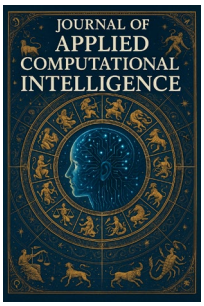
UNIVERSITY OF WEST ATTICA

eISSN: 3136-6937

Journal of Applied Computational Intelligence

Vol. 4, No. 2, April, 2026



LSTM Model Enhanced by Wavelet Decomposition and Force Displacement Stage Alignment for Interlock Estimation in Self Piercing Riveting

Maximilian Müller[1,*], Hannah Wagner[1], Lena Pichler[1]

[1] *Faculty of Mechanical and Industrial Engineering, TU Wien, Vienna, 1040, Austria*

Received 04 February 2026; Revised 09 June 2026; Accepted 10 June 2026; Published 14 June 2026

Abstract

Completed self-piercing riveting strokes contain force-displacement patterns associated with final interlock, while variation in the durations of mechanical stages shifts corresponding features along the source coordinate. This study evaluates a stage-referenced long short-term memory estimator enhanced by two-level Haar decomposition. Four process intervals are mapped to equal displacement-coordinate blocks of 24 positions, yielding a 96-position aligned force sequence. A coarse approximation and two reconstructed detail channels form the recurrent input, and hidden states at the four stage endpoints feed a scalar regression head. The numerical benchmark contains 960 independently generated curves, divided into 600 training, 180 validation and 180 test records. The reference interlock is determined by expansion and lower-sheet descriptors within a declared smooth mapping. Across three training seeds, the combined model attains mean absolute error 0.0223 mm, root mean squared error 0.0288 mm and coefficient of determination 0.9514. Its mean absolute error is 65.96% lower than that of the raw-curve LSTM and 24.85% lower than that of the aligned single-channel LSTM. Aligned ridge regression reaches 0.0156 mm, showing that this constructed target also supports a simpler, more accurate estimator. A normalized-stroke boundary shift raises combined-model error to 0.0785 mm, and halving original sample density gives 0.0241 mm; added force noise yields 0.0700 mm. These findings separate stage correspondence, multiscale representation and input-quality sensitivity within a controlled curve population. Measured riveting records with independently sectioned interlock references are necessary to assess physical transfer.

Keywords: Self-Piercing Riveting, Interlock Estimation, Force–Displacement Alignment, Wavelet Decomposition, Long Short-Term Memory

1. Introduction

Data-driven manufacturing assessment uses recorded process responses to estimate quality characteristics that are difficult to inspect continuously [1], creating a computational link between production signals and geometric outcomes. Sequence-learning methods represent the order and persistence of these responses through shared temporal features [2], supporting analysis of complete manufacturing strokes. Multiscale signal representations distinguish gradual evolution from local variations in measured curves [3], providing complementary descriptors for regression. Alignment methods establish correspondence between process intervals whose durations vary across observations [4], allowing sequence models to compare events with similar operational meaning. These ideas are relevant to self-piercing riveting, where the force response changes as material is clamped, pierced, displaced, and compressed around the rivet. The resulting

* Corresponding author: muller.ma@mec.tuwien.ac.at

http://dx.doi.org/10.68406/jaci.2026.vol4iss2no12:166-179

➢ This is an open access article under the CC BY-NC-SA 4.0 license (https://creativecommons.org/licenses/by/4.0/).

© Authors retain all copyrights.

interlock is a geometric feature of the completed connection. Estimating this feature from the full force–displacement record requires attention to both force evolution and the location of successive process stages.

A curve with a longer piercing interval has its later features at different positions from a curve with a shorter interval. Representation Learning for Variable-Length Processes [5] employs a modification of such processes as an explicit modelling factor. Uniformly sample the total displacement and save the original path to associate different mechanical phenomena with the same sequence index. A recurrent model must learn the geometric relationship and address the problem of position variation. Multiresolution learning [6] is another way to organise the trends of coarse forces and local fluctuations in separate channels. The value of this division is relative to the position of the division. Wavelet coefficients calculated before the construction of correspondence mix stage-duration effects with local signal variations. Therefore, a shared-stage coordinate can be used to organise the interpretation of the impact of scale-specific information on an interlock estimate.

Do not confuse computational assessment with the difficulty of a regression model. Matched-input benchmarking [7] aligns and matches the sequences for use in the same test observations and training budget. If most of the variations are already aligned in the input, a simple regressor can be employed to model the smooth target relationship. Reliability-oriented prediction studies [8] also consider the reporting errors caused by input perturbations in the clean-data average. At present, some research has been conducted on raw-curve LSTM, wavelet LSTM, aligned single-channel LSTM, the combined representation and aligned ridge regression. Independent Training Seeds Quantify Optimisation Variability. Perturbations in force and boundary shifts are used to study different causes of input degradation. Thus, the accuracy of aggregates can be attributed to explicit modifications in data preparation and conditions for increased sensitivity of the representation.

This study evaluates a stage-referenced wavelet LSTM for completed-stroke interlock estimation. Recorded boundaries transform unequal mechanical intervals into four aligned 24-position blocks; a two-level Haar representation separates coarse and detailed force components before recurrent encoding. Hidden states at stage endpoints support millimetre-scale scalar regression. A controlled parametric curve population varies stage length, force gain and local response while retaining an explicit reference interlock. Matched baselines and fixed-model perturbations distinguish coordinate alignment, channel construction and input-quality effects. Section 2 reviews process-curve regression, multiscale representations and temporal alignment. Section 3 defines reference generation, displacement mapping, Haar reconstruction, recurrent readout and training. Section 4 reports independent population splits, baseline rankings, stage and channel interventions, robustness and model dimensions. Section 5 summarizes three quantitative findings and the validation scope of the synthetic reference.

2 Related Work

2.1 Process Curve Based Geometric Quality Estimation

Learning from manufacturing curves requires selecting a representative that retains the information related to the output. Process-oriented regression methods [9] aim to preserve the connections among force level, displacement progression and local changes in slope. A single number is the highest point of a whole piece of music, and a series retains the positions and shapes of various alterations. All parts of a rivet record are used to locate the start of resistance, penetration, radial expansion and end compression. Analysis of Representation [10] shows the quantity of predictive information remaining after each compression step. Therefore, the current comparison uses a full-resampled curve in the simplest sequence baseline. A linear model of the aligned profile is also used to assess the contribution of coordinate preparation independently from recurrent nonlinearity. All the inputs are derived from the same force record and use the same set of boundary metadata.

The relationship of a process signal with a geometric target is determined by the origin of the target definition. Computational Benchmark Design [11] sets up reference variables and nuisance variables independently to attribute any improvement in the constructed problem to them. A family of parametric curves can be used to alter the stage duration and output-related force independently. Therefore, control can be exercised directly on the test alignment by moving the stage location independently of the reference interlock. Regression evaluation methods [12] can also provide both the mean and tail errors of the model when it is to be applied to dimensional inspection. A small number of relatively large errors can be due to a small mean error. Mean Absolute Error, Root Mean Squared Error, Signed Bias and the Upper Absolute-Error Quantile are thus used together. Together, they are the normal deviations, large dispersions and problem areas of the reference range.

2.2 Stage Alignment and Multiscale Recurrent Modelling

Stage Alignment offers an intuitive division of the input record to set up correspondence. Sequence registration methods [13] map all the process intervals to a common coordinate system prior to learning relationships among observations. A map is the arrangement of stages, and a set number of samples are collected at each stage. Thus, the

density of representation is reduced, and a short stage has the same number of model positions as a long stage. Recurrent learning [14] can add a way to combine information from the resulting sequence and keep a hidden state of previous intervals. Extraction of hidden states at various stages can also reveal the influence of earlier process information on the final result. Boundary-marker perturbation is analysed independently to quantify the impact of alignment quality under a controlled alteration of the map.

Wavelets are used to display the various components of a signal at different scales. Multiscale feature-learning methods [15] propose to provide these components separately for the model and thus can receive different learned transformations for coarse evolution and local response. The sum of the components in a full reconstruction is equal to the original signal; therefore, there is no alteration. Robust sequence analysis [16] can also be used to investigate the effect of small-scale channel perturbations on sensitivity. The two-level Haar basis has the attributes of orthogonality and a small support in this study. It can be grouped according to the selected stage sample number and averages the groups within the stage. The added value of the added-on explicit scale separation over the single-channel aligned input is shown here. Channel suppression then investigates the influence of a fitted model on each reconstructed component under a specific intervention.

3 Stage Referenced Wavelet LSTM Representation of Riveting Curves

3.1 Process Stages and Displacement Coordinate Mapping

The numerical reference is a parameterized four-stage force-displacement family. Expansion descriptor z, lower-sheet descriptor t and nuisance response are sampled independently on the unit interval. A declared smooth mapping of z and t defines interlock h in millimetres; the compact notation below uses s for the sine of twice pi times z. The force-curve generator also introduces z through the flaring-stage slope and local shape, creating an observable relationship between the process signal and reference target. Nuisance variation and force gain change the observed curve while the target mapping retains its stated descriptors. This construction defines a controlled regression benchmark for stage correspondence and representation, with physical interlock measurement reserved for independent validation.

$$h = 0.20 + 0.50z + 0.08t + 0.03s \tag{1}$$

Each complete stroke contains initial loading, penetration, flaring and final compression in that order. Relative lengths follow a Dirichlet draw with concentrations 10, 14, 14 and 10, scaled by 0.8 and increased by 0.05 for each stage. The resulting positive fractions sum to one. Total displacement spans 4.8-6.8 mm and each original force curve contains 256 equally spaced source samples. Recorded boundaries accompany each curve as process metadata. Within a stage, local coordinate u is computed from displacement d, stage start b and stage width w as shown below. This coordinate identifies mechanical progress within an event independently of its absolute duration in the full stroke.

$$u = (d - b)/w \tag{2}$$

Each recorded stage is mapped to a unit local interval and sampled at 24 left-closed, right-open positions. Four blocks therefore form a chronological 96-position aligned sequence. Linear interpolation evaluates the original 256-sample force record at the corresponding displacement coordinates. In the mapping below, F denotes the force function reconstructed on the measured displacement axis and F-tilde denotes its stage-aligned value. Training, validation and test curves use the same transformation with their recorded stage boundaries. Total stroke displacement remains available as source metadata, while the recurrent estimator receives the aligned force channels. This design makes boundary-location perturbation a distinct intervention on the input-coordinate map.

$$\tilde{F}(u) = F(b + uw) \tag{3}$$

Figure 1 is the four-process interval, the original displacement record and the common-stage coordinate of the estimator. The original boundaries and the conversion of unequal interval widths into equal blocks of 24 model positions are shown in the schematic. It can also remember the connection between a local aligned position and the source displacement to trace the transformed signal back to the original record. Initial loading and penetration are in the first two blocks; flaring and compression are in the last two. Corresponding landmarks have the same location in all strokes. All four levels of the model resolution and order are the same. Use the recorded boundary markers of the benchmark as the alignment input, and then change these markers in sensitivity tests to explore the effect of stage location.

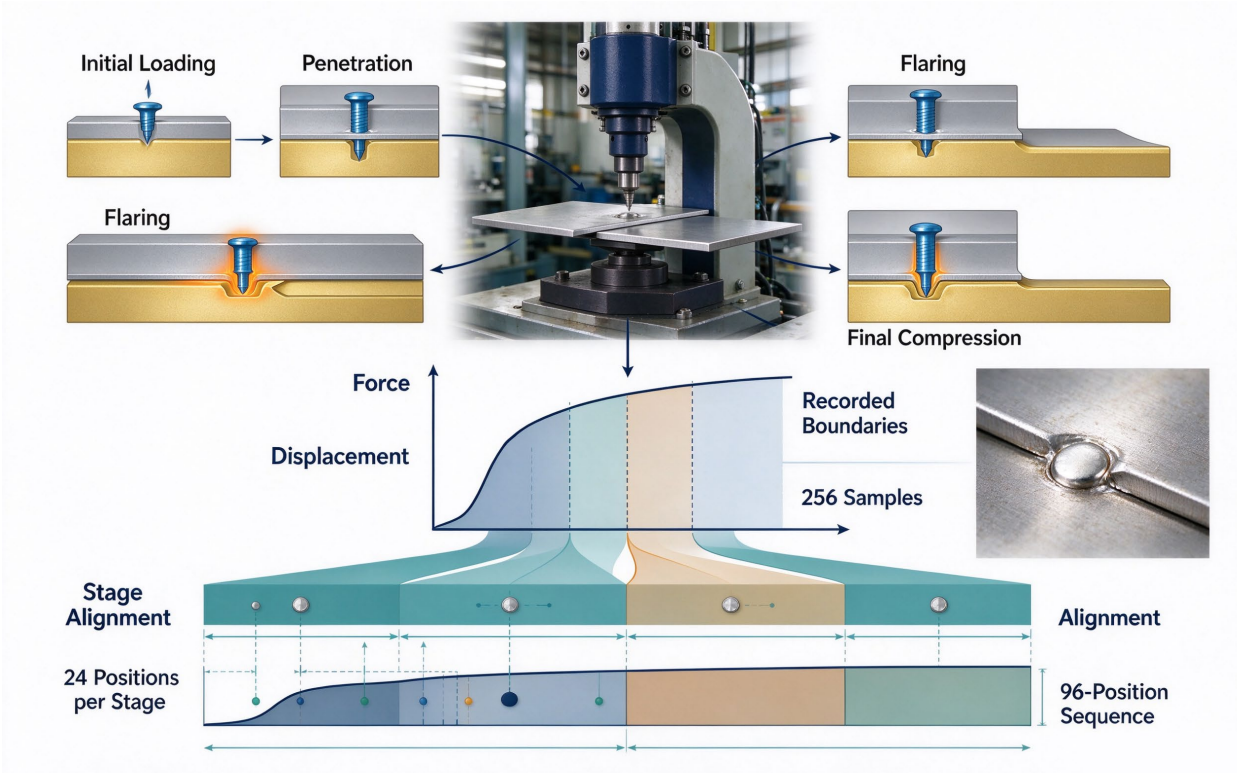


Figure 1. Process stage reference and displacement mapping for aligned force sequence construction.

3.2 Wavelet Decomposition and Recurrent Encoding

A two-level Haar decomposition is applied to the 96-position aligned force sequence. Pairwise averages, repeated at both positions, form the first approximation. Means over groups of four, repeated at each of the four positions, form the second-level coarse approximation A2. The difference between the original signal and the pair approximation gives first-level reconstructed detail D1; the difference between the pair and four-point approximations gives second-level detail D2. At every aligned position, A2, D1 and D2 sum to the original force value. Each group of four lies within one 24-position stage block. Separate channels expose coarse force evolution and shorter-range changes to the recurrent input transformation while retaining the aligned signal information.

Figure 2 is the multi-scale recurrent encoding path. The aligned curve is in the Haar branch and outputs three reconstructed channels of the same 96-position length. Normalise the means and standard deviations of the training-fitted channels before recurrent processing. An LSTM with 24 hidden units is used to process the channels sequentially. Hidden states at positions 24, 48, 72 and 96 are concatenated to keep track of the four-stage block information. A 32-unit rectified-linear hidden layer and a scalar-output layer are employed for processing the 96-dimensional descriptor. All of the stage-specific readouts are connected to the same recurrent encoder; therefore, the earlier information is still available in both the temporal state propagation and explicit stage sampling.

Input tensor X stacks the second-level approximation and two reconstructed Haar details as three force channels across 96 positions. Channel means and standard deviations are calculated only from training curves and their sequence positions; a 0.000001 minimum standard deviation stabilizes scaling. The same training-derived values are reused for validation and test curves. Single-channel LSTM baselines receive equivalently training-scaled force sequences. This preprocessing keeps comparison groups separated and lets the recurrent model assign different weights to coarse force magnitude, four-position variation and pair-level variation. Channel-suppression diagnostics subsequently alter one standardized input channel at a time in a fixed fitted network.

$$X = [A_2, D_2, D_1] \tag{4}$$

The combined estimator contains 5,921 trainable parameters: 2,784 in its 24-unit LSTM and 3,137 in the regression head. Single-channel recurrent variants contain 5,729 parameters because their input matrices receive one force channel compared with three in the combined model. All recurrent variants share hidden width, four stage-index readouts, optimizer, minibatch and epoch budget. For aligned inputs, hidden states at positions 24, 48, 72 and 96 summarize the ends of the four stage blocks. The same indices on uniformly resampled raw inputs represent quarters of total displacement, whose locations can differ from physical stage boundaries. The wavelet-only comparator applies the same Haar operations to that uniformly sampled complete-stroke sequence.

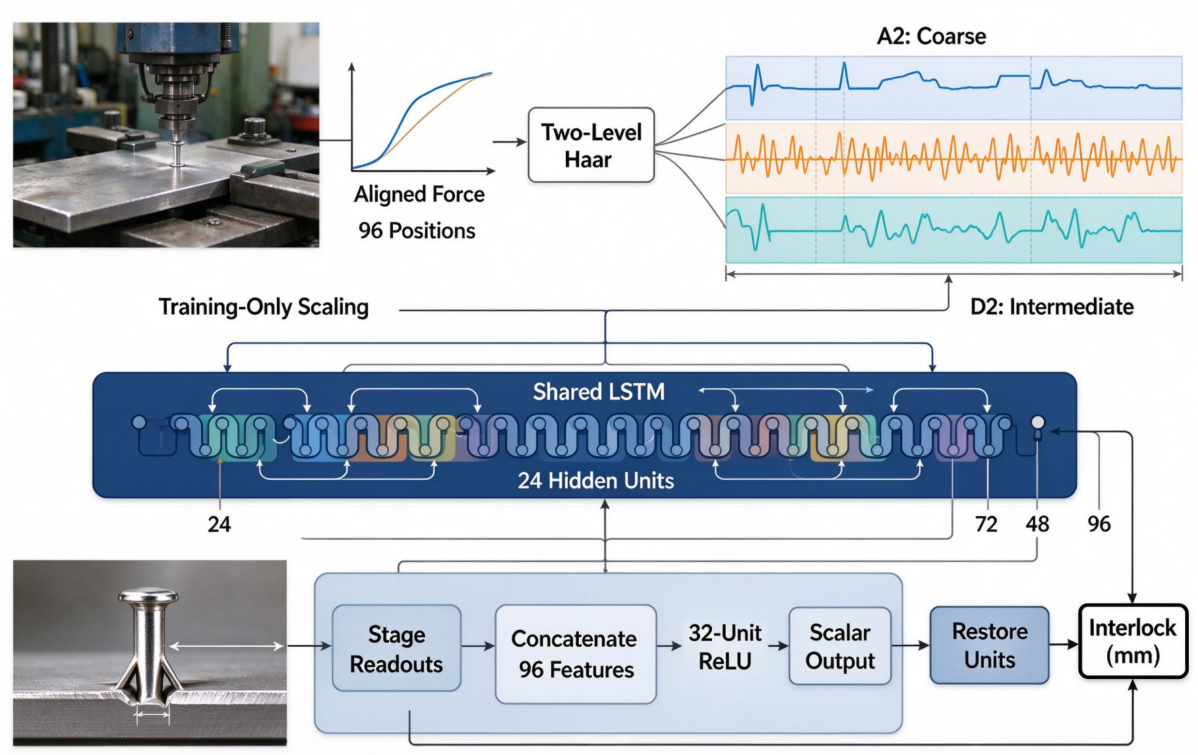


Figure 2. Haar component construction and stage specific LSTM encoding for interlock regression.

3.3 Interlock Regression and Training Organisation

The regression head estimates interlock on a training-standardized target scale. The mean and standard deviation of the 600 training interlocks are retained for validation and testing, and the inverse transformation returns every prediction to millimetres. In the expression below, H concatenates hidden states from the four stage endpoints and g is the two-layer regression head. Error, signed bias and upper-tail measures therefore share one physical unit across model variants. Fixed-model stage and channel suppression act on the standardized input tensor, while the fitted recurrent encoder, head and inverse target transformation remain unchanged. This distinction makes intervention sensitivity interpretable separately from a retrained representation.

$$\hat{h} = \mu + \sigma g(H) \tag{5}$$

The population comprises 960 independent curves generated with seed 2810. A separate seed-2811 permutation assigns 600 to training, 180 to validation and 180 to testing. Target standardization precedes optimization of mean squared prediction error. Adam uses learning rate 0.003, gradient-norm limit one and 60 epochs. Each epoch partitions the 600 training curves into ten minibatches after a new permutation. Validation mean squared error is evaluated after restoring millimetre units, and its minimum selects a checkpoint independently for fitting seeds 281, 282 and 283. Training-only input and target statistics remain fixed across the held-out partitions. This split prevents positions from a single generated stroke from crossing assessment boundaries.

$$L = E[(\hat{y} - y)^2] \tag{6}$$

The numerical curve family maintains continuity at stage transitions before observational perturbation is added. Initial loading increases from a 1 kN level through a 5 kN rise with a thickness-dependent increment. Penetration adds a 9 kN rise and a smooth nuisance-dependent variation. The flaring interval contains a rise of 6 kN plus 8 kN multiplied by the expansion descriptor, together with two sinusoidal components. Final compression continues from the preceding endpoint with a rise between 9 and 11 kN. A force gain drawn between 0.85 and 1.15 multiplies the complete profile. Independent force perturbations with standard deviation 0.08 kN are then added at the 256 original positions. The gain and stage-length variation challenge the predictor to distinguish reference-related shape changes from observational and coordinate variability.

Evaluation calculates mean absolute error, root mean squared error, signed bias, R-squared and the 95th percentile of absolute error on the same 180 test curves for each fitted seed, then averages the three seed-level metrics. The aligned ridge baseline uses all 96 aligned force samples, standardized with training-only statistics, and an unpenalized intercept. Its penalty is selected from 0.01, 0.1, 1, 10 and 100 by validation error. Paired robustness interventions add 0.4-kN force noise, shift interior normalized stage boundaries by up to 0.02, or retain every second original source sample before

interpolation. Reference interlocks and fitted combined-model parameters stay fixed during these input-quality comparisons.

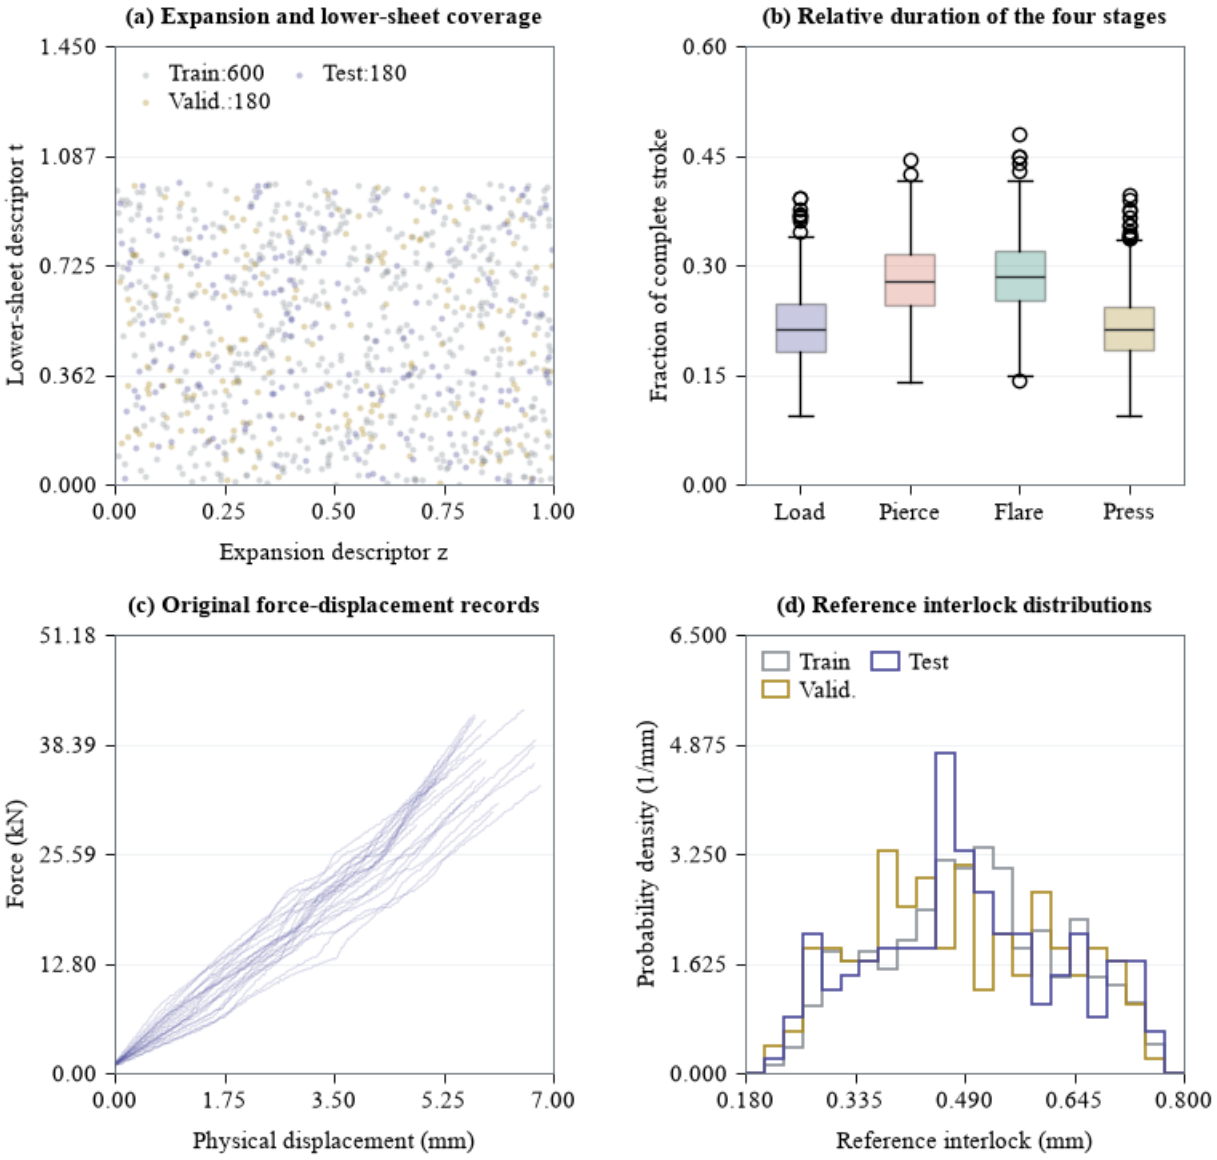


Figure 3. Parametric curve population and reference coverage. (a) Expansion and lower sheet descriptors. (b) Relative stage lengths. (c) Force displacement profiles. (d) Interlock reference distributions.

4 Alignment Effects and Reliability of Interlock Estimation

4.1 Curve Population and Prediction Baselines

Figure 3 is the numerical population, and changes in the input curves are connected to the declared interlock references. The 960 independent strokes are continuous-variable expansions of the lower-sheet, nuisance, force-gain and total-displacement descriptors. Controlled benchmark design [17] separates these factors to evaluate the impact of changes in stage extent on alignment. A reference interlock is provided by the expansion and the lower-sheet descriptor, and the stage lengths and the total force gain are independent. A set of benchmark problems was released to promote research. Order Stage Markers are added to the input metadata, and only one reference value is required for the whole stroke. Population construction is performed prior to the independent training-validation-test split, and all of the comparison predictors receive the same partition.

Figure 3(a) is the distribution of the normalised expansion and lower-sheet descriptors. Together, they specify the reference interlock range and the list of similar interlocks with different constituent descriptors. Figure 3(b) is the four relative stage lengths of the population. Modifications to the features of the characteristic force occur along the original displacement axis and retain the chronological order of the stages. Distribution Coverage Analysis [18] is used to address the above problems of different learning requirements for target coverage and coordinate variability. Independent stage-

length construction is used to align the relevant part of the force law for a single stage. The sum of the lengths of all the curves is 4.8 mm and 6.8 mm.

Figure 3(c) is an example force-displacement curve before alignment. Curves are convenient for organisation of group stages and have different force gains and stage transition features. Figure 3(d) shows the reference interlock distributions of the three data partitions. Ensure that the range of all partitions is the same as the target range. Leakage-aware preprocessing [19] only uses the training set to calculate input scaling and target scaling, then applies them to the validation and test data unchanged. The test set is not changed in model comparison and perturbation tests. Stage Markers are the input curves, and only used for the construction of the alignment map. The first requirement of the main evaluation is that they must be available, and their quality should not change in the boundary-perturbation experiment.

Figure 4 is the prediction accuracy of the five representations and regressors on the common test set. Compare the Contributions of Alignment, Explicit Scale Separation and Recurrent Modelling to the Problem. Figure 4(a) is a plot of the predicted interlocks against their required references, and individual observations with offsets and dispersion are retained. The Mean Absolute Error of the combined model is 0.0223mm and the Root Mean Square Error is 0.0288mm. Figure 4(b) shows the distribution and upper tail of absolute errors. Matched evaluation methodology [20] prompts the reporting of these distributions alongside the mean, and the 95th-percentile absolute error of the combined model is 0.0602 mm. The mean absolute errors of the raw LSTM, the wavelet-only LSTM, and the aligned single-channel LSTM are 0.0656 mm, 0.0542 mm, and 0.0297 mm, respectively. The two are compared, and Alignment is the larger of the two representation changes.

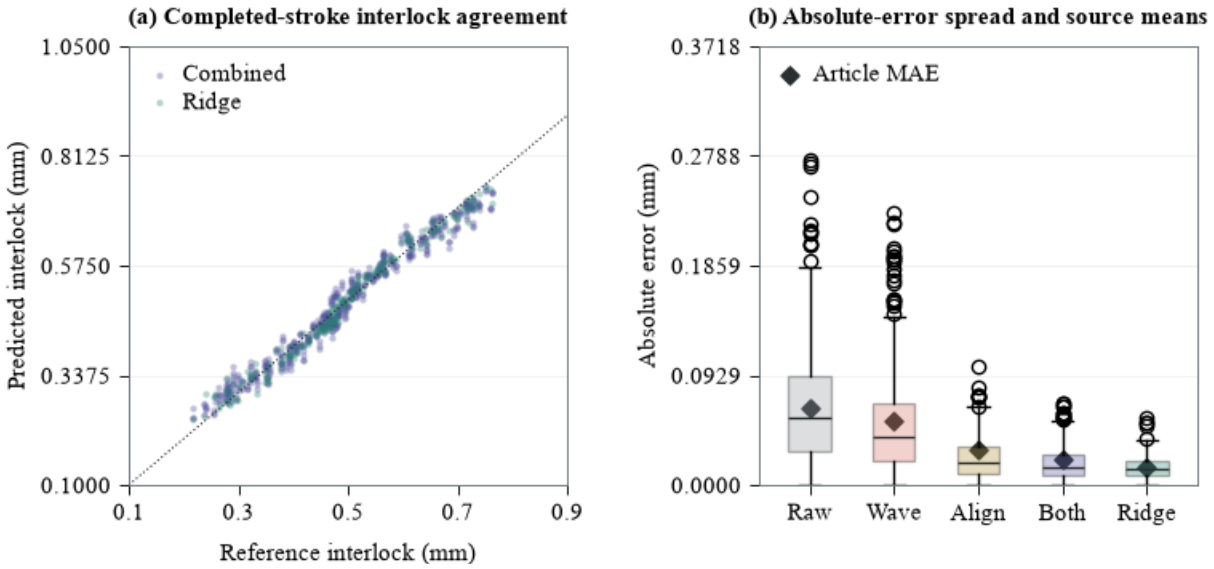


Figure 4. Baseline interlock estimation performance. (a) Predicted and reference interlock. (b) Absolute error distributions.

Aligned ridge regression achieves test MAE 0.0156 mm, RMSE 0.0197 mm and R-squared 0.9774, outperforming the combined LSTM on this declared smooth target. Baseline selection principles [21] support retaining this simpler estimator in the architecture comparison. The combined LSTM has R-squared 0.9514, compared with 0.6210 for the raw-curve LSTM and 0.9240 for the aligned single-channel LSTM. Thus stage alignment substantially improves recurrent prediction, and explicit Haar channels provide a further gain within that model family. The stronger ridge score indicates that the generated interlock mapping is well represented by a regularized linear relationship after correspondence. Its ranking is a property of this parameterized curve and target construction, motivating assessment on independently measured joints.

4.2 Stage Correspondence and Multiscale Contributions

Figure 5 shows the effect of establishing stage correspondence prior to recurrent processing. Uniform sampling of the total displacement assigns model positions according to stroke fraction, and stage alignment assigns positions according to the process stage within an interval. Temporal Registration Analysis [22] is also carried out on the shape of curves and prediction errors. Currently, the aligned single-channel LSTM has reduced the mean absolute error to 0.0297 mm from 0.0656 mm, and only one force channel and the same recurrent hidden width have been employed. Combined, the results are as low as 0.0223 mm. As a result, there is a considerable improvement in the relocation of coordinates and an additional increase due to scaling. The same start curve records and fixed test references are used for comparison. Therefore, all transformations are linked to specific changes in the organisation of input resources for the evaluated target population.

Figure 5(a) is the curve in the original normalised displacement axis, and the recorded stage changes are shown. Due to the relatively narrow stages, the local force pattern is the same at all places. Figure 5(b) is the same curve after mapping to four equal sample blocks. The corresponding intervals are now in the same position range and convenient to compare the changes in that section. Alignment diagnostics [23] are used to keep the source-to-target coordinate map and find the origin of the apparent similarity after transformation in the original signal. Rearrange does not change the order of the samples. It will be convenient to read the stage-index display. The hidden states after the third aligned block consistently summarize the information in the flaring interval, and the same numerical index in an unaligned record can appear at different positions in the process.

Figure 5(c) shows the aligned and unaligned prediction errors in the interlock reference region. As shown in the above figure, the whole increase occurs near the centre of the planned distribution, not at the ends. Figure 5(d) shows the sensitivity to the perturbation of supplied stage boundaries. An independent shift of 0.02 in the normalised stroke increases the mean absolute error and root mean squared error of the combined model to 0.0785 mm and 0.1012 mm, respectively. Input-registration assessment [24] drives the interpretation of this deterioration along with clean alignment gain. Correspondence is the mapping of markers to intervals, and therefore, shifts in markers redistribute force evidence in the fixed blocks. Perturbation is added before resampling and wavelet construction to change both sets of data. Based on the above analysis, the model's quality attribute of boundary localisation has been identified.

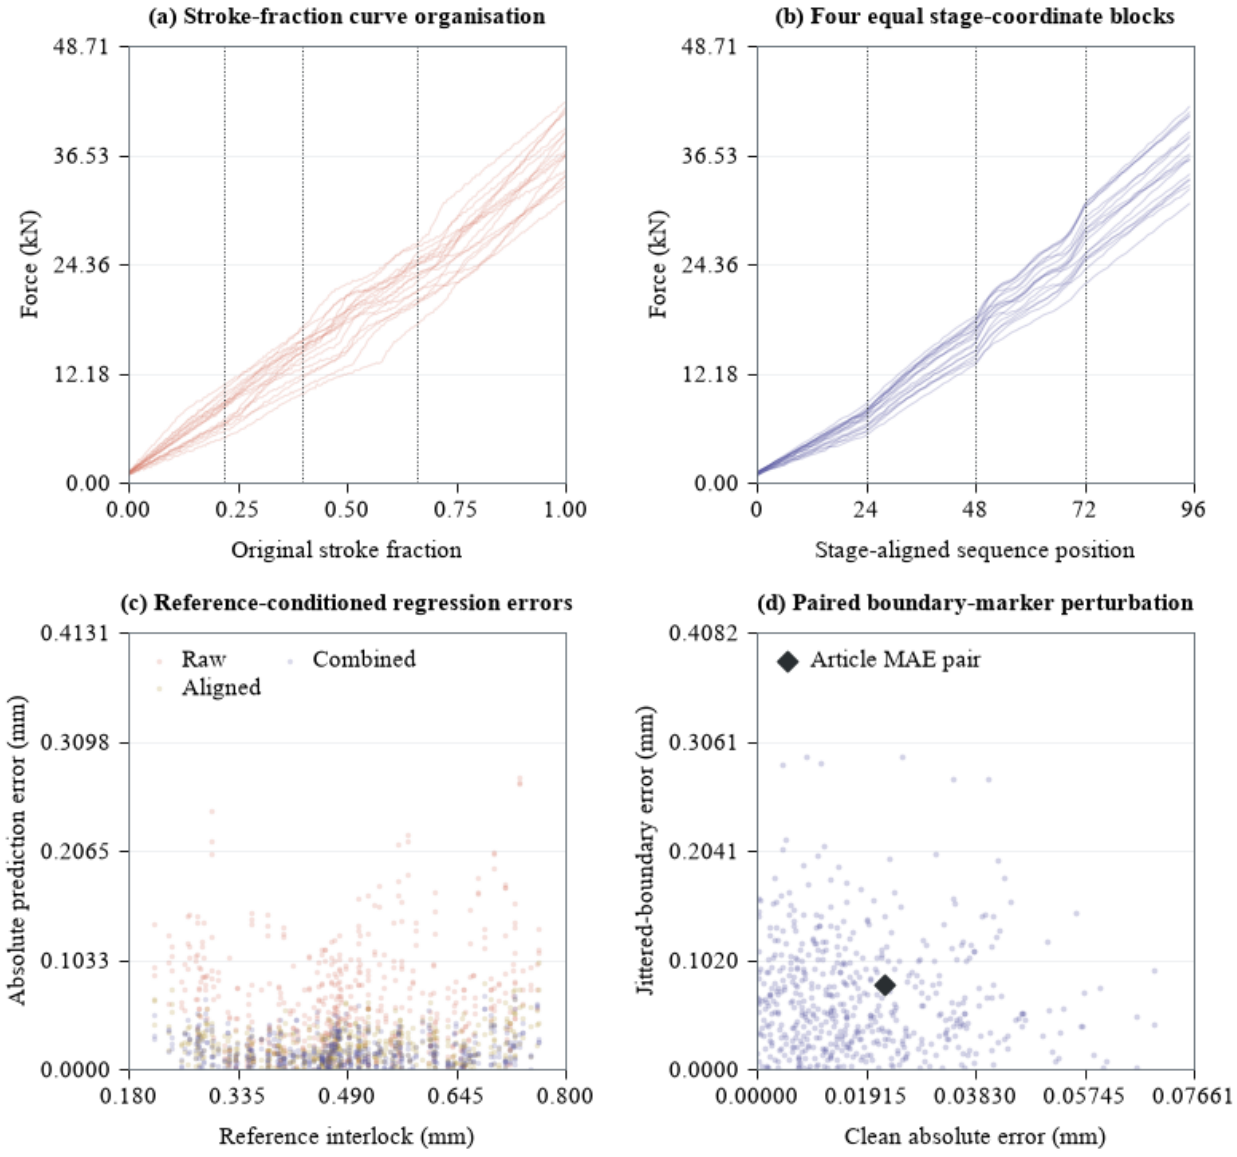


Figure 5. Stage correspondence and boundary sensitivity. (a) Original curve alignment. (b) Stage referenced curves. (c) Error over the reference range. (d) Boundary perturbation response.

Figure 6 compares fixed-model stage suppression with channel and representation analyses. For Figure 6(a), one of four standardized stage blocks is set to zero while the trained network parameters remain fixed. Suppressing the third, flaring-stage block raises MAE to 0.3175 mm; suppressing the first, second and fourth blocks gives 0.0528, 0.1052 and

0.0501 mm. Figure 6(b) records the associated signed bias and upper-tail errors, extending interpretation beyond the mean. Perturbation-based explanation methods [25] treat these scores as sensitivity of the fitted predictor to the specified replacement. The third stage contains expansion-dependent force terms in the generated reference, while zeroing an entire block also changes the joint sequence distribution. Both properties matter when interpreting this intervention.

Figure 6(c) evaluates fixed-model suppression of the approximation, second-level detail and first-level detail channels, yielding MAEs of 0.0620, 0.1396 and 0.0371 mm, respectively. The second-level detail has the strongest suppression response among the three channels. Figure 6(d) places these interventions beside separately retrained representation variants. Model interpretation research [26] motivates distinguishing a fitted-network replacement from a new model optimized with a different input, because the remaining features can be used differently after retraining. On the common test set, the combined LSTM reduces MAE by 65.96% relative to the raw LSTM and by 24.85% relative to the aligned single-channel LSTM. Force-noise sensitivity qualifies the benefit of explicit detail channels under the declared synthetic curves.

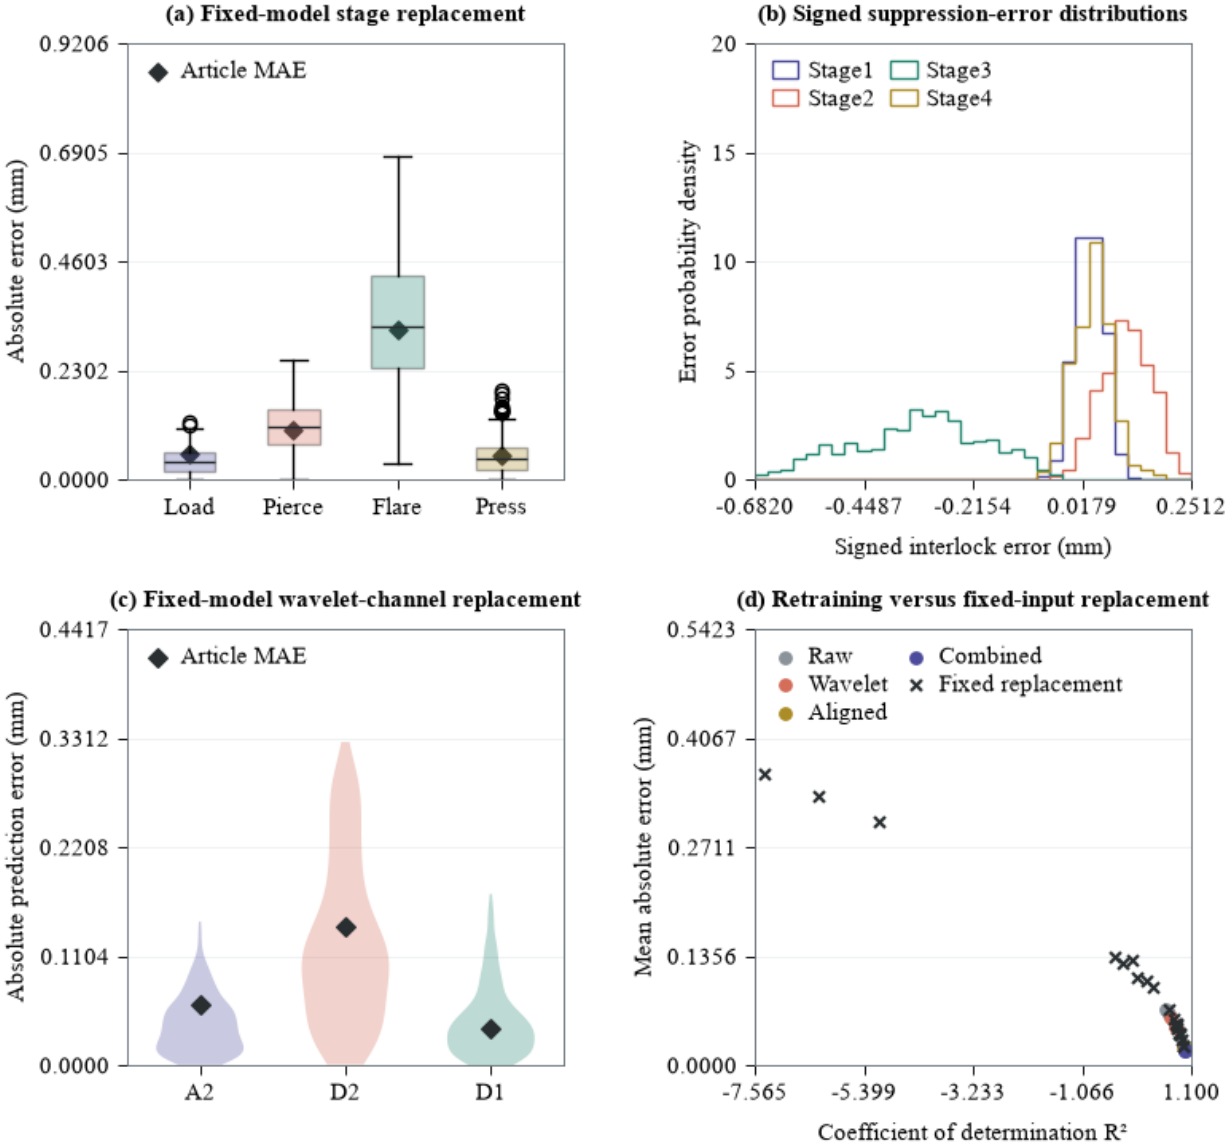


Figure 6. Stage and wavelet contributions. (a) Stage suppression error. (b) Stage suppression bias and tail error. (c) Wavelet channel suppression. (d) Retrained representation comparison.

4.3 Robustness and Computational Assessment

Table 1 presents the clean-data baseline results and the three input-perturbation results of the combined model. The columns are still the corresponding Mean Absolute Error, Root Mean Squared Error, Bias, Explained Variance and Upper-Tail Absolute Error. Reliability-oriented regression analysis [27] will be used to correct the problem of a small standard error bias and a large error dispersion. The average of the three fitted models' metrics for each recurrent row is taken to construct the test curve. Ridge Regression is used to obtain a deterministic estimator in the validation set. Perturbation rows keep the fitted combined models and the unchanged interlock references for paired inspection of

degraded input quality. Noise is added to the original 256 sample locations, boundary shifts alter the alignment map, and sparse sampling modifies the source density prior before reconstruction. The various start points are reasons for the different error effects.

Table 1. Baseline accuracy and input sensitivity for interlock estimation.

Model / input setting	MAE (mm)	RMSE (mm)	Bias (mm)	R^2	95% error quantile (mm)
Raw LSTM	0.0656	0.0806	-0.0037	0.6210	0.1510
Wavelet LSTM	0.0542	0.0663	0.0043	0.7405	0.1213
Aligned LSTM	0.0297	0.0358	0.0108	0.9240	0.0674
Aligned wavelet LSTM	0.0223	0.0288	0.0011	0.9514	0.0602
Aligned ridge	0.0156	0.0197	-0.0017	0.9774	0.0356
Force noise	0.0700	0.0877	0.0131	0.5504	0.1670
Boundary jitter	0.0785	0.1012	-0.0014	0.4008	0.1765
Sparse sampling	0.0241	0.0302	0.0013	0.9466	0.0622

Figure 7 shows the three robustness interventions along with computational and training-variability data. It is to determine which part of the input or model pipeline has been modified by an increase in the observed error. Robustness evaluation [28] keeps the data preparation and network fitting fixed during the isolation of inference-time sensitivity. The force-noise intervention increases the observation variation of fixed-stage markers. Boundary intervention keeps the force samples and modifies their mapping to the stage positions. Sparse-input intervention retains the original record and the markers but reduces the number of samples for interpolation. Therefore, the above tests explore the different failure modes of the same representation. The clean-data model is also used in the comparison to separate the impact of perturbation from changes caused by independent training initialisations. All the absolute errors are in mm at the designated reference interface.

Figure 7(a) shows the error response of an added Gaussian force perturbation with a standard deviation of 0.4 kN. The Mean Absolute Error is 0.0223 and 0.0700 mm, respectively, and the 95th-percentile absolute error is 0.0602- and 0.1670-mm. Noise-sensitive representation analysis [29] is employed to explore the reason for the increase in reconstructed detail channels. Perturbations will not affect the progress of the coarse stage via these paths. Figure 7(b) shows the results of boundary-perturbation and sparse-input. The Mean Absolute Error of the former is 0.0785mm, and that of the latter is 0.0241mm. Sampling-aware sequence evaluation [30] can differentiate between a decrease in source density and a modification of stage boundaries. Store every second original sample to obtain 128 source positions, linearly interpolate to reconstruct the profile, and save the final source endpoint when necessary. The reduction in quantity for the family of smooth curves tested here is relatively small.

Figure 7(c) reports seed-specific combined-model MAEs of 0.0216, 0.0218 and 0.0236 mm, with between-seed standard deviation 0.0011 mm. Repeated-training assessment [31] separates this fitted-model spread from the distribution of errors across held-out strokes. Wavelet-only and aligned single-channel variants have between-seed standard deviations of 0.0067 and 0.0046 mm, respectively. These values describe optimization variability under the common split and validation procedure. The increases associated with boundary perturbation and added force noise exceed the combined-model seed spread in the reported assessment. Stage-marker quality and local force variation therefore merit separate reporting when the estimator is evaluated on measured process records.

Figure 7(d) summarizes the model and preprocessing inventory. The combined recurrent input has shape 96 by 3, and its LSTM has 24 hidden units. Four 24-dimensional endpoint states enter a regression head with a 32-unit hidden layer. Computational analysis [32] distinguishes this 5,921-parameter network from the aligned ridge baseline with 97 coefficients including its intercept. Stage interpolation produces the 96-position coordinate sequence, and fixed Haar groups of two and four yield the three input channels. The ridge baseline obtains lower test error on the declared smooth target despite its smaller parameter count. These counts describe model size and operations; measured processor latency and physical sensor throughput remain distinct assessment quantities.

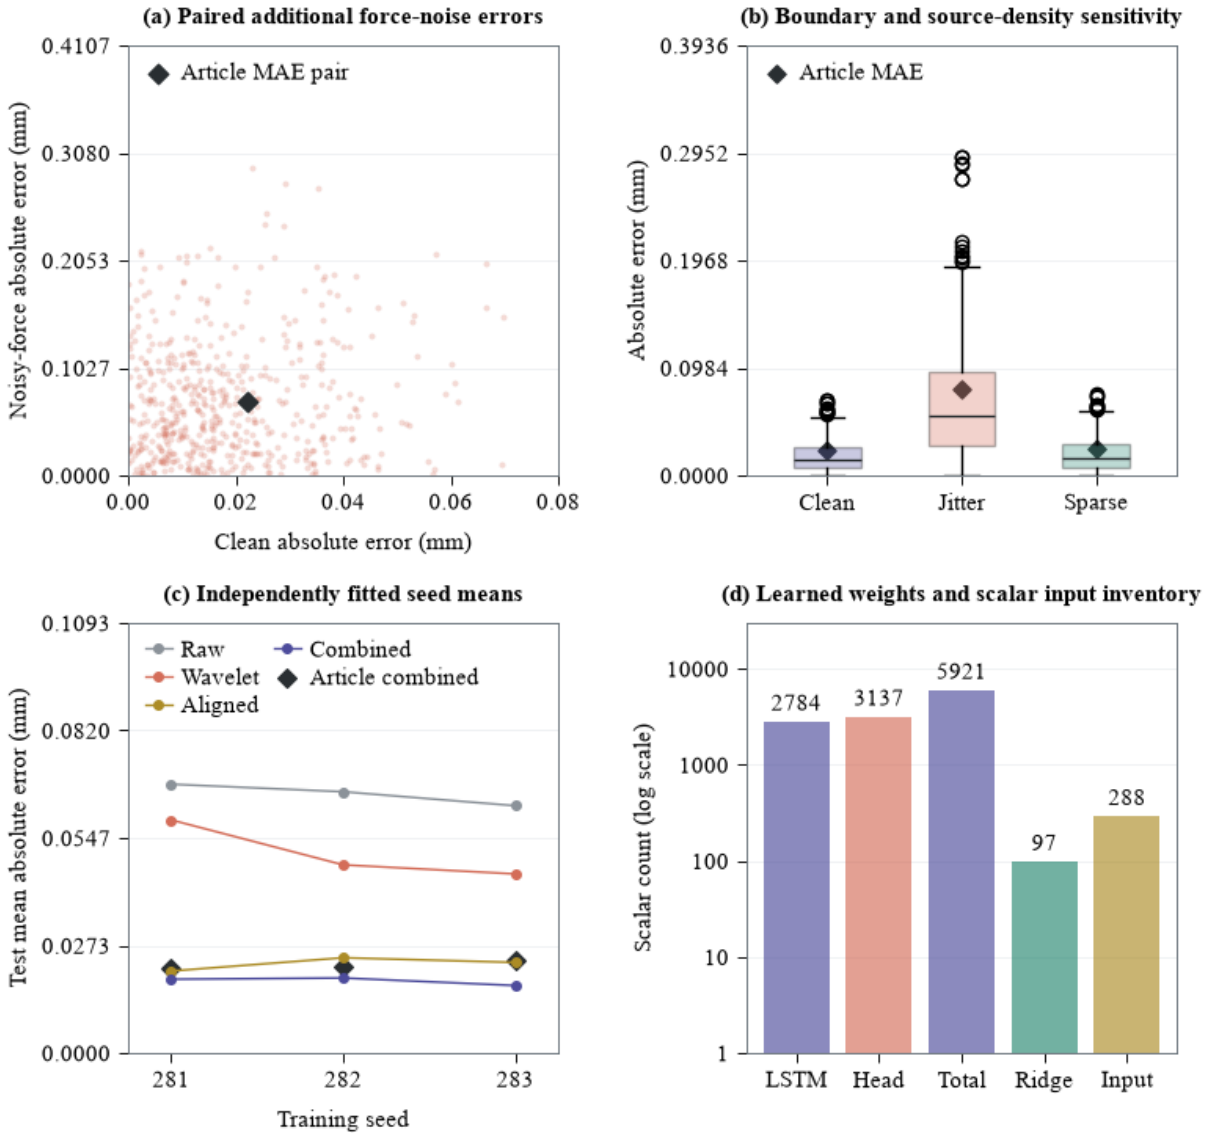


Figure 7. Robustness and computational assessment. (a) Force perturbation. (b) Boundary shifts and sparse sampling. (c) Training seed variability. (d) Model dimensions and preprocessing operations.

5 Conclusion

This study evaluated a stage-referenced wavelet LSTM for completed-stroke self-piercing-rivet interlock estimation. Four mechanical intervals are mapped to equal-length displacement blocks, and a two-level Haar transform separates coarse and detailed force channels. Recurrent states at stage endpoints feed a millimetre-scale scalar regressor. The controlled 960-curve benchmark isolates the contributions of stage correspondence, explicit scale representation and input quality.

First, on 180 independent test curves the combined model attains MAE 0.0223 mm, RMSE 0.0288 mm and R-squared 0.9514 across three seeds; its MAE is 65.96% below raw LSTM and 24.85% below aligned single-channel LSTM. Second, aligned ridge regression achieves lower MAE 0.0156 mm and RMSE 0.0197 mm, showing that the declared target supports a simpler estimator after stage correspondence. Third, shifting stage boundaries by 0.02 of normalized stroke raises combined MAE to 0.0785 mm, additional 0.4-kN force noise raises it to 0.0700 mm, and halving source-sample density gives 0.0241 mm. These results distinguish representation gains, comparator ranking and sensitivity to input quality.

The findings apply to parameterized force-displacement curves with recorded stage boundaries and interlocks generated from a smooth expansion and lower-sheet mapping. The combined model contains 5,921 parameters and processes a fixed 96-position, three-channel sequence. Fixed-model block and channel suppression identify sensitivity under the generated input distribution, while the ridge score limits claims of architecture superiority. Physical assessment should pair measured riveting strokes with independently sectioned interlock values, documented marker sources, material and stack variation, and calibrated force noise. The same stage-alignment and paired error procedures can then test whether the observed correspondence benefit transfers to production joints.

6. Declarations

6.1 Author Contributions

Maximilian Müller contributes to conceptualization, methodology, software, validation, analysis, investigation, data collection, draft preparation, manuscript editing, visualization, supervision, project administration, and funding acquisition. Hannah Wagner contributes to validation, analysis, investigation, data collection, draft preparation. Lena Pichler contributes to conceptualization, validation, analysis. All authors have read and agreed with the manuscript before its submission and publication.

6.2 Funding and Acknowledgments

This research received no specific financial support from any funding agency.

6.3 Institutional Review Board Statement

Not applicable.

6.4 Informed Consent Statement

Not applicable.

6.5 Date Availability

Not applicable.

6.6 Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

6.7 Declaration of generative AI and AI-assisted technologies

During the preparation of this manuscript, the authors utilized Grammarly for linguistic polishing, logical organization and draft revision. The generative AI tool was not involved in research design, data collection, data analysis, or the equation of core conclusions. All outputs generated by AI were reviewed, revised and verified manually by the authors. The authors take full responsibility for all content of this manuscript.

7. References

[1] Sun, X., Beghi, A., Susto, G. A., & Lv, Z. (2024). Deep learning-based quality prediction for multi-stage sequential hot rolling processes in heavy rail manufacturing. Computers & Industrial Engineering, 196, 110466. https://doi.org/10.1016/j.cie.2024.110466

[2] Mohammadi Foumani, N., Miller, L., Tan, C. W., Webb, G. I., Forestier, G., & Salehi, M. (2024). Deep learning for time series classification and extrinsic regression: A current survey. ACM Computing Surveys, 56(9), Article 217, 1–45. https://doi.org/10.1145/3649448

[3] Murad, M. M. N., Aktukmak, M., & Yilmaz, Y. (2025). WPMixer: Efficient multi-resolution mixing for long-term time series forecasting. Proceedings of the AAAI Conference on Artificial Intelligence, 39(18), 19581–19588. https://doi.org/10.1609/aaai.v39i18.34156

[4] Wang, X., Zhang, M., & Su, J. (2025). Robust temporal alignment for multivariate time series forecasting. Expert Systems with Applications, 289, 128299. https://doi.org/10.1016/j.eswa.2025.128299

[5] Wu, S., Song, S., Deng, S., Xie, W., & Shen, L. (2026). Variable-length time series classification: Benchmarking, analysis and effective spectral pooling strategy. Information Fusion, 126, 103584. https://doi.org/10.1016/j.inffus.2025.103584

[6] Zhang, Z., Pham, T. D., An, Y., Doan, N. P., Alsharari, M., Tran, V.-H., Hoang, A.-T., Vandierendonck, H., & Mai, S. T. (2025). WaveletMixer: A multi-resolution wavelet based MLP-Mixer for multivariate long-term time series forecasting. Proceedings of the AAAI Conference on Artificial Intelligence, 39(21), 22741–22749. https://doi.org/10.1609/aaai.v39i21.34434

[7] Reuel, A., Hardy, A., Smith, C., Lamparth, M., Hardy, M., & Kochenderfer, M. J. (2024). BetterBench: Assessing AI benchmarks, uncovering issues, and establishing best practices. Advances in Neural Information Processing Systems, 37, 21763–21813. https://doi.org/10.52202/079017-0685

[8] Inácio, R., Cerqueira, V., Barandas, M., & Soares, C. (2025). Mast: Interpretable stress testing via meta-learning for forecasting model robustness evaluation. Machine Learning, 114, 251. https://doi.org/10.1007/s10994-025-06881-3

[9] Kim, I., Na, J., Yun, J. P., & Lee, S. (2024). Deep feature selection framework for quality prediction in injection molding process. IEEE Transactions on Industrial Informatics, 20(1), 503–512. https://doi.org/10.1109/TII.2023.3268421

[10] Ren, L., Mo, T., Cheng, X., & Li, X. (2024). Temporal–frequency attention focusing for time series extrinsic regression via auxiliary task. IEEE Transactions on Neural Networks and Learning Systems, 35(11), 15494–15506. https://doi.org/10.1109/TNNLS.2023.3287318

[11] Wang, A., Tong, T. P., Mizera, A., & Pang, J. (2024). Benchmarking structural inference methods for interacting dynamical systems with synthetic data. Advances in Neural Information Processing Systems, 37, 135129–135185. https://doi.org/10.52202/079017-4295

[12] Terven, J., Cordova-Esparza, D.-M., Romero-González, J.-A., Ramírez-Pedraza, A., & Chávez-Urbiola, E. A. (2025). A comprehensive survey of loss functions and metrics in deep learning. Artificial Intelligence Review, 58, 195. https://doi.org/10.1007/s10462-025-11198-7

[13] Li, H., Yu, Y., Zhao, Z., & Tang, X. (2024). A time series short-term prediction method based on multi-granularity event matching and alignment. Computers, Materials & Continua, 78(1), 653–676. https://doi.org/10.32604/cmc.2023.046424

[14] Beck, M., Pöppel, K., Spanring, M., Auer, A., Prudnikova, O., Kopp, M., Klambauer, G., Brandstetter, J., & Hochreiter, S. (2024). xLSTM: Extended long short-term memory. Advances in Neural Information Processing Systems, 37, 107547–107603. https://doi.org/10.52202/079017-3417

[15] Shi, X., Hu, T., Chang, Y., Song, Y., Liang, A., & Sun, Y. (2025). Waveel: A wavelet-based ensemble learning framework for time series forecasting. EURASIP Journal on Wireless Communications and Networking, 2025, 52. https://doi.org/10.1186/s13638-025-02479-4

[16] Tian, Y., Zhao, S., & Zhang, X. (2025). Robustness and orthogonality: Time series forecasting via wavelets. Information Sciences, 717, 122328. https://doi.org/10.1016/j.ins.2025.122328

[17] Forootani, A., & Khosravi, M. (2026). Compact transformer variants for synthetic time series forecasting. Neurocomputing, 677, 133140. https://doi.org/10.1016/j.neucom.2026.133140

[18] Hiniduma, K., Byna, S., & Bez, J. L. (2025). Data readiness for AI: A 360-degree survey. ACM Computing Surveys, 57(9), Article 219, 1–39. https://doi.org/10.1145/3722214

[19] Semmelrock, H., Ross-Hellauer, T., Kopeinik, S., Theiler, D., Haberl, A., Thalmann, S., & Kowald, D. (2025). Reproducibility in machine-learning-based research: Overview, barriers, and drivers. AI Magazine, 46(2), e70002. https://doi.org/10.1002/aaai.70002

[20] Mathieu, T., & Preux, P. (2024). Statistical comparison in empirical computer science with minimal computation usage. In Proceedings of the 2nd ACM Conference on Reproducibility and Replicability (pp. 20–24). ACM. https://doi.org/10.1145/3641525.3663618

[21] Guijo-Rubio, D., Middlehurst, M., Arcencio, G., Silva, D. F., & Bagnall, A. (2024). Unsupervised feature based algorithms for time series extrinsic regression. Data Mining and Knowledge Discovery, 38(4), 2141–2185. https://doi.org/10.1007/s10618-024-01027-w

[22] Koziel, S., Pietrenko-Dabrowska, A., Wojcikowski, M., & Pankiewicz, B. (2024). Efficient calibration of cost-efficient particulate matter sensors using machine learning and time-series alignment. Knowledge-Based Systems, 295, 111879. https://doi.org/10.1016/j.knosys.2024.111879

[23] Stržinar, Ž., Pregelj, B., & Škrjanc, I. (2024). Non-elastic time series fuzzy clustering for efficient analysis of industrial data sets. Applied Soft Computing, 167, 112398. https://doi.org/10.1016/j.asoc.2024.112398

[24] d'Hondt, J. E., Li, H., Yang, F., Papapetrou, O., & Paparrizos, J. (2025). A structured study of multivariate time-series distance measures. Proceedings of the ACM on Management of Data, 3(3), 1–29. https://doi.org/10.1145/3725258

[25] Nasr, S., & Sikdar, S. (2024). IndMask: Inductive explanation for multivariate time series black-box models. In ECAI 2024 (pp. 1108–1115). IOS Press. https://doi.org/10.3233/FAIA240603

[26] Šimić, I., Veas, E., & Sabol, V. (2025). A comprehensive analysis of perturbation methods in explainable AI feature attribution validation for neural time series classifiers. Scientific Reports, 15, 26607. https://doi.org/10.1038/s41598-025-09538-2

[27] Song, Y., Li, F., Wang, Z., Zhang, B., & Zhang, B. (2024). Uncertainty quantification of data-driven quality prediction model for realizing the active sampling inspection of mechanical properties in steel production. International Journal of Computational Intelligence Systems, 17, 74. https://doi.org/10.1007/s44196-024-00451-6

[28] Park, J., & Min, M. (2025). Robustness and performance comparison of generative AI time series anomaly detection under noise. Computer Modeling in Engineering & Sciences, 145(3), 3913–3948. https://doi.org/10.32604/cmes.2025.072261

[29] Lin, J., Liu, D., Wu, T., Xu, Y., & Chen, Y. (2025). TSINet: A temporal-channel factorized mixing and spectral enhanced interactive network for time series forecasting. Information Sciences, 719, 122440. https://doi.org/10.1016/j.ins.2025.122440

[30] Michałowska, K., Goswami, S., Karniadakis, G. E., & Riemer-Sørensen, S. (2025). Multi-resolution learning with DeepONets and long short-term memory neural networks. Neurocomputing, 653, 131154. https://doi.org/10.1016/j.neucom.2025.131154

[31] Pecher, B., Srba, I., & Bielikova, M. (2024). A survey on stability of learning with limited labelled data and its sensitivity to the effects of randomness. ACM Computing Surveys, 57(1), Article 19, 1–40. https://doi.org/10.1145/3691339

[32] Wang, X., Tang, Z., Guo, J., Meng, T., Wang, C., Wang, T., & Jia, W. (2025). Empowering edge intelligence: A comprehensive survey on on-device AI models. ACM Computing Surveys, 57(9), Article 228, 1–39. https://doi.org/10.1145/3724420