

Classification of superconducting radio-frequency cavity faults of CAFE2 using machine learning

Authors: Yang, Li-juan, Peng, Jia-Yi, Qiu, Feng, He, Yuan, Ma, Jin-Ying, Xue, Zong-Heng, Jiang, Tian-Cai, Zhu, Zheng-Long, Chen, Qi, Xu, Cheng-Ye, Yu, Jing-Wei, Ma, Zhen, Luo, Di-Di, Yang, Zi-Qin, Gao, Zheng, Lie-Peng Sun, Zhang, Zhou-Li, Gui-Rong Huang, Wang, Zhi-Jun, Qiu, Feng, He, Yuan

Date: 2025-02-03T00:00:00+00:00

Abstract

Superconducting radio-frequency (SRF) cavities are the core components of SRF linear accelerators, making their stable operation considerably important. However, the operational experience from different accelerator laboratories has revealed that SRF faults are the leading cause of short machine downtime trips. When a cavity fault occurs, system experts analyze the time-series data recorded by low-level RF systems and identify the fault type. However, this requires expertise and intuition, posing a major challenge for control-room operators. Here, we propose an expert feature-based machine learning model for automating SRF cavity fault recognition. The main challenge in converting the “expert reasoning” process for SRF faults into a “model inference” process lies in feature extraction, which is attributed to the associated multidimensional and complex time-series waveforms. Existing autoregression-based feature-extraction methods require the signal to be stable and autocorrelated, resulting in difficulty in capturing the abrupt features that exist in several SRF failure patterns. To address these issues, we introduce expertise into the classification model through reasonable feature engineering. We demonstrate the feasibility of this method using the SRF cavity of the China Accelerator Facility for superheavy Elements (CAFE2). Although specific faults in SRF cavities may vary across different accelerators, similarities exist in the RF signals. Therefore, this study provides valuable guidance for fault analysis of the entire SRF community.

Full Text

Preamble

Classification of superconducting radio-frequency cavity faults of CAFE2 using machine learning

Li-Juan Yang,¹ Jia-Yi Peng,^{1,2} Feng Qiu,^{1,2,†} Yuan He,^{1,2,‡} Jin-Ying Ma,¹ Zong-Heng Xue,¹ Tian-Cai Jiang,¹ Zheng-Long Zhu,¹ Qi Chen,¹ Cheng-Ye Xu,¹ Jing-Wei Yu,¹ Zhen Ma,¹ Di-Di Luo,¹ Zi-Qin Yang,¹ Zheng Gao,¹ Lie-Peng Sun,¹ Zhou-Li Zhang,¹ Gui-Rong Huang,¹ and Zhi-Jun Wang¹

¹Institute of Modern Physics, Chinese Academy of Sciences, Lanzhou 730000, China

²School of Nuclear Science and Technology, University of Chinese Academy of Sciences, Beijing 100049, China

Superconducting radio-frequency (SRF) cavities are the core components of SRF linear accelerators, making their stable operation considerably important. However, operational experience from different accelerator laboratories has revealed that SRF faults are the leading cause of short machine downtime trips. When a cavity fault occurs, system experts analyze the time-series data recorded by low-level RF systems and identify the fault type. However, this requires expertise and intuition, posing a major challenge for control-room operators. Here, we propose an expert feature-based machine learning model for automating SRF cavity fault recognition. The main challenge in converting the “expert reasoning” process for SRF faults into a “model inference” process lies in feature extraction, which is attributed to the associated multidimensional and complex time-series waveforms.

Existing autoregression-based feature-extraction methods require the signal to be stable and autocorrelated, resulting in difficulty capturing the abrupt features that exist in several SRF failure patterns. To address these issues, we introduce expertise into the classification model through reasonable feature engineering. We demonstrate the feasibility of this method using the SRF cavity of the China Accelerator Facility for superheavy Elements (CAFE2). Although specific faults in SRF cavities may vary across different accelerators, similarities exist in the RF signals. Therefore, this study provides valuable guidance for fault analysis of the entire SRF community.

Keywords: superconducting radio-frequency cavity, fault recognition, machine learning, feature engineering, particle acceler-

Introduction

The China Initiative Accelerator Driven System (CiADS) [1], currently under construction, employs a high-power linear accelerator at its front-end to generate a 500 MeV proton beam with an intensity of 5 mA [2–4]. To verify the feasibility of a continuous wave (CW) proton beam with a current of 10 mA, the China ADS Front-End Demo Linac (CAFe) was built. In March 2021, CAFe achieved its design goal with the successful commissioning of a 10 mA, 205 kW CW proton beam at an energy of 20 MeV [5].

The synthesis and property study of superheavy nuclei is an important frontier and one of the difficulties in current nuclear physics [6–10]. Since 2021, the CAFe facility has been upgraded to CAFE2 (China Accelerator Facility for su-

perheavy Elements) for the exploration of new isotopes with an operating beam intensity of approximately 10 pμA [11, 12]. The layout of CAFE2, as shown in Fig. 1 [Figure 1: see original paper], includes both normal conducting and superconducting (SC) sections, and a new gas-filled recoil separator, SHANS2 (Spectrometer for Heavy Atoms and Nuclear Structure-2), was constructed at the end of the beam line [13].

The SC section contains a total of 23 SC half-wave resonator (HWR) cavities assembled in four cryomodules (CM1–CM4) and regulated with an individual digital low-level radio-frequency (LLRF) system, in which CM1 to CM3 are each equipped with six HWR010 cavities, while CM4 is equipped with five HWR015 cavities [5, 14–16]. HWR010 and HWR015 are two cavity types named according to their optimal β value, with their operational parameters shown in Table 1.

To meet the high demand for beam availability in the future CiADS, the research team at the Institute of Modern Physics is working diligently to enhance the reliability of various subsystems of CAFE2. However, owing to the stringent operating conditions of the SC cavities (high power, electric field, and frequency) and the extremely narrow operating bandwidth [17], cavity failures easily occur when subjected to disturbances (e.g., mechanical vibrations). Operational experiences from different accelerator laboratories have revealed that the leading causes of short machine downtime trips are SRF faults [18–20]. Rapidly identifying the causes of faults and reducing the SC cavity failure rate for stable operation of the accelerator are imperative.

When an RF fault occurs, the LLRF's data acquisition (DAQ) system simultaneously records 16 RF signals from each cavity, providing comprehensive fault information. This process is triggered when the LLRF system for any cavity in a cryomodule detects a fault condition (e.g., field fluctuation beyond the tolerance limit). Based on this data, system experts can analyze the fault types and causes to comprehend the underlying physical mechanisms. To implement appropriate measures for fault handling, the accurate and swift identification of fault patterns is essential. However, the diversity of fault modes and the similarity of fault characteristics complicate fault analysis. Although control-room operators have access to raw waveform data captured during fault occurrences, correctly interpreting the signals requires expertise. Additionally, a fault in one pattern can trigger a different pattern through several physics effects, and a fault in a single SC cavity may propagate to adjacent cavities, leading to group faults in a cryomodule. In such cases, providing near real-time fault feedback is rather crucial for control-room operators.

Identifying the offending cavity with existing software and hardware is difficult to do automatically. Traditional methods are generally limited by the requirement of expertise and cannot quickly process large amounts of fault data. In recent years, machine learning (ML) methods have made remarkable progress in pattern recognition tasks and are widely used in various fields [21]. As a data-driven algorithm, ML shows potential applications in particle accelerators,

such as beam optimization, intelligent control system, anomaly detection, and fault diagnosis [22–27]. For fault-pattern recognition of SC cavities, the challenge lies in solving a multidimensional time-series classification problem. SC cavity faults occur in milliseconds or microseconds; therefore, a high sampling rate is required to capture the signal features when the fault occurs. However, existing time-series feature-extraction models, such as long short-term memory (LSTM) and gated recurrent units (GRU), cannot process such a long sequence, and the fault information is lost in the data after downsampling [28]. Therefore, implementing feature engineering is vital for fault identification. At the Jefferson Laboratory (JLab), the Continuous Electron Beam Accelerator Facility (CEBAF) uses the autoregression (AR) method to extract features from the cavity voltage and the incident and reflected voltages, and builds an ML-based fault classification model [18]. Compared to expert results, the method achieves a classification accuracy of 82%. Research results from CEBAF indicate that the performance of the ML method in identifying abrupt faults (e.g., LLRF control trips and E-quench faults) is unsatisfactory [18, 29]. This may be primarily attributed to the limitations of AR methods in extracting non-stationary signals.

In this work, we introduce an expert knowledge-driven approach to feature engineering construction, aiming to address the limitations of existing methods in automated fault identification. We analyzed the historical data generated via the operation of CAFE2 and categorized the SC cavity faults into eight types. Based on the formation mechanism and waveform characteristics of different faults, we designed reasonable feature engineering to transform raw data into an intermediate representation that expresses the underlying data patterns. Subsequently, we evaluated the effect of feature engineering on CAFE2 through two aspects: confusion matrix and information gain, to obtain a comprehensive understanding of its impact on model performance. Finally, fault analysis was conducted on the historical data of CAFE2 operation, tallying the most prevalent fault types for each cavity. This analysis provides valuable guidance for the future maintenance and upgrade of SC cavities, enabling the development of preventive measures against common faults as well as the optimization of maintenance strategies to ensure system stability and sustained high-efficiency operation.

The remainder of this paper is organized as follows. In Section II, the method for acquiring offline fault data of the SC cavity is introduced and the criteria for labeling fault types are discussed. In Section III, the development of ML models is discussed, including the calibration of the raw data, the implementation of feature engineering, and the theory of ensemble learning methods. Finally, the performance evaluation of the aforementioned method based on 2023 operational data is presented, followed by a discussion regarding future research.

II. Data Analysis and Labeling

A. Data acquisition

For each cavity fault, the newly developed DAQ system synchronously captures timestamps and saves waveform records of 16 RF signals from each of the cavities in the cryomodule. The DAQ system comprises LLRF and EPICS (experimental physics and industrial control systems) components along with various high-level applications that collaborate to collect and store data for subsequent offline analyses and inspections.

A waveform capture module was developed to gather RF time-series signals after a fault occurrence and write them to a file for later analysis. Each of the 16 harvested waveform signals comprises 50000 points. The trigger is configured such that approximately 80% of the recorded data precede the fault, whereas 20% follow the fault. Subsequently, the collected waveform data are written to network storage and uploaded to a data server via a waveform-specific web service. Finally, all waveform-related data are backed up online indefinitely to tape daily and compressed monthly to reduce online storage (Fig. 2 [Figure 2: see original paper]).

According to the different research requirements of CAFE2, the sampling rate is typically adjusted within the range of 10 kHz to 100 kHz (based on the dominant fault pattern of the specific cavity), resulting in approximately 0.5–5 s of fault data. As the feature-engineering method proposed in this study is not affected by the sampling rate, we extract 0.5 s segments from all fault waveforms as the raw data for feature engineering, of which 80% is pre-fault information and 20% is post-fault information.

B. Data labeling

We analyzed the fault data generated by the 23 SC cavities of the CAFE2 accelerator between January 2023 and November 2023 and labeled 1932 typical samples for supervised learning. When a fault is triggered in the SC cavity, the low-level system sends 16 channels of RF signals to the data server. Notably, these 16 channel signals include 6 real measurement signals extracted using a pick-up coupler and directional coupler, as well as 10 control signals generated internally within the FPGA (e.g., feedforward signals for pulse beam compensation or calibration signals). In this work, the cavity voltage (V_c^*), incident voltage (V_f^*) and reflected voltage (V_r^*) from the 6 measurement signals and the LLRF output signal (V_{LLRF}) from the 10 FPGA signals were selected for fault analysis, as shown in Fig. 2b. Generally, the other signals can be obtained by linearly transforming these four signals. Table 2 lists key signals and parameters.

Based on these signals, we summarize eight fault modes: thermal quench (quench), helium pressure fluctuations (helium fluc), electrical quench (E-quench), flashover, microphonics, ponderomotive, LLRF trip, and single-cavity

off (cavity off), as shown in Fig. 3 [Figure 3: see original paper]. Notably, during the commissioning phase of CAFE, faults induced by transient beam-loading effects with a beam current of 10 mA are common; however, with CAFE2 operating in the CW mode at microampere-level beam currents, this fault pattern is essentially absent [30, 31]. We briefly describe the process of fault-signal analysis and labeling from the perspective of system experts.

Quenching refers to the localized overheating of the SC cavity wall, which results in the premature breakdown of superconductivity (thermal breakdown). A quench typically manifests as a rapid drop in the unloaded quality factor (Q_0) and the loaded quality factor (Q_L). When a fault occurs, the cavity's Q_L and detuning can be solved according to V_c and V_f , respectively, as shown in Eq. (1) [32].

$$\begin{cases} \frac{d|V_c|}{dt} = \omega_{0.5}\rho \cos(\theta - \phi) - \omega_{0.5}|V_c| \\ \frac{d\phi}{dt} = \Delta\omega = \omega_{0.5} \frac{\rho \sin(\theta - \phi)}{|2V_f|} - \omega_{0.5} \end{cases}$$

where ϕ and θ represent the phases of V_c and V_f , respectively; and $\omega_{0.5}$ and $\Delta\omega$ represent the half-bandwidth and detuning, respectively. Fig. 4a [Figure 4: see original paper] shows the calculated Q_L and the detuning based on the waveforms of the four fault modes in Fig. 3: quenching, helium fluc, E-quench, and flashover [33, 34]. A considerable change in the cavity Q_L was observed only for the quench patterns, whereas no change in the cavity Q_L was observed for other fault patterns.

Quenches induce changes in the heat load of the cryosystem, resulting in rapid fluctuations in helium pressure over a short period, ultimately causing the SC cavities within the cryomodule to undergo considerable detuning on the millisecond scale. When detuning exceeds the cavity bandwidth, the power source output reaches saturation and eventually triggers multiple cavity faults (Fig. 5a [Figure 5: see original paper]). Typically, helium pressure fluctuations (helium flucs) are secondary faults that are induced by quenching. However, in a few cases (e.g., SC magnet quenching or cryogenic system control logic faults), we observed simultaneous helium pressure changes in the four cryosystems without any cavities experiencing quenching (Fig. 4b). In this study, we labeled these fault patterns as helium flucs.

E-quench typically manifests as a sudden and complete loss of stored energy in the cavity. JLAB interpreted this loss as the effect of the release of numerous electrons inside the cavity, which absorbed the cavity energy. A flashover involves an FE-initiated discharge on an RF ceramic window surface [18, 33]. It typically does not cause any V_c degradation but can result in burst noise in the cavity's pick-up signal. Notably, E-quench can also be accompanied by burst noise events. The main difference between the two is that E-quench can cause total or partial gradient loss, whereas flashover does not cause such a loss [34]. Experience from CAFE2 operations suggests that when the gradient loss

exceeds 30%, E-quench may further trigger a quench fault and cause multiple cavity failures within the same cryomodule (Fig. 5b). Conversely, when the gradient loss is less than 30%, multiple cavity failures generally do not occur. Therefore, in this work, we categorize E-quench events with gradient loss less than 30% as “flashover” faults, and those with gradient loss greater than 30% as “E-quench” faults.

Ponderomotive oscillatory instabilities result from the nonlinear coupling between the electrical and mechanical modes of the cavity, which is accompanied by an accelerating gradient and detuning that begins to oscillate with increasing amplitude [35]. Based on measurements of the cavity mechanical mode transfer function, most cavities exhibit a significant mechanical mode around 125 Hz [36]. As shown in Fig. 5c, when a cavity undergoes oscillations due to the ponderomotive effect, the oscillations in the cavity can be transmitted to other cavities, resulting in a multicavity fault. Notably, the formation of ponderomotive oscillations depends on factors such as feedback parameters, Lorentz detuning coefficient, and cavity detuning [37, 38]. In this example, no ponderomotive oscillations are observed in CM2–5.

Microphonics are changes in the cavity frequency caused by connections to the external world, such as vacuum pump vibrations, at a frequency generally less than 50 Hz. Compared with ponderomotive instability, cavity detuning induced by microphonics is determined by external vibration sources, with the oscillation energy typically not exhibiting divergent growth. As shown in Fig. 5d, microphonics typically occur in multiple cavities. Notably, microphonics and helium fluxes are commonly grouped as microphonic faults [35]. In this study, we specifically distinguished between vibration-dominated and non-vibration-dominated (e.g., cryogenic system-dominated) cases. Therefore, we categorized these into two fault modes.

There are many possible causes of LLRF faults, such as electronics being affected by radiation showers in the tunnel, leading to single-event upsets that flip a bit in the digital data stream [39]. In CAFE2, the most common type of LLRF fault is triggered by the control logic inside the FPGA. As shown in Fig. 4c, around 0 ms, the DAC output suddenly drops to zero, causing a transient fluctuation in V_c and triggering a fault. We carefully checked the internal logic of the LLRF but found no issues. One possible reason is that clock glitches disturb the accumulator of the proportional-integral (PI) controller. The yellow curve in Fig. 4c shows the PI output obtained from the simulation based on the input of the PI controller, which differs from the DAC output by a fixed constant. LLRF faults are generally single-issue faults, implying that they do not cause further faults in multiple cavities. Similar to the case in CEBAF [18], we classify “cavity turn off” events triggered by external machine interlock signals as “cavity off” modes, including arc interlock or RF source interlock.

Based on the above steps, we completed data annotation and labeled a total of 1932 fault events. The distribution of sample counts for each fault type is shown in Fig. 6 [Figure 6: see original paper]. Because the first cavity to trigger

a fault can usually be determined based on the time of the fault occurrence, in this study, we focused on identifying the fault type of the source cavity.

III. Machine Learning Method

For fault pattern recognition in an SC cavity, the challenge lies in solving a multidimensional time-series classification problem. In this section, we introduce how to extract fault-related features from raw RF signals and construct a machine-learning model.

A. Data preprocessing

The cavity voltage (V_c^*), incident voltage (V_f^*), and reflected voltage (V_r^*) were selected for feature extraction and analysis (* represents the raw measurement data). Previous studies demonstrated the superior predictive capability of these three signals for fault classification and fault warning [18, 23]. Before feature extraction, it is imperative to perform calibration and normalization procedures on the raw signals. The calibration of the actual V_f and V_r is given by [40, 41]

$$\begin{cases} V_{f\text{-cali}} = XV_f \\ V_{r\text{-cali}} = YV_r \end{cases}$$

where X and Y are complex coefficients obtained by solving the linear regression equations [42].

Subsequently, the three signals were normalized relative to V_c using the following formula:

$$V_c = \frac{V_c^*}{S}, \quad V_{f\text{-cali}} = \frac{V_{f\text{-cali}}}{S}, \quad V_{r\text{-cali}} = \frac{V_{r\text{-cali}}}{S}$$

Note that $S = |V_c^*|$ represents the mean value of $|V_c^*|$ in the steady state.

B. Feature engineering

The success of ML methods often depends on data and features, with feature engineering playing a crucial role and directly affecting the performance, generalization, and interpretability of the models. Fig. 3 shows the amplitude and phase changes in V_c , V_f , and V_r recorded by the LLRF system when a fault occurs in the SC cavity. Based on the experience of experts in inferring fault types, we extracted eight features related to fault types, which were calculated from the amplitudes and phases of V_c , V_f , and V_r . These features serve as intermediate representations of the raw data and are employed as model inputs. The following section introduces the calculation methods for the eight features.

First, we introduce the thermal quenching (quenching) recognition feature Q_{id} . When a cavity quenches, its Q_L decreases rapidly [43]. Although Q_L serves as a

hallmark for distinguishing quench faults from other modes, its computational process requires solving the V_c differential equation (Eq. (1)), which is highly time consuming, whereas fault identification must be accomplished within milliseconds. Next, we introduce the quench identification features based on the cavity coefficient difference equation. Let $V_c = re^{i\phi}$ and $V_f = (1/2)\rho e^{i\theta}$. Based on the differential equation of the cavity without a beam and separating its real and imaginary parts, Eq. (4) can be obtained from [40, 42, 44]

$$\begin{cases} \dot{r}_c + r_c \omega_{0.5} = \omega_{0.5} \rho \cos(\theta - \phi) \\ r_c \dot{\phi} - r \Delta \omega = \omega_{0.5} \rho \sin(\theta - \phi) \end{cases}$$

where r_c denotes the amplitude of V_c predicted using the differential equation of the cavity. Let $\Delta\theta = \phi - \theta$; we construct a new signal $\alpha = \rho \cos \Delta\theta$. The real parts of Eq. (4) can be expressed as

$$r_c(n) = T_s \omega_{0.5} \alpha(n-1) + (1 - T_s \omega_{0.5}) r_c(n-1)$$

where T_s is the sampling period. Based on Eq. (5), we solved the values of r_c for the eight fault patterns, as depicted in each subplot in Fig. 7 [Figure 7: see original paper], which corresponds to Fig. 3.

Let $e = r_c - r$. When cavity quenching does not occur, the differential equation of the cavity can effectively describe its dynamic behavior. Therefore, the predicted amplitude value, r_c , should be consistent with the measured value, r ; that is, the error $e \rightarrow 0$. When cavity quench occurs, Q_L and $\omega_{0.5}$ change by approximately one order of magnitude, as shown in Figs. 5. In this case, the dynamic behavior of V_c no longer satisfies the coefficient difference equation above. Therefore, the predicted value r_c does not agree with the measured value r , and the error e increases sharply, as shown in Fig. 7a. In addition, some strong transient disturbances on the order of microseconds, such as the dark current triggered by the E-quench fault, can lead to large transient spikes in the error signal e , as shown in Fig. 7c. Therefore, we employed the area under the curve e as the quench fault-recognition feature, which is given by

$$Q_{id} = \int_{t=0}^{T_Q} |e| dt$$

where $T_Q = 100$ ms, as indicated by the shaded area in Fig. 7a. We calculated Q_{id} for each of the labeled 1932 samples. Fig. 9a [Figure 9: see original paper] shows that Q_{id} of the quench fault is significantly greater than that of the other faults.

Helium fluc and ponderomotive faults are prone to cavity phase detuning. We utilized a quantity related to the detuning angle, the mean phase difference

between V_c and V_f , as a learning feature to identify these two fault types, which is expressed as follows:

$$\Delta\Theta = |\text{mean}(\Delta\theta)|$$

As shown in Fig. 9b, the phase difference between the helium flux and ponderomotive faults is approximately 40° .

Both ponderomotive and microphonics are related to mechanical vibrations, as shown in Fig. 3e and Fig. 3f, where V_c exhibits significant oscillatory characteristics. Given this, we apply fast Fourier transforms (FFT) to convert the detuning signal ($\Delta\theta = \theta - \phi$) from the time domain to the frequency domain to obtain frequency-domain representations for the analysis of frequency components and spectral characteristics. Subsequently, as illustrated in Fig. 8b [Figure 8: see original paper], the main frequency of this signal, $F_{\max}$, and the ratio of the energy of the main frequency to the total energy, F_{ratio} , are extracted as features for classifying such faults. As shown in Fig. 9c, the main frequency of the ponderomotive fault is concentrated around 130 Hz, and the main frequency of microphonics faults is in the range of 25 Hz–50 Hz. The quantity F_{ratio} was calculated as follows:

$$F_{\text{ratio}} = \frac{\int_{F_{\max}-\Delta F}^{F_{\max}+\Delta F} |y_{\text{FFT}}|^2}{\int_0^{F_s/2} |y_{\text{FFT}}|^2}$$

where F_s is the sampling frequency of the waveform data, y_{FFT} is the normalized power spectral density, and ΔF is 5–10 Hz.

The flashover, E-quench, and LLRF trip faults induced a rapid change in the amplitude of V_c on the submillisecond timescale, exhibiting significant gradients at the transition points. We extracted the relative change in the amplitude of V_c as a learnable feature, denoted by E_{id} , to quantify the deviation of the transient signal from the baseline. The calculation is as follows:

$$E_{id} = \max \left(\frac{\Delta r}{\max(r) - \min(r)} \right)$$

where $\Delta r = r(n+1) - r(n)$ is the first-order difference in V_c amplitude. From the analysis of Fig. 8c and Fig. 9e, it can be observed that flashover and E-quench exhibited significantly large values of E_{id} . Simultaneously, the amplitudes of V_f for the three faults mentioned above exhibited sudden changes. We measured this change process using the first-order difference in the amplitude of V_f and extracted the position with the maximum difference as a learnable feature for the ML model. This is calculated as follows:

$$\Delta\rho_{\max} = \max(|\Delta\rho/2|)$$

where $\Delta\rho = \rho(n+1) - \rho(n)$. From the statistical results in Fig. 9e and 9f, it is evident that the E_{id} and $\Delta\rho_{\max}$ values for the flashover, E-quench, and LLRF faults are significantly larger than those for the other faults, and the distributions of E_{id} and $\Delta\rho_{\max}$ vary among the three faults. In addition to the above expert features, the following statistical features are included: changes in the rms radius of the V_c amplitude before and after the fault, noting the rms radius before the fault as r_{rms1} and that after the fault as r_{rms2} .

In the preceding section, we systematically clarified the theory and calculation methods of the designed expert features, from complex physical attributes to basic statistical features, which are of great significance to the analysis and decision-making processes as intermediate representations of the raw data. Table 3 lists simple definitions of the eight expert features. The distribution results for each feature in the 1932 labeled samples are shown in Fig. 9. As can be observed, except for the quench fault, it is challenging to distinguish other faults based on a single feature. Therefore, it is necessary to explore complex combinations of features.

In addition to the aforementioned eight expert features, we employed the AR method to explore the autocorrelations within sequential data, capturing the trends and periodicities in the signal. In the AR method, it is assumed that the current value of a time series is correlated with several past values; that is, past observations impact the current value. This autocorrelation can be controlled by the order (p) of the AR model, where p indicates the extent to which past observations affect the current values. By linearly combining past observations to predict future values, a mathematical expression for AR can be obtained as follows [18]:

$$X_t = c + \phi_1 X_{t-1} + \phi_2 X_{t-2} + \cdots + \phi_p X_{t-p}$$

where ϕ_1 and ϕ_p are the weight parameters of the model and c is a constant term. The temporal features of the signal can be obtained using the above formula to fit the amplitude and phase of V_c , amplitude of V_f , and detuning angle ($\Delta\theta$), while extracting the weight parameters obtained after fitting. Different fault types exhibited significant differences in the distribution ranges of the weight parameters. These differences can be exploited to distinguish various SC cavity faults. Subsequently, we will discuss the performance of the expert and AR features for the identification of SC cavity faults.

C. Ensemble learning models

Ensemble learning, an ML technique that combines the predictions of multiple models to improve overall performance, is widely used in various data-driven

scenarios [45]. Ensemble models mitigate the weaknesses inherent in a single algorithm by aggregating diverse predictions, resulting in improved accuracy and robustness. Moreover, ensemble learning excels in handling complex and high-dimensional data, where individual models may struggle. The diversity introduced through different learning approaches or models helps reduce overfitting and provides a more generalized and reliable solution. Furthermore, ensemble methods, such as bagging and boosting, offer versatility across a spectrum of tasks, making them adaptable to different types of datasets and problems. Overall, exploiting the collective intelligence of multiple-model position ensemble learning is a powerful technique for optimizing the predictive outcomes of ML models.

Random Forest (RFs) is a model based on decision tree classifiers, using an ensemble approach that utilizes bagging among multiple decision trees [46]. The core idea behind bagging is to create multiple subsets of the original training dataset using random sampling with replacement. Each subset is used to train a separate base model. The final prediction is obtained by aggregating the predictions of all the individual base models, thereby reducing the risk of bias and variance associated with individual trees. For regression tasks, this aggregation is usually performed by averaging the predictions, whereas for classification tasks, a majority voting mechanism is often employed. The “random” in RFs stems from the introduction of randomness in two key aspects: bootstrap sampling and feature selection. Bootstrap sampling can generate multiple differentiated subsets to train a range of base models and is fundamental in ensemble learning methods such as bagging [47]. Feature selection refers to the process of selecting a subset of relevant features to construct individual decision trees within a forest. Instead of considering all available features to determine the best split at each node, only a randomly chosen subset of features is evaluated. This random selection of features introduces variability among trees because different trees may consider different features for splits, even if they are trained on the same data, which contributes to the robustness and generalization ability of the model. In RFs, the feature selection process is controlled by the key parameter “max_features”. Besides that, the “n_estimators” parameter specifies the number of trees in the forest; more trees generally improve accuracy but increase computational cost. The “max_depth” parameter controls the maximum depth of each tree; deeper trees capture more complex patterns but may overfit the data.

eXtreme Gradient Boosting (XGBoost) is a gradient boosting algorithm known for its efficiency and excellent predictive performance [48]. Unlike bagging methods that train models independently in parallel, boosting sequentially trains boosters (such as gbtrees or gblinears), with each tree attempting to correct the errors of the previous tree with the aim of incrementally improving accuracy. The final prediction is the weighted sum of the predictions from all the individual trees. During the iterative training process, observations are assigned different weights based on their classification; misclassified observations are given more weight, whereas correctly classified observations are given less weight. This

process is achieved by focusing on the model residuals, which directs the subsequent models to focus more on hard-to-predict cases. To prevent overfitting, XGBoost applies “shrinkage” during training, meaning it does not fully trust the residuals learned by each weak learner. This is achieved by multiplying the residual value that each weak learner fits by a “learning_rate” in the range of $(0, 1]$. A lower “learning_rate” makes the model more robust to overfitting by ensuring that each tree makes only a small adjustment to the model. This typically requires more trees to reach the same level of performance as a model with a higher “learning_rate”. Therefore, there is a trade-off between “learning_rate” and “n_estimators”. Additionally, XGBoost combines parameters such as “max_depth”, “gamma”, and regularization parameters (L1 and L2) to further reduce overfitting. It also uses “subsample” and “colsample_bytree” to introduce randomness by specifying the fraction of the training data and features used for each tree, respectively. A robust model can be achieved by coordinated optimization of these parameters.

Next, we separately evaluated the performances of the two ensemble learning methods in identifying SC cavity faults.

IV. Results and Discussion

A. Data visualization

Before model training, we applied principal component analysis (PCA) to perform dimensionality reduction and visualized all samples in 2D coordinates. The results are shown in Fig. 10 [Figure 10: see original paper], where the clustering, distribution, and correlations within the data are clearly observed. This visualization aids experts in better understanding the data and uncovering potential relationships, thereby facilitating a more detailed categorization of the original dataset. Another important aspect of dimensionality-reduction visualization is the identification of outliers or anomalous points in each class to check for errors in the manual labeling process. Manual labeling requires a system expert to have considerable experience and intuition regarding SRF cavities operating with beams and to understand the complex physical mechanisms underlying the faults, for which PCA serves as a valuable auxiliary tool. Fig. 10 indicates the presence of several outliers. After verification with domain experts, corrections were made to several erroneously labeled samples. For instance, a cavity-off fault was mislabeled as an LLRF trip, a helium fault was mislabeled as a microphonics fault, another helium fault was mislabeled as a quench fault, and several helium faults were mislabeled as ponderomotive faults. Through the aforementioned scrutiny, rectifications were made to human-labeled errors, and the mislabeled samples were relabeled and used for subsequent model training.

B. Model performance evaluation

A class imbalance problem exists in the collected fault data. Random splitting (or k-fold) methods may be used when samples of a category are rare or

missing from the test set. Therefore, we used stratified k-fold cross-validation to ensure that each fold maintained the same class distribution as the original dataset. This method can be imported from the sklearn library and provides a more reliable estimate of model performance across different subsets of data. Subsequently, two ensemble learning models, RFs and XGBoost, were selected for fault-type identification.

RFs and XGBoost contain numerous hyperparameter settings that are typically optimized using the GridSearchCV method, which automatically scans the specified parameter range and returns the best hyperparameter combination. The GridSearchCV method has a high computational overhead because of the need to test all the parameter combinations. Herein, we experimented with heuristic search algorithms, such as particle swarm optimization (PSO) and genetic algorithms (GA), to determine the optimal parameters. Although the PSO method converges quickly, the performance of the model is slightly better than that obtained using the GridSearchCV method with a larger step, which may be because RFs and XGBoost are relatively tolerant to variations in certain hyperparameters. Finally, employing the hyperparameter combinations searched by GridSearchCV, XGBClassifier (learning_rate = 0.05, n_estimators = 250, max_depth = 5, min_child_weight = 5, gamma = 0.2, subsample = 0.7, colsample_bytree = 0.6), and RandomForestClassifier (n_estimators = 200, max_depth = 17, max_features = 3) are utilized to build the final models. These models were evaluated using stratified 5-fold cross-validation, and the results are presented in Table 4 as the mean and variance of the F1 scores.

Different feature combinations are tested in Table 4, including the use of AR features, expert features, and a combination of both in the three scenarios. The expert features comprise the previously mentioned Q_{id} , $F_{\max}$, F_{ratio} , E_{id} , $\Delta\rho_{\max}$, $\Delta\theta$, r_{rms1} , and r_{rms2} values. The AR features are the weight coefficients obtained by fitting the amplitude and phase of V_c , amplitude of V_f , and detuning ($\Delta\theta$) using the 3rd order AR method. As shown in Table 4, the ensemble learning method using AR features achieved an accuracy of 90% for the multiple-fault classification tasks. The corresponding accuracy using expert features was 95%, and the accuracy using a combination of the two was greater than 96%. Both ensemble learning models exhibited comparable performance while significantly outperforming the support vector machine (SVM) method. Notably, in our experiments, AR models with orders higher than three did not show significant performance improvements, and even ensemble models led to a slight decrease in classification ability. Therefore, the third-order AR coefficients were determined to be the best-performing features for the AR-based method. In addition, we considered the computational cost and found that the feature extraction time per sample using expert engineering was 0.0266 s, whereas for the third-order AR model it was 0.0378 s under the same test conditions. Thus, expert engineering was approximately 30% faster than the third-order AR model.

In conclusion, our feature-engineering scheme demonstrated significant advan-

tages in terms of both model performance and computational efficiency.

Further analyses were performed using the XGBoost model. We conducted a comprehensive analysis of the classification accuracy of the model for different categories using a confusion matrix. Confusion matrix analysis identifies a model's weaknesses, enabling targeted adjustments to parameters, feature engineering, and other aspects of model optimization. Fig. 11 [Figure 11: see original paper] (left) shows that the XGBoost model based on AR features has a lower accuracy for faults such as E-quench, flashover, and microphonics, which may be attributed to the difficulty of the AR method in capturing the signal features of these three fault types. As shown in Fig. 12 [Figure 12: see original paper], the amplitude of V_c for E-quench exhibits significant abrupt changes, leading to substantial errors at the mutation positions when the AR method is employed to fit these signals. For continuously changing signals, such as microphonics, the AR method can capture data trends. However, this trend may be insufficient to describe microphonics fault features, thereby reducing the accuracy of the model in identifying microphonics faults.

As displayed in Fig. 11 (right), the expert feature-based XGBoost model effectively addresses the challenges associated with the AR method. The introduction of expert features increases the accuracy of the model in capturing essential task-related features, thereby enhancing its applicability and performance. Subsequently, we interpreted the reasons for the improvement in the performance of the model from the perspective of feature importance analysis.

First, the multiclass problem was transformed into a binary classification problem, after which the information gain was utilized as a measure of the contribution of each feature to the model's predictions. As shown in Fig. 13 [Figure 13: see original paper], during the identification of the quench fault, the Q_{id} feature exhibits the highest contribution. For the recognition of ponderomotive and helium faults, the F_{max} feature was the most influential. For E-quench fault identification, the E_{id} feature exhibits the highest contribution. This indicates that the optimal segmentation features selected by the XGBoost model based on information gain align with the reasoning process adopted by experts during the fault analysis. Furthermore, various feature combinations have been used in the identification process for each fault, particularly for microphonics faults, which pose a major challenge for control room operators. The significance of this study is substantiated in terms of rational feature engineering and model interpretability.

C. Big data analysis of cavity faults

The trained XGBoost model was employed to analyze the historical data generated by CAFE2 during its operation. The fault data for CAFE2's daily operations are packaged into zip files, each containing four folders that store the RF signals of the fault cavities in the four cryomodules (CM1–CM4). Each fault event is named as “cavity name” + “fault time” (accurate to microseconds).

Algorithm 1 summarizes the workflow of the ML method for classifying offline fault events. Notably, the output fault time, cavity name, and fault type can be used in future collective fault analyses.

Using the ML model based on fault data from the second half of 2023, the probability of faults for a given pattern occurring in each cavity was calculated, as shown in Fig. 14 [Figure 14: see original paper], where the cavities prone to faults in this particular pattern are highlighted. The histograms reveal that the results of the AR-based and expert feature-based models were generally consistent when analyzing historical big data. Notably, the statistical results for E-quench (Fig. 14c) and microphonics (Fig. 14f) faults, the AR model identified CM3–5 as prone to E-quench and CM4–1 as prone to microphonics. After verification, the expert feature-based method classified these faults as Flashover or helium, with subtle differences observed in the corresponding cavities in Fig. 14d and 14b. Subsequently, we consulted the fault data with subject-matter experts, and their assessments concurred with the inferences made by the expert feature-based model. These findings further substantiate the generalization capability of the proposed method. Moreover, the statistical results of the AR feature-based model serve as a comparative baseline, offering an alternative perspective that reinforces the robustness of our conclusions.

As shown in Fig. 14e and 14f, CM2–2 and CM2–3 are susceptible to vibration-induced microphonics and ponderomotive faults. In subsequent operations, we increased the loop gain of the low-level system corresponding to CM2–2 and CM2–3. CM1–5, CM3–1, and CM1–2 were identified as the primary sources of E-quench and quench faults; we will reduce the acceleration gradient of these cavities in subsequent operations. In conclusion, employing ML for big data analysis is of great significance for enabling system experts to quickly identify the sources of faults and ensure the stable operation of accelerators.

D. Experience for feature engineering

This study provides a summary of fault types occurring in SRF cavities operating in the CW mode, along with discussions on fault mechanisms and feature engineering methods. Although the specific faults in SRF cavities may vary across different accelerators, there are similarities in the waveforms. Therefore, the feature engineering techniques proposed in this study offer valuable insights into the detection of faults in the SRF community.

1. For quench and helium faults, quantities such as Q_{id} and $\Delta\Theta$ can be used for identification.
2. For vibration-related faults, such as ponderomotive and microphonics faults, methods such as FFT and wavelet transforms can be employed to extract the main vibration frequency and its corresponding energy.
3. For faults involving transient changes, such as E-quench, flashover, and LLRF trips, the first-order difference can be utilized to extract abrupt change values.

4. Some statistical features, such as the root-mean-square radius (e.g., r_{rms1} and r_{rms2}), peak-to-peak value, and waveform factor can be used to describe the shape features of the waveforms.

These insights are valuable for the SRF community and aid in the development of fault detection and analysis techniques across various accelerators.

V. Future Work

Based on our expertise and ML methods, we successfully classified the SC cavity faults. The next step in this study involves several potential expansions.

1. Use of deep learning (DL) methods instead of ML methods for fault classification. ML methods rely on feature engineering, encompassing both expert and AR features that are fixed and cannot be tuned during training. Therefore, we will explore DL models to build an end-to-end model structure that combines inference and feature representation learning, using raw waveform signals as inputs with simultaneous optimization via gradient backpropagation. DL requires numerous training samples. Nevertheless, the ML model and PCA method proposed in this study can provide ample and reliable labeled samples for DL, thereby reducing manual costs.
2. Research on fault prediction algorithms. In previous studies, we found that an SC cavity experiences an unhealthy state when transitioning from a healthy to a fault state. If anomalous states can be predicted in advance and inhibitory measures can be implemented, fault-induced accelerator downtime can be avoided. Therefore, another extension of this study involves exploiting DL algorithms for the early prediction of failures.

VI. Summary and Conclusion

We proposed an expert-feature-based automatic recognition method for CAFE2 SRF cavity faults. The confusion matrix and feature importance analyses indicated that the implemented feature engineering technique was reasonable and successful. Moreover, this method is not restricted by the sampling rate and performs excellently with data collected at sampling rates of 10–100 kHz.

ML, as a data-driven method, cannot be sufficiently emphasized because of its reliance on data. Each step is crucial, from data collection and labeling to feature extraction. Based on our experience, we suggest combining various data visualization methods, such as feature distribution analysis, PCA/TSNE analysis, unsupervised clustering, and information gain, to improve the quality of data labeling and the understanding of underlying patterns, thus increasing the accuracy of the ML model. Currently, this method only works offline; therefore, its importance lies in data analysis. During the beam commissioning process, the model can serve as a good assistant for controlling room operators. During

the annual maintenance, the historical operation data analysis results provide valuable guidance for the maintenance and upgrading of SRF cavities.

VII. Acknowledgment

This work was supported by the project of Large Research Infrastructures “China initiative Accelerator Driven System” (Grant No. 2017-000052-75-01-000590), “Studies of intelligent LLRF control algorithms for superconducting RF cavities” (No. E129851YR0), and “National Natural Science Foundation of China” (No. 12205344).

References

- [1] Z.J. Wang, S.H. Liu, W.L. Chen et al., Beam physics design of a superconducting linac. *Phys. Rev. Accel. Beams* 37, 010101 (2024). <https://doi.org/10.1103/PhysRevAccelBeams.27.010101>
- [2] W.L. Zhan, Accelerator driven sustainable fission energy, in Proceedings of the 7th International Particle Accelerator Conference, Busan, Korea, 8–13 May 2016. <https://doi.org/10.18429/JACoW-IPAC2016-FRYAA03>
- [3] Y.S. Qin et al., Transfer line including vacuum differential system for a high-power windowless target. *Phys. Rev. Accel. Beams* 23, 113002 (2020). <https://doi.org/10.1103/PhysRevAccelBeams.23.113002>
- [4] S.H. Liu, Z.J. Wang, H. Jia et al., Physics design of the CIADS 25 MeV demo facility. *Nucl. Instrum. Meth. A* 843, 11-17 (2017). <https://doi.org/10.1016/j.nima.2016>.
- [5] Y. He, T. Tan, A.D. Wu et al., Operation experience at CAFe, in Oral Presentation of the 2021 International Conference on RF Superconductivity, virtual conference, 2021. Available at https://indico.frib.msu.edu/event/38/attachments/160/1298/MOOFV03_yuan
- [6] H.C. Manjunatha, N. Sowmya, P.S. Damodara Gupta et al., Investigation of decay modes of superheavy nuclei. *Nucl. Sci. Tech.* 32, 130 (2021). <https://doi.org/10.1007/s41365-021-00967-y>
- [7] S. Madhu, H.C. Manjunatha, N. Sowmya et al., Cr-induced fusion reactions to synthesize superheavy elements, *Nucl. Sci. Tech.* 35, 90 (2024). <https://doi.org/10.1007/s41365-024-01449-7>
- [8] T. Niwase, Y.X. Watanabe, Y. Hirayama et al., Discovery of new isotope ^{241}U and systematic high-precision atomic mass measurements of neutron-rich Pa-Pu nuclei produced via multinucleon transfer reactions. *Phys. Rev. Lett* 130, 132502 (2023). <https://doi.org/10.1103/PhysRevLett>.
- [9] H.B. Yang, Z.G. Gan, Y.J. Li et al., Discovery of new isotopes ^{160}Os and ^{156}W : revealing enhanced stability of the $N = 82$ shell closure on the neutron-deficient side. *Phys. Rev. Lett* 132, 072502 (2024). <https://link.aps.org/doi/10.1103/PhysRevLett.132.072502>
- [10] Zh. Wang, Zh.Zh. Ren, Predictions of the decay properties of the superheavy nuclei $^{293,294}\text{119}$ and $^{294,295}\text{120}$ (in Chinese). *Nucl. Tech.* 46, 080011 (2023). <https://10.11889/j.0253-3219.2023.hjs.46.080011>
- [11] L.Y. Gong, Z.J. Wang, W.P. Dou et al., Development of the heavy

ion RFQ for CAFE2. Nucl. Instrum. Meth. A 1058, 168819 (2024). <https://doi.org/10.1016/j.nima.>

[12] Z.G. Gan, W.X. Huang, Z.Y. Zhang et al., Results and perspectives for study of heavy and super-heavy nuclei and elements at IMP/CAS. Eur. Phys. J. A. 58, 158 (2022). <https://doi.org/10.1140/epja/s10050-022-00811-w>

[13] S.Y. Xu, Z.Y. Zhang, Z.G. Gan et al., A gas-filled recoil separator, SHANS2, at the China Accelerator Facility for Superheavy Elements. Nucl. Instrum. Meth. A 1050, 168113 (2023).

Note: Figure translations are in progress. See original paper for figures.

Source: ChinaXiv — Machine translation. Verify with original.