Applying machine learning techniques to the operation and optimization of the VENUS ECR ion source

Damon Todd

Work from machine learning and ion source teams:
Ezra Apple, Janilee Benitez, Heather Crawford, Alex Kireeff, Yue Shi Lai,
Marco Salathe, Wenhan Sun, Victor Watson

Lawrence Berkeley National Laboratory
February 07, 2025

An optimization problem

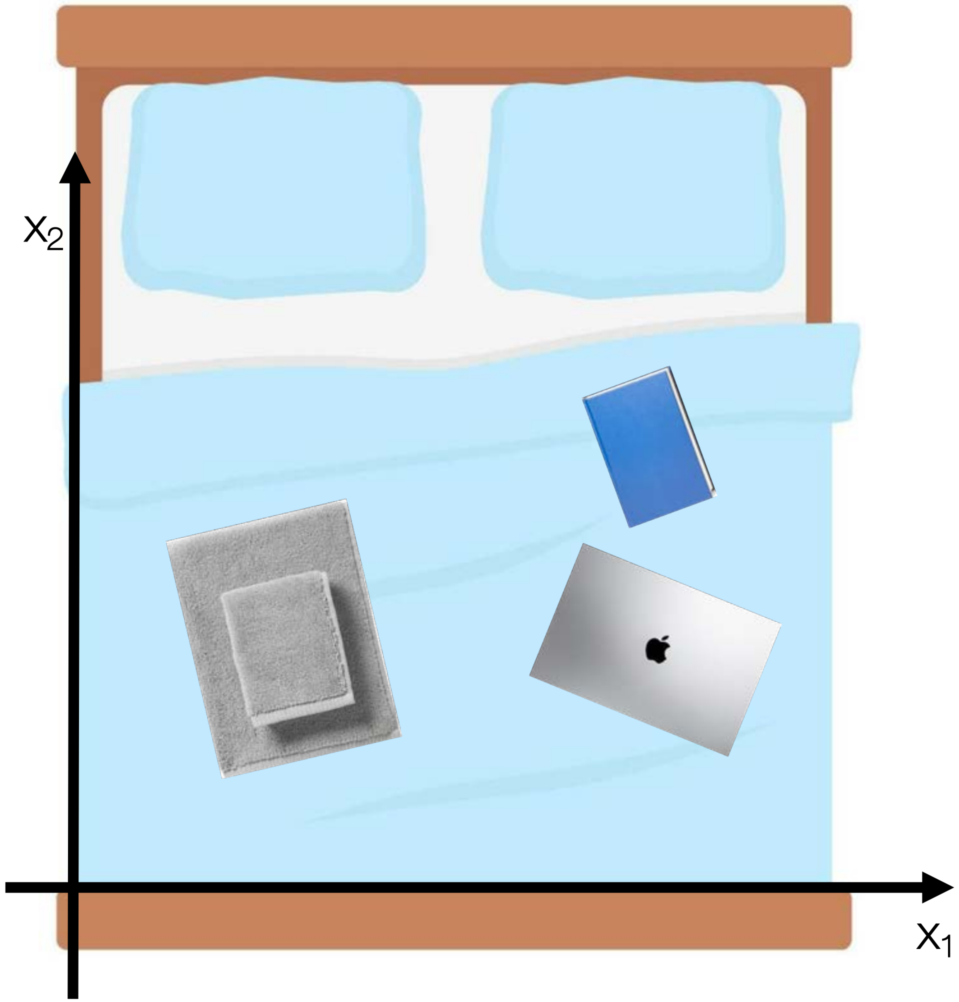

An optimization problem

Choose coordinate in an occasionally variant system that will:

- Maximize comfort "y" (temperature, softness, …)
- And (?) minimize disruption

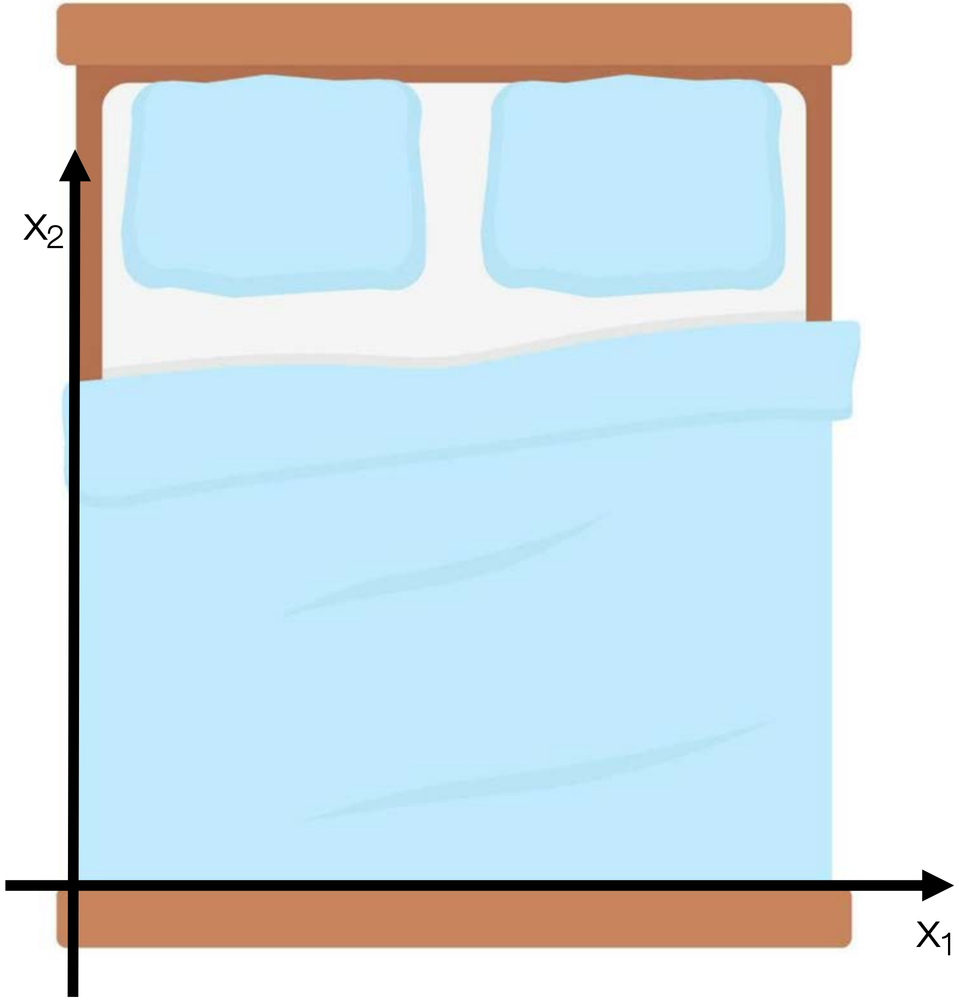

ECR ion source optimization
Choose operational space coordinate (10-20 dimensions) to:

- Maximize beam current "y"
- And maintain stability

An optimization problem

Cat optimization goal:

- Quickly assess conditions
- Predict "best" $(x_1,x_2) \rightarrow y$
- Reoptimize as conditions change

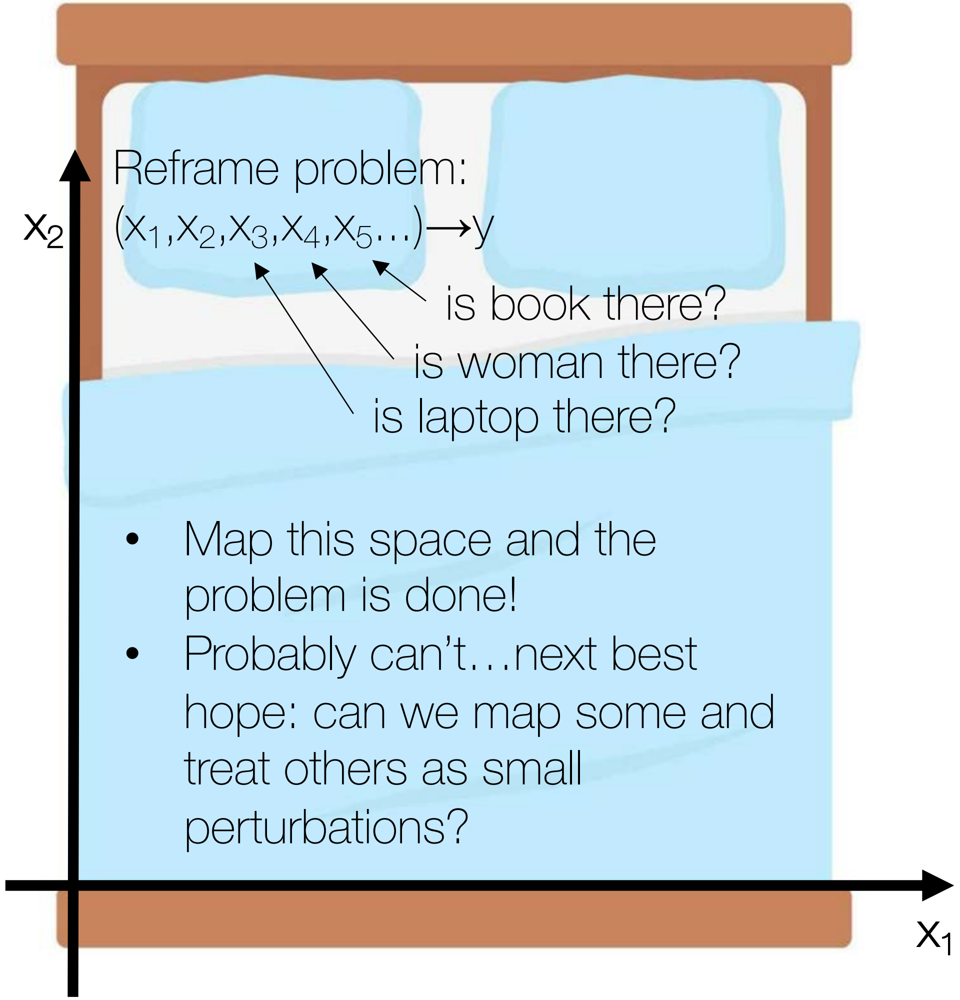

An optimization problem

Cat optimization goal:

- Quickly assess conditions
- Predict "best" $(x_1,x_2) \rightarrow y$
- Reoptimize as conditions change

ECR machine learning optimization goal:

- Quickly assess conditions
- Predict "good" control parameters: $(x_1, x_2, \ldots x_n) \rightarrow y$
- Reoptimize as conditions change

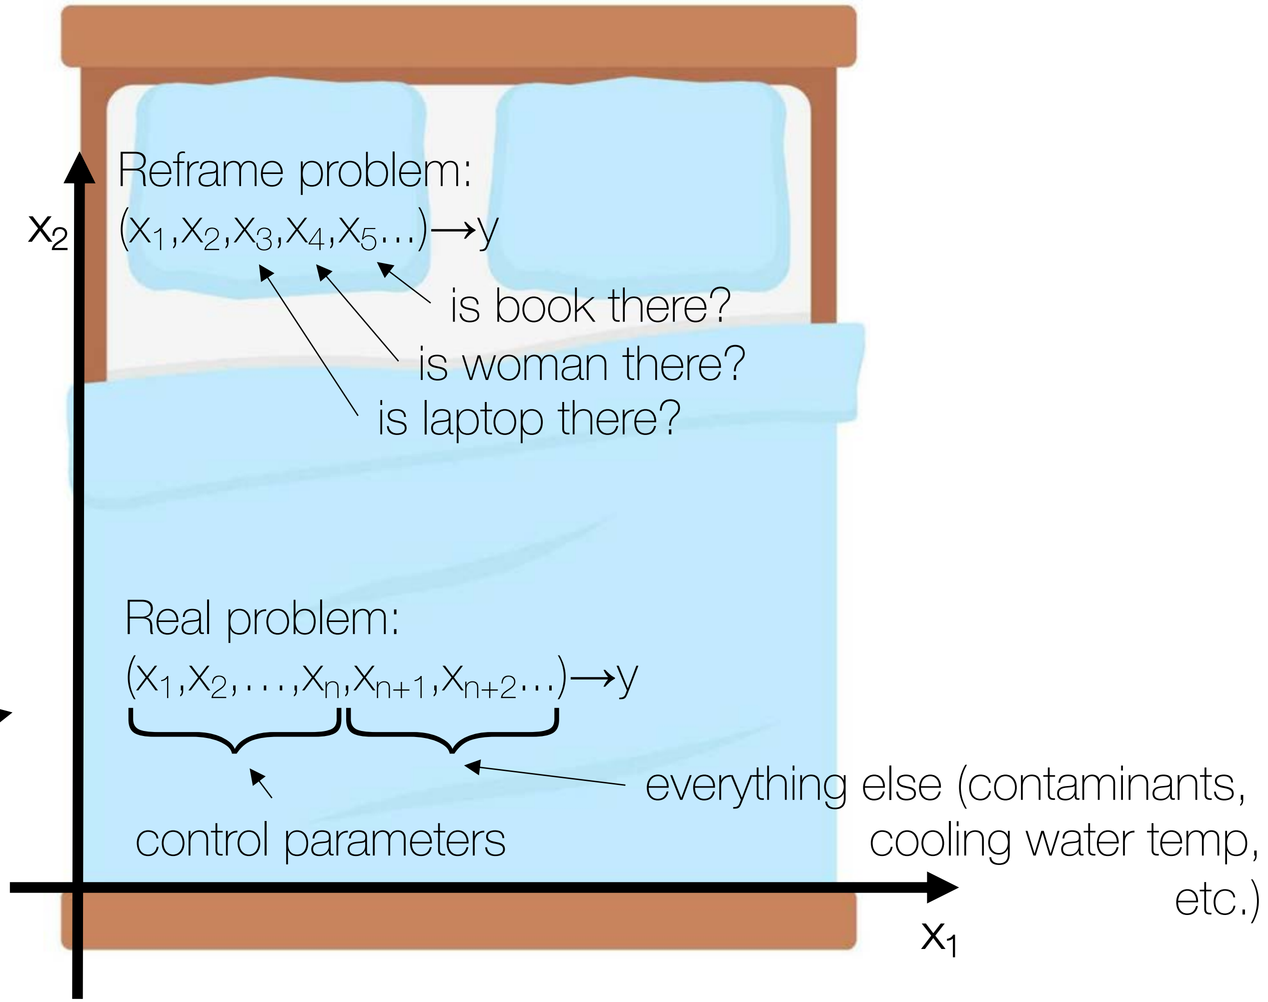

We are attempting to go at this with VENUS

VENUS:
- World's first fully-superconducting ECR ion source designed for 28 GHz operation
- Injector for LBNL's 88" Cyclotron
- RF: 28+ 2x18 GHz (up to 10+2x2 kW)
- Maximum axial fields: 4 T

Example beams:
- > 4.7 mA O^{6+}, > 20 mA He^{+} from source
- $^{197}Au^{61+}$ extracted from cyclotron
- > 2 pμA $^{48}Ca^{11+}$ and > 1.7 pμA $^{50}Ti^{12+}$ 5 MeV/u beams from Cyclotron for superheavy element research

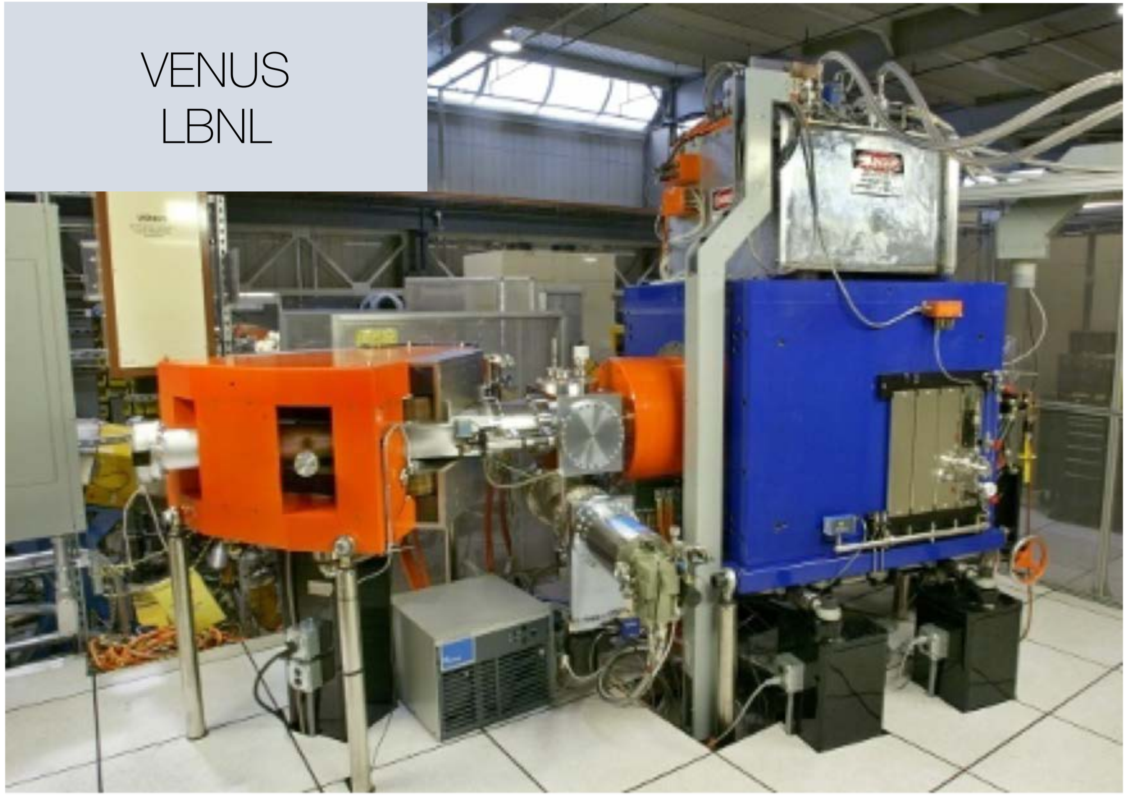

VENUS operation and data collection

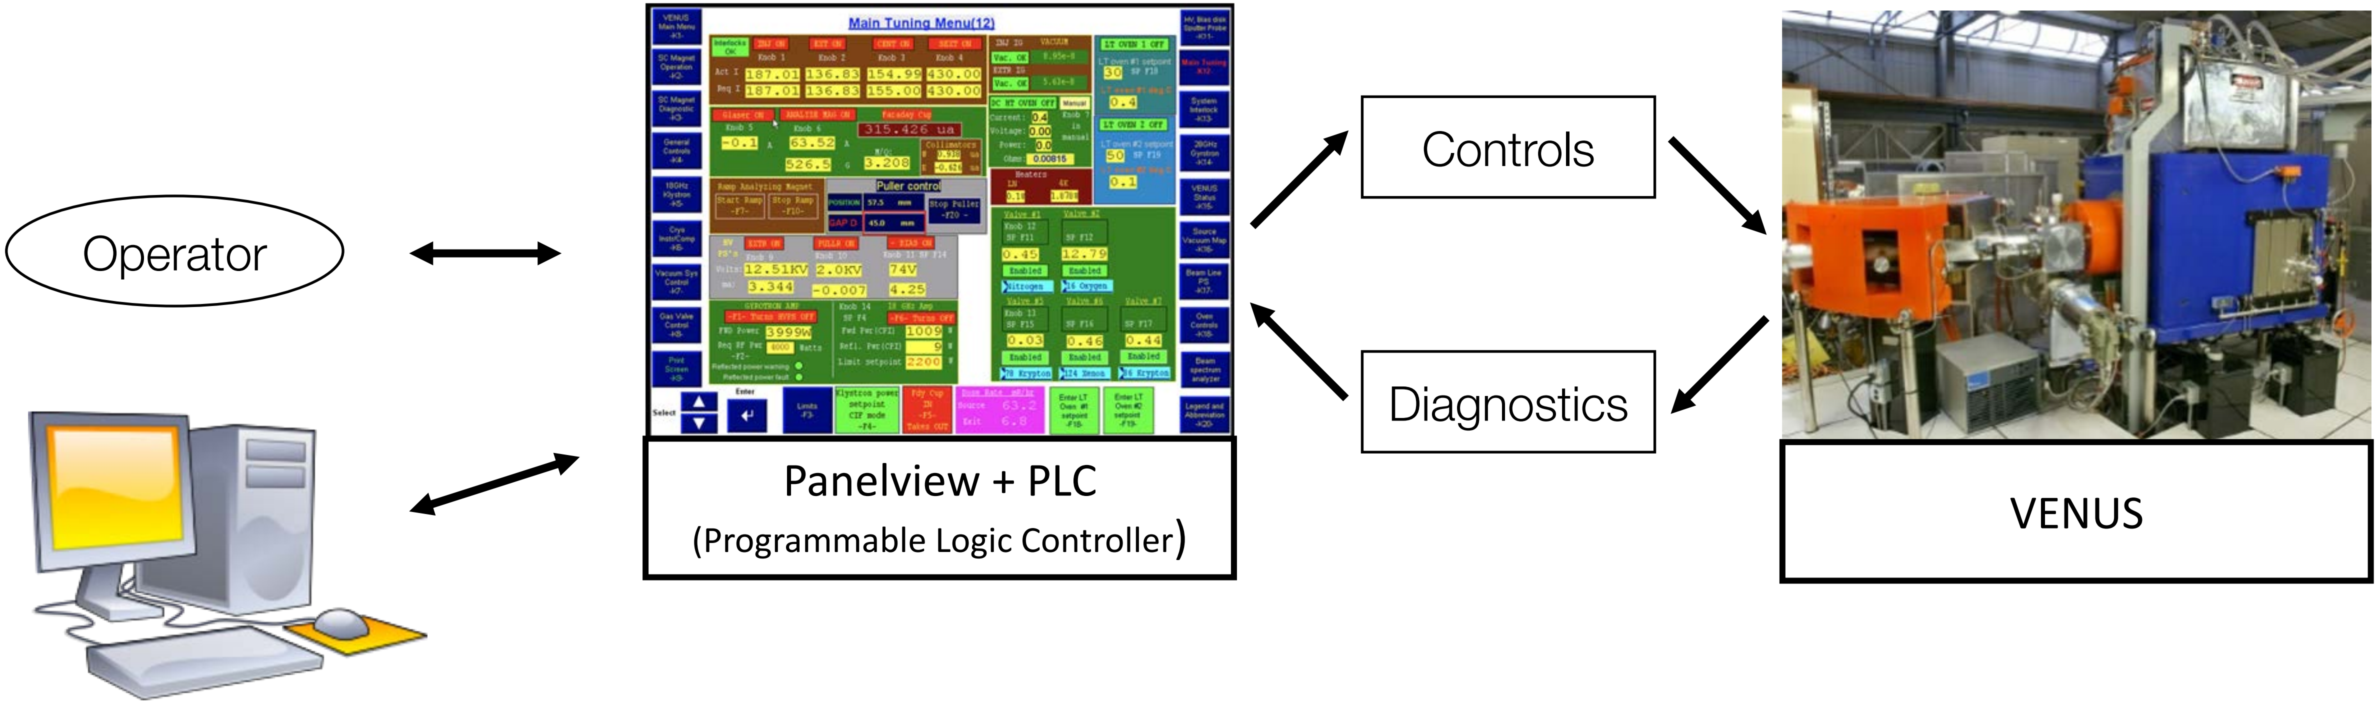

Computer control through PLC:

Advantage:	Exploit 2 decades of safely logic in PLC
Disadvantage:	Slow ~3 Hz communication

What we have done

- Collected over 4 years of all VENUS control and primary diagnostic data at 1 second intervals
- Performed thousands of hours of computer-only control and optimization of VENUS without human interaction

This work is supported by the U.S. Department of Energy, Office of Science, Nuclear Physics program under Award Numbers DE-FOA-0002490 and DE-FOA-0002875.

Problem: find extremum of an unknown function when measurements are expensive

Find maximum of unknown function:

- Measurements are expensive (time, $, etc.)
- Four points have been measured
- Where to search next?

Method:

- Fit curve between points
- Find maximum
- Repeat

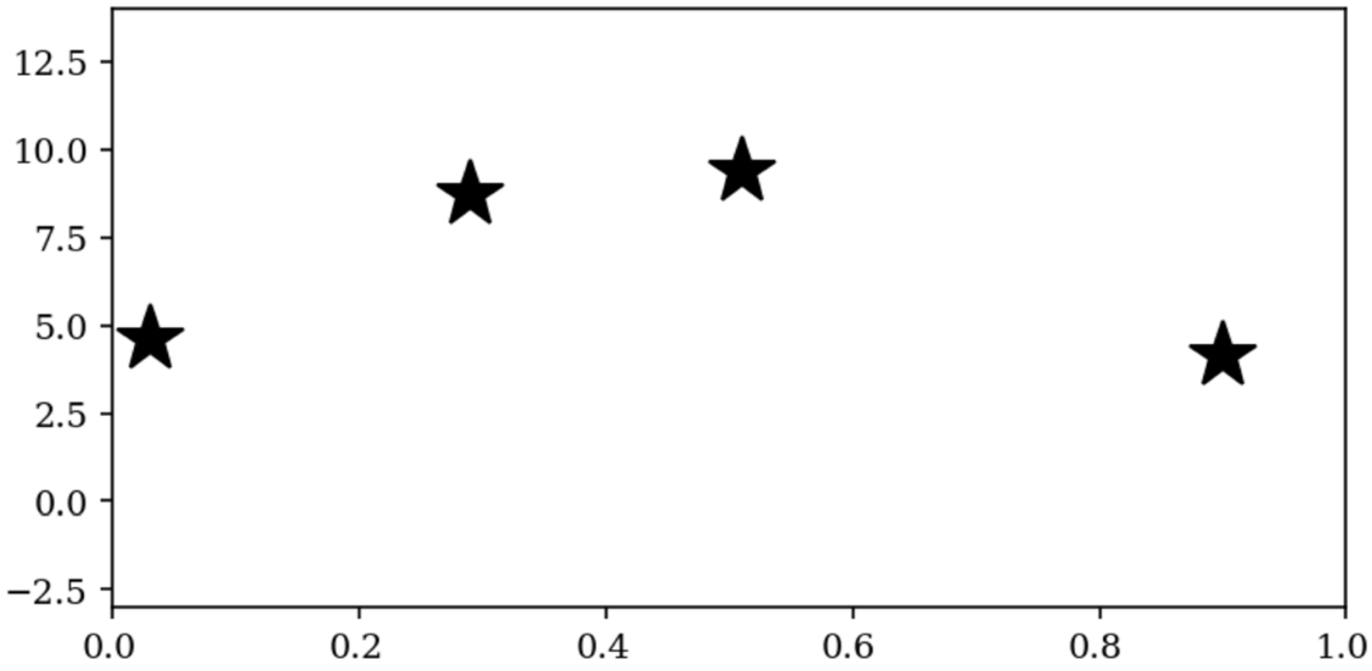

Problem: find extremum of an unknown function when measurements are expensive

Find maximum of unknown function:

- Measurements are expensive (time, $, etc.)
- Four points have been measured
- Where to search next?

Method:

- Fit curve between points
- Find maximum
- Repeat

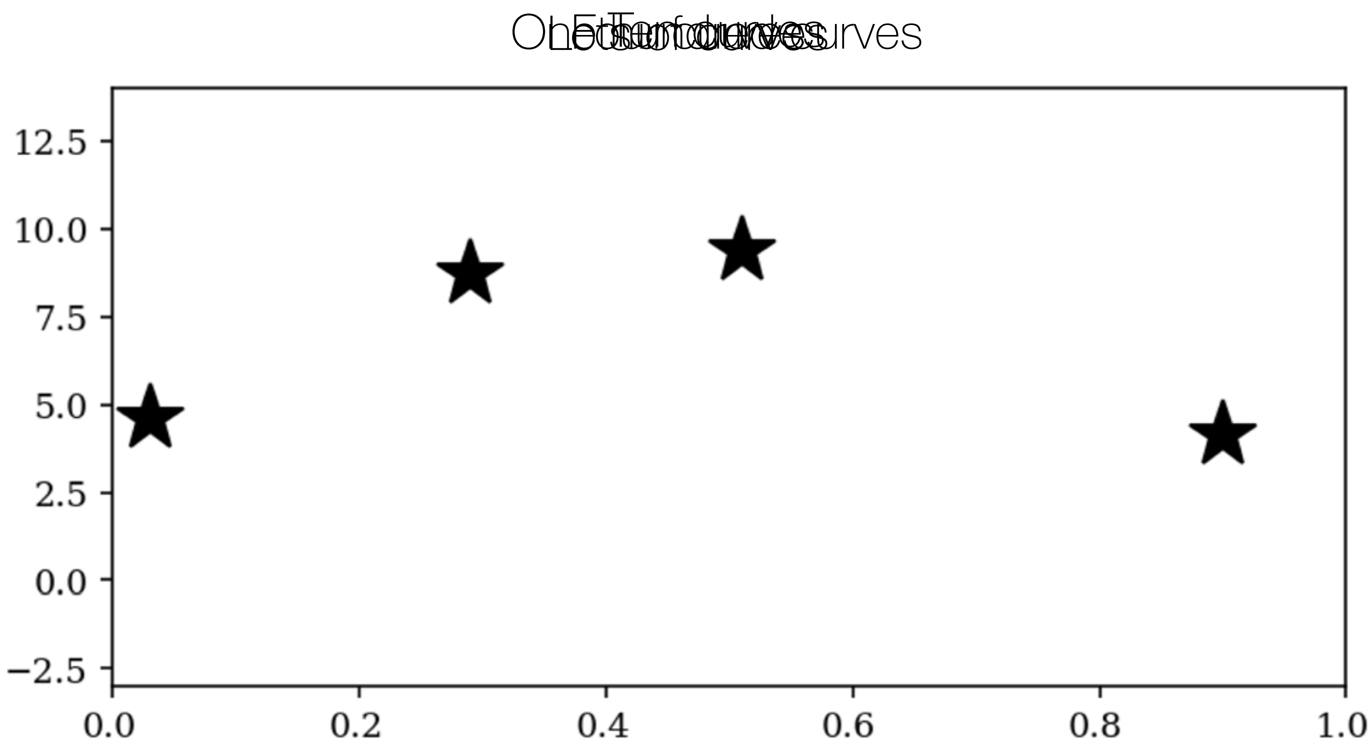

Bayesian optimization basically follows a path like this

Notes:

- Next search point has highest probability of improvement though other points may have curves with higher extrema

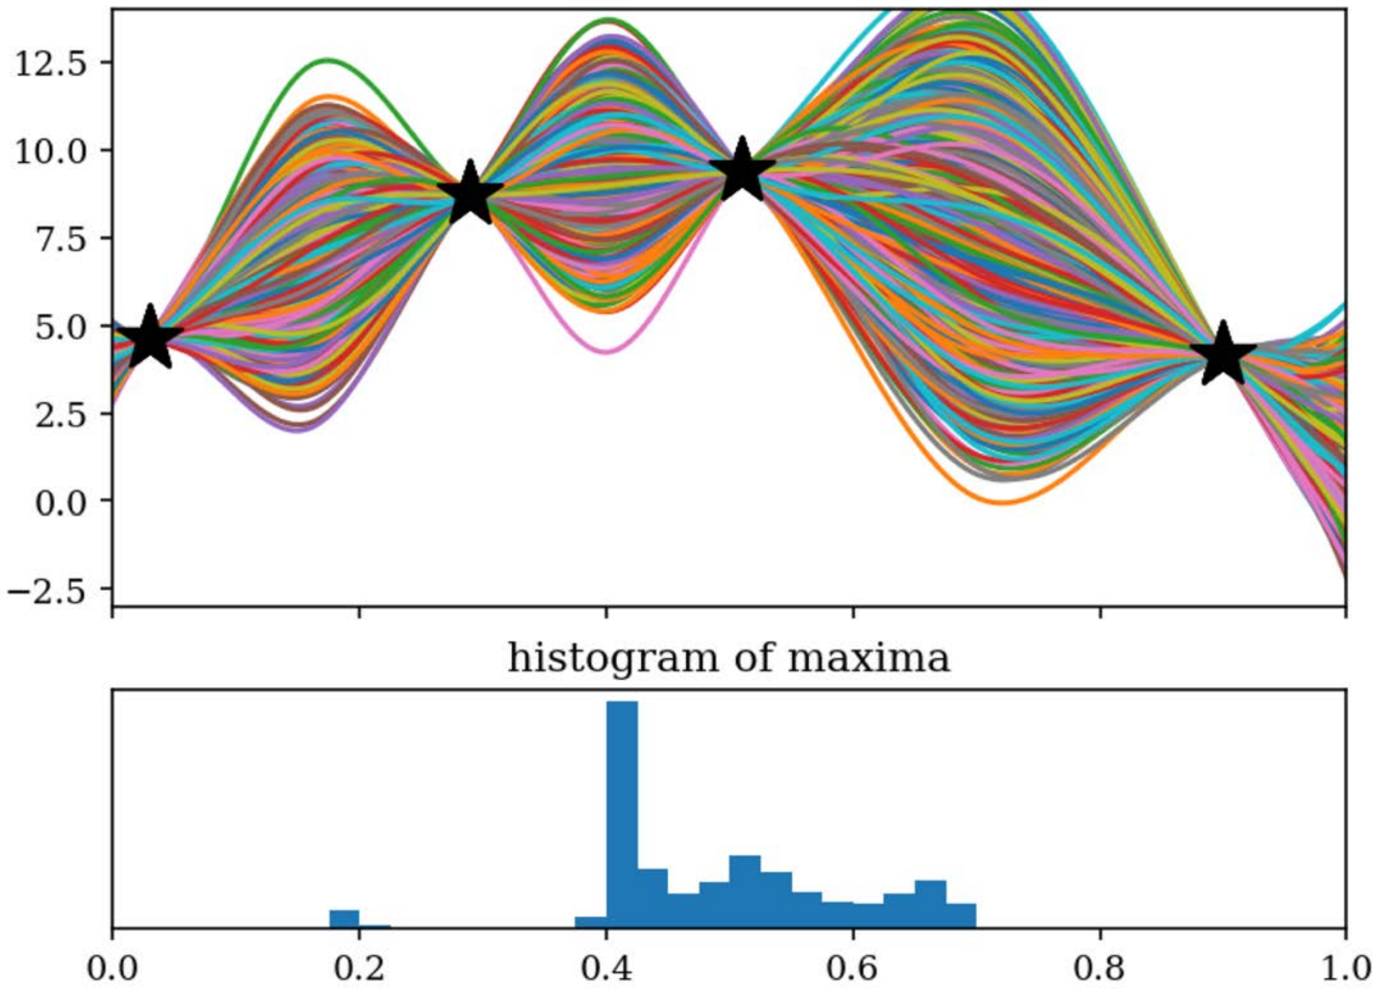

VENUS primary control and diagnostic parameters

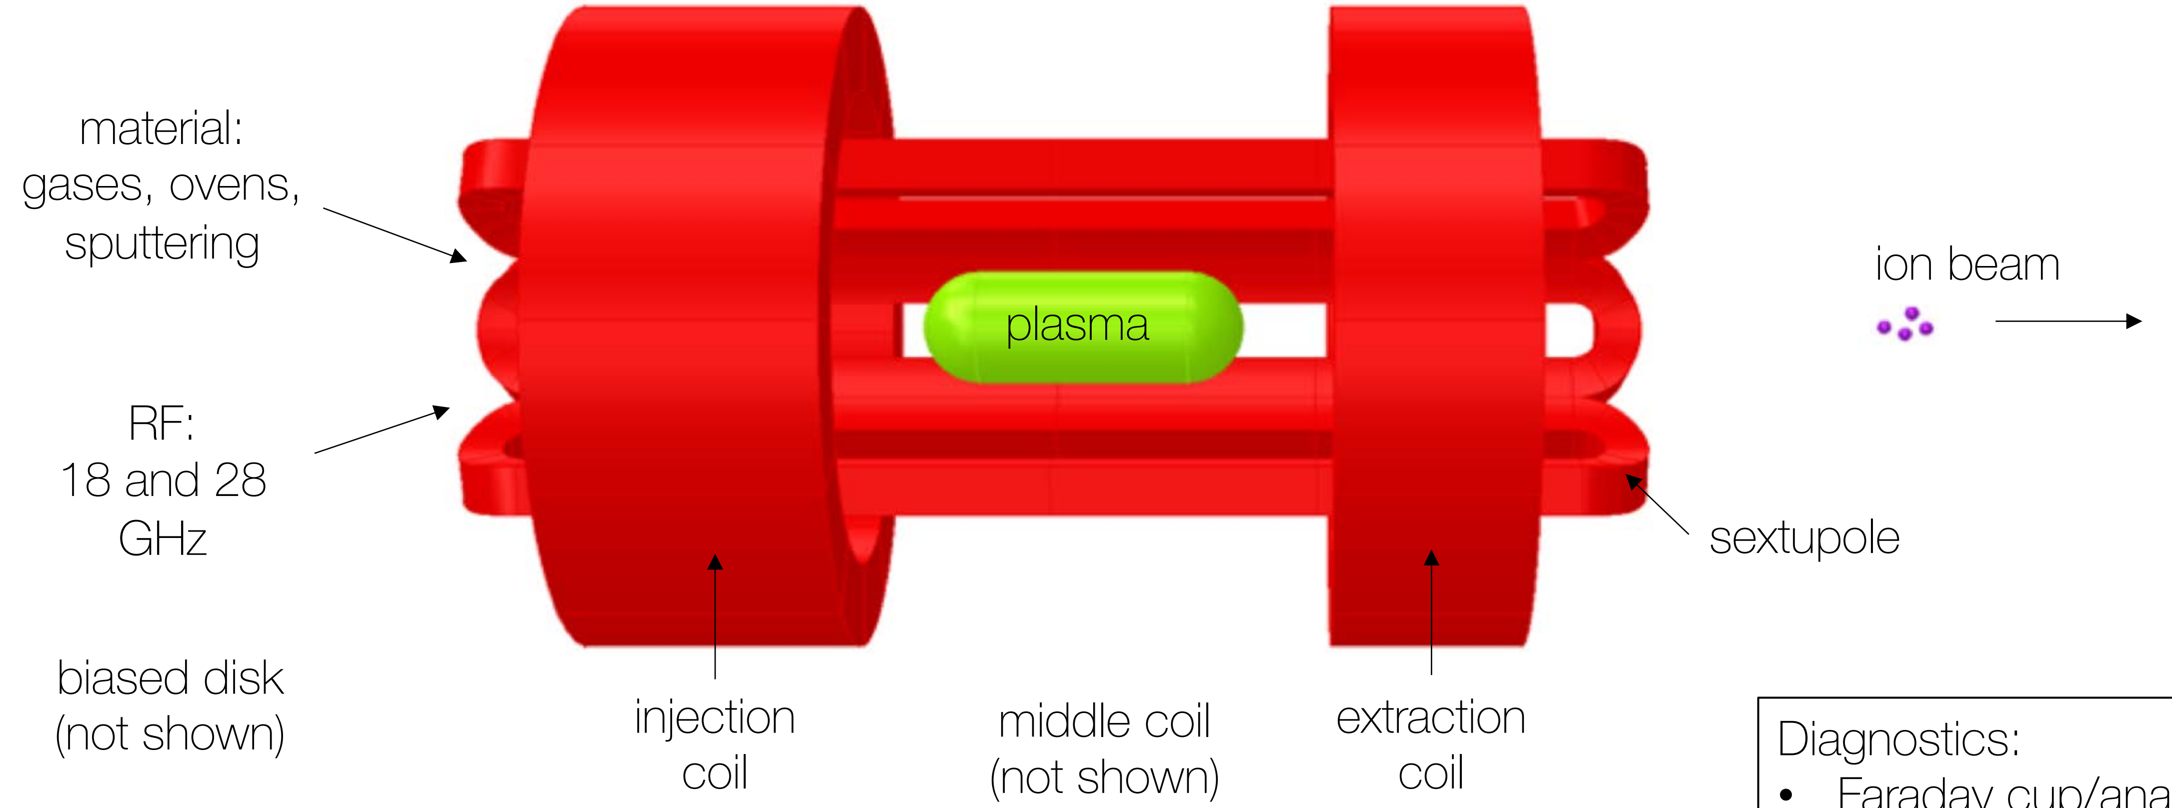

Diagnostics:
- Faraday cup/analyzing magnet
- Emittance scanner
- Drain, bias currents
- Cryostat x-ray load
- Etc.

Machine Learning: Full Bayesian Optimization of $^{124}Xe^{37+}$

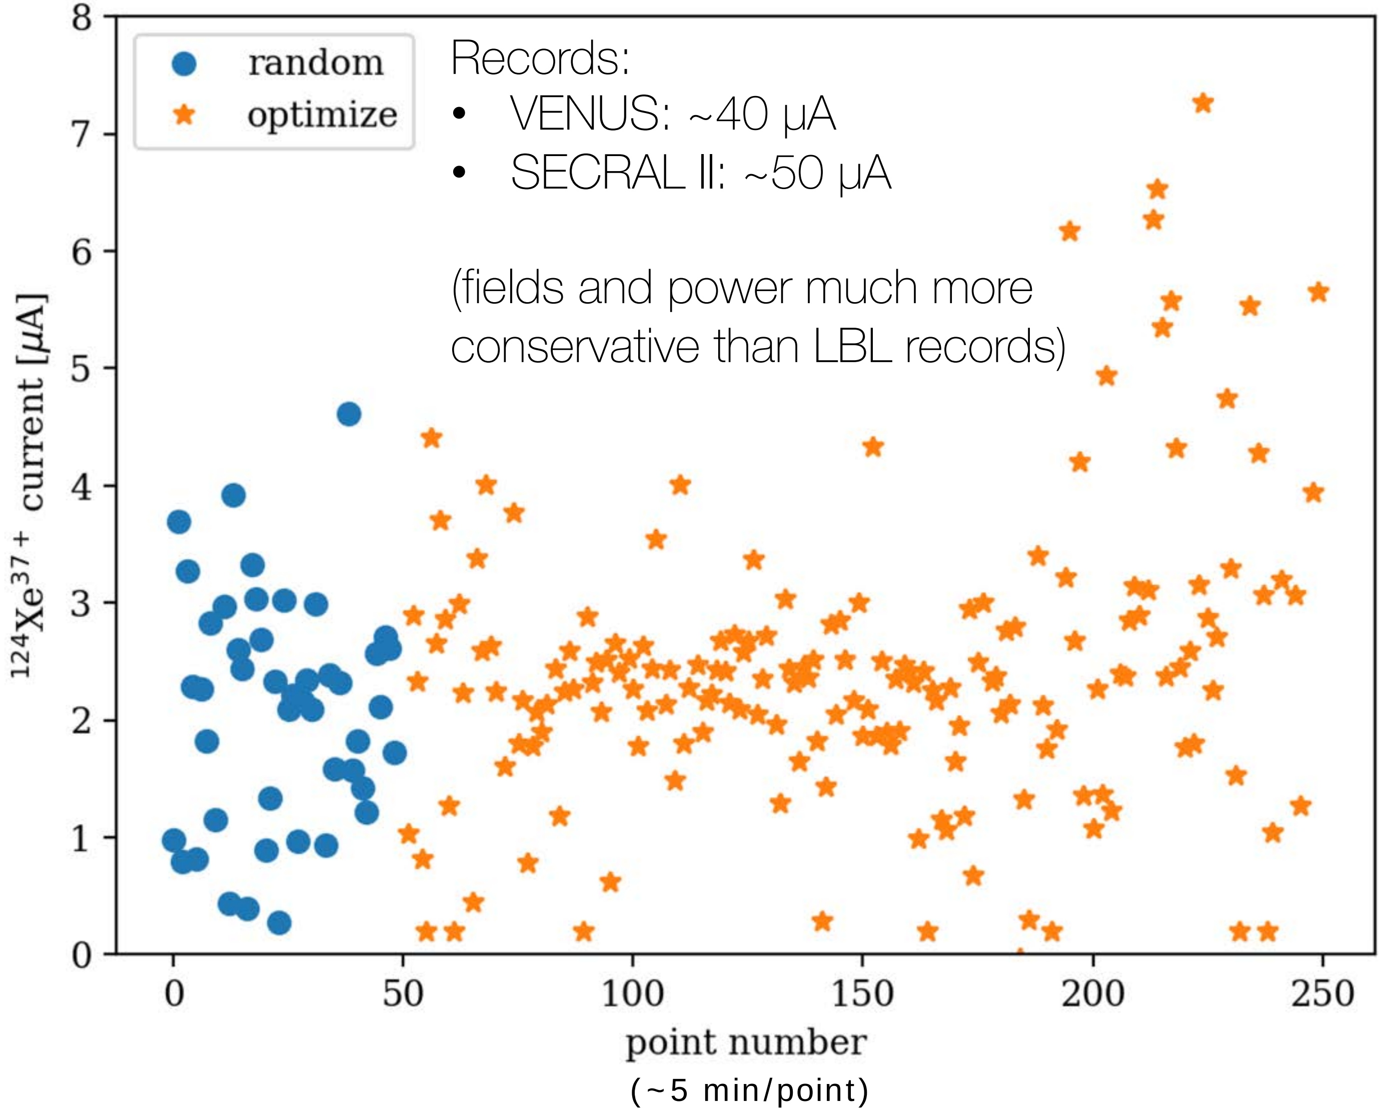

Parameter	Min	Max
Bias voltage [V]	40	105
Oxygen valve	11.6	12.5
Xenon valve	8.0	13.0
Inj coil [A]	185.6	186.0
Ext coil [A]	136.6	136.8
Mid coil [A]	152.0	152.3
Sext coil [A]	430.3	430.5
18 GHz [kW]	1.4	1.8
28 GHz [kW]	5.2	6.0

- VENUS completely under computer control
- Computer “knows” nothing about VENUS

Augmenting search based on cost

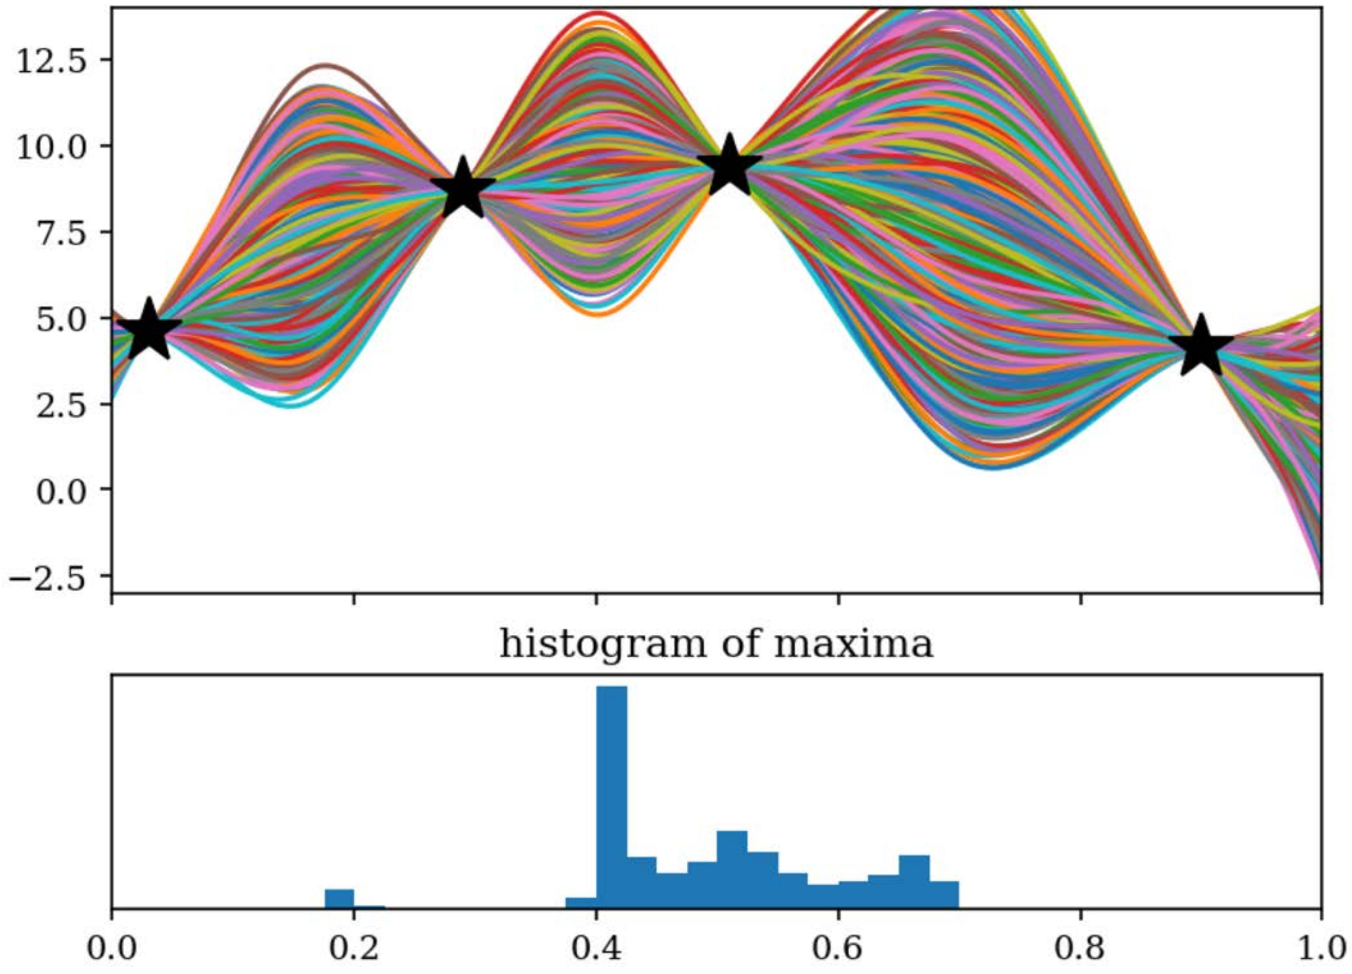

- Say cost of search increases with x
- We could create a "cost function" where we only search a large x point if the likelihood of improvement is over some threshold

For superconducting sources:

- Coils: Δt_{change} ➞ minutes
- Bias: Δt_{change} ➞ seconds

We have employed cost functions to judiciously use coil changes

Using cost function in VENUS optimizations

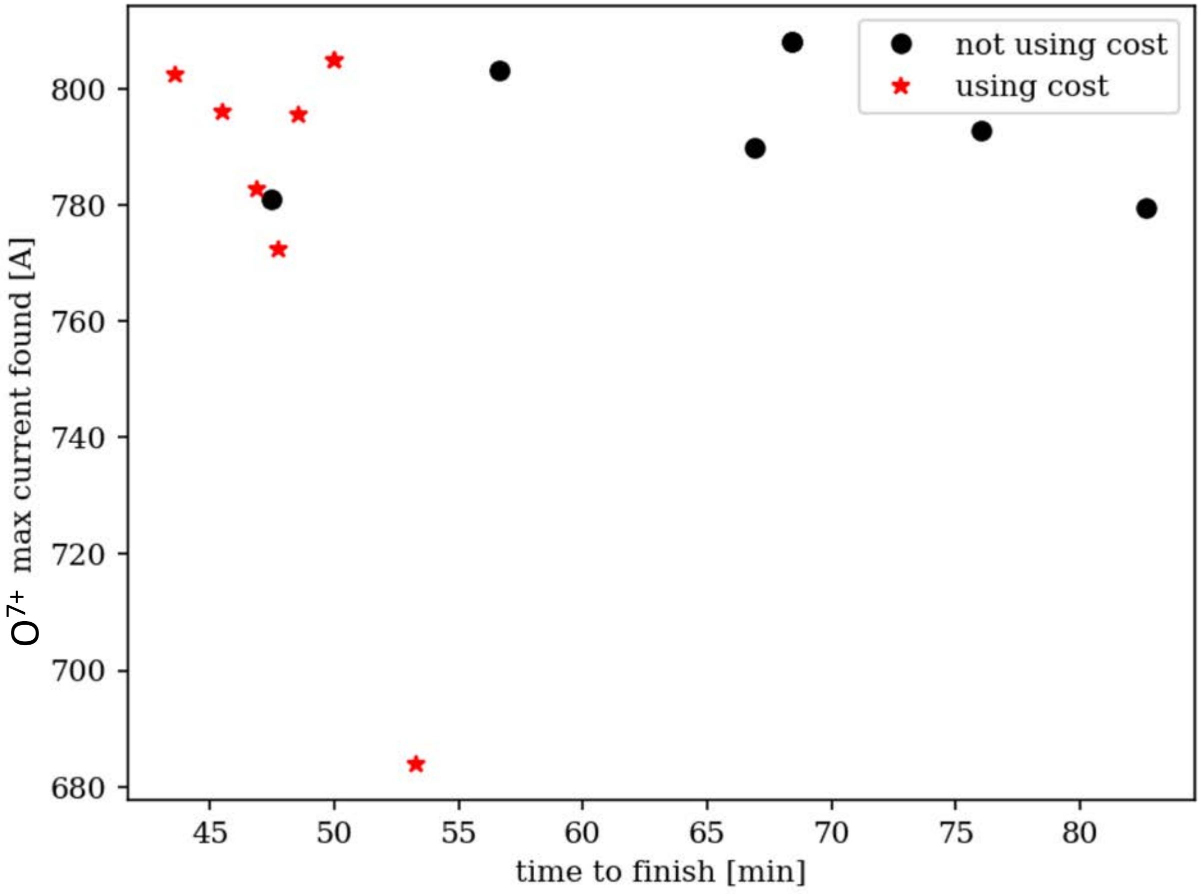

Full optimization time

Proof of principle:

- Three parameter optimization of O^{7+} (biased disk, middle solenoid, and oxygen valve) can be tens of minutes faster using cost function
- Peak current found can be lower as it is a less-thorough search

We will continue to explore this optimization technique

Optimizing a little more like a human

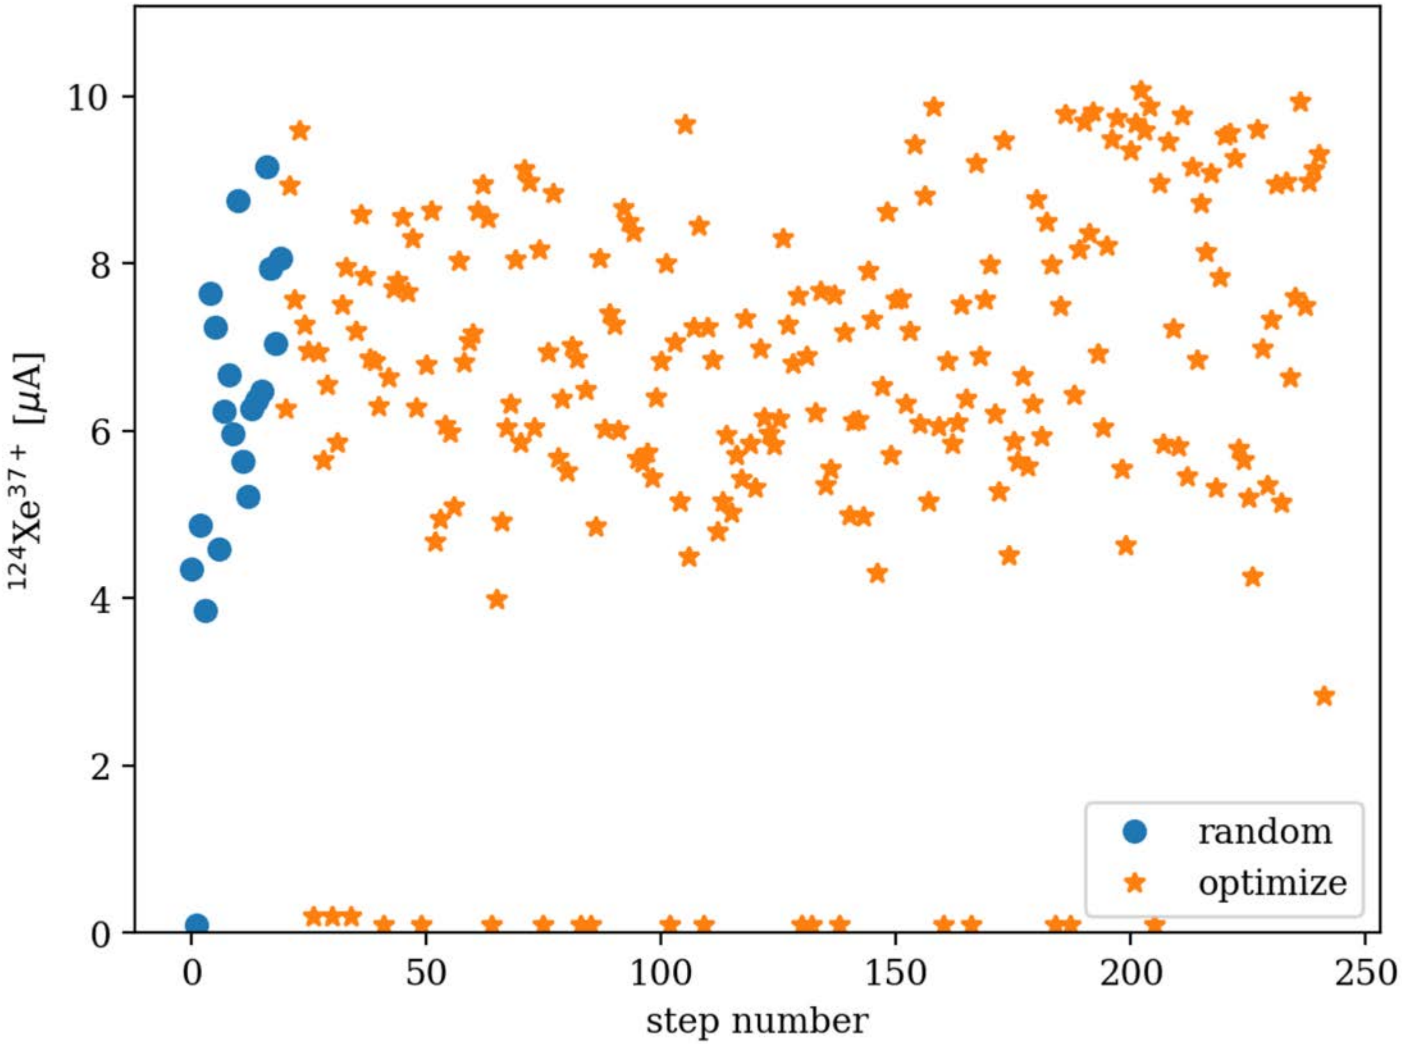

Parameter	**Min**	**Max**
Bias voltage [V]	25	65
Oxygen valve	11.5	12.0
Xenon valve	9.0	12.0
18 GHz [kW]	1.20	1.80

Many human records are achieved by using a cost-function-like approach:

- **Coils are slow, so find a pretty good solution and work from there**

Note:
When instabilities or too low pressures encountered, low current is recorded

Are there trends we can see?

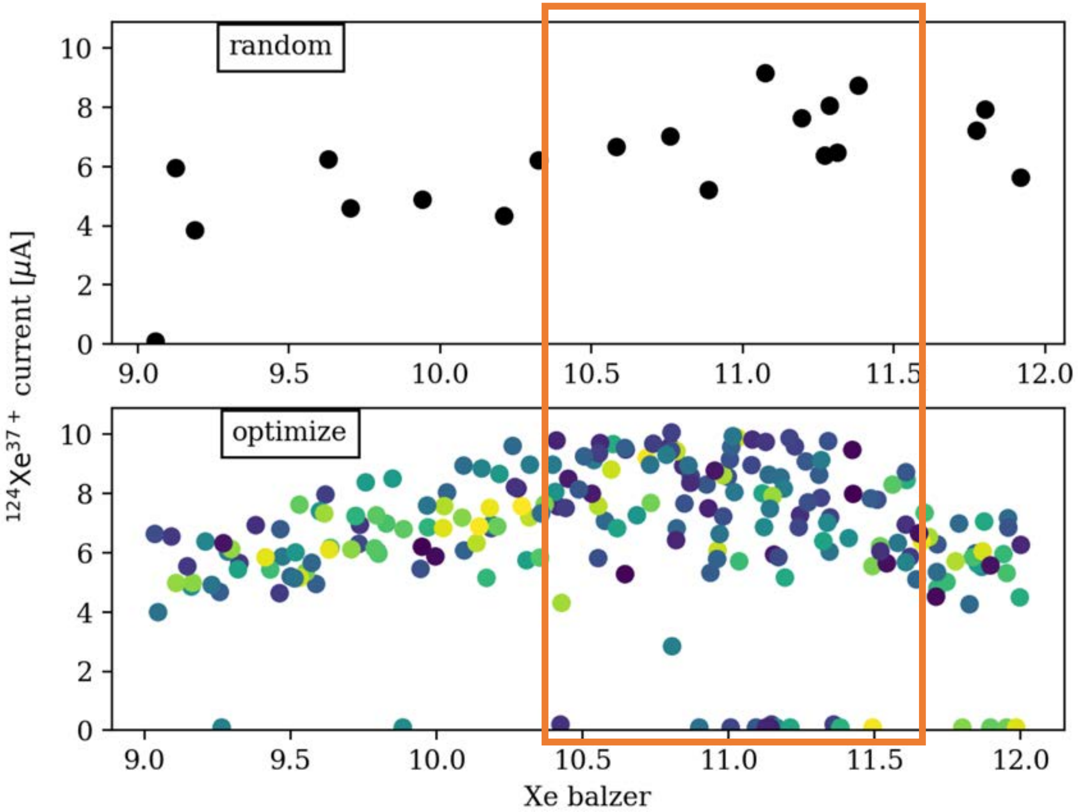

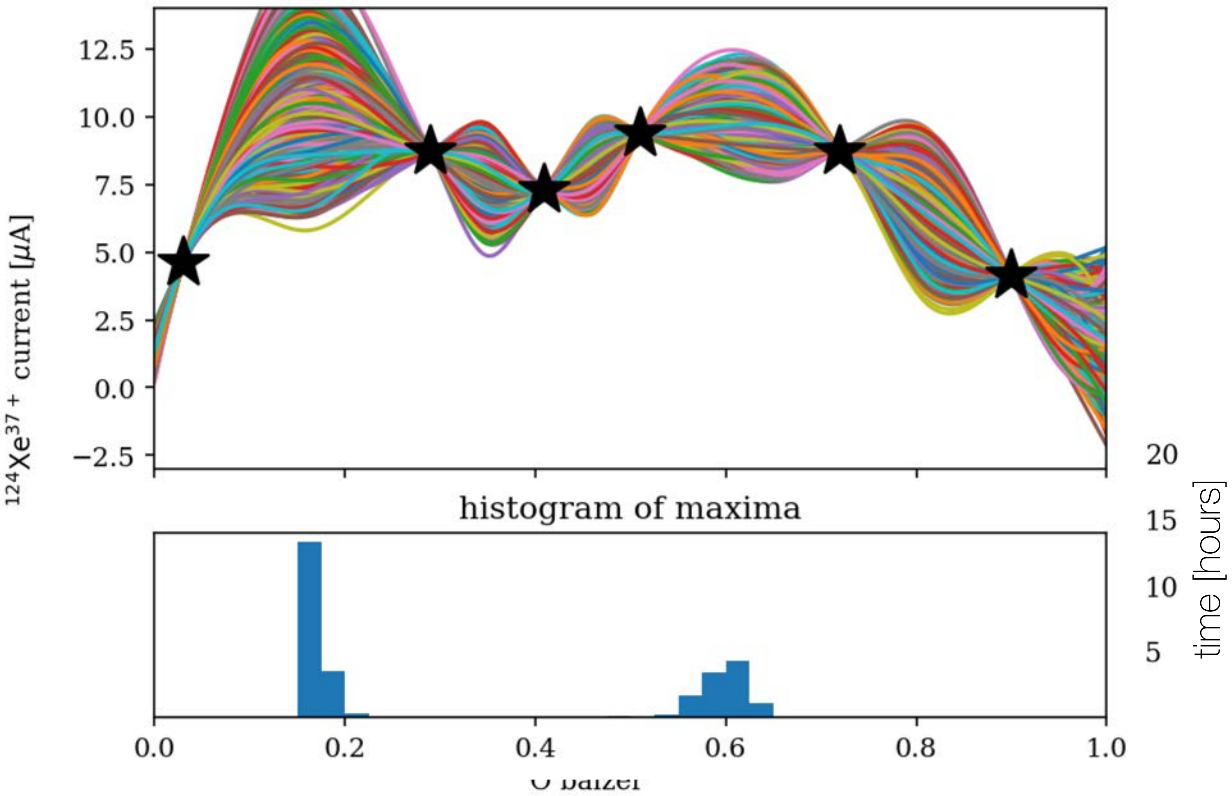

- Code not being as smart as a person would be
- Can change this by adjusting "exploration" vs. "optimization"

A deeper dive into optimization

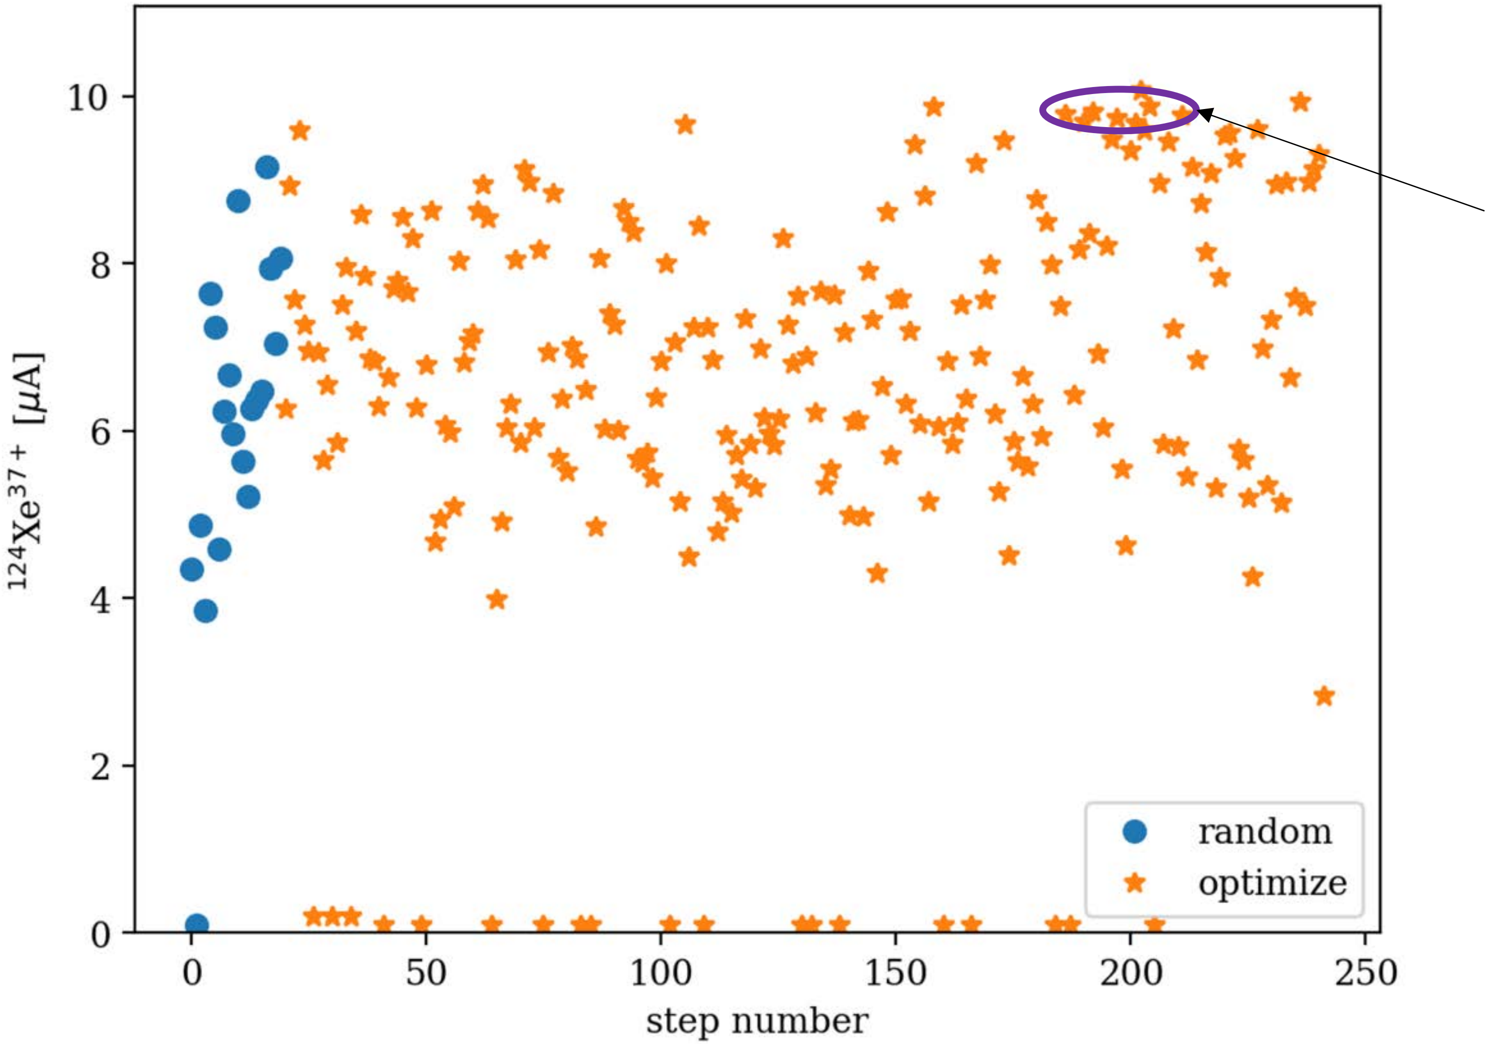

Investigate these more closely

How were these measurements performed?

- Optimization code is looking at only one charge state (this is also what human operators spend most of their time doing!)
- Sweep dipole to get charge state distribution (CSD)

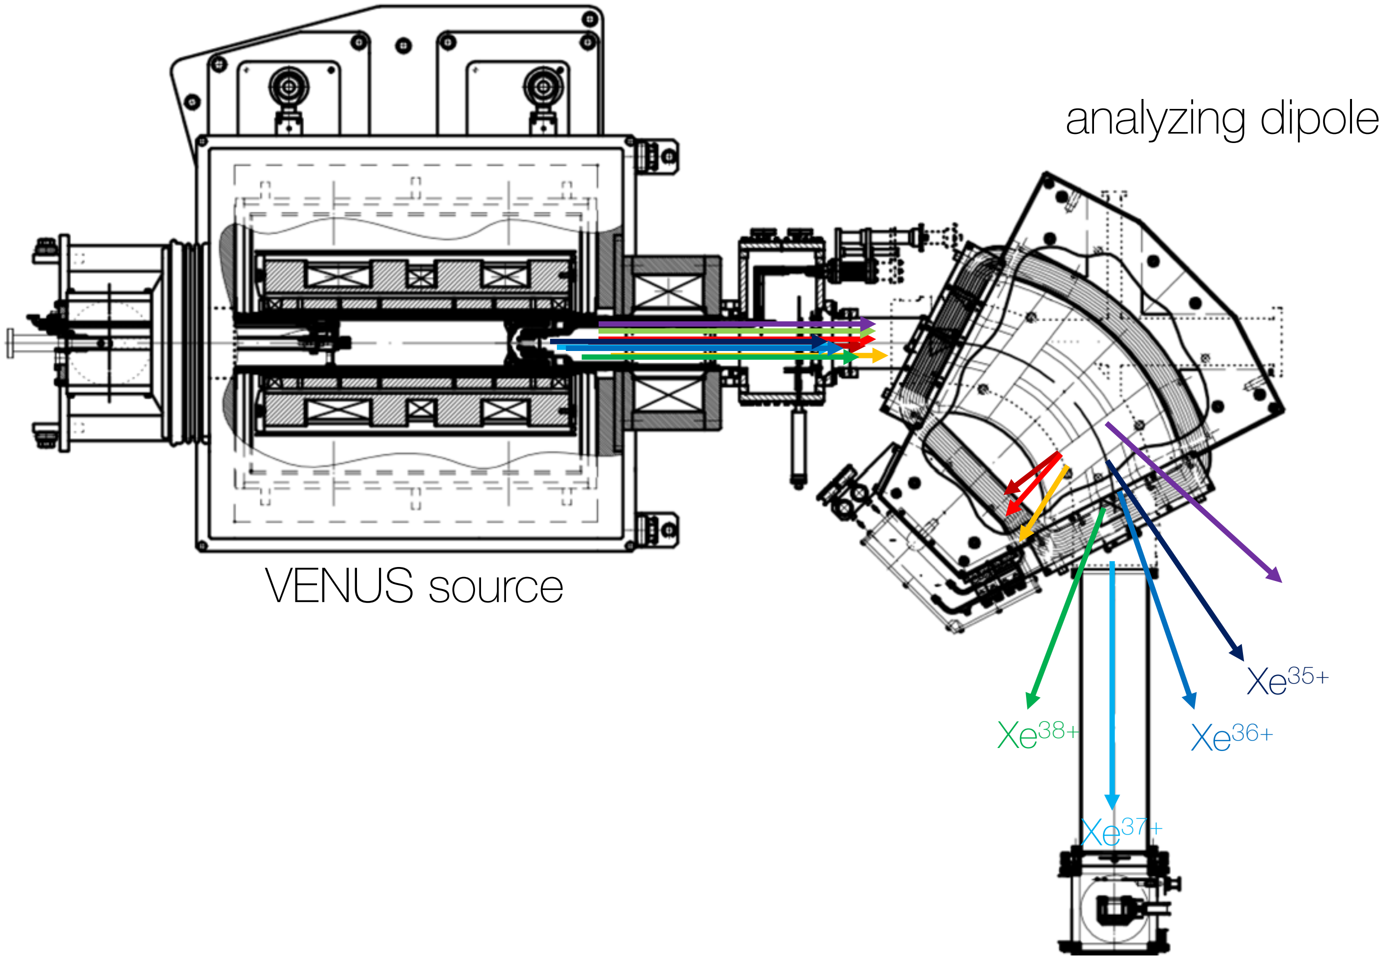

Many ways to same result

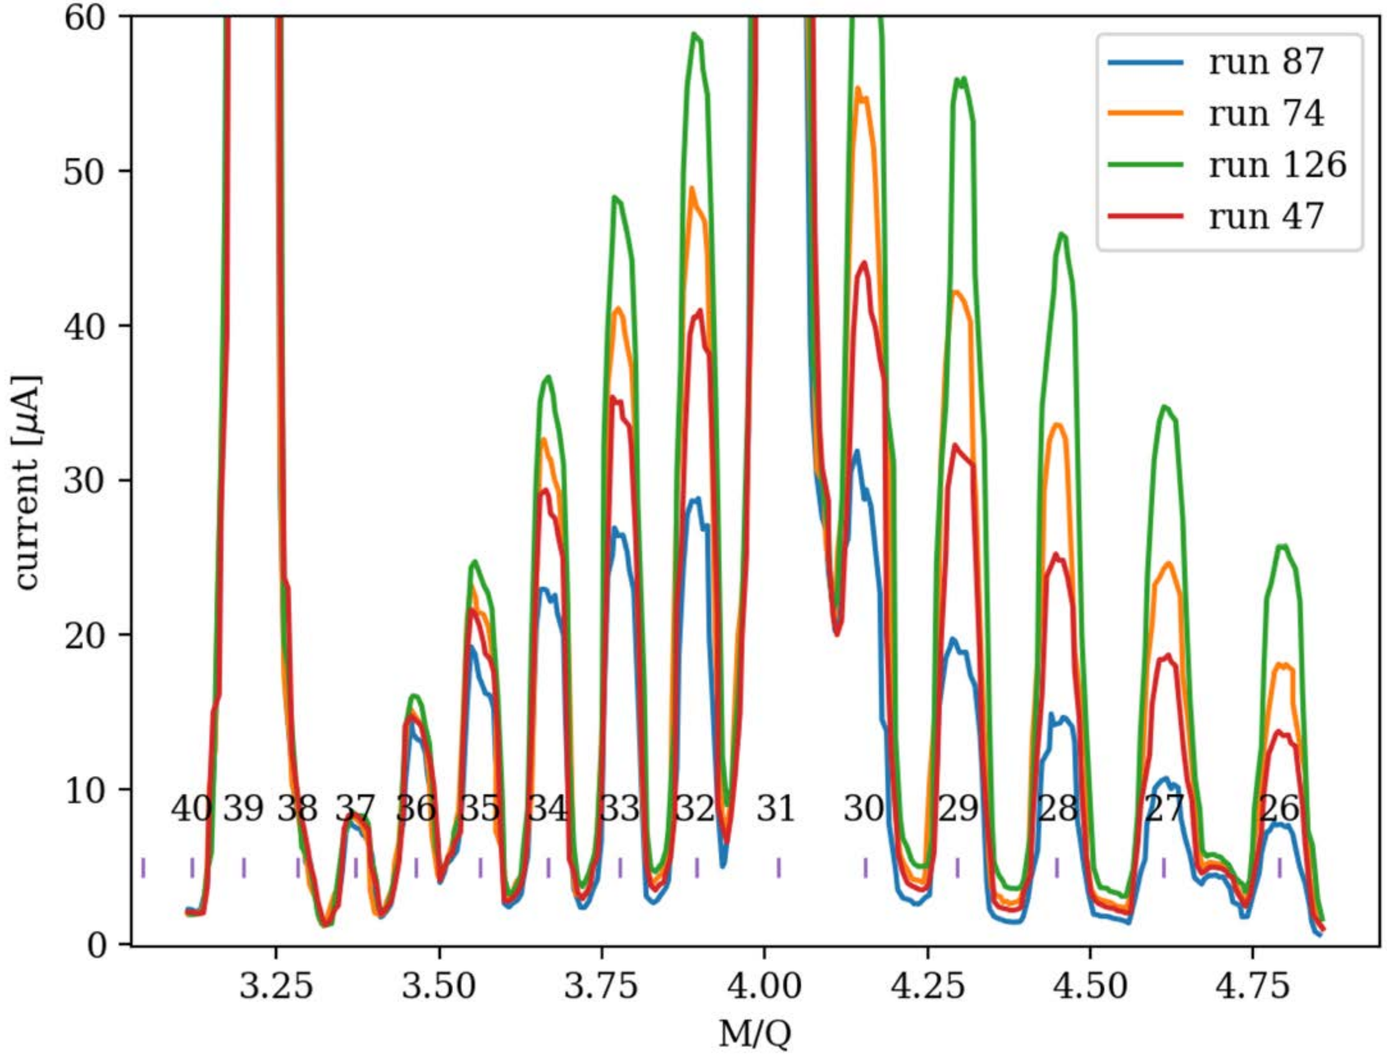

CSDs show there there are many ways to optimize a specific charge state

Takeaways:

- optimizing charge state's current without CSD knowledge is restricting
- CSD is slow: ~2-3 minutes each
- Even beam statistics are slow: ~3 Hz

Goals after initial optimizations (September 2023):

- Let computer try for record beam ✗
- Speed up data gathering ✓

Faster beam measurements

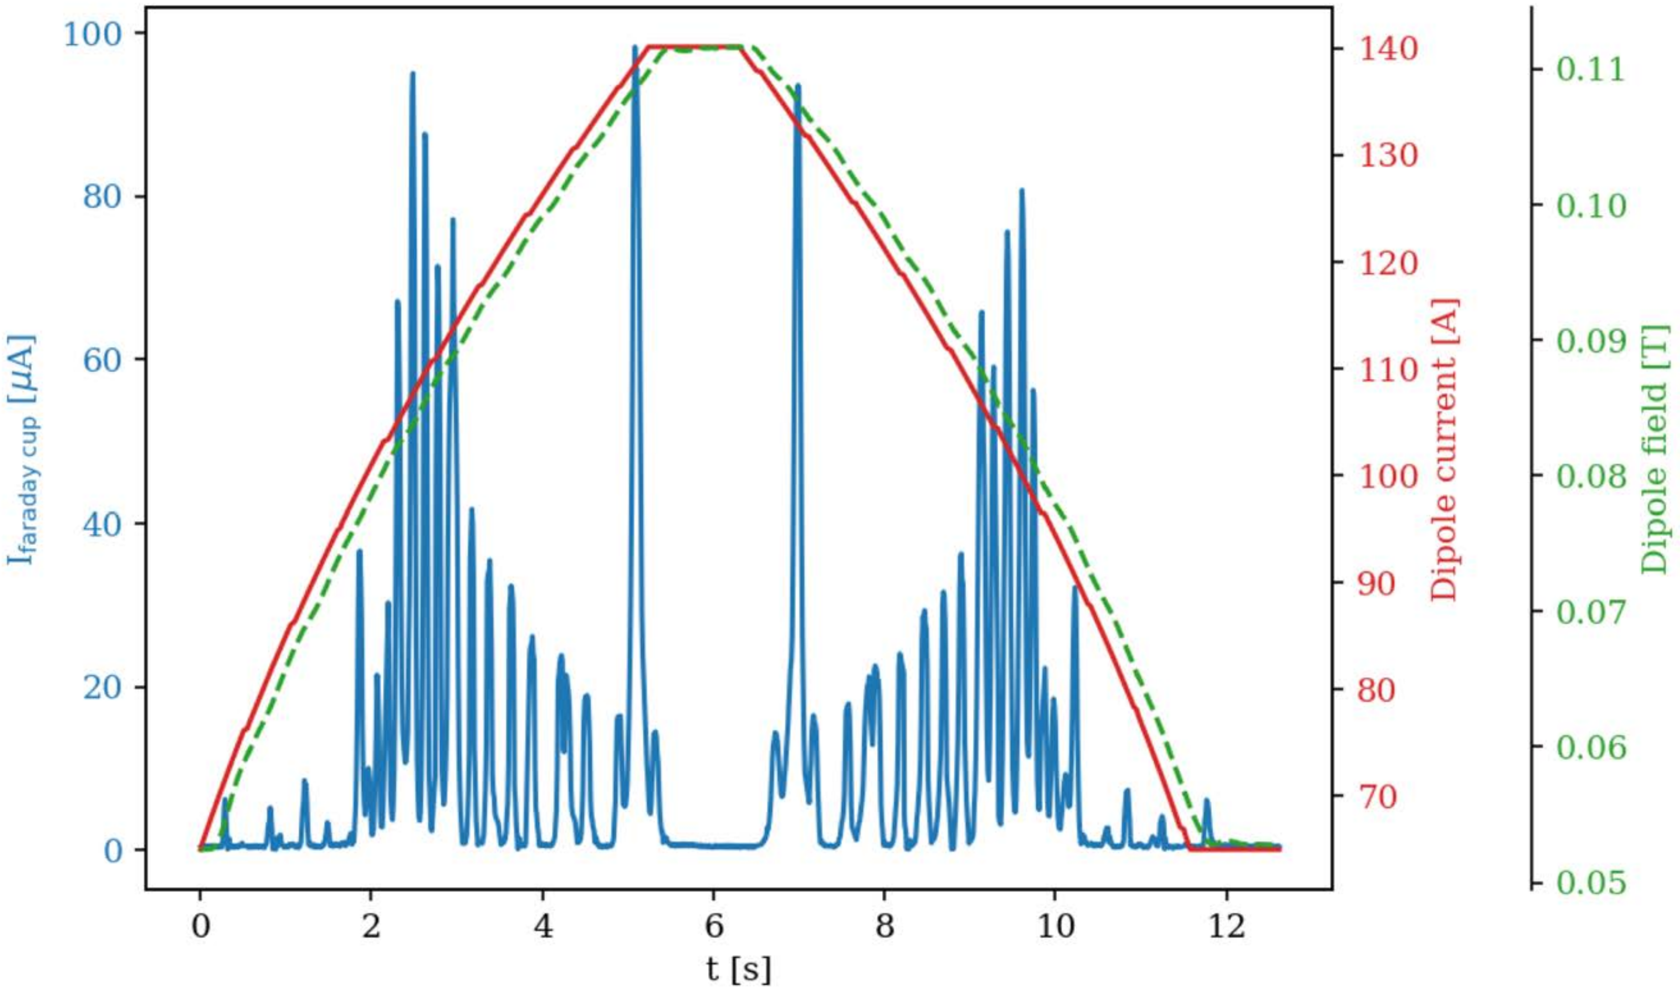

At 100 Hz:

- Set dipole current
- Read dipole's hall probe
- Read beam current

Faster CSDs

- agree well with slower ones

Unwrapping fast CSDs

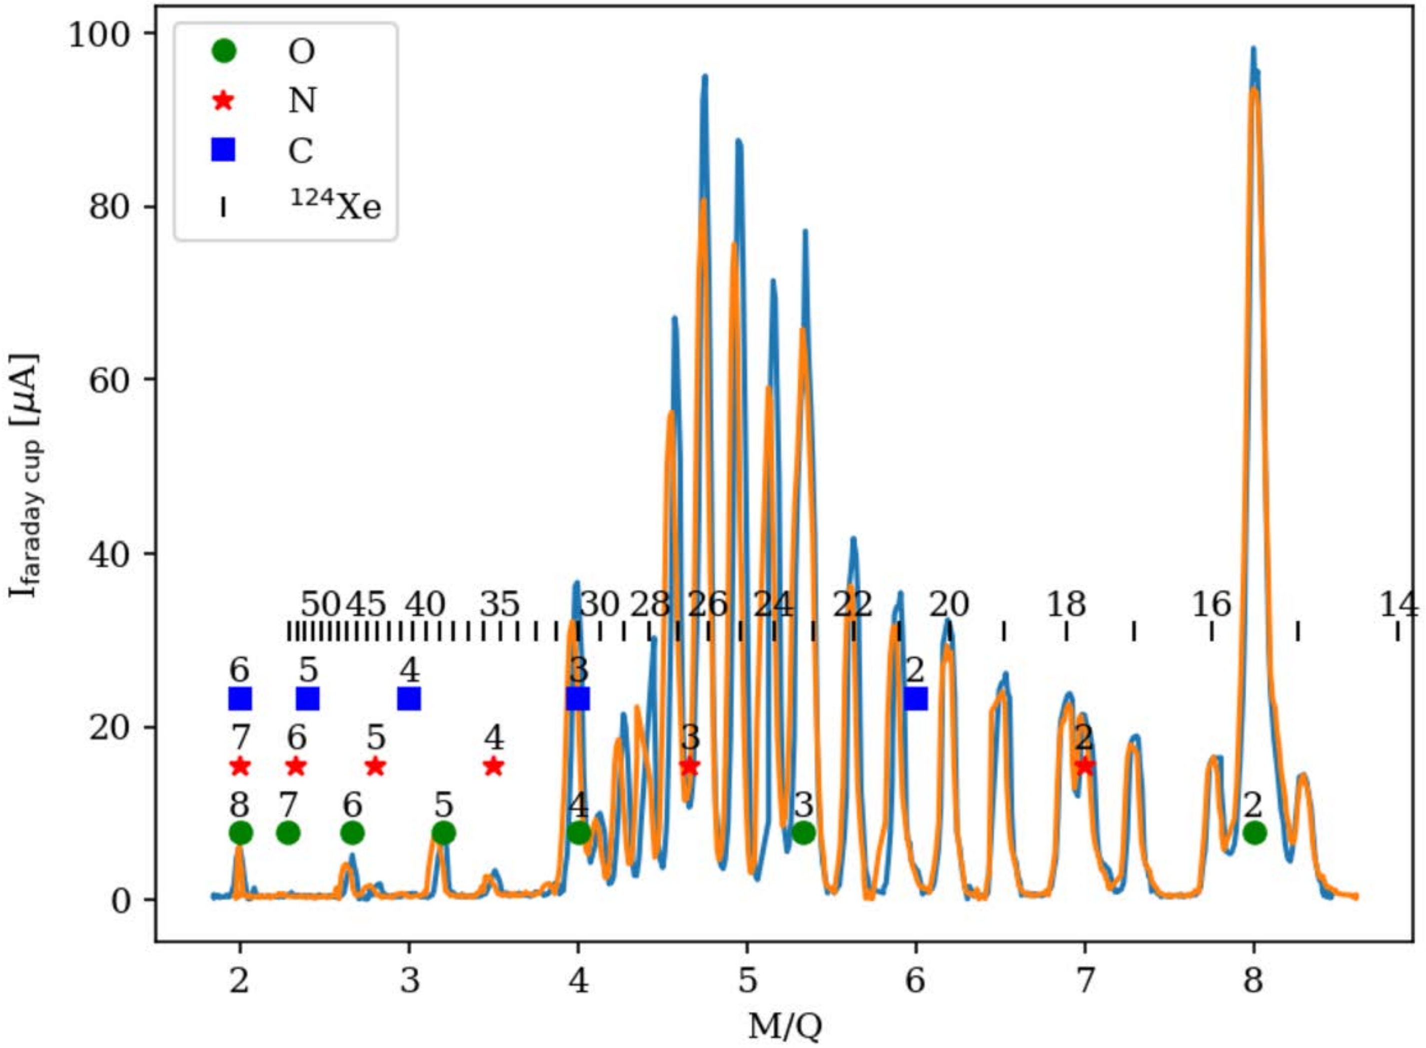

- Reduction in measured current on return, especially for high charge states
- Reduction recovers by the next CSD sweep
- Use only "increasing current" CSDs for comparison

Overbent ions are tightly spaced on beam pipe: heating

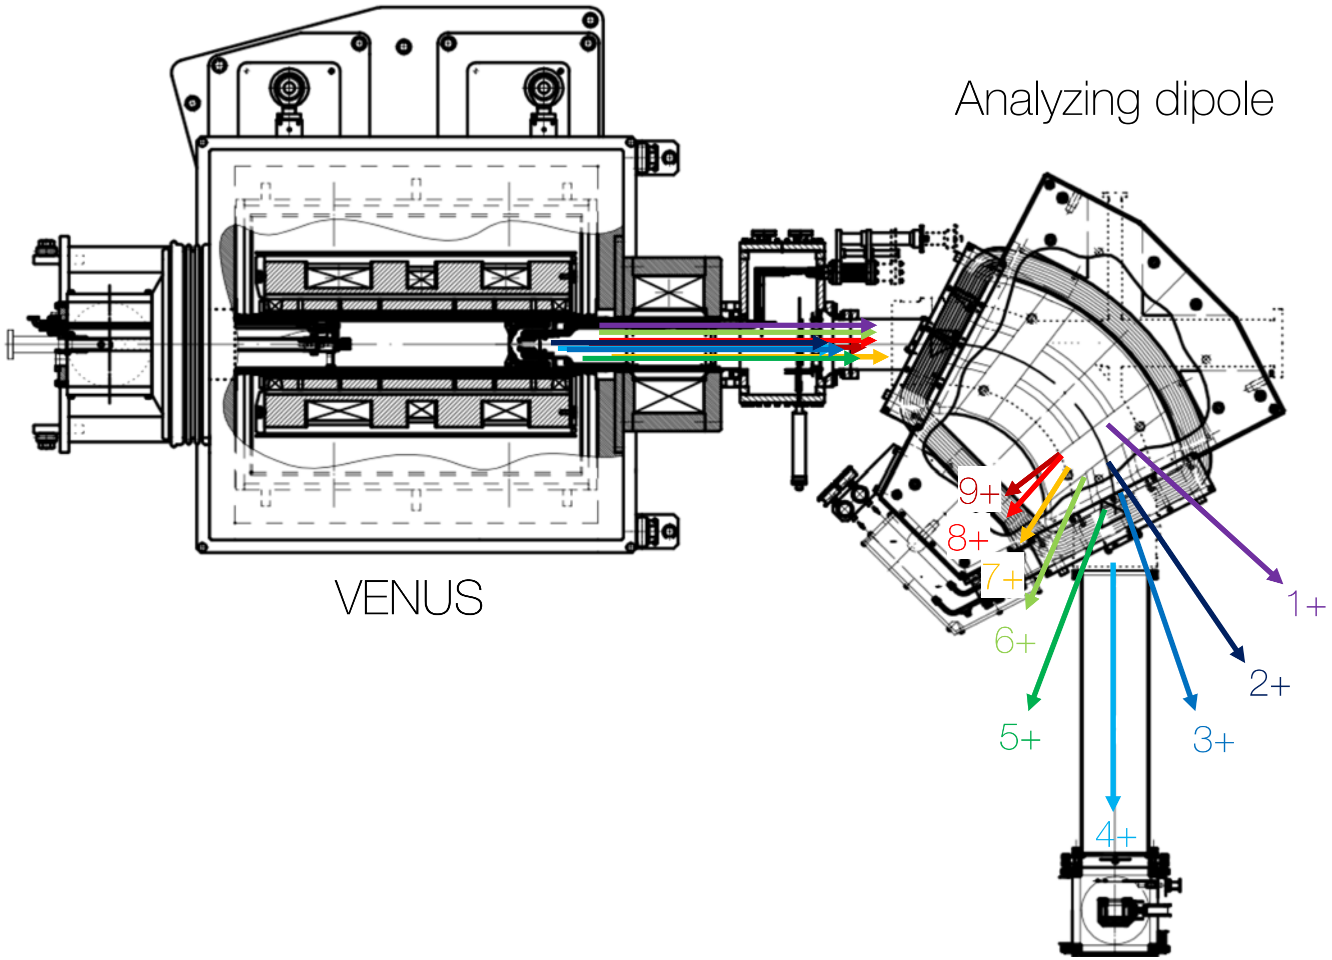

Dynamic charge state distribution: Xe off

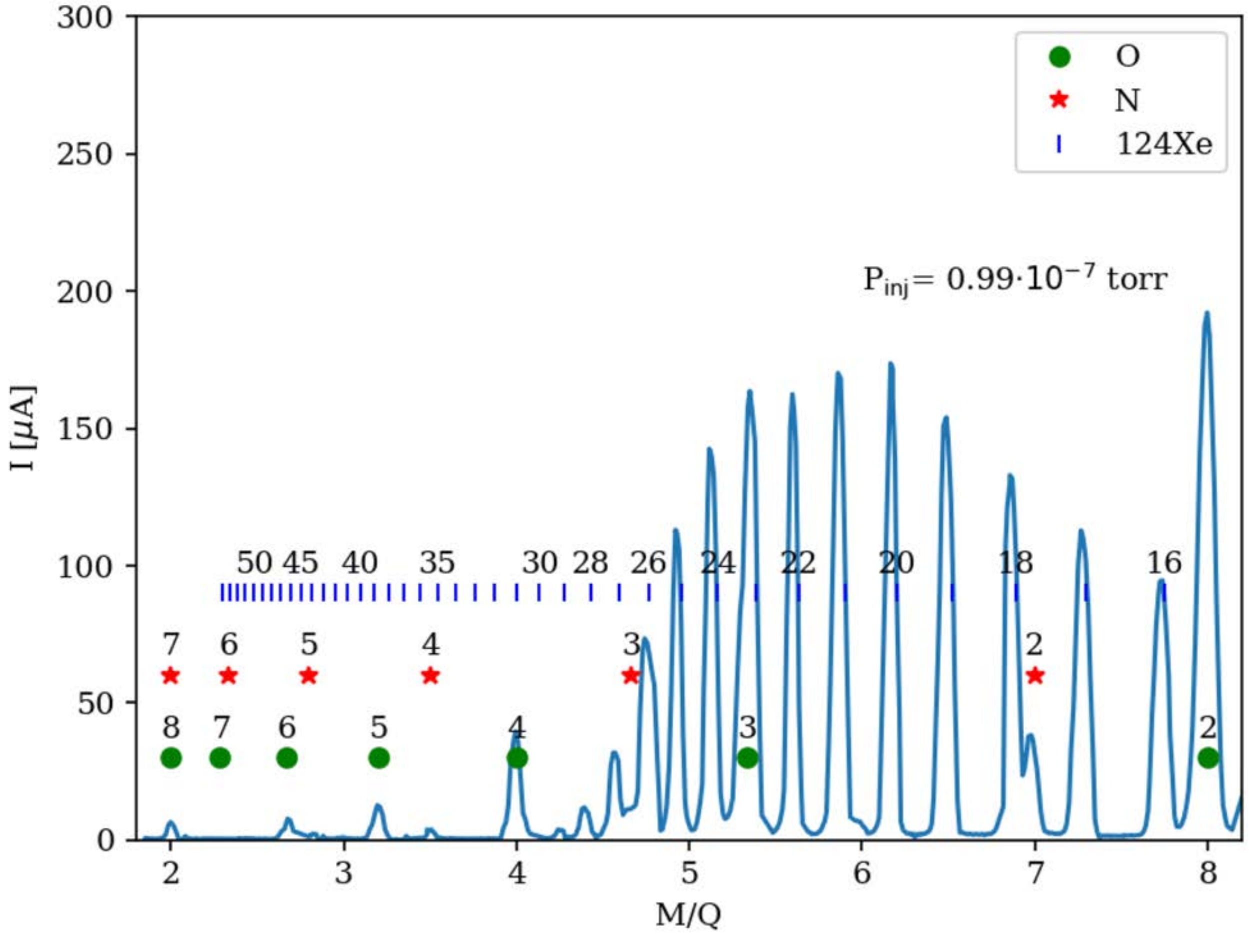

- Gradually close Xe valve
- CSD every 12 s
- Peaks occasionally dramatically reduced (instability?)

Smoothing CSD data for each species

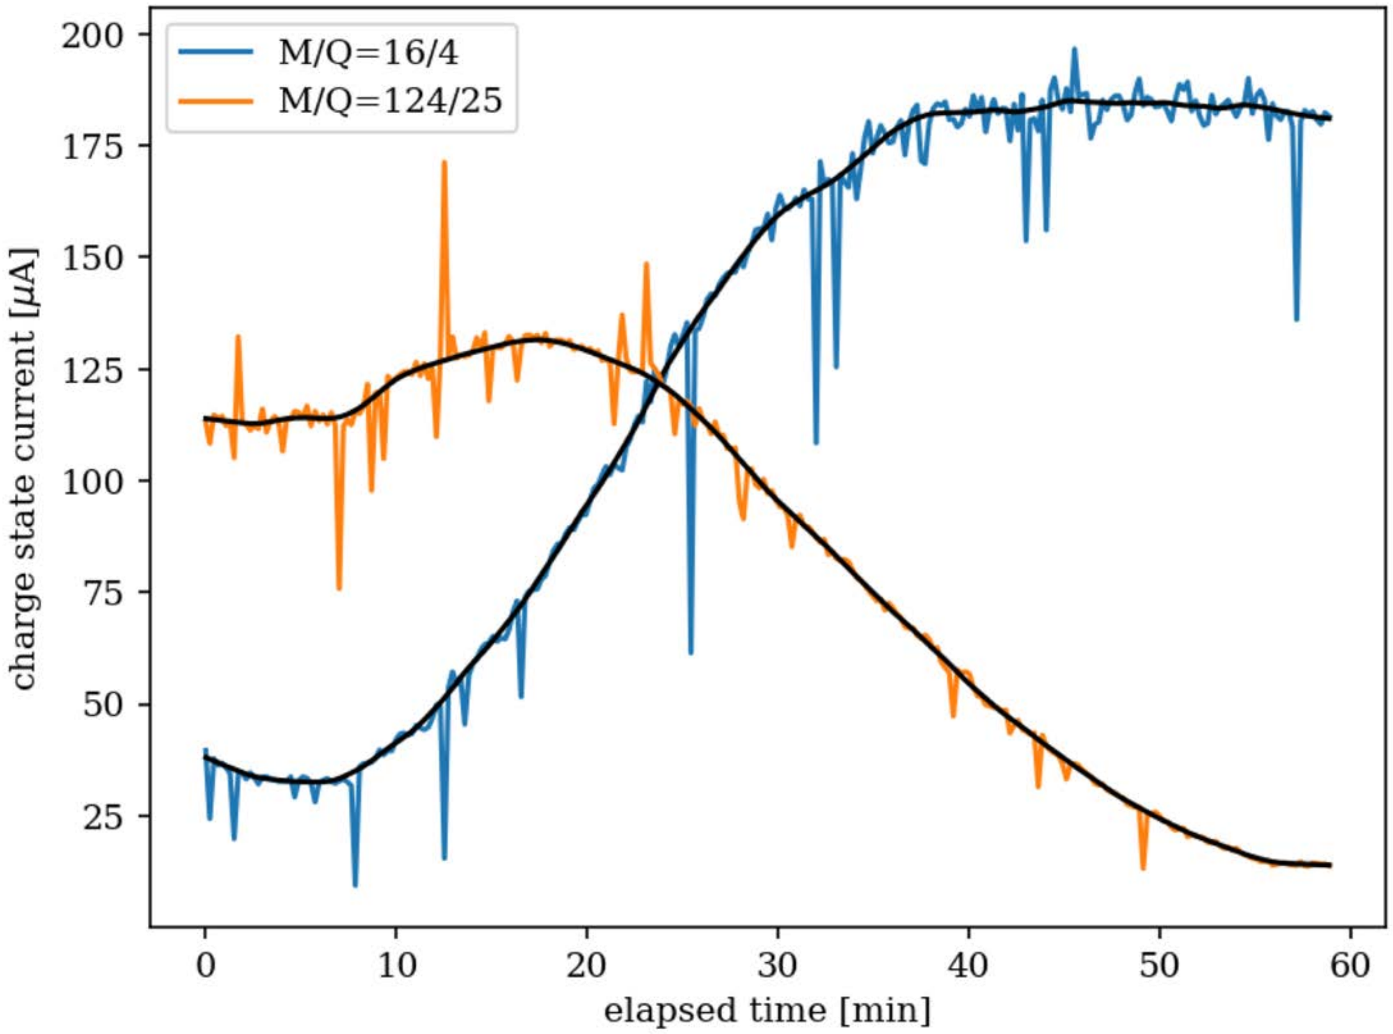

Find peaks for all species in a given CSD

Repeat for all CSDs

Smooth species data over sets by applying local regressive filter: LOWESS (locally weighted scatterplot smoothing)

Visualizing dynamic CSD information for humans

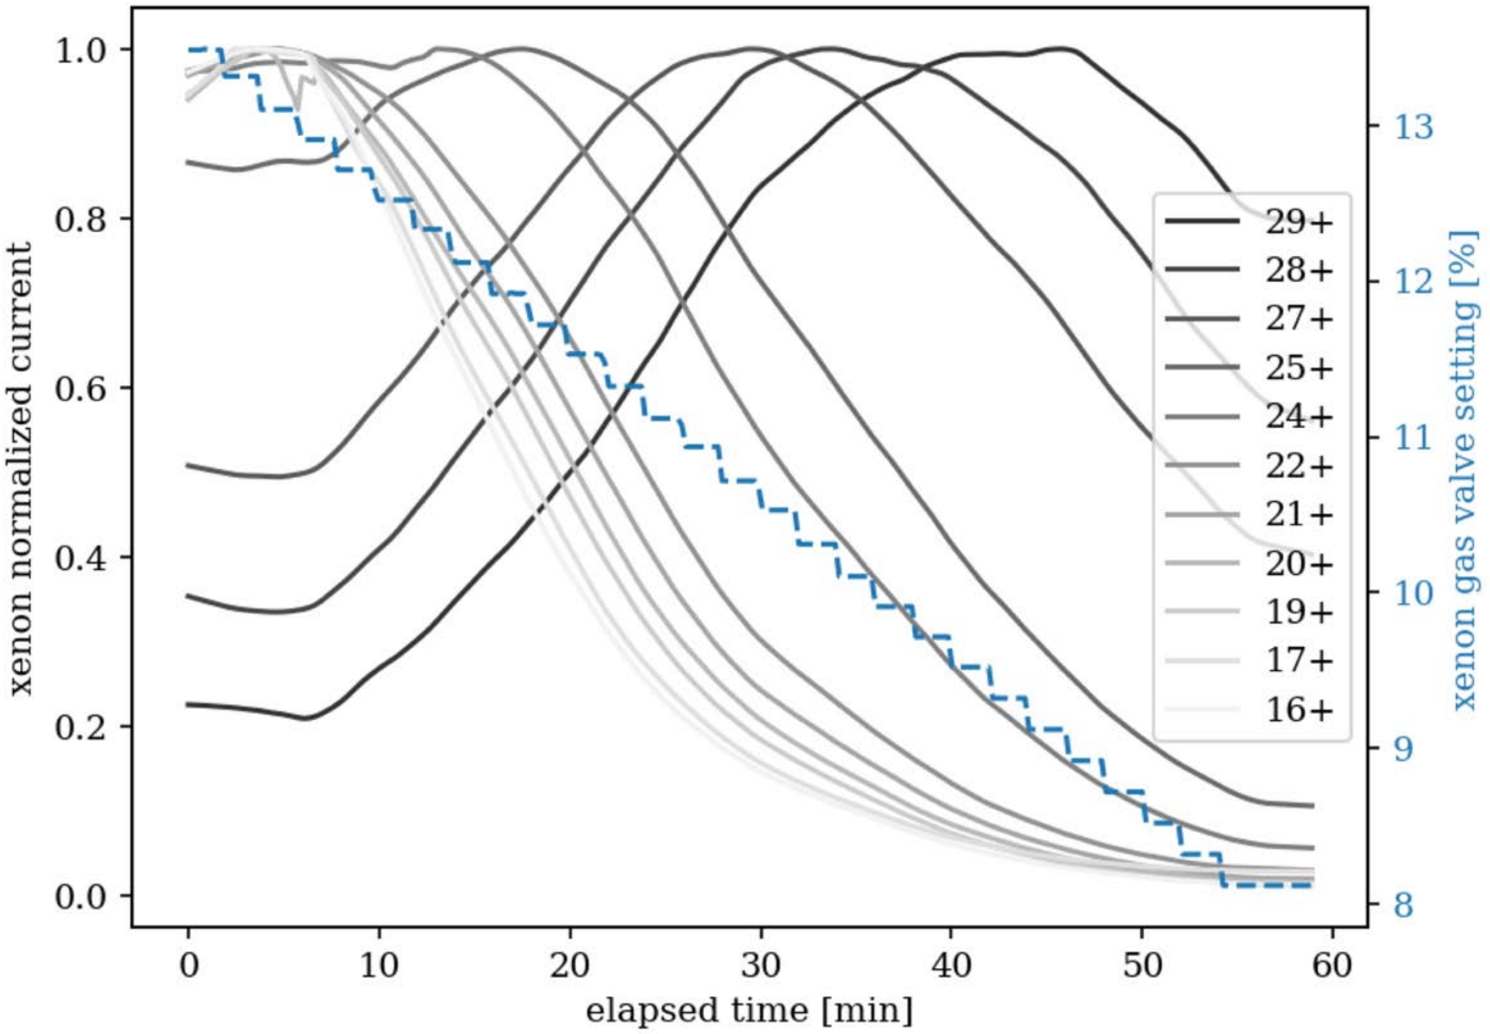

Increased charge state for ^{124}Xe plotted in darker gray

Visualizing dynamic CSD information for humans, attempt 2

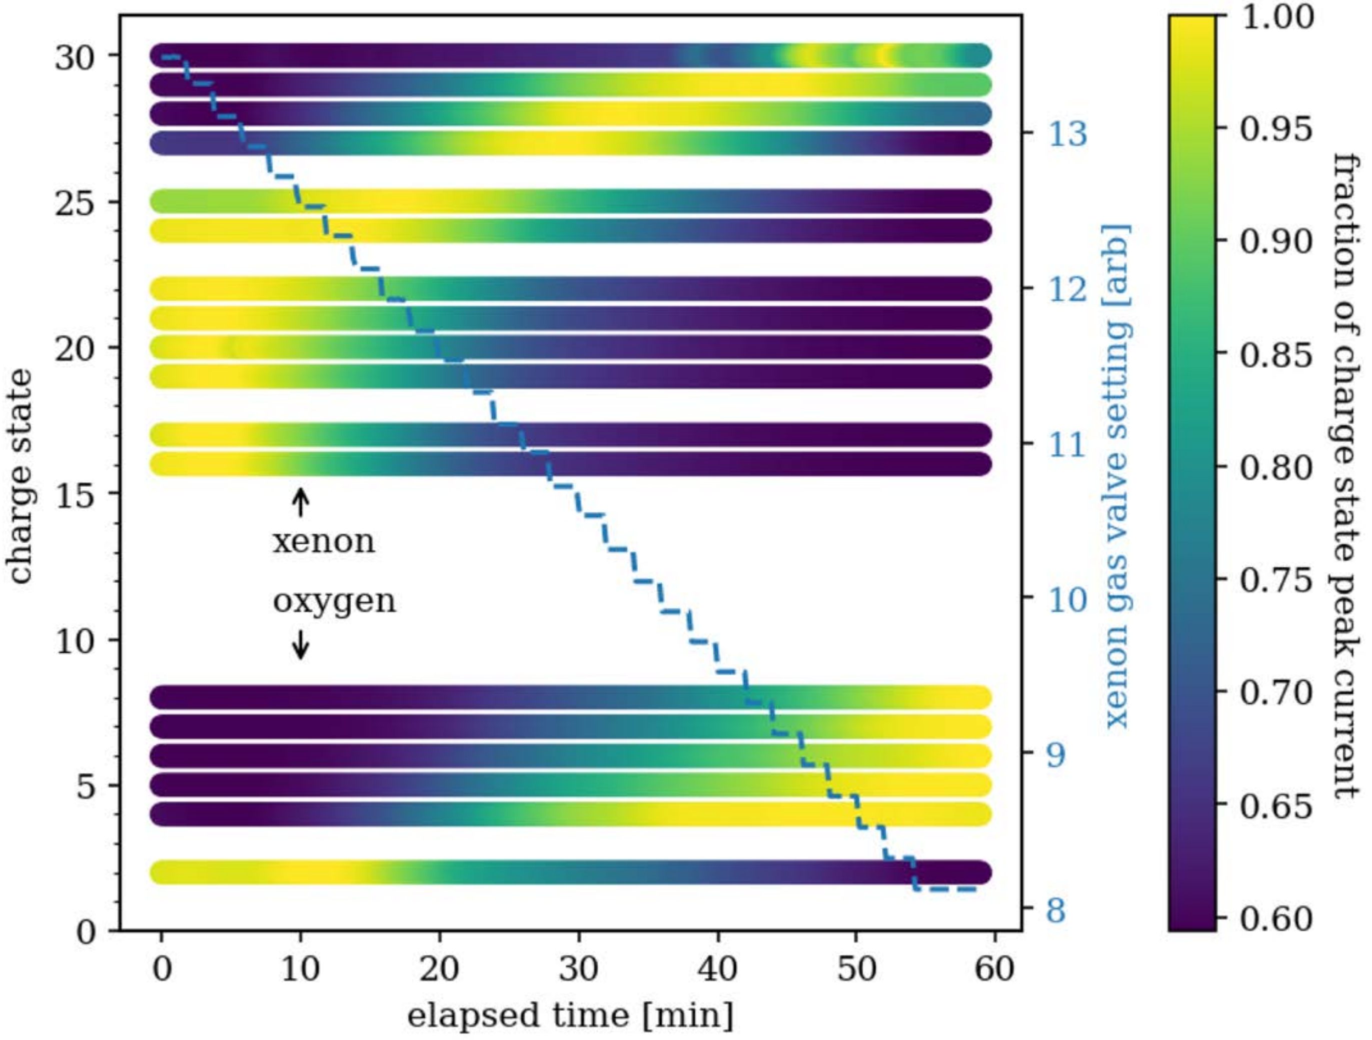

Note: this visualization problem is ours only. Machine learning can deal with multiple dimensional arrays, etc.

Increasing biased disk voltage

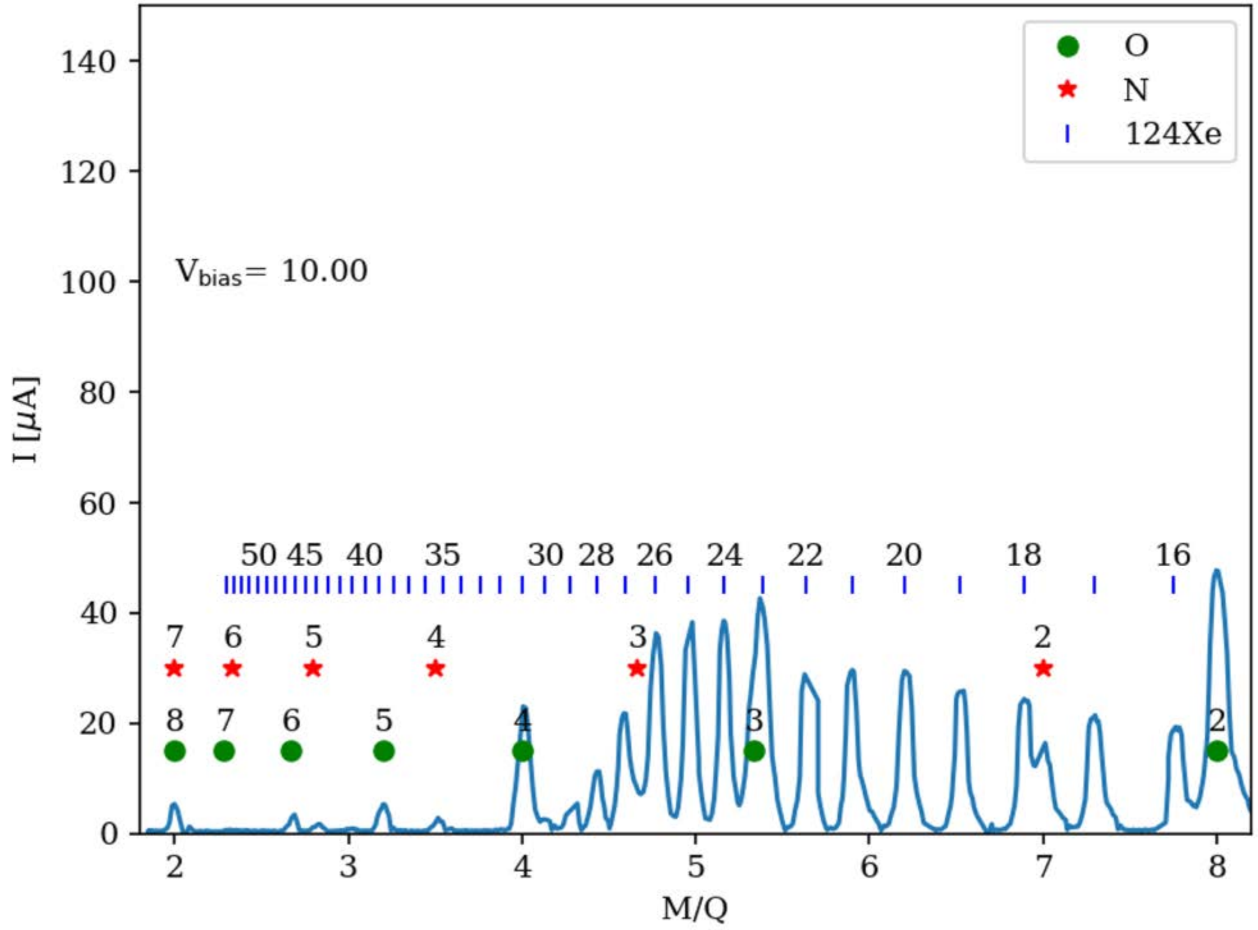

- Xe + O plasma
- Biased disk voltage is swept from 10 V to 80 V in 0.2 V steps. CSD at each step

Visualizing unstable regions

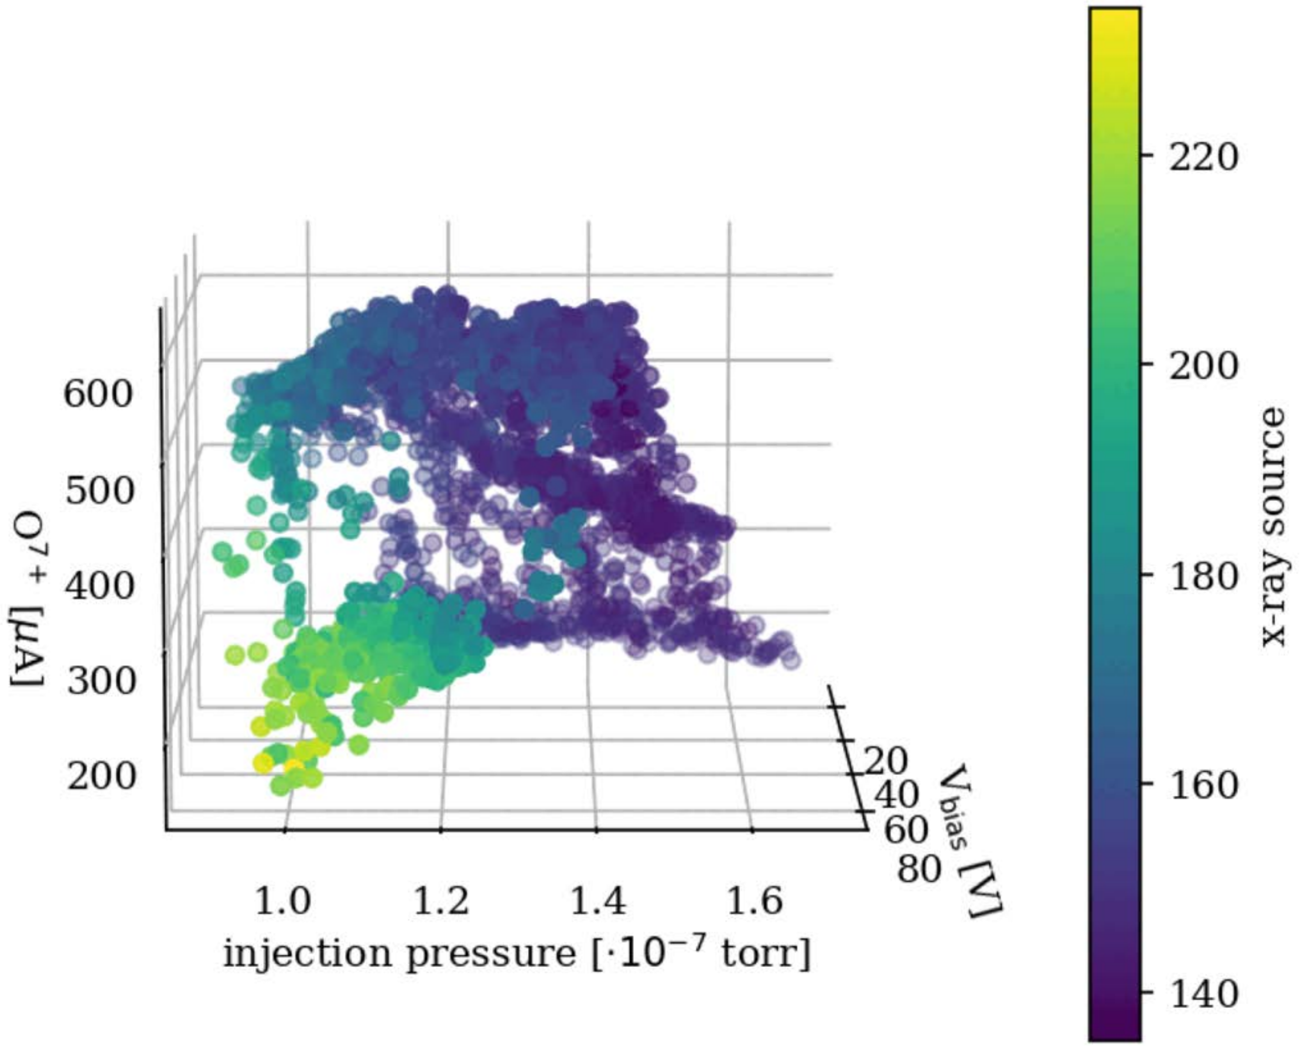

- Oxygen plasma
- Controlled Parameters:
  - biased disk voltage
  - oxygen valve
- Only plotting points where standard devation is < 5% of beam current

- 28/18 GHz = 4300/1400 W

Instability region in source pressure/biased disk voltage operation space

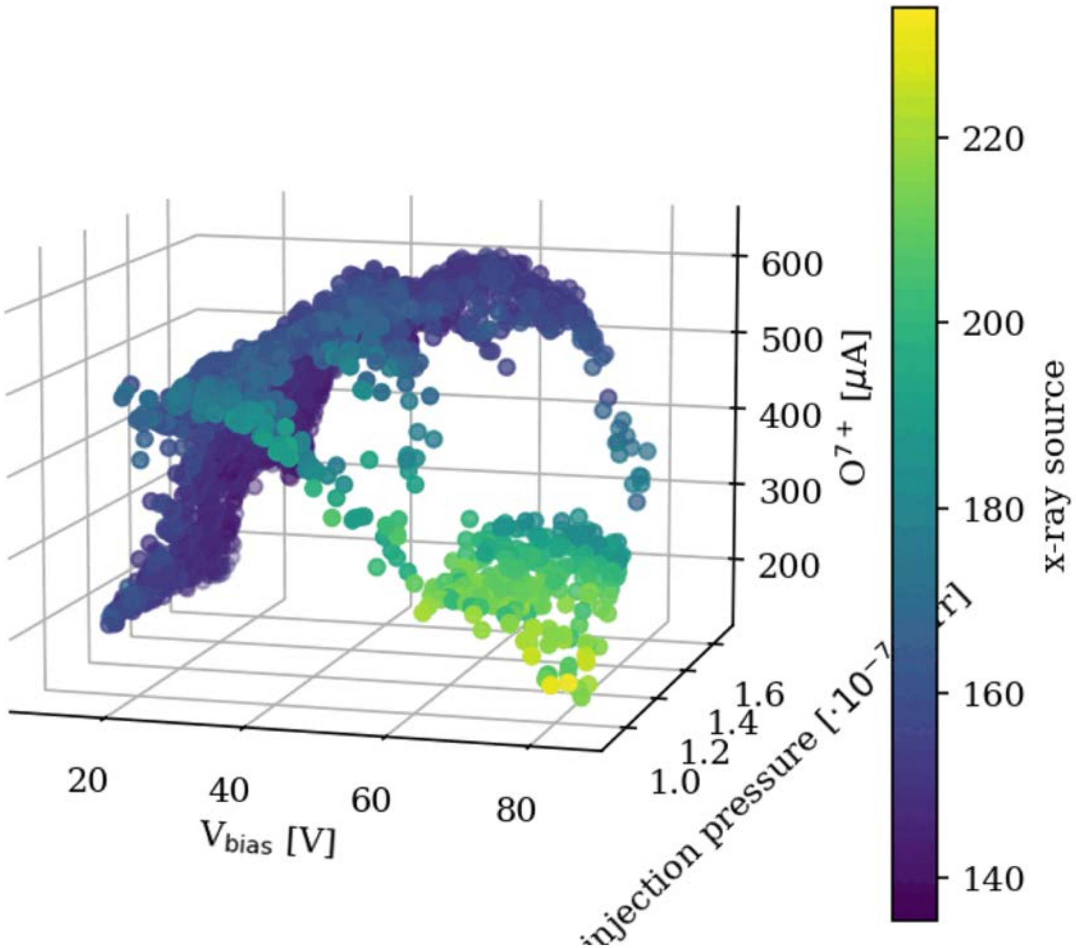

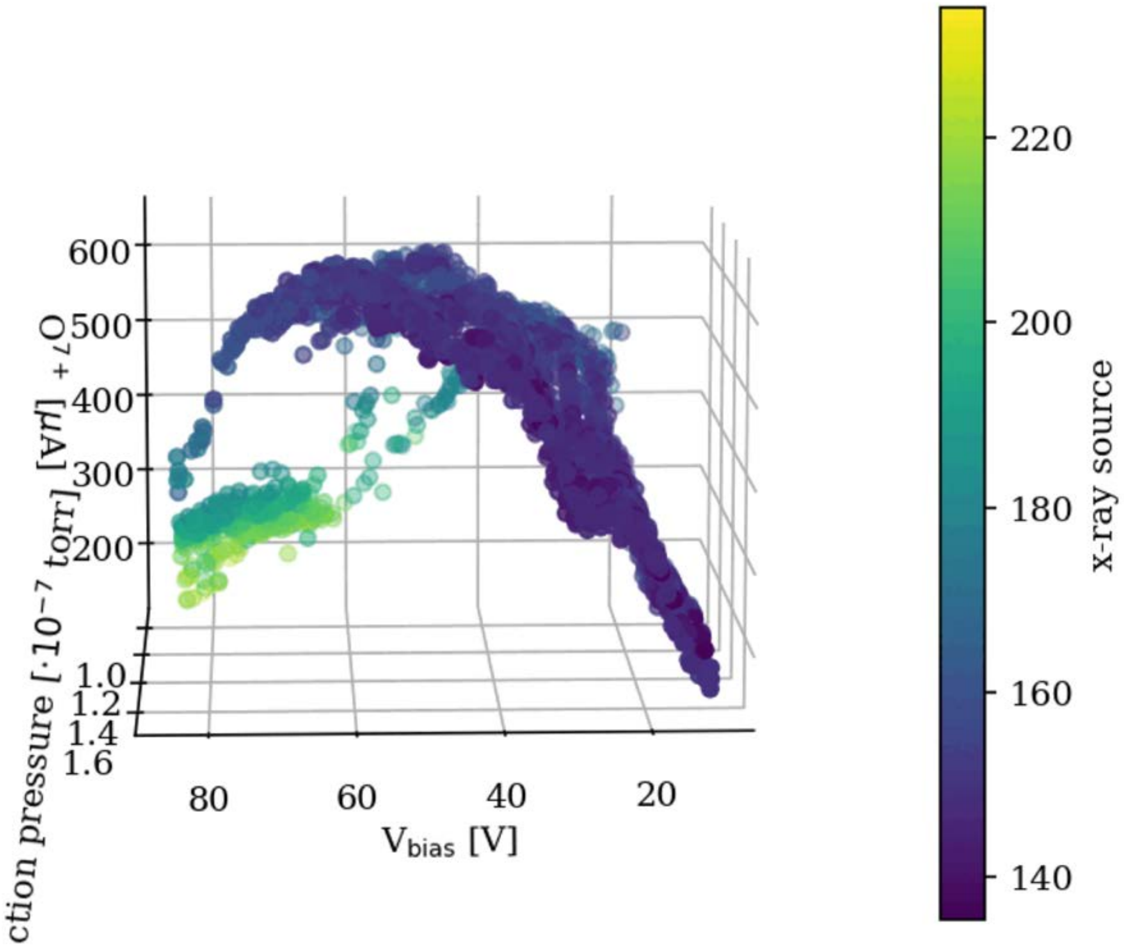

Weekend-to-weekend differences

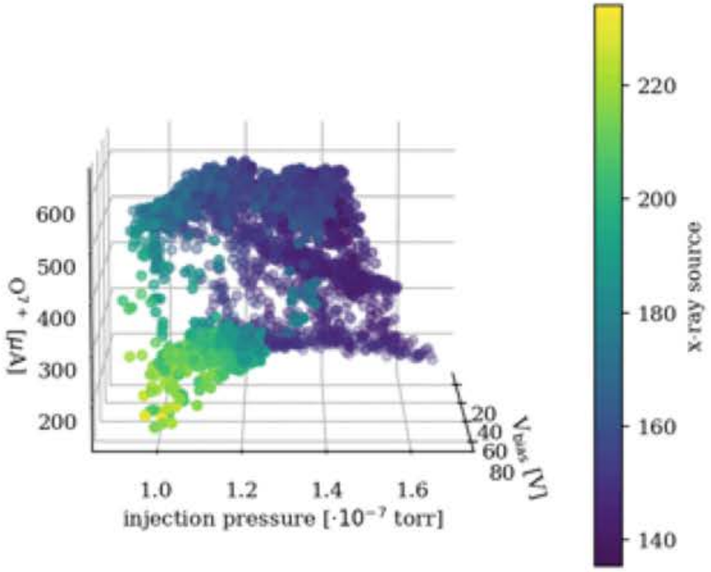

02 Sept. 2022

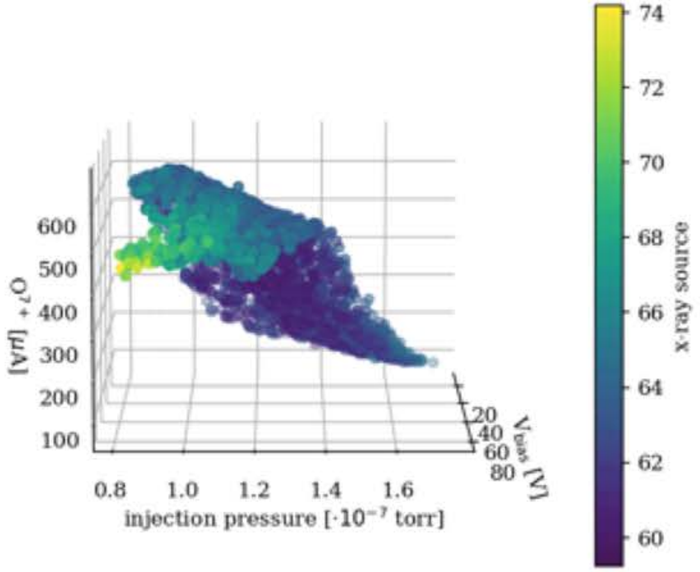

15 Jan. 2023

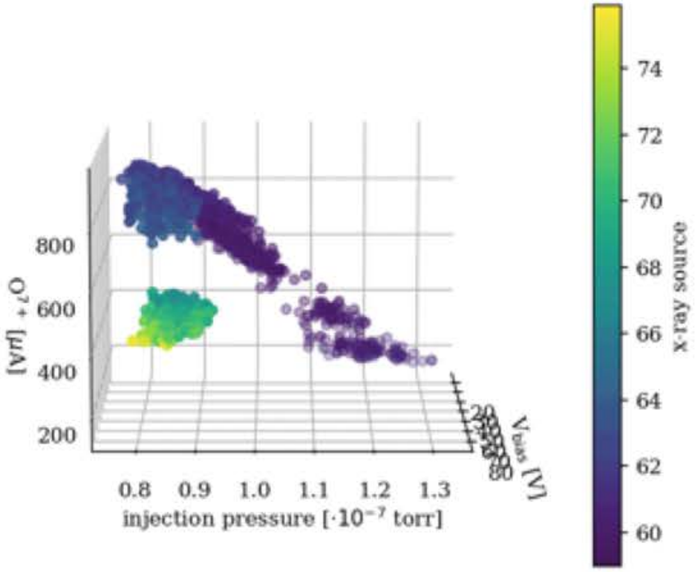

22 July 2023

Weekend-to-weekend differences

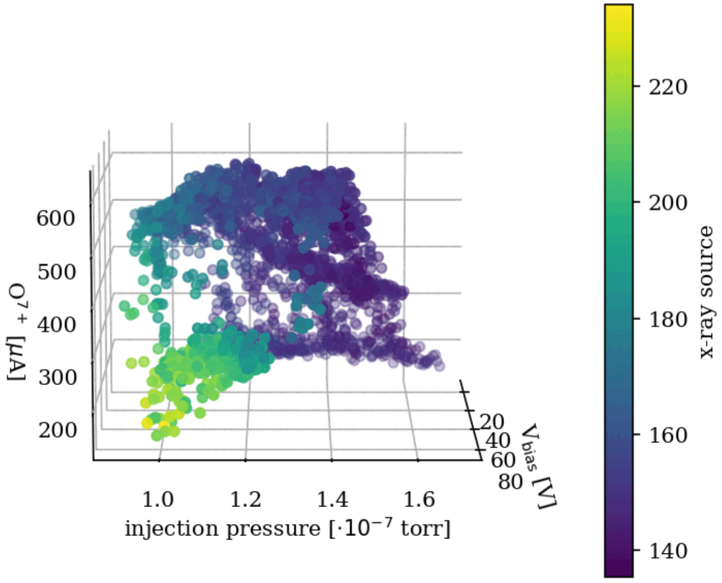

02 Sept. 2022

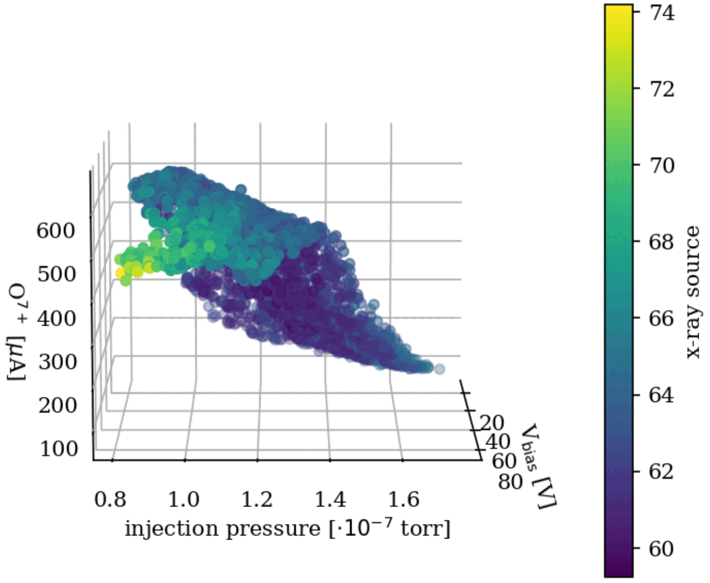

15 Jan. 2023

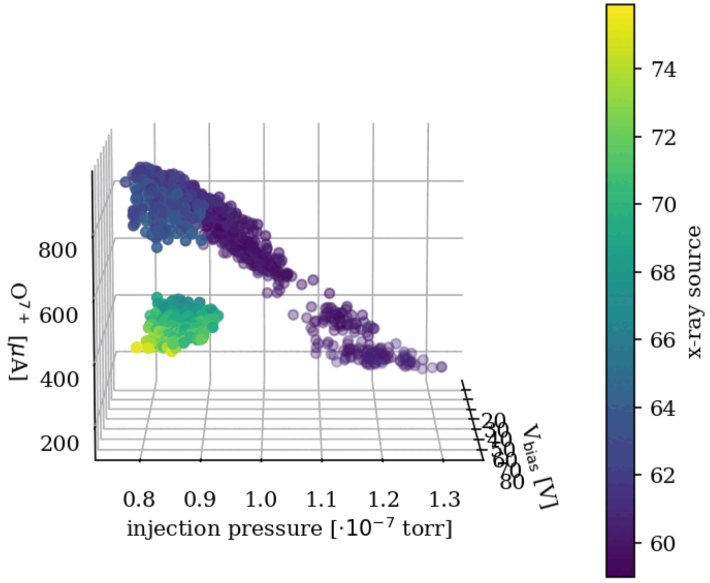

22 July 2023

Even though different, can we use information from one weekend to tell us something about another

Neural Networks

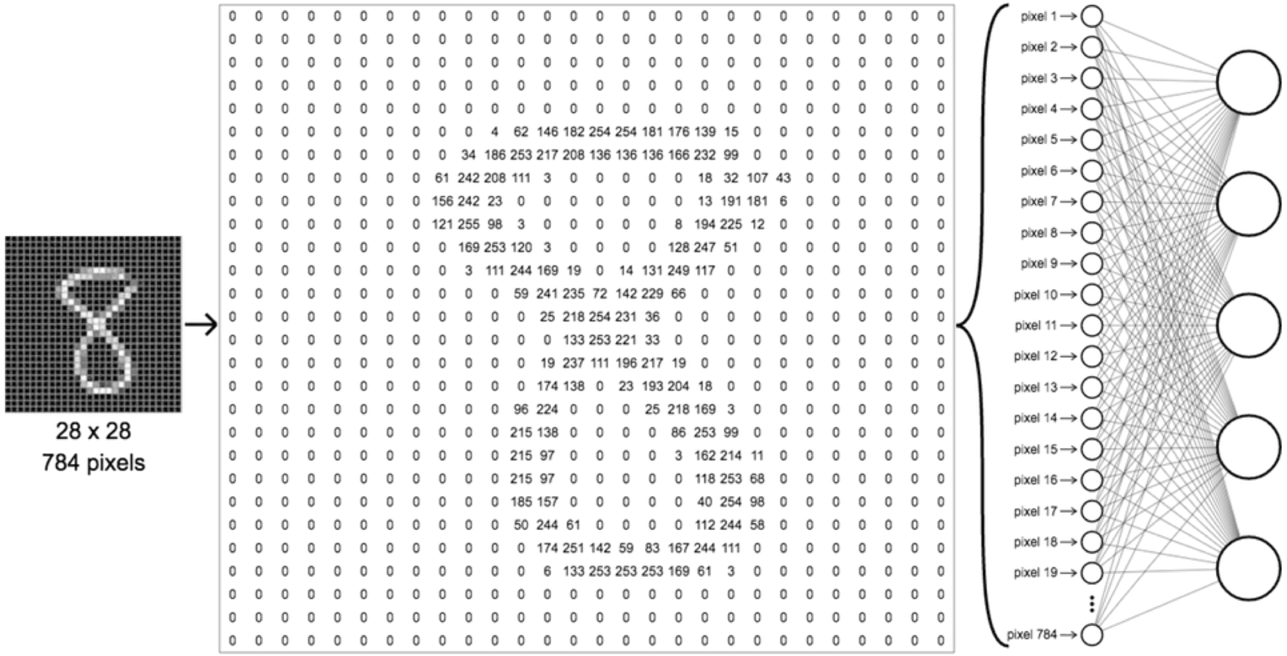

Goal: train network on some subset of number images and predict identity of new, unseen numbers

Try to apply this to VENUS:

- For given settings (coil currents, valve positions, pressures, etc.) tell me the beam current

https://medium.com/algorithms-for-life/numbre-a-number-recognizer-neural-network

Neural Network research by undergraduate Ezra Apple

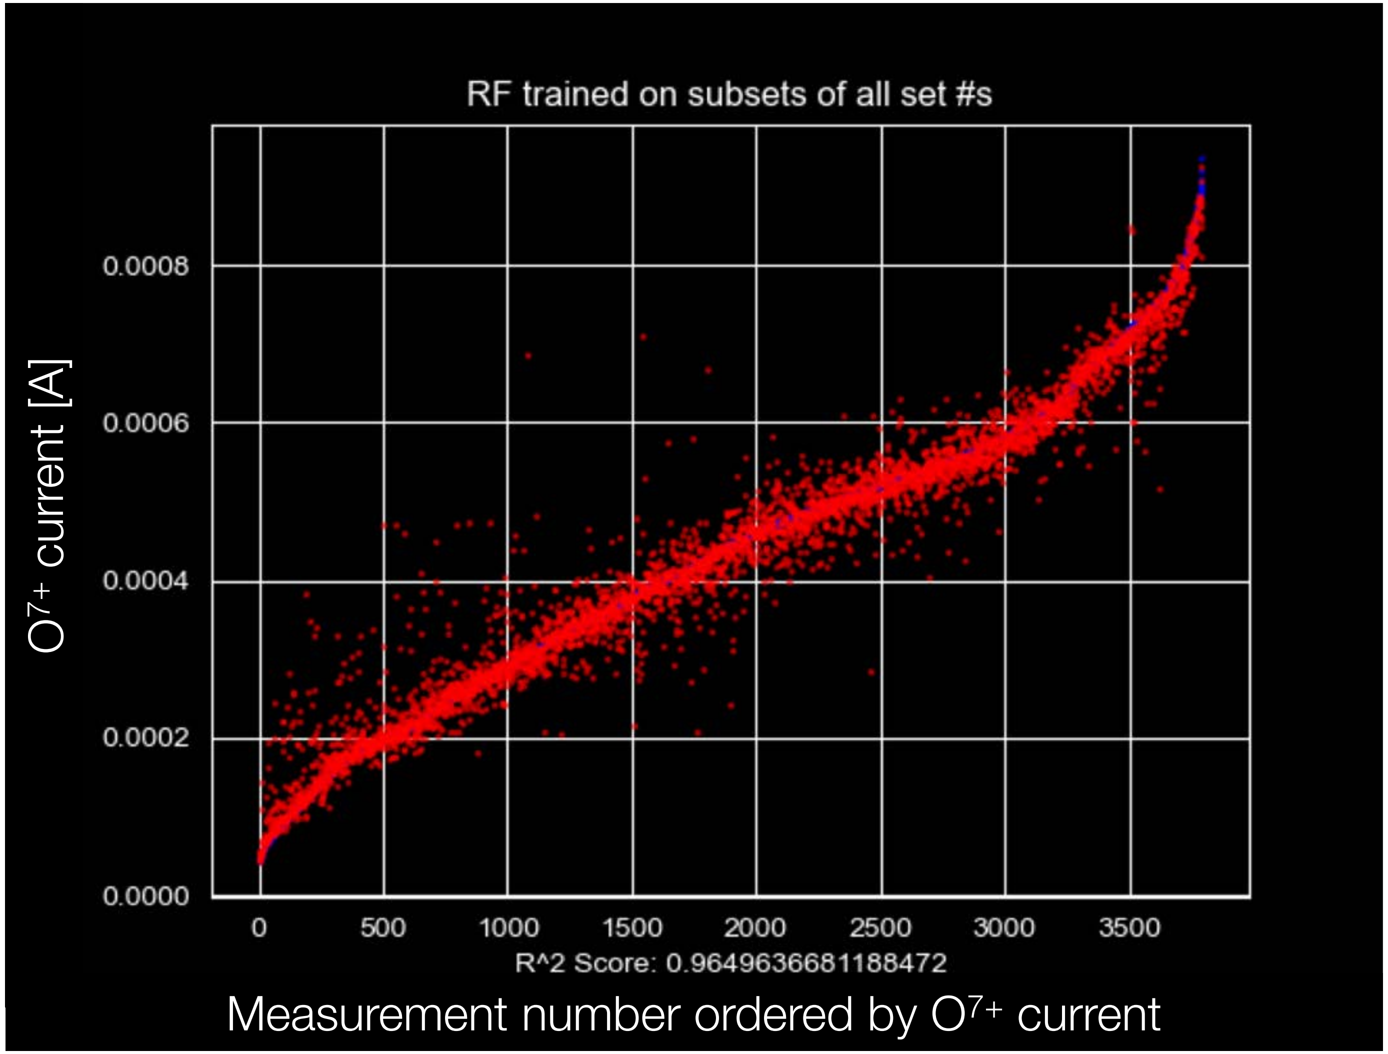

- Like number identification, train on subset of collected data
- Test ability to predict on unseen remainder of data

Neural Network does not do well with predicting in new weekend

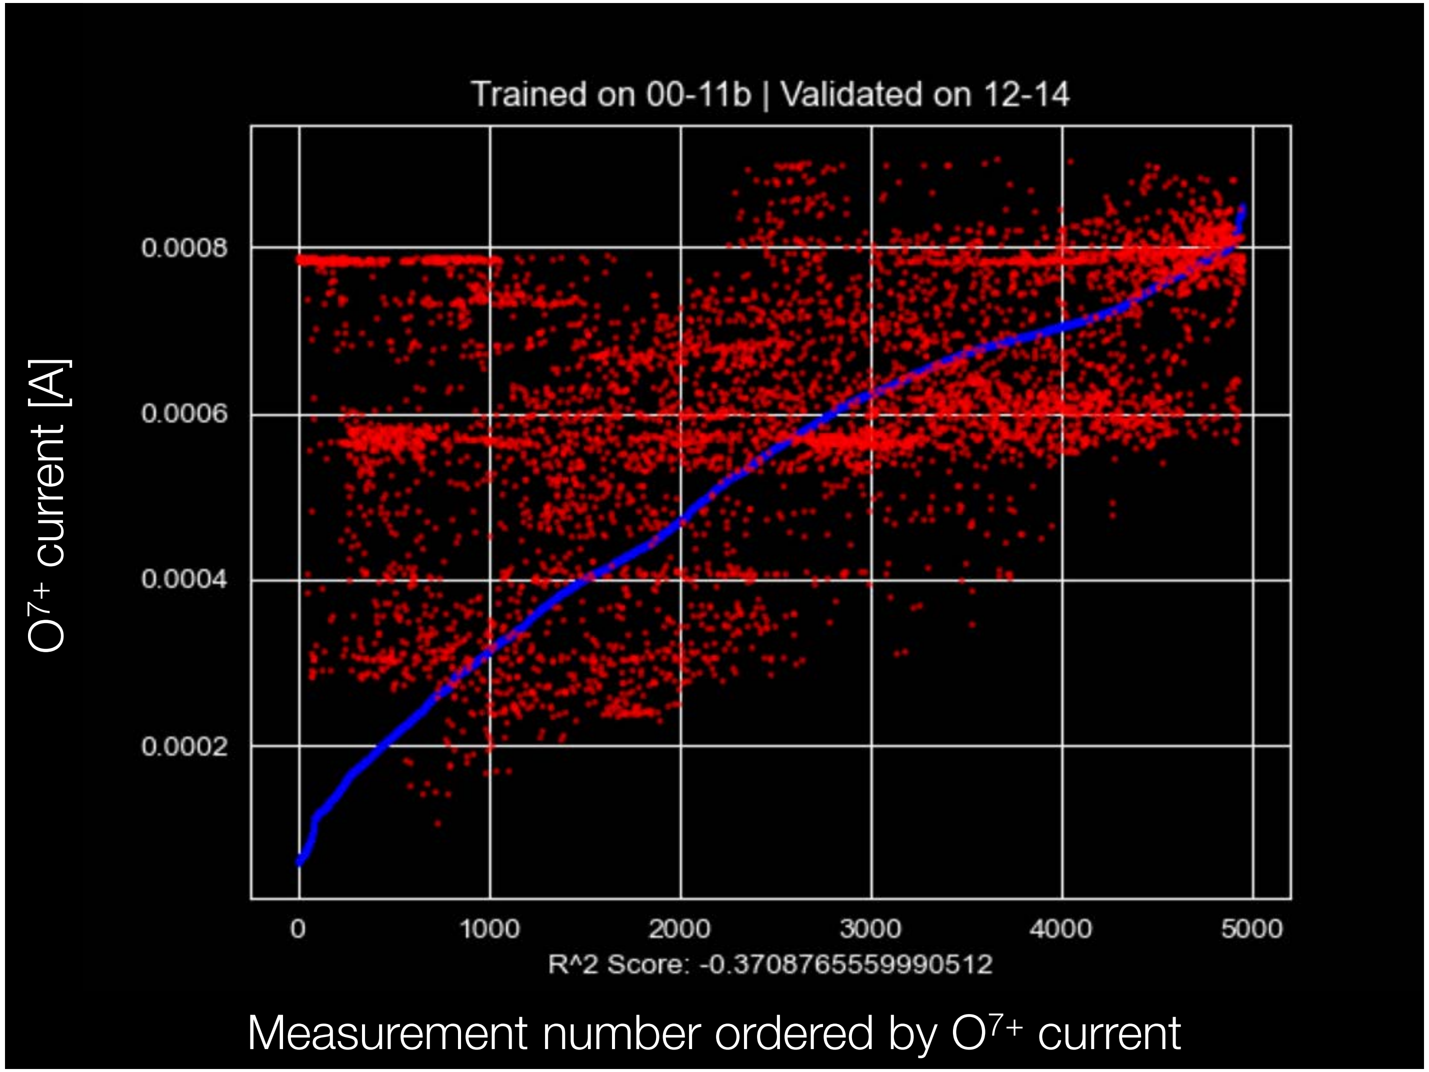

- Weekends differ enough that training on previous doesn't help with new weekend

Some hope: seeing 10% of data sets is enough to predict remainder

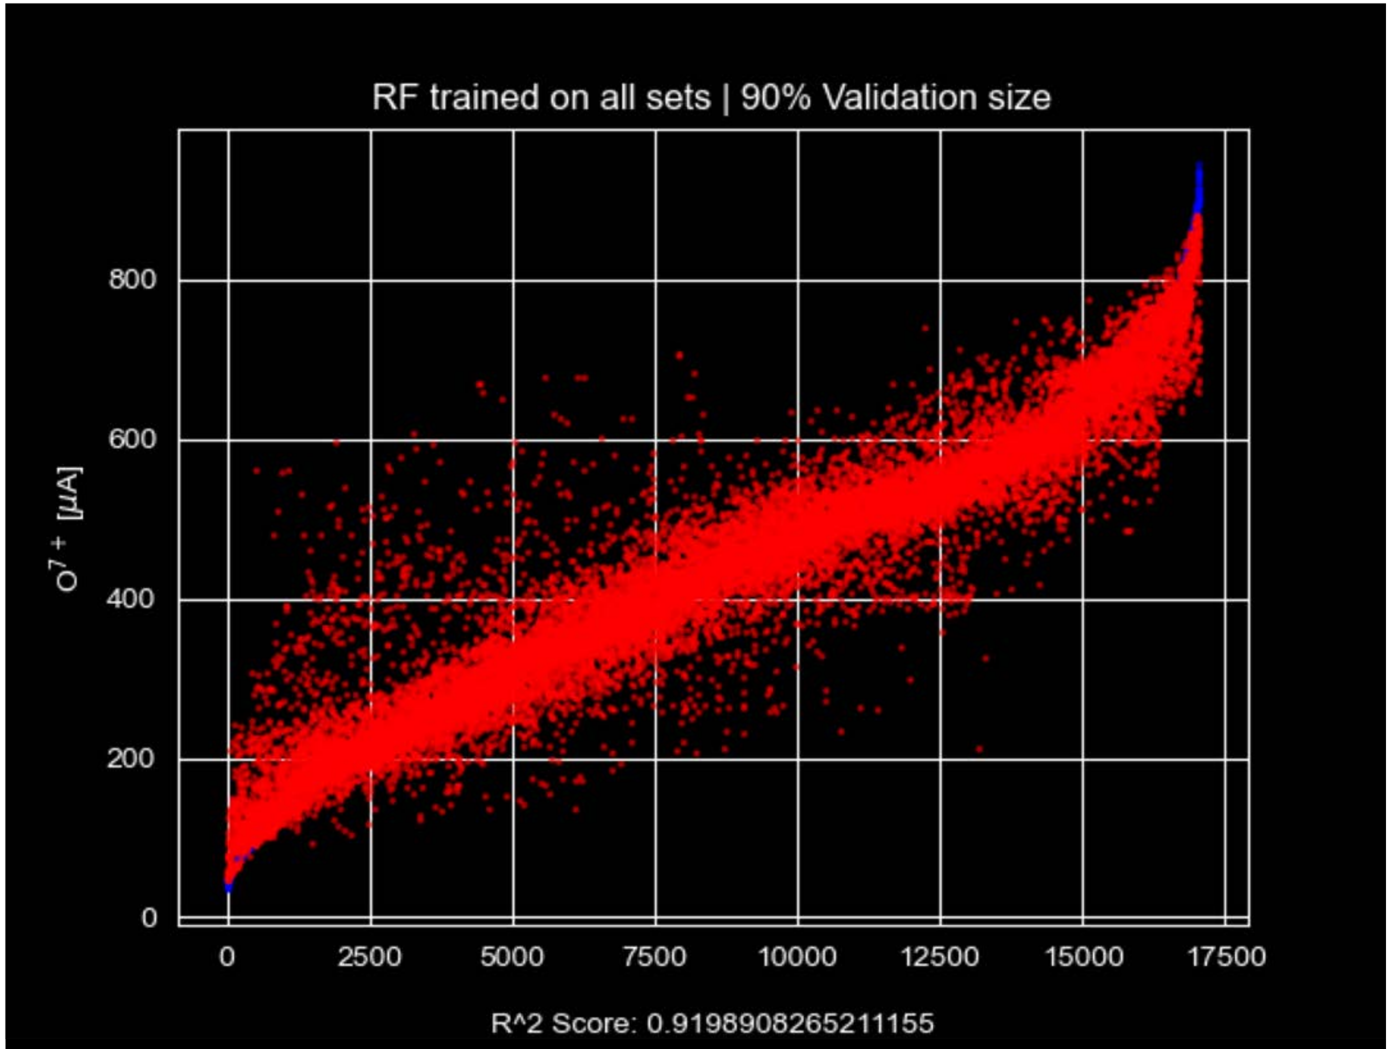

Measurement number ordered by O^{7+} current

- Training using 10% of every weekend's data
- Validation on remainder of data is pretty accurate

Preparing emittance scanners for ML optimization

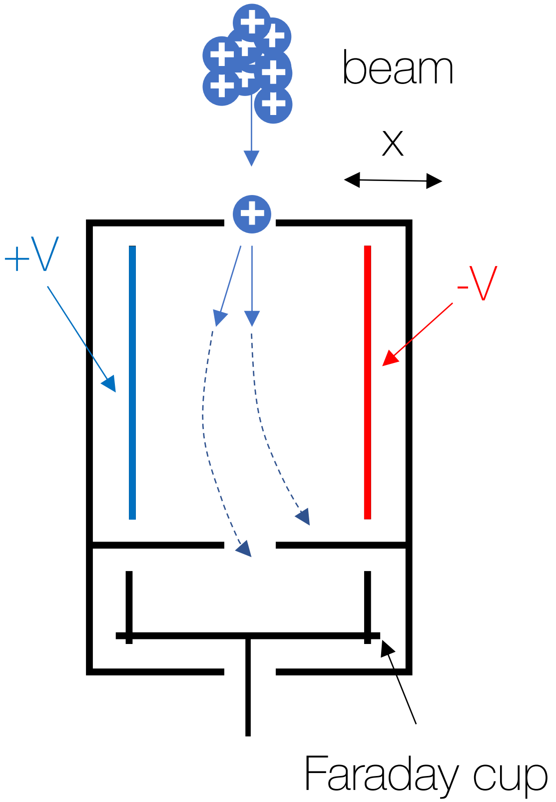

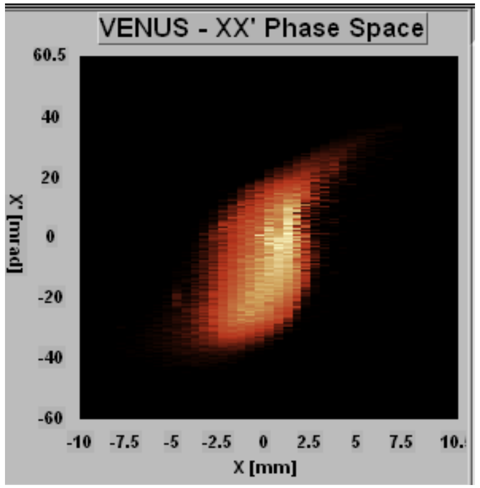

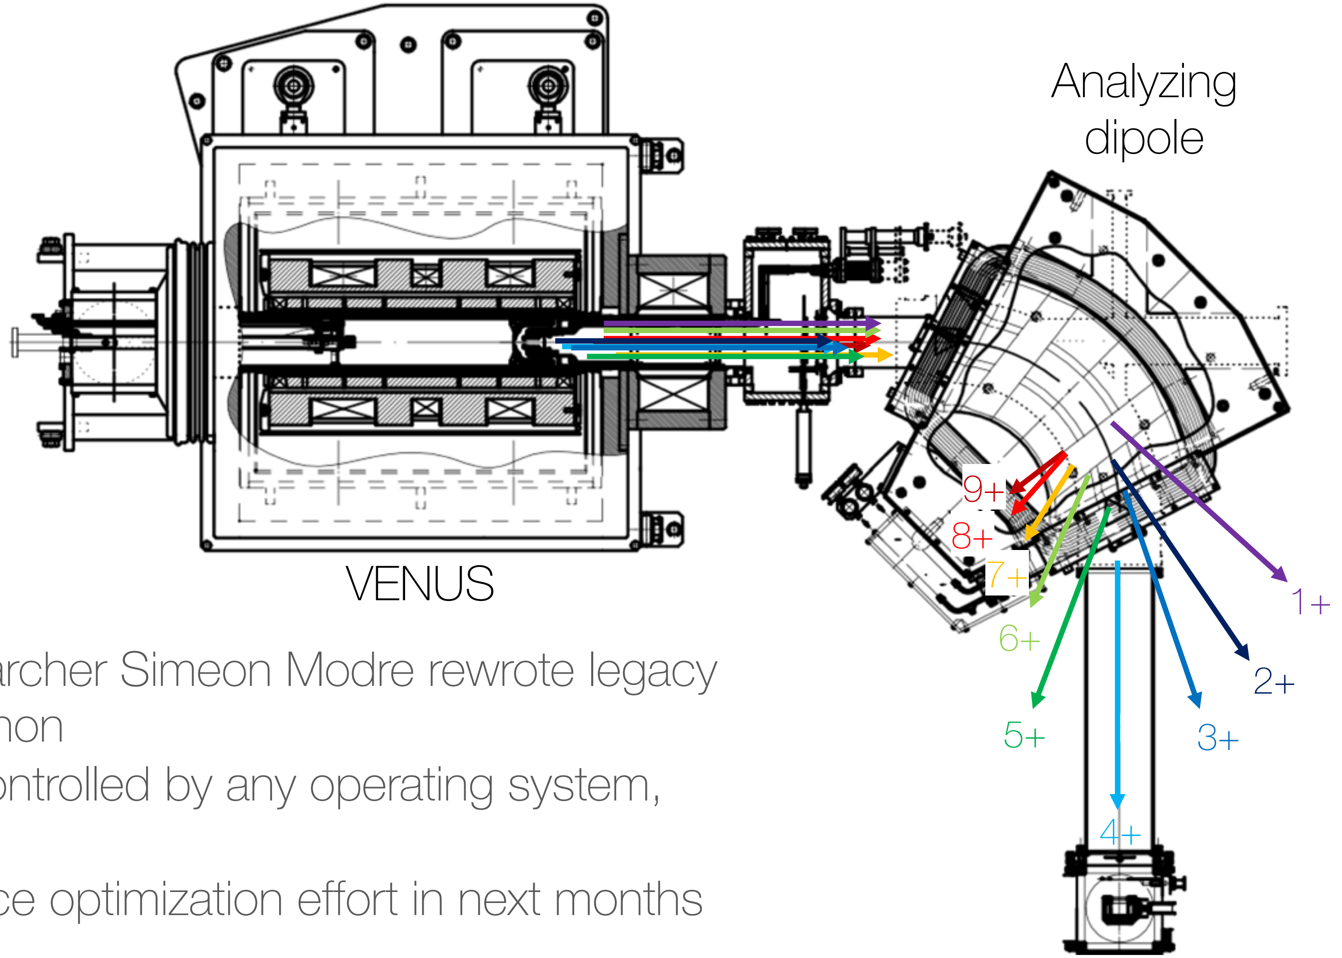

- Visiting student researcher Simeon Modre rewrote legacy LabView code in Python
- New code can be controlled by any operating system, including ML
- Will start ML emittance optimization effort in next months

Bayesian optimization in dynamic systems

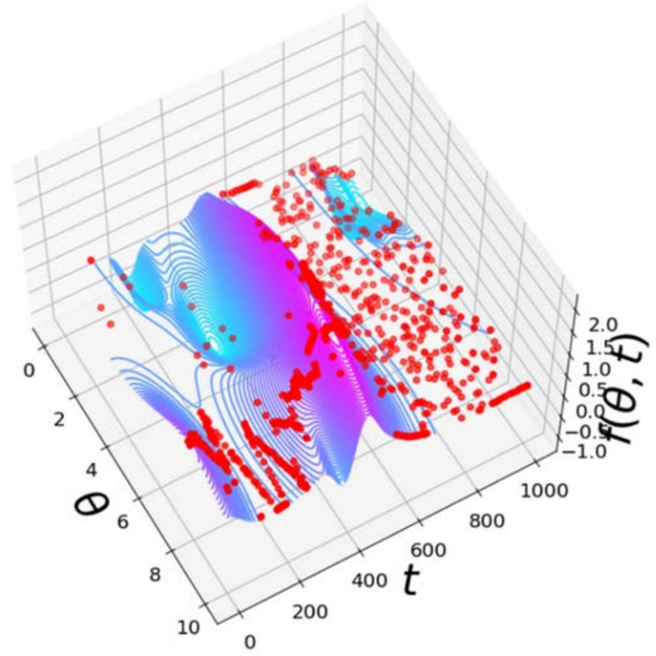

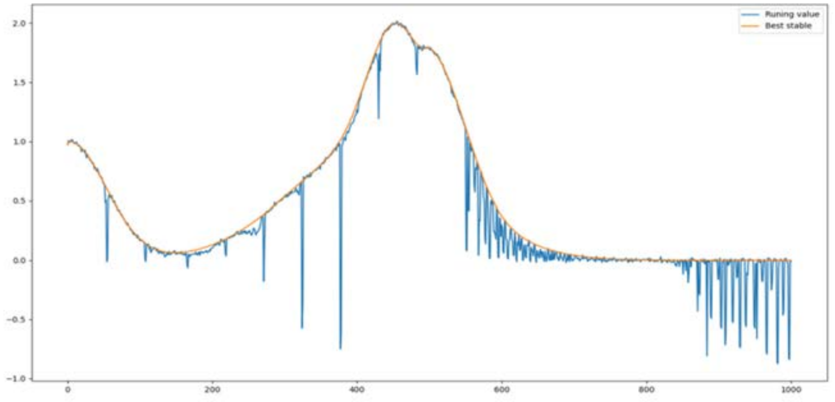

- Work done by post-doctoral researcher Victor Watson
- θ: control variable, t: time.
- Maximize current (f) and keep instability below a threshold in a dynamic system.
- Theoretical work done, and will implement with VENUS
- Goal: do this non-destructively

Reinforcement learning

https://research.google/blog/multi-task-robotic-reinforcement-learning-at-scale/

- Have robots take action
- Score result and reward or penalize robot
- Robot takes next action based on score with goal of improving

Using reinforcement learning to predict optimal bias voltage

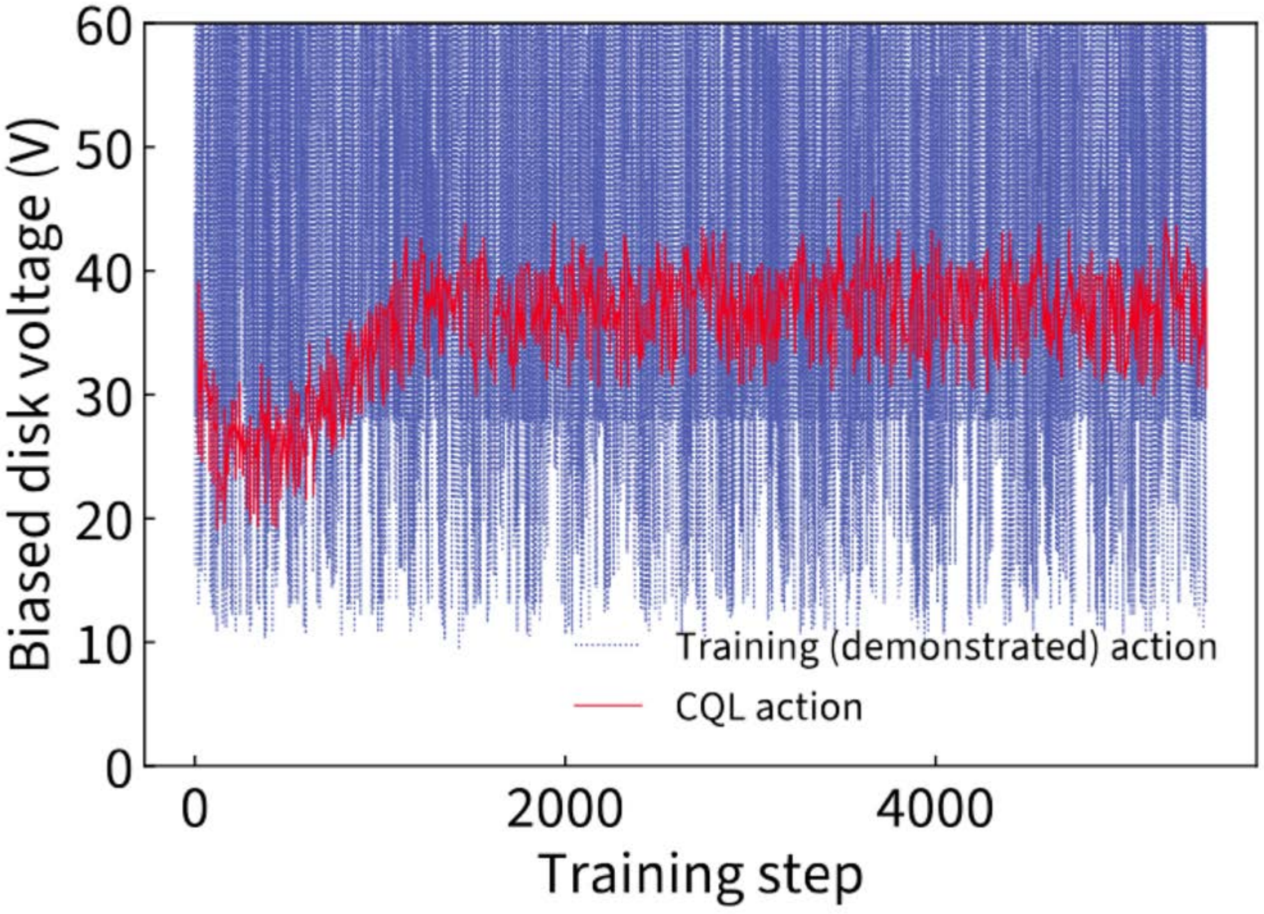

Physicist Yue Shi Lai applying Reinforcement Learning to VENUS:

Conservative Q-Learning (CQL)

- Uses offline training to pre-learn

Fast diagnostics will be designed into MARS control system

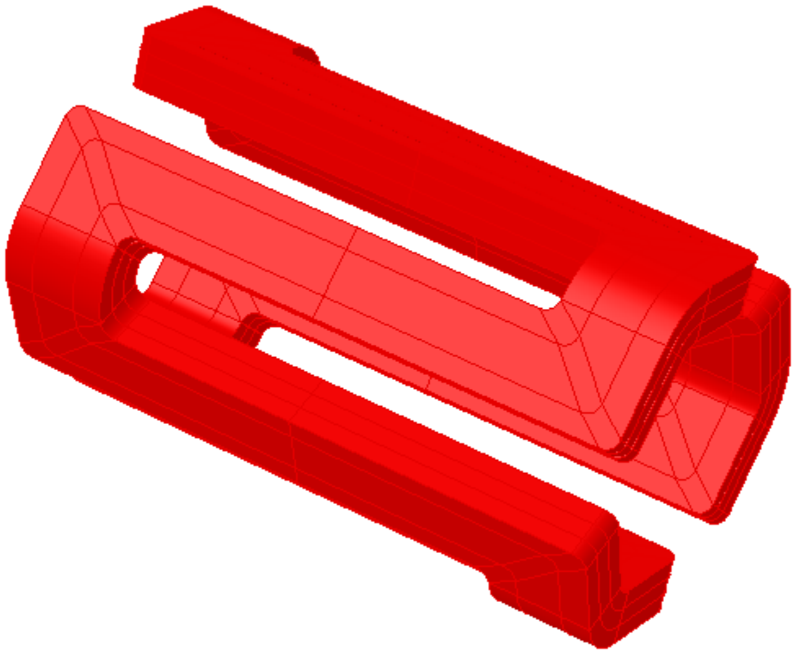

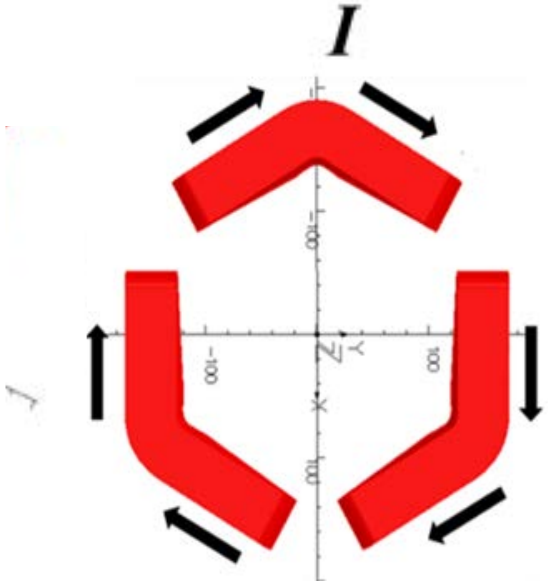

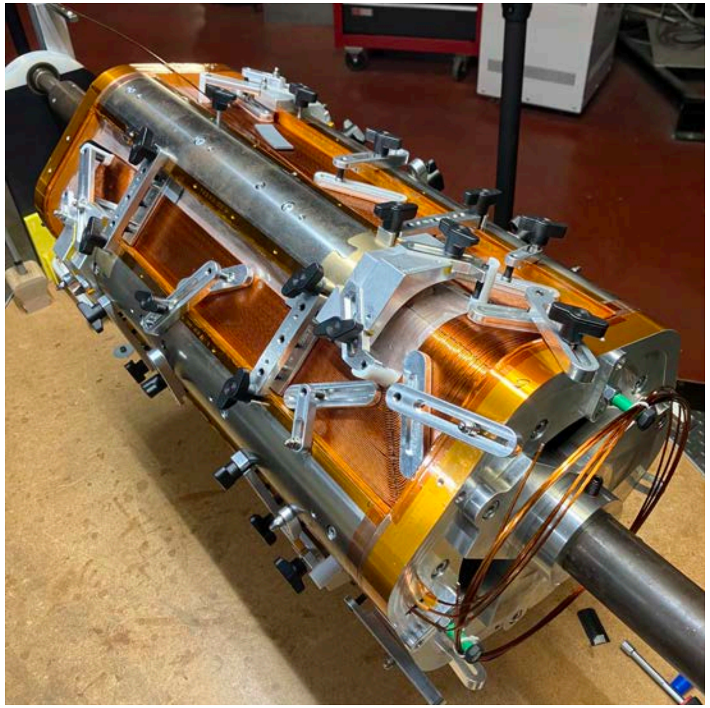

15 of 24 64-turn layers complete

- Currently winding closed loop coil for 45 GHz MARS-D ECRIS
- Think ahead: faster diagnostics (drain current, biased disk current, etc.) for MARS vs. VENUS 3 Hz. Computer control (and ML?) not going away!

Apply machine learning to metal beam optimization for ^{50}Ti beams

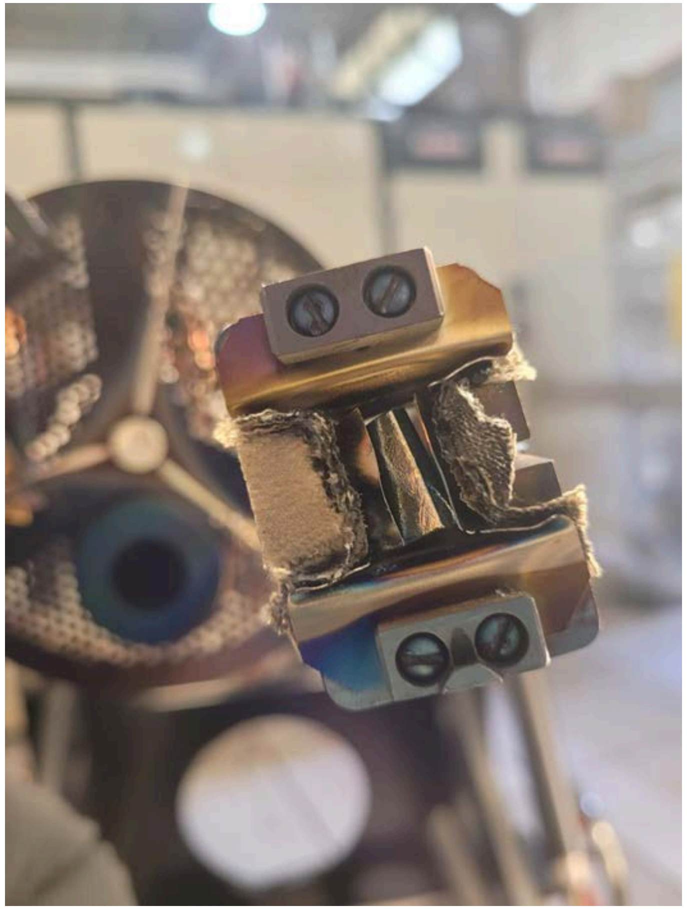

- LBNL's superheavy element research program required >1 pμA, ~5 MeV/u 50Ti in preparation for element 120 search
- Tasked with applying ML for optimization of this difficult beam
- Used "boat" ovens which sprayed Ti everywhere and beams were very unstable
- Attempted (~1 day) to use computer control

Use machine learning to optimize ^{50}Ti from this oven instead?

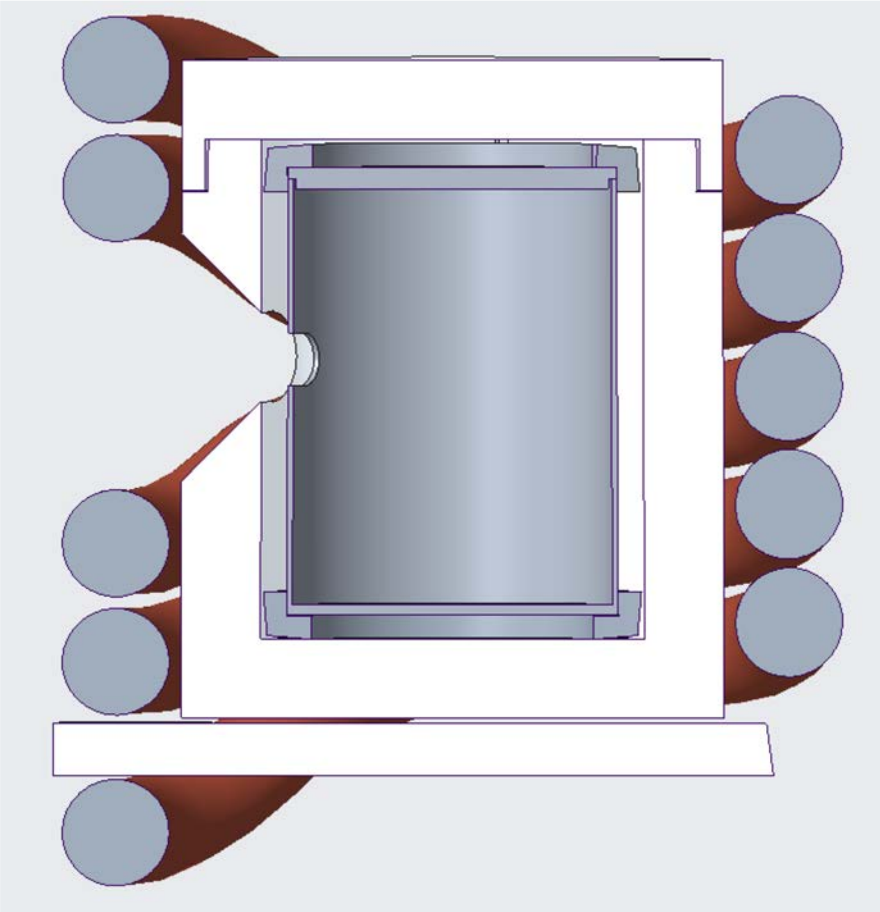

- Inductive oven aims output at plasma
- Extremely stable operation (days without adjustment)
- Produced element 116!
- Machine learning hasn't been necessary!

Machine learning effort provided non-ML operational improvements

Blessing/curse of ECRIS: ion beams produced from any materials reaching plasma that don't destroy the plasma

Adsorbed particles later desorbed:

- Residues from previous work or handling of inner surfaces
- Water, salts, etc. on all surfaces after atmosphere exposure
- Prior beam materials (e.g., metals)

Plasma chamber desorption can be sped up using plasma as heating element

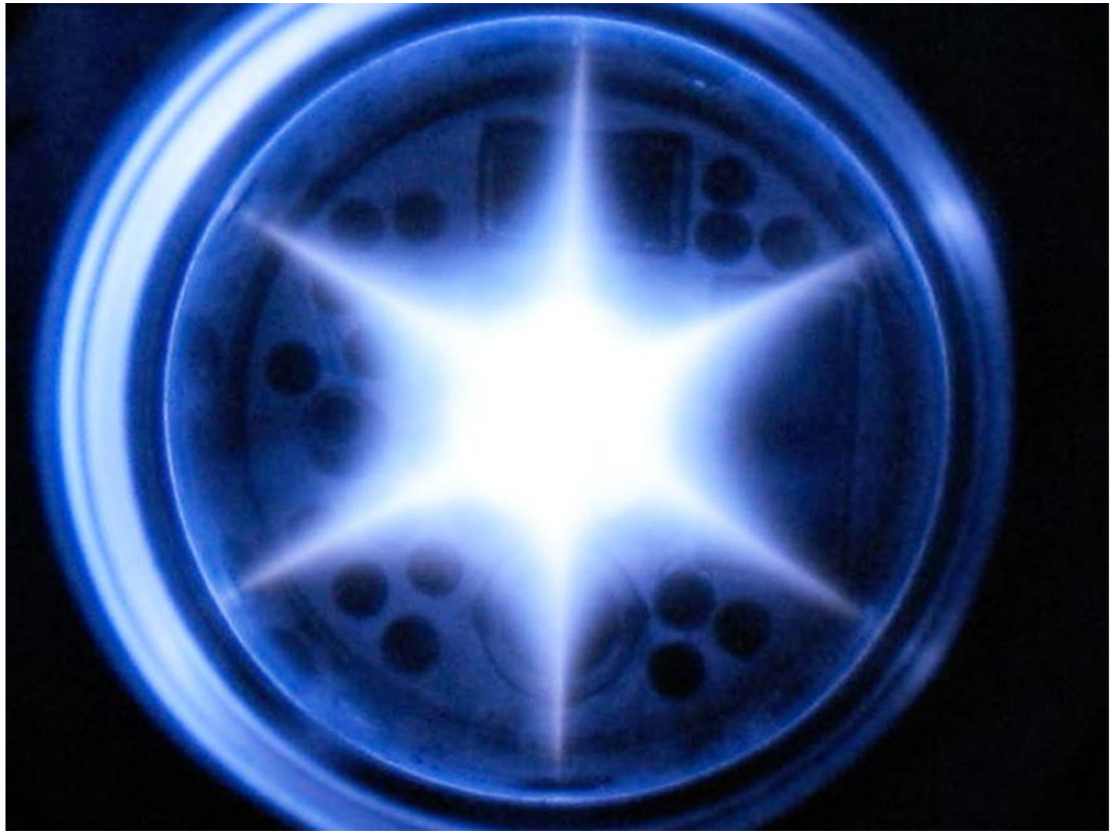

Computer-driven baking

Two-step process:

- Source to full RF power
  - Monitor pressure, currents, etc. and raise safely
  - Time significantly reduced as the computer is persistent
- "Slosh" plasma around to bake plasma chamber
  - Change confining fields and add gas when desorption rate drops
  - Every 6th change return to "base" fields and take charge state distribution to monitor

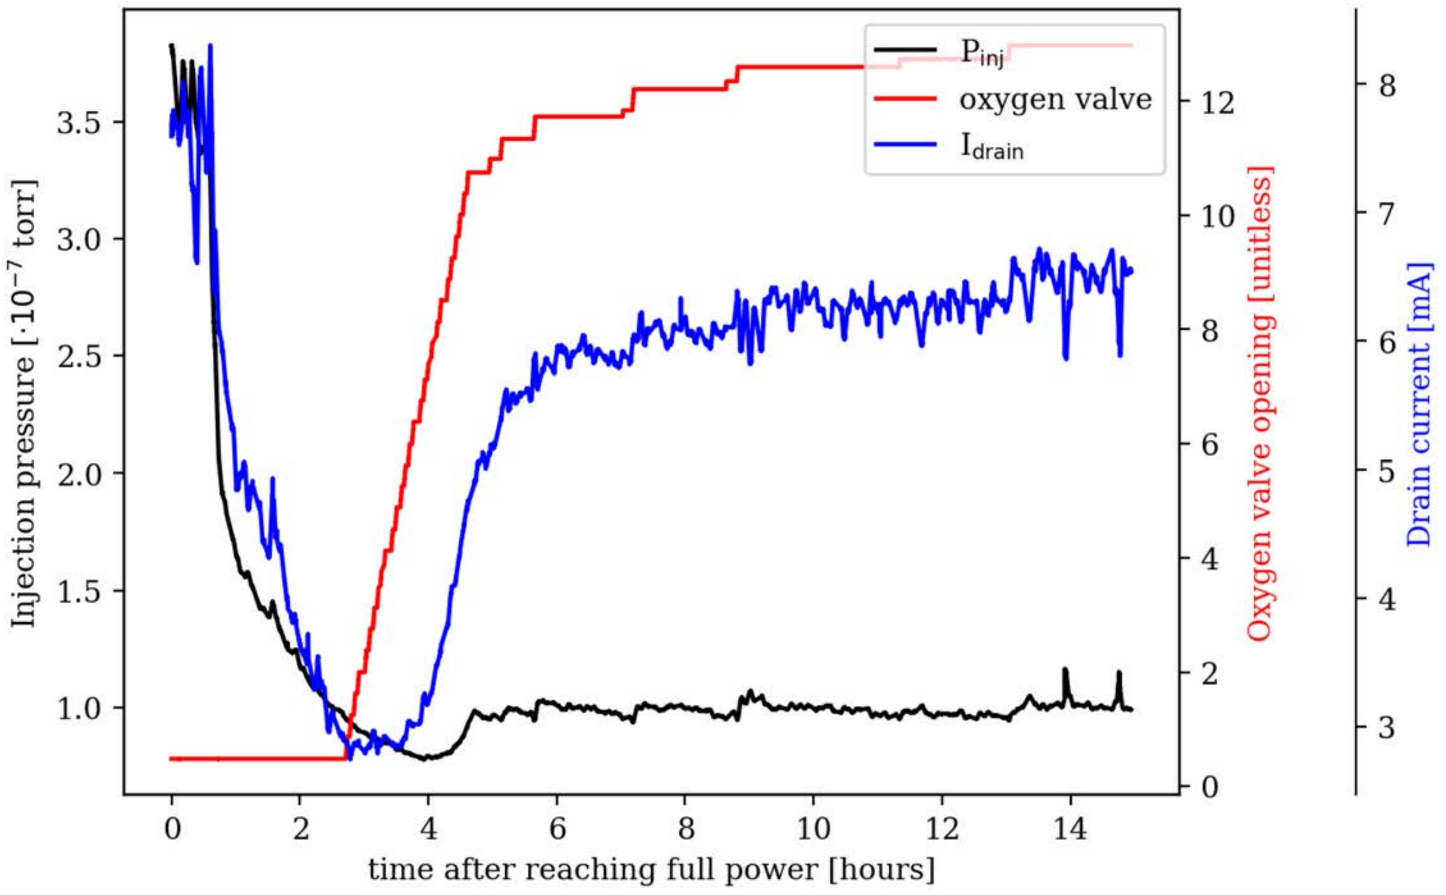

Charge state distributions (CSDs) give feedback on progress

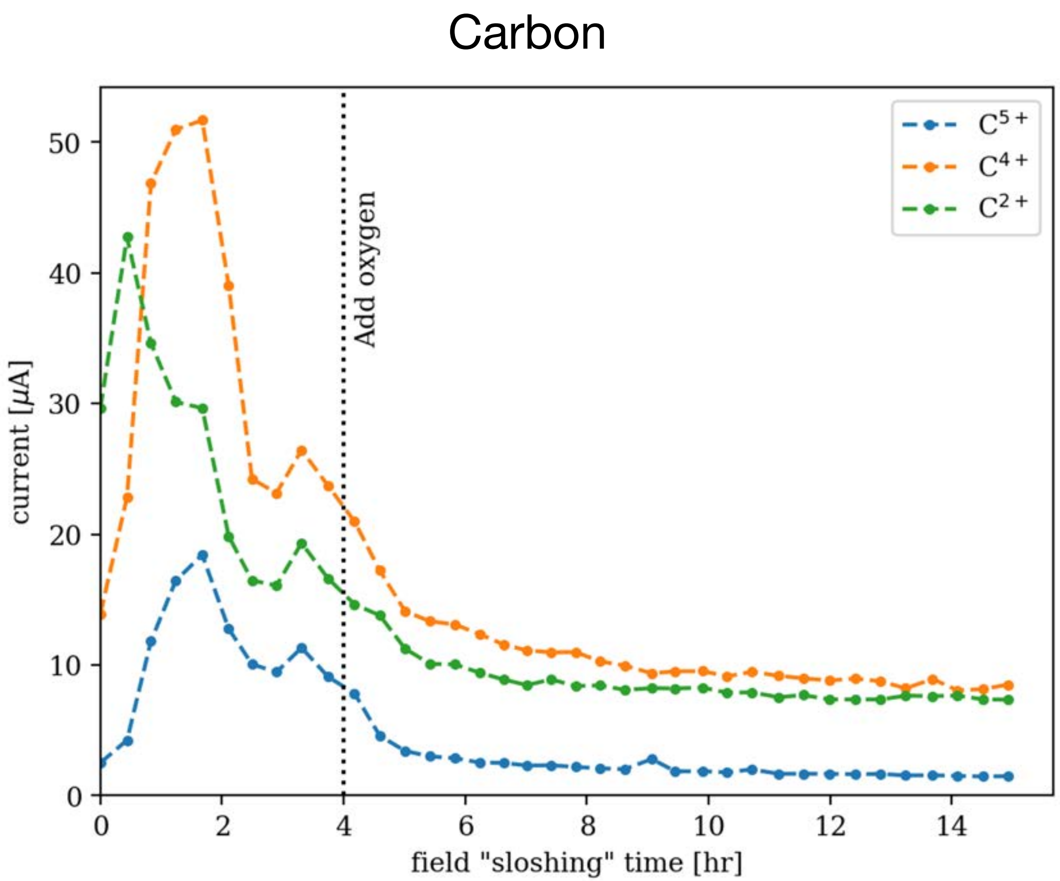

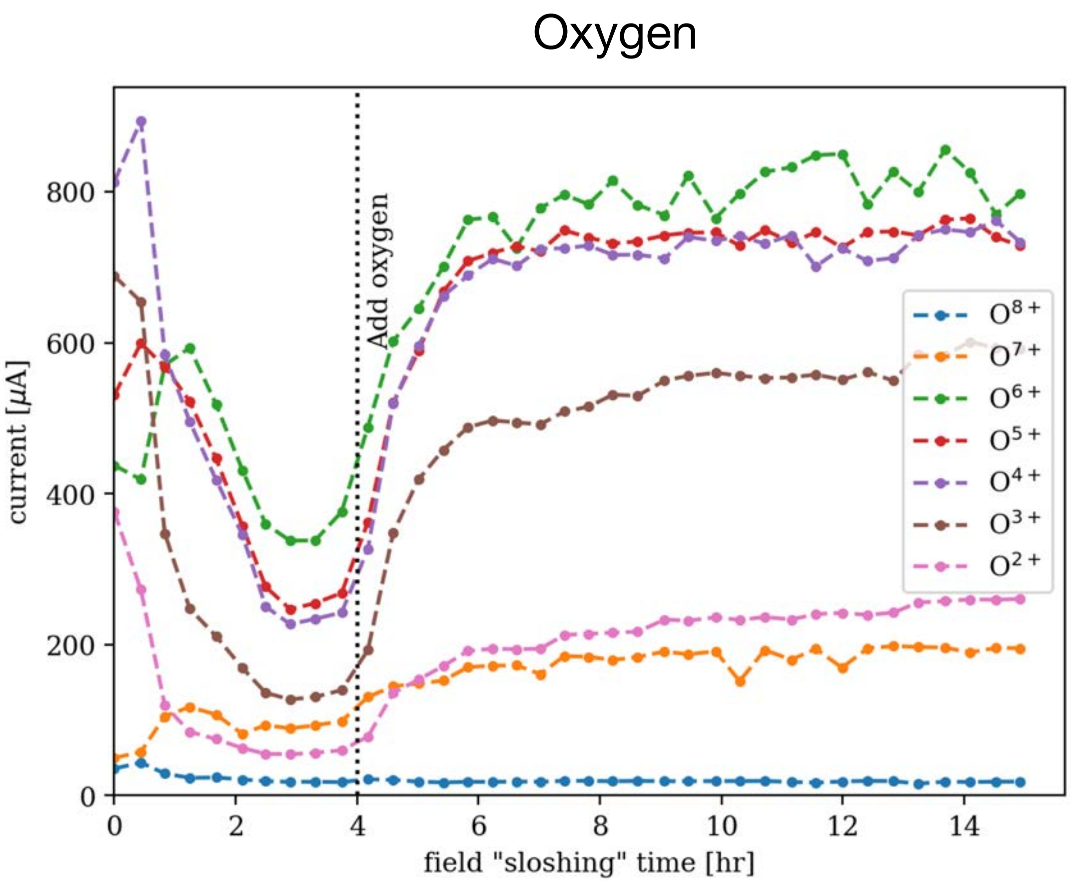

- Last ~6 bakes have been fully automatic
- Not machine learning, but data will be used to produce more efficient baking technique (ML!)

Summary and looking forward

- The addition of fast CSDs has provided hope for machine learning with ECR ion sources
  - Can be completed in less than 30 seconds
  - Performing CSDs will become part of routine running even when delivering beam to cyclotron
- Bayesian optimization in place for maximizing current with no human interaction:
  - spare chamber delivered Fall 2024, and will set ML loose for record attempt this year
- We have thousands of hours of full-computer-control of VENUS for baking and beam optimization (limited by VENUS availability---we need more time and data!)
- Next steps:
  - Optimize on emittance
  - Add ML-usable diagnostics (stability, etc.)

Thank you for your attention!