

Kriging to Kolmogorov-Arnold Network model accelerated discovery of oxygen control strategy in lead-based fast reactors

Received: 24 January 2025

Accepted: 24 September 2025

Published online: 05 November 2025

 Check for updatesShiwei Wang^{1,2}, Jiajie Chen^{1,2}, Xiaojing Liu^{1,2}  , Tengfei Zhang^{1,2}, Xiang Chai^{1,2}, Qi Lu³, Danhong Shen³ & Hui He^{1,2}  

As a cornerstone of net-zero emission strategies and Gen-IV nuclear technologies, the commercialization of Lead-based Fast Reactor (LFR) is impeded by lead-bismuth eutectic (LBE)-induced cladding corrosion. Although active oxygen control demonstrates promise under laboratory conditions, its long-term effectiveness under realistic conditions remains uncertain due to multi-physics interactions and prohibitive computational costs. To address this challenge, we introduce a high-fidelity and high-accuracy surrogate model, K2K (Kriging to Kolmogorov-Arnold Networks) with a predictor-corrector structure combined with a gradient penalty operator, thereby effectively enhancing model accuracy while mitigating non-physical extrapolations. The application of K2K enables the identification and localization of cladding failure mechanisms. Leveraging these insights, we develop a robust and comprehensive oxygen concentration control strategy, encompassing feasible concentration ranges and optimal values to support the safe, reliable, and long-term operation of LFR. Finally, we explore the potential of K2K model for analyzing multiphysics behaviors in energy systems.

Amidst the escalating global climate crisis^{1,2}, declarations such as “Net Zero Requires Nuclear Energy” and “Triple Nuclear Energy Declaration” have been successively endorsed during recent United Nations Climate Change Conferences^{3–5}. As the demand for nuclear energy intensifies, the development of fourth-generation (Gen-IV) nuclear technologies has become a pivotal strategy for mitigating climate change and advancing sustainable energy systems^{6–8}. Among Gen-IV nuclear technologies, the lead-based fast reactor (LFR) has garnered significant attention, attributed to the advantageous properties of lead-based coolants, including high boiling points, low melting points, exceptional heat capacities, and superior radiation shielding^{6,9,10}. Despite these benefits, the dissolution and corrosion of steel in lead-bismuth eutectic (LBE) present persistent challenges, constituting critical bottlenecks for ensuring the long-term structural integrity and safety of cladding materials^{11–13}.

Active oxygen concentration control has demonstrated potential in mitigating LBE-induced corrosion by promoting the formation of

protective oxide layers, as evidenced in experimental investigations and theoretical calculations^{11,14–20}. Oxygen concentration plays a crucial role in determining the formation of protective oxide layers, with the oxidation rate closely tied to environmental parameters such as LBE temperature and flow state^{12,15,21–23}. The requirement for oxygen concentration to maintain effective oxidation protection increases considerably with the rise in LBE temperature and flow rate (Fig. 1a, b)^{12,17}. LBE corrosion experiments indicate that an oxygen concentration in the range of 1e-6 to 1e-5 wt.% confers measurable benefits in mitigating corrosion under specific experimental conditions^{21,22,24}. Corrosion theories based on the oxide stability and contamination conditions provide the oxide stability boundary, serving as a guideline for oxygen control¹⁴. However, under actual LFR operating conditions, the complex interplay of LBE-induced corrosion, fuel performance, and coolant flow leads to notable deviations between the theoretical oxide stability boundary and cladding safety boundaries^{25,26}. Recent

¹School of Mechanical Engineering, Shanghai Jiao Tong University, Shanghai, China. ²Shanghai Digital Nuclear Reactor Technology Integration Innovation Center, Shanghai, China. ³Nuclear Power Institute of China, Chengdu, China.  e-mail: xiaojingliu@sjtu.edu.cn; me_hehui@sjtu.edu.cn

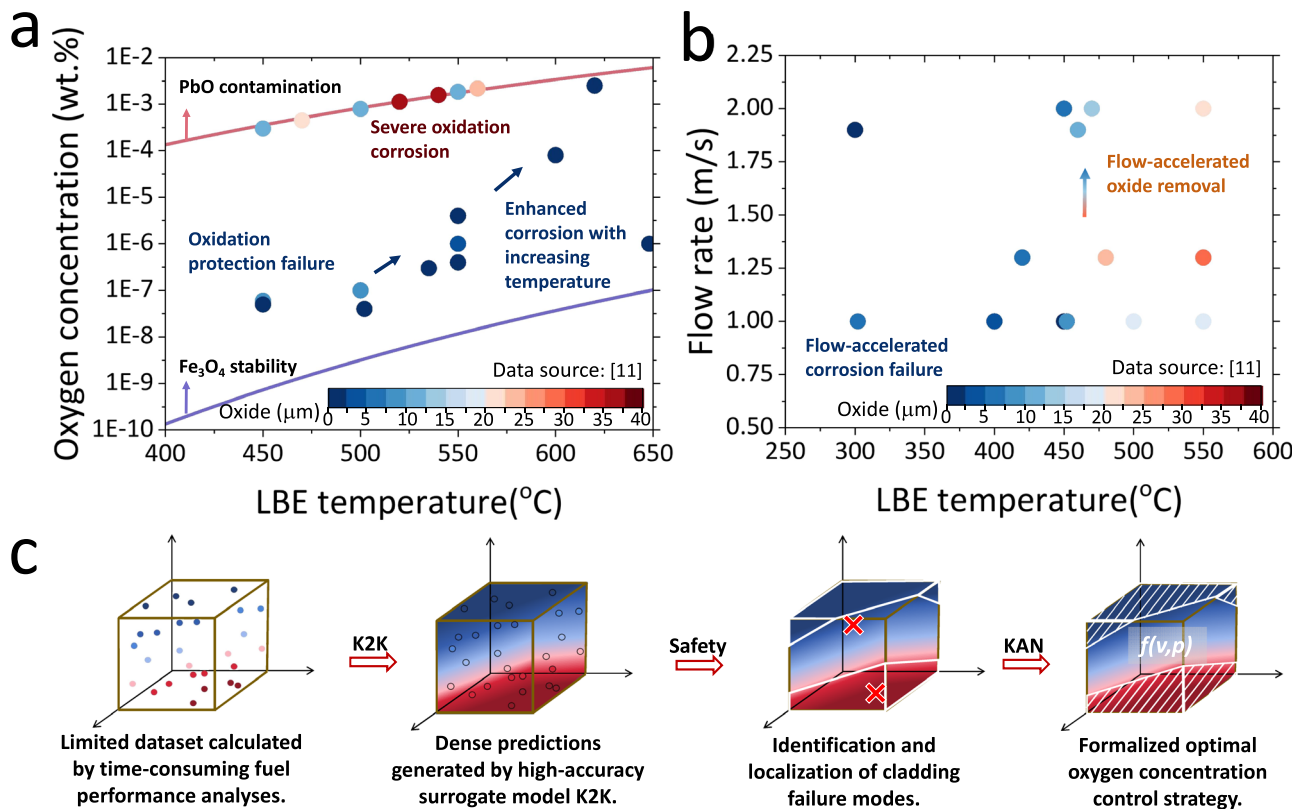


Fig. 1 | Limitations of existing experiment-oriented oxygen control strategies and conceptual framework for establishing a safety-oriented oxygen control strategy in LFR. Key experimental results with corrosion duration exceeding 1000 h are presented in **a**, **b**, with marker colour representing oxide layer thickness¹¹. **a** Existing oxygen control strategy for LFR are based on the oxidation corrosion theory, since the bioxide layer plays a crucial role in protecting nuclear fuel rod from LBE corrosion. The oxygen control window is defined by the red and blue curves, calculated from equations (S49) and (S50) in Supplementary Note B. The red curve denotes the threshold above which lead oxide deposition occur, while the blue curve represents the minimum oxygen concentration required for magnetite stability. However, experimental findings reveal that oxidation protection failure (represented in dark blue points) and excessive oxide layer growth

(represented in dark red points) may occur even within the recommended oxygen concentration range¹¹. Additionally, the oxygen concentration required for oxidation protection increases significantly with rising temperatures^{11,14–17}. **b** Liquid LBE flow corrosion experiments have demonstrated a thinning of the oxide layer with increasing flow velocity¹¹. Given the non-uniformity of temperature and flow velocity distributions within the reactor core, oxygen control strategies should account for the full spectrum of key operating conditions. **c** By constructing the K2K surrogate model, predictions of fuel performance across the primary design domain are achieved using a limited dataset. The K2K predictions enable the identification and localization of cladding failure modes. Additionally, KAN facilitates the derivation of a formalized optimal oxygen concentration control strategy.

multiphysics coupled simulations capture signs of this deviation, where oxidation corrosion-induced heat transfer degradation and thermal feedback lead to increases in cladding peak temperature, strain, and oxide layer thickness^{26,27}. The deviations in oxygen concentration boundaries, driven by the nonlinear coupling of multiphysics fuel performance parameters, highlight the limitations of existing oxygen control strategy in capturing the dynamic and multifaceted nature of real operation scenarios. Addressing this challenge necessitates an oxygen control strategy for LFR that is informed by multiphysics fuel performance models and accounts for the full spectrum of operating conditions. Such a strategy represents a crucial step toward bridging the gap between corrosion protection theories and the operational safety requirements of LFR systems.

Reproducing reactor-like conditions, encompassing neutron irradiation, fission gas pressure, and non-uniform axial heat sources, presents substantial challenges in experimental settings, making multiphysics coupling techniques critical for reconciling differences between experimental conditions and operational environments. Recent advancements, such as the development of multiphysics analysis codes for LFR cladding, exemplified by COMMA (Cladding Oxidation/Corrosion Multiphysics Modeling and Analysis), have successfully captured the spatiotemporal and multiphysical non-uniformity of cladding performance²⁶. However, multiphysics coupling approaches for fuel performance remain computationally

prohibitive. Key steps, including the solution of constitutive equations and iterative interactions across physical fields, make the exhaustive verification of oxygen concentration strategies under diverse operating conditions exceedingly time-consuming. Given the computation time of 15–20 h per case, the available sample data is limited to the scale of several hundred. To address this bottleneck, a surrogate model capable of providing high-fidelity predictions from limited sample data is indispensable. Such a model should facilitate efficient exploration of the operational parameter space without compromising accuracy. However, the constraints imposed by small datasets present a formidable barrier to advanced artificial intelligence algorithms, such as long short-term memory networks, gated neural networks, and transformers, which are prone to overfitting and instability under data-scarce conditions^{28,29}. While Physics-Informed Neural Networks (PINN) and Deep Operator Networks (DeepONet) inherently satisfy high-fidelity requirements, the extensive semi-empirical formulations, grid mappings, and complex coupling relationships in multiphysics analyses introduce formidable challenges to network construction and training^{30–32}. Thus, based on the current experience with model training under limited dataset in the energy sector, this study focuses on improving the predictive accuracy of foundational models, including Kriging, Multi-Layer Perceptrons (MLP), and Kolmogorov-Arnold Networks (KAN), aiming to achieve a balance between robustness and fidelity^{33–39}.

Essentially, the surrogate model in this context serves as a solver to multiphysical constitutive equations. Drawing inspiration from predictor-corrector methods commonly employed in differential equation solvers, the Kriging to KAN (K2K) surrogate model integrates Kriging for prediction and KAN for correction. A comprehensive evaluation of Kriging, MLP, KAN, and K2K models across metrics such as predictive accuracy, data efficiency, model capacity, hyperparameter complexity, and non-physical predictions demonstrates the clear superiority of K2K in scenarios with limited data availability. Furthermore, inspired by the discretization schemes in differential equation solvers, a gradient penalty operator is incorporated into the K2K loss function, effectively suppressing non-physical artifacts and enhancing model robustness. To further reduce data requirements, we introduce a two-stage adaptive sampling strategy tailored to the K2K framework.

Utilizing the K2K model in place of traditional multiphysics fuel performance analyses, we perform high-density predictions and feasibility assessments across the operational parameter domain (Fig. 1c). K2K effectively captures the nonlinear degradation of cladding performance arising from coupled effects of oxygen concentration and operating conditions. Building on these predictions, we identify the cladding failure boundaries encompassing oxidation protection failure, cladding thickness failure, and cladding overheating, and establish the optimal oxygen concentration determined by maximizing safety margins, offering a robust guideline for operational safety. Using KAN networks, expressions for oxygen boundaries are further derived. Leveraging the high-fidelity predictions of K2K, we develop a comprehensive oxygen concentration control strategy, delineating feasible ranges and optimal values across the primary operational space of LFR. This strategy serves as a critical reference for the design and safe operation of LFR systems, bridging theoretical analyses with practical implementation.

Results

Surrogate modelling for COMMA via K2K

Due to the high-temperature and high-radiation environment inside the reactor during operation, the fuel pellets undergo physical processes such as irradiation swelling, thermal deformation, densification, fission gas release, etc. The cladding performance is influenced by irradiation-induced elongation, thermal deformation, Fuel-Cladding Chemical Interaction (FCCI), LBE-induced corrosion, etc. However, out-of-reactor experiments can only account for LBE-induced corrosion and thermal deformation. COMMA, a comprehensive multiphysics fuel performance program for LFR, incorporates all these effects in a coupled manner to ensure the safety of LFR fuel rods²⁶. Therefore, we select COMMA as the benchmark model for this research. However, to achieve multiphysics coupling solutions, COMMA employs regression mapping algorithms, Newton-Raphson iterations, and advanced numerical methods for solving partial differential equations, such as staggered grid technique, finite element analysis, and finite volume methods, which make its computation time-consuming. To replace COMMA, the K2K model is developed and trained following the workflow outlined in Figure. S1.

Input and output parameters of K2K. Drawing upon the design parameters of the typical lead-based fast reactor ELSY and CLEAR (Table S1)^{6,40,41} and a sensitivity analysis (Table S2), oxygen concentration, linear power, and coolant flow rate are selected as K2K inputs (Table S3). In alignment with safety design criteria, ten output parameters are identified for the K2K model, encompassing maximum cladding thickness reduction, internal gas pressure, maximum cladding axial strain, maximum pellet radial strain, peak cladding temperature, peak pellet temperature, average magnetite layer thickness, minimum magnetite layer thickness, average spinel layer thickness, and minimum spinel layer thickness^{42,43}.

K2K development

By analyzing and comparing the performances of Kriging, MLP, and KAN (detailed in Supplementary Note C), and considering the characteristics of the oxygen control strategy search problem, the K2K model is constructed. Inspired by the predictor-corrector methods used in differential equation solving, the K2K model, as illustrated in Fig. S2g, employs a predictor-corrector structure. In K2K, Kriging is employed to generate initial output predictions, and KAN is used to correct the Kriging predictions. The final output of K2K is the sum of Kriging estimation value and KAN correction value. This structure enables the K2K model to combine the high fidelity of Kriging with the high capability of KAN, effectively leveraging the strengths of both components. During the training process of K2K, we employ a data trimming technique to ensure that Kriging accurately captures the overall trends while mitigating the impact of error noise introduced by dense sampling. A gradient penalty term, as defined in Eq. (1), is incorporated into the loss function to steer K2K predictions away from non-physical predictions.

$$\text{Gradient penalty} = \left| \frac{(\hat{y}_2 - \hat{y}_1) - (\tilde{y}_2 - \tilde{y}_1)}{\hat{x}_2 - \hat{x}_1} \right| \odot \text{sign} \left(\left| \frac{(\hat{y}_2 - \hat{y}_1) - (\tilde{y}_2 - \tilde{y}_1)}{\hat{y}_2 - \hat{y}_1} \right| - \nu \right) \quad (1)$$

Performance testing of K2K. To evaluate K2K's performance, comparative tests are conducted using an initial sampling dataset calculated by COMMA. K2K is benchmarked against Kriging, MLP, and KAN across multiple metrics, including model accuracy, data requirements, model capacity, training speed, hyperparameter tuning complexity, and non-physical predictions (Fig. S2). During the comparative tests, hyperparameter optimization based on NSGA-II is employed (Table. S4), aiming to reduce the influence of manual hyperparameter selection. The MSE of the comprehensive error is used to compare the accuracy of different models.

In the 50-point training and testing, the minimum Mean Squared Error (MSE) achieved by K2K is 40.3 with 95% confidence interval (CI) = [39.41, 42.40]. The initial accuracy of K2K is relatively high due to the utilization of the Kriging model, being only slightly lower than that of the original Kriging model. This is primarily attributed to the significant overfitting of KAN when the number of data points is limited. With the implementation of KAN for correction, the model error decreases rapidly as the number of data points increases. After reaching 35 data points, K2K emerges as the optimal model among the evaluated four models. At the end of training, K2K achieves a 44% and 9% reduction in error levels compared to MLP and KAN, respectively. Throughout the training process, the computational cost of K2K is approximately 2.3 times that of MLP and 1.1 times that of KAN. Owing to the rapid training efficiency of Kriging, the computational cost of K2K remains comparable to that of KAN. As for the K2K predictions, there are only <1% non-physical results (95% CI = [0.75%, 1.18%]), which are defined by any anomalous fluctuation and inconsistency in the of multiphysics parameters (Supplementary Note H). In summary, by combining the high fidelity of Kriging with the high capability of KAN through the predictor-corrector framework, K2K demonstrates superior model capacity and fidelity, enabling higher accuracy with smaller data requirements.

The advantage of K2K lies in its complementary predictor-corrector structure, which not only overcomes the limitations of Kriging accuracy but also alleviates KAN's tendency to overfit. Specifically, Kriging establishes an overall perception of the design domain, capturing the preliminary trends of parameter variations. However, it struggles to capture detailed fluctuations. By utilizing the prediction errors from Kriging to train KAN, K2K can effectively capture the complex patterns of parameter fluctuations, enabling precise correction of Kriging's errors. Benefiting from the strong capability of KAN, a

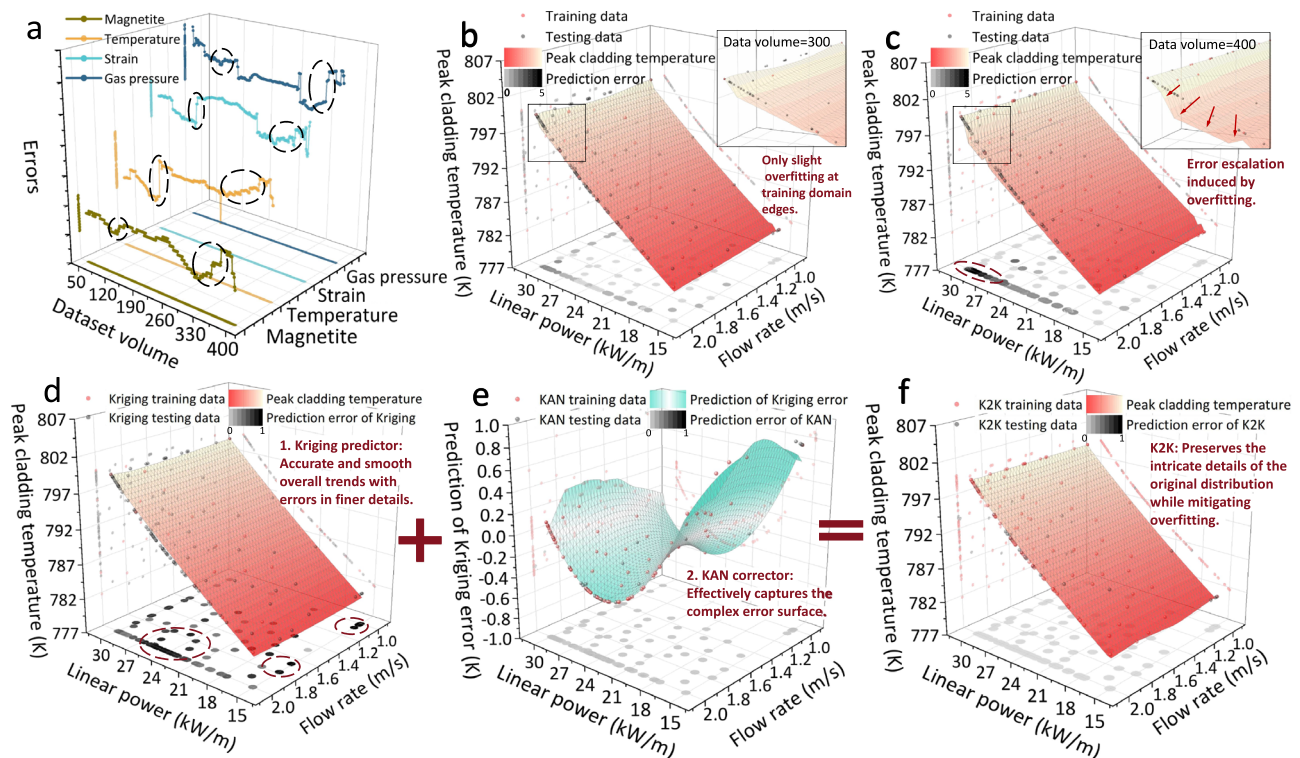


Fig. 2 | Data sensitivity of Kriging and advantages of the K2K model. **a** To clearly demonstrate the sensitivity of Kriging to data volume, Kriging is trained with dataset volume expanded incrementally from 50 to 400. Across varying prediction parameters, Kriging consistently exhibits instances of error escalation with increasing dataset volume. **b, c** The error escalation phenomenon in Kriging is further examined using peak cladding temperature datasets. As the dataset volume increases from 300 (**b**) to 400 (**c**), newly added data introduces local distortions and leads to overfitting in the surface predictions. Compared to the original Kriging model, the predictor-corrector framework and data trimming strategy in K2K mitigate these issues, enhancing both prediction accuracy and fidelity. **d–f** Taking the peak cladding temperature as an example, we demonstrate the complementary roles of Kriging and KAN in the K2K prediction process and the K2K outputs. The error distribution is illustrated in the xy-plane projection, where darker colors

indicate larger errors. **d** Initially, the Kriging predictor is trained on the original dataset. Through data trimming, a Kriging model with correct global trends and moderate overall predictive accuracy is achieved. As influenced by the true function's behavior, the error distribution exhibits localized concentrations. **e** Subsequently, the KAN corrector is trained using the error dataset. The KAN effectively captures the complex surface of error distribution, addressing the overlooked details in Kriging's predictions. **f** Finally, the combined output of the K2K model, integrating Kriging prediction and KAN correction, produces smooth functional trends with minimal errors. The resulting predictions demonstrate high fidelity, free from significant error concentrations or overfitting artifacts. In **b–f**, the coloring of the surface plots is based on the corresponding y-axis variable and shares a consistent parameter range.

small K2K configuration can achieve high accuracy without requiring a large amount of data points. Additionally, the complex structure of the error dataset can mitigate the risk of overfitting during KAN training. Figure 2 provides a more intuitive display of these advantages.

Accuracy of the trained K2K. In K2K, Kriging can provide a prior estimate of prediction errors. KAN, functioning as the correction function for Kriging, is also capable of error prediction. The self-error estimation feature of K2K facilitates guidance of adaptive sampling, enabling efficient improvements in model accuracy. Thus, based on the characteristics of K2K, we construct a two-stage adaptive sampling strategy to optimize the selection of supplementary sampling locations. As shown in Fig. 3a, after 385 iterations of adaptive sampling, the integrated MSE for the training and testing datasets decreased to 0.2 and 2.3, respectively, meeting the target accuracy of 0.2. The initial 270 iterations correspond to the first sampling stage, where the Maximum-Minimum Distance (MD) and MSE criteria are employed to guide the adaptive sampling, thereby enhancing the global predictive capability of K2K. The MD criterion focuses on the spatial filling characteristics of the data points (calculated as shown in Eq. 2), whereas the MSE criterion serves as one of the most effective guides for reducing Kriging errors. At the start of adaptive sampling, the sparsity of the initial data constrains Kriging's capacity to capture trends. Consequently, the MSE criterion dominated the early iterations of the first stage. As the

supplemental sampling in the major high-error regions concludes, the influence of the MD criterion increases. After reaching 97 iterations, the model transitions into a more stable error reduction trend with MD criterion dominated. The low decay rate of MD also indicates the effectiveness of adaptive sampling for space filling in the first stage. Thus, in the first sampling stage, error reduction follows a hyperbolic trend.

$$MD = \max_{S=1,2,\dots} \left\{ \min_{i,j=1,2,\dots} \left\{ \sqrt{(\vec{x}(i) - \vec{x}(j))^2} \right\} \right\} \quad (2)$$

The remaining 115 iterations correspond to the second sampling stage, where the sampling strategy is based on the Maximum predicted Error (ME) criterion and an improved Expected Improvement (EI) criterion. The ME and EI criterion are chosen to guide the adaptive sampling, thereby enhancing the prediction accuracy of K2K in high-complexity local regions. As illustrated in Fig. 3, the additional sampling points in this stage are concentrated in regions with low oxygen concentration or low flow rates, where cladding corrosion behavior exhibited heightened complexity. The increased dataset size induces evolutions in the hyperparameter search outcomes for the optimal KAN structure (from 3-4-10 to 3-2-4-10), leading to a more rapid decline in MSE during the later iterations of the second stage. Meanwhile,




Fig. 3 | Training process and error distribution of the K2K model. **a** During the two-stage adaptive sampling process employed for training the K2K model, the training and testing errors are represented by the purple and pink curves, respectively. The overall error of the K2K model exhibited a rapid decline in the early stages, followed by a gradual stabilization. The brown-yellow curve illustrates the slow decline of the MD value, indicating a well-distributed spatial arrangement of the sampling points. Based on the predefined MD value requirement, the MD value drops to 0.099 after the training dataset reaches 321 sampling points,

indicating that the adaptive sampling enters the second stage. The first and second stages of adaptive sampling are indicated by light pink and purple-pink backgrounds, respectively. At the end of the first stage, the integrated MSE for the training and testing datasets is 2.6 and 11.8, respectively. The spatial distribution of errors in the trained K2K is visualized through projections of oxygen concentration versus linear power (**b**) and oxygen concentration versus flow rate (**c**). Within the design domain, the error distribution remains uniform and does not exceed 5.0, with no discernible trends.

benefiting from the balance between global and local exploration provided by the EI criterion, the decay rate of MD remains around $7e-4$ per point in the second stage. No significant oversampling is observed during the second stage (Fig. S3b, c). By the conclusion of training, the maximum prediction error of K2K does not exceed 4.8, significantly lower than the modeling error of COMMA (Table S5). Furthermore, errors are uniformly distributed across the design domain (Fig. S3b, c). No signs of model overfitting in densely sampled regions or localized error surges resulting from insufficient data coverage are detected. Benefiting from the gradient penalty operator, within the design domain, K2K’s predictions of oxide layer thickness aligns precisely with the trends identified using COMMA (Fig. S4–S5)²⁶. Additionally, the K2K predicted variation trends in temperature, strain, and gas pressure closely align with the original coupled physical relationships, demonstrating comparable consistency (Fig. S6–S8). K2K predictions revealed no evidence of non-physical behavior within the design domain. These results affirm that the trained K2K model serves as a high-fidelity and high-accuracy surrogate for COMMA within the specified design domain.

To further assess the robustness of the trained K2K model, stratified 10-fold cross-validation and adversarial testing are conducted. During cross-validation, prediction errors across different folds exhibit minimal variation, with the maximum fluctuation in MSE_{total} not exceeding 0.033 for the training set and 0.22 for the testing set (Fig. S9). In the adversarial testing, under the maximum allowable input perturbation $\epsilon = [0.05, 0.005, 0.005]$, the resulting output deviation remains within 5% of the original prediction error of K2K (Fig. S10). Additionally, an independent validation set comprising 62 samples is generated using the EI and ME criteria in combination with an improved LHS method to examine the model’s susceptibility to overfitting. To evaluate the extrapolation capability, an extreme condition set is constructed by expanding input parameters by 10–15%. Results show that the MSE_{total} for the independent validation set and the extreme condition set are 2.1 and 3.4, respectively, indicating that the model demonstrates satisfactory generalization and certain extrapolation capability. The detailed testing process of K2K’s robustness is provided in Supplementary Note F.

Optimal concentration and cladding failure boundaries

Leveraging the high-fidelity and high-accuracy predictions of LFR fuel performance from trained K2K, oxygen concentration control boundaries in discrete form (Fig. S11) are identified through an initial rough estimation followed by precise localization, with a maximum

error of $1e-4$ lg.wt.%. These boundaries encompass cladding failure boundaries, feasible oxygen concentration ranges, and optimal oxygen concentrations. The cladding failure boundaries are identified by the conservative safety limits that incorporate both prediction errors of K2K and model errors of COMMA (Table S5). Based on cladding failure boundaries, feasible oxygen concentration ranges are established. The optimal oxygen concentration for various operating conditions is determined by iteratively approaching the maximum safety margin using a bisection method. Furthermore, oxygen concentration boundary formulas, constituting the final oxygen concentration control strategy, are derived from the K2K predictions using KAN network. The process of searching for and establishing the oxygen concentration control strategy is illustrated in Fig. S12.

Oxygen concentration control strategy. The KAN network achieves an accuracy with R^2 greater than 0.95 and MSE not exceeding $5.7e-4$, as shown in Fig.S13 and Table S6. In the 10-fold cross-validation, the maximum error fluctuation of the KAN network does not exceed $5.7e-5$, demonstrating good robustness. The six oxygen concentration boundary formulas generated are 81-term polynomials of linear power p and flow rate v (Table S7), with the highest powers of linear power and flow rate being 8 and 16, respectively. After formula structural transformation, the oxygen concentration boundary formula can be calculated as follows,

$$P = \frac{p}{10kW/m} \tag{3}$$

$$v = \left(\frac{0.5 \cdot v + 1}{m/s}\right)^2 \tag{4}$$

$$K_{4 \times 1} = D_{4 \times 3} \begin{pmatrix} P \\ v \\ 1 \end{pmatrix} \tag{5}$$

$$O_{4 \times 1} = L_{4 \times 5} \cdot \left(\begin{pmatrix} K_{4 \times 1} \\ 1 \end{pmatrix} \right)_{5 \times 1}^2 \tag{6}$$

$$C_O = \left(k \cdot \sum_{i=1}^4 m_i O_i + n \right) \text{lg.wt.}\% \tag{7}$$

Where \mathcal{P} is the non-dimensional linear power and \mathcal{V} is the non-dimensional flow rate. $\mathbb{D}_{4 \times 3}$ and $\mathbb{K}_{4 \times 1}$ are the input projection matrix and oxygen concentration factor, respectively. $\mathbb{L}_{4 \times 5}$ and $\mathbb{O}_{4 \times 1}$ are the oxygen boundary coefficient matrix and oxygen concentration coefficient, respectively. k and n are correction factors, and m_i is a sign coefficient taking values of +1 or -1. Tables S8 and S9 provide the values of $\mathbb{D}_{4 \times 3}$, $\mathbb{L}_{4 \times 5}$, k, n , and m_i for the six oxygen concentration boundaries. The detailed formula simplification and computational procedures are provided in Supplementary Note G. Equations (3–7) enable the calculation of corresponding failure modes, feasible oxygen concentration ranges, and optimal oxygen concentrations based on selected operating conditions. Thus, Eqs. (3–7) constitute an oxygen control strategy for LFR that ensures long-term cladding durability and operational safety under specified conditions.

Boundary formula analysis. In Fig. S14, an analysis of the term importance in the original 81-term polynomial form of oxygen concentration boundary formulas is presented. The oxygen concentration boundary formulas generated by K2K and KAN exhibit several characteristics consistent with physical phenomena: (1) Flow rate exhibits a higher power than linear power; however, the contribution of high-power flow rate terms is relatively small. This is due to the dominant influence of linear power, while flow rate primarily modulates more nuanced aspects of cladding behavior. (2) A comparison between Fig. S14a–c and S14d–f indicates that the dominant terms for the upper and lower oxygen concentration limits, as well as the optimal oxygen concentration, involve higher orders than those of cladding failure boundaries. This is because the upper and lower oxygen concentration limits are constrained by varying failure mechanisms across distinct operating scenarios, resulting in complex, non-linear transition boundaries between these regions. (3) The terms contributing most significantly to the boundary formulas are predominantly associated with linear power. This aligns with the greater sensitivity of fuel performance to linear power compared to flow rate, as evidenced in Table S2. These findings underscore the interpretability and physical consistency of oxygen concentration boundary formulas, which effectively capture the interplay between linear power, flow rate, and oxygen concentration boundaries.

Applicability testing of boundary formulas. The applicability range of the derived oxygen concentration boundary formulas is evaluated using COMMA. Employing the fallback stepwise validation methodology detailed in Supplementary Note D, a recommended usage range for Eqs. (3–7) has been established, ensuring prediction errors remain below ± 0.01 lg.wt.% (Table S10). In terms of application beyond the design domain, the boundary formulas demonstrate a robust tolerance of up to $\pm 6 - 7\%$ for variations in geometric parameters, such as inner pellet diameter, outer pellet diameter, gap size, and cladding thickness. In contrast, due to the direct impact of lattice pitch and coolant inlet temperature on heat transfer, the tolerance for variations in these parameters is limited to $\pm 2\%$ and $\pm 3\%$, respectively. For allowable errors of 0.03 and 0.1 lg.wt.%, the permissible variation ranges for fuel geometric parameters can reach up to 11% and 14%, respectively (Table S11). Validation through COMMA shows that the surrogate model errors affect the optimal oxygen concentration by no more than 0.5%. Equations (3–7) demonstrate strong extrapolation capability in regions adjacent to the design domain boundaries.

Patterns of fuel performance captured by K2K

Since the oxygen control strategy is derived from the outputs of K2K, which is trained without incorporating explicit physical principles, the preceding validation and discussion analysis are conducted solely from a numerical perspective. To provide further insight, we present a detailed examination of the intricate variations in fuel performance within the operating domain, based on K2K predictions. Particular

emphasis is placed on key parameters, including oxidation corrosion, cladding temperature, deformation, and gas pressure. These analyses aim to guide the selection of operating conditions that optimize the economic efficiency of LFR while ensuring safety. The variations in fuel performance as a function of operating condition is illustrated in Figs. 4 and S15. Further detailed analyses, from the perspective of fuel performance, are included in Supplementary Note E.

Parabolic influence of oxygen concentration induced by corrosion mode transition. Oxygen concentration serves as a decisive factor of cladding corrosion performance. A dual-layer oxide film, consisting of an outer magnetite layer and an inner protective spinel layer, forms on the cladding surface in oxygenated LBE. The effect of oxygen concentration manifests predominantly through its regulation of oxide layer thickness (Fig. 4a, b). As illustrated in Figs. 4c and S5, at low oxygen concentrations, the oxidation rate is insufficient to counteract the reductive removal effect of LBE, resulting in oxide layer penetration and the failure of oxidation protection. Complete coverage of the cladding substrate by the spinel layer is achieved only when oxygen concentrations reach approximately 4e-8 wt.%. With further increases in oxygen concentration, the oxygen diffusion rate, oxidation rate, and corrosion inhibition removal all intensify, transitioning the corrosion mode toward oxidation-dominated behavior. Consequently, oxide layer thickness exhibits a parabolic growth pattern with increasing oxygen concentration. The most significant influence of oxygen concentration on the oxide layer occurs in the range of 1.8e-9 to 1.0e-7 wt.%, accounting for approximately 80% of the total effect. Beyond the oxygen-dominated concentration threshold, the surface becomes oxygen-saturated, and further increases in oxygen concentration primarily affect the oxygen diffusion process, slowing the increment of oxide layer growth at higher concentrations. Due to the transition of corrosion mode from oxide removal dominated to oxidation corrosion dominated, the variation rate of oxide layer thickness becomes more pronounced at low oxygen concentrations. As a result, fuel performance parameters exhibit higher sensitivity to oxygen concentration under low oxygen conditions.

As oxide layer thickens, heat transfer efficiency declines, leading to an increase in peak cladding temperature. In turn, the rising temperature accelerates the corrosion rate, establishing a positive feedback loop. As shown in Fig. 4, peak cladding temperatures exhibit a parabolic trend as oxygen concentration increases. Under the combined effect of the heightened oxidation corrosion and enlarged thermal strain, the cladding axial strain also follows a parabolic increasing trend as oxygen concentration rises. Similarly, rising temperatures elevate gap gas temperatures, while deformation-induced gap volume reductions result in parabolic growth of gap gas pressure with increasing oxygen concentration. For instance, at a linear power of 23 kW/m and a flow rate of 1 m/s, when the oxygen concentration increases from 1.8e-9 to 4.0e-6 wt.%, the cladding peak temperatures, cladding axial strain, and gas pressure increase by 10.5 K, 12.3%, and 2.1%, respectively. However, as the oxygen concentration rises further to 1.3e-4 wt.%, the corresponding increments diminish to 1.2 K, 1.9%, and 0.5%. These findings highlight the K2K model's capability to precisely capture the parabolic influence of oxide layer growth and its cascading effects across multiple physical fields, offering valuable insights into the coupled dynamics of corrosion, heat transfer, and mechanical deformation in LFR systems.

Sensitivity of linear power impact to oxygen concentration. Linear power, as a critical determinant of total fuel heat generation, exerts its primary influence on fuel performance through its effect on temperature. As linear power increases, the peak cladding temperature rises significantly (Fig. 4d). This temperature elevation, coupled with the deterioration of heat transfer caused by oxide layer growth, amplifies cladding corrosion, leading to a mutually reinforcing effect.

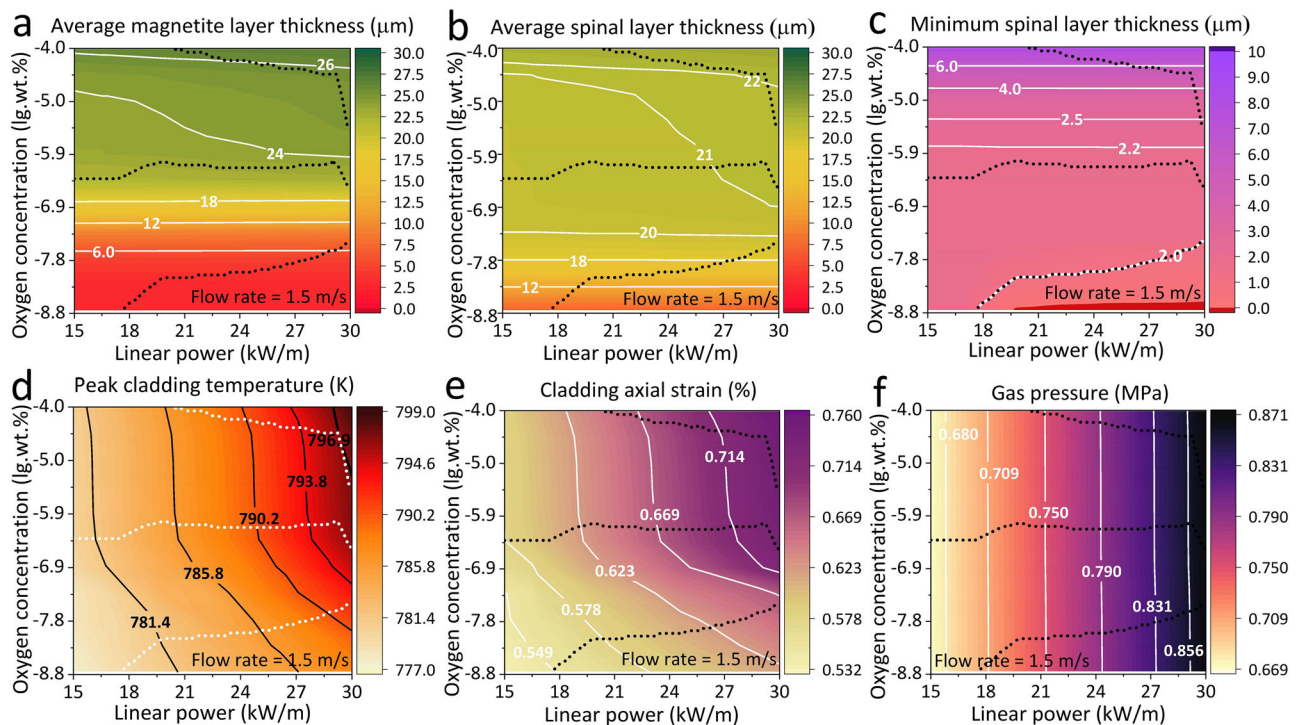


Fig. 4 | Prediction of fuel performance parameters using K2K at a flow rate of 1.5 m/s. K2K is employed to predict the variation of fuel performance parameters within the design domain as a function of linear power and oxygen concentration. The upper, lower, and middle dotted lines denote the oxygen concentration upper and lower limits determined by conservative safety limits and the optimal oxygen concentration derived for maximum safety margins, respectively. To clearly represent the range of parameter values, the contour values are not evenly spaced. **a, b** The average thicknesses of magnetite and spinel layers illustrate the evolution of cladding corrosion modes under varying operational conditions. Understanding cladding corrosion modes is critical for discerning the trend of multiphysics fuel

performance parameters. **c** The minimum spinel layer thickness serves as an indicator of oxidation protection failure, thereby defining the lower oxygen concentration limit. **d** The cladding peak temperature is related to the identification of cladding overheating, which determines the upper oxygen concentration limit. **e, f** The cladding axial strain and the gas pressure provide guidance for mitigating failure risks. The similar trends observed in **(a, b)** are attributable to the dual-layer oxidation behavior of the cladding, while the strong resemblance between **(d and e)** arises from the dominant influence of thermal deformation. These correlations highlight the high fidelity and accuracy of K2K in capturing the true physical characteristics of LFR fuel performance.

Consequently, the influence of linear power exhibits a dependence on oxygen concentration. In the low oxygen concentration range of $1.8\text{e-}8$ to $4.0\text{e-}7$ wt.%, oxide layer thickness remains relatively minimal, and the mutual amplification effect is not prominent. Within this range, cladding temperature is jointly governed by both oxygen concentration and linear power. By contrast, in the higher oxygen concentration range of $4.0\text{e-}7$ to $1.3\text{e-}4$ wt.%, oxide layer growth accelerates with increasing linear power, resulting in an approximately linear relationship between cladding temperature and linear power. As linear power rises, the oxygen-dominated concentration also increases, elevating the steady-state oxidation rate. Near the oxygen-dominated concentration, the mutual amplification effect becomes most pronounced. For instance, as linear power increases from 15 to 30 kW/m, the oxygen concentration influence zone expands by up to 21%. Furthermore, at low oxygen concentrations, the oxidation removal rate exhibits heightened sensitivity to temperature, leading to an expansion of the oxidation protection failure region (Fig. 4c). Similarly, cladding axial strain and fission gas pressure both increase with rising linear power (Fig. 4e, f). Given the strong correlation of cladding thermal deformation and fission gas release with temperature, axial strain and fission gas pressure distributions closely mirror the temperature profile as linear power increases.

Oxygen concentration dependence of flow rate impact. Coolant flow rate predominantly influences element diffusion and heat transfer, exhibiting a distinct dependence on oxygen concentration. At low oxygen concentrations, the cladding surface contains an excess of

iron, and an increased flow rate facilitates the removal of this excess, thereby enhancing the oxidation removal rate. In contrast, at high oxygen concentrations, the cladding surface becomes saturated with oxygen, and the limiting factor for oxidation rate is the element diffusion rate on the solid side. Consequently, the impact of flow rate is more pronounced at low oxygen concentrations and diminishes significantly at higher oxygen concentrations, demonstrating a clear oxygen concentration dependence (Fig. S15).

The enhanced oxidation removal rate at low oxygen concentrations also leads to a notable increase of up to 20% in the oxygen concentration required to prevent oxide layer penetration (Fig. S15c). Additionally, at low oxygen concentrations, the intensified flow rate exacerbates corrosion-induced cladding thinning, resulting in a reduction in peak cladding temperature, which also causes a slight increase in cladding axial strain and fission gas pressure due to the thermal-mechanical coupling (Fig. S15e, f).

Physical interpretation and analysis of oxygen concentration control strategy

Building upon the detailed physical interpretation of the K2K predictions, we provide an analysis regarding the trends and positions of the oxygen concentration boundary formulas within the oxygen control strategy. This analysis aims to identify the critical factors influencing boundary behavior under diverse operating conditions, enabling the detection of potential vulnerabilities in the LFR design. These insights provide a valuable reference for optimizing reactor performance and extending the operational lifespan of LFR systems.

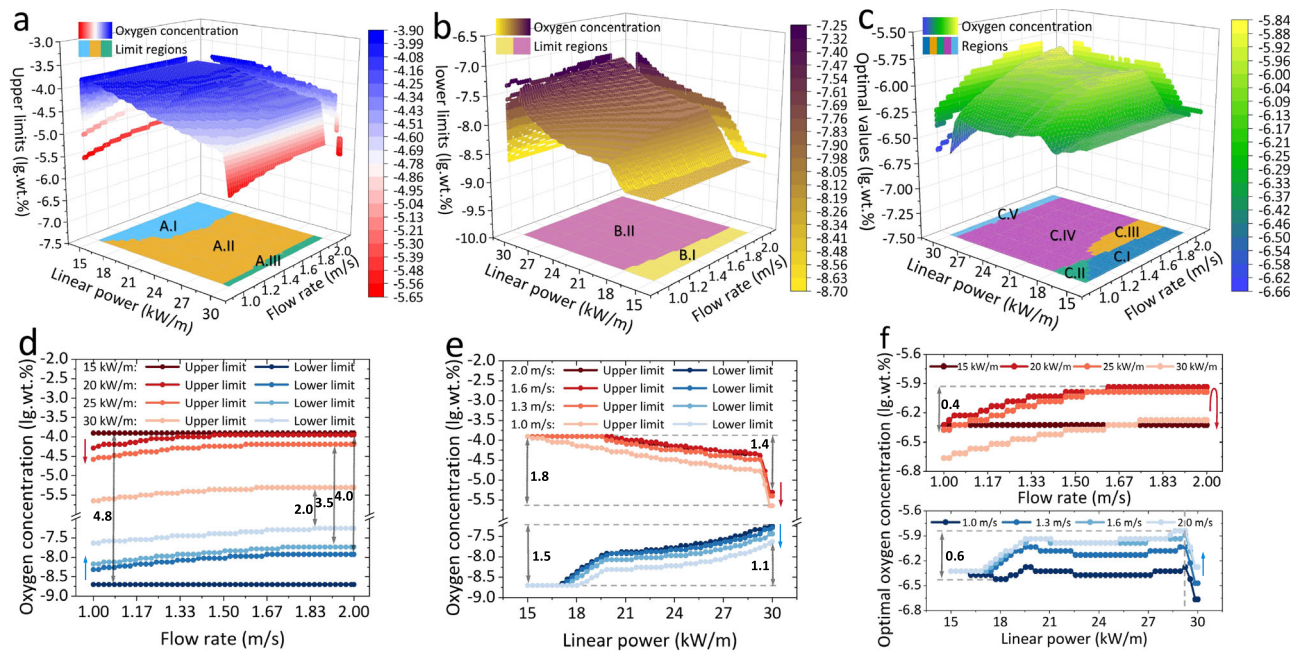


Fig. 5 | Oxygen control strategy involving oxygen concentration limits and optimal oxygen concentration under different operating conditions. **a–c** The spatial distributions of the upper limit, lower limit, and optimal oxygen concentration are illustrated using color-coded surfaces. Based on distinct constraints, the upper limit, lower limit, and optimal oxygen concentration are categorized into 3, 2, and 5 regions, respectively, as depicted in the projection on the xy-plane. **d** The influence of flow velocity on the oxygen control strategy is demonstrated using a sequence of linear power values of 15, 20, 25, and 30 kW/m. As operational conditions vary, the width of the feasible oxygen concentration range decreases from

4.8 to 2.0 lg.wt.%. The arrows in **d–f** indicate the trend of oxygen concentration boundary changes as the parameters increase. **e** The impact of linear power on the oxygen concentration limits is shown for flow velocities of 1.0, 1.3, 1.6, and 2.0 kW/m. With changes in operational conditions, the upper oxygen concentration limit decreases by a maximum of 1.8 lg.wt.%, while the lower oxygen concentration limit increases by a maximum of 1.5 lg.wt.%. **f** The variation of the optimal oxygen concentration under characteristic flow velocities and linear powers is presented. By restricting linear power and flow velocity, the width of the optimal oxygen concentration range can be constrained to only 0.4 lg.wt.%.

Cladding failure locations and trends. Cladding thickness failure occurs near saturation oxygen concentrations. As shown in Fig. S13d, a reduction in flow rate and an increase in linear power lead to elevated cladding temperatures and accelerated oxidation corrosion, thereby expanding the regions of cladding thickness failure. Oxidation protection failure emerges when the oxide layer removal rate surpasses the oxidation rate at low oxygen concentrations (Fig. S13e). A decrease in flow rate prolongs the diffusion time of elements, resulting in a linear contraction of the oxidation protection failure region with reduced flow rates. Due to the transition in corrosion modes and the parabolic nature of oxidation corrosion, the oxidation protection failure boundary exhibits a three-stage increase as linear power rises. Since higher linear power corresponds to greater heat flux, cladding overheating cases are distributed in the high-power region of 29–30 kW/m (Fig. S13f). At a linear power of 30 kW/m and a flow rate of 1.0 m/s, oxidation protection failure occurs at an oxygen concentration as low as 1.2×10^{-7} wt.%.

Reduction and upward shifting of feasible oxygen concentration range with operating conditions. As shown in Fig. 5a, the upper limits of oxygen concentration under different operating conditions can be categorized into three regions, constrained by the solubility limit of lead oxide, cladding thickness failure, and cladding overheating, respectively. With increasing linear power, cladding temperature rises, intensifying oxidation corrosion, and resulting in a progressive decline of the upper limit by up to 58% in Region A.II, followed by a rapid and steep reduction in Region A.III. Similarly, the lower oxygen concentration limits are divided into two regions, determined by oxide stability and oxidation protection failure. Influenced by flow rate and linear power, which govern element distribution and oxidation removal rates, the lower limit in Region B.II increases by approximately 2.3×10^{-8} wt.% as flow rate and linear power rise (Fig. 5b).

Consequently, the feasible oxygen concentration range narrows with increasing linear power and decreasing flow rate. As depicted in Fig. 5, operating conditions that fully satisfy the initial oxygen concentration range are confined to region C.I, characterized by low linear power and high flow rate. Enhanced flow rates improve heat transfer and oxidation removal rates, causing the oxygen concentration range to shift upward (Fig. 5d). However, higher linear power results in a contraction of the feasible range by up to 58% (Fig. 5e). Even under the design condition, the usable oxygen concentration range shrinks by approximately 22%. This deviation from the ideal scenario can be attributed to non-uniform cladding temperatures and the complex interactions between multiple physical phenomena.

Stability of optimal oxygen concentration across varying operational conditions. The optimal oxygen concentration typically lies below the oxygen-dominant region within the design domain, ranging from approximately 4.1×10^{-7} – 2.3×10^{-6} wt.%. Corresponding to the regions defined by the upper and lower limits of oxygen concentration, the distribution of optimal oxygen concentration can be categorized into four main regions (Fig. 5c). The optimal oxygen concentration exhibits a slight increase with rising flow rate and linear power (Fig. 5f). At high linear powers, the risk of cladding overheating failure sharply increases due to the positive temperature feedback from oxide layer growth under high oxygen concentrations. After the linear power rises to approximately 29 kW/m, a decrease in the optimal oxygen concentration is observed. By restricting linear power to the range of 21–27 kW/m, the optimal oxygen concentration can be maintained within the range of 3.9×10^{-7} – 1.2×10^{-6} wt.%. Considering the precision limits of current oxygen sensor technology^{44,45}, this suggests that a single optimal oxygen concentration control can effectively encompass a broad operational range, thus simplifying operational management.

Discussion

Replicating actual operating conditions and evaluating their influence on fuel performance through experimental methods presents substantial challenges. Currently, a gap exists between experimental investigations and practical applications. Although the oxygen control technology has been demonstrated to mitigate LBE corrosion under experimental conditions, the long-term performance and lifespan of LFR fuel rods under actual operating conditions remain unverified. To ensure the long-term operational safety of fuel rods, the oxygen control strategy based on the multiphysics fuel performance evaluation becomes a critical imperative for the commercialization of LFR. The necessity, complexity, and time-consuming nature of multiphysics coupling fuel performance calculations have resulted in the oxygen control strategies for operational conditions remaining in the preliminary stages of investigation. To address these challenges, we developed an oxygen control strategy for the primary operational parameter design domain of LFR based on the established K2K model.

To mitigate the high computational cost associated with multiphysics coupled fuel performance calculations for LFR and to establish a practical oxygen concentration control strategy, we propose a surrogate model, K2K. The K2K model employs a predictor-corrector framework, integrating a Kriging-based predictor with a KAN-based corrector. Comparative analyses with Kriging, MLP, and KAN algorithms underscore the superiority of K2K in addressing small-sample, high-complexity problems. To ensure physically plausible predictions, a gradient penalty operator is incorporated during model training. We implement a two-stage adaptive sampling strategy to efficiently train the K2K model. Validation against the benchmark model COMMA confirms the high fidelity and predictive accuracy of K2K.

The advantages of the K2K model primarily stem from several key features. First, the predictor-corrector architecture of K2K allows dual utilization of the sample set, significantly enhancing the efficiency of information extraction from limited data. The Kriging-based predictor provides robust stability, while the KAN-based corrector refines complex predictive details, establishing a synergistic relationship that enhances the model's expressiveness. Second, the K2K model remains a small-scale model. By reducing the number of KAN nodes, K2K avoids divergence and overfitting caused by an excessive number of trainable parameters. Third, although there are certain requirements for datasets' features during usage, the gradient penalty term captures the trend information of physical fields without the need for explicit physical information, further enhancing the model's fidelity. To adopt the K2K algorithm as a general-purpose approach, several aspects of its construction can be improved. First, the KAN corrector introduces substantial computational overhead during training, which limits the overall training speed. Developing a more lightweight yet equally expressive corrector could alleviate this issue. Second, the Kriging model demonstrates limited applicability to certain specialized datasets, such as those with extreme fluctuations. Adapting the choice of the predictor to align with dataset characteristics could further enhance predictive accuracy.

Leveraging the trained K2K model, we identify and localize the potential failure modes and regions within the operation domain, including oxidation protection failure, cladding thickness failure, and cladding overheating. Furthermore, we define the feasible oxygen concentration range and the optimal oxygen concentration under different operational conditions. Interpretable oxygen concentration boundary formulas are derived using the KAN network. Utilizing these formulas, the corresponding optimal oxygen concentration for selected operational conditions can be computed, thereby enhancing the long-term safety margin for LFR operations. In situations where maintaining stability at the optimal oxygen concentration is challenging, keeping the oxygen concentration within the feasible range effectively prevents fuel failure and reduces the risk of nuclear

accidents. This study offers an effective oxygen control strategy for the safe and reliable operation of LFR systems.

Finally, by integrating actual physical processes, we conduct a comprehensive analysis of fuel performance and oxygen control strategies from the perspectives of corrosion, temperature, deformation, and gas pressure. The results indicate that the influence of oxygen concentration follows a parabolic trend, suggesting a self-limiting nature. This inherent characteristic provides a safety buffer against unstable oxygen control during the initial stages of LFR operation. The analysis further highlights the synergistic interaction between linear power and oxygen concentration, demonstrating that the impact of linear power on fuel performance intensifies at higher oxygen concentrations. As linear power and flow rate increase, the feasible oxygen concentration range contracts up to 54%. At low oxygen concentrations, the impact of flow rate on fuel performance becomes more pronounced. Additionally, higher flow rates shift the feasible oxygen concentration range towards elevated oxygen levels.

Under practical operating conditions, the width of the feasible oxygen concentration range reflects the operational flexibility and safety margin of the reactor. The above analysis suggests that, by strategically placing oxygen measurement points and developing more stable oxygen concentration control technologies, the long-term operational safety of LFR can be markedly improved. Moreover, while the feasible oxygen concentration range exhibits significant variation across different operating conditions, the optimal oxygen concentration remains relatively stable. This stability within the design domain suggests that long-term safe operation of LFR can be achieved with minimal operational intervention through effective oxygen concentration control and operating condition design. The proposed oxygen concentration control strategy provides detailed references for the design and operation of LFR.

Methods

Multi-physical coupling analysis program COMMA

The multi-physical coupling analysis program COMMA primarily consists of three modules, comprising the fuel performance solver, the 2D steady-state LBE flow and mass transfer model named RACER, and the cladding corrosion module²⁶. The fuel performance solver is a 1.5D steady-state fuel performance calculation program, incorporating multiple solvers for neutron physics, temperature field, mechanical deformation, and fission gas release. In COMMA, to solve the coolant temperature distribution, flow state, and changes in elemental concentrations, the 2D steady-state LBE flow and mass transfer solver program, RACER, is selected. As for the cladding corrosion module, a parabolic growth model based on available space model and oxide removal rate that aligns well with experimental studies has been selected^{14,25,46,47}.

As shown in Fig. S16a, following the main computational flow of COMMA, RACER, and the cladding corrosion module are integrated into the calculation process of fuel performance. Multi-module connections are achieved through internal parameter transfers, and multi-physics convergence is obtained through multi-level iterations. In COMMA, T91 ferritic-martensitic steel is selected as cladding material for analysis, owing to its favorable high-temperature properties and demonstrated commercial applicability. COMMA has been rigorously validated against material corrosion experiments, including COLIMESTA, LINCE, CICLAD, CORRIDA, and CIEMAT, as well as fuel performance results calculated by FEMAXI-SCI-1. The relative errors remain below 3%, confirming the reliability of COMMA as a reference standard for this study¹⁹. The detailed calculations of COMMA can be referred to Supplementary Note A. The reference case for oxygen concentration strategy search is provided in Table S1, derived from typical LFR, such as ELSY and CLEAR^{6,40,41}.

Training of the K2K model

The K2K training process is systematically divided into three stages, including initial dataset construction, initial K2K development, and K2K accuracy enhancement. Firstly, sampling of initial dataset is performed using the Latin Hypercube Sampling (LHS) method, enhanced by MD criterion (Eq. 2) to avoid unreasonable LHS (Fig. S17). Due to the oxygen concentration varying across multiple orders of magnitude, the oxygen concentration is pre-processed using a logarithmic transformation with base 10.

The initial training and testing datasets consist of 50 and 30 cases, respectively. The multiphysics fuel performance at each sampling data point is obtained using COMMA. Then, the initial K2K model is trained on the initial dataset, with 50 rounds of hyperparameter optimization conducted to determine the optimal model configuration. Given the limited size of the initial dataset, additional sampling is necessary to further improve the prediction accuracy and fidelity of K2K. To this end, a two-stage adaptive sampling strategy is devised to expand the training dataset and optimize the selection of sampling locations, thereby accelerating improvements in model accuracy. In subsequent training following each adaptive sampling, 20 rounds of hyperparameter optimization are performed to update the model parameters. As the training dataset expands, the testing dataset is correspondingly augmented. Specifically, after every ten adaptive sampling iterations for the training dataset, the testing dataset is expanded once through random search guided by the MD criterion.

During the above training process, data trimming techniques are also employed to enhance training stability. Additionally, a gradient penalty term, as defined in Eq. 1, is incorporated into the loss function to steer K2K predictions away from non-physical solutions. A detailed description of the key techniques utilized during the K2K model training is provided below.

MD criterion improved LHS. For the selected three-dimensional design domain, an initial dataset is generated using an improved LHS method based on the MD criterion⁴⁸. LHS is a classic random multi-dimensional stratified sampling method⁴⁹. It effectively avoids the clustering or duplication issues that can arise with traditional Monte Carlo methods. However, LHS does not account for the spatial filling characteristics of design points. To address this limitation, the MD criterion is introduced to evaluate the spatial filling quality of the LHS datasets⁵⁰. The initial dataset is obtained using a random search strategy aimed at maximizing the spatial filling.

Gradient penalty operator. Methods to avoid non-physical predictions typically require the direct incorporation of physical information into the model training process. Notable advancements in this area have been made through algorithms such as PINN, DeepONet, and KAN^{30–32}. However, this fusion of physics and machine intelligence necessitates highly specialized expertise, complex network modeling, and considerable training costs. In these physics-informed neural network computations, the underlying function is essentially serving as a highly generalizable solver for differential equations. Inspired by the differential equation solver and the concept of redesigning loss functions of PINN, the gradient penalty operator, as expressed in Eq. (1), has been proposed to penalize the loss function. In Eq. (1), operator \odot denotes the pointwise multiplication of matrices. By inputting two sets of data values and predicted values, the model's gradient prediction error is obtained. The allowable gradient error is controlled via a relaxation factor γ , and penalties are applied when the gradient error exceeds this threshold. In the training process of K2K, the value of γ ranges from 0.2 to 0.4 and decreases as the data volume increases. The combinations of x_1 and x_2 are selected through random generation. Under well-distributed data conditions, the gradient penalty operator exhibits low sensitivity to the selection strategy of x_1 and x_2 . In the isolation tests, the use of the gradient penalty operator results in only

approximately 0.14 s of additional training time per epoch. The intention behind the gradient penalty operator is to enhance the fidelity of existing models' predictions and reduce non-physical outcomes, without relying on expert knowledge or significantly increasing training costs.

It is noteworthy that the gradient penalty operator itself imposes no constraints on the function's monotonicity or the distance between two input values. Thus, during its application, the selection of data points should be adjusted based on the problem's characteristics. For highly nonlinear problems, the effectiveness of this penalty term may be limited. In the K2K training process, the loss value calculation is performed using a validation set composed of pairs of data, where, for each data point, the nearest five data points are selected.

Data trimming. The interpolation nature of Kriging renders it highly sensitive to both data distribution and quality. As depicted in Fig. 2a–c, within high-dimensional input spaces, an increase in data points can exacerbate errors and induce non-physical hyperplane distortions. To preserve the global predictive accuracy of the Kriging model, a data trimming strategy is employed during K2K training. Data trimming regulates the volume of data used for Kriging training, mitigating overfitting and ensuring robust performance. The information from data excluded during trimming is leveraged by the KAN corrector, which captures additional intricate details. The extent of data trimming emerges as a crucial hyperparameter in the K2K training process.

Two-stage adaptive sampling strategy. In the adaptive sampling process, the strategy for selecting the locations of new sampling points is a decisive factor for the rate of error reduction (Fig. S16d). Based on the estimation-correction characteristics of the K2K model, a two-stage adaptive sampling method has been designed.

In the first stage, the priority is to enhance the spatial breadth of the data points, thereby improving the global predictive capability of the K2K model. This approach quickly leverages the predictive potential of the Kriging model, leading to a rapid decrease in the overall error of the K2K model. Consequently, the MD criterion and MSE criterion are employed to evaluate the quality of the sampling points during the first stage.

The MD criterion is a prior sampling criterion based solely on the spatial filling of data points⁴⁸. The MD criterion can evaluate the spatial distribution quality of the dataset without relying on the accuracy of outputs. However, due to its disregard for output trend information, the MD criterion may lead to sparse sampling in regions with concentrated data variation. The MSE criterion leverages the characteristic of Kriging to provide prior assessments of prediction errors, thereby enabling sampling at locations with larger prediction errors³³. The MSE for prediction value is,

$$MSE(\hat{y}(x)) = \sigma^2 (1 - r^T R^{-1} r + (1 - F^T R^{-1} r)^2 / F^T R^{-1} F) \quad (8)$$

The MSE criterion is one of the most effective criteria for rapidly improving the accuracy of Kriging. By combining the MD and MSE criterion through the multi-objective optimization algorithm NSGA-II, the sampling location search can rapidly enhance the exploration breadth and overall accuracy of the K2K model during the first stage.

Due to the limitations of Kriging's expressive capacity, once the design point filling becomes dense or the model error has reached the level of the original COMMA model, it becomes difficult to further improve Kriging's prediction accuracy. Therefore, when the MD value of the dataset decreases to 0.1 or the testing dataset error MSE_{total} decreases to 10.0, the adaptive sampling process transitions into the second stage.

In the second stage, the focus is on enhancing the K2K model's ability to predict local details, allowing it to closely replicate the

original multi-physical model COMMA. Therefore, the sampling points' fitness is evaluated using the ME criterion and an improved EI criterion during the second stage of sampling. By leveraging KAN's expressive potential, the predictive accuracy of the K2K model is further improved.

The ME criterion utilizes the estimation-correction structure of the K2K model to identify the model's weaknesses and guide sampling. Due to the predictor-corrector structure of K2K, locations with high predicted correlation value from KAN correspond to areas where Kriging's expressive capability is insufficient. Furthermore, because of the nature of the Kriging model, locations with significant prediction errors are more likely to exhibit complex parameter fluctuations or steep changes. Leveraging these characteristics, the ME criterion is established as follows,

$$ME(x) = y_{KAN}(x) \tag{9}$$

The EI criterion, based on prior predictive errors, guides sampling by evaluating the expected improvement in model accuracy from adding new points⁵¹. However, the optimal values comprise a surface of optimal oxygen concentrations under varying conditions rather than a single value. Thus, the formula for assessing improvement, $I(x)$, needs to be reconstructed. Assuming the true value $y(x)$ at point x is known, the error improvement brought by the $y(x)$ at point x can be expressed as,

$$I(x) = y(x) - \hat{y}(x) \approx y_{KAN}(x) \tag{10}$$

And it can be approximated that the correlation value estimated by the KAN model, $y_{KAN}(x)$, is roughly equal to the true error. Similarly, the modified EI criterion can be expressed as follows,

$$EI = y_{KAN} \Phi\left(\frac{y_{KAN}}{s(x)}\right) + s(x) \phi\left(\frac{y_{KAN}}{s(x)}\right) \tag{11}$$

The EI criterion serves as a sampling strategy that balances both global and local considerations, thereby preventing an excessive concentration of sampling points during the second stage. Based on the inherent errors of the COMMA model, the K2K model achieves the required accuracy and reaches the end of training process when training dataset error MSE_{total} falls to 0.2. Additionally, the effectiveness of the proposed two-stage adaptive sampling is tested based on test functions as shown in Supplementary Note C.

Construction of oxygen concentration control strategy based on K2K and KAN

As shown in Fig. S12, the construction of oxygen concentration control strategy process comprises four main steps: rough estimation, precise localization, boundary formula generation, and applicability test.

Evaluation criteria. As shown in Table S5a conservative safety limit is established based on safety design limits of LFR, along with error margins from both COMMA and K2K models. This threshold is employed in subsequent evaluations to ensure that the derived oxygen concentration control strategy robustly supports the long-term operational safety of the LFR. The feasibility of specific design points is assessed relative to this conservative safety limit. Any performance parameter exceeding these limits is considered indicative of fuel failure. Specifically, the failure boundaries for oxygen concentration are determined by comparing fuel performance parameters from individual physical fields with their corresponding conservative safety limits. The upper and lower oxygen concentrations correspond to the boundaries of the interval where no failures occur in any physical field. The optimal oxygen concentration under given operating conditions is determined by maximizing the integrated safety margin defined by Eq.

(12), thereby enhancing the reliability of the proposed control strategy.

$$\Delta_{total} = \frac{\Delta_{mag} \times \Delta_{sp} \times \Delta_{clad}}{1 \mu m} \times \frac{\Delta_{Tco} \times \Delta_{Tfi}}{10 K} \times \frac{\Delta_{ec} \times \Delta_{ef}}{0.1 \%} \times \frac{\Delta_p}{0.01 MPa} \tag{12}$$

Where Δ_{mag} , Δ_{sp} , and Δ_{clad} represent the safety margins associated with magnetite layer thickness, spinel layer thickness, and cladding substrate thickness, respectively. Δ_{Tco} and Δ_{Tfi} denote the safety margins of peak cladding temperature and peak pellet temperature, respectively. Δ_{ec} and Δ_{ef} correspond to the safety margins of maximum cladding axial strain and maximum pellet radial strain, respectively, while Δ_p represents the safety margin related to the fission gas pressure.

Rough estimation. A dense sampling of the three-dimensional design domain, defined by oxygen concentration, linear power, and flow rate, is performed using a $50 \times 50 \times 100$ grid and validated using K2K. A total of 52,932 instances of fuel failure are identified. The causes of cladding failure include reductions in cladding thickness (15,777 cases), oxidation protection failure (35,084 cases), and cladding overheating (233 cases). As shown in Fig. S11, failures related to cladding thickness are predominantly observed in regions with high oxygen concentrations, while oxidation protection failures primarily occur at low oxygen concentrations. Cladding overheating is concentrated in regions of high linear power. The failure regions expand with increasing linear power and decreasing coolant flow rate. The optimal oxygen concentration is approximately situated just below the oxygen-dominated concentration. This rough estimation provides a preliminary identification of potential failure scenarios and their spatial distributions within the operational domain.

Precise localization. Based on the rough estimation dataset, an approximate oxygen concentration boundary with a resolution of -0.044 lg.wt.% is obtained. To achieve a maximum allowable error of $1e-4$ lg.wt.%, further refinement and additional sampling are conducted near these boundary points using a bisection method. Through step-wise localization across 50×50 (linear power \times flow rate) conditions, a total of 7733 precise oxygen concentration boundary points are identified.

Boundary formula generation. Due to the complex interplay of multifactorial influences and nonlinear interactions within multiphysics coupling, deriving reasonable and effective formula structures based on human intuition is challenging. As an alternative to empirical formula structures, oxygen concentration boundary formulas are generated through the symbolic training of KAN. The inherent structure of KAN features both common and independent nodes. Thus, this characteristic enables the generated oxygen concentration boundary formulas to incorporate shared components while retaining independent elements, which align with the multi-physical characteristics of the oxygen control strategy. Consequently, predefined formula structures and manual construction are supplanted by symbolic training through the KAN network, facilitating the development of high-accuracy oxygen concentration boundary formulas. Using the oxygen concentration boundary dataset, a KAN network is designed and trained to generate the corresponding boundary formulas.

The inputs to the KAN include linear power and flow rate, while the outputs encompass the upper and lower oxygen concentration limits, the optimal oxygen concentration, and the oxygen concentration boundaries for cladding thickness failure, oxidation protection failure, and cladding overheating. In the symbolic training of KAN, the operator set selected for symbolic training consists of $[x, x^2, \sin, \tanh]$. The detailed symbolic training and term selection procedures are provided in Supplementary Note G. To improve the fit accuracy of the

oxygen concentration boundary formulas, hyperparameter optimization of the KAN network structure and training parameters is conducted using the NSGA-II algorithm. To enhance the fidelity of KAN trained formulas, the gradient penalty operator is also applied. The multi-objective optimization evaluation metric is the goodness of fit, represented by R^2 , for different boundaries. The optimization process concludes when the R^2 values for all fits reach or exceed 0.95, resulting in the final output of the oxygen concentration boundary formulas.

Applicability test. To explore the applicability of the generated formulas, the design parameters are appropriately extended in the vicinity of the design domain of the K2K model. Six primary design parameters related to fuel rods, including core inlet temperature, gap size, pellet outer diameter, pellet inner diameter, lattice pitch, and initial cladding thickness, are selected for testing. During the testing process, fuel performance data for the extended design points are obtained using the COMMA model.

Conducting extensive calculations of fuel performance in a six-dimensional design space using COMMA entails significant time costs. To reduce the time required for searching the applicability range, the testing process follows a structured approach. The applicability range testing comprises three main steps, which include design parameter variation, design parameter validation, and design parameter rollback. At the beginning of the applicability range testing, design parameter variations are initiated, followed by design parameter validation. If a set of design parameters fails to pass the validation, the design parameter rollback and design parameter validation are performed repeatedly until the renewed design parameter successfully meets the criteria. The process described above is repeated for five iterations. The design parameter exhibiting the maximum variation magnitude that satisfies the validation criteria is considered the applicability range of the formulas. It is important to note that the choice of five iterations is based on computational cost considerations, which means the maximum design parameter variation is limited to only 5%. Consequently, the applicability range obtained through this search serves as a reference range, while the actual applicability range may be larger.

Data availability

All other data supporting the findings of this study are available in the Supplementary Information file. Source data are provided in this paper. Source data are provided with this paper.

Code availability

The core code of K2K model is available via GitHub at <https://github.com/wangshiwei20000711/K2K.git>. Shiwei Wang, Jijie Chen, Xiaojing Liu, Tengfei Zhang, Xiang Chai, Qi Lu, Danhong Shen, Hui He, Kriging to KAN model accelerated discovery of oxygen control strategy in lead-based fast reactors, K2K, <https://doi.org/10.5281/zenodo.16040458>, 2025.

References

- Brown, P. T. & Caldeira, K. Greater future global warming inferred from Earth's recent energy budget. *Nature* **552**, 45–50 (2017).
- Geng, T. et al. Increased occurrences of consecutive La Niña events under global warming. *Nature* **619**, 774–781 (2023).
- Price, J., Keppo, I. & Dodds, P. E. The role of new nuclear power in the UK's net-zero emissions energy system. *Energy* **262**, 125450 (2023).
- Bodel, W., Butler, G. & Matthews, J. Nuclear energy for net zero: a strategy for action. *Dist. Heat.* **1**, 2 (2021).
- Nevitt, M. Assessing COP28: The New Global Climate Deal in Dubai. *Just Security* (2023).
- Smith, C. F. & Cinotti, L. In *Handbook of Generation IV Nuclear Reactors (Second Edition)* (ed Igor L. Pioro) 195–230 (Woodhead Publishing, 2023).
- Dittmar, M. Nuclear energy: status and future limitations. *Energy* **37**, 35–40 (2012).
- Locatelli, G., Mancini, M. & Todeschini, N. Generation IV nuclear reactors: current status and future prospects. *Energy Policy* **61**, 1503–1520 (2013).
- Buckthorpe, D. in *Structural Materials for Generation IV Nuclear Reactors 1-22* (Elsevier, 2017).
- Alemberti, A. Lead cooled fast reactors. *Reference Module in Earth Systems and Environmental Sciences*, 523–544 (2021).
- Fazio, C. et al. *Handbook on Lead-bismuth Eutectic Alloy and Lead Properties, Materials Compatibility, Thermal-hydraulics, and Technologies-2015 edition* (Organisation for Economic Co-Operation and Development, 2015).
- Gong, X., Short, M. P., Auger, T., Charalampopoulou, E. & Lambrinou, K. Environmental degradation of structural materials in liquid lead-and lead-bismuth eutectic-cooled reactors. *Prog. Mater. Sci.* **126**, 100920 (2022).
- Xie, A. et al. Microstructural evolution of the inner oxide layer of a Si-modified austenitic stainless steel exposed to oxygen-saturated lead-bismuth eutectic (LBE): Toward the origin of LBE penetration. *Corros. Sci.* **237**, 112350 (2024).
- Martinelli, L. & Balbaud-Célérier, F. Modelling of the oxide scale formation on Fe-Cr steel during exposure in liquid lead-bismuth eutectic in the 450–600 °C temperature range. *Mater. Corros.* **62**, 531–542 (2011).
- Luo, L. et al. Honeycomb structure with oxygen-poor pores at the top of magnetite layer on a martensitic steel CLAM exposed to lead-bismuth eutectic at 500 °C. *Corros. Sci.* **204**, 110410 (2022).
- Shi, H. et al. Influence of Si addition on the corrosion behavior of 9 wt% Cr ferritic/martensitic steels exposed to oxygen-controlled molten Pb-Bi eutectic at 550 and 600 °C. *Corros. Sci.* **193**, 109871 (2021).
- Ma, Z. et al. Influence of temperature on the microstructural evolution of SIMP and T91 steels in liquid lead-bismuth eutectic: Experiments and molecular dynamics simulations. *Corros. Sci.* **222**, 111411 (2023).
- Xiao, J. et al. A refined oxidation mechanism proposed for ferritic-martensitic steels exposed to oxygen-saturated liquid lead-bismuth eutectic at 400° C for 500 h. *J. Nucl. Mater.* **549**, 152852 (2021).
- Shahboub, A., Chen, P., Deng, C., Ding, W. & Zheng, M. Investigation of the anti-corrosion behavior of Fe-Cr oxide layer in LBE: a first-principles study. *J. Nucl. Mater.* **605**, 155592 (2025).
- Li, D. et al. The influence of liquid Pb-Bi on the anti-corrosion behavior of Fe 3 O 4: a first-principles study. *Phys. Chem. Chem. Phys.* **18**, 7789–7796 (2016).
- Zhu, Z. et al. Corrosion behavior of T91 steel in liquid lead-bismuth eutectic at 550° C: effects of exposure time and dissolved oxygen concentration. *Corros. Sci.* **204**, 110405 (2022).
- Maia, E. L. et al. Oxygen-enhanced dissolution corrosion of austenitic stainless steel 316 L in static lead-bismuth eutectic at 500°C. *Corros. Sci.* **236**, 112224 (2024).
- Košek, L. et al. Mechanism of localized corrosion issues of austenitic steels exposed to flowing lead with 10–7 wt.% oxygen at 480 °C up to 16,000 h. *J. Nucl. Mater.* **572**, 154045 (2022).
- Wang, W., Yang, C., You, Y. & Yin, H. A review of corrosion behavior of structural steel in liquid lead–bismuth eutectic. *Crystals* **13**, 968 (2023).
- Feng, W., Zhang, X., Cao, L., Wang, S. & Chen, H. Development of oxygen/corrosion product mass transfer model and oxidation-corrosion model in the lead-alloy cooled reactor core. *Corros. Sci.* **190**, 109708 (2021).
- Wang, S. et al. Multi-physical effects induced by spatiotemporal corrosion of T91 cladding in oxygen-controlled LBE environment. *Corros. Sci.* **233**, 112075 (2024).

27. Ji, X., Liu, X., Zhang, T. & He, H. Numerical investigation on the Multi-Physics coupling characteristics of Lead-Bismuth reactor fuel assembly based on oxidation corrosion behaviors. *Nucl. Eng. Des.* **428**, 113519 (2024).
28. LeCun, Y., Bengio, Y. & Hinton, G. Deep learning. *Nature* **521**, 436–444 (2015).
29. Koldasbayeva, D. et al. Challenges in data-driven geospatial modeling for environmental research and practice. *Nat. Commun.* **15**, 10700 (2024).
30. Raissi, M., Perdikaris, P. & Karniadakis, G. E. Physics-informed neural networks: a deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *J. Comput. Phys.* **378**, 686–707 (2019).
31. Lu, L., Jin, P., Pang, G., Zhang, Z. & Karniadakis, G. E. Learning nonlinear operators via DeepONet based on the universal approximation theorem of operators. *Nat. Mach. Intell.* **3**, 218–229 (2021).
32. Karniadakis, G. E. et al. Physics-informed machine learning. *Nat. Rev. Phys.* **3**, 422–440 (2021).
33. Krige, D. G. A statistical approach to some basic mine valuation problems on the Witwatersrand. *J. South. Afr. Inst. Min. Metall.* **52**, 119–139 (1951).
34. Diaw, A., McKerns, M., Sagert, I., Stanton, L. & Murillo, M. S. Efficient learning of accurate surrogates for simulations of complex systems. *Nat. Machine Intell.* 1–10 <https://doi.org/10.48550/arXiv.2207.12855> (2024).
35. Rumelhart, D. E., Hinton, G. E. & Williams, R. J. Learning representations by back-propagating errors. *Nature* **323**, 533–536 (1986).
36. Liu, Z. et al. Kan: kolmogorov-arnold networks. *arXiv preprint at arXiv:2404.19756* (2024).
37. Yang, Q. et al. Machine learning assisted prediction for hydrogen production of advanced photovoltaic technologies. *DeCarbon* **4**, 100050 (2024).
38. Sun, W. et al. Artificial intelligence designer for highly-efficient organic photovoltaic materials. *J. Phys. Chem. Lett.* **12**, 8847–8854 (2021).
39. Sun, W. et al. Machine learning-assisted molecular design and efficiency prediction for high-performance organic photovoltaic materials. *Sci. Adv.* **5**, eaay4275 (2019).
40. Alemberti, A. et al. European lead fast reactor—ELSY. *Nucl. Eng. Des.* **241**, 3470–3480 (2011).
41. Wu, Y. Design and R&D progress of China lead-based reactor for ADS research facility. *Engineering* **2**, 124–131 (2016).
42. Luzzi, L. et al. Application of the TRANSURANUS code for the fuel pin design process of the ALFRED reactor. *Nucl. Eng. Des.* **277**, 173–187 (2014).
43. Wang, G., Wang, Z. & Yun, D. Cladding failure modelling for lead-based fast reactors: a review and prospects. *Metals* **13**, 1524 (2023).
44. Adhi, P. M., Kondo, M. & Takahashi, M. Performance of solid electrolyte oxygen sensor with solid and liquid reference electrode for liquid metal. *Sens. Actuators B Chem.* **241**, 1261–1269 (2017).
45. Lim, J. et al. Control of dissolved oxygen in liquid LBE by electrochemical oxygen pumping. *Sens. Actuators B Chem.* **204**, 388–392 (2014).
46. Zhang, J., Li, N., Chen, Y. & Rusanov, A. Corrosion behaviors of US steels in flowing lead–bismuth eutectic (LBE). *J. Nucl. Mater.* **336**, 1–10 (2005).
47. Zhang, J. & Li, N. Oxidation mechanism of steels in liquid–lead alloys. *Oxid. Met.* **63**, 353–381 (2005).
48. Duarte, F. L. & de Lamare, R. C. Switched max-link relay selection based on maximum minimum distance for cooperative MIMO systems. *IEEE Trans. Vehicular Technol.* **69**, 1928–1941 (2019).
49. Huntington, D. & Lyrantzis, C. Improvements to and limitations of Latin hypercube sampling. *Probabilistic Eng. Mech.* **13**, 245–253 (1998).
50. Joseph, V. R. & Hung, Y. Orthogonal-maximin Latin hypercube designs. *Statistica Sinica* **18**, 171–186 (2008).
51. Jones, D. R., Schonlau, M. & Welch, W. J. Efficient global optimization of expensive black-box functions. *J. Glob. Optim.* **13**, 455–492 (1998).

Acknowledgements

The authors are grateful for the support of Natural Science Foundation of China (No: 12427811(X.L.), 52176082(H.H.), 12275261(Q.L.)), and the authors also want to appreciate the LingChuang Research Project of China National Nuclear Corporation (H.H.) and “Shu Guang” project supported by Shanghai Municipal Education Commission (X.L.) and Shanghai Education Development Foundation (21SG13(H.H.)). In addition, the authors are especially grateful to Professor Jinbiao Xiong for his guidance and advice.

Author contributions

S.W., J.C. and H.H. conceived the study. S.W., J.C., X.L., T.Z., X.C., Q.L., D.S. and H.H. designed the study. S.W. and J.C. developed the program. S.W., J.C., X.L. and H.H. performed the analyses. X.L., T.Z., X.C., Q.L. and D.S. supervised the analyses. S.W. and H.H. wrote the paper with input from all authors.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary information The online version contains supplementary material available at <https://doi.org/10.1038/s41467-025-64747-7>.

Correspondence and requests for materials should be addressed to Xiaojing Liu or Hui He.

Peer review information *Nature Communications* thanks Kuan Sun and the other anonymous reviewer(s) for their contribution to the peer review of this work. A peer review file is available.

Reprints and permissions information is available at <http://www.nature.com/reprints>

Publisher’s note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Open Access This article is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License, which permits any non-commercial use, sharing, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if you modified the licensed material. You do not have permission under this licence to share adapted material derived from this article or parts of it. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by-nc-nd/4.0/>.

© The Author(s) 2025