

A Two-Stage Machine Learning-Aided Approach for Quench Identification at the European XFEL

L. Boukela* A. Eichler*,** J. Branlard* N.Z. Jomhari*

* *Deutsches Elektronen-Synchrotron DESY, 22607 Hamburg, Germany*
(e-mail: lynda.boukela@desy.de, annika.eichler@desy.de,
julien.branlard@desy.de, nur.zulaiha.jomhari@desy.de).

** *Hamburg University of Technology, 21073 Hamburg, Germany*
(e-mail: annika.eichler@tuhh.de).

Abstract: This paper introduces a machine learning-aided fault detection and isolation method applied to the case study of quench identification at the European X-Ray Free-Electron Laser. The plant utilizes 800 superconducting radio-frequency cavities in order to accelerate electron bunches to high energies of up to 17.5 GeV. Various faulty events can disrupt the nominal functioning of the accelerator, including quenches that can lead to a loss of the superconductivity of the cavities and the interruption of their operation. In this context, our solution consists in analyzing signals reflecting the dynamics of the cavities in a two-stage approach. (I) Fault detection that uses analytical redundancy to process the data and generate a residual. The evaluation of the residual through the generalized likelihood ratio allows detecting the faulty behaviors. (II) Fault isolation which involves the distinction of the quenches from the other faults. To this end, we proceed with a data-driven model of the k-medoids algorithm that explores different similarity measures, namely, the Euclidean and the dynamic time warping. Finally, we evaluate the new method and compare it to the currently deployed quench detection system, the results show the improved performance achieved by our method.

Copyright © 2024 The Authors. This is an open access article under the CC BY-NC-ND license (<https://creativecommons.org/licenses/by-nc-nd/4.0/>)

Keywords: Fault detection, Fault isolation, Machine learning, Particle accelerators, XFEL.

1. INTRODUCTION

Fault detection plays a crucial role in ensuring the safe and optimal operation of complex systems. Traditional model-based approaches, although widely explored, have some limitations in their ability to isolate all possible faults, especially those with evolving patterns as complete models are usually hard to obtain. Machine learning (ML) approaches, on the other hand, are simpler and handle more effectively complex systems, however, their performance relies heavily on the data quantity and quality. Hybrid solutions seem to be promising although the choice of methods to integrate is challenging (Wilhelm et al., 2021). A hybrid method is thus developed for our case study, i.e., the problem of quench detection at the European X-Ray Free-Electron Laser (EuXFEL), where a fault detection method based on physical model residuals and data-driven clustering for fault isolation are combined.

The EuXFEL is the largest particle accelerator for X-ray laser generation worldwide. Along with the high-voltage power supplies, the klystrons and the waveguides, the superconducting radio-frequency cavities (SRFCs), are key components of the power transfer chain to the electron

beam. The plant's linac comprises hundreds of 1.3 GHz TESLA-type cavities, controlled and monitored with the low level radio frequency (LLRF) system and operated in a *pulsed* mode with a pulse repetition rate of 10 Hz and at an average gradient of 23.6 MV/m (Reschke et al., 2017). This enables the three self-amplified spontaneous emission undulators to generate up to 27 000 photon flashes every second. This setup is exploited by several hundred interdisciplinary users every year to carry out their experiments.

Enhancing the availability, safety and reliability of the accelerator is therefore crucial to meet the demand, minimize potential energy and financial losses and avoid the degradation of the facility. In this context, the behavior of the SRFCs is monitored through the LLRF system to detect and report the faulty events, especially quenches. These are severe faults that cause a loss of the superconductivity of the SRFCs and thus an operation downtime. If a quench goes undetected, it has the potential to generate sufficient heat that can lead to disturbances in the helium flow of the cryogenic system and to an operation interruption that could last up to 24 hours.

A quench detection system (QDS) has been deployed at the EuXFEL since its first commissioning in 2017 (Branlard et al., 2013). It relies on a statistical analysis of the pulses for indication of changes of the cavity quality factor Q_L , which is an indicating measure of the field coupling and the dissipating power in the SRFCs. The

* The authors acknowledge support from DESY (Hamburg, Germany), a member of the Helmholtz Association HGF. This work was funded in the context of the R&D program of the European XFEL. The authors would like to thank Jan Horst Karl Timm, Christian Schmidt and Nicholas Walker for their help.

QDS is, however, not robust enough, as disruptive events such as controlled detuning, field emitters, and digital glitches can affect the Q_L computation, and lead the system to trigger alarms as a result of *fake* quenches (Eichler et al., 2023). Because of the considerable number of these false alarms, and since the physical model of the SRFCs and measurements from the LLRF system are available to build a more robust quench detection method, alternative solutions have been explored and developed (Nawaz et al., 2016, 2018; Syed et al., 2021; Bellandi et al., 2021; Sulc et al., 2022; Martino et al., 2022; Branlard et al., 2022; Eichler et al., 2023). Purely data-driven techniques have also been investigated in other facilities (Tennant et al., 2020; Vidyaratne et al., 2022).

In our solution, the well-established model of the nominal behavior of the SRFCs from (Schilcher, 1998) is exploited in order to obtain a residual. To this purpose, the parity space-based approach has been retained for fault detection (Eichler et al., 2023), followed by the main contribution of this paper, the ML-assisted quench isolation. The main motive for the usage of the parity space is the robustness of the method against drifts and noise, but also due to the need of a firmware implementation. A statistical test of the residual, namely the generalized likelihood ratio (GLR), that presents a distribution that alleviates the process of threshold determination, is used to detect all occurring faults. The test outputs fault-specific patterns, and to further distinguish the quenches from other types of faults, the k-medoids clustering algorithm is employed to learn a quench model. Two clustering models are obtained with two different similarity measures, namely, the lockstep Euclidean (EUC) and the elastic dynamic time warping (DTW). The inference is achieved through thresholds and decision boundaries on the distances to the quench medoids. This approach presents advantages in terms of suitability for an online processing setup and with respect to interpretability which is an important aspect for the LLRF experts.

Section 2 presents the model-based residual generation and evaluation for fault detection. Section 3 details the quench identification through the data-driven fault isolation. The experimental evaluation setup and results are discussed in Section 4. Finally, Section 5 concludes the paper.

2. RESIDUAL-BASED FAULT DETECTION

Model-based fault detection relies on the assessment of consistency between available measurements and prior knowledge. This is achieved through residuals that can be obtained with different methods. In our case, the nonlinear parity space method is used, and it is detailed here.

2.1 The physical model of the SRFCs

Because of its structure that is able to exhibit a resonant behavior at specific frequencies, a superconducting cavity is considered as a resonator. The parameters influencing the cavity's dynamics are the detuning $\Delta\omega(t)$ and the half bandwidth $\omega_{1/2}$ which represent, respectively, the difference between the driving frequency and the resonance frequency, and how sensitive a SRFC is towards the

detuning. The cavity electromagnetic model as elaborated in (Schilcher, 1998) is given by

$$\begin{bmatrix} \dot{V}_{P,I}(t) \\ \dot{V}_{P,Q}(t) \end{bmatrix} = \begin{bmatrix} -\omega_{1/2} & -\Delta\omega(t) \\ \Delta\omega(t) & -\omega_{1/2} \end{bmatrix} \begin{bmatrix} V_{P,I}(t) \\ V_{P,Q}(t) \end{bmatrix} + 2\omega_{1/2} \begin{bmatrix} V_{F,I}(t) \\ V_{F,Q}(t) \end{bmatrix} - \omega_{1/2} \begin{bmatrix} V_{B,I}(t) \\ V_{B,Q}(t) \end{bmatrix}, \quad (1)$$

where I and Q are the in-phase and quadrature components of the different signals. $V_F(t)$ is the forward field coupled into the cavity, $V_P(t)$ is the probe signal, i.e., the field generated inside the cavity, and $V_B(t)$ is the field induced by the beam.

2.2 Residual generation

The nonlinear parity space is an analytical redundancy method that aims at generating redundant expressions using the state-space model. Successive derivations of the model must be performed in order to retain only the known variables and obtain residuals. The residuals are then used as indicators of faults presence, a residual is non-zero in presence of faults, it is equal to zero otherwise. For the problem at hand, the cavity model has been exploited to derive the residual

$$r(t) = \frac{-\dot{V}_{P,I}(t) + \omega_{1/2} [-V_{P,I}(t) + 2V_{F,I}(t) - V_{B,I}(t)]}{V_{P,Q}(t)} - \frac{\dot{V}_{P,Q}(t) + \omega_{1/2} [V_{P,Q}(t) - 2V_{F,Q}(t) + V_{B,Q}(t)]}{V_{P,I}(t)},$$

that is based on the analytical redundancy of (1) with respect to the detuning. More details regarding the choice of the residual can be found in (Nawaz et al., 2018).

2.3 Residual assessment

In order to automate the residual analysis for presence of faults, the log-likelihood ratio is used with the assumption that the residual follows a Gaussian distribution $\mathcal{N}(\mu, \sigma)$ with a mean of zero under nominal conditions, i.e., $\mu = 0$, representing the null hypothesis. While with fault occurrence, the residual is expected to exhibit a Gaussian distribution with a mean different from zero, i.e., $\mu \neq 0$, representing the non-null hypothesis. Unlike the variance, the mean is unknown and needs to be estimated, the generalized likelihood ratio is therefore applied,

$$\lambda_{\text{GLR}}(k) = \frac{K}{2} \left(\frac{1}{K} \sum_{i=k-K+1}^k r(i)^\top \right) \sigma^{-1} \left(\frac{1}{K} \sum_{i=k-K+1}^k r(i) \right),$$

where $r(i)$ is the evaluated and discretized residual, K is the size of the moving evaluation window and σ is the variance of the nominal residual. The result follows a χ^2 distribution, which gives, with a desired false alarm rate, directly the corresponding threshold on rating events as anomalous (Eichler et al., 2023).

3. CLUSTERING-BASED QUENCH ISOLATION

The statistical test helps to detect faults, it is however unable to identify their types. Data-driven methods help to address this issue. The goal in our case is to isolate




Fig. 1. Illustrative example of the DTW and EUC similarity measures, and their difference in terms of sample correspondence. Here, DTW is equal to 2.8 and EUC is equal to 18.6.

the quenches from the rest of faulty data. The k -medoids algorithm (Kaufman and Rousseeuw, 1990) is thus used to obtain quench clustering models from the statistical results. Furthermore, decision boundaries are used in addition to thresholds in order to ensure an optimal separation of the different faults during the inference.

3.1 Quench clustering models

The k -medoid method is briefly summarized in the following. Given $X = \{x_1, x_2, \dots, x_n\}$, with $x_i \in \mathbb{R}^d$, a dataset of n observations with dimension d . The k -medoids is a clustering algorithm that aims to partition X into k clusters $C = \{c_1, c_2, \dots, c_k\}$ by iteratively assigning every observation to a cluster medoid $M = \{m_1, m_2, \dots, m_k\}$, here, a medoid is a representative observation, i.e., $M \subset X$. In the same cluster, observations show high degree of similarity $S(x_1, x_2)$, while they are as dissimilar as possible in different clusters. The clustering can be achieved with different partitioning algorithms. The main input parameters are the number of clusters k , the similarity measure S and the maximum number of iterations.

Two models are built in our case, based on two different similarity measures. As base measure, the EUC is used,

$$\text{EUC}(x_1, x_2) = \sqrt{\sum_{i=1}^n (x_1^i - x_2^i)^2}.$$

Given the GLR traces of normalized quench pulses that exhibit variability, characterized primarily by stretches and shifts in time, as shown in Fig. 1, the DTW (Sakoe and Chiba, 1978)

$$\text{DTW}(x_1, x_2) = \arg \min_{i,j} \sum \text{dist}(x_1^i, x_2^j),$$

is also explored, where $i, j \in \{1, 2, \dots, d\}$ are sample indices of the observations x_1 and x_2 . Any distance (dist) can be employed for DTW, here, the Euclidean is used.

The difference between the two, as can be seen from the previous equations and from the illustrative example in Fig. 1, consists mainly in what the measures are trying to capture. While the EUC computes the similarity as the sum of the distances between each sample from observation x_1 and its time-corresponding sample in observation x_2 , DTW is the sum of distances from the optimal time-warping path, i.e., the one assigning samples of observation x_1 to the closest samples in observation x_2 . DTW tries therefore to capture the shape similarity in addition to the distance. And the quench traces in Fig. 1 are thus




Fig. 2. Traces and medoids of the two clusters obtained with the two distances, EUC and DTW.

more similar with respect to DTW that is equal to 2.8 compared to EUC found to be equal to 18.6. Data from 2021, consisting of $n = 76$ quench traces of dimension $d = 1819$, have been used to build the models. The data are first pre-processed, a maximum-based normalization is adopted in order to equalize the scale of the data without changing its inherent structure that is important for DTW. The scaling is further accompanied by a reduction, in order to retain the frame of interest from the traces. It is obtained by keeping 300 samples following the first sample greater than 0.2. This applies only to the Euclidean distance, as it is unable to capture the similarity of traces with time misalignment.

The models are then built using k -medoids. The Partitioning Around Medoids (Kaufman and Rousseeuw, 1990) algorithm is utilized. Two main patterns of the quench traces have been noticed, two clusters are therefore built, i.e., $k = 2$, and two medoids $M = \{m_1, m_2\}$ are obtained with 100 iterations. The traces and medoids of the clusters obtained with the two distances are shown in Fig. 2. With both measures, the second cluster gathers the quenches appearing at an early stage, (at the beginning of the pulse), while those in the first cluster are quenches that have occurred at a later stage (at the end of the pulse).

3.2 Inference

Given a new faulty trace $g \in \mathbb{R}^d$, in order to make decision whether it is a quench or not, its distance to the quench medoids is computed to assess its similarity with the quench clustering model. For both similarity measures, we define a set, based on the distances to the medoids, within which a trace is labeled as a quench. To estimate the sets, we set lower and/or upper thresholds on the distance to the quench medoids, in addition to decision boundaries \mathcal{H} . To estimate \mathcal{H} , a validation dataset N , with a total of 407 non-quench faulty traces, is used.

EUC-based inference: Fig. 3 depicts the distance space (distance from the medoids), obtained with EUC, of the two quench clusters from X and the validation data N . As can be seen, the two quench clusters are not distinct and present a semi-elliptical shape. In order to separate the quenches from the other faults, and given the shape of the quench distances, in addition to the thresholds, two




Fig. 3. EUC decision boundaries (\mathcal{H}_{EUC}^i and \mathcal{H}_{EUC}^o) and thresholds. The green delimits the area of the quench distances.

ellipses are used as decision boundaries, an inner boundary \mathcal{H}_{EUC}^i and an outer boundary \mathcal{H}_{EUC}^o . The two ellipses are obtained as a scaling of the main ellipse \mathcal{H}_{EUC} fitted to the quench distances through the Least Squares method, and satisfying the following equation

$$\mathcal{H}_{EUC}(s_1, s_2) = \frac{((s_1 - s_1^0) \cos(\phi) - (s_2 - s_2^0) \sin(\phi))^2}{a^2} + \frac{((s_1 - s_1^0) \sin(\phi) + (s_2 - s_2^0) \cos(\phi))^2}{b^2} = r^2,$$

where s_1 is the distance to the first cluster, s_2 is the distance to the second cluster, ϕ is the rotation angle, s_1^0 is the center at the horizontal axis, s_2^0 is the center at the vertical axis and r is the scaling factor equal to 1.

A set of quench traces based on EUC can be defined as

$$\mathcal{Q}_{EUC} = \{q \mid q \in \mathbb{R}^d, \text{EUC}(q, m_j) < \max(\text{EUC}(x_i, m_j)) + \epsilon, (r^i)^2 < \mathcal{H}_{EUC}(\text{EUC}(q, m_1), \text{EUC}(q, m_2)) < (r^o)^2, \forall x_i \in X, \forall m_j \in M\},$$

where ϵ is a small tolerance value learned empirically, r^i and r^o are the radii of \mathcal{H}_{EUC}^i and \mathcal{H}_{EUC}^o , respectively, obtained as the closest and furthest distances, $\mathcal{H}_{EUC}(\text{EUC}(x_i, m_1), \text{EUC}(x_i, m_2))$ with $x_i \in X$, minus or plus ϵ . The detection for a new trace g is therefore achieved as in the following

$$\text{Inference}_{EUC}(g) = \begin{cases} \text{Quench}, & \text{if } g \in \mathcal{Q}_{EUC} \\ \text{Other fault}, & \text{otherwise.} \end{cases}$$

DTW-based inference: Fig. 4 depicts the distance space, obtained with DTW, of both the quench clusters from X and the validation data N . Here, we can see a clear separation of the two trace types. While the quench cluster in the top left of the plot is well isolated from the non-quench faults, some overlapping can be noticed between the other faults and the second quench cluster, a decision boundary is therefore fitted as follows

$$\mathcal{H}_{DTW}(s_1) = as_1^3 + bs_1^2 + cs_1 + f,$$

with a , b , c and f being the polynomial coefficients. The decision boundary is obtained by adapting the intercept of




Fig. 4. DTW decision boundary and thresholds. The green highlights the area of the quench distances.

the cubic fit of the distances in the cluster in the bottom of the plot.

Let \mathcal{Q}_{DTW} be the set of quench traces based on DTW

$$\mathcal{Q}_{DTW} = \{q \mid q \in \mathbb{R}^d, \text{DTW}(q, m_j) < \max(\text{DTW}(x_i, m_j)) + \epsilon, \text{DTW}(q, m_j) > \min(\text{DTW}(x_i, m_j)) - \epsilon, \text{DTW}(q, m_2) < \mathcal{H}_{DTW}(\text{DTW}(q, m_1)), \forall x_i \in X, \forall m_j \in M\},$$

where ϵ is a small tolerance value. The inference for a new trace g is therefore achieved as follows

$$\text{Inference}_{DTW}(g) = \begin{cases} \text{Quench}, & \text{if } g \in \mathcal{Q}_{DTW} \\ \text{Other fault}, & \text{otherwise.} \end{cases}$$

4. EVALUATION

4.1 Setup

The EuXFEL linac is organized into 25 stations, each comprising 4 cryomodules. Within each cryomodule, there are 8 cavities, resulting in a total of 32 SRFCs per station. The data from the SRFCs consist of the forward and probe signals that correspond to the radio frequency pulses, where each pulse lasts for 1.82ms and is sampled at 1 MHz (i.e., 1819 samples). The data are captured by a snapshot recorder and saved in files in the hdf5 format. The snapshots are triggered by different protection systems, such as, klystrons, modulators, couplers, quadrupoles and cryogenics, but also the quench detection system and/or the finite state machine. Each file, identified by a unique timestamp, corresponds to a specific faulty event from an individual station, and gathers the sampled signals for 250 consecutive pulses in all the cavities of the station. Results of the residual-based detection, i.e., the GLR traces, are also saved to the same file. The ground truth for the specific faults, however, is not available. For this study, data from the first half and second half of 2022, corresponding to a total of 671 station-related snapshots, have been collected and annotated in order to evaluate the performance of the proposed approach.

We compare the proposed solution to the currently deployed QDS that relies on a statistical analysis of the quality factor Q_L of the cavities. The Q_L is computed for almost every pulse and compared to a running average

from the previous 100 pulses, and a sudden drop of the Q_L is an indicator of a quench. Results of the QDS are event-wise, i.e., the QDS pulse-by-pulse predictions are not available. In our case, an event is considered as a quench if at least one pulse in the station is identified as a quench. This is applicable to the problem at hand, as the identification of a quench results in the interruption of the radio frequency driving the affected station.

In order to illustrate the performance comparison in terms of quench identification, the area under the curve (AUC) of the receiver operating characteristics (ROC) curves is used. The *ROC-AUC* depicts the true positive rate

$$TPR = \frac{TP}{(FN + TP)},$$

as a function of the false positive rate

$$FPR = \frac{FP}{(FP + TN)},$$

where TP are the true positives (i.e. the method accurately identified a quench), TN are the true negatives (i.e. the method accurately recognized a fault was not a quench), FP are false positives (i.e. the method identified a quench that's in reality another fault), and FN are the false negatives (i.e., the algorithm failed to identify a real quench).

4.2 Implementation

The online and offline residual generation and evaluation are integrated to the trip event logger (Timm et al., 2021), an efficient tool implemented in C++ for automatic fault handling and prevention. The machine learning clustering models, along with an annotation tool, being at a testing phase so far, have been implemented in Matlab. The clustering models are deployed offline, and run daily to analyse the statistical results on the hdf5 snapshots.

4.3 Results

Results of the residual-based fault detection have shown that given the total number of 671 events, 354 were detected as faults within the SRFCs, some of them are documented in (Eichler et al., 2023). After annotation, 87 events were found corresponding to quenches. The highest number was recorded in April with a total of 21 quenches. This is because the facility was running at high energies, i.e., most cavities were operated just below the quench gradient. More detailed evaluation of the residual-based fault detection, and examples of the events that triggered alarms with the GLR, can be found in previous publications (Nawaz et al., 2018; Eichler et al., 2023).

It is worth mentioning that the current residual implementation does not include the voltage induced by the beam V_B , as additional input, as shown in (1), but its effect is compensated as presented in (Eichler et al., 2023). This becomes problematic in certain rare cases where inaccurate inferences may be triggered due to transitions of the beam turning on and off or changes in its pattern occurring within the 250 pulses captured in the snapshots. First tests where the bunch charge measurements from the toroids are added as an additional input have therefore been carried out, this will be automated and evaluated in the future.




Fig. 5. Comparison of the quench identification results obtained with the different methods.

The evaluation is based on the traces of the 354 faulty events. Results, obtained with the different methods for quench isolation, namely, EUC, DTW and the currently deployed QDS, are presented in Fig.5. The plot depicts the obtained number of TP , FN , FP , TN and the *ROC-AUC*. The two clustering-based methods, with a *ROC-AUC* equal to 0.94 for both EUC and DTW, outperform clearly the QDS having a *ROC-AUC* equal to 0.86 ($TPR=0.93$ and $FPR=0.20$). The latter generates a considerable number of fake quenches, 54 in this case, as a result of some behaviors, such as digitization failures, controlled detuning and field emitters, that lead to a drop in the quality factor Q_L , and thus a false alarm from the QDS. For instance, the fifth station has triggered many false quenches due to a software bug, caused by a run away situation resulting from conflicting feedback algorithms.




Fig. 6. False negative and false positive examples. fn1 is a quench trace missed by both EUC and DTW, while fn2 and fn3 were only missed by DTW. fp1 and fp2 were triggered by DTW, while fp3 was triggered by EUC.

Results of EUC and DTW are relatively similar, the EUC ($TPR=0.95$ and $FPR=0.07$) is however slightly better in terms of real quench detection and the DTW ($TPR=0.93$ and $FPR=0.04$) is slightly better in terms of false positives. Examples of false negatives and false positives can be seen in Fig.6. The false negatives are mostly caused by quenches exhibiting patterns relatively different from those used to learn the clustering models. The examples fn1 and fn2 are, for instance, caused by the beam signal, i.e., the quenches occurred simultaneously with other events, this will however be resolved upon the inclusion of the toroid. Example fn3, however, represents a quench captured at a time in between the two clusters learned with DTW. Similarly, the false positives are faults

having a GLR trace with a pattern close to those of the quenches used for the model learning. The false positives fp1, fp2 and fp3 were, for instance, caused by problems with other systems than the SRFs.

Unlike the complexity of DTW that is quadratic in the pulse length, the complexity of EUC is linear in the length of the frame of interest from the pulse. It's important to note that both approaches can operate on inter-pulse intervals, but using the Euclidean distance allows for intra-pulse analysis too. Both intra-pulse and inter-pulse analyses are needed and both methods perform similarly in terms of interpretability, as both make cluster assignments mostly based on the time of the quench occurrence. The EUC therefore appears to be more advantageous for an online setup, despite the possibility of deploying both methods as part of a more collaborative strategy.

5. CONCLUSION

In this paper, a hybrid method to detect quenches at the EuXFEL is presented. It is a two-stage method in which the first stage generates and evaluates residuals with the nonlinear parity space method and the generalized likelihood ratio, respectively, in order to detect faults. The quenches are subsequently distinguished from the other faults in the second stage through two different clustering models, based on the EUC and DTW similarity measures. The proposed solution has been implemented and evaluated on recent data, the obtained results depict its high performance compared to the currently deployed detection system. In the near future, the beam information will be included to the implementation of the residual generation, this will help to eliminate some false alarms. The current deployment of the proposed solution is however offline, an implementation on dedicated server and on a FPGA for an online deployment has therefore been initiated. We also aim to develop and compare other detection methods and implement a human-in-the-loop procedure, in which, all methods will be deployed, and in case of disagreement, an expert intervention will be initiated. Subsequently, with the new information, there will be an automatic retraining or adjustment of the inference strategy, in order to enhance both the detection and the identification.

REFERENCES

- Bellandi, A., Butkowski, L., Dursun, B., Eichler, A., Gummus, C., Kuntzsch, M., Nawaz, A., Pfeiffer, S., Schlarb, H., Schmidt, C., Zenker, K., and Branlard, J. (2021). Online detuning computation and quench detection for superconducting resonators. *IEEE Trans. Nucl. Sci.*, 68(4), 385–393.
- Branlard, J., Ayvazyan, V., Hensler, O., Schmidt, C., and Schlarb, H. (2013). Superconducting cavity quench detection and prevention for the European XFEL. In *14th International Conference on Accelerator & Large Experimental Physics Control Systems*, 1239–1241.
- Branlard, J., Eichler, A., Timm, J.H.K., and Walker, N. (2022). Machine learning assisted cavity quench identification at the European XFEL. In *31st International Linear Accelerator Conference*, 798–801.
- Eichler, A., Branlard, J., and Timm, J. (2023). Anomaly detection at the European X-ray free electron laser using a parity-space-based method. *Phys. Rev. Accel. Beams*, 26(1), 012801.
- Kaufman, L. and Rousseeuw, P.J. (1990). *Finding groups in data: an introduction to cluster analysis*. John Wiley & Sons, Inc, New York, NY, USA.
- Martino, G., Bellandi, A., Eichler, A., Branlard, J., Schlarb, H., Doolittle, L., Aderhold, S., Hoobler, S., Nelson, J., Porter, R.D., Zacarias, L., Benwell, A., Gonnella, D., Ratti, A., and Fey, G. (2022). Anomaly detection based quench detection system for CW operation of SRF cavities. In *31st International Linear Accelerator Conference*, 775–777.
- Nawaz, A., Pfeiffer, S., Lichtenberg, G., and Rostalski, P. (2018). Anomaly detection for the European XFEL using a nonlinear parity space method. *IFAC-PapersOnLine*, 51(24), 1379–1386. 10th IFAC Symposium on Fault Detection, Supervision and Safety for Technical Processes.
- Nawaz, A.S., Pfeiffer, S., Lichtenberg, G., and Schlarb, H. (2016). Self-organized critical control for the European XFEL using black box parameter identification for the quench detection system. In *3rd Conference on Control and Fault-Tolerant Systems*, 196–201.
- Reschke, D., Decking, W., Walker, N., and Weise, H. (2017). The commissioning of the European-XFEL linac and its performance. In *18th International Conference on RF Superconductivity*, 1–5.
- Sakoe, H. and Chiba, S. (1978). Dynamic programming algorithm optimization for spoken word recognition. *IEEE Trans. Acoust. Speech Signal Process*, 26(1), 43–49.
- Schilcher, T. (1998). *Vector sum control of pulsed accelerating fields in Lorentz force detuned superconducting cavities*. Ph.D. thesis, University of Hamburg.
- Sulc, A., Eichler, A., and Wilsken, T. (2022). A data-driven anomaly detection on SRF cavities at the European XFEL. *J. Phys. Conf. Ser.*, 2420(1), 012070.
- Syed, W.H., Eichler, A., Nawaz, A., Sharan, B., and Werner, H. (2021). Koopman-based Kalman filter for fault detection for the superconducting radio frequency cavities of the European XFEL. In *60th IEEE Conference on Decision and Control*, 6855–6860.
- Tennant, C., Carpenter, A., Powers, T., Solopova, A.S., Vidyaratne, L., and Iftikharuddin, K. (2020). Superconducting radio-frequency cavity fault classification using machine learning at Jefferson laboratory. *Phys. Rev. Accel. Beams*, 23(11), 114601.
- Timm, J., Branlard, J., Eichler, A., Schlarb, H., et al. (2021). The trip event logger for online fault diagnosis at the European XFEL. In *12th International Particle Accelerator Conference*, 3344–3346.
- Vidyaratne, L., Carpenter, A., Powers, T., Tennant, C., Iftikharuddin, K.M., Rahman, M.M., and Shabalina, A.S. (2022). Deep learning based superconducting radio-frequency cavity fault classification at Jefferson laboratory. *Front. Artif. Intell.*, 4, 718950.
- Wilhelm, Y., Reimann, P., Gauchel, W., and Mitschang, B. (2021). Overview on hybrid approaches to fault detection and diagnosis: Combining data-driven, physics-based and knowledge-based models. *Procedia Cirp*, 99, 278–283. 14th CIRP Conference on Intelligent Computation in Manufacturing Engineering.