Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Concept-Based Force-Signature Analysis of Tool-Parameter Effects in Fine Blanking

Martin Unterberg

M.Unterberg@mti.rwth-aachen.de

Manufacturing Technology Institute (MTI) of RWTH Aachen University

Daria Gelbich

Manufacturing Technology Institute (MTI) of RWTH Aachen University

Frank Schweinshaupt

Manufacturing Technology Institute (MTI) of RWTH Aachen University

Antonia Holzapfel

Institute for Data Science in Mechanical Engineering (DSME) of RWTH Aachen University

Daniyal Kazempour

Department of Computer Science, Christian-Albrechts-Universität zu Kiel

Philipp Niemietz

Manufacturing Technology Institute (MTI) of RWTH Aachen University

Peer Kröger

Department of Computer Science, Christian-Albrechts-Universität zu Kiel

Sebastian Trimpe

Institute for Data Science in Mechanical Engineering (DSME) of RWTH Aachen University

Thomas Bergs

Manufacturing Technology Institute (MTI) of RWTH Aachen University

---

Research Article

Keywords: Fine blanking, Process noise, Explainable AI, XAI, Concept extraction, Predictive quality, Deep learning

Posted Date: January 30th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-8730577/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: The authors declare no competing interests.

Concept-Based Force-Signature Analysis of Tool-Parameter Effects in Fine Blanking

Martin Unterberg[1*], Daria Gelbich[1], Frank Schweinshaupt[1], Antonia Holzapfel[2], Daniyal Kazempour[3], Philipp Niemietz[1], Peer Kröger[3], Sebastian Trimpe[2], Thomas Bergs[1,4]

[1*]Manufacturing Technology Institute (MTI), RWTH Aachen University, Campus-Boulevard 30, Aachen, 52074, Germany.
[2]Institute for Data Science in Mechanical Engineering (DSME), RWTH Aachen University, Theaterstraße 35–39, Aachen, 52062, Germany.
[3]Department of Computer Science, Christian-Albrechts-Universität zu Kiel, Christian-Albrechts-Platz 4, Kiel, 24098, Germany.
[4]Fraunhofer Institute for Production Technology (IPT), Steinbachstraße 17, Aachen, 52074, Germany.

*Corresponding author(s). E-mail(s): m.unterberg@mti.rwth-aachen.de;
Contributing authors: d.gelbich@mti.rwth-aachen.de; f.schweinshaupt@mti.rwth-aachen.de; antonia.holzapfel@dsme.rwth-aachen.de; daniyal.kazempour@informatik.uni-kiel.de; p.niemietz@mti.rwth-aachen.de; peer.kroeger@informatik.uni-kiel.de; trimpe@dsme.rwth-aachen.de; t.bergs@mti.rwth-aachen.de;

Abstract

Fine blanking part quality is highly sensitive to tool parameterization, yet cause–effect relationships are obscured by substantial process noise. This limits the utility of both simulation and scalar process metrics for tool optimization and mechanistic understanding. At the same time, high-resolution force-signature data is information-rich and can be used as evidence for learning physically plausible relationships, provided that interpretability is retained. This paper presents a data-to-knowledge pipeline combining finite element simulation, neural network prediction, and concept extraction to derive interpretable patterns from force time series and formulate mechanistic hypotheses linking tool parameters to die-roll formation. Controlled tool-parameter variation generates approximately 27,000 strokes with die-roll metrology. Neural networks predict die-roll height

and tool parameters with mean absolute percentage errors of 0.6% and 2.6%, respectively. Concept extraction (ECLAD-ts) and segment-aware attribution localize predictive information to peak-load and post-peak regimes. For die-clearance variation, piecewise-linear analysis within concept windows identifies two mechanism-sensitive slope descriptors (Cliff's $|\boldsymbol{\delta}| \geq \mathbf{0.83}$) with ordered progressions mirroring the graded die-roll response. The rising-flank slope captures clearance-dependent shearing resistance, while the post-peak decay slope reflects thermo-mechanical softening once crack growth initiates. Smaller clearance yields higher resistance, faster decay, and reduced die-roll height, consistent with forming theory. The pipeline enables phase-local assessment of tool-parameter effects under process noise and supports testable mechanistic hypotheses for simulation refinement.

Keywords: Fine blanking, Process noise, Explainable AI, XAI, Concept extraction, Predictive quality, Deep learning

1 Introduction

Discrete sheet-metal forming processes such as fine blanking enable high-volume production of functional parts with tight tolerances. Process capability and part quality are largely determined by the design of the forming tool: small changes in punch and die geometry, such as clearances, edge radii, and local features, can shift the stress state, material flow, and fracture initiation. Thus, tool design and parameterization are high-leverage decisions with direct impact on reject rates, tool life, and consequently both planned and unplanned machine downtime.

In practice, tool parameterization is guided by a combination of expert knowledge, shop-floor iteration, and finite element (FE) simulations. While FE models are indispensable for studying stress and strain fields and for testing design hypotheses, predictive power in industrial settings is constrained by (i) uncertain boundary conditions and constitutive parameters, (ii) limited observability of relevant local states, and (iii) the presence of substantial *process noise* [1]. Here, process noise is understood as the aggregate of variability and disturbances that are not explicitly controlled or modelled, caused by, e.g., batch-to-batch material variation, lubrication and friction fluctuations, machine compliance and alignment changes, temperature effects, and evolving surface conditions. Especially in highly sensitive shear- and fracture-dominated operations, process noise can dominate the scatter of observable signals and quality metrics, such that stable cause–effect relations are difficult to infer from simulations or scalar summary measures alone.

At the same time, modern sensor technology enables high-resolution process data acquisition (e.g., force–displacement curves) that is directly influenced by the underlying mechanics. Such signals are information-rich, yet high-dimensional, non-stationary across the stroke, and affected by the same process noise that limits simulation fidelity. Data-driven models can be trained to map tool parameters and process signals to quality-relevant targets, but are often operated as black boxes and therefore provide limited support for mechanistic understanding or engineering transfer. What is missing

is a systematic approach in which process data is used *as evidence* for learning physically plausible relationships between tool parameters, force-signature structure, and quality outcomes, while interpretability is retained to a degree that permits scrutiny by domain experts.

This gap is addressed in this paper through the development of a *data-to-knowledge* pipeline in which tool design, i.e., tool parametrization, is positioned as the main experimental narrative (Fig. 1). The aim of this work is *explicitly not* to identify an "optimal" parameter set. Instead, controlled variation of tool parameters is treated as an intervention that generates structured data for mechanism-oriented analysis. A modelling–data feedback loop is formulated as the intended mode of operation: FE studies are used to derive physics-informed *indications* of where tool-parameter effects may manifest in force signatures, while experimental process data is used to localize and quantify corresponding phase-local patterns under realistic process noise. On this basis, mechanistic hypotheses are motivated and a route towards refined, targeted simulations is provided, while the present work focuses on establishing and demonstrating the individual pipeline components.

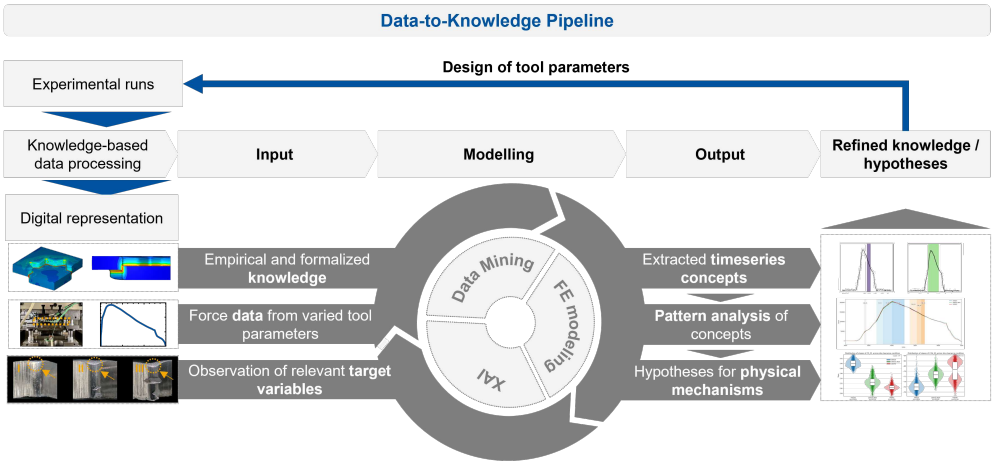


Fig. 1: Data-to-knowledge pipeline: tool parameter variation is used to generate structured experimental evidence; FE modelling provides physics-informed reference trends; CNN-based prediction is leveraged to capture discriminative signal structure; concept extraction (ECLAD-ts) and concept-centric data mining are used to derive interpretable representations and to relate force-signature patterns to tool parameters and quality criteria under process noise.

A central element of the proposed pipeline is the extraction of an interpretable representation from high-dimensional force time series. Convolutional neural networks (CNNs) have been shown to be effective for time series classification (TSC) and are therefore used here as flexible function approximators to capture discriminative structure in measured signals with respect to tool-parameter settings and quality-relevant

targets. However, the internal representations learned by CNNs are not directly accessible, which complicates their use for mechanism-oriented interpretation. For the present work, model explanation is not pursued as an end in itself. Rather, access to learned internal structure is required to derive compact, interpretable descriptors that can be related to physical mechanisms and to quality criteria such as the die-roll height.

To this end, *concept extraction* (CE) is employed to transform learned CNN representations into a set of *time series concepts* that can be analysed independently of the classifier. In a prior work, the time series CE and localization method ECLAD-ts was introduced, providing post-hoc global concept representations based on how subsets of the input are encoded at different levels of abstraction [2]. Concepts are obtained by clustering timestep-wise aggregations of CNN activation maps, and their relevance is quantified based on their impact on the prediction process. In the pipeline proposed here, these concepts are primarily interpreted as a data-centric *vocabulary of recurring motifs* and phase-local signal patterns, enabling subsequent concept-centric data mining and mechanistic interpretation.

Based on the extracted concepts, concept-centric data mining is performed to relate force-signature structure to tool parameters and quality criteria. In this context, concepts are analysed as objects that can be aggregated, compared, and correlated across tool settings and experimental runs, allowing stroke phases with concentrated information content to be identified. In addition, piecewise-linear approximations (PLA) are employed as a complementary interpretable signal representation to obtain primitives such as breakpoints, slopes, offsets, and segment durations. Segment-aware attribution analyses are applied to determine stroke phases in which tool-parameter dependencies and quality-relevant information are concentrated. The resulting patterns are plausibilized using process knowledge and, where available, triangulated against simulation trends.

The following contributions are provided within this work:

- A tool-parameterization-driven experimental design and digital data representation linking tool geometry parameters, multichannel force signals, and die-roll metrology as quality-relevant target variables.
- A data-to-knowledge pipeline combining FE-based sensitivity analysis with data-driven modeling and concept-centric analysis, enabling phase-local assessment of tool-parameter effects under process noise and supporting mechanism-oriented interpretation.
- The use of the time series concept extraction and localization method ECLAD-ts for regression, to derive concepts and importance scores from CNN internal representations for force-signal analysis.
- A concept-centric analysis workflow that combines segment-aware attribution and concept-window PLA to relate extracted concepts and phase-local slope descriptors to tool parameters and die-roll behavior.

The remainder of this paper is structured as follows. Related work is presented in Sections 2. The proposed methodology is described in Section 3. Results and discussion are presented in Section 4. Section 5 concludes with a summary of the main findings and an outlook on future work.

2 Related work and background

This section provides the domain and methodological background for the proposed pipeline. Fine blanking and the challenge of process noise are introduced in Section 2.1. Concept extraction for multivariate time series is reviewed in Section 2.2. Piecewise-linear representations of force signatures are discussed in Section 2.3.

2.1 Fine blanking and process noise

Fine blanking is a high-precision shearing process for the mass production of net-shape components, characterized by a sheared surface with a large clean-shear zone and high perpendicularity [3]. This quality is essential for functional components, particularly in safety-critical applications, as it minimizes the fracture zone and die roll. The process achieves this by inducing a state of high compressive stress in the shear zone, which suppresses crack initiation and promotes controlled plastic flow.

This state of stress is accomplished through the interplay of three key process characteristics: the use of a V-ring indenter to clamp the material, the application of high blank holder and counter forces, and a significantly reduced die clearance, typically in the range of 0.5% of the sheet thickness [4]. This combination is designed to transform the shearing operation from a brittle fracture into a guided shearing process.

Despite these well-defined process characteristics, the actual quality of the sheared surface is subject to significant variability, a phenomenon broadly termed process noise [1]. Several studies have emphasized the critical role of the material's microstructure in contributing to this noise. Inhomogeneities such as carbide accumulations and non-metallic inclusions can act as stress concentrators, promoting the formation of tears even under high compressive stress states [5, 6]. The occurrence of these tears, along with the magnitude of the die roll height, compromises the functional, load-bearing area of the component's surface. Understanding and mitigating the stochastic formation of these quality-impacting features remains a central challenge in industrial fine blanking applications.

2.2 Concept extraction for multivariate time series

A concept describes an input pattern that is consistently identifiable within the model's internal representations. Explanation methods based on concepts describe a model's global behavior by linking its predictions to the presence/absence of such patterns and quantifying their influence on the decision process, which is useful for auditing, debugging, and communicating model behavior in applied settings. These methods are typically post hoc, i.e., applied after model training.

Concept-based explanations are well established in computer vision: concept testing methods such as TCAV [7] and CAR [8] quantify a model's sensitivity to human-defined concepts, while CE methods like ACE [9] and ECLAD [10] automatically discover visual concepts. CE has been used in safety-critical and industrial settings, including medical imaging and visual inspection, underscoring its practical value [11, 12]. For multivariate time series, a concept typically corresponds to a recurring time-localized motif, e.g., a characteristic waveform shape or a transient event

that repeats across instances and aligns with an interpretable system mechanism. In this domain, concept explanations are emerging: TCAV has been adapted to structured temporal data such as EHRs and physiological signals, and generalized concept frameworks have begun to appear [13, 14]. Most recently, ECLAD-ts provides post-hoc global concept extraction for multivariate time series via latent-space clustering and concept localization [2]. To the knowledge of the authors, CE has not yet been applied to industrial forming time series, in particular fine blanking; the present work bridges this gap by extracting process-level concepts from multichannel force signals and using them for deriving mechanism-oriented descriptors under process noise.

In this work, ECLAD-ts [2] is applied to multivariate force time series to obtain concept masks and concept relevance scores. ECLAD-ts groups timesteps with similar "hidden-states" within a model into a few recurring patterns (concepts) and then scores how much each pattern affects the output. ECLAD-ts comprises three steps: (1) *Extraction of a representation for each timestep in the latent space of the network.* Activation maps from multiple layers are upsampled to the input length and concatenated to form per-timestep Local Aggregated Descriptors (LADs). (2) *Concept mining and localization.* LADs are clustered via mini-batch k-means to obtain concept centroids; each timestep is assigned to its nearest centroid, yielding a localized concept mask (expanded per channel for multivariate inputs). (3) *Concept relevance computation.* Each concept is assigned an importance score based on instance-wise sensitivity (computed via gradients) to regions belonging to that concept; larger magnitude indicates stronger influence on the output. Overall, given a trained model and dataset, ECLAD-ts returns recurring, time-localized concept patterns together with relevance measures that support global model interpretation.

2.3 Piecewise-linear representations of force signatures

Force signatures in fine blanking are complex and affected by process noise, motivating an approximation that reduces dimensionality while retaining physical interpretability. A piecewise-linear representation is suitable because distinct phases of the blanking cycle (clamping, shearing, separation, and stripping) often exhibit monotonic trends that can be summarized by linear segments (slopes and intercepts). The resulting slope descriptors provide a compact proxy of local force evolution and enable phase-consistent comparisons across strokes.

PLA summarizes a signal by k linear segments. Larger k yields a more local approximation, while smaller k yields a more global abstraction. In mass production, systematic changes such as thermal stabilization or tool wear can manifest as gradual changes in segment slopes across strokes, whereas stochastic disturbances can appear as localized deviations. In the proposed pipeline, PLA is used in two complementary ways: (i) as an anchored, stroke-consistent segmentation of the full force curve for phase-resolved attribution, and (ii) as a concept-local analysis tool to extract slope descriptors within concept windows where predictive information is concentrated.

3 Methodology

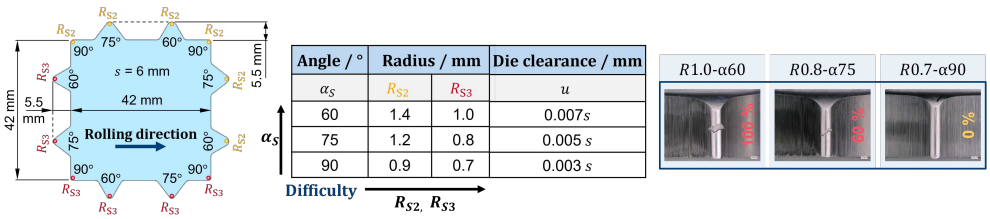
An end-to-end pipeline was used to acquire and analyze fine blanking data and to derive interpretable process descriptors. Multivariate time-series data were recorded on an industrial fine blanking press and paired with die-roll metrology (Section 3.1). Physics-informed reference trends were obtained via thermomechanically coupled FEM simulations (Section 3.2). Supervised neural networks were trained to predict die-roll height and tool parameters, and post-hoc CE was performed to obtain recurring time series patterns (Section 3.3). Predictive information was additionally localized along the stroke via anchored piecewise-linear segmentation and segment-wise attribution. Finally, PLA were fitted within concept windows to quantify phase-local differences across die-clearance conditions (Section 3.4).

3.1 Process setup

The experimental setup was designed to systematically investigate the effects of tool design parameters on the fine blanking process under conditions reflective of industrial mass production. A series of fine blanking experiments was carried out on a servomechanical press (FEINTOOL XFT 2500 SPEED). The press was operated at a stroke frequency of 50 min^{-1}, corresponding to a blanking velocity of approximately 45 mm/s in the linear range of the press kinematics.

The sheet-metal material was a cold-rolled, high-strength alloyed steel 58CrV4 (DIN EN 1.8161) from a single supplier, with a sheet thickness (s) of 6 mm. The material was supplied in an annealed state to achieve spheroidized carbides (+AC), a condition known to exhibit microstructural inhomogeneities that contribute to process noise.

A custom single-cavity fine blanking tool was utilized. The active tool elements (punch, die, guide, counter punch) were manufactured from tool steels (BÖHLER S390 and X155CrVMo12-1). A mineral oil-based, chlorine-free lubricant (FUCHS WISURA FMO 5020) was used throughout the experiments. The tool was designed to produce a component with a complex geometry featuring varied radii and angles analogous to functional gear-tooth features (Fig. 2). This geometry was selected to induce challenging shearing conditions and to promote quality-relevant surface features, including die roll.



Angle / °	Radius / mm		Die clearance / mm
α_S	R_{S2}	R_{S3}	u
60	1.4	1.0	0.007s
75	1.2	0.8	0.005 s
90	0.9	0.7	0.003 s

Fig. 2: Schematic of the part geometry used in the experiments. Key features designed to induce challenging shearing conditions are highlighted.

To generate a dataset covering a wide range of process states, a structured tool-parameter variation was performed. The study focused on key tooling parameters (active surfaces) known to influence sheared surface quality: die clearance (u), chamfer angle (α_s), and chamfer height (h_s). Three levels were considered around a reference configuration ($u_{\mathrm{U030}} = 0.005s$, $\alpha_{s,\mathrm{U030}} = 35°$, $h_{s,\mathrm{U030}} = 0.5\,\mathrm{mm}$). As detailed in Table 1, five distinct parameter sets were defined. For each parameter set, a new blanking punch was installed to reduce confounding by punch wear. Each experiment was run for one full sheet-metal coil of material (split into two sub-coils), resulting in ten processed sub-coils and approximately 27,000 recorded strokes. This procedure was selected to capture not only immediate parameter effects but also longer-horizon variability sources contributing to process noise (e.g., drift, thermal stabilization, and batch-wise material variation along the coil).

Table 1: Overview of the experiments. Five primary conditions were investigated, with each condition run on two separate sub-coils from a single parent coil.

Condition	**Sub coil**[a]	**Die clearance** [mm]	**Chamfer angle** [deg]	**Chamfer height** [mm]
U030 (Ref)	R1C6	0.03	35	0.5
	R1C3	0.03	35	0.5
U040	R1C2	0.04	35	0.5
	R1C5	0.04	35	0.5
U020	R1C1	0.02	35	0.5
	R1C4	0.02	35	0.5
CA25	R2C5	0.03	25	0.5
	R2C2	0.03	25	0.5
CH03	R2C6	0.03	35	0.3
	R2C3	0.03	35	0.3

[a] Each of the five primary conditions was initiated with a new blanking punch.

A comprehensive data acquisition system was employed to create a digital representation of the process. Piezoelectric force sensors were used to measure the punch force and the counter-holder force at multiple locations (yielding four punch channels and four counter-holder channels, as summarized in Table 4). Tool displacement during the shearing phase was measured using a laser triangulation sensor (Micro-Epsilon optoNCDT 1900), enabling a mapping of force signals to tool trajectory. All signals were sampled synchronously at 20 kHz.

Die-roll height was used as the primary quality label. The die roll is the curved rollover of material at the edge of the sheared surface (Fig. 3), formed during the initial stage of blanking due to plastic flow under high compressive stresses. From a functional standpoint, the die roll is a non-load-bearing surface that reduces the effective height of the clean-shear zone, and its minimization is a common objective in fine blanking [4]. To quantify the die roll, parts were sampled at regular intervals along production. From each sample batch, the 1st, 6th, 11th, and 16th parts were selected (preserving their production order). The profiles of three corners per part (60°, 75°,

90°) were digitized using a 3D optical profilometer (KEYENCE VR-6200). Die-roll height was determined algorithmically from the profile by identifying the transition from the curved die-roll region to the straight clean-sheared surface via the local profile angle. The vertical height at the first point where this angle fell below a threshold of 10° was defined as the die-roll height (per corner).

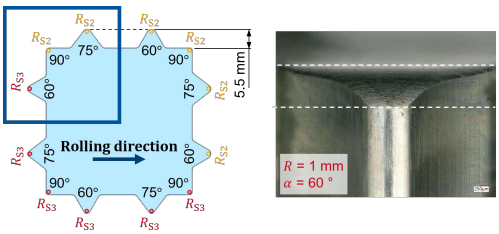


Fig. 3: Left: Part geometry indicating evaluated corner angles (60°, 75°, 90°). Right: Micrograph of a representative sheared surface illustrating the die roll.

3.2 Numerical modeling for fine blanking

To provide a physics-based reference for interpreting data-driven findings, a thermo-mechanically coupled FEM model was developed. The modeling approach followed the framework of Schweinshaupt et al. [15] and was implemented in Forge NxT 4.1 (implicit solver).

Given geometric symmetry, a quarter model was used. Tool components and the workpiece were discretized with 3D tetrahedral elements. A mesh convergence analysis based on averaged temperature and stress in the shear zone indicated that a minimum element edge length of 0.35 mm was sufficient. Therefore, the shear zone and corresponding contact surfaces were meshed with $l_{\min} = 0.35\,\text{mm}$ and coarsened away from the deformation zone. Tool components were modeled as elastic bodies. Tool and workpiece parameters were adapted from the validated framework in [15]. Since a validated material model for 58CrV4+AC was not available, parameters for the chemically similar 51CrV4 (soft annealed) were used. The workpiece was modeled as thermoviscoplastic with remeshing. Flow stress was described by a Hensel-Spittel law [16]:

$$\sigma_f(\theta, \varphi, \dot{\varphi}) = A \cdot e^{m_1\theta} \cdot \varphi^{m_2} \cdot e^{m_4/\varphi} \cdot \dot{\varphi}^{m_3}, \tag{1}$$

with coefficients listed in Table 2. Material separation was governed by the normalized Cockcroft–Latham criterion ($C_{crit} = 1.44$) [17]. Friction at the tool–workpiece interface was described by a hybrid law combining Coulomb friction (μ) with a shear friction limit (m).

A sensitivity study was conducted to derive physics-informed indications of where parameter effects can manifest in the force signature. A reference simulation was established and compared against the experimental reference configuration. Thermo-mechanical and boundary-condition parameters were varied as summarized in Table 3.

Table 2: Coefficients for the Hensel–Spittel flow curve model for 51CrV4 (soft annealed state) [18].

A	m_1	m_2	m_3	m_4
997.51	−0.000 80	0.170 08	0.013 24	−0.008 18

In addition, geometric variations in die clearance and chamfer geometry were evaluated using the experimental levels in Table 1. The resulting force-curve trends were used as qualitative reference indications for subsequent, data-driven localization and interpretation.

Table 3: Thermomechanical and boundary-condition parameters varied in the FEM sensitivity analysis (geometric variations are reported in Table 1).

Parameter	**Low**	**Reference**	**High**
Factor of flow curve coefficient A [-]	0.9 (897.76)	1.0 (997.51)	1.1 (1097.26)
Friction μ [-] / m [-]	0.05 / 0.1	0.1 / 0.2	0.2 / 0.3
Tool temperature [°C]	–	20	50

3.3 Neural network training and concept extraction

Supervised neural networks were trained as flexible function approximators to learn mappings from fine blanking signals to quality-relevant targets (die-roll height) and tool-parameter settings. Predictive performance alone is not the end goal: the trained models serve as representation learners that concentrate task-relevant signal structure in their latent space. Concept extraction is then used to expose this structure in the form of concepts that can be related to tool parameters and quality criteria.

3.3.1 Model training

Using the fine blanking data, supervised neural networks were trained under multiple input–output settings. To assess the contribution of different sensor subsets, three time-series datasets (S, M, L) were defined (Table 4). Since predictive performance improved with the full sensor set, most downstream analyses were performed using the L dataset.

Two output types were considered: (i) die-roll height (quality) and (ii) tool parameters (chamfer angle, chamfer height, and die clearance). Across the resulting input–output combinations (3 datasets × 2 targets), seven architectures were evaluated: DenseNet-121/169 [19], InceptionTime with 10 and 20 inception blocks [20], and ResNet-18/50/152 [21]. Inputs were standardized, and zero-mean Gaussian noise was applied for augmentation. Data were split 80%/20% into train/validation; the same

Table 4: Sensor-channel availability per dataset.

	Punch				Counter-holder				Position range		Combined / derived		
Dataset	f	r	ri	l	1	2	3	4	100 mm	25 mm	$F_{\mathrm{P},c}$	$F_{\mathrm{CH},c}$	F_{S}
S											✓		✓
M									✓		✓	✓	✓
L	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓

Note. Punch channels correspond to $F_{\mathrm{P},f}$, $F_{\mathrm{P},r}$, $F_{\mathrm{P},ri}$, $F_{\mathrm{P},l}$ with $F_{\mathrm{P},c} = \sum F_{\mathrm{P},\cdot}$. Counter-holder channels correspond to $F_{\mathrm{CH},1}$–$F_{\mathrm{CH},4}$ with $F_{\mathrm{CH},c} = \sum_{i=1}^{4} F_{\mathrm{CH},i}$. Shearing force: $F_{\mathrm{S}} = F_{\mathrm{P},c} - F_{\mathrm{CH},c}$.

split was used for concept extraction and visualization. Experiments were repeated with seeds $\{0, 1, 2\}$ to account for stochasticity. Optimization was performed with Adam (weight decay 10^{-3}, initial learning rate 10^{-4}) for up to 400 epochs, using early stopping on validation loss (patience 50 epochs).

3.3.2 Extraction of concepts

ECLAD-ts [2] is originally defined for classification and computes concept relevance using a wrapper $g(y) = ||y \cdot \mathbf{1}^T - \mathbf{1} \cdot y^T||$ which measures the effect of concept regions on class discrimination. For regression tasks, this wrapper was adapted to $g(y) = y$, such that the relevance directly reflects a concept's influence on the predicted scalar output.

The resulting importance scores were normalized to $[-1, 1]$ and interpreted as follows: values close to $+1$ indicate that the concepts support larger predicted values; values close to -1 indicate that they support smaller predicted values; values near 0 indicate low concept importance.

ECLAD-ts was implemented in PyTorch. The number of concepts was set to $k \in \{10, 15\}$ for mini-batch k-means. Layers used to form Local Aggregated Descriptors (LADs) for each architecture are listed in Table 5.

Table 5: Layers for CE with each model.

CE methods	ECLAD-ts
DenseNet121	'model.features.transitionn_b.conv', $n_b \in \{1, 2, 3\}$ 'model.features.denseblock4.block.15.conv2'
DenseNet169	'model.features.transitionn_b.conv', $n_b \in \{1, 2, 3\}$ 'model.features.denseblock4.block.31.conv2'
InceptionTime10	'model.inception_block.inception_layers.n_b.bottleneck', $n_b \in \{6, 7, 8, 9\}$
InceptionTime20	'model.inception_block.inception_layers.n_b.bottleneck', $n_b \in \{8, 16, 17, 19\}$
ResNet18	'model.layers.n_b.1.relu', $n_b \in \{0, 1, 2, 3\}$
ResNet50	'model.layers.n_b.relu', $n_b \in \{0.2, 1.3, 2.5, 3.2\}$
ResNet152	'model.layers.n_b.relu', $n_b \in \{0.2, 1.7, 2.35, 3.2\}$

3.4 Piecewise-linear analysis for phase-local attribution and comparison

To complement concept extraction with interpretable, phase-resolved descriptors, PLA is applied to the force signature in two distinct ways: (i) as a global, stroke-consistent segmentation for localizing predictive information via segment-wise attribution, and (ii) as a concept-local analysis tool for quantifying die-clearance effects within concept windows where relevance is concentrated. Both applications rely on an anchored segmentation strategy that ensures phase consistency across strokes.

3.4.1 Anchored piecewise-linear approximation

To enable segment-wise comparisons under process noise, breakpoints must be consistent across all strokes. An anchored PLA is therefore defined as follows. A pooled template force curve $\tilde{y}(x)$ is constructed by aggregating strokes and computing the pointwise median. Breakpoints $\{\beta_r\}_{r=1}^{K-1}$ are then estimated on this template by minimizing the sum of squared residuals:

$$\min_{\{\beta_r\}} \sum_{x} \left(\tilde{y}(x) - \hat{y}(x; \{\beta_r\})\right)^2, \tag{2}$$

where $\hat{y}$ denotes the continuous piecewise-linear approximation. Once determined, these breakpoints are held fixed for all strokes, yielding K phase-consistent segments. For each stroke j and segment k, a linear model $f_{k,j}(x) = \sigma_{k,j}x + \iota_{k,j}$ is fitted by ordinary least squares. The slope $\sigma_{k,j}$ serves as the primary descriptor of local force dynamics, summarizing the rate of force change independently of absolute force level.

3.4.2 Global segment-wise attribution

To localize predictive information along the full stroke, APLA is applied globally with a moderate number of segments covering all process phases. For each stroke and segment, a feature vector is extracted comprising slope, intercept, and summary statistics (mean, variance, autocorrelation). A gradient-boosted decision tree regressor (XGBoost [22]) is trained to predict die-roll height from these segment features under leave-one-experiment-out validation. Segment-level relevance is quantified using SHAP values [23], yielding a phase-local importance ranking that identifies stroke regions with concentrated predictive information. This global attribution view is used to corroborate concept-based findings and to verify that relevance peaks are not explained by trivial signal properties such as variance or amplitude.

3.4.3 Concept-local slope comparison

Where concept extraction identifies stroke phases with high relevance for die-roll prediction, PLA is applied within the corresponding concept windows to obtain finer-grained, mechanism-oriented descriptors. For each concept window W_C, the pooled-template approach is repeated: condition-specific median curves are aggregated into a condition-agnostic template, breakpoints are estimated within the window,

and per-stroke slopes are extracted. The resulting slope distributions are compared across die-clearance conditions using non-parametric statistics (Kruskal–Wallis, Mann–Whitney U) and effect sizes (Cliff's δ). Segments are ranked by maximum pairwise $|\delta|$ to identify where clearance yields the strongest phase-local separation. A monotonic trend descriptor based on Spearman correlation is used to assess whether slope progressions mirror the graded die-roll response across $U020 \rightarrow U030 \rightarrow U040$.

The global APLA partitions the entire force signature to identify which stroke phases carry predictive information. The concept-local PLA operates within predefined concept windows to quantify how tool-parameter effects manifest in specific, already-relevant regimes. The two analyses are complementary: global attribution localizes relevance, while concept-local comparison quantifies mechanism-sensitive differences where relevance is concentrated.

4 Results

This section follows the proposed data-to-knowledge pipeline. *First*, FEM sensitivity studies are used to derive physics-informed expectations of where tool-parameter effects should manifest in the force signature and to assess whether tool-parameter changes can translate into die-roll differences even when the corresponding global force-curve differences are visually subtle (Section 4.1). *Second*, metrology outcomes are reported to establish the observed quality response to tool parameterization (Section 4.2). *Third*, predictive performance is evaluated to verify that measured time series contain sufficient information for die-roll and tool-setting prediction, enabling concept extraction (Section 4.3). Building on this, predictive information is localized along the stroke via segment-wise attribution (Section 4.4). *Finally*, internal concepts are connected to interpretable, phase-local force descriptors by fitting piecewise-linear slopes inside concept windows and comparing them across die-clearance conditions (Section 4.5).

4.1 Numerical modeling of tool parameter variation

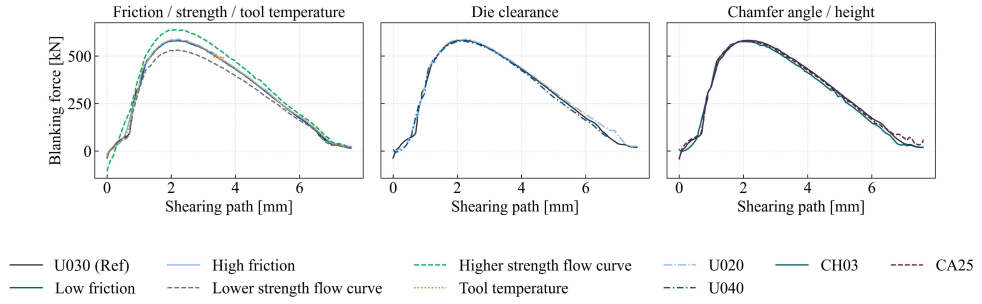


Fig. 4: Results of the FEM parameter variation study, showing the influence of different physical parameters on the simulated blanking force curve relative to the reference (U030) simulation.

Fig. 4 summarizes the FE-based parameter sensitivity of the simulated blanking force signature for representative variations in flow curve, friction, and tool temperature, as well as die clearance and chamfer geometry. Variations of the material flow curve primarily shift the force level over the high-load regime: the peak force and the surrounding stroke interval move upward/downward in an approximately proportional manner, while only minor changes of the overall curve shape are obtained. Changes induced by friction are smaller in magnitude and are mainly expressed on the rising flank and around the peak-load region. Within the investigated tool-temperature range (20–50 °C), only minor differences are observed in the simulated force response; thermo-mechanical sensitivity may become more pronounced at higher tool temperatures and under long-term thermal stabilization effects, which were not covered by the present range. These thermomechanical sensitivities provide useful context for interpreting force-level variations under process noise, but they primarily affect global amplitude rather than phase-local structure.

For die clearance, a different sensitivity pattern is obtained. Over a large portion of the stroke, including the region shortly prior to the peak, the peak itself, and the immediate post-peak region in which die-roll formation is expected, only weak differences are observed in the simulated global force curves. The most pronounced separation between die clearance configurations occurs in the late stage shortly before separation, where the force tail and its decay behavior differ. Importantly, the absence of visually salient differences in the peak-region force response does not imply an absence of die clearance effects on local mechanisms or part-quality features. Rather, it indicates that, within the present modelling setup and when global force is used as the observable, die clearance sensitivity is either concentrated in narrow stroke intervals and derivative-like characteristics or mediated by highly local phenomena that remain weakly expressed in the global signature. This contrast motivates the subsequent focus on die-clearance effects in the data-driven analysis: the phase-local, derivative-like nature of the expected sensitivity makes die clearance a suitable testbed for the proposed concept-based methods.

The FE sensitivity study was designed to provide physics-informed indications rather than an exhaustive reproduction of all real-world variability sources. Accordingly, several effects that can influence measured force signatures under industrial conditions were not fully represented. In particular, the investigated tool-temperature range (20–50 °C) does not cover elevated tool temperatures and thermal stabilization over longer production horizons. Furthermore, idealized assembly conditions were assumed, i.e., tolerances and mounting-induced misalignment (e.g., an offset punch) were not modelled, although such deviations can affect contact conditions and the force curve. Finally, frictional interactions between tool elements (e.g., punch–guide interfaces) were simplified, which may suppress sensitivity components that are present in practice.

In the proposed pipeline, these simulation trends are treated as physics-informed *indications* that guide subsequent data-driven localization and interpretation. For flow-curve and friction perturbations patterns of interest are expected around the peak-load regime (force-level shifts and phase-local shape changes). For die clearance, attention is directed towards late-stage segments as well as towards slope changes

in the peak and immediate post-peak region that may not be visually apparent in global force curves. These expectations are tested on experimental force data using concept extraction and segment-aware analyses, enabling weak or localized effects to be quantified under process noise and related to quality criteria such as die roll. Where systematic discrepancies between simulation-informed indications and data-observed patterns occur, targeted simulation refinements (e.g., elevated tool temperatures, controlled misalignment, and tool-element friction modelling) are motivated for subsequent mechanism isolation.

In addition to force-signature sensitivity, die-roll height was evaluated from the simulation results to provide a direct link to the experimental metrology outcomes reported in Section 4.2 and to assess whether tool-parameter effects that are subtle in global force curves can still translate into measurable quality differences. The simulated die-roll heights for key parameter variations are summarized in Table 6. Variations in material strength yield the largest absolute change in the simulated die roll: a higher strength flow curve results in a smaller die-roll height (0.89 mm), while a lower strength flow curve produces a larger one (1.66 mm). This inverse relationship between material strength and die-roll height is consistent with established forming theory [3] and motivates the hypothesis that part-to-part variability in material behavior during a stroke series can contribute to variability in die-roll height under industrial process noise, as examined in the experimental distributions in Section 4.2. Tool-parameter-induced differences in the simulated die-roll height are smaller in magnitude but follow tendencies that are compatible with the experimental results, while remaining sensitive to modelling assumptions and discretization. Accordingly, the simulated die-roll heights are interpreted as qualitative indications and are used to motivate analysis of the experimental force signatures, in particular for die clearance where global force-curve differences remain visually subtle.

Table 6: Simulated die-roll height for key parameter variations, corresponding to a 90° corner.

Simulation configuration	**Simulated die-roll height [mm]**
Higher strength flow curve	0.89
Lower strength flow curve	1.66
Die clearance 20 µm (U020)	1.40
Die clearance 30 µm (U030, Ref)	1.42
Die clearance 40 µm (U040)	1.56
Chamfer angle 25° (CA25)	1.34
Chamfer height 0.3 mm (CH03)	1.40

4.2 Metrology

Part quality was assessed via the die-roll height, measured at three corner geometries (60°, 75°, and 90°). Fig. 5 provides an overview of the resulting die-roll height distributions for all configurations, with corresponding descriptive statistics provided in

Table 7. Two consistent trends are observed. *First*, for all configurations, die-roll height increases with corner angle, with the sharpest corner (60°) exhibiting the smallest die-roll heights and the bluntest corner (90°) the largest, which might be explained by the different radii (see Fig. 2) and the resulting higher shearing difficulty with decreased radii of the respective geometry feature [24]. *Second*, distinct parameter effects are observed. The smaller chamfer angle (CA25) increases median die-roll height at 60° and 75° relative to the reference condition, while results at 90° are comparable. The smaller chamfer height (CH03) yields the smallest medians across corner geometries. For the experiments varying the die clearance, the largest die clearance (U040) yields elevated median die-roll heights and markedly increased scatter, most prominently at 90°, as reflected by the increased interquartile range and widened whiskers, suggesting decreased process stability under this configuration.

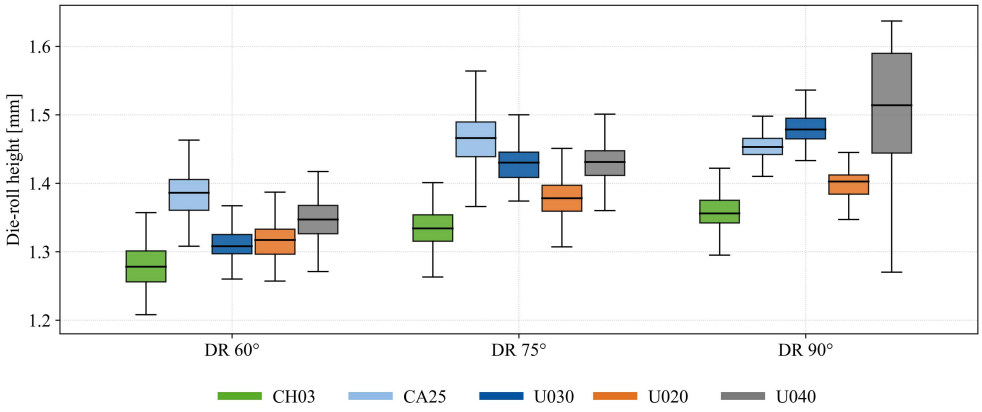


Fig. 5: Comparative boxplots of die-roll height distributions for the five experimental configurations, grouped by corner geometry.

Table 7: Median and Interquartile Range (in parentheses) for die-roll height (DR) [mm] across all experimental configurations and corner geometries.

Configuration	DR 60°	DR 75°	DR 90°
CH03	1.278 (0.045)	1.334 (0.038)	1.356 (0.033)
CA25	1.382 (0.046)	1.461 (0.052)	1.453 (0.027)
U030 (Ref)	1.308 (0.028)	1.430 (0.037)	1.478 (0.030)
U020	1.317 (0.036)	1.378 (0.038)	1.402 (0.028)
U040	1.347 (0.041)	1.431 (0.036)	1.514 (0.146)

While die-roll height results are reported for all investigated configurations, the mechanism-oriented analyses in Sections 4.4–4.5 focus on die-clearance experiments

(U020, U030, U040). This selection is motivated by two considerations: first, die clearance produces the most pronounced quality effects at the 90° corner (Table 7), where U040 exhibits both elevated median die-roll height and markedly increased scatter (IQR = 0.146 mm vs. 0.028–0.033 mm for other conditions); second, FEM indicates that die-clearance sensitivity is concentrated in derivative-like, phase-local characteristics rather than in global force-curve amplitude (Fig. 4, center panel), making it a suitable testbed for the proposed analysis methods. Extension to chamfer-parameter effects is left for future work.

Since a graded die-roll response with clearance is observed (Table 7), the subsequent analysis searches for concepts that not only separate the clearance conditions, but also exhibit an ordered progression across $U020 \rightarrow U030 \rightarrow U040$. In particular, descriptors are of interest for which the reference condition $U030$ lies between the two extremes, reflecting a graded quality effect and enabling mechanism-oriented interpretation rather than a purely categorical separation.

4.3 Predictive learnability and global concept profiles

A prerequisite for extracting interpretable concepts and downstream process knowledge is that the raw time-series measurements contain sufficient information to distinguish relevant process characteristics under industrial process noise. Supervised models are therefore evaluated for prediction of (i) die-roll height (DR) and (ii) process parameters from the acquired time series. Table 8 reports validation MAPE across datasets and model architectures; best datasets are highlighted in bold and best models per dataset are underlined.

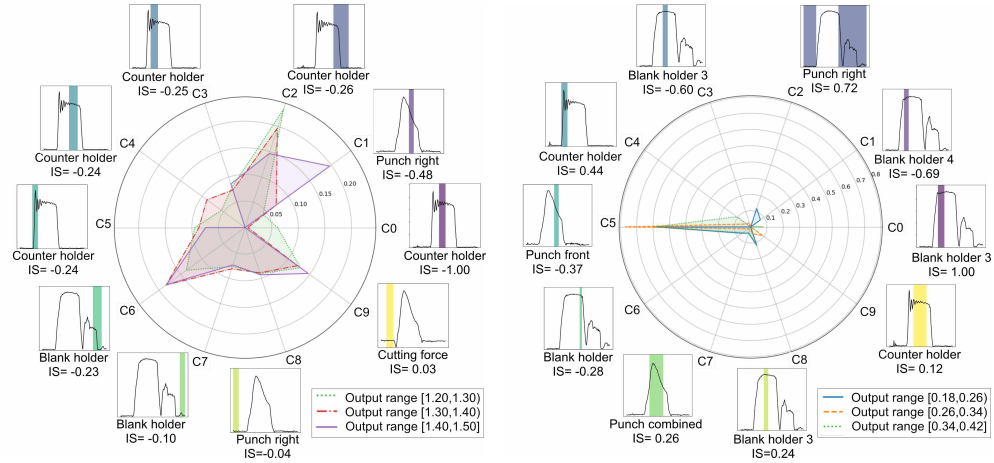
Model	DenseNet1D121	DenseNet1D169	InceptionTime10	InceptionTime20	ResNet1D152	ResNet1D18	ResNet1D50
$\mathbf{L}_{\mathrm{DR}}$	0.008 ± 0.001	0.007 ± 0.000	0.010 ± 0.002	0.007 ± 0.000	0.009 ± 0.005	0.007 ± 0.001	$\underline{0.006 \pm 0.001}$
M_{DR}	0.011 ± 0.000	0.010 ± 0.001	0.013 ± 0.002	0.010 ± 0.000	$\underline{0.008 \pm 0.000}$	0.008 ± 0.001	0.008 ± 0.001
S_{DR}	0.011 ± 0.001	0.010 ± 0.000	0.015 ± 0.001	0.011 ± 0.000	0.009 ± 0.001	0.010 ± 0.000	$\underline{0.008 \pm 0.000}$
Model	DenseNet1D121	DenseNet1D169	InceptionTime10	InceptionTime20	ResNet1D152	ResNet1D18	ResNet1D50
$\mathbf{L}_{\mathrm{params}}$	0.034 ± 0.002	0.032 ± 0.004	0.091 ± 0.099	0.049 ± 0.033	0.065 ± 0.082	$\underline{0.026 \pm 0.005}$	0.069 ± 0.095
$\mathrm{M}_{\mathrm{params}}$	0.085 ± 0.035	$\underline{0.069 \pm 0.014}$	0.078 ± 0.009	0.118 ± 0.128	0.117 ± 0.081	0.078 ± 0.019	0.140 ± 0.134
$\mathrm{S}_{\mathrm{params}}$	$\underline{0.067 \pm 0.026}$	0.076 ± 0.019	0.081 ± 0.015	0.086 ± 0.074	0.084 ± 0.076	0.081 ± 0.009	0.105 ± 0.071

Table 8: Validation MAPE on all considered datasets and model architectures. The best datasets are in bold, and the best models per dataset are underlined.

For the die-roll height prediction, the large dataset ($\mathbf{L}_{\mathrm{DR}}$) consistently yields the lowest errors across architectures, reaching MAPE $= 0.006 \pm 0.001$ with ResNet1D50 and remaining below ≈ 0.011 for all evaluated models. Similarly, for parameter prediction the best overall performance is obtained on $\mathbf{L}_{\mathrm{params}}$, with a minimum of MAPE $= 0.026 \pm 0.005$ (ResNet1D18). These results indicate that the time series contains enough information to predict both the die-roll height and the process parameters reliably, providing the empirical basis for the subsequent concept extraction and analysis.

To characterize what the networks learn beyond predictive accuracy, concept extraction was performed with $k = 10$ clusters, yielding ten concepts ($C0$–$C9$) from

a representative high-performing model (ResNet1D50 trained on the large dataset). Here, "concept presence" denotes the proportion of LADs in which a concept is detected, reported after binning the continuous regression output into target-specific output ranges. Fig. 6 summarizes, for each target, (i) concept presence across output ranges (spider plot) and (ii) exemplar concept localizations as time masks on representative input channels (corner insets), together with the concept importance score (IS). IS is derived from model gradients and captures how the prediction changes with respect to the concept across the dataset: the magnitude |IS| reflects sensitivity, while the sign indicates direction (positive increases the output in a local gradient approximation; negative decreases it). Presence is descriptive and should not be interpreted causally, but it is useful for comparing how frequently patterns occur across the output space; in practice, frequency and |IS| often complement each other but need not coincide.



(a) Die-roll height (DR) prediction for a selected corner.

(b) Process parameter prediction (die clearance).

Fig. 6: Concept presence profiles for a representative trained network (ResNet1D50 on the large dataset). The spider plot reports the fraction of the LADs in which each concept ($C0$–$C9$) is present, computed separately for three output ranges (legend). Insets show, for each concept, an example input channel with the corresponding time-localized concept mask (colored region) and the empirical importance score (IS).

For die-roll height (Fig. 6a), concept presence is dominated by one frequent concept, while the remaining concepts occur more sparsely and shift across output bins, indicating that different patterns become characteristic across die-roll height ranges. Several concepts also show large |IS| and are confined to short intervals on interpretable channels (blank-holder, punch, counter-holder). These high-importance intervals largely

fall around transitions into and out of the high-load regime, suggesting that the die-roll height is primarily determined by specific stroke phases rather than by the full time series.

For the process-parameter target (Fig. 6b), concept presence is spread across multiple concepts and changes more clearly across output bins. Unlike for the case of the die-roll height, there is no obvious "correct" stroke interval a priori for parameter recovery, so the masks and |IS| provide an indication of where the regressor extracts parameter information. Many of the high-|IS| concepts lie in the same broad stroke regions (and often the same channels) that are also important for predicting the die-roll height, suggesting temporal overlap in where the signal carries information about quality and settings. One exception is concept $C2$, whose mask includes background and stripping-phase segments; this may reflect model sensitivity to phases that are less directly tied to the shearing event and is therefore treated cautiously in the interpretation.

To reduce dependence on a single architecture and to retain interpretability of the concept-based findings, compact global concept presence profiles are reported for a high-performing ResNet1D50 model, while DenseNet1D121 is used for detailed concept overlays in the following Section 4.4. For both architectures, die-roll-relevant information is localized to the same peak-load and immediate post-peak regimes. Accordingly, the subsequent analysis is restricted to these shared, phase-consistent regions.

4.4 Localizing predictive information along the stroke

Concept presence profiles characterize which recurring patterns occur across target ranges, but they do not directly resolve *how* predictive information is distributed along the stroke. A phase-resolved view is therefore required to relate model relevance to interpretable regions of the force signature under process noise.

4.4.1 Global segment-wise attribution

Fig. 7 combines APLA-based stroke segmentation with segment-wise SHAP attribution and segment variance. This visualization tests whether model relevance is explained by trivial signal properties (e.g., high-variance or high-amplitude regions) or concentrates in specific stroke phases corresponding to regime transitions.The model was trained on the full dataset, with each experiment further split into two equally sized subsets corresponding to the processing of different coils. This setup allows systematic variations between coils to be represented in the training data, while preserving the experimental structure required for leave-one-experiment-out evaluation. Overall, the trained models exhibit stable predictive performance across leave-one-experiment-out folds; however, for two specific experiments, a pronounced degradation in both RMSE and R^2 is observed, indicating experiment-specific deviations that are not fully captured by the remaining data.

Segment-wise SHAP attributions peak in the vicinity of the force maximum and in the immediate post-peak region (segments 8–11), whereas the early loading phase and the late tail contribute substantially less. In contrast, the process-noise quantifier based on segment-wise variance reveals a different distribution: the highest variance

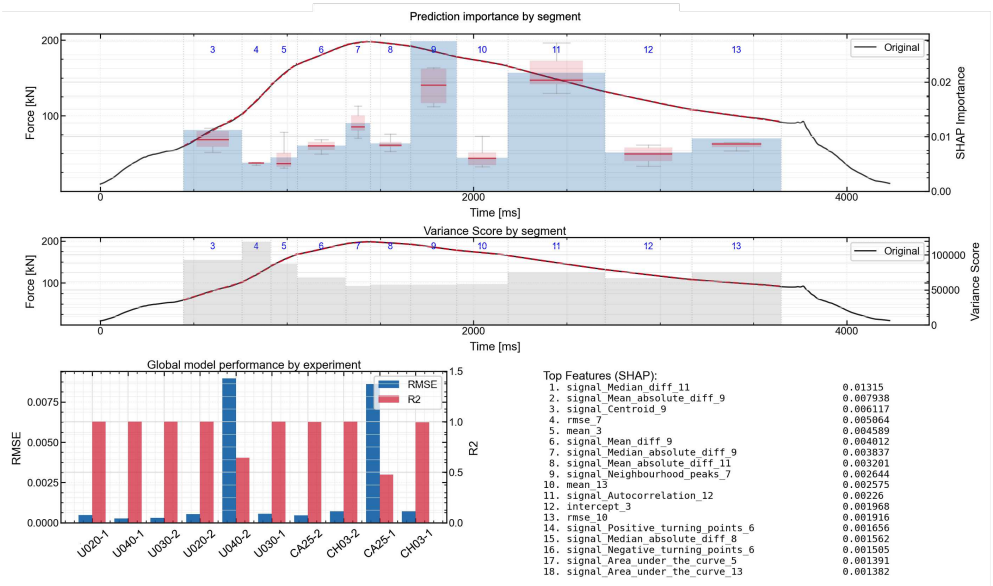


Fig. 7: Segment-based interpretation of the force signal using APLA segmentation and SHAP attribution. The top panel shows the original force–time signal with APLA-derived segments (red) and segment-wise SHAP importance for the model trained on all data(blue), and the distribution of model importance by leave-one-experiment-out evaluation (red boxplots). The middle panel reports the signal variance score per segment. The bottom panels summarize global model performance across experiments and the most influential segment-level features.

is observed predominantly in the early stroke (segments 3–4), decreases around the peak region, and increases slightly again towards the tail. Importantly, the regions of highest model relevance do not coincide with the segments of highest variance, indicating that the model does not rely on signal energy or amplitude alone, but instead exploits structurally informative stroke phases characterized by specific force-evolution patterns.

The SHAP boxplots further indicate that this segment ranking is stable under leave-one-experiment-out training: while absolute importance varies across folds, the same transition segments consistently dominate, supporting the robustness of the identified information zones across experimental conditions. The most influential features corroborate this interpretation: dominant contributions arise from difference-based measures (e.g., median and mean absolute differences), centroids, and autocorrelation, all of which capture local force evolution and temporal structure rather than absolute force levels. Together, these results indicate that quality-relevant information is encoded primarily in *how* the force changes around regime transitions, rather than in the peak force itself.

4.4.2 From segment-level attribution to concept-level inspection

The APLA-based analysis identifies *where* along the stroke predictive information is concentrated but does not reveal *which* characteristic signal patterns within these intervals are exploited by the neural network. To disentangle these patterns, individual time-series concepts extracted from a representative model are examined.

The DenseNet1D121 architecture is selected for this purpose, as it exhibits stable performance across experiments and provides a sufficiently rich concept decomposition. The set of ten concepts obtained via concept extraction is analyzed, with concept relevance values aggregated on the signal level to obtain a cumulative importance profile. From this set, two highly relevant concepts are selected and inspected by visualizing their activation on the force signal and their evolution across experiments (Fig. 8).

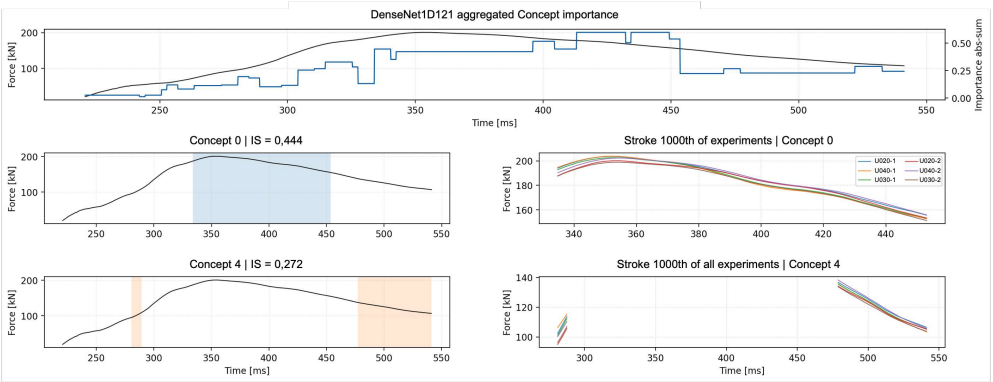


Fig. 8: Concept-level analysis for a representative DenseNet1D121 model trained on die-roll prediction. The top panel shows the force signal with aggregated concept importance. The lower panels visualize two selected high-importance concepts: left, concept overlays on the force signal; right, qualitative evolution of the corresponding force segments across all experiments.

The aggregated importance profile confirms that model attention is strongly concentrated in the peak and post-peak region, in agreement with the segment-wise SHAP analysis. Concept $C0$, the most influential concept, spans a broad interval around the force maximum and early decay, capturing the overall shape and stability of the high-load regime. Its inter-experiment comparison reveals systematic but smooth shifts in force level and decay behavior, consistent with gradual changes in process conditions across the coil.

Concept $C4$, while less dominant, is localized to narrower stroke intervals including transitions into and out of the high-load regime. Its activation highlights more localized structural differences between experiments, suggesting sensitivity to subtle changes in force evolution rather than to absolute force magnitude. Together, these concepts decompose globally relevant stroke phases into distinct, interpretable signal patterns, providing a qualitative basis for the mechanism-oriented analysis in Section 4.5.

4.5 In-depth concept analysis and phase-local die-clearance effects

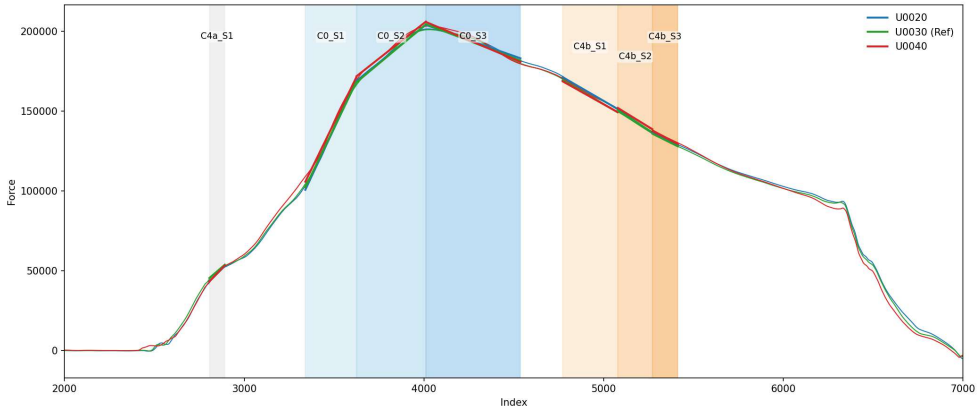


Fig. 9: Conceptual depiction of PLA segments for the two most important concepts $C0$ and $C4$ regarding die-roll height prediction.

The concept-level inspection is extended from qualitative overlays to a mechanism-oriented, phase-local comparison across die-clearance conditions. Concepts $C0$ and $C4$ were selected for this analysis as they exhibited the highest absolute importance scores for die-roll prediction among all extracted concepts (see Fig. 8, top panel). As shown in Fig. 9, these two dominant concepts are localized to physically distinct regimes of the stroke: $C0$ covers the loading and peak-load neighborhood, whereas $C4$ (in particular $C4b$) is concentrated in the post-peak decay. This localization is exploited by resolving each concept window into piecewise-linear sub-intervals and evaluating per-stroke slopes within each PLA segment. The slope is interpreted as a local proxy of the incremental force response in the respective regime, enabling a direct mapping from data-driven separations to candidate process mechanisms. Beyond maximizing separation between $U020$ and $U040$, emphasis is placed on segments whose median slopes exhibit an ordered progression across $U020 \rightarrow U030 \rightarrow U040$, consistent with the graded die-roll response.

The quantitative segment comparison yields significance (Kruskal–Wallis and Mann–Whitney p-values) for all tested segments due to the large number of pooled strokes; interpretation is therefore guided primarily by effect sizes and by the consistency of the median ordering across $U020$, $U030$ (Ref), and $U040$. For $C0_S1$, a near-maximal separation is observed between $U020$ and $U040$ ($|\delta| \approx 1$), and the medians are ordered with $U030$ between both extremes. For $C4b_S2$, a strong separation is likewise obtained ($|\delta| \approx 0.83$ for $U020$ vs. $U040$), and a monotonic ordering with clearance is supported by rank correlation. Summary statistics for the two most separating PLA segments are presented in Table 9.

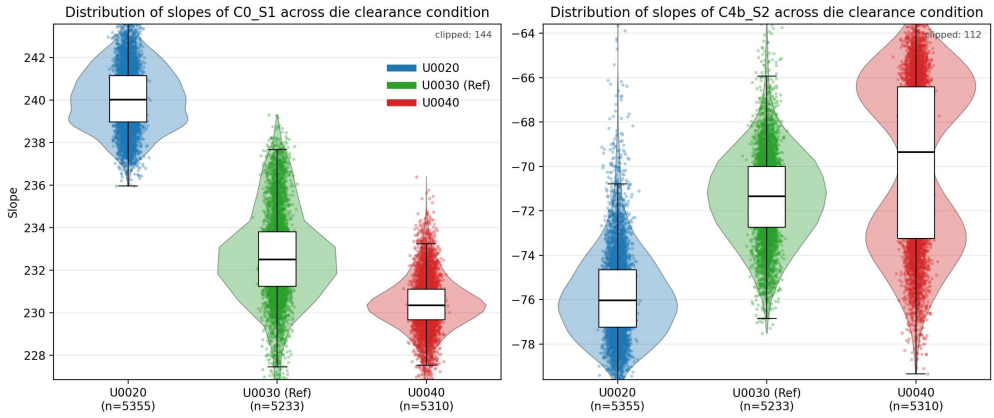


Fig. 10: Distribution of per-stroke slopes across die clearance conditions for the two segments with the highest absolute Cliff's delta $|\delta|$ between $U020$ and $U040$: the first PLA segment in $C0$ ($C0_S1$) and the second PLA segment in $C4b$ ($C4b_S2$). The reference condition $U030$ is shown between $U020$ and $U040$.

Table 9: Most separating PLA segments (ranked by max. $|\delta|$). $\delta_{20,40}$ is Cliff's delta for $U020$ vs. $U040$; ρ_S is Spearman correlation between slope and clearance. "Trend" reports median ordering across $U020 \rightarrow U030 \rightarrow U040$.

Segment	Regime	Trend with clearance	$\delta_{20,40}$	ρ_S
$C0_S1$	rising flank (peak-load onset)	decreasing	+0.999	-0.884
$C4b_S2$	post-peak decay	increasing	-0.833	+0.654

A first mechanism-sensitive separation is identified in $C0_S1$, i.e., the initial PLA segment of Concept $C0$ located on the rising flank into the high-load regime (Fig. 9). The slope distributions in Fig. 10 show a strong and ordered shift across the clearance conditions, with the reference $U030$ located between $U020$ and $U040$. In this regime, a steeper positive slope corresponds to a faster force build-up per index increment. A plausible interpretation is that $C0_S1$ captures clearance-dependent changes in the effective resistance during load transfer into the shearing zone (contact establishment and early shearing), linking Concept $C0$ to the primary force-carrying regime.

A second, complementary separation is identified in $C4b_S2$, i.e., the second PLA segment of the post-peak Concept $C4$ window (Fig. 9). In this regime, slopes are typically negative; a more negative slope indicates a faster force drop after peak load. The ordered differences across $U020$, $U030$, and $U040$ therefore suggest that clearance systematically affects post-peak unloading and separation dynamics. One hypothesis is that reduced clearance increases shearing resistance during the primary load regime (as observed in $C0_S1$), which elevates local plastic and frictional work in the shear zone. The resulting temperature rise reduces flow stress in the later stages of the stroke, leading to accelerated unloading once crack growth and contact loss set in.

This interpretation links the two concept windows mechanistically: the rising-flank separation ($C0_S1$) captures clearance-dependent resistance that preconditions the post-peak behavior quantified in $C4b_S2$.

Importantly, the two segments indicate different physical pathways by which die clearance becomes predictive. In $C0_S1$, the separation is attributed to build-up mechanics (penetration and early shearing), where the ordered slope shift is consistent with clearance-dependent resistance during the transition into the high-load regime. In $C4b_S2$, the separation is attributed to post-peak mechanics (unloading and separation), where the ordered decay slopes point to clearance-dependent boundary conditions that shape the force drop. This two-regime structure is directly visible in Fig. 9, where the identified segments occupy the peak-load rise and the post-peak decay, and is quantified in Fig. 10 by the strong, ordered slope shifts across $U020$, $U030$, and $U040$.

These findings provide an explicit link between the network's internal concepts and interpretable process effects. Concept $C0$ is positioned as the dominant carrier of information on clearance-driven shearing resistance in the primary load regime, whereas Concept $C4$ is positioned as a carrier of information on clearance-driven separation and boundary conditions that co-determine die roll formation. Notably, FEM indicated that die-clearance effects may be weak in the peak-force level but more pronounced in late-stage dynamics. The strong separation observed in $C4b_S2$ is consistent with this indication. However, the even stronger separation found on the rising flank ($C0_S1$, $|\delta| \approx 1$) was not anticipated by FEM, suggesting that the data-driven approach captures clearance-dependent effects that remain unexpressed in the simulated global force signature. The concept-local PLA analysis therefore does not only explain *where* relevant information resides, but also motivates *which* physical ingredients (e.g., frictional/thermal boundary conditions and contact loss during separation) should be strengthened to reproduce the observed clearance sensitivity in future simulation refinements.

5 Summary and outlook

Fine blanking part quality is highly sensitive to tool parameterization, yet cause–effect relationships are obscured by substantial process noise. This limits the utility of both simulation and scalar process metrics for tool optimization and mechanistic understanding. At the same time, high-resolution force-signature data is information-rich and can be used as evidence for learning physically plausible relationships, provided that interpretability is retained. This paper presents a data-to-knowledge pipeline that addresses this challenge by integrating finite element simulation, neural network prediction, and concept extraction to derive interpretable force-signature descriptors and to formulate mechanism-oriented hypotheses linking tool parameters to part quality. The pipeline was demonstrated on a large-scale industrial dataset comprising approximately 27,000 strokes with corresponding die-roll metrology.

The central finding of this work is that the combination of FE-informed sensitivity analysis, CNN-based prediction, and concept extraction with piecewise-linear analysis

yields physically meaningful, interpretable descriptors that connect tool parameterization to quality-relevant process behavior under industrial noise conditions. Three supporting findings substantiate this claim. *First*, high-resolution force-signature data contains sufficient information to reliably predict both tool parameters and die-roll height, with mean absolute percentage errors of 2.6% and 0.6%, respectively. *Second*, concept extraction (ECLAD-ts) localizes predictive information to recurring patterns concentrated in the peak-load and post-peak decay phases. *Third*, piecewise-linear analysis within these concept windows reveals that die-clearance variation manifests as statistically significant, ordered shifts in local force-signature slopes (Cliff's $|\delta| \geq 0.83$), providing an explicit link between tool parameterization and interpretable process descriptors.

From a physical standpoint, these findings support specific mechanistic hypotheses. The rising-flank slope is interpreted as a candidate descriptor of clearance-dependent shearing resistance during the transition into the high-load regime, while the post-peak decay slope is consistent with clearance-dependent changes in unloading and separation dynamics once crack growth initiates. The observation that smaller die clearance yields higher shearing resistance, faster post-peak decay, and reduced die-roll height is consistent with both FEM die-roll height simulation results and established forming theory. Notably, while FEM indicated that die-clearance effects are most pronounced in late-stage dynamics, the data-driven analysis reveals strong separation in both the rising-flank regime ($|\delta| \approx 1$) and the post-peak decay ($|\delta| \approx 0.83$), suggesting that concept-based methods capture effects not fully expressed in the simulated global force signature.

Several directions for future work emerge from this study. First, the mechanistic hypotheses generated by the pipeline can guide targeted simulation refinements, for example by strengthening frictional or thermal boundary conditions to reproduce the observed clearance sensitivity. Second, extension of the framework to other tool parameters and forming processes would assess its generalizability beyond the die-clearance case demonstrated here. Third, while the present work identifies correlational descriptors, the integration of causal inference methods offers a path toward distinguishing direct effects of tool parameters on part quality from confounded or mediated relationships.

Declarations

Funding. Funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under priority program SPP 2422 (project ID 520460745).

Conflict of interest. The authors have no competing interests.

Ethics approval and consent to participate. The authors declare that there are no conflicts with the ethical standards by Springer and the research conducted in this research paper.

Consent for publication. A consent was obtained from all individuals included in the study.

References

[1] Liewald, M., Bergs, T., Groche, P., Behrens, B.-A., Briesenick, D., Müller, M., Niemietz, P., Kubik, C., Müller, F.: Perspectives on data-driven models and its potentials in metal forming and blanking technologies. Production Engineering **16**(5), 607–625 (2022)

[2] Holzapfel, A., Posada Moreno, A.F., Trimpe, S.: Concept extraction for time series with eclad-ts. In: World Conference on Explainable Artificial Intelligence, pp. 90–112 (2025). Springer

[3] Klocke, F.: Manufacturing Processes 4: Forming. Springer, Berlin, Heidelberg (2013)

[4] Aravind, U., Chakkingal, U., Venugopal, P.: A review of fine blanking: influence of die design and process parameters on edge quality. Journal of Materials Engineering and Performance **30**(1), 1–32 (2021)

[5] Zhuang, X., Ma, S., Zhao, Z.: A microstructure-based macro-micro multiscale fine-blanking simulation of ferrite-cementite steels. International Journal of Mechanical Sciences **128**, 414–427 (2017)

[6] Mao, H., Li, S., Liu, Y., Hua, L.: An investigation on the microstructure of the fine-blanked sprocket. The International Journal of Advanced Manufacturing Technology **90**(9), 3171–3185 (2017)

[7] Kim, B., Wattenberg, M., Gilmer, J., Cai, C., Wexler, J., Viegas, F., *et al.*: Interpretability beyond feature attribution: Quantitative testing with concept activation vectors (tcav). In: International Conference on Machine Learning, pp. 2668–2677 (2018). PMLR

[8] Crabbé, J., Schaar, M.: Concept activation regions: A generalized framework for concept-based explanations. Advances in Neural Information Processing Systems **35**, 2590–2607 (2022)

[9] Ghorbani, A., Wexler, J., Zou, J.Y., Kim, B.: Towards automatic concept-based explanations. Advances in neural information processing systems **32** (2019)

[10] Posada-Moreno, A.F., Surya, N., Trimpe, S.: Eclad: extracting concepts with local aggregated descriptors. Pattern Recognition **147**, 110146 (2024)

[11] Graziani, M., Andrearczyk, V., Marchand-Maillet, S., Müller, H.: Concept attribution: Explaining cnn decisions to physicians. Computers in biology and medicine **123**, 103865 (2020)

[12] Posada-Moreno, A.F., Müller, K., Brillowski, F., Solowjow, F., Gries, T., Trimpe, S.: Scalable concept extraction in industry 4.0. In: World Conference on Explainable Artificial Intelligence, pp. 512–535 (2023). Springer

[13] Mincu, D., Loreaux, E., Hou, S., Baur, S., Protsyuk, I., Seneviratne, M., Mottram, A., Tomasev, N., Karthikesalingam, A., Schrouff, J.: Concept-based model explanations for electronic health records. In: Proceedings of the Conference on Health, Inference, and Learning, pp. 36–46 (2021)

[14] Brenner, A., Knispel, F., Fischer, F.P., Rossmanith, P., Weber, Y., Koch, H., Röhrig, R., Varghese, J., Kutafina, E.: Concept-based ai interpretability in physiological time-series data: Example of abnormality detection in electroencephalography. Computer Methods and Programs in Biomedicine **257**, 108448 (2024)

[15] Schweinshaupt, F., Stoel, T., Müller, M., Herrig, T., Bergs, T.: Thermomechanical modeling of the shearing process during fine blanking of quenched and tempered steel. Procedia Structural Integrity **61**, 214–223 (2024)

[16] Hensel, A., Spittel, T.: Kraft-und Arbeitsbedarf Bildsamer Formgebungsverfahren. Deutscher Verlag fur Grundstoffindutrie, Leipzig (1978)

[17] Oh, S.I., Chen, C.C., Kobayashi, S.: Ductile fracture in axisymmetric extrusion and drawing—part 2: Workability in extrusion and drawing. Journal of Engineering for Industry **101**(1), 36–44 (1979)

[18] Spittel, M., Spittel, T.: Steel symbol/number: 51CrV4/1.8159: Datasheet from Landolt-Börnstein, Group VIII Advanced Materials and Technologies, Volume 2C1: Metal Forming Data of Ferrous Alloys—Deformation Behaviour. Springer (2009)

[19] Huang, G., Liu, Z., Van Der Maaten, L., Weinberger, K.Q.: Densely connected convolutional networks. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 4700–4708 (2017)

[20] Ismail Fawaz, H., Lucas, B., Forestier, G., Pelletier, C., Schmidt, D.F., Weber, J., Webb, G.I., Idoumghar, L., Muller, P.-A., Petitjean, F.: Inceptiontime: Finding alexnet for time series classification. Data Mining and Knowledge Discovery **34**(6), 1936–1962 (2020)

[21] He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 770–778 (2016)

[22] Chen, T., Guestrin, C.: Xgboost: A scalable tree boosting system. In: Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, pp. 785–794. ACM, New York, NY, USA (2016)

[23] Lundberg, S.M., Erion, G., Chen, H., DeGrave, A., Prutkin, J.M., Nair, B., Katz, R., Himmelfarb, J., Bansal, N., Lee, S.-I.: From local explanations to global understanding with explainable AI for trees. Nature Machine Intelligence **2**(1), 56–67 (2020)

[24] Schmidt, R.-A.: Cold Forming and Fineblanking. Carl Hanser Verlag, Munich (2007)