Machine Learning Tools
Improve BESSY II Operation

L. Vera Ramírez, T. Birke, G. Hartmann, R. Müller, M. Ries, A. Schälicke, P. Schnizer

Helmholtz-Zentrum Berlin

luis.vera_ramirez@helmholtz-berlin.de

18th International Conference on Accelerator and Large Experimental Physics Control Systems (ICALEPCS), 21/10/2021

From https://www.helmholtz-berlin.de/forschung/quellen/bessy/bessy-in-zahlen_en.html

- Located in Berlin-Adlershof
- Circumference: 240 m
- Bending magnets: 32
- Straight sections: 16
- Beamlines: $\sim$ 45
- Electron energy: 1.7 GeV
- Nominal beam current: 300 mA
- Energy of the synchrotron radiation: 1eV to 150 keV
- Duration of the light pulse: 20 ps
- Measurement time: $\sim$ 5000 h estimated for 2021
- User visits from guest researchers: $\sim$ 3000 p.a.
- Beam + Machine Commissioning: $\sim$ 1000 h p.a.

BESSY II - Some ML Use Cases

Use case @ BESSY II	Potential gain	ML Paradigm
Measurement prediction (lifetime, vertical beam size...) [VRHM+20]	Anticipating and understanding accelerator behaviour	Supervised Learning
Booster current & injection efficiency optimisation [VRHM+20]	Automation (avoiding manual tuning)	Reinforcement Learning
Beamline raytracing inversion – *Dr. G. Hartmann & Team*	Performance improvement & automation	Supervised & Reinforcement Learning
Data compression of a photoelectron diagnostic unit with latent space labelling – *Dr. G. Hartmann & Team*	Performance improvement & knowledge generation	Unsupervised Learning
Mitigation of harmonic orbit perturbations	Improvement of beam stability in analitically non-tractable case	Supervised & Reinforcement Learning
Anomaly detection system with feature assignation	Additional support for operators	Unsupervised Learning

Table of Contents

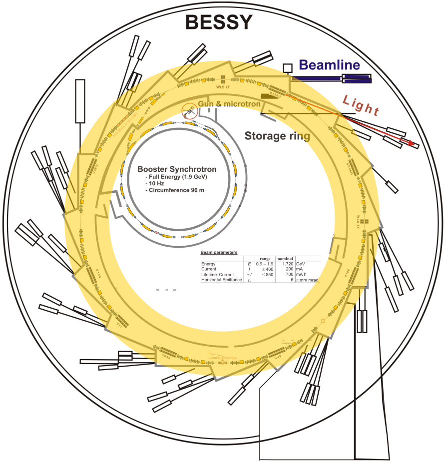

- Electron beam stability is a critical factor in order to achieve light radiation with high quality brilliance and brightness over time.
- Transverse beam dimensions at BESSY II: $100 \times 20\mu$m.
- Magnet position errors destabilize the orbit $\rightarrow$ efficiently corrected with *traditional* methods - BESSY II: **fast orbit feedback (FOFB)** ([MGL+09]).

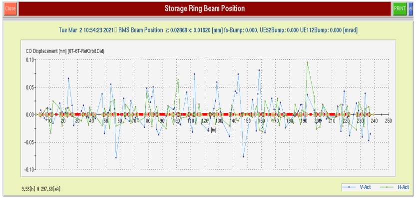

Response-matrix-based correction in a nutshell:

$$\Delta \mathbf{x} \approx S \Delta \mathbf{c}$$

with $\mathbf{x}$ relative beam position, $\mathbf{c}$ corrector magnets strength and S response matrix (calculated or measured on the accelerator). Hence, for $\mathbf{x}_{t+1} = \mathbf{0}$, we apply recursively

$$\mathbf{c}_{t+1} := \mathbf{c}_t - \alpha S^{-1} \mathbf{x}_t$$

with S^{-1} Moore-Penrose pseudoinverse of S (at BESSY $\alpha = 0.8$).

- But the environment also produces additional **perturbations** interfering with user operation:
  - *Civil noise*
  - Main power (50Hz)
  - Some imperfectly isolated magnetic sources (e.g. booster power supply at 10Hz)
  - The FOFB system itself (20-40Hz).
- Olivier Churlaud faced this problem in his MSc ([Chu16])
- Follow-up: mitigate the harmonic perturbations with a **RL agent learning the perturbation dynamics** from observation and interaction with the environment
- **Early 2020**: simulations ([Chu16], [AGTZ14])
- **Since Summer 2020**: tests at the machine with Bluesky ([ACCR19], [SMS$^+$21]) and zmq

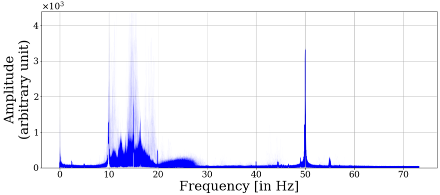

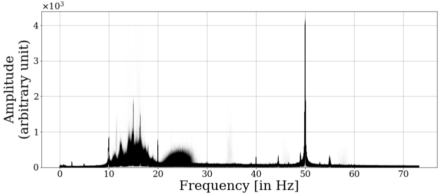

Mitigation of Harmonic Orbit Perturbations

- Model-free approach
- Towards a model-based approach
- Interpretability

Anomaly detection with feature assignation

Summary and Outlook

References

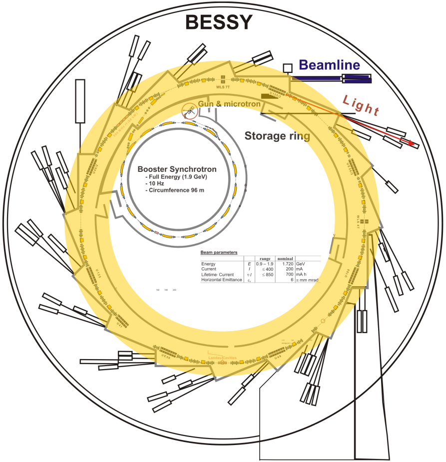

Model-free approach: Test for BPMZ5D7R, 1/3/21

- **State**: last readbacks from 6 top-correlated beam position monitors (BPMs) (Δ's w.r.t. the reference orbit in mm)
- **Action**: strength of 6 top-correlated horizontal steerers (Δ's w.r.t. the initial stregth, $\pm$ 20 mA)
- **Reward**: transformation of the next readback from reference BPM: $2e^{-20\,|x_{t+1}|} - 1$
- **Algorithm**: DDPG ([LHP$^+$16])
- **Baselines**
  - *Static steerers*: steerers set back to the initial precalculated strength
  - *FOFB*: standard FOFB system
  - *FOFB (zmq)*: FOFB reimplemented in the zmq loop.

Time domain

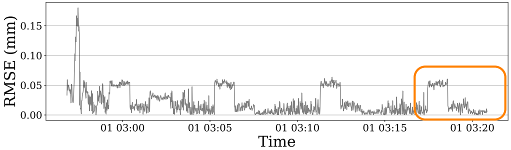

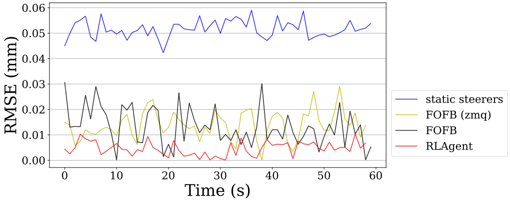

Frequency domain

Archiver beam motion variable

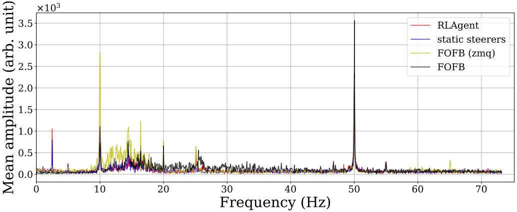

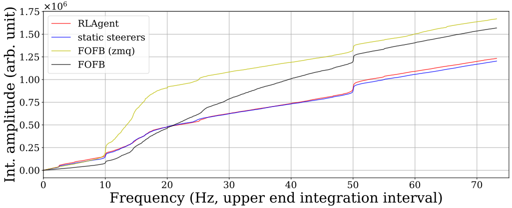

RMSE from on-the-fly data

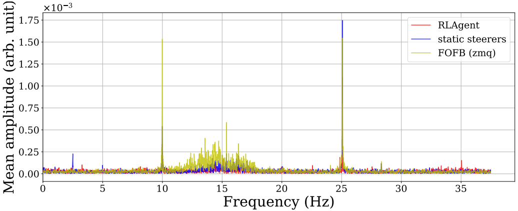

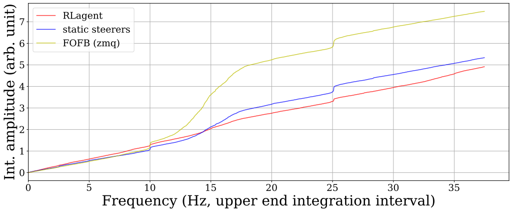

- **State**: last readbacks from 6 top-correlated beam position monitors (BPMs) (Δ's w.r.t. the reference orbit in mm)
- **Action**: strength of 6 top-correlated horizontal steerers (Δ's w.r.t. the initial stregth, $\pm$ 20 mA)
- **Reward**: transformation of the next readback from all BPMs (m = 102):

$$2e^{-20\,\mathrm{RMSE}[\mathbf{x}_{t+1}]} - 1 = 2e^{-20\sqrt{\frac{\sum_{i=1}^{m}\left(x_{t+1}^{i}\right)^2}{m}}} - 1$$

- **Algorithm**: DDPG ([LHP$^+$16])
- **Baselines**
  - *Static steerers*: steerers set back to the initial precalculated strength
  - *FOFB*: standard FOFB system
  - *FOFB (zmq)*: FOFB reimplemented in the zmq loop.

Time domain

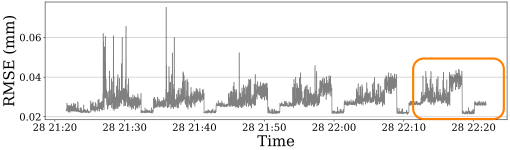

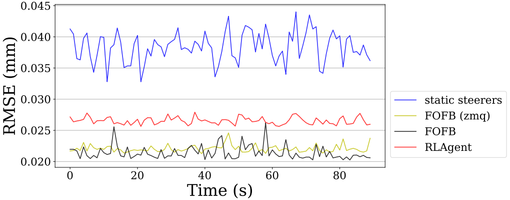

HZB Helmholtz Zentrum Berlin

Frequency domain

Archiver beam motion variable

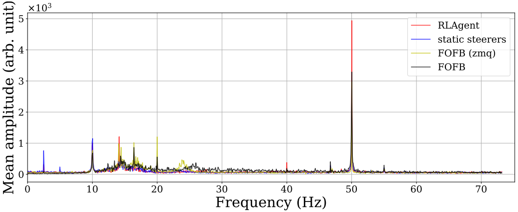

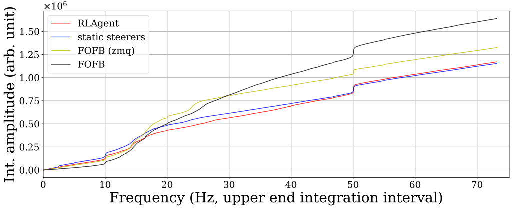

RMSE from on-the-fly data

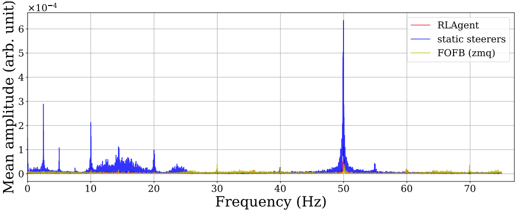

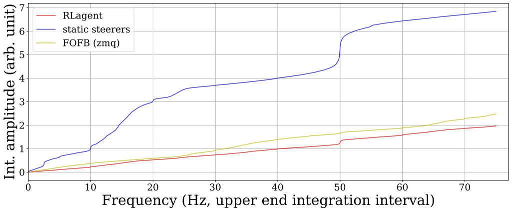

Table of Contents

Mitigation of Harmonic Orbit Perturbations

- Model-free approach
- Towards a model-based approach
- Interpretability

Anomaly detection with feature assignation

Summary and Outlook

References

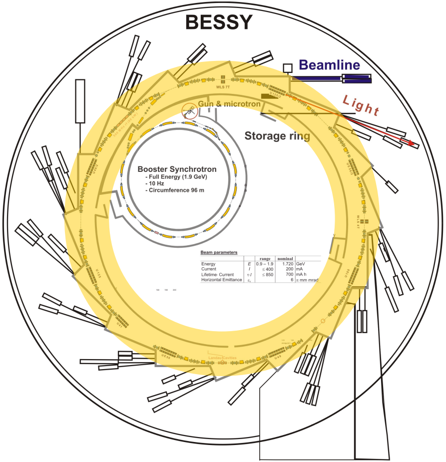

Towards a model-based approach

- Ideal set up: global response-matrix-based correction *smoothed* with ML/RL for the frequency domain.
- First approach: *learn* perturbation patterns with a **surrogate model** used for
  - combined correction with response-matrix methods
  - model-based RL
  - further analysis

$$F : \mathbb{R}^{n\times m} \times \mathbb{R}^{l} \longrightarrow \mathbb{R}^{m}$$
$$(\mathbf{x}_{t-n}, \ldots, \mathbf{x}_t, \mathbf{c}_{t+1}) \longmapsto \mathbf{x}_{t+1}$$

Input: BPMs *n = 10, m = 100* → **Dense** *10 × 2000* → **Dense** *10 × 1000* → **Dense** *10 × 500* → **Flatten & Concatenate** *5048* (with **Input**: steerers *l = 48*) → **Dense** *200* → **Output** *100*

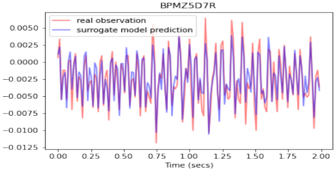

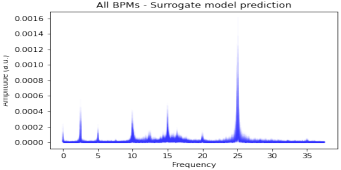

Error	Test Set Avg.	Previous BPM Data					Resp. Matrix: $\mathbf{x}_t + S\,(\mathbf{c}_{t+1} - \mathbf{c}_t)$	Model: $F(\mathbf{x}_{t-9}, \ldots, \mathbf{x}_t, \mathbf{c}_{t+1})$
		$\mathbf{x}_t$	$\mathbf{x}_{t-1}$	$\mathbf{x}_{t-2}$	$\mathbf{x}_{t-3}$	$\mathbf{x}_{t-4}$		
RMSE	0.0154	0.0274	0.0178	0.0236	0.0204	0.0229	0.0098	0.0036
R^2	0	-2.172	-0.3455	-1.3492	-0.769	-1.2223	0.5896	0.9441

→ *our surrogate model achieves an impressive prediction accuracy!*

Towards a model-based approach: Test for all BPMs, 21/6/21

$c^0_{t+1} := \mathbf{c}_t - \alpha S^{-1}\mathbf{x}_t$

repeat

$$\begin{aligned}
\mathbf{x}^k_{t+1} &:= \mathbf{x}_t + S(\mathbf{c}^k_{t+1} - \mathbf{c}_t) \\
\tilde{\mathbf{x}}^k_{t+1} &:= F(\mathbf{x}_{t-n}, \ldots, \mathbf{x}_t, \mathbf{c}^k_{t+1}) \\
\mathbf{c}^{k+1}_{t+1} &:= \mathbf{c}_t + \alpha S^{-1}\left[-(\tilde{\mathbf{x}}^k_{t+1} - \mathbf{x}^k_{t+1}) - \mathbf{x}_t\right]
\end{aligned}$$

until *convergence or k too large*

Time domain

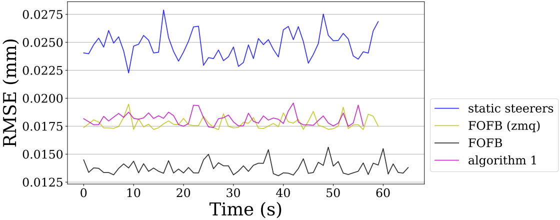

Frequency domain

Archiver beam motion variable

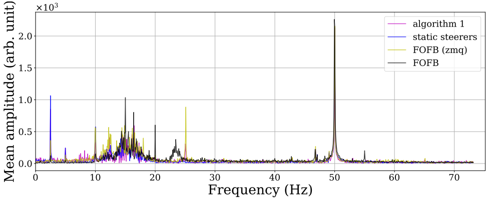

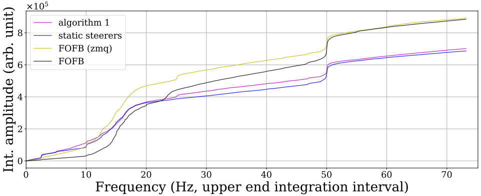

Table of Contents

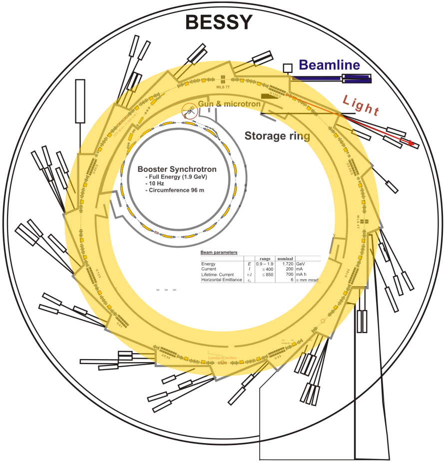

$$\mathbf{x}_{t+1} \approx F(\mathbf{x}_{t-n}, \ldots, \mathbf{x}_t, \mathbf{c}_{t+1})$$

- F is a NN with `ReLU` and linear activation functions $\rightarrow$ differentiable a.e.

$$J^c := \begin{pmatrix} \frac{\partial F^1}{\partial c^1_{t+1}} & \cdots & \frac{\partial F^m}{\partial c^1_{t+1}} \\ \vdots & \ddots & \vdots \\ \frac{\partial F^1}{\partial c^l_{t+1}} & \cdots & \frac{\partial F^m}{\partial c^l_{t+1}} \end{pmatrix}$$

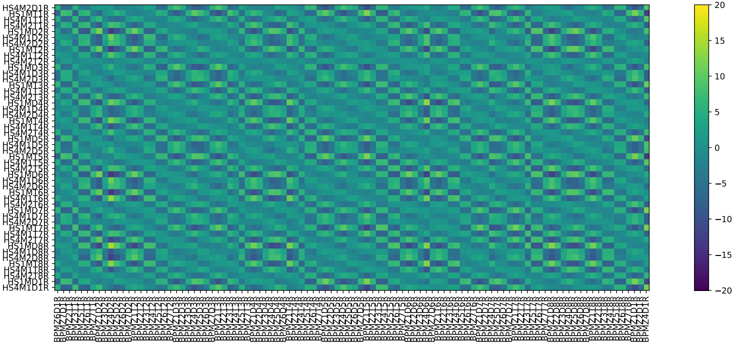

Avg. of J^c's (approx. $\mathbb{E}_{(X_t, \mathbf{c}_{t+1}) \sim \text{test data}} \left[J^c \,|_{(X_t, \mathbf{c}_{t+1})} \right]$)

HZB Helmholtz Zentrum Berlin

$$\mathbf{x}_{t+1} \approx F(\mathbf{x}_{t-n}, \ldots, \mathbf{x}_t, \mathbf{c}_{t+1})$$

- F is a NN with `ReLU` and linear activation functions $\rightarrow$ differentiable a.e.

$$J^c := \begin{pmatrix} \frac{\partial F^1}{\partial c^1_{t+1}} & \cdots & \frac{\partial F^m}{\partial c^1_{t+1}} \\ \vdots & \ddots & \vdots \\ \frac{\partial F^1}{\partial c^l_{t+1}} & \cdots & \frac{\partial F^m}{\partial c^l_{t+1}} \end{pmatrix}$$

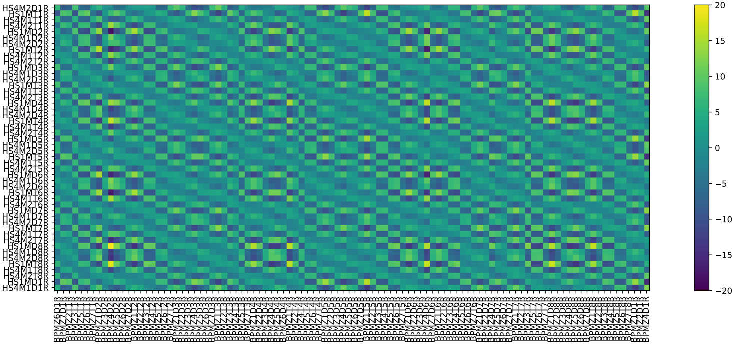

Actual measured (horizontal) response matrix S

$$\mathbf{x}_{t+1} \approx F(\mathbf{x}_{t-n}, \ldots, \mathbf{x}_t, \mathbf{c}_{t+1})$$

- F is a NN with `ReLU` and linear activation functions $\rightarrow$ differentiable a.e.

$$J^c := \begin{pmatrix} \frac{\partial F^1}{\partial c^1_{t+1}} & \cdots & \frac{\partial F^m}{\partial c^1_{t+1}} \\ \vdots & \ddots & \vdots \\ \frac{\partial F^1}{\partial c^l_{t+1}} & \cdots & \frac{\partial F^m}{\partial c^l_{t+1}} \end{pmatrix}$$

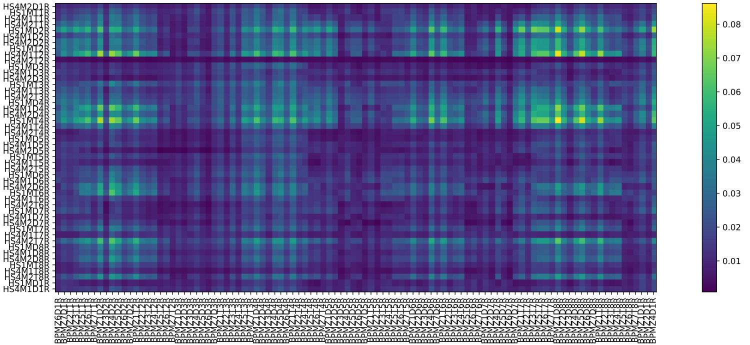

Standard deviation of J^c's

Table of Contents

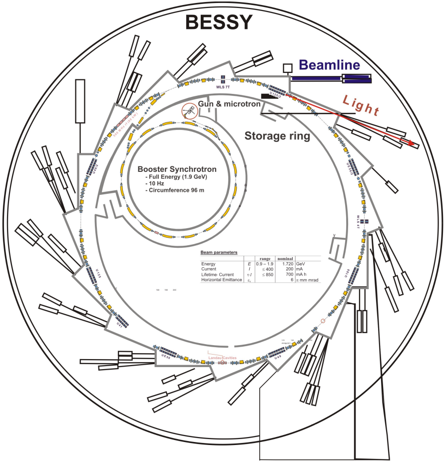

- First experiments with **Isolation Forests** ([LTZ08]) - prototype ready for roll-out in control room.
- We are developing additional **anomaly scores** that can be **individually assigned to input features** - based on [LWSG13], [CTS20], [LEC$^+$19].

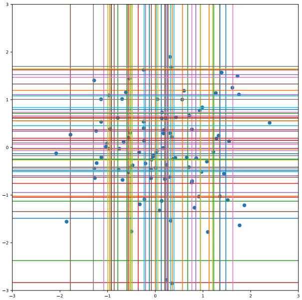

Let $\mathcal{P}_i(x) = \{N_0, N_1, \ldots, N_l\}$ be the node path of point x in i-th tree T_i and $s(N)$ the number of samples assigned to node N after training. Define the *Splitting path* of x in tree T_i $(p > 0)$ as:

$$\mathcal{S}_i(x) := \left\{ \left(1 - \frac{s(N_{j+1})}{s(N_j)} \right)^p \right\}_{N_j, N_{j+1} \in \mathcal{P}_i(x),\, j=1,\ldots,l}$$

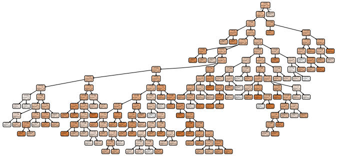

Anomaly score of point x and feature F at tree T_i:

$$a_i(F, x) := \sum_j \{ s_j \in \mathcal{S}_i(x) \mid F(N_j) = F \}$$

extended to the whole forest as: $a(F, x) := \dfrac{1}{k} \sum_{i=1}^{k} a_i(F, x)$

Anomaly detection with feature assignation

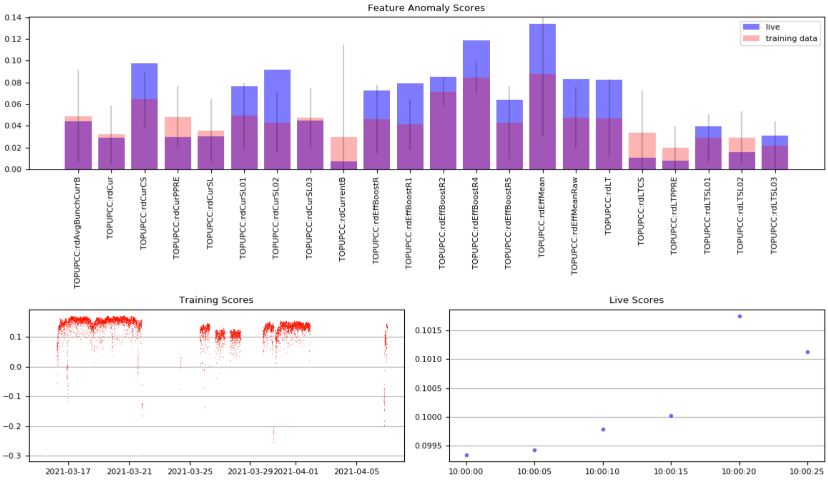

’Normal’ point between injections

High score, feature anomaly coefficients within the range w.r.t. training data.

Anomaly detection with feature assignation

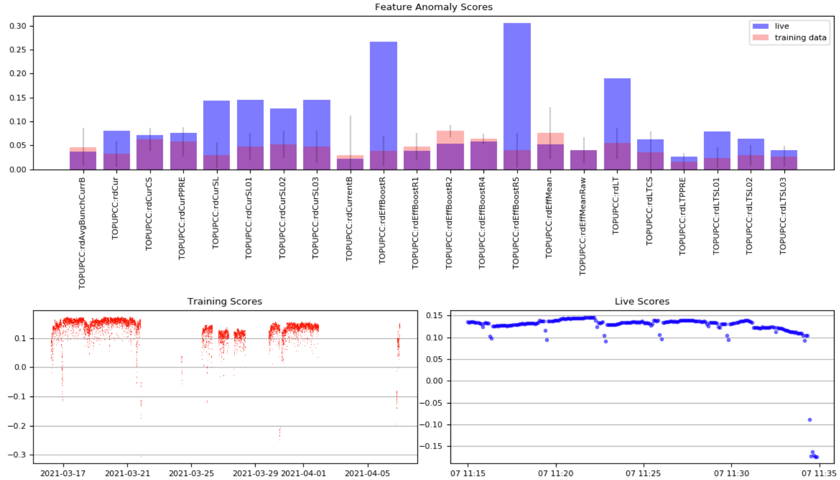

Anomalous point

Negative score; the machine was clearly in an anomalous state (decay mode - training data was only from top-up mode).

Anomaly detection with feature assignation

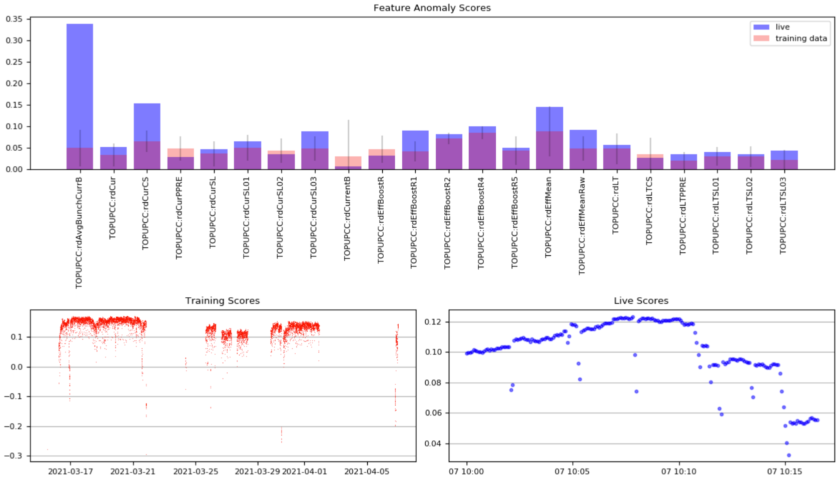

Almost-anomalous point

Lower score, since **per-bunch booster current** was also anomalous, but still positive.

Anomaly detection with feature assignation

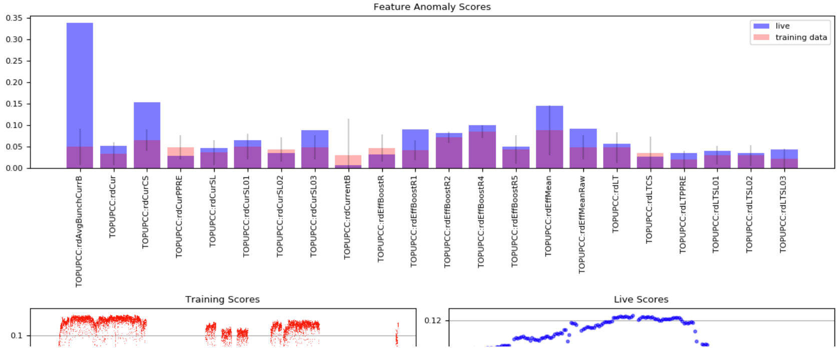

Almost-anomalous point

Lower score, since **per-bunch booster current** was also anomalous, but still positive.

Wed Apr 7 09:10:06 2021	(NO_ALARM / noackT)	TOPUPCC:swOffLTrg	Auto-Gun-Pulse OFF is disabled	NO_ALARM
Wed Apr 7 09:10:07 2021	(NO_ALARM / noackT)	TOPUPCC:swOffKSTrg	Auto-Booster-Injection OFF is disabled	NO_ALARM
Wed Apr 7 10:15:56 2021	(MAJOR / noackT)	TOPUPCC:rdAvgBunchCurrB	Bad per-Bunch Booster Current	MAJOR
Wed Apr 7 10:20:29 2021	(MINOR / noackT)	TOPUPCC:rdAvgBunchCurrB	Bad per-Bunch Booster Current	MINOR
Wed Apr 7 10:20:30 2021	(NO_ALARM / noackT)	TOPUPCC:rdAvgBunchCurrB	Bad per-Bunch Booster Current	NO_ALARM
Wed Apr 7 10:59:05 2021	(NO_ALARM / noackT)	TOPUPCC:rdAvgBunchCurrB	Bad per-Bunch Booster Current	NO_ALARM

*The model detected the feature anomaly in the per-bunch booster current **without any information about our alert intervals!***

HZB Helmholtz Zentrum Berlin

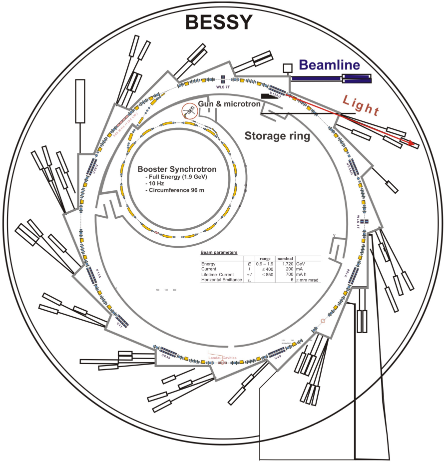

Summary and Outlook

- **Mitigation of harmonic orbit perturbations**
  - **Model-free** deep RL agent with 6 BPMs and steerers improves stability in frequency domain but not in time domain.
  - Steps towards **model-based** global optimisation: very accurate surrogate models, first combined algorithms tested.
- Anomaly detection with **feature assignation scores** developed and tested in several proof-of-concepts - ready for roll-out in control room.
- Next steps:
  - Model-based RL for harmonic orbit perturbations .
  - Interpretability: simple gradient analysis already gives interesting information - more sophisticated tools towards interpretability $\rightarrow$ Physical insights?

Table of Contents

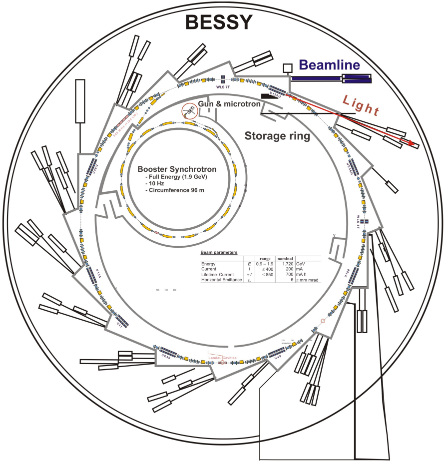

References I

HZB Helmholtz Zentrum Berlin

[AAB+15] Martin Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S. Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Ian Goodfellow, Andrew Harp, Geoffrey Irving, Michael Isard, Yangqing Jia, Rafal Jozefowicz, Lukasz Kaiser, Manjunath Kudlur, Josh Levenberg, Dandelion Mane, Rajat Monga, Sherry Moore, Derek Murray, Chris Olah, Mike Schuster, Jonathon Shlens, Benoit Steiner, Ilya Sutskever, Kunal Talwar, Paul Tucker, Vincent Vanhoucke, Vijay Vasudevan, Fernanda Viegas, Oriol Vinyals, Pete Warden, Martin Wattenberg, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. TensorFlow: Large-scale machine learning on heterogeneous systems, 2015. Software available from tensorflow.org.

[ACCR19] Daniel Allan, Thomas Caswell, Stuart Campbell, and Maksim Rakitin. Bluesky's ahead: A multi-facility collaboration for an a la carte software project for data acquisition and management. *Synchrotron Radiation News*, 32(3):19–22, 2019.

[AGTZ14] Ilya Agapov, Gianluca Geloni, Sergey Tomin, and Igor Zagorodnov. OCELOT: A software framework for synchrotron light source and FEL studies. *Nuclear instruments & methods in physics research / A*, 768:151 – 156, 2014. (c) Elsevier B.V.

[Cho15] François Chollet. keras. `https://github.com/fchollet/keras`, 2015.

References II

[Chu16] Olivier Churlaud. Localization and correction of orbit perturbations in BESSY II storage ring. Master's thesis, TU Berlin, HZB, 8 2016.

[CTS20] Mattia Carletti, Matteo Terzi, and Gian Antonio Susto. Interpretable anomaly detection with diffi: Depth-based feature importance for the isolation forest, 2020.

[LEC+19] Scott M. Lundberg, Gabriel G. Erion, Hugh Chen, Alex DeGrave, Jordan M. Prutkin, Bala Nair, Ronit Katz, Jonathan Himmelfarb, Nisha Bansal, and Su-In Lee. Explainable AI for trees: From local explanations to global understanding. *CoRR*, abs/1905.04610, 2019.

[LHP+16] Timothy P. Lillicrap, Jonathan J. Hunt, Alexander Pritzel, Nicolas Heess, Tom Erez, Yuval Tassa, David Silver, and Daan Wierstra. Continuous control with deep reinforcement learning. In *4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto Rico, May 2-4, 2016, Conference Track Proceedings*, 2016.

[LTZ08] Fei Tony Liu, Kai Ming Ting, and Zhi-Hua Zhou. Isolation forest. In *Proceedings of the 2008 Eighth IEEE International Conference on Data Mining*, ICDM '08, pages 413–422, Washington, DC, USA, 2008. IEEE Computer Society.

HZB Helmholtz Zentrum Berlin

[LWSG13] Gilles Louppe, Louis Wehenkel, Antonio Sutera, and Pierre Geurts. Understanding variable importances in forests of randomized trees. In *Proceedings of the 26th International Conference on Neural Information Processing Systems - Volume 1*, NIPS'13, pages 431–439, USA, 2013. Curran Associates Inc.

[MGL+09] R. Mueller, R. Goergen, R. Lange, I. Mueller, and J. Rahn. Introducing Fast Orbit Feedback at BESSY. In *Proc. 12th Int. Conf. on Accelerator and Large Experimental Physics Control Systems (ICALEPCS'09)*, pages 773–775. JACoW Publishing, 10 2009.

[Pla16] Matthias Plappert. keras-rl. `https://github.com/keras-rl/keras-rl`, 2016.

[SMS+21] William Smith, Roland Mueller, Pierre Schnizer, Simone Vadilonga, Sebastian Kazarski, and Luis Vera Ramirez. Status of Bluesky Deployment at BESSY II. In *this conference*, page FRBR03, 2021.

[VRHM+20] Luis Vera Ramirez, Gregor Hartmann, Tom Mertens, Roland Mueller, and Jens Viefhaus. Adding Machine Learning to the Analysis and Optimization Toolsets at the Light Source BESSY II. In *17th International Conference on Accelerator and Large Experimental Physics Control Systems*, page TUCPL01, 2020.

Appendix - Orbit Perturbations Chronology

- **Early 2020**: first RL tests with simulations (code from [Chu16] and OCELOT [AGTZ14]). A RL agent acting in time domain able to reduce perturbations also in frequency domain:

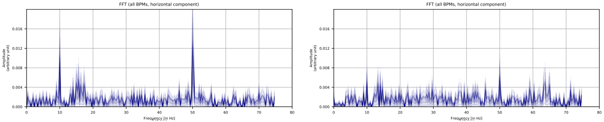

- **July 2020**: infrastructure set-up during machine commissioning and first *plausibility tests* of the Bluesky-based framework up to 20Hz.
- **September 2020**: direct zmq-communication with the mBox (fast orbit correction infrastructure), communication up to 100Hz. First learning results.
- **Since February 2021**: zmq-communication improved, stable RL-interaction loop at 75 Hz, further attemps at 150 Hz:

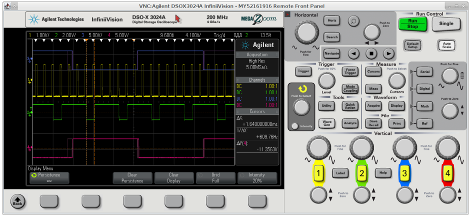

...and how is the RL agent manipulating the steerers?

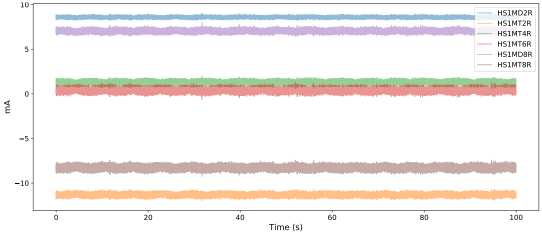

Notice that the smoothness of the RL agent performance in frequency domain is achieved in spite of (or probably due to) a high activity in the horizontal steerers...

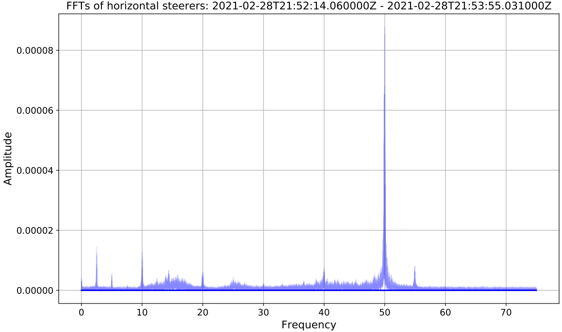

... which is a translation of the observed perturbation pattern of the orbit

- **Frameworks**: [AAB+15], [Cho15], [Pla16]

DDPG - Hyperparameters

- **Actor**: feedforward network with three hidden layers (50-20-10 neurons), `ReLU` as activation function (output with `tanh`)
- **Critic**: feedforward network with four hidden layers (50-(50+action)-20-10) neurons, `ReLU` as activation function, `Adam` as optimizer
- **Learning**: $\gamma = 0.99$, target model update rate $= 0.01$, batch size $= 32$
- **Exploration**: Ornstein-Uhlenbeck process ($\sigma = 0.05$, $\theta = 0.1$, no annealing), memory buffer with 20000 steps, short episodes (50 steps) followed by steerers random restart.

Surrogate model correction - Hyperparameters

- **Model**: architecture plotted in slide, `ReLU` as activation function, linear output, `Adam` as optimizer
- **Learning**: 16000 points for training, 4000 for test (gathered at 150 Hz), validation split 0.05, batch size $= 32$, 50 epochs.
- **Steerer randomisation**: steerer intensities sampled from $\mathcal{N}(0, 5)$ (in mA) and actually set with probability $p = 0.005$.