# Simulation informed particle accelerator`s optimizers

2025/07/10~11

ムーンショット目標10奥野プロジェクトワークショップ「核融合分野に貢献する１A重陽子線形加速器とCARA加速器の実現可能性」

Sumitomo Heavy Industries, Ltd.

DE FRANCO Andrea

**Simulation informed particle accelerator`s optimizers**


Reframing beam halo simulations

Introduction to RL and BO

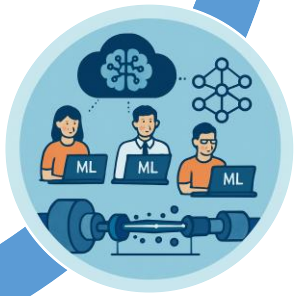


3 success stories

# Beam Simulation needs

Continuous demand for particle accelerators with higher intensities.
Control of beam losses to better than 1 in 1 millionth part starts to be necessary.

  

# Beam Simulation needs

Continuous demand for particle accelerators with higher intensities.
Control of beam losses to better than 1 in 1 millionth part starts to be necessary.

How can we simulate to this precision?
Can we just decrease PIC mesh size, increase number of tracked particle?
Is there additional interactions with beam to be included? (e-cloud, walls, dust particles, etc.)
Even if we achieve perfect codes can we perform measurements good enough to obtain realistic input distributions? Or to validate the simulations?
How do we use perfect beam simulations? Calibration and imperfection from HW becomes main issue!

# Beam Simulation needs

Continuous demand for particle accelerators with higher intensities.
Control of beam losses to better than 1 in 1 millionth part starts to be necessary.

How can we simulate to this precision?
Can we just decrease PIC mesh size, increase number of tracked particle?
Is there additional interactions with beam to be included? (e-cloud, walls, dust particles, etc.)
Even if we achieve perfect codes can we perform measurements good enough to obtain realistic input distributions? Or to validate the simulations?
How do we use perfect beam simulations? Calibration and imperfection from HW becomes main issue!

If we cannot predict beam halo in simulations to the required uncertainty, what else can we do?

# Reframe into an optimization problem

**If you cannot predict the halo, control its effects.**

Reframe the problem into an optimization of the real machine parameters with the objective to minimize the observed losses from the halo.
Black box optimizers want to explore strange parameters space, and this cannot be allowed in high intensity accelerators for machine safety reasons.
→ Inform optimizers of the well-known physics from the beam core simulations.

<u>Necessary Assumptions:</u>

- HW design with enough margins that unknown halo can be transported or lost where we want (high power slits).
- Simulation capable to predict beam core are sufficiently developed

# Optimizer of particle accelerators

**Old topic, recently growing very rapidly. Main reasons:**

- Complexity of particle accelerators keep growing, while we must keep reducing costs and timeline in any project (automation!).
- Rapid development of methods in the machine learning/AI community.

Today I want to present 3 success stories from the accelerator community using ML tools.

Simulation informed particle accelerator`s optimizers


Reframing beam halo simulations

**Introduction to RL and BO**


3 success stories

# Reinforcement Learning (RL) in a nutshell

Agent takes an action per step

Allowed actions

Agent

Agent gets a reward if gold is found.

For optimization we carefully craft the reward. For example:
Reward→ get higher output value
Punishment→ number of moves (to avoid being slow)

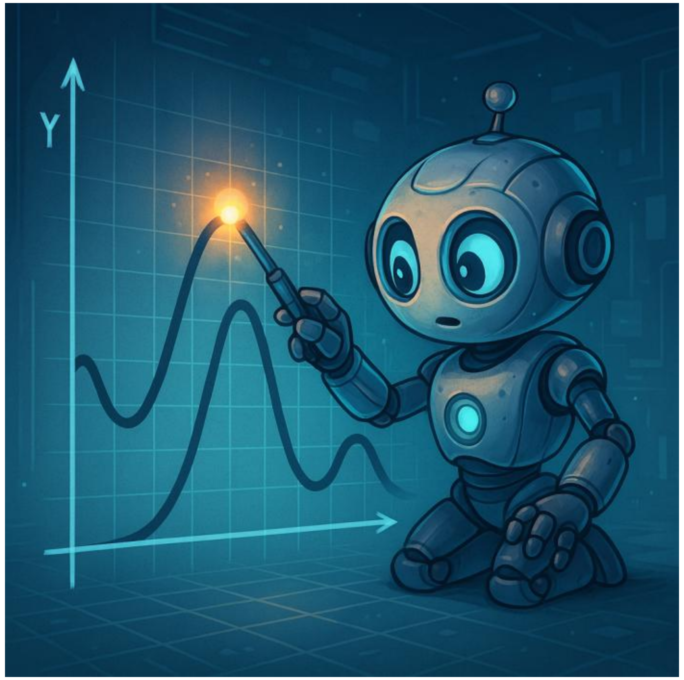


The agent training requires order of ~$10^6$ times interaction with the environment.
Often, we use simulations for training and then let trained agent work with real world environment.

# Bayesian Optimization (BO) in a nutshell

Example: we want to find what value of parameter x maximize the output y

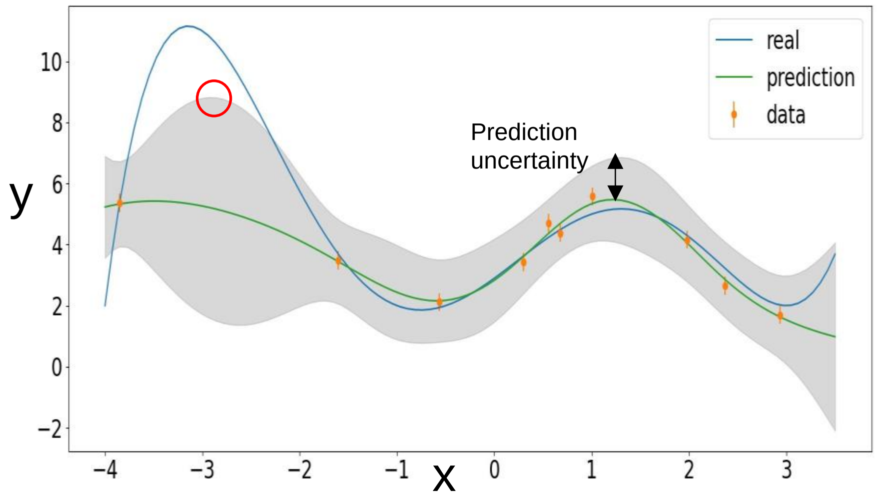


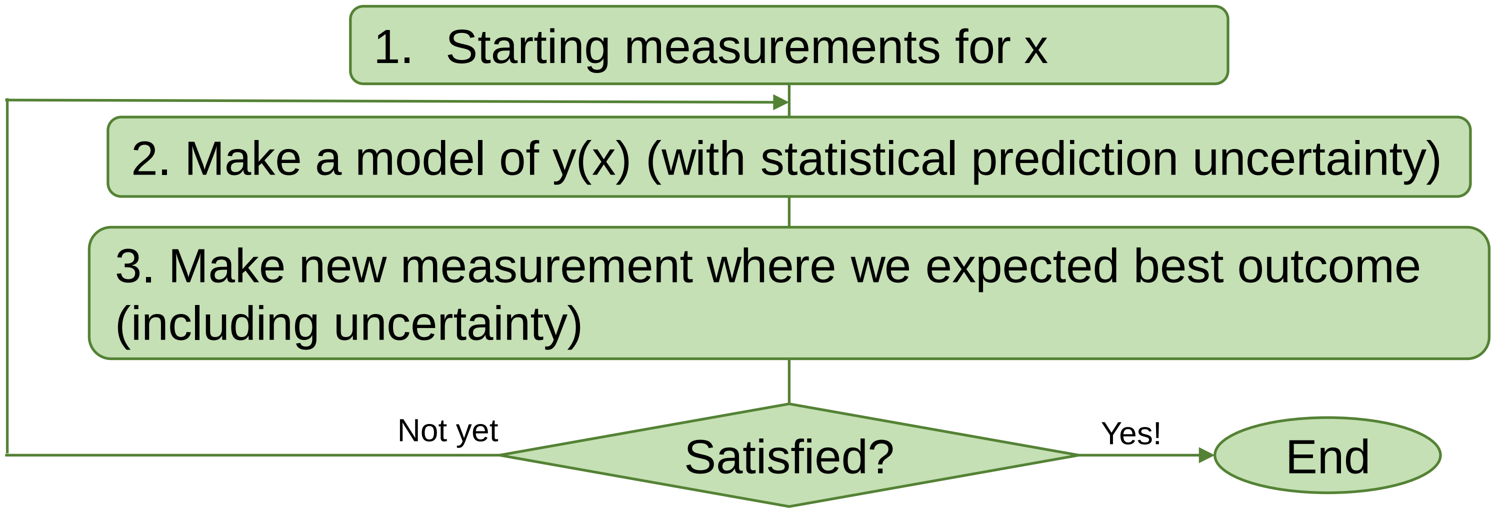


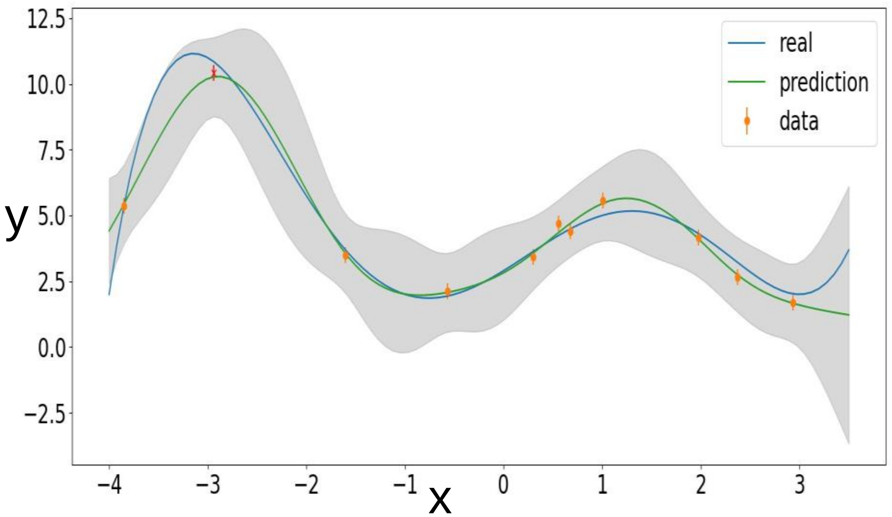


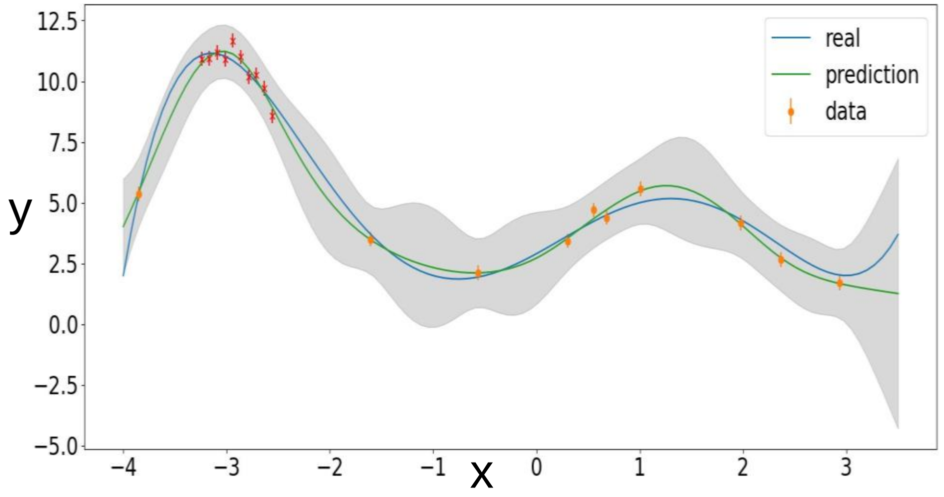

**Simulation informed particle accelerator`s optimizers**


Reframing beam halo simulations

Introduction to RL and BO


**3 success stories**

# Beam size/alignment @ ARES (DESY)

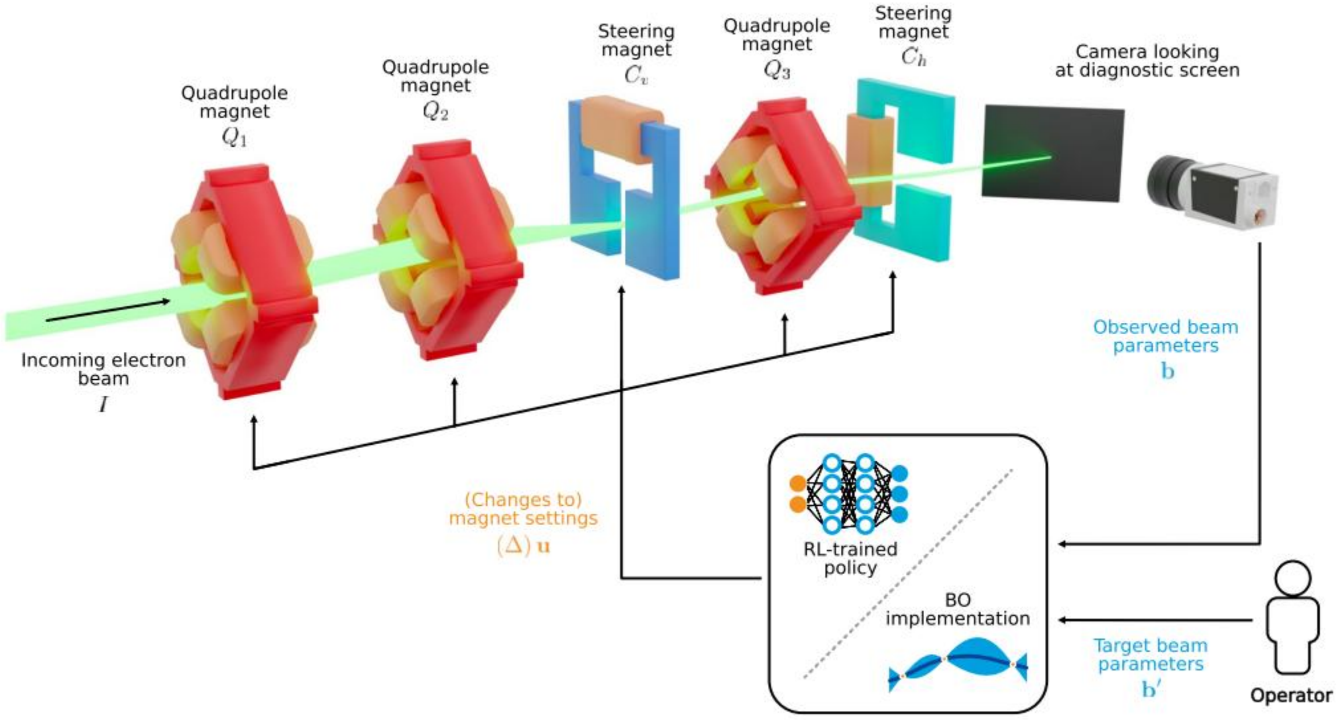


Parameters: strength of 3 quads + 2 steerers
Goal: required sizes + position on a screen
Technique: Reinforcement Learning (RL) (also compared to Bayesian Optimization (BO )

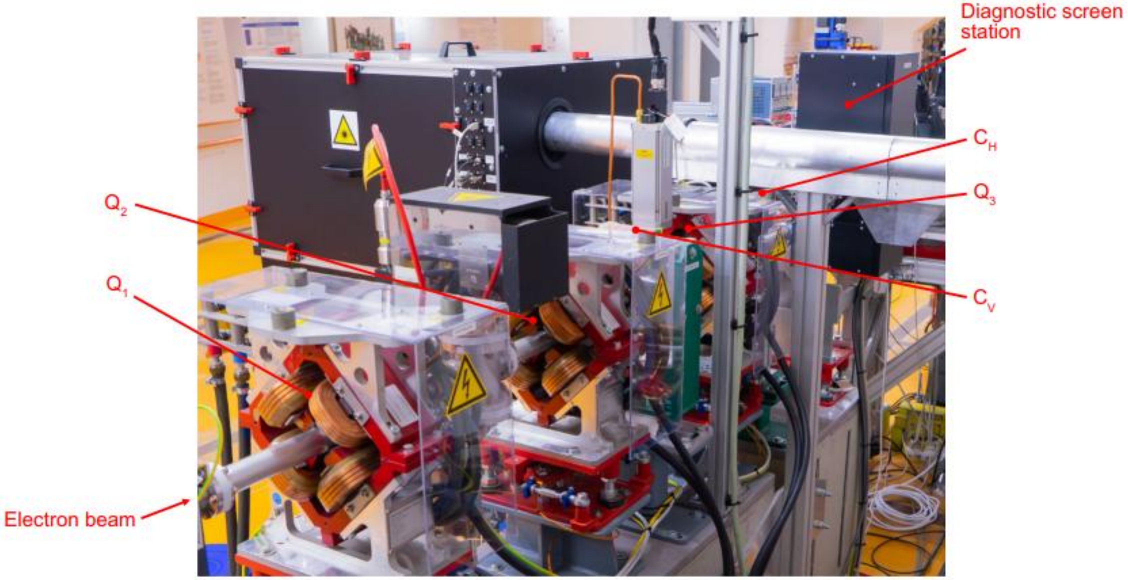


Method in a nutshell:
Train RL with simulations ($10^6$).
Random magnets tilts/misalignment are included to teach agent a close to real world method.
Let trained agent optimize in real world.

Kaiser J. et al., Sci. Rep. 14, 15733 (2024)

# RL Vs. BO

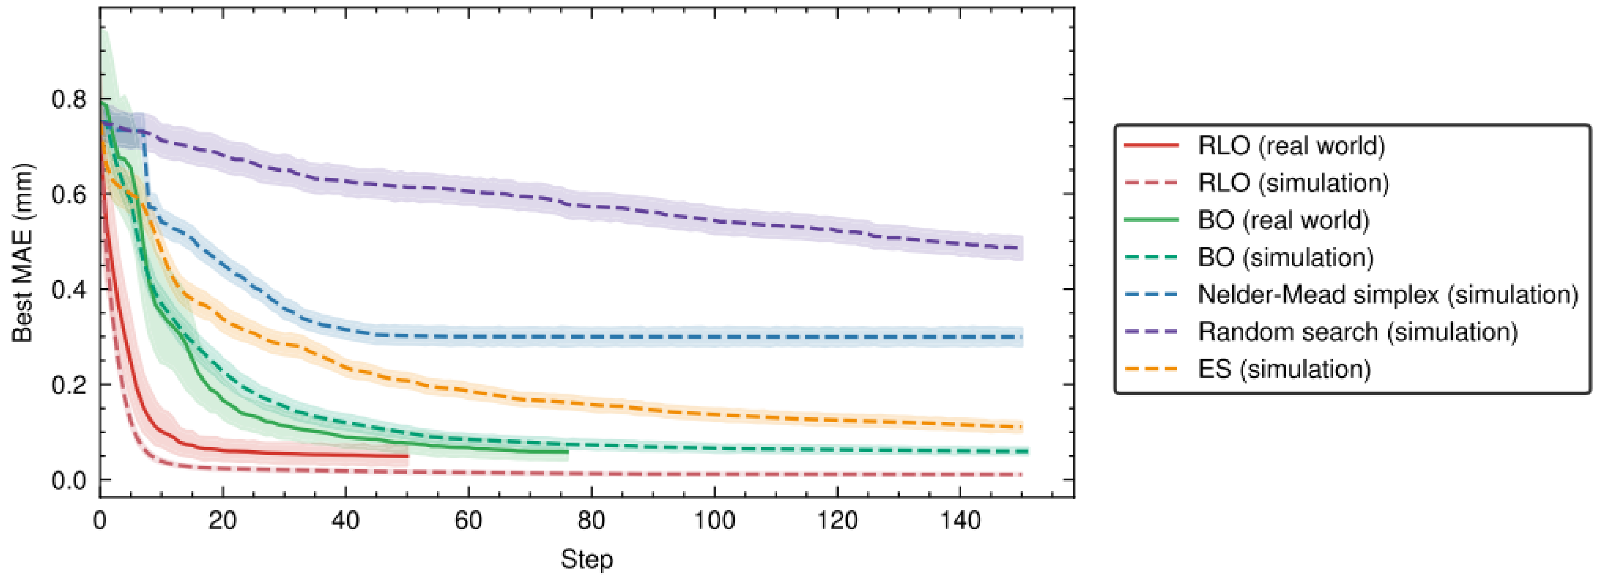


| Optimiser | Final beam difference (µm) | | Steps to target | | | Steps to convergence | |
|---|---|---|---|---|---|---|---|
| | Median | Mean | Median | Mean | Success rate | Median | Mean |
| Real world (finite diagnostic screen, beam aligned to quadrupole magnets) → Section "Real-world study" | | | | | | | |
| RLO[a,b] | 24 | 29 ± 48 | 12 | 13 ± 7 | 55% | 7 | 8 ± 4 |
| BO[a,c] | 44 | 58 ± 41 | 45 | 48 ± 12 | 45% | 13 | 17 ± 11 |
| Simulation (finite diagnostic screen, no beam alignment) → Section "Robustness in the presence of sensor blind spots" | | | | | | | |
| RLO | 4 | 13 ± 40 | 7 | 12 ± 16 | 93% | 6 | 7 ± 4 |
| BO | 38 | 61 ± 96 | 44 | 53 ± 28 | 53% | 20 | 30 ± 30 |

RL achieves better values
RL converges faster
RL & BO similar success rate

Kaiser J. et al., Sci. Rep. 14, 15733 (2024)

# RL Vs. BO

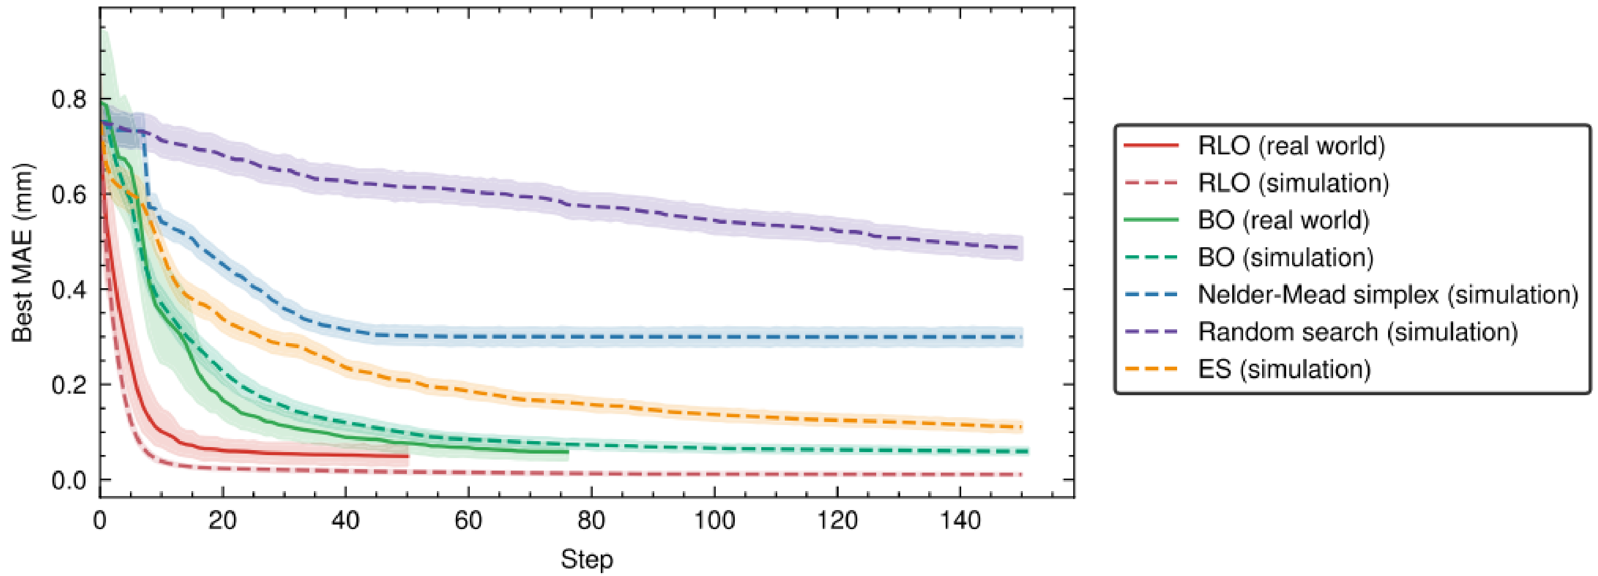


| Optimiser | Final beam difference (µm) | | Steps to target | | | Steps to convergence | |
|---|---|---|---|---|---|---|---|
| | Median | Mean | Median | Mean | Success rate | Median | Mean |
| Real world (finite diagnostic screen, beam aligned to quadrupole magnets) → Section "Real-world study" | | | | | | | |
| RLO[a,b] | 24 | 29 ± 48 | 12 | 13 ± 7 | 55% | 7 | 8 ± 4 |
| BO[a,c] | 44 | 58 ± 41 | 45 | 48 ± 12 | 45% | 13 | 17 ± 11 |
| Simulation (finite diagnostic screen, no beam alignment) → Section "Robustness in the presence of sensor blind spots" | | | | | | | |
| RLO | 4 | 13 ± 40 | 7 | 12 ± 16 | 93% | 6 | 7 ± 4 |
| BO | 38 | 61 ± 96 | 44 | 53 ± 28 | 53% | 20 | 30 ± 30 |

RL achieves better values
RL converges faster
RL & BO similar success rate
RL degraded when simu → real world

Kaiser J. et al., Sci. Rep. 14, 15733 (2024)

# Beam spot/alignment @ ARES (DESY)

Applicability to 1A class accelerator

In this example:

- 👍 Physics information from simulations
- 👍 Fast and robust convergence

- ~ Large number of simulations required for training (~$10^6$)
- ✗ Low number of parameters (~5)
- ✗ Quality of simulations affects the result (is it applicable without halo simulations?)
- ✗ Simulations should be backward differentiable (not trivial for space charge)

# Beam Halo loss min @ LIPAc/IFMIF (QST 六ヶ所)

Parameters: strength of 4 quads + 4 steerers
Goal: minimize beam halo losses (vacuum pressure)
Technique: Bayesian Optimization (BO) + Beam core simulations

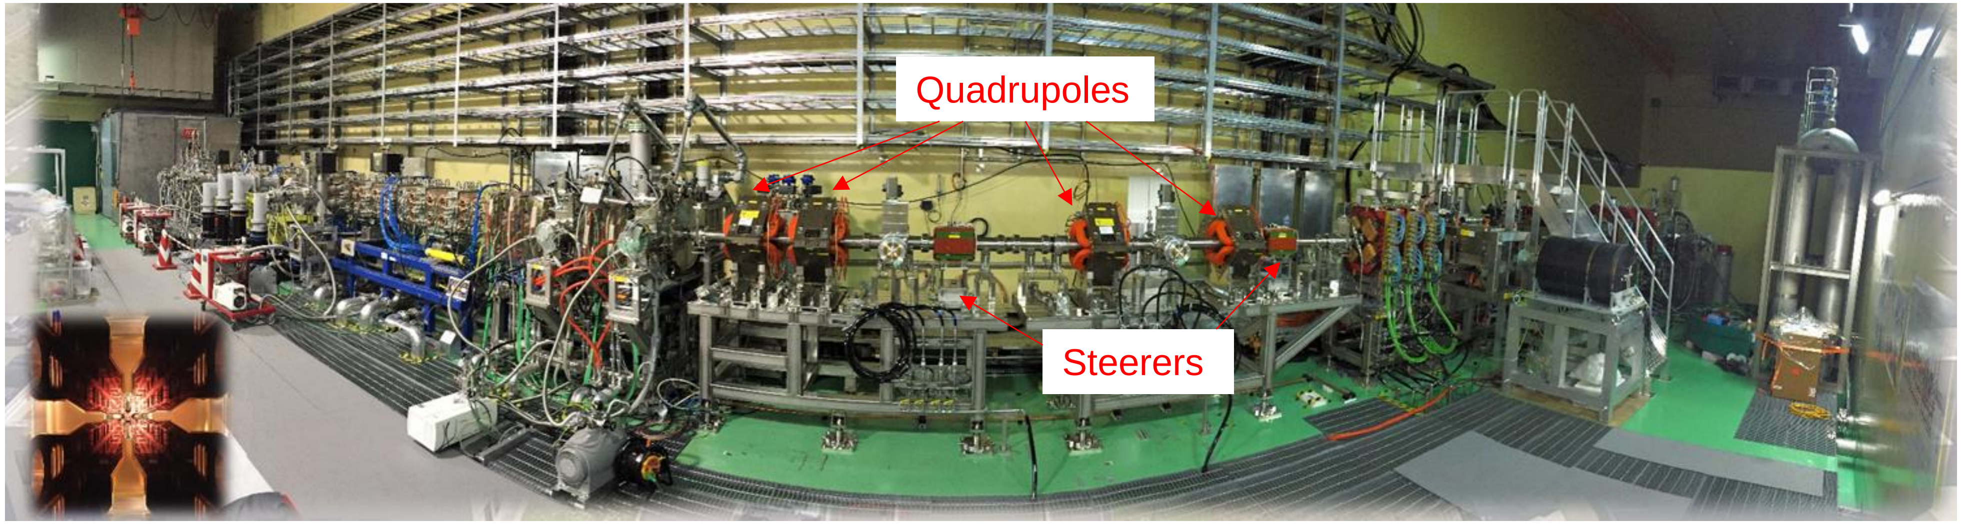


Method in a nutshell:
Use BO exploration algorithm to model beam halo losses.
Use beam core simulations to constraint the parameter space to safe settings only.
The vacuum pressure is used as indicator of beam halo losses.

De Franco et al., Malapa 2025

# Method in a nutshell

De Franco et al., Malapa 2025

Vacuum BO surrogate model

*Vac(8x magnets)*

Simulations of core losses

*CoreLoss(8x magnets)*

Acquisition function

$$UCB(Vac) - 10^2\left(\frac{1}{1 + e^{(-CoreLoss+\alpha)}}\right)$$

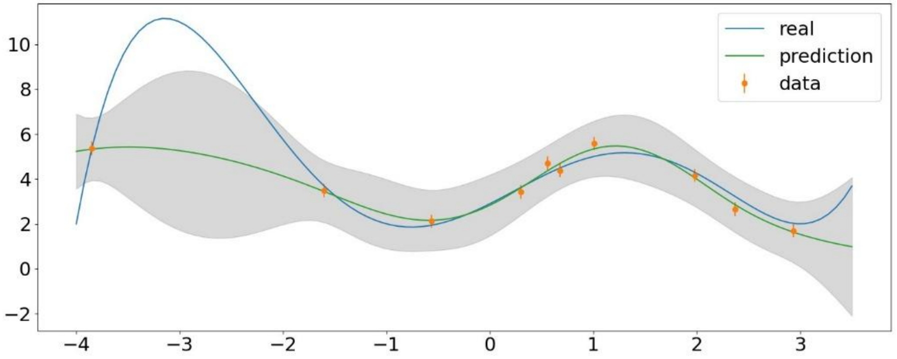


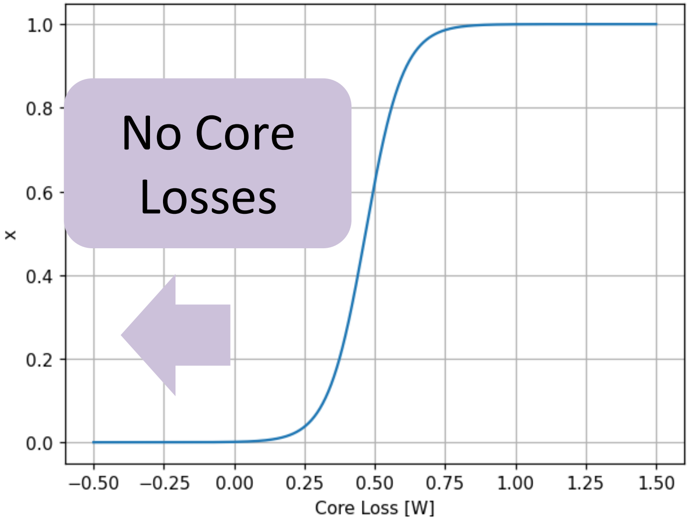

# Beam halo losses reduction- physics interpretation

Vacuum pressure dropped to ~1/3!
Physics interpretation of choices ML did:

1. Weaker focusing → larger beam → less halo formation
2. Compromise of transport between halo & core

Experimental Optics

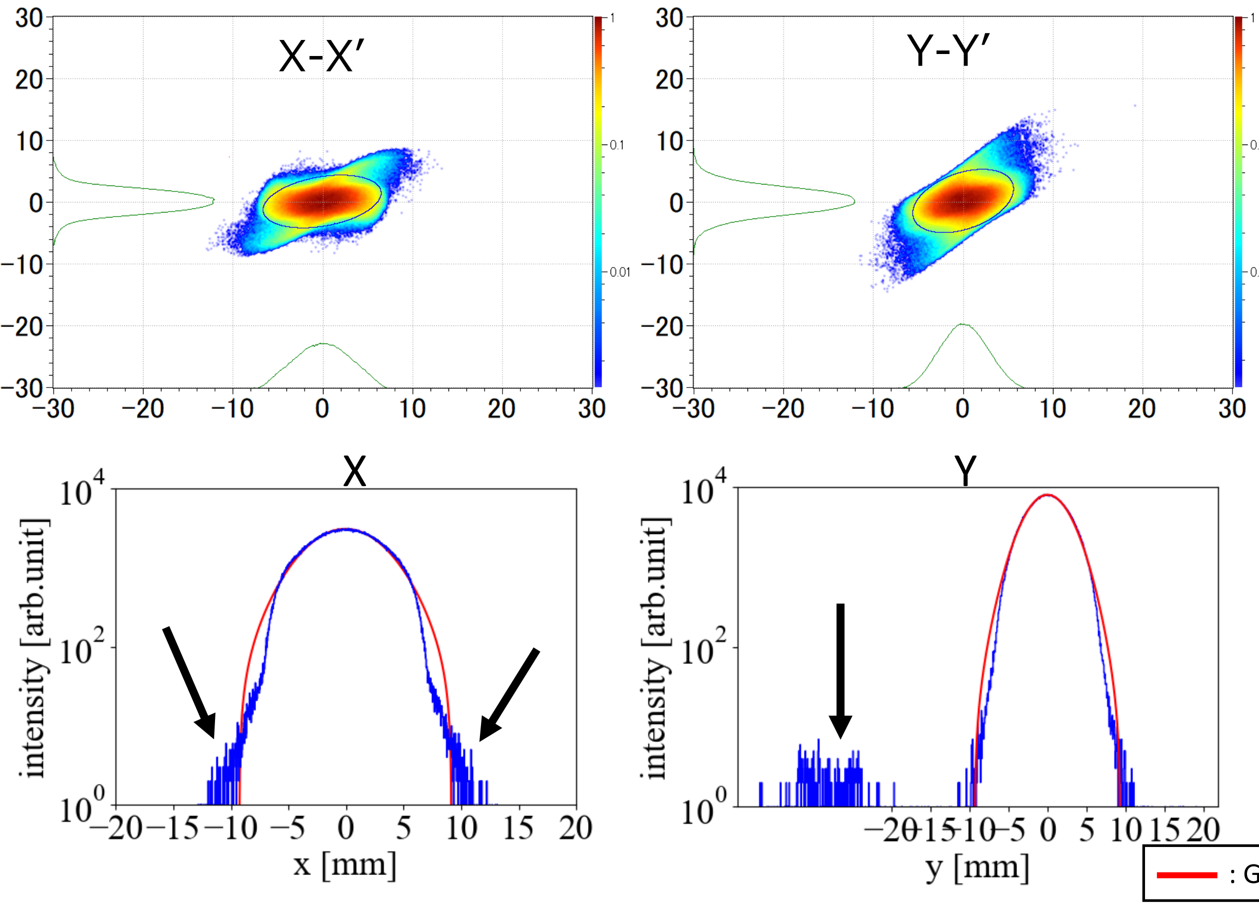


ML Optics

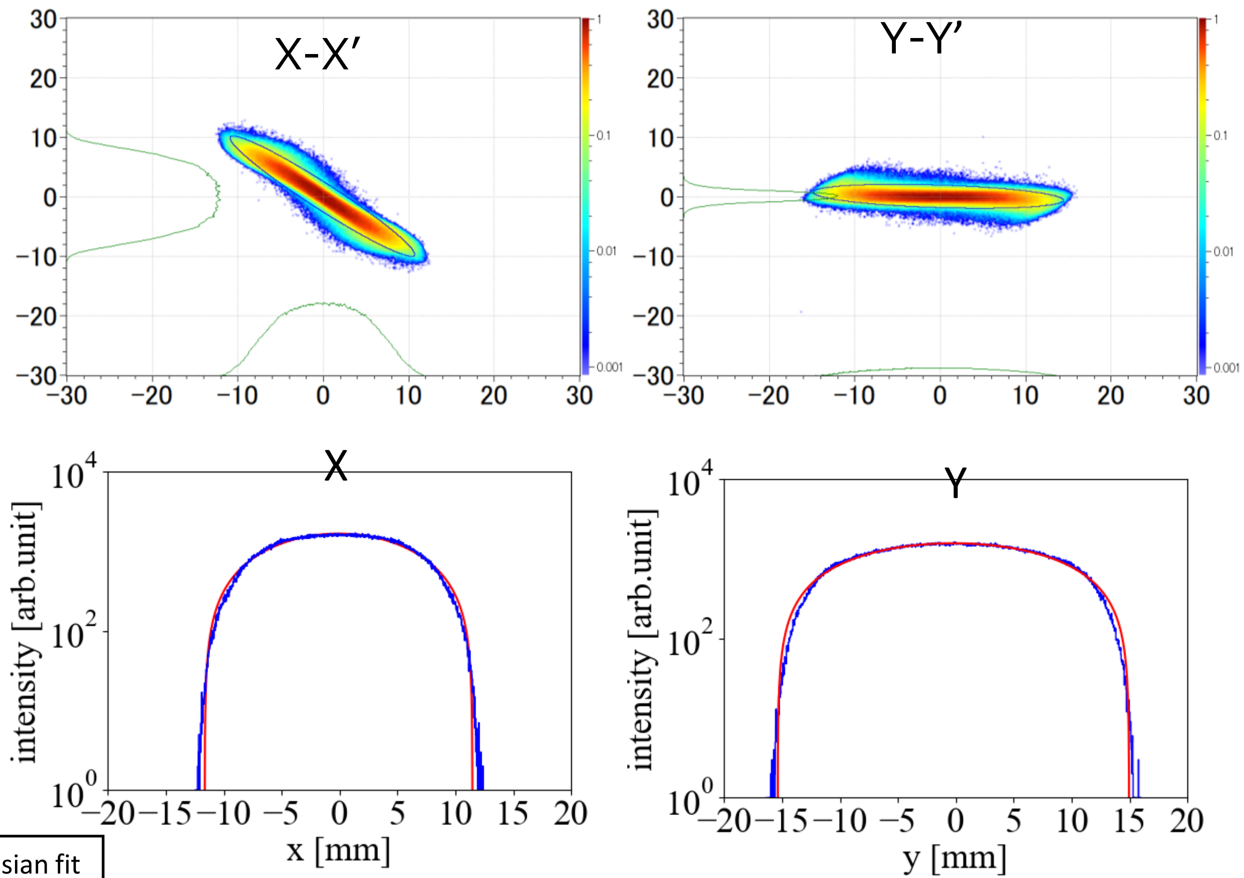


: Gaussian fit

De Franco et al., Malapa 2025

# Beam Halo loss min @ LIPAc/IFMIF (QST 六ヶ所)

Applicability to 1A class accelerator
In this example:

- 👍 Physics information from simulations
- 👍 Halo losses optimized even if simulated poorly
- 👍 Fast convergence

- ❌ Robustness over many trials not proven
- ❌ Low number of parameters (~8)

# XFEL optimization @ SACLA (RIKEN)

Parameters: strength of quads & RF phases
Goal: optimize XFEL brightness
Technique: Bayesian Optimization (BO)

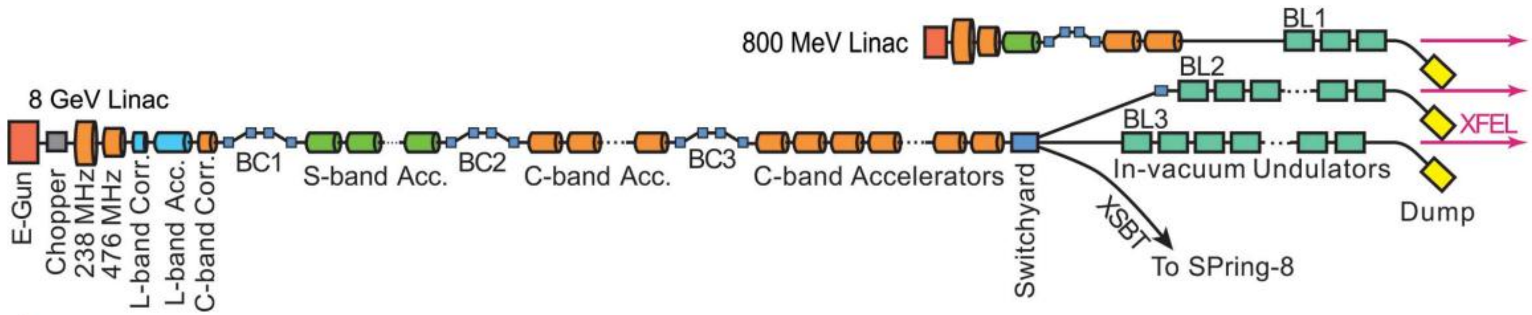


**Figure 1**
A schematic view of the SACLA accelerator.

Iwai et al., J. Synchrotron Rad. (2023). **30**, 1048-1053

**Table 1**
Tuning knobs of SACLA to be adjusted by the BO.

| Parameter type | Number of parameters |
|---|---|
| RF phases | 7 phases |
| Magnetic lens currents | 9 units |
| Quadrupole-magnet currents | 62 units |
| Steering-magnet currents | 8 units |
| Undulator gaps (*K* values) | 7 |

Method in a nutshell:
Use standard BO.
Improve robustness with smart limits/hyperparameters choice.
Too many parameters for standard BO, the problem is split into sections.

# XFEL optimization @ SACLA (RIKEN)

Achieved a significantly better result.
Robust automation becomes an essential operator`s tool.

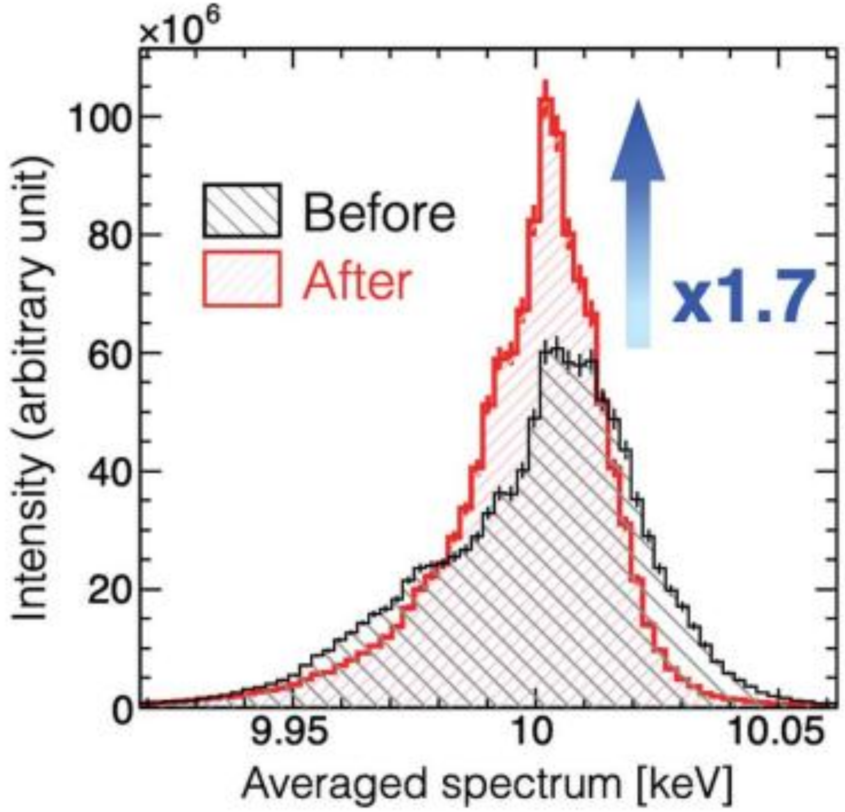


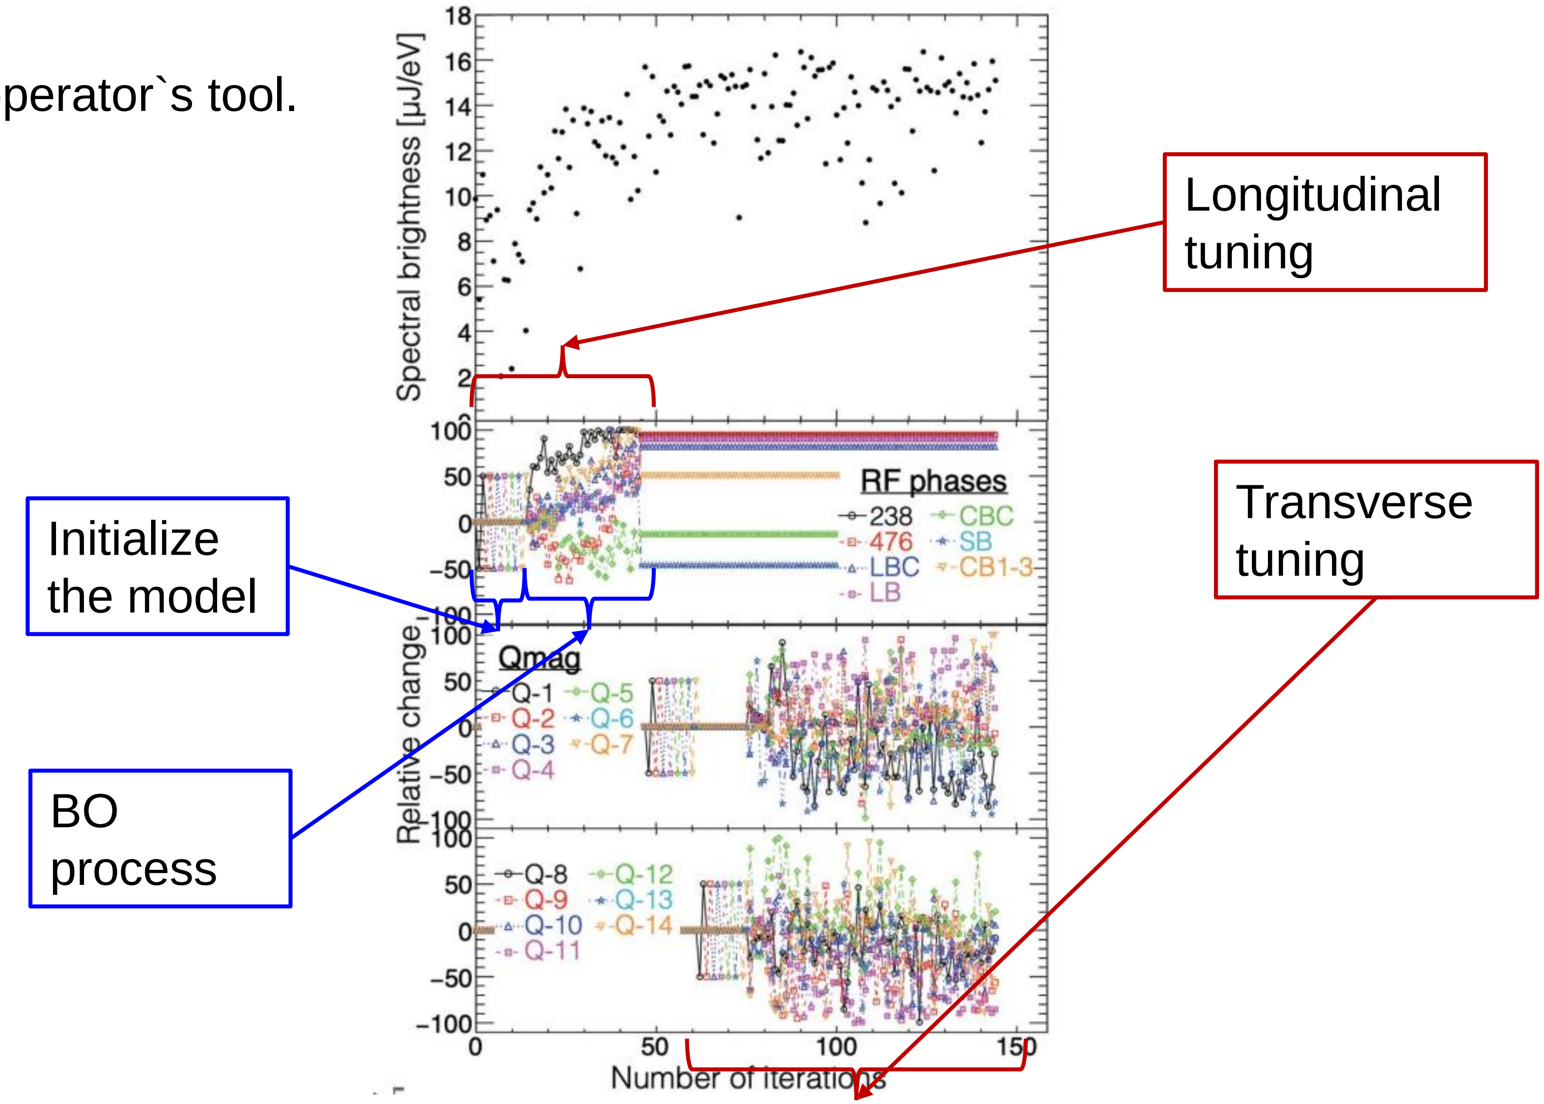


Iwai et al., J. Synchrotron Rad. (2023). **30**, 1048-1053

# XFEL optimization @ SACLA (RIKEN)

Applicability to 1A class accelerator

In this example:

- 👍 Fast and robust convergence
- 👍 Large number of parameters (but had to split in sections)

- ✕ No information from physics
- ✕ No safety functions, all settings are considered safe

Simulation informed particle accelerator`s optimizers


Reframing beam halo simulations

Introduction to RL and BO


3 success stories

# Conclusions & future works

- Developing simulations with 1/1millionth precision is challenging
- Even if successful, limited by input measurements uncertainty
- Argued a paradigm shift: simulation ~10% uncertainty + smart tuning
- Showcased recent success stories from community

**What gap is still necessary to fill to effectively control beam halo for a 1A class accelerator?**
Optimizer with many parameters + Beam core physics informed
(+fast & robust is a must for industrial applications)

This problem has many synergies with automation of procedures for industrial applications.
R&D at SHI ongoing on efficient and safe optimizers for particle accelerator that are physics informed from not necessarily perfect simulations.

Sumitomo Heavy Industries, Ltd.

All forward-looking statements regarding the company's future performance are based on information currently available to Sumitomo Heavy Industries and determined subjectively. Future performance is not guaranteed and all information related to future performance contained herein is subject to changes in business environments.