Advances in Science and Technology Research Journal

Advances in Science and Technology Research Journal, 2027, 21(1), 241–257
https://doi.org/10.12913/22998624/231656
ISSN 2299-8624, License CC-BY 4.0

Received: 2026.07.20
Accepted: 2026.11.15
Published: 2026.12.01

Custom recurrent-based friction prediction in a wear-resistant hardfacing alloy

Krzysztof Dziedzic[1*], Marcin Barszcz[1], Mychajło Paszeczko[2], Maria Skublewska-Paszkowska[1], Paweł Powroźnik[1]

[1] Faculty of Electrical Engineering and Computer Science, Lublin University of Technology, ul. Nadbystrzycka 36B, 20-618 Lublin, Poland
[2] Faculty of Mathematics and Information Technology, Lublin University of Technology, ul. Nadbystrzycka 38, 20-618 Lublin, Poland
* Corresponding author's e-mail: k.dziedzic@pollub.pl

ABSTRACT

Predicting tribological behaviour in multi-component wear-resistant alloys is challenging due to their complex microstructures and nonlinear phenomena in the contact zone. The Fe–Mn–C–B–Si–Ni–Cr coating combines hard carbide and boride phases with a ductile austenitic matrix. A custom recurrent neural network was developed to predict its friction behaviour. Tribological tests over a sliding distance of 2000 m on a high-temperature ball-on-disc tribometer (THT 1000) gave a mean coefficient of friction of 0.739 ± 0.105 at ambient temperature and 0.600 ± 0.077 at 300 °C. Because pointwise analysis could not capture the non-stationary nature of the friction signal, the recorded signals, totalling about 390,000 samples, were processed using a moving-window technique to extract statistical descriptors. A custom two-branch recurrent network, TriboRNN, was developed. The recurrent branch comprises a bidirectional gated recurrent unit followed by a long short-term memory layer with additive attention. It processes the extracted descriptors, whereas a dense branch encodes alloy characteristics and tribological test parameters. Averaged over five random initialisations, the model achieved a pooled coefficient of determination of 0.9113 ± 0.0005 and a root-mean-square error of 0.0312 ± 0.0001 on the withheld test runs. It matched the accuracy of a stacked long short-term memory baseline (0.9111 ± 0.0050) but proved substantially more reproducible, with a seed-to-seed standard deviation approximately one order of magnitude lower. The proposed method provides an effective framework for friction prediction under the limited-data conditions typical of tribological research.

Keywords: artificial intelligence, recurrent algorithms, attention mechanism, coefficient of friction, coated alloy.

INTRODUCTION

A common cause of premature failure in mechanical systems is wear of working surfaces when cooperating machine parts operate without lubricant. It generates substantial economic and energy losses in industry, driving a continuous search for methods to counteract this problem [1, 2]. Pad welding (hardfacing) is one such method that deposits a wear-resistant layer using a welding technique. That is widely applied for both regenerating worn machine parts and protecting them from further wear [3–5]. Among the materials used for wear-resistant hardfacing layers, multi-component iron-based alloys are particularly popular. One of them is the eutectic Fe–Mn–C–B–Si–Ni–Cr alloy, which combines high hardness, arising from a significant fraction of hard boride and carbide phases embedded in a ductile austenitic–martensitic matrix, with good resistance to abrasive and adhesive wear [6–8]. The tribological response of multi-component iron-based alloys is, however, inherently nonlinear and non-stationary. In a typical pin/ball-on-disc tribological test, the friction signal passes through a running-in phase and a long quasi-stationary

operating period, while simultaneously undergoing high-frequency fluctuations caused by stick–slip phenomena, asperity interactions, third-body formation in the contact zone and tribo-oxidation [6, 7, 9]. During tests at elevated temperatures, this picture becomes even more complex. Iron, chromium, and silicon oxides begin to dominate the contact, modifying both the mean COF level and its variability [7, 10]. The friction coefficient should therefore be treated as a stochastic time series with evolving statistical properties, and the classical approach of reporting the mean and standard deviation in the steady state does not capture the full information contained in the signal. This is a recurring problem in many tribological studies investigating the nature of friction and wear processes [9, 11, 12]. The present authors, too, have so far relied mainly on measures based on the means and standard deviations of the tribological quantities under study [6–8].

The development of machine learning methods now allows tribological data to be examined from a much broader perspective. ANNs (artificial neural networks), SVMs (support vector machines), and ensemble methods have been used in recent years to predict wear rate, friction coefficient, or wear mechanism type from process variables for polymer composites [13], copper-based coatings [14] and ceramic friction pairs [15]. Among deep-learning architectures, recurrent neural networks deserve particular attention, including LSTM (long short-term memory) and GRU (gated recurrent unit) networks, whose gating mechanisms effectively handle both short- and long-range dependencies in sequential data [16–18]. Attention-based recurrent architectures of this type have also been applied to biomedical sequential signals [19]. Classical machine-learning algorithms (SVM, KNN – K-Nearest Neighbours, RF – Random Forest, XGBoost – eXtreme Gradient Boosting) have been applied to COF prediction in self-lubricating composites, reaching R^2 of about 0.92 [14]. An LSTM-based model for COF prediction achieved validation and test R^2 = 0.9973 and 0.9965, respectively, on coated Al10Cu [20]. Review papers devoted to machine learning in tribology, covering tribo-informatics [21], modelling frameworks [22] and multiscale simulation [23], document a transition from classical regressors (GBDT – Gradient Boosting Decision Tree) to hybrid deep architectures that combine a sequential branch with an encoder of physical descriptors. Li [24] proposed a hybrid LSTM–MLP (Multi-Layer Perceptron) for friction prediction in mixed elastohydrodynamic lubrication (EHL), in which parallel LSTM networks encode full surface-roughness profiles. At the same time, the MLP module integrates the operating conditions. The model, trained on 888,720 simulation points, achieved MSE (Mean Squared Error) = 1.91×10^{-7}, outperforming a standard MLP (MSE = 6.72×10^{-7}) and classical linear regression (MSE = 1.56×10^{-5}). An analogous strategy of combining a sequential branch with a dense branch was applied by Sabry et al. [25] to predict the corrosion rate of FSW (Friction Stir Welding) joints of the AA6061 alloy, where the LSTM hybrid achieved a 45% reduction in mean error relative to a classical ANN. Song et al. [26] in turn proposed a two-module attention mechanism (feature attention together with time attention) layered over an LSTM, which, for non-invasive friction diagnostics in hydrodynamic seals operating in four regimes (fluid, mixed, boundary and dry), achieved R^2 = 0.853 with RMSE = 0.026, outperforming the CNN (Convolutional Neural Network)–LSTM hybrid proposed in a related study [27] for COF prediction and wear-state diagnostics. An important aspect of friction- and wear-prediction studies is the frequent need to work with limited sample sizes, typical of materials and tribological research, due to their high cost and time-consuming nature. Yu et al. [28], addressing the same problem, extended an attention-LSTM with a GAN (for data augmentation) and an SSAE autoencoder (for feature extraction), thereby enabling tool-wear prediction despite imbalanced wear-state classes. Jayasinghe and Ramezani [29] applied three complementary architectures (ANN, RNN (Recurrent Neural Network), and CNN) to epoxy composites, obtaining R^2 = 0.98 for the friction coefficient and R^2 = 0.78 for the wear rate under reciprocating sliding conditions at 2–8 Hz and 5–15 N load. The RNN branch reproduced the temporal evolution of the COF under the dominant abrasion, adhesion, fatigue, and delamination mechanisms. Another important aspect is the high-temperature tribological behaviour of metallic alloys and the use of contact-accompanying signals for real-time assessment of the wear state. Saravanan et al. [11] showed that heat treatment of the WE54 alloy significantly modifies the friction coefficient and wear rate under sliding contact at elevated temperatures, confirming the sensitivity of the tribological response to the material's structural

state and operating temperature. They identified three coexisting wear mechanisms (abrasive, oxidative, and delamination) with the wear rate rising sharply at 150 °C regardless of the heat-treatment state. Pearson et al. [12] showed that, under fretting conditions, an increase in temperature accelerates the formation of protective Fe_2O_3 and Fe_3O_4 oxides through sintering. Tian et al. [30] developed a DNN (Deep Neural Network) model coupled with FEM (Finite Element Method) for brass–stainless steel pairs, achieving a COF relative error below 10% and using acoustic signals as a non-invasive indicator of the current wear state. Liu et al. [31] showed that friction-induced noise (Si_3N_4–SiC) carries information about surface roughness. LSTM, BiLSTM (Bidirectional Long Short-Term Memory), and GRU networks achieved over 90% classification accuracy, with time–frequency-domain features being the most informative. This indicates that the full tribological information resides in the high-frequency character of the friction signal.

Two tendencies are evident in previous studies. First, most developed recurrent models learn directly from the raw COF signal, omitting the statistical structure provided by a moving-window analysis. Second, the architectures used are typically transferred from other domains without adaptation to the specifics of tribological data. Only a few works explicitly couple a temporal branch of the network with a parallel branch that encodes alloy composition and process parameters – yet it is precisely this information that allows the model to generalise across different operating conditions and material variants. The eutectic Fe–Mn–C–B–Si–Ni–Cr system, previously investigated at room temperature and at 250 °C [6, 7], has not yet been studied using a combined experimental–computational approach with AI (Artificial Intelligence) tools.

Thus, the main motivation of the present study was to develop a method that combines physical descriptors of the alloy and the process with the dynamic characteristics of the friction signal, which allows the prediction of COF evolution under the limited-data conditions typical of tribological research. The aim of the present work is to design, implement, and validate a model for predicting the friction coefficient of the Fe–Mn–C–B–Si–Ni–Cr coating under dry sliding contact at ambient temperature and 300 °C. The authors' contributions in this study are as follows: (I) extending the experimental tribological characterisation of the Fe–Mn–C–B–Si–Ni–Cr coating to a contact temperature of 300 °C, complementing earlier data obtained at ambient temperature and at 250 °C; (II) developing a feature-engineering methodology based on the moving-window technique that transforms the raw friction signal into a sequence of statistical descriptors; (III) designing and validating a custom recurrent neural network (the proposed TriboRNN architecture), that integrates dynamic descriptors with alloy and process parameters to predict the evolution of the friction coefficient. The model was trained, validated and tested on an original dataset generated by the authors on the THT 1000 tribotester.

MATERIALS AND METHODS

A general scheme of the applied research methodology, linking specimen fabrication, tribological tests, and friction coefficient prediction via the developed recurrent neural network, is shown in Figure 1. In the first stage, the Fe–Mn–C–B–Si–Ni–Cr coating was deposited by arc hardfacing, and the specimens were ground and characterised for surface roughness and layer thickness. In the second stage, ball-on-disc tribological tests were carried out at ambient temperature and at 300 °C, with continuous recording of the friction coefficient and subsequent profilometric measurement of wear. In the third stage, the recorded friction signal was processed using the moving-window technique to obtain sequences of statistical descriptors, which, together with the alloy and process parameters, were used to train and validate the proposed recurrent network.

Coating production and tribological tests

The eutectic Fe–Mn–C–B–Si–Ni–Cr coating was deposited onto flat plates of S235JR structural steel, selected for its good weldability, by semi-automatic GMAW (Gas Metal Arc Welding) on a Kemppi Pro Evolution station, using a 2.4 mm flux-cored wire with a filling level of 30–35% under a CO_2 shielding atmosphere at a welding current of 270 A and an arc voltage of 30 V. The core-mixture composition was designed on the Fe–Mn–C–B system with Si, Cr and Ni additions; the wires were manufactured and the weld-metal chemical analysis was performed at the Institute of Welding in Gliwice, and the

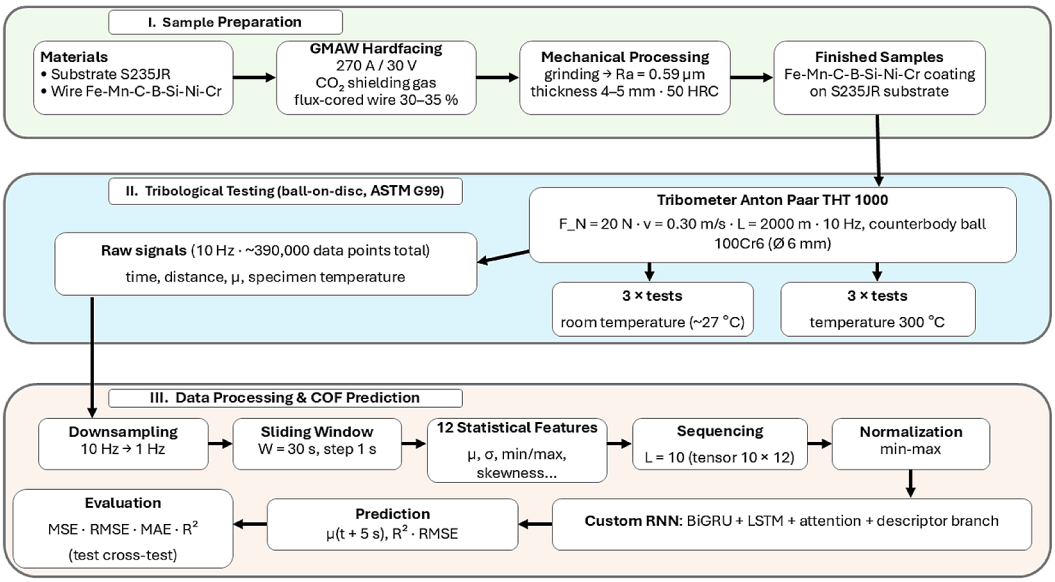


Figure 1. Pipeline of the applied research methodology

coating values given in Table 1 (C 1.63; Mn 7.29; B 1.79; Cr 16.24; Si 2.46; Ni 17.68 wt.%) come from that analysis, whereas the compositions of the S235JR substrate and of the 100Cr6 counterbody are nominal values according to the relevant standards, both materials having been supplied commercially with an inspection certificate. After pad welding, discs 8 mm in diameter and 10 mm thick were cut from the plates by electrical discharge machining, and the layer was ground with abrasive papers of progressively decreasing grit until a flat surface with a roughness of $R_a = 0.59 \pm 0.01$ µm was obtained, measured with a Surtronic 3+ stylus profilometer over a 4 mm evaluation length with a 0.8 mm cut-off per PN-ISO 4288, as the mean of ten repetitions, with a final hardfacing thickness of 4–5 mm, determined on metallographic cross-sections (phenolic-resin mounting, grinding from 120 µm, 3 µm diamond polishing, Mi1Fe/Mi20Fe etching) [6, 7]. The chemical compositions of the coating, the substrate, and the counterbody are summarised in Table 1.

The mean hardness of the hardfacing layer, taken as the mean of five Rockwell indentations, was 50 HRC [7]. Sliding-friction tests were carried out on a high-temperature THT 1000 tribotester manufactured by Anton Paar, operating in a ball-on-disc configuration. The counterbody was a 6 mm diameter 100Cr6 steel ball, loaded with a normal force of 20 N. The radius of the friction track was 6.03 mm, and the linear sliding velocity was 0.30 m/s (in accordance with ASTM G99). The tests were conducted at two contact temperatures: ambient conditions (~27 °C) and 300 °C, complementing previously published data at 250 °C [7]. The reported value of 300 °C refers to the furnace set-point; the corresponding mean sample temperature measured by a thermocouple during the test was 288 °C. For each temperature level, three independent runs were performed over a 2000 m sliding distance, yielding a total of six runs. The tribotester continuously recorded the instantaneous friction coefficient µ, the linear distance, the number of laps, and the sample temperature (at a frequency of 10 Hz), yielding approximately 65,000 measurements per test and 390,000 in total. Such a high acquisition frequency is advisable given the non-stationary nature of the friction signal, whose short-term fluctuations convey information about the current contact mechanisms. The sample temperature was monitored by a thermocouple placed

Table 1. Chemical composition (wt. %) of the eutectic coating, the S235JR substrate and the 100Cr6 counterbody

Material	C	Mn	B	Cr	Si	Ni	P	Fe
Fe–Mn–C–B–Si–Ni–Cr coating	1.63	7.29	1.79	16.24	2.46	17.68	–	Balance
S235JR substrate	0.17–0.20	≤1.40	–	–	–	–	0.045	Balance
100Cr6 ball (counterbody)	0.93–1.50	0.25–0.45	–	1.35–1.60	0.15–0.35	–	≤0.025	Balance

2 mm below the specimen surface. Wear of the disc was determined from cross-sectional profiles of the wear track recorded with the Surtronic 3+ profilometer perpendicular to the sliding direction, over a 4 mm evaluation length with a 0.8 mm cut-off per PN-ISO 4288, at ten points evenly distributed around the circumference of the track. The wear volume was obtained as the product of the track's mean cross-sectional area and length, and the wear rate was the wear volume divided by the product of the normal load and the sliding distance. Wear of the ball was determined from the wear-scar diameter measured under a metallographic microscope and converted to the volume of the removed spherical cap, with all calculations performed in the Tribometr software supplied with the THT 1000 tribometer. The coating hardness was measured on a Rockwell hardness tester, and the worn surfaces were examined using a Quanta 3D FEG scanning electron microscope.

Statistical feature extraction from a moving window

Training a recurrent network directly on the raw 10 Hz signal is costly. Adjacent samples differ only minimally, while significant tribological changes only emerge on a scale of several seconds [9, 10, 33]. For this reason, the signal is processed in two stages. In the first stage, the raw signals were resampled to 1 Hz by averaging every 10 consecutive measurements, preserving the macroscopic shape of the COF curve while removing electrical noise. Next, for each resampled signal, a window of width W = 30 s was shifted by 1 s. From each window, twelve features are computed: the mean, the median, Q_1 and Q_3 (describing the location of the COF distribution), the standard deviation, the minimum, the maximum and the range (its spread), skewness and kurtosis (shape), and the mean absolute difference between consecutive samples together with the standard deviation of those differences (the local variability of the signal). The output is a 12-dimensional feature vector updated every second. The moving-window technique applied here and the set of descriptors used (standard deviation, kurtosis, skewness) originate from machine diagnostics and condition-based maintenance, where statistical features of this kind are a standard tool for extracting information from noisy mechanical signals [32]. Let μ_k denote the value of the friction coefficient at the k-th sample of the signal resampled to 1 Hz. For a window of width W = 30 samples (30 s) ending at instant k, the central tendency and dispersion are computed according to Equations 1 and 2 [32]:

$$\bar{\mu}_k = \frac{1}{W}\sum_{j=1}^{W}\mu_{k-W+j}\,,\sigma_k = \\ = \sqrt{\frac{1}{W-1}\sum_{j=1}^{W}\left(\mu_{k-W+j}-\bar{\mu}_k\right)^2} \tag{1}$$

$$\mu_k^{min} = \min_{1\le j\le W}\mu_{k-W+j}\,,\mu_k^{max} = \\ = \max_{1\le j\le W}\mu_{k-W+j}\,,R_k = \mu_k^{max}-\mu_k^{min} \tag{2}$$

The median and quartiles, as robust measures of the position of the μ-distribution within the window, are defined in Equation 3 as empirical p-th quantiles Q_p. The shape of the in-window distribution is quantified by the standardised moments of third order (skewness) and fourth order (kurtosis), as in Equation 4. These two measures characterise, respectively, the asymmetry of the distribution and the excess tail mass associated with momentary stick–slip events [32]:

$$\tilde{\mu}_k = Q_{0.50}(x_k), Q_{1,k} = \\ = Q_{0.25}(x_k), Q_{3,k} = Q_{0.75}(x_k) \tag{3}$$

$$\gamma_{1,k} = \frac{1}{W}\sum_{j=1}^{W}\left(\frac{\mu_{k-W+j}-\bar{\mu}_k}{\sigma_k}\right)^3,\gamma_{2,k} = \\ = \frac{1}{W}\sum_{j=1}^{W}\left(\frac{\mu_{k-W+j}-\bar{\mu}_k}{\sigma_k}\right)^4-3 \tag{4}$$

The local variability of the signal inside the window is characterised by the mean absolute change between consecutive samples (mac) and by the standard deviation of the first differences, as in Equation 5 [32]:

$$mac_k = \frac{1}{W-1}\sum_{j=2}^{W}\left|\Delta\mu_j\right|,\sigma_k^{\Delta} = \\ = std\left(\{\Delta\mu_j\}_{j=2}^{W}\right),\Delta\mu_j = \\ = \mu_{k-W+j}-\mu_{k-W+j-1} \tag{5}$$

All twelve descriptors computed for the window ending at instant k are aggregated into a 12-dimensional feature vector, as shown by Equation 6. The recurrent network receives as input a

stack of the last L = 10 such vectors, forming the time tensor (7), and the prediction target (8) is the 1 Hz-resampled COF value mean μ_{k+h} measured H = 5 s after the end of the observation window. Both the targets and the feature channels undergo min–max normalisation (9) [34], computed exclusively from the statistics of the training set (denoted by the superscript tr).

$$f_k = \begin{bmatrix} \bar{\mu}_k, \sigma_k, \mu_k^{min}, \mu_k^{max}, R_k, \tilde{\mu}_k, \\ \gamma_{1,k}, \gamma_{2,k}, Q_{1,k}, Q_{3,k}, mac_k, \sigma_k^{\Delta} \end{bmatrix}^{\top} \in R^{12} \quad (6)$$

$$X_k = [f_{k-L+1}, f_{k-L+2}, \dots, f_k]^{\top} \in R^{L\times 12}, L = 10 \quad (7)$$

$$y_k = \bar{\mu}_{k+H}, H = 5\ s \quad (8)$$

$$\tilde{y}_k = \frac{y_k - y_{min}^{tr}}{y_{max}^{tr} - y_{min}^{tr}} \quad (9)$$

The recurrent network's input is a stack of L = 10 consecutive feature vectors (covering the last 10 s of moving-window outputs and, through the W = 30 s window of each vector, an effective lookback of approximately 39 s of contact history), forming a tensor of dimensions (L, 12). The prediction target is the COF value measured H = 5 s after the end of the observation window, i.e., forecasting friction 5 s ahead. This horizon is short enough for the COF to remain strongly determined by the current contact state, yet long enough to be practically relevant. Further extension of the horizon only increased the variance of the predictions.

Recurrent neural network architecture TriboRNN

The custom recurrent architecture was implemented in the TensorFlow 2.20 and Keras 3.x environment (Figure 2). The network consists of two parallel branches merged before the regression head and has approximately 36,800 trainable parameters.

The first branch processes the time series of statistical descriptors. At the input, a bidirectional GRU layer (BiGRU) with 48 hidden units returns the full sequence. The bidirectional configuration allows information from both ends of the observation window to be considered. After layer normalisation and 15% dropout, the sequence is passed through a unidirectional LSTM layer (32 units), which also returns the full sequence, followed by another layer normalisation. The final element of the branch is an additive attention mechanism (a dense layer with tanh activation followed by a softmax along the time axis), which produces a

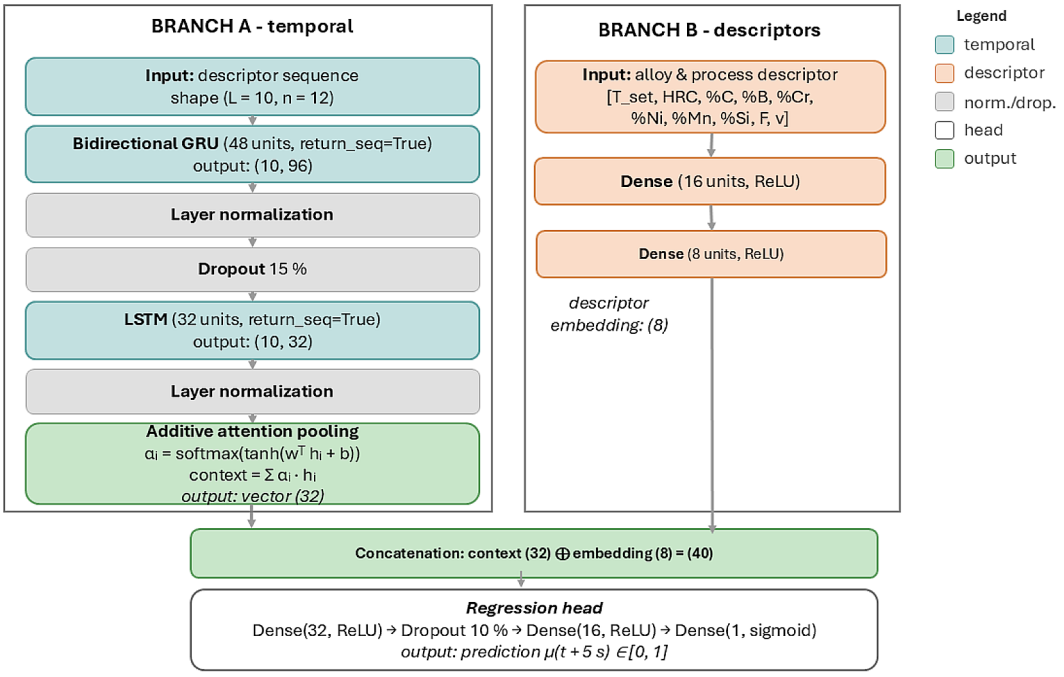


Figure 2. Detailed architecture of the proposed TriboRNN recurrent neural network. The encircled plus (⊕) denotes the concatenation of the context vector c with e_alloy. The merged vector is propagated through the regression head φ (Dense(32, ReLU) → Dropout 10% → Dense(16, ReLU) → Dense(1, sigmoid)), which returns the normalised friction-coefficient prediction ŷ ∈ [0, 1]

context vector as a weighted sum of the LSTM hidden states. As a result, the network can selectively emphasise the most informative time steps, which is not possible in classical stacked-LSTM architectures. Let $h_t \in R^{32}$ denote the hidden state of the second LSTM layer at time step t = 1, …, L, and w, b the trainable parameters of the attention mechanism; the raw alignment scores are computed in equation (10) [35]. The normalised weights α_t are obtained via a softmax along the time axis in equation (11), and the context vector c, being a weighted sum of the LSTM hidden states, is defined by equation (12):

$$e_t = tanh(w^\top h_t + b), t = 1, 2, \ldots, L \quad (10)$$

$$\alpha_t = \frac{exp(e_t)}{\sum_{j=1}^{L} exp(e_j)}, \sum_{t=1}^{L} \alpha_t = 1 \quad (11)$$

$$c = \sum_{t=1}^{L} \alpha_t h_t \in R^{32} \quad (12)$$

The second branch encodes the alloy and process parameters as a ten-dimensional vector comprising the set temperature, bulk hardness, mass fractions of C, B, Cr, Ni, Mn and Si, the normal load and the sliding velocity. In the present dataset, only the set temperature varies across the six runs; the remaining nine inputs are constant, so within this dataset, branch B effectively conditions the model on the contact-temperature regime. The full ten-dimensional encoding is retained, enabling the architecture to be directly transferred to multi-alloy and multi-condition datasets without modification. The vector passes through two dense layers (16 and 8 neurons, ReLU activation), producing an alloy- and process-descriptor embedding $e_{alloy} \in R^8$. The context vector c from branch A is then concatenated with e_{alloy} from branch B. The result propagates through the regression head φ (Dense(32, ReLU) → Dropout 10% → Dense(16, ReLU)), terminated by a single neuron with a sigmoidal activation σ that generates the normalised prediction $\hat{y} \in [0, 1]$, according to Equation 13 [34]:

$$\hat{y}_k = \sigma\left(w_o^\top \varphi\left([c\,;\, e_{alloy}]\right) + b_o\right) \quad (13)$$

Both the COF targets and the feature channels were normalised using the min–max method based on training-set statistics. The network was optimised using the Adam algorithm (initial learning rate 1×10^{-3}), with an MSE (Mean Squared Error) loss function and a batch size of 128, for a maximum of 60 epochs. Early stopping was used (patience of 8 epochs), with the learning rate halved after 4 epochs without improvement in the validation loss. The seeds of the random-number generators (NumPy, Python, TensorFlow) were fixed at 42 to ensure reproducibility. The loss function $\mathcal{L}(\theta)$ optimised during training on a training set of size N_{tr} (the number of samples in the training set) has the form (14) [34]:

$$\mathcal{L}(\theta) = \frac{1}{N_{tr}} \sum_{k=1}^{N_{tr}} \left(\tilde{y}_k - \hat{y}_k(\theta)\right)^2 \quad (14)$$

Baseline model, dataset split and evaluation metrics

As the reference model, a dual-layer (stacked) LSTM network was implemented (64 → 32 hidden units with tanh activation and return_sequences = True in the first layer; a head composed of a 16-unit dense layer with ReLU activation and a sigmoidal output layer), reproducing the architecture described in [20] for friction prediction in coated systems. Both models received the same input data and were trained with identical hyperparameters and stopping criteria. To verify the model's ability to generalise, two of the six tests (one for each temperature) were withheld in their entirety as an unseen test set during training. From the remaining four runs, 15% of the samples were set aside as a validation set for early stopping and learning rate adjustment. After processing with moving windows (length 30 s, step 1 s) and a lookback of 10, the procedure yielded 21,993 training, 3,881 validation, and 12,937 test sequences. Separating complete runs (rather than randomly shuffling samples from a single time series) provides a more rigorous assessment of the model's ability to generalise, because the test set includes complete tribological runs, including the running-in phase, that the model has never seen before. The resulting splits are 57% training, 10% validation, and 33% test (21,993/3,881/12,937 sequences), a more demanding scenario than the 70/15/15 split commonly used with randomly shuffled samples. The prediction accuracy of both the proposed TriboRNN and the baseline LSTM was assessed on the test set, after inverse min–max scaling to the original friction-coefficient

scale, using four standard regression metrics: the mean squared error (MSE), the root mean squared error (RMSE), the mean absolute error (MAE) and the coefficient of determination (R^2). Let y_k and $\hat{y}_k$ denote, respectively, the measured and predicted value of μ on the k-th sample of an evaluation set of size N. Equations (15)–(18) define the four metrics [34]. To quantify the contribution of the individual architectural components of the TriboRNN, an ablation study was additionally performed, in which each component (the bidirectional GRU front-end, the additive attention pooling, and the parallel alloy- and process-descriptor branch) was removed in turn. Because the absolute differences between variants are small, each variant was retrained across five random seeds (42–46) using the same fixed cross-test split, normalisation, and hyperparameters, and the metrics are reported as mean ± standard deviation (Table 4).

$$MSE = \frac{1}{N}\sum_{k=1}^{N}(y_k - \hat{y}_k)^2 \tag{15}$$

$$RMSE = \sqrt{\frac{1}{N}\sum_{k=1}^{N}(y_k - \hat{y}_k)^2} \tag{16}$$

$$MAE = \frac{1}{N}\sum_{k=1}^{N}|y_k - \hat{y}_k| \tag{17}$$

$$R^2 = 1 - \frac{\sum_{k=1}^{N}(y_k - \hat{y}_k)^2}{\sum_{k=1}^{N}(y_k - \bar{y})^2} \tag{18}$$

RESULTS AND DISCUSSION

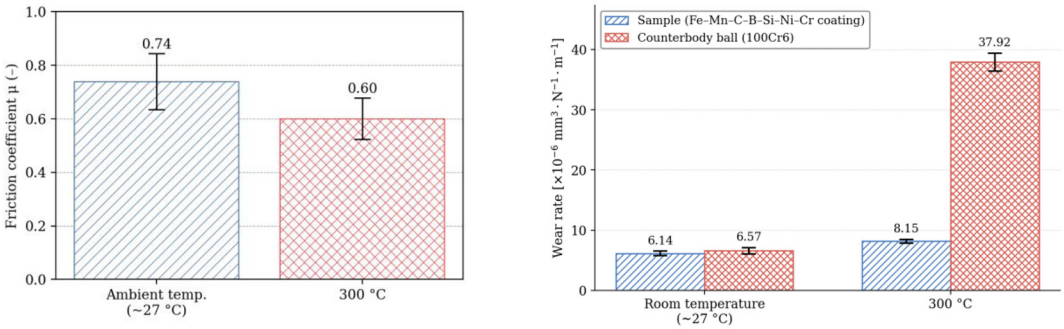
Friction coefficient and wear at ambient temperature and at 300 °C

Over the entire sliding distance of 2000 m, the mean friction coefficient was 0.739 ± 0.105 at ambient temperature and 0.600 ± 0.077 at 300 °C (Figure 3a).

The ~19% decrease is consistent with the value previously observed at 250 °C (μ = 0.59) [7]; the friction coefficient thus stabilises at a comparable level at and above 250 °C. The descriptive statistics gathered in Table 2 further show that the COF signal at 300 °C is not only lower but also more concentrated around the median (lower interquartile range, lower coefficient of variation). This is consistent with the formation of a stable tribological oxide layer that smooths the asperity interactions [7, 10].

Representative COF time histories (Figure 4) show a classical three-stage evolution: a rapid rise during the running-in phase (the first ~50 s), followed by a broad, unstable period of continued friction increase, and finally a quasi-stationary state. At 300 °C, the amplitude of the high-frequency fluctuations is markedly smaller than at ambient temperature, suggesting the formation of a continuous oxide film.

The wear rate of the specimen (coating) and of the counterbody (100Cr6 ball) is summarised in Figure 3b. At ambient temperature, the wear rate of the coating (6.14×10^{-6} mm³·N⁻¹·m⁻¹) is markedly lower than at 300 °C (8.15×10^{-6} mm³·N⁻¹·m⁻¹; ~33% increase), which indicates the effective anti-wear protection provided by the eutectic microstructure of the alloy at



a) b)

Figure 3. Tribological parameters (mean ± s.d. of three runs per temperature) a) mean friction coefficient for the eutectic Fe–Mn–C–B–Si–Ni–Cr coating tested at ambient temperature and at 300 °C, b) mean wear rate (×10⁻⁶ mm³·N⁻¹·m⁻¹) of the Fe–Mn–C–B–Si–Ni–Cr coating specimen and of the 100Cr6 ball counterbody at ambient temperature (~27 °C) and at 300 °C

Table 2. Descriptive statistics of the friction coefficient at ambient temperature and at 300 °C. Statistics are computed over all three runs at each temperature (including the run later withheld as the test set); all quantities except the sample count N are dimensionless (COF, μ [−])

Statistic	Ambient temp. (~27 °C)	300 °C
Sample count N	195,444	195,351
Mean	0.7386	0.6003
Median	0.7388	0.5995
Standard deviation	0.1047	0.0774
Variance	0.0110	0.0060
Minimum	−0.0514	−0.0372
Maximum	1.0180	0.8467
Q1 (25th percentile)	0.6773	0.5528
Q3 (75th percentile)	0.8087	0.6538
Interquartile range	0.1314	0.1010
Skewness	−1.1578	−0.5127
Kurtosis	5.8253	2.9348
Coefficient of variation (%)	14.18	12.89

non-elevated temperature. The wear of the ball (counterbody) increases more steeply with temperature, from 6.57×10^{-6} mm³·N⁻¹·m⁻¹ at ambient temperature to 37.92×10^{-6} mm³·N⁻¹·m⁻¹ at 300 °C (about a 5.8-fold increase), reflecting accelerated oxidative wear and material transfer onto the ball at elevated temperature.

Scanning electron microscopy images of the wear scars (Figure 5) reveal that the surface degradation mechanism changes with test temperature. At ambient temperature (Figure 5a), a clearly developed wear track is visible, with detached wear debris distributed along both sides of the track; the regularly spaced linear marks oriented obliquely to the sliding direction are residual grinding traces from the specimen preparation (R_a = 0.59 ± 0.01 μm) and not tribological grooves. Within the wear track itself, the surface shows plastic deformation and material transfer characteristic of micro-cutting and adhesive interaction with the counterbody. This morphology corresponds to the high friction coefficient recorded at this temperature (Figure 3a). It is typical of the abrasive mechanism observed in metal–metal systems lacking a protective oxide layer [9, 10]. At 300 °C (Figure 5b), the wear scar is clearly smoothed, and sharp abrasion is replaced by a flat, low-contrast tribological film, which indicates the formation of a sintered oxide layer (glaze layer) under elevated-temperature conditions. This phenomenon is consistent with earlier observations [12, 33] showing that an increase in temperature accelerates the formation of protective oxide layers through the nanoscale sintering of Fe_2O_3 and Fe_3O_4 particles, leading to a smooth, load-bearing oxide film that reduces friction and, under favourable conditions, also wear. An analogous

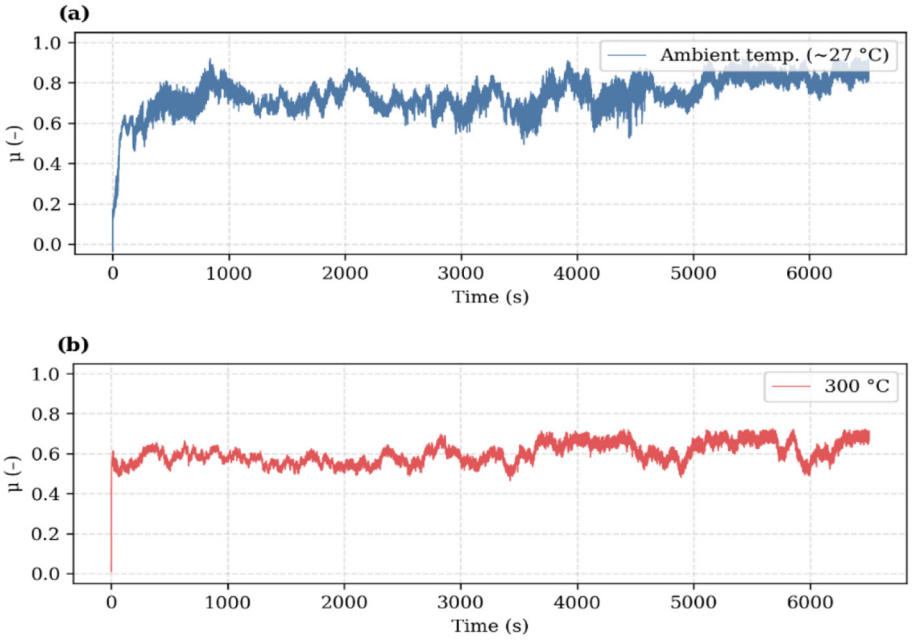


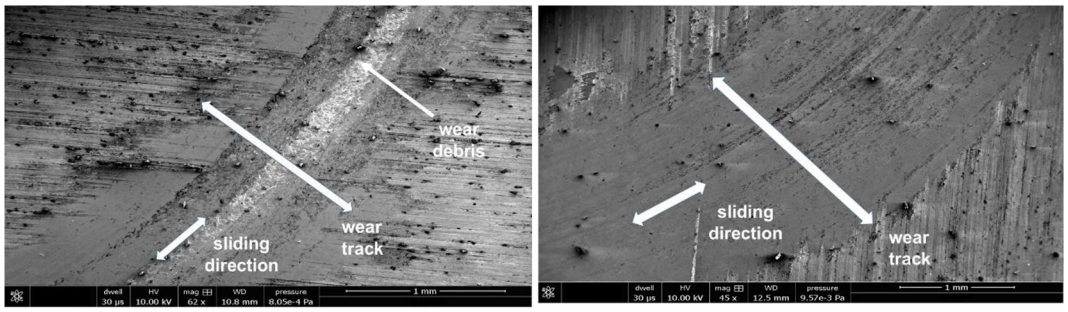
Figure 4. Time evolution of the friction coefficient for the Fe–Mn–C–B–Si–Ni–Cr coating (a) at ambient temperature and (b) at 300 °C

tribofilm-stabilisation effect was reported by Kolev et al. [20] in the Cr/nanodiamond–Al10Cu system, where the hard composite coating preserved the surface integrity after the tribological test, and is consistent with the tribofilm-mediated reduction of wear observed by Pearson et al. [12] under elevated-temperature fretting contact.

In the investigated Fe–Mn–C–B–Si–Ni–Cr system, the eutectic boride–carbide phases carry the contact load and limit abrasive penetration into the matrix, which (together with the oxide tribological film at 300 °C) results in a reduction of the friction coefficient (Figure 3a). At the same time, despite the change of the dominant wear mechanism from abrasive to oxidative, the combined wear rate of the two bodies at 300 °C increases (Figure 3b). The glaze layer reduces friction. Still, it does not compensate for the accelerated thermal and oxidative degradation of the coating material and of the counterbody.

Moving-window features

Two selected moving-window features are compared in Figure 6: the mean ± standard deviation and the range. At ambient temperature, the mean rises steeply during the running-in phase and then drifts gradually upward over the entire test duration (~6500 s), while the range remains large and noisy. At 300 °C, the mean quickly stabilises around 0.6, and the range remains low and almost constant. This observation justifies



a) b)

Figure 5. SEM images of the wear scars after the tribological test: a) at ambient temperature (~27 °C), magnification 62×; b) at 300 °C, magnification 45×. Scale bar = 1 mm in both panels

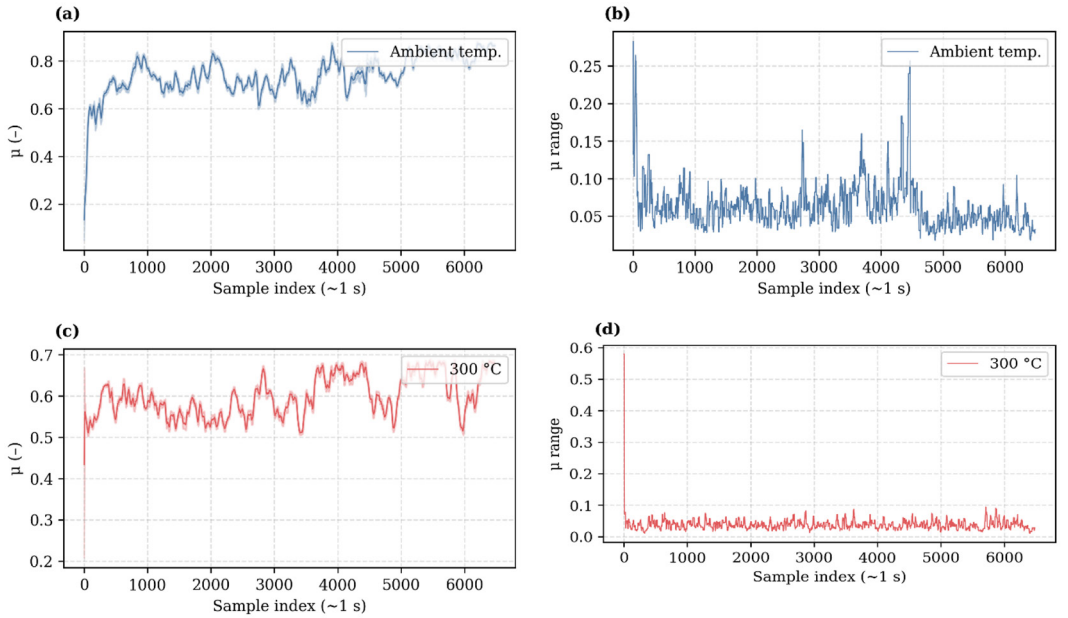


Figure 6. Statistical features extracted with the moving-window technique from one representative test at ambient temperature (a and b) and at 300 °C (c and d), where (a and c) is the in-window mean ± standard deviation and (b and d) is the in-window range (max − min) of the friction coefficient

using the descriptors as the network input. Each 12-dimensional vector compresses 30 s of contact history into a set of physically interpretable quantities, and their evolution in time reflects the current tribological conditions.

Prediction accuracy of the custom recurrent network

The learning curves of both networks (Figure 7) decrease smoothly, with the validation loss tracking the training loss throughout and no signs of overfitting — a consequence of the layer normalisation and dropout regularisation. The custom network converges within about 40 epochs (lowest validation loss at epoch 36; training was early-stopped at epoch 44), whereas the baseline continues to improve gradually up to the 60-epoch cap. Consistent with Table 3, the baseline attains a slightly lower validation loss; nevertheless, the two models are statistically equivalent on the entirely withheld test runs, indicating that this lower in-distribution validation loss does not translate into better generalisation to unseen runs.

A comparison of the measured and predicted COF values on the withheld tests (Figure 8) shows that the custom network reproduces both the slow drift and the medium-frequency fluctuations of the friction signal. At ambient temperature, the predicted curve closely follows the

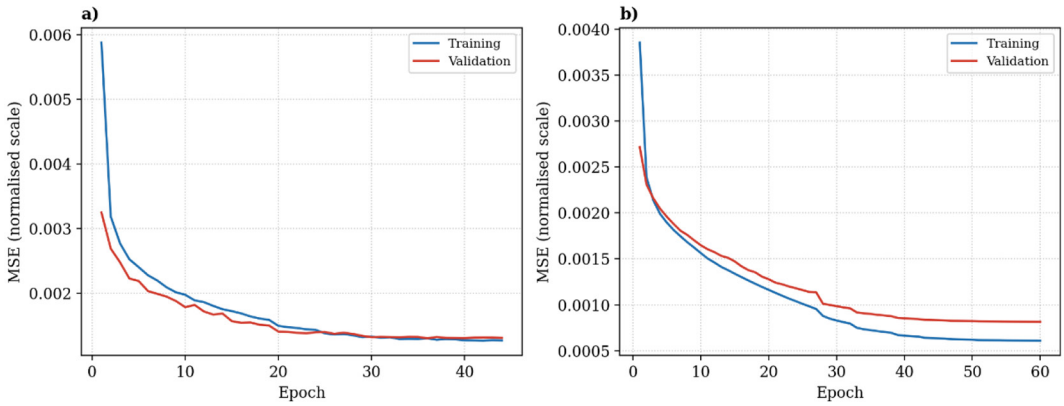


Figure 7. Training and validation losses (mean squared error in normalised scale) for the custom recurrent neural network (a) and for the baseline stacked-LSTM (b)

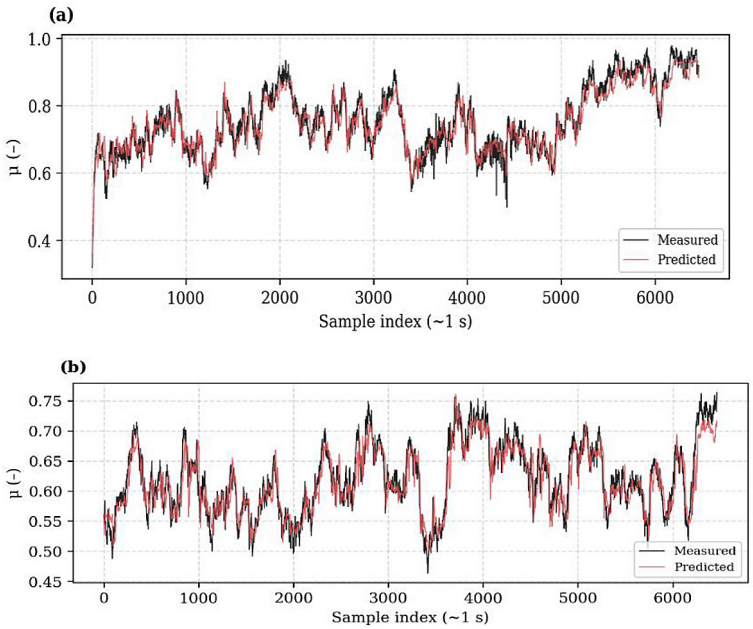


Figure 8. Measured and predicted friction coefficient on two withheld tests: at ambient temperature (a) and at 300 °C (b). The predictions are produced by the developed recurrent network with a 5s forecast horizon (representative single run, seed 42)

measured one throughout the entire running-in phase and the long increase up to μ ≈ 0.95 (R^2 = 0.869, RMSE = 0.0353). At 300 °C, where the signal is more stationary, the model achieves R^2 = 0.796 and RMSE = 0.0263. The corresponding parity plot of measured versus predicted COF on the validation and test sets (Figure 9) confirms close agreement across the full friction range.

The lower R^2 results from the smaller target variance rather than from a worse absolute fit; in absolute terms (RMSE), the fit at 300 °C is in fact tighter than at ambient temperature. The pooled R^2 reported in Table 3 is computed over the two withheld runs jointly, so its denominator includes the between-regime variance arising from the different mean COF levels at ambient temperature (≈ 0.74) and at 300 °C (≈ 0.60). This between-regime component inflates the denominator of R^2 relative to the within-regime variance, thereby raising the pooled R^2 above the per-run values for any model that captures the regime difference. The same effect is visible for the baseline stacked-LSTM (pooled R^2 = 0.9111 vs per-run 0.858 and 0.819; Table 3). The per-run coefficients (R^2 = 0.869 at ambient temperature and R^2 = 0.796 at 300 °C) therefore better characterise the genuine within-regime predictive quality. At the same time, the pooled R^2 remains useful for comparison with the single-split results reported in the literature. The per-run values quoted above correspond to the single representative run shown in Figures 7 and 8 (seed 42); the corresponding five-seed means ± s.d. are given in Table 3.

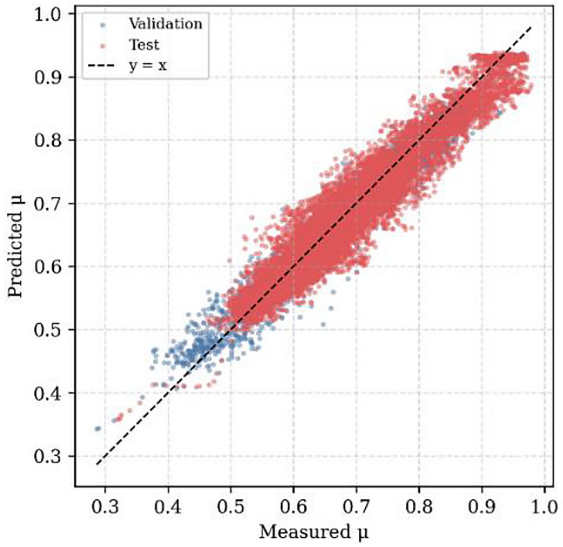


Figure 9. Parity plot of the measured vs predicted friction coefficient for the custom recurrent neural network on the validation and test sets (seed 42)

On the withheld tests, the two models are statistically indistinguishable in accuracy: pooled R^2 = 0.9113 ± 0.0005 for TriboRNN versus 0.9111 ± 0.0050 for the baseline (RMSE 0.0312 ± 0.0001 vs 0.0313 ± 0.0009). On the validation set and in the more stationary 300 °C regime, the baseline is in fact marginally better (validation R^2 0.966 vs 0.949; 300 °C R^2 0.819 vs 0.803), while TriboRNN is marginally better in the non-stationary ambient regime (R^2 0.865 vs 0.858). The most pronounced and consistent difference is in stability: the seed-to-seed standard deviation of the pooled test R^2 is about an order of magnitude smaller for TriboRNN (0.0005 vs 0.0050), i.e. the proposed architecture is markedly more reproducible. This low seed-to-seed variance, together with the interpretability of the attention weights and the extensibility provided by the parallel descriptor branch (Section 3.4), rather than a raw accuracy gain, constitutes the practical advantage of TriboRNN on this dataset. Notably, under this rigorous leave-runs-out protocol, even a plain

Table 3. Prediction accuracy of the proposed TriboRNN and the baseline stacked-LSTM, reported as mean ± standard deviation over five random seeds (42–46) on the same cross-test split

Model and subset	MSE	RMSE	MAE	R^2
TriboRNN – validation	$6.40 \times 10^{-4} \pm 0.3 \times 10^{-4}$	0.0252 ± 0.0005	0.0192 ± 0.0004	0.9486 ± 0.0022
TriboRNN – test (pooled)	$9.73 \times 10^{-4} \pm 0.1 \times 10^{-4}$	0.0312 ± 0.0001	0.0244 ± 0.0001	0.9113 ± 0.0005
TriboRNN – test, ambient	–	0.0358 ± 0.0004	0.0285 ± 0.0002	0.8646 ± 0.0032
TriboRNN – test, 300 °C	–	0.0259 ± 0.0004	0.0203 ± 0.0003	0.8030 ± 0.0065
Baseline LSTM – validation	$4.21 \times 10^{-4} \pm 0.7 \times 10^{-4}$	0.0205 ± 0.0016	0.0157 ± 0.0011	0.9660 ± 0.0054
Baseline LSTM – test (pooled)	$9.79 \times 10^{-4} \pm 0.6 \times 10^{-4}$	0.0313 ± 0.0009	0.0241 ± 0.0005	0.9111 ± 0.0050
Baseline LSTM – test, ambient	–	0.0366 ± 0.0011	0.0287 ± 0.0008	0.8583 ± 0.0088
Baseline LSTM – test, 300 °C	–	0.0248 ± 0.0006	0.0194 ± 0.0003	0.8191 ± 0.0087

stacked-LSTM achieves $R^2 \approx 0.91$, far below the $R^2 > 0.99$ reported for the same architecture under random-split evaluation [20] — confirming that the high literature values are inflated by the validation scheme rather than by the model.

Ablation study

To quantify the contribution of the individual architectural elements of the proposed TriboRNN, an ablation study was carried out. Starting from the full model, a single component was removed at a time, and the network was retrained from scratch. Because the absolute differences between variants are small, every variant was trained over five random seeds (42–46) on a fixed cross-test split, with identical feature normalisation and training hyper-parameters; the metrics are reported as mean ± standard deviation, so the seed spread reflects training noise (weight initialisation, dropout, mini-batch order) only. Three factors were considered: (i) the bidirectionality of the GRU front-end, where a unidirectional GRU replaced the BiGRU; (ii) the attention pooling, replaced by temporal average pooling over the LSTM hidden states; and (iii) the parallel alloy- and process-descriptor branch (branch B), removed to leave the temporal branch only. All variants were evaluated on the same two entirely withheld tests as the full model. Averaged over five seeds, each component removes a specific, regime-dependent weakness rather than uniformly raising accuracy (Table 4, Table 5, Figure 10). Removing the attention pooling most degrades the 300 °C regime (R^2 = 0.7953 vs 0.8030 for the full model; four of five seeds perform worse than the full model), indicating that selective time weighting is useful when the signal is stationary.

A single seed at 300 °C makes plain average pooling look competitive with attention, but this advantage disappears when averaging across five seeds. Removing the alloy- and process-descriptor branch most degrades the ambient regime (R^2 = 0.8583) and yields the lowest pooled R^2 (0.9091), confirming that jointly encoding the dynamic signal with the material and process parameters improves generalisation where mechanical wear dominates. Replacing the bidirectional GRU with a unidirectional one does not lower the mean pooled R^2 (0.9123 vs 0.9113 — the small positive difference is well within one standard deviation and is not statistically meaningful) but increases its run-to-run variance about six-fold (s.d. 0.0030 vs 0.0005) while cutting the parameter count by about 41% (36 826 → 21 658); the bidirectional front-end therefore does not improve mean accuracy on this dataset, but materially reduces seed-to-seed variability. Removing the attention pooling creates a 300 °C weakness, and removing the descriptor branch creates an ambient weakness; the full model and the unidirectional variant avoid both, but the full (bidirectional) model has about six-fold lower seed-to-seed variance, which is the basis for retaining the bidirectional front-end in the final architecture. The absolute pooled differences are small and largely within the seed spread, as expected given the demanding cross-test protocol and the fact that the 12-dimensional

Table 4. Ablation study of the proposed TriboRNN architecture (mean ± s.d. over five seeds, 42–46). Pooled metrics on the two entirely withheld tests. ΔR^2 is given relative to the full model

Variant	Trainable parameters	RMSE (test)	MAE (test)	R^2 (test)	R^2 (val.)	ΔR^2 (test) [pp]
Full (TriboRNN)	36 826	0.0312 ± 0.0001	0.0244 ± 0.0001	0.9113 ± 0.0005	0.9486 ± 0.0022	—
− Bidirectional	21 658	0.0311 ± 0.0005	0.0242 ± 0.0004	0.9123 ± 0.0030	0.9471 ± 0.0011	+0.10
− Attention	36 793	0.0313 ± 0.0003	0.0244 ± 0.0003	0.9108 ± 0.0019	0.9504 ± 0.0020	−0.05
− Descriptor (B)	36 300	0.0316 ± 0.0003	0.0246 ± 0.0002	0.9091 ± 0.0019	0.9519 ± 0.0046	−0.22

Table 5. Per-temperature breakdown of the ablation study (R^2 ± s.d. / RMSE on each withheld test, five seeds)

Variant	Ambient (~27 °C) R^2 / RMSE	300 °C R^2 / RMSE
Full (TriboRNN)	0.8646 ± 0.0032 / 0.0358	0.8030 ± 0.0065 / 0.0259
− Bidirectional	0.8655 ± 0.0061 / 0.0357	0.8074 ± 0.0043 / 0.0256
− Attention	0.8662 ± 0.0035 / 0.0356	0.7953 ± 0.0055 / 0.0264
− Descriptor (B)	0.8583 ± 0.0038 / 0.0366	0.8063 ± 0.0085 / 0.0257

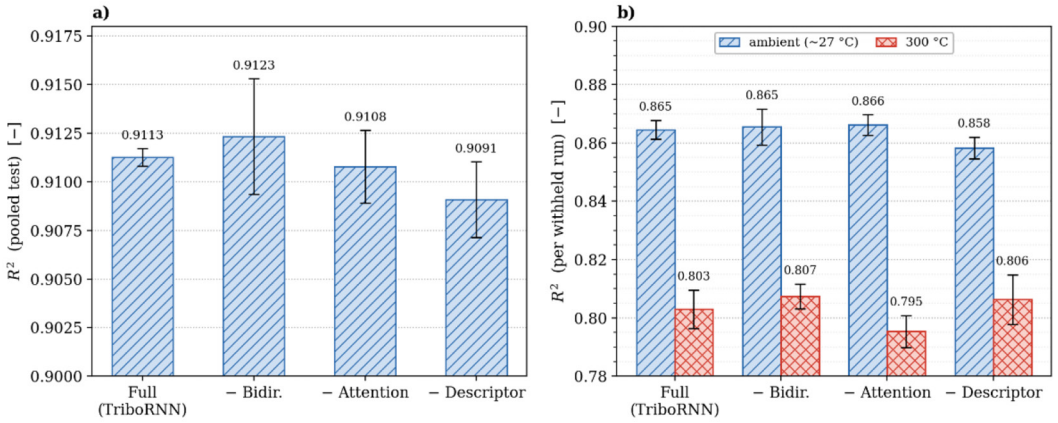


Figure 10. Multi-seed ablation (five seeds, 42–46): (a) pooled test R^2 (mean ± s.d.); (b) per-regime test R^2. Removing the attention creates a 300 °C weakness; removing the descriptor branch creates an ambient weakness. The full model and the unidirectional variant avoid both, with the full (bidirectional) model showing about six-fold lower seed-to-seed variance

descriptors already capture most tribological information; however, the consistent per-regime pattern and lower variance justify retaining all three components in the final architecture. For completeness, a stronger additive-attention variant matching the canonical form $w^T \cdot \tanh(W_a h_t + b_a)$ was also evaluated over the same five seeds; it did not improve the test accuracy (pooled $R^2 = 0.9100 \pm 0.0018$, slightly lower due to mild overfitting), so the lightweight attention was retained.

DISCUSSION

The wear results show a clear asymmetry of tribological resistance between the coating and the counterbody. At ambient temperature, the coating's wear rate is low due to the effective action of the hard eutectic phases (borides, carbides), which carry the contact load and limit microcutting of the matrix. At 300 °C, the wear rate of both bodies increases; this phenomenon is associated with intense oxidation of the counterbody and material transfer into the contact zone, accelerated by elevated temperature and by the presence of the tribological oxide layer [7, 10]. Good measurement repeatability (low standard deviation) indicates the homogeneity of the coating's eutectic microstructure. The obtained R^2 and RMSE values are comparable to results reported in recent literature on friction prediction. The stacked-LSTM model applied to the Al10Cu alloy with a chromium/nanodiamond coating achieved $R^2 > 0.99$ [20], but this result was obtained from a single run with random time-series splitting. After transferring the same architecture to the cross-test split applied here, both the baseline and the proposed network reach $R^2 \approx 0.91$ (pooled), a far more realistic measure of generalisation ability than the $R^2 > 0.99$ obtained under random splitting. According to this protocol, the two architectures are statistically equivalent in accuracy (Table 3); the proposed network is distinguished mainly by its substantially lower seed-to-seed variance. Li [24], for mixed EHL friction using full roughness profiles (LSTM–MLP, MSE = 1.91×10^{-7}), showed that the sequential branch can distinguish surfaces with identical statistical parameters (Root Mean Square, RMS) but different textures. This observation is consistent with the premise adopted here that the sequence of moving-window statistical-descriptor vectors carries substantially more information than single mean values and standard deviations. Sabry et al. [25], in predicting corrosion rate in FSW joints (LSTM–dense hybrid, 45% error reduction vs ANN), showed that the dense branch effectively encodes the structural process parameters, in analogy to branch B of the proposed network, which processes the alloy and tribometer parameters. Yu et al. [28], for tool wear with GAN-based data augmentation and an SSAE autoencoder for feature extraction, reported that an attention-LSTM can outperform a classical stacked-LSTM in small-sample regimes; in the present work, under a strict leave-runs-out protocol, the two architectures reached comparable pooled-test accuracy (Table 3), with the attention-based model offering greater training

stability and interpretability rather than a higher mean score. Song et al. [26] reported $R^2 = 0.853$ and RMSE = 0.026 for non-invasive diagnostics of hydrodynamic seals. These are values of the same order as those obtained in the present work at 300 °C ($R^2 = 0.796$, RMSE = 0.0263 on the representative single run; $R^2 = 0.8030 \pm 0.0065$ across five seeds, Table 3), despite the different material system. There is, however, an important difference relative to [20]. There, the LSTM achieved $R^2 > 0.99$ (validation: 0.9973; test: 0.9965) on a single run with random series splitting. In contrast, the cross-test scheme applied here (entire tribological runs absent from training) is more demanding and closer to real-world conditions. The stabilisation of the COF at 300 °C through the glaze layer is consistent with independent studies of high-temperature wear mechanisms. Saravanan et al. [11] reported that, for the WE54 alloy in the 25–150 °C range, abrasive, oxidative, and delamination mechanisms coexist, with the wear rate increasing with temperature regardless of heat-treatment state. In the data presented here, temperature is the operating parameter most strongly influencing the evolution of μ, which justifies its inclusion as an explicit input in the alloy- and process-descriptor branch.

Tribologically, the model predicts well in two regimes: high-friction at ambient temperature, driven by micro-cutting and adhesive transfer, and at 300 °C, where oxidation yields a smoother, lower signal. A single model handles both via the descriptor branch, enabling industrial monitoring without per-point retraining. Its main limitation is that the branch encodes only contact temperature, since composition, load, velocity, and hardness are constant across the six runs; cross-composition and cross-load generalisation are therefore not yet validated. Extending to further iron-based hardfacing systems and a broader parameter range remains future work.

CONCLUSIONS

A combined experimental–computational approach to predicting friction is presented for a eutectic Fe–Mn–C–B–Si–Ni–Cr hardfacing coating under dry sliding conditions. In ball-on-disc tests, the mean friction coefficient fell from 0.74 at ambient temperature to 0.60 at 300 °C, and varied less at the higher temperature — consistent with stabilisation by the oxide film. The pre-processing step uses a moving window to turn the raw 10 Hz signal into sequences of 12 statistical descriptors. These capture the central tendency, dispersion, distribution shape, and local variability of the COF over a 30 s window. The developed recurrent network (BiGRU + LSTM + attention with a parallel branch of alloy- and process-descriptors) achieved a pooled $R^2 = 0.911$ and RMSE = 0.0312 on entirely withheld test data. Under a fair five-seed evaluation on the same split, its accuracy is statistically equivalent to a strong stacked-LSTM baseline (pooled test R^2 0.9113 ± 0.0005 vs 0.9111 ± 0.0050). Still, the proposed model is about an order of magnitude more reproducible (seed-to-seed s.d. 0.0005 vs 0.0050) and adds interpretable attention weights and a descriptor branch, making the architecture directly extensible to multiple alloys and operating conditions. The model maintains its predictive accuracy across two qualitatively different friction regimes — mechanical (abrasive) wear at ambient temperature and oxidative wear (glaze-layer regime) at 300 °C — without retraining, confirming the value of jointly encoding the dynamic tribological signal with material and process parameters in a single architecture. In a five-seed ablation study, each component affected a different operating regime. Attention pooling gave a small gain at the stationary 300 °C regime, while the alloy- and process-descriptor branch improved both the ambient regime and the pooled accuracy. The bidirectional GRU front-end mainly reduced run-to-run variance. Without attention pooling, the model lost accuracy at 300 °C, and without the descriptor branch, it lost accuracy at ambient temperature. Both the full model and the unidirectional variant avoided these losses. Still, the full bidirectional model showed about six-fold lower seed-to-seed variance, which is why the bidirectional front-end was retained. The absolute differences are small and partly within the seed spread, so component contributions are reported as mean ± standard deviation. The proposed methodology enables reliable predictions with the limited amount of data typical of materials and tribological research. For the alloy and operating window investigated here, the methodology is a candidate approach. After validation of additional iron-based hardfacing systems, alloy compositions, contact loads, and sliding velocities, it could be applied to friction

prediction in related coatings and used in tribotesters and other mechanical systems for early-failure warning.

Friction is a nonlinear and non-stationary process, and the recurrent network learns the temporal dependencies directly from the recorded signal, without assuming their form in advance. This makes it possible to replace an assessment based on averaged values with an online prediction of the friction pair's state. The resulting solution is not limited to the alloy studied here and can be applied to other problems in machine condition monitoring. Further work is planned to extend the measurement matrix by additional repetitions and temperature levels, to introduce variable load and sliding velocity, and to cover further iron-based alloys.

REFERENCES

1. Holmberg K, Erdemir A. Influence of tribology on global energy consumption, costs and emissions. Friction, 2017;5(3):263–284. https://doi.org/10.1007/s40544-017-0183-5
2. Stachowiak GW, Batchelor AW. Engineering tribology. 4th ed. Oxford: Butterworth-Heinemann; 2014.
3. Yüksel N, Şahin S. Wear behaviour–hardness–microstructure relation of Fe–Cr–C and Fe–Cr–C–B based hardfacing alloys. Mater Des, 2014;58:491–498. https://doi.org/10.1016/j.matdes.2014.02.032
4. Lu L, Soda H, McLean A. Microstructure and mechanical properties of Fe–Cr–C eutectic composites. Mater. Sci. Eng. A, 2003;347(1–2):214–222. https://doi.org/10.1016/S0921-5093(02)00607-7
5. Szala M, Walczak M, Pałka T, Kowal M, Nowak WJ. Comparison of cavitation erosion and sliding wear resistance of welded CoCrWC and NiCrBSi hardfacings, AISI 316L stainless steel, and S235JR mild steel. Adv Sci Technol Res J, 2025;19(11):275–291. https://doi.org/10.12913/22998624/209577
6. Pashechko M, Dziedzic K, Mendyk E, Jozwik J. Chemical and phase composition of the friction surfaces Fe–Mn–C–B–Si–Ni–Cr hardfacing coatings. J. Tribol., 2018;140(2):021302. https://doi.org/10.1115/1.4038981
7. Pashechko M, Dziedzic K, Stukhliak P, Barszcz M, Borc J, Jozwik J. Wear resistance of eutectic welding coatings of iron-based Fe–Mn–C–B–Si–Ni–Cr at increased temperature. J Frict Wear, 2022;43(1):90–94. https://doi.org/10.3103/S1068366622010135
8. Dziedzic K, Jozwik J, Barszcz M, Gauda K. Wear characteristics of hardfacing coatings obtained by tungsten inert gas method. Adv Sci Technol Res J, 2019;13(4):8–14. https://doi.org/10.12913/22998624/112405
9. Torres H, Varga M, Adam K, Rodríguez Ripoll M. The role of load on wear mechanisms in high-temperature sliding contacts. Wear, 2016;364–365:73–83. https://doi.org/10.1016/j.wear.2016.06.025
10. Blau PJ. Elevated-temperature tribology of metallic materials. Tribol Int, 2010;43(7):1203–1208. https://doi.org/10.1016/j.triboint.2010.01.003
11. Saravanan S, Hattimare AR, Mahadik AB, Singh A, Kumar UN, Annamalai AR. Effect of heat treatment on high-temperature tribological behaviour of WE54 alloy: an experimental study. J Manuf Mater Process, 2025;9(9):304. https://doi.org/10.3390/jmmp9090304
12. Pearson SR, Shipway PH, Abere JO, Mayfield RAA. The effect of temperature on wear and friction of a high-strength steel in fretting. Wear, 2013;303(1–2):622–631. https://doi.org/10.1016/j.wear.2013.03.048
13. Mohammed AJ, Mohammed AS, Mohammed AS. Prediction of tribological properties of UHMWPE/SiC polymer composites using machine learning techniques. Polymers, 2023;15(20):4057. https://doi.org/10.3390/polym15204057
14. Ning H, Chen F, Su Y, Li H, Fan H, Song J, et al. Modelling and prediction of tribological properties of copper/aluminium-graphite self-lubricating composites using machine learning algorithms. Friction, 2024;12(6):1322–1340. https://doi.org/10.1007/s40544-023-0847-2
15. Yin F, He Z, Nie S, Ji H, Ma Z. Tribological properties and wear prediction of various ceramic friction pairs under seawater lubrication conditions of different medium characteristics using CNN-LSTM method. Tribol Int, 2023;189:108935. https://doi.org/10.1016/j.triboint.2023.108935
16. Hochreiter S, Schmidhuber J. Long short-term memory. Neural Comput, 1997;9(8):1735–1780. https://doi.org/10.1162/neco.1997.9.8.1735
17. Cho K, van Merrienboer B, Gulcehre C, Bahdanau D, Bougares F, Schwenk H, et al. Learning phrase representations using an RNN encoder–decoder for statistical machine translation. In: Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP); 2014 Oct 25–29; Doha, Qatar. Stroudsburg (PA): ACL; 2014, p. 1724–1734. https://doi.org/10.3115/v1/D14-1179
18. Luong T, Pham H, Manning CD. Effective approaches to attention-based neural machine translation. In: Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing (EMNLP); 2015 Sep 17–21; Lisbon, Portugal. Stroudsburg (PA): ACL; 2015, p. 1412–1421. https://doi.org/10.18653/v1/D15-1166

19. Skublewska-Paszkowska M, Powroźnik P, Rejdak R, Nowomiejska K. Application of convolutional gated recurrent units U-Net for distinguishing between retinitis pigmentosa and cone-rod dystrophy. Acta Mech Autom, 2024;18(3):505–513. https://doi.org/10.2478/ama-2024-0054
20. Kolev M, Petkov V, Lazarova R, Petkov V, Kolev K, Uzun S. Investigation of wear behaviour and LSTM-based friction prediction in Cr/nanodiamond-coated Al10Cu alloys. Alloys, 2026;5(1):8. https://doi.org/10.3390/alloys5010008
21. Xiang D, Wang S, Wang D, Hu Y, Yan L, Sim RK, et al. Applications of machine learning in tribo-informatics: current status and future perspective. Wear, 2026;598:206757. https://doi.org/10.1016/j.wear.2026.206757
22. Li Y, Liu B, Chang Z, Zhang Z, Ouyang W, Liu Z, et al. Machine learning in tribology: a review on framework, case studies, and future perspectives. J Manuf Process, 2026;159:243–270. https://doi.org/10.1016/j.jmapro.2026.01.029
23. Pastewka L, Vakis AI, Eder SJ, Aghababaei R, Almqvist A, Carbone G, et al. Modelling in tribology: recent advances, applications, and open questions. Tribol Int, 2026;218:111326. https://doi.org/10.1016/j.triboint.2025.111326
24. Li S. Deep learning prediction of EHL friction using full surface roughness profiles: an LSTM-MLP approach. Mech Mach Theory, 2026;223:106417. https://doi.org/10.1016/j.mechmachtheory.2026.106417
25. Sabry I, Mourad AHI, Elwakil M. Hybrid LSTM–dense neural network for accurate corrosion rate prediction in friction stir welded flange joints. Eng Appl Artif Intell, 2026;167:113697. https://doi.org/10.1016/j.engappai.2025.113697
26. Song Y, Li S, Bi E, Cao X. Non-invasive friction-state prediction for fluid dynamic seals via dual-attention enhanced LSTM. Measurement, 2026;276:121486. https://doi.org/10.1016/j.measurement.2026.121486
27. Song J, Dong C, Bai X, Yuan C. Friction-coefficient prediction and wear-state diagnosis of rotary-vane steering gear seals using CNN–LSTM and LightGBM. Tribol Int, 2025;212:110991. https://doi.org/10.1016/j.triboint.2024.110991
28. Yu W, Huang H, Guo R, Yang P. Tool wear prediction based on an attention long short-term memory network with small samples. Sens Mater, 2023;35(7):2321–2335. https://doi.org/10.18494/SAM4509
29. Jayasinghe R, Ramezani M. AI-powered prediction of friction and wear in functionalised epoxy-MWCNT composites. Wear, 2025;582–583:206333. https://doi.org/10.1016/j.wear.2025.206333
30. Tian Y, Zheng B, Khan M, Yang Y. Real-time prediction of wear morphology and coefficient of friction using acoustic signals and deep neural networks in a tribological system. Processes, 2025;13(6):1762. https://doi.org/10.3390/pr13061762
31. Liu X, Lu Y, Li H, Lu X, Wei Z. Surface roughness prediction from the friction-induced noise. Measurement, 2026;258:119115. https://doi.org/10.1016/j.measurement.2025.119115
32. Jardine AKS, Lin D, Banjevic D. A review on machinery diagnostics and prognostics implementing condition-based maintenance. Mech Syst Signal Process, 2006;20(7):1483–1510. https://doi.org/10.1016/j.ymssp.2005.09.012
33. Stott FH. The role of oxidation in the wear of alloys. Tribol Int, 1998;31(1–3):61–71. https://doi.org/10.1016/S0301-679X(98)00008-5
34. Goodfellow I, Bengio Y, Courville A. Deep learning. Cambridge (MA): MIT Press; 2016.
35. Bahdanau D, Cho K, Bengio Y. Neural machine translation by jointly learning to align and translate. In: 3rd International Conference on Learning Representations (ICLR); 2015 May 7–9; San Diego (CA), USA. arXiv:1409.0473. https://doi.org/10.48550/arXiv.1409.0473