Turbulence and transport in, and optimization of, LAPD mirror configurations

Phil Travis
Thesis defense
5/30/2025

Committee: Troy Carter (Chair), Jacob Bortnik, Chris Niemann, Paulo Alves

My goal: get fusion faster with mirrors machines and machine learning (ML)

- Interested in mirrors because of their mechanical simplicity (some downsides)
- Study turbulence and transport — important for all fusion devices
- Work towards automating fusion science; we can use ML to:
  - Optimize plasmas
  - Infer trends
  - Extract insight (by interrogating models)

Mirror machines operate via conservation of magnetic moment and are intrinsically unstable

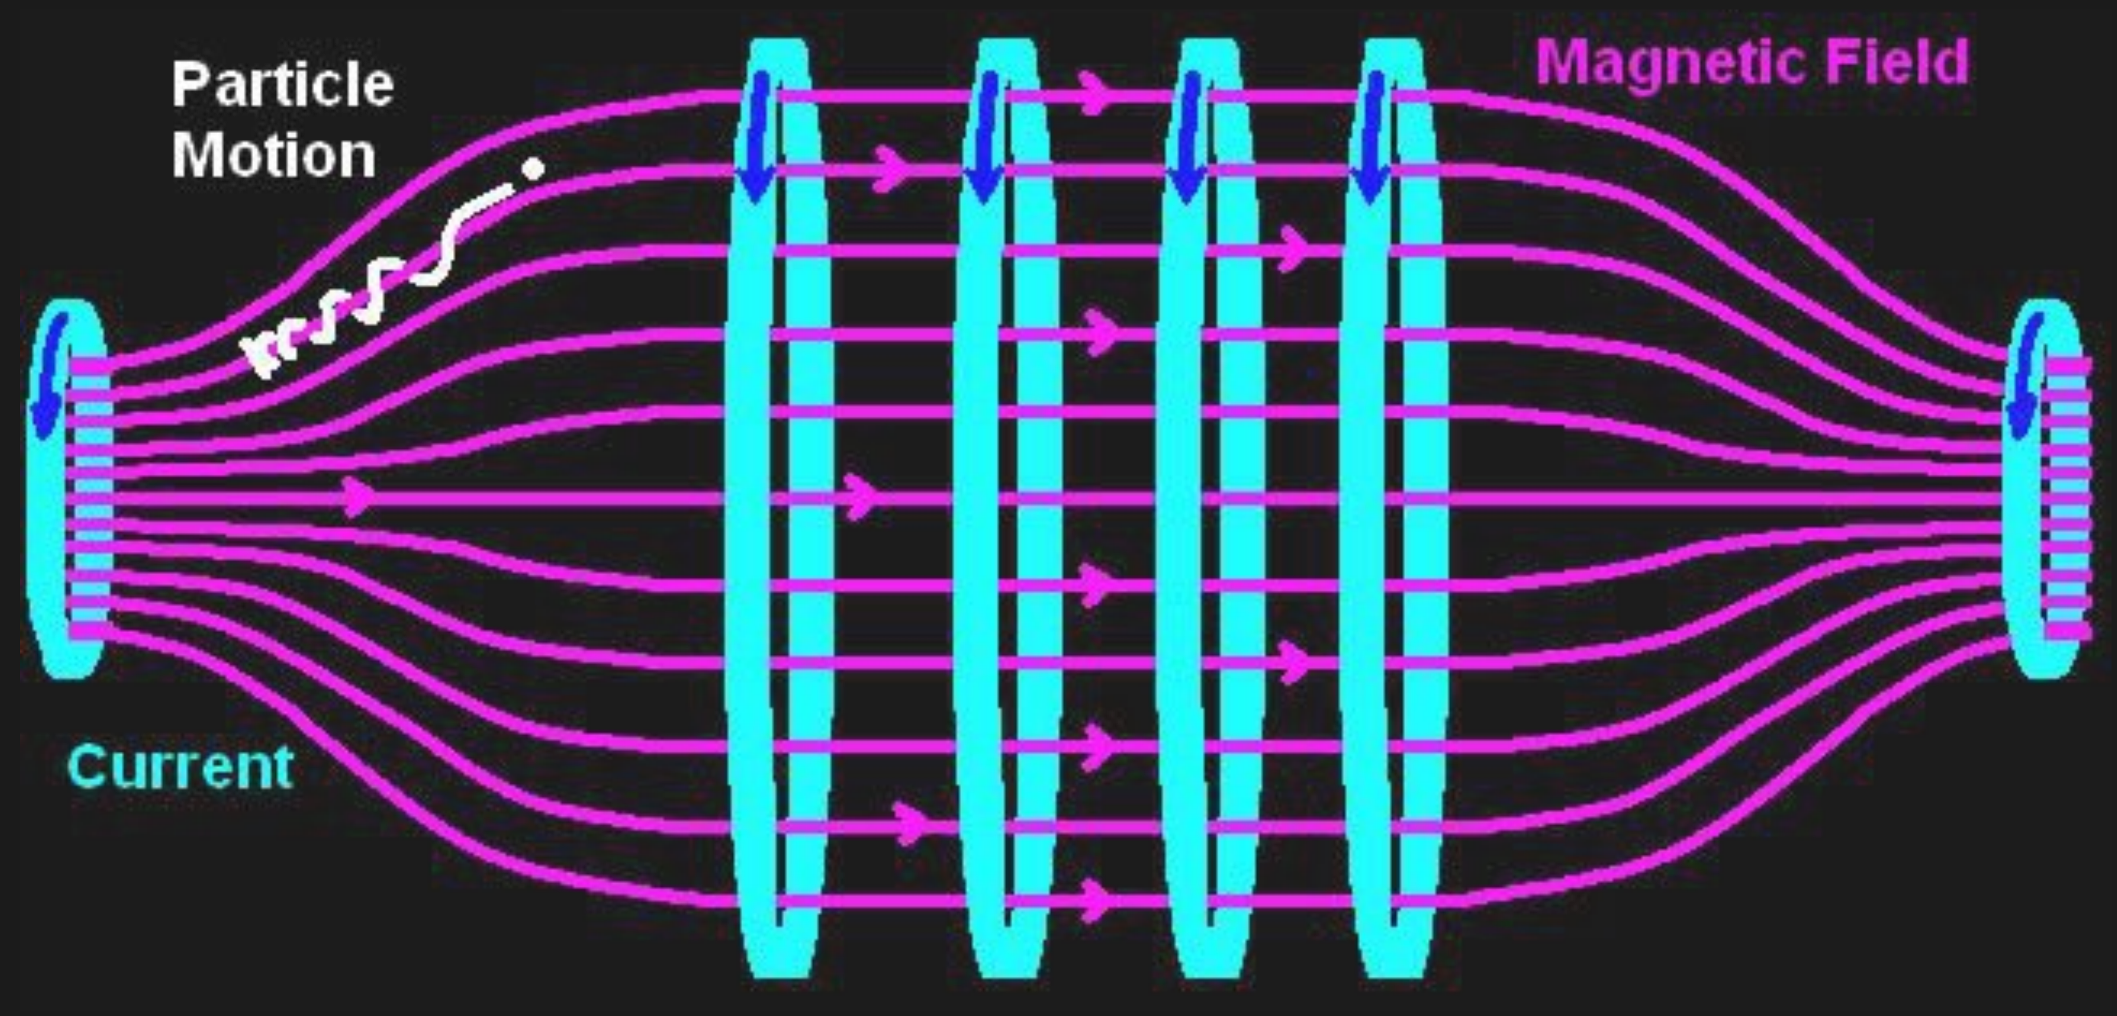

Dr . Moynihan

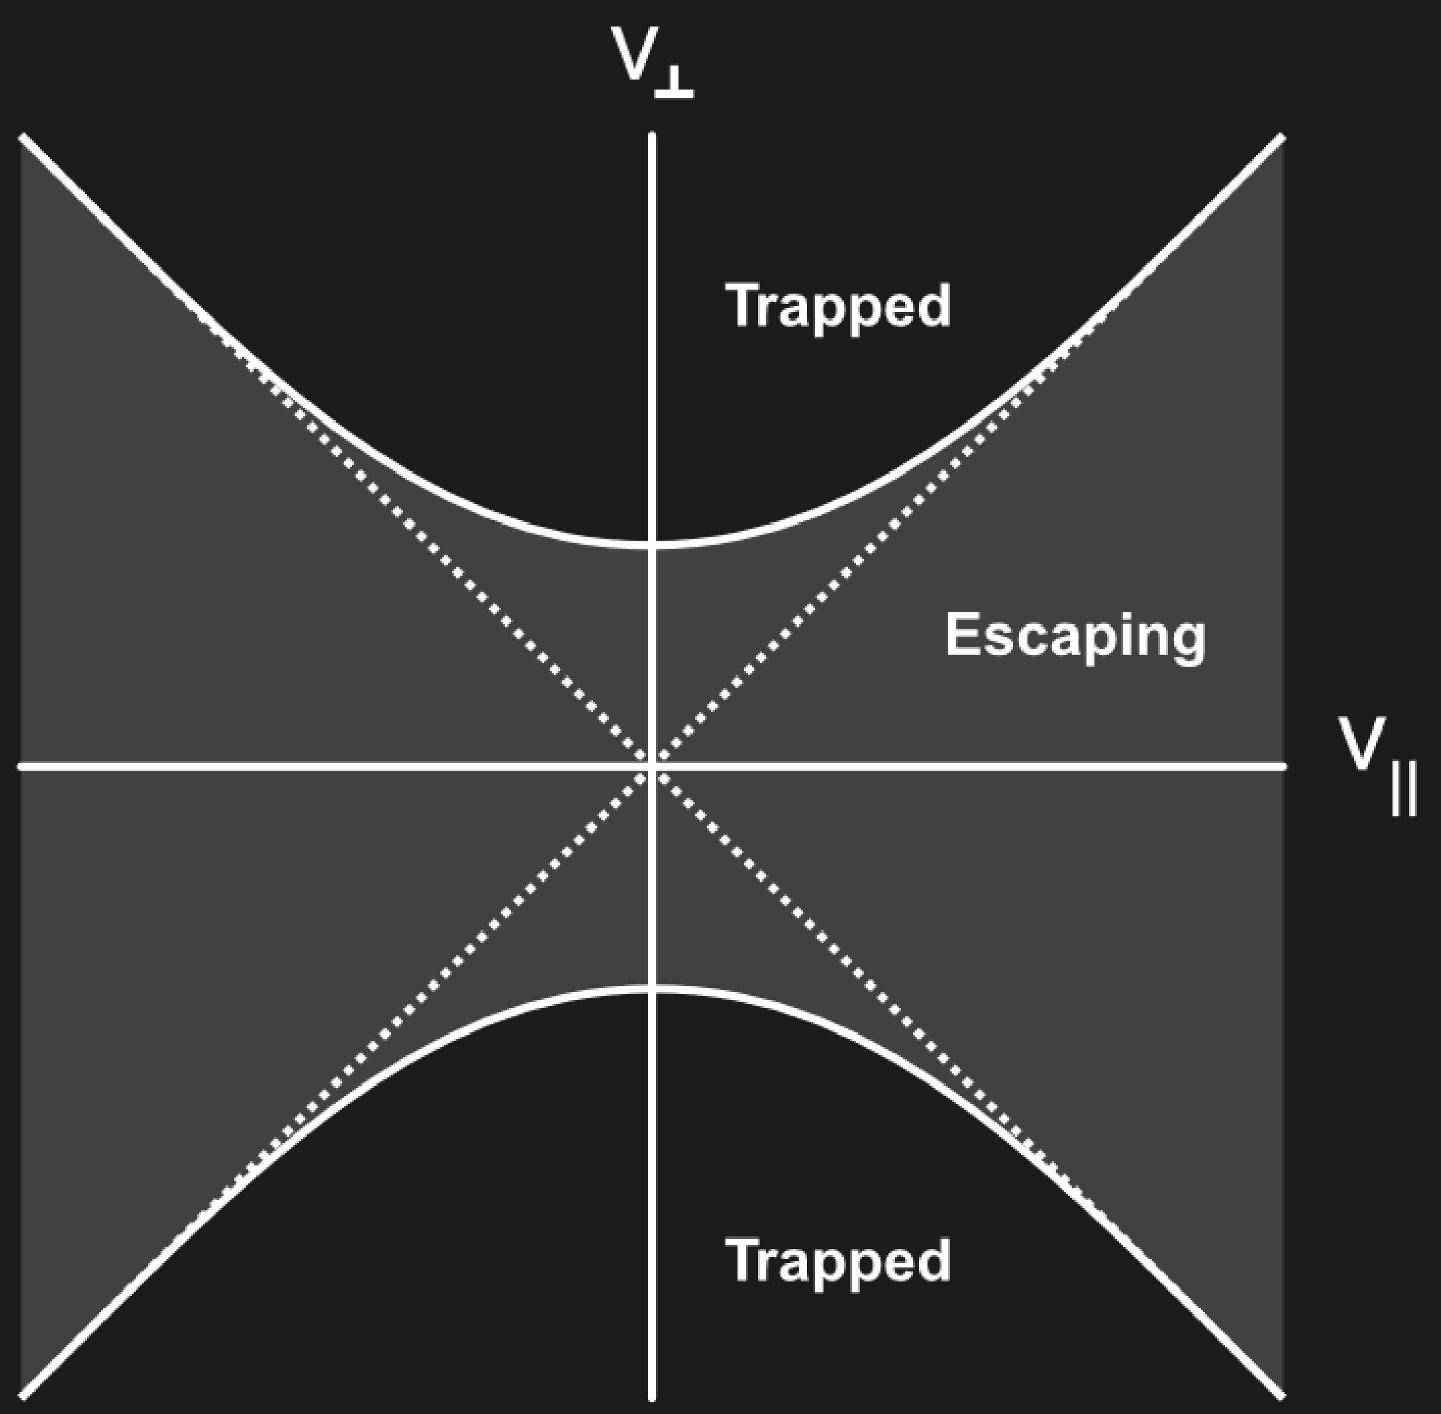

Conservation of magnetic moment: $\mu = \frac{W_\perp \uparrow}{B \uparrow}$ → $W_{||} \downarrow$ → particle reflected

conservation of energy

Mirrors suffer from the interchange instability

- Interchange instability: pressure gradient in same direction as the curvature vector
  - Historical focus of research: stabilize interchange and loss-cone instabilities
- Other instabilities are present regardless

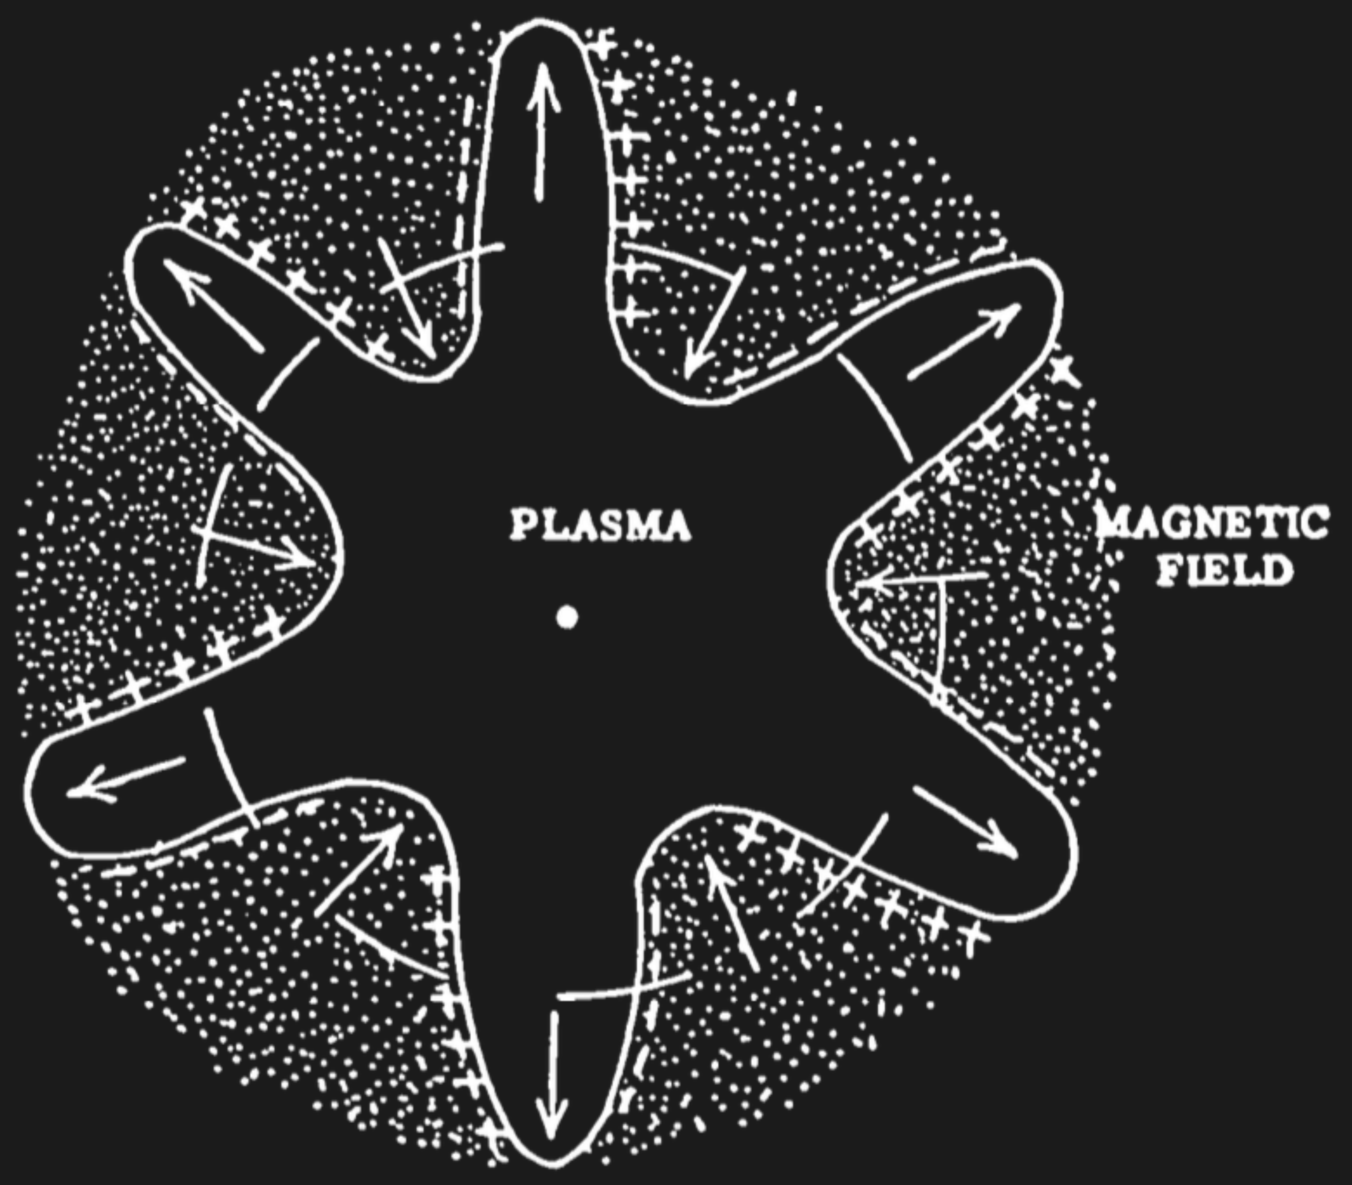

FIG. 30. Schematic illustration of the high-m flute instability in a simple mirror field, showing polarization fields and directions of unstable motion.

Post 1987

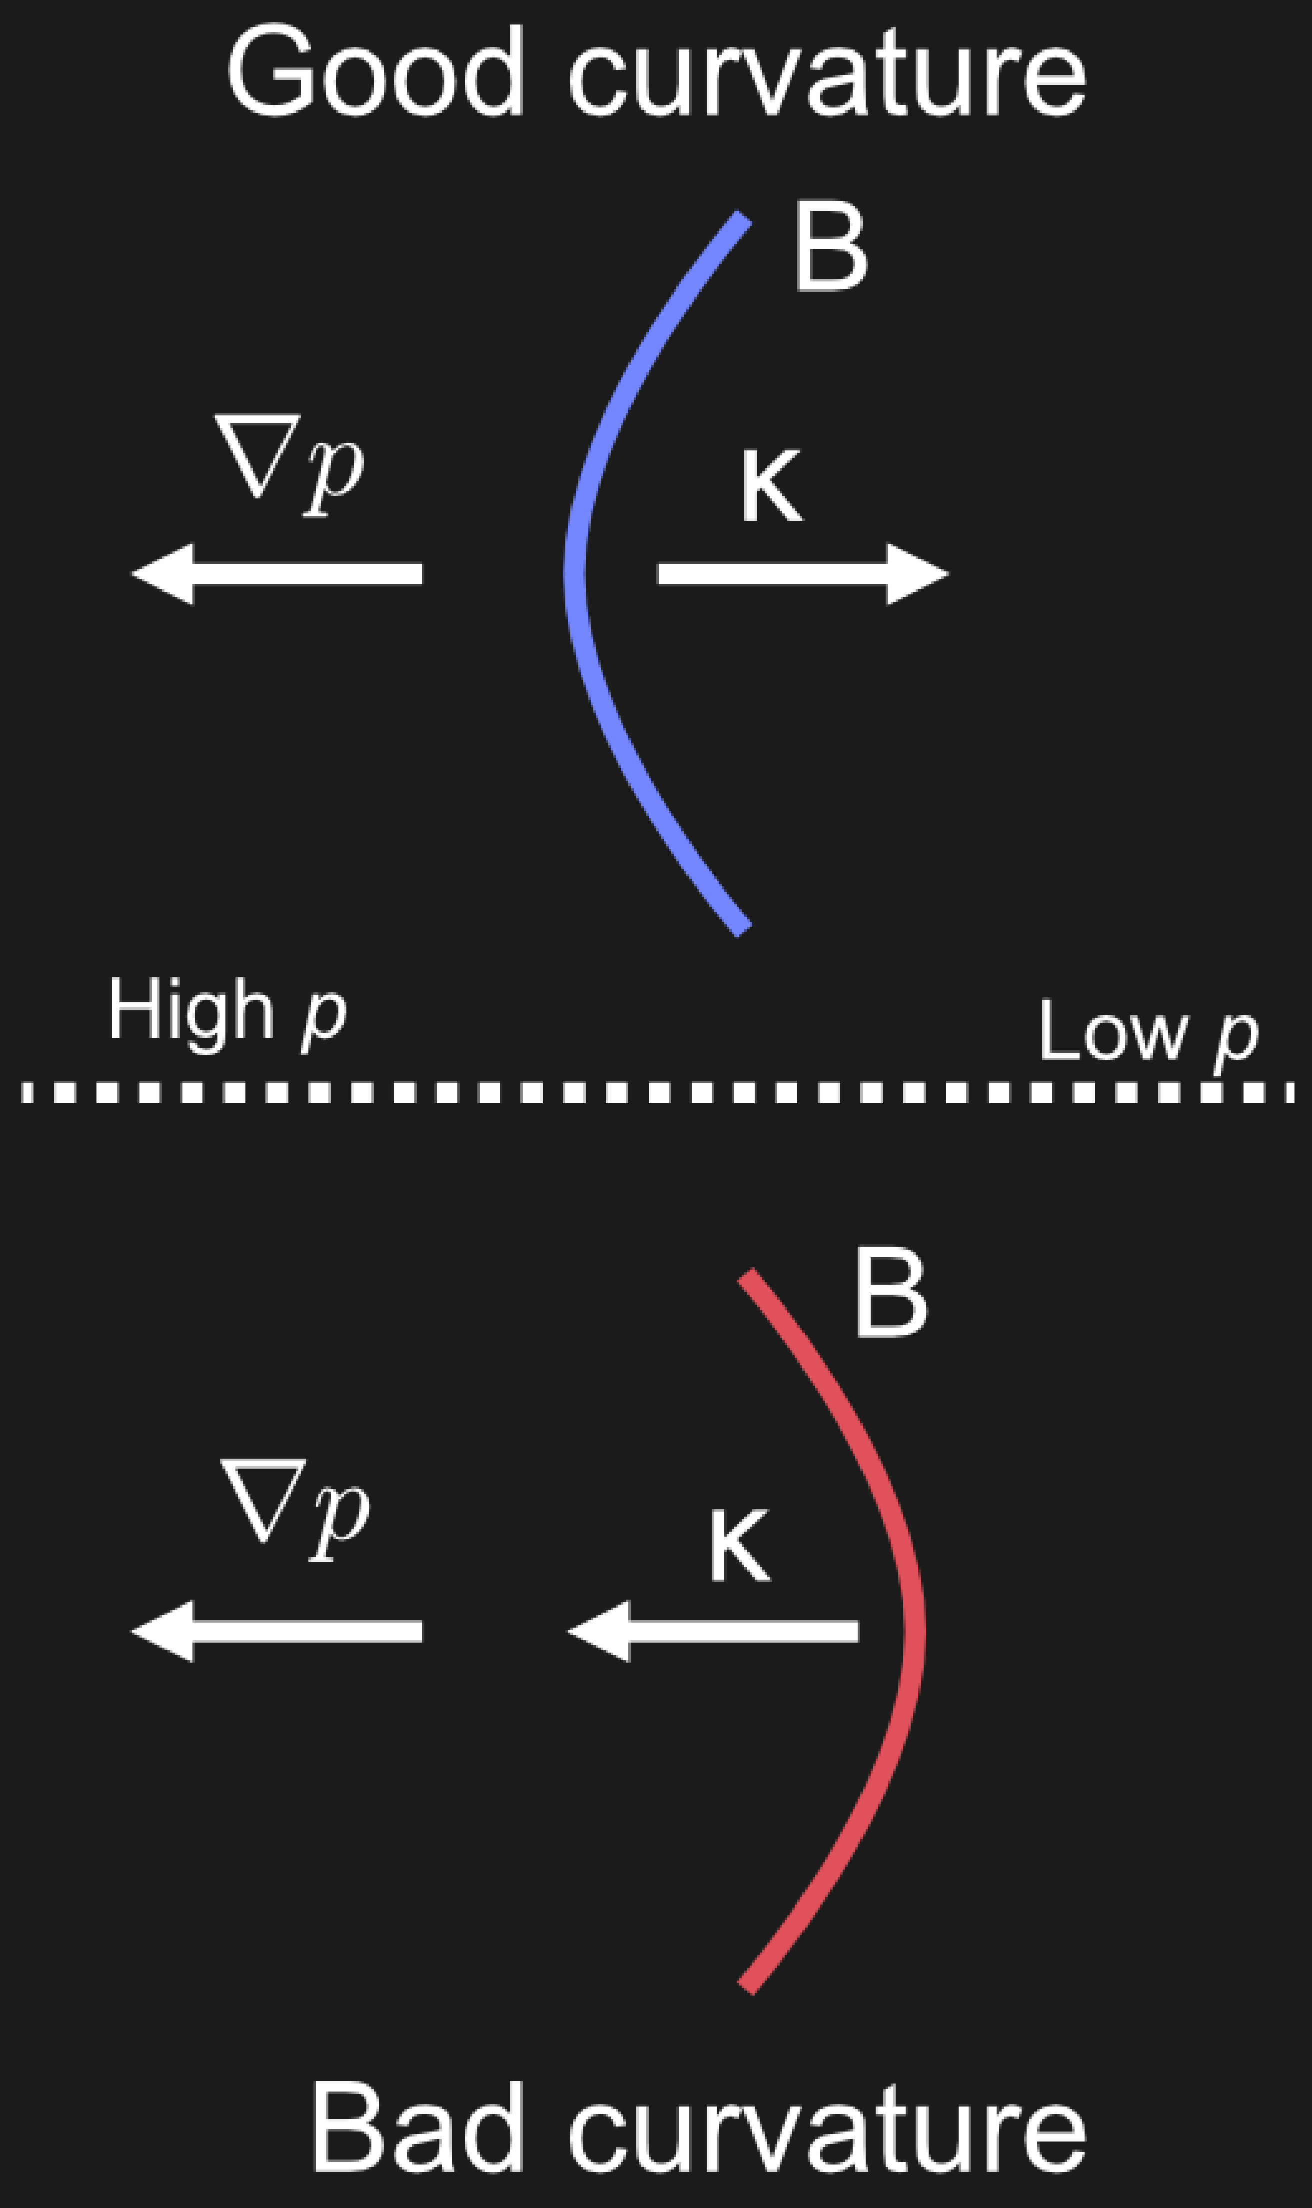

Drift waves and turbulence are ubiquitous in fusion plasmas

- Drift waves are unstable when there exists a density gradient, a background field, and finite resistivity
- “Universal” instability — see in any laboratory plasma
- Instabilities drive turbulence

—> seen in any thermal fusion plasma

- Do drift waves interact with interchange modes?

—> study on the Large Plasma Device

The Large Plasma Device (LAPD) is a flexible, accessible, basic plasma science device

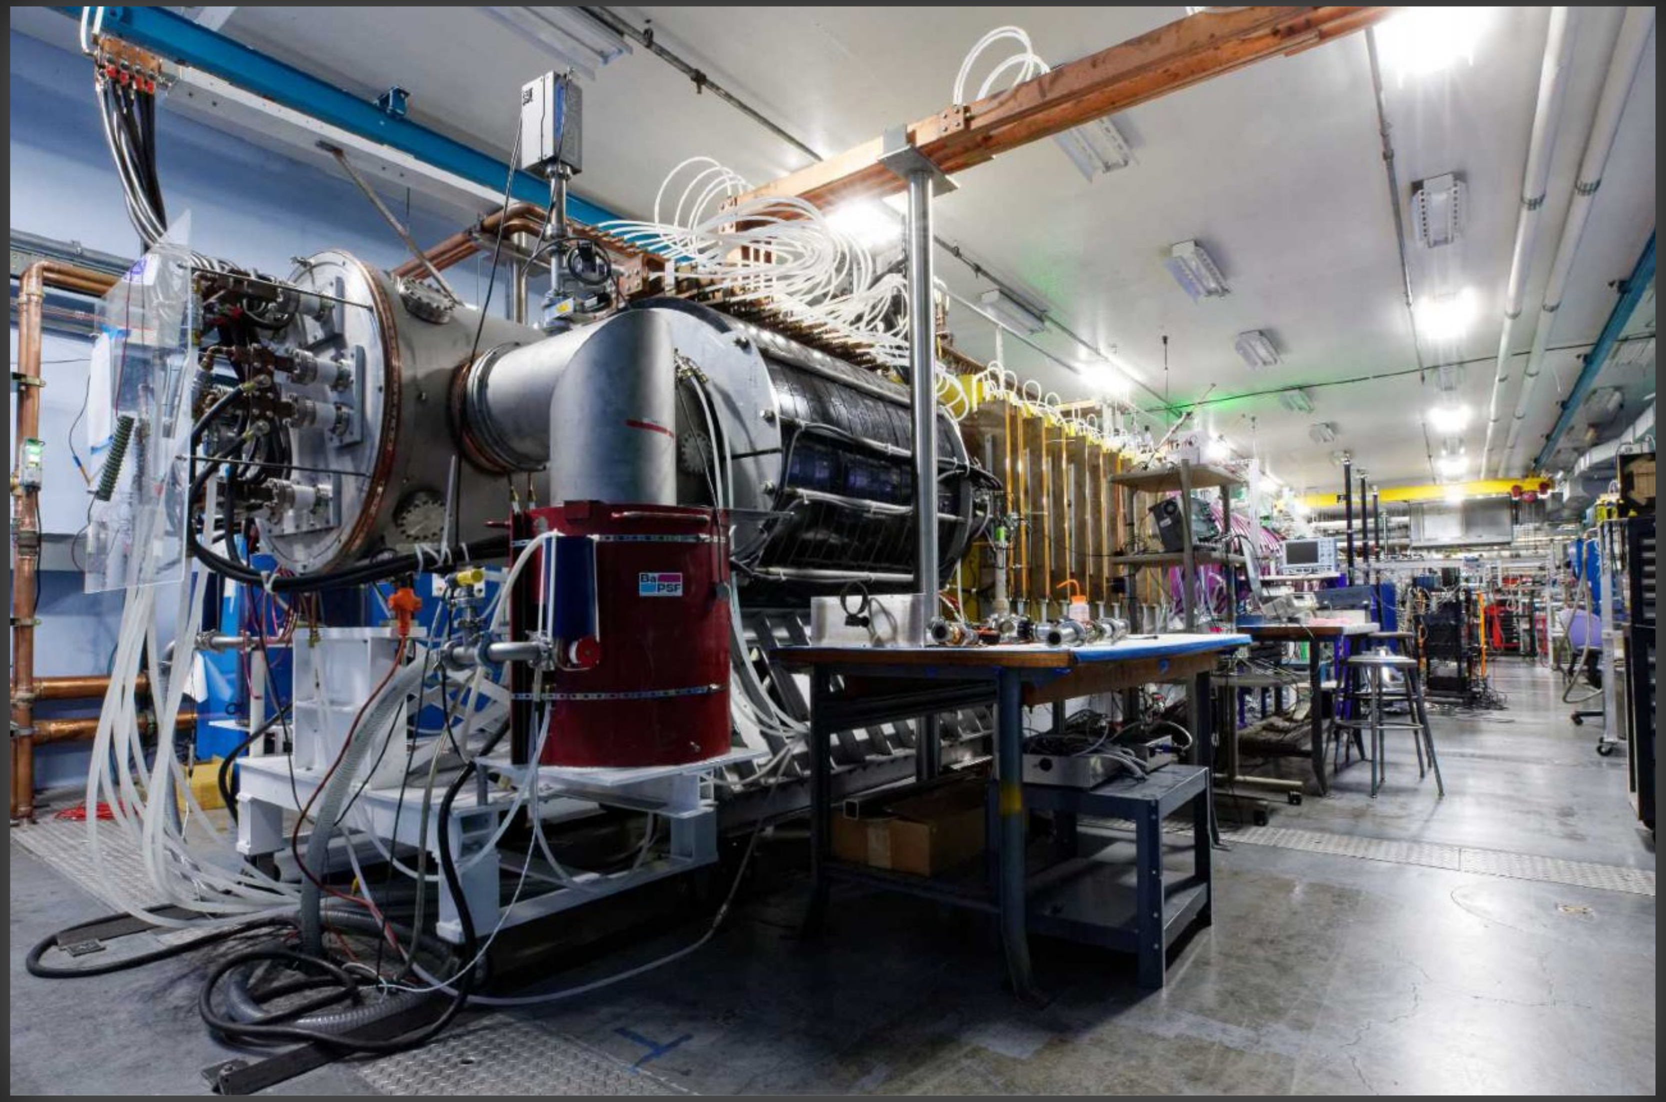

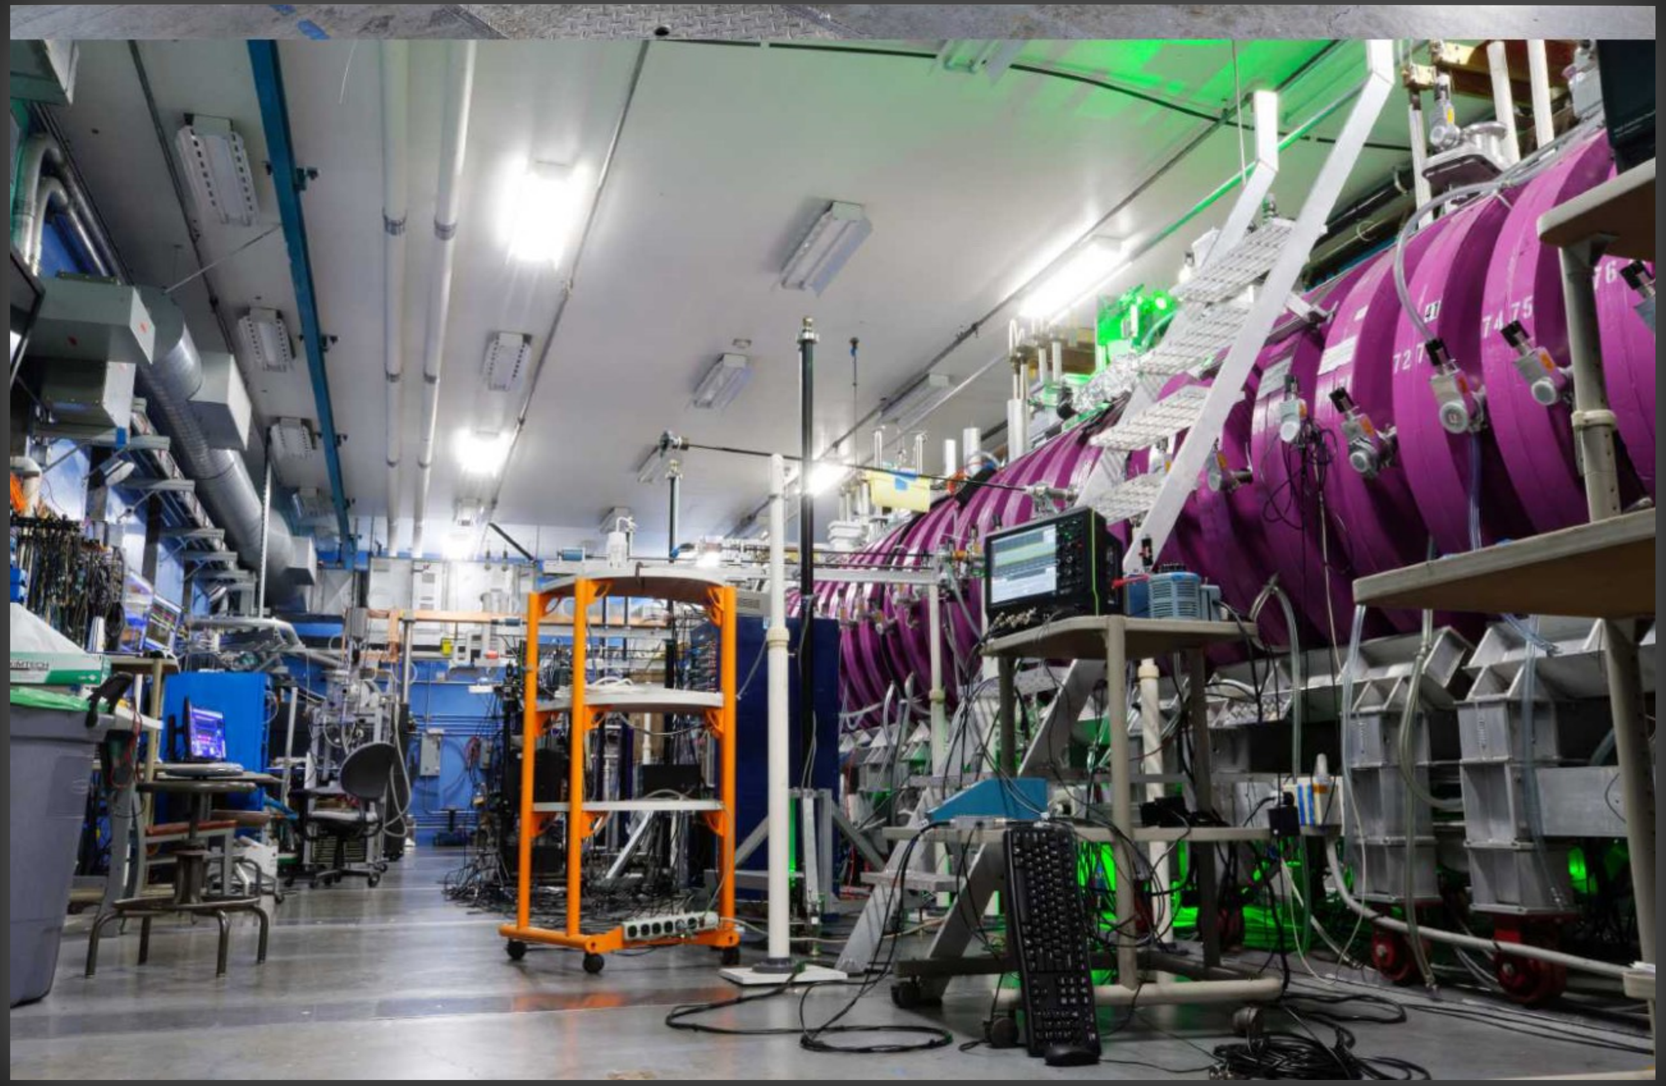

very length

such bright

wow

wow

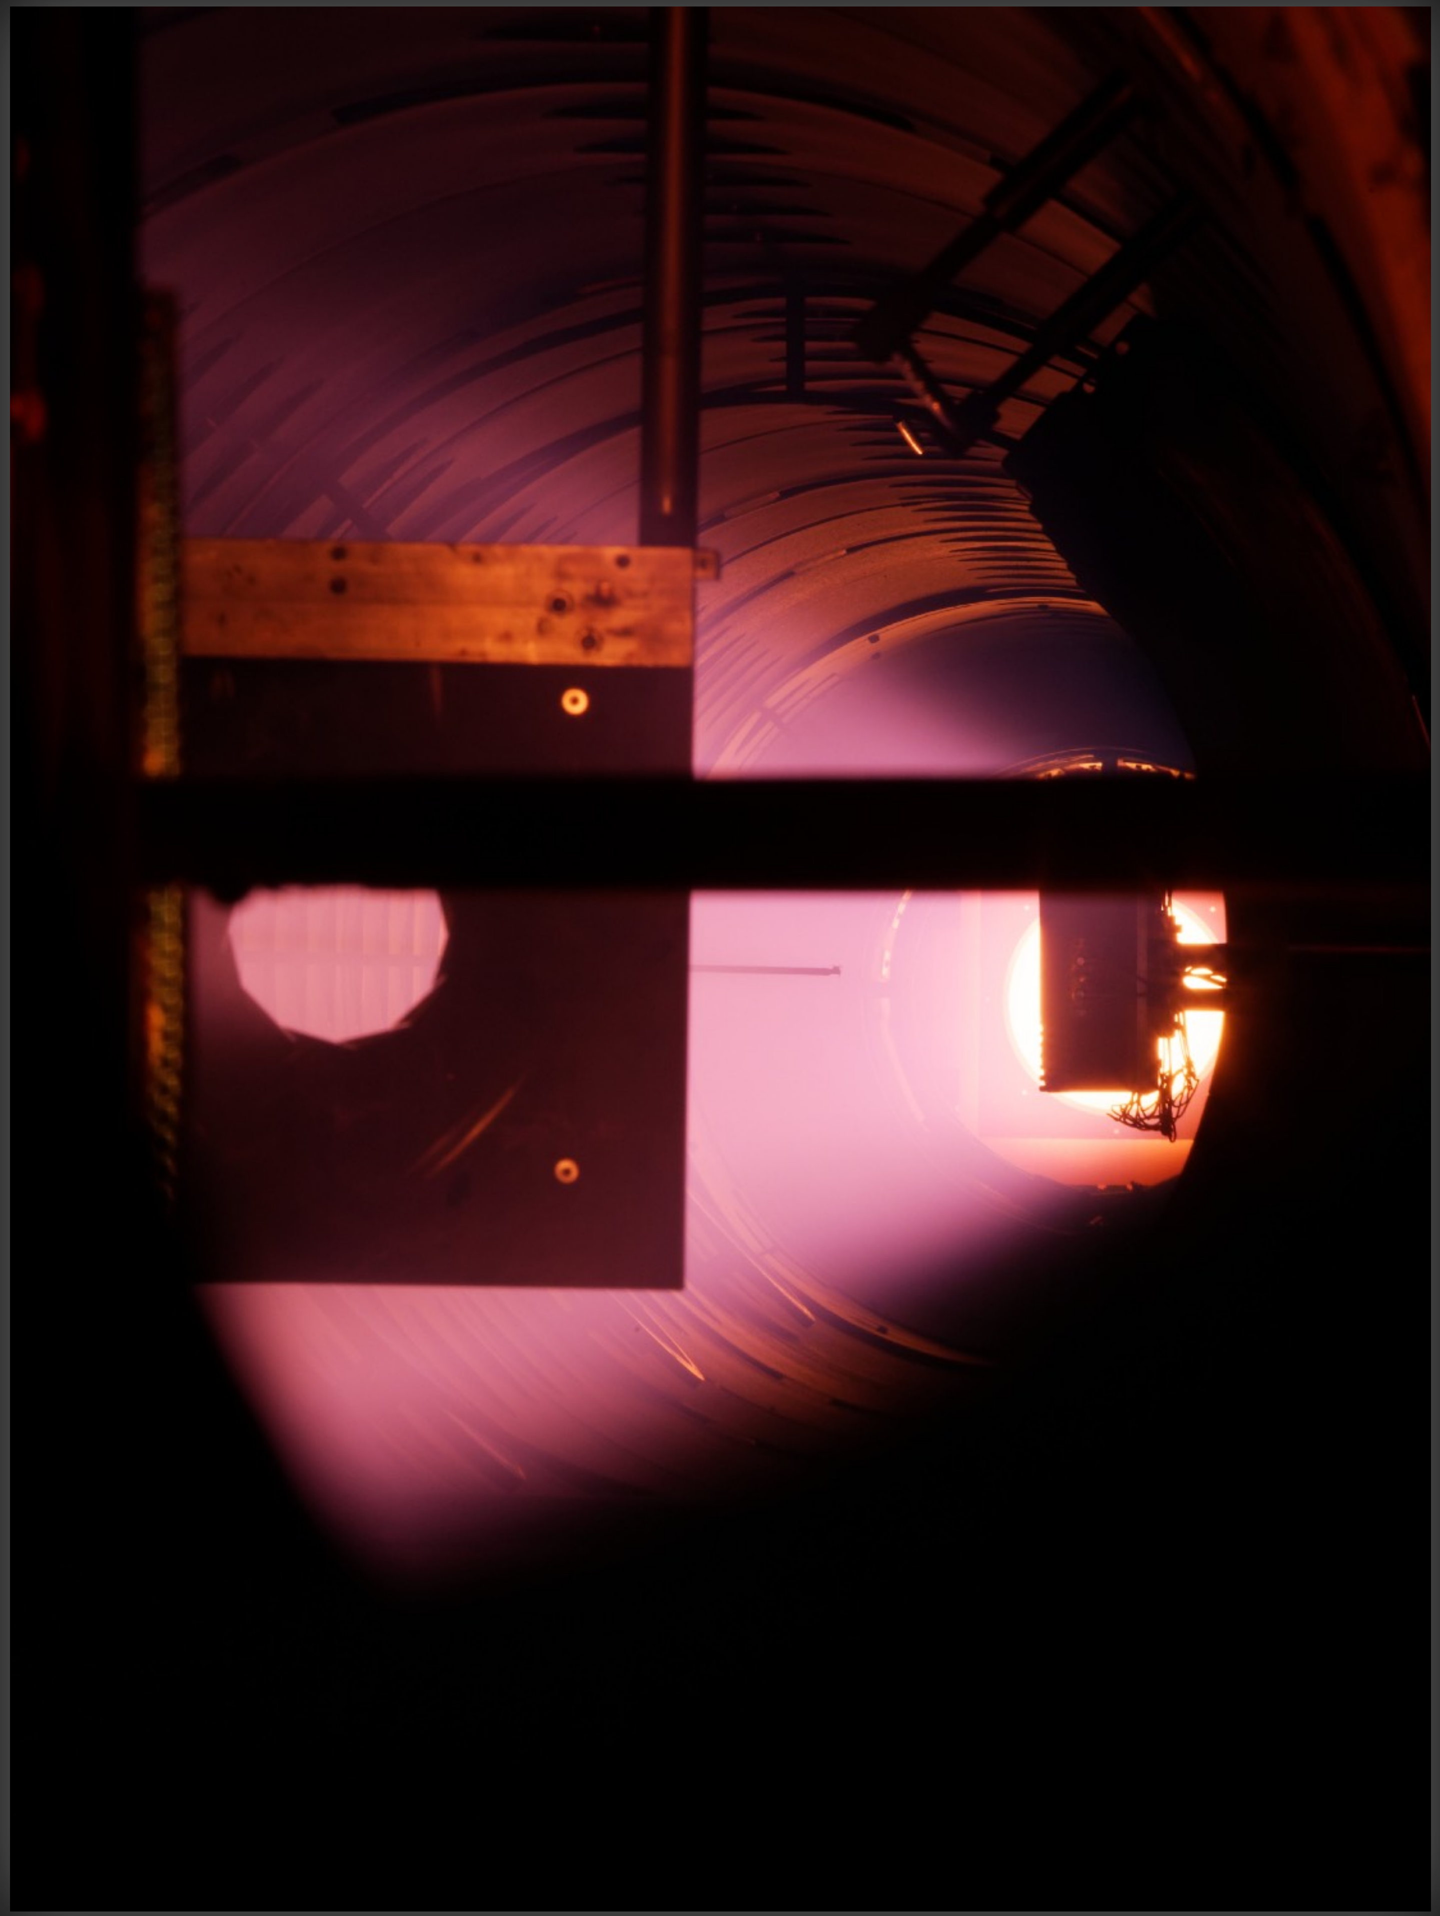

much pink

Langmuir and magnetic fluctuation (Bdot) probes are the workhorses of LAPD science

- Langmuir probes give you:
  - Density via ion saturation current: $I_{sat} \propto n_e\sqrt{T_e}$
  - Temperature via sweeps or triple probes
  - Potential via sweeps or floating potential
- Magnetic field fluctuations via Bdot
  - Useful for identifying and studying Alfvén waves
- High spatial resolution and reach: can measure (pretty much) anywhere in the LAPD

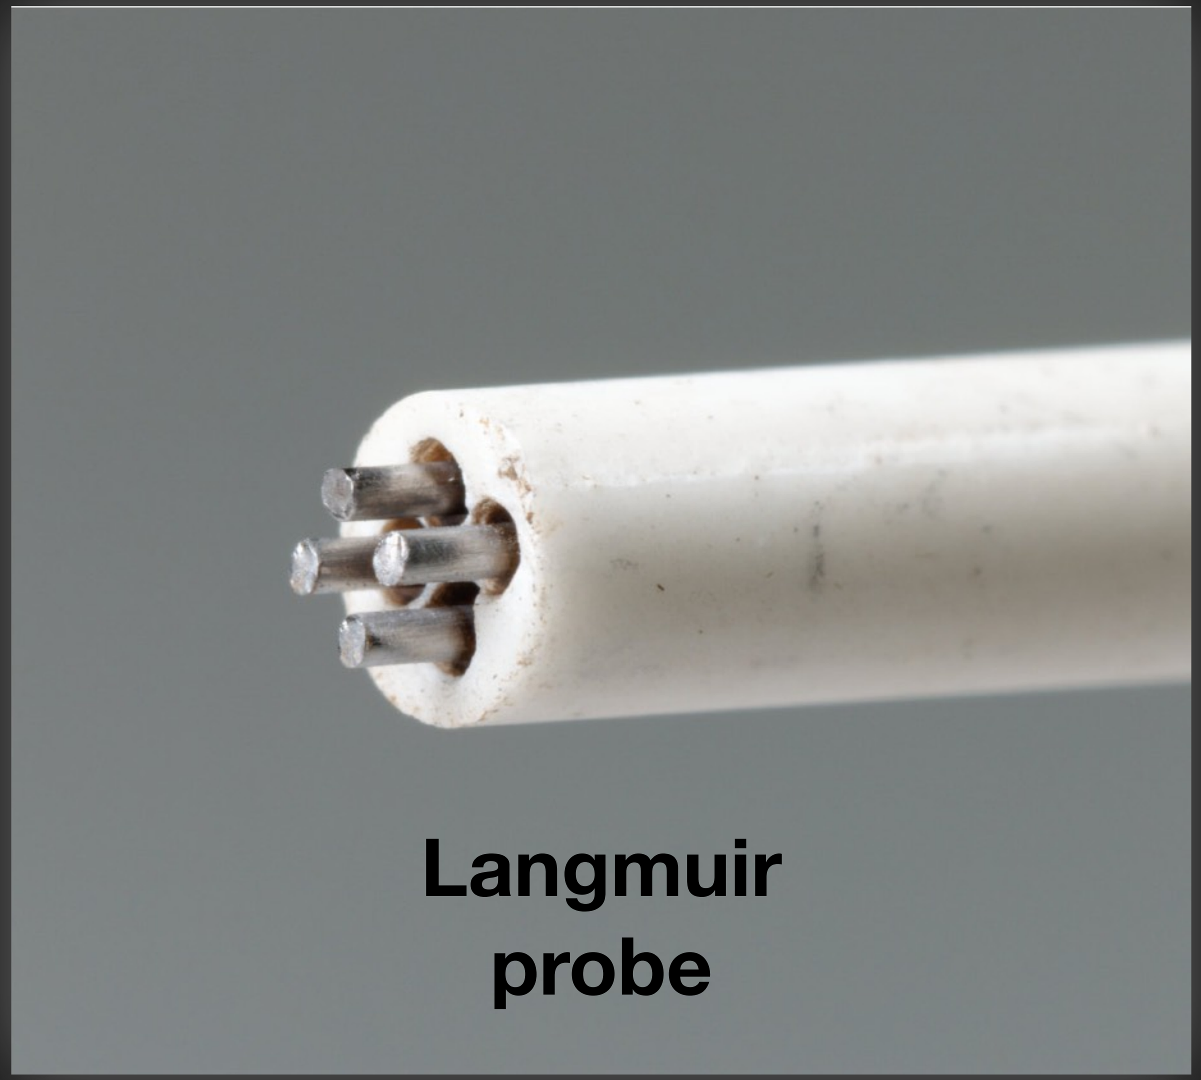

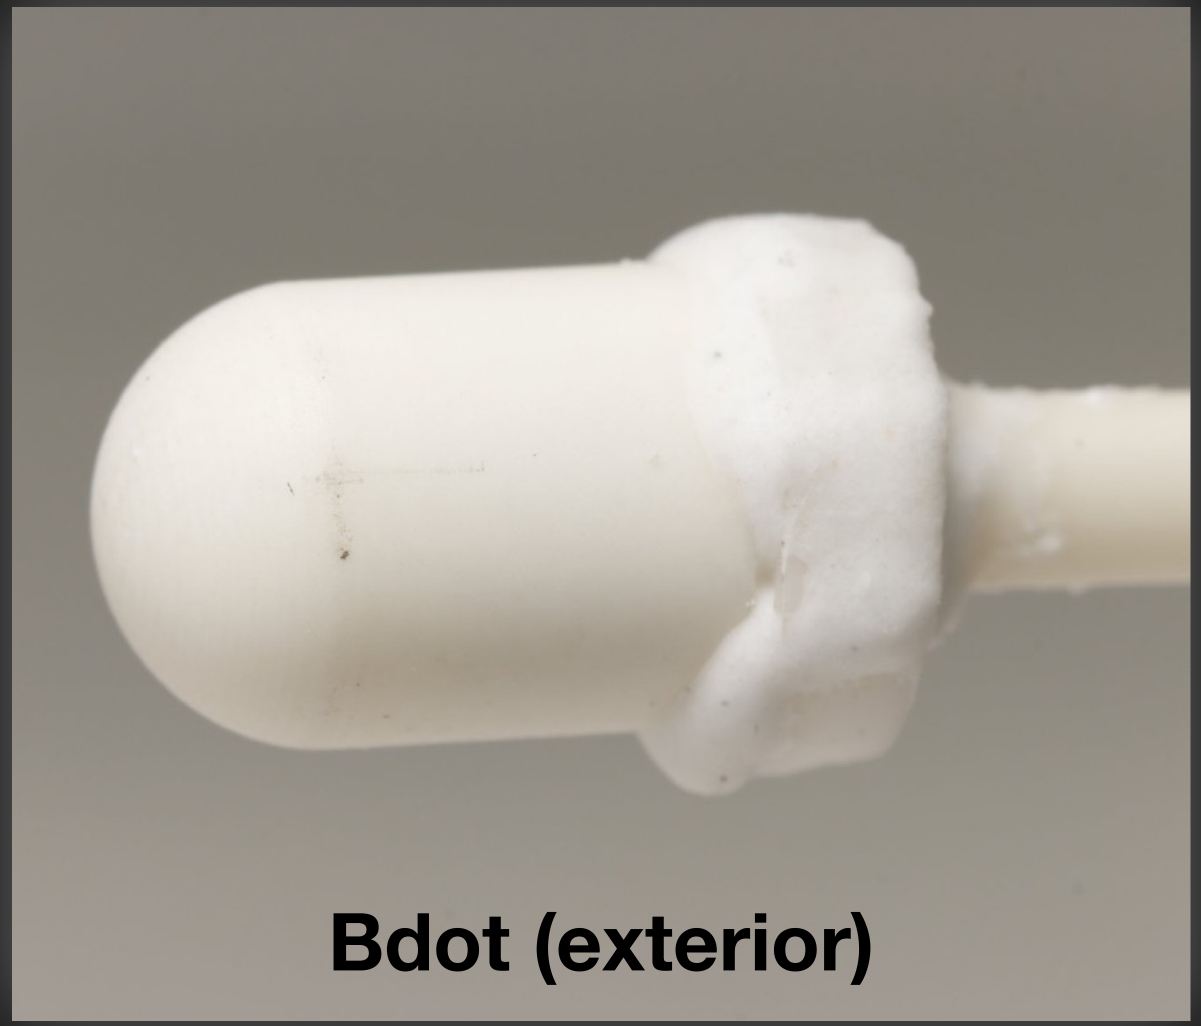

We made mirrors in the LAPD to study interchange modes and drift waves

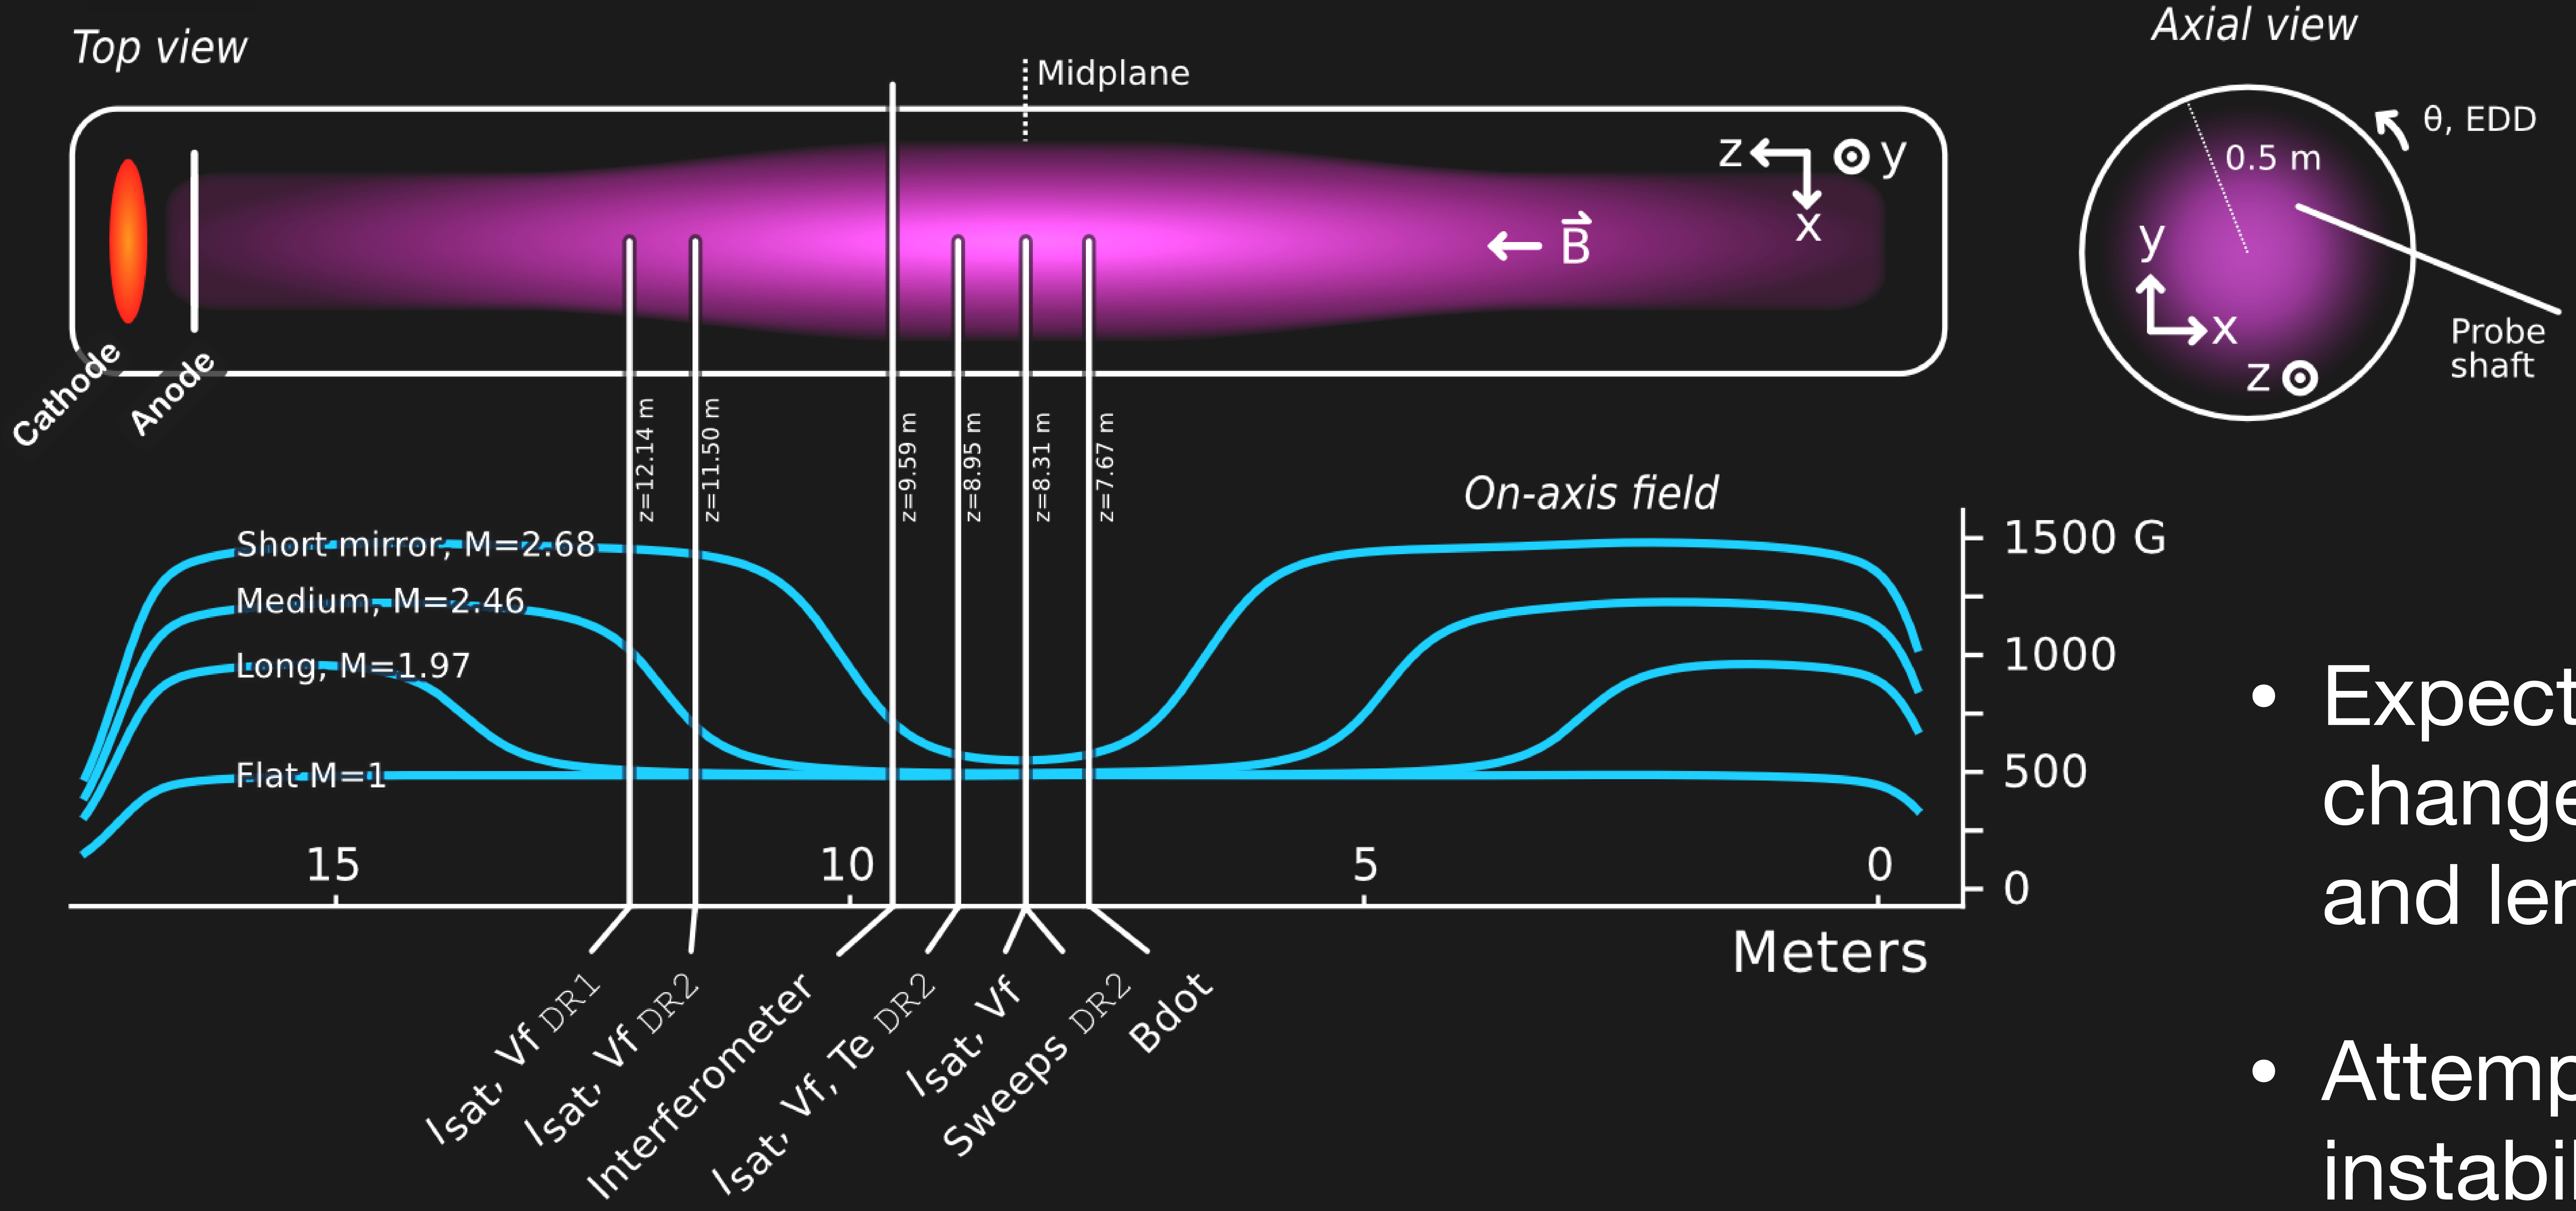

- Expect instabilities to change with mirror ratio and length
- Attempt to diagnose instabilities and modes present
- Expect low $k_{\|}$ modes — focus on central cell

Travis and Carter, JPP 2025

Changing mirror ratio: core profiles and gradient region are similar, some differences throughout

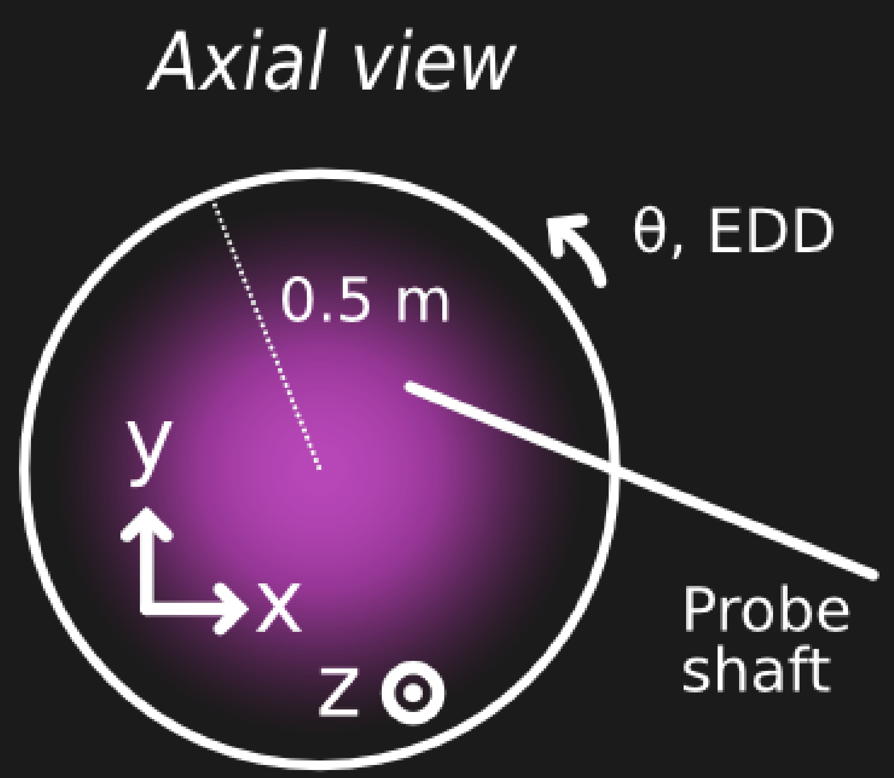

x_c: mapped cathode radius

x_{PF}: peak fluctuation point (from data)

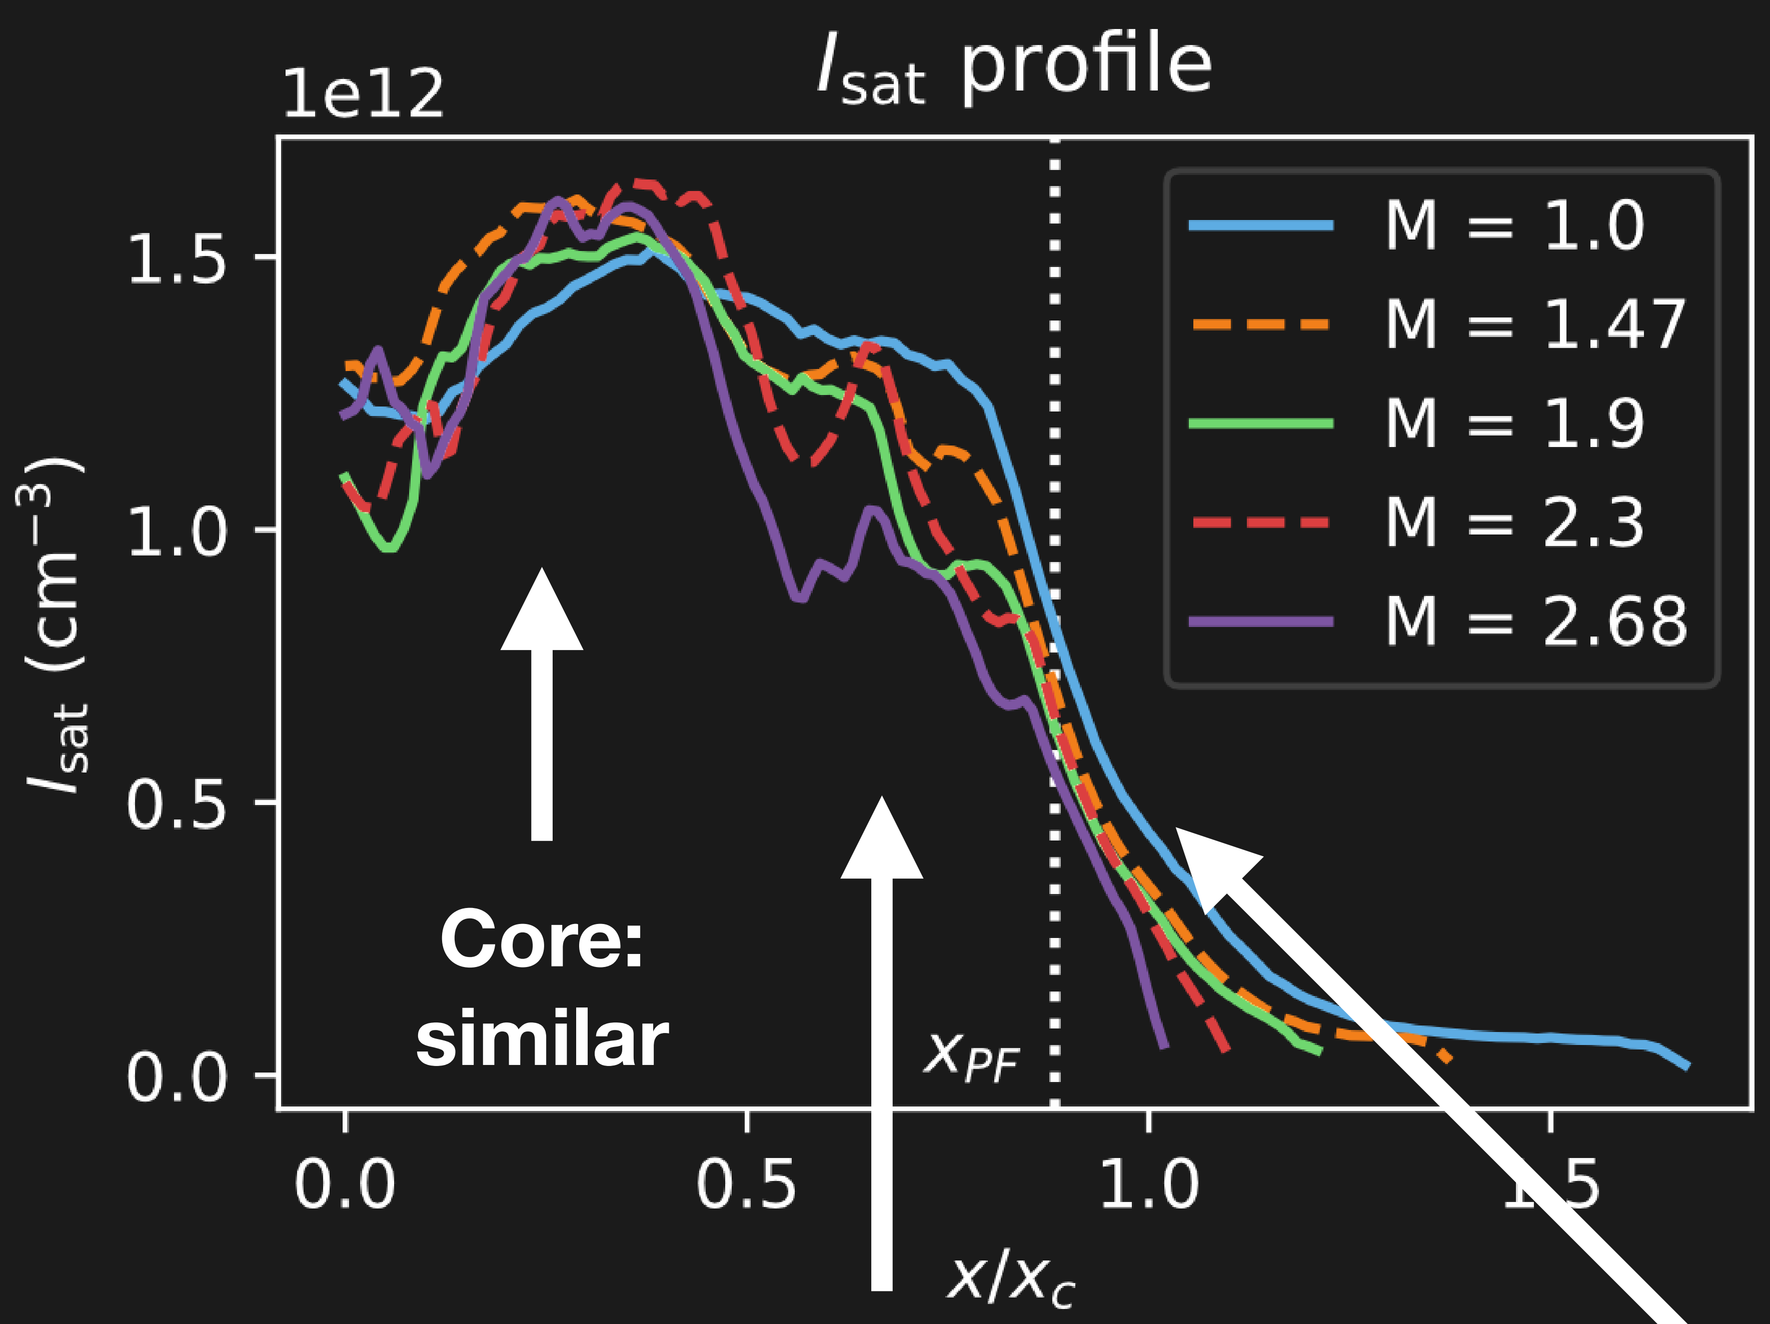

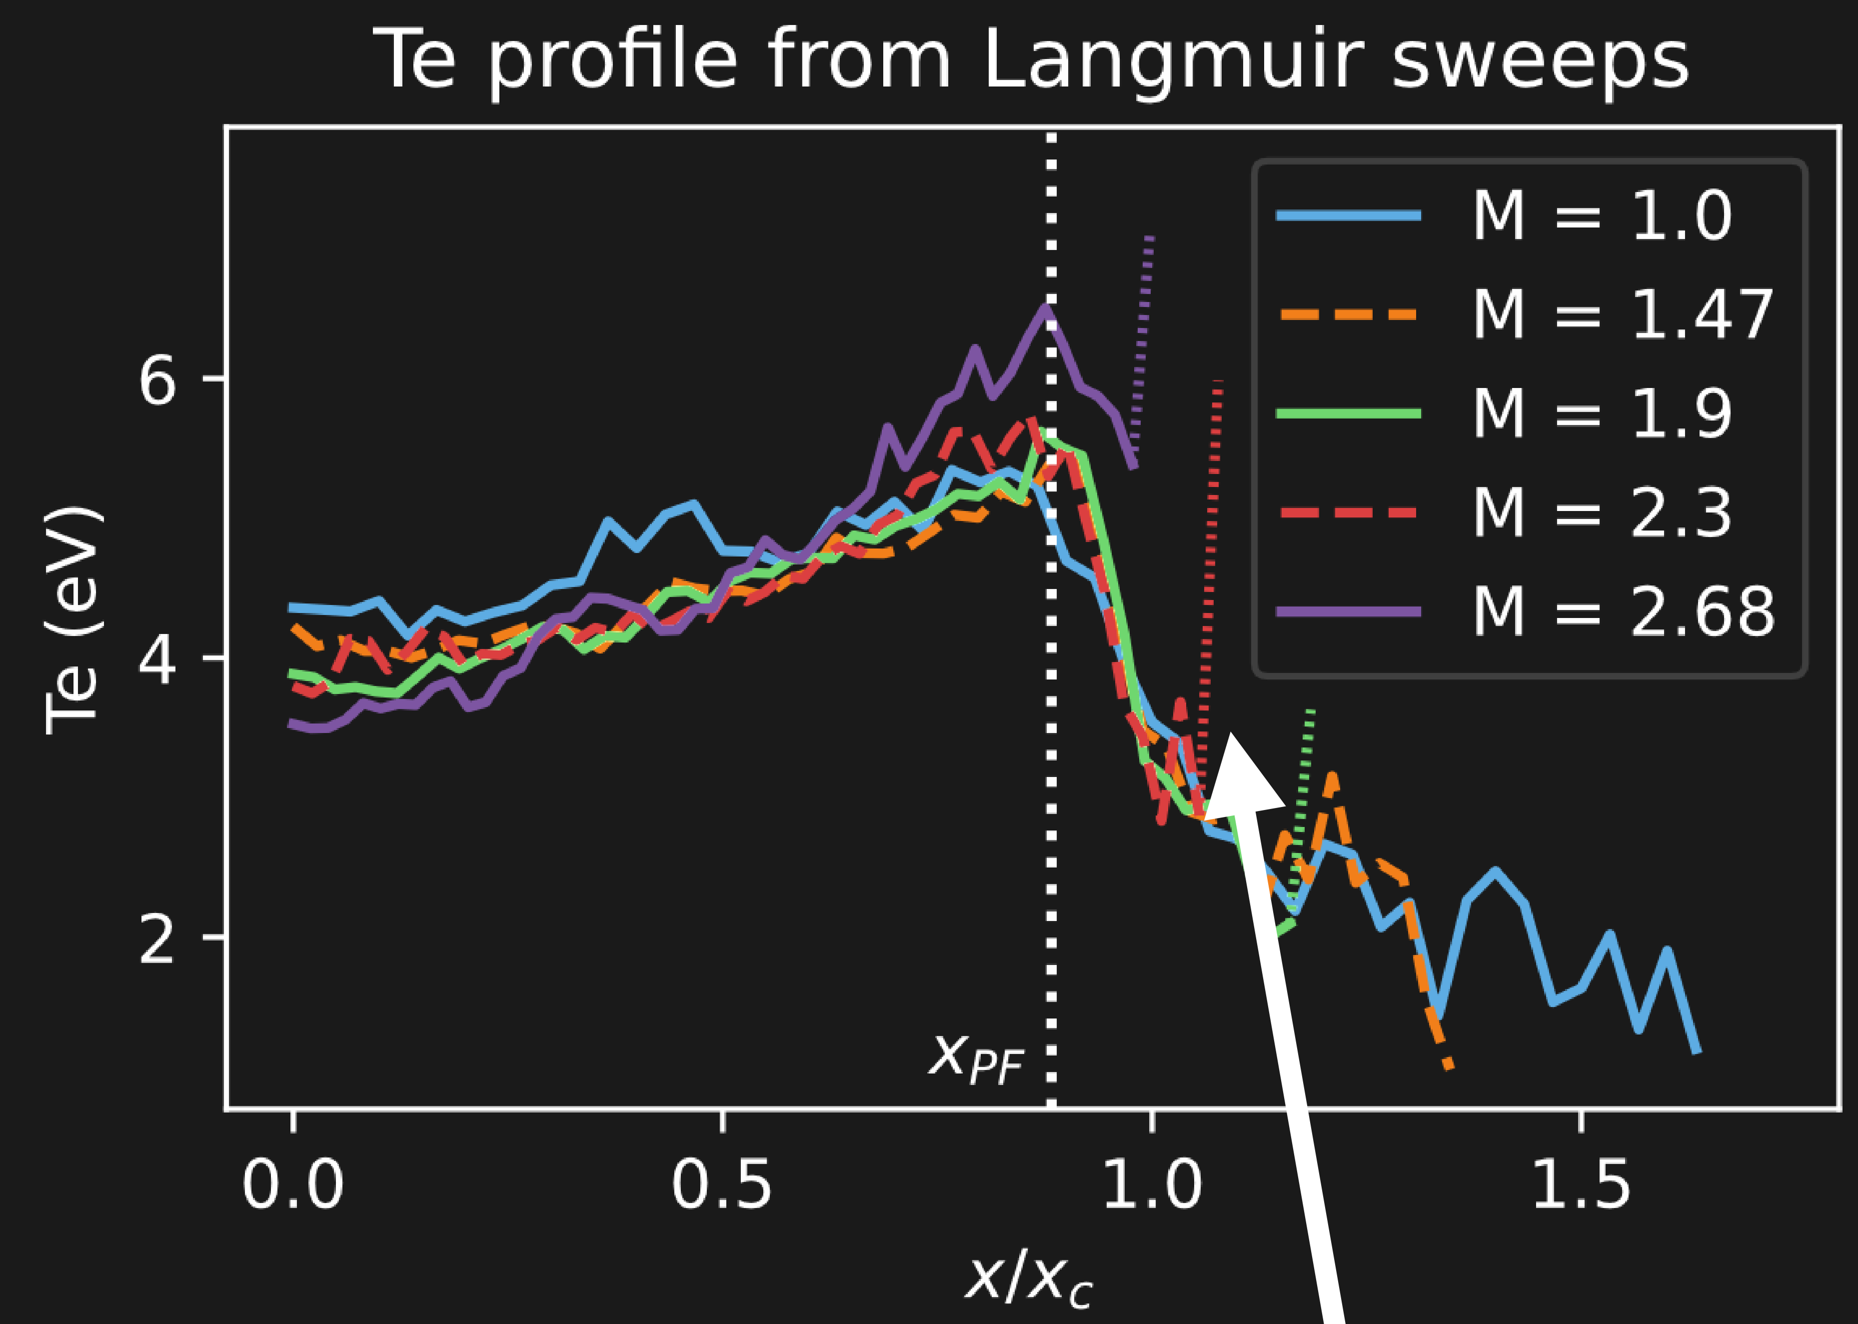

further out: different

Gradient region: similar

Ignoring dotted bits

We observe an unexpected decrease in particle flux and diffusivity

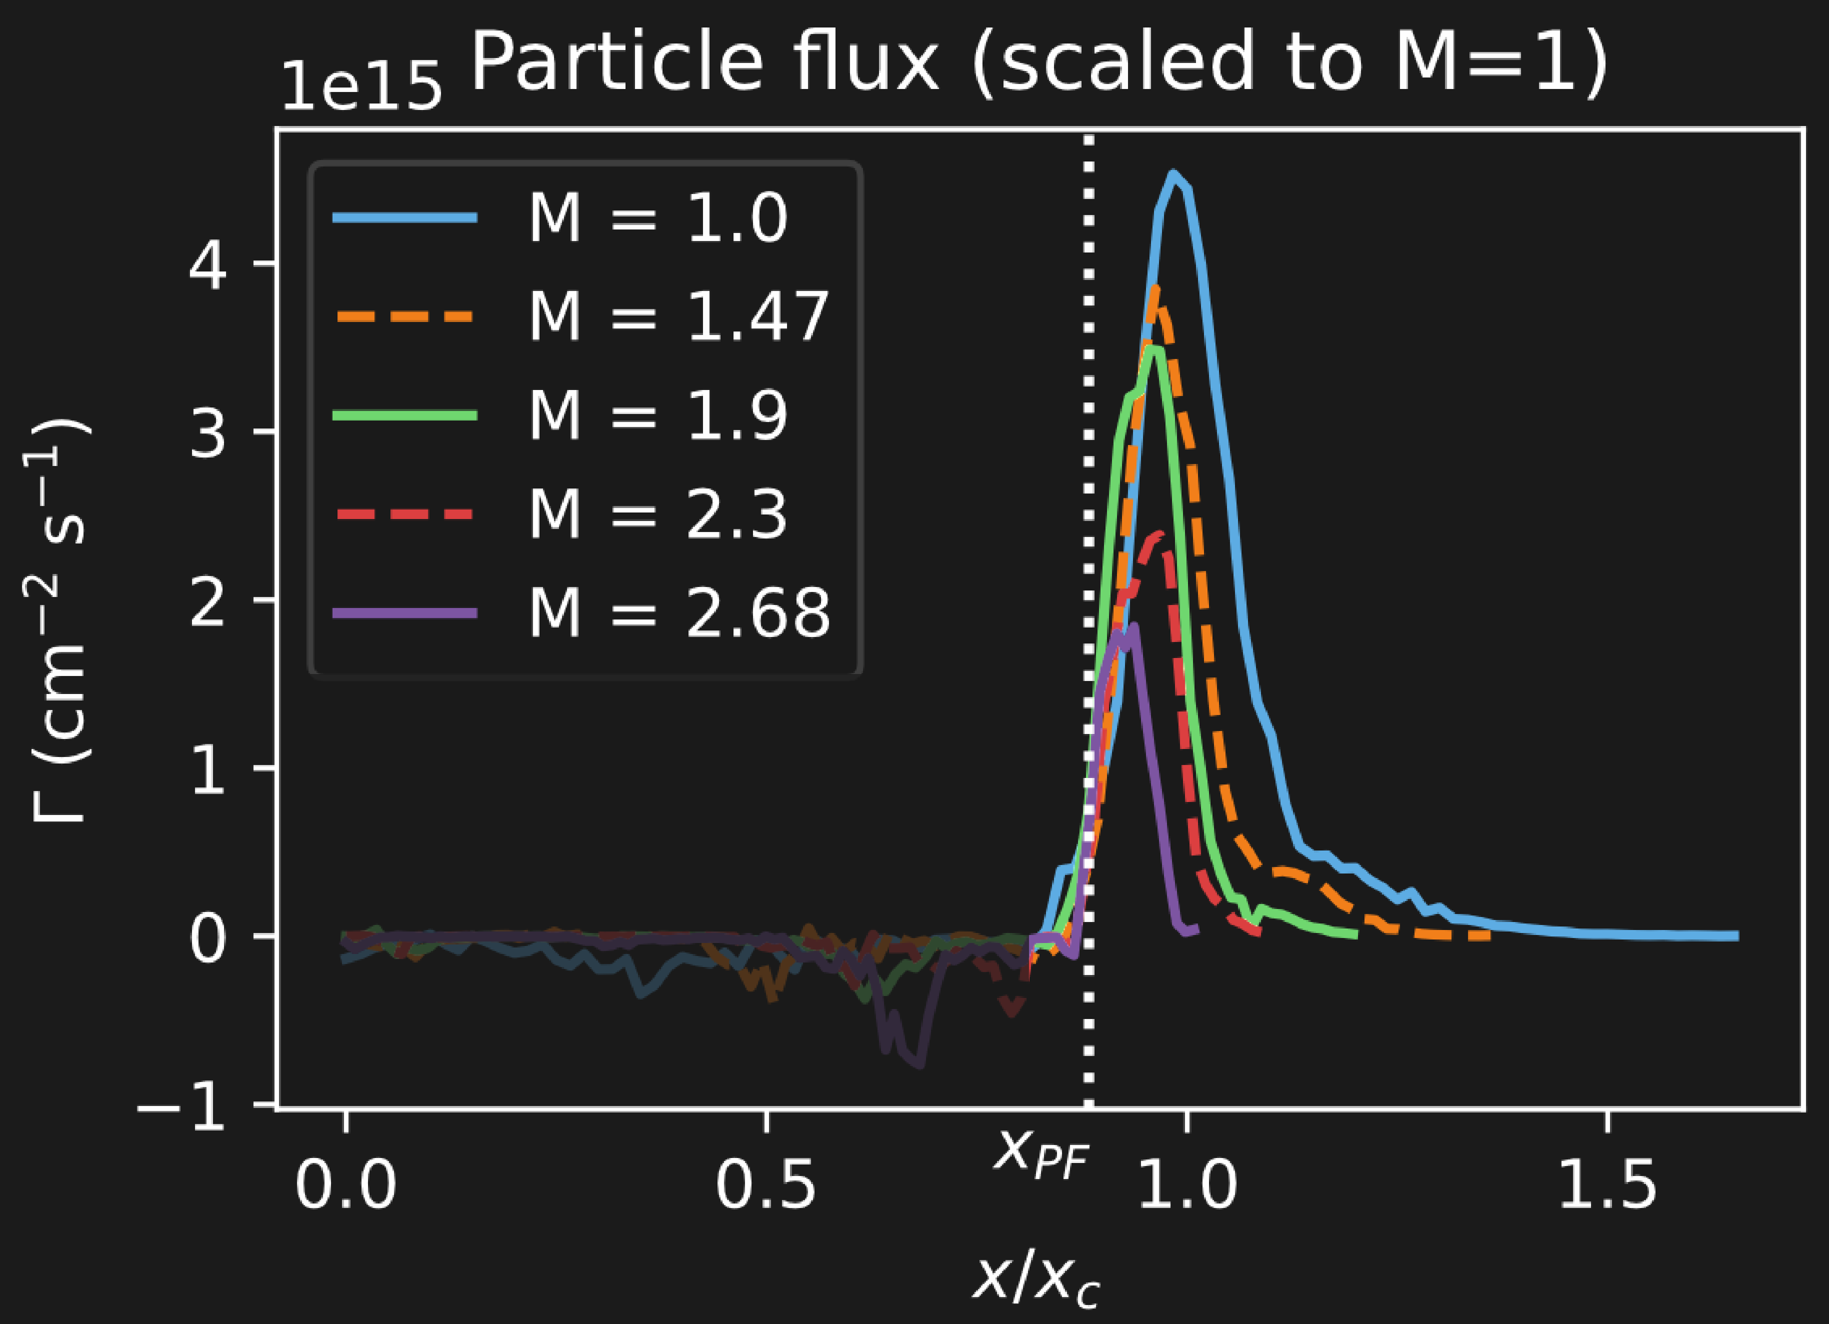

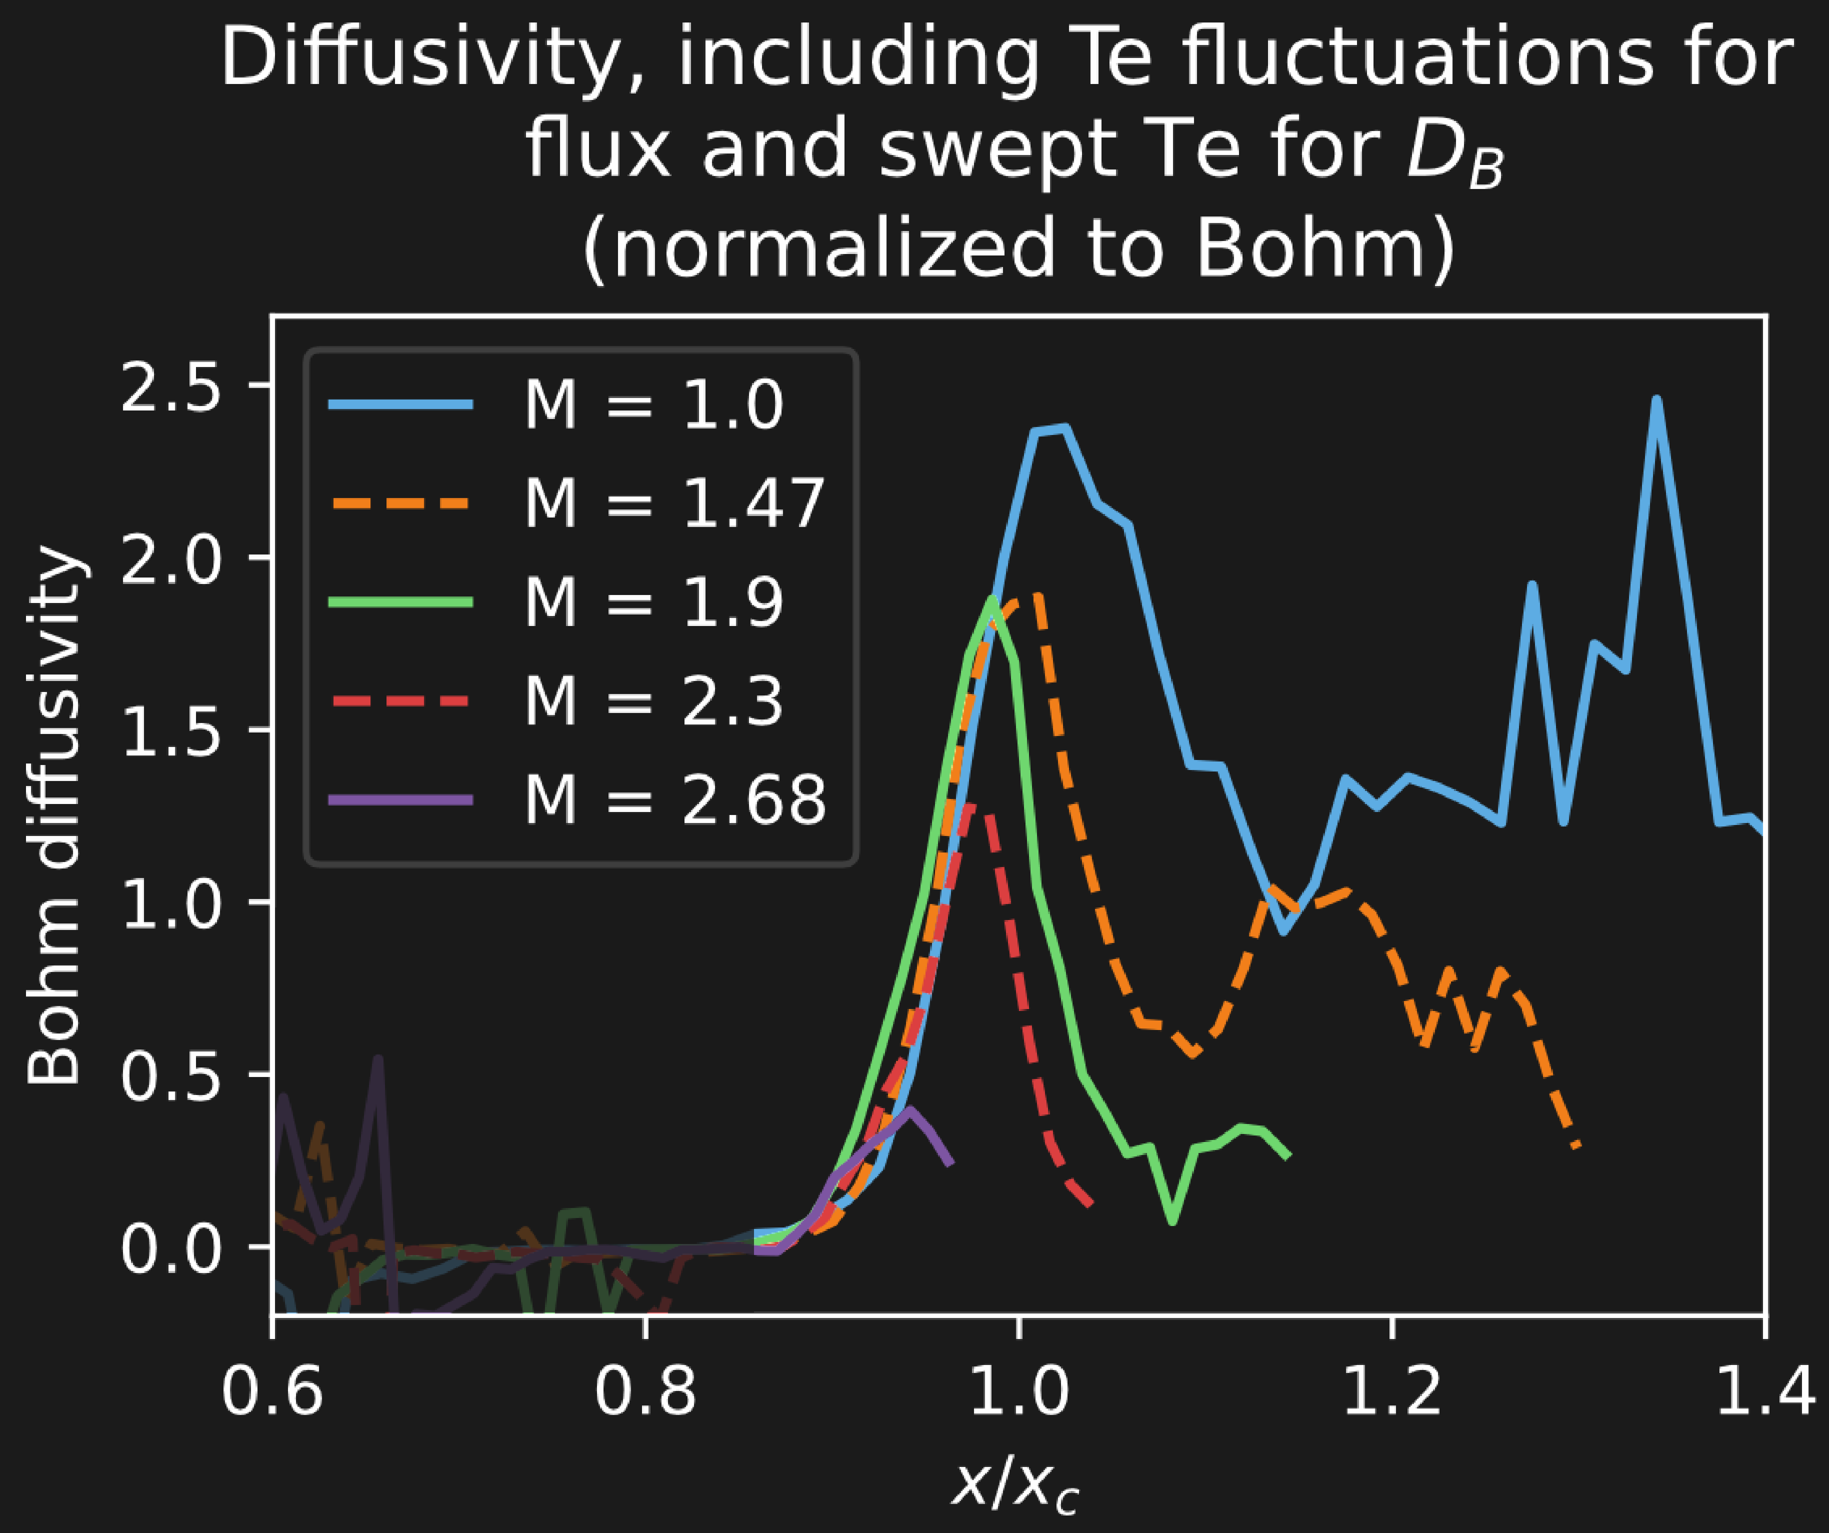

- Expect increased instability drive with increased curvature

Drift-Alfvén waves are clear on the fluctuation spectra; 3-6 kHz unclear

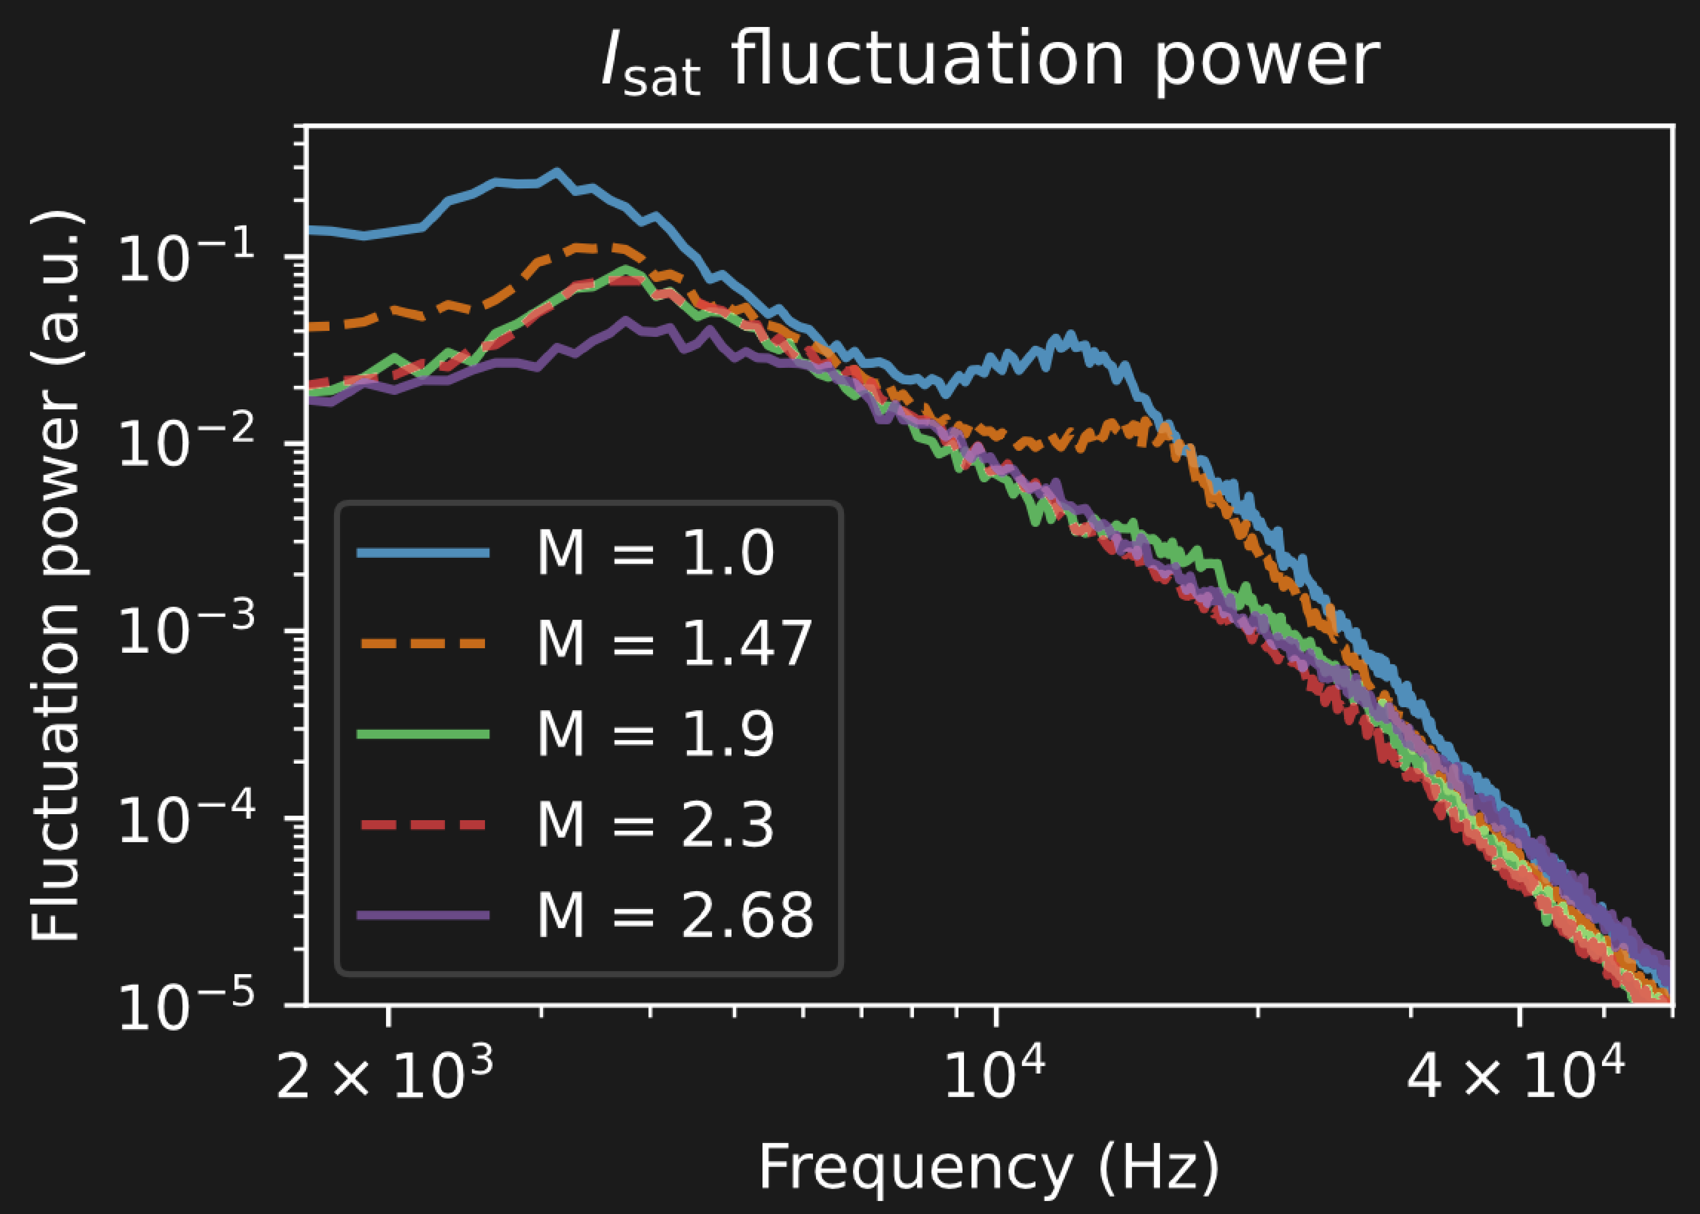

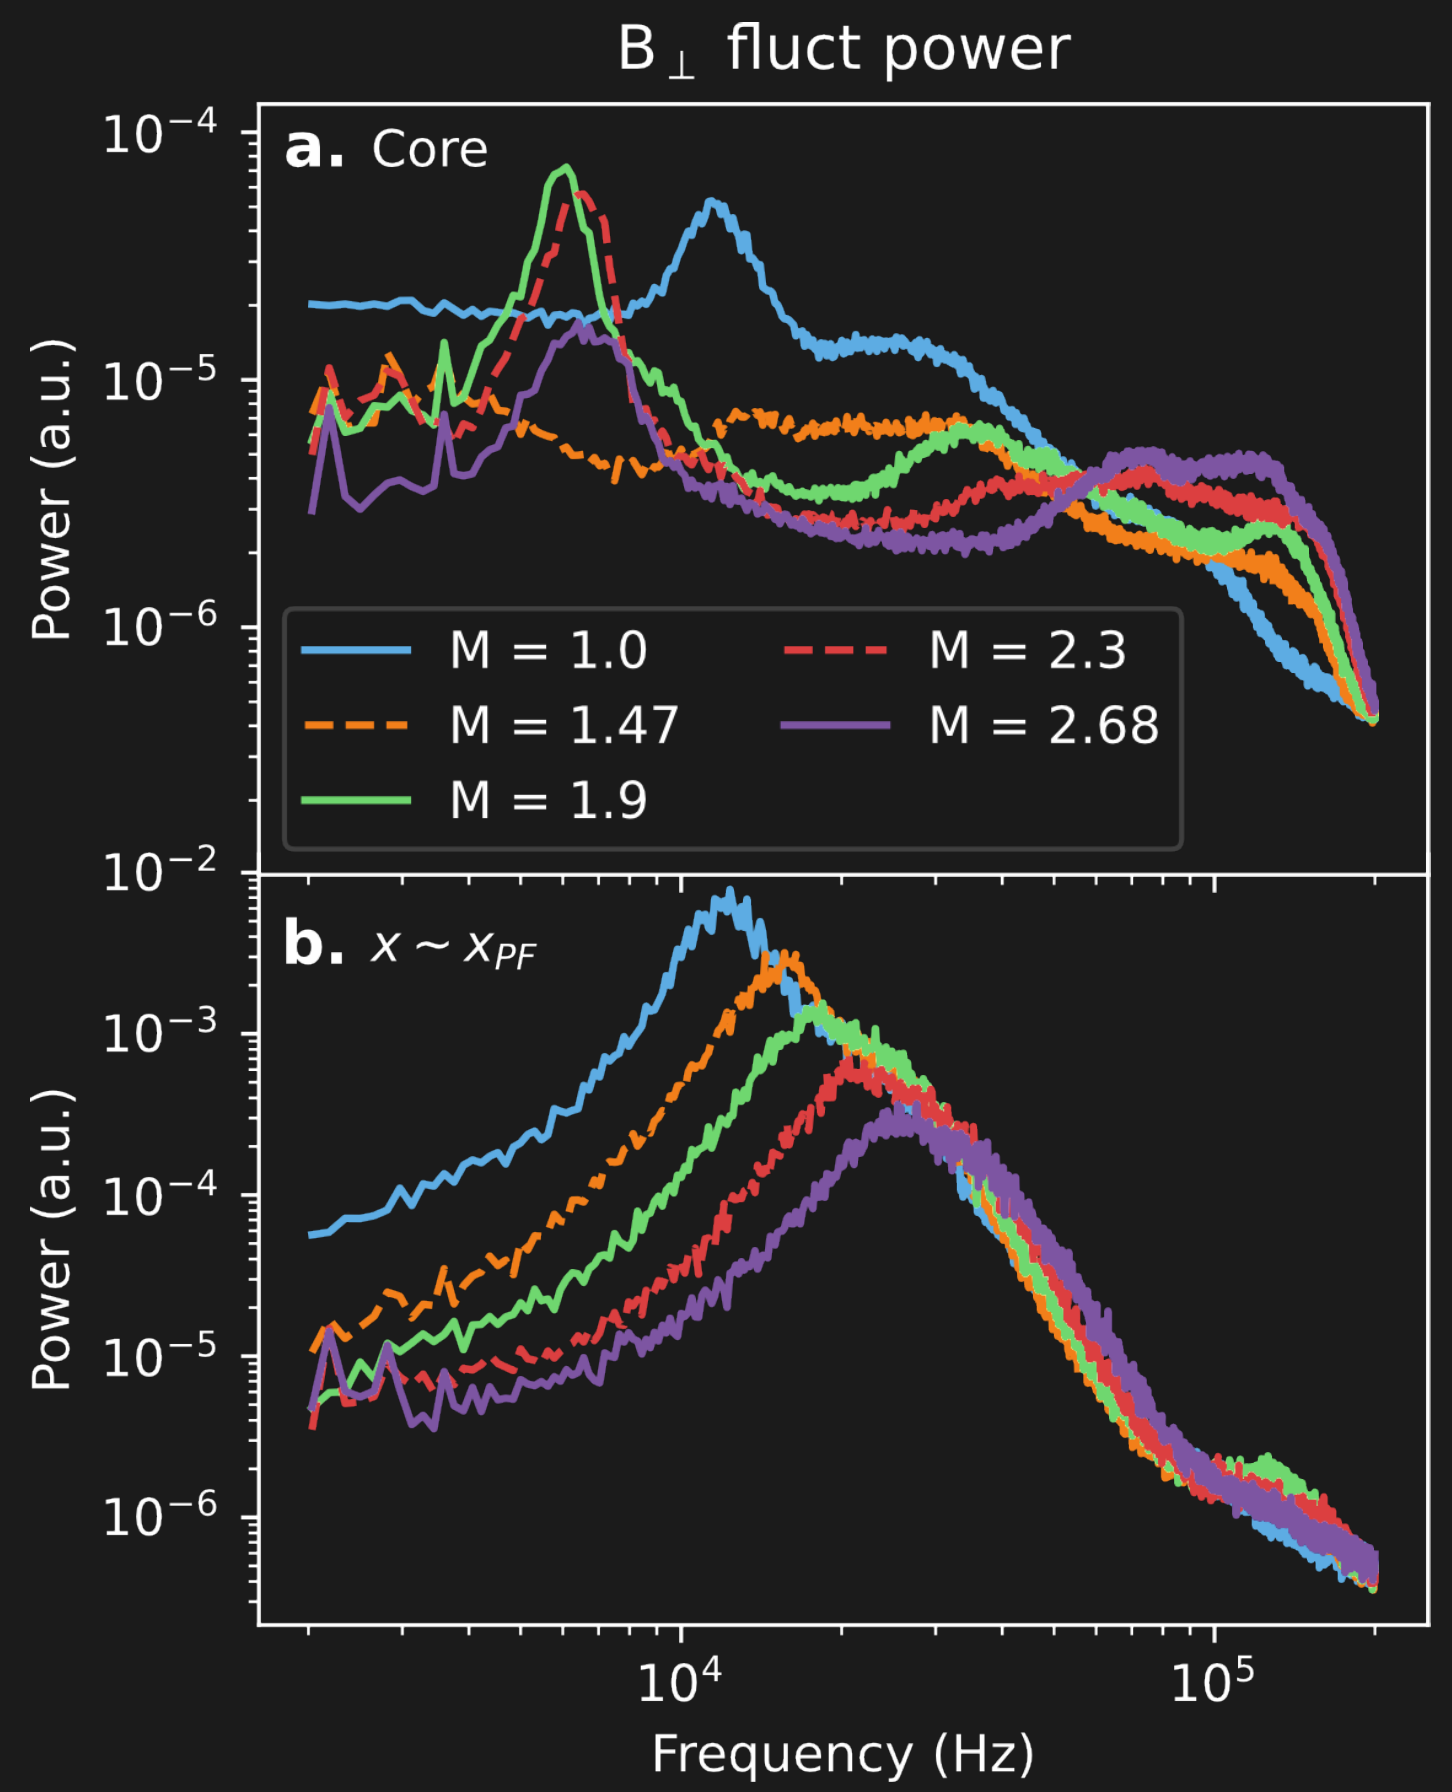

- 10+ kHz peaks: likely drift-Alfvén waves
- Peaks 3-6 kHz: open question
  - Rotational interchange, drift waves, nonlinear instability, conducting wall mode?

No interchange instability is seen in these mirrors; mysteries remain

- Performed experiments in a range of mirror ratios and lengths
- No evidence for the interchange instability (many stabilization mechanisms)
- See an unexpected decrease in particle flux and diffusivity
- To study interchange on the LAPD, likely need to explore higher-β plasmas

Broader exploration of parameter space would be beneficial

How?

Machine learning

The Large Plasma Device is an ideal machine for collecting data for ML

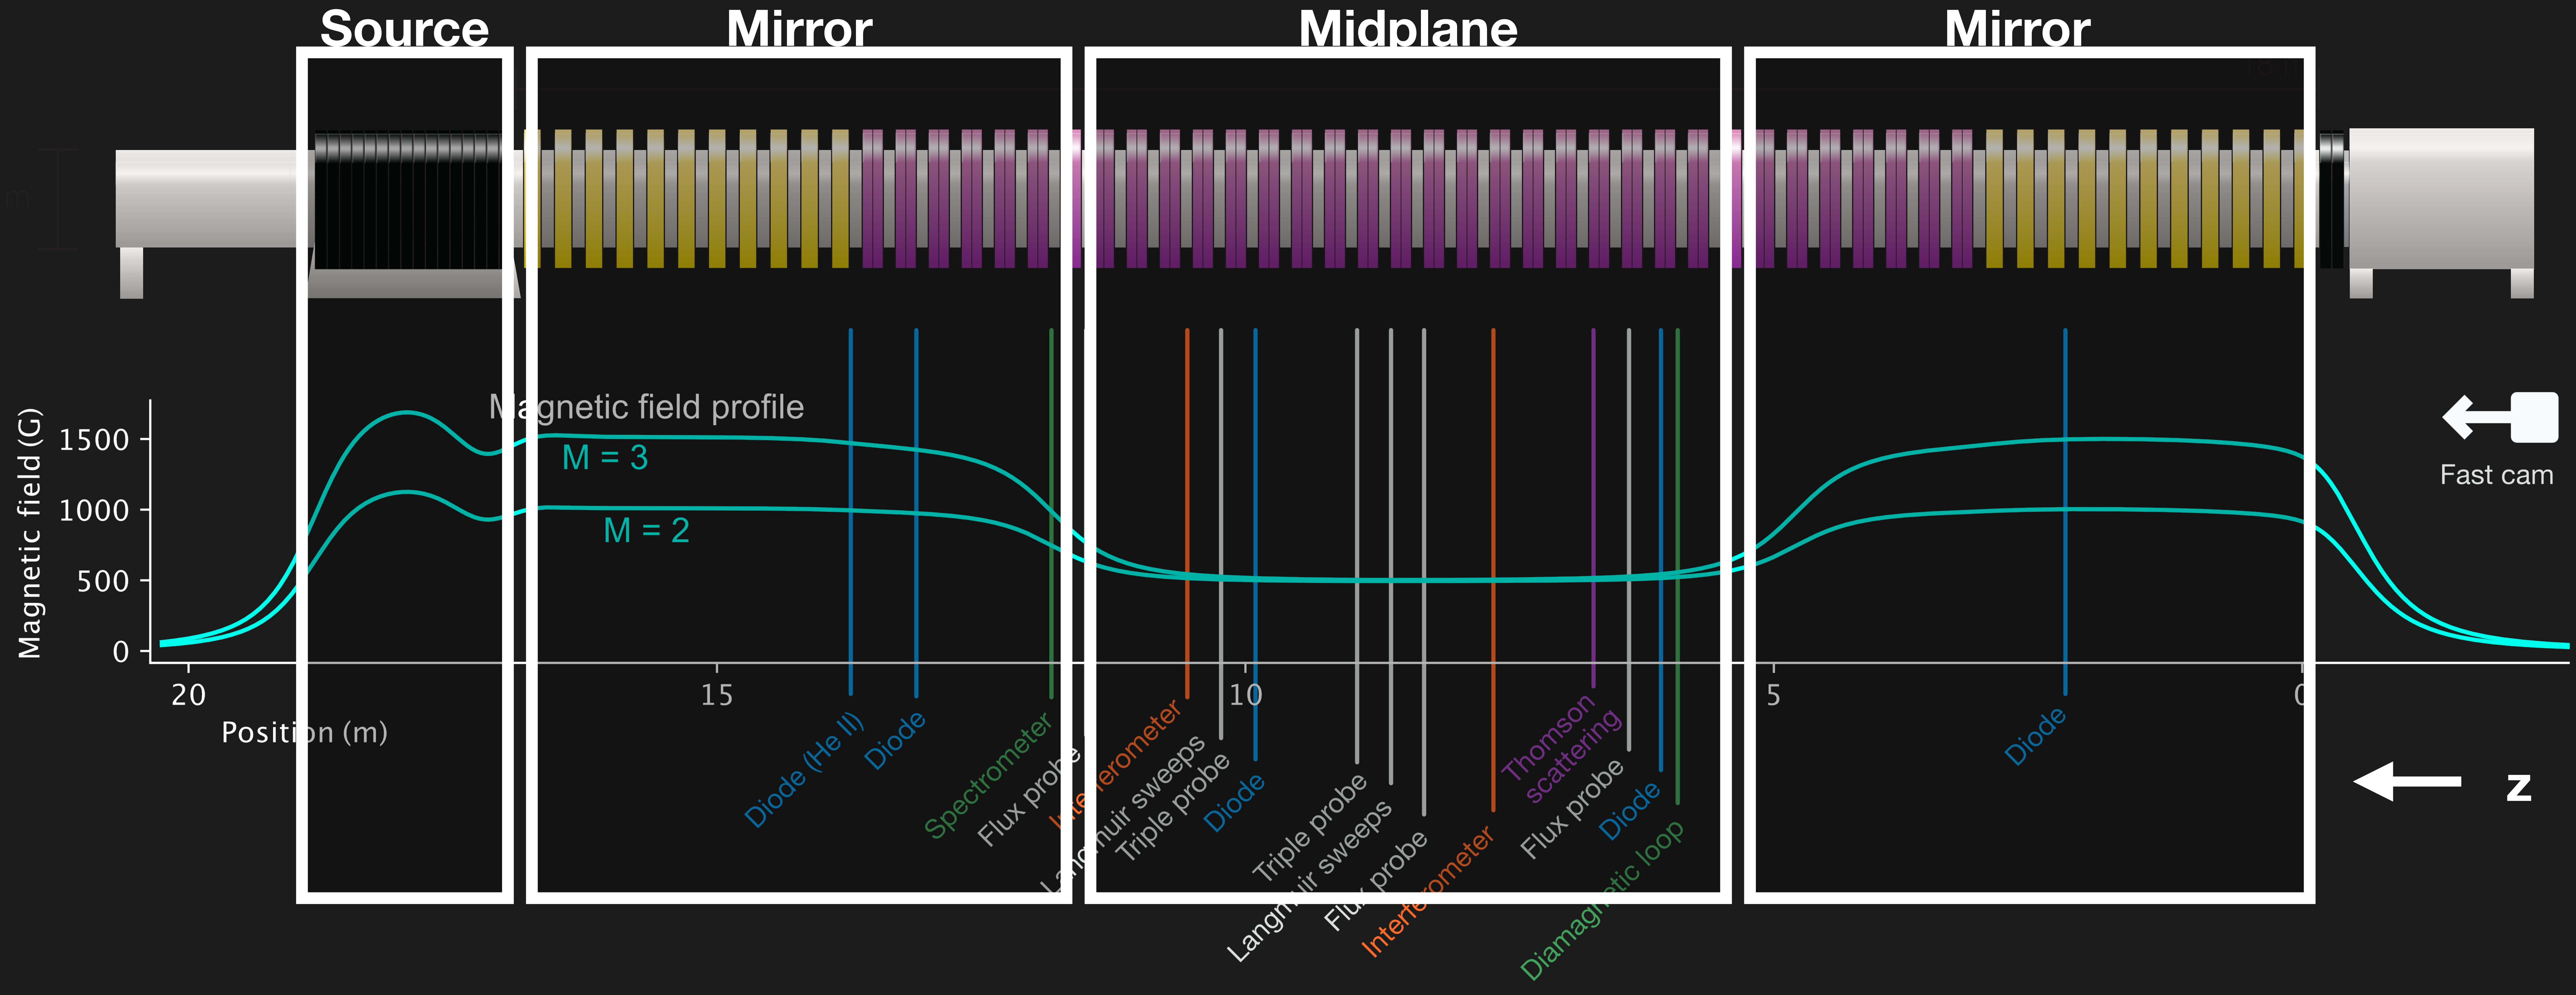

- High rep rate: 0.25-1 Hz rep rate
- Flexible machine configuration
- Great diagnostic access
- High-resolution probe measurements

For profile optimization, LAPD configurations were randomly sampled

- Goal: determine machine settings for optimal axial profiles at high density
- Machine controls and actuators have a nonlinear effect on plasmas
- Randomization is necessary but risky
- Collected 44 randomized dataruns (67 runs total)

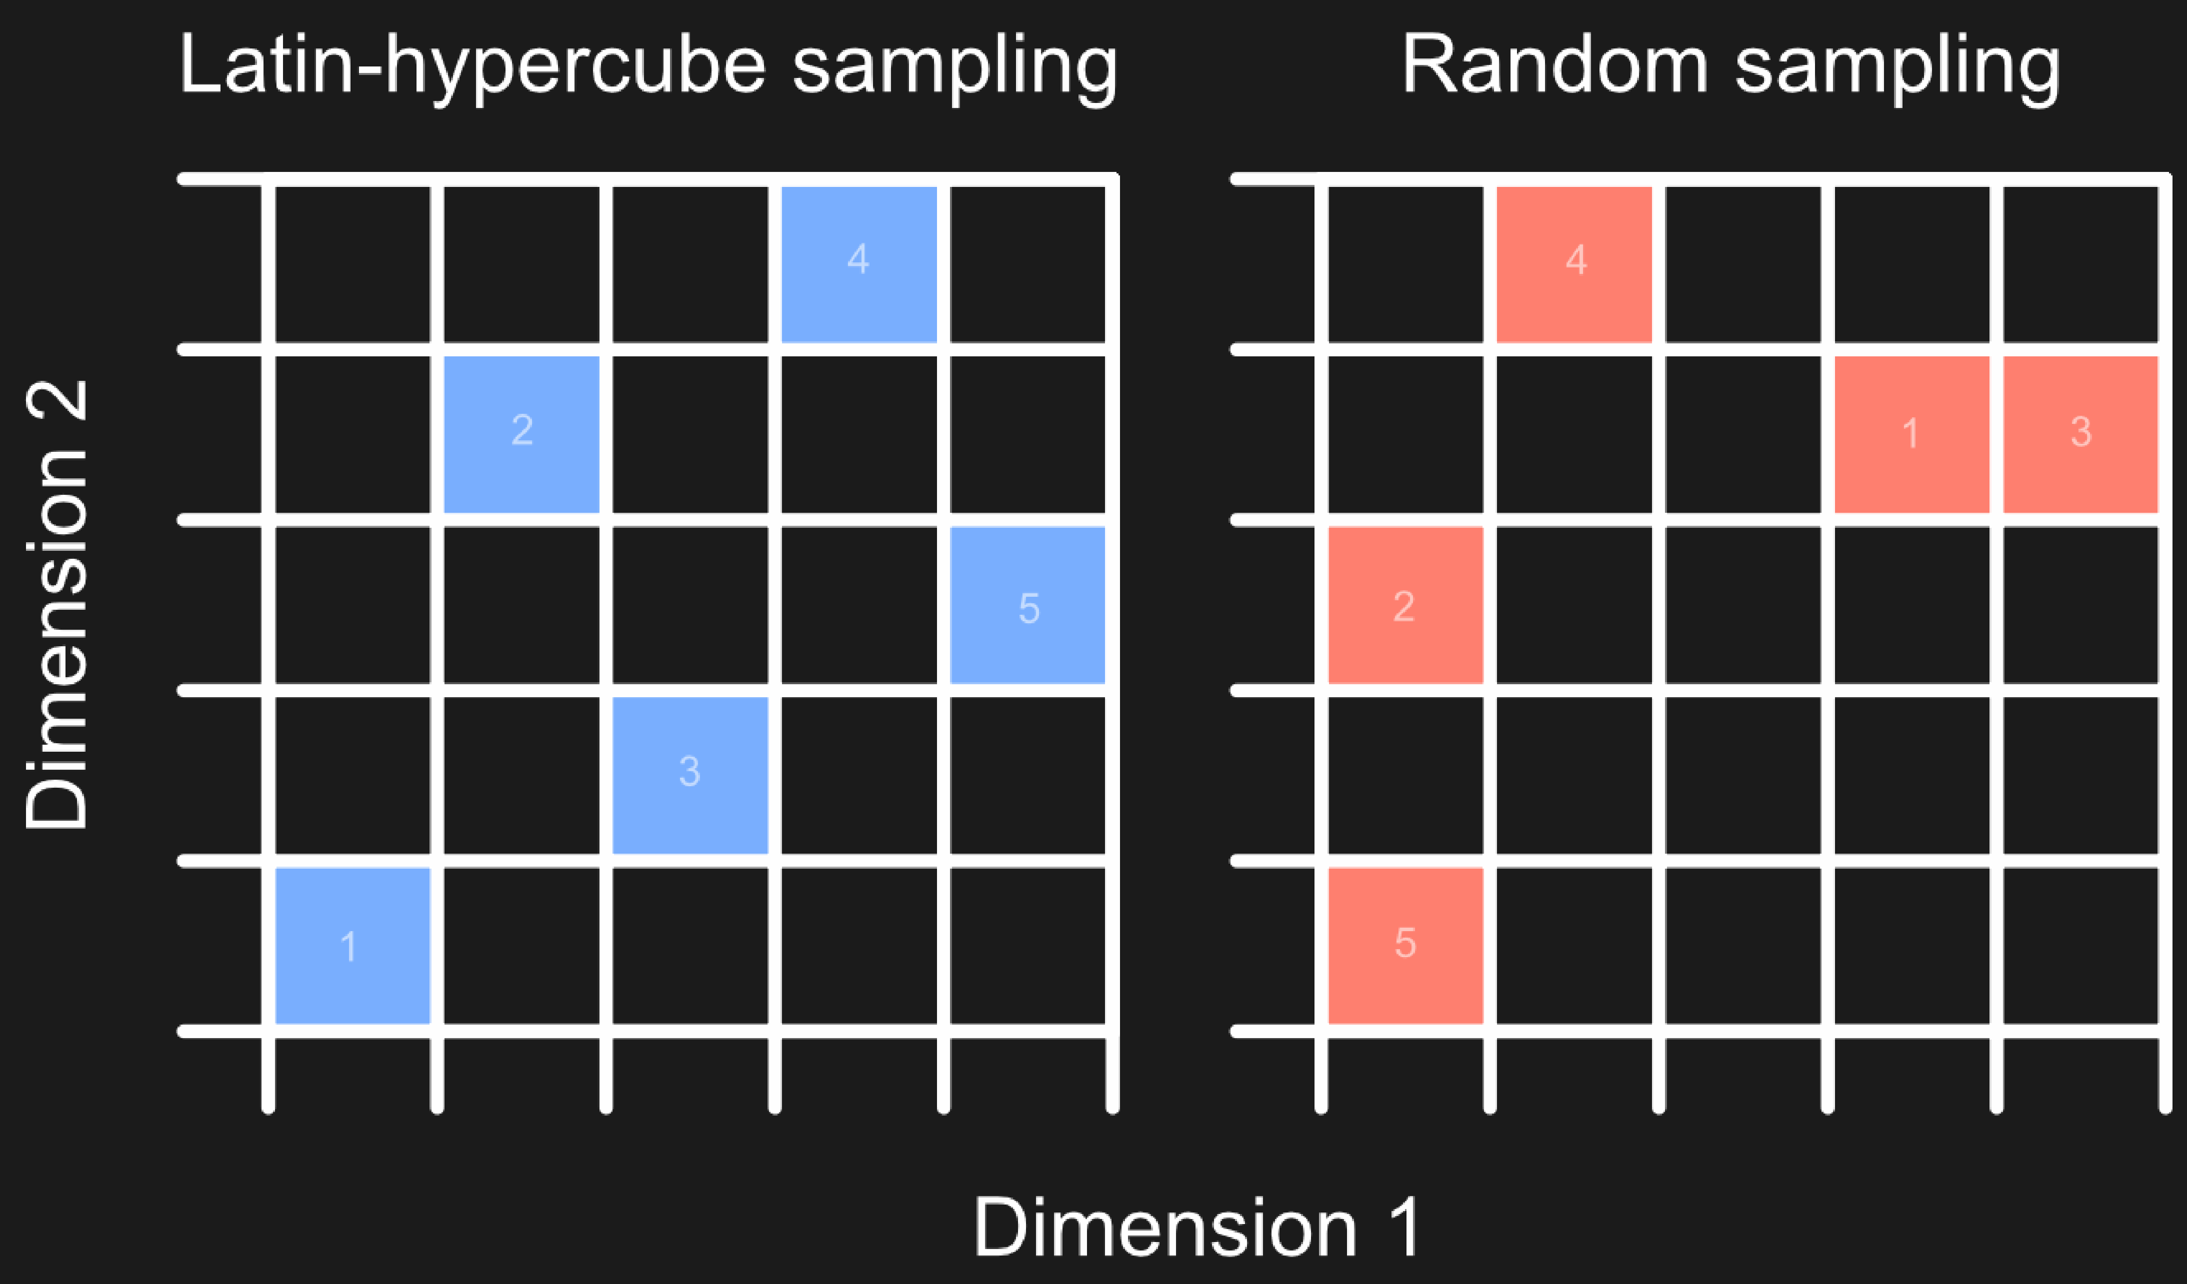

136k possible machine configurations

B_{source},
B_{mirror},
$B_{midplane}$,
Gas puff settings,
Discharge voltage

Travis, Bortnik, and Carter, arXiv:2503.09868

Neural networks (NNs) are used to learn time-averaged I_{sat}

- NNs are "universal function approximators"
  - they can fit any function given sufficient capacity (200k parameters for mine)
- NNs will learn the trends necessary to reduce error

Machine learning is just fancy curve fitting

Input: machine settings → **NN** → **Output**: time-averaged I_{sat} (10-20 ms)

Test

- Held out 8 dataruns
- Used to evaluate model on unseen machine configurations

Train

- 80% of remaining 59 runs
- Model trains on this set

Validation

- 20% of remaining 59 runs
- To prevent overfitting

Uncertainty can be quantified using the NLL loss and ensembles

$$\mathscr{L}_{\beta\text{–}\mathbf{NLL}} = \frac{1}{2}\left(\log \sigma_i^2(\mathbf{x}_n) + \frac{\left(\mu_i(\mathbf{x}_n) - y_n\right)^2}{\sigma_i^2(\mathbf{x}_n)}\right) \textbf{StopGrad}\left(\sigma_i^{2\beta}\right)$$

Model i
Example n

Penalty for large uncertainty

MSE scaled by uncertainty

Example-specific learning rate

- Break uncertainty into intrinsic randomness (aleatoric) and model-based (epistemic) uncertainty

Aleatoric uncertainty

$$\langle \sigma_i^2(\mathbf{x}) \rangle$$

Epistemic uncertainty

$$\langle \mu_i^2(\mathbf{x}) \rangle - \mu_*^2(\mathbf{x}) = \mathbf{Var}[\mu_i(\mathbf{x})]$$

$$\mu_*(\mathbf{x}) = \langle \mu_i(\mathbf{x}) \rangle$$

- The uncertainty quantification done here is uniquely thorough

Cross-validation: choice of test set can have a big impact on estimated error

- Test set 0 was hand picked for diversity
- Changing test set can dramatically change the measured error
- Test set performance improves when using ensembles
- Will use the median RMSE as a guide for estimating error

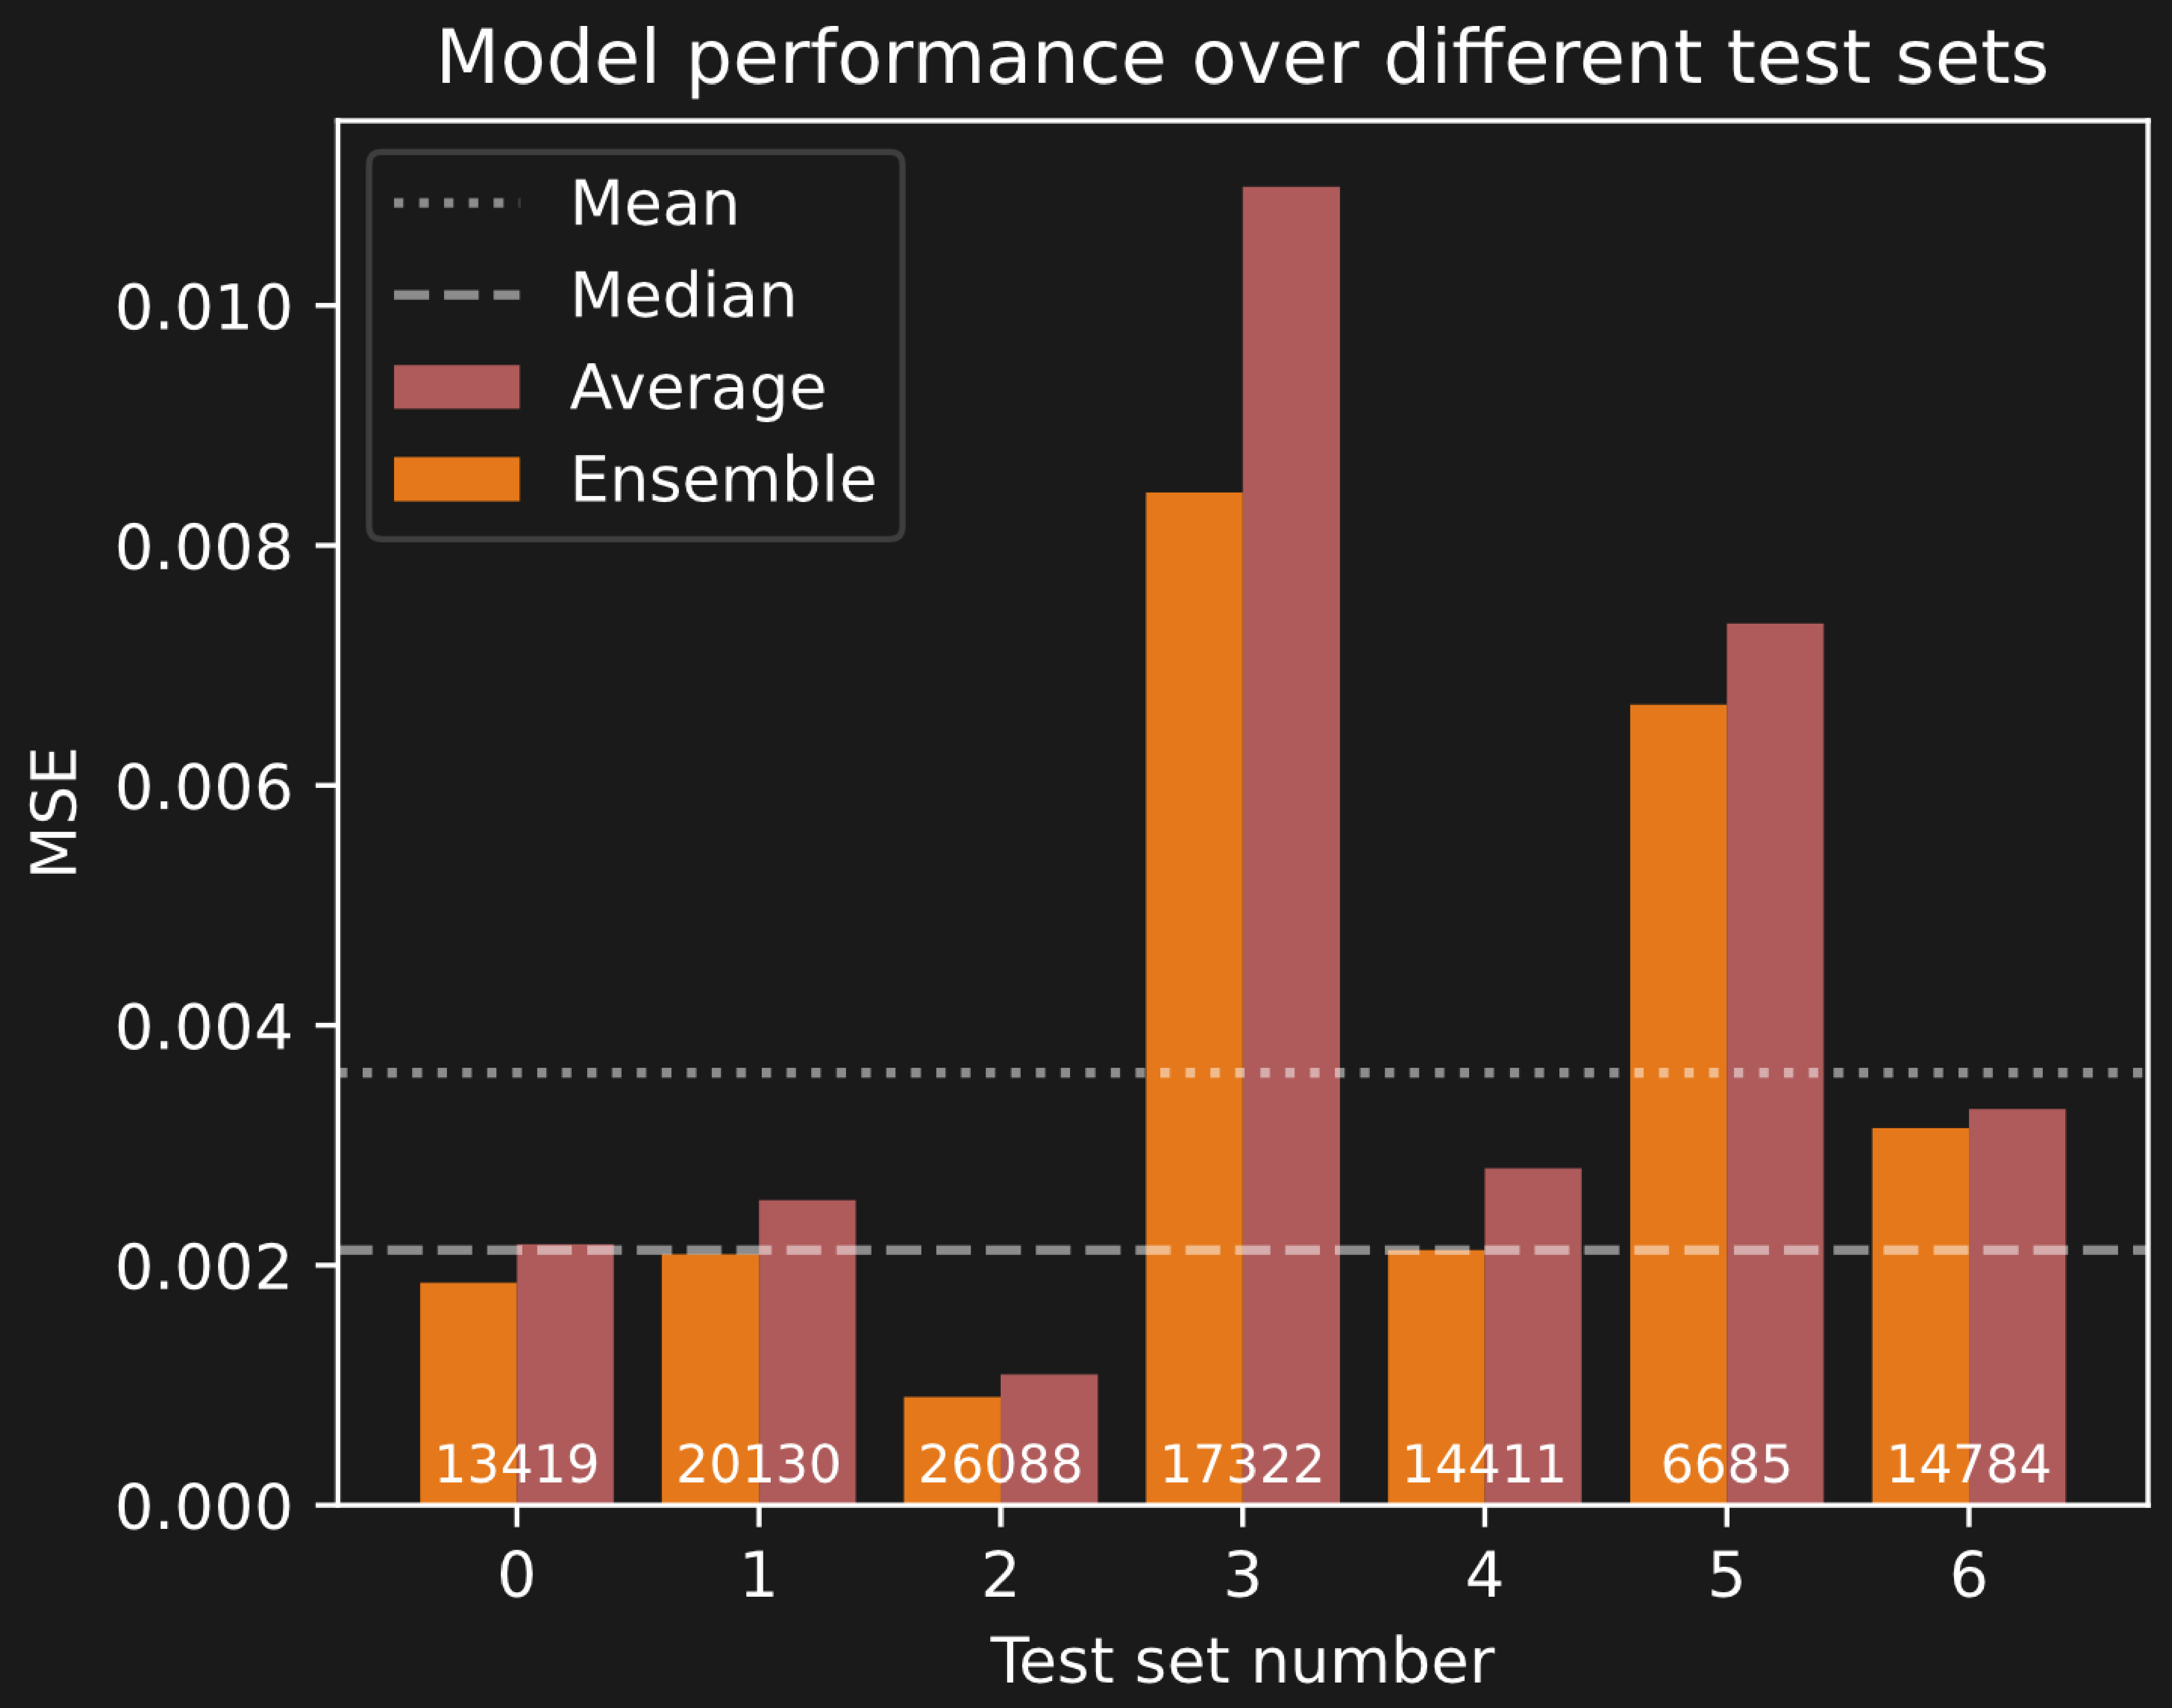

Inferred trends are intuitive and predictions agree with LAPD data

- Discharge voltage scan: agrees with intuition

Discharge voltage: effect on x and z profiles

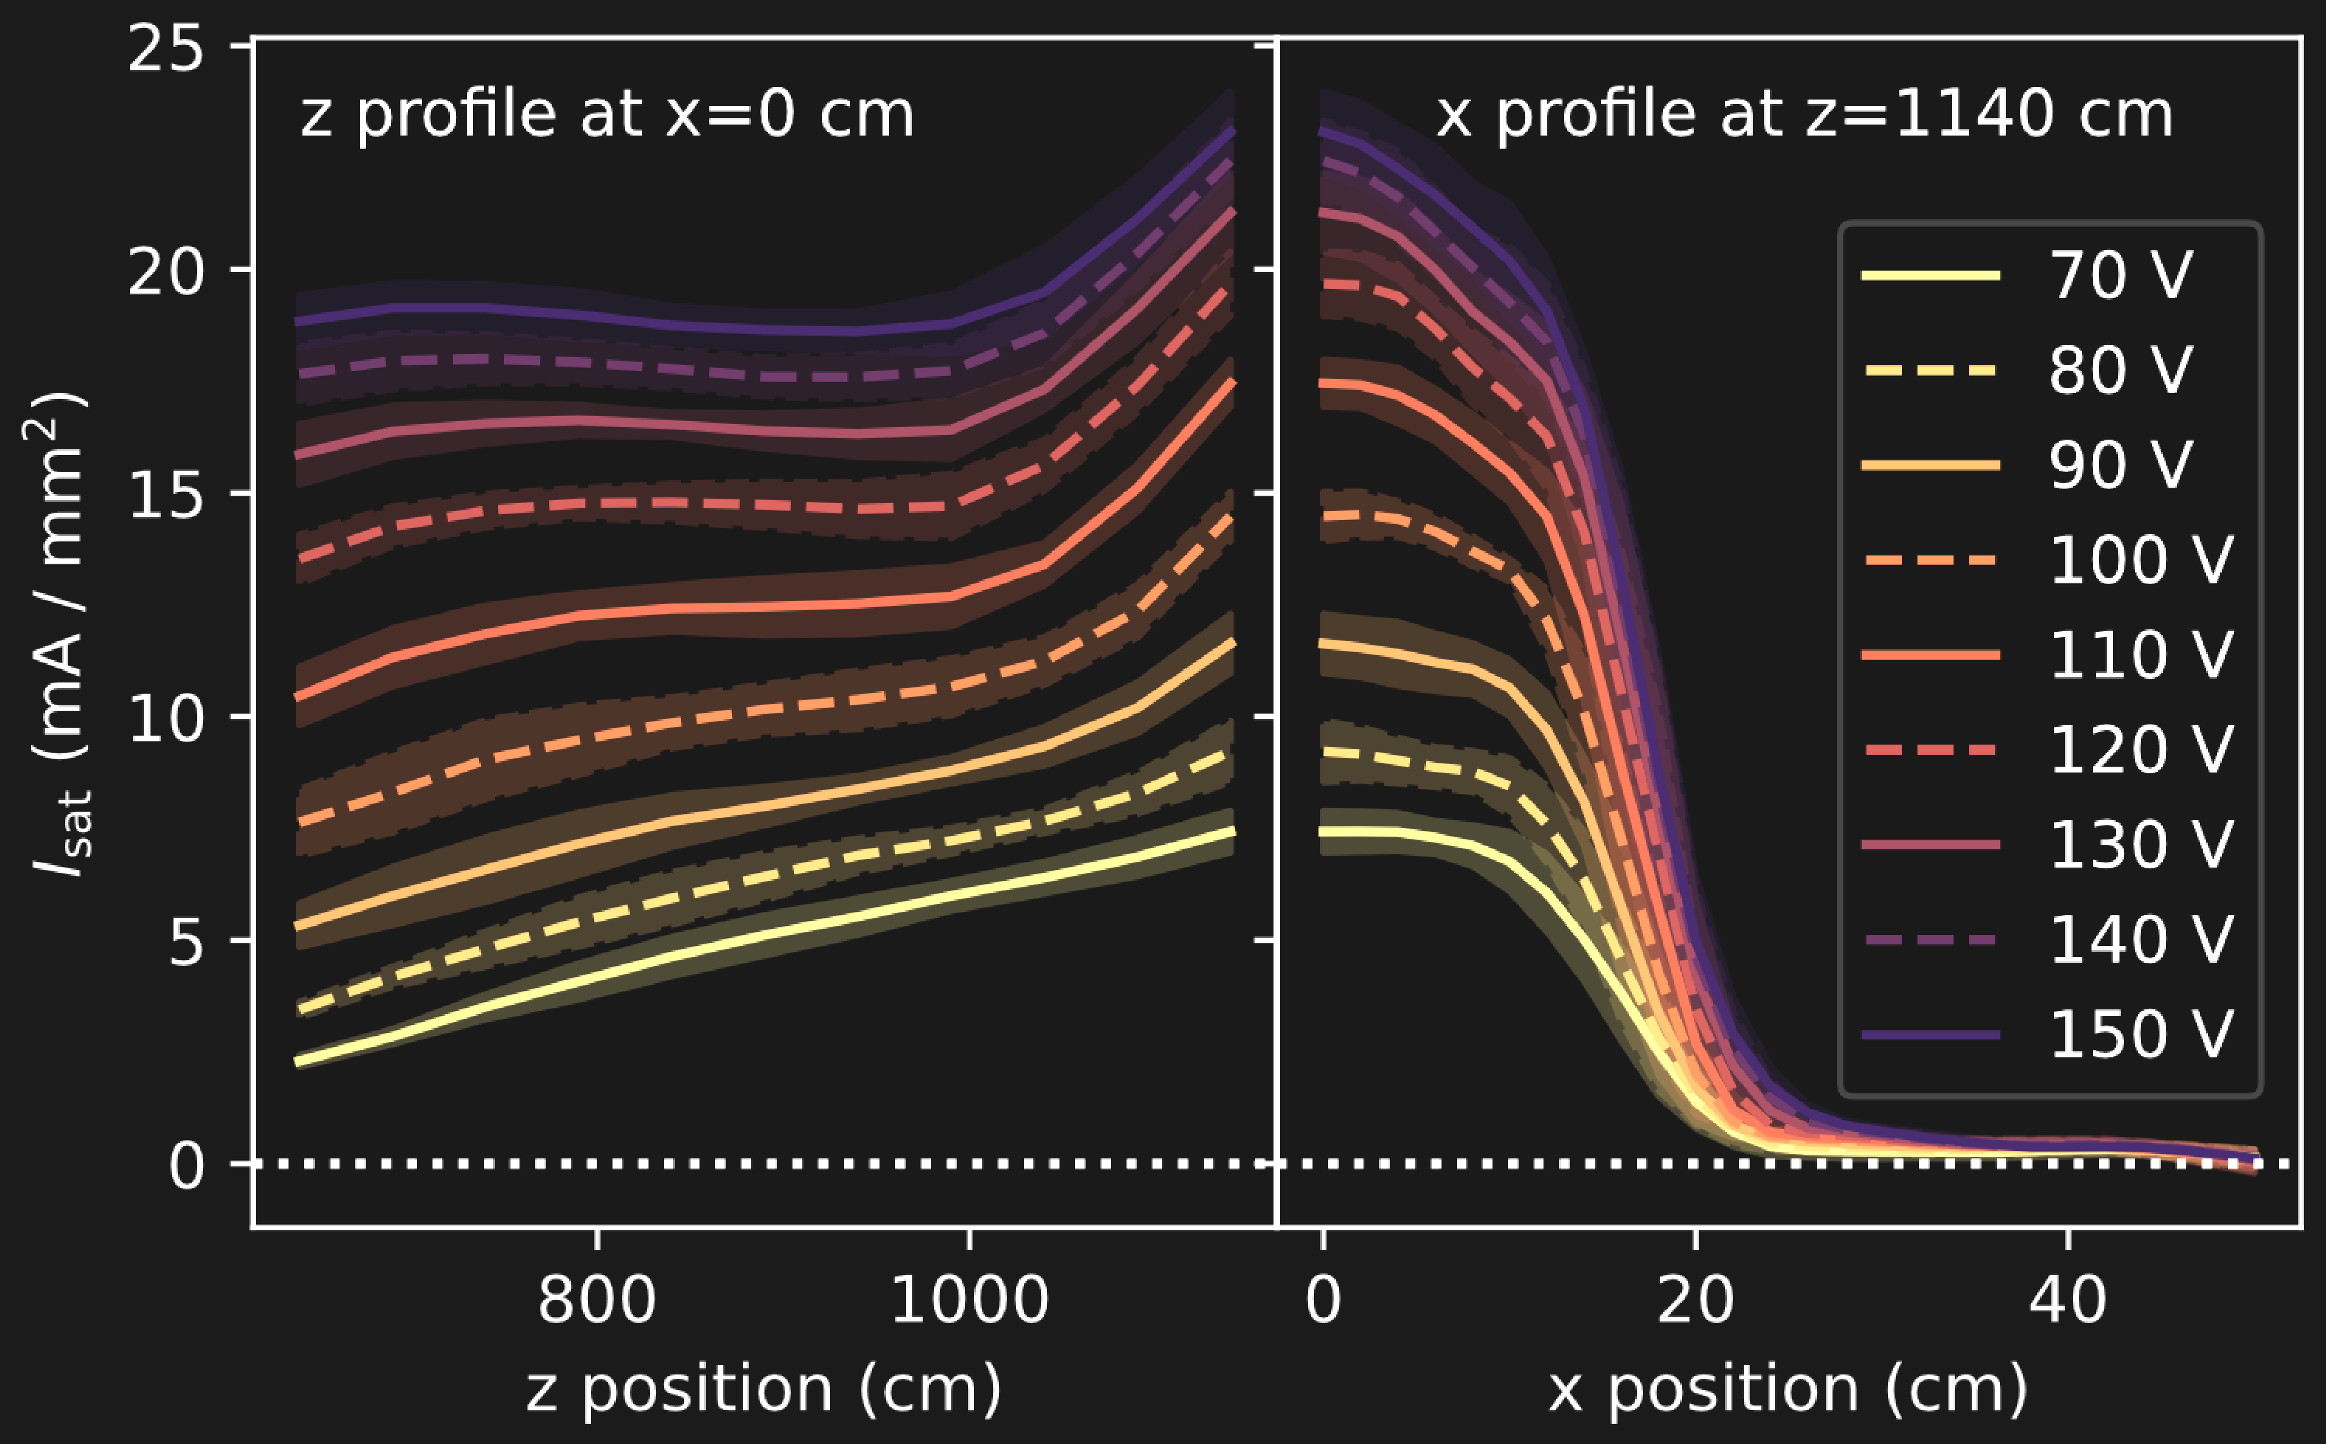

1 kG flat field, 38 ms gas puff

Measured and predicted I_{sat}

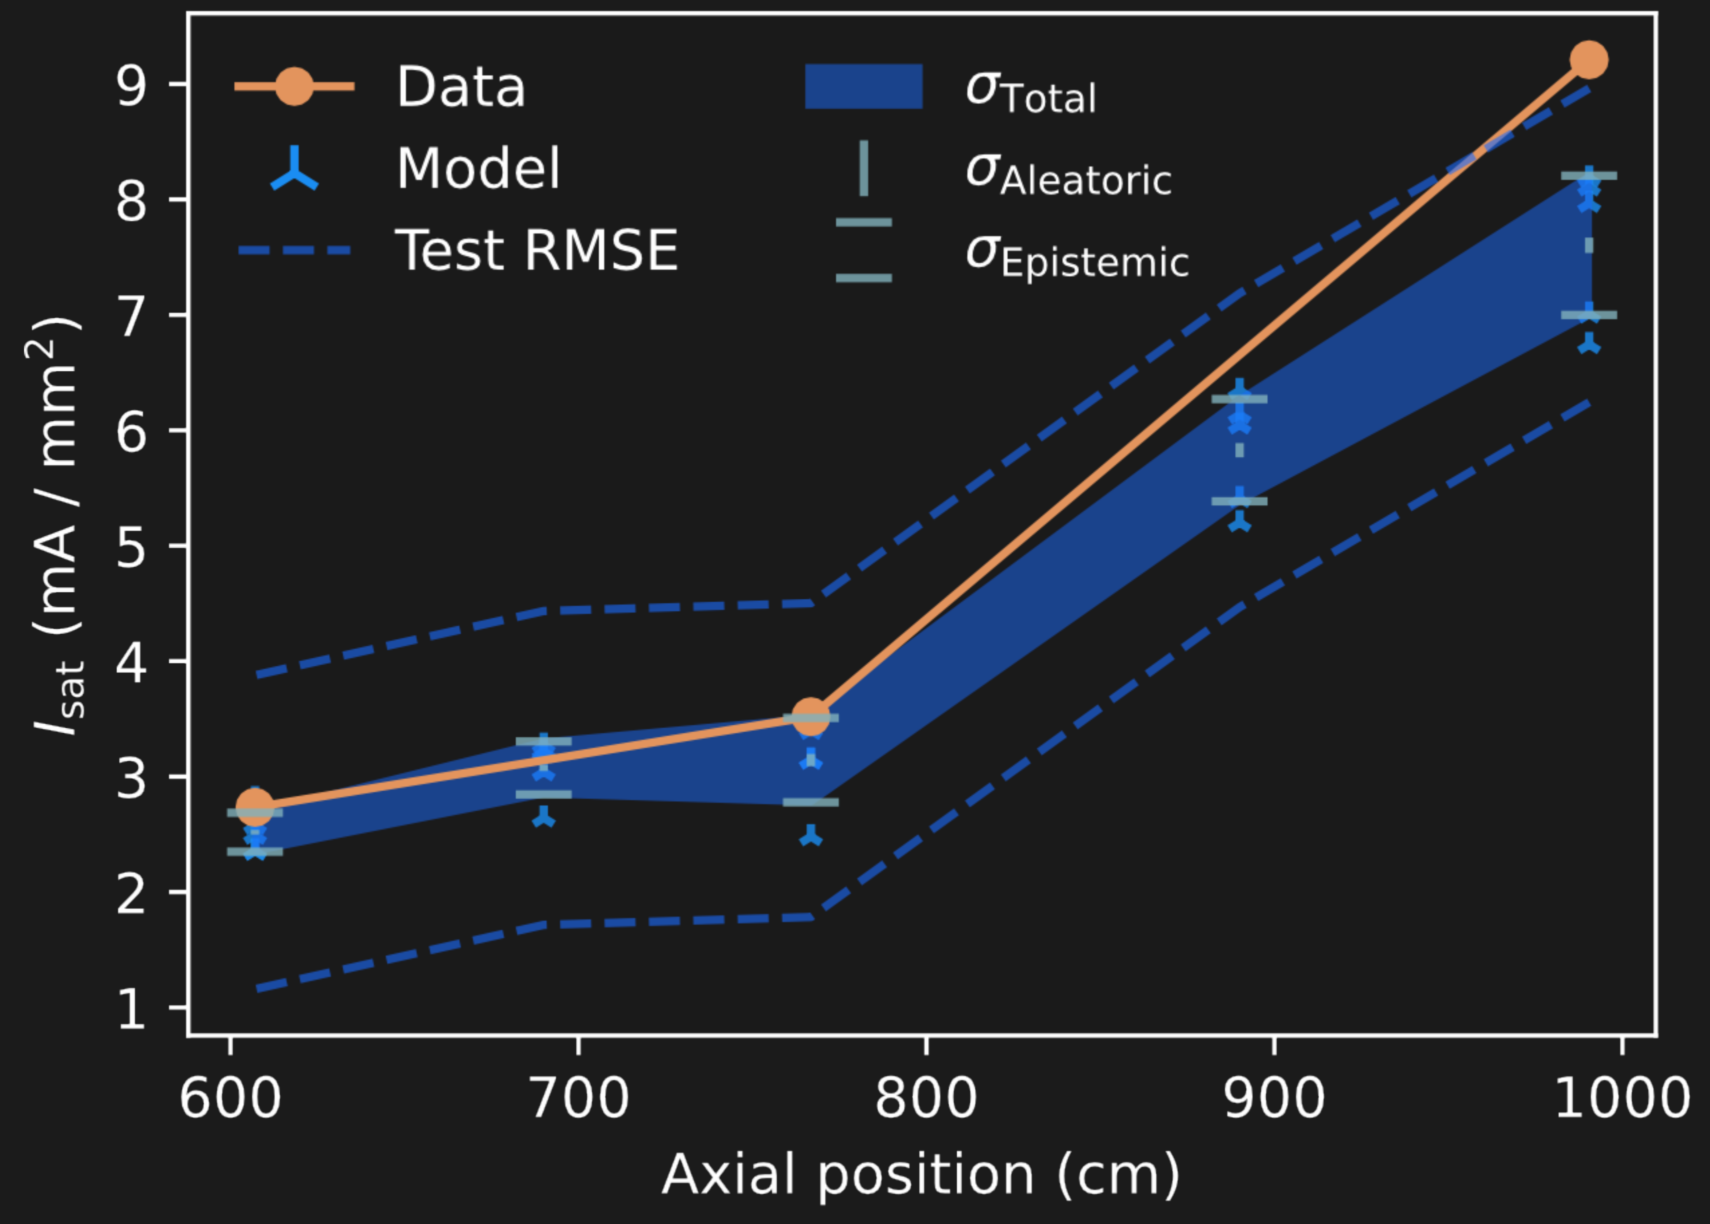

500G source, 500G mirror, 1500G midplane, 90V gas puff, 150V discharge, 38 ms gas puff

- Probes misaligned, but we can predict off-axis no problem

Model gives us optimized profile with constraints on Isat

- Important for LAPD: high densities with a flat profile
  - Comprehensive search for best and worst axial variation

$$\text{Inputs} = \underset{\text{Inputs} \neq z}{\arg\min}\, \underset{z}{\text{sd}}(I_{\text{sat}}|_{x=0})$$

  - Also constrain search for $I_{sat} > 7.5$ mA / mm²

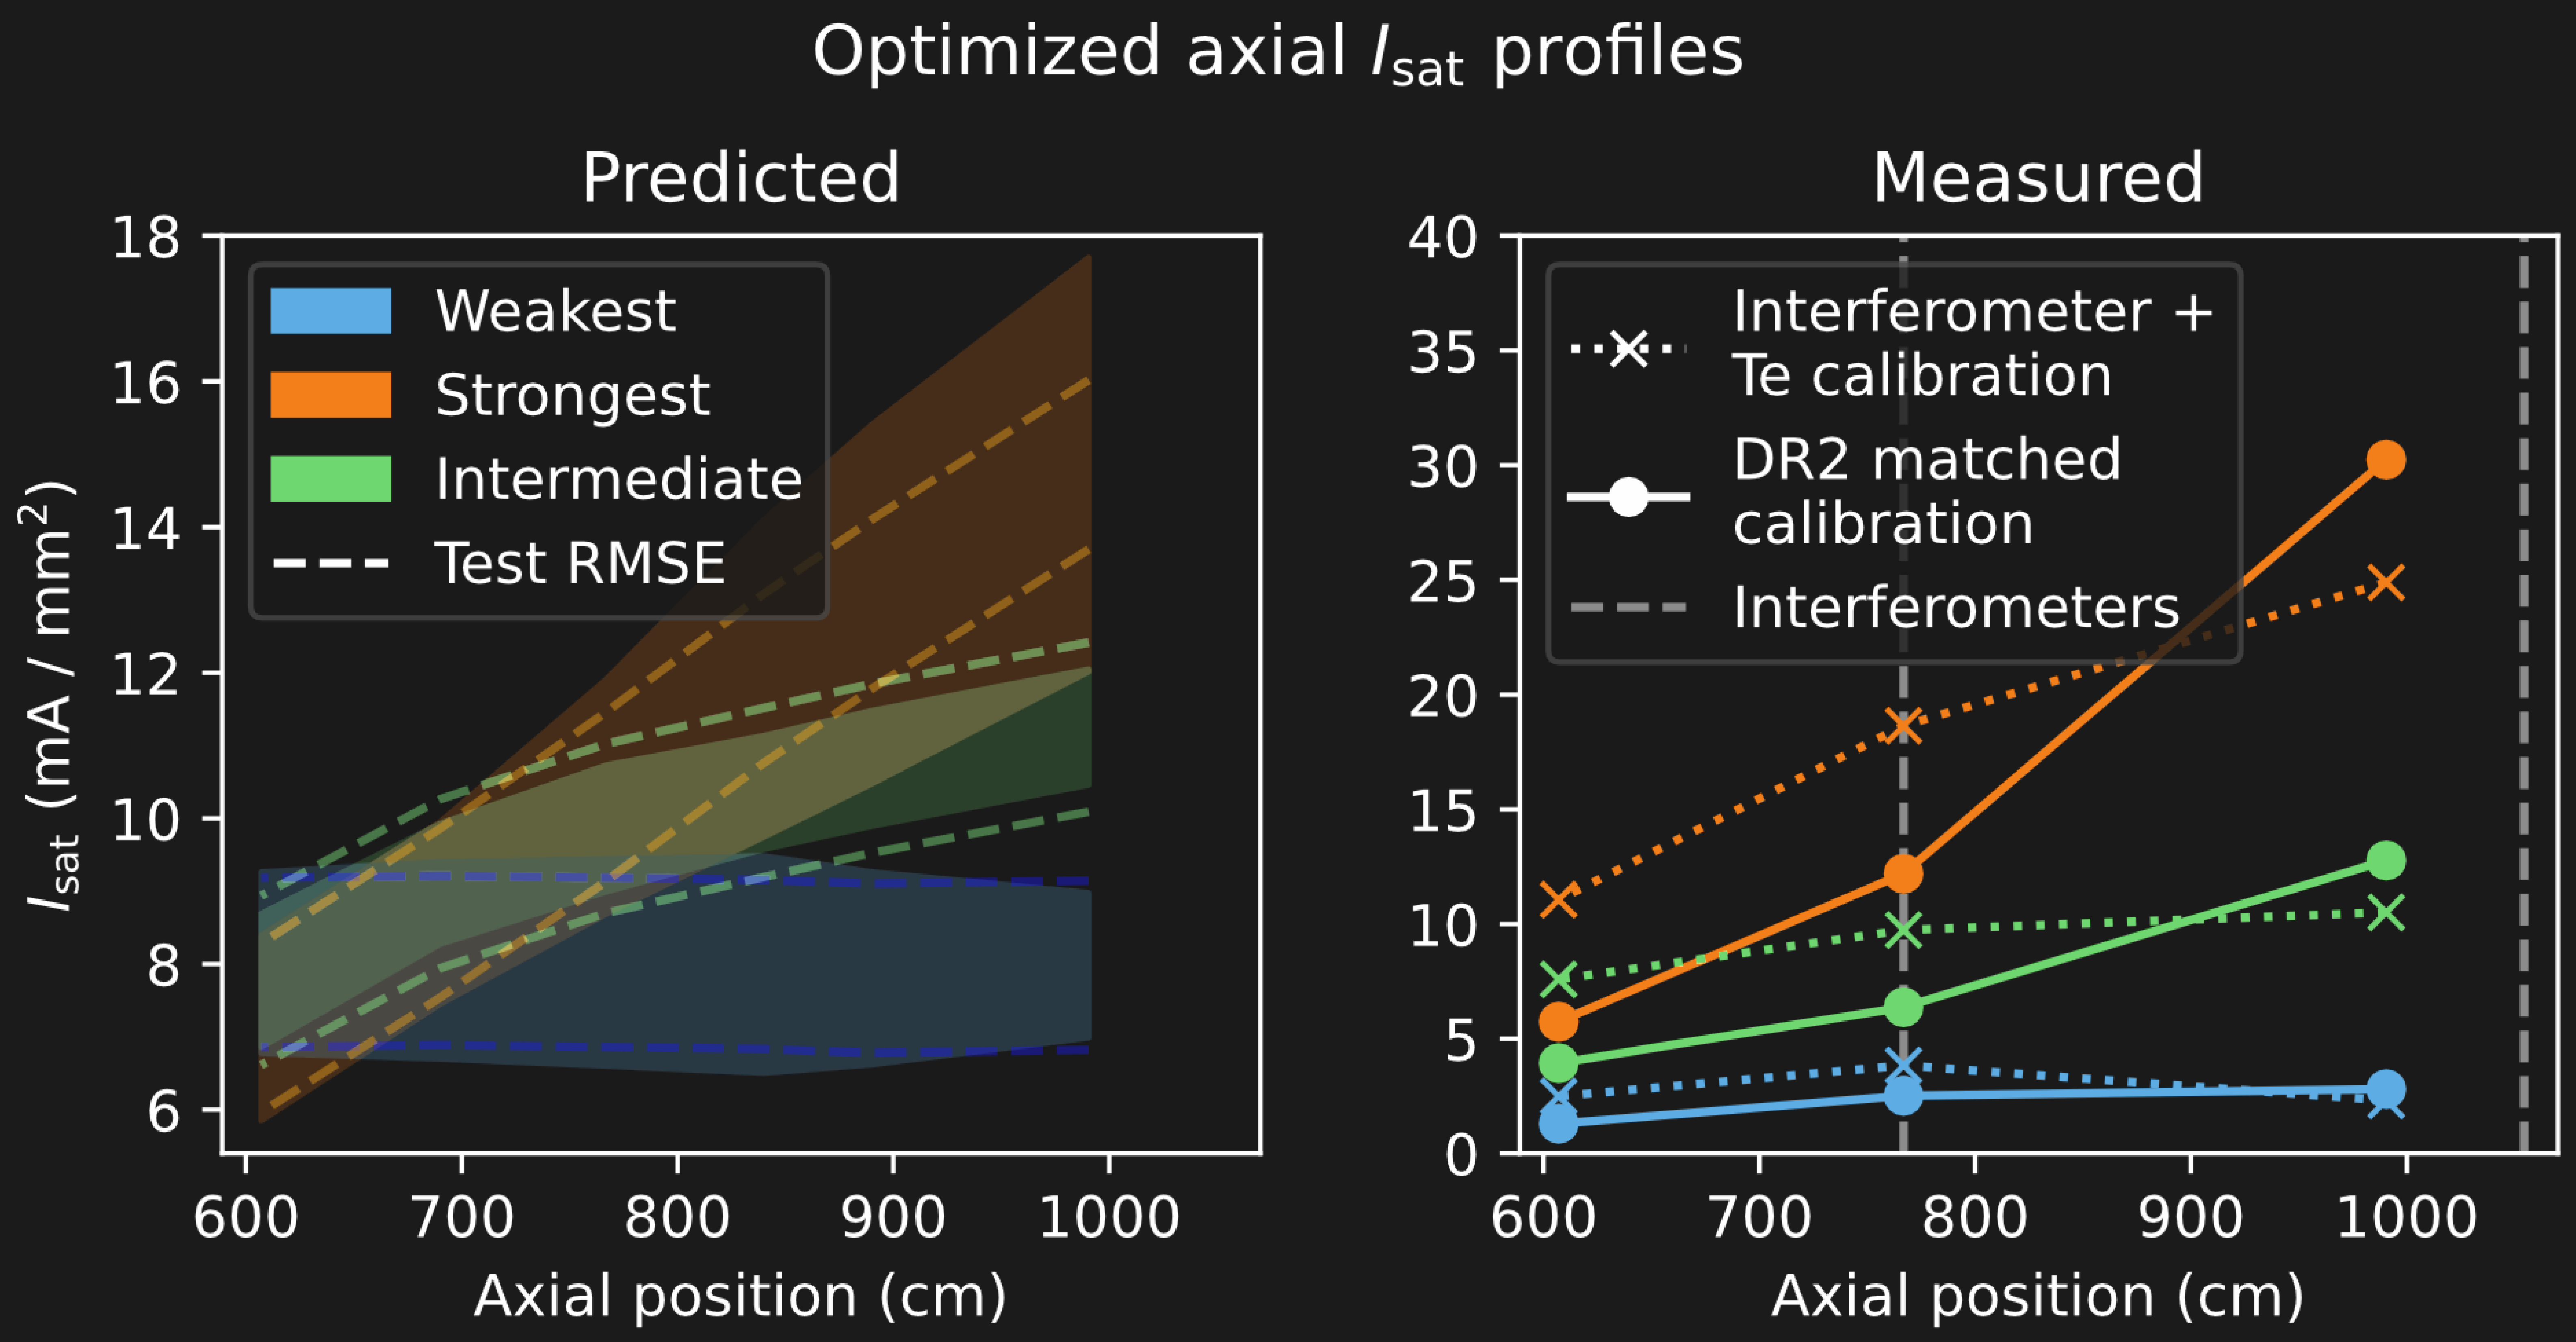

- Intermediate case: model learns trends in addition to extrema

We can predict Isat anywhere* in any* mirror configuration in the LAPD

* as long as it is reasonably within the bounds of the training data (and your standards aren't too high)

- Optimized the LAPD given any function of I_{sat}
- This work is quite novel:
  - trend inference using NNs
  - random generation of machine configurations
  - thorough uncertainty quantification

GitHub link

github.com/physicistphil/lapd-isat-predict

What if we want to reconstruct any input or diagnostic, not just I_{sat}?

Energy based models

Energy-based models learn a probability distribution over the data

$$p(x) \sim e^{-E(x)}$$ <- generative ML model

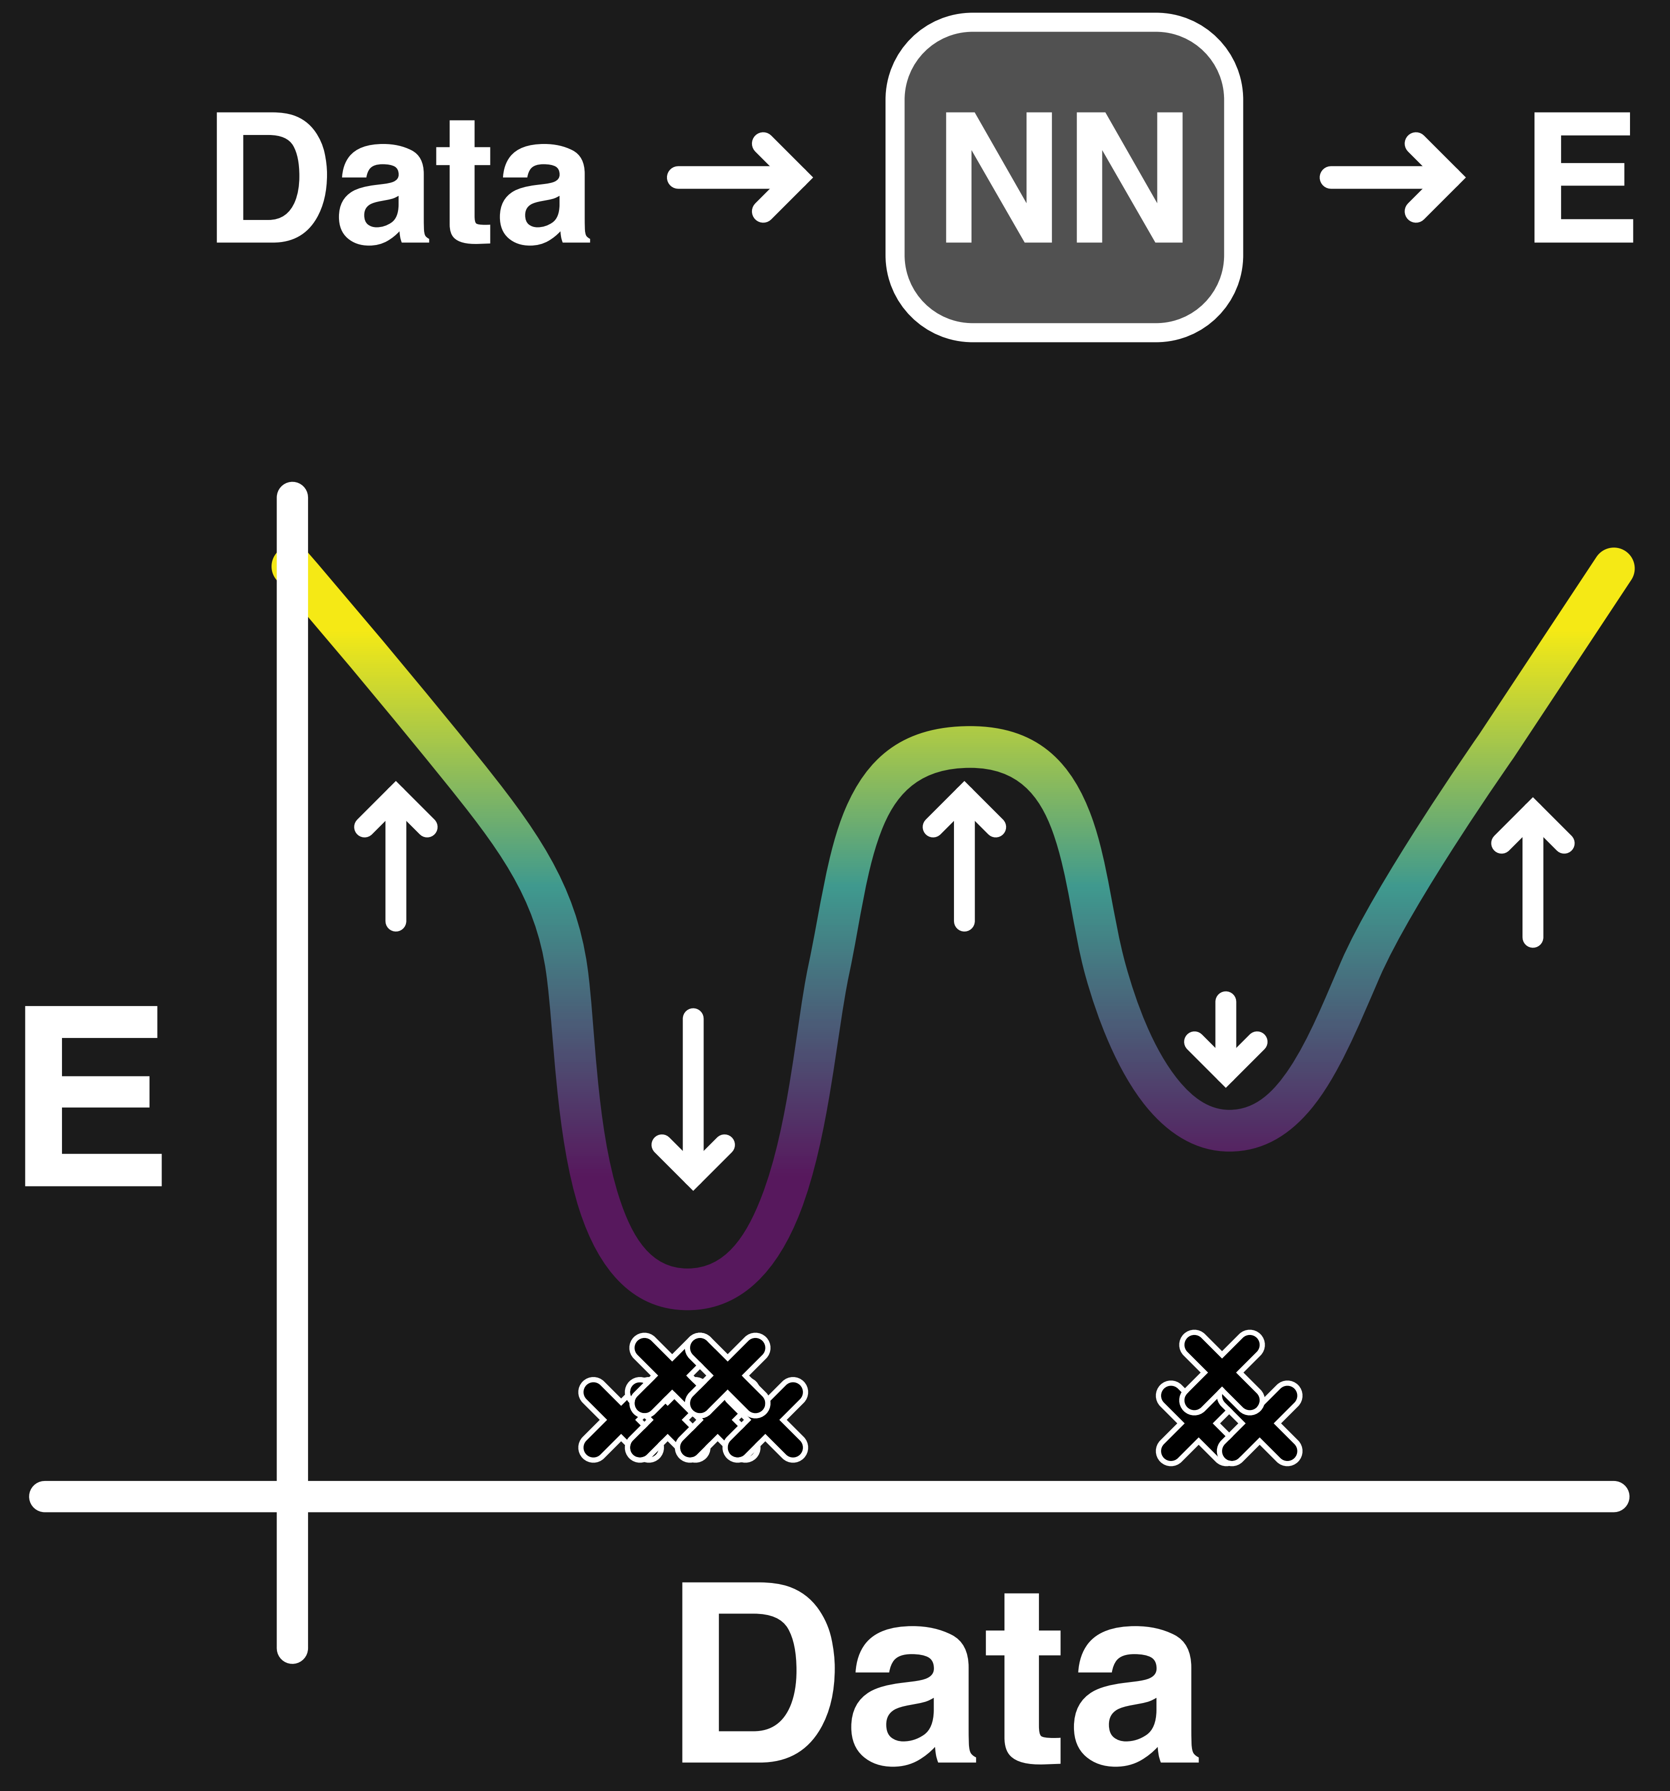

Sampled via Langevin dynamics:

$$\tilde{x}_i^{\ell} \leftarrow \tilde{x}_i^{\ell-1} - \frac{\epsilon^2}{2} \nabla_x E_\theta(\tilde{x}_i^{\ell-1}) + \epsilon \mathcal{N}(0,1)$$

Energy gradient

Gaussian noise

$\tilde{x}$: sample
ℓ: step in the MCMC chain
i: example number
ϵ: step size

- Sampling is inference!

Training an EBM molds the surface, sampling finds the minima

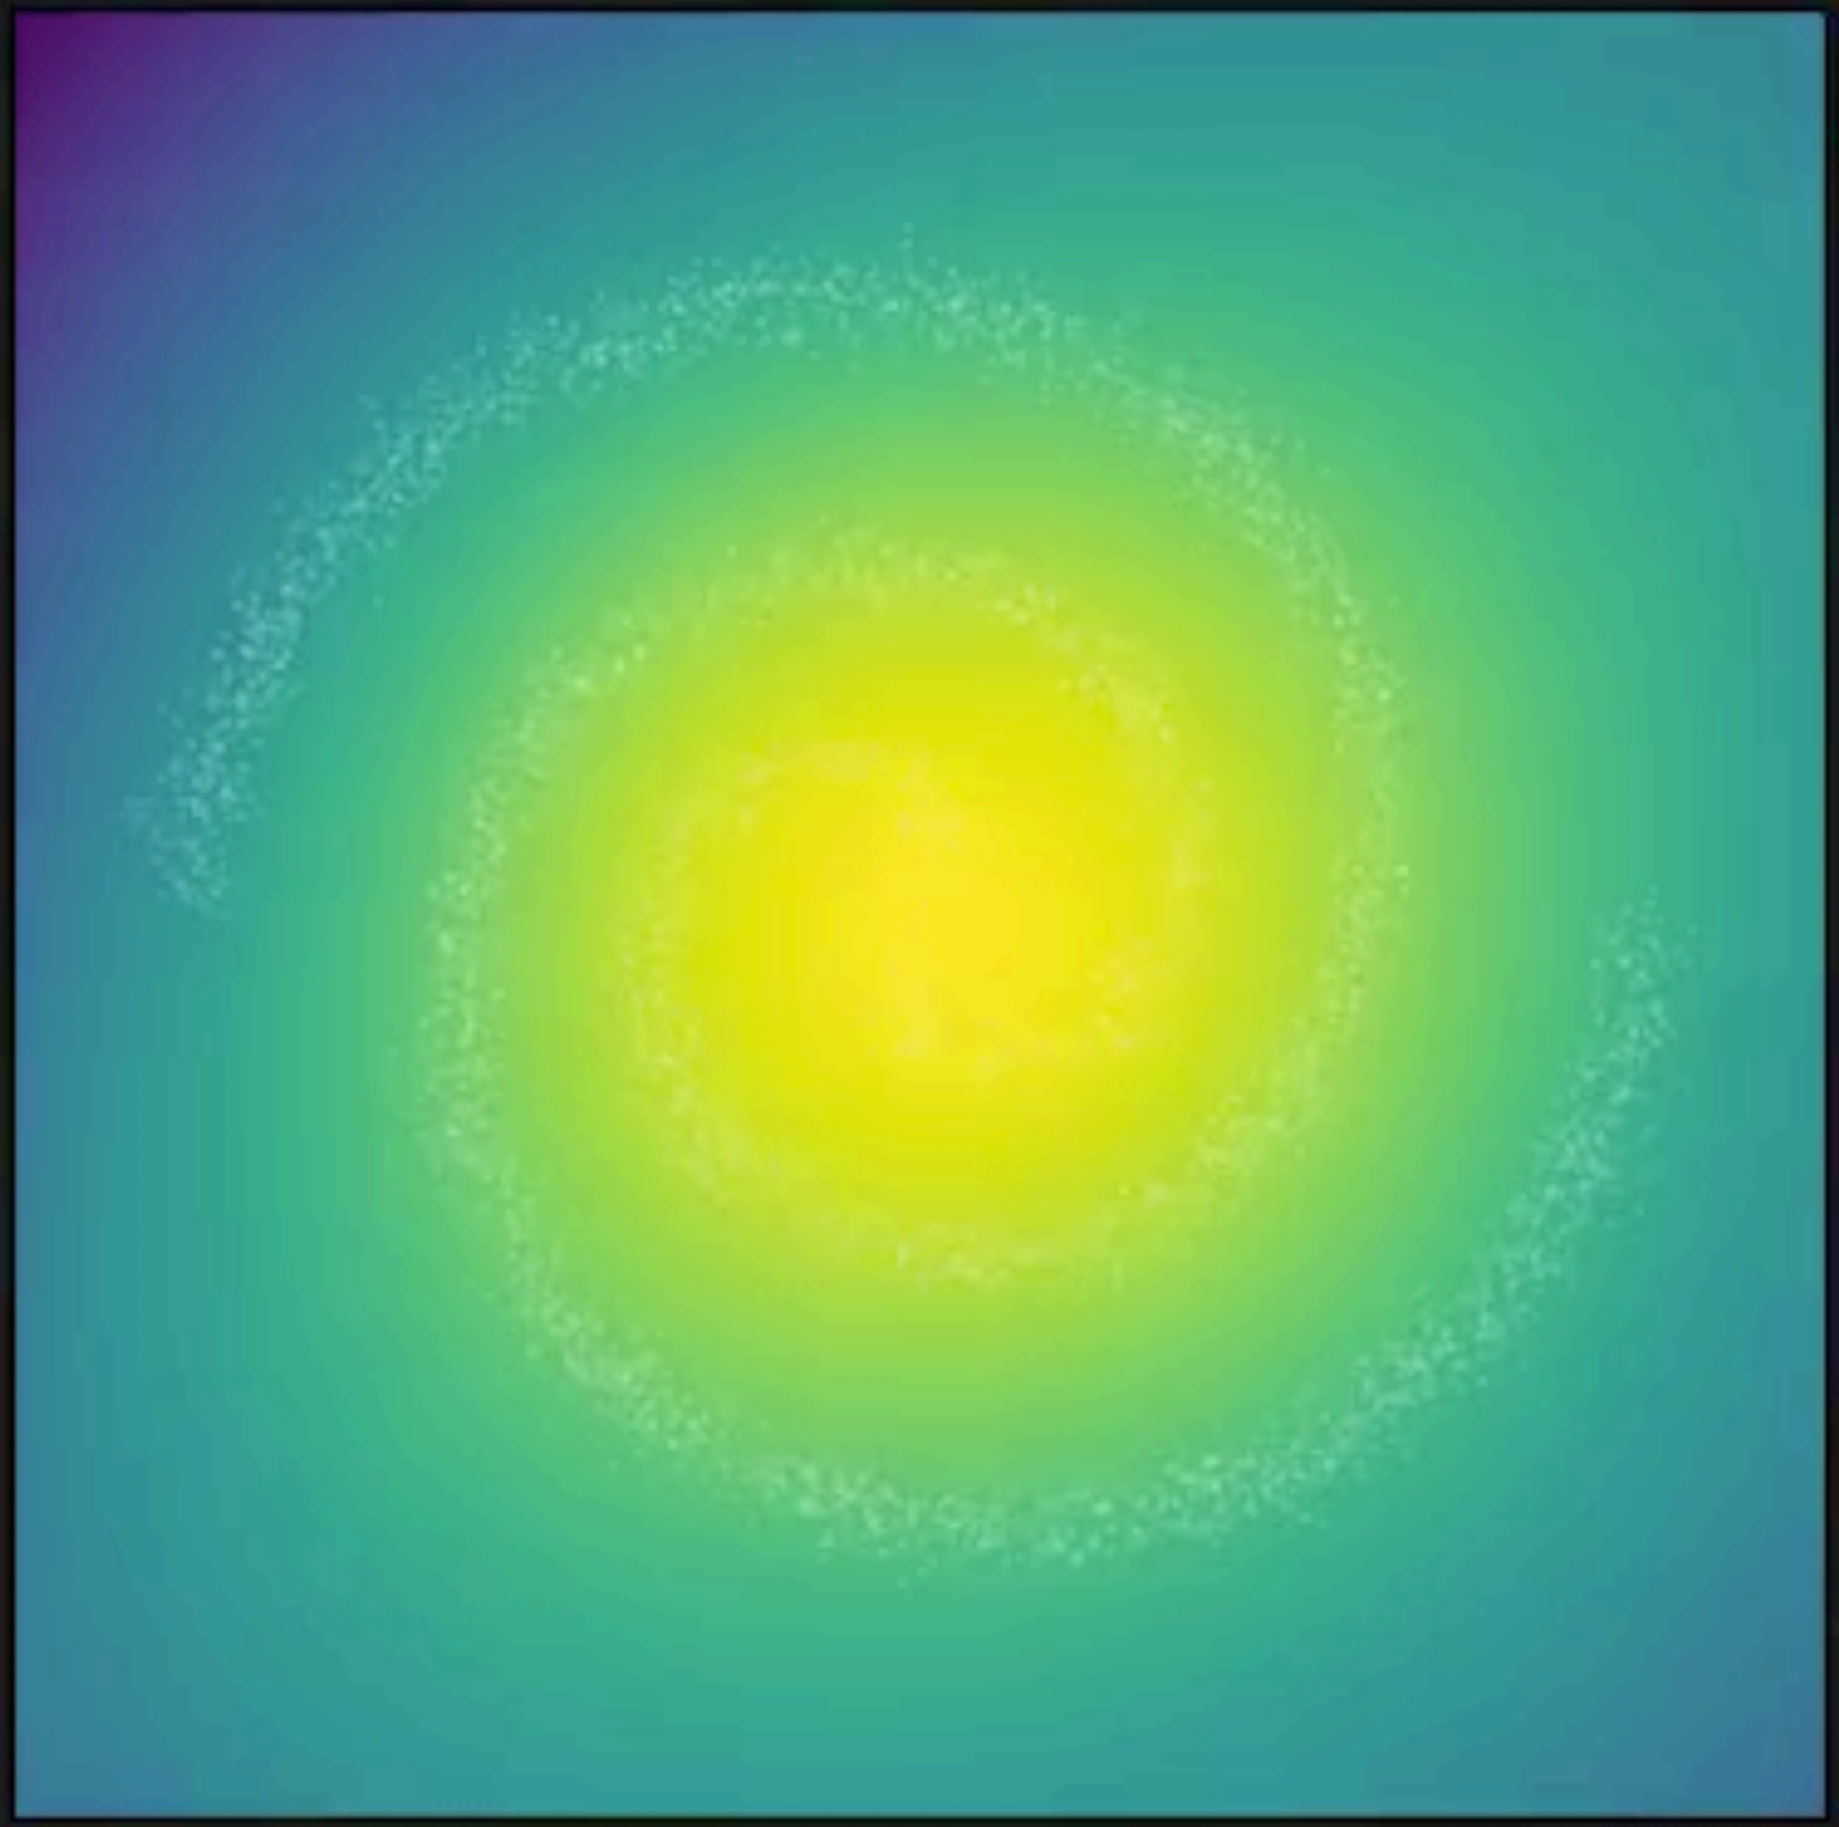

Training

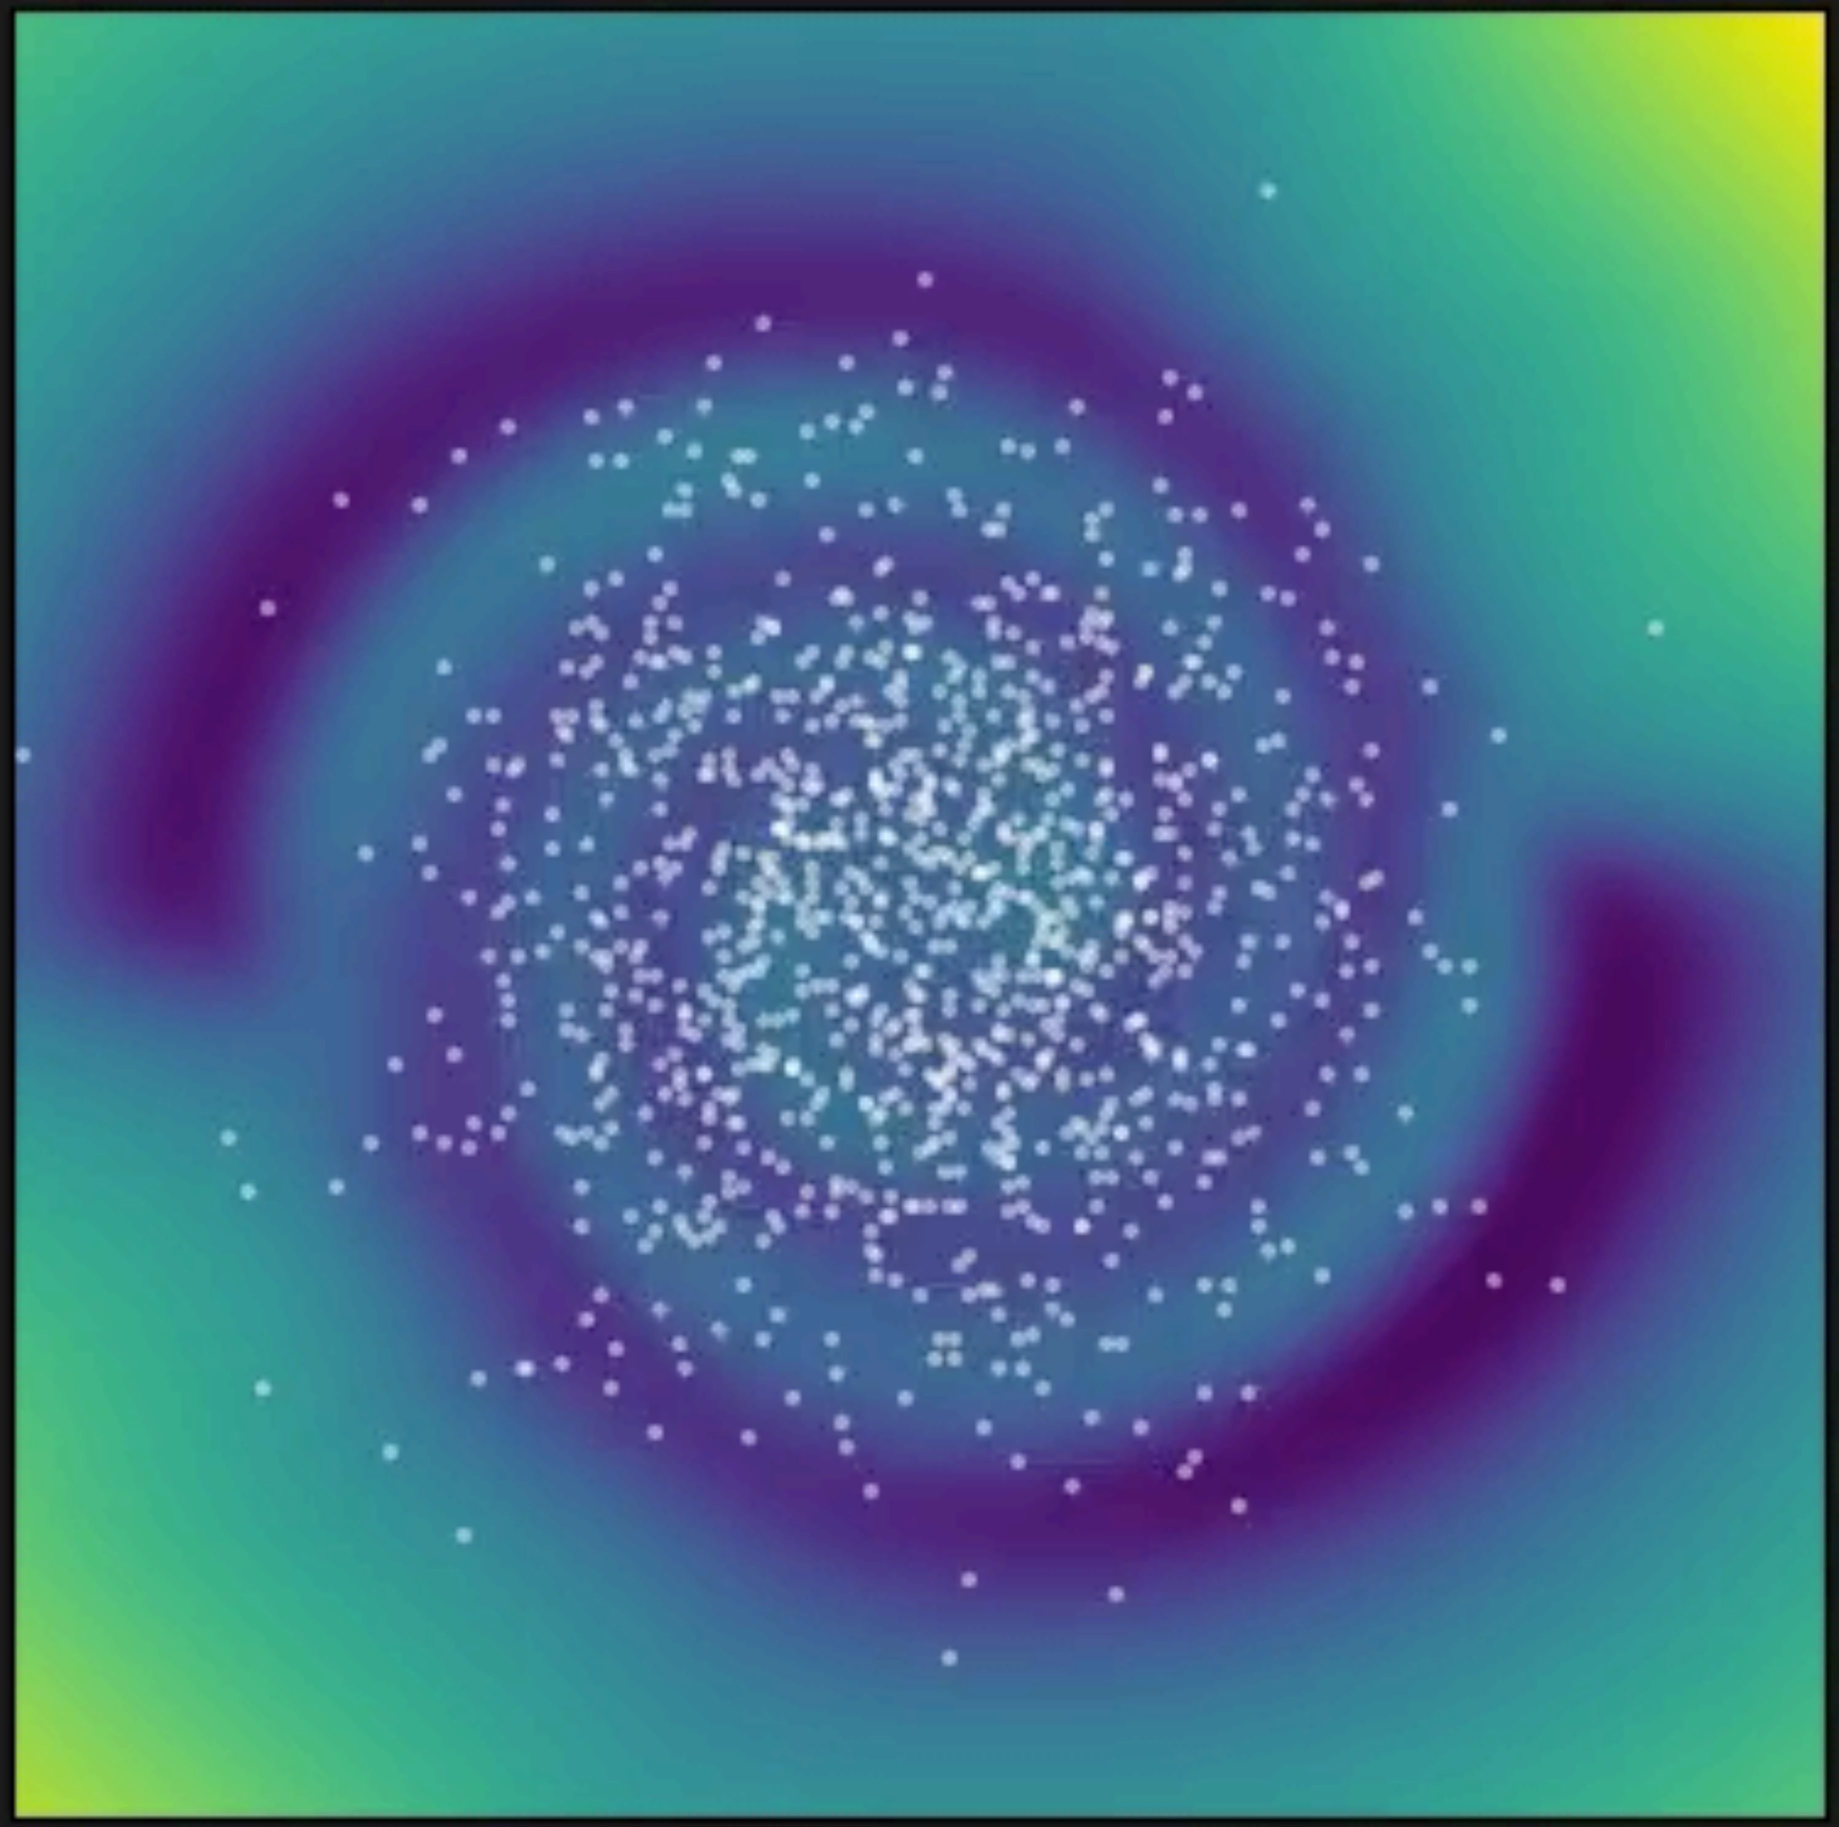

(Unconditional) Sampling

For this study: used the same dataset (with time series) and a larger model

Input: machine settings, time series data → **NN** → **Output**: energy

- 699 inputs into the model (up from 12)
  - Time series: discharge I and V, diodes, interferometer, I_{sat}
  - Magnetic field, gas info, probe positions, flags
- Model: ~14.7 million parameters
- Utilized CNNs and attention (transformer-like) blocks
- Multi-modal model: intermediate and hybrid fusion

Conditional sampling performed by freezing gradients performs poorly

$$\text{Ifo} \sim p(\text{Ifo} \mid \text{All other inputs})$$

$$\tilde{x}_i^{\ell} \leftarrow \tilde{x}_i^{\ell-1} - \frac{\epsilon^2}{2} \nabla_x E_\theta(\tilde{x}_i^{\ell-1}) + \epsilon \mathcal{N}(0,1)$$

Freeze conditional inputs on real data

- Approach used in the literature for conditional sampling
- Yields unphysical results: negative interferometer signals

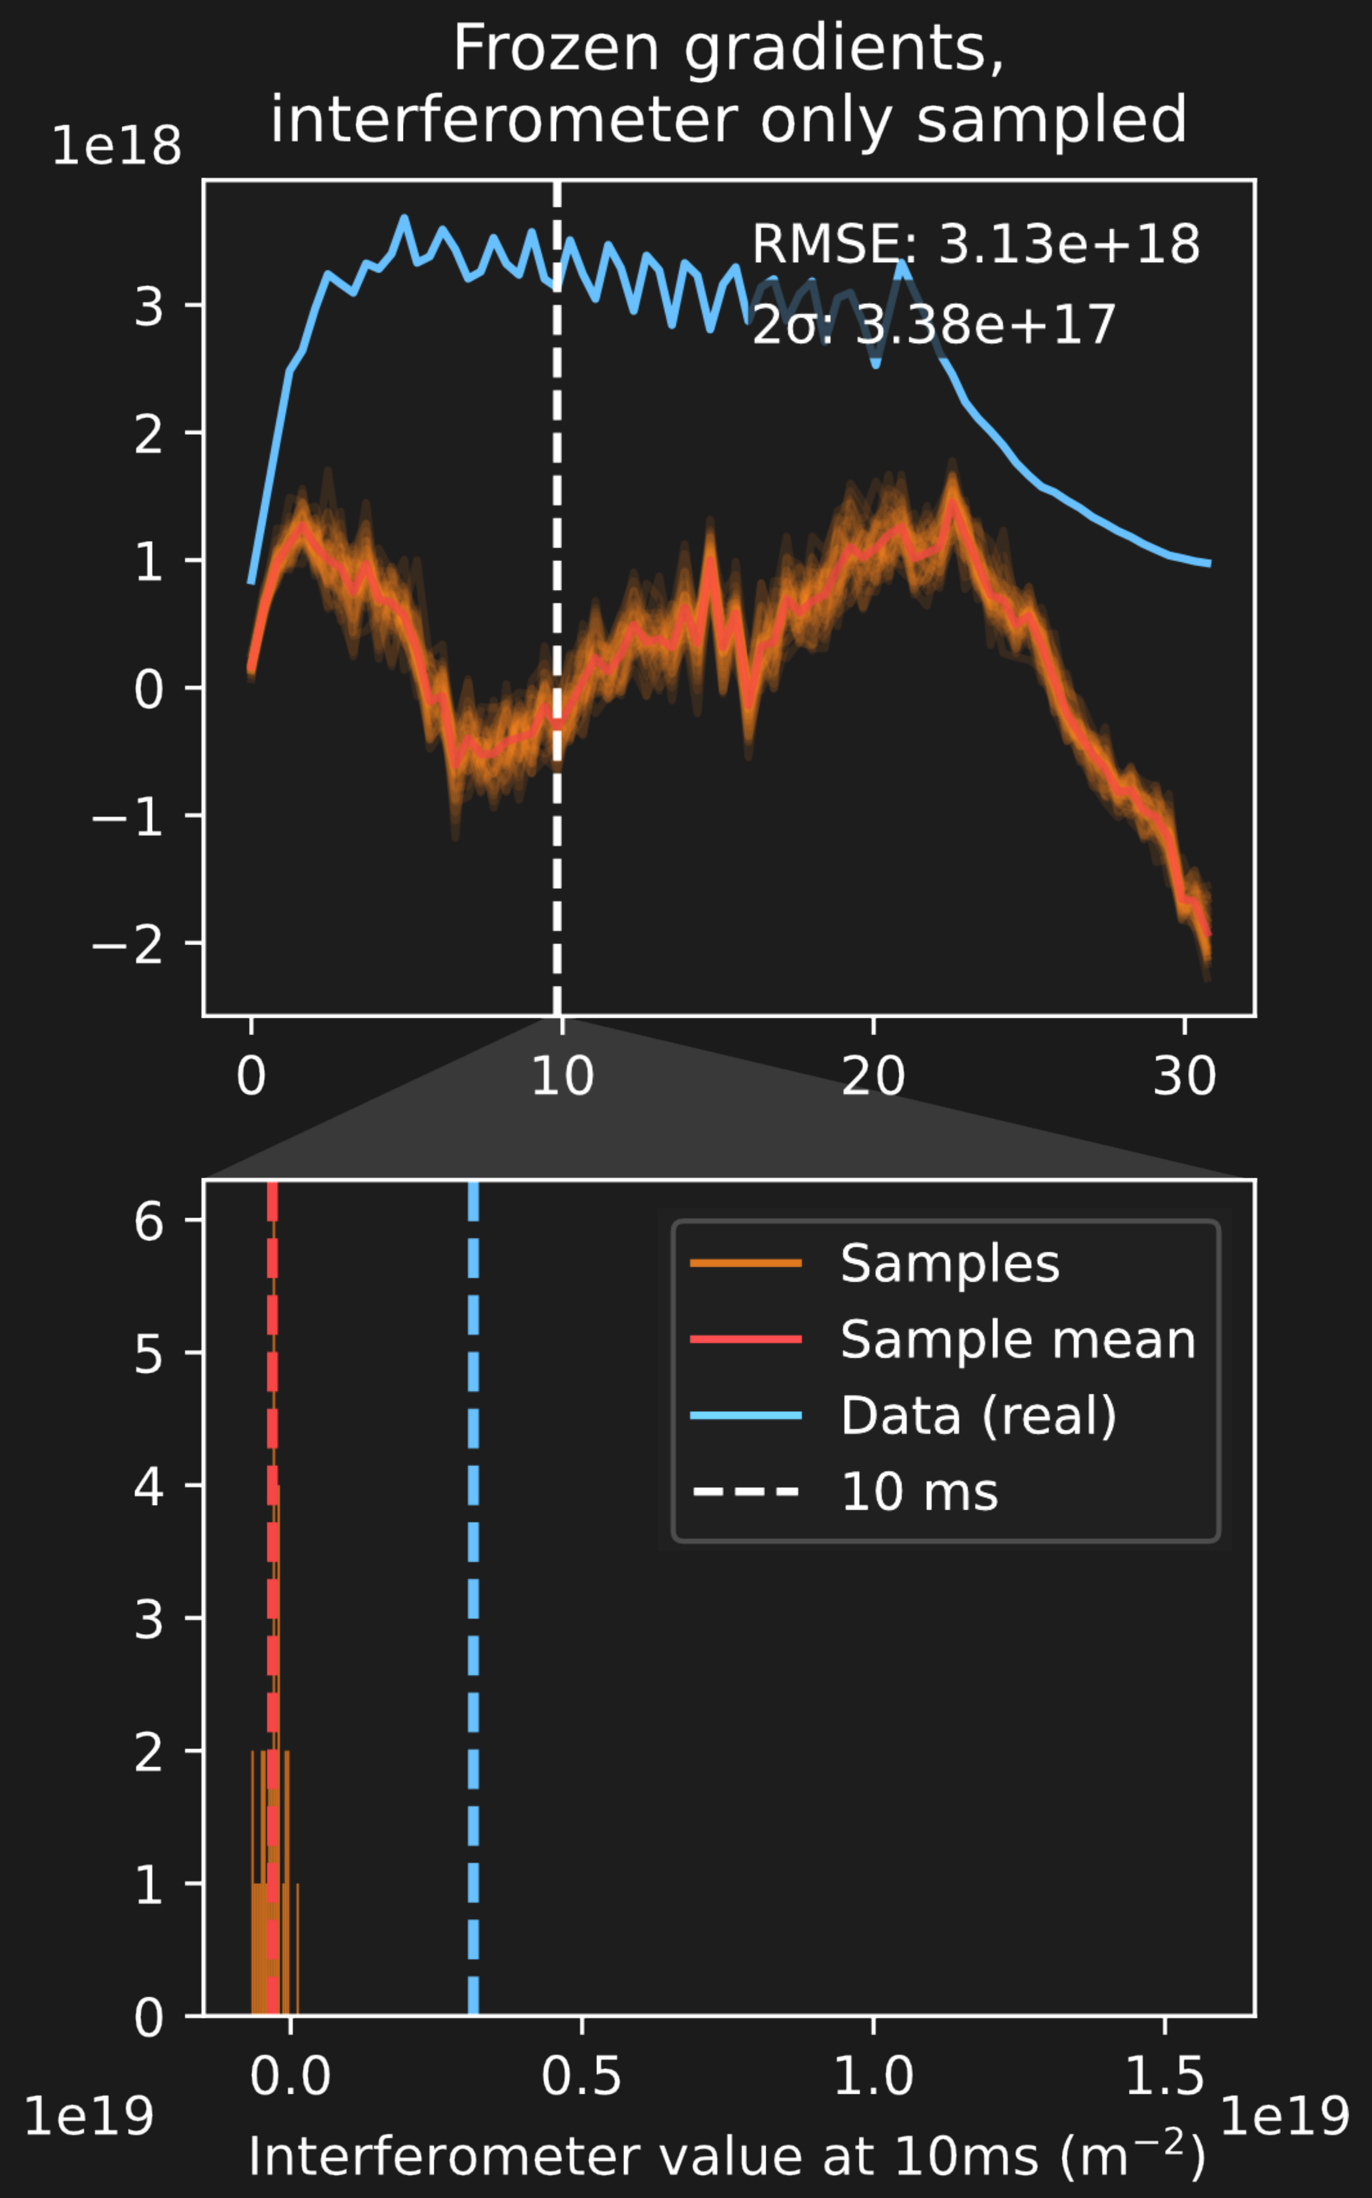

Modifying the energy function for conditional sampling works well

$$E_{\text{cond}}(\tilde{x}) = E(\tilde{x}) + F(\tilde{x}), \quad F(\tilde{x}) = \left(\frac{\tilde{x} - x_i}{2\epsilon}\right)^2$$

$p(\tilde{x}) \sim e^{-E(\tilde{x})}$ —> constraining samples via Gaussian

- Generated realistic samples
- Distribution is reasonable
- Novel method in ML community
- EBMs are composable

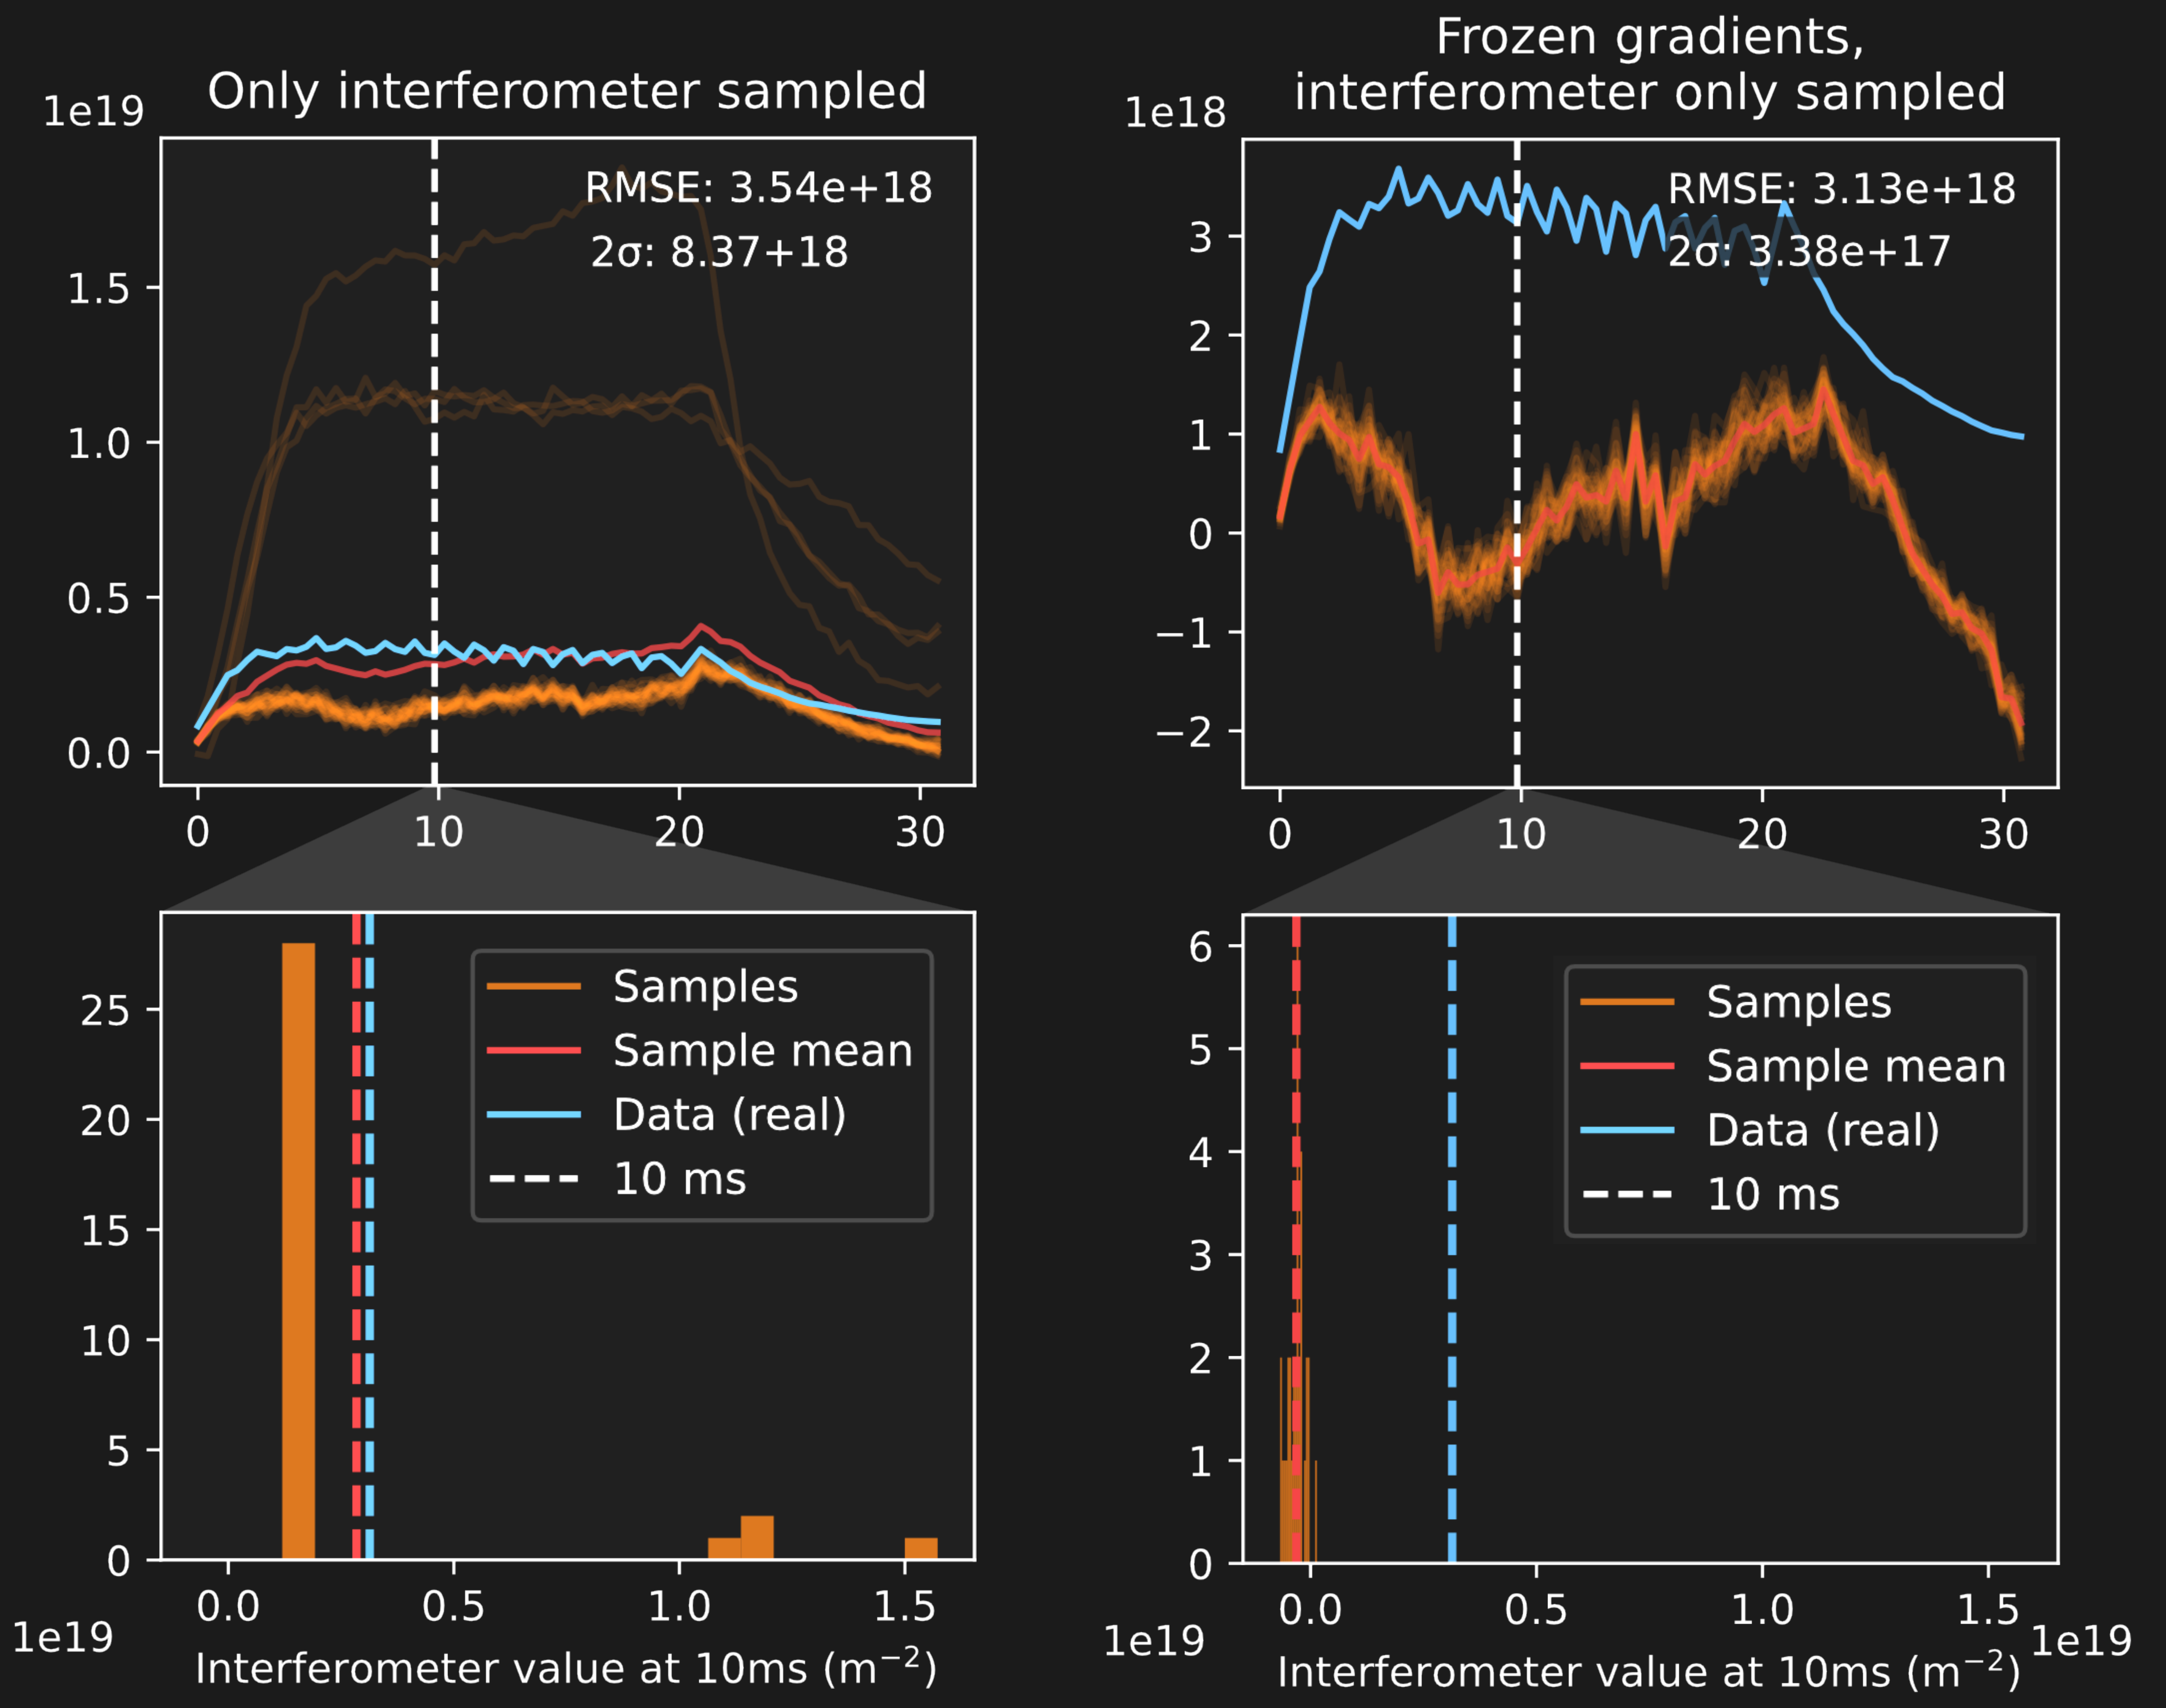

Supplying additional inputs improves diagnostic reconstruction

Given:	LAPD settings only	All signals
RMSE (test set)	4.12×10^{18}	2.91×10^{18}
RMSE (DR2_02)	3.77×10^{18}	3.54×10^{18}

Improves

- More signals improve mean prediction
- More signals better constrain the interferometer distribution
- Get a free uncertainty metric

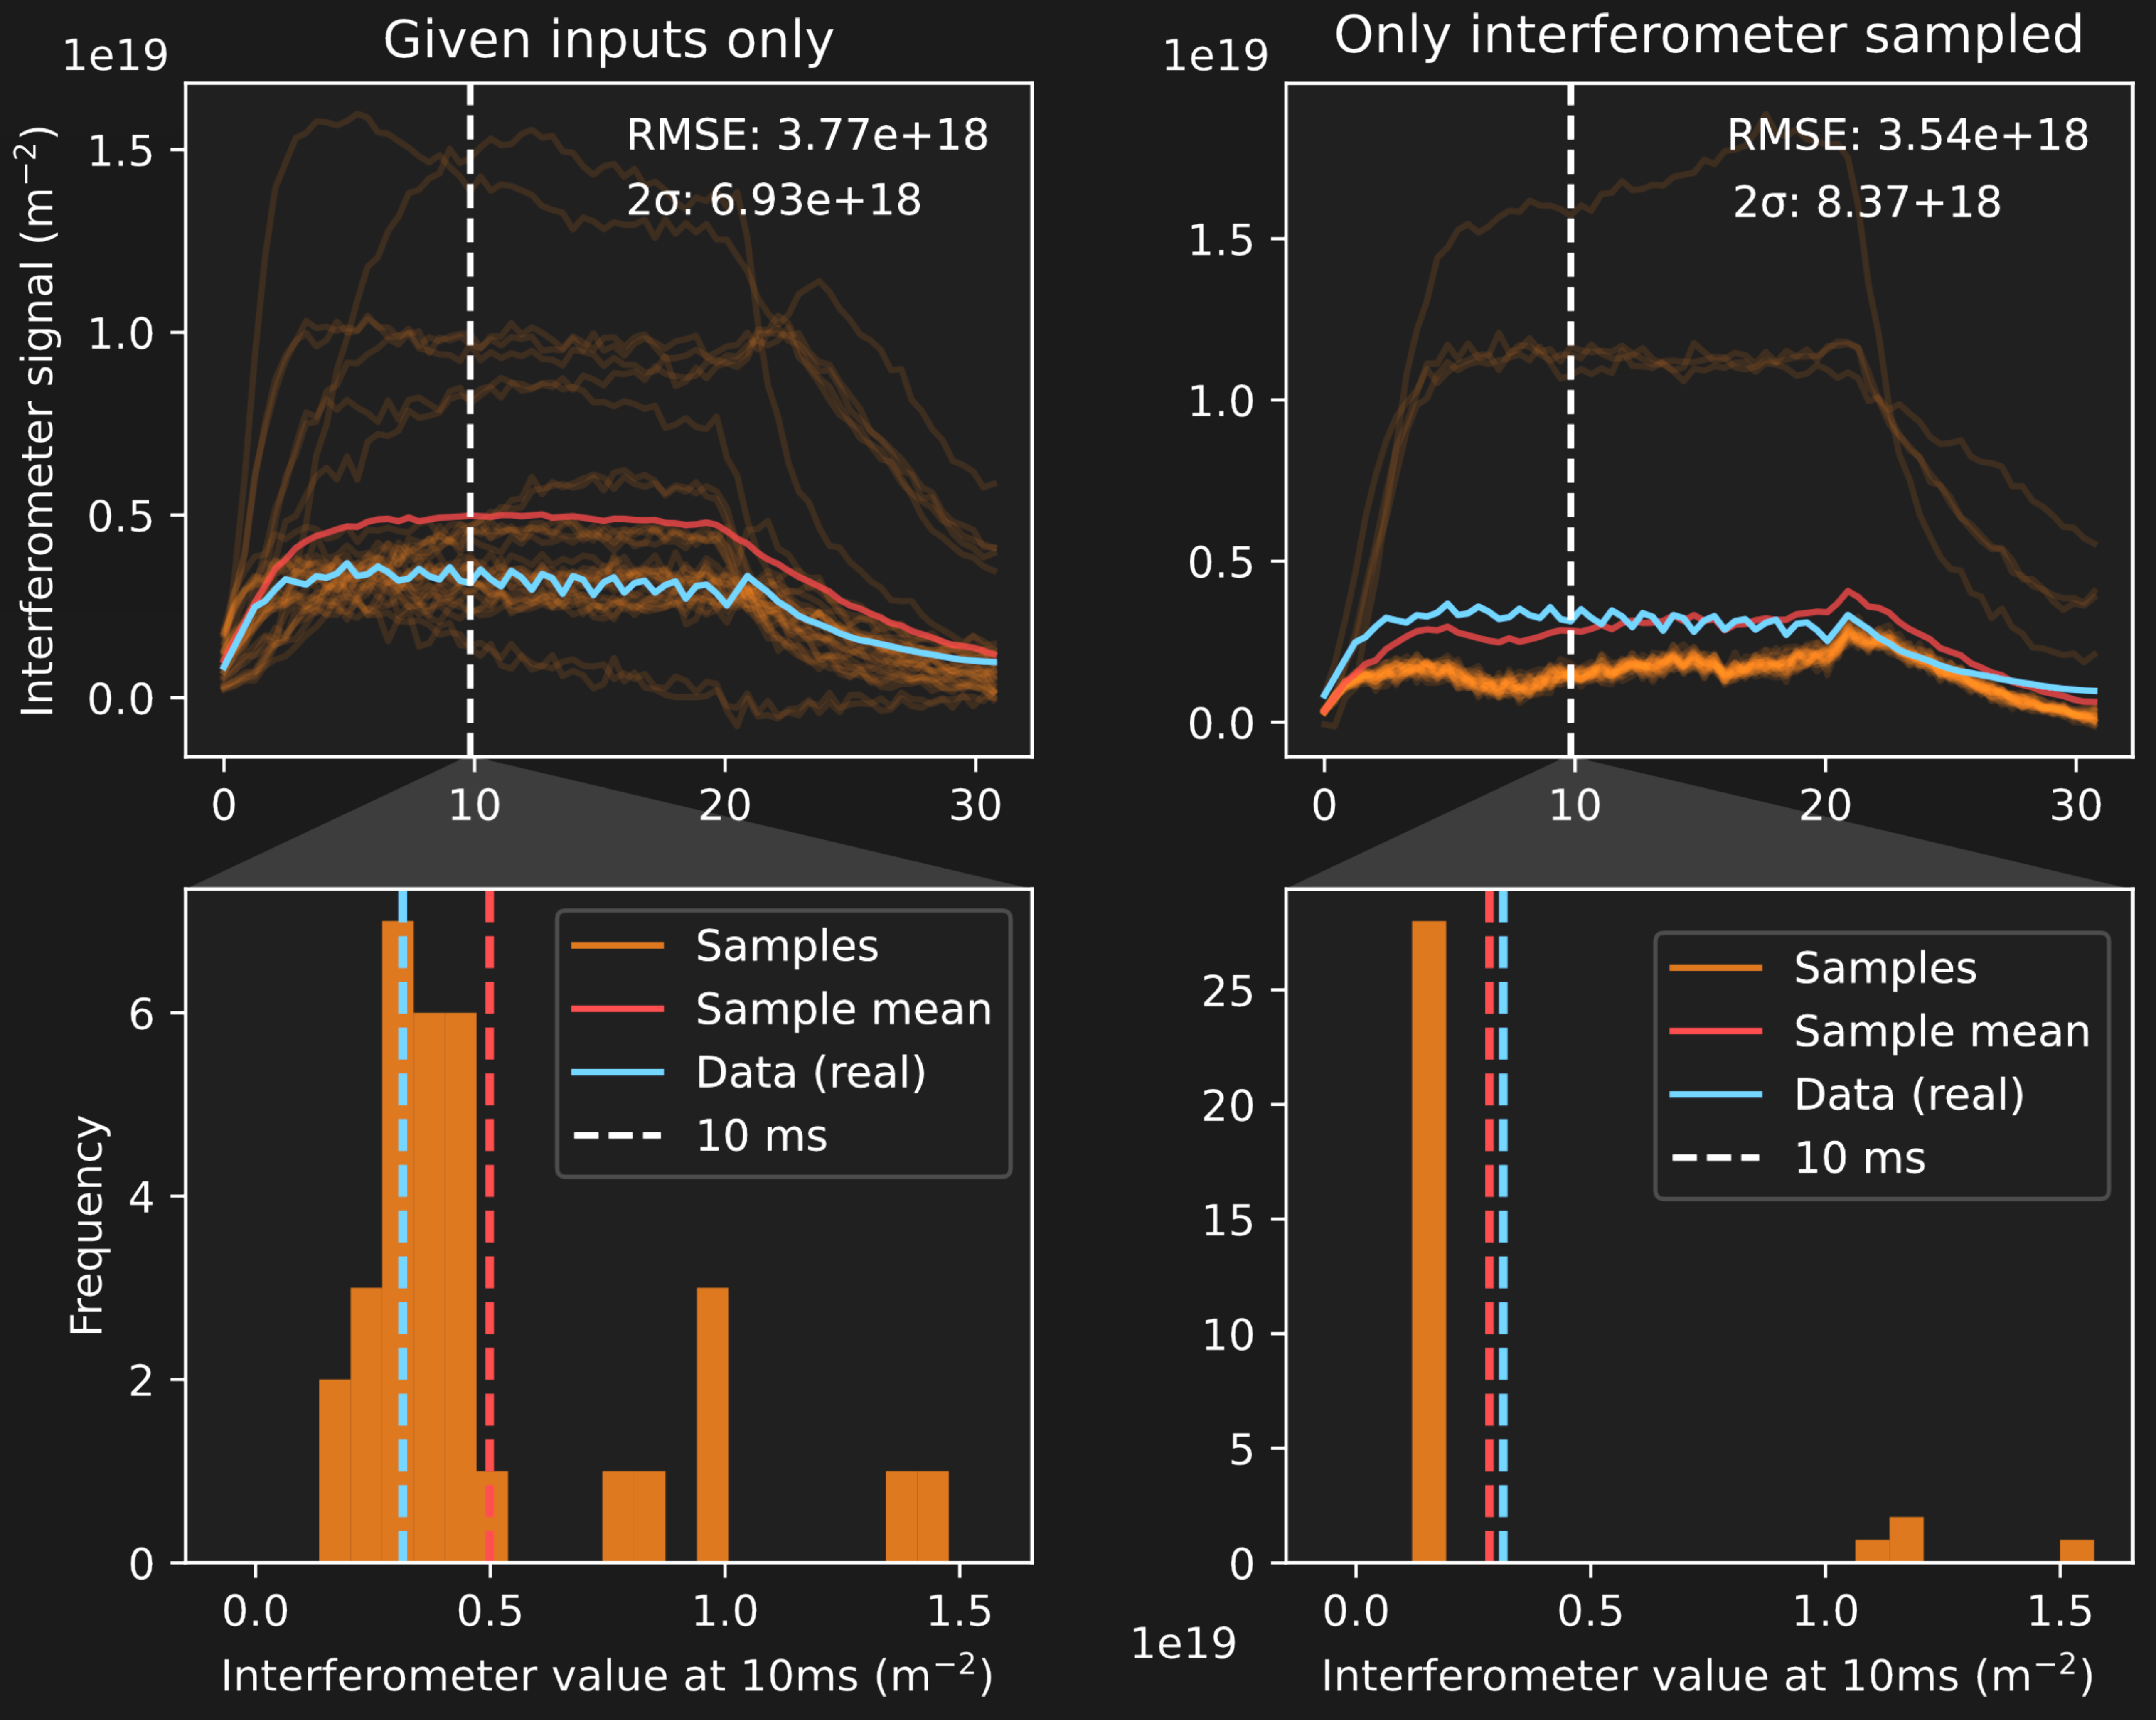

Insights may be found by directly evaluating the energy function

- Evaluating energy over probe position
- Symmetry in I_{sat} signals
- Relationships need not be invertible
- Symmetry sometimes is not observed

—> real or model issue?

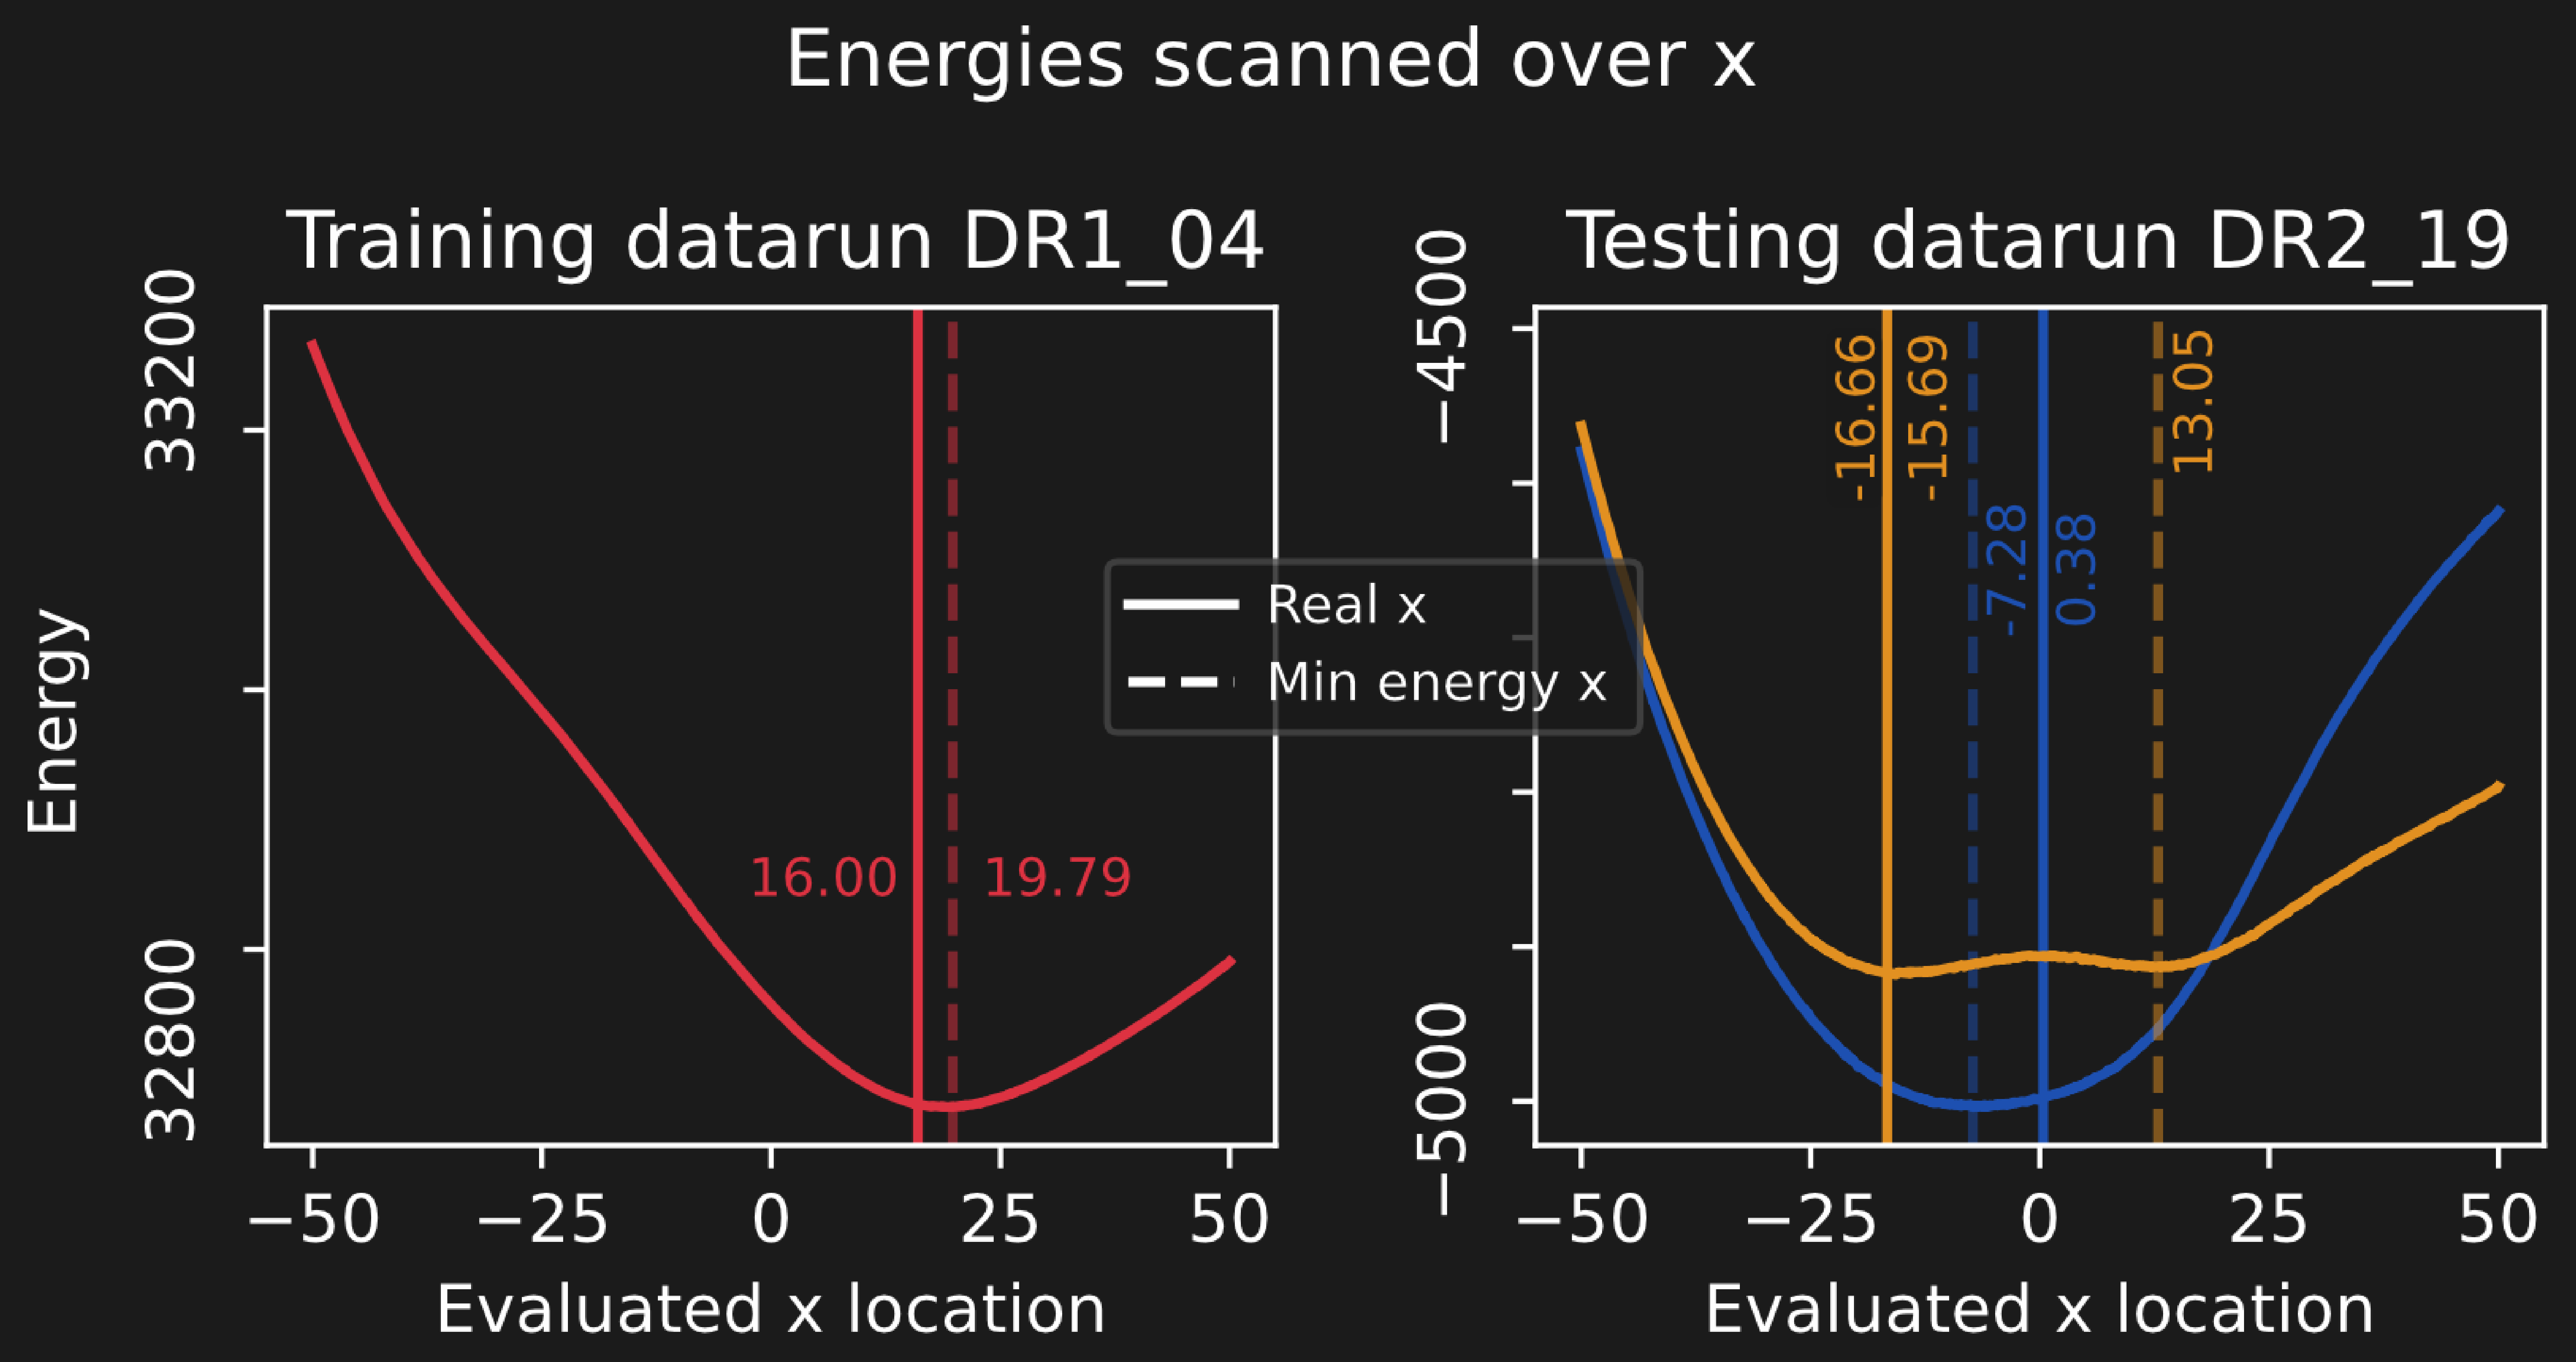

Energy-based models are an incredibly flexible way of modeling data

- Demonstrated diagnostic reconstruction with any combination of inputs
- Modified the energy function to generate good samples
- Found symmetries via direct evaluation of the energy function
- Very novel work — I’ve only found one other use of EBMs (particle physics)

- Many potential improvements:
    - more data, more diagnostics, better probe calibration (or not)
    - track cathode condition (already have a 29M+ shot dataset)
    - combine with simulations

Mirror machines and machine learning can be a faster way for fusion power

- Undertook a study of mirror turbulence, optimized the LAPD using ML, reconstructed diagnostics using EBMs

I started this PhD thinking we might be able to speed up fusion science using ML

I now see a trajectory where that's possible

- We now have a way of extracting trends and optimizing devices from data
  - May require restructuring our scientific programs
  - Can combine experiment with simulation using EBMs
- If we iterate on physics faster, we'll need to iterate our devices faster


Backup slides

Fun stats

- Data collected: 12+ TB
- Models trained: >1749
- Taxpayer dollars ~~wasted~~ utilized: ~$0.5M (thanks everyone!)
- Photos taken: 113,065 (4.9 TB)

Altair

Mirror-turb: Loss cone instabilities

- Alfvén ion cyclotron (AIC) instability: Alfvén waves coupling to the ion cyclotron motion
- Drift cyclotron loss cone (DCLC) instability:

Mirror-turb: Stabilization mechanisms for interchange

- Line-tying
- Finite Larmor radius effects
- Azimuthal flow shear
- New electrons trapped by the ambipolar potential
- We are looking at a large aspect-ratio mirror

Mirror-turb: Mirror: plasma parameters

Cathode radius (M=1)	x_c	30		cm
Machine radius	R	50		cm
Plasma length	L	~ 17		m
Primary species		He-4 1+		
Electron-helium mass ratio		1.37×10^{-4}		
Neutral pressure		$6 - 20 \times 10^{-5}$		Torr

Quantity		Core	$x = x_{\mathrm{PF}}$	Unit
Density	n_e	1.25×10^{12}	0.6×10^{12}	cm^{-3}
Ion temperature	T_i	~ 1		eV
Electron temperature	T_e	4	5	eV
Beta (total)	β	9×10^{-4}	6×10^{-4}	
Midplane magnetic field	B_{mid}	500		G
Plasma freq	Ω_{pe}	10	7.1	GHz
Ion cyclotron freq	Ω_{ci}	200		kHz
Electron cyclotron freq	Ω_{ce}	1.4		GHz
Debye length	λ_D	0.013	0.021	mm
Electron skin depth	λ_e	30	43	mm
Ion gyroradius	λ_{ci}	5.8		mm
Electron gyroradius	λ_{ce}	0.13	0.15	mm
Ion thermal velocity	$\bar{v}_i$	6.94		km/s
Electron thermal velocity	$\bar{v}_e$	1190	1330	km/s
Sound speed	c_s	13.0	13.9	km/s
Alfvén speed	v_a	$446 - 1140$	-1620	km/s
Ion sound radius	ρ_s	65	69	mm
Ion-ion collision freq	ν_{ii}	730	380	kHz
Electron-ion collision freq	ν_{ei}	6.77	2.59	MHz
Electron collision freq	ν_{ee}	9.57	3.66	MHz
Ion mean free path	$\lambda_{i,\mathrm{mfp}}$	26	50	mm
Electron mean free path	$\lambda_{e,\mathrm{mfp}}$	175	512	mm
Spitzer resitivity	η	192	146	$\mu\Omega\,\mathrm{m}$

Mirror-turb: magnetic fluctuation breakdown

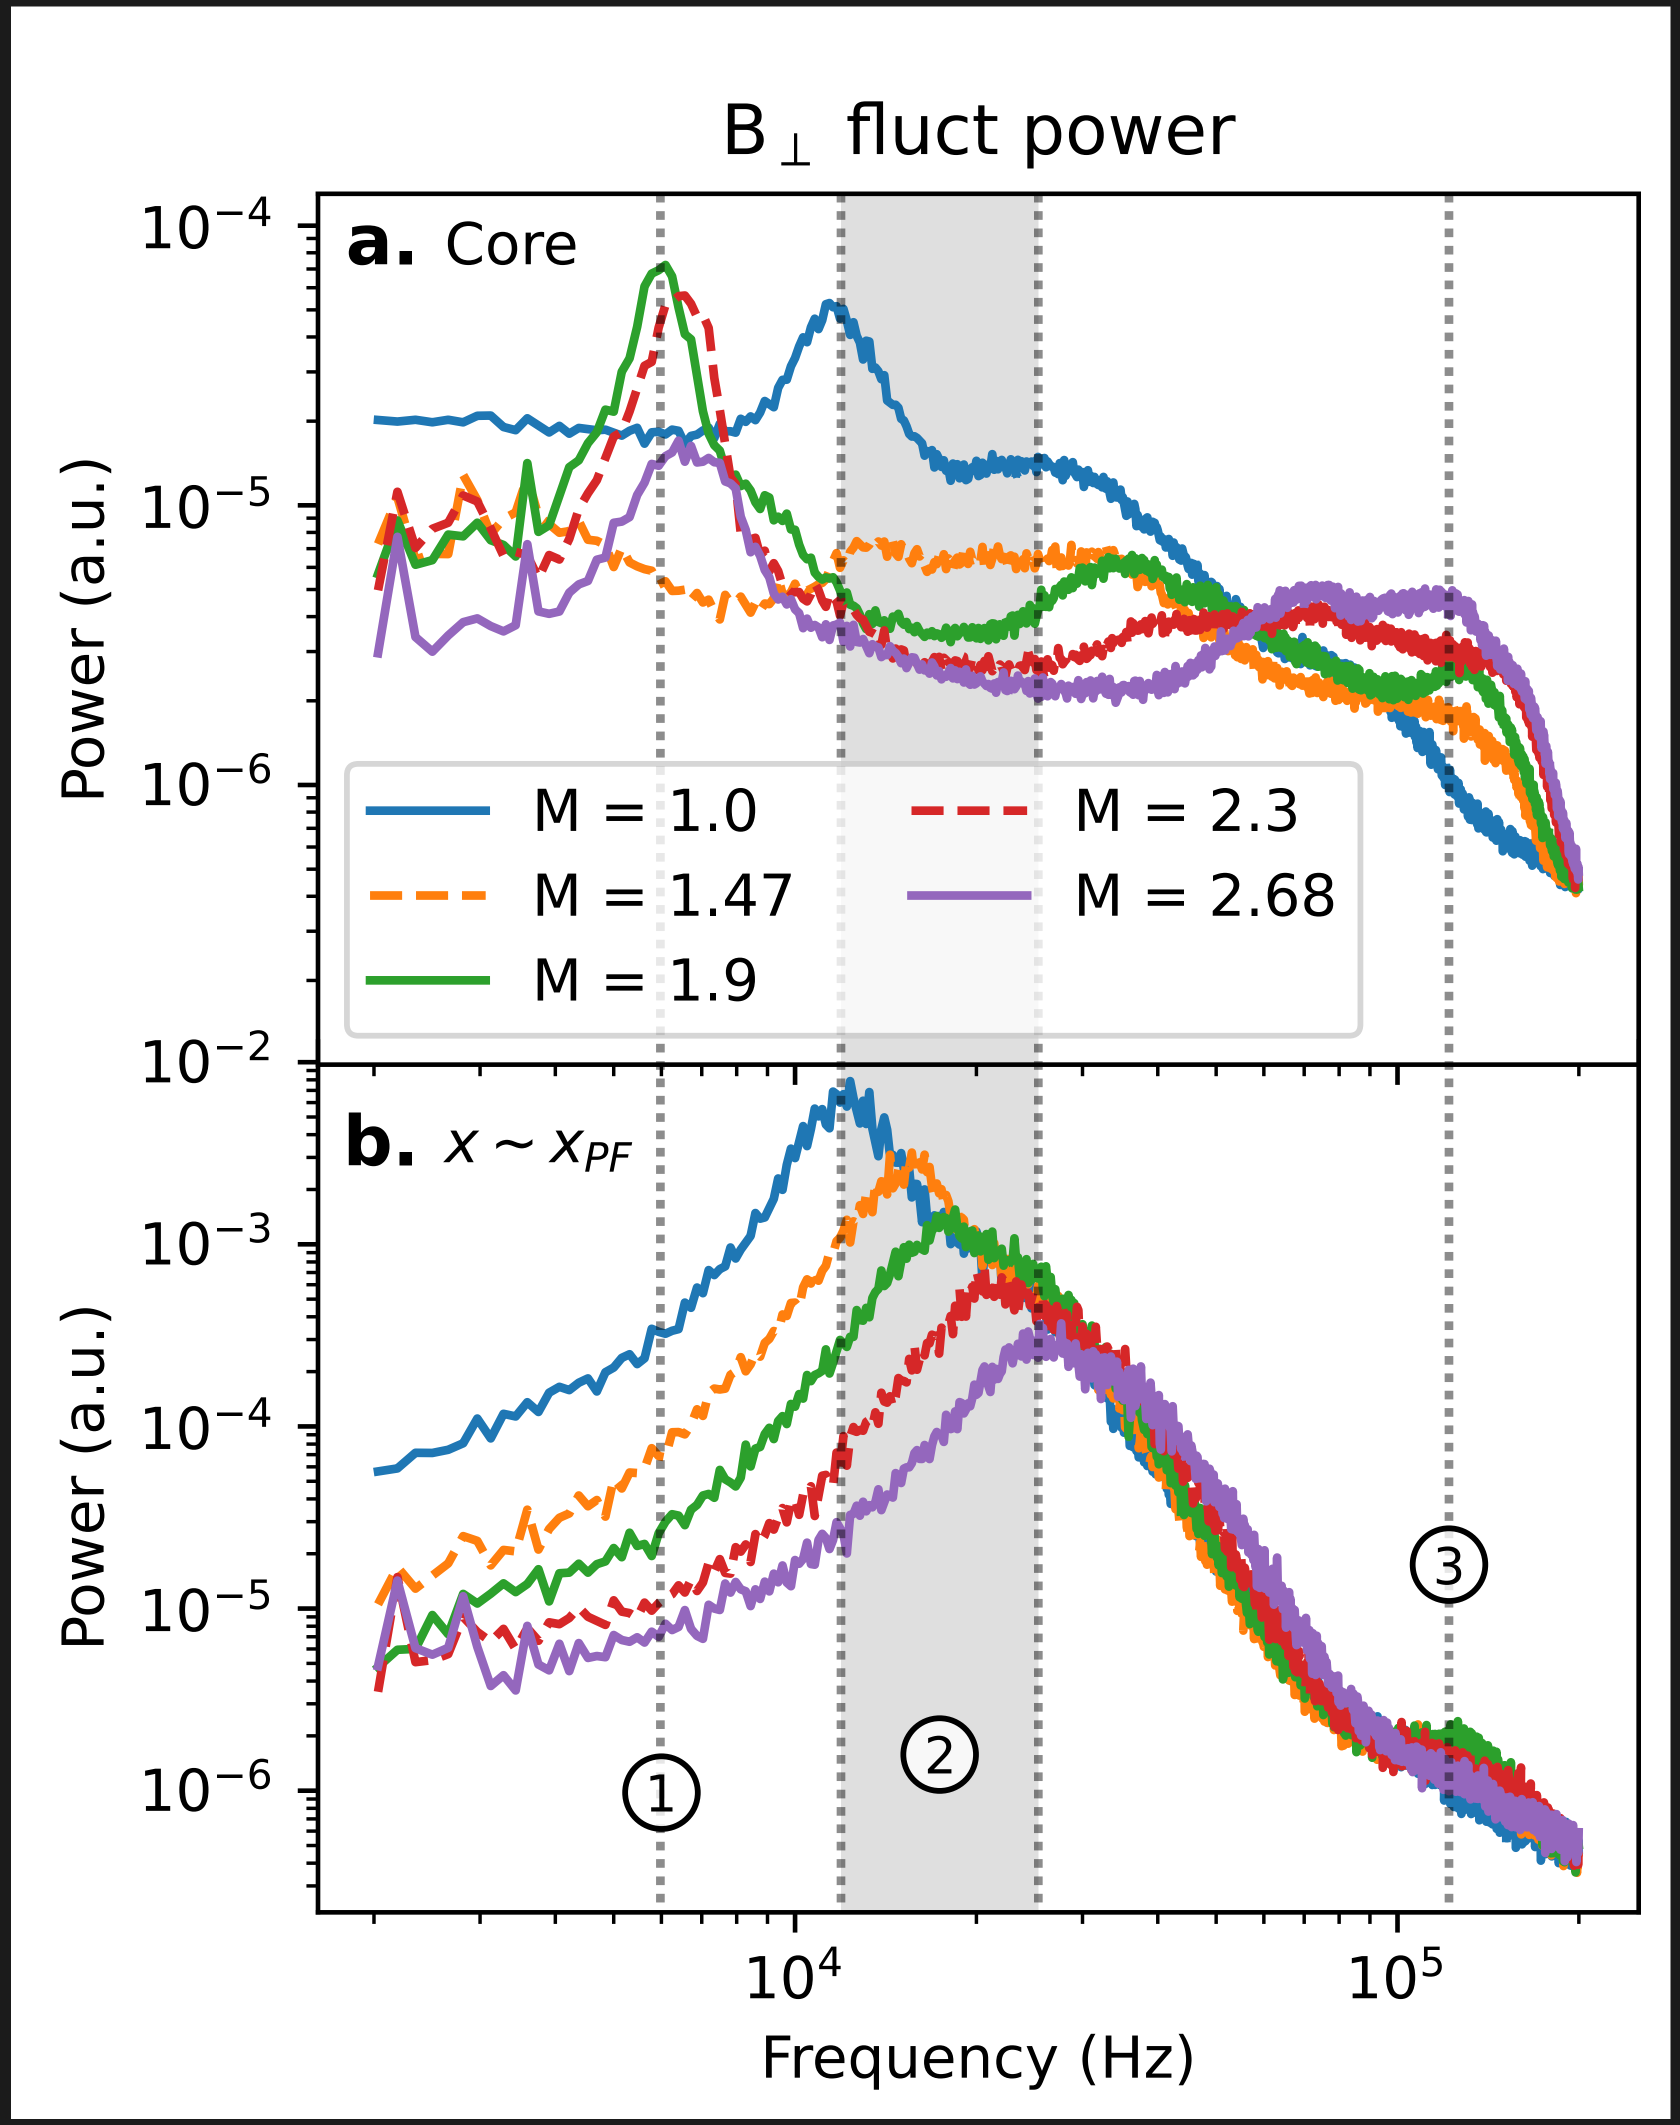

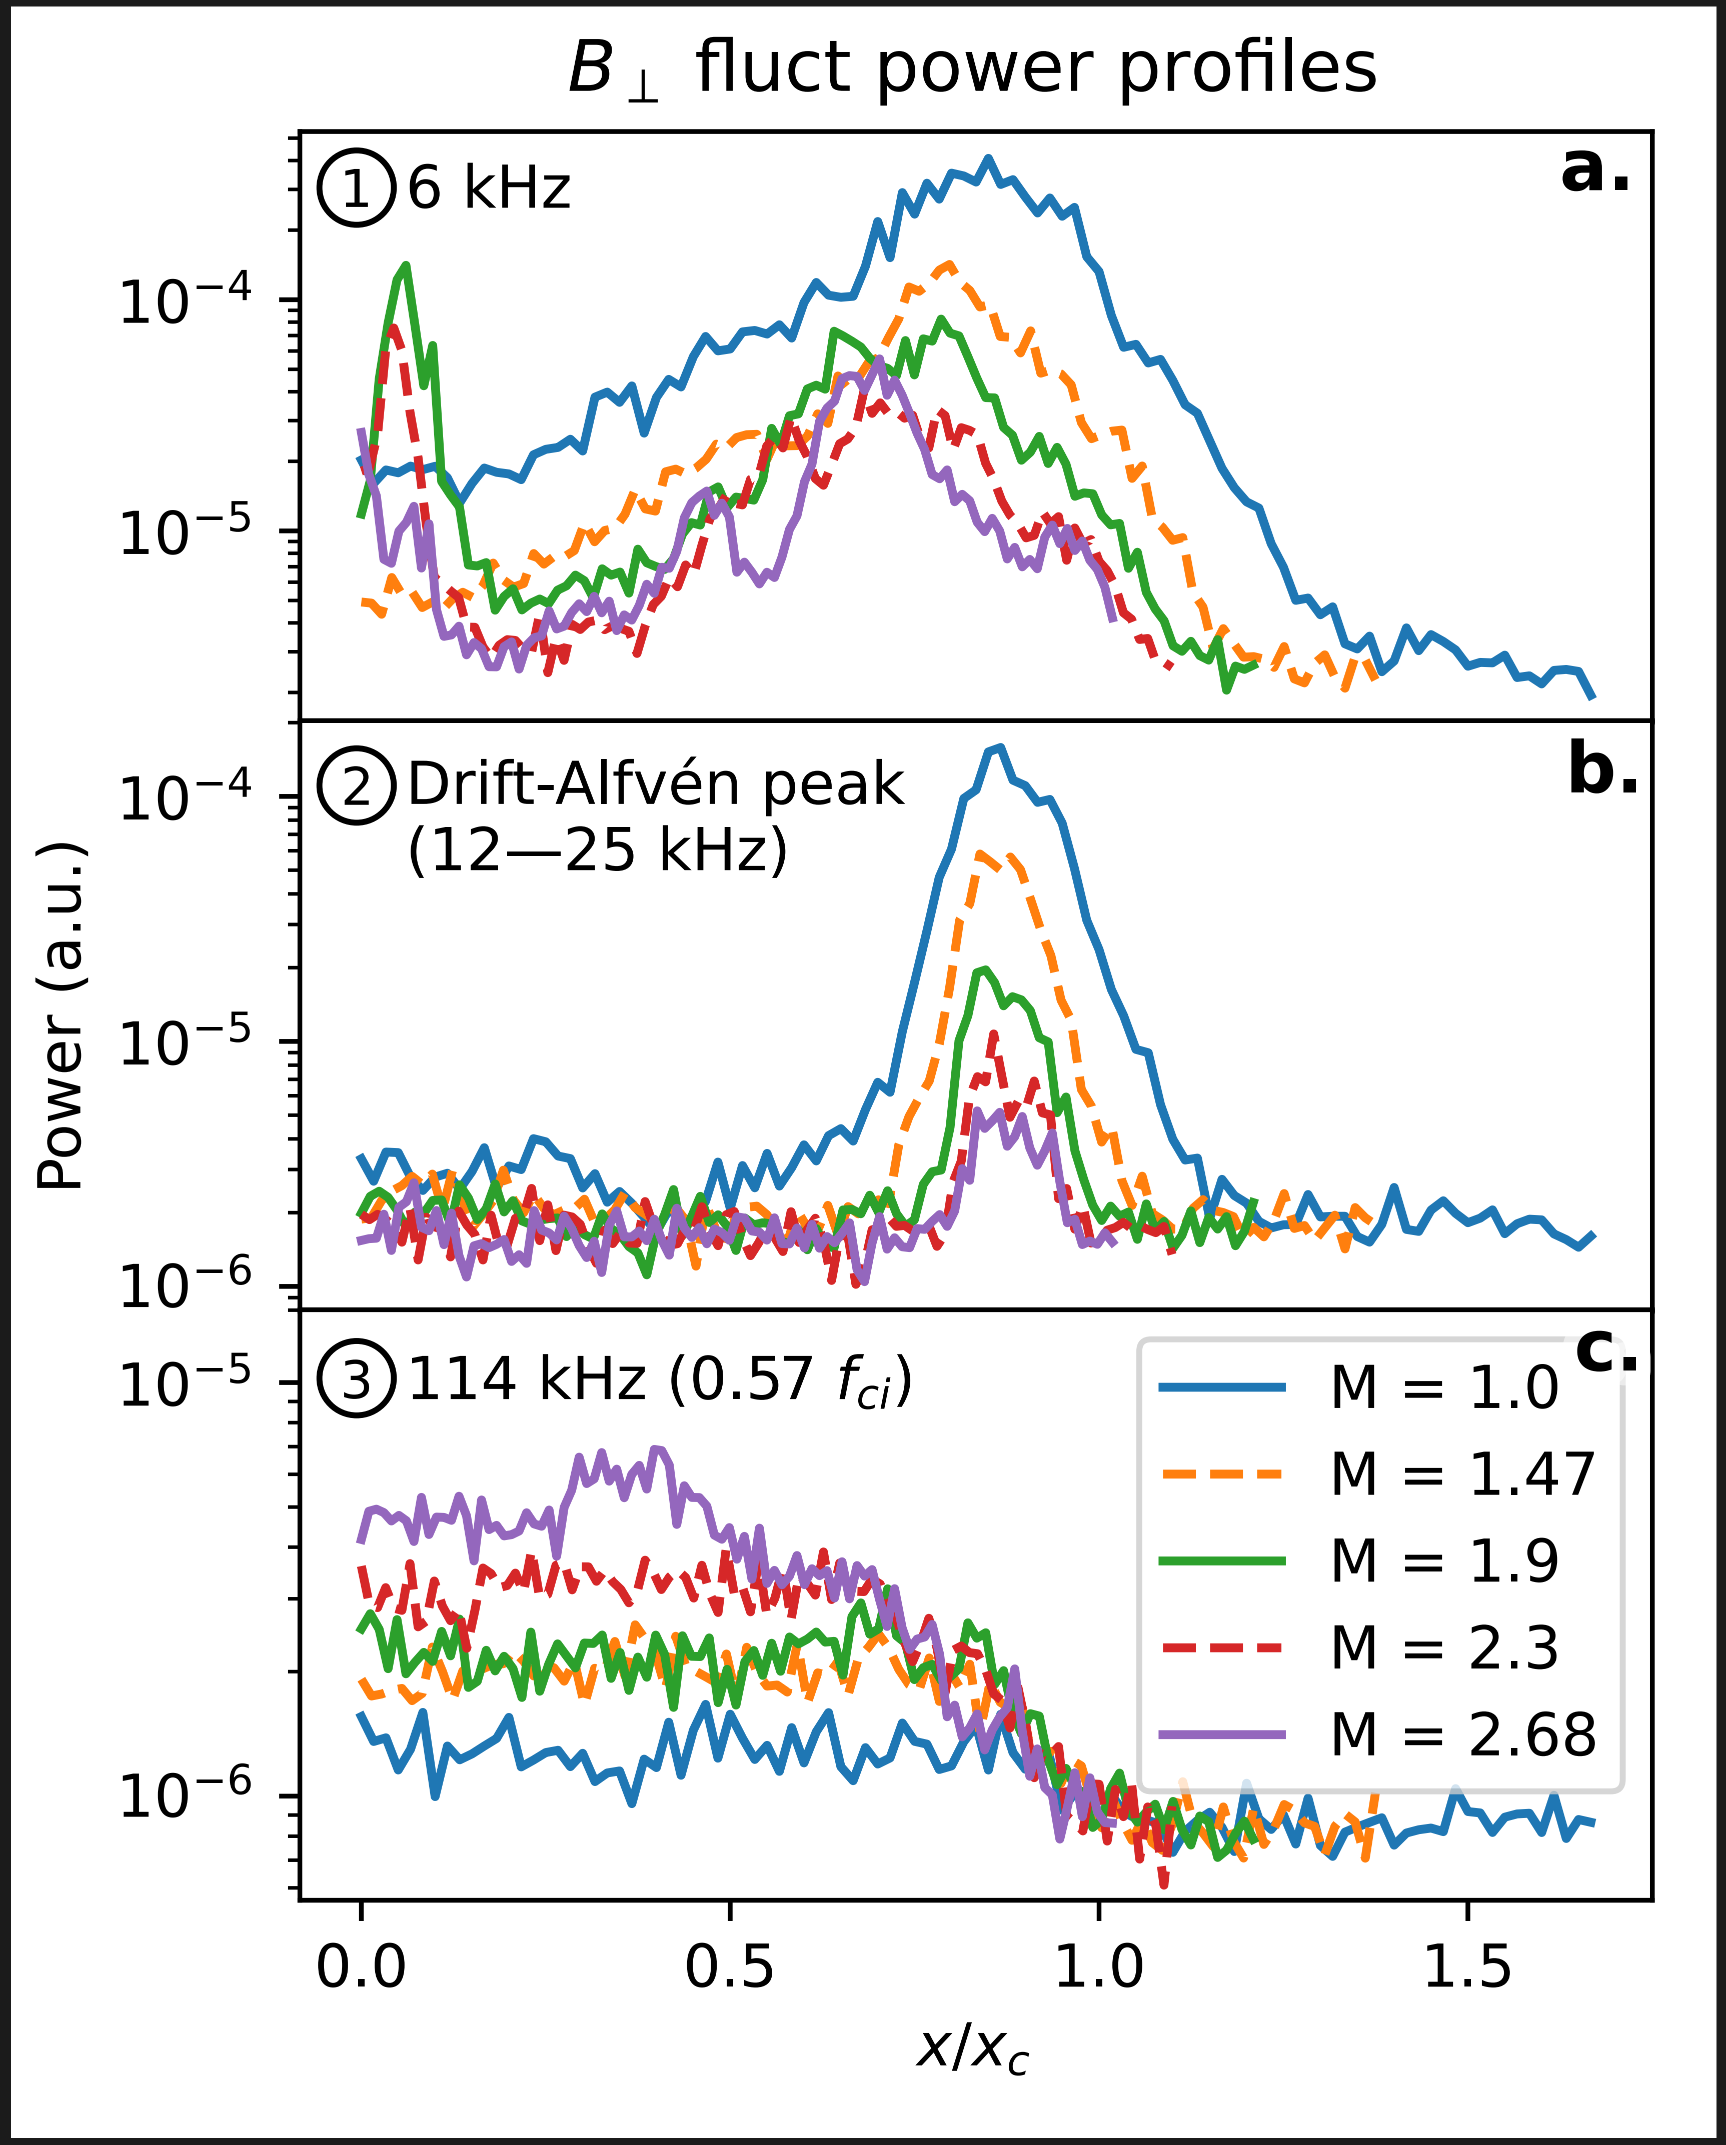

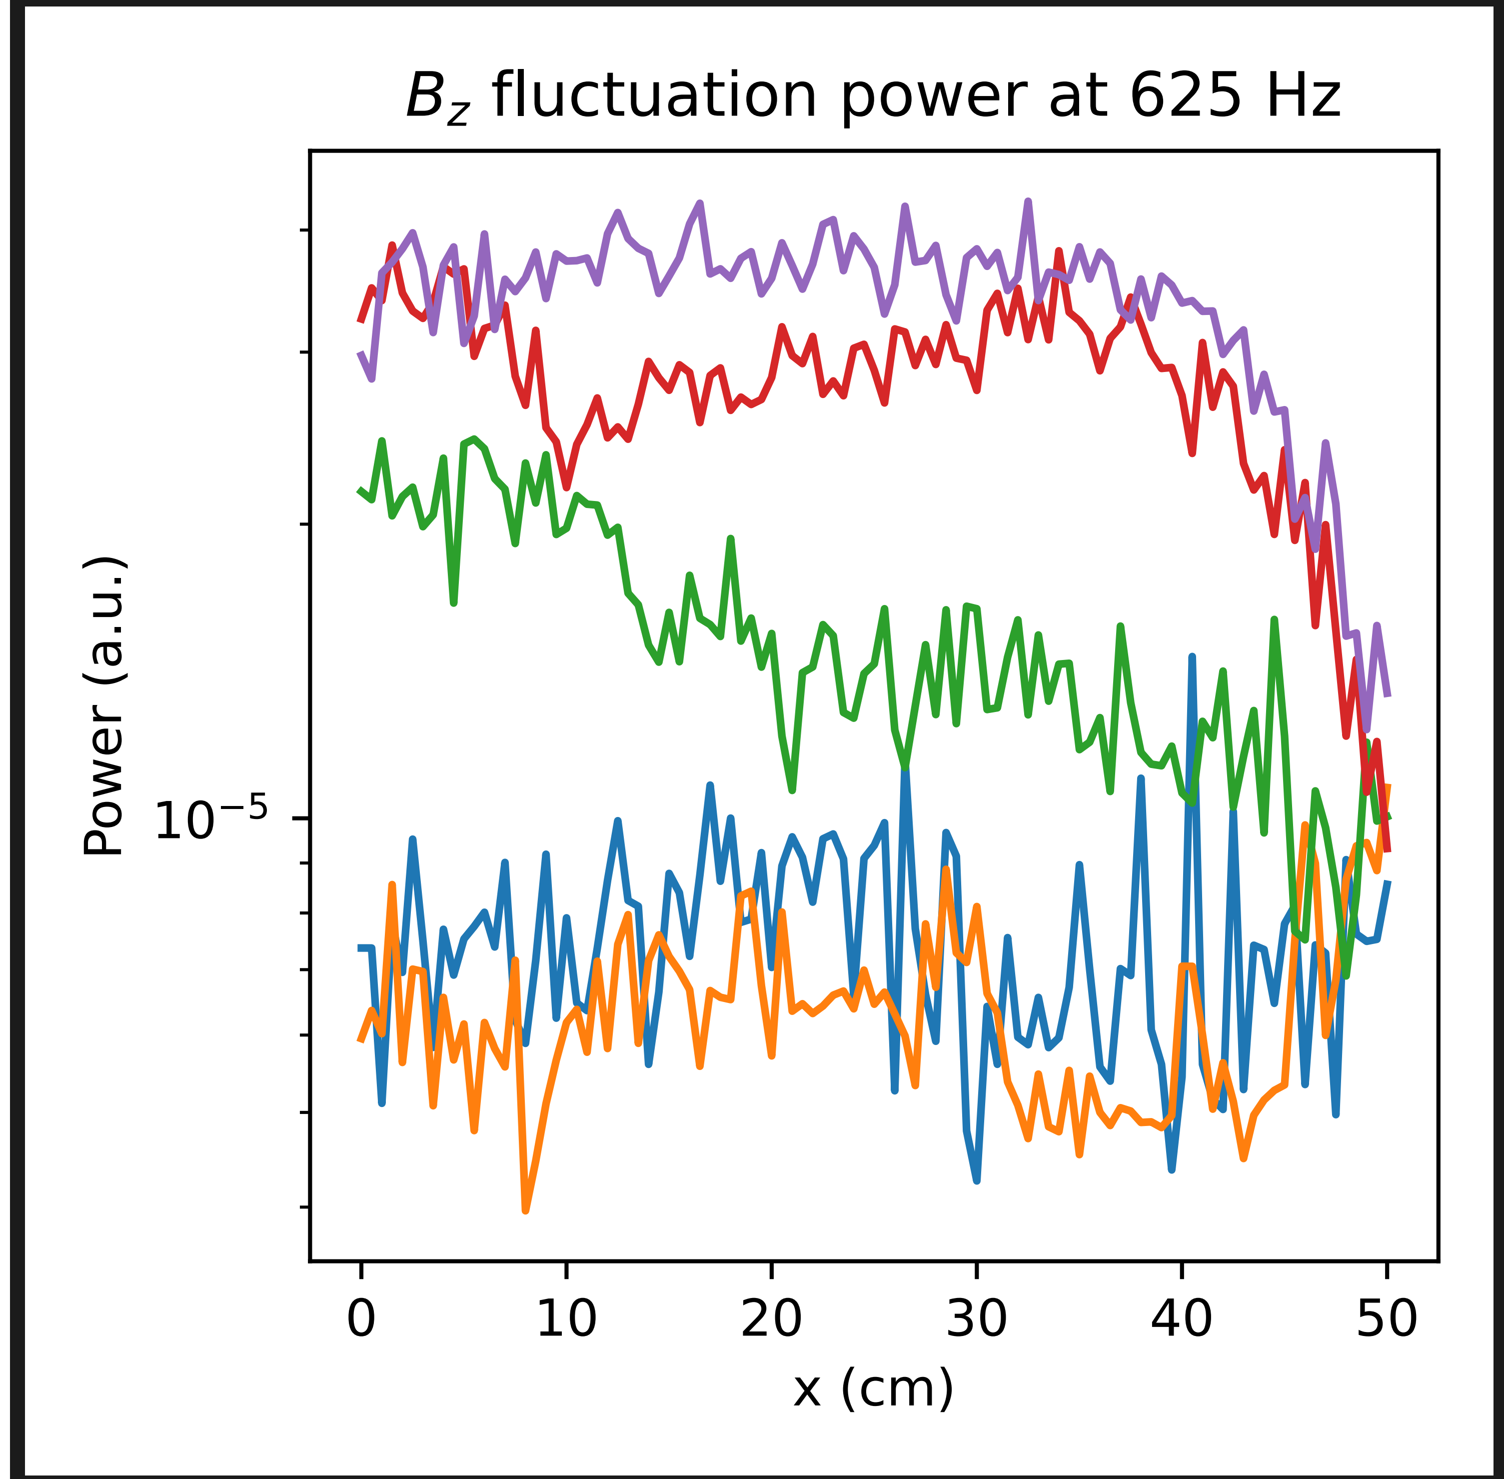

Mirror-turb: evidence for interchange

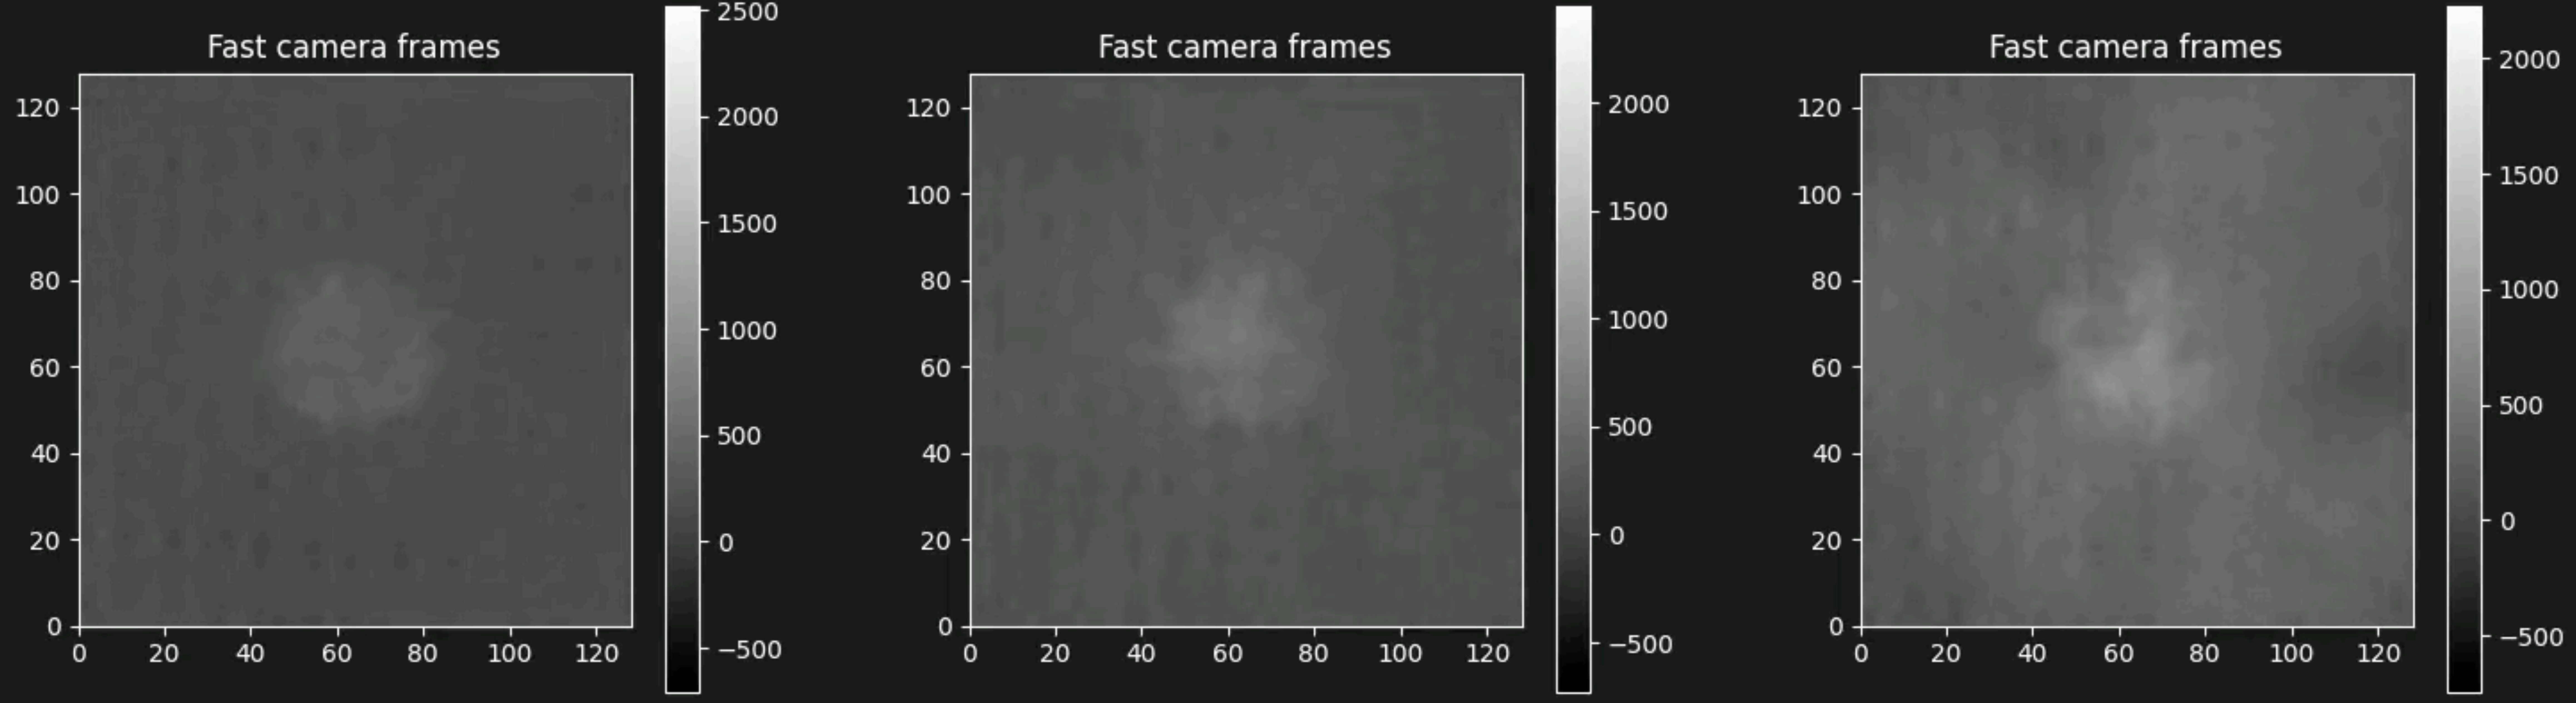

M=1 M=2.9 M=5.7

FFC, mean frame subtracted. Sample rate = 2.5 kHz

- See peak in Langmuir probe fluctuation spectra at ~2 kHz
- Temperatures get very high (> 20 eV) with short gas puff timings
  - Largely collisionless on the length scale of the mirror cell — this all points to interchange

Mirror-turb: evidence for interchange

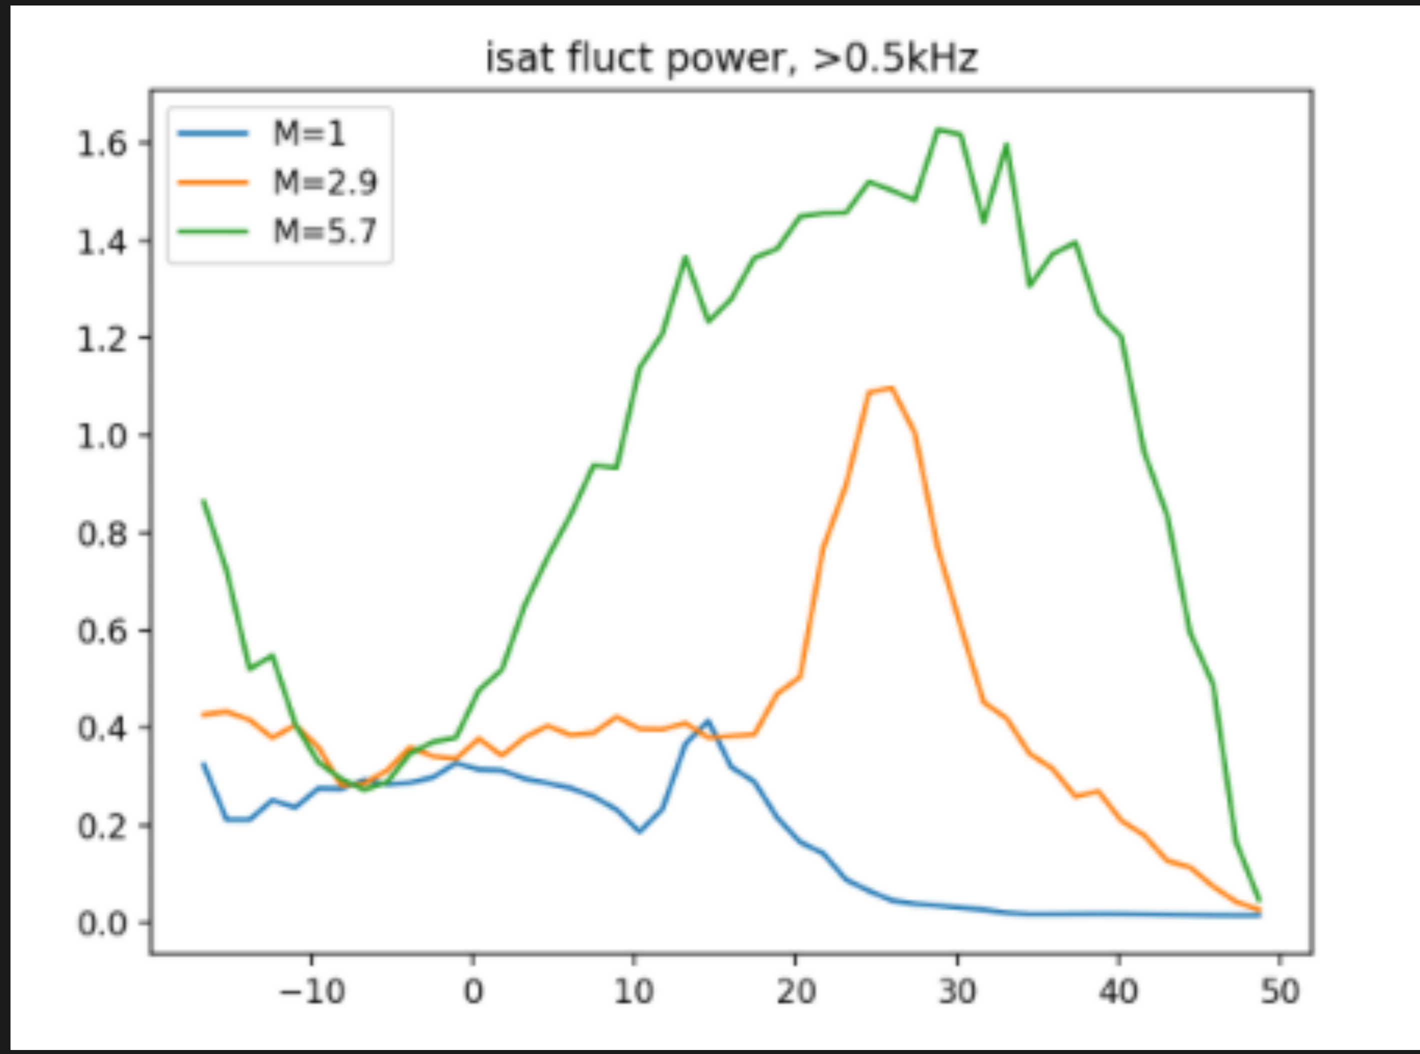

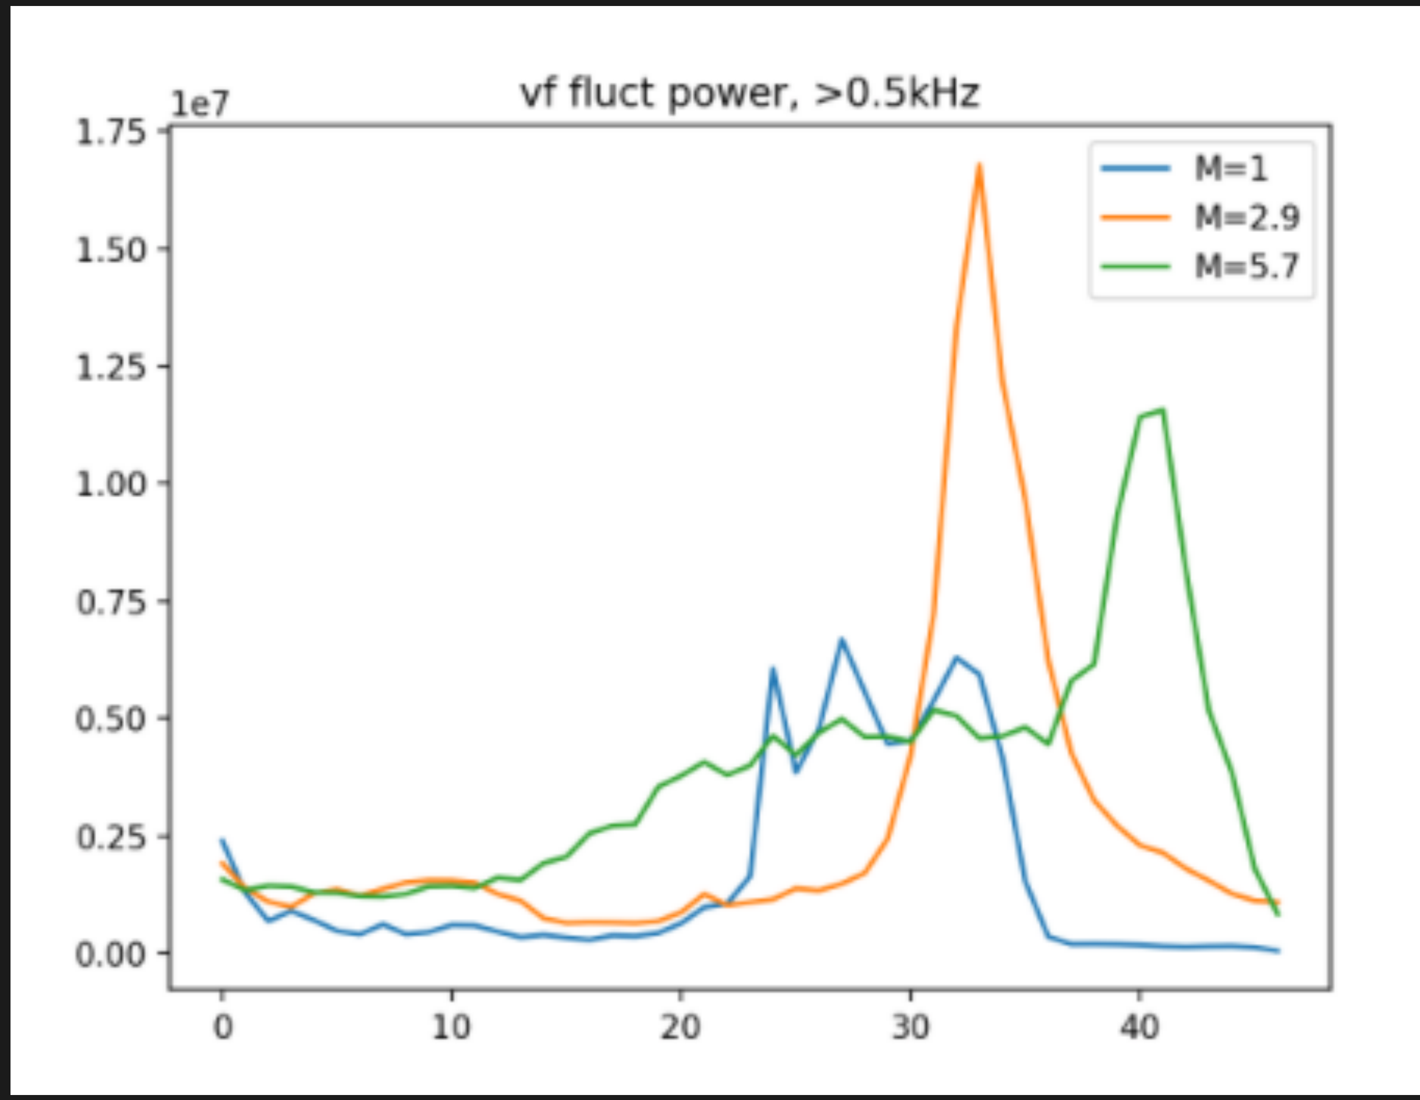

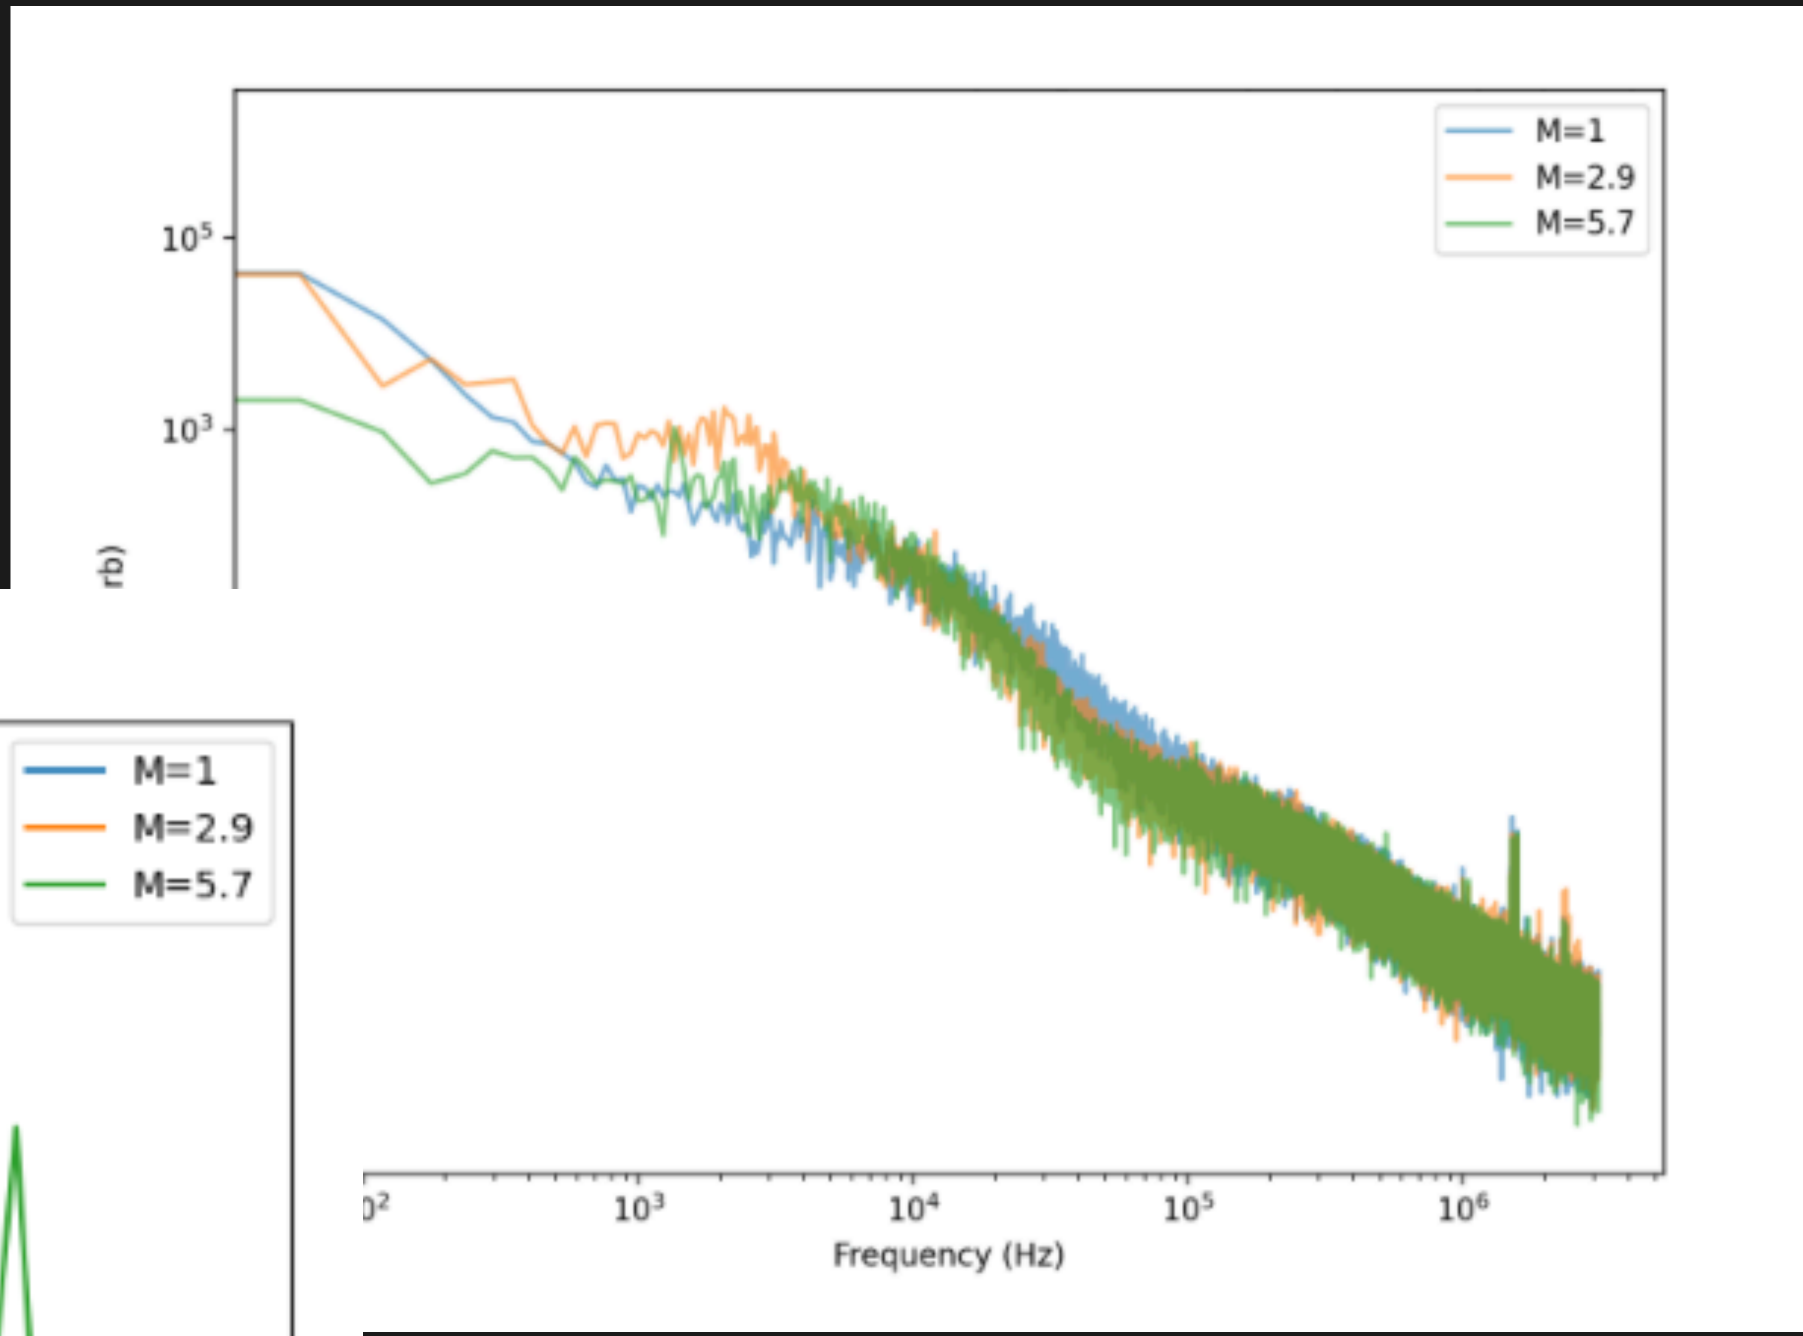

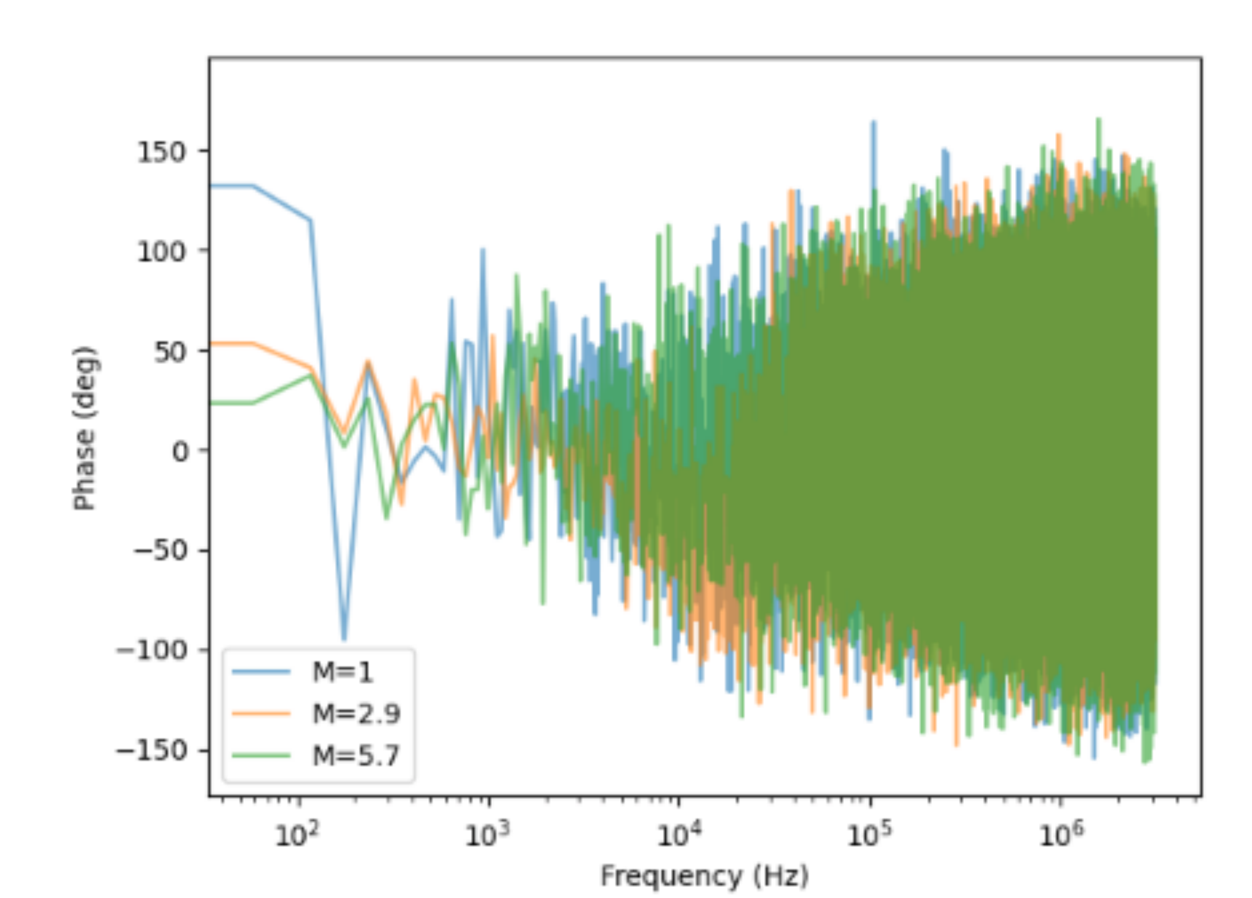

Mirror-turb: data processing

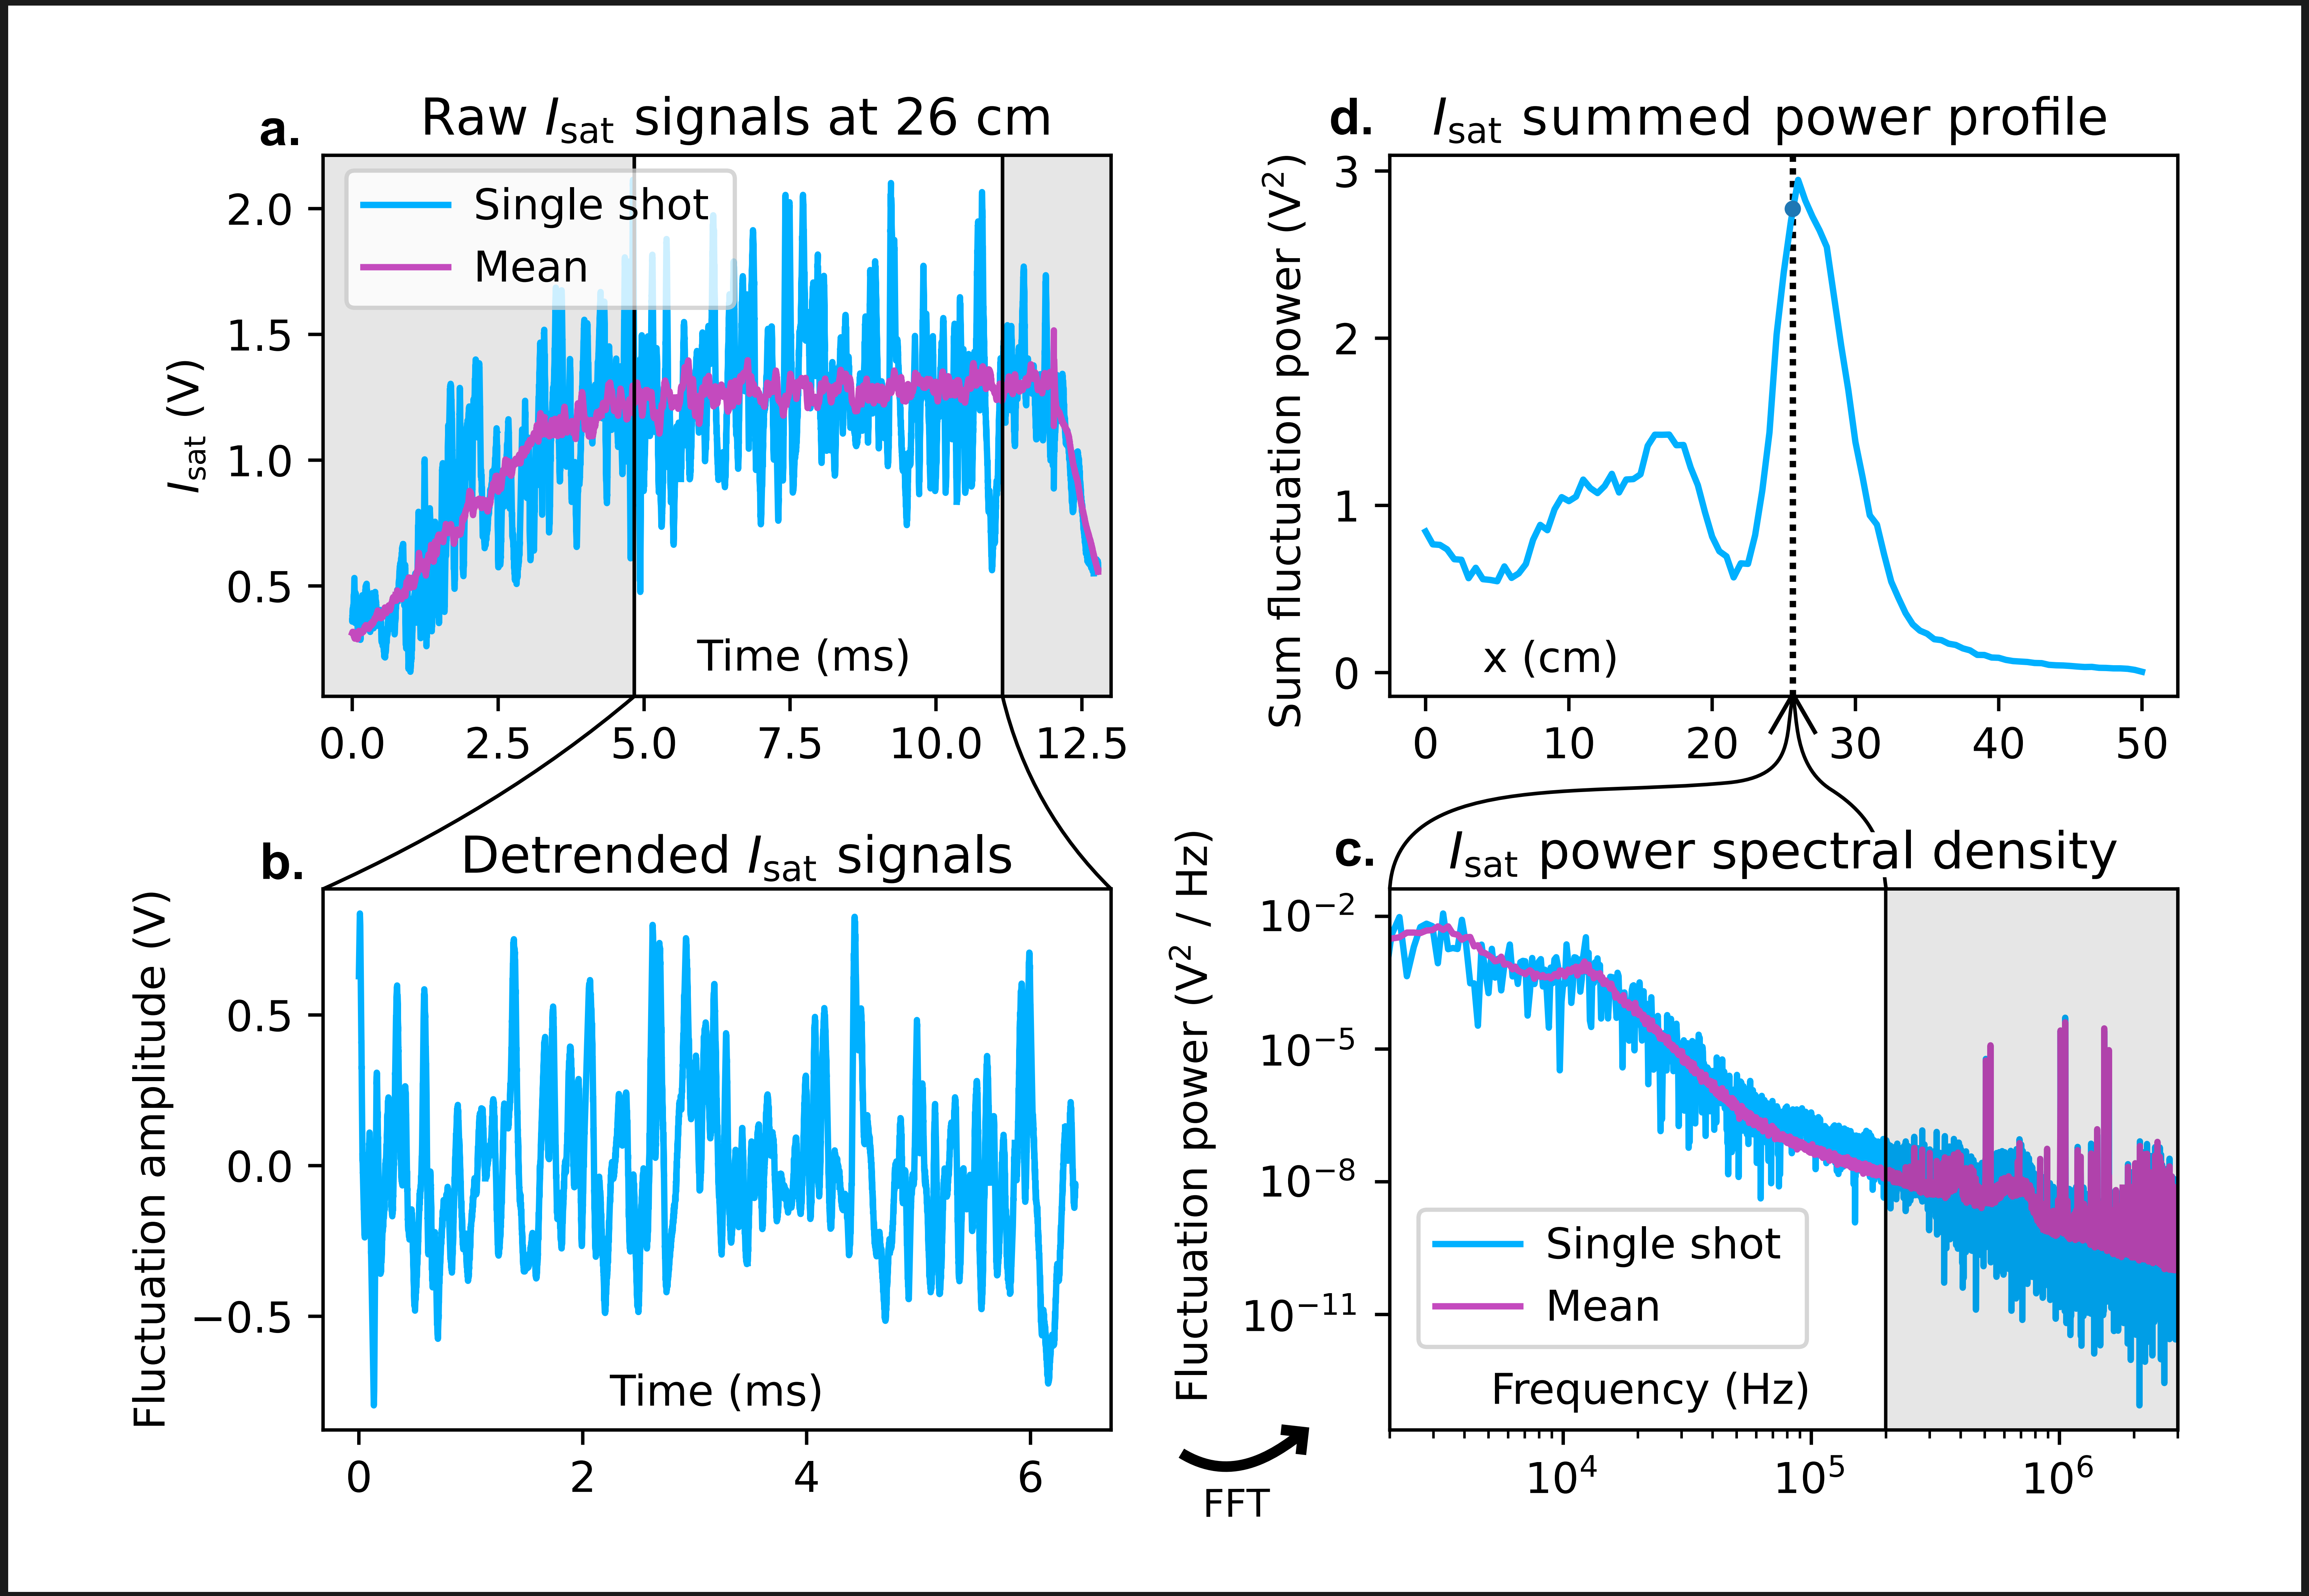

Mirror-turb: particle flux breakdown

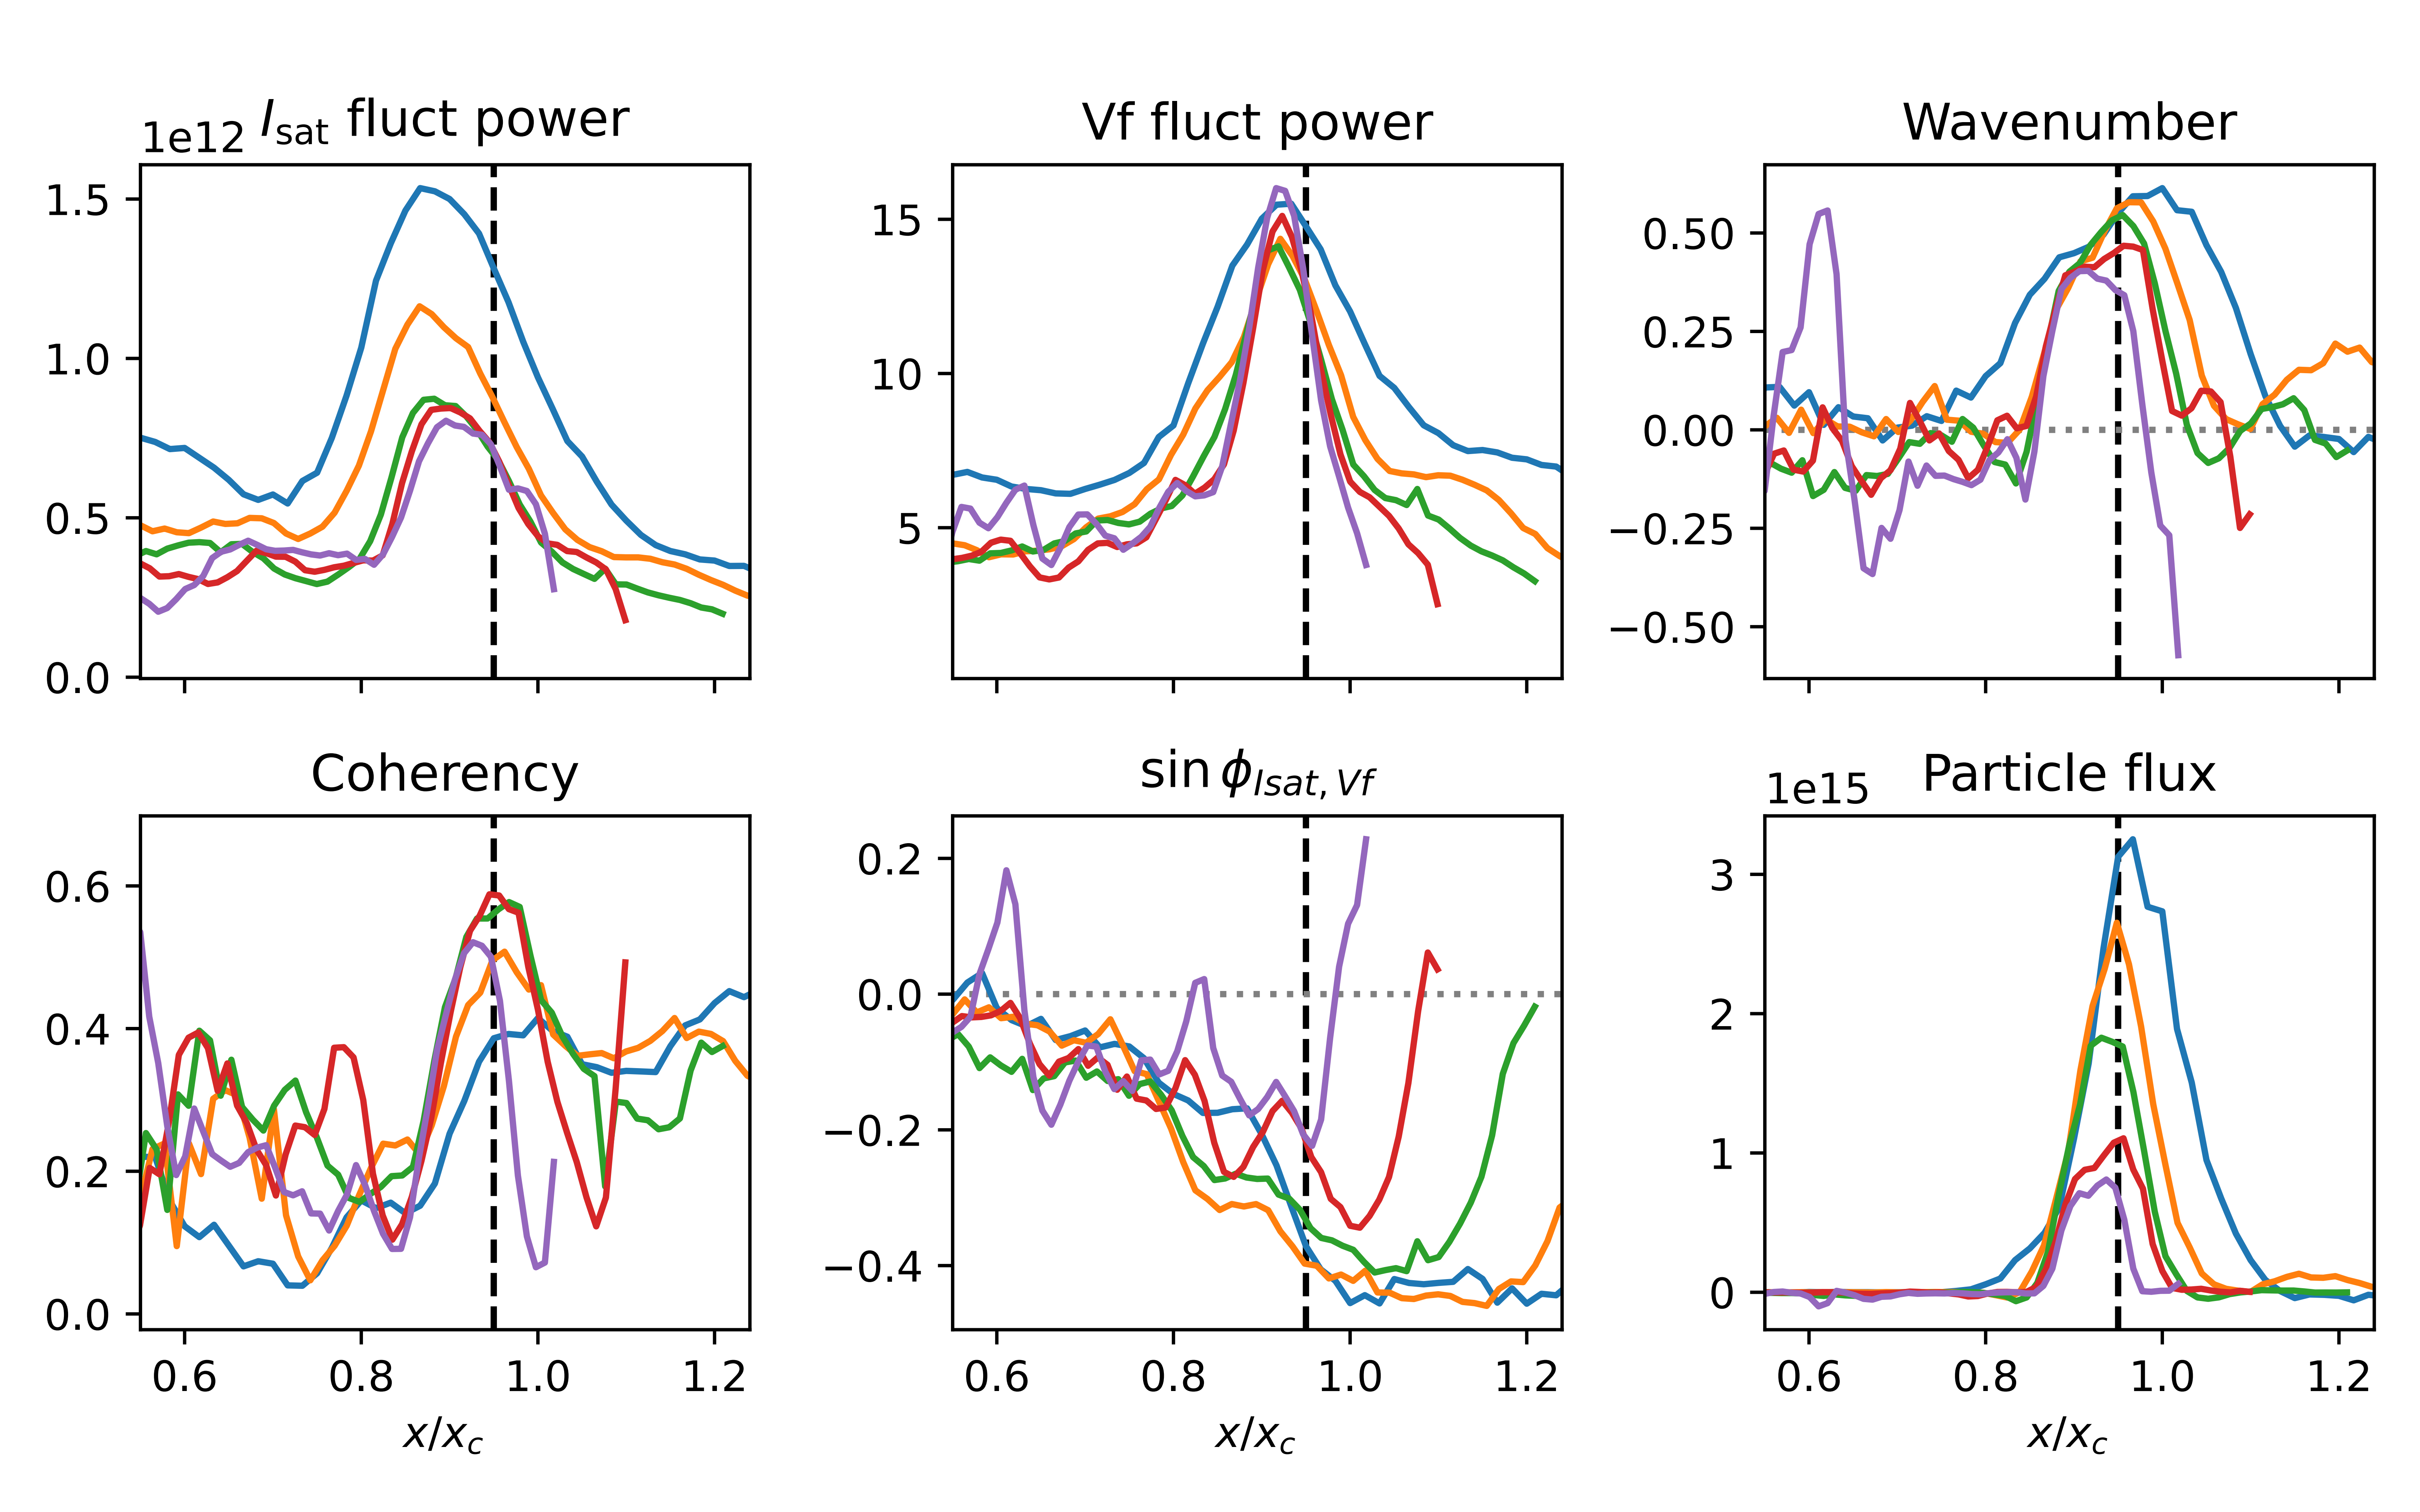

ML: data breakdown

Table 4.1: Data breakdown by class and dataset (percent)

B source (G)				B mirror (G)				B midplane (G)			
	Train	Test	All		Train	Test	All		Train	Test	All
500	4.77	0	4.29	250	4.30	8.41	4.72	250	8.25	21.01	9.55
750	3.34	12.61	4.29	500	30.49	8.41	28.23	500	43.80	8.41	40.19
1000	43.13	78.99	46.78	750	6.68	16.81	7.72	750	6.62	52.19	11.27
1250	12.59	0	11.30	1000	28.85	57.97	31.82	1000	26.36	5.78	24.26
1500	19.23	0	17.27	1250	3.34	4.20	3.43	1250	9.24	0	8.30
1750	1.91	0	1.71	1500	26.34	4.20	24.08	1500	5.73	12.61	6.43
2000	15.03	8.41	14.35								

Gas puff voltage (V)				Discharge voltage (V)				Axial probe position (cm)			
70	12.11	16.81	12.59	70	12.22	8.41	11.83	639	12.48	8.41	12.06
75	6.68	0	6.00	80	5.25	0	4.72	828	17.07	36.28	19.03
80	11.46	8.41	11.15	90	2.86	8.41	3.43	859	12.48	8.41	12.06
82	41.49	57.97	43.17	100	3.34	8.41	3.86	895	33.01	30.10	32.71
85	14.13	0	12.69	110	8.77	0	7.87	1017	12.48	8.41	12.06
90	14.13	16.81	14.40	112	20.62	0	18.52	1145	12.48	8.41	12.06
				120	3.82	8.41	4.29				
				130	0.95	0	0.86				
				140	2.86	8.41	3.43				
				150	39.30	57.97	41.20				

Gas puff duration (ms)				Vertical probe position (cm)			
38	94.27	91.59	94.00	≈ 0	36.26	46.08	37.26
< 38	5.73	8.41	6.00	$\neq 0$	63.74	53.92	62.74

ML: data info

- Collected across two run weeks about a year apart, DR1 and DR2
- 131k shots collected over 67 dataruns
- DR1 had much higher neutral pressures than DR2
- Model will perform better where there's more data

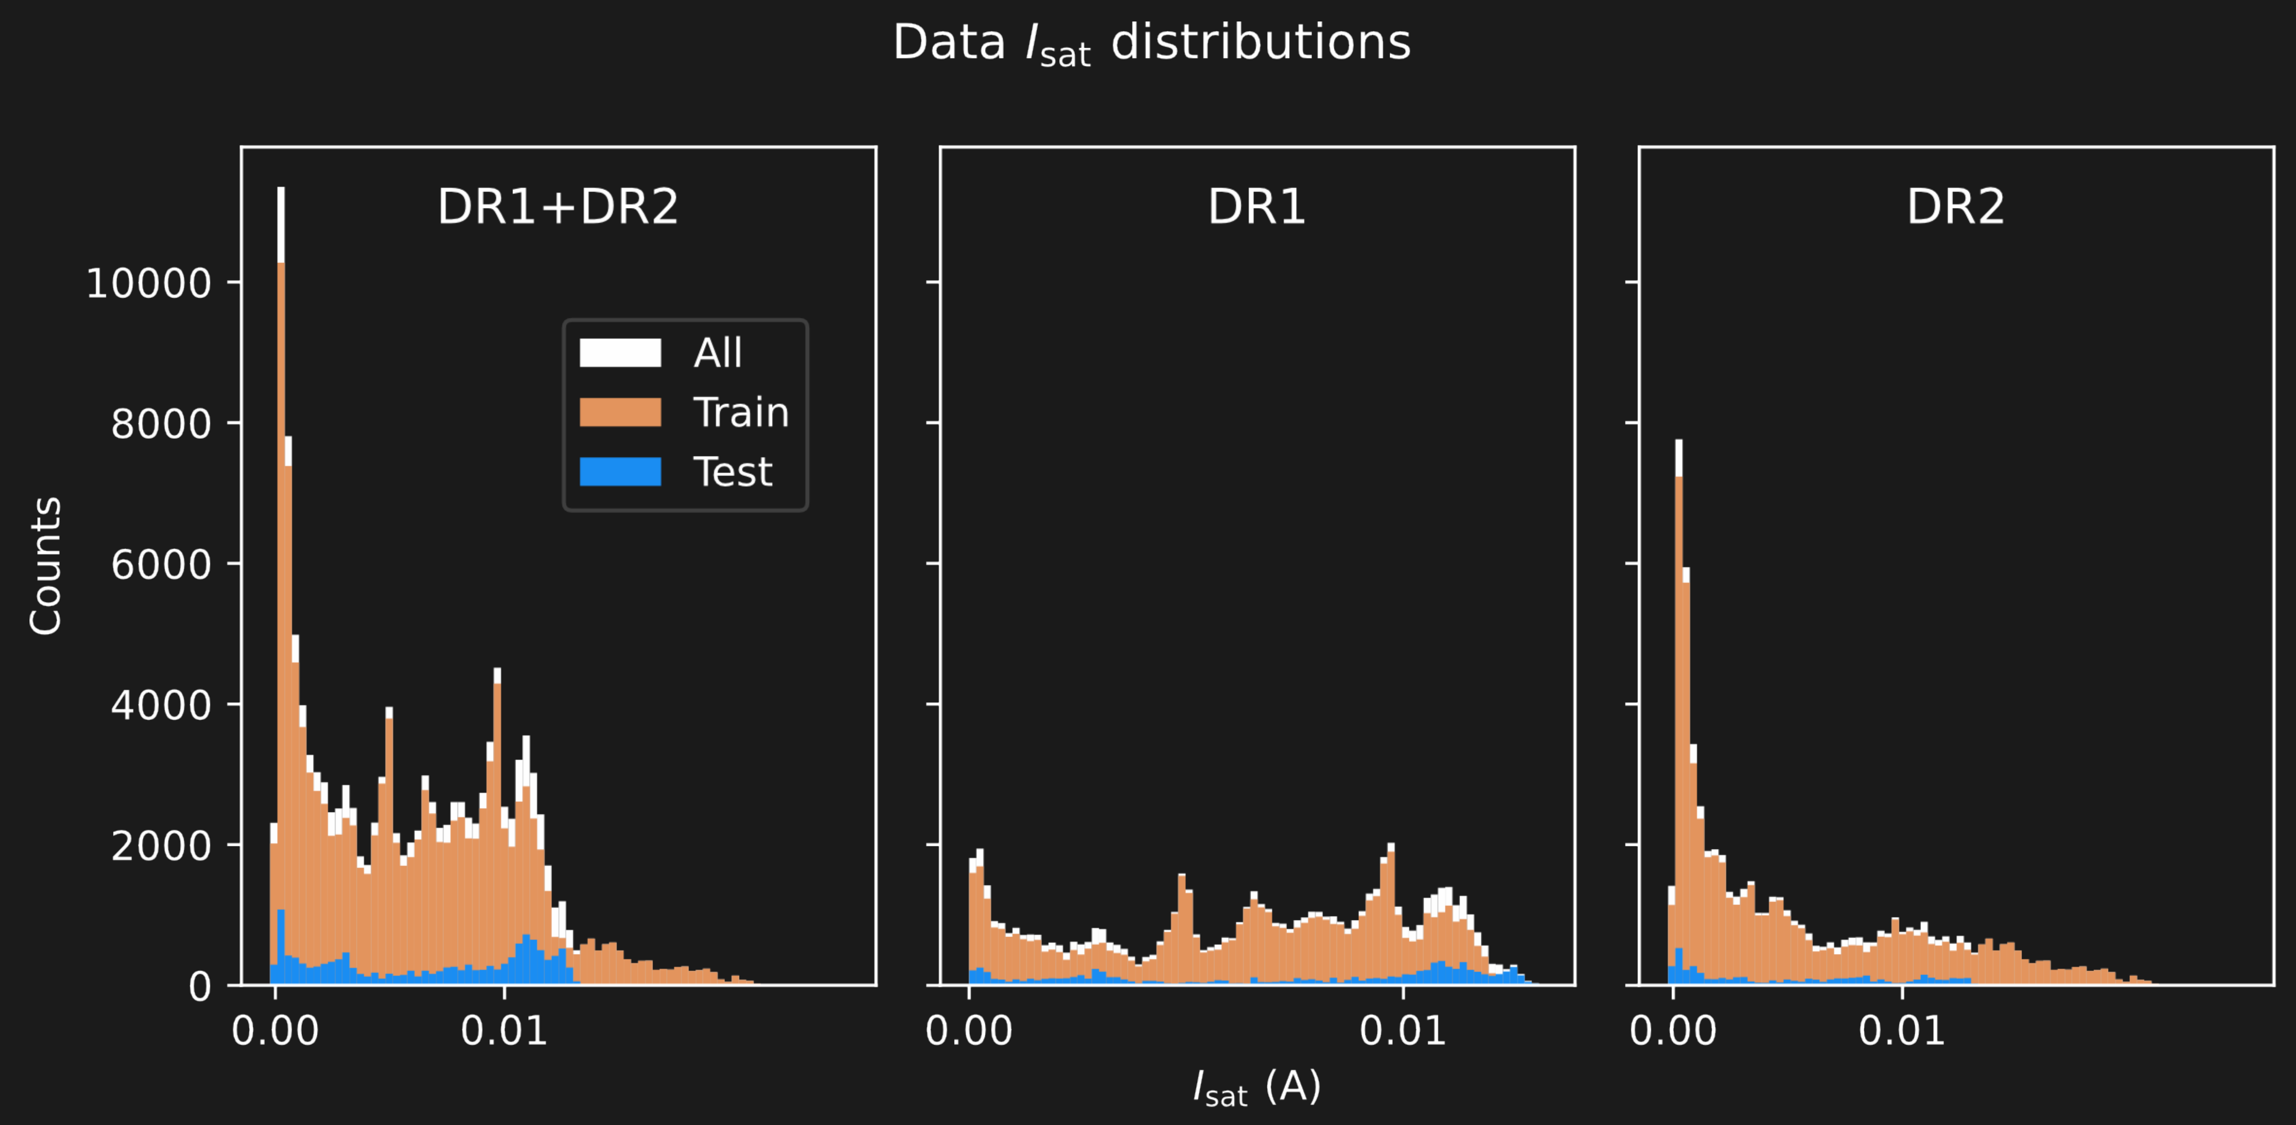

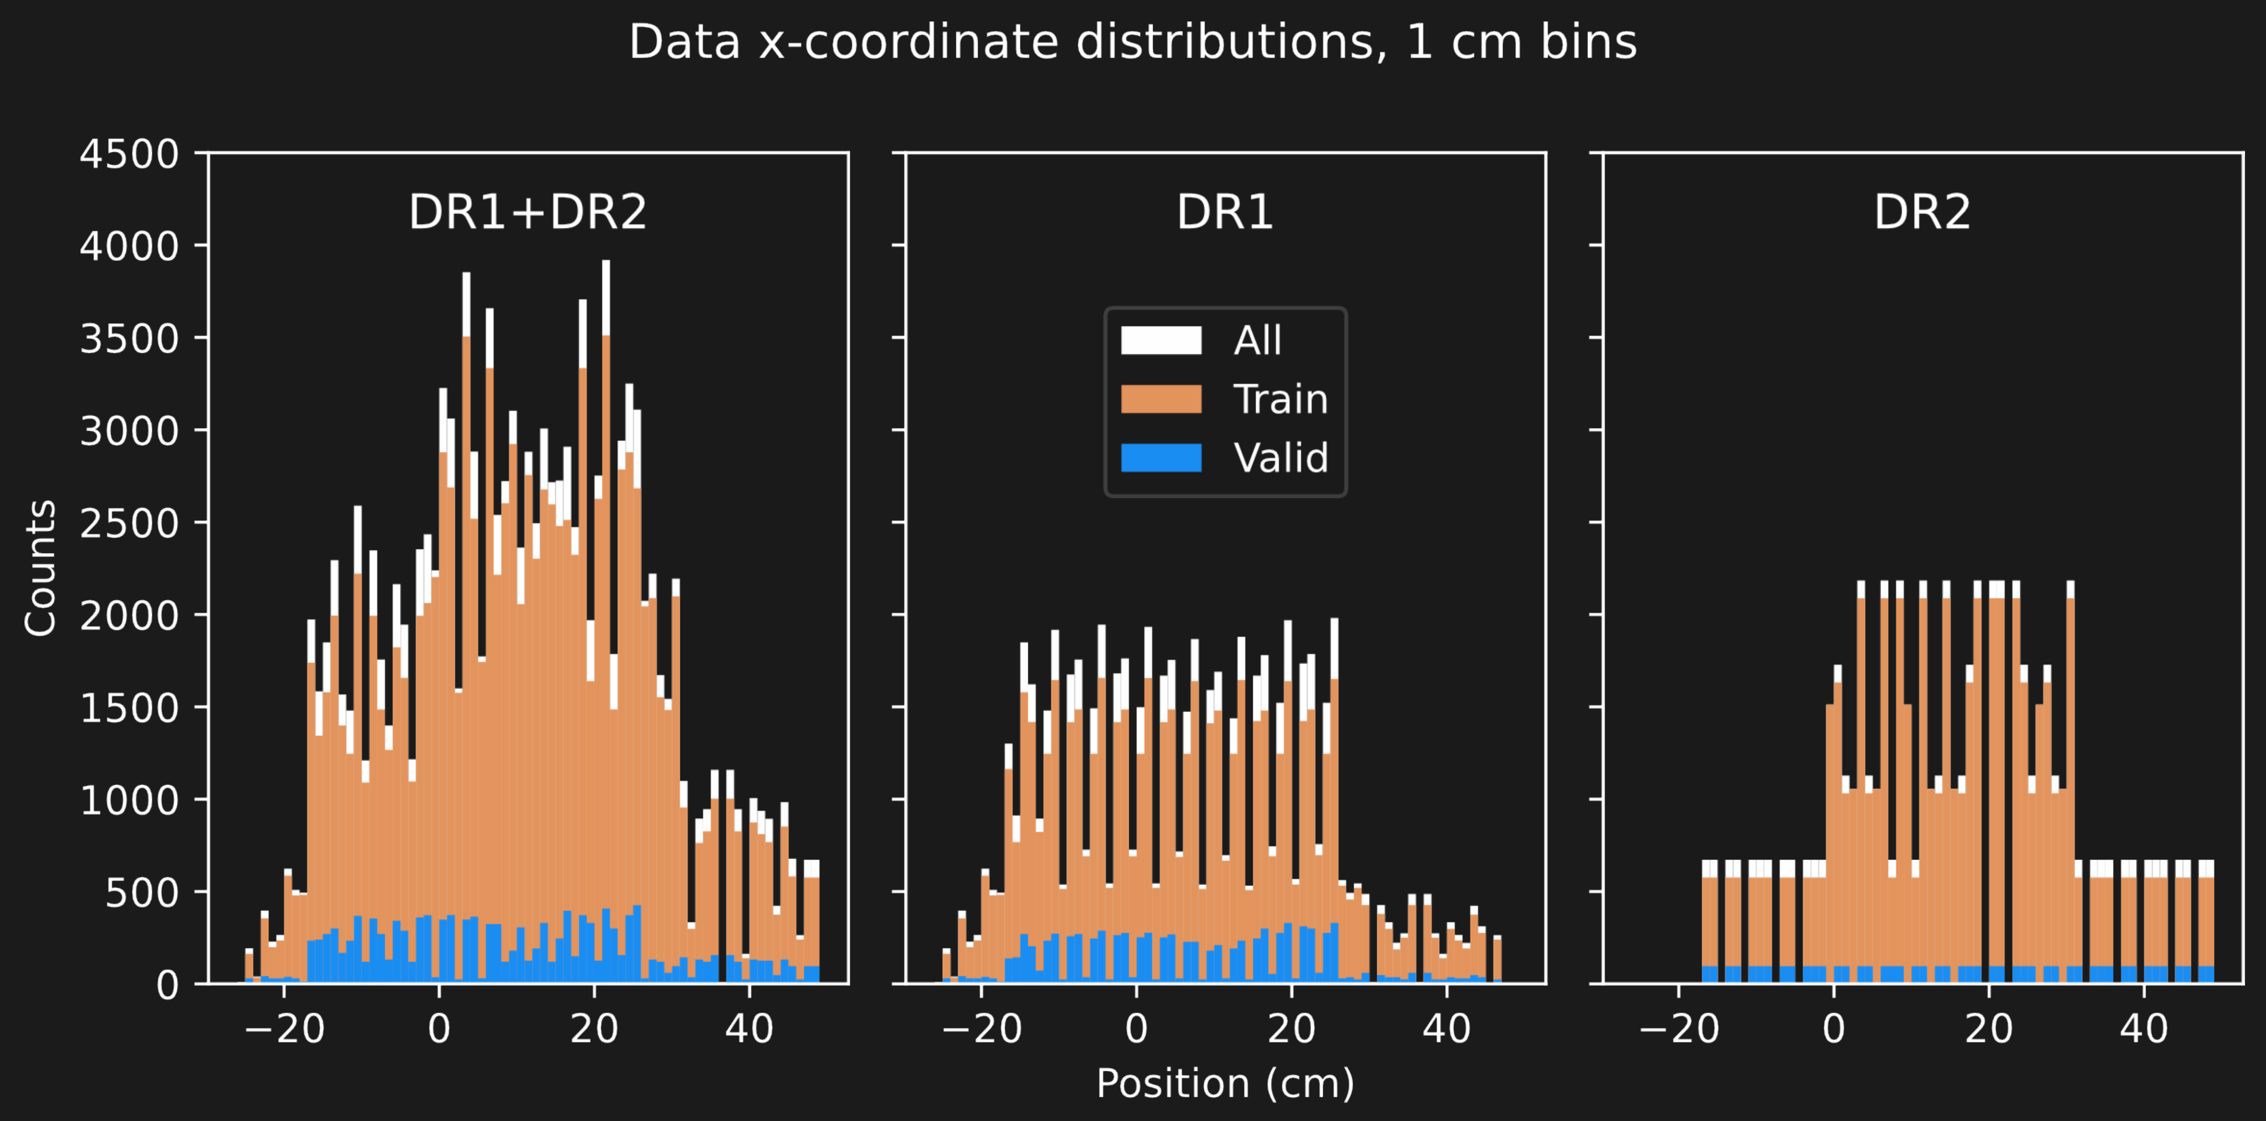

ML: isat averaging

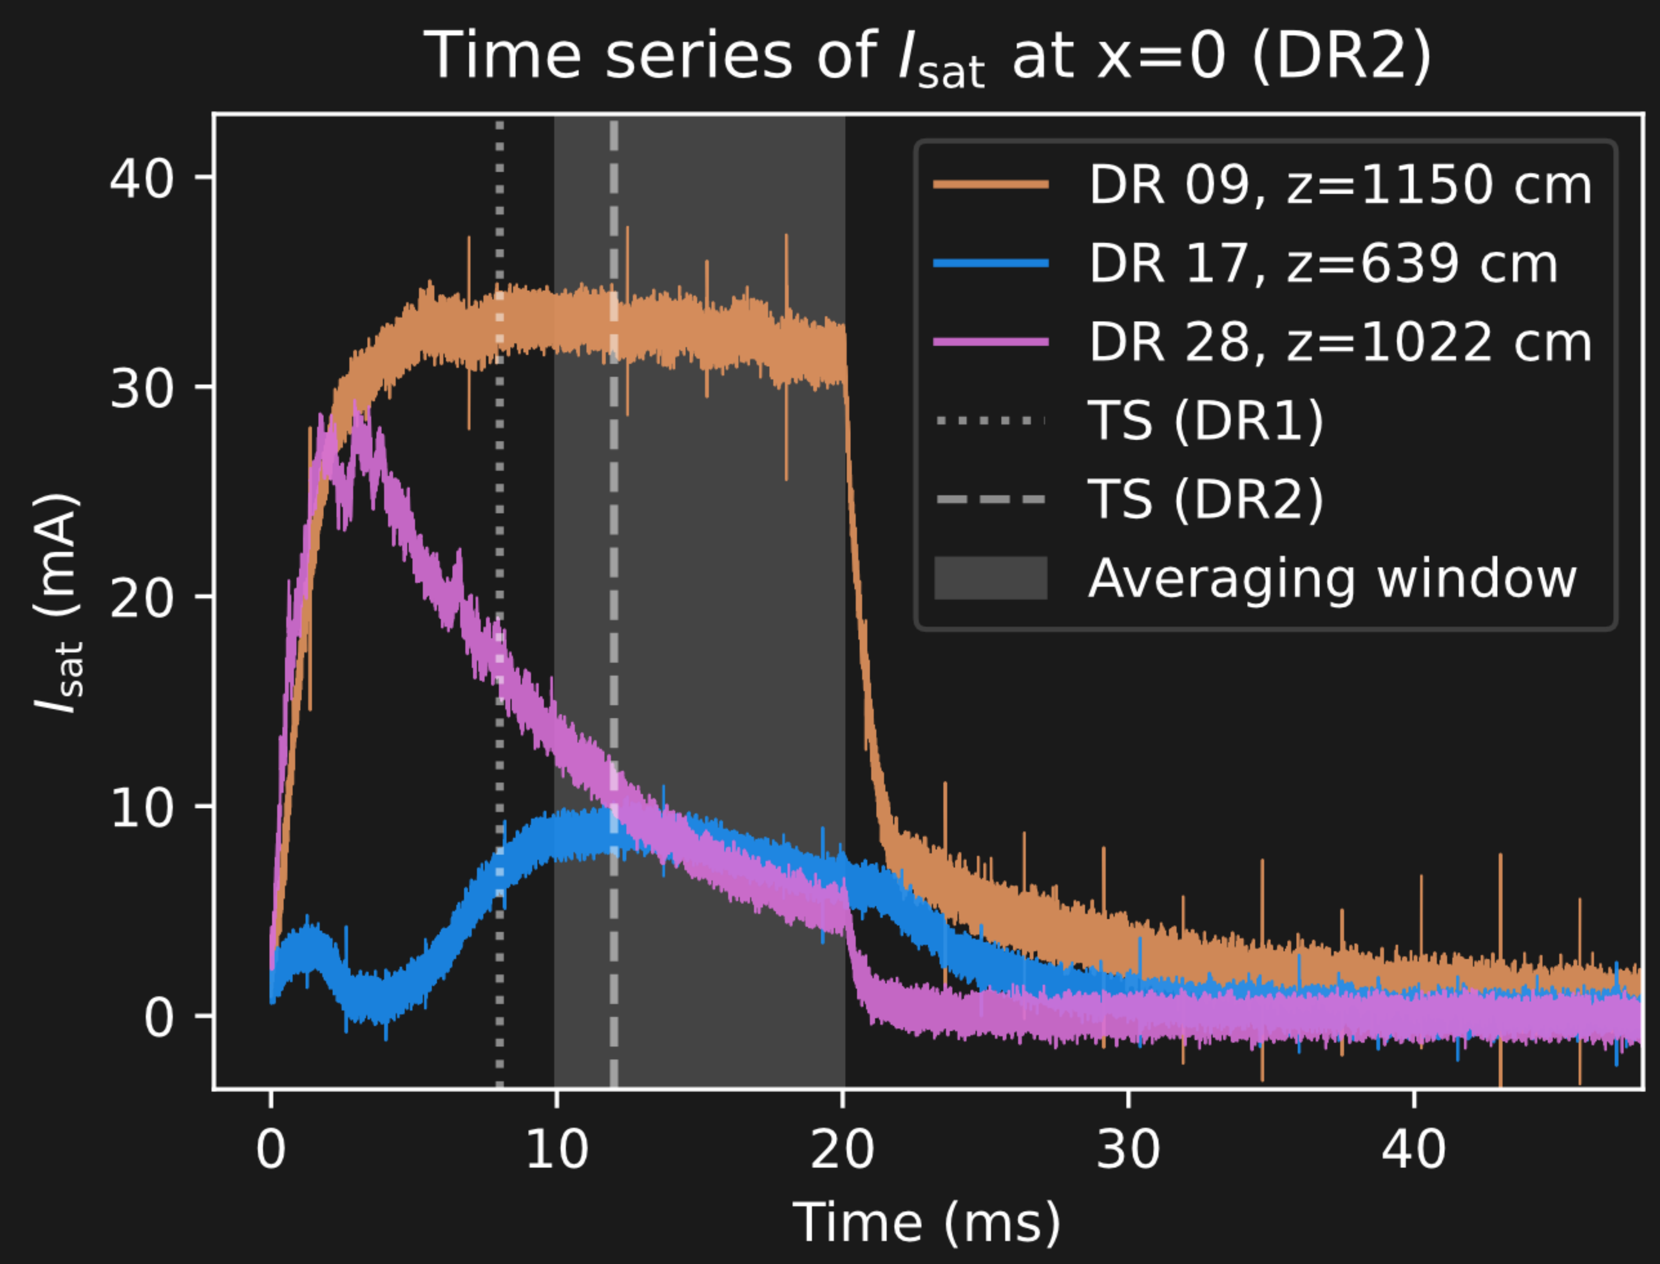

- Averaged from 10 to 20 ms after 1 kA trigger
  - Minimize complexity of the project starting out

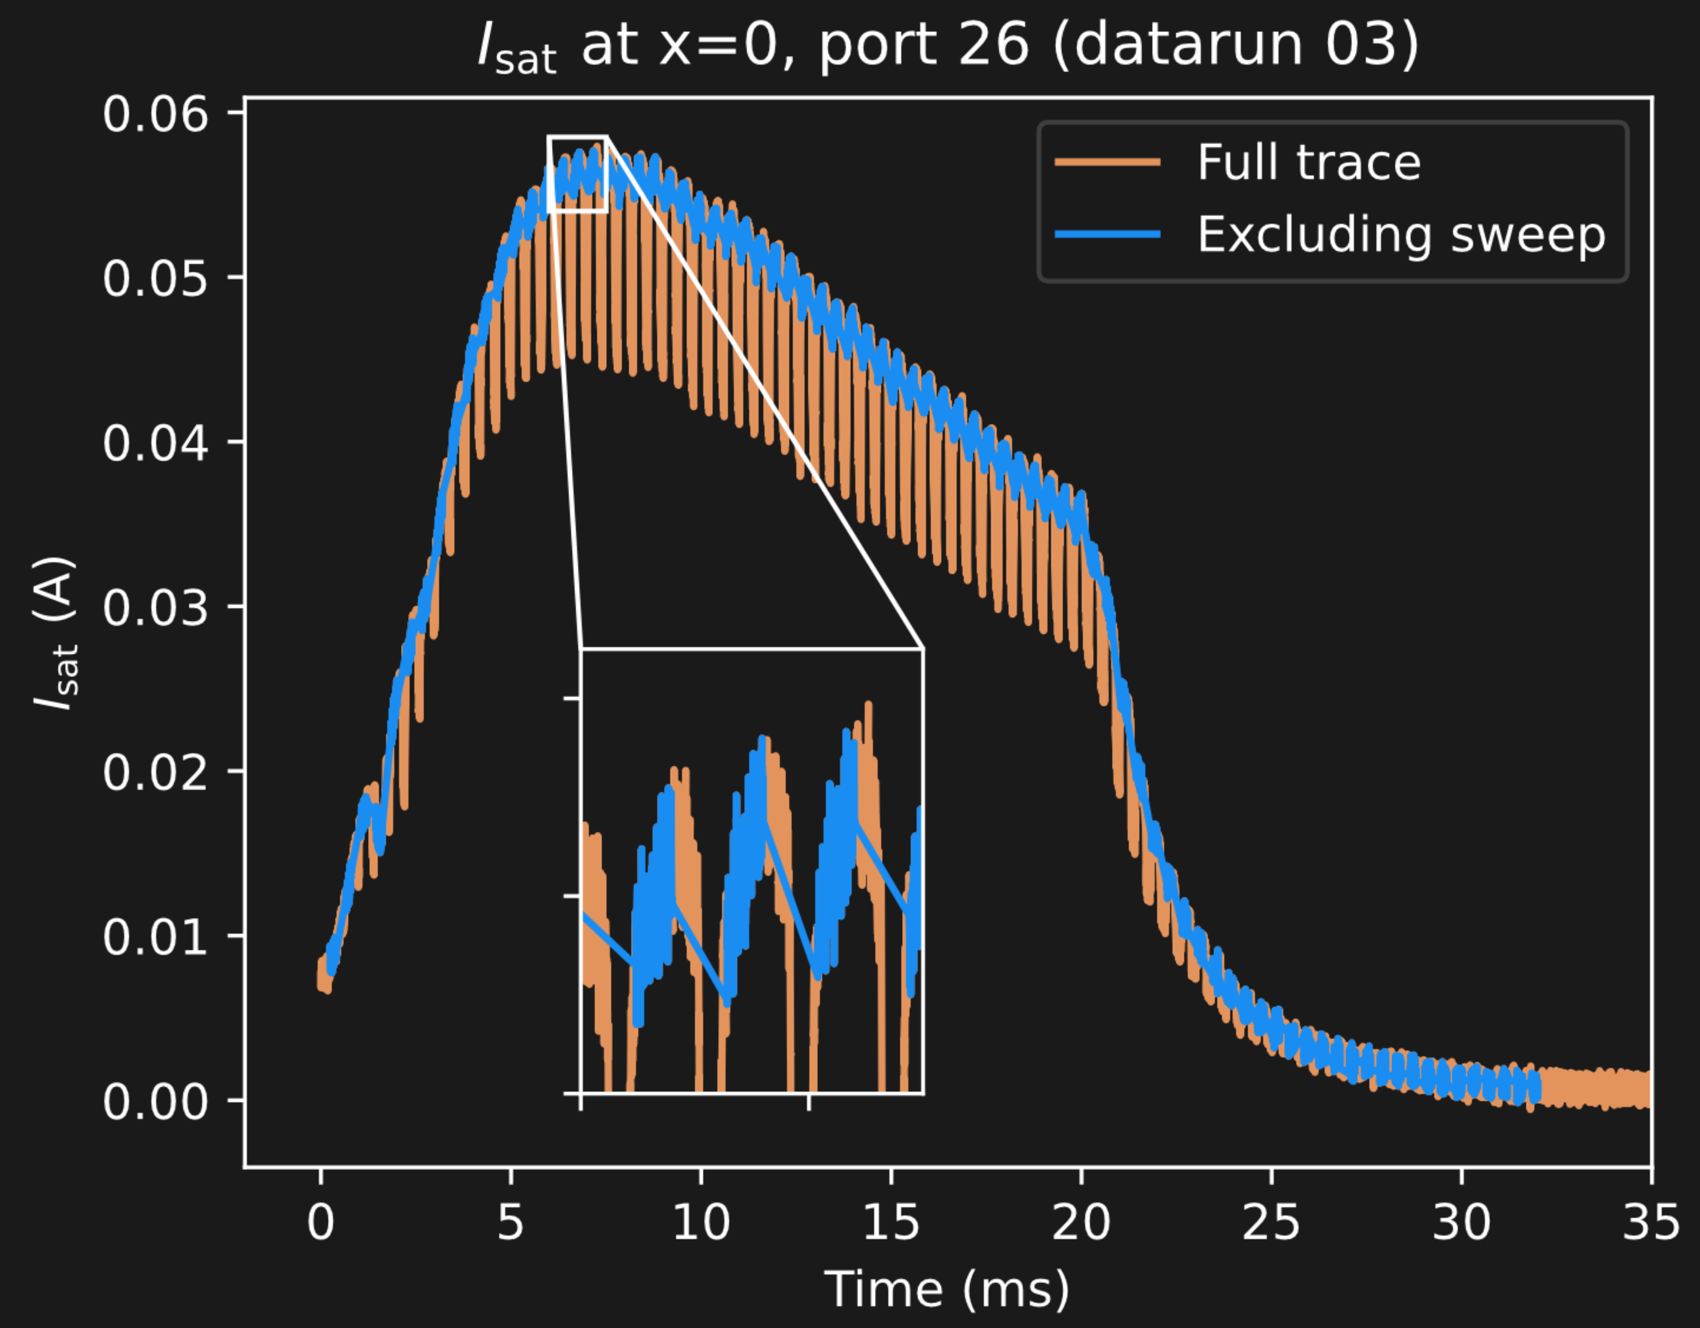

- Cleaning data is always required
  - Cut out shots that saturated the isolator or digitizer
  - Selective averaging for one of the probes

Test performance improves with more data

- Using the MSE loss function: $\mathscr{L}_{\mathbf{MSE}} = \frac{1}{m}\sum_{i=1}^{m}\left(f\left(x_i\right) - y_i\right)^2$
- 4 layers, 512 units wide
- Test set performance is improved by:
  - more dataruns in training set
  - combining both sets of dataruns
  - ensembles of models
  - larger models

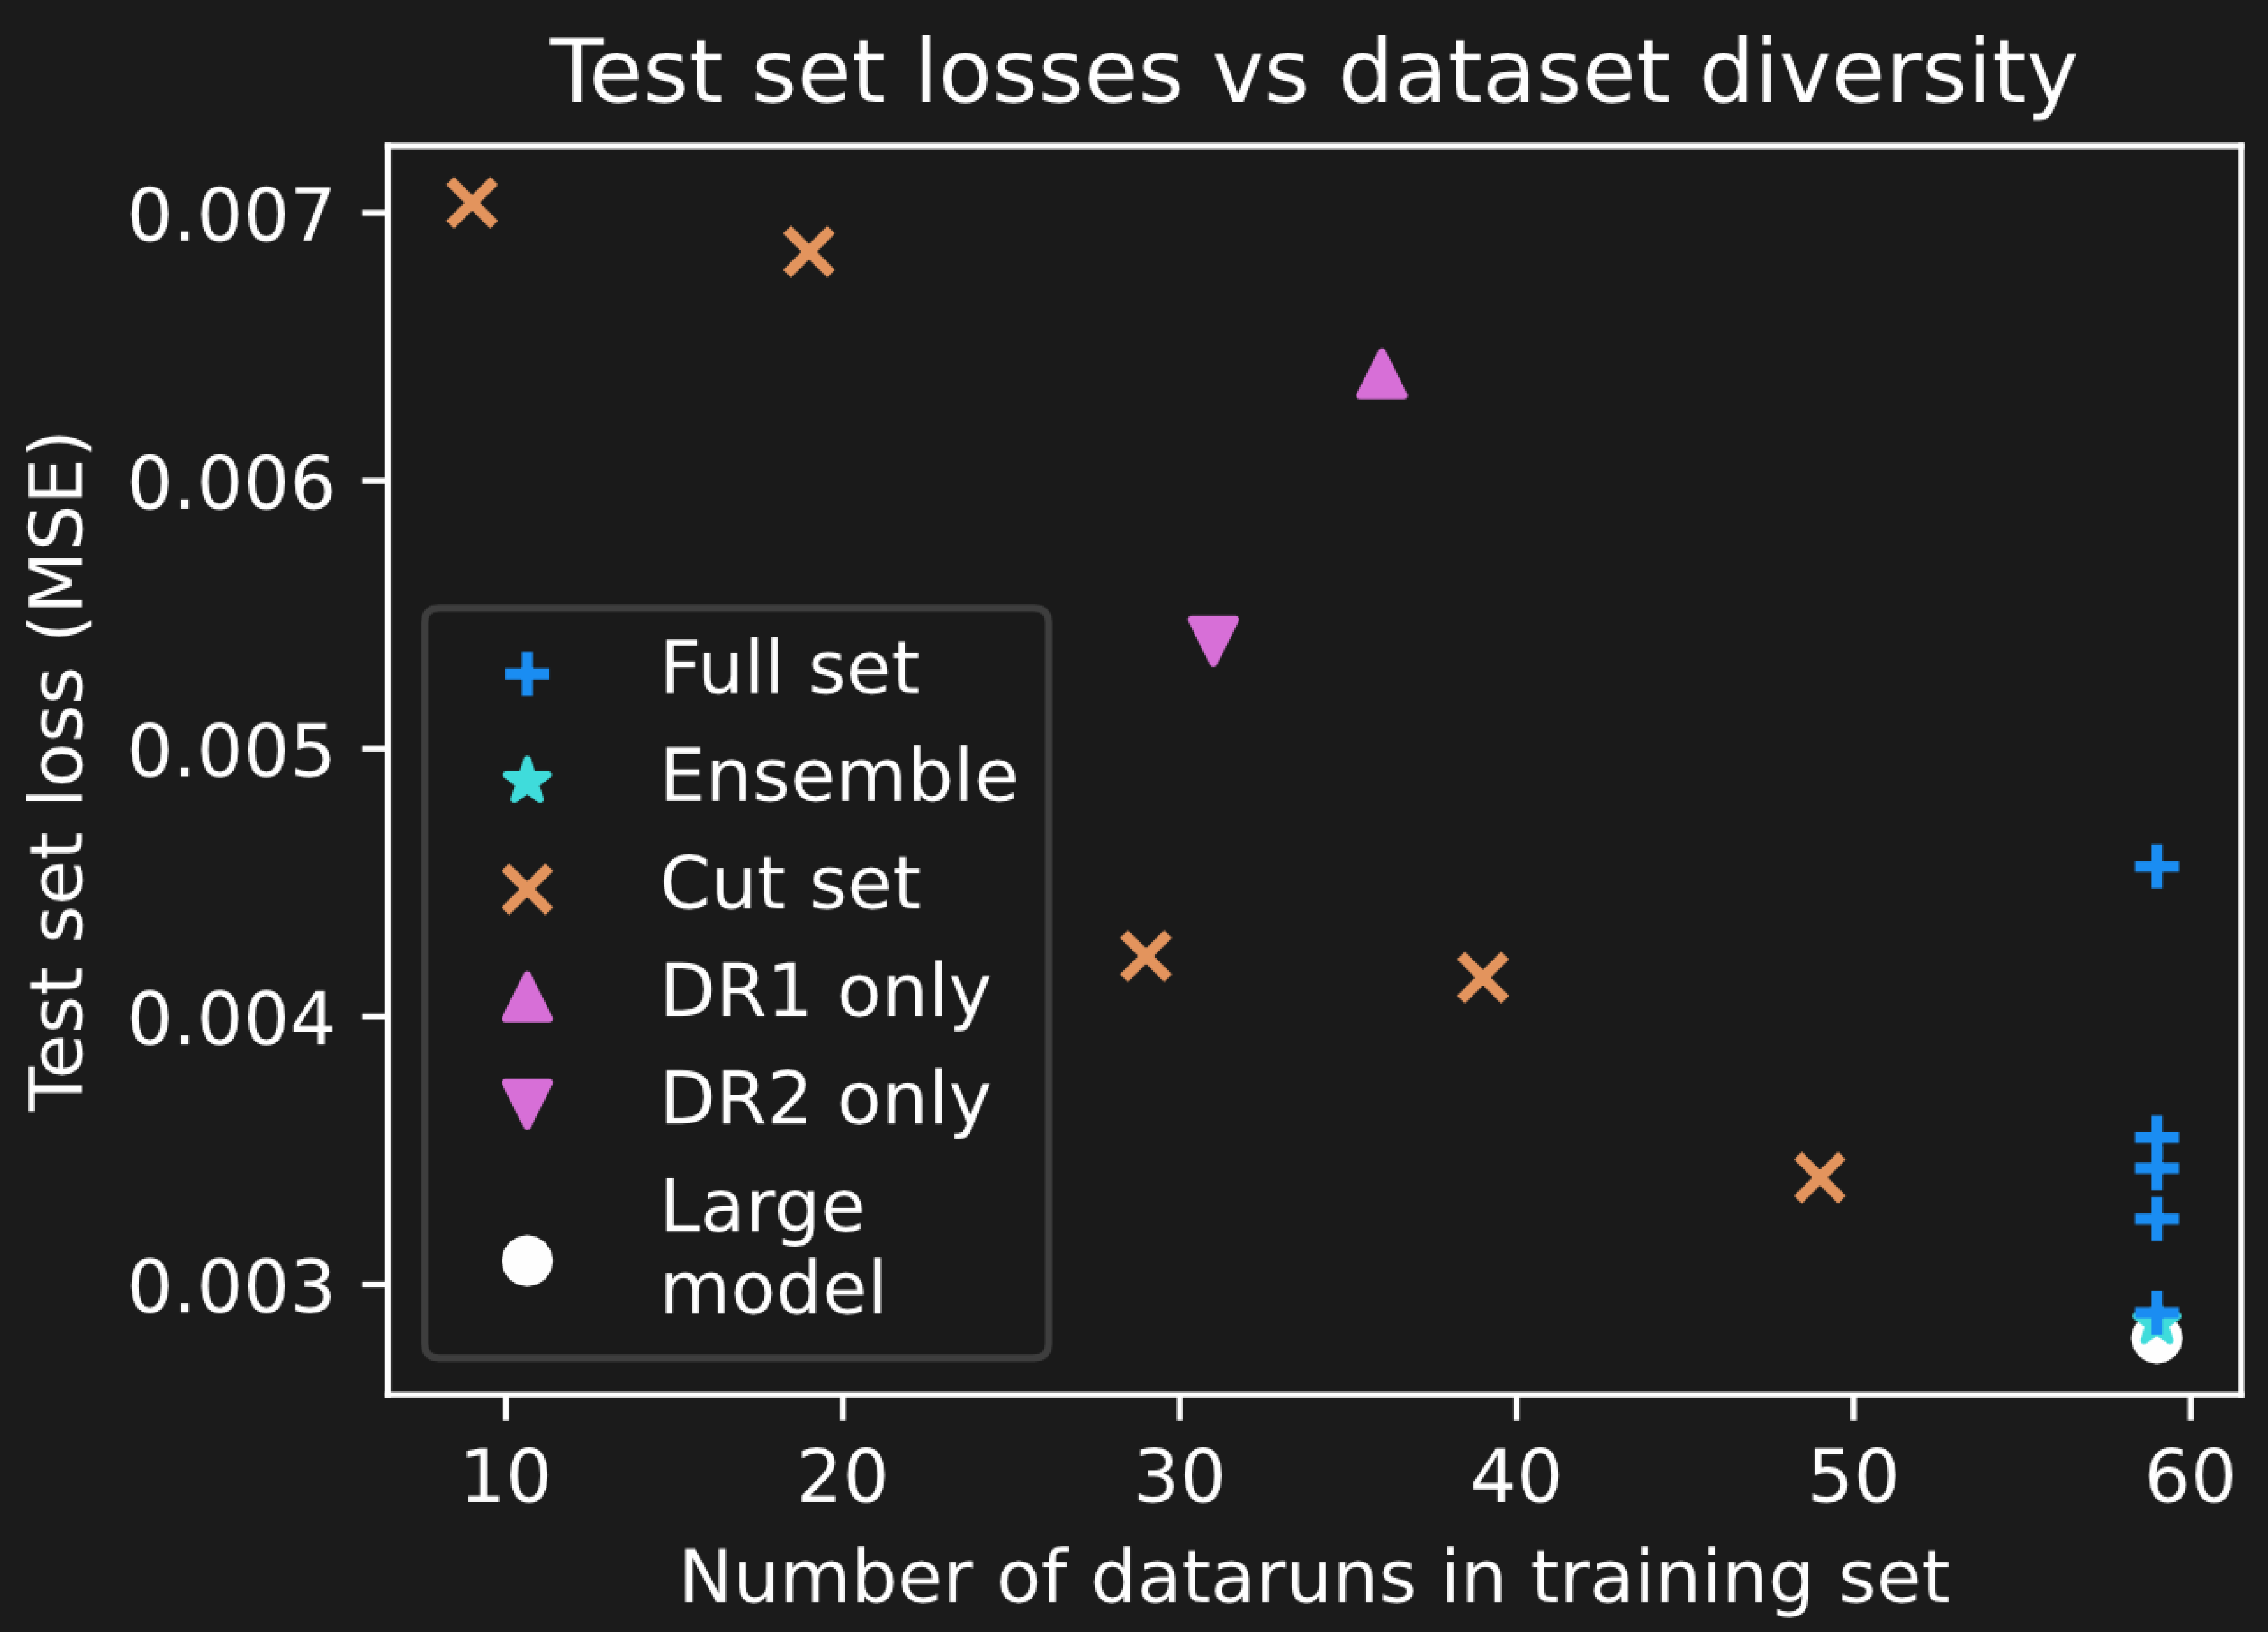

ML: Benchmarking and pipeline validation are important

- ML bugs are very insidious: nothing crashes, model performance is degraded, hard to notice
- Look for expected behavior
  - Train with zeroes for inputs
  - Train a linear model
  - Try feature engineering on the linear model (+tanh)
  - Overfit the model
- It's learning as expected

Test	MSE
Zero-input	0.036
Linear only	0.014
Linear+tanh	0.011

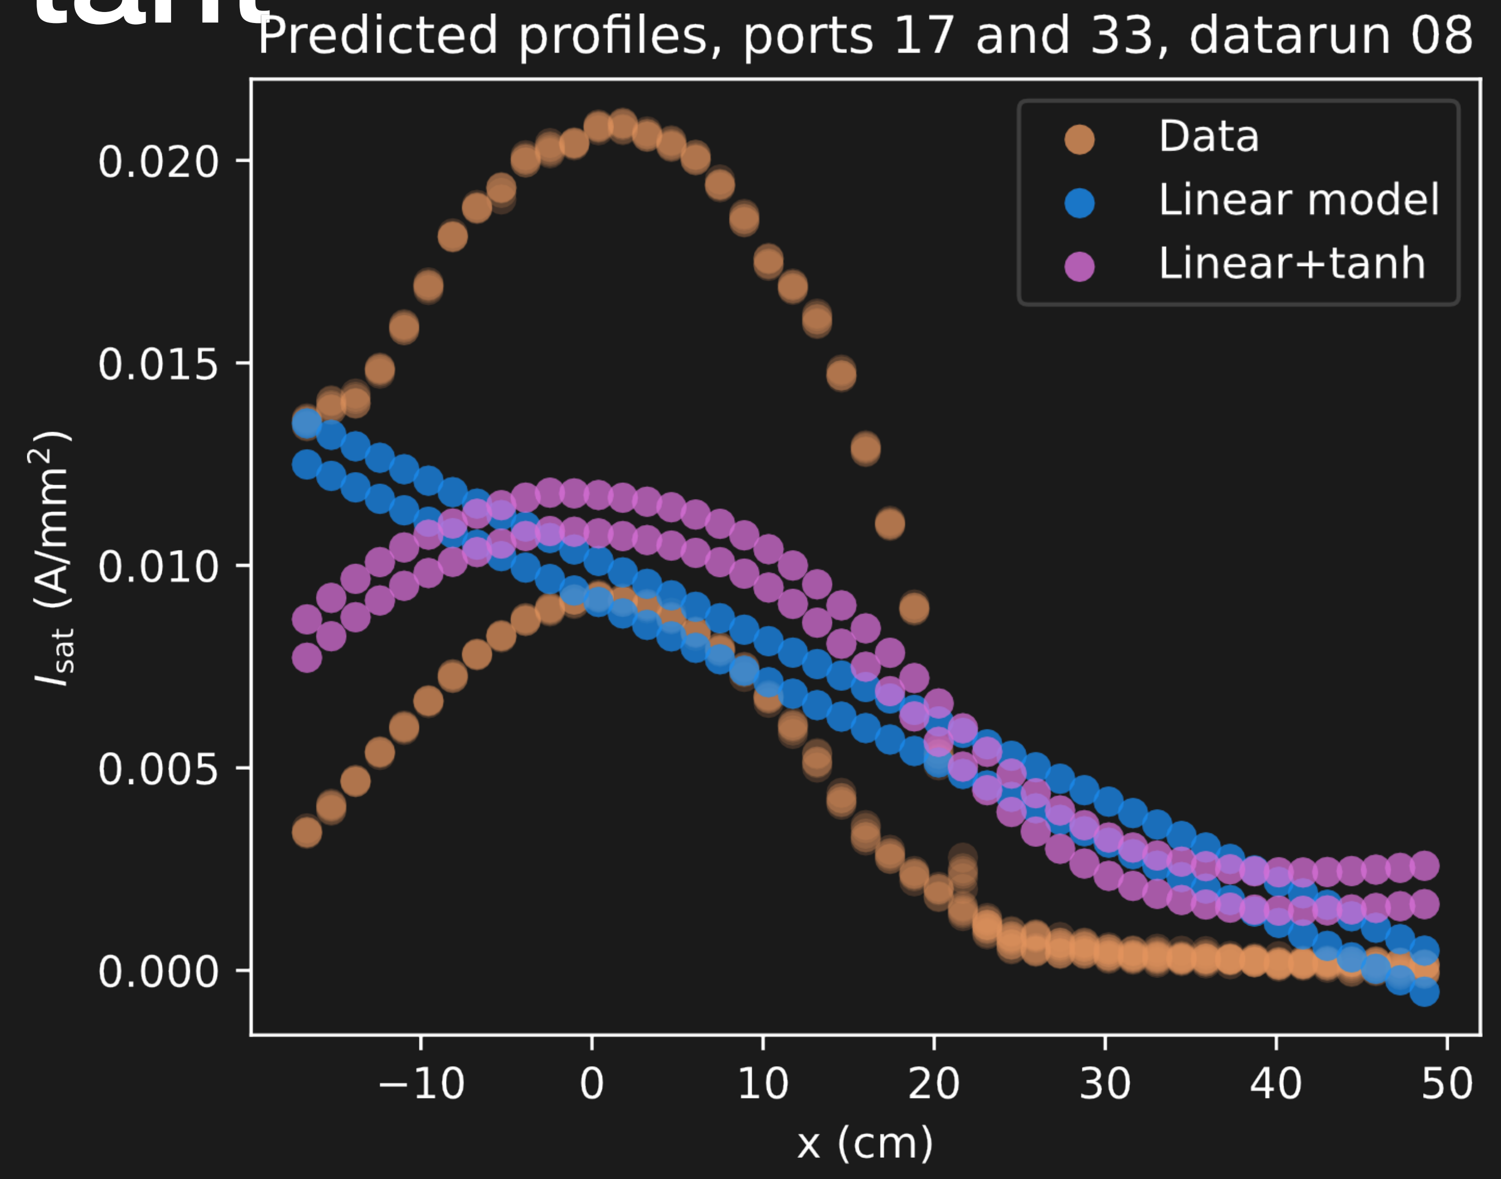

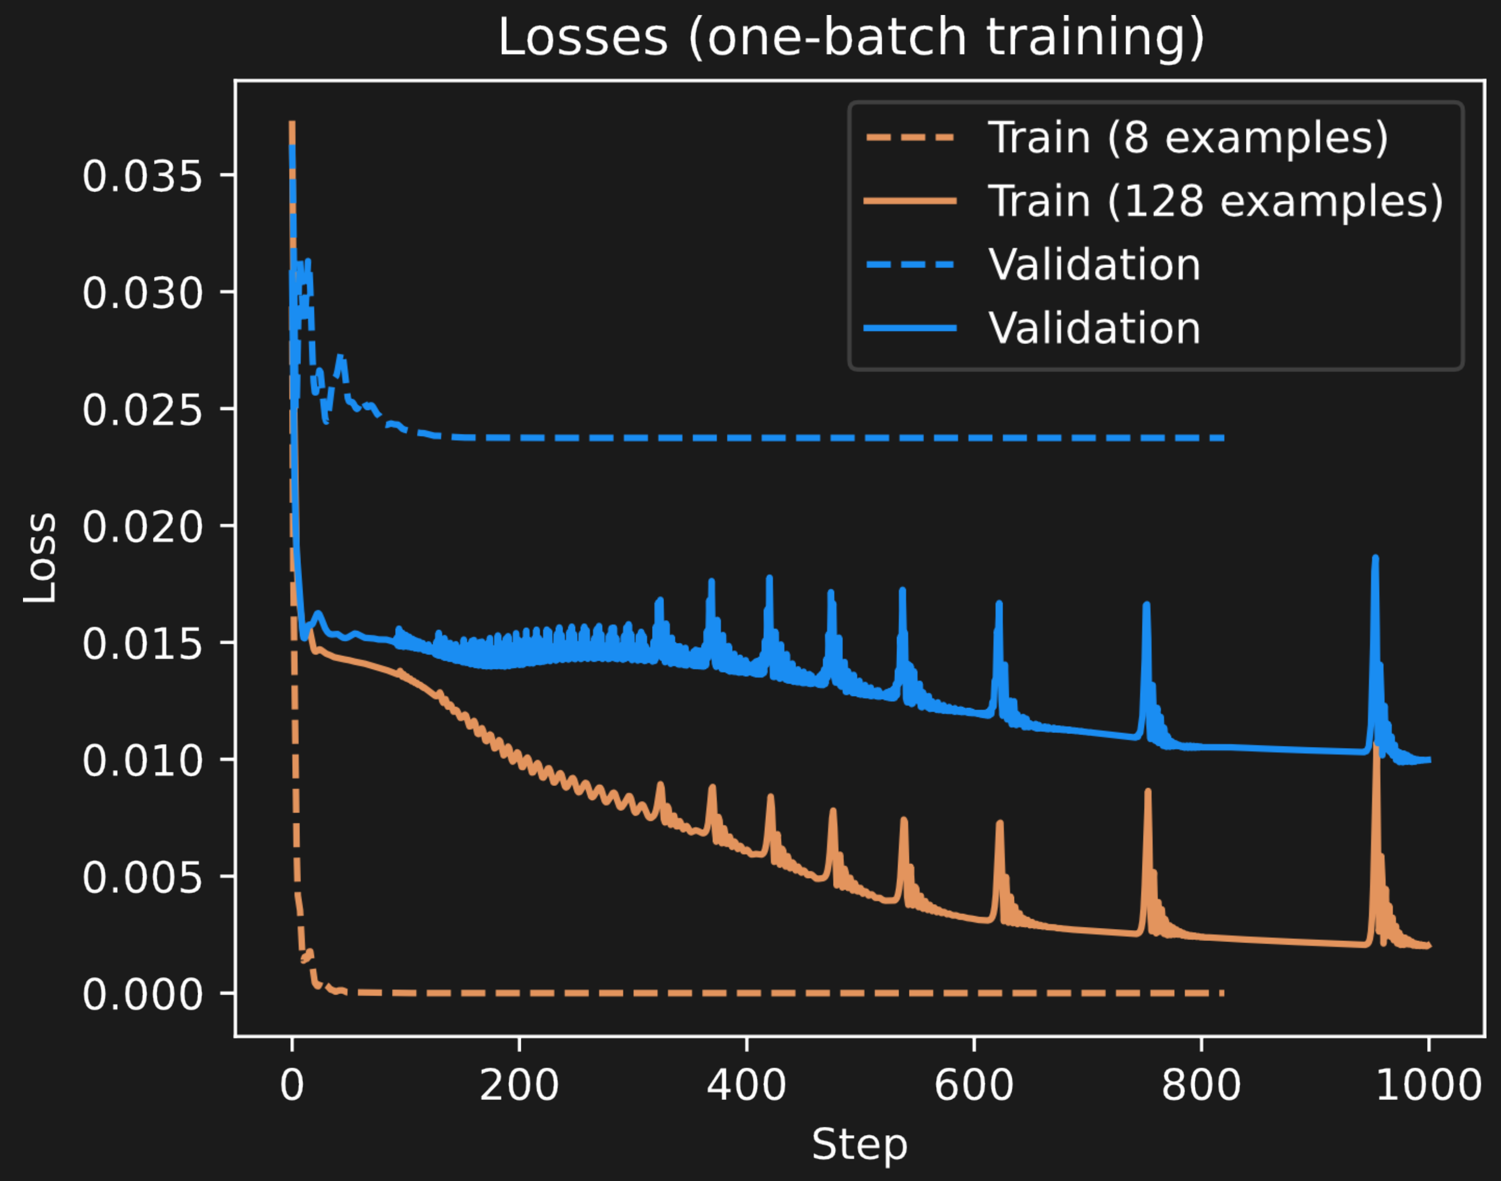

ML: z-score calibration

- z-score: squared error scaled by standard deviation
- Test set z-scores are broader than the training set
- Attempt to calibrate model through weight decay

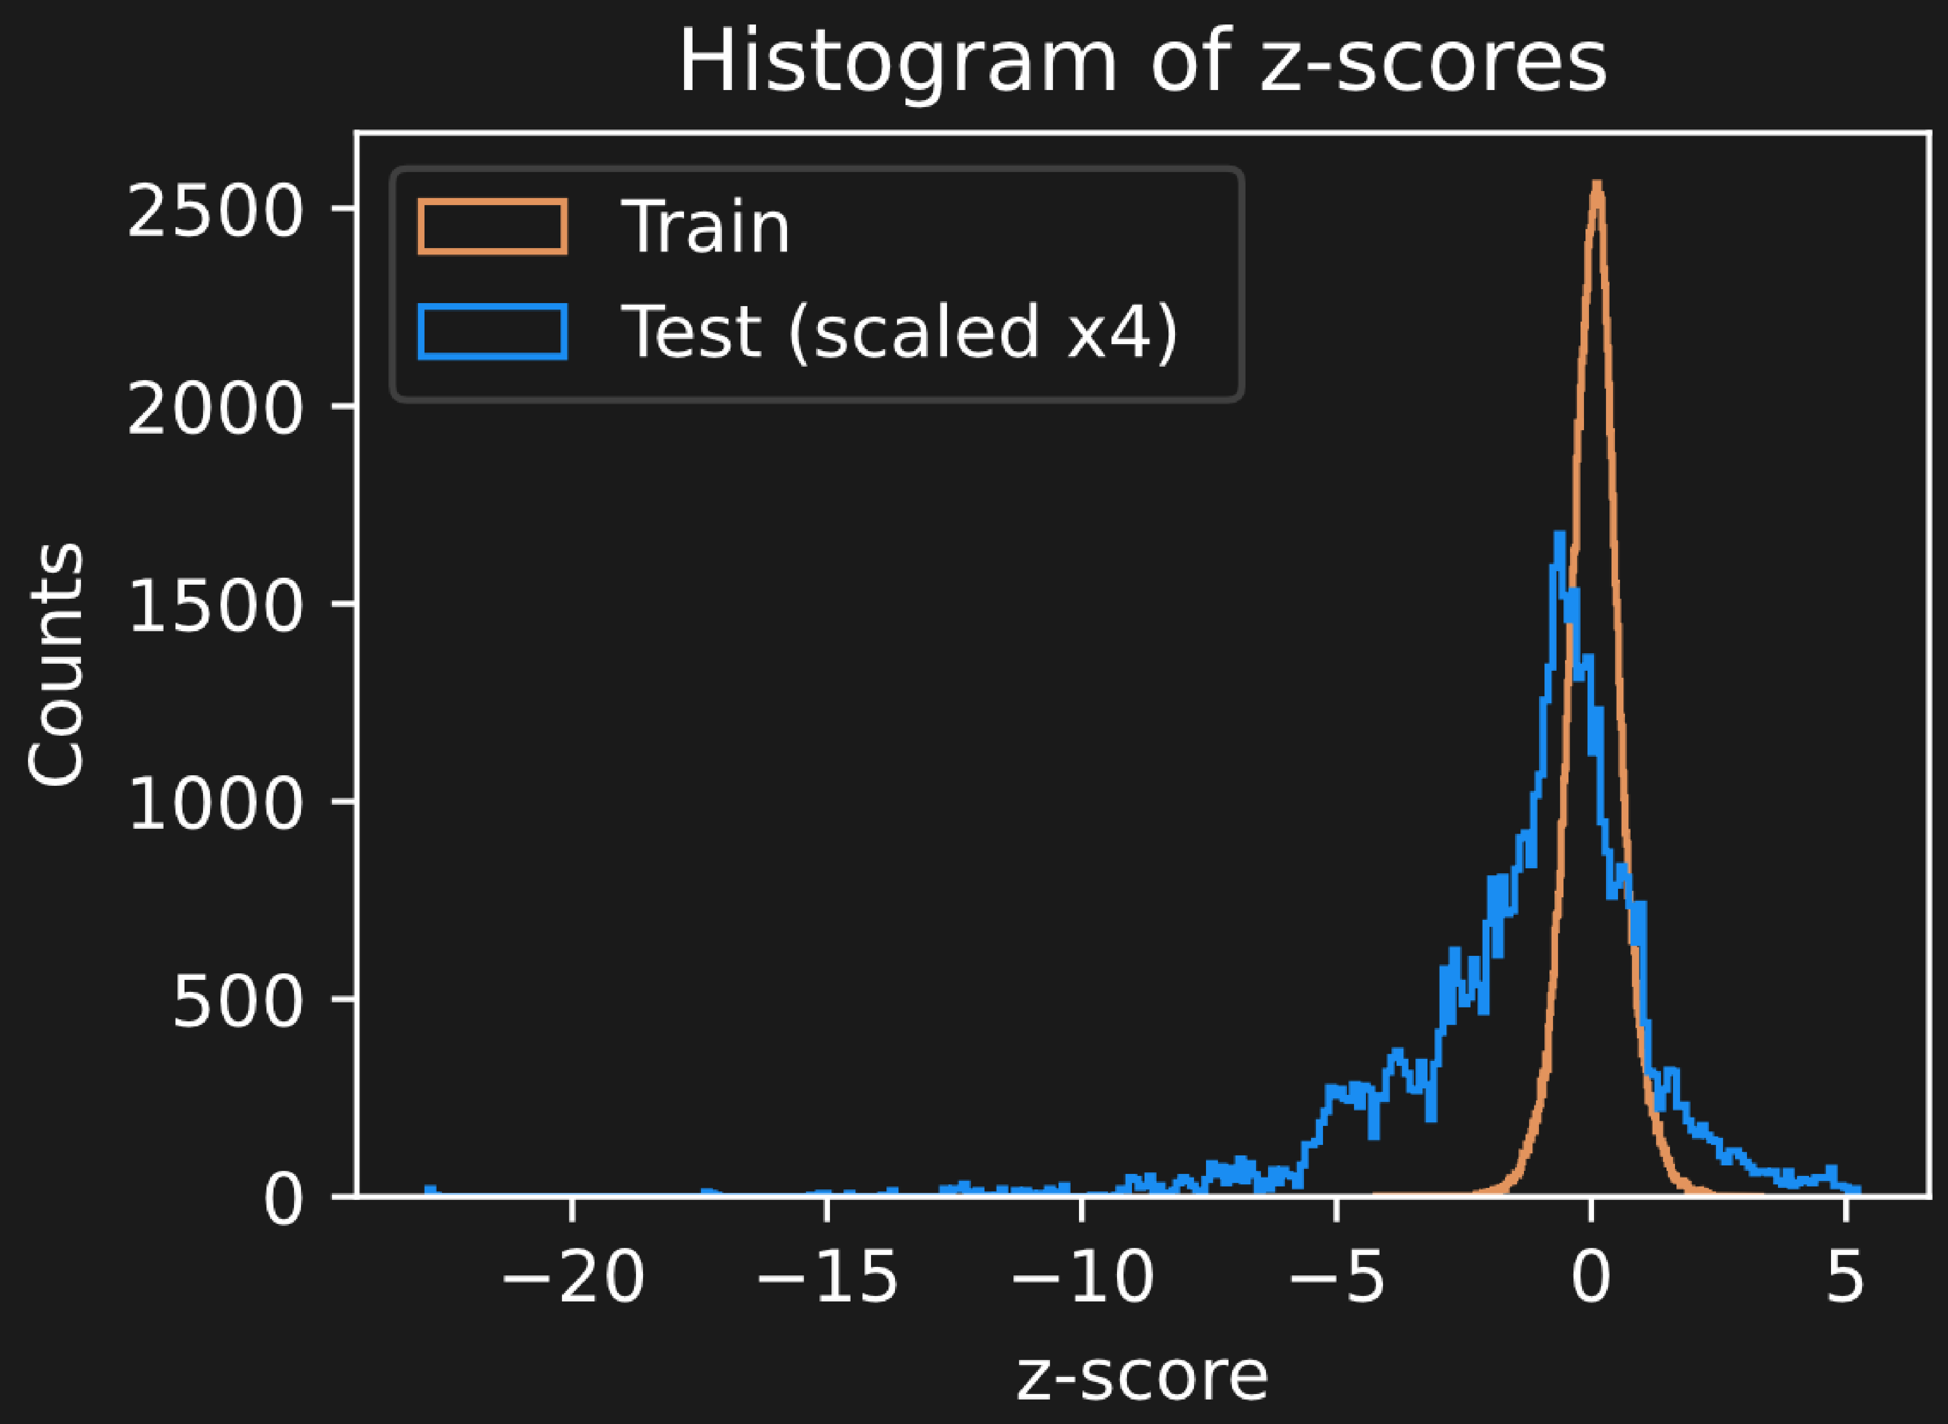

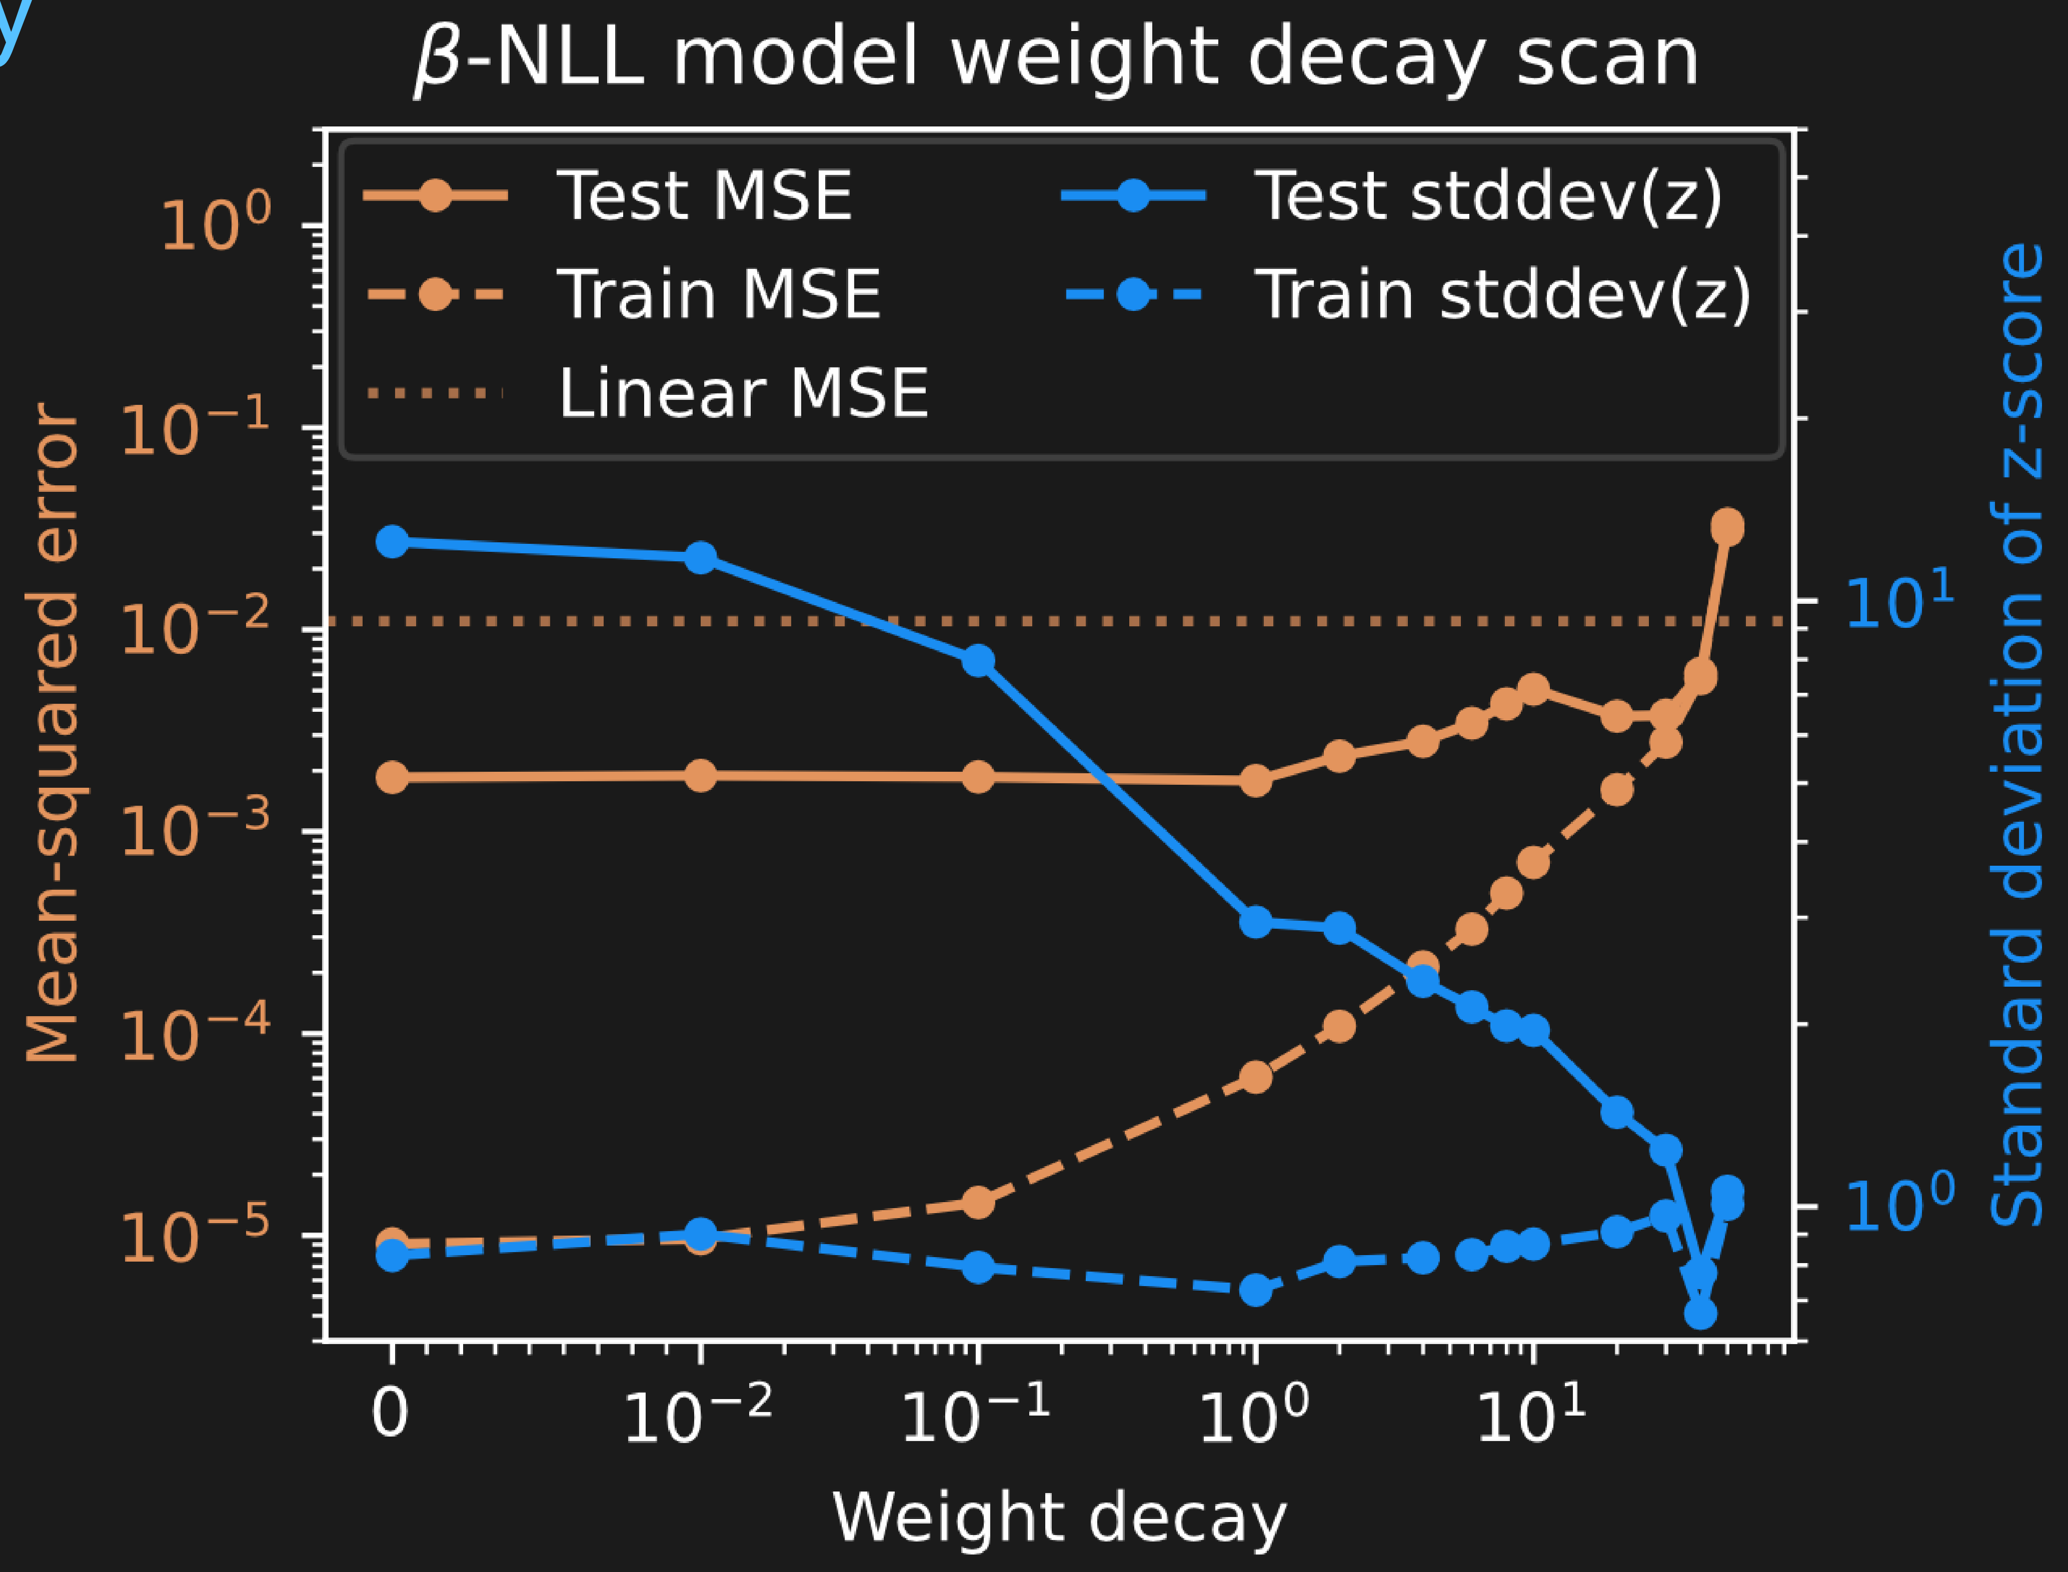

ML: nitty-gritty training details

- Leaky ReLU activations
- AdamW optimizer
- 4 layers, 256 width (occasionally 512 or 1024)
- No weight decay
- Gradient clipping (percentile and absolute)
- No other regularization
- Models take ~30 min to train for 500 epochs (no early stopping)
- All in PyTorch

ML: optimization / search parameters

Table 5.2: Machine inputs and actuators for model inference

Input or actuator	Range	Step	Count
Source field	500 G to 2000 G	250 G	7
Mirror field	250 G to 1500 G	250 G	6
Midplane field	250 G to 1500 G	250 G	6
Gas puff voltage	70 V to 90 V	5 V	5
Discharge voltage	70 V to 150 V	10 V	9
Gas puff duration	5 ms to 38 ms	8.25 ms	5
Probe x positions	-50 cm to 50 cm	2 cm	51
Probe y positions	0 cm	–	–
Probe z positions	640 cm to 1140 cm	50 cm	11
Probe angle	0 rad	–	–
Run set flag	off and on	1	2
Top gas puff flag	off and on	1	2

Table 5.3: Machine inputs and actuators for optimized axial profiles

Input or actuator	Weakest	**Weakest**	**Strongest**	Intermediate
I_{sat} constraint (mA/mm^2)	I_{sat} = any	$I_{sat} > 7.5$	$I_{sat} > 7.5$	$I_{sat} > 7.5$
Source field	750 G	1000 G	500 G	2000 G
Mirror field	1000 G	750 G	500 G	1250 G
Midplane field	250 G	250 G	1500 G	750 G
Gas puff voltage	70 V	75 V	90 V	90 V
Discharge voltage	130 V	150 V	150 V	120 V
Gas puff duration	5 ms	5 ms	38 ms	38 ms
Run set flag	on	on	on	on
Top gas puff flag	on	off	off	off

EBM: the model captures all modes of the probability distribution

- Model captures all modes of the distribution
  - GANs, VAEs can struggle

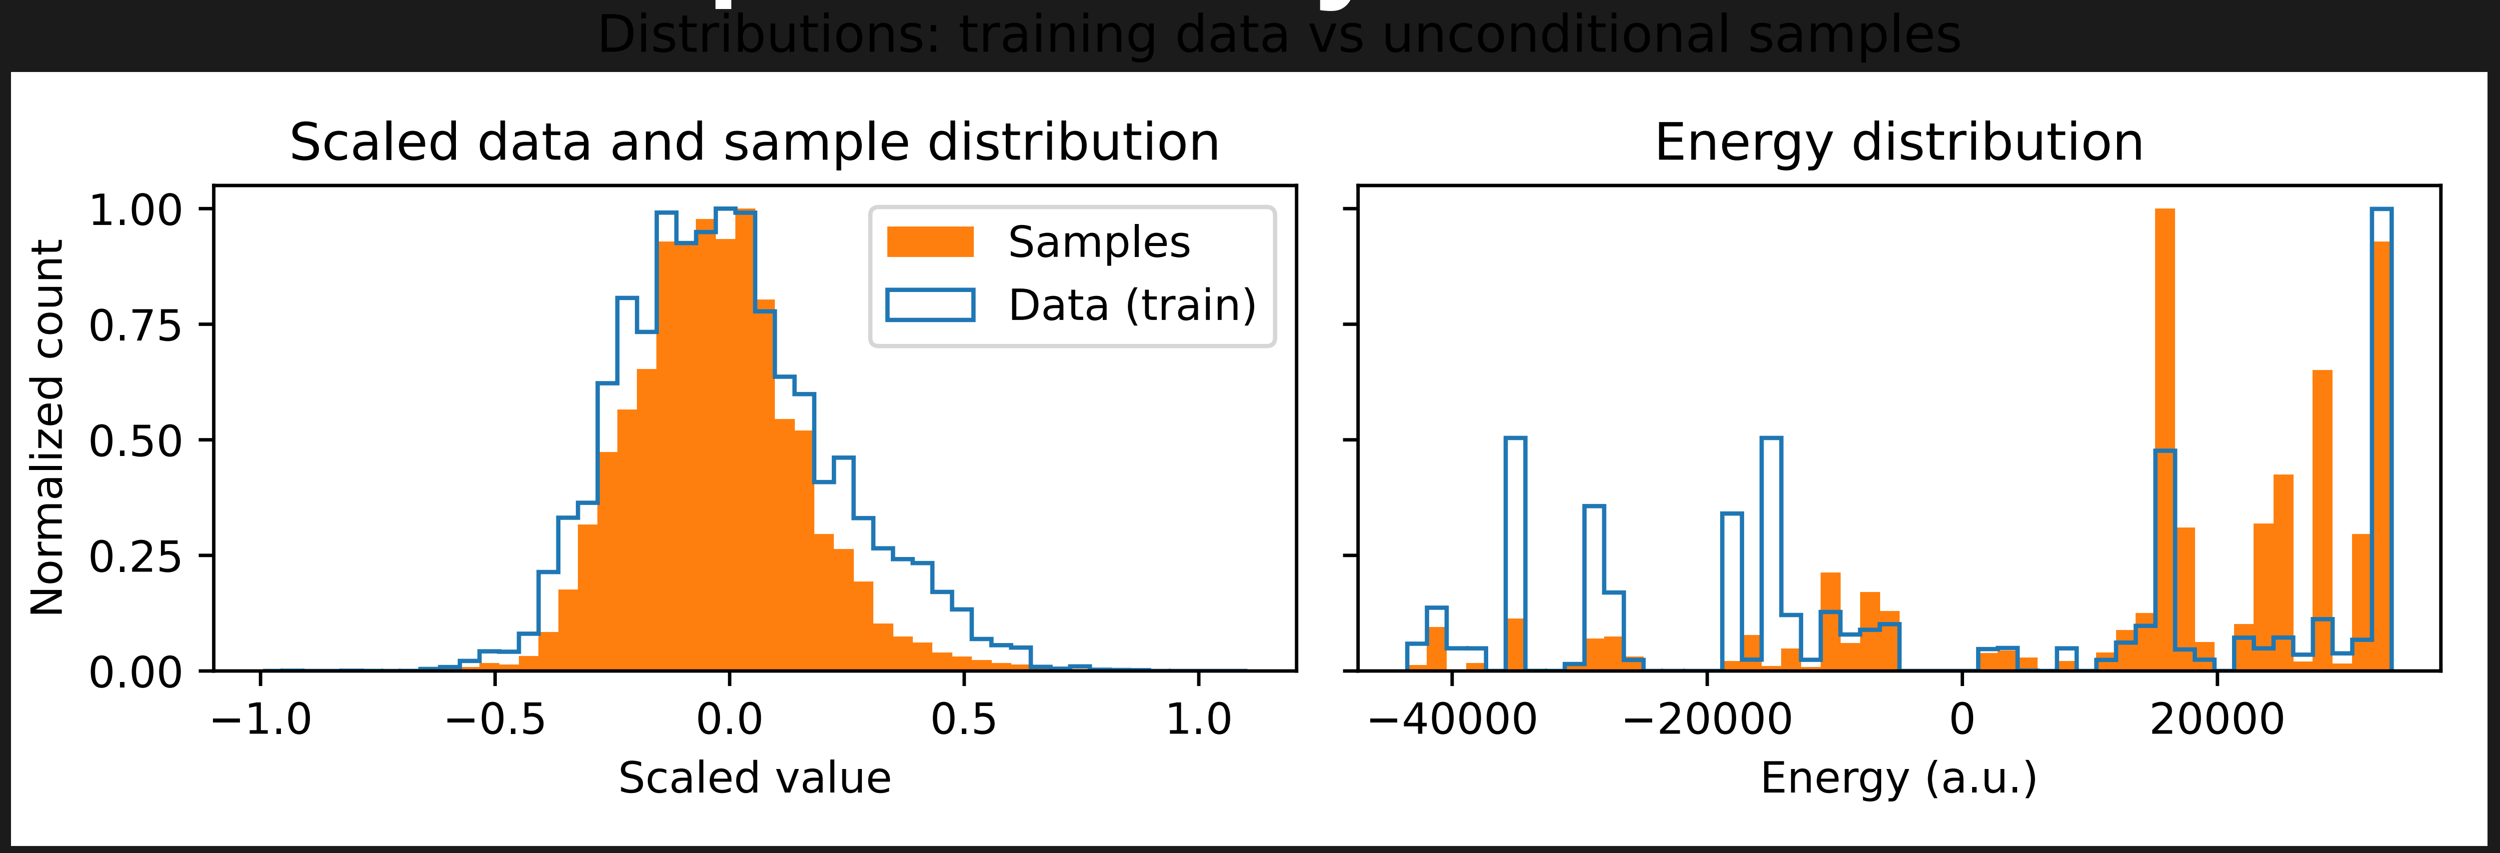

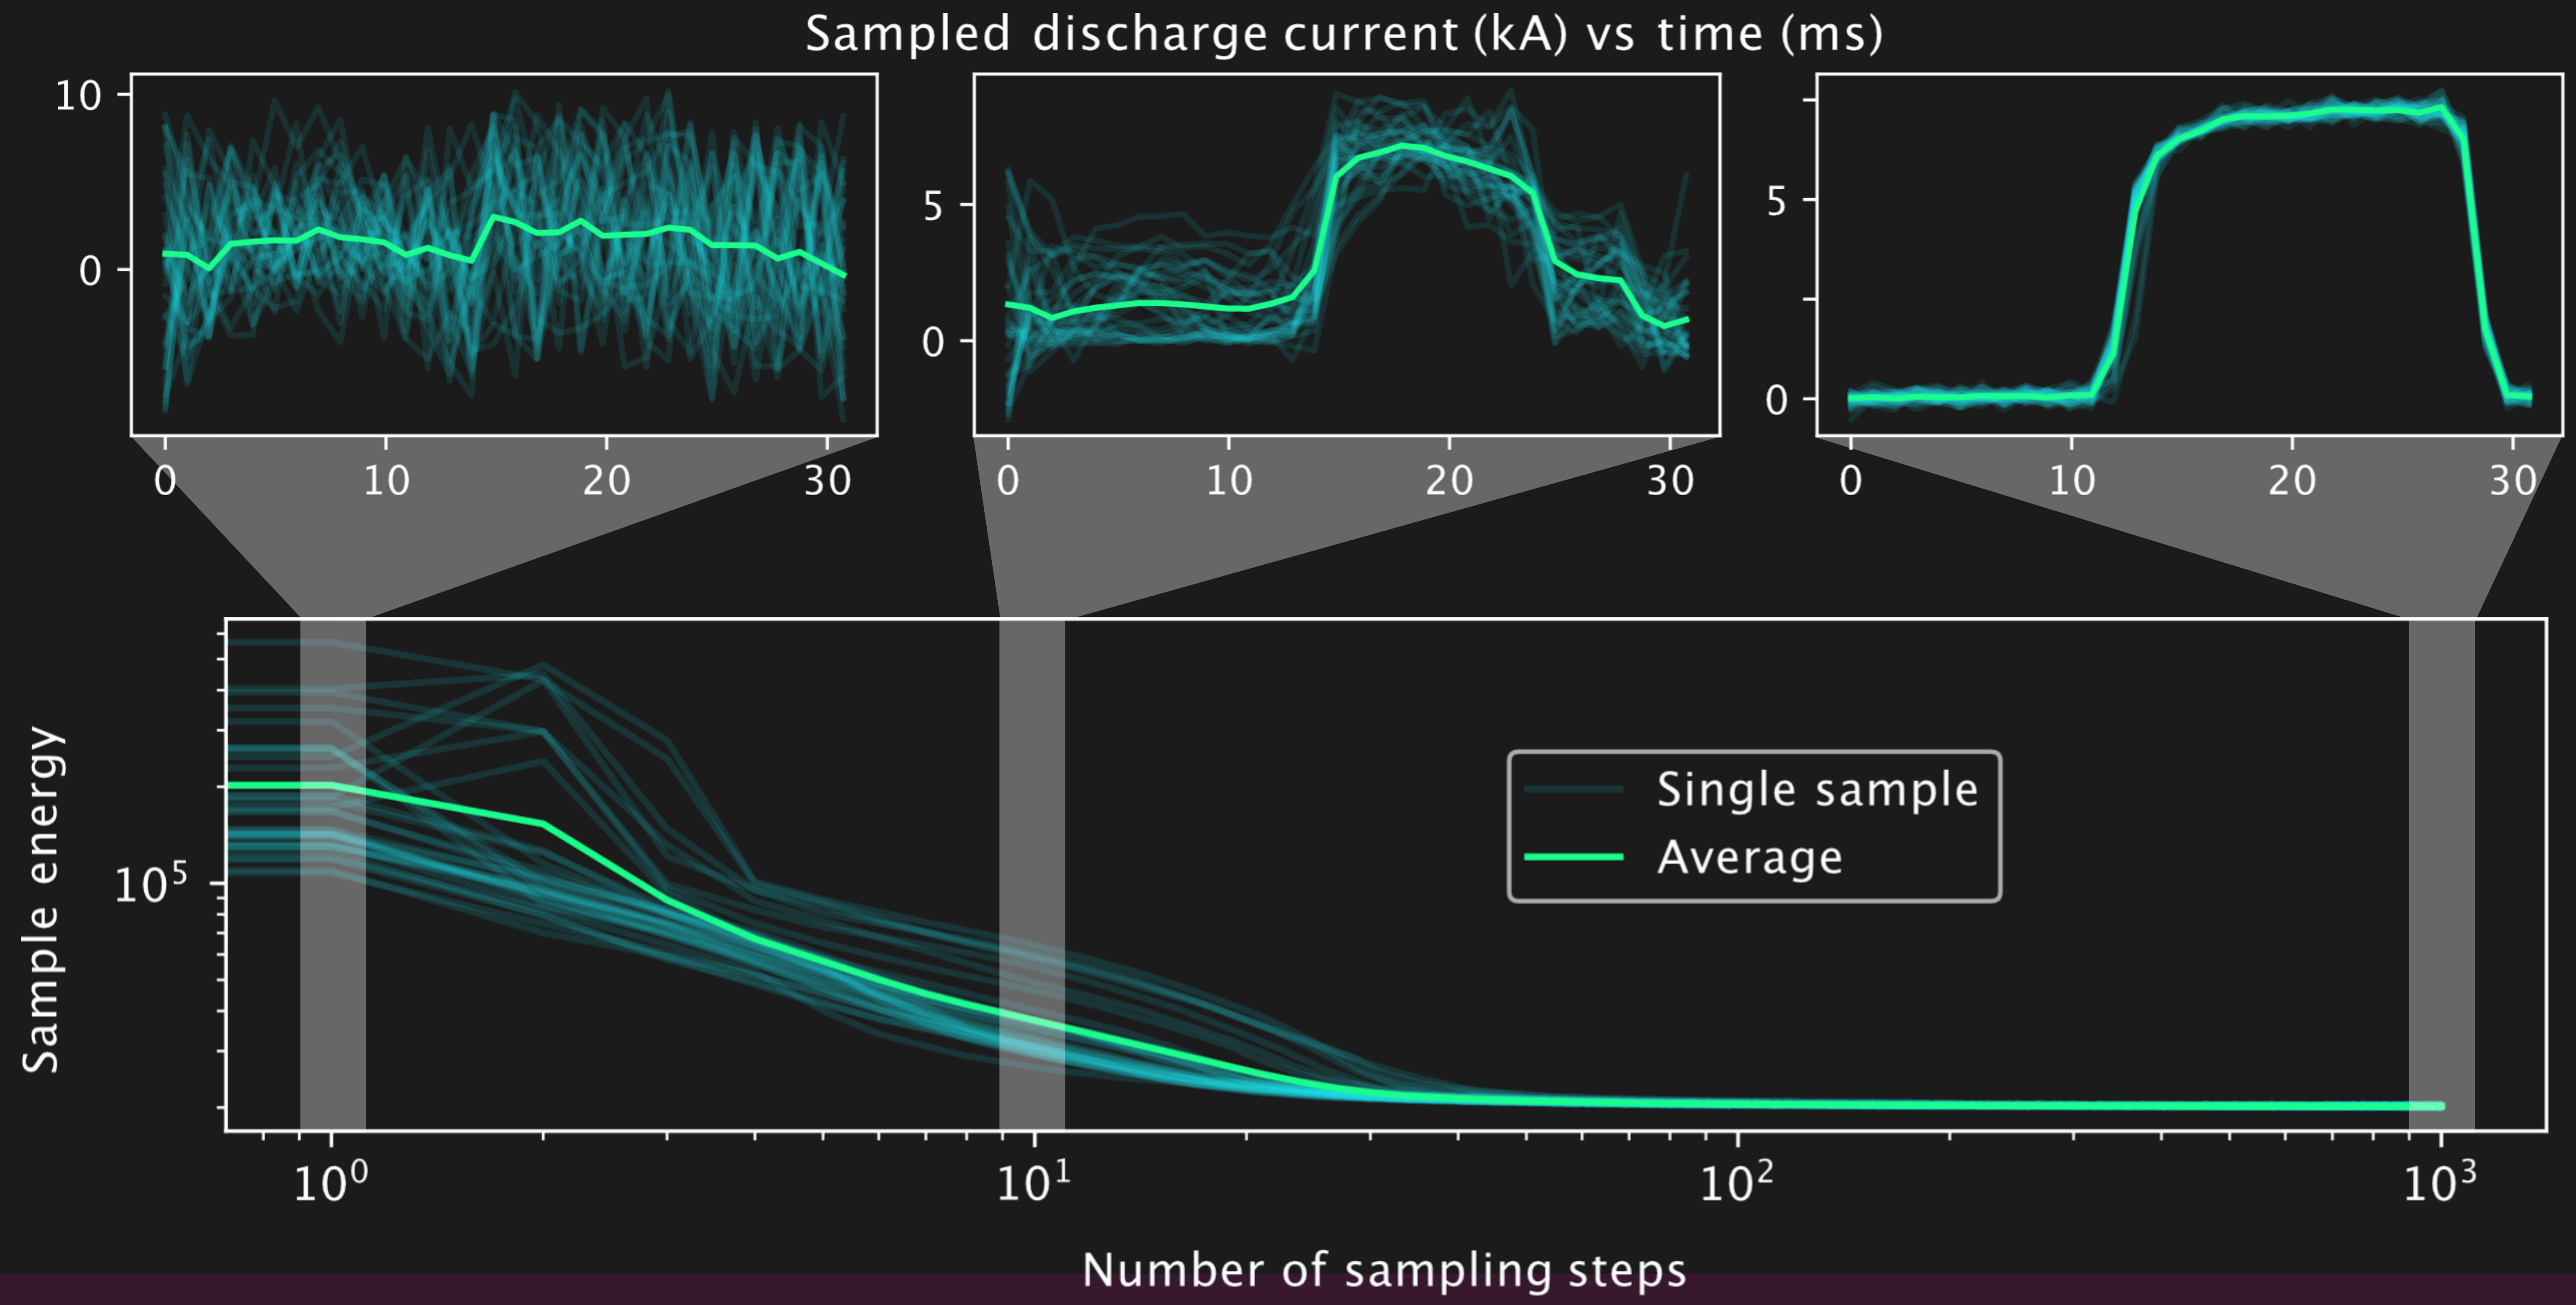

EBM: losses

$$\mathcal{L} = \mathcal{L}_{\mathrm{CD}} + \mathcal{L}_{\mathrm{KL}} + \alpha \mathcal{L}_{\mathrm{reg}}$$

$$\mathcal{L}_{\mathrm{CD}} = \frac{1}{M} \sum_i E_\theta(\tilde{x}_i^+) - E_\theta(\tilde{x}_i^L)$$

$$\mathcal{L}_{\mathrm{KL}} = \frac{1}{M} \sum_i E_{\Omega(\theta)}(E_\theta(\hat{x}_i^K) - \mathrm{NN}(X, \hat{x}_i^K))$$

$$\mathcal{L}_{\mathrm{reg}} = \frac{1}{M} \sum_i E_\theta(\tilde{x}_i^+)^2 + E_\theta(\tilde{x}_i^L)^2$$

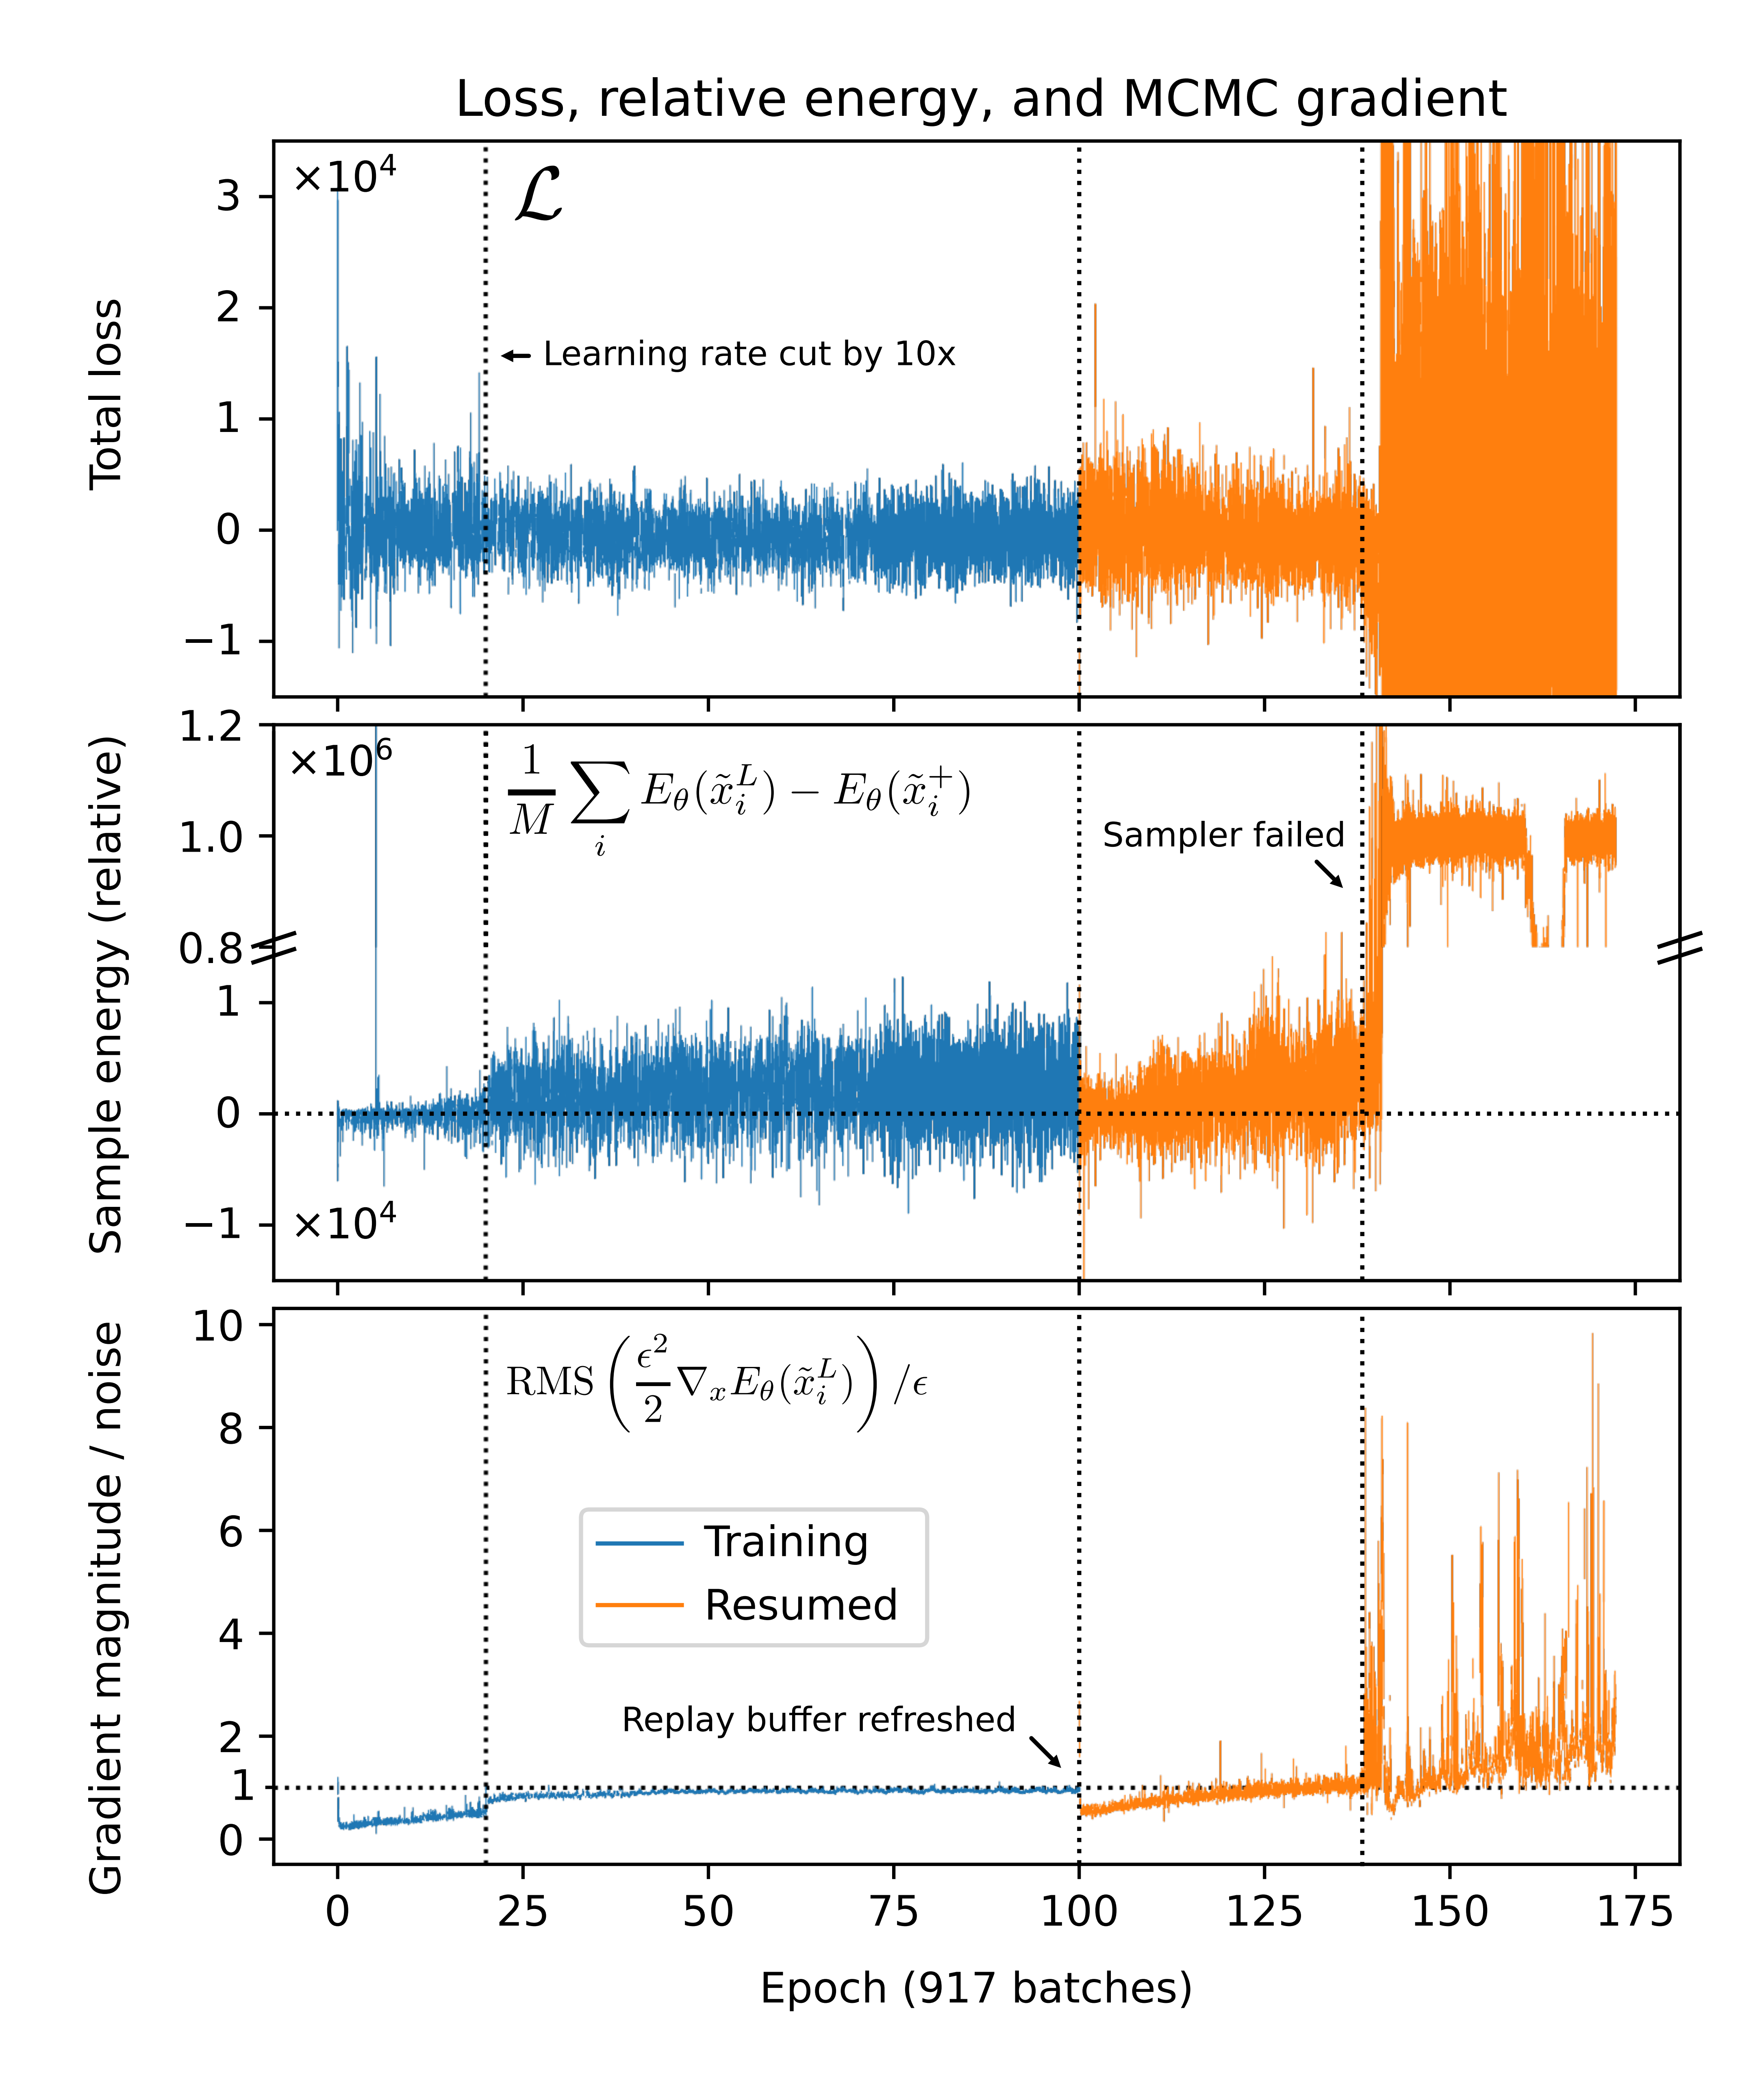

EBM: architecture

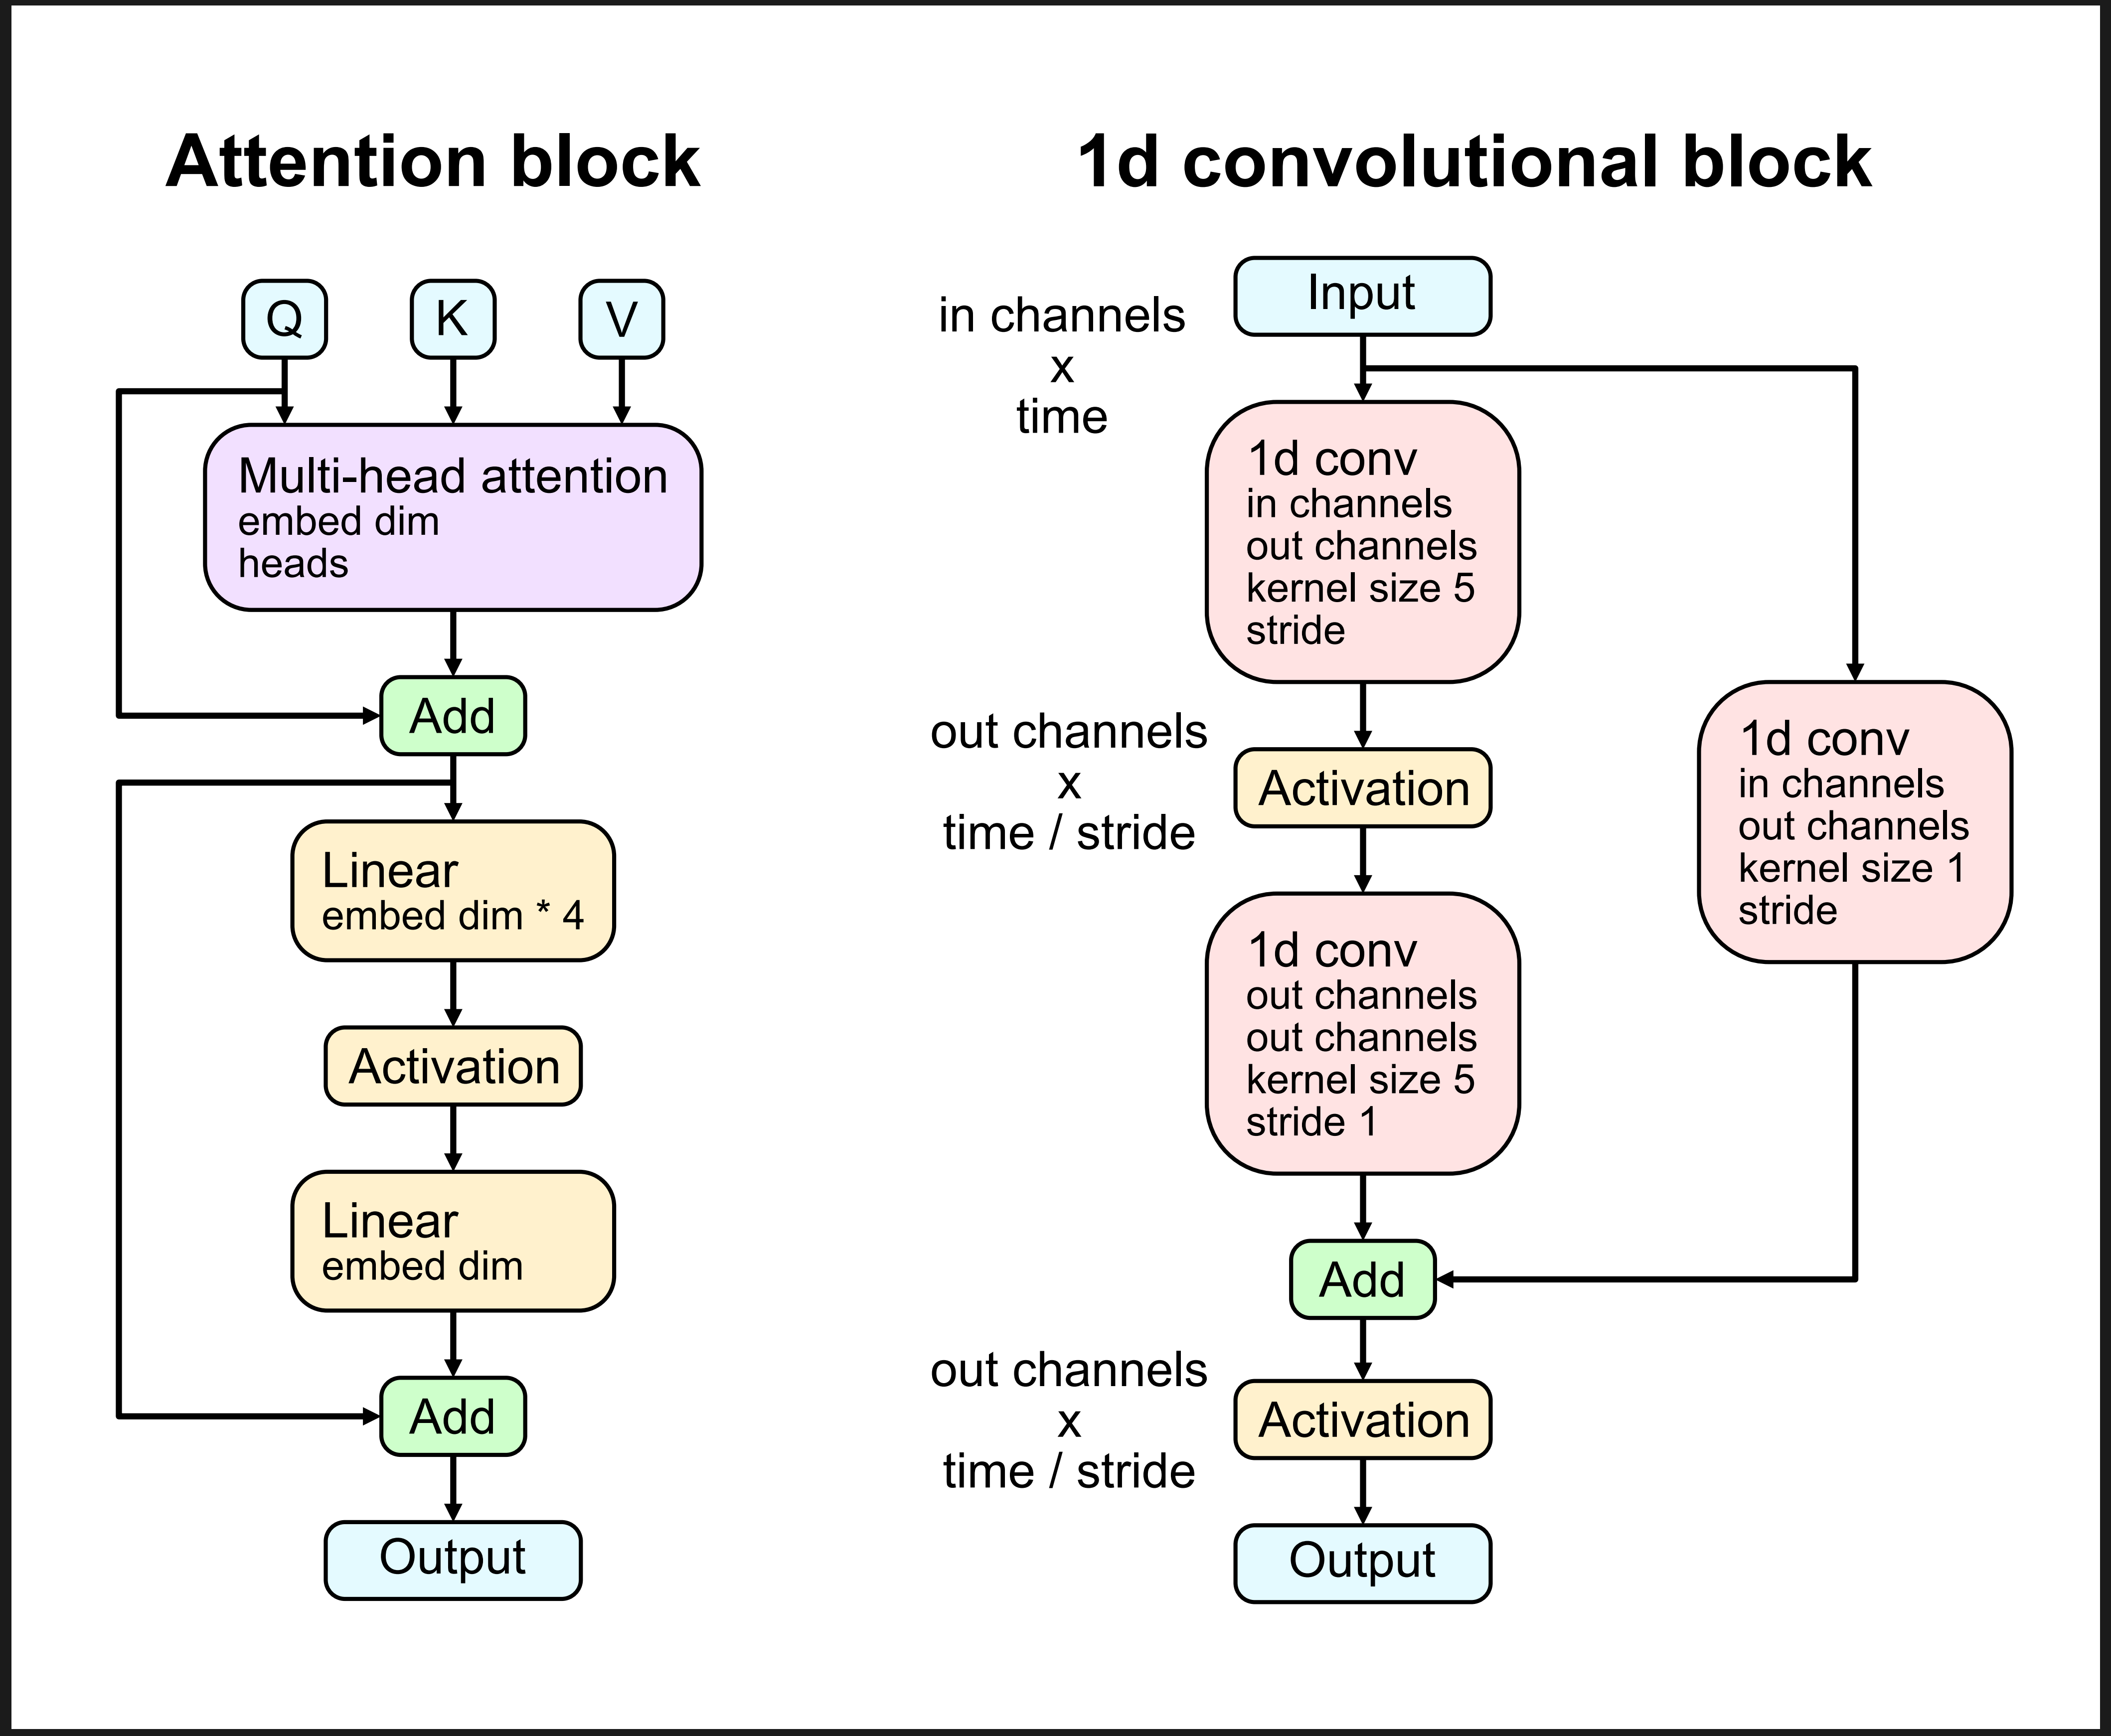

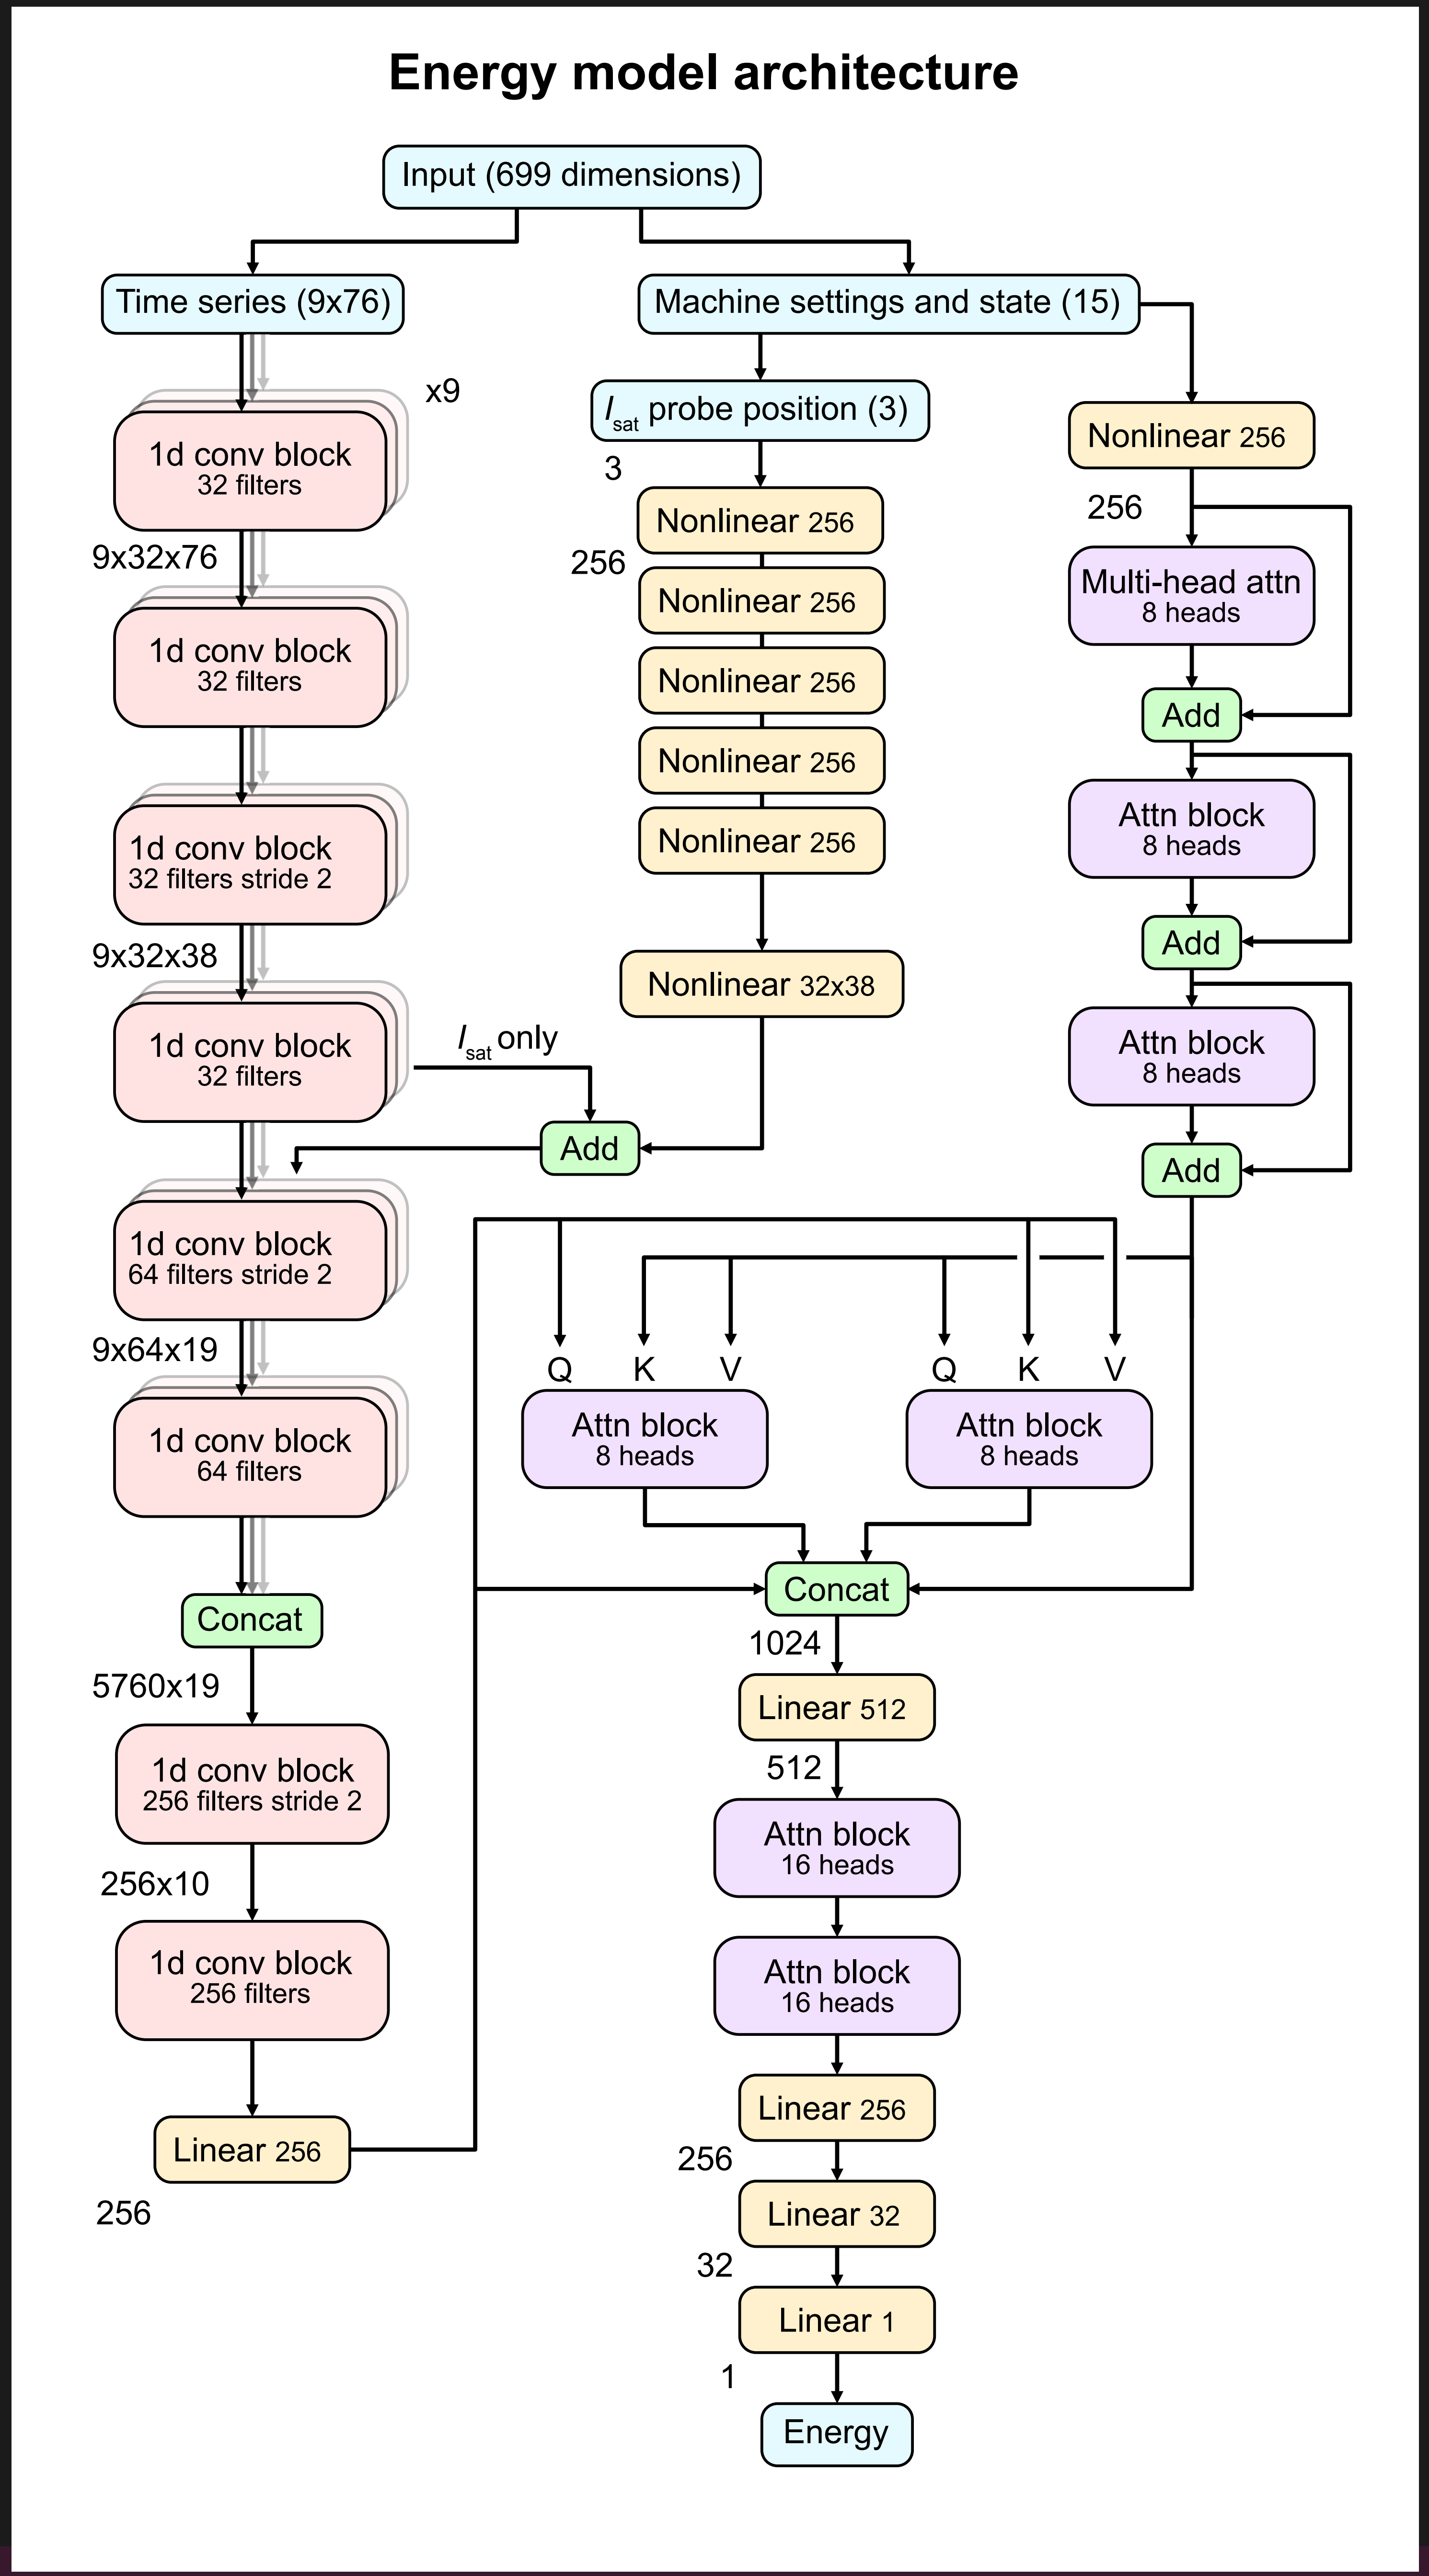

EBM: unconditional samples

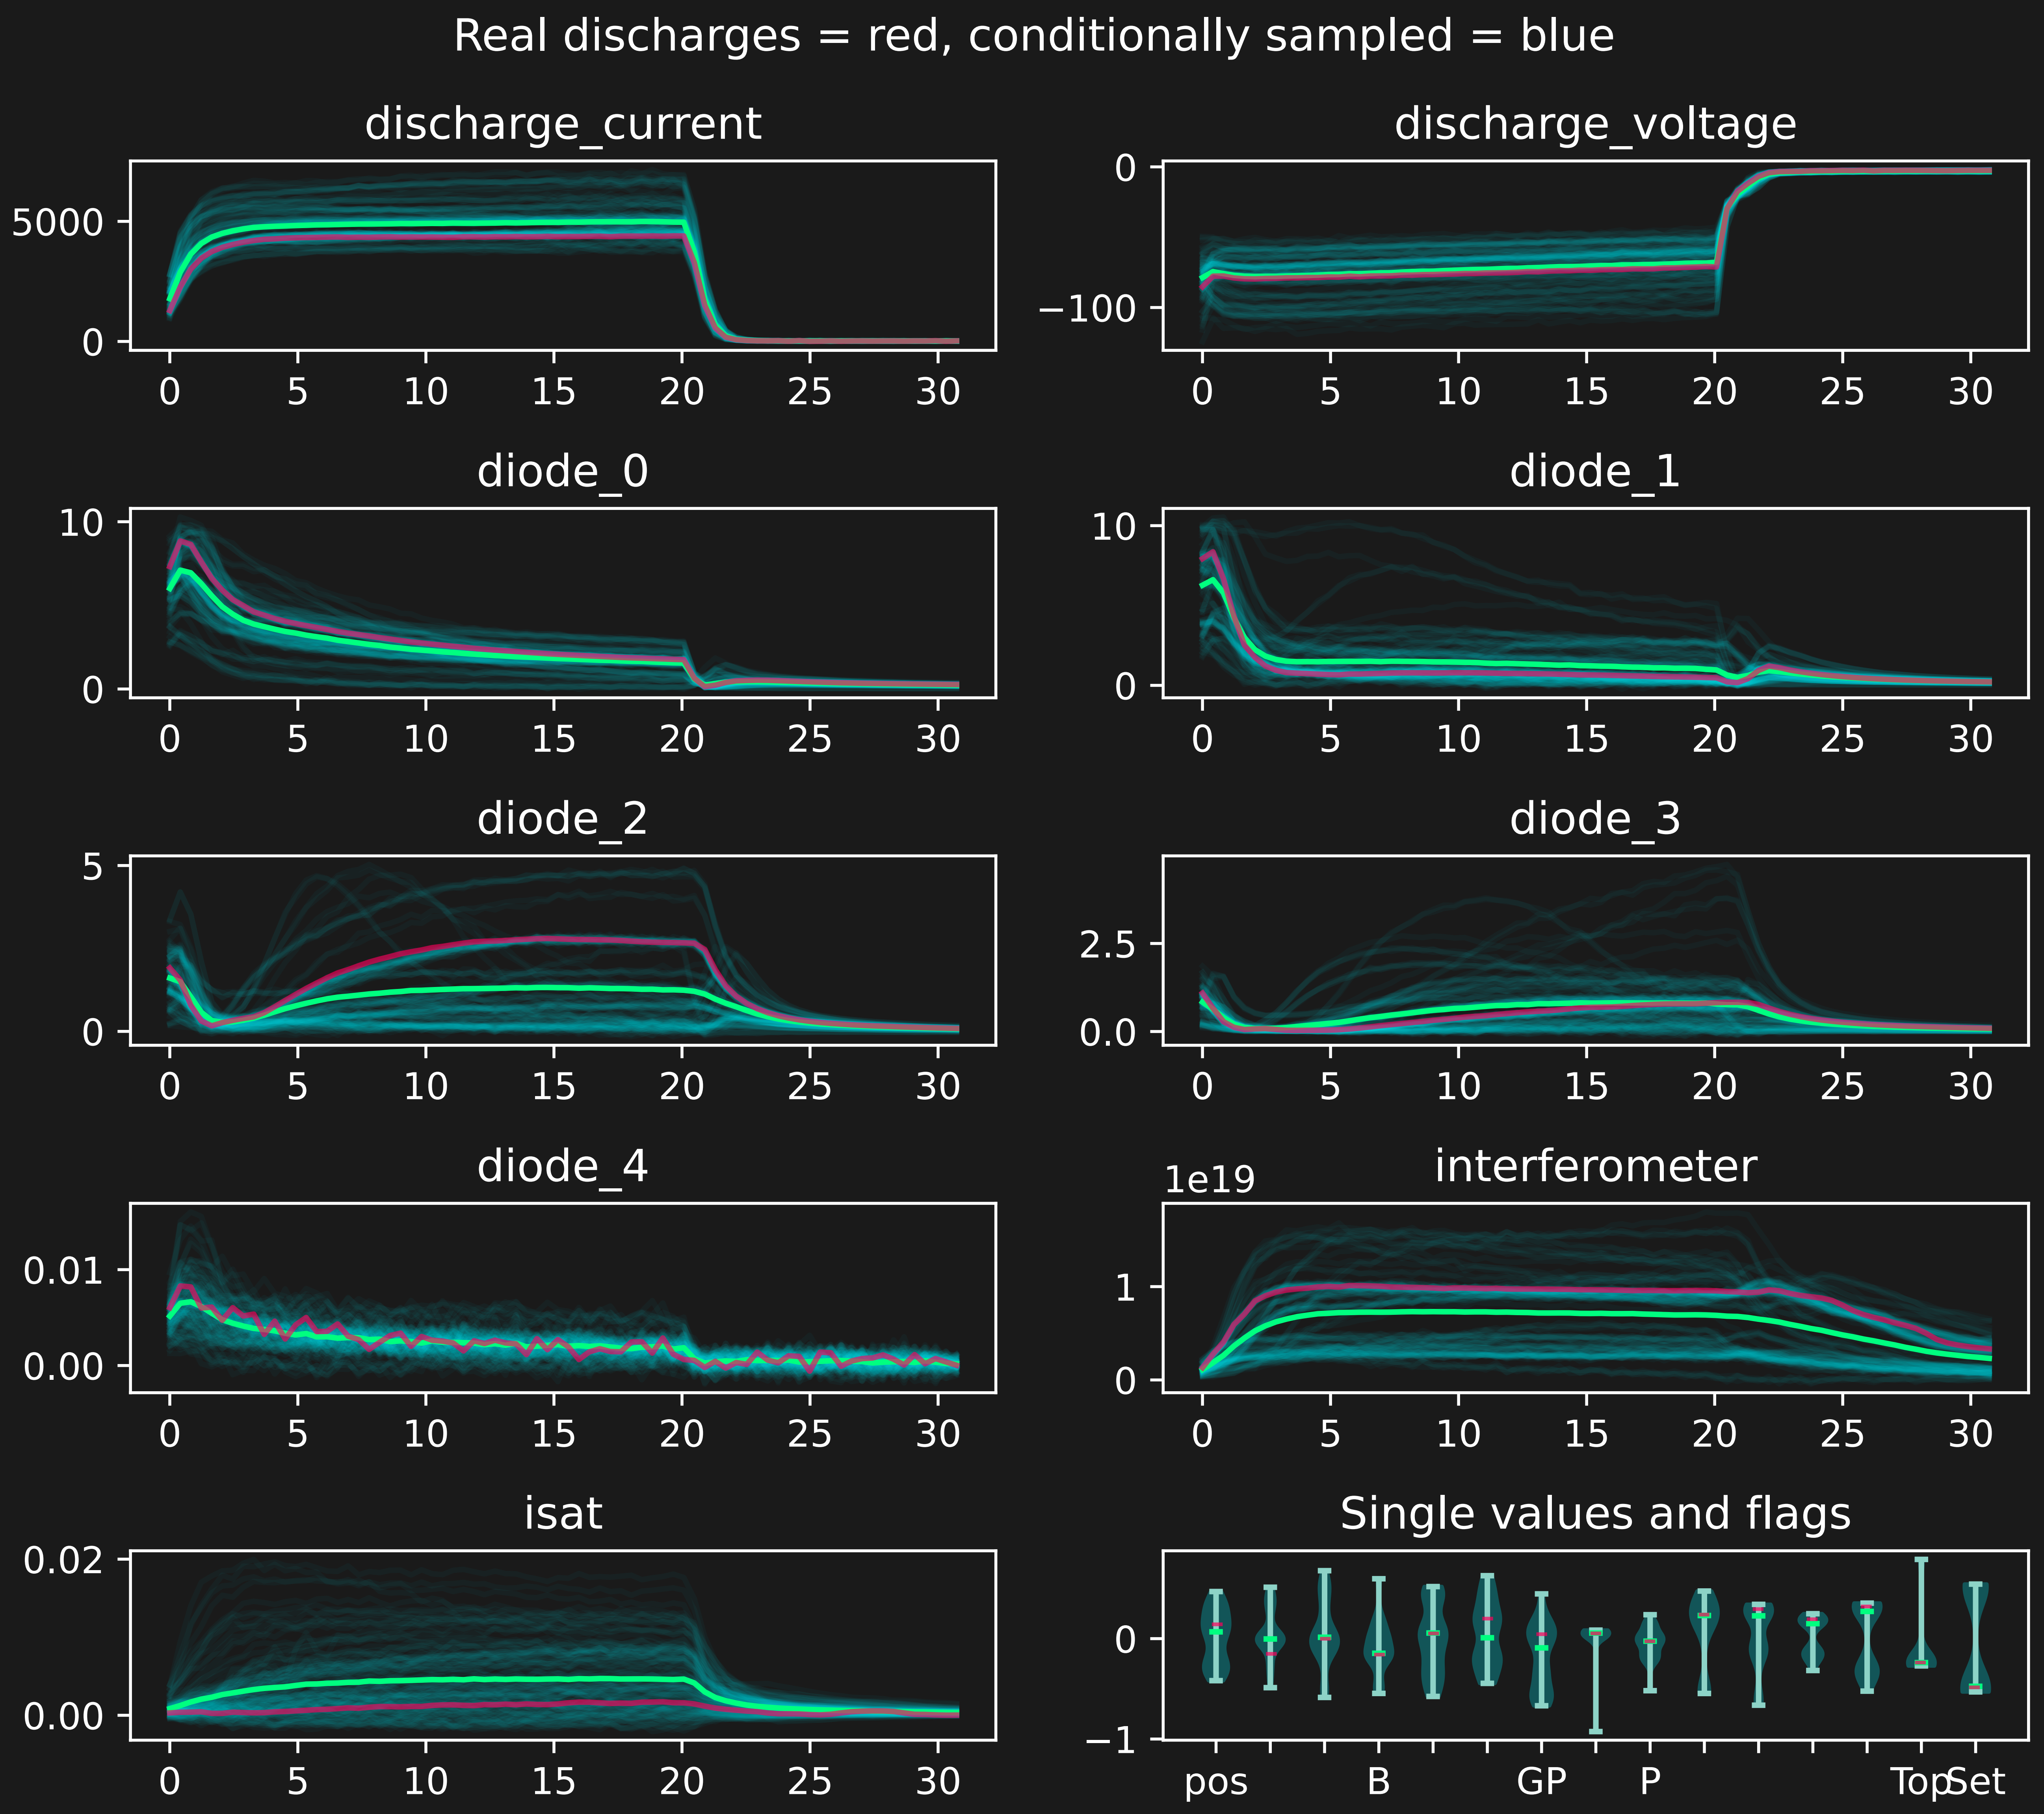

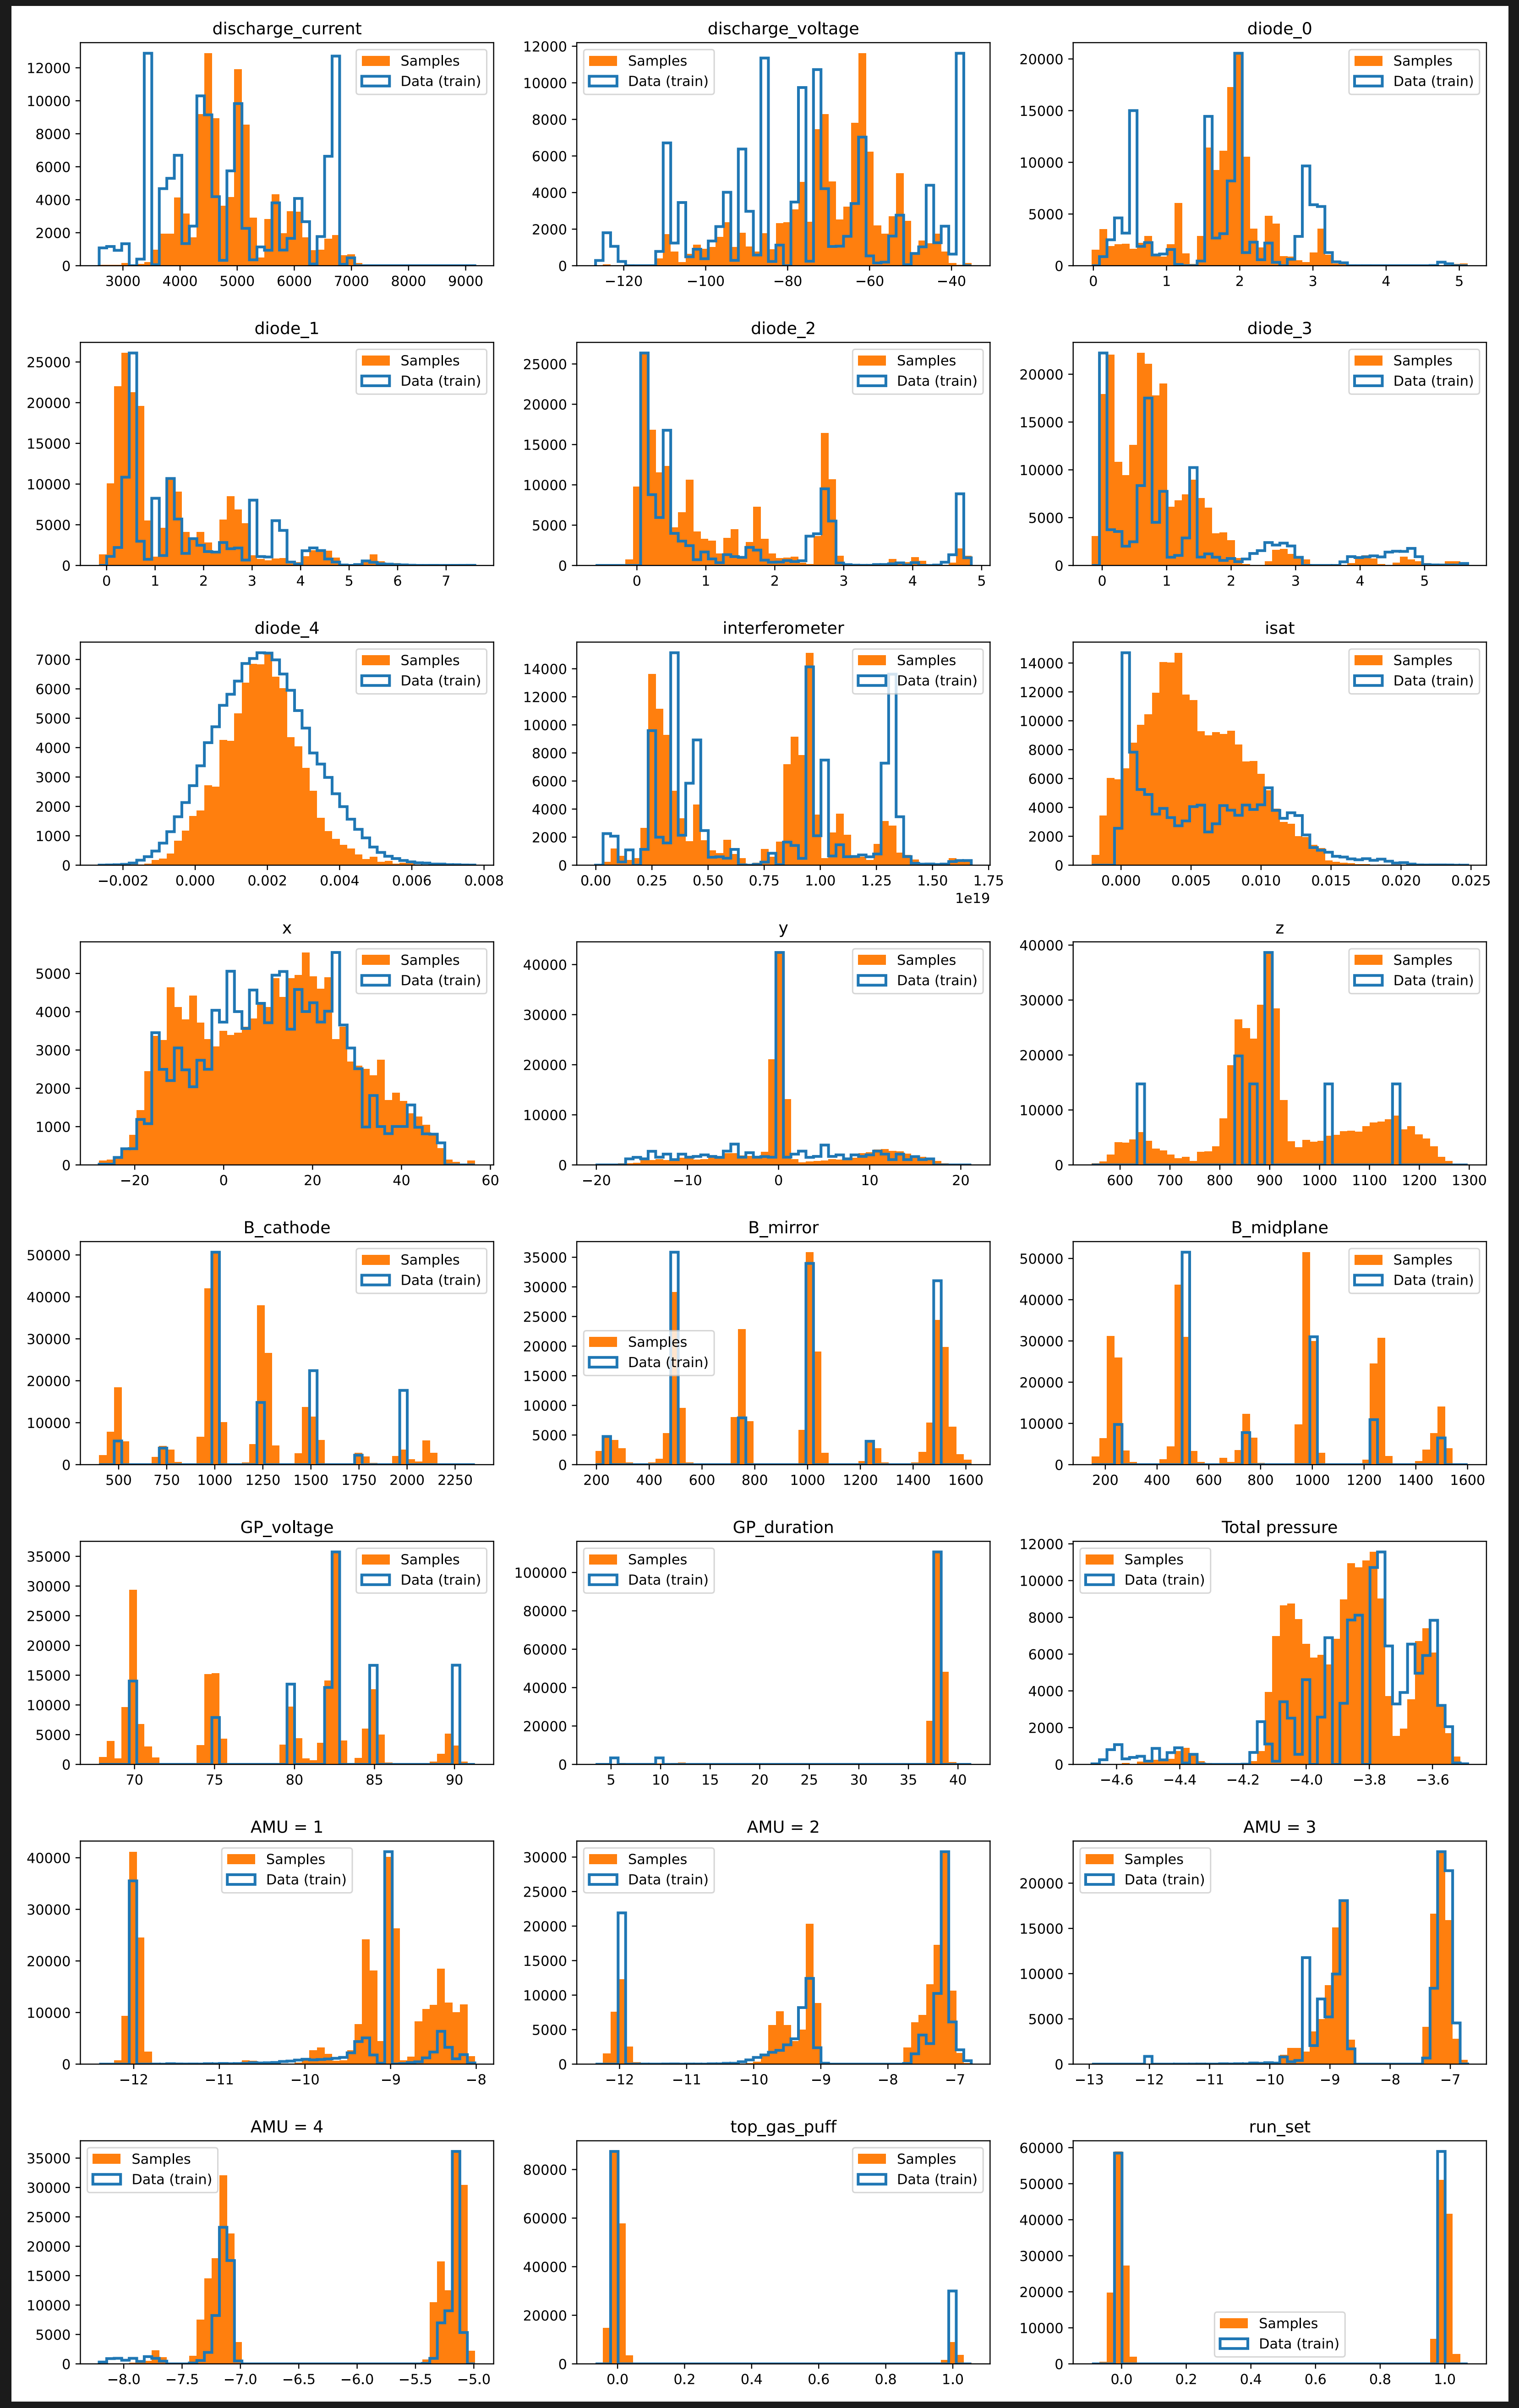

NNs are repeated matrix multiplication and a nonlinearity

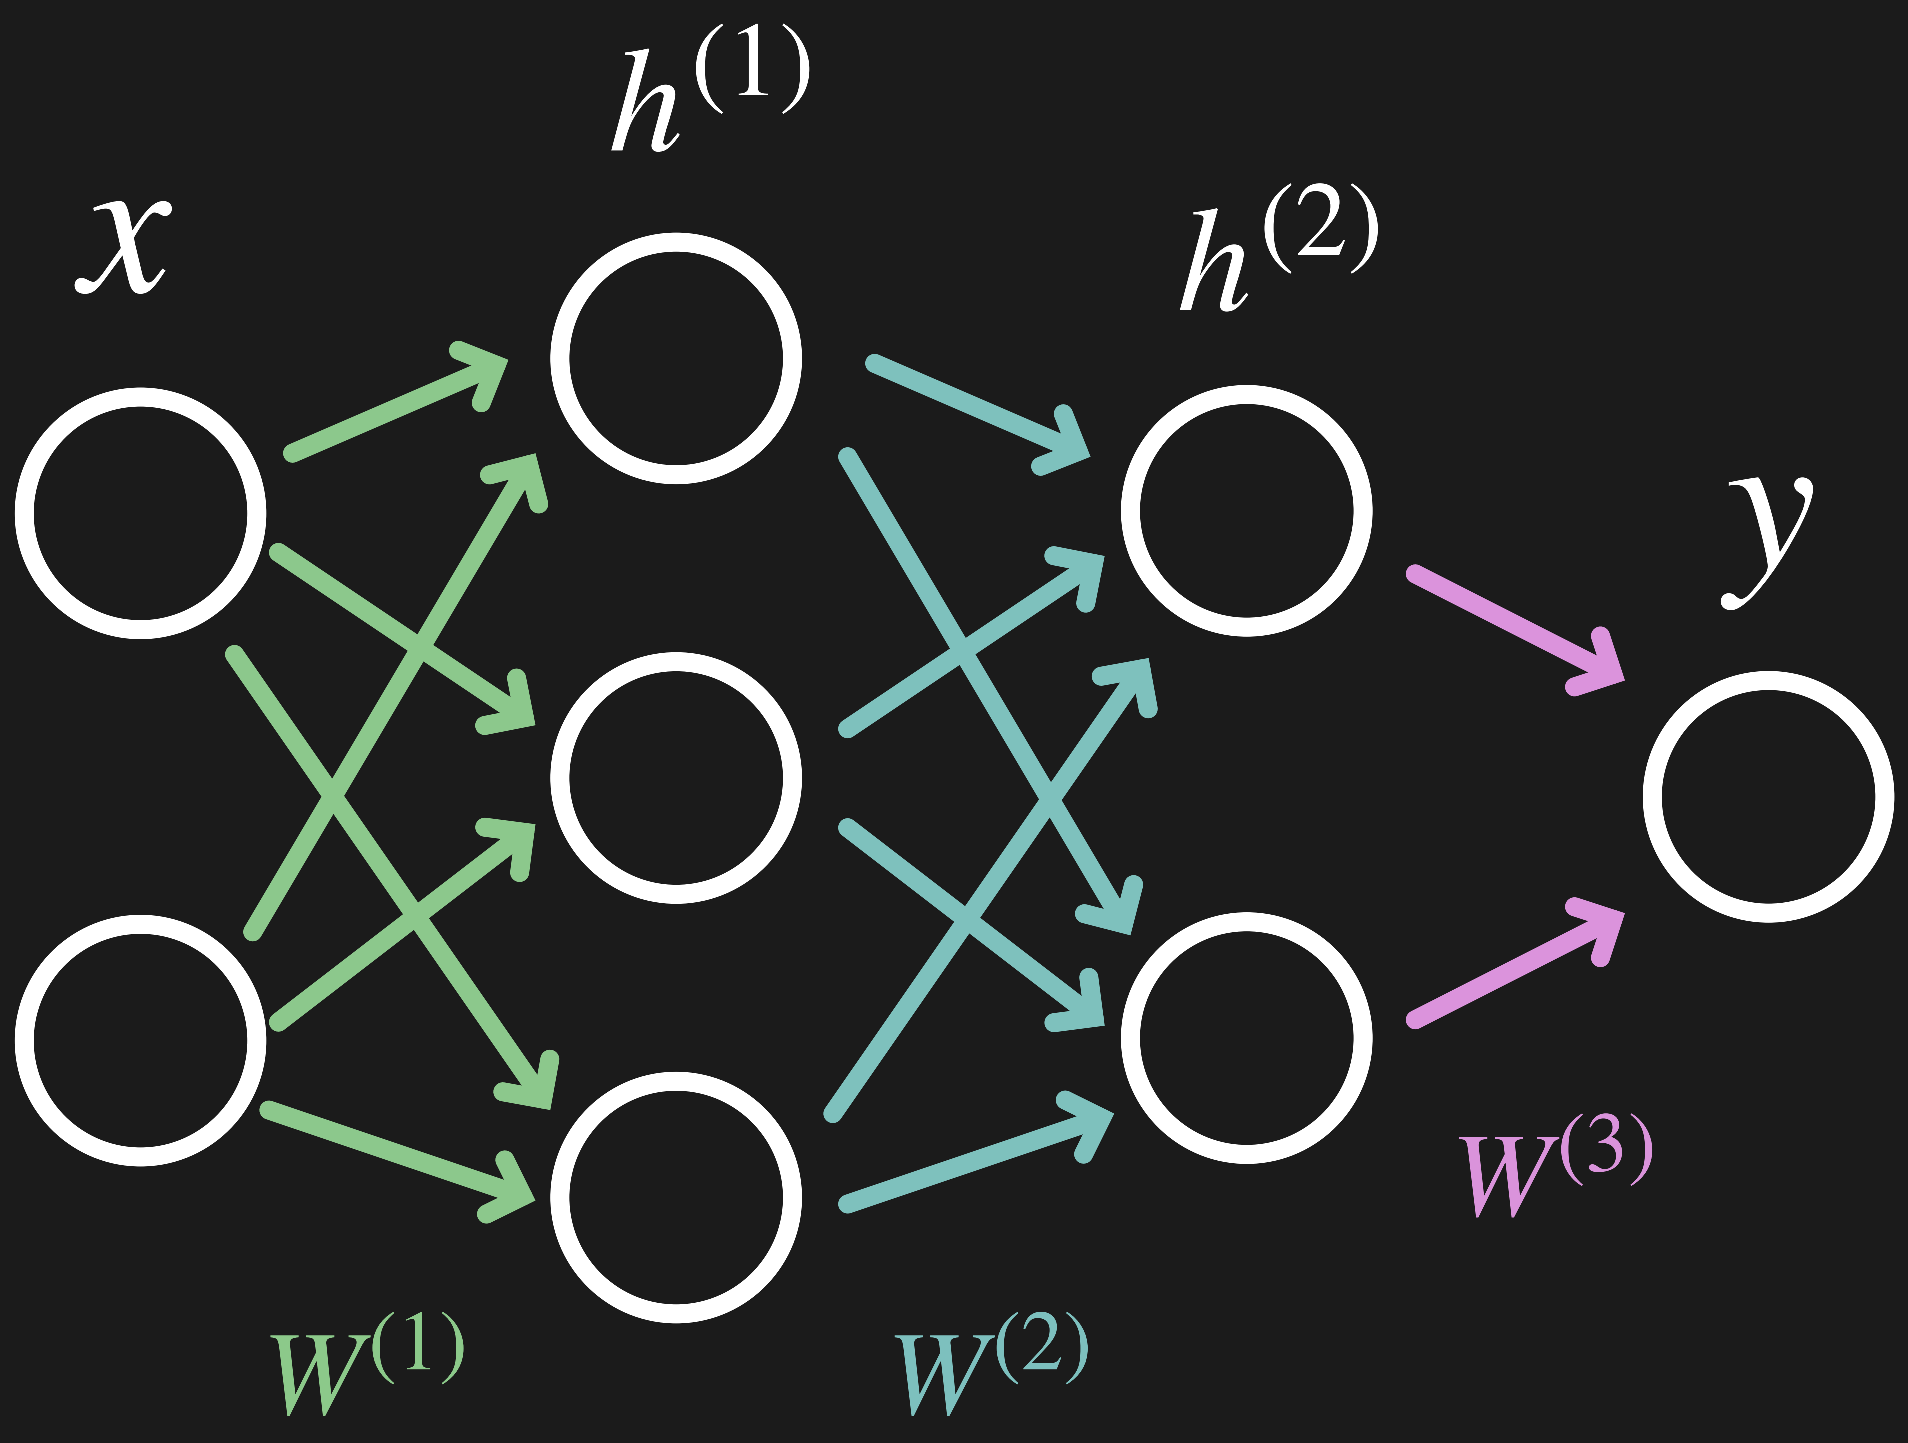

Neural network "layer"

$$h^{(1)} = g^{(1)}\left(W^{(1)}x + b^{(1)}\right)$$

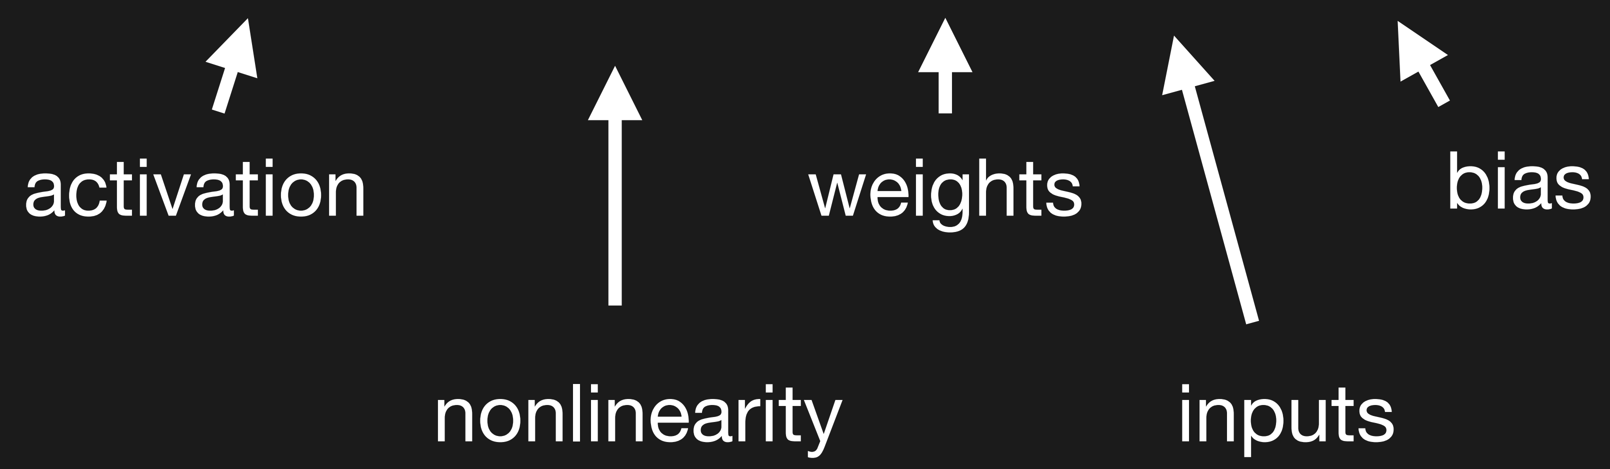

Input x: our machine configuration and probe location

Output y: time-averaged Isat value

and that is how you make sand think

NNs are trained via gradient descent over some loss function

- Update values based on gradient (I use AdamW)
  - $\vec{x} := \vec{x} - \nabla_{\vec{x}}\mathscr{C} \cdot \lambda$

cost function   step size

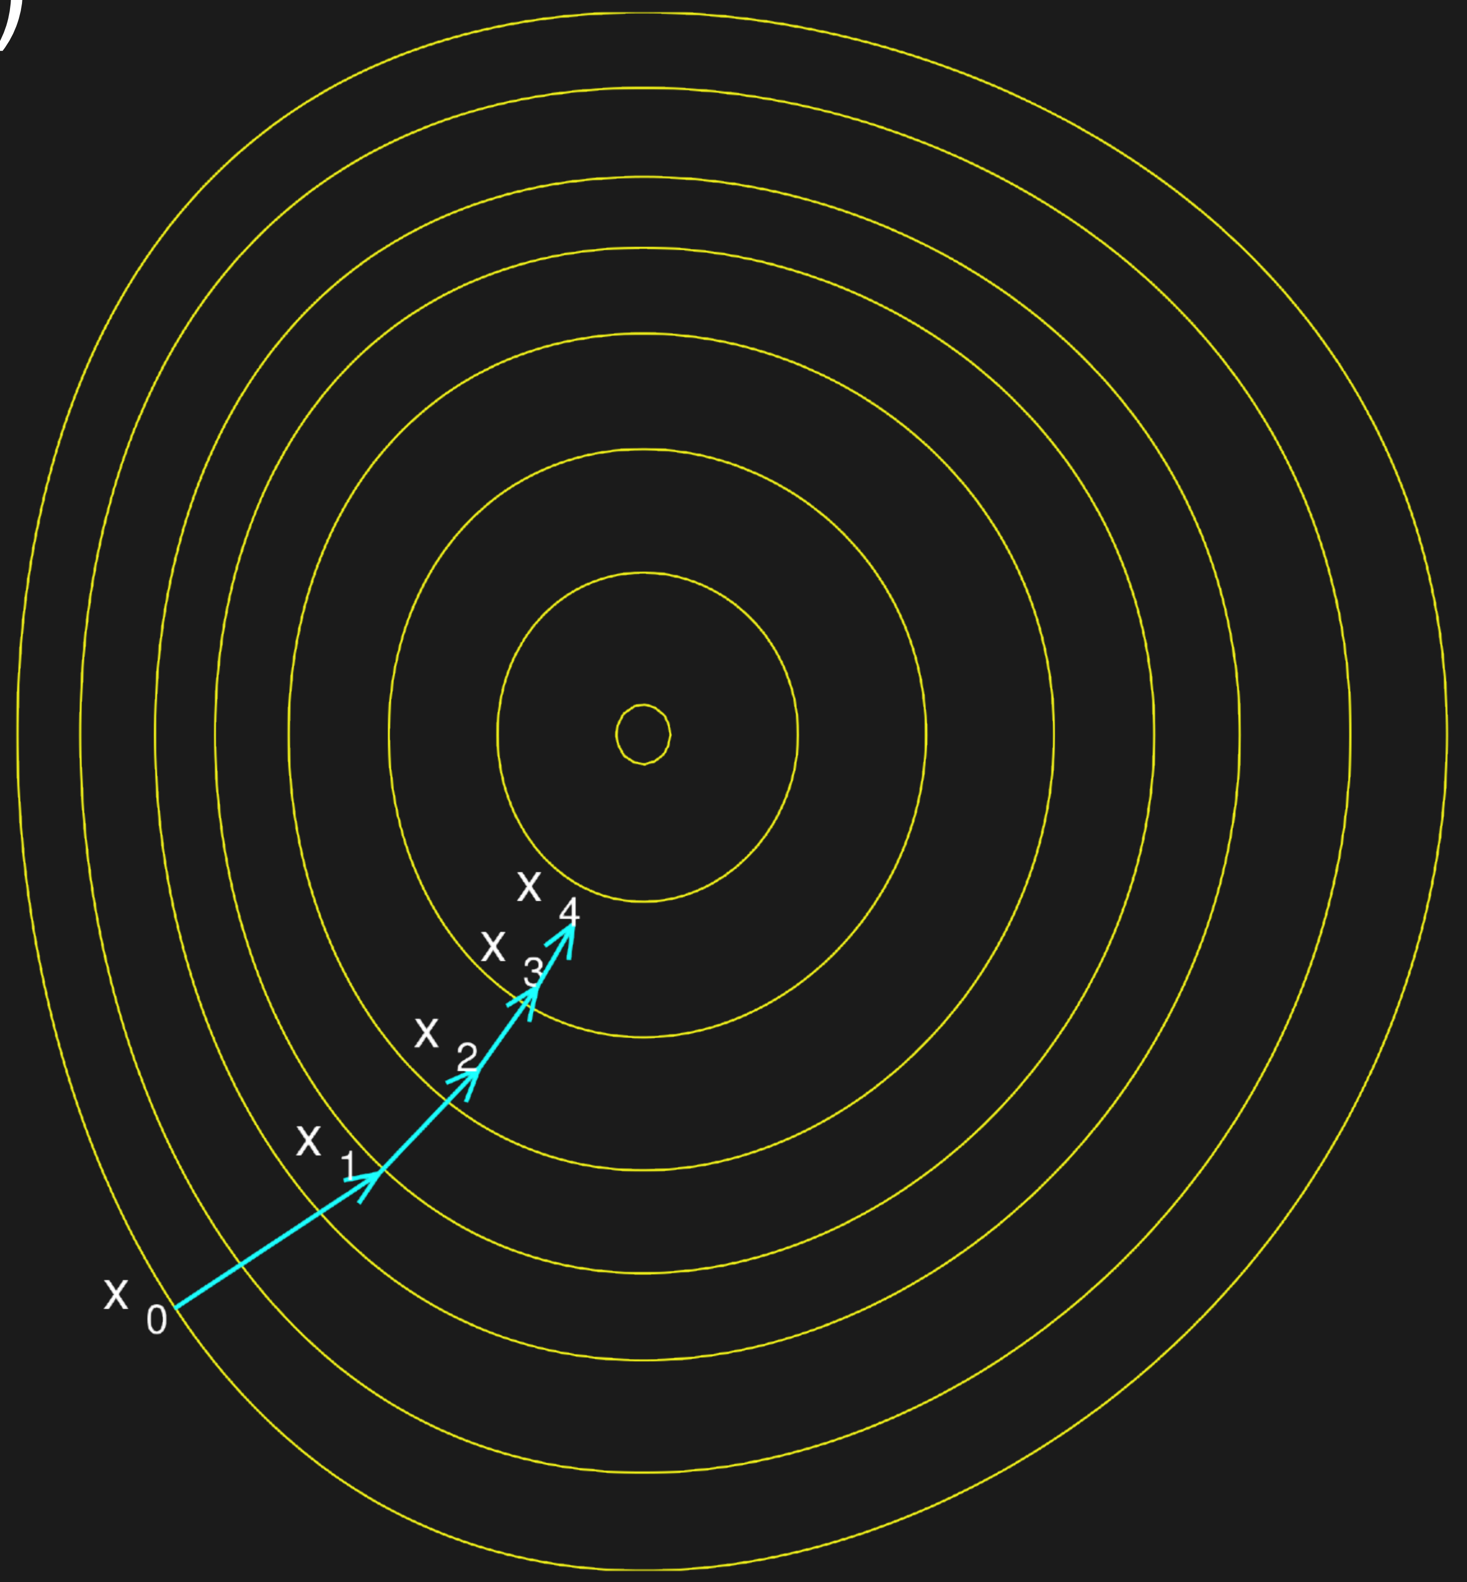

$$\mathscr{L}_{\textbf{MSE}} = \frac{1}{m}\sum_{i=1}^{m}\left(f(x_i) - y_i\right)^2$$

$$\mathscr{L}_{\beta\textbf{–NLL}} = \frac{1}{2}\left(\log \sigma_i^2(\mathbf{x}_n) + \frac{\left(\mu_i(\mathbf{x}_n) - y_n\right)^2}{\sigma_i^2(\mathbf{x}_n)}\right)\textbf{StopGrad}\left(\sigma_i^{2\beta}\right)$$