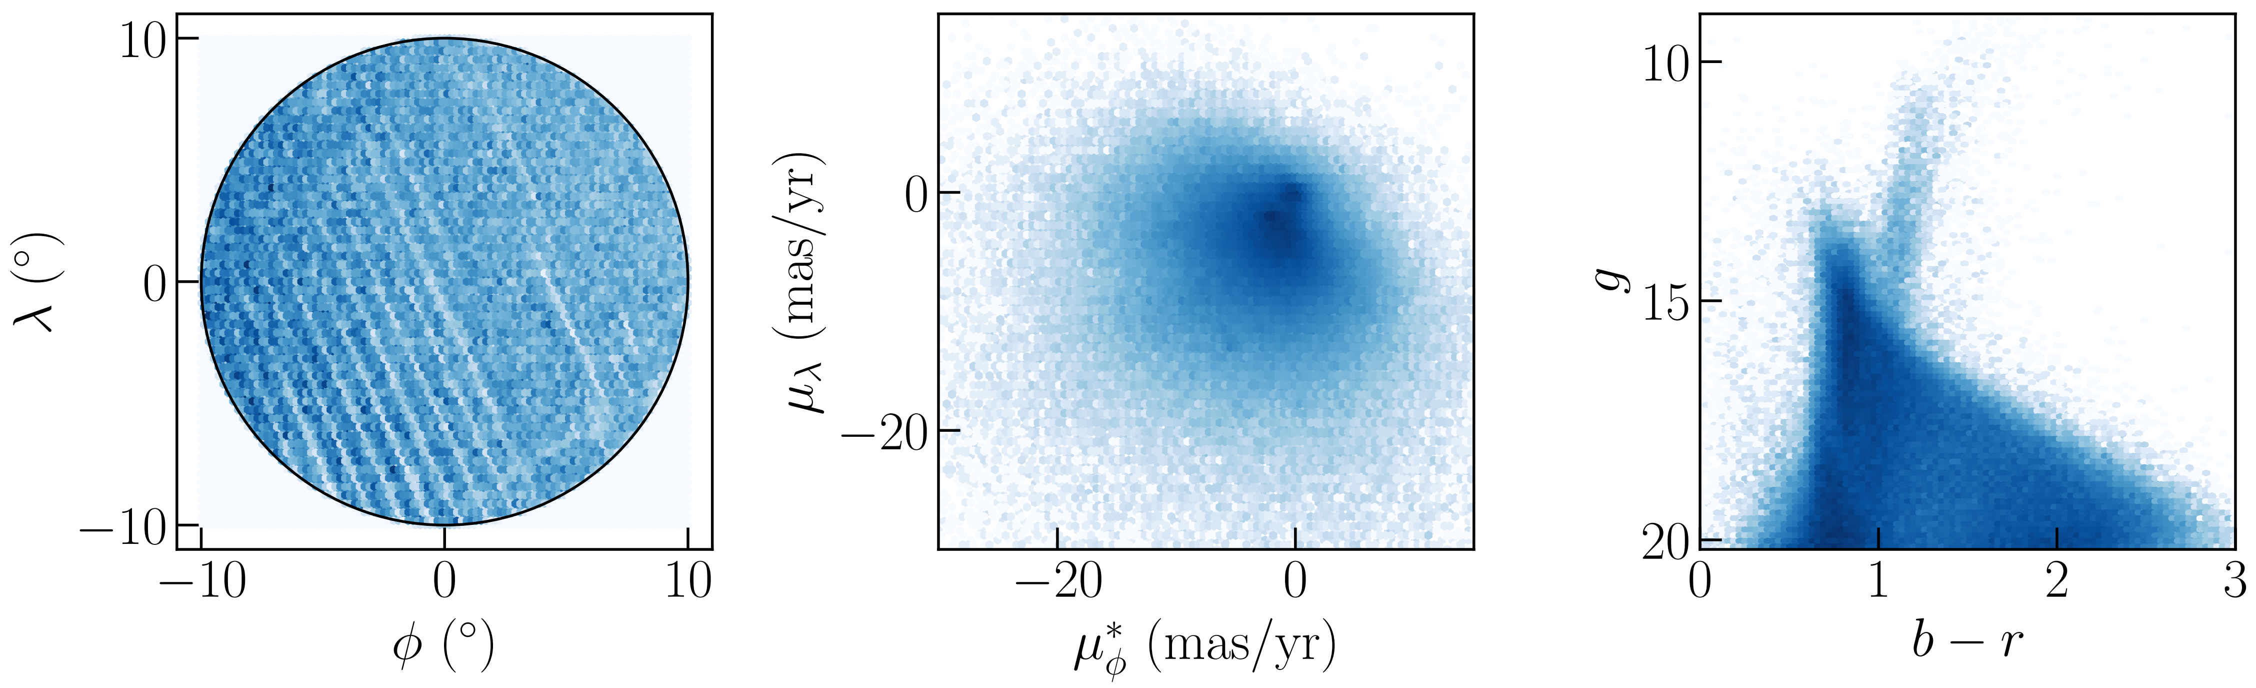

All stars in a patch of the sky containing (part of) GD-1 (ra,dec)=(148.6,24.2)

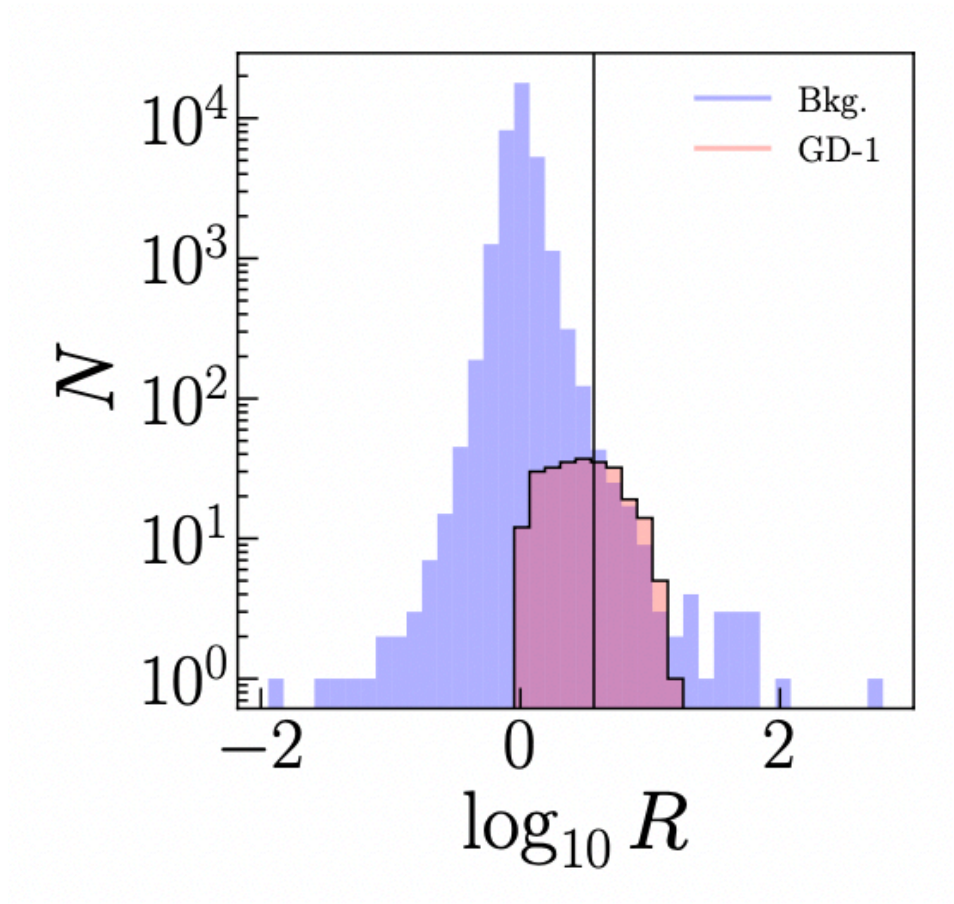

Core method — illustrated with GD-1 Stream

[DS, Buckley, Necib, Tamanas ’21]

Fully data driven, simulation independent!

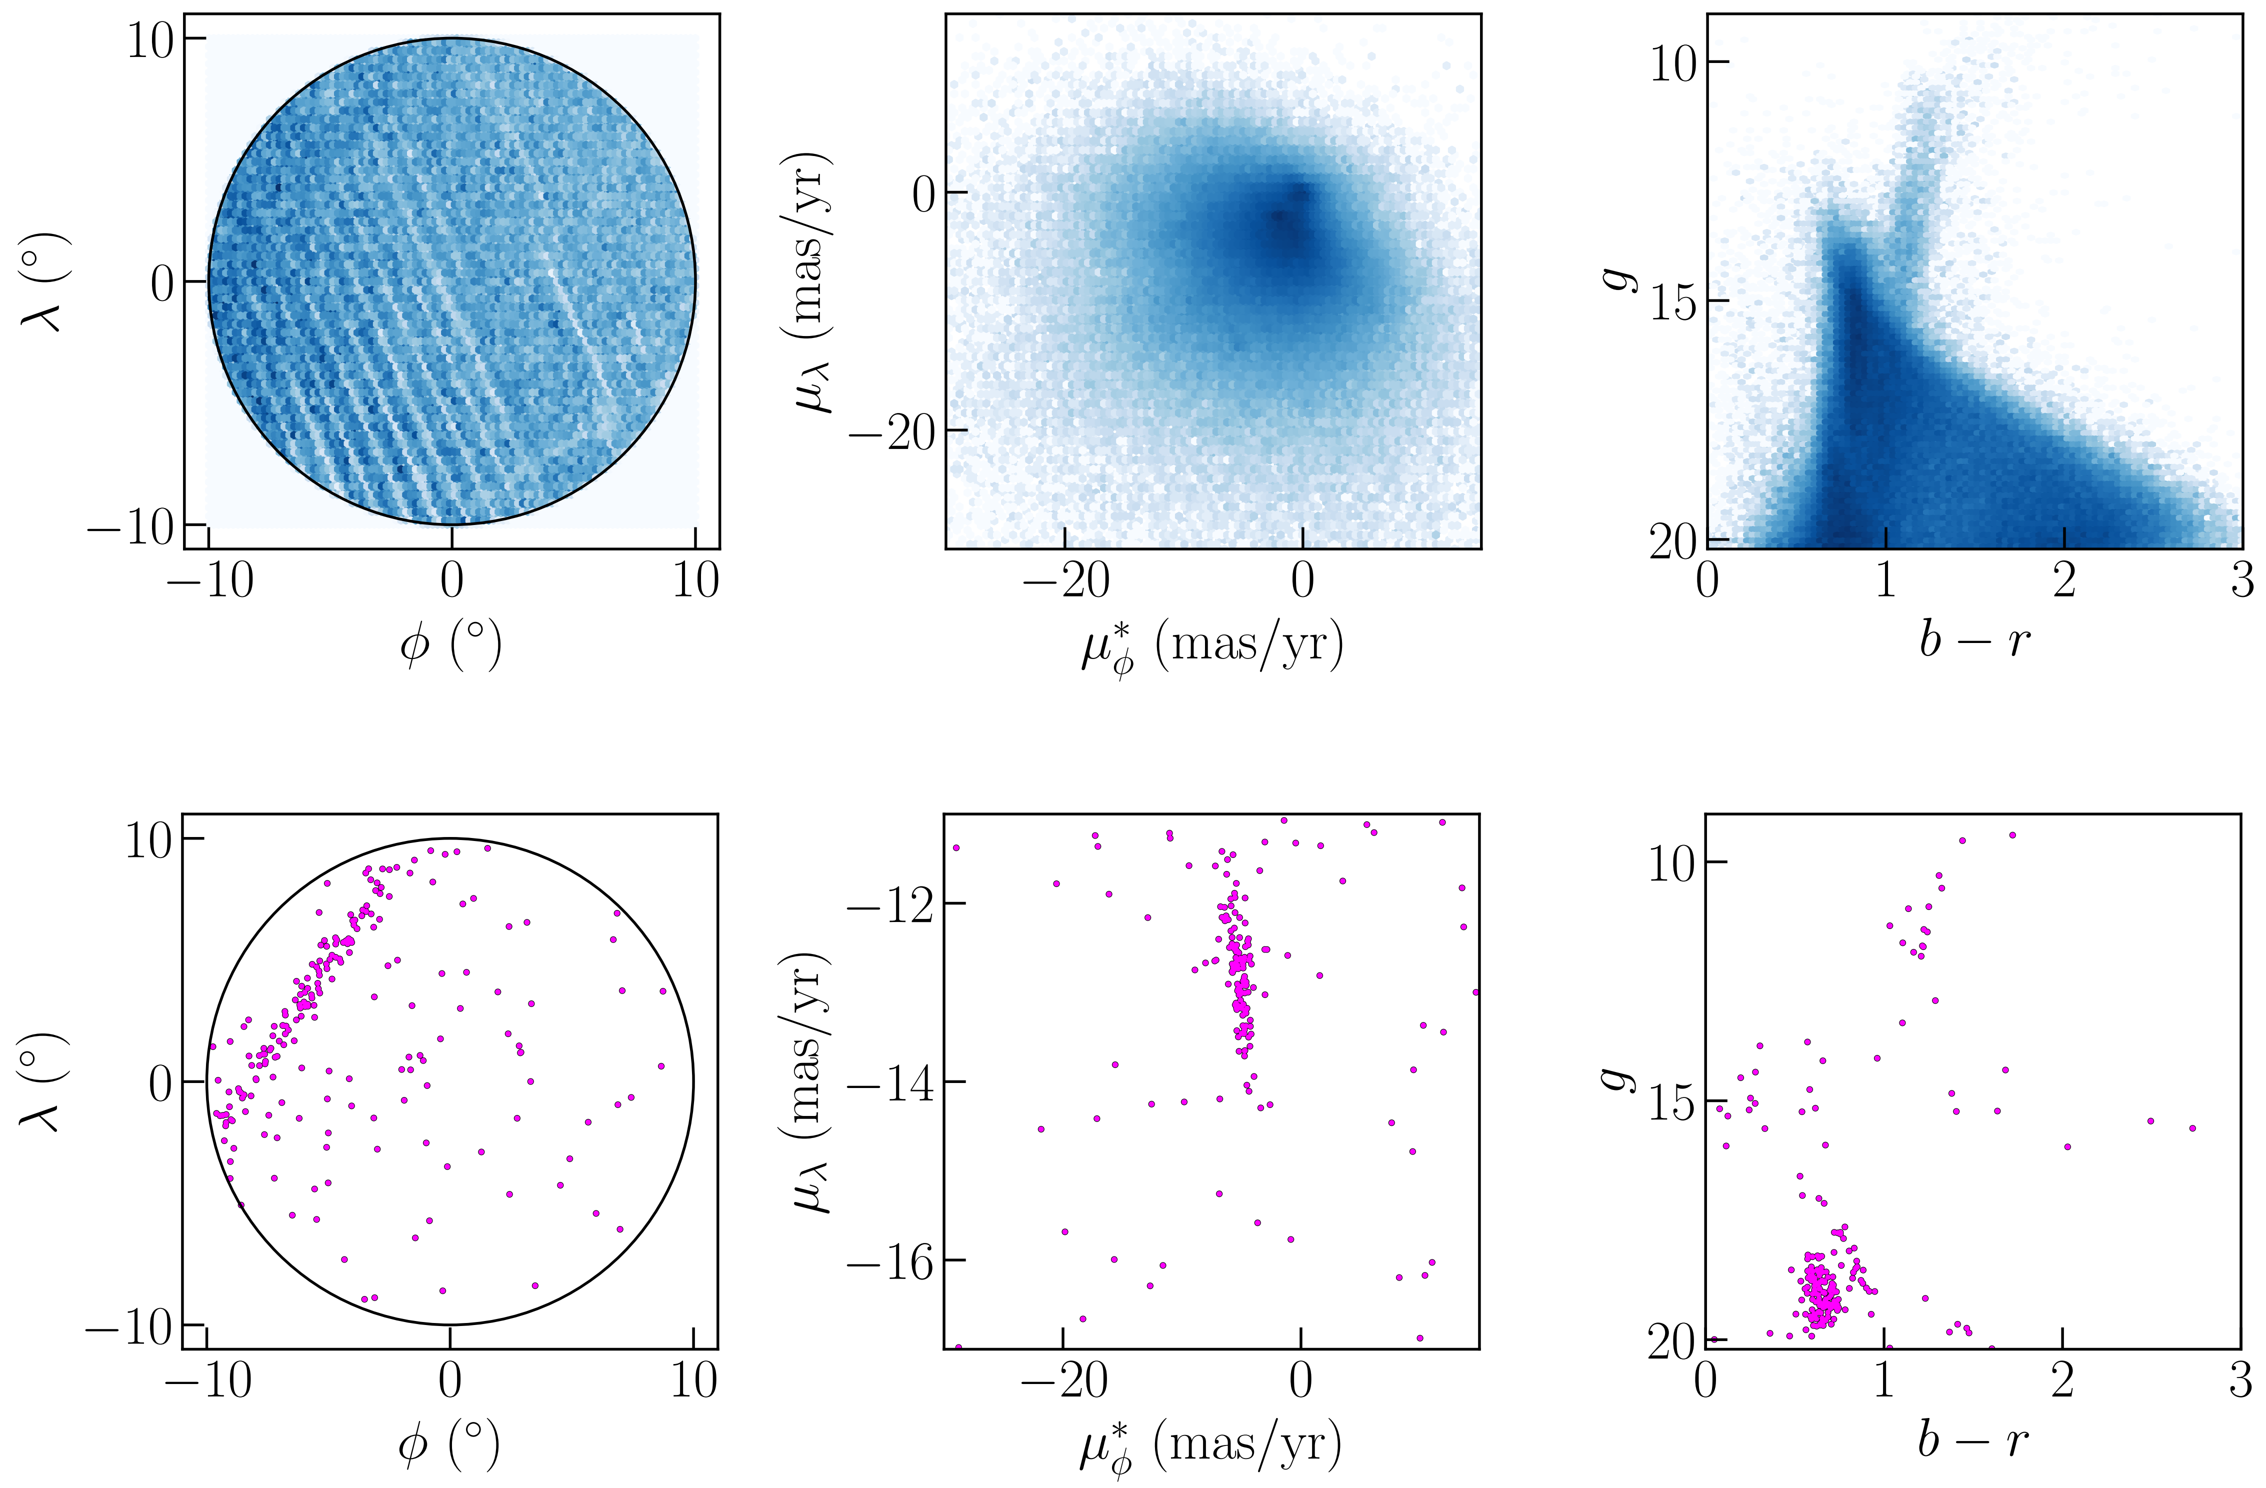

All stars in a patch of the sky containing (part of) GD-1 (ra,dec)=(148.6,24.2)

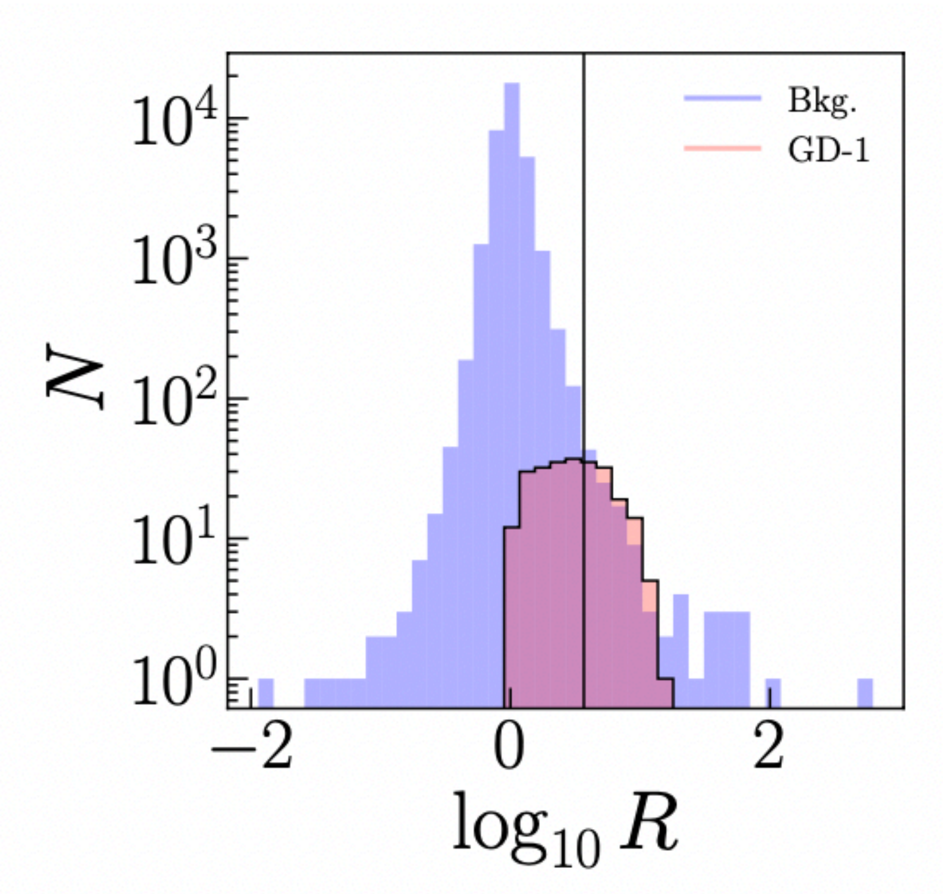

Stars in SR after cut on $R(x)$ obtained from ANODE

Core method — illustrated with GD-1 Stream

[DS, Buckley, Necib, Tamanas '21]

Fully data driven, simulation independent!

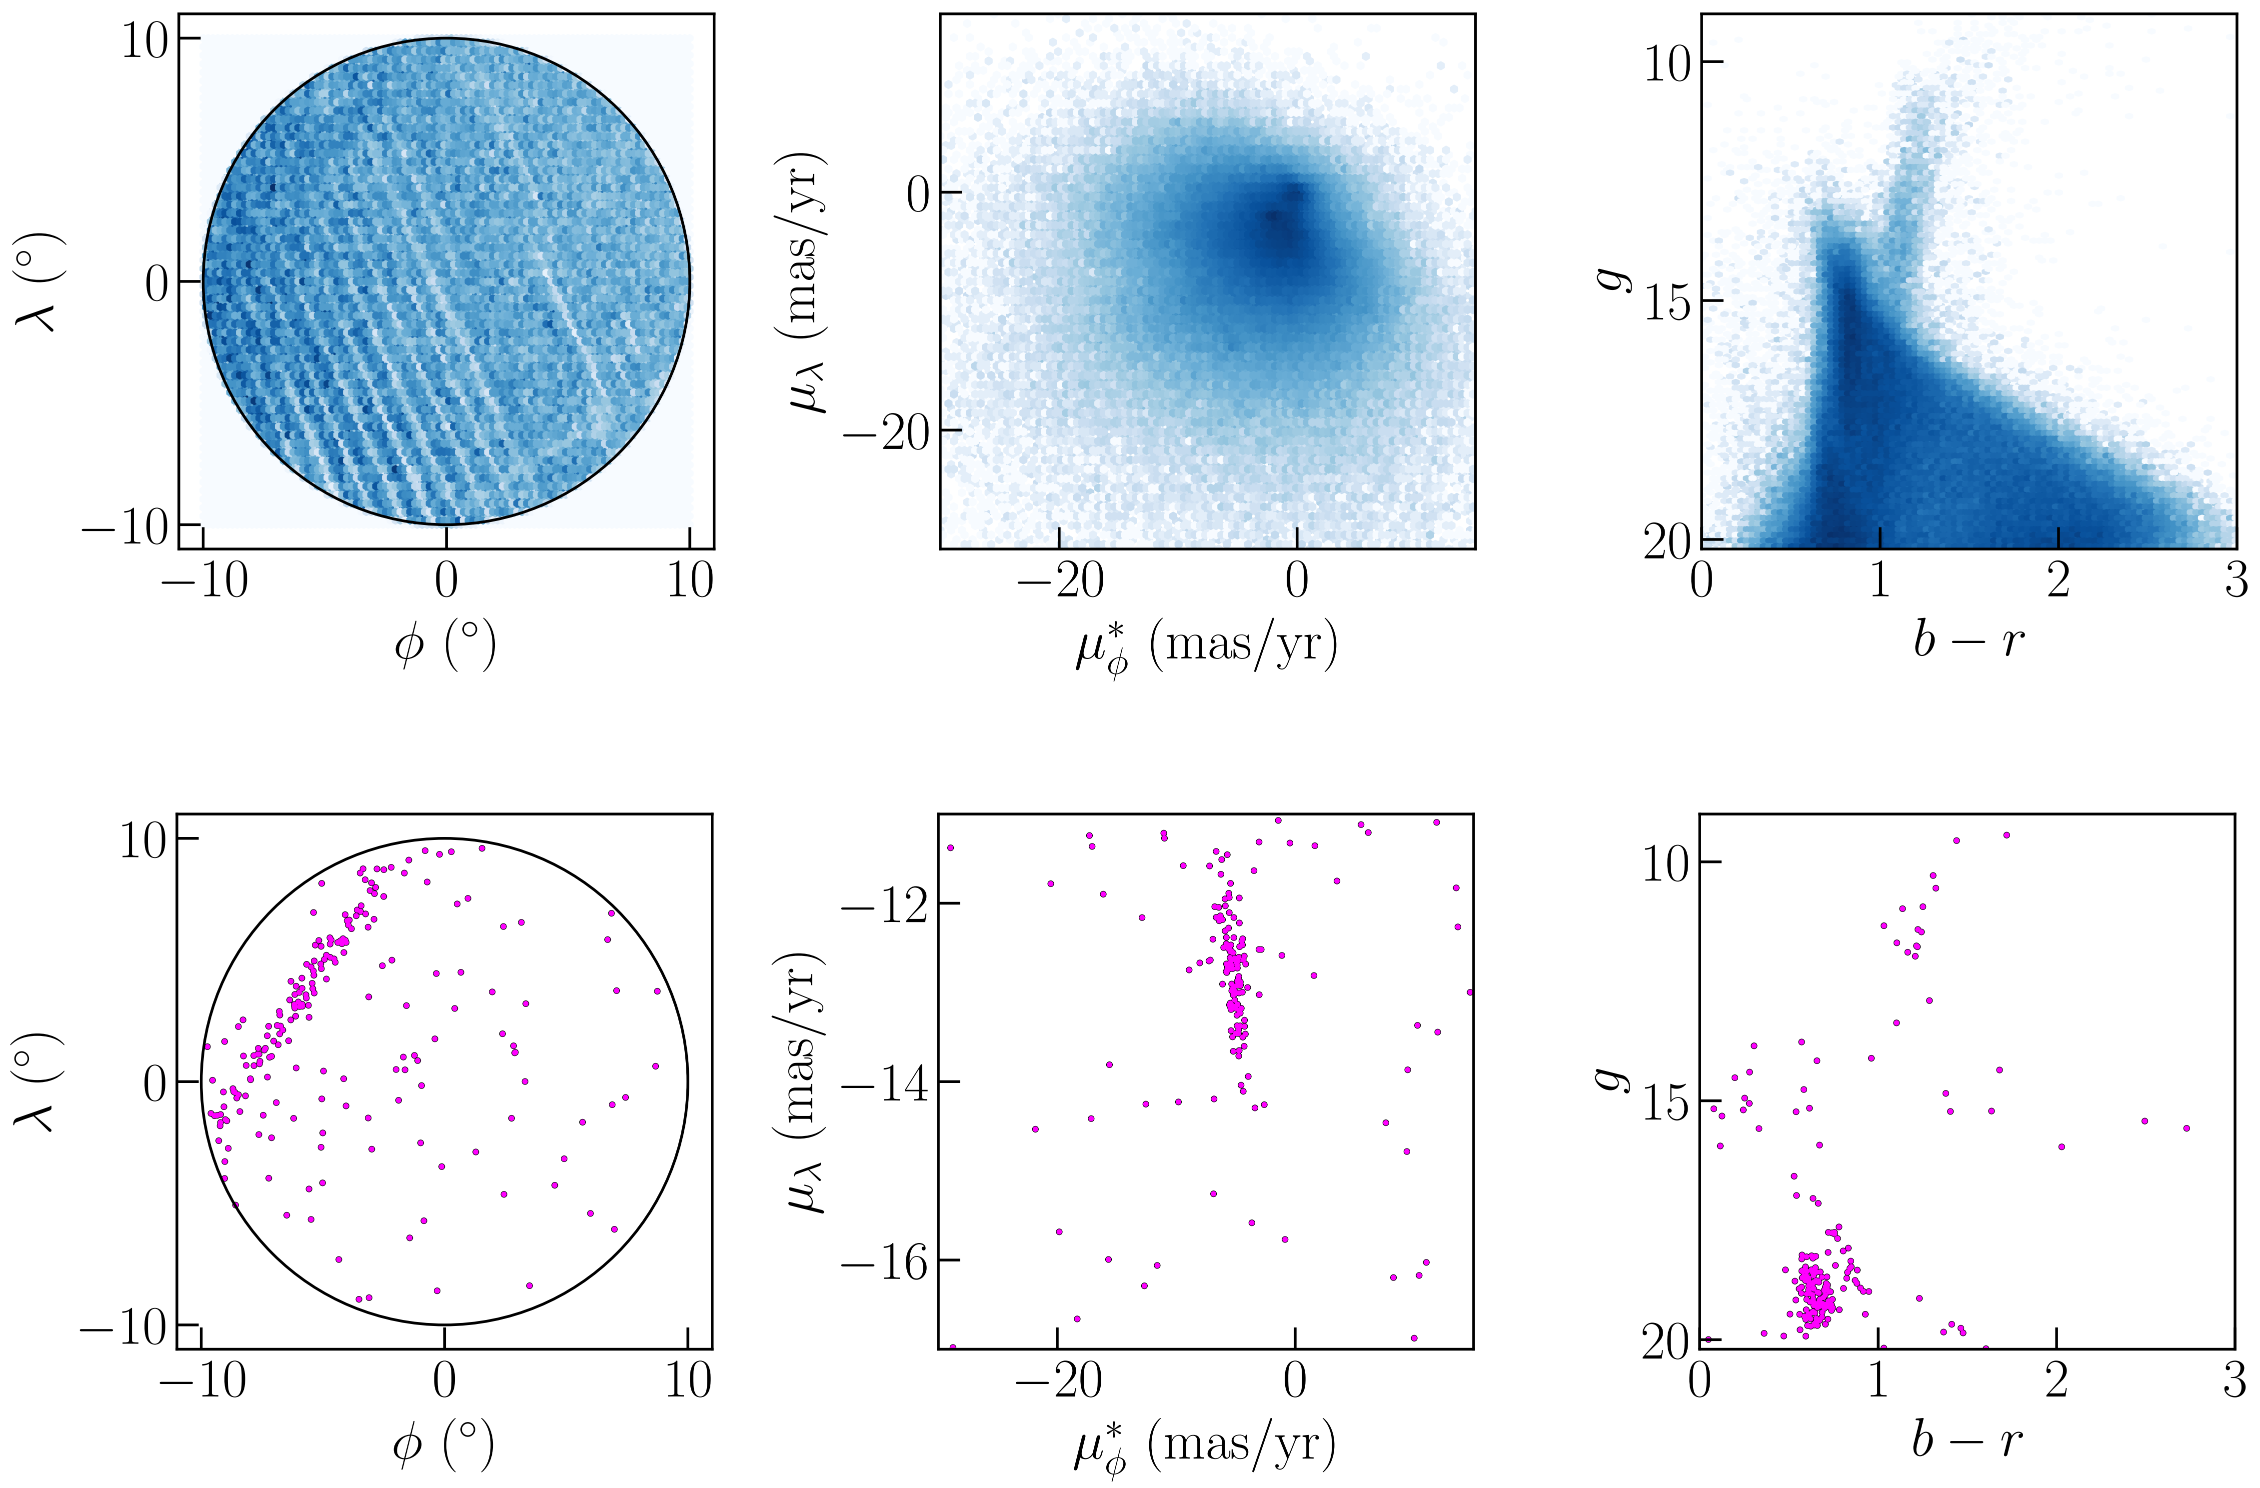

All stars in a patch of the sky containing (part of) GD-1 (ra,dec)=(148.6,24.2)

Stars in SR after cut on $R(x)$ obtained from ANODE

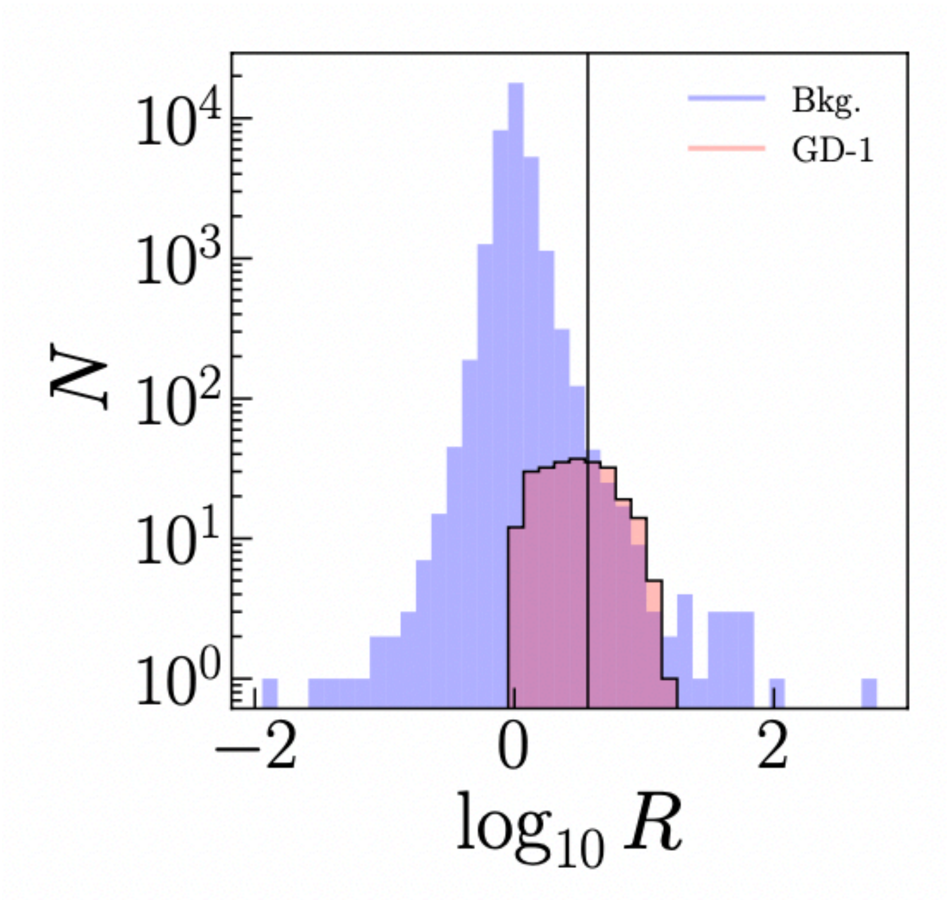

The method works!

ANODE on Gaia data

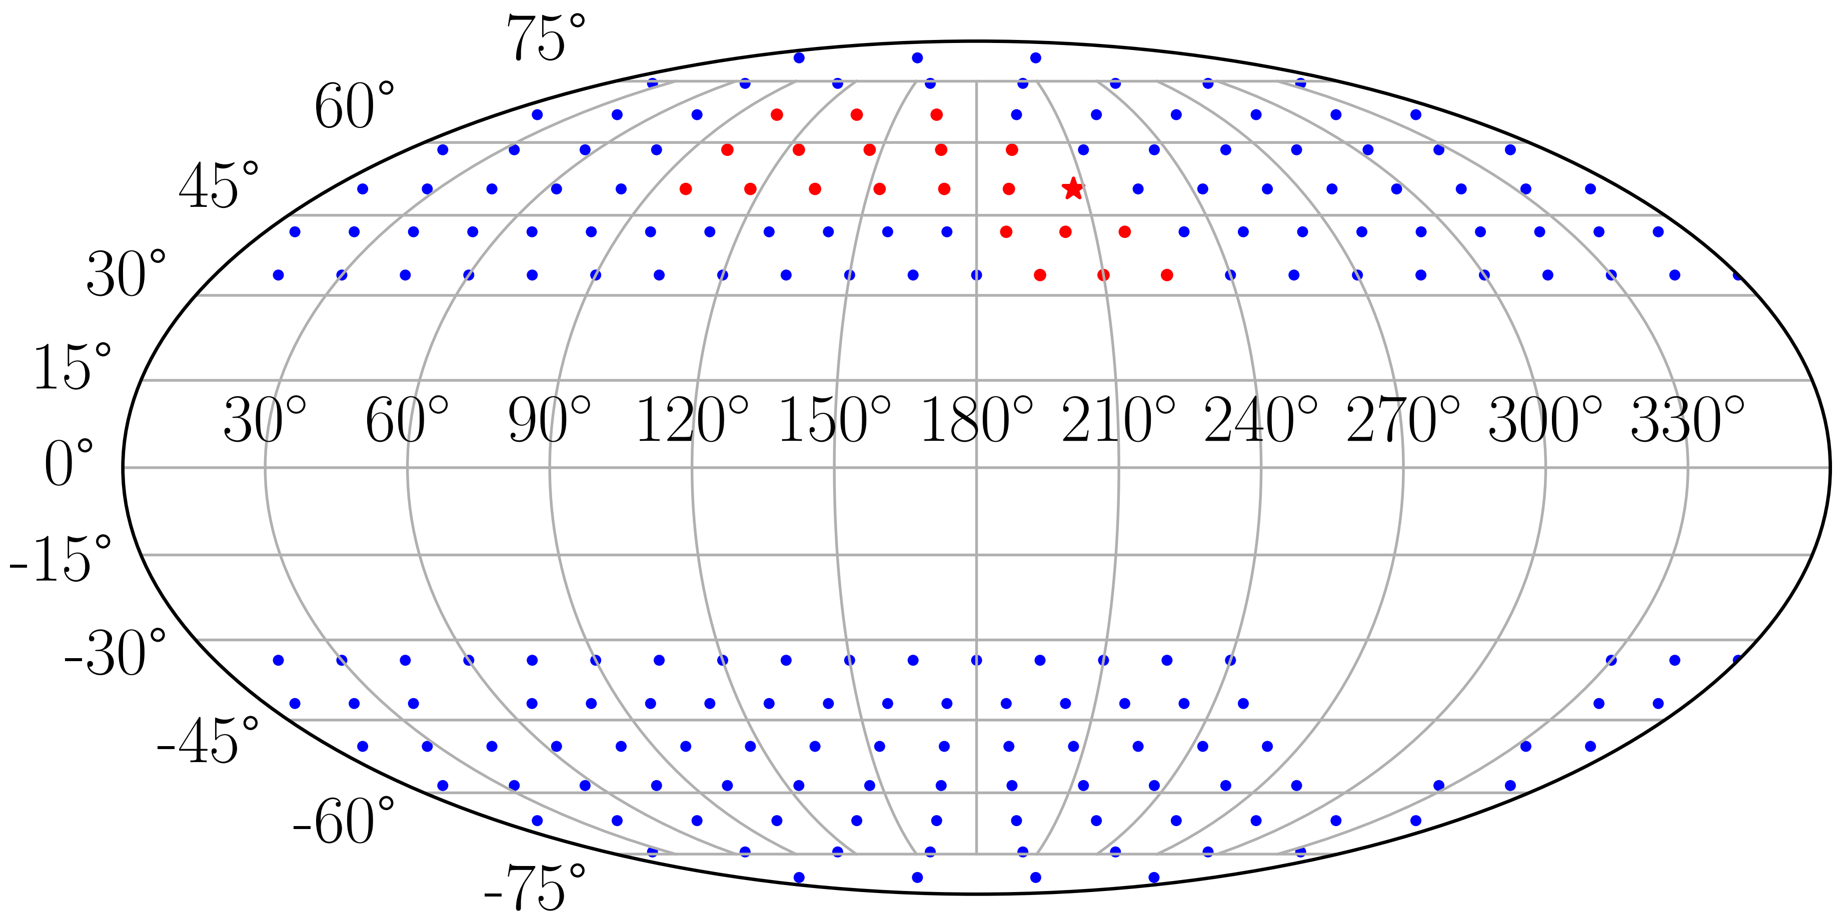

ANODE on Gaia data

- So the method works!

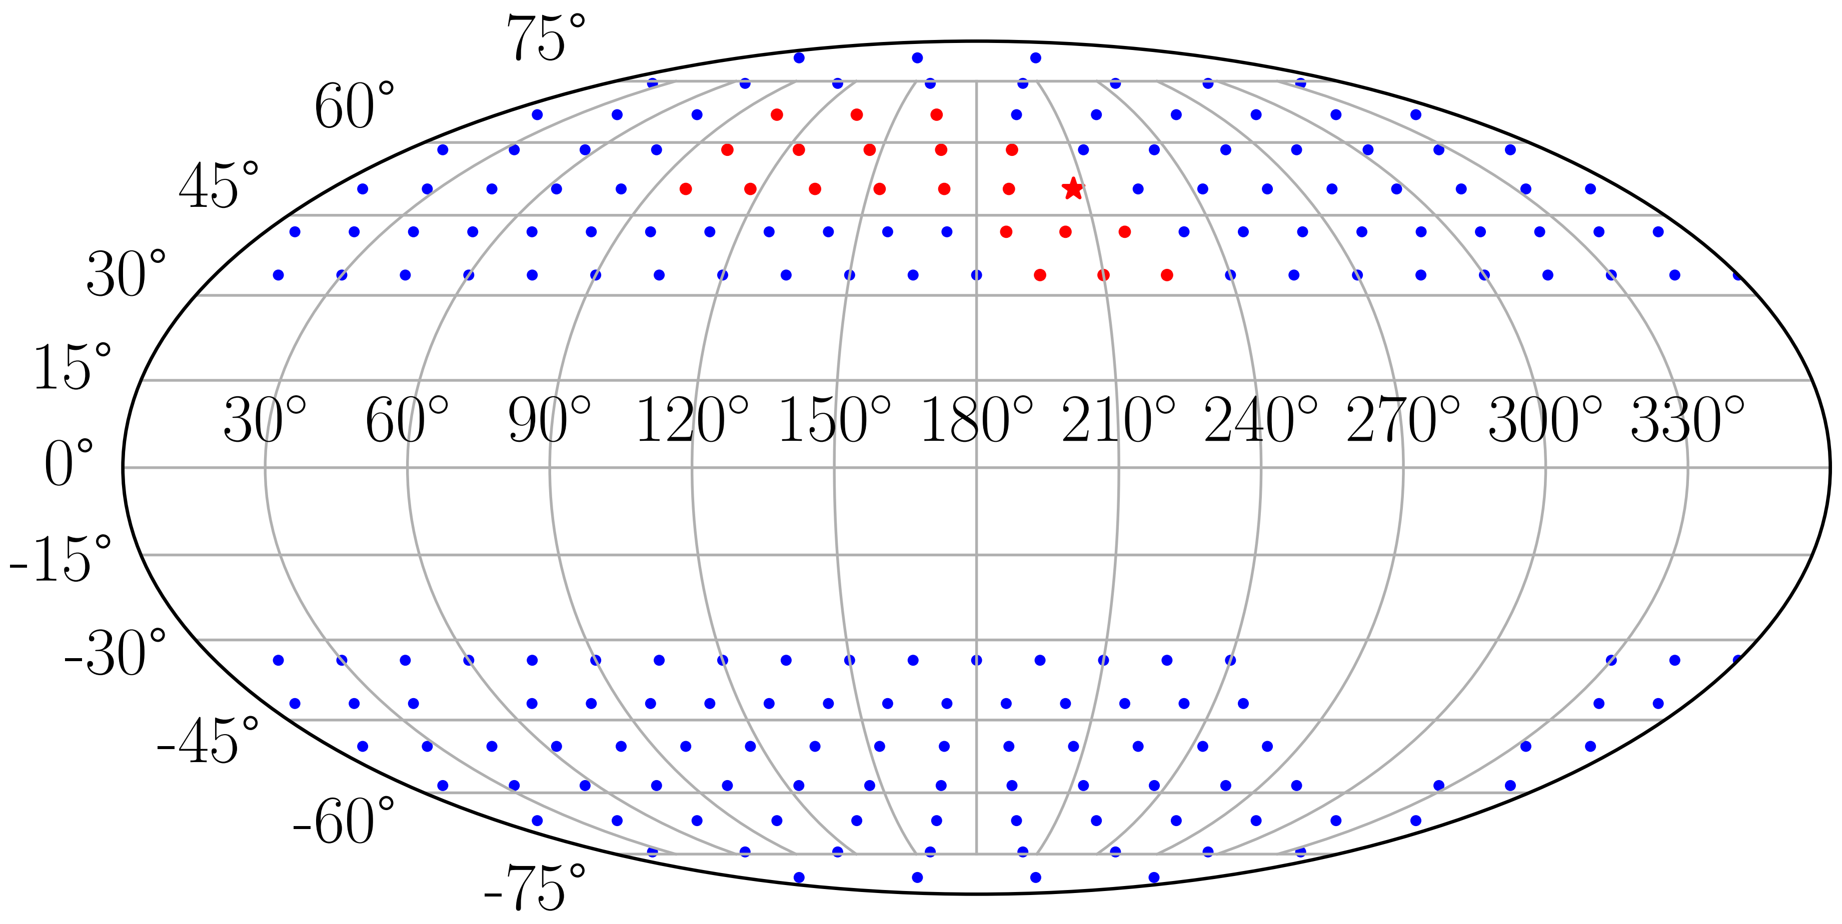

ANODE on Gaia data

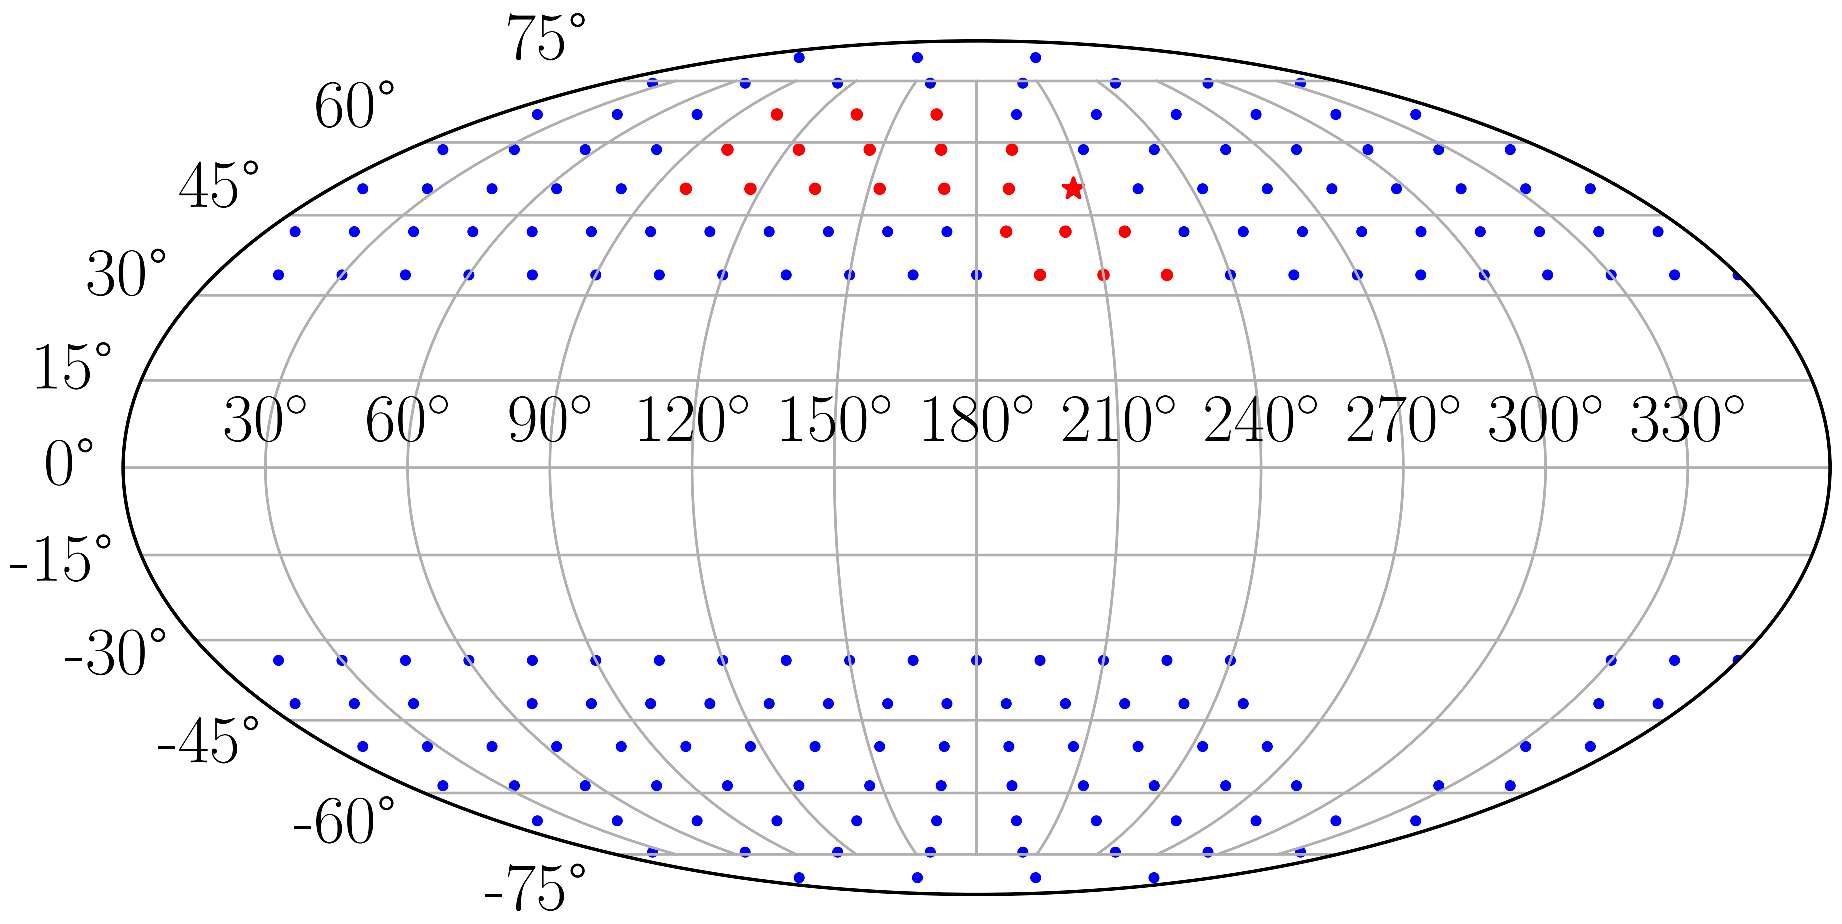

- So the method works!
- How do we apply this to the entire sky?

ANODE on Gaia data

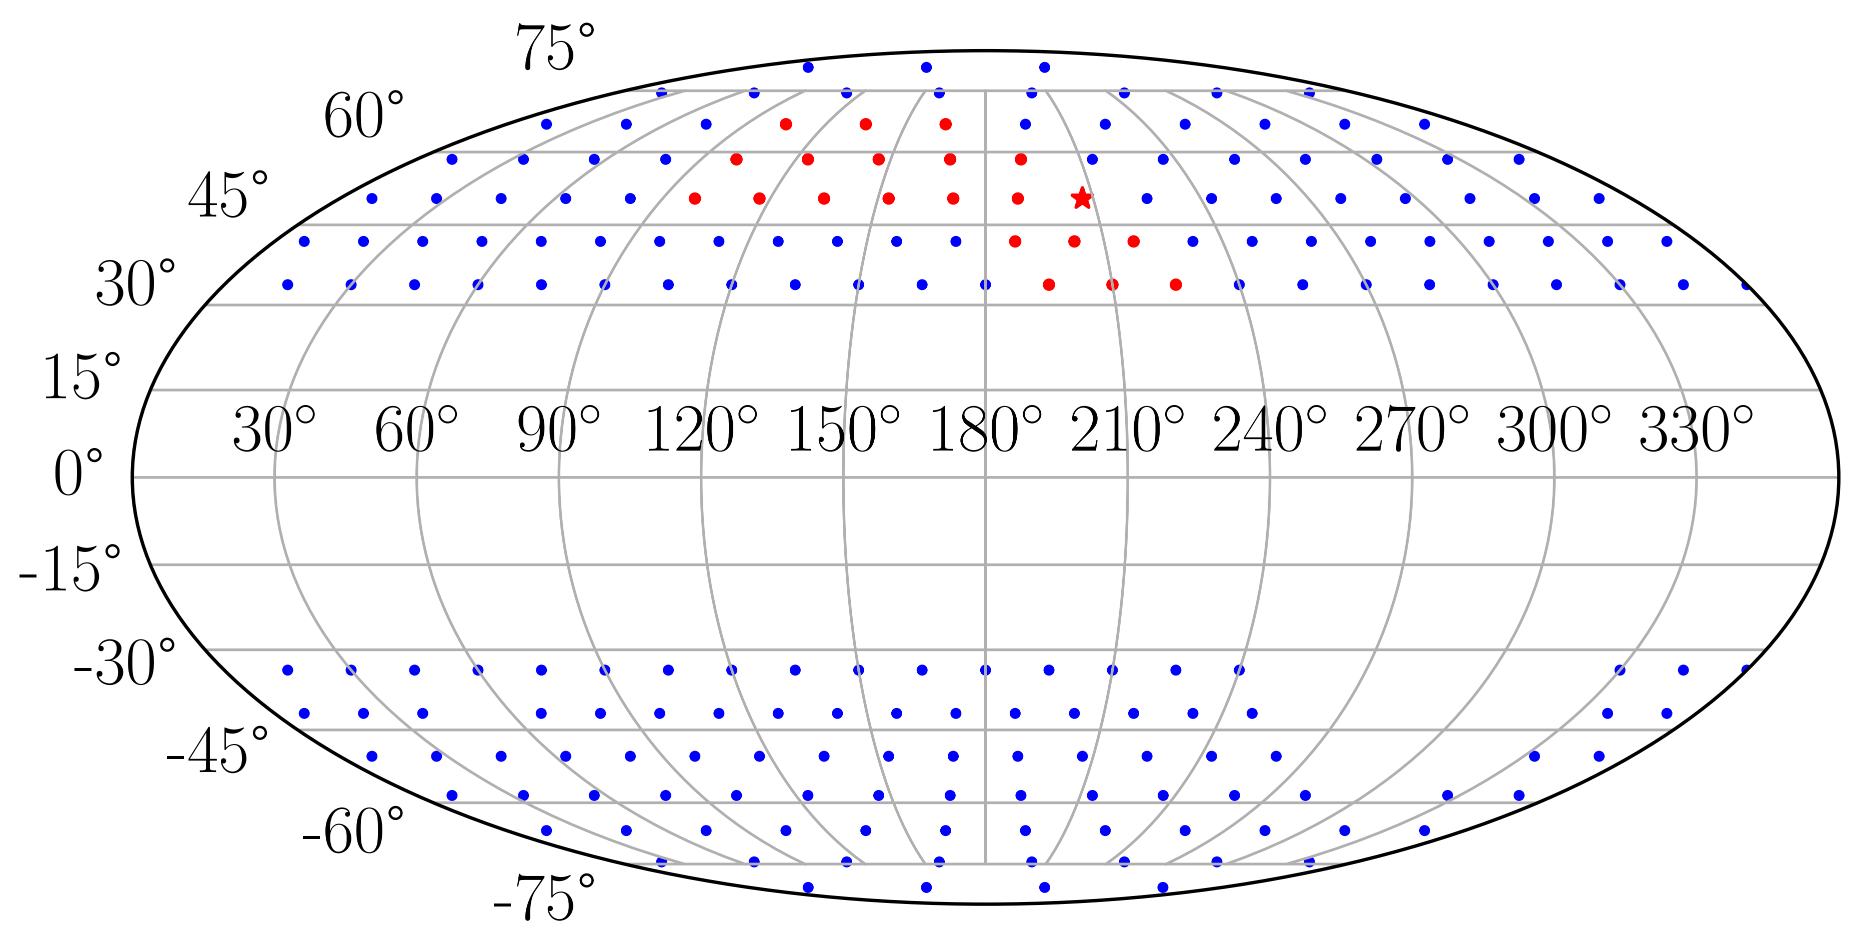

- So the method works!
- How do we apply this to the entire sky?
  - Divide sky up into overlapping 15° circular patches

ANODE on Gaia data

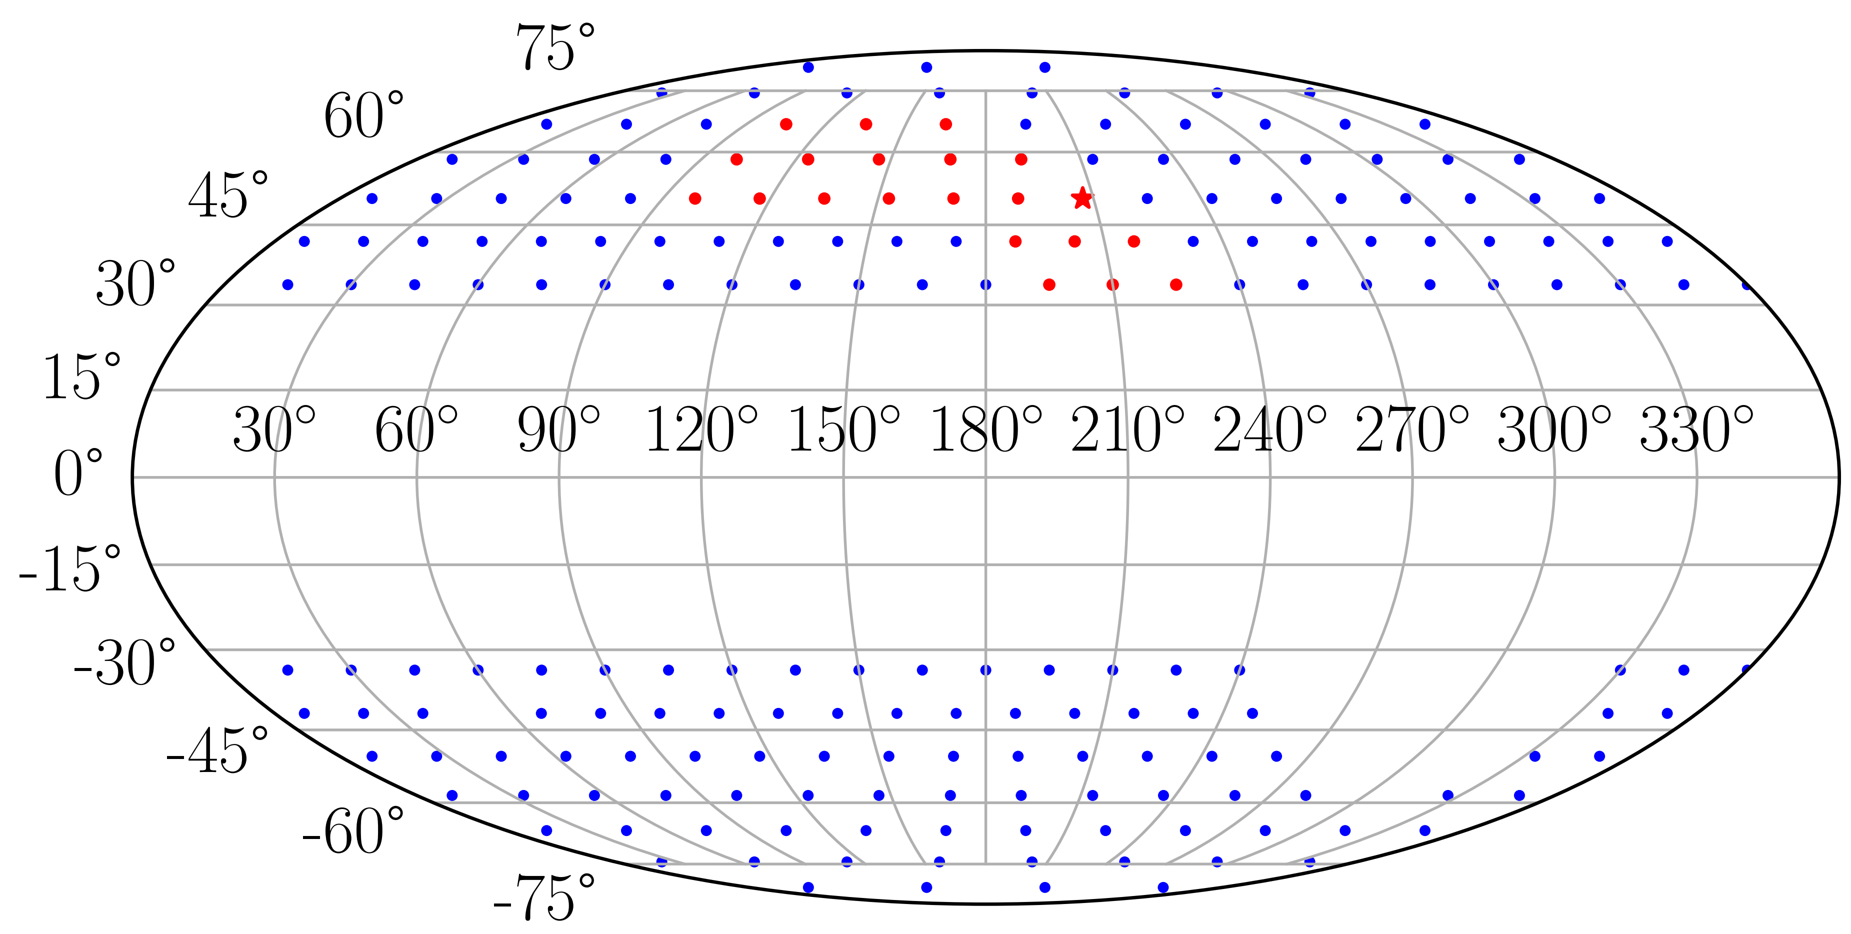

- So the method works!
- How do we apply this to the entire sky?
  - Divide sky up into overlapping 15° circular patches
  - Exclude patches too close to the disk (too many stars) and overlapping with the LMC and SMC

ANODE on Gaia data

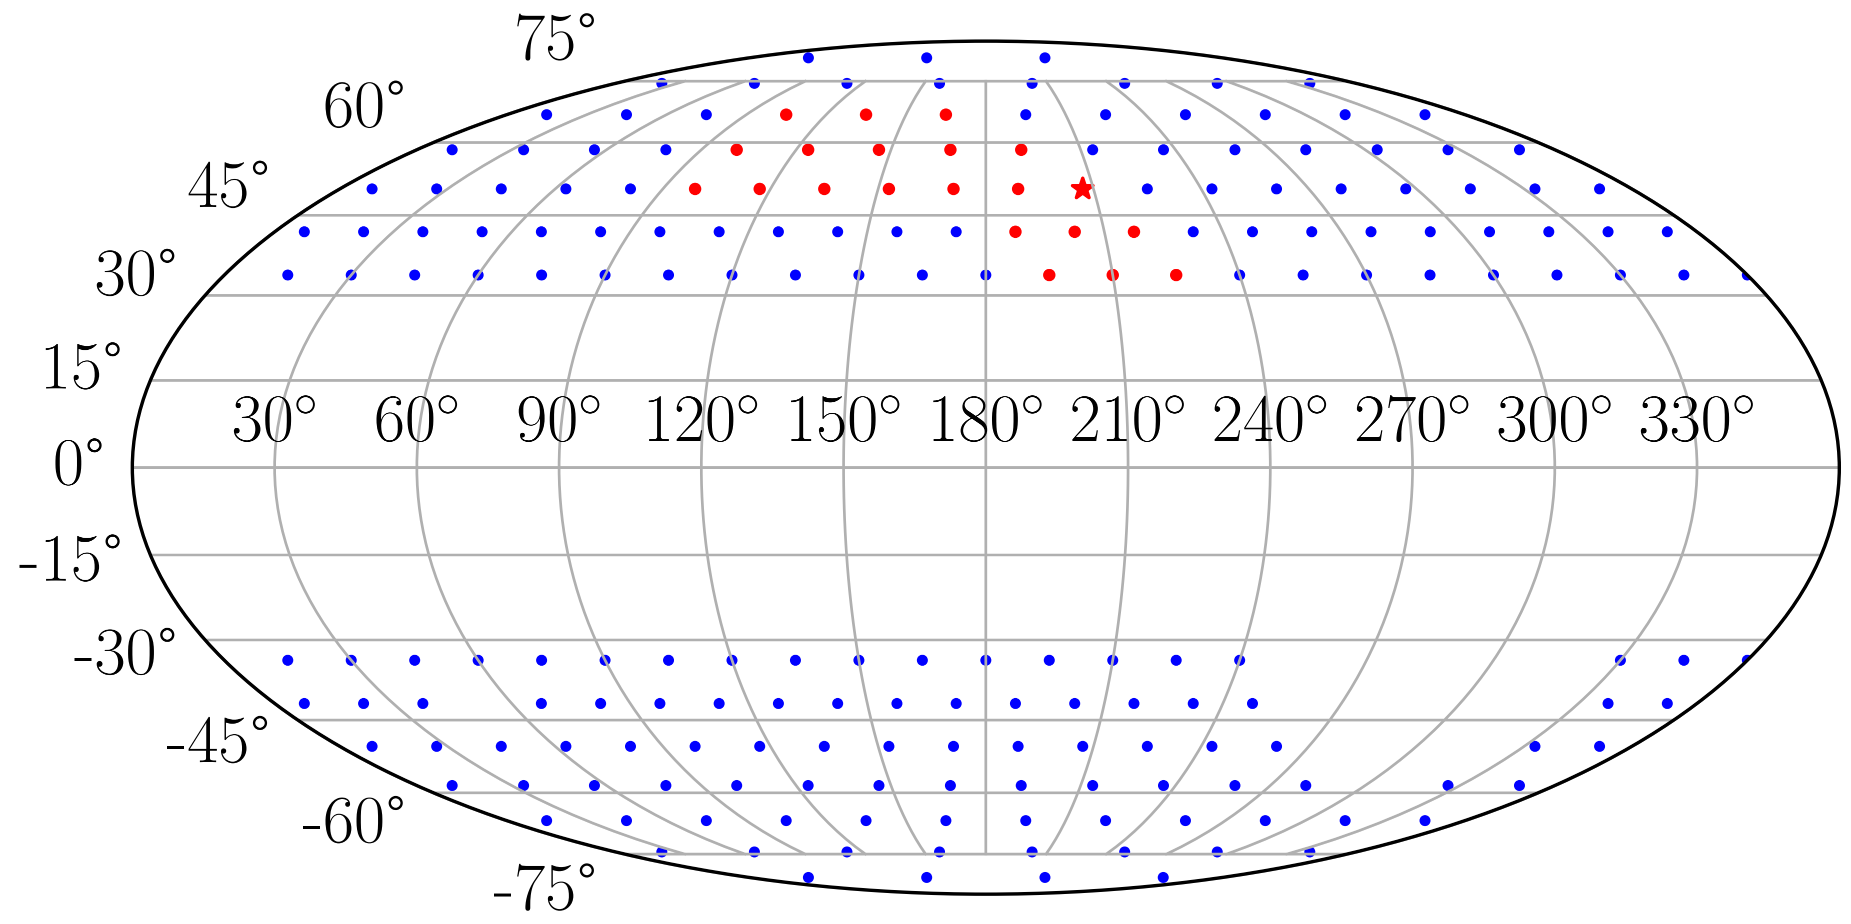

- So the method works!
- How do we apply this to the entire sky?
    - Divide sky up into overlapping 15° circular patches
    - Exclude patches too close to the disk (too many stars) and overlapping with the LMC and SMC

→ 200 patches

ANODE on Gaia data

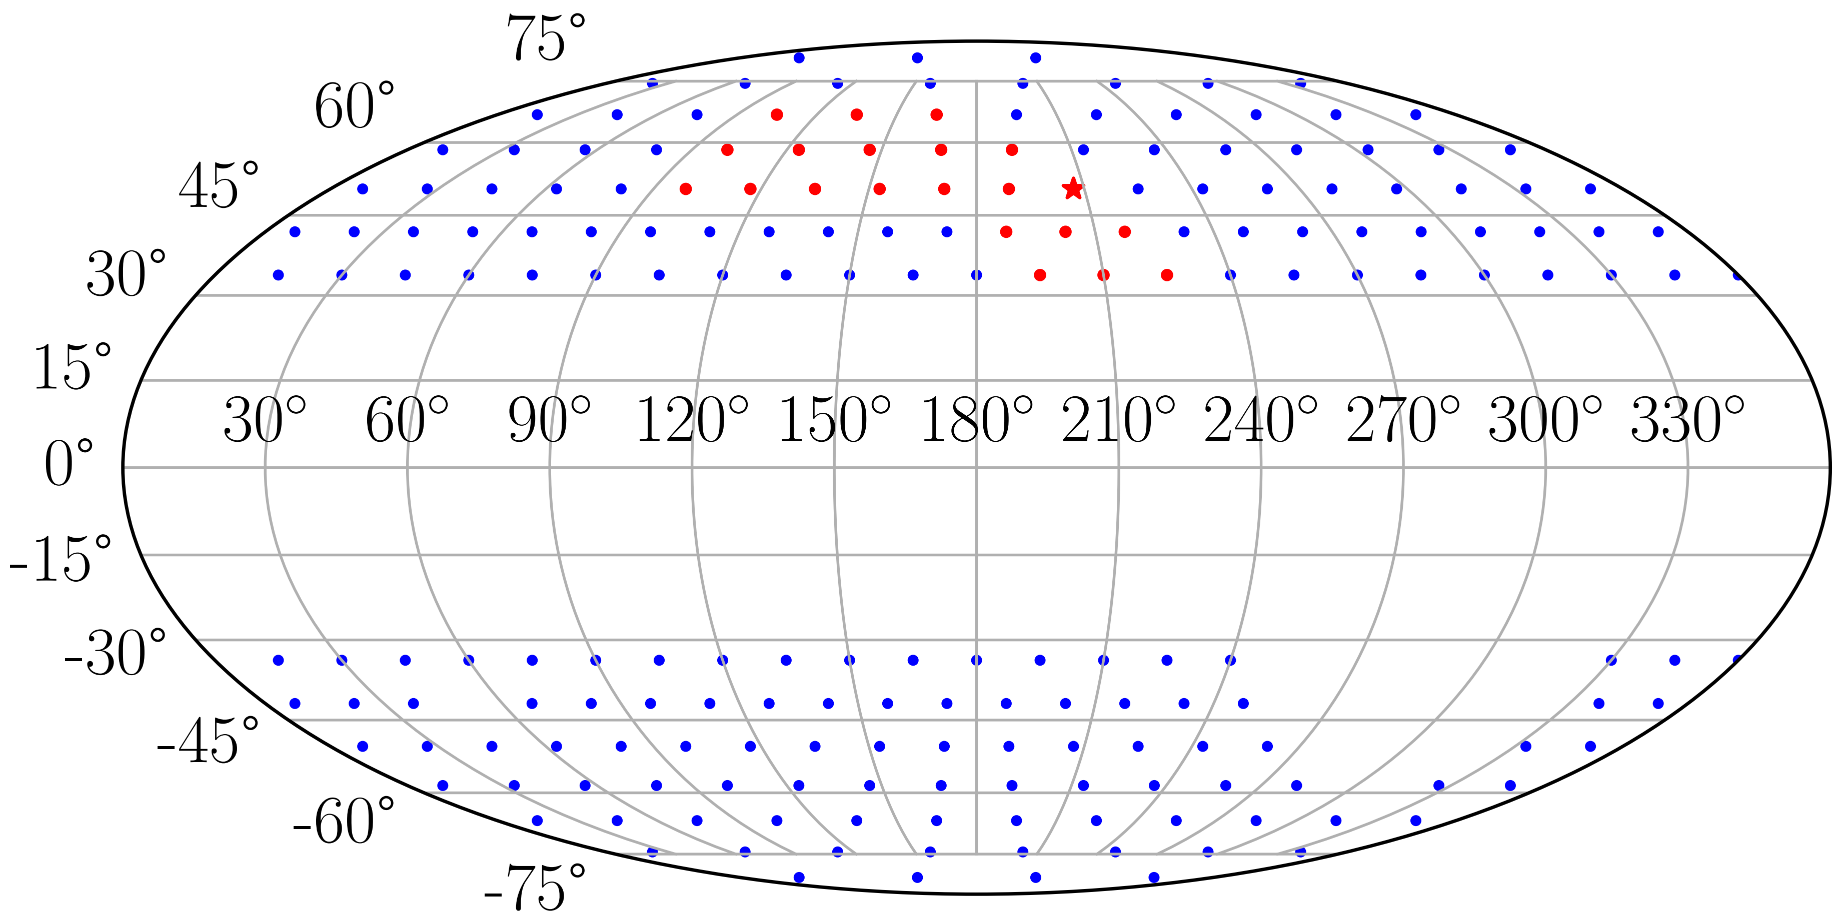

- So the method works!
- How do we apply this to the entire sky?
  - Divide sky up into overlapping 15˚ circular patches
  - Exclude patches too close to the disk (too many stars) and overlapping with the LMC and SMC

➡ 200 patches

  - Divide up each patch into overlapping SRs (width 6) in μ_α, train ANODE

ANODE on Gaia data

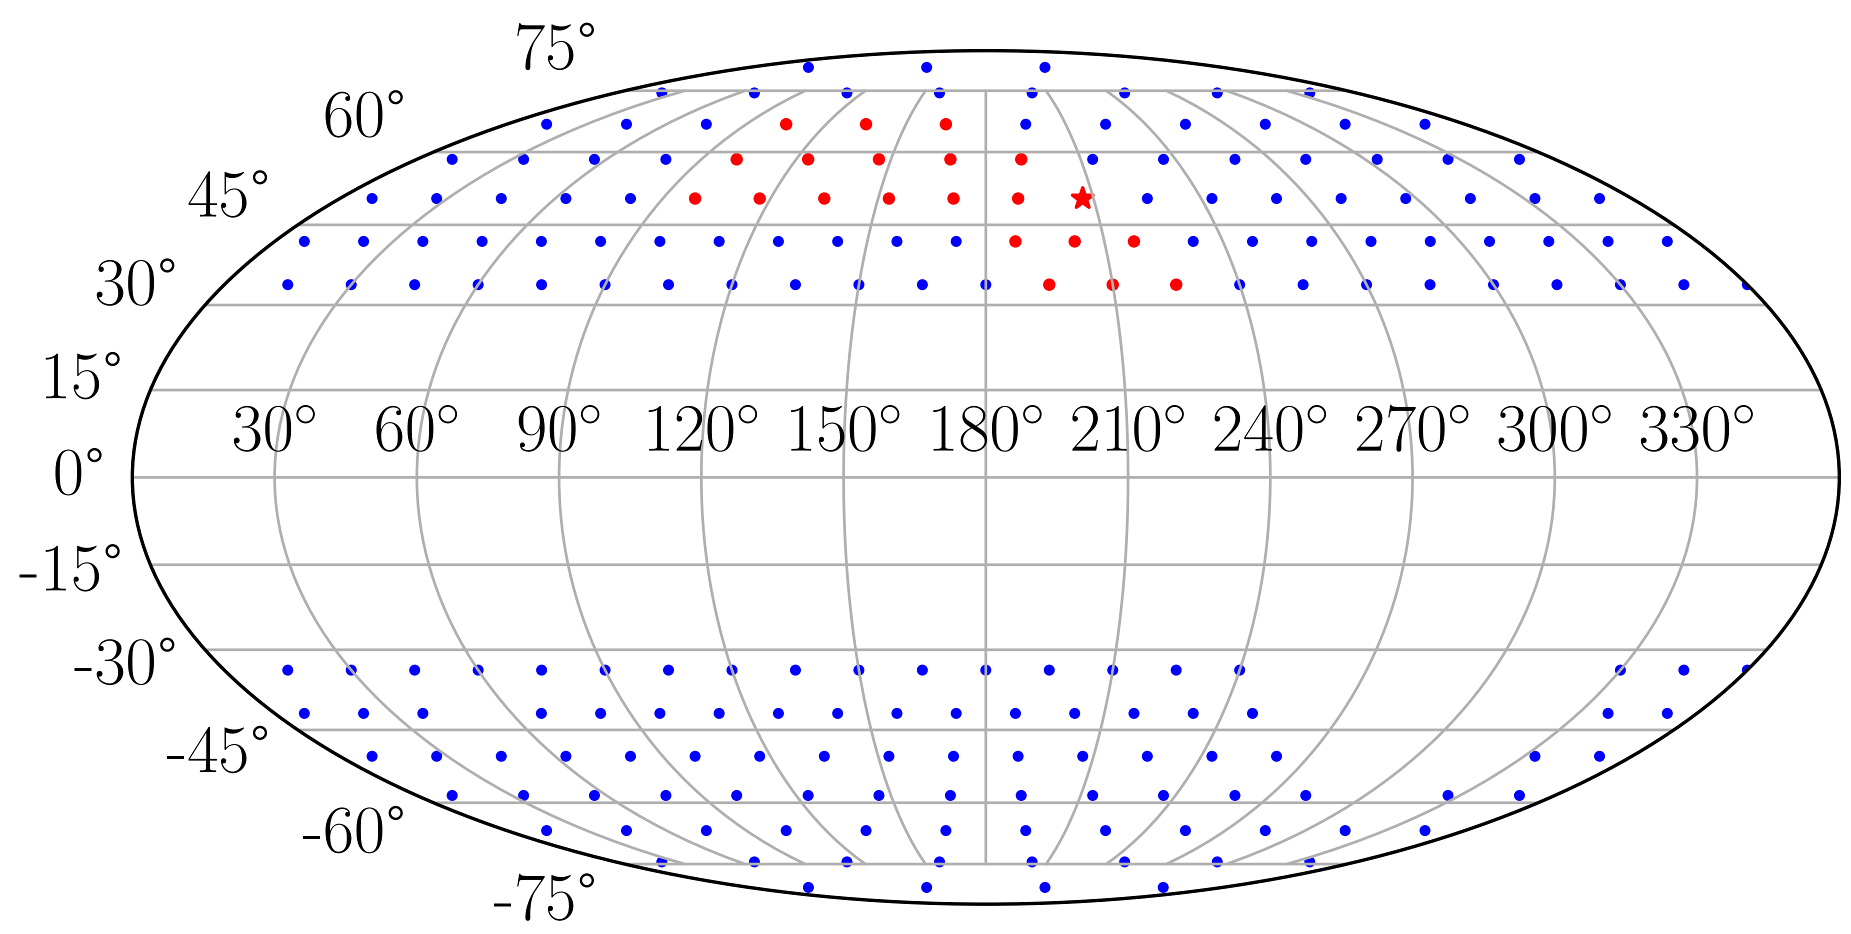

- So the method works!
- How do we apply this to the entire sky?
    - Divide sky up into overlapping 15° circular patches
    - Exclude patches too close to the disk (too many stars) and overlapping with the LMC and SMC

➡ 200 patches

- Divide up each patch into overlapping SRs (width 6) in μ_α, train ANODE

➡ ~30 SRs / patches

ANODE on Gaia data

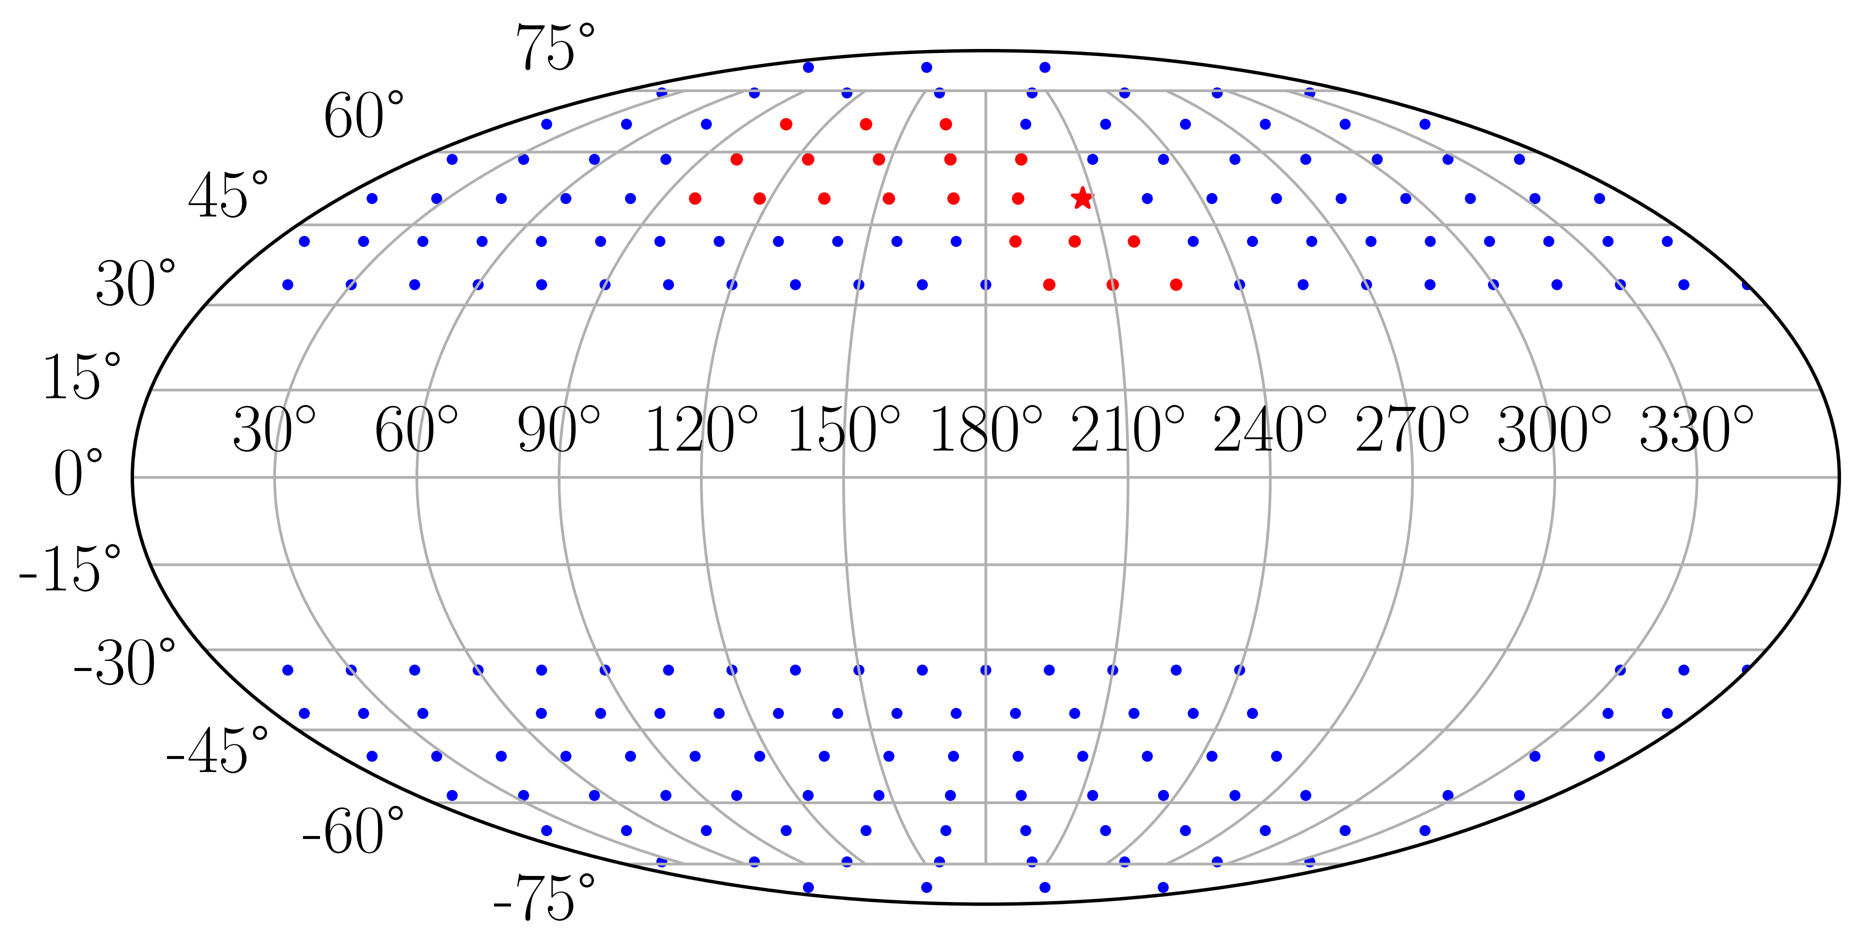

- So the method works!
- How do we apply this to the entire sky?
  - Divide sky up into overlapping 15˚ circular patches
  - Exclude patches too close to the disk (too many stars) and overlapping with the LMC and SMC

200 patches

  - Divide up each patch into overlapping SRs (width 6) in μ_α, train ANODE

~30 SRs / patches

  - Repeat for SRs in orthogonal proper motion coordinate μ_δ

ANODE on Gaia data

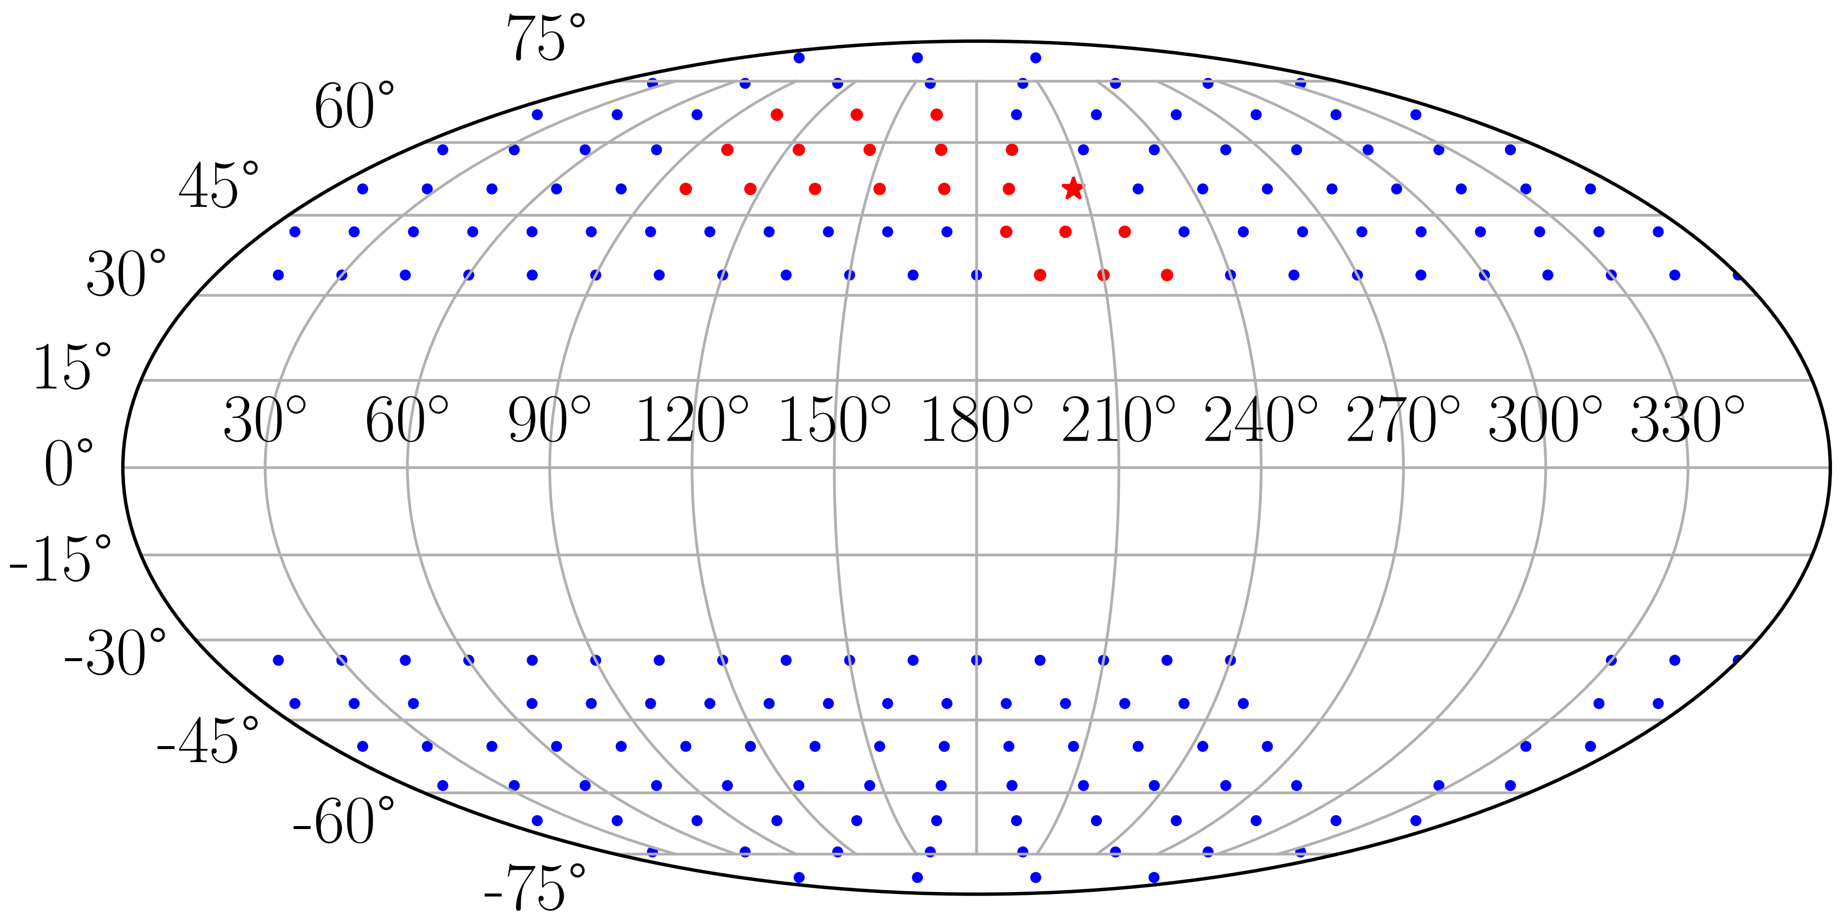

- So the method works!
- How do we apply this to the entire sky?
  - Divide sky up into overlapping 15˚ circular patches
  - Exclude patches too close to the disk (too many stars) and overlapping with the LMC and SMC

→ 200 patches

  - Divide up each patch into overlapping SRs (width 6) in μ_α, train ANODE → ~30 SRs / patches
  - Repeat for SRs in orthogonal proper motion coordinate μ_δ → x 2

ANODE on Gaia data

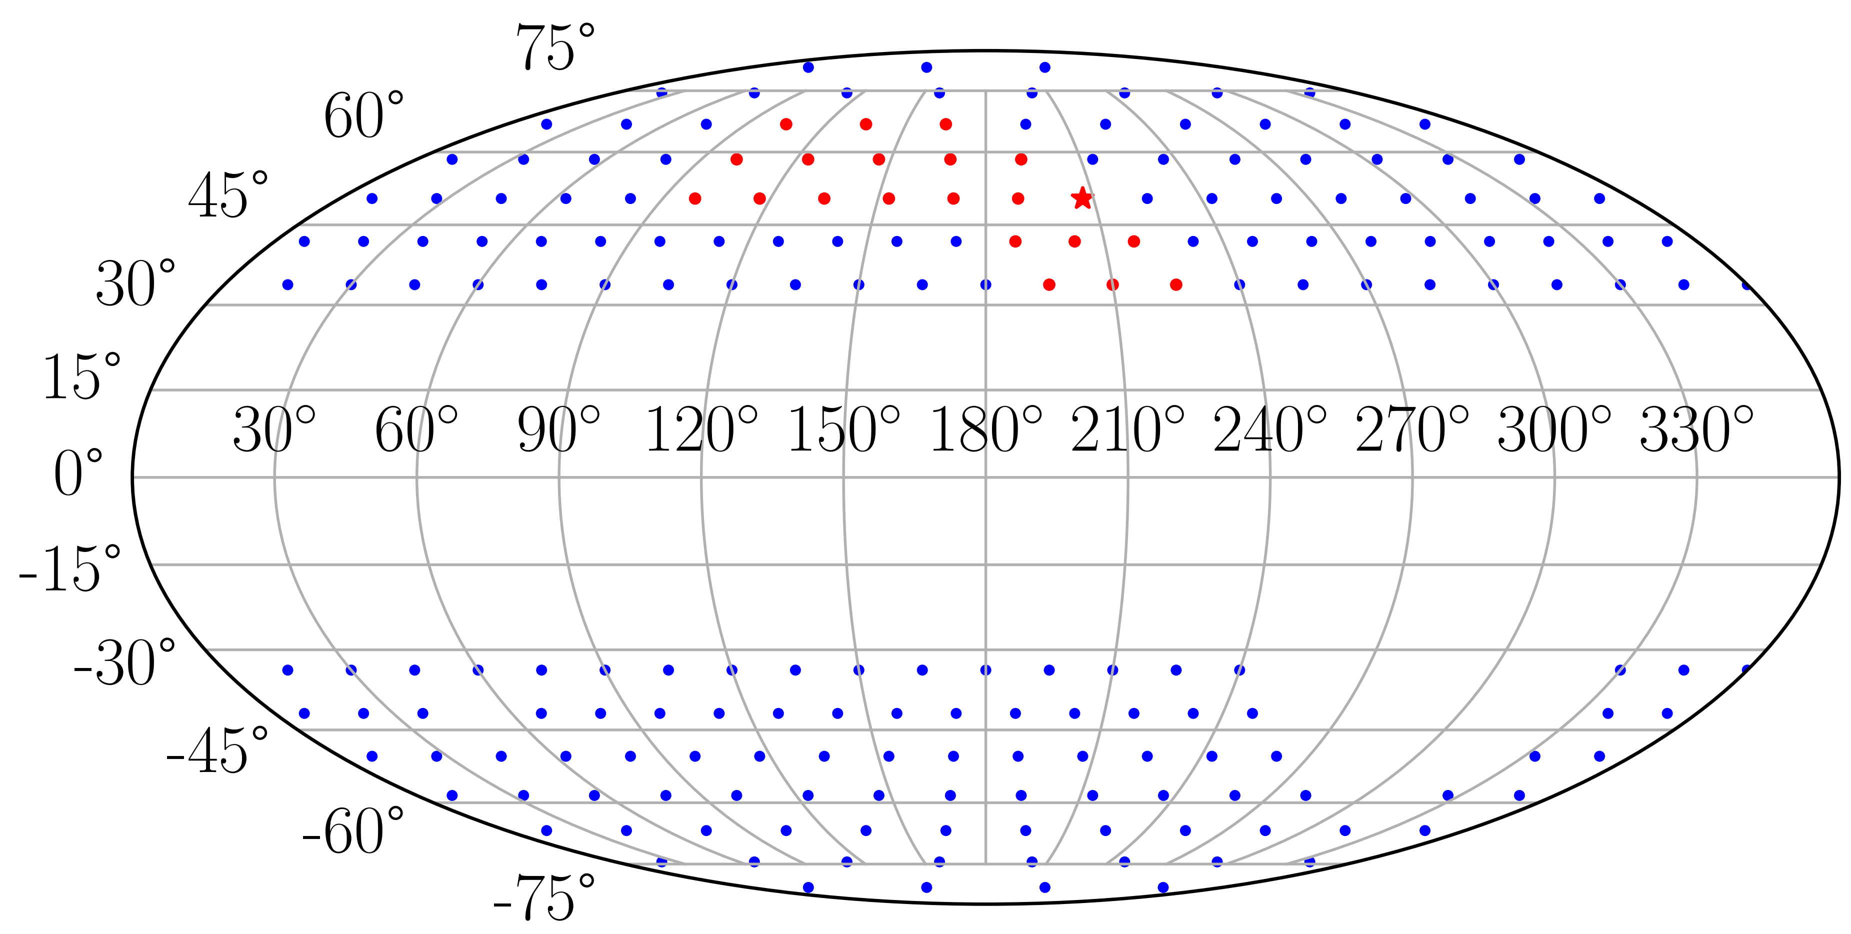

- So the method works!
- How do we apply this to the entire sky?
  - Divide sky up into overlapping 15° circular patches
  - Exclude patches too close to the disk (too many stars) and overlapping with the LMC and SMC

→ 200 patches

  - Divide up each patch into overlapping SRs (width 6) in μ_α, train ANODE → ~30 SRs / patches
  - Repeat for SRs in orthogonal proper motion coordinate μ_δ → x 2

Need to train ANODE ~12,000 times to cover the entire sky!!
Each training takes O(10h)

ANODE on Gaia data

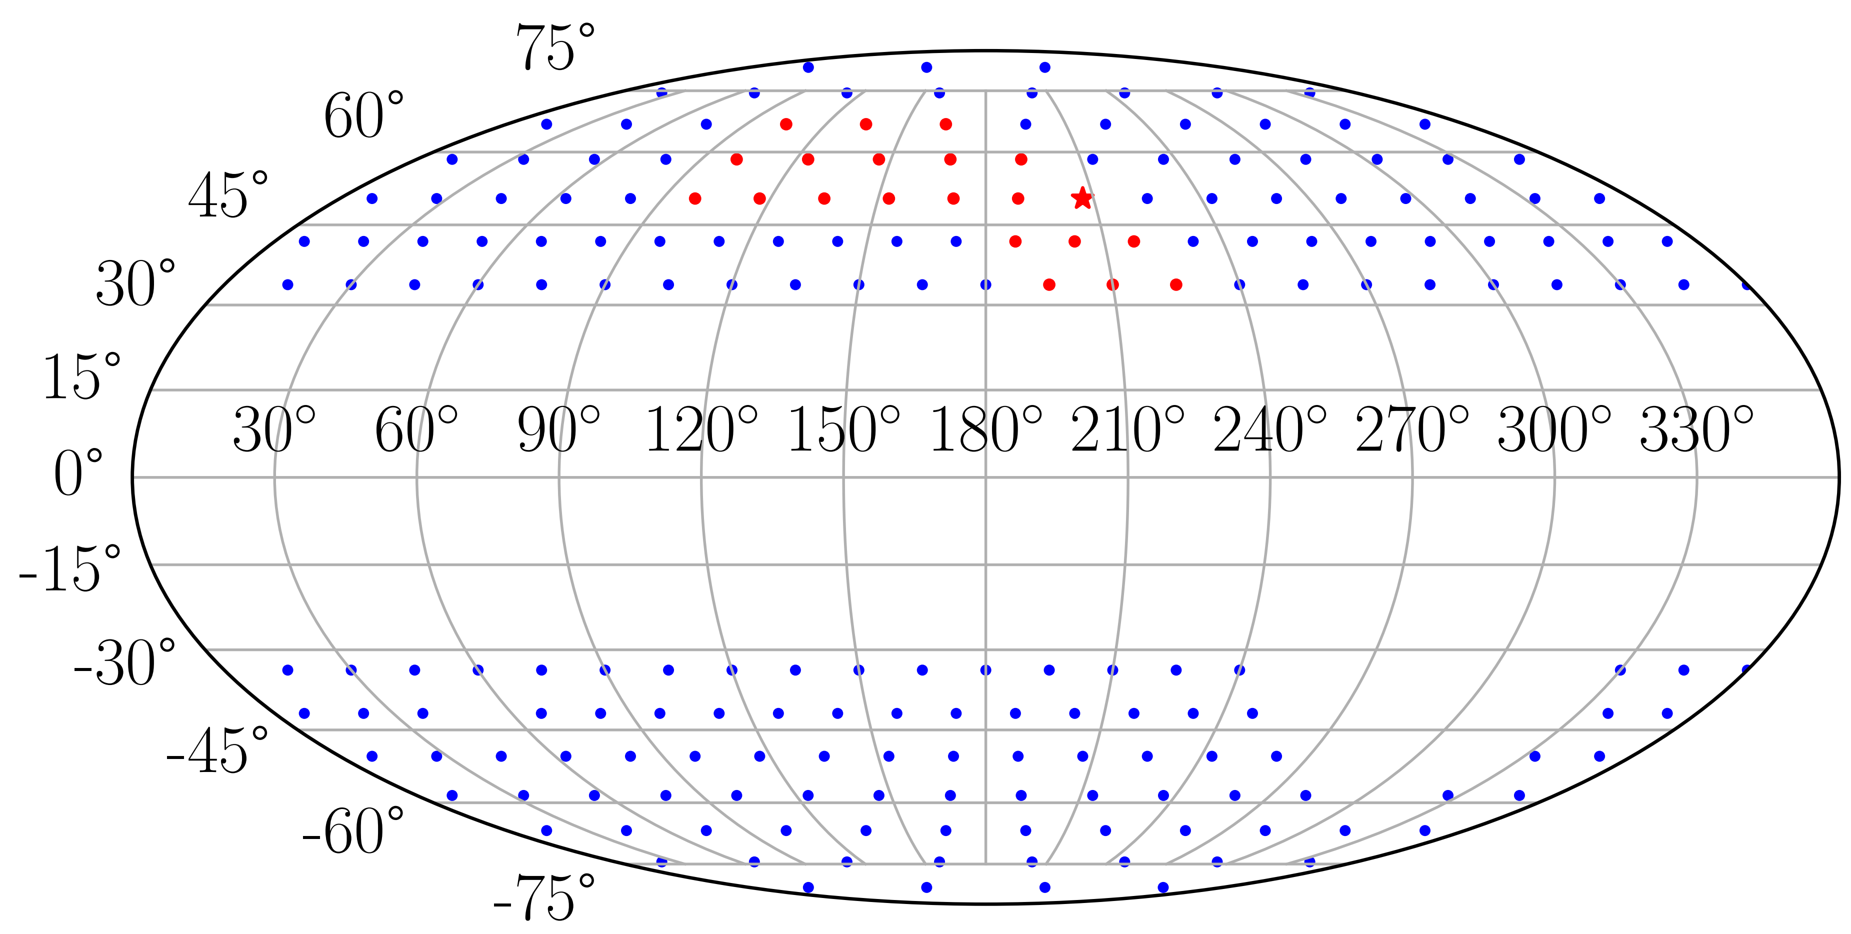

- So the method works!
- How do we apply this to the entire sky?
  - Divide sky up into overlapping 15° circular patches
  - Exclude patches too close to the disk (too many stars) and overlapping with the LMC and SMC

→ 200 patches

  - Divide up each patch into overlapping SRs (width 6) in μ_α, train ANODE → ~30 SRs / patches
  - Repeat for SRs in orthogonal proper motion coordinate μ_δ → x 2

Need to train ANODE ~12,000 times to cover the entire sky!!
Each training takes O(10h)

We ran this on the NERSC supercomputer at LBNL

ANODE on Gaia data

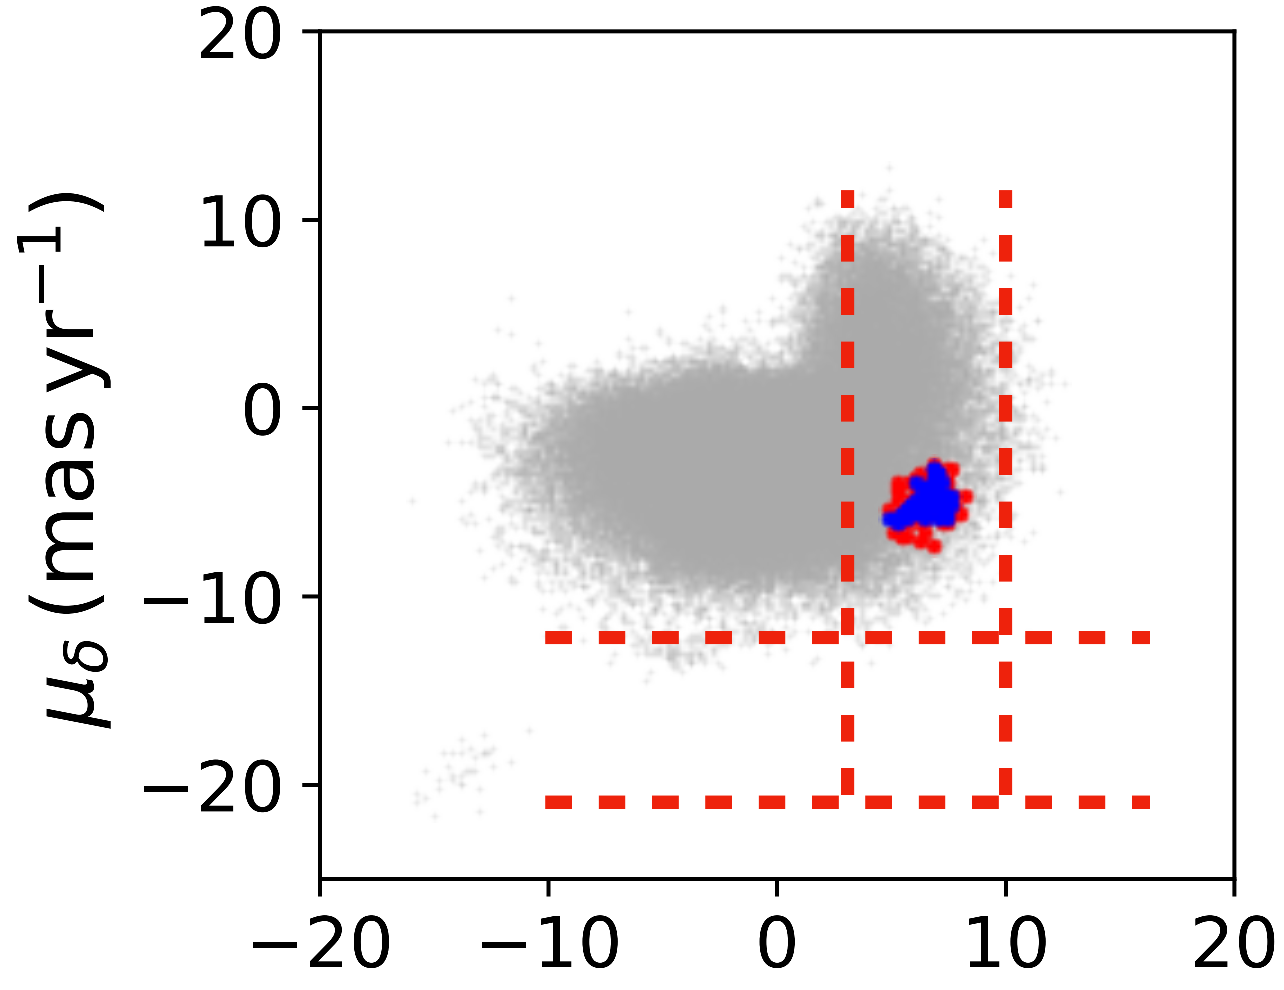

ANODE on Gaia data

- How to set the cut on R(x)?

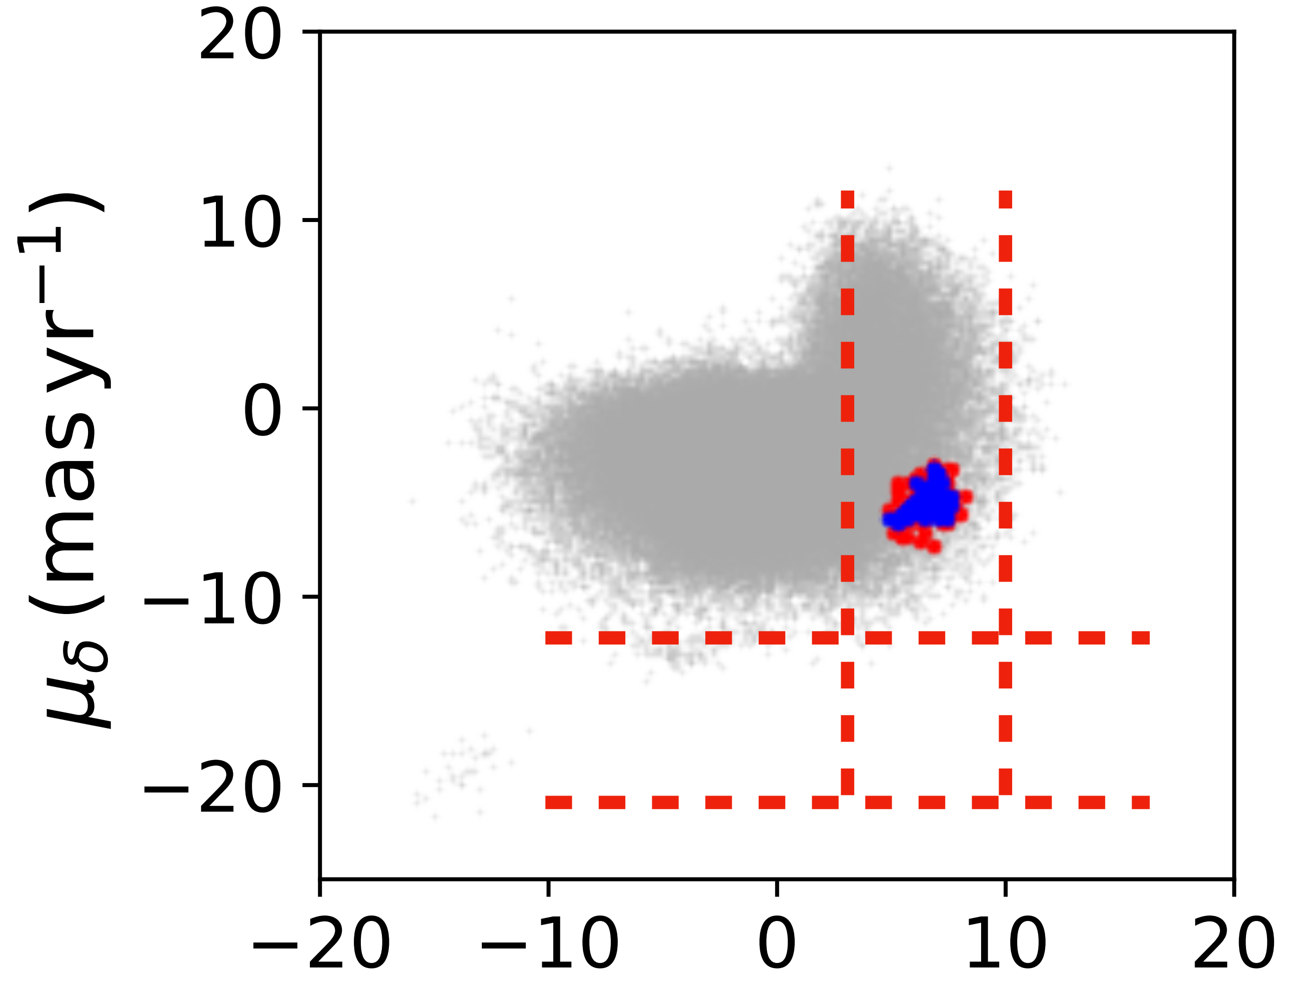

ANODE on Gaia data

- How to set the cut on R(x)?
- We found a defining a single threshold on R(x) across all SRs was insufficient to find other known streams besides GD-1.

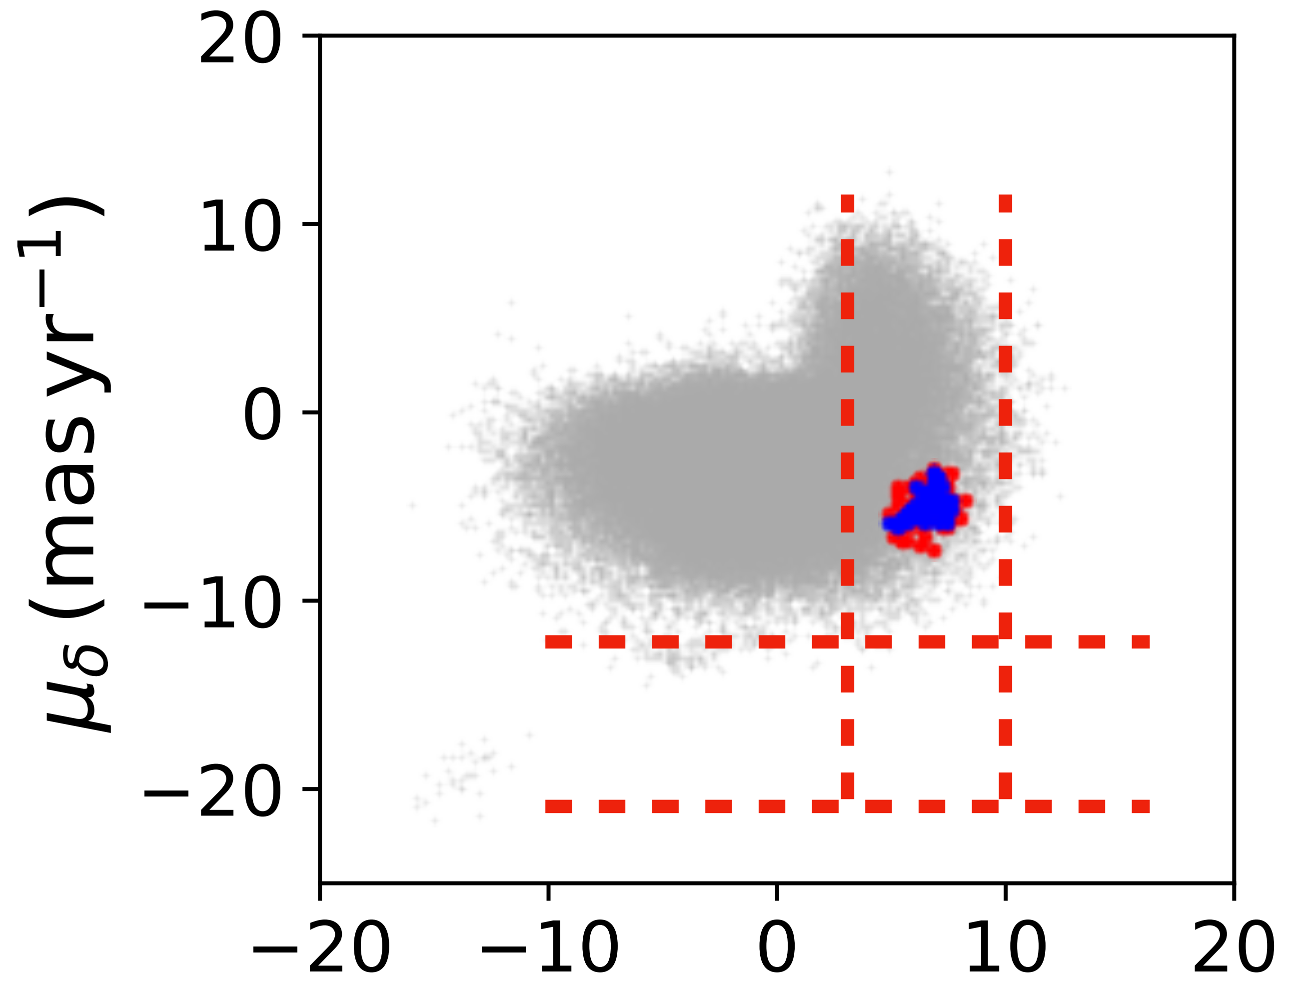

ANODE on Gaia data

- How to set the cut on R(x)?
- We found a defining a single threshold on R(x) across all SRs was insufficient to find other known streams besides GD-1.
- What worked instead was to further subdivide SRs into slices by the orthogonal proper motion => “ROIs”

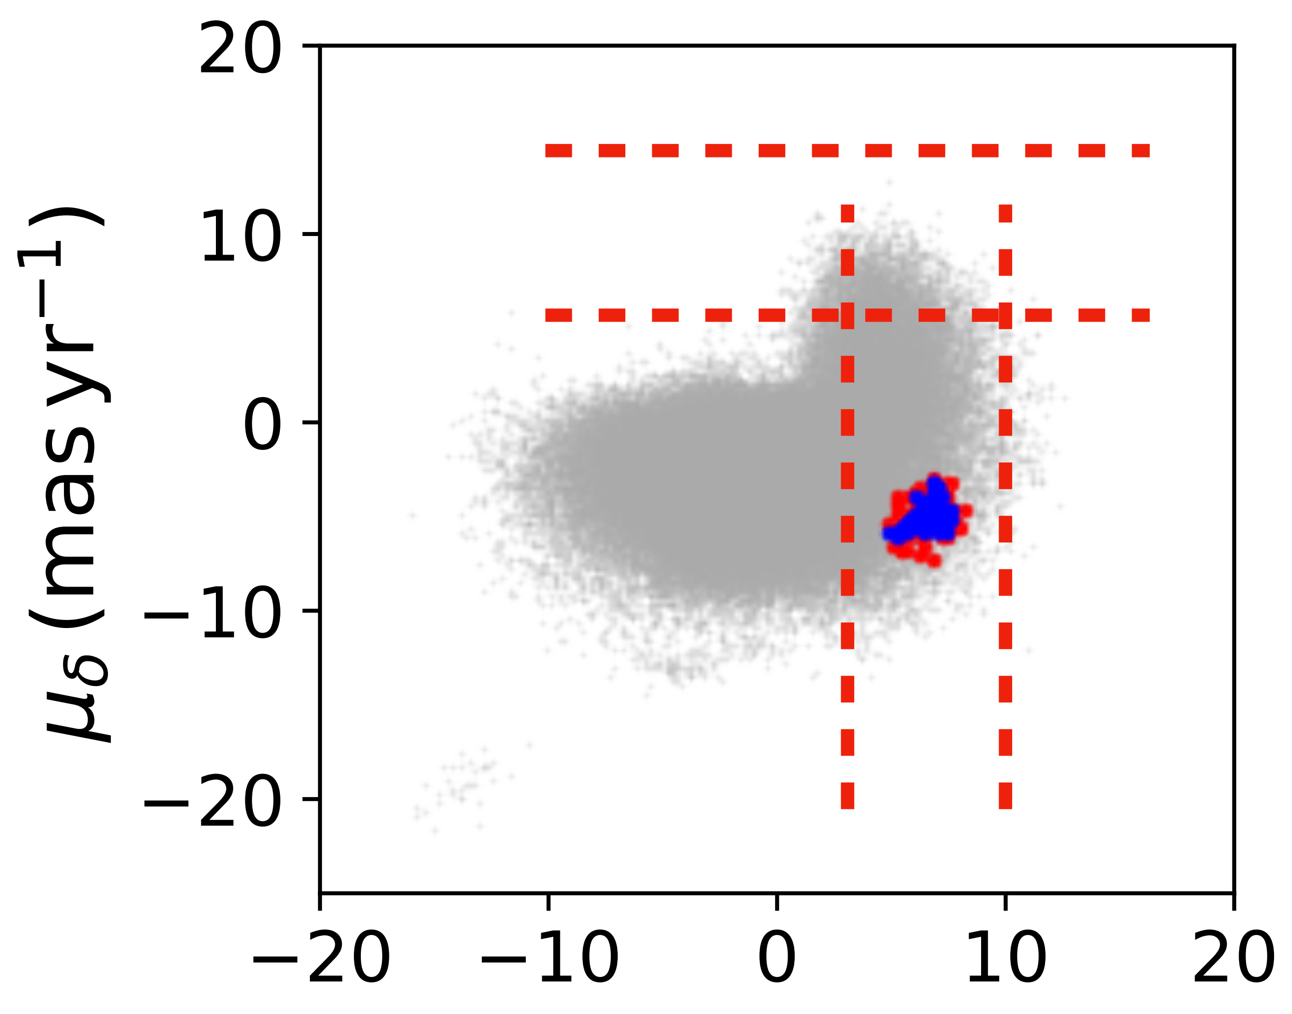

ANODE on Gaia data

- How to set the cut on R(x)?
- We found a defining a single threshold on R(x) across all SRs was insufficient to find other known streams besides GD-1.
- What worked instead was to further subdivide SRs into slices by the orthogonal proper motion => “ROIs”
- In each ROI, take the 100 highest R stars

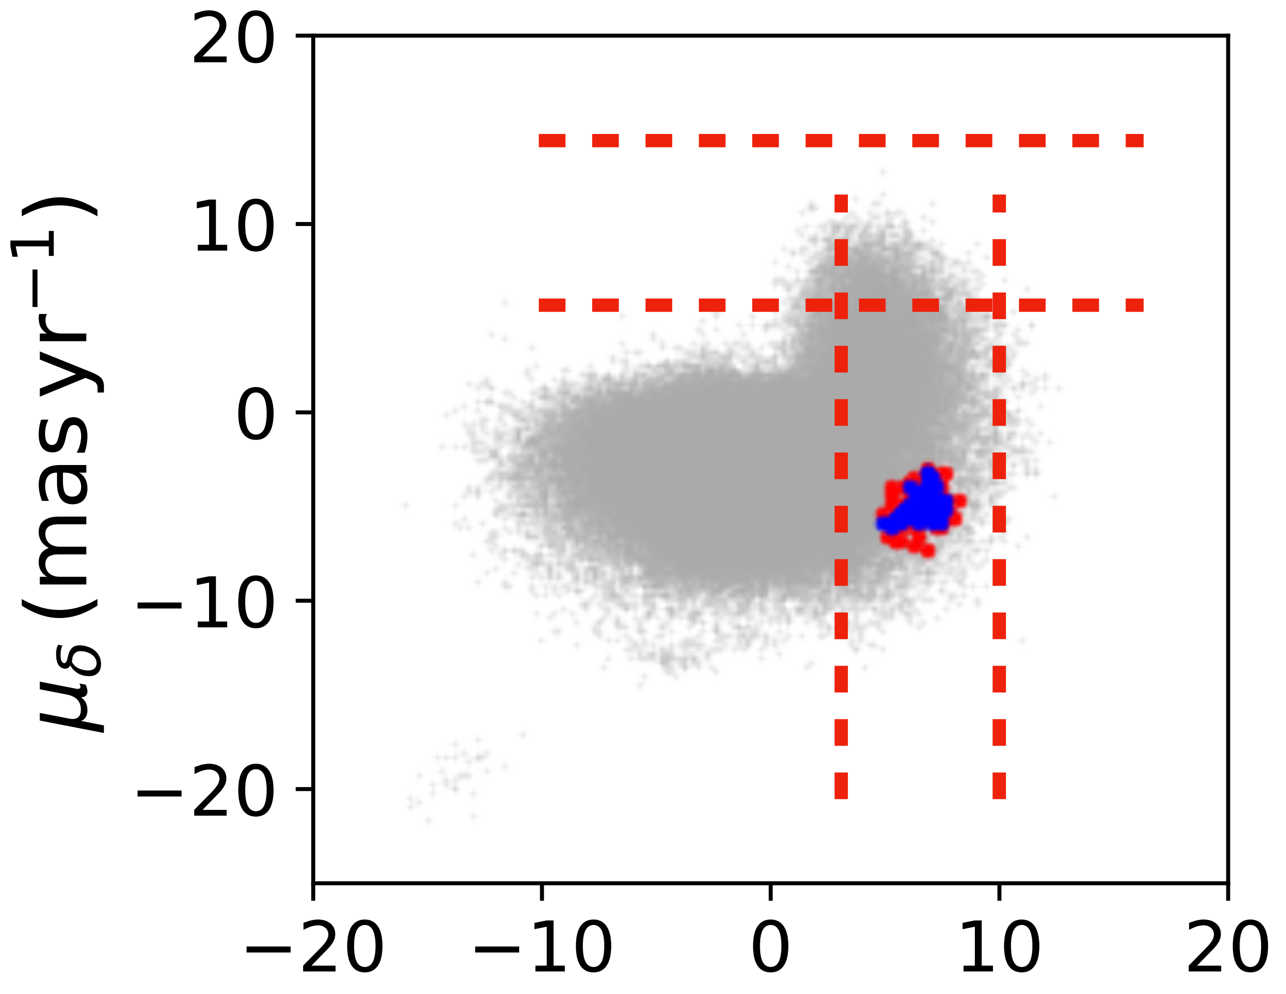

ANODE on Gaia data

- How to set the cut on R(x)?
- We found a defining a single threshold on R(x) across all SRs was insufficient to find other known streams besides GD-1.
- What worked instead was to further subdivide SRs into slices by the orthogonal proper motion => “ROIs”
- In each ROI, take the 100 highest R stars
- Increases the sensitivity to real streams, but at the cost of a bigger look elsewhere effect.

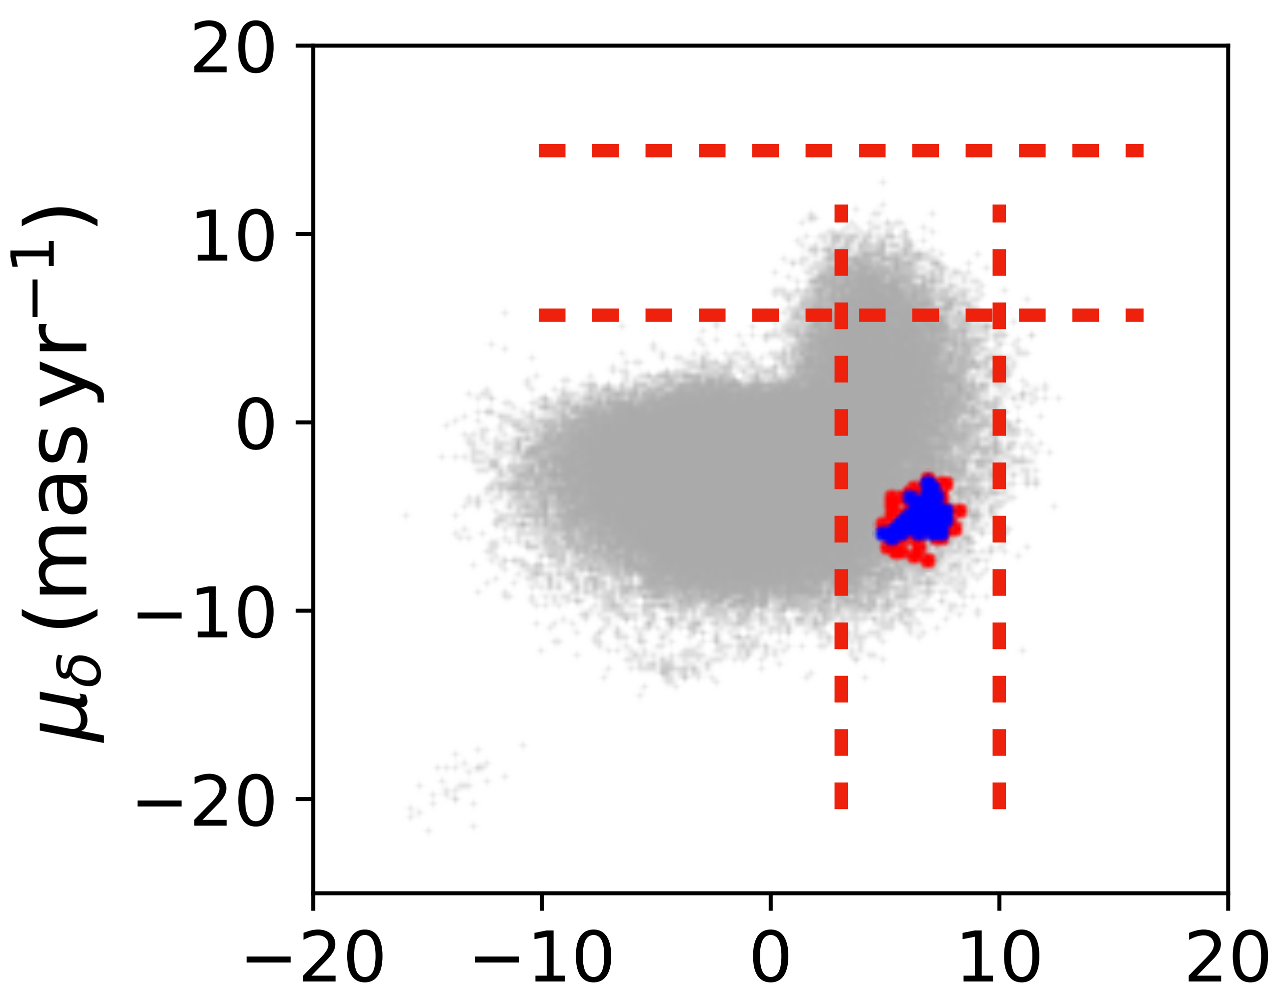

Building streams from fragments

- We end up with ~10^5 ROIs —> need an automated way to scan them for potential streams and a way to cut down on trials factor!
  - Hough transform for line finding => significance
  - Cluster together ROIs from independent runs of ANODE => build stream fragments in each patch and cut down on LEE
  - Cluster together significant stream fragments in different patches to build full stream candidate

Galaxia false positive rate

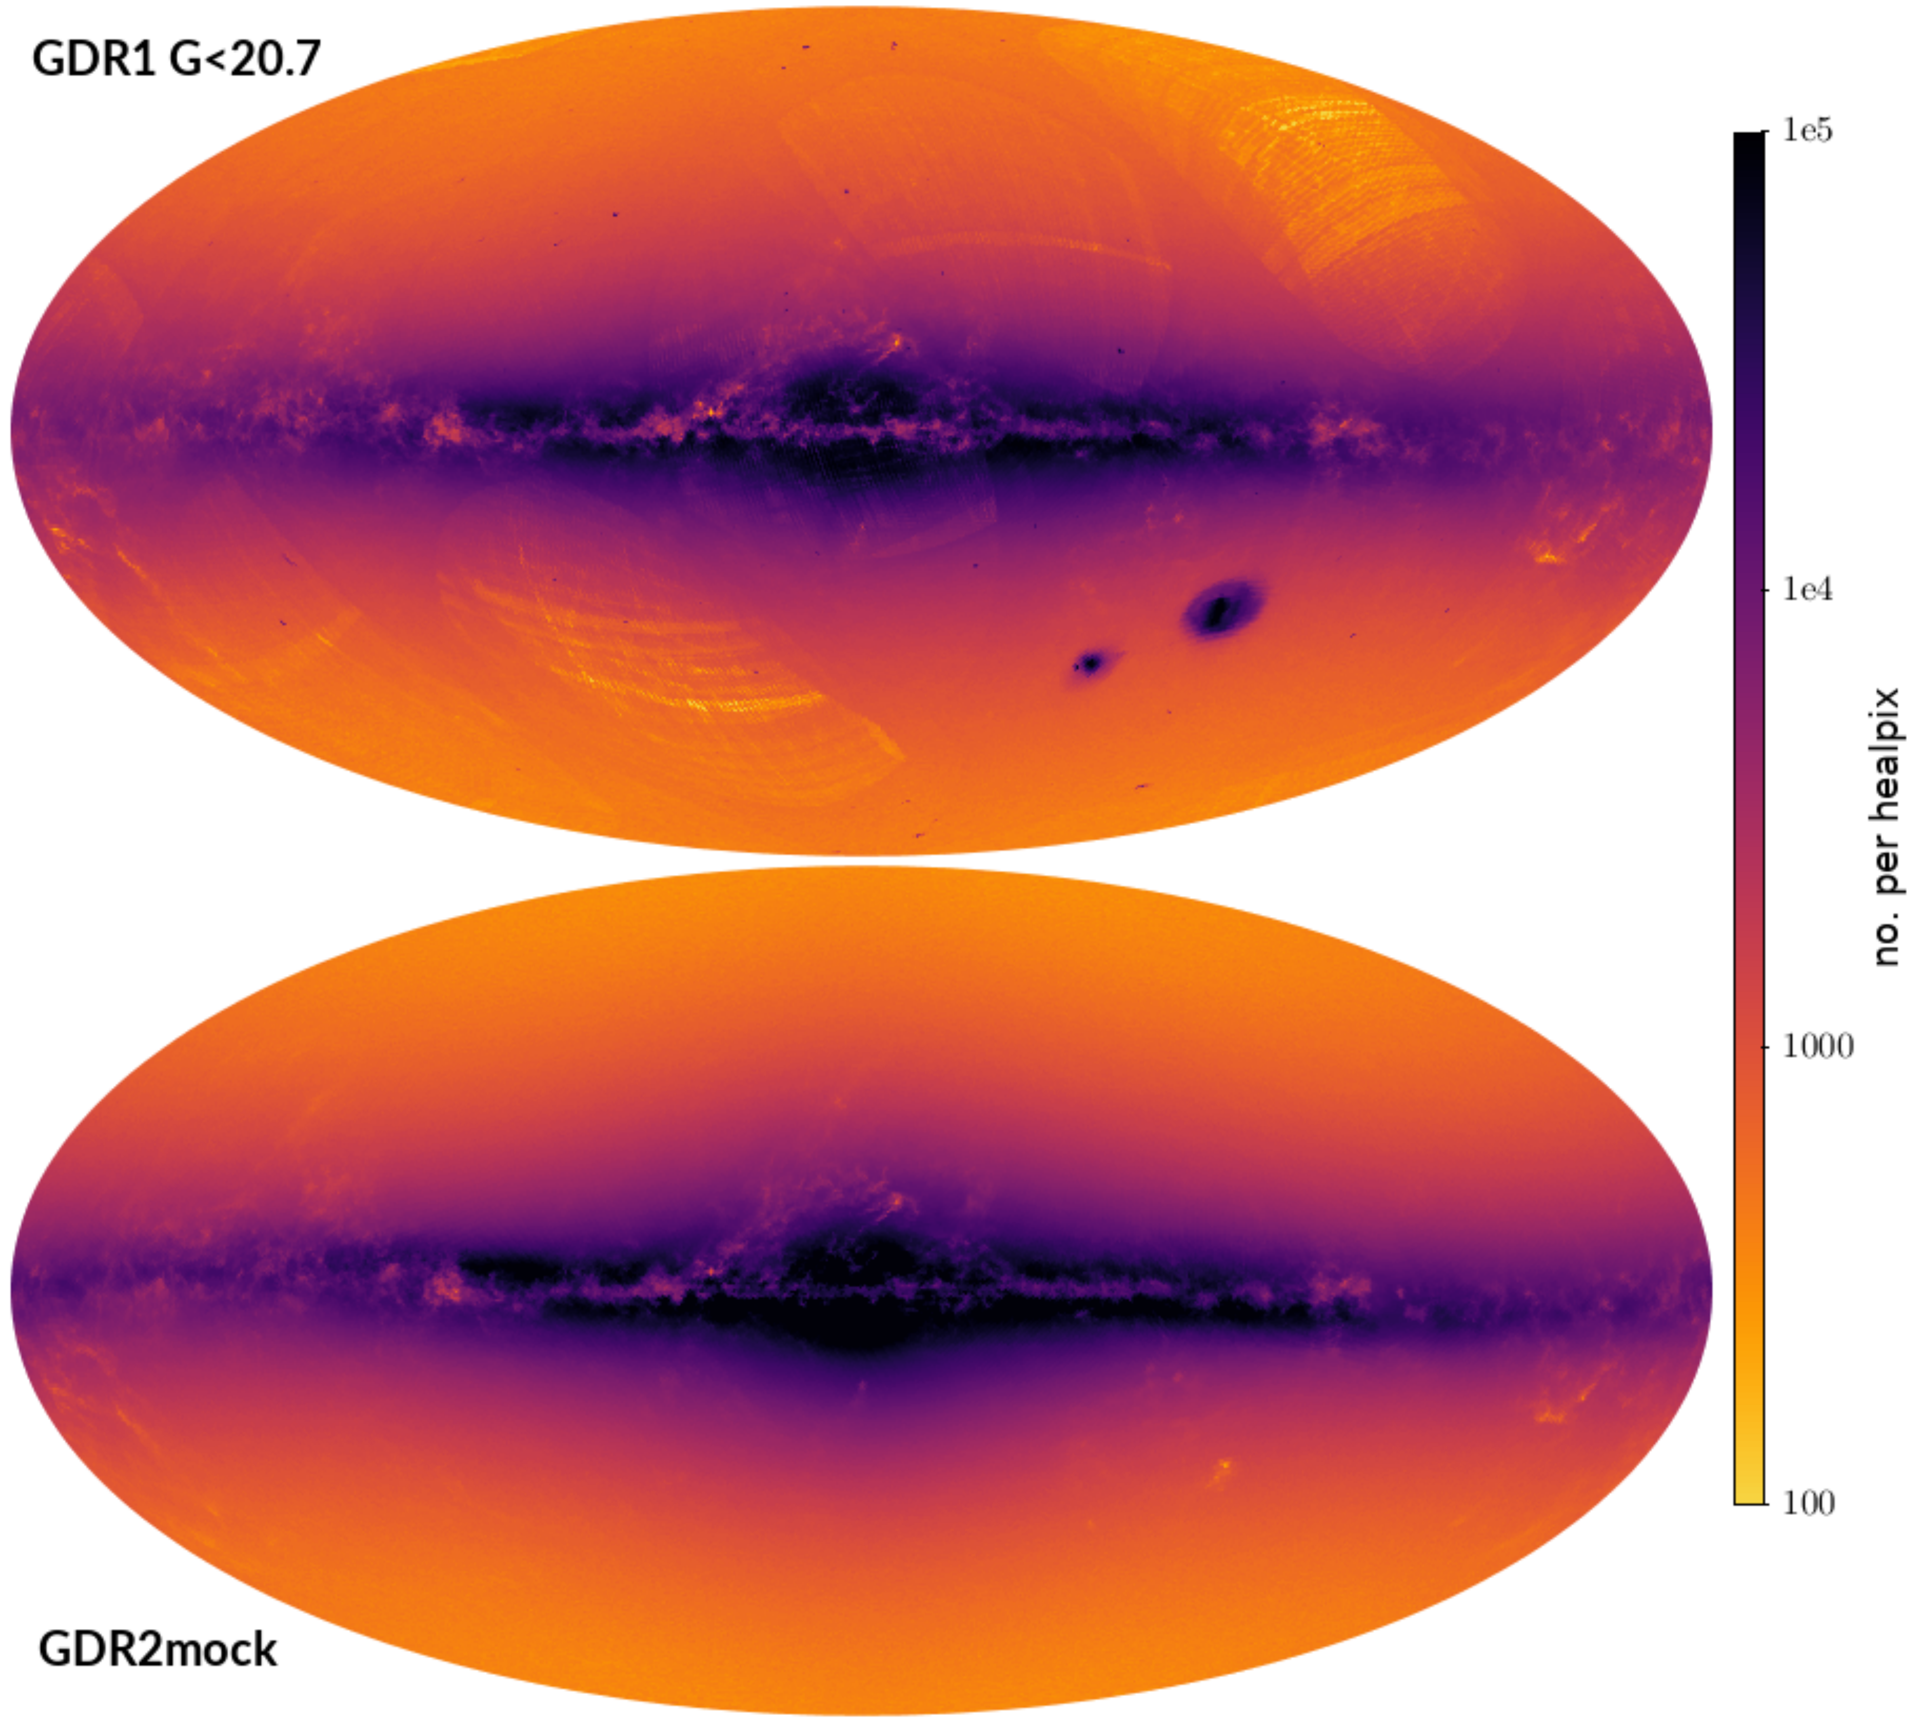

To quantify our false positive rate, we ran our full method on a semi-realistic Gaia mock catalog called Galaxia (Rybizki et al 2018) which does not have stellar streams

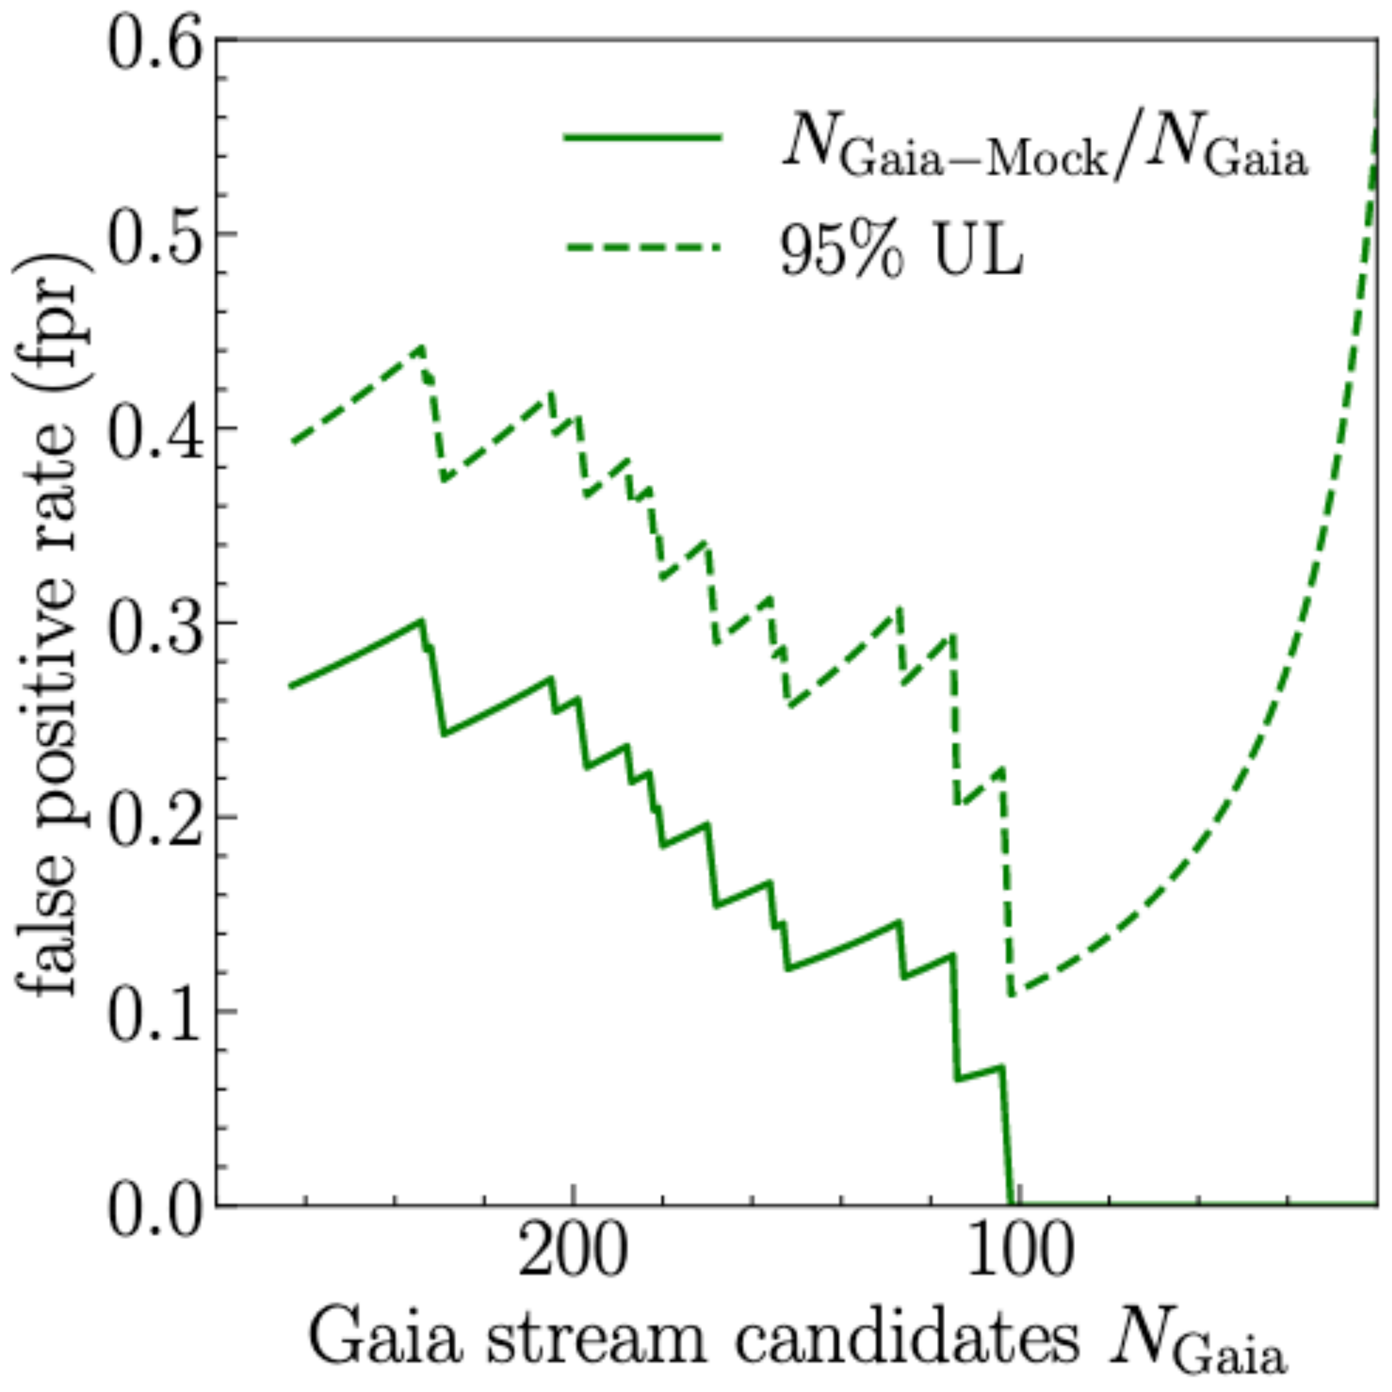

We find 100 stream candidates with a 95% UL on fpr of 11%!

Results: known streams

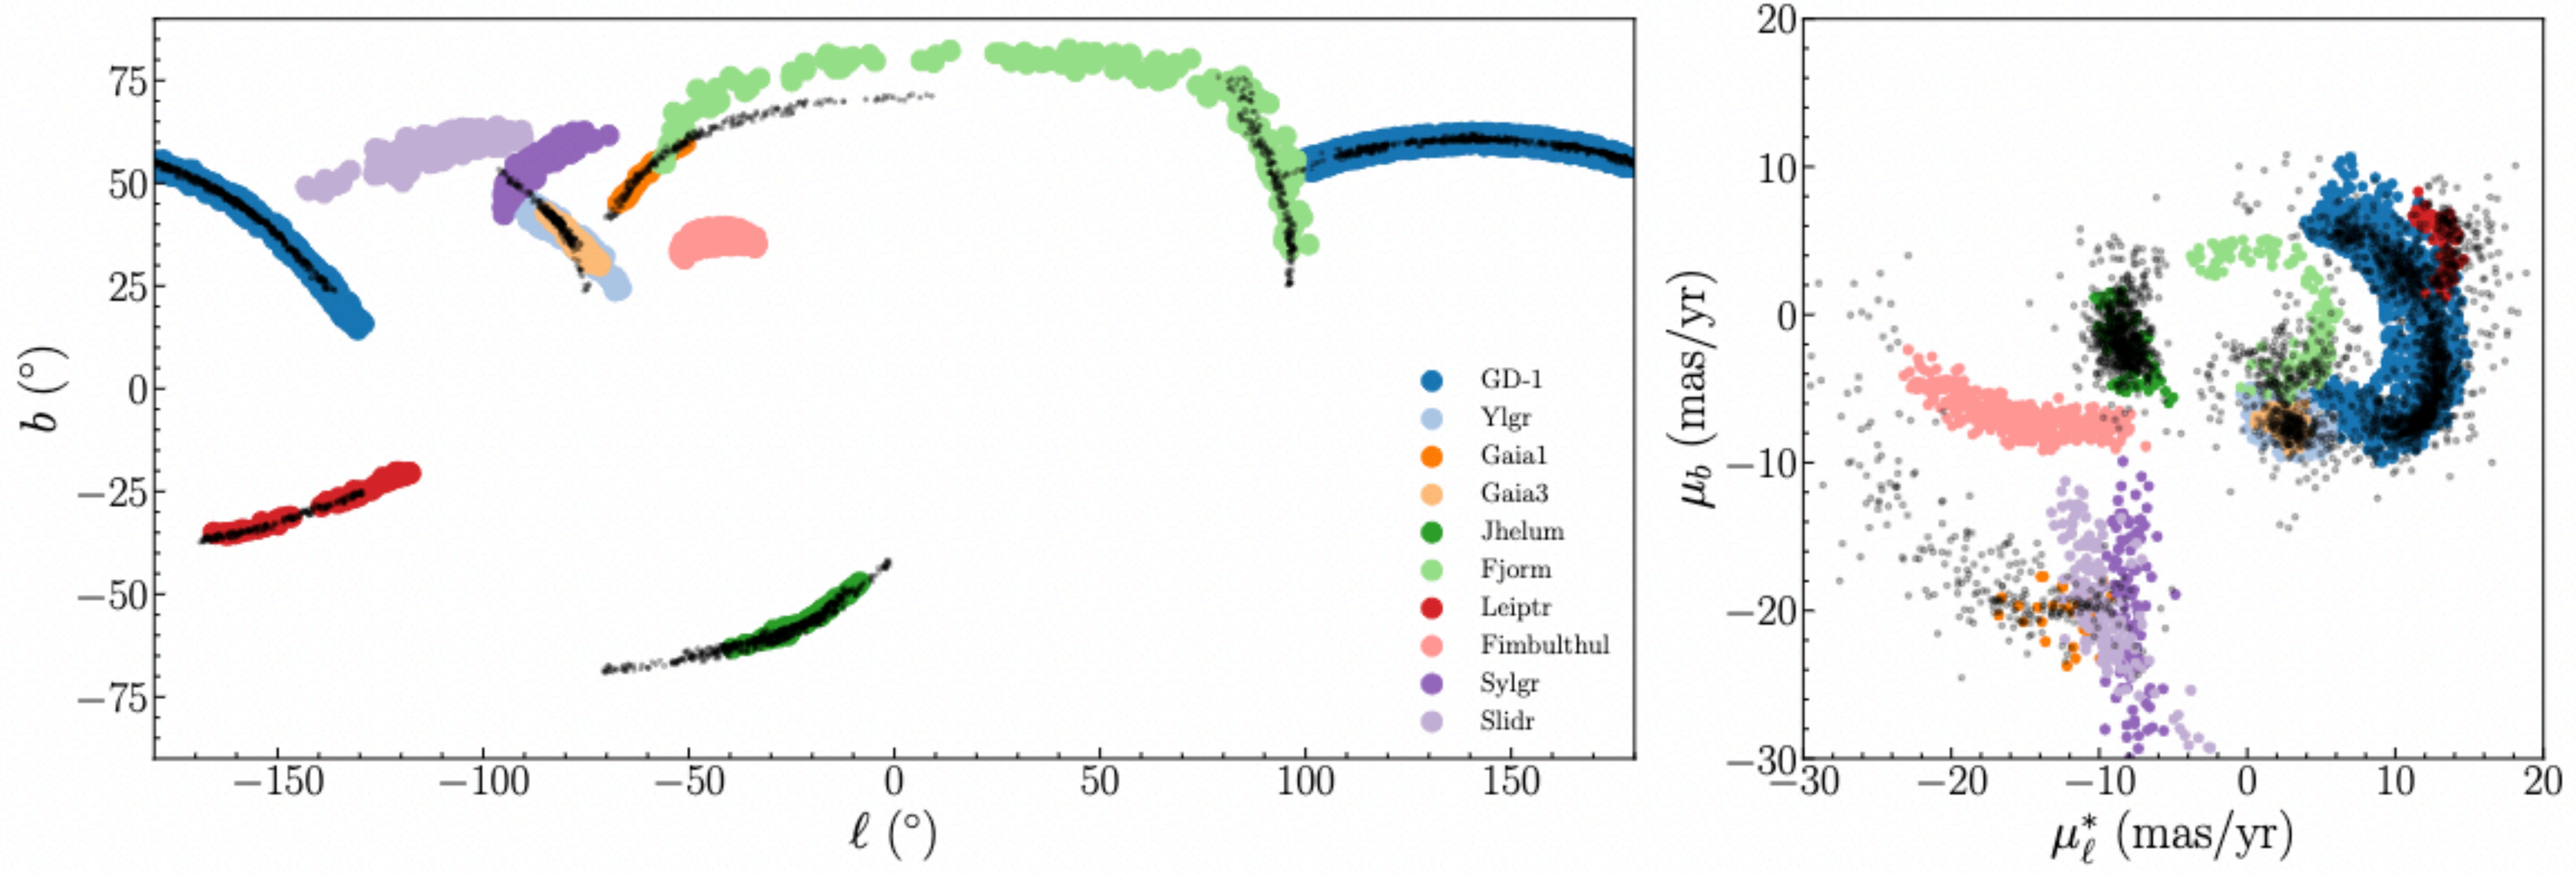

We confirm 6 previously discovered stream candidates

Others are either too wide, or have too few stars

Results: known streams

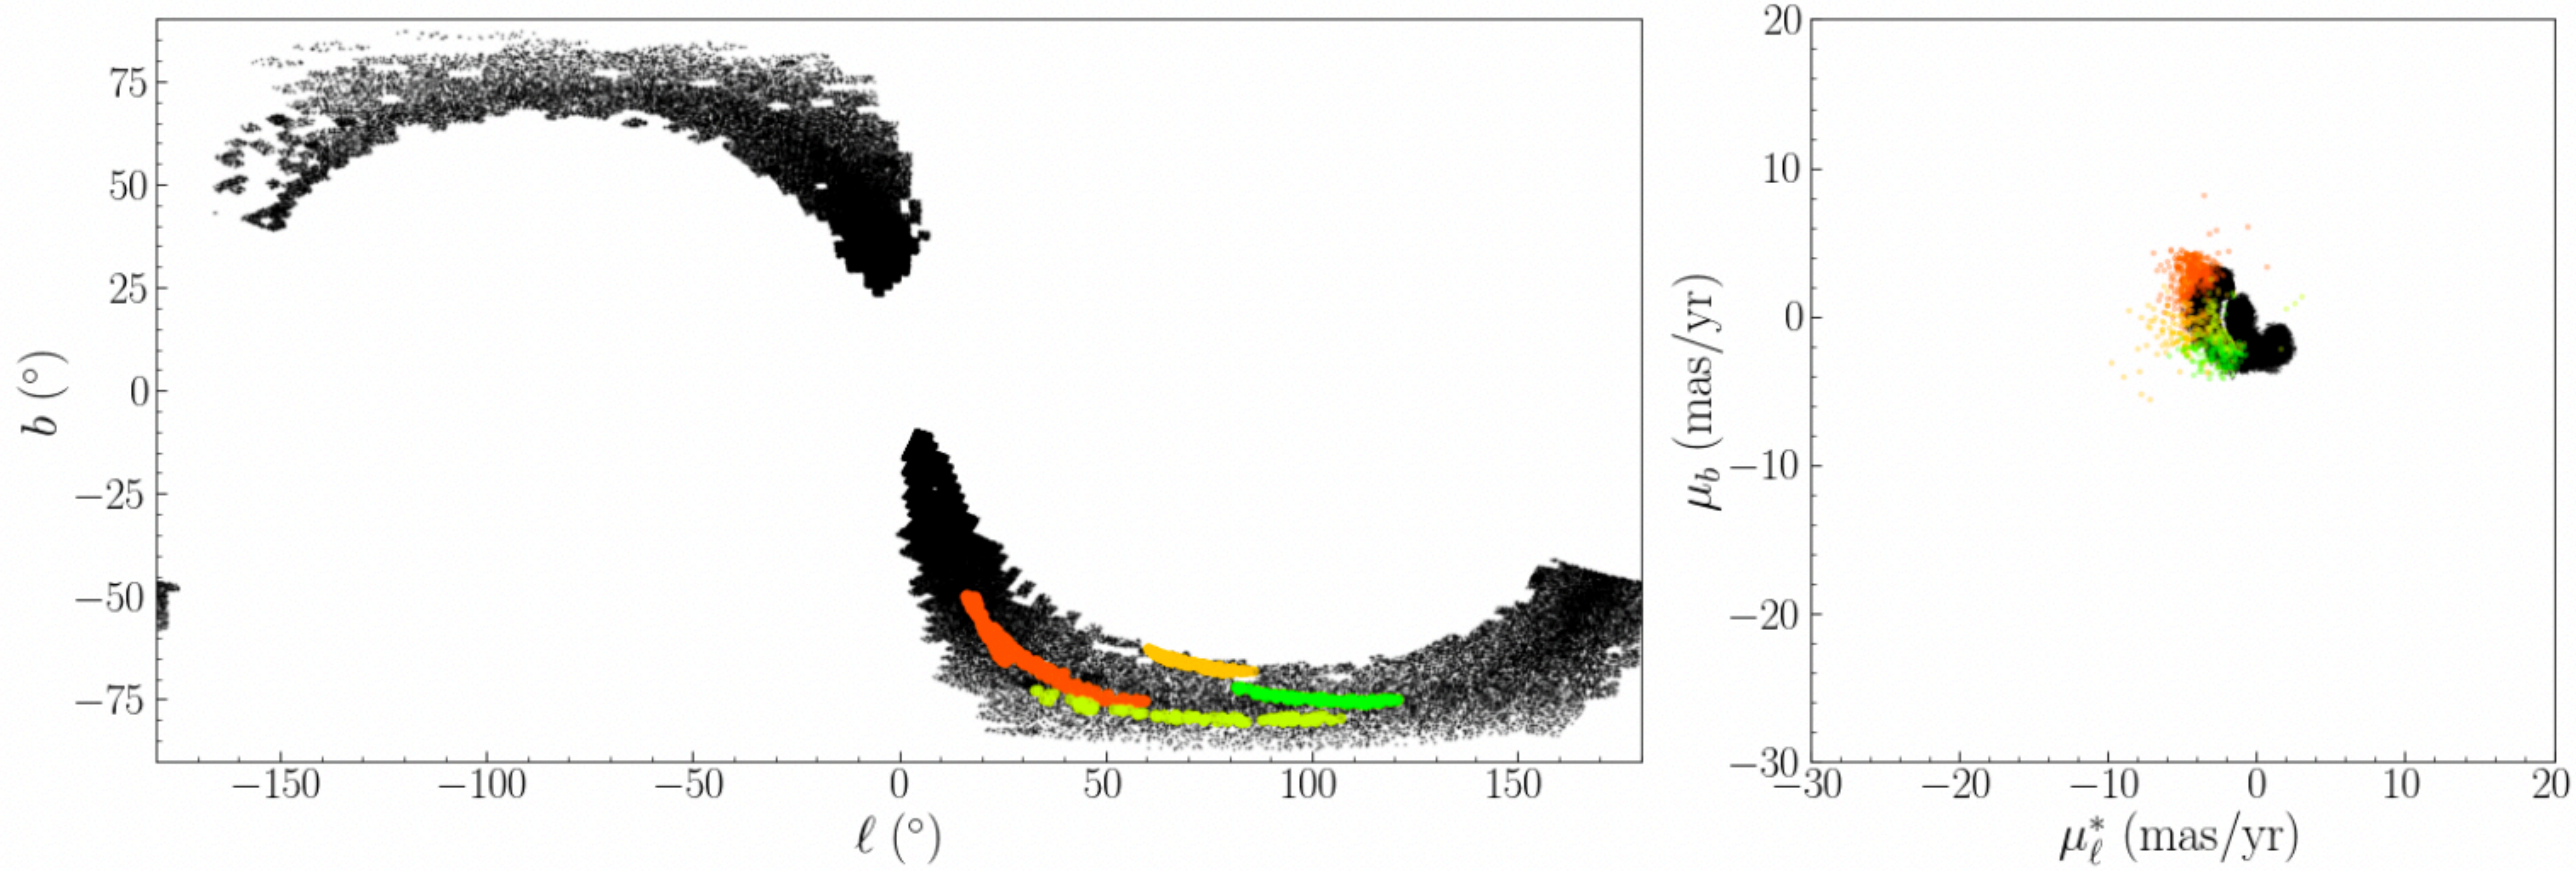

We also recover fragments of the Sagittarius Stream, despite it generally being much too wide for our narrow stream search

New stream candidates from Gaia DR2

[DS, Buckley, Necib 2303.01529]

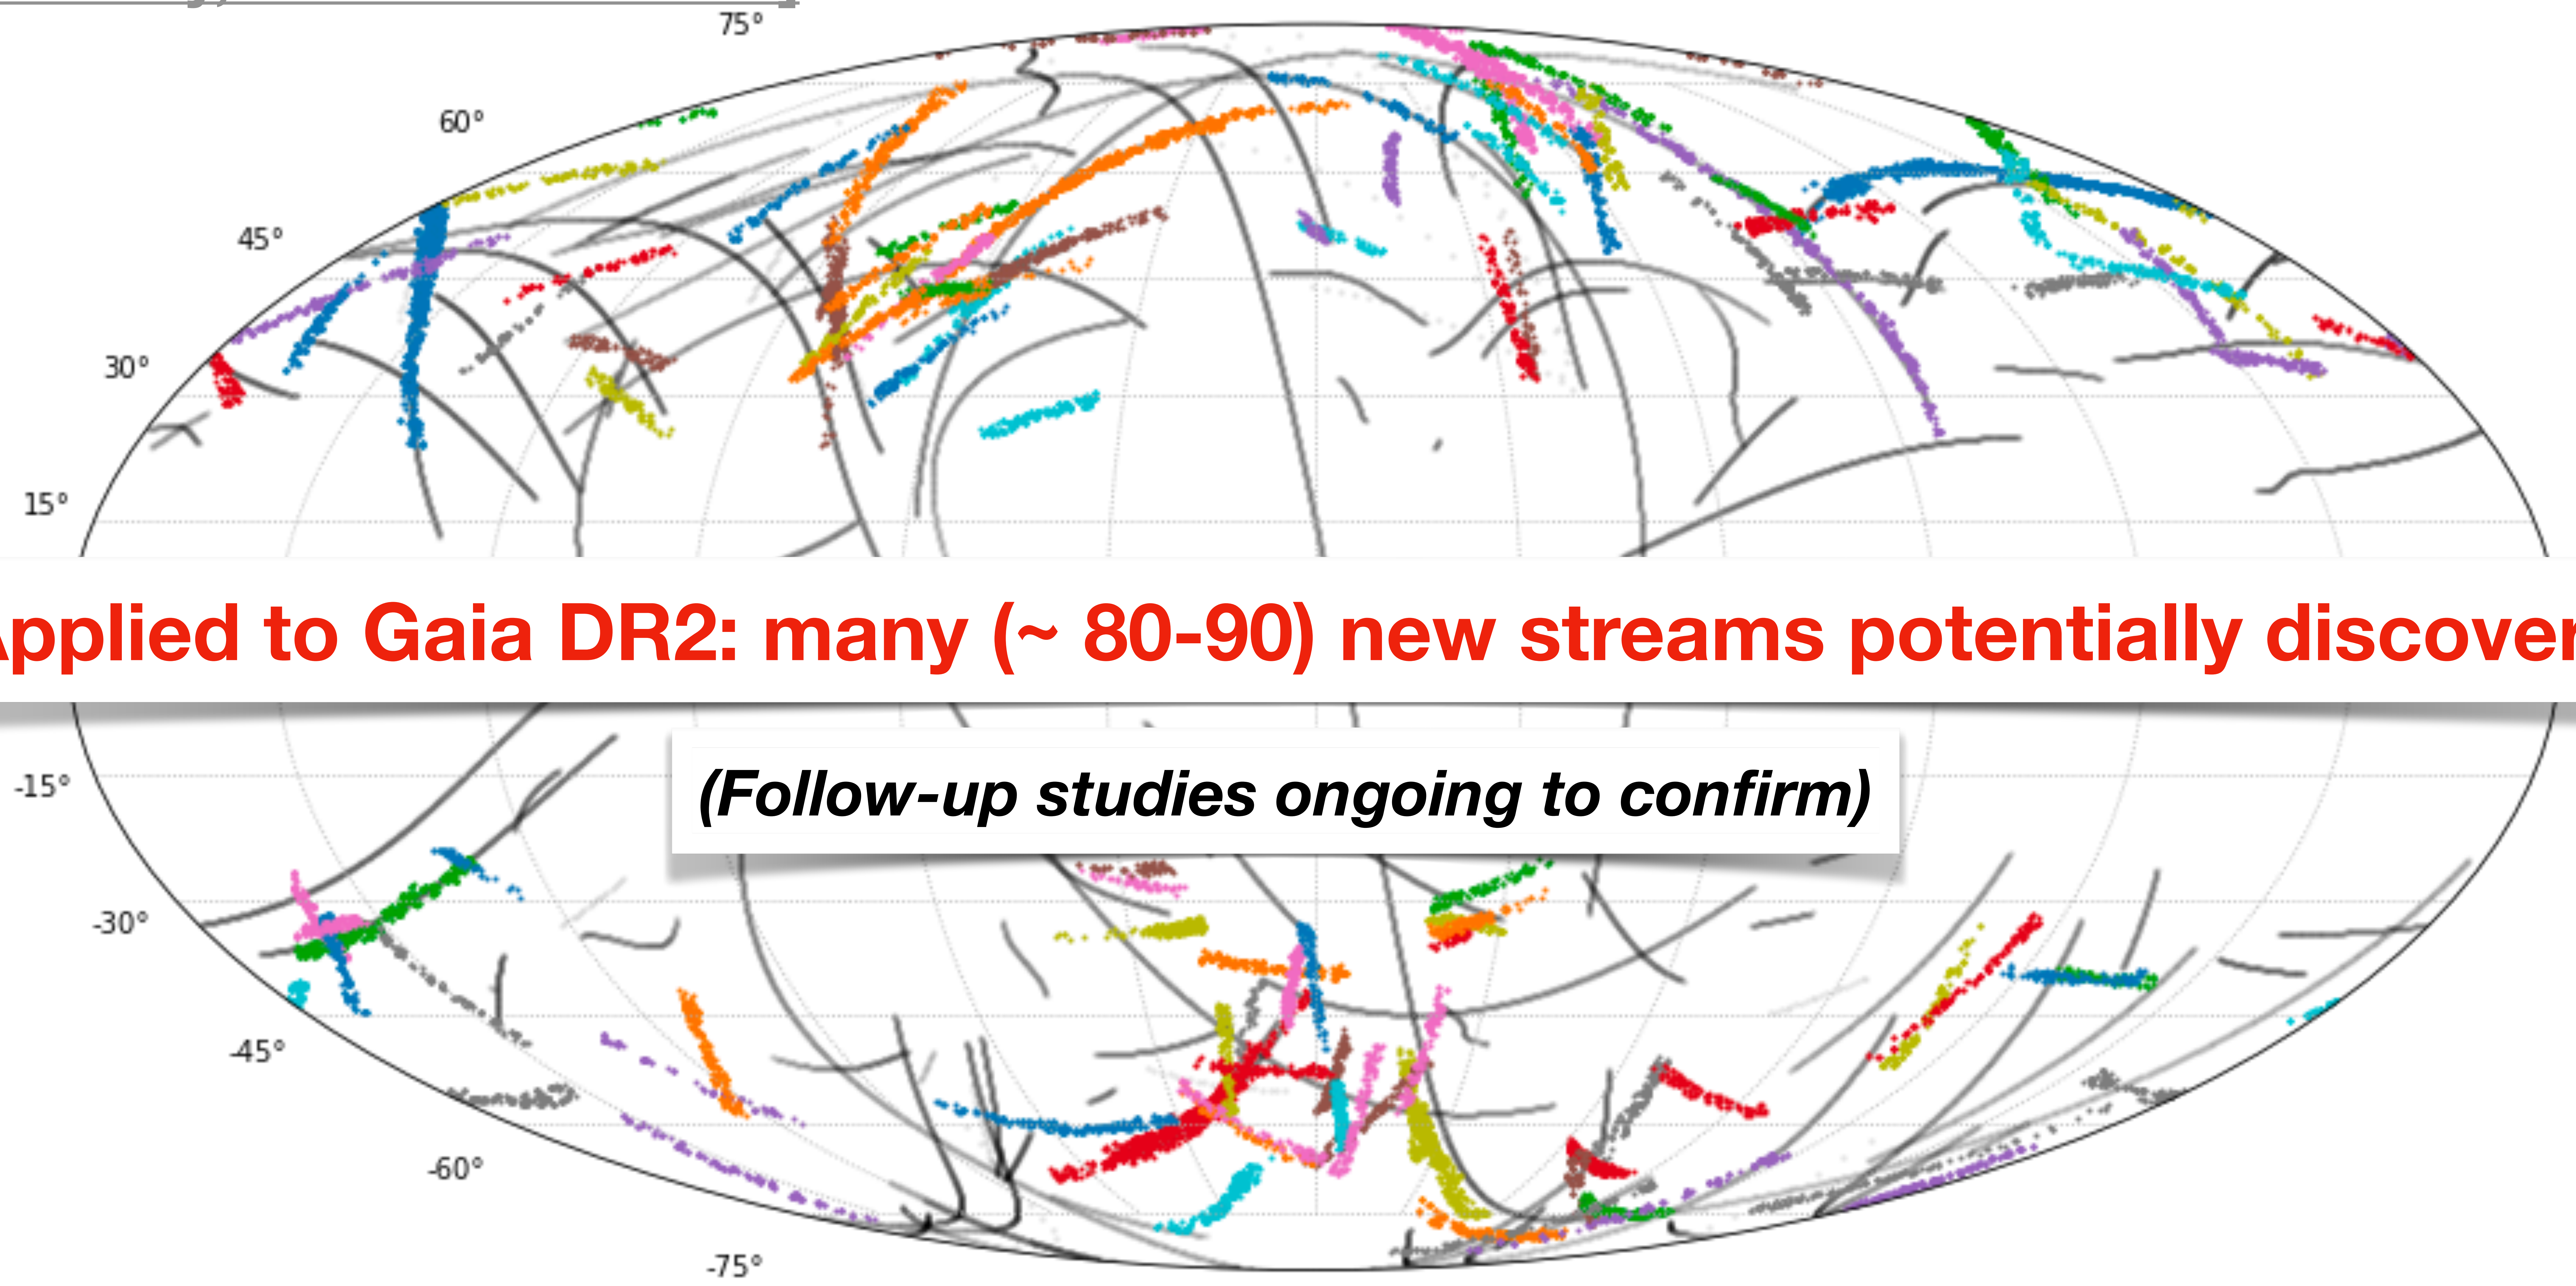

Applied to Gaia DR2: many (~ 80-90) new streams potentially discovered!

(Follow-up studies ongoing to confirm)

3. Bonus: unsupervised ML for measuring DM density with Gaia data

Mapping the local density of DM in 3d

Buckley, Lim, Putney & **DS** 2205.01129, 2305.13358
Green et al 2011.04673, 2205.02244, Naik et al 2112.07657, An et al 2106.05981

- We realized that training density estimators on the Gaia dataset could have other interesting applications
- The full 6D phase space density $p(\vec{x}, \vec{v})$ of all the stars in the Galaxy (or at least all the nearby ones) carries a wealth of information about Galactic dynamics.
- In particular, we can directly infer the mass density $\rho(\vec{x})$ of the Galaxy from knowledge of $p(\vec{x}, \vec{v})$, and from that the mass density $\rho_{DM}(\vec{x})$ of the dark matter.

Local dark matter density

Knowing the local dark matter density $\rho_{DM}(x)$ is very important for many reasons:

from 2104.07634

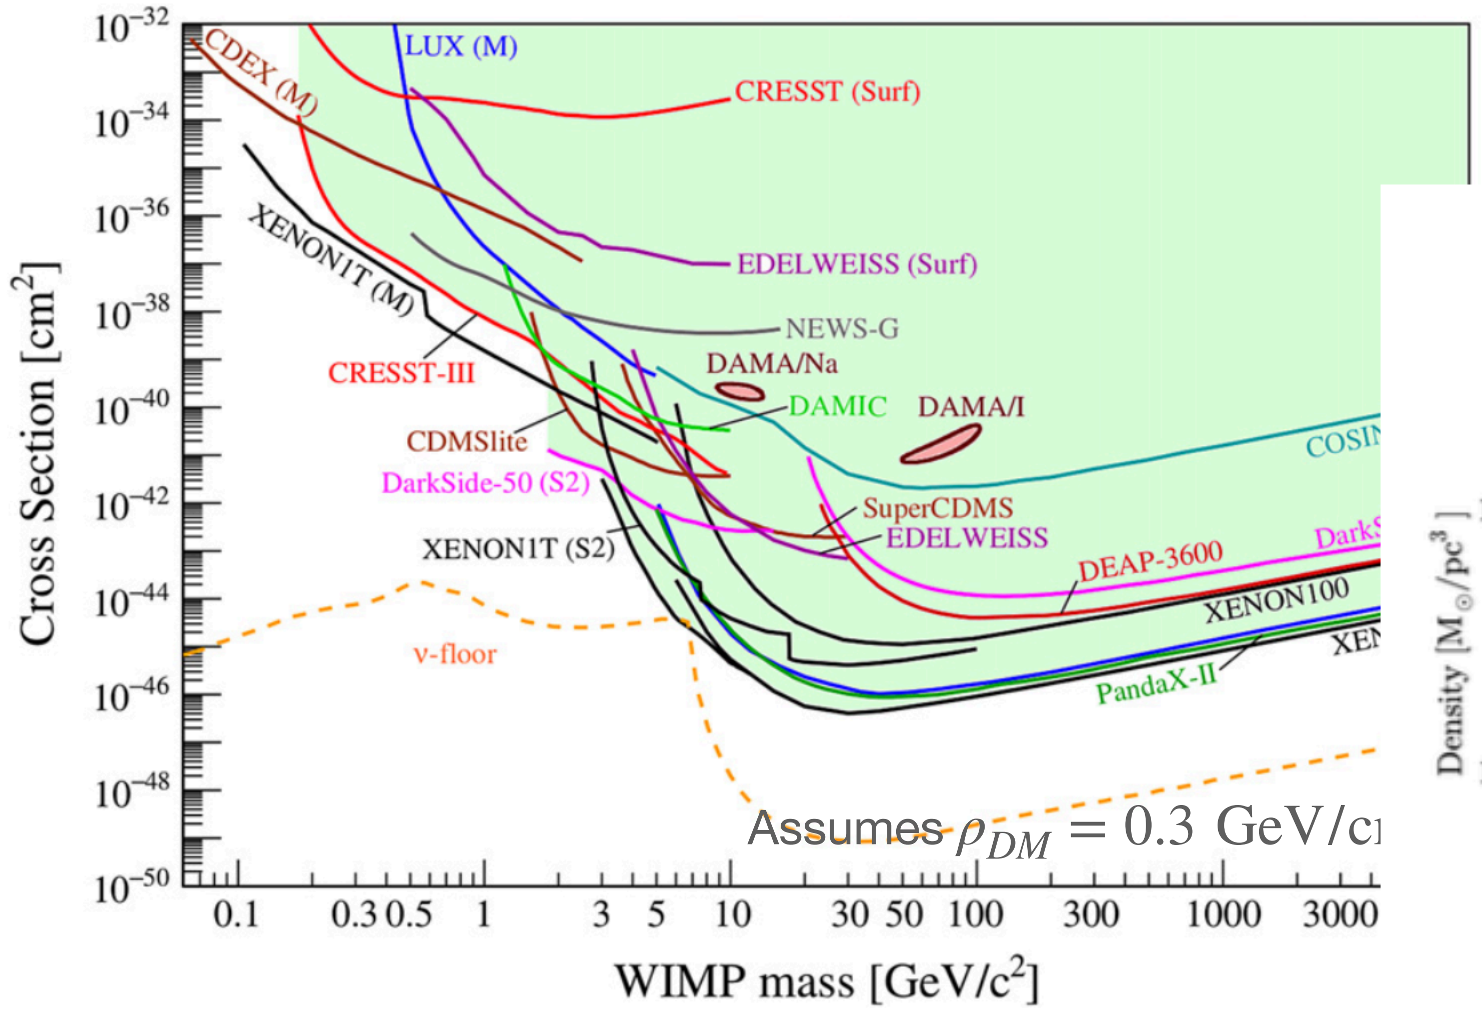

Input to direct detection experiments

from 1812.00044

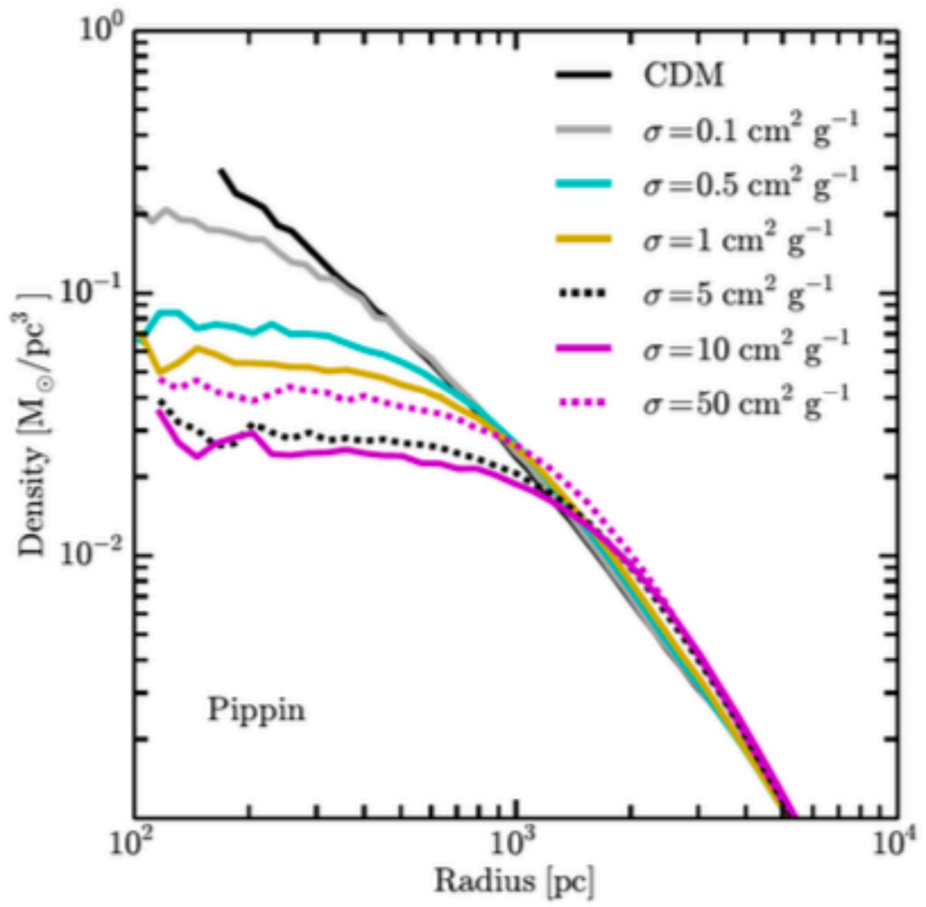

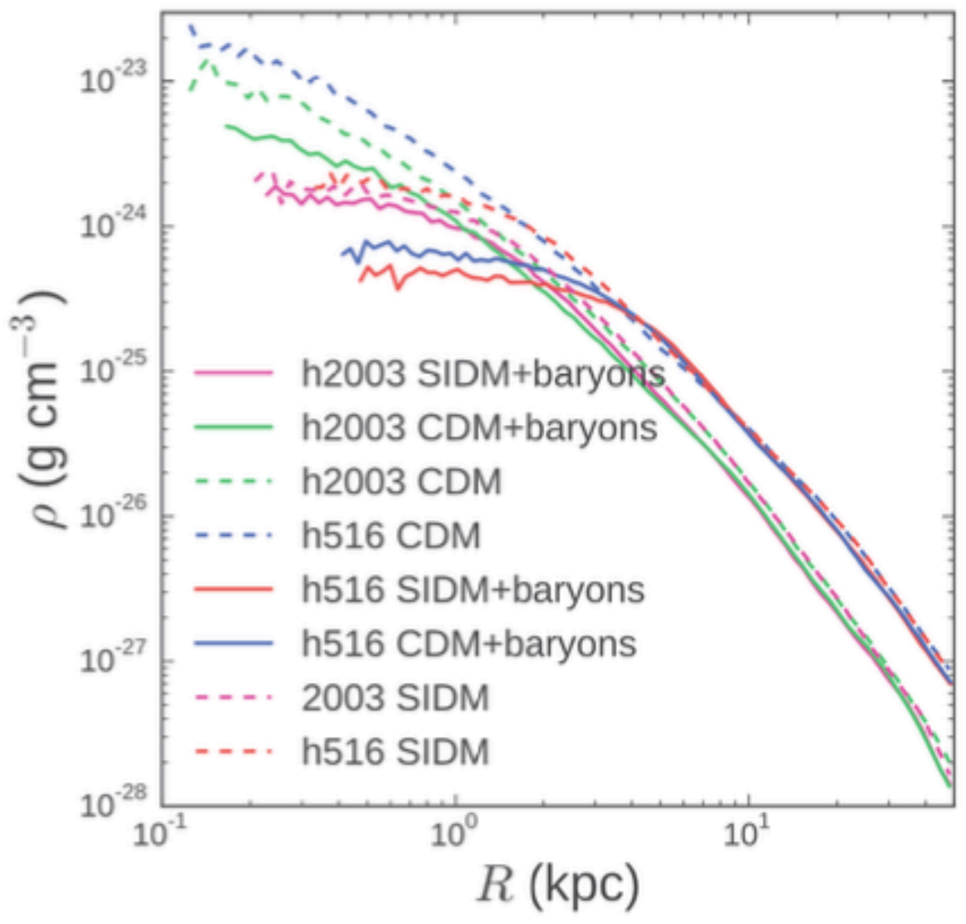

Make contact with models (NFW, Einasto, etc) and simulations, learn more about Galaxy formation and nature of dark matter

from 2001.05503

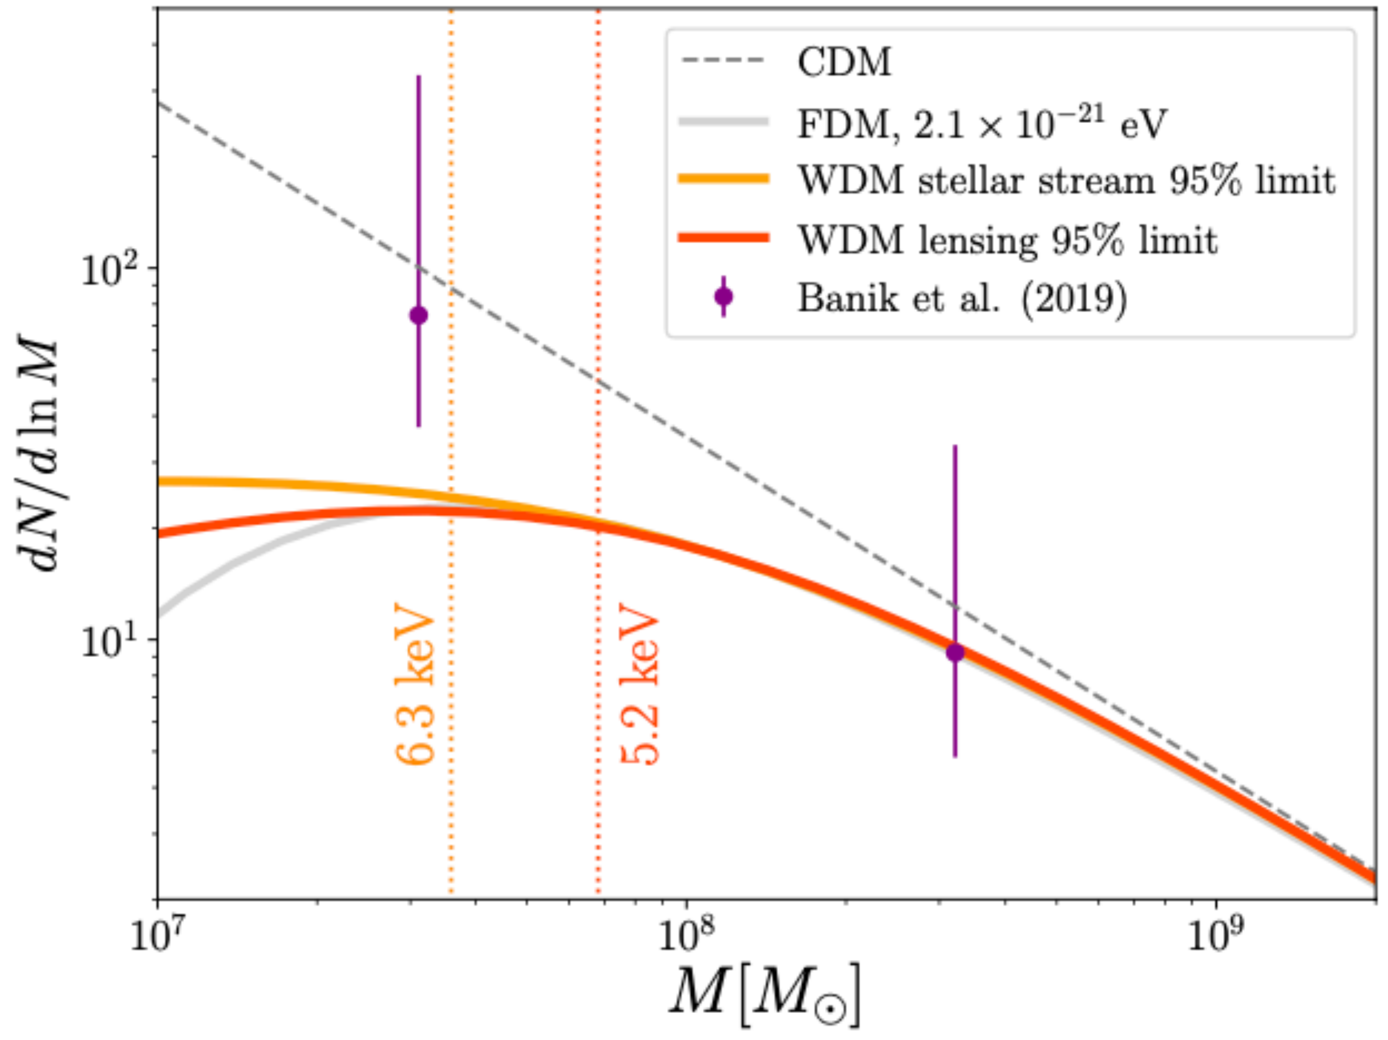

Could potentially resolve the presence of dark matter substructure

The Collisionless Boltzmann Equation

The Collisionless Boltzmann Equation

- Baryons+DM source the galactic potential $\Phi(x)$. Gravitational tracers (stars) drawn from $p(\vec{x}, \vec{v}, t)$ accelerate in response to $\Phi(x)$.

$$\frac{dp}{dt} = \left[\frac{\partial}{\partial t} + \vec{v} \cdot \frac{\partial}{\partial \vec{x}} + \vec{a}(\vec{x}) \cdot \frac{\partial}{\partial \vec{v}}\right] p = 0 \qquad \vec{a}(\vec{x}) = -\nabla \Phi(\vec{x})$$

The Collisionless Boltzmann Equation

- Baryons+DM source the galactic potential $\Phi(x)$. Gravitational tracers (stars) drawn from $p(\vec{x}, \vec{v}, t)$ accelerate in response to $\Phi(x)$.

$$\frac{dp}{dt} = \left[\frac{\partial}{\partial t} + \vec{v} \cdot \frac{\partial}{\partial \vec{x}} + \vec{a}(\vec{x}) \cdot \frac{\partial}{\partial \vec{v}}\right] p = 0 \qquad \vec{a}(\vec{x}) = -\nabla \Phi(\vec{x})$$

- Over many dynamic timescales, $p(\vec{x}, \vec{v}, t)$ equilibrates $\rightarrow p(\vec{x}, \vec{v})$

The Collisionless Boltzmann Equation

Example: 2D phase space in toy potential

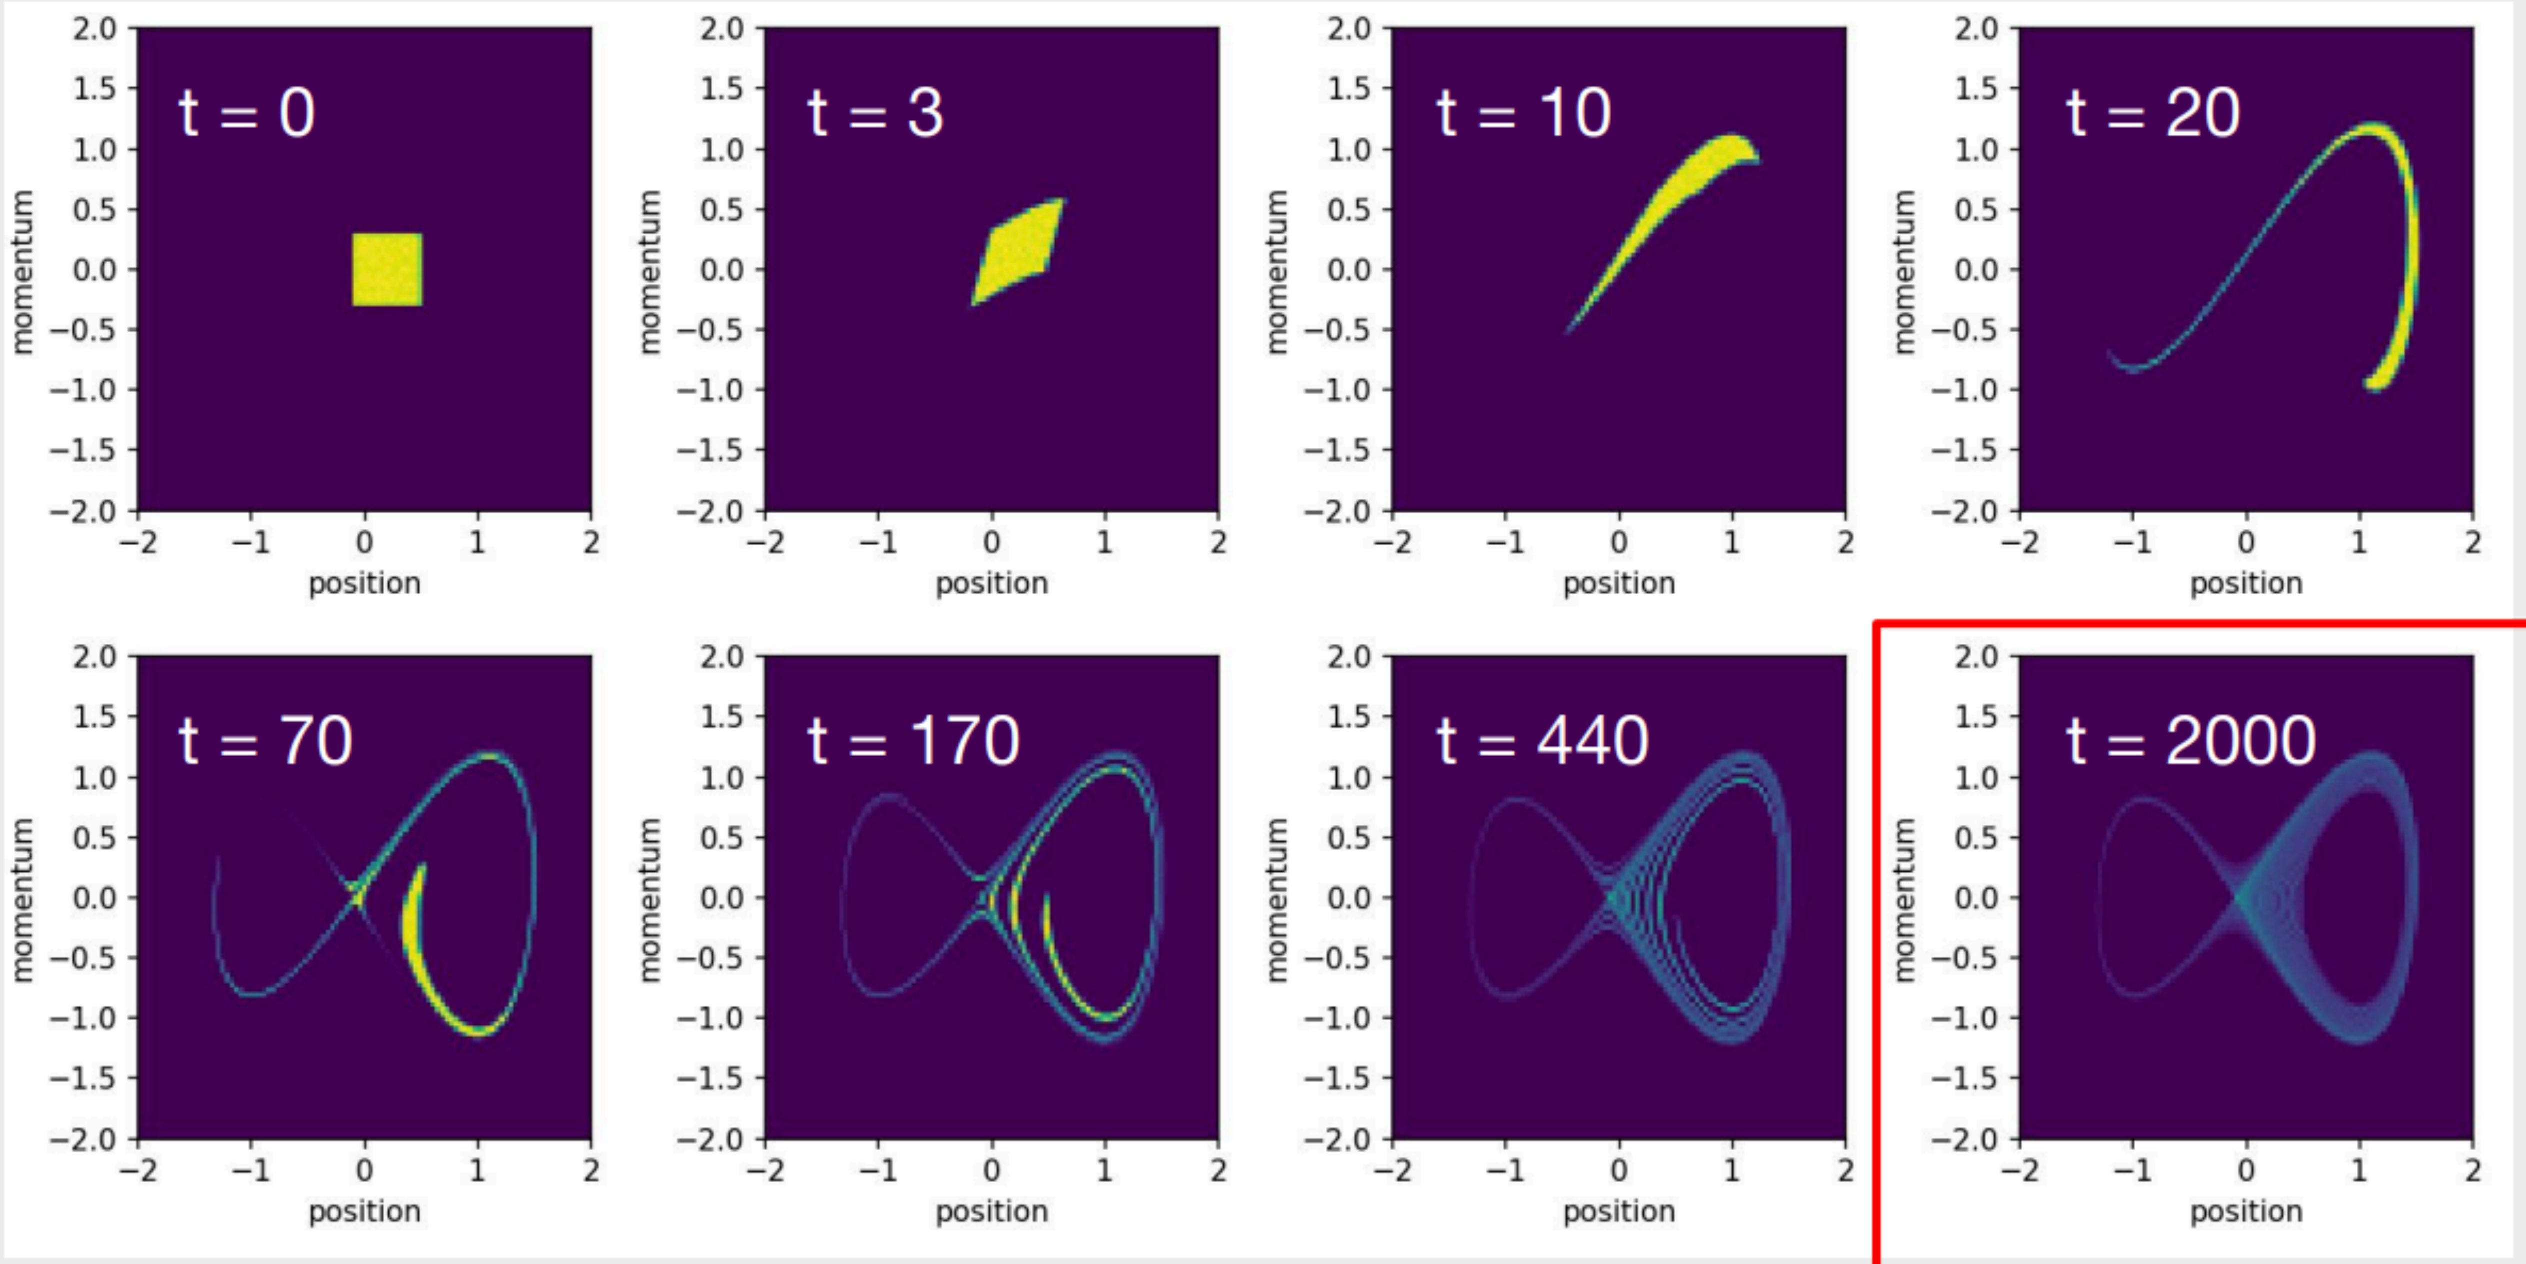

Dynamic equilibrium: f(**x**,**v**) for old stars in our galaxy should be smooth and "phase mixed"

$$\frac{df}{dt} = 0, \quad \frac{\partial f}{\partial t} = 0$$

Figure Credit: Eric Putney

(stars)

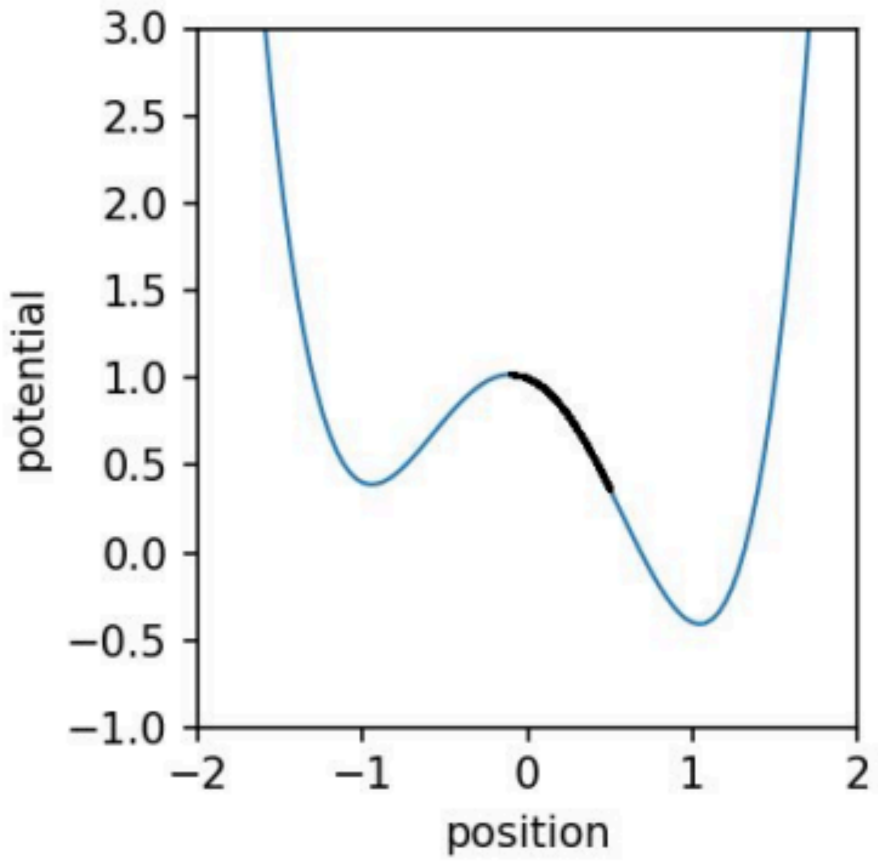

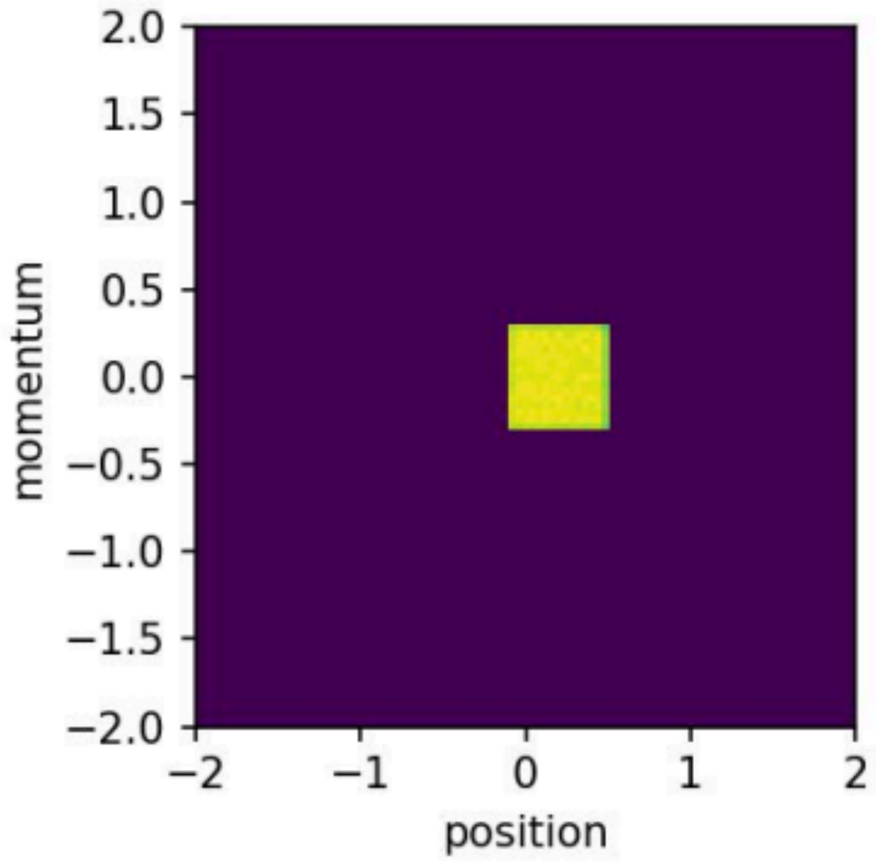

The Collisionless Boltzmann Equation

- Baryons+DM source the galactic potential $\Phi(x)$. Gravitational tracers (stars) drawn from $p(\vec{x}, \vec{v}, t)$ accelerate in response to $\Phi(x)$.

$$\frac{dp}{dt} = \left[\frac{\partial}{\partial t} + \vec{v} \cdot \frac{\partial}{\partial \vec{x}} + \vec{a}(\vec{x}) \cdot \frac{\partial}{\partial \vec{v}}\right] p = 0 \qquad \vec{a}(\vec{x}) = -\nabla\Phi(\vec{x})$$

- Over many dynamic timescales, $p(\vec{x}, \vec{v}, t)$ equilibrates $\rightarrow p(\vec{x}, \vec{v})$

The Collisionless Boltzmann Equation

- Baryons+DM source the galactic potential $\Phi(x)$. Gravitational tracers (stars) drawn from $p(\vec{x}, \vec{v}, t)$ accelerate in response to $\Phi(x)$.

$$\frac{dp}{dt} = \left[\cancel{\frac{\partial}{\partial t}} + \vec{v} \cdot \frac{\partial}{\partial \vec{x}} + \vec{a}(\vec{x}) \cdot \frac{\partial}{\partial \vec{v}}\right] p = 0 \qquad \vec{a}(\vec{x}) = -\nabla\Phi(\vec{x})$$

- Over many dynamic timescales, $p(\vec{x}, \vec{v}, t)$ equilibrates $\rightarrow p(\vec{x}, \vec{v})$

- We can use a snapshot of $p(\vec{x}, \vec{v})$ today to infer the acceleration field $\vec{a}(\vec{x})$

From phase space density to mass density

Buckley, Lim, Putney & **DS** 2205.01129, 2305.13358
Green et al 2011.04673, 2205.02244, Naik et al 2112.07657, An et al 2106.05981

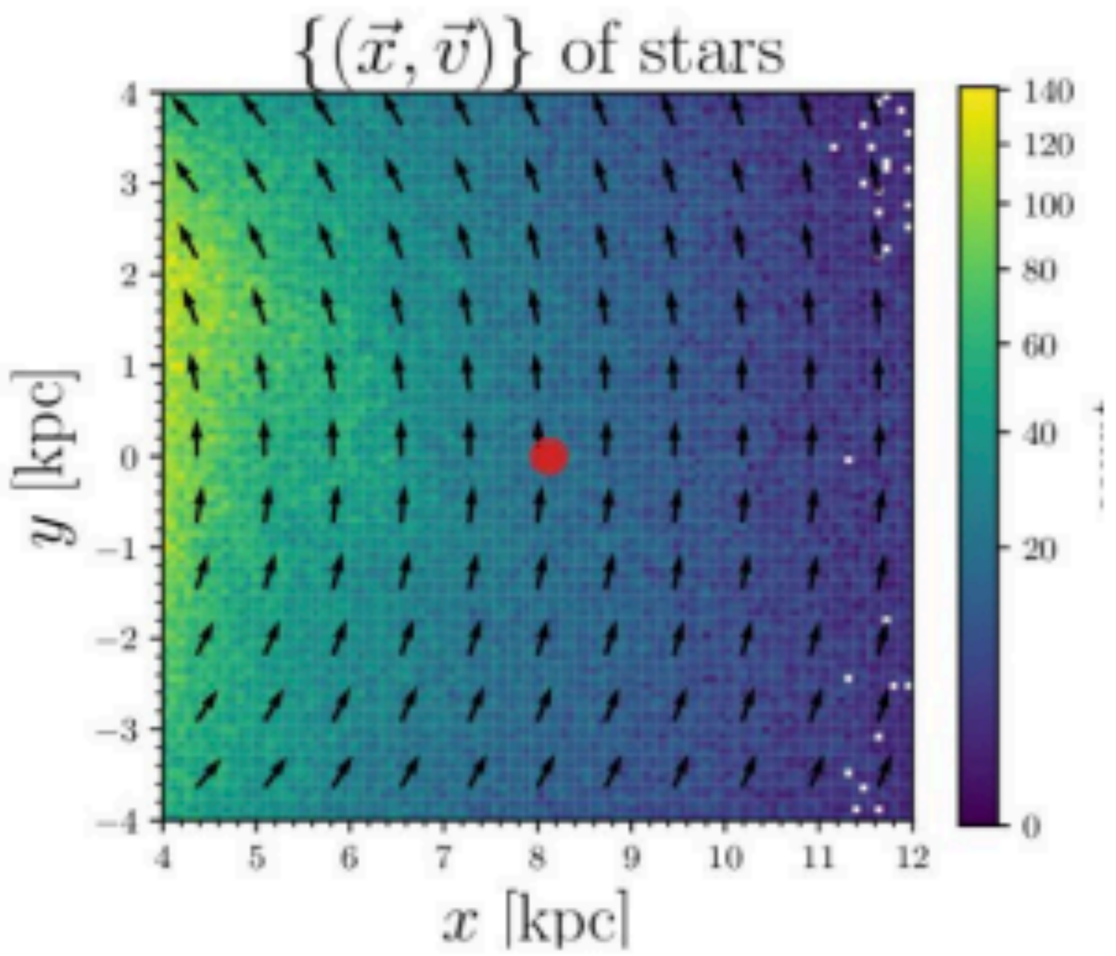

Figure credit:
S.H. Lim

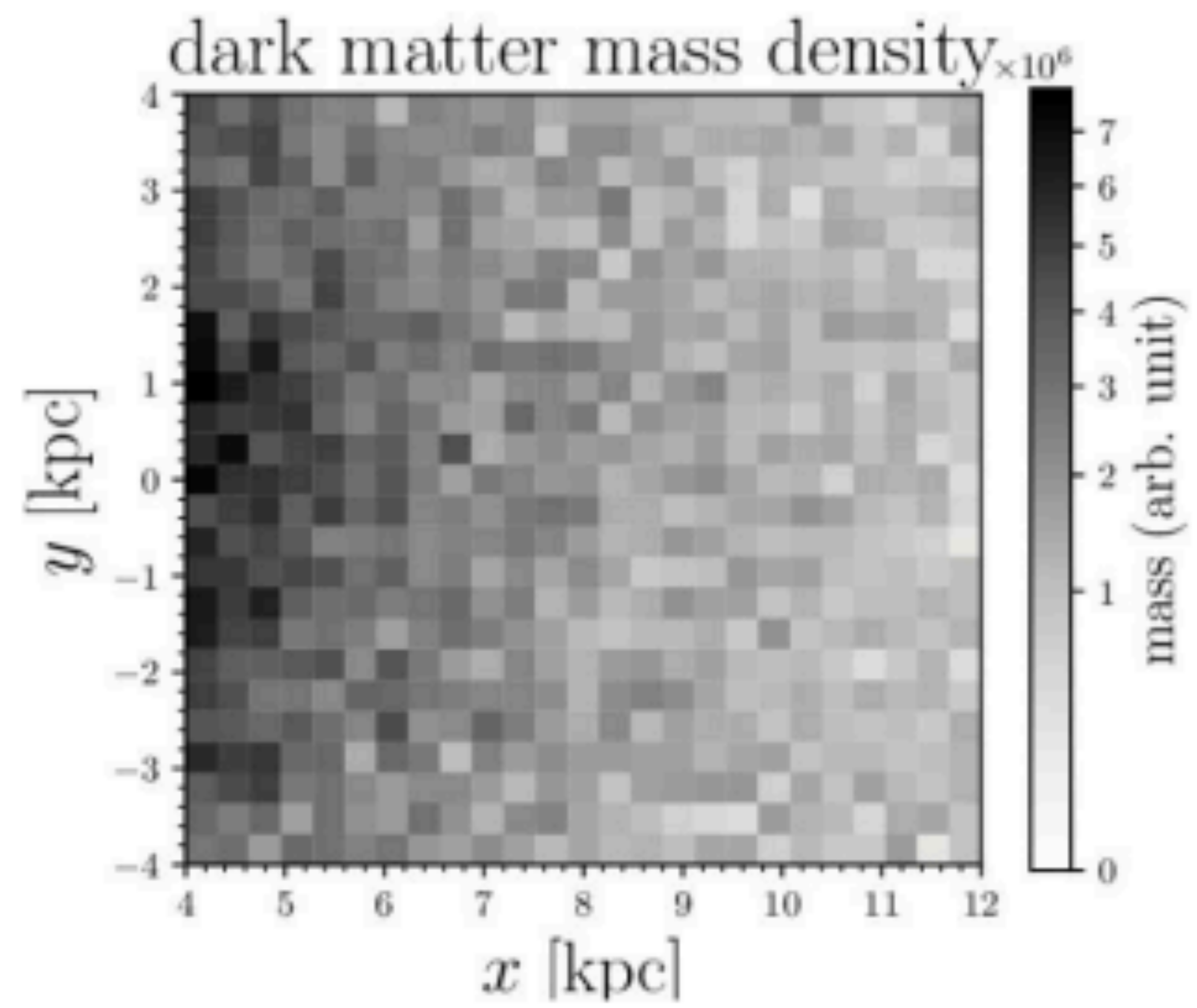

$$\left[\cancel{\frac{\partial}{\partial t}} + \vec{v}\cdot\frac{\partial}{\partial \vec{x}} + \vec{a}\cdot\frac{\partial}{\partial \vec{v}}\right] f(\vec{x},\vec{v}) = 0$$

$$\{(\vec{x},\vec{v})\} \longrightarrow f(\vec{x},\vec{v}) \longrightarrow \vec{a}(\vec{x}) \longrightarrow \rho(\vec{x})$$

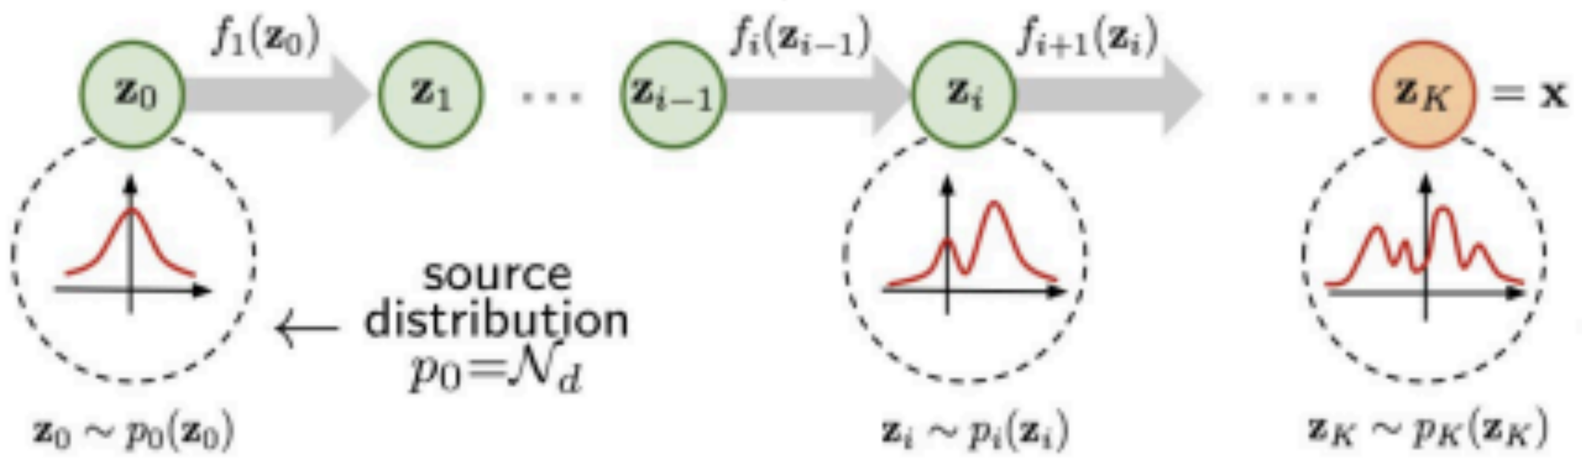

$$\rho = -\frac{\nabla\cdot\vec{a}}{4\pi G}$$

Comparison with previous approaches

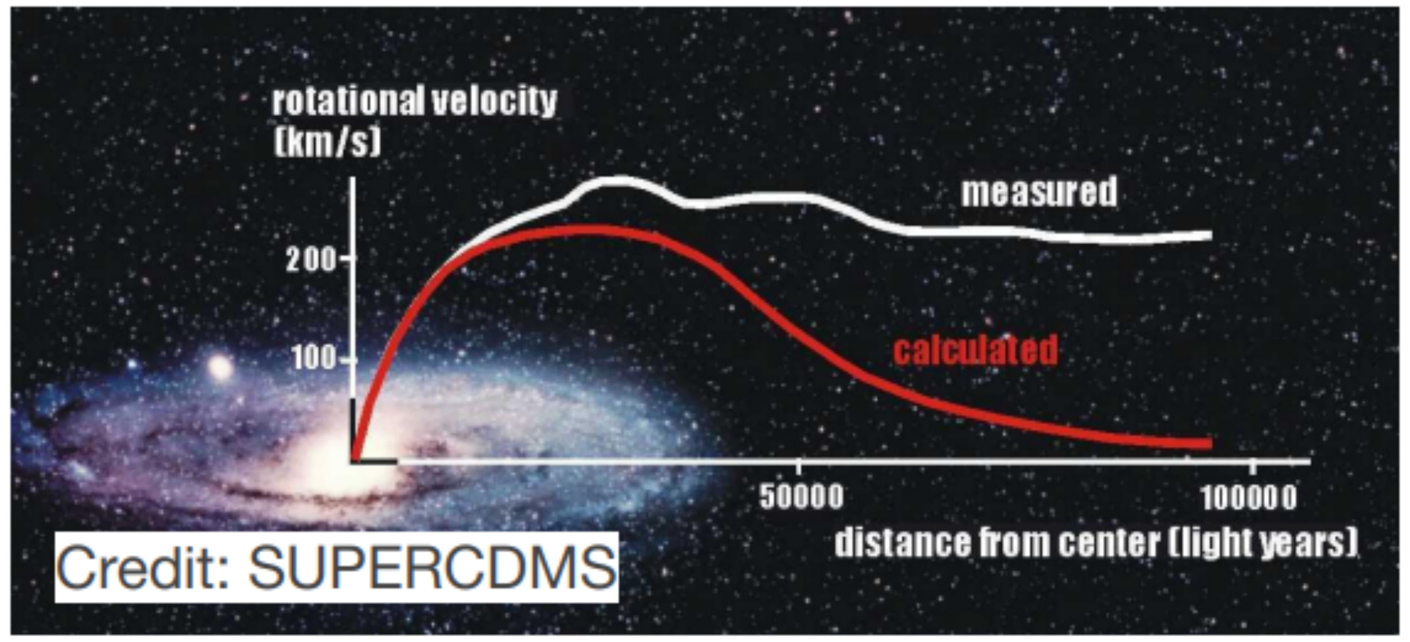

- Existing measurements typically use **Jean's equation** (second moment of Boltzmann equation) or **rotation curves**
- **They make many assumptions** (axisymmetry, reflection symmetry, simple parametric models…) and **bin the data**
- Results can seem precise but might not be accurate (biased)

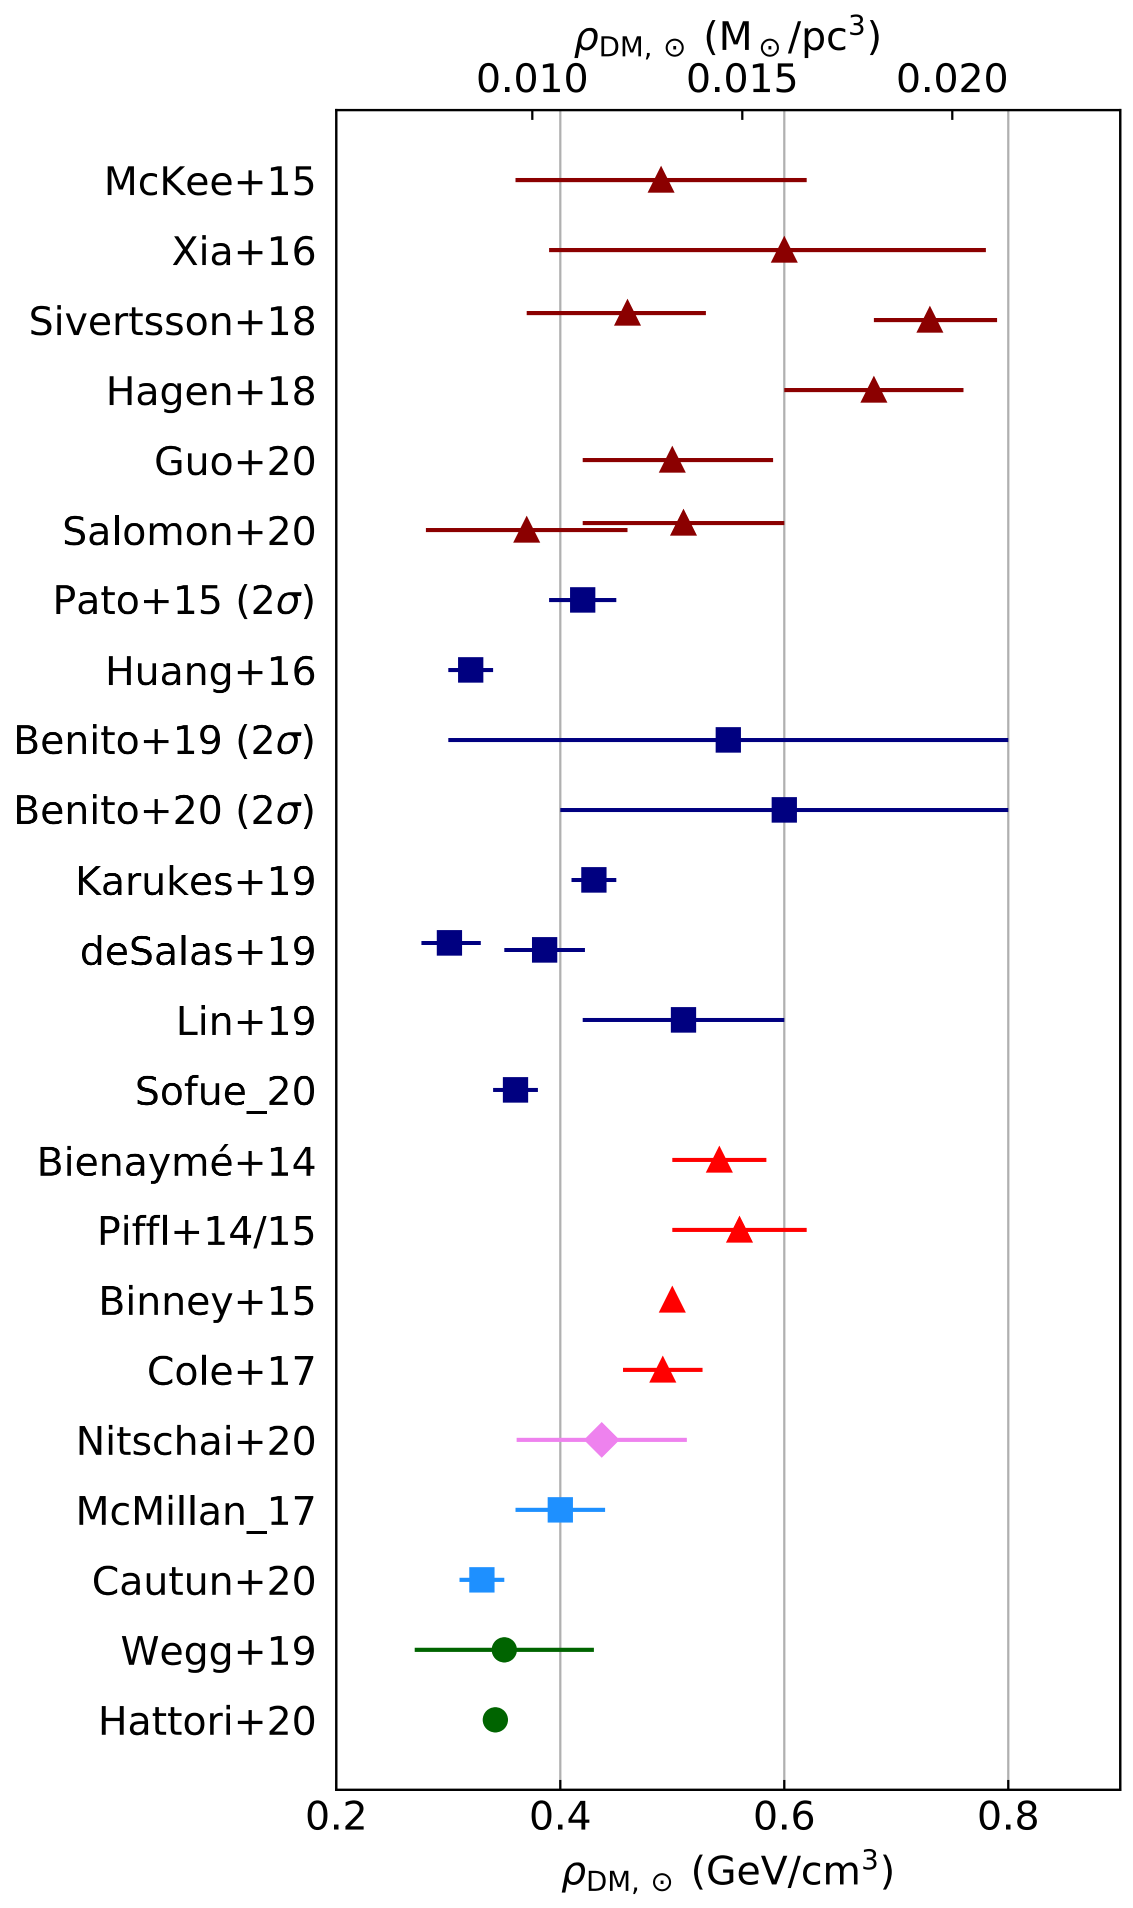

From de Salas & Widmark 2012.11477

Comparison with previous approaches

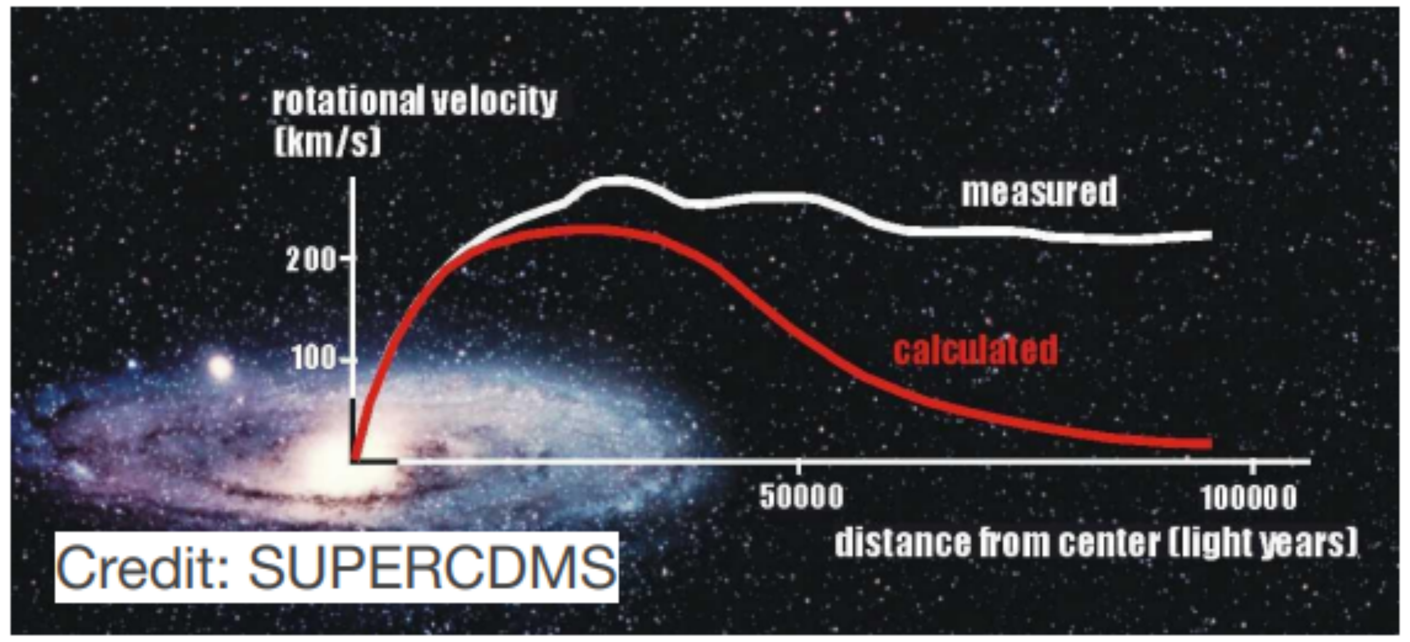

- Existing measurements typically use **Jean's equation** (second moment of Boltzmann equation) or **rotation curves**
- **They make many assumptions** (axisymmetry, reflection symmetry, simple parametric models…) and **bin the data**
- Results can seem precise but might not be accurate (biased)

Our approach using normalizing flows is model-free, does not assume symmetries, and is unbinned

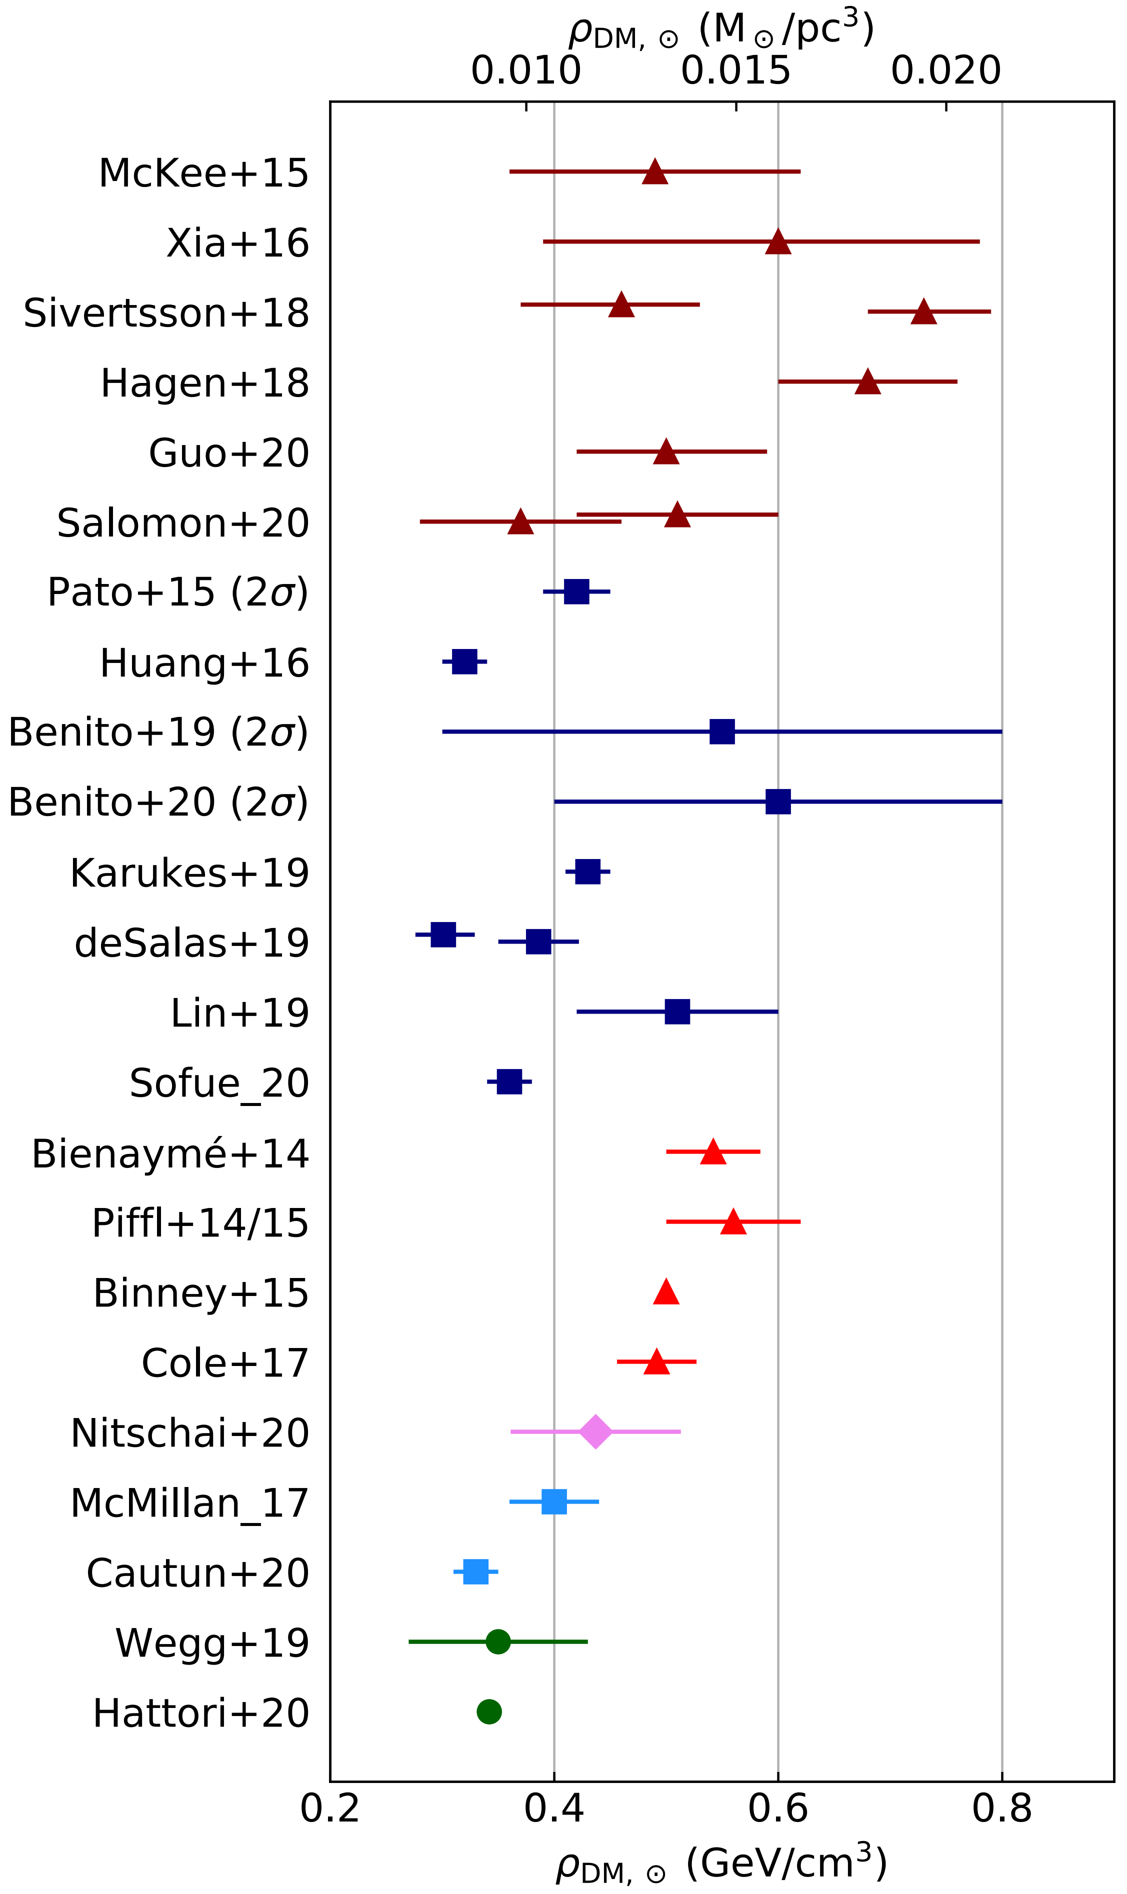

From de Salas & Widmark 2012.11477

Comparison with previous approaches

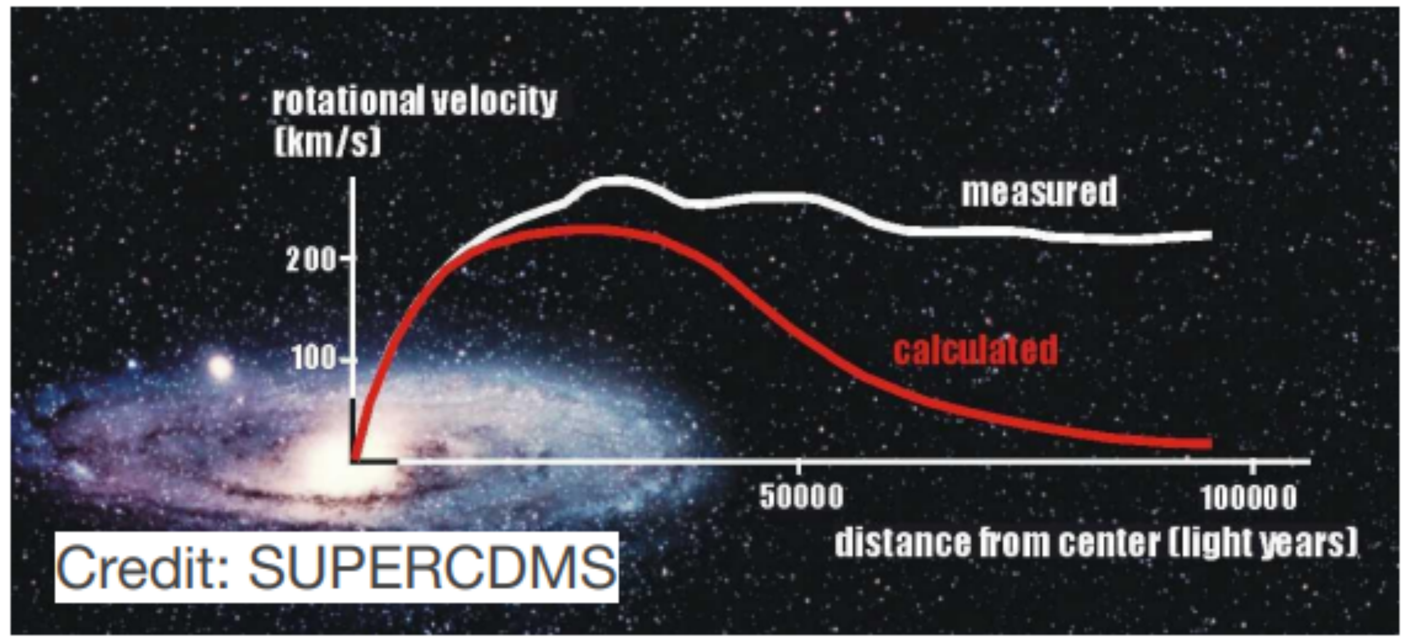

- Existing measurements typically use **Jean's equation** (second moment of Boltzmann equation) or **rotation curves**
- **They make many assumptions** (axisymmetry, reflection symmetry, simple parametric models…) and **bin the data**
- Results can seem precise but might not be accurate (biased)

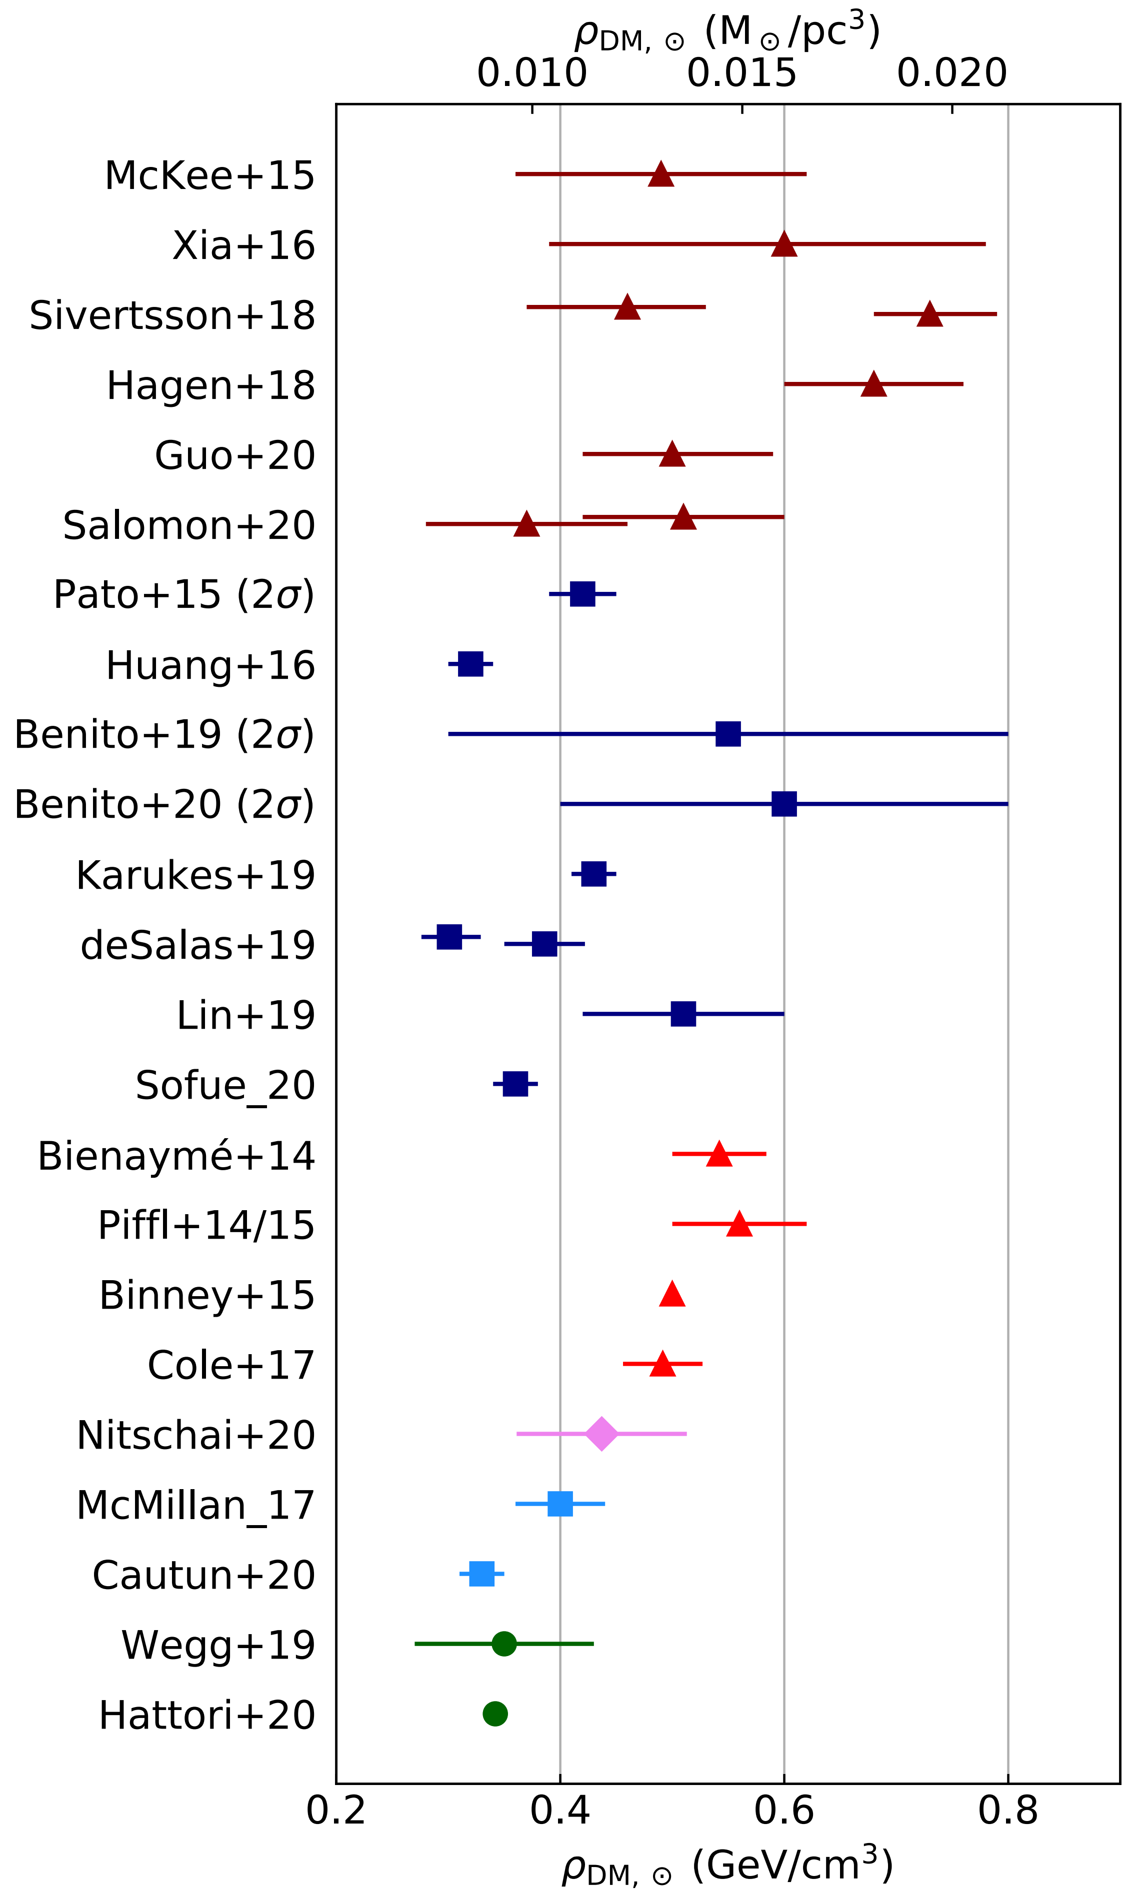

Our approach using normalizing flows is model-free, does not assume symmetries, and is unbinned

First ever fully 3d measurement of dark matter density in the solar neighborhood

From de Salas & Widmark 2012.11477

From proof-of-concept to real data

Buckley, Lim, Putney & **DS** 2205.01129, 2305.13358

- After validating our method with a realistic hydrodynamical cosmological simulation, we applied it to Gaia DR3.
- Selected stars in Gaia DR3 within 4 kpc with
  - full 6d features
  - brightness cut to ensure completeness
- dominated by **"red clump" stars** which are supposed to be a good equilibrium tracer population => **5.8M stars**

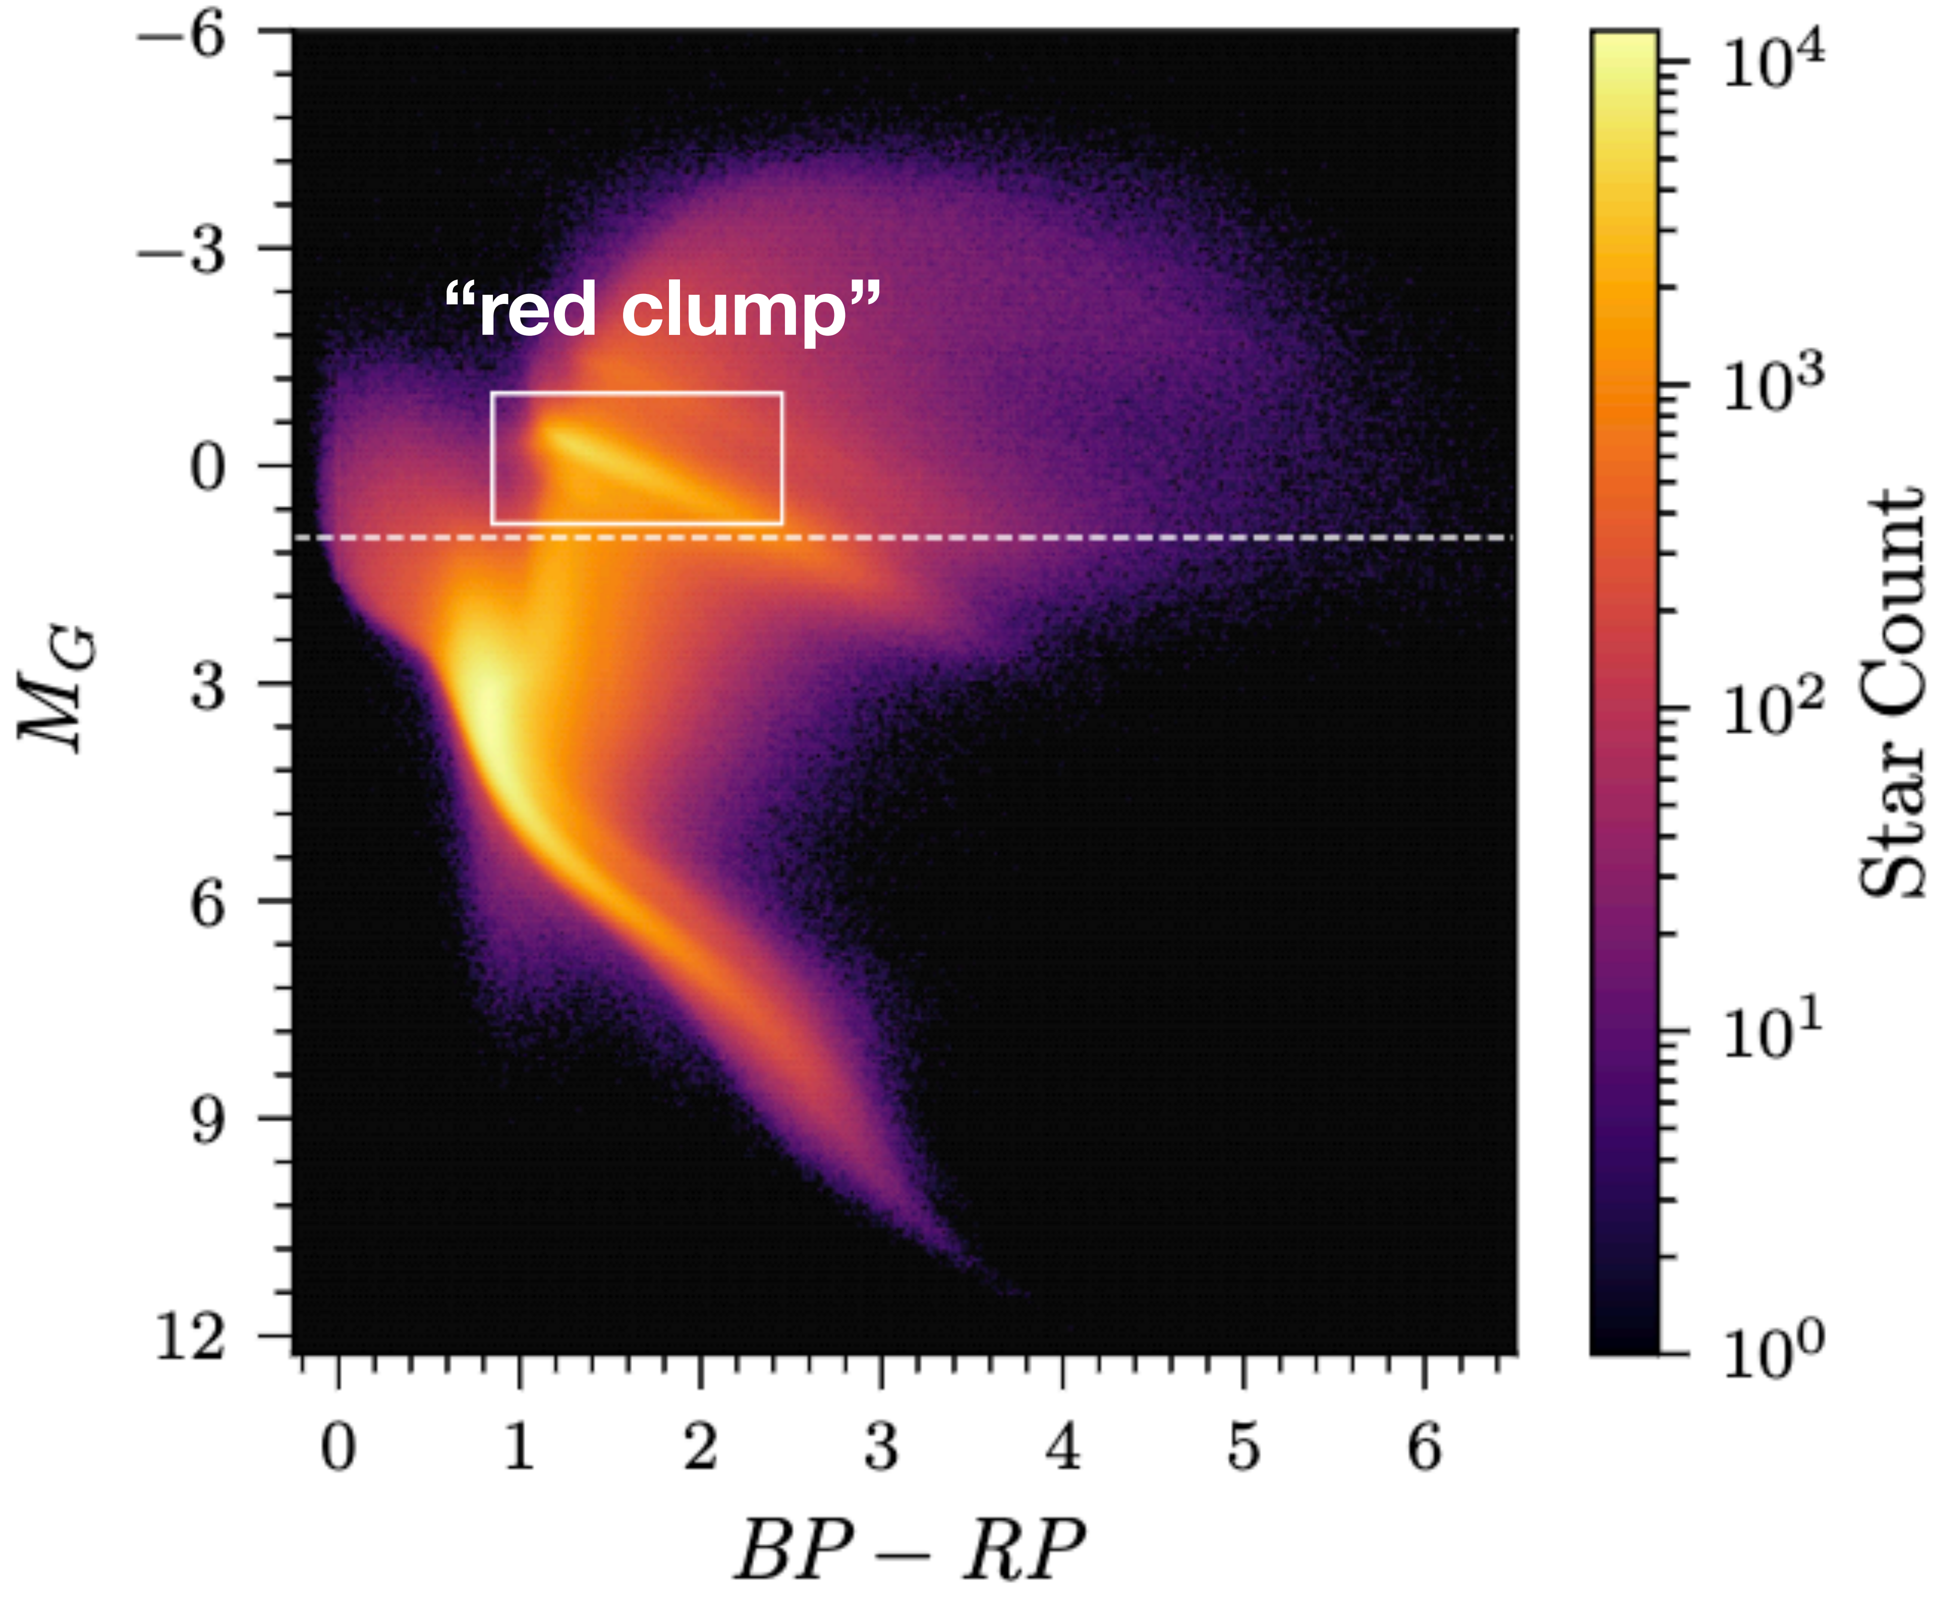

Results: density estimation

Lim, Putney, Buckley & **DS** 2305.13358

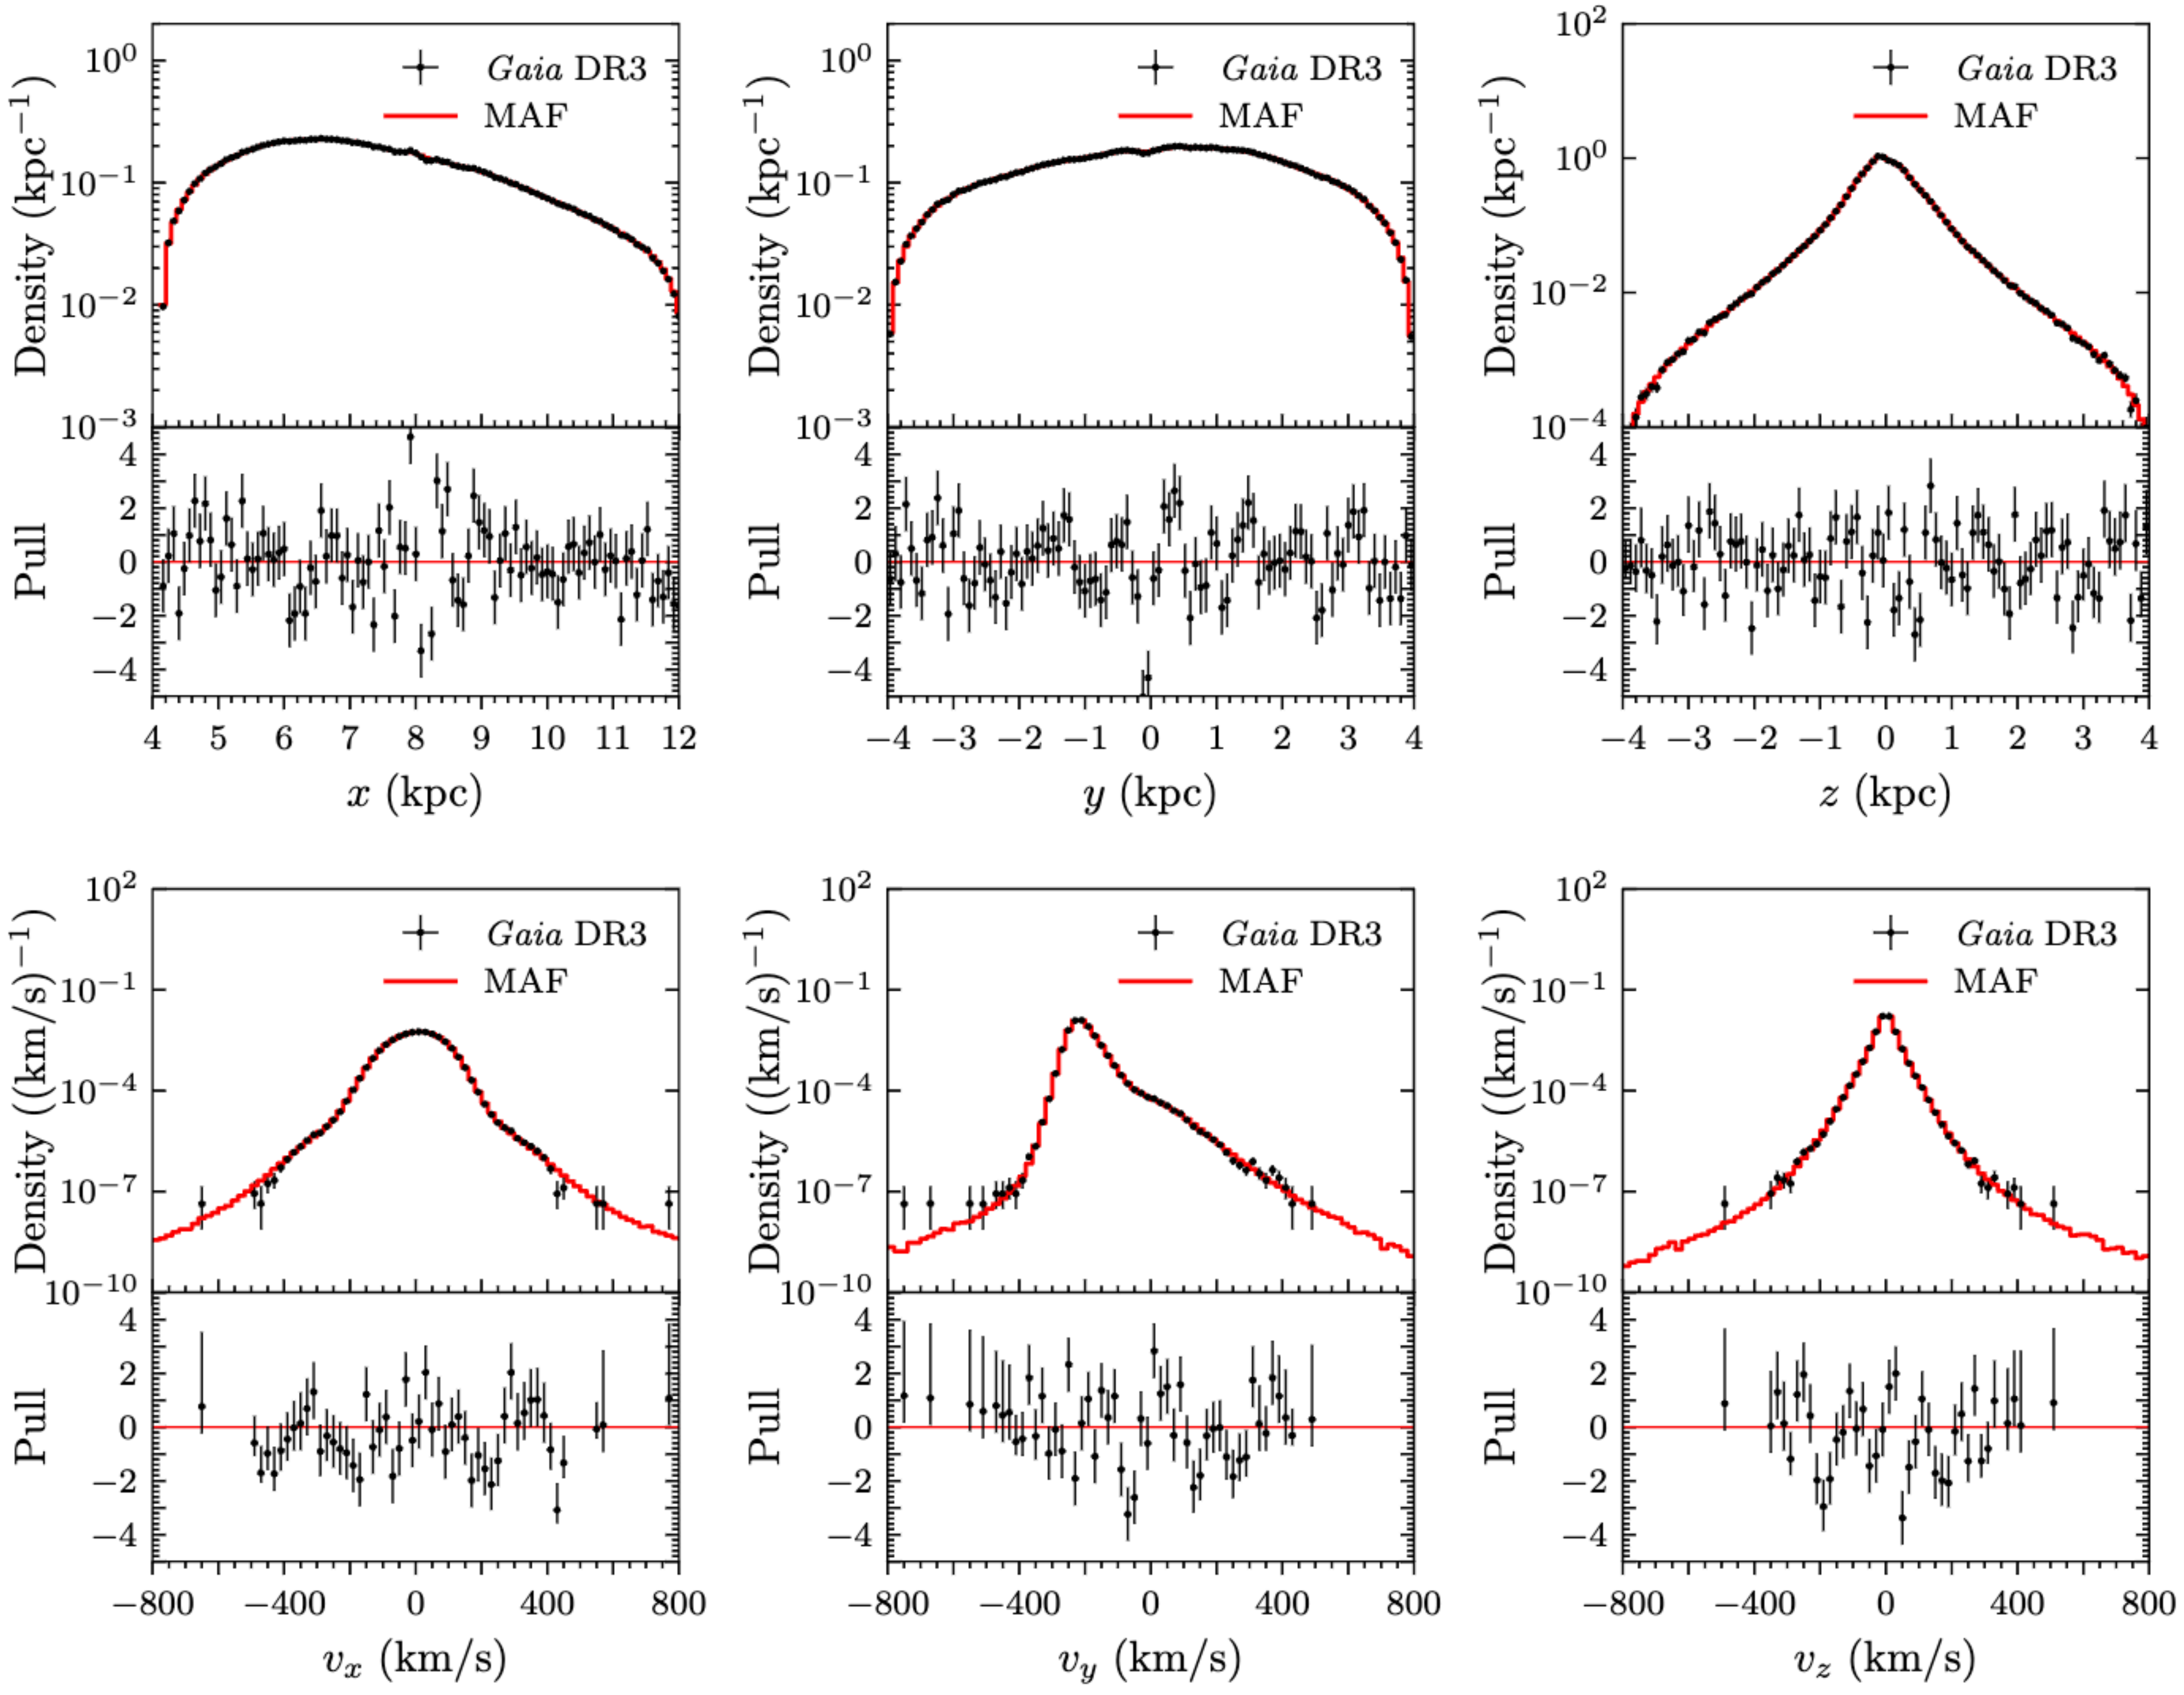

Results: accelerations

Lim, Putney, Buckley & **DS** 2305.13358

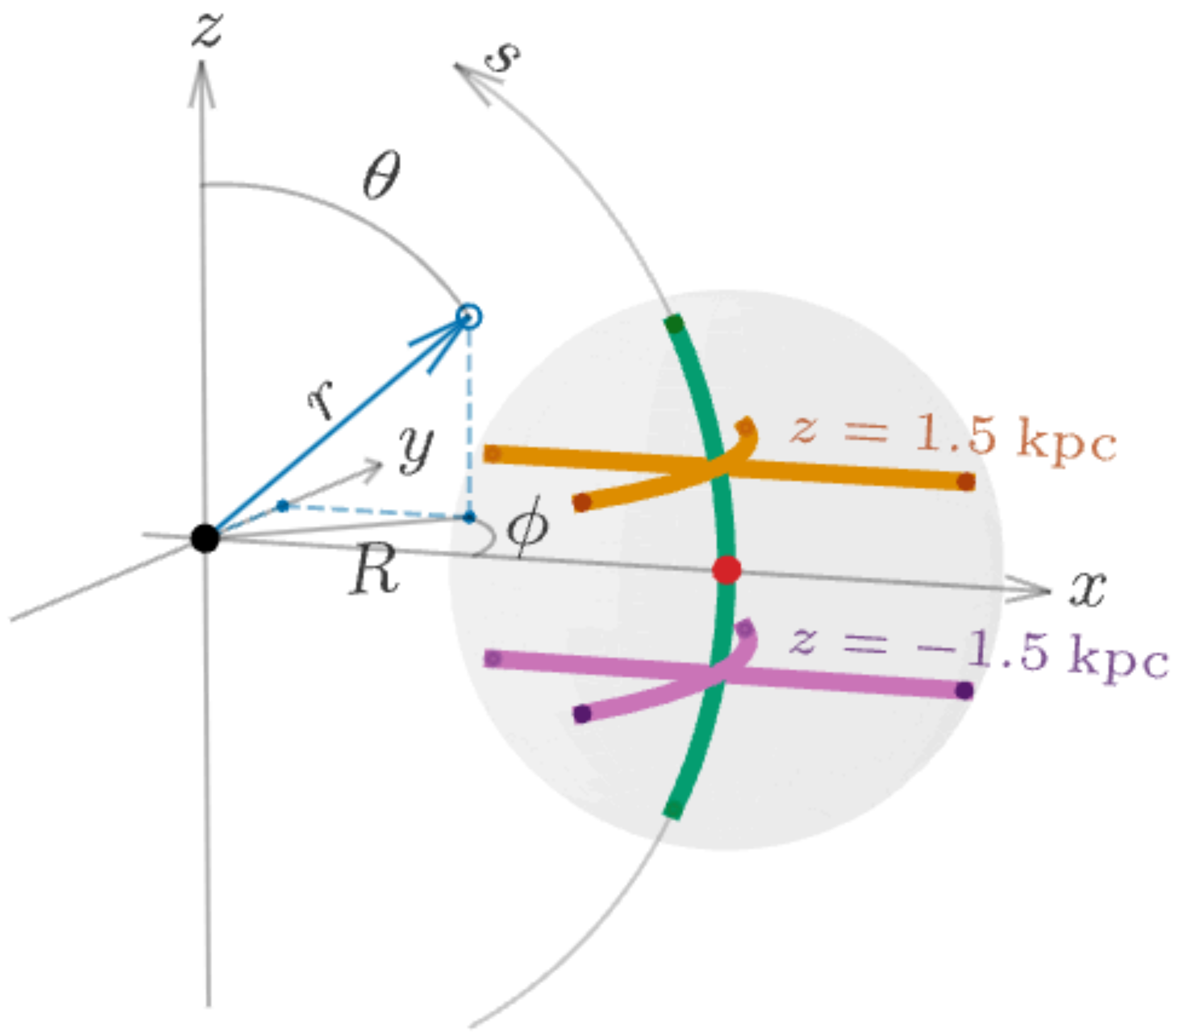

Symmetries to ~10% level:

- north-south
- azimuthal (phi)

=> Expected from dynamical equilibrium

	Gaia EDR3 [56]	**This work**
a_x $(10^{-10}\mathrm{m/s^2})$	-2.32 ± 0.16	-1.94 ± 0.22
a_y $(10^{-10}\mathrm{m/s^2})$	0.04 ± 0.16	0.08 ± 0.08
a_z $(10^{-10}\mathrm{m/s^2})$	-0.14 ± 0.19	-0.06 ± 0.08
$\lvert\vec{a}\rvert$ $(10^{-10}\mathrm{m/s^2})$	2.32 ± 0.16	1.94 ± 0.22

TABLE I: Galactic acceleration at the Solar location $\vec{a}_\odot$ in Cartesian coordinates, calculated by averaging the solution to the Boltzmann equation within a 100 pc sphere centered on the Sun. We list for comparison the acceleration at the Solar location obtained from *Gaia* DR3 quasar measurements [56].

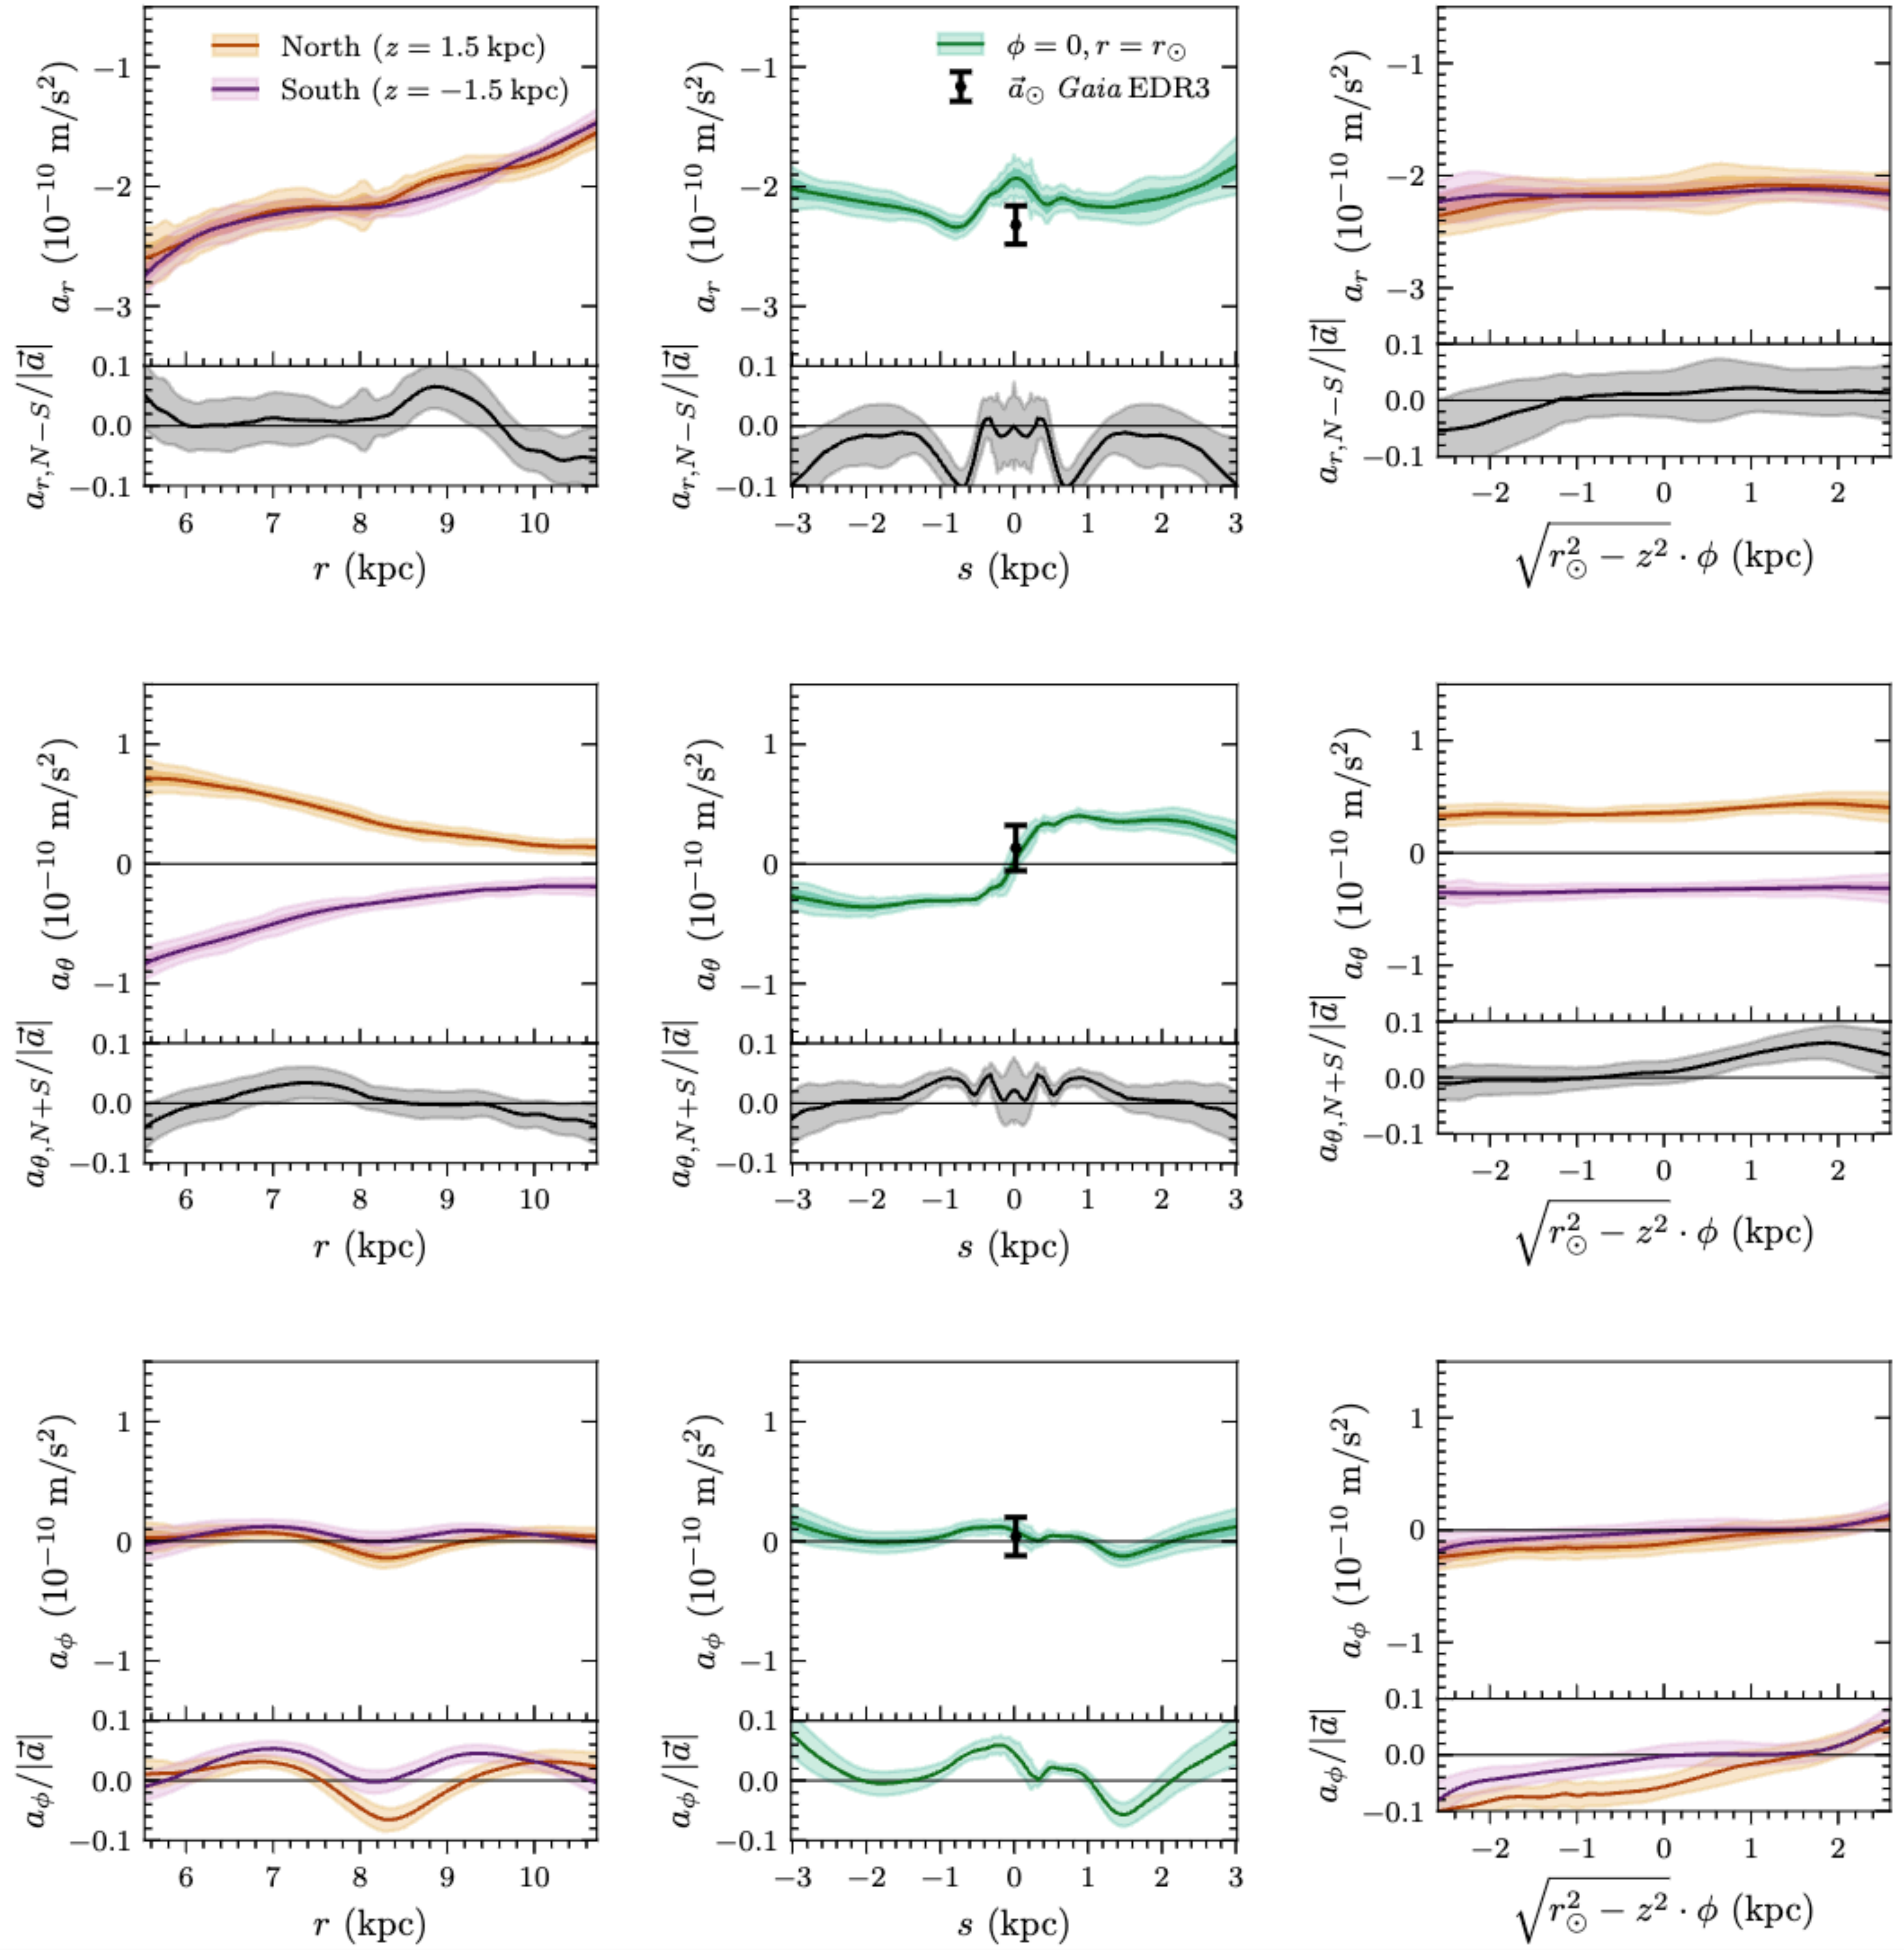

Results: mass density

Lim, Putney, Buckley & **DS** 2305.13358

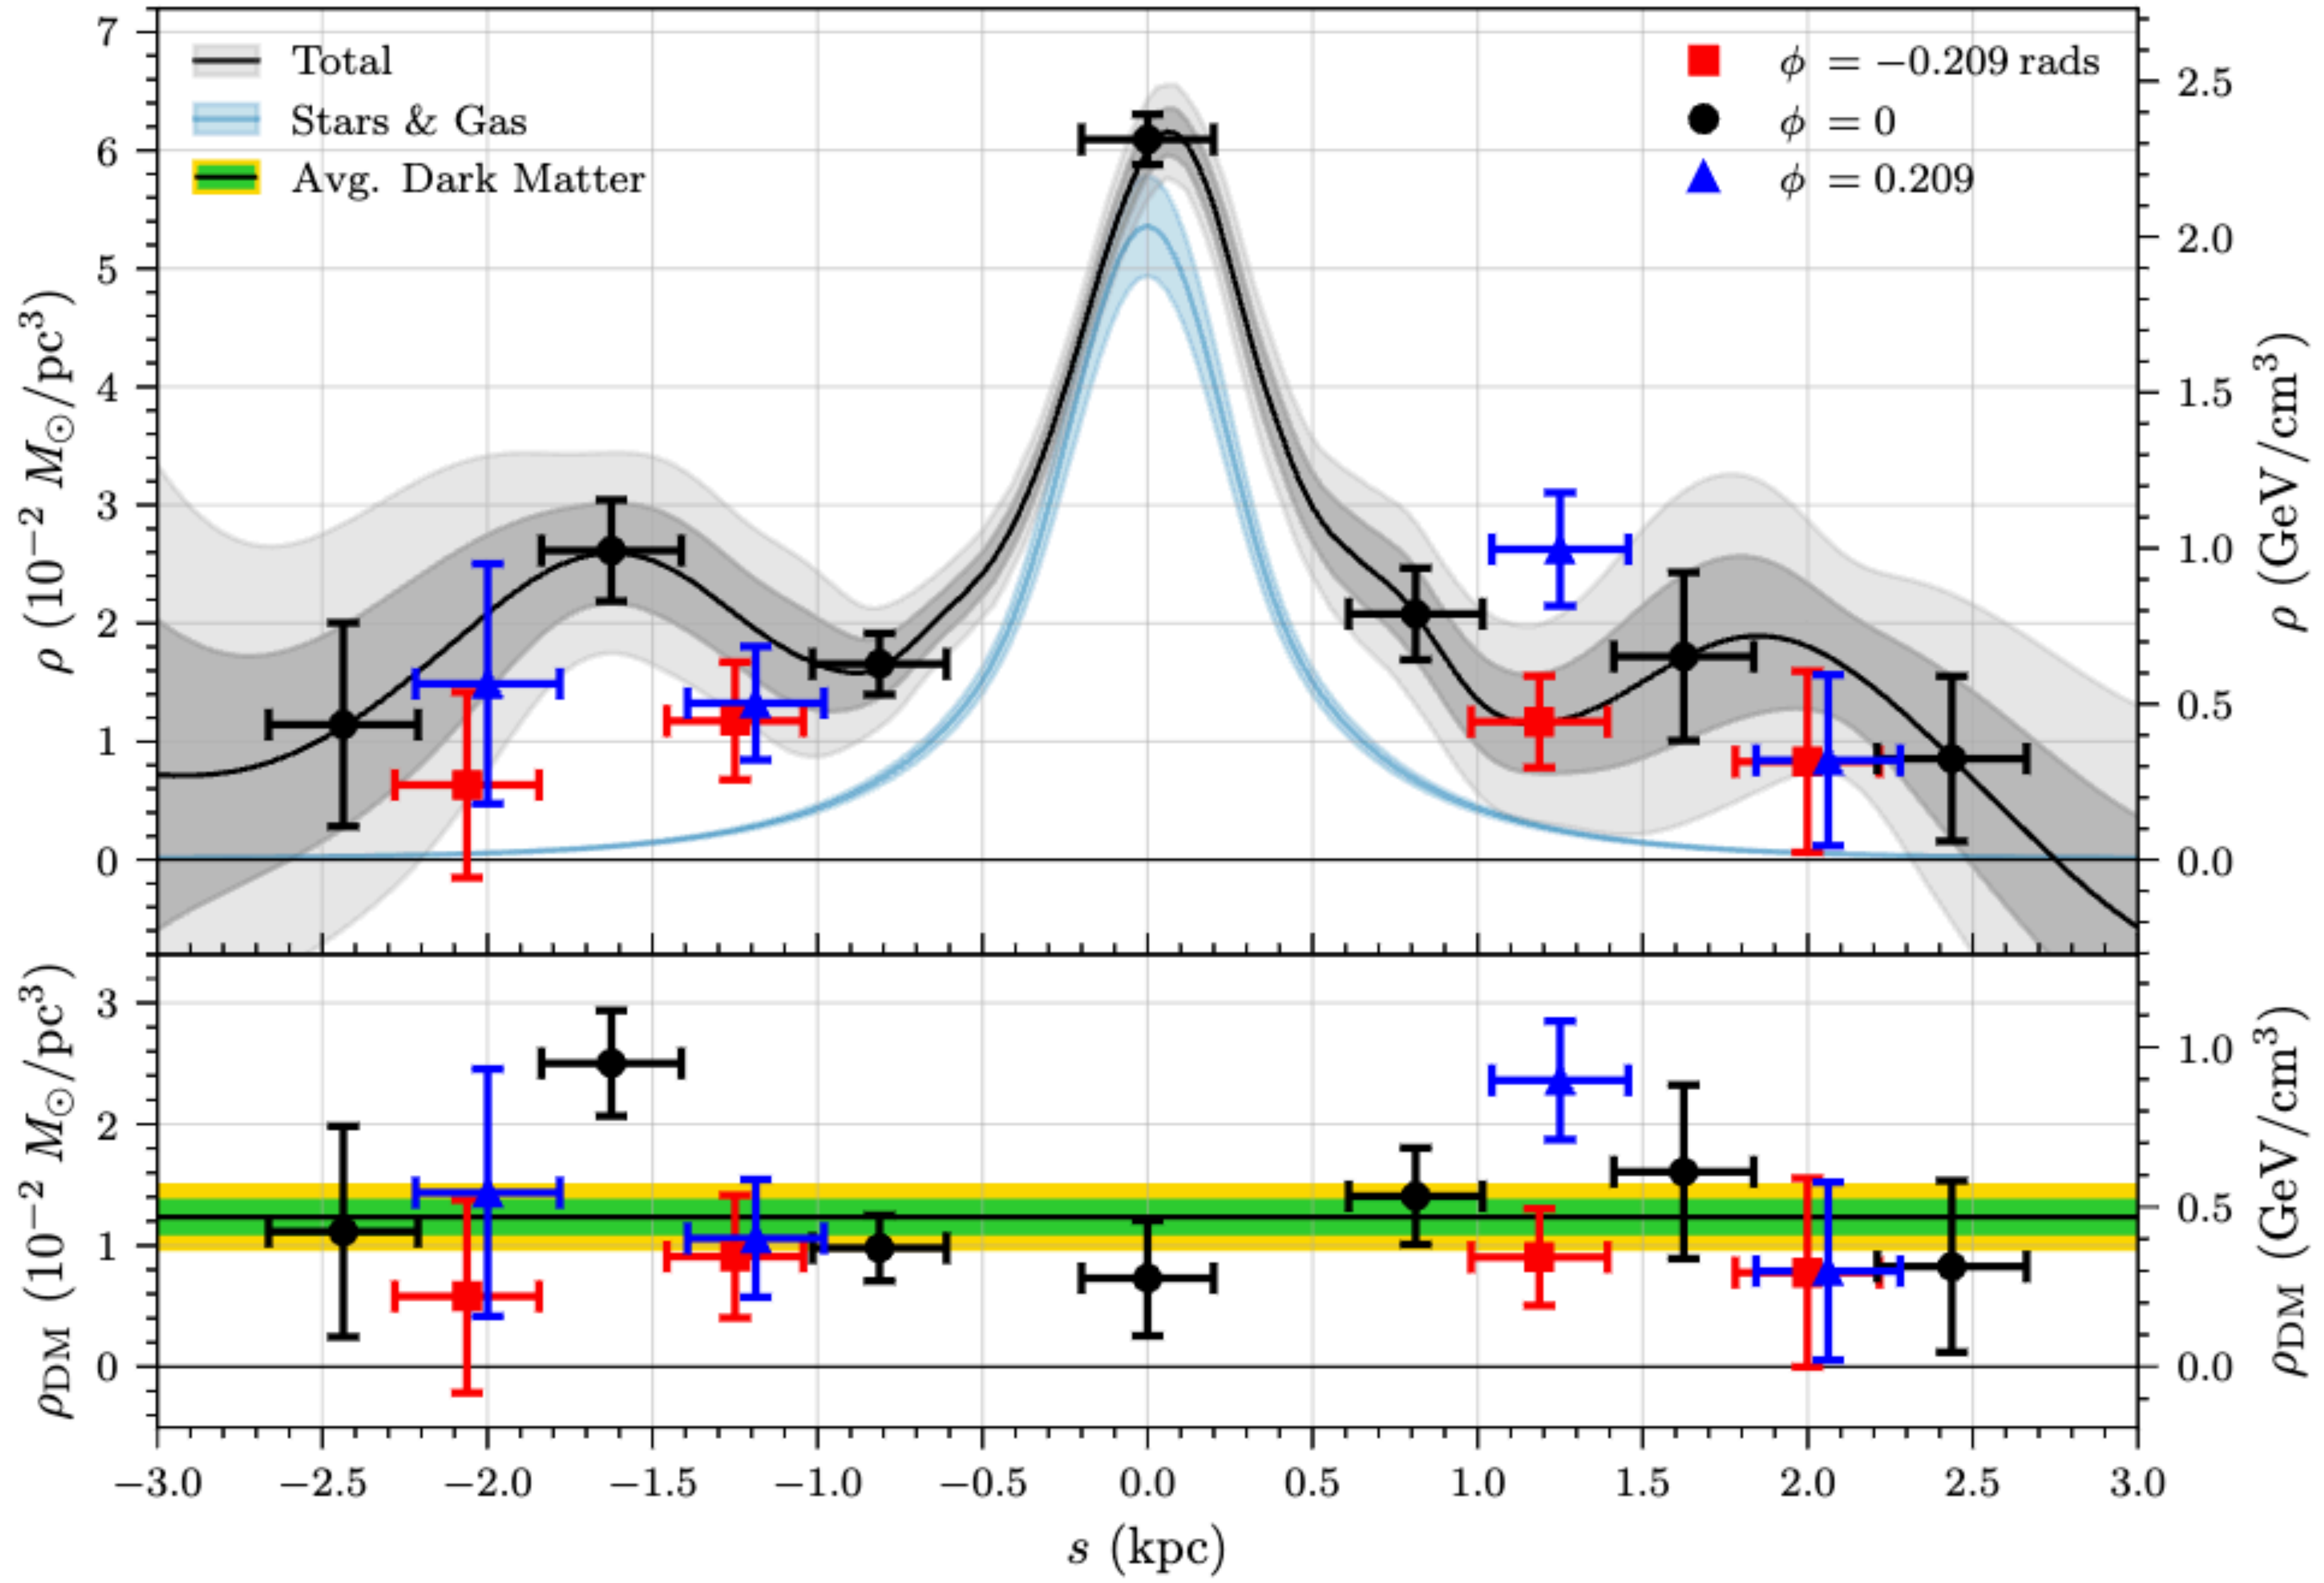

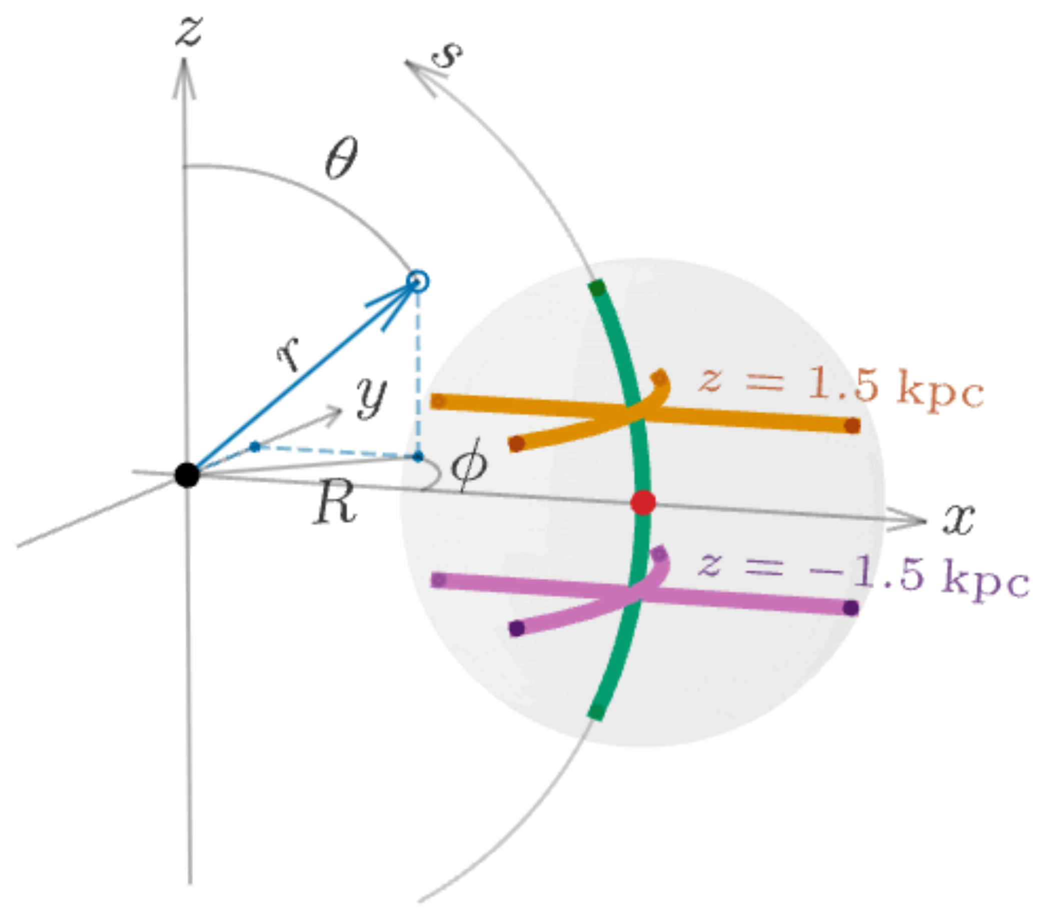

Density	(10^{-2} $M_\odot$/pc^3)	(GeV/cm^3)	χ^2_ν
$\rho_\odot$	6.17 ± 0.20	2.34 ± 0.08	
$\rho_{b,\odot}$	5.34 ± 0.42	2.03 ± 0.16	
$\rho_{\rm DM,\odot}$	0.83 ± 0.47	0.32 ± 0.18	
$\overline{\rho}_{\rm DM}(r = r_\odot)$	1.18 ± 0.14	0.47 ± 0.05	1.38

Result is consistent with nonzero, spherically symmetric DM density!

Error bars include:

- MAF training variance
- Gaia measurement error
- Finite training statistics

Results: mass density

Lim, Putney, Buckley & **DS** 2305.13358

Our result:

$$\rho_{DM}(r_\odot) = 0.47 \pm 0.05\ \mathrm{GeV/cm^3}$$

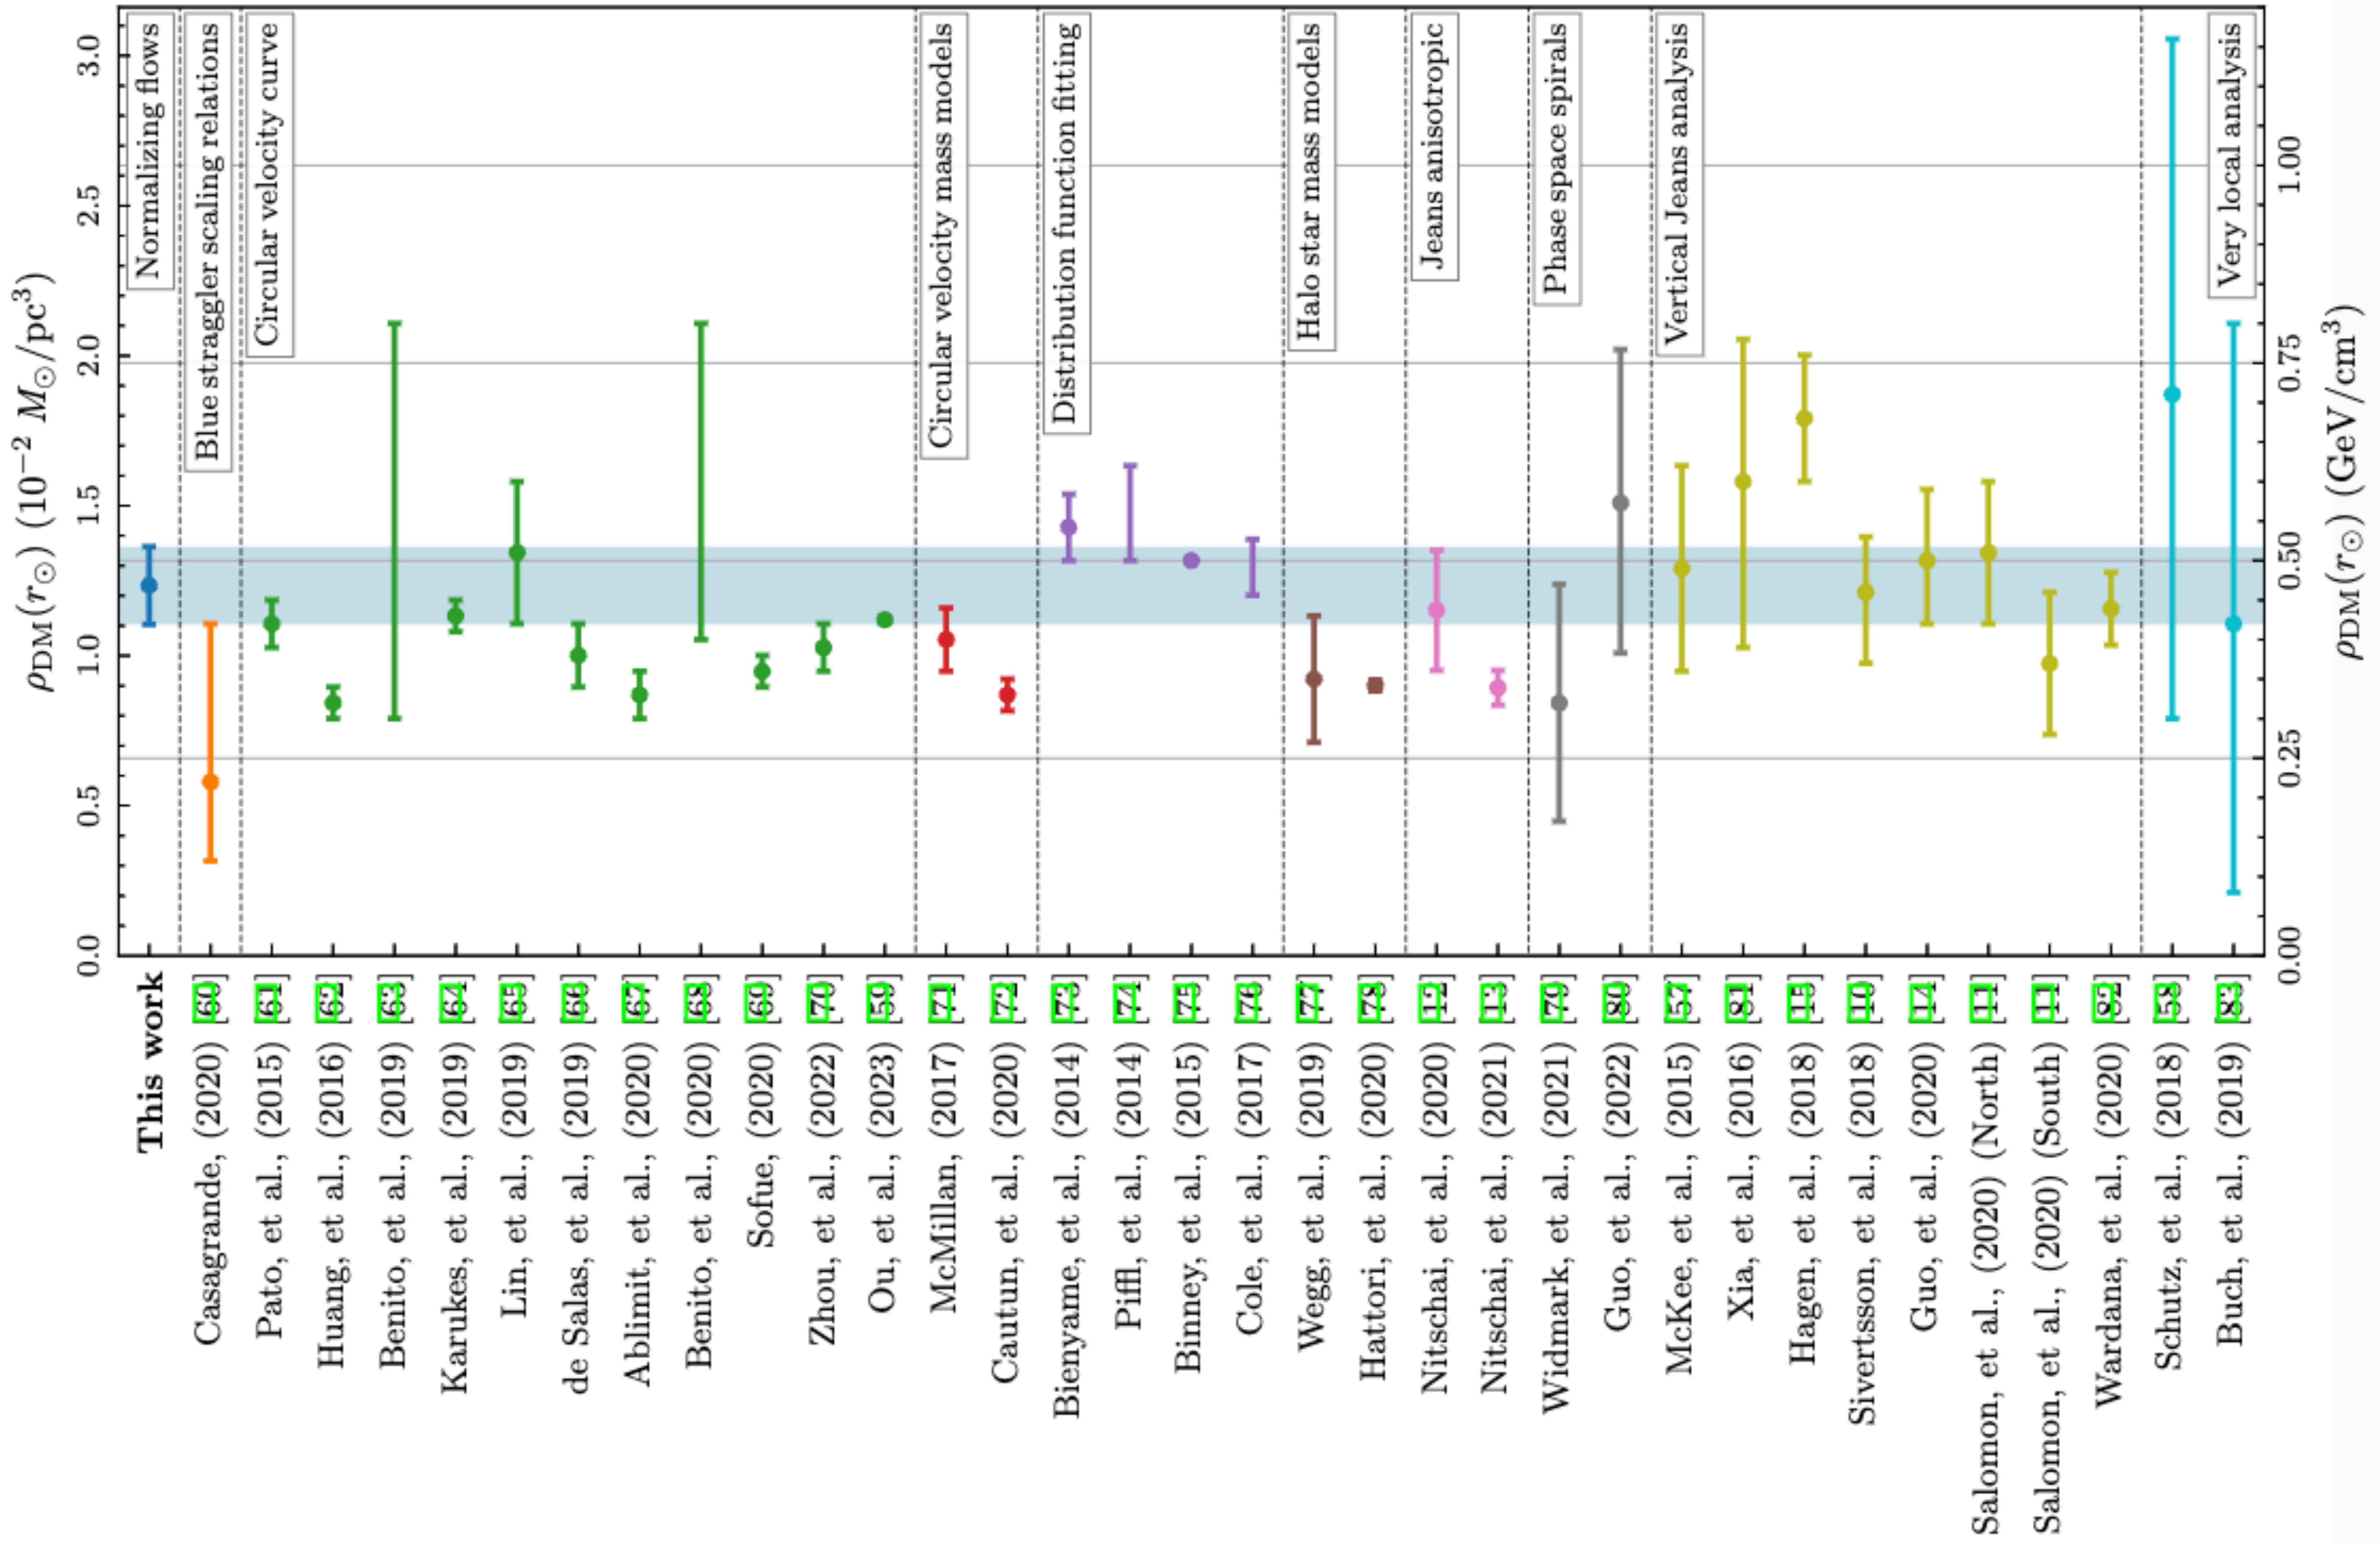

Excellent agreement with previous measurements, with hopefully more realistic error bars

Results: mass density

Lim, Putney, Buckley & **DS** 2305.13358

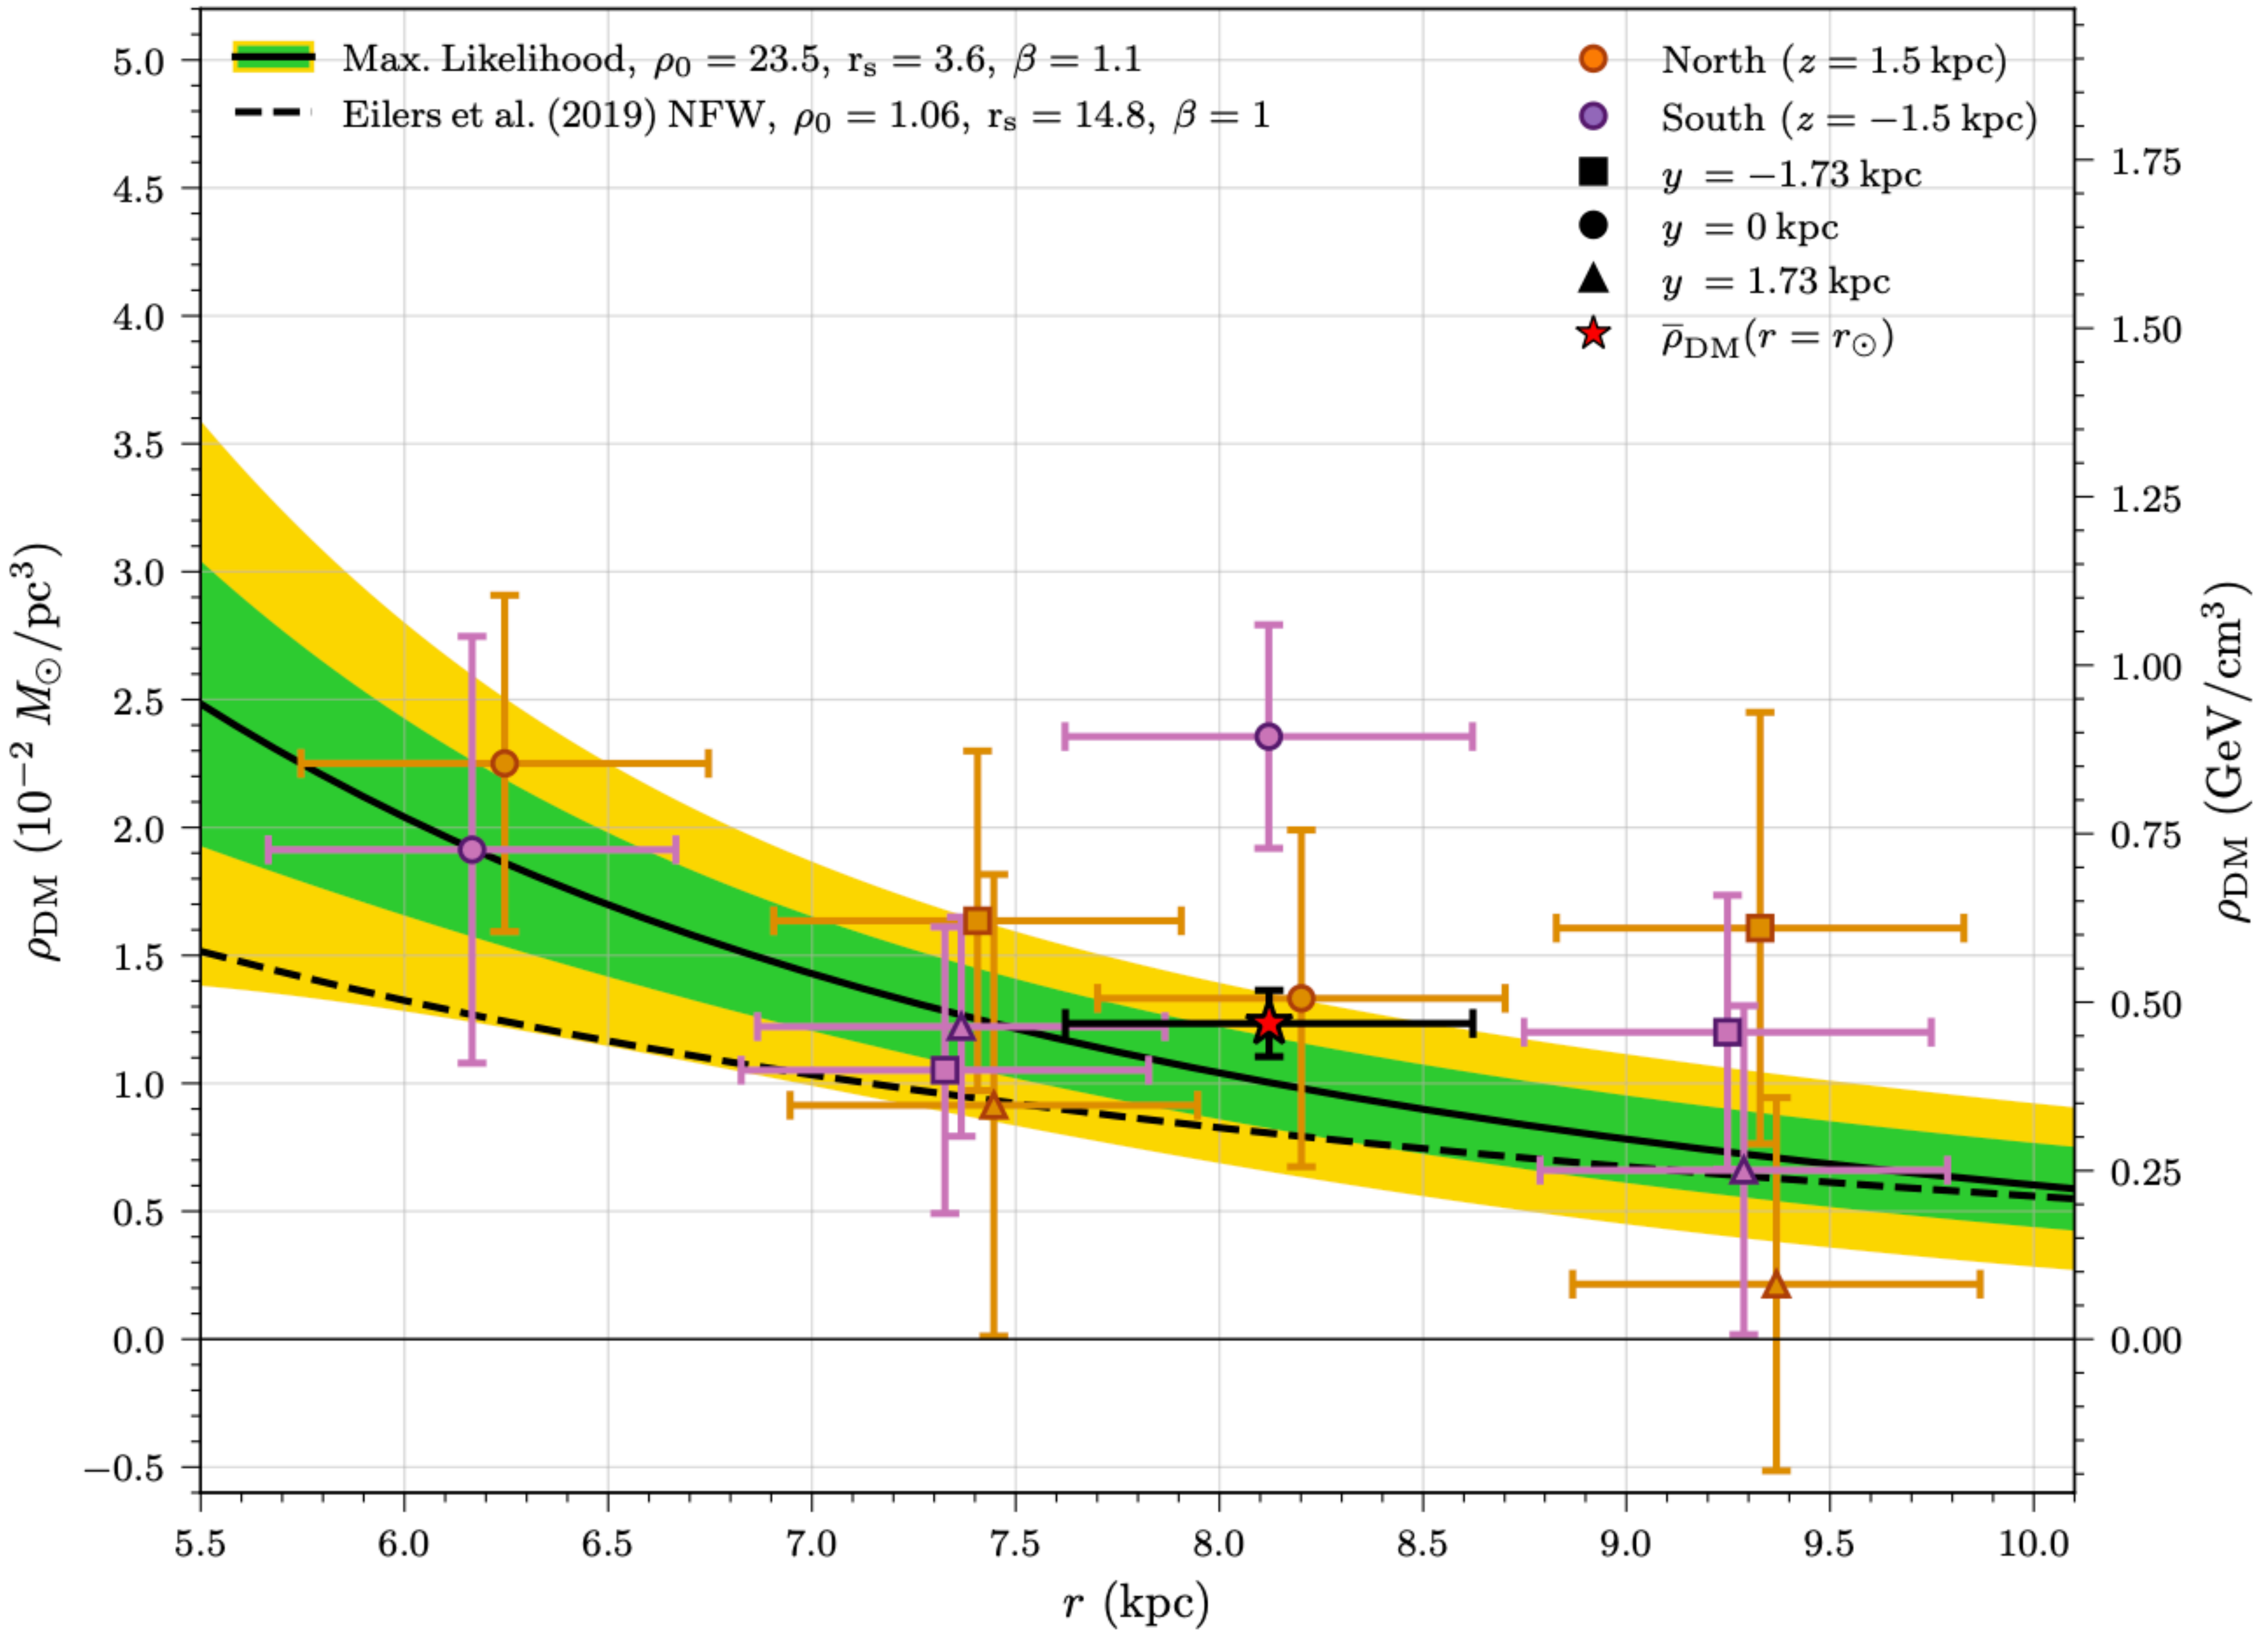

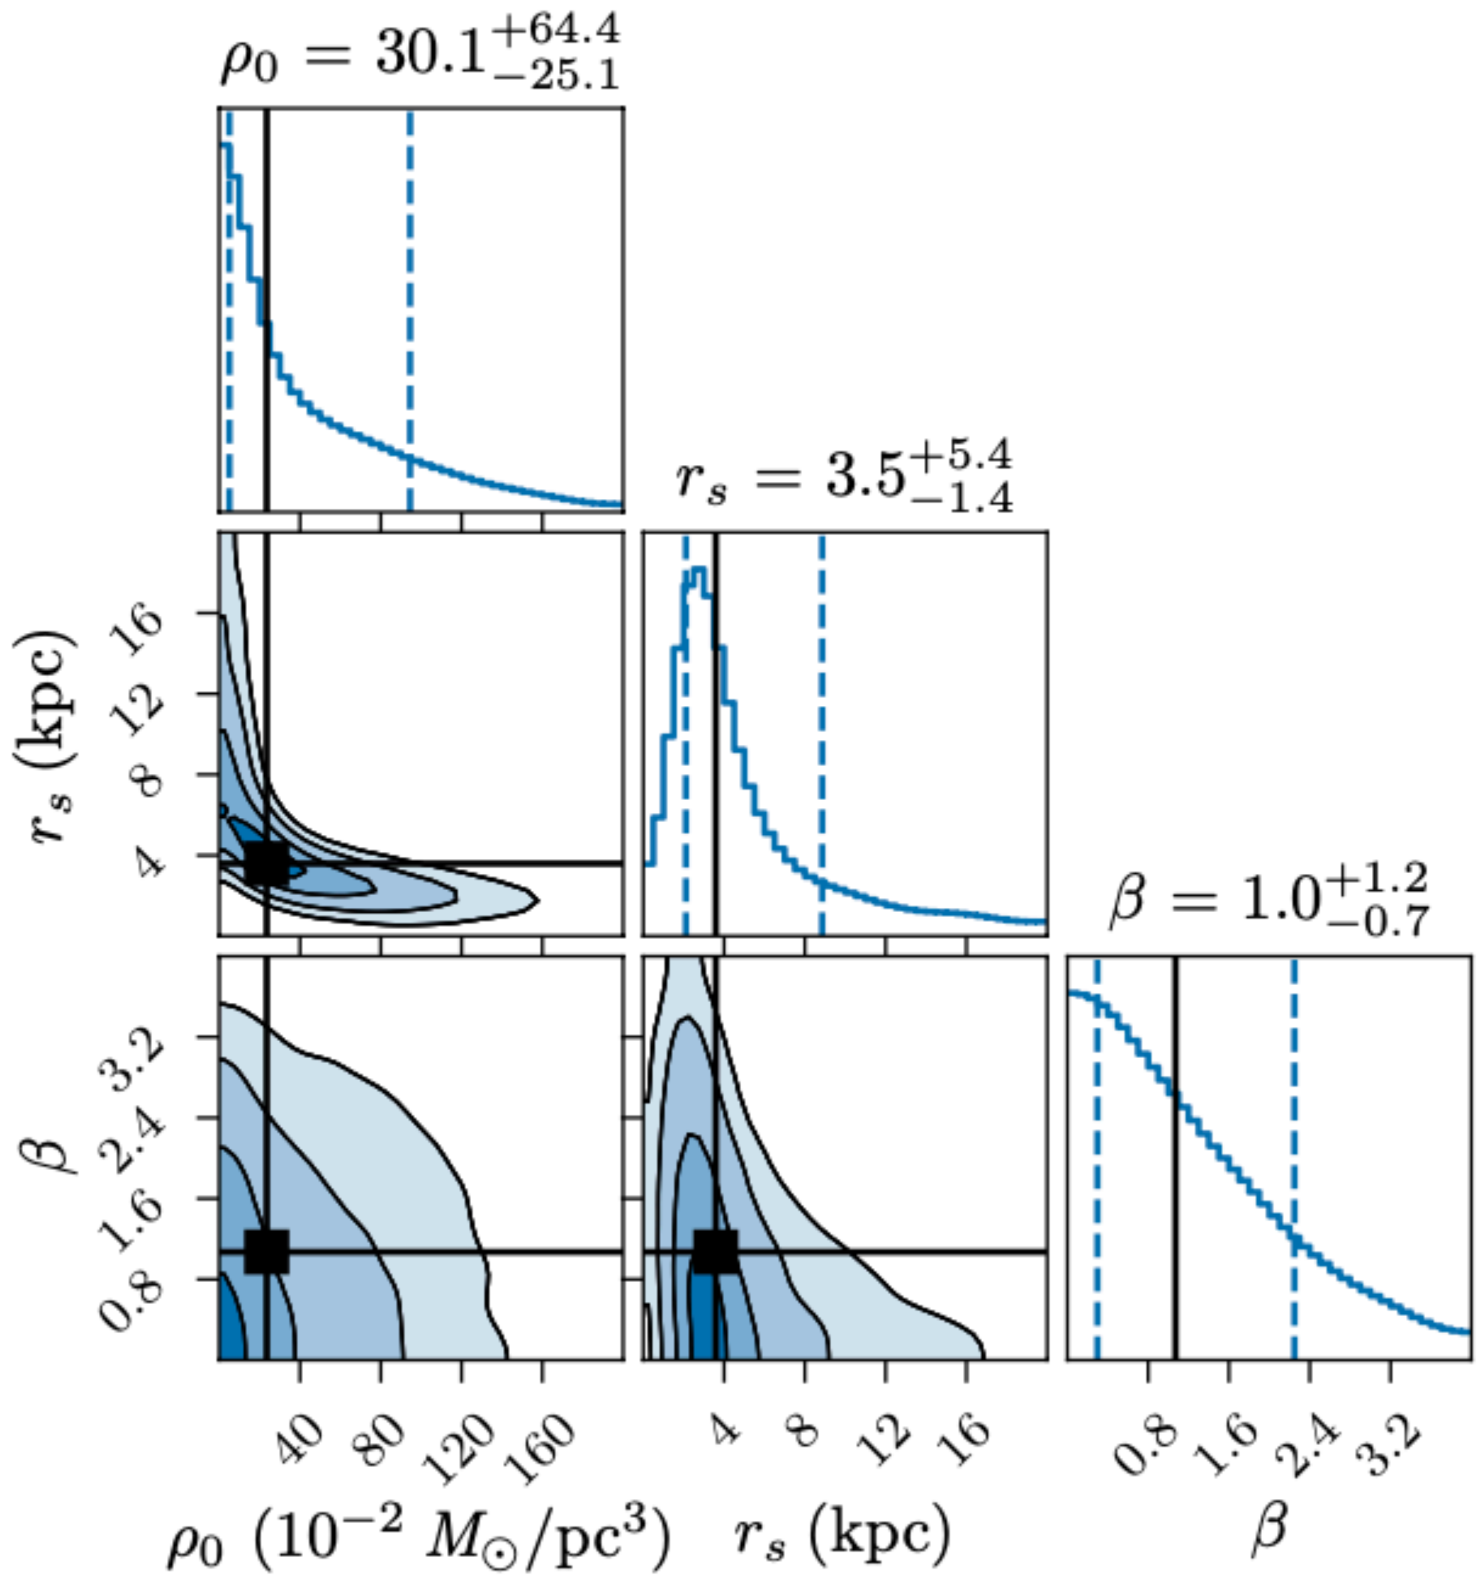

Radial profile broadly consistent with recent NFW fits

Summary and Outlook

- While countless searches for new physics have been performed at the LHC, nearly all of them are highly model-specific.
- This represents a **huge opportunity** for a new paradigm of model-agnostic search strategies.
- Motivated in part by community data challenges (LHCO2020, DarkMachines, ADC2022), theorists, experimentalists (and others!) have developed many new and exciting model-agnostic methods using the tools of modern ML.
- Some of these ideas are beginning to be ported over to ATLAS and CMS and implemented as actual analyses on real data, but much work remains to be done!
- Methods are also being ported over to the Astro domain — highlights the cross-cutting power of ML tools!

Thanks for your attention!