Variability and risk of edge-localized modes at JET using machine learning

J. Alhage, C. Haems, M. Van Damme, Y. Zhang, G. Verdoolaege,
JET Contributors, and the EUROfusion Tokamak Exploitation Team

6th IAEA Technical Meeting on Fusion Data Processing, Validation and Analysis
September 2025, Shanghai, China

GHENT UNIVERSITY

infusion

EUROfusion

Edge-localized modes

- Type of instability
- Repeating releases of energy and particles from plasma edge.
- Effect is intense heat flux:
  - Duration < 10ms
  - Waiting time < 30ms
  - Energy 2-10% plasma
- Must be regulated to prevent damage to PFC

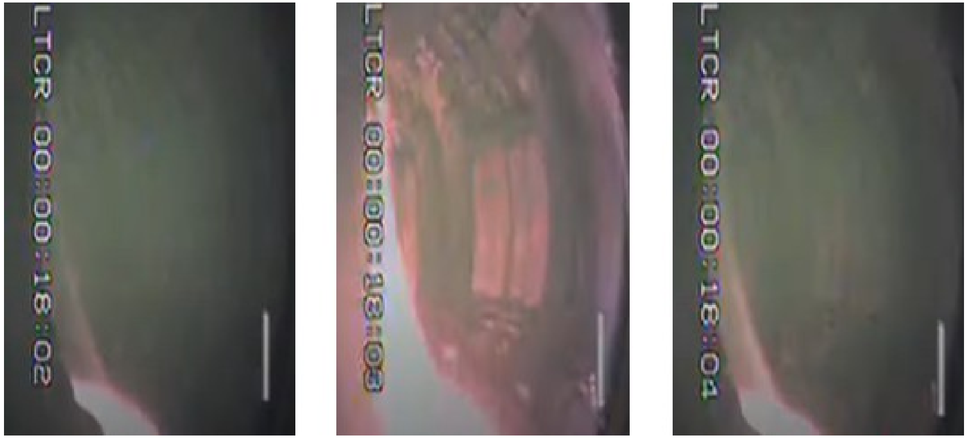

Edge-localized modes

- Behavior varies from pulse to pulse (frequency, energy)
- Depending on experiment conditions

- Motivation:
- Which, and how, machine and plasma conditions affect ELMs
- Data driven approach to understand the full variability and find areas of safe operation

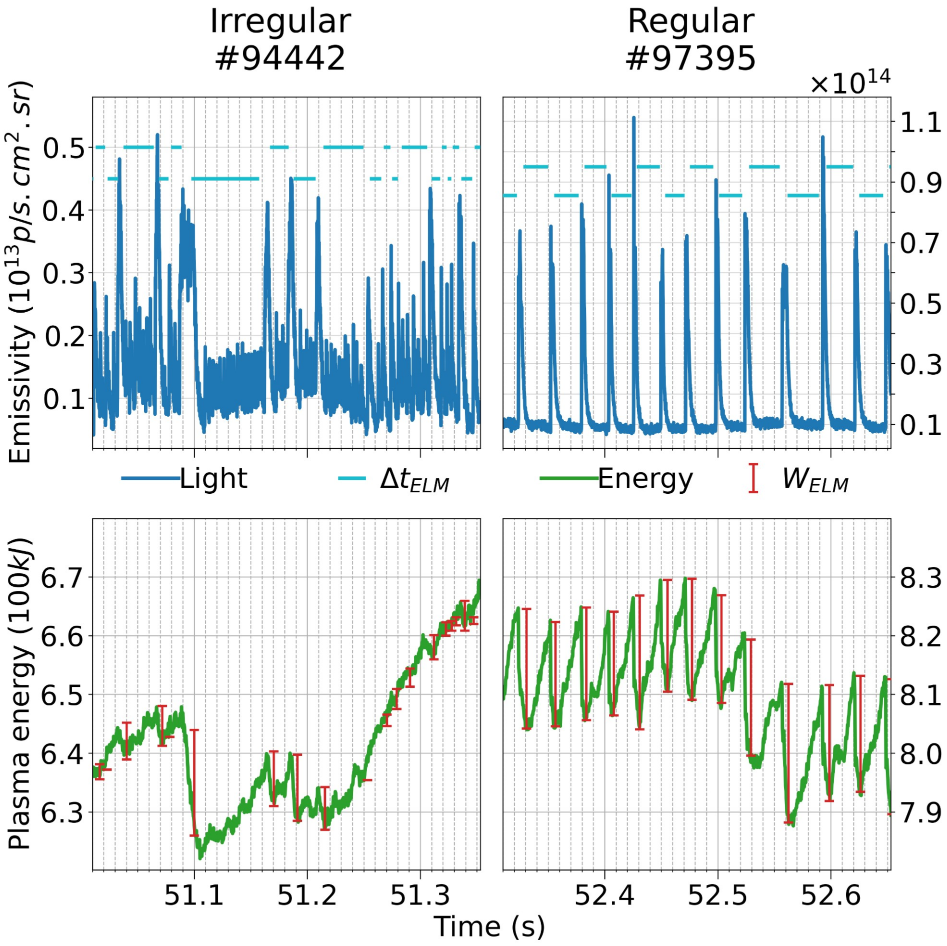

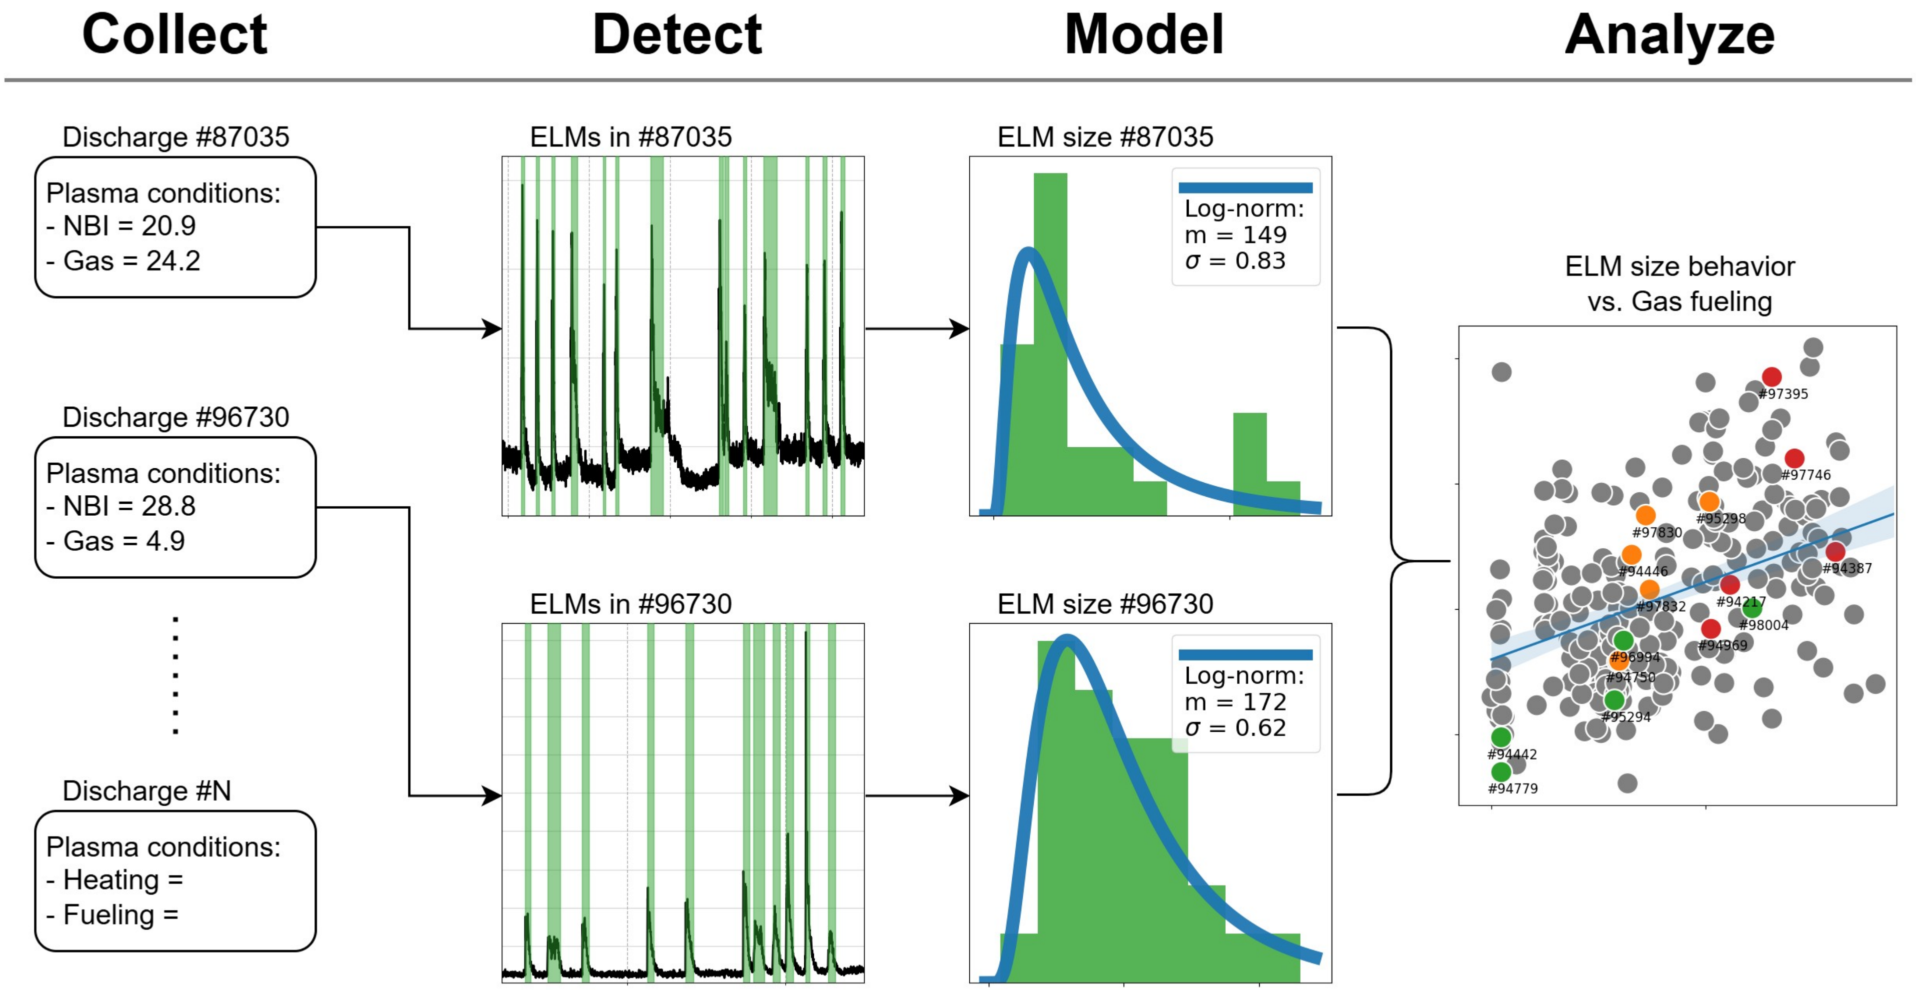
Collect
Detect
Model
Analyze
Discharge #87035
Plasma conditions:
- NBI = 20.9
- Gas = 24.2
Discharge #96730
Plasma conditions:
- NBI = 28.8
- Gas = 4.9
Discharge #N
Plasma conditions:
- Heating =
- Fueling =
ELMs in #87035
ELMs in #96730
ELM size #87035
Log-norm:
m = 149
σ = 0.83
ELM size #96730
Log-norm:
m = 172
σ = 0.62
ELM size behavior
vs. Gas fueling
#97395
#97746
#95298
#97830
#94446
#94387
#97832
#94217
#98004
#94969
#96994
#94750
#95294
#94442
#94779

Data collection

13825 ELMs manually marked for ML training

214 shots, 270 windows
JET campaign 2019-2021

Process PDB-ILW[1] data with advanced detection

2000+ windows, 160000+ ELMs from JET 2012-2020

- Operating conditions or major controls:
  I_p, B_t, P_{NBI}, P_{ICRH}, gas, pellets, plasma shape, q95
- ELM behavior:
  - Timing: BeII emissivity
  - Size (proxy): stored energy W

[1] L. Frassinetti et al 2021 Nucl. Fusion 61 016001

ELM detection

- Differentiate between events and non-events in timeseries
- Not only peaks, but regions with a specific signature

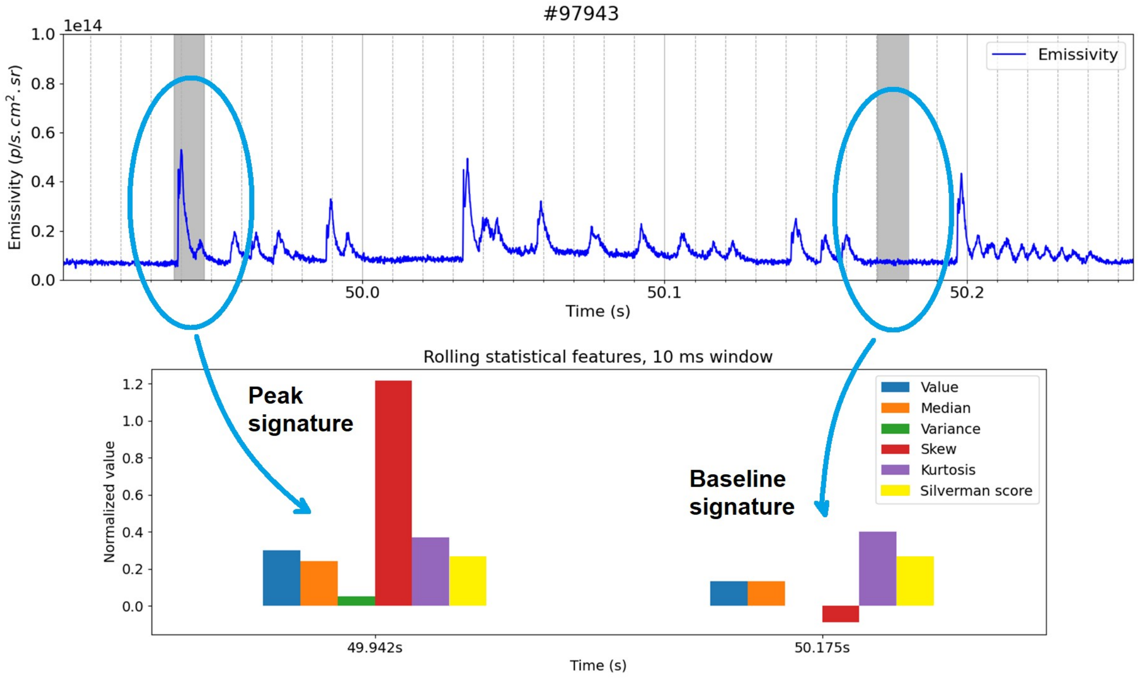

Advanced methods

- Laplacian of Gaussian filter + threshold (2^{nd} order derivative and smoothing) using kernel optimized for ELM shape

- Recurrent neural network (3 x GRU)

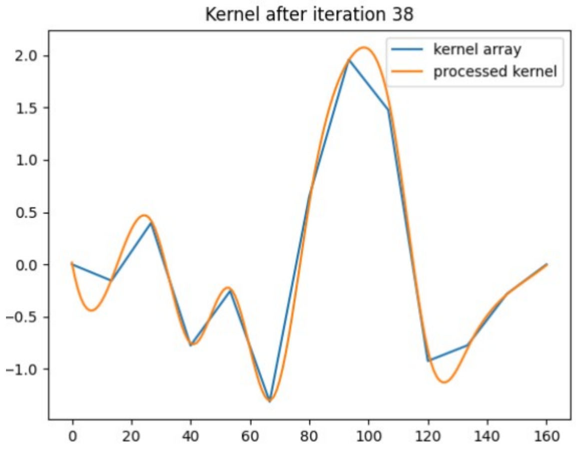

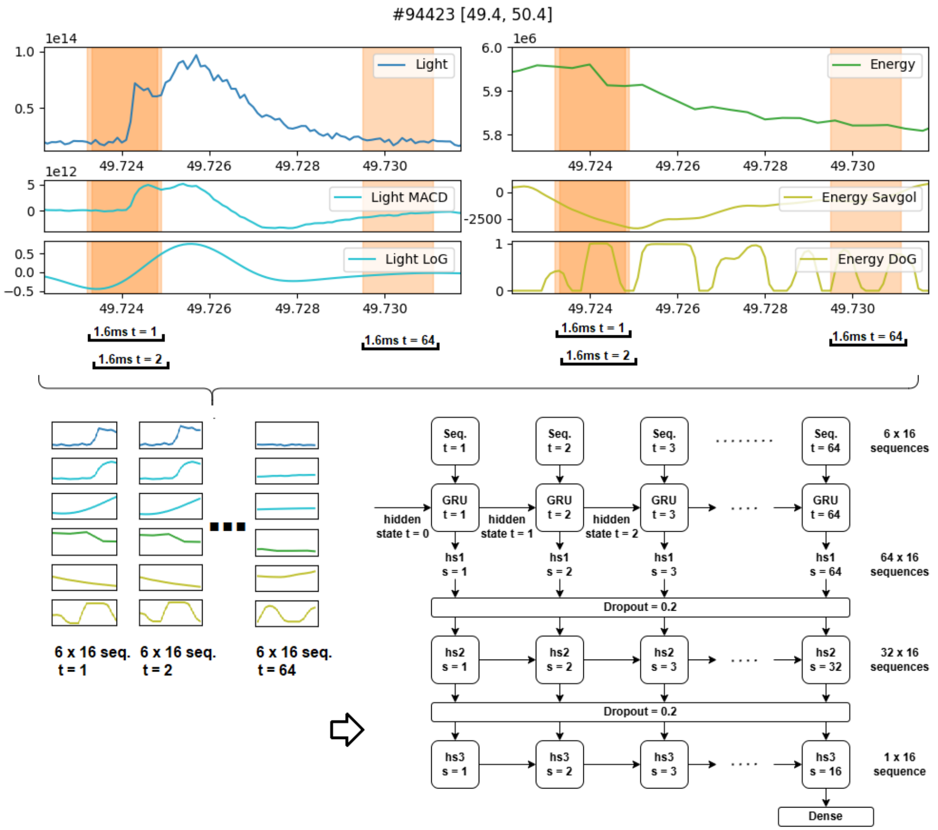

Training and evaluation

- Split data into train/test/validation
- Quantify a correct detection (TP/FP/FN)
- Implement, fine tune, and evaluate

- ELM detection pipeline with master students

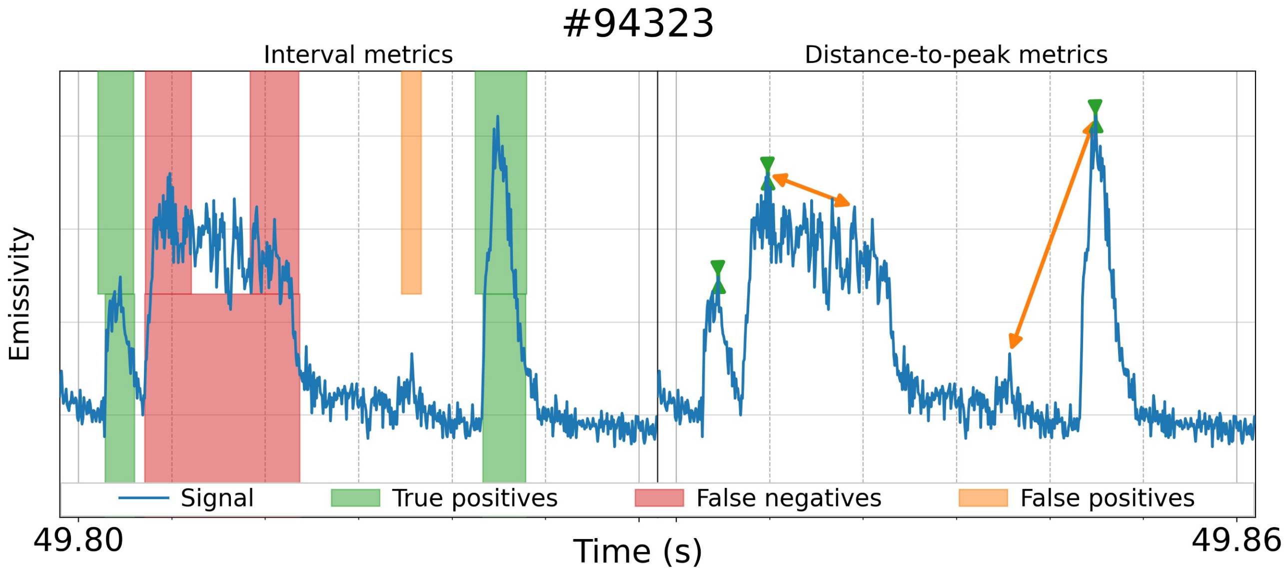

Detection performance

- Good performance
- Tested on other machines

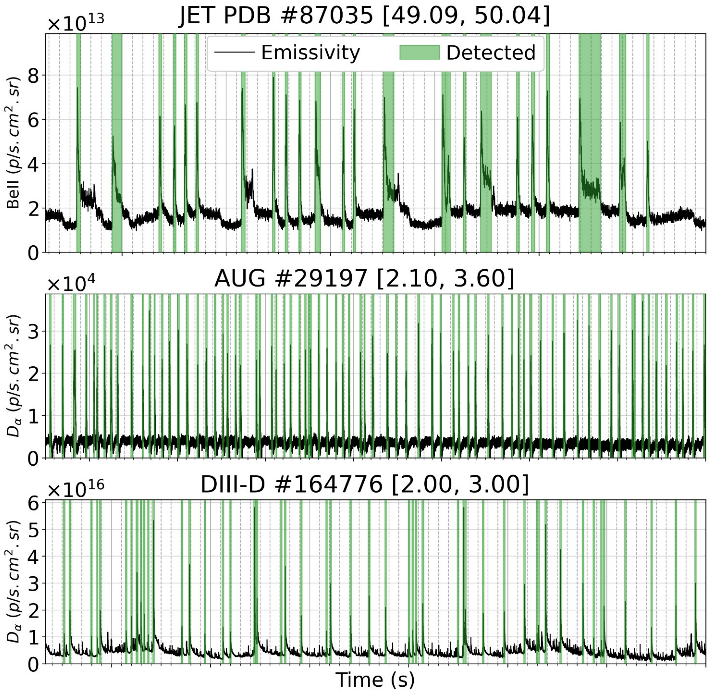

Method	TP (%)	FP (%)
RT	34.41	17.53
MACD	87.17	21.63
Rolling Z-score	89.99	16.39
LoG	86.61	12.79
LoG with doG	80.78	10.91
LoG custom kernel	82.97	9.39
Deconvolution	78.14	70.30
CNN	80.01	16.63
CNN min. filter	73.64	15.64
RNN 64,32,16,1	91.95	15.90
RNN 32,BN,1	85.20	15.69
RNN 8,1	70.66	13.33

Data summary

- Eliminate outliers
- Natural ELMs
- Remains 90000+ ELMs
  from 1388 windows
  ~ 1.5s duration
  ~ 66 ELMs per sample
- Varied operating conditions

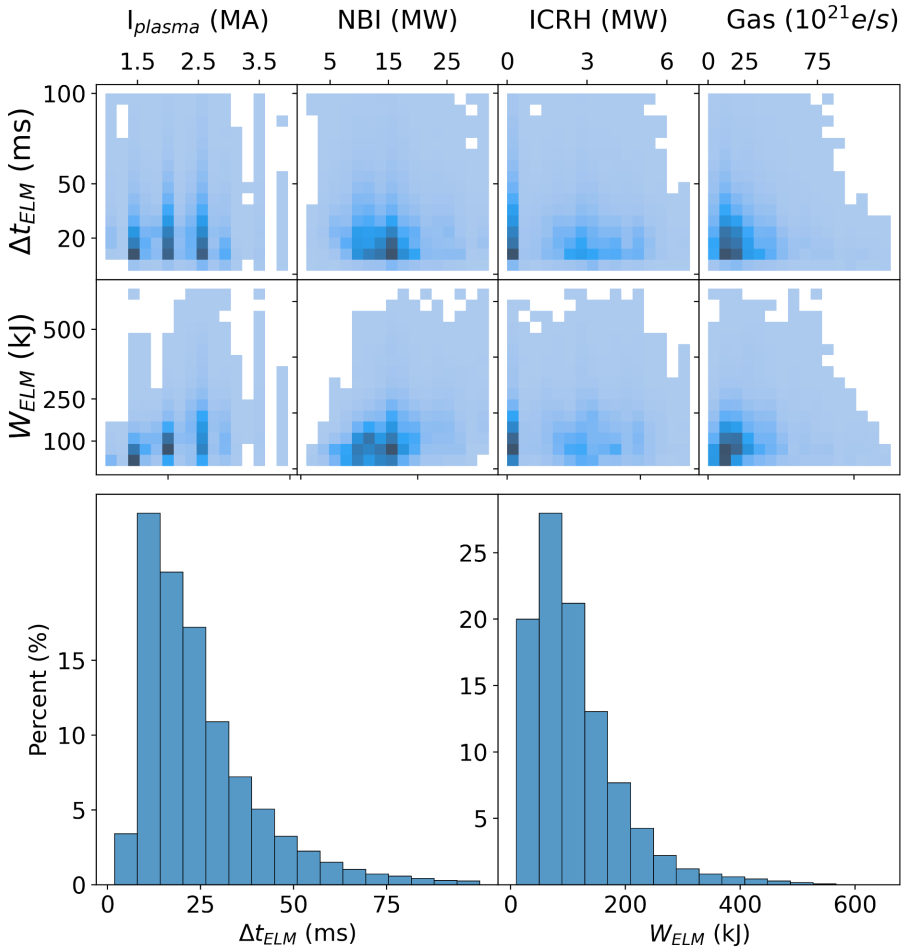

Modeling

Collective ELM behavior

- Study ELM timing and ELM losses with distributions
- More info maintained with PDF parameters

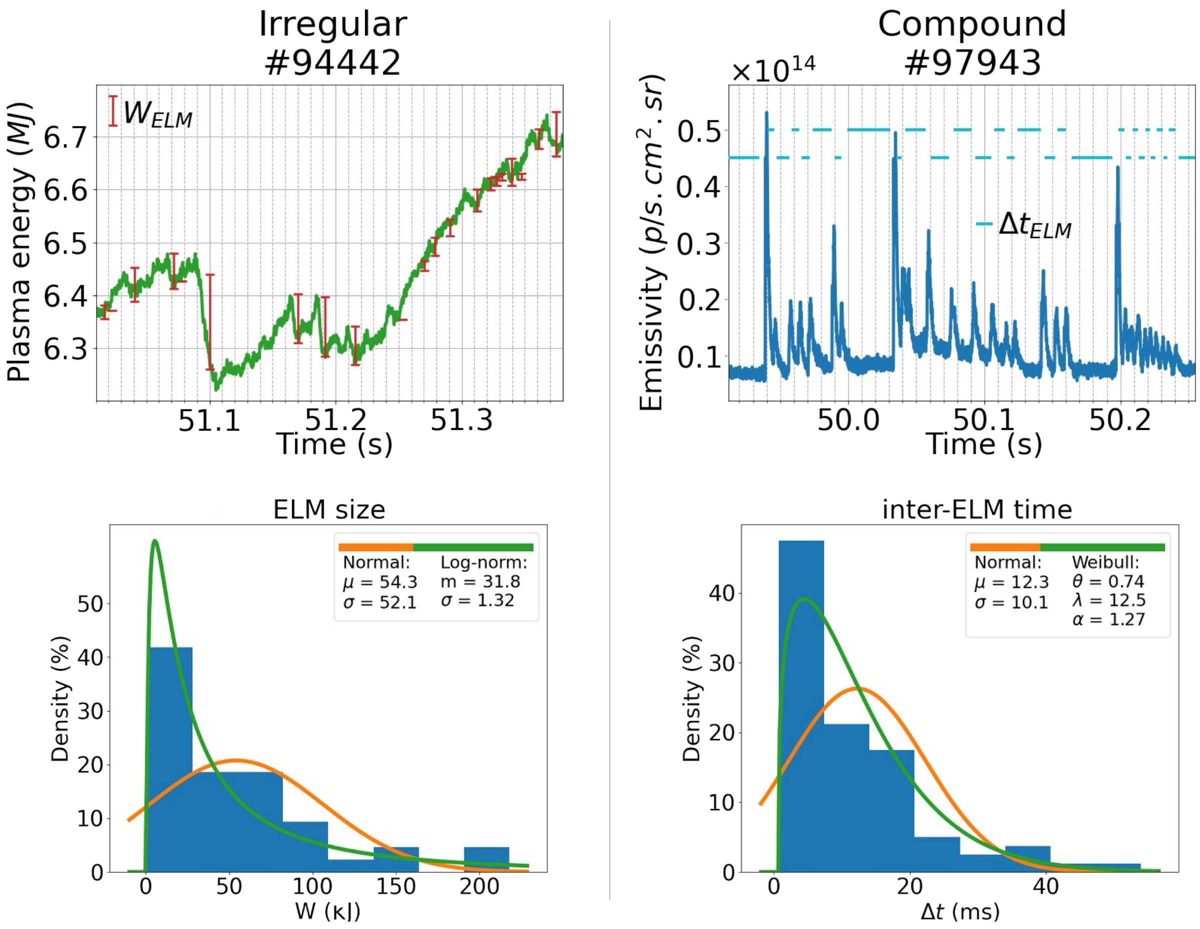

Tail behavior

Collective ELM behavior

- Study ELM timing and ELM losses with distributions
- More info maintained with PDF parameters
- Distribution tails show rare but large ELMs = risk

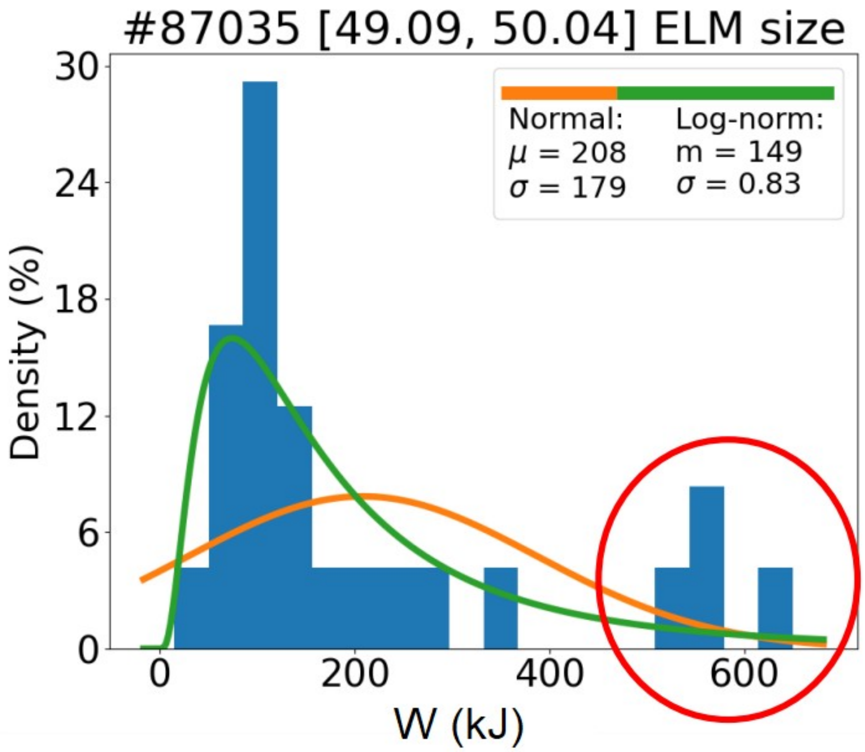

Modeling risk: tail events

- Ways to think of the risk of unforeseen large ELMs
- Tail size (top 10%)
- Tail-to-body ratio of the ELM size

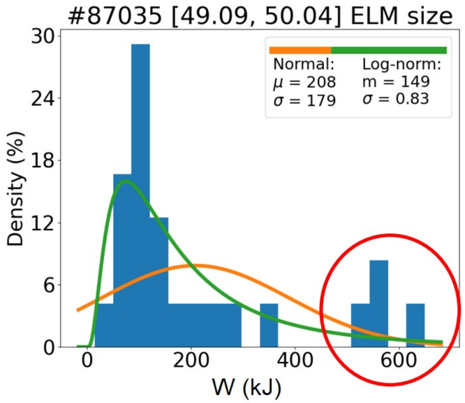

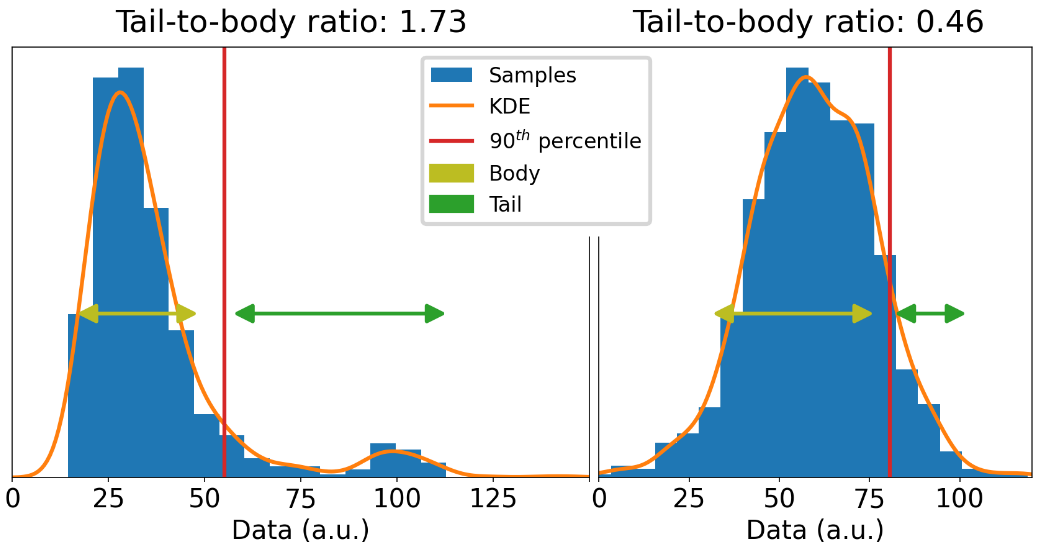

Modeling risk

- Several ELMs of varied size accumulate energy
- Response time of actuators

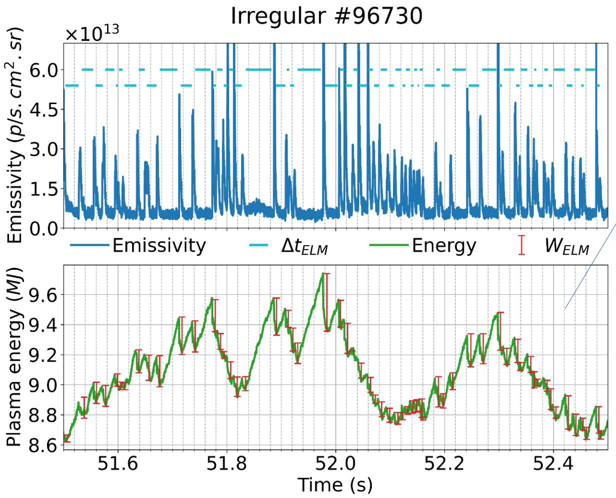

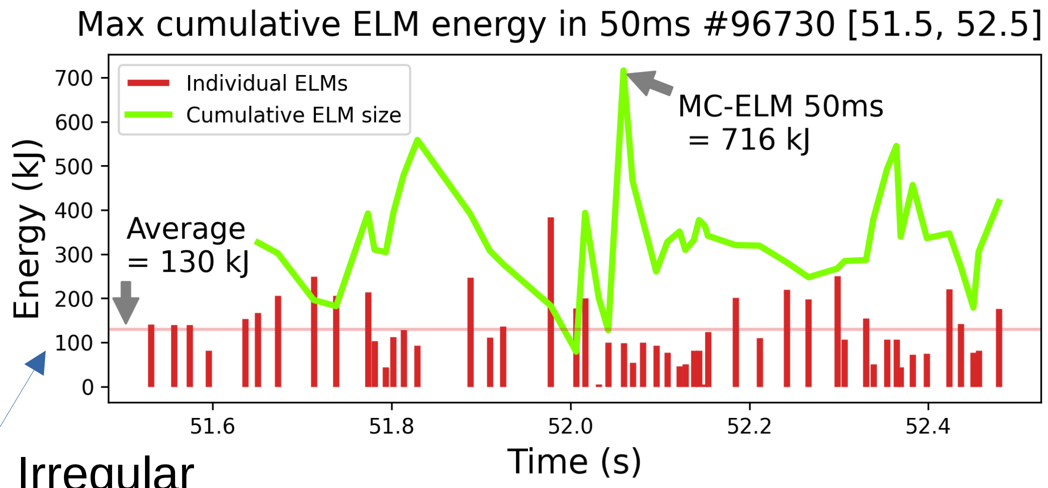

Irregular

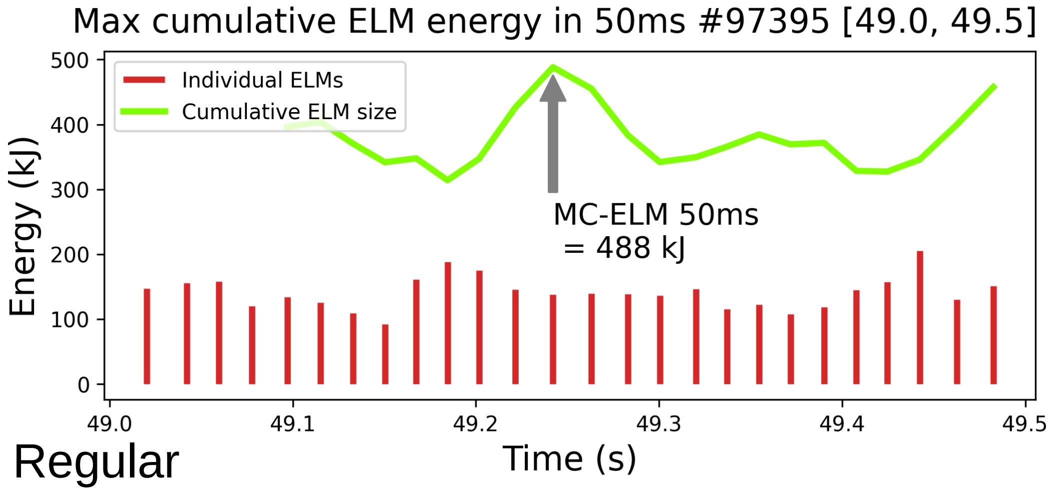

Regular

Analysis: regression maps

A R Field et al 2020 Plasma Phys. Control. Fusion 62 055010

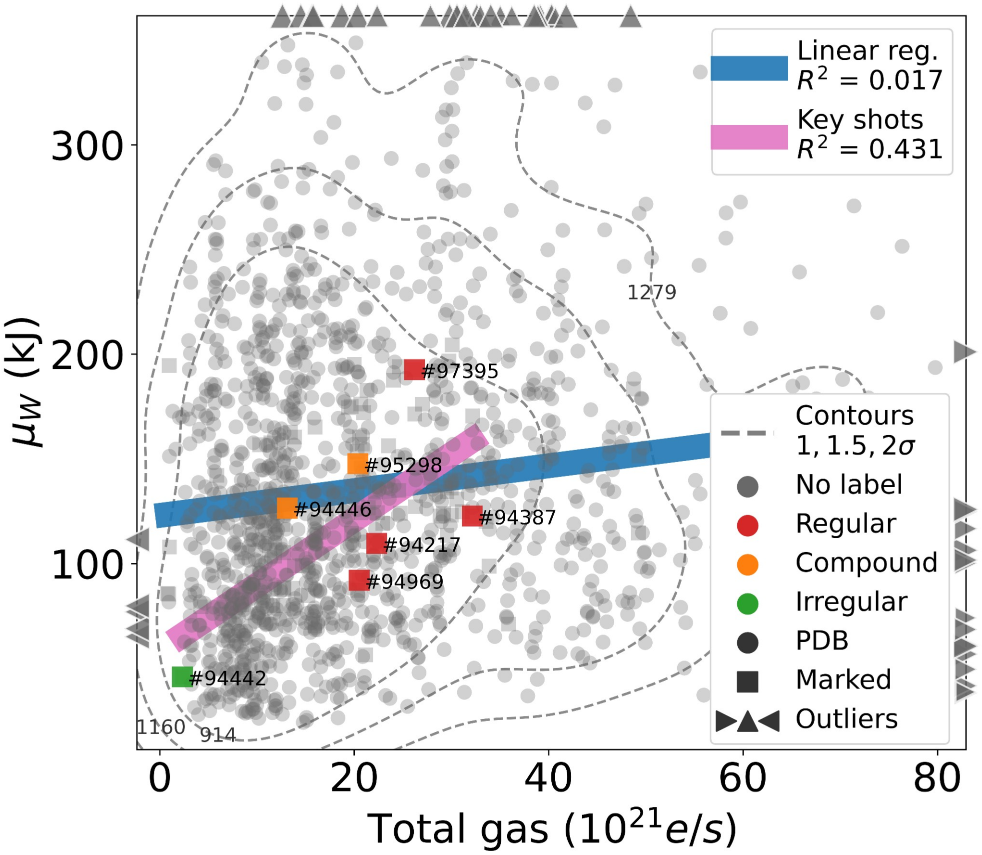

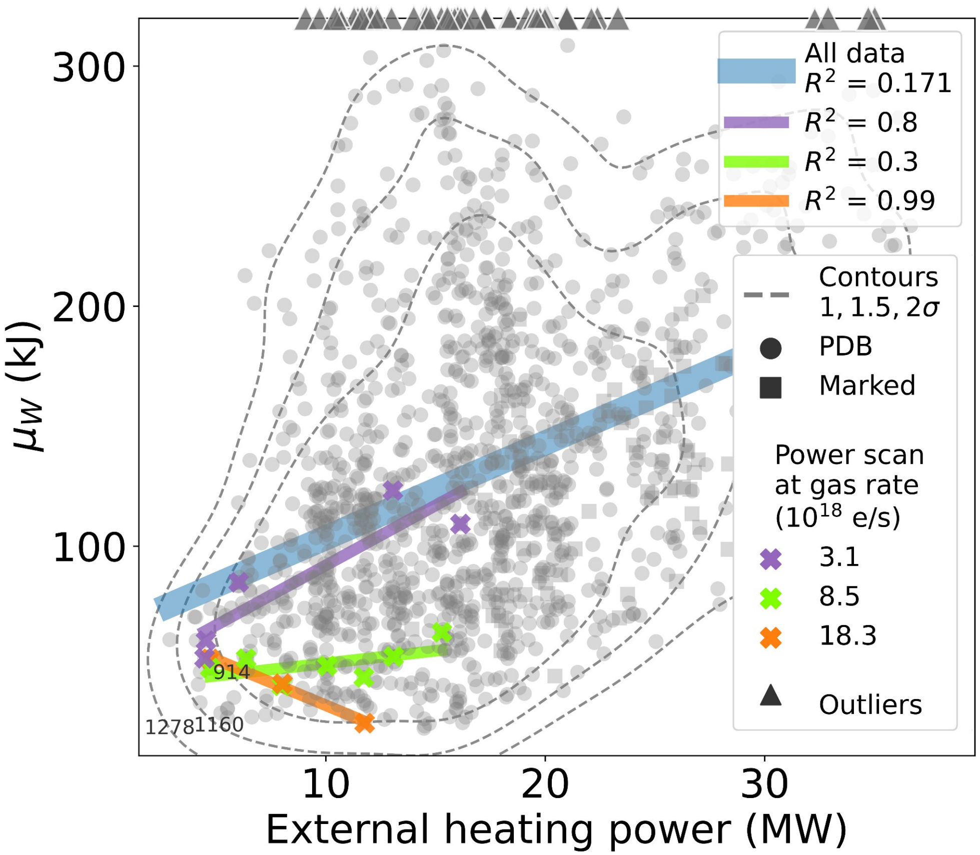

Analysis: multi-regression

- Models:
  - Linear regression
  - k-NN (k = 5%)
  - SVM regression
  - Random forests
- Reasonable fits for prediction
- Cross-validated (2/3)
- Forward selection
- Improvement in R^2

	Regression model	R2 score	±
ELM size MC 50ms	Forest	0.686	0.05
ELM size top 10%	Forest	0.64	0.072
ELM size mean	Forest	0.623	0.06
ELM size std	Forest	0.602	0.093
ELM size MC 50ms	Linear	0.565	0.049

⋮

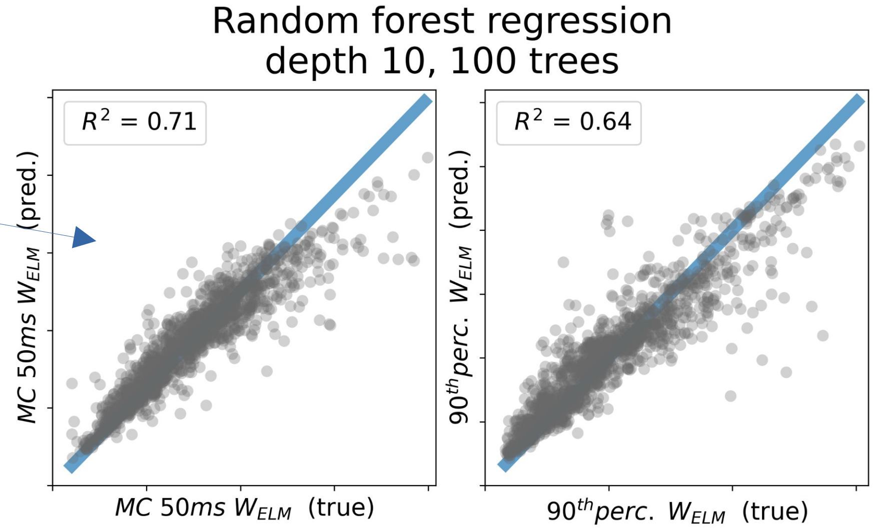

Analysis: feature importance

- Impact of each input on R^2 and percentage error
- Model-averaged feature weights

	Plasma current	Safety factor q95	Ext. heating power	ICRH power proportion	Gas fuel. throughput	Upper triangularity
ELM size MC 50ms	0.33	0.05	0.56	0	0.01	0.07
ELM size top 10%	0.48	0.03	0.02	0.02	0.08	0.24
ELM size mean	0.45	0.01	0.07	0.02	0.09	0.25
ELM size std	0.47	0.04	0.04	0.01	0.07	0.24

Summary

- Data:
  - 13000+ marked ELMs
  - Detection 90% accurate
  - 1500+ shots processed
- ELM behavior:
  - Features quantifying variability and risk
- Regression:
  - Complex relationship
  - Local effects, focus on subsets

- Explainable/interpretable data techniques

Which plasma parameters influence ELM behavior?

✔

How do plasma parameters influence ELM behavior?

~

Summary

- Data:
  - 13000+ marked ELMs
  - Detection 90% accurate
  - 1500+ shots processed
- ELM behavior:
  - Features quantifying variability and risk
- Regression:
  - Complex relationship
  - Local effects, focus on subsets

Which plasma parameters influence ELM behavior?

✔

How do plasma parameters influence ELM behavior?

~

- Explainable/interpretable data techniques
- Applicable to other machine subsystems
- Detection framework open source

→ `github.com/infusion-ugent`

Thank you

orcid.org/0009-0003-5702-939X

jerome.alhage@ugent.be

GHENT UNIVERSITY

infusion

FACULTY OF ENGINEERING AND ARCHITECTURE

EUROfusion

This work has been carried out within the framework of the EUROfusion Consortium, funded by the European Union via the Euratom Research and Training Programme (Grant Agreement No 101052200 — EUROfusion). Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Commission. Neither the European Union nor the European Commission can be held responsible for them.

JET

ELM labeling GUI

- Semi-automated, using keyboard and mouse
- Draw lines on emissivity signal, peaks above lines are ELMs if they correspond to a drop in the (high-res) plasma energy
- Can manually position cursor for compound ELMs
- Clean afterward by merging peaks, filtering small ones…

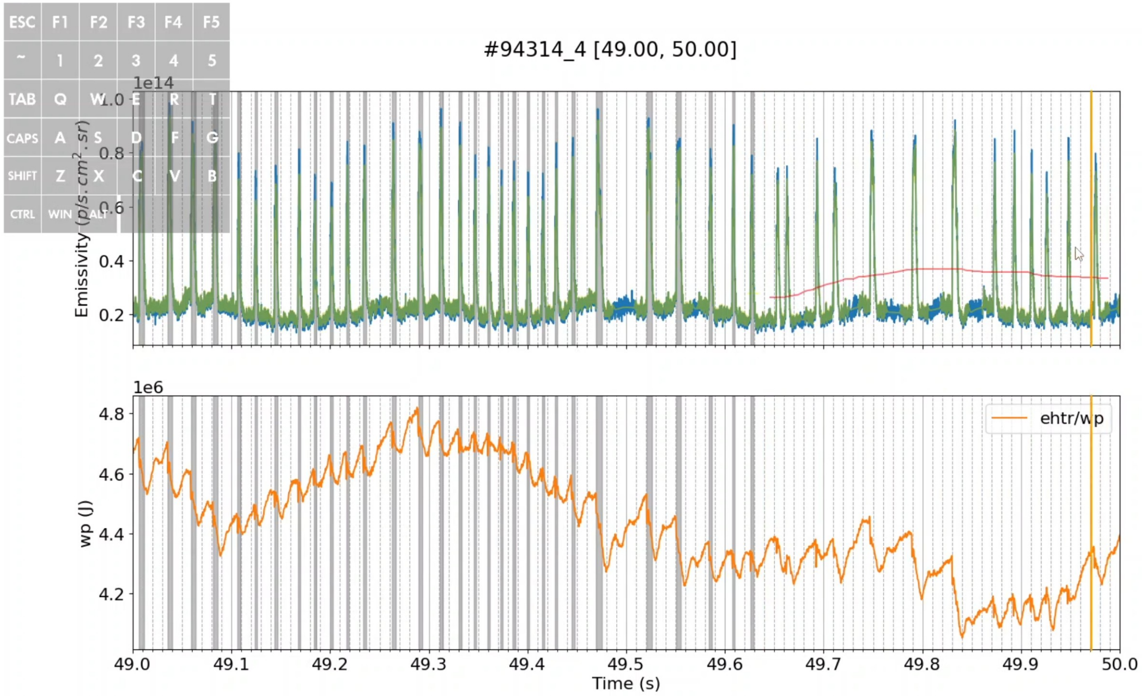

ELM size

Steps

- Calculate inter-ELM time upper limit n (n = dt 90th percentile)
- Calculate avg. ELM duration m
- Detrend Wp using rolling mean (lag = 2*n)
- Search area start is 80% of previous cycle (max: ELM start - m)
- Search area stop is 80% of current cycle (max: ELM end + n)

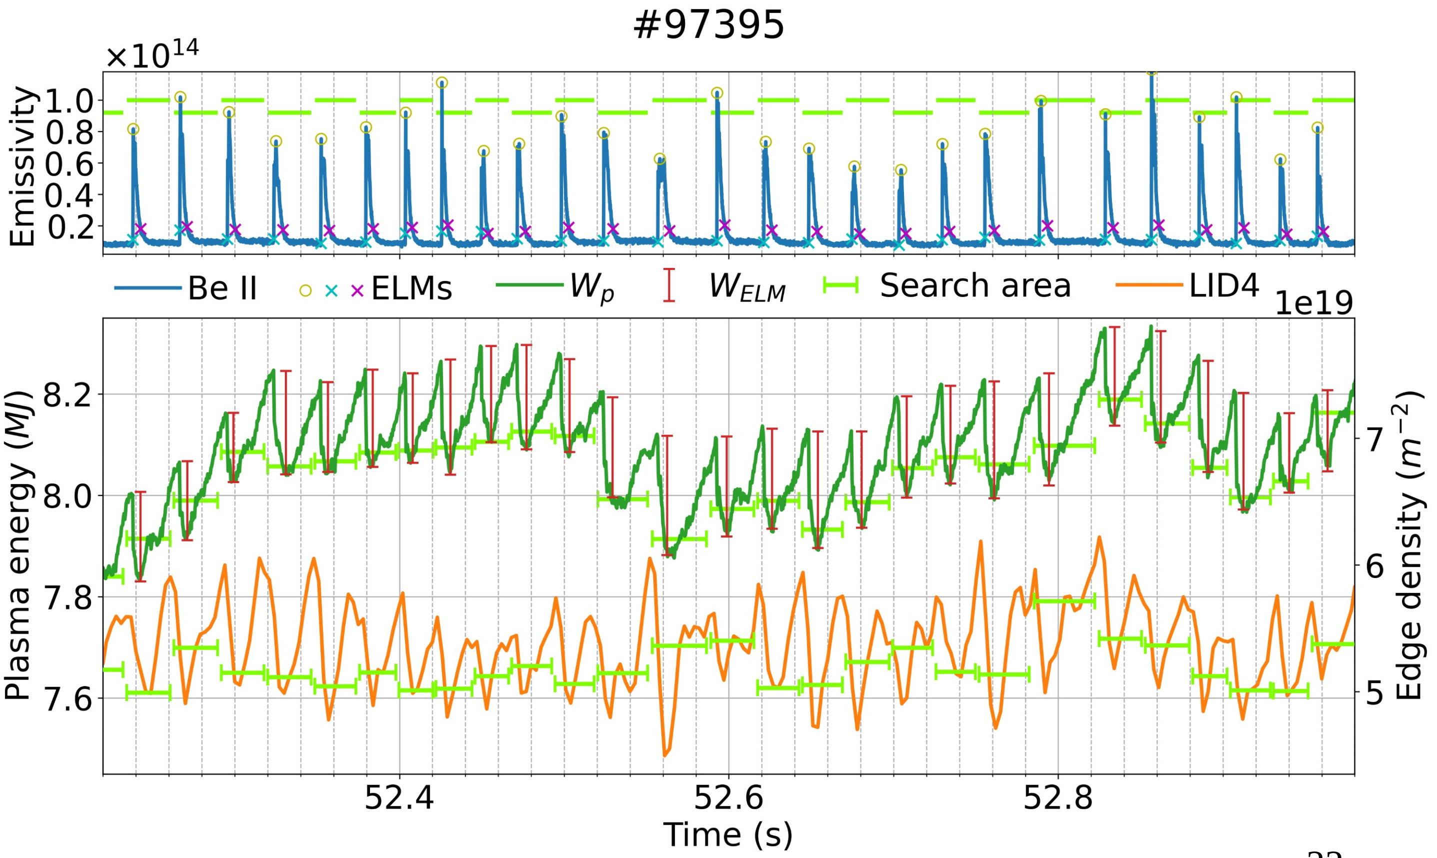

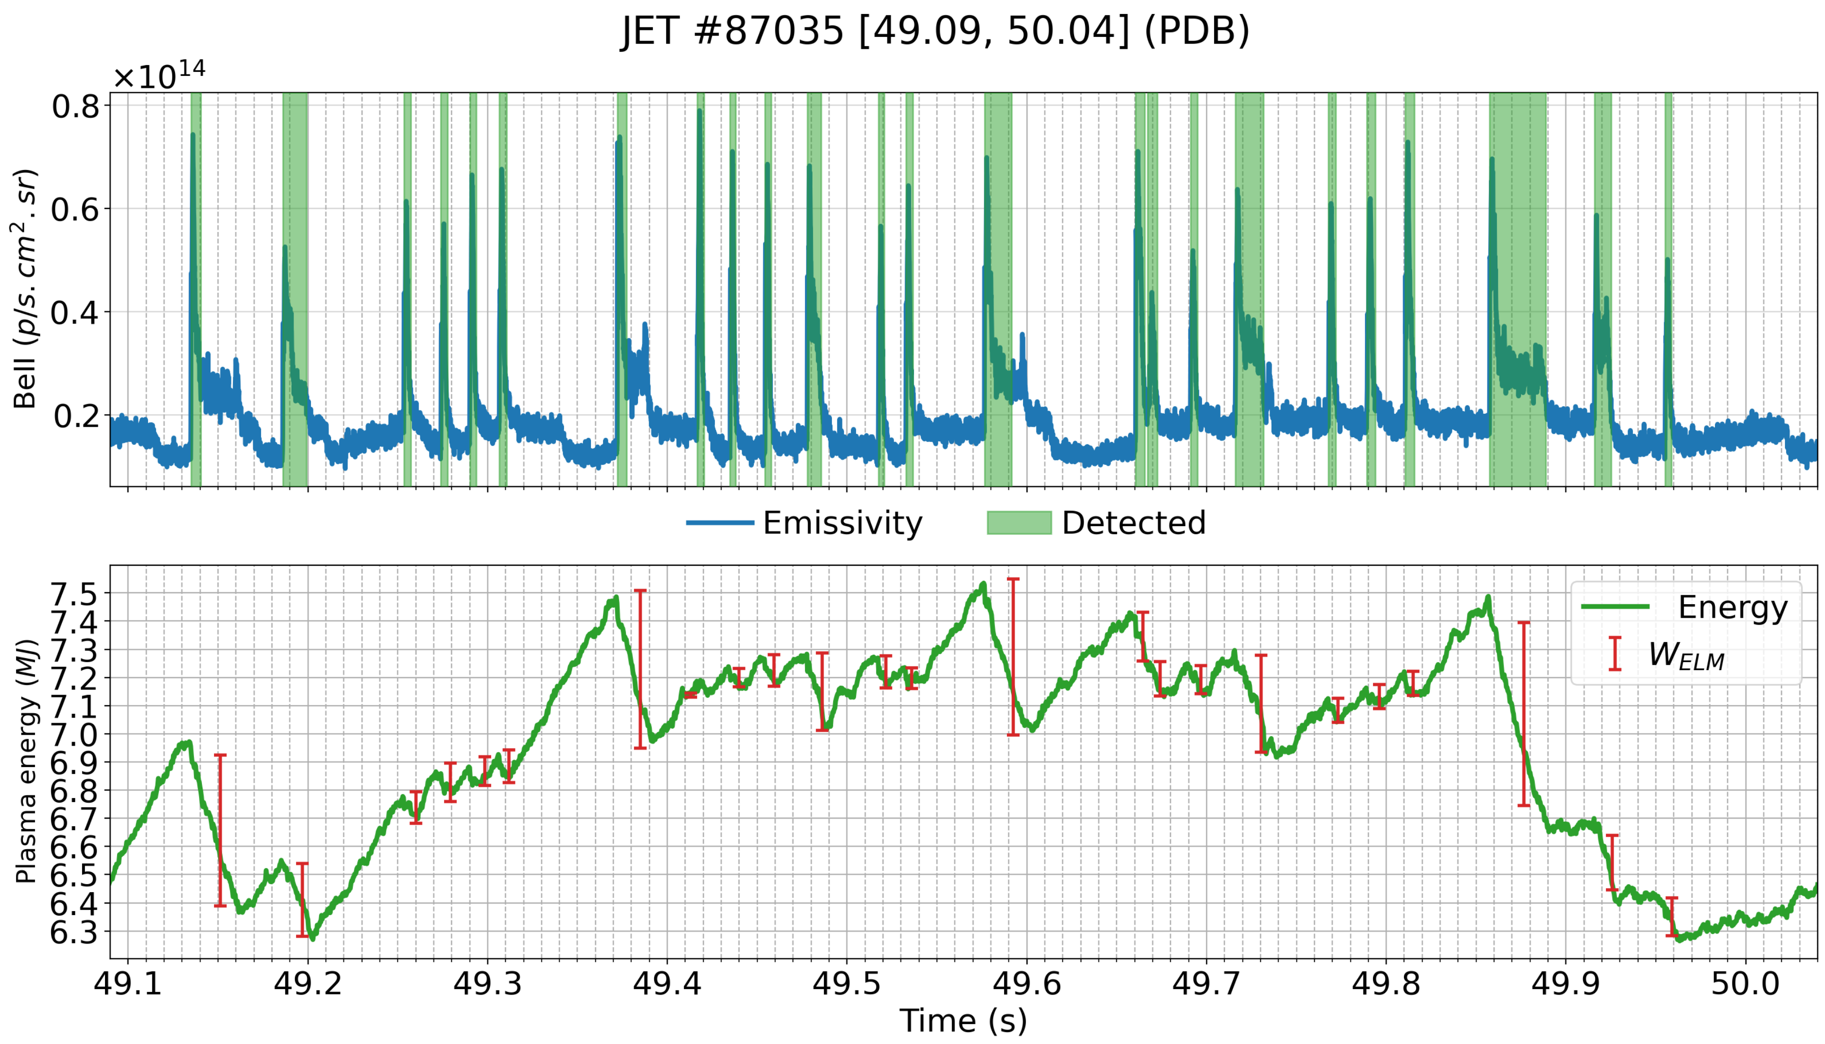

JET #87035 [49.09, 50.04] (PDB)
×10^14
Bell (p/s.cm^2.sr)
0.8
0.6
0.4
0.2
Emissivity
Detected
Plasma energy (MJ)
7.5
7.4
7.3
7.2
7.1
7.0
6.9
6.8
6.7
6.6
6.5
6.4
6.3
Energy
W_ELM
49.1
49.2
49.3
49.4
49.5
49.6
49.7
49.8
49.9
50.0
Time (s)

ELM power

How calculated

- ELM size ~ the drop in plasma stored energy Wp (green signal)
- ELM power is the drop energy (red) divided by the drop time (pink)
- Also calculate the max instantaneous ELM power (from derivative)

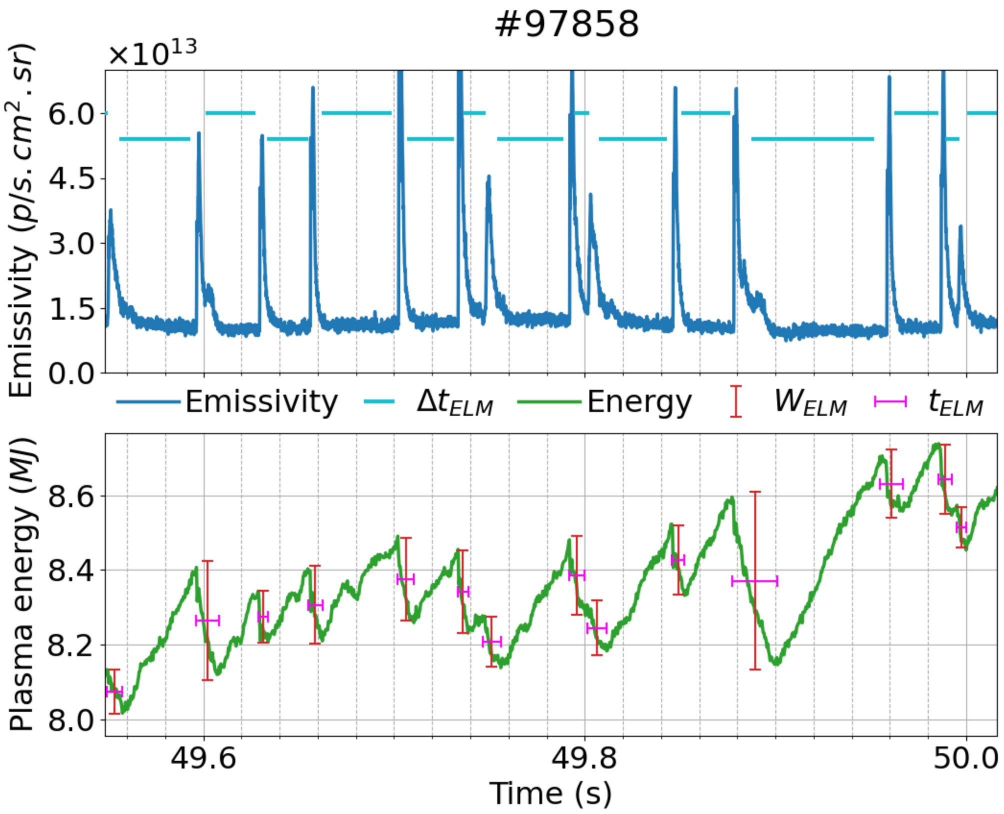

ELM power

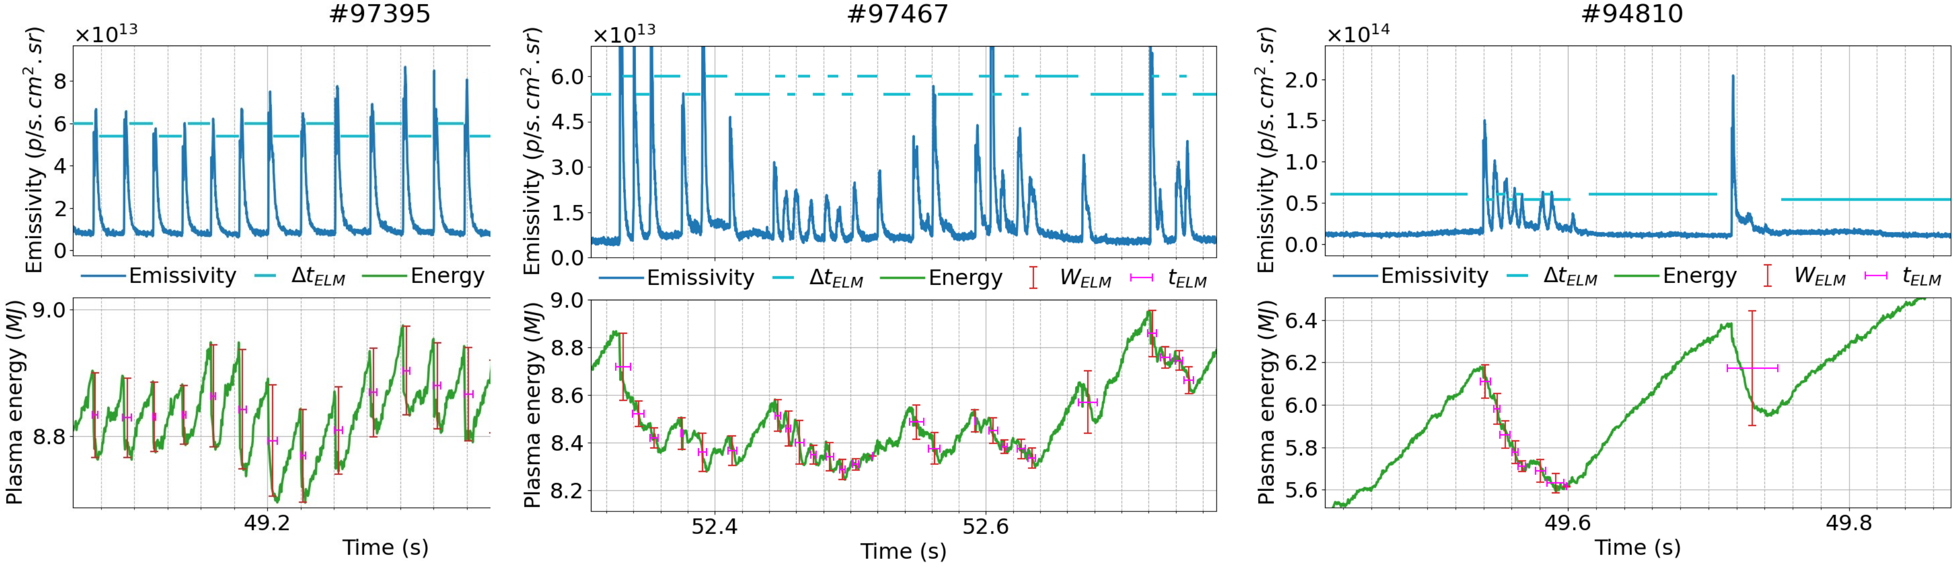

#97395
#97467
#94810
×10^13
×10^13
×10^14
Emissivity (p/s. cm^2 . sr)
Emissivity
Δt_ELM
Energy
W_ELM
t_ELM
Plasma energy (MJ)
Time (s)
49.2
52.4
52.6
49.6
49.8

ELM power

- ELM power increases with increased heating power.
- Rate of increase depends on machine parameters (ex: gas rate).

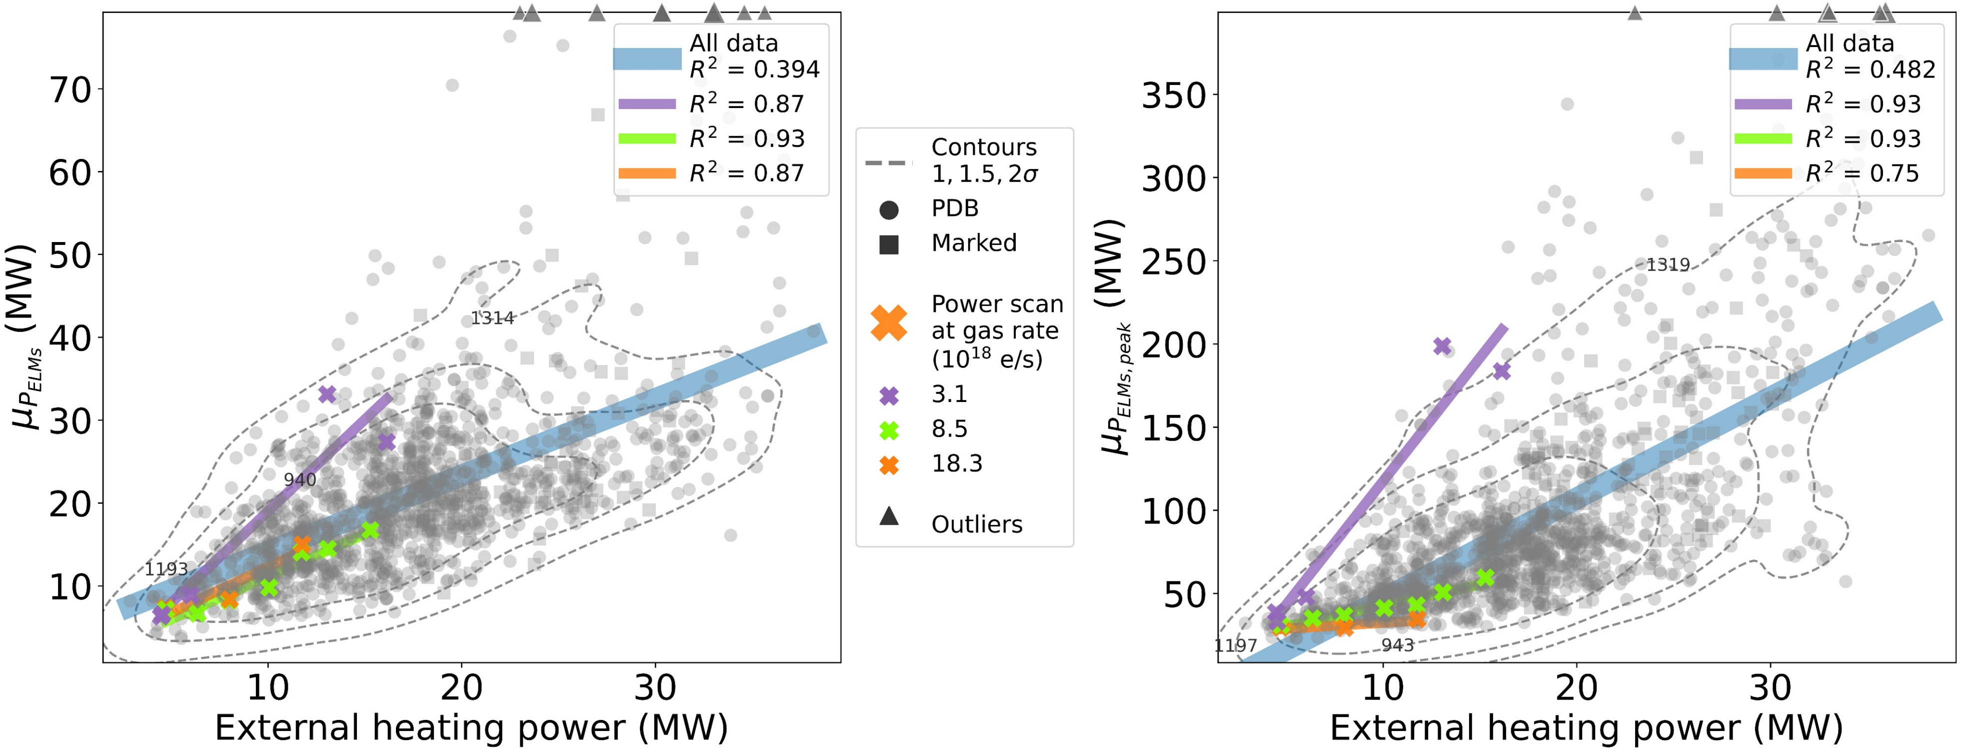

Detection benchmark

- CNN stops learning
- RNN GRU-3 still room for improvement (20 epochs)
- LoG-CK good for real-time

Method	CNN	RNN	non-NN
Training time per epoch	3 h	2 h	-
Inference time per second	3.76 s	1.63 s	approx. 50 ms
Number of parameters	812818	25916	<20

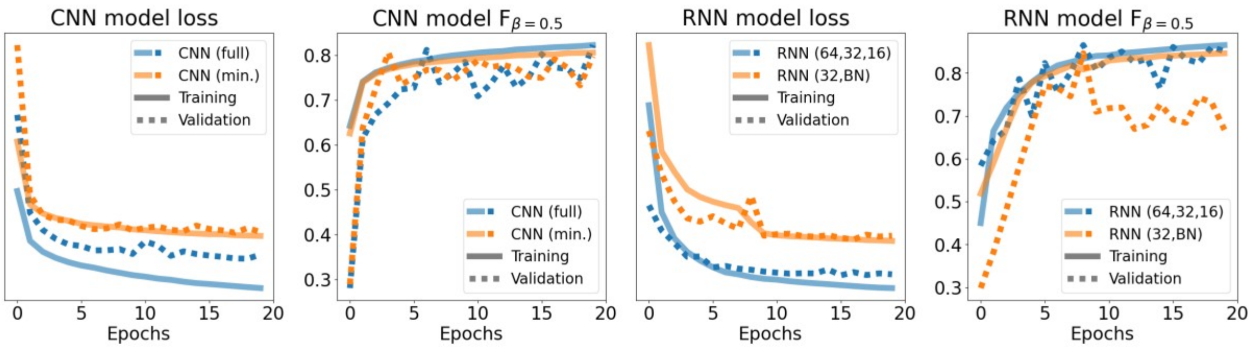

(a) CNN loss (b) CNN $F_{0.5}$ (c) RNN loss (d) RNN $F_{0.5}$

Detection preview

- RNN is proof-of-concept with great potential
- Prefer LoG-CK for simplicity, false positives

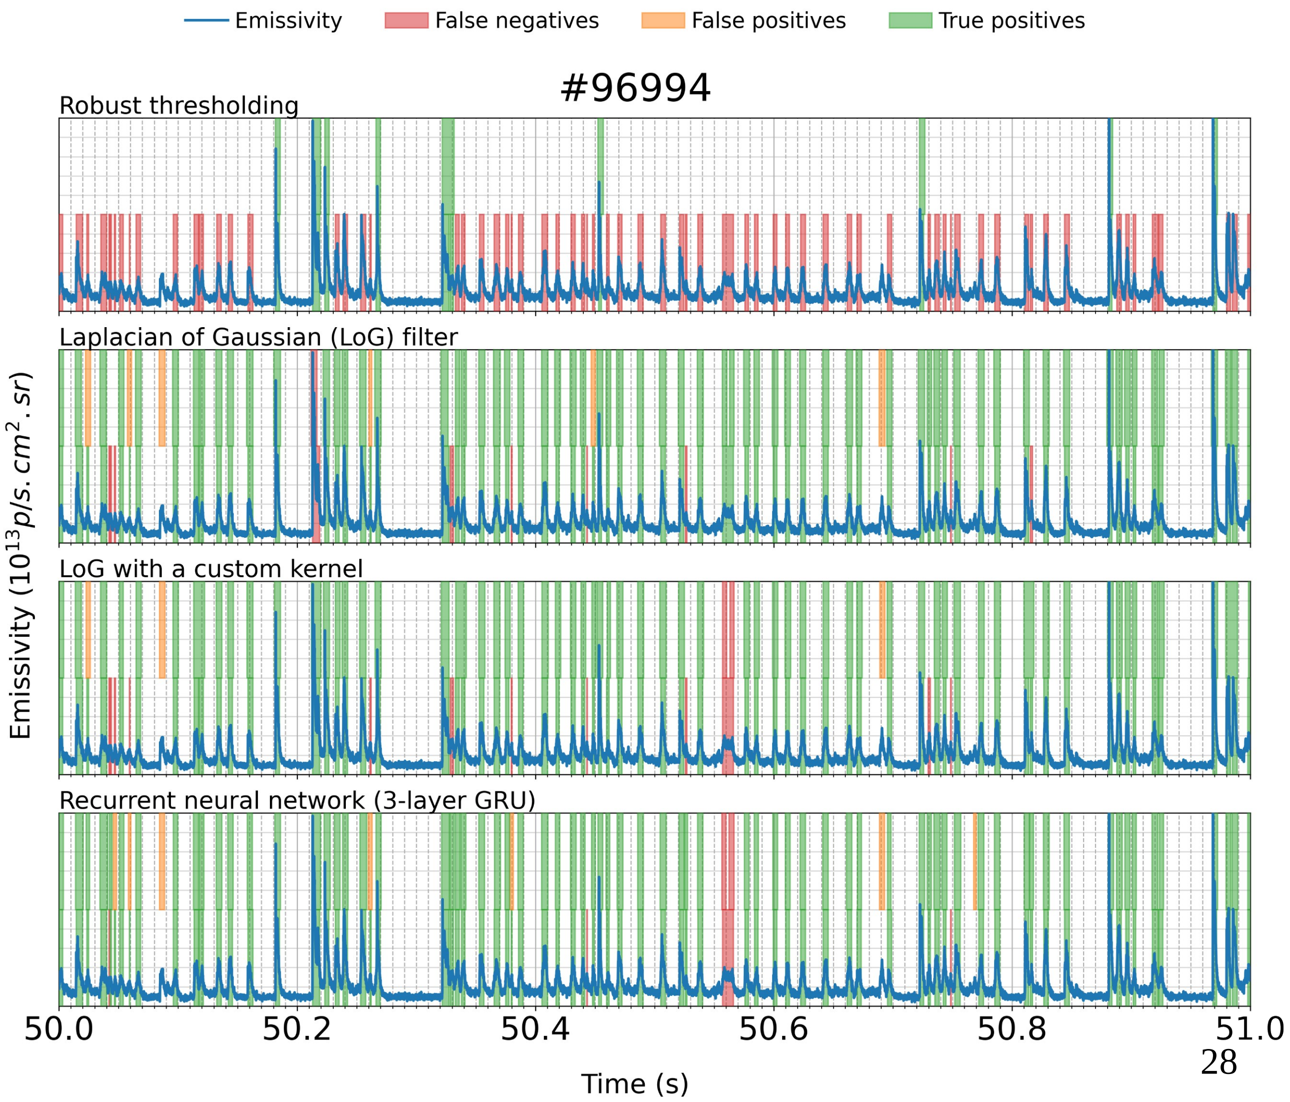

Detection breakdown

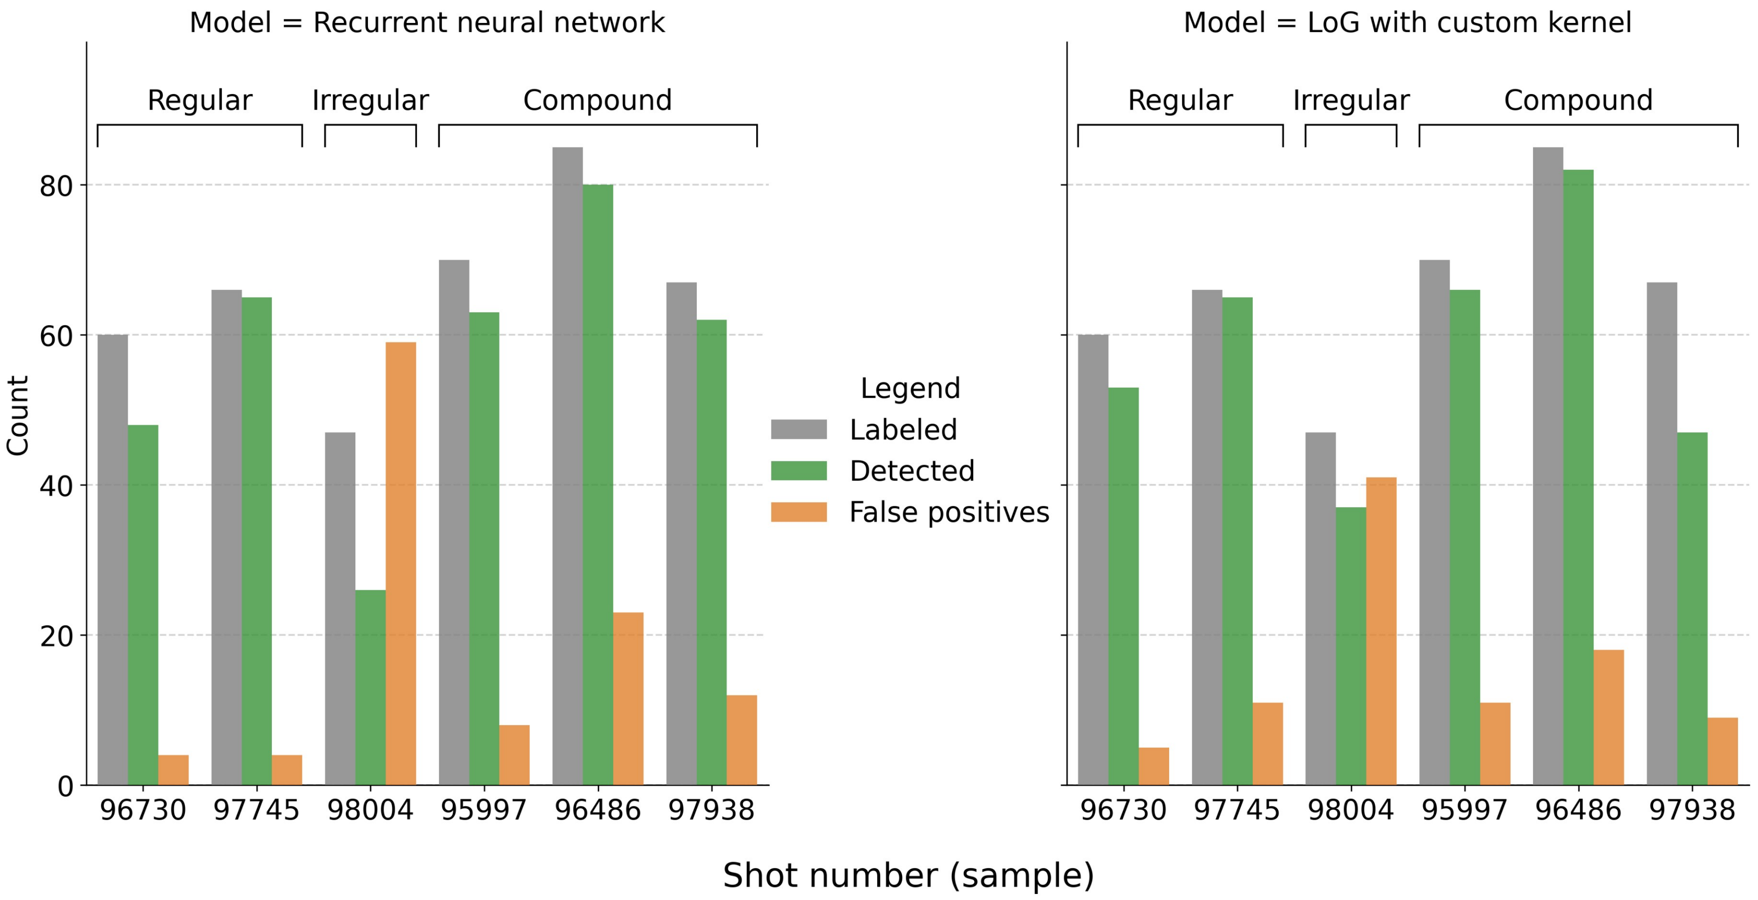

Data selection

- Distilled 90000+ ELMs from 1388 windows
  ~ 1.5s duration
  ~ 66 ELMs per sample

- Natural ELMs
- Varied operating conditions

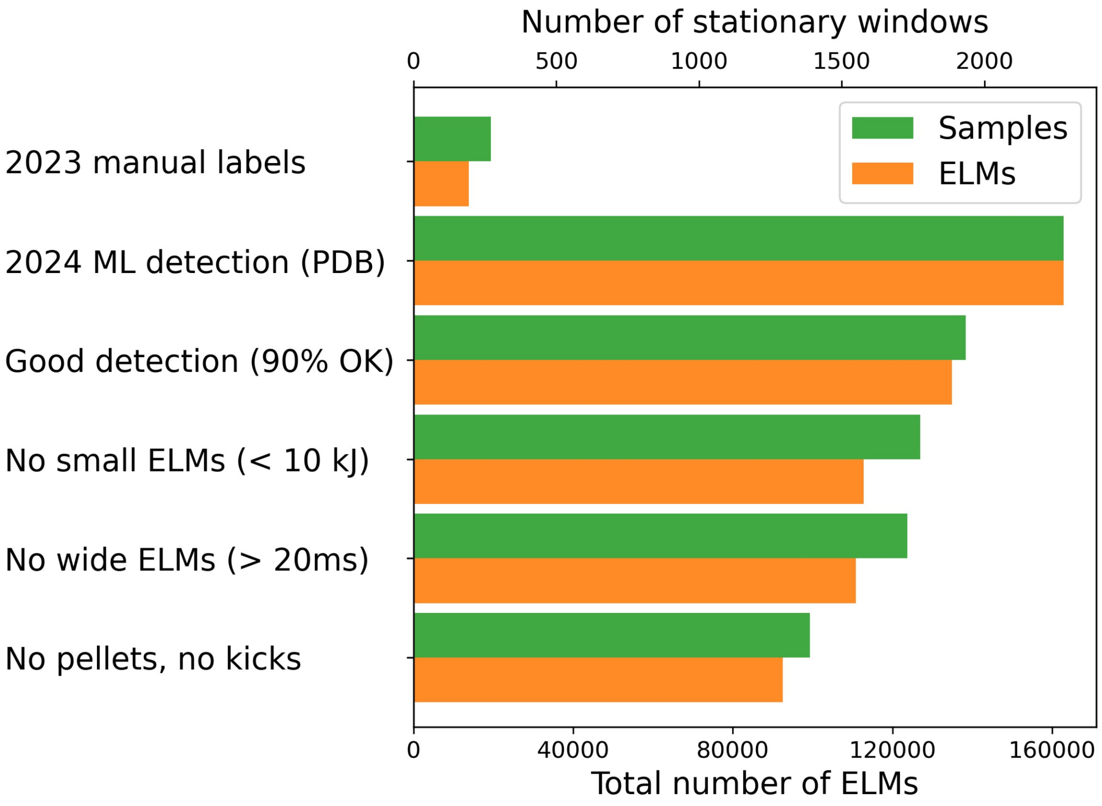

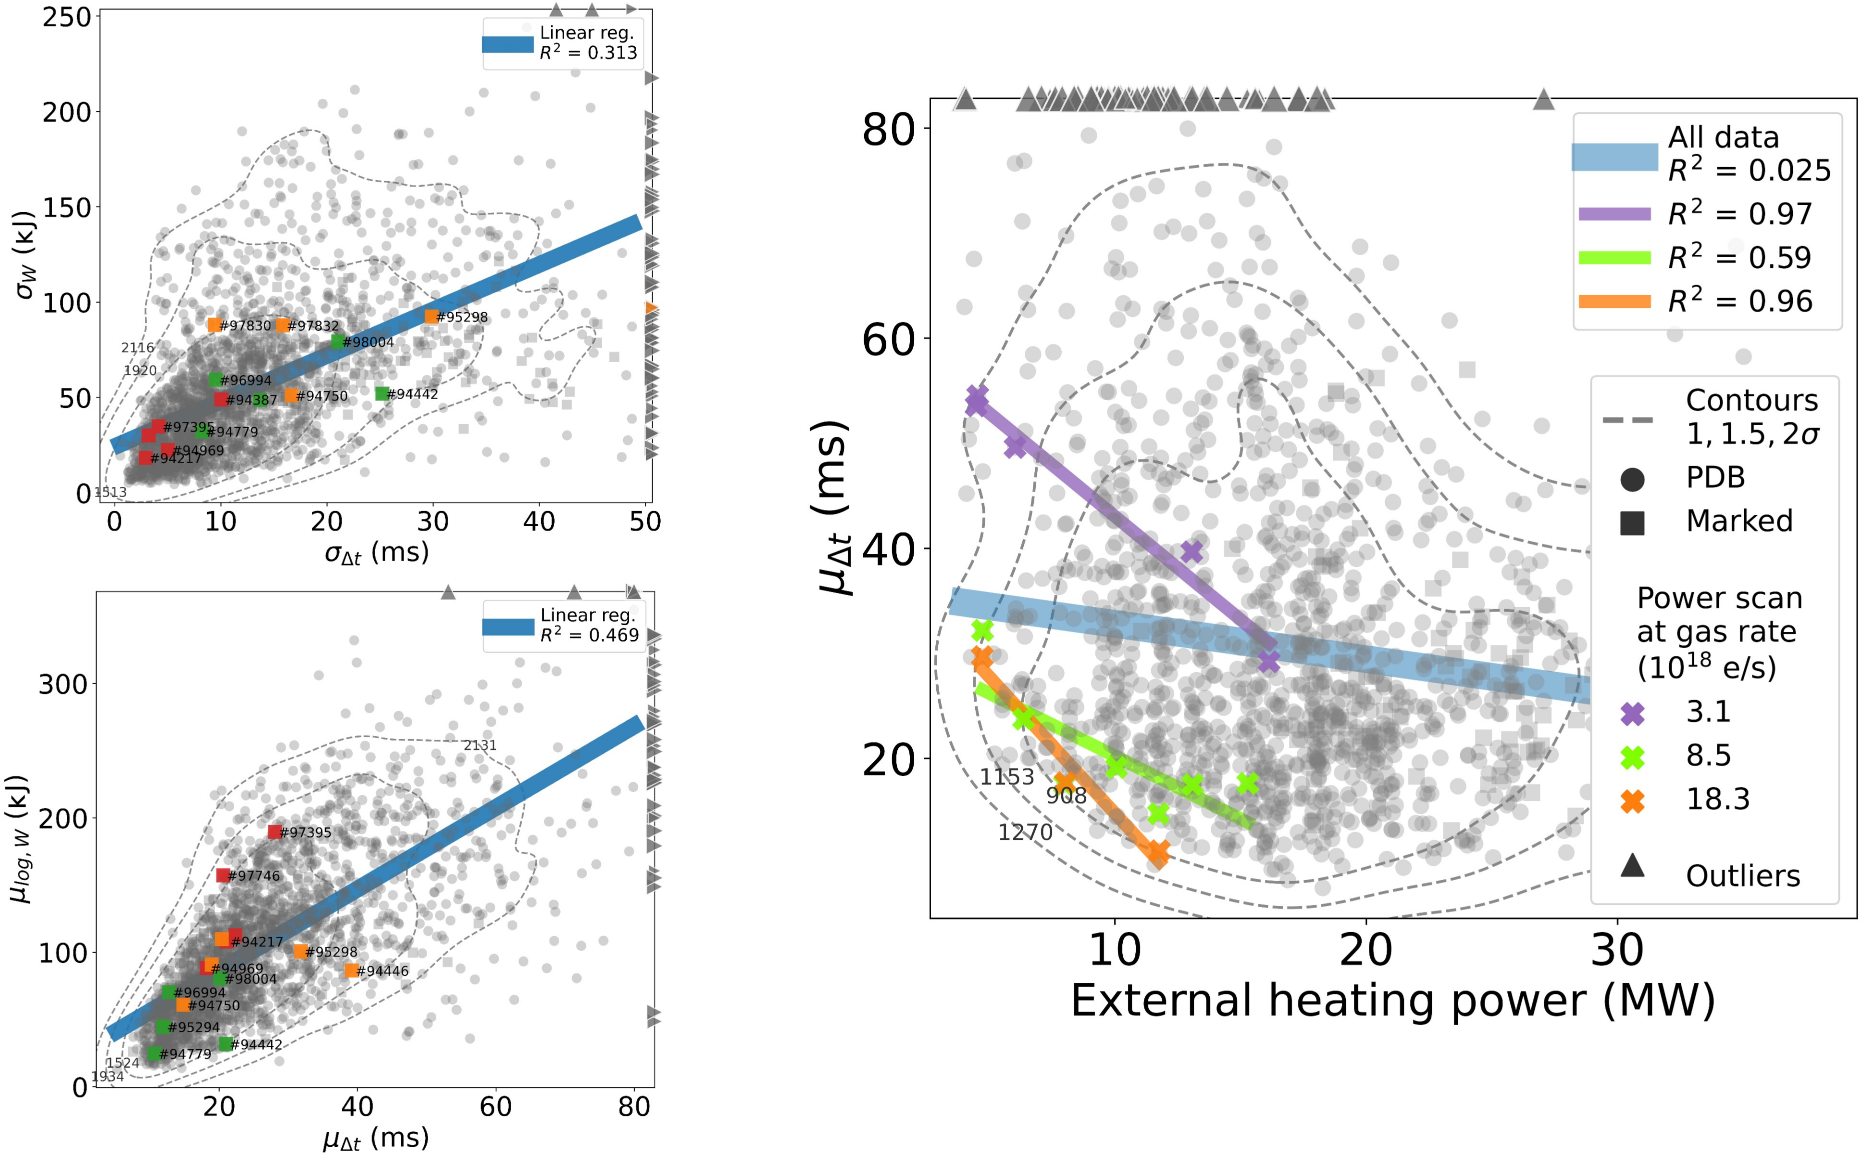
Linear reg.
R^2 = 0.313
σ_W (kJ)
$\sigma_{\Delta t}$ (ms)
#97830
#97832
#95298
#98004
#96994
#94387
#94750
#94442
#97395
#94779
#94969
#94217
2116
1920
1513
Linear reg.
R^2 = 0.469
$\mu_{log,W}$ (kJ)
$\mu_{\Delta t}$ (ms)
2131
#97395
#97746
#94217
#95298
#94969
#98004
#94446
#96994
#94750
#95294
#94779
#94442
1524
1934
All data
R^2 = 0.025
R^2 = 0.97
R^2 = 0.59
R^2 = 0.96
Contours
$1, 1.5, 2\sigma$
PDB
Marked
Power scan
at gas rate
(10^{18} e/s)
3.1
8.5
18.3
Outliers
1153
908
1270
$\mu_{\Delta t}$ (ms)
External heating power (MW)

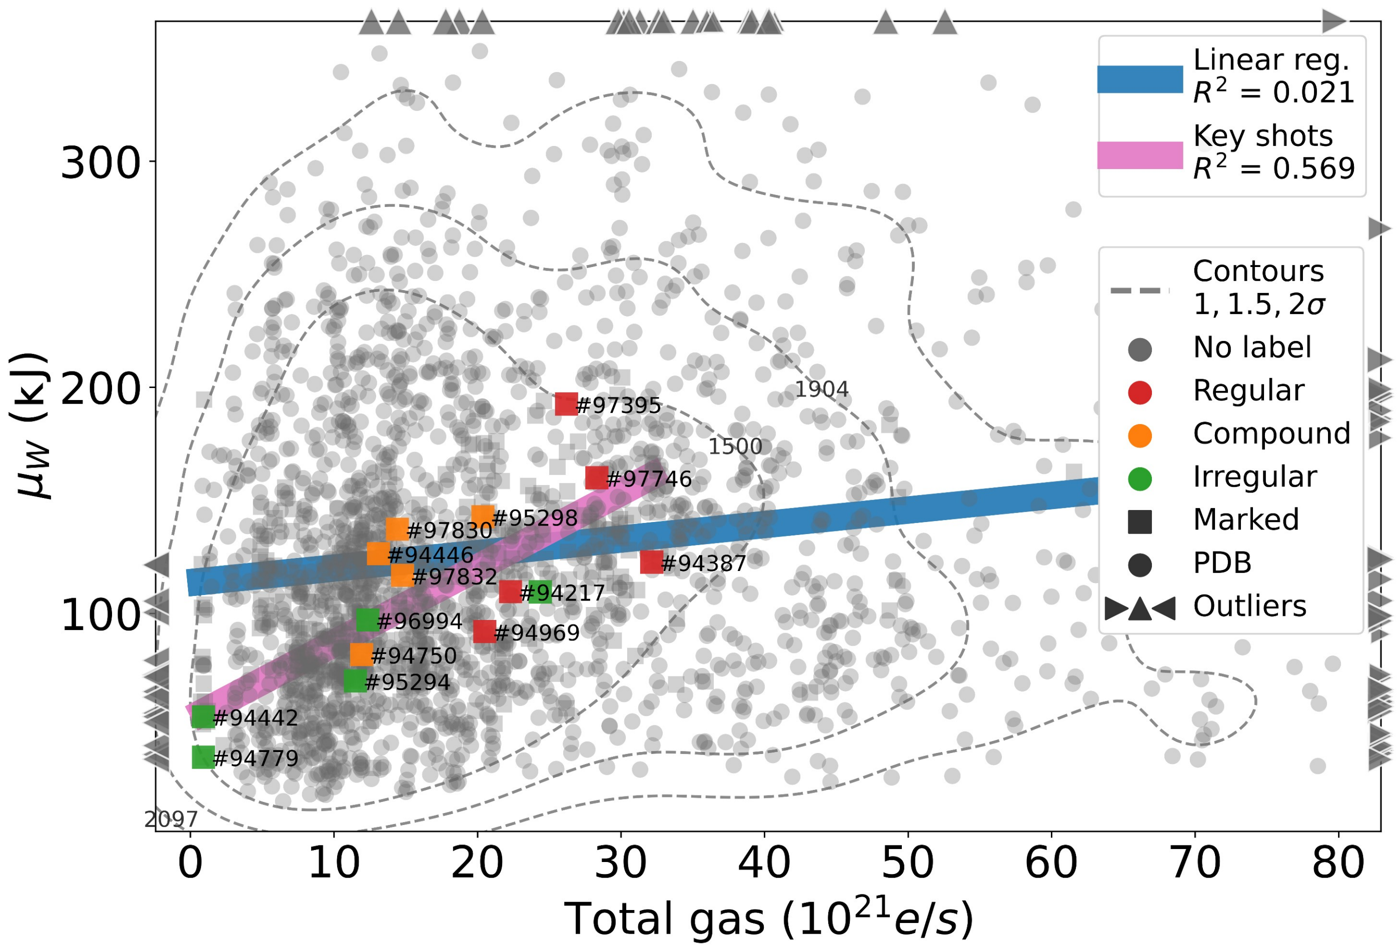
Linear reg.
R^2 = 0.021
Key shots
R^2 = 0.569
Contours
1, 1.5, 2σ
No label
Regular
Compound
Irregular
Marked
PDB
Outliers
1904
1500
2097
#97395
#97746
#95298
#97830
#94446
#97832
#94387
#94217
#96994
#94969
#94750
#95294
#94442
#94779
μ_W (kJ)
Total gas ($10^{21}e/s$)
0
10
20
30
40
50
60
70
80
100
200
300

Regression: feature selection

- Models:
  - Linear regression
  - KNN (k = 5%)
  - SVM regression
  - Random forests
- Reasonable fits for prediction
- Cross-validated (2/3)
- Forward selection
- Improvement in R^2

	Regression model	R2 score
ELM size MC 50ms	Forest	0.686
ELM size top 10%	Forest	0.64
ELM size mean	Forest	0.623
ELM size std	Forest	0.602
ELM size MC 50ms	Linear	0.565

⋮

Plasma current	Safety factor q95	Ext. heating power	ICRH power proportion	Gas fuel. throughput	Upper triangularity
0.424	0.083	0.121		0.008	0.039
0.264	0.019		0.011	0.115	0.17
0.218	0.018		0.007	0.129	0.196
0.278	0.076	0.005	0.013	0.055	0.15
0.099		0.447			0.012

Feature selection

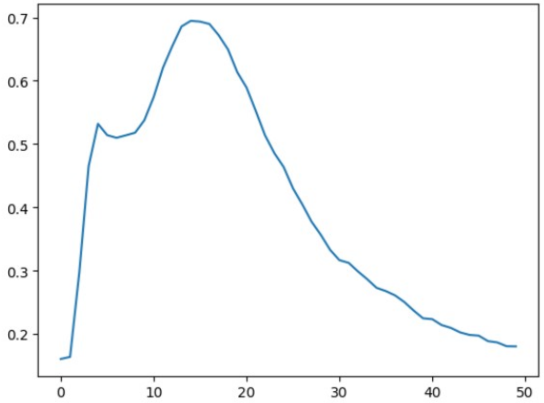

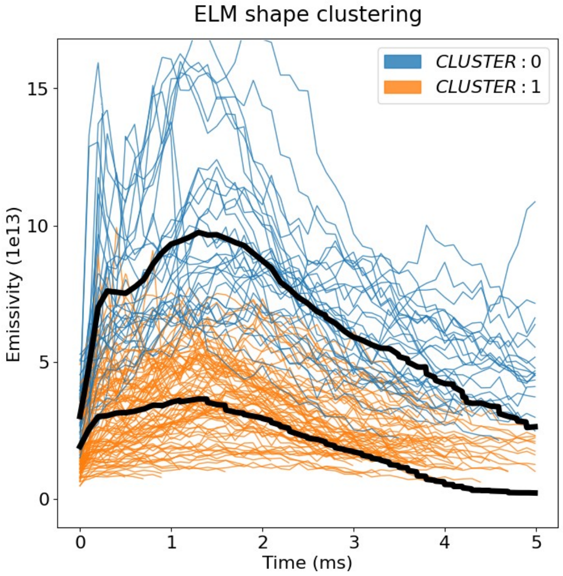

n_clusters	cluster label	ELM size (κJ)	ELM size (%)	peak width (ms)	peak height (1e13 p/s.cm2.sr)	peak skew	count
2	0	135.6	2	5	11.8	0.2	3931
2	1	81.6	1.1	3.3	4.2	0	9923
3	0	185.3	2.6	5.2	16.5	0.3	1330
3	1	108.3	1.6	4.6	8.5	0.1	4200
3	2	77.1	1	3.2	3.6	0	8324
4	0	202.6	2.7	5.2	18.7	0.4	859
4	1	121.2	1.9	5.1	10	0.1	2522
4	2	101.8	1.4	4	6.6	0.1	4061
4	3	70.1	0.9	3	3.1	0	6412
5	0	203.9	2.7	5.1	19.2	0.4	782
5	1	134.1	2.1	6	9.9	0.1	1420
5	2	102.6	1.5	4	9.5	0.2	2287
5	3	100.9	1.3	4.1	5.4	0.1	4220
5	4	64.6	0.8	2.7	2.8	-0.1	5145
6	0	204.8	2.7	5.1	19.3	0.4	763
6	1	135.8	2.1	6	10.1	0.1	1386
6	2	102.8	1.6	4	9.5	0.2	2209
6	3	86.7	1.1	3	6.1	0.1	2360
6	4	109.7	1.5	5	4.7	0	2362
6	5	64.4	0.8	2.7	2.7	-0.1	4774

n_clusters	class_variance
2	0.028%
3	0.000%
4	0.058%
5	0.063%
6	21.962%

Regression: fuzzy logic

- Predict when the tail ratio (risk) will be high:
  - Fuzzy classes (wiki)
  - Rules partition the dataset (parts with no support hidden)

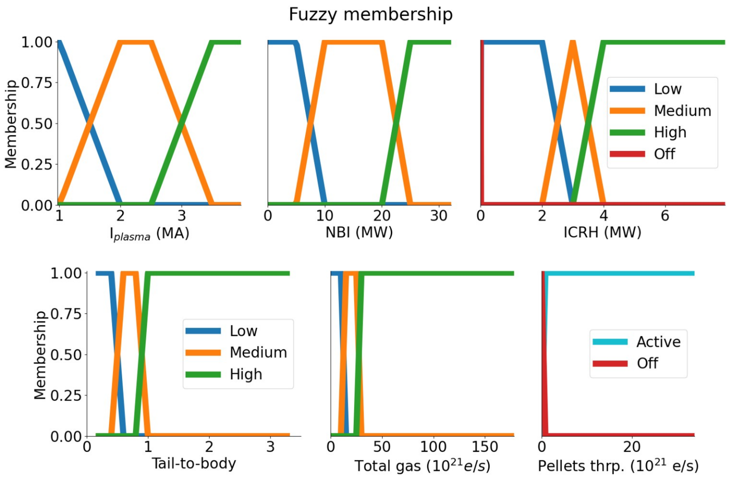

Heating			Fueling		Rule		
I_{plasma}	**NBI**	**ICRH**	**Gas**	**Pellets**	**Support**	**Degree** (%)	**Tail ratio**
High	High	High	Low	Active	36	75.0	High
High	High	High	Medium	Active	18	65.8	High
High	High	Medium	Medium	Active	19	74.5	High
....							
Medium	High	High	Medium	Off	51	67.9	Medium
Medium	High	Low	High	Off	22	74.9	Medium
Medium	Medium	Medium	Medium	Off	72	61.3	Medium
Medium	Medium	Medium	Low	Off	39	63.9	Medium
Medium	Medium	High	Low	-	53	49.6	Medium
Medium	High	Medium	Medium	-	38	45.3	Medium
....							

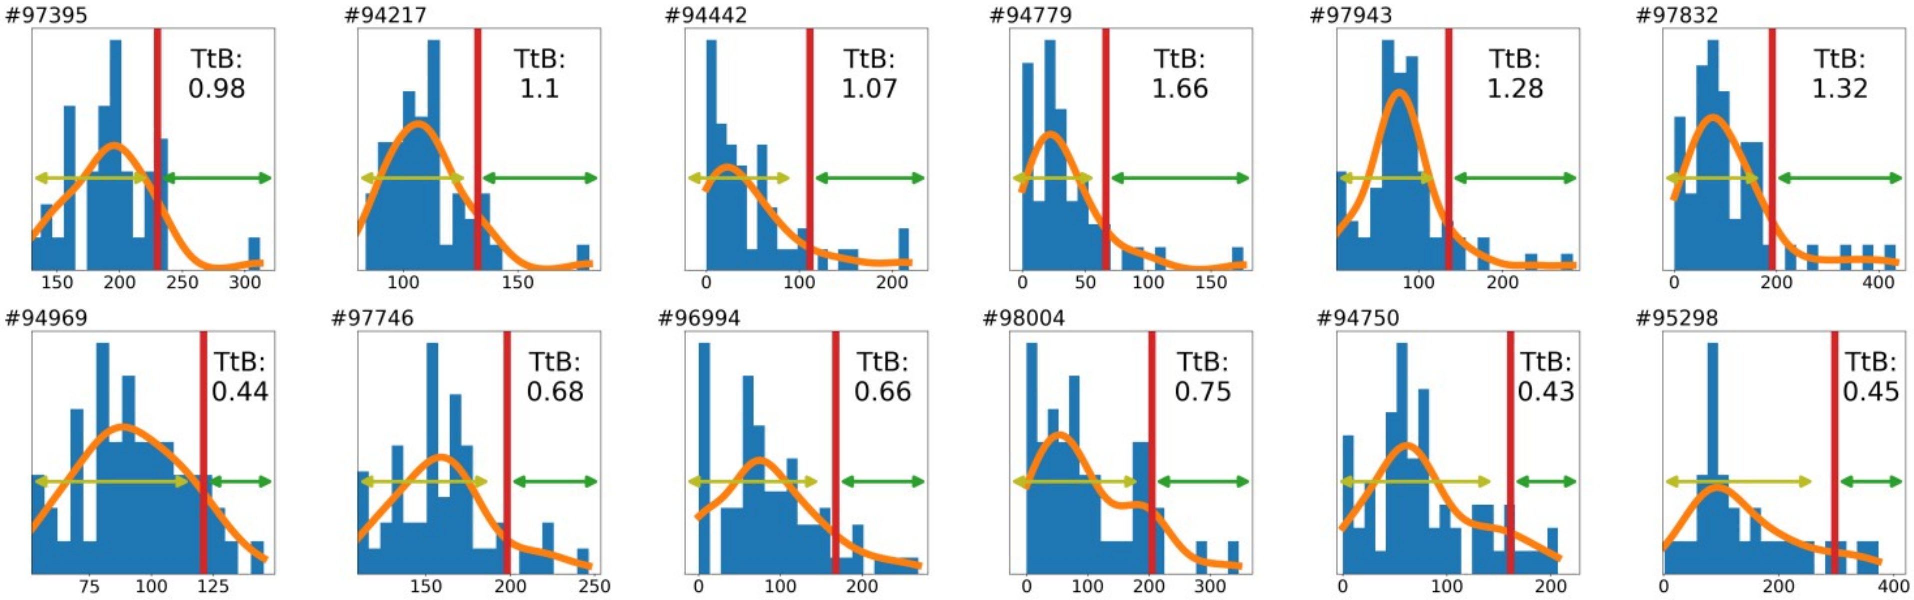

(a) Regular (b) Irregular (c) Compound

Ipla \ TTB	**High**	**Mid**	**Low**
High	52	42	42
Mid	223	618	947
Low	27	116	206